

Installing Oracle database software

So far we've decided what system we're going to host the database on, downloaded the appropriate install file for that system, and unzipped the install files into a folder to begin the installation. We'll navigate to that folder and run the setup.exe file located there. This will launch the Oracle Universal Installer program to begin the installation. Those of us with experience installing the Oracle Database from prior versions will immediately notice the installer for 11gR2 has a slightly different look. It is more like a setup program for the database than the Universal Installer we're used to from previous versions, including the first edition of this topic.

We are installing the full database, which now automatically includes the Warehouse Builder client and database components. If we had an older version of the database (10g R2 for example) that did not include the Warehouse Builder software, or if we wanted to run the client on a different workstation than where the database software is installed, then there is the option to install the Warehouse Builder by itself.

A separately downloadable install for the standalone option is available at

<http://www.oracle.com/technology/software/products/warehouse/index.html>.

Skip ahead to the section titled installing the OWB standalone software if just the Warehouse Builder software is needed.

1. The first thing the installer is going to ask us is our email address for use in being notified of critical system updates that are available. This is something new Oracle has started doing with their installs to get people thinking about critical vulnerabilities and keeping their databases properly patched. It's similar to Microsoft's Windows Update feature that keeps users notified of available patches for the Windows operating system. That's a good feature but we would need a support agreement with Oracle and a My Oracle Support login to really make use of it so we're going to skip this and move on to step 2. Be aware that it will pop up a warning dialog asking us if we really want to remain uninformed about security related issues. We'll answer yes and move along.

2. The second step asks us what installation option we'd like. We can choose from one of three, create and configure a database, install the software only, or upgrade an existing database. We'll choose the second option. The Create and Install a Database option will make some assumptions about the database that we don't want depending on what options are selected so it's easier if we just install the database separately after the software is installed.

3. For the third step, we'll choose Single Instance for the Installation Type. The other option is for installing a database as part of a RAC installation (Real Application Cluster) of clustered databases.

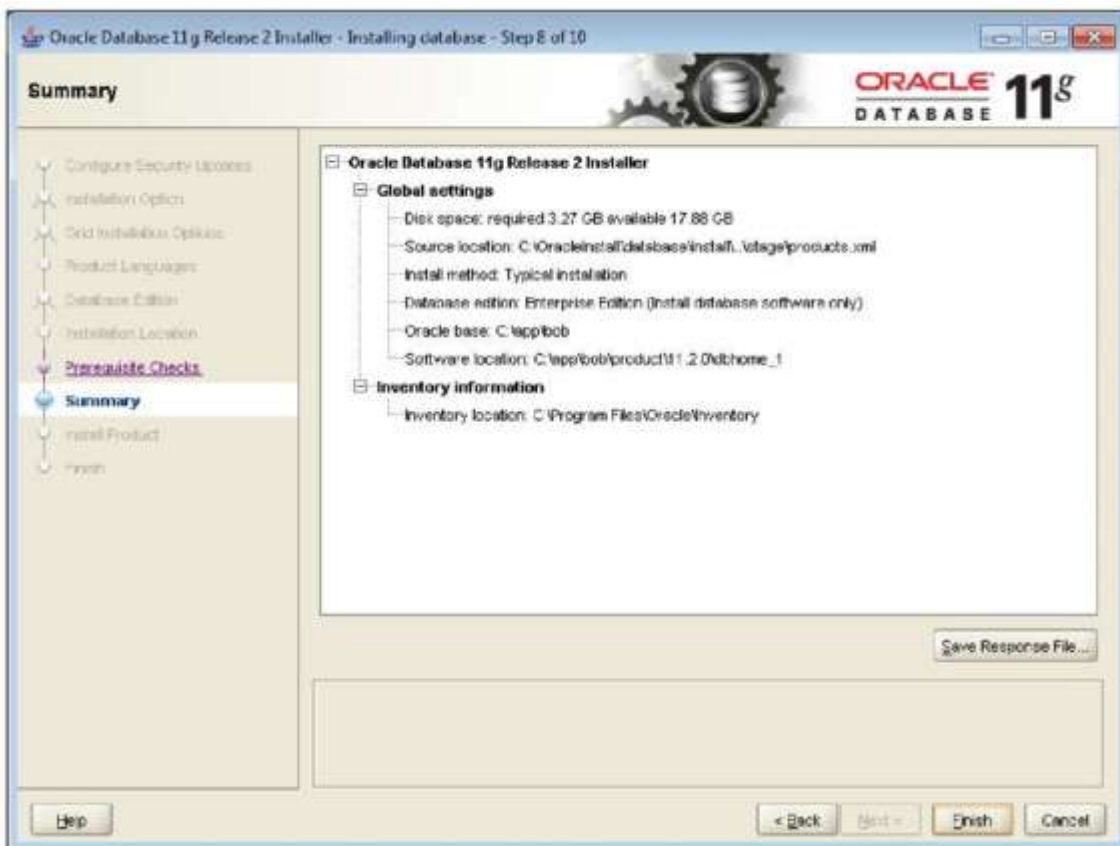
4. In step 4 we'll choose the database language, or languages, we want to install.

5. Step 5 is where we choose the edition of the database to install, Enterprise, Standard, Standard Edition One, or Personal Edition. We'll choose the Enterprise Edition so we have access to all of the advanced features we need for the Warehouse Builder.

6. Step 6 will ask us for path names for ORACLE_BASE and for the Software Location (or Oracle home location). They will have suggested paths filled in for us. It is a good idea to leave the path names as they are and only change the drive designation if we'd like to install to a different hard drive. The install program will suggest a drive for the installation, but we might have a different preference. Oracle recommends a convention for naming folders and files that they call the Optimal Flexible Architecture (OFA). This is described in topic B of the Oracle Database Installation Guide for Microsoft Windows, which can be found at the following URL: http://download.oracle.com/docs/cd/E11882_01/install.112/e10843/ofa.htm#CBBEDHEB. It is a good idea to follow their recommendations for standardization so that others who have to work with the database files will know where to find them, and to save us from problems with possible conflicts with other Oracle products we may have installed. If we keep the default folder locations intact and only change the drive letter, we will adhere to the standard.

7. Step 7 of the install is where it will conduct the prerequisite checks to ensure our system is capable of running the database. If everything succeeds we'll move right to step 8, bypassing step 7 results. We could hit the back button on Step to move back to see the results if we wanted. If anything failed, it would have displayed the results for us.

8. Step 8 is the summary screen. It will display the Global Settings as in the previous version but includes a new Inventory section which replaces the old Product Languages, Space Requirements, and New Installations sections.



9. The actual installation happens in step 9. A progress bar proceeds to the right as the installation happens and steps for Prepare, Copy Files, and Setup Files are checked off as they are done.

10. Step 10 is the conclusion and finishes up with a success message:

Your database configuration files have been installed in C:\app\ while other components selected for installation have been installed in C:\app\product\11.2.0\dbhome_1. Be cautious not to accidentally delete these configuration files.

1. We will click Close to end the installation.

Basic versus advance install

The installation method we're following here is the quickest and easiest, but makes many decisions for us that more advanced options will ask us about like creating a database and Desktop vs Server installs. For the purpose of working through the examples in this topic, we will be OK with the basic installation. But if we were installing for a production environment, we would want to read through the Oracle Database Installation Guide (<http://www.oracle.com/technology/documentation/database.html>; click on View Library to view the documentation online or click on Download to download the documentation) to familiarize ourselves with the various situations that would require us to use the more advanced installation options. This would ensure that we don't end up with a database installation that will not support our needs.

Location of install results


A good idea is to pay particular attention to the inventory location on the Step 8 summary screen, which tells us where we can find a log of the installation. The logs that the installer keeps are stored in the Oracle folder on the system drive in the following subfolder: C:\Program Files\Oracle\Inventory\logs. The files are named with the following convention: install ActionsYYYY-MM-DD_HH-MI-SSPM where YYYY is the year, MM the month, DD the day, HH the hour, MI the minutes, SS the seconds of the time the installation was performed, and PM is either AM or PM. The files will have a .log extension. This information may come in useful later to see just what products were installed. The folder also will contain any errors encountered during the installation in files with a file extension of .err and any output generated by the installer in files with a file extension of .out.

Now that the software is installed, it's time to proceed with creating a database. But there is one step we have to do first—we need to configure the listener.

Configuring the listener

The listener is the utility that runs constantly in the background on the database server, listening for client connection requests to the database and handling them. It can be installed either before or after the creation of a database, but there is one option during the database creation that requires the listener to be configured—so we'll configure it now, before we create the database.

Run Net Configuration Assistant to configure the listener. It is available under the Oracle menu on the Windows Start menu as shown in the following image:



The welcome screen will offer us four tasks that we can perform with this assistant. We'll select the first one to configure the listener, as shown here:



The next screen will ask you what we want to do with the listener. The four options are as follows:

- Add
- Reconfigure
- Delete
- Rename

Only the Add option will be available since we are installing Oracle for the first time. The remainder of the options will be grayed out and will be unavailable for selection. If they are not, then there is a listener already configured and we can proceed to the next section—creating the database. For those of us installing for the first time on our machines, we need to proceed with the configuration. The next screen will ask us what we want to name the listener. It will have LISTENER entered by default and that's a fine name, which states exactly what it is, so let's leave it at that and proceed.

The next screen is the protocol selection screen. It will have TCP already selected for us, which is what most installations will require. This is the standard communications protocol in use on the Internet and in most local networks. Leave that selected and proceed to the next screen to select the port number to use. The default port number is 1521, which is standard for communicating with Oracle databases and is the one most familiar to anyone who has ever worked with an Oracle database. So, change it only if you want to annoy the Oracle people in your organization who have all memorized the default Oracle port of 1521.

To change or not change the default listener port


Putting aside the annoyance, the Oracle people might have to suffer as there are valid security reasons why we might want to change that port number. Since it is so common, the people accustomed to working with the Oracle database aren't the only people who know that port number. Hackers looking to break into an Oracle database are going to go straight for that port number, so if we change it to something obscure, the database will be harder to find on the network for the people with malicious intent. If it does get changed, be sure to make a note of the assigned number. There also may be firewall issues that allow only certain port numbers to be open through the firewall, which means communication on any of the other port numbers would be blocked. 1521 might be allowed by default since it is common for the Oracle database. It would be a good idea to check with the network support personnel to get their recommendation.

That is the last step. It will ask us if we want to configure another listener. Since we only need one, we'll answer "no" and finish out the screens by clicking on the Finish button back on the main screen.

Creating the database

So far we have the Oracle software installed and a listener configured, but we have not created a database.

We will install a new database using Database Configuration Assistant, which Oracle provides to walk us step-by-step through the process of creating a database. It is launched from the Windows Start menu as shown in the following image:



Running this application may require patience as we have to wait for the application to load after it's selected. Depending on the system it is running on, it can take over a minute to display, during which time there is no indication that anything is happening. It may be tempting to just select it again from the Start menu because it appears it didn't work the first time, but don't as that will just end up running two instances of the program. It will appear soon. The following are steps in the creation process:

1. The first step is to specify what action to take. Since we do not have a database created, we'll select the Create a Database option in step 1. If there was a database already created, the options for configuring a database or deleting a database would be selectable. Templates can be managed with the Database Configuration Assistant application, which are files containing preset options for various database configurations. Pre-supplied templates are provided with the application, and the application has the ability to custom-build templates.

In previous versions of the database, Automatic Storage Management could be configured as well however as of 11gR2, ASM has its own configuration assistant now, ASMCA. It is Oracle's feature for databases for automatically managing the layout and storage of database files on the system. These are both topics for a more advance topic on the Oracle Database. We will be creating a database using an existing template.

2. This step will offer the following three options for a database template to select:

° General Purpose or Transaction Processing ° Custom Database ° Data Warehouse

3. We are going to choose the Data Warehouse option for our purposes. If we already had a database installed that we wanted to use for learning OWB, but that's not configured as a data warehouse, it's not a problem. We can still run OWB hosted on it and create the data warehouse schema (database user and tables), which we'll be creating as we proceed through the topic. This would be fine for learning purposes, but for production-ready data warehouses a database configured specifically as a data warehouse should be used.

4. This step of the database creation will ask for a database name. The name of the database must be one to eight characters in length. Any more than that will generate an error when trying to proceed to the next screen. This is an Oracle database limitation. The database name can also include the network domain name of the domain of the host it is running on, to further uniquely identify it. Follow the name with a period and then the domain, which itself can include additional periods. If this database is being created for business use, a good naming scheme would reflect the purpose of the database. Since we're creating this database for the data warehouse of ACME Toys and Gizmos Company, we'll choose a name that reflects this—ACME for the company name and DW for data warehouse, resulting in a database name of ACMEDW. It is important to remember this name as it will be a part of any future connections to the database.

As the database name is typed in, the SID (or Oracle System Identifier) is automatically filled in to match it. If the domain is added to the database name, the SID will stop pre-populating after the first period is entered. The end result is that the SID becomes the same as the first part of the database name.

5. This step of the database creation process asks whether we want to configure Enterprise Manager. The box is checked by default and left as it is. This is a web-based utility Oracle provides for controlling a database, and as it is very useful to have, we will want to enable it. There are two options for controlling a database: registering with Grid Control or local management. Grid Control is Oracle's centralized feature for controlling a grid, a network of loosely coupled modular hardware and software components that can be joined and rejoined together on demand to meet business needs. That is what the "g" in Oracle Database 11g stands for. If your network is not configured in a grid architecture, or you are installing on a standalone machine, then choose the local management option. It will automatically detect a Grid Control agent that is running locally, and if it doesn't find one, the Grid Control option will be grayed out anyway. In that case, you will only be able to select local management.

6. New in the 11gR2 version of the DBCA is the additional tab on this screen for the Automatic Maintenance option. This step used to be all by itself as step 12 of the install process. We'll deselect that option and move on, since we don't need that additional functionality. Automatic Maintenance

Tasks are tasks that run in predefined maintenance windows of time to perform various preconfigured maintenance operations on the database. Since the database for this topic is only for learning purposes, it is not critical that these maintenance tasks be done automatically.

7. Automatic maintenance is designed to run during preset maintenance windows, which are usually in the middle of the night. So if the database system is shut down every day, there wouldn't be a good window to run the tasks on regularly anyway. If installing in a production environment with servers that will be running 24 hours a day every day, then consider setting up the automatic maintenance to occur. Oracle provides three pre-configured maintenance tasks to choose from—collecting statistics for the query optimizer (for improving performance of SQL queries), Automatic Segment Advisor for analyzing storage space for areas that can possibly be reclaimed for use, and the Automatic SQL Tuning Advisor for automatically analyzing SQL statements for performance improvements.

If an error is encountered at some point during the database creation that indicates a listener is not configured, it simply means we started the DBCA before configuring a listener. To solve that, there is no need to exit out of the database install window, just go back and perform the listener install steps and come back here where the screen will allow us to proceed.

8. On this screen (step 5) we can set the database passwords on the system accounts using a different one for each account, or by choosing one password for all four. We're going to set a single password on all four, but for added security in a production environment, it is a good idea to make a different password for each. Click on the option to Use the Same Administrative Password for All Accounts and enter a password. This is very important to remember as these are key system accounts used for database administrative control. Oracle has decided as of this release to attempt to make us choose better, more secure passwords by popping up a warning dialog if the password we've chosen does not meet with their standards of complexity. It is only a warning however and we can answer Yes and continue.

9. This step is a combination now of the old steps 6 and 7 of the installation. It is about storage and specifying the locations where database files are to be created. We'll leave it at the default of File System for storage management. The other two options are for more advanced installations that have greater storage needs. The locations where database files are to be created can be left at the default for simplicity (which uses the locations specified in the template and follows the OFA standard for naming folders described above). A storage screen will come up where we'll be able to change the actual file locations if we want, for all but the Oracle-Managed Files option.

The Oracle-Managed Files option is provided by the database so that we can let Oracle automatically name and locate our data files. A folder location is specified on the step 6 screen, which will become the default location for any files created using this option. This is why we won't be able to change any file locations later on during the installation if this option is chosen. However, files can still be created with explicit names and locations after the database is running.

10. The next screen is for configuring recovery options. We're up to step 8 now. If we were installing a production database, we would want to make sure to use the Flash Recovery option and to enable Archiving. Flash Recovery is a feature Oracle has implemented in its database to provide a location that is managed by the database. It stores backups and files needed to recover a database in the event of disk failure. With Flash Recovery Area specified, we can recover data that would otherwise be lost in a system failure.

Enabling archiving turns on the archive log mode of the database, which causes it to archive the redo logs (files containing information that is used by the database to recover transactions in the event of a failure.) Having redo logs archived means you can recover your database up to the time of the failure, and not just up to the time of the last backup.

These recovery options will consume more disk space, but will provide a recovery option in the event of a failure. Each individual will have to make the call for their particular situation whether that is needed or not.

We'll specify Flash Recovery and for simplicity, we will just leave the default for size and location. We will not enable archiving at this point. These options can always be modified after the database is running, so this is not the last chance to set them.

11. This step is where we can have the installation program create some sample schemas in the database for our reference, and specify any custom scripts to run. The text on the screen can be read to decide whether they are needed or not. We don't need either of these for this topic, so it doesn't matter which option we choose.

12. The next screen is for Initialization Parameters. These are the settings that are put in place to define various options for the database such as Memory options. There are over 200 different parameters and to go through all of them would take much more time and space than we have here. There is no need for that at this point as there are about 28 parameters that Oracle says are basic parameters that every database installation should set. We're just going to leave the defaults set on this screen, which will set the basic parameters for us based on the amount of memory and disk space detected on our machine. We'll just move on from here. Once again, these can all be adjusted later after the database is created and running if we need to make changes.

13. The next step (step 10 of 11) is the Database Storage screen referred to earlier. Here the locations of the pre-built data and control files can be changed if needed. They should be left set to the default for simplicity since this won't be a production database. For a production environment, we would want to consider storing data files on separate partitions for performance reasons, and to minimize the impact of disk failures on the running database if something goes wrong. If all the data files are on one drive and it goes bad, then the whole database is down.

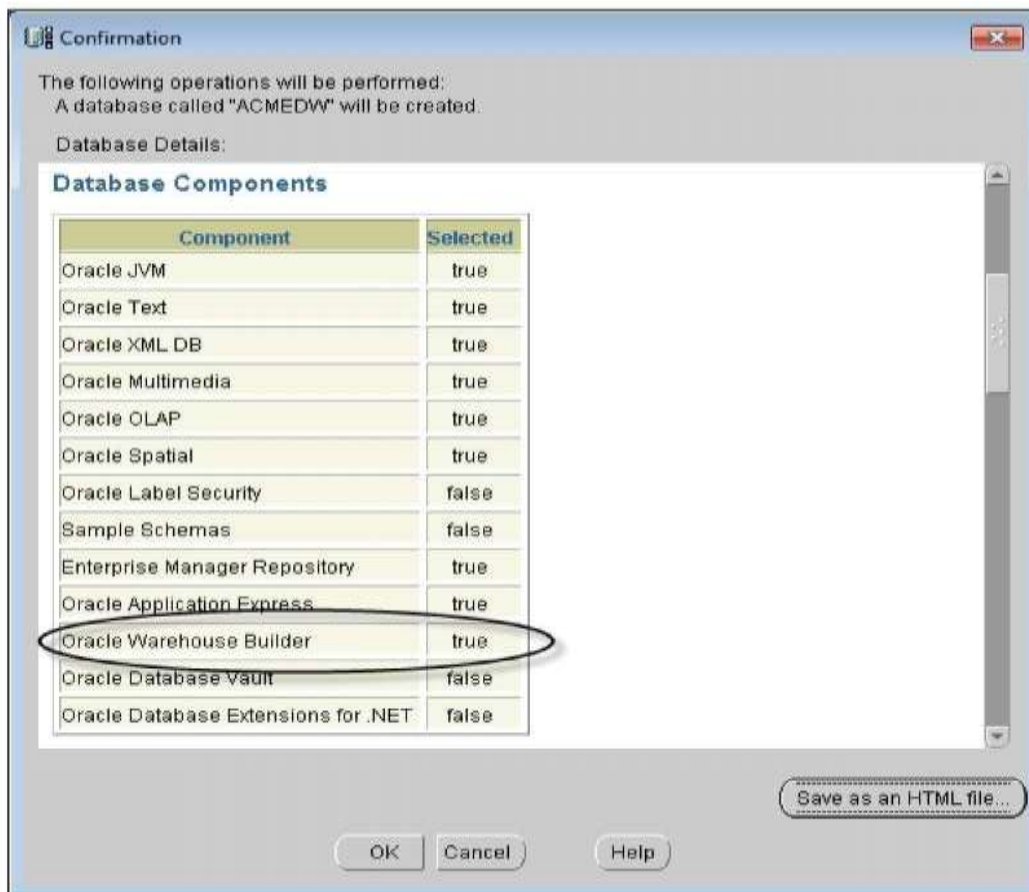
14. The final step has the following three options, and any or all can be selected for creating the database:

- ° Create the database directly
- ° Save the creation options as a template for later use
- ° Save database creation scripts that can be used later to create the database

We'll leave the first checkbox checked to go ahead and create the database.

The Next button is grayed out since this is the last screen. So click on the Finish button to begin creating the database using the selections we've just chosen. It will display a summary screen showing what options it will be using to install with. We can save this as an HTML file if we'd like to keep a record of it for future reference.

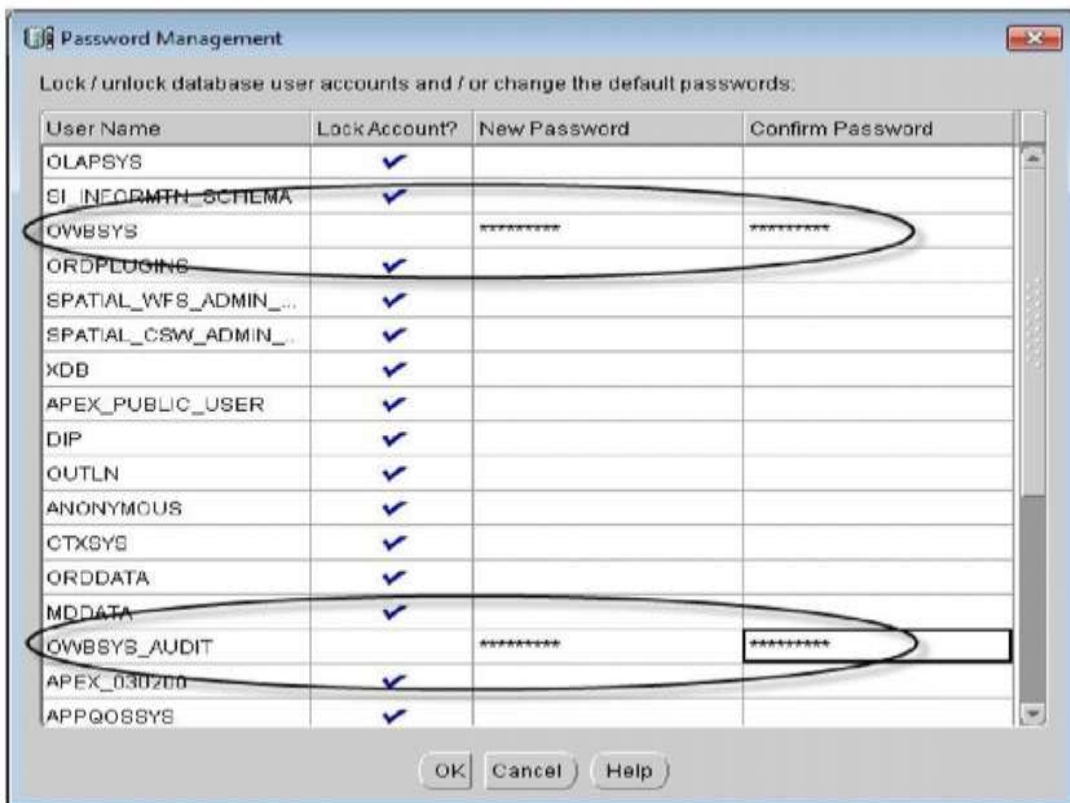
All that information will be available in the database by querying system tables later, but it's nice to have it all summarized in one file. We can scroll down that window and verify the various options that will be installed, including Oracle Warehouse Builder, which will have a true in the selected column as shown here:



We will be presented with the progress screen next that will show us the progress as it creates the database.

When the install progress screen gets to 100% and all the items are checked off, we will be presented with a screen summarizing the database configuration details. Take a screen capture of this screen or write down the details because it's good to know information on how the database is configured. Especially, we'll need the database name in later installation steps. We may see the progress screen at 100% doing nothing with apparently no other display visible. Just look around the desktop underneath other windows for the Database Configuration Screen. It's important for the next step.

On the final Database Configuration Screen, there is a button in the lower right corner labeled Password Management. We need to click on this button to unlock the schema created for OWB use. Oracle configures its databases with most of the pre-installed schemas locked, and so users cannot access them. It is necessary to unlock them specifically, and assign our own passwords to them if we need to use them. Two of them are the OWBSYS schema and the OWBSYS_AUDIT schema. These are the schemas that the installation program automatically installs to support the Warehouse Builder. They are required for running the Warehouse Builder. Click on the Password Management button and on the resulting Password Management screen, we'll scroll down until we see the OWBSYS and OWBSYS_AUDIT schemas. We'll click on the check box to uncheck it on each one (indicating we want them unlocked) and then type in a password and confirm it as shown in the following image:



Click on the OK button to apply these changes and close the window. On the Database Configuration Screen, click on the Exit button to exit out of the Database Configuration Assistant. That's it. We're done installing our first database and it's ready to use. Next, we'll discuss installing the OWB client if we want to run the client on another computer, or if we already have a 10gR2 database installed that we want to use with the Warehouse Builder.

Installing the OWB standalone software

If we are going to run the OWB client on the same computer as we just installed the Oracle database on, we don't need any more installations. That is the configuration used in this topic. The OWB client software is now installed by default with the main database installation. We can verify that by checking the Start menu entry for Oracle. We will see a submenu entry for Warehouse Builder as shown in the following image:



If we want to run the OWB client on another computer on the network, or if we have an older version of the database already installed (10g Rel 2) and want to be able to use the Warehouse Builder software with it, we'll need to continue here with the installation of the OWB client software. For all others, we can proceed to the next section on OWB — OWB components and architecture.

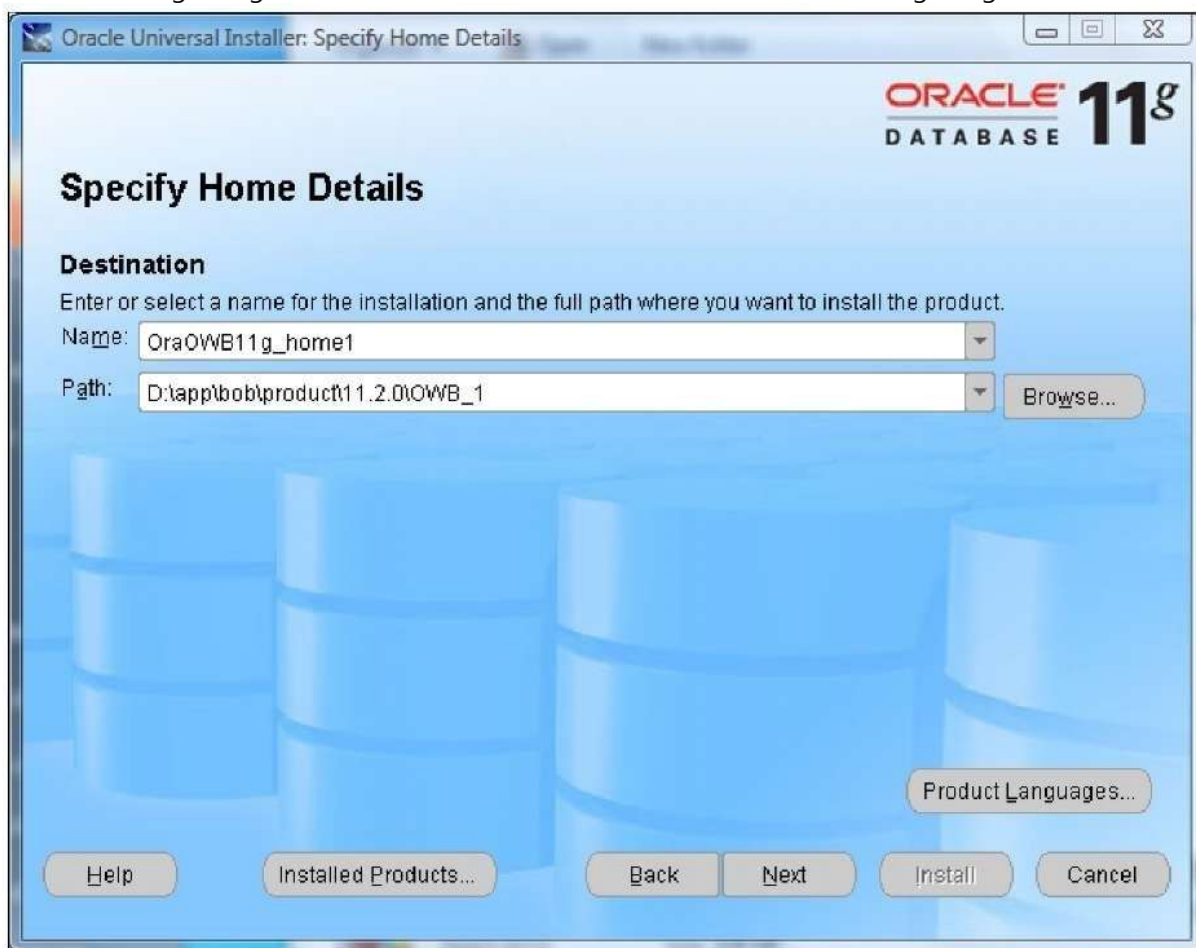
For the task of installing the standalone client, we'll need to download the OWB client install file. So we will go back to the Oracle site on the Internet. The download page is at the following URL at the time of writing: <http://www.oracle.com/technology/software/products/warehouse/index.html>. If that is not working, go to the main Oracle site and search for the Business Intelligence | Data Warehousing page where there is a link for the download of the OWB client.

Once again we'll have to accept the license agreement before the download links will become active. So we'll accept it and download the install file to the client computer on which we'll be installing the software. The Windows ZIP file is about 930MB in size so we need to make sure we

have enough room on our hard drive to store the file. We'll need at least double that amount of space because the install files will take up that much space when unzipped.

When we have downloaded the ZIP file and unzipped it to our hard drive, run setup.exe in the top-level folder to run the Oracle Universal Installer. It should look familiar. Oracle is definitely correct in calling their installer "Universal". Every Oracle database product uses that installer, so we will become very familiar with it if we have to install any more Oracle products. It is universal also in the fact that it runs on every platform that Oracle supports, and so the same interface is used no matter where we install it. The installation steps are as follows:

1. The first step it goes through is asking us for the Oracle home details. It's similar to what it asked at the beginning of the database installation as shown in the following image:



The installer will again suggest OraDb11g_home1 or something similar, but we'll change it to OraOWB11g_home1 since it's just the OWB installation and not the full database.

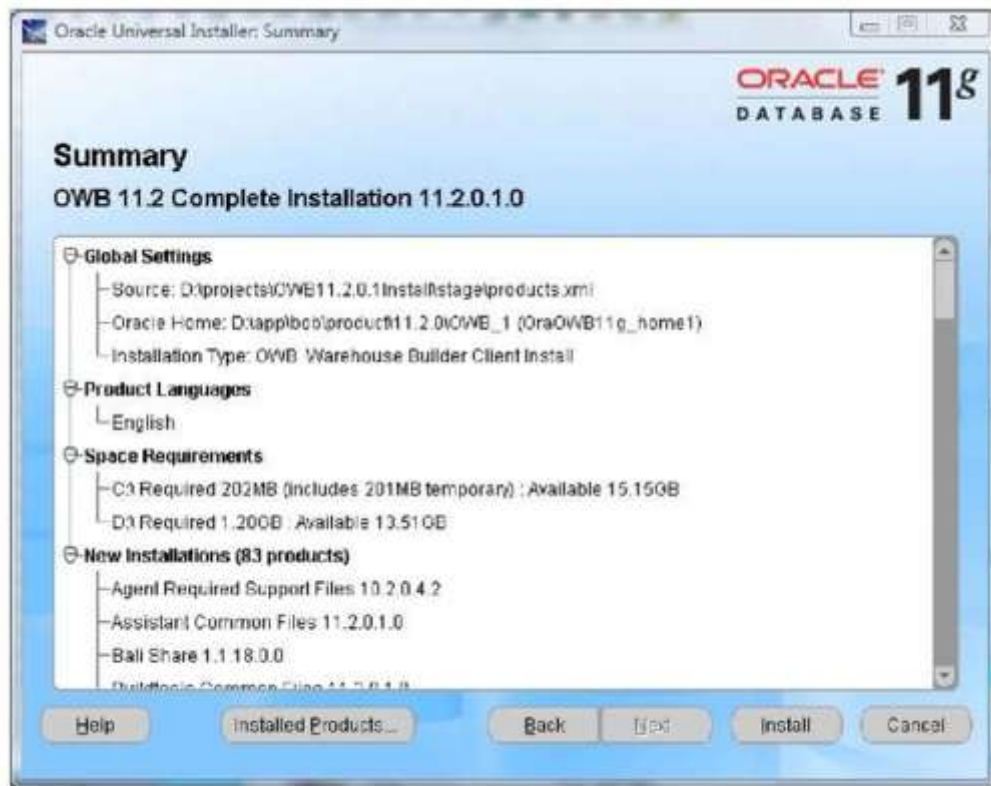
When installing the standalone OWB software, remember that it cannot be installed into the same ORACLE home as the database. It must reside in its own Oracle home folder. So if we have a database that's already installed on the same machine, we'll have to make sure the ORACLE_HOME

we specify is different. The installer will warn us if we try to specify the same one and won't let us continue until it is different.

We need to verify the installation location for the home location also. The suggested name that it provides conforms to the OFA standard just as the database installation did, so we'll want to just change the drive letter if needed. However, the bottom-most folder name can be changed if needed without violating the OFA standard. If it has a default of db_1, we can change it to OWB_1 just to be clear that it's the OWB client.

2. The second step is that email prompt again for being notified about security issues that we saw earlier in the database software install. We can just continue and answer yes to the popup warning.

3. The next and final step is the summary screen. The OWB client installation is not as complex as a full database installation, so it does not need all the additional information it asked for during the database installation. The summary screen should look similar to the following:



This summary gives us an idea of the disk space it will need, as well as the products that will be installed. If we scroll down the list, we'll see a number of other Oracle utilities and applications that it will install. We will also see items that are installed on the server as a part of the database install, but that will now be available to us on our client workstation. SQL*Plus appears there, which is the command line utility for accessing an Oracle database directly using SQL (Structured Query Language, the language used for accessing information stored in databases) among a host of other features.

Upon proceeding, the next screen will begin the installation and present us with the progress screen with a sliding bar moving to the right to indicate how far it has progressed. This is similar to what it did for the full database installation. An example of that screen is included next for reference:



Install results

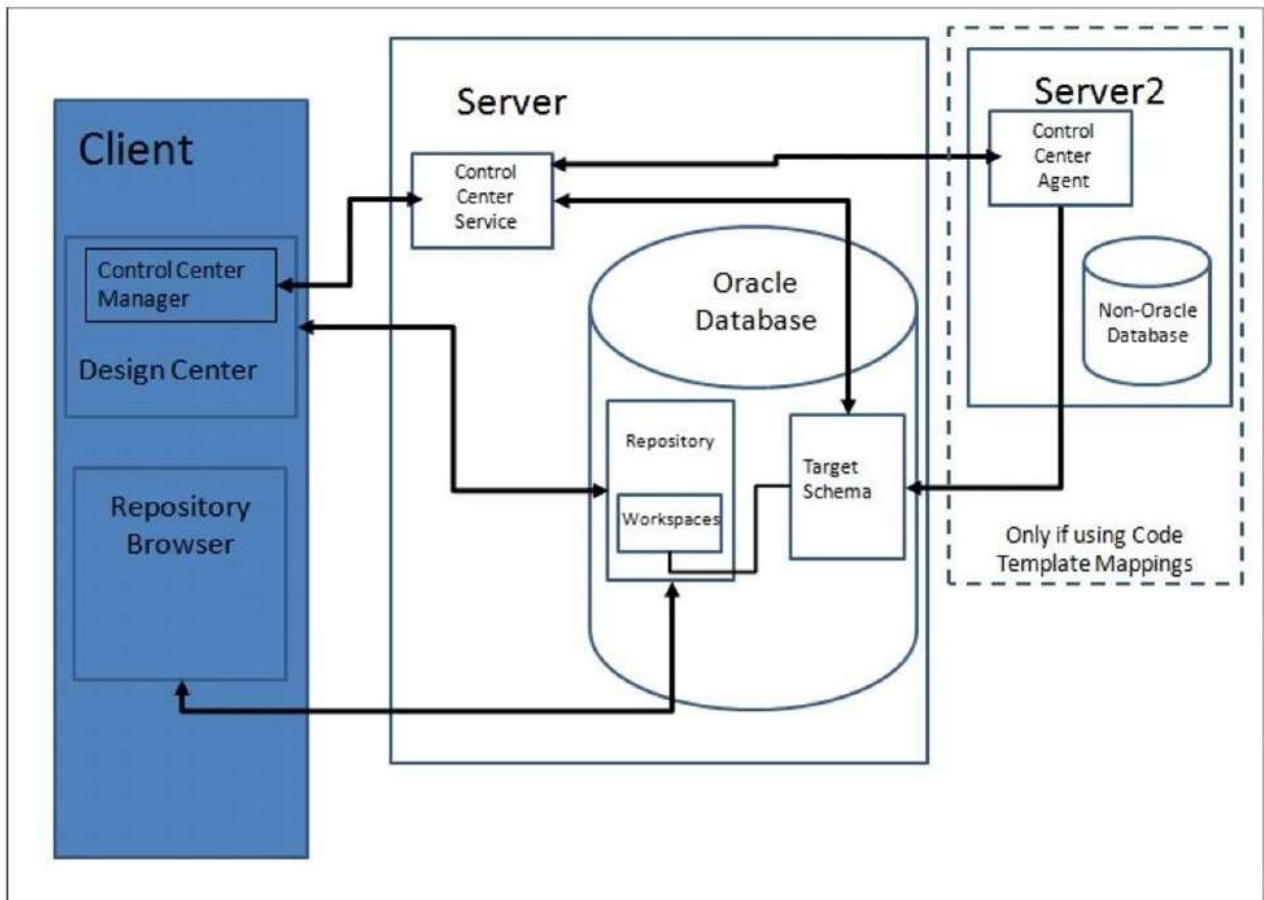
The log files with the results of the installation are stored in the same location as they are for a full database install. The universal installer will use that same folder for all its installs.

When the installation is complete, we will be presented with the final success screen and an Exit button. And as if to remind us about the universal nature of the installer, it will pop up a confirmation box asking if we really want to exit, even though for this installation there is nothing else that would be available to do on that final screen if we said no.

OWB components and architecture

Now that we've installed the database and OWB client, let's talk about the various components that have just been installed that are a part of the OWB and the architecture of OWB in the database. Then we'll perform one final task that is required before using OWB to create our data warehouse.

Oracle Warehouse Builder is composed on the client of the Design Center (including the Control Center Manager) and the Repository Browser. The server components are the Control Center Service, the Repository (including Workspaces), and the Target Schema. New for release 11gR2 is the Control Center Agent which is used by the new Code Template Mappings to communicate with a non-Oracle database. We'll be covering that in more detail in topic 10. A diagram illustrating the various components and their interactions follows:



Client and server

The previous diagram depicts a client and server, but these are really just logical notions to indicate the purpose of the individual components and are not necessarily physically separate machines. The client components are installed with the database as we've seen previously and, therefore, can run on the same machine as the database. This configuration is assumed throughout the topic.

The Design Center is the main client graphical interface for designing our data warehouse. This is where we will spend a good deal of time to define our sources and targets, and describe the extract, transform, and load (ETL) processes we use to load the target from the sources. The ETL procedures

are what we will define to carry out the extraction of the data from our sources, any transformations needed on it and subsequent loading into the data warehouse. What we will create in the Design Center is a logical design only, not a physical implementation. This logical design will be stored behind the scenes in a Workspace in the Repository on the server. The user interacts with the Design Center, which stores all its work in a Repository Workspace.

We will use the Control Center Manager for managing the creation of that physical implementation by deploying the designs we've created into the Target Schema. The process of deployment is OWB's method for creating physical objects from the logical definitions created using the Design Center. We'll then use the Control Center Manager to execute the design by running the code associated with the ETL that we've designed. The Control Center Manager interacts behind the scenes with the Control Center Service, which runs on the server as shown in the previous image. The user directly interacts with the Control Center Manager and the Design Center only.

The Target Schema is where OWB will deploy the objects to, and where the execution of the ETL processes that load our data warehouse will take place. It is the actual data warehouse schema that gets built. It contains the objects that were designed in the Design Center, as well as the ETL code to load those objects. The Target Schema is not an actual Warehouse Builder software component, but is a part of the Oracle Database. However, it will contain Warehouse Builder components such as synonyms that will allow the ETL mappings to access objects in the Repository.

The Repository is the schema that hosts the design metadata definitions we create for our sources, targets, and ETL processes. Metadata is basically data about data. We will be defining sources, targets, and ETL processes using the Design Center and the information about what we have defined (the metadata) is stored in the Repository.

The Repository is a Warehouse Builder software component for which a separate schema is created when the database is installed—OWBSYS. This is the schema we talked about unlocking during the installation discussion previously as one of the final steps in the database creation process. This will be created automatically by the 11g install, but is installed separately using scripts if we want to host the Repository on an Oracle 10g database. The explanations in this topic all assume that the Repository is hosted on an Oracle 11g database. The Oracle Warehouse Builder Installation and Administration Guide found at the following URL:

http://download.oracle.com/docs/cd/E11882_01/owb.112/e17130/toc.htm discusses the procedure for installing the Repository schema on an Oracle 10g release 2 database if needed. It also contains more detailed information about the various other configurations that can be installed along with more detailed diagrams of the various options.

The Repository will contain one or more Workspaces as shown in the previous diagram. A Workspace is where we will do our work to create the data warehouse. There can be more than one workspace defined in the Repository. A common example of how multiple workspaces can be employed is to

use different workspaces corresponding to sets of users working on related projects. We could have one workspace for development, one for testing, and one for production. The development team could be working in the development environment separately from the test team that would be working in the test environment. For our purposes at the ACME Toys and Gizmos Company, we will be working out of one workspace.

This concept of the workspace is new in the 11g releases of OWB. The Repository is created in the OWBSYS schema during the database installation. So setting up the Repository information and workspaces no longer requires SYSDBA privileges for the user to install the Repository. SYSDBA is an advanced administrative privilege that is assigned to a user in an Oracle database. This allows the user to perform tasks affecting the database and other database users that ordinary user accounts cannot do (or for that matter, other administrative accounts without SYSDBA). For security reasons, we want to restrict user accounts with SYSDBA privilege to a minimum. So it is good that we don't have to use that privilege when we install the Repository.

One final OWB component to consider is the Repository Browser on the client. It is a web browser interface for retrieving information from the Repository. It will allow us to view the metadata, create reports, and audit runtime operations. It is the only other component besides the Design Center and the Control Center Manager that the user interacts with directly.

We will have a chance to visit each one of these areas in much more detail as we progress through the design and build of our data warehouse. However, first there is one more installation step we have to take before we can begin using the Warehouse Builder. The Repository must be configured for use and a workspace must be defined.