

Technical white paper

Standardize Oracle RAC Provisioning Using HP DMA



HP Database and Middleware Automation version 10.21

Table of Contents

| | |
|---|----|
| Purpose | 2 |
| Prerequisites | 4 |
| Process Overview | 5 |
| Workflow 1: Oracle - Provision or Upgrade Grid Infrastructure | 6 |
| Workflow 2: Oracle - Provision Database Software | 8 |
| Workflow 3: Oracle - Provision Database | 10 |
| FAQs | 12 |
| Appendix A: How to meet prerequisites | 14 |
| Appendix B: Alternative methods for specifying input files | 15 |
| Appendix C: How to import a file into the software repository | 16 |
| Appendix D: How to run an HP DMA workflow | 17 |
| Appendix E: How to expose hidden parameters | 18 |

Purpose

This paper describes how to use HP Database and Middleware Automation (HP DMA) to create a repeatable, standardized “gold image” for provisioning an Oracle Grid Cluster Ready Services (CRS), Automatic Storage Management (ASM), and Real Application Clusters (RAC) database.

Tip: To provision Oracle Grid standalone environments, see *Standardize Oracle Grid Standalone Provisioning Using HP DMA*.

This document is available at: support.openview.hp.com/selfsolve/manuals

Benefits of HP DMA

HP DMA automates many of the daily administrative tasks required to manage the lifecycle of relational databases and J2EE application servers. These tasks are complex, often manual, typically time-consuming, and frequently error-prone. HP DMA improves the efficiency of these administrative tasks, enabling administrators to deliver change faster with higher quality, better consistency, and improved reliability.

HP DMA equips you to do the following:

- Define and enforce standards for software installation
- Define an installation process once and reuse it repeatedly
- Leverage knowledge enterprise-wide
- Avoid human error

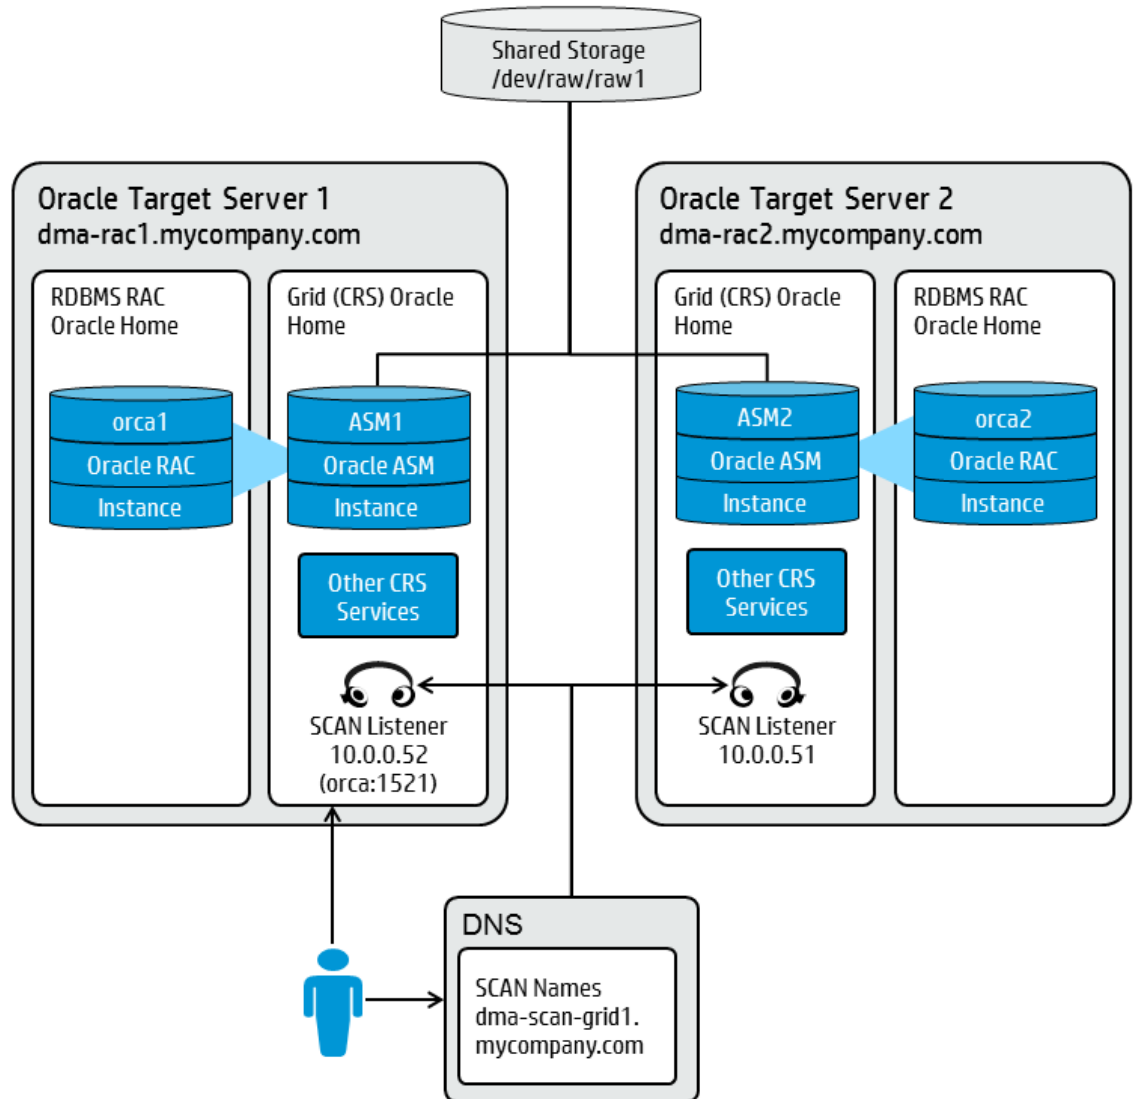
What Oracle Grid infrastructure for cluster does

The Oracle Grid infrastructure for cluster allows an Oracle database to participate as a RAC database and use common ASM storage across nodes. It enables the user to use these features:

- Start automatically with the server
- Manage the configurations of the database
- Run Oracle Restart
- Use ASM
- Manage nodes in an Oracle cluster
- Manage virtual IP addresses and SCAN virtual IP addresses

Goal

This paper describes how to create the following configuration. This example shows a RAC database that lives on two nodes of a CRS cluster:



HP DMA Versions

In the course of HP DMA development, improvements to workflows and steps could cause customers' existing deployments and customizations to break (for example, by changing steps, parameters, or mappings). In this case, HP DMA releases the upgrade with a version appended to the name. For example: Workflow Name v2

You should always use the latest (highest version number) of the HP DMA workflows and steps.

Prerequisites

Before performing the procedures in this paper, your environment must meet the following minimum requirements:


- A server running one of the following operating systems:
 - Linux (any version that is supported by Oracle and HP DMA)
 - AIX
 - Solaris

This server must meet the Oracle requirements for installing 12c (see [Oracle Database Online Documentation 12c Release 1 \(12.1\)](#) for more information).
- A shared raw disk (or disks) to be used across all of the nodes. The disk needs to be available to be mounted and used by Oracle ASM. The device cannot be formatted, but it may be partitioned.
- Storage:
 - A staging directory with 8 gigabytes available to unzip the Oracle Grid Infrastructure and Oracle Database binaries.
 - For ASM disks, a minimum of 5 gigabytes combined for logical storage (more may be required for your environment).
 - A minimum of 30 gigabytes on the partition to install Oracle Grid Infrastructure and Oracle Database Homes (more may be required for your environment).
- Networking
 - A separate NIC configured for a VLAN private interconnect on both nodes.
 - Virtual IP addresses, one per node.
 - A SCAN virtual IP address, one per cluster.
- Licenses for Oracle Database and HP DMA.

Tip: For a specific example of how to meet these general requirements for Oracle Enterprise Linux version 6, see [How to meet prerequisites](#). You can adapt these instructions for your platform.

Process Overview


Use the following HP DMA workflows to standardize the process of provisioning Oracle Grid (CRS), ASM, and database:



The following sections provide detailed information required to run each workflow.

Note: For additional information, see [How to run an HP DMA workflow](#).

Workflow 1: Oracle - Provision or Upgrade Grid Infrastructure



This section provides detailed information required to run the Oracle - Provision or Upgrade Grid Infrastructure workflow.

Solution pack

This workflow requires the HP DMA Advanced Database Provisioning Solution Pack.

Parameters to expose

None

Input parameters

When you deploy the Oracle - Provision or Upgrade Grid Infrastructure workflow, specify input parameter values for the following steps.

Note: Bold text in the following tables indicates that you must specify a value for the parameter.

Step: Gather Parameters for Provision Grid Infrastructure

| Parameter | Description | Example Value |
|---------------------------|--|--|
| ASM Diskgroup List | Required: A comma-separated list of the diskgroups that you are creating, with each diskgroup containing a comma-separated list of its associated disks. | ASMDATA(/dev/raw/raw1) |
| ASM Groups | Required: The operating system groups that manage ASM. The syntax is: ASMGroup:ASMDBA:ASMOper | oinstall:dba:oinstall |
| ASM Password | Required: The password for provisioning an Oracle database using ASM storage. The default is Manager1. | ●●● |
| CRS Base | Required: The location of the Oracle Base directory. This is where the admin directory is located. | /u01/app/grid |
| CRS Home | Required: The location where the CRS software will be installed. The default is: /u01/app/oracle/product/12.1.0/grid1 | /u01/app/oracle/product/12.1.0/grid1 |
| Oracle Software | Required: A comma-separated list of the Oracle Database software (CRS) archive files (.zip or .cpio.gz). ¹ | linuxamd64_12c_grid_1of2.zip, linuxamd64_12c_grid_2of2.zip |


¹If the file(s) are not found on the target server(s), they will be downloaded from the software repository. For additional information, see [Alternative methods for specifying input files](#).

Step: Gather Advanced Parameters for Provision Grid Infrastructure

| Parameter | Description | Example Value |
|---------------------------------|---|---|
| CLUSTER - Grid Node List | Optional: A comma-separated list of the Grid Infrastructure nodes. Must be in the format: node1-public:node1-private:node1-virtual, node2-public:node2-private:node2-virtual | racnode1:racnode1-priv:racnode1-vip, racnode2:racnode2-priv:racnode2_vip |
| CLUSTER - Scan Info | Optional: The Oracle single Client Access name and port that you will use to reference this cluster without specifying a specific node. | rac-scan:1522 |
| Download Location | Required: The location where the CRS archive has been (or will be) downloaded. ¹ | /tmp |
| Extract Location | Required: The directory location where the CRS archive has been (or will be) extracted. | /tmp |
| Trust All Certificates | Optional: If this parameter is set to True, the workflow will trust any Secure Sockets Layer (SSL) certificate used to connect to the HP DMA web service. | True |
| Upgrade to Oracle 12 | Optional: Set to True if you are upgrading an existing Oracle 11g Grid Infrastructure to Oracle 12 Grid Infrastructure. The default is False. | False |
| Web Service Password | Optional: Password for the HP DMA Discovery web service API. | ●●● |
| Web Service User | Optional: User who is capable of modifying the managed environment by using the HP DMA Discovery web service API. | |

¹For additional information, see [Alternative methods for specifying input files.](#)

Workflow 2: Oracle - Provision Database Software



This section provides detailed information required to run the Oracle - Provision Database Software workflow.

Solution pack

This workflow requires the HP DMA Database Provisioning Solution Pack.

Parameters to expose

In the workflow's Gather Advanced Parameters for Provision Oracle Software step, expose the following parameters:¹

- Cluster Nodes
- runInstaller Parameters

Input parameters

When you deploy the Oracle - Provision Database Software workflow, specify input parameter values for the following steps.²

Note: Bold text in the following tables indicates that you must specify a value for the parameter.

Step: Gather Parameters for Provision Oracle Software

| Parameter | Description | Example Value |
|------------------------|---|---|
| Oracle Account | Optional (required only if inventory does not exist): The Oracle user that will own the Oracle Home. | oracle |
| Oracle Base | Optional (required only if Oracle Home has not been set up): The fully-qualified path to the Oracle base directory where the admin directories will be located. | /u01/app/oracle |
| Oracle Home | Optional (required only if Oracle Home has not been set up): The fully-qualified pathname where the Oracle Home will be created. If the specified directory does not exist, it will be created. | /u01/app/oracle/product/12.1.0/dbhome_1 |
| Oracle Software | Optional (required only if Oracle Home has not been set up): A comma-separated list of the Oracle Database software archive files. ³ | linuxamd64_12c_database_1of2.zip, linuxamd64_12c_database_2of2.zip |

¹These parameters are hidden by default and must be exposed when you make a copy of the workflow. For additional information, see [How to expose hidden parameters](#).

²HP DMA uses the latest (highest version number) of steps. For more information see [HP DMA Versions](#).

³If the file(s) are not found on the target server(s), they will be downloaded from the software repository. For additional information, see [Alternative methods for specifying input files](#).


Step: Gather Advanced Parameters for Provision Oracle Software

| Parameter | Description | Example Value |
|--------------------------------|--|---|
| Cluster Nodes | Optional (required when provisioning a RAC database): Comma-separated list of nodes to install software on. Leave blank for non-clustered environments. ¹ | racnode1, racnode2 <i>Use the same node values as the CLUSTER - Grid Node List parameter for the Oracle - Provision or Upgrade Grid Infrastructure deployment.</i> |
| Download Location | Optional: The directory where input files already exist or to which files will be downloaded from the software repository. ² | /tmp |
| Extract Location | Optional: The directory location where the Oracle database software archives will be extracted. It will be cleaned up at end of workflow execution. If not specified, a default will be created. | /tmp |
| runInstaller Parameters | Optional: The parameters to pass to the Oracle runInstaller command. For example: -force or -ignoreSysPrereqs ² | -ignoreSysPrereq |

¹This parameter is hidden by default and must be exposed when you make a copy of the workflow. For additional information, see [How to expose hidden parameters](#).

²For additional information, see [Alternative methods for specifying input files](#).

Workflow 3: Oracle - Provision Database



This section provides detailed information required to run the Oracle - Provision Database workflow.

Solution pack

This workflow requires the HP DMA Database Provisioning Solution Pack.

Parameters to expose

In the workflow's Gather Advanced Parameters for Provision Oracle Database step, expose the ASM Password parameter.¹

Input parameters

When you deploy the Oracle - Provision Database workflow, specify input parameter values for the following steps.²

Note: Bold text in the following tables indicates that you must specify a value for the parameter.

Step: Gather Parameters for Provision Oracle Database

| Parameter | Description | Example Value |
|--------------------------|--|---|
| Database Name | Required: The name of the database to provision. | orca |
| Datafile Location | Required: The database file locations. | +ASMDATA <i>Use a plus sign (+) followed by the same value as the ASM Diskgroup List parameter for the Oracle - Provision or Upgrade Grid Infrastructure deployment.</i> |
| Oracle Account | Optional (required only if inventory does not exist): The Oracle user that will own the Oracle Home. | oracle <i>Use the same value specified for the Oracle Account parameter for the Oracle - Provision Database Software deployment.</i> |
| Oracle Base | Required: The fully-qualified path to the Oracle base directory where the admin directories are located. | /u01/app/oracle <i>Use the same value specified for the Oracle Base parameter for the Oracle - Provision Database Software deployment.</i> |

¹This parameter is hidden by default and must be exposed when you make a copy of the workflow. For additional information, see [How to expose hidden parameters](#).

²HP DMA uses the latest (highest version number) of steps. For more information see [HP DMA Versions](#).

Step: Gather Parameters for Provision Oracle Database, continued

| Parameter | Description | Example Value |
|-----------------------------|---|--|
| Oracle Home | Optional: The Oracle Home to use if more than one Oracle Home exists in the inventory file(s). | <code>/u01/app/oracle/product/12.1.0/dbhome_1</code> <i>Use the same value specified for the Oracle Home parameter for the Oracle - Provision Database Software deployment.</i> |
| Trust SSL Certificates | Optional: If this parameter is set to True, the workflow will trust any Secure Sockets Layer (SSL) certificate used to connect to the HP DMA web service. | True |
| Web Service Password | Required: Password for the HP DMA Discovery web service API. | ●●● |
| Web Service User | Required: User who is capable of modifying the managed environment by using the HP DMA Discovery web service API. | dmawebuser |

Step: Gather Advanced Parameters for Provision Oracle Database

| Parameter | Description | Example Value |
|----------------------|---|--|
| ASM Password | Optional (required when provisioning an Oracle database using ASM storage): The password used to manage ASM. ¹ | <i>Use the same value specified for the ASM Password parameter for the Oracle - Provision or Upgrade Grid Infrastructure deployment.</i> |
| Cluster Nodes | Optional (required when provisioning a RAC database): Comma-separated list of nodes where this database will run. Leave blank for non-clustered environments. | racnode1, racnode2 <i>Use the same value specified for the Cluster Nodes parameter for the Oracle - Provision Database Software deployment.</i> |

¹This parameter is hidden by default and must be exposed when you make a copy of the workflow. For additional information, see [How to expose hidden parameters](#).

FAQs

What if I don't want to meet all of the prerequisites?

If you want to override the normal system prerequisites, do the following:

1. Set the Ignore System Prerequisites value to Y before you execute the deployment:

Workflow: Oracle - Provision or Upgrade Grid Infrastructure

Step: Gather Advanced Parameters for Provision Grid Infrastructure

Parameter: Ignore System Prerequisites¹

2. Set the runInstaller Parameters value to `-ignoreSysPrereqs` before you execute the deployment:

Workflow: Oracle - Provision Database Software

Step: Gather Advanced Parameters for Provision Oracle Software

Parameter: runInstaller Parameters¹

What if I get a failure on the Oracle installer that the ASM disk is not a valid device?

The raw device must exist and be owned by `oracle` (the user who owns the Oracle Home).

1. Verify that the raw device exists and is owned by `oracle`:

```
ls -l /dev/raw/raw1
```

2. Look for a line similar to the following:

```
crw-rw----. 1 oracle oinstall 162, 1 Mar 12 14:24 /dev/raw/raw1
```

3. If the device does not exist or is not owned by `oracle`, run the following commands:

```
raw /dev/raw/raw1 /dev/sdc1
chown oracle:oinstall /dev/raw/raw1
```

What if I get an error message that the `/u01/app` directory doesn't exist?

The `/u01/app` directory must exist and be owned by `oracle` (the user who owns the Oracle Home).

1. Create the directory:

```
mkdir /u01/app
```

2. Change the ownership:

```
chown -R oracle:oinstall /u01/app
```

¹This parameter is hidden by default and must be exposed when you make a copy of the workflow. For additional information, see [How to expose hidden parameters](#).

What if I get a failure in the install log related to the virtual IP address or SCAN virtual IP address?

If you get a failure related to virtual IP addresses, check the following:

1. Ensure that the virtual IP address and SCAN virtual IP address resolve in DNS and are in `/etc/hosts`
2. Ensure that your public name, virtual IP address, and SCAN virtual IP address are all in the same subnet.

What if I get an error referencing the public, private, or virtual IP addresses?

If HP DMA gives an error referencing the public, private, or virtual IP addresses, verify that the values specified for CLUSTER - Grid Node List are in the required format:

node1-public:node1-private:node1-virtual,node2-public:node2-private:node2-virtual

Workflow: Oracle - Provision or Upgrade Grid Infrastructure

Step: Gather Advanced Parameters for Provision Grid Infrastructure

Parameter: CLUSTER - Grid Node List

Appendix A: How to meet prerequisites

To fulfill the general requirements listed in [Prerequisites](#) on an Oracle Enterprise Linux version 6 (OEL6) platform, you can use the following instructions:

1. Add a disk that already exists:

```
raw /dev/raw/raw1 /dev/sdc1
```

2. Change ownership of the disk:

```
chown oracle:oinstall /dev/raw/raw1
```

3. Edit the `/etc/rc.local` file as root using the following example:

```
#!/bin/sh
#
# This script will be executed *after* all the other init scripts.
#
touch /var/lock/subsys/local
raw /dev/raw/raw1 /dev/sdc1
chown oracle:oinstall /dev/raw/raw1
```

4. Configure YUM to point to the OEL6 repositories.

5. YUM the Oracle RPM to meet the prerequisites—based on your Oracle version:

Oracle11.2 prerequisite:

```
yum install oracle-rdbms-server-11gR2-preinstall
```

Oracle12.1 prerequisite:

```
yum install oracle-rdbms-server-12cR1-preinstall
```

Note: Adapt these instructions as appropriate for other platforms.

Appendix B: Alternative methods for specifying input files

HP DMA provides alternative methods for specifying input files, either downloaded from the software repository or stored directly on the target server(s).

Method 1: The input files are in the software repository

1. Obtain the pertinent input files.

For example, the Oracle database software archive files.

2. Import the files into the software repository. For additional information, see [How to import a file into the software repository](#).
3. When you create the workflow deployment, specify the filenames and/or directory names for the pertinent input parameters—the files should not exist on the target server(s).
4. When the deployment is executed, HP DMA determines that the input files do not exist on the target server(s) in the specified (or default) download location and then downloads them from the software repository. If the input files need to be processed (for example, unzipped), the new files are placed in the specified (or default) extract location.
5. All downloaded, extracted, and staged files are removed upon successful completion of the workflow.

Method 2: The input files are stored on each target server

1. Obtain the pertinent input files.

For example, the Oracle database software archive files.

2. Copy the files to each target server into the specified (or default) download location.
3. When you create the workflow deployment, specify the filenames and/or directory names for the pertinent input parameters.
4. When the deployment is executed, HP DMA determines that the input files exist on the target server(s) in the specified (or default) download location. If the input file needs to be processed (for example, unzipped), the new files are placed in the specified (or default) extract location.
5. All extracted and staged files (and any downloaded files) are removed upon successful completion of the workflow.

Appendix C: How to import a file into the software repository

Many HP DMA workflows are capable of downloading files from the software repository on the HP DMA server to the target server(s) where the workflow is running. The following procedure shows how to import a file into the software repository so that it can be downloaded and deployed by a workflow.

HP DMA uses the HP Server Automation (SA) Software Library as its software repository.

Tip: Be sure to use unique file names for all files that you import into the software repository.

To import a file into the SA Software Library:

1. Launch the SA Client from the Windows Start Menu.

By default, the SA Client is located in Start → All Programs → HP Software → HP Server Automation Client.

If the SA Client is not installed locally, follow the instructions under “Download and Install the HP SA Client Launcher” in the *HP Server Automation Single-Host Installation Guide* available at: support.openview.hp.com/selfsolve/manuals

2. In the navigation pane in the SA Client, select Library → By Folder. Select (or create) the folder where you want to store the file. Click **Import Software**.
3. In the Import Software dialog, browse to the file (or files) to import, select the character encoding, and then click **Open**. For the Type, select **Unknown**. Select (or browse to) the folder where you want to store the files. For the Platform, select all the operating systems. Click **Import**.
4. Click **Close** after the import completes.

Appendix D: How to run an HP DMA workflow


The following steps outline the general instructions to run an HP DMA workflow:

1. Import the specified HP DMA solution pack (if it is not already available).
2. Create a deployable copy of the workflow: Specify a name, add roles, optionally expose parameters (for additional information, see [How to expose hidden parameters](#)), and then **Save**.
3. Create a deployment: Specify a name, schedule, targets, and input parameters, and then **Save**.
4. Execute the deployment to run the workflow.
5. View the results to verify that the workflow ran correctly (refer to the online workflow documentation).

Note: For more information about running HP DMA workflows, see the *HP DMA Quick Start Tutorial* available at: support.openview.hp.com/selfsolve/manuals

Appendix E: How to expose hidden parameters

If you need to expose any hidden parameters, perform the following additional steps when you create a deployable copy of the workflow:

1. Click the Workflow tab.
2. Click the blue arrow  next to the pertinent step to expand the list of input parameters.
3. For the parameter that you want to expose, select **- User selected -** from the drop-down list. For example:



4. Repeat steps 2 and 3 for all the parameters that you would like to expose.
5. Save the copy of the workflow.

To learn more about HP Database and Middleware Automation visit
go.hp.com/dma

© Copyright 2013-2014 Hewlett-Packard Development Company, L.P. The information contained herein is subject to change without notice. The only warranties for HP products and services are set forth in the express warranty statements accompanying such products and services. Nothing herein should be construed as constituting an additional warranty. HP shall not be liable for technical or editorial errors or omissions contained herein.

Oracle® is a registered trademark of Oracle and/or its affiliates.

July 2014

