Oracle® Fusion Applications

Cloning and Content Movement Administrator's Guide

11*g* Release 8 Refresh 7 (11.1.8)

E38322-18

July 2014

How to clone an on-premise, fully provisioned instance of
Oracle Fusion Applications and perform production-to-test
data movement.

ORACLE®

Oracle Fusion Applications Cloning and Content Movement Administrator's Guide, 11*g* Release 8 Refresh 7 (11.1.8)

E38322-18

"Licensor" shall mean the copyright owner or entity authorized by the copyright owner that is granting the License.

"Legal Entity" shall mean the union of the acting entity and all other entities that control, are controlled by, or are under common control with that entity. For the purposes of this definition, "control" means (i) the power, direct or indirect, to cause the direction or management of such entity, whether by contract or otherwise, or (ii) ownership of fifty percent (50%) or more of the outstanding shares, or (iii) beneficial ownership of such entity.

"You" (or "Your") shall mean an individual or Legal Entity exercising permissions granted by this License.

"Source" form shall mean the preferred form for making modifications, including but not limited to software source code, documentation source, and configuration files.

"Object" form shall mean any form resulting from mechanical transformation or translation of a Source form, including but not limited to compiled object code, generated documentation, and conversions to other media types.

"Work" shall mean the work of authorship, whether in Source or Object form, made available under the License, as indicated by a copyright notice that is included in or attached to the work

(an example is provided in the Appendix below).

"Derivative Works" shall mean any work, whether in Source or Object form, that is based on (or derived from) the Work and for which the editorial revisions, annotations, elaborations, or other modification represent, as a whole, an original work of authorship. For the purposes of this License, Derivative Works shall not include works that remain separable from, or merely link (or bind by name) to the interfaces of, the Work and Derivative Works thereof.

"Contribution" shall mean any work of authorship, including the original version of the Work and any modifications or additions to that Work or Derivative Works thereof, that is intentionally submitted to Licensor for inclusion in the Work by the copyright owner or by an individual or Legal Entity authorized to submit on behalf of the copyright owner. For the purposes of this definition, "submitted" means any form of electronic, verbal, or written communication sent to the Licensor or its representatives, including but not limited to communication on electronic mailing lists, source code control systems, and issue tracking systems that are managed by, or on behalf of, the Licensor for the purpose of discussing and improving the Work, but excluding communication that is conspicuously marked or otherwise designated in writing by the copyright owner as "Not a Contribution."

"Contributor" shall mean Licensor and any individual or Legal Entity on behalf of whom a Contribution has been received by Licensor and subsequently incorporated within the Work.

2. Grant of Copyright License. Subject to the terms and conditions of this License, each Contributor hereby grants to You a perpetual, worldwide, non-exclusive, no-charge, royalty-free, irrevocable copyright license to reproduce, prepare Derivative Works of, publicly display, publicly perform, sublicense, and distribute the Work and such Derivative Works in Source or Object form.

3. Grant of Patent License. Subject to the terms and conditions of this License, each Contributor hereby grants to You a perpetual, worldwide, non-exclusive, no-charge, royalty-free, irrevocable (except as stated in this section) patent license to make, have made, use, offer to sell, sell, import, and otherwise transfer the Work, where such license applies only to those patent claims licensable by such Contributor that are necessarily infringed by their Contribution(s) alone or by combination of their Contribution(s) with the Work to which such Contribution(s) was submitted. If You institute patent litigation against any entity (including across-claim or counterclaim in a lawsuit) alleging that the Work or a Contribution incorporated within the Work constitutes direct or contributory patent infringement, then any patent licenses granted to You under this License for that Work shall terminate as of the date such litigation is filed.

4. Redistribution. You may reproduce and distribute copies of the Work or Derivative Works thereof in any medium, with or without modifications, and in Source or Object form, provided that You meet the following conditions:

(a) You must give any other recipients of the Work or Derivative Works a copy of this License; and

(b) You must cause any modified files to carry prominent notices stating that You changed the files; and

(c) You must retain, in the Source form of any Derivative Works that You distribute, all copyright,

    patent, trademark, and attribution notices from the Source form of the Work, excluding those

    notices that do not pertain to any part of the Derivative Works; and

(d) If the Work includes a "NOTICE" text file as part of its distribution, then any Derivative Works that

    You distribute must include a readable copy of the attribution notices contained within such NOTICE

    file, excluding those notices that do not pertain to any part of the Derivative Works, in at least one

    of the following places: within a NOTICE text file distributed as part of the Derivative Works; within

the Source form or documentation, if provided along with the Derivative Works; or,

within a display generated by the Derivative Works, if and wherever such third-party notices

normally appear. The contents of the NOTICE file are for informational purposes only and do not

modify the License. You may add Your own attribution notices within Derivative Works that You

distribute, alongside or as an addendum to the NOTICE text from the Work, provided that such

additional attribution notices cannot be construed as modifying the License.

You may add Your own copyright statement to Your modifications and may provide additional or

different license terms and condition for use, reproduction, or distribution of Your modifications, or

for any such Derivative Works as a whole, provided Your use, reproduction, and distribution of the

Work otherwise complies with the conditions stated in this License.

5. Submission of Contributions. Unless You explicitly state otherwise, any Contribution intentionally submitted for inclusion in the Work by You to the Licensor shall be under the terms and conditions of this License, without any additional terms or conditions. Notwithstanding the above, nothing herein shall supersede or modify the terms of any separate license agreement you may have executed with Licensor regarding such Contributions.

6. Trademarks. This License does not grant permission to use the trade names, trademarks, service marks, or product names of the Licensor, except as required for reasonable and customary use in describing the origin of the Work and reproducing the content of the NOTICE file.

7. Disclaimer of Warranty. Unless required by applicable law or agreed to in writing, Licensor provides the Work (and each Contributor provides its Contributions) on an "AS IS" BASIS, WITHOUT WARRANTIES OR CONDITIONS OF ANY KIND, either express or implied, including, without limitation, any warranties or conditions of TITLE, NON-INFRINGEMENT, MERCHANTABILITY, or FITNESS FOR A PARTICULAR PURPOSE. You are solely responsible for determining the appropriateness of using or redistributing the Work and assume any risks associated with Your exercise of permissions under this License.

8. Limitation of Liability. In no event and under no legal theory, whether in tort (including negligence), contract, or otherwise, unless required by applicable law (such as deliberate and grossly negligent acts) or agreed to in writing, shall any Contributor be liable to You for damages, including any direct, indirect, special, incidental, or consequential damages of any character arising as a result of this License or out of the use or inability to use the Work (including but not limited to damages for loss of goodwill, work stoppage, computer failure or malfunction, or any and all other commercial damages or losses), even if such Contributor has been advised of the possibility of such damages.

9. Accepting Warranty or Additional Liability. While redistributing the Work or Derivative Works thereof, You may choose to offer, and charge a fee for, acceptance of support, warranty, indemnity, or other liability obligations and/or rights consistent with this License. However, in accepting such obligations, You may act only on Your own behalf and on Your sole responsibility, not on behalf of any other Contributor, and only if You agree to indemnify, defend, and hold each Contributor harmless for any liability incurred by, or claims asserted against, such Contributor by reason of your accepting any such warranty or additional liability.

END OF TERMS AND CONDITIONS

APPENDIX: How to apply the Apache License to your work.

To apply the Apache License to your work, attach the following boilerplate notice, with the fields enclosed by brackets "[]" replaced with your own identifying information. (Don't include the brackets!) The text should be enclosed in the appropriate comment syntax for the file format. We also recommend that a file or class name and description of purpose be included on the same "printed page" as the copyright notice for easier identification within third-party archives.

Copyright [yyyy] [name of copyright owner]

Licensed under the Apache License, Version 2.0 (the "License");

you may not use this file except in compliance with the License.

You may obtain a copy of the License at

  http://www.apache.org/licenses/LICENSE-2.0

Unless required by applicable law or agreed to in writing, software distributed under the License is distributed on an "AS IS" BASIS, WITHOUT WARRANTIES OR CONDITIONS OF ANY KIND, either express or implied.

See the License for the specific language governing permissions and limitations under the License.

Contents

Part I   Cloning Oracle Fusion Applications

1   Introduction to Cloning

2 Cloning Procedure

Part II   Production-to-Test Data Movement

3   Perform Production-to-Test Data Movement

Part III   Appendices

A   Post-Clone Data Masking (Optional)

B   Best Practices: Install with Cloning in Mind

C   Abstract Hostnames in Detail

D   Change JPS Root Name as Needed

E   Install the BI Administration Tool

Preface

The *Oracle Fusion Applications Cloning and Content Movement Administrator's Guide* is designed for on-premise installations of Oracle Fusion Applications. It contains two major sets of instructions: 1) how to clone an environment, and 2) how to copy just the data from one environment (such as production) to another (such as test).

Audience

The audience for this guide includes experienced Oracle Fusion Applications administrators who are familiar with their own enterprise-level installation of Fusion Applications. A variety of users and roles may be involved in creating a clone of a Fusion Applications instance, including: the database administrator, the LDAP/Identity Management administrator, users familiar with the specific products, such as the "General Accountant" user for the Financials product, and network administrators.

Documentation Accessibility

For information about Oracle's commitment to accessibility, visit the Oracle Accessibility Program website at
http://www.oracle.com/pls/topic/lookup?ctx=acc&id=docacc.

Access to Oracle Support

Oracle customers have access to electronic support through My Oracle Support. For information, visit
http://www.oracle.com/pls/topic/lookup?ctx=acc&id=info or visit
http://www.oracle.com/pls/topic/lookup?ctx=acc&id=trs if you are hearing impaired.

Related Documentation

This User Guide includes a companion workbook: *The Oracle Fusion Applications Discovery Workbook for Cloning and Content Movement*. This is an Excel-based entry form that is required for creating the cloning response file used by the cloning tools.

Conventions

The following text conventions are used in this guide:

Convention	Meaning
boldface	Boldface type indicates graphical user interface elements associated with an action, or terms defined in text or the glossary.
italic	Italic type indicates book titles, emphasis, or placeholder variables for which you supply particular values.
`monospace`	Monospace type indicates commands within a paragraph, URLs, code in examples, text that appears on the screen, or text that you enter.

Part I

Cloning Oracle Fusion Applications

This section describes the steps required to clone an existing Oracle Fusion Applications instance, creating an identical second version.

Part I contains the following chapters:

- Chapter 1, "Introduction to Cloning"
- Chapter 2, "Cloning Procedure"

1

Introduction to Cloning

This chapter contains the following high-level sections:

- Section 1.1, "Understanding Oracle Fusion Applications Cloning"
- Section 1.2, "Understanding Your Source Environment"
- Section 1.3, "System Requirements for Cloning"

1.1 Understanding Oracle Fusion Applications Cloning

The cloning discussed in this book refers to the complete duplication of an Oracle Fusion Applications source environment which was initially installed using the standard provisioning tools. Cloning creates a matching destination Fusion Applications environment with its own unique set of access URLs.

> **Note:** The procedures in this book are NOT designed for Oracle Fusion Applications systems that were *installed* using OVM templates. If your source system was installed in this way, contact Oracle Support for the correct cloning documentation and procedures.
>
> If you used virtualization technology, such as Oracle Virtual Machines, to host an operating system, but performed full standard provisioning into that virtualization layer, then the procedures in this book CAN be used.

1.1.1 What Is Cloning?

In the simplest terms, the cloning solution involves creating master image copies of the Oracle Fusion Applications components, mounting them to a matching array of destination hardware, and running the cloning tools (`faclone.sh`) to "rewire" the destination environment. The cloning tools also make the following necessary changes in the destination environment:

- Password changes
- Rewiring for new database configuration
- Cleanup of in-flight processes
- Rewiring for new Fusion Applications external URLs

In more detail, cloning is the process of duplicating the Fusion Applications file system structure (comprised of servers, storage, and optional load balancer/reverse proxy) onto a fresh destination environment, while making the changes necessary to the software and server configurations to accommodate a new set of URLs for access by

end users. Simultaneously, the destination environment must keep all the internal wiring intact, allowing the various Fusion Applications components communicate just as they did on the source.

Technically, this feat is accomplished using virtualized or abstract internal endpoints. This keeps the internal wiring intact on the destination. External endpoints for the destination (the user-accessed URLs) can be defined by the user and will be modified during the cloning process.

> **Note:** Cloning, as described in this document, relies on local name resolution using /etc/hosts, and details the steps needed for that type of name resolution.

1.1.1.1 Understanding Internal Wiring and Abstract Endpoints

What exactly is meant by "internal wiring"?
When Fusion Applications is installed, the Provisioning Wizard collects information, such as: host names for Fusion Applications, Identity Management, and databases, port numbers, load balancer information, internal virtual hosts and ports, etc. These details, along with system path information, metadata file names, and more, are used by the Provisioning Framework, and the Fusion Applications components themselves, to establish intercommunication. This collection of configurations equals the "internal wiring."

To keep this wiring intact on a new and different environment, the cloning tools rely on internal endpoint abstraction (abstract endpoints). (Clearly, if the endpoints were non-abstract but had a fixed 1:1 relationship with a physical server, they could not be used on multiple environments.) These abstract endpoints have the same name for all clones, but are defined individually, each pointing to its own environment.

During the clone phase (see Section 2.7), the clone tool also validates that the internal endpoints exist in /etc/hosts only.

If, in your existing source environment, you have placed internal names in DNS, they need to be removed. You then must add them to /etc/hosts on each node of the source, where they take over the name resolution activity that was formerly handled on DNS. You should also replace any abstract hostname entries in the source DNS. This is described in Section 2.3.1.

For more information on this topic, see Appendix C, "Abstract Hostnames in Detail".

1.1.1.2 What Does Cloning Include?

The destination environment clone will be an exact replica of the source, including:

- Complete topology
- Transactional and setup data
- Oracle Identity Management components and data
- Directory structures will be identical between source and destination
- Internal host names will be the same between source and destination

Cloning relies upon some fundamental design principles that are worth discussing. The basic assumptions that go into cloning are below:

- The Fusion Applications topology and release version will be identical between source and destination
- Relies on using /etc/host files for mapping

- The destination hardware must be capable of having everything brought up before any post clone activities are done. For example, the destination hardware can't be so much smaller that you can't start all JVM processes, without any load.

1.1.1.2.1 Component Review Reviewing the high-level components of a Fusion Applications instance can help you frame and understand how the terms are used in this guide. A Fusion Applications instance has two basic parts: 1) Identity Management, and 2) Oracle Fusion Applications itself (including the transaction database). The two images below show this in a simplified way:
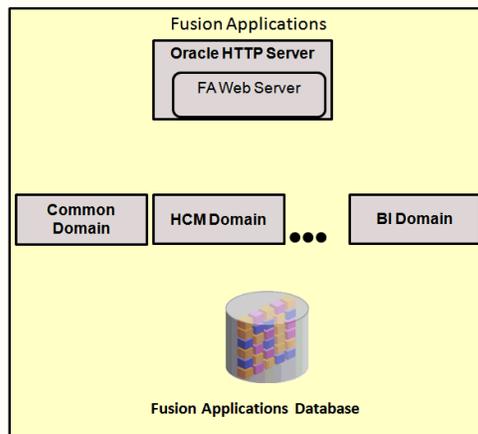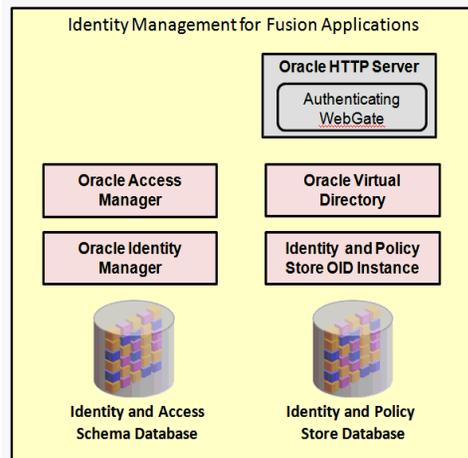
The two fundamental components are broken down as follows:

Oracle Identity Management for Fusion Applications

- **Oracle HTTP Server (OHS):** comprises the web tier for Identity Management

- **Oracle Access Manager (OAM):** located in the application tier for Identity Management

- **Oracle Identity Manager (OIM):** located in the application tier for Identity Management

- **Oracle Virtual Directory (OVD):** located in the application tier for Identity Management

- **Identity and Policy Store Oracle Internet Directory (OID) instance:** located in the application tier for Identity Management

- **Identity and Access database:** located in the data tier for Identity Management

- **Identity and Policy Store database:** located in the data tier for Identity Management (may or may not share the same database as Identity and Access database)

Fusion Applications Components:

- **Oracle HTTP Server (OHS):** comprises the web tier for Fusion Applications

- **Common Domain, HCM Domain, BI Domain...:** sit in the application tier for Fusion Applications.
  These domains represent the various product domains that are installed as a part of the normal provisioning process. Each domain contains the necessary and appropriate administration, WebLogic, and/or SOA servers as described in the installation documentation.

- **Oracle Fusion Applications database:** also called the transaction database; located in the data tier of Fusion Applications

1.1.1.3 What Changes are Permitted in Cloning?

The components listed in the Component Review are the only parts handled by the cloning process. If your environment contains additional components or extensions, any re-wiring of those must be done manually during the post-clone phase of the process.

Examples of things that the cloning process will NOT take into consideration are:

- OIF configuration (commonly used in solutions involving federation/SAML)

- Single-sign-on integrations

- HR2HR integration/co-existence

- Any code that invokes a URL external to Oracle Fusion Applications (also known as outbound messaging)

The following table describes the configurations that may change between the source and destination during cloning:

Component	Configuration	Can be changed during cloning?
Database (Identity Management and Fusion Applications)	SID	Yes
	Service Name	Yes
	Abstract Host Name	Yes
	Listener Port Number	Yes
	Install Paths	Yes
	Passwords	Yes
	Schema Names	No
	Non-RAC to RAC and vice-versa	No

Component	Configuration	Can be changed during cloning?
Oracle Identity Management	SSO External (login) URL and Port Number	Yes
	SMTP Email Server/Port	Yes
	Admin passwords	Yes
	Internal HTTP Endpoints	No
	Install Paths	No
	Abstract hostnames	No
	LDAP User, Group and `jpsroot` bases	No
	Admin usernames	No
Fusion Applications	External HTTP Endpoints (URLs/Ports)	Yes
	SMTP Email Server/Port	Yes
	Admin and APPID passwords	Yes
	Internal HTTP Endpoints	No
	Install Paths	No
	Topology must be the same as source, no consolidation	No
	Abstract hostnames	No
	Admin usernames	No
	Changes to installed product families, offerings or languages	No
	Obfuscation of transactional data	No

1.1.1.4 Terminology

Common terminology used in cloning includes:

Source Environment - An existing Fusion Applications environment that is already installed (and therefore contains the pre-requisite Identity Management for Fusion Applications) against which a copy or clone is to be made. Source environments can be in one of two states; 1) with transactional data or 2) without transactional data. An example of an environment with transactional data is one in which functional setup has been done and functional scenarios have been executed. (For example, orders were processed or an employee review cycle was completed). An example of an environment without transactional data is one in which the installation of Oracle Identity Management and the provisioning of Oracle Fusion Applications has been completed and validated, but no functional setup has been done and therefore no functional scenarios executed. The cloning process is the same in both cases.

Target Environment - The new environment that will be created as a result of the clone activity. Also called the destination environment.

Content Movement - Refers to the task of moving Fusion Applications components and/or data from one environment to another environment.

 Production to Test - A type of content movement that refers to refreshing all the transactional data in a destination environment by copying all the transactional data from a source environment

Physical Host Name -- Refers to the "internal name" of the current computer. On UNIX, this is the name returned by the `hostname` command.
Physical host name is used by Oracle Fusion Middleware to reference the local host. During installation, the installer automatically retrieves the physical host name from the current computer and stores it in the Oracle Fusion Middleware configuration metadata on disk.

Virtual Host Name -- Virtual host name is a network-addressable host name that maps to one or more physical computers via a load balancer or a hardware cluster. For load balancers, "virtual server name" is used interchangeably with virtual host name in this book. A load balancer can hold a virtual host name on behalf of a set of servers, and clients communicate indirectly with the computers using the virtual host name. A virtual host name in a hardware cluster is a network host name assigned to a cluster virtual IP. Because the cluster virtual IP is not permanently attached to any particular node of a cluster, the virtual host name is not permanently attached to any particular node either.

Abstract Host Name -- An abstract host name is an alias given to represent a physical node. It has a one-to-one relationship with a virtual host name. If your environment was installed before the release of cloning and done without the use of abstract host names, the virtual host names in your source environment will become abstract names in the destination environment. If your source environment did not make use of virtual host names, then physical host names will be used.

HTTP Endpoint - The entity on one end of an HTTP transport. An HTTP endpoint can be represented by an HTTP URL in the form of `http(s)://<hostname>:<port>`.

External HTTP Endpoint - In Fusion Applications, external HTTP endpoint define a set of URLs used by: 1) end users, to access Fusion Applications pages, and 2) by external systems and tools, to integrate with Fusion Applications.

Depending how the environment is configured, the endpoints may be located at the web tier (Oracle HTTP Server/OHS), or at a load balancer/reverse proxy (LB/RP). (In that case, the LB/RP are configured to forward the requests to the Oracle HTTP Server.)

Internal HTTP Endpoint - In Fusion Applications, internal HTTP endpoints define a set of URLs used by Fusion Applications itself for internal invocations and services. They are not used externally.

Depending how the environment is configured, the endpoints may be located at the web tier (Oracle HTTP Server/OHS), or at a load balancer/reverse proxy (LB/RP). (In that case, the LB/RP are configured to forward the requests to the Oracle HTTP Server.)

1.1.2 Use Cases: When Is Cloning Useful?

Cloning enables the copying of an existing Fusion Applications environment to a new set of servers and has many applications:

- Standing up a new environment that is a copy of the existing one for testing or development purposes

- Refreshing non-production systems with production data and production configurations/ binaries

- Using the gold-image source files as "templates" for creating multiple new environments

1.1.2.1 Other Content Movement Options and How They Differ

Standard database duplication is another kind of content movement, which is a subsection of cloning (duplicating the database structure and contents), but does not include the rest of the Fusion Applications instance.

Production-to-test content movement involves a data refresh from a production system back to a test environment. See Part II, "Production-to-Test Data Movement".

1.1.3 Roadmap: What Does the Cloning Process Entail?

Chapter 2 gives complete details on how to replicate a working Fusion Applications instance onto one or more additional environments. The process involves the following high-level steps:

- Understanding your source environment and the basic system requirements (Section 1.2 and Section 1.3)

- Performing an in-depth Discovery process regarding all aspects of the source and destination systems, collated into the detailed workbook provided (Section 2.2)

- Creating a "master image" of the source Oracle Fusion Applications components, using `.tar` or other tools (Section 2.3)

- Preparing the destination environment to receive the `.tar` copies (Section 2.4)

- Making the master images accessible to the destination (Section 2.5)

- Duplicating the databases from source to destination (Section 2.6)

- Running `faclone.sh` to extract and rewire components on the destination (Section 2.7)

- Performing validation steps (Section 2.8)

- Performing post-clone cleanup (Section 2.9).

- There are also two appendices: One on optional data masking steps (Appendix A) and one on how to install a new source environment so that the cloning process will be easier later (Appendix B)

1.2 Understanding Your Source Environment

Cloning presumes that you have a working version of Oracle Fusion Applications-- including the Fusion Applications database and Oracle Identity Management components-- installed, configured, and running on a source environment. Consider the following:

1.2.1 If Your Source Environment Has Already Been Installed

If your source environment was already installed before the cloning tools became available, then you likely did not make use of abstract host names. If no abstract names were used then the virtual host names in your source environment will become the abstract names in the destination environment. If the source environment did not make use of virtual host names, then physical host names will be used.

1.2.2 If Your Source Environment Was Installed and Virtualized

If your source environment was installed into a guest operating system, then the cloning process described in this guide still applies.

If the installation did not make use of abstract names, then the virtual host names in the source environment will become the abstract names in the destination environment. If the source environment did not make use of virtual host names, then physical host names will be used. Additionally, how .tar files of the source environment are created may differ, if you wish to clone the entire VM. This is a supported methodology, but details on copying the VMs are not provided in this guide.

When you have the copied VMs, the remainder of the process has only minor differences, detailed below:

- Copying VM images does not follow the .tar process described in this guide, and creation of the clone image files may not be necessary, as the VM image files may already include all the storage used in the environment.

- You must copy any shared storage to the destination environment separately.

- Since storage will be copied (either automatically, by copying the VM images, or manually) before the clone phase starts, the cloning response file must include the property:

   IMAGEFILE_EXTRACT=false.

   This disables automatic image file extraction during the clone phase. The clone tool will still perform all extraction verifications normally, so the response file entries for Storage (STORAGE_*) must still be completed; all other IMAGEFILE_* properties should be left blank.

1.2.3 If You Are Installing Your Source Environment Now

If you are installing your source Oracle Fusion Applications instance for the first time, you have the opportunity to plan ahead and follow best-practices during installation that will simplify cloning later. See Appendix B, for tips on how to optimize your installation process for cloning, using abstract host names and generic internal/external HTTP endpoints and LDAP endpoints.

1.3 System Requirements for Cloning

The source and the destination environments must meet the following system requirements.

> **Note:** The use of symbolic links is not supported in any tier (data-, middle-, or web tier) in either the source or the destination installation.

1.3.1 Source Requirements

Ensure that your source environment meets the following installation, storage, database, and component requirements, before beginning the cloning process. There are also requirements for the destination environment, below.

These requirements refer to the current version of the faclone.sh cloning scripts and tools.

1.3.1.1 Installation and Patch Requirements

Oracle Fusion Applications, in the correct release, was installed and configured, following the procedures in the *Oracle Fusion Applications Installation Guide*. (Check the

title of this document for the relevant version number of the software. To use this guide, the software and document versions must match.)

Important: Check the Release Notes for any required patches.

Additional installation requirements:

- **Operating system:** In this version of cloning, the source and destination environments must be installed on one of the following operating systems:

  - Linux (versions certified for Oracle Fusion Applications).
   - AIX version 7.1
   - Solaris 10

  (This requirement excludes the database, which can run on all supported OSs.)

- **Shells:** Ensure the following packages are installed on all host environments before starting the cloning procedure:
  ```
  -sh
  -bash
  -korn
  -gawk.
  ```

- **Local storage NOT supported:** The core Fusion Applications components must have been installed using shared storage. Using the local domain/storage option is not supported for Cloning.

  > **Note:**   The inventory must use shared storage or the environment cannot be cloned.
  > To check how the source environment was configured, open the `provisioning.rsp` file and search for `INSTALL_LOCALCONFIG_ENABLE`. If it is set to `=true`, then the system cannot be cloned.
  > For more information on finding and using the `provisioning.rsp` file, see Section 2.2.1.1, "Where to Find `provisioning.rsp` and `provisioning.plan`"

- **Oracle Identity Management:** The Identity Management stack must be fully contained on a single domain, and the application tier for Identity Management must be fully contained in a single node.

1.3.1.2 Source Storage Requirements

Table 1–1 shows the mandatory and optional requirements for both shared and local storage. Reminder: Provisioning of core Fusion Applications components must have been done without utilizing the local domain/storage option.

Table 1–1

Component	Shared	Local	DMZ
Database *Database duplication is currently manual, so storage options can be freely chosen.	Yes	yes	N/A
Identity Management	Yes Mandatory One root directory for Identity Management Shared Storage can be used for all IDM components (or at a minimum for the IDM Oracle_Homes, IDMDomain AdminServer and `oraInventory`	Yes Optional One root directory per Identity Management node, including DMZ	Yes Optional One root directory for Identity Management DMZ
Fusion Applications	Yes Mandatory One root directory for Fusion Applications Shared Storage can be used for all FA components (or at a minimum for the components that Provisioning installs on shared storage, in addition to the `oraInventory`	No Not supported If your source environment used local storage for Fusion Applications, the cloning tools cannot be used.	Yes Optional One root directory for Fusion Applications DMZ DMZ storage can contain the OHS Oracle_Home (and may also contain the OHS instance

1.3.1.3 Database Requirements

Database cloning involves two major categories of database: the Oracle Fusion Applications transaction database, and the Oracle Identity Management databases. Database duplication is done using RMAN or other duplication mechanisms. This is an independent process from copying the core Fusion Applications components, and has the following restrictions: 1) it must be done before the cloning scripts have been run to reconfigure the destination environment, and 2) no patching or other modifications are allowed after the environment has been frozen, to ensure that the database and Fusion Applications components remain synchronized. See Section 2.6 for details.

Note that it is not possible to selectively move data from source to destination; all data is moved. Also, during the Clone phase, in-flight transactional data, business process instances, queues, home page pop-up notifications, and jobs will be deleted. This is not optional, as the goal is to guarantee that no leftover processes access the source environment.

Successful database duplication requires the following conditions:

- Source and destination operating systems must be identical.

- Each database must be a physical copy of the source, not a logical copy created using data pump or other tools.

- Oracle Fusion Applications installation version and patch set must match requirements, and have been duplicated exactly on the destination, before database copying begins.

- Oracle Home must be registered in the destination environment's inventory.

- If RMAN is the preferred method, then the `RMAN DUPLICATE` command should be used.

- When the source database is a RAC database with multiple instances and the target is a non-RAC database or a single node RAC, specify the same database host name and port. The number of entries specified should match the number of instances that existed in the source database.

- If RAC is used, ensure that the Grid infrastructure is installed on the destination environment.

- The backup of the Fusion Applications database and the Oracle Identity Management database(s) must be a *cold* backup, with the source application down. (A backup taken while Oracle Fusion Applications is running could leave the destination environment inoperable.)

- Steps must be taken to fence off the destination server(s) and set the `job_queue` processes to 0 before starting RMAN duplication, to ensure that when the duplicate database is started it will not process jobs, and that the duplicate is isolated from the production systems.

1.3.1.4 Oracle Identity Management Component Requirements

- **Oracle Internet Directory (OID):** In a scaled-out/high-availability environment, the port numbers for all instances must be the same.

- **Oracle Virtual Directory (OVD):** In a scaled-out/high-availability environment, the port numbers for all instances must be the same, and only one OVD component is allowed per instance. (The component name is configurable, but must be the same for all instances.)

- **Endpoints:** For the Identity Management component, the internal endpoints (such as admin.mycompany.com) must use a different port from the external endpoints (such as sso.mycompany.com).

- **If the source environment was created using IDM scripts,** check the following:

  1. Check `jpsroot` in Identity Management for the CSF entry. If the CSF entry is not present, then:

  2. Go to the Enterprise Manager Cloud Control UI, and set the OVD Enterprise Manager user and password.

  3. Go to the `$OVD_INSTANCEHOME/bin` and call `opmnctl deregisterinstance` with the correct options.

  4. In Enterprise Manager, confirm the setting for OVD.

  5. In `LDAP`, confirm the CSF entry.

- **Policy and Credential Store Entries:** Check your source environment; the entry `cn=OVD,cn=CredentialStore,cn=IDMDomain,cn=JPSContext,cn=idm_jpsroot` must be present in the Policy Store. If it is not, then do the following:

  ```
  $OVD_INSTANCE/bin/opmnctl unregisterinstance

  $OVD_INSTANCE/bin/opmnctl registerinstance -adminHost
  adminvhn.mycompany.com -adminPort 7001 -adminUser weblogic
  ```

  The entry `cn=OVD` should also be present in the Credential Store.

1.3.1.5 Port Requirements

The clone environment will use the same ports as the source; therefore all ports used in the source environment must also be available on the destination.

Note: If any component is using privileged ports (< 1024), it becomes necessary to maintain `root` permissions, so the extraction of the master images on the destination environment must be performed manually by the customer as `root`.

In this case, you must set `IMAGEFILE_EXTRACT=false` on the `faclone.rsp`, so the clone tool will skip image file extraction. All other preparation and clone activities remain the same.

1.3.2 Destination Requirements

When planning the destination environment, keep the following requirements in mind:

1.3.2.1 Destination Server Requirements

- **Topology:** There is a direct mapping between source and destination environment topologies. The destination environment will have the same number of hosts/nodes and the same number of managed servers as existed in the source environment *prior* to any scale-out activities. For example, if the target environment distributed Fusion Applications across two hosts (meaning a portion of the WebLogic domains reside on one server and a portion on the other) then the target environment will also require two hosts to distribute the cloned Fusion Application across in matching fashion.

  The destination hardware cannot be so much smaller than the source so as to require a change in the topology of the primary server(s) in the source system. For example the RAM in the target environment cannot be so much smaller than the source's primary server that all of the managed servers defined in the target's primary servers cannot start up or run effectively.

- **Operating system:** Target and source nodes have the same operating system and version number.

- **Users:** The "Oracle" Operating System User that will own the destination Fusion Applications install is configured on all nodes with the same ID and Groups as the source.

- **Random Number Generator (rdng):** Because cloning does a great deal of encryption, it requires a large pool of random numbers. If this pool is depleted, the cloning process runs very slowly and appears to hang. To avoid this, rndg must be running on the destination environment before the cloning scripts are run.

- **SUDO password requirements (Linux only)**: If the source environment is scaled-out, the destination environment must have SUDO set up so that passwords are not required. For information on how to do this, see the section on "Setting Environment and Superuser Privileges for the wlsifconfig.sh Script" in the *Oracle Fusion Applications Installation Guide*.

1.3.2.2 Have JDK Installed on Destination So Clone Package Can be Extracted

This is a chicken-egg issue; the JDK resides on the source and is included in the source clone, yet it must be pre-installed on the destination to extract the clone. To manage this:

1. On the source environment, locate the repository and the JDK (usually located in `repository_location/jdk6`).

2. On the destination environment, create the `Clone_Home` directory and copy this JDK from the source into the `Clone_Home` directory.

> **Note:** If the GCC Linux Java implementation is already installed on the destination, it should not be used. Use the JDK as described above instead.

1.3.2.3 Target Storage Requirements

The storage setup on the destination environment must mirror the source, having:

- Same size requirements
- Same mount point
- Shared storage must be mounted to equivalent nodes

2

Cloning Procedure

This chapter contains the start-to-finish steps for cloning a source Oracle Fusion Applications instance onto a fresh destination environment.

2.1 Overview

Cloning Oracle Fusion Applications requires the following steps:

- **Discovery:** An in-depth notation of all aspects of the source and destination topology and configuration details, with entries typed into the *Discovery Workbook for Cloning and Content Movement* provided. See Section 2.2.

- **Creating a Master Image of the Source:** This is a manual process. If you have virtualized your installation using Oracle Virtual Machine, then you can create a master image of the entire VM. (This guide does not describe that procedure.) For standard Fusion Applications installations, this is done using `.tar`. See Section 2.3.

- **Preparing the Target Environment:** Though this step can be done out of sequence (i.e. before creating the master image), it clearly must be done before copying the images onto the destination. See Section 2.4.

- **Making Master Images Accessible to Target:** See Section 2.5.

- **Duplicating the Database(s) from Source to Target:** This is a self-contained process that must occur before running the `faclone.sh` scripts. See Section 2.6.

- **Running the faclone.sh Script:** This will "wire together" and configure the entire destination instance of Oracle Fusion Applications. See Section 2.7.

- **Performing Validation Steps:** See Section 2.8.

- **Performing Post-Clone Cleanup:** See Section 2.9.

- Appendix A, "Post-Clone Data Masking (Optional)" can be performed.

- Appendix B, "Best Practices: Install with Cloning in Mind" gives tips for installing a new source environment in such a way that cloning will be streamlined.

- Appendix C, "Abstract Hostnames in Detail" gives more theoretical understanding of the naming conventions, if needed.

- Appendix E, "Install the BI Administration Tool" gives the installation steps, in case the BI Administration tool has not yet been installed.

2.2 Discovery

The discovery phase may be the most important part of the cloning process. Here you determine all the details of your source environment (existing or planned), as well as your destination environment, and record them.

2.2.1 Using the Discovery Workbook

The *Oracle Fusion Applications Discovery Workbook for Cloning and Content Movement* is a required companion document to this User Guide. It is used to help you research and annotate every aspect of your source and destination Fusion Applications environments. **Fill in all the tabs in the Workbook that are preceded by C_; these are cloning-related.** You will then copy/paste the entries to complete the cloning response file (`faclone.rsp`) appropriately.

2.2.1.1 Where to Find `provisioning.rsp` and `provisioning.plan`

The best resource for many of the Workbook entries is the `provisioning.rsp` file. For some data, it is also necessary to refer to `provisioning.plan`.

Both files may be located in the same directory:
(`APPLICATIONS_BASE/provisioning/plan/`).
If the `.rsp` file is not in the `/plan` directory, search for
`provisioning.setup.choice.core.provisionplan.install` within
`provisioning.plan`, to see where the `.rsp` file is located.

2.2.2 Prepare for the Discovery Phase

The Workbook gives some shorthand tips on where to find things or how to enter them, but this section of the User Guide provides much more guidance.

To begin, open the Discovery Workbook and proceed through the tabs of data you are asked to collect. (Cloning-related tabs are preceded by **C_**.) They are organized as follows:

- Section 2.2.3, "General Info"

- Section 2.2.4, "Environment"

- Section 2.2.5, "Topology"

- Section 2.2.6, "Storage"

- Section 2.2.7, "WebLogic Server (WLS) Domains"

- Section 2.2.8, "Virtual Hosts"

- Section 2.2.9, "Databases"

- Section 2.2.10, "Identity Management"

- Section 2.2.11, "Email and Business Intelligence (BI)"

- Section 2.2.12, "Ports"

- Section 2.2.13, "Passwords"

The last tab is special; it automatically collates the data from the rest of the tables and organizes them for ease of use in the `faclone.rsp` file. It is:

- Section 2.2.14, "Generated RSP Entries"

2.2.3 General Info

The System Administrator responsible for Oracle Fusion Applications should know the values of the properties in this section, without having to delve too deeply into the environment itself.

The General Information table provides the baseline data for the Fusion Applications instance to be cloned.

- **Company Name:** Self-explanatory and the same for both source and destination. This entry is not officially used within the response files.

- **Fusion Applications Version:** Will be the same for source and destination.

- **Operating System:** Note the OSs used for the databases, the Middleware components, front-end servers (web servers, load balancers, Oracle HTTP Server (OHS)), and any ODI clients.

- **Environment Name:** This is the naming convention you use to identify your installation. It's for human use only; not a name that is registered in the system. Examples could be "QA1" or "Production".

- **Environment Type:** Like Environment Name, this is for your own understanding only. Annotate, for example, whether the installation is physical or virtual (i.e. using VMWare or another method).

- **HA/Scaled-Out:** By default, only one instance of each Middleware component is provisioned during a standard installation. Indicate whether the source environment has been scaled out beyond this for high-availability or not. Enter the value "`true`" or "`false`"; it will later be used in the `faclone.rsp` file for the property `FA_TOPOLOGY_SCALEOUT`.

- **Includes Customer Data:** Set to TRUE if the source contains transaction data and FALSE if not. The flag alerts the cloning tool to execute some required additional data cleanup when transactional data is present.

- **Includes Customizations:** If your enterprise has extended or altered the core Fusion Applications installation, those changes will have to be manually replicated after the cloning process is complete.

- **HTTP Domain:** Sample entry: `mycompany.com`. Enter the domain portion of the external URL for the source instance, and the one that will be used on the destination.

  Enter the value for the destination domain in the `faclone.rsp` file, for the property `FA_WEBTIER_DOMAIN_TARGET`.

- **LBR/Reverse Proxy Used for Internal HTTP Endpoints:** `TRUE` or `FALSE` value describing whether a load balancer/reverse proxy is used for internal HTTP endpoints. Check `provisioning.rsp` for the `LBR_ENABLED` entries to confirm, if necessary. (NOTE: If your organization changed the use of LBR/Reverse Proxy sometime after provisioning, then the `.rsp` file entry would not be accurate and you would need to consult the network administrator for details.)

- **LBR/Reverse Proxy Used for External HTTP Endpoints:** `TRUE` or `FALSE` value describing whether a load balancer/reverse proxy is used for external HTTP endpoints. Check `provisioning.rsp` for the `LBR_ENABLED` entries to confirm, if necessary. (NOTE: If your organization changed the use of LBR/Reverse Proxy sometime after provisioning, then the `.rsp` file entry would not be accurate and you would need to consult the network administrator for details.)

2.2.4 Environment

The OS Commands table is automatically populated based on the entry you provided for Operating System on the **General Info** tab. The parameters are required for the cloning tool, but do not need to be manually modified unless your system values differ from the defaults.

- **AWK:** The path to the `awk` tool in your system

- **EGREP:** The path to the `egrep` command in your system

- **ID:** The path to discover the `user ID` command in your system

- **IFCONFIG:** The path to the `ifconfig` utility in your system

- **MKNOD:** The path to the `mknod` command in your system,

- **OPENSSL:** The path to the OpenSSL toolkit in your system

- **SED:** The path to the stream editor utility in your system

2.2.4.1 User and Group for `OHS httpd.conf`

- **OS_FA_USER_TARGET:** The common default for this is `<oracle>`.
  To confirm the setting in your system, log in as the Unix user who owns the HTTP server for Fusion Applications, and type in the unix command : `id`.
  This will give you the correct user and group for that user.

- **OS_FA_GROUP_TARGET:** The common default for this is `<oinstall>`.
  To confirm the setting in your system, log in as the Unix user who owns the HTTP server for Fusion Applications, and type in the unix command : `id`.
  This will give you the correct user and group for that user.

- **OS_IDMOHS_USER_TARGET:** The common default for this is `<oracle>`.
  To confirm the setting in your system, log in as the Unix user who owns the HTTP server for IDM (Oracle Identity Management), and type in the Unix command : `id`. This will give you the correct user and group for that user.

- **OS_IDMOHS_GROUP_TARGET:** The common default for this is `<oinstall>`. To confirm the setting in your system, log in as the Unix user who owns the HTTP server for IDM (Oracle Identity Management), and type in the Unix command : `id`. This will give you the correct user and group for that user.

2.2.5 Topology

The Topology information is obtained through a multi-step discovery process, and depends on which domain configuration was used during the source provisioning. The midtier component information, the database information, and the Oracle Identity Management information are each discovered separately. This section includes the following:

- Section 2.2.5.1, "Topology: Find Source MidTier Abstract Host Name(s)"

- Section 2.2.5.3, "Topology: Find Source (Real) Host Names, IP Addresses, and Users"

- Section 2.2.5.2, "Topology: Find Database Abstract Host Name"

- Section 2.2.5.4, "Topology: Find Oracle Identity Management Information"

- Section 2.2.5.5, "Topology: Give Node Numbers"

- Section 2.2.5.6, "Component Assignment: Associate Nodes with Components"

- Section 2.2.5.7, "Complete Topology Table for Target"

■ Section 2.2.5.8, "Blacklisted Hosts"

> **Note:** Abstract hostnames are an important part of the cloning process and can be confusing. This section provides guidance on finding and entering abstract hostnames; for more context and examples, see Appendix C, "Abstract Hostnames in Detail".

Multiple Network Cards: Be aware that in environments with multiple network cards, only one will be configured during provisioning. Before running the cloning phase on the host with multiple network cards, the hostname that was used during provisioning must be set temporarily.

To set the hostname temporarily, issue this command as root:

```
$ hostname <name_used_during_provisioning>
```

This command does not store this change permanently and the original hostname will be restored after a reboot.

2.2.5.1 Topology: Find Source MidTier Abstract Host Name(s)

During the initial install/provisioning of Oracle Fusion Applications, the Fusion Applications *midtier* host names were defined in the Domain Topology Configuration screen. (See "The Domain Topology Configuration" in the *Oracle Fusion Applications Installation Guide*, for a refresher, if desired. See also "Appendix D," in the same guide for details about the screen, copied below.)

The Provisioning Wizard screen that lists the domain topology during installation.

****

Note the three installation options, which correspond in the `provisioning.rsp` file to three levels: **Basic**, **Medium**, and **Advanced**.

- **Basic** = **One host for all domains**. To find out if your source environment used this configuration, search `provisioning.rsp` for `#Domain Topology` and check whether `TOPOLOGY_BASIC_SELECTED= TRUE`.

  If it does, then the `TOPOLOGY_BASIC_HOST` (such as `fusionapps1.mycompany.com`) listed immediately thereafter is the **Abstract Host Name** for your entire midtier environment.

- **Medium** = **One host per domain**. To find out if your source environment used this configuration, search `provisioning.rsp` for `#Domain Topology` and check whether `TOPOLOGY_MEDIUM_SELECTED= TRUE`

  If it does, then the `TOPOLOGY_MEDIUM_COMMONDOMAIN_HOST=` listed immediately thereafter gives a path name that refers to the `provisioning.plan` file. For example:`${provisioning.setup.common.core.default.host.name}`. Follow these steps to search `provisioning.plan` and fill in the Abstract Host Name for each environment domain:

  1. Open `provisioning.plan`.

  2. Search for `${provisioning.setup.common.core.default.host.name}`.

  3. Take note of the domain names listed thereafter. In a **Medium** configuration, for example, all the `CRM` components will have the same domain assignment (such as `crmdomain1`), all the `HCM` components to another, and so on. Enter these domains as source Abstract Host Names in the Topology table

- **Advanced** = **One host per application and middleware component.** To find out if your source environment used this configuration, search `provisioning.rsp` for `#Domain Topology` and check whether `TOPOLOGY_ADVANCED_SELECTED= TRUE`.

  If it does, then all the components and products will be listed thereafter with either an Abstract host name, or a reference to `provisioning.plan` where you can search out the Abstract name.

  See a partial sample of an Advanced `provisioning.rsp` file below. Note, for example, that the value for `TOPOLOGY_ADVANCED_COMMON_ADMIN_HOST` is the Abstract name `commonhost01.mycompany.com`. The value of `TOPOLOGY_ADVANCED_ CRM_ANALYTICS` must be found by searching `provisioning.plan` for `${provisioining.setup.crm-hosts.domain.hostname}`.

  Sample (Partial) Advanced Provisioning.rsp File Entries:

```
TOPOLOGY_ADVANCED_SELECTED=true
TOPOLOGY_ADVANCED_COMMON_ADMIN_HOST=commonhost01.mycompany.com
TOPOLOGY_ADVANCED_COMMON_FUNCTIONALSETUP_
HOST=${provisioning.setup.fs-hosts.domain.hostname}
TOPOLOGY_ADVANCED_COMMON_HELPPORTAL_
HOST=${provisioning.setup.fs-hosts.domain.hostname}
TOPOLOGY_ADVANCED_COMMON_HOMEPAGE_
HOST=${provisioning.setup.fs-hosts.domain.hostname}
TOPOLOGY_ADVANCED_COMMON_ESS_
HOST=${provisioning.setup.fs-hosts.domain.hostname}
TOPOLOGY_ADVANCED_COMMON_OWLCS_
HOST=${provisioning.setup.fs-hosts.domain.hostname}
TOPOLOGY_ADVANCED_COMMON_OWLCSSIP_
HOST=${provisioning.setup.fs-hosts.domain.hostname}
TOPOLOGY_ADVANCED_COMMON_IPM_
HOST=${provisioning.setup.fs-hosts.domain.hostname}
TOPOLOGY_ADVANCED_COMMON_SES_
```

```
HOST=${provisioning.setup.fs-hosts.domain.hostname}
TOPOLOGY_ADVANCED_COMMON_SOA_
HOST=${provisioning.setup.fs-hosts.domain.hostname}
TOPOLOGY_ADVANCED_COMMON_UCM_
HOST=${provisioning.setup.fs-hosts.domain.hostname}
TOPOLOGY_ADVANCED_COMMON_WCSERVICES_
HOST=${provisioning.setup.fs-hosts.domain.hostname}
TOPOLOGY_ADVANCED_COMMON_WCSPACES_
HOST=${provisioning.setup.fs-hosts.domain.hostname}
TOPOLOGY_ADVANCED_CRM_ADMIN_HOST=crmhost01.mycompany.com
TOPOLOGY_ADVANCED_CRM_ANALYTICS_
HOST=${provisioning.setup.crm-hosts.domain.hostname}
TOPOLOGY_ADVANCED_CRM_CONTRACT_
HOST=${provisioning.setup.crm-hosts.domain.hostname}
...
```

Search your own `provisioning.plan` file and note each Abstract host name, as needed, in the Topology table.

In the destination (target) environment, the names assigned to the source Domain Topology Configuration become the destination Abstract host names and will be added to the `/etc/hosts` files on all clone nodes.

When all the Abstract host name information is entered for the midtier, then find it for the database(s) and Oracle Identity Management components. See Section 2.2.5.2 and Section 2.2.5.4.

2.2.5.2 Topology: Find Database Abstract Host Name

In provisioning.rsp, search for `#Database Configuration` and `Database Host:=` to find the abstract host name. If RAC is used, then enter additional lines in the workbook to add the correct number of RAC instances and their names.

2.2.5.3 Topology: Find Source (Real) Host Names, IP Addresses, and Users

Ping each Abstract Host Name that was identified in sections Section 2.2.5.1 and Section 2.2.5.2. This will reveal the "real" host names and IP addresses. Enter this information in the Workbook.

To find the assigned users, go to the Fusion Applications directory, and use list (`ls -l`) to find the users assigned to each of the directories. Enter all values into the Workbook.

2.2.5.4 Topology: Find Oracle Identity Management Information

Access the Oracle Identity Management information through Fusion Middleware Control user interface (FMW). See sample, below. Enter the values in the Workbook.
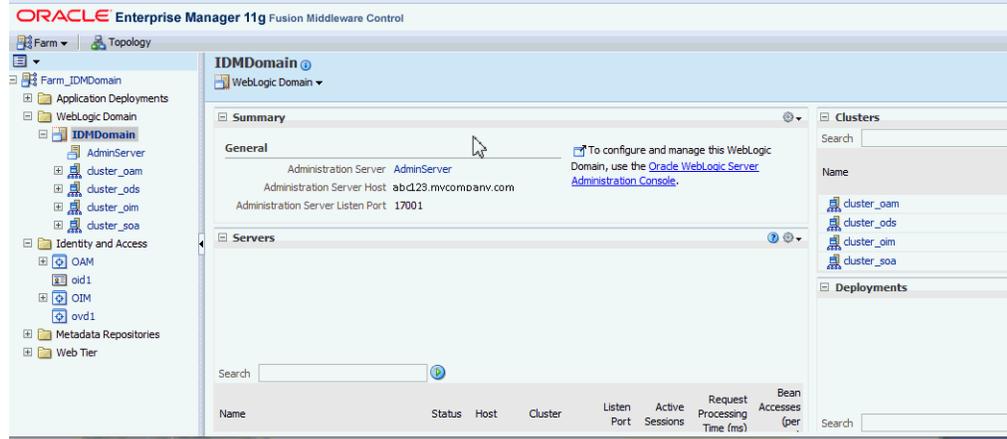
FMW Control page with details about Identity Management Domain, such as Administrator, Host, and Port.

****

2.2.5.5 Topology: Give Node Numbers

When all the host information is entered in the Topology table, then number each entry consecutively in the **Node #** column on the left. These node associations will be used to define which Oracle Fusion Applications components are installed/associated with which source hosts. See Section 2.2.5.6.

2.2.5.6 Component Assignment: Associate Nodes with Components

The Component Assignment table lists all Oracle Fusion Applications products/components that could possibly be installed in an environment. For each component that is relevant for your installation, deduce the abstract host name on which that component resides. (Ideally, best-practice naming conventions were used, so that components and host names have a logical relationship, but do not include environment-specific tags such as "test," "production," "staging," or similar.)

In a core installation, only one node (abstract host) is installed per component. If your source installation was subsequently scaled out over multiple hosts, then enter the additional node numbers in a comma-separated list in the HA-Scale-out Node column. See an example in the sample tables, below.

Sample Topology Table (Partial)

Node #	Abstract Host Name
1	fusionapps1.mycompany.com
2	faweb1.mycompany.com
3	fusionidm.mycompany.com
4	sso.mycompany.com
5	fadb1.mycompany.com
6	fadb2.mycompany.com
7	idmdb.mycompany.com
8	fusionapps2.mycompany.com
9	faweb2.mycompany.com

Sample Component Assignment Table (Partial)

Node #	Component	HA-Scale-out
5	FADB	6
7	IDMDB	
7	OIDDB	
3	OIDINST	
3	IDMOHS_INST	
4	IDMDOMAIN_ADMIN	
3	IDMDOMAIN_SOA (...)	
2	FA_WEBTIER	9
1	COMMON_ADMIN	8

2.2.5.7 Complete Topology Table for Target

Enter the information for the planned destination environment:

- **Target Host Name:** The (real) host names of the machines to be used in the destination environment. Note that you will map abstract host names onto these machines in Section 2.4.1.

- **Target User:** Enter planned target user names. They may or may not match the source.

- **Target IP Address:** Enter the IP addresses that are planned in the destination environment for each component.

2.2.5.8 Blacklisted Hosts

Determine whether there are any connections to servers other than the servers running Fusion Applications components. Examples could include:

- A mis-configuration using an alternate host name, to connect to a Fusion Applications host

- External integrations, such as Active Directory or other Business Intelligence resources

- Database Links

- Customization

Because these endpoints could potentially lead to a production system, enter them into the Blacklist so that the `faclone.sh` tool will stop any unintended connections. (You will then, in post-clone, have to update these connections manually to point to alternate resources.)

2.2.5.8.1   ODI Host Blacklist

The script `ODI_blacklist_hosts.sql`, found in the `/discover/scripts` directory, will generate a list of non-Fusion Application database sources or targets that have been created in the source ODI environment. These need to be added to the blacklist to prevent accidental connections from the cloned environment.

Connect to the Fusion Database as the SYS user and run the script.

This will generate a list of host names that should be added to the Topology Worksheet in the Discover Workbook, under the Blacklist section.

If no additional database sources have been set up previously, the script will not generate an output.

2.2.5.8.2 Post-Clone Steps for Oracle Identity Federation (OIF)

Depending on your environment, you may need to change some settings of the OIF Federations, as follows:

1. Point your browser to the IDMDomain Fusion Middleware Control (also known as IDMDomain Enterprise Manager). Example: `http://adminserverhost:port/em`

2. Login with your IDMDomain credentials.

3. Expand `Farm_IDMDomain`.

4. Expand **Identity and Access**.

5. Select **OIF**. (This may have a longer name, but starts with OIF.)

6. Open the menu **Oracle Identity Federation**.

7. Select **Administration > Federations**.

8. Update all URLs of the trusted providers.

9. Select **Administration > Data Stores**.

10. Update the settings of the Federation Data Store.

2.2.6 Storage

The storage tables are divided into Shared storage, DMZ, and Node storage. Note that not all entries may be relevant for your particular installation. The default values shown in the Workbook are the most common and are entered for convenience, but can be edited if your environment differs.

2.2.6.1 Image Files

- **IMAGEFILE_EXTRACT:** Extraction of image files required. This is set to true or false.

- **IMAGEFILE_POOL_DIR:** Directory location of all of the image files created in preparation for cloning.

- **IMAGEFILE_LOCAL_1:** Local directory for runtime (domain) configuration (optional, local storage not supported at this time).

- **IMAGEFILE_LOCAL_2:** Local directory for runtime (domain) configuration (optional, local storage not supported at this time).

- **IMAGEFILE_SHARED_FA:** Location of Fusion Applications image file.

- **IMAGEFILE_DMZ_FA:** If a DMZ is present, this will be the location of the FA Image files and the field IMAGEFILE_SHARED_FA should be blank.

- **IMAGEFILE_SHARED_IDM:** Location of IDM image file >

- **IMAGEFILE_DMZ_IDM:** If a DMZ is present, this will be the location of the IDM Image files and the field IMAGEFILE_SHARED_IDM should be blank.

- **IMAGEFILE_UNTAR_COMMAND:** The command utilized by your system to untar a file - the default is `"tar -xvpf"`

2.2.6.2 Shared Storage

This section describes how to fill the columns, not the rows.

- **Abstract Mount Point (FA Shared):** To find this, search `provisioning.rsp` for `#Installation Locations` and find `INSTALL_APPHOME_DIR`. The abstract mount point is usually the directory that contains `INSTALL_APPHOME_DIR` and is defined in `/etc/fstab`.

- **Abstract Mount Point (IDM Shared):** Check the FMW Control.

- **Source User (FA Shared and IDM Shared):** Enter the OS user that owns the mount point.

- **Source Group (FA Shared and IDM Shared):** Enter the OS group that owns the mount point.

- **Source Real Mount Point:** This field is for informational purposes only and is not used during Cloning. Use it to provide an alternative mount point (to be used for administrative purposes; will be different between source and target) or to provide the remote NFS directory that corresponds to the abstract mount point at the source.

- **Target Real Mount Point:** This field is for informational purposes only and is not used during Cloning. Use it to provide an alternative mount point (to be used for administrative purposes. it will be different between source and target) or to provide the remote NFS directory that corresponds to the abstract mount point at the target.

2.2.6.3 DMZ Storage

To discover whether DMZ is used, search the `provisioning.rsp` file for `WEBTIER_DMZINSTALL_ENABLE=`. If `TRUE`, then the DMZ mount point information can be found by going to the DMZ host listed, and checking the `oraInventory` there.

2.2.6.4 Node Storage

Copy the Node #s from the **Topology** table and enter the relevant storage information in each column.

- **Local Storage Directory:** Leave row blank. Source must have been provisioned without using local storage.

- **Uses FA Shared Storage?:** Enter `TRUE` or `FALSE` for each node.

- **Uses IDM Shared Storage?:** Enter `TRUE` or `FALSE` for each node.

- **Uses FA DMZ Shared Storage?:** Enter `TRUE` or `FALSE` for each node.

- **Uses IDM DMZ Shared Storage?:** Enter `TRUE` or `FALSE` for each node.

2.2.6.5 Directories

Below are tips on how to find the Full Paths for each relevant directory:

- **IDM-related entries:** Log in to the FMW Control for IDM and go through the Topology for each component. The Oracle Homes and Instance Homes for each IDM component are listed and can be entered into the Workbook.
  Fields should be left blank if not used.

  `IDM_INSTALL_APPHOME_DIR:` This is only used if IDM was installed as part of provisioning (Rel 7 or later). This is your products directory similar to `/u01/oracle/products`

  `IDM_INSTALL_APPCONFIG_DIR:` This is only used if IDM was installed as part of provisioning (Rel 7 or later). This is your config directory similar to

```
/u01/oracle/config
```

IDM_INSTALL_LOCALCONFIG_DIR: This is only used if IDM was installed as part of provisioning (Rel 7 or later). This is your local config directory similar to `/u01/oracle/config`

- **FA Base, Home, and OHS/Webtier:** Search `provisioning.rsp` for INSTALL_APPHOME_DIR (Base, and OHS/webtier), and INSTALL_APPCONFIG_DIR (Instance Home), which then includes the Domains as subdirectories. If the INSTALL_LOCALCONFIG_DIR entry contains values, indicating that local storage is used, then the OHS/webtier would be located there instead. In this release, local storage is not supported for cloning.

- **FA BI Instance (when used):** If Business Intelligence (BI) is used, the directory path is either the LOCALCONFIG_DIR or INSTALL_APPCONFIG_DIR, followed by `/BIInstance`.

- **IDM_KEYSTORE_DIR:** The value can have two different options, depending on the install type.
  If the environment was installed using the IDM Provisioning tool, the default value is : `${IDM_INSTALL_APPCONFIG_DIR}/keystores`.
  This is only applicable to environments that were originally provisioned on Release 7 or later.

  If the environment was installed following the EDG option, the default value is : `empty`.

  The easiest way to find out the correct value is to check the source environment to see where it is located.

- **IDM_NODEMANAGER_DIR:** The value can have two different options depending on the install type.

  If the environment was installed using the IDM Provisioning tool, the default value is : `${IDM_INSTALL_APPCONFIG_DIR}/keystores`. This is only applicable to environments that were originally provisioned on Release 7 or later.

  If the environment was installed following the EDG option, the default value is : `${IDM_OIM_MWHOME}/wlserver_10.3/common/nodemanager`.

  The easiest way to find out the correct value is to check the source environment to see where it is located.

2.2.6.6 Inventories

Use the `orainst.loc` files on each Oracle Home, and search to find the respective `oraInventory`.

2.2.7 WebLogic Server (WLS) Domains

This tab includes information about the Admin and Managed servers in both the Fusion Applications (FA) and Identity Management (IDM) components.

2.2.7.1 Fusion Applications Domains

This section describes the columns rather than the rows.
Provide the Admin Server path and port for each Fusion Applications domain.

■ **AdminServerPath:** This is the path to the domain directory, in the format
`<FA Instance home/domains/<abstract host name of the topology component>/Domain name>`.
For example, if the instance home is `u01/app/fa/instance`, and the abstract hostname for COMMON Admin is `fusionapps.mycompany.com`, then the Admin Server path for Common Domain would be:
`u01/app/fa/instance/domains/fusionapps.mycompany.com/CommonDomain`

■ **AdminServerPort:** This can be found in `provisioning.rsp` under `TOPOLOGY ADVANCED <topology component> PORT`. For example:

```
TOPOLOGY_ADVANCED_COMMON_ADMIN_PORT=17001
TOPOLOGY_ADVANCED_CRM_ADMIN_PORT=19001
TOPOLOGY_ADVANCED_FIN_ADMIN_PORT=17401
TOPOLOGY_ADVANCED_HCM_ADMIN_PORT=19401
```

2.2.7.2 Identity Management (IDM) Domains

Provide the Admin Server path and port for the Identity Management domains as well.

■ **AdminServer Path:** The path to the domain home can be found from Fusion Middleware Control.

■ **AdminServer Port:** This can be found in provisioning.rsp under IDENTITY OIM ADMINPORT.

2.2.7.3 Identity Management (IDM) Managed Servers

Provide the server name prefix for each type of managed server. The server name is found in the WebLogic (WLS) Console under Servers. For example: If the WLS Console shows the server name for OAM is `wls_oam1`, then enter `wls_oam` in the Workbook as the prefix for OAM.

2.2.8 Virtual Hosts

This tab is set up assuming that different virtual hosts are used for Fusion Applications core components, and for Identity Management components.

2.2.8.1 FA HTTP Endpoints

The Internal Names for both source and destination are the same. To find most endpoint information, use `provisioning.rsp` as a first resource.

Search `provisioning.rsp` for `#Load Balancer Configuration`, and if set to `TRUE`, use the subsequent entries to fill in the workbook.

If `#Load Balancer Configuration` is `FALSE`, then search for `#Virtual Host Configuration`. Three sections, `IP`, `Port`, and `Name` are listed. Search `#Webtier Configuration` and discover which `WEBTIER_MODE` is used in your installation (`IP`, `Port`, or `Name`). Then refer to the individual listings for the entries.

For example, if my `WEBTIER_MODE=IP`, then **FA Internal Name** would correspond to `VIRTUALHOST_IP_FIN_INTHOST` in the `#Virtual Host Configuration - IP` section. This method can be used to discover all the Source HTTP endpoint entries. Destination/Target entries are based on your plans for that environment, and may match the source.

■ FA Internal Name

■ FA Internal Port

- Source FA External Name

- Source FA External Port

- Source FA Ext. Protocol

- FA OHS Internal Port (w/LBR)

- FA OHS External Port (w/LBR)

- Target FA External Name

- Target FA External Port

- Target FA External Protocol

- *In Topology Mgr?*: Note this is for internal use and should not be changed.

- *Topology Mgr Domain Name*: Note this is for internal use and should not be changed.

2.2.8.2 IDM HTTP Endpoints

Search `provisioning.rsp` for `IDENTITY_OIM_INTERNAL_ENDPOINT_URL`. The resulting value (for example: `http://idminternal.mycompany.com:7777`) is used for the Internal Name, Internal Port, and Protocol. Search for `IDENTITY_OIM_EXTERNAL_ENDPOINT_URL` for the corresponding external entries. Note that the internal and external endpoint ports should be different.

For Identity Management, the default Topology Mgr Domain Name is IDMDomain, but it is not *required* to be named in this way. If your organization used a different name, edit this value in the Workbook as needed.

2.2.8.3 LDAP Abstract Names

Most of the LDAP Abstract Name and Port information can be found in `provisioning.rsp`. If OVD is used, then you must check the OVD Adapter in ODSM to find information about Oracle Internet Directory.

- **Policy Store (OID):** Search `provisioning.rsp` for `#Access and Policy Management` Within that section, locate `OAM_OPSS_HOST` and `OAM_OPSS_PORT` to identify the Abstract Host and Abstract Port.

You can search `provisioning.rsp` for `#Identity Management Configuration` to find the `IDENTITY_HOST` (Abstract name) and `IDENTITY_PORT` (Abstract Port) for one of the Identity Stores.

To determine which one you are dealing with, check `IDENTITY_SERVERTYPE`. The value will be `OID` or `OVD`.

- **OID Identity Store:** If `OID` was the `IDENTITY_SERVERTYPE`, enter its value here. In this case, `OVD` is not relevant and can be skipped.

- **OVD Identity Store (if used):** If `OVD` was used, enter the `OVD` information. If `OVD` is active, then it is used to reach `OID` and you check the OVD Adapter in Oracle Directory Services Manager (ODSM) to locate the information about the OID Identity Store.

2.2.8.4 Virtual IPs (VIPs)

The Enterprise Deployment Guides use virtual IPs in their sample environments. List any virtual IPs and virtual hostnames that may have been created in your source system for Identity Management and Fusion Applications. Add the source IP

addresses to the *Source Virtual IP* column. The target IP address should go in the *Target Virtual IP* column.

Note that the clone process requires that VIPs match the same IP address as the one specified in the **Topology** tab, **Topology** table, *Target IP Address* column. If that is not the case, you must change the IPs in the /etc/hosts file for the target system (temporarily) for the duration of the cloning process. You can revert to the desired VIPs once cloning is complete. See details in Section 2.4.3.3, "Set Up /etc/hosts File for Internal Endpoints".

2.2.8.5 Additional `/etc/hosts` File Entries

List any additional /etc/hosts file entries in the source system that should be added to the /etc/hosts in the target system. This includes any other hostnames used in the source environment that are not in the abstract hostnames Topology tab or in the FA/IDM HTTP Endpoints or Virtual IPs sections.

2.2.9 Databases

The database administrator can usually provide the information on this Workbook tab, or most information can be found in the provisioning.rsp file. Note that the system requirements for database duplication are described in Section 1.3.1.3, "Database Requirements". Also, in addition to the Discovery tables, there are additional tips for understanding and encompassing the source database structure, listed in Section 2.6.1, "Full Discovery of Source Database Environment".

2.2.9.1 FA DB Table

Search provisioning.rsp for #Database Configuration and use the subsequent entries to fill out the FA DB table entries in the Workbook. For example, see these sample .rsp entries:

```
DATABASE_ENTERPRISE_SELECTED=true
DATABASE_ENTERPRISE_SYSDBA_USERNAME=sys
DATABASE_ENTERPRISE_SYSDBA_PASSWORD=<encrypted key>
DATABASE_HOST=fusiondb.mycompany.com
DATABASE_PORT=1521
DATABASE_SERVICENAME=fusiondb
DATABASE_RAC_SELECTED=false
DATABASE_RAC_SYSDBA_USERNAME=
DATABASE_RAC_SYSDBA_PASSWORD=
DATABASE_RAC_SERVICENAME=fusion
DATABASE_RAC_INSTANCE1_HOST=
DATABASE_RAC_INSTANCE1_NAME=
DATABASE_RAC_INSTANCE1_PORT=
DATABASE_RAC_INSTANCE1_ROW=Instance1
```

The database administrator needs to provide the Directory entries that were used when installing and configuring the database.

When installing the Fusion Applications database on the source, it was required to create a file directory for the following database directories:

```
Directory APPLCP_FILE_DIR
Directory APPLLOG_DIR
Directory KEYFLEXCOMBFILTER
Directory FUSIONAPPS_PROV_RECOVERY_DIR
Directory OTBI_DBINSTALL_DUMP_DIR
```

To rediscover those file directory paths on the source environment, use sqlplus to run the following SQL command against the Fusion Applications database as sysdba:

```
select DIRECTORY_PATH from dba_directories where directory_name=<DIRECTORY_NAME>
```

where `<DIRECTORY_NAME>` is the database directory (such as `APPLCP_FILE_DIR`). Repeat for each directory. Enter the resulting paths in the source column. For the destination, you can choose to use matching paths or create a different path. You will manually create the directories on the destination in .

2.2.9.2 IDM and (Optional) OID DB Tables

Search `provisioning.rsp` for `#IDM Database Configuration` and use the subsequent entries to fill out the Identity Management DB entries in the Workbook. The Identity Management database administrator must provide the `IDM DB Prefix` used in the schemas.

Typically, Oracle Internet Directory (OID) DB information matches the Identity Management (IDM) DB entries. The Fusion Middleware Control (FMW) interface also includes information about the IDM Database(s). Otherwise, consult the Identity Management database administrator.

2.2.9.3 Whitelisted DB Links

If your destination databases have any database links, the cloning tool will report a validation error and will not continue unless the database link uses a connect string that is whitelisted: `\u0009\u0009`

Enter all whitelisted connect strings in `DB_WHITELISTED_DBLINK_CONNECTSTRING_LIST` separated by `','`.
For example: `DB_WHITELISTED_DBLINK_CONNECTSTRING_LIST=devdb,testdb`

2.2.10 Identity Management

Some of the Container information is located in `provisioning.rsp` and can be located there (as noted below). All the information can also be located in the Oracle Directory Services Manager (ODSM) and tips for locating the sections are listed as well.

2.2.10.1 LDAP Containers

- **IDM_IDSTORE_TYPE:** This specifies the identity store that is used in the source environment for IDM. If you provisioned your environment using the Enterprise Development Guide, this will be OVD.

- **FA_IDSTORE_TYPE:** This specifies the identity store that is used in the source environment for IDM. This can be either OID or OVD. This can be found in the `provisioning.rsp` in the `IDENTITY_SERVERTYPE` field.

- **Base DN:** This is the DN that appears after the `cn=Users`.

- **User Base DN:** Search `provisioning.rsp` for `IDENTITY_USERBASEDN=` (for example: `cn=Users,dc=mycompany,dc=com`)

- **Group Base DN:** Search `provisioning.rsp` for `IDENTITY_GROUPBASEDN=` (for example: `cn=Groups,dc=mycompany,dc=com`)

The remaining values must be located in the Oracle Directory Services Manager (ODSM).

- **APPID User Container DN:** Look under the User tree node in ODSM.

- **APPID Group Base DN:** Look under the Group tree node in ODSM

- **FusionUsers Container DN:** Look under the User tree node in ODSM. If your system has not had any users created, this entry will not be present in ODSM. In that case, leave it blank in the Discovery Worksheet.

- **FusionGroups Container DN:** Look under the User tree node in ODSM.

- **FA Domains JPSRoot:** The best place to find this is in the response file `OAM_OPSS_JPSROOTNODE`.

- **IDMDomain JPSRoot:** Best place to find this is in the IDMDomain Enterprise Manager. Go to **Farm_IDMDomain**, then **WebLogic Domains**, then **IDMDomain**. Right-click and go to **Security**, then **Security Provider Configuration**. The JSP Root is the `cn=XXXXXX` portion of the **Location URL** under **Security Stores.**

- **IDM Administrators Group DN:** Go to ODSM and get the entire DN for the IDM Administrators Group.

2.2.10.2 IDM Users

Go to the ODSM for all entries. Enter the user name, i.e. the value after **cn=**, as it appears in the ODSM.

- **OAAM admin user (LDAP):** Look for `cn=oaamadmin` in the Users tree node of ODSM.

- **OAM admin user (LDAP):** This is the user name used to log into the OAM Console, normally `oamadmin`. To check, look for `cn=oamadmin` in the Users tree node of ODSM. This user should be part of the **OAMAdministrators** group.

- **OblixAnonymous user (LDAP):** Look for `cn=OblixAnonymous` in ODSM in the Base DN tree.

- **OIM LDAP user:** Look for `cn=oimldap` in ODSM in the *systemids* tree. This user should be the user in OIMAdministrators Group.

- **OIM admin user (LDAP):** This is the user name used to log into the OIM Console, normally `xelsysadm`. To check, look for `cn=xelsysadmin` the Users tree node of ODSM.

- **IDM super user (LDAP):** Look for `cn=weblogic_idm` in the Users tree node of ODSM. This should be the administrator user ID for the Weblogic server and should be stored in OID.

- **OAM LDAP user:** Look for `cn=oamldap` in the Users tree node of ODSM.

- **IDM Policy Store ReadWrite (LDAP):** IDMPolicyRWUser if your environment was installed using the IDM Provisioning tool.

- **IDM Policy Store ReadOnly (LDAP):** IDMPolicyROUser if your environment was installed using the IDM Provisioning tool.

- **OVD ROOT user:** The OVD Root User will not appear in ODSM. It was defined during OVD installation and is used to log into OVD as administrator.

- **OVD EMADMIN user:** Check FMW Control under the **Agent-Monitored Targets** menu, and check the OVD Agent for the details.

- **IDMDomain WLS admin:** Should be the administrator of the Weblogic server, which is stored in weblogic-embeded LDAP. Typically is `weblogic`.

- **IDMDomain WLS nodemanager admin:** Check the `config.xml` file in the IDM domain and search for `<node-manager-username>`. You can also look in the WLS

Console under the **Security** tab. Select **Advanced** at the bottom of the page to see the listing.

- **OIM LDAP IT Resource user:** Usually `oimldap`. To check, got to the Oracle Identity Manager (OIM) Console, select **Advanced**, then **Manage IT Resource**. On the popup window click the **Search** button. Click on **Directory Server**. This is the OIM LDAP IT Resource User.

- **OAM LDAP Data Source user (LDAP):** Usually `oamldap`. To check, go to the Oracle Access Manager (OAM) interface, choose the **System Configuration** tab, then **Data Sources**, and select **OIMIDSTORE** and check the **Bind DN** value.

- **OVD--> OID Adapter User:** Go to ODSM, connect to OVD.

- **IDMDomain WLS Authenticator (LDAP):** Can be obtained by going to the WLS console in any of the domains, then going to **Security**, then **my realm**, then **Providers**, then **OVD Authenticator**, then **Provider Specific**. and obtaining user name portion of the `Principal` property.

- **IDMDomain EM Monitoring User:** Go to the Enterprise Manager for IDM Domain and click on **Farm** on the top-left corner. Select **Agent-Monitored Targets**. A list of agent-monitored targets (OID, OVD) will appear. Click the **Configure** button next to any of them. The property is the `WebLogic Monitoring User Name`.

- **IDMDomain Credential Bootstrap User:** Usually `orcladmin` OID user (if your environment was installed manually following the EDG) or the `IDMPolicyRWUser` if your environment was installed using the IDM provisioning tool.

2.2.10.3 FA Users

All the FA User information can be found in the `provisioning.rsp` file, or optionally in ODSM.

- **FA Super User (LDAP):** Search `provisioning.rsp` for `IDENTITY_SUPERUSER`.

- **FA IDStore Read-Only (LDAP) :** Search `provisioning.rsp` for `IDENTITY_RO_USER`.

- **FA IDStore ReadWrite (LDAP):** Search `provisioning.rsp` for `IDENTITY_USERDN`.

- **FA Policy Store ReadWrite (LDAP):** Search `provisioning.rsp` for `OAM_OPSS_USER`.

- **FA Policy Store ReadOnly (LDAP):** This is usually `PolicyROUser`, which can be verified by going to ODSM and looking in the Users tree node.

- **FA Domains nodemanger admin:** The nodemanager admin user.

- **FA Domains WLS Authenticator (LDAP):** The FA Domains WLS Authenticator user can be obtained by either 1) checking the `provisioning.rsp` file for `IDENTITY_RO_USERDN`, or 2) going to the WLS console in any of the domains, then going to **Security**, then **my realm**, then **Providers**, then **OVD Authenticator**, then **Provider Specific**. and obtaining user name portion of the `Principal` property.

- **FA IDStore Admin:** This is normally `orcladmin`.

- **FADomains Credential Bootstrap User:** This will usually be the `PolicyRWUser`.

2.2.10.4 OAM

Log in to the **Oracle Access Manager** and select **System Configuration** tab, then expand the **Access Manager Settings** tab, then choose **OAM Agents** to find the values for both web gates. Fusion Applications and Identity Management web gates use the same mode.

- **Webgate Mode:** Open either webgate entry and check whether the **Security** mode is set to `Open`, `Simple`, or `Cert`. Enter the value in the Workbook.

- **FA WebGate Name:** When you select **OAM Agents**, the table on the resulting page lists both the FA Webgate and IDM Webgate names in the Name column.

- **IDM WebGate Name:** When you select **OAM Agents**, the table on the resulting page lists both the FA Webgate and IDM Webgate names in the Name column.

- **Coherence Cluster Address Target:** This information can be found in the `OAMConfig.xml` file.

2.2.10.5 Credential Store: CSFName

The CSF (User) names for the IDM and FA domains are both found in the FMW Control.

For IDM Domain: Log in to FMW, and under **WebLogic Domain**, select **IDMDomain**. Right-click **Security>Credentials** to see the Credentials page. Open each listed folder, locate the key you are looking for, then click **Edit** to open the window that displays the User Name. If it's not there, then leave the row empty.

For FA Domain: Under **WebLogic Domain**, select, for example, **Common Domain**. Select **Security** and look for **maps** (such as Oracle Patching) and **key**. Click **Edit** to open the window that displays the User Name.

2.2.11 Email and Business Intelligence (BI)

In this version of the cloning tool, the Email entries are not used and can be ignored.

2.2.11.1 Email

This section is for SMTP information.

- **Email SMTP Host:** The SMTP host used for sending email notifications. This may be external to the Fusion Applications instance, such as a corporate address for order notifications, incident reporting, and so on. Enter the correct information for both the source and the destination environments.

  You will enter the value for the destination SMTP host in the `faclone.rsp` file, for the property `FA_SMTP_HOST_TARGET`.

- **Email SMTP Port:** The SMTP port associated with the SMTP host, for both source and destination environments.
  Enter the value for the destination SMTP port in the `faclone.rsp` file for the property `FA_SMTP_PORT_TARGET`.

- **Email SMTP User:** The SMTP user associated with the SMTP server, for both source and destination environments.
  Enter the value for the destination SMTP port in the `faclone.rsp` file, for the property `FA_SMTP_PORT_TARGET`

2.2.11.2 Business Intelligence

Use the Oracle BI Administration Tool to open the current RPD, find the following values, and enter them in the worksheet. For help with installing the tool, see *Appendix E, "Install the BI Administration Tool"*.

For **CRM BI Broker APPID User** and **URL**:

1. On the BI Admin Tool, open the latest RPD from your Fusion Applications environment and locate the **Physical** column on the screen.

2. Expand the item named
   `oracle.apps.crm.model.analytics.applicationModule.CrmAnalyticsAM_`
   `CrmAnalyticsAMLocal`

3. Double-click the "Connection Pool" subitem to display its properties.

4. On the **Connection Pool** popup window **General** tab, locate the field called *User name*. This normally contains the value `FUSION_APPS_BI_APPID`.

   This is the "CRM BI Broker APPID" value to be used in the Discovery Workbook

5. Click on the **Miscellaneous** tab and locate the field named **URL**. This is the value of the "CRM BI Broker URL" to be used in the Discovery Workbook

For **HCM BI Broker APPID User** and **URL**:

■ Repeat the steps used above for CRM, but use the item called
   `oracle.apps.hcm.model.analytics.applicationModule.HcmTopModelAnalyticsG`
   `lobalAM_HcmTopModelGlobalAMLocal` in the **Physical** column.

For **FSCM BI Broker APPID User** and **URL**:

■ Repeat the steps used above, but use the item called
   `oracle.apps.fscm.model.analytics.applicationModule.FscmTopModelAM_`
   `FscmTopModelAMLocal`in the **Physical** column.

> **Note:** The RPD password must be changed, which can be done in the BI Administration Tool.

2.2.12 Ports

The tables in the Ports tab of the Workbook are divided into Identity Management components and Fusion Applications components. Note that not every component may be relevant for your particular installation.

2.2.12.1 Identity Management Component Ports

Look for port information in various places, depending on the component type.

■ **IDM OHS:** Check the `httpd.conf` file in the IDM web server `config` directory.

■ **IDM OHS SSL:** Check the `ssl.conf` file in the IDM web server `config` directory.

■ **IDM OHS OPMN Local/Remote:** Check the `opmn.xml` file in the IDM web server `config/OPMN` directory.

■ **OID Components:** Check the OID `config` directory to find the following properties file: `/config/OPMN/opmn.ports.prop`.

■ **OVD Components:** Check the OVD `config` directory to find the following properties file: `/config/OPMN/opmn.ports.prop`.

■ **ODS Components:** Check the ODS `config` directory to find the following properties file: `/config/OPMN/opmn.ports.prop`.

■ **IDM Managed Servers and Admin Server:** Check the WebLogic Server Console (WLS) to find the port information in the user interface.

■ **Node Manager Properties:** If using Node Manager, check the `nodemanager.properties` file in the `wlserver_10.3/common/nodemanager` directory.

- **Proxy** and **Coherence**: Check the `oam-config.xml` file located in the `<IDMDomain>/config/fmwconfig` directory.

  Note that **SOA Coherence** is almost always set to the default value of `9778`. To double-check, go to `<IDMDomain>/bin`, find `set.DomainEnv.sh`, and search for the listing.

2.2.12.2 Fusion Applications Component Ports

Find the port information for almost all the Oracle Fusion Applications components by logging in to the WLS Console and checking the user interface. Two exceptions are listed below:

- **Node Manager Properties:** For Fusion Applications components, node manager properties can be found in the directory `wlserver_10.3/common/nodemanager/<abstracthost>` . For example: `wlserver_10.3/common/nodemanager/admin-apps.oracleoutsourcing.com`

  Every host has its own node manager process with its own directory, but typically all hosts use the same port number, so checking one node manager directory may be sufficient.

- **Informatica (IR or IIR):** This component is used only when Customer Relationship Management (CRM) products are installed. If applicable, look for this port number by searching `provisioning.rsp` for *Informatica* or *IR* or *IIR*, or search the Informatica host for `InformaticaIR/env/envc`.

2.2.13 Passwords

This tab is informational only; do not enter values in the fields! These are the passwords that will be required during cloning.

2.2.14 Generated RSP Entries

This tab organizes all your entries and presents them so they are easy to use. The Generated RSP Entries tab collates the data entered in all the other tabs and tables, and generates the entries and values that are needed by the final`faclone.rsp` file. When discovery is finished, transfer the generated RSP entries to the `.rsp` file as follows:

1. Locate the sample template that will become your `.rsp` file for cloning. Look for `$CLONE_HOME/bin/faclone.rsp.template.R7`. Save it as a new file, with the name `faclone.rsp`. This template, now renamed, includes both sample parameters and non-editable internal information used by the cloning tool.

2. Replace the sample parameters and values on the new `faclone.rsp` template with the Generated Cloning RSP Entries by copying each segment and pasting into your `faclone.rsp`.

3. Save the file; it will be used when the cloning scripts are run.

2.3 Create a Master Image of Source

For standard installations, master images from the source environment should be obtained using TAR and are done on a component-by-component basis. Be sure that all necessary patches have been applied to the source environment, as detailed in Section 1.3.1.3.

> **Note:** The databases (the Oracle Fusion Applications database and the Oracle Identity Management database(s)) are handled in a separate process (see Section 2.6). The master image does not include the source databases.

Before obtaining the images, it is recommended to validate the name resolution settings in the source environment. See Section 2.3.1. The image creation instructions are in Section 2.3.2.

> **Note:** If you have virtualized your provisioned Fusion Applications instance, then this section of the guide will not apply. You can request a copy of the *Oracle VM User's Guide* for information on importing existing VMs, instead.

2.3.1 Prepare Source Environment Name Resolution

The `faclone.sh` scripts rely on local name resolution (using the `/etc/hosts` file) for internal HTTP endpoints and abstract host names, to keep the internal wiring intact on the destination environment. If you are installing the source system following the Best Practices Appendix B, then this step will be streamlined and unnecessary. But if you already have a source installation that did not make use of abstract host names, then some preparation is required on both the source and the destination environments.

See Section 1.1.1.1, "Understanding Internal Wiring and Abstract Endpoints" to review the concepts, if necessary.

If, in your existing source environment, you have placed internal names in DNS, they need to be removed. You then must add them to `/etc/hosts` on each node of the source, where they take over the name resolution activity that was formally handled on DNS. You should also replace any abstract host name entries in the source DNS.

These steps are described in the following sections:

- Part 2.3.1.1, "Verify Whether DNS Entries for HTTP Internal Endpoints Exist on Source"

- Part 2.3.1.2, "Move Internal HTTP Endpoint Entries to `/etc/hosts` If Necessary"

- Part 2.3.1.3, "Remove Abstract Host Names from DNS if Necessary"

- Part 2.3.1.4, "Remove Virtual IPs from DNS if Necessary"

2.3.1.1 Verify Whether DNS Entries for HTTP Internal Endpoints Exist on Source

To validate:

- **Verify that no DNS entries exist for FA or IDM Internal Names**
  Check *Discovery Workbook: Virtual Hosts* tab, *FA Internal Name* column and *IDM Internal Name* column for the references.

- **Verify that these entries are all present in the /etc/hosts file** on each node at the source environment.

- **Ping these entries** to verify that they resolve to the correct nodes. They should all point either to the FA OHS node, the IDM OHS node or the Load Balancer/Reverse proxy.

If everything validates, then no further action is required on the source system.

If there are entries in the DNS for internal endpoints, then continue with Section 2.3.1.2.

2.3.1.2 Move Internal HTTP Endpoint Entries to `/etc/hosts` If Necessary

If your source environment has the internal HTTP endpoints set up in DNS, we recommend:

- Removing the DNS entries for the internal http endpoints.
  These corollate with the entries in the *Discovery Workbook: Virtual Hosts* tab, *FA Internal Name* column and *IDM Internal Name* column.

- Add the entries to the /etc/hosts file on each node at the source environment.

- Ping these entries to verify that they resolve to the correct nodes (they should all point either to the FA OHS node, the IDM OHS node or the Load Balancer/Reverse proxy.

2.3.1.3 Remove Abstract Host Names from DNS if Necessary

If your source environment was installed without using the Best Practices Appendix B, then the source host names will be used as the abstract host names and used as-is in the internal wiring of Fusion Applications. (In Section 2.4.3.3 you will add these now-abstract host name entries to the/etc/hosts files at the destination environment, so they will resolve correctly.)

It is desirable to further abstract these names by removing their entries from the source DNS. To do so:

- Add these host names to the /etc/hosts file for all servers at the source environment.
  Check *Discovery Workbook: Topology* tab, *Abstract Host Name* column for the references.

- Remove these CNAMES from DNS.

- Create new DNS CNAMES for the source servers.

2.3.1.4 Remove Virtual IPs from DNS if Necessary

If your Fusion Applications source environment was set up using Virtual IPs (VIPs), as per the instructions in the various *Oracle Fusion Applications Enterprise Deployment Guide*s, the corresponding host names must also be added to the /etc/hosts file for cloning. To do so:

- Ensure the host names for these virtual IPs are in the /etc/hosts file for all servers at the source environment.
  Check *Discovery Workbook, Virtual Hosts* tab, Virtual IPs (VIPs)

- Remove these CNAMES from DNS.

2.3.2 Create Master Image (.tar) Files

Table 2–1 describes the clone image files to be obtained from the source environment. We recommend using options `czvpf`, which compact the images and maintain permissions.

> **Note:** All mount point and directory structures for the image files must be the same between source and target.

The image creation instructions use the property names in the faclone.rsp file related to image files. Ensure that those properties are defined correctly in your .rsp file, as they will be used by the faclone.sh scripts to extract the files to the appropriate locations.

The Image files all require a checksum to be generated. Here are example commands for capturing the Identity Management (IDM) Shared directory:

```
cd $IMAGEFILE_POOL_DIR
cd ..
tar czvpf <image_file_name_with_extension> $STORAGE_SHARED_IDM
cksum <image_file_name_with_extension>
```

Table 2–1   Image File Details

Directory	Availability	Includes	May Also Include	TAR File Name	Root Directory
IDM Shared directory	Always	IDM Oracle Home IAM Oracle Home SOA Oracle Home IDMDomain AdminServer IDM oraInventory	IDMDomain managed servers AS instances for OID and OVD Webtier Oracle Home and instances for OHS	IMAGEFILE_SHARED_IDM	STORAGE_SHARED_IDM
FA Shared directory	Always	FA Base directory FA oraInventory	FA Webtier Oracle Home and instances under FA Scaled-out instances and domain managed servers	IMAGEFILE_SHARED_FA	STORAGE_SHARED_FA
IDM DMZ directory	Optional	IDM Webtier Oracle Home	Instances for IDM OHS	IMAGEFILE_DMZ_IDM	STORAGE_DMZ_IDM
FA DMZ directory	Optional	FA Webtier Oracle Home	Instances for FA OHS	IMAGEFILE_DMZ_FA	STORAGE_DMZ_FA
Local Storage	Optional	Instances for OVD, OID, OHS IDM domains, managed servers FA domains, managed servers BI Instance		IMAGEFILE_LOCAL_*	STORAGE_LOCAL_*

> **Note:**   In certain circumstances, you may want to copy the files manually, instead of creating TAR files and letting faclone.sh extract them. This can be accomplished by setting the faclone.rsp property EXTRACT_IMAGEFILE = false, which will skip the extraction of TAR files to the destination environment.

2.4 Prepare the Target Environment

After ensuring that your destination environment meets the core system requirements laid out in Section 1.3.2, "Destination Requirements", there are some additional steps to set up the destination network. They include:

- Section 2.4.1, "Configure DNS Entries for External Endpoints"
- Section 2.4.2, "Configure Load Balancer/Reverse Proxy Settings"
- Section 2.4.3, "Set Up Target Nodes and Storage"
- Section 2.4.4, "Test Connectivity"

2.4.1 Configure DNS Entries for External Endpoints

All external endpoints for the destination environment must be set up in DNS so they can be accessed by the end users.

Configure the following CNAMES on your DNS server.

DNS CNAME for	External HTTP Endpoint	Points At
Common		
BI	**Discovery Workbook: Virtual Hosts** tab	**Discovery Workbook: Topology** tab
Financial		
CRM	*Target FA External Name* column	External LBR/RP or nodes containing FA_WEBTIER
HCM		
Procurement		
Procurement (Supplier Portal)		
Projects		
SCM		
IDM	*Target IDM External Name* column	External LBR/RP or nodes containing IDMOHS_INST

> **Note:** Remember that *internal* HTTP endpoints should *not* be set up on DNS, but configured in /etc/hosts (see Section 2.4.3.3).

2.4.2 Configure Load Balancer/Reverse Proxy Settings

If using a load balancer (LBR) or reverse proxy (RP) as the front-end of Fusion Applications, then configure the following external and internal mappings.

2.4.2.1 External HTTP Endpoints

These should be configured at the LBR/RP that provides external access to Fusion Applications (for end-users and external integrations.

External LBR Mapping for	Port (LBR/RP)	Maps to (Node)	Maps to (Port)
Common			
BI	**Discovery Workbook: Virtual Hosts** tab	**Discovery Workbook: Topology** tab	**Discovery Workbook: Virtual Hosts** tab
Financial			
CRM	*Target FA External Port* column	Nodes containing `FA_WEBTIER`	*FA OHS External Port* column
HCM			
Procurement			
Procurement (Supplier Portal)			
Projects			
SCM			
IDM	*Target IDM External Port* column	Nodes containing `IDMOHS_INST`	*IDM OHS External Port* column

Note: Fusion Applications configures SSL to terminate at the LBR/RP, so you may also have to configure certificates as appropriate, on your LBR/RP.

2.4.2.2 Internal HTTP Endpoints

The default configuration when using the Load Balancer option during Fusion Applications provisioning is for the source environment to also have internal endpoints at the load balancer. In this case, you must also create appropriate mappings. Note that internal and external LBR/RPs may be different.

Internal LBR Mapping for	Port (LBR/RP)	Maps to (Node)	Maps to (Port)
Common			
BI	**Discovery Workbook: Virtual Hosts** tab	**Discovery Workbook: Topology** tab	**Discovery Workbook: Virtual Hosts** tab
Financial			
CRM	*Target FA Internal Port* column	Nodes containing `FA_WEBTIER`	*FA OHS Internal Port* column
HCM			
Procurement			
Procurement (Supplier Portal)			
Projects			
SCM			
IDM	*Target IDM Internal Port* column	Nodes containing `IDMOHS_INST`	*IDM OHS Internal Port* column

2.4.2.3 LDAP Endpoints

In a highly-available or scaled-out topology, the load balancer is used to route requests to the various instances of Oracle Internet Directory and Oracle Virtual Directory.

	LDAP LBR Mapping for	Port (LBR)	Maps to (Node)	Maps to (Port)
Main Discovery Workbook Location		**Identity Management** Tab, **LDAP Abstract Names** table	**Topology** Tab	**Ports** Tab
	OID	*Abstract Port* column	Nodes containing `OID_INST`	OID
	OVD	*Abstract Port* column	Nodes containing `OVD_INST`	OVD

2.4.3 Set Up Target Nodes and Storage

Complete the destination setup in the following sections.

2.4.3.1 Verify Target Server Setup

Ensure that:

- Target nodes have the same operating system and version number as the source nodes

- There is a 1:1 mapping between source and destination nodes; i.e the same number of nodes on each, and the same memory and disk space requirements will apply.

- An OS user (the "oracle" user) is configured on all nodes with the same ID and the same Groups. This user will own the destination Fusion Applications instance.

2.4.3.2 Verify System Requirements are Met for Installation

Ensure all the prerequisites for operating system, network, and storage, are satisfied for the destination environment. Check the following documentation for additional guidance:

- *Oracle Fusion Applications Installation Guide*

- *Oracle Fusion Middleware Enterprise Deployment Guide for Oracle Identity Management (Oracle Fusion Applications Edition)*

2.4.3.3 Set Up `/etc/hosts` File for Internal Endpoints

See Appendix C, "Abstract Hostnames in Detail" for more background information on this topic.

In the `/etc/hosts` file for each destination node, add the following entries:

Category	Hostname (Discovery Workbook)	Map to IP (Discovery Workbook)
Abstract host names	**Topology** tab, **Topology** table, *Abstract Host Name* column	**Topology** tab, **Topology** table, Target IP Address
Abstract LDAP names	**Virtual Hosts** tab, **LDAP Abstract Names** table, *Abstract Name* column	**Virtual Hosts** tab, **LDAP Abstract Names** table, *IP Endpoint - Target Environment* column
FA Internal HTTP Endpoints	**Virtual Hosts** tab, **FA HTTP Endpoints** table, FA Internal Name column	**Topology** tab, Node containing IDMOHS_INST or Internal LBR/RP
IDM Internal HTTP Endpoints	**Virtual Hosts** tab, **IDM HTTP Endpoints** table, IDM Internal Name column	**Topology** tab, Node containing FA_WEBTIER or Internal LBR/RP

Category	Hostname (Discovery Workbook)	Map to IP (Discovery Workbook)
Blacklisted host names	**Topology** tab, **Blacklisted hosts** table, Host name column	**Topology** tab, **Blacklisted hosts** table, Target IP Address (or 127.0.0.1)
Virtual IPs (VIPs)	**Virtual Hosts** tab, **Virtual IPs (VIPs)** table, Virtual hostname column	During Cloning: use the local IP of the host the VIP is pointing to, as in the **Topology** tab, **Topology** table, *Target IP Address* column. After Cloning: if a different IP address is to be used, use the VIP as in the **Virtual Hosts** tab, **Virtual IPs** (VIPs) table, *Target IP* column
Additional entries	**Virtual Hosts** tab, **Additional /etc/hosts Entries** table, Abstract Name column	**Virtual Hosts** tab, **Additional /etc/hosts Entries** table, Target IP column

2.4.3.4 Verify Target Storage Setup

The storage setup on the destination environment must mirror the setup at the source, i.e:

- Same size requirements
- Same mount point
- Shared storage must be mounted to equivalent nodes

Verify all storage is mounted according to the Storage tab in the Discovery Workbook, including:

- Identity Management Shared storage
- Identity Management DMZ storage (if applicable)
- Fusion Applications Shared Storage
- Fusion Applications DMZ storage (if applicable)
- Local Identity Management or Fusion Applications storage for each node in the topology (if applicable)

2.4.4 Test Connectivity

Ensure all DNS, LBR and `/etc/hosts` setup has been done correctly by performing the following tests from each one of the nodes in the destination environment:

2.4.4.1 Test Hostnames

Use the ping command to test if the endpoints are set correctly in your environment:

Ping from	Hostname to Ping	Should resolve to
All nodes	All host names from Section 2.4.3.3	As defined in that section

Ping from	Hostname to Ping	Should resolve to
FA DMZ node (if applicable)	Abstract host names for node containing all Fusion Applications WLS domains	As defined
IDM DMZ node (if applicable)	Abstract host names for node contain IDMDomain	As defined
All nodes	All endpoints defined in Section 2.4.2.1	As defined in that section (OHS or LBR, depending on topology)

2.4.4.2 Test Load Balancer/ Reverse Proxy Mapping

To test the load balancer/reverse proxy (LBR/RP) mappings, you must start a listener process on the node to which the load balancer/reverse proxy is forwarding requests. This can be done using the command `"nc -l <target port number>"` for each port you want to test.

LBR Mapping Type	Telnet from	Hostname and Port to Telnet to (from Discovery Workbook)	Should Resolve to (Discovery Workbook: Topology Tab)
Fusion Applications External LBR/RP	All nodes except DMZ	**Virtual Hosts** tab, **FA Virtual Hosts** table, *Target FA external Name/Port* columns	Nodes containing `FA_WEBTIER`
Identity Management External LBR/RP	All nodes except DMZ	**Virtual Hosts** tab, **IDM Virtual Hosts** table, *Target IDM external Name/Port* columns	Nodes containing `IDMOHS_INST`
Fusion Applications Internal LBR	All nodes except DMZ	**Virtual Hosts** tab, **FA Virtual Hosts** table, *Target FA internal Name/Port* columns	Nodes containing `FA_WEBTIER`
Identity Management Internal LBR	All nodes except DMZ	**Virtual Hosts** tab, **IDM Virtual Hosts** table, *Target IDM internal Name/Port* columns	Nodes containing `IDMOHS_INST`
LDAP LBR	All nodes except DMZ	**Identity Management** tab, **LDAP Abstract Names** table	Nodes containing `OID_INST`, `OVD_INST`

2.5 Make Master Images Accessible to Target

All the master files created in Section 2.3.2 should be copied to a shared location, visible to all nodes in the destination environment. In the `faclone.rsp` file, this location is defined as `IMAGEFILE_POOL_DIR`.

For DMZ nodes, where the shared location will not be accessible, copy the respective DMZ image file (`IMAGEFILE_SHARED_FA` and/or `IMAGEFILE_DMZ_IDM`) to the same location on the DMZ node. If the DMZ nodes also use local storage, copy the local storage image files (`IMAGEFILE_LOCAL_*`) as well.

We recommend using `checksum` to validate the file copy to the destination environment.

2.6 Duplicate the Databases from Source to Target

Ensure that database system requirements have been met (see Section 1.3.1.3). Note that you should have completed the Discovery phase for Database information as well (see Section 2.2.9).

Database duplication includes the following steps:

- Section 2.6.1, "Full Discovery of Source Database Environment"
- Section 2.6.2, "Duplicate the Source Database Using Your Preferred Tools"
- Section 2.6.3, "Post-Duplication Activities"
- Section 2.6.4, "Prepare Database Files that the Clone Tool Will Use for Validation"
- Section 2.6.5, "Database Validation and Final Steps"

2.6.1 Full Discovery of Source Database Environment

Section 2.2.9 gives tips for filling out the Discovery Workbook tables related to the Oracle Fusion Applications transaction database and the Oracle Identity Management database(s). In addition, collaborate with the database administrator to attain full familiarity with the following additional database details.

All this information will be used to prepare a destination environment that exactly matches the source, and then post-duplication, to verify that the two environments are identical in content and semantics.

In the Discovery Workbook, you will have identified:

- Whether the source is RAC, and what version, patch sets, and patches were applied
- The databases and instance names to be duplicated
  Use `v$database` and `gv$instance` to discover. Note that if RAC is used, you can search from any one of the hosts in the database cluster.
- Hosts on which the databases are installed
- Host operating system, version, and patch levels
- Oracle database Homes involved, their patch levels, and the space taken by each Home

Also identify the follow database-specific details:

- Identify whether ASM storage is used. If it is, then find the version, patch sets, and patches applied.
- Check if the SIDs are present in the `/etc/oratab` or `/var/opt/oracle/oratab`, as appropriate.
- Identify the Oracle Parameter files. Determine if initialization files (`init<SID>ora`) are in use, and if an `spfile` is in use.
- Identify the Database character set.
- Identify the Archive log mode.
- Identify the number of `redo` log groups and files.
- Identify the Flashback mode.
- Determine the number of data files and the number of temporary files. Not the space that they consume.

- Identify flash recovery area, and the storage allocated to it.

- Determine the time a last full (level 0) backup was taken.
  Use your preferred technology, to simplify the duplication process.

- Determine whether any tablespaces should be skipped.

- Determine the directory objects in the source database, and the location to which they point.

- List the database links in the source database, the corresponding TNS name from `tnsnames.ora`, if appropriate, and the hosts or IP address to which they point.

- Determine any local hosts entries (`/etc/hosts`, for example) that have been created.

- Determine the jobs and schedules that may have been created.

- Identify the number of job-queue processes in the `job_queue_processes` initialization parameter.

2.6.2 Duplicate the Source Database Using Your Preferred Tools

Database duplication is done in whatever method your enterprise uses: RMAN backup (along with installing the Oracle RDBMS Server binaries), file system copy, storage replication, VM snapshot, etc. Adhere to the requirements when duplicating the source database and mounting it to the destination environment:

- For RAC installations, ensure that the Grid Infrastructure is installed on the destination.

- Before duplicating, shut down the source Fusion Applications web tier and application tier, as well as the Identity Management web tier and application tier, and ensure that all in-flight transactions have been completed.

- Shut down the source database cleanly (no abort). The clone/copy must be taken cold.

- Remember that the topology and operating systems must be identical between source and destination.

2.6.3 Post-Duplication Activities

When the database has been duplicated onto the destination environment, bring it up and verify that the contents match the source.

Take care, before starting the destination database:

- Ensure the host is fenced off, so the database links do not point to the source system.

- Ensure that the initialization parameter (`job_queue_processes`) is set to `0`, so that the database scheduler and job queue system are disabled.

When you have verified that the destination database has been duplicated and mounted correctly, then configure it for full use. This means:

- Turn off archive log mode, if needed.

- Turn off flashback mode, if needed.

- Register the destination database with the local listener.

■ Manually create the file directory paths for the database directories, as defined in the Discovery Workbook (see Section 2.2.9, "Databases").
The relevant directories are:

```
Directory APPLCP_FILE_DIR
Directory APPLLOG_DIR
Directory KEYFLEXCOMBFILTER
Directory FUSIONAPPS_PROV_RECOVERY_DIR
Directory OTBI_DBINSTALL_DUMP_DIR
```

Then, as `sysdba`, update the dba directories on the database using sqlplus to run the following SQL command against the Fusion Applications destination database:

```
create or replace directory <DIRECTORY_NAME> as '<DIRECTORY_PATH>'
```

where `<DIRECTORY_NAME>` is the name of the directory as specified in the Setting column (such as `APPLOG_DIR`) and `<DIRECTORY_PATH>` is the value specified in the Target column of the Discovery Workbook.

■ Temporarily switch off any database password policies for all IDM and FA databases on the target environment. Those policies can be turned back on after cloning is complete.

■ Change the passwords for all schemas listed in Column A of the Discovery Workbook **C_Passwords** tab.

■ If it is a RAC database, then convert the single instance into a RAC multi-instance database.
For more information on duplicating databases in a RAC environment, and on converting a single-instance database to a RAC database, see *Oracle® Real Application Clusters Installation Guide for Linux and UNIX*,11*g* Release 2 (11.2).

2.6.4 Prepare Database Files that the Clone Tool Will Use for Validation

In addition to the administrator verifying and validating that the databases are duplicated properly, the cloning tool also performs validation steps to ensure that the source and destination database environments are equivalent. This automated validation requires files with information about the two environments, which the cloning tool will compare.

Obtaining the files and getting them into the correct location is a two-step process that the administrator performs on both the source and the target databases:

1. Run a discovery script on the each database (source and target). This generates two files per database. (Note that if you have a RAC database, the script is run on the first instance only.)

2. Place the generated files in the discover output directory (typically located in `${CLONE_HOME}/discover`.

If the database administrator is not responsible for the cloning process, then provide the script to the database administrator with the script instructions that follow.

The `discoverEnv.sh` script is located in the `${CLONE_HOME}/discover/scripts` directory.

2.6.4.1 Run the `discoverEnv.sh` Script on Each Database

The following activities are done by the database administrator for the Fusion Applications and Identity Management databases on the source and the destination environments:

1. Create temporary directories to hold the discover scripts and the output files. For example:

```
mkdir -p /some-path/clone/discover/scripts
mkdir -p /some-path/clone/discover/output
```

2. Place the discoverEnv.sh script in the scripts directory.

3. Ensure that the script has execute permissions turned o:

```
cd /some-path/clone/discover/scripts
chmod u+x ./discoverEnv.sh
```

> **Note:** Perform these steps on each of the servers where the fusion applications and Identity Management are available. For RAC this should be done only on the host where the first instance runs.

4. Run the discoveryEnv.sh script for each database.

   For help using the discoverEnv.sh command, the syntax is:
   `discoverEnv.sh {-h | --help}`

   The syntax to run the discoveryEnv.sh script is :

```
discoverEnv.sh [ {-t | --tmp} tmpdir ] [ {-d | --debugLevel} level ]
 {-e | --env | --environment}  {source|target} {-i | --sid} sid {-n |
--serviceName} service-name {-o | --oh | --oracleHome} oracle-home  {
-u |--userName} username  {-l | --directoryLocation} output-directory
```

   For example:

```
cd /some-path/clone/discover/scripts
./discoverEnv.sh --env target --sid fatarg1 --serviceName
fatarg1.mycompany.com --oh oracle-home-of-database -u system
--directoryLocation /some-path/clone/discover/output
```

   Where (conventions):

   - Command line options are of the form `--switch switch-value`

   - Optional switches are surrounded by brackets [ ]

   - Mandatory switches are surrounded by curly braces { } if there are multiple ways of specifying that switch. If there are is only one way of specifying it then it is not surrounded by {}. For example : { --env | --environment } {source|target} and --userName sysdba-username

   - Where there are multiple choices, they are separated by the pipe symbol "|"

   Parameters:

Valid Switches	Value
`--tmp tmpdir`	Directory where script writes temporary files before writing the final output files

Valid Switches	Value
`--env` or `--environment`	Specify `source` or `target`, depending on which environment the script is run.
`--sid  sid`	The SID of the database, as would be identified in the clone scripts. Note: If the instance name (noted in the Discovery Workbook) differs from the SID, then use the instance name.
`--serviceName service-name`	The service name of the database instance.
`--oracleHome $ORACLE_HOME`	Specify the `ORACLE_HOME` path for the database instance for which the `discoverEnv.sh` script is being run.
`--userName sysdba-username`	Specify the database account with `sysdba` privileges. Valid values include "/", sys, and so on.
`--directoryLocatio n ${CLONE_ HOME}/discover/out put`	The output file directory created for this instance.

2.6.4.2 Place Generated Files In Correct Directory

The script generates two files for each run (for each database instance) and places them in the directory defined by the `--directoryLocation` command line switch. Provide these files to the application administrator who performs the cloning.

> **Note to Application Administrator:** Before running the cloning scripts, ensure that the files are placed in the `$CLONE_HOME/discover/output` directory.
>
> There should be two files generated for each database involved. Thus there should be at least 8 files total: two for Identity Management, two for Fusion Applications, for both source and target. When there are separate databases for Identity Management and OID, there would be a total of 12 files (6 from source and 6 from target).
>
> RAC instances do not affect the file number; the script is run only on the first instance.

2.6.5 Database Validation and Final Steps

As an administrator, to finish verifying that the destination database is up, running, and identical to the source:

- Check that all instances (single or RAC) have been registered with the listeners and are accessible remotely.

- Ensure all database instances are running.

- Ensure that jobs relevant only to the source environment have been dropped (preferable) or disabled.

- Ensure locations defined in the directory object are present, and have the appropriate permissions. In the case of RAC, ensure that it is in a shared file system or location, as matches the source.

- Ensure the number of tablespaces and data files matches the source.
  (Note: if you skipped source-specific tablespaces, don't forget to adjust your count to accommodate.

- Finally, shut the database(s) down, ensuring all instances are down, and take a complete cold backup. This represents your master copy of the destination environment, and can be used for recovery, in case of need.

If the preceding sections in this Guide are complete, you can continue with running the cloning script (Section 2.7).

2.7 Run `faclone.sh` to Extract and Rewire Components on Target

Running the cloning scripts on the destination environment for each Fusion Applications component is known as the "clone phase" of the process. Before the clone phase, ensure that:

- You have completed the Discovery Workbook

- You have completed the `faclone.rsp` file

- You have completed all preceding steps in this guide (met all system requirements, created master image files, prepared the destination environment, made master images available, and duplicated the databases.

- You have extracted the FAClone Kit to a location shared among all nodes in the topology with read-write permission (except the database and DMZ nodes). We will refer to this location as `$CLONE_HOME`

The clone phase is fully automated and accomplishes the following:

- Extracting master TAR files

- Re-wiring the mid tier to the new destination databases

- Changing the passwords for all components (see details below)

- Performing cleanup tasks, such as removing running job instances, clearing log files, etc.

- Re-wiring the midtier and web tier for the new external endpoints

- Starting up components

- Performing initial validation of individual components

During the clone phase, no access to the source environment is needed. This phase implements fail-safe measures to ensure that no access from the clone to the source environment takes place during or after the clone process.

Understand Password Handling for Cloning

Part of cloning is entering passwords for the target environment, which cannot be the same as the source environment. You can define these passwords in one of two ways:

- As an independent activity before running the *faclone.sh* script, or

- As the first prompted action when running the script.

To enter the passwords independently, use the following syntax:

```
/path-to-bin/faclone.sh -p /path-to-rsp-file/response_file.rsp generate
passwordfile
```

You will be prompted to enter passwords which are labeled according to Column D of the **C_Passwords** tab in the workbook.

Otherwise, when you start `faclone.sh` the first time, the script will prompt you for passwords then. Check Column D of the **C_Passwords** tab in the workbook for any needed clarification.

When the passwords have been entered, an encrypted password file is generated and stored on the same node as the response file.

2.7.1 Recovering from Issues While Running `FAClone.sh`

If issues occur while running the script, you can identify the problem, resolve it, and re-run the script.

In the event of a failure of the cloning process, you will be prompted to run what is known as "fail clean up". Once prompted, you will have three choices:

1. Leave the clone in this state for a period of time while you gather the information required to debug the issue (i.e., do not press any key). Once you proceed, debug information will be lost. After obtaining the debug information, you can proceed with fail cleanup and prepare to restart the cloning process

2. Proceed by hitting any key and the fail cleanup process will be run.

3. Hit **Ctrl-C** to exit the clone. This leaves the system in a state where the fail cleanup will be detected when the clone is restarted and it will automatically be run with no prompt to the user. This is an alternate way to take the required time to gather the debug information prior to restarting the clone.

To identify the cause of the error, find the latest log available in the `CLONE_LOG` directory and look for error messages. Additional logs to check include WLS server logs, OPMN logs, node manager logs, and so on.

Once found, an error can be addressed as follows:

- If the fix requires a response file or password file change, you may have to clean up the target environment and the CLONE_HOME before running `faclone.sh` again. See *Section 2.7.1.1, "Making Changes to Response or Password Files After Having Run the Script"*.

- If the fix requires a manual change to a configuration on the target environment (mid-tier or web tier), we recommend cleaning up the target environment and restarting the process. See *Section 2.7.1.2, "Cleaning Up the Target Environment and CLONE_HOME to Restart the Script"*

- If the fix requires an OS-level parameter change only or a simple database parameter change (with NO changes to the response file, password file or configuration changes in the mid-tier or web-tier), you can simply re-run `faclone.sh` with no additional changes.

2.7.1.1 Making Changes to Response or Password Files After Having Run the Script

If the issue requires a change to the response file or password files after you have already run `faclone.sh`, you must restart the clone process from the beginning, including cleaning up the target environment and the `CLONE_HOME`. This is necessary because some of the new values in the response/password file may have already been written to the clone image, so simply changing them in the response file is not enough to revert those changes already made.

After making changes to the response file or password file, follow the procedure outlined in Section 2.7.1.2.

To make file changes:

Response File:

Simply change the value in the file directly. For consistency, don't forget to update the Discovery Workbook accordingly.

Password File:

This file is encrypted, so changing the value manually is not an option. To regenerate the file:

1. Find the password file in the same directory you previously specified for the response file when you ran `faclone.sh` (`-p` or `--responsefile` option). The password file name is defined by the response file property `CLONE_PASSWORDFILE_NAME`.

2. Rename or delete the file.

3. Next time `faclone.sh` is run, the tool will prompt again for the password values.

2.7.1.2 Cleaning Up the Target Environment and CLONE_HOME to Restart the Script

In some situations, including changes to the response or password files, you must restart the clone process from the beginning, and need to make manual changes to the target environment. To do so:

1. Ensure that `faclone.sh` is not running. If it is, wait for it to finish or kill the process manually.

2. Ensure there are no running processes on the target environment, including any mid-tier process (OPMN, OID, OVD, BI, OHS, WLS, node manager) as well as databases.

3. Clean up the target environment by deleting all Fusion Applications (FA), Identity Management (IDM) and database (DB) directories.

4. Restore the database clones to the target environment, leaving the target environment in the same state as the end of Section 2.6.

5. Clean up the `CLONE_HOME` by deleting the contents of the following directories (don't delete the directories themselves; just leave them empty):

   - backup (`CLONE_BACKUP`)

   - log (`CLONE_LOG`)

   - phaseguards (`CLONE_PHASEGUARDS`)

   - temp (`CLONE_TMP`)

6. Re-run `faclone.sh`.

2.7.2 How the Clone Phase Works

The automated clone scripts are run on the following nodes in the destination environment:

- Oracle Internet Directory (OID)

- Oracle Virtual Directory (OVD)

- Identity Management Domain Admin Server

- Fusion Applications Common Domain AdminServer

- Fusion Applications Business Intelligence (BI)

- Fusion Applications Web Tier

If one or more of these components run on the same node, `faclone.sh` will only run once on each mode. For example:

- If OID and OVD are located on the same node, since both are processed in the `IDMMT` mode, `faclone.sh` will detect this automatically and process both without the need to run it twice.

- If the FA CommonDomain AdminServer and BI are on the same node, since both are processed in the `FAMT` mode, `faclone.sh` will detect this automatically and process both without the need to run it twice.

- If the FA CommonDomain AdminServer and the FA Web Tier are located on the same node, they will still require two separate runs of `faclone.sh` (one in `FAMT` mode and one in `FAWT` mode) since they are processed in different modes of the clone phase.

Note that in *version 1*, `faclone.sh` only supports running on the primary nodes of scaled-out environments. Thus, while it can still run against scaled-out environments, only the primary nodes will be cloned completely.

2.7.3 Executing the Clone Phase Scripts

To execute the clone phase, run the `faclone.sh` tool located in `$CLONE_HOME/bin`, as follows

```
faclone.sh clone <mode> --responseFile <response file full path>
```

Table 2–2 describes where, how, and in what order to run the `faclone.sh` scripts.

Running `faclone.sh` may take several minutes, depending on the mode chosen. While it is running, you can follow its progress on the screen. You can also view the log files generated in the `$CLONE_HOME/log` directory. Upon completion of each mode on a node, you will see a message indicating that mode has completed successfully.

Table 2–2   Clone Phase

Order	Mode	Runs On	Details
1	idmmt	All Primary IDM midtier nodes in the following order: 1. OID (primary node) 2. OVD (primary node) 3. IDMDomain AdminServer (primary node) If one or more of the above are located on the same node, you should still follow the same order e.g. if OID and OVD are on the same node and the IDMDomain AdminServer is on a separate one, run `faclone.sh` first on the OID/OVD node (both will be processed automatically), then run it on the IDMDomain AdminServer node. If you attempt to run the `IDMMT` mode in a different order, `faclone.sh` will display an error message. This mode requires node manager to be running on all WLS nodes before starting it.	Extracts the IDM MidTier clone images, and performs re-wiring, password change, cleanup, server startup, and validation

Table 2–2   (Cont.) Clone Phase

Order	Mode	Runs On	Details
2	idmwt	The primary IDM Web tier node.	Extracts the IDM Web Tier clone images, and performs re-wiring, password change, cleanup, server startup and validation
		If this node is located on the DMZ and the node is not able to access the shared storage where `$CLONE_HOME` is, copy the FAClone Kit to the DMZ node and run it from there.	
		The clone phase for the Web Tier checks whether the IDMDomain can be accessed via HTTP from the Web Tier node, so be sure you have completed the clone phase in `IDMMT` mode before starting it in `IDMWT` mode.	
3	famt	**NOTE: This mode requires node manager to be running on all WLS nodes before starting it.**	Extracts the Fusion Applications Mid Tier clone images and performs re-wiring, password change, cleanup, server startup and validation
		All primary Fusion Applications midtier nodes in the following order:	
		1. CommonDomain AdminServer (aka Primordial node)	
		2. BI (primary node)	
		If both are located on the same node, you only have to run the `FAMT` mode of the clone phase once.	
		If you attempt to run the `FAMT` mode in a different order, `faclone.sh` will display an error message.	
		The `FAMT` mode validates the successful completion of the `IDMMT` and `IDMWT` modes, when it's started.	
4	fawt	The primary FA Web tier node.	extracts the FA Web Tier clone images and performs rewiring, password change, cleanup, server startup and validation
		If this node is located on the DMZ and the node is not able to access the shared storage where $CLONE_HOME is, copy the FAClone Kit to the DMZ node and run it from there.	
		The Clone phase for the Web Tier checks whether the FA Domains can be accessed via HTTP from the Web Tier node, so make sure you have completed the Clone phase in FAMT mode before starting it in IDMWT mode.	

2.8 Perform Validation Steps

The clone tool performs automatic validation on each component during the clone phase. Those automated processes are detailed in Table 2–3, and can also be triggered manually with corresponding validation modes, as desired.

The Fusion Applications Administrator should also perform a final environment validation of the entire stack. Table 2–4 gives a high-level list of such tasks.

Table 2–3   Validation Steps

Automated Task	Detail	Manual Validation Mode
Validate IDM DB	- Test connection to IDM DB with all schema users -Basic data validation	`idmdb` May be 1 or 2 databases
Validate IDM Mid Tier	- Test direct connectivity to OAM Console, OIM console, EM, WLS console - Test all servers are up	`idmmt` validates the IDM Mid Tier environment using the destination configuration from the response file
Validate IDM Web Tier	- Test connectivity to IDM consoles via LBR / OHS - Validate SSO for IDM - Leverage existing IDM validation code	`idmwt` Validates the IDM Web Tier environment using the destination configuration from the response file
Validate FA DB	- Test connection to FA DB with all schema users - Basic data validation	`fadb`
Validate FA Mid Tier	- Test direct connectivity to EM, WLS consoles - Test all servers are up	`famt` Validates the FA Mid Tier environment using the destination configuration from the response file
Validate FA Web Tier	- Test connectivity to FA via LBR / OHS - Validate SSO for FA	`fawt` Validates the FA Web Tier environment using the destination configuration from the response file

Table 2–4 provides a high-level list of functional validation steps that should be performed by an experienced Oracle Fusion Applications administrator.

Table 2–4   Technical Stack Validation

Component	Task	Expected Outcome
OID	Use an LDAP client (such as JXplorer, Apache Directory Studio) to connect to OID at the destination host and port specified in the Discovery Workbook. If using a Load Balancer for LDAP, use the Load Balancer host and port to connect.	You should be able to connect successfully and see the objects in the LDAP tree, including users and groups.
OVD	Use an LDAP client (such as JXplorer, Apache Directory Studio) to connect to OVD at the destination host and port specified in the Discovery Workbook. If using a Load Balancer for LDAP, use the Load Balancer host and port to connect.	You should be able to connect successfully and see the objects in the LDAP tree, including users and groups.
WLS Console (Identity Management	Connect to the WLS Console of the IDMDomain using a browser (point directly at the WLS AdminServer port).	You should be able to connect successfully and check the status of the servers in the domain.

Table 2–4   (Cont.) Technical Stack Validation

Component	Task	Expected Outcome
EM Console (Identity Management)	Connect to the EM Console of the IDMDomain using a browser (point directly at the WLS AdminServer port).	You should be able to connect successfully and check the status of the Fusion Middleware components.
OIM Console	Connect to the OIM Console using a browser (point directly at the OIM server port).	You should be able to connect successfully and see users and groups.
OAM Console	Connect to the OAM Console using a browser (point directly at the AdminServer server port).	You should be able to connect successfully.
Identity Management OHS	Connect to the external endpoint host and port for the IDM OHS (point at the OHS host/port)	You should be able to connect successfully and see the default HTTP Server page.
Identity Management OHS (via LBR)	Connect to the external endpoint host and port for the load balancer for IDM (point at the LBR if your environment has one).	You should be able to connect successfully and see the default HTTP Server page.
SSO	Connect to any/all of the IDM Consoles (WLS, EM, OIM, OAM) using a browser (point at the OHS host/port or LBR if your environment has one).	You should be redirected to the OAM SSO login screen. After logging in you should see the console main page.
WLS Console (Fusion Applications)	Connect to the WLS Console of each Fusion Applications Domain using a browser (point directly at the WLS AdminServer port).	You should be able to connect successfully and check the status of the servers in the domain.
EM Console (Fusion Applications)	Connect to the EM Console of each Fusion Applications Domain using a browser (point directly at the WLS AdminServer port).	You should be able to connect successfully and check the status of the Fusion Middleware components (including BI, WebCenter, ESS, SOA) as well all the Fusion Applications managed servers for each domain.
SSO/Home Page (Fusion Applications)	Open the Fusion Applications Homepage using a browser (point at the CommonDomain OHS host/port or LBR if your environment has one);	You should be redirected to the OAM SSO login screen. After logging in you should see the Fusion Applications Homepage and no error messages. If you see error message, attempt to refresh the page as they may simply be timeouts since this is the first time the page is being accessed.

Table 2–4   (Cont.) Technical Stack Validation

Component	Task	Expected Outcome
Functional Setup	Navigate to Setup and Maintenance page using the Navigator Menu.	You should see the Setup and Maintenance page with no error messages.
Scheduled Jobs	Navigate to the Scheduled Jobs page using the Navigator Menu. *	You should see the Scheduled Jobs page with no error messages. There should be no running or pending jobs as all jobs should have been cleaned up during the clone phase.
Reports and Analytics	Navigate to the Reports and Analytics page using the Navigator Menu. *	You should see the Reports and Analytics page with no error messages. Click on the folders to display the available reports.

* If the Scheduled Jobs or Reports and Analytics links do not appear in the Navigator Menu, you will have to enable them in Functional Setup:

■ Navigate to **Setup and Maintenance** using the Navigator Menu.

■ Use the Search box to search for the **Work Menu**. The results will display on the right and should include **Manage Menu Customizations**.

■ Click on **Go to Task** (next to Manage Menu Customizations).

■ Find the **Scheduled Jobs** or **Reports and Analytics** menu items on the tree and make sure they are configured as visible/rendered.

■ For more information see the*Fusion Applications Extensibility Guide*.

2.9 Post-Clone Cleanup

When cloning and validation is complete, you may need to do the following:

■ **Change the RPD password:** See Section 2.9.1.

■ **Change OID and OVD passwords for Oracle Identity Management:** Because the passwords for these components are stored in the cwallet.sso files, they are not updated by the faclone scripts and must be handled manually. See Section 2.9.2.

■ **Rewire IIR using the Functional Setup Manager:** See Section 2.9.3.

■ **Re-spin essbase cubes** for Financial and Customer Relations Management (CRM) products, if applicable (see Section 2.9.4). If these products are used, this step is mandatory.

■ **Enable the OIM Job Scheduler:** See Section 2.9.5.

■ **Perform data masking.** See Appendix A. This step is optional.

■ **Rewire or reintegrate any extensions or customizations** that might have been done on your particular source environment, which are not replicated by the clone tool. (As needed.)

If your environment contains components in addition to those described in Section 1.1.1.2.1, "Component Review", then you must manually re-wire those components in the post-clone phase.

For example, if your environment uses federation (OIF) for Single-sign-on between Fusion Applications and other applications protected by Oracle Identity Management software, then you must manually re-wire the OIF components, based on what portion(s) of the applications outside Fusion Applications will be cloned. Likewise, if you have enabled HR2HR integration, then those integration components must be rewired as well. Your environment may contain additional examples.

- **Revert OIF Endpoints (if applicable):** If you use OIF to control front-end access for Fusion Applications, then you should have added the OIF endpoints to the blacklisted hosts. You will now revert the OIF configuration using it as the Fusion Applications front end to be able to log into Fusion Applications. See Section 2.9.6.

2.9.1 Change the RPD Password

Follow these steps to change the repository password using the BI Administration Tool, then publish the modified repository in Fusion Middleware Control. (To install the BI Administration tool, see Appendix E.)

Change the password:

1. Open the repository in the Administration Tool in offline mode.
2. Select **File**, then select **Change Password**.
3. Enter the current (old) password.
4. Enter the new password and confirm it. The repository password must be longer than five characters and cannot be empty.
5. Click OK, then save and close the repository.

Publish the repository:

1. Open a Web browser and log in to Fusion Middleware Control from the computer where the updated repository is located.
2. In the navigation tree, expand **Business Intelligence** and then click **coreapplication** to display the Business Intelligence Overview page.
3. Display the Repository tab of the Deployment page.
4. Click **Lock and Edit Configuration**.
5. Click **Browse** next to Repository File. Then, select the updated repository file and click **Open**.
6. Enter the new (updated) repository password in the **Repository Password** and the **Confirm Password** fields.
7. Click **Apply**, then click **Activate Changes.**
8. Return to the Business Intelligence Overview page and click **Restart**.

2.9.2 Change OID and OVD Passwords

To update the OID and OVD passwords in the Fusion Middleware Control (or Enterprise Manager) these steps must be performed manually:

1. The password may be present in the `IDMLCM` provisioning subtree: (`IDMLCM_HOME/provisioning/idm-provisioning-build/config/adminpassword`). This file should be deleted.

2. Go to the primary OID node and check for the presence of the file

   `<OID_INST_DIR>/config/OPMN/opmn/admin.password.txt`
   If this file is present, you should delete it completely.

3. Run the following command:

   `<OID_INST_DIR>/bin/opmnctl updatecomponentregistration -componentType OID`
   `-componentName <OID_COMPONENT_NAME> -Sport <OID_SSL_PORT> -Port <OID_PORT>`

   `<OID_INST_DIR>` is the directory of the OID instance.
   `<OID_COMPONENT_NAME>` is the component name for OID.

   This will display the abstract hostname of the IDM Domain host and the username of the domain administrator. You are prompted for the domain admin password. Upon success, the message `Command succeeded` is displayed.

4. Go to the primary OVD node and check for the presence of the file:

   `<OVD_INST_DIR>/config/OPMN/opmn/admin.password.txt`
   If this file is present, you should delete it completely.

5. Run the following command:

   `< OVD_INST_DIR>/bin/opmnctl updatecomponentregistration -componentType OVD`
   `-componentName < OVD_COMPONENT_NAME> -Sport <OID_SSL_PORT> -Port <OID_PORT>`

   As above, you are prompted to enter the domain admin password for OVD, and should receive the `Command Succeeded` message.

2.9.3 Rewire Informatica Identity Resolution (IIR) in the Functional Setup Manager

If you have installed the Customer Relations Management (CRM) product in your environment, then IIR must be manually rewired using the Functional Setup Manager. To do so:

1. From the Navigator menu, select Setup and Maintenance. Select the All Tasks tab.

2. In the resulting Search panel, enter "Manage Server Configuration" in the Name field. When Manage Server Configuration is displayed at the bottom of the page, click its "Go to task" icon in the fourth column.

   The Manage Data Quality Server Configurations page is displayed, with another Search panel at the top.

3. Leave all search fields blank and click "Search."

   Four results are displayed at the bottom of the page: "Batch Cleanse Server," "Realtime Cleanse Server", "Advanced Batch Match Server", and "Realtime and Batch Server."
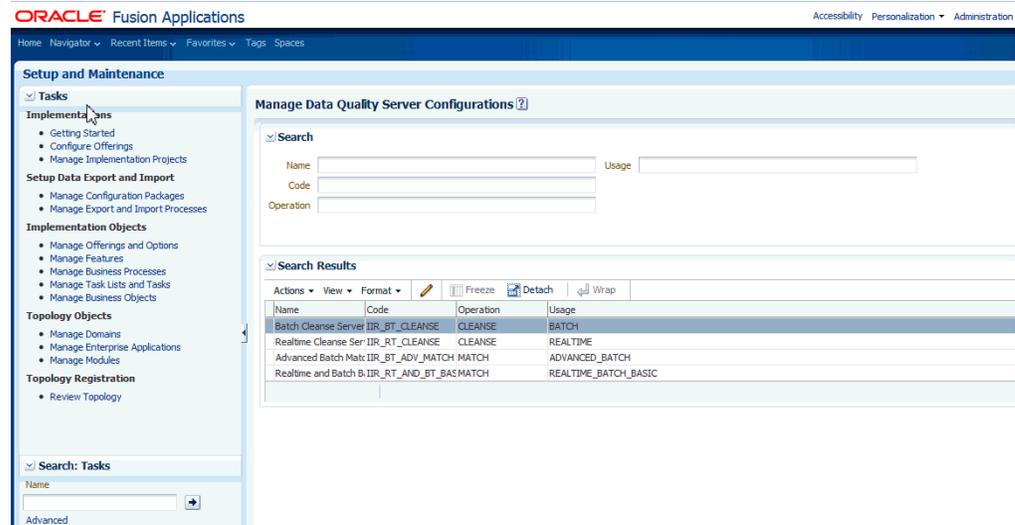
Functional Setup Manager page, as described in surrounding text.

****

For each of these results, you will need to edit the server address to match your target environment. If the port is different the default, or if no port number is displayed, then port information will be needed as well. To do this:

4. Select one of the results (i.e. Batch Cleanse Server) and click "Edit."

5. On the resulting page, enter the following values:

- **Server Address:** IP of destination Admin host where IIR will be running

- **Server Port:** enter the appropriate value if needed, or accept the default.

6. Save edits and repeat the process for the other three pages.

2.9.4 Re-Spinning Essbase Cubes for CRM and FSCM

The Customer Relations Management and Financial product families of Fusion Applications use multi-dimensional Essbase cubes as part of the reporting and analytics functionality.

After completion of cloning, these Essbase cubes must be recreated against the cloned Fusion Applications database, to confirm that the re-wiring was successfully completed, and to ensure that the cubes exactly correlate to the transactional data in the cloned database. This process is called 're-spinning'.

2.9.4.1 Re-spinning for Customer Relations Management (CRM)

If your cloned environment uses the CRM product family, follow these steps.

1. Log into the Fusion Applications Home Page.

2. From the **Navigator** drop-menu, select **Territories and Quotas**.

3. From the **Task** list on the left, select **Enable Dimensions and Metrics**.

4. Create a New Stage environment, or Edit an existing Stage environment (depending on whether a Stage environment has already been previously created).

   ■ To create new: Click on the **Create Stage Environment** button then move to step 5.
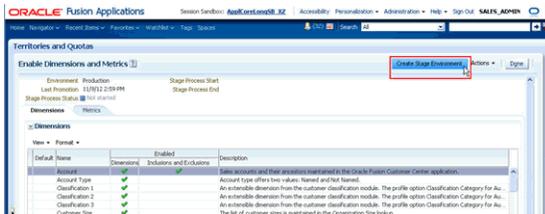
   ■ To edit existing: Click on the **View Stage Environment** button.

   ■ Click the **Edit** button.
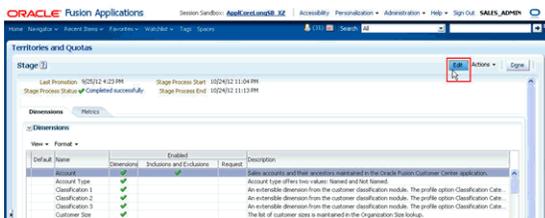
5. Start the **Stage synchronization process** (background ESS Job) which populates the dimensional data used to define territory coverages, as well as the Essbase cube used to support Territory metrics and analytics.
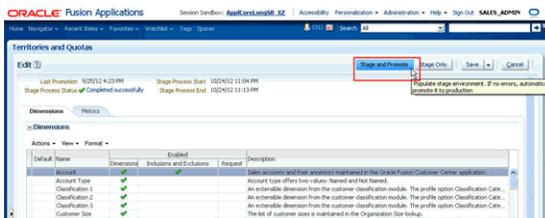
6. Click OK at the informational pop-up. Note: The UI will initially show the status for the Staging process as **Scheduled for later,** but will change to **Processing** when

the job is actually running. You can click on the **Stage process** status link to view the ESS Job details as well as download the log.

7. When the job completes, the dimension data and Essbase cube should be loaded and available for use in Territory alignment.
The **Last Promotion** value will be updated to the latest date-time when the process completed the promotion. The S**tage Process Status** should return back to *Not started* in preparation from the next cycle. The button at the top should also change back from **View Stage Environment** to **Create Stage Environment** to prepare for the next cycle.

2.9.4.2 Re-spinning for Financials (FSCM)

If your cloned environment uses the FSCM product family, follow these steps.

(Note that to re-spin the FSCM cubes, the user performing the process must have the "General Accountant" role assigned, and the menu option for **Tools / Scheduled Processes** (within the main **Navigator** section) must be un-hidden.)

1. Generate a list of the cubes that need to be rebuilt by logging into the Fusion Applications database through a SQL+ session and running the following command:

```
SELECT fksiv.name chart_of_accounts_name, glc.user_period_set_name calendar_
name
FROM gl_balances_Cubes cubes,
gl_calendars glc,
fnd_kf_str_instances_vl fksiv
WHERE cubes.calendar_id = glc.calendar_id
AND cubes.chart_of_accounts_id = fksiv.structure_instance_number AND
fksiv.application_id = 101;
```

2. Run the following SQL code to delete the existing cube details:

```
delete from gl_balances_cubes;
delete from gl_essbase_duplicate_parents;
COMMIT;
```

3. Log in to the Fusion Applications Home Page.

4. From the **Navigator** drop-menu, select **Scheduled Processes**' from the **Tools** section.

5. Select **Schedule New Process**.

6. In the search window, click the drop-down arrow and scroll to the bottom of the list. Select **Search**.

7. Enter **Create Cube**, and then **Search**. Select the **Create Cube** action in the results set, then click OK, and then OK again

8. Using the output from Step 1, enter the details of the **Chart of Accounts**: *Name*, *Calendar* and *Starting Period*, and then click **Close**.
You can select the option to be notified of the process completion if you prefer.
Repeat this process for each cube that was listed in the output from Step 1.

2.9.5 Enable the OIM Job Scheduler

The clone tool disables the OIM job scheduler to ensure that no jobs run during the cloning process. When cloning is complete, you must manually enable the job scheduler.

1. Disable jobs that send email notifications, as follows:

- Log in to the OIM console (`http://<host>:</port>/oim`).

- Go to Advanced and click on "Search Scheduled Jobs."

  Look in the right pane for the following jobs:

  ```
  Disable/Delete User After End Date
  Password Expiration Task
  Password Warning Task
  ```

- Click each job, and its resulting page, select the "Disable" button.

2. Enable the OIM Job Scheduler, as follows:

- Log in to the WebLogic Console of the IDM Domain.
  (`http://<host>:</port>/console`.

- On the left pane, select the **Environment | Servers** link and click the name of the OIM managed server (usually `oim_server1` or `wls_oim1`).

- Select the **Configuration | Server Start** tab in the right pane. (Page will be read-only). The property called 'Arguments' includes the following:
  `-Dscheduler.disabled=true`, which disables the OIM job scheduler.
  Removing this option enables the scheduler, as follows.

- Click the **Lock and Edit** button on the left pane.

- On the "Arguments" box, de-select the property `-Dscheduler.disabled=true`, and click Save.

- Click the **Activate Change** button on the left pane.

- Return to the main Servers page and restart the OIM managed server for the changes to take effect.

2.9.6 Enabling/Disabling Oracle Federation

To disable Oracle Identity Federation

1. In a browser, go to the OAM Console, at:
`http://ADMINVHN.mycompany.com:7001/oamconsole`

2. Log in as the Oracle Access Manager user identified in the provisioning response file. (When provisioning IDM, a variety of users were automatically created, and three of the user names could be edited to more user-friendly titles within the response file. If you followed this editing process, then the following user names were potentially affected:
`#IDStore UserNames Configuration`
`IDSTORE_OAMADMINUSER=oamadmin`

```
IDSTORE_OAMSOFTWAREUSER=oamLDAP
IDSTORE_OIMADMINUSER=oimLDAP)
```

3. Navigate to **Policy Configuration -> Shared Components -> Authentication Schemes -> FAAuthScheme**.

4. Set the Challenge Method to `FORM`.

5. Set the Authentication Module to `SaaSModule`.

6. Set the Challenge URL to `/pages/oamLogin.jsp`.

7. Set the Context Type to `customWar`.

8. Set the Context Value to `/fusion_apps`.

9. Set the Challenge Parameters field with the following entries:
```
federationEnabled=false
ssoChooserEnabled=false
fedSSOEnabled=false
initial_command=NONE
TAPPartnerId=OIFDAPPartner
TAPChallengeURL=https://SSO.mycompany.com:443/fed/user/spoam11g
```

The user should be able to log into Fusion Applications.

Subsequently, the user can re-enable OIF if desired.

To enable OIF:

1. In a browser, go to the OAM Console, at:
`http://ADMINVHN.mycompany.com:7001/oamconsole`

2. Log in as the Oracle Access Manager user identified in the provisioning response file. (When provisioning IDM, a variety of users were automatically created, and three of the user names could be edited to more user-friendly titles within the response file. If you followed this editing process, then the following user names were potentially affected:
```
#IDStore UserNames Configuration
IDSTORE_OAMADMINUSER=oamadmin
IDSTORE_OAMSOFTWAREUSER=oamLDAP
IDSTORE_OIMADMINUSER=oimLDAP)
```

3. Navigate to **Policy Configuration -> Shared Components -> Authentication Schemes -> FAAuthScheme**.

4. Set the Challenge Method to `FORM`.

5. Set the Authentication Module to `SaaSModule`.

6. Set the Challenge URL to `/pages/oamLogin.jsp`.

7. Set the Context Type to `customWar`.

8. Set the Context Value to `/fusion_apps`.

9. Set the Challenge Parameters field with the following entries:
```
federationEnabled=true
ssoChooserEnabled=false
fedSSOEnabled=true
initial_command=NONE
TAPPartnerId=OIFDAPPartner
TAPChallengeURL=https://SSO.mycompany.com:443/fed/user/spoam11g
```

Part II

Production-to-Test Data Movement

This section describes all the steps in moving application data from one Fusion Applications instance to another matching instance.

This part contains the following chapter:

-

3

Perform Production-to-Test Data Movement

This chapter contains the following high-level sections:

- Section 3.1, "Introduction"
- Section 3.2, "Prerequisites"
- Section 3.3, "Fill Out the Discovery Workbook"
- Section 3.4, "Reconcile Identity Management Data (IDM)"
- Section 3.5, "Move Fusion Applications Data"

3.1 Introduction

This chapter contains the start-to-finish steps for transferring data from a source Oracle Fusion Applications instance onto an existing destination Fusion Applications instance. This is separate from the cloning procedures described in Part I.

3.1.1 What is Production-to-Test Data Movement?

"Production-to-test" is the movement of application data from a source to a target Fusion Applications installation. Although a common use case is the refreshing of a test database with production data, the same tools could be used to move data between any two environments (production, staging, testing, etc.). In this chapter, "production" is assumed to be source, and "test" is assumed to be the target.

3.1.1.1 What is Moved in Production to Test?

There are two phases in moving data in a Fusion Applications installation: 1) moving the Identity Management Identity and Policy Store data, and 2) moving data from the Fusion Applications transaction database(s). At a high level, the following are moved:

- Identity Management Policy Store data (application and system policies, but not credentials and keys)
- Identity Management Identity Store data (not including AppID and user passwords)
- Fusion Applications transaction data and the crawl index stored in SES
- File attachments stored in UCM (such as orders, agreements)
- ADF Customizations (such as Flex Fields), SOA and ESS customizations stored in MDS
- Business Intelligence (BI) Web Catalog and RPD
- ODI repository

- WebCenter contents

Production-to-test movement replaces most of the target database with production data; a small category of data on the test/target system is preserved, as required by the system. When the content is moved, the target environment is reconfigured and rewired. All long-running processes on the target are stopped and purged, in order to prevent the non-production system from sending emails and alerts to real users, as if it were the production system.

3.1.1.2 Terminology

Common terminology used in production-to-test data movement includes:

Source Environment - In data movement, the source environment is a fully provisioned Fusion Applications environment with data that will be replicated to another existing environment. The source environment may be used for production, thus the term "production-to-test."

Target Environment - The target environment (which may be used for testing) is a matching Fusion Applications instance to the source. It will have its transaction data overwritten by the source data.

Content Movement - A general term that refers to the task of moving Fusion Applications components and/or data from one environment to another environment.

Abstract Host Name - An abstract host name is an alias given to represent a physical node. It has a one-to-one relationship with a virtual host name. If your environment was installed before the release of cloning and done without the use of abstract host names, the virtual host names in your source environment will become abstract names in the destination environment. If your source environment did not make use of virtual host names, then physical host names will be used.

3.1.2 Roadmap: What Does Production to Test Data Movement Entail?

Production-to-test data movement requires the following steps:

- **Fulfill Prerequisites** and download the production-to-test tools. See Section 3.2.

- **Complete Discovery:** Fill in the in-depth notation of the source and destination topology and configuration details, with entries typed into the P2T tabs of the *Discovery Workbook for Cloning and Content Movement*, provided. See Section 3.3.

- **Move Identity Management data** using a five-step process. See Section 3.4. This step must be completed before moving the Fusion Applications data.

- **Move Fusion Applications data** from production to test, while also exporting and re-imported selected test data that must be preserved. See Section 3.5.

3.2 Prerequisites

The following assumptions are made for production-to-test data movement:

- Source and target systems must be identical in terms of product version, initial patches, deployment topology, and configurations. The same applies to their respective databases.
  Note: There are required patches for production-to-test that need to be applied only to the target system. At that point, the patching for the two systems will no longer be identical.

- Both systems were set up following the same set of instructions.

> **Note:** The procedures in this book are NOT designed for Oracle Fusion Applications systems that were *installed* using OVM templates. If your source system was installed in this way, contact Oracle Support for the correct production-to-test documentation and procedures.
>
> If you used virtualization technology, such as Oracle Virtual Machines, to host an operating system, but performed full standard provisioning into that virtualization layer, then the procedures in this book CAN be used. Both source and target systems must match.

- The OS version and configurations are identical in both environments.

- Internal hostnames are identical in both environments.

- The directory paths and structures are identical in both environments.

- Both source and target environments are available for access over SSH.

- The host and port of both OID stores are accessible for data movement.

- "The name values for "IDM_JPSROOT" and "FS_JPSROOT" values must be identical between source and target systems. If they are not, refer to Appendix D, "Change JPS Root Name as Needed".

3.2.1 System Requirements

Versions: Both production and test installations must be on matching versions of **Oracle Fusion Applications.** Check the title page of this guide for the correct software version; to use this guide, the software and guide versions must match.

The starting versions of the two environments must be identical in terms of patching. The additional patches listed for production to test can be applied to the target system only.

Operating system: The source and destination environments must be installed on Linux only (versions certified for Oracle Fusion Applications). (This requirement excludes the database, which can run on all supported OSs.)

3.2.1.1 Required Patches for Production-to-Test on Target Environment

There are patches specific to production-to-test that must be installed on the Identity Management and Fusion Applications servers. Check the Release Notes for the current list of patch numbers to be installed.

3.2.2 Obtaining the Production-to-Test Tools

If Fusion Applications is installed on multiple servers, you can install the production-to-test kit in shared storage with identical mapping from all the servers of the Fusion Applications environment. (This includes the Identity Management environment).

Download and unzip the patch for this release of P2T to `$ORACLE_BASE`

Run the following command:

```
java -jar FAP2T_11.1.7.0.0_generic.jar
```

3.3 Fill Out the Discovery Workbook

The discovery phase may be the most important part of the data movement process. Here you determine all the relevant details of your source and destination environments, and record them. Note that the details required for production-to-test data movement are different than those for Cloning, and have their own tabs in the Workbook.

3.3.1 Using the Discovery Workbook

The *Oracle Fusion Applications Discovery Workbook for Cloning and Content Movement* is a required companion document to this User Guide. It is used to help you research and annotate every aspect of your source and destination Fusion Applications environments. Fill in the P2T tabs in the Workbook; you will then copy/paste the entries to complete the `p2t.rsp` response file appropriately.

3.3.1.1 Where to Find `provisioning.rsp` and `provisioning.plan`

The best resource for many of the Workbook entries is the `provisioning.rsp` file. For some data, it is also necessary to refer to `provisioning.plan`.

Both files may be located in the same directory: (`APPLICATIONS_BASE/provisioning/plan/`).
If the `.rsp` file is not in the `/plan` directory, search for `provisioning.setup.core.provisionplan.install` within `provisioning.plan`, to see where the `.rsp` file is located.

3.3.2 Prepare for the Discovery Phase

The Workbook gives some shorthand tips on where to find things or how to enter them, but this section of the User Guide provides much more guidance.

To begin, open the Discovery Workbook and proceed through the three tabs of data you are asked to collect. They are organized as follows:

- Section 3.3.3, "P2T Identity Management"

- Section 3.3.4, "P2T Fusion Applications"

- Section 3.3.5, "P2T Passwords"

The last tab is special; it automatically collates the data from the rest of the tables and organizes them for ease of use in the `p2t.rsp` response file. It is:

- Section 3.3.6, "Generated P2T RSP Entries"

3.3.3 P2T Identity Management

There are three tables in the P2T Identity Management tab. The following sections give tips on finding the correct values for each row in the tables.

3.3.3.1 IDM Database Information (Source and Target)

The IDM database administrator should know the host names, service names, port numbers, and schema names for the OID and OIM on the target and source environments. Enter in the appropriate tabs. When a field is marked N/A, this entry will not be needed by the P2T script, and can be omitted.

The IDM database configuration accepts either the database SID or Service name. If the SID name is provided, then it will use SID name irrespective of Service name.

Therefore, if you want to use the Service name, ensure the SID name has no value (left blank in the Workbook).

3.3.3.2 IDM Midtier Information (Production/Source)

- **OID Hostname:** Enter the physical host name for the server where the OID resides on both source and target.

- **OID Port:** If you need to locate this information, perform a file system search for the `ports.properties` file: `$OID_INSTANCE/config/OPMN/opmn/ports.prop`. Search for `/oid1/oid1_nonSSLPort=` to find the number.

- **OVD Port:** If you need to locate this information, perform a file system search for the listeners xml file: `$OVD_INSTANCE/config/OVD/ovd1/listeners.os_xml`. Search for `<ldap id="LDAP Endpoint" version="1">`; the port number is listed immediately below.

- **JPS Config Directory:**

  This section proceeds in several parts.

  TARGET: To fill in the Target value, find `fmwconfig` on IDM domain home. (`$IDM_DOMAIN_HOME/config/fmwconfig`). Enter this value in the Target column.

  SOURCE: This entry is unusual and requires several steps. The entry in the Source column is a pointer to a temporary directory you will create on the Target. Into this directory, you will copy three items from the Source `fmwconfig`, as follows:

  1. Create a temporary directory on the target, such as `/tmp/config/fmwconfig>`. Enter this value in the Source column.

  2. Search the source system for IDM domain home: `$IDM_DOMAIN_HOME/config/fmwconfig`. In this directory are 1)the `bootstrap` directory, 2) `jps-config.xml`, and 3) `jps-config-jse.xml`.

  3. Copy these three items into the directory that was created on the target in step 1, such as `/temp/config/fmwconf`.

  Note: When you have completed the whole P2T process, delete the `/temp` directory from the target environment.

- **IDM Admin Server Path/ IDM_DOMAIN_HOME:** The path to this domain home can be found from Fusion Middleware Control, if needed.

3.3.3.3 IDM Midtier Information (Test/Target)

- **IDM Java Home Directory:** This is typically located in the OIM Middleware home.

- **OIM Middleware Home:** Log in to the FMW Control for IDM and go through the Topology for each component. The Oracle Homes and Instance Homes for each IDM component are listed and can be entered into the Workbook.

  Search FMW control in the same way for **OIM Oracle Home**, **OID Oracle Home**, **OIM Managed Server Hostname**, **OIM Managed Server Port**, and **OIM Managed Server Name**, as well.

- **OID Instance Directory:** Enter the path where OID instance resides.

- **OID Instance Name:** If you need to find this value, open the `$OID_INSTANCE/config/OPMN/opmn/opmn.xml` file and search for `ias-instance id="oid1" name="oid1"`

- **OID Component Name:** If you need to find this value, open the `$OID_INSTANCE/config/OPMN/opmn/opmn.xml` file and search for `<ias-component id="oid1" type="OID">`.
  (Be sure that the type is equal to OID; there are other `ias-component-id`s available.)

- **OVD Hostname:**If you need to find this value, return to the listener file. (Perform a file system search for the listeners xml file:
  `$OVD_INSTANCE/config/OVD/ovd1/listeners.os_xml`.) Search for `<ldap id="LDAP Endpoint" version="1">`; the hostname is listed immediately below:
  `<port>6501</port>`
  `<host>LDAPHOST1.mycompany.com</host>`

- **IDM super user (LDAP):** Look for `cn=weblogic_idm` in the Users tree node of ODSM.

- **IDStore Admin User Name:** `cn=oracladmin`. This value is almost always used; if your enterprise changed this value, enter the change. Otherwise, just remove the brackets from the sample value.

- **OAM Admin User Name:** This is the user name used to log into the OAM Console, normally `oamadmin`. To check, look for `cn=oamadmin` in the Users tree node of ODSM. This user should be part of the OAMAdministrators group.

- **OIM Admin User Name:** This is the user name used to log into the OIM Console, normally `xelsysadm`. To check, look for `cn=xelsysadmin` the Users tree node of ODSM.

To find all the following **JPS root** values, log on to the ODSM.

- **FA JPSROOT:** Usually, the provisioning process assigns the name `fa_jpsroot`, or `jpsroot_fa`, or `FAPolicies` (depending on the version you've installed), but it could be given a unique name by your company. To check this value: in ODSM, select the **Data Browser** tab, and check the listed values.
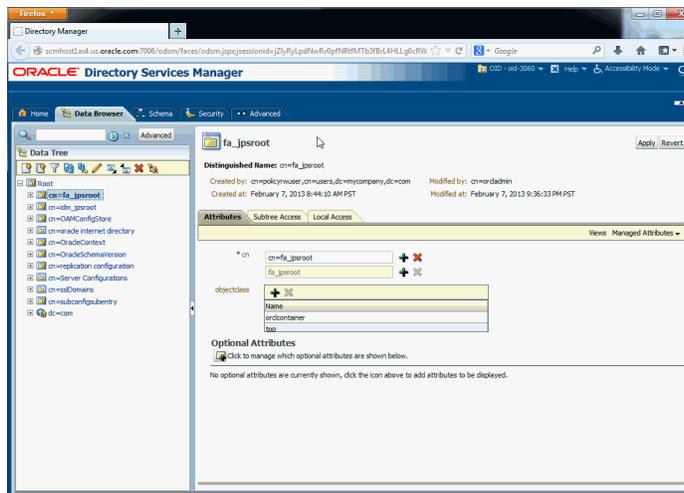
Finding jpsroot information in ODSM, as described in text.

****

- **FA Domain under JPSRoot:** When you've located the FA JPSROOT in ODSM, expand the tree to find the FA Domain.

- **IDM_JPSROOT:** Usually assigned `idm_jpsroot` or `jpsroot_idm`. Check the same Data Browser tab in ODSM to find the correct value for your installation.

- **IDM Domain under JPSROOT:** Expand the listing for `idm_jpsroot` in the ODSM to see the IDM Domain.

- **Base DN:** Look in the ODSM Data Browser data tree and expand the `dc=` to find the full value. The Base DN is everything above the `cn=Users`.

- **Replication ID:** In ODSM Data Browser data tree, expand `cn=replication configuration` until you see `orclreplicaid=`. Highlight this entry to see the full details. The Replication ID is the Distinguished Name at the top of the page (everything after `orclreplicaid=`).

  The **Replication Hostname** is listed at the bottom of the same page in `orclReplicaURI`.

- **TEST_RESET_PWD:** Set value to FALSE. False means that the user passwords from source will not be carried over to the target system. True means that they will, but this is not recommended.

3.3.4 P2T Fusion Applications

The five tables in this tab of the Discovery Workbook include:

- Section 3.3.4.1, "FA DB"

- Section 3.3.4.2, "FA Common Information"

- Section 3.3.4.3, "FA Test/Target Information"

- Section 3.3.4.4, "FA BI Test/Target Information"

- Section 3.3.4.5, "FA BI (Prod/Source) Information"

3.3.4.1 FA DB

The database administrator should be able to enter correct values for the source target environments in this table. Note that if no data pump directories exist, they must be created on the database server of the target system.

3.3.4.2 FA Common Information

- **FA Base Directory (APPLTOP):** If you need to find this value, search `provisioning.rsp` for `INSTALL_ APPHOME_DIR`.

- **FA Java Home:** Check the Fusion Applications installation directory to find/verify the `jdk` directory.

- **Common Domain Home Directory:** This is the path to the domain directory, in the format `<FA Instance home>/domains/<abstract host name of the topology component>/<Domain name>`.

  For example, if the instance home is `/u01/app/fa/instance`, and the abstract hostname for COMMON Admin is `fusionapps.mycompany.com`, then the Admin Server path for Common Domain would be:
  `/u01/app/fa/instance/domains/fusionapps.mycompany.com/CommonDomain`

- **FA Super User Name:** If you need to find this value, search `provisioning.rsp` for `IDENTITY_SUPERUSER`.

3.3.4.3 FA Test/Target Information

- **P2T Working Directory:** Enter the directory you created when extracting the production-to-test tools. See Section 3.2.2 for an example of the P2T Working Directory.

- **Common Domain Host Name:** Search the `provisioning.rsp` file for the `#Domain Topology` to get the CommonDomain hostname.

- **T3 URL Entries:** For all the T3 URL entries, search the `provisioning.rsp` file for the `#Domain Topology`. This will list each host name and port; concatenate them to create the full entry, using the format: `t3://<hostname>:<port>`.
  NOTE: If you do not have all products installed, and therefore domain does not exist, use NONE as a value. Do NOT delete or leave empty.

  This applies to *Common Domain T3 URL*, *CRM Domain T3 URL*, *HCM Domain T3 URL*, *SCM Domain T3 URL*, *FIN Domain T3 URL*, *Project Domain T3 URL*, *Procurement Domain T3 URL*, and *IC Domain T3 URL*.

- **SES and ESS Entries:** Log on to the Common Domain Admin Console. Go to Servers to find the SES and ESS information. This applies to *Common Domain SES (Secure Search Server) Hostname*, *Common Domain SES (Secure Search Server) Port Number*, and *Common Domain ESS Server Name*.

3.3.4.4 FA BI Test/Target Information

- **BI Machine OS User Name:** this is the user that installed the Business Intelligence domain on the BI server.

- **BI Domain Home Directory:** This is the path to the domain directory, in the format `<FA Instance home>/domains/<abstract host name of the topology component>/<Domain name>`.

- **BI Admin Server Host Name** and **BI Admin Server Port:** go to OHS `moduleconf` directory and view `FusionVirutalHost_bi.conf`. The BI Admin Server hostname and port can be found under **Connect** roots for `BI Weblogic: <Location /Console>`.

- **Broker Hostname** and **Broker Port:** information: For all entries, access `/u01/oracle/fa/config/CommonDomain_webtier/config/OHS/ohs1/moduleconf`

  For **Financial Broker** information, open the file `FusionVirtualHost_fin.conf`. Search for`#Internal virtual host for fin<VirtualHost fininternal.mycompany.com:20603 >` This gives you the Financial Broker Port and Hostname.

  For **CRM Broker** information, open the file `FusionVirtualHost_crm.conf`. Search for `#Internal virtual host for crm <VirtualHost crminternal.mycompany.com:20615 >`.

  For **HCM Broker** information, open the file `FusionVirtualHost_hcm.conf`. Search for `#Internal virtual host for hcm <VirtualHost hcminternal.mycompany.com:20619 >`.

- **FA DB Host Names** and **Ports:** For single-instance environments, the database administrator can fill this in. In the case of a RAC installation, you must enter all instances of the database in an escape semi-colon-separated list.
  For example:1521\;1522.

3.3.4.5 FA BI (Prod/Source) Information

- **BI RPD Directory:** Use the format
  `<FA Instance home>/BIInstance/bifoundation/OracleBIServerComponent/`
  `coreapplication_obis1/repository>`

- **BI Domain Home Directory:** This is the path to the domain directory, in the
  format `<FA Instance home>/domains/<abstract host name of the topology`
  `component>/<Domain name>`.

- **BI Admin Server Hostname** and **Port:** Search the `provisioning.rsp` file for the
  `#Domain Topology`. This will list each host name and port.

- **Broker Hostname** and **Broker Port** information: For all entries, access
  `/u01/oracle/fa/config/CommonDomain_webtier/config/OHS/ohs1/moduleconf`

  For **Financial Broker** information, open the file `FusionVirtualHost_fin.conf`.
  Search for `#Internal virtual host for fin<VirtualHost`
  `fininternal.mycompany.com:20603 >` This gives you the Financial Broker Port
  and Hostname.

  For **CRM Broker** information, open the file `FusionVirtualHost_crm.conf`. Search
  for `#Internal virtual host for crm <VirtualHost`
  `crminternal.mycompany.com:20615 >`.

  For **HCM Broker** information, open the file `FusionVirtualHost_hcm.conf`. Search
  for `#Internal virtual host for hcm <VirtualHost`
  `hcminternal.mycompany.com:20619 >`.

- **UCM Weblayout Directory:** Use the format `<FA Instance home>/<abstract host`
  `name>/CommonDomain/ucm/cs/weblayout/>`.

- **UCM Vault Directory:** Use the format `<FA Instance home>/<abstract host`
  `name>/CommonDomain/ucm/cs/vault/>`

- **BI Webcat Directory:** Use the format `<FA Instance home>/<abstract host`
  `name>/CommonDomain/ucm/cs/catalog/>`.

3.3.5 P2T Passwords

This tab is informational only; do not enter values in the fields! These are the
passwords that will be required during production-to-test movements.

3.3.6 Generated P2T RSP Entries

This tab organizes all your entries and presents them so they are easy to use. The
Generated RSP Entries tab collates the data entered in all the other tabs and tables, and
generates the entries and values as they should be entered in the `p2t.rsp` file. The file
is located in `$working_dir/utilhome/bin/p2t.rsp`. When discovery is finished,
transfer the generated RSP entries to the .rsp file as follows:

1. Locate the p2t.rsp sample file in the package of production-to-test materials.

2. Replace the contents with the Generated P2T RSP Entries by copying all and
   pasting into the sample file.

3. Search for properties ending in `_PASSWORD` and `_PWD`. These are password entries
   for the target environment and are not auto-filled, since they are not entered into
   the Workbook for security reasons. Manually enter each password value for the
   target environment into the `p2t.rsp` file.

4. Save the file; it will be used when the production-to-test scripts are run.

3.4 Reconcile Identity Management Data (IDM)

In production-to-test for Identity Management, the application users and roles are migrated from source to target, but the passwords are not. Therefore, the system administrator would need to set new passwords on the target system for each newly migrated user who did not already exist on the target.

To prepare Identity Management data is done in the following steps:

- Section 3.4.1, "Validate on the Target OID Server"
- Section 3.4.2, "Disable Reconciliation Jobs on Target OIM"
- Section 3.4.3, "Reconcile OID Directories"
- Section 3.4.4, "Re-enable Reconciliation Jobs on Target OIM"
- Section 3.4.6, "Run Validation Script"

Note: If OID and OIM are on same host, all the steps can be executed from that server. If they are on different hosts, then the OID-related steps must be done on the OID server, and the OIM and OPSS (Policy and Identity Store) steps on the OIM host.

Before beginning any of the steps, it is necessary to complete the Discovery Workbook (Section 3.3). Using the Generated P2T Response File tab in the Workbook, you also modify the `p2t.rsp` file, located in `$WORKDIR/utilhome/bin directory`. This file will be used throughout the production-to-test process on both Identity Management and Fusion Applications.

3.4.1 Validate on the Target OID Server

In this step, the production-to-test tool connects to the source and target OIDs, based on the entries in `p2t.rsp`, to ensure that the host, port and credential information are all correct.

1. Shut down the Fusion Applications stack. The Fusion Applications database must be up.

2. Run the following step on the OID server:

   ```
   cd $WORKDIR/utilhome/bin
   ./bmIDM.sh validate
   ```

3. Fix any errors until validation is complete.

3.4.2 Disable Reconciliation Jobs on Target OIM

This step disables reconciliation jobs on the target OIM, before data movement begins. Sample reconciliation jobs would include: "LDAP User Create and Update Reconciliation", "LDAP Role Create and Update Reconciliation", "LDAP Role Membership Reconciliation", "LDAP User Delete Reconciliation", "LDAP Role Delete Reconciliation", etc.

Please execute the following steps on the target OIM:

```
cd $WORKDIR/utilhome/bin
./bmIDM.sh oimpre
```

3.4.3 Reconcile OID Directories

This step compares the source and target OID directories and reconciles any discrepancies by merging any differing attributes from the source to the target.

Execute the following steps on the target OID server:

```
cd $WORKDIR/utilhome/bin
./bmIDM.sh oid
```

3.4.4 Re-enable Reconciliation Jobs on Target OIM

When the OID directories have been reconciled, it is possible to restart the reconciliation jobs. Execute the following steps on the target OIM server:

```
cd $WORKDIR/utilhome/bin
./bmIDM.sh oimpost
```

3.4.5 Move Identity and Policy Store Data

This step migrates application-specific policies and system policies from the source to the target.

```
cd $WORKDIR/utilhome/bin
./bmIDM.sh opss
```

3.4.6 Run Validation Script

Execute the following step to check that source and target system are valid. You should also check the `p2t_validate.log` file on the working directory.

```
cd $WORKDIR/utilhome/bin
./bmIDM.sh validate
```

When validation is complete, then:

1. Restart the target Identity Management stack.

2. Start the Fusion Applications stack that was shut down in section Section 3.4.1. All the domains and managed servers must restart successfully. Access critical admin console and application URLs to log in successfully, but do not do any functional testing or transactions until the next phase is complete Section 3.5, "Move Fusion Applications Data".

3.5 Move Fusion Applications Data

Production-to-test movement for the transaction data includes the following steps. For each command, run `preverify` and correct any errors until preverify passes, then execute `run`.

- Section 3.5.1, "Run Scripts to Pack Source Files"

- Section 3.5.2, "Transfer Files to Target Servers"

- Section 3.5.3, "Export Specific Data from Target (Test) System"

- Section 3.5.4, "Replace Target Data with Source Data"

- Section 3.5.5, "Move Packed Data and Clean Up In-Flight Transactions"

3.5.1 Run Scripts to Pack Source Files

The production-to-test script must be installed on the production (source) server. You run packing scripts on the primordial Fusion Applications server and the Business Intelligence server.

When the packing is complete, the following files will reside on the Production working directory: `weblayout.tgz`, `vault.tgz`, `webcatalog.tgz`, `obirpd.tgz`, and `rpdAttributes`. DO NOT change the file/directory names!

3.5.1.1 Core Files

Run the following commands to pack core files on the primordial production server:

```
$WORKING_DIR/utilhome/bin/packData.sh <preverify or run>
```

Note: The "preverify" commands validate the environment and connectivity only.

3.5.1.2 Business Intelligence (BI) Files

Run the following commands to pack Business Intelligence files on the BI production server. The script also connects to the BI Admin Server to extract the password of the credential map "`oracle.bi.enterprise`" and key "`repository`" in the Production server

```
$WORKING_DIR/utilhome/bin/packDataBI.sh <preverify or run>
```

Note: The "preverify" commands validate the environment and connectivity only.

3.5.2 Transfer Files to Target Servers

Transfer the packed files from Section 3.5.1 in the following way:

To the target/test primordial Fusion Applications server, transfer `weblayout.tgz`, `vault.tgz`, and `webcatalogTr.tgz`.

To the target/test BI server working directory, transfer `obirpd.tgz`, `rpdAttributes`

On the target BI Server working directory:

```
mkdir unpackrpd
cd unpackrpd
tar xzf ../obirpd.tgz
```

Note: The file/directory names must be exactly as documented.

3.5.3 Export Specific Data from Target (Test) System

The step preserves some of the data on the target system which will be re-imported after the production data is migrated.

1. Under the Fusion Applications working directory, run the command `mkdir datapumpdir_db` .

2. Ensure that the Functional Setup Server_1 in the Common domain is running, and the SOA servers in all domains are up and running.

3. Run the command

   ```
   $WORKING_DIR/utilhome/bin/generateData.sh <preverify or run>
   ```

   At the end of this process, all the Fusion Applications servers are automatically stopped.

For reference only, the following table lists the tasks accomplished in this step.

Topology Manager	This includes environment-specific topology data that must be preserved on Test, for example, external endpoint information for each deployed domain.
ExportFMWSchemas	The following schemas are backed up: `FUSION_IPM`, `FUSION_OTBI`, `FUSION_BIPLATFORM`, `FUSION_ORASDPLS`, `FUSION_ORASDPXDMS`, `FUSION_ORASDPSDS`, `FUSION_ORASDPM`
SaveExternalWorkflowUrls	Backup front-end host URLs (Human Workflow external OHS configurations
exportAdfConfig	Backup `adf-config.xml`

3.5.4 Replace Target Data with Source Data

Perform a database backup of the source data, using whatever method your enterprise prefers: RMAN backup (along with installing the Oracle RDBMS Server binaries), file system copy, storage replication, VM snapshot, etc.

Adhere to the following requirements when duplicating the source database and mounting it to the destination environment:

- For RAC installations, ensure that the Grid Infrastructure is installed on the destination.

- Before duplicating, shut down the source Fusion Applications Web tier and application tier, as well as the Identity Management Web tier and application tier, and ensure that all in-flight transactions have been completed.

- Shut down the source database cleanly (no abort). The clone/copy must be taken cold.

- Remember that the topology and operating systems must be identical between source and destination.

3.5.5 Move Packed Data and Clean Up In-Flight Transactions

This step imports data from Section 3.5.1, Section 3.5.2, and Section 3.5.3 back in to the target (test) database. All long-running processes will be stopped and purged, to prevent the non-production system from sending emails or notifications to real users as if it were a production system.

There are three scripts to be run on the primordial Fusion Applications server, and two Business Intelligence scripts to run on the BI server.

- While Fusion Applications is shut down, apply changes to the test environment:

  ```
  $WORKING_DIR/utilhome/bin/applyDataOffline.sh <preverify or run>
  ```

- Restart the servers in the Fusion Applications stack:

  ```
  $WORKING_DIR/utilhome/bin/startAllServersMT.sh <preverify or run>
  ```

  (It is also possible to use your own script/process to start the Fusion Applications stack.)

- While Fusion Applications is running, apply changes to the test environment:

  ```
  $WORKING_DIR/utilhome/bin/applyDataOnline.sh <preverify or run>
  ```

- If Business Intelligence is installed on a different server, run the following command from the BI server

```
$WORKING_DIR/utilhome/bin/applyDataBI.sh <preverify or run>
```

■ If Informatica is installed, run the following command from the server where it is installed:

```
$WORKING_DIR/utilhome/bin/applyDataIIR.sh <preverify or run>
```

3.5.5.1 Validate

After completing the Fusion Applications production-to-test steps, restart the Fusion Applications stack again. All domains and managed servers must restart successfully. The system is ready for functional testing.

Part III

Appendices

Part III includes the following Appendices used in Cloning:

- Appendix A, "Post-Clone Data Masking (Optional)"

- Appendix B, "Best Practices: Install with Cloning in Mind"

- Appendix C, "Abstract Hostnames in Detail"

- Appendix E, "Install the BI Administration Tool"

Production-to-test data movement uses the following Appendix:

- Appendix D, "Change JPS Root Name as Needed"

A

Post-Clone Data Masking (Optional)

Although Data Masking is described in other guides, this section gives targeted tips on performing the process after cloning. It is an optional step.

A.1 Introduction

This document covers the procedures for configuring and running Data Masking for Oracle Fusion Applications. For reference and more detail, see:

- The Data Masking section of the "Securing Oracle Fusion Applications" chapter in the *Oracle Fusion Applications Administrator's Guide*.

- For Data Masking Tutorial:
  http://www.oracle.com/webfolder/technetwork/tutorials/obe/em/emgc10205/data_masking/datamask.htm

- Oracle Data Masking Overview Demo:
  http://www.youtube.com/watch?v=dcldnaukpcI

A.1.1 Data Masking Requirements

Post-clone data masking requires the following:

- Oracle Fusion Applications 11*g*

- Oracle Database Release 2 for Oracle Fusion Applications

- Oracle Enterprise Manager Grid Control 11*g* or Oracle Enterprise Manager Cloud Control 12*c* or Oracle Enterprise Manager Database Control for Oracle Database

- The Oracle Data Masking Pack

A.1.2 Data Masking Format Library

The Data Masking Format Library contains the data format and functions for masking. This library must be installed in the target database where data masking will be run. Oracle provides 2 SQL scripts for installing the Data Masking Format Library. These SQL scripts can be found in the database ORACLE_HOME:

```
$ORACLE_HOME/sysman/admin/emdrep/sql/db/latest/masking/dm_fmtlib_pkgdef.sql
$ORACLE_HOME/sysman/admin/emdrep/sql/db/latest/masking/dm_fmtlib_pkgbody.sql
```

A.1.3 Data Masking Definitions

Oracle Fusion Applications provides out-of-the-box masking definitions for each product family. These definitions specify the list of sensitive database tables and columns, along with the data formats to be used to mask these columns. The masking

templates can be found on Oracle Fusion Applications nodes in the Oracle Fusion Applications Middleware Home (`APPLICATIONS_BASE/fusionapps`)

- For Oracle Enterprise Manager 11*g*:

  ```
  APPLICATIONS_BASE/fusionapps/atgpf/sysman/dataMasking/Mask_Oracle_Fusion_
  Applications_1.0_EM_11.1.0.1.0_Combined_Template.xml
  ```

- For Oracle Enterprise Manager 12*c*:

  ```
  APPLICATIONS_BASE/fusionapps/atgpf/sysman/dataMasking/ADM_Oracle_Fusion_
  Applications_1.0_EM_12.1.0.1.0_Combined_Template.xml
  ```

  ```
  APPLICATIONS_BASE/fusionapps/atgpf/sysman/dataMasking/Mask_Oracle_Fusion_
  Applications_1.0_EM_12.1.0.1.0_Combined_Template.xml
  ```

A.1.4 Preliminary Steps

Before using Data Masking, perform the following on the target database:

1. Install the Data Masking Format Library in the target database by executing the following 2 SQL scripts using SQL*Plus. Connect as `SYS` or a user that can create packages in the `DBSNMP` schema.

    - `$ORACLE_HOME/sysman/admin/emdrep/sql/db/latest/masking/dm_fmtlib_pkgdef.sql`

    - `$ORACLE_HOME/sysman/admin/emdrep/sql/db/latest/masking/dm_fmtlib_pkgbody.sql`

2. Change all the temp spaces used for data masking to **Auto Extend** and Increment by 150MB and Maximum File Size to 32767MB. The amount of additional space you require depends on the amount of data being masked. Data masking takes up approximately two times the size of the largest table being masked.

3. Ensure that sufficient free space is available for the database before executing the masking job.

4. The user executing the data masking script must have the `dba` role.

A.2 Importing Data Masking Definitions

Before using Data Masking, import the out-of-the-box Data Masking Definitions provided by Oracle Fusion Applications.

### A.2.1 For Oracle Enterprise Manager Grid Control 11*g*

1. Login to Oracle Enterprise Manager Grid Control as `SYSMAN`.

2. Click the **Targets** tab, then click the **Database** subtab.

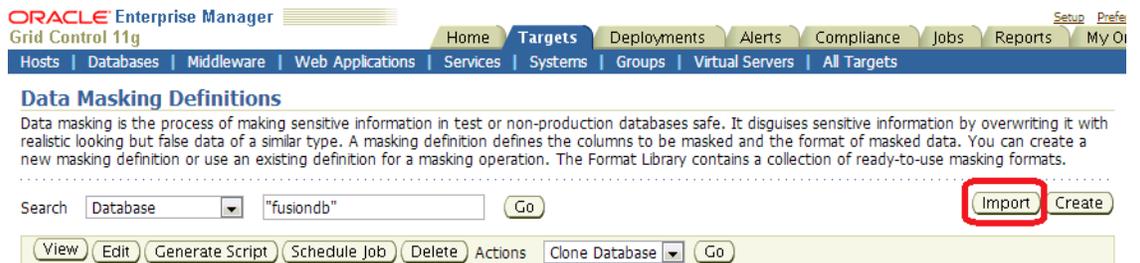3. Select the target database where data masking will be run.

EM Grid Control interface, Targets tab selected, showing databases.

****

4. Click the **Schema** tab, then click **Definitions** in the Data Masking section.
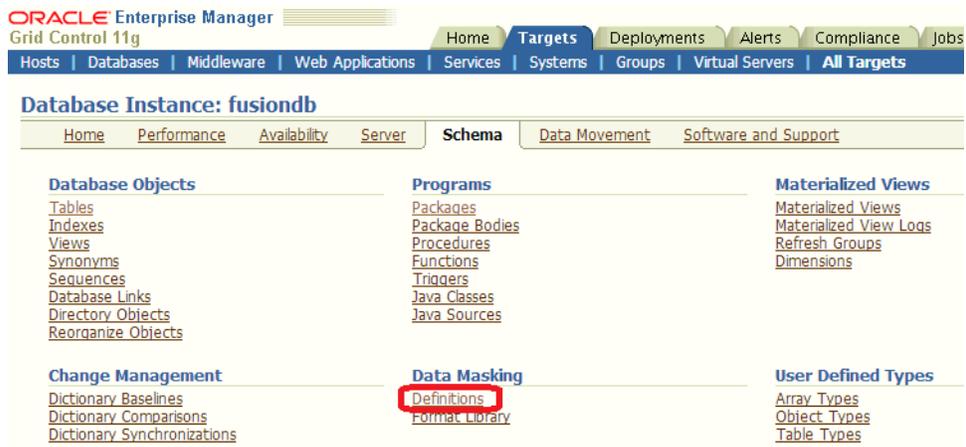
5. Click **Import**.

EM Grid Control, Targets tab, showing where to import data masking definitions for a Fusion Applications database.

****

6. Make sure the masking template XML file is accessible from the browser by copying the XML file to the machine from which you are browsing.

This XML file is available on Oracle Fusion Applications nodes at: `APPLICATIONS_ BASE/fusionapps/atgpf/sysman/dataMasking/Mask_Oracle_Fusion_ Applications_1.0_EM_11.1.0.1.0_Combined_Template.xml`.
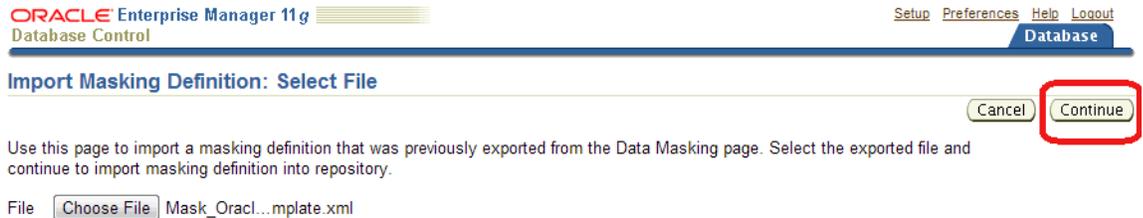
7. Edit the XML file and change the database name in line 3 to the target database name. For example, change

```
<TARGET_NAME>database</TARGET_NAME>
```

to

```
<TARGET_NAME>fusiondb</TARGET_NAME>
```
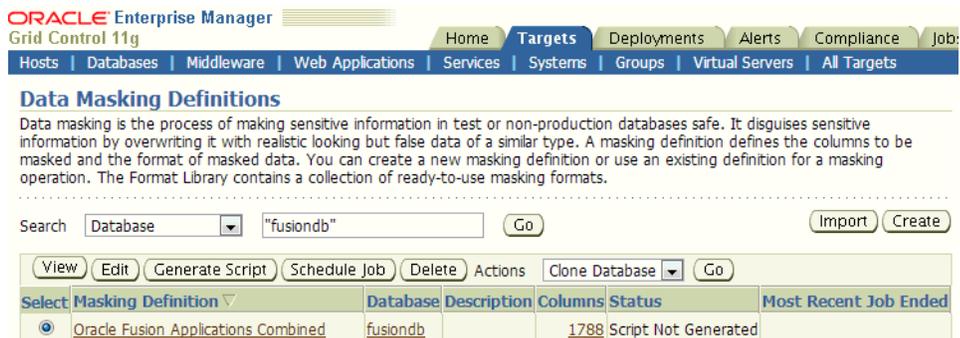
8. In the browser, click **Choose File** and select the Data Masking Template XML file that you edited in Step 7.

9. Click **Continue**.

10. When the Masking Definition is imported successfully it is shown in the Masking Definition table.

11. Click the **View** button to view the tables, columns and format in the Masking Definition.

12. Click the **Edit** button to customize the Masking Definition.

13. Click **Generate Script** to generate the SQL script for running the data masking job. This might take some time. When the script is generated, the data masking job can be scheduled to run on the target database by clicking **Schedule Job**.

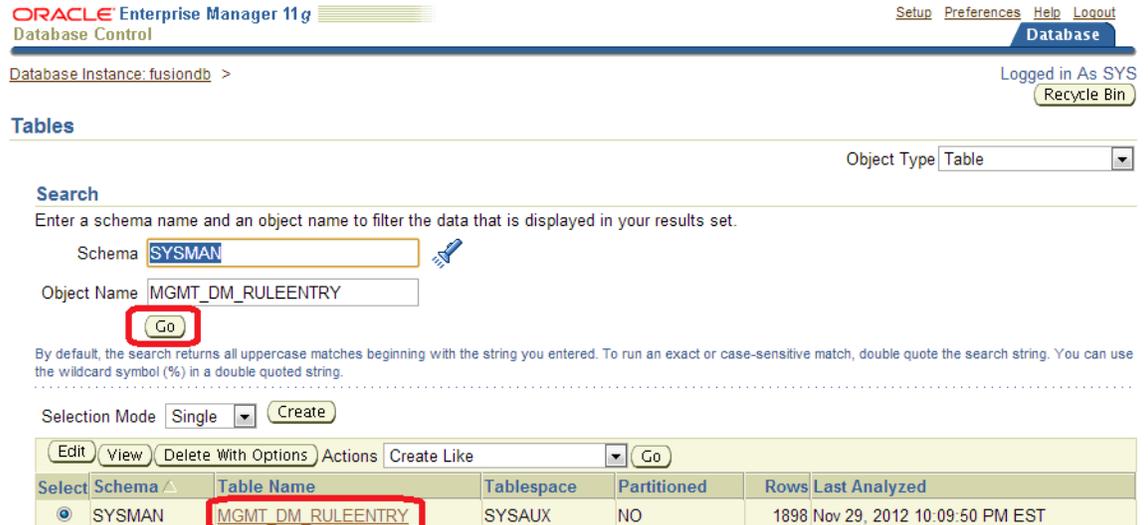## A.2.2 For Oracle Enterprise Manager 11*g* Database Control

1. Login to Oracle Enterprise Manager Database Control as SYS in SYSDBA role.

2. Modify the SYSMAN.MGMT_DM_RULEENTRY table and increase FIXED_STRING column size to 40:

   a. Click the **Schema** tab, then click **Tables** under the Database Objects heading.

EM Database Control interface, showing Tables selected under Database objects.

****

b. Enter SYSMAN in the Schema field and MGMT_DM_RULEENTRY in the Object Name field and click **Go**.

c. In the search result table, click the **MGMT_DM_RULEENTRY** link.

Database Control, choosing the MGMT_DM_RULEENTRY table to edit.

****

d. Click **Edit**.

e. Change FIXED_STRING column size to 40 and click **Apply.**

Where the actual edits to MGMT_DM_RULEENTRY are made. Columns for editing include a select button, followed by Name, Data Type, Size, Scale, Not Null checkbox, Default Value, and Encrypted checkbox.

****

3. Increase the heap size of the database console:

    a. Stop the database console by running this command:

       ```
       $ORACLE_HOME/bin/emctl stop dbconsole
       ```

    b. Increase the heap size by running this command:

       ```
       $ORACLE_HOME/bin/emctl config dbconsole -heap_size 1024m -max_perm_size
       512m
       ```

    c. Restart the database console by running this command:

       ```
       $ORACLE_HOME/bin/emctl start dbconsole
       ```

4. Login to Oracle Enterprise Manager Database Control as SYS in SYSDBA role.

5. Click the **Schema** tab, then click **Definitions** under the Data Masking heading.
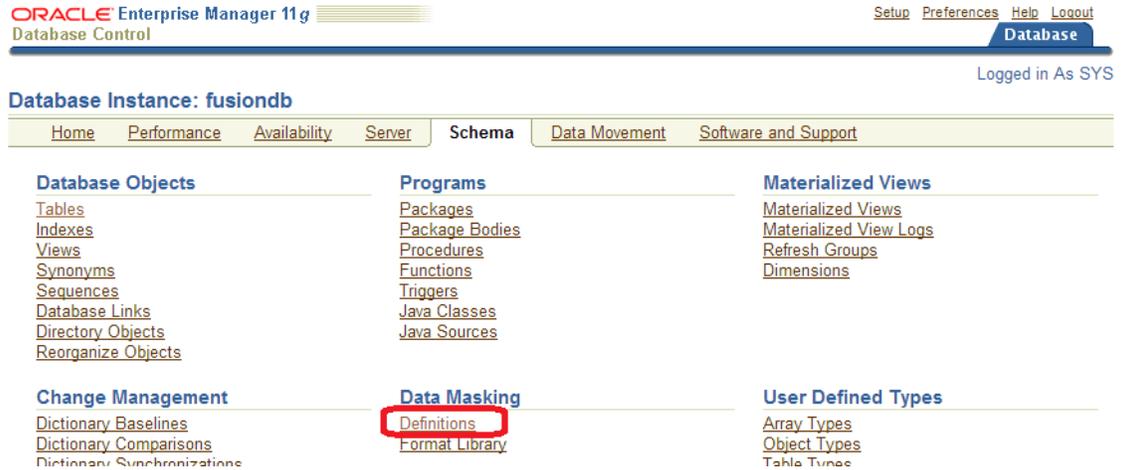
6. Click **Import**.
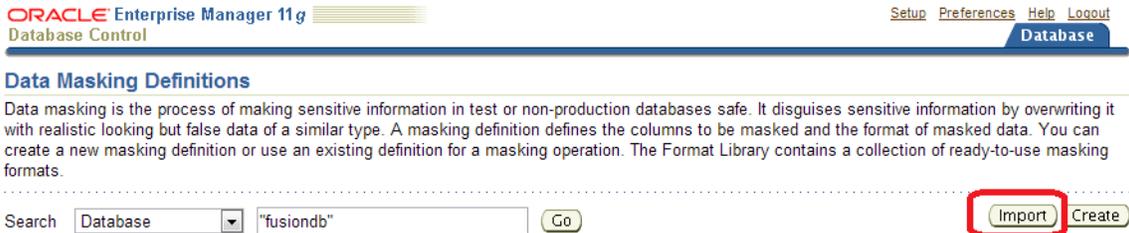
7. Make sure the masking template XML file is accessible from the browser by copying the XML file to the machine from which you are browsing.

   This XML file is available on Oracle Fusion Applications nodes at: `APPLICATIONS_ BASE/fusionapps/atgpf/sysman/dataMasking/Mask_Oracle_Fusion_ Applications_1.0_EM_11.1.0.1.0_Combined_Template.xml`.

8. Edit the XML file and change the database name in line 3 to the target database name. For example, change

   `<TARGET_NAME>database</TARGET_NAME>`

   to

   `<TARGET_NAME>fusiondb</TARGET_NAME>`

9. In the browser, click **Choose File** and select the Data Masking Template XML file that you edited in Step 8.

10. Click **Continue**.

11. When the Masking Definition is imported successfully it is shown in the Masking Definition table.

12. Click the **View** button to view the tables, columns and format in the Masking Definition.

13. Click the **Edit** button to customize the Masking Definition.

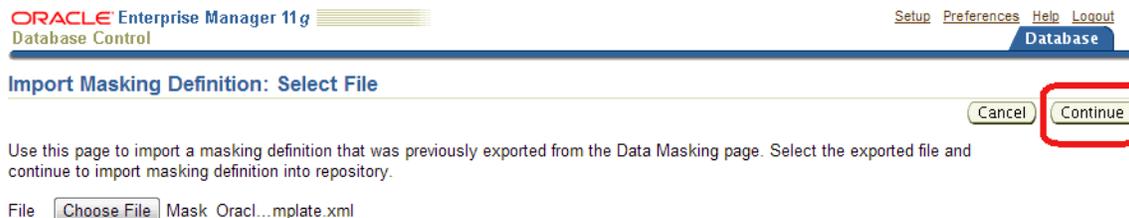14. Click **Generate Script** to generate the SQL script for running the data masking job. This might take some time. When the script is generated, the data masking job can be scheduled to run on the target database by clicking **Schedule Job**.

A.2.3 For Oracle Enterprise Manager Cloud Control 12*c*

1. Using SQL*Plus on the target database, connect as SYS in SYSDBA role.

2. Grant select_catalog_role and DBA to the FUSION user.

3. Login to Oracle Enterprise Manager Cloud Control as sysman.

4. Deploy the TDM package:

   a. Navigate to **Enterprise**, then **Jobs**, then **Job Activity**.

   b. From the OS Command list, select **Deploy Test Data management packages** and click **Go**.

   c. Enter the job name, click **Add**, then add the database where masking is to be run.

   d. Click **Parameters**. From the list, select **Fusion Driver**.

   e. Click **Credential** and enter the FUSION credentials. This user must have all the privileges specified in Step 2.

   f. Click **Submit**.

   g. When the job completes, verify that it succeeded. Review the job details and make sure there are no errors.

5. Import the Application Data Model (ADM):

   a. Navigate to **Enterprise**, then **Quality Management**, then **Data Discovery and Modelling**.

   b. Click **Action**, then **Import**. Provide the ADM XML file, ADM_Oracle_Fusion_ Applications_1.0_EM_12.1.0.1.0_Combined_Template.xml, the ADM name, and the FUSION database. Make sure the ADM XML file is accessible from the browser by copying the XML file to the machine from which you are browsing.

      This XML file is located on Oracle Fusion Applications nodes at: APPLICATIONS_BASE/fusionapps/atgpf/sysman/dataMasking/ADM_Oracle_ Fusion_Applications_1.0_EM_12.1.0.1.0_Combined_Template.xml.

    c. If prompted for database credentials, provide the `FUSION` user credentials.

> **Note:** You might see one or more warnings indicating that duplicate sensitive types were not imported; these warnings can be safely ignored.

6. Create a verification job:

    a. Navigate to **Enterprise**, then **Quality Management**, then **Data Discovery and Modelling**.

    b. Click on the ADM you just imported and click **Verify**.

    c. Click **Create Verification Job**.

    d. Provide a job name and job description.

    e. Click **New Credential** and provide the `FUSION` credentials.

    f. Schedule the job to start immediately and click **Submit**.

7. Check the job status:

    a. When the verification job completes, navigate to **Enterprise**, then **Job Activity**, then select **All Job Status**.

    b. Click **Go** and check that your job completed successfully.

8. Import the masking template file:

    a. Navigate to **Enterprise**, then **Quality Management**, then **Data Masking Definition**. The drop-down list contains data masking definitions and formats.

    b. Import the masking template file, `Mask_Oracle_Fusion_Applications_1.0_ EM_12.1.0.1.0_Combined_Template.xml`. Make sure this XML file is accessible from the browser by copying the XML file to the machine from which you are browsing.

    This XML file is located on Oracle Fusion Applications nodes at: `APPLICATIONS_BASE/fusionapps/atgpf/sysman/dataMasking/Mask_Oracle_ Fusion_Applications_1.0_EM_12.1.0.1.0_Combined_Template.xml`.

    c. Provide the masking definition name.

    d. For the ADM name, specify the name of the `FUSION` ADM you created in Step 5.

    e. For the database name, specify the `FUSION` database.

9. Select the mask definition and click the **Generate Script** button.

> **Note:** Script generation can take a few hours, so consider running it on a terminal that will be available for a while to allow the job to complete. If you run it on a laptop and have to disconnect it during execution, you will lose the browser session that was generating the script.

10. When the job completes, you can view the impact report and save or view the script.

11. Submit the masking script for execution by clicking **Submit Job**.

A.3 Additional Options

This section describes additional data masking options.

A.3.1 Modifying Data Masking Definition

You can modify the Data Masking Definition by adding and deleting column(s) to the Masking Definition and/or changing the masking format of existing definition. You can modify the Data Masking Definition in the following ways:

- Select the definition and click **Edit**. Columns in the masking definition and its masking format are shown.

- Click **Add** to add a table column and format to the definition.

- To remove column(s), select the column(s) and click **Remove**.

- To modify a column format, click the **Format** icon

A.3.2 Generating the Masking Script

When a data masking definition is newly created or imported or modified, the SQL masking script must be generated before data masking can be run on the target database. To generate the script:

1. Select the definition and click **Generate Script**. This may take some time to complete.

2. After the script is generated successfully, you can schedule a data masking job to run on a target database by clicking the **Schedule Job** button. You can save the script for viewing or to run the script on the target database manually. You can also view the Impact Report.

A.3.3 Customizing Mask Formats

To customize Masking Formats, click the **Format Library** link under the Masking Definition table.
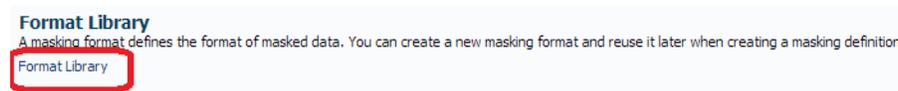
You can view, edit, and delete existing formats. You can also create or import new formats or create a new format based on an existing format using the Create Like feature.

A.3.4 Frequently Asked Questions

Does Masking remove nulls?

Masking will try to preserve null values. When masking a column only the non-null values get masked; null values are left alone.

How are dependent objects such as Intermedia Index and Materialized views masked after the base tables are masked?

Dependent objects are masked automatically by the masking tool.

How does the shuffle format work?

Shuffle format will shuffle the distinct values so the distribution will change. For example, if you had 35 Married, 30 Single, 20 Divorced and 15 Widowed people and you masked them by shuffling, after masking you might have 35 Single, 30 Divorced, 20 Widowed and 15 Married. All the records that had the same value as each other before masking will have the same value after masking. In general almost all masking formats will mask the same values in a table consistently.

Does Masking maintain Data Distribution? Is it prone to inference based attacks?

Masking does not maintain data distribution but it does mask values in a single column consistently. If there were 5 people with the same salary then they will have their salaries masked to the same values. This behavior is true for all mask formats. In this particular case one of the 5 people with the same salary can infer the salaries of the others based on their salary. This is called an inference based attack. Preventing inference based attacks in a generic way is a hard problem. For example there are other such cases like the CEO of a company can usually be identified because he won't have a manager, etc. The only way to avoid inference based attacks is to generalize the rows by taking people in certain salary ranges (or with special circumstances) and mask them to a fixed value (like an average salary range). There is currently no support in masking for generalization. One extreme approach to generalization is to mask all salaries to a fixed value to prevent any chance of inference. Note that inference based attacks can involve using multiple pieces of information in multiple columns and tables to infer a person's identity. For instance, people who had a given medical condition and who are under the age of 30, in a particular department, and who are male may be just one person and it may be possible to identify their row in a table.

For security reasons it is extremely desirable to have a mask function that does not maintain any distribution. In such a case it would not be possible to handle any denormalization rules on the column as the same salary in different rows can be changed to different mask values.

How will related product families can be masked?

For instance, Oracle Fusion Customer Relationship Management masking requires Oracle Fusion Applications Customer Data Management masking. Similarly, Oracle Fusion Human Capital Management and Oracle Fusion Financials require Oracle Fusion Applications Customer Data Management to be masked. An Oracle Fusion Applications database contains all schemas for all product families, even if the customer is not using some of the schemas. Masking unused schemas containing empty tables is not a problem.

How will sensitive information in BLOB, CLOB, columns be masked?

LOB columns can be masked using one of the following formats: Fixed String, Null, or Fixed Number.

What is the difference between a User Defined Function, Post Processing Function, and Pre/Post Masking Script?

A user defined function (UDF) takes the original value, row ID, and column name and generates the mask value. A single-column format can be a combination of one of more formats, including UDFs.

A **post-processing function** (PPF) is a special case of user-defined function. A PPF is called after the mask value is generated using one of the other formats. The PPF will take the generated mask value and further modify it to produce the actual mask value. For example, if the masking format used Random Number (1000,10000) and a post-processing fuction (checksum), a number between 1000 and 10000 is

generated and this value is fed into the PPF. The PPF can compute the checksum and append it to the original number and return the new mask value. There can be only one PPF to a column. Additionally, PPF cannot be the only format for a column; a PPF has to have some other format preceding it.

A post-masking script is not a masking format but a SQL script that executes after all masking completes. It can be use to recompute aggregated columns after the detailed data is masked. This ensures aggregated and masked columns are consistent and the totals match.

A pre-masking script is also a SQL script but it runs before the start of masking. Both the pre- and post- masking scripts execute in the same connection.

B

Best Practices: Install with Cloning in Mind

The procedures in this Appendix contain recommended steps for creating a Fusion Applications environment with the goal of making the cloning process as simple as possible. Some of the recommendations include:

- **Use generic host names and URLs.** Refrain from using environment-specific descriptors, such as "test," "dev," "prod," etc., as these descriptors will get copied to the destination environment. Pay attention to this fact when creating domain names, host names, SIDs/service names, and directory structures.

- **Install without using the local domain/storage option.** That is, use only shared storage, and do not put environment-specific labels in your directory structures or mount points.

This section gives recommendations for the following aspects of Fusion Applications installation:

B.1 Planning the Environment

For demonstration purposes, the sample environment contains all product families, and uses four servers, total: two for Identity Management, and two for Fusion Applications, with web tier and midtier sharing the same host. You can modify the principles of this recommendation to fit your own host and product requirements.

When cloning, remember that the destination topology is only as flexible as the source topology.

B.1.1 Component Allocation

Recommended component allocation is as follows:

Component	Abstract Host Name
Identity Management Database(s)	idmdb.mycompany.com
Identity Management MidTier	idmmidtier.mycompany.com
Identity Management WebTier	idmwebtier.mycompany.com
Fusion Applications Database	fusiondb.mycompany.com
Fusion Applications WebTier	fusionapps.mycompany.com
CommonDomain	fusionapps.mycompany.com

Component	Abstract Host Name
BIDomain	fusionapps.mycompany.com
CRMDomain	fusionapps.mycompany.com
ICDomain	fusionapps.mycompany.com
SCMDomain	fusionapps.mycompany.com
HCMDomain	fusionapps.mycompany.com
FinancialDomain	fusionapps.mycompany.com
ProjectsDomain	fusionapps.mycompany.com
ProcurementDomain	fusionapps.mycompany.com

B.1.2 Host Recommendations

It is highly recommended that all hosts use the same operating system and install user/group (with same IDs). The entries in the table are examples only. To ensure a successful install, follow the minimum hardware and software requirements in the installation documentation

Abstract Host Name	Real Host Name	OS Group
idmdb.mycompany.com	HOST1	oinstall
idmmidtier.mycompany.com	HOST2	oinstall
fusiondb.mycompany.com	HOST3	oinstall
fusionapps.mycompany.com	HOST4	oinstall

B.1.3 Storage Mount Points and Directories

To simplify for cloning, do not use local storage. Do not use environment labels, such as "production" or "test," in your mount points. Do not use soft links to install any of the components. The examples in the table use /u01 as a shared directory among all the hosts in this setup.

Storage Mount	Cloned to Target?	Local or Shared (Source)	Local or Shared (Target)	Mount Point	Directory
Fusion Applications Shared	Yes	Shared	Shared	/u01	fa
Fusion Applications Database	Yes	Shared	Shared or Local	/u01	fusiondb
Identity Management MidTier/WebTier	Yes	Shared	Shared or Local	/u01	idm
Identity Management Database(s)	Yes	Shared	Shared or Local	/u01	idmdb

Storage Mount	Cloned to Target?	Local or Shared (Source)	Local or Shared (Target)	Mount Point	Directory
JDK	Yes	Shared	Shared or Local	/u01	jdk6
Repository/Staging (includes Provisioning)	Optional	Shared	Shared	/u01	repository
Response Files	No	Shared	-	/u01	response_files
Temporary Backup	No	Shared	-	/u01	backup

B.2 OS, Provisioning Framework, and JDK Recommendations

These are preparatory steps before the main installation.

Preparation Point	Instructions
Prepare the OS	For each host: ■ Make sure the OS user specified in the Host table above exists, can log in and has access to a graphical interface (X-forwarding works but is extremely slow, VNC is recommended) ■ Make sure the OS Group specified in the Host table above exists and is the primary group of the OS User ■ Make sure there is no `oraInst.loc` file in the `/etc` directory ■ Set the `JAVA_HOME` environment variable to `/u01/jdk6` and the `PATH` to include it `export JAVA_HOME=/u01/jdk6 export PATH=$JAVA_HOME/bin:$PATH`
Before installing the Provisioning Framework	Complete the following tasks: 1. Download the zip files from E-Delivery 2. Create directory `/u01/repository` 3. Extract the installer files from E-Delivery into `/u01/repository` 4. Create directory `/u01/backup` 5. Create directory `/u01/response_files`
Install the JDK	Unzip the `jdk6.zip` file located at `/u01/repository/installers/jdk` to `/u01`, which creates the directory `/u01/jdk6`
Install the Provisioning Framework	Follow the documentation normally: ■ Install the Provisioning Framework to `/u1/faprov` ■ Make sure there is no `/etc/oraInst.loc` file and choose the option to use a local inventory and choose to create the `oraInventory` directory under `/u01` (this will be the same `oraInventory` that will be used later when provisioning) ■ Save response and summary files to `/u01/response_files`

B.3 Identity Management Installation Recommendations

There are separate instructions for installing the database and installing the core components of Identity Management.

B.3.1 Installing Identity Management Databases

For cloning, use special steps to assign abstract host names. This is done in several stages, as described below, and is repeated for every Identity Management tier.

Preparation Point	Instructions
Before installing the Identity Management database(s) on the Identity Management Database Host	Complete the following tasks to create an abstract host name for the Identity Management database host. Because the Provisioning Framework uses the given host name information in subsequent configurations, use a temporary name to simplify for cloning. 1. Edit the `/etc/hosts` file and add an entry for `idmdb.mycompany.com` at the top. 2. Temporarily change the host name to `idmdb.mycompany.com`: `hostname idmdb.mycompany.com` 3. Start a new session (logout and log back in). if using VNC, kill the current VNC session and start a new one so the new host name is properly picked up.
Run the database installer and create the database and listener	Follow the documentation normally, with special attention to the following: ▪ Make sure there is no `/etc/oraInst.loc` file and choose the option to use a local inventory and choose to create the oraInventory directory under `/u01/idmdb` ▪ When providing the host name, always use the abstract host name `idmdb.mycompany.com` ▪ Use a port different from 1521 for the listener (e.g. 1522) as 1521 will be used for the Oracle Fusion Applications Database. Also use a different listener name e.g. `LISTENER_IDMDB`. This will enable target environments to have both Identity Management and Fusion Applications databases share a single host if needed ▪ Make sure everything gets installed under `/u01/idmdb`, including the directories for data files. This will simplify cloning later.
After completing the Identity Management database install	Complete the following tasks: 1. Stop all database processes (instance, listener, database control, etc.). 2. Change the host name back to the original one. 3. Start a new session (logout and log back in). if using VNC, kill the current VNC session and start a new one so the new host name is properly picked up. 4. Start the database instance and listener. 5. Verify that everything is working:
Subsequently	Complete the following tasks: 1. Complete the remaining steps for the database e.g. RCU 2. Stop all database processes (instance, listener, database control, etc.) 3. Take a full cold backup of the idmdb directory using this command: `tar -czvpf <version>idmdb_afterRCU.tgz /u01/idmdb`

B.3.2 Installing Identity Management Web- and Mid Tiers

Preparation Point	Instructions

Before installing the Identity Management web tier and mid tier	Complete the following tasks to create an abstract host name for the Identity Management web and mid tier host. Because the Provisioning Framework uses the given host name information in subsequent configurations, use a temporary name to simplify for cloning.
	
	1. Edit the `/etc/hosts` file and add the following entries pointing to the server itself):
	
	`idmmidtier.mycompany.com` (at the top)
	
	`idmwebtier.mycompany.com`
	
	`oididstore.mycompany.com`
	
	`policystore.mycompany.com`
	
	`idstore.mycompany.com`
	
	`idminternal.mycompany.com`
	
	`sso.mycompany.com`
	
	2. Add an entry for `idmdb.mycompany.com` if it's not already there (pointing to the Identity Management database server).
	
	3. Temporarily change the host name to `idmmidtier.mycompany.com`:
	
	`hostname idmmidtier.mycompany.com`
	
	4. Start a new session (logout and log back in). If using VNC, kill the current VNC session and start a new one so the new host name is properly picked up.

Install/configure Identity Management web and mid tiers	Follow the documentation normally, with special attention to the following: ■ Make sure there is no `/etc/oraInst.loc` file and choose the option to use a local inventory. Choose to create the `oraInventory` directory under `/u01/idm`. ■ Manually create a file called `oraInst.loc` under `/u01` with the following text: `inventory_loc=/u01/idm/oraInventory inst_group=oinstall` ■ When running `config.sh`, post-patching instructions, `idmconfigtool` and `LDAPSync` post-install: Use the alias `idmmidtier.mycompany.com` when referring to the host name for the Identity Management mid-tier components Use `idmdb.mycompany.com` when referring to the host name for Identity Management databases Use `idstore.mycompany.com` when referring to the host name for OVD Use `policystore.mycompany.com` when referring to the host name for OID during policy store configuration Use `oididstore.mycompany.com` when referring to the host name for OID during identity store configuration Use the virtual host `sso.mycompany.com` when referring to the HTTP URLs (web tier) for OAM and OIM ■ When creating domains and instances, make sure they are all located in a directory that is under `/u01/idm` (for example, `/u01/idm/instance`). ■ When choosing port numbers, use ports outside of the ranges that may conflict with the default ports used by Oracle Fusion Applications (in cases where Identity Management and Oracle Fusion Applications are cloned to the same server): For the WebLogic Administration Server, use port 17001 (or something outside the 7XXX range which is used by the Oracle Fusion Applications CommonDomain) For Node Manager, use port 5566 (or something other than 5556 which is the default port for the Oracle Fusion Applications Node Managers) For ODSM use port 17006 (or something outside the 7XXX range which is used by the Oracle Fusion Applications CommonDomain)

Preparation Point	Instructions
After installing/configuring Identity Management web and mid tiers	Complete the following tasks:
	1. Stop all processes for mid tier and web tier components (IDMDomain, nodemanager, Oracle Internet Directory (OID), OHS, ODSM, etc.)
	2. Change the host name back to the original one.
	3. Start a new session (logout and log back in). If using VNC, kill the current VNC session and start a new one so the new host name is properly picked up.
	4. Start the components and verify they are all working by using the browser to connect to:
	WLS Console, Enterprise Manager and OAM Console
	ODSM (and verifying access to the Identity and Policy stores)
	OIM Console (via HTTP Server)
Subsequently	Complete the following tasks:
	1. Stop all processes including the Identity Management databases
	2. Take a full cold backup of the `idmdb` and `idm` directories (the config process has written to the database as well, so both have to be in sync) using this command:
	`tar -czvpf <version>idmdb_starter.tgz /u01/idmdbtar -czvpf <version>idm_starter.tgz /u01idm /u01/jdk6`

B.4 Fusion Applications Installation Recommendations

There are separate instructions for installing the database and installing the core components of Fusion Applications.

B.4.1 Installing Fusion Applications Database

For cloning, use special steps to assign abstract host names. This is done in several stages, as described below, and is repeated for every Fusion Applications tier.

Preparation Point	Instructions
Before installing the Fusion Applications database(s) on the Fusion Applications database Host	Complete the following tasks to create an abstract host name for the Fusion Applications database host. Because the Provisioning Framework uses the given host name information in subsequent configurations, use a temporary name to simplify for cloning.
	1. Edit the `/etc/hosts` file and add an entry for `fusiondb.mycompany.com` at the top.
	2. Temporarily change the host name to `fusiondb.mycompany.com`:
	`hostname fusiondb.mycompany.com`
	3. Start a new session (logout and log back in). if using VNC, kill the current VNC session and start a new one so the new host name is properly picked up.
Run the database installer and create the database and listener	Follow the documentation normally, with special attention to the following:
	■ Make sure there is no `/etc/oraInst.loc` file and choose the option to use a local inventory and choose to create the `oraInventory` directory under `/u01/fusiondb`
	■ When providing the host name, always use the alias `fusiondb.mycompany.com`
	■ You may use regular port 1521 for the listener as conflicts have already been avoided by using 1522 for the Identity Management database listener
	■ Make sure everything gets installed under `/u01/fusiondb`, including the directories for datafiles. This will simplify cloning later

Preparation Point	Instructions
After completing the Fusion Applications database install	Complete the following tasks:
	1. Stop all database processes (instance, listener, database control, etc.).
	2. Change the host name back to the original one.
	3. Start a new session (logout and log back in). if using VNC, kill the current VNC session and start a new one so the new host name is properly picked up.
	4. Start the database instance and listener. Verify that everything is functioning normally by connecting remotely with SQL*Plus.
Subsequently	Complete the following tasks:
	1. Complete the remaining steps for the database e.g. RCU.
	2. Stop all Fusion Applications database processes.
	3. Take a full cold backup of the fusiondb directory using this command: `tar -czvpf <version>fusiondb_afterRCU.tgz /u01/fusiondb`

B.4.2 Provisioning Fusion Applications

For cloning, use special steps to assign abstract host names.

Preparation Point	Instructions
Before Provisioning	On each Oracle Fusion Applications Host:
	1. Edit the `/etc/hosts` file and add an entry for `fusionappsX.mycompany.com` at the top (where X is a number depending on how many Oracle Fusion Applications hosts you plan to have; if you have only one then make that entry `fusionapps.mycompany.com`).

	2. Add entries for the internal and external virtual hosts as well:
	`fainternal.mycompany.com`
	`ofaexternal.mycompany.com`

	3. Temporarily change the host name to `fusionappsX.mycompany.com` (where X is a number depending on how many Oracle Fusion Applications hosts you plan to have; if you have only one then make that `fusionapps.mycompany.com`): `hostname fusionapps.mycompany.com`
	4. Start a new session (logout and log back in). if using VNC, kill the current VNC session and start a new one so the new host name is properly picked up.
Run Provisioning	Follow the documentation normally, with special attention to the following:
	■ Run the Provisioning Framework from the primordial host (the one with the `CommonDomain` Administration Server).
	■ Add the `invPtrLoc` option when invoking `provisioningWizard.sh` and `runProvisioning.sh`. Point it at the `oraInst.loc` file at /u01/faprov.
	■ Install Oracle Fusion Applications to `/u01/fa`.
	■ When providing the hostname, always use the alias `fusionappsX.mycompany.com` (where X is a number depending on how many Oracle Fusion Applications hosts you plan to have; if you have only one then make that `fusionapps.mycompany.com`).
	■ In the Virtual Hosts screen, select IP-based and provide the following for host names: Internal: `fainternal.mycompany.com` External: `faexternal.mycompany.com`
	■ In the LBR screen, leave un-selected (you can add a LBR later by changing the `/etc/hosts` entries for `fainternal` and `faexternal`)

Preparation Point	Instructions
After Provisioning	Complete the following tasks:
	1. Stop all Oracle Fusion Applications processes (you may leave the database running).
	2. Change the host name of each Oracle Fusion Applications host back to the original one.
	3. Start a new session (logout and log back in). if using VNC, kill the current VNC session and start a new one so the new host name is properly picked up.
	4. Start Oracle Fusion Applications again and verify that it is functioning correctly.
	5. Run the Post-Install steps from the documentation.
Subsequently	Complete the following tasks:
	1. Stop all processes including Oracle Fusion Applications, the Oracle Fusion Applications Database, Identity Management, and Identity Management databases
	2. Take a full cold backup of the entire environment using this command:
	`tar -czvpf <version>_idmdb_complete.tgz /u01/idmdbtar -czvpf <version>idm_complete.tgz /u01/idm /u01/jdk6tar -czvpf <version>fusiondb_complete.tgz /u01/fusiondbtar -czvpf <version>fa_complete.tgz /u01/faprov /u01/fa /u01/oraInventory /u01/jdk6`

C

Abstract Hostnames in Detail

The instructions below should be followed during the Discover phase of Cloning process to help obtain the correct abstract hostnames for Fusion Applications, Identity Management and Databases.

Abstract hostnames are the source hostnames (whether physical or virtual) that are present in Oracle Fusion Applications directory paths, configuration files and metadata, and which will be maintained across clones. They must be added to the `/etc/hosts` file in each one of the target environment nodes so that those references resolve to the destination environment itself, not to the source.

For example: When MyCompany provisioned their source environment, the hostname used in the Provisioning wizard was *fahost.mycompany.com*. Thus, their environment will contain *fahost.mycompany.com* in many places, such as directory paths, configuration files, and metadata. When they clone this environment, the clones will also contain *fahost.mycompany.com* in the same places, as the cloning script does not edit these settings. In the destination environment, *fahost.mycompany.com* is an abstract hostname and must be added to the `/etc/hosts` file in each one of the target environment nodes.

Note: abstract hostnames are not the only hostnames that must be added to `/etc/hosts` on the target environment. This requirement also exists for internal HTTP endpoints and VIPs used for failover, which are also discussed in this document.

This chapter includes the following aspects of how to discover abstract hostnames for cloning:

- Section C.1.1, "Discovering Fusion Applications Nodes"
- Section C.1.2, "Discovering Identity Management Nodes"
- Section C.1.3, "Discovering Database Nodes"

C.1 How to Discover Abstract Hostnames for Cloning

There are separate processes for the Fusion Applications components and the Identity Management components.

C.1.1 Discovering Fusion Applications Nodes

In Fusion Applications nodes, the easiest way to determine the abstract hostname is to check the `provisioning.rsp` file in the source environment. The hostnames used in the Provisioning wizard (Domain Topology Configuration and Web Tier Configuration screens) are saved to the response file in the `#Domain Topology` and `#Webtier Configuration` sections. These are the abstract names used in the Discovery Workbook for their respective nodes.

During the Clone phase, the cloning script uses Fusion Applications abstract hostnames to determine the path to each node's nodemanager startup script, when run in the `famt` mode.

Table C–1   Fusion Applications Example: Simple

Fusion Applications installed on host with hostname	fahost.mycompany.com
Hostname provided on the Provisioning wizard (Domain Topology)	fahost.mycompany.com
Hostname provided on the Provisioning wizard (Web Tier)	fahost.mycompany.com
Abstract hostname for node (must be added to `/etc/hosts`)	fahost.mycompany.com

C.1.1.1 Special Case: Web Tier and Fusion Applications on Same Node but Different Hostname Were Provided

If the web tier is installed on the same node as the rest of Fusion Applications, the entry for the web tier hostname in the provisioning response file (`WEBTIER_HOST`) is normally the same as for the Fusion Applications node. If for any reason it's different (yet they are in the same node), use the Fusion Applications hostname as the abstract name for that node and add the web tier hostname to the "Additional `/etc/hosts` File Entries" table in the "Virtual Hosts" tab of the Discovery Workbook. Both hostnames will have to be added to the `/etc/hosts` file before running the cloning scripts.

Table C–2   Single Node Install with Two Hostnames Special Case

Hostname as returned by the `hostname` command	fahost.mycompany.com
Hostname provided on the Provisioning wizard (Domain Topology)	fahost.mycompany.com
Hostname provided on the Provisioning wizard (Web Tier)	fawebtiervirtual.mycompany.com
Abstract hostname for node (must be added to `/etc/hosts`)	fahost.mycompany.com
Additional `/etc/hosts` file entries (all other hostnames	fawebtiervirtual.mycompany.com

C.1.1.2 Special case: Hostnames Provided during Provisioning Differ from Actual Hostnames

Normally, the hostname in the provisioning response file will match the physical hostname, as returned by the UNIX `"hostname"` command. If a source environment setup was done so that the hostname as returned by the UNIX `hostname` command and the hostname provided to the Provisioning Wizard were not the same, then there are two abstract names and both must be added to the `/etc/hosts` file.

In this case, the hostname provided to the Provisioning Wizard (and found in the `provisioning.rsp` file) is used in the Discovery Workbook as the Abstract Hostname of the respective node. The hostname as returned by the Unix `hostname` command should be added to the "Additional /etc/hosts File Entries" table in the "Virtual Hosts" tab.

Table C–3   Single Node Install with Web Tier All Using Different Names

Hostname as returned by the `hostname` command	fahost.mycompany.com
Hostname provided on the Provisioning wizard (Domain Topology)	favirtual.mycompany.com
Hostname provided on the Provisioning wizard (Web Tier)	fawebtiervirtual.mycompany.com
Abstract hostname for node (must be added to `/etc/hosts`)	favirtual.mycompany.com
Additional `/etc/hosts` file entries (all other hostnames	fahost.mycompany.com fawebtiervirtual.mycompany.com

C.1.2 Discovering Identity Management Nodes

For Identity Management (as installed in the *Oracle Identity Management Enterprise Deployment Guide*), abstract hostnames may vary depending on the topology and the hostnames/VIPs provided during the installation and configuration. As a rule of thumb, abstract hostnames for Identity Management nodes match the hostnames as returned by the Unix `hostname` command.

Identity Management abstract hostnames are not used by the cloning scripts; however they are used by the clone tool during validation, when "idm" mode is used. The tool will use the abstract hostname of the OAM and OVD nodes to validate the environment's OAM settings. Thus it is important to match the instructions below.

In any case, all Identity Management hostnames, as returned by the UNIX `hostname` command, and all virtual hostnames used during the Identity management installation must be added to the `/etc/hosts` file. Therefore, if any are not listed in the Abstract Hostname column of the Topology tab of the Discovery Workbook, ensure that they are added to the "Additional `/etc/hosts` file entries" table in the "Virtual Hosts" tab.

> **Note:**   if the servers where IDM and Fusion Applications were installed have more than one network card, and the additional network cards are registered in DNS with a different name, these additional names may have ended up in some of the configuration files, even if the 'hostname' command returned the correct one.
>
> More specifically, this has happened with the EMAGENT component and the unexpected hostname turned up on *targets.xml*.
>
> These additional hostnames must also be added to /etc/hosts on the target.

C.1.2.1 Abstract Hostname of the OAM Node

Ensure the abstract hostname of the node containing OAM is the same as displayed on the OAM Console (System Configuration tab/ Common Configuration/ Server Instances area) in the field labeled "Host". If they don't match, running the cloning tool in **validate idm** mode may fail. However, it does not affect the actual clone creation.

C.1.2.2 Abstract hostname of the OVD Node

Ensure the abstract hostname of the node containing OVD is the same as displayed on the OAM Console (System Configuration tab/ Common Configuration/ Data

Sources/User Identity Stores/ OIMIDStore page), in the field labeled "Location". If they don't match, running the cloning tool in **validate idm** mode may fail. However, it does not affect the actual clone creation.

C.1.2.2.1 Special case: OVD and OAM share the same node but have different abstract hostnames

In this case, use the OAM abstract hostname (Section C.1.2.1) as the Abstract Hostname of the respective node in the Discovery Workbook. Add the OVD abstract hostname (Section C.1.2.2) to the "Additional /etc/hosts file entries" table in the "Virtual Hosts" tab of the Discovery Workbook. In this case, running the cloning tool in **validate idm** mode may fail since the OVD abstract hostname does not exactly match the entry in the OAM configuration.However, it does not affect the actual clone creation.

C.1.2.2.2 Special case: Other components share the same node with different hostnames used during install

If any other components share the same node with OVD or OAM, follow the notes above to ensure the abstract hostnames match the appropriate hostname for OAM or OVD. For all other cases, simply use the hostnames as returned by the UNIX hostname command.

Table C–4   Single Node Install Identity Management Install

Hostname as returned by the `hostname` command	idmhost.mycompany.com
Hostnames provided during install and configuration	idstore.mycompany.com
	policystore.mycompany.com
	ldaphost1.mycompany.com
	webhost1.mycompany.com
	idmhost.mycompany.com
	idmhost
Hostname as displayed on the OAM Console	idmhost.mycompany.com
Abstract hostname for node	idmhost.mycompany.com
(must be added to `/etc/hosts`)	(matches the hostname displayed on the OAM Console)
Additional `/etc/hosts` file entries (all other hostnames	idstore.mycompany.com
	policystore.mycompany.com
	ldaphost1.mycompany.com
	webhost1.mycompany.com
	idmhost

Table C–5   Two-Node Identity Management Install

	Hostname as returned by the `hostname` command	idmhost1.mycompany.com
Node 1 (OVD and OID)	Hostnames provided during install and configuration	idstore.mycompany.com
		policystore.mycompany.com
		ldaphost.mycompany.com
		idmhost1.mycompany.com
		idmhost1
	Hostname as displayed on the OAM Console	idstore.mycompany.com
	Abstract hostname for node (must be added to `/etc/hosts`)	idstore.mycompany.com (matches the hostname displayed on the OAM Console
	Additional `/etc/hosts` file entries (all other hostnames	idmhost1.mycompany.com
		policystore.mycompany.com
		ldaphost.mycompany.com
		webhost.mycompany.com
		idmhost1
	Hostname as returned by the `hostname` command	idmhost2.mycompany.com
Node 2 (WLS and OHS)	Hostnames provided during install and configuration	webhost.mycompany.com
		idmhost2.mycompany.com
		idmhost2
	Hostname as displayed on the OAM Console	idmhost2.mycompany.com
	Abstract hostname for node (must be added to `/etc/hosts`)	idmhost2.mycompany.com (matches the hostname displayed on the OAM Console)
	Additional `/etc/hosts` file entries (all other hostnames	webhost.mycompany.com
		idmhost2

C.1.3 Discovering Database Nodes

Database abstract hostnames are the ones used during Identity Management and Fusion Applications installation. Database abstract hostnames are the only ones that will be changed by the clone tool.

When the cloning scripts are run, the tool uses the entries on the tables (located in the Databases tab of the Discovery Workbook) to rewire all database connections. Database information (hostname, port, SID, Service Name) from the source environment is replaced with the destination entries provided, so it's important that the source information you enter on the Discovery Workbook match what's in your source environment.

C.1.3.1 OID Database

The OID Database abstract hostname (or multiples, in the case of RAC databases) is the one used during OID installation. On an existing environment it can be identified by inspecting the tnsnames.ora file located in $OID_INSTANCE/config.

C.1.3.2 Identity Management (IDM) Database

The Identity Management database abstract hostname (or multiples, in the case of RAC databases) is the one used during installation of Oracle Identity and Access Management (OIM/OAM), as well as Fusion Applications provisioning.

It can be obtained by inspecting the provisioning.rsp file, or checking the WLS Console for the IDMDomain data sources (all data sources) and Fusion Applications domains data sources (OWSM data sources only).

C.1.3.3 Fusion Applications (FA) Database

The Fusion Applications Database abstract hostname (or multiples, in the case of RAC databases) is the one used during Fusion Applications provisioning and can be obtained by inspecting the provisioning.rsp file.

It can be identified on an existing environment by inspecting the FA domain data sources (with the exception of the OWSM and BI data sources, which instead point at the IDM Database and BI server respectively).

These abstract hostnames should be added to the respective node's Abstract Hostname column in the Topology tab as well as the Source DB Instances row for each database type (OID, IDM, FA) in the Databases tab of the Discovery Workbook.

D

Change JPS Root Name as Needed

When performing production-to-test data movement, it may be necessary to change JPS Root, if you find during Discovery that the `IDM_JPSROOT` name or the `FA_JPSROOT` name are not identical between the source and target systems.

To fix this, use the sample `jpsroot.properties` file that was included in the `$WORKDIR` in the P2T package. On the target system, you will enter target values for all the properties, as annotated in the Discovery Workbook. Enter the source values for `IDM_JPSROOT` and/or `FA_JPSROOT`, depending on what was inconsistent. You then run the appropriate scripts below (for FA and/or IDM, as needed). This will create a matching jpsroot name on the target and will re-associate the target Security store to newly created jpsroot.

Access and Edit JPS Root Name

To change JPS root:

1. Locate the `jpsroot.properties` file under `$WORK_DIR/utilhome/bin` and enter the target values. See the Sample File, below, for the format.
   The values were annotated in the Discovery Workbook, in the P2T tabs.

2. For `IDM_JPSROOT` and `FA_JPSROOT`, enter values that match the source system.

3. **If there was a discrepancy in the IDM_JPSROOT name:**
   Run the following script on the target system to apply the changes and re-associate the Security store with the new value:

   ```
   cd $WORK_DIR/utilhome/bin
   ./createJpsRoot.sh jpsroot.properties IDM
   ```

   Restart IDM stack.

4. **If there was a discrepancy in the FA_JPSROOT name:**
   Run the following script on the target system to apply the changes and re-associate the Security store with the new value:

   ```
   cd $WORK_DIR/utilhome/bin
   ./createJpsRoot.sh jpsroot.properties FA
   ```

   Restart Fusion Applications stack.

Sample File

```
JAVA_HOME=/
```

```
# Working Directory
WORKDIR=

# IDM WLS Home
IDM_BASE=

### IDM Domain directory
IDM_DOMAIN_DIR=

# Enter the absolute path of the APPLTOP directory.
APPLTOP=

# destination hostname
OID_HOSTNAME=
# IDM Home
OID_HOME=

#destination port
OID_PORT=

# IDM WLS ADMIN Server User Name
IDM_WLS_USER=

# IDM WLS Admin User Password
IDM_WLS_PASSWORD=

# IDM T3 URL
IDM_T3URL=

# password for cn=orcladmin
ADMIN_PASSWORD=

# Enter DC base - Example: dc=mycompany.com
DC=dc=mycompany,dc=com

# Enter IDM new jpsroot name
IDM_JPSROOT=

# Enter Fusion Applications new jpsroot name
FA_JPSROOT=

# FA WLS ADMIN Server User Name
FA_WLS_USER=

# FA WLS Admin User Password
FA_WLS_PASSWORD=

# Enter Fusion Apps Domains Directory. Specify 'NONE' if this domain is not in
use.
COMMON_DOMAIN_
DIR=/u01/oracle/fa/config/domains/fahost01.mycompany.com/CommonDomain
CRM_DOMAIN_DIR=
HCM_DOMAIN_DIR=
SCM_DOMAIN_DIR=
FIN_DOMAIN_DIR=
BI_DOMAIN_DIR=
PRC_DOMAIN_DIR=
PRJ_DOMAIN_DIR=
IC_DOMAIN_DIR=
```

```
# Enter the URL of all the domains of Oracle WebLogic Server domain in the form of
t3://hostname:port. Specify 'NONE' if this domain is not in use.
COMMONDOMAIN_T3_URL=t3://
CRMDOMAIN_T3_URL=t3://
HCMDOMAIN_T3_URL=t3://
SCMDOMAIN_T3_URL=t3://
PRJDOMAIN_T3_URL=t3://
FINDOMAIN_T3_URL=t3://
PRCDOMAIN_T3_URL=t3://
ICDOMAIN_T3_URL=t3://
BIDOMAIN_T3_URL=t3://

# Enter all FA domains URL except CommonDomain seperated by comma - no comma in
the end - no blank space anywhere
FA_T3_URLS=
```

E

Install the BI Administration Tool

This appendix gives a quick version of the Business Intelligence Administration tool installation steps.

Installing the Tool

The Repository, or RPD, is the file that contains the metadata for the BI Server in Fusion Applications. This includes database connections, tables, joins, and the structures by which these are presented to the report writer. In order to read or make changes to the RPD file, the BI Administration Tool must be installed.

System Prerequisites

- **64-bit Windows OS** is required to run the BI Administration tool.

- **The BI Admin tool version** must match the Fusion Applications version you have installed.

Locate and Install the Software

There are two different ways to obtain the correct version of the tool:

1. From the Oracle Fusion Applications BIEE Analytics Home, download the 64-bit Windows Oracle BI Client Installer.

2. OR: Copy it from the BI Server from this path:
   `FA_HOME/fusionapps/bi/clients/biserver/biee_client_install_x64.exe`
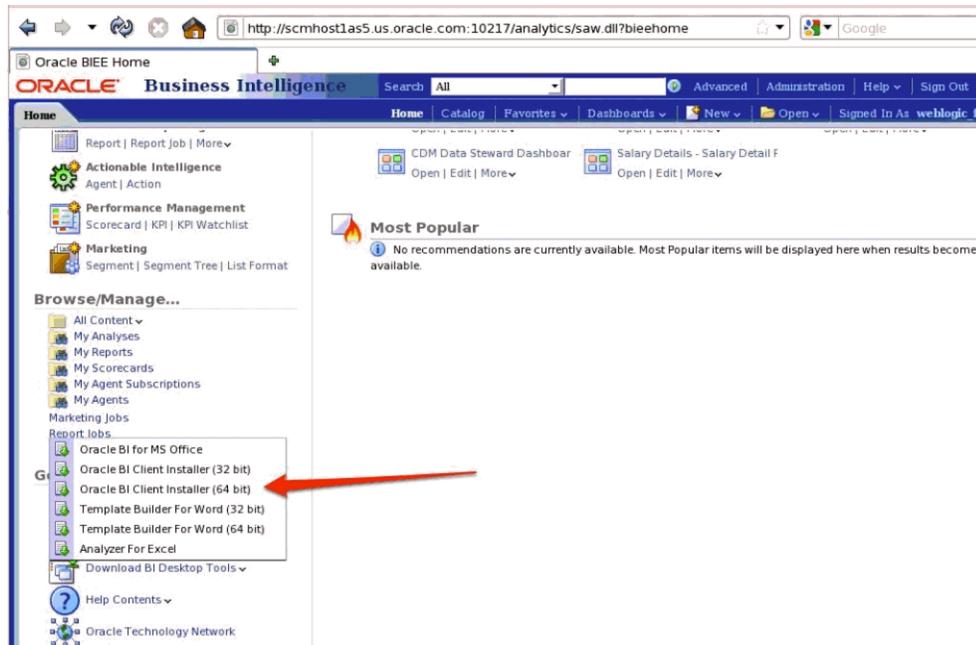
Once you have downloaded, run the biee_client_install_x64.exeexecutable to install the tool.

Set up ODBC Connection to Fusion Applications

To set up the ODBC connection from the Windows machine, you must locate the correct port on the Fusion Applications BI Server. Typically it is **10206**, but if ports were changed as part of the FA install, perform the following steps:

1. Log in to the BI Server's **Enterprise Manager Console.**

2. Expand **Business Intelligence** within the **Farm_BIDomain** and select **coreapplication**. Select **Overview** and the sub tab **Processes**.

3. Expand the BI Servers, and make a note of the port number.

Armed with the correct port, you are ready to create a new ODBC connection.

1. Select Oracle BI Server Driver on the Windows machine.

2. Enter the Server name for the host that contains the BI Domain. Click Next. (If necessary, add an entry to the `windows/system32/drivers/etc` file, so that the windows machine can resolve the host name.)

3. On the resulting screen, enter the BI Server port (either 10205 or derived above) and the user name/password with which you plan to connect. Click Next.

Connect to RPD in BI Admin Tool

The user could now connect to the BI Server in On-Line mode and could update the live RPD. CAUTION! Depending on network speed and the size of the RPD, this can be both slow and dangerous. Updating a live RPD while Fusion Applications users and reports are working with it is not advised.

A better solution is to copy the RPD file from the BI Server to a local or shared drive the Windows client can access, updating the RPD in off-line mode, and then publishing it through the BI Server's Enterprise Manager console.

The RPD files are stored in the following path on the BI Server: `BIInstance/bifoundation/OracleBIServerComponent/coreapplication_obis1/repository`.

To find the correct version:

1. Log in to Enterprise Manager, go to the **Business Intelligence / coreapplication** section, and select the **Deployment / Repository** tabs.
   This will identify the current version of the RPD.

2. Copy that version locally or to the shared directory, to ensure you are working with the most current RPD.