

# Oracle9iAS Personalization

Administrator's Guide

Release 9.0.1

July 2001

Part No. A87539-02

**ORACLE®**

Copyright © 2001, Oracle Corporation. All rights reserved.

The Programs (which include both the software and documentation) contain proprietary information of Oracle Corporation; they are provided under a license agreement containing restrictions on use and disclosure and are also protected by copyright, patent, and other intellectual and industrial property laws. Reverse engineering, disassembly, or decompilation of the Programs is prohibited. [

The information contained in this document is subject to change without notice. If you find any problems in the documentation, please report them to us in writing. Oracle Corporation does not warrant that this document is error free. Except as may be expressly permitted in your license agreement for these Programs, no part of these Programs may be reproduced or transmitted in any form or by any means, electronic or mechanical, for any purpose, without the express written permission of Oracle Corporation.

If the Programs are delivered to the U.S. Government or anyone licensing or using the programs on behalf of the U.S. Government, the following notice is applicable:

**Restricted Rights Notice** Programs delivered subject to the DOD FAR Supplement are "commercial computer software" and use, duplication, and disclosure of the Programs, including documentation, shall be subject to the licensing restrictions set forth in the applicable Oracle license agreement. Otherwise, Programs delivered subject to the Federal Acquisition Regulations are "restricted computer software" and use, duplication, and disclosure of the Programs shall be subject to the restrictions in FAR 52.227-19, Commercial Computer Software - Restricted Rights (June, 1987). Oracle Corporation, 500 Oracle Parkway, Redwood City, CA 94065.

The Programs are not intended for use in any nuclear, aviation, mass transit, medical, or other inherently dangerous applications. It shall be the licensee's responsibility to take all appropriate fail-safe, backup, redundancy, and other measures to ensure the safe use of such applications if the Programs are used for such purposes, and Oracle Corporation disclaims liability for any damages caused by such use of the Programs.

Oracle is a registered trademark, and Oracle9i is a trademark or registered trademark of Oracle Corporation. Other names may be trademarks of their respective owners.

---

---

# Contents

<b>Send Us Your Comments .....</b>	<b>v</b>
<b>Preface.....</b>	<b>vii</b>
<b>1 Overview</b>	
<b>Installation Overview .....</b>	<b>1-1</b>
<b>2 Requirements</b>	
<b>User Prerequisites.....</b>	<b>2-1</b>
<b>Database and Application Server Requirements .....</b>	<b>2-1</b>
<b>Performance Recommendation.....</b>	<b>2-2</b>
<b>Workflow Requirements .....</b>	<b>2-2</b>
<b>Java 1.2.2 Requirements.....</b>	<b>2-2</b>
<b>Environment Variables .....</b>	<b>2-2</b>
<b>3 Installing and Configuring Oracle9i and Oracle9iAS</b>	
<b>Installing Oracle9i .....</b>	<b>3-1</b>
<b>Installing Oracle9i Application Server.....</b>	<b>3-4</b>
<b>4 Installing Oracle Workflow</b>	
<b>Installing Oracle Workflow 2.6.....</b>	<b>4-1</b>
<b>Configuring Email Notifications (Optional).....</b>	<b>4-4</b>

<b>5</b>	<b>Installing OP</b>	
	Installing OP Components.....	5-1
	Manual Steps .....	5-6
	Test the OP Installation .....	5-8
	Deinstalling OP .....	5-9
<b>6</b>	<b>OP Security</b>	
	User Access.....	6-1
<b>7</b>	<b>Configuring Oracle9iAS Personalization</b>	
	Changing OP Configuration Parameters .....	7-1
	Data Synchronization.....	7-7
<b>8</b>	<b>OP Schemas</b>	
	Mining Table Repository.....	8-2
	Recommendation Engine .....	8-9
	Mining Object Repository .....	8-10
<b>9</b>	<b>Initial Data Collection</b>	
	Creating and Removing Seed Data .....	9-1
	Preparing to Build Models with Seed Data.....	9-2
<b>A</b>	<b>Installation and Configuration Considerations</b>	
	RE Installation and Configuration .....	A-1
	MTR Installation and Configuration.....	A-4

---

---

# Send Us Your Comments

**Oracle9iAS Personalization Administrator's Guide, Release 9.0.1**

**Part No. A87539-02**

Oracle Corporation welcomes your comments and suggestions on the quality and usefulness of this document. Your input is an important part of the information used for revision.

- Did you find any errors?
- Is the information clearly presented?
- Do you need more information? If so, where?
- Are the examples correct? Do you need more examples?
- What features did you like most?

If you find any errors or have any other suggestions for improvement, please indicate the document title and part number, and the chapter, section, and page number (if available). You can send comments to us in the following ways:

- [DARWINDOC@us.oracle.com](mailto:DARWINDOC@us.oracle.com)
- FAX: 781-684-7738 Attn: Oracle9iAS Personalization Documentation
- Postal service:  
Oracle Corporation  
Oracle9iAS Personalization Documentation  
200 Fifth Avenue  
Waltham, Massachusetts 02451  
U.S.A.

If you would like a reply, please give your name, address, telephone number, and (optionally) electronic mail address.

If you have problems with the software, please contact your local Oracle Support Services.



---

---

# Preface

This manual describes how to install the Oracle9iAS Personalization (OP) software and how to perform other administrative functions. See the README.htm file for platform-specific installation information and any late-breaking changes.

## Intended Audience

This manual is intended for anyone planning to install and run Oracle9iAS Personalization— either a database administrator or a system administrator.

## Structure

This manual contains the following chapters:

Chapter 1	Provides an overview of the installation process for OP.
Chapter 2	Describes generic requirements and prerequisites.
Chapter 3	Gives an overview of installing Oracle9i and Oracle9iAS.
Chapter 4	Describes how to install Oracle Workflow.
Chapter 5	Describes how to install OP and how to initialize the database tables used by OP.
Chapter 6	Describes security measures.
Chapter 7	Describes OP configuration.
Chapter 8	Describes the OP schemas.
Chapter 9	Describes initial data collection.
Appendix A	Describes additional installation and configurations (tuning).

---

## Where to Find More Information

The documentation set for Oracle9iAS Personalization at the current release consists of the following documents:

- README.htm, on the Oracle9iAS Personalization CD; this file contains platform-specific installation instructions.
- *Oracle9iAS Personalization Release Notes*, Release 9.0.1.
- *Oracle9iAS Personalization Administrator's Guide*, Release 9.0.1 (includes installation instructions that are the same across all platforms) (this document).
- *Getting Started with Oracle9iAS Personalization*, Release 9.0.1.
- *Oracle9iAS Personalization Recommendation Engine API Programmer's Guide*, Release 9.0.1. A programmer's manual for programming the recommendation engines in real time.
- *Oracle9iAS Personalization Recommendation Engine Batch API Programmer's Guide*, Release 9.0.1. A programmer's manual for obtaining bulk recommendations.

### Related Manuals

For more information about the database underlying OP, see:

- *Oracle9i Administrator's Guide*
- *Oracle9i Application Server Installation Guide* (the appropriate version for your operating system).

### Requirements

OP documentation is distributed on the same CD that OP is distributed on. Documentation is provided in PDF and HTML formats.

After OP is installed, the OP documentation can be read by opening the following URL using either Netscape or Internet Explorer:

`http://server/opDoc/op.901/index.htm`

where *server* is name of the system where OP is installed.

You can read or print the documentation directly from the CD or from your browser.



---

To view the PDF files, you will need

- Adobe Acrobat Reader 3.0 or later, which you can download from [www.adobe.com](http://www.adobe.com).

To view the HTML files, you will need

- Netscape 4.x or later, or
- Internet Explorer 4.x or later

## Documentation Accessibility

Oracle's goal is to make our products, services, and supporting documentation accessible to the disabled community with good usability. To that end, our documentation includes features that make information available to users of assistive technology. This documentation is available in HTML format, and contains markup to facilitate access by the disabled community. Standards will continue to evolve over time, and Oracle is actively engaged with other market-leading technology vendors to address technical obstacles so that our documentation can be accessible to all of our customers. For additional information, visit the Oracle Accessibility Program web site at <http://www.oracle.com/accessibility/>.

### Accessibility of Code Examples in Documentation

JAWS, a Windows screen reader, may not always correctly read the code examples in this document. The conventions for writing code require that closing braces should appear on an otherwise empty line; however, JAWS may not always read a line of text that consists solely of a bracket or brace.

## Conventions

In this manual, Windows refers to the Windows 95, Windows 98, and the Windows NT operating systems.

The SQL interface to Oracle9i is referred to as SQL. This interface is the Oracle9i implementation of the SQL standard ANSI X3.135-1992, ISO 9075:1992, commonly referred to as the ANSI/ISO SQL standard or SQL92. In examples, an implied carriage return occurs at the end of each line, unless otherwise noted. You must press the Return key at the end of a line of input.

---

The table below shows the conventions followed in this manual and their meanings:

Convention	Meaning
. . . . . .	Vertical ellipsis points in an example mean that information not directly related to the example has been omitted.
...	Horizontal ellipsis points in statements or commands mean that parts of the statement or command not directly related to the example have been omitted
<b>boldface text</b>	Boldface type in text indicates a term defined in the text, the glossary, or in both locations.
< >	Angle brackets enclose user-supplied names.
[ ]	Brackets enclose optional clauses from which you can choose one or none.

---

# Overview

Oracle9iAS Personalization (OP) permits Web applications to collect and store customer data in an Oracle database, build models, and make recommendations on the basis of the models.

This manual provides the generic instructions for installing Oracle9iAS Personalization using Oracle Universal Installer (OUI). Platform-specific information is in the README.htm file for each platform.

This manual is organized as follows:

- Overview (Chapter 1)
- Requirements (Chapter 2)
- Installing the Oracle9i database and Oracle9i Application Server (Chapter 3)
- Installing Oracle Workflow (Chapter 4)
- Installing OP (Chapter 5)
- Security (Chapter 6)
- Configuring OP (Chapter 7)
- OP Schemas (Chapter 8)
- Initial data collection (Chapter 9)

## Installation Overview

These are the major steps required to install OP; perform them in this order:

1. Make sure that your operating system satisfies the requirements described in the README.htm file on the OP CD. Verify that your Web site will support OP.
2. Install and configure the Oracle9i database and Oracle9i Application Server.

3. Install the Workflow component of Oracle Management and Integration.
4. Install OP and its components, and initialize required tables.

---

# Requirements

This chapter describes the generic OP requirements and prerequisites.

Additional requirements and recommendations may be specified during installation of each component.

---

**Note:** Platform-specific system requirements are described in the README.htm file for each platform.

---

## User Prerequisites

The person installing OP is expected to be familiar with:

- Oracle database administration; some database operations require DBA privileges
- Web server installation and configuration

and is expected have the following permissions:

- an operating system user account within the DBA group
- root access

## Database and Application Server Requirements

OP requires both Oracle9i and Oracle9i Application Server. The default configuration is to install Oracle9i and Oracle9iAS on different systems. When OP is installed, the components that require Oracle9iAS are installed on the system where Oracle9iAS is installed; all other OP components, including documentation, are installed on the system where Oracle9i is installed.

# Performance Recommendation

If you plan to use OP in an application that has a large amount of data relative to the available temporary space on your system, you will not get satisfactory performance without partitioning. Suppose that

- P is the average profile size
- C is the number of customers
- T is the available temporary space on your system in bytes

If  $54 * P^2 * C > 6 * T$ , then we recommend that you purchase Oracle9i partitioning.

# Workflow Requirements

Oracle Workflow 2.6.1 (Oracle9i Management and Integration 9.0.1) is required.

# Java 1.2.2 Requirements

To be able to run Java 1.2.2 properly on your system, install the patches recommended for your platform.

# Environment Variables

Ensure that the environment variables in the table below are set appropriately in the environment of the Oracle database administrator login account. The table shows the name of each environment variable, its function and syntax, its definition, and an example showing what it looks like with real values substituted where appropriate. Oracle9iAS requires a somewhat different list of variables; for details, see "Before Installing Oracle9iAS" in Chapter 3.

Variable	Detail	Definition
ORACLE_HOME	Function	Specifies the directory containing the Oracle software.
	Syntax	<i>directory_path</i>
	Example	\$ORACLE_BASE/product/9.0.0
ORACLE_SID	Function	Specifies the Oracle system identifier.
	Syntax	A string of numbers and characters that begins with a letter. Oracle Corporation recommends a maximum of eight characters. For more information, see the <i>Oracle 9i Installation Guide for Sun SPARC Solaris</i> Release 9.0.0 Beta.

Variable	Detail	Definition
ORACLE_BASE	Example	SAL1
	Function	Specifies the base of the Oracle directory structure for Optimal Flexible Architecture (OFA) compliant databases.
	Syntax	<i>directory_path</i>
PATH	Example	<i>/u01/app/oracle</i>
	Function	Used by the shell to locate executable programs; must include \$ORACLE_HOME/bin.
	Syntax	Colon-separated list of directories: <i>directory1:directory2:directory3</i>
	Example	<i>/bin:/usr/bin:/usr/local/bin:/usr/bin/X11:\$ORACLE_HOME/bin:\$HOME/bin:.</i> <b>Note:</b> The period adds the current working directory to the search path.
LD_LIBRARY_PATH	Function	Used by the shared library loader ( <i>ld.so.1</i> ) at runtime to find shared object libraries. See the <i>ld.so.1</i> man pages for details.
	Syntax	Colon-separated list of directories: <i>directory1:directory2:directory3</i>
	Example	<i>/usr/dt/lib:\$ORACLE_HOME/lib</i>





---

# Installing and Configuring Oracle9i and Oracle9iAS

The instructions in this chapter cover only the highlights of installing and configuring Oracle9i and Oracle9iAS for use with OP. Oracle9i and Oracle9iAS are installed on different systems.

## Installing Oracle9i

This section summarizes installing Oracle9i. For full information, see *Oracle9i Administrator's Guide*.

---

**Note:** Oracle9i and Oracle9iAS are installed on different systems.

---

Use Oracle Universal Installer (OUI) to install Oracle9i Database 9.0.1.0.0:

1. Start OUI as described in *Oracle9i Administrator's Guide*.
2. Install the Enterprise Edition of Oracle9i. With the Enterprise Edition, installation includes all the required components plus creating and configuring a starter database.

See "Required Database Components", below, for a list of the database components required for OP.

OP runs on EE with and without partitioning. Partitioning is required if you want to deal with large datasets. See "Performance Recommendation" in Chapter 2 for more information.

3. When prompted, select General Purpose, Global Database Name, and Database File Location. Select appropriate database character set.

4. Click **Install**.
5. When prompted, run script `$ORACLE_HOME/root.sh` while logged in as root.

### Required Database Components

The following database components are required for OP:

- Enterprise Edition Options 9.0.1.0.0
  - Oracle Partitioning 9.0.1.0.0 (optional; see "Performance Recommendation" in Chapter 2.)
- Oracle Net Services 9.0.1.0.0
  - Oracle Net Listener 9.0.1.0.0
  - Oracle Net Protocol Support 9.0.1.0.0
- Oracle Management Server 9.0.1.0.0
  - Oracle HTTP Server Extensions 9.0.1.0.0
  - Oracle Mod PL/SQL Gateway 3.0.9.0.7
  - Apache Module for Oracle Servlet Engine 9.0.1.0.0

### Configuring the Database

1. When installation is finished, a message announces that database creation is completed, and the Password Management window appears.
2. **Password Management:** This Password Management window displays the default passwords for SYS and SYSTEM. Change these passwords here.
  - SYS: `change_on_install` (this is the default; change this password here)
  - SYSTEM: `manager` (this is the default; change this password here)
3. Next, edit the `init.ora` file for your Oracle database instance to set additional startup parameters. Typically this file will be named `init<$ORACLE_SID>.ora` and can be found in the `$ORACLE_HOME/admin/<db_name>/pfile` directory. Go to the end of the `init<$ORACLE_SID>.ora` file and add:

```
aq_tm_processes = 1
utl_file_dir = /tmp/
```

`job_queue_processes = 2` (This parameter may already be present; if it is, leave its value set as it is. If it is not there, add it and set it to at least 2.)

The meaning of these parameters is as follows:

- `aq_tm_processes` enables time monitoring of queue messages.
  - `utl_file_dir` lets you specify a directory that Oracle should use for PL/SQL file I/O.
  - `job_queue_processes` is the maximum number of jobs that can be created for execution of jobs; this value should be at least 2.
4. Follow these steps to create a new server parameter file using the updated `init<$ORACLE_SID>.ora` file. (See *Oracle 9i Administrator's Guide* for more information.)

- a. Start SQL\*Plus using the following command:

```
$ sqlplus /nolog
```

- b. Connect to database as SYS user:

```
SQL> connect / as sysdba
SQL> CREATE SPFILE = 'ORACLE_HOME/dbs/newspfile<$ORACLE_SID>.ora'
FROM PFILE='ORACLE_HOME/dbs/init<$ORACLE_SID>.ora';
```

Use the value of `ORACLE_HOME` in the above SQL statement.

- c. Shut down the database.

```
SQL> shutdown
```

- d. Copy the new server parameter file created in step 4b to the existing one:

```
ORACLE_HOME/dbs/spfile<$ORACLE_SID>.ora
```

5. Start up the database so that it can read the new parameters.

```
$ sqlplus /nolog
SQL> connect / as sysdba
SQL> startup
```

Once database starts, exit sqlplus:

```
SQL> exit
```

6. You can test your connection to the database with the following command:

```
$ sqlplus system/<password>@<net_service_name>
```

Your `tnsnames.ora` file contains `net_service_name` (TNSname).

7. Once the database connection is established, you can list the parameters with the following command:

```
SQL> show parameters
```

and verify that `aq_tm_processes`, `utl_file_dir`, and `job_queue_processes` have the values given in step 3.

## Installing Oracle9i Application Server

This section summarizes installing Oracle9iAS for use with OP. For complete installation information, see the appropriate version of *Oracle9i Application Server Installation Guide* for your operating system.

---

---

**Note:** Oracle9i and Oracle9iAS are installed on different systems.

---

---

## Before Installing Oracle9iAS

Before you install Oracle9iAS, perform these steps:

1. Set the required environment variables: `ORACLE_HOME`, `ORACLE_TERM`, `DISPLAY`, `TMP`, and `TNS_ADMIN`.
2. Create required UNIX account and groups.
3. Preinstall any component-specific components on the middle tier and preconfigure an origin database.
4. Oracle9iAS requires an active database connection. OUI uses this connection to add database objects to the origin database. The origin database can be the database that holds OP database objects, but it is not mandatory that the origin database be the one that holds OP. The origin database is typically located on a database server tier. It can also be located in the same tier where Oracle9iAS is installed.

## Installing Oracle9iAS

Follow these steps to install Oracle9iAS:

1. Start OUI as described in *Oracle9i Application Server Installation Guide*. (Don't start OUI in the CD directory because you will have to swap CDs.)
2. **File Locations:** Specify the path of source file (CD or staging area) and destination (Oracle Home).
3. **UNIX Group Name:** Provide the UUNIX group name (for example, dba).
4. **Installation Types:** Select the Enterprise Edition option (4.28GB).  
Ensure that there is sufficient disk space (4.28GB+) and swap space (800MB) available on the machine where you are installing Oracle9iAS. (OUI will detect if there is not enough disk space.)
5. **Component Configuration and Startup:** Select Oracle HTTP Server in non-SSL mode.

OP requires the Oracle HTTP Server component to be installed.

Other deselected components in this screen will be installed but not configured or started later on. Only HTTP Server will be started after the installation completes.

HTTP Server can be accessed using the following URLs:

Non SSL Mode (executed at install time): `http://host_name:7778`

SSL mode: (executed at install time): `http://host_name:80` and  
`http://host_name:443`

6. **Origin Database Connection Information:** Provide host name, port number, and service name where Oracle9iAS is installed (for example, myhost, 1521, myhost.mycompany.com)
7. **Apache Listener Configuration for Oracle9iAS Portal:** Use the default setting.
8. **Origin Database User Information:** Provide the origin database "sys" account password.
9. **Wireless Edition repository information:** Use default settings unless you need to configure Wireless Edition repository.
10. **Wireless Edition schema information:** Use default settings.
11. **Enter SYSTEM Password for Wireless Edition:** Use default settings.

12. **Summary:** Examine the product list that will be installed; if it is correct, click the **Install** button to start installation.

There are five CDs for Oracle9iAS installation. Use the **Browse** button to change and locate the right disk number when OUI prompts you on the screen during the installation.

13. After installation is complete, OUI prompts you to run the `root.sh` script. Log in as root, change directory to `$ORACLE_HOME`, and run `root.sh`.
14. Start the origin database (if it has been shut down during the installation).
15. **Configuration Tools:** The Configuration Tools screen lists the configuration tools for all installed components. You should see at least Oracle HTTP Server being started here. Others will depend on the components you selected at the beginning of the installation. Verify that the list of configuration tools is correct, and click **Next** to continue.
16. **End of Installation.**

There are some post-installation steps related to component-specific tasks. For details, refer to the *Oracle9i Application Server Installation Guide*.

---

# Installing Oracle Workflow

The directions in this chapter represent the steps that are typically required to install Oracle Workflow on a UNIX platform; they are included here for your convenience. These directions may not be appropriate in all circumstances and for all platforms. If you encounter problems or if these directions are not appropriate for your site, see the *Oracle Workflow Guide*, release 2.6.1, for complete directions.

Oracle Workflow is installed on the system where Oracle9i is installed.

## Installing Oracle Workflow 2.6

Follow these steps to install Oracle Workflow 2.6:

1. Start Oracle Universal Installer (OUI) as described in *Oracle9i Administrator's Guide* for your platform.
2. **Welcome:** The **Welcome** window appears. Click **Next**.
3. **File Locations:** Click **Next**.
4. **Available Products:** Choose **Oracle9i Management and Integration 9.0.1.0.0** and click **Next**.
5. Oracle Workflow is a component of Oracle Integration Server, which is a component of Oracle9i Management and Integration Product. Select **Custom installation**, and click **Next**.
6. Find **Workflow** component under Oracle Integration Server and select it. Deselect other options that may be highlighted for install. Click **Next**.
7. **Summary:** Displays selected products that will be installed. Click **Install** to start the installation process.

8. **Workflow Configuration Assistant:** Workflow Configuration Assistant starts automatically during Workflow installation, and prompts you for the following Workflow configuration information:

- **Workflow Account:** `owf_mgr` (default)
- **Workflow Password:** This field is empty; provide the password here.
- **SYS Password:** Enter the new password you provided in step 2 of "Configuring the Database", in Chapter 3.
- **SYSTEM Password:** Enter the new password you provided in step 2 of "Configuring the Database", in Chapter 3.
- **Install Option:** Select **Install**.
- **Language Selection:** (default)
- **Connect Method:** Select **Local**.
- **Connect String:** Enter `net_service_name` (check `tnsnames.ora` file).

Click **Submit**. This completes Workflow configuration.

9. If Oracle HTTP Server is not already started, log in as root and start it. To test whether it is already started, start a browser and point it to your hostname:

```
http://<hostname>/
```

where `<hostname>` is the name or IP address of the system on which you are installing these components.

You'll see a page with the title "Oracle HTTP Server Components" and a list of server components. If you do not see this page, it means Oracle HTTP Server is not running, and you must start it. To do so, log in as root and enter:

```
# $ORACLE_HOME/Apache/Apache/bin/apachectl startssl
```

If Oracle HTTP Server is correctly started, you will see the following:

```
./apachectl start ssl:http started
```

If Oracle HTTP Server fails to start, check the log files to learn the reason.

Exit as root and return to your role as Oracle database administrator.

10. Configure a Database Access Descriptor (DAD) for Oracle Workflow as follows:

Point your browser to your `<hostname>`, and among the components listed,

- `select mod_plsql`.

Then click the **Gateway Database Access Descriptor Settings** link:



- **Database Access Descriptors:** Administer DAD Entries — on that page, click **Add Default** (blank configuration).
- **Database Access Descriptor:** Create DAD Entry — enter the following:
  - **Database Access Descriptor Name:** (enter any name; we suggest wf)
  - **Schema Name:** Leave this empty.
  - **Oracle User Name:** Leave empty.
  - **Oracle Password:** Leave empty.
  - **Oracle Connect String:** (the same Connect String that you specified in step 8)
  - **Default (Home) Page:** wfa\_html.home

Ensure that Oracle User Name and Oracle Password are blank, and that all other fields have their default settings.

Click **OK** (at the top).

You'll see WF added to the sample DADs listed. Click **Close**. **Exit** the browser.

11. Edit the file `$ORACLE_HOME/Apache/Apache/conf/httpd.conf` to add the following aliases:

```
Alias /OA_JAVA/ "<$ORACLE_HOME>/jlib/"
Alias /OA_MEDIA/ "<$ORACLE_HOME>/wf/java/oracle/apps/fnd/wf/icons/"
Alias /OA_DOC/ "<$ORACLE_HOME>/wf/doc/"
```

Be sure to substitute the full path of `$ORACLE_HOME` in each alias.

12. Restart Oracle HTTP Server:

- Log in as root.
- Stop Oracle HTTP Server:

```
# $ORACLE_HOME/Apache/Apache/bin/apachectl stop
```

- Wait 15 seconds.
- Start Oracle HTTP Server:

```
# $ORACLE_HOME/Apache/Apache/bin/apachectl startssl
```

If Oracle HTTP Server is correctly started, you will see the following:

```
./apachectl start ssl:http started
```

- Exit root.

13. To test workflow installation/configuration, point your browser to

`http://<yourhost>/pls/<wf-DAD>`

for `<wf-DAD>`, substitute the value for the Database Access Descriptor name that you chose in step 10.

To log in to Workflow Administrator, you'll be prompted for user name for Workflow Account) and password (Workflow Password) (these values are from step 8). This brings up the `wfa_html.home` Oracle Workflow page.

14. Now update the Global Workflow Preferences so that the workflow administrator is the only user with the `privlwdgw` to administer the workflow.

- On the `wfa_html.home` page, click **Global Workflow Preferences**.
- Click **Update** to make the fields editable.
- Click the drop-down arrow of the Workflow Administrator field.
- The **Find** window appears. In the **Find** field, enter your Workflow Account user name. Click **Find**. The DAD that you created appears, with a link. Click the link and it takes you to the finished page.
- Click **OK**. **Close** browser.

## Configuring Email Notifications (Optional)

You can configure Workflow email notifications so that OP can send an email message when a build, deployment, or report completes. Email is sent to the email address that is entered in the email notification fields on the **Create Build Schedule**, **Create Deployment Schedule**, and **Create Report Schedule** pages. The message can be sent to individual users and/or mailing lists. Configuring email notifications requires the following two steps:

1. Configure `sendmail` on the system where OP is installed.
2. Configure the Workflow Notification Mailer engine.

### Configure sendmail

You must configure `sendmail` to run on the system where OP is installed. You should verify that the `sendmail` daemon is running before you try to use email notifications. Consult your UNIX system administrator for details about configuring `sendmail`.

## Configure the Workflow Notification Mailer Engine

The exact details depend on the operating system that you are using. The general steps are as follows:

1. Log in as root. Create a UNIX account, for example, `opmail`, to run the Workflow mailer. You may want to add a group before creating this user account. The following steps are required on a Solaris system to add a group and a user (commands to add a group and a user may differ on other UNIX platforms; consult your UNIX system administrator):

- a. Add a group:

```
groupadd <opmail_group_name>
```

where `<opmail_group_name>` is a string of 8 or fewer lowercase alphabetic and numeric characters.

- b. Create a user account for the new user:

```
useradd -g <opmail_group_name> -d <opmail_user_home_directory>  
-s /bin/csh -c "Workflow Mailbox" -m <opmail_user_name>
```

where `<opmail_user_home_directory>` is the home directory for the new user and `<opmail_user_name>` is the login name of the new user. The login name must be a string of 8 or fewer characters; the only characters allowed in the login name are alphabetic characters, numeric characters, period (`.`), underscore (`_`), and hyphen (`-`).

- c. Set the password for the new user. You can choose any password allowed on your system.

```
passwd <opmail_user_name>
```

2. Log in to the system as `opmail`.
3. Create a directory named `mailFolders`.
4. Edit the `.login` file for `opmail` as follows:
  - a. Set the `ORACLE_HOME` and `ORACLE_SID` variables.
  - b. Add `$ORACLE_HOME/bin` and `/usr/lib` to the `PATH` variable.
  - c. Add the `LD_LIBRARY_PATH` variable set to `$ORACLE_HOME/lib`.
5. Log out and log in again.

6. Verify that `sqlplus` works from the `opmail` account. You can use the following command:

```
sqlplus /nolog
```

7. Log in to the user account that owns the Oracle software on your system.
8. Change directory to `$ORACLE_HOME/wf/res`. This directory contains the `wfmail.cfg` file.

The `wfmail.cfg` file is generic, and must be customized for your system. Edit `wfmail.cfg` to customize it for your system, carefully checking all the configuration settings for your system.

9. Again log in to the system as `opmail`.
10. Start the mailer with the following command:

```
$ORACLE_HOME/bin/wfmail.csh -f $ORACLE_HOME/wf/res/wfmail.cfg &
```

11. You will see two mailer processes (owned by the user `opmail`) start up. You can verify that the two mailer notification processes are running with the following command on Solaris:

```
ps -ef | grep wfmail
```

For more information, see the README file for your operating system and the documentation for Oracle Workflow.

---

## Installing OP

Before installing Oracle9iAS Personalization (OP), install Oracle9i database and Oracle9iAS (Chapter 3) and Oracle Workflow 2.6.1 (Chapter 4).

Then, follow the steps described in this chapter to install OP.

The installation of OP is divided into two top-level steps, as follows:

- The first step of the installation is generally performed by the database administrator on the back-end Oracle9i database. This part entails creating database accounts and schemas, loading stored procedures into the database, and installing necessary software and libraries. The components that are installed in the first step are the Mining Object Repository (MOR), Mining Table Repository (MTR), and Recommendation Engine (RE). Each of these three components requires separate schemas and tablespaces; for the MOR there is also a temporary tablespace dedicated to running data mining model builds.
- The second step of the installation is generally performed by the application server administrator on the system where the Oracle9i Application Server is installed. This step installs the Java libraries for runtime support of the Recommendation Engine API and the Recommendation Engine Batch API, along with REAPI demo. (The REAPI demo is a servlet that works with the demo movie data that can optionally be loaded into the MTR. The servlet can be used to verify that all the OP components are configured and working.)

Note that each step has manual steps that you must perform after the OUI installation completes.

### Installing OP Components

OP is installed on two different systems.

The following components, referred to as the "OP Application Server components," are installed on the system where Oracle9iAS is installed:

- Recommendation Engine API
- Recommendation Engine Batch API
- `redemo servlet`

All other OP components, referred to as the "OP database components" (including the Administrative UI and OP documentation), are installed on the system where Oracle9i is installed.

## Installing OP Database Components

Follow these steps to install the OP database components:

1. Log in to your Oracle DBA's operating system account on the system where Oracle9i is installed.
2. If any OP components are already installed on this system, deinstall them (see "Deinstalling OP", below).
3. In a shell, issue the following command to start Oracle Universal Installer (OUI):

```
<cdrom>/Disk1/runInstaller
```

where `<cdrom>` points to the directory where your cd is mounted. Be sure your `DISPLAY` environment variable is set appropriately.

4. **Welcome:** The **Welcome** screen appears. Click **Next**.
5. **File Locations:** Displays source and destination. Confirm that the correct source and destination paths are specified. The destination should be the Oracle Home directory. Then click **Next**.
6. **Available Products:** Select Oracle9iAS Personalization Database Components. Click **Next**.
7. **Installation Types:** Complete or Custom. Choose Custom only if you want to install certain components, for example, if you want to install everything except the MTR on one system and then install the MTR on a different system. Click **Next**.
8. **Available Product Components:** This page lists the components; check the ones you want to install. If you are doing a complete installation, everything is checked by default. Click **Next**.

9. **Database Information:** Provide the following information for your Oracle database, as established while installing the database. The fields may be filled in with default values. Confirm that these values are appropriate for your database or enter new values.

- **ORACLE\_SID:**
- **TNSName:**
- **INET Hostname:** This is the hostname or IP address of the system containing the database.
- **Port:**

Click **Next**.

10. **DB System Account Information:** Provide the following and then click **Next**:

- **DBA Name:** `system` (This is the default; provide a user name for a user with DBA privileges.)
- **DBA Password:** `manager` (This is the default; provide a password for a user with DBA privileges.)
- **Workflow Admin Name:** `owf_mgr` (This is the default; provide the Workflow administrator name created during Workflow installation.)
- **Workflow Admin Password:** `owf_mgr` (This is the default; provide the Workflow administrator password created during Workflow installation.)
- **OPWFBase Password:** The OPWFBase user is created during MOR configuration. Create the OPWFBase password here.
- **SYS Password:** Enter the password for SYS user (provided after database installation).

11. **OP Tablespace Data File Directory:** Specify the path and directory for OP tablespace data files. The default path is displayed as:

```
$ORACLE_HOME/oradata/<ORACLE_SID>
```

By clicking the **Browse** button, you can browse and assign another path for data file directory. If you accept the default path, confirm that this path exists. If it does not exist, OUI prints a message saying that the path does not exist, and your installation will fail. Click **Next**.

12. **MOR Configuration Information:** Provide the following and then click **Next**:

- **MOR User Name:** OPMOR (default)

- **MOR Password:** MORPW (default; change this password here)
  - **MOR Tablespace Name:** OPMOR (default)
  - **MOR Data File Size:** 512M (default)  
**Important:** Sizes must be specified in megabytes. If, for example, you want 1 gigabyte, you must convert it to megabytes and enter 1024M. (For 2 gigabytes, however, you cannot specify 2048M, because 2047M is Oracle's limit.)
  - **MOR Temp Tablespace Name:** OPMOR\_TEMP\_TS (default)
  - **MOR Temp Data File Size:** 1024M (default)
13. **RE Configuration Information:** Provide the following and then click **Next**.
- **RE User Name:** RE (default)
  - **RE Password:** REPW (default; change this password here)
  - **RE Tablespace Name:** RE (default)
  - **RE Data File Size:** 512M (default)
14. **MTR Configuration Information:** Provide the following and then click **Next**:
- **MTR User Name:** DEMO\_MTR (default)
  - **MTR User Password:** MTRPW (default; change this password here)
  - **MTR Tablespace Name:** DEMO\_MTR (default)
  - **MTR Data File Size:** 512M (default)
15. **Populated MTR Schema:** Choose whether you want to install an MTR schema populated with demo data. The default selection is **Yes**, which results in an MTR schema populated with the demo data. Select **No** if you want to use your own data (you will get the same schema structure, but without data). (See Chapter 8 and Chapter 9.) Click **Next**.
16. **Summary:** Lists elements to be installed. Click **Install**.
17. **Install:** OUI starts installing. (If this is the first installation of OP, you will see several messages such as "RE doesn't exist" — this is normal. OP creates all necessary files and tables.)



18. **Configuration Tools:** There is a tool for each component being configured; use the vertical scroll bar on the right to see the names of all the tools:
  - **RE Create Wizard:** Once this tool completes successfully, it creates an RE user and RE schema on the specified database.
  - **MTRDemo Create Wizard:** When it completes successfully, it creates an MTR user and MTR schema populated with demo data on the specified database.
  - **MOR Create Wizard:** When it completes successfully, it creates an MOR user and an MOR schema and OPWFBASE user and OPWFBASE schema on the specified database.
19. **End of Installation:** Click **Exit**.
20. You may have to stop and restart the database, especially if you are reinstalling OP.

---

**Note:** Whenever you stop and restart the database, you should check to see if Listener is still running. If Listener is not running, you should restart it.

---

## Installing OP Application Server Components

Follow these steps to install the OP Application Server components:

1. Log in to your Oracle DBA's operating system account on the system where Oracle9iAS is installed.
2. If any OP components are already installed on this system, deinstall them (see "Deinstalling OP", below).
3. In a shell, issue the following command to start Oracle Universal Installer (OUI):

```
<cdrom>/Disk1/runInstaller
```

where <cdrom> points to the directory where your cd is mounted. Be sure your DISPLAY environment variable is set appropriately.

4. **Welcome:** The **Welcome** screen appears. Click **Next**.
5. **File Locations:** Displays source and destination. Confirm that the correct source and destination paths are specified. The destination should be the Oracle Home directory. Then click **Next**.

6. **Available Products:** Select Oracle9iAS Personalization *iAS* Components. Click **Next**.
7. **Installation Types:** Complete or Custom. We recommend that you select Complete. Click **Next**.
8. **Available Product Components:** This page lists the components; check the ones you want to install. REAPI Demo Test Servlet should be checked. If you are doing a complete installation, everything is checked by default. Click **Next**.
9. **Summary:** Lists elements to be installed. Click **Install**.
10. **Install:** OUI starts installing. OP creates all necessary files and tables.
11. **End of Installation:** Click **Exit**.

## Manual Steps

After completing an OP installation and configuration that includes an MOR, there are a few things you need to do manually on both systems where you've installed OP.

### Manual Steps for Oracle9i System

Perform the following steps on the system where Oracle9i is installed.

1. Log in to your Oracle DBA's operating system account on the system where Oracle9i is installed.
2. Update `jserv.conf` file: Go to  
`$ORACLE_HOME/Apache/Jserv/etc/`  
and append the file  
`$ORACLE_HOME/dmt/conf/jservConf.add`  
to `jserv.conf` file.
3. Update `jserv.properties`: Append the file  
`$ORACLE_HOME/dmt/conf/jservProperties.add`  
to `jserv.properties` file.

4. In the `jserv.properties` file, update the `DISPLAY` parameter to point to the display at the host that is running X Windows by replacing the name of the X Windows server host name in `wrapper.env=DISPLAY=<hostname>:0.0`

Verify that there is only one entry configuring

`wrapper.env=DISPLAY=<hostname>`. If you have more than one entry, comment the others out.

The system that is running the `Xdisplay` must allow the use of its display by the `java awt` toolkit with the command

```
xhost + <system_where_Jserv_runs>
```

**Note:** You need X Windows to render images related to the OP Administrative UI. Ideally, X Windows runs on the machine on which you have installed OP. If this is inconvenient or impractical, X Windows must run on a machine that is connected to the machine on which you installed OP.

5. Stop and restart the database.

---

**Note:** Whenever you stop and restart the database, you should check to see if Listener is still running. If Listener is not running, you should restart it.

---

6. Log in as root and restart the Oracle HTTP Server server with these commands:

- Stop Oracle HTTP Server:

```
$ORACLE_HOME/Apache/Apache/bin/apachectl stop
```

- Wait 15 seconds.

- Start Oracle HTTP Server in ssl mode:

```
$ORACLE_HOME/Apache/Apache/bin/apachectl startssl
```

If the command succeeds, you see the following line:

```
./apachectl start ssl:http started
```

- Exit root.

7. To start Oracle Workflow engine, as Oracle DBA, use the following command:

```
/usr/bin/csh $ORACLE_HOME/dmt/install/WF/opwfengine.csh start <MOR_schema>
```

You will be prompted to enter the following:

- Workflow user name (default was `owf_mgr`; use the user name you provided in step 10 earlier in this chapter)

- Workflow password (default was `owf_mgr`; use the password you provided in step 10 earlier in this chapter)
- Database TNS name

## Manual Steps for Oracle9iAS System

Perform the following steps on the system where Oracle9iAS is installed.

1. Log in to your Oracle DBA's operating system account on the system where Oracle9iAS is installed.
2. Update `jserv.conf` file on the system where Oracle9iAS is installed: Go to

```
$ORACLE_HOME/Apache/Jserv/etc/
```

and append the file

```
$ORACLE_HOME/dmt/conf/oprocMgr/reapiJservConf.add
```

to `jserv.conf` file.

3. Log in as root and restart the Oracle HTTP Server with these commands:

- Stop Oracle HTTP Server:

```
$ORACLE_HOME/Apache/Apache/bin/apachectl stop
```

- Wait 15 seconds.
- Start Oracle HTTP Server in ssl mode:

```
$ORACLE_HOME/Apache/Apache/bin/apachectl startssl
```

If the command succeeds, you see the following:

```
./apachectl start ssl:http started
```

- Exit root.

## Test the OP Installation

To confirm that OP is properly installed, do the following:

- Verify that you can run the OP Administrative UI by pointing your browser to the URL `http://<hostname>/OP/Admin`, where `<hostname>` is the name of the server on which Oracle9i is installed.
- Verify that you can run the REAPI demo and use it to get a recommendation, as described in *Getting Started with Oracle9iAS Personalization*. To start REAPI demo,

point your browser to the URL `http://<hostname>/redemo`, where `<hostname>` is the name of the server on which Oracle9iAS is installed.

## Deinstalling OP

If there is a copy of OP installed on either of your systems, deinstall it from both systems before installing a new version.

## Deinstalling Oracle9i OP Components

Follow these steps to deinstall the OP components that are installed on the system where Oracle9i is installed:

1. Log in to your Oracle DBA system account on the system where Oracle9i is installed.
2. Stop Oracle Workflow with the command
3. Make sure no OP users are connected to the database. To find out whether there are any OP user sessions, connect to the database as SYSTEM user and execute the following SQL statement:

```
/usr/bin/csh $ORACLE_HOME/dmt/install/WF/opwfengine.csh stop <MOR_schema>
```

```
SQL> select username from v$session where status='ACTIVE' or status='INACTIVE';
```

and close the connections or kill the existing sessions for OP users.

4. All OP processes are now stopped or removed, and you are ready to deinstall OP. Log in to your Oracle DBA system account.
5. Set the DISPLAY environment variable if necessary.
6. Run the following command to start Oracle Universal Installer:
7. **Welcome:** Click **Deinstall Products**.
8. **Inventory:** This window displays the installed products in the available Oracle home.
9. Select **Oracle9iAS Personalization Database Components 9.0.0.1.0**, and click **Remove**. A confirmation window appears; click **Yes**.
10. **Remove:** The **Remove** window appears; it displays a progress bar.

```
<cdrom>/Disk1/runInstaller
```

11. You may want to remove data files corresponding to MOR and RE tablespaces for the deinstalled OP. MTR schema and associated data files that are not removed when you deinstall. If you want to install an MTR with the same user name, the current MTR will be dropped and replaced with the new one.
12. Remove the `$ORACLE_HOME/dmt` directory.

## Deinstalling Oracle9iAS OP Components

Follow these steps to deinstall the OP components that are installed on the system where Oracle9i is installed:

1. Log in to your Oracle DBA system account on the system where Oracle9iAS is installed.
2. Set the DISPLAY environment variable if necessary.
3. Run the following command to start Oracle Universal Installer:  

```
<cdrom>/Disk1/runInstaller
```
4. **Welcome:** Click **Deinstall Products**.
5. **Inventory:** This window displays the installed products in the available Oracle home.
6. Select **Oracle9iAS Personalization iAS Components 9.0.1.0.0**, and click **Remove**. A confirmation window appears; click **Yes**.
7. **Remove:** The **Remove** window appears; it displays a progress bar.
8. Remove the `$ORACLE_HOME/dmt` directory.

---

## OP Security

Oracle9iAS Personalization follows the Oracle Data Server security policy. OP implements the following measures:

- Any data that needs to be protected is encrypted using standard Oracle security tools.
- Transmission of sensitive data is encrypted between various OP components.
- The DBA grants access to qualified users.
- OP uses secure SSL technology to access Oracle HTTP Server.
- To invoke Workflow Engine background processes, the user needs to know the Workflow administrator user name and password.

### User Access

The DBA is responsible for creating database users and setting up proper privileges for them so that they can access OP.

OP controls user access by providing two database roles:

- **OP\_ADMIN:** This role allows the user to access the OP product as an OP Administrator. This means having access to the entire OP product, including the Message Viewer.
- **OP\_REPORT:** This role allows the user to view OP reports and access the Report Viewer URL (<http://<hostname>/OP/Admin/ReportsPage>).

By default, the MOR Schema owner is the OP Administrator, and thus has the full privileges to access the OP product.

For new OP users, the DBA needs to grant either OP\_ADMIN or OP\_REPORT to the new user as follows:

```
SQL> grant OP_ADMIN to <new_user>;
```

```
SQL> grant OP_REPORT to <new_user>;
```

For existing database users, the DBA must grant either role to allow them access to OP product at the appropriate level.



---

# Configuring Oracle9iAS Personalization

After you have installed OP and verified that the installation is correct, you can specify certain configuration parameters for OP.

This chapter lists the external OP configuration parameters and their default values, and indicates the parameters that you can change and how to change them. This chapter also describes data synchronization and how to configure it.

You should read this chapter in conjunction with Appendix A, which describes how to determine appropriate parameter values for your installation.

All OP configuration parameters reside on the system where Oracle9i is installed.

## Changing OP Configuration Parameters

These parameters are divided into three categories:

- Values that are changed using SQL Plus, indicated by Y in the **Change** column in the summary tables.
- Values that you should *not* change, indicated by N in the **Change** column in the summary tables.
- Values that are changed using the OP Administration UI, indicated by UI in the **Change** column in the summary tables.

## RE Configuration Parameters

Table 7-1 lists the RE configuration parameters, their data types, their default values, and a description for each. These parameters can be found in the RE\_CONFIGURATION table.

**Table 7-1 Recommendation Engine Configuration Parameters**

Parameter Name	Data Type	Default Value	Description	Change
LOG_LEVEL	int	2	0=OFF, 1=INTERNAL ERROR plus Error and Warning, 2=All errors logged for 1 plus notifications, 3=All errors logged for 2 plus trace	Y
RE_TRACE	int	0	0=OFF, 1=LOW (detail), 2=MODERATE, 3=HIGH	Y
REAPIRT_TRACE	int	0	0=OFF, 1=LOW (detail), 2=MODERATE, 3=HIGH	Y
REAPIDEMO_TRACE	int	0	0=OFF, 1=LOW (detail), 2=MODERATE, 3=HIGH	Y
UTIL_TRACE	int	0	0=OFF, 1=LOW (detail), 2=MODERATE, 3=HIGH	Y
REAPIBATCH_TRACE	int	0	0=OFF, 1=LOW (detail), 2=MODERATE, 3=HIGH	Y
TimeoutInterval	int	1800	Session timeout interval (in seconds)	UI
TimeoutFlag	int	1	Session timeout indicator (1=TRUE, 0=FALSE)	UI
DataSyncInterval	int	1800	Interval on which to synchronize customer data (in seconds)	UI
SyncCustomerNavigationalData	int	0	Is customer navigational data synchronized (boolean)	N
SyncCustomerRatingData	int	0	Is customer rating data synchronized (boolean)	N
SyncVisitorNavigationalData	int	0	Is visitor navigational data synchronized (boolean)	N
SyncVisitorRatingData	int	0	Is visitor rating data synchronized (boolean)	N
SyncPurchasingData	int	0	Is customer purchasing data synchronized (boolean)	N
SyncDemographicData	int	0	Is customer demographic data synchronized (boolean)	N

**Table 7–1 Recommendation Engine Configuration Parameters**

Parameter Name	Data Type	Default Value	Description	Change
ConnectionPoolSize	int	128	Java connection pool limit per proxy.	Y <sup>1</sup>

<sup>1</sup> Requires restart of OP.

## MOR Configuration Parameters

Table 7–2 describes the configuration parameters for the OP Mining Object Repository (MOR). The table shows their data types, their default values, and a description for each. These parameters can be found in the MOR\_CONFIGURATION table.

The table includes a column titled **Change**. If the value in this column is "N," do not change the parameter. If the value in this column is "Y," the value of the parameter must be changed to a value suitable for your environment. The description of these parameters includes the instruction "change value on install."

**Table 7–2 MOR Configuration Parameters**

Parameter	Data Type	Value	Description	Change
MOR_USERNAME	string	<user name>	User name for Admin UI; change value on install	N
MOR_PASSWORD	string	<password>	Password for Admin UI; change value on install	N
MOR_DBALIAS	string	<alias>	Alias for the MOR database; change value on install	N
MOR_SCHEMA	string	< schema>	MOR schema name	N
WF_SCHEMA	string	<workflow schema>	Workflow schema name	N
MOR_HOST_URL	string	<hostname>	MOR hostname; change value on install	N
MOR_PORT	string	<port>	MOR port; change value on install	N

**Table 7–2 MOR Configuration Parameters**

Parameter	Data Type	Value	Description	Change
MOR_SID	string	isab900	MOR system ID; change value on install	N
MOR_VERSION	string	9.0.1	MOR version number	N
scheduleItemGracePeriod	int	60	Number of minutes a scheduled item must have been past due for it to cause an error	Y
browserToBuyerConversion <sup>1</sup>	int	20	Maximum number of reports of this type to keep in log	UI
recommendationEffectiveness	int	20	Maximum number of reports of this type to keep in log	UI
mostActiveCrossSellProducts	int	20	Maximum number of reports of this type to keep in log	UI
buildEvents	int	20	Maximum number of events of this type to keep in log	UI
deployEvents	int	20	Maximum number of events of this type to keep in log	UI
reportEvents	int	20	Maximum number of events of this type to keep in log	UI
LOG_LEVEL	int	2		Y
ALGS_TRACE	int	0	0=OFF, 1=LOW (detail), 2=MODERATE, 3=HIGH	Y
DMAPI_TRACE	int	0	0=OFF, 1=LOW (detail), 2=MODERATE, 3=HIGH	Y
PAR_TRACE	int	0	0=OFF, 1=LOW (detail), 2=MODERATE, 3=HIGH	Y

**Table 7–2 MOR Configuration Parameters**

Parameter	Data Type	Value	Description	Change
TNB_TRACE	int	0	0=OFF, 1=LOW (detail), 2=MODERATE, 3=HIGH	Y
UI_TRACE	int	0	0=OFF, 1=LOW (detail), 2=MODERATE, 3=HIGH	Y
UTIL_TRACE	int	0	0=OFF, 1=LOW (detail), 2=MODERATE, 3=HIGH	Y
WFJAVA_TRACE	int	0	0=OFF, 1=LOW (detail), 2=MODERATE, 3=HIGH	Y
MAIL_PREFERENCE <sup>2</sup>	string	MAILHTML	Alternative value: MAILTEXT	Y
NLS_LANGUAGE <sup>3</sup>	string		Supported values: AMERICAN, FRENCH, GERMAN, ITALIAN, SPANISH, BRAZILIAN PORTUGUESE, JAPANESE, KOREAN, SIMPLIFIED CHINESE	Y
NLS_TERRITORY <sup>4</sup>	string		Supported values: AMERICA, FRANCE, GERMANY, ITALY, SPAIN, BRAZIL, JAPAN, KOREA, CHINA	Y

<sup>1</sup> The term "browserToBuyer" was changed to "visitorToCustomer"; the two terms mean the same thing. The original term persists in some of the code.

<sup>2</sup> The email notifications sent by OP support either plain text or HTML format. The default setting in the MOR configuration table is HTML. If the incoming mail server that you are using supports only plain text, set this parameter to MAILTEXT.

<sup>3</sup> The default value for NLS\_LANGUAGE after OP installation is that this is empty (blank). However, if it is edited, OP verifies that the value is supported and uses it instead of the database default language for the server code (it does not affect PL/SQL messages however). If it is not specified, the database default language is used. If neither these values nor the database default indicates one of the supported languages, OP workflow defaults to AMERICAN.

<sup>4</sup> The default value for NLS\_TERRITORY after OP installation is that this is empty (blank). However, if it is edited, OP verifies that the value is supported and uses it instead of the database default language for the server code (it does not affect PL/SQL messages however). If it is not specified, the database default language is used. If neither these values nor the database default indicates one of the supported languages, OP workflow defaults to AMERICA.

MTR Configuration Parameters

Table 7–3 describes the configuration parameters for the OP Mining Table Repository (MTR). The table shows their data types, their default values, and a description for each. These parameters can be found in the MTR\_CONFIGURATION table in MTR schema.

These parameters allow selecting different types of data to be synchronized to the MTR. At the end of an OP session, MTR synchronization adds data collected in the RE (during the session) to the data already stored in the MTR. In order for data synchronization to take place, the MTR must be configured to allow the various types of data to be synchronized.

Table 7–3 MTR Configuration Parameters<sup>1</sup>

Parameter	Data Type	Value	Description	Change
MTR_ALLOW_SYNC_DEMOGRAPHIC	boolean	T	Allow demographic data to be synchronized to MTR	Y
MTR_ALLOW_SYNC_NAVIGATION	boolean	T	Allow navigational data to be synchronized to MTR	Y
MTR_ALLOW_SYNC_PURCHASING	boolean	T	Allow purchasing data to be synchronized to MTR	Y

**Table 7–3 MTR Configuration Parameters<sup>1</sup>**

Parameter	Data Type	Value	Description	Change
MTR_ALLOW_SYNC_RATING	boolean	T	Allow rating data to be synchronized to MTR	Y
ALLOW_MTR_SYNC_VISITOR_NAVIGATION	boolean	T	Allow visitor navigation data to be synchronized to MTR	Y
MTR_ALLOW_SYNC_VISITOR_RATING	boolean	T	Allow visitor rating data to be synchronized to MTR	Y

<sup>1</sup> Change these values only when the MTR is not being used.

## Data Synchronization

Data synchronization moves user-specific data that is collected in the RE during a session to permanent storage, that is, to the appropriate table in the mining table repository (MTR). Session and recommendation data are always synchronized; other kinds of data are synchronized according to how the RE Farm and MTR are configured. See "Data to Synchronize", later in this chapter, for configuration instructions. Customer data and visitor data are copied to the appropriate MTR tables. (There is one set of MTR tables for customer data and a different set for visitor data.)

Data is synchronized every `DataSyncInterval`. `DataSyncInterval` is a configuration parameter that is specified for an RE Farm. Data synchronization is performed only for users whose sessions are inactive. A session is *inactive* if there has been no activity for a specified period of time or if the session has been explicitly closed. Note that a user can have more than one session at any time. A customer ID is deleted from `RE_PROFILE_DATA` only when all of the customer's sessions are inactive.

After the data is copied to the MTR, the data is purged (deleted) from the RE tables. Data that cannot be synchronized for some reason (for example, data that has an invalid item ID) is also purged.

Data is collected in the `RE_CURRENT_SESSION_DATA` table and the `RE_RECOMMENDATION_DETAIL` table. The data source type of the data determines the MTR table to which data is copied.

Table 7–4 shows the four data source types and the MTR table for each.

**Table 7–4 Data Synchronization for RE\_CURRENT\_SESSION\_DATA**

DATA_SOURCE_TYPE	MTR Table
1 (demographic) for customers only	MTR_CUSTOMER
2 (purchasing) for customers only	MTR_PURCHASING_DETAIL
3 (rating) for visitors and customers	MTR_RATING_DETAIL
4 (navigational) for visitors and customers	MTR_NAVIGATIONAL_DATA

RE\_RECOMMENDATION\_DETAIL data is copied to the MTR\_RECOMMENDATION\_DETAIL table and appropriate RE\_ACTIVE\_USER data is copied to MTR\_SESSION table. RE\_PROFILE\_DATA is updated in the MTR\_CUSTOMER table.

## Configuring Data Synchronization

You specify data synchronization in two places: you specify the synchronization interval for a Farm and exactly what data to synchronized for a specific MTR connection.

### Synchronization Interval

In the **Farms** tab of the Administrative UI, select a farm, select **Edit**, and then click the **Advanced Settings** button. Specify an appropriate data synchronization interval for the selected farm. (You can also specify the timeout interval here.)

The default synchronization interval is 300 seconds (5 minutes). The synchronization interval should be adjusted based on the number of users of the application. If there are many users and the synchronization interval is long, the REs will fill with data.

### Data to Synchronize

In order for data synchronization to take place, the MTR must allow that type of data to be synchronized. These rules are specified when you install OP.

You configure the MTR connection using the OP Administrative UI. At the top of the **Farms** tab, click **Options**, click **MTR database connections**, click **Edit**, and finally click the **Sync settings** button. The synchronization settings for this MTR are displayed. To change a setting click the appropriate checkbox.



The types of data that are allowed to be synchronized are indicated by a checkmark in the corresponding checkbox. If a selection is greyed out, the configuration of the MTR does not allow synchronization of that type of data.

By default, *no* data is synchronized. You can choose to allow synchronization of any type of data for which the MTR allows synchronization. Any changes apply only to the current MTR connection.



---

## OP Schemas

Oracle9iAS Personalization (OP) uses several database schemas as follows:

- **Mining Object Repository (MOR)**, with tables containing mining objects (deployable packages, reports, schedule items, etc., as well as package build results such as rule tables, etc.)
- **Mining Table Repository (MTR)**, holding mining tables (bin boundaries, customer information, hot picks, taxonomy, etc.)
- **RE Schema**, with tables supporting current session information, model rules
- **Site database** holding demographic information

All OP schemas reside on the system where Oracle9i is installed.

To see a small example of the MOR, MTR, and RE schemas that are correctly populated, install the REAPI Demo, as described in Chapter 5, and examine the tables there. Alternatively, you can install an unpopulated MTR when you install OP. You can examine the schema of the unpopulated MTR and populate it with your own data.

Note that OP uses a fixed schema for the MTR. By "fixed," we mean that the MTR must be populated with tables matching OP table and column names.

Before you can obtain recommendations, you must create and deploy a package. You cannot create a package until there is some data available in the MTR. There are two ways to populate an MTR:

- Collect data using the REProxyRT method `addItem` or `addItems`.
- Populate with existing data, i.e., convert from existing data that was collected by your Web application and stored in an Oracle database.

## Mining Table Repository

The OP MTR consists of several tables and views. They are listed below.

MTR_ATTR_ID_BIN_BOUNDARY	TABLE
MTR_ATTR_NAME_TO_ID_MAP	TABLE
MTR_BIN_BOUNDARY	TABLE
MTR_CATEGORY	TABLE
MTR_CONFIGURATION	TABLE
MTR_CUSTOMER	TABLE
MTR_CUSTOMER_NAV_DETAIL	TABLE
MTR_CUSTOMER_RATING_DETAIL	TABLE
MTR_HOTPICK	TABLE
MTR_HOTPICK_GROUP	TABLE
MTR_INTERNAL_CONFIGURATION	TABLE
MTR_ITEM	TABLE
MTR_NAVIGATION_DETAIL	VIEW
MTR_PROFILE_DATA	VIEW
MTR_PROXY	TABLE
MTR_PURCHASING_DETAIL	TABLE
MTR_RATING_DETAIL	VIEW
MTR_RECOMMENDATION_DETAIL	TABLE
MTR_SCHEAM_VERSION	VIEW
MTR_SESSION	TABLE
MTR_TAXONOMY	TABLE
MTR_TAXONOMY_CATEGORY	TABLE
MTR_TAXONOMY_CATEGORY_ITEM	TABLE
MTR_VISITOR_NAV_DETAIL	TABLE
MTR_VISITOR_RATING_DETAIL	TABLE

Certain tables that make up the OP MTR must be populated with data specific to your Web site in accordance with the MTR schema. Other tables, such as the tables associated with sessions and recommendations, are automatically populated by OP.

The rest of this section describes the schemas for the MTR tables; tables that you must populate are described in detail.

## Item Table

The item table contains a list of all the individual items that a Web site deals with. When OP returns a recommendation, it returns the ID and type of the item; the item table provides more information. The item table is usually mapped to the catalog tables in the site database. The schema for MTR\_ITEM has four fields; they are listed below, in order, with their data types.

ID	NUMBER PK
ITEM_TYPE	VARCHAR2(30) PK
LABEL	VARCHAR2(150)
DESCRIPTION	VARCHAR2(4000)

## Bin Boundaries

The model building algorithms in Oracle9iAS Personalization require discrete data. All numerical data must be **discretized** or binned before the data is used to build a model. In OP, discretization is performed in a transformation step before model build. The value ranges for discretization (the bin boundaries) must be specified in order for OP to perform discretization.

In release 9.0.1 of OP, the bin boundaries must be explicitly specified in the bin boundaries table.

Discretization is performed by joining the input data and the bin boundaries table.

Categorical data in cases where there are a large number of distinct values should also be binned. If you bin categorical data, bin boundaries must be specified as for numerical data

In summary, OP requires all numerical data to be binned, and high cardinality categories should also be binned.

When you create bins of numeric values, specify the bounds for each bin (upper and lower values); when you create bins of categorical data, specify the items in each bin. To map several values to the same bin, use several records with the same bin numbering.

The table MTR\_BIN\_BOUNDARY has seven fields; they are listed below, in order, with their data types:

DATA_SOURCE_TYPE	NUMBER(3)
ITEM_TYPE	VARCHAR2(30)
ATTRIBUTE_NAME	VARCHAR2(30)

LOWER_VALUE	NUMBER
UPPER_VALUE	NUMBER
STRING_VALUE	VARCHAR2 ( 60 )
BIN_NUMBER	NUMBER ( 15 )

### Examples of Specifying Bin Boundaries

The following examples illustrate how to specify bin boundaries.

Consider movie rating data on a scale of 1 - 5. Suppose that you want to discretized ratings as follows:

- 1 and 2 are in bin number 1
- 3 is in bin number 2
- 4 and 5 are bin number 3

You should insert the following entries into the bin boundaries table:

```
(3, 'MOVIE', 'VALUE', 1, 2.1, NULL, 1),
(3, 'MOVIE', 'VALUE', 3, 3.1, NULL, 2),
(3, 'MOVIE', 'VALUE', 4, 5.1, NULL, 3).
```

The range of the bin includes all values that are greater than or equal to the lower value and strictly less than the upper value. The data source type for rating is 3 and string value is set to NULL for numeric data.

The following bin boundary table entries discretize martial status, a categorical attribute:

```
(1, 'NONE', 'MARITAL_STATUS', NULL, NULL, 'Single', 1),
(1, 'NONE', 'MARITAL_STATUS', NULL, NULL, 'Divorced', 2),
(1, 'NONE', 'MARITAL_STATUS', NULL, NULL, 'Separated', 2),
(1, 'NONE', 'MARITAL_STATUS', NULL, NULL, 'Married', 3),
(1, 'NONE', 'MARITAL_STATUS', NULL, NULL, 'Widowed', 4)
```

The data source type is 1 and the item type is NONE for demographic data. Lower value and upper value are NULL for categorical data.

## Taxonomy

In OP, a taxonomy refers to the structural organization of items in a company's catalog or site. Typically the catalog and/or the site has a hierarchical structure like

a tree or collection of trees (branching from broader groups at the top all the way down to individual items at the leaves).

Items can belong to more than one category and to different levels of the structure. The structure of the OP taxonomy is a DAG (direct acyclic graph), which can contain multiple top-level nodes. The different portions of the taxonomy can be disconnected too. Any node can connect to any other node but there cannot be a path that connects a node's child back to the node itself.

OP also supports multiple taxonomies (different ways of organizing the items).

The taxonomy is implemented using a group of tables (they are all specified by the customer at installation time):

- **MTR\_TAXONOMY:** Lists the different taxonomies used by the site. The schema for this table has three fields; they are listed below, in order, with their data types:

ID	NUMBER PK
NAME	VARCHAR2(150)
DESCRIPTION	VARCHAR2(4000)

- **MTR\_TAXONOMY\_CATEGORY:** Specifies which categories belong to the different taxonomies. (A category can belong to multiple taxonomies; however, for a given taxonomy, there can be only one instance of any category.) The schema for this table has four fields; they are listed below, in order, with their data types:

TAXONOMY_LEVEL	NUMBER PK
TAXONOMY_ID	NUMBER PK
PARENT_ID	NUMBER PK
CHILD_ID	NUMBER PK

- **MTR\_TAXONOMY\_CATEGORY\_ITEM:** Specifies which items go with a given taxonomy, category pair. The schema for this table has four fields; they are listed below, in order, with their data types:

CATEGORY_ID	NUMBER PK
TAXONOMY_ID	NUMBER PK
ITEM_ID	NUMBER PK
ITEM_TYPE	VARCHAR2(30) PK

- **MTR\_CATEGORY:** Specifies the different categories used by the site. The schema for this table has three fields; they are listed below, in order, with their data types.

ID	NUMBER PK
NAME	VARCHAR2(150)
DESCRIPTION	VARCHAR2(4000)

### Samples of the MTR Taxonomy Tables

The REAPI Demo includes a taxonomy; you can examine the demo MTR to see examples of all of these tables.

## Customer Table

The MTR\_CUSTOMER table contains demographic information about the customer. Some customer attributes are common to all OP applications and some can be tailored to your application. The common attributes are customer ID, name, creation date, gender, age, marital status, personal income, whether or not the customer is the head of household, household income, household size, and whether the customer rents or owns.

You can specify up to 50 attributes specific to your Web application. These variable attributes are all strings.

The schema of the MTR\_CUSTOMER table has the following fields. They are listed below, in order, with their data types.

ID	VARCHAR2(32)
NAME	VARCHAR2(80)
CREATION_DATE	DATE
GENDER	VARCHAR2(10)
AGE	NUMBER(3)
MARITAL_STATUS	VARCHAR2(20)
PERSONAL_INCOME	NUMBER
IS_HEAD_OF_HOUSEHOLD	CHAR(1)
HOUSEHOLD_INCOME	NUMBER
HOUSEHOLD_SIZE	NUMBER(2)
RENT_OWN_INDICATOR	VARCHAR2(30)
ATTRIBUTE1	VARCHAR2(150)



ATTRIBUTE2	VARCHAR2 ( 150 )
ATTRIBUTE3	VARCHAR2 ( 150 )
ATTRIBUTE4	VARCHAR2 ( 150 )
...	
ATTRIBUTE49	VARCHAR2 ( 150 )
ATTRIBUTE50	VARCHAR2 ( 150 )

## Hot Picks

Hot picks are used by some Web sites; for example, the daily specials might be hot picks. Information about hot picks is stored in two MTR tables, as follows:

- **MTR\_HOTPICK\_GROUP** lists the distinct hot picks groups used by the site. There is one record for each group. Each record contains a group ID, the group name (label), and a brief description of the group. The schema for this table has three fields; they are listed below, in order, with their data types.

ID	NUMBER PK
LABEL	VARCHAR2 ( 150 )
DESCRIPTION	VARCHAR2 ( 400 )

- **MTR\_HOTPICK** lists the items in each hot picks group, arranged according to group ID. Each record consists of a group ID, an item ID, and an item type. The schema for this table has three fields; they are listed below, in order, with their data types.

ITEM_ID	NUMBER
ITEM_TYPE	VARCHAR2 ( 30 )
GROUP_ID	NUMBER

A hot pick group can also contain categories. In this case, the item-type is set to CATEGORY and item ID is set to the appropriate ID value in the MTR\_CATEGORY table.

## Detail Tables

Several tables in the MTR store the details of various activities.

- **MTR\_CUSTOMER\_NAV\_DETAIL** stores the navigation data corresponding to a customer session. This table is populated with data collected in the RE.

- **MTR\_CUSTOMER\_RATING\_DETAIL** stores rating data for customers. This table is populated using the data collected in the RE during data collection.
- **MTR\_PURCHASING\_DETAIL** stores purchasing data on a per-session basis. Typically this data is collected by the Web application.
- **MTR\_RECOMMENDATION\_DETAIL** stores the results of recommendation requests. The data stored in this table is used to generate reports on the performance of OP.
- **MTR\_VISITOR\_NAV\_DETAIL** stores the navigation data corresponding to a visitor session. This table is populated with data collected in the RE.
- **MTR\_VISITOR\_RATING\_DETAIL** stores rating data for visitors. This table is populated using the data collected in the RE during data collection.

## Miscellaneous MTR Tables

The following tables are used internally by OP:

- **MTR\_ATTR\_NAME\_TO\_ID\_MAP** is used to speed up package building.
- **MTR\_CONFIGURATION** and **MTR\_INTERNAL\_CONFIGURATION** stores configuration information.
- The **MTR\_SESSION** table stores information about the session that OP creates internally on behalf of the application.
- **MTR\_ATTR\_ID\_BIN\_BOUNDARY** is a materialized view of the join of the **BIN\_BOUNDARIES** table and the **ATTR\_NAME\_TO\_ID\_MAP** table. It is used when transforming data during package builds.
- **MTR\_PROXY** is used to set up proxies for new items. When a new item is introduced, there may not be enough detail information about it to build packages; OP uses data about a similar existing product.

## Recommendation Engine

The RE schema stores current session data. The data is synchronized back to the MTR automatically. The RE includes the following tables (partial list):

ATTR_ID_BIN_BOUNDARY	TABLE
HOTPICK	TABLE
HOTPICK_GROUP	TABLE
I_I_ANCECEDENT	TABLE
I_I_RULE	TABLE
P_I_CATEGORY_RULES	TABLE
P_I_ITEM_RULES	TABLE
RE_ACTIVE_USER	TABLE
RE_CONFIGURATION	TABLE
RE_CURRENT_SESSION_DATA	TABLE
RE_DEPLOYABLE_PACKAGE	TABLE
RE_DEPLOYABLE_PKG_CONTENTS	TABLE
RE_ERROR_TABLE	TABLE
RE_INTERNAL_CONFIGURATION	TABLE
RE_LOG	TABLE
RE_MESSAGE_LOG	TABLE
RE_PROFILE_DATA	TABLE
RE_RECOMMENDATION_DETAIL	TABLE
RE_SCHEMA_ACCESS	TABLE
TAXONOMY_CATEGORY	TABLE
TAXONOMY_CATEGORY_ITEM	TABLE
TAXONOMY_TRANS_CLOSURE	TABLE

HOT\_PICKGROUP and HOTPICK are copies of the corresponding tables in the MTR.

RE\_CURRENT\_SESSION\_DATA holds all the data collected using the data collection methods. This data is written back to the MTR using data synchronization.

RE\_PROFILE\_DATA stores the historical profiles of active users. When a user is detected, the profile of that user is loaded from the MTR to this table.

RE\_RECOMMENDATION\_DETAIL is the source of data for the corresponding table in the MTR. The data is synchronized back to the MTR.

ATTR\_ID\_BIN\_BOUNDARY is a copy of the corresponding table in the MTR.

RE\_CONFIGURATION and RE\_INTERNAL\_CONFIGURATION store the configuration parameters for the RE.

RE\_DEPLOYABLE\_PACKAGE keeps track of the deployable package that is currently deployed in the RE.

RE\_LOG records events occurring in the RE.

RE\_ACTIVE\_USER stores information about all users who are currently active in the system. Data from this table is used to populate the session table in the MTR.

All other tables are used internally by the RE.

## Mining Object Repository

Much of the work that is done by OP uses MOR tables and views. The MOR includes the following tables (partial list):

MOR_VISITOR_TO_BROWSER_REPORT	TABLE
MOR_CONFIGURATION	TABLE
MOR_CROSS_SOLD_ITEMS_REPORT	TABLE
MOR_DEPLOYABLE_PACKAGE	TABLE
MOR_EFFECTIVENESS_REPORT	TABLE
MOR_EMAIL_ADDRESS	TABLE
MOR_ERROR_TABLE	TABLE
MOR_INTERNAL_CONFIGURATION	TABLE
MOR_MESSAGE_LOG	TABLE
MOR_MINING_MODEL	TABLE
MOR_MINING_RESULT	TABLE
MOR_MTR_CONNECTION	TABLE
MOR_RECOMMENDATION_ENGINE	TABLE
MOR_RECOMMENDATION_REPORT	TABLE
MOR_RECOMMENDATION_STRATEGY	TABLE
MOR_RE_FARM	TABLE
MOR_SCHEDULE_EVENT	TABLE
MOR_SCHEDULE_ITEM	TABLE
MOR_SCHEMA_ACCESS	TABLE
MOR_TAXONOMY_TRANS_CLOSURE	TABLE

MOR\_TRANS\_SUPERVISED\_RESULT

TABLE



---

## Initial Data Collection

Model building requires data. If you have data collected already, the data can be used to populate the MTR tables or the MTR schema can be mapped to the existing data via views. However, if you have no data, you might want to use the REAPI methods `addItem` and `addItems` to collect data. Data collection occurs in the Oracle9iAS Personalization Recommendation Engine (RE). For an RE to be up and running, there must be a package deployed in that RE. However, in order to build and deploy a package, you must have data in the MTR. To put it simply, you can't collect data unless you have enough data to build a package. You resolve this problem by populating the MTR with seed data and then using the seed data to build and deploy an initial package.

You create seed data on the system where Oracle9i is installed.

### Creating and Removing Seed Data

OP includes a script to populate an MTR with seed data and also another script to remove the seed data from the MTR once enough data is collected in the MTR. The following two scripts are in the directory `$ORACLE_HOME/install/dbscripts`:

- `insertMTRSeedRecords.sql`: Populates an unpopulated MTR with seed data.
- `removeMTRSeedRecords.sql`: Removes the seed data from the MTR once the initial model is built.

The OP administrator can run these scripts when appropriate.

## Preparing to Build Models with Seed Data

Follow these steps to populate the tables required for model building; many of these steps are described in *Getting Started with Oracle9iAS Personalization*:

1. Install OP with an unpopulated MTR.
2. Configure synchronization settings in the MTR\_CONFIGURATION table appropriately.
3. Create an MTR connection and allow synchronization for the data source types that you intend to collect.
4. Log in using the MTR account that you created during install, and run `insertMTRSeedRecords.sql`.
5. Create package using the default selections.
6. Create an RE Farm with at least one RE in it. Use the advanced settings to specify the customer data synchronization interval.
7. Build the package.
8. Deploy the package to the RE Farm that you created in step 6.
9. Logged in using the MTR account that you created during install, run `removeMTRSeedRecords.sql`.
10. Populate or map the following tables with data for your application:
  - **MTR\_ITEM**: This table is the site catalog. When the collected data is saved in the MTR, data corresponding to non-existent items are thrown away as part of the validation check. So, if the MTR\_ITEM table is empty, the collected data is thrown away.
  - **MTR\_CUSTOMER**: The customer-IDs are expected to be managed by the site. Hence the MTR\_CUSTOMER table should be populated or mapped with at least the ID attribute before the data is written to the MTR. The data values for other attributes can be collected via the data collection API.
  - **MTR\_ATTR\_NAME\_TO\_ID\_MAP** and **MTR\_BIN\_BOUNDARY**: These tables should be populated with the mapping information that maps attribute names to numeric identifiers and the bin boundaries, respectively.

You must specify bin boundaries for all data source types that you plan to collect.



11. After these tables are populated the OP administrator must run the following script, which populates the MTR\_ATTR\_ID\_BIN\_BOUNDARY table:

```
populate_mtr_attr_id_bin_boundary.sql
```

Now collect data from your web site. After you collect data for an appropriate amount of time, you can build models based on the collected data.



---

## Installation and Configuration Considerations

---

This appendix contains information about tuning and configuring OP for your installation. You should read this appendix in conjunction with Chapter 7.

### RE Installation and Configuration

Installation and configuration of a recommendation engine (RE) must be tailored to the expected number of active users that it will support. The RE in this context refers to a single engine on a single database instance. If multiple engines are installed on the same database instance, the configuration parameters will require adjustment.

Many factors go into the optimization of an RE. Some of these are set by the installation procedure, while others are techniques that may be used by the DBA. Configuration options fall into two broad categories: System availability and performance. System availability consists of settings required by the RE to handle user load without failure. Performance includes those that help maximize throughput.

The system availability settings for RE configuration are dependent on the number of anticipated users. Given this, it is possible to estimate the approximate system resource requirements and make database configuration recommendations. Since the REAPI maintains a connection pool of user connections, which can be reused, the number of required connections depends on how well user requests are being satisfied by the RE. That is, if for some reason there is a slowdown in the RE causing connection links to be held longer in the REAPI connection pool, the number of connections will tend to increase. As the number of connections increases, the number of actual database sessions increases. Each connection in the REAPI connection pool represents a database session.

The maximum number of connections in the REAPI connection pool is a configurable parameter in each RE. If this limit is exceeded, it may indicate that there are performance issues that need to be addressed other than simply increasing the size of the connection pool.

Each user’s client session results in database activity in the RE schema. As such, first configure the database to handle the number of anticipated users. Depending on the amount of available memory and CPUs in the system the RE database is installed on, it may be possible to support 50-100 users in a dedicated server environment. In this environment, each user connection to the database would require its own dedicated Oracle server for database access. As the number of users extends beyond 100, it may be more appropriate to use Oracle’s Multi-Threaded Server environment where database connections are pooled and serviced by shared database servers. The DBA responsible for the RE must decide whether the dedicated or shared server environment is used.

REAPI performance may be affected by several factors. On the client side, the REAPI runs in the JServer environment. Sufficient memory and CPU must be available to the client to handle the throughput for the active users. Communication with the RE from the REAPI clients is implemented through JDBC connections over Oracle’s SQL\*Net network. As the number of users grows, so does the demand on the network.

**RE Required Parameters**

The recommendation engine requires certain database parameters to be set to a minimum value, as follows:

```
JOB_QUEUE_PROCESSES=2
```

**Suggested RE Database Parameters (All Load Levels)**

The parameters listed below, while not necessary, are strongly recommended. Recommended sizes are also indicated:

BUFFER_POOL_KEEP	50 @ 8192 bytes ea
SORT_AREA_SIZE	819200 bytes
SORT_AREA_RETAINED_SIZE	819200 bytes

## Suggested RE Database Parameters (Variable Load Level)

The table below suggests guidelines for database configuration parameters based on number of projected users. The table shows, for a specified number of users, whether multi-threaded servers (MTS) are recommended, and recommended values for the number of MTS dispatchers, shared MTS servers, sessions, the size of large pool, and the size of shared pool:

Users	MTS	MTS Dispatchers	MTS Shared Servers	Sessions	Large Pool	Shared Pool
100	No	N/A	N/A	100	Default	Default
1000	Yes	4	20	300	250M	100M
2000	Yes	8	30	600	500M	120M
3000	Yes	10	40	900	1000M	140M
4000	Yes	12	50	1200	1500M	160M

## Suggested RE Configuration Parameters

The following table suggests guidelines for recommendation engine configuration settings based on the number of project users. These parameters are set in the RE schema table RE\_CONFIGURATION. The table shows, for a specified number of users, the recommended connection pool size and data synchronization interval:

Users	ConnectionPoolSize (number of connections)	DataSyncInterval (in seconds)
100	128	300
1000	256	300
2000	512	300
3000	1024	180
4000	2048	180

# MTR Installation and Configuration

The Mining Table Repository (MTR) database holds customer and product data. Data mining models are built based on this data. If the MTR has been configured by the administrator to permit data synching from the RE, data collected in the RE will be copied to the MTR at scheduled intervals. Customer profile data is also copied from the MTR into the RE when a customer begins a user session. All data transfer between the MTR and the RE is done using database links.

The following table offers guidelines for database configuration parameters based on the number of projected users. The table shows, for a specified number of users, whether multi-threaded servers (MTS) are recommended, and recommended values for the number of MTS dispatchers, the number of MTS servers, the number of sessions, and the size of the large pool:

Users	MTS	MTS Dispatchers	MTS Servers	Sessions	Large Pool
100	No	N/A	N/A	100	Default
1000	Yes	2	10	100	25M
2000	Yes	3	20	200	50M
3000	Yes	4	20	300	100M
4000	Yes	50	30	400	150M