

Oracle9iAS Personalization

Release Notes

Release 9.0.1

July 2001

Part No. A87537-01

These release notes describe Oracle9iAS Personalization, a new product with the Oracle9i Application Server (Oracle9iAS). This document is organized as follows:

- About Oracle9iAS Personalization 9.0.1, page 1.
- Where to Find More Information, page 2.
- For System Administrators: Installation and Administration, page 5.
- Bugs and Limitations, page 8.

See the README file on the CD for late-breaking information about the product, any new bugs, and platform-specific installation and administration information.

About Oracle9iAS Personalization 9.0.1

Oracle9iAS Personalization (OP) provides real-time personalization for Web sites using an integrated recommendation engine that is embedded in an Oracle database.

OP is based on data mining technology and modeling. It builds a predictive model of customer preferences using Web-based behavioral data collected by a Web site as well as demographic data.

The behavioral data it collects can include customers' historical and current navigation and preferences, both stated and implied. OP then builds a model from the data collected (and stored in Oracle database tables) and makes real-time personalized recommendations based on the model.



Oracle9iAS Personalization consists of the following components:

- **Recommendation Engines:** Recommendation engines collect and process data from the calling web application and return personalized recommendations to it.
- **Recommendation Engine Farms:** A collection of related recommendation engines. The recommendation engines in a farm make recommendations based on the currently deployed package.
- **Package:** An object created using the Administrative UI to schedule and configure model builds. Packages define the build settings and other attributes necessary to building data mining models.
- **Mining Object Repository (MOR):** Maintains mining metadata and mining model results as defined by the OP data mining schema; provides core mining algorithms and administrative functions for OP users.
- **Mining Table Repository (MTR):** Contains the fixed schema and data used for building packages
- **REAPI:** A collection of classes that enables a web application written in Java to collect and preprocess data that is used to build OP models and to obtain recommendations from OP (that is, to score using OP models).
- **RE Batch API:** A collection of classes that enables a web application written in Java to request Oracle9iAS Personalization-style recommendations in bulk mode.

Use the Administrative UI to manage Oracle9iAS Personalization functions such as creating and managing recommendation engines and farms, creating and managing packages, scheduling builds and deployments of packages, and scheduling reports. Use the REAPI to collect web-based customer data and obtain recommendations.

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 information and any late-breaking news about the product.
- *Oracle9iAS Personalization Release Notes*, Release 9.0.1 (this document).
- *Oracle9iAS Personalization Administrator's Guide*, Release 9.0.1 (includes installation instructions that are the same across all platforms).
- *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 the name of the system where OP database components are 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`.

To view the HTML files, you will need

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

Online Help The Oracle9iAS Personalization Administrative UI includes online help, available by clicking the Help button in the upper right corner of each page of the Administrative UI.

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 a 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.

For System Administrators: Installation and Administration

For technical support services, contact Customer Support at

<http://www.oracle.com/support/>

For information on installing OP, see the administrator's guide and the README file distributed with the OP documentation.

Installing OP on Two Servers

The recommended installation for OP is to install Oracle9iAS and certain OP components on one server, and install the Oracle9i database and other OP components on a different server .

- On the Oracle9iAS server, install the OP REAPI, REAPI Batch, and the REAPI demo.
- On the Oracle9i database server, install all other OP components, including the OP Administrative UI, the MTR, the MOR, the RE, and documentation (both online help and user manuals).

The single-system installation (less desirable from the performance perspective) is described in `<$ORACLE_HOME>/dmt/conf/oprocmgr/README.txt`.

Installing Additional REs, MORs, or MTRs

Installing an Additional RE: You cannot install a second recommendation engine (RE) on the same database using Oracle Universal Installer. You need to run the following script from the installation area with the required arguments:

```
$ORACLE_HOME/dmt/install/RE/createRE.sh DBAusername/DBApassword@TNSname  
REusername REpassword tableSpaceName dataFileDirectory dataFileSize  
sys_password hostname:port:SID
```

Installing an Additional MOR: You cannot install a second Mining Object Repository (MOR) on the same database using Oracle Universal Installer. Unlike the MTR and RE schemas, there is no script to add a second MOR on a database. The OP permits only one MOR on a database.

Installing an Additional MTR: You cannot install a second Mining Table Repository (MTR) on the same database using Oracle Universal Installer. You can create a second MTR that is not populated with data, as shown below:

```
$ORACLE_HOME/dmt/install/dbscripts/createMTR.sh  
DBAusername/DBApassword@TNSname MTRusername MTRpassword  
dataFileDirectory dataFileSize TablespaceName
```

Performance and Data Sizing

The table below shows the performance results on three datasets. The build time is linear in the number of customers and quadratic in the average profile size. We obtained the results shown here on a Sun Enterprise 450 (4 x UltraSPARC-II 400MHz) with 4096 megabytes of memory. Based on these numbers, it is possible to estimate the build time for any arbitrary dataset; see the subsection on Data Sizing, below. Definitions of the column headings follow the table.

Dataset	# of Cust	# of Items	Avg. Profile Size	Avg. Rating Profile	Avg. Purch. Profile	Avg. Nav. Profile	Build Time	# of Rules Agg. Model	# of Rules Cross-Sell Model
1	1000	50k	50	8	17	25	1 min	9152	155
2	5000	50k	50	8	27	36	3 min	166588	37
3	1000	50k	100	16	34	50	5 min	313154	2997

The column headings have the following meanings:

- **Dataset:** This is simply an identifier for the dataset.
- **# of Cust:** The number of registered customers, which is the number of records in the MTR_CUSTOMER table.
- **# of Items:** The number of items, which is the number of records in the MTR_ITEM table.
- **Avg. Profile Size:** The number of items in each customer's profile; this is the sum of average rating profile, average purchasing profile, and average navigation profile.
- **Avg. Rating Profile:** The average number of rating items in each customer's profile. Can be computed by dividing the number of records in MTR_RATING_DETAIL by the number of distinct CUSTOMER_ID's in MTR_RATING_DETAIL.
- **Avg. Purch. Profile:** The average number of purchasing items in each customer's profile. Can be computed by dividing the number of records in MTR_PURCHASING_DETAIL by the number of distinct CUSTOMER_ID's in MTR_PURCHASING_DETAIL.
- **Avg. Nav. Profile:** The average number of navigation items in each customer's profile. Can be computed by dividing the number of records in MTR_NAVIGATION_DETAIL by the number of distinct CUSTOMER_IDs in MTR_NAVIGATION_DETAIL.
- **Build Time:** The total time taken to build the aggregated model and the cross-sell model.

Data Sizing When you install OP, you are asked to specify temp space for the MOR. Use the following formula to calculate a rough estimate for the number of bytes of MOR temp space required:

$$(54 * P^2 * C) / 2$$

where P is the average profile size and C is the number of customers. Each item pair generated during a build takes 54 bytes, which is the space needed to store a record with two item columns and a count column. For the model build to run, the available temp space should be more than the estimated temp space requirement for the P and C values of the dataset.

Password Encryption and Decryption

Because of password encryption and decryption, there are two restrictions on passwords for OP users:

- OP users must not enter passwords with trailing blanks.
- User passwords are limited to 30 or fewer characters, the standard Oracle limit.

Bugs and Limitations

This section describes known bugs and limitations in Oracle9iAS Personalization release 9.0.1.

Correct Name of Product: The correct name of the product is "Oracle9iAS Personalization." In some instances the product may be incorrectly referred to as "Oracle9i Personalization."

OUI Should Verify Enough Hard Disk Space: OUI does not verify that there is enough hard disk space for data files.

OUI Should Check for Invalid Input: Certain errors in input are not caught, which may result in invalid results.

Default Deployment Selection Can Generate Error: The default selection for the deployment schedule is "deploy at a future time," which is typically within a few minutes of the current time. This is fine if the package has already been built, but it generates an error if the package has not yet been built. Workaround: Use the default deployment selection only if you know that the package has been built; otherwise select "deploy after every build."

Sorting by ID and Sorting by Type Do Not Work: For both sessionful and sessionless REAPIs, for recommendation content, sorting by ID and sorting by Type do not work.

Filter by Drop-Down List Should Display Only Applicable Selections: On the event log page, the drop-down list for "Filter by" displays selections that are not applicable to the particular event log page.

Deployment Times on RE Farms Page and Event Log Do Not Match: The time displayed for deployment should be the same (and is not) on the Recommendation Engine Farms page and the Event Log page.

Recommend_Top_Items_Batch Failure: Error message "ORA-00913 — Too Many Values" is displayed with `REProxyBatch.recommendTopItems()` if `TuningSettings.profileUsage` is set to "EXCLUDE". Workaround: None.

Resizing Window Results in Blank Screen and "Data Missing" Message: Running the Administrative UI with Netscape, if you resize the window, it disappears and you get a "Data Missing" message. Workaround: Instead of immediately returning the data when processing a POST, you can return a page that is a redirect to the real page you want to display. An alternative workaround: Use IE browser.

Corrupted Characters: The `DateFieldBean` javascript validation window displays broken Italian and Brazilian Portuguese characters when you enter an invalid date, for example, 11/12//01. Workaround: None.

Corrupted Messages in Message Log: Some messages appear in the message viewer without proper parameter substitution. For example, "The

Schema {0} is incorrect" should read, for example, "The Schema Scott is incorrect."

Schedule and Log Grids Lose Sort Order upon Application of Filter:

When "Filter-by" is applied, the sort order gets reset to the default column, requiring the user to re-apply the sort.

Result Tables for REAPI Batch: When a result table is created, its name is stored in the BATCH_RESULT_TABLE_NAME table (one such table per schema). This is done for ease of bookkeeping and cleanup.

Reports Need New Titles: The titles of reports are in some cases misleading.

