

Oracle® Healthcare Foundation

Quick Reference Guide for Informatica

Release 7.2.1

F11018-01

October 2018

This guide provides the configuration details and the process for data loading.

Prerequisite

Set the database properties according to the Oracle Healthcare Foundation configuration recommendations. For the configuration recommendations, see *Oracle Healthcare Foundation Configuration Guide*.

Terminology Loaders

Configuration File

Modify the HLI_ParameterFile.prm, in the */InfaSrcFiles* directory on the Informatica server, to configure the Terminology Loader.

Parameter Name	Description
VAR_SOURCE_FOLDER_NAME	Location of Terminology Loader data files.
VAR_ARCHIVE_FOLDER_NAME	Location where Terminology Loader data files are archived after successful data load.
PARALLEL_DEGREE	Parallel degree for the session level parallelism.

ETL Execution

Before executing the ETL, make sure that the Terminology Loaders source files are available in the source directory.

You can execute the Terminology Loader ETLs by using:

- **Scheduler:** For information, see [Executing Terminology Loader ETLs Using the Scheduler](#).
- **Informatica Workflow Manager:** For information, see [Executing Terminology Loader ETLs Using the Informatica Workflow Manager](#).

Executing Terminology Loader ETLs Using the Scheduler

You can schedule ETLs on any enterprise scheduler to match the frequency (daily, weekly, and so on) of the source data acquisition.

To schedule ETLs:

-
-
1. Connect to the Terminology Loaders Work Repository, and select **Designer**.
 2. Expand **EHA_HLI_Integration** project > **Global** > **Packages** > **MASTER_EXECUTION_PLAN** > **Scenarios**
 3. Expand the desired scenario.
 4. Right-click **Scheduling**, and select **New Scheduling** to build the schedule.

Executing Terminology Loader ETLs Using the Informatica Workflow Manager

Executing a Single Terminology Loader ETL

1. Open the Workflow Manager and connect to the Warehouse Integration Loader repository.
2. Expand the **HLI_MASTER_DATA_MANAGEMENT** folder > **Workflows**.
3. Right-click the workflow to execute, and select **Start Workflow**.

Executing all Terminology Loader ETLs

1. Open the Workflow Manager and connect to the Warehouse Integration Loader repository.
2. Expand the **HLI_MASTER_DATA_MANAGEMENT** folder > **Workflows**.
3. Right-click **wf_SIL_HLI_MASTER_LOAD**, and select **Start Workflow**.

System Monitoring and Error Handling

- ETL execution error logging:

When a particular source record fails during data processing, and gets rejected, the ETL logs this exception into the **HDI_HLI_ETL_EXCPTN_LOG_G** error table. Every exception logged has a message and description. If any session fails, check the Informatica session log.

- ETL execution error correction:

ETLs do not fail for any data issues if the file structure is maintained. If an ETL fails for any database related issues, correct the issue, note the step that failed in the **wf_SIL_HLI_MASTER_LOAD** workflow, and follow the appropriate instructions:

- Failed in the **HLI_Command** task - Check the **HLI_Archive.log** file in the **\$PMSourceFileDir** directory.
- Failed before or in the **HLI_Command** task - Run the **wf_SIL_HLI_MASTER_LOAD** workflow.

- Failed after the HLI archive - In the wf_SIL_HLI_MASTER_LOAD workflow, right-click **wklt_SIL_HLI_FILE_TO_STAGE_LOAD** worklet, and select **Start Workflow From Task**.
- Failed after the worklet wklt_SIL_HLI_FILE_TO_STAGE_LOAD - Run the failed ETL and remaining ETLs individually according to load order.

Warehouse Integration Loaders

The Warehouse Integration Loaders include a comprehensive set of ETLs that provide a data integration solution to process data from diverse source systems through the Interface Tables to the Healthcare Data Warehouse.

Prerequisites

Update Global Settings

Review and update the following configurations in the HDI_ETL_GLBL_PARAM_G table of the HMC schema:

Parameter Name	Description
ALLOW_INVALID_CODE	Controls the way data is loaded to the target when the source code is not resolved against HDM_CD_REPOSITORY. Yes - Loads the record to the target with an NAV value for the HDM code attributes. No - Rejects the record.
ALLOW_INVALID_REFERENCES	Controls the way data is loaded to the target when the source reference is not resolved against the HDM table. This is applicable for optional references only. Yes - Loads the record to the target with an NAV value for the HDM reference attributes. No - Rejects the record.
VERSIONING	Applicable for the incremental ETLs only when a changed record is from the source system. Yes - Creates a new version of the record. No - Overwrites the existing record.
RULE_IDS_TO_BE_REPROCESSED	List of Rule IDs, separated by commas, to be automatically reprocessed if the ETL fails. For example, -9998, -9993. NULL indicates no rules should be reprocessed.
MAX_REPROCESS_CNT	Number of times a record can be reprocessed before rejecting it. You must configure this in conjunction with RULE_IDS_TO_BE_REPROCESSED.
PARALLEL_DEGREE	Parallel degree for the session level parallelism.

Configure Rules and Terminology

- Review and enable or disable the default data validation rules in the HDI_ETL_RL_G table of the HMC schema.
- Review the seeded terminology standardization configurations in the HDI_ATRB_CD_SYS_LKUP_G and HDM_ATRB_CD_SYS_LKUP_G tables in the HDI and HDM schemas respectively.

You must disable the seeded configurations if the terminology standardization feature is not used or the seeded configurations are not relevant.

- Review the seeded terminology validation configurations in the HDI_ATRB_CD_TYP_LKUP_G and HDM_ATRB_CD_TYP_LKUP_G tables in the HDI and HDM schemas respectively.

You must disable the seeded configurations if the terminology validation feature is not used or the seeded configurations are not relevant.

- (Optional) Configure the late arriving data rules in the HDI_RL_REFRNTL_INTGRTY_G table in the HMC schema.

For more information, see *Oracle Healthcare Foundation Administrator's Guide*.

Schedule Load Plans

A load plan is a group of related workflows. The load plans created in Warehouse Integration Loader consists of multiple command tasks grouped together as per the dependencies, which executes the workflows. You can modify the load plans to disable, delete, insert, or modify the command task. When you modify the load plan, make sure you delete or add dependencies between workflows.

Load Plan	Description
wf_WIL_MDM_LOAD_PLAN	Includes the master data, rules, and configuration ETLs.
wf_WIL_INITIAL_LOAD_PLAN	Includes the HCD or CDM specific ETLs for the initial load in the Party Not Available mode.
wf_WIL_INCREMENTAL_LOAD_PLAN	Includes the HCD or CDM specific ETLs for the incremental load in the Party Not Available mode.
wf_WIL_INITIAL_LOAD_PLAN_PARTY_AVLBL	Includes the HCD or CDM specific ETLs for the initial load in the Party Available mode.
wf_WIL_INCREMENTAL_LOAD_PLAN_PARTY_AVLBL	Includes the HCD or CDM specific ETLs for the incremental load in the Party Available mode.

You can schedule ETLs on any enterprise scheduler to match the frequency (daily, weekly, and so on) of the source data acquisition.

Before executing or scheduling the load plan:

- Create the following environment variables, and associate them with the integration service.

- **PMUSER:** Repository user who executes the workflow.

- **PMPASS:** Encrypted repository user password.

Restart the integration service for the variables to take effect.

- Assign the integration service to load plans.

You can schedule any of the load plans as per your execution frequency.

Setting Up the Environment for the First ETL Execution

- Gather statistics on all Interface Tables.
- Disable the automatic statistics gathering option on the HDM schema.
- Delete and lock the statistics on the HDM schema.

- Set the parallel degree in the HDI_ETL_GLBL_PARAM_G table.
- Review and modify the global parameters.

Executing Master Data Management ETLs

The Master Data Management (MDM) ETLs load data to the MDM tables, such as User, Data Source, Code System, Code Repository, Cross Map, and other Code related tables. The MDM ETLs are grouped under the DI_MASTER_DATA_MANAGEMENT folder in the Warehouse Integration Loader repository.

You can choose to run the MDM ETLs as needed only when there is a change.

Executing Rules and Configurations ETLs

You can create late arriving data rules, data validation rules, and terminology configurations when they are needed, and run the ETLs accordingly. These ETLs are grouped under the DI_MASTER_DATA_MANAGEMENT folder in the Warehouse Integration Loader repository.

Executing the Initial Load ETLs for Transaction Entities

The initial load ETLs process historical data and perform data loading in an optimal way.

Setting Up the Environment

1. Drop the foreign keys according to the guidelines in the *Oracle Healthcare Foundation Administrator's Guide*.
2. Make sure that the HDI data load is complete.
3. Make sure that the source data does not contain multiple versions of the same record. The initial load ETLs cannot handle multiple versions.
4. Process any additional rules and configurations, and load the master data tables to the HDM tables.

Executing ETLs

When you use an enterprise scheduler, ETLs start automatically at the scheduled time or after the HDI data load is complete. You can monitor the progress of ETLs through the PowerCenter Workflow Monitor.

Restarting Failed ETLs

When an ETL fails during the data loading process:

1. Clean up the partially loaded data manually.
 - When a target table (for example, HDM_ENC) is loaded by a single source table (for example, HDI_ENC), use the following query to clean up the partially loaded HDM data.

For example,

```
DELETE FROM HDM_ENC WHERE REQUEST_ID= (SELECT REQUEST_ID FROM
HMC.HDI_ETL_LOAD_DT_G WHERE WRK_FLOW_NM= 'wf_SIL_DI_HDI_ENC' );
```

COMMIT;

- When a target table (for example, HDM_PRTY) is loaded by multiple sources (for example, HDI_PT, HDI_SVCPRV, and so on), delete the partially loaded data of the failed ETL. Assuming that the HDI_PT ETL has processed the data to the HDM tables but the HDI_SVCPRV ETL failed while processing data to HDM_PRTY, delete the partially data loaded by the HDI_SVCPRV ETL.

For example,

```
DELETE FROM HDM_PRTY WHERE INTEGRATION_ID LIKE 'SVCPRV%' AND  
REQUEST_ID= (SELECT REQUEST_ID FROM HMC.HDI_ETL_LOAD_DT_G WHERE  
WRK_FLOW_NM='wf_SIL_DI_HDI_SVCPRV');
```

COMMIT;

2. Restart the ETL.

Each workflow consists of following structure:

```
s_SIL_DI_Create_Rule_Param_File [Start]  
  3 tasks to capture and update load summary details  
s_SIL_DI_<HDI table>_SQLVALID  
  3 tasks to capture and update load summary details  
s_SIL_DI_<HDI Table>_<HDM Table> [Master session]  
-> Point 1  
  3 tasks to capture and update load summary details  
s_SIL_DI_<HDI Table>_<HDM Table> [Child session 1]  
-> Point 2  
  3 tasks to capture and update load summary details  
s_SIL_DI_<HDI Table>_<HDM Table> [Child session 2]  
-> Point 3  
  3 tasks to capture and update load summary details  
s_SIL_DI_HDI_ETL_LOAD_DT_G [Load Date]  
-> Point 4  
  3 tasks to capture and update load summary details  
s_SIL_DI_HDI_ETL_EXCPTN_LOG_TMP [Exception]  
-> Point 5 [End]  
  3 tasks to capture and update load summary details
```

When ETL fails...	What to do:
Between Start and Point 1	Restart the workflow from the beginning.
Between Child Session 1 and Point 2	Restart the workflow from the Child Session 1.
Between Child Session 2 and Point 3	Restart the workflow from the Child Session 2.
Between Load Date and Point 4	Restart the workflow from the Load Date.
Between Exception and End	Restart the workflow from the Exception.

When ETL fails...	What to do:
At assignment or command task	<p>Create a parameter file under the \$PMSourceFileDir/\$PMFolderName/ directory, and restart the workflow from the subsequent session that updates the load summary, or directly update the job details in the JB_EXCTN table.</p> <p>The following is the sample parameter file (when a_enc_hdm_enc or c_enc_hdm_enc fails):</p> <pre>Name: s_enc_hdm_enc_stat.prm [Global] \$\$workflow_name=wf_SIL_DI_HDI_ENC \$\$session_name=s_SIL_DI_HDI_ENC_HDM_ENC \$\$session_status=SUCCEEDED \$\$source_success_count=15 \$\$target_success_count=12</pre> <p>where, \$\$source_success_count and \$\$target_success_count are the sum of records across all source and target instances, and read from Workflow Monitor for the session under consideration.</p>

- , Create a parameter file under the **\$PMSourceFileDir/\$PMFolderName/** directory, and restart the workflow from the subsequent session that updates the load summary, or directly update the job details in the JB_EXCTN table.

The following is the sample parameter file (when a_enc_hdm_enc or c_enc_hdm_enc fails):

```
Name: s_enc_hdm_enc_stat.prm
[Global]
$$workflow_name=wf_SIL_DI_HDI_ENC
$$session_name=s_SIL_DI_HDI_ENC_HDM_ENC
$$session_status=SUCCEEDED
$$source_success_count=15
$$target_success_count=12
```

where, \$\$source_success_count and \$\$target_success_count are the sum of records across all source and target instances, and read from Workflow Monitor for the session under consideration.

For multiple sessions at the same level, restart the ETL from the level task from where the multiple sessions originate.

Executing the Incremental Load ETLs for the Transaction Entities

The incremental ETLs are grouped under the DI_HDWF_INCREMENTAL_LOAD and DI_HDWF_INCREMENTAL_LOAD_PARTY_AVAILABLE folders in the Warehouse Integration Loaders repository.

Setting Up the Environment

- For the first incremental load, create the foreign key indexes on the HDM table according to guidelines in the *Oracle Healthcare Foundation Administrator's Guide*.
- Schedule an automatic statistics gathering for the HDM and HDI schema using the stale option.
- Make sure that the HDI data load is complete.
- Process necessary rules and configurations, and load the master data tables to the HDM tables.

Running ETLs

When you use an enterprise scheduler, ETLs start automatically at the scheduled time or after the HDI data load is complete. You can monitor the progress of ETLs through the PowerCenter Workflow Monitor.

Restarting Failed ETLs

1. Correct the partially loaded data.

When an incremental ETL fails, correct the partially loaded data in the HDM schema before restarting the ETL. The Warehouse Integration Loaders provide a common ETL, `wf_m_SIL_DI_Restartability`, which corrects the HDM data. Before running this ETL, configure the failed session in the `HDI_ETL_GLBL_PARAM_G` table under the parameter `RESTART_SESSION`.

Sample Update Command:

```
UPDATE <HMC SCHEMA NAME>.HDI_ETL_GLBL_PARAM_G SET PARAM_VAL = '<Failed Session Name>' WHERE PARAM_NM='RESTART_SESSION';  
  
COMMIT;
```

When multiple sessions fail, correct the data of each failed HDM table by running the `wf_m_SIL_DI_Restartability` ETL for each failed session.

Only master and child sessions loads data to the HDM tables, and other sessions are used for tracking date, processing rules, handling exceptions, and so on.

2. Restart the ETL

Each workflow consists of following structure:

```
s_SIL_DI_Create_Rule_Param_File [Start]  
    3 tasks to capture and update load summary details  
s_SIL_DI_<HDI table>_SQLVALID  
    3 tasks to capture and update load summary details  
s_SIL_DI_<HDI Table>_<HDM Table> [Master session]  
-> Point 1  
    3 tasks to capture and update load summary details  
s_SIL_DI_<HDI Table>_<HDM Table> [Child session 1]  
-> Point 2  
    3 tasks to capture and update load summary details  
s_SIL_DI_<HDI Table>_<HDM Table> [Child session 2]  
-> Point 3  
    3 tasks to capture and update load summary details  
s_SIL_DI_HDI_ETL_LOAD_DT_G [Load Date]  
-> Point 4  
    3 tasks to capture and update load summary details  
s_SIL_DI_HDI_ETL_EXCPTN_LOG_TMP [Exception]  
-> Point 5 [End]  
    3 tasks to capture and update load summary details
```

When the ETL fails:

- Between Exception and End: Restart the workflow from the Exception.

For multiple sessions at the same level, restart the ETL from the level task from where the multiple sessions originate.

Monitoring and Logging

You can monitor the progress of the ETL execution through the PowerCenter Workflow Monitor.

Load Summary

The Warehouse Integration Loaders capture the ETL execution summary in the JB table of the enterprise schema and the EXCPTN table.

The execution statistics includes component (Warehouse Integration Loaders and Application Toolkit), job group name (package name), job name (interface name), job type (initial or incremental), job start date, job end data, job status (running, succeeded, or failed), source record count, and target record count.

The exception statistics includes exception type (reject, reprocess, or suspend) and the exception count.

ETL Exceptions

The ETL logs any record that fails for mandatory rules, data validation rules, late arriving data rules, or data quality rules during the data loading process, in the HDI_ETL_EXCPTN_LOG_TMP table in the HMC schema.

Any reprocessing rules, the records that can be reprocessed in the subsequent loads are retained in the HDI_ETL_EXCPTN_LOG_TMP table, and all other exceptions are moved to the HDI_ETL_EXCPTN_LOG_G table.

Application Toolkit Loaders

The Application Toolkit Loaders include a comprehensive set of ETLs that provide a data integration solution to process data from Healthcare Data Warehouse to Healthcare Common Data mart (HCD).

Updating Configurations

- Review and update the configurations in the HCD_ETL_ENTY_SELCTN_PARAM_G table of the HMC schema as needed.

Parameter Name	Description
MSTR_ENTY_NM	Entity which uses the parameterized attribute value.
ATTRIB_NM	Name of the attribute to be parameterized.
ATTRIB_VAL	Value of the parameterized attribute.

Parameter Name	Description
ENTERPRISE_ID	Unique identifier for an enterprise in a multi-enterprise environment.

- Set the degree of parallelism in the HCD_GLBL_PARAM_G table in the HMC schema to parallelize the SQL executions.

Parameter Name	Default Value	Description
PARALLEL_QUERY_NO	1	This parameter is used for performance enhancement for the initial load ETLs. Set this value to an exponential value of 2, that is, 1, 2, 4, 8, 16, and so on. For Exadata environment, set this value to 16 or 32. For non-Exadata environment, set this value to 2 or 4.
HDMF_SCHEMA_NAME	HDM	Specifies the Healthcare Data Warehouse schema name.
HMC_SCHEMA_NAME	HMC	Specifies the ETL configuration schema name.
PROXY_USER_ENABLED	NO	Indicates whether the Application Toolkit ETLs have to be executed through a proxy user. Update this parameter to YES if you want ETLs to be executed through a proxy user.
HCD_SCHEMA_NAME	HCD	Specifies the HCD schema name.

Scheduling Workflows

Identify the list of HCD tables and the corresponding ETLs that process data as needed. Follow an optimal load sequence considering the ETL dependencies, data volume, and resource availability.

You can schedule ETLs on any enterprise scheduler to match the frequency (daily, weekly, and so on) of the source data acquisition.

You can create multiple ETL groups based on the frequency of data changes to be processed in a sequence.

- ETL - Initial
- ETL - Incremental

Scheduling Load Plans

The load plan created in HCD consists of multiple command tasks grouped together as per the dependencies, which executes the workflows. You can modify the load plans to disable, delete, insert, or modify the command task. When you modify, make sure you delete or add dependencies between workflows.

Load Plan	Description
wf_HCD_HIERARCHY_LOAD_PLAN	Includes the hierarchy ETLs.
wf_HCD_INITIAL_LOAD_PLAN	Includes the HCD ETLs for the initial or bulk load.
wf_HCD_INCREMENTAL_LOAD_PLAN	Includes the HCD ETLs for the incremental load.

Before executing or scheduling the load plan:

- Create the following environment variables with the help of an administrator, and associate them with the integration service:

- **PMUSER:** Repository user who executes the workflow.
- **PMPASS:** Encrypted repository user password.

Restart the integration service for the variables to take effect.

- Assign the integration service to load plans.

You can schedule any of the load plans as per your execution frequency.

Executing ETLs Using Informatica Workflow Manager

1. Open the Workflow Manager and connect to the HCD Loader repository.
2. Expand the required folder.
3. Expand **Workflows**.
4. Right-click the workflow to execute, and select **Start Workflow**.

Setting Up the Environment for the First ETL Execution

- Gather statistics on all HDM tables.
- Disable the automatic statistics gathering option on the HDM schema.
- Delete and lock the statistics on the HCD schema.
- Set the parallel degree in the HDI_ETL_GLBL_PARAM_G table.
- Review and modify the global parameters.
- Create a folder **HCD_HIERCHARY_LOAD** under the <infa_home>/server/infa_shared/SrcFiles/, and provide read and write permissions by executing the `chmod 755` command.

Executing Initial Load ETLs

The initial load ETLs process historical data and perform data loading in an optimal way.

Setting Up the Environment

- Drop the foreign keys according to guidelines in the *Oracle Healthcare Foundation Administrator's Guide*.
- Make sure that the HDM data load is complete.
- Make sure that the selection or inline code configurations are valid in the HCD_ETL_ENTY_SELCTN_PARAM_G entity.

Running ETLs

When you use an enterprise scheduler, ETLs start automatically at the scheduled time or after the HDM data load is complete. You can monitor the progress of ETLs through the PowerCenter Workflow Monitor.

Restarting failed ETLs

When an ETL fails during the data loading process, rerun the workflow from the failed session.

Executing Incremental Load ETLs

The incremental ETLs process the data acquired on a regular basis.

Setting Up the Environment

1. For the first incremental load, create the foreign key indexes on the HDM table according to guidelines in the *Oracle Healthcare Foundation Administrator's Guide*.
2. Schedule an automatic statistics gathering for the HDM and HCD schema using the stale option.
3. Make sure that the HDM data load is complete.

Executing ETLs

When you use an enterprise scheduler, ETLs start automatically at the scheduled time or after the HDM data load is complete. You can monitor the progress of ETLs through the PowerCenter Workflow Monitor.

Restarting Failed ETLs

When an ETL fails during the data loading process, rerun the workflow from the failed session.

Monitoring and Logging

You can monitor the progress of the ETL execution through the PowerCenter Workflow Monitor.

For the load summary details, see [Load Summary](#).

Cohort Data Mart Loaders

Updating Configurations

- Review and update the configurations in the C_LOAD_PARAM table of the HDM schema as needed.
- Set the degree of parallelism in the C_LOAD_CONFIG table in the CDM schema to parallelize the SQL executions.

Scheduling Workflows

Use the following workflows in the CDM_ETL folder to load data into CDM:

- wf_CDM_INITIAL_LOAD_PLAN - Use this workflow in the scheduler for full or initial load execution.

- wf_CDM_INCREMENTAL_LOAD_PLAN - Use this workflow in the scheduler for incremental load execution.

Executing ETLs Using Informatica Workflow Manager

1. Open the Workflow Manager and connect to the CDM Loader repository.
2. Expand the CDM_ETL folder > **Workflows**.
3. Right-click the workflow to execute, and select **Start Workflow**.

Setting Up the Environment for the First ETL Execution

- Gather statistics on all HDM tables.
- Review and modify the global parameters.

Executing the Initial Load ETLs

CDM has a high level workflow (wf_CDM_INITIAL_LOAD_PLAN) to load data in the initial load.

Setting Up the Environment

Make sure that the HDM data load is complete.

Executing ETLs

When you use an enterprise scheduler, ETLs start automatically at the scheduled time or after the HDM data load is complete. You can monitor the progress of ETLs through the PowerCenter Workflow Monitor.

Note: When upgrading to OHF 7.2.1 or later, you must truncate and reload the CDM tables for which the ETL is optimized.

Restarting Failed ETLs

When an ETL fails during the data loading process, rerun the workflow from the failed job.

Executing Incremental Load ETLs

CDM has a high level workflow (wf_CDM_INCREMENTAL_LOAD_PLAN) for incremental ETLs to process the data acquired on a regular basis.

Setting Up the Environment

Ensure that the HDM data load is complete.

Executing ETLs

When you use an enterprise scheduler, ETLs start automatically at the scheduled time or after the HDM data load is complete. You can monitor the progress of ETLs through the PowerCenter Workflow Monitor.

Restarting Failed ETLs

When an ETL fails during the data loading process, rerun the workflow from the failed job.

Monitoring and Logging

You can monitor the progress of the ETL execution through the PowerCenter Workflow Monitor.

Documentation Accessibility

For information about Oracle's commitment to accessibility, visit the Oracle Accessibility Program website at

<http://www.oracle.com/pls/topic/lookup?ctx=acc&id=docacc>.

Access to Oracle Support

Oracle customers that have purchased support have access to electronic support through My Oracle Support. For information, visit

<http://www.oracle.com/pls/topic/lookup?ctx=acc&id=info> or visit

<http://www.oracle.com/pls/topic/lookup?ctx=acc&id=trs> if you are hearing impaired.

Oracle Healthcare Foundation Quick Reference Guide for Informatica, Release 7.2.1
F11018-01

Copyright © 2018, Oracle and/or its affiliates. All rights reserved.

This software and related documentation are provided under a license agreement containing restrictions on use and disclosure and are protected by intellectual property laws. Except as expressly permitted in your license agreement or allowed by law, you may not use, copy, reproduce, translate, broadcast, modify, license, transmit, distribute, exhibit, perform, publish, or display any part, in any form, or by any means. Reverse engineering, disassembly, or decompilation of this software, unless required by law for interoperability, is prohibited.

The information contained herein is subject to change without notice and is not warranted to be error-free. If you find any errors, please report them to us in writing.

If this is software or related documentation that is delivered to the U.S. Government or anyone licensing it on behalf of the U.S. Government, then the following notice is applicable:

U.S. GOVERNMENT END USERS: Oracle programs, including any operating system, integrated software, any programs installed on the hardware, and/or documentation, delivered to U.S. Government end users are "commercial computer software" pursuant to the applicable Federal Acquisition Regulation and agency-specific supplemental regulations. As such, use, duplication, disclosure, modification, and adaptation of the programs, including any operating system, integrated software, any programs installed on the hardware, and/or documentation, shall be subject to license terms and license restrictions applicable to the programs. No other rights are granted to the U.S. Government.

This software or hardware is developed for general use in a variety of information management applications. It is not developed or intended for use in any inherently dangerous applications, including applications that may create a risk of personal injury. If you use this software or hardware in dangerous applications, then you shall be responsible to take all appropriate fail-safe, backup, redundancy, and other measures to ensure its safe use. Oracle Corporation and its affiliates disclaim any liability for any damages caused by use of this software or hardware in dangerous applications.

Oracle and Java are registered trademarks of Oracle and/or its affiliates. Other names may be trademarks of their respective owners.

Intel and Intel Xeon are trademarks or registered trademarks of Intel Corporation. All SPARC trademarks are used under license and are trademarks or registered trademarks of SPARC International, Inc. AMD, Opteron, the AMD logo, and the AMD Opteron logo are trademarks or registered trademarks of Advanced Micro Devices. UNIX is a registered trademark of The Open Group.

This software or hardware and documentation may provide access to or information about content, products, and services from third parties. Oracle Corporation and its affiliates are not responsible for and expressly disclaim all warranties of any kind with respect to third-party content, products, and services unless otherwise set forth in an applicable agreement between you and Oracle. Oracle Corporation and its affiliates will not be responsible for any loss, costs, or damages incurred due to your access to or use of third-party content, products, or services, except as set forth in an applicable agreement between you and Oracle.