

**Oracle® Retail Active Retail Intelligence**  
Operations Guide  
Release 13.0

April 2008

Copyright © 2003, 2008, Oracle. All rights reserved.

Primary Author: Archana Kishore

The Programs (which include both the software and documentation) contain proprietary information; 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, except to the extent required to obtain interoperability with other independently created software or as specified by law, 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. This document is not warranted to be 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.

If the Programs are delivered to the United States Government or anyone licensing or using the Programs on behalf of the United States Government, the following notice is applicable:

U.S. GOVERNMENT RIGHTS Programs, software, databases, and related documentation and technical data delivered to U.S. Government customers are "commercial computer software" or "commercial technical data" pursuant to the applicable Federal Acquisition Regulation and agency-specific supplemental regulations. As such, use, duplication, disclosure, modification, and adaptation of the Programs, including documentation and technical data, shall be subject to the licensing restrictions set forth in the applicable Oracle license agreement, and, to the extent applicable, the additional rights set forth 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 we disclaim liability for any damages caused by such use of the Programs.

Oracle, JD Edwards, PeopleSoft, and Siebel are registered trademarks of Oracle Corporation and/or its affiliates. Other names may be trademarks of their respective owners.

The Programs may provide links to Web sites and access to content, products, and services from third parties. Oracle is not responsible for the availability of, or any content provided on, third-party Web sites. You bear all risks associated with the use of such content. If you choose to purchase any products or services from a third party, the relationship is directly between you and the third party. Oracle is not responsible for: (a) the quality of third-party products or services; or (b) fulfilling any of the terms of the agreement with the third party, including delivery of products or services and warranty obligations related to purchased products or services. Oracle is not responsible for any loss or damage of any sort that you may incur from dealing with any third party.

## Value-Added Reseller (VAR) Language

- (i) the software component known as **ACUMATE** developed and licensed by Lucent Technologies Inc. of Murray Hill, New Jersey, to Oracle and imbedded in the Oracle Retail Predictive Application Server – Enterprise Engine, Oracle Retail Category Management, Oracle Retail Item Planning, Oracle Retail Merchandise Financial Planning, Oracle Retail Advanced Inventory Planning and Oracle Retail Demand Forecasting applications.
- (ii) the **MicroStrategy** Components developed and licensed by MicroStrategy Services Corporation (MicroStrategy) of McLean, Virginia to Oracle and imbedded in the MicroStrategy for Oracle Retail Data Warehouse and MicroStrategy for Oracle Retail Planning & Optimization applications.
- (iii) the **SeeBeyond** component developed and licensed by Sun Microsystems, Inc. (Sun) of Santa Clara, California, to Oracle and imbedded in the Oracle Retail Integration Bus application.
- (iv) the **Wavelink** component developed and licensed by Wavelink Corporation (Wavelink) of Kirkland, Washington, to Oracle and imbedded in Oracle Retail Store Inventory Management.
- (v) the software component known as **Crystal Enterprise Professional and/or Crystal Reports Professional** licensed by Business Objects Software Limited (“Business Objects”) and imbedded in Oracle Retail Store Inventory Management.
- (vi) the software component known as **Access Via**<sup>TM</sup> licensed by Access Via of Seattle, Washington, and imbedded in Oracle Retail Signs and Oracle Retail Labels and Tags.
- (vii) the software component known as **Adobe Flex**<sup>TM</sup> licensed by Adobe Systems Incorporated of San Jose, California, and imbedded in Oracle Retail Promotion Planning & Optimization application.
- (viii) the software component known as **Style Report**<sup>TM</sup> developed and licensed by InetSoft Technology Corp. of Piscataway, New Jersey, to Oracle and imbedded in the Oracle Retail Value Chain Collaboration application.
- (ix) the software component known as **WebLogic**<sup>TM</sup> developed and licensed by BEA Systems, Inc. of San Jose, California, to Oracle and imbedded in the Oracle Retail Value Chain Collaboration application.
- (x) the software component known as **DataBeacon**<sup>TM</sup> developed and licensed by Cognos Incorporated of Ottawa, Ontario, Canada, to Oracle and imbedded in the Oracle Retail Value Chain Collaboration application.



---

---

# Contents

<b>Preface .....</b>	<b>ix</b>
Audience .....	ix
Related Documents.....	ix
Customer Support.....	ix
<b>1 Introduction .....</b>	<b>1</b>
How and When to Use this Guide.....	1
<b>2 Process Overview .....</b>	<b>3</b>
ARI Shutdown.....	3
Metadata Modification.....	3
Rule Construction and Modification.....	4
Import Export Tool (IET) .....	4
Code Generation .....	4
ARI Start.....	5
Summary .....	6
<b>3 Process Details .....</b>	<b>7</b>
ARI Logs.....	7
Understanding Log Types .....	7
Setting the Log Level.....	7
Reviewing the Log.....	7
Purging the Log.....	7
DBA_JOBS Queue.....	7
Scheduler.....	8
Exception Validation Engine (EVE).....	8
Code Generation .....	8
History Purge .....	9
Administrator Groups.....	9
<b>4 Product Integration.....</b>	<b>11</b>
Metadata.....	11
Oracle Retail Merchandise System (RMS).....	11
Multilanguage Support .....	12
Key ARI Tables Related to Internationalization .....	12
Presentation Interface .....	13
Oracle Retail Data Warehouse .....	14
RDW Report Setup .....	14
Metadata Setup .....	14
Integration with Oracle Retail Workspace .....	14
<b>A Appendix: Architectural Reference.....</b>	<b>17</b>
Analyst Process Definition.....	17
1. Set Up and Maintain Metadata .....	17

2. User/Group Maintenance .....	18
3. Schedule Maintenance.....	18
4. Exception Type Maintenance .....	18
5. Event Type Maintenance .....	18
System Process Management .....	19
1. ARI Daily Management .....	19
2. Table Manager and Queue Process Builder .....	19
3. Trigger Builder .....	20
4. Scan Code Builder.....	20
5. Validation Procedure Generator .....	20
6. Evaluation Procedure Generator .....	20
Exception Candidate Detection.....	21
1. Real-Time Monitor.....	22
2. Periodic Monitor .....	22
3. Scanner Builder .....	22
4. Data Warehouse Monitoring.....	22
5. External API Monitor .....	22
6. Trickle Data Monitor .....	22
Candidate Validation and Event Creation .....	23
1. Validate Engine .....	23
2. Realm Queue Process .....	24
3. Exception Creation and Validation .....	24
4. Event Creation and Evaluation.....	24
User-Initiated and Automated Event Resolution .....	25
1. Alert Viewer .....	25
2. Event Management.....	26
3. Exception Revalidation .....	26
4. Event Reevaluation.....	26
5. Schedule Reevaluation.....	26
6. Action Execution.....	26
Single Sign-on Overview .....	27
What is Single Sign-On?.....	27
What Do I Need for Oracle Single Sign-On?.....	27
Can Oracle Single Sign-On Work with Other SSO Implementations? .....	27
Oracle Single Sign-on Terms and Definitions.....	27
What Single Sign-On is not.....	29
How Oracle Single Sign-On Works.....	29
Installation Overview.....	31
User Management.....	32
Setting up ARI for Single Sign-on.....	32
Configuring formsweb.cfg.....	32
Creating a RAD Entry .....	33

<b>B Appendix: API List</b> .....	<b>34</b>
Error and Activity Logging .....	34
Changing the Log Level.....	34
Stopping and Starting the Backend .....	34
Starting the Master Processes.....	34
Starting EVE Only.....	34
Stopping the Master Processes.....	34
Stopping EVE Only.....	35
Stopping All ARI Processes .....	35
Code Generation .....	35
Scheduler.....	35
Starting the Scheduler .....	36
Stopping the Scheduler .....	36
Signal-Driven Scheduler Signaler.....	36
EVE (Exception Validation Engine).....	36
Starting EVE .....	36
Stopping EVE .....	36
Periodic Purges.....	36
Event Purge .....	37
Event History Purge .....	37
ARI Alert Notification API.....	37
End User Cases.....	38
Architecture .....	38
Implementation.....	38
<b>C Appendix: ARI Options</b> .....	<b>40</b>
EVE_NUM_THREADS .....	40
EVE_QUEUE_REFRESH_INTERVAL .....	40
INTERNAL_SCHEMA.....	40
MASTER_SCHEMA .....	40
MAX_EVENT_RECURSION.....	40
REEVAL_STATUS_LOCKOUT .....	40
PRIMARY_LANGUAGE_NUMBER.....	41
ANALYST_ADMIN_GROUP_ID.....	41
CLOSE_EVENT_REALM_ID .....	41
CLOSE_EVENT_REALM_ID .....	41
ERROR_ADMIN_GROUP_ID.....	41
EVENT_INSTANCE_PARM_ID.....	41
EXCEPTION_CREATE_DATE_PARM_ID .....	41
EXCEPTION_CREATE_DATE_USER_ID.....	41
LOG_LEVEL.....	41
FORWARD_GENERATION_HOURS .....	41
RDW_LINK.....	42

RDW_OWNER.....	42
RDW_PREFIX.....	42

---

---

# Preface

Oracle Retail Operations Guides are designed so that you can view and understand the application's 'behind-the-scenes' processing. Including such information as the following:

- Key system administration configuration settings
- Functional integration dataflow across the enterprise

## Audience

Anyone who has an interest in better understanding the inner workings of the Active Retail Intelligence system can find valuable information in this guide. There are three audiences in general for whom this guide is written:

- System analysts and system operation personnel:
  - Who are looking for information about Active Retail Intelligence processes internally or in relation to the systems across the enterprise
  - Who operate Active Retail Intelligence on a regular basis
- Integrators and implementation staff who have the overall responsibility for implementing Active Retail Intelligence into their enterprise
- Business analysts who are looking for information about processes and interfaces to validate the support for business scenarios within Active Retail Intelligence and other systems across the enterprise

## Related Documents

For more information, see the following documents in the Oracle Retail Active Retail Intelligence Release 13.0 documentation set:

- Oracle Retail Active Retail Intelligence Installation Guide
- Oracle Retail Active Retail Intelligence Release Notes
- Oracle Retail Active Retail Intelligence User Guide
- Oracle Retail Active Retail Intelligence Online Help
- Oracle Retail Merchandising Implementation Guide

## Customer Support

<https://metalink.oracle.com>

When contacting Customer Support, please provide the following:

- Product version and program/module name
- Functional and technical description of the problem (include business impact)
- Detailed step-by-step instructions to re-create
- Exact error message received



---

---

# Introduction

This guide is designed to explain the administrative processes that support the operation of Oracle Retail Active Retail Intelligence (ARI).

## How and When to Use this Guide

None of the processes described here should be undertaken until after ARI installation is complete. Appropriate administration of ARI is critical to its successful operation, and although many of the tasks, once in a production, are primarily the responsibility of the database administrator, significant collaboration between the database administrator and business analyst is required throughout the rule definition process. Database administrators should read and understand this guide after installation and before doing any additional ARI work. Business analysts should read and understand this guide before creating any ARI rules.

One of the first things to do after reading and understanding this guide is to make sure the ARI Options table values are all set correctly (see Appendix C for details of the options).



---

---

## Process Overview

As discussed in the Overview sections of the online help, ARI is only a tool that helps with several steps of a larger process, which, for context, is reviewed here. This larger process, to implement a new business process, or existing process modification, involves first gathering requirements and conducting user walkthroughs, designing possible implementation strategies, reviewing impact, choosing a strategy and conducting additional walkthroughs. This is followed by development of rules, test implementations and user acceptance. Finally, the rules are moved into production. The sub-process relevant to ARI administration is essentially the same in development, test and production.

More than a single, specific process, the sub-process relevant to database administrators is a set of processes that can be linked in a number of different ways depending on the operational environment. Factors such as, whether the database is shut down regularly or not, and how the shutdown occurs (shutdown immediate or otherwise), will impact exactly how these sub-processes should be implemented.

### ARI Shutdown

ARI uses the DBA\_JOBS queue to control the background processes essential to exception detection and validation. ARI also uses generated code in these detection processes. (A more detailed architectural overview can be found in Appendix A.) Before making *any* DDL changes in any ARI schema or in any schema monitored by ARI, including recompiling packages, alter tables, creating new functions, or even running the ARI code generator, it is critical that the Exception Validation Engine (EVE) be halted from processing exception candidates. (Stopping EVE can be done using one of the many stop methods described along with the other APIs in Appendix B.)

EVE uses persistent variables to control multi-threading and other aspects of its operation. Because of this, EVE should be stopped before bringing down the database and restarted when the database is brought back up. If EVE is not stopped, either due to an unexpected shutdown or shutdown before EVE is cleared from the DBA\_JOBS queue (EVE may take a few minutes to finish processing after the shutdown script is run), the shut down request should be sent immediately after the database is restarted. Once it has cleared the DBA\_JOBS queue, EVE can then be restarted.

### Metadata Modification

As a rule, metadata should be synchronized with the actual DDL and other objects (Oracle Forms, for example) it describes at all times. Because such instantaneous synchronization is essentially impossible, the next best thing is that these tasks are performed in immediate sequence. This is not terribly difficult since the DDL and Forms should change very little during production, and when they do, non-administrative users usually must be logged out while the DDL changes are taking place.

---

**Note:** While the DDL and metadata changes are made, EVE should NOT be running.

---

In series with actual DDL changes and metadata modification, the code generator must also be run immediately while database users are still logged out. DDL changes can invalidate generated code, but the generated code can only be corrected accordingly after the metadata is changed, so this third thing, code generation, must be done in sequence with its precedents. EVE should not be running during code generation either, as will be discussed.

A possible exception to this rigorously limited circumstance under which metadata can be changed is when new actions or functions are added to the system. By the time you have gone through development and test, you should have the expectation that migration of the new packages or Forms will be routine. Creating the metadata for these objects, before shutting down EVE, and doing the other two steps in the process:

- (a) Actually installing the changes and
- (b) Running the code generator.

If you choose to make modifications to the metadata, other than the times when EVE is shutdown, it is critical that you only add or change the metadata that you just added and actually add the new code prior to running the metadata generator.

## Rule Construction and Modification

Rules can be constructed or modified at any time except while the code generator is running. New rules appear in the system only after the code generator is run. Modifications involving an exception or event end date require that the code generator be run before they will take effect. Other modifications, such as the linking between exceptions and event, or schedule to either exceptions or events are dynamic, and will take effect as soon as they are applied. There is no significant administrative task here, but this process is highly dependent on the code generation process that it is worth highlighting. Rule construction drives the need for code generation, so database administrators and business analysts may like to stay coordinated about how and when the new rules are likely to be created.

## Import Export Tool (IET)

IET is a tool for copying rules into and out of an ARI instance. It allows ARI consultants to supply retailers with pre-packaged rules. It allows retailers, to replicate rules between different testing and production instances. This is preferable to manually recreating rules within each ARI instance. Rules are exported to .xml documents in a database-independent form. From these documents the rules can be imported to a different ARI instance, where IET will attempt to resolve metadata references within the rule to corresponding metadata in the new ARI instance. ARI consultants can also package up supporting actions and data that will be added to new ARI instance with the rule during import. Select Enter to create blank line under the last line of text. This is where you will insert the new graphic.

## Code Generation

Two processes require code generation. Rule construction/modification, and metadata changes. After either of these occurs, code generation is required. Code generation is required immediately after metadata changes (except in the instance of new metadata where delayed code generation may be deferred), and in a reasonable timeframe after rule construction/modification. (For details on executing code generation, see the API details in Appendix B.)

One of the issues with when to run code generation is how soon it needs to be run after rule construction or modification. The code generator builds supporting code for

exceptions and events for all new exceptions and events configured to start before some future time; specifically, now (the time the code generator is started) plus the number of hours defined by the FORWARD\_GENERATION\_HOURS option. (For forward generation details reference Appendix C on configuring ARI options).

Obviously code cannot be generated constantly, so forward generation is done in anticipation of a business analyst defining an exception or event several hours, or even a few days, before they actually want it to take effect. Generally, significant planning and design effort are put into defining exceptions and events, so planning well enough ahead that, the code generator need not be run immediately is typical in a production environment. Certainly the code generator can be run at any time that EVE is stopped, but ideally in a production environment it is run at an off-hours time so that any issues can be handled with minimal impact on the production environment.

Exceptions and events can be linked even after code generation, both linked to each other and to schedules. However, rules cannot be modified after code generation, and the consequence of forward generation is that changes to effective dates do not take effect until the code generator is run again. An exception or event set up to start immediately will only start after the code generator is run. Once scheduled to stop in three weeks, but you want to stop it now instead, can be stopped as soon as you set the end date and run the code generator again.

Both of these cases can be handled as they occur by simply stopping EVE and running the generator on demand, but a typical configuration might be to forward generate by 30 hours and run the generator daily at approximately the same time (as a scheduled job even) every day. This typical configuration does not preclude on-demand generation as well (alternatives and other issues are in Appendix C).

## **ARI Start**

Once the first rules are built and the code generator run, EVE and the other continuously running processes (Scheduler) should be active whenever the database is running. Running the appropriate ARI start program can start all of these programs. If only EVE is shut down (perhaps during a code generation run) and in need of starting, then a different program can be used to start EVE only (see Appendix B).

## Summary

The following diagram summarizes 4 possible options (vertically) for performing key routine ARI tasks. The critical processes for which EVE is not running and users are logged out are highlighted in darker boxes, specifically Change Metadata and Modify DDL. The left-hand column shows the status of the DBA\_JOBS queue each step of the way.

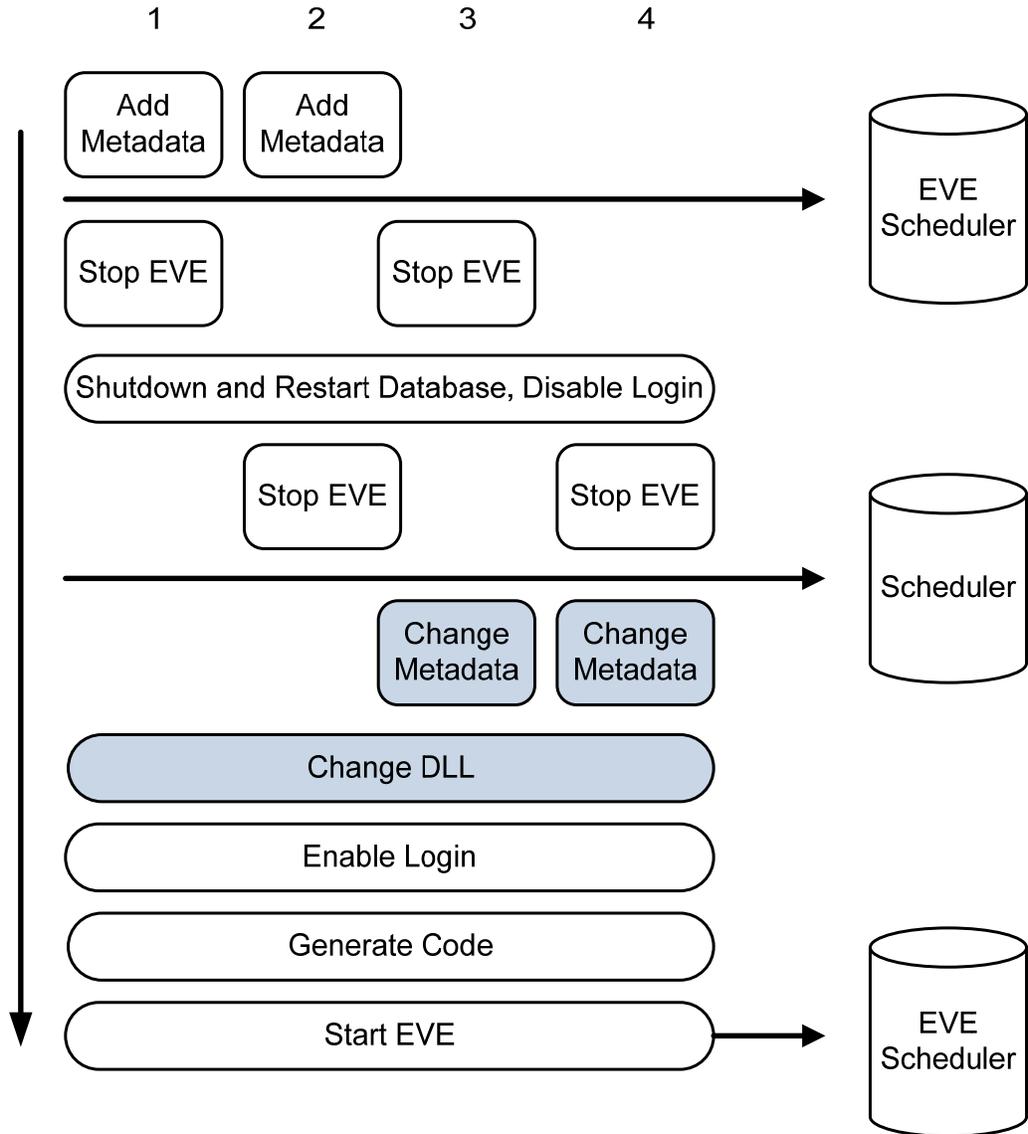


Diagram: DBA\_JOBS Queue

---

---

## Process Details

ARI has several continuously running processes whose progress should be checked periodically. The processes write to an activity log and an error log table, `ARI_ACTIVITY_LOG` and `ARI_ERROR_LOG` respectively. The continuously running processes and the additional processes they spawn are queued in the `DBA_JOBS` queue. Both business analysts and database administrators will need to be familiar with these tables and processes, and at least the database administrator, with some escalation procedure to the business analyst, will want to review the tables and processes on at least a daily basis.

### ARI Logs

#### Understanding Log Types

`ARI_ERROR_LOG` should usually be empty. When it is not, then the errors should be reviewed. Certain types of errors may be non-critical, but a repeated number of such errors may indicate some action is required. Everything written to the error log is also in `ARI_ACTIVITY_LOG` (the activity log), but the activity log contains many more details of what the system is doing and can be helpful for troubleshooting.

#### Setting the Log Level

ARI backend processes have message levels assigned to log entries. Level 1 messages are the most important, with level 3 messages being at a fine detail level. Typically we only run level 2 at Oracle Retail, but level 1 may be sufficient for production. The log level is cached for a given database session, so to change it requires starting a new session. In the case of EVE or the Scheduler this means stopping them, resetting the `LOG_LEVEL` option on `ARI_OPTIONS` (see Appendix C) and creating a new database session when restarting them.

#### Reviewing the Log

When the error log is empty the activity log usually does not need review (unless things are not happening as the way it should be, in which case lack of recent entries in either log may indicate ARI is not running). When any errors have been debugged the logs can be truncated.

#### Purging the Log

There is no automated process for purging the error and activity logs; however, these tables can become quite large if they are not truncated periodically. This truncation is at the discretion of the business analyst and database administrator reviewing the logs. In production daily truncation after review, it should be a standard procedure.

#### DBA\_JOBS Queue

`DBA_JOBS` will usually contain a scheduler job, an EVE job and a number of validation threads (spawned by EVE) and a number of batch scans and event reevaluation jobs (spawned by the scheduler). Broken jobs indicate something not working correctly.

Broken validation threads typically mean that EVE is not going to be spawning as many threads as it could because it will be waiting for the broken thread to return. After troubleshooting, broken jobs should be removed (DBMS\_JOB.REMOVE) and, if any of the jobs were validation threads, EVE should be shutdown and restarted.

## Scheduler

The scheduler lives in the DBA\_JOBS queue. Every five minutes it determines what ARI schedules are due for execution, and generates batch exception scans and event reevaluation blocks, and inserts them into the DBA\_JOBS queue. If an exception has multiple schedules and is already being scanned by a function still in queue, that exception will not be added to the new scan if it is scheduled again. The scheduler should remain in the DBA\_JOBS whenever the database is running after ARI is installed. It is not harmful for the schedule to be out of the queue, but scheduled jobs will get backed up until it is restarted and will not run on schedule.

Multiple instances of this job should not be running. You must decide from a scheduling standpoint, whether to remove it at shutdown (and again at startup if the typical shutdown is a shutdown immediate) or whether just to leave it in the queue past shutdown (it should start up when the database restarts). In the latter case the DBA\_JOBS should be checked daily to make sure it is still an active job (this can be validated by an entry in ARI\_ACTIVITY\_LOG approximately every five minutes indicating the scheduler is running looking for things to do).

## Exception Validation Engine (EVE)

EVE uses the REALM\_QUEUE\_CTL table to track which realms it is currently processing and which ones to do next. EVE creates validation threads that work on individual realms. These threads get added to DBA\_JOBS and when they are complete they tell EVE they are finished so that even can updated the control table and create a new thread on a different realm. Throughout all of this, EVE and its validation threads write to the log extensively.

EVE and its validation threads should be removed from the DBA\_JOBS queue before shutdown, or, if they are not because the shutdown is immediate, they should be removed on restart. This is necessary because EVE uses memory persistent information to control its threading process. The issue is primarily one of conserving system resources, though other undesirable behaviors can also result from having the too many additional validation threads that EVE creates when not properly stopped completely either before shutdown or immediately after restart. Before shutting down or running the code generator, it is important to check the DBA\_JOBS queue to make sure that EVE is actually done; since its shutdown procedure may not be immediate.

## Code Generation

Please note that for code generation, EVE should not be running and the metadata and DDL and Forms must all be synchronized when run. Provided these conditions are met it can be run anytime, though it is generally recommended that it be scheduled periodically and during off-hours to reduce the impact of any code generation errors that might occur. The code generator logs heavily. For troubleshooting it is sometimes helpful to stop both the scheduler during code generation just to simplify reading the log activity.

## History Purge

The history purge process should be scheduled often, perhaps even daily. Unless it produces an error it should not require any special attention. History purge removes events older than the history retention days specified in the event definition.

## Administrator Groups

ARI uses special user groups, an error group and an analyst group. These groups must always have at least one user each. Both the key database administrator and business analyst should be added to each group before any rules are put into production. If not these users, then someone who will be involved in the daily administration of ARI *must be added before building any rules*. The error group is the assignment target group for all events that enter an error state; the analyst group is the assignment target group for all events that are assigned to otherwise empty groups.

Both groups may also be useful for additional analyst and error monitors that you might like to create with ARI (or that Oracle Retail might in some future release deliver or help build from standard templates during a consulting engagement). The analyst group might be used to build a security function for a custom modification to enforce application security on access to the ARI administrative forms.

---

**Note:** Such access level modifications are allowed in spite of the general caveat that all custom modifications of ARI are unsupported – such a modification is not strictly a functional modification.

---



---

---

## Product Integration

ARI is a tool intended for use with one or more applications, so integration with other products is an inevitable post-installation configuration requirement for useful production operation. Generic integration issues include creating metadata for the systems ARI will work with, linking ARI into the main application presentation interface of these systems (or not), and multi-language support. Some specific assistance is provided to simplify integration with Oracle Retail's Merchandise System (RMS) and the Oracle Retail Data Warehouse (RDW).

### Metadata

Once the seed data has been installed, metadata needs to be created for the systems that ARI will work with. Metadata is the foundation of all ARI rule building, so it is critical that it be accurate. Metadata maintenance is a database administrator task that should become part of the routine that goes with changing Form specifications or table definitions. Fortunately, this does not occur very often in a production system, so it should not be a particularly time consuming task.

Metadata maintenance must be done in sync with DDL changes. Users should not be in the system trying to work with ARI Alerts during the gap in time between when DDL changes are made and when the system is updated. Furthermore, EVE cannot be running during this time. The exception to this rule is when metadata is being added to the system either initially or even after production use of the system has begun. In this case, the metadata can be added during production hours provided the code generator is not run until after the DDL changes are made. This means for initial setup of metadata the entire process can be done during production (and you do not have to log users out for 3 days while you do it).

---

---

**Note:** It is not necessary that the entire system be described in metadata, but the metadata that is described needs to be accurate.

---

---

This can be a time saver since you only need to create metadata for the rules you plan to build. There may be some tables in the system that you will require to select data from or monitor, and hence there will be more Forms and PL/SQL actions that you will not need. Analysis before implementing an ARI rule should clearly identify what physical data entities and actions will be needed. So creating appropriate metadata, if it wasn't already created for another ARI rule, could be treated as part of the rule development process.

### Oracle Retail Merchandise System (RMS)

For owners of the Oracle Retail Merchandise System, Oracle Retail has pre-defined the metadata for all of the tables in the system. This metadata can be populated (provided it is the first metadata added after running the seed data scripts that are required in the ARI Installation Guide) by executing in order the scripts provided: `mts_parm_type.sql`, `mts_realm.sql` and `mts_parm.sql`.

## Multilanguage Support

The technical infrastructure of ARI supports languages other than English. The software can efficiently handle multiple languages. Tables have been added to ARI to accommodate internationalization. The retailer sets up the user's language preferences on the ARI\_USER\_ATTRIB table. At login, the LANG column on the ARI\_USER\_ATTRIB table determines the user's language setting and displays the code string associated with it. ARI has a fail/safe mechanism built into the code. If the user's preferred language is not found, then ARI displays English in the user interface.

---

**Note:** A retailer has the two options below regarding internationalization when installing the application. See the RMS Installation Guide for the procedures related to each.

---

- English and multiple secondary languages
- Install English first and then update with a translated language (fully translated non-English installation)

### Key ARI Tables Related to Internationalization

Several tables were created to handle displayable text that can also be translated. If the retailer creates a new form, a new menu, or a new object on a form, then the retailer will need to populate these tables with the corresponding information. If the retailer customizes the information in any of the tables ARI\_FORM\_ELEMENTS, ARI\_FORM\_ELEMENTS\_LANGS, ARI\_MENU\_ELEMENTS, or ARI\_MENU\_ELEMENTS\_LANGS, the base\_ind field in customized records must contain 'N'. Any record with BASE\_IND=N will be preserved in a temp table during future patches.

#### ARI\_FORM\_ELEMENTS

This table is used for screen display and holds the master list of items for all forms whose labels/prompts are translated. This information will always be in English. The BASE\_IND=Y means that the item is part of the base Oracle Retail code set. BASE\_IND=N indicates that the item was added as part of retailer customization. Anything with the BASE\_IND=N will be preserved at upgrade time on the FORM\_ELEMENTS\_TEMP, but the retailer is responsible for moving the data back to FORM\_ELEMENTS.

#### ARI\_FORM\_ELEMENTS\_LANGS

This table is used for screen display. This table holds translated values for labels/prompts on forms. This information is in a language that is defined on the lang column of the ARI\_USER\_ATTRIB table. All users see data from this table; the retailer may customize the text of a given field. The access key for a button is defined by filling in the DEFAULT\_ACCESS\_KEY field. At runtime, that character will be marked in the string, and function as the access key. Any time the retailer changes the DEFAULT\_LABEL\_PROMPT or DEFAULT\_ACCESS\_KEY, the BASE\_IND should be updated to N because it is not part of the base language translations provided by Oracle Retail. Anything with the BASE\_IND=N will be preserved at upgrade time on the FORM\_ELEMENTS\_LANGS\_TEMP, but the retailer is responsible for moving the data back to FORM\_ELEMENTS\_LANGS.

#### ARI\_MENU\_ELEMENTS

This table is used for screen display. This table holds the master list for all menus whose items are translated. This information will always be in English. The access key for a

menu option is defined by using the ampersand (&) before the character that is the access key in the default description. The `BASE_IND=Y` means that the item is part of the base Oracle Retail code set. `BASE_IND=N` indicates that the item was added as part of retailer customization. Anything with the `BASE_IND=N` will be preserved at upgrade time on the `MENU_ELEMENTS_TEMP`, but the retailer is responsible for moving the data back to `MENU_ELEMENTS`.

### **ARI\_MENU\_ELEMENTS\_LANGS**

This table is used for screen display. This table holds the values for all menus whose items are translated. This information will be in a language that is defined on the `lang` table. Even English language users see data from this table, as the retailer may customize the text of a given menu option. Any time the retailer changes the `LANG_LABEL`, the `BASE_IND` should be updated to `N` because it is not part of the base language translations provided by Oracle Retail. Anything with the `BASE_IND=N` will be preserved at upgrade time on the `MENU_ELEMENTS_LANGS_TEMP`, but the retailer is responsible for moving the data back to `MENU_ELEMENTS_LANGS`.

### **ARI\_FORM\_MENU\_LINK**

This table is used for screen display. This table holds the intersection of form and menu files, mapping each form to the menu that it displays.

## **Presentation Interface**

A key feature of ARI is to be able to be notified in your operational applications that an alert has occurred in ARI. To that end ARI provides an API set that will enable you to put a button on the toolbar of your Forms applications. The button can represent whether any new (undeferred) alerts exist for the user logged into the database, and pressing it will launch the ARI Alert Viewer. For non-Forms applications it is possible to bypass the Forms library and simply access the PL/SQL procedure that does all of the work. (For details on the API, consult Appendix B.)

ARI has a startup form (`aristart.fmb`) that allows access to the ARI forms. This form is primarily intended for administrative, full functionality access users to get into ARI. The integrated end-user entry point into the system will typically be through a notification button on the toolbar of the applications that ARI is monitoring, as just described. However, if that entry point is not used, you likely will want to add role-based security to the menus in the `ARISTART` form to restrict end-user access.

---

**Note:** The ARI launch/notification button already exists in RMS, and the API interface is already installed.

---

To make the interface operational, simply drop the PL/SQL portion of the interface (`ARI_INTERFACE_SQL`) from the RMS product schema and replace it with a synonym to that same package in the ARI schema.

## Oracle Retail Data Warehouse

ARI has monitoring and metadata methods specific to the Oracle Retail Data Warehouse. It is not a typical Oracle transaction system and so the integration with this product has been done above and beyond typical ARI functionality. Except for creating a database link from the ARI Master schema to the Oracle Retail Data Warehouse instance, no special configuration of ARI is necessary to use this feature. If you do not have the RDW, nothing special need be done.

### RDW Report Setup

The RDW Interface accesses RDW report data that is stored in a Datamart. The Datamart functionality allows report data to be stored in Oracle tables that can be accessed by ARI. Reports that are to be run for ARI must be set up to store their data in a Datamart schema. This schema must either be located on the ARI database, or on another database (such as the RDW database) which is accessible through a database link from the ARI Master schema. Moreover, although in a typical RDW installation some reports may be run and cached periodically while many others will be generated on demand, reports monitored by ARI are generally scheduled. This RDW report scheduling usually involves some manual synchronization with the exception's report-monitoring schedule, configured in ARI. All of these factors lead us to recommend that reports intended for monitoring by ARI be set up under a single user in a single project, even if some of these reports are redundant with others that already exist in the RDW.

### Metadata Setup

Metadata auto-discovery will only search for reports owned by a single Datamart schema, in a single database, which are stored in the ARI options RDW\_OWNER, RDW\_LINK. Auto-discovery stores these option settings in the SCHEMA and DB\_LINK fields on the table REALM, and it is these fields which are used when searching for cached reports during exception monitoring. Therefore it is possible to manually create metadata for reports stored in a different schema and/or database. In addition, if the ARI option RDW\_PREFIX is set, auto-discovery will only search for Datamart tables beginning with this prefix. During auto-discovery, the RDW Interface must search for Datamart tables created for a report, in other words, the report must already have been run and stored a non-empty result set to the Datamart before auto-discovery can be executed.

## Integration with Oracle Retail Workspace

The Oracle Retail Workspace installer prompts you to enter the URL for your supported Oracle Retail applications. However, if a client installs a new application after Oracle Retail Workspace is installed, the retail-workspace-page-config.xml file needs to be edited to reflect the new application.

The file as supplied comes with all appropriate products configured, but the configurations of non-installed products have been "turned off". Therefore, when "turning on" a product, locate the appropriate entry, set "rendered" to "true", and enter the correct URL and parameters for the new application.

The entry consists of the main URL string plus one parameter named "config". The value of the config parameter is inserted by the installer. Somewhere in the installer property files there is a value for the properties "deploy.retail.product.rms.url" and "deploy.retail.product.rms.config".

For example, suppose ARI was installed on mycomputer.mycompany.com, port 7777, using a standard install and rms configured with the application name of "ari13intsso ". If you were to access ARI directly from your browser, you would type in:

```
http://mycomputer.mycompany.com:7777/forms/frmservlet?config=ari13intsso
```

The entry in the retail-workspace-page-config.xml after installation would resemble the following:

```
<secure-work-item id="ari"
    display-string="{confMsgs.ariTitle}"
    rendered="true"
    launchable="true"
    show-in-content-area="false">
  <url>http://mycomputer.mycompany.com:7777/forms/frmservlet</url>
  <parameters>
    <parameter name="config">
      <value>ari13intsso</value>
    </parameter>
  </parameters>
</secure-work-item>
```

---

**Note:** See Appendix A: Architectural Reference for information on single sign-on and how to set up ARI with single sign-on.

---

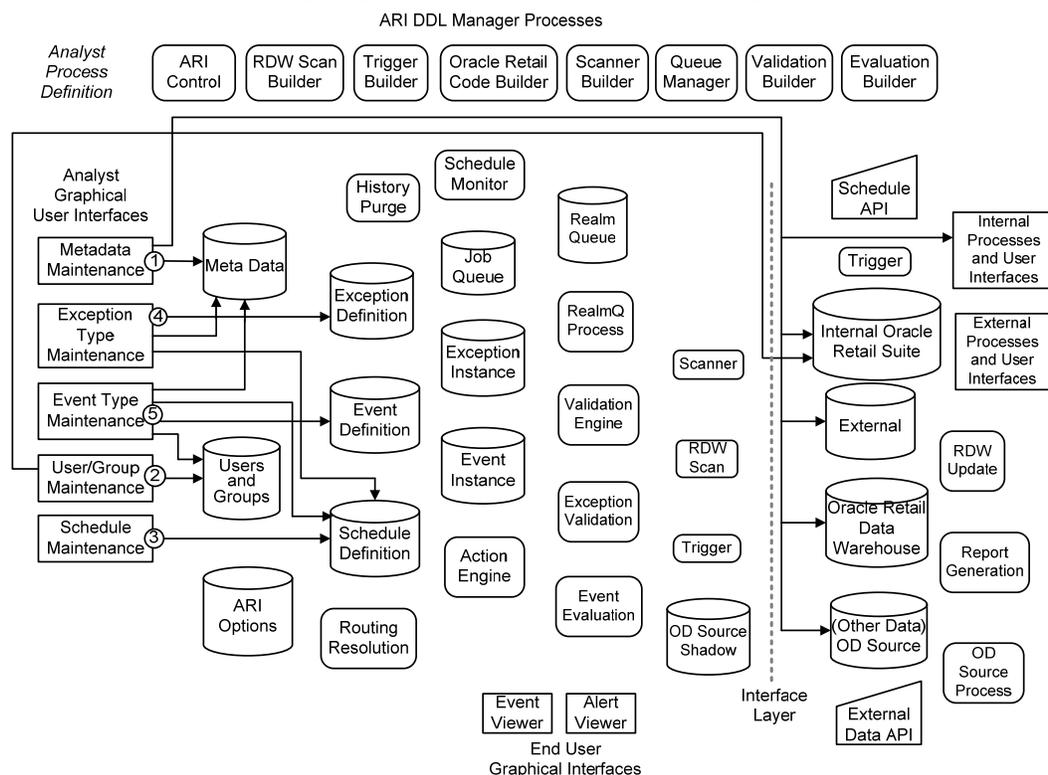


## Appendix: Architectural Reference

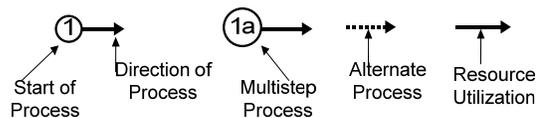
This section provides diagrams and descriptions of the processes that drive ARI. These processes are grouped according to the main functional uses of ARI.

### Analyst Process Definition

Before ARI will perform any functions, several processes must be completed to tell ARI what to do. The following diagram and description walk through these processes.



- ① Set up and maintain meta data
- ② Set up and maintain Users and Groups
- ③ Set up and maintain schedules
- ④ Define Exception Types
- ⑤ Define Event Types



#### Process Diagram: Analyst Process Definition

Descriptions of the processes shown in the diagram follow. The numbers in the headings correspond to numbered processes in the diagram.

#### 1. Set Up and Maintain Metadata

Database administrators keep metadata and Oracle DDL synchronized. In a production environment, the DDL should seldom, if ever, change. So, once set up, maintaining the

metadata is not a time-consuming task. Metadata for new functions/actions to support ARI processes defined by business analysts can be set up and maintained by either a DBA or a business analyst. Metadata for Data Warehouse reports and external systems (those using the API) is setup and maintained by business analysts. Like functions and actions, these reports and external systems are probably defined in metadata on-demand, as needed to support specific ARI processes.

## **2. User/Group Maintenance**

Performed by business analysts, user and group maintenance supports event assignment and event supervision assignment. Users entered here must be defined as Oracle users in the database.

## **3. Schedule Maintenance**

Schedules indicate when a task will occur. These are maintained by business analysts and assigned to exception and event definitions to determine when periodic exception scans and automatic event re-evaluations will occur.

## **4. Exception Type Maintenance**

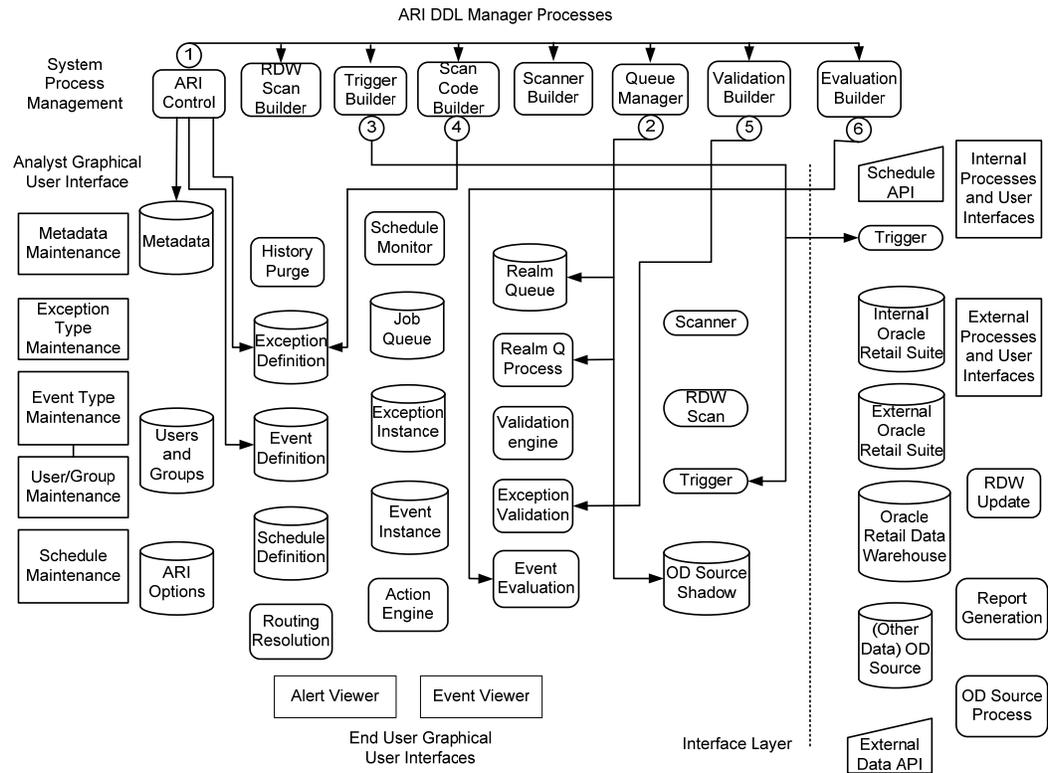
Business analysts maintain exception definitions. They define when to monitor for specific data conditions, what the minimal actionable data conditions are, and what event(s) should be created to help resolve the exception.

## **5. Event Type Maintenance**

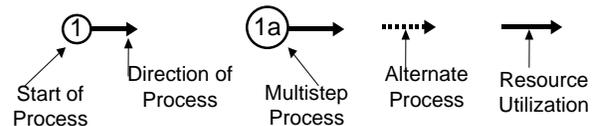
Also performed by business analysts, event definition maintenance determines the presentation of exception-related information in the form of alerts. An event contains an alert definition and an exception link, plus information about which users are notified, what actions are available to resolve the event, and what new data conditions will determine that the event is, in fact, resolved.

# System Process Management

Using analysts' definitions, ARI prepares to monitor exceptions and present events.



- ① ARI Daily Management
- ② Table Manager and Queue Process Builder
- ③ Table Trigger Builder and Generator
- ④ Where Clause Builder
- ⑤ Candidate Validation Procedure Generator
- ⑥ Event Evaluation Procedure Generator



## Process Diagram: System Process Management

Descriptions of the processes shown in the diagram follow. The numbers in the headings correspond to numbered processes in the diagram.

### 1. ARI Daily Management

The daily ARI control program is run at the beginning of the batch window. After all users are cleared, it shuts down the currently active ARI processes, rebuilds exception and event management code (processes 2-6), and then restarts the ARI processes.

### 2. Table Manager and Queue Process Builder

This process creates a realm queue table for every monitored realm (if one does not already exist) and deletes those realm queue tables no longer needed. It also builds/drops (as appropriate) realm queue procedures, one per queue, to manage queued data. Finally, it builds the shadow tables that hold data input from external (non-Oracle) data systems.

### **3. Trigger Builder**

Drops and creates triggers to monitor real-time exceptions on the Oracle Retail suite that is part of the same DDL as ARI. Also drops/creates triggers on the external source shadow tables that make externally sources data not real-time monitored but trickle-monitored. (The data is moved into the queue real-time as it is received).

### **4. Scan Code Builder**

The scan code builder prepares appropriate clauses to select only rows of interest (based on the exception definition) from tables that are monitored via periodic scanning. This data is stored with the exception definition for easy retrieval by the Scanner builder.

### **5. Validation Procedure Generator**

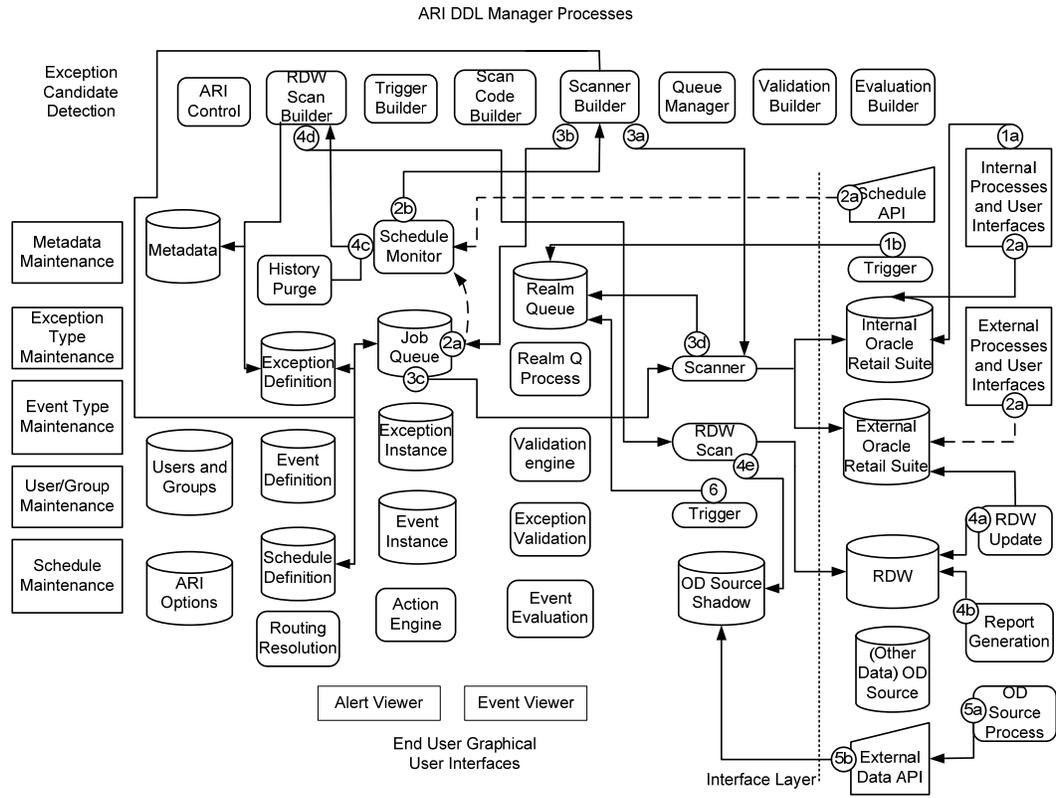
Exception candidates must be validated to determine whether an exception exists. This involves selecting additional information from other parameters and processing the conditions on those parameters. The procedure generator creates this business logic in a static package.

### **6. Evaluation Procedure Generator**

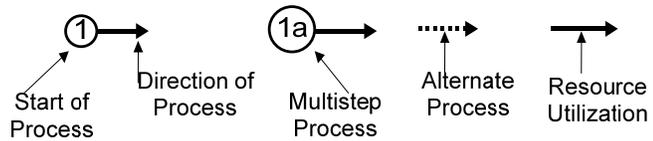
Once an exception candidate is known to exist, it can create one or more events. Underlying data conditions may change the way an event should be presented (what actions are available) or who should be dealing with it. The reevaluation logic is encapsulated in a static package.

# Exception Candidate Detection

To minimize performance impact on other systems, ARI first identifies candidate exceptions by imposing some of the exception conditions on the data set being monitored. This is done in order to filter out those that could be exceptions, based on complete condition evaluation after additional information is fetched, from those that could not be data states of interest, irrespective of what other information might be gathered up during the exception process. Candidates are queued and processed as systems resources are available.



- ① Real-Time Monitor
- ② Periodic Monitor
- ③ Schedule Monitor
- ④ Data Warehouse Monitor
- ⑤ External API Monitor
- ⑥ Trickle Monitor



## Process Diagram: Exception Candidate Detection

Descriptions of the processes shown in the diagram follow. The numbers in the headings correspond to numbered processes in the diagram.

## 1. Real-Time Monitor

When a user makes changes to the database, triggers write candidates directly into their corresponding realm queue.

## 2. Periodic Monitor

When changes occur in the database that cannot be monitored in real-time due to high volume of changes type considerations or just because the change is time-based (that is, the data is aging), periodic monitoring must be used. The actual monitoring occurs either when a signal is sent to the schedule monitor from the schedule API or when the schedule monitor is periodically prompted by the Job Queue, finds a schedule in need of execution, and runs all of the monitors associated with that schedule.

## 3. Scanner Builder

The scanner builder assembles the scan code for all exceptions scheduled to occur simultaneously, and places a job in the Job Queue to run those scans. As threads become available for processing, the Job Queue initiates the scans that populate data representing exception candidates into the realm queue tables.

## 4. Data Warehouse Monitoring

On a periodic basis, data in the Data Warehouse is updated and reports are re-run. To keep the data warehouse monitoring in sync with this schedule, in addition to kicking off the scanner builder, the schedule monitor kicks off the RDW scan builder when required. The scan builder dumps report data directly into a shadow table.

## 5. External API Monitor

This API is used to feed in data from an external (non-Oracle Retail) system. As the API is called from the external system it writes data to an appropriate shadow table within ARI.

## 6. Trickle Data Monitor

The trickle data monitor takes data as it arrives from non-Oracle systems via an appropriate monitor, and pre-filters it, putting the candidates into the appropriate realm queue. Although Data Warehouse monitoring is periodic, the idea of trickle monitoring is that, the data is processed as soon as it arrives, which in the case of external systems, depends on where in the external system the ARI External API is implemented, not on any scheduling done with ARI. Indeed, the same API could be used with the Data Warehouse, but for convenience the link to this Oracle Retail system is provided in the form of a periodic link, which should be sufficient if synchronized with the periodic updates of the Data Warehouse and its monitored reports.



## **2. Realm Queue Process**

The Job Queue starts these processes when threads become available. They read candidate exceptions and call the generated validation procedures for each candidate. After each candidate, the time-slice allocation is checked. When that allocation is exceeded, realm queue processing is terminated and the validation engine is notified.

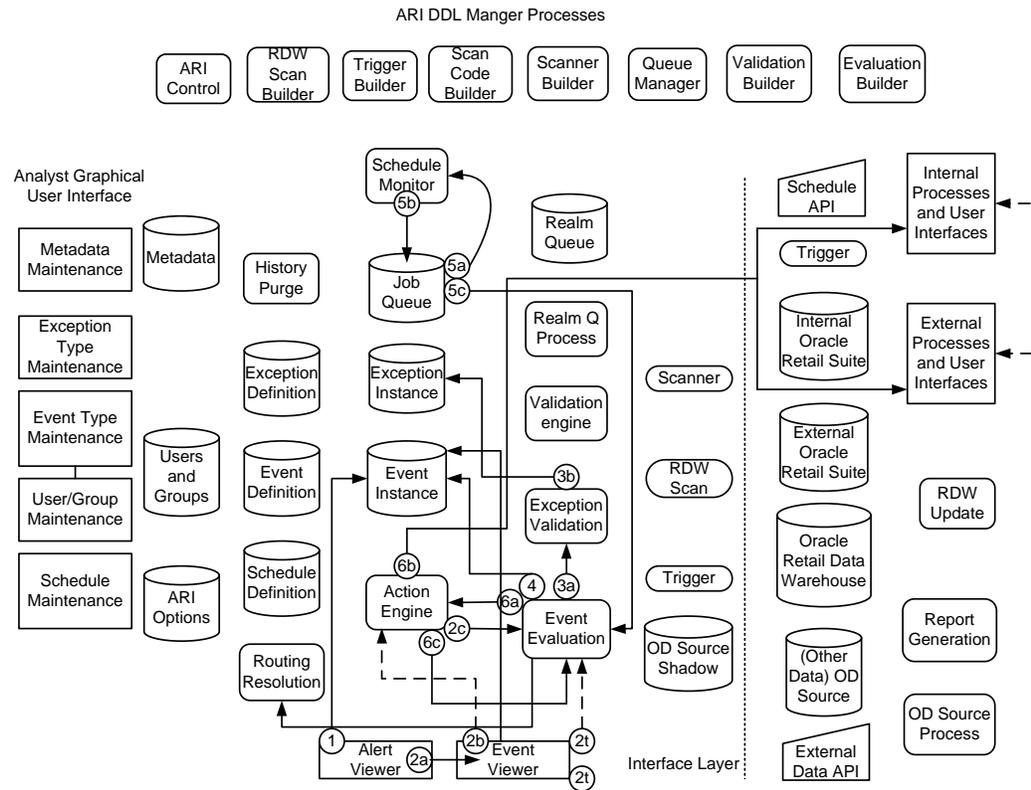
## **3. Exception Creation and Validation**

Valid candidates become exception instances, and then event evaluation is invoked to determine which events should be created or, if they already exist, re-evaluated.

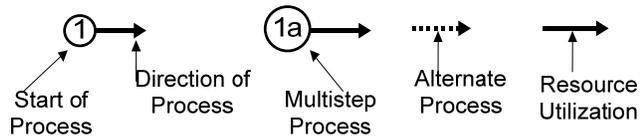
## **4. Event Creation and Evaluation**

Event instances for each event type linked to the exception instances type, must be either refreshed, if an event with the same key values already exists, or created. Evaluation occurs to determine alert routing. Then, the instance is created and, if the appropriate rules are met for a particular instance, an action may be taken automatically.

# User-Initiated and Automated Event Resolution



- ① Alert and Filter Management
- ② Event Management
- ③ Exception Revaluation
- ④ Event Revaluation
- ⑤ Schedule Driven Revaluation
- ⑥ Action Process



## Process Diagram: User Initiated and Automated Event Resolution

Descriptions of the processes shown in the diagram follow. The numbers in the headings correspond to numbered processes in the diagram.

### 1. Alert Viewer

End users initially view events as alerts, which are just event summaries, in the Alert Viewer. This viewer allows sorting and categorization by event type, state, priority, occurrence data, and even whether the user has deferred (marked as having been viewed) the alert. In the Alert Viewer, events are grouped by type state and priority with a count of how many of each type of state and priority exist.

## 2. Event Management

The Event Viewer shows all events of a single type and state at any given time, and also shows all the details of those events. The user initially is shown the most recently refreshed event data (which may or may not be very recent, depending on whether it is done periodically as defined in the event type). The user may choose to reevaluate the events before proceeding, since re-evaluation may show that some events are already resolved. The user may also drill into the monitored systems to view even more information than what is presented on the event. In addition, the user may take an action to resolve the event or move it along in its process. If an action is taken, the event is still reevaluated beforehand to ensure that the event is still in the state in which the action is valid.

## 3. Exception Revalidation

Because the data source for many event parameters is an exception, before an event is reevaluated, its associated exception instances are revalidated. This process is just like validation, except that the exception itself is treated like a new candidate, and is removed from the event and deleted if it is found to be invalid.

## 4. Event Reevaluation

Once the exceptions are revalidated, event reevaluation refreshes the event parameters and processes its defining rules to determine whether the event is resolved, and should be closed, whether it is in the same state as it was in, and to whom it should be assigned and with what priority. If the reevaluation is user-requested either directly or as a consequence of the user attempting to take an action, and if the state and user assignment are unchanged, then the action can be executed, or if user requested, the user may continue to view and act on the event. Because the Event Viewer shows only a single state at a time, events that change state, although the same user may still own them, will not continue to be seen in the view of the event's previous state.

## 5. Schedule Reevaluation

Evaluation can be user initiated via one of the previously described processes and may also occur on a scheduled basis. As with scheduled scans for exception candidates, the Job Queue notifies the scheduler to look for tasks that need to be done, which in turn places the tasks in to the Job Queue. As threads become available, these evaluation jobs, specific to an event type, are started. Each job loops through all open instances of their event type and initiates the reevaluation process.

## 6. Action Execution

Actions are always executed after an evaluation occurs. If an action is user-requested, the event is always evaluated first. Alternatively, as a result of an evaluation, whether prompted by a user request or a schedule, an action may be taken automatically if the event is so defined. After the action is taken, the event is always reevaluated, which may prompt another action, and so on. A well-defined event should never get into an infinite loop, but it is theoretically possible. Therefore, checks are in place to prevent an infinite evaluation-action-evaluation loop from occurring.

## Single Sign-on Overview

### What is Single Sign-On?

Single Sign-On (SSO) is a term for the ability to sign onto multiple web applications via a single user ID/Password. There are many implementations of SSO – Oracle currently provides three different implementations: Oracle Single Sign-On (OSSO), Java SSO (with the 10.1.3.1 release of OC4J) and Oracle Access Manager (provides more comprehensive user access capabilities).

Most, if not all, SSO technologies use a session cookie to hold encrypted data passed to each application. The SSO infrastructure has the responsibility to validate these cookies and, possibly, update this information. The user is directed to log on only if the cookie is not present or has become invalid. These session cookies are restricted to a single browser session and are never written to a file.

Another facet of SSO is how these technologies redirect a user's Web browser to various servlets. The SSO implementation determines when and where these redirects occur and what the final screen shown to the user is.

Most SSO implementations are performed in an application's infrastructure and not in the application logic itself. Applications that leverage infrastructure managed authentication (such as deploying specifying "Basic" or "Form" authentication) typically have little or no code changes when adapted to work in an SSO environment.

### What Do I Need for Oracle Single Sign-On?

The nexus of an Oracle Single Sign-On system is the Oracle Identity Management Infrastructure installation. This consists of the following components:

- An Oracle Internet Directory (OID) LDAP server, used to store user, role, security, and other information. OID uses an Oracle database as the back-end storage of this information.
- An Oracle Single Sign-On servlet, used to authenticate the user and create the OSSO session cookie. This servlet is deployed within the infrastructure Oracle Application Server (OAS).
- The Delegated Administration Services (DAS) application, used to administer users and group information. This information may also be loaded or modified via standard LDAP Data Interchange Format (LDIF) scripts.
- Additional administrative scripts for configuring the OSSO system and registering HTTP servers.

Additional OAS servers will be needed to deploy the business applications leveraging the OSSO technology.

### Can Oracle Single Sign-On Work with Other SSO Implementations?

Yes, OSSO has the ability to interoperate with many other SSO implementations, but some restrictions exist.

### Oracle Single Sign-on Terms and Definitions

#### Authentication

Authentication is the process of establishing a user's identity. There are many types of authentication. The most common authentication process involves a user ID and password.

## Dynamically Protected URLs

A “Dynamically Protected URL” is a URL whose implementing application is aware of the OSSO environment. The application may allow a user limited access when the user has not been authenticated. Applications that implement dynamic OSSO protection typically display a “Login” link to provide user authentication and gain greater access to the application’s resources.

## Identity Management Infrastructure

The Identity Management Infrastructure is the collection of product and services which provide Oracle Single Sign-on functionality. This includes the Oracle Internet Directory, an Oracle HTTP server, and the Oracle Single Sign-On services. The Oracle Application Server deployed with these components is typically referred as the “Infrastructure” instance.

## MOD\_OSSO

mod\_osso is an Apache Web Server module an Oracle HTTP Server uses to function as a partner application within an Oracle Single Sign-On environment. The Oracle HTTP Server is based on the Apache HTTP Server.

## Oracle Internet Directory

Oracle Internet Directory (OID) is an LDAP-compliant directory service. It contains user ids, passwords, group membership, privileges, and other attributes for users who are authenticated using Oracle Single Sign-On.

## Partner Application

A partner application is an application that delegates authentication to the Oracle Identity Management Infrastructure. One such partner application is the Oracle HTTP Server (OHS) supplied with the Oracle Application Server. OHS uses the MOD\_OSSO module to configure this functionality.

All partner applications must be registered with the Oracle Single Sign-On server. An output product of this registration is a configuration file the partner application uses to verify a user has been previously authenticated.

## Realm

A Realm is a collection users and groups (roles) managed by a single password policy. This policy controls what may be used for authentication (for example, passwords, X.509 certificates, and biometric devices). A Realm also contains an authorization policy used for controlling access to applications or resources used by one or more applications.

A single OID can contain multiple Realms. This feature can consolidate security for retailers with multiple banners or to consolidate security for multiple development and test environments.

## Statically Protected URLs

A URL is considered to be “Statically Protected” when an Oracle HTTP server is configured to limit access to this URL to only SSO authenticated users. Any attempt to access a “Statically Protected URL” results in the display of a login page or an error page to the user.

Servlets, static HTML pages, and JSP pages may be statically protected.

## What Single Sign-On is not

Single Sign-On is NOT a user ID/password mapping technology.

However, some applications can store and retrieve user IDs and passwords for non-SSO applications within an OID LDAP server. An example of this is the Oracle Forms Web Application framework, which maps OSSO user IDs to a database logins on a per-application basis.

## How Oracle Single Sign-On Works

Oracle Single Sign-On involves a couple of different components. These are:

- The Oracle Single Sign-On (OSSO) servlet, which is responsible for the back-end authentication of the user.
- The Oracle Internet Directory LDAP server, which stores user IDs, passwords, and group (role) membership.
- The Oracle HTTP Server associated with the web application, which verifies and controls browser redirection to the OSSO servlet.
- If the web application implements dynamic protection, then the web application itself is involved with the OSSO system.

### Statically Protected URLs

When an unauthenticated user accesses a statically protected URL, the following occurs:

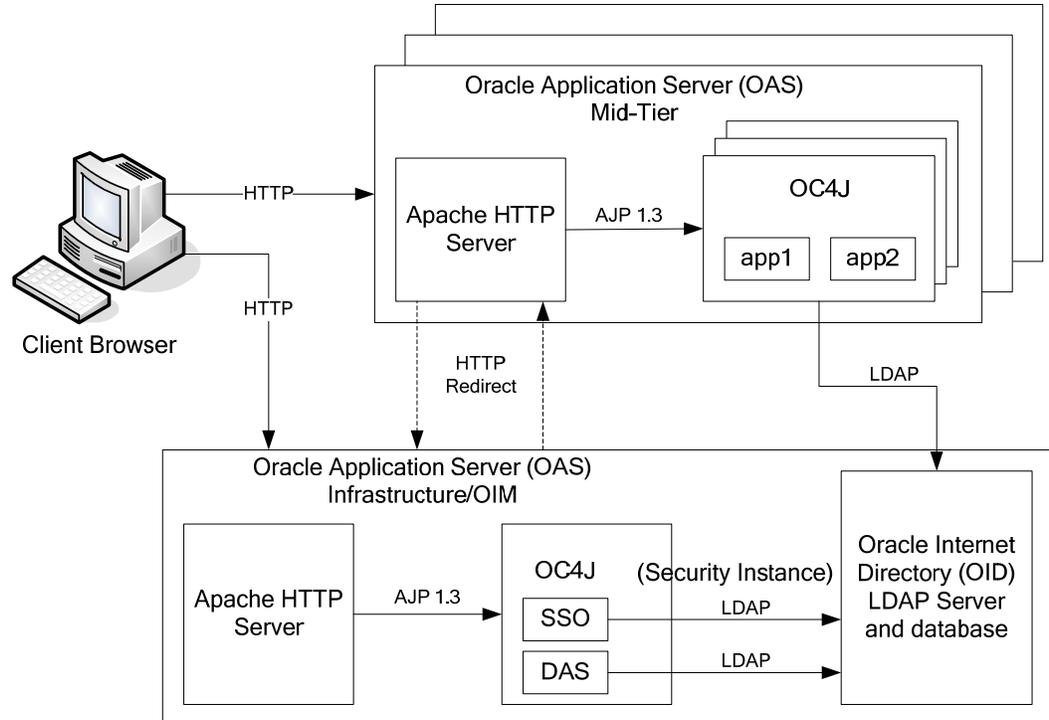
1. The Oracle HTTP server recognizes the user has not been authenticated and redirects the browser to the Oracle Single Sign-On servlet.
2. The OSSO servlet determines the user must authenticate, and displays the OSSO login page.
3. The user must sign in via a valid user ID and password. If the OSSO servlet has been configured to support multiple Realms, a valid realm must also be entered. The user ID, password, and realm information is validated against the Oracle Internet Directory LDAP server.
4. The OSSO servlet creates and sends the user's browser an OSSO session cookie. This cookie is never persisted to disk and is specific only to the current browser session. This cookie contains the user's authenticated identity. It does NOT contain the user's password.
5. The OSSO servlet redirects the user back to the Oracle HTTP Server, along with OSSO specific information.
6. The Oracle HTTP Server decodes the OSSO information, stores it with the user's session, and allows the user access to the original URL.

## Dynamically Protected URLs

When an unauthenticated user accesses a dynamically protected URL, the following occurs:

1. The Oracle HTTP server recognizes the user has not been authenticated, but allows the user to access the URL.
2. The application determines the user must be authenticated and sends the Oracle HTTP server a specific status to begin the authentication process.
3. The Oracle HTTP Server redirects the user's browser session to the OSSO Servlet.
4. The OSSO servlet determines the user must authenticate, and displays the OSSO login page.
5. The user must sign in via a valid user ID and password. If the OSSO servlet has been configured to support multiple Realms, a valid realm must also be entered. The user ID, password, and realm information is validated against the Oracle Internet Directory LDAP server.
6. The OSSO servlet creates and sends the user's browser an OSSO session cookie. This cookie is never persisted to disk and is specific only to the current browser session. This cookie contains the user's authenticated identity. It does NOT contain the user's password.
7. The OSSO servlet redirects the user back to the Oracle HTTP Server, along with OSSO specific information.
8. The Oracle HTTP Server decodes the OSSO information, stores it with the user's session, and allows the user access to the original URL.

## Single Sign-on Topology



## Installation Overview

Installing Oracle Single Sign-On consists of installing the following components:

1. Installing the Oracle Internet Directory (OID) LDAP server and the Infrastructure Oracle Application Server (OAS). These are typically performed using a single session of the Oracle Universal Installer and are performed at the same time. OID requires an Oracle relational database and if one is not available, the installer will also install this as well.  
The Infrastructure OAS includes the Delegated Administration Services (DAS) application as well as the OSSO servlet. The DAS application can be used for user and realm management within OID.
2. Installing additional OAS 10.1.2 midtier instances for the Oracle Retail applications, such as RMS, that are based on Oracle Forms technologies. These instances must be registered with the Infrastructure OAS installed in step 1).
3. Installing additional application servers to deploy other Oracle Retail applications and performing application specific initialization and deployment activities.

### Infrastructure Installation and Configuration

The Infrastructure installation for OSSO is dependent on the environment and requirements for its use. Deploying an Infrastructure OAS to be used in a test environment does not have the same availability requirements as for a production environment. Similarly, the Oracle Internet Directory (OID) LDAP server can be deployed in a variety of different configurations. See the *Oracle Application Server Installation Guide* and the *Oracle Internet Directory Installation Guide* for more details.

### OID User Data

Oracle Internet Directory is an LDAP v3 compliant directory server. It provides standards-based user definitions out of the box.

The current version of Oracle Single Sign-On only supports OID as its user storage facility. Customers with existing corporate LDAP implementations may need to synchronize user information between their existing LDAP directory servers and OID. OID supports standard LDIF file formats and provides a JNDI compliant set of Java classes as well. Moreover, OID provides additional synchronization and replication facilities to integrate with other corporate LDAP implementations.

Each user ID stored in OID has a specific record containing user specific information. For role-based access, groups of users can be defined and managed within OID. Applications can thus grant access based on group (role) membership saving administration time and providing a more secure implementation.

### OID with Multiple Realms

OID and OSSO can be configured to support multiple user Realms. Each realm is independent from each other and contains its own set of user IDs. As such, creating a new realm is an alternative to installing multiple OID and Infrastructure instances. Hence, a single Infrastructure OAS can be used to support many development and test environments by defining one realm for each environment.

Realms may also be used to support multiple groups of external users, such as those from partner companies. For more information on Realms, see the *Oracle Internet Directory Administrators Guide*.

## User Management

User Management consists of displaying, creating, updating or removing user information. There are two basic methods of performing user management: LDIF scripts and the Delegate Administration Services (DAS) application.

### OID DAS

The DAS application is a web based application designed for both administrators and users. A user may update their password, change their telephone number of record, or modify other user information. Users may search for other users based on partial strings of the user's name or ID. An administrator may create new users, unlock passwords, or delete users.

The DAS application is fully customizable. Administrators may define what user attributes are required, optional or even prompted for when a new user is created.

Furthermore, the DAS application is secure. Administrators may also what user attributes are displayed to other users. Administration is based on permission grants, so different users may have different capabilities for user management based on their roles within their organization.

### LDIF Scripts

Script based user management can be used to synchronize data between multiple LDAP servers. The standard format for these scripts is the LDAP Data Interchange Format (LDIF). OID supports LDIF script for importing and exporting user information. LDIF scripts may also be used for bulk user load operations.

### User Data Synchronization

The user store for Oracle Single Sign-On resides within the Oracle Internet Directory (OID) LDAP server. Oracle Retail applications may require additional information attached to a user name for application-specific purposes and may be stored in an application-specific database. Currently, there are no Oracle Retail tools for synchronizing changes in OID stored information with application-specific user stores. Implementers should plan appropriate time and resources for this process. Oracle Retail strongly suggests that you configure any Oracle Retail application using an LDAP for its user store to point to the same OID server used with Oracle Single Sign-On.

## Setting up ARI for Single Sign-on

To set up Forms for Single Sign-on, the Forms framework must know and/or be configured to use SSO. To do this, the Forms framework configuration file `formsweb.cfg` must be configured to enable SSO and the mid-tier HTTP Server must be registered with the Oracle Single Sign-On server. In addition, the Forms framework uses Resource Access Descriptor (RAD), to map OSSO user IDs to Database connect strings.

### Configuring `formsweb.cfg`

For each Forms application instance there are two attributes in the Forms framework configuration file `formsweb.cfg` that control OSSO behavior:

Name	Value	Description
<code>ssoMode</code>	<code>true/false</code>	Enables/disables OSSO
<code>ssoDynamicResourceCreate</code>	<code>true/false</code>	enables/disables the dynamic RAD entry creation

## Creating a RAD Entry

There are three ways by which a RAD entry (mapping an OSSO user ID to a Database connect string) may be created:

- Administrator Created  
An administrator uses the Delegated Administration Services (DAS) web application that comes with the infrastructure server and that can be launched using the URL `http://<host>:port/oiddas`.
- User Created  
The user can dynamically create a RAD entry when the Forms framework prompts the user for information. This however requires that the `ssoDynamicResourceCreate` attribute be set to true. If a RAD already exists, the user may also create additional RADs via the DAS application.
- LDIF Script  
More information about how to use an LDIF script to create a RAD entry may be found by accessing Metalink document 244526.1.

---

---

## Appendix: API List

### Error and Activity Logging

There are PL/SQL APIs related to error and activity logging that can be incorporated into scripts or used in other contexts.

### Changing the Log Level

Within a particular database session, calling `ari_error_logging.set_log_cutoff_level` can change the log level. The set value will override the ARI options setting within that session unless the function is called again. This can be a useful debugging tool.

### Stopping and Starting the Backend

The ARI Backend includes these master processes: the Scheduler and EVE. These processes live in the `DBA_JOBS` queue from which they run periodically. The Scheduler also creates Batch Scans that it adds to `DBA_JOBS`, and EVE creates Validator Threads. EVE manages the Validator Threads, and the Batch Scans manage themselves (by simply disappearing from the queue when they are done executing). There are several start and stop methods for the master processes that may be added to database start and stop scripts or executed on a regular schedule, as detailed in the Process Overview.

### Starting the Master Processes

`ari_control_sql.start_ari` calls the start methods for EVE and the Scheduler. It must be run as the ARI master schema owner to ensure that all of the processes will have appropriate privileges when they execute. This is typically placed in a start up script, or, if you prefer to let the Scheduler simply live in the queue at all times, it may only be executed very rarely with the second start method (Start EVE) being the one that is scheduled.

### Starting EVE Only

`ari_control_sql.start_eve` calls the start method for EVE. It must be run as the ARI master schema owner to ensure that EVE will have appropriate privileges when it runs.

### Stopping the Master Processes

`ari_control_sql.stop_ari` calls the stop methods for EVE and the Scheduler. It must be run as the ARI master schema user to ensure that the jobs are actually removed from `DBA_JOBS`. This is typically placed in the database shutdown script, though often daily bouncing just uses a shutdown immediate, so it is placed in the database start script instead to ensure appropriate shutdown before EVE and the other master processes are restarted.

---

---

**Note:** The Validator Threads must have completed and exited the queue (which can have a slight delay after the Stop Process is run) before EVE is restarted.

---

---

When EVE is terminated it signals the Validator Threads to abort and exit the queue, but it can (rarely) take the threads a few minutes to acknowledge the signal. This delay may

be either that, (a) they are queued but not yet running, in which case they must first start running as other threads are killed and become available, or (b) they are in the process of opening a cursor.

Also note that Batch Scans are not removed by this script, which is generally what is desired since Batch Scans removed from the DBA\_JOBS queue mean that those exceptions contained in the scan will not be scanned again until their next scheduled date. Batch Scans can and should survive database outages. Individual scans can be removed using DBMS\_JOB.REMOVE if their execution is meant to be halted (the corresponding database server process will have to be killed as well).

Finally note that if all you typically do is start and stop EVE, leaving the other master processes running, then you will want to schedule the second stop method (that stops EVE only).

## Stopping EVE Only

`ari_control_sql.stop_eve` calls the stop method for EVE. It must be run as the ARI master schema owner to ensure that it will remove the process from DBA\_JOBS correctly.

## Stopping All ARI Processes

`ari_control_sql.stop_all_ari` is the same as stopping the master processes, but also kills all of the Batch Scans in the queue. In production this is almost never used, because you do not want to lose those scans, this is more of a development utility. During development an undesirable scan or set of scans may be created by mistake, in which case, killing all processes may be the simplest thing to do, rather than removing them one at a time. Often this is done in conjunction with a quick database shutdown and restart to save the trouble of finding the appropriate database server processes to kill, which may be necessary since jobs in the queue once started run even if they are removed from the queue.

## Code Generation

Along with EVE, this is the most critical process to ARI. The code generator must be run as the ARI master schema user to ensure that all of the code can be created properly. Also, it must only be run when EVE and its associated Validator Threads have all been removed from the DBA\_JOBS queue. Other operational issues associated with code generation are discussed throughout the Process Overview, such as scheduling it on a periodic basis or running it on demand. Re-summarizing, code generation is tied up very closely with DDL changes and EVE. Whereas users cannot be logged in during DDL changes, they can usually be logged in during code generation. The exception in the above case is when DDL changes in the metadata impact any database objects used by any exceptions or events. In such a case, before the user logins are enabled, the code generator must be run again, after the metadata changes are made.

## Scheduler

The scheduler checks to see which scheduler are currently due to execute and creates Batch Scans, and anonymous Event Revalidation Blocks to be added to the DBA\_JOBS queue. The scheduler is typically started and stopped via one of the many ARI start/stop control processes, but those methods can be accessed stand-alone for custom scripting etc. The scheduler also contains an API for sending signals to set the next execute date for signal-driven schedule to the current date and time, thus causing the schedule to be executed the next time the scheduler checks schedules due for execution (approximately once every five minutes).

## Starting the Scheduler

`ari_auto_scheduler.start_scheduler` adds the `check_schedule` procedure (the Scheduler) to the `DBA_JOBS` queue. This procedure must be run as the ARI master user to ensure that `check_schedule` has the appropriate privileges when it executes. Typically this program is run from one of the many ARI start/stop methods and is not run stand-alone.

## Stopping the Scheduler

`ari_auto_scheduler.stop_scheduler` removes the `check_schedule` procedure from the `DBA_JOBS` queue. This procedure must be run as the ARI master user to ensure that `check_schedule` is correctly removed. Typically this program is run from one of the many ARI start/stop methods and is not run stand-alone.

## Signal-Driven Scheduler Signaler

This API can be added to other programs so that exception detection and/or event reevaluation can be attached to some other thing happening in the operational system. Each signal driven schedule has a signal text string. This string is passed in the appropriate argument of the `ari_schedule_sql.accept_signal` function to update the next execution date of the schedule associated with the signal text to the current date and time. This then causes processes associated with the schedule to be executed when the scheduler checker next executes.

---

---

**Note:** This function contains a commit, so likely placement of it in another procedure is just after a successful completion and commit of the procedure.

---

---

## EVE (Exception Validation Engine)

EVE is the key threading process that governs ARI exception processing. It is typically started and stopped through one of the many ARI start/stop control methods, but the direct APIs for this functionality exist as well.

## Starting EVE

`ari_eve_process.start_eve` adds the `schedule_jobs` procedure (EVE itself) to `DBA_JOBS`. This procedure must be run as the ARI master user to ensure that it has the appropriate privileges when it is executed by the `DBA_JOBS` queue manager process. Typically this program is run from one of the many ARI start/stop methods and is not run stand-alone.

## Stopping EVE

`ari_eve_process.stop_eve` removes the `schedule_jobs` procedure (EVE itself) from the `DBA_JOBS` queue, and signals to the running process that it should terminate. The EVE process sends terminate signals to its Validator Threads and the EVE process self terminates. This procedure must be run as the ARI master user to ensure that it has the appropriate privileges when it executes in the queue. Typically this program is run from one of the many ARI start/stop methods and is not run stand-alone.

## Periodic Purges

ARI uses a significant amount of data in two special areas, tracking events and tracking event history. Closed events that do not have a time-to-live (meaning almost all events – see the Online Help for reasons to create non-zero time-to-live events) are immediately removed from the system once they are closed. Events that do have a non-zero time-to-

live must be periodically purged or they will remain in the system in a closed state indefinitely. Event History is retained after the event itself is closed for at least as long as it is specified in the event type definition, and is only removed by running the history purge process.

## Event Purge

Because there will typically not be many events with a non-zero time-to-live, the need to periodically purge them is small, so the appropriate function `ari_event_instance_sql.delete_expired` need not be scheduled very often. On the other hand, because not many records will need purge, it should not take very long if run often, so it could equally be scheduled to run daily or weekly or even less often depending on the number and kinds of events in your system.

## Event History Purge

Event history should probably be purged daily or at least weekly. Times to execute will vary based on volume of events, but presumably in full production the reason for setting a history interval is to somehow regulate table sizes. For this to work correctly the history tables themselves will need the periodic purging of event history that is older than the number of retention days specified in the event type definitions. The function to schedule for history purging is `ari_event_hist_sql.purge_expired`.

## ARI Alert Notification API

The purpose of this API is to enable displaying a button on the toolbar of an application other than ARI itself. The ARI Alert Viewer button serves two functions: as a gateway for launching the alert viewer and as a visual indicator of the overall contents of the alert viewer. Launching the alert viewer simply means opening that Form. The visual indicator is that the button changes appearance depending on whether or not new alerts exist.

The API is implemented as a Forms library, a package and a table. The Forms library may not be appropriate if the application needing to interface with ARI is not a Forms application, but the substitution of that part of the API for a non-Forms application is an implementation issue not addressed by the product.

## End User Cases

From the end-user standpoint this API has the following cases:

1. If the user is an ARI user with no new alerts the button should be displayed on the toolbar with the mailbox closed icon (ari\_ari0.ico).
2. If the user is an ARI user with new alerts the button should be displayed on the toolbar with the mailbox open icon (ari\_ari6.ico).
3. If the user is not an ARI user the button is not displayed.
4. In either scenario 1 or 2 the button should be pressed to attempt to launch the alert viewer. (This will display an error message if ARI is not truly installed.)
5. Switching from either scenario 1 or 2 to scenario 3 (removing a user as an ARI user) the button will remain until the Form containing the button is closed and re-launched (which should give scenario 3). Attempting to launch the alert viewer from this button after being removed as a user will give an error. This seems annoying but is acceptable because removing an ARI user is very difficult unless the user has no events, and almost never happens in production unless a user has no events. Moreover, by simply exiting and re-launching the form the error is self-correcting.
6. Switching from 3 to 1 or 2 the button will not appear until the Form containing the button is closed and re-launched.
7. Switching between 1 and 2 the icon should change accordingly when focus is changed from the window and back again. Minimizing and then reselecting a window is one way to do this.

## Architecture

ARI provides an API that gets installed in the external Forms application. This API is like a socket. The socket is only able to accept specific values and the ARI plug is constrained to only deliver those specific values.

ARI also provides a tester module that can be configured to simulate all three end-user cases so the external application can test application behavior for all potential plug outputs/ARI interactions. The tester configured to drive the third end-user case (so that no user is an ARI user) is essentially the safety cap mode that will be deployed if the external application, in spite of installing the socket, is never deployed with ARI. This is the socket that is already installed in Oracle Retail's Merchandise Transaction System as mentioned in the Process Details section of this document.

When ARI is implemented the safety cap is replaced with a plug into ARI. This plug must have been tested in ARI to conform to the API specification that returns only three modes of operation corresponding to the first three end-user cases (the fourth case is an extension of one of the first two).

## Implementation

Putting this all together, the socket part of the API is the Forms library. ariiflib.pll was installed in RMS. (That version is full compatible with the other components of the API and does not need to be replaced with the latest version ariif90.pll that is shipped with ARI for use with any other (non-RMS) applications.) The tester part is the stub version of the ari\_interface\_sql package and the ari\_interface\_test table. To plug this into ARI, the non-ARI application simply drops its ari\_interface\_sql package and uses a synonym to the ARI master schema version of this package instead.

---

**Note:** The localized modification of this package body for more sophisticated button icon display criteria, or entirely different return values (requires more API modification) is allowed in spite of statements elsewhere in this documentation that no modifications of this product are supported.

---

To install the socket the external application must add a button to the toolbar (or wherever they want a button) and must invoke the `GET_IS_ARI_USER` function to determine, based on whether the user is an ARI user, whether to display the button or enable its operation. This function should be installed so that the decision whether to display/enable the button is made with some frequency, depending on whether users are actively being added to and removed from the list of ARI users. Adding and removing users is non-trivial so the required frequency for such an action is minimal.

On some event-driven or periodic basis, including immediately after the button is displayed/enabled, the external application must invoke the `REFRESH_ALERT_ICON` function to update the button's displayed icon. The recommended implementation of this function is when the window focus changes (in the Forms When-Window-Activate trigger).

On pressing the alert viewer button the `LAUNCH_ALERT_VIEWER` procedure must be called. If ARI is installed on the same version of Forms as the application with which it will be working, this is a simple `OPEN_FORM` call. Otherwise, this is a remote call through the plug to tell ARI to launch, so the actual launching is external to the non-ARI application.

---

---

## Appendix: ARI Options

The following ARI configuration options are on the ARI options table. These configuration options are critical settings that affect the performance and processing of ARI. However, once they are set, they can largely be ignored. These options are set by a database administrator using Oracle SQL\*Plus doing simple option value updates based on option name.

### **EVE\_NUM\_THREADS**

This is the number of realm queue processes, the validation engine should run at any one time. Generally this is the number of CPUs in the machine + 1, but not more than 32.

### **EVE\_QUEUE\_REFRESH\_INTERVAL**

This is the time slice in minutes given to each realm queue process. Ten minutes is recommended. It may need to be shorter to get reasonable real-time notification of exceptions if there are some very large queues and some very small ones. On the other hand, if the number of large queues is small enough, the entire time slice may only be used by one or two queues. If so, several processes will always be available for other processes, so a longer limit may improve the efficiency of processing those larger queues.

### **INTERNAL\_SCHEMA**

This is the name of the ARI user that is supposed to be created with the main ARI schema. This user has several explicit grants as described in the installation guide. It is in this schema that the ARI generated code and tables reside.

### **MASTER\_SCHEMA**

This is the main ARI schema. All of the static ARI code and tables are owned by this user, as described in the ARI Installation Guide.

### **MAX\_EVENT\_RECURSION**

This value determines how many times an event will reevaluate itself as the result of auto-actions occurring before it raises an error. The value depends on the complexity of the states, an auto-action structure of the defined event types. Five is a reasonable value and no adjustment should be necessary except in the unlikely case of an event that needs to reevaluate legitimately (not though a mistake in the event definition) more than 5 times is created.

### **REEVAL\_STATUS\_LOCKOUT**

This is the number of minutes between attempts to reevaluate events that could not be reevaluated either because they are locked or because their data values from remote data sources could not be fetched (e.g. because of a down database link).

## PRIMARY\_LANGUAGE\_NUMBER

This is the number that identifies the language that will be considered the primary language for a multi-language ARI implementation. The primary language is also the only language in a single language implementation. The default value is based on the assumption that customers will use value 1 to represent their primary language.

## ANALYST\_ADMIN\_GROUP\_ID

This is here for reference by the ARI programs and should not be changed.

## CLOSE\_EVENT\_REALM\_ID

This is here for reference by the ARI programs and should not be changed unless the metadata for the close event action is altered by some mistake.

## CLOSE\_EVENT\_REALM\_ID

This is here for reference by the ARI programs and should not be changed unless the metadata for the close event action is altered by some mistake.

## ERROR\_ADMIN\_GROUP\_ID

This is here for reference by the ARI programs and should not be changed.

## EVENT\_INSTANCE\_PARM\_ID

This is here for reference by the ARI programs and should not be changed.

## EXCEPTION\_CREATE\_DATE\_PARM\_ID

This is here for reference by the ARI programs and should not be changed.

## EXCEPTION\_CREATE\_DATE\_USER\_ID

This is here for reference by the ARI programs and should not be changed.

## LOG\_LEVEL

Determines the log level of messages to be written to the ARI\_ACTIVITY\_LOG table. Valid levels are 1, 2 and 3. 3 is the most verbose logging and 1 only logs at a very high level. The default setting is 2, though this also may be unnecessarily verbose in a production environment.

## FORWARD\_GENERATION\_HOURS

The value is number of hours and determines how far ahead the event and exception support code is generated before the definitions specify they are to become active. A detailed discussion exists in an earlier section on code generation. The default option is 30 hours allows some leeway if the code generator is schedule to run daily.

An alternative to routinely scheduled generation, supplemented by the occasional on-demand generation, might be to generate *only* at the business analyst's request, but the business analyst must then be careful that everything is a new start within the generation window or else it may not get generated. Starting something far in the future either requires a large generation window or a periodic generation so that, as the start time

approaches, the analyst does not risk forgetting that the code generator (if it is not run periodically) needs to be run.

The downside of a large generation window is that, by generating far into the future, you cannot continue to modify a definition during the several days preceding its activation as you could with a 30-hour window and a daily code generation. To render such future-generated, no longer modifiable rules inactive, set the end date equal to the start date and regenerate code before the start date passes.

---

---

**Note:** However, in production, modifications before deployment should not be necessary since all such issues should have been worked out in development.

---

---

Hence the definitions going into the production instance should have a solid foundation and a large generation window may be a perfectly fine approach. The issue is more significant in a development environment.

## RDW\_LINK

This option should contain the name of a database link to the database containing the schema named in the RDW\_OWNER option (see below). This database link must be accessible to the ARI master schema and users of the Metadata form. The link should connect to a user with the “select any table” system privilege. If the RDW\_OWNER schema is located in the same database as the ARI Master schema, then this option should be null.

## RDW\_OWNER

This option should contain the name of the RDW user who will own reports generated for ARI.

## RDW\_PREFIX

This option should contain the prefix that will be used to identify the tables in the Datamart schema that will contain cached report data. This option should be null if no common prefix is being used for Datamart tables.