ORACLE

FUSION MIDDLEWARE
IDENTITY MANAGEMENT

An Oracle White Paper
August 2011

# Oracle Identity Analytics Sizing Guide

ORACLE

## Disclaimer

The following is intended to outline our general product direction. It is intended for information purposes only, and may not be incorporated into any contract. It is not a commitment to deliver any material, code, or functionality, and should not be relied upon in making purchasing decisions. The development, release, and timing of any features or functionality described for Oracle's products remains at the sole discretion of Oracle.

## Introduction

Oracle Identity Analytics is a feature rich identity and access governance solution that empowers users with advanced identity analytics and dashboards, so organizations can proactively monitor, analyze, review and govern user access in order to mitigate risk, build transparency and satisfy compliance mandates.

This document outlines an estimate of hardware and software requirements for deploying Oracle Identity Analytics. A basic deployment scenario is described and scaling recommendations are provided. These recommendations should be considered as guidance while planning product deployment.

Assumptions made in this document are:

- A highly available environment is desired.
- RDMS specific best practices for high availability, backup and recovery are being followed.
- Load balancing specifics, software and hardware, is beyond the scope of this document.

## Architecture Overview

Oracle Identity Analytics is a Java™ 2 Platform, Enterprise Edition (J2EE platform) web application. The J2EE platform consists of a set of industry-standard services, APIs, and protocols that provide the functionality for developing multi-tiered, web-based, enterprise applications. The division of tiers allows Oracle Identity Analytics to scale according to customer's performance demands. Oracle Identity Analytics uses the J2EE specification to build a flexible, scalable and fault-tolerant cross-platform solution. The main tiers of Oracle Identity Analytics are:

- **The Presentation tier** - A web server layer rendering JSPs, JavaScript, XML etc. to present a UI accessible through various supported web browsers,
- **The Logic tier** - A J2EE application server forms the middle tier where all business logic of Oracle Identity Analytics is implemented, and
- **The Data tier** - The data tier usually consists of a standalone or clustered RDBMS environment utilizing Java Database Connectivity (JDBC) to integrate with the logic tier.

The Oracle Identity Analytics application resides on an application server and the central repository of application data resides on a database server. Figure 1 illustrates the architecture of Oracle Identity Analytics. Figure 2 represents sample architecture for deploying Oracle Identity Analytics.

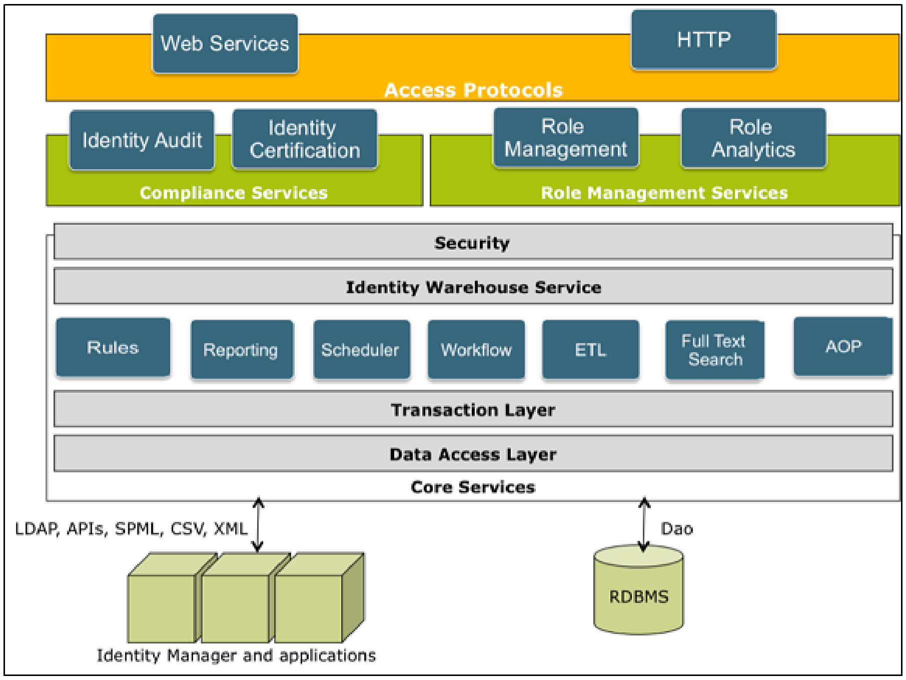

Figure 1. Oracle Identity Analytics Technical Architecture

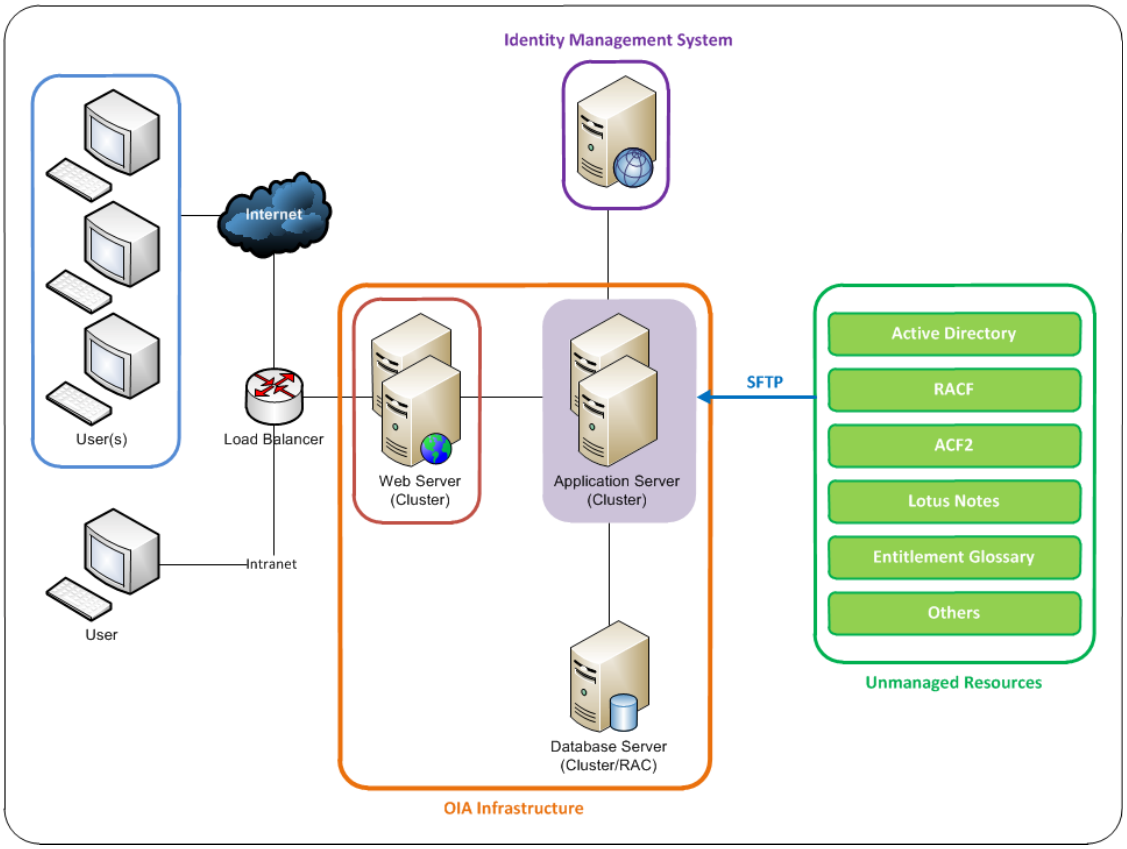

Figure 2. Sample Architecture for Deploying Oracle Identity Analytics

Typical Oracle Identity Analytics deployments comprise of the following components:

- A clustered web server load balanced using a load balancing router. End-users including administrators interact with Oracle Identity Analytics through these web servers.
- A clustered J2EE application server on which Oracle Identity Analytics is deployed.
- Oracle Identity Analytics uses a RDBMS as its data repository. Depending on the dataset size, the database server can be a standalone or clustered, as depicted in this sample architecture, the database is clustered. For optimized performance, the application servers and RDBMS are co-located, for example within the same subnet.
- In most cases, Oracle Identity Analytics is integrated with an Identity Management System. The integration with supported Identity Management Systems is beyond the scope of this document.

- Oracle Identity Analytics utilizes flat files from target systems such as RACF, AD, ACF2 etc. to build its Identity Warehouse. Typically, the target systems drop flat files on a shared location using SFTP, which is subsequently imported using Oracle Identity Analytics import process. Such target systems can be classified as unmanaged resources.

Additional infrastructure such as Single Sign-on servers, proxy servers etc. are not considered as part of the deployment.

# Deployment Considerations

Oracle Identity Analytics performance depends on the load faced and response characteristics of each tier discussed in the previous section. Performance affecting factors are identified and discussed in the following sections. These factors should be considered during deployment planning.

## Oracle Identity Analytics Web Client

The number of concurrent users accessing the system directly affects the web client performance. Performance is also affected by the activities being performed within each user session i.e. role provisioning, attestation, (Segregation of Duties) SoD monitoring, reporting & dashboarding, etc. Concurrent users and their system activities largely affect Central Processing Unit (CPU) and memory requirements of the application server.

## Oracle Identity Analytics Server

The Oracle Identity Analytics server is a J2EE application server that uses J2EE technologies for interaction with end-users, target systems, database repository etc. Following are some areas of server operation that need to be considered during Oracle Identity Analytics sizing.

- Oracle Identity Analytics Import Process - Import jobs are created to populate the Oracle Identity Analytics Identity Warehouse. Data can be imported from a text file or by using direct connections to provisioning systems. Oracle Identity Analytics inserts or updates data in the warehouse, and archives all of the data feeds.

  Importing a large data set can impose resource constraints on the application server e.g. CPU and memory usage, and the database e.g. an increase of the table-space size containing Oracle Identity Analytics repository.

- Oracle Identity Analytics Identity Certification - Identity certification is the process of reviewing user entitlements to ensure that users have not acquired entitlements that they are not authorized to have. Certifications can be scheduled to run on a regular basis to meet compliance requirements. Managers use the Identity Certification module to review their employees' entitlements to access applications and data. Based on changes reported by Oracle Identity Analytics, managers can authorize or revoke employee access, as needed.

Attestation of a data set of large user entitlements can affect Oracle Identity Analytics performance caused by resource constraints on the application server and database.

- Oracle Identity Analytics Identity Audit Process - The Identity Audit module is designed to detect segregation of duties (SoD) violations. A segregation of duties violation is a violation whereby a user account, a user attribute, or a role has been assigned two entitlements that should not be held in combination.

  While the identity certification module enables managers to certify or revoke access of users, the identity audit module has a detection mechanism that monitors users' actual access to resources and captures any violations on a continuous basis. The software can also be programmed to conform to audit policies and report exceptions. It provides a summary of all exceptions, which helps security analysts, executives, or auditors accept or mitigate the exceptions.

  In Oracle Identity Analytics, audit rules define violations. Audit rules are collected together to create an audit policy. User accounts and business structures are then scanned for audit policy violations. User accounts, user attributes, and roles that violate an identity audit policy are flagged and tracked until the violation is resolved.

## Baseline Deployment

A baseline deployment of Oracle Identity Analytics for both the application server configuration and the database configuration is shown below:

TABLE 1. APPLICATION SERVER CONFIGURATION

| | |
|---|---|
| CPU | Two Quad-Core (Minimum) |
| JVM Heap Size | 16 GB (minimum) |

TABLE 2. DATABASE SERVER CONFIGURATION

| | |
|---|---|
| CPU | Two Six-Core (Minimum) |
| RAM | 16 GB (minimum) |

Notes:

- Both Application server and database server operating system are 64-bit (required)

- WebLogic Managed Server (recommended)

- Java 64-bit JVM (required)

- Clustering might require load balancing routers and special networking configurations
- The database configuration may be RAC or non-RAC. It's recommended that customers follow the Oracle Maximum Availability architecture (MAA) recommendations for database high availability

## Deployment Characteristics

The baseline configuration above supports the deployment characteristics shown in the table below.

**TABLE 3. DEPLOYMENT CHARACTERISTICS**

| | |
|---|---|
| Number Global Users | 1000000 |
| OIA Users | 1000 |
| Accounts Per User | 100 |
| Resources | 500 |
| Roles | 5000 |
| Roles Requests Per Day | 200 |
| Policies | 5000 |
| Certifications Per Period | 6000 |
| Concurrent Operations | 100 |
| Identity Manager Integration | None |
| OIA Clustering | No |
| Scheduled Background Jobs | Normal Scheduling |
| Average GUI Response Time | 5 Seconds |
| CPU Utilization | 60% |

## Possible Deployment Scaling Factors

Using the above baseline configuration and deployment characteristics as a starting point, installations can scale their deployments both vertically and horizontally based on variety of factors:

- Anticipated global user population and OIA login users new growth
- Projected increase in  number of concurrent operations, role requests and bulk operations
- Estimated growth of user attribute changes, number of roles and role membership changes
- Change in background job scheduling, both current job frequency increases and additional jobs being brought on line
- Increased volume in Certifications per period
- Integration with Identity Management Systems such as  Oracle Identity Manager
- High availability requirements (clustering)

The items detailed above are just some of the more important factors affecting changes to the Oracle Identity Analytic's deployment baseline.

### Application Server JVM Heap and RAM (Vertical Scaling)

In addition to  changes in the possible deployment scaling factors mentioned above, increasing the JVM heap size in the Application Server baseline server might be required if JVM GC diagnostics show JVM garbage collection activity impacting server performance. Increasing the Application Server baseline server's RAM might be required if OS paging or other OS functions impact Application Server performance.

### Application Server CPU (Horizontal Scaling)

Changes to the Application Server baseline server CPU infrastructure might be required if:

- Concurrent operations increase in number and are take longer to complete
- There are increases in background job scheduling requirements
- Current background jobs are taking longer to complete, or not completing in expected time
- GUI response time degrades
- Average CPU utilization above 60%

Possible solutions to Application Server CPU scaling are:

- Upgrading baseline Application Server CPU count and/or CPU core count
- Create Oracle Identity Analytic cluster and add additional baseline Application Server nodes to the cluster

## Oracle Identity Analytics Database Size Calculation

The following steps can be used to estimate DB size requirements for Oracle Identity Analytics deployment on an Oracle DB Server:

### Calculate Account Objects

To calculate the number of account objects, substitute the corresponding values into the following formula:

> Total Account Objects = Number of Resource Types x Number of Resources x Accounts per Resources

### Calculate Policy Objects

To calculate the number of policy objects, substitute the corresponding values into the following formula:

> Total Policy Objects = Number of Resource Types x Number of Resources x Policies per Resources

Note: Each policy version is an object, add number of policy versions during calculation of policy objects.

### Calculate Total Number of Objects

To calculate the total number of objects:

> Total Number of Objects = Global User Objects + Total Account Objects + Total Policy Objects + Total Role Objects + Total Request Objects

Note: Each role version is an object, add number of role versions during calculation of role objects.

### Determine Database Disk Space Requirement

Objects are typically 120 KB in size, each report and certification is about 4 MB of data, and each Identity Audit (SoD) violation is about 500 KB in size. To calculate the approximate object disk space, substitute the corresponding values into the following formula:

> Approximate Object Disk Space = Total Number of Objects x Size Per Object
>
> Total DB Size = Approximate Object Disk Space + (Number of Certifications Annually x Size per Report) + (Number of Reports Annually x Size per Report) + (Avg. Number of SoD Violations x Size per Violation)

As the number of accounts per user grows, the disk space increases exponentially. The space required depends on the following:

- Number of global users
- Number of accounts per user
- Number of resource types and resources

If Oracle Database Server is being utilized as Oracle Identity Analytics data repository, the automated snapshots using journaling and checkpoint systems add extra hard disk space requirements. Such data recovery constraints must also be factored into the database hard disk free-space requirements when sizing Oracle Identity Analytics implementation.

# Oracle Database Tuning

## Sample Instance Configuration Parameters

The following sample configuration parameter settings are based on a server with four CPUs (64 bit) and 8 or 20 gigabytes (GB) RAM.  SGA, PGA size are limited by the underlying operating system restrictions on the maximum available memory in some platforms. See Support Note: Oracle Database Server and the Operating System Memory Limitations [ID 269495.1]

Note: In the Table below, ASMM denotes the Automatic Shared Memory Management feature available in Oracle Database 10g onward. It automatically distributes the memory among various subcomponents to ensure the most effective memory utilization. You should set the processes parameter to accommodate the Oracle Identity Analytics connection pool requirements configured in the file(s) dataaccess-context.xml and oim-11g-context.xml if Oracle Identity Manager integration is enabled.

TABLE 4. SAMPLE CONFIGURATION PERIMETERS

| PARAMETER | SUGGESTED INITIAL SETTING FOR ORACLE DATABASE 11G |
|---|---|
| db_block size | 8192 |
| Memory Target | Using Automatic Memory Management feature in Oracle Database 11g, the MEMORY_TARGET and MEMORY_MAX_TARGET parameters can be used to manage the SGA and PGA together. Recommended value is 6 GB. For maximum value, use the following formula:<br>MEMORY_TARGET/MEMORY_MAX_TARGET=Total Memory X 80% or 20GB (which ever is greater)<br>Assuming that the computer has the Oracle Identity Analytics database  as the primary consumer.<br>When considering MEMORY_TARGET for managing the database memory components, SGA_TARGET and PGA_AGGREGATE_TARGET can be left unallocated, which is 0. |
| sga_content | If you use ASMM available in Oracle Database 10g onward, then the SGA components can be managed by specifying the SGA_TARGET and SGA_MAX_SIZE parameters. PGA is managed |

| | |
|---|---|
| | separately through PGA_AGGREGATE_TARGET. Use any one of the two memory management approaches depending on the Oracle Database version:<br>• MEMORY_TARGET available in Oracle Database 11g<br>• SGA_TARGET/PGA_AGGREGATE_TARGET ratio available in Oracle Database 10g onward<br>Use Oracle ASMM. Minimum value is 4 GB. For maximum value, use the following formula:<br>SGA_TARGET = Total Memory X 80% X 60% or 16 GB (which ever is greater)<br>Assumptions:<br>• An overall memory cap of 20 GB for the Oracle Identity Analytics database to run<br>• The computer has the Oracle Identity Analytics database  as the primary consumer |
| sga_max_size | 10 GB |
| pga_aggregate_target | Minimum value is 2 GB. For maximum value, use the following formula:<br>PGA_TARGET = Total Memory X 80% X 40% or 4 GB (which ever is greater)<br>Assuming that the computer has the Oracle Identity Analytics database as the primary consumer. |
| db_keep_cache_size | 800 MB |
| log_buffer | 15 MB |
| cursor sharing | FORCE |
| open_cursors | 2000 |
| session_cached_cursors | 800 |
| qyery_rewrite_integrated | TRUSTED |
| db_file_multiblock_read_count | 16 |
| db_writer_process | 2 |
| Processes | Based on connection pool settings (see note above) |

Note: These memory parameter values are ball-park figures. As a database administrator, you can also refer to the memory advisors to manage and tune the database.

## Redo-Log File

Depending on the number of certifications configured in Oracle Identity Analytics, the volume of database transactions and commits during certification periods can be high. In those situations, it is recommended you use multiple redo-log files. The total allocated redo-log space should be 1GB to 2GB.

Oracle recommends use of at least three redo log groups with redo log members of at least 500 MB in size. The multiplexing and the exact number of members and disk space for each member can be considered in accordance with the planning for failure

### Other Oracle Database Tuning Resources

See Oracle Identity Manager Performance Tuning and Best Practices Guide (chapter 24).

## Conclusion

Oracle Identity Analytics is a scalable, cost-effective, secure and flexible identity and access governance solution that could be tailored as per requirements. Oracle Identity Analytics 11*g* offers significant architectural innovations that allow the solution to scale horizontally and vertically and provide high performance and availability. The combination of Oracle Identity Analytics' innovative use of standards-based technologies and Oracle Database ensure that customers following the sizing guidelines achieve the optimal deployment.

Oracle is committed to developing practices and products that help protect the environment

Oracle Identity Analytics Sizing Guide
August 2011
Author: Mike Dugan, Gustavo Faerman

Oracle Corporation
World Headquarters
500 Oracle Parkway
Redwood Shores, CA 94065
U.S.A.

Worldwide Inquiries:
Phone: +1.650.506.7000
Fax: +1.650.506.7200

oracle.com

Hardware and Software, Engineered to Work Together