



An Oracle Technical White Paper

August 2017

Oracle Identity Governance 12cPS3 Sizing Guide

Introduction	3
Deployment Architecture Overview	4
Sizing Guidelines	6
Sizing Considerations for OIG	8
UI Operations	8
Provisioning.....	8
Reconciliation	9
Certification	9
Request.....	10
OIG Middle Tier Sizing Internals	10
Sizing Considerations for Database Tier	11
Sizing Guidelines for DB Storage Capacity.....	11
Sizing Guidelines for Workload Capacity – H/W Resources	12
OIG DB Sizing internals	12
Factors impacting OIG DB growth	13
Sizing Calculator	14
Usage Parameters.....	14
Test System Configuration	19
Deployment Size	20
Small, Medium and Large Size Deployments	20
Very Large Size Deployments	21
Sizing for Non Production Environments	21
Deployment Types	22
Identity Administration Only Installations	22
Identity Administration & Provisioning Installations	22
Identity Administration & Certification Installations.....	22
Identity Administration, Provisioning and Certification Installations	22
Sizing for Connector Server	23
Recommendations	23

INTRODUCTION

This document is designed to provide capacity planning and environment sizing estimates based on the computing power required for deploying Oracle Identity Governance (OIG) 12cPS3 components. The document is based on the reference deployment provided in the Deployment Architecture Overview section. A simplified version of the architecture described in the Oracle Identity Management Enterprise Deployment guide. The document focuses on the Oracle Identity Governance components that are significant for sizing.

The Sizing Considerations section explains the various factors that can affect the hardware required for your OIG installation.

The *Deployment Templates and Deployment Type* section provides recommendations for the most common deployments. It allows you to extend sizing for special considerations using sizing parameters.

These sizing guidelines are based on the in-house sizing tests and are intended to be used as a starting point in determining the configuration for a given deployment.

LEGAL DISCLAIMER - Expected performance characteristics are based on laboratory test implementations and are provided as is without warranty of any kind. The entire risk of the use of information provided herein remains with the recipient and in no event shall Oracle be liable for any direct, consequential, incidental, special, punitive or other damages including without limitation, damages for loss of business profits, business interruption or loss of business information.

DEPLOYMENT ARCHITECTURE OVERVIEW

The Diagram below shows simplified deployment architecture of Oracle Identity Governance. Emphasis is on the essential processes to be sized for an OIG deployment.

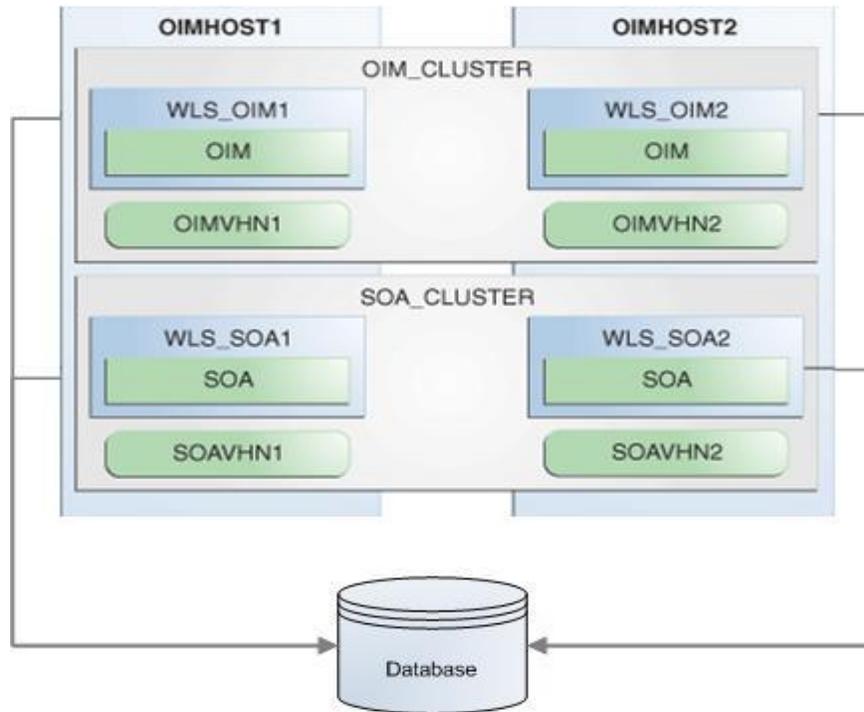


Fig1: Simplified OIG Deployment Architecture

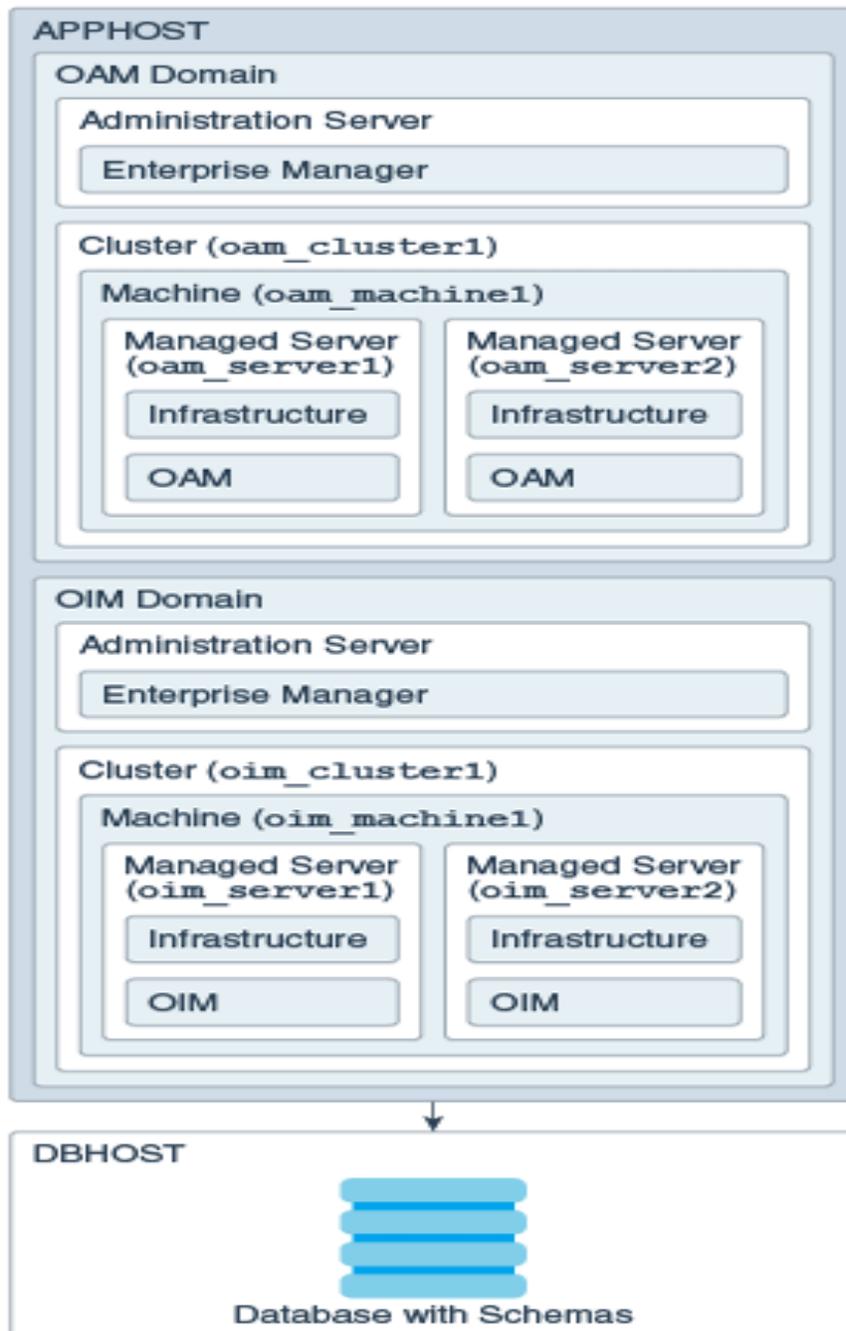


Fig2: Simplified OIG Deployment

From a sizing perspective OIG and the Database are the most important processes. OIG and SOA are recommended to be installed in clustered deployments. This document helps determine how much computational capacity is required for a specific installation to meet its usage. Capacity may be materialized using multiple physical servers, virtual servers, and clusters.

Other administrative processes that are also part of OIG deployment are: Weblogic Admin Server and BI Publisher etc, are considered as fixed cost processes and not addressed by this sizing guide. Minimum hardware requirement for these processes is given in their respective document. These processes may be deployed alongside OIG/SOA processes after augmenting hardware with recommended capacity.

High-availability, disaster recovery and peak load variation requirements are subjective topics. These are not considered in base-line sizing recommended by this document. Hardware should be augmented to meet these requirements. Disaster recovery typically means that the production configuration is duplicated in a different location

High availability can be configured with various levels of service assurance. Depending on acceptable response time and throughput requirements, the customer should decide how many additional nodes are required in the topology.

SIZING GUIDELINES

There is no single guideline in infrastructure sizing. Many factors influence the ultimate outcome including

- Network performance
- Deployment complexity
- Number and type of resources
- Database performance.
- Operating System performance

Customers should also note that additional tuning may be required, depending upon the selected hardware, to get optimal performance. Please refer performance tuning guide for the same.



Fig2: OIG Sizing Engineering

SIZING CONSIDERATIONS FOR OIG

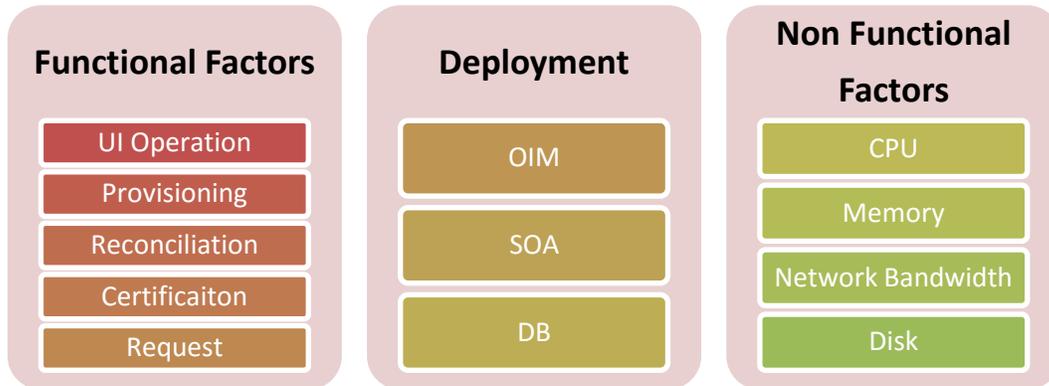


Fig3: Factors Effecting Sizing in OIG

There are 5 main functional areas of OIG which affect sizing in OIG.

- UI Operations
- Provisioning
- Reconciliation
- Certification
- Request

We need to size the deployment to handle maximum load in any of these areas at a given point of time.

UI OPERATIONS

Users need to log into OIG to perform most of the operations like User Life Cycle Management, Role Life Cycle Management, Organization Life Cycle Management, Self Service Operations, Admin Operations, Searches and Requesting/Approving/Certifying Access.

Impact of UI operations on sizing depends upon these factors:

- *Number of concurrent users logged into the system at a same time*
- *No of screens navigated by the user*
- *Think time – Average Time spend on a single screen*

PROVISIONING

Provisioning is the process by which operations such as user creation, modification or deletions initiated by OIG are communicated to a target system via OIG connectors.

Following operations will generate load on the provisioning engine:

- *Self Service Operation – Self Profile Update, Self Password Change, Catalog Based Request for Account, Role or Entitlement*

- *Admin/Help Desk Operation –*
 - *User Profile Update*
 - *Password Reset*
 - *Account , Role and Entitlement provisioning operation*
- *Access Policy driven provisioning operations. Inclusive of SOD checks for preventive validation.*
- *Changes coming into a system via Trusted reconciliation will generate downstream provisioning*
- *Closed-loop remediation*

RECONCILIATION

Reconciliation is the process by which operations such as user creation, modification or deletions started on the target system are communicated to Oracle Identity Governance. The reconciliation process compares the entries in Oracle Identity Governance and the target system repository, determines the difference between the two repositories and applies the latest changes to Oracle Identity Governance.

Impact of reconciliation on sizing depends upon these factors:

- *Number of authoritative sources*
- *Number of target resources*
- *Number of accounts per resource*
- *Frequency of reconciliation*
- *Degree of change in all resources (authoritative and target)*
- *Amount of pre- or post-processing of data involved e.g. transformation and validation logic*

CERTIFICATION

Access Certification is an on-going process that allows organizations to manage risk by identifying, certifying, and modifying user access to vital corporate resources and demonstrate compliance with corporate policies and key industry and government regulations. Certification feature affects the following areas in the system:

- **Increase in Backend Operations**
Certification creation is a backend schedule job, so it will increase the load on the server when the certification campaign is initiated.
- **Increase in User Concurrency**
The number of users concurrently logged into the system will increase towards the end of the campaign. As most of the people will tend to complete their certification campaign only at the last moment, customers will see maximum concurrency in the system during the end of certification campaign cycle.
- **Increased Load on SOA Server**
The Certification campaign will increase the load on the SOA server as every certification action generates a SOA task flow.
- **Increased Load on Provisioning Engine.**
There will be some load on the provisioning engine when a certification request will get translated into closed-loop remediation.
- **Growth in Database Size**
There will be substantial data growth due to certification generated entities, user actions and audit data.

Impact of Certification on sizing depends upon these factors:

- *Frequency of Certification*

- *Type and Complexity of certification*
- *Customizations of certifications*
- *Number of Managers, Data Owners, resource owner and Total Number of Users in the System.*
- *Total number of resources (Application Instance, Business Roles, Entitlements) in the system*
- *Average Number of resources (Application Instance, Business Roles, Entitlements) per user*

REQUEST

Oracle Identity Governance leverages the rich workflow capabilities of BPEL and Human Workflow components of the Oracle SOA Suite, referred to as SOA hereafter, to carry out approvals and workflow routing. End users can request access to resources, roles and entitlements using OIG catalog. There are two areas in the Middle Tier which will be impacted by request operations:

OIG Server – OIG processing is required to generate request payload and send it to the SOA server

SOA Server - Handles the incoming payload through the BPEL process and calls appropriate human tasks and business logic.

Following areas will impact the sizing considerations at both layers:

- *Number of request generated during the maximum usage time.*
- *Size of the payload data generated in OIG and processed by SOA.*
- *Number of workflows/BPEL processes deployed in SOA.*
- *Complexity of workflows – serial, parallel, notification, escalation, external calls etc.*
- *Number of people concurrently accessing SOA directly to perform an action.*

OIG MIDDLE TIER SIZING INTERNALS

Every OIG operation mentioned above consumes system resources in terms of CPU cycles and Memory.

CPU

We estimated the CPU requirements for each functional area in OIG by putting maximum load in the system that can meet our SLA and recorded the units of CPU required for performing high load operations.

The formula used to calculate CPU Core required for each functional area is as follows:

Middle Tier Cores = \sum [(CPU unit cycle required for per unit operation) * (Total number of operations)*(CPU Load factor)]

Where – The total number of transactions per feature is a derived value based on the inputs to the question in the Sizing calculator; The CPU load factor is based on input variables in our sizing calculator.

UI concurrency, requests and certifications also factor in dependent components such as SOA for CPU usage.

MEMORY

Memory requirements are primarily driven from two factors:

- Type of Deployment
- Concurrency

DISK SPACE

Middle tier generally requires static configuration data and binaries which need static allocation of disk space.

SIZING CONSIDERATIONS FOR DATABASE TIER

A key challenge for Oracle Database administrators is to effectively design and manage system storage, especially to accommodate performance, capacity and future growth requirements.

The goal of this paper is to present the sizing/capacity planning for an Oracle Identity Governance deployment based on the tests conducted in-house, and provide sizing guidelines and best practices based on the observations made during internal benchmarking runs.

There are two focus areas for DB Sizing:

- Storage—Database size
- Workload—CPU, memory, and storage performance (IOPS/MBPS)

SIZING GUIDELINES FOR DB STORAGE CAPACITY

When sizing storage capacity for OIG DB, you need to consider the following components:

- Database files
- Online REDO log files
- UNDO data
- Archive log files
- Backup/flash data

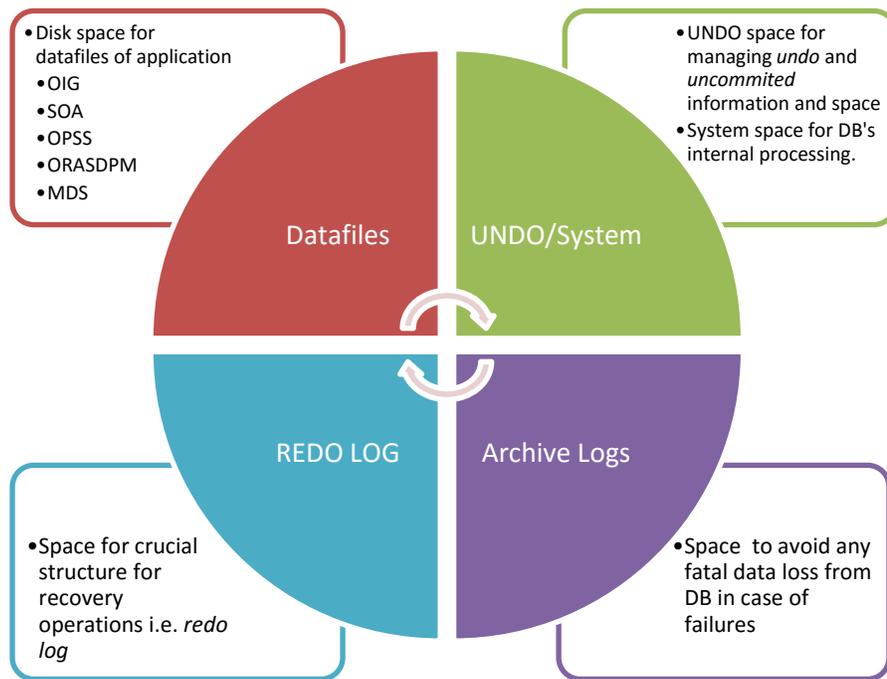


Fig4: DB Sizing Factors

DB components	Space Requirements	Remarks
Database Files (Datafiles)	As per OIG Sizing calculator	Space required for datafiles associated with its tablespaces.
Temporary Table space	IDM's Repository Creation Utility creates(RCU) TEMP tablespace for every component by default.	Based on database workload. If application operations involve Sequential I/O then you may need a larger allocation.
UNDO/ SYSTEM Space	As per DBA's recommendation	The customer's DBA should review the application nature/load.
REDO Logs	As per FMW Tuning Guide's recommendation for OIG	OOTB size is 50 MB per REDO Group OIG recommends 500M to 1G depending on the loads.
Archive Logs	DB_Size * 1.0	The customer's DBA should review this and alter based on the environment.
Backup	DB_Size * 2.0	The customer's DBA should review this and alter as per the environment.

***DB_Size = Datafiles + Temp Table Space + UNDO/SYSTEM Space + REDO LOGS**

SIZING GUIDELINES FOR WORKLOAD CAPACITY – H/W RESOURCES

There are other significant factors at DB layer which also contribute in DB sizing like CPU, memory, and storage performance (IOPS/MBPS). OIG is an enterprise product which relies heavily on a database as primary repository. Oracle is a database with rich features providing strong RDBMS capabilities so it needs high resources from a CPU/memory perspective. The following functional areas have a high impact on the database resources:

- Reconciliation
- Certification
- Provisioning
- Request catalog
- Customization which cause DB hits

OIG DB SIZING INTERNALS

Every OIG operation contributes to DB growth. We started loading the system until it reaches the maximum stress level where it can support defined SLAs and recorded the DB size trending for each functional area.

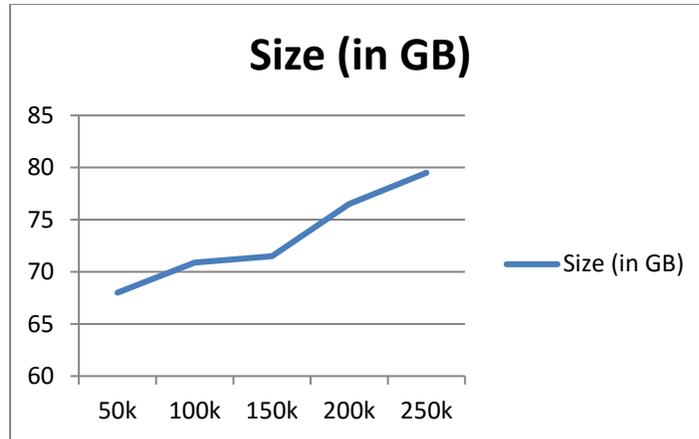
The formula used to calculate DB Size is as follows:

OIG DB = [(∑ Sum of DB objects space consuming high space) + (∑ Sum of space consumed by rest of DB objects)]

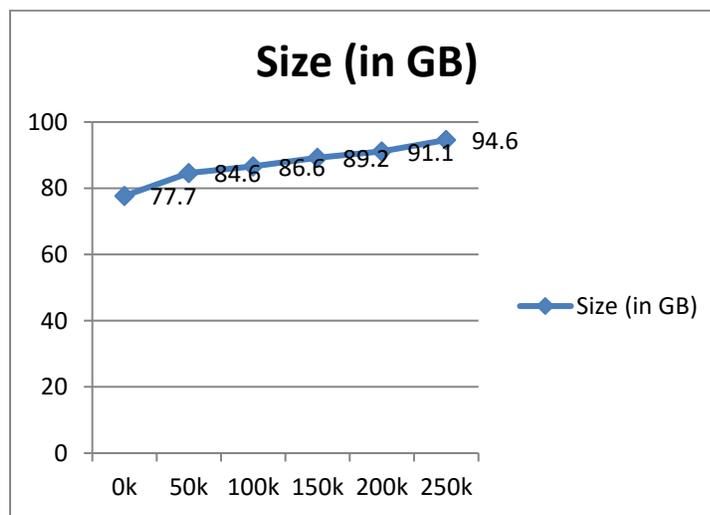
All the operations impact disk space consumed by dependent components such as SOA, MDS, and OPSS.

Below is the DB growth trending we noticed in our sizing runs.

- Trusted reconciliation



- Target reconciliation



FACTORS IMPACTING OIG DB GROWTH

1. **Data load on system** - As the number of users, roles, accounts and entitlements in the system grow, it influences the response time and resource requirements.
2. **Audit Level** – is one of the significant factors in the DB growth. If we make audit level more granular, the DB growth/size will be high. This document provides guidelines based on the default audit setting, i.e., resource-form.

Customers can control auditing granularity by setting the relevant auditing level parameters. At its finest level, the auditing engine will audit the following:

- User profile changes

- User configuration changes
- Role profile change
- Account definition changes
- Account data changes
- SOD validations

From a sizing perspective, customers must consider the impact to performance and data growth when setting audit levels.

3. **Reconciliation/Certification Frequency** – The higher the reconciliation/certification frequency, the greater the data growth.
4. **Archival/Purge policy** - It is important to follow OIG's archival and/or purge recommendations for all the entities whose purge solution is available like reconciliation, request, provisioning tasks, audit, and orchestration.

SIZING CALCULATOR

An Interactive sizing calculator based on Microsoft Excel spreadsheet has been built based on the formulas defined from the analysis of the test environment results. The calculator is included in this zip file.

The end user needs to supply answers to the business questions related to their environment to calculate the sizing results.

USAGE PARAMETERS

The end user is expected to fill the following business questions in the calculator:

Section1: User

- Total number of active users in the system

Consider External and Internal Population before filling the sheet

External Users will not be there in all target systems. We do not need to reconcile them.

***For Example,** if your total user base is 10M and out of it only 500K users are internal users like employees, contractors and rest of the users are consumers or external users who do not have access to internal corporate resources then we need to supply following parameters in the calculator as a sample:*

Number of users = 10000000

Maximum number of users per target = 400000

Average number of users per target = 50000

- Concurrent logins at peak usage time
- Average number of resource / account /app instance per user

This is the average number of target system a user is provisioned to.

- Average number of roles per user
- Average number of entitlements per user

This is the average number of entitlements a user has in the target systems.

For Example, if there are two target systems like AD and LDAP, each user in the system has 4 AD groups, 2 LDAP roles on an average then the value of this field is 6.

Section2: Resource/Targets

Illustration –

Following is an example to understand how to fill data in this section:

Five target systems with the following data:

Target System1 = {10K entitlements, 20K users, 100 changes in target per day}

Target System2 = {20K entitlements, 30K users, 200 changes in target per day}

Target System3 = {30K entitlements, 50K users, 300 changes in target per day}

Target System4 = {40K entitlements, 70K users, 400 changes in target per day}

Target System5 = {50K entitlements, 90K users, 500 changes in target per day}

The following inputs should be provided to the calculator:

- Total number of targets = 5
- Maximum active users in a target = 90000
- Average number of entitlements per target = 30000
- Average active user population per target = 52000
- % Changes per day on each target system to be consumed by OIG = 10

This is average change in percentage which directly occurs on a single target which needs to be consumed by OIG every day.

Section 3: Roles

Let's take an example to understand how to fill data in this section:

If there are 20 access policies in the system and each policy has an average 2 IT resources and 5 child data associated with it then the following inputs should be provided to the calculator:

- Total number of business roles defined in OIG = 20

This is the total no of roles which will be used to define access policies.

- Average resource and entitlement associated per business Role = 7

This is the number of Application Instance and Entitlements associated per access policy.

Section4: Organization

- Total number of Organizations

This is total number of Organizations in OIG.

Section5: Self Service/Admin/Help Desk Operations

- Number of admin operations per day

This is average number of Admin Operations like Create/Update/Enable/Disable Users and accounts/Role/Organization, Password Reset etc.

- Number of self service operations per day.

This is average number of self service operations like Self Registration, Profile Modify and Password Management.

Section6: Request

- Average number of requests and approvals per day.

This is total number of access requests and request approval generated by the system in a day.

- Maximum number of requests per hour during peak time.

This is the maximum number of access requests generated in one hour during worst case situation.

Section7: Certification

- Certification feature is enabled or disable (Y for Enabled and N for Disabled)

- Frequency of certification campaign in a Year/Bi-Yearly/Quarterly

Number of times certification campaign is triggered.

If an organization runs quarterly certifications then the value of this field will be 3.

If an organization runs half yearly certifications then the value of this field will be 2.

If an organization runs yearly certifications then the value of this field will be 1.

- Types of Certifications:-
- Do you do manager certification?
- Number of managers in the system
- Do you do data/Resource owner certification?
- How many data/Resources owners are there in the system?

Section8: System SLAs

Average acceptable memory utilization (in %)

Average acceptable CPU utilization (in %)

Under a steady load of transactions what percentage of CPU/Memory usage is acceptable is stated as CPU/Memory load factor.

Section9: Data Retention Policy

- Number of days after which request data will be purged.
- Number of days after which task data will be purged.
- Number of days after which recon data will be purged.
- Number of days after which certification data will be purged.

Customer's data retention policies play a very important role to predict the data volume and OIG enables customers to set it according to their data policy.

Section10: Business Growth Pattern

- Business growth per year in percentage.

This is the percentage increase in active users per year.

- No of years you want to predict DB growth?

Section11: Identity Audit - Detective

- Detect existing compliance issues by launching detective scan for a list of violations against users.
- Define rules and policies to perform SOD checks.

This is not a complete list of variables. Some installations include use cases with a very high volume of self-service operations, while other may contain full user reconciliation that runs very often. For such extreme use-cases the recommended option is to estimate hardware requirements based on the relevant performance testing and measurements for the solution.

TEST SYSTEM CONFIGURATION

Below Table provide details for the major hardware and software system components used in the test system configuration.

Test Configuration – Hardware Components	
OIG/SOA Servers	<ul style="list-style-type: none">• 2 Intel machines• Intel(R) Xeon(R) CPU 5148 @ 2.33GHz processors• 32 GB RAM, 4M Cache, 2.33 GHz• 2 Dual core CPU per machine i.e. 4 CPU• NIC Card - 2 Broadcom NetXtreme II BCM5708 1000Base-T PCI-X 64-bit 133MHz
Database server	<ul style="list-style-type: none">• 1 Intel machines• Intel(R) Xeon(R) CPU E5410 @ 2.33GHz• 16 GB 6 4M Cache, 2.33 GHz• 4 Dual core CPU per machine i.e. 8 CPU• NIC Card - 2 Broadcom NetXtreme II BCM5708 1000Base-T PCI-X 64-bit 133MHz
Storage	<ul style="list-style-type: none">• SAS disks• scsi0 Channel• Vendor: FUJITSU Model: MBB2147RC

Test Configuration – Software Components	
OIG/SOA Servers	<ul style="list-style-type: none">• Oracle Enterprise Linux 5• Weblogic 12.2.1.3• Oracle HTTP Server 12.2.1.3
Database server	<ul style="list-style-type: none">• Oracle Enterprise Linux 5• Oracle RDBMS 12.2.1.3• Single instance
Monitoring Tools	<ul style="list-style-type: none">• JVM utilities like JConsole• JProfiler• FMW Enterprise manager• Oracle Database Enterprise manager• OS utilities like vmstat etc• DMS (Dynamic monitoring service)• AWR (Automatic workload repository)• JRockit utilities like JMRC/JFR• Custom made data collection tools• Custom made DB sizing framework

DEPLOYMENT SIZE

SMALL, MEDIUM AND LARGE SIZE DEPLOYMENTS

Metric	Small Size	Medium Size	Large Size (All Internal)	Large Size (Internal + External)
Total Number of Active users in the system	50K	150K	500K	500K (150K Internal + 350K External)
Concurrent logins at peak usage time	25	100	250	250
Average number of Resource per user	2	5	10	10
Average number of Roles per user	2	5	10	2
Average number of Entitlements per user	5	10	20	5
Total Number of target resources	5	15	30	30
Average no of Entitlements per resource	10K	20K	30K	30K
Maximum active users in a Single Target	10K	100K	300K	100K
Average active user population per target	5K	50K	100K	50K
% Changes per day on each target system	1	1	1	1
Total Number of Business Roles defined in OIG	5	20	40	40
Average Resource/Entitlement Associated per Business Role	5	5	5	5
Total Number of Organizations	1	5	20	20
Self Service Operations per day	100	500	1000	2500
Number of Admin Operations Per day	40	200	400	400
Number of requests per day	50	200	1000	200
Number of Approvals per day	20	100	300	100
Certification Feature is Enabled or Disable	No	Yes	Yes	Yes
Frequency of Certification Campaign in a Year	N/A	2	3	3
Do you do manager certification?	N/A	Y	Yes	Yes
Number of Managers in the system	N/A	20K	50K	20K
Do you do data owner certification?	N/A	No	Yes	Yes
How many Data Owners are there in the system?	N/A	N/A	20K	5K

Sizing Recommendations

		Small Size	Medium Size	Large Size (All Internal)	Large Size (Internal + External)
OIG/SOA					
	CPU	4 Core	8 Core	20 Core	20 Core
	OIG Heap	4 GB	4 GB	4 GB	4 GB

	SOA Heap	1 GB	4 GB	8 GB	8 GB
	OIG/SOA Nodes	1Node / 2 Node	2 Nodes	5 Nodes	5 Nodes
Database					
	CPU	2 Core	6 Core	12	12
	Memory	4 GB	6 GB	6 GB	6 GB
	OIG Disk Space [1 st Year]	20 GB	103 GB	~3 TB	~ 692 GB
	OIG Disk Space [3 rd Year, With 5% Y2Y Growth]	~30GB	~150GB	~4.9 TB	~ 1 TB
	Database Nodes	1 Node	1 Node	3 Node RAC	3 Node RAC

VERY LARGE SIZE DEPLOYMENTS

The results are based on the lab results to calculate the hardware requirements. However, when you have a very large deployment, the formulas may not work. Installation with multi-million users or very high number of active users should be considered as a large deployment.

Use of CPU is mostly a fraction of the number of operations performed in the system. As such an installation may need special considerations if it expects very high concurrency, very complex pre-/post-processing logic or very high frequency / volume of reconciliation events.

One way to size such a deployment is to execute performance tests and derive recommendations from it. These systems also need to be monitored for any unusual resource usages.

SIZING FOR NON PRODUCTION ENVIRONMENTS

Following table provide minimum sizing recommendations for Non-Production environments:

Type	OIG/SOA	DB
DEVELOPMENT	Single Node, 2 Core CPU, 4 GB Heap OIG, 1 GB Heap SOA	Single Node, 4 GB RAM
QA	Two Node, 2 Core CPU, 4 GB Heap OIG Per Node, 1 GB Heap SOA Per Node	Single Node, 4 GB RAM
UAT (User Acceptance /Performance and Load Testing Environments)	Replica of Production	Replica of Production

DEPLOYMENT TYPES

IDENTITY ADMINISTRATION ONLY INSTALLATIONS

When a deployment is only addressing identity Administration use-cases, it has less overhead for integrating with other applications. Also, the number of screens used by a logged in user will be limited to only user, role and organization management. Also, there may be less reconciliation post processing requirements. Such Deployments are mostly common in a deployment where OIG is deployed for external users. In such deployments self registration, self profile update and password reset/forgot password operations are very high.

IDENTITY ADMINISTRATION & PROVISIONING INSTALLATIONS

When a deployment is addressing both identity Administration use-cases, and provisioning use-cases it includes more back-office processing to integrate with other applications. Also, the number of screens used by a logged in user will depend on supported target system types.

IDENTITY ADMINISTRATION & CERTIFICATION INSTALLATIONS

When a deployment is addressing identity Administration and certification use-cases but not provisioning use cases, it has less overhead for integrating with other applications because most of the applications are modeled as disconnected resources. But certification creation puts a substantial load on the back-office processing.

IDENTITY ADMINISTRATION, PROVISIONING AND CERTIFICATION INSTALLATIONS

When a deployment is addressing all use-cases including identity Administration, provisioning and certification then it requires maximum resources. In such deployment, the number of screens used by a logged in user will be maximum to include user, role, organization management, request and certification. There will be a substantial load on back-office processing for certification creation and reconciliation.

Follow the table below to fill the calculator for various installation types:

Metric Section /Deployment Type	IDM Only	IDM and Provisioning	IDM and Certification	IDM, Provisioning and Certification
Users	Yes	Yes	Yes	Yes
Resource/Targets	No	Yes	Yes	Yes
Roles	Yes	Yes	Yes	Yes
Organization	Yes	Yes	Yes	Yes
Self Service /Admin/Help Desk Operations	Yes	Yes	Yes	Yes
Request	No	Yes	No	Yes
Certification	No	No	Yes	Yes
Minimum Memory Recommendations				
OIG (in GB)	4	4	4	4
SOA (in GB)	1	2	3	4

SIZING FOR CONNECTOR SERVER

An identity connector server is required when an identity connector bundle is not directly executed within your application. By using one or more identity connector servers, the ICF architecture permits your application to communicate with externally deployed identity connector bundles. Identity connector servers are available for Java™ and Microsoft .NET Framework applications.

A single connector server can support multiple ICF connectors, and these ICF connectors may be of different connector types. A single ICF connector can be used to communicate with multiple targets.

Following table provide minimum sizing recommendations for Java and .NET Connector Server in production:

	Java Connector Server	.NET Connector Server
CPU	2 Core CPU	2 Core CPU
RAM	4 GB	4 GB

RECOMMENDATIONS

1. It is recommended that customers keep a 10-20% spare capacity to allow the deployment to handle out-of-bound situations, rather than operating the deployment at capacity.
2. Admin Server contains Weblogic admin applications, Enterprise Manager and Diagnostic Monitoring Services (DMS). Some additional capacity is required for it in each deployment. Typically Admin Server requires minimum 1GB Heap but it is recommended to allocate 2 Core CPU with 2GB heap for medium and large scale deployments.
3. Deployments should factor in Disaster Recovery scenarios where the cluster node(s) may be required to handle the load of the failed nodes.
4. Core sizing given in this document is based on hardware component specification given in the Test System Configuration section. You should adjust the output by comparing the Test System Configuration with the actual H/W used for the following factors:
 - Processor Chipset
 - Network Bandwidth
 - RAM Type
 - DISK TYPE
5. For Performance Tuning recommendations please see *Oracle® Fusion Middleware Performance and Tuning Guide*.
6. For DB and Archival Policy recommendations please see *Oracle® Fusion Middleware Administrator's Guide for Oracle Identity Governance*.
7. For High Availability and EDG recommendations please see *Oracle® Fusion Middleware High Availability Guide*.



Oracle Identity Governance Sizing
Guide12cPS3

August 2017

Authors: Kamal Narayan

Oracle Corporation

World Headquarters

500 Oracle Parkway

Redwood Shores, CA 94065

U.S.A.

Copyright © 2014, Oracle and/or its affiliates. All rights reserved. This document is provided for information purposes only and the contents hereof are subject to change without notice. This document is not warranted to be error-free, nor subject to any other warranties or conditions, whether expressed orally or implied in law, including implied warranties and conditions of merchantability or fitness for a particular purpose. We specifically disclaim any liability with respect to this document and no contractual obligations are formed either directly or indirectly by this document. This document may not be reproduced or transmitted in any form or by any means, electronic or mechanical, for any purpose, without our prior written permission.

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

SOFTWARE. HARDWARE. COMPLETE.