



# Best Practices for Database Consolidation



A Guide for Implementation  
Including MAA Reference Architectures

January 2026 | Version [5.1]

Copyright © 2026, Oracle and/or its affiliates

Public

## Disclaimer

This document in any form, software, or printed matter, contains proprietary information that is the exclusive property of Oracle. Your access to and use of this confidential material is subject to the terms and conditions of your Oracle software license and service agreement, which has been executed and with which you agree to comply. This document and information contained herein may not be disclosed, copied, reproduced, or distributed to anyone outside Oracle without prior written consent of Oracle. This document is not part of your license agreement, nor can it be incorporated into any contractual agreement with Oracle or its subsidiaries or affiliates.

This document is for informational purposes only and is intended solely to assist you in planning for the implementation and upgrade of the product features described. 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 in this document remains at the sole discretion of Oracle.

Due to the nature of the product architecture, it may not be possible to safely include all features described in this document without risking significant destabilization of the code.

## Table of Contents

---

<b>Disclaimer</b>	<b>2</b>
<b>Executive Overview</b>	<b>4</b>
<b>Overview of Database Consolidation</b>	<b>5</b>
<b>Why Exadata for Consolidation</b>	<b>11</b>
<b>Establishing the Foundation for Consolidation</b>	<b>16</b>
<b>Oracle AI Database 26ai</b>	<b>18</b>
<b>Exadata Exascale</b>	<b>19</b>
<b>MAA Reference Architectures</b>	<b>19</b>
<b>Applying MAA to Database Consolidation</b>	<b>22</b>
<b>Measuring Capacity Utilization Before Consolidation</b>	<b>28</b>
<b>Resource Management for Consolidation</b>	<b>32</b>
<b>Implementing Database Consolidation</b>	<b>53</b>
<b>Conclusion</b>	<b>60</b>
<b>References</b>	<b>63</b>

## Executive Overview

Information Technology (IT) organizations are under pressure to manage costs, increase agility, ensure availability, reduce security risks, and deliver the proper level of system performance to meet business objectives. Consolidation of databases and other systems has emerged as a common strategy used by IT organizations to meet these objectives. This paper outlines best practices for consolidation of Oracle databases, including use of key technologies including O/S virtualization, resource management, Oracle's Multitenant database capabilities, Oracle Exadata Database Machine, Oracle Exadata Database Service on Cloud@Customer, Oracle Exadata Database Service in Oracle Cloud and in Multicloud (DB@Azure, DB@Google Cloud, DB@AWS), and Oracle Autonomous Database.

The Oracle Exadata Database platform provides the ideal platform for the consolidation of databases at lower cost, with simplified management, industry-leading availability, enterprise-grade security, and performance capabilities that can be employed and managed to meet business objectives. Exadata incorporates the full range of Oracle Maximum Availability Architecture (MAA) best practices and is combined with standardized operational practices, enabling customers to achieve the highest levels of reliability possible in a consolidated database environment.

Exadata can be deployed in traditional on-premises environments, in the Oracle Public Cloud, in Multicloud environments (Azure, Google Cloud Platform, and AWS), and the Oracle Cloud@Customer model. Database consolidation, as outlined in this paper, applies to all of these deployment models. In addition, Oracle Autonomous Database also runs on Exadata and automatically incorporates the principles outlined in this document, allowing customers to consolidate databases with greater ease.

This paper will explore these topics in detail to give IT organizations the best possible guidance for consolidating Oracle databases.

## Overview of Database Consolidation

This paper outlines best practices for the consolidation of Oracle databases to achieve the primary business objectives of managing costs, improving manageability of databases, ensuring availability, improving security, and delivering the proper level of performance for business processes that rely on Oracle databases. While many of the recommendations included in this document can be applied to any platform, special attention will be given to Oracle Exadata and how it delivers on these business objectives on-premises and in the Cloud.

### Database Consolidation Defined

Database consolidation involves placing multiple databases onto a single set of computing infrastructure. Database consolidation is similar to but different from server consolidation, which involves consolidating multiple physical servers into a single physical server running multiple Virtual Machines. The primary purpose of database consolidation is to reduce the cost of database infrastructure by taking advantage of more powerful servers that have come onto the market, which include hundreds of processor cores in a single server. Database consolidation may involve any of the following configurations:

- Multiple Databases on a Single Physical Server
- Multiple Databases on a Single Virtual Machine
- Multiple Databases on a Cluster of Physical Servers
- Multiple Databases on a Cluster of Virtual Machines

Database consolidation allows more databases to run on fewer servers, which reduces infrastructure and also reduces operational costs. Each physical server represents an Information Technology asset or “Configuration Item” (CI) that must be placed into a data center rack, connected, secured, and maintained. Each Virtual Machine represents another Configuration Item (CI) that also must be deployed, connected (virtually), secured, and maintained. Database Consolidation applies in both on-premises and Cloud environments.

Oracle databases can also be consolidated into a cluster of physical servers or virtual machines. Oracle Database clustering capabilities include active/passive clustering, as well as active/active clustering using Oracle Real Application Clusters. Clustering is used to increase the availability of Oracle databases, enable rolling patching, and provide for scalability across multiple physical servers or virtual machines.

Oracle databases can also be further consolidated as pluggable databases into Container Databases. The Oracle Multitenant option allows 4096 pluggable databases to be

managed under a single Container Database, with a single backup, single disaster recovery configuration, and single clustering configuration.

## **Database Consolidation Goals**

While database consolidation is intended to reduce the cost of computing infrastructure and the cost of Information Technology operations, database consolidation must also consider the following factors to meet business requirements:

- Reduce Cost
- Simplify Administration
- Improve Security and Isolation
- Meet Availability Goals
- Deliver Required Performance

**Cost reduction** is the primary business driver of database consolidation. The cost-saving advantages of consolidation apply in traditional on-premises data centers, as well as when using Cloud services. Consolidation of databases enables higher utilization of computing assets. Consolidation also enables dynamic, automatic reallocation of computing resources using the Resource Management capabilities of the Oracle database.

**Simplified Administration** of databases in a consolidated environment begins with standardization and a reduction in the variety of configurations that need to be maintained. Database consolidation enables enterprises to manage many databases as one, reducing the management of physical servers, virtual machines, and even the number of container databases with Oracle Multitenant. Simplifying administration enables enterprises to focus efforts on ensuring the availability of databases and business applications.

**Improved Security and Isolation** of the database are achieved in a consolidation environment through a reduction in the number of points of vulnerability, an increase in standardization across databases, and a reduction in the administrative workload required to secure those databases. Database consolidation can also raise the stakes in security, since an attacker can potentially gain greater access to more data than in a non-consolidated environment. Oracle Database includes a wide array of security capabilities itself, and these capabilities are enhanced as part of Exadata to further mitigate security vulnerabilities.

**Meeting Availability Goals** of databases and applications as dictated by business requirements is critical in a consolidation environment due to the fact that multiple

databases reside on the same infrastructure. This paper outlines the proper configuration to harden systems against failures, deploy multiple layers of isolation to reduce the scope of impact (or “blast radius”) of failures, as well as ensuring fault tolerance and proper operational practices that can achieve continuous availability during maintenance events, upgrades, or in response to failure of infrastructure components.

**Delivering the Required Performance** of databases and applications is a key goal of any consolidation environment. The performance of one database should not be impacted by issues that arise on a separate database. Databases and the business applications serviced by those databases must receive the proper amount of computing resources assigned to those databases and business applications. The Oracle Database Resource Manager (DBRM) and Exadata I/O Resource Manager (IORM) provide the resource management tools necessary to provide the required level of control to ensure the proper level of database and application performance.

### **The Role of Virtual Machines in Database Consolidation**

While virtual machines (VMs) are universally used for server consolidation, they play a specific role in the consolidation of Oracle databases. Virtual Machines allow a single physical server to simultaneously run multiple Operating System (O/S) images, keeping the content of those O/S images isolated from each other. Each O/S is allocated its own disk, memory, CPU, and virtual network interfaces, and each O/S has its separate system processes and is isolated from other O/S images running on the same physical server. The primary benefit of Virtual Machines is to provide isolation at the O/S layer when multiple systems are consolidated onto the same physical server.

Virtual Machine (VM) technology plays an important role in database consolidation to provide isolation where necessary for security purposes. For example, Virtual Machines can be used to provide separation between Development and Test environments, or between multiple production databases. Virtual Machines can also be used to address isolation for data security requirements, including PCI (Payment Card Industry) data and HIPAA (Healthcare Insurance Portability and Accountability Act) data, including the use of distinct VLAN-tagged networks for each VM Cluster.

Virtual Machines also provide advantages during system software maintenance by isolating groups of databases during upgrades, as well as providing more flexible fallback methods, such as O/S snapshots. However, each Virtual Machine brings additional administrative overhead, including installation, patching, security, and management. Oracle recommends against the approach of deploying one database per VM due to the increased administrative overhead. Rather, Oracle recommends a more judicious use of Virtual Machines, including deployment of small numbers of VMs per physical machine,

with each VM hosting multiple databases, container databases, and/or pluggable databases.

### **The Role of Oracle Real Application Clusters in Database Consolidation**

Oracle Real Application Clusters (RAC) provides key features for availability and scalability in a database consolidation environment. Oracle Real Application Clusters (RAC) allows an Oracle database to span the servers of a cluster and actively operate on ALL servers in the cluster simultaneously in what is often referred to as active/active clustering. Oracle Real Application clustering involves grouping two or more servers (or Virtual Machines) that work together cooperatively under the control of software known as clusterware (such as Oracle Clusterware). Oracle RAC allows databases to scale “horizontally” by adding servers (or VMs) to the cluster. Oracle RAC also provides the ability to continue operating during server (or VM) maintenance or when a server (or VM) failure occurs. Oracle RAC allows a database to continue operation in the event of:

- Server Hardware Maintenance
- O/S Maintenance
- VM Hypervisor Maintenance
- Server Hardware Failure
- O/S Failure (ex, Kernel Panic)
- Database Software Updates
- Oracle Grid Infrastructure Updates

Oracle RAC facilitates hardware and software maintenance, including quarterly security patches, database software fixes, and Oracle Grid Infrastructure Updates. Oracle RAC provides for high availability in combination with Oracle Application Continuity, which allows user sessions to continue operating (without loss of in-flight transactions) during server failures by moving any active sessions to other surviving servers within the Oracle RAC cluster.

Non-RAC databases can be used in a database consolidation environment for lower-tier, non-critical databases that do not require Oracle RAC for scalability. However, these databases will not have protection from the events outlined above and may represent a greater operational burden on administrators due to this lack of clustering benefits. Even non-critical databases should use Oracle Real Application Clusters, if possible, simply to reduce administrative burden.

### **The Importance of Resource Management in Consolidation**

Oracle Database Resource Manager (DBRM) and I/O Resource Manager on Exadata systems (IORM) are critical components in the operation of any Oracle database and are even more critical in database consolidation. The Oracle Resource Managers (DBRM and IORM) provide the controls necessary to apportion resources across databases and prevent resource conflicts between databases that can result in performance issues of one database impacting other databases on the same system. Without DBRM and IORM, administrators use a combination of tools that are more difficult to manage and don't provide the same degree of control over resource allocations for Oracle Databases. The Oracle Resource Managers (DBRM and IORM) allow administrators to:

- Allocate and manage system resources for databases
- Ensure a proper level of database & business application performance
- Ensure the stability of databases and systems

Oracle Database Resource Manager and I/O Resource Manager (for Exadata systems) enable database administrators to allocate system resources (CPU, memory, processes, I/O, and network) to databases and manage those allocations over time with changing business demands. The Oracle Resource Managers allow system resources to be re-allocated easily or even dynamically to deliver the level of performance needed for business applications. Properly configured resource management also improves system stability by preventing resource starvation caused by misbehaving applications, databases, or systems.

The Oracle database also provides I/O resource management for non-Exadata systems, although without the deep integration found on Exadata. Please refer to the Resource Management section of this document for more information on this topic.

## **Methods for Database Consolidation**

There are 3 methods for consolidating Oracle databases onto physical machines or Virtual Machines, as discussed below. Virtual Machines are used to provide isolation at the O/S layer, while consolidation refers to multiple databases running within a Virtual Machine, within a non-virtualized physical server, or a cluster comprised of physical servers or Virtual Machines. Multiple databases can be consolidated using one of these methods:

- Multiple Databases per Physical Server, Virtual Machine, or Cluster
- Oracle Multitenant
- Schema Consolidation


These methods can be used separately or in combination to consolidate databases onto a set of physical machines or virtual machines.

## Multiple Databases per Physical or Virtual Machine

Oracle databases can be consolidated onto a single physical machine or Virtual Machine, and multiple Oracle RAC (Oracle Real Application Cluster) databases can be consolidated onto a cluster of physical machines or a cluster of Virtual Machines. These can be the same or different versions of the Oracle database and can use either shared or dedicated Oracle software (shared or dedicated ORACLE\_HOMES). The Instance Caging feature of Oracle Database Resource Manager (DBRM) enables the allocation of CPU across multiple consolidated databases on a single server or cluster of servers. The I/O Resource Manager (IORM) features of Exadata extend these resource management features into the storage layer of Exadata systems. Please refer to the related section of this document for more information on Oracle Resource Manager.

## Database Consolidation Using Oracle Multitenant

Oracle Multitenant provides the efficiency gains achieved previously with schema consolidation (see below for more information), while also providing the needed isolation between applications. Oracle Multitenant introduced the concept of Pluggable Databases (PDBs) that reside within a Container Database (CDB). Pluggable Databases can be plugged into one Container Database, then moved to another Container Database through “unplug” and “plug” operations. This architecture greatly increases the number of databases that can reside on a single server, as shown below.



**Figure 1: Consolidation using Virtual Machines compared to Database Virtualization using Oracle Multitenant**

Oracle internal tests have demonstrated that Oracle Multitenant <sup>1</sup>increases consolidation density compared to single-instance databases running on dedicated Virtual Machines. In

<sup>1</sup> Starting Oracle Database 21c, multitenant is the only option

Figure 1, 123 databases were consolidated in the same server using Oracle Multitenant, while the same server would only service 9 databases using the single database per Virtual Machine approach. The single database per VM approach quickly reached the CPU utilization limit (75%), whereas 123 databases were consolidated onto the server using Oracle Multitenant before reaching the transaction response time limit. Oracle Multitenant also reduces administrative burden by allowing many databases to be managed as a single database.

## Schema Consolidation

Schema consolidation is an approach that was used before the advent of Oracle Multitenant to combine schemas from multiple databases into a single database, making more efficient use of system resources and enabling a much higher consolidation density. While schema consolidation makes more efficient use of computing resources, this approach often requires application changes because it does not provide the necessary degree of isolation that Oracle Multitenant provides between applications and does not facilitate isolation for maintenance operations such as upgrades. Oracle Multitenant has become the preferred solution for high-density consolidation because it addresses the complexities found with schema consolidation. Schema consolidation is a useful approach in some circumstances, such as combining multiple databases belonging to a single application, but it will not be addressed further in this document.


## Why Exadata for Consolidation

[Oracle Exadata for on-premises, Oracle Cloud, and multicloud](#) is an engineered system purpose-built to provide optimal performance, availability, and manageability for Oracle Databases. Its scalable architecture and advanced software capabilities make it ideally suited as a standard database platform for the consolidation of Oracle databases. Exadata addresses each of the business goals outlined earlier, including:

- Reduce Cost
  - Simplify Administration
  - Improve Security and Isolation
  - Meet Availability Goals
  - Deliver Performance for mission-critical applications
- While Oracle databases can certainly be consolidated onto non-Exadata platforms, there are significant advantages to using Exadata for consolidation.

## Reducing Costs with High Consolidation Density

Oracle’s benchmark testing has shown that substantially higher consolidation densities can be achieved with Exadata than with similarly sized non-Exadata systems. The workload in these tests included from 150 to 300 OLTP databases running Oracle Real Application Clusters (Oracle RAC) on an Exadata X8-2 Quarter Rack with Oracle Multitenant. Exadata delivered transaction response times under the defined 3-millisecond cap in these tests, as the number of databases was increased from 150 to 300, as seen in Figure 2. Newer generations of Exadata, such as X11M, will exceed the results shown in this test.



**Figure 2: Exadata Consolidation Density and Performance**

Exadata offers high I/O bandwidth, high IOPS (Input/Output Operations Per Second), and high throughput rates, combined with offloading of database processing into storage to deliver higher consolidation density than traditional platforms. Exadata Elastic Expansion provides the ability to add storage with zero downtime as more databases are added to the system or when data volumes grow.

## **Simplification through Standard Configuration & Management Practices**

Thousands of Oracle customers run Exadata, and standardized practices have been developed for the management and administration of Exadata. Standardization decreases the complexity of the operational environment, improves security, enhances availability, and reduces costs. Exadata is a single platform that can run workloads of any scale and complexity, making Exadata the only single platform that can be used for all workloads. Oracle Enterprise Manager integrates the Exadata plugin for detailed observability and manageability for Exadata systems.

## **Enhanced Security with Exadata**

The Exadata platform includes a number of features that deliver improved security out of the box, without placing additional burden on system administrators. Consolidation of databases results in fewer systems to secure, and Exadata provides a more secure platform by default. The capabilities delivered with Exadata include:

- Defense In-Depth
- Principle of Least Privilege
- Pre-Integrated Updates & Security Fixes
- Minimal Attack Surface
- Pre-Scanned Software
- Automated Intrusion Detection Environment (A.I.D.E.)

Oracle follows the approach of defense in depth in the security design of Exadata, which includes strong authentication, encryption capabilities, database users and roles, and system-level (O/S) security. Oracle implements the principle of “least privilege”, enabling customers to assign separate administrative roles, such as for infrastructure administration, and multiple subdivisions of database administrator roles.

Exadata includes security fixes that are integrated with each quarterly release, monthly security releases, and emergency fixes to address zero-day vulnerabilities. These updates are pre-integrated across the entire stack, including O/S, Virtual Machines, Clusterware, Logical Volume Manager, drivers, storage software, networking, and the database software itself. Exadata incorporates a minimal attack surface by only including the software components required specifically to run the Oracle database. Oracle conducts multiple security scans of each Exadata software release using industry-leading security scanners. These scans include malware scanning, vulnerability scanning, and configuration compliance. Customers should also run their own scans on deployed systems to ensure against any deviations from the delivered configurations.

Exadata also includes the Advanced Intrusion Detection Environment (A.I.D.E.) as of the 19.1 Exadata system software release. AIDE automatically detects modification of critical files by checking SHA256 hash signatures against a database of expected values. For more information, please refer to the Exadata Security Guide:

<https://docs.oracle.com/en/engineered-systems/exadata-database-machine/dbmsq/exadata-security-overview.html#GUID-4D178313-47E7-451B-8AF8-1FEAC2CD38B2>

## **Ensuring Availability with Exadata and MAA**

Exadata is the leading platform for Maximum Availability Architecture (MAA) due to extensive development investment in the Exadata platform to achieve the highest levels of availability, data protection, and ensure the best performance for Data Guard and Oracle GoldenGate. These capabilities include fast node death detection, automatic storage failure detection and rebalance, automatic network fault detection, as well as fast detection and reconfiguration during instance failure. Oracle has made continuous improvements in all areas of availability with each Exadata software release. The high performance of Exadata contributes to both higher consolidation density as well as the HA and DR features, such as faster Oracle RAC instance recovery and DR standby recovery. Please refer to the Exadata MAA white paper for more information.

(<https://www.oracle.com/a/tech/docs/maa-data-sheet.pdf> )

## **Exadata Performance Optimized for Density**

The high-performance capabilities of Exadata are often used to deliver high performance, but those same performance advantages are also used to drive higher consolidation density that cannot be achieved with other platforms. Exadata reduces or eliminates most common bottlenecks and optimizes database transactions through smart Flash, smart high-bandwidth internal InfiniBand/RoCE<sup>2</sup> network, and smart storage that offloads Oracle database processes and SQL processing into the Exadata storage layer. Exadata uniquely understands the database workloads and can differentiate between high priority and lower priority I/O operations, ensuring latency-sensitive operations such as commits are not delayed by operations like backups or bulk operations.

While Exadata is often known for delivering the highest levels of performance, Exadata can also be used to deliver performance consistent with a configured service level through the use of resource management. Exadata brings comprehensive resource management

---

<sup>2</sup> Depending on the system. All Exadata systems starting X8M are on RoCE (RDMA over Converged Ethernet) internal network

capabilities under unified control. Performance of databases can be easily managed on the Exadata platform through simple allocation and management of system resources such as CPU, memory, processes, and I/O. Exadata extends the Oracle database-level resource management capabilities to include the Exadata storage network and I/O resources as well, allowing customers to manage performance and deliver the desired level of performance required by the business. Please refer to the section of this document on Resource Management for a discussion of resource management capabilities generally, as well as specifically for Exadata environments.

### **Exadata Deployment Models**

Exadata can be deployed using the following deployment models, all of which use the same underlying Exadata technology as follows:

- Exadata On-Premises
- Exadata Cloud@Customer
- Exadata Cloud Infrastructure (Oracle Cloud, Microsoft Azure, Google Cloud Platform, and AWS)

All of these deployment models bring the same benefits in terms of managing costs, simplifying operations, ensuring security, meeting availability goals, and providing the needed level of database performance. Deployment in the Oracle Cloud and Cloud@Customer brings additional capabilities that further simplify operations through Oracle's Cloud automation and Cloud Service offerings. The Cloud deployment also enhances security further through the use of mandatory data encryption and Cloud network security.

## Establishing the Foundation for Consolidation

Consolidation of databases relies on a solid foundation of standardized system configurations and operational practices to meet the business goals of cost containment, simplified operations, and improved security, while delivering the required level of application performance. Standardized configurations must also provide sufficient flexibility to meet the needs of the business.

In this section, we will address a number of critical configuration standards that allow organizations to meet these goals. However, these are not intended to be comprehensive recommendations for all configuration settings for each component. Please refer to the additional information section for a complete set of recommendations in each area.

### **Exadata Cloud Infrastructure and Exadata Cloud@Customer**

Exadata Cloud Infrastructure and Exadata Cloud@Customer, running either Exadata Database Service or Autonomous Database, incorporate the best-practice recommendations included in this section, providing the foundation needed for consolidation. Both of these offerings include infrastructure management services as well as automation needed to accomplish the business goals of database consolidation.

### **Highly Available Storage**

Consolidation of databases increases the performance and availability requirements for storage associated with those systems. Any downtime of storage has the potential to impact many databases at once. For Exadata systems, HIGH redundancy storage (triple mirroring) should be used in any consolidation environment to provide improved protection against hardware failures, as well as to allow continued operation during maintenance. All Exadata systems in Oracle Cloud are configured with high redundancy storage (disk groups when using ASM, Vaults when using [Exascale](#)).

### **Virtual Machine Configuration**

In cases where virtual machines (VMs) are used, each physical server can be configured with 1-50<sup>3</sup> Virtual Machines to provide isolation where necessary without adding excessive complexity and administrative overhead. Virtual Machines are optional in on-premises Exadata configurations but are integral to the configuration of Exadata Cloud Infrastructure and Exadata Cloud@Customer, where they provide administrative isolation between Oracle and customer operations teams. In the case of multi-VM configurations,

---

<sup>3</sup> Starting Exadata X10M and Exadata System Software 24.1, on-premises physical database server can host 50 VMs. For cloud deployments, reference the cloud service data sheets.

it is important to understand the maximum number of Virtual Machines that can be configured for each deployment type. For more information on the use of Virtual Machines with Exadata, please refer to the following:

KVM (Kernel Virtual Machine): <https://www.oracle.com/a/tech/docs/exadata-kvm-overview.pdf>

## Memory Configuration

While this section does not include a comprehensive treatment of server configuration, special attention should be paid to the topics below in a database consolidation environment. Memory management is especially important in a database consolidation environment due to the potential for impact across multiple databases. Customers should be using the following memory management settings of Linux and the Oracle database:

- Linux Huge Pages
- Oracle Database USE\_LARGE\_PAGES parameter

A portion of system memory should be allocated to Huge Pages in Linux in proportions designed for the specific workload type (OLTP or Data Warehouse), and the Oracle SGA should be placed within the Huge Pages area. Specifying USE\_LARGE\_PAGES=ONLY in the Oracle database will ensure the SGA is placed in the Huge Pages area. While this prevents the affected database instance from starting, the purpose is to prevent memory starvation that can potentially impact other databases on the system. Please refer to MOS [KB150354](#) and MOS [KB151310](#) for more details on the use of Linux Huge Pages. Starting Oracle Database Release 19.27 and Oracle Database Release 23.8, small pages are not allowed on Exadata (MOS [KB174916](#))

## Multitenant Container Databases

Each physical or virtual machine running Oracle Multitenant databases should include provisions for at least 2 Container Databases (CDBs) to facilitate upgrades. Pluggable Databases (PDBs) can be moved between CDBs during upgrade, or an entire CDB can be upgraded to reduce administrative complexity. One CDB on each server (physical or virtual) is used for the current database version, while the 2<sup>nd</sup> CDB is used for the next database version during upgrade. Databases can be moved from one container to the next during upgrade using relocate operations.

## Configuration Checking Tools

Oracle provides a number of configuration checking tools that are consistently updated with the latest best practices configuration checks. These tools include the following:

- DB Security Assessment Tool (DBSAT)
- exachk

The Oracle Database Security Assessment tool (DBSAT) is used for checking security vulnerabilities of Oracle databases. The exachk (EXA Check) tool is a holistic health check of the entire Exadata system, evaluating storage, network operating system, Grid Infrastructure, and database. It highlights when key Exadata MAA configuration practices or software vulnerabilities are exposed in the target system.

### **Data Guard & Active Data Guard (Read Replicas)**

Oracle Data Guard has provided industry-leading Disaster Recovery (or “standby” database) capabilities for more than two decades. Data Guard standby databases are kept in sync with a “primary” source (production) database through the automatic application of database transactions contained in redo logs. Data Guard standby databases normally consume a fraction of the resource consumption found on the source (or “primary”) database until the database is activated during a disaster scenario.

While Data Guard standby databases are intended primarily for use during disaster recovery, Active Data Guard (ADG) enables a standby database to be actively queried as a “Read Replica” while the database is being updated from the stream of redo changes generated by the primary database. Data Guard also provides additional protection, especially with automatic block repair of physical data corruption.

### **Exadata Infrastructure & DB Version Interoperability**

The Exadata Infrastructure layer consists of the software layers that underlie databases running on the Exadata system. Oracle Databases, Container Databases, and Pluggable Databases are installed on top of these infrastructure layers. Multiple Oracle Database versions are supported with each release of the underlying Exadata Infrastructure layers. Version interoperability is outlined in MOS Note # [KB153930](#) and # [KB114881](#) (cloud versions).

### **Oracle AI Database 26ai**

Oracle AI Database 26ai offers several features that enhance database consolidation, enabling organizations to streamline operations, reduce costs with less infrastructure, and improve resource utilization.

### **Converged Database Capabilities**

Oracle 26ai supports diverse data types (relational, JSON, graph, vectors, text) and workloads (AI, analytics, and Online transaction processing) within a single database engine. Both capabilities eliminate the need for separate specialized databases, enabling varied workloads on one Exadata platform, reducing complexity and infrastructure costs.

### **AI and automation features**

Oracle AI Database 26ai integrates AI capabilities like AI Vector Search and can be combined with existing relational business data to run real-time queries on both structured and unstructured data. This capability provides consolidation opportunities to host different data models in one database engine.

## **Exadata Exascale**

Introduced with Exadata System Software 24ai, release 24.1, Exadata Exascale is a new intelligent data architecture that combines the performance, availability, scalability, and security of Exadata with hyper-elasticity, multi-tenancy, and resource pooling.

Exascale moves storage management out of the database servers and into storage servers enabling organizations to consolidate tens to thousands of Oracle Databases on a single cloud infrastructure, supporting various data types (relational, document, AI vectors, text, etc.) and workloads (AI, analytics, and mission-critical transactional). Some additional features enhancing consolidation include:

### **Centralized VM Hosting**

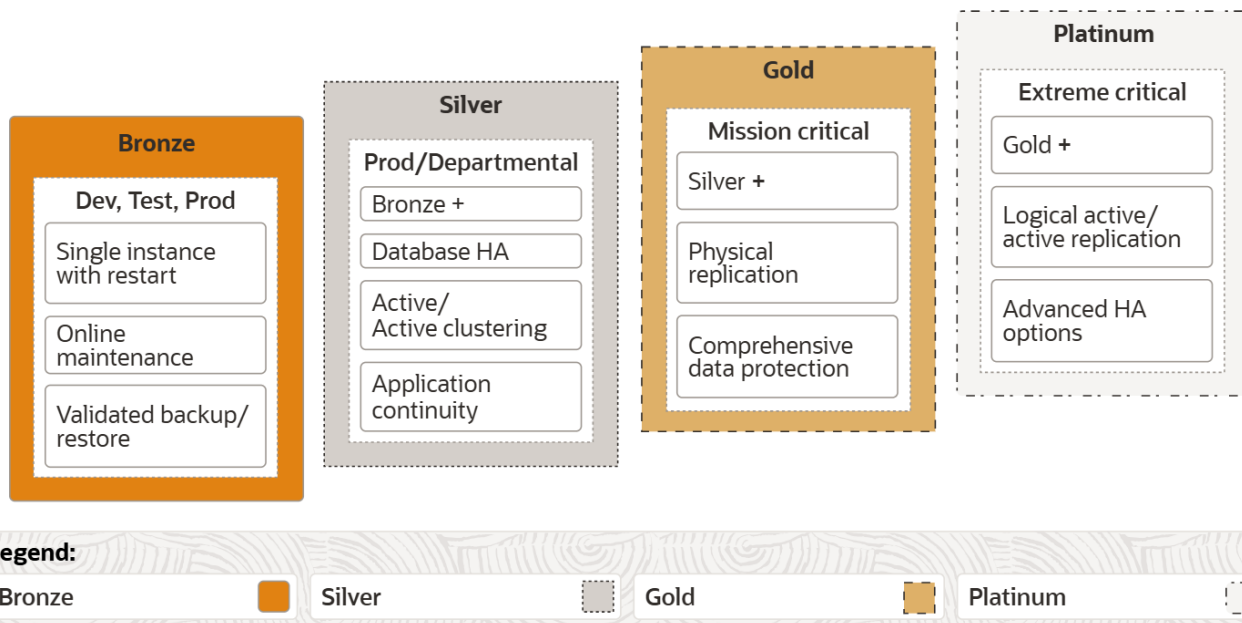
Exadata Exascale's multitenant architecture allows Oracle Databases to run within virtual machines whose VM images are hosted on shared, centralized Exascale storage, eliminating the storage-space limitation of local disk storage. This centralized approach aids in consolidation specifically, server and database consolidation by enabling efficient sharing of resources across multiple tenants and workloads, reducing the need for dedicated hardware.

### **Intelligent Snapshots and Thin Provisioning**

Exadata Exascale utilizes snapshots and thin provisioning to significantly improve storage efficiency and database cloning capabilities. Thin provisioning allows for the creation of clones (including those from live databases or snapshots) that only consume storage space for the changes made to them, rather than a full copy. In a consolidation environment, this feature dramatically reduces storage requirements, especially for development and testing environments.

## **MAA Reference Architectures**

Consolidation of databases onto a common set of infrastructure raises the stakes in terms of availability. Oracle Maximum Availability Architecture (MAA) best practices define four HA reference architectures that address the complete range of availability and data protection required to address these needs in a database consolidation environment. The architectures, or HA tiers, are designated PLATINUM, GOLD, SILVER, and BRONZE, delivering the service levels described in Figure 3 below.



**Figure 3: HA and Data Protection Service Levels**

Each tier uses a different MAA reference architecture to deploy the optimal set of Oracle HA capabilities that reliably achieve a given service level at the lowest cost and complexity. These architectures explicitly address all types of unplanned outages, including data corruption, component failure, system and site outages, as well as planned outages due to software maintenance, migrations, or database upgrades.

**The Bronze** MAA reference architecture provides basic database service at the lowest possible cost. A reduced level of high availability and data protection is accepted in exchange for reduced cost and implementation complexity. This architecture may be suitable for databases used for testing, development, and less critical production applications and databases. The Bronze architecture uses the high availability capabilities included in Oracle Enterprise Edition. Bronze defaults to the Oracle Database single-instance or multitenant architecture. Oracle Restart or Oracle Clusterware high availability capabilities are used to restart a failed instance, database server, or any relevant managed service. In the worst-case scenario of a complete site outage, there is additional time required to perform these tasks at a secondary location, which may result in days of downtime.

In the Bronze reference architecture, a local backup within the same data center is always recommended for the fastest recovery. Oracle also recommends maintaining a second copy of backups in a remote data center to protect against site outages and disasters. You can use Oracle Cloud Database Backup Service to maintain a cloud-based backup of on-premises databases.

**Silver** reference architecture provides an additional level of availability for databases that require minimal or zero downtime in the event of database instance or server failure, as well as many types of planned maintenance, such as operating system, Database, and Grid Infrastructure software updates, including required security updates. The MAA Silver architecture is built on the foundation of MAA Bronze architecture and adds active-active clustering technology using Oracle Real Application Clusters (Oracle RAC) for minimal or zero downtime in the event of database instance or server failover, as well as zero database downtime for the most common planned maintenance events. RMAN provides database-optimized backups to protect data and restore availability should an outage prevent the cluster from being able to restart. Application Continuity provides a reliable replay of in-flight transactions, masking outages from users. In addition, gold-image patching and upgrades are recommended to standardize updates and minimize downtime during software maintenance activities across the Oracle Database, Grid Infrastructure, and Exadata software infrastructure using Oracle Fleet Patching and Provisioning (FPP) for on-premises deployments and the Oracle Exadata Fleet Update service in the Oracle Cloud.

**Gold** substantially raises the service level for business-critical applications by providing low-downtime solutions for all unplanned outages and major database software upgrades. Gold adds database-aware replication technologies using Oracle Active Data Guard to synchronize one or more replicas of the production database, providing data protection and availability. Far Sync extends real-time data protection and availability to ensure support for long distances between production and disaster recovery sites. Database-aware replication greatly increases HA, disaster recovery, and data protection beyond what is possible with storage replication technologies. This also reduces cost while improving return on investment by enabling active use of all replicas at all times.

**Platinum** builds upon the Gold tier of protection by adding Oracle Golden Gate to provide bi-directional capabilities that enable zero downtime upgrades and migrations. Global Data Services in the Platinum Tier provides automated service management and workload balancing for replicated database environments. In addition, the platinum tier includes optional configurations for Edition-Based Redefinition (EBR), which enables application upgrades to be performed online, and sharding, which provides intelligent horizontal partitioning of data across a farm of independent databases using a common key (i.e., shard key). While each of these technologies requires additional effort to implement at the platinum tier, these technologies deliver substantial value for the most critical applications where downtime and data loss are not an option.

The MAA reference architectures provide a standard infrastructure optimized for Oracle Database that enables enterprises to select the level of HA appropriate for different

service-level requirements. Standardization reduces cost and simplifies the movement of databases from one HA tier to the next, or from one hardware platform to another, should business requirements change. It should be noted that all tiers also include Oracle Enterprise Manager Cloud Control for configuration, monitoring, alerting, and management in addition to the specific set of Oracle HA capabilities that are described for a given tier.

Please refer to the MAA architecture documentation (<https://www.oracle.com/webfolder/technetwork/tutorials/architecture-diagrams/high-availability-overview/high-availability-reference-architectures.html>), Exadata MAA architecture presentation (<https://www.oracle.com/a/tech/docs/exadata-maa.pdf>) and Exadata Cloud MAA architecture presentation ([Oracle Cloud: Maximum Availability Architecture Presentation](https://www.oracle.com/a/tech/docs/cloud-maa-overview.pdf) or <https://www.oracle.com/a/tech/docs/cloud-maa-overview.pdf>) for further details.

## Applying MAA to Database Consolidation

Service level expectations in a non-consolidated environment are quite different compared to a consolidated environment. For example, with a stand-alone database used by a developer or a department, the level of disruption caused by a system down event is limited to a smaller user base. However, when this same database is consolidated with 100 other databases, the impact is much greater if each database supports different departments and user communities. The level of disruption to the enterprise of an outage that impacts the consolidated database is magnified by 100 times, making HA and data protection a higher priority. The principles outlined in this section apply equally to on-premises database deployments on traditional hardware architectures, as well as to Exadata deployed in on-premises, Cloud@Customer, or OracleCloud environments.

## Virtual Machines for Isolation

On systems that use Virtual Machines, Oracle recommends using approximately 2-50<sup>4</sup> Virtual Machines per physical server at most in order to minimize the additional administrative overhead that comes with each Virtual Machine. Virtual Machines can be used to isolate groups of databases according to the environment (isolating DEV from TEST) or to isolate databases to meet governance requirements (such as PCI data from non-PCI data). Virtual Machines can also be used to isolate a subset of databases from others to reduce the impact of upgrades and other O/S level system maintenance tasks. Oracle Real Application Clusters allow user sessions to be relocated to another node of a

---

<sup>4</sup> Starting Exadata X10M and Exadata System Software 24.1, on-premises physical database server can host 50 VMs

Virtual Cluster, enabling zero downtime maintenance. Oracle RAC therefore eliminates the approach of deploying one database per VM for maintenance isolation, eliminating the administrative overhead associated with large numbers of VMs.

### **Patching and Upgrades in Consolidated Environments**

Patching and upgrades are a particular challenge in a consolidated environment. With multiple databases and business applications relying on a common set of hardware and a common software stack, it is critical to provide the proper level of isolation to facilitate upgrades. For example, a single upgrade at the CDB level affects all PDBs within that container. This ability to “manage many as one” significantly reduces administrator time but increases the need for coordination of upgrades. Therefore, databases should be grouped according to SLA and software update cycles to ease coordination. Capabilities, including Grid Infrastructure Rolling Patching, Oracle Real Application Cluster Rolling Patching, Data Guard Standby First Patching, Data Guard Transient Logical Rolling Upgrade, or Oracle GoldenGate, can also be used to reduce downtime to near zero or zero to upgrade an entire CDB in a single operation.

To upgrade individual PDBs within a CDB, a new CDB can be created with the upgraded version, and the PDB(s) can be relocated to the new CDB. This results in downtime during the PDB movement operation, but provides the flexibility that may be required to address specific application needs. Although more complex to operate, Oracle GoldenGate also provides the option to upgrade PDBs with minimal or zero downtime by replicating between different CDBs and between databases of different versions.

### **Application Continuity for Availability**

Consolidation of databases involves multiple databases running on each physical server and/or Virtual Machine. Oracle Real Application Clusters (Oracle RAC) and Active Data Guard (ADG) provide for high availability for both planned and unplanned outages. Oracle Database also includes the ability to provide continuous application service availability with the following capabilities:

- Transparent Application Failover (TAF)
- Transparent Application Continuity (TAC)

Sessions are drained & rebalanced across nodes of a Oracle RAC cluster or Active Data Guard standby database for planned maintenance, while in-flight (active) sessions can be relocated across nodes of the Oracle RAC cluster using TAF or TAC. The automatic relocation of sessions across Oracle RAC nodes also occurs in the event of physical server or Virtual Machine failures. Databases in a consolidation environment that require high availability should be configured to enable these capabilities. Please refer to the following

Continuous Availability white paper for more information on session drain and rebalance, as well as the use of TAF and TAC:

<https://www.oracle.com/docs/tech/database/continuous-service-for-apps-on-atpd.pdf>

### **Considerations for Multiple Databases per Server**

Consolidation can be accomplished by the placement of multiple non-PDB (non-Multitenant<sup>5</sup>) databases on each physical server, Virtual Machine, or Cluster of physical machines or virtual machines. These databases will not benefit from the advantages of Oracle Multitenant, such as reduction in background processes, reduced shared memory usage, dynamic sharing of system resources, and administrative benefits from managing many-as-one. Multiple non-Multitenant databases can share a single installation of Oracle software, in what is known as a shared Oracle Home (\$ORACLE\_HOME) configuration. Systems should include provision for at least 2 Oracle Homes to facilitate upgrades.

### **Multitenant Considerations**

Oracle Multitenant uses the Oracle Maximum Availability Architecture (MAA) to address the HA and data protection requirements of consolidated environments. In addition to MAA's customary objectives for HA and data protection, some objectives are specific to a multitenant architecture context:

**Manage Many-as-One Simplicity.** MAA best practices must easily scale the management of large consolidated environments so that HA and data protection objectives are achieved while realizing maximum cost benefit (capital and operating costs).

**Isolation.** MAA best practices must prevent problems that impact a single PDB from impacting the availability of other PDBs in the CDB.

On the surface, the first two objectives appear contradictory. Isolation for “n” environments often results in “n” different environments that must be individually managed. Oracle MAA can leverage the multitenant architecture so that HA is achieved with minimal compromise in isolation while achieving substantial benefits from being able to manage-as-one.

Oracle Multitenant enables the management of a CDB as a single database regardless of how many PDBs it contains. A CDB with 100 PDBs can achieve a 100 to 1 reduction for many maintenance and operational tasks: a single RMAN backup, a single Oracle Active Data Guard standby for disaster recovery, and a single database to upgrade or patch.

---

<sup>5</sup> Starting Oracle Database 21c, multitenant is the only option

Oracle MAA also provides the flexibility to manage a PDB in isolation from other PDBs when appropriate. For example:

If an individual PDB must be restored, RMAN can do so without impacting other PDBs that are open in the same CDB. Note that a PDB that has just been plugged in should be backed up immediately after the plugin operation to ensure that it can be recovered should a problem arise before the next regularly scheduled CDB backup.

If fast point-in-time recovery is required, the first release of Oracle Multitenant (in 12cR1) enabled using Flashback Database at the CDB level. Oracle Database 12c Release 2 introduced the ability to use Flashback Database on an individual PDB without impacting the availability of other PDBs.

If an individual PDB experiences a problem that the administrator believes has the potential to impact other PDBs in the CDB, the suspect PDB can be easily unplugged and plugged into another CDB, where the issue can be resolved in isolation. An unplug/plugin operation also enables the flexibility of applying a patch without impacting other PDBs running in the original CDB. Once the problem is resolved, the PDB can be left in the new CDB, or the PDB can be relocated back to the original CDB after the new patches have been installed.

If one PDB is corrupted on the primary CDB database and there's an existing standby database where the PDB is healthy, a single PDB failover can be executed by unplugging and plugging the PDB from the standby to a new CDB. In Oracle Database 12c Release 2, this functionality is automated through the Data Broker command line interface. Refer to My Oracle Support Note: Unplugging a Single Failed PDB from a Standby Database and Plugging into a New Container (Doc ID 2088201.1).

Oracle Resource Manager will prevent a PDB from consuming more than its assigned share of system resources if there is a sudden spike in workload, a bug, or some other event that changes consumption patterns.

## Managing Unplanned Outages

Table 1 identifies various unplanned outages that can impact a database in a multitenant architecture. It also identifies the Oracle HA solution to address that outage that is available in each of the HA tiers described earlier in this paper.

**Table 1: Unplanned outage matrix for multitenant architecture**

EVENT	SOLUTIONS FOR BRONZE, SILVER, GOLD & PLATINUM	RECOVERY WINDOWS (RTO)	DATA LOSS (RPO)
	BRONZE: Database Restart	Minutes after server restart	Zero

<b>Instance failure</b>	SILVER, GOLD, PLATINUM:  Oracle RAC  Application Continuity	Seconds with Oracle RAC  Zero with Application Continuity	Zero
<b>Permanent DB node failure</b>	BRONZE: Restore and recover	Hours to Day	Zero
	SILVER, GOLD, PLATINUM:  Oracle RAC  Application Continuity	Seconds with Oracle RAC and Exadata Fast Node Death Detection.  Zero with Application Continuity	Zero
<b>Storage failure</b>	ALL Tiers: Automatic Storage Management on Exadata and with Exadata HIGH redundancy (triple-mirroring)	Zero database downtime	Zero
<b>Physical Data Corruption</b>	BRONZE, SILVER: Basic protection. Some corruptions require restoration and recovery of the PDB or entire CDB	Hours for restore and recovery needed	Varies depending on the time of last backup.
	GOLD, PLATINUM: Comprehensive corruption protection and Auto Block Repair with Oracle Active Data Guard	Zero with Active Data Guard automatic block repair for physical corruptions.  Seconds to minutes if corruption occurs due to lost writes when using Data Guard Fast Start Failover (FSFO) for repair.	Zero with Automatic Block Repair for Physical Corruptions.  Zero with SYNC or FAR SYNC  Seconds with ASYNC or when encountering lost writes on primary.
<b>Logical Data Corruption</b>	ALL Tiers: Logical failures resolved by flashback drop, flashback table, flashback transaction, flashback query, and undo.	Varies based on detection time	Varies based on the extent of failure
	BRONZE, SILVER: Some logical block corruptions require restoration and recovery of PDB or entire CDB.	Hours to day	Varies based on the extent of failure
	GOLD, PLATINUM: Standby prevents and detects logical corruptions during redo apply with DB_BLOCK_CHECKING enabled. Once detected, a failover can be completed in seconds with minimal downtime.	Seconds after the manual failover is executed	Seconds
<b>Cluster, Database complete site failure</b>	BRONZE, SILVER: Restore and recover	Hours to Day	Since the last backup. For Oracle Cloud, it can be 30 minutes or less since the last archive backup.
	GOLD, PLATINUM: Fail over to secondary using Oracle Active Data Guard Fast Start Failover	Seconds to a couple of mins	Zero to Near Zero
<b>Single PDB</b>	BRONZE, SILVER: Restore and recover PDB	Minutes to hours	Recovery point as of last backup.  For Oracle Cloud, it can be 30 minutes or less since the last archive backup.

<b>Failure</b>	GOLD, PLATINUM: Fail over to CDB secondary or perform PDB Failover ( <a href="#">Note 2088201.1</a> )	Seconds to minutes	Zero for Data Guard and PDB failover
<b>Performance degradation</b>	ALL Tiers: Database Resource Manager to isolate impact	Degraded service on the affected database.	Zero

While it is possible to remedy many issues with an individual PDB without impact on other PDBs, there are situations in which isolation is required. For example, you may need to apply a patch to the Oracle Home being used by the CDB. In this scenario, it is recommended that you create a new Oracle Home and CDB, then unplug the PDB and plug it into the new CDB. You can then resolve the problem with 100% certainty with zero impact to other PDBs in the original CDB. Once the problem is resolved, you can unplug/plug the PDB back into its original CDB or leave it in the new CDB until some future point in time.

### **Managing Planned Maintenance**

From a planned maintenance perspective, all of the traditional solutions available to non-CDBs will work in a multitenant architecture environment. Additionally, there are cases where the administrator can decide if the maintenance should be done to just one PDB or all PDBs in the same container. Oracle provides the flexibility to handle either situation. Oracle HA solutions for planned maintenance are provided in Table 2 below.

**Table 2: PLANNED MAINTENANCE MATRIX FOR MULTITENANT ARCHITECTURE**

EVENT	SOLUTIONS FOR BRONZE, SILVER, GOLD AND PLATINUM	EXPECTED DOWNTIME
<p><b>Dynamic and Online Resource Provisioning</b></p> <p>or</p> <p><b>Online reorganization and redefinition</b></p>	<p>ALL Tiers: Online Reorganization and Redefinition of select objects within each PDB</p> <p>Documentation: <a href="#">Dynamic and Online Resource Provisioning</a> and <a href="#">Online Reorganization and Redefinition</a></p>	<p>Zero</p>
<p><b>Database and Grid Infrastructure Software Updates</b></p>	<p>ALL Tiers: The entire CDB can be online-patched if relevant</p> <p>SILVER, GOLD, PLATINUM: The entire CDB can leverage Oracle RAC rolling upgrade with Application Continuity</p> <p>GOLD, PLATINUM: With Gold and Platinum, apply software changes to the standby first by following standby-first patch recommendations to evaluate the software changes as an additive precaution, then apply Oracle RAC rolling upgrades with Application Continuity on the primary.</p>	<p>Zero</p> <p>Seconds to minutes</p> <p>Zero</p> <p>Zero by relocating services</p> <p>Zero application outage</p>
<p><b>Database Upgrades (e.g., 18c to 19c or 26ai)</b></p>	<p>ALL Tiers: <a href="#">PDB can relocate to CDB with targeted software release</a></p> <p>GOLD, PLATINUM: The entire CDB can leverage Active Data Guard database rolling upgrade for patchsets and major database releases</p> <p>PLATINUM: Upgrade CDB to the target release and then switchover connections to the new upgraded CDB/PDB</p>	<p>Estimated seconds to an hour with no datafile copy option. PDB metadata still needs to be upgraded after relocation.</p> <p>Seconds to Minutes</p> <p>Zero downtime</p>
<p><b>Application upgrades</b></p>	<p>PLATINUM: <a href="#">Edition-Based Redefinition</a> requires developers to design to leverage this feature</p> <p>PLATINUM: PDB can switch over to GoldenGate replica with the targeted application changes</p> <p>Documentation: <a href="#">Online Application Maintenance and Upgrades</a></p>	<p>Reduce software sprawl and risk. Reduce OPEX and increase overall stability and availability.</p>

## Measuring Capacity Utilization Before Consolidation

This section addresses the measurement of capacity utilization before migrating to a consolidation environment rather than methods for ongoing monitoring of capacity usage. Monitoring implies ongoing operational support, related tools (such as Oracle Enterprise Manager), and implementation of those tools, which is beyond the scope of this document. Proper measurement of system resource utilization is critical to resource planning and configuration of Resource Management, which will be covered in detail later in this document.

It is critical to note that database performance is closely related to system capacity allocated to the database, but these are ultimately separate aspects of system management. Performance can often be improved by allocating more resources to a database, while performance tuning of the database can often reduce resource usage. Therefore, performance and capacity are linked but must be addressed separately. This section addresses the amount of resources (or capacity) assigned to each database and does not address database and application performance tuning.

Oracle Automatic Workload Repository (AWR) provides detailed data for analysis of capacity usage for Oracle databases and the systems that host Oracle databases for databases using the Oracle Diagnostic Pack. While AWR is a valuable tool for this analysis, this section provides a “least common denominator” approach that can be used in virtually any environment.

Measuring capacity utilization begins before each database is migrated to the consolidation environment. It is important to understand the database’s average and peak resource consumption at the source while planning migration to a consolidated environment. Capacity utilization should be tracked (monitored) after migration to the target consolidation environment to ensure capacity allocations are adjusted as needed when workloads change. It is always recommended to validate performance and capacity needs in a test environment before moving databases into production.

Capacity usage measures should be gathered at 3 levels:

**System Level.** The servers and storage should be monitored to see if CPU, memory, and storage utilization (storage space and I/O) are at acceptable levels. Customers planning to consolidate more databases onto a set of servers and storage should also monitor the available capacity or headroom.

**Database (or CDB) Level.** Databases should be evaluated to determine resource usage patterns. Databases should also be evaluated to determine whether the Resource Manager (if used currently) is restricting or throttling them. Waits for CPU or I/O due to Resource Manager constraints may indicate the need to increase resources assigned to the Database or Container Database. Monitoring of database performance is also important to ensure proper resource allocation, or if tuning is required.

**PDB monitoring (Multitenant only).** Each PDB should be monitored in the same way as a database. If actual resource usage surpasses the resources guaranteed by the CDB Resource Plan or individual PDB settings, then performance will likely degrade as more PDBs are added to the CDB. If the anticipated performance

degradation is not acceptable, additional PDBs should not be introduced into the CDB.

Key capacity indicators are highlighted in the table below.

**Table 3: Key CAPACITY indicators**

MONITOR /ADMINISTER	GUIDELINE
CPU	<p><b>System (Physical Server or Virtual Machine):</b></p> <p>Monitor the system's CPU utilization, using available tools or using the "Host CPU Utilization" metric from v\$sysmetric_history. If CPU utilization is near 100%, use OS tools or the "OS Load" metric from v\$sysmetric_history to determine the degree of over-subscription on the system. Excessive over-subscription should be avoided, as specified in Table 4, Guidelines for Over-Subscription by Environment. Use Instance Caging to avoid excessive loads.</p> <p>Database CPU utilization on existing systems is used as a baseline for consolidating databases onto the new systems. Databases moved from legacy hardware will often require fewer CPU resources on an Exadata system. However, the degree of CPU savings when moving to Exadata will vary depending on the specific workload of each database. Therefore, the existing system CPU utilization is used as a baseline for the initial deployment. The following query can be used on databases with Resource Plans enabled to determine the CPU resource consumption for the database, including the running and waiting sessions:</p> <pre> select to_char(begin_time, 'HH24:MI') time,        sum(avg_running_sessions) avg_running_sessions,        sum(avg_waiting_sessions) avg_waiting_sessions from v\$rsrcmgrmetric_history group by begin_time order by begin_time; </pre> <p>The average running sessions from this query determine the CPU_COUNT required. However, if Average Waiting Sessions is greater than zero with any regularity and performance is not acceptable, the CPU_COUNT should be raised higher.</p> <p><b>Database or CDB:</b></p> <p>If Instance Caging, CDB Resource Manager, or per PDB CPU management is enabled, monitor the amount of CPU that the instance used and the amount of CPU it needed but was prevented from using. Use v\$rsrcmgrmetric_history, as described in MOS note <a href="#">KB137827</a>. The sum of 'avg_running_sessions' across all Consumer Groups and PDBs specifies the number of CPUs used. The sum of 'avg_waiting_sessions' across all Consumer Groups and PDBs specifies the throttling performed by Resource Manager due to insufficient CPU. This value corresponds to the additional number of CPUs needed.</p> <p>For CDBs, evaluate the available CPU capacity or headroom to determine if more PDBs can be added. If CDB Resource Manager or per PDB management is enabled, calculate the total amount of CPU needed by summing 'avg_running_sessions' and 'avg_waiting_sessions' from v\$rsrcmgrmetric_history across all Consumer Groups and PDBs. If this sum is less than CPU_COUNT for the entire CDB instance, then the gap between this sum and the instance CPU_COUNT is the CPU headroom on the CDB, and you can consider adding more PDBs. If this sum is greater than the instance CPU_COUNT, then this CDB is already operating at capacity and has no headroom. Adding more PDBs will only increase contention for the CPU. Therefore, additional PDBs should only be added for BRONZE CDBs if the existing PDBs and their applications can tolerate some drop in performance. Monitoring should be done during peak hours. For more details, see MOS note <a href="#">KB137827</a>.</p> <p><b>PDBs:</b></p>

	<p>Monitor the actual CPU usage and CPU wait time of each PDB. Use <code>v\$rcmgrmetric_history</code> to see how many CPUs a PDB used, using the 'avg_running_sessions' metric. Compare the actual CPU usage with the PDB's guaranteed CPU, which is based on the PDB's shares divided by the total number of shares across all PDBs that are open on the instance. For example, if the PDB has 1 share and the total shares are 5, then the PDB is guaranteed 1/5th of CPU_COUNT. If the PDB's actual CPU usage is above its guaranteed CPU usage, then the PDB's performance may suffer if additional PDBs are added. Use <code>v\$rcmgrmetric_history</code> to view the PDB's CPU wait time, using the 'avg_waiting_sessions' metric. If non-zero, the PDB's performance could be improved by increasing its shares or utilization limit in the CDB Resource Plan. For more details, see MOS note <a href="#">KB137827</a>.</p>
<p><b>Memory</b></p>	<p><b>System (Physical Server or Virtual Machine):</b></p> <p>No paging should be seen. The <code>vmstat</code> command or other tools can be used to monitor for zero or very low page-in and page-out rates.</p> <p>Memory should never be oversubscribed. On Linux systems, <code>/proc/meminfo</code> can be evaluated to determine the total system memory allocated. Compare the memory allocated at the O/S level to the memory allocated to Oracle databases and client processes. SGA usage can be computed from each database's <code>SGA_TARGET</code> parameter. PGA usage can be determined using the statistics below.</p> <p><b>Database or CDB:</b></p> <p>SGA should reside in HugePages to avoid impacting the performance of the database or other databases on the system. Setting <code>USE_LARGE_PAGES=ONLY</code> forces SGA into HugePages and prevents databases from starting if the HugePages space is not sufficient. Otherwise, SGA size simply needs to be captured for capacity utilization purposes.</p> <p>PGA usage is highly dependent on <code>PGA_AGGREGATE_TARGET</code> and can be monitored via <code>v\$pgastat</code> or using monitoring tools. PGA size will be hard-limited to <code>PGA_AGGREGATE_LIMIT</code> (if configured), or the "maximum PGA allocated" statistic will show the maximum PGA that the instance has allocated since its startup. The statistic "total PGA allocated" is the amount of PGA that the instance has allocated currently. Monitoring these values enables the DBA to know how much PGA is actually being used by the instance. They should be compared to <code>PGA_AGGREGATE_TARGET</code> to see if the database is exceeding it and by how much. <code>PGA_AGGREGATE_LIMIT</code> will impose a hard limit on PGA usage by reducing PGA where possible without disruption, as well as terminating sessions using excessive PGA.</p>
<p><b>I/O</b></p>	<p>For Exadata storage cells. Monitor Exadata I/O metrics using the script in MOS note 1337265.1. These metrics can also be viewed using Enterprise Manager.</p> <p>Monitor the Exadata FlashCache hit ratio for OLTP I/Os and evaluate increases in FlashCache allocations for the database if hit ratios are unacceptable. If the latency is still not acceptable, then consider changing the IORM objective from "auto" to "balanced" or "low latency".</p> <p>Monitor the actual disk IO utilization across databases using the <code>DB_IO_UTIL_SM + DB_IO_UTIL_LG</code> metrics. This comparison shows which databases are using the disks.</p> <p>Compare <code>FC_IO_RQ_R</code> with <code>FC_IO_RQ_R_MISS</code> to calculate the OLTP flash hit ratio. If the current performance should remain the same, make sure the flash has headroom by monitoring its flash hit ratio before adding more databases to use flash. For example, if performance is acceptable at an 80% or higher flash hit ratio, performance may deteriorate if more databases are added and the flash hit ratio drops below 80%.</p> <p><b>For the database or CDB:</b></p> <p>Monitor I/O wait events, using AWR. If the problematic wait event is "db file sequential read", then monitor and tune the database's flash cache hit rate (see above), the storage server's disk latency (see above), and the database's IORM throttle time for small requests (see below). If the problematic wait event is "cell smart table scan", then monitor and tune the database's throttle time (see below).</p> <p>If IORM is enabled, monitor the average IORM throttle time per request, using the <code>DB_IO_WT_SM_RQ</code> and <code>DB_IO_WT_LG_RQ</code> metrics. If the wait times are large and the database performance is unsatisfactory, either the database's allocation/share or utilization limit in the inter-database plan needs to be increased. If the wait times are large across all databases, then the disks have reached their maximum capacity. In this case, no new databases should be added unless the existing databases can tolerate a drop in performance.</p> <p>For CDBs, monitor the actual disk IO utilization across PDBs using the <code>PDB_IO_UTIL_SM + PDB_IO_UTIL_LG</code> metrics. This comparison shows which PDBs within the CDB are most heavily using the disks.</p> <p><b>For the PDB:</b></p>

	If IORM is enabled, monitor the average IORM throttle time per request, using the PDB_IO_WT_SM_RQ and PDB_IO_WT_LG_RQ metrics. If the wait times are large and the PDB performance is unsatisfactory, either the PDB's share or utilization limit in the CDB Resource Plan needs to be increased. If the wait times are large across all PDBs, then the allocation and/or limits for the CDB in the inter-database IORM plan need to be increased. Until then, no new PDBs should be added unless the existing PDBs can tolerate a drop in performance.
<b>Database performance indicators</b>	Use Automatic Workload Repository and/or Enterprise Manager for the database load profile, and drill down to see top wait events. From here, drill down into specific waits.  Use ASH analytics to view a breakdown of CPU usage by consumer group or PDB.
<b>Application performance indicators</b>	Use Automatic Workload Repository and/or Enterprise Manager for the database load profile, and drill down to see top wait events. From here, drill down into specific waits.  Use ASH analytics to view a breakdown of CPU usage by consumer group or PDB.

## Resource Management for Consolidation

Database Consolidation is the placement of multiple databases onto shared computing and storage infrastructure, primarily for the purposes of reducing infrastructure and operational costs. However, shared infrastructure typically raises performance concerns among those responsible for the business applications that rely on those shared resources. How do business owners know their applications are receiving the proper amount of system resources required to deliver acceptable performance? The answer is resource management.

Oracle database includes a number of powerful tools for managing resources that allow multiple databases to share a single set of compute infrastructure, while still ensuring the proper amount of resources required to support critical business functions. These capabilities allow IT organizations to properly manage resources without the administrative overhead of deploying multiple physical servers or Virtual Machines. Exadata extends these resource management capabilities into the Exadata storage tier, providing comprehensive management of the complete set of resources that databases rely upon. The resource management approach outlined in this chapter applies to Exadata on-premises, Exadata Cloud Infrastructure, and Exadata Cloud@Customer using the same settings discussed in the sections below.

### Resource Management Defined

Computing resources in either a Cloud environment or a traditional data center environment represent a significant cost to any organization. Proper management of those computing resources controls cost, ensures performance of business applications, and prevents performance issues in one business application from impacting another. Oracle's resource management tools include the ability to manage system resources so that each database (and therefore each business application) receives the desired

minimum amount of system resources, and each database does not exceed the amount of allocated resources. Oracle allows control over the following resources:

- Compute (CPU)
- Memory
- Processes (Database Sessions & Parallel Query Processes)
- I/O Operations and Bandwidth
- Exadata Storage & Interconnect Network
- Storage Space (Disk and/or Flash)
- Exadata Smart FlashCache
- Exadata RDMA Memory

Compute, memory, and process controls are provided by the Oracle Database Resource Manager (DBRM) across all platforms, while I/O and storage-related controls are provided by the Exadata I/O Resource Manager (IORM) specifically on Exadata platforms. Oracle also offers I/O rate limit controls for non-Exadata environments, but Exadata storage includes deep integration with the Oracle database, allowing much more robust control over storage and I/O.

Resource Management protects databases from what is known as the “noisy neighbor” problem, where one application (or database) can consume more than the desired amount of resources. For example, runaway SQL statements within one database could impact other databases sharing the same compute and storage resources. Oracle Resource Manager prevents one database (including runaway SQL statements within that database) from impacting other databases on the same system. Resource Manager Consumer Groups can also limit the impact of runaway SQL within an individual database.

Exadata Network Resource Management is automatic and does not require administrator intervention to manage network traffic on the Exadata storage and cluster interconnect network. This feature automatically prioritizes critical network messages over less critical messages to ensure the performance of Exadata systems.

## **Resource Starvation (Performance Issues)**

Resource starvation occurs when a database attempts to use more resources than allowed by the resource plan. The most common causes of resource starvation are application coding errors (including SQL coding errors), accidentally dropped or otherwise missing indexes, and incorrect optimizer statistics. These issues result in what are called “runaway” SQL statements, which can cause the database to consume (or attempt to consume) more than the amount of CPU, I/O, and other system resources allowed by the Resource Plan. The affected database is therefore constrained by Resource Manager, which may impact ALL users of the database, but will not impact other databases on the same system. Resource Manager Consumer Groups can be created to limit the impact of runaway SQL statements within a database.

When resource starvation occurs due to runaway SQL statements or other unexpected heavy workloads, DBRM and IORM are not the root cause of performance issues. Rather, the runaway application SQL statement or workload is the cause, and DBRM and IORM are simply imposing the defined resource limits to prevent the database from impacting the performance of other databases on the system.

## **Resource Manager Improves Availability**

The Oracle Resource Managers (DBRM and IORM) ensure databases do not consume more than the provisioned amount of system resources. As outlined in this section, these resource controls govern the use of CPU, memory, O/S processes, and I/O. Resource Manager prevents applications and databases from consuming excessive resources, which can result in availability issues and potential downtime of those systems. Proper resource management, as outlined in this section, can protect against application failures such as connection storms, memory leaks in PL/SQL applications, and runaway application SQL statements.

## **Hierarchy of Resource Allocations**

The resources of a system are subdivided through successive layers in a hierarchy according to the business needs of each application. Resources can be allocated and managed down to the level of granularity required to control costs and ensure databases are isolated from each other to prevent “noisy neighbor” issues, where one database, application service, microservice, job task, or user affects others in the system. The hierarchy of resource allocation includes the following:

- Physical Servers
- Virtual Machines
- Databases or Container Databases

- Pluggable Databases
- Consumer Groups
- Users and Jobs

Systems running Oracle databases or a group of Oracle databases may be comprised of multiple physical servers. Each physical server includes a given amount of resources (CPU, Memory, and effective process limits) that can be subdivided into multiple Virtual Machines. The resources of a Virtual Machine can be subdivided and allocated across one or more databases or container databases. In the case of Container Databases, the resources assigned to each can also be subdivided and allocated to 1 or more Pluggable Databases. Finally, the resources within one database can be allocated to various consumer groups, which are effectively individual users, groups of users, job classes, etc.

### **Over-Subscription of Resources**

When system resources are carefully divided among databases, administrators tend to allocate more than the resources needed by a given database at any specific point in time. Resources are often allocated for PEAK processing periods but are underutilized during periods of lower demand. For example, retail businesses often allocate resources for peak seasons, tax-related businesses allocate resources based on user needs during tax filing season, and financial accounting systems have resources allocated to meet end-of-month, end-of-quarter, and end-of-fiscal-year processing needs.

Over-subscription is a technique for taking advantage of the peaks and valleys of processing requirements that may differ between applications. Over-subscription effectively means allocating MORE resources than are physically present on a system under the assumption that users and databases will not experience simultaneous peak demands for those resources. In this section, we will explore the topic of over-subscription, including guidelines for each resource type, guidelines by workload, and by environment, such as development, test, production, and disaster recovery.

### **Over-Subscription Among Pluggable Databases**

Container Databases provide a mechanism to allow over-subscription of ALL resources among the Pluggable Databases within the Container without causing system stability issues. Performance of databases may suffer, but the operating system will not be impacted because the contained databases do not exceed what can be supported at the system level.

### **Over-Subscription of Memory**

Memory should never be over-subscribed at the VM level, database level, or between Container Database due to the severe performance and stability issues that will result,

including node evictions, node failures, and extreme paging/swapping. Memory should be carefully divided among Virtual Machines, then subdivided in turn between databases, and container databases. Operating stability or severe performance issues can result due to paging of memory. Pluggable databases within a container can use memory over-subscription without such performance impacts.

### **Over-Subscription of Processes**

Each operating system image (whether running on bare metal or within a Virtual Machine) can only support a certain number of O/S processes based on the compute resources (processor cores) assigned. Oracle recommends a maximum in the range of 64 database sessions for each processor core assigned to the database. The optimal number of sessions for a particular database will vary based on workload. Over-subscription of processes assumes some databases will be idle while others are active, such as multiple test, QA, or development databases residing on a single server, Virtual Machine, or Cluster.

### **Over-Subscription of Compute and I/O**

It is possible to over-subscribe compute and I/O resources assigned at the Virtual Machine, database, or Container Database level without causing system stability issues. However, this represents a trade-off between capacity utilization and potential application performance impact. Operating systems are able to effectively manage the prioritization of critical work without suffering stability issues when computing (processor cores) use a moderate level of over-subscription.

### **Over-Subscription by Criticality**

Over-subscription of resources ensures full use of available resources, but does not provide adequate control over resource usage. For example, ALL available CPUs within a bare metal server or a Virtual Machine can be allocated to 2 databases. If one database experiences high levels of activity while the other database has less activity, the over-provisioning ensures all of the resources can be utilized. However, if BOTH databases experience high levels of activity simultaneously, the databases will be in contention with each other for those resources, and one or both databases will suffer performance issues that are difficult to control. In extreme circumstances, the server can become unstable or unresponsive due to over-committed system resources.

For any mission-critical or “tier-1” systems, over-subscription should be used judiciously to avoid system performance issues. As shown in the table below, memory should only be over-subscribed when using Pluggable Databases because the Container guards against over-subscription at the system level. CPU, processes, and I/O may be over-subscribed depending on business needs, especially for less critical environments. Table 4 (below)

shows 5 systems with 4 classification levels. Both TEST and QA are considered the same level of criticality in this example.

**Table 4: Guidelines for Over-Subscription by Environment**

CRITICALITY	ENVIRONMENT	CPU	PROCESSES	MEMORY	I/O
1	Critical Production	Not to exceed 75% of CPU Threads	64 x Cores	NO	Limit 2x I/O Share
2	Non-Critical Production	Equal CPU Threads	64 x Cores	NO	Limit 2x I/O Share
3	Test	2 x CPU Threads	128 x Cores	NO	Limit 2x I/O Share
4	QA	2 x CPU Threads	128 x Cores	50% (PDB only)	Limit 4x I/O Share
5	Development	2 x CPU Threads	128 x Cores	50% (PDB only)	Limit 4x I/O Share

### Enabling Database Resource Manager (DBRM)

While it is possible to create detailed resource management plans using DBRM, the vast majority of non-CDB database deployments should simply use the Default Plan. Enabling the default resource management plan is done by executing the following command in each non-CDB database:

```
SQL> ALTER SYSTEM SET
      RESOURCE_MANAGER_PLAN = 'DEFAULT_PLAN'
      SID = '*'
      SCOPE=both;
```

Container Databases use the DEFAULT\_CDB\_PLAN as shown in the following command. This setting covers all Pluggable Databases included under the Container Database.

```
SQL> ALTER SYSTEM SET  
  
    RESOURCE_MANAGER_PLAN = 'DEFAULT_CDB_PLAN'  
  
    SID = '*'  
  
    SCOPE=both;
```

The above commands assume an SPFILE is being used for all databases (as indicated by the SCOPE=both clause), which sets the resource manager plan in memory of each instance as well as in the spfile. The SID='\*' clause applies the setting to all instances in a non-RAC or Oracle RAC environment. Use of SPFILE is considered “best practice” for managing Oracle databases.

### Enabling PDB Performance Profiles

The DEFAULT\_CDB\_PLAN at the container level is sufficient in cases where the “allocation method” is used at the Pluggable Database level. However, taking full advantage of the resource-sharing capabilities of Pluggable Databases requires the use of PDB Performance Profiles.

```
DECLARE  
  
    v_plan VARCHAR2(30) := 'cdb_profile_plan';  
  
BEGIN  
  
    DBMS_RESOURCE_MANAGER.clear_pending_area;  
  
    DBMS_RESOURCE_MANAGER.create_pending_area;  
  
  
    DBMS_RESOURCE_MANAGER.create_cdb_profile_directive(  
  
        plan                => v_plan,  
  
        profile              => 'xxsmall',  
  
        shares               => 1,  
  
        utilization_limit    => 4,  
  
        parallel_server_limit => 4);  
  
  
    DBMS_RESOURCE_MANAGER.validate_pending_area;
```

```
DBMS_RESOURCE_MANAGER.submit_pending_area;  
END;  
/
```

The above example creates the PDB Performance Profile named “xxsmall” and the pluggable database is altered to use this PDB Performance Profile as follows:

```
ALTER SESSION SET CONTAINER=pdb1;  
ALTER SYSTEM SET DB_PERFORMANCE_PROFILE=xxsmall SCOPE=BOTH;
```

All performance profiles required by PDBs within a Container Database should be created and made available to the Pluggable Databases within that Container. The Performance Profile can be included in a lockdown profile (in Oracle 12cR2 and high releases) in cases where PDBs are managed by a separate administrator other than the CDB administrator.

### **Pluggable Database Lockdown Profiles**

Large-scale deployments of Pluggable Databases often involve two layers of administrators, with a “fleet administrator” who manages the Container Database level, and one or more administrators who separately manage Pluggable Databases within each Container Database. Lockdown Profiles (available in Oracle 12cR2 and later releases) allow the CDB administrator to control resources allocated to Pluggable Databases, preventing PDB administrators from allocating beyond the approved amount of resources for each database.

### **Enabling Exadata I/O Resource Manager (IORM)**

On an Exadata VM cluster or in a bare metal configuration, all databases share dedicated storage servers which include flash storage and Exadata RDMA Memory. By default, the databases are given equal priority with respect to these resources. The Exadata software uses a first come, first served approach for query processing.

When there is contention for I/O resources, Exadata I/O Resource Manager (IORM) schedules I/O by immediately issuing some I/O requests and queuing others. I/O requests are serviced immediately for workloads that are not exceeding their resource allocation, according to the resource plans. I/O requests are queued for workloads exceeding their resource allocation. Queued I/Os are serviced according to the priorities in the resource plans when the workload no longer exceeds its resource allocation. When the storage server is operating below capacity and there is no contention for I/O resources, IORM does not queue I/O requests, and lets a workload exceed its resource allocation to consume available I/O resources. The IORM “objective” option specifies the optimization mode for IORM. The objectives available are ‘auto’, ‘balanced’, ‘high\_throughput’, and ‘low\_latency’. By default, IORM is enabled for all environments and uses the ‘auto’ I/O resource

objective. The 'auto' objective causes IORM to determine the best mode on active workloads and resource plans.

In case there is a need to change the IORM objective, the below example shows the use of the 'dcli' command to set the IORM Plan Objective to 'low\_latency' across all Exadata storage servers in the system. Alternatively, it is also possible to execute the commands on each storage server using the CELLCLI utility. To deliver consistent overall system performance, ensure every storage server in the cluster uses the same IORM configuration settings.

```
dcli-g ~/cell_group
    -l root
    cellcli -e alter iormplan objective = low_latency;
```

### **Exadata IORM Plans and Profiles**

The Exadata I/O Resource Manager (IORM) uses PLANS for individual databases, or PROFILES assigned to multiple databases with similar resource requirements. IORM allows control of the following resources in an Exadata system:

- Share of I/O relative to other databases
- Limit of I/O available on the system
- Flash Cache Limit (upper limit)
- Flash Cache Size (reservation of space)
- XRMEM (Exadata RDMA Memory) Limit (upper limit)
- XRMEM (Exadata RDMA Memory) Size (reservation)

The SHARE clause gives each database a minimum share of I/O resources on the system relative to the total number of databases on that system. For example, if 10 databases are each assigned share=10, each database receives at least 10% of the total I/O resources. Databases can consume MORE than the assigned share of resources if other databases are idle.

While SHARE sets the minimum resources available to a database, the LIMIT clause sets the maximum I/O resources for databases. Assigning LIMIT to each database results in greater consistency of resource allocation, resulting in more consistent performance. However, the use of LIMIT prevents databases from using idle resources, resulting in under-utilized IT assets.

IORM plans for individual databases use a plan name equal to the DB\_UNIQUE\_NAME of the database. Profiles are DBPLANS that include the type=profile clause to designate the

directive as a profile rather than a plan for an individual database, as shown in the example below:

The following Exadata dcli command creates or alters an IORM plan for the database named ORCL across all storage servers in the configuration. IORM Plans are based on the DB\_UNIQUE\_NAME of each database.

```
dcli -g ~/cell_group
      -l root cellcli
      -e alter iormplan DBPLAN=((name=ORCL, share=10,
limit=21))
```

The following Exadata dcli command uses type=profile to create or alter an IORM profile named “myprofile” across all storage servers in the configuration.

```
dcli -g ~/cell_group
      -l root cellcli
      -e alter iormplan DBPLAN=((name=myprofile, share=10,
limit=21,
      type=profile))
```

Every database using the same IORM profile will receive the same SHARE, LIMIT, FlashCache, and XRMEMCache allocations. It is recommended to assign databases to a small number of profiles (or “shapes”) such as SMALL, MEDIUM, and LARGE for manageability. See the section on Resource Manager Shapes for more discussion of this concept.

## Setting the Database Performance Profile (Exadata IORM)

In an Exadata environment, databases (including non-PDB and PDB) can use an IORM Performance Profile, which eliminates the need to set the IORM Plan for each database individually. All databases using the same profile will receive the same IORM resource allocations set in the IORM Profile (see the following section).

```
SQL> ALTER SYSTEM SET  
  
        DB_PERFORMANCE_PROFILE = 'myprofile'  
  
        SID = '*'  
  
        SCOPE=both;
```

The DB\_PERFORMANCE\_PROFILE setting in the database must simply match the profile name set in the Exadata storage IORM plan as outlined in the previous section.

## I/O Resource Management for Non-Exadata Platforms (PDB only)

For Oracle database 12c Release 2 and above using Oracle Multitenant on non-Exadata platforms, the following settings provide basic I/O Resource Management capabilities for Pluggable Databases within a Container Database:

- MAX\_IOPS
- MAX\_MBPS

While these database parameter settings provide some I/O Resource Management capabilities, it is important to note that non-Exadata storage does not include Exadata Smart FlashCache, and there is no substitute in non-Exadata systems for the IORM settings that govern Flash Cache Size and Flash Cache Limit.

It is important to note that non-Exadata systems do not have the storage network resource management features of Exadata (found even in the “basic” I/O Resource Plan) that ensure time-critical I/O is prioritized above non-critical I/O. Examples of time-critical I/O events include log file sync, log file parallel write, file header writes, data block writes performed during checkpoints, and data block writes performed to free space in the buffer cache. Criticality of I/O is related to the object being written as well as activity within the database. This capability depends on the integration of hardware with database software, which is only possible in Exadata.

## O/S Process Controls

To prevent excessive numbers of database processes, database administrators should target an active session count less than or equal to the CPU\_COUNT for each database. In addition, administrators should target an upper bound on the total session count for all databases within a physical server or Virtual Machine to 64 times the total cores in the physical server or Virtual Machine (32 times CPU\_COUNT in a server with 2 hyper-threads per core). High process counts on any operating system can lead to memory paging, CPU contention, and even node instability. High numbers of database connections do not always result in higher application throughput and can increase contention and decrease application performance. The number of sessions should also be decreased when using Parallel Query to devote more system resources to parallelism.

In addition to the overall count of processes in a system, the arrival rate of connections in what is known as a “login storm” can also result in system stability issues as the O/S attempts to service those connections. By using one or more of the following techniques, the overall process count and connection arrival rates can be reduced to improve performance and stability of the system.

**Configure Oracle listeners** to throttle incoming connections to avoid logon storms after a database node or instance failure.

**Limit the number of processes** and connections to the database servers by using connection pools and setting the maximum number of connections to a value slightly above the estimated active working sessions.

**Avoid the expensive allocation** and de-allocation of processes by eliminating dynamic connection pools by setting MIN and MAX connections to be the same.

Oracle listeners can be configured to throttle incoming connections to avoid logon storms after a database node or instance failure. The connection rate limiter feature in the Oracle Net Listener enables a database administrator (DBA) to limit the number of new connections handled by the listener. When this feature is enabled, Oracle Net Listener imposes a user-specified maximum limit on the number of new connections handled by the listener every second.

Depending on the configuration, the rate can be applied to a collection of endpoints or a specific endpoint. This feature is controlled through two listener.ora configuration parameters as follows:

`CONNECTION_RATE_<listener_name>=number_of_connections_per_second` sets the global rate that is enforced across all listening endpoints that are rate-limited. When this

parameter is specified, it overrides any endpoint-level numeric rate values that might be specified.

RATE\_LIMIT indicates that a particular listening endpoint is rate-limited. The parameter is specified in the ADDRESS section of the listener endpoint configuration. When the RATE\_LIMIT parameter is set to a value greater than 0, the rate limit is enforced at that endpoint level.

Example: Throttle new connections to prevent logon storms.

```
APP_LSNR= (ADDRESS_LIST=
  (ADDRESS= (PROTOCOL=tcp) (HOST=) (PORT=1521) (RATE_LIMIT= 10))
  (ADDRESS= (PROTOCOL=tcp) (HOST=) (PORT=1522) (RATE_LIMIT=20))
  (ADDRESS= (PROTOCOL=tcp) (HOST=) (PORT=1523))
)
```

In the preceding example, the connection rates are enforced at the endpoint level. A maximum of 10 connections is processed through port 1521 every second. The connections through port 1522 are limited to 20 every second. Connections through port 1523 are not limited. When the rate is exceeded, a TNS-01158: Internal connection limit reached is logged. Please refer to the Oracle Net Services Reference Guide for more information.

### Idle Session and Idle Blocking Session Timeouts

Oracle Database Resource Manager (DBRM) can be used to control the timeout of idle sessions and idle blocking sessions to ensure overall system stability. The following settings are used to control both idle sessions as well as idle sessions that are blocking other sessions:

- MAX\_IDLE\_TIME is used to control idle sessions in general (blocking or not)
- MAX\_IDLE\_BLOCKER\_TIME is used to control timeouts for blocking sessions

The MAX\_IDLE\_TIME setting can also be controlled using a database parameter or as a Resource Plan Directive. The MAX\_IDLE\_BLOCKER\_TIME can only be set as a Resource Plan Directive. Oracle recommends the use of MAX\_IDLE\_BLOCKER\_TIME rather than MAX\_IDLE\_TIME in cases where applications commonly leave idle sessions in the system, or where applications do not easily tolerate the termination of idle sessions.

### Database & PDB Resource Controls

Resources for each database are controlled using a core set of controls within each Pluggable Database (PDB) or stand-alone (non-PDB) database. DBRM (and IORM in Exadata) uses these settings to ensure proper management of available system resources to ensure database performance at the desired level, as well as to ensure system stability.

Table 5 shows the set of core inter-database resource management controls and the enhancements in these controls from release to release from 12cR1 through 26ai.

**Table 5: Database Resource Controls by Release (19c – 26ai)**

	19C	26AI
RESOURCE_MANAGER_PLAN	✓	✓
CPU_COUNT	✓	✓
SGA_TARGET	✓	✓
PGA_AGGREGATE_TARGET	✓	✓
PARALLEL_MAX_SERVERS	✓	✓
SESSIONS	✓	✓
PGA_AGGREGATE_LIMIT	✓	✓
MAX_IDLE_BLOCKER (Resource Plan Directive)	✓	✓
DB_PERFORMANCE_PROFILE	✓	✓
MAX_IDLE_TIME	✓	✓

As seen in the table above and detailed above, new resource management settings can be set in Oracle 19c, and in Oracle 26ai. These settings improve the manageability of resources and address availability issues that might arise due to resource starvation at the database or system level.

The PGA\_AGGREGATE\_LIMIT setting was introduced in 12cR1 to provide control over the total amount of PGA memory allocated by all sessions in the database. This setting improves overall system availability by limiting excessive memory allocations in the PGA, which can arise from a variety of conditions, including overallocation of memory in PL/SQL code.

The DB\_PERFORMANCE\_PROFILE setting can be used in Oracle Database 19c and Oracle AI Database 26ai to allow the use of IORM profiles, which allow a single profile to control resources for many databases. In earlier releases, IORM plans are set for each individual database according to the DB\_UNIQUE\_NAME of the database.

The MAX\_IDLE\_TIME setting is enabled for automatic termination of idle database sessions, which improves database and system availability by terminating idle sessions that are holding system resources. Some applications don't tolerate the termination of

idle sessions well. Use of MAX\_IDLE\_BLOCKER\_TIME might be preferable in those circumstances.

The MAX\_IDLE\_BLOCKER\_TIME setting was introduced in earlier versions of Oracle to enable automatic termination of idle database sessions that are blocking other sessions, which improves database and system availability by terminating idle sessions that are also blocking other sessions.

The CPU\_COUNT is an instance-level setting that specifies the number of threads (on hyperthreaded systems) that can be used by the instance. Each Oracle Real Application Clusters (Oracle RAC) database includes one or more instances across nodes of a cluster. It is important to note that the other parameters discussed in this section (with the exception of DB\_PERFORMANCE\_PROFILE) are also instance-level settings.

The memory and O/S process controls above should be used even in non-consolidated environments to prevent overallocation of system resources, which can lead to system availability issues. Overallocation of PGA memory (such as in a PL/SQL application) can lead to memory starvation. Applications that cause connection storms over-allocate O/S processes and lead to starvation of resources at the O/S level.

For Pluggable Databases, these settings can also be configured in a “lockdown profile” at the CDB level, preventing individual PDB database administrators from modifying these settings. This capability is important in large-scale consolidation environments with multiple database administrators.

### **Intra-Database Resource Management - Consumer Groups**

While this document deals primarily with provisioning and management of resources ACROSS databases, the Oracle Database Resource Manager also provides the ability to manage resources within a single database. Resource Manager Consumer Groups can be used to assign system resources to users, sessions, or connections to those consumer groups using various rules and policies, and manage the use of those resources. For more information on intra-database Resource Management, please refer to MOS Note# [KB165271](#) (Master Note: Overview of Oracle Resource Manager and DBMS\_RESOURCE\_MANAGER, Doc ID [KB165271](#)).

## Runaway Queries

Resource Manager Consumer Groups provide the ability to manage sessions that consume more than the allowed amount of resources, such as sessions executing “runaway” queries. This capability is important in any environment that allows ad-hoc queries, such a Data Warehouse and Analytic systems. Resource Manager allows the following actions on any session or query consuming excessive resources:

- Switch Group
- Cancel SQL
- Kill Session
- Log Only

Resource Manager Plan Directives are used to switch a session to a different (usually lower priority) consumer group, cancel the specific SQL statement, kill (cancel) the user session, or simply LOG the event for later analysis by administrators. While addressing runaway queries is an important topic, it is beyond the scope of this document. Please refer to MOS Note# [KB165271](#) (Master Note: Overview of Oracle Resource Manager and DBMS\_RESOURCE\_MANAGER, Doc ID [KB165271](#)) for more information.

## Using Resource Controls

The following table provides a summary of system resources and the controls for those resources in Oracle Database Resource Manager (DBRM) and the I/O Resource management capabilities of Exadata, as well as settings for non-Exadata systems.

**Table 6: Controlling System Resources**

Resource	DB (Stand-alone)	CDB	PDB
<b>CPU</b>	<b>Instance Caging Method</b> CPU_COUNT (with a minimum of 2) RESOURCE_MANAGER_PLAN = 'DEFAULT_PLAN'	RESOURCE_MANAGER_PLAN = 'DEFAULT_CDB_PLAN'	<b>Control via Parameters:</b> CPU_COUNT <b>Control via CDB Resource Plan:</b> SHARE = share of CPU resources per PDB UTILIZATION_LIMIT = Limit of resources for a PDB
<b>Memory</b>	SGA_TARGET PGA_AGGREGATE_TARGET PGA_AGGREGATE_LIMIT (defaults to 2X target)	<b>CDB Level:</b> Applies to the entire container.  SGA_TARGET PGA_AGGREGATE_TARGET PGA_AGGREGATE_LIMIT (defaults to 2X target)	<b>Control via Parameters:</b> SGA_TARGET PGA_AGGREGATE_TARGET PGA_AGGREGATE_LIMIT (defaults to 2X target) <b>Control via CDB Resource Plan:</b> SHARE = share of CPU resources per PDB

			UTILIZATION_LIMIT = Limit of resources for a PDB
<b>Processes and Sessions</b>	<p><b>Database Settings:</b></p> <p>SESSIONS</p> <p>MAX_IDLE_TIME</p> <p>MAX_IDLE_BLOCKER_TIME</p> <p>PARALLEL_MAX_SERVERS</p> <p>MAX_IDLE_BLOCKER_TIME</p> <p>PARALLEL_TARGET_SERVERS</p> <p><b>SQL Net Controls:</b></p> <p>CONNECTION_RATE_listener_name</p> <p>RATE_LIMIT</p>	<p><b>CDB Settings</b></p> <p>SESSIONS</p> <p>MAX_IDLE_TIME</p> <p>MAX_IDLE_BLOCKER_TIME</p> <p>PARALLEL_MAX_SERVERS</p> <p>MAX_IDLE_BLOCKER_TIME</p> <p>PARALLEL_TARGET_SERVERS</p> <p><b>SQL Net Controls:</b></p> <p>CONNECTION_RATE_listener_name</p> <p>RATE_LIMIT</p>	<p><b>Control via Parameters:</b></p> <p>SESSIONS</p> <p>MAX_IDLE_TIME</p> <p>MAX_IDLE_BLOCKER_TIME</p> <p>PARALLEL_MAX_SERVERS</p> <p>MAX_IDLE_BLOCKER_TIME</p> <p>PARALLEL_TARGET_SERVERS</p> <p><b>Control via CDB Resource Plan:</b></p> <p>SHARE = share of CPU resources per PDB</p> <p>UTILIZATION_LIMIT = Limit of resources for a PDB</p>
<b>I/O Resource Management</b>	<p><b>Exadata</b></p> <p>IORM Objective = 'auto'</p> <p>IORM Share – Specifies a share of I/O resources relative to other databases on the system.</p> <p>IORM Limit – Places an upper bound (in %) on I/O used by the database.</p> <p>FlashCache Min – Minimum amount of FlashCache for the specified database.</p> <p>FlashCache Limit – Can use up to this size, depending on other activity in the system.</p> <p>FlashCache Size – Reserves this space for the database.</p> <p>XRMemCache Min – Minimum amount of XRMEM Cache for the specified database</p> <p>XRMemCache Limit – Can use up to this size, depending on other activity in the system.</p> <p>XRMemCache Size – Reserves this space for the database</p> <p><b>Non-Exadata Systems</b></p> <p>12cR2 and above – PDB Only (NOT on Exadata)</p> <ul style="list-style-type: none"> <li>» MAX_IOPS</li> <li>» MAX_MBPS</li> </ul>		
<b>Network</b>	<p><b>Exadata</b></p> <p>Network resource management is automatic – no settings required</p> <p><b>All Platforms (including Exadata)</b></p> <p>Activate Oracle RAC databases on a minimum number of instances necessary for the workload.</p>		

Use of parameters to control PDB compute resources (CPU\_COUNT) gives greater control over those resources compared to the use of shares and limits in a CDB Resource Plan.

## Standardized DB Resource Shapes

The use of standardized resource shapes greatly simplifies resource management. Standardized shapes use a common definition for the relationship among resources. All resources in a system can be allocated as simple functions of the CPU count used for each database running on that system. For example, if a database is allocated 10% of the CPU, it should also receive 10% of the memory, 10% of the I/O, and 10% of the available FlashCache in Exadata systems.

## Shapes

Oracle recommends establishing 2 types of resource shapes that contain the full range of resource allocations by workload type. The same overall ratios will be applied to each, but with several variations needed to meet the resource demands of each. The recommended types of resource shapes are:

- Data Warehouse (DW)
- Online Transaction Processing (OLTP)

Data Warehouse workloads typically require larger PGA (Process Global Area) and smaller SGA (System Global Area), whereas OLTP workloads require larger SGA and smaller PGA. Likewise, Data Warehouse workloads tend to have fewer concurrent user sessions, with heavier use of Parallel Query processing. OLTP workloads are normally “mixed” workloads with primarily non-parallel operations, and a smaller percentage of Parallel Query work.

While the overall size of the DW and OLTP shapes is nearly identical in terms of CPU, Memory, Process, and I/O allocations, the distribution of memory and processes differs as shown in the tables below.

## Multiples

The standardized shapes discussed above establish common ratios between all of the system resources. These shapes are then scaled using multiples to allocate resources to databases. Using standardized shapes and multiples in this manner simplifies system resource allocations, while still allowing administrators to provision the proper amount of resources for each application.

## Database Resource Shape Examples

The standardized database resource shapes outlined in this section will be used later during the bin-packing process to determine the size of the “bins” (Server Resource Shapes and Container Databases) as well as to arrange databases into those bins. The CPU\_COUNT required for each database should be used to establish the proportion of all other system resources allocated to each database. Note that these are DATABASE resource shapes as opposed to Virtual Machine or Bare Metal shapes. Database resource

shapes are layered on top of underlying Virtual Machine or Bare Metal shapes, and one Virtual Machine or Bare Metal Exadata server may include more than 1 database.

Tables 7 and 8 show standardized database resource shapes for OLTP and DW workloads using the following Exadata configuration:

- Exadata Cloud@Customer X11M
- 8 Database Servers
- 190 OCPUs per Database Server (380 vCPU per Database Server)
- 8 Virtual Machines per Server
- 8 Exadata Storage Servers (27.2 TB FlashCache per Storage Server)

These Database Resource Shapes apply to stand-alone databases as well as at the Container Database level. While these same database resource shapes can be used with Pluggable Databases (PDBs), it is also possible to use SHARE and LIMIT settings for greater flexibility and greater resource sharing with PDBs. Please refer to the section on PDB Resource Shapes for more information. Each database or container database should be configured to run on a minimum of 2 database nodes using Oracle Real Application Clusters (Oracle RAC) to provide high availability. It is important to note that memory and process settings are per-node settings, whereas IORM applies to databases regardless of the number of nodes.

**Table 7: OLTP DB Resource Shapes – Exadata Cloud@Customer X11M**

Shape	CPU_COUNT	Nodes	Total	Memory				Processes		IORM		
				DB Memory	SGA	PGA Target	PGA Limit	Sessions	PQ Processes	Share	Limit %	FlashCache Limit
OLTP Small 2V	4	VM 2	8	30 GB	15 GB	7.5 GB	15 GB	128	4	8	5%	3.7 TB
OLTP Medium 2V	8	VM 2	16	60 GB	30 GB	22 GB	30 GB	256	8	16	5%	7.4 TB
OLTP Large 2V	16	VM 2	32	120 GB	60 GB	30 GB	60 GB	512	16	32	5%	14.9 TB
OLTP 2X Large 2V	32	VM 2	64	240 GB	120 GB	60 GB	120 GB	1024	32	64	8%	29.8 TB
OLTP 3X Large 2V	48	VM 2	96	344 GB	172 GB	86 GB	172 GB	1536	48	96	12%	44.8 TB
OLTP 3X Large 4V	48	VM 4	192	344 GB	172 GB	86 GB	172 GB	1536	48	192	25%	89.6 TB
OLTP 6X Large 2V	96	VM 2	192	694 GB	347 GB	172 GB	347 GB	3072	96	192	25%	89.6 TB
OLTP 12X Large 4V	96	VM 4	384	694 GB	347 GB	172 GB	347 GB	3072	96	384	50%	179.2 TB
OLTP 15X Large 4V	124	VM 4	492	1390 GB	695 GB	347 GB	695 GB	3072	96	384	50%	179.2 TB
OLTP 24X Large 8V	124	VM 8	992	1390 GB	695 GB	347 GB	695 GB	3072	96	768	100%	358.4 TB

As outlined in the previous section, Data Warehouse (DW) workloads typically use a smaller SGA than OLTP systems, with more memory allocated to the PGA to optimize for aggregations. Additionally, DW workloads typically involve fewer users and heavier use of parallelism. Consequently, DW systems use lower session limits, with more PQ processes than OLTP systems. In all of these example shapes, Oracle RAC is used with a minimum of 2 instances (or nodes) per Oracle RAC cluster for scaling and high availability.

**Table 8: DW DB Resource Shapes – Exadata Cloud@Customer X11M**

Shape	CPU_COUNT	Nodes	Total	Memory				Processes		PDB & IORM		
				DB Memory	SGA	PGA Target	PGA Limit	Sessions	PQ Processes	Share	Limit %	FlashCache Limit
DW Small 2V	4	VM 2	8	30 GB	10 GB	10 GB	20 GB	32	16	8	5%	3.7 TB
DW Medium 2V	8	VM 2	16	60 GB	20 GB	20 GB	40 GB	64	32	16	5%	7.4 TB
DW Large 2V	16	VM 2	32	120 GB	40 GB	40 GB	80 GB	128	64	32	5%	14.9 TB
DW 2X Large 2V	32	VM 2	64	240 GB	80 GB	80 GB	160 GB	256	128	64	8%	29.8 TB
DW 3X Large 2V	48	VM 2	96	350 GB	120 GB	120 GB	240 GB	384	192	96	12%	44.8 TB
DW 3X Large 4V	48	VM 4	192	350 GB	120 GB	120 GB	230 GB	384	192	96	25%	89.6 TB
DW 6X Large 2V	96	VM 2	192	690 GB	230 GB	230 GB	460 GB	768	384	192	25%	89.6 TB
DW 12X Large 4V	96	VM 4	384	690 GB	230 GB	230 GB	460 GB	768	384	384	50%	179.2 TB
DW 15X Large 4V	124	VM4	496	1390 GB	460 GB	460 GB	920 GB	768	384	576	75%	268.8 TB
DW 24X Large 8V	124	VM 8	992	1390 GB	460 GB	460 GB	920 GB	768	384	768	100%	358.4 TB

Use of extremely low values for UTILIZATION\_LIMIT (below 5%) should be avoided to prevent thrashing on the system. The above examples show the lowest utilization limits at 5%, with higher values being equal to 2X the share value.

### Resource Sharing Between Pluggable Databases

Using Oracle Multitenant, resources within a Container Database can be shared across Pluggable Databases residing within the Container. Sharing of compute resources within a CDB is controlled by the following settings of the resource plan for each PDB rather than using the strict CPU\_COUNT control (also known as Instance Caging) outlined in the previous section:

- Share
- Utilization Limit
- Parallel Max Servers
- Parallel Servers Target

The SHARE of compute resources for a database is relative to other databases within the container. For example, if there are 10 Pluggable Databases within a container, where each has a SHARE value equal to 10, each PDB receives at least 10% of the CPU resources of the Container. While SHARE provides a minimum allocation of available resources, Pluggable Databases within a Container can use MORE CPU resources if other databases are not active and consuming those resources.

The UTILIZATION\_LIMIT for a database places an upper bound on the amount of compute resources available to a PDB. The primary purpose of UTILIZATION\_LIMIT is to prevent or limit the variability in resources available to a database, which can lead to variability in end-user performance. Utilization limits will prevent the use of idle system resources, which typically reduces overall system utilization and does not allow users to take advantage of those resources, but will provide a more consistent end-user experience.

The PARALLEL\_MAX\_SERVERS setting places a limit on the number of parallel server processes that can be utilized by a Pluggable Database. The PARALLEL\_SERVERS\_TARGET setting establishes a target for shrinking the number of parallel servers when demand is low.

### **Pluggable Database Resource Shapes**

As noted previously, Pluggable Databases can use the allocation method outlined above for stand-alone databases and Container Databases, or Pluggable Databases can use PDB Performance Profiles as detailed in the table below. Using either method, Pluggable Databases receive a portion of the resources allocated to the Container Databases. Hard allocations of CPU, memory, and processes result in under-utilization of system resources overall because the PEAK amount of resources must be allocated to each database. Using the Shares & Limits method, the resources are shared between PDBs, with limits placed on the resource consumption of each Pluggable Database.

**Table 9: PDB Performance Profiles – Exadata X11M**

Profile Name	Nodes	Compute & Memory & Process Shares			IORM Shares		
		Shares	Utilization Limit	Parallel Server Limit	Share	Limit %	FlashCache Limit
PDB XX Small 2	VM 2	1	5%	4%	1	5%	9 TB
PDB Extra Small 2	VM 2	2	8%	8%	2	5%	1.8 TB
PDB Small 2	VM 2	4	16%	16%	4	5%	3.7 TB
PDB Medium 2	VM 2	8	32%	33%	8	5%	7.4 TB
PDB Large 2	VM 2	16	64%	66%	16	5%	14.9 TB
PDB 2X Large 2	VM 2	32	100%	100%	32	8%	29.8 TB
PDB 3X Large 2	VM 2	48	100%	100%	48	16%	44.8 TB
PDB 3X Large 4	VM 4	48	100%	100%	96	32%	89.6 TB
PDB 6X Large 2	BM 2	96	100%	100%	96	32%	89.6 TB
PDB 12X Large 4	BM 4	96	100%	100%	192	64%	179.2 TB
PDB 24X Large 8	BM 8	96	100%	100%	384	100.00%	358.4 TB

In the example shown in Table 9, each PDB resource shape is given a “share” of resources, as well as utilization limits equal to DOUBLE the amount of the share in what is known as a PDB Performance Profile. For example, in the “PDB XX Small” Performance Profile, a share of “1” is approximately 2% of the 48 cores on the configured Virtual Machine. LIMIT should not be set lower than 5% to prevent thrashing of system resources. Using the “Shares & Limits” method for PDBs is identical for OLTP and DW workloads because the appropriate resource allocations are assigned at the Container Database level. PDBs are simply taking a share (and limit) of the resources assigned to the container.

Notice that IORM shares and limits for PDBs follow the same settings for stand-alone databases, since Exadata I/O Resource Manager is independent of the compute tier.

### Implementing Database Consolidation

This chapter outlines the process for implementing database consolidation using the techniques outlined in this document. Consolidating databases follows the same process as when consolidating onto Exadata, Exadata Cloud@Customer, or Exadata Cloud Infrastructure. Implementing database consolidation should follow a standardized process that determines resource capacity needs for each database and arranges databases into groups for simplified management. The following steps outline the process:

1. Inventory Databases to be Consolidated
2. Gather Resource Utilization & Growth Metrics
3. Map Resource Requirements to New Platform
4. Determine Database Isolation Requirements & Methods

5. Select Consolidation Method for Databases
6. Group Databases into HA Tiers
7. Perform Bin-Packing of Databases into Resource Shapes
8. Create Resource Plans for Databases

### **Inventory of Databases to be Consolidated**

Database consolidation begins with an inventory of all databases to be consolidated. The following information should be collected for databases that are candidates for consolidation:

- Database Name
- Database Version
- Physical location, region, or data center
- Server Name
- Server Platform O/S
- Server O/S Version
- Associated Application(s)
- High Availability Tier (Bronze, Silver, Gold, Platinum)
- Environment (DEV, TEST, QA, Production, DR, etc.)

This list of candidate databases will be used throughout the consolidation planning process to determine which databases will be consolidated together on the same physical and/or virtual servers.

### **Gather Resource Utilization & Growth Metrics**

Resource utilization metrics are required for each database in the list of consolidation candidates. These metrics can be collected from the Oracle Automatic Workload Repository (AWR) for each database:

- CPU Utilization
- Processes & Sessions
- Memory
- Storage Space
- I/O

Metrics collected on existing systems are the same metrics collected for capacity monitoring and planning purposes. Please refer to the section of this document on measuring capacity utilization for more information on how to collect these metrics for planning migration to a database consolidation environment. Current system resource utilization analysis shows the current state as well as some amount of historical resource utilization, but growth can be more complex.

All organizations experience growth in data volumes and processing workloads. This growth results in higher demand for storage space, computing power, system memory, and I/O rates. Historical growth rates can often be used for forward planning, but certainly do not reflect anticipated changes in business conditions, changes in application functionality, or other external forces that impact capacity. For these reasons, it is often necessary to look beyond historical system statistics during the capacity planning process.

The planning horizon refers to the period for which capacity planning is being conducted. Most organizations perform capacity planning on a 3-5 year basis, and purchase hardware or acquire needed services for that period of time. For example, capacity planning might show a 10% compounded annual growth over 5 years, or 61% growth in total over that period. Year-to-year growth is not always consistent due to changes in business climate, corporate acquisitions, application functional changes, regulatory changes, and other factors. However, historical growth rates in prior years can often be used to forecast future growth rates.

Systems that are hitting resource limits do not provide an accurate view of resource utilization. Anytime demand is higher than supply, the true level of demand is unknown and will be higher than the supply of resources. Customers should be aware of this phenomenon and take this into account during capacity planning.

### **Map Resource Utilization to New Platform Characteristics**

Resource utilization is gathered on an existing system, but will typically be applied to a separate system during the consolidation process. Newer computing platforms will normally have faster processors, more memory, faster I/O, greater I/O throughput, and increased capacity and performance compared to older systems. Each system vendor uses their own methodology for determining system capacity requirements for a given set of applications or databases. Oracle uses system performance metrics known as M-Values to track performance differences across multiple generations of servers.

Consolidation onto a new system might also be a different architecture than older systems, such as when moving from traditional server & storage environments to Exadata. Platform architectural differences often have an impact on system capacity requirements for databases. For example, Exadata typically requires less relative compute power

compared to traditional systems due to the offloading of database processing into the Exadata storage tier. This is referred to as the Exadata Database Server Optimization and varies according to the workload and degree of storage offload for that workload. The benchmark shown earlier in this document showed a 2X greater consolidation density, or 50% lower compute resource consumption.

**Pent-up demand** is a phenomenon that should always be taken into account when mapping resource requirements to a new platform. This phenomenon is especially prevalent in Data Warehouse or Analytics systems, where the use of those systems is often limited by the performance of those systems. Improved performance of these systems often leads to increased use of those systems, resulting in unanticipated growth in resource requirements. This is especially true when moving Data Warehouse and Analytics databases to the Exadata platform.

### **Determine Isolation Requirements & Method**

All IT organizations have requirements to provide some degree of isolation of databases from others. The most common isolation requirements are:

- Environment (DEV, TEST, QA, PROD, DR, etc.)
- Data Governance (PII, PCI, HIPAA, etc.)
- Upgrade & Patching

Non-production environments (DEV, TEST, QA, etc.) typically must be separated from production environments, and data that is subjected to special governance regulations, such as Payment Card Industry (PCI) requirements, typically must be separated from data that is not subjected to those regulations. Two isolation methods are available for Oracle databases on Exadata systems:

- Physical Isolation
- Virtual Machine Isolation
- Isolation by Container Database

Although physically separate machines provide complete isolation between environments, Virtual Machines provide many of the same benefits within the same physical machine. Non-production databases (DEV, TEST, QA, etc.) should be isolated from production databases (PROD and DR) using physically separate machines. Isolation for data governance purposes can also be achieved with either physical or virtual machine isolation.

Use of multiple Container databases within a physical server, Virtual Machine, or physical/virtual cluster provides a further level of database isolation. Multiple Container

Database may reside on the same O/S, but can use different Oracle DBMS patches or version levels.

### **Select Consolidation Method for Databases**

The reference architecture for each HA tier supports all consolidation methods and deployment models. However, detailed best practices for the migration of standalone databases to a consolidated environment are influenced by the consolidation method used.

**Multiple Databases per Server** using Oracle RAC combined with Oracle Resource Manager is the MAA best practice for Oracle Database 12c and later releases. Oracle RAC provides availability and scalability, while Resource Manager provides isolation of resources.

**Oracle Multitenant**<sup>6</sup> is applicable for Oracle Database 19c and later releases. The multitenant architecture delivers the highest possible consolidation density and improves the manageability of large numbers of databases.

MAA best practices are equally relevant to physical or virtual environments. Schema consolidation has largely been replaced by Oracle Multitenant, but it is still an option if customers choose to combine multiple applications or microservices into a single database. The resulting database is then consolidated with others using one of the methods above.

When stand-alone databases are migrated to Oracle Multitenant, the process is as simple as converting each database into a PDB and moving that database to an appropriate CDB for the given HA tier.

### **Group Databases into HA Tiers**

All databases should be classified according to one of the HA Tiers outlined earlier in this document, including business tolerance for data loss (recovery point objective, or RPO) and downtime (recovery time objective, or RTO) for each database that is a candidate for consolidation. Any dependencies that might exist between databases should also be well understood to ensure those databases are properly grouped to meet business objectives.

The MAA best practice is to consolidate databases that have similar RTO (Recovery Time Objective) and RPO (Recovery Point Objective) according to the standard set of HA tiers described in the preceding section. Note that in cases where there are dependencies

---

<sup>6</sup> Starting Oracle Database 21c, multitenant is the only option

between databases, each database is assigned to the HA tier appropriate for the database having the most stringent HA requirement.

The process of grouping databases into standard HA tiers according to RTO and RPO requirements accomplishes three objectives:

Standardization on a limited set of HA tiers reduces complexity and achieves economies of scale.

Efficient consolidation by avoiding over-investment or unnecessary complexity in HA infrastructure and processes. For example, it would be inefficient to consolidate a Bronze Tier database (where RTO and RPO can be achieved by restoring a backup) with Silver or Gold MAA Tier databases, where RTO and RPO requirements dictate additional HA infrastructure.

Establishing the HA and data protection component of the [Service Catalog for DBaaS](#). The service catalog describes the services that an IT organization provides its user community - developers, architects, and end-users - with specifics on how database services are delivered and managed.

### **Perform Bin-Packing of Databases into Resource Shapes**

Once the isolation requirements, consolidation method, and HA Tier groupings (outlined above) have been completed, databases can then be arranged into Resource Shapes in a process known as bin-packing. The process of bin-packing involves determining which databases should be consolidated together into what is known as a Resource Shape (or “bin” in this context), which can be one of the following configurations:

- Physical Server
- Virtual Machine
- Physical Server Cluster
- Virtual Machine Cluster
- Container Databases

The bin-packing process determines the size requirement and configuration of each Resource Shape (or “bin”), including the size of Virtual Machines and the number of physical or Virtual Machines combined into a cluster. The bin-packing process should always be performed in order from largest to smallest, where size is measured by the compute resources assigned to each database, not by storage volume. The volume of storage space required for a Resource Shape is determined AFTER the bin-packing process is completed, based on the size of the associated databases. The required amount of storage is then allocated to each Resource Shape (physical server, Virtual

Machine, or physical/virtual cluster) according to the total size of databases assigned to that Resource Shape.

Any databases to be consolidated as Pluggable Databases (PDB) within a Container Database (CDB) should all belong to the same MAA Tier (Bronze, Silver, Gold, Platinum).

The underlying hardware may have a capacity limit that prevents all PDBs assigned to a particular HA tier from being consolidated into a single server or cluster of servers. Performance requirements must be assessed to determine the suitability of a database as a candidate for database consolidation, and to match existing managed servers that you want to consolidate with the servers that you will consolidate to.

It is often necessary to support more than one Oracle Database release/patch set at a time. For example, one CDB may be at Oracle Database 19c Release 1 (19.5.0.0) and a second CDB at Oracle AI Database 26ai. This enables an individual PDB to be ‘unplugged’ from an Oracle 19.5.0.0 CDB and ‘plugged’ into the Oracle 23.5.0.24 CDB to implement an upgrade

Bronze Tier of databases typically includes the largest number of databases eligible for consolidation and are the most likely candidates for use of Oracle Multitenant. Oracle recommends starting any consolidation effort with Bronze Tier databases due to the lower criticality of those applications and higher return on investment.

Container Databases represent a “bin” as well, with databases packed into the Container. Resources for a Container Database are sized to contain all of the Pluggable Databases the CDB will include. The CDB provides isolation from a version/patching perspective, but also serves to constrain resources used by the PDBs under the CDB.

### **Resource Manager Planning for Databases**

The Oracle Database Resource Manager (DBRM) and I/O Resource Manager (IORM) on xadata serve to ensure databases and the applications they service receive the proper amount of system resources to meet business needs.

- Individual Database Resource Plans
- Resource Profiles for Groups of Databases

Individual Resource Manager PLANS are created for any larger, more critical databases that require closer attention and more active management. Resource Manager PROFILES are used for smaller, less critical databases that often have more uniform resource requirements. Please refer to the Resource Manager section of this document for more information on the use of the Oracle Resource Managers (DBRM and IORM).

An overview of the consolidation planning process is described in Table 10.

**Table 10: Consolidation Planning**

Key Steps for Consolidation	Recommendations	Benefits
<b>1. Inventory databases to be consolidated</b>	The list of databases to be consolidated includes location, platform, application, availability tier, and other characteristics that guide consolidation decisions.	Determines the overall scope of the consolidation effort.
<b>2. Gather resource utilization &amp; growth</b>	Current system resource utilization (CPU, memory, I/O, etc.) provides a basis for resource planning on the consolidated database platform. Growth also must be taken into account to ensure the system meets the resources needs of the business. Resource utilization also determines Resource Manager configuration.	Needed for bin-packing exercise in step #7. Resource Management Plans are defined for each database and/or group of databases.
<b>3. Map Resource Requirements to New Platform</b>	Consult system vendors to determine the relative performance differences between older and newer generations of hardware.	New computing platforms typically deliver improved performance and higher capacity than older systems. Changes in system architecture, such as moving to Exadata, can also affect system capacity requirements.
<b>4. Determine Database Isolation Requirements &amp; Methods</b>	Isolation can be done by placement on separate physical machines or into separate Virtual Machines. Development and Test databases are normally isolated from Production databases. Databases subjected to special governance rules (PII, PCI, HIPAA, etc.) are normally isolated from other databases.	Meet business requirements for isolation, while minimizing complexity related to ongoing management of large numbers of physical servers or Virtual Machines.
<b>5. Select Consolidation Method for Databases</b>	Oracle Multitenant Container Databases (CDB) deliver the highest consolidation density possible. Multiple databases can also be consolidated onto Resource Shapes (single servers or VMs, or clusters of physical or virtual servers) without Oracle Multitenant. Consolidating onto Exadata provides a higher degree of consolidation density in conjunction with either approach. A combination of Multitenant and stand-alone databases can also be used.	Higher consolidation density reduces system infrastructure costs as well as operational costs for maintaining systems and databases.
<b>6. Group Databases into HA Tiers</b>	Follow MAA Tiers (PLATINUM, GOLD, SILVER, BRONZE). All databases within a CDB should be of the same tier.	Reduce operational expense and lower risk through standardization.
<b>7. Perform Bin-Packing of Databases into Resource Shapes</b>	Create appropriately sized Resource Shapes or “bins”. Half Racks up to one full rack (Cloud and Cloud@Customer), or two full multi-rack systems (on-premise only) for a typical resource shape.  Assign each resource shape to an HA Tier. Since each tier has a different HA architecture, there is only one HA tier per resource shape.  Pack databases into “bins” starting from largest to smallest.	Bin-packing exercise determines Resource Shape configurations.
<b>8. Resource Manager Plans for Databases</b>	Individual Resource Manager plans are created for each of the largest or most critical databases. Resource Manager profiles are used for smaller and less critical databases.	Ensures that the proper amount of system resources is provided to each database according to business needs. Resource Management also serves to ensure system stability is not jeopardized by faulty applications or databases.

## Conclusion

Database consolidation is a useful approach to meeting the goals of the modern enterprise, and is designed to balance 5 key business goals, including:

- Reducing **costs**
- **Simplifying** administration

- Improving **security**
- Ensuring database and application **availability**
- Delivering the levels of **performance** required for business applications


This document outlined the strategy for meeting these goals using Oracle Database, Virtual Machines, Oracle Multitenant, and Exadata, whether it is deployed in on-premises, Oracle Cloud, or in the Cloud@Customer model. In all of these deployment models, Exadata provides a higher level of consolidation density than other platforms, which reduces the cost of the overall platform. Administration is simplified through the consolidation of databases, but also through the use of the Exadata platform, which is delivered as an integrated solution of servers, storage, networking, operating system, virtual machines, clusterware, logical volume manager, and the Oracle database software.

Oracle Multitenant provides an even higher level of consolidation density (in addition to the high density provided by Exadata), reducing costs further, while providing simplified “manage many-as-one” administration to large-scale consolidation environments.

Oracle AI Database 26ai’s converged, multimodel engine supports various data types such as relational, document, AI vectors, graph, and text which enhances workloads of all types, like AI Vector Search, analytics, and mission-critical transaction processing, reducing the need for multiple specialized and proprietary databases. Exadata Exascale enhances consolidation with its intelligent data architecture, leveraging pooled compute and storage resources for hyper-elastic scaling. This allows tens to thousands of databases, including multiple versions, to run efficiently on a single Exadata environment, simplifying management, lowering operational costs, and improving resource utilization. Features like AI Smart Scan and database-integrated provisioning of thin clones further optimize performance and agility, making Exadata Exascale an ideal platform for enterprise-wide database consolidation.

Exadata Exascale is a new intelligent data architecture that delivers the best of Exadata and the best of Cloud. Combining the performance, availability, scalability, and security of Exadata, with hyper-elasticity, multi-tenancy, and resource pooling, Exascale is the next-generation software and Exadata Cloud architecture powering extreme performance for AI, analytics, and mission-critical workloads at any scale. Exascale’s capability of intelligent snapshots and thin provisioning of databases dramatically reduces storage requirements in a dense consolidation environment.

Oracle’s Maximum Availability Architecture (MAA) was outlined in this document, including the application of MAA in the context of database consolidation. Providing the required level of database and application availability is a key goal for consolidation environments, and this document outlines how to address those business needs.



Security is improved through database consolidation by the reduction in complexity of multiple systems and often divergent platforms with varying management practices. Security is enhanced even further in deployments on the Exadata platform, which is delivered in a “secure by default” configuration, includes a pre-scanned system stack, and the Advanced Intrusion Detection Environment (AIDE) implementation to ensure against attack.

Finally, any consolidation effort must deliver the proper level of performance required by business applications. Systems need to be architected to achieve consolidation without sacrificing performance, and Exadata has demonstrated performance capabilities that can be leveraged to provide consolidation as well as the performance necessary.

## References

My Oracle Support Documents

Doc ID	Legacy ID	Document Title
<a href="#">KB137827</a>	1338988.1	Scripts and Tips for Monitoring CPU Resource Manager
TBD	1337265.1	Tool for Gathering I/O Resource Manager Metrics: metric_iorm.pl
<a href="#">KB153930</a>	888828.1	Exadata Database Machine and Exadata Storage Server Supported Versions
<a href="#">KB150354</a>	361468.1	HugePages on Oracle Linux 64-bit
<a href="#">KB151310</a>	401749.1	Oracle Linux: Shell Script to Calculate Values Recommended Linux HugePages / HugeTLB Configuration
<a href="#">KB174916</a>	3081878.1	Starting DBRU 19.27 and 23.8 Small Pages Are Not Allowed for RDBMS SGA In Exadata
<a href="#">KB165271</a>	1484302.1	Master Note: Overview of Oracle Resource Manager and DBMS_RESOURCE_MANAGER
<a href="#">KB151455</a>	1585184.1	Using PROCESSOR_GROUP_NAME to bind a database instance to CPUs or NUMA nodes on Linux
<a href="#">KB137827</a>	1338988.1	Scripts and Tips for Monitoring CPU Resource Manager
<a href="#">KB74841</a>	1070954.1	Oracle Exadata Database Machine exachk or HealthCheck

## White Papers & Presentations

### Exadata Security Guide

<https://docs.oracle.com/en/engineered-systems/exadata-database-machine/dbmsq/index.html#Oracle%C2%AE-Exadata-Database-Machine>

### MAA best practices

[Oracle Maximum Availability Architecture Best Practices](#)

### Virtual Machines with Exadata

<https://www.oracle.com/a/tech/docs/exadata-kvm-overview.pdf>

## Continuous Availability White Paper

<https://www.oracle.com/docs/tech/database/continuous-service-for-apps-on-atpd.pdf>

## Advantages of Exascale

<https://blogs.oracle.com/exadata/post/advantages-of-exadata-exascale>

## CONNECT WITH US

Call +1.800.ORACLE1 or visit [oracle.com](https://www.oracle.com).

Outside North America, find your local office at [oracle.com/contact](https://www.oracle.com/contact).

 [blogs.oracle.com](https://blogs.oracle.com)

 [facebook.com/oracle](https://facebook.com/oracle)

 [twitter.com/oracle](https://twitter.com/oracle)

Copyright © 2026, 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.

Intel and Intel Xeon are trademarks or registered trademarks of Intel Corporation. All SPARC trademarks are used under license and are trademarks or registered trademarks of SPARC International, Inc. AMD, Opteron, the AMD logo, and the AMD Opteron logo are trademarks or registered trademarks of Advanced Micro Devices. UNIX is a registered trademark of The Open Group. 0120

Best Practices for Oracle Database Consolidation  
August 2025

Author: Christian A. Craft, Maruti Sharma

Contributing Authors: Glen Hawkins, Bob Thome, Lawrence To, Tina Rose, Dan Norris, Kevin Diehl

