ORACLE

Maximizing Availability with Oracle Database

May 2024, Version 6.0

Public

Purpose statement

This document provides an overview of the high availability and disaster recovery features of the Oracle Database in the context of Oracle's Maximum Availability Architecture reference tiers. It is intended solely to help assess the business and technical benefits of adapting and configuring applications and databases to best meet Recovery Time Objective (RTO) and Recovery Point Objective (RPO) goals via high availability and data protection solutions and best practices.

The intended audience is anyone responsible for the maintenance and lifecycle of applications (ranging from critical to the development and test systems) that utilize the Oracle Database as part of the architecture. While database administration knowledge is helpful in understanding some of the more profound concepts, most of this document can be read by anyone with an understanding of essential software and database operations as well as high availability and disaster recovery architecture.

Disclaimer

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, timing, and pricing 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.

ORACLE

Table of Contents

Introduction **5**

The High Availability Challenge **5**

Types of Downtime and Risks **5**

Data Availability and Corruption Protection 6

Human Error Protection 6

Protection from Physical Data Corruption 6

Considering Planned Downtime 6

Oracle Database High Availability **6**

Oracle's 3 High Availability Design Principles 7

Oracle Maximum Availability Architecture (MAA) **8**

Addressing both Planned and Unplanned Downtime **9**

The Bronze Tier: The Basics of a Single Instance HA Environment 9

Backup and Recovery – Oracle Recovery Manager 10

Fast Recovery Area 10

Real-Time Data Protection – Zero Data Loss Recovery Appliance 11

Oracle Database Zero Data Loss Autonomous Recovery Service 11

Recovery from Logical Corruption: Oracle Flashback Technology 12

Online Data Reorganization and Redefinition 13

The Silver Tier: Active/Active Database Clustering 13

Server: Oracle Real Application Clusters 14

Transparent Failover: Application Continuity 15

Storage: Automatic Storage Management 16

Oracle Globally Distributed Database 16

The Gold Tier: Physical Replication, Zero Data Loss, Fast Failovers 18

Real-time Data Protection and Availability – Oracle Data Guard 18

High Availability with Zero Data Loss Across Any Distance: Active Data Guard 19

Increase ROI by Offloading Workloads to an Active Data Guard Standby Database 20

The Platinum Tier: Highest Uptime for all Outages, Zero Data Loss 22

Active-Active HA: GoldenGate 22

Online Application Upgrades: Edition-Based Redefinition (EBR) 24

Managing Oracle Database High Availability Solutions **24**

Oracle Enterprise Manager 24

Full Stack Disaster Recovery 25

Global Data Services 25

Conclusion **26**

Appendix **27**

New High Availability Feature Highlights in Oracle Database 23ai 27

New High Availability Feature Highlights in Oracle Database 21c 28

High Availability Feature Highlights in Oracle Database 19c 29

Introduction

Enterprises use Information Technology (IT) to gain competitive advantages, reduce operating costs, enhance customer communication, and increase management insight into their business. Thus, enterprises increasingly depend on their IT infrastructure and its continuous availability. Application downtime and data inaccessibility directly translate into lost productivity and revenue, dissatisfied customers, and damage to corporate reputation.

A basic approach to building high availability infrastructures is to deploy redundant and often idle hardware and software resources supplied by disparate vendors. This approach is usually expensive yet falls short of service level expectations due to the lost integration of components, technological limitations, and administrative complexity. In contrast, Oracle provides customers with comprehensive and integrated high availability technologies to reduce cost, maximize return on investment through the productive use of all high availability resources, and improve the quality of service to users.

This paper examines the types of outages affecting IT infrastructures and presents Oracle Database technologies that comprehensively address them. These technologies, integrated into Oracle's Maximum Availability Architecture (MAA), reduce or avoid unplanned downtime, enable rapid failure recovery, and minimize planned downtimes.

This paper introduces the concepts of Oracle Maximum Availability Architecture, including the high availability features and disaster recovery features utilized to protect the Oracle Database from downtime and data loss in terms of performance, scalability, functionality, and ease-of-use - including Real Application Clusters (RAC), Automatic Storage Management, Globally Distributed Database, Recovery Manager, Data Guard and Active Data Guard, Oracle Golden Gate, and Edition-Based Redefinition (EBR).

The High Availability Challenge

Designing, implementing, and managing a high availability (HA) architecture that achieves all business objectives under real-world constraints is difficult. Many technologies and services from different suppliers offer to protect your business from data loss and downtime - who can you trust?

From Oracle's perspective, HA encompasses many vital aspects in addition to the primary goal of preventing downtime. Key dimensions of a comprehensive HA architecture include:

- **Data availability:** ensuring access to data to prevent business interruption.
- **Data protection:** preventing data loss that compromises the viability of the business.
- **Performance & scalability:** delivering adequate response time for efficient business operations.
- **Cost:** reducing deployment, management, and support costs to conserve corporate resources.
- **Risk:** consistently achieving required service levels over a long period as the business evolves with no costly surprises or disappointments.

A successful HA implementation begins with understanding the service levels the business requires along each of these dimensions. This guides essential decisions on technology and determines the appropriate level of investment in the future HA architecture.

Types of Downtime and Risks

When considering different HA solutions, one must understand the various risks and types of downtime that impact an application. These downtime and risk events routinely fall into two categories, planned downtime, and unplanned outages. Examples of planned downtime would routinely consist of patching, upgrading, application updates (i.e., new application version), or migrating to a new platform or hardware. Likewise, unplanned outages might consist of

server instance outages, site disasters (i.e., floods, long-term power outages, or fire), recovery from human error, or data corruption.

When most IT teams responsible for applications and their associated infrastructure consider these events, they consider options to reduce Recovery Point Objective (RPO) and Recovery Time Objective (RTO). RPO and RTO are considered two of the most important parameters when creating a disaster recovery and data protection plan. Proper planning involves careful consideration per application to determine acceptable thresholds for each, which will play a big part in choosing the correct HA architecture.

Data Availability and Corruption Protection

Data availability is about avoiding and mitigating data failures: business-critical data loss, damage, or corruption. Data failures are due to one or a combination of causes or events, such as storage subsystem failures, site failures, human errors, and data corruption. Their multifaceted events often make data failures challenging to identify and diagnose. Subsequent sections examine the HA technologies included in the Oracle Database that help diagnose, prevent, mitigate, and recover from data failure.

Human Error Protection

Human errors are a leading cause of downtime; hence, sound risk management must include measures to prevent and remediate human errors. For example, an incorrect WHERE clause may cause an UPDATE to affect more rows than intended. The Oracle Database provides powerful capabilities that help administrators prevent, diagnose and recover from such errors. It also includes features for end-users to recover from problems directly, speeding the recovery of lost and damaged data.

Protection from Physical Data Corruption

Physical data corruption is created by faults in any of the components of the Input/Output (I/O) stack. When Oracle issues a write, this database I/O operation is passed to the operating system's code. The write goes through the I/O stack: from the file system to the volume manager to the device driver to Host-Bus Adapter to the storage controller to the NVRAM cache, and finally to the disk drive where the data are written. Hardware failures or bugs in any of these components can result in invalid or corrupt data being written to disk. This corruption could damage internal Oracle control information or application/user data – either of which can be catastrophic to the functioning of the database.

Considering Planned Downtime

Planned downtime is typically scheduled to give administrators a window to perform system and/or application maintenance. During these maintenance windows, administrators take backups, repair or add hardware components, upgrade or patch software packages, and modify application components, including data, code, and database structures. Oracle has recognized the need to minimize or eliminate planned downtime while performing these system and maintenance activities. Oracle Database enables planned maintenance with minimal downtime for a production database using -of-place gold-image patching. Other options include rolling upgrades using a synchronized copy of the production database or bi-directional replication between two copies of the production database to migrate from one version to the next with zero downtime.

Oracle Database High Availability

Oracle has been working hard for decades, helping IT departments worldwide solve high availability (HA) challenges by designing and implementing comprehensive HA capabilities into the Oracle database. This innovation results in HA solutions that give actual competitive advantages to enterprises by helping them achieve their service level objectives in the most cost-effective manner.

Oracle Database's high availability capabilities address the full range of planned and unplanned outages. Oracle builds and delivers database-aware HA capabilities that are deeply integrated with core internal features of the database. This results in cost-effective solutions that reduce business risk and achieve exceptional data protection, availability, performance, and return on investment. Oracle Database High Availability capabilities are flexible, enabling you to choose the appropriate level of HA, and are adaptable to support your business objectives today, efficiently, and in the future.

Oracle's 3 High Availability Design Principles

1. *Leverage Oracle Database internals for maximum data protection. Knowledge and control of its internal algorithms and data structures, including database block structure and redo format, enables Oracle to build intelligent, unique-to-Oracle data protection solutions. For instance, because it can detect corruption in a database immediately, Oracle Data Guard prevents the propagation of physical corruption, logical intra-block corruption, and logical corruption caused by lost-writes. Active Data Guard goes a step further, automatically repairing physical on-disk corruption that can occur at either the primary or standby database transparent to the user. Similarly, Recovery Manager (RMAN) performs Oracle-aware physical and logical block validation, ensuring valid backups.*
2. *Deliver application-integrated high availability. Providing high availability and data protection using cold failover clusters and storage-centric mirroring solutions is inadequate for comprehensive protection and fast recovery. Oracle Real Application Clusters (Oracle RAC) enable a single Oracle Database to run on a cluster of database servers in an active-active configuration. Performance is easy to scale out through online provisioning of additional servers – users are active on all servers, and all servers share access to the same Oracle Database. High availability is maintained during unplanned outages and planned maintenance by transitioning users on the server that is out of service to other servers in the Oracle RAC cluster that continue to function. Outages ultimately impact the availability of an application, and unlike storage-centric solutions, Oracle High Availability technologies are designed to operate at the business object level – e.g., repairing tables or recovering specific transactions. Therefore, the most crucial recent enhancement to Oracle's High Availability solution set in this context is Application Continuity (AC). AC masks many outages from end users and applications by replaying the failed in-flight transactions after a server or site failover has occurred – transparent to the application.*
3. *Provide an integrated, automated, and open architecture with a high return on investment. HA features built into the Oracle Database require no separate integration or installs. Upgrades to new versions are greatly simplified, eliminating the painful and time-consuming process of release certification across multiple vendors' technologies. Also, all the features can be managed via the unified Oracle Enterprise Manager Cloud Control management interface. Oracle builds automation into every step, preventing common mistakes typical in manual configurations. For example, customers can easily choose to automatically failover to a standby database if the production database becomes offline, automatically remove and archive backups for effective space management; and automatically repair physical block corruptions. Oracle High Availability solutions are inherently active – avoiding idle components that only function when a failure occurs. All Oracle RAC nodes are active. Active Data Guard standby systems support read-only applications, data extracts, and fast incremental backups. There is never a question of whether it will start and how long it will take after a failure occurs to resume service: all Oracle High Availability components are already active, performing valuable work, and enabling continuous user validation, ensuring they are ready for prime-time.*

Oracle Maximum Availability Architecture (MAA)

Oracle Maximum Availability Architecture (MAA) is a set of best practices blueprints for the integrated use of Oracle High Availability (HA) technologies (see Figure 1).

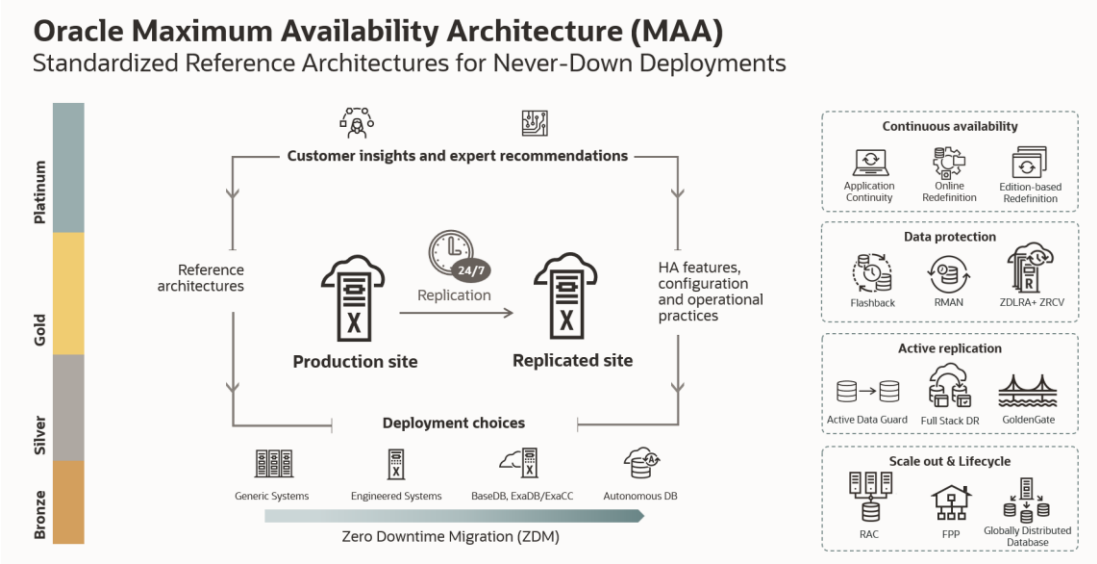

Figure 1: Oracle's High Availability Technologies and the Oracle Maximum Availability Architecture

For over a decade, MAA best practices have been created and maintained by a team of Oracle engineers that continually validate the integrated use of Oracle Database high availability features. Ongoing real-world customer experience is also constantly fed back into the validation process performed by the MAA team, spreading lessons learned to other customers and evolving these MAA blueprints to accommodate additional use cases.

MAA includes best practices for critical infrastructure components, including servers, storage, and network, combined with configuration and operational best practices for the Oracle High Availability capabilities deployed on it. MAA resources (oracle.com/goto/maa) are continually updated and extended.

Given that all applications do not have the same high availability and data protection requirements, MAA best practices describe standard architectures designed to achieve different service level objectives. Details are provided in *Oracle Maximum Availability Architecture Blueprints for reduced planned and unplanned downtime for the On-Premises, Exadata-based, or Cloud-base Oracle Database*.[1]

Over the years, Oracle MAA has evolved in multiple directions. For example, Oracle MAA on Engineered Systems now provides the MAA best practices and blueprint recommendations as part of those Engineered Systems, such as the Oracle Exadata Database Machine. For Oracle Database Services in the Oracle Cloud, Oracle MAA is not only integrated into the deployment. For example, the Oracle Cloud, especially the platform as a service offering (PaaS), is operated following those standards that have ensured maximum availability for many of Oracle's customers for decades.

1 https://www.oracle.com/a/tech/docs/maa-onpremises-overview.pdf

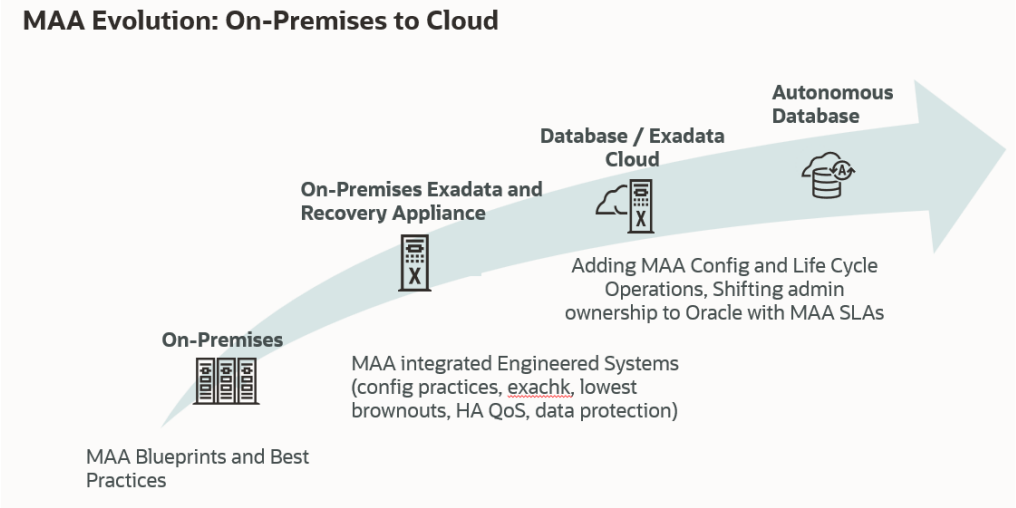

Figure 2: MAA Evolution from On-Premises into the Oracle Cloud

Finally, Oracle MAA has become the new de facto high availability standard. In the absence of any other comprehensive literature on this subject, Oracle MAA acts as general guidance for any database operation that would want to meet the highest level of availability, as MAA blueprints consider and discuss the various failure scenarios that can affect any database. For Oracle Databases, Oracle MAA goes a step further in providing a solution based on Oracle's integrated High Availability features which will be discussed in more detail in the remainder of this paper.

Thus, Oracle MAA does not only address Oracle customers that want to improve their database availability but also non-Oracle databases and especially future Oracle customers that would like to review failure scenarios and get an idea about what type of failures and planned maintenance operations need to be covered. In this context, Oracle MAA is also an exciting topic for application developers, as it provides guidance on which failures the application may have to tailor to and which failures an application can ignore, or even better, for which failures the application can rely on Application Continuity to keep them completely transparent.

ADDRESSING BOTH PLANNED AND UNPLANNED DOWNTIME

Hardware faults, which cause server failures, are essentially unpredictable and result in application downtime when they eventually occur. Likewise, a range of data availability failures, including storage corruption, data corruption, site outage, and human error, can often result in unplanned downtime disrupting productivity and the overall business. Last but not least, patching and other planned maintenance operations can severely impact the availability of the database if downtime is required, which sometimes can span a day or more.

The following sections have been organized by MAA tier. They are designed to provide an overview of how Oracle's High Availability features can help tackle any disruptive use cases discussed above, whether they fall into the planned maintenance or unplanned outage category. Oracle's MAA blueprints hierarchically build on one another, so everything in the bronze tier is carried over to the silver tier, which is likewise represented in the gold and platinum tiers. You will see this represented as we outline these HA solutions for each tier which maps to a specific set of RPO and RTO requirements mapping to the needs of a particular application along with the associated end-users and businesses that depend on it.

The Bronze Tier: The Basics of a Single Instance HA Environment

In some cases, such as development, test, or non-critical systems that may run a single instance of the Oracle Database, a certain amount of downtime for planned and unplanned outages may be acceptable as long as recovery

is possible within a specified time timeframe. For example, the acceptable recovery timeframe for an unplanned data corruption event may be a few hours or an hour of downtime in cases where a human mistake corrupts the underlying data as long as the ability to return data to a previous state within that time period is available. These HA requirements for the Bronze MAA tier are outlined in the table below:

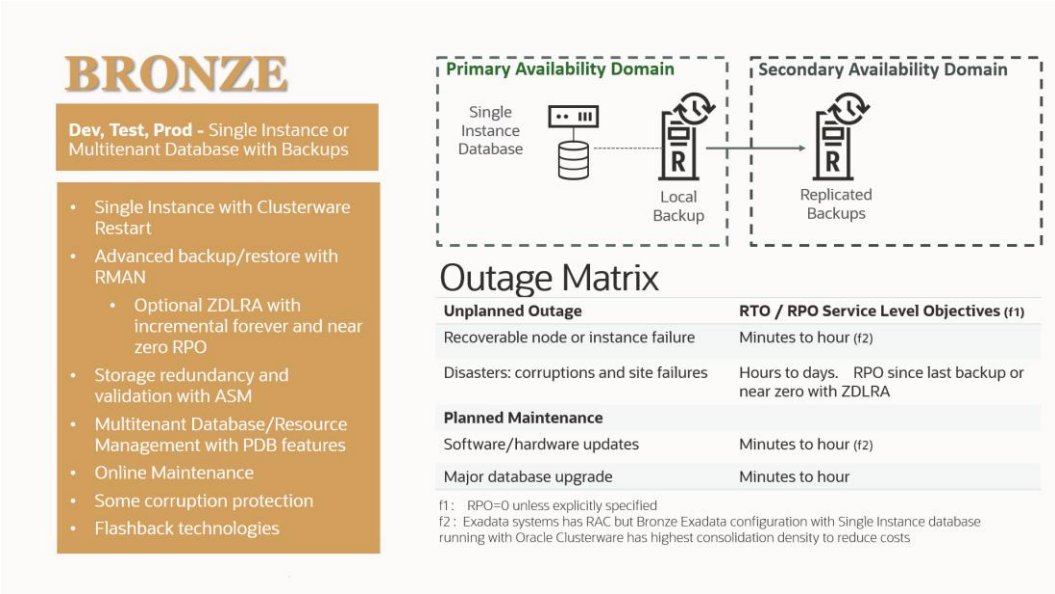

Figure 3: Bronze Tier RTO and RPO Levels of Protection

The HA technologies and solutions below represent Oracle's HA technologies that can maintain the RTO and RPO levels above when set up and configured as indicated in our tiers.

Backup and Recovery – Oracle Recovery Manager

In addition to prevention and recovery technologies, every IT organization must implement a complete data backup procedure to respond to multiple failure scenarios. Oracle provides best-of-breed, Oracle-aware tools to efficiently backup and restore data and to recover data up to the time just before a failure occurred. Oracle supports backups to disk, tape, and cloud storage. This wide range of backup options allows users to deploy the best solution for their environment. The following sections discuss Oracle's disk, tape, cloud backup technologies, and the Data Recovery Advisor.

Oracle Recovery Manager (RMAN)

Recovery Manager (RMAN) manages database backup, restore, and recovery processes. RMAN maintains configurable backup and recovery policies and keeps historical records of all database backup and recovery activities. Large databases can include hundreds of files, making backup challenging without an Oracle-aware solution. Missing even one critical file can render the entire database backup useless, and incomplete backups may go undetected until needed in an emergency. RMAN ensures that all files required to restore and recover a database successfully are included in database backups. During the backup and restore processes, RMAN validates all data to ensure that corrupt blocks are not propagated. If corrupt blocks are found during a restore operation, RMAN automatically relies on file(s) from a previous backup as necessary for a successful recovery.

For more details on Oracle's RMAN, refer to *oracle.com/goto/rman*.

Fast Recovery Area

A vital component of the Oracle Database backup strategy is the Fast Recovery Area (FRA), a location on a file system or ASM disk group for all recovery-related files and activities for an Oracle Database. All the files required to recover a

database from media failure can reside in the FRA, including control files, archived logs, data file copies, and RMAN backups. Oracle automatically manages space in the FRA. One or more databases may share a single FRA.

Real-Time Data Protection – Zero Data Loss Recovery Appliance

The Zero Data Loss Recovery Appliance (ZDLRA) is an innovative optional data protection solution integrated with RMAN and the Oracle Database. It eliminates data loss exposure and dramatically reduces data protection overhead on production servers across the enterprise. The Recovery Appliance efficiently protects all databases in the data center with a massively cloud-scale architecture, ensures end-to-end data validation, and fully automates the management of the entire data protection lifecycle for all Oracle Databases through the unified Enterprise Manager Cloud Control interface.

The Recovery Appliance is an integrated hardware and software appliance with substantial technical innovation that standardizes backup and recovery processes for Oracle Databases across the entire data center. The appliance offers the following unique advantages.

- It eliminates data loss by using proven Data Guard technology to transmit redo records, the fundamental unit of transactional changes within a database. Protected databases ship redo to the Recovery Appliance as soon as it is generated, eliminating the requirement to take archived log backups at a production database. The granularity and real-time nature of this unique level of protection allow databases to be protected up to the last sub-second of data.
- Minimal impact backups – The Recovery Appliance's Delta Push technology offloads backup operations from production databases using a proper incremental-forever backup strategy. After a full backup, protected databases send RMAN incremental backups to the Recovery Appliance. RMAN block change tracking sends deltas, resulting in effective source-side deduplication by only sending unique changes. Delta Push eliminates recurring full backups and reduces bandwidth utilization. In addition, all overhead from RMAN backup deletion/validation/maintenance operations and tape backups are offloaded to the Recovery Appliance.
- Any point-in-time restore using Delta Store technology. The Recovery Appliance validates, compresses, indexes, and stores the incoming deltas. The deltas are the foundation of virtual full database backups, essentially space-efficient pointer-based representations of full physical backups as of an incremental backup point-in-time. When the time comes for a restore operation, Delta Store efficiently recreates a full physical backup from the appropriate incremental backup point. Archived log backups stored by the appliance are then used to roll forward to the desired time. The Delta Store eliminates traditional restore's typical production server overhead and applies successive incremental backups. The scalability and performance of the underlying Exadata-based hardware architecture further optimize the performance of the restore operation.

The Recovery Appliance is the ideal solution for enterprise backup and any point-in-time recovery for Oracle Databases. It is also the ideal disaster recovery solution for Oracle Databases that support applications with recovery time objectives that a restore from backup can achieve. Oracle Data Guard and Active Data Guard, discussed in the following sections, are the solutions for applications with more aggressive recovery time objectives that can only be achieved by fast failover to a running copy of the production database

For more details on Zero Data Loss Recovery Appliance (ZLDRA), refer to *http://www.oracle.com/recoveryappliance*.

Oracle Database Zero Data Loss Autonomous Recovery Service

For those running their applications in Oracle Cloud Infrastructure (OCI), Oracle has introduced the Oracle Database Zero Data Loss Autonomous Recovery Service bring the powerful data protection capablitilies of ZDLRA to the cloud. Oracle Database Zero Data Loss Autonomous Recovery Service is a fully-managed data protection service for Oracle databases running on Oracle Cloud Infrastructure (OCI). Unique, automated capabilities protect the Oracle Database changes in real time, validate backups without production database overhead, and enable fast, predictable recovery to

any point in time. The service is directly integrated with various Oracle Database service offerings including Oracle Autonomous Database and Oracle Exadata Databsae Service.

For more details on Oracle Database Zero Data Loss Autonomous Recovery Service, refer to https://www.oracle.com/database/zero-data-loss-autonomous-recovery-service.

Recovery from Logical Corruption: Oracle Flashback Technology

Human errors happen. Oracle Database Flashback Technologies provide a unique and rich set of data recovery solutions that enable reversing human errors by selectively and efficiently undoing the effects of a mistake. Before Flashback, it might take minutes to damage a database but hours to recover. With Flashback, the time required to recover from an error depends on the work done since the error occurred. Recovery time does not rely on the database size, a capability unique to the Oracle Database that becomes necessary as database sizes continue to grow. Flashback supports recovery at all levels, including the row, transaction, table, and entire database.

Flashback is easy to use: the entire database can be recovered with a single short command instead of following a complex procedure. It also provides fine-grained analysis and repair for localized damage, e.g., when the wrong customer order is deleted. In addition, Flashback can repair more widespread damage while avoiding the need for the extended downtime, e.g., all of yesterday's customer orders have been deleted. The following sub-sections walk through some of the critical features of Flashback.

Flashback Database

To restore an entire database to a previous point-in-time, the traditional method is to restore the database from an RMAN backup and recover to the point-in-time before the error. This can take time proportional to the (ever-growing) size of the database resulting in hours or even days of recovery time using traditional methods.

In contrast, Flashback Database, using Oracle-optimized flashback logs, can quickly restore an entire database to a specific point in time without the need for prolonged recovery time. Flashback Database is fast because it restores changed blocks only. Flashback Database can restore a whole database in minutes via a simple command.

Flashback Table

When logical corruption is limited to one or a set of tables, Flashback Table allows the administrator to recover the affected tables to a specific point-in-time. A query such as:

FLASHBACK TABLE orders, order_items TIMESTAMP time

This will undo updates to the orders and order_items tables made after the specified time.

Flashback Query

Using Oracle Flashback Query, administrators can query any data at some point in time in the past. This powerful feature can view and logically reconstruct corrupted data that may have been deleted or changed inadvertently. For example, a simple query such as:

SELECT * FROM emp AS OF TIMESTAMP time WHERE…

This displays rows from the emp table as of the specified time (a timestamp obtained via a TO TIMESTAMP conversion). Administrators can use Flashback Query to identify and resolve logical data corruption. This functionality can also be built into an application to provide its users with a quick and easy mechanism to undo erroneous changes to data without contacting their database administrator.

Flashback Transaction

Often, data failures take time to be identified. Additional 'good' transactions may have been executed on data logically corrupted by an earlier 'bad' transaction. In this situation, the administrator must analyze changes made by the 'bad' transaction and any other (dependent) transactions that subsequently modify the same data to ensure that undoing the 'bad' transaction preserves the original, correct state of the data. This analysis can be laborious, especially for complex applications.

Flashback Transaction enables an administrator to Flashback a single 'bad' transaction and, optionally, all dependent transactions with a single PL/SQL operation. Alternatively, an administrator can use Oracle Enterprise Manager Cloud Control to identify and Flashback the necessary transactions.

Flashback Time Travel

Flashback Time Travel automatically tracks every transactional change made to database data and maintains a secure, and easily accessible, archive of historic data. It maintains a history of the evolution of the tables' schema. By maintaining the history of the transactional changes to a table, and its schema, Flashback TimeTravel enables you to perform operations such as Flashback Queries (AS OF and VERSIONS) onthe table to view the history of the changes made during the transaction time.

As Flashback Time Travel tracks and stores transactional changes to a table, it is useful forcompliance with record stage policies and audit reports. With the release of Oracle Database 23ai, numerous enhancements were made to improve the performance of Flashback Time Travel.

For more details on Flashback, refer to *oracle.com/goto/flashback*.

Online Data Reorganization and Redefinition

Online data and schema reorganization improves overall database availability and reduces planned downtime by allowing users full access to the database throughout the reorganization process. For example, adding columns with a default value does not affect database availability or performance. Many data definition language (DDL) maintenance operations allow administrators to specify timeouts on lock waits to maintain a highly available environment while performing maintenance operations and schema upgrades. Also, indexes can be created with the INVISIBLE attribute so the Cost-Based Optimizer (CBO) ignores them, although DML operations maintain them. Once an index is ready for production, a simple ALTER INDEX statement will make it visible to the CBO.

Online Table Redefinition

As business requirements evolve, the applications and databases supporting the business go through a similar evolution process. Through the strategic use of the DBMS_REDEFINITION package (also available in Oracle Enterprise Manager) – administrators can reduce downtime in database maintenance by allowing changes to a table structure while continuing to support an online production system. Administrators using this API enable end users to access the original table, including insert/update/delete operations, while the maintenance process modifies an interim copy of the table. The interim table is routinely synchronized with the original table, and once the maintenance procedures are complete, the administrator performs the final synchronization and activates the newly structured table.

For more details on Online Data Reorganization and Redefinition, refer to
https://www.oracle.com/database/technologies/high-availability/online-ops.html

The Silver Tier: Active/Active Database Clustering

As applications become more critical to the overall business, a high availability solution that provides near-instantaneous recovery from unexpected outages becomes necessary. This becomes true for almost any application that provides an internal critical function or interfaces with external customers and partners. With both RPO and RTO decreasing, HA architecture must ensure that databases can handle underlying infrastructure failures and any unexpected disruptions to the database instance itself while reducing the data loss potential. Routinely, there is also a desire to ensure all infrastructure is utilized. Thus there is a desire to ensure redundant instances are active and are being used to handle workload providing additional performance and scalability benefits.

With those needs in mind, Oracle has provided the silver MAA tier, which expands the Bronze MAA tier capabilities by adding active-active clustering, automatic storage management (ASM), and Application Continuity. These technologies continue to evolve, but the RPO and RTO thresholds can be seen below in Figure 4: Silver Tier RTO and RPO Levels of Protection. In addition, it provides an alternative architecture for maximizing fault tolerance via the

horizontal partitioning that Oracle Globally Distributed Database provides as an option in the Oracle Database architecture.

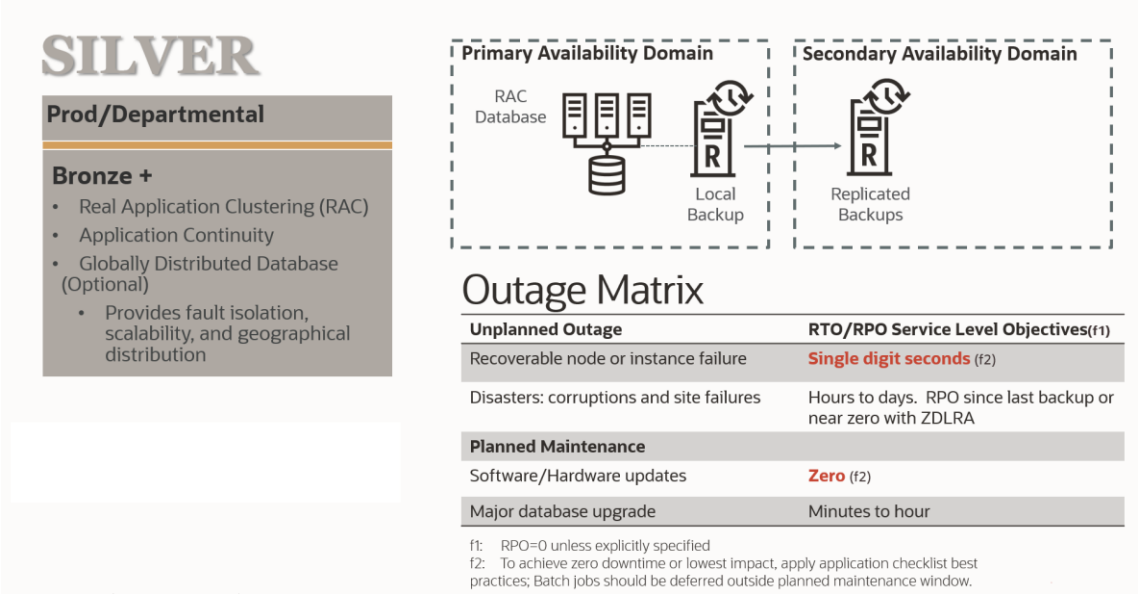

Figure 4: Silver Tier RTO and RPO Levels of Protection

Server: Oracle Real Application Clusters

Server availability is related to ensuring uninterrupted access to database services despite the unexpected failure of one or more machines hosting the database server, which could happen due to hardware or software faults. Oracle Real Application Clusters (RAC) can effectively protect against such failures.

Oracle Real Application Clusters (RAC) is Oracle's premier shared everything database clustering technology. Oracle Database with the Oracle RAC option enables multiple database instances to run on different servers in the cluster against a shared set of data files that comprise a database. The database spans multiple hardware systems yet appears as a single unified database to the application.

The Oracle RAC architecture extends availability and scalability benefits to all applications, specifically:

- Fault tolerance within the server pool, especially for compute failures. Since the nodes run independently, losing one or more does not impact other nodes. This architecture also allows a group of nodes to transparently be put online or taken offline while the rest of the system continues to provide database services.

- Flexibility and cost-effectiveness, to the degree that a system can scale to any desired capacity as business needs change. Oracle RAC allows users to add nodes to the system as capacity needs increase while reducing costs by avoiding the more expensive and disruptive upgrade path of replacing an existing monolithic system with a larger one.

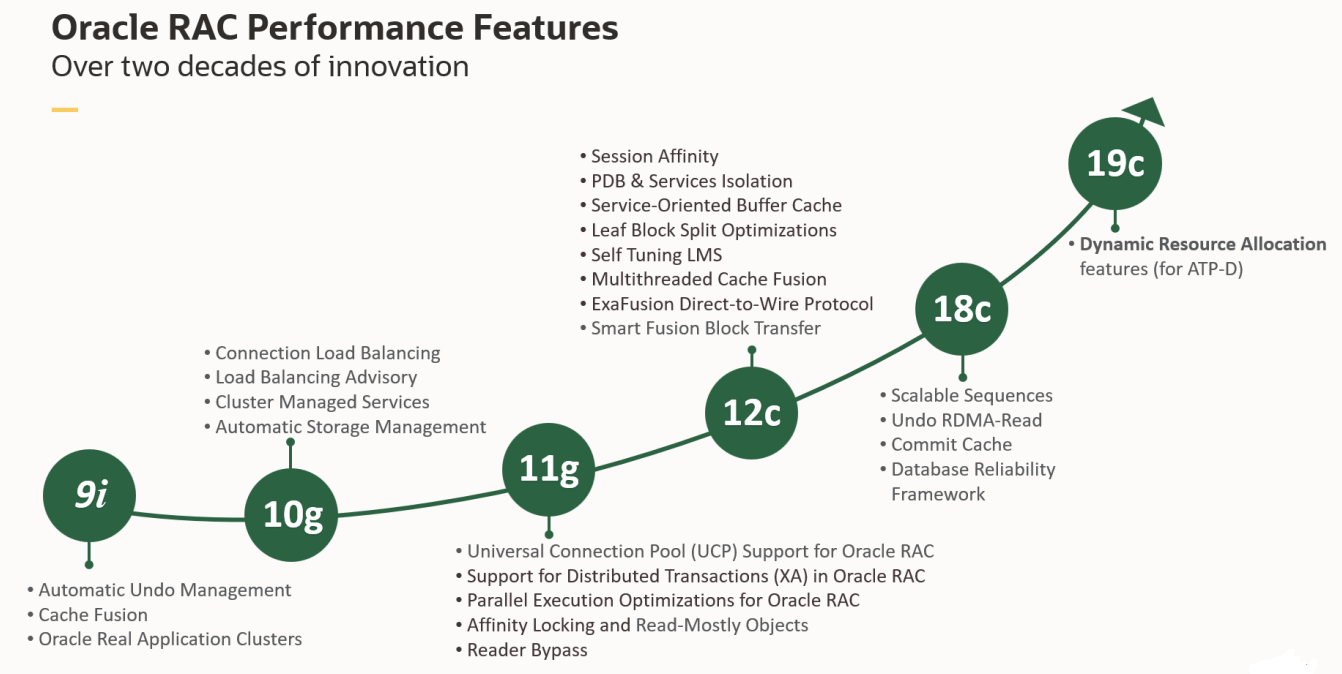

Figure 5: Oracle RAC Performance Innovations

Rolling patch Upgrades using Oracle RAC

Oracle supports the application of patches to the nodes of a Real Application Cluster (RAC) system in a rolling fashion, maintaining the database available throughout the patching process. To perform the rolling upgrade, one of the instances is quiesced and patched while the other instance(s) in the server pool continue in service. This process repeats until all instances are patched. The rolling upgrade method can be used for Patch Set Updates (PSUs), Critical Patch Updates (CPUs), one-off database, and diagnostic patches using OPATCH, operating system upgrades, and hardware upgrades.

For more details on Oracle Real Application Clusters (RAC), refer to: http://oracle.com/goto/rac.

Fleet Patching and Provisioning (FPP) for Oracle RAC

Fleet Patching and Provisioning (FPP) is a feature of Oracle Grid Infrastructure (GI) that significantly simplifies provisioning, patching, and upgrading RAC and single instance databases across large-scale deployments. It provides standardization by utilizing a gold image for out-of-place patching across a fleet of databases. The FPP Server's components are managed automatically by GI, simplifying the configuration. FPP maintains a space-efficient repository of gold images providing standardized software homes that can be provisioned to any number of target machines. Any number of homes can be provisioned from a given gold image, and FPP maintains lineage information so that the provenance of deployed software is always known. Gold images can be organized into series, allowing you to create groupings that track the evolution of a release. In addition, a notification system informs interested parties when a new image is available in a given series.

For more details on Fleet Patching & Provisioning, refer to
https://www.oracle.com/database/technologies/rac/fpp.html.

Transparent Failover: Application Continuity

It is complex for application developers to mask database session outages; as a result, errors and timeouts are often exposed to end users leading to frustration and lost productivity. Oracle Database includes Application Continuity

(AC). This capability intends to mask database outages from the application by catching failed transactions (in-flight or DML transactions including), reconnecting the application to another node in an Oracle RAC cluster or via Oracle Active Data Guard, and replaying the failed transaction. Application Continuity performs these steps beneath the application so that the outage simply appears in the application as a slightly delayed execution.

Application Continuity supports OCI, ODP.NET unmanaged, JDBC Thin on XA, Tuxedo, and SQL*Plus clients. By supporting the relocation or stopping of database services, Application Continuity made it easy to migrate existing connections to another database instance even if Oracle Connection Pools were not used.

With the Oracle Database 19c release, *Transparent Application Continuity (TAC)* was introduced, which tracks and records sessions and transactional states with full transparency. At the same time, the core Application Continuity framework has been enhanced to further assist with the outages that come as a side-effect of planned maintenance operations. Therefore, AC (with TAC) now drains sessions during planned maintenance so that the server that hosts applications can shut down for maintenance purposes in the least disruptive manner, making it an ideal fully integrated solution to ensure end-users of applications are not impacted by both planned maintenance and unexpected outage events. Oracle Database 23ai introduced additional enhancements to Application Continuity including support for Oracle Active Data Guard *DBMS_ROLLING* (covered under Active Data Guard in subsequent sections) allowing admiistrators to fully benefit from the session draining capabilities of Application Continuity when doing rolling upgrades.

Storage: Automatic Storage Management

Automatic Storage Management (ASM) is the Oracle Database's purpose-built file system and volume manager. For Oracle databases, ASM simplifies both the file system and volume management. In addition to simplifying storage management, ASM improves file system scalability, performance, and database availability. These benefits hold for single-instance and Oracle Real Application Cluster (RAC) databases.

ASM is designed to maximize database availability with minimal need for manual configuration. For example, ASM provides automatic mirror reconstruction, resynchronization (self-healing), and rolling upgrades. ASM also supports dynamic and online storage reconfiguration. Customers realize significant cost savings and achieve lower total cost of ownership because of features such as just-in-time provisioning and clustered pool of storage, which make ASM ideal for database consolidation without additional licensing fees.

The concept of "database-oriented storage management" comes with a new type of Disk Group called Flex Disk Groups. With Flex Disk Groups, all files belonging to an individual database or a PDB, in the case of multitenancy, are collectively identified with a new ASM object called a File Group. A File Group logically contains the files associated with a single database. It hence simplifies operations that address all files of a given database at once, as command syntax referring to a File Group refers to all the files belonging to the File Group.

To improve database availability, ASM also supports Extended Disk Groups, which build the foundation for the Oracle Extended Clusters architecture that can now be found on Exadata Database Machines. This extends an Oracle RAC cluster's availability beyond a single data center by deploying RAC clusters across two closely located data centers. The design uses ASM mirroring across the data centers so that availability is in case of a complete failure of one or more data centers within close proximity.

For more details on ASM, refer to *oracle.com/goto/asm*.

Oracle Globally Distributed Database

Oracle Globally Distributed Database is a scalability, availability fault isolation, and geo-distribution feature for OLTP applications that distributes and replicates data across a pool of discrete Oracle databases. Each database in the elastic pool is referred to as a shard. Globally Distributed Database is built on a shared-nothing horizontal partitioning

architecture in which the databases do not share storage or rely on cluster software. Oracle Globally Distributed Database provides some benefits for web-scale applications:

- **Linear scalability**. OLTP applications designed for Globally Distributed Database can elastically scale (data, transactions, and users) to any level, on any platform, simply by deploying new shards on additional standalone servers. The performance scales linearly as shards are added to the pool because each shard is entirely independent of other shards.
- **Extreme Data Availability**. Globally Distributed Database eliminates a single point of failure (shared disk, SAN, clustering, etc.) and provides strong fault isolation. The unavailability or slowdown of a shard due to an unplanned outage or planned maintenance affects only the users of that shard; it does not affect the availability or performance of the application for users of other shards. Each shard may run a different release of the Oracle Database as long as the application is backward compatible with the oldest running version – making it simple to maintain the availability of an application while performing database maintenance.
- **Data Sovereignty and Data Proximity via Geographic Data Distribution**. Globally Distributed Database makes it possible to locate different parts of the data in different countries or regions – thus satisfying regulatory requirements where data has to be located in a specific jurisdiction. It also supports storing particular data closer to its consumers.

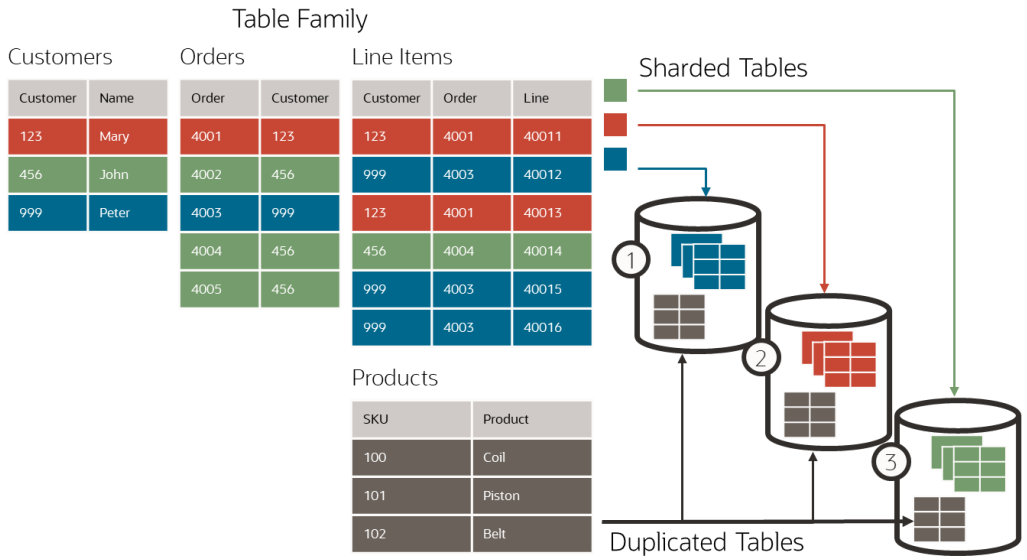

Figure 6: Oracle Globally Distributed Database at a Glance

Sharded Database and Shards

Shards are independent Oracle Databases hosted on database servers with their local resources - CPU, memory, and disk. No shared storage is required across the shards. A sharded database is a collection of shards forming one logical database. Shards can all be placed in one region (data center [s]) or different regions. A region in the context of Globally Distributed Database represents a data center or multiple data centers in close network proximity.

Shards are replicated for High Availability (HA) and Disaster Recovery (DR) with Oracle replication technologies such as Active Data Guard. For HA, the standby shards can be placed in the same region where the primary shards are placed. For DR, the standby shards are located in another region. Oracle Globally Distributed Database supports three automatically configured replication options: Data Guard, Active Data Guard, or Oracle GoldenGate.

With the release of Oracle Database 21c and 23ai, Oracle Globally Distributed Database has been enhanced significantly including support for a Sharding Advisor to help transition databases into a sharded schema and RAFT replication which provides built-in replication with support for active-active external consistency and data protection within a Sharding Group.

For more details on Oracle Globally Distributed Database, refer to: www.oracle.com/database/distributed-database/

The Gold Tier: Physical Replication, Zero Data Loss, Fast Failovers

While RTO requirements are routinely optimal for most with the introduction of RAC, recovering from data corruption is still required for many critical applications central to business functions. In addition, if there is a requirement for a remote site data center to protect from more significant site disasters such as floods, power outages, fire, or other natural disasters, a solution will be required to keep those sites synchronized to ensure recovery can be addressed in seconds preventing an outage from the application end-user perspective even in the case of these more significant outage events. There is often a need for near-zero downtime with planned maintenance activities such as migrating to a new platform, whether new hardware or an entirely new deployment platform such as the Oracle Cloud. These requirements are addressed in the Oracle MAA Gold Tier; please see Figure 7 below for this tier providing details on utilizing Active Data Guard in critical HA Architectures across multiple data centers potentially spanning long distances.

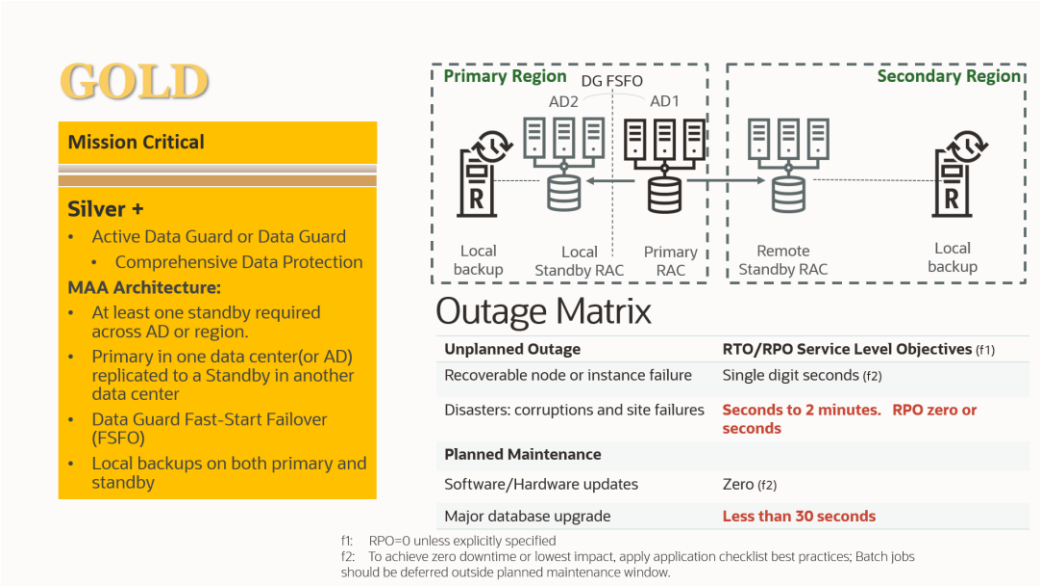

Figure 7: Gold Tier RTO and RPO Levels of Protection

Real-time Data Protection and Availability – Oracle Data Guard

Enterprises must protect their critical data and applications against events taking an entire cluster or data center offline. Human error, data corruption, or storage failures can make a cluster unavailable. Natural disasters, power outages, and communications outages can affect the availability of an entire site.

The Oracle Database offers a variety of data protection solutions that can safeguard an enterprise from costly downtime due to cluster or site failures. Frequently updated and validated local and remote backups constitute the foundation of an overall High Availability strategy. However, the complete restoration of a multi-terabyte backup can

take longer than the enterprise can afford to wait, and the backups may not contain the most up-to-date versions of data.

For these reasons, enterprises often maintain one or more synchronized replicas of the production database in separate data centers. Oracle provides several solutions that can be used for this purpose. Oracle Data Guard and Active Data Guard are optimized to protect Oracle data, providing high availability and disaster recovery.

Data Guard is a comprehensive solution to eliminate single points of failure for mission-critical Oracle Databases. It prevents data loss and downtime simply and economically by maintaining one or more synchronized physical replicas (standbys) of a production database (primary). Administrators can choose manual or automatic failover to these standby databases if the primary database is unavailable. Client connections can quickly and automatically failover to the standby and resume service.

Data Guard achieves the highest level of data protection through its deep Oracle Database integration, strong fault isolation, and Oracle-aware data validation. System and software defects, data corruption, and administrator errors that affect a primary database are not mirrored on the standby.

Last but not least, Data Guard provides a choice of either asynchronous (near zero data loss) or synchronous (zero data loss) protection. Asynchronous configurations are simple to deploy, with no performance impact on the primary, regardless of the distance that separates primary and standby databases. Synchronous transport, however, will affect performance and thus imposes a practical limit to the distance between the primary and standby database. Performance is affected because the primary database does not proceed with the next transaction until the standby acknowledges that changes for the current transaction are protected. Waiting for acknowledgment increases as the distance between primary and standby increases, directly affecting application response time and throughput. As described in subsequent sections, those effects can be mitigated using Oracle Fast or Far Sync.

High Availability with Zero Data Loss Across Any Distance: Active Data Guard

Active Data Guard represents a superset of the Data Guard functionality that includes advanced data protection, high availability capabilities, and features that increase return on investment (ROI) in disaster recovery systems. Several key capabilities of Oracle Active Data Guard are described below.

Active Data Guard Automatic Block Repair

Block-level data loss usually results from intermittent I/O errors and memory corruptions written to disk. When Oracle Database reads a block and detects corruption, it marks the block as corrupt and reports the error to the application. No subsequent read of the block will be successful until the block is recovered manually unless you are using Active Data Guard.

With Active Data Guard, block media recovery happens automatically and transparently. Active Data Guard repairs physical corruption on a primary database using a good version of the block retrieved from the standby. Conversely, corrupt blocks detected on the standby database are automatically repaired using the good version from the primary database.

Active Data Guard Far Sync: Zero Data Loss at any Distance

Active Data Guard Far Sync provides zero data loss protection for a production database by maintaining a synchronized standby database located at any distance from the primary location, without impacting database performance and with minimal cost or complexity.

A far sync instance (a new type of Data Guard destination) receives changes synchronously from a primary database and forwards them asynchronously to a remote standby (see Figure 8: Active Data Guard Far Sync – Zero Data Loss Protection at any Distance below) so that production can occur as quickly as needed, whether that be manual or automatic to the remote standby database with zero data loss.

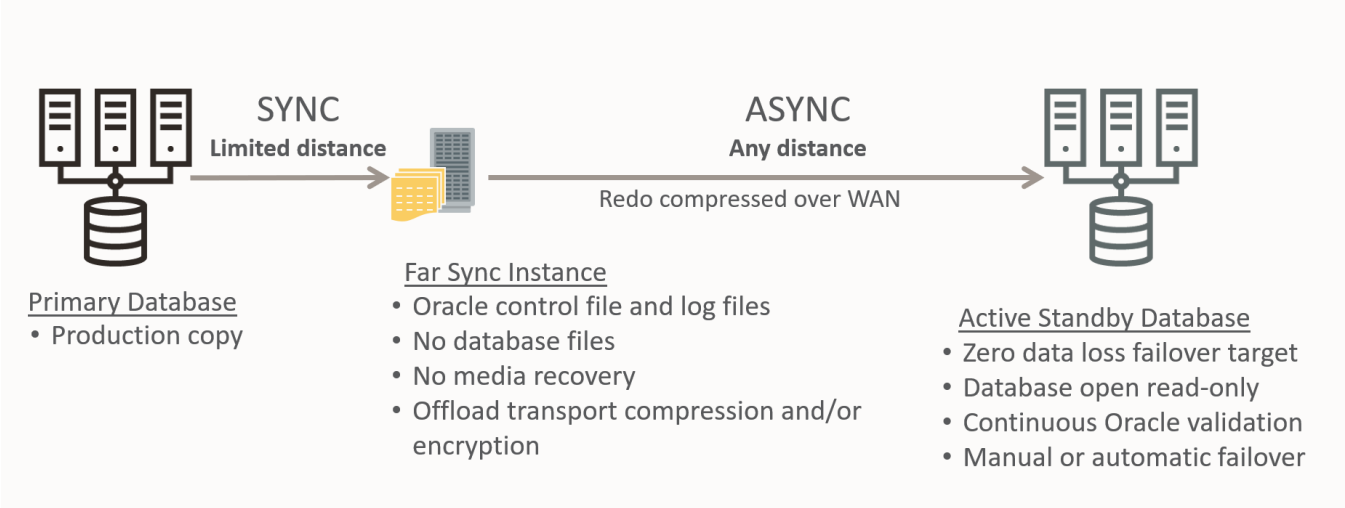

Figure 8: Active Data Guard Far Sync – Zero Data Loss Protection at any Distance

A far sync instance is a lightweight entity that manages a control file and log files. It requires a fraction of a standby database's CPU, memory, and I/O resources. It does not keep user data files, nor does it run recovery. Its purpose is to transparently relieve a primary database from serving remote destinations. A far sync instance can save network bandwidth by performing transport compression using the Oracle Advanced Compression option.

For example, consider an asynchronous Active Data Guard configuration with a primary in New York and a standby in London. Upgrade to zero data loss simply by using Active Data Guard to deploy a far sync instance within synchronous replication distance of New York (less than 150 miles). There is no disruption to the existing environment or requirement for proprietary storage, specialized networking, more database licenses, or complex management.

Increase ROI by Offloading Workloads to an Active Data Guard Standby Database

Active Data Guard enables the offloading of read-only and read-mostly reporting applications, ad-hoc queries, data extracts, and so on to an up-to-date physical standby database while providing disaster protection. Active Data Guard relies on a unique, highly concurrent apply process for the best performance while enforcing the same read consistency model for read-mostly access on the standby as executed on the primary database. No other physical or logical replication solution provides this capability. This makes it attractive to offload read-mostly workloads to an active standby, eliminating the cost of idle redundancy.

Active Data Guard (19c and beyond) allows for DML operations on the read-only standby to be redirected to the primary database, enabling even more reporting applications (even those requiring occasional writes) to use an Active Data Guard standby database.

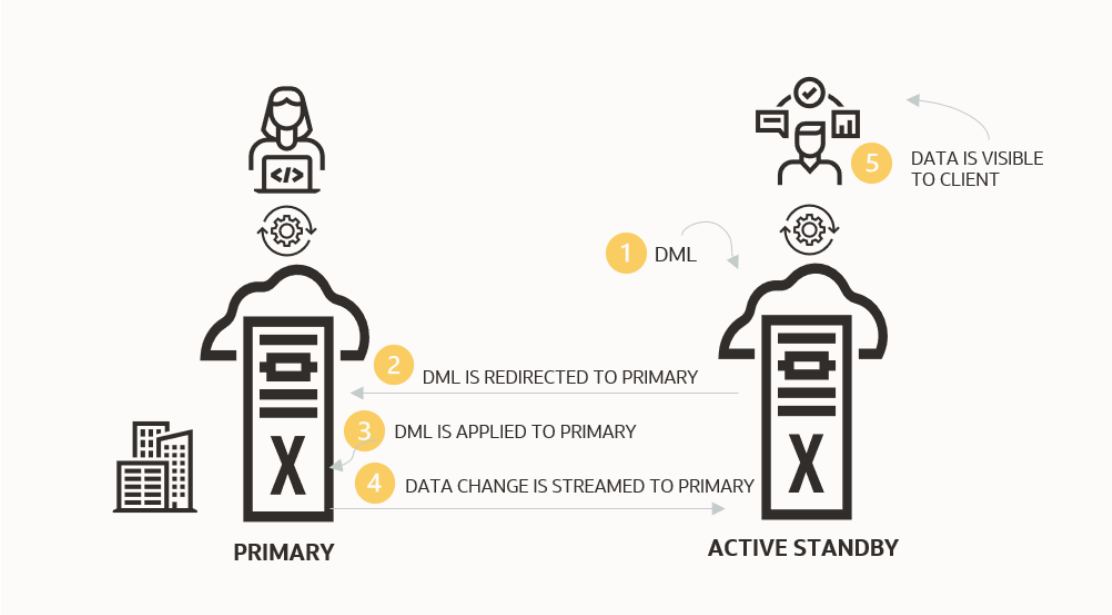

Figure 9: DML Redirect with Oracle Active Data Guard

In this context, it might also be worthwhile mentioning that the In-Memory data store can also be enabled on the standby database to improve the performance of these reports while Multi-Instance redo apply (MIRA) is enabled.

Database Rolling Upgrades using Active Data Guard

Although the database rolling upgrade process described above effectively reduces planned downtime, it is a manual procedure with many steps and is thus error-prone. This creates a reluctance to use the rolling upgrade process, making users accept longer downtimes associated with traditional upgrade methods. Traditional upgrade methods also increase risk because maintenance is performed on the production database BEFORE being sure of the outcome is possible.

Database Rolling Upgrades (DBMS_ROLLING) using Active Data Guard solves this problem by replacing forty-plus manual steps required to perform a rolling database upgrade with three PL/SQL packages that automate much of the process. This automation helps minimize planned downtime and reduce risk by implementing and thoroughly validating all changes on a complete production replica before moving users to the new version.

You can use this capability for database version upgrades starting from the first patchset of Oracle Database 12c. It should be noted that Oracle Database 23ai now includes Application Continuity support for Active Data Guard Database Rolling Upgrades with DBMS_ROLLING.

Data Guard Multitenant Support

Oracle Multitenant enables an Oracle Database to function as a container database (CDB). A CDB consolidates multiple pluggable databases (PDB), a portable collection of schemas, schema objects, and non-schema objects. With Data Guard, administrators have been able to configure disaster recovery at the CDB-level meaning that all PDBs would be protected in a single CDB configuration. This provides a significant level of protection with minimal configuration as hundreds of PDBs could be protected by a single Data Guard configuration. This also means that CDB/PDB placement should be carefully considered as switchover and failovers would be applicable to all PDBs within a specific CDB for role transitions and it is not possible to failover or switchover a single PDB with a CDB-level configuration.

With Oracle Database 21c, Oracle introduced Oracle Data Guard support for PDB-level configurations. This means that each PDB in a CDB may be a primary or standby depending on the configuration for that specific PDB Data Guard configuration. This maximizes flexibility to switchover or failover a single PDB if the PDB shows health or performance issues for example while not impacting the other PDBs running in the CDB. It should be noted that there are still many limitations related to

Active Data Guard features (none are supported at this time) or protection modes (only maximum performance mode is currently supported) while using a PDB-level configuration. With that in mind, administrators deciding between the two Data Guard Multitenant configurations should carefully consider their RTO and RPO requirements when weighing which option to go with. By default, Data Guard utilizes the CDB-level configuration.

Platform Migration, Systems Maintenance, and Data Center Moves

Data Guard also offers some flexibility for primary and standby databases to run on systems having different operating systems or hardware architectures, providing a straightforward method for platform migration with minimal downtime.[2] Data Guard can also be used to easily migrate to ASM and/or to move from single instance Oracle Databases to Oracle RAC, as well as for data center moves, with minimal downtime and risk.

For more details on Data Guard and Active Data Guard, refer to *http://www.oracle.com/goto/dataguard*

The Platinum Tier: Highest Uptime for all Outages, Zero Data Loss

The Platinum MAA tier provides reference blueprints that utilize Oracle's top level of high availability features to reduce the RTO time for database upgrades, patch sets, and even application upgrades to zero by introducing full active-active replication with Oracle GoldenGate. In addition, it provides the alternative option of utilizing Edition-based Redefinition to seamlessly upgrade your application when schema and other changes are required to the underlying database, as is often necessary for significant application upgrades. The sections below run through the full breadth of the Platinum MAA tier reference solution in more detail.

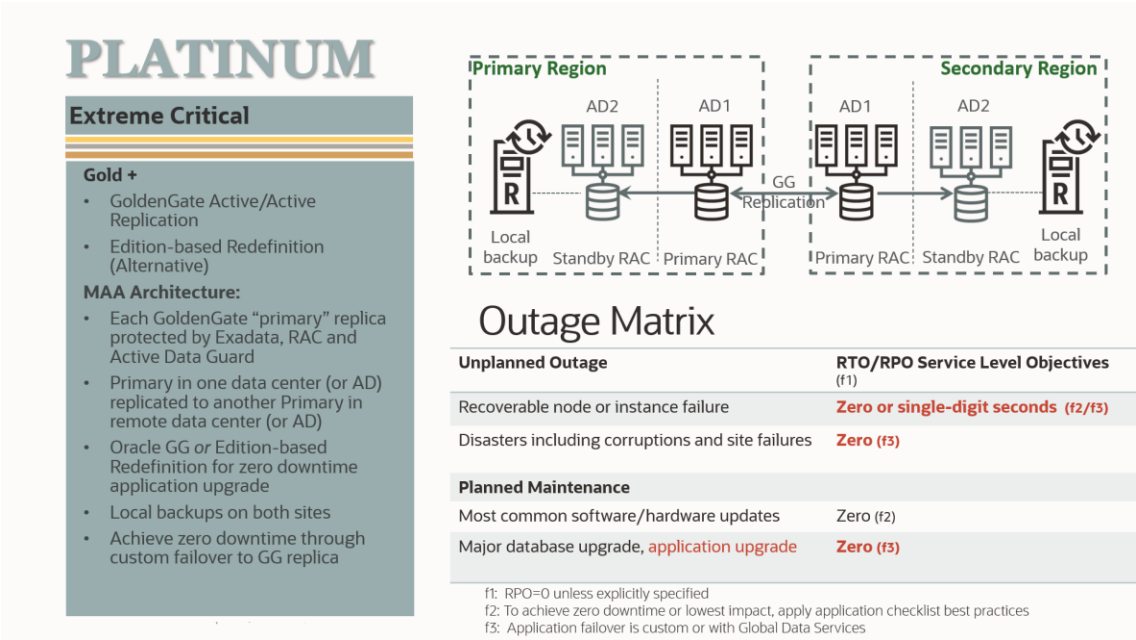

Figure 10: Platinum Tier RTO and RPO Levels of Protection

Active-Active HA: GoldenGate

Data Guard physical replication is optimized for a specific purpose – simple, transparent, one-way physical replication for optimal data protection and availability with specialized protection for data corruption with its bidirectional auto-

2 See MOS Note 413484.1 for details on platform combinations supported in a Data Guard configuration.

repair capability. Oracle GoldenGate, in contrast, is a feature-rich logical replication product with advanced features that can supplement Active Data Guard to support multi-master replication, hub and spoke deployment, subset replication, and data transformation, providing customers with flexible options to address their replication requirements fully. GoldenGate also supports replication between various heterogeneous hardware platforms and database management systems beyond Oracle.

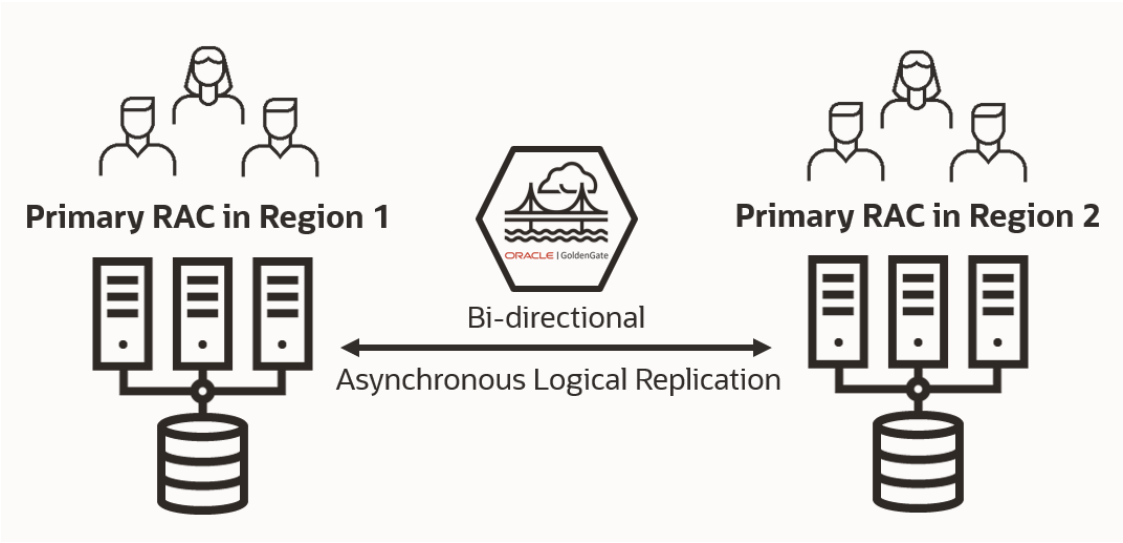

Figure 11: Oracle GoldenGate – Active-Active Bi-Directional Replication

Applications can use GoldenGate with minimal modification or special handling. GoldenGate can be configured, for example, to capture changes for an entire database, a set of schemas, or individual tables. Databases using Oracle GoldenGate technology can be heterogeneous – e.g., a mix of Oracle, DB2, SQL Server, etc. These databases may be hosted on different platforms – e.g., Linux, Solaris, Windows, etc. Participating databases can also maintain other data structures using GoldenGate to transform the data into the appropriate format. All these capabilities enable large enterprises to simplify their IT environment by making GoldenGate a single standard for replication technology.

Active – Active HA

In a GoldenGate active-active configuration, the source and destination databases are available for reading and writing, yielding a distributed configuration where any workload can be balanced across any participating database. This provides high availability and data protection should an individual site fail. It also provides an excellent way to perform zero downtime maintenance – by implementing changes in one replica, synchronizing it with a source database operating at the prior version, and then gradually transitioning users with zero downtime to the replica operating at the new version.

Because users in a GoldenGate active-active configuration can update different copies of the same table anywhere, update conflicts may result from changes made to the same data element in different databases simultaneously. Oracle GoldenGate provides various options for avoiding, detecting, and resolving conflicts. These options can be implemented globally, on an object-by-object basis, based on data values and filters, or through event-driven criteria, including database error messages.

Over the past few releases, Oracle GoldenGate has introduced many new features such as self-describing trial files for simplified user experience, automatic heartbeat with real-time end-to-end replication lag, support for big data, and support for new databases and enhanced monitoring, performance, and integration with invisible column support, Data Pump and Clusterware integration.

Zero Downtime Maintenance using Oracle GoldenGate

Oracle GoldenGate is the most flexible method for reducing or eliminating planned downtime. Its heterogeneous replication can support virtually any platform migration, technology refresh, database upgrade, and application

upgrades that change back-end database objects with minimal downtime. GoldenGate logical replication can keep databases on different platforms or versions synchronized. This enables changes to be implemented on a production copy, then synchronized with the old version. Once validated, users are switched to the copy running at the new version or on the new platform. GoldenGate one-way replication does require some downtime while all users are disconnected from the old version and reconnect to the new. GoldenGate bidirectional replication using conflict resolution enables gradual migration of users from the old version for zero downtime.

For more details on Oracle GoldenGate, refer to: oracle.com/goto/goldengate.

Online Application Upgrades: Edition-Based Redefinition (EBR)

Oracle Database's Edition-Based Redefinition (EBR) feature allows the online upgrade of an application with uninterrupted availability. When the installation of the upgrade is complete, the pre-upgrade application and the post-upgrade application can be used at the same time. This means that an existing session can continue to use the pre-upgrade application until its user ends it, while all new sessions use the post-upgrade application. Once all sessions that use the pre-upgrade application end, the old edition can be retired. Thus the application as a whole enjoys a hot rollover from the pre-upgrade version to the post-upgrade version. With the introduction of Edition-based Redefinition, a new scope has been introduced -- an edition:

- Code changes are installed in the privacy of a new edition.
- Data changes are made safely by writing only to new columns or tables not seen in the old edition. An editioning view exposes a different projection of a table into each edition so each sees just its columns.
- A cross-edition trigger propagates data changes made by the old edition into the new edition's columns or (in hot-rollover) vice-versa.

Not only can EBR be utilized in custom applications, but it is routinely also used by Oracle Applications such as E-Business Suite.

For more details on Oracle EBR, refer to *https://www.oracle.com/database/technologies/high-availability/ebr.html*

Managing Oracle Database High Availability Solutions

With a comprehensive High Availability solution, as discussed in the MAA reference tiers, you would also expect a single pane of glass solution to monitor, diagnose, and manage your Oracle Database environment. Likewise, providing complete load balancing control is critical, particularly in active-active configurations. To address these needs, Oracle provides Oracle Enterprise Manager Cloud Control as a monitoring, diagnostics, and management platform and Global Data Services for load balancing. The sections below describe these solutions in more detail.

Oracle Enterprise Manager

Oracle Enterprise Manager (OEM) Cloud Control is the management interface for an entire Oracle environment for one or more data centers. Cloud Control delivers centralized management functionality for the complete Oracle IT infrastructure, including systems running Oracle and non-Oracle technologies. With a broad set of administration, configuration management, provisioning, end-to-end monitoring, diagnostics, and security capabilities, Oracle Enterprise Manager Cloud Control reduces the cost and complexity of managing complex environments. In addition, OEM helps customers maintain their required IT infrastructure service levels.

The latest release of Oracle Enterprise Manager Cloud Control includes key High Availability capabilities as follows:

- OEM offers a High Availability Console that integrates monitoring of various High Availability areas (e.g., clustering, backup & recovery, replication, standby databases, and disaster recovery), provides overall High Availability configuration status, and initiates appropriate operations.

- OEM directly integrates Fleet Patching & Provisioning, the most optimized solution for gold image patching and upgrades across your Oracle Database fleet for both Single Instance and RAC deployments.

For more details on Oracle Enterprise Manager Cloud Control, refer to: *https://www.oracle.com/database/technologies/high-availability/em-maa.html*

Full Stack Disaster Recovery

While Oracle Maximum Availability Architecture focuses primarily on the Oracle Database, the entire application system stack must be considered in terms of disaster recovery. This level of planning to protect the whole stack is critical when setting up a remote data center where a switchover or failover of just the database would result in network latency challenges risking application performance impact. Full Stack Disaster Recovery (Full Stack DR) is an Oracle Cloud Infrastructure (OCI) disaster recovery orchestration and management service that provides comprehensive disaster recovery capabilities for all layers of an application stack, including infrastructure, middleware, database, and application. The service fully integrates with Oracle MAA features like Oracle Data Guard when a protection group (the logical representation of an application system) includes an Oracle Database cloud service such as Oracle Autonomous Database or Oracle Exadata Database Service.

Full Stack Disaster Recovery assures comprehensive business continuity from various data center outages to ensure that organizations have minimal impact from region-wide outages, Availability Domain (AD), or regional outages. Full Stack DR is flexible enough to easily integrate with various Oracle platforms, non-Oracle applications, and infrastructure. The service generates, runs, and monitors disaster recovery plans for services and applications deployed in your tenancy. Full Stack DR operates at the service level, so there is no impact on other services running in your tenancy. You can customize the disaster recovery plans generated by the service based on your specific needs.

During the execution of a full stack switchover or failover, Full Stack DR actively monitors the progress of Full Stack DR operations and allows operators to take corrective actions if errors occur during an operation with the option to resume or cancel the orchestration process. In addition, operators can also validate and monitor business continuity readiness and compliance by periodically running Full Stack DR Prechecks.

For more details on OCI Full Stack Disaster Recovery, refer to: https://www.oracle.com/cloud/full-stack-disaster-recovery/

Global Data Services

Many customers have offloaded read-only and read-mostly workloads to their Active Data Guard Standby replicas requiring a load balancer to handle these read-only requests automatically and transparently. In addition, Oracle GoldenGate replication also enables distributing of workloads over multiple databases, both within and across data centers. In replicated multi-data center architectures, dynamic, transparent, automated load balancing and high availability are challenging to implement and operate.

Global Data Services (GDS) addresses those challenges by extending the familiar notion of Database Services to span multiple database instances in near and far locations. GDS extends RAC-like failover, service management, and service load balancing to replicated database configurations (see Figure 12: Global Data Services for Failover and Load Balancing Across Datacenters). GDS provides inter- and intra-region load balancing across replicated databases. For example, it can distribute load across a reader farm composed of standby instances and even direct read traffic to the primary if conditions warrant it. GDS is intended for applications that are replication-aware.

Global Data Services (GDS) benefits include:

- Higher availability by supporting service failover across local and global databases
- Better scalability by providing load balancing across multiple databases
- Better Manageability via centralized administration of global resources

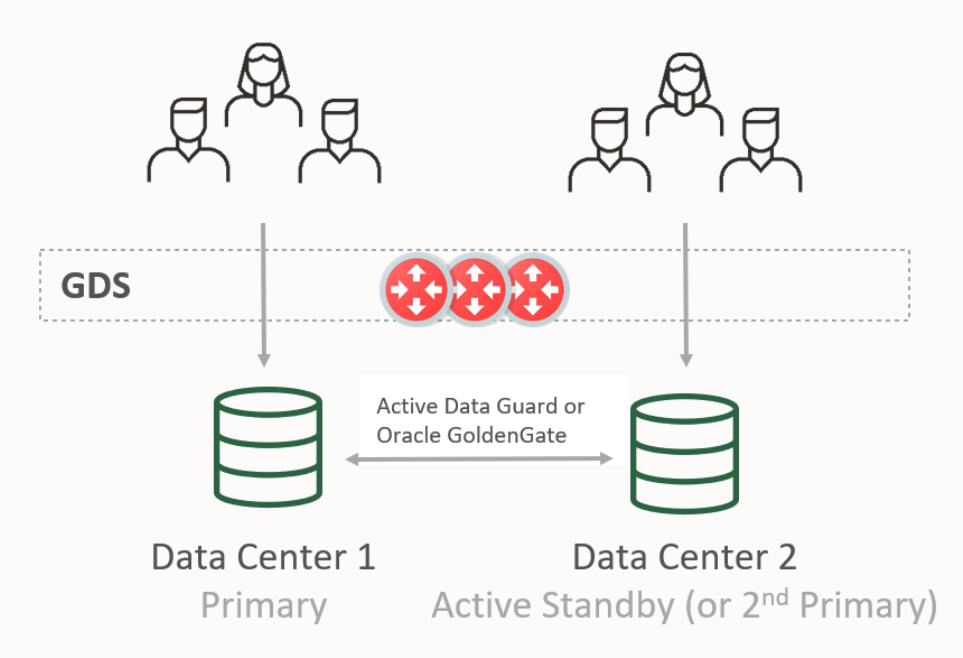

Figure 12: Global Data Services for Failover and Load Balancing Across Datacenters

The GDS example in Figure 13 above depicts replicated databases using ADG and OGG, both local and remote, in a GDS Configuration. Read Write Service runs on the Master database (DB01). Upon the failure of the Master, GDS will failover the global service to another available database (in this case, DB02)

With Active Data Guard, GDS supports:

- Automatic role-based services integrated with Data Guard role transitions.
- Replication lag-based routing

With Active Data Guard and GoldenGate, GDS supports:

- Service failover and load balancing across replicated databases in local and remote data centers.
- Region-based routing

With GoldenGate, GDS supports failover and load balancing for local and remote data centers. When Active Data Guard and Oracle GoldenGate allow offloading production workloads to the replication assets, GDS enables better replica utilization, yielding better performance, scalability, and availability.

Conclusion

Successful enterprises deploy and operate highly available technology infrastructures to protect critical data and information systems. The Oracle Database is at the core of many mission-critical information systems, responsible for the information technology infrastructure's availability, security, and reliability. Building on decades of innovation, Oracle Database continues improving its world-class availability and data protection solutions to maximize data and application availability in the event of planned maintenance activities and unexpected failures.

Oracle's MAA best practices empower customers to achieve their high availability goals by deploying resources and technology commensurate to their requirements and constraints. These best practices enable customers to attain High availability on various platforms and deployments. MAA applies to database deployments on low-cost commodity servers, where availability and performance are enhanced by horizontal scalability, or to the Oracle Cloud, where these HA solutions are automatically configured and maintained depending on your selected cloud option (i.e.,

Autonomous Database). MAA also applies to high-end, storage, and general-purpose servers. Lastly, Oracle's engineered systems are built from the ground up following MAA. Customers seeking extreme performance with maximum availability deploy Oracle Exadata Database Machines as the core of their database-centric IT infrastructure. The same deep understanding of IT infrastructure and database technology that underlies Oracle's MAA best practices, with proven success in thousands of global, mission-critical deployments, also underlies Oracle Exadata Database Machines, which provides the foundation for the Oracle Cloud as well.

Oracle's High Availability solutions have widespread customer adoption and continue to be a critical differentiator when choosing a database technology to support the 24x7 uptime requirements of today's businesses. Review Oracle High Availability customer success stories across various industry verticals worldwide at oracle.com/goto/availability.

Appendix

New High Availability Feature Highlights in Oracle Database 23ai

FEATURE	DESCRIPTION OF NEW OR ENHANCED FUNCTIONALITY IN ORACLE DATABASE 23AI
Data Guard & Active Data Guard	• Oracle Data Guard now provides the capability to decrypt redo operations in hybrid cloud disaster recovery configurations where the cloud database is encrypted with Transparent Data Encryption (TDE) and the on-premises database is not. To enable this feature, Oracle Database introduces the TABLESPACE_ENCRYPTION initialization parameter, which enables you to control the automatic encryption of tablespaces in both the primary and standby databases, for on-premises sand Oracle Cloud Infrastructure (OCI) environments. For example, an on-premises database can be unencrypted and an OCI database can be encrypted.
Application Continuity & Transaction Guard	• Application Continuity and draining of database sessions are supported when you perform a rolling upgrade using DBMS_ROLLING. This is the strategic approach for applying non-rolling patches and for performing major upgrades for the Oracle Autonomous Database. This feature supports the upgrade from a major release to a major release, for example, Oracle Database 19c to Oracle Database 23ai. • Transaction Guard is an application-independent infrastructure that enables recovery of work from an application perspective. With Transaction Guard, each logical transaction may map to single or multiple server-side transactions. Persisting each logical transactions, as part of a commit, introduces overheads in normal transaction operation. Database Native Transaction Guard enhances existing Transaction Guard and does not require persistence in a separate table.
Real Application Clusters	• Local Rolling Database Maintenance provides uninterrupted database availability during maintenance activities (such as patching) for Oracle RAC and Oracle RAC One Node databases. This significantly improves the availability of your databases without causing extra workload on other cluster nodes. • Single-Server Rolling Database Maintenance provides uninterrupted database availability during maintenance activities (such as patching) for single-node Oracle RAC or Oracle RAC One Node databases. This significantly improves the availability of your single-node databases without expanding them to a multi-node cluster.
Oracle Flashback	• Flashback Time Travel can automatically track, and archive, transactional changes to tables. Flashback Time Travel creates archives of the changes made to the rows of a table and stores the changes in history tables. It also maintains a history of the evolution of the tables' schema. By maintaining the history of the transactional changes to a table, and its schema, Flashback Time Travel enables you to perform operations such as Flashback Queries (AS OF and VERSIONS) on the table to view the history of the changes made during the transaction time.

Oracle Globally Distributed Database (formerly Oracle Sharding)	• Raft replication provides built-in replication for Oracle Globally Distributed Database without requiring configuration of Oracle GoldenGate or Oracle Data Guard. Raft replication is logical replication with consensus-based (RAFT) commit protocol, which enables declarative replication configuration and sub-second failover. • ODP.NET supports Sharding split partition set events in core and managed drivers. A split partition set is an operation on the Globally Distributed Database performed when moving data within a specified super Sharding key to a different shard space.

For more details on Oracle Database 23ai New High Availability features, refer to: https://docs.oracle.com/en/database/oracle/oracle-database/23/nfcoa/high-availability.html

New High Availability Feature Highlights in Oracle Database 21c

FEATURE	DESCRIPTION OF NEW OR ENHANCED FUNCTIONALITY IN ORACLE DATABASE 21C
Data Guard & Active Data Guard	• Data Guard for Multitenant Portable Databases enables the choice of replication and failover/switchover at either the PDB level or default CDB level allowing administrators to choose the granularity of their disaster recovery configuration. • Using a single command, the Data Guard Broker now enables users to create and add a Far Sync instance to a Data Guard Broker configuration. • The far sync instance can fully be utilized in Maximum Performance mode in both normal configurations and when fast-start failover (FSFO) is enabled.
Real Application Clusters	• The Global Cache Service (LMS) process is vital to operating an Oracle Real Application Clusters (Oracle RAC) Database. Cache Fusion Hardening helps ensure that the critical LMS process remains running despite discrepancies between instances that would otherwise lead to LMS and database instance failures.
Oracle Globally Distributed Database (formerly known as Oracle Sharding)	• Sharding Advisor is a standalone command-line tool that helps you redesign a database schema to efficiently migrate an existing, non-sharded Oracle Database to an Oracle sharding environment. Sharding Advisor analyzes your existing database schema and produces a ranked list of possible sharded database designs. • Oracle Globally Distributed Database backup and recovery operations are centralized using new commands in the GDSCTL utility. You can define a backup policy for a sharded database as a whole and restore one or more shards, or the entire sharded database, to the same point in time. Configured backups are run automatically, and you can define a schedule to run backups during off-peak hours. • Convert existing databases running the same application into a sharded database without modifying the database schemas or the application. The databases can be geographically distributed and have some differences in their schemas. • If any shards are unavailable during query execution, the enhanced multi-shard query attempts to find alternate shards to operate on. The query resumes without issuing a failure condition. Bulk data loading and DML can operate on multiple shards simultaneously.
Application Continuity	• Application Continuity Protection Check (ACCHK) provides guidance on the level of protection for each application that uses Application Continuity and assists you in increasing protection if required. • Oracle Database invokes planned failover at points where the database knows that it can replay the session using Application Continuity and that the session is not expected to drain. Session Migration is an automatic solution that Oracle Database uses for relocating sessions during planned maintenance for batch and long-running operations that are not likely to be complete in the specified drain timeout period.

	• The reset session state feature clears the session state the application sets when the request ends. The RESET_STATE database service attribute cleans up dirty sessions so that the applications, which use these sessions after cleanup, cannot see the state of these sessions.

For more details on Oracle Database 21c New High Availability features, refer to: https://docs.oracle.com/en/database/oracle/oracle-database/21/nfcon/performance-and-high-availability.html

High Availability Feature Highlights in Oracle Database 19c

FEATURE	DESCRIPTION OF NEW OR ENHANCED FUNCTIONALITY IN ORACLE DATABASE 19C
Data Guard & Active Data Guard	• Active Data Guard - DML operations on the Read Only Standby can be redirected to the Primary database to allow some reporting applications that make infrequent writes to run on the ADG Standby. • Active Data Guard - You can enable the Oracle Database In-Memory Column Store and simultaneously use Data Guard Multi-Instance Redo Apply on an Active Data Guard standby database. • Data Guard - Without impacting your current environment, you can test how fast-start failover will work by using the observe-only mode of fast-start failover. • Data Guard - The process of flashing back a physical standby to a point in time that was captured on the primary is simplified by automatically replicating restore points from the primary to the standby.
Application Continuity	Transparent Application Continuity (TAC) is introduced, which is fully automated and transparently tracks and records sessions and transactional state, and thus recoverable outages are hidden from users.
Real Application Clusters	The new architecture called *Oracle Cluster Domain* frees individual clusters to dedicate all their resources to the database or application as management tasks like deployment, storage management, and performance monitoring are delegated to run on a pre-defined Domain Services Cluster.
Oracle Globally Distributed Database (formerly known as Oracle Sharding)	Globally Distributed Database in Oracle 19c allows multiple table families in the same Globally Distributed Database deployment. Globally Distributed Database in multitenant databases has been enhanced to allow for multiple shards in a single CDB, and a global sequence number concept has been introduced to assist with key generation.

For more details on Oracle Database 19c New High Availability features, refer to: https://docs.oracle.com/en/database/oracle/oracle-database/19/newft/new-features.html

ORACLE

Connect with us

Call +**1.800.ORACLE1** or visit **oracle.com**. Outside North America, find your local office at: **oracle.com/contact**.

blogs.oracle.com facebook.com/oracle twitter.com/oracle