Oracle Maximum Availability Architecture

Maximum Availability with Oracle Database 19c

WHITE PAPER / DECEMBER 17, 2019

ORACLE

PURPOSE STATEMENT

This document provides an overview of features and enhancements included in release Oracle Database 19c. It is intended solely to help you assess the business benefits of utilizing Oracle Maximum Availability Architecture (MAA) to plan the High Availability (HA) architecture for the Oracle Database.

DISCLAIMER

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

TABLE OF CONTENTS

INTRODUCTION

Enterprises use Information Technology (IT) to gain competitive advantages, reduce operating costs, enhance communication with customers, and increase management insight into their business. Thus, enterprises are becoming increasingly dependent on their IT infrastructure and its continuous availability. Application downtime and data unavailability 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 often 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 their return on investment through productive use of all High Availability resources, and improve quality of service to users.

In this paper, we examine the types of outages that affect IT infrastructures, and present Oracle Database technologies that comprehensively address those outages. These technologies, integrated into Oracle's Maximum Availability Architecture (MAA), reduce or avoid unplanned downtime, enable rapid recovery from failures, and minimize planned downtimes.

This paper describes new High Availability features and enhancements made in Oracle Database 19c in terms of performance, functionality, and ease-of-use - including Real Application Clusters (RAC), Automatic Storage Management, Sharding, Recovery Manager, Data Guard and Active Data Guard, Oracle Secure Backup, and Edition-Based Redefinition (EBR).

Oracle Database 19c not only represents a leap in database technology compared to the database versions most commonly used by Oracle's customers today (commonly Oracle Database 11g Release 2 (i.e. 11.2.0.4) and even for that matter, Oracle Database 12c Rel. 2), it is also the first long term support release since the naming and release schedule change. Last but not least, the amount of new features introduced with Oracle Database 19c represents a good balance between innovation and stabilization in that the right amount of time and effort was spent to enhance and stabilize features that have already been introduced with Oracle Database 12c Release 2 and 18c.

THE HIGH AVAILABILITY CHALLENGE

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

In Oracle's perspective, HA encompasses a number of important aspects in addition to the main 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**: 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 of time as the business evolves with no costly surprises or disappointments.

A successful HA implementation begins with the understanding of the service levels required by the business along each of these dimensions. This guides important 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, it is critical that one understands 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 perhaps a migration to a new platform or hardware. Likewise, unplanned outages might consist of server instance outages, site disasters (i.e. flood, long-term power outage, or fire), recovery from human error, or data corruption.

When most IT teams responsible for applications and their associated infrastructure consider these events, they begin looking at options to reduce Recovery Point Objective (RPO) and Recovery Time Objective (RTO) in regards to how much time will be required to handle one of these types of events which could historically take even days or weeks in some cases depending on the severity of these events for the unprepared. Both RPO and RTO are considered two of the most important parameters when putting together a disaster recovery and/or data protection plan and need to be considered carefully per application to determine acceptable thresholds of 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: the loss, damage, or corruption of business-critical data. 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 difficult 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, a good 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 a set of powerful capabilities that help administrators prevent, diagnose and recover from such errors. It also includes features for end-users to directly recover from problems, speeding 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 file system to volume manager to 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 provide administrators with 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 12*c* enables planned maintenance to be performed online to the production version of the database, or in rolling fashion using a synchronized copy of the production database, or using bi-directional replication between two copies of the production database to migrate from one version to the next with zero downtime. The following sections address these capabilities.

ORACLE DATABASE HIGH AVAILABILITY

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

Oracle Database 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 unique levels of 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 efficiently support your business objectives today 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 at the earliest opportunity, Oracle Data Guard prevents propagation of physical corruption, logical intra-block corruption, and logical corruptions 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. RMAN enables a backup once, incremental forever strategy that only backs up changed blocks, providing implicit source-side deduplication that is more efficient than an external de-duplication appliance. RMAN also provides fine-grained, efficient recovery of individual blocks instead of entire data files. Another unique-to-Oracle example of data protection is the Flashback capability to undo database changes at a level of granularity appropriate to the scope of the error, be it the entire database, a table, or an individual transaction, without requiring a full database restore.

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) enables 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. The most important enhancement to Oracle's High Availability set of solution in this context therefore is Application Continuity (AC), a new capability first made available in Oracle Database 12*c*. 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. AC works with Oracle RAC (One Node) and (Active) Data Guard. With the release of Oracle Database 19c, Transparent Application Continuity provides enhances transparency to the customer regarding its coverage. Last but not least, Oracle's High Availability solutions go beyond unplanned outages. All types of database maintenance can be performed either online or in rolling fashion for minimal or zero downtime. Active Data Guard (ADG) standby systems are easily dual-purposed as test systems, reducing risk by ensuring all changes are fully tested on an exact copy of the production database before they are applied to the production environment.

3. *Provide an integrated, automated, and open architecture with 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 fail over to a standby database if the

production database becomes offline; to automatically remove and archive backups for effective space management; and to 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, and Oracle GoldenGate supports read-write workloads with conflict resolution distributed across replicated copies of an Oracle Database in an update-anywhere architecture. 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 useful 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).

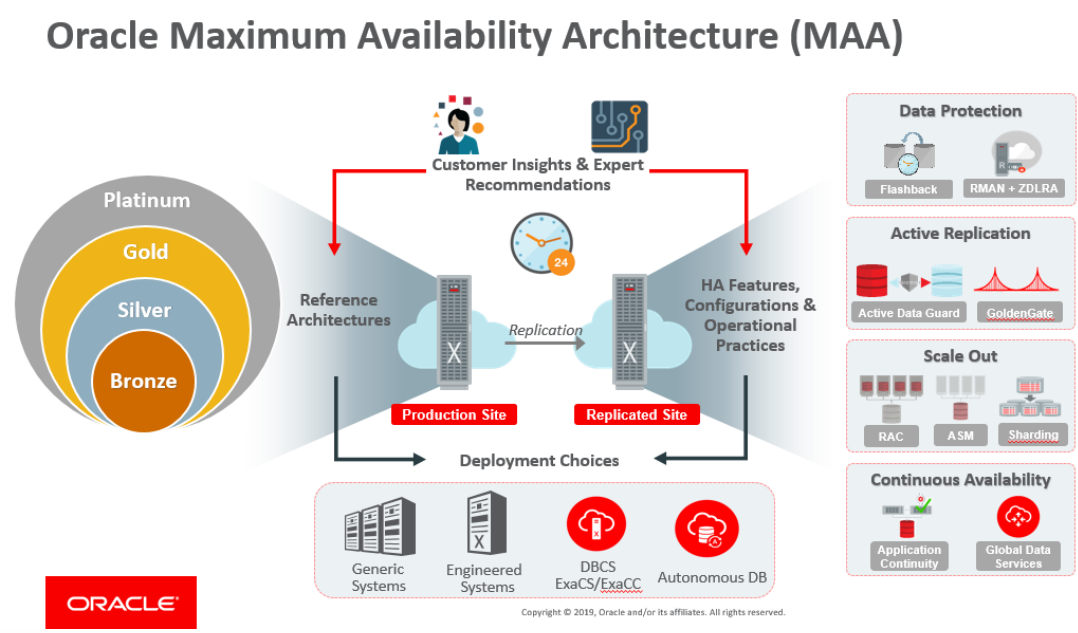

Figure 1: Oracle's High Availability Technologies and the Oracle Maximum Availability Architecture

For over a decade, MAA best practices have been created and maintained by a team of Oracle engineers that continually validate the integrated the 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 offerings, is operated following those standards that have ensured maximum availability for many of Oracle's customers for decades.

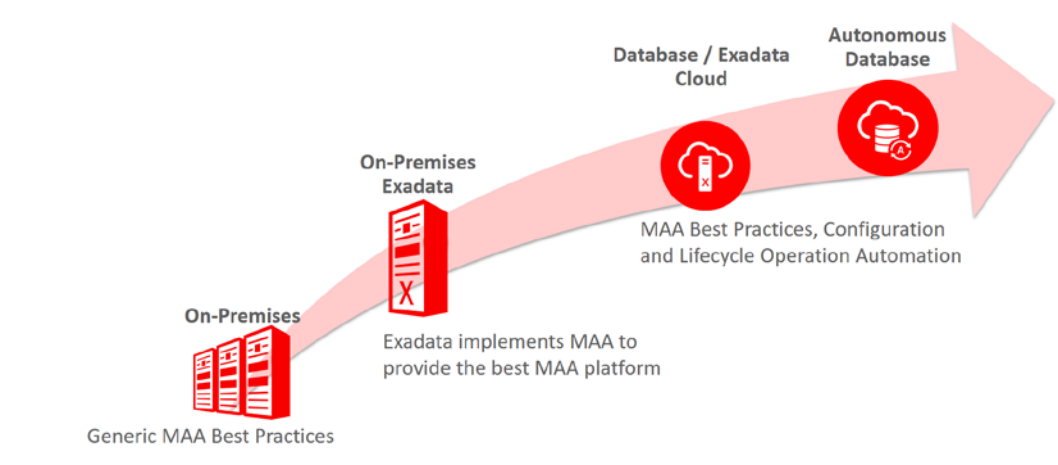

Figure 2: MAA Evolution from On-Premises into the Oracle Cloud

Last but not least, Oracle MAA has evolved to be the new de facto High Availability standard. In the absence of any other comprehensive literature on this subject, Oracle MAA acts as a general guidance for any database operator 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 that it also provides 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 database and especially future Oracle customers (see Figure 3: Oracle MAA is for Everyone! below) 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 interesting 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.

1 https://www.oracle.com/a/tech/docs/maa-overview-onpremise-2019.pdf

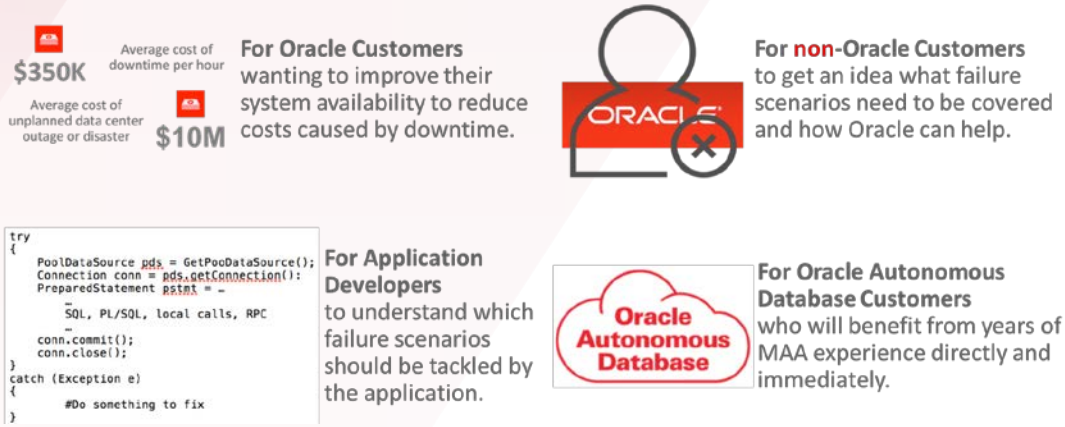

Figure 3: Oracle MAA is for everyone!

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 and are designed to provide an overview of how Oracle's High Availability features can help to tackle any of the disruptive use cases discussed above whether they fall into the planned maintenance or unplanned outage category. Oracle's MAA blue prints build on one another in a hierarchical fashion, 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 specific application along with the associated end-users and business 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 both planned and unplanned outages may be acceptable as long as recovery is possible within a specified timeframe. For example, the acceptable timeframe of recovery for an unplanned data corruption event may be few hours or an hour of downtime may be acceptable in cases where a human mistake occurs corrupting 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:

Bronze - Single Instance Oracle Database

	Events	Downtime (RTO)	Data Loss Potential (RPO)
Unplanned Outages	Recoverable Database instance failure	Minutes	Zero
	Recoverable server failure	Minutes to hour	Zero
	Data corruptions, unrecoverable instance, server, database or site failures	Hours to days	Since last backup or Near Zero with RA
Planned Maintenance	Online file move, reorganization/redefinition, and certain patches	Zero	Zero
	Hardware or operating system maintenance and database patches that cannot be done online	Minutes to hours	Zero
	Database upgrades: patch sets and full database releases	Minutes to hours	Zero
	Platform migrations	Hours to a day	Zero
	Application upgrades that modify back-end database objects	Hours to days	Zero

Figure 4: Bronze Tier RTO and RPO Levels of Protection

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

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, to tape, and to cloud storage. This wide range of backup options allows users to deploy the best solution for their particular environment. The following sections discuss Oracle's disk, tape, and 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 very 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 successfully restore and recover a database 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.

Other recent 18c/19c RMAN enhancements to provide increased performance and ease-of-use include:

- RMAN support for multi-section backup of image copies and incremental backups.
- A new RMAN capability has been introduced to recover a standby database using "FROM SERVICE" to perform an Active Data Guard synchronization refreshing the standby DB with a single RMAN command.
- Quick synchronization of a standby database with the primary database using simple RMAN command: RECOVER DATABASE FROM SERVICE.
- Direct support for SQL statements by the RMAN command line (CLI) – no SQL keyword or quotes needed.
- Enhanced feature integration with Data Guard to allow Far-Sync database creation, validation and repair of standby database blocks that were invalidated due to primary data changed using NOLOGGING.
- RMAN supports the multitenant architecture. The familiar BACKUP DATABASE / RESTORE DATABASE command now backs up / restores the Multitenant Container Database (CDB), including all its Pluggable Databases (PDBs).

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

Fast Recovery Area

A key component of 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. A single FRA may be shared by one or more databases.

ORACLE SECURE BACKUP (OSB)

Oracle Secure Backup (OSB) is Oracle's enterprise-grade media management solution for both database and file system data. Oracle Secure Backup delivers scalable, centralized backup management for distributed, heterogeneous IT environments, by providing:

- Oracle database integration with Recovery Manager (RMAN) supporting versions Oracle Database 10g Release 2 to Oracle Database 19c with optimized performance achieving 25-40% faster backups than comparable media management utilities with up to 10% less CPU utilization
- Faster data transfer from Exadata and/or Oracle Database Appliance (ODA) to media servers by leveraging RDS / RDMA (Reliable Datagram Sockets over Remote Data Memory Access) over InfiniBand (IB)
- File system data protection: UNIX / Windows / Linux servers
- NAS data protection leveraging the Network Data Management Protocol (NDMP)
- Supports cloud storage target devices and disk based devices in addition to tape libraries
- Staging devices for rule-based migration of duplication: Disk to Tape or Disk to Cloud
- Advanced Software Compression

For more details on Oracle Secure Backup resources refer to https://www.oracle.com/database/technologies/high-availability/secure-backup.html.

ORACLE DATABASE BACKUP CLOUD SERVICE

Oracle Database Backup Cloud Service is a low cost offsite storage backup solution for storing backups in Oracle cloud. This service securely backs up Oracle Databases that are deployed on-premises or Oracle cloud using RMAN to the cloud. The data is encrypted and securely transmitted over HTTPS/SSL. Backup data is then stored in multiple copies in the cloud for high availability and can be accessed anytime for restore and validation. The encryption keys are kept with the customer. The data can be optionally replicated to another cloud datacenter for disaster recovery. The backup data can be used to instantiate database instances in the cloud using UI for test/dev or DR purposes.

Oracle Database Backup Service Cloud module supports all major platform and all supported Oracle Databases. Administrators can use Oracle Enterprise Manager 13c, RMAN CLI or 3rd party software like Cloudberry to perform backup & recovery management.

For more details on Oracle Database Backup Cloud Service, refer to cloud.oracle.com/database_backup.

REAL-TIME DATA PROTECTION – ZERO DATA LOSS RECOVERY APPLIANCE

The Zero Data Loss Recovery Appliance (ZDLRA) is an innovative data protection solution that is completely 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 easily 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 that includes 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 transmit 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 allows 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 true incremental-forever backup strategy. Protected databases send RMAN incremental backups to the Recovery Appliance after an initial full backup. RMAN block change tracking is used to send 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, which are essentially space-efficient pointer-based representations of physical full backups as of an incremental backup point-in-time. When the time comes for a restore operation, Delta Store efficiently recreates a physical full backup from the appropriate incremental backup point. Archived log backups stored by the appliance are then used to roll forward to the exact point in time desired. The Delta Store eliminates typical production server overhead of traditional restore and apply of successive incremental backups. The performance of the restore operation is further optimized by the scalability and performance of the underlying Exadata-based hardware architecture.
- End-to-end data validation as deltas are received combined with on-disk background validation of existing backups. Logical and physical validation using deep knowledge of Oracle block structure provides a level of protection un-matched by other backup solutions.
- Secure replication of backups between Recovery Appliances. This protects against potential outages of a Recovery Appliance and provides disaster protection against site outages. Deltas and redo can also be sent directly from a protected database to a remote Recovery Appliance for disaster protection.
- Low cost, autonomous, 24x7 tape archival without impacting production database servers. The Recovery Appliance comes pre-installed with Oracle Secure Backup (OSB) media management software. It supports a 16Gb Fibre Channel Adapter on each compute server within the appliance so that OSB can connect directly to tape hardware without costly third party tape backup agents or specialized media servers.
- End-to-End visibility and management of the data protection life-cycle using Enterprise Manager Cloud Control. Beginning from the time the backup is created by RMAN on the database, to the time it is stored on disk, on tape, and/or replicated to another appliance in a remote data center. All backup locations are tracked by the Recovery Appliance catalog. Any RMAN restore and recovery operation can retrieve the most appropriate backups wherever they reside.

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 that have recovery time objectives that can be achieved by a restore from backup. 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.

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 depend on the database size, a capability unique to the Oracle Database that becomes a necessity as database sizes continue to grow. Flashback supports recovery at all levels including the row, transaction, table, and the 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 still avoiding the need for long periods of downtime, e.g., all of yesterday's customer orders have been deleted. The following sub-sections walk through some of the key features of Flashback.

Flashback Query

Using Oracle Flashback Query, administrators are able to query any data at some point-in-time in the past. This powerful feature can be used to 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 for example 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 Versions Query

Flashback Versions Query enables administrators to retrieve different versions of a row across a specified time interval instead of a single point-in-time. For instance, a query such as:

```
SELECT * FROM emp VERSIONS BETWEEN TIMESTAMP time1 AND time2 WHERE…
```

This displays each version of the row between the specified timestamps, including the transactions that operated on the row. The administrator can pinpoint when and how data has changed, providing great utility in both data repair and application debugging.

Flashback Transaction Query

Logical corruption may also result when an erroneous transaction changes data in multiple rows or tables. Flashback Transaction Query allows an administrator to see all the changes made by a specific transaction. For instance, a query such as:

```
SELECT * FROM FLASHBACK_TRANSACTION_QUERY WHERE XID = transactionID
```

This shows changes made by this transaction. It also produces the SQL statements necessary to undo (flashback) the transaction (where transactionID may be obtained via a Flashback Versions Query). This precision tool empowers the administrator to efficiently pinpoint and resolve logical corruptions in the database in relation to a transaction.

Flashback Transaction

Often, data failures take time to be identified. Additional ‘good’ transactions may have executed on data logically corrupted by an earlier ‘bad’ transaction. In this situation, the administrator must analyze changes made by the ‘bad’ transaction and by any other (dependent) transactions that subsequently modified 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 flash back a single ‘bad’ transaction, and optionally, all of its dependent transactions, with a single PL/SQL operation. Alternatively, an administrator can use Oracle Enterprise Manager Cloud Control to identify and flash back the necessary transactions.

Flashback Table

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

```
FLASHBACK TABLE orders, order_items TIMESTAMP time
```

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

Flashback Drop

Getting back an erroneously dropped table used to require restore, recovery, export/import, and re-creation of all associated table attributes. With Flashback Drop, dropped tables can be easily recovered, via a FLASHBACK TABLE <table> TO BEFORE DROP statement. This restores the dropped table and all of its indexes, constraints, and triggers, from the Recycle Bin (logical container for dropped objects).

Flashback Database

To restore an entire database to a previous point-in-time, the traditional method is to restore the database from a RMAN backup and recover to the point-in-time prior to 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 a matter of minutes via a simple command like:

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 has no effect on database availability or performance. Many data definition language (DDL) maintenance operations allow administrators to specify timeouts on lock waits in order 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 they are still maintained by DML operations. 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. Recent enhancements to Online Table redefinition in Oracle Database 12*c* include:

- Online Data File Move and Online Partition Move ensures data availability during maintenance operations
- Online redefinition of tables with VPD policies with new parameter copy_vpd_opt in start_redef_table.
- Single command redefinition with new REDEF_TABLE procedure.
- Improved sync_interim_table performance, improved resilience of finish_redef_table with better lock management, and better availability for partition redefinition with only partition-level locks, and improved performance by logging changes for only the specified partitions.

As of the Oracle Database 12.2 *release*, Online Redefinition is enhanced to support even the largest and busiest databases by allowing to resume at point of failure, running the process without acquiring exclusive DDL locks, track changes for faster rollback and support for binary XML storage changes, BFILE and invisible columns.

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, the need for a high availability solution that provides near instantaneous recovery from unexpected outages becomes a necessity. 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 as well as any unexpected disruptions to the database instance itself while further reducing the data loss potential. Routinely, there is also a desire to ensure all of the infrastructure is being utilized as well and thus there is a desire to ensure redundant instances are active and are being utilized to handle workload providing additional performance and scalability benefits.

With those needs in mind, Oracle has provided the silver MAA tier which expands on the Bronze MAA tier capabilities by adding active-active clustering, automatic storage management (ASM) and Application Continuity. All of these technologies continue to evolve but the RPO and RTO thresholds can be seen below in Figure 5: Silver Tier RTO and RPO Levels of Protection.

Silver – High Availability with Fast Failover

	Events	Downtime	Data Loss Potential
Unplanned Outages	Database instance failure	Zero	Zero
	Recoverable server failure	Seconds	Zero
	Data corruptions, database unable to restart, site failure	Hours to days	Since last backup, or Near-zero with ZDLRA redo transport
Planned Maintenance	Online file move, reorganization/redefinition, and patching	Zero	Zero
	Hardware or operating system maintenance and database software updates that cannot be done online	Zero	Zero
	Database upgrades: patch sets and full database releases	Minutes to hours	Zero
	Platform migrations	Hours to a day	Zero
	App upgrades that modify back-end database objects	Hours to days	Zero

Figure 5: 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 fault. Oracle Real Application Clusters (RAC) can provide the most effective protection 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 and 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, the failure of 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 gives users the flexibility 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.

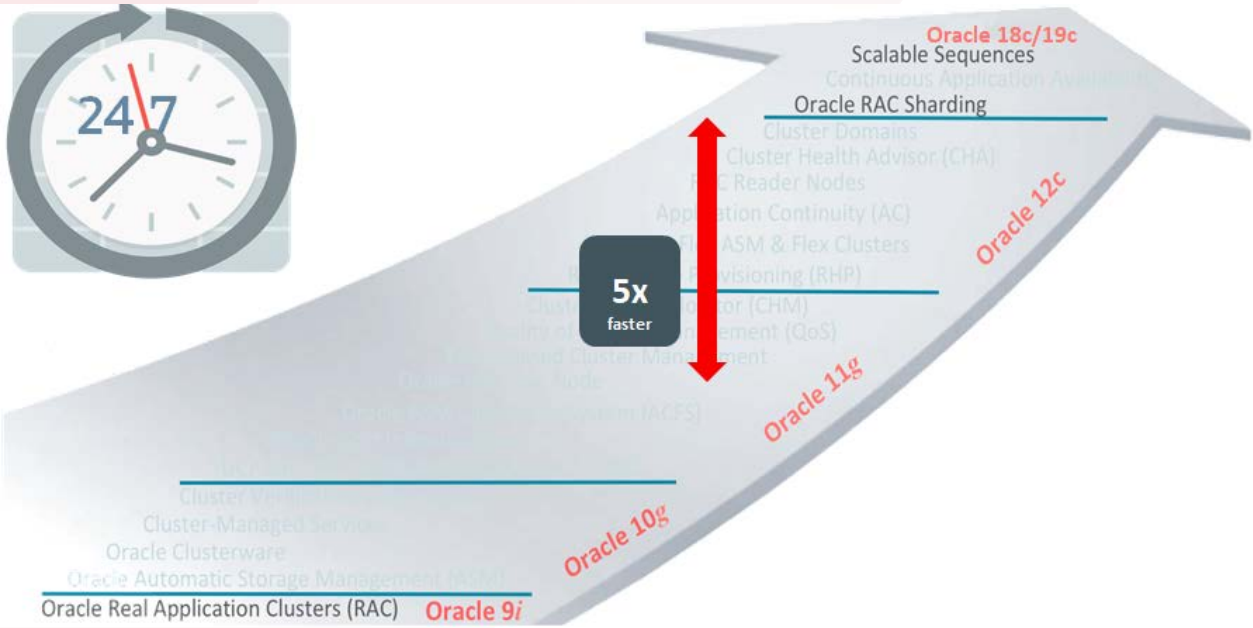

Figure 6: Oracle RAC Downtime Improvements and HA Features

Oracle RAC has provided twenty years of innovation and over the last three releases has especially been enhanced in the areas of planned and unplanned downtime prevention as well as scalability. It is therefore no surprise that Oracle RAC 18c, continued in Oracle Database 19c, has set new records in scalability and recoverability considering outage times incurred by a failure or a maintenance task. Just to provide a simple example, Oracle RAC performance has been improved by 5-times between Oracle 11g Release 2 and Oracle Database 19c for high contention workloads, while the average brownout time has gone down by up to 6-times.

Many features have contributed to those improvements over the last three releases (Oracle Database 12c Rel. 2, Oracle Database 18c, and Oracle Database 19c). *Node Weighting,* introduced with Oracle Database 12c Rel. 2), is only one example. Node Weighting, or "Smart Fencing" considers the workload hosted in the cluster during a fencing operation with the goal to let the majority of the work to survive in the event everything else is equal. The relatively recent addition of Autonomous Health Framework which continuously monitors and analyzes the production database estate and can be hosted on the Domain Service Cluster (DSC) instead of on each production cluster is yet another feature to be mentioned here.

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 greatly simplifies provisioning, patching and upgrading RAC as well as single instance databases across large scale deployments. It provides standardization via the utilization of a gold image for out of place patching across a fleet of databases. Components that form the FPP Server 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 in to 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 12*c* introduced Application Continuity (AC), a capability that 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 Data Guard and replaying the failed transaction so that it will come to an successful end from an application perspective. Application Continuity performs these steps beneath the application so that the outage simply appears in the application as a slightly delayed execution.

In Oracle Database 12.2, Application Continuity was enhanced to support OCI, ODP.NET unmanaged, JDBC Thin on XA, Tuxedo and SQL*Plus clients. By supporting relocation or stopping of services of a database, Application Continuity made it easy to migrate existing connections to another database instance even if Oracle Connection Pools were not used.

With the recent release of Oracle Database 18c and 19c, *Transparent Application Continuity (TAC)* was introduced which tracks and records session and transactional state 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. AC (with TAC) therefore 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.

STORAGE: AUTOMATIC STORAGE MANAGEMENT

Automatic Storage Management (ASM) is a purpose-built file system and volume manager for the Oracle Database. 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 both single-instance databases as well as for 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 and resynchronization (self-healing) and rolling upgrades. ASM also supports dynamic and on-line 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 new type of Disk Groups called Flex Disk Groups. With Flex Disk Groups, all files belonging to an individual database or in the case of multitenancy, a PDB, are collectively identified with a new ASM object called a File Group. A File Group logically contains the files associated with a single database and 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.

Another critical feature from the storage management and consolidation perspective is quota management. Without the means of providing quota management, a single database can consume all the space in a particular Disk Group. Flex Disk Groups therefore offer a new feature called Quota Groups.

A Quota Group is a logical container specifying the amount of Disk Group space that one or more File Groups are permitted to consume. As an example, Quota Group A contains File Groups DB1 and DB2, whereas Quota Group B contains File Group DB3. The databases in Quota Group A are then limited by the specification of available space in that Quota Group so as not to consume all of the Disk Group space.

For the purpose of improving database availability, ASM also provides support Extended Disk Groups, which build the foundation for the Oracle Extended Clusters architecture that can now also be found on Exadata Database Machines. This provides the capability to extend 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 datacenters so that availability is case of a complete failure of one or more data center within close proximity.

Extended Disk Groups eliminated one previous limitation of Oracle Extended RAC Clusters. Historically, a Disk Group in an Extended RAC implementation could have at most two Failure Groups, each in a different datacenter. However, given that Extended Disk Groups are an extension of Flex Disk Groups, this allows multiple Failure Groups within a single datacenter or site. This means that more than one copy of a file's extent can exist, enabling mirroring within a datacenter, as well as across datacenters.

Following the path of optimized storage management "with the database in mind", Oracle ASM 18c introduced the long awaited ASM Database Clones. The advantage of ASM database clones, when compared with storage array-based replication, is that ASM database clones replicate databases rather than generic files or blocks of physical storage. Storage array or file system-based replication in a database environment requires coordination between database objects being replicated with the underlying technology doing the replication. With ASM Database Clones, the administrator does not need to worry about the physical storage layout.

With the recent release of the Parity Protection feature in Oracle ASM 19c, the solution has been enhanced further providing yet another important feature designed to drive down the total cost of storage management. Parity Protection is an additional option which allows for write-once files such as archive logs and backup sets. Prior to Parity Protection, file protection could be set to unprotected, mirror, and high protection only. Parity protection requires a minimum of three regular (not quorum) failure groups in a flex disk group for its use. If there are three or four failure groups when the parity file is generated, then each parity extent set will have two data extents. In this scenario, the redundancy overhead is reduced by 50% over two-way mirrored files.

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

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

While RTO requirements are routinely optimal for most with the introduction of RAC, the need to recovery from data corruption is still a requirement for many critical applications that are central to business functions. In addition, if there is a requirement for a remote site data center to protect from larger site disasters such as floods, power outages, fire, or other natural disasters, a solution will be required to keep those sites synchronized in order to ensure recovery can be addressed in seconds preventing an outage from the application end-user perspective even in the case of these larger outage events. Often, there is also a need for near zero downtime with planned maintenance activities such as migrating to a new platform whether that be new hardware or perhaps an entirely new deployment platform such as the Oracle Cloud. All of these requirements are addressed in the Oracle MAA Gold Tier, please see Figure 7 below for this tier providing details on how to utilize Active Data

Guard in critical HA Architectures across multiple data centers potentially spanning long distances.

Unplanned Outages and Planned Maintenance

Gold – Comprehensive HA and Data Protection

	Events	Downtime	Data Loss Potential
Unplanned Outages	Database instance failure	Zero	Zero
	Recoverable server failure	Seconds	Zero
	Data corruptions, database unable to restart, site failure	Zero* to seconds	Near-zero if ASYNC Zero if SYNC
Planned Maintenance	Online file move, reorganization/redefinition, and patching	Zero	Zero
	Hardware or operating system maintenance and database software updates that cannot be done online	Zero	Zero
	Database upgrades: patch sets, full database releases	Seconds	Zero
	Platform migrations	Seconds	Zero
	Application upgrades that modify database objects	Hours to days	Zero

* Downtime for physical block corruptions can be zero due to Active DG auto-block repair

Figure 7: Gold Tier RTO and RPO Levels of Protection

REAL-TIME DATA PROTECTION AND AVAILABILITY – ORACLE ACTIVE DATA GUARD

Enterprises need to protect their critical data and applications against events that can take an entire cluster or data center offline. Human error, data corruptions or storage failures can make a cluster unavailable. Natural disaster, 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 restore 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 both 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 either 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 to 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 to 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 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. The time spent waiting for acknowledgement increases as the distance between primary and standby increases, directly affecting application response time and throughput. Those effects can be mitigated using Oracle Fast or Far Sync as described in subsequent sections.

Active Data Guard represents a superset of the Data Guard functionality that includes a number of advanced capabilities for data protection and high availability, as well as 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, as well as memory corruptions that get 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.

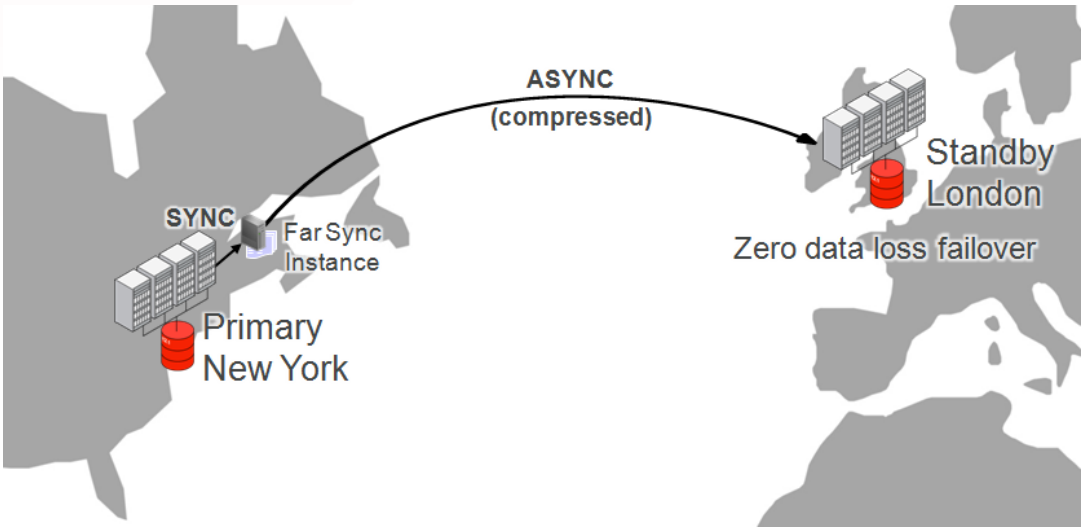

Figure 8: Active Data Guard Far Sync – Zero Data Loss Protection at any Distance

A far sync instance is a light-weight entity that manages a control file and log files. It requires a fraction of the CPU, memory, and I/O resources of a standby database. It does not keep user data files, nor does it run recovery. Its only purpose is to transparently relieve a primary database from serving remote destinations. A far sync instance can also 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 nor is there any requirement for proprietary storage, specialized networking, more database licenses, or complex management.

INCREASE ROI BY OFFLOADING WORKLOADS TO AN ACTIVE DATA GUARD 19C 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 it is enforced 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 introduces new and unprecedented capabilities in this regard. It now allows for DML operations on the read-only standby to be redirected to the primary database which enables even more reporting applications (even those that require occasional writes) to use an Active Data Guard standby database.

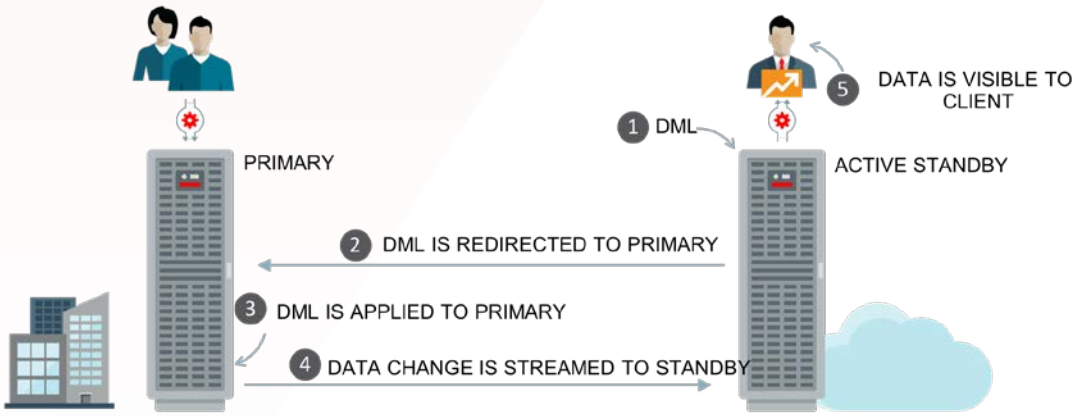

Figure 9: DML Redirect with Oracle Active Data Guard 19c

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 is enabled.

Data Guard Standby-First Patch Assurance

Data Guard Standby-First Patch Assurance enables the physical standby to support different software patch levels between a primary and standby databases for the purpose of applying and validating Oracle patches in rolling fashion.[2] Eligible patches include:

- Patch Set Update, Critical Patch Update, Patch Set Exception, and Oracle Database bundled patch, and full release upgrades.
- Oracle Exadata Database Machine bundled patch, Exadata Storage Server Software patch.

Database Rolling Upgrades using Data Guard

The transient logical database rolling upgrade process uses a Data Guard physical standby database to install a complete Oracle Database patch set (i.e. Oracle 11.2.0.1 to 11.2.0.3), or major release (i.e. Oracle 11.2 to 12.1), or perform other types of maintenance that change the logical structure of a database. The process begins with a primary and physical standby database. The standby is upgraded first as usual, except

2 See MOS Note 1265700.1 for more information on Standby-First Patch Apply eligible patches.

in the case Data Guard logical replication (SQL Apply) is used on a temporary basis to synchronize across old and new versions. Unlike Redo Apply, logical replication uses SQL to replicate across versions and thus is unaffected by differences in physical redo structure that may exist between different Oracle releases.

A switchover moves the production to the new version on the standby database after the upgrade and resynchronization with the original primary is complete. The original primary is then flashed back to the point where the upgrade process began and converted to a physical standby of the new primary. The physical standby is mounted in a new Oracle home, upgraded and resynchronized using redo generated by the new primary (a second catalog upgrade is not required).

Database Rolling Upgrades using Active Data Guard

Although the database rolling upgrade process described above is very effective at reducing planned downtime, it is a manual procedure with many steps and thus error-prone. This creates reluctance to use the rolling upgrade process that results in users accepting longer downtimes associated with traditional upgrade methods. Traditional upgrade methods also increase risk because maintenance is performed on the production database BEFORE it is possible to be certain of the outcome.

Database Rolling Upgrades using Active Data Guard, introduced in Oracle Database 12*c*, 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 replica of production before moving users to the new version.

You can use this capability for database version upgrades starting from the first patchset of Oracle Database 12*c*.[3] You can use it for other database maintenance tasks with Oracle Database 12*c*.[4]

Platform Migration, Systems Maintenance, Data Center Moves

Data Guard also offers some flexibility for primary and standby databases to run on systems having different operating system or hardware architectures, providing a very simple method for platform migration with minimal downtime.[5] 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.

ADDITIONAL NEW FEATURES IN ACTIVE DATA GUARD AND DATA GUARD

Building on the great capabilities above, the following smaller enhancements have also been made in Data Guard and Active Oracle Data Guard 19c:

- Dynamically change the fast-start failover target without disabling fast-start failover.
- Test how fast-start failover would work by using the observe-only mode of fast-start failover.
- 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 primary to the standby.
- When flashback or point-in-time recovery is performed on the primary database, a standby that is in mounted mode can automatically follow the same recovery procedure performed on the primary.

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

3 You must still the Transient Logical Standby upgrade when upgrading from Oracle Database 11g to Oracle Database12c, or from Oracle Database 12.1 to the first patchset of Oracle Database 12.1.

4 Maintenance tasks include: partitioning non-partitioned tables, changing BasicFiles LOBs to SecureFiles LOBs, moving CLOB-stored XMLType to binary XML-stored, altering tables to be OLTP-compressed.

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

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 an alternative architecture for maximizing fault tolerance via the horizontal partitioning that Oracle Sharding provides and introduces the option of utilizing Edition-based Redefinition to seamless upgrade your application when schema and other changes are required to the underlying database as is often required in major application upgrades. The sections below run through the full breadth of the Platinum MAA tier reference solution below in more detail.

Unplanned Outages and Planned Maintenance

Platinum – Comprehensive HA and Data Protection

	Events	Downtime	Data Loss Potential
Unplanned Outages	Database instance failure	Zero	Zero
	Recoverable server failure	Seconds	Zero
	Data corruptions, database unable to restart, site failure	Zero* to seconds	Near-zero if ASYNC Zero if SYNC
Planned Maintenance	Online file move, reorganization/redefinition, and patching	Zero	Zero
	Hardware or operating system maintenance and database patches that cannot be done online but are qualified for RAC rolling install	Zero	Zero
	Database upgrades: patch sets, full database releases	Zero with GG	Zero
	Platform migrations	Seconds	Zero
	Application upgrades that modify database objects	Zero	Zero

* Zero with Active Data Guard and auto-block repair of physical data corruptions

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-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 flexible options to fully address their replication requirements. GoldenGate also supports replication between a broad range of heterogeneous hardware platforms and database management systems beyond Oracle.

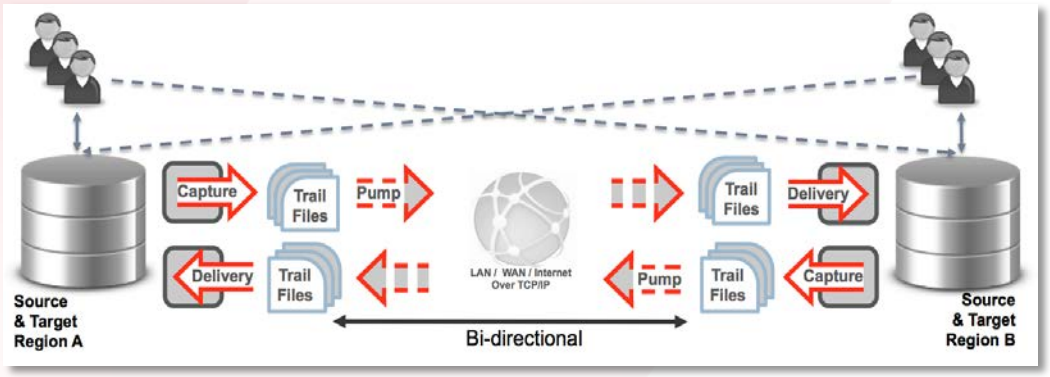

Figure 11: Oracle GoldenGate – Active-Active Bi-Directional Replication

Applications can use GoldenGate with minimal modification or special handling. GoldenGate can be configured, for example, to capture changes for an entire database, or 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 in different platforms – e.g. Linux, Solaris, Windows, etc. Participating databases can also maintain different 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, both 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 at the same time. Oracle GoldenGate provides a variety of 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, DataPump and Clusterware integration. GoldenGate Cloud Services in the cloud which supports active/active bi-directional replication both for cloud deployments and in the hybrid model between on-premises and cloud.

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 many application upgrades that change back-end database objects, with minimal or zero downtime. GoldenGate logical replication is able to keep databases on different platforms or versions synchronized. This enables changes to be implemented on a copy of production, 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.

ORACLE SHARDING

Oracle Sharding 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. Sharding is built on a shared-nothing horizontal partitioning architecture in which the databases do not share storage or rely on cluster software. Oracle Sharding provides a number of benefits for web-scale applications:

- **Linear scalability**. OLTP applications designed for Oracle sharding can elastically scale (data, transactions and users) to any level, on any platform, simply by deploying new shards on additional stand-alone servers. Performance scales linearly as shards are added to the pool because each shard is completely independent from other shards.
- **Extreme Data Availability**. Oracle Sharding 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 either 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 availability of an application while performing database maintenance.
- **Data Sovereignty and Data Proximity via Geographic Data Distribution**. Sharding 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 certain jurisdiction. It also supports storing particular data closer to its consumers.

Figure 12: Oracle Sharding Benefits at a Glance

Sharded Database and Shards

Shards are independent Oracle Databases that are hosted on database servers which have their own 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 (datacenter[s]) or can be placed in different regions. A region in the context of Oracle Sharding represents a datacenter or a multiple datacenters that are 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 Sharding supports three automatically configured replication options: Data Guard, Active Data Guard, or Oracle GoldenGate.

Application Connection and Data-Dependent Routing

A Sharding Key is used for routing the database connection requests at a user session level during connection checkout. Based on this information, a connection is established to the relevant shard which contains the data pertinent to the given sharding_key. Once the session is connected to a shard, all SQL queries and DMLs are supported and executed in the scope of the given shard and require no modification.

Upon the first connection to a given shard, the sharding key range mapping is collected from the shards to dynamically build the shard topology cache, a routing map, which is cached in the client. This allows subsequent requests using sharding keys within the cached range to be routed directly to the shard, thereby eliminating extra network-hops and decreasing the latency for high volume OLTP applications.

Oracle Sharding continues to evolve to ensure it can handle any use case where this level of fault isolation, performance, and scalability are required. With the release of Oracle Sharding 19c, the following new features have been included:

1. Multiple Shards in a Single Multitenant Database
2. Scalable Cross-Shard Query Coordinators
   - Horizontally scalable cross-shard query coordinators can improve performance and availability of read-intensive cross-shard queries. A Shard Catalog can be protected by one or more Active Data Guard standby databases. The primary and all the read-only standby Shard Catalogs can be used as cross shard query coordinator.
3. Middle Tier Sharding
   - The application middle tier can also be sharded to provide affinity to database sharding. Affinitive grouping of middle tiers with database shards is at times referred to as swim lanes. In such deployments, the application's front end routing tier can call a REST API (provided by a sharded database Middle Tier Routing Service), by-passing the sharding key, to retrieve the swim lane details which include the shard name to help route the request to the appropriate shard. This provides better fault isolation, cache locality, scalability and reduction in database connections used by the middle tier.

For more details on Oracle Sharding, refer to: www.oracle.com/database/technologies/high-availability/sharding.html

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 of the application. 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 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 new tables not seen by the old edition. An editioning view exposes a different projection of a table into each edition so each sees just its own 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 is routinely also used by Oracle Applications such as E-Business Suite which has been utilizing EBR since Oracle Database 12.2. With the recent release of Oracle Database 18c and 19c, EBR now enables attribute value constructors for PL/SQL records and collections as well.

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, it is critical to provide full control of load balancing, particularly in active-active configurations. In order 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 Cloud Control 13c 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, while helping customers maintain their required IT infrastructure service levels.

The latest release of Oracle Enterprise Manager Cloud Control (13c) includes key High Availability capabilities as follows:

- It offers a High Availability Console that integrates monitoring of various High Availability areas (e.g. clustering, backup & recovery, replication, disaster recovery), provides overall High Availability configuration status and initiates appropriate operations.
- The Maximum Availability Architecture Configuration Advisor page allows you to evaluate the configuration and identify solutions for protection from server, site, storage, human and data corruption failures, enabling workflows to implement Oracle recommended solutions.
- It enables further MAA automation by enabling migration of databases to ASM and conversion of single instance databases to Oracle RAC with minimum downtime.
- It supports management of the Oracle Secure Backup administrative server and Oracle Secure Backup File System backup/restore and reporting.
- It provides direct integration with Fleet Patching & Provisioning which is 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

COMPLETE SITE FAILOVER: ORACLE SITE GUARD

Oracle Site Guard is part of Oracle Enterprise Manager Cloud Control, and extends automation of disaster recovery to the rest of the Oracle stack. Oracle Site Guard enables administrators to automate complete site failover. Site Guard eliminates the need for specialized skill sets by relieving IT staff of the burden of manually executing complex failover operations, thus reducing the likelihood of human error that can lead to extended downtime and data loss. Site Guard orchestrates the coordinated failover of Oracle Fusion Middleware, Oracle Databases, and is extensible to include other data center components. Site Guard integrates with underlying replication mechanisms that synchronize primary and standby environments and protect mission critical data such as Oracle Data Guard for Oracle data, and storage replication for file system data external to the Database.

For more details on Oracle Site Guard, refer to: *https://www.oracle.com/database/technologies/high-availability/site-guard.html*

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 transpartently. In addition, Oracle GoldenGate replication also enables distributing workloads over multiple databases, both within and across datacenters. In replicated multi-data center architectures, dynamic, transparent, and automated load balancing and high availability are difficult to implement and operate.

Global Data Services (GDS), introduced in Oracle Database 12*c*, 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 13: 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.

In addition to your existing Oracle Databases, GDS requires one or more Global Services Manager (GSMs), and a GDS Catalog Database. Each region has its own GSM (plus replicas for HA), which is a server with specialized software that monitors database load and availability and directs

workload appropriately. To the application layer (the clients using the database services), the GSM looks like a listener. The GDS Catalog is a database (one for the whole GDS framework, but replicated for HA) that hosts the metadata required for GDS to operate, in a manner similar to the RMAN Catalog's hosting of backup metadata.

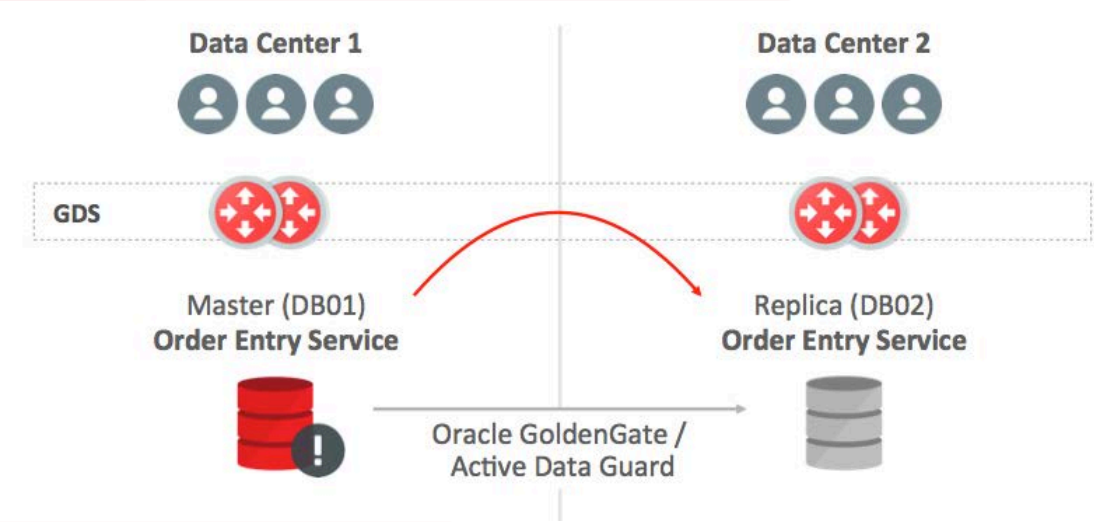

Figure 13: Global Data Services for Failover and Load Balancing Across Datacenters

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. At the core of many mission critical information systems is the Oracle Database, responsible for the availability, security, and reliability of the information technology infrastructure. Building on decades of innovation, Oracle Database 19c continues to improve its world-class availability and data protection solutions to maximize data and application availability in the event of both planned maintenance activities and of 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 a range of 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 be configured and maintained depending on your selected cloud option (i.e. Autonomous Database). MAA also applies to high-end, storage and general purpose servers. Last, but not least, 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 FEATURES IN ORACLE DATABASE 12C

FEATURE	DESCRIPTION OF NEW OR ENHANCED FUNCTIONALITY IN ORACLE DATABASE 12*c*
Application Continuity	Protects applications from database session failures due to instance, server, storage, network or any other related component. Application Continuity re-plays affected "in-flight" requests so that the failure of a RAC node appears to the application as a slightly delayed execution.
Flex ASM	Increases database (instance) availability, facilitation cluster-based database consolidation, by enabling inter-node storage failover and reducing ASM-related resource consumption by up to 60%.
ASM Disk Scrubbing	Checks for logical corruptions and repairs them automatically, in both normal and high-redundancy disk groups. This complements the health checks that RMAN performs during backup and recovery.
Data Guard Fast Sync	Allows a standby to acknowledge the primary database as soon as it receives redo in memory, without waiting for disk I/O to a standby redo log file.
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 with minimal cost or complexity.
Global Data Services (GDS)	Extends Database Services to span multiple database instances in near and far locations. GDS extends RAC-like failover, service management, and service load balancing to a set of replicated databases.
Oracle Secure Backup (OSB)	Faster performance in NUMA (Non-Uniform Memory Access) environments. Increased data transfer rates over InfiniBand (IB) by leveraging of RDS/RDMA instead of TCP / IP. Improved network utilization by load balancing network interfaces.
RMAN and the multitenant architecture	The BACKUP DATABASE / RESTORE DATABASE command now backs up / restores the Multitenant Container Database (CDB), including all its Pluggable Databases (PDBs). RMAN commands can also be applied to individual PDBs, including full backup and restore, using the keyword PLUGGABLE.
Cross-platform	RMAN backup and restore across different platforms for efficient tablespace and database migration.
Other Recovery Manager (RMAN) enhancements	Can recover the most recent or an older version of an individual database table from a backup; tables can be recovered in-place or to a different tablespace. Multi-section backup of image copies and incremental backups. Quick synchronization of a standby database with the primary database using a command. Direct support for SQL statements by the RMAN command line – no SQL keyword needed.
Online Move functionality	Online Data Move enables moving a data file while users are accessing its data, Online Partition Move supports online, multi-partition redefinition in a single session.

Online Table Redefinition enhancements	Single command redefinition. Improved sync_interim_table performance, improved resilience of finish_redef_table with better lock management, better availability for partition redefinition with only partition-level locks, and improved performance by logging changes for only the specified partitions
Upgrades with Active Data Guard	Replaces dozens of steps required to perform a rolling database upgrade with 3 PL/SQL packages that automate much of the process. Minimizes planned downtime and risk by implementing and thoroughly validating all changes on a complete replica of production before moving users to the new version.

APPENDIX: NEW HIGH AVAILABILITY FEATURES IN ORACLE DATABASE 12*C* R2

FEATURE	DESCRIPTION OF NEW OR ENHANCED FUNCTIONALITY IN ORACLE DATABASE 12*c* Release 2
Oracle Sharding	Oracle Sharding is a scalability and availability feature for custom-designed OLTP applications that enables distribution and replication of data across a pool of discrete Oracle Databases that share no hardware or software.
Data Guard / Active Data Guard Enhancements	Many ease of management features, multi-instance redo apply for improved recovery using RAC, support for in-memory column store in standby, ability to run analytical queries and AWR reports which otherwise fail due to read-only standby, auto-repair of standby blocks which were invalidated due to nologging operation in primary, ability to encrypt standby database with no downtime, improved automatic block repair and a bunch of improvements in Oracle Data Guard Broker.
RMAN enhancements	Support for Oracle Sharding, ability to RECOVER TABLE to another schema, many cross-platform enhancements, support for space efficient Sparse Database backups, ability to perform DUPLICATE using backups that are encrypted with non-auto login wallet, additional support for Data Guard enhancements with Far Sync standby creation, duplicate for standby from a standby and repair standby data that got invalidated due to primary nologging operation.
Application Continuity	Support for OCI, ODP.NET unmanaged, JDBC Thin on XA, Tuxedo and SQL*Plus.
Flashback Database	Now supports PDB level Flashback Database operation.
Automatic Storage Management	Cluster Domains, Database-oriented storage management and extreme availability. With new diskgroup type Flex Diskgroups enables easier quota management, redundancy changes and ability to easily and dynamically create database clones for test/dev or production databases. New extended diskgroup to support Extened RAC up to 3 sites.

APPENDIX: NEW HIGH AVAILABILITY FEATURES IN ORACLE DATABASE 18C

FEATURE	DESCRIPTION OF NEW OR ENHANCED FUNCTIONALITY IN ORACLE DATABASE 18C
Oracle Sharding	User-defined sharding, support for PDBs as shards, support of GoldenGate replication with sharding, optimizer enhancements for multi-shard queries are some of the capabilities in Oracle Database 18c.
Data Guard / Active Data	» Global Temporary Table creation is supported with standby databases.

Guard Enhancements	» You can even do DML operations on standby which gets redirected to primary without ACID compromise. » Preservation of buffer cache during role change » No-logging enhancements with two new modes to choose from performance or availability, » Multi-instance Redo Apply support with In-Memory Database » Multi-Instance Redo Apply support with Block Change Tracking
RMAN enhancements	Multitenant PDB backups are made usable after that PDB is plugged into another CDB. PDB cloning to another CSB capability using RMAN DUPLICATE has been added. Encryption and decryption of database during backups has been introduced. You can refresh the standby database from either the primary database or a backup using a single RECOVER command. Oracle RMAN cloud module now supports Oracle Cloud Infrastructure Archive Storage Classic where you can backup and keep it there for longer time with a very low cost of $0.001/GB per month.
Application Continuity	Transparent Application Continuity (TAC) is introduced which is fully automated and transparently tracks and records session 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 its resources to the database or application as management tasks like deployment, storage management, performance monitoring is delegated to run on a pre-defined Cluster called the Domain Services Cluster.
Automatic Storage Management	Customers can now convert to the Flex Disk Group and take advantage of the enhanced management capabilities of Flex Disk Group like (a) modifiable redundancy at individual database file level via File Groups (b) snapshot capabilities and (c) quota management at the database level for consolidated environments Support with bidirectional snapshots and even better integration with Oracle Data Guard when using ACFS to store data files. Customers can additionally utilize ACFS tagging feature to add custom tags to their data and retrieve tags using a command line or using tagging API calls directly from their application.

APPENDIX: NEW HIGH AVAILABILITY FEATURES IN ORACLE DATABASE 19C

FEATURE	DESCRIPTION OF NEW OR ENHANCED FUNCTIONALITY IN ORACLE DATABASE 19C
Oracle Sharding	Sharding in Oracle 19c now allows for multiple table families in the same Sharding deployment. Sharding 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.
Data Guard / Active Data Guard Enhancements	» Active Data Guard - DML operations on the Read Only Standby can be redirected to the Primary database to allow for 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 use Data Guard Multi-Instance Redo Apply at the same time on an Active Data Guard standby database. » Data Guard - You can dynamically change the fast-start failover target without disabling fast-start failover

- 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 primary to the standby.
- Data Guard - When flashback or point-in-time recovery is performed on the primary database, a standby that is in mounted mode can automatically follow the same recovery procedure performed on the primary.

RMAN enhancements	Multitenant PDB backups are made usable after that PDB is plugged into another CDB. PDB cloning to another CSB capability using RMAN DUPLICATE has been added. Encryption and decryption of database during backups has been introduced. You can refresh the standby database from either the primary database or a backup using a single RECOVER command. Oracle RMAN cloud module now supports Oracle Cloud Infrastructure Archive Storage Classic where you can backup and keep it there for longer time with a very low cost of $0.001/GB per month.
Application Continuity	Transparent Application Continuity (TAC) is introduced which is fully automated and transparently tracks and records session 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 its resources to the database or application as management tasks like deployment, storage management, performance monitoring is delegated to run on a pre-defined Cluster called the Domain Services Cluster.
Automatic Storage Management	Customers can now convert to the Flex Disk Group and take advantage of the enhanced management capabilities of Flex Disk Group like (a) modifiable redundancy at individual database file level via File Groups (b) snapshot capabilities and (c) quota management at the database level for consolidated environments Support with bidirectional snapshots and even better integration with Oracle Data Guard when using ACFS to store data files. Customers can additionally utilize ACFS tagging feature to add custom tags to their data and retrieve tags using a command line or using tagging API calls directly from their application.

ORACLE CORPORATION

Worldwide Headquarters
500 Oracle Parkway, Redwood Shores, CA 94065 USA

Worldwide Inquiries
TELE + 1.650.506.7000 + 1.800.ORACLE1
FAX + 1.650.506.7200
oracle.com

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

Integrated Cloud Applications & Platform Services

White Paper **Maximum Availability with Oracle Database 19c**Maximum Availability with Oracle Database 19cMaximum Availability With Oracle Database 19c
December 2019
Author: Glen Hawkins

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