

Oracle Maximum
Availability Architecture

Oracle MAA Reference Architectures

Oracle Database High Availability On-Premises and in the Cloud

ORACLE WHITE PAPER | FEBRUARY 2016



ORACLE®



Disclaimer

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

Table of Contents

Introduction	1
High Availability Reference Architectures – On Premises and Cloud	2
Bronze: Development, Test and Production Databases	3
Bronze Summary	4
Silver: Departmental Databases	5
Oracle Real Application Clusters (Oracle RAC)	5
Oracle Data Guard	6
Decision Point – Oracle RAC or Data Guard	7
Silver Summary	8
Gold: Business Critical Databases	9
Gold Summary	10
Platinum: Mission Critical Databases	10
Application Continuity	11
Active Data Guard Far Sync	11
Oracle GoldenGate	12
Edition-Based Redefinition	12
Platinum Summary	13
Corruption Prevention, Detection, and Auto-Repair	14
Hybrid Cloud	15
Oracle MAA Cloud Roadmap	15
Conclusion	16



Introduction

Enterprises are under intense pressure to do more with less, reduce risk and increase agility. The aggressive consolidation of information technology (IT) infrastructure and deployment of Database as a Service (DBaaS) on public and private clouds is a strategy that many enterprises are pursuing to accomplish these objectives.

Database consolidation drives cost savings by dramatically improving system utilization and reducing management overhead. DBaaS drives cost savings and increased agility through the standardization of I.T. infrastructure and processes. The Cloud enhances these benefits by enabling a more efficient utility model of computing.

All of the above initiatives, however, also incur business risk by amplifying the impact of downtime and data loss. The failure of a standalone environment used by a single developer or small work group is usually of limited impact. The failure of a critical application running in a traditional standalone environment will be immediately felt by the business, but other applications can continue to run unaffected. In contrast, an outage of a consolidated environment supporting an organization's entire development staff or multiple applications used by numerous departments would have a crippling effect on the business. Equally crippling would be an interruption in service at a cloud provider where such applications are running.

The Oracle Maximum Availability Architecture (Oracle MAA) prescribes four HA reference architectures that provide the requisite level of standardization for DBaaS while addressing the complete range of availability and data protection required by enterprises of all sizes and lines of business. All reference architectures are based upon a common platform able to be deployed on-premises or on cloud. This approach makes Oracle MAA simpler and less risky to move to the cloud.

This paper describes Oracle MAA reference architectures in detail and the service level requirements that they address. It is most appropriate for a technical audience: Architects, Directors of IT and Database Administrators responsible for designing and implementing DBaaS and moving to the cloud.

High Availability Reference Architectures – On Premises and Cloud

Oracle MAA best practices define four HA reference architectures that address the complete range of availability and data protection required by enterprises of all sizes and lines of business. The reference architectures are designated PLATINUM, GOLD, SILVER, and BRONZE described in the Figure 1. Each uses an optimal set of Oracle HA capabilities to reliably achieve a given service level at the lowest cost and complexity. They address all types of unplanned outages including data corruption, component failure, system and site outages, as well as planned downtime due to maintenance, migrations, or other purposes.



Figure 1: Oracle Database High Availability Reference Architectures

MAA reference architectures are based upon a common infrastructure optimized for Oracle Database that enables customers to dial-in the level of HA appropriate for different service level requirements. This makes it simple to move a database from one HA tier to the next should business requirements change, or from one hardware platform to another, or from on-premises to the Oracle Cloud.

Bronze Reference Architecture is appropriate for databases where simple restart or restore from backup is 'HA enough'. Bronze uses HA capabilities included Oracle Enterprise Edition. Bronze defaults to the Oracle single tenant architecture (a single pluggable database in a container database) for the simplicity and agility of unplug/plug for upgrades and migrations. Oracle Multitenant is an option for database consolidation (multiple pluggable databases in a single container to reduce operational expenses by managing many databases as one, and to reduce capital costs by increasing consolidation density. Bronze relies upon Oracle-optimized backups using Oracle Recovery Manager (RMAN) to provide data protection and restore availability should an outage prevent the database from being able to restart. A local backup is always recommended for fastest recovery. Oracle also recommends maintaining a second copy of backups in a remote data center to protect against site outages. The Oracle Database Backup Cloud Service can be used to perform cloud-based backup of on-premises databases for this purpose.

Oracle DBaaS Enterprise (PaaS) cloud service combined with the Oracle Backup Cloud Service are the minimum cloud services required for databases deployed on the Oracle Cloud to attain a Bronze level of service.

Silver Reference Architecture is designed for databases that can't afford to wait for a cold restart or a restore from backup should there be an unrecoverable database outage. Silver begins with the same functionality of Bronze and adds capabilities that provide a choice of two different patterns for enhancing availability.

- » The first pattern uses Oracle RAC active-active clustering technology for minimal or zero downtime in the event of database instance or server failure as well as for many types of planned maintenance. Just as with Bronze, RMAN provides database-optimized backups to protect data and restore availability should there be a complete cluster outage. This first pattern requires Oracle Enterprise Edition, Oracle RAC, and Enterprise Manager life-cycle, management, diagnostic and tuning packs for on-premises databases. Oracle DBaaS Extreme Performance (PaaS) cloud service or Exadata cloud service combined with the Oracle Database Backup Cloud Service are the minimum required for databases deployed on the Oracle Cloud.
- » The second pattern uses Data Guard to maintain a local but separate synchronized copy of the production database for HA. While Data Guard does not provide the same transparency to outages or scalability for read-write workloads as Oracle RAC, a local Data Guard copy does provide HA for a broader set of database outages, including data corruptions, human error, cluster outages, network faults and database upgrades. Data Guard synchronous replication with automatic database failover also provides an additional level of data protection beyond that provided by Oracle RAC alone. Data Guard is an included feature of Oracle Enterprise Edition; there are no additional licensed products or cloud services beyond those used for Bronze databases in order to attain Silver service level requirements.

Gold Reference Architecture is well suited for service level requirements that cannot tolerate site outages. It begins with the first architecture pattern for the Silver - Oracle RAC, and adds a remote synchronized copy of the production database using Active Data Guard. This provides a combination of the two architecture patterns presented in Silver, the transparent HA and scalability of Oracle RAC combined with real-time data protection and availability of Active Data Guard.

On-premise systems require Active Data Guard, Oracle RAC and select Enterprise Manager packs to achieve Gold service level requirements. Cloud deployments require DBaaS Extreme Performance (PaaS) or Exadata Cloud services and the Oracle Database Backup cloud service.

Platinum Reference Architecture is designed to make unplanned outages and planned maintenance transparent to the user. Platinum includes Application Continuity for reliable replay of in-flight transactions; Active Data Guard Far Sync for zero data loss protection at any distance; Oracle GoldenGate for zero downtime upgrades and migrations; and Edition-Based Redefinition for zero downtime application upgrades. The Platinum reference architecture delivers substantial value for the most critical applications where downtime and data loss are not an option. Platinum requires the same on-premise products as Gold plus Oracle GoldenGate, and the same cloud services as Gold plus Oracle GoldenGate cloud service.

The following sections describe all reference architectures in greater detail.

Bronze: Development, Test and Production Databases

The Bronze tier (Figure 2) provides basic database service at the lowest absolute cost. A reduced level of HA and data protection is accepted in exchange for reduced cost and implementation complexity.




Figure 2: Bronze HA Reference Architecture

Bronze is based upon single instance Oracle Database with Oracle Restart for auto-restart following recoverable outages. When a machine becomes unusable or the database unrecoverable, the recovery time objective (RTO) is a function of how quickly a replacement system can be provisioned or a backup restored. In a worst case scenario of a complete site outage there will be additional time required to perform these tasks at a secondary location.

Oracle Recovery Manager (RMAN) is used to perform regular backups of the Oracle Database. The potential for data loss, also referred to as the recovery point objective (RPO), if there is an unrecoverable outage is equal to the data generated since the last backup was taken. Database backups should also be replicated to a remote location or on the Cloud to be used for DR should a disaster strike the primary data center.

Bronze utilizes the following capabilities included with Oracle Database Enterprise Edition:

- » [Oracle Restart](#) automatically restarts the database, the listener, and other Oracle components after a hardware or software failure or whenever a database host computer restarts.
- » Oracle corruption protection checks for physical corruption and logical intra-block corruptions. In-memory corruptions are detected and prevented from being written to disk, and in many cases can be repaired automatically. For details see [Preventing, Detecting, and Repairing Block Corruption for the Oracle Database](#).
- » [Automatic Storage Management \(ASM\)](#) is an Oracle-integrated file system and volume manager that includes local mirroring to protect against disk failure.
- » [Oracle Flashback Technologies](#) provide fast error correction at a level of granularity that is appropriate to repair an individual transaction, a table, or the full database.
- » [Oracle Recovery Manager \(RMAN\)](#) enables low-cost, reliable backup and recovery
- » [Online maintenance](#) includes online redefinition and reorganization for database maintenance, online file movement, and online patching.

Bronze Summary

Table 1 summarizes RTO and RPO service level requirements for the Bronze reference architecture.

TABLE 1: BRONZE RECOVERY TIME (RTO) AND POTENTIAL DATA LOSS (RPO)

Event	Downtime- RTO	Potential Data Loss -RPO
Disk failure	Zero	Zero
Machine and recoverable database instance failures	Minutes	Zero
Data corruption and unrecoverable database outages	Hours	Since last backup
Site outages	Hours to days	Since last backup
Database reorganization, file move, eligible one-off patches	Zero	Zero
Hardware and software maintenance and patching	Minutes	Zero
Database upgrades (patch-sets and full releases)	Minutes to Hour	Zero
Application upgrades that modify back-end database objects	Hours	Zero

Bronze Requirements: On premises deployment requires Oracle Enterprise Edition and Enterprise Manager Express. Cloud deployments require a minimum of Oracle Enterprise DBaaS (PaaS) and Oracle Database Backup Cloud services. If Oracle Multitenant is used for database consolidation then on-premise deployment also requires a license for Oracle Multitenant and cloud deployment requires a minimum of Oracle High Performance DBaaS (Paas).

Silver: Departmental Databases

The silver reference architecture (Figure 3) is designed for databases that can't afford to wait for a cold restart or a restore from backup should there be an unrecoverable database outage. Silver begins with the same functionality of Bronze and adds capabilities that offer customers a choice of two different patterns for additional HA. The first pattern uses Oracle RAC to enable automatic failover to a second active Oracle instance for HA. The second pattern uses Data Guard database replication with automatic failover to a completely separate synchronized copy of the production database for HA. In both cases the failover instance or replicated copy is local to the primary. Similar to Bronze, a local backup copy is always recommended for fastest recovery from unrecoverable outages. Oracle also recommends maintaining a second copy of backups in a remote data center to protect against site outages.

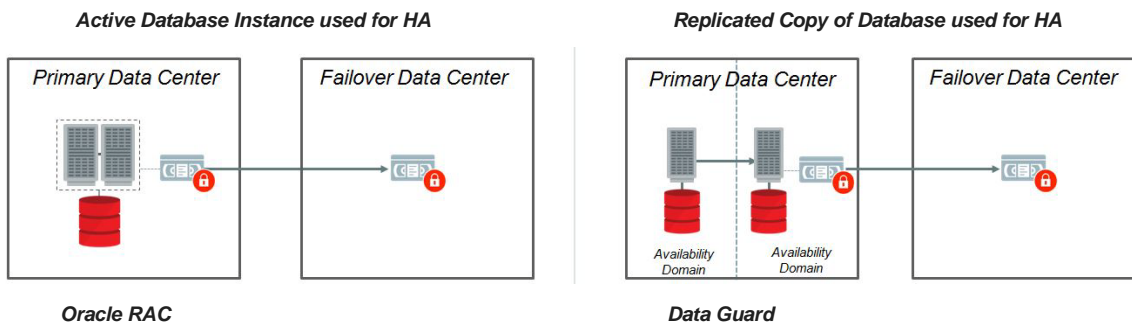


Figure 3: Silver HA Reference Architecture

Silver includes the HA components described in the following sections.

Oracle Real Application Clusters (Oracle RAC)

[Oracle RAC](#) improves application availability within a data center should there be an outage of a database instance or the server on which it runs. Server failover with Oracle RAC is immediate. There is a very brief brownout before service is resumed on surviving instances and users from the down instance are able to reconnect. Downtime is also eliminated for planned maintenance tasks that can be performed in rolling fashion across Oracle RAC nodes. Users complete their work and terminate their sessions on the node where maintenance is to be performed. When they reconnect they are directed to a database instance already running on another node.

A quick review of how Oracle RAC works helps to understand its benefits. There are two aspects: Oracle Database instances and the Oracle Database itself.

- » A database *instance* is defined as a set of server processes and memory structures running on a single node (or server) which make a particular database available to clients. The Oracle RAC configuration in Figure 3 has two database instances.
- » The *database* is a particular set of shared files (data files, index files, control files, and initialization files) that reside on persistent storage, and together can be opened and used to read and write data. The Oracle RAC configuration in Figure 3 has a single Oracle Database.
- » Oracle RAC uses an active-active architecture that enables multiple database instances, each running on different nodes, to simultaneously read and write to the same Oracle database.

The active-active architecture of Oracle RAC provides a number of advantages:

- » Improved high availability: If a server or database instance fails, connections to surviving instances are not affected; connections to the failed instance are quickly failed over to surviving instances that are already running and open on other servers in the cluster.
- » Scalability: Oracle RAC is ideal for high volume applications or consolidated environments where scalability and the ability to dynamically add or reprioritize database workload across more than a single server are required. An individual database may have instances running on one or more nodes of a cluster. Similarly, a database service may be available on one or more database instances. Additional nodes, database instances, and database services can be provisioned online.
- » Reliable performance: Oracle Quality of Service (QoS) can be used to allocate capacity for high priority database services to deliver consistent high performance in database consolidated environments. Capacity can be dynamically shifted between workloads to quickly respond to changing requirements.
- » HA during planned maintenance: High availability is maintained by implementing changes in a rolling manner across Oracle RAC nodes. This includes hardware, OS, or network maintenance that requires a server to be taken offline; software maintenance to patch the Oracle Grid Infrastructure or database; or if a database instance needs to be moved to another server to increase capacity or balance the workload.

Oracle RAC is the recommended MAA best practice for server HA. RAC One Node, however, is an acceptable lower cost alternative to Oracle RAC if a slightly lesser level of HA is acceptable and scalability is not a requirement.

- » RAC One Node is an active-passive failover technology. While it is built upon the same infrastructure as Oracle RAC, RAC One Node has only one database instance open at a time during normal operation. If the server hosting the open instance fails, RAC One Node automatically starts a new database instance on a second node to quickly resume service.
- » RAC One Node provides several advantages over alternative active-passive clustering technologies. In a RAC One Node configuration, Oracle Database HA Services, Grid Infrastructure, and database listeners are always running on the second node. At failover time only the database instance and database services need to start, improving availability by reducing the time required to resume service.
- » RAC One Node also provides the same advantages for planned maintenance as Oracle RAC. RAC One Node allows two active database instances during periods of planned maintenance to allow graceful migration of users from one node to another with zero downtime. Maintenance is performed in a rolling manner across nodes while database services remain available to users at all times.

Oracle Data Guard

[Oracle Data Guard](#) synchronizes one or more physical copies (standby databases) to eliminate single point of failure for a production database (the primary database). Data Guard provides an alternative architecture pattern to Oracle RAC in the Silver reference architecture. Whereas Oracle RAC enables multiple Oracle instances (running on separate compute nodes) to share access to the same Oracle Database, Data Guard maintains synchronization of completely separate Oracle Databases each having their own Oracle instance.

Data Guard provides the following capabilities:

- » Data Guard synchronous replication and Maximum Availability protection mode are used to provide zero data loss protection required for an HA solution. Data Guard transmits changes made on a primary database to a standby database in real-time. Changes are transmitted directly from the log buffer of the primary to minimize propagation delay and overhead, and to completely isolate the standby database from corruptions that can occur in the I/O stack of a production database.
- » The primary database and its standby copy are deployed locally in the same data center but are isolated from each other to the greatest extent possible to avoid single point of failure. A silver reference architecture in the Oracle Cloud for example, deploys the primary and its standby in different Availability Domains. Each Availability Domain has its own power, cooling, network, servers, and storage.
- » In addition to the physical isolation of an Availability Domain, Data Guard performs continuous Oracle validation to ensure that corruption is not propagated from the primary to the standby database. It detects physical and logical intra-block corruptions that can occur independently at either primary or standby databases. It is also unique in enabling run-time detection of silent lost-write corruptions (lost or stray writes that are acknowledged by the I/O subsystem as successful). For more details see [My Oracle Support Note 1302539.1 - Best Practices for Corruption Detection, Prevention, and Automatic Repair](#)
- » Data Guard Fast-Start Failover provides automatic database failover. A Data Guard standby is a running Oracle database, it does not need to be restarted to transition to the primary role. An automatic database failover can [complete in less than 60 seconds](#), even on heavily loaded systems. Fast-Start Failover provides HA by eliminating the delay required for an administrator to be notified and respond to an outage.
- » Data Guard uses role-specific database services and the same [Oracle client notification framework](#) used by Oracle RAC to ensure that applications quickly drop their connections to a failed database and automatically reconnect to the new primary database. Role transitions can also be executed manually using either a command line interface or Oracle Enterprise Manager.
- » Data Guard performs complete, one-way physical replication of an Oracle Database with the following characteristics: high performance, simple to manage, support for all data types, applications, and workloads such as DML, DDL, OLTP, batch processing, data warehouse, and consolidated databases. Data Guard is closely integrated with Oracle RAC, ASM, RMAN and Oracle Flashback technologies.
- » Primary and standby systems are exact physical replicas, enabling backups to be offloaded from the primary to the standby database. A backup taken at the standby can be used to restore either the primary or standby database. This provides administrators with flexible recovery options without burdening production systems with the overhead of performing backups.
- » Standby databases can be used to upgrade to new Oracle Patch Sets (for example, patch release 11.2.0.2 to 11.2.0.4) or new Oracle releases (for example, release 11.2 to 12.1) in a rolling manner. This is done by first upgrading the standby and then switching production to run on the new version. Total downtime is limited to the time required to switch the standby database to the primary production role and transition users to the new primary after maintenance has been completed.

For more background on why Oracle recommends database replication using Data Guard or Active Data Guard and not storage-based remote mirroring solutions (for example, SRDF, Hitachi TrueCopy, and so on) refer to an in-depth discussion in [Oracle Active Data Guard vs. Storage Remote Mirroring](#).

Decision Point – Oracle RAC or Data Guard

Oracle RAC active-active clustering technology provides the best HA for machine or database instance failures and for all planned maintenance that can be performed using RAC-rolling maintenance. Choose the architecture pattern with Oracle RAC if this best addresses your service level requirements.

Silver Requirements (Oracle RAC): On premises deployment using clustering for HA requires Oracle Enterprise Edition, Oracle RAC, Oracle Multitenant (optional for database consolidation) and Enterprise Manager life-cycle management, diagnostic and tuning packs. Cloud deployments require a minimum of Oracle Extreme Performance DBaaS (PaaS) and Oracle Database Backup Cloud services.

Data Guard does not provide the same transparency to outages or scalability for read-write workloads as Oracle RAC, however, a local Data Guard copy does provide HA for a broader set of database outages including data corruptions, human error, cluster outages, power and network outages (if appropriately isolated from the primary system) and database upgrades. Data Guard is an included feature of Oracle Enterprise Edition. While there are no additional licensed products or cloud services beyond that needed for Bronze databases in order to achieve Silver service levels, there is additional configuration required for Data Guard Fast-Start Failover. Choose Data Guard if your service level requirements are better served by addressing a broader range of potential outages and if lower licensing costs and infrastructure requirements are preferred (Data Guard does not require shared storage or inter-node communication required by active-active clustering).

Silver Requirements (Data Guard): On premises deployment using replication for HA requires Oracle Enterprise Edition (includes Data Guard) and Enterprise Manager Express. Cloud deployments require a minimum of Oracle Enterprise DBaaS (PaaS) and Oracle Database Backup Cloud services. Similar to Bronze, if Oracle Multitenant is used for database consolidation then on-premise deployment also requires a license for Oracle Multitenant and cloud deployment requires a minimum of Oracle High Performance DBaaS (Paas).

Silver Summary

Table 2 summarizes RTO and RPO service level requirements for each of the Silver architecture patterns for various unplanned and planned outages. All differences from the bronze tier are bolded in green.

TABLE 2: SILVER RECOVERY TIME (RTO) AND POTENTIAL DATA LOSS (RPO)

Event	Downtime – RTO		Potential Data Loss - RPO	
	Cluster	Replication	Cluster	Replication
Disk failure	Zero	Zero	Zero	Zero
Machine and recoverable database instance failures	Zero to seconds	Seconds	Zero	Zero
Data corruption and unrecoverable database outages, availability domain outages (power, network, etc)	Hours	Seconds	Since last backup	Zero
Site outages	Hours to days	Hours to days	Since last backup	Since last backup
Database reorganization, file move, eligible one-off patches	Zero	Zero	Zero	Zero
Hardware and software maintenance and patching	Zero	Seconds	Zero	Zero
Database upgrades (patch-sets and full releases)	Minutes to Hour	Seconds	Zero	Zero
Application upgrades (modify database objects)	Hours	Hours	Zero	Zero

Service levels that require the best of both patterns – HA and/or scalability of RAC and the broader fault coverage of Data Guard – will utilize the Gold or Platinum reference architectures described in the following sections.

Gold: Business Critical Databases

The Gold reference architecture is for service level requirements that cannot tolerate site failure. It builds upon Silver by using Oracle RAC with a remote Active Data Guard standby database to eliminate single point of failure and provide HA and data protection from all types of unplanned outages including data corruptions, database failures, and site outages. The standby database is also used to reduce planned downtime for database upgrades.

An overview of the Gold tier is shown in Figure 4. As pictured, the Active Data Guard copy in the failover data center is sized for less capacity than required to run production workloads at the same service levels as in the primary data center. In such cases the standby database can be pre-configured for Oracle RAC and scaled up when needed. The standby system must be of equivalent capacity as the primary system only if there is an expectation of an identical service immediately upon failover. The remote standby database eliminates the requirement to replicate backups across sites. Local backups are taken directly from each copy of the database. Since Data Guard and Active Data Guard standby databases are physical copies of the production database, backups from either the production or standby copies can be used to restore the other.




Figure 4: Gold HA Reference Architecture

Active Data Guard is a superset of capabilities that are provided by Data Guard. The Gold reference architecture uses the following advanced features of Active Data Guard:

- » Choice of zero or near-zero data loss protection regardless of the distance separating primary and standby Active Data Guard. Far Sync uses a light-weight forwarding mechanism to enable zero data loss failover even when primary and standby databases are hundreds or thousands of miles apart, without impacting primary database performance. Far Sync is simple to deploy and transparent to operate. Far Sync can also be used in combination with the Oracle Advanced Compression Option to enable off-host transport compression to conserve network bandwidth.
- » Offload of read-only workload to an Active Data Guard standby database open read-only while replication is active. An up-to-date active standby database is ideal for offloading ad-hoc queries and reporting workloads from the production database. This increases ROI in standby systems and improves performance for all workloads by utilizing capacity that would otherwise be idle. It also provides continuous application validation that standby databases are ready to support production workload should an outage occur.
- » Fast incremental backups from the standby database using an RMAN block change tracking file. Fast incremental backups complete up to 20x faster than traditional incremental backups.

- » Automatic repair of block-level corruption caused by intermittent random I/O errors that can occur independently at either primary or standby databases. Active Data Guard retrieves a good copy of the block from the opposite database to perform the repair. No application changes are required and impact of the corruption is transparent to the user.

Gold Summary

RTO and RPO service level requirements addressed by Gold are summarized in Table 4. Areas of improvement beyond that provided by Silver are bolded in green.

TABLE 4: GOLD RECOVERY TIME (RTO) AND POTENTIAL DATA LOSS (RPO)

Event	Downtime- RTO	Potential Data Loss - RPO
Disk failure	Zero	Zero
Machine and recoverable database instance failures	Zero to seconds	Zero
Data corruption and unrecoverable database outages	Seconds to minutes	Zero to seconds*
Site outages	Seconds to minutes	Zero to seconds*
Database reorganization, file move, eligible one-off patches	Zero	Zero
Hardware and software maintenance and patching	Minutes	Zero
Database upgrades (patch-sets and full releases)	Seconds to minutes	Zero
Application upgrades that modify back-end database objects	Hours	Zero

**Zero data loss requires Far Sync, seconds if asynchronous transport is used*

Gold Requirements: On premises deployment requires Oracle Enterprise Edition, Oracle RAC, Active Data Guard, Oracle Multitenant (optional for database consolidation) and Enterprise Manager life-cycle management, diagnostic and tuning packs. Cloud deployment requires a minimum of Oracle Extreme Performance DBaaS (PaaS) or Exadata cloud services. Gold also utilizes Oracle Database Backup Cloud services.

Platinum: Mission Critical Databases

The Platinum reference architecture (Figure 5) builds upon Gold by deploying an extra level of redundancy and several advanced HA capabilities. Platinum is ideal for applications that have extremely low, if any, tolerance for downtime or data loss.

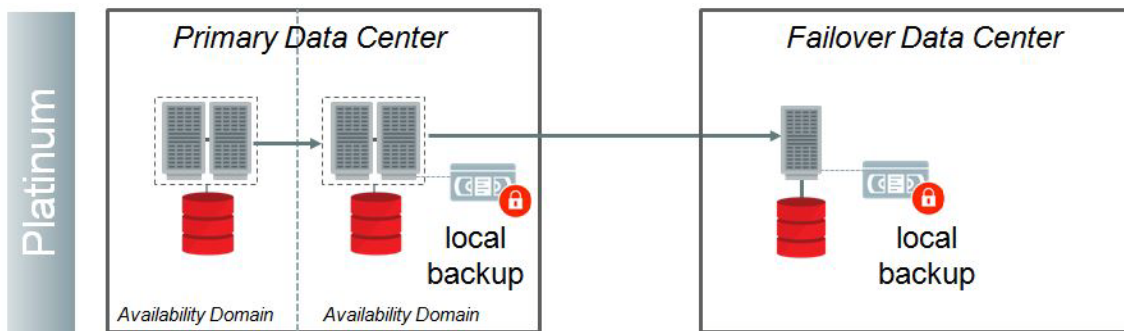


Figure 5 - Platinum HA Reference Architecture

Platinum adds a standby RAC database in a separate Availability Domain (separate power, cooling, network, servers and storage) for fast, automatic, zero data loss failover of existing application tier connections should there be a local outage of the primary RAC database. Platinum masks the impact of outages to applications and users by automatically replaying in-flight transactions following RAC instance failures or failover to the local Active Data Guard copy. Platinum has the same remote standby database included with Gold. It may use either asynchronous replication cascading from the local standby database (as shown) or it can use Far Sync to enable zero data loss failover across the wide area network. Platinum also uses Oracle GoldenGate and Edition-Based Redefinition to enable zero downtime maintenance, migrations, and application upgrades.

Additional details on capabilities of the Platinum reference architecture include:

Application Continuity

[Application Continuity](#) protects applications from database session failures due to instance, server, storage, network, or any other related component, and even complete database failure. Application Continuity re-plays affected “in-flight” requests so that the failure appears to the application as a slightly delayed execution, masking the outage to the user. Transaction replay will occur on surviving nodes following a RAC instance failure, or at the local Active Data Guard standby following an automatic database failover.

Active Data Guard Far Sync

Data Guard and Active Data Guard are the only Oracle-aware replication technologies that offer zero data loss failover for Oracle Database. Network latency between primary and standby sites, however, will affect database throughput and response time when synchronous transport is used. As the distance between sites increases, so will round-trip network latency and its impact on database performance. Because primary and secondary data centers are often separated by long distances, zero data loss failover is impractical to implement for many DR configurations.

Active Data Guard Far Sync eliminates prior limitations by enabling zero data loss failover even when primary and standby databases are hundreds or thousands of miles apart without impacting primary database performance. It achieves this by using a light-weight forwarding mechanism that is both simple to deploy and transparent to Oracle Active Data Guard failover or switchover operations. Far Sync, when used in combination with the Oracle Advanced Compression Option also enables off-host transport compression to conserve network bandwidth.

Oracle GoldenGate

GoldenGate enables logical replication to maintain a synchronized copy (target database) of the production database (source database). The target database contains the same data, but is a different database from the source (for example, backups are not interchangeable). GoldenGate logical replication is a more sophisticated process that has a number of prerequisites that do not apply to Data Guard physical replication. In return for these prerequisites GoldenGate provides unique capabilities to address advanced replication requirements. Refer to [MAA Best Practices: Oracle Active Data Guard and Oracle GoldenGate](#) for additional insights on the tradeoffs of each replication technology and requirements that may favor the use of one versus the other, or the use of both technologies in a complementary manner.

The Platinum reference architecture uses GoldenGate bi-directional replication to implement zero downtime maintenance and migrations. In such a scenario:

- » Maintenance is first implemented at a target database.
- » Source and target are synchronized across versions using GoldenGate replication.
- » Once the new version or platform is synchronized and stable, GoldenGate bi-directional replication enables users to gradually migrate to the new platform with zero downtime. Users naturally terminate their connections to the database operating at the prior version when their work is complete and new connections are directed to the database running the new version. GoldenGate bi-directional replication keeps old and new versions in sync throughout the migration process. This also provides for a fast fall back option should any unanticipated issues arise with the new version as workload grows.

Though the process is not trivial, GoldenGate bi-directional replication may also be used for application upgrades that modify back-end database objects. Developer-level knowledge of the database objects being modified or added by the new release is required in order to enable GoldenGate to replicate across versions. Implementing cross-version mapping using GoldenGate replication is required for each new release of the application.

Bi-directional replication can also be used to increase availability service levels when a continuous read-write connection to multiple copies of the same data is required. It is important to note that bi-directional replication is not application transparent. It requires conflict detection and resolution when changes are made to the same record at the same time in multiple databases. It also requires careful consideration of the impact of different failure states and replication lag.

It is also important to note that GoldenGate replication is an asynchronous process that is not able to provide the same zero data loss protection as Data Guard and Active Data Guard. This is one reason why Data Guard and Active Data Guard are used by the Silver, Gold and Platinum reference architectures to provide HA during unplanned outages since there is an assumption that an HA event should result in zero data loss. GoldenGate replication may be used in place of Data Guard or Active Data Guard where zero data loss protection is not a concern.

There is no concern for data loss during planned maintenance when GoldenGate is used as long as the production copy of the database is protected by a Data Guard standby.

Edition-Based Redefinition

[Edition-Based Redefinition](#) (EBR) enables online application upgrades that require changes to database objects that would otherwise require the database to be offline. EBR enables all changes to be implemented while the previous version of the application and the database remain online. When the upgrade installation is complete, the pre-upgrade application and the post-upgrade application can be used at the same time against the same copy of the

Oracle Database. Existing sessions can continue to use the pre-upgrade application until their users decide to end them, and new sessions can use the post-upgrade application. When there are no longer any sessions using the pre-upgrade version of the application it can be retired.

EBR enables online application upgrades in the following manner:

- » 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 editing view exposes a different projection of a table into each edition to allow each to see just its own columns.
- » A cross-edition trigger propagates data changes made by the old edition into the new edition's columns, and vice-versa.

Similar to GoldenGate zero downtime application upgrades, the use of EBR requires deep knowledge of the application and a non-trivial effort on the part of the developer to incorporate it. Unlike GoldenGate, there is a one-time investment to utilize EBR. From that point forward minimal effort is required to use EBR for subsequent new releases of the application. EBR has proven that it can be implemented even for the most complex applications, for example, Oracle E-Business Suite 12.2 uses EBR for online patching. EBR is a feature included with Oracle Database at zero additional cost.

Platinum Summary

RTO and RPO service level requirements addressed by the Platinum reference architecture are summarized in Table 5.

TABLE 5: PLATINUM RECOVERY TIME (RTO) AND POTENTIAL DATA LOSS (RPO)

Event	Downtime - RPO	Potential Data Loss – RTO
Disk failure	Zero	Zero
Machine and recoverable database instance failures	Zero to seconds	Zero
Data corruption and unrecoverable database outages, availability domain outages (power, network, etc)	Seconds	Zero
Site outages	Seconds	Zero
Database reorganization, file move, eligible one-off patches	Zero	Zero
Hardware and software maintenance and patching	Zero	Zero
Database upgrades (patch-sets and full releases)	Zero to seconds	Zero
Application upgrades that modify back-end database objects	Zero to seconds	Zero

Platinum Requirements: On premises deployment requires Oracle Enterprise Edition, Oracle RAC, Active Data Guard, Oracle GoldenGate, Oracle Multitenant (optional for database consolidation) and Enterprise Manager life-cycle management, diagnostic and tuning packs. Cloud deployments require a minimum of Oracle Extreme Performance DBaaS (PaaS) or Exadata cloud services and Oracle GoldenGate cloud service. Platinum also utilizes Oracle Database Backup Cloud services.

Corruption Prevention, Detection, and Auto-Repair

Common across all reference architectures is the ability of the Oracle Database to prevent, detect, and automatically repair data corruption. Each check is unique to Oracle Database using specific knowledge of Oracle data block and redo structures to increase HA and protect data by preventing the spread of data corruptions that can occur in memory or due to faults in the I/O stack. These capabilities are summarized in Table 6. The column 'Type' in Table 6 indicates when validations for physical and logical corruption are performed.

- » Manual checks are initiated by the administrator or at regular intervals by a scheduled job that performs periodic checks.
- » Runtime checks are automatically executed on a continuous basis by background processes while the database is open.
- » Background checks are run on a regularly scheduled interval, but only during periods when resources would otherwise be idle.

TABLE 6: CORRUPTION PREVENTION, DETECTION, AND AUTO-REPAIR

Type	Reference Architecture (s)	Capability	Physical Block Corruption	Logical Block Corruption
Manual	All	Dbverify, Analyze	Physical block checks	Logical checks for intra-block and inter-object consistency
Manual	All	RMAN	Physical block checks during backup and restore	Intra-block logical checks
Runtime	Silver – Pattern 2, Gold and Platinum	Data Guard, Active Data Guard	Physical block checking at standby Strong isolation between primary and standby eliminates single point of failure Automatic database failover	Detect lost-write corruption, auto shutdown and failover Intra-block logical checks at standby
Runtime	Gold and Platinum	Active Data Guard	Automatic repair of physical corruptions	
Runtime	All	Oracle block checksum and block checking	In-memory block and redo checksum	In-memory intra-block logical checks
Runtime	All	ASM	Automatic corruption detection and repair using local extent pairs	
Runtime	All – Exadata only	Exadata	HARD checks on write	HARD checks on write
Background	All – Exadata only	Exadata	Automatic HARD Disk Scrub and Repair	

Note that HARD validation and the Automatic Hard Disk Scrub and Repair are unique to Exadata storage. HARD validation ensures that Oracle Database does not write physically corrupt blocks to disk. Automatic Hard Disk Scrub and Repair inspects and repairs hard disks with damaged or worn out disk sectors (cluster of storage) or other physical or logical defects periodically when there are idle resources. Exadata sends a request to ASM to repair the bad sectors by reading the data from another mirror copy. By default the hard disk scrub runs every two weeks.

Hybrid Cloud

Each of the MAA Reference Architectures described above apply to both on-premises and cloud deployment. There is one more category of deployment often referred to as hybrid cloud, where the cloud is used for backup and for disaster recovery for on-premises production systems.

- » The Oracle Database Backup Service provides cost-effective, secure, cloud-based backups for on-premises databases. Cloud backups can be also used for disaster recovery when service level requirements can tolerate downtime (the time to provision new database instances, restore backups and validate that the system is ready for production) and data loss (data generated since last backup is not protected).
- » Oracle Data Guard and Active Data Guard can be used for Disaster Recovery between an on-premise production database and the Oracle Cloud (hybrid cloud). This is supported following the procedures documented in [Disaster Recovery to the Oracle Cloud - Production on Premises, DR in the Cloud](#). Data Guard and Active Data Guard are appropriate for service level requirements that have little tolerance for downtime or data loss. Each enables fast failover to a synchronized copy of the production database. Each uses asynchronous database replication to provide near-zero data loss protection across any distance or synchronous replication for zero data loss protection for local DR (usually within 100 miles of separation). Alternatively, Active Data Guard Far Sync (12c) enables zero data loss failover across any distance for databases.
- » The Oracle GoldenGate Cloud Service is a new offering planned for first half of calendar 2016 that can provide DR using logical replication as an alternative to Data Guard and Active Data Guard. Logical replication offers much increased flexibility for replicating between a source database and its DR copies (different Oracle releases, hardware platforms and architectures) and thus can address a greater range of requirements. Note that GoldenGate is an asynchronous replication process, it does not offer the same zero data loss options provided by Data Guard and Active Data Guard.

Oracle MAA Cloud Roadmap

Capabilities required to fully implement each of the reference architectures described in this paper for on-premises systems are all available today. In most cases functionality is supported in both Oracle Database 11g and 12c. Oracle 12c is required in several cases, such as with Active Data Guard Far Sync, Application Continuity, and Edition-Based Redefinition.

There are, however, several aspects of Oracle cloud services required to fully deploy all reference architectures that are not yet generally available. These services are planned for release during calendar year 2016. Table 6 summarizes the roadmap for these cloud services and the reference architectures where they apply.

TABLE 6: ORACLE MAA CLOUD ROADMAP – PLANNED FOR CALENDAR YEAR 2016

Planned Functionality Required by Reference Architecture	Bronze	Silver Cluster	Silver Replication	Gold	Platinum	Hybrid
Backups automatically replicated to remote site	H1	H1	H1	-	-	-
Automated deployment of Data Guard/Active Data Guard for HA across cloud Availability Domains (PaaS)	-	-	H1	-	H1	-
Automated deployment of Data Guard/Active Data Guard between Cloud data centers for DR (PaaS)	-	-	-	H2	H2	-
GoldenGate bi-directional replication within a cloud data center for zero downtime maintenance (PaaS)	-	-	-	-	H2	-
Automated deployment of Data Guard/Active Data Guard for hybrid Cloud DR (replaces current manual blueprint)	-	-	-	-	-	H1
GoldenGate uni-directional replication between on-premises and cloud (PaaS) as an alternative to Active Data Guard for DR if no zero data loss requirement	-	-	-	-	-	H1

Note: 1 is first half of calendar 2016, H2 is second half of calendar 2016

Conclusion



Enterprises need solutions that address the full continuum of requirements for data protection and availability. Oracle MAA best practices define four HA reference architectures: BRONZE, SILVER, GOLD and PLATINUM. Each MAA reference architecture uses the optimal set of Oracle HA capabilities based upon a common platform to reliably achieve a given service level requirement whether databases are deployed on-premises, public cloud, or a hybrid cloud.



Oracle Corporation, World Headquarters
500 Oracle Parkway
Redwood Shores, CA 94065, USA

Worldwide Inquiries
Phone: +1.650.506.7000
Fax: +1.650.506.7200

CONNECT WITH US

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

Integrated Cloud Applications & Platform Services

Copyright © 2016, Oracle and/or its affiliates. All rights reserved. This document is provided *for* information purposes only, and the contents hereof are subject to change without notice. This document is not warranted to be error-free, nor subject to any other warranties or conditions, whether expressed orally or implied in law, including implied warranties and conditions of merchantability or fitness for a particular purpose. We specifically disclaim any liability with respect to this document, and no contractual obligations are formed either directly or indirectly by this document. This document may not be reproduced or transmitted in any form or by any means, electronic or mechanical, for any purpose, without our prior written permission.

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

Intel and Intel Xeon are trademarks or registered trademarks of Intel Corporation. All SPARC trademarks are used under license and are trademarks or registered trademarks of SPARC International, Inc. AMD, Opteron, the AMD logo, and the AMD Opteron logo are trademarks or registered trademarks of Advanced Micro Devices. UNIX is a registered trademark of The Open Group. 0116