

Oracle Maximum
Availability Architecture

Maximize Availability

with Oracle Database 12c Release 2

ORACLE WHITE PAPER | NOVEMBER 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.

Disclaimer	0
Introduction	3
The High Availability Challenge	4
Oracle Database High Availability	4
Innovation in Oracle Database 12c	4
Oracle Database High Availability Design Principles	5
Oracle Maximum Availability Architecture	7
Addressing Unplanned Downtime	8
Server: Oracle Real Application Clusters	8
Transparent Failover: Application Continuity	9
Storage: Automatic Storage Management (ASM)	9
Data Availability and Corruption Protection	10
Backup and Recovery – Oracle Recovery Manager	11
Backup to Tape – Oracle Secure Backup (OSB)	13
Real-Time Data Protection – Zero Data Loss Recovery Appliance	15
Recovery from Logical Corruption: Oracle Flashback Technology	16
Real-time Data Protection and Availability – Oracle Data Guard	18
High Availability with Zero Data Loss across Any Distance: Active Data Guard	19
Active-Active HA: GoldenGate	22
Complete Site Failover: Oracle Site Guard	23
Addressing Planned Downtime	23
Online System Reconfiguration	23
Online Data and Application Change	24
Online Application Upgrades: Edition-Based Redefinition	25



Online Patching	25
Rolling Patch Upgrades using Oracle RAC	25
Data Guard Standby-First Patch Assurance	25
Database Rolling Upgrades using Data Guard	26
Database Rolling Upgrades using Active Data Guard	26
Platform Migration, Systems Maintenance, Data Center Moves	26
Zero Downtime Maintenance using Oracle GoldenGate	27
Managing Oracle Database High Availability Solutions	27
Global Data Services	27
Conclusion	29
Appendix: New High Availability Features in Oracle Database 12c	30
Appendix: New High Availability Features in Oracle Database 12c Release 2 on Oracle Cloud	31



Introduction

Oracle Database 12c Release 2, the latest generation of the world's most popular database is now available in the Oracle Cloud.

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 a High Availability infrastructure 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 loose 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 discusses Oracle Database 12c features, including Application Continuity, Global Data Services, and Active Data Guard Far Sync, which improve application recovery, support global database services, and extend zero-data-loss protection to a global scale, respectively. It describes enhancements to Oracle Database 12c Release 2 (Oracle Database 12.2) performance, functionality, and ease-of-use to existing High Availability features including Real Application Clusters, Automatic Storage Management, Recovery Manager, Data Guard and Active Data Guard, Oracle Secure Backup, and Edition-Based Redefinition. This paper also describes the Zero Data Loss Recovery Appliance, which revolutionizes data protection, and recovery for Oracle customers.

The High Availability Challenge

Designing, implementing, and managing a High Availability 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?

Oracle believes that High Availability encompasses a number of important aspects in addition to the main goal of preventing downtime. Key dimensions of a comprehensive High Availability architecture includes:

- » 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.

Successful High Availability begins with understanding 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 High Availability architecture.

Oracle Database High Availability


Oracle has been hard at work for over three decades helping IT solve High Availability challenges by designing comprehensive High Availability capabilities integrated into the database. This innovation results in High Availability solutions that give true competitive advantages to enterprises, by helping them achieve service level objectives for high availability 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 High Availability 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.

Innovation in Oracle Database 12c

The Multitenant architecture of Oracle Database 12c, delivered groundbreaking technology for database consolidation and cloud computing. It helps drive down IT costs by enabling a true 'manage-as-one' architecture for consolidation and virtualization of the database tier. The Multitenant architecture also makes extreme high availability a fundamental requirement when database consolidation is applied to business-critical applications. By definition, database consolidation is an exercise of 'putting all eggs in one basket.' The more successful you are at driving down cost through consolidation, the more eggs are in a single basket, and the greater is the operational and financial impact to the business should an outage occur.

The High Availability (HA) capabilities in Oracle Database 12c are designed to provide the extreme level of availability required for consolidating databases. This includes support for multitenant architecture across all



Oracle High Availability features, new levels of redundancy, transparent failover of in-flight transactions, zero-data loss disaster protection at any geographic distance. The Oracle Multitenant architecture represents the next-generation in database technology, and long-standing and time-proven Oracle High Availability design principles are ready from day one to provide the extreme availability required by consolidated environments.

New in Oracle Database 12.2 on Oracle Cloud is Oracle Sharding, a scalability and availability feature for custom-designed OLTP applications that enables distribution and replication of data across a pool of discrete Oracle Databases. This pool of databases is presented to the application as a single logical database. Applications elastically scale (data, transactions and users) to any level, on any platform, simply by adding additional databases (shards) to the pool.

Oracle Sharding provides superior run-time performance and simpler life-cycle management compared to home-grown deployments that use a similar approach to scalability. It also provides the advantages of an enterprise DBMS, including: relational schema, SQL, and other programmatic interfaces, support for complex data types, online schema changes, multi-core scalability, advanced security, compression, high-availability, ACID properties, consistent reads, developer agility with JSON, and much more.

Oracle Database High Availability Design Principles

Oracle Database High Availability relies on a set of tightly integrated features built within the database kernel. Oracle's vision for High Availability is guided by three principles, described next.

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. 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 does fine-grained, efficient recovery of individual blocks instead of entire data files. Another unique-to-Oracle example of data protection is the ability of Flashback technologies to undo database changes at a level of granularity appropriate to the scope of the error, be it the entire database, or a table, or an individual transaction, without requiring a full database restore.

Deliver application-integrated high availability

Providing High Availability and data protection using cold failover clusters or at the raw bits level as done by storage-centric 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. Oracle solutions enable granular recovery and thus very efficient, with no disruption to the availability of applications using unaffected portions of the database. Oracle also allows making structural changes to a table while others are accessing and updating it, via the Online Redefinition feature. Application Continuity, a new capability in Oracle Database 12c, masks many outages from end users and applications by replaying the session after a server or site failover has occurred, transparent to the application.

Oracle 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. Data Guard 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.

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, 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. Oracle's active High Availability architecture provides high ROI and the same time it minimizes risk of failure. There is never a question of if it will start and how long it will take after a failure occurs to resume service: all Oracle High Availability components are already started, already performing useful work, enabling continuous user validation that they are ready for prime-time.

Oracle Maximum Availability Architecture

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

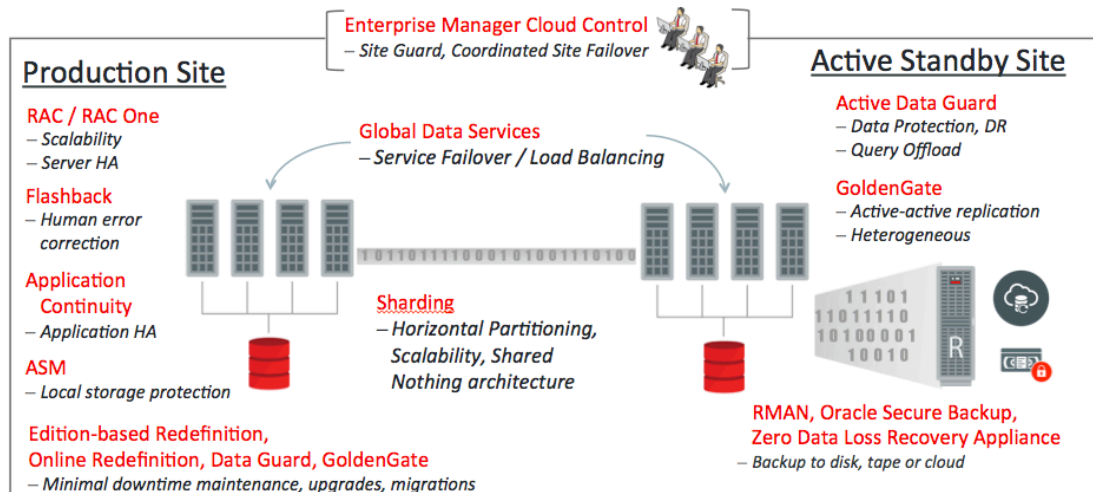


Figure 1: Oracle's High Availability Technologies and the Oracle Maximum Availability Architecture

MAA best practices are created and maintained by a team of Oracle developers that continually validate the integrated use of Oracle Database High Availability features. Real-world customer experience is also integrated into the validation performed by the MAA team, spreading lessons learned to other customers.

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 MAA Reference Architectures – The Foundation for Database as a Service*.¹

The remainder of this document examines High Availability capabilities for Oracle Database 12c Release 2 in greater depths.

¹ <http://www.oracle.com/technetwork/database/availability/maa-reference-architectures-2244929.pdf>

Addressing Unplanned Downtime

Hardware faults, which cause server failure, are essentially unpredictable, and result in application downtime when they eventually occur. Likewise, a range of data availability failures, including storage corruption, site outage and human error, also cause unplanned downtime. In this section we discuss how Oracle's High Availability solutions address these fundamental categories of failures in order to prevent and mitigate unplanned downtime.

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 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 computer failures. Since the nodes run independently, the failure of one or more does not affect other nodes. This architecture also allows a group of nodes to be transparently put online or taken offline, while the rest of the system continues to provide database services.
- » Flexibility and cost effectiveness in capacity planning, so 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, reducing costs by avoiding the more expensive and disruptive upgrade path of replacing an existing monolithic system with a larger one. Oracle RAC supports enables near-linear scaling without any changes to your application.

Building on those capabilities, new in Oracle Database 12.2 on Oracle Cloud, Oracle RAC has improved scalability, availability and performance. They are achieved using service-oriented buffer cache access, pluggable database and service isolations and near-zero downtime reconfigurations. Massive parallel query RACs which can be obtained by overlaying Hadoop Cluster (HDFS) with an Oracle Flex Cluster to access data in Hadoop via SQL and perform cross-data analysis. Moreover, you can use read-only workloads on read-mostly leaf node instances (RAC Reader Nodes) that can be scaled to hundreds of nodes with no delay in accessing updated data and without impacting OLTP performance. Node Weighting is another new feature in Oracle Database 12.2 on Oracle Cloud that considers the workload hosted in the cluster during fencing operation by letting the majority of the work to survive in the event everything else is equal. Autonomous Health Framework continuously monitors and analysis and is hosted on Domain Service Cluster (DSC) instead of production cluster. DSC is also used for Rapid Home Provisioning.

Oracle RAC also supports the new multitenant architecture, and in addition to providing server HA, Oracle RAC software stack² is also the ideal shared infrastructure for database consolidation.

For more information see Real Application Clusters resources on OTN (oracle.com/goto/rac).

² Oracle Grid Infrastructure including Oracle ASM / ACFS and Oracle Clusterware, and the Oracle Database with the Oracle Real Application Clusters option, constitute the Oracle Database RAC software stack.

Transparent Failover: Application Continuity

It is complex for application development 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 12c introduced Application Continuity, a capability in Oracle Database 12c masks outages by recovering the database session following unplanned outages. Application Continuity performs this recovery beneath the application so that the outage appears to the application as a slightly delayed execution.

In Oracle Database 12.2 on Oracle Cloud, Application Continuity supports 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 makes it easy to migrate existing connections to another database instance even if Oracle Connection Pools are not used.

Storage: Automatic Storage Management (ASM)

Oracle Automatic Storage Management (ASM) is the underlying (clustered) volume manager technology used by the Oracle Database and Oracle ASM Cluster File System (ACFS) that enables storing and managing any type of data on shared storage. Through its low cost, ease of administration and high performance, ASM is the storage technology of choice for Oracle Databases.

For performance and high availability, ASM stripes and mirrors everything. Intelligent mirroring capabilities allow administrators to define 2- or 3-way mirrors to protect data. When a read operation identifies that a corrupt block exists on disk, ASM automatically relocates the valid block from the mirrored copy to an uncorrupted portion of the disk. Administrators can also use the ASMCMD utility to manually relocate specific blocks. When disk failures occur, system downtime is avoided by using the data available on the mirrored disks. If the failed disk is permanently removed from ASM, the underlying data is striped or rebalanced across the remaining disks for continued high performance.

Flex ASM, a feature of Oracle Database 12c increases database (instance) availability by enabling inter-node storage failover and reducing ASM-related resource consumption by up to 60%. Flex ASM facilitates cluster based database consolidation, as it ensures that database instances running on a particular server will continue to operate, should the ASM instance on that server fail.

ASM disk scrubbing, a feature of Oracle Database 12c, 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.

New in Oracle Database 12.2 on Oracle Cloud, ASM introduces a new style of storage management called "Database Oriented Storage Management" using a new type of diskgroup called ASM "Flex Diskgroup". These diskgroups enable a new concept of "File Group" which is a logical container of files in a diskgroup that belong to an individual database or PDB. This allows operations to be targeted against all the files belong to a database collectively. Flex Diskgroups enable quota management for databases, redundancy change and split mirror capabilities for devops. New diskgroup type "Extended Diskgroup" can be used to support extended RAC with up to 3 sites.

For more details on ASM, visit oracle.com/goto/asm.



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: storage subsystem failure, site failure, human error, and corruption. Their multifaceted causes often make data failures difficult to identify and diagnose. This and subsequent sections examine the High Availability 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 good risk management must include measures to prevent and remediate human error. For example, an incorrect WHERE clause may cause UPDATE to affect many more rows than intended. Oracle Database 12c 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.

A good way to prevent costly human errors is to restrict users' access scope to just data and services they need. The Oracle Database provides security tools to flexibly control user access by authenticating users and allowing administrators to grant users only those privileges required to perform their duties.

Previously, a backup administrator, for example, would be granted broad SYSDBA privileges, with the consequent security exposure. Oracle Database 12c introduced privileges such as SYSDG and SYSBACKUP to support separation of duties and finer scope definition for database administration. SYSDG is for Data Guard activities such as configuration, monitoring, and effecting role change. SYSBACKUP is for Recovery Manager (RMAN) activities such as backing up or restoring a database.

We discuss in sections below Backup and Flashback technologies to recover from human errors.

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. We discuss Oracle's comprehensive set of solutions to protect data from corruption in the next pages.

Detect and Prevent Physical and Logical Intra-Block Corruption

For comprehensive corruption protection Oracle MAA recommends deploying Data Guard combined with appropriate parameter settings that enable key corruption checks, including block header checks, full-block checksums, and lost-write verification (physical and logical block checking). Active Data Guard also provides automatic repair of physical block corruption detected on a primary database using a good copy from the active standby, and vice versa. These settings will affect performance and therefore need to be tested before

introducing them to production. See My Oracle Support Note 1302539.1 for more detail on each parameter.³ See the MAA whitepaper, *Preventing, Detecting, and Repairing Block Corruptions* for a complete discussion of this topic.⁴

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 backup and restore, 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.

RMAN offers a choice of compression levels: BASIC is included in the Oracle Database Enterprise Edition while LOW, MEDIUM and HIGH levels are available as part of the Oracle Advanced Compression Option (ACO). The compression ratio and CPU usage vary from highest to lowest in the following order: HIGH, BASIC, MEDIUM and LOW. Therefore, the HIGH compression level will achieve the best compression ratio while also requiring the most CPU overhead.

RMAN Active Duplicate functionality creates a clone or physical standby database over the network without the use of backups. Data file copies are written directly to the destination database. In Oracle Database 12c, the workload is moved to the destination server via auxiliary channels, relieving resource bottlenecks on the source (usually, production) database server. New for Oracle Database 12c, Active Duplicate Cloning can use RMAN compression and multi-section capabilities to further increase performance. Unused block compression happens automatically. Administrators can, as before, also configure RMAN to apply binary compression, if network traffic is a bottleneck. Capability to perform Table level recovery using simple syntax RECOVER TABLE was introduced in Oracle Database 12c.

New in Oracle Database 12.2 on Oracle Cloud, Table Recovery has been improved to allow recovery across schemas for flexibility and to facilitate more use cases. Moreover, disk space verification is performed before the process is started to avoid mid-way interruption. Space efficient level 0 (full) backups for Sparse databases (such as PDB Snapshot etc.) is introduced in Oracle Database 12.2 which now allows to perform backup, recovery and duplication of delta change of Sparse database instead of the whole base (backing) files using SPARSE syntax.

³ MOS Note 1302539.1 explains the protection/performance tradeoffs of these parameters.

⁴ <http://www.oracle.com/technetwork/database/availability/corruption-bestpractices-12c-2141348.pdf>

Cross-platform Backup and Restore

New with Oracle Database 12c, RMAN Cross-platform functionality enables backup and restore across different platforms,⁵ for efficient tablespace and database migration. On the source platform, BACKUP creates backup sets of user tablespaces, including Data Pump metadata dump file, in read-only mode. RESTORE on the destination platform automatically performs data file endian conversion and plugs-in tablespaces. To minimize read-only impact, we recommend taking incremental backups that are then converted and applied to restored data files. Only the final incremental backup need be taken while tablespaces are in read-only mode.

Oracle Database 12.2 for cloud, now supports cross-platform import of PDB from another CDB and also from non-CDB. Additional cross-platform enhancements include support for encrypted tablespace migrations, over the network restore for efficient data transfers.

RMAN support for Oracle Multitenant

RMAN also 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). RMAN commands can also be applied to individual PDBs, including full backup and restore, using the keyword PLUGGABLE. For example, the following simple RMAN script can be run for Point-in-time Recovery of a pluggable database:

```
RMAN> RUN
  {SET UNTIL TIME 'SYSDATE-3';
  RESTORE PLUGGABLE DATABASE <PDB>;
  RECOVER PLUGGABLE DATABASE <PDB>;
  ALTER PLUGGABLE DATABASE <PDB> OPEN RESETLOGS; }
```

RMAN also supports efficient cloning of the container database including all or some (user-specified) pluggable databases.

Other RMAN Enhancements Available with Oracle Database 12c

RMAN can now recover individual database tables from backup, via a simple RECOVER TABLE command. This recovers one or more tables (the most recent or an older version) from an RMAN backup. Tables can be recovered in-place or to a different tablespace. Optionally, RMAN can create a Data Pump dump file of the table(s). This functionality replaces an error-prone manual process and improves the Recovery Time Objective (RTO). It extends the range of recovery where Flashback is not applicable, for example when a dropped table has been purged out of the Recycle Bin, or when the desired point to recover is outside the window given by the UNDO_RETENTION parameter.

Other RMAN enhancements in Oracle Database 12c and Oracle Database 12.2 on Oracle Cloud to provide increased performance and ease-of-use include:

- » RMAN support for multi-section backup of image copies and incremental backups.

⁵ Cross-platform incremental backups are supported for Linux on earlier releases as described in [MOS Note 1389592.1](#). Traditionally, moving a database across platforms required either import/export or cross-platform transportable tablespaces procedures, seriously affecting application availability.

- » 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.
- » DUPLICATE support for auto-login wallet based encrypted backups to allow interruption free duplication
- » 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.

For more information see Oracle's RMAN resources on OTN (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 filesystem 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.

In addition to a location, the FRA is also assigned a quota. If multiple databases are sharing a single FRA, each will have its own quota and the size of the FRA will be the sum of database quotas. When new backups are created in the FRA and there is insufficient space (per the assigned quota) to hold them, backups and archived logs that are not needed to satisfy the RMAN retention policy (or that have already been backed up to tape), are deleted automatically to reclaim space. The FRA also notifies the administrator (via the alert log) when disk space used is nearing its quota and no additional files can be deleted. The administrator can add more disk space, back up files to tape to free up disk space for the FRA, or change the retention policy.

Data Recovery Advisor

Many data outages can be mitigated based on accurate analysis of errors and trace files that are present prior to an outage. The Data Recovery Advisor (DRA) can proactively run database health checks that verify physical integrity, identify possible precursors to a database outage, and alert the administrator. The administrator can get recovery advice and perform preventive actions to fix the problem before it results in system downtime. When critical business data are damaged, the DRA assists the database administrator to ensure a safe and fast recovery under pressure, by quickly and thoroughly evaluating recovery and repair options. As it is tightly integrated with other Oracle High Availability features such as Data Guard and RMAN, the DRA is able to identify which recovery options are feasible given the specific conditions. These options are presented to the administrator, ranked from least to most potential data loss. The DRA can also automatically implement the best recovery option(s) or just serve as a guide for manual recovery by the administrator.

Backup to Tape – Oracle Secure Backup (OSB)

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

- » Recovery Manager (RMAN) integration, supporting versions Oracle Database 10g to Oracle Database 12c, that can increase backup performance by 25 – 40% over comparable products.
- » File system data protection for UNIX, Windows, and Linux servers, as well as Network Attached Storage (NAS) protection via the Network Data Management Protocol (NDMP).
- » Policy-based fine-grained control, including backup encryption and key management, tape duplication, and rotating tapes between different locations (vaulting).

- » The Oracle Secure Backup environment may be managed using the command line, the OSB web tool or Oracle Enterprise Manager.

The following enhancements in the latest release, OSB 12.1, are ideal for Exadata environments:

- Automatic device discovery
 - Oracle Secure Backup simplifies the process of configuring devices by automatically discovering devices that are attached to media servers and then configuring them as devices in the administrative domain. See "[Overview of Automatic Device Discovery](#)".
- Backup and restore operations using disk
 - Oracle Secure Backup provides integrated disk backup and restore capabilities. You can use disk as a backup target or a restore source for both file-system and Oracle Database backup and restore operations. Storing backups on disk provides improved performance for file-system backup and restore operations. Tape device utilization can be improved by first creating backups on disk and then moving them to tape.
 - You can create multiple instances of a particular backup. Oracle Secure Backup distinguishes the metadata that defines the backup from the actual backup data. A backup image refers to the metadata that includes device-independent information about the backup. The actual data that is backed up is called a backup image instance. A single backup image can have multiple backup image instances. See "[About Disk Pools](#)".
- Domain SSL Certificate Renewal and Management
 - Oracle Secure Backup provides a method for renewing domain SSL certificates and specifying a new lifetime of the domains SSL certificates. The SSL certificates can be renewed automatically or manually for increased security. See [Managing Certificates with obcm](#).

Oracle Secure Backup Cloud Module

Oracle Secure Backup Cloud Module is the SBT library for RMAN to backup to Amazon S3 object storage. Using this module, you can send backups directly to Amazon S3. The OSB Cloud module can back up all supported versions of Oracle Database.⁶ Administrators can continue to use their existing backup tools – Enterprise Manager, RMAN scripts, etc. – to perform cloud backups. See OSB resources on OTN for more information (oracle.com/goto/osb).

Oracle Database Backup Cloud Service

Oracle Database Backup Cloud Service is a low cost offsite storage backup solution for storing backups offsite. It is an Oracle Cloud's Platform as a Service (PaaS) offering which enables Oracle Databases that are deployed on-premises or Oracle cloud to backup the data using RMAN by encrypting and securely transmitting 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

⁶ The OSB Cloud module uses the RMAN media management interface, which seamlessly integrates external backup libraries with RMAN for all database backup and recovery operations.

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 Enterprise Manager 13c, RMAN CLI or 3rd party software like Cloudberry to perform backup & recovery management. For more information, refer to cloud.oracle.com/database_backup.

Real-Time Data Protection – Zero Data Loss Recovery Appliance

The Zero Data Loss Recovery Appliance is an innovative data protection solution that is completely integrated with RMAN and the Oracle Database.⁷ 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 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 unmatched 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.

⁷ <http://www.oracle.com/recoveryappliance>

- » 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.
- » Cloud-Scale Data Protection. The Recovery Appliance introduces the concept of a protection policy, which defines recovery window goals that are enforced on a per-database basis on the appliance and tape, if present. Using protection policies, databases across the enterprise can be easily grouped by recovery service tier.
- » 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.
- » Modern Cloud Scale Architecture. The Recovery Appliance is built on a massively scalable, highly redundant, fault tolerant, storage architecture. As more and more databases within an enterprise are moved to the recovery appliance, compute and storage servers are easily added to provide a simple, no-downtime, scale-out data protection cloud to support ongoing business growth. The base configuration consists of 2 compute servers and 3 storage servers providing up to 94 TB of usable capacity for incoming backups. Up to 18 storage servers can be added to a rack to increase usable capacity to a maximum of 580TB as needs grow. When the first rack is full, additional racks can then be connected via InfiniBand. Up to 18 fully configured racks can be connected together providing 10+ PB of usable capacity. With 10% change rate assumed, this equates to 100s of PB of effective backup can be stored in this single logical Recovery Appliance.

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

Recovery from Logical Corruption: Oracle Flashback Technology


Human errors happen. Oracle Database Flashback Technologies are 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 it. With Flashback, the time required to recover from an error is depends on the work done since the error was made. Recovery time does not depend on the database size, a capability that becomes a necessity as database sizes continue to grow, and that is unique to the Oracle Database. 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. Flashback also provides fine-grained analysis and repair for localized damage, e.g., when the wrong customer order is deleted. Flashback can also repair more widespread damage while still avoiding long downtimes, e.g., all of yesterday's customer orders have been deleted.

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...
```



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...
```

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
```

shows changes made by this transaction and 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.

Flashback Transaction

Often, data failures take time to be identified, and 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 an Enterprise Manager wizard to identify and flash back the necessary transactions.

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 takes time proportional to the (ever growing) size of the database – hours or even days.

In contrast, Flashback Database, using Oracle-optimized flashback logs, can quickly restore an entire database to a specific point-in-time. Flashback Database is fast because it restores changed blocks only. Flashback Database can restore a whole database in minutes via a simple command like:

```
FLASHBACK DATABASE TO TIMESTAMP time
```



No complicated recovery procedures are required and there is no need to restore backups. Flashback Database drastically reduces the downtime required for database point-in-time recovery. Also, Flashback Database integrates with Data Guard to support Data Guard's Snapshot Standby and the reinstatement of the previous primary after a failover (see also the Data Guard section).

In Oracle Database 12c, Flashback Database operation can be performed only at the CDB (container) level - which rewinds all the PDBs. In Oracle Database 12.2 on Oracle Cloud, Flashback Database can now be performed at the PDB level. For example, the following command rewinds only the PDB pdb1 and other PDBs are not impacted.

```
FLASHBACK PLUGGABLE DATABASE pdb1 TO RESTORE POINT before_update;
```

Both Normal Restore Points and Guaranteed Restore Points are supported.

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
```

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).

For more details, see Flashback resources OTN (oracle.com/goto/flashback).

Real-time Data Protection and Availability – Oracle 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 downtimes 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 a synchronized physical replica (standby) of a production database (primary). Administrators can choose either manual or automatic failover to a standby database 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.

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. Fast Sync and Active Data Guard Far Sync are two new capabilities for Oracle Database 12c that address this limitation (see the Active Data Guard section for information on Far Sync).

Fast Sync

Fast Sync provides an easy way of improving performance in synchronous zero data loss configurations. 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. This reduces the impact of synchronous transport on primary database performance by shortening the total round-trip time between primary and standby. Fast Sync is included with Data Guard.

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

Active Data Guard is a superset of 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 are described below.

Increase ROI by Offloading Workloads to an Active Data Guard Standby

Active Data Guard enables the offloading of read-only 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 best performance, while also enforcing the same read consistency model for read-only access at the standby as is enforced at the primary database. No other physical or logical replication solution does this. This makes it attractive to offload read-only workloads to an active standby, eliminating the cost of idle redundancy.

There are also many reporting applications that would be eligible to use a read-only database except for the requirement that they write to global temporary tables and/or access unique sequences. Active Data Guard includes new capabilities with Oracle Database 12c to allow writes to global temporary tables and access to unique sequences at an active standby. This further expands the number of reporting applications that can be offloaded from a primary database. No other physical or logical replication solution can provide all of these capabilities: each alternative solution is deficient in one or more areas compared to Active Data Guard. Active Data Guard is an option for Oracle Database Enterprise Edition.

In Oracle Database 12.2 on Oracle Cloud, for increased ROI, you can now run analytics queries that create ADTs on primary which then gets replicated to the standby and applied which eliminates read-only database errors. Similarly, you can run AWR reports on standby, SQL Plan Analyzer etc. on the standby.

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

Active Data Guard Far Sync is the capability introduced in Oracle Database 12c which 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 2). Production can be quickly failed over, manually or automatically, to the remote standby database with zero data loss.



Figure 2: 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 save network bandwidth by performing transport compression using the Oracle Advanced Compression option.

Consider an asynchronous 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.

In Oracle Database 12.2 on Oracle Cloud, many major enhancements are added to Data Guard to address ease of use, increased protection and performance.

- New syntaxes are introduced to create standby databases seamlessly via DBCA and EM CLI commands.
- You can also create a Far Sync instance using RMAN DUPLICATE command.

- REST interface is now supported with Data Guard Broker for easy operations. Scripting made simple with Data Guard Broker DGMGRL command line interface. Broker has been enhanced to automatically handle PDB failover/relocation to a new CDB instead of manual unplug & plug procedure. Broker is also enhanced for multitenant environment to support individual PDB-level failovers.
- Until Oracle Database 12c, only one RAC node of a standby can be used for recovery. From Oracle Database 12.2 on Oracle Cloud release, all the nodes can utilize their resources to do the recovery using Multi-Instance Redo Apply which significantly reduces recovery time.
- Oracle Database 12.2 now supports In Memory Column (IMC) Store with Standby Databases. You can create IMC tables and columns for analytics on Active Data Guard standby database and can populate with data that can be different than the production database.
- During failover/switchover scenario, read-only connections to the standby are not disconnected and they can continue to perform their read-operations. Transparent client connection failover using session draining is now supported with broker and preserves application connections to the standby during role changes.
- Standby database blocks are invalidated when a NOLOGGING operation is performed on the primary. In Oracle Database 12.2 on Oracle Cloud, the blocks can be recovered from the primary by simply running RMAN commands using new NOLOGGED BLOCK syntax.
- Additional Oracle Database 12.2 on Oracle Cloud enhancements for Data Guard and Active Data Guard includes
 - Standby-first encryption to encrypt databases with no downtime,
 - Automatic update of password files in all standby databases when the primary site password is changed,
 - Faster failovers,
 - Improved automatic block repair,
 - Fast Start Failover in Maximum Protection mode,
 - Zero data loss failovers in any mode with storage failures,
 - Network flood control .

See Data Guard and Active Data Guard resources on OTN (oracle.com/goto/dataguard).

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-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. Oracle GoldenGate, in contrast, is a feature-rich logical replication product with advanced features that 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.

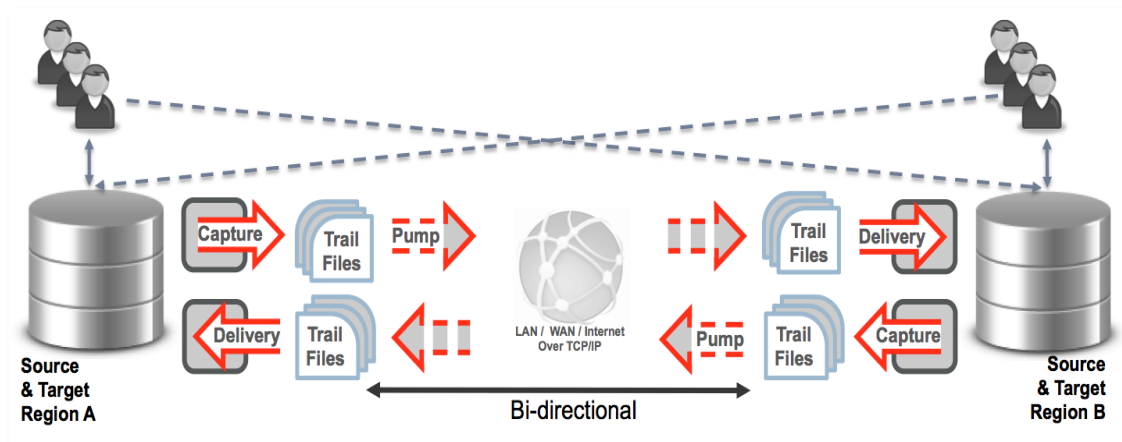



Figure 3: Oracle GoldenGate – Active-Active Bi-Directional Replication

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.



Oracle GoldenGate Oracle Database 12.2 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, 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 in the cloud and in the hybrid model between on-premises and cloud.

For more information, see GoldenGate resources on OTN (oracle.com/goto/goldengate).

GoldenGate Cloud Service

Oracle GoldenGate Cloud Service is a secure, high performance, scalable data replication solution for moving data in real-time to public cloud. Oracle GoldenGate Cloud Service complements your onpremises data replication strategy by extending data replication to cloud. You can use the familiar Oracle GoldenGate Service interface to now perform in-cloud data replication and can easily integrate new cloud solutions with existing Oracle GoldenGate Service deployments.

This new on-demand platform allows you to capture data from on-premises and deliver that data to the cloud with minimum impact to source systems. All with no up-front capital cost, or hardware expenses and in an environment with simplified configuration and management. Building on top of the elastic compute and storage offered by Oracle Public Cloud, you can run any workload on Oracle GoldenGate Cloud Service and easily expand your environment when your application needs to grow.

For more information, go to cloud.oracle.com/goldengate


Complete Site Failover: Oracle Site Guard

Oracle Site Guard is part of Oracle Enterprise Manager Cloud Control 12c, 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; Oracle Data Guard for Oracle data, and storage replication for file system data external to the Database.

Addressing 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 12c 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.

Online System Reconfiguration



Oracle supports dynamic online system reconfiguration for all components of your Oracle hardware stack. Oracle's Automatic Storage Management (ASM) has built-in capabilities that allow the online addition or removal of ASM disks. When disks are added or removed from an ASM Diskgroup – Oracle automatically rebalances the data across the new storage configuration while storage, database, and application remain online. Real Application Clusters (RAC) provides powerful online reconfiguration capabilities. Administrators can dynamically add and remove clustered nodes without any disruption to the database or the application. Oracle also supports the dynamic addition or removal of CPUs on SMP servers that have this online capability. Finally, Oracle's dynamic shared memory tuning capabilities allow administrators to grow and shrink the shared memory and database cache online. With automatic memory tuning capabilities, administrators can let Oracle automate the sizing and distribution of shared memory according to Oracle's analysis of memory usage characteristics. Oracle's extensive online reconfiguration capabilities support administrators' ability to not only minimize system downtime due to maintenance activities – but to also enable enterprises to scale capacity on demand.

Online Data and Application Change

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, 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 Data File Move and Online Partition Move

Oracle Database 12c has the ability to move a data file while users are accessing its data, via command ALTER DATABASE MOVE DATAFILE. This capability maintains data availability during maintenance operations. This capability is useful to move infrequently accessed datafiles to lower-cost storage. Another example of use is to move a database from non-ASM to ASM storage.

Online Partition Move, a new capability for Oracle Database 12c, makes it easier to compress online. It supports online, multi-partition redefinition in a single session.

Online Table Redefinition

As business requirements evolve, so too do the applications and databases supporting the business. Through the strategic use of the DBMS_REDEFINITION package (also available in 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. Enhancements to Online Table redefinition in Oracle Database 12c include:

- » 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.

In Oracle Database 12.2 on Oracle Cloud, 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.

Online Application Upgrades: Edition-Based Redefinition

Oracle Database's Edition-Based Redefinition 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. Edition-based Redefinition introduces a scope -- 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 crossedition trigger propagates data changes made by the old edition into the new edition's columns, or (in hot-rollover) vice-versa.

In Oracle Database 12.2 on Oracle Cloud, all E-Business Suite Oracle Database 12.2 patches are applied online using Edition-Based Redefinition.

Online Patching

Online patching, which is integrated with OPatch, provides the ability to patch the processes in an Oracle instance without bringing the instance down. Each process associated with the instance checks for patched code at a safe execution point, and then copies the code into its process space.

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.

Data Guard Standby-First Patch Assurance

Data Guard Standby-First Patch Assurance (Oracle Database 11.2.0.1 onward) enables physical standby to support different software patch levels between a primary and standby database for the purpose of applying and validating Oracle patches in rolling fashion.⁸ 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 (Oracle 11.2.0.1 to 11.2.0.3), or major release (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 in this 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 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, a new capability for Oracle Database 12c, 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 12c.⁹ You can use it for other database maintenance tasks with Oracle Database 12c.¹⁰

Platform Migration, Systems Maintenance, Data Center Moves

⁸ See [MOS Note 1265700.1](#) for more information on Standby-First Patch Apply eligible patches.

⁹ You must still the Transient Logical Standby upgrade when upgrading from Oracle Database 11g to Oracle Database 12c, or from Oracle Database 12.1 to the first patchset of Oracle Database 12.1.

¹⁰ 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.



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.¹¹ 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. Oracle GoldenGate offers the most flexibility for platform migration between heterogeneous platforms with minimal or zero downtime.

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.

Managing Oracle Database High Availability Solutions

Oracle Enterprise Manager Cloud Control 13c is the management interface for an Oracle environment. 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, and security capabilities, Oracle Cloud Control reduces the cost and complexity of managing complex environments, while helping customers maintain their required IT infrastructure service levels.

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.

Global Data Services

Many customers have offloaded read-only and read-mostly workloads to their Active Data Guard Standby replicas. 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.

¹¹ See [MOS Note 413484.1](#) for details on platform combinations supported in a Data Guard configuration.

Global Data Services (GDS), a new capability introduced in Oracle Database 12c, 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 4). 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. The GSMs and the GDS Catalog act in concert with new GDS functionality in Oracle Database 12c.

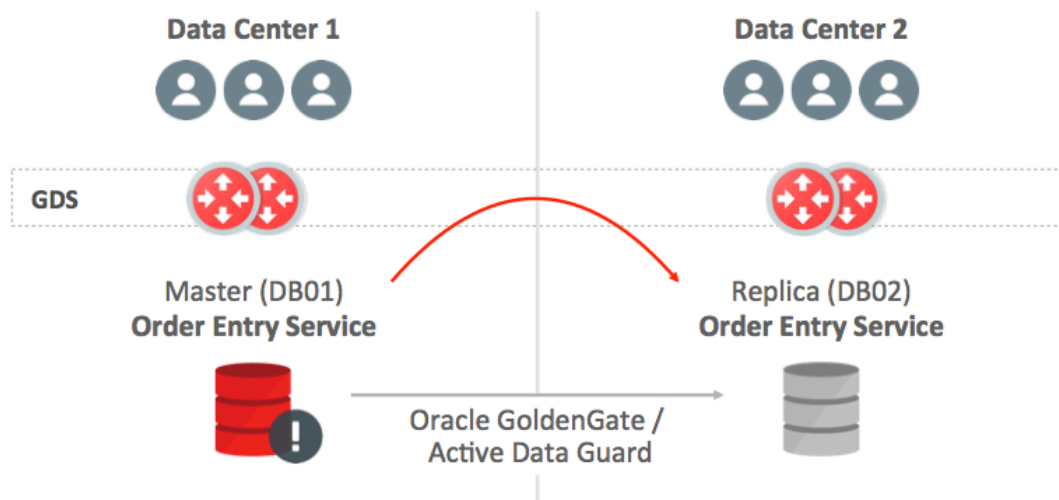



Figure 4: Global Data Services for Failover and Load Balancing Across Datacenters

In the GDS example in figure 4, depicts replicated databases, 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 upon 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 12c 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. 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.

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 12c
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 12c Release 2 on Oracle Cloud

FEATURE	DESCRIPTION OF NEW OR ENHANCED FUNCTIONALITY IN ORACLE DATABASE 12c 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 Extended RAC up to 3 sites.

**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**Hardware and Software, Engineered to Work Together**

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. 1116



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