Oracle® Database

High Availability Overview

10*g* Release 2 (10.2)

B14210-02

July 2006

ORACLE®

Oracle Database High Availability Overview, 10g Release 2 (10.2)

B14210-02

Primary Author: Immanuel Chan

Contributors: Andrew Babb, Tammy Bednar, Barb Lundhild, Rahim Mau, Valarie Moore, Ashish Ray, Vivian Schupmann, Michael T. Smith, Lawrence To, Douglas Utzig, James Viscusi, Shari Yamaguchi

Contents

3 Determining Your High Availability Requirements

4 High Availability Architectures

5 High Availability Best Practices

Index

Preface

This book introduces you to Oracle's approach for a highly available database environment. It provides an overview of high availability and helps you to determine your high availability requirements. It describes the Oracle database products and features that are designed to support high availability and describes the primary database architectures that can help your business achieve high availability.

This preface contains these topics:

- Audience
- Documentation Accessibility
- Related Documents
- Conventions

Audience

This book is intended for chief technology officers, information technology architects, database administrators, system administrators, network administrators, and application administrators who perform the following tasks:

- Plan data centers
- Implement data center policies
- Maintain high availability systems
- Plan and build high availability solutions

Documentation Accessibility

Our goal is to make Oracle products, services, and supporting documentation accessible, with good usability, to the disabled community. To that end, our documentation includes features that make information available to users of assistive technology. This documentation is available in HTML format, and contains markup to facilitate access by the disabled community. Accessibility standards will continue to evolve over time, and Oracle is actively engaged with other market-leading technology vendors to address technical obstacles so that our documentation can be accessible to all of our customers. For more information, visit the Oracle Accessibility Program Web site at

http://www.oracle.com/accessibility/

Accessibility of Code Examples in Documentation

Screen readers may not always correctly read the code examples in this document. The conventions for writing code require that closing braces should appear on an otherwise empty line; however, some screen readers may not always read a line of text that consists solely of a bracket or brace.

Accessibility of Links to External Web Sites in Documentation

This documentation may contain links to Web sites of other companies or organizations that Oracle does not own or control. Oracle neither evaluates nor makes any representations regarding the accessibility of these Web sites.

TTY Access to Oracle Support Services

Oracle provides dedicated Text Telephone (TTY) access to Oracle Support Services within the United States of America 24 hours a day, seven days a week. For TTY support, call 800.446.2398.

Related Documents

For more information, see the Oracle database documentation set. These books may be of particular interest:

- *Oracle Data Guard Concepts and Administration*

- *Oracle Database Oracle Clusterware and Oracle Real Application Clusters Administration and Deployment Guide*

- *Oracle Database Backup and Recovery Advanced User's Guide*

- *Oracle Database Administrator's Guide*

Many books in the documentation set use the sample schemas of the seed database, which is installed by default when you install Oracle. Refer to *Oracle Database Sample Schemas* for information on how these schemas were created and how you can use them yourself.

Oracle High Availability Best Practice white papers can be downloaded at

http://www.oracle.com/technology/deploy/availability/htdocs/maa.htm

Conventions

The following text conventions are used in this document:

Convention	Meaning
boldface	Boldface type indicates graphical user interface elements associated with an action, or terms defined in text or the glossary.
italic	Italic type indicates book titles, emphasis, or placeholder variables for which you supply particular values.
monospace	Monospace type indicates commands within a paragraph, URLs, code in examples, text that appears on the screen, or text that you enter.

1

Overview of High Availability

This chapter contains the following sections:

- Introduction to High Availability
- What is Availability?
- Importance of Availability
- Causes of Downtime
- What Does This Book Contain?

Introduction to High Availability

Databases and the Internet have enabled worldwide collaboration and information sharing by extending the reach of database applications throughout organizations and communities. This reach emphasizes the importance of high availability in data management solutions. Both small businesses and global enterprises have users all over the world who require access to data 24 hours a day. Without this data access, operations can stop, and revenue is lost. Users, who have become more dependent upon their solutions, now demand service-level agreements from their Information Technology (IT) departments and solutions providers. Increasingly, availability is measured in dollars, euros, and yen, not just in time and convenience.

Enterprises have used their IT infrastructure to provide a competitive advantage, increase productivity, and empower users to make faster and more informed decisions. However, with these benefits has come an increasing dependence on that infrastructure. If a critical application becomes unavailable, then the entire business can be in jeopardy. Revenue and customers can be lost, penalties can be owed, and bad publicity can have a lasting effect on customers and a company's stock price. It is critical to examine the factors that determine how your data is protected and maximize the availability to your users.

What is Availability?

Availability is the degree to which an application, service, or functionality is available upon user demand. Availability is measured by the perception of an application's end user. End users experience frustration when their data is unavailable, and they do not understand or care to differentiate between the complex components of an overall solution. Performance failures due to higher than expected usage create the same havoc as the failure of critical components in the solution.

Reliability, recoverability, timely error detection, and continuous operations are primary characteristics of a highly available solution:

- **Reliability**: Reliable hardware is one component of a high availability solution. Reliable software—including the database, Web servers, and application—is just as critical to implementing a highly available solution.

- **Recoverability**: There may be many choices in recovering from a failure if one occurs. It is important to determine what types of failures may occur in your high availability environment, and how to recover from those failures in the time that meets your business requirements. For example, if a critical table is accidentally deleted from the database, what action should you take to recover it? Does your architecture provide the ability to recover in the time specified in a service level agreement (SLA)?

- **Timely error detection**: If a component in your architecture fails, then fast detection is another essential component in recovering from a possible unexpected failure. While you may be able to recover quickly from an outage, if it takes an additional 90 minutes to discover the problem, then you may not meet your SLA. Monitoring the health of your environment requires reliable software to view it quickly and the ability to notify the DBA of a problem.

- **Continuous operations**: Continuous access to your data is essential when very little or no downtime is acceptable to perform maintenance activities. Activities such as moving a table to another location within the database, or even adding additional CPUs to your hardware, should be transparent to the end user in a high availability architecture.

More specifically, a high availability architecture should have the following traits:

- Be transparent to most failures

- Provide built-in preventative measures

- Provide proactive monitoring and fast detection of failures

- Provide fast recoverability

- Automate the recovery operation

- Protect the data so that there is minimal or no data loss

- Implement the operational best practices to manage your environment

- Provide the high availability solution to meet your SLA

Importance of Availability

The importance of high availability varies among applications. However, the need to deliver increasing levels of availability continues to accelerate as enterprises re-engineer their solutions to gain competitive advantage. Most often, these new solutions rely on immediate access to critical business data. When data is not available, the operation can cease to function. Downtime can lead to lost productivity, lost revenue, damaged customer relationships, bad publicity, and lawsuits.

If a mission-critical application becomes unavailable, then the enterprise is placed in jeopardy. It is not always easy to place a direct cost on downtime. Angry customers, idle employees, and bad publicity are all costly, but not directly measured in currency. On the other hand, lost revenue and legal penalties incurred because SLA objectives are not met can easily be quantified. The cost of downtime can quickly grow in industries that are dependent upon their solutions to provide service.

Other factors to consider in the cost of downtime are the maximum tolerable length of a single unplanned outage, and the maximum frequency of allowable incidents. If the event lasts less than 30 seconds, then it may cause very little impact and may be barely

perceptible to end users. As the length of the outage grows, the effect may grow exponentially and result in a negative impact on the business. When designing a solution, it is important to take into account these issues and to determine the true cost of downtime and the cost of added availability. An organization should then weigh the cost of downtime and balance it with the expected availability improvement. High availability solutions are effective insurance policies.

Oracle provides a range of high availability solutions that fit every organization regardless of size. Small workgroups and global enterprises alike are able to extend the reach of their critical business applications. With Oracle and the Internet, applications and their data are now reliably accessible everywhere, at any time.

Causes of Downtime

One of the challenges in designing a high availability solution is examining and addressing all the possible causes of downtime. It is important to consider causes of both unplanned and planned downtime when designing a fault tolerant and resilient IT infrastructure. Planned downtime can be just as disruptive to operations, especially in global enterprises that support users in multiple time zones.

Table 1–1 describes the outage categories and provides examples of each outage type.

Table 1–1   Causes of Downtime

Category	Outage Type	Description	Examples
Unplanned	Computer failure	A computer failure outage occurs when the system running the database becomes unavailable because it has shut down or is no longer accessible.	Database system hardware failure Operating system failure Oracle instance failure Network interface failure
	Storage failure	A storage failure outage occurs when the storage holding some or all of the database contents becomes unavailable because it has shut down or is no longer accessible.	Disk drive failure Disk controller failure Storage array failure
	Human error	A human error outage occurs when unintentional or malicious actions are committed that cause data within the database to become logically corrupt or unusable. The service level impact of a human error outage can vary significantly depending on the amount and critical nature of the affected data.	Dropped database object Inadvertent data changes Malicious data changes
	Data corruption	A data corruption outage occurs when a hardware or software component causes corrupt data to be read or written to the database. The service level impact of a data corruption outage may vary, from a small portion of the database (down to a single database block) to a large portion of the database (making it essentially unusable).	Operating system or storage device driver, host bus adapter, disk controller, or volume manager error causing bad disk read or writes Stray writes by operating system or other application software

Table 1–1   (Cont.) Causes of Downtime

Category	Outage Type	Description	Examples
	Site failure	A site failure outage occurs when an event causes all or a significant portion of an application to stop processing or slow to an unusable service level. A site failure may affect all processing at a data center, or a subset of applications supported by a data center.	Extended site-wide power failure Site-wide network failure Natural disaster making a data center inoperable Terrorist or malicious attack on operations or the site
Planned	System changes	Planned system changes occur when performing routine and periodic maintenance operations and new deployments. Planned system changes include any scheduled changes to the operating environment that occur outside the organizational data structure within the database. The service level impact of a planned system change varies significantly depending on the nature and scope of the planned outage, the testing and validation efforts made prior to implementing the change, and the technologies and features in place to minimize the impact.	Adding/removing processors to/from an SMP server Adding/removing nodes to/from a cluster Adding/removing disks drives or storage arrays Changing configuration parameters Upgrading/patching system hardware and software Upgrading/patching Oracle software Upgrading/patching application software System platform migration Database relocation
	Data changes	Planned data changes occur when there are changes to the logical structure or physical organization of Oracle database objects. The primary objective of these changes is to improve performance or manageability.	Table definition changes Adding table partitioning Creating and rebuilding indexes

Oracle offers high availability solutions to help avoid both unplanned and planned downtime, as well as recover from failures. Chapter 2 discusses each of these high availability solutions in detail.

What Does This Book Contain?

Choosing and implementing the architecture that best fits your availability requirements can be a daunting task. This architecture must:

- Encompass redundancy across all components

- Provide protection from computer failures, storage failures, human errors, data corruption, and site disasters

- Recover from outages as quickly and transparently as possible

- Provide solutions to eliminate or reduce planned downtime

- Provide consistent high performance

- Be easy to deploy, manage, and scale

To help you select the most suitable architecture for your organization, this book describes several high availability architectures and provides guidelines for choosing the one that best meets your requirements. Knowledge of the Oracle Database server, Oracle Real Application Clusters and Oracle Data Guard terminology is required to understand the configuration and implementation details.

Chief technology officers and information technology architects can benefit from reading the following chapters:

- Chapter 3, "Determining Your High Availability Requirements"
- Chapter 4, "High Availability Architectures"

Database administrators and network administrators can find useful information in the following chapters:

- Chapter 2, "Oracle Database High Availability Solutions"
- Chapter 4, "High Availability Architectures"

Oracle High Availability Best Practice white papers can be downloaded at

http://www.oracle.com/technology/deploy/availability/htdocs/maa.htm

2

Oracle Database High Availability Solutions

Oracle Database 10*g* offers an integrated suite of high availability solutions that increase availability and eliminate or minimize both planned and unplanned downtime. These solutions also help enterprises maintain 24x7 business continuity:

- Oracle High Availability Features
- Oracle High Availability Solutions for Unplanned Downtime
- Oracle High Availability Solutions for Planned Downtime
- High Availability and Grid Computing
- High Availability Management

Oracle High Availability Features

Oracle provides the following features for high availability:

- Oracle Real Application Clusters
- Oracle Data Guard
- Oracle Streams
- Oracle Flashback Technology
- Automatic Storage Management
- Recovery Manager
- Flash Recovery Area
- Oracle Security Features
- Fast-Start Fault Recovery
- LogMiner
- Hardware Assisted Resilient Data (HARD) Initiative

Oracle Real Application Clusters

Oracle Real Application Clusters (RAC) allows the Oracle database to run any packaged or custom application unchanged across a set of clustered servers. This capability provides the highest levels of availability and the most flexible scalability. If a clustered server fails, the Oracle database will continue running on the surviving servers. When more processing power is needed, another server can be added without interrupting user access to data.

RAC enables multiple instances that are linked by an interconnect to share access to an Oracle database. In a RAC environment, the Oracle database runs on two or more systems in a cluster while concurrently accessing a single shared database. The result is a single database system that spans multiple hardware systems yet appears as a single unified database system to the application. This enables RAC to provide high availability, scalability, and redundancy during failures within the cluster. RAC accommodates all system types, from read-only data warehouse (DSS) systems to update-intensive online transaction processing (OLTP) systems.

High availability configurations have redundant hardware and software that maintain operations by avoiding single points-of-failure. To accomplish this, the Oracle Clusterware is installed as part of the RAC installation process. Oracle Clusterware is a portable solution that is integrated and designed specifically for the Oracle database. In a RAC environment, Oracle Clusterware monitors all Oracle components (such as instances and listeners). If a failure occurs, Oracle Clusterware will automatically attempt to restart the failed component. Other non-Oracle processes can also be managed by Oracle Clusterware. During outages, Oracle Clusterware relocates the processing performed by the inoperative component to a backup component. For example, if a node in the cluster fails, Oracle Clusterware will cause client processes running on the failed node to reconnect and resume running on a surviving node.

The Oracle Clusterware requires two files, the Oracle Cluster Registry (OCR) and the voting disk. To avoid single points-of-failure, the Oracle Clusterware automatically maintains redundant copies of these files. Oracle Clusterware also enables you to replace a damaged copy of the OCR online. Oracle's recovery processes quickly re-master resources, recover partial or failed transactions, and rapidly restore the system.

RAC provides the following benefits:

- Ability to tolerate and quickly recover from computer and instance failures

- Fast, automatic, and intelligent connection and service relocation and failover

- Rolling patch upgrades for qualified one-off patches

- Rolling release upgrades of Oracle Clusterware

- Load balancing advisory

- Runtime connection load balancing

- Flexibility to scale up processing capacity using commodity hardware without downtime or changes to the application

- Comprehensive manageability integrating database and cluster features

> **See Also:** *Oracle Database Oracle Clusterware and Oracle Real Application Clusters Administration and Deployment Guide*

Oracle Data Guard

Oracle Data Guard provides a comprehensive set of services that create, maintain, manage, and monitor one or more standby databases to enable production Oracle databases to survive failures, disasters, errors, and data corruption. Data Guard maintains these standby databases as transactionally consistent copies of the production database. Then, if the production database becomes unavailable due to a planned or an unplanned outage, Data Guard can switch any standby database to the production role, thus greatly reducing the downtime caused by the outage. The failover of data processing from the production to the standby database can be completely automatic and done without any human intervention, thereby reducing the

management costs associated with the Data Guard configuration. Data Guard can be used with traditional backup, restore, and clustering solutions to provide a high level of data protection and data availability.

A Data Guard configuration consists of one production database and one or more physical or logical standby databases. The databases in a Data Guard configuration are connected by Oracle Net and may be dispersed geographically. There are no restrictions on where the databases are located if they can communicate with each other. For example, you can have a standby database in the same building as your primary database to help manage planned downtime and two or more standby databases in other locations for use in disaster recovery.

Oracle Data Guard provides the following benefits:

- Maintains real-time, transactionally consistent database copies to provide protection against unplanned downtime and disaster

- Complete data protection against computer failures, human errors, data corruption, and site failures

- Reduces planned downtime for hardware and system upgrades, and Oracle patch set and database upgrades

- Automatic detection and resolution of missing data following temporary loss of connectivity between the primary and standby database

- Multiple levels of data protection and performance to balance data availability against system performance requirements

- Allows efficient use of system resources by diverting reporting and backup operations from the production database to standby databases

- Ability to diverge a Redo Apply standby database for reporting or testing purposes and resynchronize it with primary database once complete

- Managed and automatic role transition and application notification to minimize planned and unplanned downtime

- Automatic resynchronization of a failed primary database following a failover

- All systems managed as a single configuration for simplified administration

> **See Also:** *Oracle Data Guard Concepts and Administration*

Oracle Streams

Oracle Streams enables the propagation and management of data, transactions, and events in a data stream, either within a database or from one database to another. Streams provides a set of elements that allow users to control what information is put into a data stream, how the stream is routed from node to node, what happens to events in the stream as they flow into each node, and how the stream terminates.

Streams can be used to replicate a database or a subset of a database. This enables users and applications to simultaneously update data at multiple locations. If a failure occurs at one of the locations, then users and applications at the surviving sites can continue to access and update data.

Streams can be used to build distributed applications that replicate changes at the application level using message queuing. If an application fails, then the surviving applications can continue to operate and provide access to data through locally maintained copies.

Streams provides granularity and control over what is replicated and how it is replicated. It supports bidirectional replication, data transformations, subsetting, custom apply functions, and heterogeneous platforms. It also gives users complete control over the routing of change records from the primary database to a replica database.

As with Streams, Oracle Data Guard in SQL Apply mode can capture database changes, propagate them to destinations, and apply the changes at these destinations. Although Streams and Data Guard in SQL Apply mode share much of the same underlying technologies for high availability, Data Guard in SQL Apply mode is easier to implement and manage than a Streams-based high availability solution.

Oracle Streams provides the following benefits:

- Data protection by maintaining a full or partial remote copy of the database

- Achieves little or no downtime during database upgrade or maintenance operations such as migrating a database to a different platform or character set, modifying database objects to support upgrades to user-created applications, and applying an Oracle software patch

- Data replication by capturing DML and DDL changes made to database objects and replicating these changes to one or more other databases

- Event management and notification by enqueuing messages or capturing events, propagating the messages and events through queues, and dequeuing and applying or acting upon the message or event

- Supports heterogeneous platforms across databases within the configuration

- Allows character sets to differ between replicas

- Permits fine-grained control of data sharing

> **Note:** The increased flexibility and capability of Oracle Streams over Oracle Data Guard with SQL Apply requires more investment and expertise to build and maintain an integrated high availability solution.

See Also: *Oracle Streams Concepts and Administration*

Oracle Flashback Technology

Flashback technology provides a set of features to view and rewind data back and forth in time. The flashback features offer the capability to query past versions of schema objects, query historical data, perform change analysis, and perform self-service repair to recover from logical corruption while the database is online.

Flashback technology provides a SQL interface to quickly analyze and repair human errors. Flashback provides fine-grained analysis and repair for localized damage such as deleting the wrong customer order. Flashback technology also enables correction of more widespread damage, yet does it quickly to avoid long downtime. Flashback technology is unique to the Oracle Database and supports recovery at all levels including row, transaction, table, tablespace, and database.

Flashback technology includes the following features:

- Oracle Flashback Query

- Oracle Flashback Versions Query

- Oracle Flashback Transaction Query

- Oracle Flashback Table

- Oracle Flashback Drop

- Oracle Flashback Database

- Oracle Flashback Restore Points

> **See Also:**
>
> - *Oracle Database Backup and Recovery Advanced User's Guide*
>
> - *Oracle Database SQL Reference*

Oracle Flashback Query

Oracle Flashback Query provides the ability to view the data as it existed in the past by utilizing the Automatic Undo Management system to obtain metadata and historical data for transactions. Undo data is persistent and will survive a database malfunction or shutdown. The unique features of Flashback Query not only provide the ability to query previous versions of tables, they also provide a powerful mechanism to recover from erroneous operations.

Uses of Flashback Query include:

- Recovering lost data or undoing incorrect, committed changes. For example, rows that have been deleted or updated can be immediately repaired even after they have been committed.

- Comparing current data with the corresponding data at some time in the past. For example, using a daily report that shows the changes in data from yesterday, it is possible to compare individual rows of table data, or find intersections or unions of sets of rows.

- Checking the state of transactional data at a particular time, such as verifying the account balance on a certain day.

- Simplifying application design by removing the need to store certain types of temporal data. Using a Flashback Query, it is possible to retrieve past data directly from the database.

- Applying packaged applications, such as report generation tools, to past data.

- Providing self-service error correction for an application, enabling users to undo and correct their errors.

Oracle Flashback Versions Query

Oracle Flashback Versions Query is an extension to SQL that can be used to retrieve the versions of rows in a given table that existed in a specific time interval. Oracle Flashback Versions Query returns a row for each version of the row that existed in the specified time interval. For any given table, a new row version is created each time the `COMMIT` statement is executed.

Flashback Versions Query is a powerful tool for the DBA to run analysis to determine what happened. Additionally, application developers can use Flashback Versions Query to build customized applications for auditing purposes.

Oracle Flashback Transaction Query

Oracle Flashback Transaction Query provides a mechanism to view all changes made to the database at the transaction level. When used in conjunction with Flashback

Versions Query, it offers a fast and efficient means to recover from a user or application error. Flashback Transaction Query increases the ability to perform online diagnosis of problems in the database by returning the user that changed the row, and performs analysis and audits on transactions.

Oracle Flashback Table

Oracle Flashback Table enables users to recover a table to a previous point in time. It provides a fast, online solution for recovering a table or set of tables that has been erroneously modified by a user or application. In most cases, Flashback Table alleviates the need for administrators to perform more complicated point-in-time recovery operations. Even after a flashback, the data in the original table is not lost; it can later be reverted back to the original state.

Oracle Flashback Drop

Dropping objects by accident has always been a problem for users and DBAs alike. Historically, there is no easy way to recover dropped tables, indexes, constraints, or triggers. Oracle Flashback Drop provides a safety net when dropping objects. When a user drops a table, Oracle automatically places it into the Recycle Bin. The Recycle Bin is a virtual container where all dropped objects reside. Users can continue to query data in a dropped table.

Oracle Flashback Database

Oracle Flashback Database provides a more efficient alternative to database point-in-time recovery. With Oracle Flashback Database, current datafiles can be reverted to their contents at a past time. The result is much like the result of a point-in-time recovery using datafile backups and redo logs, but it is not necessary to restore datafiles from backup, or to re-apply as many individual changes in the redo logs as required in conventional media recovery.

Enabling Oracle Flashback Database provides the following benefits:

- Eliminate the time to restore a backup when fixing human error that has a database-wide impact.

- Allows standby databases to use real-time apply to synchronize with the primary database because humans errors can be quickly undone.

- Allows quick standby database reinstantiation after a database failover.

Oracle Flashback Restore Points

When an Oracle Flashback recovery operation is performed on the database, the DBA must determine the point in time—identified by the System Change Number (SCN) or timestamp—to which the data can later be flashed back. Oracle Flashback restore points are user-defined labels that can be substituted for the SCN or transaction time used in Flashback Database, Flashback Table, and Recovery Manager (RMAN) operations. Furthermore, a database can be flashed back through a previous database recovery and open resetlogs by using guaranteed restore points. Guaranteed restore points allow major database changes—such as database batch jobs, upgrade, or patch—to be quickly undone by ensuring that the undo required to rewind the database is retained.

Using a combination of Oracle Data Guard, Flashback restore points and RMAN incremental backups, a physical standby database can be opened temporarily in read/write mode for development, reporting, or testing purposes. The physical standby database can then be resynchronized as an updated physical standby

database by flashing back to the restore point and applying a recent incremental backup from the primary database.

Using Oracle Flashback restore points provides the following benefits:

- Provides the ability to quickly cancel planned database changes that produced undesirable results, such as a failed batch job or application upgrade

- Can be used in conjunction with Oracle Data Guard and RMAN incremental backups to quickly resynchronize a read/write clone database with the primary database

> **See Also:** *Oracle Data Guard Concepts and Administration*

Automatic Storage Management

Automatic Storage Management (ASM) provides a vertically integrated file system and volume manager directly in the Oracle kernel, resulting in:

- Significantly less work to provision database storage

- Higher level of availability

- Elimination of the expense, installation, and maintenance of specialized storage products

- Unique capabilities for database applications

For optimal performance, ASM spreads files across all available storage. To protect against data loss, ASM extends the concept of SAME (stripe and mirror everything) and adds more flexibility in that it can mirror at the database file level rather than the entire disk level.

More importantly, ASM eliminates complexities associated with managing data and disks; it simplifies the processes of setting up mirroring, adding disks, and removing disks. Instead of managing hundreds and possibly thousands of files (as in a large data warehouse), DBAs using ASM create and administer a larger-grained object—the disk group—which identifies the set of disks that will be managed as a logical unit. Automation of file naming and placement of the underlying database files save DBAs time and ensures adherence to standard best practices.

The ASM native mirroring mechanism (2-way or 3-way) is an option that is used to protect against storage failures. With ASM mirroring, an additional level of data protection can be provided with the use of failures groups. A failure group is a set of disks sharing a common resource (disk controller or an entire disk array) whose failure can be tolerated. Once defined, an ASM failure group will intelligently place redundant copies of the data in separate failure groups to ensure that the data will be available and transparently protected against the failure of any component in the storage subsystem.

ASM provides the following benefits:

- Provides the ability to mirror across drives and storage arrays

- Automatically re-mirrors from a failed drive to remaining drives

- Automatically rebalances stored data when disks are added or removed while the database remains online

- Allows for operational simplicity in managing a database storage grid

> **See Also:** *Oracle Database Administrator's Guide*

Recovery Manager

Recovery Manager (RMAN) is an Oracle utility to manage the backup and, more importantly, the recovery of the database. It eliminates operational complexity while providing superior performance and availability of the database.

Recovery Manager determines the most efficient method of executing the requested backup, restoration, or recovery operation and then submits these operations to the Oracle database server for processing. Recovery Manager and the server automatically identify modifications to the structure of the database and dynamically adjust the required operation to adapt to the changes.

RMAN provides the following benefits:

- Automated channel failover on backup and restore operations

- Automatic failover to a previous backup when the restore operation discovers a missing or corrupt backup

- Automated creation of new database and temporary files during recovery

- Automated recovery through a previous point-in-time recovery—recovery through resetlogs

- Block media recovery enables the datafile to remain online while fixing the block corruption

- Fast incremental backups using block change tracking

- Merge incremental backups into image copies in the background providing up-to-date recoverability

- Optimized backup and restore of required files only

- Retention policy ensures that relevant backups are retained

- Resumable backup and restore of previously failed operations

- Automatic backup of the control file and the server parameter file ensuring that backup metadata is available in times of database structural changes as well as media failure and disasters

- Online backup does not require the database to be placed into hot backup mode

> **See Also:** *Oracle Database Backup and Recovery Advanced User's Guide*

Flash Recovery Area

The flash recovery area is a unified storage location for all recovery-related files and activities in an Oracle database. After this feature is enabled, all RMAN backups, archive logs, control file autobackups, and datafile copies are automatically written to a specified file system or automatic storage management disk group, and the management of this disk space is handled by RMAN and the database server.

Making a backup to disk is faster because using the flash recovery area eliminates the bottleneck of writing to tape. More importantly, if database media recovery is required, then datafile backups are readily available. Restoration and recovery time is reduced because you do not need to find a tape and a free tape device to restore the needed datafiles and archive logs.

The flash recovery area provides:

- Unified storage location of related recovery files

- Management of the disk space allocated for recovery files to simplify database administration tasks

- Fast, reliable disk-based backup and restoration

  See Also: *Oracle Database Backup and Recovery Advanced User's Guide*

Oracle Security Features

The best protection against human errors is to prevent their occurrence. The best way to prevent human errors is to restrict user access to data and services to only those they truly need to perform their business functions. Oracle provides a wide range of security tools to control user access to application data by authenticating users and then enabling administrators to grant users only those privileges required to perform their duties.

In addition, the security model of the Oracle database provides the ability to restrict data access at a row level using the Virtual Private Database feature, thereby further isolating users from data that they do not need to access.

Oracle security features include:

- Authentication control to validate the identities of entities using the networks, databases, and applications

- Authorization control to provide limits to user access and actions linked by user identities and roles

- Access control to objects, providing protection regardless of the entity seeking to access or alter them

- Auditing control to monitor and gather data about specific database activities, investigate suspicious activity, deter users (or others) from inappropriate activities, and detect problems with authorization or access control implementation

- Security policy management using profiles

- Encryption of data residing within the database and backups, or transferred to and from databases

  See Also: *Oracle Database Security Guide*

Fast-Start Fault Recovery

Oracle provides fast and predictable recovery from system faults and database failures. The Fast-Start Fault Recovery technology included in the Oracle database automatically bounds database recovery time upon startup by using its self-tuned checkpoint processing. This makes recovery time fast and predictable, and improves the ability to meet service level objectives. Oracle's Fast-Start Fault Recovery can reduce recovery time on a heavily-laden database from tens of minutes to a few seconds.

Fast-Start Fault Recovery features include:

- Predictable, bounded recovery from computer failures

- Database checkpointing is self-tuning to maintain desired recovery time objective

  See Also: *Oracle Database Backup and Recovery Advanced User's Guide*

LogMiner

Oracle log files contain useful information about the activities and history of the Oracle database. Log files contain all data necessary to perform a database recovery, and also record all changes made to the data and metadata within the database.

LogMiner is a fully relational tool that allows redo log files to be read, analyzed, and interpreted using SQL. Analysis of the log files with LogMiner can be used to:

- Track or audit changes to data

- Provide supplemental information for tuning and capacity planning

- Retrieve critical information for debugging complex applications

- Recover deleted data

LogMiner features include:

- Pinpoint when a logical corruption to the database—such as errors made at the application level—may have occurred

- Determine the necessary actions to perform fine-grained recovery at the transaction level

- Performance tuning and capacity planning through trend analysis

- Perform post-auditing

> **See Also:** *Oracle Database Utilities*

Hardware Assisted Resilient Data (HARD) Initiative

Oracle9*i* introduced the Hardware Assisted Resilient Data (HARD) Initiative, a program designed to prevent data corruption before it happens. Data corruption is very rare, but when it happens, it can have a catastrophic effect on a database, and therefore a business.

Under the HARD Initiative, Oracle works with selected system and storage vendors to build operating system and storage components that can detect corruption early and prevent corrupted data from being written to disk. The key approach is block checking where the storage subsystem validates the Oracle block contents.

To use HARD validation, all datafiles and log files are placed on HARD-compliant storage. The user must also enable the HARD validation feature on the storage, using the vendor-provided interface. When Oracle writes data to the storage, the storage system validates the data. If the data appears to be corrupted, then the write is either rejected with an error, or it is accepted with an error logged by the storage in the internal logs.

Figure 2–1   Oracle Data Validation
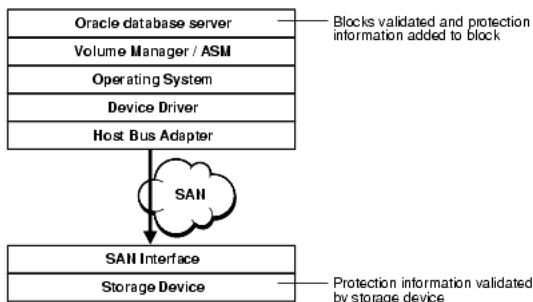
Storage vendors may choose to implement some or all of the checks in their implementation. Also, each vendor's implementation is unique and their control interfaces may have different features. Please check with the HARD initiative page for the latest vendor and implementation information.

`http://www.oracle.com/technology/deploy/availability/htdocs/HARD.html`

Oracle High Availability Solutions for Unplanned Downtime

Oracle provides high availability solutions for all types of unplanned downtime:

- Computer Failures

- Storage Failures

- Human Errors

- Data Corruption

- Site Failures

Table 2–1 describes the various Oracle high availability solutions for unplanned downtime along with the recovery time that can be attained with each solution.

Table 2–1   Oracle High Availability Solutions for Unplanned Downtime

Outage Type	Oracle Solution	Benefits	Recovery Time
Computer failures	Fast-Start Fault Recovery	Tunable and predictable cache recovery	Minutes to hours[1]
	RAC	Automatic recovery of failed nodes and instances, fast connection failover, and service failover	No downtime[2]
	Data Guard	Fast Start Failover and fast connection failover	Seconds to 5 minutes
	Oracle Streams	Online replica database	No downtime[2]
Storage failures	ASM	Mirroring and online automatic rebalance	No downtime
	RMAN with flash recovery area	Fully managed database recovery and managed disk-based backups	Minutes to hours
	Data Guard	Fast Start Failover and fast connection failover	Seconds to 5 minutes
	Oracle Streams	Online replica database	No downtime[2]
Human errors	Oracle security features	Restrict user access as prevention	No downtime
	Oracle Flashback technology	Fine-grained and database-wide rewind capability	< 30 minutes[3]
	LogMiner	Log analysis	Minutes to hours
Data corruption	HARD	Corruption prevention within a storage array	No downtime
	RMAN with flash recovery area	Online block media recovery and managed disk-based backups	Minutes to hours
	Data Guard	Automatic validation of redo blocks before they are applied, execute fast failover to an uncorrupted standby database	Seconds to 5 minutes
	Oracle Streams	Online replica database	No downtime[2]

Table 2–1 (Cont.) Oracle High Availability Solutions for Unplanned Downtime

Outage Type	Oracle Solution	Benefits	Recovery Time
Site failures	RMAN	Fully managed database recovery and integration with tape management vendors	Hours to days
	Data Guard	Fast Start Failover and fast connection failover	Seconds to 5 minutes[4]
	Oracle Streams	Online replica database	Seconds to 5 minutes[4]

[1] Recovery time consists largely of the time it takes to restore the failed system.

[2] Database is still available, but portion of application connected to failed system is affected.

[3] Recovery time for human errors depend primarily on detection time. If it takes seconds to detect a malicious DML or DLL transaction, it typically only requires seconds to flashback the appropriate transactions. Longer detection time usually leads to longer recovery time required to repair the appropriate transactions. An exception is undropping a table, which is literally instantaneous regardless of detection time.

[4] Recovery time indicated applies to database and existing connection failover. Network connection changes and other site-specific failover activities may lengthen overall recovery time.

Computer Failures

A computer failure outage occurs when the system running the database becomes unavailable because it has shut down or is no longer accessible. Downtime caused by computer failures can be reduced by employing rapid database recovery upon startup, or avoided by using cluster technology or data mirroring techniques.

Oracle offers the following high availability solutions to address computer failures:

- Fast-Start Fault Recovery

- Oracle Real Application Clusters

- Oracle Data Guard

- Oracle Streams

For information on the benefits and attainable recovery time for each solution, see Table 2–1.

Storage Failures

A storage failure outage occurs when the storage holding some or all of the database contents becomes unavailable because it has shut down or is no longer accessible. Downtime caused by storage failures can be reduced by keeping disk-based backups, copies, or replicas of the database, or avoided by using storage mirroring.

Oracle offers the following high availability solutions to address storage failures:

- Automatic Storage Management

- Recovery Manager

- Flash Recovery Area

- Oracle Data Guard

- Oracle Streams

For information on the benefits and attainable recovery time for each solution, see Table 2–1.

Human Errors

A human error outage occurs when unintentional or malicious actions are committed that cause data within the database to become logically corrupt or unusable. The service level impact of a human error outage can vary significantly depending on the amount and critical nature of the affected data. The best protection against a human error outage is to prevent humans errors from occurring where possible, and when prevention is not possible, to detect and undo the errors quickly.

Oracle offers the following high availability solutions to address human errors:

- Oracle Security Features
- Oracle Flashback Technology
- LogMiner

For information on the benefits and attainable recovery time for each solution, see Table 2–1.

Data Corruption

A data corruption outage occurs when a hardware or software component causes corrupt data to be read or written to the database. The service level impact of a data corruption outage may vary, from a small portion of the database (down to a single database block) to a large portion of the database (making it essentially unusable). If not prevented—or quickly detected and repaired—data corruption can disrupt the entire database or cause key business data to be lost.

Oracle offers the following high availability solutions to prevent—or detect and repair—data corruption:

- Hardware Assisted Resilient Data (HARD) Initiative
- Recovery Manager
- Flash Recovery Area
- Oracle Data Guard
- Oracle Streams
- Block checking `db_block_checking` and block checksumming with `db_block_checksum`
- `DBVERIFY`, `ANALYZE`, and `DBMS_REPAIR` detection tools

For information on the benefits and attainable recovery time for each solution, see Table 2–1.

Site Failures

A site failure outage occurs when an event causes all or a significant portion of an application to stop processing or slow to an unusable level. A site failure may affect all processing at a data center, or a subset of applications supported by a data center. Downtime caused by a site failure can be minimized by keeping copies or replicas of the database updated in real time.

Oracle offers the following high availability solutions to address site failures:

- Recovery Manager
- Oracle Data Guard
- Oracle Streams

For information on the benefits and attainable recovery time for each solution, see Table 2–1.

Oracle High Availability Solutions for Planned Downtime

Planned downtime can be just as disruptive to operations as unplanned downtime. This holds especially true for global enterprises that need to support users in multiple time zones, or for those that need to provide 24x7 Internet access to their customers.

Planned downtime usually becomes necessary when performing routine operations, periodic maintenance, and new deployments. Routine operations include frequent maintenance tasks such as backup, performance tuning, user management, security enhancements, and batch operations. Periodic maintenance—such as patching or reconfiguring the system—may be necessary to update the database, application, operating system, middleware, or network on occasion. New deployments include major upgrades or new rollouts of the hardware, database, application, operating system, middleware, or network.

When the volume of data stored in a database becomes very large, such maintenance operations that require planned downtime may become quite time consuming. It thus becomes very important that these maintenance operations be performed without affecting the users' access to the data.

Oracle provides the following high availability solutions to address planned downtime:

- For system changes:
    - Dynamic Resource Provisioning
    - Rolling Upgrades
- For data changes:
    - Online Reorganization and Redefinition

Dynamic Resource Provisioning

Oracle continues to broaden support for dynamic reconfiguration of the database, enabling it to adapt to changes in hardware demands without any service interruptions. The Oracle database dynamically accommodates various changes to hardware and database configurations:

- Add and remove processors from an SMP server
- Add and remove nodes and instances in an Oracle Real Application Cluster (RAC) environment
- Dynamically grow and shrink its shared memory allocation and automatically tune memory online using Automatic Shared Memory Management
- Add and remove database disks online without disturbing database activities using Automatic Storage Management (ASM)
- Add and remove storage arrays online without disturbing database activities using ASM
- Automatically rebalance I/O load across the database storage using ASM
- Move datafiles online when adding or dropping disks using ASM, which automatically rebalances database storage whenever the storage configuration is changed

These capabilities provide no-cost system changes and capacity on-demand provisioning, both of which are fundamental requirements of enterprise Grid computing.

Memory and storage management have improved significantly with the advent of Automatic Shared Memory Management and Automatic Storage Management (ASM). By setting the SGA_TARGET parameter to a nonzero value, the shared pool, large pool, Java pool, Streams pool, and buffer cache can automatically and dynamically resize as needed. ASM automates and simplifies the layout of datafiles, control files, and log files. Database files are automatically distributed across all available disks, and database storage is rebalanced whenever the storage configuration changes, including adding and removing disks or storage arrays. ASM provides redundancy through the mirroring of database files, and provides optimal performance by automatically distributing database files across all available disks. Rebalancing of the database storage automatically occurs whenever the storage configuration changes.

Another type of dynamic reconfiguration occurs when Oracle polls the operating system to detect changes in the number of available CPUs and reallocates internal resources. In addition, almost all initialization parameters can be changed without shutting down the instance. Simply use the ALTER SESSION statement to change the value of a parameter during a session, or the ALTER SYSTEM statement to change the value of a parameter in all sessions of an instance for the duration of the instance.

> **See Also:** *Oracle Database Concepts* and *Oracle Database Administrator's Guide* for information on Automatic Shared Memory Management and Automatic Storage Management

Rolling Upgrades

The Oracle database continues to reduce downtime required for system, software, and application upgrades. Oracle provides the following benefits:

- Zero downtime for system and hardware upgrades with RAC

- Zero downtime for operating system upgrades with RAC

- Zero downtime for qualified one-off database patches with RAC

- Zero downtime for storage migration with ASM

- Minimum downtime for system or cluster upgrades with Data Guard

- Minimum downtime for patchset or database upgrades with Data Guard

- Minimum downtime for database upgrade with Transportable Tablespace

- Minimum downtime for platform migration using Transportable Tablespace and potentially Data Guard

- Minimum downtime for database upgrade with Oracle Streams

- Minimum downtime for platform migration with Oracle Streams

Table 2–2 describes the various Oracle high availability solutions for planned downtime along with the recovery time that can be attained with each solution and their known considerations. For all cases, extensive testing is highly recommended prior to performing any rolling upgrade.

Table 2–2   Oracle High Availability Solutions for Planned Downtime

Maintenance Type	Oracle Solution	Description	Recovery Time	Considerations
System and hardware upgrades	RAC	To avoid downtime: 1. Dynamically redirect connections and services to a different instance. 2. Shut down target instance. 3. Upgrade target node while other nodes and instances are still available. 4. Start node and instance. Repeat on another node.	No downtime	Need to check for system restrictions. Need to check if the database and clusterware versions are certified with the new system and hardware changes.
Operating system upgrade	RAC	To avoid application downtime: 1. Dynamically redirect connections and services to a different instance. 2. Shut down target instance. 3. Upgrade operating system on target node while other nodes and instances are still available. 4. Start node and instance. Repeat on another node.	No downtime	Need to check if the database and the clusterware versions are certified for both operating system patch releases.
Oracle one-off patches	RAC	"One-off" patches—or interim patches—to database software are usually applied to implement known fixes for software problems, or to apply diagnostic patches to gather information on a problem. Such patch application is often performed during a schedule maintenance outage. Oracle provides the capability to do rolling patch upgrades with RAC with little or no database downtime using the `opatch` command-line utility. A RAC rolling upgrade enables at least some instances of the RAC installation to be available during the scheduled outage required for patch upgrades. Only the RAC instance that is currently being patched must be disabled. The other instance remains available. This means that the impact on the application downtime required for scheduled outages is further reduced. The Oracle `opatch` utility enables the user to apply the patch successively to the different instances in a RAC installation.	No downtime	Rolling upgrade is only available for patches that are certified for rolling upgrades. Typically, patches that can be installed in a rolling upgrade include: ■ Patches that do not affect the contents of the database, such as the data dictionary ■ Patches not related to RAC inter-node communication ■ Patches related to client-side tools such as SQL*Plus, Oracle utilities, development libraries, and Oracle Net ■ Patches that do not change shared database resources, such as datafile headers, control files, and common header definitions of kernel modules RAC cannot be used for rolling upgrade of patch sets.
CRS upgrades	RAC	All upgrades to Oracle Clusterware can be done in a rolling fashion.	No downtime	

Table 2–2 (Cont.) Oracle High Availability Solutions for Planned Downtime

Maintenance Type	Oracle Solution	Description	Recovery Time	Considerations
Storage migration[1]	ASM	ASM enables you to add all disks in one storage array and subsequently drop all disks from another array. ASM will automatically rebalance and migrate data to the new storage while the database remains operational.	No downtime	Before removing the source storage array, ensure that the rebalancing is complete.
System and cluster upgrades, and ASM upgrades that are not rolling upgradeable	Data Guard	For system upgrades that are not rolling upgradable with RAC due to system restrictions or cluster firmware upgrades that require downtime, upgrade the standby first and then leverage Data Guard to switch over to a physical or logical standby database: 1. Issue Data Guard Switchover (only downtime component: optimally seconds to minutes). 2. Shut down initial primary database (now standby). 3. Execute system and cluster upgrade steps. 4. Restart as standby database and allow recovery to synchronize. 5. Optionally issue Data Guard Switchover to return to original database.	Seconds to minutes	For fastest switchover, the standby database should be using real-time apply and synchronized prior to the switchover operation. This is the best approach if RAC rolling upgrade is not possible.
Patchset and database upgrades	Data Guard using SQL Apply	Leverage Data Guard using SQL Apply to upgrade an Oracle database: 1. Set up SQL Apply (logical standby database). 2. Upgrade logical standby database to new release. 3. Disconnect applications. 4. Execute Data Guard switchover. 5. Reconnect applications to the new primary database. 6. Shut down initial primary database (now logical standby database). 7. Execute database software upgrade steps. 8. Restart the standby database and allow recovery to synchronize. 9. Optionally issue Data Guard Switchover to return to the original database.	Seconds to minutes	Only supported for Oracle database versions 10.1.0.3 and higher. SQL Apply has some data type restrictions. For more information, see *Oracle Data Guard Concepts and Administration*. This is the best approach if RAC rolling upgrade is not possible and there are no data type restrictions.

Table 2–2   (Cont.) Oracle High Availability Solutions for Planned Downtime

Maintenance Type	Oracle Solution	Description	Recovery Time	Considerations
Database upgrades and platform migration	Transportable tablespace	Transporting a database requires only copying datafiles and integrating the tablespace structural information. Tablespaces can even be transported between databases from different releases. With Oracle database 10*g*, tablespaces can be transported across platforms. To perform a database upgrade or platform migration: **1.** Create and prepare a separate database using the target release. **2.** Transport tablespace from primary database to target database. Only copy datafiles from the source to target if the databases are not on the same storage device. **3.** Prepare and open the new production database. If the target database resides on a separate host but on the same platform, create a physical standby database from the initial primary database co-located with the target database. After a Data Guard Switchover, transport the tablespaces from the source to the target without incurring the file transfer time as part of the downtime.[2]	Minutes to hours	Transportable tablespace has limitations and restrictions in regard to character sets, opaque types, and system tablespace objects. Unlike previous solutions, the steps are not automated. Transportable tablespaces do provide the following benefits: ■ Provides an easier and more efficient means for content providers to publish structured data and distribute to customers running Oracle on a different platform ■ Simplifies the distribution of data from a data warehousing environment to data marts that are often running on smaller systems with a different platform ■ Enables the sharing of read-only tablespaces across a heterogeneous cluster This is the best rolling upgrade approach if both of the following are true: ■ RAC rolling upgrade, SQL Apply or Streams upgrade approach is not possible. ■ The time to run upgrade or migration scripts is considerably greater than the time to export and import the metadata between source and target databases.

Table 2–2   (Cont.) Oracle High Availability Solutions for Planned Downtime

Maintenance Type	Oracle Solution	Description	Recovery Time	Considerations
Database upgrades and platform migration	Oracle Streams	Like Data Guard using SQL Apply, Oracle Streams can capture database changes, propagate them to destinations, and apply the changes at these destinations. Oracle Streams is optimized for replicating data and can capture changes locally in the online redo log as it is written or remotely at a downstream database. The captured changes can then be propagated asynchronously to replica databases, or the redo can be sent to the downstream database synchronously or asynchronously and applied in real time by the downstream capture process. These optimizations can reduce latency and enable the replicas to lag the primary database by no more than a few seconds. Unlike Data Guard using SQL Apply, Oracle Streams enables updates on the replica and provides support for heterogeneous platforms with different database releases. Therefore, Oracle Streams may provide the fastest approach for database upgrades and platform migration.	Seconds to minutes to hours	Oracle Streams also has data type limitations and restrictions, such as for advanced queue and object types. But in some cases you can work around them by creating shadow tables on the source database. You can create a trigger on tables with unsupported data types to capture and propagate changes to tables with supported data types. Those changes will be replicated by way of Streams to the target database. You can customize apply to apply the changes to the original tables in the target database. Oracle Streams implementations will require additional investment for setup and configuration because it is designed to be a more flexible architecture. This is the best approach if you are currently using Oracle Streams.
Platform migration to another platform with same endianness	Transportable Database	Transportable Database, a new feature in Oracle Database 10*g* Release 2 (10.2), is the recommended method for migrating an entire database to another platform that has the same endian format.	Varies	The downtime required for a platform migration when using Transportable Database is determined by the time needed to: ■ Place the source database in read-only mode ■ Convert all data files to the new platform format ■ Transfer all data files from the source system to the target system You can significantly minimize this time by using a storage infrastructure that can make the data files available to the target system without the need to physically move the files.

[1]   An example is migration from traditional storage to low cost storage

[2]   For more information, refer to the best practices white papers available at
http://www.oracle.com/technology/deploy/availability/htdocs/maa.htm.

See Also:

- *Oracle Data Guard Concepts and Administration* for more information on using Data Guard with SQL Apply to upgrade an Oracle Database

- *Oracle Database Concepts* and *Oracle Database Administrator's Guide* for more information on transportable tablespace

- The best practices white papers on rolling upgrades at

  http://www.oracle.com/technology/deploy/availability/htdocs/maa.htm

Online Reorganization and Redefinition

One way to enhance availability and manageability is to allow users full access to the database during a data reorganization operation. The Online Reorganization and Redefinition feature in Oracle Database 10*g* offers administrators significant flexibility to modify the physical attributes of a table and transform both data and table structure while allowing users full access to the database. This capability improves data availability, query performance, response time, and disk space utilization—all of which are important in a mission-critical environment and make the application upgrade process easier, safer, and faster.

This online architecture provides the following benefits:

- Tables can be reorganized and redefined online

  - Any physical attribute of the table can be changed online. The table can be moved to a new location, partitioned, and converted from one organization (such as heap-organized) to another (such as index-organized).

  - Many logical attributes can also be changed. Column names, types, and sizes can be changed. Columns can be added, deleted, or merged. One restriction is that the primary key of the table cannot be modified.

- All index operation can be performed online

  - Indexes can be created online and analyzed simultaneously. Online repair of the physical guess component of logical rowids (used in secondary indexes and in the mapping table for index-organized tables) can also be used.

  - An index-organized table and secondary indexes can be reorganized online to eliminate the reorganization maintenance window. Secondary indexes support efficient use of block hints (physical guesses). Invalid physical guesses of logical rowids stored in secondary indexes on index-organized table can also be repaired online.

  - An index-organized table or table partition can be reorganized without rebuilding its secondary indexes, resulting in a short reorganization maintenance window.

- Online move of a partitioned table

- Online reorganization support for advanced queues, clustered tables, materialized views, and abstract data types (objects)

Depending on the type of online reorganization that is required, the following types of data reorganization can be performed using the DBMS_REDEFINITION package or the SQL CREATE/ALTER TABLE and INDEX commands:

- Modify table storage parameters

- Move table to a different tablespace

- Add support for parallel queries

- Add or drop partitioning support

- Recreate table to avoid fragmentation

- Change from table to IOT or from IOT to table

- Add or drop a column

- Transform a column using a function

- Create indexes online

- Rebuild indexes online

- Coalesce indexes online

- Move index-organized tables online

- Copy dependent objects (such as triggers, constraints, and indexes)

- Convert LONG and LONG RAW columns to a LOB

- Use a unique key as an alternative to a primary key or rowid

- Specify columns to order data by

- Change a table without recompiling stored procedures

- Online segment shrink

- Reorganize a single partition

- Reorganize advanced queue and clustered tables

- Reorganize a table containing an ADT

- Retain and clone statistics

- Copies check and not null constraints

- Copies dependent objects for nested tables

> **See Also:** *Oracle Database Administrator's Guide*

## High Availability and Grid Computing

The traditional Oracle database configuration consists of a monolithic server connected to a monolithic storage array. Trends in computing technology are enabling a new IT architecture called Grid computing. Grid computing is a new computing architecture that effectively pools large numbers of servers and storage into a flexible, on-demand computing resource for all enterprise computing needs. Technology innovations such as low-cost blade servers, small and inexpensive multi-processor servers, modular storage technologies, and open source operating systems such as Linux provide the raw materials for the Grid. By harnessing these technologies and leveraging the Grid technology available in Oracle Database 10*g*, enterprises can deliver extremely high quality of service to their users while vastly reducing their expenditures on IT.

Oracle Database 10*g* captures the cost advantages of Grid enterprise computing without sacrificing performance, scalability, security, manageability, functionality, or system availability. A Database Server Grid is a collection of commodity servers connected together to run on one or more databases. A Database Storage Grid is a

collection of low-cost modular storage arrays combined together and accessed by the computers in the Database Server Grid.

Figure 2–2 illustrates the Database Server Grid and Database Storage Grid in a Grid enterprise computing environment.

Figure 2–2   Grid Computing Environment
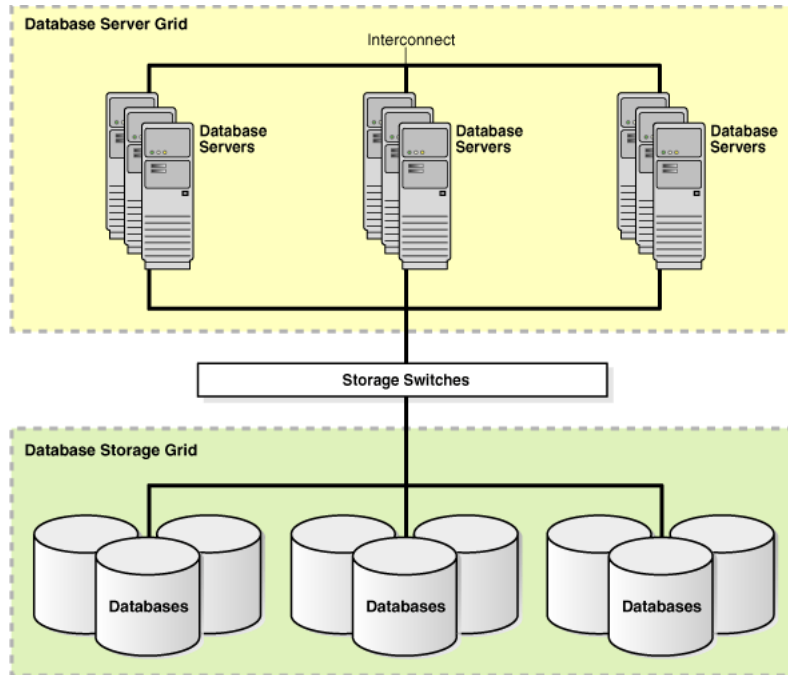
This section covers the following topics:

- Database Server Grid
- Database Storage Grid
- Resilient Low-Cost Storage Initiative

Database Server Grid

The availability of low-cost and reliable blade servers, small multi-processor servers, and inexpensive open-source operating systems such as Linux, has made it possible to build a Database Server Grid that is highly available, scalable, flexible, and manageable.

Oracle Real Application Clusters (RAC) is the technology that enables a Database Server Grid by providing a single database that spans multiple low-cost servers yet appears to the application as a single, unified database system. RAC provides flexibility to dynamically provision resources and services within the Grid as computing needs change, and to add systems to the Grid as capacity demands increase. In addition, RAC provides protection from system failures by automatically recovering the processing of a failed node by any of the surviving systems running the database, and facilitating the reconnection of clients and redistribution of load affected by the failed system.

Database Storage Grid

The availability of low-cost ATA disk-based storage arrays and low-cost storage networks has made it possible to use a Database Storage Grid with the Oracle database at very low cost. A DBA can use the Automatic Storage Management (ASM) interface to specify the disks within the Database Storage Grid that ASM should manage across all server and storage platforms. ASM partitions the disk space and evenly distributes the data storage throughout the entire storage array. Additionally, ASM automatically redistributes the data storage as storage arrays are added or removed from the Database Storage Grid.

When used with Oracle Data Guard, Oracle Flashback, and Recovery Manager (RMAN), the Database Storage Grid can achieve high availability and performance for the Oracle database. Oracle Data Guard protects against site disasters and data corruption by automatically maintaining a standby database. Oracle Flashback protects against human errors using snapshot capabilities, allowing a DBA to rewind a table or the entire database to a specific moment in time before the error occurred. RMAN provides full and incremental backup to disk or tape for archival and protection against data corruption, while the Flash Recovery Area fully manages all disk-based recovery related files and database activities.

Resilient Low-Cost Storage Initiative

To help customers successfully deploy a Database Storage Grid, Oracle has launched the Resilient Low-Cost Storage Initiative. This initiative is analogous to the Oracle Storage Compatibility Program (OSCP) but with a focus on low-cost storage. Its goal is to ensure that low-cost storage performance can be manageable, reliable, and optimal.

This initiative partners Oracle with storage vendors to validate that their low-cost storage arrays meet the minimum price, functionality, and performance requirements for use in a Database Storage Grid. Vendors use Oracle-supplied tools to characterize their storage array performance. After verifying that all requirements can be satisfied, the storage array is certified for this initiative and the vendor—together with Oracle—produces a Best Practice white paper to instruct Oracle customers on how to optimally configure the storage array for a Database Storage Grid environment.

> **See Also:**
>
> - Oracle Resilient Low-Cost Storage Initiative Web site at
>
>   `http://www.oracle.com/technology/deploy/availability`
>   `/htdocs/lowcoststorage.html`
> - Oracle Storage Compatibility Program (OSCP) Web site at
>
>   `http://www.oracle.com/technology/deploy/availability`
>   `/htdocs/oscpf.html`

High Availability Management

Oracle Enterprise Manager Grid Control is an HTML-based user interface that provides the administrator with complete monitoring across the entire Oracle technology stack—business applications, application servers, databases, and the E-Business Suite—as well as non-Oracle components within the Grid. If a component within the Grid becomes unavailable or experiences performance problems, an alert is automatically generated to the Enterprise Manager console to inform the administrator so appropriate action can be taken.

The components of the Oracle Enterprise Manager Grid Control include:

- Oracle Management Service (OMS) - a J2EE Web application that renders the user interface for the Grid Control Console, works with all Management Agents to process monitoring information, and uses the Management Repository as its persistent data store

- Oracle Management Agents - processes deployed on each monitored host to monitor all targets on the host, communicate that information to the Management Service, and maintain the host and its targets

- Oracle Management Repository - schema in the Oracle database that contains all available information about administrators, targets, and applications managed by Enterprise Manager

Communication between the console, the OMS, and Oracle Management Agents is done through HTTP. SSL can also be enabled to allow secure communications between tiers within firewall-protected environments. The Management Agent uploads collected monitoring data to the OMS, which in turn loads the data into the Management Repository. Changes in a target state (such as an availability state change) result in an alert being generated to the Enterprise Manager Console.

Using Oracle Enterprise Manager Grid Control, an administrator can:

- Monitor architecture components and be alerted when a failure occurs

- View overall system status, such as the number of nodes in the database cluster and their current status

- View alerts aggregated across all instances

- Set thresholds for alert generation on a database cluster-wide basis

- Monitor performance metric across all instances

- Perform database cluster-wide operations such as backup and recovery

- Viewing hardware and operating system information in the entire Grid as a whole

- Interconnect monitoring of cluster databases

> **See Also:**
>
> - *Oracle Enterprise Manager Concepts* and *Oracle Enterprise Manager Administrator's Guide* for more information Oracle Enterprise Manager
>
> - The best practices white papers on configuring Enterprise Manager for high availability at
>
>   http://www.oracle.com/technology/deploy/availability
>   /htdocs/maa.htm

3

Determining Your High Availability Requirements

This chapter includes the following topics:

- Why It Is Important to Determine High Availability Requirements
- Analysis Framework for Determining High Availability Requirements
- High Availability Architecture Requirements

Why It Is Important to Determine High Availability Requirements

Since high availability is a critical issue for any modern day enterprise, an enterprise that is designing and implementing a high availability strategy must perform a thorough analysis and have a complete understanding of the business drivers that require high availability. Implementing high availability may involve critical tasks such as:

- Retiring legacy systems
- Investment in more sophisticated and robust systems and facilities
- Redesign of the overall IT architecture to adapt to this high availability model
- Redesign of business processes
- Hiring and training of personnel

Higher degrees of availability reduce downtime significantly. An analysis of the business requirements for high availability and an understanding of the accompanying costs enables an optimal solution that both meets business managers' needs for a highly available system and be within the financial and resource limitations of the business. This chapter provides a simple framework that can be used effectively to evaluate the high availability requirements of a business.

Analysis Framework for Determining High Availability Requirements

The elements of this analysis framework are:

- Business Impact Analysis
- Cost of Downtime
- Recovery Time Objective
- Recovery Point Objective

Business Impact Analysis

A rigorous business impact analysis identifies the critical business processes within an organization, calculates the quantifiable loss risk for unplanned and planned IT outages affecting each of these business processes, and outlines the impacts of these outages. It takes into consideration essential business functions, people and system resources, government regulations, and internal and external business dependencies. This analysis is done using objective and subjective data gathered from interviews with knowledgeable and experienced personnel, reviewing business practice histories, financial reports, IT systems logs, and so on.

The business impact analysis categorizes the business processes based on the severity of the impact of IT-related outages. For example, consider a semiconductor manufacturer, with chip design centers located worldwide. An internal corporate system providing access to human resources, business expenses and internal procurement is not likely to be considered as mission-critical as the internal e-mail system. Any downtime of the e-mail system is likely to severely affect the collaboration and communication capabilities among the global R&D centers, causing unexpected delays in chip manufacturing, which in turn will have a material financial impact on the company.

In a similar fashion, an internal knowledge management system is likely to be considered mission-critical for a management consulting firm because the business of a client-focused company is based on internal research accessibility for its consultants and knowledge workers. The cost of downtime of such a system is extremely high for this business. This leads us to the next element in the high availability requirements framework: cost of downtime.

Cost of Downtime

A well-implemented business impact analysis provides insights into the costs that result from unplanned and planned downtimes of the IT systems supporting the various business processes. Understanding this cost is essential because this has a direct influence on the high availability technology chosen to minimize the downtime risk.

Various reports have been published, documenting the costs of downtime across industry verticals. These costs range from millions of dollars for each hour of brokerage operations and credit card sales, to tens of thousands of dollars for each hour of package shipping services.

While these numbers are staggering, the reasons are quite obvious. The Internet has brought millions of customers directly to the businesses' electronic storefronts. Critical and interdependent business issues such as customer relationships, competitive advantages, legal obligations, industry reputation, and shareholder confidence are even more critical now because of their increased vulnerability to business disruptions.

Recovery Time Objective

A business impact analysis, as well as the calculated cost of downtime, provides insights into the recovery time objective (RTO), an important statistic in business continuity planning. It is defined as the maximum amount of time that an IT-based business process can be down before the organization starts suffering significant material losses. RTO indicates the downtime tolerance of a business process or an organization in general.

The RTO requirements are proportional to the mission-critical nature of the business. Thus, for a system running a stock exchange, the RTO is zero or very near to zero.

An organization is likely to have varying RTO requirements across its various business processes. Thus, for a high volume e-commerce Web site, for which there is an expectation of rapid response times and for which customer switching costs are very low, the Web-based customer interaction system that drives e-commerce sales is likely to have an RTO close to zero. However, the RTO of the systems that support backend operations such as shipping and billing can be higher. If these backend systems are down, then the business may resort to manual operations temporarily without a significantly visible impact.

A systems statistic related to RTO is the network recovery objective (NRO), which indicates the maximum time that network operations can be down for a business. Components of network operations include communication links, routers, name servers, load balancers, and traffic managers. NRO impacts the RTO of the whole organization because individual servers are useless if they cannot be accessed when the network is down.

Recovery Point Objective

Recovery point objective (RPO) is another important statistic for business continuity planning and is calculated through an effective business impact analysis. It is defined as the maximum amount of data an IT-based business process may lose before causing detrimental harm to the organization. RPO indicates the data-loss tolerance of a business process or an organization in general. This data loss is often measured in terms of time, for example, 5 hours or 2 days worth of data loss.

A stock exchange where millions of dollars worth of transactions occur every minute cannot afford to lose any data. Thus, its RPO must be zero. Referring to the e-commerce example, the Web-based sales system does not strictly require an RPO of zero, although a low RPO is essential for customer satisfaction. However, its backend merchandising and inventory update system may have a higher RPO; lost data in this case can be re-entered.

High Availability Architecture Requirements

Using the high availability analysis framework, a business can:

1. Complete a business impact analysis

2. Identify and categorize the critical business processes that have the high availability requirements

3. Formulate the cost of downtime

4. Establish RTO and RPO goals for these various business processes.

This enables the business to define service level agreements (SLAs) in terms of high availability for critical aspects of its business. For example, it can categorize its businesses into several high availability tiers:

- Tier 1 business processes have maximum business impact. They have the most stringent high availability requirements, with RTO and RPO close to zero, and the systems supporting it need to be available on a continuous basis. For a business with a high-volume e-commerce presence, this may be the Web-based customer interaction system.

- Tier 2 business processes can have slightly relaxed high availability and RTO/RPO requirements. The second tier of an e-commerce business may be their supply chain / merchandising systems. For example, these systems do not need to maintain extremely high degrees of availability and may have nonzero RTO/RPO

values. Thus, the high availability systems and technologies chosen to support these two tiers of businesses are likely to be different from those of tier 1 processes.

- Tier 3 business processes may be related to internal development and quality assurance processes. Systems supporting these processes need not have the rigorous high availability requirements of the other tiers.

The next step for the business is to evaluate the capabilities of the various high availability systems and technologies and choose the ones that meet its SLA requirements, within the guidelines as dictated by business performance issues, budgetary constraints, and anticipated business growth.

Figure 3–1 illustrates this process.

Figure 3–1 Planning and Implementing a Highly Available Enterprise
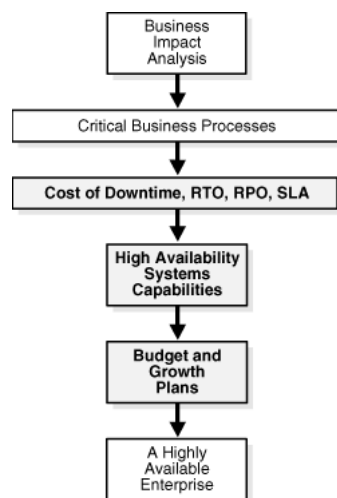
The following sections provide further details about this methodology:

- High Availability Systems Capabilities
- Business Performance, Budget and Growth Plans

> **See Also:** "Choosing the Correct High Availability Architecture" on page 4-9

High Availability Systems Capabilities

A broad range of high availability and business continuity solutions exists today. As the sophistication and scope of these systems increase, they make more of the IT infrastructure, such as the data storage, server, network, applications, and facilities, highly available. They also reduce RTO and RPO from days to hours, or even to minutes and seconds. Increased availability often comes with an increased cost, and on some occasions, with an increased impact on systems performance. Higher availability does not always equate higher cost, however, and the high availability approach to satisfying business requirements may differ for a legacy system.

Organizations need to carefully analyze the capabilities of these high availability systems and map their capabilities to the business requirements to ensure they have an optimal combination of high availability solutions to keep their business running. Consider the business with a significant e-commerce presence as an example.

For this business, the IT infrastructure supporting the system that customers encounter, the core e-commerce engine, needs to be highly available and

disaster-proof. The business may consider clustering for the Web servers, application servers and the database servers serving this e-commerce engine. With built-in redundancy, clustered solutions eliminate single points of failure. Also, modern clustering solutions are application-transparent, provide scalability to accommodate future business growth, and provide load-balancing to handle heavy traffic. Thus, such clustering solutions are ideally suited for mission-critical high-transaction applications.

If unplanned and planned outages occur, the data that supports the high volume e-commerce transactions must be protected adequately and be available with minimal downtime. This data should not only be backed up at regular intervals at the local data centers, but should also be replicated to databases at a remote data center connected over a high-speed, redundant network. This remote data center should be equipped with secondary servers and databases readily available, and be synchronized with the primary servers and databases. This gives the business the capability to switch to these servers at a moment's notice with minimal downtime if there is an outage, instead of waiting for hours and days to rebuild servers and recover data from backed-up tapes. Factors to consider when planning a remote data center include the network bandwidth and latency (distance) between sites, as well as usage consideration (such as whether the sites are fully or partially staffed). These factors should be used to determine whether remote data centers are feasible and their location in relation to the primary data center.

Maintaining synchronized remote data centers is an example where redundancy is built along the entire system's infrastructure. This may be expensive; however, the mission-critical nature of the systems and the data it protects may warrant this expense. Considering another aspect of the business, the high availability requirements are less stringent for systems that gather clickstream data and perform data mining. The cost of downtime is low, and the RTO and RPO requirements for this system could be a few days, because even if this system is down and some data is lost, that will not have a detrimental effect on the business. While the business may need powerful computers to perform data mining, it does not need to mirror this data on a real-time basis. Data protection may be obtained by simply performing regularly scheduled backups, and archiving the tapes for offsite storage.

For this e-commerce business, the back-end merchandising and inventory systems are expected to have higher high availability requirements than the data mining systems, and thus they may employ technologies such as local mirroring or local snapshots, in addition to scheduled backups and offsite archiving.

The business should employ a management infrastructure that performs overall systems management, administration and monitoring, and provides an executive dashboard. This management infrastructure should be highly available and fault-tolerant.

Finally, the overall IT infrastructure for this e-commerce business should be extremely secure, to protect against malicious external and internal electronic attacks.

Business Performance, Budget and Growth Plans

High availability solutions must also be chosen keeping in mind business performance issues. For example, a business may use a zero-data-loss solution that synchronously mirrors every transaction on the primary database to a remote database. However, considering the speed-of-light limitations and the physical limitations associated with a network, there will be round-trip-delays in the network transmission. This delay increases with distance, and varies based on network bandwidth, traffic congestion, router latencies, and so on. Thus, this synchronous mirroring, if performed over large WAN distances, may impact the primary site performance. Online buyers may notice

these system latencies and be frustrated with long system response times; consequently, they may go somewhere else for their purchases. This is an example where the business must make a trade-off between having a zero data loss solution and maximizing system performance.

High availability solutions must also be chosen keeping in mind financial considerations and future growth estimates. It is tempting to build redundancies throughout the IT infrastructure and claim that the infrastructure is completely failure-proof. Although higher availability does not always equate higher cost, going overboard with such solutions may not only lead to budget overruns, it may lead to an unmanageable and unscalable combination of solutions that are extremely complex and expensive to integrate and maintain.

A high availability solution that has very impressive performance benchmark results may look good on paper. However, if an investment is made in such a solution without a careful analysis of how the technology capabilities match the business drivers, then a business may end up with a solution that does not integrate well with the rest of the system infrastructure, has annual integration and maintenance costs that easily exceed the upfront license costs, and forces a vendor lock-in. Cost-conscious and business-savvy CIOs must invest only in solutions that are well-integrated, standards-based, easy to implement, maintain and manage, and have a scalable architecture for accommodating future business growth.

4

High Availability Architectures

This chapter describes high availability architectures in an Oracle environment. It includes the following sections:

- Oracle Database High Availability Architectures
- Choosing the Correct High Availability Architecture
- Assessing Other Architectures

Oracle Database High Availability Architectures

Oracle Database 10*g* provides a full range of capabilities to protect from all causes of system downtime, both planned and unplanned. Table 4–1 shows the outage types and the Oracle database capabilities and features that most effectively prevent, tolerate, or repair each outage type.

Table 4–1   Oracle Database High Availability Architectures

Outage Type	Database Capabilities and Features
Unplanned	
Computer failures	■ Oracle Database 10g with RAC ■ Oracle Database 10g with Data Guard ■ Oracle Database 10g with Streams ■ Fast-Start Fault Recovery
Storage failures	■ Automatic Storage Management ■ Recovery Manager ■ Flash Recovery Area
Human errors	■ Oracle Security Features ■ Oracle Flashback Technology ■ LogMiner
Data corruption	■ Block Checking ■ Block Checksumming ■ Hardware Assisted Resilient Data (HARD) Initiative ■ Oracle Database 10g with Data Guard ■ Oracle Database 10g with Streams ■ Recovery Manager ■ Flash Recovery Area

Table 4–1   (Cont.) Oracle Database High Availability Architectures

Outage Type	Database Capabilities and Features
Site failures	■ Oracle Database 10g with Data Guard
	■ Oracle Database 10g with Streams
	■ Recovery Manager
Planned	
Data changes	■ Online Reorganization and Redefinition
	■ Oracle Database 10g with Data Guard
	■ Oracle Database 10g with Streams
System changes	■ Automatic Storage Management
	■ Dynamic Resource Provisioning
	■ Rolling patch updates and system upgrades using Oracle Database 10g with RAC
	■ Rolling release upgrades and system upgrades using Oracle Database 10g with Data Guard or Oracle Database 10g with Streams
	■ Platform Migrations and Database Upgrades with Transportable Tablespace

This section describes the following top database architectures that address various high availability business needs:

■ Oracle Database 10g

Oracle Database 10*g* running a single database on a standalone computer contains significant high availability features and capabilities. For more information, see Chapter 2, "Oracle Database High Availability Solutions".

■ Oracle Database 10g with RAC

Oracle Real Application Clusters (RAC) builds upon the features and capabilities of Oracle Database 10*g*. RAC comprises several Oracle instances running on many clustered computers that access a shared database residing on shared disk. RAC combines the processing power of these multiple interconnected computers to provide system redundancy, scalability, and high availability. Application scale in a RAC environment to meet increasing data processing demands without changing the application code. In addition, allowing maintenance operations to occur on a subset of components in the cluster while the application continues to run on the rest of the cluster can reduce planned downtime.

■ Oracle Database 10g with Data Guard

Oracle Data Guard builds upon the features and capabilities of Oracle Database 10*g*. Data Guard maintains up to nine standby databases—each of which is a real-time copy of the production database—to protect against all threats: computer failures, storage failures, human errors, data corruption, and site failures. If a failure occurs on the production (primary) database, data processing can fail over to one of the standby databases (which will become the new primary database). In addition, planned downtime for maintenance can be reduced because production processing can quickly and easily switch over from the current primary database to a standby database, and then back again.

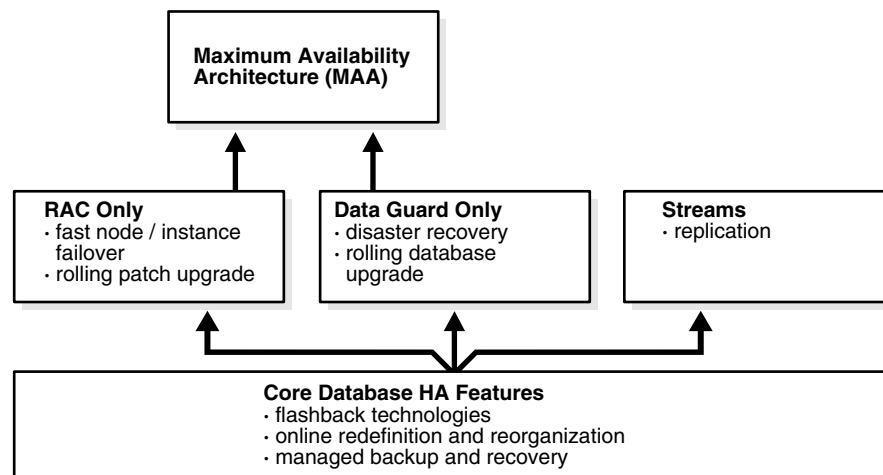■ Oracle Database 10g with RAC and Data Guard - MAA

Maximum Availability Architecture (MAA) combines the scalability and availability advantages of RAC with the site protection capabilities of Data Guard. An MAA environment consists of a site containing a RAC production database, along with a second site containing a cluster that either hosts both logical and physical standby databases, or at least one physical or logical standby database. This architecture provides the most comprehensive set of solutions for both unplanned and planned outages because it inherits the capabilities and advantages of both the Oracle Database 10*g* with RAC and Oracle Database 10*g* with Data Guard architectures.

- Oracle Database 10g with Streams

  Oracle Streams enables the propagation and management of data, transactions, and events in a data stream either within a database, or from one database to another. Similar to Oracle Database 10*g* with Data Guard, Oracle Database with Streams can capture database changes, propagate them to destinations, and apply the changes at these destinations. Using a Streams environment may require additional administrative overhead, but it offers increased flexibility that might be required to meet specific business requirement.

Oracle provides a wide array of high availability architectural solutions. The Oracle Database 10*g* architecture contains many availability features and assets that are used by all other architectures and is the starting point for most customers. Oracle Database 10*g* with RAC, Oracle Database 10*g* with Data Guard, and Oracle Database 10*g* with Streams provide additional high availability capabilities in addition to the Oracle Database 10*g* capabilities. MAA incorporates both RAC and Data Guard advantages and represents the architecture with maximum availability. Choosing an architecture with more availability features does not necessarily lead to higher costs. As a matter of fact, RAC technology and GRID computing enable a more available and resilient architecture to be attained with lower total cost of ownership than most legacy high availability features. Figure 4–1 illustrates the hierarchy of the different high availability architectures.

Figure 4–1   Hierarchy of High Availability Architectures

The following sections provide further details on the various Oracle database high availability architectures:

- Oracle Database 10g

- Oracle Database 10g with RAC

- Oracle Database 10g with Data Guard

- Oracle Database 10g with RAC and Data Guard - MAA

- Oracle Database 10g with Streams

Oracle Database 10*g*

Oracle provides high availability features that can be used in any of the database architectures. These features make the standalone database on a single computer attractive and available:

- Oracle Flashback Technology

- Automatic Storage Management

- Fast-Start Fault Recovery

- Recovery Manager

- Flash Recovery Area

- Online Reorganization and Redefinition

Oracle Database 10*g* with RAC

Oracle Database 10*g* with RAC architecture uses Real Application Clusters and is an inherently high availability system.The clusters that are typical of RAC environments can provide continuous service for both planned and unplanned outages. RAC build higher levels of availability on top of the standard Oracle features. All single instance high availability features, such as flashback technologies and online reorganization, apply to RAC as well.

In addition to the standard Oracle features, RAC exploits the redundancy that is provided by clustering to deliver availability with n - 1 node failures in an n-node cluster. All users have access to all nodes as long as there is one available node in the cluster.

This architecture provides the following benefits:

- Fast node (measured in minutes) and instance failover (measured in seconds)

- Integrated and intelligent connection and service failover across various instances

- Planned node, instance, and service switchover and switchback

- Rolling patch upgrades

- Rolling release upgrades of Oracle Clusterware

- Multiple active instance availability and scalability across multiple nodes

- Comprehensive manageability that integrates database and cluster features

- Extensive cluster and application services that allows the database and application services to be restarted or relocated in case of failures

Figure 4–2 shows Oracle Database 10*g* with RAC architecture.

Figure 4–2   Oracle Database 10g with RAC Architecture
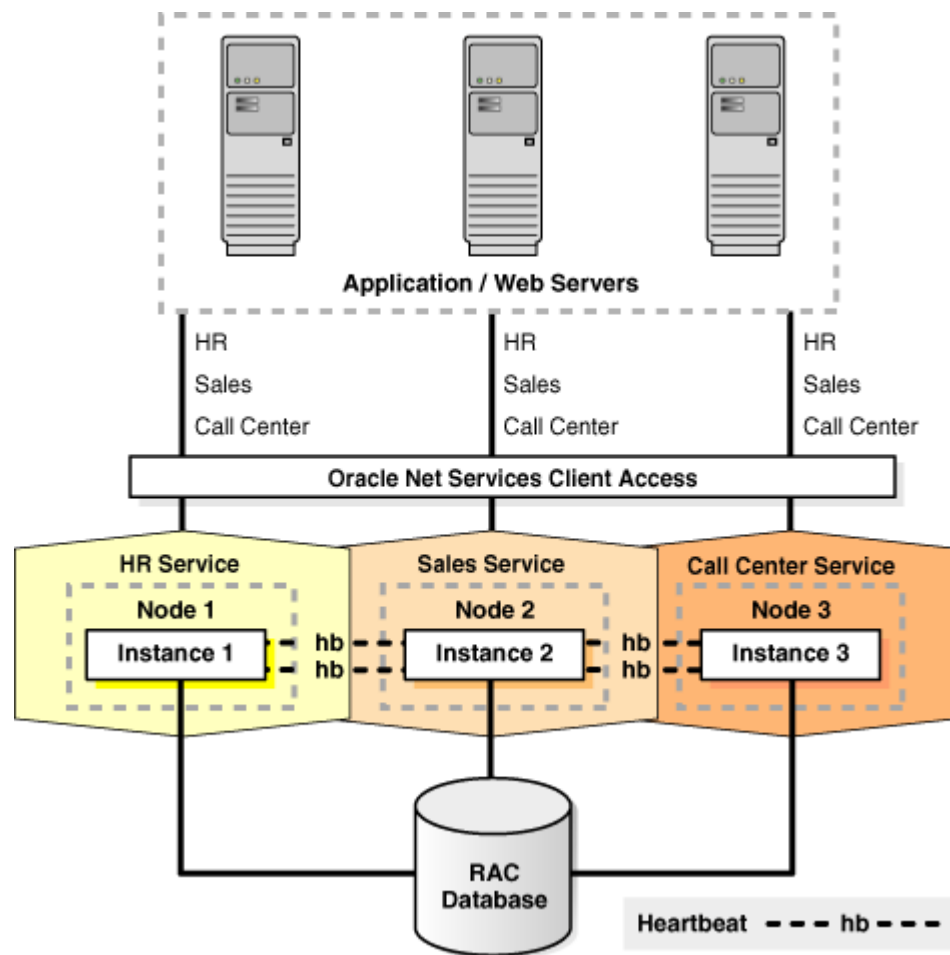
Oracle Database 10g with Data Guard

Oracle Data Guard ensures high availability, data protection, and disaster recovery for enterprise data. Data Guard provides a comprehensive set of services that create, maintain, manage, and monitor one or more standby databases to enable Oracle databases to survive disasters and data corruption. Data Guard maintains these standby databases as transactionally consistent copies of the production database. If the production database becomes unavailable due to a planned or unplanned outage, Data Guard can switch any standby database to the production role, minimizing the downtime associated with the outage. Data Guard can be used with traditional backup, restoration, and cluster technologies to provide a high level of data protection and availability. With Data Guard, administrators can optionally improve production database performance by diverting resource-intensive backup and reporting operations to standby systems.

Using a backup copy of the primary database, it is possible to create up to nine standby databases and integrate them in a Data Guard configuration. Once created, Data Guard automatically maintains each standby database by transmitting redo data from the primary database and applying the redo to the standby database. Similar to a primary database, a standby database can be either an Oracle single-instance or RAC database.

A standby database can be either a physical standby database or a logical standby database. A physical standby database provides a physically identical copy of the primary database, with on disk database structures that are identical to the primary database on a block-for-block basis. A physical standby database is synchronized with the primary database through Redo Apply, which recovers the redo data received from the primary database and applies it to the physical standby database. A physical standby database can be used for business purposes other than disaster recovery on a limited basis.

Physical standby databases provide these advantages:

- Protection from user errors and logical corruption

- Protection from disasters and site failures if located remotely

- Fast site and database failover (less than 1 minute to five minutes)

- Fast-start failover provides the ability to automatically, quickly, and reliably fail over to a designated, synchronized standby database in the event of primary database failure

- Fast site and database planned switchovers for maintenance

- Using Flashback Database, a Redo Apply standby database can diverge for reporting or testing purposes and resynchronize with its primary database once complete

- Backups can be taken from the physical standby database instead of the production database, relieving the load on the production database

- Read-only capability, resulting in better use of system resources

- Greater support for fast application notification and application callouts resulting in better full-stack application failover

A logical standby database can be used for other business purposes in addition to disaster recovery. Users can access a logical standby database for queries and reporting purposes. Using a logical standby database, it is possible to upgrade Oracle database software and patch sets with minimal downtime. Therefore, a logical standby database can be concurrently used for data protection, reporting, and database upgrade purposes.

In addition to disaster recovery and data protection, logical standby databases provide the following benefits:

- Enable the standby database to be open for normal operations with both read-only and read/write accessibility

- Enable additional objects to be built and maintained

- Enable rolling database upgrades of the production database

A recommended configuration for many cases includes both physical and logical standby databases. They can reside on the same database computer or cluster, but they should be remote from the production database. The physical standby database can be reserved for failovers in case of disaster, and the logical standby database can continue to be used for reporting. The physical standby database provides a faster apply technology because redo logs do not have to be converted to SQL.

Figure 4–3 shows the production database at the primary site and the standby databases at the secondary site.

Figure 4–3   Oracle Database 10g with Data Guard Architecture on Primary and Secondary Sites
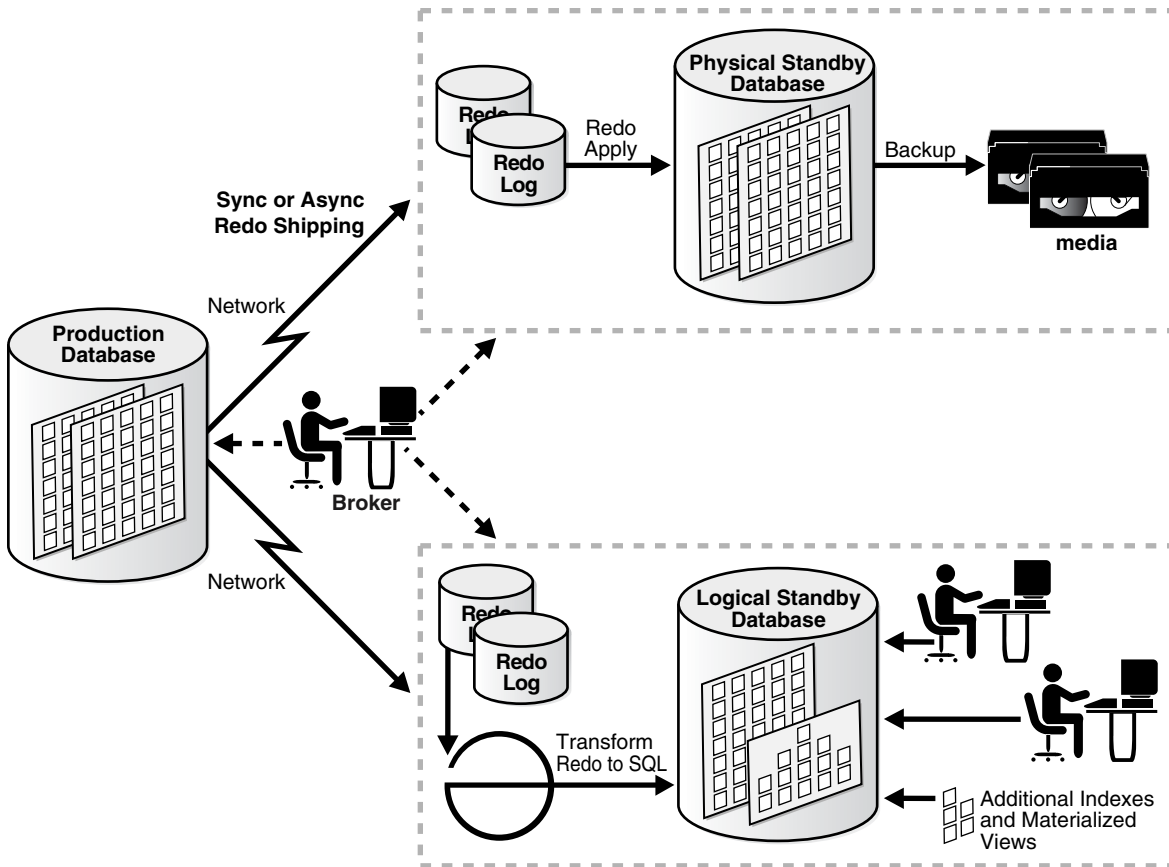
See Also:

- *Oracle Data Guard Concepts and Administration* for more information about datatypes supported by logical standby databases

- The papers about standby databases at

  http://www.oracle.com/technology/deploy/availability/htdocs/maa.htm

Oracle Database 10*g* with RAC and Data Guard - MAA

RAC and Data Guard provide the basis of Oracle Database 10*g* - MAA. Maximum Availability Architecture (MAA) provides the most comprehensive architecture for reducing downtime for scheduled outages and preventing, detecting, and recovering from unscheduled outages. The recommended MAA has two identical sites. The primary site contains the RAC primary database, and the secondary site contains a RAC standby database.

Identical site configuration is recommended to ensure that performance is not sacrificed after a failover or switchover. Symmetric sites also enable processes and procedures to be kept the same between sites, making operational tasks easier to maintain and execute.

MAA encompasses RAC, Data Guard, and a set of recommended best practices for configuring and managing the architecture as well as recovering from various outages. For more information, visit the MAA Web site at:

http://www.oracle.com/technology/deploy/availability/htdocs/maa.htm

Figure 4–4 provides an overview of Oracle Database 10*g* with RAC and Data Guard - MAA.

Figure 4–4   Oracle Database 10g with RAC and Data Guard - MAA
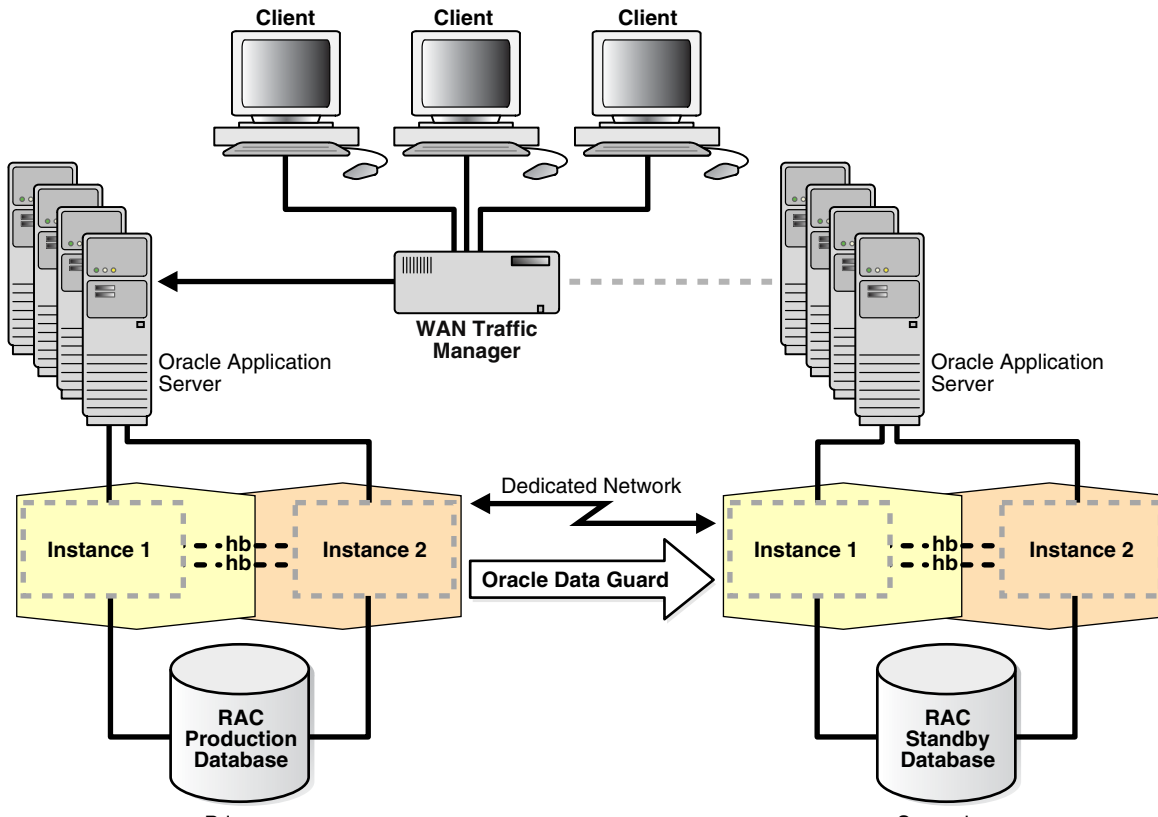
Oracle Database 10*g* with Streams

Oracle Streams is meant for information sharing and distribution. It can also provide an efficient and highly available architecture.

Oracle Database 10*g* with Streams provides granularity and control over what is replicated and how it is replicated. It supports bidirectional replication, data transformations, subsetting, custom apply functions, and heterogeneous platforms. It also gives users complete control over the routing of change records from the primary database to a replica database. The capture of data changes can be performed at the primary database or downstream at a replica database. This enables users to build hub and spoke network configurations that can support hundreds of replica databases.

Oracle Database 10*g* with Streams should be evaluated if one or more of the following conditions are true:

- A full active/active site configuration is required, including bidirectional changes

- Site configurations are on heterogeneous platforms

- Different character sets are required between the primary database and its replicas

- Fine control of information and data sharing are required

- More investment and expertise to build and maintain an integrated high availability solution is available

For disaster recovery, Data Guard is Oracle's recommended solution.

Figure 4–5 shows Oracle Database 10*g* with Streams with local capture running at the primary database.

Figure 4–5   Oracle Database 10g with Streams
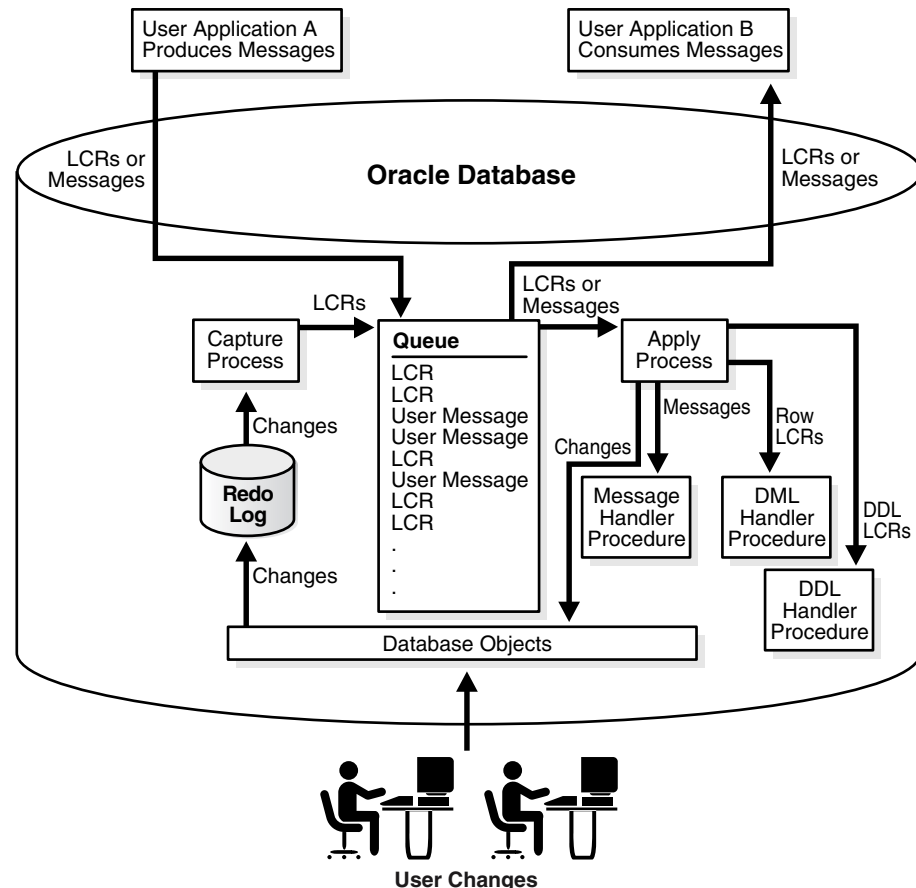
Choosing the Correct High Availability Architecture

This section summarizes the advantages of the high availability architectures discussed in this chapter and provides guidelines for you to choose the correct high availability architecture for your business.

Oracle Database 10*g* with RAC and Oracle Database 10*g* with Data Guard are the most common Oracle high availability architectures, and each provides very significant high availability advantages. MAA provides the most redundant and robust high availability architecture. It prevents, detects, and recovers from different outages to meet stringent RTO and RPO requirements, as well as preventing or minimizing downtime for maintenance. Oracle Database 10*g* with Streams is an alternative high availability solution, but it requires more customization and administrative effort, and may not be as transparent to the application.

The baseline high availability architecture is Oracle Database 10*g*. Consider using:

- Oracle Database 10*g* with RAC if:

- Maximum Recovery Time Objective (RTO) for instance or node failure is in seconds or minutes

- Database scalability beyond one instance or node is required

■ Oracle Database 10*g* with Data Guard if:

- Maximum RTO for instance or node failure is in seconds to minutes or more

- Maximum RTO for data corruption or site failure is less than 1 minute to five minutes

■ MAA if:

- Projected planned maintenance is in hours or less for each year

- Both RAC and Data Guard are required

■ Oracle Database 10*g* with Steams if active/active replicated system or heterogeneous solution is required

Table 4–2 identifies the additional capabilities provided by the architectures that build upon Oracle Database 10*g*.

Table 4–2   Additional Capabilities of High Level Oracle High Availability Architectures

Oracle High Availability Architecture	Key Characteristics and Additional Capabilities
Oracle Database 10g with RAC	Transparent to application Fast repair for human error Fast failover for computer failure and storage failure Scalability beyond a single system Reduced downtime for computer maintenance
Oracle Database 10g with Data Guard	Transparent to application Fast repair for human error Fast failover for computer failure, storage failure, and data corruption Protection from site failure Reduced downtime for computer or site maintenance
Oracle Database 10g with RAC and Data Guard - MAA	Transparent to application Fast repair for human error Fast failover for computer failure, storage failure, and data corruption Protection from site failure Scalability beyond a single system Reduced downtime for computer or site maintenance
Oracle Database 10g with Streams[1]	Fast repair for human error Replica database(s) available for read/write use Provides heterogeneous platform support Fast failover for computer failure and storage failure Protection from site failure Reduced downtime for computer or site maintenance

[1]   Requires planning and overhead to make solution robust

Table 4–3 shows the attainable recovery times for all types of unplanned downtime for each Oracle high availability architecture.

Table 4–3   Attainable Recovery Times for Unplanned Outages

Outage Type	Oracle Database 10*g*	RAC	Data Guard	MAA	Streams
Computer failure	Minutes to hours[1]	No downtime[2]	Seconds to five minutes	No downtime	No downtime
Storage failure	No downtime[3]	No downtime[3]	No downtime[3]	No downtime[3]	No downtime[3]
Human error	< 30 minutes[4]	< 30 minutes[4]	< 30 minutes[4]	< 30 minutes[4]	< 30 minutes[4]
Data corruption	HARD prevents data corruption[5] Potentially hours[6]	HARD prevents data corruption[5] Potentially hours[6]	HARD prevents data corruption[5] Seconds to five minutes	HARD prevents data corruption[5] Seconds to five minutes	HARD prevents data corruption[5] Seconds to five minutes
Site failure	Hours to days	Hours to days	Seconds to five minutes[7]	Seconds to five minutes[8]	No downtime[7]

[1] Recovery time consists largely of the time it takes to restore the failed system.

[2] Database is still available, but portion of application connected to failed system is temporarily affected.

[3] Storage failures are prevented by using ASM with mirroring and its automatic rebalance capability.

[4] Recovery time for human errors depend primarily on detection time. If it takes seconds to detect a malicious DML or DDL transaction, it typically only requires seconds to flashback the appropriate transactions. Longer detection time usually leads to longer recovery time required to repair the appropriate transactions. An exception is undropping a table, which is literally instantaneous regardless of detection time.

[5] Not all types of data corruption are prevented. For the most recent information about the HARD initiative, refer to http://www.oracle.com/technology/deploy/availability/htdocs/HARD.html.

[6] Recovery time depends on the age of the backup used for recovery and the number of log changes scanned to make the corrupt data consistent with the database.

[7] Recovery time indicated applies to database and existing connection failover. Network connection changes and other site-specific failover activities may lengthen overall recovery time.

[8] The portion of any application connected to the failed system is temporarily affected.

Table 4–4 shows the attainable recovery times for all types of planned downtime for each Oracle high availability architecture.

Table 4–4   Attainable Recovery Times for Planned Outages

Outage Type	Oracle Database 10*g*	RAC	Data Guard	MAA	Streams
System changes - Dynamic Resource Provisioning	No downtime	No downtime	No downtime	No downtime	No downtime

Table 4–4   (Cont.) Attainable Recovery Times for Planned Outages

Outage Type		Oracle Database 10*g*	RAC	Data Guard	MAA	Streams
System changes - Rolling Upgrades	System level upgrade	Minutes to hours	No downtime	Seconds to five minutes	No downtime	No downtime
	Cluster or site wide upgrade	Minutes to hours	Minutes to hours	Seconds to five minutes	Seconds to five minutes	No downtime[1]
	Storage migration	No downtime[2]	No downtime[2]	No downtime[2]	No downtime[2]	No downtime[2]
	Database one-off patch	Minutes to hours	No downtime[3]	Seconds to five minutes	No downtime[3]	No downtime
	Database patchset and version upgrade	Minutes to hours	Minutes to hours	Seconds to five minutes	Seconds to five minutes	No downtime[1]
	Platform migration	Minutes to hours	Minutes to hours	Minutes to hours	Minutes to hours	No downtime[1]
Data changes - Online Reorganization and Redefinition		No downtime	No downtime	No downtime[4]	No downtime[4]	No downtime[4]

[1] The portion of any application connected to the failed system is temporarily affected.

[2] ASM automatically rebalances stored data when disks are added or removed while the database remains online. For storage migration, you will require both storage arrays to be leveraged by ASM temporarily.

[3] For qualified one-off patches only

[4] Tables can be reorganized online using the `DBMS_REDEFINITION` package. However, the online changes are not supported by SQL Apply or data capture, and therefore the effects of this subprogram are not visible on the logical standby or replica database. For more information, see *Oracle Data Guard Concepts and Administration* or *Oracle Streams Replication Administrator's Guide*.

Assessing Other Architectures

There are other Oracle and non-Oracle high availability and enterprise computing architectures. This section focuses on the most common variants.

Table 4–5 describes common alternative high availability architecture, their disadvantages, and the recommended Oracle high availability architectures.

Table 4–5   Comparison of High Availability Architectures

Alternative Architecture	Disadvantages	Recommended Oracle Architecture
Single instance database on hardware cluster	▪ Node failure requires restart, storage remastering, and reconnect for all application connectors ▪ No protection from disaster and site failure ▪ No protection against data corruption beyond Oracle Database 10*g* capabilities ▪ Limited ability to reduce downtime for system rolling upgrades ▪ Inability to reduce downtime for Oracle upgrades ▪ Under-utilized hardware resources ▪ No database server scalability beyond one node ▪ Inability to offload database activities such as backup or reporting	1. Oracle Database 10g with RAC 2. Oracle Database 10g with Data Guard 3. Oracle Database 10g with RAC and Data Guard - MAA
Remote mirrored single instance database	▪ High network utilization ▪ No protection against data corruption beyond Oracle Database 10*g* capabilities ▪ Site failure requires instance restart, storage remastering, and reconnect for all applications connections ▪ No database server scalability beyond primary node(s) ▪ Inability to offload database activities such as backup or reporting ▪ Inability to reduce downtime for rolling upgrades ▪ Customization required	1. Oracle Database 10g with Data Guard 2. Oracle Database 10g with RAC and Data Guard - MAA
RAC database in a stretch cluster configuration	▪ No protection against data corruption beyond Oracle Database 10*g* capabilities ▪ Limited protection against site failures with regional impact ▪ Limited ability to reduce downtime for rolling upgrades ▪ High network utilization ▪ Limited by distance between nodes in cluster ▪ As network latency increases, application performance can be impacted significantly	1. Oracle Database 10g with Data Guard 2. Oracle Database 10g with RAC and Data Guard - MAA
RAC database with standby database on same site	▪ No protection from site failures	Oracle Database 10g with RAC and Data Guard - MAA
Single instance database with standby database on same site	▪ No protection from site failures ▪ Limited ability to reduce downtime for rolling upgrades	Oracle Database 10g with RAC and Data Guard - MAA

Table 4–6 describes common traditional enterprise computing architectures, their disadvantages, and the recommended Oracle enterprise computing architectures.

Table 4–6   Comparison of Enterprise Computing Architectures

Traditional Architecture	Disadvantages	Recommended Architecture
Monolithic database server	■ High cost ■ Fixed scalability ■ Not flexible to changing capacity and resource demands	Database Server Grid
Monolithic storage array	■ High cost ■ Fixed scalability ■ Not flexible to changing capacity and resource demands	Database Storage Grid

See Also:

■ "High Availability and Grid Computing" on page 2-21

■ Oracle Resilient Low-Cost Storage Initiative Web site at

  `http://www.oracle.com/technology/deploy/availability`
  `/htdocs/lowcoststorage.html`

5

High Availability Best Practices

Choosing and implementing the architecture that best fits the availability requirements of a business can be a daunting task. This architecture must encompass appropriate redundancy, provide adequate protection from all types of outages, ensure consistent high performance and robust security, while being easy to deploy, manage, and scale. Needless to mention, this architecture should be driven by well-understood business requirements.

To build, implement and maintain such an architecture, a business needs high availability best practices that involve both technical and operational aspects of its IT systems and business processes. Such a set of best practices removes the complexity of designing a high availability architecture, maximizes availability while using minimum system resources, reduces the implementation and maintenance costs of the high availability systems in place, and makes it easy to duplicate the high availability architecture in other areas of the business.

An enterprise with a well-articulated set of high availability best practices that encompass high availability analysis frameworks, business drivers and system capabilities, will enjoy an improved operational resilience and enhanced business agility. Oracle offers various white papers that provide technical details on its various high availability technologies, along with best practice recommendations on configuring and using such technologies. Oracle high availability white papers can be downloaded at

http://www.oracle.com/technology/deploy/availability/htdocs/maa.htm

Index