

Oracle® Database
High Availability Overview
11g Release 1 (11.1)
B28281-01

July 2007

Oracle Database High Availability Overview, 11g Release 1 (11.1)

B28281-01

Copyright © 2005, 2007, Oracle. All rights reserved.

Contributors: Andrew Babb, Tammy Bednar, Larry Carpenter, Immanuel Chan, Ray Dutcher, Craig Foch, Frank Kobylanski, Barb Lundhild, Rahim Mau, Joe Meeks, Valarie Moore, Michael Nowak, Ashish Ray, Viv Schupmann, Michael T. Smith, Lawrence To, Douglas Utzig, James Viscusi, Shari Yamaguchi

The Programs (which include both the software and documentation) contain proprietary information; they are provided under a license agreement containing restrictions on use and disclosure and are also protected by copyright, patent, and other intellectual and industrial property laws. Reverse engineering, disassembly, or decompilation of the Programs, except to the extent required to obtain interoperability with other independently created software or as specified by law, is prohibited.

The information contained in this document is subject to change without notice. If you find any problems in the documentation, please report them to us in writing. This document is not warranted to be error-free. Except as may be expressly permitted in your license agreement for these Programs, no part of these Programs may be reproduced or transmitted in any form or by any means, electronic or mechanical, for any purpose.

If the Programs are delivered to the United States Government or anyone licensing or using the Programs on behalf of the United States Government, the following notice is applicable:

U.S. GOVERNMENT RIGHTS Programs, software, databases, and related documentation and technical data delivered to U.S. Government customers are "commercial computer software" or "commercial technical data" pursuant to the applicable Federal Acquisition Regulation and agency-specific supplemental regulations. As such, use, duplication, disclosure, modification, and adaptation of the Programs, including documentation and technical data, shall be subject to the licensing restrictions set forth in the applicable Oracle license agreement, and, to the extent applicable, the additional rights set forth in FAR 52.227-19, Commercial Computer Software--Restricted Rights (June 1987). Oracle USA, Inc., 500 Oracle Parkway, Redwood City, CA 94065.

The Programs are not intended for use in any nuclear, aviation, mass transit, medical, or other inherently dangerous applications. It shall be the licensee's responsibility to take all appropriate fail-safe, backup, redundancy and other measures to ensure the safe use of such applications if the Programs are used for such purposes, and we disclaim liability for any damages caused by such use of the Programs.

Oracle, JD Edwards, PeopleSoft, and Siebel are registered trademarks of Oracle Corporation and/or its affiliates. Other names may be trademarks of their respective owners.

The Programs may provide links to Web sites and access to content, products, and services from third parties. Oracle is not responsible for the availability of, or any content provided on, third-party Web sites. You bear all risks associated with the use of such content. If you choose to purchase any products or services from a third party, the relationship is directly between you and the third party. Oracle is not responsible for: (a) the quality of third-party products or services; or (b) fulfilling any of the terms of the agreement with the third party, including delivery of products or services and warranty obligations related to purchased products or services. Oracle is not responsible for any loss or damage of any sort that you may incur from dealing with any third party.

Contents

Preface	vii
Audience	vii
Documentation Accessibility	vii
Related Documents	viii
Conventions	viii
What's New in High Availability?	ix
1 Overview of High Availability	
1.1 Introduction to High Availability	1-1
1.2 What is Availability?	1-1
1.3 Importance of Availability	1-2
1.4 Causes of Downtime	1-3
1.5 What Does This Book Contain?	1-6
2 Oracle Database High Availability Features and Products	
2.1 Oracle High Availability Features and Solutions for Unplanned Downtime	2-1
2.1.1 Fast-Start Fault Recovery	2-2
2.1.2 Oracle Real Application Clusters and Oracle Clusterware	2-2
2.1.3 Oracle Data Guard	2-4
2.1.4 Oracle Streams	2-6
2.1.5 Transportable Technologies	2-9
2.1.6 Oracle Flashback Technology	2-10
2.1.7 Automatic Storage Management	2-13
2.1.8 Recovery Manager	2-14
2.1.9 Oracle Secure Backup	2-15
2.1.10 Data Recovery Advisor	2-16
2.1.11 Flash Recovery Area	2-17
2.1.12 Oracle Security Features	2-17
2.1.13 LogMiner	2-18
2.1.14 Hardware Assisted Resilient Data (HARD) Initiative	2-19
2.1.15 Hang Manager	2-19
2.1.16 Data Block Corruption Prevention and Detection Parameters	2-20
2.1.17 Oracle HA Solutions and Recovery Times for Unplanned Downtime	2-20
2.2 Oracle High Availability Features and Solutions for Planned Downtime	2-22

2.2.1	Dynamic Resource Provisioning	2-23
2.2.1.1	Dynamic Reconfiguration of the Database	2-23
2.2.1.2	Autotuning Memory Management	2-23
2.2.1.3	Automated Distribution of Datafiles, Control Files, and Log Files.....	2-24
2.2.2	Oracle High Availability Solutions and Recovery Times for Planned Downtime..	2-24
2.2.2.1	Avoiding Downtime During Operating System and Hardware Upgrades	2-25
2.2.2.2	Using Oracle Data Guard for System and Cluster Upgrades and Migrations.	2-26
2.2.2.3	Oracle Interim Database Patches.....	2-27
2.2.2.4	Online Patching	2-28
2.2.2.5	Upgrading Oracle Clusterware	2-28
2.2.2.6	Upgrading Automatic Storage Management (ASM)	2-28
2.2.2.7	Storage Migration	2-29
2.2.2.8	Patch Set and Database Upgrades.....	2-29
2.2.2.9	Platform Migration Across Same Endian Format Platforms	2-30
2.2.2.10	Platform Migration Across Different Endian Format Platforms	2-32
2.2.3	Online Reorganization and Redefinition	2-33
2.2.4	Online Application Maintenance and Upgrades	2-35
2.2.4.1	Oracle Streams for Rolling Upgrades	2-35
2.2.4.2	DDL With the WAIT Option.....	2-35
2.2.4.3	ENABLE, DISABLE and FOLLOWS Clauses for CREATE TRIGGER.....	2-35
2.2.4.4	Enhanced ADD COLUMN Functionality	2-35
2.2.4.5	Finer Grained Dependencies	2-36
2.2.4.6	Invisible Indexes	2-36
2.2.4.7	Materialized View Logging Control.....	2-36
2.2.4.8	Dependent PL/SQL Recompilation After Online Table Redefinition.....	2-36
2.3	Optimizing Grid Computing and Disaster Recovery Solutions	2-37
2.3.1	Grid Computing.....	2-37
2.3.2	Database Server Grid	2-38
2.3.3	Database Storage Grid	2-38
2.3.4	Disaster Recovery Solutions with Better Standby Database Usage	2-38
2.3.4.1	Real-Time Query Physical Standby Databases	2-39
2.3.4.2	Web Scale Using Standby Reader Farms	2-40
2.4	Optimizing Manageability.....	2-41
2.4.1	Intelligent Infrastructure.....	2-42
2.4.2	Change Assurance	2-43
2.4.3	Oracle Enterprise Manager Grid Control.....	2-43

3 Determining Your High Availability Requirements

3.1	Why It Is Important to Determine High Availability Requirements	3-1
3.2	Analysis Framework for Determining High Availability Requirements	3-1
3.2.1	Business Impact Analysis	3-2
3.2.2	Cost of Downtime.....	3-2
3.2.3	Recovery Time Objective	3-2
3.2.4	Recovery Point Objective.....	3-3
3.2.5	Manageability Goal	3-3
3.2.6	Total Cost of Ownership (TCO) and Return On Investment (ROI)	3-3
3.3	High Availability Architecture Requirements.....	3-4

3.3.1	High Availability Systems Capabilities.....	3-5
3.3.2	Business Performance, Budget, and Growth Plans	3-6

4 High-Availability Architectures and Solutions

4.1	Oracle Database High-Availability Architectures	4-1
4.1.1	Oracle Database	4-2
4.1.2	Oracle Database with Oracle Clusterware (Cold Cluster Failover)	4-3
4.1.3	Oracle Database with Oracle Real Application Clusters	4-6
4.1.4	Oracle Database with Oracle RAC on Extended Distance Clusters	4-8
4.1.5	Oracle Database with Data Guard	4-12
4.1.5.1	Overview of Single Standby Database Architectures	4-13
4.1.5.2	Overview of Multiple Standby Database Architectures	4-15
4.1.6	Oracle Database with Oracle Clusterware and Data Guard	4-19
4.1.7	Oracle Database with Oracle RAC and Data Guard	4-20
4.1.8	Oracle Database with Streams	4-21
4.2	Choosing the Correct High-Availability Architecture	4-24
4.3	Integrating Application Server High Availability	4-30
4.3.1	Oracle Application Server High Availability Architectures	4-30
4.3.2	Redundant Architectures.....	4-31
4.3.3	High Availability Services in Oracle Application Server	4-31
4.4	Integrating High Availability for All Applications.....	4-31

5 MAA and High Availability Best Practices

Index

Preface

This book introduces you to Oracle best practices for deploying 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 2 Day + Real Application Clusters Guide*
- *Oracle Clusterware Administration and Deployment Guide*
- *Oracle Real Application Clusters Administration and Deployment Guide*
- *Oracle Database Backup and Recovery 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 about how these schemas were created and how you can use them yourself.

Oracle High Availability Best Practice white papers can be downloaded from

<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.
<i>italic</i>	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.

What's New in High Availability?

Oracle Database 11g Release 1 (11.1) offers an integrated suite of high availability (HA) solutions and Maximum Availability Architecture (MAA) best practices that provide:

- High availability to eliminate or minimize both planned and unplanned downtime
- Unmatched data protection
- Flexible deployment options that enhance overall performance and scalability
- Reduced cost using low cost platforms, without compromising HA
- Maximum return on investment by fully using all HA components for productive work at all times

For an overview of all of the high availability features available with the Oracle Database, see the "High Availability" chapter in *Oracle Database Concepts*.

For a list of new high availability features introduced in Oracle Database 11g Release 1 (11.1), see the "Availability" section in Chapter 1 of the *Oracle Database New Features Guide*.

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?](#)

1.1 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 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 availability to your users.

1.2 What is Availability?

Availability is the degree to which an application, service, or functionality is available upon demand. Availability is measured by the perception of an application's end user. End users experience frustration when their data is unavailable or not performing within certain expectations, 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. A related concept is Resilience. For example, low-cost commodity hardware, combined with software functionality such as Oracle RAC, can be used to implement a very reliable system, because the *resilience* of an Oracle RAC database allows processing to continue even though individual servers may fail.
- **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:** Provide the ability for continuous access to your data. This 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:

- Tolerate failures such that processing continues with minimal or no interruption
- Be transparent to—or tolerant of—system, data, or application changes
- Provide built-in preventative measures
- Provide proactive monitoring and fast detection of failures
- Provide fast recoverability
- Automate the detection and recovery operation
- Protect the data so that there is minimal or no data loss
- Implement the operational best practices to manage your environment
- Ability to achieve SLAs (for example, Recovery Time and Recovery Point) for the lowest possible total cost of ownership.

1.3 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.

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. Alternatively, frequent outages, even if short in duration, may similarly disrupt business operations. When designing a solution, it is important to understand the true cost of downtime to understand how the business can benefit by availability improvements.

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.

1.4 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 unplanned outage types and provides examples of each type.

Table 1–1 Causes of Unplanned Downtime

Type	Description	Examples
Computer failure	A computer failure outage occurs when the system running the database becomes unavailable because it has crashed or is no longer accessible.	<ul style="list-style-type: none"> ■ 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.	<ul style="list-style-type: none"> ■ 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.	<ul style="list-style-type: none"> ■ File deletion (at the file system level). Dropped database object ■ Inadvertent data changes ■ Malicious data changes

Table 1–1 (Cont.) Causes of Unplanned Downtime

Type	Description	Examples
Data corruption	<p>A corrupt block is a block that has been changed so that it differs from what Oracle Database expects to find. Block corruptions fall under two categories: physical and logical block corruptions. In a physical corruption, which is also called a media corruption, the database does not recognize the block at all: the checksum is invalid, the block contains all zeros, or the header and footer of the block do not match. In a logical corruption, the contents of the block are logically inconsistent. Examples of logical corruption include corruption of a row piece or index entry.</p> <p>Block corruptions can also be divided into interblock corruption and intrablock corruption. In intrablock corruption, the corruption occurs within the block itself and can be either physical or logical corruption. In an interblock corruption, the corruption occurs between blocks and can only be logical corruption.</p> <p>A data corruption outage occurs when a hardware, software or network component causes corrupt data to be read or written. 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).</p>	<ul style="list-style-type: none"> ■ Operating system or storage device driver ■ Host bus adapter ■ Disk controller ■ Volume manager error causing bad disk read or writes ■ Software defects

Table 1–1 (Cont.) Causes of Unplanned Downtime

Type	Description	Examples
Lost Writes	<p>A lost write is another form of data corruption, but it is much more evasive to detect and repair quickly. A data block lost write occurs when:</p> <ul style="list-style-type: none"> ■ An I/O subsystem acknowledges the completion of the block write, while in fact the write did not occur in the persistent storage. On a subsequent block read on the primary database, the I/O subsystem returns the stale version of the data block, which might be used to update other blocks of the database, thereby corrupting it. ■ The write I/O completed but it was written somewhere else, and a subsequent read operation returns the stale value. ■ A read I/O from one cluster node returns stale data after a write I/O on another node. For example, this could occur if an NFS caching policy is incompatible with Oracle RAC. 	<ul style="list-style-type: none"> ■ Operating system or storage device driver ■ Host bus adapter ■ Disk controller ■ Volume manager error ■ Other application software ■ NFS write visibility across cluster
Hang or slow down	<p>Hang or slow down occurs when the database or application are unable to process transactions because of a resource or lock contention. Perceived hang can be caused by lack of system resources.</p>	<ul style="list-style-type: none"> ■ Database or application deadlocks ■ Runaway processes that consume system resources ■ Log on storms or system faults ■ Combination of application peaks with lack of system or database resources ■ Archive log destination or flash recovery area destinations become full
Site failure	<p>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.</p>	<ul style="list-style-type: none"> ■ 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

Table 1–2 describes planned outage types and provides examples of each types.

Table 1–2 Causes of Planned Downtime

Type	Description	Examples
System and database changes	<p>Planned system changes occur when performing routine and periodic maintenance operations and new deployments.</p> <p>Planned system changes include any scheduled changes to the operating environment that occur outside the organizational data structure within the database.</p> <p>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.</p>	<ul style="list-style-type: none"> ■ 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 ■ Moving from 32 bits to 64 bits ■ Migrating to cluster architecture ■ Migrating to new storage
Data changes	<p>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.</p>	<ul style="list-style-type: none"> ■ Table definition changes ■ Adding table partitioning ■ Creating and rebuilding indexes
Application Changes	<p>Planned application changes may include data changes as well as schema and programmatic changes. The primary objective of these changes is to improve performance, manageability, and functionality.</p>	<ul style="list-style-type: none"> ■ Application upgrades

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.

1.5 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 and tolerance from computer failures, storage failures, human errors, data corruption, lost writes, system hangs or slowdown, 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
- Achieve SLA's at the lowest possible total cost of ownership

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 and Solutions"](#)

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

- [Chapter 2, "Oracle Database High Availability Features and Products"](#)
- [Chapter 4, "High-Availability Architectures and Solutions"](#)

Oracle High Availability Best Practice recommendations can be found in the *Oracle Database High Availability Best Practices* and in the white papers that can be downloaded from:

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

Oracle Database High Availability Features and Products

Oracle Database offers an integrated suite of high availability solutions that increase availability and eliminate or minimize both planned and unplanned downtime. These solutions help enterprises maintain business continuity 24 hours a day, seven days a week. However, the Oracle high availability solutions go beyond reducing downtime by providing solutions to increase system utilization on the primary and secondary systems and helping to improve overall performance, scalability, and manageability.

The chapter contains the following sections that outline the key impacts of the Oracle high availability features on businesses and applications:

- [Oracle High Availability Features and Solutions for Unplanned Downtime](#)
- [Oracle High Availability Features and Solutions for Planned Downtime](#)
- [Optimizing Grid Computing and Disaster Recovery Solutions](#)
- [Optimizing Manageability](#)

2.1 Oracle High Availability Features and Solutions for Unplanned Downtime

Oracle provides the following features for high availability:

- [Fast-Start Fault Recovery](#)
- [Oracle Real Application Clusters and Oracle Clusterware](#)
- [Oracle Data Guard](#)
- [Oracle Streams](#)
- [Transportable Technologies](#)
- [Oracle Flashback Technology](#)
- [Automatic Storage Management](#)
- [Recovery Manager](#)
- [Oracle Secure Backup](#)
- [Data Recovery Advisor](#)
- [Flash Recovery Area](#)
- [Oracle Security Features](#)
- [LogMiner](#)

- [Hardware Assisted Resilient Data \(HARD\) Initiative](#)
- [Hang Manager](#)
- [Data Block Corruption Prevention and Detection Parameters](#)

Also, the [Oracle HA Solutions and Recovery Times for Unplanned Downtime](#) section on page 2-20 provides a summary of the key HA solutions that address different types of unplanned downtime along with the recovery time that can be attained with each solution.

See Also:

- The "High Availability" chapter in *Oracle Database Concepts*. for an overview of the high availability features
- The "Availability" section in the *Oracle Database New Features Guide*. for a list of all the new high availability features introduced in Oracle Database 11g Release 1 (11.1)

2.1.1 Fast-Start Fault Recovery

Oracle provides fast and predictable recovery from system faults and database failures. The Fast-Start Fault Recovery technology included in 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 instance, database, and computer failures
- Database checkpointing is self-tuning to maintain desired recovery time objective

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

2.1.2 Oracle Real Application Clusters and Oracle Clusterware

Oracle Real Application Clusters (Oracle RAC) and Oracle Clusterware allow the Oracle Database to run any packaged or custom application across a set of clustered servers. This capability provides the highest levels of availability and the most flexible scalability. If a clustered server fails, Oracle Database continues running on the surviving servers. When more processing power is needed, you can add another server without interrupting access to data.

Oracle RAC enables multiple instances that are linked by an interconnect to share access to an Oracle database. In an Oracle 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. This enables Oracle RAC to provide high availability and redundancy during failures in the cluster. Oracle RAC accommodates all system types, from read-only data warehouse (DSS) systems to update-intensive online transaction processing (OLTP) systems.

Oracle Clusterware is software that, when installed on servers running the same operating system, enables the servers to be bound together to operate as if they are one server, and manages the availability of user applications and Oracle databases. Oracle Clusterware also provides all of the features required for cluster management,

including node membership, group services, global resource management, and high availability functions:

- For high availability, you can place Oracle databases (single-instance or Oracle RAC databases), and user applications (Oracle and non Oracle) under the management and protection of Oracle Clusterware so that they restart on process failure or fail over to another node on node failure.
- For cluster management, Oracle Clusterware presents multiple independent servers as if they are a single-system image or one virtual server. This single virtual server is preserved across the cluster for all management operations, enabling administrators to perform installations, configurations, backups, upgrades, and monitoring functions once. Then, Oracle Clusterware automatically distributes the execution of these management functions to the appropriate nodes in the cluster.

Oracle Clusterware is a requirement for using Oracle RAC and it is the only clusterware that you need for most platforms on which Oracle RAC operates. Although the Oracle Database continues to support select third-party clusterware products on specified platforms, using Oracle Clusterware provides the benefit of dispensing with proprietary vendor clusterware and using an integrated software stack from Oracle that provides disk management with Oracle ASM to data management with the Oracle Database and Oracle RAC. In addition, Oracle Database features, such as Oracle Services, use the underlying Oracle Clusterware mechanisms to provide their capabilities.

Oracle Clusterware requires two clusterware components: a voting disk to record node membership information and the Oracle Cluster Registry (OCR) to record cluster configuration information. The voting disk and the OCR must reside on shared storage. The Oracle Clusterware requires that each node be connected to a private network by way of a private interconnect.

Together, Oracle RAC and Oracle Clusterware provide the following benefits:

- Ability to tolerate and quickly recover from computer and instance failures
- Ability to apply Oracle Clusterware upgrades, patch sets, and interim patches in a rolling fashion

For example, upgrading Oracle Clusterware from Oracle 10g to Oracle 11g, patching Oracle Clusterware from Oracle 10.2.0.3 to 10.2.0.4, and patching Oracle Clusterware from Oracle 10.2.0.2 Bundle 1 to Oracle 10.2.0.2 Bundle 2

- Rolling upgrades for system and hardware changes
- Rolling patch upgrades for some interim patches
- Service relocation and, when you perform additional Fast Application Notification (FAN) and client configuration, applications can receive FAN events and react immediately to achieve fast, automatic, and intelligent connection and failover
- Load balancing advisory
- Runtime connection load balancing with Oracle JDBC, OCI, and ODP.NET
- Flexibility to increase processing capacity using commodity hardware without downtime or changes to the application
- Comprehensive manageability integrating database and cluster features
- Scalability across database instances

Specifically, Oracle Clusterware provides the following benefits

- Automatically restarts failed Oracle processes
- Automatically manages and fails over Oracle Virtual IP (VIP) on another node in the cluster on node failures
- Automatically restarts resources from failed nodes on surviving nodes
- For Oracle RAC databases, all Oracle processes are under the control of Oracle Clusterware by default; for Oracle single-instance databases, you can configure the Oracle processes into a resource group that is under the control of Oracle Clusterware
- For Oracle and non Oracle applications, Oracle Clusterware provides an application programming interface (API) that enables you to control other Oracle processes with Oracle Clusterware, such as restart or react to failures and certain rules
- Manages node membership and prevents split brain syndrome in which two or more instances attempt to control the database
- Provides rolling release upgrades of Oracle Clusterware, with no downtime for applications

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

2.1.3 Oracle 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 corruptions. Data Guard maintains standby databases as transactionally consistent copies of the primary (production) database. Then, if the primary database becomes unavailable because of a planned or an unplanned outage, Data Guard can switch any standby database to the primary role, minimizing the downtime associated with the outage. Data Guard can be used with traditional backup, restoration, and cluster techniques to provide a high level of data protection and data availability.

With Data Guard, administrators can optionally improve primary database performance by offloading resource-intensive backup and reporting operations to standby systems.

An Oracle Data Guard configuration consists of one primary database and one or more standby databases. Using a backup copy of the primary database, you can create up to nine standby databases and incorporate them in a Data Guard configuration. Once created, Data Guard automatically maintains each standby database by transmitting redo data from the primary database and then applying the redo to the standby database.

Similar to a primary database, a standby database can be either a single-instance Oracle database or an Oracle RAC database.

A standby database can be a physical standby database, a snapshot standby database, or a logical standby database, and a Data Guard configuration can include any combination of these types of standby databases:

- Physical standby database
 - Provides a physically identical copy of the primary database, with datafiles that are identical to the primary database. The database schema, including indexes, are

the same. A physical standby database is kept synchronized with the primary database, though Redo Apply, which recovers the redo data received from the primary database and applies the redo to the physical standby database.

A physical standby database can be used for business purposes other than disaster recovery. Starting in Oracle Database 11g Release 1, the physical standby database can be opened with read-only access while redo is being applied to the standby database. This mode, referred to as real-time query standby, allows users to access an up-to-date physical standby database for queries at any time. Also, you can convert a physical standby database into a snapshot standby database.

- Snapshot standby database

A snapshot standby database is an updatable standby database that you create from a physical standby database. A snapshot standby database receives and archives redo data received from the primary database, but the snapshot standby database does not apply redo data from the primary database while the standby is open read/write. Because of this, the snapshot standby typically diverges from the primary database over time. Moreover, local updates to the snapshot standby database cause additional divergence.

Redo data is not applied until you convert the snapshot standby database back into a physical standby database, and after all local updates to the snapshot standby database are discarded. With a single command, you can revert a snapshot standby back to a physical standby database, at which time the changes made to the snapshot standby state are discarded, and Redo Apply automatically resynchronizes the physical standby database with the primary database using the redo that was archived.

- Logical standby database

Contains the same logical information as the primary database, although the physical organization and structure of the data can be different. The logical standby database is kept synchronized with the primary database though SQL Apply, which transforms the data in the redo received from the primary database into SQL statements and then executing the SQL statements on the standby database.

A key benefit of a logical standby database is that you can create significant auxiliary structures to optimize the reporting workload, including structures that could have a prohibitive effect on the primary database's transactional response time. A logical standby database can have its data physically reorganized into a different storage type with different partitioning, having many different indexes, having on-demand refresh materialized views created and maintained, and it can be used to drive the creation of data cubes and other OLAP data views.

A logical standby database can be used for other business purposes in addition to disaster recovery requirements. This allows users to access a logical standby database for queries and reporting purposes at any time. Also, using a logical standby database, you can upgrade Oracle Database software and patch sets with almost no downtime. Thus, a logical standby database can be used concurrently for data protection, reporting, and database upgrades.

Oracle Data Guard provides the following benefits:

- Maintenance of real-time, transactionally consistent database copies to provide protection against unplanned downtime and disaster.
- Data protection and fast repair against computer failures, human errors, data corruption, lost writes, and site failures.

- Automatic failover with flexible data protection levels to support all network configurations and business requirements.
- Faster redo application, redo transport, and role transitions with various enhancements.
- Reduction of planned downtime for system changes, some platform migrations, hardware and system upgrades, and Oracle patch set and database upgrades (see also [Table 2–1, "Outage Types and Oracle High Availability Solutions for Unplanned Downtime"](#)).
- Multiple levels of data protection and performance to balance data availability against system performance requirements.
- Support for the real-time query capability of physical standby databases provides better primary database performance and better use of standby resources, because real-time query:
 - Diverts read-only queries and transactions from the primary database to standby databases.
 - Diverts backup operations from the primary database to standby databases.
- Support for both physical standby databases (real-time query) and logical standby databases to provide more efficient use of system resources by diverting more querying and reporting functions from the primary database to standby databases (with the logical standby databases providing greater flexibility for any activity that requires access to a standby database that is open read/write).
- Support for snapshot standby database for reporting or testing (cloning) purposes and automatic resynchronization with the primary database once reporting or testing has completed.
- Managed and automatic role transition and application notification to minimize planned and unplanned downtime.
- Automatic or automated resynchronization of a failed primary database following a failover.
- Management of all systems as a single configuration for simplified administration.
- Increased flexibility for Data Guard configurations where the primary and standby systems may have different CPU architectures, operating systems (for example, Windows and Linux), operating system binaries (32-bit and 64-bit), and Oracle database binaries (32-bit and 64-bit); this is subject to restrictions that are defined in MetaLink note 413484.1.

See Also: Metalink note 413484.1 at <https://metalink.oracle.com/> for the latest capabilities and restrictions

See Also: *Oracle Data Guard Concepts and Administration*

2.1.4 Oracle Streams

Oracle Streams is a very flexible and powerful database feature to implement fine-grained replication, multimaster replication, many-to-one replication, data transformation, hub and spoke replication, and message queuing.

Comparing Streams and Data Guard

Streams is designed from the ground up for information sharing. It enables highly customized replication strategies to satisfy the many varied uses of data replicated to a target database. These same capabilities can also make Streams a useful technology for addressing HA/DR requirements and for minimizing planned downtime during upgrades to new database releases and patch sets.


Data Guard is designed for simple, one-way replication of an entire database expressly for maintaining a synchronized copy that can assume the primary role in the event of a failure. Data Guard Redo Apply (physical standby) best exemplifies this notion of *simplicity*, as a disaster recovery solution that is both datatype and application agnostic, and able to scale to very high levels of performance. While Data Guard also provides capabilities that enable a standby database to offload the primary database of the overhead of performing backups, queries, and reports, these capabilities are ancillary to Data Guard's primary mission, and are provided to increase the ROI of customer HA/DR investments. You can use Data Guard SQL Apply (logical standby) to get additional value from a Data Guard configuration by minimizing planned downtime during upgrades to new database releases and patch sets.

Streams Messaging and Information Flow

Oracle Streams enables information sharing. Using Streams, each unit of shared information is called a message, and you can share these messages in a stream. The stream can propagate information within a database or from one database to another.

For example, [Figure 2-1](#) shows a multimaster configuration where all sites are directly connected to all other sites participating in the replication environment. This enables data to be replicated between all locations in a near realtime manner.

Figure 2-1 Oracle Streams Multimaster Configuration




Another example is the Oracle Streams 1-N, or *hub-and-spoke configuration* in which changes made at the primary or hub location can be propagated to the remote or spoke locations in a near real-time manner.

Although it is possible to configure a hub-and-spoke configuration for bidirectional replication, you may prefer to restrict updates to a single location, the hub, as shown in [Figure 2-2](#). In query intensive environments, you can still balance the load between multiple locations, with fast local access, while updates are restricted to the hub. By offloading reporting to the spoke locations, you improve performance at the hub, or primary OLTP location. This type of configuration is easier to implement than multimaster replication because it is not necessary to establish connectivity between all

locations in the replication environment and it is not necessary to implement a conflict resolution strategy.

Figure 2–2 Information Dissemination with Oracle Streams (1-N configuration)




The stream routes specific information to specific destinations. The result is a feature that provides greater functionality and flexibility than traditional solutions for capturing and managing messages, and sharing the messages with other databases and applications. Streams provides the capabilities needed to build and operate distributed enterprises and applications, data warehouses, and high availability solutions. You can use all of the capabilities of Oracle Streams at the same time. If your business requirements change, then you can implement a new capability of Streams without sacrificing existing capabilities.

As with any Oracle Streams configuration, there are three phases: capture, stage (propagate), and consume (apply). Using Oracle Streams, you control what information is put into a stream, how the stream flows or is routed from database to database, what happens to messages in the stream as they flow into each database, and how the stream terminates. By configuring specific capabilities of Streams, you can address specific requirements. Based on your specifications, Streams can capture, stage, and manage messages in the database automatically, including, but not limited to, data manipulation language (DML) changes and data definition language (DDL) changes. You can also put user-defined messages into a stream, and Streams can propagate the information to other databases or applications automatically. When messages reach a destination, Streams can consume them based on your specifications.

The following figure shows the Oracle Streams information flow.

Figure 2–3 Oracle Streams Information Flow




With Streams, you can create a local or remote copy of a production database. In the event of human error or a catastrophe, the copy can be used to resume processing. You can use Streams to configure flexible high availability environments.

You can use the features of Oracle Streams to achieve little or no database down time during database upgrade and maintenance operations. Maintenance operations include migrating a database to a different platform, migrating a database to a

different character set, modifying database schema objects to support upgrades to user-created applications, and applying an Oracle software patch.

Figure 2–4 shows an application that explicitly enqueues and dequeues messages through Oracle Streams Advanced Queuing as a method of sharing information with business partners or customers with different messaging systems. Once enqueued, messages can be transformed and propagated as desired, before being dequeued to a business partner’s application that is a non-database oriented messaging systems.

Figure 2–4 Oracle Streams Message Queuing



Benefits of Using Oracle Streams

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 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; bidirectional replication environment, in which exactly two databases share the replicated database objects and data, is possible.
- Event management and notification by enqueueing messages or capturing events, propagating the messages and events through queues, and dequeuing and applying or acting upon the message or event (as shown in Figure 2–4)
- Supports heterogeneous platforms across databases within the configuration
- Allows character sets to differ between replicas
- Permits fine-grained control of data sharing

Note: Although Oracle Streams requires some initial implementation investment, it is well worth the effort because of the high flexibility you get with Streams.

See Also: *Oracle Streams Concepts and Administration*

2.1.5 Transportable Technologies

Transportable technologies provides transportable database and transportable tablespace:

- **Transportable database** is used to transport an entire database (user data and the Oracle dictionary) to a new platform with the same endian format. Transportable database permits a minimal downtime migration to a new platform by avoiding

the time-consuming method of unloading all user data from the source database and loading it into the target database.

- **Transportable tablespaces** is a feature designed to move a subset of one database into another, even among platforms that differ in endian format. The cross-platform capability of transportable tablespaces can be used to migrate all user data within a database to a new platform with a different endian format. Leveraging transportable tablespaces in this manner permits a minimal downtime migration to a new platform by avoiding the time-consuming method of unloading all user data from the source database and loading it into the target database.

See Also: *Oracle Database Administrator's Guide* for details about how to move or copy tablespaces to another database, including details about transporting tablespaces across platforms

2.1.6 Oracle Flashback Technology

Flashback technology provides a set of features to switch between views of the data as it existed at different points in time. Using flashback features you can query past versions of schema objects and historical data. You can also perform change analysis and 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 Oracle Database and supports recovery at all levels including row, transaction, table, tablespace, and database.

Most of the flashback features use undo data while other features, such as Flashback Database and Block Media Recovery, use flashback logs:

- **Undo tablespace**—A dedicated tablespace that stores only undo information when the database is run in automatic undo management mode.
- **Flashback data archive**—An archive that is stored in a tablespace and contains transactional changes to every record in a table for the duration of the record's lifetime. The archived data can be retained for much longer duration than the retention period offered by an undo tablespace.
- **Flashback logs**—Oracle-generated logs used to perform flashback operations. The database can only write flashback logs to the flash recovery area. Flashback logs are written sequentially and are not archived. They cannot be backed up to disk.

The following list describes the Flashback features:

- **Oracle Flashback Query**

Oracle Flashback Query provides the ability to view the data as it existed in the past by using the Automatic Undo Management system to obtain metadata and historical data for transactions. Undo data is persistent and survives 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.

See Also: *Oracle Database Advanced Application Developer's Guide*

- **Oracle Flashback Versions Query**

Oracle Flashback Versions Query is an extension to SQL that you can use 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 the sources of problems. Additionally, application developers can use Flashback Versions Query to build customized applications for auditing purposes.

See Also: *Oracle Database Advanced Application Developer's Guide*

- **Oracle Flashback Transaction**

Oracle Flashback Transaction is a new feature in Oracle Database 11g Release 1 that can easily back out a transaction and its dependent transactions. The `DBMS_FLASHBACK.TRANSACTION_BACKOUT()` procedure rolls back a transaction and its dependent transactions while the database remains online. This recovery operation uses undo data to create and execute the compensating transactions that return the affected data to its original state. You can query the `DBA_FLASHBACK_TRANSACTION_STATE` view to see the current state of a transaction with respect to whether the transaction has been backed out using dependency rules or forced out by either:

- Backing out nonconflicting rows
- Applying undo SQL

Oracle Flashback Transaction increases availability during logical recovery by easily and quickly backing out a specific transaction or set of transactions and their dependent transactions, with one command while the database remains online.

See Also: *Oracle Database Advanced Application Developer's Guide*

- **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 human or application error. Flashback Transaction Query increases the ability to

perform online diagnosis of problems in the database by returning the database user that changed the row, and performs analysis and audits on transactions.

See Also: *Oracle Database Advanced Application Developer's Guide*

- **Oracle Flashback Table**

Oracle Flashback Table enables you 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 modified by a human or application error. 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.

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

- **Oracle Flashback Drop**

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

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

- **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 labels you can define 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 the Oracle Flashback restore points feature provides the following benefits:

- Provides the ability to quickly restore to a consistent state, to a point in time that was prior to a planned operation that has gone awry (for example, a failed batch job, an Oracle software upgrade, or an application upgrade).
- Allows the snapshot standby to be resynchronized with the production database.
- Allows for a quick mechanism to restore a test or cloned database back to its original state.

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

- **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 restoring data from

data file backups and executing point-in-time database recovery. However, Flashback Database skips the data file restoration and most of the application of redo data.

Enabling Oracle Flashback Database provides the following benefits:

- Eliminates the time to restore a backup when fixing human error that has a database-wide impact.
- Because human errors can be quickly undone, it allows standby databases to use real-time apply to synchronize with the primary database.
- Allows quick standby database reinstatement after a database failover.

See Also:

- *Oracle Database Backup and Recovery User's Guide*
- *Oracle Database SQL Language Reference*

- **Block Recovery Using Flashback Logs**

Starting with Oracle Database release 11.1, block recovery can optionally retrieve a more recent copy of a data block from the flashback logs to reduce recovery time. Furthermore, a corrupted block encountered during instance recovery does not result in instance recovery failing. The block is automatically marked as corrupt and added to the RMAN corruption list in the `V$DATABASE_BLOCK_CORRUPTION` table. You can subsequently issue the `RMAN RECOVER BLOCK` command to fix the associated block.

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

- **Flashback Data Archive**

An archive that is stored in a tablespace and contains transactional changes to every record in a table for the duration of the record's lifetime. The archived data can be retained for much longer duration than the retention period offered by an undo tablespace.

See Also: *Oracle Database Advanced Application Developer's Guide*

2.1.7 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 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 called a disk group. The disk group identifies the set of disks that are 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 protects against storage failures. With ASM mirroring, an additional level of data protection can be provided with the use of failure 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 intelligently places redundant copies of the data in separate failure groups. This ensures that the data is available and transparently protected against the failure of any component in the storage subsystem.

ASM provides the following benefits:

- Provides the ability to mirror and stripe 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 database storage
- Manages OCR and voting disks
- Provides local read capability, which gives better performance in an extended cluster
- Supports very large databases
- Supports ASM rolling upgrades
- Supports finer granularity in tuning and security
- ASM Fast Mirror Resync, which provides fast repair after a temporary disk failure

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

2.1.8 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

- Fast backup and restore operations with intrafile and interfile parallelism
- Enhanced security with Virtual Private Catalog
- Lower space consumption when creating a database over the network by eliminating staging areas
- 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
- Ability to resume 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 User's Guide*

2.1.9 Oracle Secure Backup

Oracle Secure Backup (OSB) is a centralized tape backup management solution providing performant, heterogeneous data protection in distributed UNIX, Linux, Windows, and Network Attached Storage (NAS) environments. By protecting file system and Oracle database data, OSB provides a complete tape backup solution for your IT environment.

OSB is tightly integrated with Recovery Manager (RMAN) to provide the media management layer for RMAN, supporting releases since Oracle9i. With optimized integration points, OSB and RMAN provide the fastest most efficient tape backup for the Oracle database.

You can backup distributed servers to local and remote tape devices from a central OSB administrative server using backup policies, calendar based scheduling for *lights out* operations, or on-demand backup for immediate requirements. With its highly scalable client/server architecture, OSB provides local and remote data protection, leveraging SSL for secure intradomain communication and two-way server authentication.

The following list describes the key benefits of OSB:

- Optimized tape backup for the Oracle database by backing up only the currently used blocks and increasing backup performance by 10% to 25%.
- Policy based management allows backup administrators to exercise precise control over the backup domain.
- Dynamic drive sharing for increased tape resource use.
- Heterogeneous storage area network (SAN) support allowing NAS, UNIX, Windows, and Linux to share tape drives and media.
- File system backup at the file, directory, file system or raw partition level with full, incremental and offsite backup scheduling.
- Integrated with Oracle Enterprise Manager, providing an intuitive, familiar interface.
- Backup encryption to tape.

- Broad tape-device support for new and legacy tape devices in SAN and SCSI environments.
- Network Data Management Protocol (NDMP) support for highly efficient backup of NAS filers.
- Scalable, low-cost licensing model reduces IT costs and operational considerations.

See Also: *Oracle Secure Backup Administrator's Guide*

2.1.10 Data Recovery Advisor

Data Recovery Advisor, a new feature in Oracle Database 11g release 1 (11.1), automatically diagnoses persistent (on disk) data failures, presents appropriate repair options, and runs repair operations at your request.

Note: Note that the initial release of Data Recovery Advisor does not support Oracle RAC. In addition, while you can use Data Recovery Advisor when managing a primary database in a Data Guard configuration, you cannot use Data Recovery Advisor to troubleshoot a physical standby database. Data Recovery Advisor only takes the presence of a standby database into account when recommending repair strategies if you are using Enterprise Manager 11g Grid Control.

Data Recovery Advisor includes the following functionality:

- **Failure Diagnosis**

The first symptoms of database failure are usually error messages, alarms, trace files and dumps, and failed health checks. Assessing these symptoms can be complicated, error-prone and time-consuming. Data Recovery Advisor automatically diagnoses data failures and informs you about them.
- **Failure Impact Assessment**

After a failure is diagnosed, you must understand its extent and assess its impact on applications before devising a repair strategy. Data Recovery Advisor automatically assesses the impact of a failure and displays it in a easily understood format.
- **Repair Generation**

You typically have several repair options available, offering trade-offs in recovery time and potential data loss. If there are multiple failures present, you must also determine the best sequence of repair steps. In some situations it can be advantageous to consolidate repairs. Data Recovery Advisor does all this for you, automatically determining the best repair options.
- **Repair Feasibility Checks**

Before presenting repair options, Data Recovery Advisor validates them with respect to the specific environment and availability of media components required to complete the proposed repair. The feasibility check is fast and validates if the required backups are available. The actual contents of these backups will be validated during repair.
- **Repair Automation**

If you accept the suggested repair option, Data Recovery Advisor automatically performs the repairs, verifies that the repair was successful, and closes the appropriate failures.

- Validation of Data Consistency and Database Recoverability

Data Recovery Advisor can validate the consistency of your data, as well as backups and redo stream, whenever you choose.

- Early Detection of Corruption

Through Health Monitor, you can schedule periodic runs of Data Recovery Advisor diagnostic checks to detect, analyze, and repair data failures before a database process executing a transaction discovers the corruption and signals an error. Early warnings can limit the damage caused by corruption.

- Integration of Data Validation and Repair

Data Recovery Advisor is a single tool for data validation and repair.

See Also: “Diagnosing and Repairing Failures with the Data Recovery Advisor” in *Oracle Database Backup and Recovery User’s Guide*

2.1.11 Flash Recovery Area

The flash recovery area is a unified storage location for all recovery-related files and activities in 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, which simplifies database administration tasks
- Fast, reliable disk-based backup and restoration
- Ability to backup and restore the entire flash recovery area
- Ability to tolerate failures to the flash recovery area

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

2.1.12 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 only those data and services truly needed to perform business functions. Oracle provides a wide range of security tools to control access to application data by authenticating database users and then enabling administrators to grant them only those privileges required to perform their duties.

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

Oracle security features include:

- Authentication control to validate the identities of entities using networks, databases, and applications. Network sessions between databases, such as redo transport sessions, are also authenticated.
- Authorization control to provide limits to access and actions linked by database 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* and *Oracle Data Guard Concepts and Administration*

2.1.13 LogMiner

Oracle log files contain useful information about the activities and history of the Oracle database. Log files contain all data necessary to perform 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
- Provide additional browser-based simplification to help troubleshoot and resolve logical failures

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*

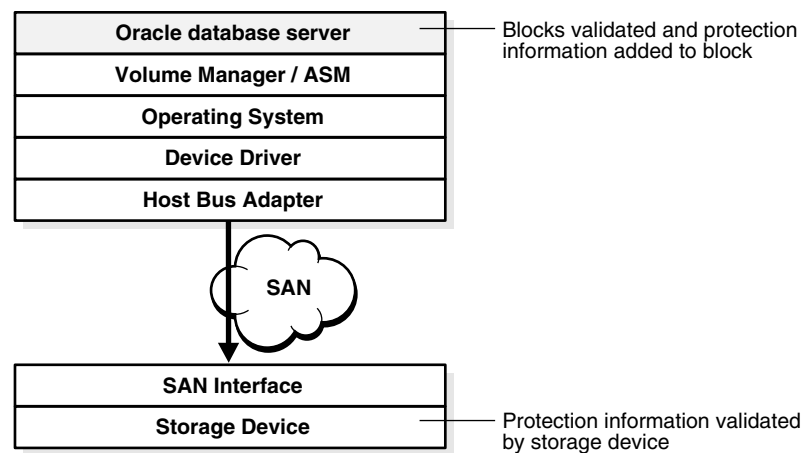
2.1.14 Hardware Assisted Resilient Data (HARD) Initiative

The Hardware Assisted Resilient Data (HARD) Initiative is a joint initiative between Oracle and hardware vendors to prevent data corruptions from being written out to disk. 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.

Any datafiles and log files located on HARD-compliant storage is protected. You 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 2–5 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.

See Also: The HARD initiative page for the latest vendor and implementation information at <http://www.oracle.com/technology/deploy/availability/htdocs/HARD.html>

2.1.15 Hang Manager

The Hang Manager, a new feature in Oracle Database 11g Release 1, simplifies database management by enabling you to more quickly resolve database hangs. It automatically detects, analyzes, and dumps diagnostic information for hangs in Oracle Database environments, including Oracle RAC and ASM databases.

The Hang Manager is enabled by default in Oracle RAC databases and ASM instances. Thus, you can use Hang Manager to follow a hang from the database instance to the underlying ASM instance.

Active entities that attempt to obtain restrictive access to shared resources or request services from other Oracle Database processes, sessions, and transactions are in danger

of hanging. A hang chain is a chain of processes with each one waiting on a resource held by the next, with a single process serving as the root of the hang.

Hangs in Oracle Database can cost a great deal in terms of system unavailability. Specifically, hangs lead to the following problems:

- **Extended system outages.** These outages may occur frequently before a fix is found, which adds to the total downtime.
- **Difficult analysis:** Analyzing the hang to determine where the problem lies can be lengthy, complex, and prone to error.

See Also: *Oracle Database Reference* for more information about the views used to determine hangs

2.1.16 Data Block Corruption Prevention and Detection Parameters

Prior to Oracle Database 11g, RMAN-detected block corruptions were recorded in `V$DATABASE_BLOCK_CORRUPTION`. In Oracle Database 11g, several database components and utilities, including RMAN, can now detect a corrupt block and record it in that view. Oracle Database automatically updates this view when block corruptions are detected or repaired (for example, using block media recovery or data file recovery). The benefit is that the time it takes to discover block corruptions is shortened.

In addition, you can use the `DB_ULTRA_SAFE` initialization parameter to automatically configure the appropriate data protection block checking level in the database. The performance impact may vary depending on the application and available system resources, but the effect can vary from 1% to 10%.

The `DB_ULTRA_SAFE` initialization parameter:

- Controls the setting of other related initialization parameters, including `DB_BLOCK_CHECKING`, `DB_BLOCK_CHECKSUM`, and `DB_LOST_WRITE_PROTECT`
- Controls other data protection behavior in the Oracle Database, such as requiring ASM to perform sequential mirror writes

By making it possible to detect data corruptions in a timely manner, these features provide critical high availability benefits for the Oracle database.

See Also: *Oracle Database Reference* for more information about these views and initialization parameters

2.1.17 Oracle HA Solutions and Recovery Times for Unplanned Downtime

Oracle provides high availability solutions to prevent, tolerate and reduce downtime for all types of unplanned 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. The table shows how the features discussed in [Section 2.1.1](#) through [Section 2.1.16](#) can be used to address various causes of unplanned downtime. Also, see [Table 4–4](#) on page 4-29 for a summary of the attainable recovery times for all types of unplanned downtime for each Oracle high-availability architecture.

Table 2–1 Outage Types and Oracle High Availability Solutions for Unplanned Downtime

Outage Type	Oracle Solution	Benefits
Computer Failures	Oracle Real Application Clusters and Oracle Clusterware	<ul style="list-style-type: none"> Automatic recovery of failed nodes and instances. Fast application notification with integrated Oracle client failover.
Computer Failures	Fast-Start Fault Recovery	<ul style="list-style-type: none"> Tunable and predictable cache recovery from computer failures.
Computer Failures	Oracle Data Guard	<ul style="list-style-type: none"> Fast-start failover and fast application notification with integrated Oracle clients.
Computer Failures	Oracle Streams	<ul style="list-style-type: none"> Online replica database resume processing.
Storage Failures	Automatic Storage Management	<ul style="list-style-type: none"> Mirroring and online automatic rebalance places redundant copies of the data in separate failure groups.
Storage Failures	Oracle Data Guard	<ul style="list-style-type: none"> Fast-start failover and fast application notification with integrated Oracle clients
Storage Failures	Recovery Manager with Flash Recovery Area	<ul style="list-style-type: none"> Fully managed database recovery and managed disk-based backups
Storage Failures	Oracle Streams	<ul style="list-style-type: none"> Online replica database resumes processing.
Human Errors	Oracle Security Features	<ul style="list-style-type: none"> Restrict access as prevention
Human Errors	Oracle Flashback Technology	<ul style="list-style-type: none"> Fine-grained and database-wide rewind capability
Human Errors	LogMiner	<ul style="list-style-type: none"> Redo log analysis
Data Corruption	Hardware Assisted Resilient Data (HARD) Initiative	<ul style="list-style-type: none"> Corruption prevention within a storage array.
Data Corruption	Data Block Corruption Prevention and Detection Parameters Database initialization settings such as DB_ULTRA_SETTINGS, DB_BLOCK_CHECKING, DB_BLOCK_CHECKSUM	<ul style="list-style-type: none"> Different levels of block corruption prevention and detection at the database level.
Data Corruption	Data Recovery Advisor and Recovery Manager with Flash Recovery Area	<ul style="list-style-type: none"> Data Recovery Advisor automatically detects data corruptions and advises you about the best recovery plan. RMAN online block-media recovery time in Oracle Database 11g is faster, because RMAN can now use flashback logs to restore a more current copy of the data block for recovery.
Data Corruption	Oracle Data Guard	<ul style="list-style-type: none"> Fast-start failover and fast application notification with integrated Oracle clients.
Data Corruption	Oracle Streams	<ul style="list-style-type: none"> Online replica database resumes processing.

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

Outage Type	Oracle Solution	Benefits
Lost writes	<p>Oracle Data Guard, Recovery Manager, and the <code>DB_LOST_WRITE_PROTECT</code> initialization parameter</p> <p>Also, setting <code>DB_ULTRA_SETTINGS</code> to <code>DATA_ONLY</code> or <code>DATA_AND_INDEX</code> automatically enables <code>DB_LOST_WRITE_PROTECT</code></p>	<ul style="list-style-type: none"> ■ <code>DB_LOST_WRITE_PROTECT</code> initialization parameter provides lost write detection. ■ If a lost write is detected on the primary database either by the physical standby database or during media recovery of the primary database, recovery is stopped to preserve the consistency of the database. However, data loss is expected. ■ If a lost write is detected on the standby database, you can restore the affected file and restart managed recovery if the lost write is isolated and the hardware problem is corrected. <p>Note: Lost writes can corrupt the entire database. In many cases, you need to rebuild the affected database after resolving the hardware issue.</p>
Lost writes	Hardware Assisted Resilient Data (HARD) Initiative	<ul style="list-style-type: none"> ■ Detection and prevention of lost writes resulting from stray or misdirected writes. Some vendors in the HARD program have not implemented this additional protection. For the most comprehensive lost write protection, use Oracle Data Guard. ■ For the most comprehensive lost write protection, use Oracle Data Guard and set either the <code>DB_ULTRA_SETTING</code> parameter (to <code>DATA_ONLY</code> or <code>DATA_AND_INDEX</code>) or set the <code>DB_LOST_WRITE_PROTECT</code> parameter (to <code>TYPICAL</code> or <code>FULL</code>) on both the primary and standby databases.
Hangs or slow down	Hang Manager and Oracle Enterprise Manager	<ul style="list-style-type: none"> ■ Hang Manager monitors for database hangs and attempts to resolve them automatically. ■ Oracle Enterprise Manager can be configured to detect application or response time slowdown and react to these SLA breaches. For example, you can configure the Enterprise Manager Beacon to monitor and detect application response times.
Site Failures	Oracle Data Guard	<ul style="list-style-type: none"> ■ Fast-start failover and fast application notification with integrated Oracle clients
Site Failures	Oracle Streams	<ul style="list-style-type: none"> ■ Online replica database resumes processing.
Site Failures	Recovery Manager	<ul style="list-style-type: none"> ■ Fully managed database recovery and integration with Oracle Secure Backup.

2.2 Oracle High Availability Features and 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 Internet access to customers 24 hours a day, seven days a week.

In the past, planned downtime became necessary when performing periodic maintenance or when migrating to new deployments. Periodic maintenance—such as patching or reconfiguring the system—may be necessary to update the database, application, operating system, middleware, or network. New deployments include

major upgrades or new rollouts of the hardware, database, application, operating system, middleware, or network.

Oracle provides the following high availability solutions to eliminate or reduce planned downtime for system and database changes, data changes, and application changes:

- For system and database changes, see [Dynamic Resource Provisioning](#) on page 2-23
- For data changes, see [Online Reorganization and Redefinition](#) on page 2-33
- For application changes, see [Online Application Maintenance and Upgrades](#) on page 2-35

2.2.1 Dynamic Resource Provisioning

This section describes dynamic resource provisioning under the following topics:

- [Dynamic Reconfiguration of the Database](#)
- [Autotuning Memory Management](#)
- [Automated Distribution of Datafiles, Control Files, and Log Files](#)

2.2.1.1 Dynamic Reconfiguration of the Database

Oracle continues to broaden support for dynamic reconfiguration of the database, enabling it to adapt to changes in hardware demands without any service interruptions. 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 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 the 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
- Change almost all initialization parameters without shutting down the instance by using the `SQL*Plus 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

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

2.2.1.2 Autotuning Memory Management

Oracle Database 11g release 1 (11.1) introduces two new memory management initialization parameters, `MEMORY_TARGET` and `MEMORY_MAX_TARGET` that enable

automatic management of the system global area (SGA), program global area (PGA), and other memory required to run Oracle Database.

Note: MEMORY_MAX_TARGET is the value up to which MEMORY_TARGET can grow dynamically. If these initialization parameters are left at their default values (0), then Oracle Database does not autotune memory. If one parameter is set to a nonzero value and other is not set, then Oracle Database internally sets both parameters to the nonzero value.

Oracle Database uses a noncentralized policy to free and acquire memory in each subcomponent of the SGA and the PGA. Oracle Database autotunes memory by prompting the operating system to transfer granules of memory from less needy to more needy components. The granularity of the memory transfer is dependent on the current free memory and the amount of memory the operating system needs to maintain a basic level of service.

Note: Automatic memory management with the MEMORY_TARGET and MEMORY_MAX_TARGET initialization parameters is supported on Linux, Windows, Solaris, HP-UX, and AIX. Check the *Oracle Database Concepts* and *Oracle Database Administrator's Guide* for more information about all supported platforms.

2.2.1.3 Automated Distribution of Datafiles, Control Files, and Log Files

ASM automates and simplifies the layout of datafiles, control files, and log files. Database files are automatically distributed across all available disks. 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 striping database files across available disks.

See Also: *Oracle Database Concepts* and *Oracle Database Storage Administrator's Guide* for more information about Automatic Storage Management

2.2.2 Oracle High Availability Solutions and Recovery Times for Planned Downtime

Oracle provides high availability solutions to prevent, tolerate and reduce downtime for all types of unplanned failures.

Table 2–2 describes the various Oracle high availability solutions for planned downtime along with the outage time that can be attained with each solution and their known considerations. In all cases, Oracle recommends you perform extensive testing prior to performing any rolling upgrade. Also, see Table 4–5 on page 4-29 for a summary of the attainable recovery times for all types of planned downtime for each Oracle high-availability architecture.

Table 2–2 Oracle High Availability Solutions for Planned Downtime

Maintenance Type	Oracle Recommended Solution	Solution Description	Outage Time
Operating system and hardware upgrades	Oracle Real Application Clusters and Oracle Clusterware	Section 2.2.2.1	No downtime
Oracle interim patches	Oracle Real Application Clusters (Oracle RAC)	Section 2.2.2.3	No downtime ¹
Online Patches	Online Patching	Section 2.2.2.4	No downtime
Oracle Clusterware upgrades and patches	Cluster Ready Services (CRS)	Section 2.2.2.5	No downtime
ASM upgrades	Automatic Storage Management	Section 2.2.2.6	No downtime
Storage migration ²	Automatic Storage Management	Section 2.2.2.7	No downtime
Migrating to ASM or migrating a single-instance database to Oracle RAC	Oracle Data Guard	Section 2.2.2.2	Seconds to minutes
Patch set and database upgrades	Oracle Data Guard using SQL Apply and logical standby databases	Section 2.2.2.8	Seconds to minutes
Platform Migration Across Windows and Linux Platforms	Oracle Data Guard	Section 2.2.2.8	Seconds to minutes
Platform Migration across same endian format platforms	Transportable Database	Section 2.2.2.9	Minutes to hours
Platform migration across different endian format platforms	Transportable Tablespace	Section 2.2.2.10	Minutes to hours
Application upgrades	Online Application Maintenance and Upgrades	Section 2.2.4	

¹ Patches that cannot be applied by performing a rolling upgrade can be applied with the `MINIMIZE_DOWNTIME` option of the `OPatch` utility to reduce the availability impact of the patch application.

² An example is migration from traditional storage to low-cost storage

See Also:

- *Oracle Data Guard Concepts and Administration* for more information about using Data Guard with SQL Apply to upgrade an Oracle database
- *Oracle Database Concepts* and *Oracle Database Administrator's Guide* for more information about transportable tablespace
- The best practices white papers on rolling upgrades at <http://www.oracle.com/technology/deploy/availability/htdocs/maa.htm>

2.2.2.1 Avoiding Downtime During Operating System and Hardware Upgrades

Using Oracle RAC is the recommended solution for avoiding downtime during system and hardware upgrades.

If you cannot perform the upgrade using Oracle RAC, then Oracle Data Guard and physical standby databases are the recommended solution. See [Section 2.2.2.2](#) for more information.

Oracle RAC Solution Description

Perform the following steps:

1. Stop the application service.
This implicitly redirects connections off of the target instance when using FAN.
2. Shut down target instance or instances with the `IMMEDIATE` option.
3. Shut down and disable Oracle Clusterware.
Disable is used to prevent startup at boot time.
4. Perform maintenance.
5. Enable and start Oracle Clusterware.
This step implicitly starts the database instances.
6. Start the application service.
This step implicitly redirects connections back on to the target instance when using FAN.
7. Repeat all steps on the next node.

Additional Considerations

In addition, verify the following:

- Ensure the planned maintenance can be done in a rolling fashion from an operating system perspective.
- Ensure the database and clusterware versions are certified with the new system and hardware changes.

See Also: Your operating system-specific Oracle Real Application Clusters Installation Guide

2.2.2.2 Using Oracle Data Guard for System and Cluster Upgrades and Migrations

Oracle Data Guard and physical standby databases are the recommended solution for performing system and cluster upgrades that are not upgradeable using Oracle RAC rolling upgrades. Oracle Data Guard is also recommended for migrations to ASM, Oracle RAC, 64-bit systems, Windows to Linux or Linux to Windows, or same processor architecture platforms.

For example, Oracle Data Guard can be used for system upgrades that cannot be upgraded using Oracle RAC rolling upgrades due to system restrictions or if migrating to ASM, from a noncluster environment to Oracle RAC, to a different platform with the same endian format or to a different platform with the same processor architecture.

In general, you first upgrade the physical standby database and then perform a Data Guard switchover to the physical standby database, as follows:

1. Upgrade the system or change the physical standby database system to your target environment.

For example, you can convert the standby database from a single-instance database to an Oracle RAC database by using ASM, without any impact on the primary database. Then, restart the standby database, ensure that it matches your target environment, and wait for Redo Apply to finish applying all redo data to the standby database.

2. Perform a Data Guard switchover—optimally the switchover should take only seconds to minutes.
3. Shut down the original primary database (now the standby database).
4. Upgrade or make system changes to the original primary database.
5. Restart it as a standby database and allow recovery to synchronize the databases.
6. Optionally, perform a Data Guard switchover to return the standby database to the primary database role.

Additional Considerations

- 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 Oracle RAC rolling upgrade or online patching is not possible.
- The conversion from 32 to 64 bit is automatic if you are applying an Oracle Database patch set or doing an Oracle Database upgrade at the same time. If you are upgrading only the operating system, you may need to perform additional post-upgrade steps that are described in Metalink note 414043.1. Also, see the *Oracle Database Upgrade Guide* for more information about upgrades.

See Also: *Oracle Data Guard Concepts and Administration*

2.2.2.3 Oracle Interim Database Patches

Using Oracle RAC is the recommended solution for avoiding downtime when applying Oracle interim database patches.

If you cannot apply patches using Oracle RAC, then Oracle Data Guard and physical standby databases are the recommended solution. See [Section 2.2.2.2](#) for more information.

Solution Description

Oracle interim (one-off) patches to database software are usually applied to implement known fixes for software problems, or to apply diagnostic patches to gather information about a problem. Applying patches is often performed during a schedule maintenance outage.

Oracle provides the capability to do rolling patch upgrades with Oracle RAC with little or no database downtime using the `opatch` command-line utility.

An Oracle RAC rolling upgrade enables all but one of the instances of the Oracle RAC installation to be available during the scheduled outage. This means that the impact on the application downtime required for scheduled outages is further reduced. The Oracle `opatch` utility enables you to apply the patch successively to the different instances in an Oracle RAC installation.

Additional Considerations

Performing a rolling upgrade is possible only 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 Oracle RAC internode 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

Do not use Oracle RAC to perform rolling upgrades of patch sets.

See Also: Your operating system-specific Oracle Real Application Clusters Installation Guide

2.2.2.4 Online Patching

Using Online Patching is the recommended solution for avoiding downtime when an online patch is available.

Solution Description

Online patches are a special type of interim patch that can be applied while the instance remains online.

Oracle provides the capability to do online patching with any Oracle database using the `opatch` command-line utility.

Additional Considerations

- Oracle provides online patches when the changed code is small in scope and complexity, such as with diagnostic patches or small bug fixes.
- Oracle provides online patches when the patch does not change shared memory structures in the System Global Area (SGA), or other critical internal code structures.
- Applying an online patch increases memory consumption on the system because each Oracle process uses more memory from the Program Global Area (PGA) during the patch application. You need to take your memory requirements into consideration before you begin applying an online patch. Each online patch is unique and the memory requirements are patch specific. As is always the case, the best practice is to apply the patch on your test system first. Doing so also enables you to assess the effect of the online patch on your production system and estimate any additional memory usage.

See Also: *Oracle Universal Installer and OPatch User's Guide* for information about online patching and OPatch, and see *Oracle Database Upgrade Guide* for an overview of rolling upgrades and rolling patches

2.2.2.5 Upgrading Oracle Clusterware

Performing rolling upgrades of the Oracle Clusterware using Cluster Ready Services (CRS) software is the recommended solution for avoiding downtime when upgrading Oracle Clusterware.

Solution Description

All upgrades to Oracle Clusterware can be performed in a rolling fashion.

See Also: Your operating system-specific Oracle Clusterware installation guide

2.2.2.6 Upgrading Automatic Storage Management (ASM)

Upgrades ASM in a rolling fashion is the recommended solution for upgrading ASM.

Solution Description

All upgrades starting with Oracle Database 11g (and later releases) can be performed in a rolling fashion.

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

2.2.2.7 Storage Migration

Using ASM is the recommended solution for performing storage migrations.

Solution Description

ASM enables you to add all disks in one storage array and subsequently drop all disks from another array. ASM automatically rebalances and migrates data to the new storage while the database remains operational.

Additional Considerations

Before removing the source storage array, ensure that the rebalancing is complete.

See Also: The chapter about performing ASM Data Migration in the *Oracle Database Backup and Recovery User's Guide*

2.2.2.8 Patch Set and Database Upgrades

Oracle Data Guard using SQL Apply is the recommended solution for performing patch set and database upgrades. [Section 2.2.2.8.1](#) describes this solution.

Oracle Streams is the recommended solution for performing database upgrades when there are data type restrictions that prevent you from performing a rolling upgrade with SQL Apply. [Section 2.2.2.8.2](#) describes this solution.

2.2.2.8.1 Solution Description for Database Upgrades Using Data Guard and SQL Apply

Follow these steps to leverage Data Guard using SQL Apply to upgrade an Oracle database:

1. Upgrade logical standby database to the new release and evaluate the change.
2. Ensure that SQL Apply has applied all redo data to the logical standby database
3. Disconnect applications.
4. Perform Data Guard switchover.
5. Reconnect applications to the new primary database.
6. Shut down the original primary database (now the logical standby database).
7. Execute database software upgrade steps on the new standby database.
8. Restart the standby database and allow recovery to synchronize.
9. Optionally perform a Data Guard Switchover to return to the original database.

Additional Considerations

SQL Apply rolling upgrades are only supported for Oracle Database versions 10.1.0.3 and higher. For complete information, see the chapter about using SQL Apply to upgrade the Oracle Database in *Oracle Data Guard Concepts and Administration*.

SQL Apply has some data type restrictions. For more information, see the appendix about data type and DDL support on a logical standby database in *Oracle Data Guard Concepts and Administration*. If there are data type restrictions, consider using Oracle

Streams. See [Section 2.2.2.8.2, "Solution Description for Database Upgrades That Cannot Use SQL Apply"](#) on page 2-30 for more information.

Oracle Data Guard is the best approach if performing a Oracle RAC rolling upgrade is not possible and there are no data type restrictions.

See Also: The "Rolling Database Upgrades Using Data Guard SQL Apply" white paper available at <http://www.oracle.com/technology/deploy/availability/htdocs/maa.htm>

2.2.2.8.2 Solution Description for Database Upgrades That Cannot Use SQL Apply Streams is similar in function to Data Guard SQL Apply but provides added flexibility if your database includes data types that SQL Apply does not support. Streams enables a method that can be used to work around the existence of unsupported data types and still effect a rolling database upgrade with minimal downtime.

The following high-level steps describe how to perform a database upgrade:

1. Before you begin the upgrade process, see *Oracle Streams Concepts and Administration* for information about how to perform a database upgrade on a database that has user-defined types.
2. Create a duplicate database. (Ideally the replica will start out as a physical standby database that is up to date.)
3. Activate and upgrade the database to the later version.
4. Enable Oracle Streams replication.
5. During the upgrade of the replica, the source database continues ahead. Once the replica is caught up, perform a switchover.

See Also: *Oracle Streams Concepts and Administration* for complete information about online database upgrade with Oracle Streams

2.2.2.9 Platform Migration Across Same Endian Format Platforms

Consider the following approaches when performing platform migrations across same endian format platforms

- Oracle Data Guard (physical standby database) is the recommended solution for performing platform migration across Linux and Windows platforms. This solution is described in [Section 2.2.2.2](#) on page 2-26.
- If cross-platform physical standby database is not available for the platform combination to be migrated, then use transportable database. This solution is described in [Section 2.2.2.9.1, "Solution Description for Platform Migration Using Transportable Database"](#) on page 2-31.
- If transportable database cannot perform the migration quickly enough, then use Oracle Streams as described in [Section 2.2.2.9.2, "Solution Description for Platform Migration Using Oracle Streams"](#) on page 2-31

2.2.2.9.1 Solution Description for Platform Migration Using Transportable Database

Transportable database should be used for platform migration only when cross-platform physical standby database or logical standby database is not supported for the platform combination in question¹.

For example, if you want to move from Windows x86-64 to Linux x86-64, it is best to use cross-platform standby database instead of transportable database. There is less downtime (simply the time it takes to switchover) and it is possible to run the standby database on the new platform for a period of time to ensure that everything is working as planned.

The high-level steps (with target system conversion) are as follows:

1. Place the source database in read/only mode
2. Run `RMAN CONVERT DATABASE` command
3. Move files to the target system
4. Run RMAN generated script to convert datafiles to target platform format
5. Run RMAN generated script to complete the migration

When using transportable database, the downtime required for a platform migration 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.

- Invalidate and recompile all PL/SQL using SQL scripts `utlirp.sql` and `utlrip.sql`.

See Also: The “Platform Migration using Transportable Database” white paper available at <http://www.oracle.com/technology/deploy/availability/htdocs/maa.htm>

2.2.2.9.2 Solution Description for Platform Migration Using Oracle Streams Oracle Streams enables updates on the multiple masters and provides support for heterogeneous platforms with different database releases. Therefore, Oracle Streams may provide the fastest approach for database upgrades and platform migration.

Oracle Streams 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 are 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.

¹ Beginning with Oracle Database release 11g, the primary and standby systems in a Data Guard configuration can have different CPU architectures, operating systems (for example, Windows and Linux), operating system binaries (32-bit and 64-bit), and Oracle database binaries (32-bit and 64-bit). For the latest capabilities and restrictions, see Metalink note 413484.1 at <https://metalink.oracle.com/>.

Oracle Streams implementations require additional administrative effort for testing, setup, and configuration because Streams is designed to be a more flexible architecture.

The following high-level steps describe how to perform a platform migration with Oracle Streams:

1. Set up the Streams environment on the source database.
2. Instantiate the replica database (aka target database) using the new target version or on the target platform.
3. Setup the Streams environment on the target database.
4. Enable Streams to propagate all changes made on the source database to the target database to completely synchronize the target database with the source.
5. Connect users to target database and shutdown source database.
6. Remove Streams configuration.

See Also: *Oracle Streams Concepts and Administration*

2.2.2.10 Platform Migration Across Different Endian Format Platforms

Consider the following approaches when performing platform migrations on different endian format platforms:

- Transportable tablespace is the recommended solution for performing platform migration across different endian format platforms and reduces downtime significantly. See the "[Solution Description for Transportable Tablespace](#)" section on page 2-32 for more details.
- Oracle Data Pump is the simplest of all the approaches. See the *Oracle Database Utilities* for complete information about using Oracle Data Pump.
- For planned downtime of potentially seconds, consider using Oracle Streams as described in [Section 2.2.2.9.2, "Solution Description for Platform Migration Using Oracle Streams"](#) on page 2-31.

Solution Description for Transportable Tablespace

Migrating a database to a new platform using a different endian format with transportable tablespaces requires the following high level steps:

1. Create a new, empty database on the target platform.
2. Import objects required for transport operations from the source database into the target database.
3. Export transportable metadata for all user tablespaces from the source database.
4. Transfer datafiles for user tablespaces to the target system.
5. Use RMAN to convert the datafiles to the target system's endian format.
6. Import transportable metadata for all user tablespaces into the target database.
7. Import the remaining database objects and metadata (that were not moved by the transport operation) from the source database into the target database.

If the target database is being moved to a new location (for example, to a new data center) during the migration, then create a physical standby database from the original 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.

Additional Considerations

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.

Perform a platform migration using transportable tablespaces if all of the following are true:

- The source and target platforms have different endian formats.
- The time required to perform a full Data Pump Export and Import does not fit in the maintenance window.

See Also: The best practices white paper titled: “Oracle Database 10g Release 2 Best Practices: Platform Migration using Transportable Tablespaces” available at <http://www.oracle.com/technology/deploy/availability/htdocs/maa.htm>

2.2.3 Online Reorganization and Redefinition

One way to enhance availability and manageability is to allow customary user access to the database during a data reorganization operation. The Online Reorganization and Redefinition feature in Oracle Database offers administrators significant flexibility to modify the physical attributes of a table and transform both data and table structure while allowing customary user access to the database. This capability improves data availability, query performance, response time, and disk space utilization. All of these 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)

- Fast `ADD COLUMN` with default value (does not need to update all rows to default value)
- Invisible Indexes speed application migration and testing:
 - Speeds up migration with explicit hints, then drops when finished
 - Prevents premature use of newly created indexes
 - Tests effects of `DROP INDEX`, making the index *visible* if needed, thus there is no need for an index rebuild
- Online index build with *no* pause to DML (no exclusive DML locks are required)
- No recompilation of dependent objects when online redefinition does not logically affect objects (for example, when columns are added to tables, or when procedures are added to packages)
- Easier to execute table DDL operations online (there is an option to wait for active DML operations instead of aborting)
- Support for redefinition of tables that have materialized views or materialized view logs

Depending on the type of online reorganization that is required, you can perform the following types of data reorganization 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
- Re-create 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*

2.2.4 Online Application Maintenance and Upgrades

The following sections describe features that can significantly reduce (or eliminate) the application downtime required to make changes to an application's database objects.

2.2.4.1 Oracle Streams for Rolling Upgrades

Consider using Oracle Streams for fast rolling upgrades. However, note that while Oracle Streams upgrades can achieve little or no database down time, your ability to configure this solution will require some operational investment. See [Section 2.1.4, "Oracle Streams"](#) on page 2-6 and Oracle Streams Concepts and Administration for more information.

2.2.4.2 DDL With the WAIT Option

Data definition language (DDL) commands require exclusive locks on internal structures. If DDL commands are issued, these locks may not be available causing the statement to immediately fail even though the DDL could have possibly succeeded sub-seconds later. DDL specified with the `WAIT` option resolves this issue.

DDL with the `WAIT` option is the new `DEFAULT`. The wait time is specified instance-wide (in the initialization parameter file) and can be modified on a session level.

DDL commands specified with the `WAIT` option gives you more flexibility to define grace periods for such commands to succeed instead of raising an error right away, thus requiring additional application logic to handle such errors.

See Also: *Oracle Database Administrator's Guide*

2.2.4.3 ENABLE, DISABLE and FOLLOWS Clauses for CREATE TRIGGER

New states (`ENABLE` and `DISABLE`) and ordering (`FOLLOWS`) are introduced for triggers to control the firing of triggers. With this feature, the `CREATE TRIGGER` statement can be created in a disabled state to validate successful compilation before enabling. In addition, the trigger order can be controlled with the `FOLLOWS` clause.

These additional states allow greater administrative control for triggers.

See Also: *Oracle Database Advanced Application Developer's Guide*

2.2.4.4 Enhanced ADD COLUMN Functionality

Default values of columns are maintained in the data dictionary for columns specified as `NOT NULL`.

Adding new columns with `DEFAULT` values and `NOT NULL` constraint no longer requires the default value to be stored in all existing records. This not only enables a schema modification in sub-seconds and independent of the existing data volume, it also consumes no space.

See Also: *Oracle Database Administrator's Guide*

2.2.4.5 Finer Grained Dependencies

In releases prior to Oracle Database 11g, metadata records mutual dependencies between objects with the granularity of the whole object. For example, PL/SQL unit P depends on PL/SQL unit Q or that view V depends on table T. This means that dependent objects were sometimes invalidated when there was no logical requirement to do so. For example, if view V depends only on columns C1, C2, and C3 in table T and a new column, C99, is added, the validity of view V is not logically affected. Nevertheless, in earlier releases, V was invalidated by the addition of column C99.

Oracle Database 11g records dependency metadata at a finer level of granularity so that the addition of C99 does not invalidate view V. Similarly, if procedure P depends only on elements E1 and E2 in package PKG, then if element E99 is added to PKG, procedure P is not invalidated. (In Oracle Database 10g, this change to PKG would invalidate procedure P.)

By reducing the consequential invalidation of dependent objects in response to changes in the objects they depend upon, application availability is increased. The benefit is felt both in the development environment and when a live application is parsed or upgraded. The benefit occurs when an Oracle Database patch set is applied because changes to schema objects are required to be compatible and, therefore, not cause consequential invalidations.

See Also: *Oracle Database Advanced Application Developer's Guide*

2.2.4.6 Invisible Indexes

An invisible index is an alternative to making an index unusable or even to drop it. An invisible index is maintained for any DML operation but is not used by the optimizer unless you explicitly specify the index with a hint.

Applications often have to be modified without being able to bring the complete application offline. Invisible indexes enable you to leverage temporary index structures for certain operations or modules of an application without affecting the overall application. Furthermore, invisible indexes can be used to test the removal of an index without dropping it right away, thus enabling a grace period for testing in production environments.

See Also: *Oracle Database Administrator's Guide*

2.2.4.7 Materialized View Logging Control

Oracle Database 11g Release 1 has added session-level control for materialized view logs. The capture of changes for materialized views (materialized view logs) can be disabled for an individual session while logging continues for changes made by other sessions. This feature reduces Application patching downtime.

2.2.4.8 Dependent PL/SQL Recompilation After Online Table Redefinition

This feature minimizes the need to recompile dependent PL/SQL packages after an online table redefinition. If the redefinition does not logically affect the PL/SQL packages, recompilation is not needed. This optimization is turned on by default.

This feature reduces the time and effort to manually recompile dependent PL/SQL after an online table redefinition. This also includes views, synonyms, and other table dependent objects (with the exception of triggers) that are not logically affected by the redefinition.

2.3 Optimizing Grid Computing and Disaster Recovery Solutions

With Grid Computing and new standby database capabilities, you can leverage and scale your existing system infrastructure. For the primary database, this implies that all hardware resources are leveraged for performance and scalability. For secondary or disaster recovery systems, system and database resources can be used for real-time query or reporting serving a production purpose while in standby database role. With Oracle Database 11g Release 1 (11.1), Oracle Data Guard can be integral part of your IT operations and application business.

This section covers the following topics:

- [Grid Computing](#)
- [Database Server Grid](#)
- [Database Storage Grid](#)
- [Disaster Recovery Solutions with Better Standby Database Usage](#)

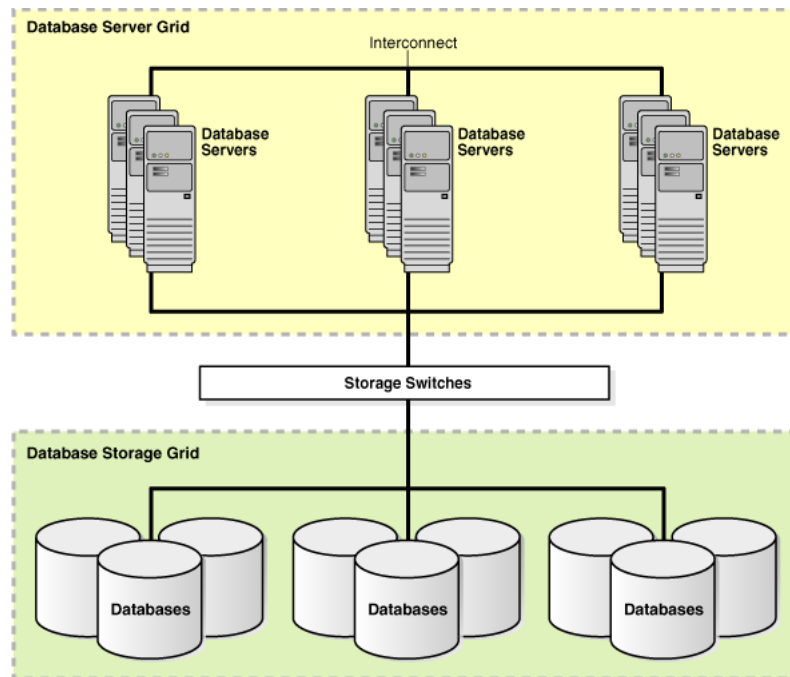
2.3.1 Grid Computing

Grid computing is a computing architecture that effectively pools large numbers of servers and storage into a flexible, on-demand computing resource for all enterprise computing needs.

Oracle Database 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.

With the Database Server and Storage Grid, you can build standby and testing Hubs to leverage a pool of system resources. The system resources can be dynamically allocated and deallocated depending on various priorities. For example, if the production database fails over to one of the standby databases in the standby hub, it will acquire more system and storage resources while the testing resources may be temporarily starved. With Grid technologies, you can enable high level of utilization and low TCO without sacrificing business requirements.

[Figure 2–6](#) illustrates the Database Server Grid and Database Storage Grid in a Grid enterprise computing environment.

Figure 2–6 Grid Computing Environment

2.3.2 Database Server Grid

The availability of low-cost and reliable blade servers, small multiprocessor 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 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. Oracle 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, Oracle 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.

2.3.3 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 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.

2.3.4 Disaster Recovery Solutions with Better Standby Database Usage

Beginning in Oracle Database 11g, standby databases can be used for dynamic IT and application requirements in addition to providing disaster recovery. The real-time query feature in Oracle Data Guard enables you to use physical standby databases for

other useful work during normal operations, in addition to providing a disaster-recovery solution.

The following sections describe the Oracle Data Guard features that help you to leverage physical standby databases for additional business purposes:

- [Real-Time Query Physical Standby Databases](#)
- [Web Scale Using Standby Reader Farms](#)

2.3.4.1 Real-Time Query Physical Standby Databases

Data Guard Redo Apply (physical standby database) has proven to be a popular solution for disaster recovery due to its relative simplicity, high performance, and superior level of data protection. Beginning with Oracle Database 11g, a physical standby database can be open read-only while redo apply is active. This means that you can run queries and reports against an up-to-date physical standby database without compromising data protection or extending recovery time in the event a failover is required. This makes every physical standby database able to support productive uses even while in standby role.

To enable real-time query, open the database in read-only mode and then issue the `ALTER DATABASE RECOVER MANAGED STANDBY` statement. Note that the `COMPATIBLE` parameter must be set to 11.0.0 on both the primary and physical standby databases. Using this feature is totally transparent to applications.

Real-time query provides an ultimate high availability solution because it:

- Supports Oracle RAC on the primary and standby databases
 - Real-time query works on both single-instance and Oracle RAC physical standby databases. Although Redo Apply can be running on only one Oracle RAC instance, you can have all of the instances running in read-only mode while Redo Apply is running on one instance.

- Queries return transactionally consistent results that are very close to being up-to-date with the primary database

Depending on any delay settings or apply rates, the standby database can be lagging seconds behind the primary database. The queries will always be transactionally consistent and will represent a consistent view of the last committed transaction at that time.

- Allows fast switchovers or failovers because the redo generated by the primary database while the standby database was open read-only has already been applied to the standby database, making it immediately available to assume the primary database role
- Enables you to use fast-start failover to allow for automatic fast failover in the case the primary database fails

Note: Transactions that attempt to modify a physical standby database running with real-time query enabled will fail with an error.

See Also: *Oracle Data Guard Concepts and Administration* for complete information about using real-time query

2.3.4.2 Web Scale Using Standby Reader Farms

Beginning with 11gR1, you can use both physical standby databases (using real-time query) and logical standby databases to deploy a *reader farm*. An example of such a configuration is provided [Figure 2–7](#), complete with the use of Data Guard fast-start failover to automatically fail over should the primary database fail. Note that all standby databases in the reader farm automatically recognize the new primary database after such a fast-start failover occurs.

Because a Data Guard configuration can support multiple standby databases, customers have used this capability to boost read performance of the most demanding web applications beyond what the underlying system and storage architecture can support. This provides a relatively low-cost method of scaling out using a Grid architecture where I/O is the driving factor.

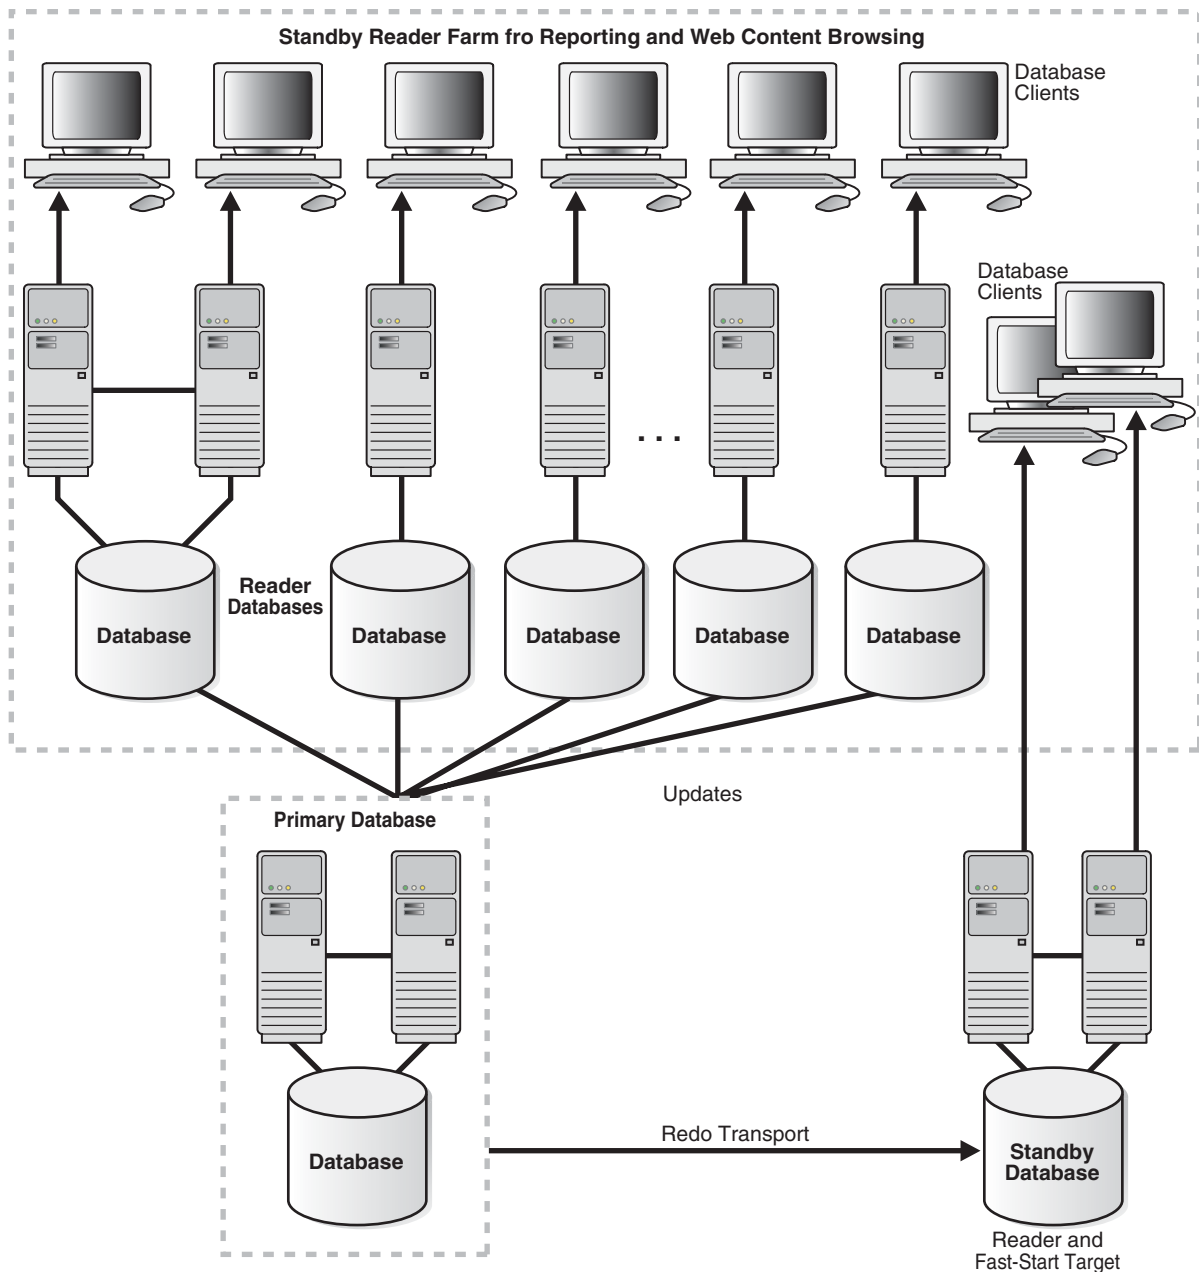
The concept is straightforward—a single primary database that supports read/write transactions, and multiple standby databases that provide read-only access to web users. Such an approach scales read performance linearly as additional standby databases are added. It is also an effective way to isolate faults, because problems that impact one standby database are isolated from the other standby databases in the configuration.

The benefits of creating a reader farm of physical standby databases include the following:

- Fault isolation
- High performance with physical standby databases and Redo Apply
- Seamless support for all DDL and data types using Redo Apply
- All reader databases are kept up-to-date with changes made to the primary database
- Automatic, zero or minimal data loss failover capability
- Management as a unified configuration through Grid Control
- Scale-out using single writer database and n reader databases.

[Figure 2–7](#) shows a good example of how you can leverage Oracle Data Guard, physical standby databases and real-time query to provide the flexibility you need to grow your business quickly, while still providing disaster recovery. In the configuration, the primary database transmits redo data to multiple standby database, one of which is also enabled for fast-start failover for automatic, zero or minimal data loss failover.

Figure 2-7 Standby Database Reader Farms



If a fast-start failover is triggered in the Data Guard configuration in [Figure 2-7](#):

- Automatic failover occurs to the designated standby database
- All standby databases accept data from new primary database
- You can perform a switchover at a convenient time in the future to return all databases to their original roles

2.4 Optimizing Manageability

Complex environments demand coordinated configuration changes, system upgrades and new application roll-outs. Manageability in Oracle Database 11g has improved

dramatically to automate and simplify operations in high availability architectures, and represents a major milestone in the drive toward self-managing Oracle databases.

This section contains these topics:

- [Intelligent Infrastructure](#)
- [Change Assurance](#)
- [Oracle Enterprise Manager Grid Control](#)

2.4.1 Intelligent Infrastructure

Oracle Database has a sophisticated self-management infrastructure that allows the database to learn about itself and use this information to adapt to workload variations or to automatically remedy any potential problem. The self-management infrastructure includes the following:

- Automatic Workload Repository

The Automatic Workload Repository (AWR) is a built-in repository that contains performance statistics used by Oracle Database for problem detection and self-tuning purposes. At regular intervals, Oracle Database makes a snapshot of vital statistics and workload information and stores them in the AWR. The data contained in the snapshots is then analyzed by the Automatic Database Diagnostic Monitor (ADDM). See the *Oracle Database Performance Tuning Guide* for information about the AWR.
- Automatic Maintenance Tasks

By analyzing the information stored in the AWR, the database can identify the need to perform routine maintenance tasks. The automated maintenance tasks infrastructure (known as “AutoTask”) enables Oracle Database to automatically schedule such operations. AutoTask schedules automatic maintenance tasks to run in a set of Oracle Scheduler windows known as maintenance windows. Maintenance windows are those windows that are members of the Oracle Scheduler window group `MAINTENANCE_WINDOW_GROUP`. See the *Oracle Database Administrator’s Guide* and the *Oracle Database 2 Day DBA* for more information.
- Fault diagnosability infrastructure

Oracle Database includes an advanced fault diagnosability infrastructure for preventing, detecting, diagnosing, and resolving problems. The problems that are targeted are critical errors such as those caused by database code bugs, metadata corruption, and customer data corruption. This includes:

 - The automatic diagnostic repository (ADR), which is a file-based repository for database diagnostic data such as traces, the alert log, health monitor reports, and more. It has a unified directory structure across multiple instances and multiple products.
 - The incident packaging services that a DBA can use to automatically and easily gather all diagnostic data (traces, health check reports, SQL test cases, and more) pertaining to a critical error and package the data into a zip file suitable for transmission to Oracle Support.

See the *Oracle Database Administrator’s Guide* for more information about these components.
- Server generated alerts

For problems that cannot be resolved automatically and require administrators to be notified (such as running out of space) the Oracle Database provides

server-generated alerts. Oracle Database can monitor itself and send out alerts to notify you of any problem and provide recommendations on how the reported problem can be resolved. This ensures quick problem resolution and helps prevent potential failures.

- Advisor framework

Oracle Database includes a number of advisors for different subsystems in the database to automatically determine how the operation of the corresponding subcomponents could be further optimized. The SQL Tuning Advisor and the SQL Access Advisor, for example, provide recommendations for running SQL statements faster. Memory advisors help size the various memory components without resorting to trial-and-error techniques. The Segment Advisor handles space-related issues, such as recommending wasted-space reclamation and analyzing growth trends, while the Undo Advisor guides you in sizing the undo tablespace correctly. See the *Oracle Database 2 Day DBA* for more information about using advisors.

- Hang Manager

The Hang Manager is an Oracle Database infrastructure that can detect hangs, analyze them, and then obtain the required diagnostic data from Oracle. The Hang Manager is enabled by default in single-instance databases, Oracle RAC databases, and Automatic Storage Management (ASM) instances. Thus, you can use Hang Manager to follow a hang from the database instance to the underlying ASM instance.

2.4.2 Change Assurance

Oracle Database 11g Release 1 introduces automatic capture and replay of workloads before and after changes so that you can analyze the impact of a database or a SQL change:

- Database Replay

The Database Replay feature addresses enables you to perform real-world testing by capturing the actual database workload on the production system and replaying it on the test system. It also provides analysis and reporting to highlight potential problems (for example, errors encountered and divergence in performance) and recommend ways to remedy the problems.

- SQL Performance Analyzer

SQL performance regression is always a concern during system changes such as database upgrades, initialization parameter changes, and addition or dropping of indexes. The SQL Performance Analyzer feature alleviates this concern by providing an easy way to assess the impact of a change on the performance of SQL statements by comparing and contrasting their response times before and after the change. SQL Replay enables you to capture the SQL workload from the source system, such as the production database, and to replay it on the test system where the change has been applied.

See Also: *Oracle Database Performance Tuning Guide*

2.4.3 Oracle Enterprise Manager Grid Control

By reducing the amount of human intervention required to execute routine and repetitive tasks, services become more stable, reliable, and available. This is particularly important when administrators need to manage very large numbers of systems as efficiently as possible.

Oracle Enterprise Manager Grid Control is an HTML-based 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. If a component within the fast application notification becomes unavailable or experiences performance problems, an alert is automatically generated to Grid Control to inform the administrator so appropriate action can be taken.

The components of Grid Control include:

- Oracle Management Service (OMS)
The OMS is now a set of J2EE applications that renders the interface for Grid Control, works with all Management Agents to process monitoring information, and uses the Management Repository as its persistent data store.
- Oracle Management Agents
These are processes deployed on each monitored host to monitor all targets on the host, communicate that information to OMS, and maintain the host and its targets.
- Oracle Management Repository
This is a schema in Oracle Database that contains all available information about administrators, targets, and applications managed by Grid Control.

Communication between Grid Control, the OMS, and Oracle Management Agents is done through HTTP. Also, you can enable SSL 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 Grid Control.

Using 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
- Interconnect monitoring of cluster databases

See Also:

- *Oracle Enterprise Manager Grid Control Installation and Basic Configuration* and *Oracle Enterprise Manager Concepts* for more information about Oracle Enterprise Manager Grid Control
- The best practices white papers on configuring Enterprise Manager for high availability at

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

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](#)

3.1 Why It Is Important to Determine High Availability Requirements

Any enterprise that is designing and implementing a high availability strategy must begin by performing a thorough analysis of the business drivers that require high availability. Implementing high availability may involve critical tasks such as:

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

An analysis of business requirements for high availability combined with an understanding of the level of investment required to implement different high availability solutions enables the development of a high availability architecture that will achieve both business and technical objectives. This chapter provides a simple framework that can be used effectively to evaluate the high availability requirements of a business.

3.2 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](#)
- [Manageability Goal](#)
- [Total Cost of Ownership \(TCO\) and Return On Investment \(ROI\)](#)

3.2.1 Business Impact Analysis

A rigorous business impact analysis identifies the critical business processes in 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 fabrication plants located worldwide. Semiconductor manufacturing is an intensely competitive business requiring huge financial investment that is amortized over high production volumes. The human resources applications used by plant administration is unlikely to be considered as mission critical as the applications that control the manufacturing process in the plant. Failure of the applications supporting the fabrication process will affect production levels and have a direct impact on financial results of 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*.

3.2.2 Cost of Downtime

A complete business impact analysis provides the insight needed to quantify the costs of unplanned and planned downtime. Understanding this cost is essential because this helps prioritize your high availability investment and has a direct influence on the high availability technologies chosen to minimize the downtime risk.

Various reports have been published, documenting the costs of downtime across industry verticals. Examples include costs that 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.

These numbers are staggering and the reasons are obvious. The Internet can connect the business directly to millions of customers. Application downtime can disrupt this connection, cutting off a business from its customers. In addition to lost revenue, downtime can have an equally negative effect on other critical and interdependent business issues such as customer relationships, competitive advantages, legal obligations, industry reputation, and shareholder confidence.

3.2.3 Recovery Time Objective

The business impact analysis will determine your recovery time objective (RTO). RTO is defined as the maximum amount of time that an IT-based business process can be down before the organization starts suffering unacceptable consequences (financial losses, impact to customer satisfaction, reputation, and so on). RTO indicates the downtime tolerance of a business process or an organization in general.

The RTO requirements are driven by 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.

3.2.4 Recovery Point Objective

The business impact analysis also determines your recovery point objective (RPO). RPO is 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 reentered.

3.2.5 Manageability Goal

A *manageability goal* is more subjective than either the RPO or the RTO. It results from an objective evaluation of the skill sets and management resources available in an organization, and the degree to which the organization can successfully manage all elements of a high availability architecture. In a fashion similar to how RPO and RTO measures an organization's tolerance toward downtime or data loss, your manageability goal measures the organization's tolerance to complexity in the IT environment. To the extent that less complexity is a requirement, simpler methods of achieving high availability are preferred over methods that may be more complex to manage, even if the later could attain more aggressive RTO and RPO objectives. Having a good understanding of manageability goals helps organizations differentiate between what is possible and what is practical to implement.

3.2.6 Total Cost of Ownership (TCO) and Return On Investment (ROI)

Understanding TCO and ROI is essential to selecting an high availability architecture that also achieves the business goals of your organization. TCO includes all costs such as acquisition, implementation, systems, networks, facilities, staff, training, and support, over the useful life of the solution chosen. Likewise, the ROI calculation captures all of the financial benefits that accrue to a given high availability architecture.

For example, consider a high availability architecture in which IT systems/storage at a remote standby site remains idle has no other business use that can be served by its standby systems. The only return on investment for the standby site is the cost of downtime avoided by its use in a failover scenario. Contrast this with a different high availability architecture that enables IT systems and storage at the standby site to be used productively while in the standby role (for example, for reports or for offloading the primary system of the overhead of end-user queries), that makes the standby system a production system in its own right. The return on investment of such an

architecture includes both the cost of downtime avoided and the financial benefits that accrue to its use as a production system while in the standby database role.

3.3 High Availability Architecture Requirements

Using the high availability analysis framework, a business can:

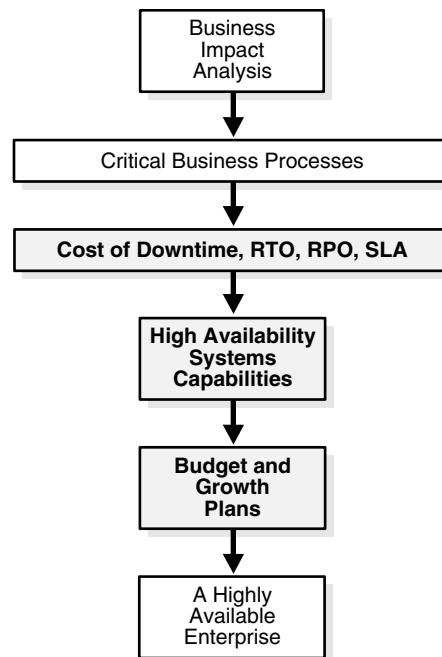
1. Complete a business impact analysis
2. Identify and categorize the critical business processes that have the high availability requirements
3. Formulate the cost of downtime
4. Establish utilization, RTO, and RPO goals for these various business processes.
5. Understand your goals for manageability, TCO, and ROI

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 business processes into several high availability tiers:

- Tier 1 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 processes that have slightly relaxed high availability and RTO/RPO requirements. The second tier of an e-commerce business may be their supply chain and 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 the tier 1 processes.
- Tier 3 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 4-24

3.3.1 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. With Oracle Grid infrastructure, higher availability can equate to lower cost, greater scalability, and more complete utilization of system resources. 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 does 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.

3.3.2 Business Performance, Budget, and Growth Plans

High availability solutions must also be based on 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 are round-trip-delays in the network transmission. These delays increase with distance and vary 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. Conversely, if the business drivers justify the investment to avoid making this tradeoff, a multisite architecture can be implemented that places a synchronous zero data loss standby site in close proximity to the primary site and a second asynchronous standby site located up to thousands of miles away.

High availability solutions must also be based on 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 to extremes with such solutions may lead to budget overruns or an unmanageable and unscalable combination of solutions that is extremely complex and expensive to integrate and maintain.

A high availability solution that has very impressive performance benchmark results may look good in theory. 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 up-front implementation costs
- Forces a vendor lock-in

Cost-conscious and business-savvy decision makers 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.

High-Availability Architectures and Solutions

The Maximum Availability Architecture (MAA) is Oracle's best practices blueprint. It is based on proven Oracle high-availability technologies and recommendations. The goal of the MAA is to remove the complexity in designing the optimal high-availability architecture by providing configuration recommendations and tuning tips to get the most out of your architecture and Oracle features.

This chapter describes the various high-availability architectures in an Oracle environment and helps you to choose the correct architecture for your organization.

It includes the following sections:

- [Oracle Database High-Availability Architectures](#)
- [Choosing the Correct High-Availability Architecture](#)
- [Integrating Application Server High Availability](#)
- [Integrating High Availability for All Applications](#)

4.1 Oracle Database High-Availability Architectures

The following sections provide an overview of the Oracle Database high-availability architectures:

- [Oracle Database](#)
- [Oracle Database with Oracle Clusterware \(Cold Cluster Failover\)](#)
- [Oracle Database with Oracle Real Application Clusters](#)
- [Oracle Database with Oracle RAC on Extended Distance Clusters](#)
- [Oracle Database with Data Guard](#)
- [Oracle Database with Oracle Clusterware and Data Guard](#)
- [Oracle Database with Oracle RAC and Data Guard](#)
- [Oracle Database with Streams](#)

All of these architectures must leverage the MAA best practices.

The comparison of the different architectures highlighting their benefits and considerations is described in [Choosing the Correct High-Availability Architecture](#) on page 4-24.

Once you have chosen an architecture, you can then implement it using the operational and configuration best practices described in the MAA white papers and the *Oracle Database High Availability Best Practices*. These best practices are required to maximize the full benefits of each architecture. See [Chapter 5, "MAA and High](#)

[Availability Best Practices](#)" for more information about the best practices documentation.

4.1.1 Oracle Database

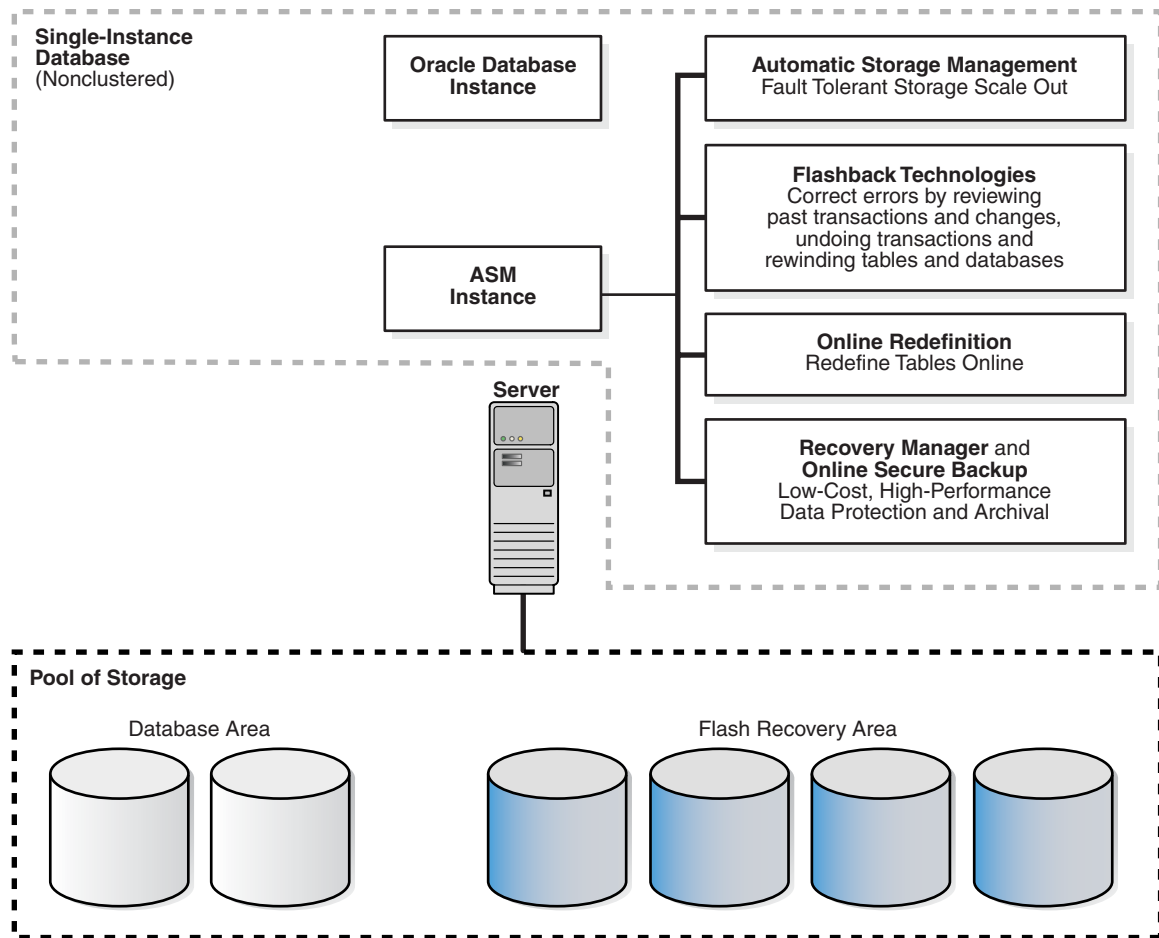
Oracle Database is a single-instance, noncluster database. Although this architecture does not have the node or database redundancy, there are numerous high-availability features that can be used in this architecture and any subsequent database architectures. These features make the standalone database on a single computer attractive and available for certain failures and planned maintenance activities.

Oracle recommends that you leverage the following Oracle features for this architecture. This is the base foundation for subsequent high-availability architectures.

- [Fast-Start Fault Recovery](#) bounds and optimizes instance and database recovery times.
- [Automatic Storage Management](#) tolerates storage failures and optimizes storage performance and usage.
- [Oracle Flashback Technology](#) optimizes logical failure repair. Oracle recommends that you use automatic undo management with sufficient space to attain your desired undo retention guarantee, enable Flashback Database and allocate sufficient space and I/O bandwidth in the flash recovery area.
- [Recovery Manager](#) optimizes local repair of data failures. Oracle recommends that you create and store the local backups in the flash recovery area.
- [Flash Recovery Area](#) manages local recovery related files.
- [Online Reorganization and Redefinition](#) allows for dynamic data changes.
- [Oracle Security Features](#) prevent unauthorized access and changes.
- [Hardware Assisted Resilient Data \(HARD\) Initiative](#) detects and prevents data corruptions and stray or misdirected writes (that result in a lost write to the intended location).
- [Data Recovery Advisor](#) provides intelligent advise and repair of different data failures
- [Data Block Corruption Prevention and Detection Parameters](#) detects and prevents some corruptions and lost writes.
- [Dynamic Resource Provisioning](#) allows for dynamic system changes.
- [Online Patching](#) allows for dynamic database patches of typical diagnostic patches
- [Oracle Secure Backup](#) provides a centralized tape backup management solution.

[Figure 4–1](#) shows a basic, single-node Oracle Database that includes an ASM instance.¹ This architecture takes advantage of several high-availability features, including Flashback Database, Online Redefinition, Recovery Manager, and Oracle Secure Backup.

¹ Single-instance databases can use clustered ASM (Storage GRID) or nonclustered ASM.

Figure 4–1 Single-Node, Nonclustered Oracle Database with an ASM Instance

4.1.2 Oracle Database with Oracle Clusterware (Cold Cluster Failover)

Oracle Clusterware is software that manages the availability of user applications and Oracle databases. The servers on which you want to run Oracle Clusterware must be running the same operating system.

Many high-availability architectures today use clusters alone to provide some rudimentary node redundancy and automatic node failover. However, when you use Oracle Clusterware, there is no need or advantage to using third-party clusterware.

Oracle Clusterware provides a number of benefits over third-party clusterware:

- Oracle Clusterware enables you to use an entire software solution from Oracle, avoiding the cost and complexity of maintaining additional cluster software.

By reducing the number of combinations of software that you need to coordinate and support, you can increase the manageability and availability of your system software.
 - Oracle Clusterware provides seamless integration with, and migration to, Oracle RAC and Oracle Data Guard.
- [Section 4.1.7](#) on page 4-20 describes how you can achieve the highest level of availability with Oracle RAC and Oracle Data Guard
- Oracle Clusterware includes all of the features required for cluster management, including node membership, group services, global resource management, and

high-availability functions such as managing third-party applications, event management, and Oracle notification services that enable Oracle clients to reconnect to the new primary database after a failure.

- Oracle Clusterware uses a private network and a voting disk to detect and resolve *split brain*² scenarios.

With Oracle Clusterware you can provide a *cold failover cluster* to protect an Oracle instance from a system or server failure. The basic function of a cold failover cluster is to monitor a database instance running on a server, and if a failure is detected, to restart the instance on a spare server in the cluster. Network addresses are failed over to the backup node. Clients on the network experience a period of lockout while the failover takes place and are then served by the other database instance once the instance has started. Also, you can use the Oracle Clusterware ability to relocate applications and application resources (using the `CRS_RELOCATE` command) as a way to move the workload to another node so you can perform planned system maintenance on the production server.

The cold cluster failover solution with Oracle Clusterware provides these additional advantages over a basic database architecture:

- Automatic recovery of node and instance failures in minutes
- Automatic notification and reconnection of Oracle integrated clients³
- Ability to customize the failure detection mechanism.

For example, you can use your favorite application query in the database check action. Providing application-specific failure detection means Oracle Clusterware can fail over not only during the obvious cases such as when the instance is down, but also in the cases when, for example, an application query is not meeting a particular service level.

- High availability functionality to manage third-party applications
- Rolling release upgrades of Oracle Clusterware

The operation of an Oracle Clusterware cold failover cluster is depicted in [Figure 4-2](#) and [Figure 4-3](#). These figures show how you can use the Oracle Clusterware framework to make both the Oracle database and your custom applications highly available.

[Figure 4-2](#) shows a configuration that uses Oracle Clusterware to extend the basic Oracle Database architecture and provide cold cluster failover. In the figure, the configuration is operating in normal mode in which Node 1 is the active instance connected to the Oracle Database that is servicing applications and users. Node 2 is connected to Node 1 and to the Oracle Database, but it is currently standby mode.

² Network splits, commonly referred to as split brains, occur when nodes on each side of the cluster cannot see the nodes on the other side of the cluster.

³ Oracle Clusterware sends the service events and FAN-integrated clients automatically react to those events.

Figure 4–2 Oracle Database with Oracle Clusterware (Before Cold Cluster Failover)

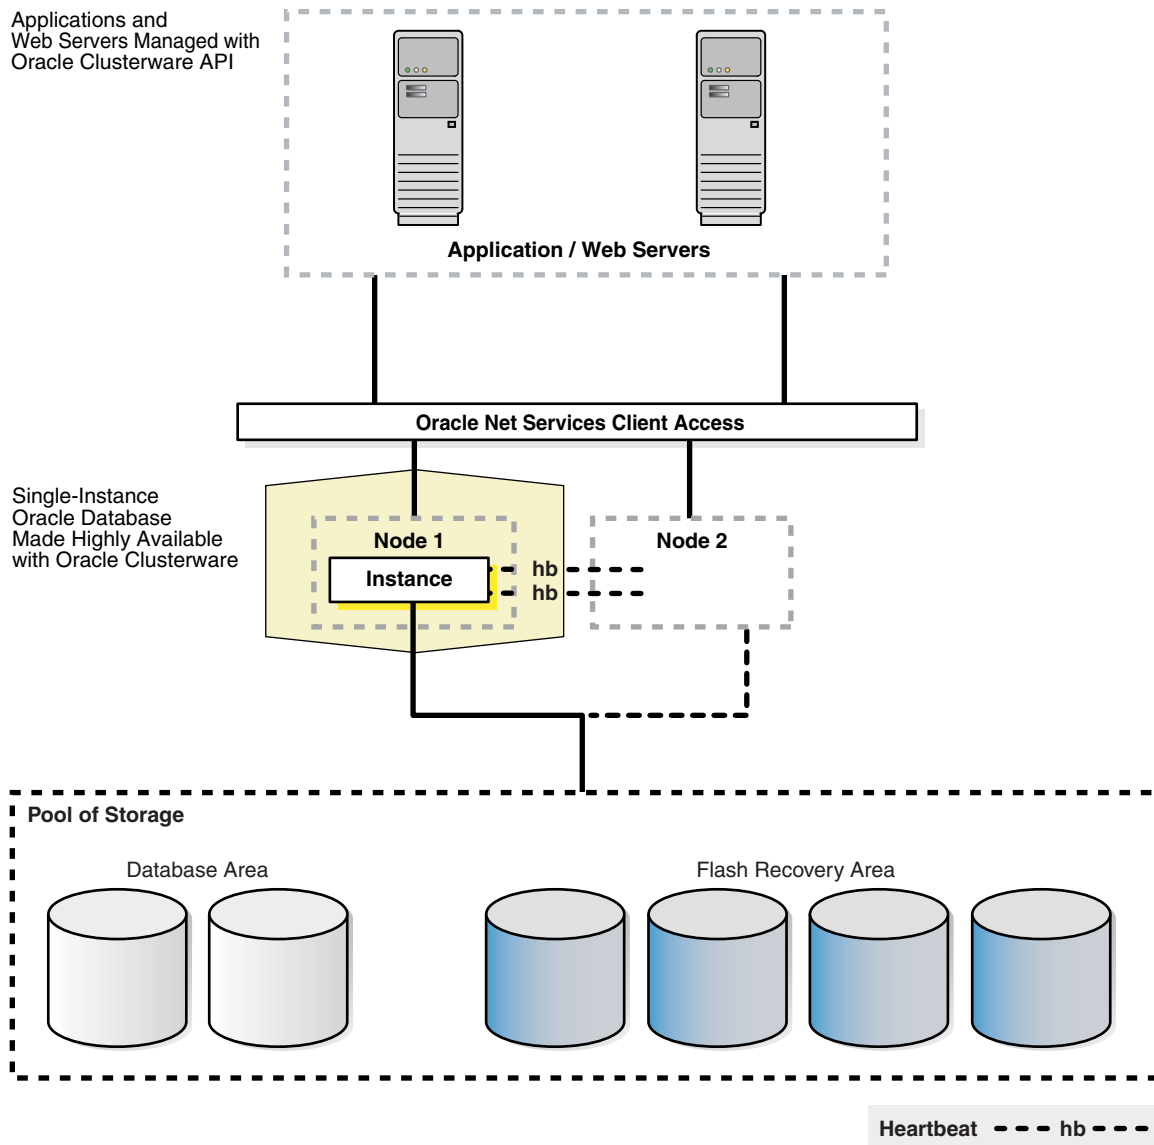
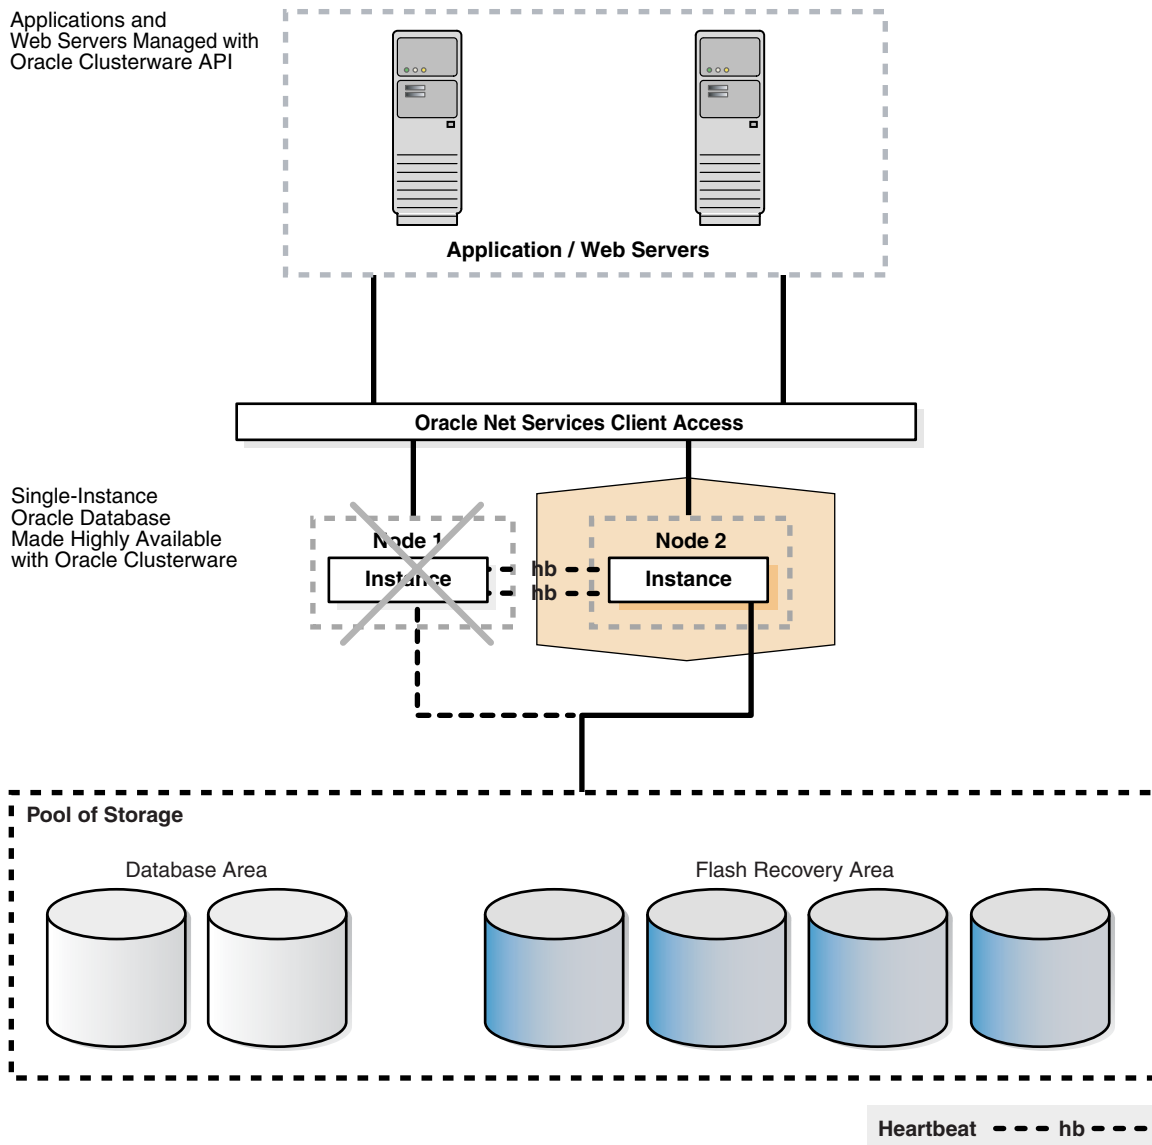


Figure 4–3 the Oracle Clusterware configuration after a cold cluster failover has occurred. In the figure, Node 2 is now the active instance connected to the Oracle Database and servicing applications and users. Node 1 is connected to Node 2 and to the Oracle Database but Node 1 is currently idle, in standby mode.

To provide this transparent failover capability, Oracle Clusterware requires a virtual IP address for each node in the cluster. With Oracle Clusterware you also define an *application* virtual IP address so users can access the application independently of the node in the cluster where the application is running. You can define multiple application VIPs, with generally one application VIP defined for each application running. The application VIP is tied to the application by making it dependent on the application resource defined by Cluster Ready Services (CRS).

Figure 4–3 Oracle Database with Oracle Clusterware (After Cold Cluster Failover)



Note: Neither Oracle Enterprise Manager or Oracle Universal Installer (OUI) provide configuration support for Oracle Clusterware. To configure an Oracle Clusterware environment, follow the step-by-step instructions in your platform-specific Oracle Clusterware installation guide.

4.1.3 Oracle Database with Oracle Real Application Clusters

An architecture that combines the Oracle Database with Real Application Clusters (Oracle RAC) is inherently a highly available system. Unlike a traditional monolithic database server that is expensive and is not flexible to changing capacity and resource demands, Oracle RAC combines the processing power of multiple interconnected computers to provide system redundancy, scalability, and high availability.

The clusters that are typical of Oracle RAC environments can provide continuous service for both planned and unplanned outages. Oracle RAC builds higher levels of

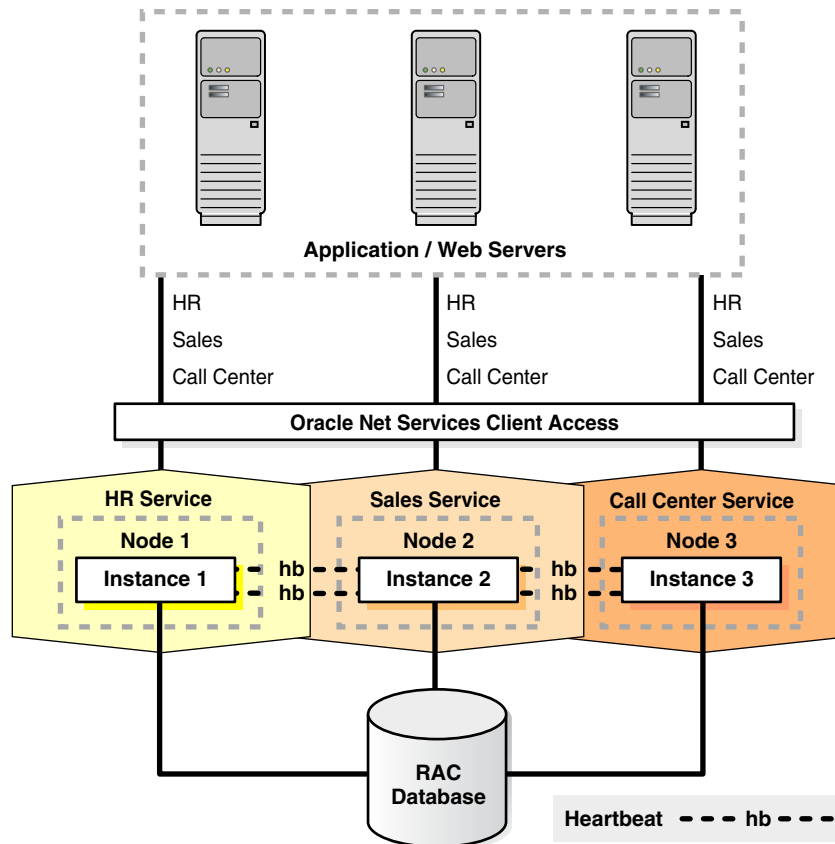
availability on top of the standard Oracle features. All single instance high-availability features, such as the Flashback technologies and online reorganization, also apply to Oracle RAC. Applications scale in an Oracle RAC environment to meet increasing data processing demands without changing the 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 RAC exploits the redundancy that is provided by clustering to deliver availability with $n - 1$ node failures in an n -node cluster. Unlike the cold cluster model where one node is completely idle, all instances and nodes can be active to scale your application.

The Oracle Database with Oracle RAC architecture provides the following benefits over a traditional monolithic database server and the cold cluster failover model:

- Scalability across database instances
- Flexibility to increase processing capacity using commodity hardware without downtime or changes to the application
- Ability to tolerate and quickly recover from computer and instance failures (measured in seconds)
- Rolling upgrades for system and hardware changes
- Rolling patch upgrades for some interim patches
- Fast, automatic, and intelligent connection and service relocation and failover
- Load balancing advisory and runtime connection load balancing
- Comprehensive manageability integrating database and cluster features

[Figure 4-4](#) shows the Oracle Database with Oracle RAC architecture.

Figure 4–4 Oracle Database with Oracle RAC Architecture

4.1.4 Oracle Database with Oracle RAC on Extended Distance Clusters

The Oracle Database with Oracle RAC architecture is designed primarily as a scalability and availability solution that resides in a single data center. It is possible, under certain circumstances, to build and deploy an Oracle RAC system where the nodes in the cluster are separated by greater distances. This architecture is referred to as an *extended distance cluster*.

An Oracle RAC extended distance cluster is an architecture that provides extremely fast recovery from a site failure and allows for all nodes, at all sites, to actively process transactions as part of single database cluster. For example, if a customer has a corporate campus the extended Oracle RAC configuration could consist of individual Oracle RAC nodes being located in separate buildings. Oracle RAC on an extended distance cluster provides greater high availability than a local Oracle RAC cluster, but it may not fit the full disaster recovery requirements of your organization.

When the two data centers are located relatively close, extended distance clusters can provide great protection for some disasters, but not all. You should do an analysis to determine if both sites are likely to be affected by the same disaster. For example, if the extended cluster configuration is set up properly, it can provide protection against disasters such as a local power outage, an airplane crash, or server room flooding. However, they cannot protect against comprehensive disasters such as earthquakes, hurricanes, and regional floods that affect a greater area. (For disaster recovery, use the architecture described in [Section 4–9, "Oracle Database with Oracle RAC and Data Guard - MAA"](#).)

The advantages to using Oracle RAC on extended distance clusters include:

- Ability to fully use all system resources without jeopardizing the overall failover times for instance and node failures
- Extremely rapid recovery if one site should fail
- All of the Oracle RAC benefits listed in [Section 4.1.3](#) on page 4-6

Note: While this architecture can be effective and it has been successfully implemented, you should implement it only in the environments (distance, latency, and degree of protection) recommended in this discussion.

When configuring the extended cluster architecture, Oracle recommends that you:

- Use ASM normal or high redundancy so that a storage array failure does not affect the application and database availability.

Beginning with Oracle Database Release 11g, ASM includes a preferred read capability that ensures that a read I/O accesses the local storage instead of unnecessarily reading from a remote failure group. When you configure ASM failure groups extended distance clusters, you can specify that a particular node read from a failure group extent that is closest to the node, even if it is a secondary extent. This is especially useful in extended distance clusters where remote nodes have asymmetric access with respect to performance, thus leading to better usage and lower network loading.

See Also: *Oracle Database Storage Administrator's Guide* for information about configuring preferred read failure groups with the `ASM_PREFERRED_READ_FAILURE_GROUPS` initialization parameter.

- Add a third voting disk to a third site

Because most extended distance cluster have only two storage systems (one at each site). During normal processing, each node writes and reads a disk heartbeat at regular intervals, but if the heartbeat cannot complete the node exits, generally causing a node restart. Thus, the site that houses the majority of the voting disks is a potential single point of failure for the entire cluster. For availability reasons, you should add a third site that can act as the arbitrator in case either one of the sites fail or a communication failure occurs between the sites.

To build an Oracle RAC database on an extended distance cluster environment, you must:

- Configure one set of nodes at Site A.
- Configure another set of nodes at Site B.
- Use a fast, redundant dedicated connection between the nodes (or buildings) for Oracle RAC cross-instance communication.

You can optionally configure Dense Wavelength Division Multiplexing (referred to as DWDM, or Dark Fiber) to allow communication to occur between the sites without using repeaters and to allow greater distances between the sites. However, the disadvantage is that Dark Fiber can be prohibitively expensive.

- Use server-based or array-based mirroring to host all of the data on both sites and keep it synchronously mirrored. Oracle recommends server-based mirroring using ASM to internal mirror across the two storage arrays. Implementing mirroring with ASM provides an active/active storage environment in which system write

I/Os are propagated to both sets of disks, making them appear as a single set of disks independent of location.

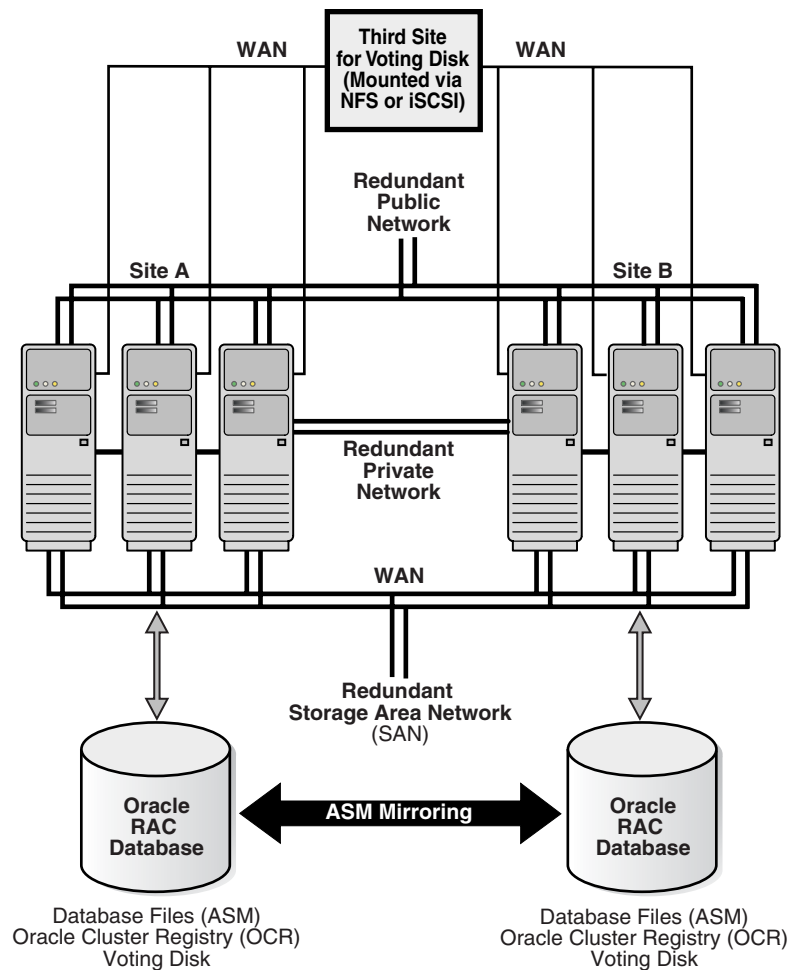
The ASM volume manager provides flexible server-based mirroring redundancy options. You can choose to use external redundancy to defer the mirroring protection function to the hardware RAID storage subsystem. The ASM normal and high-redundancy options allow two-way and three-way mirroring, respectively.

- Configure a third site for a voting disk⁴

Figure 4–5 shows an Oracle RAC extended distance cluster for a configuration that has multiple active instances on six nodes at two different locations: three nodes at Site A and three at Site B. The public and private interconnects, and the Storage Area Network (SAN) are all on separate dedicated channels, with each one configured redundantly. For availability reasons, the Oracle Database is a single database that is mirrored at both of the sites. Also, to prevent a full cluster outage if either site fails, the configuration includes a third voting disk on an inexpensive, low-end standard Network File System (NFS) mounted device.

⁴ If you have an extended distance cluster and do not configure a third site, you must make one of the sites the primary site and make the other site a secondary site. Then, if the primary site fails, you must manually restart the secondary site. See the white paper about using standard NFS to support a third voting disk on a stretch cluster configuration that is available on the Oracle Real Application Clusters Web site at <http://www.oracle.com/technology/products/database/clustering/index.html>

Figure 4-5 Oracle RAC On an Extended Distance Cluster



Finally, consider the following when implementing this architecture:

- High internode and interstorage latency can have a major effect on performance and throughput. Performance testing is mandatory to assess the impact of latency. In general, distances of 50 km or less are recommended.
- Network, storage, and management costs will increase
- Write performance incurs the overhead of network latency
- Because this is a single database without Oracle Data Guard, there is no protection from data corruption or data failures
- A third site is recommended for high availability to be leveraged as another location for the voting disk (quorum disk) and as an arbitrator in case of connection issues between the nodes

See Also: The white paper about extended distance clusters on the Oracle Real Application Clusters Web site at <http://www.oracle.com/technology/products/database/clustering/>

4.1.5 Oracle Database with Data Guard

Oracle Data Guard is a high availability and disaster-recovery solution that provides very fast automatic failover (referred to as fast-start failover) in the case of database failures, node failures, corruption, and media failures. Furthermore, the standby databases can be used for read-only access and subsequently for reader farms, for reporting purposes, and for testing and development purposes.

While traditional solutions (such as backup and recovery from tape, storage based remote mirroring, and database log shipping) can deliver some level of high availability, Data Guard provides the most comprehensive high availability and disaster recovery solution for Oracle databases.

Data Guard provides a number of advantages over traditional solutions, including the following:

- Fast, automatic or automated failover for data corruptions, lost writes, and database and site failures
- Protection against data corruptions and lost writes on the primary database
- Reduced downtime with Data Guard rolling upgrade capabilities
- Ability to offload primary database activities, such as backups, queries or reporting without sacrificing RTO and RPO
- Site failures do not require instance restart, storage remastering, or application reconnections
- Transparent to applications
- Effective network utilization

In addition, for data resident in Oracle databases, Oracle Data Guard, with its built in zero data loss capability, is more efficient, less expensive and better optimized for data protection and disaster recovery than traditional remote mirroring solutions. Oracle Data Guard provides a compelling set of technical and business reasons that justify its adoption as the disaster recovery and data protection technology of choice, over traditional remote mirroring solutions.

The following list summarizes the advantages of using Oracle Data Guard compared to using remote mirroring solutions:

- **Better Network Efficiency**—With Oracle Data Guard, only the redo data needs to be sent to the remote site. However, if a remote mirroring solution is used for data protection, typically you must mirror the database files, the online redo logs, the archived redo logs and the control file. If the flash recovery area is on the source volume that is remotely mirrored, then you must also remotely mirror the flashback logs. This means that compared to Data Guard, a remote mirroring solution sends each change many more times to the remote site.
- **Better Performance**—Data Guard only transmits writes to the redo logs of the primary database, whereas remote mirroring solutions must transmit these writes and every write I/O to data files, additional members of online log file groups, archived redo log files, and control files. Data Guard is designed so that it does not affect the Oracle DBWR process that writes to data files, because anything that slows down DBWR impacts database performance. However, remote mirroring solutions do impact DBWR performance because they subject all DBWR writes to network and disk I/O induced delays inherent to synchronous, zero-data-loss configurations. Compared to mirroring, Data Guard provides better performance and is more efficient, Data Guard always verifies the state of the standby database and validates the data before applying redo, and Data Guard enables you to use

the standby database for updates while it continues to protect the primary database.

- **Better suited for WANs**—Remote mirroring solutions based on storage systems often have a distance limitation due to the underlying communication technology (Fibre Channel, ESCON) used by the storage systems. In a typical example, the maximum distance between these two boxes connected in a point-to-point fashion and running synchronously can be only 10 km. Using specialized devices this distance can be extended to 66 km. However, when the standby data center is more than 66 km apart, you have to use a series of repeaters and converters from third-party vendors. These devices convert ESCON/Fibre Channel to the appropriate IP, ATM or SONET networks.
- **Better resilience and data protection**—Oracle Data Guard ensures much better data protection and data resilience than remote mirroring solutions, because corruptions introduced on the production database probably can be mirrored by remote mirroring solutions to the standby site, but corruptions are eliminated by Data Guard. For example, if a stray write occurs to a disk, or there is a corruption in the file system, or the Host Bus Adaptor corrupts a block as it is written to disk, then a remote mirroring solution may propagate this corruption to the DR site. Because Data Guard only propagates the redo data in the logs, and the log file consistency is checked before it is applied, all such external corruptions are eliminated by Data Guard.
- **Higher Flexibility**—Data Guard is implemented on top of pure commodity hardware. It only requires a standard TCP/IP-based network link between the two computers. There is no fancy or expensive hardware required. It also allows the storage to be laid out in a different fashion from the primary. For example, you can put the files on different disks, volumes, file systems, and so on.
- **Better Functionality**—Data Guard, with its full suite of data protection features (Redo Apply for physical standby databases and SQL Apply for logical standby databases, multiple protection modes, push-button automated switchover and failover capabilities, automatic gap detection and resolution, GUI-driven management and monitoring framework, cascaded redo log destinations), is a much more comprehensive and effective solution optimized for data protection and disaster recovery than remote mirroring solutions.
- **Higher ROI**—Businesses have to ensure that they are getting as much value as possible from their IT investments, and no IT infrastructure is sitting idle. Data Guard is designed to allow businesses get something useful out of their expensive investment in a disaster-recovery site. Typically, this is not possible with remote mirroring solutions.

The recommended high availability and disaster-recovery architectures that leverage Oracle Data Guard are described in the following sections:

- [Overview of Single Standby Database Architectures](#)
- [Overview of Multiple Standby Database Architectures](#)

4.1.5.1 Overview of Single Standby Database Architectures

A single standby database architecture consists of the following key traits and recommendations:

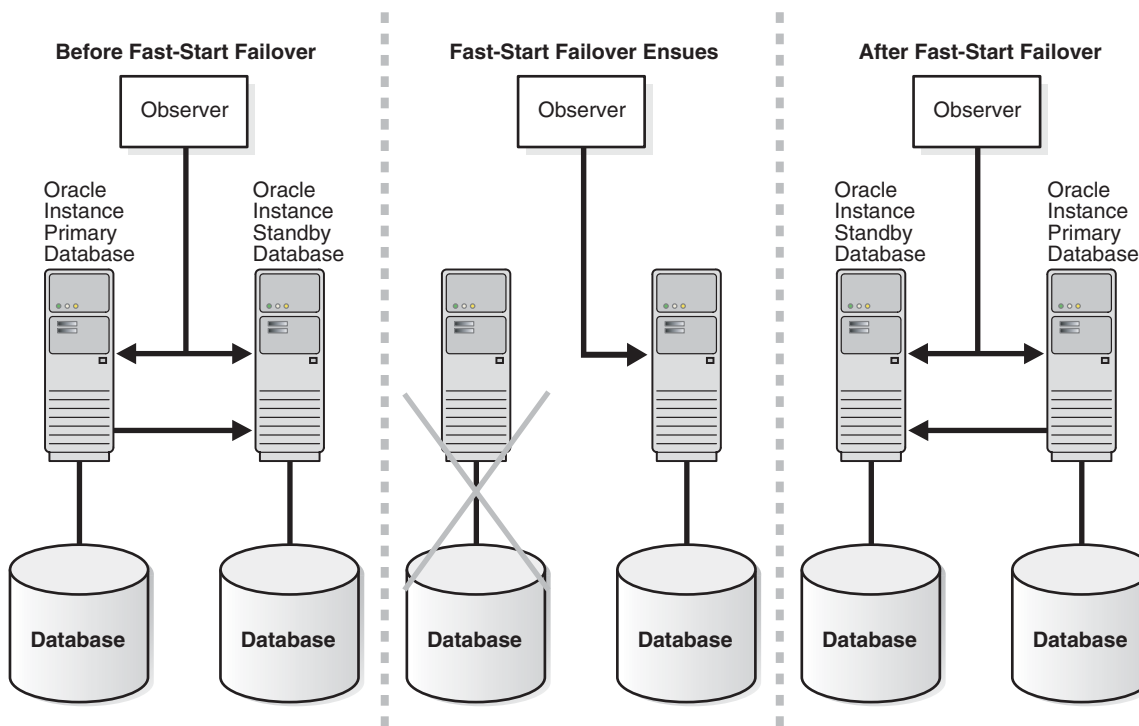
- Primary database resides in Site A.
- Standby database resides in Site B. If zero data loss is required with minimum performance impact on the primary database, the best practice is to locate the

secondary site within 200 miles from the primary database. Note, however, that the synchronous redo transport does not impose any physical distance limitation.

- Fast-start failover is recommended to provide automatic failover without user intervention and bounded recovery time. If the primary database uses the asynchronous redo transport, configure your maximum data loss tolerance or the Data Guard broker's `FastStartFailoverLagLimit` property to meet your business requirements. The Observer (thin client watchdog) resides in the application tier and monitors the availability of the primary database. The observer is described in more detail in *Oracle Data Guard Broker*.
- Use a physical standby database if read-only access is sufficient.
- Evaluate logical standby databases if additional indexes are required for reporting purposes and if your application only uses data types supported by logical standby database and SQL Apply.

Figure 4–6 shows the relationships between the primary database, target standby database, and the observer before, during, and after a fast-start failover occurs.

Figure 4–6 Relationship of Primary and Standby Databases and the Observer During Fast-Start Failover



The following list describes examples of Data Guard configurations using single standby databases:

- A national energy company uses a standby database located in a separate facility 10 miles away from its primary data center. Outages or data loss that could impact customer service and safety are avoided by using Data Guard synchronous transport and automatic failover (fast-start failover).
- An infrastructure services provider to the telecommunication industry utilizes a single standby database located over 400 miles away from the primary configured for synchronous redo transport, enabling zero data loss failover for maximum data protection and high availability.

- A telecommunications provider uses asynchronous redo transport to synchronize a primary database on the west coast of the United States, with a standby database on the east coast, over 2,200 miles away. This enables the provider to use existing data centers that are geographically isolated, offering a unique level of high availability.
- A global manufacturing company used Data Guard to replace storage-based remote mirroring and maintain a standby database at its recovery site 50 miles away from the primary site. Data Guard provides more comprehensive data protection and its more efficient network utilization means there is plenty of headroom to grow without incurring the additional expense of upgrading their network.

4.1.5.2 Overview of Multiple Standby Database Architectures

This architecture is identical to the single-standby database architecture that was described in [Section 4.1.5.1](#), except that there are multiple standby databases in the same Data Guard configuration. The following list describes some implementations for a multiple standby database architecture:

- Continuous and transparent disaster or high-availability protection in case of an outage at the primary database or the targeted standby database
- Reader farms or look up databases
- Reporting databases
- Regional reporting or reader databases for better response time
- Synchronous transport transmits to a more *local* standby database, and asynchronous transport transmits to a more *remote* standby database to provide optimum levels of performance and data protection
- Testing and development clones using snapshot standby databases
- Rolling upgrades

Note that it is possible to convert a physical standby database to a logical standby database or to a snapshot standby database, or you can create additional logical standby databases or snapshot standby databases:

- **Transient logical standby databases** can be used to minimize downtime for database upgrades. Using transient logical standby databases is helpful in Data Guard architectures where there are no logical standby databases.

In a multiple standby database environment, you can create a transient logical standby database temporarily (for planned maintenance) and then convert it back to the physical standby database role. For example, you can use transient logical standby databases to minimize downtime for database upgrades, when required. There is no need to create a separate logical standby database to perform upgrades. The high-level steps for rolling upgrades with a transient logical standby database are as follows:

1. Start performing a rolling database upgrade with the physical standby database.
2. Temporarily convert the physical standby database to a logical standby database to perform the upgrade. (Note that data type restrictions are limited for the short window of time required to perform an upgrade.)
3. Revert the logical standby database back to the physical standby database role.

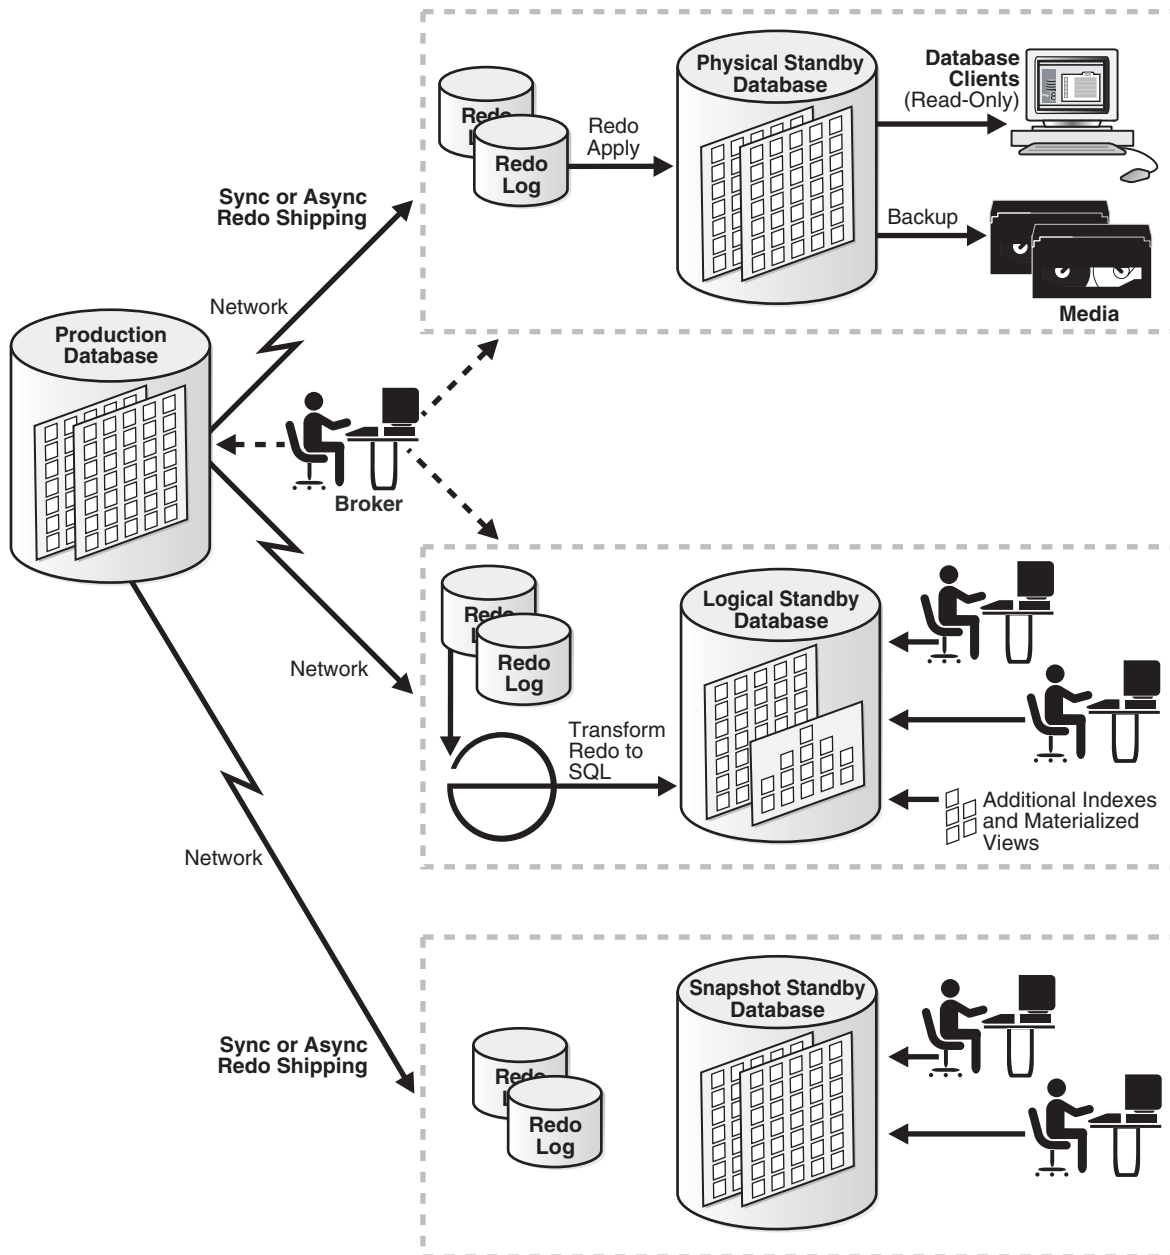
See Also: *Oracle Data Guard Concepts and Administration* or the *Oracle Database High Availability Best Practices* for step-by-step instructions about performing a rolling upgrade with a transient logical standby database

- **Snapshot standby databases** can be used as a clone or a test database to test new functionality and new releases. The snapshot standby database continues to receive and queue redo data so data protection and RPO are not sacrificed.

Snapshot standby databases diverge from the primary database over time because redo data from the primary database is not applied when it is received. Redo Apply does not apply the redo data until you convert the snapshot standby database back into a physical standby database, and all local updates that were made to the snapshot standby database are discarded. Although the local updates to the snapshot standby database cause additional divergence, the data in the primary database is fully protected by means of the redo logs that are located at the standby site.

[Figure 4-7](#) shows the production database at the primary site and multiple standby databases at secondary sites. Also, see [Figure 2-7, "Standby Database Reader Farms"](#) on page 2-41 for another example of a multiple standby database environment.

Figure 4–7 Oracle Database with Data Guard Architecture on Primary and Multiple Standby Sites

**See Also:**

- *Oracle Data Guard Concepts and Administration* for more information about the various types of standby databases and to find out what datatypes are supported by logical standby databases
- The white papers about Oracle Data Guard and standby databases at

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

The following list describes examples of Data Guard configurations using multiple standby databases:

- A world-recognized financial institution uses two remote physical standby databases for continuous data protection after failover. If the primary system should fail, the first standby database becomes the new primary. The second standby database automatically receives data from the new primary, insuring that data is protected at all times.
- A nationally recognized insurance provider in the U.S. maintains two standby databases in the same Data Guard configuration, one physical and one logical standby database. Their strategy further mitigates risk by maintaining multiple standby databases, each implemented using a different architectures - Redo Apply and SQL Apply.
- A world-recognized e-commerce site utilizes multiple standby databases - a mix of both physical and logical databases - both for disaster recovery purposes and to scale-out read performance by provisioning multiple logical standby databases using SQL Apply.
- A global provider of information services to legal and financial institutions uses multiple standby databases in the same Data Guard configuration to minimize downtime during major database upgrades and platform migrations.

Also, for large data centers where there is a need to support many applications with Data Guard requirements, you can build a Data Guard hub to reduce the total cost of ownership.

With the Database Server and Storage Grid, you can build standby and testing Hubs that leverage a pool of system resources. The system resources can be dynamically allocated and deallocated depending on various priorities. For example, if the primary database fails over to one of the standby databases in the standby hub, the new primary database acquires more system and storage resources while the testing resources may be temporarily starved. With the Oracle Grid technologies, you can enable a high level of utilization and low TCO, without sacrificing business requirements.

A Data Guard hub can consists of:

- Several standby databases in an Oracle RAC environment residing in a cluster of servers, called a grid server
- Leveraging the storage grid

The premise of the standby hub is that it provides higher utilization with lower cost. The probability of failing over all the databases at the same time is unlikely. Thus, when there is a failover, you can prioritize the system resources to production activity and allocate new system resources in a grid for the standby database functions. At the time of role transition, more storage and system resources can be allocated toward that application.

For example, a Data Guard hub could include multiple databases and applications that are supported in a Grid server and storage architecture. This configuration consists of a central resource supporting 10 applications and databases in the grid compared to managing 10 separate system or storage units in a non-grid infrastructure.

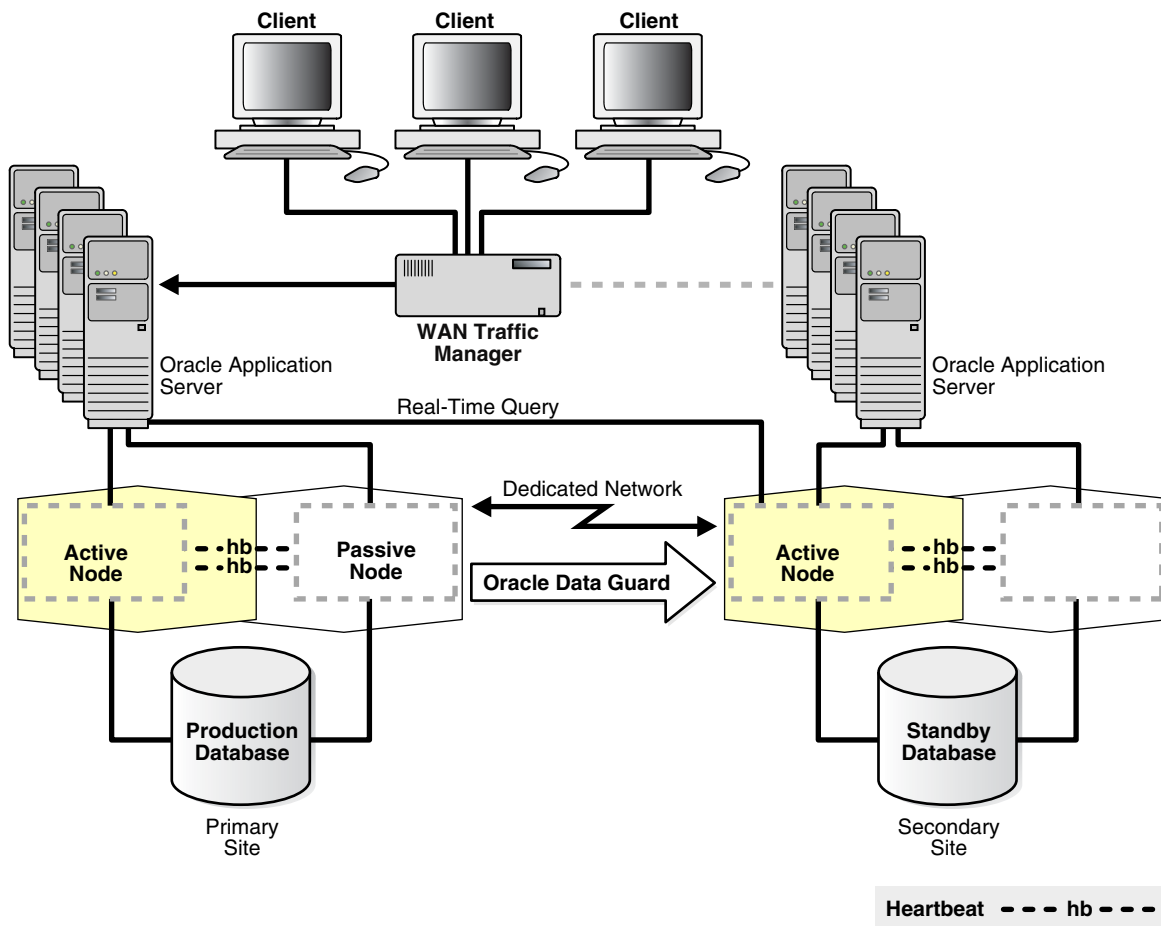
Another possible configuration might be a testing hub consisting of snapshot standby databases. With the snapshot standby database hub, you can leverage the combined storage and server resources of a Grid instead of building and managing individual servers for each application.

4.1.6 Oracle Database with Oracle Clusterware and Data Guard

If your business does not require the scalability and additional high availability benefits provided by Oracle RAC, but you still need all the benefits of Oracle Data Guard and cold cluster failover, then this architecture is a good compromise. With Oracle 11g, Oracle Clusterware cold cluster failover combined with Oracle Data Guard makes a tightly integrated solution in which failover to the secondary node in the cold cluster failover is transparent and does not require you to reconfigure the Data Guard environment or perform additional steps.

Figure 4–8 shows an Oracle Clusterware and Oracle Data Guard architecture that consists of a primary and a secondary site. Both the primary and secondary sites contain Oracle application servers, two database instances, and an Oracle Database.

Figure 4–8 Oracle Clusterware (Cold Cluster Failover) and Oracle Data Guard



In Figure 4–8:

- The application servers on the secondary site are connected to the WAN traffic manager by a dotted line to indicate that they are not actively processing client requests at this time. The application server on the secondary site can be active and processing client requests such as queries if the standby database is a physical standby database with real-time query enabled, or if it is a logical standby database.
- Oracle Data Guard transmits redo data from the primary database to the secondary site to keep the databases synchronized.

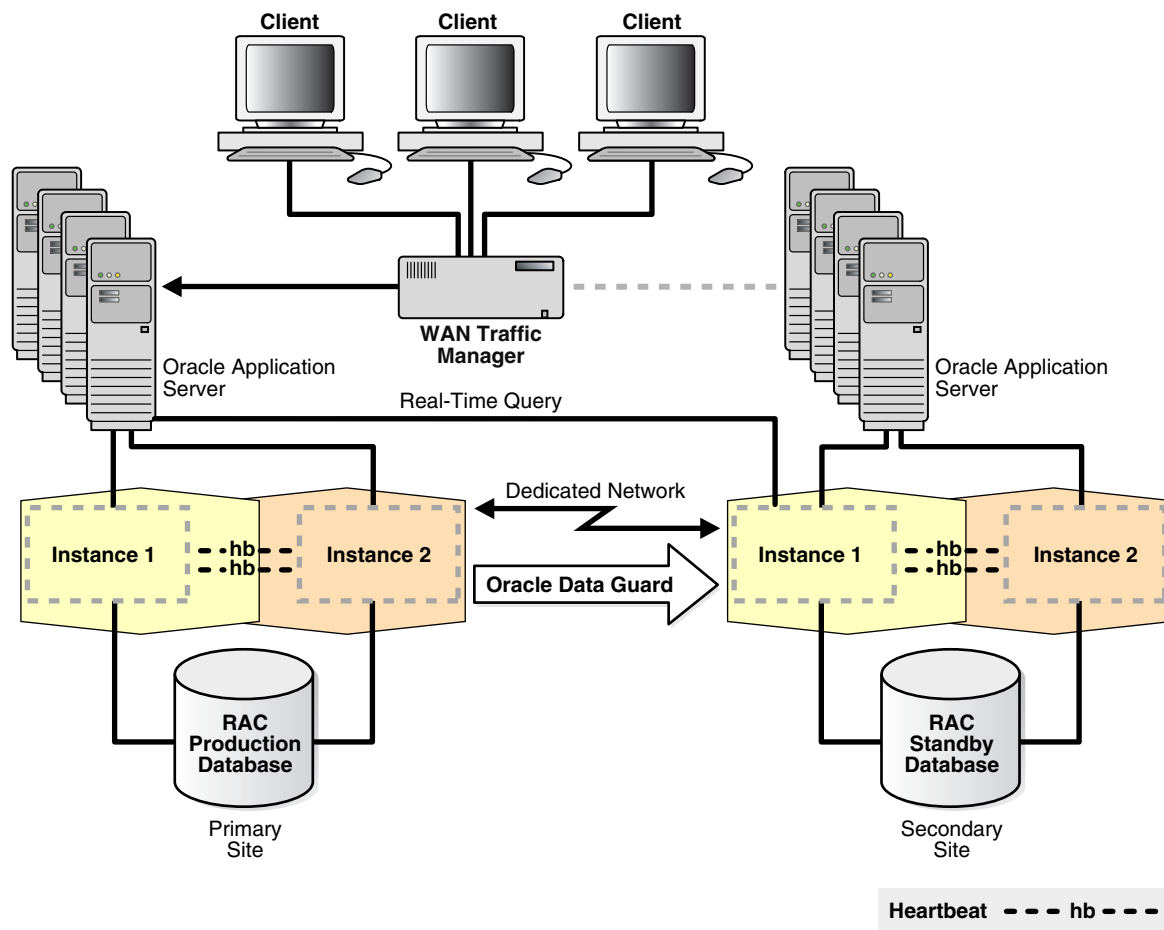
- Oracle Clusterware manages the availability of both the user applications and Oracle databases.
- Oracle Clusterware provides tolerance of node failures, while Data Guard provides additional protection against data corruptions, lost writes, and database and site failures. (See [Oracle Database with Data Guard](#) on page 4-12 for a complete description.)
- Although Cold Cluster Failover is not shown in [Figure 4-8](#), you can configure it by adding a passive node on the secondary site.

4.1.7 Oracle Database with Oracle RAC and Data Guard

You can achieve the highest level of availability when using Oracle RAC and Oracle Data Guard without application changes. These Oracle features provide the most comprehensive architecture for reducing downtime for scheduled outages and preventing, detecting, and recovering from unscheduled outages. This architecture combines the benefits of both Oracle RAC and Data Guard and it is the recommended architecture for Maximum Availability Architecture (MAA).

To protect against site failures, the MAA recommends Oracle RAC and Data Guard reside on separate systems (clusters) and data centers. [Figure 4-9](#) shows the recommended MAA configuration, with Oracle Database, Oracle RAC, and Data Guard. Configuring symmetric sites is recommended to ensure that each site can accommodate the performance and scalability requirements of the application after any role transition. Furthermore, operational practices across role transitions is simplified when the sites are symmetric.

Figure 4–9 Oracle Database with Oracle RAC and Data Guard - MAA



4.1.8 Oracle Database with Streams

Like Oracle Data Guard in SQL Apply mode, Oracle Streams can capture database changes, propagate them to destinations, and apply the changes at these destinations. Streams is optimized for replicating data. Streams can capture changes at a source database, and the captured changes can be propagated asynchronously to replica databases. A logical copy configured and maintained using Streams is called a *replica*, not a logical standby database, because it provides many capabilities that are beyond the scope of the normal definition of a standby database.

You might choose to use Streams to configure and maintain a logical copy of your production database. Although using Streams might require additional work, it offers increased flexibility that might be required to meet specific business requirements.

Oracle Database 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.

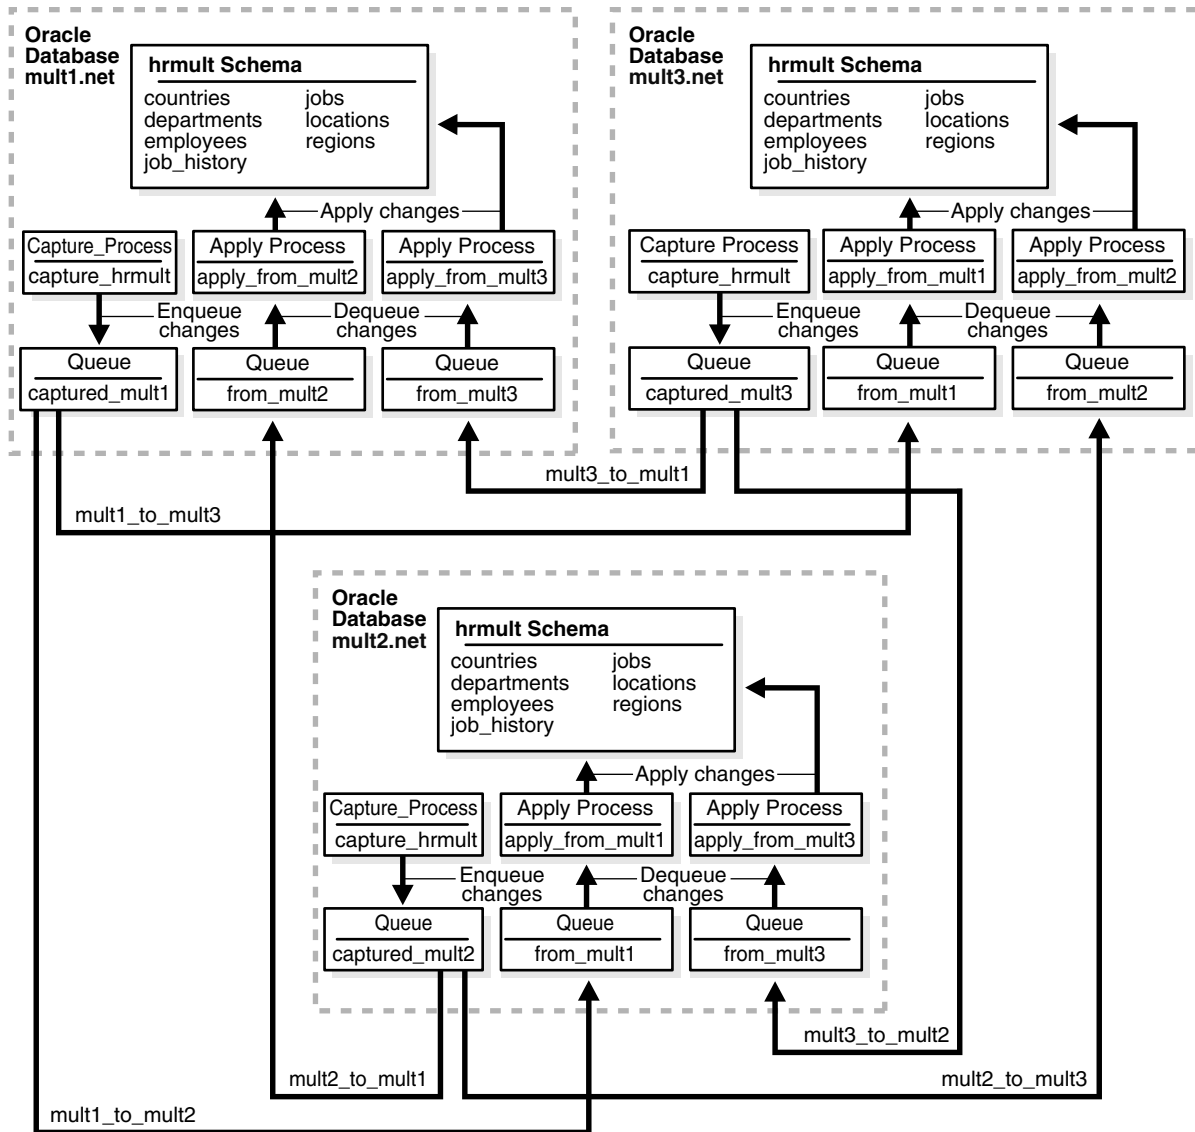
Consider using Oracle Database with Streams if one or more of the following conditions are true:

- Updates are required on both sites or databases, and the changes need to be propagated bidirectionally
- 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

Figure 4–10 shows a sample Oracle Database using Streams to replicate data for a schema among three Oracle databases. DML and DDL changes made to tables in the hr schema are captured at all databases in the environment and propagated to each of the other databases in the environment.

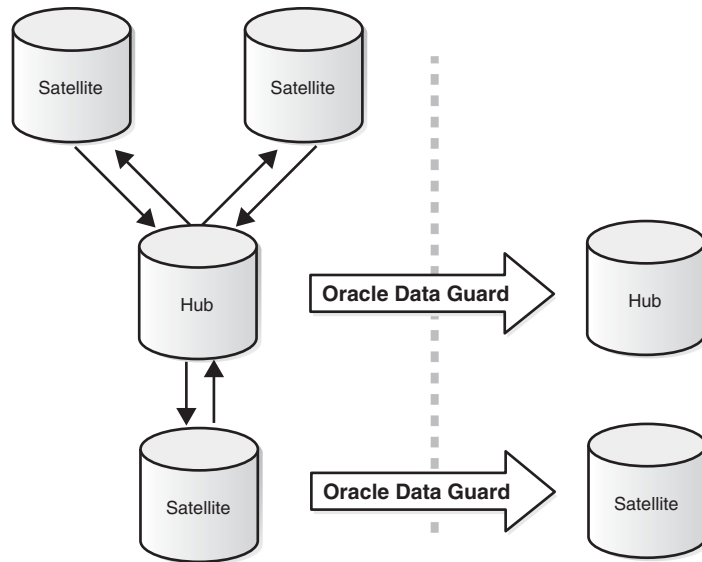
See Also: *Oracle Streams Replication Administrator's Guide* for complete information about constructing multiple-source replication environments using Streams

Figure 4–10 Oracle Database with Streams Architecture That Shares Data From Multiple Databases



You can configure Streams with Data Guard to provide protection for the individual databases in the configuration. Figure 4–11 shows a hub and spoke network configuration in which Oracle Data Guard is providing additional data protection for the hub and one of the satellites.

Figure 4–11 Oracle Streams Hub and Spoke Network Configuration



4.2 Choosing the Correct High-Availability Architecture

This section summarizes the advantages of the different high-availability architectures and provides guidelines for you to choose the correct high-availability architecture for your business.

[Chapter 3, "Determining Your High Availability Requirements"](#) describes how the high-availability requirements for the business plus its allotted budget determine the appropriate architecture. The key factors include:

- Recovery time objective (RTO) and recovery point objective (RPO) for unplanned outages and planned maintenance
- Manageability Overhead (MO)
- Total Cost of Ownership (TCO) and Return On Investment (ROI)

For example, [Table 4–1](#) provides some insight into the probability of different outages during unplanned and planned activities. The data is derived from actual user experiences and from Oracle service requests.

Table 4–1 Frequency of Outages

Activity	Outage
Media or disk failures	High
Application patches	High
Application failures	High
Logical or user failures that manipulate logical data (DMLs and DDLs)	High
Data corruptions (hardware or software induced)	Medium
Computer failures	Medium
Database patches	Medium
Hardware patches and upgrades	Low
Operating system patches and upgrades	Low

Table 4–1 (Cont.) Frequency of Outages

Activity	Outage
Database or application upgrades	Low
Database failures	Low
Platform migrations	Very low
Site failures	Very low

Table 4–2 recommends architectures based on your business requirements for RTO, RPO, MO, scalability, and other factors.

Table 4–2 High-Availability Architecture Recommendations

Consider Using ..	Business or Application Impact ...
Oracle Database with Oracle Clusterware (Cold Cluster Failover)	<ul style="list-style-type: none"> ▪ Maximum RTO for instance or node failure is in minutes ▪ MO is low ▪ ROI is low ▪ Rolling upgrade and patch capabilities for Oracle Clusterware with zero database downtime.
Oracle Database with Oracle Real Application Clusters	<ul style="list-style-type: none"> ▪ Maximum RTO for instance or node failure is zero for the database¹ ▪ MO is medium ▪ ROI is high ▪ Rolling upgrade for system, clusterware, operating system and some Oracle interim patches ▪ Database scalability beyond one instance or node
Oracle Database with Oracle RAC on Extended Distance Clusters	<ul style="list-style-type: none"> ▪ All of the business benefits of Oracle Real Application Clusters ▪ MO is high² ▪ ROI is medium ▪ Additional protection from data center failure with special considerations that are documented in "Oracle Database with Oracle RAC on Extended Distance Clusters" on page 4-8
Oracle Database with Data Guard	<ul style="list-style-type: none"> ▪ Maximum RTO for instance or node failure is in seconds to minutes ▪ Maximum RTO for data corruptions, database, or site failures is in seconds to minutes ▪ MO is low ▪ ROI is high ▪ Rolling upgrade for system, clusterware, database, and operating system ▪ Offload read-only, reporting, testing and backup activities to the standby database
Oracle Database with Oracle Clusterware and Data Guard	<ul style="list-style-type: none"> ▪ All of the business benefits of Oracle Clusterware (Cold Cluster Failover) and Oracle Data Guard ▪ MO is low ▪ ROI is medium

Table 4–2 (Cont.) High-Availability Architecture Recommendations

Consider Using ..	Business or Application Impact ...
Oracle Database with Oracle RAC and Data Guard	<ul style="list-style-type: none"> ■ All of the business benefits of Oracle RAC and Oracle Data Guard ■ MO is medium ■ ROI is high
Oracle Database with Streams	<ul style="list-style-type: none"> ■ Maximum RTO for instance or node failure is in seconds to minutes ■ Maximum RTO for data corruption, cluster, database, or site failures is in seconds to minutes ■ MO is high² ■ ROI is high ■ Rolling upgrade for system, clusterware, operating system, database and application ■ Replication between multiple updatable databases or support for heterogeneous platforms ■ Flexible propagation and management of data, transactions, and events ■ With Oracle RAC integration, database scalability is possible

¹ Database is still available, but a portion of the application connected to the failed system is temporarily affected.

² Architectures for which the MO is "High" might require additional time and expertise to build and maintain, but offer increased flexibility and capabilities required to meet specific business requirements.

[Table 4–3](#) identifies the additional capabilities provided by the architectures that build on the Oracle Database and attempts to label each architecture with its greatest strengths.

Table 4–3 Additional Capabilities of High Level Oracle High-Availability Architectures

Oracle High-Availability Architecture	Key Characteristics and Additional Capabilities
<p>Oracle Database (Base Architecture) The foundation for all high-availability architectures</p>	<ul style="list-style-type: none"> ▪ Fast-Start Fault Recovery bounds and optimizes instance and database recovery times to minutes. ▪ Automatic Storage Management tolerates storage failures and optimizes storage performance and utilization. ▪ Oracle Flashback Technology optimizes logical failure repair. ▪ Recovery Manager optimizes local repair of data failures using local backups. ▪ Oracle Secure Backup provides a centralized tape backup management solution. ▪ Flash Recovery Area manages local recovery related files automatically. ▪ Online Reorganization and Redefinition allows for dynamic data changes. ▪ Oracle Security Features prevents unauthorized access and changes. ▪ Hardware Assisted Resilient Data (HARD) Initiative detects and prevents data corruptions and stray or misdirected writes (that result in a lost write to the intended location). ▪ Data Block Corruption Prevention and Detection Parameters detects and prevents some corruptions and lost writes. ▪ Dynamic Resource Provisioning allows for dynamic system changes. ▪ Online Patching allows for dynamic database patches of typically diagnostic patches. ▪ Data Recovery Advisor diagnoses persistent (on disk) data failures, presents appropriate repair options, and runs repair operations at your request. Support is for single-instance databases only.
<p>Oracle Database with Oracle Clusterware (Cold Cluster Failover)</p>	<ul style="list-style-type: none"> ▪ All of the benefits of Oracle Database ▪ Automatic and fast failover for computer failure ▪ Minimum rolling upgrade capabilities for system, clusterware, and operating system¹
<p>Oracle Database with Oracle Real Application Clusters High availability, scalability, and foundation of server database grids</p>	<ul style="list-style-type: none"> ▪ All of the benefits of Oracle Database ▪ Scalability beyond a single system ▪ Automatic recovery of failed nodes and instances ▪ Fast application notification (FAN) with integrated Oracle client failover ▪ Rolling upgrade for system, clusterware, operating system and some Oracle interim patches¹
<p>Oracle Database with Oracle RAC on Extended Distance Clusters Database Grid with site failure protection</p>	<ul style="list-style-type: none"> ▪ All of the benefits of Oracle Database ▪ Protection from site failure

Table 4–3 (Cont.) Additional Capabilities of High Level Oracle High-Availability Architectures

Oracle High-Availability Architecture	Key Characteristics and Additional Capabilities
<p>Oracle Database with Data Guard Simplest high availability, data protection, and disaster-recovery solution</p>	<ul style="list-style-type: none"> ■ All of the benefits of Oracle Database ■ Automatic and fast failover for computer failure, storage failure, data corruption, for configured ORA- errors or conditions and database failures ■ Protection from site failure ■ Rolling upgrade for system, clusterware, database, and operating system² ■ Offload backups to the standby database ■ Offload read and reporting workload to the standby database ■ Only comprehensive lost write protection
<p>Oracle Database with Oracle Clusterware and Data Guard Simple high availability solution with added data and disaster recovery protection.</p>	<ul style="list-style-type: none"> ■ The sum of benefits of Oracle Clusterware with Data Guard
<p>Oracle Database with Oracle RAC and Data Guard Best high-availability, data protection and disaster-recovery solution with scalability built in</p>	<ul style="list-style-type: none"> ■ The sum of benefits of Oracle RAC with Data Guard
<p>Oracle Database with Streams³ Bidirectional replication and information management</p>	<ul style="list-style-type: none"> ■ Replica database (or databases) are available for read/write use ■ Provides heterogeneous platform support ■ Fast failover for computer failure and storage failure ■ Protection from site failure ■ Minimizes downtime for computer or site maintenance and database and application upgrades

¹ Rolling upgrades with Oracle Clusterware and Oracle RAC incur zero downtime.

² Rolling upgrades with Oracle Data Guard incur minimal downtime.

³ The initial investment to build a robust solution is well worth the long-term flexibility and capabilities that Streams delivers to meet specific business requirements.

Table 4–4 shows the recovery time including detection and client failover time of an integrated Oracle client, whenever relevant. You should adopt the MAA best practices to achieve the optimal recovery time and configuration. Oracle High Availability Best Practice recommendations can be found in the *Oracle Database High Availability Best Practices* and in the white papers that can be downloaded from:

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

Table 4–4 Attainable Recovery Times for Unplanned Outages

Outage Type	Oracle Database	Cold Cluster	Oracle RAC and RAC on Extended Distance Clusters	Data Guard	Oracle RAC and Data Guard	Streams
Computer failure	Minutes to hours ¹	Minutes	No downtime ²	Seconds to a minute	No downtime ²	No downtime ²
Storage failure	No downtime ³	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	< 30 minutes ⁴
Data corruption	Potentially hours ⁵	Potentially hours ⁵	Potentially hours ⁵	Seconds to a minute	Seconds to a minute	Seconds to a minute
Site failure	Hours to days	Hours to days	No downtime ² if the outage affects one building Hours to days if the outage affects building	Seconds to a minute ⁶	Seconds to a minute ⁶	No downtime ⁷

- ¹ Recovery time consists largely of the time it takes to restore the failed system
- ² Database is still available, but a portion of the application connected to the 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 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.
- ⁵ 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.
- ⁶ Recovery time indicated applies to database and existing connection failover. Network connection changes and other site-specific failover activities may lengthen overall recovery time.
- ⁷ The portion of any application connected to the failed system is temporarily affected. You can configure the failed application connections to fail over to the replica.

Table 4–5 compares the attainable recovery times of each Oracle high-availability architecture for all types of planned downtime.

Table 4–5 Attainable Recovery Times for Planned Outages

System Change or Data Change	Outage Type	Oracle Database	Oracle RAC	Data Guard	MAA	Streams
System change - Dynamic Resource Provisioning	--	No downtime	No downtime	No downtime	No downtime	No downtime
System change - Rolling Upgrade	System level upgrade	Minutes to hours	No downtime	Seconds to five minutes	No downtime	No downtime
System change - Rolling Upgrade	Cluster or site wide upgrade	Minutes to hours	Minutes to hours	Seconds to five minutes	Seconds to five minutes	No downtime ¹

Table 4–5 (Cont.) Attainable Recovery Times for Planned Outages

System Change or Data Change	Outage Type	Oracle Database	Oracle RAC	Data Guard	MAA	Streams
System change - Rolling Upgrade	Storage migration	No downtime ²	No downtime ²	No downtime ²	No downtime ²	No downtime ²
System change - Rolling Upgrade	Database one-off patch	Minutes to an hour	No downtime ³	Seconds to five minutes	No downtime ³	No downtime
System change - Rolling Upgrade	Database patch set and version upgrade	Minutes to hours	Minutes to hours	Seconds to five minutes	Seconds to five minutes	No downtime ¹
System change - Rolling Upgrade	Platform migration	Minutes to hours	Minutes to hours	Minutes to hours	Minutes to hours	No downtime ¹
Data change	Online Reorganization and Redefinition	No downtime	No downtime	No downtime ⁴	No downtime ⁴	No downtime ⁴

¹ Applications (or a portion of an application) connected to the system that is being maintained may be temporarily affected.

² ASM automatically rebalances stored data when disks are added or removed while the database remains online. For storage migration, you are required to leverage both storage arrays 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 database or replica database. For more information, see *Oracle Data Guard Concepts and Administration* or *Oracle Streams Replication Administrator's Guide*.

4.3 Integrating Application Server High Availability

The Oracle Application Server provides flexible and automated high availability solutions for Oracle Application Server to ensure that applications that you deploy on Oracle Application Server meet the required availability to achieve your business goals. The solutions introduced in this book are described in detail in the *Oracle Application Server High Availability Guide*.

This section contains the following topics:

- [Oracle Application Server High Availability Architectures](#)
- [Redundant Architectures](#)
- [High Availability Services in Oracle Application Server](#)

4.3.1 Oracle Application Server High Availability Architectures

Oracle Application Server provides high availability and disaster recovery solutions for maximum protection against any kind of failure with flexible installation, deployment, and security options. The redundancy of Oracle Application Server local high availability and disaster recovery originates from its redundant high availability architectures.

At a high level, Oracle Application Server local high availability architectures include several active-active and active-passive architectures for the OracleAS middle-tier and

the OracleAS Infrastructure. Although both types of solutions provide high availability, active-active solutions generally offer higher scalability and faster failover, although, they tend to be more expensive as well. With either the active-active or the active-passive category, multiple solutions exist that differ in ease of installation, cost, scalability, and security.

Building on top of the local high-availability solutions is the Oracle Application Server disaster recovery solution, Oracle Application Server Guard. This unique solution combines the proven Oracle Data Guard technology in the Oracle Database with advanced disaster recovery technologies in the application realm to create a comprehensive disaster recovery solution for the entire application system. This solution requires homogenous production and standby sites, but other Oracle Application Server instances can be installed in either site as long as they do not interfere with the instances in the disaster recovery setup. Configurations and data must be synchronized regularly between the two sites to maintain homogeneity.

4.3.2 Redundant Architectures

Oracle Application Server provides redundancy by offering support for multiple instances supporting the same workload. These redundant configurations provide increased availability either through a distributed workload, through a failover setup, or both.

From the entry point to an Oracle Application Server system (content cache) to the back end layer (data sources), all the tiers that are crossed by a request can be configured in a redundant manner with Oracle Application Server. The configuration can be an active-active configuration using OracleAS Cluster or an active-passive configuration using OracleAS Cold Failover Cluster.

4.3.3 High Availability Services in Oracle Application Server

Oracle Application Server provides different features and topologies to support high availability across the its stack. This includes solutions that extend across both the OracleAS middle-tier and the OracleAS Infrastructure tier.

The *Oracle Application Server High Availability Guide* describes the following high availability services in Oracle Application Server in detail:

- Process death detection and automatic restart
- Configuration management
- State replication
- Server load balancing and failover
- Backup and recovery
- Disaster recovery

4.4 Integrating High Availability for All Applications

A highly available and resilient application requires that every component of the application must be highly available or tolerate failures and changes. For example, a highly available application must analyze every component that affects the application including the network topology, application server, application flow and design, systems, and the database configuration and architecture. This book has focused primarily on the database high availability solutions.

See the high availability solutions and recommendations for Oracle Application Server, Enterprise Manager and Applications on the MAA Web site at:

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

MAA and 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, enjoys an improved operational resilience and enhanced business agility.

Once you have chosen a high-availability architecture as described in [Chapter 4](#), you can implement it using the MAA best practices documentation. The common set of best practices is described in the MAA and high availability white papers located on the MAA OTN Web site and in the *Oracle Database High Availability Best Practices* documentation:

- **MAA and high availability best practices white papers**

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 from

<http://www.oracle.com/technology/deploy/availability/htdocs/maa.htm>
- ***Oracle Database High Availability Best Practices***

This book provides detailed best practice recommendations and information. It can help you to configure a new high-availability environment, or migrate an existing configuration to create a redundant, reliable system without sacrificing simplicity and performance.

The book's step-by-step method of describing the 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.

Index

A

- advisor framework, 2-43
- ALTER DATABASE RECOVER MANAGED STANDBY statement
 - enabling real-time query, 2-39
- applications
 - defining a virtual IP address, 4-5
 - online maintenance and upgrades, 2-35
- architecture
 - recommendations, 4-25
- architectures
 - extended distance Oracle RAC, 4-8
 - Oracle Application Server, 4-30
 - single-instance Oracle Database (noncluster), 4-2
- ASM_PREFERRED_READ_FAILURE_GROUPS
 - initialization parameter, 4-9
- automatic failover
 - Oracle Data Guard, 4-14
- automatic maintenance tasks, 2-42
- automatic storage management
 - description, 2-13
- Automatic Storage Management (ASM)
 - failure groups, 4-9
 - Fast Mirror Resync, 2-14
 - server-based mirroring, 4-9
 - storage migration, 2-29
 - using normal or high redundancy, 4-9
 - volume manager, 4-10
 - with Database Storage Grid, 2-38
- Automatic Workload Repository (AWR), 2-42
- availability
 - definition, 1-1

B

- backing out a transaction, 2-11
- backups
 - Oracle Secure Backup, 2-15
- block recovery
 - using Flashback logs
 - Flashback technologies
 - block recovery using Flashback logs, 2-13
 - block recovery using Flashback logs, 2-13

C

- Cluster Ready Services (CRS)
 - avoiding downtime during upgrades, 2-28
- cold cluster failover, 4-19
 - Oracle Clusterware and Data Guard, 4-19
 - with Oracle Clusterware, 4-3
- cold failover cluster
 - described, 4-4
- computer failure, 1-3
- corruptions
 - prevention and detection of block corruption, 2-20

D

- Dark Fiber
 - Dense Wavelength Division Multiplexing (DWDM), 4-9
- data block corruption
 - prevention and detection parameters, 2-20
- data corruption, 1-4
- data corruptions
 - detecting, 2-20
- Data Guard
 - benefits, 2-4
 - comparing to Streams, 2-7
 - hub architecture, 4-18
- Data Guard hub, 4-18
- Data Recovery Advisor, 2-16
- Database Replay, 2-43
- Database Server Grid, 2-37
 - description, 2-38
- Database Storage Grid, 2-37
 - description, 2-38
- database upgrades
 - using Oracle Streams, 2-29
- databases
 - applying Oracle interim patches, 2-27
 - server grid, 2-37
- DB_BLOCK_CHECKING initialization
 - parameter, 2-20
- DB_BLOCK_CHECKSUM initialization
 - parameter, 2-20
- DB_LOST_WRITE_PROTECT initialization
 - parameter, 2-20

- DB_ULTRA_SAFE initialization parameter, 2-20
- DBA_FLASHBACK_TRANSACTION_STATE
 - view, 2-11
- DBMS_FLASHBACK.TRANSACTION_BACKOUT()
 - procedure, 2-11
- Dense Wavelength Division Multiplexing (DWDM or Dark Fiber), 4-9
- downtime
 - causes, 1-3
 - cost, 3-2
- DWDM
 - Dense Wavelength Division Multiplexing., 4-9

E

- endian format platforms
 - avoiding downtime during migration of different, 2-32
 - avoiding downtime during migration of same, 2-30
- extended distance clusters
 - architecture, 4-8
 - configuring a third site for a voting disk, 4-10
- external redundancy
 - ASM server-based mirroring, 4-10

F

- failovers
 - fast-start, 4-14
 - multiple standby database architecture, 4-15
 - single standby database architecture, 4-13
- failure groups
 - specifying in an extended cluster, 4-9
- failures
 - computer, 1-3
 - probability, 4-24
 - site, 1-5
 - storage, 1-3
- Fast Mirror Resync
 - ASM, 2-14
- fast-start failovers
 - single standby database failover, 4-14
- fault diagnosability infrastructure, 2-42
- flash recovery area
 - description, 2-17
- Flashback Data Archive, 2-13
- Flashback Database
 - description, 2-12
- Flashback Drop
 - description, 2-12
- flashback logs
 - used by Flashback features, 2-10
- Flashback Query
 - description, 2-10
- Flashback Restore Points
 - description, 2-12
- Flashback Table
 - description, 2-12
- Flashback technology

- block recovery using Flashback logs, 2-13
- Flashback Transaction
 - description, 2-11
- Flashback Transaction Query
 - description, 2-11
- Flashback Versions Query
 - description, 2-11
- frequency of outages, 4-24

G

- grid computing, 2-37
 - Database Server Grid, 2-37
 - Database Storage Grid, 2-37
- grids
 - server and storage, 2-37

H

- Hang Manager, 2-19, 2-43
- hang or slow down, 1-5
- HARD initiative, 2-19
- Hardware Assisted Resilient Data (HARD)
 - initiative, 2-19
- hardware RAID storage subsystem
 - deferring mirroring to, 4-10
- hardware upgrades
 - avoiding downtime during, 2-25
- high availability
 - applications, 4-30
 - architecture, 1-2
 - business impact analysis, 3-2
 - importance, 1-2
- high availability architectures
 - comparison, 4-26
- high availability solution
 - characteristics, 1-2
- hub-and-spoke configuration
 - Oracle Streams, 2-7
- human errors, 1-3

I

- intelligent infrastructure, 2-42

L

- logical standby databases, 2-5
 - transient, 4-15
- lost writes, 1-5

M

- Manageability Overhead (MO), 4-24
- Maximum Availability Architecture
 - benefits, 4-20
- memory
 - automatic management of, 2-24
- memory advisors, 2-43
- MEMORY_MAX_TARGET initialization
 - parameter, 2-23

- MEMORY_TARGET initialization parameter, 2-23
- migrating storage
 - avoiding downtime, 2-29
- mirroring
 - deferring to RAID storage subsystem, 4-10
- multiple standby database architecture
 - failovers, 4-15
- multiple standby database environment
 - using transient logical standby, 4-15
- multiple standby databases
 - Data Guard hub, 4-18

N

- nodes
 - virtual IP addresses, 4-5

O

- Observer
 - fast-start failover, 4-14
- one-off patches, 2-27
- online application maintenance and upgrades, 2-35
- online maintenance
 - application, 2-35
- online reorganization
 - description, 2-33
- opatch utility
 - patch upgrades for Oracle RAC, 2-27
- operating systems
 - requirements for Oracle Clusterware, 4-3
- Oracle Application Server
 - high availability architectures, 4-30
- Oracle Clusterware
 - advantages over third-party clusterware, 4-3
 - cold cluster failover, 4-3
 - cold failover cluster, 4-4
 - configured with Data Guard, 4-19
 - configuring a third site for a voting disk, 4-10
- Oracle Clusterware (CRS)
 - avoiding downtime when upgrading
 - upgrades
 - Oracle Clusterware, 2-28
- Oracle Data Guard
 - configured with Oracle Clusterware, 4-19
 - multiple standby database architecture, 4-15
 - single standby database architecture, 4-13
- Oracle Database
 - basic architecture, 4-2
 - with an Oracle RAC extended distance cluster, 4-8
 - with Oracle Clusterware (cold cluster failover), 4-3
- Oracle Database with Data Guard architecture
 - benefits, 4-12
- Oracle Database with Oracle RAC and Data Guard - MAA
 - benefits, 4-20
- Oracle Database with Oracle RAC architecture
 - benefits, 4-6

- Oracle Database with Streams architecture
 - benefits, 4-21
- Oracle Enterprise Manager Grid Control, 2-44
- Oracle interim (one-off) patches, 2-27
 - avoiding downtime during, 2-27
- Oracle Management Agents
 - Oracle Enterprise Management Grid Control, 2-44
- Oracle Management Repository
 - Oracle Enterprise Manager Grid Control, 2-44
- Oracle Management Service (OMS)
 - Oracle Enterprise Manager Grid Control, 2-44
- Oracle RAC
 - extended distance clusters, 4-8
 - Storage Area Network (SAN), 4-10
- Oracle Secure Backup (OSB)
 - overview, 2-15
- Oracle Steams
 - performing platform migrations, 2-29
- Oracle Streams
 - 1-N or hub-and-spoke configuration, 2-7
 - performing database upgrades, 2-29
- outages
 - frequency, 4-24

P

- physical standby databases, 2-4
 - real-time query, 2-39
- planned activities
 - probability of failure during, 4-24
- platform migrations
 - using Oracle Streams, 2-29
- preferred read failure groups
 - specifying ASM, 4-9
- probability
 - of different failures during unplanned and planned activities, 4-24
- program global area (PGA)
 - automatic management, 2-24

R

- Real Application Clusters
 - benefits, 2-3
- real-time query standby databases, 2-5, 2-39
- Recovery Manager (RMAN)
 - description, 2-14
- recovery point objective (RPO)
 - description, 3-3, 4-24
- recovery time objective (RTO)
 - description, 3-2, 4-24
- restore points
 - Flashback, 2-12
- Return On Investment (ROI), 4-24
- rollback
 - transactions, 2-11
- rolling upgrades
 - using transient logical standby, 4-15

S

- secure communications
 - between tiers in grid control environments, 2-44
- segment advisor, 2-43
- server generated alerts, 2-42
- server grid, 2-37
- server-based mirroring
 - ASM, 4-9
- servers
 - Oracle Clusterware requirements, 4-3
- service level agreements, 3-4
- single standby database architecture
 - failovers, 4-13
- site failure, 1-5
- SLAs, 3-4
- Snapshot Standby Database, 2-5
- snapshot standby database
 - in a multiple standby database environment, 4-16
- SQL Access Advisor, 2-43
- SQL Performance Analyzer, 2-43
- SQL Tuning Advisor, 2-43
- SSL
 - enabling for secure communications, 2-44
- standby database hub
 - example configurations, 4-18
- standby databases
 - logical standby, 4-15
 - snapshot standby in a multistandby database environment, 4-16
- standby reader farms, 2-40
- storage
 - migration, 2-29
 - mirroring to RAID, 4-10
- Storage Area Network (SAN)
 - extended distance clusters, 4-10
- storage failure, 1-3
- storage grid, 2-37
- Streams
 - comparing to Data Guard, 2-7
 - description, 2-6
- system global area (SGA)
 - automatic management, 2-24
- system upgrades
 - avoiding downtime during, 2-25

T

- tape backups
 - with OSB, 2-15
- thin client watchdog
 - observer for fast-start failover, 4-14
- third site
 - for a voting disk, 4-10
- Total Cost of Ownership (TCO), 4-24
- transactions
 - backing out with Flashback Transaction, 2-11
- transportable database
 - for unplanned downtime, 2-9
- transportable tablespaces
 - for unplanned downtime, 2-10

- transportable technologies
 - for unplanned downtime, 2-9

U

- Undo Advisor, 2-43
- undo data
 - used by flashback features, 2-10
- unplanned activities
 - probability of failure during, 4-24
- unplanned downtime
 - transportable tablespaces, 2-10
 - transportable technologies, 2-9
- upgrades
 - application, 2-35
 - database, 2-29
 - with logical standby databases, 4-15

V

- V\$DATABASE_BLOCK_CORRUPTION view, 2-20
- virtual IP address
 - defining for applications, 4-5
 - Oracle Clusterware, 4-5
- volume manager
 - ASM, 4-10
- voting disk
 - configuring a third site, 4-10

W

- web scalability
 - using standby reader farms, 2-40