

Oracle® Database

High Availability Overview



18c
E84003-01
December 2019

The Oracle logo, consisting of a solid red square with the word "ORACLE" in white, uppercase, sans-serif font centered within it.

ORACLE®

Oracle Database High Availability Overview, 18c

E84003-01

Copyright © 2005, 2019, Oracle and/or its affiliates. All rights reserved.

Primary Author: Virginia Beecher

Contributing Authors: Lance Ashdown, Tulika Das, Viv Schupmann, Janet Stern, Lawrence To

Contributors: Ahmed Abbas, Andrew Babb, Hermann Baer, Tammy Bednar, Peter Belknap, Janet Blowney, Larry Carpenter, Immanuel Chan, Dib Chatterjee, Tim Chien, Donna Cooksey, Saurav Das, Mark Dilman, Ray Dutcher, Richard Exley, Craig Foch, Stephan Haisley, Glen Hawkins, Ameet Kini, Frank Kobylanski, Rene Kundersma, Bryn Llewellyn, Barb Lundhild, Rahim Mau, Patricia McElroy, Joe Meeks, Markus Michalewicz, Valerie Moore, Dan Norris, Michael Nowak, Darryl Presley, Hector Pujol, Ashish Ray, Mark Scardina, Jia Shi, Michael T. Smith, Vinay Srihari, Andrew Steinorth, Hubert Sun, Lawrence To, Douglas Utzig, James Viscusi, Tak Wang, Shari Yamaguchi

This software and related documentation are provided under a license agreement containing restrictions on use and disclosure and are protected by intellectual property laws. Except as expressly permitted in your license agreement or allowed by law, you may not use, copy, reproduce, translate, broadcast, modify, license, transmit, distribute, exhibit, perform, publish, or display any part, in any form, or by any means. Reverse engineering, disassembly, or decompilation of this software, unless required by law for interoperability, is prohibited.

The information contained herein is subject to change without notice and is not warranted to be error-free. If you find any errors, please report them to us in writing.

If this is software or related documentation that is delivered to the U.S. Government or anyone licensing it on behalf of the U.S. Government, then the following notice is applicable:

U.S. GOVERNMENT END USERS: Oracle programs, including any operating system, integrated software, any programs installed on the hardware, and/or documentation, delivered to U.S. Government end users are "commercial computer software" pursuant to the applicable Federal Acquisition Regulation and agency-specific supplemental regulations. As such, use, duplication, disclosure, modification, and adaptation of the programs, including any operating system, integrated software, any programs installed on the hardware, and/or documentation, shall be subject to license terms and license restrictions applicable to the programs. No other rights are granted to the U.S. Government.

This software or hardware is developed for general use in a variety of information management applications. It is not developed or intended for use in any inherently dangerous applications, including applications that may create a risk of personal injury. If you use this software or hardware in dangerous applications, then you shall be responsible to take all appropriate fail-safe, backup, redundancy, and other measures to ensure its safe use. Oracle Corporation and its affiliates disclaim any liability for any damages caused by use of this software or hardware in dangerous applications.

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

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

This software or hardware and documentation may provide access to or information about content, products, and services from third parties. Oracle Corporation and its affiliates are not responsible for and expressly disclaim all warranties of any kind with respect to third-party content, products, and services unless otherwise set forth in an applicable agreement between you and Oracle. Oracle Corporation and its affiliates will not be responsible for any loss, costs, or damages incurred due to your access to or use of third-party content, products, or services, except as set forth in an applicable agreement between you and Oracle.

Contents

Preface

Audience	vii
Documentation Accessibility	vii
Related Documents	viii
Conventions	viii

1 Overview of High Availability

What Is High Availability?	1-1
Importance of Availability	1-2
Cost of Downtime	1-3
Causes of Downtime	1-3
Roadmap to Implementing the Maximum Availability Architecture	1-8

2 High Availability and Data Protection – Getting From Requirements to Architecture

High Availability Requirements	2-1
A Methodology for Documenting High Availability Requirements	2-2
Business Impact Analysis	2-3
Cost of Downtime	2-4
Recovery Time Objective	2-4
Recovery Point Objective	2-5
Manageability Goal	2-5
Total Cost of Ownership and Return on Investment	2-5
Mapping Requirements to Architectures	2-6
Oracle MAA Reference Architectures	2-7
Bronze Reference Architecture	2-8
Silver Reference Architecture	2-8
Gold Reference Architecture	2-8
Platinum Reference Architecture	2-9
High Availability and Data Protection Attributes by Tier	2-9

3 Features for Maximizing Availability

Oracle Data Guard	3-2
Oracle Active Data Guard	3-6
Oracle Data Guard Advantages Over Traditional Solutions	3-8
Data Guard and Planned Maintenance	3-9
Data Guard Redo Apply and Standby-First Patching	3-10
Data Guard Transient Logical Rolling Upgrades	3-10
Rolling Upgrade Using Oracle Active Data Guard	3-11
Oracle GoldenGate	3-12
Best Practice: Oracle Active Data Guard and Oracle GoldenGate	3-14
When to Use Oracle Active Data Guard	3-15
When to Use Oracle GoldenGate	3-15
When to Use Oracle Active Data Guard and Oracle GoldenGate Together	3-16
Recovery Manager	3-17
Oracle Real Application Clusters and Oracle Clusterware	3-19
Benefits of Using Oracle Clusterware	3-20
Benefits of Using Oracle Real Application Clusters and Oracle Clusterware	3-21
Oracle RAC Advantages Over Traditional Cold Cluster Solutions	3-21
Oracle RAC One Node	3-24
Oracle Automatic Storage Management	3-24
Fast Recovery Area	3-26
Corruption Prevention, Detection, and Repair	3-27
Data Recovery Advisor	3-30
Oracle Flashback Technology	3-31
Oracle Flashback Query	3-33
Oracle Flashback Version Query	3-33
Oracle Flashback Transaction	3-34
Oracle Flashback Transaction Query	3-34
Oracle Flashback Table	3-34
Oracle Flashback Drop	3-35
Restore Points	3-35
Oracle Flashback Database	3-35
Flashback Pluggable Database	3-36
Block Media Recovery Using Flashback Logs or Physical Standby Database	3-36
Flashback Data Archive	3-37
Oracle Data Pump and Data Transport	3-37
Oracle Replication Technologies for Non-Database Files	3-38
Oracle ASM Cluster File System	3-39
Oracle Database File System	3-40
Oracle Solaris ZFS Storage Appliance Replication	3-40

Oracle Multitenant	3-41
Oracle Sharding	3-44
Oracle Restart	3-44
Oracle Site Guard	3-45
Online Reorganization and Redefinition	3-45
Zero Data Loss Recovery Appliance	3-46
Fleet Patching and Provisioning	3-46
Enabling Continuous Service for Applications	3-47
Continuous Application Service	3-47
Edition-Based Redefinition	3-48
Learn About Other High Availability Features	3-49

4 Oracle Database High Availability Solutions for Unplanned Downtime

Outage Types and Oracle High Availability Solutions for Unplanned Downtime	4-1
Managing Unplanned Outages for MAA Reference Architectures and Multitenant Architectures	4-6

5 Oracle Database High Availability Solutions for Planned Downtime

Oracle High Availability Solutions for Planned Maintenance	5-1
High Availability Solutions for Migration	5-2

6 Operational Prerequisites to Maximizing Availability

Understand Availability and Performance SLAs	6-1
Implement and Validate a High Availability Architecture That Meets Your SLAs	6-2
Establish Test Practices and Environment	6-2
Configuring the Test System and QA Environments	6-2
Performing Preproduction Validation Steps	6-4
Set Up and Use Security Best Practices	6-5
Establish Change Control Procedures	6-6
Apply Recommended Patches and Software Periodically	6-6
Execute Disaster Recovery Validation	6-7
Establish Escalation Management Procedures	6-7
Configure Monitoring and Service Request Infrastructure for High Availability	6-8
Run Database Health Checks Periodically	6-8
Configure Oracle Enterprise Manager Monitoring Infrastructure for High Availability	6-9
Configure Automatic Service Request Infrastructure	6-10

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](#)
- [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

For information about Oracle's commitment to accessibility, visit the Oracle Accessibility Program website at <http://www.oracle.com/pls/topic/lookup?ctx=acc&id=docacc>.

Access to Oracle Support

Oracle customers that have purchased support have access to electronic support through My Oracle Support. For information, visit <http://www.oracle.com/pls/topic/>

[lookup?ctx=acc&id=info](http://www.oracle.com/pls/topic/lookup?ctx=acc&id=info) or visit <http://www.oracle.com/pls/topic/lookup?ctx=acc&id=trs> if you are hearing impaired.

Related Documents

Knowledge of Oracle Database, Oracle RAC, and Data Guard concepts and terminology is required to understand the configuration and implementation details described in this book. For more information, see the Oracle Database documentation set. These books may be of particular interest:

- *Oracle Database Administrator's Guide*
- *Oracle Clusterware Administration and Deployment Guide*
- *Oracle Real Application Clusters Administration and Deployment Guide*
- *Oracle Automatic Storage Management Administrator's Guide*
- *Oracle Data Guard Concepts and Administration*
- *Oracle Database Backup and Recovery User'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 Database. See *Oracle Database Sample Schemas* for information about using these schemas.

Also, you can download the Oracle MAA best practice white papers at <http://www.oracle.com/goto/maa>.

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.

1

Overview of High Availability

See the following topics to learn what high availability and why it is important. Then follow the roadmap to implementing a Maximum Availability Architecture.

- [What Is High Availability?](#)
Availability is the degree to which an application and database service is available.
- [Importance of Availability](#)
The importance of high availability varies among applications. Databases and the internet have enabled worldwide collaboration and information sharing by extending the reach of database applications throughout organizations and communities.
- [Cost of Downtime](#)
The need to deliver increasing levels of availability continues to accelerate as enterprises reengineer their solutions to gain competitive advantage. Most often, these new solutions rely on immediate access to critical business data.
- [Causes of Downtime](#)
- [Roadmap to Implementing the Maximum Availability Architecture](#)
Oracle high availability solutions and sound operational practices are the key to successful implementation of an IT infrastructure. However, technology alone is not enough.

What Is High Availability?

Availability is the degree to which an application and database service is available.

Availability is measured by the perception of an application's user. Users experience frustration when their data is unavailable or the computing system is not performing as expected, 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 disruption as the failure of critical components in the architecture. If a user cannot access the application or database service, it is said to be unavailable. Generally, the term downtime is used to refer to periods when a system is unavailable.

Users who want their systems to be always ready to serve them need high availability. A system that is highly available is designed to provide uninterrupted computing services during essential time periods, during most hours of the day, and most days of the week throughout the year; this measurement is often shown as 24x365. Such systems may also need a high availability solution for planned maintenance operations such as upgrading a system's hardware or software.

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 applications—is just as critical to implementing a highly available solution. A related characteristic is

resilience. For example, low-cost commodity hardware, combined with software such as Oracle Real Application Clusters (Oracle RAC), can be used to implement a very reliable system. The resilience of an Oracle RAC database allows processing to continue even though individual servers may fail. For example, the Oracle RAC database allows processing to continue even though individual servers may fail.

- **Recoverability:** Even though there may be many ways to recover from a failure, it is important to determine what types of failures may occur in your high availability environment and how to recover from those failures quickly in order to meet 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 essential to recover from the unexpected failure. Although 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 database administrator of a problem.
- **Continuous operation:** Providing continuous access to your data is essential when very little or no downtime is acceptable to perform maintenance activities. Activities, such as moving a table to another location in the database or even adding CPUs to your hardware, should be transparent to the 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 preventive measures
- Provide active monitoring and fast detection of failures
- Provide fast recoverability
- Automate detection and recovery operations
- Protect the data to minimize or prevent data loss and corruptions
- Implement the operational best practices to manage your environment
- Achieve the goals set in SLAs (for example, recovery time objectives (RTOs) and recovery point objectives (RPOs)) for the lowest possible total cost of ownership

Importance of Availability

The importance of high availability varies among applications. 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 now demand service-level agreements from their information technology (IT) departments and solution providers, reflecting the

increasing dependence on these solutions. 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. The business might lose revenue, incur penalties, and receive bad publicity that has a lasting effect on customers and on the company's stock price.

It is important to examine the factors that determine how your data is protected and maximize availability to your users.

Cost of Downtime

The need to deliver increasing levels of availability continues to accelerate as enterprises reengineer 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 on their solutions to provide service.

Other factors to consider in the cost of downtime are:

- The maximum tolerable length of a single unplanned outage
If the event lasts less than 30 seconds, then it may cause very little impact and may be barely perceptible to users. As the length of the outage grows, the effect may grow exponentially and negatively affect the business.
- The maximum frequency of allowable incidents
Frequent outages, even if short in duration, may similarly disrupt business operations.

When designing a solution, it is important to recognize the true cost of downtime to understand how the business can benefit from availability improvements.

Oracle provides a range of high availability solutions to 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 data are reliably accessible everywhere, at any time.

Causes of Downtime

One of the challenges in designing a high availability solution is examining and addressing all of 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 as unplanned downtime, especially in global enterprises that support users in multiple time zones.

The following table describes unplanned outage types and provides examples of each type.

Table 1-1 Causes of Unplanned Downtime

Type	Description	Examples
Site failure	<p>A site failure may affect all processing at a data center, or a subset of applications supported by a data center.</p> <p>The definition of site varies given the contexts of on-premises and cloud.</p> <ul style="list-style-type: none"> • Site failure - entire regional failure • Data center - entire data center location • Availability domain - isolated data center within a region with possibly many other availability domains • Fault domain - isolated set of system resources within an Availability Domain or data center <p>Typically, each site, data center, availability domain, and fault domain has its own set of isolated hardware, DB compute, network, storage, and power.</p>	<ul style="list-style-type: none"> • Extended sitewide power failure • Sitewide network failure • Natural disaster makes a data center inoperable • Terrorist or malicious attack on operations or the site
Clusterwide failure	<p>The whole cluster hosting an Oracle RAC database is unavailable or fails. This includes:</p> <ul style="list-style-type: none"> • Failures of nodes in the cluster • Failure of any other components that result in the cluster being unavailable and the Oracle database and instances on the site being unavailable 	<ul style="list-style-type: none"> • The last surviving node on the Oracle RAC cluster fails and the node or database cannot be restarted • Both redundant cluster interconnections fail or Clusterware failure • Database corruption so severe that continuity is not possible on the current database server • Clusterware and hardware-software defects preventing availability or stability.
Computer failure	<p>A computer failure outage occurs when the system running the database becomes unavailable because it has failed or is no longer available. When the database uses Oracle RAC then a computer failure represents a subset of the system (while retaining full access to the data).</p>	<ul style="list-style-type: none"> • Database system hardware failure • Operating system failure • Oracle instance failure
Network failure	<p>A network failure outage occurs when a network device stops or reduces network traffic and communication from your application to database, database to storage, or any system to system that is critical to your application service processing.</p>	<ul style="list-style-type: none"> • Network switch failure • Network interface failure • Network cable failures
Storage failure	<p>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 available.</p>	<ul style="list-style-type: none"> • Disk or flash drive failure • Disk controller failure • Storage array failure

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

Type	Description	Examples
Data corruption	<p>A corrupt block is a block that was changed so that it differs from what Oracle Database expects to find. Block corruptions can be categorized as physical or logical:</p> <ul style="list-style-type: none"> In a physical block corruption, which is also called a media corruption, the database does not recognize the block at all; the checksum is invalid or the block contains all zeros. An example of a more sophisticated block corruption is when the block header and footer do not match. In a logical block corruption, the contents of the block are physically sound and pass the physical block checks; however, the block can be logically inconsistent. Examples of logical block corruption include incorrect block type, incorrect data or redo block sequence number, corruption of a row piece or index entry, or data dictionary corruptions. <p>Block corruptions can also be divided into interblock corruption and intrablock corruption:</p> <ul style="list-style-type: none"> In an intrablock corruption, the corruption occurs in the block itself and can be either a physical or a logical block corruption. In an interblock corruption, the corruption occurs between blocks and can only be a logical block corruption. <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 application or database (down to a single database block) to a large portion of the application or database (making it essentially unusable).</p>	<ul style="list-style-type: none"> Operating system or storage device driver failure Faulty host bus adapter Disk controller failure Volume manager error causing a bad disk read or write Software or hardware defects
Human error	<p>A human error outage occurs when unintentional or other actions are committed that cause data in the database to become incorrect 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.</p>	<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
Lost or stray writes	<p>A lost or stray write is another form of data corruption, but it is much more difficult to detect and repair quickly. A data block stray or lost write occurs when:</p> <ul style="list-style-type: none"> • For a lost write, an I/O subsystem acknowledges the completion of the block write even though the write I/O 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. • For a stray write, the write I/O completed but it was written somewhere else, and a subsequent read operation returns the stale value. • For an Oracle RAC system, a read I/O from one cluster node returns stale data after a write I/O is completed from another node (lost write). For example, this occurs if a network file system (NFS) is mounted in Oracle RAC without disabling attribute caching (for example, without using the <code>noac</code> option). In this case, the write I/O from one node is not immediately visible to another node because it is cached. <p>Block corruptions caused by stray writes or lost writes can cause havoc to your database availability. The data block may be physically or logically correct but subsequent disk reads will show blocks that are stale or with an incorrect Oracle Database block address.</p>	<ul style="list-style-type: none"> • Operating system or storage device driver failure • Faulty host bus adapter • Disk controller failure • Volume manager error • Other application software • Lack of network file systems (NFS) write visibility across a cluster • Software or hardware defects

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

Type	Description	Examples
Delay or slowdown	A delay or slowdown occurs when the database or the application cannot process transactions because of a resource or lock contention. A perceived delay can be caused by lack of system resources.	<ul style="list-style-type: none"> • Database or application deadlocks • Runaway processes that consume system resources • Logon storms or system faults • Combination of application peaks with lack of system or database resources. This can occur with one application or many applications in a consolidated database environment without proper resource management. • Archived redo log destination or fast recovery area destination becomes full • Oversubscribed or heavily consolidated database system

The following table describes planned outage types and provides examples of each type.

Table 1-2 Causes of Planned Downtime

Type	Description	Examples
Software changes	<ul style="list-style-type: none"> • Planned periodic software changes to apply minor fixes for stability and security • Planned annual or bi-annual major upgrades to adopt new features and capabilities 	<ul style="list-style-type: none"> • Software updates, including security updates to operating system, clusterware, or database • Major upgrade of operating system, clusterware, or database • Updating or upgrading application software
System and database changes	<ul style="list-style-type: none"> • Planned system changes to replace defected hardware • Planned system changes to expand or reduce system resources • Planned database changes to adopt parameter changes • Planned change to migrate to new hardware or architecture 	<ul style="list-style-type: none"> • Adding or removing processors or memory to a server • Adding or removing nodes to or from a cluster • Adding or removing disks drives or storage arrays • Replacing any Field Replaceable Unit (FRU) • Changing configuration parameters • System platform migration • Migrating to cluster architecture • Migrating to new storage

Table 1-2 (Cont.) Causes of Planned Downtime

Type	Description	Examples
Data changes	Planned data changes to the logical structure or physical organization of Oracle Database objects. The primary objective of these changes is to improve performance or manageability.	<ul style="list-style-type: none"> • Table definition changes • Adding table partitioning • Creating and rebuilding indexes
Application changes	Planned application changes can include data changes and schema and programmatic changes. The primary objective of these changes is to improve performance, manageability, and functionality.	Application upgrades

Oracle offers high availability solutions to help avoid both unplanned and planned downtime, and recover from failures. [Oracle Database High Availability Solutions for Unplanned Downtime](#) and [Oracle Database High Availability Solutions for Planned Downtime](#) discuss each of these high availability solutions in detail.

Roadmap to Implementing the Maximum Availability Architecture

Oracle high availability solutions and sound operational practices are the key to successful implementation of an IT infrastructure. However, technology alone is not enough.

Choosing and implementing an architecture that best fits your availability requirements can be a daunting task. Oracle Maximum Availability Architecture (MAA) simplifies the process of choosing and implementing a high availability architecture to fit your business requirements with the following considerations:

- Encompasses redundancy across all components
- Provides protection and tolerance from computer failures, storage failures, human errors, data corruption, lost writes, system delays or slowdowns, and site disasters
- Recovers from outages as quickly and transparently as possible
- Provides solutions to eliminate or reduce planned downtime
- Provides consistent high performance and robust security
- Provides Oracle Engineered System and cloud options to simplify deployment and management and achieve the highest scalability, performance, and availability
- Achieves SLAs at the lowest possible total cost of ownership
- Applies to On-Premise, Oracle Public Cloud, and hybrid architectures consisting of parts on-premise and part in the cloud
- Provides special consideration to Container or Oracle Multitenant, Oracle Database In-Memory, and Oracle Sharding architectures

To build, implement, and maintain this type of architecture, you need to:

1. Analyze your specific high availability requirements, including both the technical and operational aspects of your IT systems and business processes, as described in [High Availability and Data Protection – Getting From Requirements to Architecture](#).
2. Familiarize yourself with Oracle high availability features, as described in [Features for Maximizing Availability](#).
3. Understand the availability impact for each MAA reference architecture, or various high availability features, on businesses and applications, as described in [Oracle Database High Availability Solutions for Unplanned Downtime](#), and [Oracle Database High Availability Solutions for Planned Downtime](#).
4. Use operational best practices to provide a successful MAA implementation, as described in [Operational Prerequisites to Maximizing Availability](#).
5. Choose a high availability architecture, as described in [Mapping Requirements to Architectures](#).
6. Implement a high availability architecture using Oracle MAA resources, which provide technical details about the various Oracle MAA high availability technologies, along with best practice recommendations for configuring and using such technologies, such as Oracle MAA best practices white papers, customer papers with proof of concepts, customer case studies, recorded web casts, demonstrations, and presentations.

Oracle MAA resources are available at <http://www.oracle.com/goto/maa>.

2

High Availability and Data Protection – Getting From Requirements to Architecture

See the following topics to learn how Oracle Maximum Availability Architecture provides a framework to effectively evaluate the high availability requirements of an enterprise.

- [High Availability Requirements](#)
Any effort to design and implement a high availability strategy for Oracle Database begins by performing a thorough business impact analysis to identify the consequences to the enterprise of downtime and data loss, whether caused by unplanned or planned outages.
- [A Methodology for Documenting High Availability Requirements](#)
- [Mapping Requirements to Architectures](#)
The business impact analysis will help you document what is already known. The outcome of the business impact analysis provides the insight you need to group databases having similar RTO and RPO objectives together.

High Availability Requirements

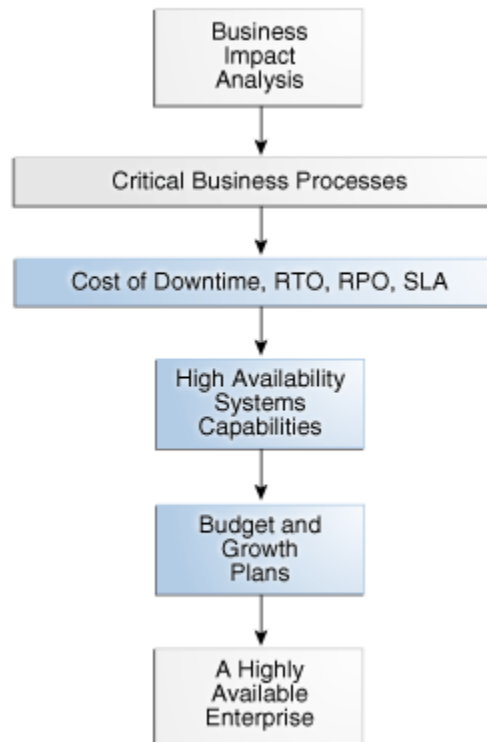
Any effort to design and implement a high availability strategy for Oracle Database begins by performing a thorough business impact analysis to identify the consequences to the enterprise of downtime and data loss, whether caused by unplanned or planned outages.

The term "business impact" is intended to be agnostic of whether the enterprise is a commercial venture, government agency, or not-for-profit institution. In all cases, data loss and downtime can seriously impact the ability of any enterprise to perform its functions. Implementing high availability may involve critical tasks such as:

- Retiring legacy systems
- Investing in more capable and robust systems and facilities
- Redesigning the overall IT architecture and operations to adapt to this high availability model
- Modifying existing applications to take full advantage of high availability infrastructures
- Redesigning business processes
- Hiring and training personnel
- Moving parts or an entire application or database into the Oracle Public Cloud
- Balancing the right level of consolidation, flexibility, and isolation
- Understanding the capabilities and limitations of your existing system and network infrastructure

By combining your business analysis with an understanding of the level of investment required to implement different high availability solutions, you can develop a high availability architecture that achieves both business and technical objectives.

Figure 2-1 Planning and Implementing a Highly Available Enterprise



A Methodology for Documenting 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 and Return on Investment](#)

- [Business Impact Analysis](#)

The business impact analysis categorizes the business processes based on the severity of the impact of IT-related outages.

- [Cost of Downtime](#)

A complete business impact analysis provides the insight needed to quantify the cost of unplanned and planned downtime.

- **Recovery Time Objective**
The business impact analysis determines your tolerance to downtime, also known as the recovery time objective (RTO).
- **Recovery Point Objective**
The business impact analysis also determines your tolerance to data loss, also known as a recovery point objective (RPO).
- **Manageability Goal**
A manageability goal is more subjective than either the RPO or the RTO. You must make an objective evaluation of the skill sets, management resources, and tools available in an organization, and the degree to which the organization can successfully manage all elements of a high availability architecture.
- **Total Cost of Ownership and Return on Investment**
Understanding the total cost of ownership (TCO) and objectives for return on investment (ROI) are essential to selecting a high availability architecture that also achieves the business goals of your organization.

Business Impact Analysis

The business impact analysis categorizes the business processes based on the severity of the impact of IT-related outages.

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
- Outlines the effects of these outages
- Considers essential business functions, people and system resources, government regulations, and internal and external business dependencies
- Is based on objective and subjective data gathered from interviews with knowledgeable and experienced personnel
- Reviews business practice histories, financial reports, IT systems logs, and so on

For example, consider a semiconductor manufacturer with chip fabrication plants located worldwide. Semiconductor manufacturing is an intensely competitive business requiring a huge financial investment that is amortized over high production volumes. The human resource applications used by plant administration are unlikely to be considered as mission-critical as the applications that control the manufacturing process in the plant. Failure of the applications that support manufacturing affects production levels and have a direct impact on the financial results of the company.

As another example, 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.

Similarly, an e-commerce company is highly dependent on customer traffic to its website to generate revenue. Any disruption in service and loss of availability can dampen customer experience and drive away customers to the competition. Thus, the company needs to ensure that the existing infrastructure can scale and handle spikes

in customer traffic. Sometimes, this is not possible using on-premise hardware and by moving the cloud the company can ensure their systems always remain operational.

Cost of Downtime

A complete business impact analysis provides the insight needed to quantify the cost of unplanned and planned downtime.

Understanding this cost is essential because it helps prioritize your high availability investment and directly influences the high availability technologies that you choose to minimize the downtime risk.

Various reports have been published, documenting the costs of downtime in different industries. 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. The Internet and Cloud 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 negatively affect customer relationships, competitive advantages, legal obligations, industry reputation, and shareholder confidence.

Recovery Time Objective

The business impact analysis determines your tolerance to downtime, also known as the recovery time objective (RTO).

An 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, customer dissatisfaction, reputation, and so on). RTO indicates the downtime tolerance of a business process or an organization in general.

RTO requirements are driven by the mission-critical nature of the business. Therefore, for a system running a stock exchange, the RTO is zero or near to zero.

An organization is likely to have varying RTO requirements across its various business processes. A high volume e-commerce website, 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 of zero or close to zero. However, the RTO of the systems that support back-end operations, such as shipping and billing, can be higher. If these back-end systems are down, then the business may resort to manual operations temporarily without a significant visible impact.

Some organizations have varying RTOs based on the probability of failures. One simple class separation is local failures (such as single database compute, disk/flash, network failure) as opposed to disasters (such as a complete cluster, database, data corruptions, or a site failure). Typically, business-critical customers have an RTO of less than 1 minute for local failures, and may have a higher RTO of less than 1 hour for disasters. For mission-critical applications the RTOs may indeed be the same for all unplanned outages.

Recovery Point Objective

The business impact analysis also determines your tolerance to data loss, also known as a recovery point objective (RPO).

The RPO is the maximum amount of data that an IT-based business process can lose without harm to the organization. RPO measures 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, zero, seconds, hours, or days of data loss.

A stock exchange where millions of dollars worth of transactions occur every minute cannot afford to lose any data. Therefore, its RPO must be zero. The web-based sales system in the e-commerce example does not require an RPO of zero, although a low RPO is essential for customer satisfaction. However, its back-end merchandising and inventory update system can have a higher RPO because lost data can be reentered.

An RPO of zero can be challenging for disasters, but it can be accomplished with various Oracle technologies protecting your database, especially Zero Data Loss Recovery Appliance.

Manageability Goal

A manageability goal is more subjective than either the RPO or the RTO. You must make an objective evaluation of the skill sets, management resources, and tools available in an organization, and the degree to which the organization can successfully manage all elements of a high availability architecture.

Just as RPO and RTO measure an organization's tolerance for downtime and data loss, your manageability goal measures the organization's tolerance for complexity in the IT environment. When 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 latter could attain more aggressive RTO and RPO objectives. Understanding manageability goals helps organizations differentiate between what is possible and what is practical to implement.

Moving Oracle databases to Oracle Cloud can reduce manageability cost and complexity significantly, because Oracle Cloud lets you to choose between various Maximum Availability Architecture architectures with built-in configuration and life cycle operations. With Autonomous Database Cloud, database life cycle operations, such as backup and restore, software updates, and key repair operations are automatic.

Total Cost of Ownership and Return on Investment

Understanding the total cost of ownership (TCO) and objectives for return on investment (ROI) are essential to selecting a 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 your chosen high availability solution. Likewise, the ROI calculation captures all of the financial benefits that accrue for a given high availability architecture.

For example, consider a high availability architecture in which IT systems and storage at a remote standby site remain idle, with no other business use that can be served by the standby systems. The only return on investment for the standby site is the costs

related to 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 off-loading the overhead of user queries or distributing read-write workload from the primary system). The return on investment of such an architecture includes both the cost of downtime avoided and the financial benefits that accrue to its productive use, while also providing high availability and data protection.

Enterprises can also reduce TCO for growing infrastructure needs by moving workloads to the cloud rather than making an upfront capital investment in building a new data center. The major economic appeal is to convert capital expenditures into operational expenditures, and generate a higher ROI.

Mapping Requirements to Architectures

The business impact analysis will help you document what is already known. The outcome of the business impact analysis provides the insight you need to group databases having similar RTO and RPO objectives together.

Different applications, and the databases that support them, represent varying degrees of importance to the enterprise. A high level of investment in high availability infrastructure may not make sense for an application that if down, would not have an immediate impact on the enterprise. So where do you start?

Groups of databases by similar RTO and RPO can be mapped to a controlled set of high availability reference architectures that most closely address the required service levels. Note that in the case where there are dependencies between databases, they are grouped with the database having the most stringent high availability requirement.

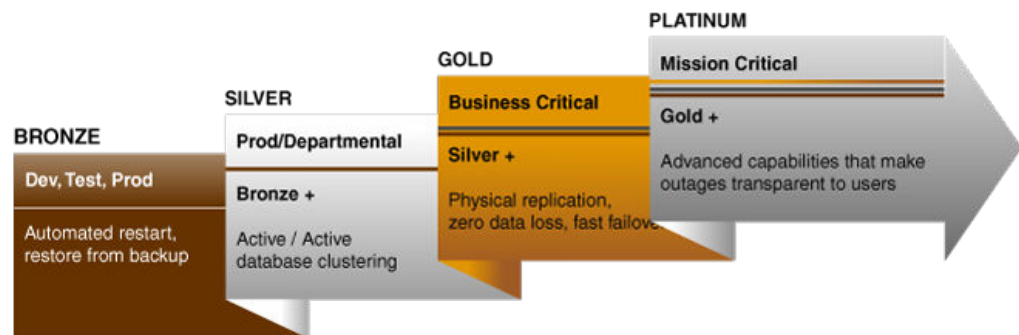
- [Oracle MAA Reference Architectures](#)
Oracle MAA best practices define high availability reference architectures that address the complete range of availability and data protection required by enterprises of all sizes and lines of business.
- [Bronze Reference Architecture](#)
The Bronze tier is appropriate for databases where simple restart of a failed component (e.g. listener, database instance, or database) or restore from backup is "HA and DR enough."
- [Silver Reference Architecture](#)
The Silver tier provides an additional level of high availability for databases that require minimal or zero downtime in the event of database instance or server failure, as well as most common planned maintenance events, such as hardware and software updates.
- [Gold Reference Architecture](#)
The Gold tier raises the stakes substantially for business-critical applications that cannot tolerate high RTO and RPO for any disasters such as database, cluster, corruptions, or site failures. Additionally, major database upgrades or site migrations can be done in seconds.
- [Platinum Reference Architecture](#)
The Platinum tier introduces several new Oracle Database capabilities, including Oracle GoldenGate for zero-downtime upgrades and migrations.
- [High Availability and Data Protection Attributes by Tier](#)
Each MAA reference architecture delivers known and tested levels of downtime and data protection.

Oracle MAA Reference Architectures

Oracle MAA best practices define high availability reference architectures that address the complete range of availability and data protection required by enterprises of all sizes and lines of business.

The Platinum, Gold, Silver, and Bronze MAA reference architectures, or tiers, are applicable to on-premises, private and public cloud configurations, and hybrid cloud. They deliver the service levels described in the following figure.

Figure 2-2 Oracle MAA Reference Architectures



Each tier uses a different MAA reference architecture to deploy the optimal set of Oracle high availability capabilities that reliably achieve a given service level at the lowest cost and complexity. The tiers explicitly address all types of unplanned outages, including data corruption, component failure, and system and site outages, as well as planned outages due to maintenance, migrations, or other purposes.

Container databases (CDBs) using Oracle Multitenant can exist in any tier, Bronze through Platinum, providing higher consolidation density and higher TCO. Typically, the consolidation density is higher with Bronze and Silver tiers, and there is less or zero consolidation when deploying a Platinum tier.

Oracle Database In-Memory can also be leveraged in any of the MAA tiers. Because the In-Memory column store is seamlessly integrated into Oracle Database, all of the high availability benefits that come from the MAA tiers are inherited when implementing Oracle Database In-Memory.

Oracle Engineered Systems can also exist in any of the tiers. Integrating Zero Data Loss Recovery Appliance (Recovery Appliance) as the Oracle Database backup and recovery solution for your entire data center reduces RPO and RTO when restoring from backups. Leveraging Oracle Exadata Database Machine as your database platform in the MAA reference architectures provides the best database platform solution with the lowest RTO and brownout, along with additional Exadata MAA quality of service.



See Also:

[High Availability Reference Architectures](#)

[Oracle Exadata Database Machine: Maximum Availability Architecture and MAA Best Practices for Oracle Exadata Database Machine](#)

<http://www.oracle.com/goto/maa> for MAA white paper "Oracle Database In-Memory High Availability Best Practices"

Bronze Reference Architecture

The Bronze tier is appropriate for databases where simple restart of a failed component (e.g. listener, database instance, or database) or restore from backup is "HA and DR enough."

The Bronze reference architecture is based on a single instance Oracle Database using MAA best practices that implement the many capabilities for data protection and high availability included with every Oracle Enterprise Edition license. Oracle-optimized backups using Oracle Recovery Manager (RMAN) provide data protection, and are used to restore availability should an outage prevent the database from restarting. The Bronze architecture then uses a redundant system infrastructure enhanced by Oracle's technologies, such as Oracle Restart, Recovery Manager (RMAN), Zero Data Loss Recovery Appliance, Flashback technologies, Online Redefinition, Online Patching, Automatic Storage Management (ASM), Oracle Multitenant, and more.

Silver Reference Architecture

The Silver tier provides an additional level of high availability for databases that require minimal or zero downtime in the event of database instance or server failure, as well as most common planned maintenance events, such as hardware and software updates.

The Silver reference architecture adds a rich set of enterprise capabilities and benefits, including clustering technology using either Oracle RAC or Oracle RAC One Node. Also, Application Continuity provides a reliable replay of in-flight transactions, which masks outages from users and simplifies application failover.

Gold Reference Architecture

The Gold tier raises the stakes substantially for business-critical applications that cannot tolerate high RTO and RPO for any disasters such as database, cluster, corruptions, or site failures. Additionally, major database upgrades or site migrations can be done in seconds.

The Gold tier also reduces costs while improving your return on investment by actively using all of the replicas at all times.

The Gold reference architecture adds database-aware replication technologies, Oracle Data Guard and Oracle Active Data Guard and Oracle GoldenGate, which synchronize one or more replicas of the production database to provide real time data protection and availability. Database-aware replication substantially enhances high availability

and data protection (corruption protection) beyond what is possible with storage replication technologies. Oracle Active Data Guard Far Sync is used for zero data loss protection at any distance.

Platinum Reference Architecture

The Platinum tier introduces several new Oracle Database capabilities, including Oracle GoldenGate for zero-downtime upgrades and migrations.

Edition Based Redefinition lets application developers design for zero-downtime application upgrades. You can alternatively design applications for Oracle Sharding, which provides extreme availability by distributing subsets of a database into highly available shards, while the application can access the entire database as one single logical database.

Each of these technologies requires additional effort to implement, but they deliver substantial value for the most critical applications where downtime is not an option.

High Availability and Data Protection Attributes by Tier

Each MAA reference architecture delivers known and tested levels of downtime and data protection.

The following table summarizes the high availability and data protection attributes inherent to each architecture. Each architecture includes all of the capabilities of the previous architecture, and builds upon it to handle an expanded set of outages. The various components included and the service levels achieved by each architecture are described in other topics.

Table 2-1 High Availability and Data Protection Attributes By MAA Reference Architecture

MAA Reference Architecture	Unplanned Outages (Local Site)	Planned Maintenance	Data Protection	Unrecoverable Local Outages and Disaster Recovery
Bronze	Single Instance, auto-restart for recoverable instance and server failures. Redundancy for system infrastructure so that single component failures such as disk, flash, and network should not result in downtime.	Some online, most off-line	Basic runtime validation combined with manual checks	Restore from backup, potential to lose data generated since the last backup. Using Zero Data Loss Recovery Appliance reduces the potential to lose data to zero or near zero.

Table 2-1 (Cont.) High Availability and Data Protection Attributes By MAA Reference Architecture

MAA Reference Architecture	Unplanned Outages (Local Site)	Planned Maintenance	Data Protection	Unrecoverable Local Outages and Disaster Recovery
Silver	HA with automatic failover for instance and server failures	Most rolling, some online, few offline	Basic runtime validation combined with manual checks	Restore from backup, potential to lose data generated since the last backup. Using Zero Data Loss Recovery Appliance reduces the potential to lose data to zero or near zero. In-flight transactions are preserved with Application Continuity.
Gold	Comprehensive high availability and disaster recovery	All rolling or online	Comprehensive runtime validation combined with manual checks	Real-time failover, zero or near-zero data loss
Platinum	Zero application outage for Platinum ready applications	Zero application outage	Comprehensive runtime validation combined with manual checks	Zero application outage for Platinum-ready applications, with zero data loss. Oracle RAC, Oracle Active Data Guard, and Oracle GoldenGate complement each other, providing a wide array of solutions to achieve zero database service downtime for unplanned outages. Alternatively, use Oracle Sharding for site failure protection, because impact on the application is only on shards in failed site rather than the entire database. Each shard can be configured with real-time failover, zero or near-zero data loss, or zero application outage for Platinum-ready applications. In-flight transactions are preserved, with zero data loss.



See Also:

Oracle Maximum Availability Architecture

<http://www.oracle.com/goto/maa>

3

Features for Maximizing Availability

Familiarize yourself with the following Oracle Database high availability features used in MAA solutions.

- [Oracle Data Guard](#)
Oracle Data Guard ensures high availability, data protection, and disaster recovery for enterprise data.
- [Oracle GoldenGate](#)
Oracle GoldenGate is Oracle's strategic logical replication solution for data distribution and data integration.
- [Best Practice: Oracle Active Data Guard and Oracle GoldenGate](#)
While Oracle Active Data Guard and Oracle GoldenGate are each capable of maintaining a synchronized copy of an Oracle database, each has unique characteristics that result in high availability architectures that can use one technology or the other, or both at the same time, depending upon requirements.
- [Recovery Manager](#)
Recovery Manager (RMAN) provides a comprehensive foundation for efficiently backing up and recovering the database. RMAN eliminates operational complexity while providing superior performance and availability of the database.
- [Oracle Real Application Clusters and Oracle Clusterware](#)
Oracle RAC and Oracle Clusterware enable Oracle Database to run any packaged or custom application across a set of clustered servers.
- [Oracle RAC One Node](#)
Oracle Real Application Clusters One Node (Oracle RAC One Node) is a single instance of an Oracle RAC database that runs on one node in a cluster.
- [Oracle Automatic Storage Management](#)
Oracle ASM provides a vertically integrated file system and volume manager directly in the Oracle Database kernel.
- [Fast Recovery Area](#)
The fast recovery area is a unified storage location for all recovery-related files and activities in Oracle Database.
- [Corruption Prevention, Detection, and Repair](#)
- [Data Recovery Advisor](#)
Data Recovery Advisor automatically diagnoses persistent (on-disk) data failures, presents appropriate repair options, and runs repair operations at your request.
- [Oracle Flashback Technology](#)
Oracle Flashback technology is a group of Oracle Database features that let you view past states of database, database objects, transactions or rows or to rewind the database, database objects, transactions or rows to a previous state without using point-in-time media recovery.
- [Oracle Data Pump and Data Transport](#)

- [Oracle Replication Technologies for Non-Database Files](#)
Oracle ASM Cluster File System, Oracle Database File System, and Oracle Solaris ZFS Storage Appliance Replication are the Oracle replication technologies for non-database files.
- [Oracle Multitenant](#)
Oracle Multitenant is the optimal database consolidation method. The multitenant architecture combines the best attributes of each of the previous consolidation methods without their accompanying tradeoffs.
- [Oracle Sharding](#)
Oracle Sharding is a scalability and availability feature for applications explicitly designed to run on a sharded database.
- [Oracle Restart](#)
Oracle Restart enhances the availability of a single-instance (nonclustered) Oracle database and its components.
- [Oracle Site Guard](#)
- [Online Reorganization and Redefinition](#)
One way to enhance availability and manageability is to allow user access to the database during a data reorganization operation.
- [Zero Data Loss Recovery Appliance](#)
The cloud-scale Zero Data Loss Recovery Appliance, commonly known as Recovery Appliance, is an engineered system designed to dramatically reduce data loss and backup overhead for all Oracle databases in the enterprise.
- [Fleet Patching and Provisioning](#)
Fleet Patching and Provisioning maintains a space-efficient repository of software, more precisely "gold images," which are standardized software homes that can be provisioned to any number of target machines.
- [Enabling Continuous Service for Applications](#)
Applications achieve continuous service when planned maintenance, unplanned outages, and load imbalances of the database tier are hidden.
- [Learn About Other High Availability Features](#)
In addition to the Oracle Database features for high availability listed above, you may be interested in the following resources.

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 outages of any kind, including natural disasters and data corruptions. A Data Guard standby database is an exact replica of the production database and thus can be transparently utilized in combination with traditional backup, restoration, flashback, and cluster techniques to provide the highest possible level of data protection, data availability and disaster recovery. Data Guard is included in Oracle Enterprise Edition.

A Data Guard configuration consists of one primary database and one or more standby databases. A primary database can be either a single-instance Oracle database or an Oracle RAC database. Similar to a primary database, a standby database can be either a single-instance Oracle database or an Oracle RAC

database. Using a backup copy of the primary database, you can create up to 30 standby databases that receive redo directly from the primary database. Optionally you can use a cascaded standby to create Data Guard configurations where the primary transmits redo to a single remote destination, and that destination forwards redo to multiple standby databases. This enables a primary database to efficiently synchronize many more than 30 standby databases if desired.

Note:

Oracle Active Data Guard is an extension of basic Data Guard providing advanced features that off-load various types of processing from a production database, extend zero data loss protection over any distance, and that enhance high availability. Oracle Active Data Guard is licensed separately from Oracle Database Enterprise Edition. Oracle Active Data Guard is discussed more completely in [Oracle Active Data Guard](#).

There are several types of standby databases. Data Guard physical standby database is the MAA best practice for data protection and disaster recovery and is the most common type of standby database used. A physical standby database uses Redo Apply (an extension of Oracle media recovery) to maintain an exact, physical replica of the production database. When configured using MAA best practices, Redo Apply uses multiple Oracle-aware validation checks to prevent corruptions that can impact a primary database from impacting the standby. Other types of Data Guard standby databases include: snapshot standby (a standby open read/write for test or other purposes) and logical standby (used to reduce planned downtime).

Benefits of Using Data Guard

- Continuous Oracle-aware validation of all changes using multiple checks for physical and logical consistency of structures within an Oracle data block and redo, before updates are applied to a standby database. This isolates the standby database and prevents it from being impacted by data corruptions that can occur on the primary system.
- Transparent operation: There are no restrictions on the use of Data Guard physical standby for data protection. Redo Apply supports all data and storage types, all DDL operations, and all applications (custom and packaged applications), and guarantees data consistency across primary and standby databases.
- Highest performance: Fast redo transport for best recovery point objective, fast apply performance for best recovery time objective. Multi-instance redo apply provides Oracle RAC scalability for redo apply, eliminating bottlenecks of a single database server. Redo apply can essentially scale up to available CPU, I/O, and network across your Oracle RAC cluster. An observed redo apply rate of 3500 MB per second (12 TB/hour) on 8 node RAC Exadata.
- Fast failover to a standby database to maintain availability should the primary database fail for any reason. Failover is either a manual or automatic operation depending on how Data Guard is configured.
- Integrated client notification framework to enable application clients to connect to a new primary database after a failover occurs.
- Automatic or automated (depending upon configuration) resynchronization of a failed primary database, quickly converting it to a synchronized standby database after a failover occurs.
- Choice of flexible data protection levels to support all network configurations, availability and performance SLAs, and business requirements.

- Management of a primary and all of its standby databases as a single configuration to simplify management and monitoring using either the Data Guard Broker command-line interface or Oracle Enterprise Manager Cloud Control.
- Data Guard Broker greatly improves manageability with additional features for comprehensive configuration health checks, resumable switchover operations, streamlined role transitions, support for cascaded standby configurations, and user-configurable thresholds for transport and apply lag to automatically monitor the ability of the configuration to support SLAs for recovery point and recovery time objectives at any instant in time.
- Efficient transport to multiple remote destinations using a single redo stream originating from the primary production database and forwarded by a cascading standby database.
- Snapshot Standby enables a physical standby database to be open read/write for testing or any activity that requires a read/write replica of production data. A snapshot standby continues to receive but does not apply updates generated by the primary. When testing is complete, a snapshot standby is converted back into a synchronized physical standby database by first discarding the changes made during the open read/write, and then applying the redo received from the primary database. Primary data is always protected. Snapshot standby is particularly useful when used in conjunction with Oracle Real Application Testing (workload is captured at the production database for replay and subsequent performance analysis at the standby database-an exact replica of production).
- Reduction of planned downtime by utilizing a standby database to perform maintenance in rolling fashion. The only downtime is the time required to perform a Data Guard switchover; applications remain available while the maintenance is being performed. (See [When to Use Oracle Active Data Guard and Oracle GoldenGate Together](#) and [Oracle High Availability Solutions for System and Software Maintenance](#) for more details).
- Increased flexibility for Data Guard configurations where the primary and standby systems may have different CPU architectures or operating systems subject to limitations defined in My Oracle Support note [413484.1](#).
- Efficient disaster recovery for a container database (CDB). Data Guard failover and switchover completes using a single command at a CDB level regardless of how many pluggable databases (PDBs) are consolidated within the CDB.
- Enables a specific administration privilege, SYSDG, to handle standard administration duties for Data Guard. This new privilege is based on the least privilege principle, in which a user is granted only the necessary privileges required to perform a specific function and no more. The SYSDBA privilege continues to work as in previous releases.
- The Oracle Database In-Memory column store is supported on standby databases in an Active Data Guard environment.
- Further improves performance and availability of Data Warehouses in a Data Guard configuration by tracking information from

NOLOGGING

operations so they can be repaired with the new RMAN command

```
RECOVER DATABASE
      NOLOGGING
```

- Improves the impact multiple SYNC transport destinations have on the primary database through the use of a new parameter

```
DATA_GUARD_SYNC_LATENCY
```

. This parameter defines the maximum amount of time (in seconds) that the Primary database must wait before disconnecting subsequent destinations after at least one synchronous standby has acknowledged receipt of the redo.

- Data Guard Broker improves manageability by supporting destinations of different Endianness than the primary in addition to enhancing management of alternate destinations.
- Data Guard improves protection and Return To Operations (RTO) and Recovery Point Objectives (RPO) through multiple features including:
 - Multi Instance Redo Apply (MIRA) provides scalable redo apply performance across Oracle RAC instances reducing RTO for even higher production OLTP or batch workloads
 - Compare primary and standby database blocks using the new

```
DBMS_DBCOMP
```

package to help identify lost writes so they can be resolved efficiently.

- Fast Start Failover (FSFO) has the robustness of highly available zero data loss configurations with support for Maximum Protection mode while giving the flexibility of multiple observers and multiple failover targets for high availability in any configuration. FSFO can also be configured to automatically fail over to the standby with the detection of a lost write on the primary .
- RPO is improved with no data loss failovers after a storage failure in ASYNC configurations and Data Guard Broker support for Application Continuity, improving the user experience during Data Guard role transitions.
- Oracle Data Guard Broker further improves the management of databases by supporting destinations of different endianness than the primary in addition to enhancing management of alternate archive destinations when the primary destination is unavailable.
- Oracle Data Guard Database Compare tool compares data blocks stored in an Oracle Data Guard primary database and its physical standby databases. Use this tool to find disk errors (such as lost write) that cannot be detected by other tools like the DBVERIFY utility. (new in Oracle Database 12c Release 2)
- Oracle Data Guard Broker supports multiple automatic failover targets in a fast-start failover configuration. Designating multiple failover targets significantly improves the likelihood that there is always a standby suitable for automatic failover when needed. (new in Oracle Database 12c Release 2)

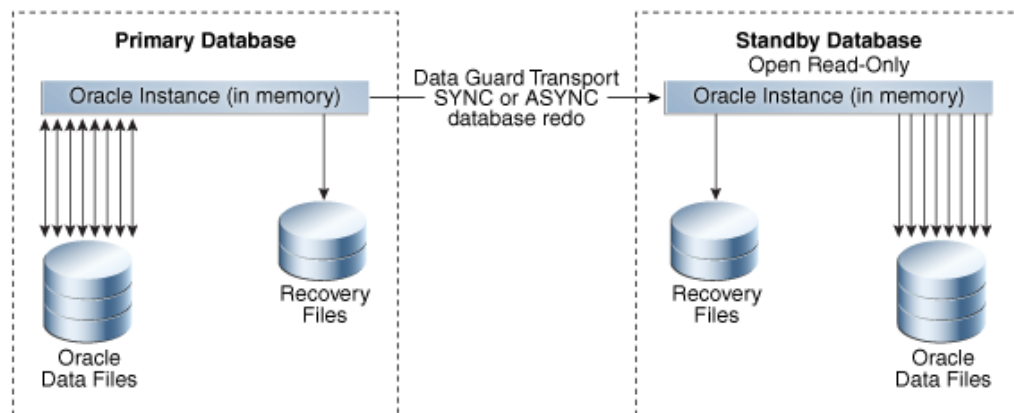
- [Oracle Active Data Guard](#)
Oracle Active Data Guard is Oracle's strategic solution for real time data protection and disaster recovery for the Oracle database using a physical replication process.
- [Oracle Data Guard Advantages Over Traditional Solutions](#)
Oracle Data Guard provides a number of advantages over traditional solutions.
- [Data Guard and Planned Maintenance](#)

Oracle Active Data Guard

Oracle Active Data Guard is Oracle's strategic solution for real time data protection and disaster recovery for the Oracle database using a physical replication process.

Oracle Active Data Guard also provides high return on investment in disaster recovery systems by enabling a standby database to be open read-only while it applies changes received from the primary database. Oracle Active Data Guard is a separately licensed product that provides advanced features that greatly expand Data Guard capabilities included with Oracle Enterprise Edition.

Figure 3-1 Oracle Active Data Guard Architecture



Oracle Active Data Guard enables administrators to improve performance by offloading processing from the primary database to a physical standby database that is open read-only while it applies updates received from the primary database. Offload capabilities of Oracle Active Data Guard include read-only reporting and ad-hoc queries (including DML to global temporary tables and unique global or session sequences), data extracts, fast incremental backups, redo transport compression, efficient servicing of multiple remote destinations, and the ability to extend zero data loss protection to a remote standby database without impacting primary database performance. Oracle Active Data Guard also increases high availability by performing automatic block repair and enabling High Availability Upgrades automation.

Note:

Oracle Active Data Guard is licensed separately as a database option license for Oracle Database Enterprise Edition. All Oracle Active Data Guard capabilities are also included in an Oracle Golden Gate license for Oracle Enterprise Edition. This provides customers with the choice of a standalone license for Oracle Active Data Guard, or licensing Oracle GoldenGate to acquire access to all advanced Oracle replication capabilities.

Benefits of Oracle Active Data Guard

Oracle Active Data Guard inherits all of the benefits previously listed for Data Guard, plus the following:

- Improves primary database performance: Production-offload to an Oracle Active Data Guard standby database of read-only applications, reporting, and ad hoc queries. Any application compatible with a read-only database can run on an Oracle Active Data Guard standby. Oracle also provides integration that enables the offloading of many Oracle E-Business Suite Reports, PeopleTools reporting, Oracle Business Intelligence Enterprise Edition (OBIEE), and Oracle TopLink applications to an Oracle Active Data Guard standby database.
- DML global temporary tables and the use of sequences at the standby database significantly expands the number of read-only applications that can be off-loaded from production databases to an Oracle Active Data Guard standby database.
- The unique ability to easily scale read performance using multiple Oracle Active Data Guard standby databases, also referred to as a Reader Farm.
- Production-offload of data extracts using Oracle Data Pump or other methods that read directly from the source database.
- Production-offload of the performance impact from network latency in a synchronous, zero data loss configuration where primary and standby databases are separated by hundreds or thousands of miles. Far sync uses a lightweight instance (control file and archive log files, but no recovery and no data files), deployed on a system independent of the primary database. The far sync instance is ideally located at the maximum distance from the primary system that an application can tolerate the performance impact of synchronous transport to provide optimal protection. Data Guard transmits redo synchronously to the far sync instance and far sync forwards the redo asynchronously to a remote standby database that is the ultimate failover target. If the primary database fails, the same failover command used for any Data Guard configuration, or mouse click using Oracle Enterprise Manager Cloud Control, or automatic failover using Data Guard Fast-Start Failover executes a zero data loss failover to the remote destination. This transparently extends zero data loss protection to a remote standby database just as if it were receiving redo directly from the primary database, while avoiding the performance impact to the primary database of WAN network latency in a synchronous configuration.
- Production-offload of the overhead of servicing multiple remote standby destinations using far sync. In a far sync configuration, the primary database ships a single stream of redo to a far sync instance using synchronous or asynchronous transport. The far sync instance is able to forward redo asynchronously to as many as 29 remote destinations with zero incremental overhead on the source database.
- Data Guard maximum availability supports the use of the

NOAFFIRM

redo transport attribute. A standby database returns receipt acknowledgment to its primary database as soon as redo is received in memory. The standby database does not wait for the Remote File Server (RFS) to write to a standby redo log file.

This feature provides increased primary database performance in Data Guard configurations using maximum availability and SYNC redo transport. Fast Sync isolates the primary database in a maximum availability configuration from any

performance impact due to slow I/O at a standby database. This new FAST SYNC feature can work with a physical standby target or within a far sync configuration.

- Production-offload of CPU cycles required to perform redo transport compression. Redo transport compression can be performed by the far sync instance if the Data Guard configuration is licensed for Oracle Advanced Compression. This conserves bandwidth with zero incremental overhead on the primary database.
- Production-offload and increased backup performance by moving fast incremental backups off of the primary database and to the standby database by utilizing Oracle Active Data Guard support for RMAN block change tracking.
- Increased high availability using Oracle Active Data Guard automatic block repair to repair block corruptions, including file header corruptions, detected at either the primary or standby, transparent to applications and users.
- Increased high availability by reducing planned downtime for upgrading to new Oracle Database patch sets and database releases using the additional automation provided by high availability Upgrade.
- Connection preservation on an Active Data Guard standby through a role change facilitates improved reporting and improves the user experience. The connections pause while the database role changes to a primary database and resume, improving the user experience.
- The Oracle Enterprise Manager Diagnostic tool can be used with Active Data Guard to capture and send performance data to the Automatic Workload Repository, while the SQL Tuning Advisor allows primary database SQL statement tuning to be offloaded to a standby database.
- Active Data Guard support for the Oracle Database In-Memory option enables reporting to be offloaded to the standby database while reaping the benefits the In-Memory option provides, including tailored column stores for the standby database workload.

Oracle Data Guard Advantages Over Traditional Solutions

Oracle Data Guard provides a number of advantages over traditional solutions.

- Fast, automatic or automated database failover for data corruptions, lost writes, and database and site failures, with recovery times of potentially seconds with Data Guard as opposed to hours with traditional solutions
- Zero data loss over wide area network using Oracle Active Data Guard Far Sync
- Offload processing for redo transport compression and redo transmission to up to 29 remote destinations using Oracle Active Data Guard Far Sync
- Automatic corruption repair automatically replaces a physical block corruption on the primary or physical standby by copying a good block from a physical standby or primary database
- Most comprehensive protection against data corruptions and lost writes on the primary database
- Reduced downtime for storage, Oracle ASM, Oracle RAC, system migrations and some platform migrations, and changes using Data Guard switchover
- Reduced downtime for database upgrades with Data Guard rolling upgrade capabilities

- Ability to off-load primary database activities—such as backups, queries, or reporting—without sacrificing the RTO and RPO ability to use the standby database as a read-only resource using the real-time query apply lag capability, including Database In-Memory column support
- Ability to integrate non-database files using Oracle Database File System (DBFS) or Oracle Automatic Storage Management Cluster File System (Oracle ACFS) as part of the full site failover operations (see [Oracle Replication Technologies for Non-Database Files](#))
- No need for instance restart, storage remastering, or application reconnections after site failures
- Transparency to applications
- Transparent and integrated support (application continuity and transaction guard) for application failover
- Effective network utilization
- Database In-Memory support
- Integrated service and client failover that reduces overall application RTO
- Enhanced and integrated Data Guard awareness with existing Oracle technologies such as Oracle RAC, RMAN, Oracle GoldenGate, Enterprise Manager, health check (orachk), DBCA, and Fleet Patch and Provisioning

For data resident in Oracle databases, 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. 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.

Data Guard and Planned Maintenance

Data Guard standby databases can be used to reduce planned downtime by performing maintenance in a rolling fashion. Changes are implemented first at the standby database. The configuration is allowed to run with the primary at the old version and standby at the new version until there is confidence that the new version is ready for production. A Data Guard switchover can be performed, transitioning production to the new version or same changes can be applied to production in a rolling fashion. The only possible database downtime is the time required to perform the switchover.

There are several approaches to performing maintenance in a rolling fashion using a Data Guard standby. Customer requirements and preferences determine which approach is used.

- [Data Guard Redo Apply and Standby-First Patching](#)
Beginning with Oracle Database 10g, there has been increased flexibility in cross-platform support using Data Guard Redo Apply.
- [Data Guard Transient Logical Rolling Upgrades](#)
- [Rolling Upgrade Using Oracle Active Data Guard](#)
Rolling database upgrade using Oracle Active Data Guard provides a simpler, automated, and easily repeatable method for reducing planned downtime than represented by the manual Transient Logical rolling upgrade procedure.

Data Guard Redo Apply and Standby-First Patching

Beginning with Oracle Database 10g, there has been increased flexibility in cross-platform support using Data Guard Redo Apply.

In certain Data Guard configurations, primary and standby databases are able to run on systems having different operating systems (for example, Windows and Linux), word size (32bit/64bit), different storage, different Exadata hardware and software versions, or different hardware architectures. Redo Apply can also be used to migrate to Oracle Automatic Storage Management (ASM), to move from single instance Oracle databases to Oracle RAC, to perform technology refresh, or to move from one data center to the next.

Beginning with Oracle Database 11g Release 2 (11.2), Standby-First Patch Apply (physical standby using Redo Apply) can support different database software patch levels between a primary database and its physical standby database for the purpose of applying and validating Oracle patches in a rolling fashion. Patches eligible for Standby-First patching include:

- Database Release Updates (RUs) or Release Update Revisions (RURs)
- Database Patch Set Update (PSU)
- Database Critical Patch Update (CPU)
- Database bundled patch

Standby-First Patch Apply is supported for certified database software patches for Oracle Database Enterprise Edition 11g Release 2 (11.2) and later.

In each of the types of planned maintenance previously described, the configuration begins with a primary and physical standby database (in the case of migration to a new platform, or to ASM or Oracle RAC, the standby is created on the new platform). After all changes are implemented at the physical standby database, Redo Apply (physical replication) is used to synchronize the standby with the primary. A Data Guard switchover is used to transfer production to the standby (the new environment).

See Also:

My Oracle Support Note [413484.1](#) for information about mixed platform combinations supported in a Data Guard configuration.

My Oracle Support Note [1265700.1](#) for more information about Standby First Patch Apply and the README for each patch to determine if a target patch is certified as being a Standby-First Patch.

Data Guard Transient Logical Rolling Upgrades

There are numerous types of maintenance tasks that are unable to use Redo Apply (physical replication) to synchronize the original version of a database with the changed or upgraded version. These tasks include:

- Database patches or upgrades that are not Standby-First Patch Apply-eligible. This includes database patch-sets (11.2.0.2 to 11.2.0.4) and upgrade to new Oracle Database releases (12.2 to 18c).
- Maintenance must be performed that modifies the physical structure of a database that would require downtime (for example, adding partitioning to non-partitioned tables, changing Basicfile LOBs to Securefile LOBs, changing XML-CLOB to Binary XML, or altering a table to be OLTP-compressed).

All of the previous types of maintenance can be performed in a rolling fashion using a Data Guard standby database by using Data Guard SQL Apply (logical replication) to synchronize the old and new versions of the database. Prior to Oracle Database 11g this required creating a logical standby database, performing the maintenance on the logical standby, resynchronizing the standby with the primary, and then switching over. Additionally if a physical standby was being used for disaster recovery, then a new physical standby database would have to be created from a backup of the production database at the new version. This represented a number of logistical and cost challenges when upgrading a multi-terabyte database.

Beginning with Oracle Database 11g, database rolling upgrades can use a new procedure called Transient Logical that begins and ends with a physical standby database. SQL Apply is only used during the phase when Data Guard is synchronizing across old and new versions. A new logical standby database does not need to be created if there is already a physical standby in place. A new physical standby database does not need to be created from a backup of the production database at the new version after the maintenance is complete. Similar to the traditional process of upgrading a Data Guard configuration having an in-place physical standby, the original primary is upgraded or changed using redo from the new primary database and Redo Apply (a single catalog upgrade migrates both primary and standby databases to the new Oracle release).

Transient Logical upgrades require that the primary database be at Oracle Database 11g release 1 (11.1) or later and that the database meet the prerequisites of SQL Apply.

Oracle provides a Bourne shell script that automates a number of the manual steps required by the Transient Logical rolling upgrade process.

Databases that use Oracle Database Vault can be upgraded to new Oracle Database releases and patch sets by using Oracle Data Guard database rolling upgrades (transient logical standby only).

 **See Also:**

<http://www.oracle.com/goto/maa> for Oracle MAA white paper “Oracle Database Rolling Upgrades: Using a Data Guard Physical Standby Database”

Rolling Upgrade Using Oracle Active Data Guard

Rolling database upgrade using Oracle Active Data Guard provides a simpler, automated, and easily repeatable method for reducing planned downtime than represented by the manual Transient Logical rolling upgrade procedure.

Rolling upgrade using Oracle Active Data Guard transforms the 42 or more steps required by the manual procedure into several easy-to-use DBMS_ROLLING PL/SQL packages. Rolling upgrades performed using the DBMS_ROLLING PL/SQL package are supported on a multitenant container database (CDB).

A rolling upgrade using Oracle Active Data Guard:

- Generates an upgrade plan with a configuration-specific set of instructions to guide you through the upgrade process.
- Modifies parameters of the rolling upgrade.
- Configures primary and standby databases participating in the upgrade.
- Performs switchover of the production database to the new version. Switchover is the only downtime required.
- Completes the upgrade of the old primary and any additional standby databases in the Data Guard configuration and resynchronizes with the new primary.

Rolling upgrade using Oracle Active Data Guard has the following additional benefits:

- Provides a simple specify-compile-execute protocol
 - Catches configuration errors at the compilation step
 - Runtime errors are detected during execution
- The state is kept in the database
 - Enables a reliable, repeatable process
- Runtime steps are constant regardless of how many databases are involved
- Handles failure at the original primary database
- Enables data protection for the upgraded primary at all times



See Also:

<http://www.oracle.com/goto/maa> for Oracle MAA white paper “Oracle Database Rolling Upgrades: Using a Data Guard Physical Standby Database”

Oracle Data Guard Concepts and Administration

Oracle GoldenGate

Oracle GoldenGate is Oracle's strategic logical replication solution for data distribution and data integration.

Oracle GoldenGate offers a real-time, log-based change data capture and replication software platform. The software provides capture, routing, transformation, and delivery of transactional data across heterogeneous databases in real time.

Unlike replication solutions from other vendors, Oracle GoldenGate is more closely integrated with Oracle Database while also providing an open, modular architecture ideal for replication across heterogeneous database management systems. This combination of attributes eliminates compromise, making Oracle GoldenGate the

preferred logical replication solution for addressing requirements that span Oracle Database and non-Oracle Database environments.

A typical environment includes a capture, pump, and delivery process. Each of these processes can run on most of the popular operating systems and databases, including Oracle Database. All or a portion of the data can be replicated, and the data within any of these processes can be manipulated for not only heterogeneous environments but also different database schemas, table names, or table structures. Oracle GoldenGate also supports bidirectional replication with preconfigured conflict detection and resolution handlers to aid in resolving data conflicts.

Oracle GoldenGate logical replication enables all databases in an Oracle GoldenGate configuration, both source and target databases, to be open read-write. This makes it a key component of MAA for addressing a broad range of high availability challenges for zero downtime maintenance, cross platform migration, and continuous data availability, specifically:

- **Zero or near zero downtime maintenance.** In this architecture, Oracle GoldenGate provides greater flexibility than the capabilities provided by Data Guard. Oracle GoldenGate source and target databases can have a different physical and logical structure, can reside on different hardware and operating system architectures, can span wide differences in Oracle Database releases (for example, 12.1 to 18c), or be a mix of Oracle and non-Oracle systems. This allows for the modernization of 24x7 servers and allows new Oracle features to be implemented without impacting the availability of the databases. Maintenance is first performed on a target database while production runs on the source. After the maintenance is complete, production can be moved to the source all at once, similar to a Data Guard switchover. Optionally, bidirectional replication can be used to gradually move users over to the new system to create the perception of zero downtime. In either case, Oracle GoldenGate replication can be enabled in the reverse direction to keep the original source database synchronized during a transition period, making it simple to effect a planned fall-back to the previous version if needed, with minimal downtime and no data loss.
- **Zero or near-zero downtime migrations when a Data Guard solution is not applicable.** Platform or database migrations can be carried out using Oracle GoldenGate as the data synchronization method between the old and new systems. Once the database has been instantiated on another host, Oracle GoldenGate is configured to replicate changes from the production database. A guaranteed restore point can be created on the migrated database so that after user testing the database can be flashed back, and Oracle GoldenGate can apply any outstanding data changes from the production database before moving the application users to the new database, similar to a snapshot standby database. If desired, bi-directional replication can also be configured from the migrated database back to the production database for use as a fallback solution.
- **Zero or near-zero downtime application upgrades.** Application upgrades that modify back-end database objects typically result in significant planned downtime while maintenance is being performed. Oracle GoldenGate replication enables data transformations that map database objects used by a previous version of an application to objects modified by the new version of an application. This enables database maintenance to be performed on a separate copy of the production database without impacting the availability of the application. After the maintenance is complete and Oracle GoldenGate has finished synchronizing old and new versions, users can be switched to the new version of the application.
- **Read-write access to a replica database while it is being synchronized with its source database.** This is most often used to offload reporting to a copy of a

production database when the reporting application requires a read-write connection to database in order to function. This is also relevant to disaster recovery environments where the nature of the technology used for the application tier requires an active read-write connection to the DR database at all times in order to meet recovery time objectives.

- **Active-Active replication.** Oracle GoldenGate supports an active-active multi-directional configuration, where there are two or more systems with identical sets of data that can be changed by application users on either system. Oracle GoldenGate replicates transactional data changes from each database to the others to keep all sets of data current.
- Seamless moves between Oracle Real Application Clusters (RAC) nodes in the event of database instance failure or during applicable maintenance operations. This ability provides high availability with Oracle GoldenGate and it is possible to patch and upgrade the Oracle GoldenGate software on one or more nodes in the cluster without affecting the node where Oracle GoldenGate is currently running. Then at a predetermined time, Oracle GoldenGate can be switched to one of the upgraded nodes. The switch is done without reconfiguring Oracle GoldenGate because configuration information is shared across the Oracle RAC cluster.



See Also:

[Oracle GoldenGate Documentation](#)

<http://www.oracle.com/goto/maa> for Oracle MAA Oracle GoldenGate white papers

Best Practice: Oracle Active Data Guard and Oracle GoldenGate

While Oracle Active Data Guard and Oracle GoldenGate are each capable of maintaining a synchronized copy of an Oracle database, each has unique characteristics that result in high availability architectures that can use one technology or the other, or both at the same time, depending upon requirements.

Examples of MAA Best Practice guidelines are as follows:

- [When to Use Oracle Active Data Guard](#)
Use Oracle Active Data Guard when the emphasis is on simplicity, data protection, and availability.
- [When to Use Oracle GoldenGate](#)
Use Oracle GoldenGate when the emphasis is on advanced replication requirements not addressed by Oracle Active Data Guard.
- [When to Use Oracle Active Data Guard and Oracle GoldenGate Together](#)
Oracle Active Data Guard and Oracle GoldenGate are not mutually exclusive. The following are use cases of high availability architectures that include the simultaneous use of Oracle Active Data Guard and Oracle GoldenGate.

When to Use Oracle Active Data Guard

Use Oracle Active Data Guard when the emphasis is on simplicity, data protection, and availability.

- Simplest, fastest, one-way replication of a complete Oracle database.
- No restrictions: Data Guard Redo Apply supports all data and storage types and Oracle features; transparent replication of DDL
- Features optimized for data protection: Detects silent corruptions that can occur on source or target; automatically repairs corrupt blocks
- Synchronized standby open read-only provides simple read-only offloading for maximum ROI
- Transparency of backups: A Data Guard primary and standby are physically exact copies of each other; RMAN backups are completely interchangeable
- Zero data loss protection at any distance, without impacting database performance
- Minimizing planned downtime and risk using standby first patching, database rolling upgrades, and select platform migrations
- Reduce risk of introducing change by dual purposing a DR system for testing using Data Guard Snapshot Standby
- Integrated automatic database and client failover
- Integrated management of a complete configuration: Data Guard Broker command line interface or Oracle Enterprise Manager Cloud Control

When to Use Oracle GoldenGate

Use Oracle GoldenGate when the emphasis is on advanced replication requirements not addressed by Oracle Active Data Guard.

- Any requirement where the replica database must be open read/write while synchronizing with the primary database
- Any data replication requirements such as multimaster and bidirectional replication, subset replication, many-to-one replication, and data transformations.
- When data replication is required between endian format platforms or across-database major versions.
- Maintenance and migrations where zero downtime or near zero downtime is required. Oracle GoldenGate can be used to migrate between application versions, for example, from Application 1.0 to Application 2.0 without downtime.
- Database rolling upgrades where it is desired to replicate from new version down to the old version for the purpose of fast fall-back if something is wrong with the upgrade.
- Zero downtime planned maintenance where bidirectional replication is used to gradually migrate users to the new version, creating the perception of zero downtime. Note that bidirectional replication requires avoiding or resolving update conflicts that can occur on disparate databases.

When to Use Oracle Active Data Guard and Oracle GoldenGate Together

Oracle Active Data Guard and Oracle GoldenGate are not mutually exclusive. The following are use cases of high availability architectures that include the simultaneous use of Oracle Active Data Guard and Oracle GoldenGate.

- An Oracle Active Data Guard standby is utilized for disaster protection and database rolling upgrades for a mission critical OLTP database. At the same time, Oracle GoldenGate is used to replicate data from the Data Guard primary database (or from the standby database using Oracle GoldenGate ALO mode) for ETL update of an enterprise data warehouse.
- Oracle GoldenGate subset replication is used to create an operational data store (ODS) that extracts, transforms, and aggregates data from numerous data sources. The ODS supports mission critical application systems that generate significant revenue for the company. An Oracle Active Data Guard standby database is used to protect the ODS, providing optimal data protection and availability.
- Oracle GoldenGate bidirectional replication is utilized to synchronize two databases separated by thousands of miles. User workload is distributed across each database based upon geography, workload, and service level using Global Data Services (GDS). Each Oracle GoldenGate copy has its own local synchronous Data Guard standby database that enables zero data loss failover if an outage occurs. Oracle GoldenGate capture and apply processes are easily restarted on the new primary database following a failover because the primary and standby are an exact, up-to-date replica of each other.
- An Oracle Active Data Guard standby database used for disaster protection is temporarily converted into an Oracle GoldenGate target for the purpose of performing planned maintenance not supported by Data Guard. For example, a Siebel application upgrade requiring modification of back-end database objects which require comprehensive testing before switching users over to the new system.
- Oracle Active Data Guard is used to protect a production environment when a major database version upgrade is required offering zero or near-zero downtime (for example, Oracle 18c to 19c.) A second primary/standby environment is created using the new database version, and Oracle GoldenGate is used to replicate data from the production environment to the copy with one-way or bidirectional replication. When Oracle GoldenGate has completed synchronizing the old and new environments, production is switched to the new environment and the old environment is decommissioned. This provides zero or minimal downtime depending upon configuration, eliminates risk by providing complete isolation between the old and new environment, and avoids any impact to data protection and availability SLAs if problems are encountered during the upgrade process.

 **See Also:**

<http://www.oracle.com/goto/maa> for Oracle MAA Best Practices white paper
“Transparent Role Transitions With Oracle Data Guard and Oracle GoldenGate”

Recovery Manager

Recovery Manager (RMAN) provides a comprehensive foundation for efficiently backing up and recovering the database. RMAN eliminates operational complexity while providing superior performance and availability of the database.

RMAN 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. RMAN and the server automatically identify modifications to the structure of the database and dynamically adjust the required operation to adapt to the changes.

RMAN is the standard interface to backup and restore from Recovery Appliance, local disk (ZFS storage), tape, and cloud object store.

RMAN provides the following benefits:

- Support for Oracle Sharding - RMAN support for every independent database (shard)
- Enhancement for Sparse Databases - allows backup and restore to operate on

SPARSE

backup sets and or image copies

- Over the Network Standby Database repair of

NONLOGGED

operation - new syntax for validation and repair on Standby -

VALIDATE/RECOVER .. NONLOGGED BLOCK;

- RMAN DUPLICATE

feature enhanced to support creation of Far Sync from Primary and backup

- RMAN DUPLICATE

Using Encrypted Backups - RMAN enhanced support non Auto-login wallet based encrypted backups with a new

SET

command - enables interrupt-free cloning

- Support for cross-platform backup and restore over the network
- Network-enabled restoration allows the

RESTORE

operations to copy data files directly from one database to another over the network

- Simplified table restoration with the

RECOVER TABLE

command

- Support for Oracle Multitenant, including backup and recovery of individual pluggable databases
- Support for cross-platform Oracle Multitenant, including backup and recovery of individual PDBs
- Automatic channel failover on backup and restore operations
- Automatic failover to a previous backup when the restore operation discovers a missing or corrupt backup
- Automatic creation of new database files and temporary files during recovery
- Automatic recovery through a previous point-in-time recovery—recovery through reset logs
- Block media recovery, which enables the data file 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 a virtual private recovery catalog
- Merger of incremental backups into image copies, providing up-to-date recoverability
- Optimized backup and restoration of required files only
- Retention policy to ensure that relevant backups are retained
- Ability to resume backup and restore operations in case of failure
- Automatic backup of the control file and the server parameter file, ensuring that backup metadata is available in times of database structural changes and media failure and disasters
- Easily reinstantiate a new database from an existing backup or directly from the production database (thus eliminating staging areas) using the

DUPLICATE

command.

 **See Also:**

Oracle Database Backup and Recovery User's Guide

Oracle Real Application Clusters and Oracle Clusterware

Oracle RAC and Oracle Clusterware enable 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, then 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, 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, enabling Oracle RAC to provide high availability and redundancy during failures in the cluster. Oracle RAC accommodates all system types, from read-only data warehouse 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 the databases and applications restart when a process fails or so that a failover to another node occurs after a 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. 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. Oracle Clusterware is the only clusterware that you need for most platforms on which Oracle RAC operates. Although Oracle Database continues to support third-party clusterware products on specified platforms, using Oracle Clusterware provides these main benefits:

- Dispenses with proprietary vendor clusterware
- Uses an integrated software stack from Oracle that provides disk management with local or remote Oracle Automatic Storage Management (Oracle Flex ASM) to data management with Oracle Database and Oracle RAC

- Can be configured in large clusters, called an Oracle Flex Cluster.

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. Oracle Clusterware requires that each node be connected to a private network over a private interconnect.

- [Benefits of Using Oracle Clusterware](#)
Oracle Clusterware provides the following benefits.
- [Benefits of Using Oracle Real Application Clusters and Oracle Clusterware](#)
Together, Oracle RAC and Oracle Clusterware provide all of the Oracle Clusterware benefits plus the following benefits.
- [Oracle RAC Advantages Over Traditional Cold Cluster Solutions](#)
Oracle RAC provides many advantages over traditional cold cluster solutions, including the following.

Benefits of Using Oracle Clusterware

Oracle Clusterware provides the following benefits.

- Tolerates and quickly recovers from computer and instance failures.
- Simplifies management and support by means of using Oracle Clusterware together with Oracle Database. By using fewer vendors and an all Oracle stack you gain better integration compared to using third-party clusterware.
- Performs rolling upgrades for system and hardware changes. For example, you can apply Oracle Clusterware upgrades, patch sets, and interim patches in a rolling fashion.

When you upgrade to Oracle Database 12c, Oracle Clusterware and Oracle ASM binaries are installed as a single binary called the Oracle Grid Infrastructure. You can upgrade Oracle Clusterware in a rolling manner from Oracle Clusterware 10g and Oracle Clusterware 11g; however, you can only upgrade Oracle ASM in a rolling manner from Oracle Database 11g release 1 (11.1).

- Automatically restarts failed Oracle processes.
- Automatically manages the virtual IP (VIP) address. When a node fails, the node's VIP address fails over to another node on which the VIP address can accept connections.
- Automatically restarts resources from failed nodes on surviving nodes.
- Controls Oracle processes as follows:
 - For Oracle RAC databases, Oracle Clusterware controls all Oracle processes by default.
 - For Oracle single-instance databases, Oracle Clusterware enables you to configure the Oracle processes into a resource group that is under the control of Oracle Clusterware.
- Provides an application programming interface (API) for Oracle and non-Oracle applications 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.
- Using server weight-based node eviction allows for aligning the choice of which node gets evicted in case of certain failures in the cluster with business requirements, ensuring that the most important workload is kept alive for as long as possible, assuming an equal choice between servers.
- Provides the ability to perform rolling release upgrades of Oracle Clusterware, with no downtime for applications.

Benefits of Using Oracle Real Application Clusters and Oracle Clusterware

Together, Oracle RAC and Oracle Clusterware provide all of the Oracle Clusterware benefits plus the following benefits.

- Provides better integration and support of Oracle Database by using an all Oracle software stack compared to using third-party clusterware.
- Relocate Oracle Service automatically. Plus, when you perform additional fast application notification (FAN) and client configuration, distribute FAN events so that applications can react immediately to achieve fast, automatic, and intelligent connection and failover.
- Detect connection failures fast and automatically, and remove terminated connections for any Java application using Oracle Universal Connection Pool (Oracle UCP) Fast Connection Failover and FAN events.
- Balance work requests using Oracle UCP runtime connection load balancing.
- Use runtime connection load balancing with Oracle UCP, Oracle Call Interface (OCI), and Oracle Data Provider for .NET (ODP.NET).
- Distribute work across all available instances using load balancing advisory.
- You can configure a database so that Oracle Clusterware is aware of the CPU requirements and limits for the given database. Oracle Clusterware uses this information to place the database resource only on servers that have a sufficient number of CPUs, amount of memory, or both.
- Allow the flexibility to increase processing capacity using commodity hardware without downtime or changes to the application.
- Provide comprehensive manageability integrating database and cluster features.
- Provide scalability across database instances.
- Implement Fast Connection Failover for nonpooled connections.

Oracle RAC Advantages Over Traditional Cold Cluster Solutions

Oracle RAC provides many advantages over traditional cold cluster solutions, including the following.

- Scalability across database instances
- Flexibility to increase processing capacity using commodity hardware without downtime or changes to the application

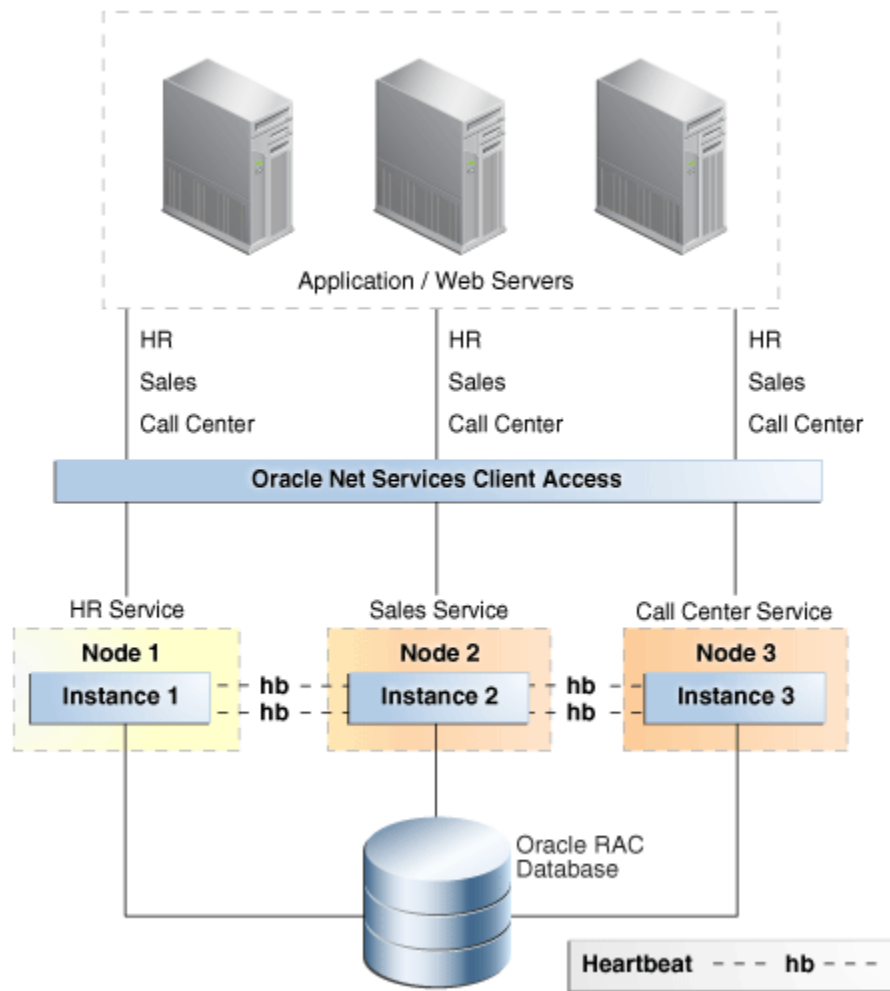
- Ability to tolerate and quickly recover from computer and instance failures (measured in seconds)
- Application brownout can be zero or seconds compared to minutes and hours with cold cluster solutions
- Optimized communication in the cluster over redundant network interfaces, without using bonding or other technologies

Oracle Grid Infrastructure and Oracle RAC make use of Redundant Interconnect Usage that distributes network traffic and ensures optimal communication in the cluster. This functionality is available starting with Oracle Database 11g Release 2 (11.2.0.2). In previous releases, technologies like bonding or trunking were used to make use of redundant networks for the interconnect.

- Rolling upgrades for system and hardware changes
- Rolling patch upgrades for some interim patches, security patches, CPUs, and cluster software
- Fast, automatic, and intelligent connection and service relocation and failover
- Comprehensive manageability integrating database and cluster features with Grid Plug and Play and policy-based cluster and capacity management
- Load balancing advisory and run-time connection load balancing help redirect and balance work across the appropriate resources
- Oracle Quality of Service (QoS) Management for policy-based run-time management of resource allocation to database workloads to ensure service levels are met in order of business need under dynamic conditions. This is accomplished by assigning a service to a server pool where the database is running. Resources from the pool are used to make sure the required capacity is available.
- Oracle Enterprise Management support for Oracle ASM and Oracle ACFS, Grid Plug and Play, Cluster Resource Management, Oracle Clusterware and Oracle RAC Provisioning and patching.
- SCAN (Single Client Access Name) support as a single name to the clients connecting to Oracle RAC that does not change throughout the life of the cluster, even if you add or remove nodes from the cluster.

The following figure shows Oracle Database with Oracle RAC architecture. This figure shows Oracle Database with Oracle RAC architecture for a partitioned three-node database. An Oracle RAC database is connected to three instances on different nodes. Each instance is associated with a service: HR, Sales, and Call Center. The instances monitor each other by checking "heartbeats." Oracle Net Services provide client access to the Application/web server tier at the top of the figure.

Figure 3-2 Oracle Database with Oracle RAC Architecture



Note:

After Oracle release 11.2, Oracle RAC One Node or Oracle RAC is the preferred solution over Oracle Clusterware (Cold Cluster Failover) because it is a more complete and feature-rich solution.

See Also:

- Oracle RAC Administration and Deployment Guide*
- Oracle Clusterware Administration and Deployment Guide*

Oracle RAC One Node

Oracle Real Application Clusters One Node (Oracle RAC One Node) is a single instance of an Oracle RAC database that runs on one node in a cluster.

This feature enables you to consolidate many databases into one cluster with minimal overhead, protecting them from both planned and unplanned downtime. The consolidated databases reap the high availability benefits of failover protection, online rolling patch application, and rolling upgrades for the operating system and Oracle Clusterware.

Oracle RAC One Node enables better availability than cold failover for single-instance databases because of the Oracle technology called online database relocation, which intelligently migrates database instances and connections to other cluster nodes for high availability and load balancing. Online database relocation is performed using the Server Control Utility (SRVCTL).

Oracle RAC One Node provides the following:

- Always available single-instance database services
- Built-in cluster failover for high availability
- Live migration of instances across servers
- Online rolling patches and rolling upgrades for single-instance databases
- Online upgrade from single-instance to multiple-instance Oracle RAC
- Better consolidation for database servers
- Enhanced server virtualization
- Lower cost development and test platform for full Oracle RAC
- Relocation of Oracle RAC primary and standby databases configured with Data Guard. This functionality is available starting with Oracle Database 11g Release 2 (11.2.0.2).

Oracle RAC One Node also facilitates the consolidation of database storage, standardizes your database environment, and, when necessary, enables you to transition to a full, multiple-instance Oracle RAC database without downtime or disruption.

Oracle Automatic Storage Management

Oracle ASM provides a vertically integrated file system and volume manager directly in the Oracle Database kernel.

This design provides several benefits, 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, Oracle ASM spreads files across all available storage. To protect against data loss, Oracle ASM extends the concept of SAME (stripe and mirror everything) and adds more flexibility because it can mirror at the database file level rather than at the entire disk level.

More important, Oracle ASM simplifies the processes of setting up mirroring, adding disks, and removing disks. Instead of managing hundreds or possibly thousands of files (as in a large data warehouse), database administrators using Oracle 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 administrators time and ensure adherence to standard best practices.

The Oracle ASM native mirroring mechanism (two-way or three-way) protects against storage failures. With Oracle ASM mirroring, you can provide an additional level of data protection 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. After it is defined, an Oracle 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.

By using Oracle ASM, you can:

- Mirror and stripe across drives and storage arrays.
- Automatically remirror from a failed drive to remaining drives.
- Automatically rebalance stored data when disks are added or removed while the database remains online.
- Support Oracle database files and non-database files using Oracle Automatic Storage Management Cluster File System (Oracle ACFS).
- Allow for operational simplicity in managing database storage.
- Manage the Oracle Cluster Registry (OCR) and voting disks.
- Provide preferred read capability on disks that are local to the instance, which gives better performance for an extended cluster.
- Support very large databases.
- Support Oracle ASM rolling upgrades.
- Improve availability and reliability using the Oracle ASM disk scrubbing process to find and repair logical data corruptions using mirror disks.
- Support finer granularity in tuning and security.
- Provide fast repair after a temporary disk failure through Oracle ASM Fast Mirror Resync and automatic repair of block corruptions if a good copy exists in one of the mirrors.
- Provide disaster recovery capability for the file system by enabling replication of Oracle ACFS across the network to a remote site.
- Patch the Oracle ASM instance without impacting the clients that are being serviced using Oracle Flex ASM. A database instance can be directed to access Oracle ASM metadata from another location while the current Oracle ASM instance it is connected to is taken offline for planned maintenance.

- Monitor and manage the speed and status of Oracle ASM Disk Resync and Rebalance operations.
- Bring online multiple disks simultaneously and manage performance better by controlling resync parallelism using the Oracle ASM Resync Power Limit. Recover faster after a cell or disk failure, and the instance doing the resync is failing; this is made possible by using a Disk Resync Checkpoint which enables a resync to resume from where it was interrupted or stopped instead of starting from the beginning.
- Automatically connect database instances to another Oracle ASM instance using Oracle Flex ASM. The local database instance can still access the required metadata and data if an Oracle ASM instance fails due to an unplanned outage.
- Use flex diskgroups to prioritize high availability benefits across multiple databases all using the same diskgroup. Some of the key HA benefits are file extent redundancy, rebalance power limit, and rebalance priority. With flex diskgroups, you can set different values for the above features for different databases, resulting in prioritization across multiple databases within one diskgroup.
- Use flex diskgroups to implement `quoto_groups` across multiple databases sharing one diskgroup which helps in space management and protection.
- Use flex diskgroups to create point-in-time database clones using the ASM split mirror feature.
- Use preferred reads with stretch clusters to improve performance by affinitizing reads to a site.

**See Also:**

Oracle Automatic Storage Management Administrator's Guide

Fast Recovery Area

The fast 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, archived redo log files, control file autobackups, flashback logs, and data file copies are automatically written to a specified file system or Oracle ASM disk group, and the management of this disk space is handled by RMAN and the database server.

Performing a backup to disk is faster because using the fast recovery area eliminates the bottleneck of writing to tape. More important, if database media recovery is required, then data file 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 data files and archived redo log files.

The fast recovery area provides the following benefits:

- 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

 **See Also:**

Oracle Database Backup and Recovery User's Guide

Corruption Prevention, Detection, and Repair

Data block corruptions can be very disruptive and challenging to repair. Corruptions can cause serious application and database downtime and data loss when encountered and worse yet it can go undetected for hours, days and even weeks leading to even longer application downtime once detected. Unfortunately, there is not one way to comprehensively prevent, detect, and repair data corruptions within the database because the source and cause of corruptions can be anywhere in memory, hardware, firmware, storage, operating system, software, or user error. Worse yet, third-party solutions that do not understand Oracle data block semantics and how Oracle changes data blocks do not prevent and detect data block corruptions well. Third party remote mirroring technologies can propagate data corruptions to the database replica (standby) leading to a double failure, data loss, and much longer downtime. Third party backup and restore solutions cannot detect corrupted backups or bad sectors until a restore or validate operation is issued, resulting in longer restore times and once again potential data loss.

Oracle MAA has a comprehensive plan to prevent, detect, and repair all forms of data block corruptions including physical block corruptions, logical block corruptions, stray writes, and lost writes. These additional safeguards provide the most comprehensive Oracle data block corruption prevention, detection, and repair solution. Details of this plan are described in the My Oracle Support note "Best Practices for Corruption Detection, Prevention, and Automatic Repair - in a Data Guard Configuration (Doc ID 1302539.1)."

The following outlines block corruption checks for various manual operational checks and runtime and background corruption checks. Database administrators and the operations team can incorporate manual checks such as running Oracle Recovery Manager (RMAN) backups, RMAN "check logical" validations, or running the `ANALYZE VALIDATE STRUCTURE` command on important objects. Manual checks are especially important to validate data that are rarely updated or queried.

Runtime checks are far superior in that they catch corruptions almost immediately or during runtime for actively queried and updated data. Runtime checks can prevent corruptions or automatically fix corruptions resulting in better data protection and higher application availability. A new background check has been introduced in Exadata to automatically scan and scrub disks intelligently with no application overhead and to automatically fix physically corrupted blocks.

Table 3-1 Summary of Block Corruption Checks

Checks	Capabilities	Physical Block Corruption	Logical Block Corruption
Manual checks	Dbverify, Analyze	Physical block checks	Logical intra-block and inter-object consistency checks
Manual checks	RMAN	Physical block checks during backup and restore operations	Intra-block logical checks
Manual checks	ASM Scrub	Physical block checks	Some logical intra-block checks
Runtime checks	Oracle Active Data Guard	<ol style="list-style-type: none"> 1. Continuous physical block checking at standby during transport and apply 2. Strong database isolation eliminates single point database failure 3. Automatic repair of block corruptions, including file block headers in Oracle Database 12c Release 2 4. Automatic database failover 	<ol style="list-style-type: none"> 1. With DB_LOST_WRITE_PROTECT enabled, detection of lost writes (11.2 and higher). With 11.2.0.4 and Data Guard broker, ability to shutdown the primary when lost writes are detected on the primary database. 2. With DB_BLOCK_CHECKING enabled on the standby, additional intra-block logical checks
Runtime checks	Database	With DB_BLOCK_CHECKSUM, in-memory data block and redo checksum validation	<p>With DB_BLOCK_CHECKING, in-memory intra-block check validation</p> <p>Starting in Oracle Database 18c, and with Shadow Lost Write Protection enabled, Oracle tracks system change numbers (SCNs) for tracked data files and enables early lost write detection. When lost writes are detected, an error is returned immediately.</p> <p>See Shadow Lost Write Protection description following this table.</p>

Table 3-1 (Cont.) Summary of Block Corruption Checks

Checks	Capabilities	Physical Block Corruption	Logical Block Corruption
Runtime checks	ASM and ASM software mirroring (inherent in Exadata, Supercluster, and Zero Data Loss Recovery Appliance)	Implicit data corruption detection for reads and writes and automatic repair if good ASM extent block pair is available during writes	.
Runtime checks	DIX + T10 DIF	Checksum validation from operating system to HBA controller to disk (firmware). Validation for reads and writes for certified Linux, HBA and disks.	.
Runtime checks	Hardware and Storage	Limited checks due to lack of Oracle integration. Checksum is most common.	Limited checks due to lack of Oracle integration. Checksum is most common
Runtime checks	Exadata	Comprehensive HARD checks on writes	HARD checks on writes
Background checks	Exadata	Automatic HARD disk scrub and repair. Detects and fixes bad sectors.	.

Shadow Lost Write Protection

New in Oracle Database 18c, shadow lost write protection detects a lost write before it can result in a major data corruption. You can enable shadow lost write protection for a database, a tablespace, or a data file without requiring an Oracle Data Guard standby database. Shadow lost write protection provides fast detection and immediate response to a lost write, thus minimizing the data loss that can occur in a database due to data corruption.

 **See Also:**

Oracle Database Reference for more information about the views and initialization parameters

My Oracle Support Note [1302539.1](https://support.oracle.com/epos1/docs?id=1302539.1)

Data Recovery Advisor

Data Recovery Advisor automatically diagnoses persistent (on-disk) data failures, presents appropriate repair options, and runs repair operations at your request.

You can use Data Recovery Advisor to troubleshoot primary databases, logical standby databases, physical standby databases, and snapshot standby databases.

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 an easily understood format.
- **Repair generation**

Even if a failure was diagnosed correctly, selecting the correct repair strategy can be error-prone and stressful. Moreover, there is often a high penalty for making poor decisions in terms of increased downtime and loss of data. Data Recovery Advisor automatically determines the best repair for a set of failures and presents it to you.
- **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, including restoring files directly from the primary or standby database to complete the proposed repair.
- **Repair automation**

If you accept the suggested repair option, Data Recovery Advisor automatically performs the repair, 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, and 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 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.

 **Note:**

Data Recovery Advisor only supports single-instance databases. Oracle RAC databases are not supported.

 **See Also:**

Oracle Database Backup and Recovery User's Guide for information about Data Recovery Advisor supported database configurations.

Oracle Flashback Technology

Oracle Flashback technology is a group of Oracle Database features that let you view past states of database, database objects, transactions or rows or to rewind the database, database objects, transactions or rows to a previous state without using point-in-time media recovery.

With flashback features, you can:

- Perform queries to show data as it looked at a previous point in time
- Perform queries that return metadata that shows a detailed history of changes to the database
- Recover tables or rows to a previous point in time
- Automatically track and archive transactional data changes
- Roll back a transaction and its dependent transactions while the database remains online
- Undrop a table
- Recover a database to a point-in-time without a restore operation

Other than the flashback database feature, most Oracle Flashback features use the Automatic Undo Management (AUM) system to obtain metadata and historical data for transactions. They rely on undo data, which are records of the effects of individual transactions. For example, if a user runs an UPDATE statement to change a salary from 1000 to 1100, then Oracle Database stores the value 1000 in the undo data.

Undo data is persistent and survives a database shutdown. By using flashback features, you can use undo data to query past data or recover from logical damage. Besides using it in flashback features, Oracle Database uses undo data to perform these actions:

- Roll back active transactions
- Recover terminated transactions by using database or process recovery
- Provide read consistency for SQL queries

Oracle Flashback can address and rewind data that is compromised due to various human or operator errors that inadvertently or maliciously change data, cause bad installations and upgrades, and result in logical errors in applications. These problems

use features such as flashback transaction, flashback drop, flashback table, and flashback database.

- [Oracle Flashback Query](#)
Oracle Flashback Query (Flashback Query) provides the ability to view data as it existed in the past by using the Automatic Undo Management system to obtain metadata and historical data for transactions.
- [Oracle Flashback Version Query](#)
Oracle Flashback Version Query is an extension to SQL that you can use to retrieve the versions of rows in a given table that existed at a specific time interval.
- [Oracle Flashback Transaction](#)
Oracle Flashback Transaction backs out a transaction and its dependent transactions.
- [Oracle Flashback Transaction Query](#)
Oracle Flashback Transaction Query provides a mechanism to view all of the changes made to the database at the transaction level.
- [Oracle Flashback Table](#)
Oracle Flashback Table recovers a table to a previous point in time.
- [Oracle Flashback Drop](#)
Although there is no easy way to recover dropped tables, indexes, constraints, or triggers, Oracle Flashback Drop provides a safety net when you are dropping objects.
- [Restore Points](#)
When an Oracle Flashback recovery operation is performed on the database, you must determine the point in time—identified by the system change number (SCN) or time stamp—to which you can later flash back the data.
- [Oracle Flashback Database](#)
Oracle Flashback Database is the equivalent of a fast rewind button, quickly returning a database to a previous point in time without requiring a time-consuming restore and roll forward using a backup and archived logs.
- [Flashback Pluggable Database](#)
You can rewind a PDB to a previous SCN. The `FLASHBACK PLUGGABLE DATABASE` command, which is available through SQL or Recovery Manager, is analogous to `FLASHBACK DATABASE` in a non-CDB.
- [Block Media Recovery Using Flashback Logs or Physical Standby Database](#)
After attempting to automatically repair corrupted blocks, block media recovery can optionally retrieve a more recent copy of a data block from the flashback logs to reduce recovery time.
- [Flashback Data Archive](#)
The Flashback Data Archive is stored in a tablespace and contains transactional changes to every record in a table for the duration of the record's lifetime.

 **See Also:**

Oracle Database Development Guide

Performing Flashback and Database Point-in-Time Recovery, Using Flashback Database and Restore Points, and Performing Block Media Recovery in the *Oracle Database Backup and Recovery User's Guide*

Oracle Database PL/SQL Packages and Types Reference

Oracle Database Backup and Recovery Reference

Oracle Flashback Query

Oracle Flashback Query (Flashback Query) provides the ability to view 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 were deleted or updated can be immediately repaired even after they were committed.
- Comparing current data with the corresponding data at some time in the past. For example, by 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. By using Flashback Query, it is possible to retrieve past data directly from the database.
- Applying packaged applications, such as report generation tools, to past data.
- Providing self-service error correction for an application, enabling users to undo and correct their errors.

Oracle Flashback Version Query

Oracle Flashback Version Query is an extension to SQL that you can use to retrieve the versions of rows in a given table that existed at a specific time interval.

Oracle Flashback Version 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.

Oracle Flashback Version Query is a powerful tool that database administrators (database administrators) can use to run analysis to determine the source of problems. Additionally, application developers can use Oracle Flashback Version Query to build customized applications for auditing purposes.

Oracle Flashback Transaction

Oracle Flashback Transaction backs 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 whether the transaction was 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 quickly backing out a specific transaction or set of transactions and their dependent transactions. You use one command to back out transactions while the database remains online.

Oracle Flashback Transaction Query

Oracle Flashback Transaction Query provides a mechanism to view all of the changes made to the database at the transaction level.

When used in conjunction with Oracle Flashback Version Query, it offers a fast and efficient means to recover from a human or application error. Oracle 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.

Oracle Flashback Table

Oracle Flashback Table recovers a table to a previous point in time.

It provides a fast, online solution for recovering a table or set of tables that were changed by a human or application error. In most cases, Oracle Flashback Table alleviates the need for administrators to perform more complicated point-in-time recovery operations. The data in the original table is not lost when you use Oracle Flashback Table because you can return the table to its original state.

Oracle Flashback Drop

Although there is no easy way to recover dropped tables, indexes, constraints, or triggers, Oracle Flashback Drop provides a safety net when you are dropping objects.

When you drop a table, it is automatically placed 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.

Restore Points

When an Oracle Flashback recovery operation is performed on the database, you must determine the point in time—identified by the system change number (SCN) or time stamp—to which you can later flash back the data.

Oracle Flashback restore points are labels that you can define to substitute for the SCN or transaction time used in Flashback Database, Flashback Table, and Oracle Recovery Manager (RMAN) operations. Furthermore, a database can be flashed back through a previous database recovery and opened with an

```
OPEN RESETLOGS
```

command by using guaranteed restore points. Guaranteed restore points allow major database changes—such as database batch jobs, upgrades, or patches—to be quickly undone by ensuring that the undo required to rewind the database is retained.

Using the restore points feature provides the following benefits:

- The ability to quickly restore to a consistent state, to a time before a planned operation that has gone awry (for example, a failed batch job, an Oracle software upgrade, or an application upgrade)
- The ability to resynchronize a snapshot standby database with the primary database
- A quick mechanism to restore a test or cloned database to its original state

Oracle Flashback Database

Oracle Flashback Database is the equivalent of a fast rewind button, quickly returning a database to a previous point in time without requiring a time consuming restore and roll forward using a backup and archived logs.

The larger the size of the database, the greater the advantage of using Oracle Flashback Database for fast point in time recovery.

Enabling Oracle Flashback Database provides the following benefits:

- Fast point in time recovery to repair logical corruptions, such as those caused by administrative error.
- Useful for iterative testing when used with Oracle restore points. A restore point can be set, database changes implemented, and test workload run to assess impact. Oracle Flashback Database can then be used to discard the changes and return the database to the original starting point, different modifications can be

made, and the same test workload run a second time to have a true basis for comparing the impact of the different configuration changes.

- Data Guard uses Oracle Flashback Database to quickly reinitialize a failed primary database as a new standby (after a failover has occurred), without requiring the failed primary to be restored from a backup.
- Flashback database operates at the CDB level or the PDB level.

Flashback Pluggable Database

You can rewind a PDB to a previous SCN. The `FLASHBACK PLUGGABLE DATABASE` command, which is available through SQL or Recovery Manager, is analogous to `FLASHBACK DATABASE` in a non-CDB.

Flashback PDB protects an individual PDB against data corruption, widespread user errors, and redo corruption. The operation does not rewind data in other PDBs in the CDB.

You can use

```
CREATE RESTORE POINT ... FOR PLUGGABLE
      DATABASE
```

to create a PDB restore point, which is only usable within a specified PDB. As with CDB restore points, PDB restore points can be normal or guaranteed. A guaranteed restore point never ages out of the control file and must be explicitly dropped. If you connect to the root, and if you do not specify the

```
FOR PLUGGABLE
      DATABASE
```

clause, then you create a CDB restore point, which is usable by all PDBs.

A special type of PDB restore point is a clean restore point, which you can only create when a PDB is closed. For PDBs with shared undo, rewinding the PDB to a clean restore point is faster than other options because it does not require restoring backups or creating a temporary database instance.

Block Media Recovery Using Flashback Logs or Physical Standby Database

After attempting to automatically repair corrupted blocks, block media recovery can optionally retrieve a more recent copy of a data block from the flashback logs to reduce recovery time.

Automatic block repair allows corrupt blocks on the primary database to be automatically repaired as soon as they are detected, by using good blocks from a physical standby database.

Furthermore, a corrupted block encountered during instance recovery does not result in instance recovery failure. 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. In addition, the RMAN

```
RECOVER BLOCK
```

command restores blocks from a physical standby database, if it is available.

Flashback Data Archive

The Flashback Data Archive 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 a much longer duration than the retention period offered by an undo tablespace, and used to retrieve very old data for analysis and repair.

Oracle Data Pump and Data Transport

Oracle Data Pump technology enables very high-speed movement of data and metadata from one database to another. Data Pump is used to perform the following planned maintenance activities:

- Database migration to a different platform
- Database migration to pluggable databases
- Database upgrade

The Data Pump features that enable the planned maintenance activities listed above are the following:

- Full transportable export/import to move an entire database to a different database instance
- Transportable tablespaces to move a set of tablespaces between databases

See Also:

Transporting Data

Oracle Replication Technologies for Non-Database Files

Oracle ASM Cluster File System, Oracle Database File System, and Oracle Solaris ZFS Storage Appliance Replication are the Oracle replication technologies for non-database files.

Table 3-2 Oracle Replication Technologies for Non-Database Files

Technology	Recommended Usage	Comments
Oracle ASM Cluster File System	Recommended to provide a single-node and cluster-wide file system solution integrated with Oracle ASM, Oracle Clusterware, and Oracle Enterprise Manager technologies. Provides a loosely coupled full stack replication solution when combined with Data Guard or Oracle GoldenGate.	<p>Oracle ACFS establishes and maintains communication with the Oracle ASM instance to participate in Oracle ASM state transitions including Oracle ASM instance and disk group status updates and disk group rebalancing.</p> <p>Supports many database and application files, including executables, database trace files, database alert logs, application reports, BFILEs, and configuration files. Other supported files are video, audio, text, images, engineering drawings, and other general-purpose application file data.</p> <p>Can provide near-time consistency between database changes and file system changes when point-in-time recovery happens</p> <p>Can be exported and accessed by remote clients using standard NAS File Access Protocols such as NFS and CIFS.</p>
Oracle Database File System	Recommended for providing stronger synchronization between database and non-database systems.	<p>Can be integrated with the database to maintain complete consistency between the database changes and the file system changes</p> <p>All data stored in the database and can be used with Oracle Active Data Guard to provide both disaster recovery and read-only access</p> <p>Can take advantage all of the Oracle database features</p>
Oracle Solaris ZFS Storage Appliance Replication	Recommended for disaster recovery protection for non-database files, and specifically for Oracle Fusion Middleware critical files stored outside of the database.	<p>Replicates all non-database objects, including Oracle Fusion Middleware binaries configuration</p> <p>Can provide near time consistency between database changes and file system changes when point-in-time recovery happens</p>

- [Oracle ASM Cluster File System](#)
Oracle ASM Cluster File System (ACFS) is a multiplatform, scalable file system, and storage management technology that extends Oracle Automatic Storage

Management (Oracle ASM) functionality to support customer files maintained outside of Oracle Database.

- [Oracle Database File System](#)

Oracle Database File System (DBFS) takes advantage of the features of the database to store files, and the strengths of the database in efficiently managing relational data, to implement a standard file system interface for files stored in the database.

- [Oracle Solaris ZFS Storage Appliance Replication](#)

The Oracle Solaris ZFS Storage Appliance series supports snapshot-based replication of projects and shares from a source appliance to any number of target appliances manually, on a schedule, or continuously.

Oracle ASM Cluster File System

Oracle ASM Cluster File System (ACFS) is a multiplatform, scalable file system, and storage management technology that extends Oracle Automatic Storage Management (Oracle ASM) functionality to support customer files maintained outside of Oracle Database.

Oracle ACFS supports many database and application files, including executables, database trace files, database alert logs, application reports, BFILEs, and configuration files. Other supported files are video, audio, text, images, engineering drawings, and other general-purpose application file data.

Oracle ACFS takes advantage of the following Oracle ASM functionality:

- Oracle ACFS dynamic file system resizing
- Maximized performance through direct access to Oracle ASM disk group storage
- Balanced distribution of Oracle ACFS across Oracle ASM disk group storage for increased I/O parallelism
- Data reliability through Oracle ASM mirroring protection mechanisms

Oracle ACFS Replication, similar to Data Guard for the database, enables replication of Oracle ACFS file systems across the network to a remote site, providing disaster recovery capability for the file system. Oracle ACFS replication captures file system changes written to disk for a primary file system and records the changes in files called replication logs. These logs are transported to the site hosting the associated standby file system where background processes read the logs and apply the changes recorded in the logs to the standby file system. After the changes recorded in a replication log are successfully applied to the standby file system, the replication log is deleted from the sites hosting the primary and standby file systems.

An additional feature of Oracle ACFS is that it offers snapshot-based replication for generic and application files, providing an HA solution for disaster recovery and Test/Development environments. Oracle Databases stored in ACFS can leverage Oracle Multitenant and ACFS snapshot technologies to create quick and efficient snapshot clones of pluggable databases.

Oracle Data Guard and Oracle ACFS can be combined to provide a full stack high availability solution with Data Guard protecting the database with a standby database and Oracle ACFS replicating the file system changes to the standby host. For planned outages the file system and the database remain consistent to a point in time with zero data loss.

 **See Also:**

[Oracle ACFS ASM Cluster File System: What is it and How to use it](#)

<http://www.oracle.com/goto/maa> for Oracle MAA white paper “Full Stack Role Transition - Oracle ACFS and Oracle Data Guard”

Oracle Database File System

Oracle Database File System (DBFS) takes advantage of the features of the database to store files, and the strengths of the database in efficiently managing relational data, to implement a standard file system interface for files stored in the database.

With this interface, storing files in the database is no longer limited to programs specifically written to use BLOB and CLOB programmatic interfaces. Files in the database can now be transparently accessed using any operating system (OS) program that acts on files. For example, extract, transform, and load (ETL) tools can transparently store staging files in the database.

Oracle DBFS provides the following benefits:

- Full stack integration recovery and failover: By storing file system files in a database structure, it is possible to easily perform point-in-time recovery of both database objects and file system data.
- Disaster Recovery System Return on Investment (ROI): All changes to files contained in DBFS are also logged through the Oracle database redo log stream and thus can be passed to a Data Guard physical standby database. Using Oracle Active Data Guard technology, the DBFS file system can be mounted read-only using the physical standby database as the source. Changes made on the primary are propagated to the standby database and are visible once applied to the standby.
- File system backups: Because DBFS is stored in the database as database objects, standard RMAN backup and recovery functionality can be applied to file system data. Any backup, restore, or recovery operation that can be performed on a database or object within a database can also be performed against the DBFS file system.

 **See Also:**

[Database File System \(DBFS\)](#)

Oracle Solaris ZFS Storage Appliance Replication

The Oracle Solaris ZFS Storage Appliance series supports snapshot-based replication of projects and shares from a source appliance to any number of target appliances manually, on a schedule, or continuously.

The Oracle Solaris ZFS Storage Appliance series supports the following use cases:

- **Disaster recovery:** Replication can be used to mirror an appliance for disaster recovery. In the event of a disaster that impacts the service of the primary appliance (or even an entire data center), administrators activate the service at the disaster recovery site, which takes over using the most recently replicated data. When the primary site is restored, data changed while the disaster recovery site was in service can be migrated back to the primary site, and normal service is restored. Such scenarios are fully testable before a disaster occurs.
- **Data distribution:** Replication can be used to distribute data (such as virtual machine images or media) to remote systems across the world in situations where clients of the target appliance would not ordinarily be able to reach the source appliance directly, or such a setup would have prohibitively high latency. One example uses this scheme for local caching to improve latency of read-only data (such as documents).
- **Disk-to-disk backup:** Replication can be used as a backup solution for environments in which tape backups are not feasible. Tape backup might not be feasible, for example, because the available bandwidth is insufficient or because the latency for recovery is too high.
- **Data migration:** Replication can be used to migrate data and configuration between Oracle Solaris ZFS Storage appliances when upgrading hardware or rebalancing storage. Shadow migration can also be used for this purpose.

The architecture of Oracle Solaris ZFS Storage Appliance also makes it an ideal platform to complement Data Guard for disaster recovery of Oracle Fusion Middleware. Oracle Fusion Middleware has a number of critical files that are stored outside of the database. These binaries, configuration data, metadata, logs and so on also require data protection to ensure availability of the Oracle Fusion Middleware. For these, the built-in replication feature of the ZFS Storage Appliance is used to move this data to a remote disaster recovery site.

Benefits of the Oracle Solaris ZFS Storage Appliance when used with Oracle Fusion Middleware include:

- Leverages remote replication for Oracle Fusion Middleware
- Provides ability to quickly create clones and snapshots of databases to increase ROI of DR sites



See Also:

[Oracle ZFS Storage Appliance Software](#)

Oracle Multitenant

Oracle Multitenant is the optimal database consolidation method. The multitenant architecture combines the best attributes of each of the previous consolidation methods without their accompanying tradeoffs.

Oracle Multitenant helps reduce IT costs by simplifying consolidation, provisioning, upgrades and more. This new architecture allows a container database (CDB) to hold many pluggable databases (PDBs). To applications, these PDBs appear as a standalone database, and no changes are required to the application in order to access the PDB. By consolidating multiple databases as PDBs into a single CDB, you

are provided with the ability to manage "many as one". The flexibility remains to operate on PDBs in isolation should your business require it.

Oracle Multitenant is fully compliant with and takes direct advantage of high availability features such as Oracle Real Application Clusters, Oracle Data Guard, and Oracle GoldenGate, just like any non-container database (non-CDB), meaning it can be used in any of the Oracle MAA reference architectures. Grouping multiple PDBs with the same high availability requirements into the same CDB ensures that all of those PDBs and their applications are managed and protected with the same technologies and configurations.

Benefits of Using Oracle Multitenant

- High consolidation density - Many PDBs can be stored in a single CDB. These PDBs share background processes and memory structures letting you run more PDBs than you would non-CDBs, because the overhead for each non-CDB is removed or reduced. You can store up to 4095 PDBs in a CDB. Each PDB can also have a different character set from other PDBs within the same CDB, as long as the CDB root character set is a superset of all of the PDBs' character sets. Logical standby databases also support this mix of character sets to allow rolling upgrades with a transient logical standby database.
- Online provisioning operations, including clones, refreshable clones, and PDB relocation - A PDB can be unplugged from one CDB and plugged into another. A PDB can also be cloned either into the same CDB or into a different CDB. Cloning can be used to create a "gold image" or seed database for DBaaS or SaaS environments. This PDB can then be rapidly cloned to easily set up database environments for new customers.
 - Near Zero Downtime PDB Relocation – This feature significantly reduces the downtime of relocating a PDB from one CDB to another by using clone functionality. The source PDB remains open and functional while the relocation takes place. The application outage is reduced to a very short window while the source PDB is brought to a consistent state, and the destination PDB is synchronized and brought online. This functionality also takes advantage of another new feature, Listener Redirects, which allows you to keep the same connect descriptor for applications and connect to the destination PDB even after it has been relocated.
 - Online provisioning and cloning – Clones of PDBs can be created without requiring the source PDB to be placed in read only-mode. The source PDB can be left in read-write mode and accessible to applications for the duration of the clone operation.
 - Refreshable Clone PDB – Clones of PDBs can be created in such a way as to be refreshed with changes with changes made to the source PDB applied either automatically at set intervals or manually. For a clone to be refreshable it must remain in read-only mode. The clone can be converted into an ordinary PDB by opening it read-write. Refreshable clones are well suited to be used as test masters for Exadata storage snapshots.
- New patching and upgrade options -When you upgrade or patch a CDB, all of the PDBs in that container are also upgraded or patched. If you need isolation, you can unplug a PDB and plug it into a CDB at a later version.
- Database backup and recovery - By consolidating multiple databases as PDBs, operations such as backup and disaster recovery are performed at the container level. Oracle Multitenant also provides the flexibility to backup and restore individual PDBs with no impact to other running PDBs in the same CDB.

- Operation with Oracle Data Guard - Data Guard configurations are maintained at the CDB level. When a Data Guard role transition (either failover or switchover) is performed, all PDBs are transitioned to the new primary database. There is no need to create or manage multiple Data Guard configurations for each PDB as would be required for single databases. Existing tools such as Data Guard Standby First Patching and Data Guard Transient Logical Rolling Upgrade can still be used to reduce downtime and are performed at the container level, so all PDBs will be maintained in a single operation.
 - PDB Migration with Data Guard Broker – The Data Guard broker has been enhanced to provide automation for migrating PDBs from one CDB, either the primary database or the standby database, to another CDB. This can be used for straight migration of a PDB from one CDB to another running at either at the same version or a CDB running at a higher version to start the upgrade process. This automation can also be used to affect a single PDB failover by using the PDBs files at a standby database to plug into a different CDB at the same version.
 - Subset Standby - A subset standby enables users of Oracle Multitenant to designate a subset of the PDBs in a CDB for replication to a standby database. This provides a finer granularity of designating which standby databases will contain which PDBs.
- Operation with Oracle GoldenGate - All of functionality provided by Oracle GoldenGate also exists for Oracle Multitenant. GoldenGate also provides the flexibility to operate at the PDB level, allowing replication to occur for a subset of the PDBs in a CDB. GoldenGate can be used for minimal to zero downtime upgrades either at the CDB level or at an individual PDB level.
- Resource management - Just as Oracle Resource Manager can control resource utilization between single databases, it can also control resource utilization of individual PDBs in a container. This can ensure that a single PDB does not access more than its assigned share of system resources. You can specify guaranteed minimums and maximums for SGA, buffer cache, shared pool, and PGA memory at the PDB limit.
- Operation with Oracle Flashback Database - If fast point-in-time recovery is required, the initial release of Oracle Multitenant enables using Flashback Database at the CDB level. Oracle Multitenant enables Flashback Database to be used on an individual PDB without impacting the availability of other PDBs. Flashback Database can be performed at the CDB level which will flashback all of the PDBs in the container. Individual PDBs can be flashed back using the Flashback Pluggable Database feature. When flashing back an individual PDB all other PDBs remain unaffected.
- Data Guard Broker PDB Migration or Failover - In multitenant broker configurations, you may need to move a Production PDB from one container database to another container database that resides on the same system. You may also need to failover a PDB from a Data Guard Standby database to a new production container database when the production PDB has failed but the container database and all other PDBs function normally. Using the new Data Guard Broker command, `MIGRATE PLUGGABLE DATABASE`, you can easily move a single PDB from one container database to another, or failover a single PDB from a Data Guard standby to a new production container database. (new in Oracle Database 12c Release 2)

 **See Also:**

Oracle Multitenant Administrator's Guide

[Best Practices for Oracle Database Consolidation: A Guide for Implementation Including MAA Reference Architectures](#)

Oracle Sharding

Oracle Sharding is a scalability and availability feature for applications explicitly designed to run on a sharded database.

Oracle sharding enables distribution and replication of data across a pool of Oracle databases that share no hardware or software. The pool of databases is presented to the application as a single logical database. Applications elastically scale (data, transactions, and users) to any level, on any platform, simply by adding additional databases (shards) to the pool. Scaling up to 1000 shards is supported.

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

 **See Also:**

Using Oracle Sharding

Oracle Restart

Oracle Restart enhances the availability of a single-instance (nonclustered) Oracle database and its components.

Oracle Restart is used in single-instance environments only. For Oracle Real Application Clusters (Oracle RAC) environments, the functionality to automatically restart components is provided by Oracle Clusterware.

If you install Oracle Restart, it automatically restarts the database, the listener, and other Oracle components after a hardware or software failure or whenever the database's host computer restarts. It also ensures that the Oracle components are restarted in the proper order, in accordance with component dependencies.

Oracle Restart periodically monitors the health of components—such as SQL*Plus, the Listener Control utility (LSNRCTL), ASMCMD, and Oracle Data Guard—that are integrated with Oracle Restart. If the health check fails for a component, Oracle Restart shuts down and restarts the component.

Oracle Restart runs out of the Oracle Grid Infrastructure home, which you install separately from Oracle Database homes.

Integrated client failover applications depend on role based services and Fast Application Notification events, managed by Oracle clusterware, to alert the application to failures. Single instance databases must have Oracle Restart to achieve integrated client failover.

 **See Also:**

Oracle Database Administrator's Guide for information about installing and configuring the Oracle Restart feature

Oracle Site Guard

Oracle Site Guard is a disaster-recovery solution that enables administrators to automate complete site switchover or failover.

Oracle Site Guard orchestrates and automates the coordinated failover of Oracle Fusion Middleware, Oracle Fusion Applications, and Oracle Databases. It is also extensible to include other data center software components.

Oracle Site Guard integrates with underlying replication mechanisms that synchronize primary and standby environments and protect mission critical data. It comes with a built-in support for Oracle Data Guard for Oracle database, and Oracle Sun ZFS. Oracle Site Guard can also support other storage replication technologies.

 **See Also:**

Oracle Enterprise Manager Oracle Site Guard Administrator's Guide

Online Reorganization and Redefinition

One way to enhance availability and manageability is to allow 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 user access to the database. This capability improves data availability, query performance, response time, and disk space usage. All of these are important in a mission-critical environment and make the application upgrade process easier, safer, and faster.

Use Oracle Database online maintenance features to significantly reduce (or eliminate) the application downtime required to make changes to an application's database objects

**See Also:**

Redefining Tables Online in *Oracle Database Administrator's Guide*

Zero Data Loss Recovery Appliance

The cloud-scale Zero Data Loss Recovery Appliance, commonly known as Recovery Appliance, is an engineered system designed to dramatically reduce data loss and backup overhead for all Oracle databases in the enterprise.

Integrated with Recovery Manager (RMAN), the Recovery Appliance enables a centralized, incremental-forever backup strategy for large numbers of databases, using cloud-scale, fault-tolerant hardware and storage. The Recovery Appliance continuously validates backups for recoverability.

Recovery Appliance is the MAA-preferred backup and recovery appliance because:

- Elimination of data loss when restoring from Recovery Appliance
- Minimal backup overhead
- Improved end-to-end data protection visibility
- Cloud-scale protection
- Integrates very well with all MAA reference architectures including Oracle Sharding tier

**See Also:**

[Zero Data Loss Recovery Appliance Documentation](#)

Fleet Patching and Provisioning

Fleet Patching and Provisioning maintains a space-efficient repository of software, more precisely "gold images," which are standardized software homes that can be provisioned to any number of target machines.

Any number of homes can be provisioned from a given gold image, and Fleet Patching and Provisioning maintains lineage information so that the provenance of deployed software is always known. Gold images can be organized into series, allowing you to create groupings that track the evolution of a release, with different series for different tailored solutions such as Oracle Database patch bundles for specific applications. A notification system informs interested parties when a new image is available in a given series. Fleet Patching and Provisioning is a feature of Oracle Grid Infrastructure. The components that form the Fleet Patching and Provisioning Server are managed automatically by Oracle Grid Infrastructure.

Fleet Patching and Provisioning can provision databases, clusterware, middleware, and custom software. Fleet Patching and Provisioning offers additional features for creating, configuring, patching and upgrading Oracle Grid Infrastructure and Oracle Database deployments. These capabilities simplify maintenance, reducing its risk and impact, and provide a roll-back option if changes need to be backed out. Additional

capabilities include provisioning clusters and databases onto base machines, and simple capacity on demand by growing and shrinking clusters and Oracle RAC databases. All of these operations are performed with single commands which replace the numerous manual steps otherwise required. All commands and their outcomes are recorded in an audit log. All workflows allow customization to support the unique requirements of any environment.

The key benefits of Fleet Patching and Provisioning are:

- Enables and enforces standardization
- Simplifies provisioning, patching and upgrading
- Minimizes the impact and risk of maintenance
- Increases automation and reduces touch points
- Supports large scale deployments

See Also:

Fleet Patching and Provisioning and Maintenance in *Oracle Clusterware Administration and Deployment Guide*

[Oracle Fleet Patching and Provisioning \(FPP\) Introduction and Technical Overview](#)

Enabling Continuous Service for Applications

Applications achieve continuous service when planned maintenance, unplanned outages, and load imbalances of the database tier are hidden.

A combination of application best practices, simple configuration changes, and an Oracle Database deployed using MAA best practices ensures that your applications are continuously available.

See the following topics for more information about continuous service for application.

- [Continuous Application Service](#)
Oracle provides a set of features that you can choose from to keep your application available during planned events, unplanned outages and load imbalances.
- [Edition-Based Redefinition](#)
Planned application changes may include changes to data, schemas, and programs. The primary objective of these changes is to improve performance, manageability, and functionality. An example is an application upgrade.

Continuous Application Service

Oracle provides a set of features that you can choose from to keep your application available during planned events, unplanned outages and load imbalances.

You can think of these features as an insurance policy protecting your applications from service interruptions. The best features are those that are fully transparent to your application, so that your application developers can focus on building functionality rather than infrastructure, and features that continue to protect the application when it changes in the future.

- Draining and Rebalancing Sessions for Planned Maintenance

When planned maintenance starts, sessions that need to be drained from an instance, PDB, or database are marked to be drained. Idle sessions are released gradually. Active sessions are drained when the work executing in that session completes. Draining of sessions is in wide use with Oracle connection pools and mid-tiers configured for Fast Application Notification (FAN). Starting with Oracle Database 18c, the database itself drains sessions when PDBs and instances are stopped or relocated. Draining is always the best solution for hiding planned maintenance. Failover solutions such as Application Continuity are the fallback when work will not drain in the time allocated

- **Transparent Application Failover**
Transparent Application Failover (TAF) is a feature dating back to Oracle8i. Following an instance failure, TAF creates a new session, and can replay queries back to where they were before the failure occurred. Starting with Oracle Database 12c Release 2, TAF can restore the initial session state before queries are replayed.
- **Application Continuity**
Application Continuity (AC) hides outages, starting with Oracle Database 12c Release 1 for thin Java-based applications, and Oracle Database 12c Release 2 (12.2.0.1) for OCI and ODP.NET based applications. Application Continuity rebuilds the session by recovering the session from a known point which includes session states and transactional states. Application Continuity rebuilds all in-flight work. The application continues as it was, seeing a slightly delayed execution time when a failover occurs. The standard mode for Application Continuity is for OLTP-style pooled applications.
- **Transparent Application Continuity**
Starting with Oracle Database 18c, Transparent Application Continuity (TAC) transparently tracks and records session and transactional state so that the database session can be recovered following recoverable outages. This is done with no reliance on application knowledge or application code changes, allowing Transparent Application Continuity to be enabled for your applications. Application transparency and failover are achieved by consuming the state-tracking information that captures and categorizes the session state usage as the application issues user calls.



See Also:

MAA white paper [MAA for Applications with the Oracle Database](#)

Ensuring Application Continuity in *Oracle Real Application Clusters Administration and Deployment Guide*

Edition-Based Redefinition

Planned application changes may include changes to data, schemas, and programs. The primary objective of these changes is to improve performance, manageability, and functionality. An example is an application upgrade.

Edition-based redefinition (EBR) lets you upgrade the database component of an application while it is in use, thereby minimizing or eliminating downtime. To upgrade an application while it is in use, you must copy the database objects that comprise the database component of the application and redefine the copied objects in isolation.

Your changes do not affect users of the application—they can continue to run the unchanged application. When you are sure that your changes are correct, you make the upgraded application available to all users.



See Also:

Using Edition-Based Redefinition in *Oracle Database Development Guide*

Learn About Other High Availability Features

In addition to the Oracle Database features for high availability listed above, you may be interested in the following resources.

- Overview of High Availability in CNCPT
- The latest high availability feature descriptions
 - [Oracle Database Features App](#)
 - Availability in *Oracle Database New Features Guide*

4

Oracle Database High Availability Solutions for Unplanned Downtime

Oracle Database offers an integrated suite of high availability solutions that increase availability.

These solutions also **eliminate or minimize** both planned and unplanned downtime, and help enterprises maintain business continuity 24 hours a day, 7 days a week. However, Oracle's high availability solutions not only go beyond reducing downtime, but also help to improve overall performance, scalability, and manageability.

- [Outage Types and Oracle High Availability Solutions for Unplanned Downtime](#)
Various Oracle MAA high availability solutions for unplanned downtime are described here in an easy to navigate matrix.
- [Managing Unplanned Outages for MAA Reference Architectures and Multitenant Architectures](#)
High availability solutions in each of the MAA service-level tiers for the MAA reference architectures and multitenant architectures are described in an easy to navigate matrix.

Outage Types and Oracle High Availability Solutions for Unplanned Downtime

Various Oracle MAA high availability solutions for unplanned downtime are described here in an easy to navigate matrix.

The following table shows how the features discussed in the referenced (hyperlinked) sections can be used to address various causes of unplanned downtime. Where several Oracle solutions are listed, the MAA recommended solution is indicated in the Oracle MAA Solution column.

Table 4-1 Outage Types and Oracle High Availability Solutions for Unplanned Downtime

Outage Scope	Oracle MAA Solution	Benefits
Site failures	Oracle Data Guard and Continuous Application Service (MAA recommended)	<ul style="list-style-type: none"> • Integrated client and application failover • Fastest and simplest database replication • Supports all data types • Zero data loss by eliminating propagation delay • Oracle Active Data Guard <ul style="list-style-type: none"> – Supports read-only services and DML on global temporary tables and sequences to off-load more work from the primary – Allows small updates to be redirected to the primary enabling read-mostly reports to be offloaded to standby • Database In-Memory support
	Oracle GoldenGate	<ul style="list-style-type: none"> • Flexible logical replication solution (target is open read/write) • Active-active high availability (with conflict resolution) • Heterogeneous platform and heterogeneous database support • Potential zero downtime with custom application failover
	Recovery Manager, Zero Data Loss Recovery Appliance and Oracle Secure Backup	<ul style="list-style-type: none"> • Fully managed database recovery and integration with Oracle Secure Backup • Recovery Appliance <ul style="list-style-type: none"> – provides end-to-end data protection for backups – reduces data loss for database restores – Non-real-time recovery
Instance or computer failures	Oracle Real Application Clusters and Oracle Clusterware and Continuous Application Service (MAA recommended)	<ul style="list-style-type: none"> • Integrated client and application failover • Automatic recovery of failed nodes and instances • Lowest application brownout with Oracle Real Application Clusters
	Oracle RAC One Node and Continuous Application Service	<ul style="list-style-type: none"> • Integrated client and application failover • Online database relocation migrates connections and instances to another node • Better database availability than traditional cold failover solutions

Table 4-1 (Cont.) Outage Types and Oracle High Availability Solutions for Unplanned Downtime

Outage Scope	Oracle MAA Solution	Benefits
	Oracle Data Guard and Continuous Application Service	<ul style="list-style-type: none"> • Integrated client and application failover • Fastest and simplest database replication • Supports all data types • Zero data loss by eliminating propagation delay • Oracle Active Data Guard <ul style="list-style-type: none"> – Supports read-only services and DML on global temporary tables and sequences to off-load more work from the primary – Allows small updates to be redirected to the primary enabling read-mostly reports to be offloaded to standby • Database In-Memory support
	Oracle GoldenGate	<ul style="list-style-type: none"> • Flexible logical replication solution (target is open read/write) • Active-Active high availability (with conflict resolution) • Heterogeneous platform and heterogeneous database support • Potential zero downtime with custom application failover
Storage failures	Oracle Automatic Storage Management (MAA recommended)	Mirroring and online automatic rebalancing places redundant copies of the data in separate failure groups.
	Oracle Data Guard (MAA recommended)	<ul style="list-style-type: none"> • Integrated client and application failover • Fastest and simplest database replication • Supports all data types • Zero data loss by eliminating propagation delay • Oracle Active Data Guard supports read-only services and DML on global temporary tables and sequences to off-load more work from the primary • Database In-Memory support
	Recovery Manager with Fast Recovery Area, and Zero Data Loss Recovery Appliance (MAA recommended)	Fully managed database recovery and managed disk and tape backups
	Oracle GoldenGate	<ul style="list-style-type: none"> • Flexible logical replication solution (target is open read/write) • Active-active high availability (with conflict resolution) • Heterogeneous platform and heterogeneous database support • Potential zero downtime with custom application failover

Table 4-1 (Cont.) Outage Types and Oracle High Availability Solutions for Unplanned Downtime

Outage Scope	Oracle MAA Solution	Benefits
Data corruption	<p>Corruption Prevention, Detection, and Repair (MAA recommended)</p> <p>Database initialization settings such as DB_BLOCK_CHECKING, DB_BLOCK_CHECKSUM, and DB_LOST_WRITE_PROTECT</p>	Different levels of data and redo block corruption prevention and detection at the database level
Data corruption	<p>Oracle Data Guard (MAA recommended)</p> <p>Oracle Active Data Guard Automatic Block Repair</p> <p>DB_LOST_WRITE_PROTECT initialization parameter</p>	<ul style="list-style-type: none"> • In a Data Guard configuration with an Oracle Active Data Guard standby <ul style="list-style-type: none"> – Physical block corruptions detected by Oracle at a primary database are automatically repaired using a good copy of the block retrieved from the standby, and vice versa – The repair is transparent to the user and application, and data corruptions can definitely be isolated • With MAA recommended initialization settings, Oracle Active Data Guard and Oracle Exadata Database Machine, achieve most comprehensive full stack corruption protection. • With DB_LOST_WRITE_PROTECT enabled <ul style="list-style-type: none"> – A lost write that occurred on the primary database is detected either by the physical standby database or during media recovery of the primary database, recovery is stopped to preserve the consistency of the database – Failing over to the standby database using Data Guard will result in some data loss – Data Guard Broker's PrimaryLostWrite property supports SHUTDOWN and CONTINUE, plus FAILOVER and FORCEFAILOVER options, when lost writes are detected on the primary database. See <i>Oracle Data Guard Broker</i> – DB_LOST_WRITE_PROTECT initialization parameter provides lost write detection • Shadow lost write protection detects a lost write before it can result in major data corruption. You can enable shadow lost write protection for a database, a tablespace, or a data file without requiring an Oracle Data Guard standby database. Note the impact on your workload may vary.

Table 4-1 (Cont.) Outage Types and Oracle High Availability Solutions for Unplanned Downtime

Outage Scope	Oracle MAA Solution	Benefits
	Dbverify, Analyze, Data Recovery Advisor and Recovery Manager , Zero Data Loss Recovery Appliance , and ASM Scrub with Fast Recovery Area (MAA recommended)	<p>These tools allow the administrator to execute manual checks to help detect and potentially repair from various data corruptions.</p> <ul style="list-style-type: none"> • Dbverify and Analyze conduct physical block and logical intra-block checks. Analyze can conduct inter-object consistency checks. • Data Recovery Advisor automatically detects data corruptions and recommends the best recovery plan. • RMAN operations can <ul style="list-style-type: none"> – Conduct both physical and inter-block logical checks – Run online block-media recovery using flashback logs, backups, or the standby database to help recover from physical block corruptions • Recovery Appliance <ul style="list-style-type: none"> – Does periodic backup validation that helps ensure that your backups are valid – Allows you to input your recovery window requirements, and alerts you when those SLAs cannot be met with your existing backups managed by Recovery Appliance • ASM Scrub detects and attempts to repair physical and logical data corruptions with the ASM pair in normal and high redundancy disks groups.
Data corruption	<p>Oracle Exadata Database Machine and Oracle Automatic Storage Management (MAA recommended)</p> <p>DIX + T10 DIF Extensions (MAA recommended where applicable)</p> <p>Oracle GoldenGate</p>	<ul style="list-style-type: none"> • If Oracle ASM detects a corruption and has a good mirror, ASM returns the good block and repairs the corruption during a subsequent write I/O. • Exadata provides implicit HARD enabled checks to prevent data corruptions caused by bad or misdirected storage I/O. • Exadata provides automatic HARD disk scrub and repair. Detects and fixes bad sectors. • DIX +T10 DIF Extensions provides end to end data integrity for reads and writes through a checksum validation from a vendor's host adapter to the storage device • Flexible logical replication solution (target is open read/write). Logical replica can be used as a failover target if partner replica is corrupted. • Active-active high availability (with conflict resolution) • Heterogeneous platform and heterogeneous database support
Human errors	Oracle security features (MAA recommended)	Restrict access to prevent human errors

Table 4-1 (Cont.) Outage Types and Oracle High Availability Solutions for Unplanned Downtime

Outage Scope	Oracle MAA Solution	Benefits
Delays or slow downs	Oracle Flashback Technology (MAA recommended)	<ul style="list-style-type: none"> Fine-grained error investigation of incorrect results Fine-grained and database-wide or pluggable database rewind and recovery capabilities
	Oracle Database and Oracle Enterprise Manager Oracle Data Guard (MAA recommended) and Continuous Application Service	<ul style="list-style-type: none"> Oracle Database automatically monitors for instance and database delays or cluster slow downs and attempts to remove blocking processes or instances to prevent prolonged delays or unnecessary node evictions. Oracle Enterprise Manager or a customized application heartbeat 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. Then, after a certain threshold expires, Enterprise Manager can call the Data Guard <p>DBMS_DG.INITIATE_FS_FAILOVER</p> <p>PL/SQL procedure to initiate a failover. See the section about "Managing Fast-Start Failover" in <i>Oracle Data Guard Broker</i>.</p> <ul style="list-style-type: none"> Database In-Memory support
File system data	Oracle Replication Technologies for Non-Database Files	Enables full stack failover that includes non-database files

Managing Unplanned Outages for MAA Reference Architectures and Multitenant Architectures

High availability solutions in each of the MAA service-level tiers for the MAA reference architectures and multitenant architectures are described in an easy to navigate matrix.

If you are managing many databases in DBaaS, we recommend using the MAA tiers and Oracle Multitenant as described in [Oracle MAA Reference Architectures](#).

The following table identifies various unplanned outages that can impact a database in a multitenant architecture. It also identifies the Oracle high availability solution to address that outage that is available in each of the MAA reference architectures.

Table 4-2 Unplanned Outage Matrix for MAA Reference Architectures and Multitenant Architectures

Event	Solutions by MAA Architecture	Recovery Window (RTO)	Data Loss (RPO)
Instance Failure	BRONZE: Oracle Restart	Minutes if instance can restart	Zero
	SILVER: Oracle RAC (see Oracle Real Application Clusters and Oracle Clusterware) or Oracle RAC One Node, and Continuous Application Service	Seconds with Oracle RAC, minutes with Oracle RAC One Node	Zero
	GOLD: Oracle RAC (see Oracle Real Application Clusters and Oracle Clusterware and Continuous Application Service)	Seconds	Zero
	PLATINUM: Oracle RAC (see Oracle Real Application Clusters and Oracle Clusterware) and Continuous Application Service	Zero Application Outage	Zero
Permanent Node Failure (but storage available)	BRONZE: Restore and recover	Hours to Day	Zero
	SILVER: Oracle RAC (see Oracle Real Application Clusters and Oracle Clusterware) and Continuous Application Service	Seconds	Zero
	SILVER: Oracle RAC One Node and Continuous Application Service	Minutes	Zero
	GOLD: Oracle RAC (see Oracle Real Application Clusters and Oracle Clusterware) and Continuous Application Service	Seconds	Zero

Table 4-2 (Cont.) Unplanned Outage Matrix for MAA Reference Architectures and Multitenant Architectures

Event	Solutions by MAA Architecture	Recovery Window (RTO)	Data Loss (RPO)
	PLATINUM: Oracle RAC (see Oracle Real Application Clusters and Oracle Clusterware) and Continuous Application Service	Seconds	Zero
Storage Failure	ALL: Oracle Automatic Storage Management	Zero downtime	Zero
Data corruptions	BRONZE/SILVER: Basic protection Some corruptions require recover restore and recovery of pluggable database (PDB), entire multitenant container database (CDB) or non-container database (non-CDB)	Hour to Days	<ul style="list-style-type: none"> • Since last backup if unrecoverable • Zero or Near Zero with Recovery Appliance
	GOLD: Comprehensive corruption protection and Auto Block Repair with Oracle Active Data Guard	<ul style="list-style-type: none"> • Zero with auto block repair • Seconds to minutes if corruption due to lost writes and using Data Guard Fast Start failover. 	Zero unless corruption due to lost writes
	PLATINUM: Comprehensive corruption protection and Auto Block Repair with Oracle Active Data Guard Oracle GoldenGate replica with custom application failover	<ul style="list-style-type: none"> • Zero with auto block repair • Zero with Oracle GoldenGate replica 	Zero when using Active Data Guard Fast-Start Failover and Oracle GoldenGate

Table 4-2 (Cont.) Unplanned Outage Matrix for MAA Reference Architectures and Multitenant Architectures

Event	Solutions by MAA Architecture	Recovery Window (RTO)	Data Loss (RPO)
Human error	ALL: Logical failures resolved by flashback drop, flashback table, flashback transaction, flashback query flashback pluggable database, and undo.	Dependent on detection time but isolated to PDB and applications using those objects.	Dependent on logical failure
	All: Comprehensive logical failures impacting an entire database and PDB that requires RMAN point in time recovery (PDB) or flashback pluggable database	Dependent on detection time	Dependent on logical failure
Database unusable, system, site or storage failures, wide spread corruptions or disasters	BRONZE/SILVER: Restore and recover	Hours to Days	<ul style="list-style-type: none"> • Since last database and archive backup • Zero or near zero with Recovery Appliance
	GOLD: Active Data Guard Fast-Start Failover and Continuous Application Service	Seconds	Zero to Near Zero
	PLATINUM: Oracle GoldenGate replica with custom application failover	Zero	Zero when using Active Data Guard Fast-Start Failover and Oracle GoldenGate
Performance Degradation	ALL: Oracle Enterprise Manager for monitoring and detection, Database Resource Management for Resource Limits and ongoing Performance Tuning	No downtime but degraded service	Zero

5

Oracle Database High Availability Solutions for Planned Downtime

Planned downtime can be just as disruptive to operations as unplanned downtime. This is especially true for global enterprises that must support users in multiple time zones, or for those that must provide Internet access to customers 24 hours a day, 7 days a week.

See the following topics to learn about keeping your database highly available during planned downtime.

- [Oracle High Availability Solutions for Planned Maintenance](#)
Oracle provides high availability solutions for all planned maintenance.
- [High Availability Solutions for Migration](#)
Oracle MAA recommends several solutions for reducing downtime due to database migration.

Oracle High Availability Solutions for Planned Maintenance

Oracle provides high availability solutions for all planned maintenance.

The following table describes the various Oracle high availability solutions and their projected downtime for various maintenance activities.

Table 5-1 Oracle High Availability Solutions for Planned Maintenance

Maintenance Event	High Availability Solutions with Target Outage Time
Dynamic and Online Resource Provisioning, or Online reorganization and redefinition	Zero application and database downtime for <ul style="list-style-type: none">• Changing initialization parameters dynamically• Renaming and relocating datafiles online• Automatic memory management tuning• Online reorganization and redefinition (managing tables and managing indexes) See the Oracle Database Administrator Guide, Oracle Database Reference Guide (to evaluate which parameters on dynamic), and Online Data Reorganization and Redefinition
Operating system software or hardware updates and patches	Zero database downtime with Oracle RAC and Oracle RAC One Node Rolling or Fleet Patching and Provisioning Zero application downtime with Continuous Availability - Application Checklist for Continuous Service for MAA Solutions Seconds to minutes database downtime with Standby-First Patch Apply and subsequent Data Guard Switchover

Table 5-1 (Cont.) Oracle High Availability Solutions for Planned Maintenance

Maintenance Event	High Availability Solutions with Target Outage Time
Oracle interim or diagnostic software updates or patches	<p>Zero downtime with Online Patching</p> <p>Zero database downtime with Oracle RAC and Oracle RAC One Node.</p> <p>Zero application downtime with Continuous Availability - Application Checklist for Continuous Service for MAA Solutions</p>
Oracle Database or Grid Infrastructure quarterly updates under the Critical Patch Update (CPU) program, or Oracle Grid Infrastructure release upgrades	<p>Zero database downtime with Oracle RAC and Oracle RAC One Node Rolling.</p> <p>Zero application downtime with Continuous Availability - Application Checklist for Continuous Service for MAA Solutions</p> <p>Seconds to minutes downtime with Standby-First Patch Apply and subsequent Data Guard Switchover</p>
Oracle Database Release Upgrade (for example, Oracle Database 11g to 12.2 or 12.2 to 19c)	<p>Seconds to minutes downtime with Data Guard transient logical or DBMS_ROLLING solution</p> <p>Zero downtime with Oracle GoldenGate</p> <p>See Automated Database Upgrades using Oracle Active Data Guard and DBMS_ROLLING for 12.2 and higher database releases or Database Rolling Upgrade using Data Guard for older releases.</p>
Exadata storage or Exadata switch software updates	<p>Zero downtime using Exadata</p> <p style="text-align: center;">patchmgr</p> <p>See Maintenance Guide for Exadata Database Machine</p>
Database Server or Oracle RAC cluster changes (add node, drop node, adjust CPU or memory size of the database server)	<p>Some hardware changes like adjusting CPU can be done online without restarting the database server. Refer to the hardware specific documentation.</p> <p>If the change is not online, then</p> <p>Zero database downtime with Oracle RAC and Oracle RAC One Node Rolling.</p> <p>Zero application downtime with Continuous Availability - Application Checklist for Continuous Service for MAA Solutions</p> <p>Seconds to minutes downtime with Standby-First Patch Apply and subsequent Data Guard Switchover</p>
Application upgrades	<p>Zero downtime with Edition Based Redefinition</p> <p>Zero downtime with Oracle GoldenGate</p> <p>See Edition Based Redefinition and Oracle GoldenGate documentation</p>

High Availability Solutions for Migration

Oracle MAA recommends several solutions for reducing downtime due to database migration.

The following table describes the high availability solutions for migration at a high level.

Table 5-2 High Availability Solutions for Migration

Maintenance Event	High Availability Solutions with Target Outage Time
Migrate the database to a different server or platform	<p>Seconds to minutes downtime with Oracle Data Guard for certain platform combinations</p> <p>Zero downtime with Oracle GoldenGate</p> <p>Data Guard always supports primary and standby combinations on the same platform. For heterogeneous platforms, Refer to Data Guard Support for Heterogeneous Primary and Physical Standbys in Same Data Guard Configuration (Doc ID 413484.1)</p>
Migrate database to an incompatible character set	<p>Zero downtime with Oracle GoldenGate</p> <p>See Character Set Migration</p>
Migrate to pluggable databases to another container database	<p>Seconds to minutes downtime with Pluggable Database Relocate (PDB Relocate)</p> <p>See Relocating a PDB</p>
Migrate to new storage	<p>Zero Downtime with Oracle Automatic Storage Management if storage is compatible</p> <p>with Oracle Data Guard for certain platform combinations</p> <p>Zero Downtime with Oracle GoldenGate</p>
Migrate database from a single-instance system to an Oracle RAC cluster	<p>Zero Downtime with Oracle RAC when applicable. See Adding Oracle RAC to Nodes with Oracle Clusterware Installed</p> <p>Seconds to minutes downtime with Oracle Data Guard for certain platform combinations</p> <p>Zero Downtime with Oracle GoldenGate</p>

6

Operational Prerequisites to Maximizing Availability

Use the following operational best practices to provide a successful MAA implementation.

- [Understand Availability and Performance SLAs](#)
- [Implement and Validate a High Availability Architecture That Meets Your SLAs](#)
- [Establish Test Practices and Environment](#)
- [Set Up and Use Security Best Practices](#)
- [Establish Change Control Procedures](#)
- [Apply Recommended Patches and Software Periodically](#)
- [Execute Disaster Recovery Validation](#)
Disaster recovery validation is required to ensure that you meet your disaster recovery service level requirements such as RTO and RPO.
- [Establish Escalation Management Procedures](#)
- [Configure Monitoring and Service Request Infrastructure for High Availability](#)
To maintain your High Availability environment, you should configure the monitoring infrastructure that can detect and react to performance and high availability related thresholds before any downtime has occurred.
- [Check the Latest MAA Best Practices](#)
The MAA solution encompasses the full stack of Oracle technologies, so you can find MAA best practices for Oracle Fusion Middleware, Oracle Fusion Applications, Oracle Applications Unlimited, Oracle Exalytics, Oracle Exalogic, Oracle VM, and Oracle Enterprise Manager Cloud Control on the MAA pages.

Understand Availability and Performance SLAs

Understand and document your high availability and performance service-level agreements (SLAs):

- Understand the attributes of High Availability and various causes of downtime as described in [Overview of High Availability](#).
- Get agreement from line of business, upper management, and technical teams on HA and performance service level agreements as described in [High Availability Requirements](#), and [A Methodology for Documenting High Availability Requirements](#).

Implement and Validate a High Availability Architecture That Meets Your SLAs

When you have agreement on your high availability and performance service level requirements:

- **Map** the requirements to one of the Oracle MAA standard and validated MAA reference architectures, as described in [High Availability and Data Protection – Getting From Requirements to Architecture](#)
- **Evaluate** the outage and planned maintenance matrices relevant to your referenced architecture in [Oracle Database High Availability Solutions for Unplanned Downtime](#) and [Oracle Database High Availability Solutions for Planned Downtime](#)
- **Learn** about the database features required to implement your MAA architecture in [High Availability Architectures](#)

Establish Test Practices and Environment

You must **validate** or **automate** the following to ensure that your target high availability SLAs are met:

- All software update and upgrade maintenance events
- All repair operations, including those for various types of unplanned outages
- Backup, restore, and recovery operations

If you use Oracle Data Guard for disaster recovery and data protection, Oracle recommends that you:

- Perform periodic switchover operations, or conduct full application and database failover tests
- Validate end-to-end role transition procedures by performing application and Data Guard switchovers periodically

A good test environment and proper test practices are essential prerequisites to achieving the highest stability and availability in your production environment. By validating every change in your test environment thoroughly, you can proactively detect, prevent, and avoid problems before applying the same change on your production systems.

These practices involve the following:

- [Configuring the Test System and QA Environments](#)
- [Performing Preproduction Validation Steps](#)

Configuring the Test System and QA Environments

The test system should be a replica of the production MAA environment (for example, using the MAA Gold reference architecture.) There will be trade offs if the test system is not identical to the MAA service-level driven standard reference architecture that you plan to implement. It's recommended that you execute functional, performance, and availability tests with a workload that mimics production. Evaluate if availability

and performance SLAs are maintained after each change, and ensure that clear fallback or repair procedures are in place if things go awry, while applying the change on the production environment.

With a properly configured test system, many problems can be avoided, because changes are validated with an equivalent production and standby database configuration containing a full data set and using a workload framework to mimic production (for example, using Oracle Real Application Testing.)

Do not try to reduce costs by eliminating the test system, because that decision ultimately affects the stability and the availability of your production applications. Using only a subset of system resources for testing and QA has the tradeoffs shown in the following table, which is an example of the MAA Gold reference architecture.

Table 6-1 Tradeoffs for Different Test and QA Environments

Test Environment	Benefits and Tradeoffs
Full Replica of Production and Standby Systems	<p>Validate:</p> <ul style="list-style-type: none"> • All software updates and upgrades • All functional tests • Full performance at production scale • Full high availability
Full Replica of Production Systems	<p>Validate:</p> <ul style="list-style-type: none"> • All software updates and upgrades • All functional tests • Full performance at production scale • Full high availability minus the standby system <p>Cannot Validate:</p> <ul style="list-style-type: none"> • Functional tests, performance at scale, high availability, and disaster recovery on standby database
Standby System	<p>Validate:</p> <ul style="list-style-type: none"> • Most software update changes • All functional tests • Full performance--if using Data Guard Snapshot Standby, but this can extend recovery time if a failover is required • Role transition • Resource management and scheduling--required if standby and test databases exist on the same system
Shared System Resource	<p>Validate:</p> <ul style="list-style-type: none"> • Most software update changes • All functional tests <p>This environment may be suitable for performance testing if enough system resources can be allocated to mimic production. Typically, however, the environment includes a subset of production system resources, compromising performance validation. Resource management and scheduling is required.</p>
Smaller or Subset of the system resources	<p>Validate:</p> <ul style="list-style-type: none"> • All software update changes • All functional tests • Limited full-scale high availability evaluations <p>Cannot Validate:</p> <ul style="list-style-type: none"> • Performance testing at production scale

Table 6-1 (Cont.) Tradeoffs for Different Test and QA Environments

Test Environment	Benefits and Tradeoffs
Different hardware or platform system resources but same operating system	Validate: <ul style="list-style-type: none"> • Some software update changes • Limited firmware patching test • All functional tests unless limited by new hardware features • Limited production scale performance tests • Limited full-scale high availability evaluations



See Also:

Oracle Database Testing Guide

Performing Preproduction Validation Steps

Pre-production validation and testing of hardware, software, database, application or any changes is an important way to maintain stability. The high-level pre-production validation steps are:

1. Review the patch or upgrade documentation or any document relevant to that change. Evaluate the possibility of performing a rolling upgrade if your SLAs require zero or minimal downtime. Evaluate any rolling upgrade opportunities to minimize or eliminate planned downtime. Evaluate whether the patch or the change qualifies for Standby-First Patching.

Note:

Standby-First Patch enables you to apply a patch initially to a physical standby database while the primary database remains at the previous software release (this applies to certain types of software updates and does not apply to major release upgrades; use the Data Guard transient logical standby and DBMS_ROLLING method for patch sets and major releases). Once you are satisfied with the change, then perform a switchover to the standby database. The fallback is to switchback if required. Alternatively, you can proceed to the following step and apply the change to your production environment. For more information, see "Oracle Patch Assurance - Data Guard Standby-First Patch Apply" in My Oracle Support Note 1265700.1 at <https://support.oracle.com/CSP/main/article?cmd=show&type=NOT&id=1265700.1>

2. Validate the application in a test environment and ensure the change meets or exceeds your functionality, performance, and availability requirements. Automate the procedure and be sure to also document and test a fallback procedure. This requires comparing metrics captured before and after patch application on the test and against metrics captured on the production system. Real Application Testing may be used to capture the workload on the production system and replay it on the test system. AWR and SQL Performance Analyzer may be used to assess performance improvement or regression resulting from the patch.

Validate the new software on a test system that mimics your production environment, and ensure the change meets or exceeds your functionality,

performance, and availability requirements. Automate the patch or upgrade procedure and ensure fallback. Being thorough during this step eliminates most critical issues during and after the patch or upgrade.

3. Use Oracle Real Application Testing and test data management features to comprehensively validate your application while also complying with any security restrictions your line of business may have. Oracle Real Application Testing (a separate database option) enables you to perform real-world testing of Oracle Database. By capturing production workloads and assessing the impact of system changes on these workloads before production deployment, Oracle Real Application Testing minimizes the risk of instabilities associated with system changes. SQL Performance Analyzer and Database Replay are key components of Oracle Real Application Testing. Depending on the nature and impact of the system change being tested, and on the type of system on which the test will be performed, you can use either or both components to perform your testing.

When performing real-world testing there is a risk of exposing sensitive data to non-production users in a test environment. The test data management features of Oracle Database help to minimize this risk by enabling you to perform data masking and data subsetting on the test data.

4. If applicable, perform final pre-production validation of all changes on a Data Guard standby database before applying them to production. Apply the change in a Data Guard environment, if applicable.
5. Apply the change in your production environment.

 **See Also:**

[Data Guard Redo Apply and Standby-First Patching and Data Guard Transient Logical Rolling Upgrades](#)

[Converting a Physical Standby Database into a Snapshot Standby Database and Performing a Rolling Upgrade With an Existing Physical Standby Database](#) in Oracle Data Guard Concepts and Administration

[Oracle Database Rolling Upgrades: Using a Data Guard Physical Standby Database](http://www.oracle.com/goto/maa) on <http://www.oracle.com/goto/maa>

[Oracle Patch Assurance - Data Guard Standby-First Patch Apply \(Doc ID 1265700.1\)](#)

Set Up and Use Security Best Practices

Corporate data can be at grave risk if placed on a system or database that does not have proper security measures in place. A well-defined security policy can help protect your systems from unwanted access and protect sensitive corporate information from sabotage. Proper data protection reduces the chance of outages due to security breaches.



See Also:

Oracle Database Security Guide.

Establish Change Control Procedures

Institute procedures that manage and control changes as a way to maintain the stability of the system and to ensure that no changes are incorporated in the primary database unless they have been rigorously evaluated on your test systems, or any one of the base architectures in the MAA service-level tiers.

Review the changes and get feedback and approval from your change management team.

Apply Recommended Patches and Software Periodically

By periodically testing and applying the latest recommended patches and software versions, you ensure that your system has the latest security and software fixes required to maintain stability and avoid many known issues. Remember to validate all updates and changes on a test system before performing the upgrade on the production system.

Furthermore, Oracle health check tools such as

`orachk`

(supporting Non-Engineered Systems and Oracle Database Appliance) and

`exachk`

(supporting Engineered Systems such as Oracle Exadata Database Machine, Exalogic, Zero Data Loss Recovery Appliance, and Big Data Appliance) provide Oracle software upgrade advice, critical software update recommendations, and patching and upgrading pre-checks, along with its system and database health checks and MAA recommendations.

 **See Also:**

"Oracle Recommended Patches -- Oracle Database" in My Oracle Support Note 756671.1 at <https://support.oracle.com/CSP/main/article?cmd=show&type=NOT&id=756671.1>

"Exadata Database Machine and Exadata Storage Server Supported Versions" in My Oracle Support Note 888828.1 at <https://support.oracle.com/CSP/main/article?cmd=show&type=NOT&id=888828.1>

"ORAchk - Health Checks for the Oracle Stack" in My Oracle Support Note 1268927.2 at <https://support.oracle.com/CSP/main/article?cmd=show&type=NOT&id=1268927.2>

"Oracle Exadata Database Machine exachk or HealthCheck" in My Oracle Support Note 1070954.1 at <https://support.oracle.com/CSP/main/article?cmd=show&type=NOT&id=1070954.1>

Execute Disaster Recovery Validation

Disaster recovery validation is required to ensure that you meet your disaster recovery service level requirements such as RTO and RPO.

Whether you have a standby database, Oracle GoldenGate replica, or leverage database backups from Zero Data Loss Recovery Appliance (Recovery Appliance), ZFS Storage, or another third party, it is important to ensure that the operations and database administration teams are well prepared to failover or restore the database and application any time the primary database is down or underperforming. The concerned teams should be able to detect and decide to failover or restore as required. Such efficient execution during disasters will significantly reduce overall downtime.

If you use Data Guard or Oracle GoldenGate for high availability, disaster recovery, and data protection, Oracle recommends that you perform regular application and database switchover operations **every three to six months**, or conduct full application and database failover tests.

Periodic RMAN cross checks, RMAN backup validations, and complete database restore and recovery are required to validate your disaster recovery solution through backups. Inherent backup checks and validations are done automatically with the Recovery Appliance, but periodic restore and recovery tests are still recommended.

Establish Escalation Management Procedures

Establish escalation management procedures so repair is not hindered. Most repair solutions, when conducted properly are automatic and transparent with the MAA solution. The challenges occur when the primary database or system is not meeting availability or performance SLAs and failover procedures are not automatic as in the case with some Data Guard failover scenarios. Downtime can be prolonged if proper escalation policies are not followed and decisions are not made quickly.

If availability is the top priority, perform repair and failover operations first and then proceed with gathering logs and information for Root Cause Analysis (RCA) after

the application service has been reestablished. For simple data gathering, use the Trace File Analyzer Collector (TFA).

 **See Also:**

MAA web page at <http://www.oracle.com/goto/maa>

My Oracle Support note 1513912.2 “TFA Collector - Tool for Enhanced Diagnostic Gathering” at [1513912.2](#)

Configure Monitoring and Service Request Infrastructure for High Availability

To maintain your High Availability environment, you should configure the monitoring infrastructure that can detect and react to performance and high availability related thresholds before any downtime has occurred.

Also, where available, Oracle can detect failures, dispatch field engineers, and replace failed hardware components such as disks, flash cards, fans, or power supplies without customer involvement.

- [Run Database Health Checks Periodically](#)
Oracle Database health checks are designed to evaluate your hardware and software configuration and MAA compliance to best practices.
- [Configure Oracle Enterprise Manager Monitoring Infrastructure for High Availability](#)
You should configure and use Enterprise Manager and the monitoring infrastructure that detects and reacts to performance and high availability related thresholds to avoid potential downtime.
- [Configure Automatic Service Request Infrastructure](#)
In addition to monitoring infrastructure with Enterprise Manager in the Oracle high availability environment where available, Oracle can detect failures, dispatch field engineers, and replace failing hardware without customer involvement.

Run Database Health Checks Periodically

Oracle Database health checks are designed to evaluate your hardware and software configuration and MAA compliance to best practices.

All of the Oracle health check tools will evaluate Oracle Grid Infrastructure, Oracle Database, and provide an automated MAA scorecard or review that highlights when key architectural and configuration settings are not enabled for tolerance of failures or fast recovery. For Oracle's engineered systems such as Exadata Database Machine, there may be hundreds of additional software, fault and configuration checks.

Oracle recommends periodically (for example, monthly for Exadata Database Machine) downloading the latest database health check, executing the health check, and addressing the key FAILURES, WARNINGS, and INFO messages. Use

```
exachk
```


for Engineered Systems such as Oracle Exadata Database Machine, Exalogic, Zero Data Loss Recovery Appliance, and Big Data Appliance, and use

orachk

for Non-Engineered Systems and Oracle Database Appliance.

Furthermore, it is recommended that you run the health check prior to and after any planned maintenance activity.

You must **evaluate**:

- Existing or new critical health check alerts prior to planned maintenance window
- Adding any new recommendations to the planned maintenance window after testing
- Existing software or critical software recommendations

 **See Also:**

My Oracle Support Note 1268927.2 "ORAchk - Health Checks for the Oracle Stack" at <https://support.oracle.com/CSP/main/article?cmd=show&type=NOT&id=1268927.2>

My Oracle Support Note 1070954.1 "Oracle Exadata Database Machine exachk or HealthCheck" at <https://support.oracle.com/CSP/main/article?cmd=show&type=NOT&id=1070954.1>

Configure Oracle Enterprise Manager Monitoring Infrastructure for High Availability

You should configure and use Enterprise Manager and the monitoring infrastructure that detects and reacts to performance and high availability related thresholds to avoid potential downtime.

The monitoring infrastructure assists you with monitoring for High Availability and enables you to do the following:

- Monitor system, network, application, database and storage statistics
- Monitor performance and service statistics
- Create performance and high availability thresholds as early warning indicators of system or application problems
- Provide performance and availability advice
- Established alerts and tools and database performance
- Receive alerts for engineered systems hardware faults

 **See Also:**

Oracle Enterprise Manager Cloud Control Advanced Installation and Configuration Guide for information about detecting and reacting to potential problems and failures

MAA Best Practices for Enterprise Manager at <http://www.oracle.com/goto/maa>

Configure Automatic Service Request Infrastructure

In addition to monitoring infrastructure with Enterprise Manager in the Oracle high availability environment where available, Oracle can detect failures, dispatch field engineers, and replace failing hardware without customer involvement.

For example, Oracle Automatic Service Request (ASR) is a secure, scalable, customer-installable software solution available as a feature. The software resolves problems faster by using auto-case generation for Oracle's Solaris server and storage systems when specific hardware faults occur.

 **See Also:**

See "Oracle Automatic Service Request" in My Oracle Support Note 1185493.1 at <https://support.oracle.com/CSP/main/article?cmd=show&type=NOT&id=1185493.1>

Check the Latest MAA Best Practices

The MAA solution encompasses the full stack of Oracle technologies, so you can find MAA best practices for Oracle Fusion Middleware, Oracle Fusion Applications, Oracle Applications Unlimited, Oracle Exalytics, Oracle Exalogic, Oracle VM, and Oracle Enterprise Manager Cloud Control on the MAA pages.

MAA solutions and best practices continue to be developed and published on <http://www.oracle.com/goto/maa>.

Index

A

Active Data Guard
See Oracle Active Data Guard

analysis
determining high availability requirements, [2-1](#)

applications
failover, [3-8](#)

architectures
manageability, [2-5](#)
requirements, [2-1](#)
roadmap, [1-8](#)

automatic block repair, [3-36](#)

automatic corruption repair, [3-8](#)

availability, [1-1](#)
about, [1-1](#)
roadmap, [1-8](#)

B

backing out a transaction, [3-34](#)

block recovery
using Flashback logs, [3-36](#)

C

components
integrated with Oracle Restart, [3-44](#)

computer failure, [1-3](#)

corruptions
automatic repair, [3-8](#)
prevention and detection, [3-27](#)

D

data corruptions, [1-3](#)
detecting, [3-27](#)
prevention and detection parameters, [3-27](#)

data distribution
Oracle GoldenGate, [3-12](#)

Data Guard, [3-2](#)
about, [3-2](#)
benefits, [3-2](#), [3-8](#)

data integration
Oracle GoldenGate, [3-12](#)

Data Recovery Advisor, [3-30](#)

data-loss tolerance, [2-5](#)

DBA_FLASHBACK_TRANSACTION_STATE
view, [3-34](#)

DBMS_FLASHBACK.TRANSACTION_BACKOUT() procedure, [3-34](#)

disk group
administering with Oracle ASM, [3-24](#)

downtime, [1-3](#)
causes, [1-3](#)
cost, [1-3](#)
reducing, [3-8](#)
See also planned downtime

F

failovers
applications, [3-8](#)
services, [3-21](#)

failure group
administering with Oracle ASM, [3-24](#)
Oracle ASM, [3-24](#)

failures
computer, [1-3](#)
site, [1-3](#)
storage, [1-3](#)

Fast Connection Failover
for nonpooled connections, [3-21](#)

Fast Mirror Resync
Oracle ASM, [3-24](#)

fast recovery area
about, [3-26](#)
benefits, [3-26](#)

flashback
PDB, [3-36](#)

flashback logs
block recovery using, [3-36](#)

Flashback technology, [3-31](#)
See also Oracle Flashback technology

H

hangs or slow down, [1-3](#)

high availability, [1-1](#)
 about, [1-1](#)
 determining requirements, [2-1](#)
 single-instance databases, [3-44](#)
 See also availability
 human errors, [1-3](#)

I

instance failure, [3-21](#)
 interblock corruption, [1-3](#)
 intrablock corruption, [1-3](#)

L

load balancing
 advisory, [3-21](#)
 run-time connection, [3-21](#)
 load balancing advisory, [3-21](#)
 logical corruption, [1-3](#)
 lost writes, [1-3](#)

M

Maximum Availability Architecture
 See Oracle Maximum Availability Architecture (MAA)
 media corruption
 physical corruption, [1-3](#)
 mirroring
 Oracle ASM native, [3-24](#)

N

network bonding, [3-21](#)

O

offloading database activity, [3-8](#)
 online redefinition, [3-45](#)
 Oracle Active Data Guard, [3-6](#)
 benefits of standby databases, [3-6](#)
 Oracle ASM Cluster File System (ACFS), [3-39](#)
 Oracle Automatic Storage Management (Oracle ASM)
 about, [3-24](#)
 benefits, [3-24](#)
 failure group, [3-24](#)
 Fast Mirror Resync, [3-24](#)
 native mirroring, [3-24](#)
 Oracle Automatic Storage Management Cluster File System (Oracle ACFS), [3-24](#)
 Oracle Call Interface (OCI), [3-21](#)
 Oracle Clusterware, [3-19](#), [3-20](#)
 cold cluster failover, [3-21](#)

Oracle Data Guard
 See Data Guard
 Oracle Data Provider for .NET (ODP.NET), [3-21](#)
 Oracle Database File System (DBFS), [3-8](#), [3-40](#)
 Oracle Enterprise Manager, [3-21](#)
 Oracle Flashback Data Archive, [3-37](#)
 Oracle Flashback Database, [3-35](#)
 Oracle Flashback Drop, [3-35](#)
 Oracle Flashback Query, [3-33](#)
 Oracle Flashback Table, [3-34](#)
 Oracle Flashback technology, [3-31](#)
 block recovery using Flashback logs, [3-36](#)
 Oracle Flashback Transaction, [3-34](#)
 Oracle Flashback Transaction Query, [3-34](#)
 Oracle Flashback Version Query, [3-33](#)
 Oracle GoldenGate
 about, [3-12](#)
 Oracle Maximum Availability Architecture (MAA),
[1-8](#)
 about, [1-8](#)
 roadmap, [1-8](#)
 Oracle Multitenant, [3-41](#)
 Oracle Quality of Service (QoS) Management,
[3-21](#)
 Oracle RAC One Node, [3-21](#)
 Oracle Real Application Clusters, [3-19](#)
 Oracle Real Application Clusters (Oracle RAC)
 benefits, [3-21](#)
 Oracle Restart, [3-44](#)
 Oracle Sharding
 overview, [3-44](#)
 Oracle Solaris ZFS Storage Appliance
 Replication, [3-40](#)
 Oracle UCP run-time connection load balancing,
[3-21](#)
 outages
 types of, [1-3](#)

P

patching
 rolling, [3-21](#)
 physical corruption, [1-3](#)
 Plug and Play, [3-21](#)
 policy-based cluster management, [3-21](#)

R

Recovery Manager (RMAN)
 about, [3-17](#)
 benefits, [3-17](#)
 recovery point objective (RPO)
 about, [2-5](#)
 recovery time objective (RTO)
 about, [2-4](#)

- reference architectures
 - data protection attributes, [2-9](#)
 - gold, [2-8](#)
 - high availability attributes, [2-9](#)
 - overview, [2-7](#)
 - platinum, [2-9](#)
 - silver, [2-8](#)
- reference architectures
 - bronze, [2-8](#)
- replication
 - Oracle GoldenGate, [3-12](#)
- restore points
 - Oracle Flashback, [3-35](#)
- return on investment (ROI), [2-5](#)
- roadmap to Maximum Availability Architecture (MAA), [1-8](#)
- rollback
 - transactions, [3-34](#)
- rolling patching, [3-21](#)
- rolling upgrades, [3-21](#)
- run-time connection load balancing, [3-21](#)

S

- scalability
 - with Oracle RAC, [3-21](#)
- SCAN, [3-21](#)
- security
 - Oracle ASM, [3-24](#)
 - RMAN, [3-17](#)
- services
 - failover, [3-21](#)
- single-instance databases
 - Oracle Restart, [3-44](#)
- site failure, [1-3](#)

- standby databases
 - benefits, [3-2](#)
- storage
 - failures, [1-3](#), [3-24](#)
 - Oracle ASM protection, [3-24](#)

T

- total cost of ownership (TCO), [2-5](#)
- transactions
 - backing out with Flashback Transaction, [3-34](#)

U

- unplanned downtime, [1-3](#)
 - causes, [1-3](#)
- upgrades
 - rolling, [3-21](#)

V

- virtual IP (VIP) address
 - managed by Oracle Clusterware, [3-20](#)

W

- workload
 - offloading, [3-8](#)
- workload management, [3-21](#)

Z

- zero data loss, [3-8](#)