

Preventing, Detecting, and
Repairing Block Corruption:
Oracle Database 11g

*Oracle Maximum Availability Architecture White Paper
May 2012*

Maximum Availability Architecture

Oracle Best Practices For High Availability

Executive Overview	2
Introduction	2
How Corruption Manifests Itself	3
Physical and Logical Corruptions	4
Intrablock and Interblock Corruptions.....	4
Corruption Prevention, Detection, and Repair.....	5
Preventing Widespread Block Corruption	5
Detecting and Monitoring Block Corruption.....	8
Detecting Block Corruptions using Grid Control.....	8
Detection and Diagnosis Using Data Recovery Advisor	9
Detecting Block Corruptions using Data Guard	9
Detecting Block Corruptions using SQL*Plus	10
Recovering from Data Corruptions	10
Use Data Recovery Advisor	10
Use Active Data Guard.....	11
Oracle Active Data Guard and Automatic Block Repair.....	11
Physical Standby Databases.....	11
Use RMAN and Block Media Recovery.....	12
Perform a Data Guard Role Transition.....	12
Use RMAN and Data File Media Recovery	13
Conclusion	14
References	15

Executive Overview

In today's information-driven enterprises, business critical information must be highly available, from internal company-sensitive data to applications that manage crucial partner and customer relationships. Outages are costly in lost revenue and damaged reputation, and can be crippling to the business.

Block corruptions are a common source of database outages. A database block is corrupted when its content has changed from what Oracle Database expects to find. If not prevented or repaired, block corruption can bring down the database and possibly result in the loss of key business data.

This Oracle Maximum Availability Architecture (MAA) white paper presents the essential tools, techniques, and best practices to help you safeguard data and prevent corruption-related database outages. MAA is Oracle's comprehensive best practices blueprint for implementing Oracle high availability technologies. The MAA best practices are described in a series of technical white papers and documentation to assist in designing, implementing, and managing optimal high availability architectures. The series of MAA white papers are available at the [Oracle Maximum Availability Architecture website](#) [1].

In general, Oracle Database 11g Release 2 (11.2) is the database release under discussion in this white paper. This paper also applies to Oracle Exadata environments.

Introduction

A data block is corrupted when it is not in a recognized Oracle Database format, or its contents are not internally consistent. Data block corruption can damage internal Oracle control information or application and user data, leading to crippling loss of critical data and services.

Block corruptions may affect only a single block or a large portion of the database (making it essentially unusable). While relatively rare, corruptions are inevitable, but Oracle provides a complete set of technologies to prevent and mitigate block corruptions.

Block corruptions are inevitable . . . Corruptions stem from many different hardware or software defects. For example, corruption can be caused by a faulty disk or disk controller, an errant bit-flip on a disk, or a bug in the operating system, storage area network (SAN), or storage system. Although the best way to address block corruption is to prevent it from happening in the first place, all corruptions cannot be completely prevented.

But you can protect your data . . . While you cannot prevent all block corruptions, there is a comprehensive set of data protection solutions that you can implement to address most of them. Oracle offers sophisticated solutions—such as:

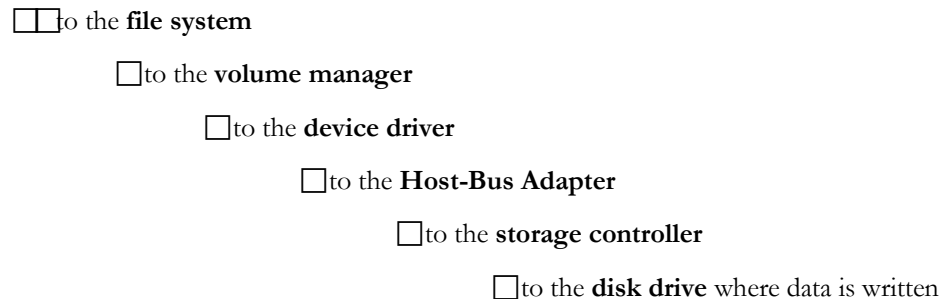
- Oracle Data Guard
- Data Recovery Advisor
- Oracle Flashback
- Oracle Recovery Manager (RMAN)
- Automatic Diagnostic Recovery (ADR)
- Oracle Secure Backup
- The MAA Advisor component of Oracle Enterprise Manager Grid Control
- Exadata Storage

These features offer database optimized ways to protect your data and ensure high availability of your data and hence, the application.

Oracle’s integrated and database-aware technologies are fundamentally different from generic hardware-centric solutions. Third-party generic solutions are limited by their lack of internal knowledge of the Oracle database blocks and the very transactions that they aim to protect. For instance, storage mirroring works at the bits-and-bytes level with no understanding of database block formats and consistency. In contrast, Oracle solutions provide database aware solution that are comprehensive, efficient, and effective protection at the level of business data and services, by leveraging internal knowledge of the data structures themselves and the integration of high availability features throughout the database.

How Corruption Manifests Itself

When Oracle issues a write operation, it moves through the following I/O sequence:



Hardware failures or bugs in any layer can result in corrupt data being written to disk, or good data not written to disk (termed a “lost write”) yet reported as written to Oracle.

Physical and Logical Corruptions

Data corruption can manifest itself as a physical or a logical corruption:

- **Physical Corruption** of a block manifests as an invalid checksum or header, or when the block contains all zeroes. When that happens, the database will not recognize the block as a valid Oracle block, regardless of its content. A physical corruption is also called a media corruption.
- **Logical Corruption** happens when a data block has a valid checksum, etc., but the block contents are logically inconsistent. Logical block corruption can also occur when the structure below the beginning of the block (below the block header) is corrupt. In this case, the block checksum is correct but the block structures may be corrupt. Logical corruption can also result from a lost write¹, in which a version of the block that was sent to disk somehow was not actually written. The result is that the version of that block on disk is older than the version in the buffer cache. Lost writes are usually caused by bugs in the operating system or hardware.

For more information, see My Oracle Support Note 840978.1.

Intrablock and Interblock Corruptions

The data blocks to which we refer are Oracle data blocks, which are comprised of multiple operating system blocks that make up the database. The data blocks are stored on disk, but are also temporarily stored in the buffer cache in memory. Thus, corruptions do not always appear on disk and can be related to memory and transient in nature.

- For **intrablock corruption**, the corruption occurs in the block itself and can be either a physical or a logical corruption.
- For **interblock corruption**, the corruption occurs between blocks and can only be a logical corruption.

¹ A **lost write** is a write I/O to persistent storage that the database believes has occurred based on information from the I/O subsystem, when in fact the write has not occurred. Lost writes are typically due to bugs in the operating system or hardware.

Corruption Prevention, Detection, and Repair

The Oracle Database corruption prevention, detection, and repair capabilities are built on internal knowledge of the data and transactions it protects, and on the intelligent integration of its comprehensive high availability solutions.

The MAA recommendation to achieve the most comprehensive data corruption prevention and detection is to use Oracle Data Guard with physical standby databases and configure the `DB_BLOCK_CHECKING`, `DB_BLOCK_CHECKSUM`, and `DB_LOST_WRITE_PROTECT` initialization parameters on the Data Guard primary and standby databases.

In addition, Oracle Automatic Storage Management (Oracle ASM) provides disk mirroring to protect against disk failures.

Once the corruption is detected, Oracle Data Guard, block media recovery², and data file media recovery³ can be used to recover the data. Database-wide logical corruptions caused by human or application errors can be undone with Oracle Flashback Technologies. Tools are also available for proactive validation of logical data structures, for example, the SQL*Plus `ANALYZE TABLE` command detects interblock corruptions.

Ultimately, your corruption recovery strategy is only as good as the extent to which you have tested it. Whatever method is employed to prevent or recover from corruptions, all recovery strategies should be thoroughly and regularly tested.

The remainder of this white paper describes the methods and best practices for using these solutions. By detecting and repairing corruptions early, the cause can be investigated and resolved before it becomes widespread.

Preventing Widespread Block Corruption

The best medicine for addressing block corruption is prevention. Block checking is the key approach to preventing widespread corrupted data. The Oracle Database validates and adds

² **Block media recovery** is a technique for restoring and recovering corrupt data blocks while data files are online. If only a few blocks are corrupt, then block media recovery may be preferable to data file media recovery. For automatic block media recovery to work, a physical standby database must be in real-time query mode, which requires an Oracle Active Data Guard license.

³ **Data file media recovery** repairs a lost or damaged current data file or control file. It can also recover changes lost when a tablespace goes offline without the `OFFLINE NORMAL` option. Media recovery is the application of redo or incremental backups to a data block or backup data file.

protection information to the database blocks. Oracle provides superior corruption and prevention through a variety of configurations, technologies, and methodologies.

This section discusses basic database configuration options to protect against block corruption. The MAA best practice is to implement these features in every environment as a preventative measure.

- **Use Oracle Data Guard**

Oracle Data Guard is the Oracle MAA recommended data availability solution, and it is the best solution for protecting against data loss, corruptions, and lost writes.

Data Guard maintains a copy of your data in a standby database that is continuously updated with changes from the production database. Data Guard validates all changes before they are applied to the standby database, preventing physical corruptions that occur in the storage layer from causing data loss and downtime.

Starting in Oracle Database 11g Release 2 (11.2), the primary database automatically attempts to repair the corrupted block in real time by fetching a good version of the same block from an Active Data Guard physical standby database. Automatic block repair requires licensing Oracle Active Data Guard⁴.

- **Set the Oracle Database generic block-corruption parameters**

For the most comprehensive block-corruption protection set the `DB_BLOCK_CHECKSUM` and `DB_BLOCK_CHECKING` and `DB_LOST_WRITE_PROTECT` parameters as follows.

Configure the following on the Data Guard primary database:

- `DB_BLOCK_CHECKSUM=FULL`
- `DB_BLOCK_CHECKING=FULL` or `MEDIUM`
- `DB_LOST_WRITE_PROTECT=TYPICAL`
- Enable Flashback Technologies for fast point-in-time recovery from logical corruptions most often caused by human error and for fast reinstatement of the new standby database following failover.

⁴ Oracle Active Data Guard is a database option license for Oracle Enterprise Edition. An Active Data Guard license is required for the primary database and each Active Data Guard physical standby database included in the Data Guard configuration.

Configure the following on the Data Guard standby database:

- `DB_BLOCK_CHECKSUM=FULL`
- `DB_BLOCK_CHECKING=FULL` or `MEDIUM`
- `DB_LOST_WRITE_PROTECT=TYPICAL`
- Enable Flashback Technologies for fast point-in-time recovery from logical corruptions most often caused by human error.
- Use Active Data Guard to enable Automatic Block Repair (Active Data Guard release 11.2 and later).

Important Note: Each of the above data protection features will impact performance to varying degrees and thus require testing before introducing to production. See My Oracle Support Note 1302539.1 for additional discussion of protection/performance tradeoffs for each parameter.

- **Implement a backup and recovery strategy with Recovery Manager (RMAN)**

Implement a good backup and recovery strategy to ensure protection from media corruption. Oracle Recovery Manager (RMAN) is the optimal solution to backup Oracle Databases. RMAN knows which files require a backup, but most importantly, it knows which files are needed for recovery. For more information about RMAN, see the *Oracle Database Backup and Recovery User's Guide* [2].

When reading data for backup or restore, RMAN verifies block checksums. By default, RMAN stops the backup operation as soon as it encounters a block corruption. Also, you can validate individual blocks on demand with the `VALIDATE DATAFILE . . . BLOCK` command. For more information about the `VALIDATE` commands, see the [Oracle Database Backup and Recovery User's Guide](#) documentation [2].

To get the strongest possible corruption protection with your RMAN backups use the `CHECK LOGICAL` option when performing backups. This will check the block content itself for logical corruptions, including missing row pieces and mismatched indexes.

Detecting and Monitoring Block Corruption

If corrupt data is written to disk or if a component failure causes good data to become corrupt after it is written, then it is critical to detect the corrupted blocks as soon as possible⁵.

The easiest and best way to monitor the database for errors and alerts is using the Oracle Grid Control Home page in Oracle Enterprise Manager. The Grid Control home page shows the availability of *all* discovered targets.

You can also view, diagnose, and repair corruption in the database using the Data Recovery Advisor. This feature automatically diagnoses persistent (on-disk) data failures, presents appropriate repair options, and runs repair operations at your request.

In addition, you can use Data Guard to monitor and detect corruptions by ensuring that the changes from the Primary database are successfully being sent and applied to the standby.

Detecting Block Corruptions using Grid Control

A major benefit of Oracle Grid Control is its ability to manage components across the entire application stack, from the host operating system to a user or packaged application. Oracle Grid Control treats each of the layers in the application as a target. Targets—such as databases, application servers, and hardware—can then be viewed along with other targets of the same type, or can be grouped by application type. You can also review all targets in a single view from the HA Console. See the [Oracle Database High Availability Best Practices](#) for more tips and recommendations [3].

- Use alerts to monitor the alert log for errors.

In Grid Control, alerts are generated by a combination of factors and are defined on specific metrics. A metric is a data point sampled by a Management Agent and sent to the Oracle Management Repository. In the Alert Log Metric group, set Oracle Grid Control to monitor the alert log for ‘Alert’ and ‘Block Corruption’.

- Use Grid Control events to monitor Data Guard system availability.

Set Grid Control metrics ‘Data Guard Status’ and ‘Apply Lag’ to monitor the availability of logical and physical Data Guard configurations.

⁵ When an error occurs, Oracle Database writes to the alert log a chronological record of messages and errors, including all internal errors (ORA-600) and block corruption errors (ORA-1578).

- If you choose to monitor the alert log directly, periodically check the log for any ORA- errors and any corruption warnings on the primary and standby databases.
- Query the V\$DATABASE_BLOCK_CORRUPTION view.

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

Detection and Diagnosis Using Data Recovery Advisor

The Data Recovery Advisor automatically diagnoses data failures, determines and presents appropriate repair options, and performs repair operations at the user's request. Data Recovery Advisor improves the manageability and reliability of an Oracle database. You can use Data Recovery Advisor to troubleshoot primary databases, logical standby databases, and snapshot standby databases.

Use Data Recovery Advisor to repair block corruptions, undo corruptions, and data dictionary corruptions. The Data Recovery Advisor integrates with the Enterprise Manager Support Workbench (Support Workbench), with the Health Monitor, and with RMAN to display data corruption problems, assess the extent of each problem (critical, high priority, low priority), describe the impact of a problem, recommend repair options, conduct a feasibility check of the customer-chosen option, and automate the repair process.

For more information about Data Recovery Advisor, see the "Diagnosing and Repairing Failures with the Data Recovery Advisor" section in the Oracle Database Backup and Recovery User's Guide [2].

Detecting Block Corruptions using Data Guard

Data Guard can detect physical corruptions in two ways, by the apply process stopping due to a corrupted block in the redo stream and when it detects a *lost write*. If you are not using Grid Control to manage and monitor your Data Guard configuration you can manually monitor the standby apply lag.

- Query the V\$DATAGUARD_STATS view. The Apply Lag and Transport Lag are specified in seconds and the granularity for refreshing this view is 1 minute. Monitor this view every 15 minutes on the standby database. Note that the V\$DATAGUARD_STATS view reads the control file each time you execute this query.
- Query the V\$RECOVERY_PROGRESS view. The "Last Applied Redo" timestamp will also give you an idea about how apply is lagging.

Data Guard will also detect a Lost Write at the Primary or Standby databases when enabled as recommended. See My Oracle Support Note 1302539.1 “Best Practices for Corruption Detection, Prevention, and Automatic Repair - in a Data Guard Configuration” for up to date information on resolving lost write and other types of database corruptions.

Detecting Block Corruptions using SQL*Plus

Data file corruptions and particularly interblock corruptions, can be detected manually by issuing the `ANALYZE TABLE <tablename> VALIDATE STRUCTURE CASCADE SQL*Plus` statement. After determining the corruptions, the table can be re-created or another action can be taken. See “Validating Tables, Indexes, Clusters, and Materialized Views” in the [Oracle Database Administrator's Guide](#) for more information [4].

Recovering from Data Corruptions

The recovery process begins when you either suspect or discover a corruption. Once a corrupt block is found, Oracle provides various techniques for recovering from most block corruptions. This section describes techniques for recovering data blocks. In general:

- Data Recovery Advisor is the simplest way to diagnose and repair most problems.
- Active Data Guard can automatically repair corrupt data blocks in a primary or standby database.
- RMAN block media recovery can repair individual corrupted blocks by retrieving the blocks from good backups.
- Data Guard switchover or failover to a standby database.
- Data file media recovery with RMAN

Whatever method you use to recover corrupted blocks, you first need to analyze the type and degree of corruption to perform the recovery. Implementing the optimal techniques to prevent and prepare for data corruptions can save time, effort, and stress when dealing with the possible consequences—lost data and downtime.

The following list prescribes the MAA best practices using a step-by-step process for resolving most corruptions, including stray and lost writes.

Use Data Recovery Advisor

The Data Recovery Advisor enables you to perform restore operations and recovery procedures or use Flashback Database as follows:

- Perform block media recovery of data files that have corrupted blocks
- Perform point-in-time recovery of the database or selected tablespaces
- Rewind the entire database with Flashback Database

- Completely restore and recover the database from a backup

You can use the Oracle Enterprise Manager Database Control (Support Workbench) or Grid Control to guide you through the recovery. The [Oracle Database 2 Day DBA](#) documentation provides details on how to use the GUI interface for Data Recovery Advisor [5].

If you prefer RMAN command line, the Data Recovery Advisor commands include `LIST FAILURE`, `ADVISE FAILURE`, `REPAIR FAILURE`, and `CHANGE FAILURE`. See the chapter about “Diagnosing and Repairing Failures with Data Recovery Advisor” in the [Oracle Database Backup and Recovery User’s Guide](#) [2].

If the Data Recover Advisor fixes the problem, there is no need to continue with any further recovery methods in this paper. However, continue to periodically check the alert log for any ORA errors and any corruption warnings on the primary and standby databases. While the database is operational and corruptions are detected, they will be recorded as ORA-600 or ORA-01578 in the alert log.

Use Active Data Guard

The following sections describe the two options for using a standby database to repair block corruption on the primary database: extracting data form a physical standby database or automatically repairing the corruption with Oracle Active Data Guard.

If the corruption is widespread, you may choose to failover or switchover to the standby database while repairs to the original primary database are made so that it can be reinstated as a new standby.

Oracle Active Data Guard and Automatic Block Repair

Starting in Oracle Database 11g Release 2 (11.2), the primary database automatically attempts to repair the corrupted block in real time by fetching a good version of the same block from a physical standby database. This capability is referred to as automatic block repair, and it allows corrupt data blocks to be automatically repaired as soon as the corruption is detected. Automatic block repair reduces the amount of time that data is inaccessible due to block corruption. It also reduces block recovery time by using up-to-date good blocks in real-time, as opposed to retrieving blocks from disk or tape backups, or from Flashback logs.

Automatic block repair requires the use of the Oracle Active Data Guard option.

You can use an Oracle Active Data Guard standby database for automatic repair of data corruptions detected by the primary database. Additionally if the corruption is discovered on an Active Data Guard physical standby database the corruption will be automatically repaired with a good block from the Primary. Both of these operations are transparent to the applications.

Physical Standby Databases

You can use a Data Guard physical standby database to repair data file wide block corruption on the primary database by replacing the corrupted data files with good copies from the standby. Once the files are restored on the primary database, archived logs are applied to make it consistent with the rest of the database.

For more information, see the [Oracle Database Data Guard Concepts and Administration](#) documentation [6].

Use RMAN and Block Media Recovery

Block media recovery recovers one block or a set of data blocks marked "media corrupt" in a data file by using the RMAN RECOVER BLOCK command. When a small number of data blocks are marked media corrupt and require media recovery, you can selectively restore and recover damaged blocks rather than whole data files. Block media recovery minimizes redo application time and avoids I/O overhead during recovery. It also enables affected data files to remain online during recovery of the corrupt blocks. The corrupt blocks, however, remain unavailable until they are completely recovered.

Use block media recovery when:

- A small number of blocks require media recovery and you know which blocks need recovery.
- Blocks are marked corrupt (you can verify this with the RMAN VALIDATE CHECK LOGICAL command).
- The backup file for the corrupted data file is available locally or can be retrieved from a remote location.

Do not use block media recovery to recover from user errors or software bugs that cause logical corruption where the data blocks are intact.

If a significant portion of the data file is corrupt, or if the extent of the corruption is unknown, then use either RMAN to restore the file from a backup, switch to an on-disk image copy, or switchover to your Data Guard standby database..

When corruption is detected, recover the block through the EM Restore and Recovery Wizard or directly with RMAN. For example, to recover a specific corrupt block using RMAN block media recovery:

```
RMAN> RECOVER BLOCK DATAFILE 7 BLOCK 3;
```

After a corrupt block is repaired, the row identifying this block is deleted from the V\$DATABASE_BLOCK_CORRUPTION view.

For more information about RMAN's block media recovery, see the [Oracle Database Backup and Recovery User's Guide](#) [2].

Perform a Data Guard Role Transition

If the primary database corruption is widespread due to a bad controller or other hardware or software problem, then you may want to failover or switchover to the standby database while repairs to the primary database server are made. Use Data Guard switchover or failover for data corruption or data failure when:

- The database is down or when the database is up but the application is unavailable because of data corruption or failure, and the time to restore and recover locally is long or unknown.
- Recovering locally takes longer than the business service-level agreement or RTO.

For more information about Data Guard failovers and switchovers, see the [Oracle Database Data Guard Concepts and Administration](#) documentation.

Use RMAN and Data File Media Recovery

When none of the above methods can be used and you do not have a Data Guard Physical standby, you must use traditional media recovery in which a backup copy of one or more files is restored, and then changes from archived log files are applied to the restored files to bring the database back to the current time or a prior point-in-time.

Data file media recovery affects an entire data file or set of data files for a database by using the RMAN RECOVER command. When a large or unknown number of data blocks are marked "media corrupt" and require media recovery, or when an entire file is lost, you must restore and recover the applicable data files.

For more information on data file recovery, see the [Oracle Database Backup and Recovery User's Guide](#) [2].

Conclusion

With the integrated data protection technologies and the MAA best practices, Oracle provides the most comprehensive and sophisticated suite of data protection, disaster recovery, and storage management capabilities.

The key approach to detecting and preventing corrupted data is to perform the following MAA Best Practices.

- Use Oracle Data Guard
- Set the Oracle Database block corruption detection parameters
- Implement a backup and recovery strategy with Recovery Manager (RMAN)

These high availability solutions are integrated with Oracle Database and understand the underlying data structures, thus offering a new class of intelligent data protection and disaster recovery.

References

1. *Oracle Maximum Availability Architecture*
<http://www.oracle.com/goto/maa>
2. *Oracle Database Backup and Recovery User's Guide*
http://download.oracle.com/docs/cd/E11882_01/backup.112/e10642/toc.htm
3. *Oracle Database High Availability Best Practices*
http://download.oracle.com/docs/cd/B28359_01/server.111/b28282/monitor.htm#g1011041
4. *Oracle Database Administrator's Guide*
http://download.oracle.com/docs/cd/E11882_01/server.112/e17120/toc.htm
5. *Oracle Database 2 Day DBA*
http://download.oracle.com/docs/cd/E11882_01/server.112/e10897/toc.htm
6. *Oracle Database Data Guard Concepts and Administration*
http://download.oracle.com/docs/cd/E11882_01/server.112/e17022/toc.htm



Preventing, Detecting, and Repairing Block

Corruption: Oracle Database 11g

May 2012

Contributing Authors: Timothy Chien, Cristobal
Pedregal-Martin, Viv Schupmann, Larry M.
Carpenter

Oracle Corporation
World Headquarters
500 Oracle Parkway
Redwood Shores, CA 94065
U.S.A.

Worldwide Inquiries:
Phone: +1.650.506.7000
Fax: +1.650.506.7200
oracle.com



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

Copyright © 2010, Oracle and/or its affiliates. All rights reserved. This document is provided for information purposes only and the contents hereof are subject to change without notice. This document is not warranted to be error-free, nor subject to any other warranties or conditions, whether expressed orally or implied in law, including implied warranties and conditions of merchantability or fitness for a particular purpose. We specifically disclaim any liability with respect to this document and no contractual obligations are formed either directly or indirectly by this document. This document may not be reproduced or transmitted in any form or by any means, electronic or mechanical, for any purpose, without our prior written permission.

Oracle is a registered trademark of Oracle Corporation and/or its affiliates. Other names may be trademarks of their respective owners.