

Oracle Active Data Guard versus Storage Remote Mirroring

An Oracle Technical Brief

May, 2023, Version 3.0
Copyright © 2023, Oracle and/or its affiliates
Public

Executive Summary

A critical objective for any enterprise is the protection of corporate assets, including data. Enterprises are equally concerned about the impact of application downtime when databases become unavailable due to outages caused by data corruptions, component, system, or site failures, failures due to software defects, or human errors. Increased cost, lost revenue, negative publicity, and regulatory non-compliance are just the beginning of many negative consequences resulting from data loss and downtime.

While a bullet-proof disaster recovery solution is the ultimate protection for enterprise data, businesses often assign this a lesser priority since disaster recovery infrastructure is expensive and rarely used. That leads to under-investment in disaster recovery or the deployment of solutions that provide inadequate protection and little confidence that they will work when called upon.

Oracle Active Data Guard fundamentally changes what enterprises expect from a disaster recovery solution for the Oracle Database. It provides the best comprehensive data protection, availability, and disaster recovery while effectively maximizing the benefits of the standby systems and the return on investment (ROI) by offloading production workload and critical activities to standby systems. Oracle Data Guard's deep integration with Oracle Database enables automatic failover for unplanned outages and database rolling upgrades that minimize downtime and risk when performing planned maintenance. That makes Oracle Data Guard a comprehensive solution for high availability and disaster recovery.

Oracle Active Data Guard also eliminates the inherent risks of generic data protection offered by storage-centric solutions such as array-based remote mirroring. Only Active Data Guard provides continuous real-time database block- and redo-level validation and automatic repair of physical block corruptions so that the disaster recovery system is ready for failover when needed.

This brief is intended for IT Managers, Database Administrators, and Architects evaluating disaster recovery solutions for the Oracle Database. It describes why customers prefer Oracle Data Guard and Active Data Guard to traditional disaster recovery solutions based on storage technologies.

Table of contents

Introduction: Active Data Guard and Data Guard	4
Why Data Guard Provides the Best Data Protection	5
Superior Isolation, Bandwidth Efficient	5
Continuous Oracle Data Validation	5
Automatic Repair	5
Lower Cost, High ROI	6
Remote Storage Mirroring – None of the Above	7
Why Data Guard Provides Better Availability (HA)	8
Fast, Automatic Failover	8
Database Rolling Maintenance	8
Better Data Protection means Better HA	9
Less risk: You Know it's Working.	9
Ease of Use	9
What about Storage Consistency Groups?	9
I/O Consistent Crash Point versus Transactional Consistency	9
How to Achieve Globally Consistent Point in Time using Oracle Technologies	10
Summary	11

List of figures

Figure 1: Oracle Data Guard and Active Data Guard Use Cases	4
Figure 2: Oracle Data Guard Architecture	5
Figure 3: Storage Remote Mirroring Architecture	7

Introduction: Active Data Guard and Data Guard

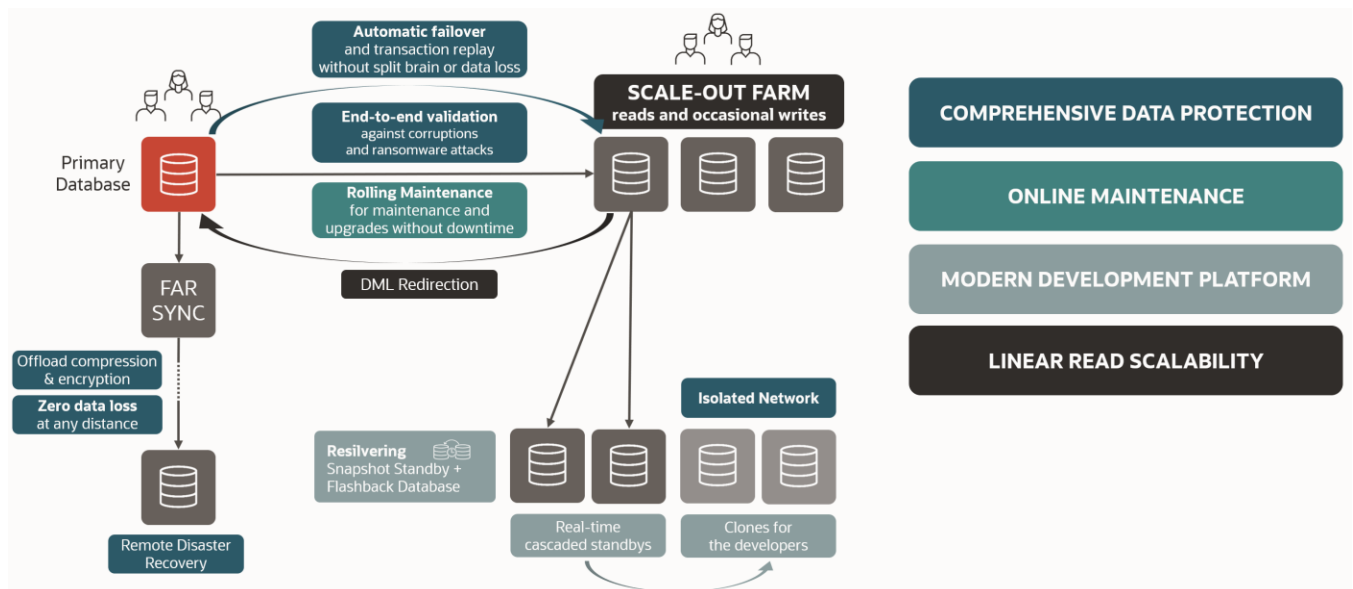
Oracle Data Guard evolves and perfects its capabilities release after release. Oracle Managed Standby, the precursor to Oracle Data Guard, first appeared in Oracle 7. It offered basic archive log shipping capabilities that required complementary scripts to maintain a synchronized replica of a production database. Oracle Data Guard was introduced with Oracle 9i and represented a significant technological evolution, eliminating the need for external scripts and providing complete management, monitoring, and automation software to create and maintain one or more replicas (standby databases) of the production database (primary).

Standby databases protect Oracle Database from failures, disasters, human error, and data corruption. Production applications can quickly resume work at the standby database if the primary becomes unavailable. Oracle Data Guard adds significant high availability (HA) features, making it a comprehensive solution for HA/DR optimized for the Oracle Database. Oracle Active Data Guard was introduced with Oracle Database 11g Release 1 to provide essential extensions to basic Oracle Data Guard functionality that further enhanced data protection, availability, and return on investment in standby systems. Oracle Active Data Guard is a database option that inherits all Data Guard capabilities adding numerous advanced features. The introduction of Oracle Multitenant Architecture in Oracle Database 12c extends all of Oracle Active Data Guard’s benefits to consolidated database environments, whether on-premises or in the Cloud.

Over the years, Oracle Data Guard has introduced innovative features for zero-data loss protection at any distance, complex topologies, rolling upgrades, and data manipulation at the standby site.

The last version, Oracle Database 23c, further improves the adoption of the Multitenant architecture with Oracle Data Guard per Pluggable Database and other enhancements for observability, manageability, and availability.

Figure 1: Oracle Data Guard and Active Data Guard Use Cases



Why Data Guard Provides the Best Data Protection

Some customers still use storage-based remote mirroring (array mirroring) to protect the Oracle Database with a replica. Storage mirroring is a sophisticated technology promoted as a generic infrastructure solution that makes a simple promise – whatever is written to primary storage will also be written to a mirrored copy at the remote site. Keeping this promise, however, can have disastrous consequences for data protection and availability when the data written to primary storage is corrupt due to failures such as hardware or software defects that checksum algorithms cannot detect.

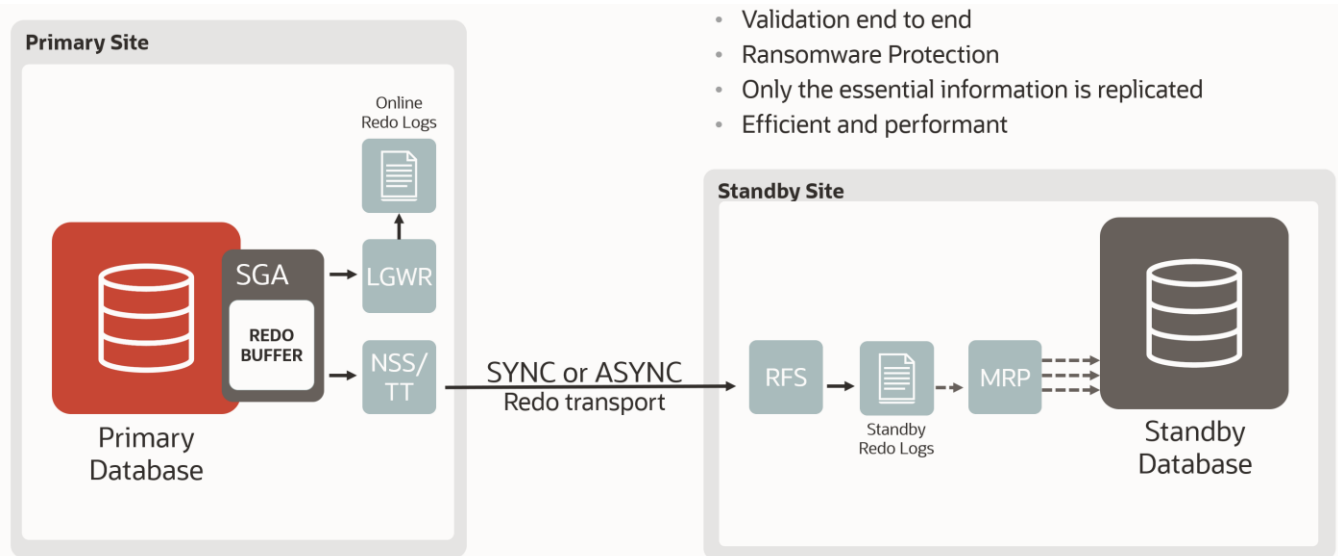
Modern cloud vendors also realize this and have stepped away from storage mirroring in favor of Oracle Data Guard to protect Oracle Relational Databases, which clearly indicates the benefits of Oracle Data Guard.

Oracle Data Guard and Active Data Guard provide better data protection and availability than is possible using storage technologies alone. Most enterprises have been replacing storage mirroring for Oracle databases with Oracle Active Data Guard for their business-critical databases for the following reasons:

Superior Isolation, Bandwidth Efficient

Simply stated, it is architecturally impossible for generic infrastructure solutions based upon storage mirroring to provide the same level of data protection as Oracle Data Guard. Oracle Data Guard is a lightweight Oracle-aware solution tightly integrated with the Oracle Database that provides superior isolation between the production database (the primary) and its standby database(s). Isolation from faults that can impact the primary copy is the most critical aspect of data protection.

Figure 2: high-level Oracle Data Guard Architecture



Oracle Data Guard replicates only the information needed to recover Oracle transactions (redo), representing a small percentage of the total write volume generated by an Oracle database. Oracle Data Guard replicates data directly from the memory of the primary database, ensuring that the standby is isolated from corruptions that the I/O stack can introduce.

Continuous Oracle Data Validation

An Oracle Data Guard standby is an independent hardware system. Oracle Database uses media recovery to apply the changes to the standby database to maintain a synchronized physical replica of the primary. Oracle Database recovery processes perform continuous validation as they apply changes to the standby. This validation uses knowledge of Oracle redo and data block structures to check for physical data corruption, logical intra-block corruption and lost writes to ensure the highest level of isolation between primary and standby.

Automatic Repair

Active Data Guard makes block-level corruptions invisible to users with no changes to the application code. Block-level corruptions are caused by intermittent random I/O errors that can occur independently at the primary or standby databases due to faulty hardware, controllers, or even alpha particles. Under normal operation, when an Oracle Database reads a block and detects corruption, it marks the block as corrupted and reports the error to the application. No

subsequent read of the block is successful until the database administrators manually recover the block - unless when using Active Data Guard. Active Data Guard automatically repairs physical block corruption at a primary database by retrieving a non-corrupted version of the block(s) from the active standby. Both high availability and data protection are always maintained.

Lower Cost, High ROI

Data Guard and Active Data Guard significantly reduce cost and increase return on investment compared to storage mirroring along several dimensions:

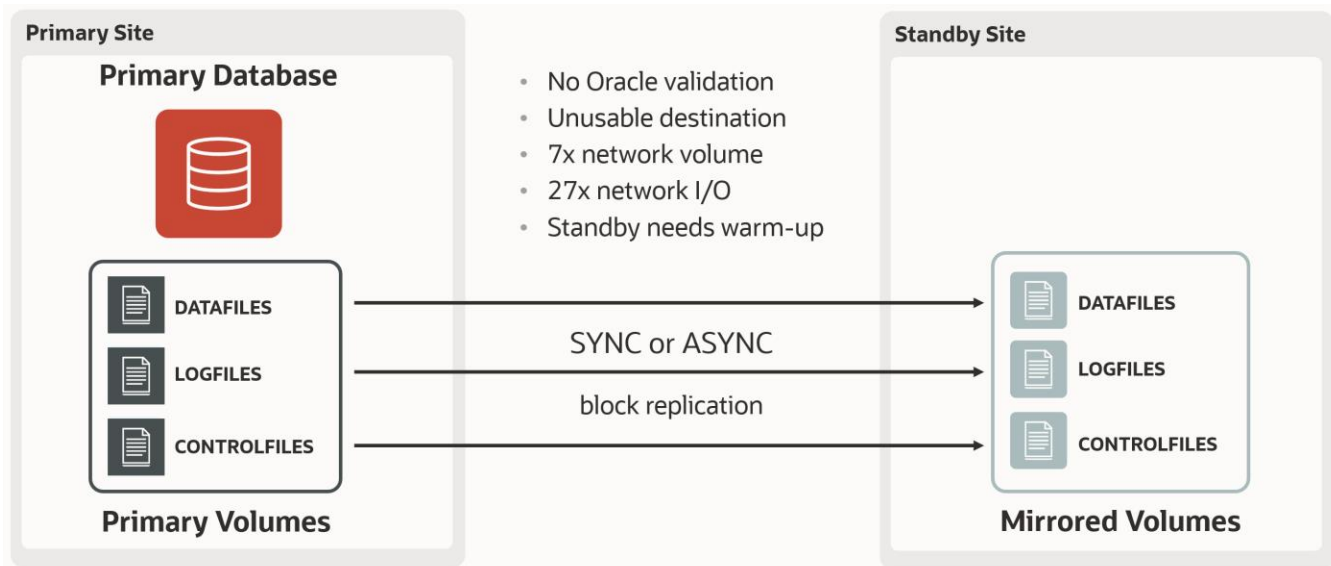
- Oracle Data Guard doesn't require additional licenses besides Oracle Database Enterprise Edition. Storage-based remote mirroring is often a premium-priced add-on to the storage cost.
- Oracle Data Guard's reduced network volume reduces bandwidth requirements significantly and can reduce the overall network cost.
- Oracle Data Guard's integration with Oracle Database and management tools reduces administrative costs.
- Advanced features included with the Active Data Guard option provide higher ROI by offloading the read-only workloads to the standby databases, allowing occasional writes on the standby, enhanced data protection, and increased high availability.
- Oracle Data Guard with Maximum Availability Architecture best practices can achieve very low end-to-end failover timings in seconds (lower RTO) thanks to Transparent Application Continuity and with zero or near zero data loss (zero to near-zero Recovery Point Objective (RPO)). Many storage-based failover solutions have a longer recovery time, with potential data loss.

Remote Storage Mirroring – None of the Above

Unlike Oracle Data Guard, remote storage mirroring has no awareness of the Oracle Database. The fact that it is a generic tool for block-level replication requires it to replicate a much higher volume of data than Oracle Data Guard to maintain real-time protection. That is due to two characteristics inherent to storage remote mirroring:

- Remote storage mirroring must replicate every write made to primary volumes (writes to data files, redo log files, archive log files, flashback log files, control files, TEMP files, etc.).
 - Writes to sort or work areas (TEMP) are useless for replication but contribute heavily to the overhead.
 - Oracle ASM automatically initiates rebalance operations after storage configuration changes, such as adding, dropping, or resizing disks. Storage replication can negatively impact ASM rebalance operations: ASM will rebalance the blocks without changing their content. However, these blocks are still replicated at the storage level, even though they don't change from a database perspective.
- Remote storage mirroring increases the replicated data volume because it must replicate the entire block (or potentially a bigger sector size), even if only a tiny portion of the data has changed.

Figure 3: High-level Storage Remote Mirroring Architecture



Storage mirroring doesn't perform any Oracle validation – it only does rudimentary physical checksums providing limited isolation between mirrored storage copies. Weak isolation and no Oracle validation allow physical corruptions and administrative errors (e.g., accidental deletion of data files or log files) to be faithfully replicated to the remote storage, making both copies unusable.

Even storage vendors acknowledge these limitations and often encourage complementary use of storage snapshot technology to perform point-in-time recovery when corruptions are replicated to the remote volumes.

Oracle believes differently. When a block corruption occurs, the primary database will not replicate it. Point-in-time recovery and the accompanying downtime and data loss will not happen in the first place. Oracle Data Guard understands native block and redo structures and uses that knowledge to perform continuous validation before applying the changes to the standby database. Unlike remote storage mirroring, Oracle Data Guard isolates problems happening on the primary database so that they do not affect the standby database. For this reason, Oracle Data Guard dramatically reduces downtime and data loss.

Why Data Guard Provides Better Availability (HA)

Remote storage mirroring cannot provide the same level of high availability as Data Guard and Active Data Guard, at least not for Oracle Databases.

Fast, Automatic Failover

Oracle Data Guard includes an automatic failover capability called Fast-Start Failover. For failover and switchover operations, applications can use the same client failover infrastructure as Oracle Real Application Clusters to resume their transactions at the new primary database. Oracle Data Guard Fast-Start Failover ensures that recovery point objectives (zero or a maximum allowable data loss threshold) are met. It safeguards against split-brain conditions (where two independent databases function as primary). Oracle Data Guard includes intelligent automation to reinstate the failed primary database as a synchronized standby, quickly returning the configuration to a protected state.

In contrast, remote storage mirroring has no Oracle-integrated capability to automate database failover and application redirection to a standby database. Remote storage mirroring requires time-consuming reconfiguration and start-up procedures to arrive at a state comparable to that of an Oracle Data Guard standby before the failure occurs. For example, the standby database volumes must be mounted when the primary database fails before starting the new primary database. These additional steps increase downtime and the risk of something else going wrong, which can lengthen the outage period.

Database Rolling Maintenance

Oracle Data Guard provides high availability while performing planned maintenance.

- Upgrades and many other types of maintenance can happen at the standby database.
- Once implemented and thoroughly tested with the new version, the standby database can transition to the primary role and serve the production workload.
- Total downtime is limited to the time required for the switchover.
- The standby can also be used to fully validate the new release before the switchover is performed without compromising data protection. That is not possible using array mirroring.

Performing maintenance in a rolling fashion and seamlessly using a standby database for preproduction testing is impossible using array mirroring.

Active Data Protection means Higher Availability.

The characteristics of Oracle Data Guard that result in better data protection - strong isolation and continuous Oracle data validation – also result in higher availability. Oracle Data Guard avoids the negative consequences of inadequate data protection.

Oracle Active Data Guard further extends Oracle Data Guard's advantages. Read-only workloads running on an Oracle Active Data Guard standby database implicitly provide continuous user and application-level validation of the standby readiness to failover. That is impossible with storage mirroring, where the only way to validate the replica is to stop the remote mirroring process and open the Oracle Database. That limits storage mirroring use cases to static, point-in-time validation that also reduces the protection level during the tests. Storage mirroring users often discover problems at their remote site at the most inopportune time – after the primary database failure.

Ease of Use

Determining whether it's easier to use Oracle Active Data Guard or remote storage mirroring is genuinely a matter of perspective.

When comparing the different approaches, it's critical to consider which solution can accomplish business objectives for data protection and availability.

From one perspective, remote storage mirroring can be perceived as easier to use because storage administration handles the configuration of volumes and operation of the mirroring process and recovery on behalf of the database administrators (DBA). The storage administrator uses the management interface provided by the storage vendor or provider to accomplish these tasks. Storage mirroring often uses storage volume groups – so multiple Oracle Databases can share a single replicated volume group. If storage mirroring is already in place, it can be simpler to continue using established processes and practices rather than introducing something new. It also provides a single mechanism for replicating any data type between sites – email, transactions, or sensitive personal data.

From a different perspective, Oracle Data Guard is easier to use because the database administrators control the replication and recovery process tightly integrated with other Oracle HA capabilities (e.g., Oracle Real Application Clusters, Automatic Storage Management, Recovery Manager, Flashback). This control makes a standby database immediately available for different recovery tasks. The database administrators can use the same management interface – Oracle Enterprise Manager Cloud Control, that provides integrated monitoring, diagnostics, and management of their Oracle environment. A database administrator can easily monitor critical Oracle Data Guard functions and execute database failover in seconds with a single mouse click.

What about Storage Consistency Groups?

A storage consistency group is a composite of volumes and volume groups with special properties to maintain the write consistency across multiple storage arrays. In an Oracle Database context, a consistency group ensures crash-tolerant consistency for Oracle Database files that span multiple volumes. Without this feature, remote storage mirroring of Oracle databases would not be possible because the remote mirrors would be inconsistent. Consistency groups provide a basic level of consistency to storage arrays without intrinsic knowledge of the nature of the replicated data. Consistency groups provide a crash-consistent copy when using array mirroring. That ensures that changes on the remote storage are written in the same order as the source.

Storage vendors often sell consistency groups to achieve global point-in-time consistency when multiple databases and applications are interdependent. Undoubtedly, storage consistency groups provide an essential value in this regard. However, it is incorrect to believe that consistency groups provide application or transactional level consistency for a set of databases.

I/O Consistent Crash Point versus Transactional Consistency

Consistency groups provide storage-level consistency but not necessarily transaction consistency across databases or between the database and filesystem data.

Achieving application consistency across multiple databases is hard without a transaction monitor coordinating the recovery across them. That still requires additional reconciliation at a transaction level.

A similar outcome occurs when using consistency groups to recover a mix of databases and non-database files. Each database recovers or rolls back the transactions to a different point than the filesystem recovery point.

The most straightforward solution for transactional consistency is two-phase commit / distributed transactions.

Interestingly, these protocols function without consistency groups - demonstrating that consistency groups are not helpful if global point-in-time consistency from an application perspective is the desired objective.

How to Achieve Globally Consistent Point in Time using Oracle Technologies

Oracle Database offers several options to achieve transactional consistency after disaster recovery. Each addresses the inherent shortcomings of consistency groups to meet the application requirements.

- File system data can be placed into the Oracle Database using Database Filesystem (DBFS). Once in the database, Oracle Data Guard replicates all the content and can fail over with transactional consistency.
- Oracle Database with Active Data Guard Far Sync enables zero data loss failover at any distance. Concerns for point-in-time consistency become a moot point when failover is a zero data loss event.
- Oracle has published a support note: Recovery for Global Consistency in an Oracle Distributed Database Environment (Doc ID 1096993.1). This support note describes achieving global point-in-time consistency across multiple independent databases. There are drawbacks - it doesn't address file system data, nor does it have the simplicity of storage consistency groups. Still, unlike storage consistency groups, it achieves transactional consistency when multiple databases are involved.
- After a failover to the disaster recovery site, a database administrator can determine the recovery point of the file system data, then uses Oracle Flashback Database to rewind the participating databases to the same point in time.

Summary

The objective of this paper has been to clearly and objectively illustrate why enterprises gain substantial benefits from Oracle Data Guard and Oracle Active Data Guard architectures to provide superior data protection and availability for the Oracle Database compared to array mirroring. IT managers are often caught between infrastructure teams that seek generic solutions due to perceived simplicity. In contrast, application and database administrators teams seek optimized solutions for a specific purpose providing better protection. Oracle Data Guard and Oracle Active Data Guard offer a common ground to address the complete range of business requirements with a simple-to-use standard infrastructure optimized to protect the Oracle Database.

When evaluating disaster recovery solutions, it is essential to focus on the main objectives:

- Having the highest degree of confidence that the data is safe from problems that can impact the primary database.
- Ensuring that service resumes within the required time frame.

Remote storage mirroring is architecturally limited by the degree of fault isolation it can provide and application knowledge it can apply to data protection and HA. Storage-centric solution proponents often promote consistency groups as the reason to dismiss these shortcomings of storage-based remote mirroring. The facts show that doing so places data and HA at risk, reduces ROI, and fails to achieve transactional consistency across multiple databases.

Oracle Data Guard exceeds the data protection and availability offered by storage array mirroring. Oracle Active Data Guard provides an additional return on investment and high availability benefiting the business and giving extra confidence that the standby will work when needed.

Connect with us

Call **+1.800.ORACLE1** or visit **oracle.com**. Outside North America, find your local office at: **oracle.com/contact**.

 blogs.oracle.com

 facebook.com/oracle

 twitter.com/oracle

Copyright © 2023, Oracle and/or its affiliates. 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, Java, MySQL, and NetSuite are registered trademarks of Oracle and/or its affiliates. Other names may be trademarks of their respective owners.