



Oracle Autonomous Database

Technical Brief Series: The Industries First Self-Repairing Database

July, 2020
Copyright © 2020, Oracle and/or its affiliates
Public

PURPOSE STATEMENT

This document focuses on the **Self-Repairing** attributes of the Oracle Autonomous Database. Autonomous Database is an Oracle Database service offering that runs in the Oracle Cloud. Self-Repairing, combined with Self-Driving and Self-Securing attributes, comprise the 3 key categories of autonomous capabilities within the Oracle Autonomous Database.

The initial sections of the paper are appropriate for business-level audiences. The details that follow may be more useful for DBAs and IT managers who are unfamiliar with the more recent Self-Repairing capabilities of Oracle Autonomous Database. This document is part of a series of Oracle Autonomous Database white papers. Details on the Self-Driving and Self-Securing capabilities of Oracle Autonomous Database are provided in separate Oracle white papers within this series.

The **“Introduction”** and **“What is an Autonomous Database?”** sections of this document are intentionally common to all of the Oracle Autonomous Database white papers in this series.

DISCLAIMER

This document is for informational purposes only and is intended solely to assist you in planning for the implementation and upgrade of the product features described. It is not a commitment to deliver any material, code, or functionality, and should not be relied upon in making purchasing decisions. The development, release, and timing of any features or functionality described in this document remains at the sole discretion of Oracle.

TABLE OF CONTENTS

Purpose Statement	2
Disclaimer	2
Introduction	4
What is an Autonomous Database?	4
Understanding the Need for Self-Repair	5
Oracle's Leadership in Self-Repairing Capabilities	6
Why Is a Self-Repairing Database Important?	6
Applicable environments for Self-Repairing (Autonomous) Databases	6
What Can a Self-Repairing Cloud Database Do?	7
Self-Repairing in Conjunction with Other Autonomous Capabilities	10
Autonomous monitoring and response: self-repair based on machine learning	10
How is the Autonomous monitoring and response Different?	10
Issues Addressed by the Autonomous monitoring and response?	11
How autonomous monitoring and response addresses these issues?	11
Conclusion	12

INTRODUCTIONⁱ

Relational databases have made tremendous improvements in performance, availability and security over the past couple of decades. They can run up to 100x faster; can be configured for zero data loss; and have hardened security capabilities that can protect against malicious internal and external threats. These attributes have been enhanced by cloud databases and infrastructure services that deliver elastic scalability and provisioning for real-time agility and growth. Although a great many databases will remain on-premises due to regulatory and data residency requirements or just to satisfy data access latency for highly connected systems that are difficult to move to cloud, many Database workloads that were deemed too large or “mission critical” to run outside corporate data centers just a few years ago now run in public clouds. In addition, capabilities such as database resource deployment, monitoring and management can also be automated, leading to greater operational efficiencies and cost savings. So, what’s missing? The degree of manual intervention required to manage today’s cloud databases and all of the above attributes inhibits true Database as a Service – as a utility, or driver-less offering, if you will. As a result, enterprises are unable to realize the full operational and financial benefits of the cloud.

WHAT IS AN AUTONOMOUS DATABASE?ⁱⁱ

There is understandably an element of confusion that arises when talking about “automatic” versus “autonomous” capabilities. A process for database backup, failover or resizing that can be accomplished **automatically** is still not **autonomous** if a database administrator has to respond to an alert, make decisions and click a few buttons (or type a few commands) in order to initiate the automated activity.

A more dramatic example is when an alert related to a component outage or performance degradation appears automatically on a management console, but doesn’t provide sufficient information to diagnose the problem, determine its root cause or offer a definitive recommendation for resolution. The automation literally stops with the alert. What happens next and how long it takes until resolution is unclear.

By contrast an autonomous database combines the dynamic agility of the cloud with the intelligent responsiveness of applied, adaptive machine learning. The design goal is to minimize or eliminate human labor – and associated human error – and ensure data safety and optimal performance. Businesses will find that autonomous capabilities can further help IT staff improve efficiencies by enabling them to focus on higher value activities in lieu of mundane, time-consuming tasks. This is significant considering that up to 75% of IT budgets are spent on manual database management.ⁱⁱⁱ An autonomous database can help organizations transform IT operations into a modern cloud model that lowers operating expenses, eliminates costly downtime and ultimately enables them to innovate more while using fewer resources.

Oracle Autonomous Database is designed to deliver the above benefits across 3 primary categories, all accomplished with minimal to zero human intervention.

- **Self-driving:** The Autonomous Database automates database and infrastructure provisioning, management, monitoring, backup, recovery and tuning.
- **Self-securing:** The Autonomous Database is more secure than a manually operated database because it automatically protects itself from internal and external vulnerabilities and attacks. The Oracle Cloud provides continuous threat detection, while the Autonomous Database automatically applies all security updates online and provides “always on”, end-to-end encryption. This preventative approach is critical because 85% of security breaches today occur after a CVE (common vulnerability and exposure) alert has been issued.^{iv}

- **Self-repairing:** The Autonomous Database provides preventative protection against all unplanned and planned downtime – and rapid, automatic recovery from outages without downtime. Autonomous Database availability and performance management is taken to the next level thru the use of AI-based autonomy that integrates multiple areas of diagnostics and enables analysis and action to be taken at runtime to minimize or eliminate operational disruption.




Image 1. Autonomous Database Components in Oracle Cloud

UNDERSTANDING THE NEED FOR SELF-REPAIR

Organizations of all sizes are becoming well versed with the risks associated with downtime and data loss, as well as performance bottlenecks that can affect data access. Over 90% of businesses have experienced an unplanned data center outage in the past 2 years.^v By some estimates, the average cost of database downtime is nearly \$8,000 per minute and the average outage lasts for approximately 90 minutes, which is about \$720,000 per incident.^{vi} This figure may not seem dramatic until we remember that organizations incur multiple outages per year, representing multiple hours or even days of downtime – which adds up to millions of dollars annually. And these are just averages. An outage of just a few minutes can cost larger, global organizations – banks, retailers, E-commerce sites – tens of millions of dollars. One doesn't have to imagine what these disruptions can do to an organization's reputation and customer good will, let alone its financial position and overall viability.

Each minute spent manually troubleshooting, analyzing, repairing, configuring, testing and bringing systems back on-line is precious. The time-pressure increases when these activities must be coordinated across multiple IT teams. This is where proactive, self-repairing capabilities can have a significant impact on database availability and performance, and the company's bottom line.

ORACLE'S LEADERSHIP IN SELF-REPAIRING CAPABILITIES

There is a broad range of events that can cause database downtime, including:

- **Component, storage and server failures, database crashes and even site-wide or regional outages due to a natural or man-made disaster.**
- **Data corruption that can cause incomplete backups or render the data useless.**
- **Human error, which plays a significant role in many cases – whether it's a database table that was dropped, a cable that was accidentally unplugged or a tape that was lost.**
- **Planned downtime for patching, upgrades, and maintenance – which represents an increasingly disproportionate percentage of overall downtime for many growing organizations.**

Oracle has successfully addressed all of these causes of downtime and disruption in on-premises environments for decades with the Oracle Maximum Availability Architecture (Oracle MAA). Oracle MAA is a set of advanced technologies and best practices that can be deployed to handle any service level requirement – with solutions ranging from periodic backups to zero data loss / zero downtime replication-based disaster recovery. The MAA portfolio is also available in the Oracle cloud database services and has been enhanced with automated functionality that minimizes, and in many cases, eliminates human intervention. Oracle MAA, along with advances in AI/ML based database systems and data center automation, represents the core of the Self-Repairing capabilities of Oracle Autonomous Database. No other vendor offers the comprehensive autonomous self-repairing database capabilities that Oracle does – on-premises or in the cloud.

WHY IS A SELF-REPAIRING DATABASE IMPORTANT?

A self-repairing database is more reliable than a manually operated database. Preventing or recovering from outages and removing bottlenecks are very often manual-intensive processes where human reaction time delays problem identification, diagnosis and repair. Such delays affect ongoing business transactions, user experience, revenue and margins. A self-repairing database has the ability to absorb any downtime or performance impacting event with zero human intervention.

By replacing manual repair and recovery tasks with extensive automation, Oracle Autonomous Database can eliminate human errors and reduce administration cost by up to 80%. Oracle also offers a 99.995% uptime guarantee for the Autonomous Database, which translates to less than 2.5 minutes of downtime per month in worst case scenarios. Unlike other cloud providers, Oracle's uptime guarantee includes planned maintenance and all other common sources of downtime in its calculations. There are NO unreasonable exclusions (for example, as seen on other cloud vendors, if the uptime guarantee were to exclude an outage caused by a bug in the cloud database engine).

APPLICABLE ENVIRONMENTS FOR SELF-REPAIRING (AUTONOMOUS) DATABASES

Oracle Autonomous Database is an Oracle Database based database cloud service offering. While many automated capabilities of Oracle Database are available both on-premises and in the cloud, a number of autonomous elements are unique to the Oracle cloud service offerings. Examples include autonomous patching, resource allocation and scaling, backup and availability SLAs.

Although Oracle Autonomous Database is a cloud-only offering, enterprises that must keep data behind corporate firewalls to meet data sovereignty, regulatory or data access latency requirements can run the Autonomous Database on-premises using Exadata Cloud@Customer. Autonomous Database on Exadata Cloud@Customer can be deployed on-premises, and delivers all of the capabilities of Autonomous Database in public cloud from within the enterprise's data center.

The Autonomous Database can be deployed in a hybrid cloud or all-cloud model; for example, when multiple databases are deployed for production and test environments or as primary and standby systems in a disaster recovery scenario.

There are no workload restrictions associated with the Autonomous Database or its self-repairing capabilities. This includes transaction processing, mixed workloads that involve transaction and batch processing and reporting, as well as analytic workloads associated with data warehouses and data lakes.

WHAT CAN A SELF-REPAIRING CLOUD DATABASE DO?

Integral to Oracle Autonomous Database is the Oracle Maximum Availability Architecture (MAA). It prevents (when possible) and provides automatic recovery from any type of outage, eliminating downtime and data loss for any environment in the cloud – from test databases to mission-critical environments. Most MAA technologies are de-facto industry-standards for protecting and repairing Oracle Database environments.

Differentiated autonomous technologies are included with each subsequent level of protection, beginning with backup as the first line of defense and evolving to real-time replication for global enterprise environments that cannot tolerate any downtime or data loss. The technologies and capabilities described below are not individual database options for the Autonomous Database as they are in on-premises Oracle Database environments. Rather, they are technologies that are automatically deployed and activated as needed based on the service level chosen by the enterprise at the time the service is initiated. For more details, please refer to [Oracle Maximum Availability Architecture in Oracle Cloud Infrastructure](#).




Image 2. Oracle Maximum Availability Architecture

- **Automated Backup:** Autonomous Database runs nightly automated backups to object storage in the Oracle Cloud, where data is protected from site outages with triple redundancy across multiple data centers in the cloud and accessible 24/7. Restores from cloud backups are also automated.
- **Key differentiators:**
 - Oracle Recovery Manager (RMAN), the same technology used to back up and restore on-premises Oracle databases, automatically manages the backup and recovery process.
 - RMAN provides unique knowledge of database file formats and recovery processes. It is deeply Oracle-aware and automatically performs physical and logical block validation, which ensures that corrupt data blocks are not propagated from an on-premises database to a backup copy in the cloud. You will always be able to completely restore ALL of your data from backups stored in the Oracle Cloud.
- **Automated Server Availability:** Oracle Real Application Clusters (RAC) is the industry-standard for database server high availability. It is designed to enable a single Oracle database to run across multiple nodes in a cluster in an active-active fashion for both high availability and dynamic scalability. In the event that any node in the cluster fails, sessions and users from the failed node are automatically and transparently moved to another node in the cluster, avoiding downtime and disruption to users. Oracle RAC is the definitive high availability solution for Oracle databases within a single Availability Domain (data center) in the Oracle Cloud.
- **Key differentiators:**
 - No other clustering solution in the industry can absorb and recover from database server or instance failures as quickly and non-disruptively as Oracle RAC.
 - Application Continuity, a capability available for RAC, enables any transactions that were “in-flight” at the time of an outage to be automatically replayed so that the application does not experience service disruption. No other vendor offers this capability.
 - Oracle RAC is supported only in the Oracle Cloud, not with any other cloud that may otherwise support Oracle Database.
- **Automated Disaster Recovery:** At startup, the Autonomous Database automatically establishes a triple-mirrored scale-out configuration in one regional cloud datacenter, with an optional full standby copy in another region for disaster protection. For disaster recovery between primary and standby copies, Data Guard provides a comprehensive replication solution that eliminates any single point of failure for mission-critical Oracle databases – from component failures to database crashes to site outages. It prevents data loss and downtime simply and economically by maintaining a synchronized physical replica (standby) of a production database (primary) in the Oracle Cloud. Databases protected using Data Guard automatically fail over to a standby database if the primary database becomes unavailable for any reason. Client connections then quickly and automatically fail over to the up-to-date standby and service is resumed.
- **Key differentiators:**
 - Data Guard has a deep understanding of Oracle data structures, which enables complete fault isolation and data validation and prevents any corrupt data from being committed on the standby database. These capabilities can’t be guaranteed by alternative technologies such as remote storage mirroring.
 - Active Data Guard extends Data Guard capabilities by automatically repairing any bad data blocks.

- Active Data Guard enables databases in the Oracle Cloud to be configured for zero data loss and automatic failover across any distance. This avoids the performance degradation that occurs with synchronized database replication (zero data loss configurations) over long distances when using competitive solutions.
 - Active Data Guard also improves resource utilization and ROI by enabling read-only workloads such as reporting, ad-hoc queries, data extracts and backups to be offloaded to an up-to-date physical standby database, without compromising disaster protection.
- **Automated Maintenance and Upgrades:** Up to 72% of IT time is spent on maintenance versus innovation. Autonomous Database offers many capabilities to eliminate planned downtime, which is increasingly problematic for growing organizations that have to support more servers, databases and applications – and consequently, spend more time on upgrades, patching and migrations.
- **Key Differentiators:**
 - Oracle RAC enables software updates to be applied in a rolling fashion across nodes of the cluster. It keeps the application online during updates to the database and upgrades to Oracle Clusterware, OS, VM, hypervisor, and firmware.
 - Active Data Guard automates the process for rolling database upgrades. It eliminates over 40 manual steps and reduces risk by thoroughly validating all changes on a replica of the production database before moving users to a new version.
 - Databases usually require multiple patch sets per year, which can translate to several hours of labor per database. The Autonomous Database applies patches automatically as they are released, taking a monumental DBA effort out of the patching process. Imagine the savings in administrative time and effort for global organizations with tens or hundreds of databases.
- **Automated Recovery from Human Errors:** Approximately 80% of unplanned outages are due to people and process issues. Flashback Database is a capability within Autonomous Database that automates recovery from human errors. It has been called a “rewind button” that enables the database to be rolled back and viewed at a previous point in time for error investigation and rapid recovery.
- **Key Differentiators:**
 - FlashBack Database eliminates the need for an expensive, time-consuming database restore.
 - Administrators can back out a transaction, update incorrect entries to a table, or if needed, rewind the entire database to any point in time since the last backup.
- **Optimal Database Cloud Infrastructure:** The Autonomous Database runs on Oracle Database Exadata Cloud infrastructure, which is based on the on-premises Exadata Database Machine. The Exadata-based cloud infrastructure includes all of the above Oracle MAA capabilities and provides compute, storage and networking resources that are integrated and optimized for Oracle Database.
- **Key Differentiators:**
 - As the embodiment of “enterprise class” the Exadata-based Cloud Infrastructure provides scalability, performance and availability for Autonomous Database that is unavailable with any other cloud solution.

- This proven architecture supports thousands of installations worldwide, including the world's most mission-critical databases run by Fortune 100 companies

SELF-REPAIRING IN CONJUNCTION WITH OTHER AUTONOMOUS CAPABILITIES

Oracle Autonomous Database is designed with a holistic approach and comprises components that are integrated and work together. For example, self-repairing capabilities such as backups, recovery, patching and failover (within a database cluster or to a standby) are accomplished by self-driving attributes of the database which automates such processes. Similarly, self-securing capabilities such as end-to-end encryption, data validation and corruption identification, isolation and repair are “always on” for backups and replication-based disaster recovery, which are part of the self-repairing capabilities.

AUTONOMOUS MONITORING AND RESPONSE: SELF-REPAIR BASED ON MACHINE LEARNING

The Oracle Autonomous Database monitoring and response is a collection of existing and next-generation integrated tools that complements Oracle MAA and is also key to the self-repairing capabilities of Autonomous Database. As with Oracle MAA, many of the monitoring and response capabilities have been available for on-premises Oracle Database environments, have been extended in the Oracle Cloud, have the additional advantage of leveraging data at scale.

Oracle Autonomous Database monitoring and response components combine self-driving and self-repairing capabilities to provide a monitoring and management framework that works autonomously to keep Oracle Databases healthy and running, while minimizing or eliminating human reaction time. Together, these tools identify situations that are potentially fatal to the database system, prevent these issues whenever possible, and take corrective action to resolve them rapidly and intelligently. Oracle monitoring and response components have been enhanced and are driven by applied machine learning to improve database runtime availability and performance with maximum efficiency. The components are lightweight and typically consume less than 3% of CPU resources at runtime as compared to on-premises databases running without autonomous monitoring and response.

HOW IS THE AUTONOMOUS MONITORING AND RESPONSE DIFFERENT?

Many health monitor solutions leverage machine learning from an anomaly detection perspective, but fail to deliver on response. The resulting manual triage efforts and floods of notifications do not scale in runtime environments. Autonomous monitoring and response goes past anomaly detection and symptoms and deeper into root cause analysis. This deeper *applied and adaptive* machine learning approach brings data science, algorithms and domain expertise together. It is based on thousands of hours of machine learning using a knowledge base developed from thousands of customer deployments for over a decade. These results have been duplicated, analyzed and resolved in Oracle lab environments and overseen by subject matter experts from Oracle technical staff.

Other machine learning frameworks require extensive training for customers to understand and become educated in using AI technologies. Oracle's approach is designed so that machine learning results and recommendations can be applied autonomously in seconds. A key goal and differentiator is to ensure that detection, analysis, and action are accomplished at runtime, minimizing operational disruption.

ISSUES ADDRESSED BY THE AUTONOMOUS MONITORING AND RESPONSE?

Oracle monitoring and response addresses availability and performance issues for single node and clustered database environments. These issues are typically handled by system and database administrators, including those associated with installation, patching, upgrades and resource availability of OS, hardware and database resources. Oracle monitoring and response automation assists in resolving both categories of issues at runtime, by autonomously monitoring and managing hardware resources as well as the database stacks.

Availability issues can be grouped into server and database categories.

- **Server availability:** These problems can cause a server to be evicted from a RAC cluster and shut down all database instances running on that server, resulting in a database outage. For example: memory stress caused by a server running out of free physical memory; network congestion or network failures due to a network card failure or a cable pull.
- **Database availability:** These issues can cause a database or one of its instances to become unresponsive and ultimately unavailable. For example: runaway queries or hangs that deny database resources and CPU to other sessions; software bugs or configuration changes.

Performance issues can be grouped into database server and client-caused categories.

- **Database Server Performance:** These issues can result in lower than optimum performance of database servers. For example: deviations from configuration best practices, bottlenecked resources, poorly constructed SQL or known bugs that can be fixed with patches.
- **Database Client-Driven Performance:** These client-side issues can also affect database performance. For example: Bottlenecks due to oversubscription of CPUs or memory, and other scenarios where client demand exceeds server or database capacity.

The above examples are not comprehensive, but they illustrate the range of operational runtime issues addressed by Oracle Autonomous Database in areas of availability and performance. For issues that require Oracle Support Service (OSS), Autonomous Database improves the interaction between customers and Oracle Support by collecting relevant information required by OSS to quickly resolve the issue.

HOW AUTONOMOUS MONITORING AND RESPONSE ADDRESSES THESE ISSUES?

Oracle monitoring and response components work with each other 24x7 in daemon mode. This integration provides a continuous, collaborative framework that addresses the above performance and availability issues.

Autonomous Diagnostics and Triage Function

- Continuously monitors database systems and over 120 key processes. It collects OS metrics and generates diagnostic views of clusters and their hosted databases in real-time to serve as a data feed for other Autonomous Database monitoring and response components.
- Establishes baseline statistics and maintains best practice configurations automatically, before and after any new installation, patch or upgrade.

- Maintains compliance with best practices and flags vulnerabilities to known high-impact issues. For example, it automatically checks and confirms that any patches released by Oracle support are properly deployed. The customer does not have to invoke this process.
- Discovers and predicts potential cluster and database problems in RAC environments, and provides corrective actions to prevent severe performance or availability issues. It uses over 150 signals, synchronized every second, between the OS and the database to get to the root cause and recommend appropriate action, such as increasing memory or CPU resources.
- Accelerates issue diagnosis, triage and resolution for the problems that involve Oracle Support Services. Analyzes all the trace files and alert logs and takes appropriate action. It uses machine learning for anomaly detection and only collects the right logs to see which files are needed for resolution.
- Discovers duplicate bugs, correlated issues and prioritizes them based on customer impact. It discovers bugs from over 400 Oracle products and uses machine learning to identify the bug – and even the developer – to help simplify or bypass the entire triage process.
- Autonomously monitors performance and manages resources to meet SLAs. It can take action to rebalance workloads based on resource usage or potential bottlenecks and can automatically decide where a new node should be deployed.
- Autonomously preserves database availability and performance by detecting and resolving database hangs. It looks at sessions that are not making progress, including other sessions impacted by the bottleneck. It identifies the session that is causing the slowdown and issues a command to quarantine or terminate the problematic session.

CONCLUSION

No databases that run on-premises or in cloud environments today are 100% autonomous – but that is the goal toward which the industry is headed. To further the evolution of cloud databases toward this true utility model, Oracle introduced the Autonomous Database service, running on Oracle Database in the Oracle Cloud. Autonomous Database minimizes or eliminates human labor using self-driving, self-securing and self-repairing functionality. Two key areas that comprise the self-repairing capabilities of the Autonomous Database are the Oracle Maximum Availability Architecture (MAA) and its monitoring and response function. Oracle MAA combines advanced technologies, best practices and autonomous monitoring and response functions to proactively protect against and rapidly recover from all unplanned and planned outages. MAA leverages a range of capabilities, based on the enterprise's SLA requirements – from automated backups to touchless zero data loss replication across any distance. Oracle MAA works with monitoring and response, which combines run-time diagnostics and root cause analysis with applied, adaptive machine learning. Autonomous Database prevents and repairs a range of database and cluster-related problems in order to maintain Oracle Database availability and performance SLAs. This collection of self-repairing capabilities offered by the Oracle Autonomous Database is unmatched by any other cloud (or on-premises) database in the industry.

-
- i The “Introduction” is intended to be common for each of the three Oracle Autonomous Database White Papers that focus on Self-Driving, Self-Securing and Self-Repairing attributes.
 - ii The “What is an Autonomous Database” section is intended to be common for each of the three Oracle Autonomous Database White Papers that focus on Self-Driving, Self-Securing and Self-Repairing attributes.
 - iii IDC Perspective, “Oracle’s Autonomous Database: AI-Based Automation for Database Management and Operations”, Feb. 2018
 - iv Verizon - 2018 Data Breach Investigation Report
 - v <https://www.healthitoutcomes.com/doc/beware-the-high-cost-of-data-center-outages-0001>
 - vi DB Maestro, “Five Ways to Prevent Database Downtime,” (<https://www3.dbmaestro.com/blog/5-ways-to-prevent-database-downtime>)

CONNECT WITH US

Call +1.800.ORACLE1 or visit [oracle.com](https://www.oracle.com).
Outside North America, find your local office at [oracle.com/contact](https://www.oracle.com/contact).

 blogs.oracle.com

 facebook.com/oracle

 twitter.com/oracle

Copyright © 2020, 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.

This device has not been authorized as required by the rules of the Federal Communications Commission. This device is not, and may not be, offered for sale or lease, or sold or leased, until authorization is obtained.

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. 0120

Oracle Autonomous Database
July 2020

