

Fault detection & resolution

Faults, including hardware failures and software errors, can and do occur in any computer system. Autonomous Databases are continually monitored for the full range of faults that might occur, and resolution is automatically launched in response.

Autonomous Database runs on fully redundant Exadata hardware that can survive hardware failures without service interruption. Oracle's Cloud Operations team automatically dispatches a hardware technician to address any hardware failures.

Software faults are often avoided through proactive health monitoring and resolution to address issues before a system is impacted.

The majority of faults are monitored and resolved automatically without requiring customers to file service requests.

Automated Optimization

Optimization of a database refers to making optimal use of the resources assigned to that database. Database optimization in an Oracle database occurs at three levels, including the service, system, and application schema level as outlined in this section.

Workload-optimized autonomous database use cases

Autonomous Database includes two workload-specific optimizations:

Autonomous Database for analytics and data warehousing (ADW)

Autonomous Database is optimized for data warehouse, data mart, and analytic workloads for efficiency and simplicity. The data is automatically stored in Hybrid Columnar Compression (HCC) format. The columnar format is automatically used in Exadata Flash Cache to accelerate analytic operations. In addition to being more space efficient, Hybrid Columnar Compression format provides more optimal data access for analytics.

Storage indexes are automatically created both on disk and in Flash to help prune out any unnecessary data from data scans (see the section below for more information). Oracle Database Result Cache is also enabled by default for all SQL statements, so workloads with repetitive SQL (such as BI dashboards etc.) will benefit from accessing the results straight from memory rather than re-executing the same statement.

Autonomous Database for transaction processing (ATP)

Autonomous Database is optimized for transaction processing and workloads that include a mixture of transaction processing and operational reporting. The data is automatically stored in a row format to speed up transaction processing where each transaction is only interested in a small number of records (often one or two records/rows).

The fastest way to find an individual record within a table is via an index, which is why ADB not only supports manually created indexes but can also take advantage of automatic indexing.

Autonomous Database relies on Exadata and uses the underlying features of Exadata, such as SQL offloading and Exadata Smart FlashCache, to meet the needs of the workload. The automatic features differ depending on the intended workload type.

Automatic system & storage optimization

Optimization of the database at the system level is completely under the control of Oracle Autonomous Database and is done without requiring input from users of the database.

Automatic schema level optimization

Autonomous Database applies a range sophisticated optimization techniques at the application schema level, such as:

- Automatic optimizer statistics
- Indexing for data integrity constraints
- Automatic storage indexes
- Automatic secondary indexing

Automatic optimizer statistics

Oracle Database uses a cost-based optimizer, which relies on statistics to determine the optimal SQL execution plan. Out of date (or “stale”) optimizer statistics is a major source of SQL performance issues. Oracle Autonomous Database eliminates the need to manually gather optimizer statistics by automatically gathering statistics in a number of ways. When data is bulk loaded into an Autonomous Database, statistics are automatically gathered as part of the load operation. During DML operations (insert, update, or delete statements), critical statistics are automatically maintained. Autonomous Databases may also use high-frequency statistics gathering jobs to adjust any stale optimizer statistics.

Data integrity constraints

Indexes are normally created to enforce data integrity. Simply defining these constraints results in indexes being created automatically. While index creation for integrity constraints is not unique to Autonomous Database, this is an important component of the overall database optimization. These constraints and related indexes are used for SQL optimization.

Automatic Storage Indexes

Optimization in Autonomous Database considers the Storage Index feature of the underlying Exadata platform. The Exadata software examines offloaded SQL fragments to determine the relationship between those predicates and data values within blocks of data held by each Exadata storage cell. The Storage Index feature, in many cases, eliminates the need for other types of indexes, especially those indexes created to support reporting workloads.

Automatic indexing

Beyond the use of indexes to ensure data integrity, additional indexing may be required for performance reasons.

Autonomous Database includes the automatic creation and management of these additional indexes. SQL statements are evaluated against existing indexes, and Autonomous Database determines whether additional indexes might be necessary for optimal performance. Autonomous Database automatically evaluates the benefits of new indexes and will test the change (automatically and independently) before implementing those indexes.

Automatic Indexing also monitors the usage of the indexes it implements. If an index is no longer useful, it will be automatically removed.

Manual schema-level optimization flexibility

Autonomous Database allows manual optimization at the schema level that will not be altered by the system. Manual schema-level optimization also includes the ability to implement even specialized indexing approaches such as function-based indexes, bitmap indexes, and composite indexes. This capability ensures customers have the widest possible range of capabilities to optimize business applications.

Simplified Administration

Autonomous Database brings advanced automation to the operation of Oracle databases. It provides a greatly simplified experience in comparison to legacy systems or other third-party cloud services. In this and the following sections, we will explore what administrative functions are automated and what controls customers have over the service.

Automated administrative functions

Autonomous Database automates virtually all administrative functions for Oracle databases, which normally consume a great deal of time and effort by database administrators, system administrators, and other IT professionals.

Customers retain the degree of control necessary to deliver database services to meet business demands, while taking advantage of this unprecedented level of automation. Oracle provides a set of robust tools to help customers more easily adopt Autonomous Database.

Database Actions

Database Actions is a web-based interface that provides development, data tools, and administration and monitoring features for Oracle Autonomous Database.

The main features include executing SQL statements and scripts, creating Data Modeler diagrams, developing RESTful web services, and managing JSON collections, as well as using the Data Load, Catalog, Data Insights, Business Models, and Data Transforms tools to load data from local and remote sources, view data in tables and views, and organize, analyze, and transform data.

Application development SDKs

Although application development SDKs are not typically used for database administration, admins often use these tools for automating common tasks or application maintenance jobs. Autonomous Database provides native support for Oracle Call Interface (OCI), ODBC, and JDBC OCI Connections. This allows Autonomous Database to support many tools on the market including all popular development tools and frameworks.

Cloud Orchestration tool integration

Autonomous Database supports common cloud orchestration tools such as Terraform, which allows users to manage, version, and persist IT infrastructure programmatically using the “infrastructure as code” model. The Oracle Cloud Infrastructure Provider (a Terraform Provider) is offered under an open-source license.

Performance analysis tools

Autonomous Database includes automatic (real-time) statistics, automatic indexing, and automatic SQL plan management, relieving developers and DBAs from common and tedious performance tuning. It also gives developers and DBAs access to the same in-depth performance analysis tools found in Oracle Database to provide insight into the performance of SQL and applications using Autonomous Database. Autonomous Database performance analysis tools also include the following:

- Cloud User Interface
- Performance Hub
- Oracle Automatic Workload Repository (AWR)
- SQL Monitor

Pay-As-You-Go Scaling

Autonomous Database resources (compute and storage) are billed on a per second basis with the minimum time being one minute. To ensure customers only pay for the resources they need, when they need them, Autonomous Database includes the ability to instantly scale system resources online in order to meet the needs of the application and business. Scaling can be done manually via the Cloud User Interface or via the REST APIs (scripting) or automated via the built-in Auto-Scale feature.

Autoscaling feature

The Autoscaling feature provides the ability to automatically scale within pre-defined boundaries in response to workload demands in real-time. The feature begins with a baseline OCPU configuration and will automatically scale CPUs up to three times the baseline when the workload increases. The baseline can be modified either manually or via REST API, and Autoscaling will operate from the established baseline. The additional CPUs are only used when necessary, and the Autonomous Database instantly returns to the baseline CPU when the workload no longer requires the additional resources.

Service scale settings

Unlike other cloud services, Autonomous Database does not require users to scale using predefined hardware shapes or configurations. Users can scale up or down the number of CPU cores and/or the storage space allocated to their configuration when needed.

The number of CPU cores and storage space is set during the initial creation of the database and can be changed at any time as desired.

Automatic scripted scaling

Advanced users might consider using the REST API to scale Autonomous Database services through automated scripts to eliminate the need for manual intervention. Automatic scripted scaling can also be used in combination with the auto-scale feature to provide dynamic scaling to better meet business needs regarding performance, capacity, and cost.

Best-in-class security

Information security has become an even more critical topic in recent years due to increases in cyber security threats and breaches. Autonomous Database is built upon the foundation of Oracle Cloud Infrastructure, which is an enterprise-grade cloud service, delivering the highest possible security standards in the industry.

Users of Autonomous Database are responsible for creating database users and schema owners. Oracle allows tens, hundreds, or even thousands of schemas in a single database. Multiple applications, application modules, or microservices can share a single database, but still achieve sufficient isolation. Of course, customers can also choose to deploy as many Autonomous Databases as needed to address the requirements of application development teams.

Application developers or development DBAs use schema owner accounts to create objects used by the application, including tables, indexes, triggers, stored procedures, etc. We will explore the role of the development DBA in greater detail in the section that follows. Customers can use the ADMIN user to reset passwords of any schema owner account when necessary.

Best practice security configuration

Systems running Autonomous Database are secured using best practices for security at each level, including virtual machines, O/S, drivers, Exadata storage, Oracle clusterware, Real Application Clusters, and Oracle Database. Autonomous Databases are continually scanned to ensure compliance with current best practice security

configuration. If anomalies are detected, changes are automatically implemented without customer intervention. Autonomous Database also includes Oracle Data Safe, which provides comprehensive tools to ensure data security.

Automatic security updates

Security fixes are automatically applied as soon as possible, normally on a quarterly basis. Autonomous Database uses the Exadata product stack, which includes security fixes at each level from virtual machine and device drivers through Oracle Database. The stack is scanned for security issues using industry leading security scanners, and fixes are integrated into the stack before being applied to the Autonomous Database. Any emergency security updates are also the responsibility of Oracle and are applied automatically.

Automatic Database encryption

Oracle Autonomous Database uses Oracle's Transparent Data Encryption (TDE) technology as a standard (non-optional) configuration. TDE provides data-at-rest encryption for the Oracle database. Database backups are encrypted as well.

Encryption key management

Encryption/decryption keys are managed automatically as part of Oracle Autonomous Database service without requiring customer intervention. Encryption keys are used to control encryption of data within the database, as well as encryption of network communication. Autonomous Database provides two options for Transparent Data Encryption (TDE) to encrypt your database:

Oracle-Managed Master Encryption Keys on Autonomous Database: By default, Autonomous Database uses Oracle-managed encryption keys. Using Oracle-managed keys, Autonomous Database creates and manages the encryption keys that protect your data, and Oracle handles rotation of the TDE master key.

Customer-Managed Encryption Keys on Autonomous Database: If your organization's security policies require customer-managed encryption keys, you can configure Autonomous Database to use an Oracle Cloud Infrastructure Vault master encryption key.

Network access control

Customers can specify an access control list (ACL) to block all IP addresses that are not in the ACL list from accessing the database. After you specify an access control list, the Autonomous Database only accepts connections from addresses on the access control list, and the database rejects all other client connections.

Auditing database access

Customers can audit database access and other activities using Oracle Unified Auditing capabilities. Customers can view audit information using the OCI Service Console under GOVERNANCE (AUDIT). Oracle Data Safe also provides comprehensive tools for assessment of end security, monitoring user actions, and meeting security compliance requirements.

Complete Data Protection

Data protection is automatically configured once an Autonomous Database is created. Oracle Autonomous Database provides robust data protection due to the underlying architecture based on Exadata and Oracle database backup/recovery capabilities.

Autonomous Data Guard

Autonomous Data Guard enables a standby database for each ADB instance. It protects data against unforeseen disaster scenarios by automatically failing over to a standby database when the primary database goes down. If a disaster were to occur, an automatic failover would be automatically triggered (no user action needed) by the

Autonomous Database when a user is unable to connect to their primary database for a few minutes. Since this is an automated action, we allow auto-failovers to succeed only when we can guarantee no data loss will occur.

High Availability

Autonomous Database includes all the built-in redundancy and resiliency features that are part of the underlying Exadata platform. Autonomous Database uses a high availability (HA) configuration as a non-optional, default configuration.

Autonomous Database extends these HA capabilities to the application tier using Oracle Application Continuity. With Application Continuity, the end-user experience is improved by masking many outages, planned and unplanned, without the application developer needing to attempt to recover the request.

Standard database backups

The standard database backup uses a weekly full and daily incremental backup strategy with a 60-day recovery window configured by default. This means that full backups are taken weekly, with incremental backups once per day. Backup of REDO logs are also included to give point-in-time recovery to any time or system change number (SCN) within the backup window.

Supplemental database backups

Customers can also supplement the standard backups as needed, such as for compliance purposes. Additional backups can be taken either manually or via a REST API and stored in Oracle's object storage.

Database recovery

The Autonomous Database Console provides a simple interface to execute database recovery without specialized knowledge. The console lists the available backups within the defined recovery window, and the customer simply chooses the desired backup to restore or specifies the desired point in time for recovery. Oracle Autonomous Database automatically performs the database recovery using the appropriate level of resources allocated to the service, such as CPU cores, storage, parallelism, etc.

Unlike traditional on-premises databases, customers only need to perform recovery in order to reset the database to an earlier point in time. With Autonomous Database, physical corruption will normally be detected by the health framework and recovery initiated (if necessary) as part of the automated repair process. Customers are not required to intervene to execute recovery to repair databases corrupted by failure of infrastructure.

Autonomous Database includes full support for Oracle's Flashback technologies, including Flashback Database, Flashback Table, and Flashback Query. Flashback is often used as an alternative to database recovery, especially for recovering individual tables, such as when a rogue user or transaction deletes or modifies data improperly.

Common Use Cases

Departmental data warehouse: improving analytics for the line of business

Customers can deploy a self-service departmental data warehouse to consolidate multiple enterprise systems, spreadsheets, and third-party data sources into a trusted, maintainable, and integrated dashboard. Integrated self-service data tools allow users to load and transform data with drag and drop, generate business models, quickly discover anomalies, and build machine learning models.

Enterprise data warehouse: simplify and modernize an enterprise data warehouse

Answer more complex questions using all data. Simplify your enterprise data warehouse to support multimodal, converged data with autonomous capabilities.

Low-code application development

Build scalable, secure, data-driven applications with minimal coding using APEX, a low-code development tool that is built into Autonomous Database.

Migrate custom applications

Optimize and extend your custom and ISV applications for improved performance, availability, and security, while reducing management complexity. Reduce total cost of ownership with autonomous operations.

Migrate Oracle applications (JDE, Peoplesoft, Siebel)

Optimize and extend your Oracle applications, such as PeopleSoft, JD Edwards, and Siebel, while reducing management complexity. Get more value from your applications while reducing cost.

Database consolidation

Reduce the cost of database infrastructure and administration by consolidating multiple databases onto a single set of computing infrastructure using fractional OCPU and GB storage allocation on a dedicated infrastructure.

Deploy mission-critical applications with autonomous services in your data center

Provide secure, governed, and high-performance self-service databases while IT controls the budgeting, capacity planning, availability, security, and governance of data.

Build real-time mixed workload applications with large-scale IoT data

Deliver billions of inserts and retrievals a second with powerful analytics for key-value transactions, such as IoT data, without the operational complexity of single-purpose databases.

Migration Tools and Services

There are multiple options available to move your data into Autonomous Database. These options are listed below.

- **Oracle Database Migration Service:** A fully managed cloud service that simplifies moving established databases from on-premises, third-party, or Oracle Cloud to Autonomous Database.
- **Oracle GoldenGate:** Oracle's data replication tool for one time data migration as well as data replication with change data capture.
- **Zero Downtime Migration (ZDM):** A service with a command line interface that you install and run on a host that you provision. The server where the Zero Downtime Migration software is installed is called the Zero Downtime Migration service host. You can run one or more database migration jobs from the Zero Downtime Migration service host.
- **Oracle Data Pump:** A database utility supporting high-speed bulk data and metadata movement between Oracle databases and Autonomous Database.
- **Oracle SQL* Loader:** A database utility to load data from external files into Oracle Database.

Conclusion

Oracle's Autonomous Database is available in the Oracle public cloud as well as on-premises in customer data centers for customers who cannot move to the public cloud. Oracle Autonomous Database automates virtually all operations DBA functions, allowing customers to focus on building and deploying applications that more effectively meet business requirements. Automation layers in Oracle Cloud automatically detect and correct issues much faster and more accurately than even the most seasoned professional can accomplish using traditional manual methods.

Oracle Autonomous Database is built on Oracle Cloud Infrastructure, which keeps systems constantly updated with the latest fixes and security patches and gives developers immediate access to the latest innovations in Oracle Database. Autonomous Database delivers the high performance and cost-effective operation customers require for their most demanding and mission-critical applications.

Resources

- oracle.com/autonomous-database
- Autonomous Database for Free
- Documentation
- Customer case studies
- White papers and reports

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 © 2022, 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

Disclaimer: If you are unsure whether your data sheet needs a disclaimer, read the revenue recognition policy. If you have further questions about your content and the disclaimer requirements, e-mail REVREC_US@oracle.com.