

Oracle Autonomous Data Warehouse: The Future of Data Warehousing in the Cloud

See how it's helping businesses get more value from their data faster.



The current global climate has brought the need for organizational agility, rapid decision-making, and resilience to the fore. It has accelerated an already rapid shift to the cloud and underscored the need for organizations to break down silos that have historically prevented them from sharing data internally and with partners and regulators.

The move to data-centric decision-making has been an imperative for many years. IDC's "[2019 Digital Business Research](#)" study found that 91% of organizations have adopted, or plan to adopt, a digital-first business strategy driven by customer experience, process efficiency, and new revenue. [McKinsey has determined](#) that data-driven business are 23 times as good at acquiring customers, 6 times as effective at retaining customers, and 19 times as likely to be profitable as a result.

Yet many companies are still hamstrung by culture, organizational silos, and outmoded information systems. Whereas two-thirds of the respondents to a [survey by the Business Application Research Center](#) said information is highly valued for decision-making at their company, 58% also said that at least half of management decisions are based on gut instinct.

In many cases, it is increasingly difficult to find a data warehouse to power business intelligence and predictive analytics, due to cost, complexity, and scalability. The legacy data warehouses that many companies run on today are expensive to operate, with on-premise hardware and software costs often exceeding US\$1 million annually. Traditional scale-up systems are also limited in their ability to support the compute and storage demands of the tenfold growth in digital data that IDC [has projected](#) will occur between 2016 and 2025.

Cloud computing has shattered these barriers, by making powerful data warehouse capabilities available to anyone. Instead of six-month installation and testing cycles, these engines can be set up in minutes and scale almost limitlessly.

The result is that a massive shift of database management systems from on-premise deployment to the cloud is under way. **Gartner predicts** that by 2022, 75% of all databases will be deployed on or migrated to a cloud platform. Cloud platforms remove much of the administrative and operational overhead of managing large datastores and support large ecosystems of third-party software and data providers. Customers pay only for what they use.

The emergence of these low-cost, highly capable cloud databases and analytics tools has also lowered the barriers to capturing and using data in a data lake or a data warehouse. Cloud storage costs are lower, pricing is scalable, and organizations can avoid much of the costly data administration overhead of on-premise platforms. **Panoply estimates** that the current adoption of cloud-based data warehouses is about 10 times that of on-premise solutions.

A Next Generation of Cloud

Not all cloud offerings are the same, however. Although cloud infrastructure is still a relatively new concept to many people, it's worth noting that the first commercial services were launched nearly 15 years ago. In those early days, cloud providers that wanted to quickly expand their range of software options often repurposed on-premise applications for their platforms. These have become their legacy systems. Amazon Web Services' Redshift, for example, is based on Postgres 8.0.2, which was introduced in 2005.

The database market has evolved significantly since then. As the popularity of specialized databases such as NoSQL, graph, and columnar has grown, cloud platform providers have responded by adding more engines to their portfolio. Google Cloud Platform, for example, has 8 discrete database systems, and Amazon has at least 10. Each is unique and requires different operating environments and skill sets. Some scale better than others, and each has unique data integration requirements. This creates administrative overhead that organizations have been moving to the cloud to avoid.



***Gartner predicts that by 2022,
75% of all databases will be deployed on
or migrated to a cloud platform.***

Recent years have also seen the emergence of cloud native databases such as Snowflake. Although these solutions provide better scalability than the early “lifted-and-shifted” legacy engines, most don’t run on premise or they integrate poorly with an organization’s existing databases. That’s bad news for the **nearly 70%** of enterprises that are running database workloads in their own data center or in private or hybrid clouds as well as the growing number of firms that are building out distributed processing networks to support the coming edge computing wave.

Even cloud native data warehouses may require extensive configuration and expertise to patch, update, and optimize. This negates one of the greatest values of cloud data warehouses, which is that they permit users to build “sandbox” environments to test ideas and experiment with analytical tools. Their ability to do so is limited if the warehouse requires extensive manual administration.





“[I] strongly recommend Exadata X8M as a platform for next generation enterprise applications.”

— David Floyer, CTO, Wikibon

The Value of Autonomy

Infrastructure-as-a-service platforms rooted in decade-old architectures can be thought of as the first generation of the public cloud. Changing customer demands have given birth to a second generation, built and optimized specifically for enterprises that need to run their most demanding workloads predictably and securely. Use of a second-generation cloud is the precursor of a digital-first strategy that enables enterprises to focus on managing their business rather than their IT infrastructure.

In 2018 Oracle introduced the concept of an autonomous database. This new design redefines database management, using machine learning and automation to reduce human labor; cut error rates; and ensure high levels of reliability, security, and operational efficiency.

Oracle Autonomous Database and Oracle Autonomous Data Warehouse are tightly integrated with the underlying Oracle Exadata Database Machine. This is a scale-out high-performance class of database servers that deliver rapid deployment, high levels of elasticity, and automated tuning and administration.

Wikibon CTO David Floyer **has called** the combination of Oracle’s X8M infrastructure and Autonomous Database “the most profound update in Oracle’s history” and “strongly recommends Exadata X8M as a platform for next generation enterprise applications.”

Among the more remarkable features of Autonomous Database is its ability to apply patches and upgrades without incurring downtime and to tune the engine for optimal performance based on the workload. Machine learning algorithms automatically optimize queries for performance with preconfigured resource profiles for different types of users. Self-tuning and integration with Exadata hardware yields performance that is up to 14 times as fast as Amazon Redshift’s, according to Oracle benchmarks.

Oracle’s converged database architecture enables one engine to run a wide variety of database workloads, including NoSQL, graph, columnar analytics, and text. Machine learning determines which processing model is most appropriate for the data.

Oracle’s second-generation cloud brings the benefits of automation, hardware integration, and massive scalability together to enable customers to deploy a full-blown data warehouse in as little as 15 seconds. Customers can double or triple the number of cores allocated to the data warehouse in seconds and scale back down when data-intensive processing is completed. Automation permits them to run powerful data analytics and build low-code applications without having to invest in data center infrastructure or additional IT staffing. Depending on their needs, they can choose a flexible serverless configuration or a dedicated hosting environment.

Autonomous Data Warehouse Cloud is also tightly integrated with Oracle’s transaction database engines, which are the most widely used in the commercial world. Cloud and on-premise Oracle data management workloads are 100% compatible with each other, enabling customers to leverage existing investments and skills in a hybrid cloud environment. In contrast, AWS Redshift customers must completely rework their code and realign their applications, because there is no on-premise version.





“In our SKY, Oracle Cloud is always there. Together, we multiply the power of information.”

— Alberto Camardelli, CIO, SKY

Autonomous Data Warehouse Cloud is open and works with all popular business analytics tools. It also includes rich data integration capabilities, the result of Oracle’s decades of experience in connecting to applications and other datasources.

Oracle has become the performance and low-cost leader in cloud data warehousing. In recent benchmark tests, Viscosity North America determined that Autonomous Data Warehouse provided a 40% price/performance advantage over Redshift.

It’s one reason SKY Brasil chose Autonomous Data Warehouse running on Oracle Cloud Infrastructure to support real-time, sophisticated marketing analytics to reach its 4.7 million customers across a vast and diverse landscape. The autonomous database was up and running in the cloud 90% faster than the previous on-premise environment. SKY Brasil has also realized 60% cost savings on infrastructure.

That means SKY Brasil’s IT team can dedicate more resources to data modeling instead of infrastructure management. Marketing and product teams can launch and scale up campaigns at any time without worrying about system availability or capacity.

“In our SKY, Oracle Cloud is always there,” says Alberto Camardelli, SKY’s CIO. “Together, we multiply the power of information.”

Automation is an important reason Unior Group, a Slovenia-based diversified global manufacturing and services company, chose Autonomous Data Warehouse and Oracle Analytics running on the Gen 2 Oracle Cloud Infrastructure to derive insights from its manufacturing operations.

“We were producing 200,000 items every day but couldn’t see where we should be focusing our business objectives,” says CIO Rok Planinšec. “Getting insight was taking months.”

The Oracle products’ self-tuning and self-patching capabilities now make it possible for Unior to manage 85 million data records with no database administrators.

“From the moment we started using Oracle Autonomous Data Warehouse, we were back in control of our production and our future,” Planinšec says.

As customers become savvier about data warehouses in the cloud, more and more of them are turning to cloud native solutions that combine their nearly limitless scale with operational simplicity and features that leverage their existing skill sets. Oracle Autonomous Data Warehouse is the future of data warehousing in the cloud.

To learn more about Oracle Autonomous Data Warehouse, [go here](#).

