THE FUTURUM GROUP

The Enterprise-wide Benefits of a Comprehensive Data Platform

Why the Oracle Data Platform delivers a compelling end-to-end data management environment for the entire enterprise

AUTHOR

RON WESTFALL
Senior Analyst + Research Director

IN PARTNERSHIP WITH

ORACLE®

APRIL 2023

Section I: Data platform requirements and the challenges with current approaches

In recent years many cloud database providers have been touting their variations of a “data platform” and its perceived benefits for organizations of all sizes. From our view, the following high-profile examples that merit additional scrutiny include:

- **Snowflake: One Platform That Powers the Data Cloud**
  Snowflake asserts that you should execute your most critical workloads on top of Snowflake’s multi-cluster shared data architecture in a fully managed platform that capitalizes on the near-infinite resources of the cloud.

- **AWS offers a wide range of data platforms, just like their cloud databases, that combine AWS and third-party offerings.**
  Specifically, Industrial Data Platform (Consumer Packaged Goods, Manufacturing), Customer Data Platform (Retail, Media & Entertainment), Big Data Platform; others offered as services through AWS or partners (Streaming Data Platform, Data Security Platform, etc.)

- **MongoDB: Developer Data Platform**
  MongoDB touts that it offers the developer data platform that provides the services and tools necessary to build distributed applications fast, at the performance and scale that users demand.

- **Microsoft: Intelligent Data Platform**
  Microsoft spotlights moving beyond the cost and complexity of point solutions with a unified data and AI platform. Adapt rapidly, add layers of intelligence to apps, generate predictive insights, and govern all your data—wherever it resides.

- **Google: Cloud Architecture Framework**
  Google Cloud emphasizes that businesses want to analyze data and generate actionable insights from that data. Google Cloud aims to provide customers with various services that can help them through the entire data lifecycle, from data ingestion through reports and visualization.

What do these data platforms/architectures have in common? They speak of “seamless” data movement from one data model to another; they refer to the ability to support “any type of data” and “any type of workload” at any time and in any location. They are referred to as a nearly universal panacea that these vendors are uniquely equipped to deliver for customers today. However, when we look underneath the covers, it becomes rather clear that the most vocal proponents of a “data platform” have not arrived at a complete data platform. For instance, we identify that they are providing only fractional and insufficient solutions for what organizations require to collect, process, analyze, and store their data—and to make intelligent, fully explainable, automated predictions from it. “Data Platform” can be a fundamentally misleading premise and sales and marketing term in the industry.

Cloud-based data platforms are intended to help organizations get more value from their data, share information with everyone that needs it, and reduce costs. The theory is that they provide more integration and lower costs than on-premises environments where application and data silos lead to low utilization rates and high expenses. Unfortunately, across the solutions we examined, theory doesn’t match reality with cloud-based data platforms because they typically are almost as inefficient as on-premises deployments. Why? Because they run many different services, each of which is overprovisioned to avoid scaling-driven downtime. The fact that we see that they don’t work well with many business-critical applications doesn’t help either.

Each vendor offers solutions that they call data platforms, but we find that they are remarkably different. Many are limited in scope, focusing on a single portion of an organization’s end-to-end data needs (e.g., Snowflake). We’ve seen environments that only focus on analytics with relational data that don’t address data integration, transactional databases, and high availability capabilities required by many crucial applications. As a result, IT organizations have the unenviable task of trying to patch together a single-purpose environment with other services to meet their organization’s end-to-end requirements.

At the other end of the spectrum are data platform offerings which have an enormous set of piece parts that must be integrated in different ways to meet the needs of each individual workload (e.g., AWS). As the do-it-yourself (DIY) integrator, you're left to wonder whether this is a good use of your time and whether you've put them together correctly. It's like building a kit car and having an extra bushing left over—did the supplier give me an extra part or did I forget to install it? You will only find out after a lot of time and effort. Both the limited availability and piece-part approaches require the consumption of numerous services, extensive integration, lots of internal resources or highly paid system integrators, and fragmented management that drives total costs higher.

We find that what organizations of all sizes and in all industries are looking for is ultimately both more and less. They want more capabilities than the limited functionality data platform, and they want less complexity and costs of DIY integration required by both approaches. They would prefer to use a single data platform that meets their requirements to integrate and consistently manage data from across the enterprise.

From our perspective, a data platform should be a comprehensive, open, and integrated suite of services that meets the end-to-end data needs of an organization's portfolio of applications and use cases. It should include high levels of automation that make it easy to manage and use while providing high performance and security. Accordingly, a successful data platform must deliver across three main design principles:

- **Comprehensive:** The data platform should provide a comprehensive set of capabilities that support any app, any use case, and data type. It should comfortably handle OLTP, analytics, machine learning (ML), and mixed workloads with relational, document, spatial, graph, and other types of data. Also, their management should be as automated as possible to allow scarce IT resources to focus on innovation that improves enterprise business outcomes.
- **Open:** The data platform should be modular, enabling organizations to use its comprehensive set of capabilities as-is and extend or replace core elements with services developed by other suppliers, the open-source community, or their own development organization. It does not lock customers into the capabilities offered by a single vendor, but rather allows them to easily add new functionality as business needs evolve.
- **Integrated:** The data platform should minimize the number of separate services, integration points, and data transformations required to meet an enterprise's full portfolio of workflows. Only when it does this will a data platform be able to help reduce complexity, improve operational efficiency, lower costs, and increase access to data from across the enterprise.

Section II: Meeting Enterprise-wide Challenges with Oracle Data Platform

Oracle Data Platform brings together a comprehensive set of proven data services in an open, modular, and integrated fashion to help organizations meet their end-to-end data needs. In our assessment, Oracle Data Platform meets the three critical design principles required for success and provides the high performance and availability customers need to run business-critical applications. It reduces complexity by integrating versatile services for data integration, transactional and analytical databases, data lakes and lakehouses, analytics, and application development. Furthermore, its open and modular approach allows organizations to select which core components they want to use, incorporate third-party capabilities, or develop their own.

The complete Oracle Data Platform for Enterprise-wide Data and Applications

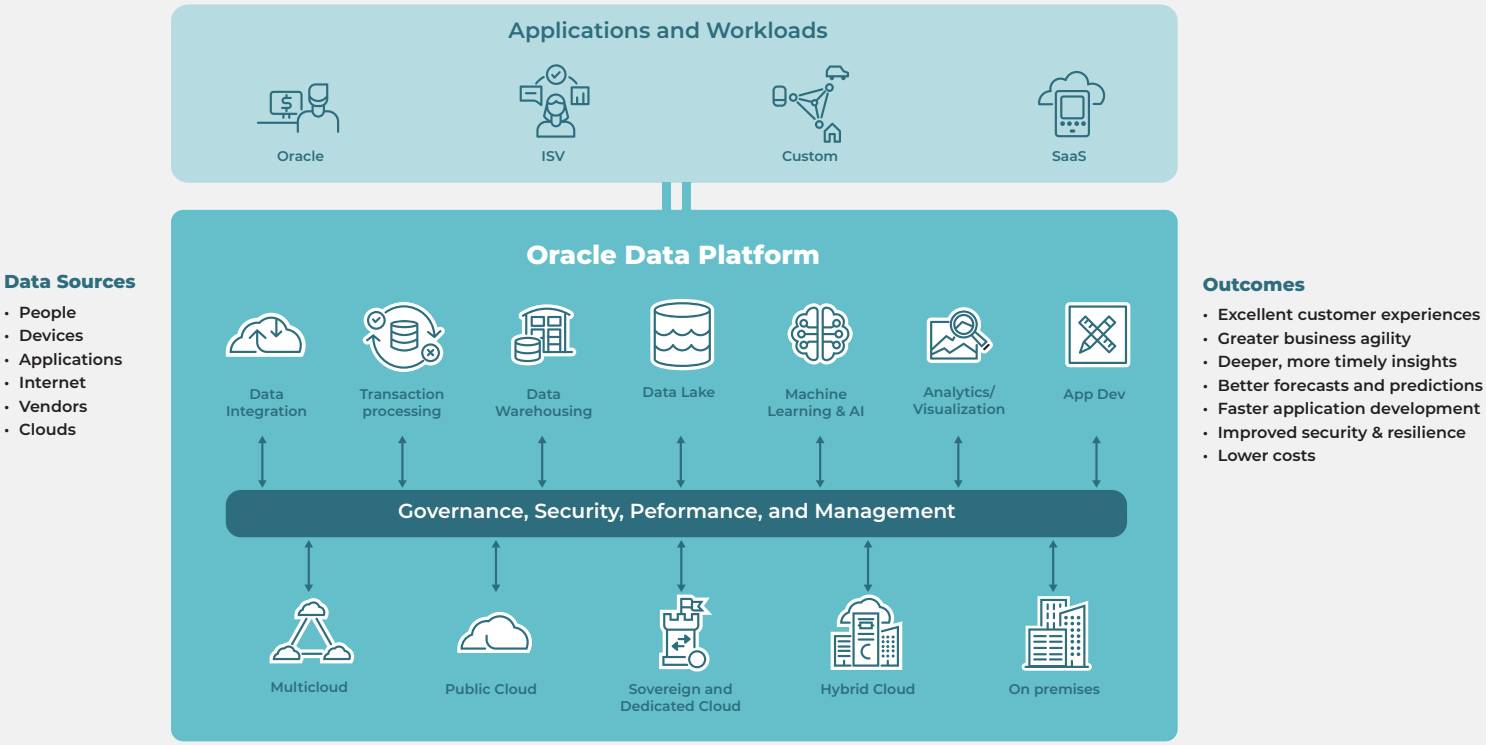

A critical component of Oracle Data Platform is the converged database capabilities of Oracle Autonomous Database which supports multiple types of data, workloads, and development environments with built-in AI/ML in a fully automated database service. It helps organizations consolidate data from across the enterprise and share it with everyone that needs it. Autonomous Database also automates many manual management tasks such as database provisioning, tuning, and scaling—freeing scarce resources to help the enterprise move their business objectives forward. Under the covers, Autonomous Database runs on Oracle Exadata systems in Oracle Cloud Infrastructure (OCI) and enables extremely efficient consumption of database resources.

We'll go into this in more detail later in this analysis, but Autonomous Database on dedicated Exadata systems in OCI enable organizations to run data platform workloads with extreme performance, scalability, and versatility. Exadata systems include more than 60 unique capabilities for Oracle databases that allow them to run OLTP, analytics, and consolidated workloads with the highest possible performance. In addition, Exadata systems in OCI use AMD EPYC™ processors to provide high parallelism and performance for database workloads. These systems start with as few as 252 available database processing cores and can scale up to more than 4,000 cores to support any sized OLTP, analytics, and consolidated workload environment. We see that the unique ability of a single system to deliver industry-leading capabilities for the most performance and scale-sensitive enterprise workloads shows that Autonomous Database on AMD EPYC processor-powered Oracle Exadata systems provide the most versatile engine to support a comprehensive data platform.

However, performance isn't everything. Enterprises prioritize maximizing value, and our analysis shows that Oracle Data Platform offers a compelling value proposition for all organizations. It provides comprehensive, easy-to-use capabilities that include data integration OLTP, data warehouses, data lakes or lakehouses, AI/ML processing, and analytics. Oracle Data Platform allows organizations to meet their end-to-end data needs with fewer services, dramatically reduces the number of integration points, and with Autonomous Database's automated tuning, scaling, management, and security management reduces both administrative and consumption costs.

Reviewing the Competitive Landscape

Let's examine a couple of the data platforms listed above to pinpoint how we see they compare to the complete Oracle Data Platform.

Snowflake

Snowflake claims that it delivers a "single, global platform that powers the Data Cloud" and that "Snowflake is uniquely designed to connect businesses globally, across any type or scale of data and many different workloads and unlock seamless data collaboration." However, Snowflake has realized that customers want more than just running analytics on stale data. They want to analyze real-time data as transactions occur. In the last six months the company attempted to add a fledgling OLTP capability to its service called "Unistore," which supposedly unites transactional and analytical data together in a single platform. This is still unavailable, and our understanding is that customers don't want to develop new apps on a proprietary platform. Once developed there, it's stuck there. It could end up being even worse than paying exorbitant data egress fees. The fact is we view that Snowflake still lacks support for graph, spatial, document, key value, blockchain, and has no built-in machine learning. The company is just now trying to address ML by embarking on the acquisition of Myst AI for time series forecasting.

We believe that what Snowflake really means by "cloud data platform" is that organizations can obtain a small list of still-not-integrated capabilities that don't come close to a complete data platform. Moreover, since Snowflake lacks a converged database, moving from transactions to analytics (Amazon Aurora to Snowflake for example) requires an Extract Transform Load (ETL) process, which incurs additional data movement costs and the inherent security risks via surface area exposures. In other words, it's not free and takes time—and poses unnecessary risks.

Moreover, Snowflake doesn't offer customers deployment choices. It's only available in the public cloud and not on their own cloud either—customers can run Snowflake on AWS, Google Cloud Platform, or Microsoft Azure with their own separate cost structures and licensing agreements. Of course, it's not available on-premises or in hybrid deployments. Customers whose data cannot go in the public cloud for security or locality regulations are out of luck.

AWS

From our perspective, AWS hasn't yet met a cloud database it doesn't like, showing how far the lack of a converged database can go. Hence, 16 cloud database offerings—each specialized to support a singular data model: For instance, Aurora for transactions, Redshift for data warehousing, DocumentDB for document, DynamoDB for NoSQL, RDS for relational, Neptune for graph, Timestream for time series, SageMaker for ML—you get the idea. Plus, AWS offers a whole slew of ancillary services that store data, connect, and create workflows across these specialized databases—all separately chargeable and requiring expertise, time and effort to execute.

In addition to these disparate cloud databases, the data platforms that AWS offers are unintegrated combinations of AWS services and third-party applications available through the AWS Marketplace. It's up to the customer to integrate the wide array of pieces into a solution.

Although AWS' focus is on the public cloud, we notice that the company made what can be characterized as a half-hearted attempt at an on-premises offering with AWS Outposts. For example, AWS' flagship cloud databases—Aurora and Redshift—are not available on Outposts. Once again, customers whose data can't go to the public cloud are out of luck.

Snowflake and AWS

We observe that both Snowflake and AWS demonstrate the limitation of their approaches to architecting and delivering an actual complete data platform. They profess to offer customers unlimited flexibility and choice—but when it comes to deployment choice it's: "Use my service in a public cloud, and if you need anything that I don't offer natively, you are welcome to access those services only if they run in the same public cloud and you sign up for them, but they control their prices—we don't."

All things considered, it's less of a "seamless" data platform in the cloud and more of an aggregation of unrelated elements under a single marketing umbrella designed to masquerade as a so-called complete solution.

Oracle Cloud Infrastructure and the Complete Oracle Data Platform

To start, we note that Oracle is the only hyperscaler cloud provider to offer a full stack solution from infrastructure to database and enterprise-grade, AI-injected SaaS applications using a common data model. When it comes to the data platform Oracle provides these key attributes:

- **Most Comprehensive:** Oracle Database combines every standard database type and model into one all-inclusive converged database with a single data copy. This enables all database data to be addressable by every type of database including OLTP, OLAP–data warehousing, key-value, JSON–document, XML–object, time series, spatial, graph, ML, data lakehouse, and blockchain. No data needs to be replicated, stored on duplicate storage and put through ETL into a specialized database type to be queried. Oracle Databases also come with Oracle APEX so organizations can develop secure, high-performance low-code applications with no extra cost or integration. Equally, for open source, MySQL HeatWave combines transactions, analytics, ML, and lakehouse capabilities into one database.

- **Complete Integration:** Oracle's portfolio goes beyond simply connecting different database types to each other with connectors and prepackaged ETLs. From our perspective, being the most integrated means that different database types are fused together synergistically at the source code level. There are no ETLs or multiple data copies. Data is accessible and capable of being queried by different database types without being copied, moved, or ETLed. For example, JSON data can be queried through the JSON API, MongoDB API, or SQL. Additionally, for open source, MySQL HeatWave includes transactions, analytics, ML, and lakehouse in one database—no ETLs required. Data is analyzed in real-time and as such there is no stale data.

- **Most User Friendly:** Offers all-encompassing simplicity for all developers and users whether they be novice, average, or expert.

- **Thorough Security:** We observe that Oracle is the industry security standard that others hope to achieve with data platform services in the cloud. Fundamentally, it goes beyond simply encrypting the data to including automated intrusion detection and threat remediation, end-to-end encryption, hardened multi-tenancy, not allowing suspicious data into the database, identity management, and much more. Specifically, the Zero Data Loss Autonomous Recovery Service even provides the ability to protect and then recover Oracle Database data up to the last transaction to mitigate any outages, errors, or cyberattacks that occur.

- **Robust Automation:** From our view, Oracle's complete data platform offers extreme levels of automation. Oracle Autonomous Database is completely self-driving, self-tuning, self-optimizing, and self-securing. It enables databases to automatically scale up to meet workload requirements and scale back down to reduce costs—all without impacting availability. It's designed to make the database administrator's (DBA) life much easier. MySQL HeatWave automates many of the most time-consuming admin tasks through MySQL Autopilot as well.

- **Highest Performance:** Oracle's database portfolio performance claims are consistently supported by real-world published and easily repeatable benchmarks (i.e., GitHub) and industry-hardened in the most critical business environments around the world.

- **Most Available:** We find that Oracle Autonomous Database and Exadata Database Service are the only cloud platforms to provide the high levels of availability afforded by Oracle Real Application Clusters and Exadata infrastructure. Component and node-level failures are transparent, as are servicing, upgrades, patching, and scaling—all of which incur downtime on other platforms. MySQL HeatWave also enables transparent scaling that does not impact availability.

- **Broadest Deployment Choices:** Oracle Database is available on virtually every platform imaginable—even AWS RDS. Autonomous Database is available on OCI and in customer data centers through Oracle Exadata Cloud@Customer and OCI Dedicated Region. MySQL HeatWave is available on OCI, AWS, and Azure as well as OCI Dedicated Region. Oracle databases can even be accessed from applications running in 12 different Azure regions, delivering what are identified as the highest levels of performance and availability with no egress fees.

- **Best Price/Performance and ROI:** Independent analysis has consistently demonstrated that customers using Autonomous Database and MySQL HeatWave gain significant savings and the highest return on their investments (i.e., GitHub).

There are sales and marketing assertions and then there is production hardened reality backed by thousands of satisfied customers. We believe the complete Oracle Data Platform fits into the latter category.

The Benefits of Using Oracle Data Platform for All Your Workloads

As we see it, Oracle Data Platform's alignment with the core design principles of comprehensiveness, openness, and integration enables it to meet the full range of customer application and data requirements. It helps organizations achieve a diverse set of benefits that include increasing application uptime, providing better customer experiences, and improving the accuracy of business forecasts.

We advocate that Oracle Data Platform is also the best approach to provide the business agility and levels of innovation that today's enterprises desire. The platform's ability to integrate data from diverse sources and other clouds, process transactions quickly, run analytics on transactional data, and powerful yet easy to use analytics enable new insights and efficiencies at all levels of the organization.

We also find that Oracle Data Platform offers greater performance and resiliency for mission-critical applications than other so-called data platforms. In particular, Autonomous Database's self-managing capabilities combined with Exadata's built-in fault tolerance and advanced disaster recovery capabilities provide availability SLAs of up to 99.995%, which are head and shoulders above single-purpose, non-integrated databases and data platforms running on generic cloud infrastructure.

As a result, Oracle Data Platform delivers solid business-centric differentiation for customers across different types of workloads, from OLTP to data warehousing and data lakes/lakehouses to analytics. Oracle's solution is further differentiated by its database platform optimization, global availability and price consistency, and the ability to support both traditional packaged applications and modern, microservices-based applications. Taking into account these capabilities, Oracle Data Platform goes a long way toward fulfilling the end-to-end data needs of an enterprise.

Section III: Getting More Value and Simplifying Operations with Data with Oracle Autonomous Database

Integral to the Oracle Data Platform vision is delivering core data management capabilities using Oracle Autonomous Database, MySQL HeatWave, and Oracle Big Data Services for Hadoop-based data lakes. Given the broad usage of Oracle Database in organizations around the world, it should be no surprise that Autonomous Database will provide this core capability for many organizations, so we're going to focus on it for the rest of this analysis.

Autonomous Database is the combination of Oracle Database, AI/ML-based automation, and Exadata systems running in OCI regions (41 and counting), Dedicated Region, and on Exadata Cloud@Customer systems located in enterprise data centers or co-location facilities such as Equinix. It combines the best converged database capabilities to support all data types, workloads, and development styles with the best database automation. Autonomous Database enables organizations to efficiently run databases workloads of any type, scale, and criticality.

We expect that most Oracle Database customers will eventually use Autonomous Database in OCI or hybrid cloud environments since it automates performance optimization with automated tuning and scaling; management with automated patching, upgrades, and maintenance; and security with proven built-in capabilities and automated implementation of security best practices. You can build an Oracle Data Platform solution with Exadata Database Service, but we discern the Autonomous Database clearly provides more value.

Several key Autonomous Database benefits that everyone should be aware of include:

- **Pay-per-use economics:** Autonomous Database offers built-in, online scaling of database consumption so that you can let it automatically adjust resources to match workload needs. Autonomous Database looks at the queries it's executing and then scales consumption up when demand peaks and back down to save money when demand subsides. When you combine this with per-second consumption billing, customers have what amounts to a pay-per-use consumption model. In addition, Exadata's extreme performance equates to faster completion of tasks and lower costs. In the cloud, time is money.
- **High availability:** Autonomous Database provides built-in capabilities that enable it to provide mission-critical high availability in OCI with up to a 99.995% uptime SLA—higher than what's available from virtually all cloud-based data platform providers. It monitors thousands of metrics and uses ML to identify and resolve potential issues before they impact customers. If something does go wrong, built-in capabilities such as redundant servers and storage with automated fail over and Autonomous Data Guard will help keep databases running.

- **Simplified, automated security:** Autonomous Database encryption is always on, there is strict separation of roles when managing databases, and in dedicated environments there is isolation of data and database processes from other users. Included with Autonomous Database is Oracle Data Safe, which provides a centralized monitoring and auditing of database security capabilities, redacting sensitive data, and performing both user and configuration risk assessments for both cloud and on-premises Oracle databases.

In addition to these capabilities, Autonomous Database on dedicated Exadata systems in OCI helps improve developer productivity by eliminating the need for them to manually provision and secure databases or worry about excess cloud consumption charges. Autonomous Database provides developers with a fully self-service development environment where security and maximum consumption levels are managed centrally, but they have full independence to perform database lifecycle operations for their databases.

The Benefits of Using Oracle Autonomous Database in Oracle Data Platform

We believe Autonomous Database provides the critical database capabilities that organizations need in a comprehensive data platform. Consolidating databases with Autonomous Database allows organizations to create a single source of truth that empowers data-driven decisions at all levels of the organization, enables deeper business insights, and improves operational efficiency.

Moreover, we think that being able to process that data with advanced analytics and in-database AI and ML algorithms will lead to innovative new ways to create even greater business insights and differentiated customer experiences.

Finally, we see Autonomous Database's making IT organizations less dependent on scarce expertise, reducing the chance of human error, and avoiding the risks from downtime and cyberattacks.

Section IV: Exadata Cloud Infrastructure: Optimizing the Impact of Autonomous Database

Autonomous Database runs only on Exadata systems and benefits from their unique approach to delivering scalable performance. First deployed over 13 years ago, Exadata was one of the first database platforms to separate database compute and storage resources – something that other data platform vendors have recently been touting as freshly innovative. More importantly, Exadata was the first to distribute database functionality across both scale-out database servers and scale-out intelligent storage servers, liberating database servers from processing low-level SQL statements for analytics queries so that they could process more OLTP queries and intricate analytics.

Exadata systems in OCI power the versatility and automation of Autonomous Database with leadership performance, scale, and availability for OLTP, analytics, AI/ML, and consolidated workloads. Coupling this with the ability to independently scale database compute and storage makes Exadata systems the most versatile way to support a comprehensive data platform.

For OLTP, Exadata systems deliver industry-leading 19 microsecond SQL Read IO latency, which is critical for maximizing application performance—vastly superior to the "sub-millisecond" and "low single-digit millisecond" latencies reported in other cloud database environments. The inclusion of AMD EPYC processors in Exadata database servers equips systems with 2.5x the number of cores and 1.8x the number of SQL Read IOPS than previous generation solutions. These increases mean that database workloads can process more simultaneous transactions and easily scale consumption to meet demands without impacting the performance of other databases.

Similarly, Exadata supports analytics workloads of any scale with database scan rates that scale linearly as you add intelligent storage servers. Entry-level configurations provide over 150 TB of usable database capacity that is backed by triple redundant storage, with a scale-out system providing over 3 PB of usable database capacity and 30PB of data warehousing space. Exadata's ability to offload data-intensive scan processing to storage servers delivers more than 200 GB/second of throughput on entry-level configurations and nearly 3 TB/second of throughput in scaled-out environments. This provides orders of magnitude more throughput for business analytics and ML model training than platforms that require all data to be loaded from generic cloud storage to generic cloud servers over generic cloud networks before it can be analyzed.

Furthermore, the availability of numerous AMD EPYC processor cores supports complex business, graph, and spatial analytics along with AI/ML models running on the Exadata database servers. High levels of parallelism allow organizations of any size to easily employ these advanced technologies without increasing the number of different services they must integrate, secure, and manage.

Finally, the expansive nature of Exadata systems and their ability to run all database workloads with all data types make it the perfect platform for database consolidation. Database consolidation is critical in the cloud because it enables greater data sharing and increased operational efficiency—two goals for anyone looking to implement a data platform. This stands in stark contrast to clouds like AWS which currently offer 16 different database services which organizations must attempt to integrate to meet their end-to-end data needs. The use of stand-alone databases forces data integration into the application layer, increasing application complexity, requiring ETLs and a multitude of specialized skills, effectively making it more difficult to administer the fragmented management and security landscape.

The Benefits of Exadata Cloud Infrastructure with AMD EPYC Processors

From our perspective, implementing Autonomous Database on Exadata Cloud Infrastructure delivers not only substantial performance improvements but also equally large economic benefits to customers. We see the higher levels of database consolidation enabled with Autonomous Database on Exadata as reducing the infrastructure requirements and associated costs needed to support a database fleet. Within this context, we view Oracle's selection of AMD EPYC processors for powering Exadata database servers in OCI as a major contributing factor in optimizing these capabilities. Some of the business-oriented benefits that we see accruing for Oracle Data Platform customers from this approach include:

- **Streamlined Operations:** Exadata's scale-out design leverages AMD EPYC processors to process more concurrent transactions and faster analytics. This allows organizations to meet Line of Business demands without incurring downtime or excessive costs and to achieve greater business agility when deploying new applications. For example, retailers and travel operators can more easily meet seasonal spikes in demand, banks can detect fraudulent transactions, and manufacturing companies can integrate real-time IoT data to predict equipment failures and maintain continuous operations.

- **Environmental Gains:** Autonomous Database on Exadata Cloud infrastructure allows organizations to run their database fleets with fewer servers and less energy consumption. The use of AMD EPYC processors enhances environmental gains by allowing even more databases to be consolidated on fewer database servers, and AMD EPYC processors consume less power per core than previous CPUs used in Exadata, improving the energy savings. Furthermore, integrated features such as graph, spatial, and AI/ML processing results in fewer servers being used to support otherwise stand-alone services. These gains are magnified in OCI due to Oracle's commitment to use 100% renewable energy in OCI Regions by 2025, a goal they've already achieved in their European OCI Regions.

- **Lower Costs:** The same database consolidation and efficiency factors that lead to environmental gains are also key factors in reducing costs. We believe that in the cloud, where per-second consumption costs are fixed, running workloads faster means that you're saving money. Exadata enables this with its high performance—and Autonomous Database takes savings to another level with the automatic scaling of resource consumption.

- **Enhanced Customer Experiences:** Today's consumer and business users don't want to wait. They want a comprehensive set of information presented quickly so it's easy to make decisions. Exadata systems in OCI help organizations deliver better customer experiences because they can rapidly provide individualized and context-sensitive information derived from multiple types and sources of data. We see customer experiences being enhanced with recommender systems and real-time analytics in location-based advertising, financial services, and EdTech environments.

- **Deeper and More Timely Insights:** When people have more computing power available for analytics, they either run the same analyses faster, run them on more data, or use more computationally- intensive approaches to uncover new insights. We believe that in most cases, analysts and business users want to do all of these at once. Exadata's impressive performance and scalability make running analytics faster and with more data a given. In addition, the scale-out Exadata architecture and high core-count AMD EPYC processors allow organizations to generate new types of insights using advanced analytics and AI/ML capabilities without adding complexity to their IT environment. We've seen these types of capabilities used in everything from fraud prevention to utility demand forecasts, and equipment failure predictions in financial services, telecommunications, and manufacturing environments.

From our perspective, even a small Exadata configuration with 252 AMD EPYC processor cores in database servers will enable customers to affordably consolidate a wide array of database workloads. This is enhanced with independent scaling of database compute and storage resources so customers can configure cost-effective solutions to match their current needs and incrementally grow them as demands change. Moreover, we find that Autonomous Database on dedicated Exadata systems in OCI sets up Oracle Data Platform to power innovation across the entire database ecosystem.

Section V: Beyond Oracle Data Platform

While Oracle offers a compelling set of data platform capabilities and benefits, we would be remiss not to mention other Oracle offerings that extend the value of Oracle Data Platform.

From an overall cloud perspective, we have witnessed the level of functionality available in OCI grow into a robust offering that meets organizational needs for full-stack solutions and provides leadership capabilities in a number of critical areas. Oracle's growth rate in terms of adding new cloud functionality is unrivaled in the industry and includes multi-cloud connections with Azure and AI/HPC capabilities with NVIDIA.

For instance, MySQL is the leading open-source database and Oracle MySQL HeatWave has burst onto the scene with built-in automation, industry-leading in-memory query acceleration and analytics, and integrated AI/ML capabilities. In OCI, MySQL HeatWave runs on AMD EPYC processors and delivers performance that is 6.5x to 7x faster than Amazon Redshift and Snowflake at a fraction of the cost. MySQL HeatWave is an integral part of Oracle Data Platform, both for its accelerated in-memory capabilities as well as an open-source based core database service.

The Oracle VMware Cloud Solution has also taken center stage as customers look to move workloads that are virtualized on premises with VMware to the cloud. The Oracle solution, which also runs on AMD EPYC processor-powered servers, makes it easy and affordable to move workloads to OCI. We note that there are many organizations that use both Exadata and VMware in their data centers, and they can easily replicate a similar environment in OCI with only minor adjustments and without retraining key staff.

Other areas where we've seen Oracle take a leadership position in deploying high-performance AMD EPYC processor-based solution are with their Big Data Service for Hadoop and Spark processing; bare metal and flexible VMs used for running packaged, open source, and self-developed applications; and for high performance computing with or without NVIDIA GPUs. In addition, we believe that Oracle's support for confidential computing, which uses application-transparent capabilities that are built into AMD EPYC processors to isolate and encrypt data and applications, will provide greater trust and peace of mind in today's security conscious world.

Section VI: Conclusion and Recommendations

In summary, we believe today's jumbled data platform landscape does not adequately support customer's end-to-end needs. The Oracle Data Platform's comprehensive, open, and integrated approach uniquely helps organizations meet their end-to-end data needs.

The key takeaways on data platform for IT decision-makers include:

- **Oracle Data Platform is Comprehensive:** Oracle Data Platform provides a full suite of open, modular, and extensible capabilities from data integration through OLTP, data warehousing, data lakes and lakehouses, AI/ML, and analytics. It can be used with any type of data, helps integrate them into a single source of truth for an organization, and supports deployments in customer data centers with OCI Regions and with hybrid cloud solutions. As a result, organizations can meet their end-to-end data needs with fewer services, less integration, and reduced management.

- **Oracle Autonomous Database Offers Differentiated Capabilities and Greater Value:** Many Oracle Data Platform customers will use Oracle Autonomous Database as their core database, allowing them to take advantage of automated provisioning, tuning, scaling, maintenance, and security capabilities. Organizations will benefit from higher database performance, lower costs, greater availability, scaling flexibility, and enhanced business resiliency.

- **Exadata systems Power a Highly Differentiated OLTP and Analytics Database Cloud:** With AMD EPYC processors in Oracle Exadata database servers in OCI, customers have 19 microsecond SQL Read latency and 2.5x more cores for more transactions with 1.8x more SQL Read I/Os than the previous-generation Exadata system combined with scale-out capabilities providing enhanced OLTP performance for more simultaneous transactions. In addition, Exadata provides extreme database scan rates that can scale to nearly 3 TB/second, providing orders of magnitude more throughput than database services running on generic cloud platforms.

We recommend that database and IT decision makers rank Oracle Data Platform high on their consideration list when looking at data platforms. They should compare Oracle Data Platform to data platform offerings from AWS, Microsoft, Snowflake, and others and see how they stack up in their ability to support diverse workloads with high availability. We also suggest that they carefully evaluate where they need to run their workloads to determine whether the platforms that they are considering can meet their data residency and isolation requirements.

We also suggest that Chief Sustainability Officers and those considering the environmental impacts of their decisions carefully consider how Oracle Data Platform and OCI can help them meet their environmental impact goals with more efficient database consolidation and Oracle's commitment to use 100% renewable energy in OCI by 2025—a goal they've already achieved in Europe.

IMPORTANT INFORMATION ABOUT THIS REPORT

CONTRIBUTORS

Ron Westfall
Senior Analyst + Research Director

PUBLISHERS

Daniel Newman
CEO
The Futurum Group

Shelly Kramer
President + Chief Evangelist
The Futurum Group

INQUIRIES

Contact us if you would like to discuss this report and The Futurum Group will respond promptly.

CITATIONS

This paper can be cited by accredited press and analysts, but must be cited in-context, displaying author's name, author's title, and "The Futurum Group." Non-press and non-analysts must receive prior written permission by The Futurum Group for any citations.

LICENSING

DISCLOSURES

The Futurum Group provides research, analysis, advising, and consulting to many high-tech companies, including those mentioned in this paper. No employees at the firm hold any equity positions with any companies cited in this document.

ABOUT ORACLE

Oracle offers integrated suites of applications plus secure, autonomous infrastructure in the Oracle Cloud. For more information about Oracle (NYSE: ORCL), please visit us at **oracle.com**.

ABOUT THE FUTURUM GROUP

The Futurum Group provides in-depth research and insights on global technology markets using advisory services, custom research reports, strategic consulting engagements, digital events, go-to-market planning, and message testing, along with creating, distributing, and amplifying rich media content where all stakeholders read, watch, and listen. Our industry analysts conduct comprehensive qualitative and quantitative assessments of technology solutions, business issues, market drivers, and end-user demand dynamics across industry sectors. We work closely with your Analyst Relations, Product, and Marketing teams to establish and amplify your subject matter expertise and thought leadership while furthering your brand and product recognition by creating, distributing, and amplifying rich content where all stakeholders read, watch, and listen. **Read our disclaimer statement here**.

CONTACT INFORMATION

The Futurum Group I **futurumgroup.com** and **futurumresearch.com**
720-258-5591 I info@futurumgroup.com
Twitter: @**TheFuturumGroup** and @**FuturumResearch**