

ESG SHOWCASE

Extracting Maximum Value from Your Data with Oracle Autonomous Data Warehouse

Date: December 2020 **Authors:** Mark Peters, Principal Analyst; and Mike Leone, Senior Analyst

ABSTRACT: Modern cloud data warehouses are prerequisite foundations for a data-driven future. They're being embraced because of their invariably easy on-ramps, agility, and limitless scale. *But* how can organizations most effectively leverage these modern cloud data warehouses to better address the real-time needs of their business, sharpen their insights, improve their predictive capabilities and ultimately, gain competitive advantage? The answer is to put a heavier focus on autonomous operations, reliability, a perfect cost model, and an ability to support what's next, whether that means broadening end-user access, integrating more data, or incorporating data science. Oracle is doing all that; going beyond “vanilla” cloud data warehousing and adding innovation to make sure that organizations get what they want, when they want it, and at the right price.

Overview

Enterprise data warehouses (EDWs) have served as the foundation of insight-driven organizations, delivering timely analysis and reporting of structured data, handling large analytics workloads, and supporting the high levels of concurrency that organizations demand. But challenges with architectural rigidity, complexity, and cost often force organizations to re-think their EDW strategies and to instead pursue modern data warehouse approaches anchored in the public cloud. These approaches promise to reduce data center footprints, increase organizational agility, improve operational efficiency, and right-size costs based on usage.

But—as early adopters of these technologies are finding out—not all cloud-enabled modern data warehouses are created equal. In fact, when ESG asked research survey respondents about the likelihood of making a change in their approach to data warehousing—specifically, if they would be changing their vendor—38% of organizations said they were considering changing to, or were evaluating, alternate solutions.¹ While that percentage encompasses both those who are leveraging a traditional EDW on-premises and those in the cloud, it speaks volumes about the drive to find the right platform...since these platforms are ones upon which organizations are increasingly staking their success.

Enter the Oracle Autonomous Data Warehouse

Oracle's Autonomous Data Warehouse (ADW) re-thinks the whole model for compute-intensive data warehouse workloads. While Oracle has kept its core database “DNA” intact, it has also built an entirely new experience for cloud databases: easy-to-use, fully autonomous, and instantly scalable. There's no more complexity. No more human error. No more over-provisioning to meet potential bursts of demand. And no more managing Oracle licensing cost models. This is a new Oracle with a new solution, available in its Oracle Cloud Infrastructure (OCI) whether that's deployed in the public cloud or on-premises on Exadata Cloud@Customer or Dedicated Region Cloud@Customer. Whichever way it is used, it delivers a right-sized data warehouse platform that can meet organizations' real-time operational, performance, availability, security,

¹ Source: ESG Master Survey Results, [The State of Data Analytics](#), August 2019.

and cost needs. ADW certainly grabs the headlines—but it also delivers a capability-rich platform that can offer organizations an integrated path to “data-driven enlightenment.”

On Being Autonomous

The “Autonomous” in ADW is not just one or two features. Once deployed, ADW is 100% self-driving; it leverages machine learning and internal best practices to eliminate the complexities of operating, maintaining, securing, tuning, and scaling databases. ADW’s Machine Learning is directly built-in (not a separate set of services or options, as some competitors have), which means its users can run inference directly on databases, build algorithms, train models, and leverage built-in ML data integrity. To IT organizations, this sophisticated and extensive autonomy translates to an end to the management complexities and error-potential of manual procedures that plague EDW operations, as well as a gain in consistent optimized performance, high availability, ultra-reliability, and security peace of mind.

Getting started is a breeze: a multi-terabyte database can be ready in seconds to minutes. But this simplicity masks the additional power of Oracle ADW. The underlying infrastructure is powered by Oracle’s proven Exadata Database Machine platform, offering the highest scalability and performance for running Oracle Database (and, according to Oracle, is used by 86% of Fortune’s Global 100). 99.5% availability is achieved through a redundant architecture and fully automated backups, while Oracle Real Applications Cluster (RAC) is used to seamlessly overcome compute node failures and deliver effectively 24/7 uptime year-round. The self-securing system applies patches and upgrades in real time to automatically maintain an ideal security posture. And organizations gain complete flexibility in how the platform scales to meet their needs with full elasticity, by independently scaling compute, storage, or both as real-time dynamic conditions dictate.

The importance of “autonomous” cannot be overstated. ADW is the easiest-to-use database that Oracle has ever released. It also delivers more than ‘just’ self-management; it ensures levels of performance, availability, flexibility, and security that other well-known self-managed data warehouses—such as Snowflake and AWS Redshift—simply cannot.

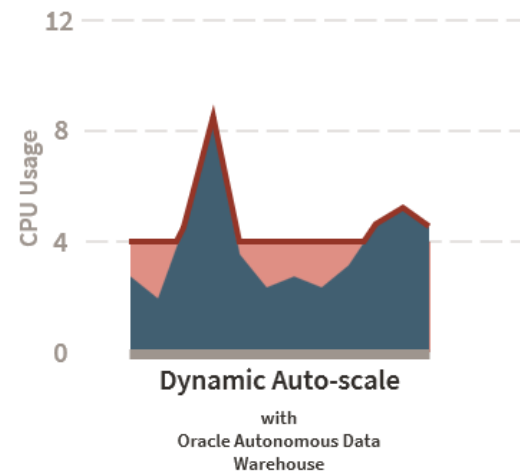
A Note on the Competition: While other solutions exist in the cloud data warehouse market today and also promise simplicity, scale, and performance, the devil truly is in the details. In most cases outside of Oracle, DBA expertise is required and tuning takes time. Before being production-ready, DBAs and operations teams must precisely define compute and storage requirements to ensure that the right amount of resources are provisioned. Furthermore, most alternative cloud data warehouses leverage generic, off-the-shelf hardware in cloud environments, so if ultra-performance and high levels of concurrency are required, users should be prepared to either be underwhelmed or pay significantly more to support it. High uptime and availability are achievable, but either will require higher tier support or lack critical capabilities in areas such as non-disruptive online patching. Moreover, when it comes to scale, fine-tuned and independent resource scalability can be a challenge, often leading to initial overprovisioning and therefore higher (and unnecessary) costs.

Scale and Cost Savings

Autonomous features are only part of the story of cost savings with ADW. Scalability—and its impact on cost—is another key differentiator for ADW. Organizations want to achieve the nirvana of true cloud elasticity: consuming precise amounts and types of resources based on current demand. When done right, that means leveraging a pay-as-you-go cloud cost model that is tightly aligned to both the platform and workloads. ADW is sized to the exact compute and storage required to support the organization’s data warehouse. If the demand of the workload increases, ADW auto-scales instantaneously to meet that demand, adjusting CPU and I/O resources accordingly while still executing queries at full performance—all the time. If demand subsides, the resources scale back down (and shut off completely when the system is idle) with the continuing ability to instantly restart if needs arise. Oracle’s cost model is aligned to the exact level of consumption at any given moment and bounded by an established maximum (3x the “baseline”) for complete cost governance. In other words, Oracle eliminates the need for organizations to over-provision and/or to scale in fixed sizes. And critical to this “all-

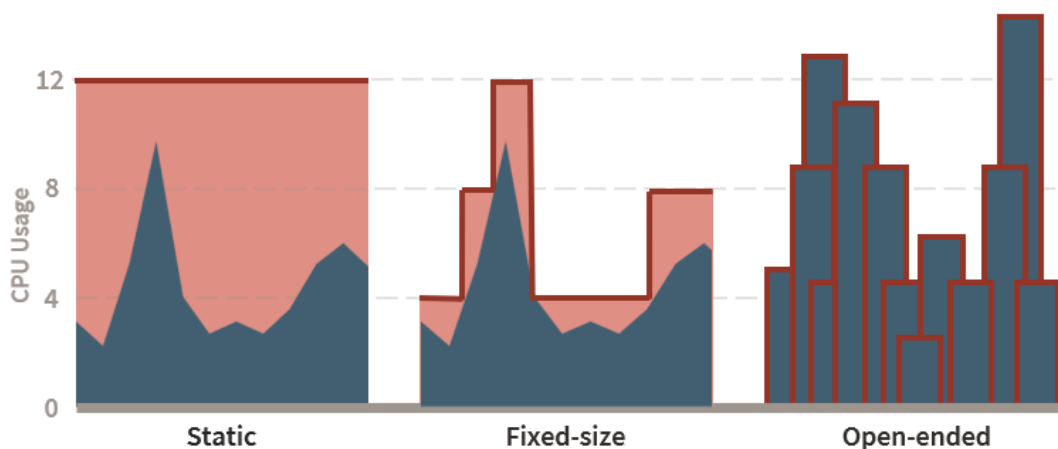
optimized, all-the-time, all automatically” approach is that all scaling—whether up or down—occurs online *while the application is continuously running*. This “Dynamic Auto-scale” ability (see graphic) is not only transparently dynamic and optimized but can result in cost savings that verge on a “how did they do that?” category. But they are real.

A Note on the Competition: While pay-as-you-go is an appealing cost model for cloud adopters, the typical competitive reliance on a more- or less rigid underlying cloud infrastructure means organizations must pick an initial “T-shirt” size and/or fixed building blocks to support incremental growth thereafter. If an organization needs 30 CPUs for its data warehouse, competitive technologies may only offer “nearby” configurations (perhaps a cluster of 4 nodes of 8 CPUs each) while ADW supports exactly the number of CPUs required.



Beyond the unit-size constraints, what further separates ADW from its competitors is scaling. ADW scales instantly to meet spikes in workloads. Most competitors who support scaling use a “bolt-on” approach of adding new nodes or new clusters to their data warehouse environment. As a result, choosing something other than Oracle ADW effectively means organizations must regularly “pick their poison” when it comes to budgeting for running a data warehouse! There are three options, but all come with potentially awkward—and sub-optimal—tradeoffs.

- “Static” – Pay for a static number of CPUs and licenses based on the highest projected peak load regardless of how much is consumed. In traditional IT parlance this is called “over-provisioning,” while in the CFO’s office, it is more likely to simply be called “over-spending.” This was always the default for traditional on-premises operations.
- “Fixed-size” – Pay for building blocks of overprovisioned resources in “steps” based on current or expected demand, which can lead to unpredictable (and likely higher) latency, and a lag-behind effect when scaling. Aside from the non-linear scaling, this is akin to deploying hope as a strategy.
- “Open-ended” – Pay for processing times based on the size of data volumes and complexity of queries across those volumes. Likely the worst of all worlds, since it relies on having an open corporate checkbook.



While it is all too easy to assert that this whole data warehouse discussion is just about fluidity and business decisions, the underlying infrastructure and application also still matter greatly. Abstracting the resource consumption of cloud data warehouses does not inherently obviate all the laws of IT technology and vendor approaches to them; “Caveat emptor” still applies. As a specific example, Oracle Database supports real-world data warehouse scenarios where existing data is often updated; Oracle’s database architecture has no corresponding expansion in the amount of storage used. Contrast this to some competitors that use a write-once approach and thus deal with updates by simply duplicating data to make the required updates to the existing values (with a corresponding increase in the storage footprint and, of course, storage costs). Others demand compromise, while Oracle ADW’s dynamic and transparent approach makes for a compelling, responsive, optimized, complete, and cost-effective platform.

Complete Platform

Oracle understands that the data warehouse serves as a critical entry point to all things data in the cloud. Oracle ADW is part of a *converged database service* that is tightly integrated across the entire pipeline to support the lifecycle of data. That means more than just data warehousing. A converged database means multi-model, multi-workload, and multi-tenant. Organizations can embrace multi-model by turning to a single data management platform for a document database doing puts and gets with JSON, a graph database for complex relational analysis, and a traditional relational database for transaction processing (just to name a few). Organizations can embrace unified and autonomous management of a single engine to support several workloads and applications infused with data. And they can embrace multi-tenancy through containerization and orchestration of databases to provide isolation, scale-out, and transactional consistency across all database types. In addition, beyond offering ADW on public cloud Oracle Cloud Infrastructure, Oracle can provide ADW in customers’ data centers via its Exadata Cloud@Customer (since 2016) and Dedicated Region Cloud@Customer offerings. There is, for example, currently no comparable “Snowflake at Customer,” and AWS Outposts with Redshift is a recent introduction.

A Note on the Competition: Of the competitors that offer ostensibly complete platforms upon which to stake a data-driven organization’s future, any excitement can quickly fall by the wayside as deciphering exercises consume more time than expected—never mind the time needed for ongoing management and maintenance—with other cloud providers. Customers must comb through what feels like hundreds of services to find the exact combination that’s right for them. Then they must be stitched together in a perfect way to satisfy end-user and application requirements. As soon as they’re “done,” it is time to start again, as IT and organizational needs change. This can quickly become an operational nightmare; just because a vendor offers what seems to be the right tools with specialized services for each data type and workload, does not—often cannot—equate to customer success when what’s needed is a holistic view to turn the data into information.

The Bigger Truth

While the technology underlying Oracle ADW is crucial, what truly matters is what it delivers to users. “Paradigm shifts” are frequently proclaimed by IT vendors but are far less often actually produced. ADW is a genuine market shift—for Oracle itself, for Oracle versus its competitors, and most of all for customers. In the realm of cloud data warehousing, it is perhaps best described by mixed-metaphor references to well-known stories: It offers its users a Goldilocks approach...*but* one where the porridge is always and instantly the right size and temperature, and one where there are no nasty surprise wolves—pricing or fixed-size T-shirts—at the door! It allows customers’ and DBAs’ focus to rightfully shift away from operations and back to extracting maximum value from their data and improving organizational outcomes, knowing that Oracle ADW is not just automatically, but autonomically, changing the game for them by handling, optimally and with real-time flexibility, the management drudgery.



All trademark names are property of their respective companies. Information contained in this publication has been obtained by sources The Enterprise Strategy Group (ESG) considers to be reliable but is not warranted by ESG. This publication may contain opinions of ESG, which are subject to change. This publication is copyrighted by The Enterprise Strategy Group, Inc. Any reproduction or redistribution of this publication, in whole or in part, whether in hard-copy format, electronically, or otherwise to persons not authorized to receive it, without the express consent of The Enterprise Strategy Group, Inc., is in violation of U.S. copyright law and will be subject to an action for civil damages and, if applicable, criminal prosecution. Should you have any questions, please contact ESG Client Relations at 508.482.0188.



Enterprise Strategy Group is an IT analyst, research, validation, and strategy firm that provides market intelligence and actionable insight to the global IT community.



www.esg-global.com



contact@esg-global.com



508.482.0188