



ESG SHOWCASE

Oracle's Modern Data Platform Strategy: Autonomous, Integrated, Secure

Date: August 2021 **Authors:** Mark Peters, Practice Director; and Mike Leone, Senior Analyst

Oracle's 'Data-centric' Strategy is both Vital *and* Valuable

The Traditional Sub-optimal Data Approach

It's not hard to list the key drivers and prerequisites of contemporary IT. Gather any group of CIOs and you can guarantee your white board will soon contain something like this: hybrid/multi-cloud, as-a-service and on-demand models, cybersecurity, dev ops, employee and customer experience, autonomous operations, app flexibility, integration, and so on. Beyond that, however, everyone will also swiftly acknowledge that data is what underpins everything. They will call it the DNA of their organizations, the new global oil, their business crown jewels, or whatever. But—and this is *a*, or most probably *the*, key challenge facing IT today—all those prerequisite elements of IT are invariably developed in silos and only layered and bolted together later, all while trying to manage exponential data growth. This means that, despite the categorical acceptance of the paramount importance of data, traditional strategies around it are all-too-often developed in isolation, and thus are not tightly integrated, instead forcing the IT equivalent of “Band-Aids” to be used to address the fact that they were not designed to work together from the ground up. Quite simply, this is putting the cart before the horse. Or focusing on the tool instead of the job. It's why we can end up blindly proliferating new single-purpose databases, tools, apps or approaches, mistakenly assuming that “new” equates to “better.”

Indeed, why is it that we assume that “cloud” equates to “modern” and “optimal” when very often it is just moving what we already do, and how we do it, to a different place? Organizations can fall into the trap of believing that simply migrating to the cloud will better enable a “data-driven future.” Moving to the cloud to some extent may well be necessary for success, but it is certainly not likely to be sufficient in and of itself. Instead of an ideal business destination, the cloud can end up as the beginning of a challenging journey.

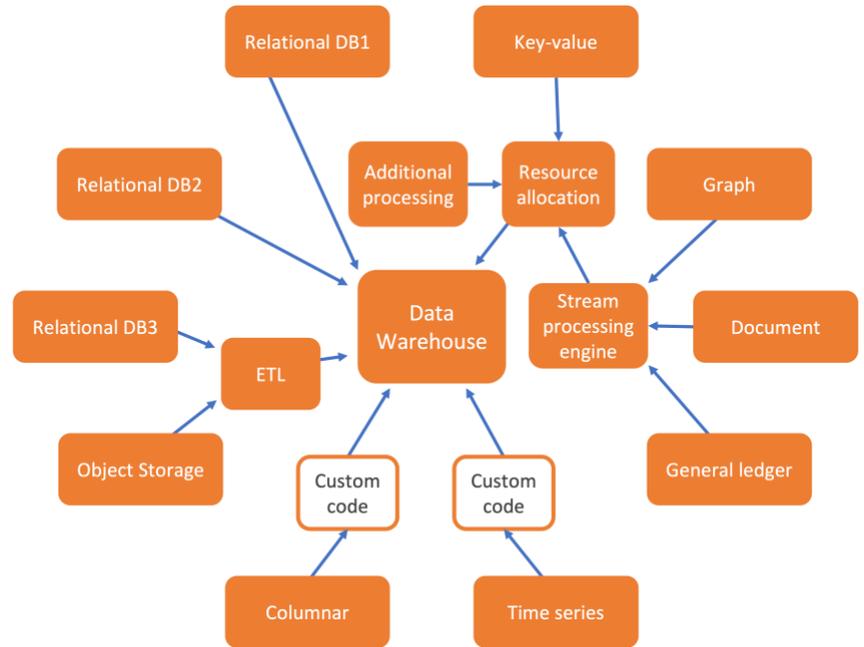
Oracle's 'Data-centric' Solution

What would clearly be better is a set of IT solutions, applications and capabilities that are based upon an integrated and flexible data foundation, a description that fits Oracle's over-arching modern data platform strategy perfectly. From consumption to security and data protection, across all types of clouds, apps, and workloads, Oracle's modern data platform strategy coalesces everything about and around the data, with a comprehensive, integrated data ecosystem of applications, development capabilities, and data management tools (including Oracle's converged Autonomous Database and MySQL Database Service with HeatWave). This is a logical and optimum approach to simplify not just the data lifecycle, but also the speed, efficacy, and breadth of a fluid and tightly aligned intelligence platform that handles multiple workloads, data types, and models. In simple terms, Oracle's data-centric, converged database approach with Autonomous Database helps provide the shortest line between IT's responsibilities and its success. Other vendors may “fix”

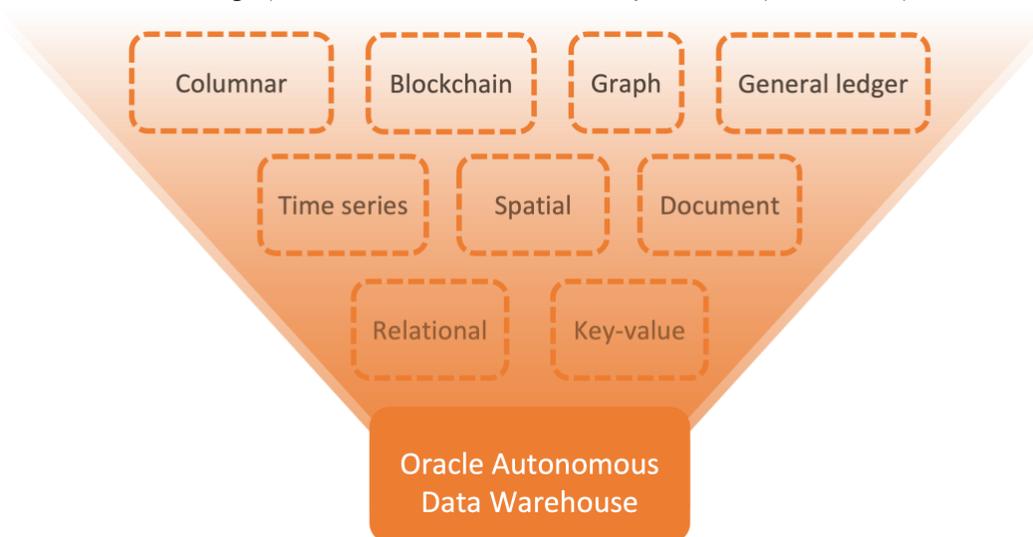
one area or another of IT, and looked at in isolation, they may do a great job for a particular use case or with individual security, storage, or app dev tools, or perhaps with business continuity, a blockchain cloud service, or whatever—but then it falls to IT organizations (whether in house or in the cloud) to assemble the piece parts, and the lack of an integrated and complete data foundation simply makes building things harder, less efficient, or nearly impossible.

Oracle's Data Platform: From Complex to Elegant

Other cloud providers offer myriad services¹ to meet myriad database and data requirements and demands. But there can be obvious challenges and operational inefficiencies in managing such disparate services. An example of the extent of the potential complexity is shown in the adjacent diagram, which is not an uncommon experience for large organizations that rely on several database types (and invariably even more database instances) to satisfy all their various business unit and data requirements. On top of these disparate data services are the services required to ensure the proper data flows with different stream data processing and ETL services. Notice that in some cases, certain database services can't even connect to the data warehouse, forcing organizations to rely on other integrations that introduce unnecessary cost, complexity, latency, and risk. Taking a multi-tool approach to try to self-build a converged database is analogous to asking a homeowner in need of light to figure out how to build and operate an electric grid instead of just flipping a light switch.



Most cloud providers would agree that the first step to effectively analyze massive amounts of disparate data is to consolidate it all into an enterprise data warehouse (EDW) in the cloud. Oracle's integrated data platform approach (summarized in the graphic below) does that inherently and thus prioritizes operational efficiency and simplicity. Why



require already overburdened staff to manage disparate data services when a *converged database* can deliver the management simplicity and efficiency desired by everyone throughout the IT and data teams? Oracle's converged database approach is both multi-model (with native support for different types of data, such as relational,

¹ AWS, for instance, lists 11 on its [website](#), in addition to a range of ETL services.

OLAP, JSON, spatial, blockchain, and more) *and* multi-workload (with optimizations for different use cases such as analytics, document management, machine learning, graph analysis, etc.). With its MySQL Database Service with HeatWave, Oracle effectively combines all the analytics use cases into a single cloud data warehouse with query acceleration that can deliver more performance at lower cost than the competition.² And, with its Autonomous Database, Oracle effectively combines dozens of database services into a single, self-operating entity. Having all this data reside in a common database provides many benefits, including:

- **Reduced Costs** – Just a single database needs to be installed and maintained, with Oracle providing machine learning and property graph capabilities integrated into the database.
- **Enhanced Performance** – Data is made immediately available, and there are no time delays due to data transfer and transformation.
- **Enhanced Reliability** – Data is stored in the self-repairing Autonomous Data Warehouse, providing automated protection from downtime with 99.995% SLAs. Furthermore, eliminating manual management, data transfers, and data manipulation ensures that no errors impact the data flow.
- **Improved Security** – A single database means limited copy contagion, fewer vulnerability and exposure points, easier compliance with data privacy regulations, and a unified security model.
- **Simplified Management** – Administering a single database for all workloads reduces management complexity and the number of different skill sets required.
- **Faster Deployment** – Developers don't have to write code for integrating data from multiple sources or propagating it across numerous systems.

The logic of Oracle's comprehensive data platform strategy is compelling. We'll now turn to examine three areas where Oracle's differentiated strategy can help optimize IT more broadly.

Evolution of Data Landscapes

The Challenge

Today, data systems are more passive than active. Customers store data in certain locations and when they want to look at that data to gain insights, they ask a question. If they don't ask the question the right way, they won't get the answer they're looking for, or, worse, they'll get the wrong answer. With the amount of data, and the pace at which it is generated, this can be akin to finding a needle in a haystack (perhaps more apropos, it's like looking for a "sharp pointed, potentially metallic, object" in a whole county of piled grasses and wheat). Another dimension to this challenge is the dynamic nature and variety of data. Today, organizations are forced to look at their data and ask themselves to identify the data type. Where is it? What system or application is most applicable to store and analyze this data? This has led to an inordinate proliferation of databases, tools, and services to satisfy myriad use cases across the business, causing operational burdens (such as ETL, which involves moving and re-formatting data from one database or tool to another) that create bottlenecks across the

² While ESG has not independently verified Oracle's comparative claims regarding analytics use cases (such as being up to 1400 times faster than AWS Aurora and 6.5x faster than an optimized AWS Redshift with AQUA data warehouse deployment, with substantial cost savings in each case), the scale of difference deserves consideration, and the benchmark testing details and scripts are publicly available on GitHub for anyone to review or replicate.

whole data lifecycle. Organizations need help in reining in the complexity of a dynamic data ecosystem to deliver more value to the business faster.

Oracle's Stance—Lake House

Oracle's approach is to simplify the data management lifecycle, from collection and integration to insight with a "lake house" that combines a data lake and data warehouse together with a supporting data catalog, data integration services, and managed open source services, including Apache Spark and Apache Hadoop. Both the data lake and converged database support multiple data types and different workloads, so organizations do not have to design a different system for each data type and workload. Instead, data is stored where it makes sense based on use case, data value, and the tools currently being used to process it—and can be easily moved if and as needs evolve. Data scientists and business analysts can access any data they need to build machine learning models or create the right reports, using the commercial or open source tools of their choice. Oracle's intent is to free organizations from the data access burden by putting data into the right place and letting the system evolve over time to meet changing business needs (scale, performance, etc.), regardless of data type, model, or workload.

From Automatic to Autonomous

The Challenge

Whether it's a traditional data center or a cloud data center (either on- or off-premises), the simple fact is that IT is getting too big and complex, and business requirements are moving too fast, for any human to truly get their management and operational arms around it. For sure, the move to all sorts of cloud operations has been swift and significant, but any assumption that cloud operations inherently confer a degree of optimized control that is not possible on-premises is simply erroneous. The main reason moving to the cloud isn't a panacea for IT's woes is that the scale, growth, and variety of data doesn't change. Databases, storage, apps, containers, and virtual machines are all simultaneously data vessels themselves as well as tools to enable improved data-centricity, but they must work together.

There are more applications, more people, and more data distributed across dynamic environments (cloud and traditional), all of which require expertise to manage, maintain, optimize, and secure. The pressure on operations teams to ensure the accessibility, performance, security, and reliability of data-centric applications has never been greater. Moreover, most people are still operating in the cloud almost identically to the way they do/did on-premises. They put a box around their application requirements, select a "t-shirt size" based on existing (or, more often, estimated peak) resource requirements, and essentially only change that shirt size when their workload hits a bottleneck. The same goes for database tuning. It's done when the database is first set up, then left alone until performance hits a wall. While cloud providers are evolving to deliver right-sized resources via platform-as-a-service consumption models and managed service offerings, the need for a fundamental re-thinking of a company's data platform is becoming more apparent. Automation is a first step to help eliminate mundane tasks but, by itself, is not sufficient to meet the lofty expectations organizations have.

Oracle's Stance—Autonomy

Oracle recognizes that an assumption that organizations have enough staff to handle all infrastructure-related planning, tasks, challenges, and problems is incorrect. There simply are not enough experts to manage thousands of interconnected systems and services today *and* exponential data growth. The complexity of managing and maintaining diverse and dynamic environments is beyond any individual's ability to comprehend. To achieve the desired accelerated improvement

in management and operations that is increasingly the rule not the exception, Oracle is driving a move from automation to autonomy as the foundational principle of its flagship Autonomous Database.

Oracle's approach to its modern data platform sees a different role for the cloud than its competition. Oracle does not believe in putting operational burdens on the limited number of technical experts who are trying to manually manage, scale, tune, patch, or upgrade the underlying infrastructure. Oracle does not believe in putting boundaries around applications and workloads. *Oracle believes in letting the system learn about the workloads, evolve as needed, and resolve its resource requirements for itself.* Virtually every industry is benefitting from embedded AI and ML, so why should data infrastructure management be different? Oracle Cloud Infrastructure relies on built-in intelligence to adapt autonomously. It recommends, scales up or down, self-corrects, self-secures, and self-tunes, freeing technical experts from the ongoing operational burdens they would otherwise suffer.

Ending Security Risk Tradeoffs

The Challenge

Cybersecurity may be on every IT person's mind, but that is often as much about the awkward tradeoffs it demands as it is about any assurance of data security. Organizations are forced to continually balance data safety and business risk; breaches will always happen to someone else or way in the future. In the meantime, organizations may opt to not implement a deep system patch because it will require a reboot and downtime, which are unacceptable to the business (and guarantee an immediate negative impact). Data privacy, sensitivity, and protection matter, whether for commercial upside, common decency, or regulatory and compliance adherence. It is not realistic to hand this off carte-blanche to a public cloud where users are traditionally on their own (as they would be on-premises) using point tools and technicians to ensure data is protected, masked, etc.

Oracle's Stance—Always-on Security

Oracle's approach to data security is simple and attractive: make the functionality uniform, end-to-end, built-in, and autonomous. To users and IT organizations alike, that makes security invisible and always on, and, as a result, risk tradeoffs essentially disappear. Keep data protected, secured, and accessible from the moment it arrives; automatically track who's accessing what and making changes; and automate hot patching to preclude tradeoffs between updated security, data protection, and application availability. Quite simply, it's security without the burden of continual technical work, as the system is autonomously responsible.

The Bigger Truth

A compelling and valuable overall data management strategy is ultimately not about how many petabytes you can store or how fast you can move them. It is about optimally (efficiently, autonomously, and securely) having all your data integrated, available, accessible, and useable, regardless of where in a hybrid cloud ecosystem it sits, and how or by whom/what it is used. Oracle's original DNA was all about databases for mission-critical applications rather than underlying infrastructure. That's 180 degrees different from much of its competition, which explains its data-centric focus. While other cloud database providers (the likes of Amazon and Snowflake, but many others too) continue on a specialized cloud database path, Oracle instead offers an integrated cloud data platform that extends and complements its core database and has extended the database itself well beyond just being converged. It is autonomous, elastic, secure, and available on-premises or in the cloud using Oracle Exadata (its high-performing Oracle Database platform that incorporates unique features not available elsewhere) or its MySQL Database Service with HeatWave for rapidly growing MySQL OLAP and OLTP workloads. Indeed, while the term most traditionally associated with Oracle to-date has been "database," looking at its

modern data platform, the future history view is just as likely to emphasize “data-centric” as Oracle’s signature industry contribution.

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