



# Oracle Database@Azure Profoundly Changes the Multicloud Game

February 2024

**Written by:** Marc Staimer, Sr. Analyst, Wikibon & President, Dragon Slayer Consulting

## Introduction

This research explores why organizations need or want to use multicloud architectures and the obstacles that have inhibited or blocked that adoption. It then evaluates how Oracle Database@Azure empowers organizations to nullify these obstacles while gaining game changing advantages. It will uncover how Oracle Exadata Cloud Infrastructure powered by AMD EPYC™ processors allows organizations to run Oracle Database workloads in Azure with the same high speed, availability, and cost-efficiency as in OCI.

## Premises

IT organizations have been demanding multicloud flexibility for nearly a decade. They recognize every cloud service provider (CSP) has both strengths and weaknesses and there's enormous value to running application workloads on the most appropriate cloud service. This potentially allows them to effectively leverage individual CSP unique capabilities and performance while ultimately lowering overall costs. This is analogous to what large and midsize organizations have been doing for their on-premises IT infrastructure.

Whereas the multicloud strategy makes enormous sense conceptually, practically it has been an entirely different matter. There are problematic, costly obstacles that can be difficult to overcome.

[S&P Global Market Intelligence](#) revealed that approximately 98% of businesses use two or more clouds. It also exposed the exceedingly tough obstacles small, medium, and even large enterprises continue to face that effectively slow or even prevent them from using multiple clouds. This is especially true for business-critical applications. Considering that, only approximately 31% of enterprise workloads have made it into public clouds today per that S&P Global Market Intelligence report. The question is why do these obstacles to multicloud adoption still exist?

There are several possible reasons. Some are based on incentives, a.k.a. the "carrot". Some are based on disincentives, a.k.a. the "stick." Some are just speed-of-light limitations. The next section will discuss these obstacles in more detail.

## Multicloud Obstacles

- 1) Cloud Service Provider (CSP) lock-ins
- 2) Too many ill-suited single-purpose database services for business-critical workloads
- 3) Multicloud integration deficiencies
- 4) Speed-of-light latency over distance
- 5) Dissimilar processes and ecosystems across CSPs
- 6) Complicated/costly multicloud HA, DR, and business continuity
- 7) Disruptive and time-consuming database migrations from on-premises to the cloud

### *Cloud Service Provider (CSP) Lock-ins.*

Amazon Web Services (AWS), Microsoft Azure (Azure), and Google Cloud Platform (GCP), have worked diligently to make their public clouds as sticky as possible, commonly with unique tools and services. It starts with the “carrots” that positively incentivize customers to use their cloud services. There’s nothing wrong with carrots to attract and keep customers. CSPs are constantly adding unique, proprietary management and platform tools – observability is one example – to make their clouds “stickier.”



However, they also use “stick” disincentives that make it difficult and quite expensive to move data or application workloads out of their clouds. Perhaps the biggest stick is very onerous egress fees. When data is moved, copied, migrated, or even just read from one CSP cloud to another cloud or on the customer’s premises, there is a large bill that accompanies it. This is a very strong lock-in especially when the customer cannot afford or justify those egress fees.

### *Too Many Single-Purpose Database Services that are Ill-Suited for Business-Critical Workloads.*

Most CSPs offer separate database services for each type of database model. These specialized databases create additional overhead. For example, analyzing data from a transactional database requires that data be copied, stored, transformed, and loaded into a separate analytics database service. That’s a lot of extra work and cost for multiple database services, duplicate storage, ETL service, and time. That time sink means the transactional data can’t be analyzed in real-time. That’s just two database models. Rinse and repeat for AI machine learning (ML), blockchain, time series, JSON, XML, spatial or graphic, data lake access, and vector store/search. The customer ends up with a lot of unnecessary costs for non-real-time analysis.

These specialized databases often don’t run on high-availability platforms, include built-in replication, deliver high-performance, or use consistent security models. Furthermore, the combination of multiple ETL processes and the fact that some of these databases only offer a partial implementation of the ACID (atomicity, consistency, isolation, and durability) properties poses a challenge for business-critical workloads.

The better alternative is a converged database service that unites multiple database models into a single database with a single copy of the data. And just as notably, ACID, high-availability, performance, and scaling so fundamental to business-critical workloads regardless of the database models can also be provided. Using one converged database for different models means limited to no ETLs. Converged databases allows customers to employ the same security and data protection policies across all their database models. When that converged database offers business-critical availability, performance, and scalability it immensely simplifies IT environments needed to support crucial applications while cutting total database costs by approximately 66% – based on previous Wikibon research by Senior Analyst David Floyer.

There simply are not very many business-critical converged database cloud services. There’s the Oracle Database cloud services – comprehensive converged database models that include OLTP, OLAP, JSON, XML, time series, block-chain, spatial/graphic, data lakes, vector, AI-ML, and more. Another, albeit not quite as converged, is the Oracle MySQL HeatWave cloud service available on both OCI and AWS. That’s it. AWS, Google, Snowflake, and Databricks are attempting to develop some convergence to their database services.

AWS tries to do this by integrating Aurora (transactional database) with Redshift (data warehouse), but it still requires two separate database service fees and duplicate storage. They just included their ‘Glue’ ETL service

with it. However, no matter how it's marketed, fundamentally Aurora and Redshift are two separate databases with two separate costs. Snowflake has been developing its own transactional database called Unistore and purchased Neeva, an AI-ML software supplier. Unistore is quite nascent and has quite a way to go in its development cycle to be business-critical. And the acquired AI-ML capabilities are not yet integrated into Snowflake data warehouse. Both are separate services. Databricks also acquired an AI-ML company that is not integrated into the Databricks data lake.

What this means is that most CSPs currently lack a complete business-critical converged database service. That's a real problem considering that most enterprises worldwide depend on the business-critical converged Oracle Database. Especially, the top companies including 10 of the top 10 banking institutions, 10 of the top 10 communication firms, 10 of the top 10 drug and food companies, 9 of the top 10 automotive firms, 9 of the top 10 healthcare organizations, and 8 of the top 10 retail operations all use Oracle Database running on Exadata infrastructure (more on that later). To move their business-critical applications to the AWS, GCP, or traditional Azure clouds means a major restructuring of their operations. These organizations would then need to refactor their applications, restructure their databases, contract for several database model services, and a lot more storage than they're currently consuming. All of this while finding new ways to provide the business-critical high-availability that Oracle's known for. These are major – non-trivial – obstacles that typically require a multi-year cloud migration, which still might not succeed.

#### *Multicloud Integration Deficiencies*

Customers can run different applications and databases in different clouds and can even run the same application and databases in different clouds. But when the customer needs or wants to run their application in one cloud and their database in a different cloud for business, functionality, or performance reasons, there are severe problems. Integrating two environments into a single application stack across clouds create monumental performance, networking, identity management, and security challenges.

Cloud customers demand a high level of seamless integration between their business-critical applications, trusted application support services, and trusted data management services. This is what they experience on-premises. They want to develop and run their applications in the same easy integrated way they do on-premises and not be restricted to the limited services of a single cloud. That translates into consistent tooling, management, and software in their multicloud environment, but unfortunately what works in one cloud does not necessarily work in the others.

Too frequently customers are required to replace either their trusted applications or trusted data management services when moving to a public cloud because those services seldom exist in the same cloud. In Azure for example, until recently the only native database option for business-critical applications was Azure SQL. As previously stated, the vast majority of enterprise organizations rely on Oracle Database and, while customers could run it on top of IaaS platforms in Azure, they couldn't get some of the business-critical features they use and depend on in their own data centers. And while these customers were also running applications on Azure, moving the Oracle databases supporting business-critical ones to Azure SQL is a massive and expensive effort that few have been willing to make.

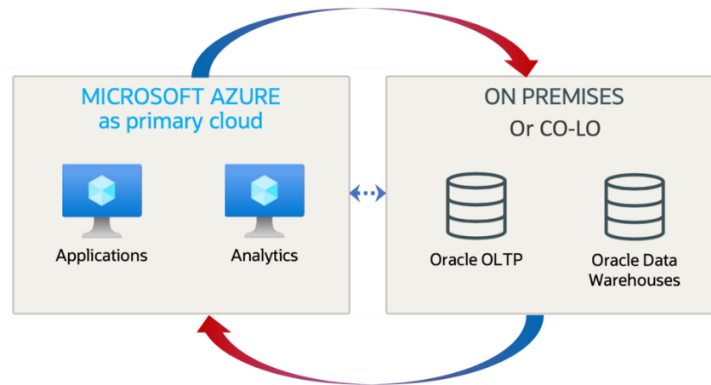


Figure 1: How Azure Cloud Customers Historically Integrated Azure Apps with Oracle Databases

### Speed-of-Light Latency

Light is fast, but it doesn't have infinite speed. Many types of applications are sensitive to the latency between when a request is made to read or write data and when it completes. Transactional applications are particularly sensitive to the accumulated latencies so it's typical to find application servers, database systems, and storage sitting right next to each other. As the network distance between an application and its database or data storage increases so do latency and application response times.



The latency an application experiences is primarily driven by round-trip circuit distance over a metropolitan area network (MAN) or wide area network (WAN). Circuit distances beyond 200 kilometers (km) are generally not acceptable for business-critical applications since round-trip latencies come in at about 5.2 milliseconds<sup>1</sup>, and many applications require latencies that are lower than that.

5.2 milliseconds may not seem like a lot until it's compared to on-premises latencies of a couple of hundred microseconds ( $\mu$ s) or less between business-critical applications and their data. Studies have shown that application response time correlates directly with positive user experiences, enhanced user productivity, time-to-actionable-insights, time-to-action, time-to-market, and time-to-revenues. High response times for business-critical applications can be incredibly costly, whereas low ones can be a major competitive advantage.

There are three ways to provide low latency with clouds, all of which require data and application processing to be located close together. Customers can either run the app and database in the same cloud data center, bring the cloud application and processing capabilities into their data center using platforms supplied by cloud providers, or have the applications and databases in different data centers that are located extremely close to each other. With multicloud being defined as using services and resources from multiple CSPs, all three approaches can technically provide the desired result depending on the latency requirements.

<sup>1</sup> Speed-of-light travels at pretty close to 200 km/millisecond (ms) in optical fiber. WAN equipment adds approximately a ms, plus some fiber adjustments adds an additional 10% or so. The total roundtrip latency equals roughly 5.2 ms @ 200 km. That's just the network latency.

### Multiple Cloud Vendors in the Same Data Center

Having multiple cloud vendors in the same data center and being able to directly access the other's resources provides the lowest possible network latency. This is the model that Oracle and Microsoft are following with Oracle Database@Azure.

### Bringing Cloud Platforms into Customer Data Centers

Having systems from multiple CSPs in a customer data center can also provide very low latencies since networking distances are quite short. Data doesn't flow out of a cloud data center – i.e., no egress fees – and the on-premises cloud resources easily connects with non-cloud application workloads, data sources, and systems located in the customer data center. However, you may not be able to run all the cloud services available in the CSP's public cloud data center, and IT teams must manage the networking, access control, and user IDs across the environments.

Oracle has bucked that trend by providing their Cloud@Customer services. More will be discussed on these services further on in this research document.

### Connecting Multiple Data Centers

Distributing an application stack across two cloud data centers will work as long as the application workloads can tolerate the additional latency and assuming the network linking the locations is short enough that its latency is under 2 milliseconds. However, customers will still generally have to manage the networking, access management, and ID federation. They'll moreover ring up costly egress fees when data is moved or read out of one or more CSP data centers. Oracle and Microsoft also support this approach in a dozen locations around the world with the Oracle Interconnect for Azure.

This approach is also applicable between a cloud data center and a corporate data center or colocation site with the same distance limitations and management caveats. For the purposes of this research, this approach falls under hybrid cloud, not multicloud.

### *Dissimilar Processes and Ecosystems across CSPs*

Every CSP has different processes for accessing, orchestrating, configuring, monitoring, managing, changing, and paying for their services. The user interfaces are different. Monitoring, troubleshooting, and escalation procedures are different. Support and SLAs are different. Identity and access management are different. Discounts and purchase commitments are vastly different.

Running in multiple clouds entails learning, knowing, and using similar but different processes in all clouds. This entails more expertise, training, personnel, human errors, and, of course, higher costs.

### *Complicated and Costly Multicloud HA, DR, and Business Continuity*

The use of several non-converged specialized databases instead of a business-critical converged database makes implementing high availability (HA), disaster recovery (DR), and business continuity extremely complex and costly in multicloud environments. Every database service, their duplicate storage, ETL services, and data pathing must be carefully mapped and replicated across multiple regions. Recoveries and restores are just as complicated, time consuming, and costly in time and treasure. Using multiple clouds can help, but users can expect disruption and downtime before everything is back up and fully operational in the event of an outage.

The HA, DR, business continuity complexity, and cost increases exponentially when application workloads, databases, and storage are multicloud. The proliferation of failure points are especially problematic when troubleshooting. Application workloads that worked acceptably before a failure suddenly do not after a recovery. Determining the root cause of that failure takes patience and time. Time that the application is down.

*Disruptive and Time-Consuming Database Migrations from On-Premises to the Cloud*

Moving from a business-critical, highly converged database to multiple specialized, single-purpose cloud databases is a highly convoluted process. The more database models that are being used, the more convoluted the process becomes. Every application workload and application to database connection has to be modified, altered, or re-written to fit the parameters and sidestep the limitations of the specialized databases. This a non-trivial, pain-staking, time-consuming exercise. Few developers want to take on this huge task, which is why it's a major obstacle to cloud adoption.

What makes it a multicloud obstacle is when the application workloads are in a different cloud than the databases it uses. Adding to the daunting databased migrations tasks come the additional time consuming tasks of application connection, testing, network management, performance tuning, troubleshooting, DR, etc. These formidable additional task makes multicloud that much more daunting.

**How Oracle Database@Azure Removes Multicloud Obstacles**

Oracle executive management has made it crystal clear that Oracle Cloud Infrastructure (OCI) is laser focused on meeting customers where they are – both in terms of location and their cloud journey. That means not forcing them to lift and shift everything or refactor application stacks to utilize a completely different architecture in the cloud. That sounds like an excellent promise. It is ultimately more important to see how Oracle’s actions align with their promise.

For multicloud, Oracle has implemented multiple paths. None of them are mutually exclusive. The first path is with the Oracle Modern Data Platform, which opens up integrations with both Oracle and non-Oracle applications, databases, analytics, and data sources. As shown below, Oracle’s done that with a boatload of vendors, including fierce competitors. More integrations are on the way.

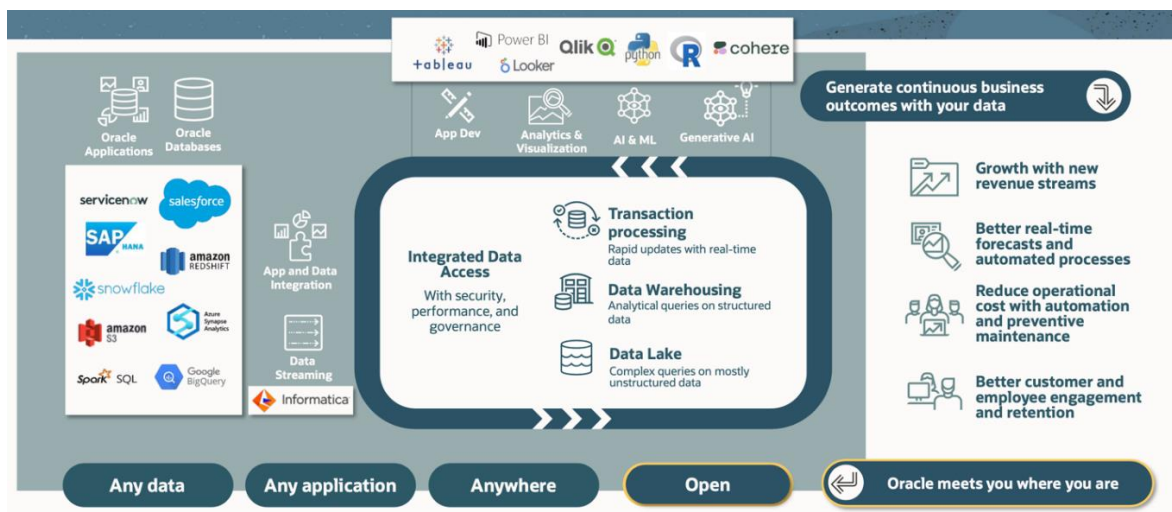


Figure 2: Oracle Modern Data Platform

The second path is to move the cloud to the customer’s premises. Oracle has done this with Exadata Cloud@Customer, Compute Cloud@Customer, and OCI Dedicated Region – a complete OCI region on the customer’s premises that offers the exact same 100+ services available in Oracle’s public cloud. As discussed above, this can offer the extremely low latency and high performance that business-critical applications need in Distributed Cloud environments.



Figure 3: Oracle Platforms in Customer Data Centers

The third path is to run some key OCI database services in other CSP clouds. A great example of that is making the extremely popular and fast-growing Oracle MySQL HeatWave available natively in the AWS cloud. MySQL HeatWave is a converged OLTP database with the integrated HeatWave analytics engine, AutoML, data lake processing, and soon vector store and search available—all with highly-automated processes managed by the ML-driven MySQL Autopilot.

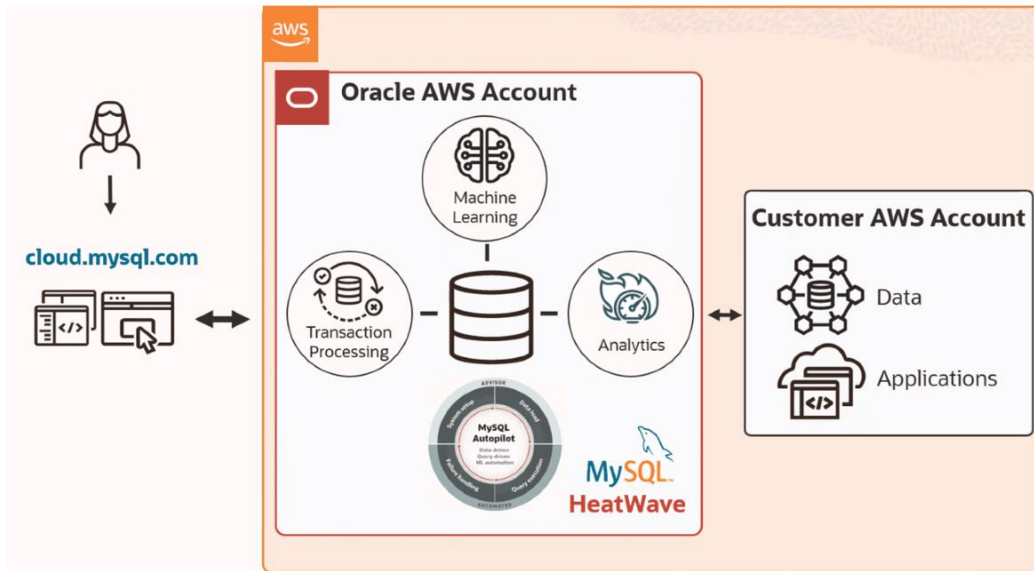


Figure 4: Oracle MySQL HeatWave in AWS

The fourth path to meet customers where they are is to partner with other CSPs. No partnership has been more crucial to this than the one between OCI and Microsoft Azure. The OCI and Azure partnership has resulted in several programs to eliminate multicloud obstacles for customers who want their trusted Azure applications to work with Oracle’s industry leading databases in the cloud.



The first partnering effort between Oracle and Microsoft was on the high-speed Oracle Interconnect for Azure, which provides high-performance interconnects with 2 ms roundtrip latency between pairs of OCI and Azure data centers. The interconnect allows running both database workloads – including Exadata Database Service, Oracle Autonomous Database, and MySQL HeatWave – as well as non-database related workloads with private endpoint. This initial multicloud partnership removed several of the obstacles to multicloud including:

1. Simplifying interconnect setup and identity federation.
2. Eliminating data egress fees between OCI and Azure over Interconnect.
3. Enabling Azure applications to use the fully managed, enterprise trusted Oracle business-critical converged database services running in OCI.
4. Providing collaborative support services.

These noteworthy improvements in multicloud transparency were a major step forward. However, Oracle and Microsoft quickly learned that customers wanted to combine the capabilities from both clouds with much lower latency and higher levels of integration. In effect, they wanted local access to Oracle’s Exadata Database Service and Autonomous Database in Azure data centers, the low latency of such an environment would provide, and the ability to pay a single company for all of their consumption. To their credit, Oracle and Microsoft listened and the result is Oracle Database@Azure.

Running Oracle Database services on Exadata infrastructure in Microsoft Azure regional data centers – branded as Oracle Database@Azure – is the partnership’s most recent and quite significant advancement in multicloud usability. It’s a major step up in the partnership, a true game-changer. It places Oracle Exadata platforms running the latest Oracle Database and Autonomous Database instances physically inside Azure data centers and directly accessible from Azure services. The same Exadata powered Oracle Database services found in OCI regions and on Exadata Cloud@Customer are now native Azure services to Azure users and provide them with the same converged database capabilities and high availability features enterprises take for granted. And since they’re based on the same OCI region AMD EPYC processor powered Exadata platforms, customers see the same extremely low-latency and unparalleled high-performance.



### Why Oracle Database@Azure is So Game-Changing

The game-changing benefits of Oracle Database@Azure start with much reduced latency. By placing OCI's high-speed architecture and the Exadata database platform within the Azure cloud, latency is reduced by an order of magnitude from 2 ms to  $\leq 200\mu\text{s}$ . The latency between OCI Oracle Database services and Azure resources such as Compute is the same as any two native Azure resources.

That reduced latency correlates with much faster application response times. Those faster application response times translates into considerable productivity gains, better work quality, higher morale, lower turnover, reduced costs, faster time-to-actionable-insights, time-to-market, leading to faster time-to-unique-revenues.

#### Exadata's Unmatched Database Performance



But why did Oracle put Exadata inside Azure data centers? Couldn't they just run Oracle Database on standard Azure cloud infrastructure and achieve the same result. In a word, no.

The performance, scale, and availability that Oracle Database achieves when running on Exadata are the second, third, and fourth game changing benefits delivered by Oracle Database@Azure.

Exadata is designed in conjunction with Oracle Database. It combines ideal database hardware with database aware system software and automated management.

Exadata's unique features reduce SQL read latency, a critical factor for transactional applications, to less than  $19\ \mu\text{s}$  – 50x faster than the 1ms read latency for block IOs in Azure. It uses Exadata Smart Scan to offload data intensive analytics and AI/ML queries to intelligent storage servers, eliminating the need to copy all data to database servers for processing and providing up to 2,880 GB/s of scan throughput. And, Exadata supports Oracle Real Application Clusters (RAC), which not only enable database environments to scale from 2 to 1000's of CPU cores, but also enable the database platform to provide high availability so Oracle Database environments continue to run when individual components or servers fail – a necessity for business-critical workloads. In addition, Azure customers were increasingly requesting Oracle Autonomous Database, which runs exclusively on Exadata, so the choice to deploy Exadata platforms inside Azure data centers was obvious.

Exadata is by far the fastest and best place to run Oracle Database – including Oracle Autonomous Database – so it's no surprise that Oracle and Microsoft would want to have it inside Azure data centers. Exadata's ideal database hardware is the result of more than 15 years of co-engineering with Oracle Database, Oracle Linux, and partners like AMD – whose EPYC processors power Exadata Cloud Infrastructure database servers. AMD EPYC processors provide exceptional time to results for business-critical applications while delivering optimized performance for database workloads. The platform includes Exadata internal RDMA networking between database servers and storage servers, high-speed network-based synchronization, massively parallel processing of analytics and AI Vector search primitives in intelligent storage servers, and automated prioritization of latency-sensitive IO operations.

Exadata Cloud Infrastructure in Azure data centers uses AMD EPYC processors to provide 252 to 4,032 cores of processing power in database servers to support thousands of concurrent transactions, complex business analytics, graph and spatial analyses, in-memory processing, and more. Based on previous Wikibon research, no other combinations of processing, networking, and storage comes anywhere near Exadata's low latencies, high

IOPS, and high throughput for Oracle Database services. That performance does not come at the sacrifice of capacity.

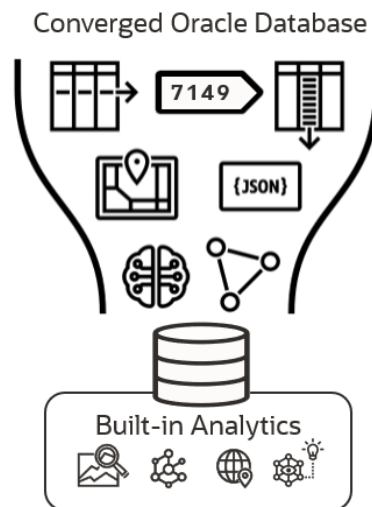
Each Exadata system in Oracle Database@Azure supports from 190 TB to 4.2 PB of usable storage capacity, which can be for any mix of the Oracle Database models and data types. The intelligent storage servers use their CPU cores to encrypt and decrypt data (all data in OCI must be encrypted) and locally run SQL queries to support analytics, ML model building, and AI vector search. When combined with Oracle Database Hybrid Columnar Compression, each Exadata system can support up to 30 PB of data warehouses.

Clearly, using Exadata in Azure data centers gets rid of the multicloud speed-of-light latency obstacles – not just for relational databases but for ones that need JSON, spatial, graph, vector, or any other converged database capabilities found in Oracle Database.

Oracle Database – the Business-Critical Converged Database for Enterprises

The Oracle Database has more than 5 dozen capabilities that are only available on the Exadata platform, that specifically accelerates database performance, increase database scaling, improve availability, and make the platform easier to use and manage. These include support for Autonomous Database – the most advanced and self-driving database available.

Oracle Database is the most complete business-critical converged database available today. It provides OLTP, OLAP (data warehouse), JSON (documents), XML (object), machine learning (ML), block-chain, spatial, graphics, time series, and data lakehouse from a single database service. This high level of versatility helps organizations reduce the complexity of their IT environment, lower costs by using resources more efficiently, and develop robust applications that can easily be extended with new data types and functionality.



A current topic of investigation for many customers is how to leverage large language models (LLMs), and specifically how to do so with the greatest accuracy and without exposing their proprietary information to competitors. Upcoming Oracle Database 23c AI Vector Search capabilities will enable enterprise data—structured and unstructured—to augment generative AI large language model (LLM) interactions in a process called retrieval augmented generation (RAG).

RAG is used to improve generative AI LLMs response accuracy and relevance to business questions by letting customers use non-public information to guide responses and reduce LLM “hallucinations.” While this research won’t go into specifically how RAG works, using it successfully requires databases to search thousands or millions of high-dimensional “vectors,” each of which include hundreds or thousands of values. Being able to rapidly sort the vectors that most closely match the initial prompt and retrieve the original information they point to is crucial to the ability of RAG to improve the accuracy and relevance of LLM responses in business environments.

Oracle Database AI Vector Search implements this in a sophisticated, flexible, and highly performant way that takes advantage of Exadata’s intelligent storage servers and the high core counts provided by AMD EPYC processors to let Oracle Database@Azure users leverage their corporate data stored in Oracle databases in Azure data centers.

As a converged database, the Oracle Database lets a single data copy be used for all database models thus eliminating most ETLs and mitigating the transformations required. Bear in mind that little to no data originates in many of these database models. In contrast, using a single-purpose vector database means that data must be copied from somewhere else, vector embeddings created, and then placed in the vector database. That translates into a non-trivial time sink and costs that escalate as data going into the vector database scales. And must be repeated ongoing as new data is added. Since most enterprise data that can be used to guide RAG already resides in Oracle Databases, storing vector embeddings in Oracle Database, alongside the source data for those vectors, means no copies, no extra storage, no data movement, no time delays, and no extra costs.

Oracle Autonomous Database Services



Oracle Autonomous Database is still the only self-provisioning, self-configuring, self-tuning, self-scaling, self-healing, self-securing (and more) database in the market today. Making these automated capabilities available in Azure is the fifth game changing benefit of Oracle Database@Azure

Oracle Autonomous Databases is built on Oracle Database. It’s the ultimate managed database service and incorporates more than four decades of Oracle best practices and AI techniques to allow it to quickly respond to changing workload needs.

Oracle Autonomous Database’s extremely high automation levels and ability to scale consumption up and down in real time, without interrupting operations, has made it the acknowledged enterprise leader in high-performance, simple, advanced cloud database services. By taking advantage of the power of Oracle Exadata in Azure data centers, Azure customers will now be able to take advantage of Autonomous Database’s benefits with business-critical application workloads running in Azure data centers.

Tighter Integration Between Azure Trusted Applications and Oracle Database@Azure

With Oracle Database@Azure, the OCI database resources and services running on Oracle Exadata are now natively accessible to applications in Azure data centers. The simplicity of using and managing multicloud resources is the sixth game changing benefit delivered by Oracle Database@Azure.

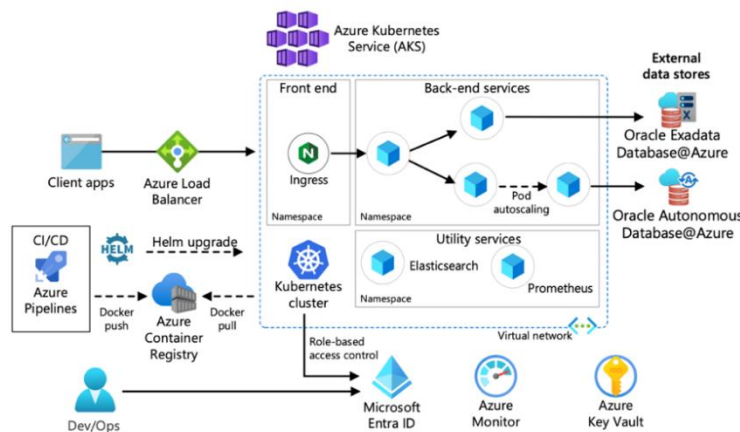


Figure 5: Example of an App Running on Azure Kubernetes Service Using Oracle Database@Azure

There’s no need to cross-connect OCI and Azure clouds. Applications access Oracle Database@Azure services using an Azure Virtual Network, Azure native identity service or a federated identity provider, and standard Azure APIs. Customers monitor database metrics, audit logs, events, logging data, and telemetry natively in Azure.

Microsoft and Oracle are delivering a unified collaborative support. Best of all, customers get a single bill, MACC applicable, and with no extraneous charges. Oracle Database@Azure looks and feels as if it is a native service of Azure while integrating seamlessly with Azure applications, support, and business practices.

The end result of providing Azure native access, monitoring, and management of Oracle Database@Azure services is that the obstacles created by having dissimilar processes and ecosystems across CSPs are eliminated.

Oracle Database@Azure HA, DR, and Business Continuity

Oracle Database@Azure comes with built-in with HA, DR, and business continuity that are enabled by running Oracle Database on Exadata. The level of these capabilities is far higher than in other clouds and represents the seventh game changing benefit customers can achieve with Oracle Database@Azure.

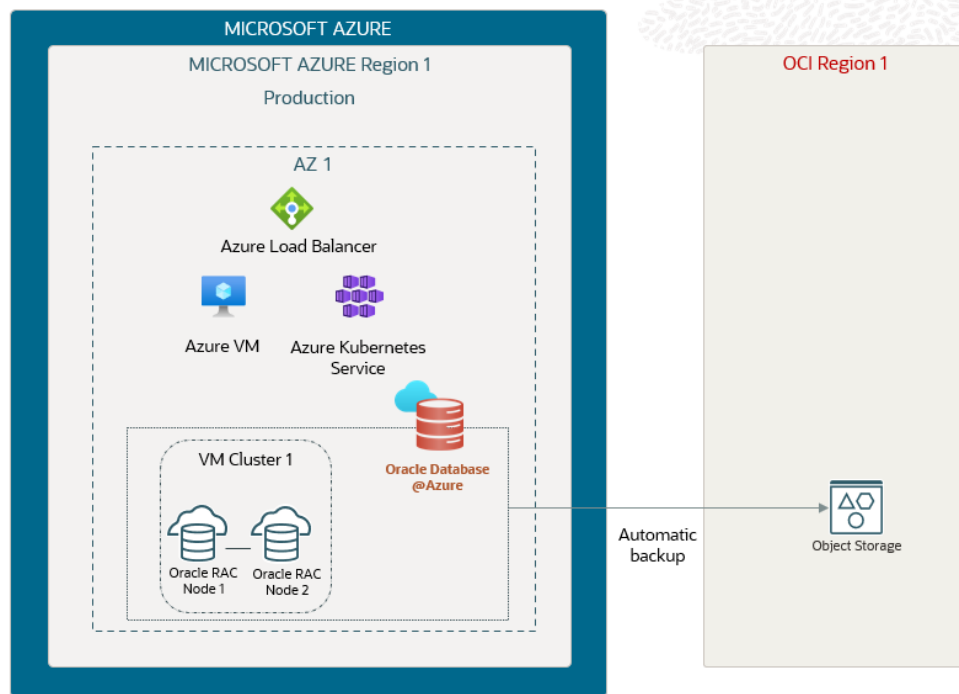


Figure 6: RAC High Availability Architecture in Oracle Database@Azure for departmental workloads

It starts with Oracle RAC, which allows customers to run their databases non-stop while scaling and patching or when faced with hardware outages. RAC distributes a customer’s Oracle Database across an Exadata VM cluster and allows the number of cores per server and the number of servers to scale up or down in an online manner. This means that customers using Exadata Database Service in Azure data centers can scale consumption to meet changing workloads, can patch and upgrade software, and can have database servers and storage serviced without incurring application-level downtime. RAC is an essential part of meeting customer requirements for high availability and scalability for their business-critical applications. In the cloud, RAC is only supported on OCI – which includes the Exadata Cloud Infrastructure located in Azure data centers. In addition, every Oracle Database@Azure database is automatically backed up to object storage in at least one OCI region. Remember, there are no data egress fees for database traffic between Azure and OCI.

That’s just the beginning. To provide even greater availability, Oracle Database@Azure can use Oracle Active Data Guard within or between Azure cloud regions. Set up one region as a standby, or run workloads in each region while using the secondary region for DR.

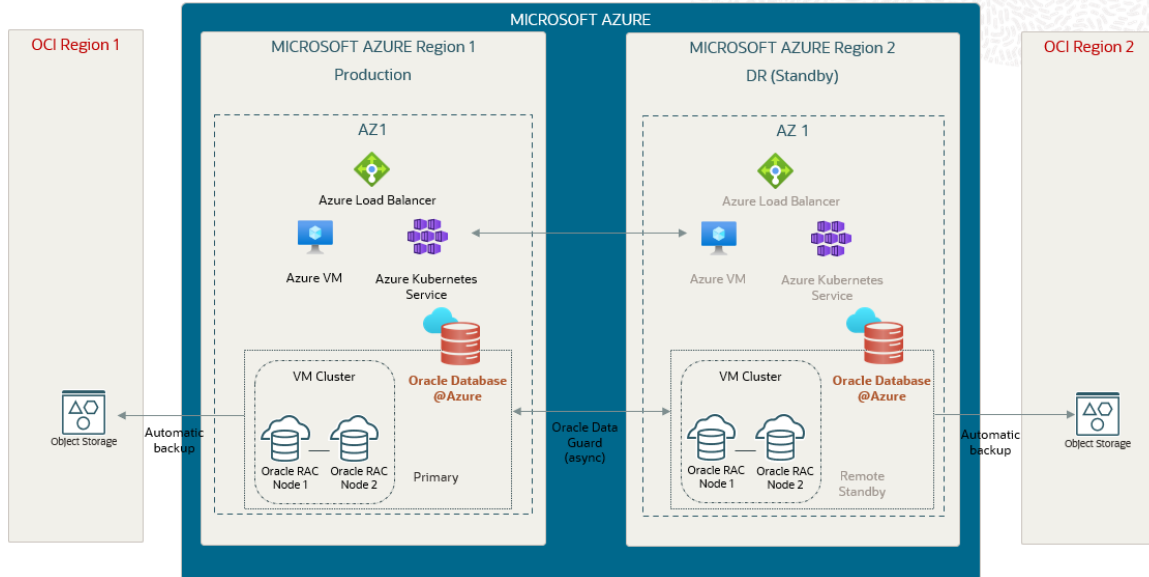


Figure 7: RAC High Availability + DR (Business-Critical) in Oracle Database@Azure

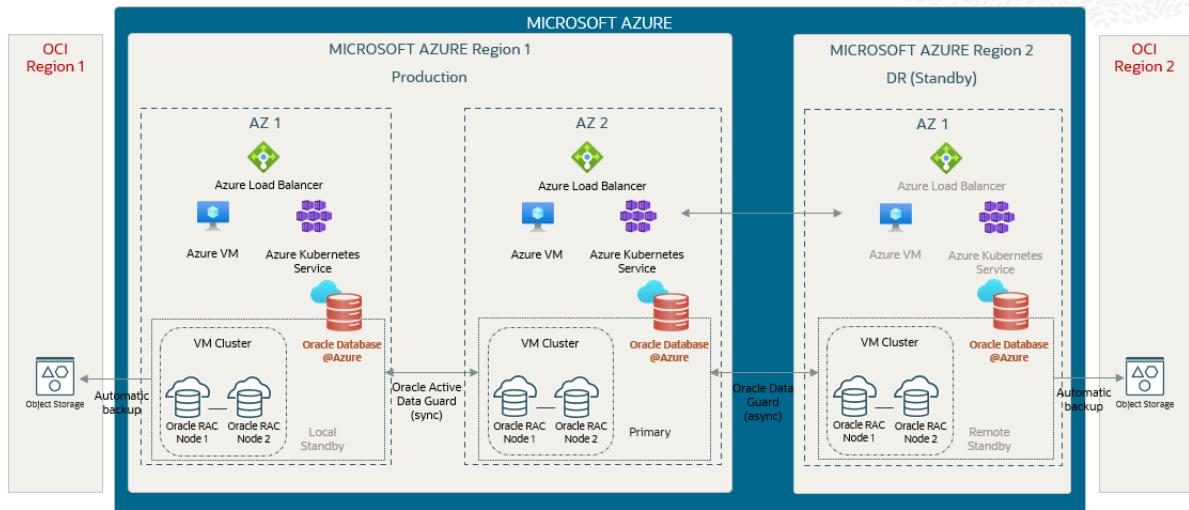


Figure 8: RAC High Availability + Additional DR (Business-Critical) in Oracle Database@Azure

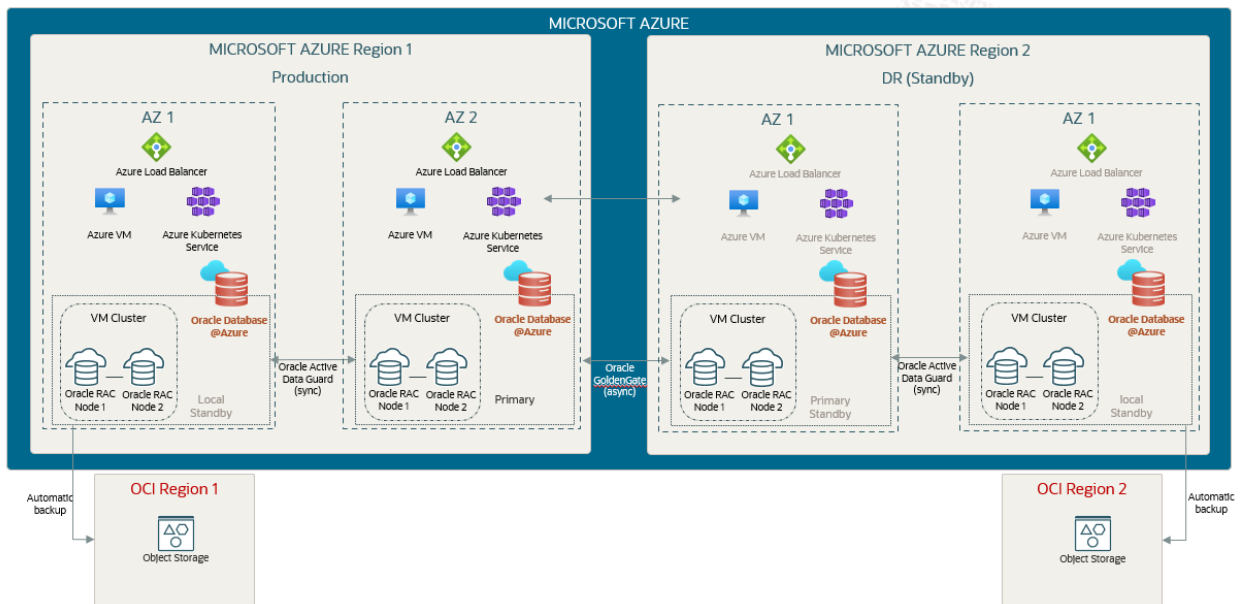


Figure 9: RAC High Availability + Full DR and Business Continuity in Oracle Database@Azure

The combination of Exadata’s high availability design, its intelligent software, and the implementation of Oracle Database@Azure means that HA, DR, and business continuity obstacles are eliminated.

#### Non-Disruptive Data Migration On-Premises to Oracle Database@Azure

Another multicloud adoption obstacle Oracle Database@Azure removes is the complexity of migrating databases from the customer premises to the Azure cloud. Oracle Zero Downtime Migrations (ZDM) uses Data Guard to orchestrate those migration without disruptions to the databases. Being able to easily migrate on-premises Oracle databases to Azure is the eighth game changing benefit achievable with Oracle Database@Azure.

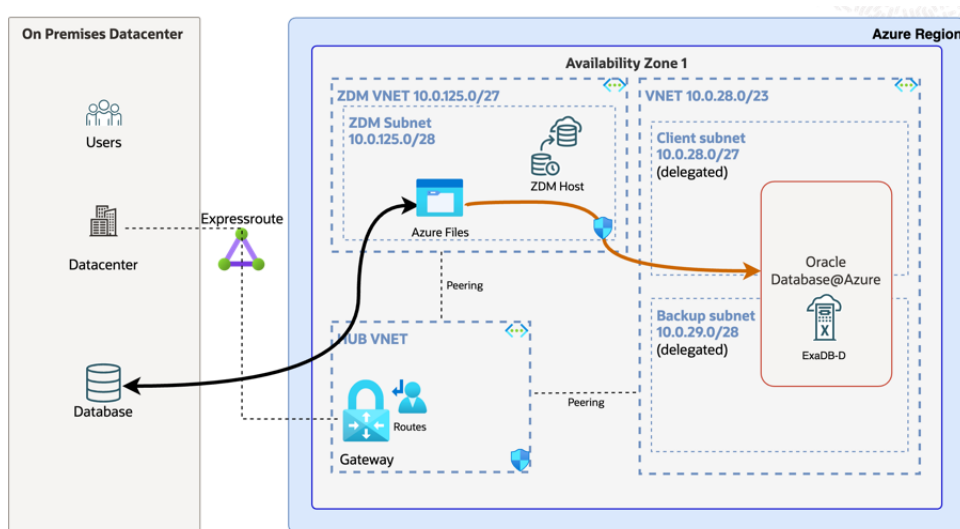


Figure 10: Oracle Zero Downtime Migration

### Simpler and Lower Costs

Multicloud environments have historically been complicated to purchase and pay for, involving payments to at least two CSPs and often additional third parties. The simplicity of buying and paying for Oracle Database@Azure and the cost savings it provides represent its ninth game changing benefit.

Like all cloud services, Oracle Database@Azure has a cost component driven by how long the processing takes. The faster a given job runs the less it costs. The Exadata platform with high-performance AMD EPYC processors is the fastest platform for running Oracle Database in the cloud. Its speed allows customers to run jobs in less time and with less processing cost, improves employee productivity which saves money throughout the organization, and enables both more concurrent end-user customers and more in-depth analysis of corporate data – both of which can lead to increased revenue.

With Oracle Database@Azure, the prices for Oracle Database services in Azure are the same as in OCI data centers and customers are eligible to receive Oracle Support Rewards on the amount they spend there. The innovative Oracle Support Rewards program enables customers to get a \$0.25 to \$0.33 credit toward their tech license support costs for each dollar they spend on OCI services – irrespective of whether the OCI cloud services are delivered in OCI or Azure data centers.


Like with OCI, Oracle Database@Azure customers have the choice of license included pricing which gives them access to all Oracle Database functionality or bring-your-own-license (BYOL) pricing which lets them use existing licenses with low-cost hardware-only consumption. Organizations with Unlimited License Agreements (ULAs) can save even more by bringing those licenses to the Azure cloud and using them in Oracle Database@Azure. Customers with ULAs don't have to purchase new licenses or pay additional license fees unless they choose to use a service or feature that is not covered in their ULA.

The biggest financial benefit for many customers may be the ability to purchase Oracle Database@Azure through a private offer on the Azure Marketplace, draw down their Microsoft Azure Consumption Commitment (MACC), and get one invoice from Microsoft that includes charges for Oracle Database@Azure and other Azure services.

Oracle Database@Azure will lower costs for those customers who need Azure trusted applications integrated with Oracle trusted databases. For these customers, this is a simple decision.

### Conclusion

Oracle Database@Azure is a true multicloud game changer. This latest program in the OCI and Azure partnership combined with their previous programs, eliminates, or at least greatly mitigates, customer perceived obstacles to adopting multiple clouds.



Multicloud Obstacle		Oracle Database@Azure Benefit
1	CSP lock-ins with high data egress fees	Eliminated since all data is in the same data center
2	Many single-purpose database services are required	Exadata Database Service and Autonomous Database running on Exadata Cloud Infrastructure
3	Multicloud integration deficiencies	Azure native access to Oracle databases running on Oracle Database@Azure
4	Speed-of-light latency issues over distance	Reduced by two orders of magnitude
5	Dissimilar processes and ecosystems across CSPs	Eliminated with integration of Oracle Database@Azure monitoring, observability, and management into Azure
6	Complicated/costly multicloud HA, DR, and business continuity	Simplified and automated based on 40 years of Oracle experience
7	Disruptive and time-consuming database migrations from on-premises to the cloud	Migrate database workloads “as-is” with automated tools that do not require application and database refactoring

All multicloud experiences should be this transparent. Oracle Database@Azure is indeed a multicloud game changer that’s showing the industry how multicloud should be done!

#### For More Information Oracle Database @ Azure

Go to:

[Oracle Database@Azure](#)