



Your Multicloud Advantage:

5 Benefits Hyperscaler Partnerships Deliver Now

Innovative deals among major cloud providers can pay off for you. Here's how.



Table of Contents

Are you ready for today's multicloud?	3
The open cloud	4
Five questions to ask to optimize your multicloud strategy	6
1. Is our data in the right place to deliver maximum value?	7
2. Will AI change how we use our data, and thus our cloud strategy?	8
3. How will regional needs impact cloud functionality?	9
4. Are we confident about securing our optimized multicloud?	10
5. Do we feel comfortable with cloud business continuity and disaster recovery?	11
Bring Oracle Database to your preferred cloud	12

Are you ready for today's multicloud?

Most enterprises have used multiple clouds for years. But having accounts with a variety of providers is very different from building an IT strategy designed to capitalize on the benefits of today's multicloud offerings. Without such a plan, sharing critical business data across applications is complicated and difficult—and in many cases, costs more than the results are worth.

The newest hyperscaler partnerships enable organizations to maintain their database and cloud technology preferences—no more compromising one or the other. That can improve business outcomes. Here's how.



The open cloud

Many enterprises drifted into using multiple clouds, either through M&As or because of decisions made over time by different work groups and departments. Maybe your sales team invested in an ecommerce platform from one cloud provider that required using a database running on-premises. On the other side of the company, perhaps your developers decided to implement a platform from a different cloud provider that suits their collaborative style and uses data from the same local database.


Your company is using multiple clouds, but in a way that requires extract, transform, and load (ETL) processes that create separate copies of data on-premises and in each hyperscale walled garden. This can cause a variety of problems, including security risks and stale information being used for decision-making in an era of constant data generation.

Why ETL slows you down

Extract, transform, and load is the process of gathering data from multiple sources and combining it to support reporting and analysis. Extraction involves identifying required data, copying it from its sources, and transporting it to the target data store. Because the extracted data is in its original form, it needs to be mapped, validated, deduplicated, and transformed to be usable.

That all takes time, and once data is pulled from its source, it's static, so reports based on it typically won't reflect the most current information. ETL can also be an error-prone process that's difficult to scale, which is why batch processing is sometimes used. If sensitive data is involved, moving it may add compliance and security risks. Setting up, maintaining, and managing complex direct network connections is one way to avoid batch ETL, but these links introduce complexity, latency, and administrative overhead.

And of course, all that adds costs. Hence, the data used by applications in walled gardens often isn't as complete as you'd like, which can lead to missed business opportunities. For these reasons, data-driven organizations are working to minimize use of ETL. That's a big driver for a more deliberate multicloud strategy, as we'll discuss.





Fortunately, the notion of an “open cloud” is within reach thanks to providers turning competition into cooperation. These partnerships deliver choice and flexibility, meaning that while it’s likely—even expected—that enterprises will work with multiple cloud providers, they need no longer sacrifice the timeliness or security of their data to select the best tool for the job.

A multicloud plan is particularly important for organizations that run compute- and data-intensive workloads in the cloud or that want to make data sources continuously accessible across applications. By employing a smart mix-and-match approach to data, workloads, and cloud providers, you can develop a custom solution to drive success while helping maximize security and laying the groundwork for a sustainable long-term technology strategy.

Your organization likely has many of the pieces in place already.
Now how do you make everything work seamlessly together?

Five questions to ask to optimize your multicloud strategy

Especially for early cloud adopters and those with decentralized IT purchasing, the multicloud reality tends to be more of a patchwork than a seamless experience. It's often also more expensive than it needs to be, with duplicative services and strategic yet unintegrated data stores running across a combination of on-premises data centers and clouds. Connections between those data stores and cloud-based apps can be slow and expensive to set up and use.

These pitfalls were often unavoidable—until recently. Now, with innovative agreements between Oracle and AWS, Oracle and Microsoft, and Oracle and Google Cloud, a data-driven enterprise multicloud strategy can deliver a seamless experience that helps improve processes, uncover hidden insights, and ultimately save money. A complete instance of Oracle Cloud Infrastructure (OCI), along with its associated database services, now resides in AWS, Azure, and Google Cloud and can be used with services from those cloud providers.

1. Is our data in the right place to deliver maximum value?
2. Will AI change how we use our data, and thus our cloud strategy?
3. How will regional needs impact cloud functionality?
4. Are we confident about securing our optimized multicloud?
5. Do we feel comfortable with cloud business continuity and disaster recovery?



1 Is our data in the right place to deliver maximum value?

A common hurdle enterprises face is poor alignment between application requirements, data availability, and provider capabilities. However the disconnect happened, there needs to be a plan for getting these applications secure access to complete and current data.

For example, a retailer tracks its ecommerce website activity, including user sentiment, in a cloud database. But transaction data is stored in an on-premises database, while data for the ecommerce site itself lives with a separate provider and is optimized for performance. Because these databases exist in silos, analysis requires an ETL process and then manual correlation. Interconnectivity hurdles, including data egress costs, latency, and data inconsistencies between parameters such as formats and time scales, further hinder analysis.

With new agreements between hyperscalers, Oracle houses its cloud offerings directly within AWS, Google Cloud, and Microsoft Azure data centers. These “coresident” services allow you to select and run best-of-breed technology, including your Oracle databases, as a native service of your cloud vendor, taking millisecond latencies down to microseconds and addressing security, collaboration, and governance issues. Now our retailer can run its enterprise database right within the cloud that hosts its ecommerce app.

Among surveyed organizations, 75% have increased their investments in data lifecycle management due to GenAI, according to Deloitte’s 2024 State of Generative AI in the Enterprise Q3 report.¹

¹ [Deloitte 2024 semiconductor industry outlook](#)

2 Will AI change how we use our data, and thus our cloud strategy?

There are a lot of reasons to choose a hyperscale cloud partner. Perhaps you've been working with AWS for years or Azure offers familiar development tools that lead to quick ROI. Maybe your ecommerce team found what it needed in Google Cloud or Oracle Cloud Infrastructure had the right combination of low cost and compatibility with your strategic data stores. Whatever your reasons, you ended up creating islands of data. That may not have been an issue in the past since there's lots of flexible computing power and other resources in each cloud to analyze select data. As such, data egress fees may not have been a big concern—processing was local and exporting results didn't drive substantial charges.

Generative AI will change that calculus. By giving GenAI analytics tools access to only the data stored with your chosen cloud provider, you may miss the nuances that AI is great at uncovering. The time lag and management headaches involved in moving data out of each cloud and into wherever you do the bulk of your analytics—and, potentially, AI fine-tuning—may also be an issue because now you must manage multiple instances of the same data.

And the problem also applies the other way around:

As those cloud apps get their own AI smarts, their analysis will benefit from access to more complete data, and that's where both data egress fees and data latency can present big problems. For instance, if you're running ecommerce analysis in Google Cloud, those results would likely be made richer by access to supply chain, inventory, and sales data stored in your ERP and CRM systems. A multicloud solution that can bring that data to where it's needed while maintaining it in systems where governance and access controls are already worked out can solve a lot of problems.

The bottom line: For an organization with ambitious AI plans—or that simply wants to maximize use of its data—being able to run its trusted database anywhere, as a service and with the same features and pricing, is a game changer.

3 How will regional needs impact cloud functionality?

There are many reasons for maintaining data in separate regions around the world. Different locations may have specific regulatory requirements and restrictions. Latency between regions may impact performance for some workers, or a company may have far-flung edge devices. All these factors create demands that, if not attended to, can have a negative impact on user experience, data analysis, compliance, and much more.

For instance, an app with a global user base may need to tackle latency issues by using cloud regions in North America, Europe, and Asia, and it may also need to maintain a regional customer database to address, for example, GDPR requirements for data residency and processing.

Some providers have networks of cloud regions that can help address strict regulatory guidelines. There are also database options for migrating data into local regions, with your choice of provider, that can solve a multitude of problems.

A key multicloud benefit: Freedom of choice



Organizations generally select a primary cloud provider for consolidated security, billing, and other logistical reasons. But exclusivity may create a situation where it's difficult to follow a best-of-breed approach or fully use the data you've spent years accumulating and managing.



A policy of using open standards with an eye on compatibility and portability allows companies to take a “best-of-cloud” approach. Similarly, developing applications using containers provides for portability should you decide to move applications between clouds or your own data centers.



A consolidated, universally compatible cloud-enabled database also gives you more flexibility. It allows for timely and unified data access for applications while helping minimize costs, complexity, and security and compliance issues.

4 Are we confident about securing our optimized multicloud?

The more data that's maintained in one primary system with a single governance and data protection model, the better—especially when that platform can continually identify and remediate threats, reduce the chance of human errors, and automate backup and recovery.

With a multicloud strategy, security and interoperability should go hand in hand. Think controlling access to highly sensitive data, encrypting communications among apps, and preventing single points of operational failure. In fact, a cloud service provider's specific suite of security and monitoring capabilities may drive optimization decisions, particularly in instances where specific workloads better align with a vendor's tools. An example is when [zero trust packet](#) routing is desirable or necessary.

However, the more players in a multicloud setup, the broader the potential risk due to complexity, scope, and capability gaps. Simplified management and centralized tools for identity and access management and resource monitoring can help IT teams simultaneously strengthen and streamline security.

Consolidating data into a trusted central cloud repository can help too. Each cloud provider offers different security tools and practices. The goal is to meet your organization's requirements in as simple a way as possible. In a perfect world, that adds up to unified security and governance frameworks. However, the real world is often more complicated, so an ideal situation may be to use one provider for a security model framework and one for a governance framework for strategic data.



5 Do we feel comfortable with cloud business continuity and disaster recovery?

Even if you're running the most optimized, secure computing environment possible, fires, power outages, hurricanes, malware, and other unexpected events happen. Hyperscalers have extremely advanced disaster recovery and redundancy capabilities, and organizations can use their multicloud's inherent diversification to minimize risk.

Doing so effectively, however, takes planning. Do your providers' data storage and failover connectivity capabilities provide layers of redundancy, helping ensure there's no single point of failure, for example? In the case of provider partnerships, it's important to consider how layering disaster recovery features from each can deliver protection that's greater than the sum of its parts. Today's multicloud options even make planned disaster recovery across multiple providers feasible.

The two types of partnerships

A multicloud approach works best when data flows smoothly between databases and applications. In the era of the open cloud, major tech companies have created partnerships to grant customers choice and flexibility. The two most important types of data partnerships are as follows:



Interconnection between providers

A direct private connection between two cloud providers supports low-latency, high-throughput data transfers across databases and applications without any adjustments to architecture or platforms.



Integrated databases at cloud providers

Now that OCI instances are available within AWS, Azure, and Google Cloud, the Oracle data management systems you use, such as Oracle Autonomous Database and Exadata, are also available as services from within each provider's data centers, along with other Oracle cloud services. The performance and ease of use allow for a broad array of use cases that were previously difficult or impossible.

Bring Oracle Database to your preferred cloud

To help you get the most out of your multicloud strategy, Oracle has partnered with major cloud providers to embed Oracle Database running on Oracle Cloud Infrastructure within [AWS](#), [Google Cloud](#), and [Microsoft Azure](#). It's a provisioned service, available from each provider's marketplace. Now you can enjoy the flexibility of running your mission-critical workloads in your preferred cloud providers' facilities and potentially accelerate the migration of your Oracle databases to the cloud. The result is a highly connected, low-latency experience that meets you where you are.

Accelerate innovation and cloud migration with the power and flexibility to deploy OCI Oracle Database services in any cloud. Combine the best of the cloud with your data to quickly build and modernize applications.

Learn more

Connect with us

Call +1.800.ORACLE1 or visit oracle.com/apac

Find your local office at oracle.com/contact

Copyright © 2025, Oracle and/or its affiliates. This document is provided for information purposes only, and the contents hereof are subject to change without notice. This document is not warranted to be error-free, nor subject to any other warranties or conditions, whether expressed orally or implied in law, including implied warranties and conditions of merchantability or fitness for a particular purpose. We specifically disclaim any liability with respect to this document, and no contractual obligations are formed either directly or indirectly by this document. This document may not be reproduced or transmitted in any form or by any means, electronic or mechanical, for any purpose, without our prior written permission. Oracle and Java are registered trademarks of Oracle and/or its affiliates. Other names may be trademarks of their respective owners.

