


Innovation White Paper

New Oracle

Compute Cloud@Customer

Showcases Oracle's Customer Advocacy



Delivering the
simplicity, agility, and
elasticity of a cloud-
based deployment

WHITE PAPER

By Bob Evans

Cloud Wars Founder
Acceleration Economy Co-Founder

ORACLE



New Oracle Compute Cloud@Customer Showcases Oracle's Customer Advocacy



By Bob Evans
Cloud Wars Founder
Acceleration Economy Co-Founder



Building on its 63% cloud-infrastructure revenue growth for FY23, Oracle has launched a new offering delivering rack-scale compute services from inside a customer's data center and compatibility with key Oracle Cloud Infrastructure (OCI) compute, storage, and networking services.

Called Oracle Compute Cloud@Customer, the new offering broadens Oracle's market-leading capabilities in distributed cloud and clearly delivers on chairman Larry Ellison's recent high-profile advocacy for providing cloud customers with unfettered choice in choosing the best cloud architectures, deployment locations, and services to realize their desired business outcomes.

That overarching philosophy from Ellison and the entire Oracle leadership team—including Oracle executive vice president and chief corporate architect Edward Screven—is positioning Oracle as the cloud provider whose products and services are built around what customers want today and will need tomorrow in the fast-changing world of digital business. This approach aligns well with the increasingly dispersed nature of enterprise operations – something that Oracle calls a distributed cloud – which the new Compute Cloud@Customer directly addresses.

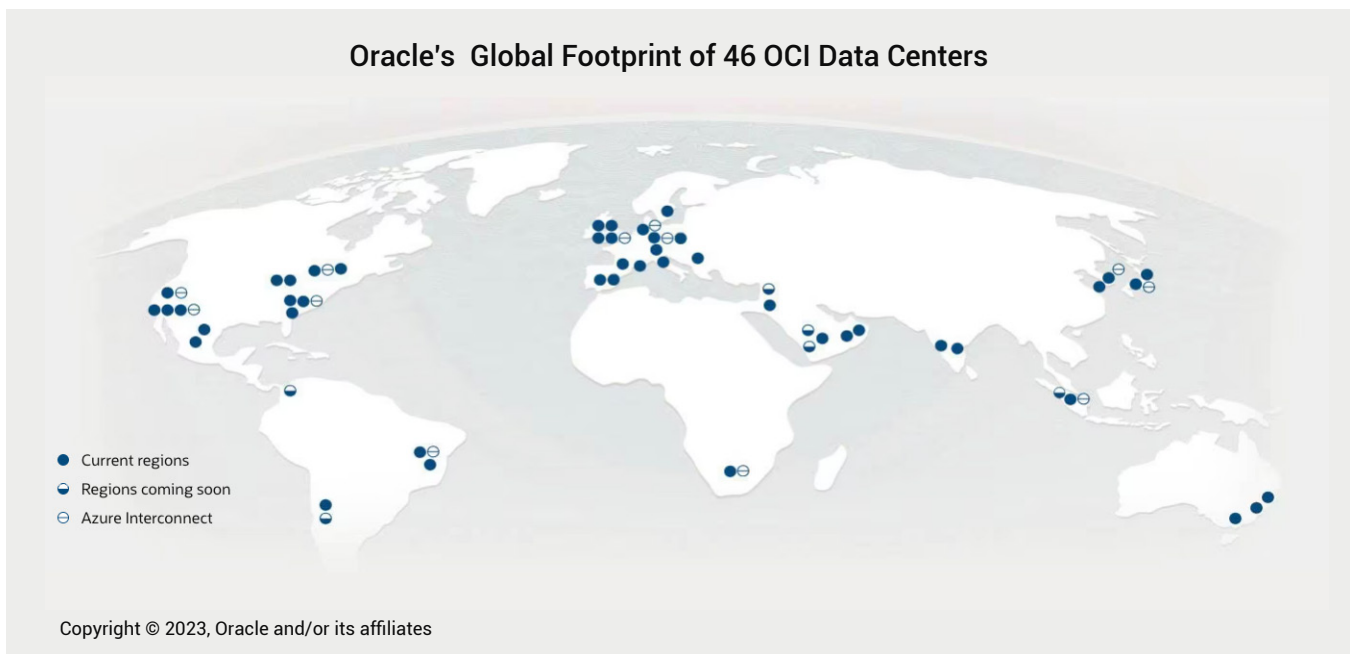
Oracle is taking a dramatically different approach to cloud infrastructure from that offered by the 3 original global hyperscalers : Microsoft, AWS, and Google Cloud; Oracle is now the fourth.

Oracle says OCI is a Gen 2 cloud that offers differentiated performance and value beyond what its competitors can offer.

Oracle has 46 OCI regions around the world and is rapidly adding more to achieve Ellison's goal of having OCI data centers in every major city in the world. To do this, Oracle is building a large number of relatively small data centers offering the full range of OCI services with uniform pricing across all public cloud regions as well as OCI Dedicated Region sites – which add substantially to the above number. In contrast, its competitors are building a relatively small number of massive data centers that don't all offer the same services at the same price to customers across regions.

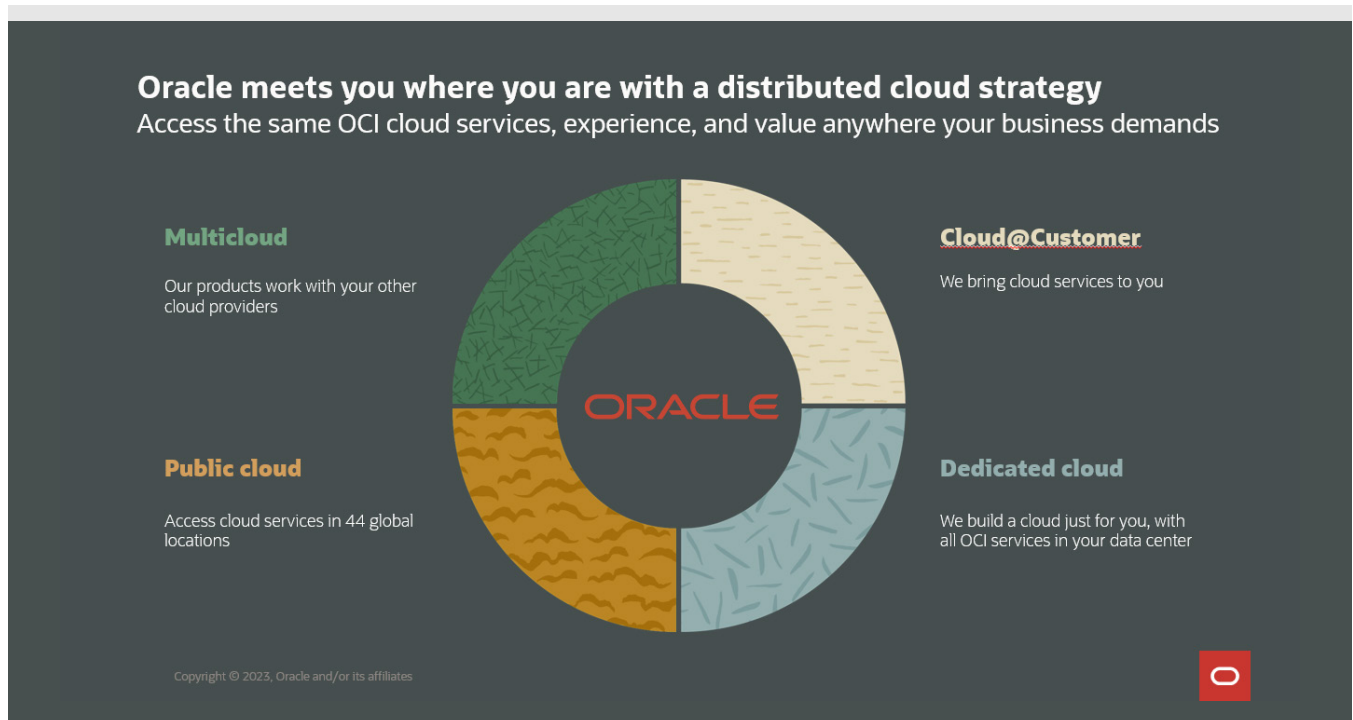
OCI is much faster and less expensive than its competitors because it is using newer technology in innovative ways. Remote Direct Memory Access (RDMA) lets Oracle create high-performance OCI Superclusters for AI and Exadata systems for running Oracle Database that run rings around other offerings in the marketplace. In the cloud time is money; faster means that customers can innovate more quickly and less expensive lets them significantly reduce costs.

With the advanced technology used within OCI, it offers unmatched cybersecurity with everything from always-on encryption to centralized database security posture monitoring and isolated management networks that can reduce the chance of cyberattacks.



The rise of the Oracle Distributed Cloud

These public cloud capabilities comprise the core of Oracle's distributed cloud. It also includes not only OCI services, multi-cloud deployments, and the ability to deliver services wherever organizations need them with Cloud@Customer and OCI Dedicated Region options in their data centers.



Oracle's overall Cloud@Customer portfolio has achieved great success among customers that want the flexibility, scalability, and economics of the cloud but with the technology located within their own data centers to meet data residency or low latency requirements. Those systems are running in 60 countries and are being used by customers in a wide range of industries.

Perhaps the best-known of Oracle's distributed cloud platforms is Exadata Cloud@Customer, which Oracle began offering in 2017. The recently announced Exadata Cloud@Customer X10M lets organizations run Autonomous Database and Exadata Database Service as a managed OCI service – all while benefiting from cloud automation and consumption-based pricing.

Oracle also offers OCI Dedicated Region which can be installed in customer data centers and can run all of Oracle's infrastructure and platform services as well as SaaS applications. OCI Dedicated Region starts at 12 racks and can scale up to virtually any size an organization might need.

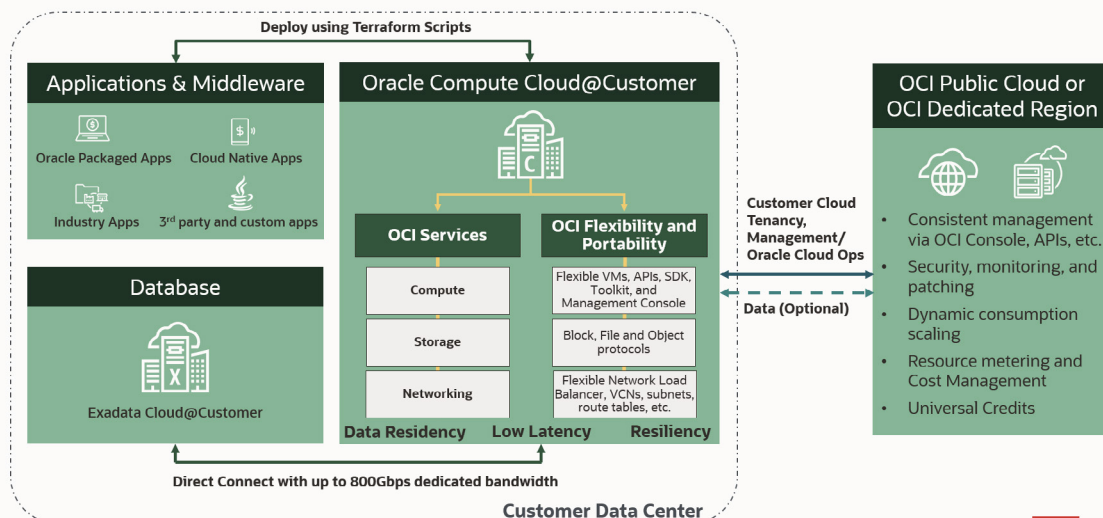
With Compute Cloud@Customer, Oracle is reacting directly to customers requests—well, more and more they are “demands”—for an option that lets them run OCI compute and storage services in their data centers – reducing costs through consumption-based pricing, and offloading system administration to Oracle. Compute Cloud@Customer starts at roughly one-quarter of a rack to complement Oracle's extremely successful Exadata Cloud@Customer platform which starts at the same scale. It's much smaller than an OCI Dedicated Region which makes it extremely attractive to smaller organizations and those that need highly distributed cloud resources.

However, let's be clear, while Compute Cloud@Customer and Exadata Cloud@Customer are scaled down from an OCI data center or Dedicated Region, they are still incredibly powerful. Both systems start out with multiple servers based on 4th generation AMD EPYC™ processors and over 150 terabytes of usable storage, let customers configure the amount of compute and storage they need, and then pay consumption only on the amount of resources they actually use. Oracle's systems are much more powerful than anything offered by OCI's primary competitors: AWS, Google Cloud, and Microsoft. I'll offer some details on that shortly.

What makes Compute Cloud@Customer tick?

Oracle Compute Cloud@Customer is a fully managed, rack-scale infrastructure that lets customers run applications and middleware on premises using OCI compute, storage, and networking services. By running the same services that are available in OCI regions, customers get the automation, simplicity, agility, and elasticity of a cloud-based deployment.

Oracle Compute Cloud@Customer—OCI compute, storage, and networking in a rack



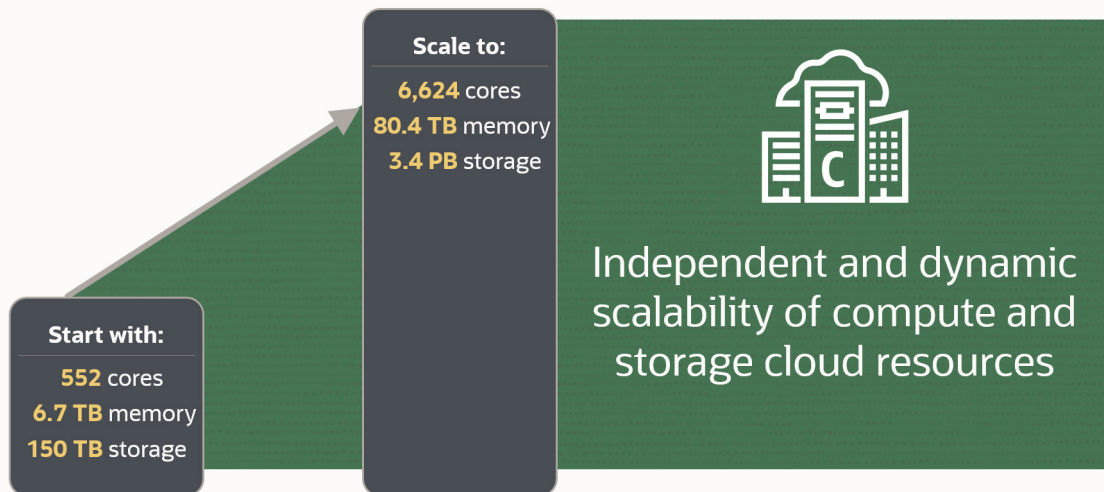
Copyright © 2023, Oracle and/or its affiliates



As noted earlier, Compute Cloud@Customer is built using 4th generation AMD EPYC processors and high-performance storage that supports block, file, and object protocols. In the base system, there are two 96-core AMD EPYC processors in each of three compute servers and 150 TB of high-availability storage capacity.

Once Oracle reserves 8 processing cores per server and the memory needed to run the hypervisor—for which they don't charge. This entry system lets customers deploy a mixture of flexible VM shapes with up to 96 cores and 960 GB of memory per VM with a total of 552 processor cores and 6.7 TB of memory on three compute servers. Customers can then expand this platform's infrastructure in similarly sized chunks to get 2,208 processor cores and 26.8 TB of memory in a rack, or 6,624 cores and more than 80 TB of memory in a 3-rack configuration. As with Exadata Cloud@Customer, compute and storage scale independently, allowing customers to align configurations to their applications and database needs. There's a distinct advantage for Oracle when Compute Cloud@Customer's rack-scale capabilities are compared to the competition—as we'll see later.

Compute Cloud@Customer scalability



Copyright © 2023, Oracle and/or its affiliates



But it's not just about how many cores customers can pack in a rack. Oracle's software has been tuned to get the most out of the high core count and performance of the AMD EPYC processors. Oracle's approach benefits both scale-out cloud native applications that can use hundreds or thousands of cores and traditional packaged applications that benefit from scale-up capabilities. With nearly 1 TB/sec of DDR5 memory throughput per server (almost 3 TB/sec in a base system and 12 TB/sec in a full

rack), customers will be able to run a mix of memory bandwidth intensive analytical applications and latency-sensitive transactional applications at peak performance.

For full-stack applications using Oracle Database, Compute Cloud@Customer can directly connect to Autonomous Database and Exadata Database Service running on Exadata Cloud@Customer platforms in the data center. Up to 800 Gbit/second of directly connected dedicated bandwidth minimizes application to database latency and eliminates contention with data center networking. The latest Exadata Cloud@Customer X10M platform uses the same 96-core 4th generation AMD EPYC processors found in Compute Cloud@Customer so full stack OLTP and OLAP environments run as fast as possible.

9 Things Customers can do with Oracle Compute Cloud@Customer

When asked by Oracle what type of capabilities this new offering should have, customers said they wanted this new rack-scale option to be able to handle a range of requirements with consistent and complete availability across the globe. Here's a list of 9 breakthrough features and attributes available with Compute Cloud@Customer, along with some breakdowns of how those stack up versus competitors.



1. *Develop once, run anywhere with superb performance*

Cloud-native apps can be developed with the new service and can run, with zero modification, across the entire range of Oracle's distributed-cloud environments. The services—APIs, SDKs, and Terraform scripting—found on Compute Cloud@Customer are also available on OCI, making it easy for customers across every industry to quickly deploy applications wherever they're needed. Then scale them up to 2,208 processors in a rack—4X AWS Outposts—and up to a total of more than 6,624 processing cores and 3.4 PB of storage in three racks.

2. *Deploy at scale*

The elastic nature of Compute Cloud@Customer lets customers scale up cloud-native apps as well as move and run existing traditional apps from Oracle or other third-party providers by leveraging flexible VMs with up to 96 cores and 960 GB of memory per VM—that's up to 3X the compute of other cloud on-premises platforms.

3. *Modernize full application stacks*

Customers can optimize the performance of their applications by running OCI compute and storage services in their data centers and connecting to Oracle Database services running on Exadata platforms for the highest end-to-end performance. Customers with direct connections to Exadata Cloud@Customer have reported up to a 50% reduction in latencies between applications and databases. Applications can then be containerized on Compute Cloud@Customer to achieve more agile deployments.

4. *Simplify compliance with data residency regulations*

The Compute Cloud@Customer platform and the OCI services running on it have been designed to enable customers to easily meet all data-residency requirements, including data locality as well as where it's processed. With Cloud@Customer on-premises options customers can deploy data and applications wherever needed to stay within designated national borders. This builds on Oracle's proven remote management and Oracle Operator Access Control that Oracle has been delivering with Exadata Cloud@Customer. These capabilities are used by banks and government institutions worldwide to meet ever evolving data residency and security requirements..

5. *Reduce costs with dynamic scalability*

Customers pay only for what they use—no paying for consumption on everything installed even if they're using only 20%. Customers pay a low cost per month for the Compute Cloud@Customer platform and then pay the same per-second consumption cost for compute and storage as on OCI. In addition, customers can scale up resources during peak times and down when not needed. Simple metering is a key feature that OCI's competitors such as AWS Outposts don't offer, and that creates a big problem for enterprises with variable workloads—pretty much everyone. Many

on-premises cloud platforms charge the same price whether customers consume services on the platform or not—which is not the case with Oracle. In any case, Oracle's costs are lower than those from AWS at full consumption and customers can save 50%, 60%, or 70% at moderate consumption rates.

6. *Deploy close to your customers*

For customers in telecommunications, gaming, Internet of Things (IoT), or have end users that require quick responses and real-time data processing then where platforms are located is critical. Historically, customers had to purchase systems and then deploy and manage them across geographies or be stuck with the “closest cloud data center” which could introduce long delays. With Compute Cloud@Customer, customers can deploy OCI services where they are needed, not have to manage the platforms, and get the lower, consumption-based costs needed to succeed.

7. *Achieve global consistency*

In a time when every company is looking to vigorously control expenses, Oracle Cloud@Customer offers a big benefit to customers that AWS Outposts can't match. Oracle offers uniform pricing for OCI services around the world, in OCI regions, multicloud environments, dedicated cloud, and Cloud@Customer deployments. This is a big win for customers that are looking to expand internationally since they can plan their budgets regardless of location. Customers can deploy Cloud@Customer platforms in more than 60 countries without needing local system administration experts since Oracle handles that.

8. *Simplify management*

Under the hood, customers can take advantage of a wide range of new technologies to help their organizations run faster, spend less money, and achieve greater agility.

- Portable infrastructure as code: Workloads located wherever they need to be
- Unified user ID management across Oracle's entire distributed cloud: OCI Public Cloud, Dedicated Region, Exadata Cloud@Customer , and Compute Cloud@Customer
- Seamless updates and patches: Fully automated process defined by the customer
- Continuous operations: Compute Cloud@Customer will continue to run if its connection to the home OCI Region or Dedicated Region is not available. Certain functions like onboarding new users or billing won't happen until control plane connectivity is restored, but most data lifecycle management operations continue uninterrupted.
- Single pane of glass management: Unified console and consistent tools to deploy and manage local copies of software and data, track usage and system health, protect data, and perform other lifecycle tasks. Enterprises such as telecom companies that have with

operations around the globe can then manage their distributed cloud environment from a central location using the OCI control plane to reduce complexity and cost.

- Cost-efficient disaster recovery: Secondary sites used for disaster recovery can operate with significantly lower consumption costs than primary sites since keeping remote data up to date does not require much processing.

9. *Lower energy consumption*

96-core 4th generation AMD EPYC processors use substantially less power per core than other processors—both when running workloads and when idle. AMD reports that a configuration with 2000 VMs takes 29% less power than on a similar configuration based on other x86 processors, which means 29% lower costs for electricity and cooling.

In the design of Exadata X10M, Oracle engineers managed to deliver on a remarkable design goal: The power usage of a Oracle Exadata X10M quarter-rack system is 27% lower than that of the previous generation with the same number of processing cores and memory. This includes lowering the electricity bill for 384 cores of database servers by 50%

How does Oracle Compute Cloud@Customer stack up versus competitors?

A big part of Oracle's promise to customers is that its cloud products and services are designed to not only deliver customers unlimited choice and flexibility, but also levels of performance, speed, and lower costs that none of its competitors can match. Here's a comparative-feature chart put together by Oracle that shows the dramatic differences in what customers get with Oracle Compute Cloud@Customer versus what other vendors are offering—and as you check out the differentiation, here are a few highlights to consider:

- Oracle Compute Cloud@Customer offers up to 2208 cores in a rack, which is about 4X the number that AWS Outposts offers. More cores means better overall performance
- Oracle offers the same services and uniform pricing in all cloud regions worldwide. AWS does not.
- Oracle says that Microsoft's Azure Stack Hub is "not really a cloud model" because customers are required to purchase a system from a hardware vendor and then run Microsoft's software on that hardware. Unlike Oracle's pure cloud model, Azure Stack Hub doesn't offer single-vendor support and simple pricing.

Oracle Compute Cloud@Customer delivers superior distributed cloud capabilities

	Oracle Compute Cloud@Customer	AWS Outposts	Azure Stack Hub	Google Distributed Cloud Edge
Cores per rack	2,208 cores	576 cores	1,536 cores	288 cores
Cores per VM	96 cores	48 cores	32 cores	48 cores
Storage capacity	3.4 petabytes	435 terabytes	1.5 petabytes	19.8 terabytes
Storage services (block, object, and file)	✓	No File	No File	No Object
Storage flexibility (change protocol on demand)	✓	No	No	No
Exadata Cloud@Customer Direct Connect	✓	No	No	No
Consistent worldwide pricing	✓	No	No	No
Cost for enterprise support	Included	Additional	Additional	Additional
Cost per core	\$53 per month	\$143 per month	No cloud model	\$84 per month*

* Includes an additional \$12,500 per month for required Premium Support

Copyright © 2023, Oracle and/or its affiliates



Final thoughts

Oracle offered 4 key takeaways about the differentiated value of its Oracle Compute Cloud@Customer—and I'll share this in a moment—but I want to mention that Oracle's biggest takeaway might be that founder and chairman Larry Ellison has fully embraced being a staunch and unwavering advocate for customers as they move to the cloud. This advocacy covers everything from a relentless focus on best-in-class technology to unmatched price-performance to simplified procurement and pricing to centralized management. Not to mention distributed cloud in any location for everyone!

As a result, customers can be assured that Oracle's cloud products, services, and go-to-market policies will keep those attributes at the top of Oracle's priority list.

And on the specific case of Oracle Compute Cloud@Customer, here are four points I mentioned previously that Oracle believes best exemplify the business value that can be created with this latest addition to the Oracle distributed-cloud family.



Compute Cloud@Customer

Copyright © 2023, Oracle and/or its affiliates

Key takeaways

1

Run the same OCI services anywhere and easily manage your distributed cloud environment from a single control plane

2

Dynamically scale up to meet peak requirements or scale down to reduce costs

3

Simultaneously run existing applications and new cloud native ones, with high availability

4

Meet data residency, security, and latency requirements



To sum up, Oracle Compute Cloud@Customer provides a compelling application and middleware platform with superior capabilities and lower costs than competitive cloud on-premises offerings from AWS, Azure, and Google. Enterprise customers with data locality or latency requirements should look into the option of deploying cloud on-premises to gain cloud automation, flexibility, elasticity, and consumption-based economics. For this, there is no better tandem than Oracle Exadata Cloud@Customer and Oracle Compute Cloud@Customer.