



Oracle Exadata Cloud@Customer X10M

EPYC Innovation and Performance Far Beyond a Hardware Refresh

by, Marc Staimer

September 2023

Research Premise

Platform vendors have historically hyped product refresh releases as the greatest innovation since the invention of sliced bread. But when looking past the hyperbole, the refresh generally boils down to the latest CPU, media, and interfaces with some minor software tweaks. The hype will emphasize the increase in cores, faster media, interconnect, and a noticeable, albeit not overwhelming, performance increase. Performance increases generally run well under 100% even when cores double. Hardly revolutionary. More of a 'meh'.

Oracle's latest Exadata Cloud@Customer X10M release is considerably different. It's far beyond EPYC. And yes, that's a deliberate play on words. Although AMD 4th Gen EPYC™ processors play an important role in the Oracle Exadata Cloud@Customer X10M performance increases, it is just one of several innovations that

provide a total that is greater than the sum of its parts. Just look at the remarkable results as compared to the previous generation Exadata Cloud@Customer X9M:

- ⇒ **360%** faster queries.
- ⇒ **300%** higher transaction throughput.
- ⇒ **240%** faster in-memory analytics scans.
- ⇒ **≤ 17μs** SQL IO latencies representing a **10.5%** improvement over their already established leadership performance.

These are huge performance increases and substantially better than a refresh. This latest Exadata X10M Cloud@Customer release makes it crystal clear that Oracle's on-premises cloud database service has turned its extremely wide competitive performance lead into a chasm. Using a race car analogy, Oracle has lapped the field multiple times.

This Wikibon research examines how Oracle achieved these gains and what it means to customers.

- Fundamental user problems solved by the extraordinary performance of Exadata Cloud@Customer X10M.
- How it compares to the other competitive on-premises cloud database services.
- Why CFOs need to pay attention to how it substantially improves their bottom line.

In Wikibon's view, Exadata Cloud@Customer X10M delivers an impressive database performance advantage in combination with Oracle Autonomous Database and will continue to win new customers and expand Oracle's cloud database market share.

How Exadata Cloud@Customer X10M Provides Unique Extraordinary Performance



Achieving the aforementioned performance gains required innovations in the Oracle Database, Exadata software, Linux Kernel, and the AMD 4th Gen EPYC CPUs. The Exadata X10M Cloud@Customer hardware and system software is cooperatively co-engineered with the Oracle Database, open-source Oracle Linux distribution, AMD 4th Gen EPYC capabilities, and RDMA networking. Oracle Database alone delivers many features, functions, benefits, and performance unique to the Exadata Cloud@Customer X10M platform. The Exadata software is distinct from any other platform in that it is specifically designed to get the best performance out of all Oracle Database iterations including Exadata Database Service and Autonomous Database, and it's the only platform to run Autonomous Database services in customer data centers.

Exadata Cloud@Customer X10M Software Enhancements Versus X9M

- Leverages the latest Oracle Linux OS and kernel (Oracle Linux 8.7 and UEK6), enhanced to deliver automated performance scaling for the new very large core count 4th Gen AMD EPYC CPUs.
- Significantly boosts 'Smart Scan' to support greater parallelism and scale for more complex queries.
- Runs decryption and decompression algorithms significantly faster so storing and using secured data does not negatively impact performance.
- Adds storage index and columnar cache persistence that ensures consistent analytic workload performance after storage server maintenance.
- Provides open monitoring that transmits alerts and metrics to popular observability platforms.
- Supports enhanced Oracle Database algorithms including machine learning, graph, spatial, in-memory, and Blockchain tables.

Exadata Cloud@Customer X10M Hardware Enhancements Versus X9M

Bold = increases

Database Server

- Scale-out 2-socket database server featuring **AMD 4th Gen EPYC (Genoa) 96-core CPUs** per socket.
 - **3X more cores**
 - Up to **3.6x faster** database queries
 - Each core is **10% faster** than the previous generation for transaction processing
 - **20% faster** per core performance for analytics
- Up to 3TB DDR5 DRAM
 - **1.5X to 2X capacity increase**
- Greater DDR5 memory throughput
 - **2.5X faster** memory
 - Up to **2.4 faster** in-memory analytics scan
- Additional client connectivity
 - From **3 to 5 NIC cards per database server**

Internal RDMA Fabric

- 100 Gb/s Remote Direct Memory Access (RDMA) over Converged Ethernet (RoCE):
 - No change

Storage Server

- Scale-out 2-socket storage server featuring **AMD 4th Gen EPYC 32-core CPUs** per socket
 - **2X more cores**
 - Up to **2.6x faster** storage server queries
 - **30% faster** per core Smart Scan
 - Up to **35% faster** extreme flash decryption
 - Up to **25% more** columnar cache capacity in flash based on improve compression algorithms
- **1.25 TB DDR5 Exadata RDMA Memory (XRMEM) DRAM Cache in storage servers**
 - **XRMEM** replaced the discontinued Intel Optane of the previous X9M while helping lower SQL IO latency to under 17 microseconds
- Three tiers of storage
 - **XRMEM**, NVMe Flash, HDD
- EF storage server
 - Flash capacity for permanent storage increases from 51.2 TB to **122.88 TB**
- HC storage server
 - Flash capacity for caches increases from 25.6 TB to **27.2 TB**
 - HC storage hard drive capacity increases from 18 TB to **22 TB** per drive, giving each storage server 264 TB of raw storage capacity and 80 TB of usable capacity when used with triple redundancy

Fundamental User Database Problems Exadata Cloud@Customer X10M Solves

With the latest shining object being generative AI and large language models (LLM), the press appears to have forgotten that the vast majority of mission-critical applications rely on databases. Databases are essential tools for businesses everywhere from banking, finance, ecommerce, logistics, healthcare, manufacturing, energy, media, entertainment, and much more. However, there are problems with many database implementations. Problems such as complicated integration between database types and workloads, overall complexity, limited scalable performance, and database sprawl with resulting silos.

Exadata Cloud@Customer X10M provides considerable relief from those persistent user database problems and several others that IT organizations simply live with. These include low DBA and application user productivity; DBA turnover; training costs; stale data that reduces actionable insights; late-to-market products and services; shrinking market share; and sliding into a state of irrelevance.

How Exadata Cloud@Customer X10M Solves Those Persistent User Database Problems

Eliminating complicated database integration and overall database complexity

Making complicated multi-database integrations effortless starts with the Oracle Database. Oracle Database converges OLTP, data warehouse, machine learning (ML), blockchain, JSON, XML, spatial, graphic, time series, data lakes, and more. With Oracle Database you have one copy of the data that doesn't need to be moved around. But that's just the start.

Reducing the amount of specialized human expertise and manual administration needed to implement, operate, manage, and troubleshoot database environments lets organizations dramatically reduce human error, time waste, and DBA turnover. Exadata Cloud@Customer makes tremendous progress in this direction by eliminating the need for IT staffs to administer database infrastructure and consolidating operations and administration for individual databases. Oracle Autonomous Database completes the simplification by automating database administration to such a degree that manual efforts are almost eliminated. Autonomous Database only runs on the Oracle Cloud Infrastructure (OCI) public cloud regions, Exadata Cloud@Customer X10M, and OCI Dedicated Region—a dedicated complete OCI region with identical hardware and software capabilities to OCI public cloud regions but on the customer's premises.

Oracle Autonomous Database is the first and currently only autonomous databases in the market today. It's built by combining Oracle's more than 4 decades of best practices for running mission-critical databases with and AI based optimizations that make it self-implementing, self-driving, self-healing, self-tuning, and self-securing. There is nothing else like it in the market today.

The Exadata Cloud@Customer X10M is additionally designed to simplify and accelerate implementations, operations, management, and troubleshooting, all of which are handled by OCI. It's purpose built for Oracle Database and Oracle Autonomous Database, which is based on Oracle Database.

Exadata Cloud@Customer X10M uniquely and automatically accelerates all Oracle Database workloads. It does so with exclusive Exadata hardware capabilities including:

- Offloading of SQL and JSON operations from the database servers to the storage servers
- Up to 1.25 TB of Exadata RDMA DDR5 very low latency Memory (XRMEM) per storage server
- Smart Flash Cache
- Storage Indexes
- Columnar Flash Cache
- Hybrid Columnar Compression (HCC) that delivers on average 10:1 data reduction
 - Versus a best case of approximately 2:1 for any other platform
 - Data does not have to be rehydrated to be queried
- I/O Resource Management
- Network Resource Management
- In-Memory Fault Tolerance
- Exafusion Direct-to-Wire Protocol

That's just some of the 60+ Exadata Cloud@Customer X10M unique capabilities designed to simplify and accelerate all Oracle databases.

Taking the 'limited' out of limited scalable performance while conquering database sprawl

All non-OCI or non-Cloud@Customer cloud database service, whether in the public cloud or on-premises, have severe scalable performance limitations. Take Amazon Web Services (AWS) as an example. Their Aurora cloud database services—MySQL or PostgreSQL—are limited to io1, which is 64,000 IOPS. Their Oracle cloud database service is limited to io2 or 256,000 IOPS. That's it. It is not much.

That means their customers have to 'shard' their databases when they need more performance. Each shard is a separate instance licensed service with its own infrastructure costs. When the customer requires 10-

100 shards, that's 10-100 licensed services, instances, and infrastructure, including storage. This database sprawl becomes exceedingly complicated to operate and manage the bigger it gets. The complexity and cost escalate when customers need to have that data queried by a Redshift data warehouse. It may require a data pipeline from each database shard plus a separate S3 bucket to transform the data for Redshift. That example is not unique to AWS, Microsoft Azure, and Google Cloud Platform (GCP) have similar limitations.

Oracle Exadata Cloud@Customer X10M changes the game entirely. It is purpose-built for extraordinary performance without necessitating the complexity or cost of sharding. That performance translates into supporting the highest OLTP databases transaction rates at the lowest read latency—less than 17 microseconds. It also enables huge cost savings with extensive database consolidation because many databases can share the same high -performance Exadata infrastructure.

The Exadata Cloud@Customer X10M starts at a quarter rack. The quarter rack comes with 2 database servers with 384 AMD 4th Gen EPYC cores, up to 6TB of DDR5 memory, and 3 storage servers with 192 AMD 4th Gen EPYC cores that together provide 3.75 TB of DDR5 based XRMEM and 240 TB of usable, triply redundant storage capacity.

Quarter rack max read IOPS based on 8K SQL I/O goes up to 5.6 million. The quarter rack max writes IOPS based on 8K SQL I/O goes up to 2.748 million. A full rack goes up to 25.2 million SQL 8K Read IOPS with a max of 6 racks. The math says that's well over 100 million 8K SQL read IOPS and well over 50 million write IOPS. It's way more and way easier to use than trying to run a high-performance dataset on something like an AWS Outposts that uses underpowered, 4- to 5-year old, vanilla servers and storage.

Elastic configurations range from a Quarter Rack, 2 database and 3 storage servers, up to a total of 16 servers in the initial rack and up to a total of 32 database and 64 storage servers across multiple racks.

Performance is massively higher and much more scalable than any other cloud database service platform in the public cloud or public cloud on the customer's premises. That translates into higher database consolidation and fewer cloud database licenses or BYOLs. Simply put customers can do more while spending less or run the same workloads on a smaller system.

What Exadata Cloud@Customer X10M Massive Performance Increases Means

It boils down to application response times. As application response times decrease all sorts of good things start to happen. Things such as reduced costs, increased revenues, and increased profits. The how is based on research. Most mission critical applications require databases, which are often the response time bottleneck. Exadata Cloud@Customer X10M removes that bottleneck with prejudice.

Productivity Gains

The extremely low Exadata Cloud@Customer X10M database read/write latencies and high scan throughput translate into much faster application response times, especially for mission-critical transaction processing and analytics applications. Faster response times directly translate into both cost savings and revenue. According to IBM's research on "[The Economic Value of Rapid Response Time](#)," as application response time decreased productivity increased. The ideal point is called the Doherty Threshold defined as:

"When the application and its users interact at a pace that ensures that neither has to wait on the other, productivity soars, the cost of the work done on the application's computer infrastructure tumbles, users get more satisfaction from their work, and their quality improves."

That threshold is often achieved at approximately 400 milliseconds (.4 seconds) application response times for the users.

Lower application response times translates into higher morale and lower turnover costs. The productivity gains are astonishing as seen in chart 1. Productivity more than doubles at the novice skill set, more than triples at the average skill set, and more than quadruples at the expert skill set. Those are substantial productivity gains. Chart 2 shows how to convert that into estimated monthly cost savings.




Chart 1: Transaction Productivity Based on Response Time




Chart 2: Cost Savings as Application Response Time Decrease

IBM’s research in Chart 2 clearly shows that application response time decreases increase monthly savings. The largest savings occur at the Doherty threshold. More users equate into more monthly savings.

Increased Revenues and Profits

Faster application response times have a major positive impact on revenues and profits. Higher productivity begets faster-time-to-project completions. That in turn begets faster-time-to-market. Faster-time-to-market begets faster-unique-revenues and profits. Those are revenues and profits that would not be realized with a later market entry, and they compound over time. The total amounts will likely exceed the cost savings while the total dwarfs all costs of Exadata Cloud@Customer X10M.

What About Amazon, Azure, and Google Cloud On-Premises Service Offerings?

Without fanfare, they don’t compare very well at all. Starting with Amazon.

Amazon On-Premises Cloud Database Services

Amazon offers a generic hardware platform for cloud database services on-premises called Outposts. Outposts has only a small subset of the AWS cloud Relational Database Service (RDS) offerings. None are Amazon’s premium cloud database services, such as Aurora and Redshift.

The generic hardware is in no way optimized for database services. In fact, it is not the same hardware Amazon uses in their own public cloud. The Outposts servers that AWS recommends for database operations are based on 5-year-old technology that was launched in their public cloud in [July 2018](#). The only RDS cloud database services available for Outposts are the relatively modest in performance and scalability Microsoft SQL Server, open-source databases MySQL, and PostgreSQL. That’s it. Amazon does not even support the Oracle Database Standard Edition 2 on Outposts.

Microsoft Azure On-Premises Cloud Database Services

Azure’s cloud database services on-premises are available in their Azure Stack platforms. Those platforms come in three flavors: Azure Stack Hub, Azure Stack HCI, and Azure Stack Edge. Their on-premises cloud database services are only slightly different from Amazon’s. They provide similarly modest performance and scalable cloud database services with Azure SQL, MySQL, and PostgreSQL.

Each of the Azure Stack offerings are not a turnkey system from Azure. The hardware must be purchased from Azure partners such as Dell, HPE, Lenovo, Supermicro, or other Azure partners. Dell through APEX and

HPE through GreenLake, provide cloud like subscriptions for their hardware. Once again, no matter who you get it from, the Azure Stack hardware is fairly generic, not what Azure uses in their public cloud, and not optimized or co-engineered in any way for any of the databases. One other sticking point is the need for one dedicated operator at a minimum to operate and manage the Azure Stack. That does not include the database administrator needed for whichever Azure Stack cloud database service customers elect to use.

Google Cloud Platform (GCP) On-Premises Cloud Database Services

GCP does not currently offer any on-premises cloud database services. Nor do they provide any hardware platforms directly to customers that are capable of running databases. GCP only provides a service called Anthos that must be deployed on hardware that customers must buy from Google partners. Anthos is a multi-cloud, hybrid cloud, and multi-site container orchestration software platform. It does not bring the GCP cloud on-premises. It orchestrates and moves around containers apps in GCP, on-premises, in other clouds, and on the edge. It's a limited offering and not comparable in any way shape or form to Exadata Cloud@Customer X10M.

Comparing Performance

Oracle publishes the performance that Exadata Cloud@Customer X10M delivers. Amazon and Azure do not publish any performance results for their on-premises cloud database services. It is easy to see why. Both are dependent on general-purpose, off-the-shelf server, network, and storage hardware. None of this hardware is co-engineered with or optimized for databases. They simply cannot come close to the Oracle Database performance or functionality running on the latest Exadata, even if they supported Oracle Database as-a-service on their on-premises cloud platform.

Then how can performance be compared? The best way is to compare Exadata Cloud@Customer X10M's performance versus Amazon and Azure cloud database services performance in their public cloud. Then compare it versus a couple of popular tier 1 all-flash-array (AFA) vendors on-premises. Granted, that would seem to give an advantage to the storage vendor because it does not add in the database query processing time. As will become clearly evident, this storage vendor advantage doesn't matter. Exadata Cloud@Customer X10M performance blows them all away.

Exadata Cloud@Customer X10M Performance vs Primary Public Cloud Competitors

Based on the fact that only Amazon and Azure provide on-premises cloud platforms that run cloud database services and publish their performance, they are Exadata Cloud@Customer X10M's primary cloud competitors¹.

Exadata Cloud@Customer X10M's performance for a single full rack is orders of magnitude better than its cloud competitors. It delivers:

- **50x** (5,000%) lower latency than either Amazon or Azure using their all-flash performance.
 - Exadata Cloud@Customer X10M's latency is **≤17 μs**.
 - AWS RDS **≤ 1000 μs**.
 - Azure SQL **≤ 1000 μs**.
- In one rack, **> 98x** (98,438%) more IOPS than Amazon can deliver to a few cloud database services².
 - It increases to **315x** or 315,000% more IOPS than most Amazon RDS cloud database services³.
 - **25.2 million** read IOPS vs. **256 thousand** read IOPS io2 or **80 thousand** io1.
 - Azure does not publish their Azure SQL IOPS.
 - However, to estimate their max 8K IOPS, divide their max 21 GBps throughput by 8K.
 - Estimated max IOPS comes out at 2.63 million 8K IOPS.
- For throughput of one rack.
 - **30x** (3,000%) greater throughput than Azure SQL.
 - **630 GBps** vs. **21 GBps**.

¹ GCP does not publish Cloud SQL (SQL Server, MySQL, and PostgreSQL) database performance. However, based on the GCP architecture, Cloud SQL performance would be very similar to Amazon RDS and Azure SQL.

² Based on Amazon io2 limited to 256,000 Flash Storage IOPS to a given instance.

³ Based on Amazon io1 limited to 64,000 Flash Storage IOPS to a given instance.

- **63x** (6,300%) more throughput than Amazon RDS.
 - **630 GBps vs. 10 GBps.**

It's important to remember those performance advantages are based on an Exadata Cloud@Customer X10M single full rack. And it scales to 6 full racks.

Exadata Cloud@Customer X10M's Performance vs Primary On-Premises Storage Competitors

Exadata Cloud@Customer X10M single system full rack performance is also considerably better than its All-flash Array (AFA) competitors. It delivers:

- For latency:
 - **8.8x** (880%) better latency than Pure Storage FlashArray//XL full system all-flash performance.
 - **17µs vs. 150µs.**
 - **3.3x** (330%) better latency than Dell PowerMax 8500 single rack all-flash performance.
 - **17µs vs. 60µs.**
- For read IOPS:
 - **33.16x** (3316%) more IOPS than Pure Storage FlashArray//XL full system all-flash performance.
 - **25.2 million vs. 760 thousand.**
 - **3.36x** (336%) more IOPS than Dell PowerMax 8500 single rack all-flash performance.
 - **25.2 million vs. 7.5 million.**
- For throughput:
 - **17.5** (1750%) more throughput than Pure Storage FlashArray//XL full system all-flash performance.
 - **630 GBps vs. 36 GBps.**
 - **3.6x** (36%) more throughput than Dell PowerMax 8500 single rack all-flash performance.
 - **630 GBps vs. 175 GBps.**

Keep in mind that the Dell PowerMax storage IOPS and throughput are based on 4K payloads. Whereas the Exadata Cloud@Customer X10M IOPS and throughput are based on typical 8K SQL payloads. That means the Oracle advantage is at least 2x larger than these numbers indicate. The Pure Storage FlashArray//XL is calculated from 32KB payload IOPS, which doesn't help databases with random IO requests to non-contiguous 8K blocks. Regardless, Exadata Cloud@Customer X10M's performance is massively better, and that does not include the Oracle Database or Autonomous Database faster processing. As previously noted, the Exadata X10M Cloud@Customer includes database processing. Dell and Pure Storage only measure storage system performance. The Exadata X10M Cloud@Customer literally blows away their performance.

CFOs Should Pay Attention to How Exadata Cloud@Customer X10M Improves Their Bottom Line

Exadata Cloud@Customer X10M's incredibly low latency advantage significantly reduces application response times. As the IBM research shows, reduced application response times equates into much higher productivity, higher morale, lower turnover costs, faster-time-to-market, faster-time-to-unique-revenues, as well as increased savings and profits.

Exadata Cloud@Customer X10M's massive IOPS advantage enables bigger transactional databases and considerable database consolidation. Consolidation can reduce database licensing costs and supporting hardware infrastructure costs.

Autonomous Database running on Exadata Cloud@Customer X10M reduces DBA workloads and the need for Oracle Database expertise, skills, knowledge, and experience. That saves considerable personnel costs over time.

Exadata Cloud@Customer X10M's huge throughput advantage accelerates analytics, AI machine learning, and data ingestion. That in turn accelerates time-to-actionable-insights, time-to-market, and time-to-unique-revenues and profits.

Depending on each customer's circumstances, Exadata Cloud@Customer X10M's cost can be substantially lower than the gains to the bottom line.

It doesn't matter what the other cloud, server, or storage vendors offer. They cannot come close to the performance of Exadata Cloud@Customer X10M. It is that performance that makes Exadata Cloud@Customer X10M the much better value why CFOs need to pay attention.

Conclusion

Wikibon concludes that the Oracle Exadata Cloud@Customer X10M delivers a substantial performance leap over the previous X9M generation and beyond the reach of competitive cloud on-premises offerings. That leap takes advantage of the latest core and performance gains from the AMD 4th Gen EPYC processors in both the database servers and for the first time in the storage servers. But it does not stop there. Oracle adds much faster XREM than Intel Optane in the storage servers and extensive co-engineered Exadata software innovations that leverage those AMD EPYC CPUs and XREM.

The performance results demonstrate Exadata Cloud@Customer X10M has no peer in the on-premises cloud database services. And it's not close at all.

Exadata Cloud@Customer X10M's performance changes the game at a fundamental level. It enables considerably:

- Reduced application response times.
- Higher employee productivity, quality of work, morale, and retention.
- Lower employee turnover, training, and cost.
- Lower database and infrastructure costs via consolidation.
- Faster-time-to-market.
- Faster-time-to-actionable-insights
- Faster-time-to-unique-revenues and profits.

When it comes to on-premises cloud database services, or even do-it-yourself (DiY) database implementations, there is just no comparison to Oracle Exadata Cloud@Customer X10M. They all come up woefully short.

Wikibon recommends Oracle Exadata Cloud@Customer X10M for medium, large, and extra-large database users looking to take advantage of cloud economics with data sovereignty and much greater performance. The decision should be simple and highly rewarding.

More Information

[Oracle Exadata Cloud@Customer X10M](#)
[Oracle Autonomous Database](#)