



RESEARCH NOTE

Oracle Exadata Cloud Service takes elasticity to a new level

Embracing Persistent Memory and AI scale

Executive Summary

Trigger

Oracle's Exadata Cloud Service is getting a major performance boost. The X8M generation of Exadata, first unveiled for on-premises implementations a year ago, is now live on Oracle Cloud Infrastructure (OCI) as well as in Exadata Cloud@Customer and Dedicated Region Cloud@Customer. To recap, X8M is a major technology shift for Exadata. It introduced persistent memory (PMEM) and made a major architectural change in the internal networking fabric, moving from InfiniBand to Remote Direct Memory Access (RDMA) over 100 Gbs Converged Ethernet. The results are the doubling of transaction processing performance and reduction of latency by an order of magnitude. And with new access through Oracle public and hybrid cloud services, customers have more options when it comes to sizing and elasticity, with maximum capacity for analytics rising to data lake country at 25 PBytes, after compression.

Our Take

There is little surprise that Exadata X8M is now available as part of Oracle Cloud Infrastructure. Ever since X8M was announced for on-premises deployments, the question for cloud wasn't whether but when. Nonetheless, there is a major – but logical – new component of X8M that will be unique to cloud: Oracle is changing how Exadata capacity in the cloud is made available. It is transitioning from standard rack sizes to variable, fully elastic combinations of compute and storage nodes – and in the cloud, it is raising the ceiling to data lake territory. As would be expected from a managed cloud Database-as-a-Service (DBaaS), software patching and updating are automatically handled by Oracle; but a key differentiator is that patching takes place while the service is still running. While we don't expect most customers to use Exadata as a data lake, the boost in effective capacity will be especially valuable for running AI/machine learning models that are especially ravenous for data.

A generational change comes to Oracle Exadata Cloud

Exadata has been one of the bright spots of Oracle's business, with Exadata Database Machine sales jumping 15% in the most recent quarter. In the cloud, a significant percentage of Oracle Exadata Cloud customers are net new customers to the Exadata platform; they are customers who want Oracle, not a third party, to run and manage their service. On-premises, the latest X8M generation, announced a year ago, have hit the inflection point, with most sales in the past quarter being the new RoCE-enhanced model. This is the generation of system that is now available in Oracle Exadata Cloud Service.

X8M is not just an incremental upgrade. With Oracle Exadata X8M, Oracle turbocharged transaction processing while also migrating the platform to the mainstream of industry standards for networking fabric and hypervisor.

With the jump to cloud, Oracle is folding in a new twist: The Oracle Exadata Cloud Service based on X8M takes advantage of the cloud's near-infinite resources to better deliver scalability and elasticity than its predecessors did. Before this, Exadata – both on-premises and in the cloud, was in effect packaged as standard T-Shirt sizes. You had extra small as a base system, followed by small (1/4 rack), medium (1/2 rack), and large (full rack). Traditionally, the sweet spot has been the quarter rack, which is configured with 2 compute nodes and 3 storage nodes.

Going forward, those arbitrary T-shirt sizes in the cloud fall by the wayside. Instead, you can start at two compute nodes and three storage nodes (which equates to the traditional 1/4 rack), but from there you can increase incrementally up to 32 compute (1600 OCPUs) and 64 storage nodes (3,072 cores). In terms of overall capacity, the Exadata X8M cloud service new upper limits are huge. It translates to a maximum capacity of 25 PBytes after compression, which is data lake territory. We have not come across any other cloud database or data warehousing services approaching that upper limit.

Oracle Exadata Cloud Service X8M behaves as a cloud-native service, and you can consume it with the expected options that native cloud services should offer: On-demand up-and-down scaling of OCPU consumption allows you to meet immediate database needs without interrupting database operations, enabling pay-per-use consumption that minimizes costs. X8M's high performance further reduces costs since faster processing equals less billable time in the cloud; you subscribe to only the compute and storage that you need. You can scale up cores within a server, or entire servers readily. At launch, you will be also able to scale active cores inside an Exadata X8M cloud server up or down; we expect that scaling down of physical servers in the cloud service will come later.

Given that Exadata's original design (as an engineered system) already separated compute from storage, it was born ready from Day 1 to operate in cloud-native mode. So, Oracle didn't have to redesign Exadata for cloud-native operation. Instead, two things were upgraded:

- Making the service fully elastic with the ability to add storage servers for data warehouses, and database compute servers for OLTP or both for massive enterprise workload consolidation.
- The upper limit of capacity growing to data lake size is a logical follow-on, given that it is more practical to deliver such capacity in a public cloud environment where enterprises do not have to arrange floor space and boost capital budgets.

The future is RoCE

Until now, one last piece of “premium” architecture remained on Oracle Exadata: InfiniBand, which was originally much faster and more reliable than Ethernet. These are differences where the gap has now been more than made up. While InfiniBand is a standard, Ethernet is the more widely adopted standard, and it’s where the brunt of technology innovation and availability of networking equipment is. Today, 100 Gbs Ethernet links are commonplace in cloud data centers, easily matching or outpacing InfiniBand, while packet loss issues have been, for all practical purposes, eliminated.

That set the stage for X8M to embrace 100 Gbs Ethernet, with InfiniBand protocol compatibility. And in turn, that enabled expansion of the use of Remote Direct Memory Access (RDMA), which as the name implies, provides direct access to memory bypassing the overhead of the operating system and I/O stack. The network implementation is RDMA over Converged Ethernet (RoCE – pronounced “Rocky”). Combined, the impact of the changes is that Oracle can tap into the mainstream of network technology development, providing access to a much broader selection of network card and switches. For Exadata, that will allow faster innovation and a more cost-competitive platform. These benefits are in turn passed on to customers, as they obtain faster performance, increased scalability, and elasticity at the same price.

Practical impact on OLTP and Analytics

Given that Oracle designed Exadata as a database consolidation platform that handles mixed workloads, the new X8M cloud service has sweeteners for both ends of the spectrum: PMEM and RoCE for turbocharged transaction processing and higher ceilings on scalability for analytics. While we don’t expect most Exadata customers to use it as a data lake, the boost in capacity will make it very useful for running AI/machine learning models that are especially ravenous for data.

What’s notable is that, while Intel unveiled what eventually became its Optane line of persistent memory 5+ years ago, it has taken awhile, not only for Intel to deliver product to market, but also for data platform providers to effectively incorporate it in their system designs. That is why Oracle, along with SAP, have largely been alone so far in capitalizing on PMEM. The trick is in the implementation. PMEM can be implemented in a memory slot, which maximizes performance but makes storage volatile, or an SSD Flash slot, which is non-volatile but delivers slower performance. And then there is the mode or protocol that is used; the best performance for application software comes with AppDirect mode, but that requires changes either to the application itself or to code that abstracts it. That’s where Oracle has done the legwork.

Additionally, Oracle has implemented Remote Direct Memory Access (RDMA), which allows access to PMEM without going through the overhead of the operating system and I/O stack. The result is more than double (2.5x) the IOPS performance and 10x faster latency to a low of 19 microseconds. The key beneficiary is transaction processing. Thanks to improved performance, it can enable Exadata customers to more economically size transaction compute workloads.

At the other end of the scale is, literally, scale. In the cloud, Oracle has drastically raised the upper limit for Exadata to multi-PByte range, which will be useful for analytics. We don't expect many customers to hit the 25-PByte limit, because at that scale, we are looking at data lakes that are probably going to be utilizing cloud object storage. But it increases the flexibility for Exadata Cloud Service to handle data-intensive analytics, which will be pivotal for running data-hungry machine learning models. For all practical purposes, Exadata cloud customers won't have to worry about hitting the wall in capacity.

Takeaways

A higher ceiling in the cloud

As noted above, Exadata X8M was already released on premises, so the extension to the Oracle Public Cloud wasn't a huge surprise. We already knew the architecture of the latest edition of Exadata. Nonetheless, changing the elasticity and shape sizing for the new cloud edition is the major new development – and a very logical one at that. It takes what was originally an “engineered system” that already separated storage from compute and provided the flexibility to best take advantage of the elasticity of the cloud. Having a high ceiling for data will prove especially useful when running machine learning models, which are ravenous for data.

Differentiating on commodity technology

Over the years, the equivalent of Moore's Law for all elements of IT infrastructure kicked in – for compute, multicore emerged as the practical limits of chip and microcircuit density were reached, and for storage, the continual price/performance trends are making Flash and DRAM memory more affordable. Over the years, Oracle has used faster and faster x86 server architectures and continually increased the quantity of storage (at all tiers) while keeping pricing level. With X8M, Exadata's journey to commodity architecture is complete – the embrace of 100 Gbs Ethernet and RoCE. A related change is the shift from Xen to KVM as the hypervisor.

With commodity technology as the floor, differentiation comes by how that technology is packaged and integrated. In Exadata X8M, the prime example is how Oracle optimized use of

non-commodity technology (PMEM) with commodity technology (RoCE). Oracle is one of the few data platform providers that has mastered PMEM to date. It made the design choices on how to package persistent memory and implement it in software. RoCE provides the last piece of the puzzle – delivering high-speed access to high-speed storage without the overhead of the operating system and I/O stack. Many of the underlying technology building blocks might be commodity, but the way they are implemented in X8M is not, and that includes specialized algorithms co-engineered with Oracle Database that optimize performance and automate management in the cloud.

With Exadata posting strong sales and the same architecture now available across all three deployment models, the company has a significant opportunity to grow market share in the hotly contested cloud database market in 2021.

Author

Tony Baer, Principal, dbInsight

tony@dbinsight.io

Twitter @TonyBaer

About dbInsight

dbInsight LLC® provides an independent view on the database and analytics technology ecosystem. dbInsight publishes independent research, and from our research, distills insights to help data and analytics technology providers understand their competitive positioning and sharpen their message.

Tony Baer, the founder and principal of dbInsight, is a recognized industry expert on data-driven transformation. *Analytics Insight* named him one of the [2019 Top 100 Artificial Intelligence and Big Data Influencers](#). His combined expertise in both legacy database technologies and emerging cloud and analytics technologies shapes how technology providers go to market in an industry undergoing significant transformation. His regular ZDnet “*Big on Data*” posts are read 25,000 – 30,000 times monthly.

Author

Tony Baer, Principal, dbInsight

tony@dbinsight.io

Twitter @TonyBaer



Oracle Exadata Cloud Service takes elasticity to a new level

About dbInsight

dbInsight LLC® provides an independent view on the database and analytics technology ecosystem. dbInsight publishes independent research, and from our research, distills insights to help data and analytics technology providers understand their competitive positioning and sharpen their message.

Tony Baer, the founder and principal of dbInsight, is a recognized industry expert on data-driven transformation. *Analytics* named him as one of its Top 100 influencers for [data](#) and [cloud](#) in 2019 and 2020. *Analytics Insight* named him one of the [2019 Top 100 Artificial Intelligence and Big Data Influencers](#). His combined expertise in both legacy database technologies and emerging cloud and analytics technologies shapes how technology providers go to market in an industry undergoing significant transformation. His regular ZDnet “*Big on Data*” posts are read 25,000 – 30,000 times monthly.

dbInsight® is a registered trademark of dbInsight LLC.