

# Oracle turbocharges Exadata

---

# Ovum view

## Summary

Oracle's second refresh of the Exadata product line this year promises to more than double transaction processing performance. The new model, numbered Oracle Exadata X8M, gets the performance boost by introducing a new Persistent Memory (PMEM) tier based on Intel's new Optane line, and a shift in networking from InfiniBand to a converged 100 Gigabit (GbE) Ethernet backbone with InfiniBand emulation. Announced at Oracle OpenWorld, the new Exadata X8M model will be priced at parity with the existing X8 line and should allow existing customers to run their applications and databases unchanged. As the use of persistent memory improves IOPS performance (input/output per second), the new model will be most attractive to customers using Exadata primarily for online transaction processing (OLTP).

## Making the Optane jump

Oracle has now joined SAP as being among the first major household names in data platforms to place their bets on next-generation persistent memory storage (SAP announced that Optane was part of its next-generation HANA plans several months ago). Oracle's and SAP's moves shouldn't be surprising given the hype that surrounded persistent memory storage when public discussion emerged about five years ago. When Intel announced what eventually was branded Optane (it was code-named 3DXpoint back then), it was hailed as a new miracle storage tier that could deliver almost the performance of memory at almost the price point of Flash. So of course, the question was, when would the other shoe drop – that is, when would the household brand names start to actually incorporate this technology into their products?

It's taken some time for data platform providers to get up to speed on persistent memory. The lesson learned, as it turned out, was on *how* Optane is implemented. You can't just simply drop Optane into a storage slot and automatically get game-changing performance. First, you need the right form factor: Optane can be slotted into solid state disk or memory form factors, but the high performance only comes if the memory path is chosen. Secondly, while you can run software without change in Optane, performance improves when you implement in AppDirect mode, which requires modifications to software.

Oracle has internalized these lessons in the new X8M release, and for Exadata, has added yet another trick. It runs in RDMA (Remote Direct Memory Access) mode, an industry standard for direct access to memory without going through the overhead of the operating system. This isn't Oracle's first experience with RDMA in Exadata, but until now, it was primarily used for system-level functions such as internode server coordination or to speed up large data transfers. In the X8M model, Oracle has become the first to implement RDMA for direct access to persistent memory. And that's where the big OLTP performance gains come. Oracle has benchmarked the new X8M model at 16 million IOPS, which is roughly 2.5x that of the 6 million IOPS for the existing Exadata X8 generation. Latency is roughly 10x lower, with Oracle clocking the system at less than 19 microseconds.

## Moving toward more mainstream architecture

A key design pillar of Exadata has always been optimized hardware. A decade ago, when the product line emerged, InfiniBand was the standard of choice for high performance, offering many times the throughput of Ethernet. Also, the nature of Ethernet's design carried the risk of packet loss. Since then

the industry has moved on, as 100-gigabit (GbE) Ethernet has become common and new converged configurations have introduced compatibility with other protocols such as InfiniBand. The result is that today's Ethernet has achieved parity with InfiniBand, and carries the advantage of being a much more widely implemented networking standard.

As noted earlier, a key to the new Exadata model's performance is its use of RDMA; the industry supports that with an emerging RDMA-over-Converged-Ethernet (RoCE) standard that has now been widely adopted by most major network card and switch makers. In the new X8M model, Oracle implements RoCE, providing the best of both worlds: backward compatibility with Exadata InfiniBand implementations on an Ethernet-based stack that has become the de facto industry standard. And with RoCE, the threat of packet loss is largely eliminated.

Another standards-related choice is the shift toward using new VMs. Until now, Oracle has relied on the Xen hypervisor; going forward, it will be utilizing the more widely adopted KVM. As a more popular standard, there has been far more development and support; for instance, RoCE and Persistent Memory are only supported on KVM. It also supports double the "guest" memory footprint (1.5 TBytes) per server.

## This release targets OLTP

Clearly, the major differentiator for this release is on transaction processing, thanks to the implementation of RoCE and persistent memory. Admittedly, with the new KVM hypervisors supporting larger memory footprints, there may be some spillover benefits for analytics. But, aside from the addition of persistent memory, a new networking backplane (that is backwards compatible with all Exadata software enhancements), and the shift to a new hypervisor, the X8M model of Exadata still carries the same automatic tiered storage features that have made it a mixed workload machine. The racks are configured with up to 3 PBytes of disk, 920 TBytes of NVMe Flash, and 27 TBytes of persistent memory. And for now, Oracle is offering the X8M at price parity with the current X8 generation line.

The bottom line is, yes, the big difference in the X8M is with OLTP, but it will continue to handle the mixed workloads for which Exadata is known.

## Appendix

### Further reading

"Oracle Exadata X8 extends analytic query footprint into Hadoop and Spark territory," (June 2019)

"Oracle extends Autonomous Database to transaction processing," INT002-000155 (August 2018)

"Oracle bakes security into its DNA," INT003-000287 (November 2018)

"Oracle's second-generation cloud is designed to be enterprise grade," INT003-000329 (February 2019)

### Author

Tony Baer, Principal, dbInsight LLC

[tony@dbinsight.io](mailto:tony@dbinsight.io)

## Ovum Consulting

We hope that this analysis will help you make informed and imaginative business decisions. If you have further requirements, Ovum's consulting team may be able to help you. For more information about Ovum's consulting capabilities, please contact us directly at [consulting@ovum.com](mailto:consulting@ovum.com).

## Copyright notice and disclaimer

The contents of this product are protected by international copyright laws, database rights and other intellectual property rights. The owner of these rights is Informa Telecoms and Media Limited, our affiliates or other third party licensors. All product and company names and logos contained within or appearing on this product are the trademarks, service marks or trading names of their respective owners, including Informa Telecoms and Media Limited. This product may not be copied, reproduced, distributed or transmitted in any form or by any means without the prior permission of Informa Telecoms and Media Limited.

Whilst reasonable efforts have been made to ensure that the information and content of this product was correct as at the date of first publication, neither Informa Telecoms and Media Limited nor any person engaged or employed by Informa Telecoms and Media Limited accepts any liability for any errors, omissions or other inaccuracies. Readers should independently verify any facts and figures as no liability can be accepted in this regard - readers assume full responsibility and risk accordingly for their use of such information and content.

Any views and/or opinions expressed in this product by individual authors or contributors are their personal views and/or opinions and do not necessarily reflect the views and/or opinions of Informa Telecoms and Media Limited.

## CONTACT US

[www.ovum.com](http://www.ovum.com)

[askananalyst@ovum.com](mailto:askananalyst@ovum.com)

## INTERNATIONAL OFFICES

Beijing

Dubai

Hong Kong

Hyderabad

Johannesburg

London

Melbourne

New York

San Francisco

Sao Paulo

Tokyo

