Ovum

Oracle Exadata X8 extends analytic query footprint into Hadoop and Spark territory

Publication Date: 12 Jun 2019 | Product code: INT002-000226

Tony Baer

Ovum view

Summary

With the release of Intel's Xeon E8 generation, Oracle is refreshing the Exadata line with X8 models that significantly increase storage density, and with security-related enhancements to the processor, boost database processing performance. But the most significant introduction in the new X8 line is expansion of Exadata's footprint. Thanks to the declining cost of magnetic storage, Oracle is adding a low-cost storage tier for historical data that would otherwise be archived offline. For Oracle customers, this could make it more economical to keep some of their longer-term historical query loads on Exadata rather than move them to Hadoop and Spark.

Adding a new historical data tier

The overlapping and convergence of data platforms is a trend that Ovum first identified nearly five years ago, with the latest Exadata release proving that this trend is still very much alive and well. Oracle is introducing a new third option to the Exadata storage server line, the Exadata X8-2 Extended (XT), which could claim some historical analytics workloads that would otherwise go to Hadoop. Alongside the all-flash and mixed flash/disk models, the new XT model provides lower-cost storage for cool (infrequently accessed) data, with 168 TBytes per server. For the new XT model, customers only pay for storage. The storage server software, which provides features such as smartscan for speeding up queries, is optional.

While this would not necessarily replace data lakes, the Exadata storage server will be useful for providing low-cost access to historical data. Several years back, we attended an Oracle OpenWorld session where the Exadata customer described their big data strategy: they kept two to three years of data on Exadata and used Hadoop for querying data up to years seven to ten. This is the use case addressed by the XT model. By eliminating the cost of licensing software and only using commodity hard disk storage (no flash), Oracle is seeking to level the playing field with Hadoop, where the customer only pays for storage and maintenance/support.

The Oracle database platform and cloud-based Autonomous Data Warehouse (ADW) already has some big data-related capabilities. For instance, the Oracle database supports extended data types, such as JSON and CSV, while ADW can automatically schematize Parquet files. However, these capabilities are not necessarily meant to substitute for Hadoop implementations or data lakes built on cloud object storage that are suited for programmatic queries. Instead, this addition to the Exadata storage line provides a sweet spot for Exadata to run queries on much larger troves of historical transaction data that can continue to be accessed more efficiently as database tables rather than file systems or object stores. Because Exadata is a pillar of Oracle's Autonomous Database cloud services, we would like to see Oracle consider adding this option to ADW in the future.

Beating all-flash servers at their own game

The new X8 Exadata storage servers have taken a balance-of-system approach to boosting performance and adding density. The scale-out two-socket storage servers have added 60% more cores, which offload database processing, while boosting density to 14-TByte disks. This is a much higher jump than the previous generation, which by comparison added 20% more capacity. For the

© Ovum. All rights reserved. Unauthorized reproduction prohibited.

all-flash model, Oracle claims 60% faster performance than the previous X7 Exadata generation. More significant, however, are the performance improvements to the hybrid disk/flash model, which Oracle claims can match or beat all-flash storage arrays. Admittedly, at first glance, the claim that a hybrid array could outperform an all-flash array would appear to violate some law of physics. The real differential, however, is that the Exadata software optimizes workload balancing through pushdown to a smart storage tier. Combined with balance-of-system hardware performance improvements, the new Exadata model more than holds its own against advanced all-flash network-attached storage that requires data to be pushed to compute nodes, which can prove to be a bottleneck.

Another enhancement in Exadata X8 comes courtesy of Intel. When Spectre and Meltdown malware hit last year, the onus for taking corrective action fell on software providers such as Oracle. In this case, Oracle implemented the patches at the OS and database tiers, which exacted hits on performance in the 5% to 20% range, depending on workload type. In the new Intel Xeon E8 processors, these fixes have been implemented in silicon, where they belong (it's a lot more efficient there). That's not all. The new E8 processors also sport 15% clockspeed increases, which Oracle fully leverages in the new Exadata refresh.

Doubling down on intrusion detection, bringing incorporating ML into Exadata

With database security a growing concern, intrusion detection highlights the 20 or so enhancements that are also coming in the new Exadata models. Among the highlights are enhanced intrusion detection capabilities and new safeguards for data erasures, as well as new controls over IP addresses or subnet masks that can access Exadata through RESTful interfaces.

Additionally, the latest Exadata release incorporates machine learning (ML) capabilities to assist several performance-management and monitoring related features. These include applying ML to resource management that start by identifying normal consumption patterns and outliers. When Exadata detects unusual resource drains, ML is used in performance monitoring to help identify the source. This is applied as part of the automated performance monitoring capabilities that are new to this Exadata release. In addition, Exadata is applying ML to implement automatic indexing with Oracle Database 19c. Based on technology from Oracle Autonomous Database, the entire process is automatic and improves database performance and eliminates manual index-tuning tasks for critical applications.

In Ovum's *2019 Trends to Watch: Big Data* (see Appendix) report, we predicted that Oracle's foray for incorporating machine learning into the running of databases would set off a trend. Oracle's latest ML-related enhancements to Exadata are proof that such capabilities are entering the mainstream.

Appendix

Further reading

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

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

© Ovum. All rights reserved. Unauthorized reproduction prohibited.

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

2019 Trends to Watch: Big Data, INT002-000201 (December 2018)

Author

Tony Baer, Principal Analyst, Information Management

tony.baer@ovum.com

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.

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

ovum.informa.com
askananalyst@ovum.com

INTERNATIONAL OFFICES

Beijing
Boston
Chicago
Dubai
Hong Kong
Hyderabad
Johannesburg
London
Melbourne
New York
Paris
San Francisco
Sao Paulo
Shanghai
Singapore
Sydney
Tokyo