

ANALYST INSIGHT

Oracle Exalytics In-Memory Machine: Engineered for Speed-of-Thought Analytics

Reference Code: OI00141-059

Publication Date: November 2011

Author: Madan Sheina

SUMMARY

Ovum view

One of the biggest product announcements at 2011's Oracle OpenWorld user conference was Oracle Exalytics In-Memory Machine, the latest addition to the "Exa"-branded suite of Oracle-Sun engineered software-hardware systems. The product (formerly codenamed BI Machine) was reportedly branded personally by Oracle CEO Larry Ellison, which makes it a high-profile product announcement. Oracle Exalytics reflects several technical trends in analytics today – in-memory, parallel database processing, and self-serve business intelligence (BI). These are converging to enable what Oracle calls "speed-of-thought" analytics, which Oracle markets and sells as a standalone appliance. However, Ovum believes Oracle Exalytics works best if paired with the company's Exadata Database Machine and its new Oracle Big Data Appliance. Oracle has engineered another integrated hardware-software system, though the base components are hardly new technologies, and some have had to be modified with parallel and in-memory capabilities in mind. Customers will have to wait until early 2012 to get their hands on the new In-Memory Machine. However, Ovum would have liked to see more business context and use cases to justify an Exalytics investment. How many organizations really want, or are positions to take advantage of, analytics that purport to work faster than employees' minds think or business processes can react?

ANALYSIS

Lightning fast, response driven, and in memory are hallmark qualities of Oracle Exalytics

The primary design principle of Oracle Exalytics is to enable fast and easy ad hoc analysis across large end-user communities using an in-memory processing engine. Speed of thought and instant response are the hallmarks of its functionality, making it highly applicable to a range of ad hoc, what-if analysis and forecasting and realtime planning applications.

The in-memory capabilities are key to enabling what Oracle calls a highly interactive and visual analytic experience for end users. Oracle claims that Exalytics introduces a new user interface designed to handle end-user queries "at the speed of thought." Peeling away the marketing speak, this means the product is designed to trawl through dense data sets (regardless of query, location, and device type and including desktops, laptops, tablets, and even smartphones) and provide results almost instantaneously. Oracle claims that the algorithmic speed of the system means it can respond to queries as they are being typed – which is, in many ways, similar to how Google offers suggested searches based on a partially typed phrase.

The promise of delivering data analytics faster than users can type will likely have many business use cases. However, Ovum believes Oracle, like many other analytics vendors pushing in-memory capabilities, has not articulated the key use cases for Oracle Exalytics.

Oracle Exalytics is impressively engineered, though based on existing and modified Oracle products

First, the basics: what exactly is Oracle Exalytics? It is the latest addition to Oracle's "Exa"-branded family of pre-engineered software-hardware systems – in which the company integrates its software on its own hardware. Now that Oracle has its own line of hardware (through the acquisition of Sun), it appears to have fully embraced the "engineered systems" model as a way to sell both hardware and software while maintaining higher margins. Oracle Exalytics' immediate "engineered" siblings currently include the Exalogic Elastic Cloud middleware machine, the Exadata Database Machine, and the Oracle SPARC SuperCluster T4-4 general-purpose servers. Ovum predicts more additions to the Exa family in the next year. These are marketed and sold as complete appliances; Oracle does not allow customers to install and configure the granular components themselves. That said, customers can purchase Oracle's BI software components independently.



Although Oracle Exalytics is a new product, the architecture is built on several existing Oracle products: parallelized versions of its TimesTen in-memory database and Oracle Essbase OLAP Server (a specialized in-memory version), together with an optimized version of the Oracle BI Foundation Suite (OBI EE 11g for standard enterprise-grade BI query, analysis, reporting, dashboarding, and other visualizations). Most, if not all, of these software products have been modified to run in-parallel and in-memory data-processing architectures. All of these combine to deliver query optimization, complex multidimensional analysis and planning calculations, and enterprise-wide BI scale, respectively, through a revamped user interface that is designed for "speed-of-thought" analytics.

The secret sauce is to be found in the innards of the appliance – built on an impressive Sun Fire server (a four-socket, 40-core Intel Xeon E7-4800 processor implemented as a rack-mountable box), which features 1TB of dynamic random access memory (DRAM) and two InfiniBand 40Gbps connection ports to other Oracle products, notably Oracle Exadata. (The system can also connect natively to the Oracle Exalogic middleware machine – 10GbE – all of the Oracle-Sun and StorageTek disk and tape storage arrays, and EMC-affiliated storage).

Oracle Exalytics configuration emphasizes in-memory analytics through an adaptive DRAM cache. DRAM is 10× more expensive per gigabyte than Flash, but it has better speed. A single terabyte of DRAM might not sound that big. However, Oracle claims that after TimesTen columnar compression rates of a 5X and 10X factor, the cache can be expanded up to 5 to 10TB of user-accessible data with scanning times of 20GB per second. That allows users to explore its full DRAM capacity in about five seconds.

Oracle has not yet disclosed exact pricing for Oracle Exalytics or specified a shipment date, but Ovum expects early 2012. Ovum believes Oracle will go with a named-user scheme until deployment numbers make it cost effective to implement processor-based pricing. Oracle Exalytics is delivered as a pre-engineered box; the system is not sold without the hardware, and the in-memory software is not available separately.

This is where things get slightly complicated. Oracle, in fact, sells the system as three SKUED products – hardware, software, and an in-memory analytics component. However, customers cannot buy a single SKU and build out the rest of their systems using (third-party) best-of-breed components. They need to buy all three together. Oracle also offers installation services, though the nature of the product means this is a minimal effort.



Oracle Exalytics is Oracle's first (predictable) in-memory analytics play

This is certainly Oracle's first big response to in-memory plays of rival BI vendors such as SAP and IBM, as well as specialized plays such as QlikTech and Tableau Software. In particular, in-memory analytics spotlights the need to support "Agile Analytics." Oracle claims that Oracle Exalytic's 1TB cache of on-board memory means there are virtually no round trips to disk when processing analytic tasks.

There are two in-memory components that enable this in Oracle Exalytics: a new adaptive, self-tuning, in-memory cache built from scratch by Oracle that optimizes which data and processing are shifted to memory, and an optimized in-memory database (courtesy of its acquired TimesTen technology). The cache decides which information should be stored in memory based on the ongoing workload, particularly the repetitive nature of the queries.

Admittedly, TimesTen's main language is SQL and it was designed for OLTP processing. However, since it came under Oracle's wing, it has been enhanced with columnar compression algorithms and analytic functions such as "rank" and "sub query" to make it more suitable for multidimensional analysis. There is also a new "Summary Advisor" component built into the in-memory engine that directs "hot" data to memory based on usage-tracking metrics. Clearly Oracle has designed Oracle Exalytics as a flexible and highly adaptive in-memory cache that continuously looks at optimal ways of updating and distributing BI query loads.

Oracle is not exactly pitching Oracle Exalytics as a rival to QlikView and Tableau, which come with much more lightweight footprints, but there are certain similarities. However, Ovum believes comparing Oracle Exalytics to these two systems is akin to Oracle giving users a hammer to swat a fly. While the QlikView engine is designed for fast processing and response for basic BI analysis and reporting against modest data cubes, Oracle Exalytics seems to be targeted at more extreme (larger) data sets and workload environments. Big Data sources such as weblogs, social media feeds, smart meters, sensors, and other devices that generate massive volumes of data are fair game for Oracle Exalytics. Also, Oracle Exalytics' ability to support planning and what-if analysis in memory could be a bonus. Comparing the products may therefore seem unjustified, though Oracle's pricing for its system will play a big part in its competitive positioning and success in the in-memory analytics solution market.

Oracle Exalytics works best when paired with Exadata for big and fast analytics, but at what cost?

No Oracle-engineered system is built as an island. As such, Oracle Exalytics is designed to work comfortably in concert with Oracle Exadata to enable a scalable and fast data warehousing and



analytics solution. Users benefit from using both systems, and Oracle extols the virtue of pairing Oracle Exalytics with Oracle Exadata. Ovum believes this is both a strength and a weakness of Oracle's strategy.

The Exadata-Exalogic combination can feasibly pull a large, complex analytic query from one hour down to minutes through optimized storage and in-memory processing. However, customers have to license both systems (as separate SKUs) to benefit from the full performance and scale potential. Most of the beta customers that Oracle spotlighted at OpenWorld used both. Also at OpenWorld, Oracle emphasized that Oracle Exalytics performed much better when used with an Oracle Exadata machine. This, of course, implies that customers would benefit from implementing the two alongside one another. Add to that equation the newly announced Oracle Big Data Appliance (which includes an open source distribution of Apache Hadoop, Oracle NoSQL Database, Oracle Data Integrator Application Adapter for Hadoop, Oracle Loader for Hadoop, and an open source distribution of R - for Big Data analysis), and the trio of Oracle Exa-boxes aimed at data warehousing and analytics could provide some confusing overlaps. Ovum believes Oracle needs to do a better job of explaining the interplay between Oracle Exalytics, Oracle Exadata, and Oracle Big Data Appliance – in particular highlighting analytic scenarios in which companies would use combinations of these appliances.

Although the combined scale and performance of all three systems is certainly appealing, Oracle has designed its system to separate analytic workloads across multiple appliances. The question for companies looking to rationalize their infrastructure is: "why not collapse the in-memory capabilities inside Oracle Exadata?" This also means there is a lot of potential for caching (and cache management) between Oracle Exadata and Oracle Exalytics. For example, only if the data resides in Oracle Exalytics' DRAM memory will it provide ultimate query performance. If not, Oracle's InfiniBand connectivity to Oracle Exadata will be used to move data into Oracle Exalytics' in-memory cache. Admittedly, caching has proved to be an effective way of speeding up predictable queries that perform complex calculations and aggregations against semi-predictable data sets. However, it can be a tricky proposition to manage and synchronize highly volatile caches to ensure integrity and currency.

A bigger implication for customers relying on multiple appliances from the same vendor for high-performance analytics is, of course, vendor lock-in. Customers might be left with no option to use their own preferred software components, even if these are from third-party vendors. Similarly, as hardware advances, organizations will be restricted on their existing platforms of choice and most will probably not be willing to accommodate new components to their appliances.

Oracle Exalytics gives Oracle a competitive stake in in-memory analytics

Oracle is entering a fiercely competitive analytic appliance market as a way to sell both hardware and software while maintaining higher margins. Although its Oracle Exadata database machine is aimed at data warehousing appliances such as Teradata, Oracle Exalytics will compete more directly with rival in-memory-based analytic appliances such as SAP HANA, which has been on the market for more than a year. It will also be up against IBM, which has had its own suite of data warehousing appliances and bought Netezza in September 2010. IBM also has a similar push towards integrated software-hardware appliances optimized for specific workloads – part of its "Smarter Systems" initiative. One big difference between Oracle Exalytics and SAP HANA lies in their respective hardware strategies; Oracle Exalytics customers have to use Sun, while SAP HANA customers have more options.

Oracle is a latecomer in terms of its in-memory offerings, which perhaps explains the lower-key tone of the Oracle Exalytics launch at the OpenWorld conference in October 2011. Oracle has experience with in-memory technology with its TimesTen technology, but this is now a crowded market. Ovum believes the company must sharpen its marketing, particularly around use cases, and perhaps more importantly, get its pricing strategy right, if it is to have any chance of success. Ovum will watch Oracle with interest during the next year to determine its true commitment to selling in-memory technologies, which have the potential to challenge the company's core relational database business. This will probably not be the case, but it gives customers a further option for complementing traditional BI and analytic strategies grounded in relational database architectures.

A final caveat: under the surface, is the Oracle Exalytics architecture really that modern?

Exalytics bundles in some mature and well-established technologies, and is the latest part of Oracle's strategy to parallelize its entire data processing stack. In fact, all of the new "Exa"-branded Oracle-engineered systems are parallelized configurations of servers, networks, and storage, while the software is a virtual machine, operating system, database, and middleware. Parallelization of both, at the same time, is key. However, parallelism aside, Ovum believes that under the surface Oracle Exalytics does not necessarily represent a ground-breaking architecture. TimesTen was designed 15 years ago as an OLTP database, not an analytic or columnar database. (Oracle acquired it in 2005.) Another component, Oracle Essbase, is a 20-plus-year-old OLAP server that was never originally architected for in-memory processing; customers still have to pre-calculate and pre-aggregate data.



However, Oracle Exalytics is a "dot.one" release, and as with any such release, Ovum expects improvements and tightening of the architecture over time. In the meantime there are plenty of trade-offs and considerations for both Oracle and non-Oracle users to think about before they add Oracle Exalytics to their analytic strategies.



APPENDIX

Author

Madan Sheina, Lead Analyst, Information Management

madan.sheina@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.

Disclaimer

All Rights Reserved.

No part of this publication may be reproduced, stored in a retrieval system or transmitted in any form by any means, electronic, mechanical, photocopying, recording or otherwise, without the prior permission of the publisher, Ovum (an Informa business).

The facts of this report are believed to be correct at the time of publication but cannot be guaranteed. Please note that the findings, conclusions and recommendations that Ovum delivers will be based on information gathered in good faith from both primary and secondary sources, whose accuracy we are not always in a position to guarantee. As such Ovum can accept no liability whatever for actions taken based on any information that may subsequently prove to be incorrect.