

REPORT REPRINT

A new HOAP? Hybrid operational-analytic processing and the future of the database market

MATT ASLETT, JAMES CURTIS

04 DEC 2017

Data from 451 Research's Total Data Market Monitor indicates that databases designed to support a combination of operational and analytical processing workloads will quickly become mainstream, at least for new application projects.

THIS REPORT, LICENSED TO MEMSQL, DEVELOPED AND AS PROVIDED BY 451 RESEARCH, LLC, WAS PUBLISHED AS PART OF OUR SYNDICATED MARKET INSIGHT SUBSCRIPTION SERVICE. IT SHALL BE OWNED IN ITS ENTIRETY BY 451 RESEARCH, LLC. THIS REPORT IS SOLELY INTENDED FOR USE BY THE RECIPIENT AND MAY NOT BE REPRODUCED OR RE-POSTED, IN WHOLE OR IN PART, BY THE RECIPIENT WITHOUT EXPRESS PERMISSION FROM 451 RESEARCH.



©2017 451 Research, LLC | WWW.451RESEARCH.COM

451 Research has previously identified the emergence of a new breed of database providers with products that are positioned for a combination of operational and analytical workloads, as well as the systems of intelligence workloads that they are used for. Data from 451 Research's Total Data Market Monitor suggests that these databases that are designed to support hybrid operational and analytic processing (HOAP) will quickly become mainstream in the coming years – at least for new application projects.

THE 451 TAKE

The blending of operational and analytical systems continues to add value for many organizations. And while hybrid systems may not be an ideal fit for every firm, there are many reasons they do make sense. Beyond the reduction in maintaining a separate transactional and analytical system, hybrid databases enable organizations to carry out analytics on incoming operational data, taking advantage of the 'transaction window,' which, if done right, could be incredibly lucrative. While most existing database applications do not take advantage of hybrid functionality, and will continue to account for the majority of database revenue for many years, we expect HOAP workloads to rapidly account for a significant proportion of incremental database revenue, and that supporting them will come to be expected in any mainstream operational database product or service.

It has become an accepted best practice over the past 40 years or so that analytics should be performed on data stored in a separate database from that used to support operational, transactional systems. While there are data management benefits to be gained from this approach, its origins lie not in architectural elegance but rather the need to avoid the performance limitations of traditional systems, which made it impossible to support high volumes of database reads and writes in the same environment.

For the most part, this remains true today – with databases tuned for online transaction processing (OLTP), the default choice for operational workloads and online analytic processing (OLAP), the default choice for analytic workloads. While most database products can be used to support both OLTP and OLAP workloads, they will typically be tuned to support one or the other, and for performance reasons, would rarely be used to support both at the same time.

However, over the past five years, we have seen the emergence of a new breed of relational and non-relational database vendors that claim to have improved performance by taking advantage of hardware, memory and processor functionality to an extent that allows them to support operational and analytical workloads in the same instance. The incumbent database vendors are also in on the act – adding in-memory columnar engines to their existing row-based databases in order to support hybrid workloads.

In recent years, multiple terms have arisen to describe these database workloads, the most popular of which is probably 'hybrid transactional and analytical processing' (HTAP), as coined by research and advisory company Gartner. It is often assumed that the term 'transactional' implies ACID (Atomic, Consistent, Isolated, Durable) transactional integrity, which would exclude most NoSQL non-relational databases for this category. While this is not necessarily the case, we prefer to use the term 'operational databases' to avoid confusion. Hence, our use of the acronym HOAP.

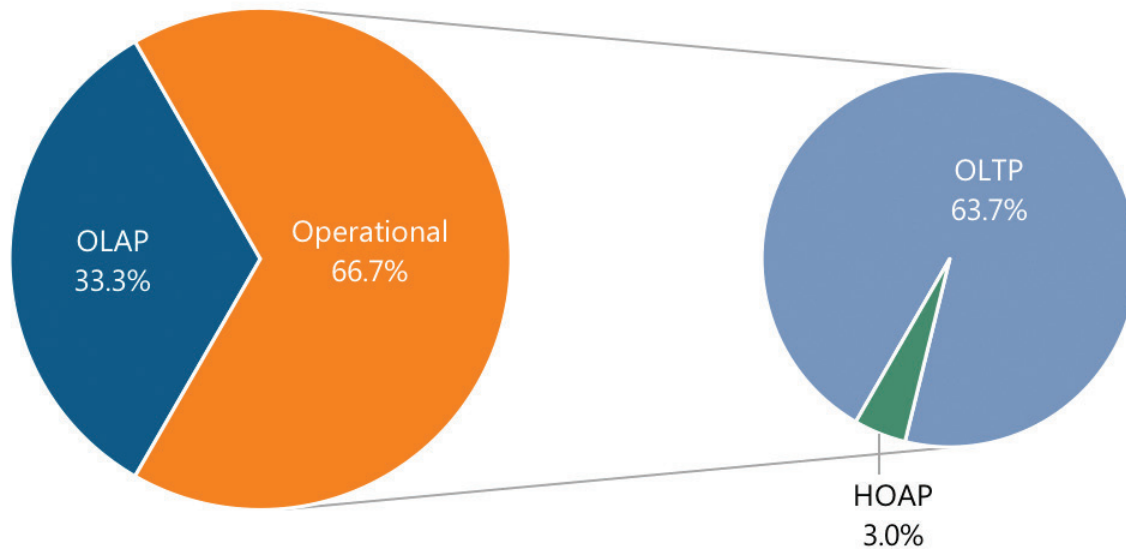
CURRENT ADOPTION

Adoption of databases capable of supporting HOAP is small compared to single-purpose deployments for transactional or analytical systems, which remain the de facto choice. Data from 451 Research's Total Data Market Monitor supports this – indicating that HOAP workloads accounted for only 1.6% of total database revenue in 2016, compared to 63.4% for OLTP workloads, and 35% for OLAP workloads.

We do see use growing, and expect it to continue to expand at a greater rate than it has up to this point. Our Total Data Market Monitor estimates indicate that HOAP workloads will grow to account for 3% of total database revenue in 2021, compared to 63.7% for OLTP workloads, and 33.3% for OLAP workloads.

2021 - Total Revenue

Source: 451 Research Total Data Market Monitor



As the chart above illustrates, while HOAP supports analytics on operational data, it is important to note that we do not see this as a direct replacement for OLAP workloads (for data marts and data warehousing). Instead, the primary use case is real-time analytics on operational data used to identify events from business applications before the data is transformed and loaded into a data warehouse/data mart.

Although advancements in database technologies and hardware are enablers, perhaps the greatest driver is that enterprises are becoming more and more comfortable with analytics, and are surprisingly open to adopting advanced analytics, particularly machine learning, to support systems of engagement with systems of intelligence (such as rules engines, decisioning systems, recommendation engines, natural language processing, and image recognition). With this as a context, organizations are zeroing in on the so-called 'transaction window' and realizing that it presents a significant opportunity – and once it's gone, it's gone for good.

Carrying out an analytic action, such as recommending an additional product during a retail purchase, is often used as a typical example of real-time analytics. However, another action takes place when the credit card is processed and the issuing company checks for fraudulent activity – in other words, it's possible to take multiple analytic actions on a single transaction. Other scenarios could include sensor or device readings in IoT. The transaction window makes real-time analytics especially important, because it encourages organizations to act quickly and decisively, with the greatest amount of analytics capacity possible, to avoid squandering the opportunity.

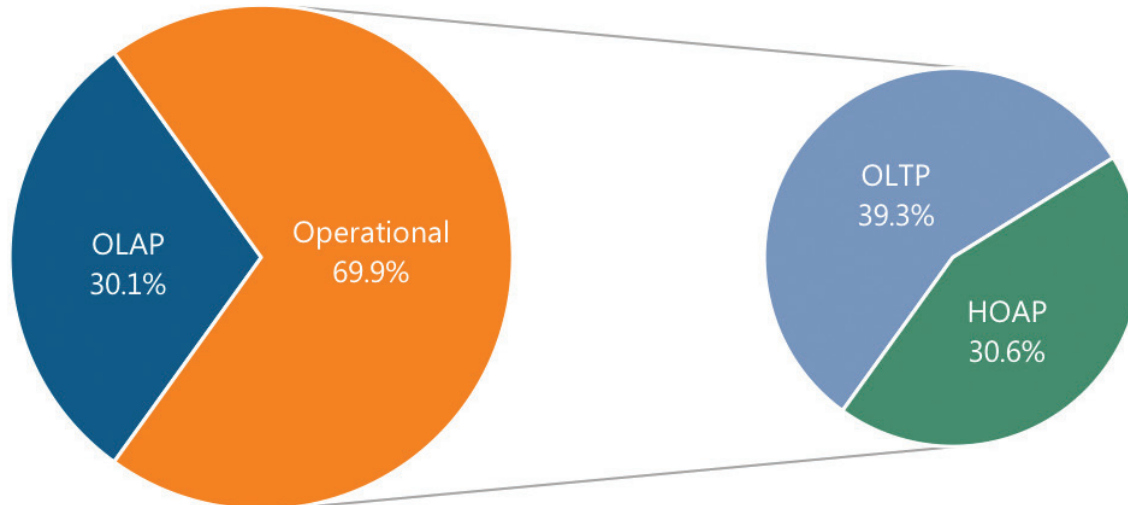
The need for real-time analytics within a transaction window is an important consideration for new application development projects and is likely to drive increased adoption of databases capable of supporting HOAP. As such, the illustration above is somewhat misleading in terms of the potential impact of HOAP, in that total database revenue is weighted in favor of existing applications that take advantage of traditional OLTP and OLAP database workloads.

GROWTH OPPORTUNITY

If we zero in on incremental revenue, we arguably get a better picture of the impact HOAP will have on database adoption choices in the immediate future. Our Total Data Market Monitor data shows that HOAP workloads accounted for 10.1% of incremental database revenue in 2016, compared to 59.6% for OLTP workloads, and 30.3% for OLAP workloads. By 2021, however, we estimate that HOAP workloads will account for 30.6% of incremental database revenue, compared to 39.3% for OLTP workloads, and 30.1% for OLAP workloads.

2021 - Incremental Revenue

Source: 451 Research Total Data Market Monitor



The suitability of a hybrid database often depends on the types of workloads and the nature of the business that organizations are involved in, but enterprises that previously thought a hybrid system would not work for them are reconsidering. One reason is that database technologies continue to improve in many areas, such as in-memory capabilities, hardware improvements like faster processors, cost reductions for memory and greater availability of SSDs. Hybrid systems can also mean having to maintain fewer systems, as well as reduced time for ETL processing or moving data from one system to the other.

PRODUCTS AND SERVICES

So which database vendors have products that would be considered suitable for HOAP workloads? SAP is a pioneer in this space with HANA, while the initial list of vendors targeting simultaneous operational and analytic workloads heavily overlaps with the new breed of NewSQL relational database vendors, including the likes of MemSQL, VoltDB, NuoDB, Clustrix and Splice Machine.

However, as noted above, the incumbent database vendors, including Oracle, IBM, Microsoft, and most recently Actian, have also introduced capabilities that would enable them to support HOAP workloads, while the various NoSQL database providers (MongoDB, DataStax, Redis Labs and Aerospike, for example) and in-memory data grid/cache providers (such as Pivotal, GridGain and GigaSpaces) are also increasingly being positioned for a combination of operational and analytical workloads. Another vendor is InterSystems with its IRIS Data Platform that leverages the company's Cache database to deliver hybrid workloads.

Indeed, as more database processing vendors have positioned to support HOAP (especially the incumbent vendors such as Oracle, IBM and Microsoft) and enterprise applications evolve to take advantage of HOAP functionality, the bigger question is whether it makes sense in the longer term to continue to make a distinction for HOAP at all.

Our chart above indicates that in 2021, HOAP workloads will represent more than 40% of new operational workloads. Certainly one could make a case that by that stage, the ability to support HOAP will be expected and assumed in any mainstream operational database, and that the applications considered OLTP will simply be those that don't take advantage of the hybrid functionality that is delivered as standard.