

REPORT REPRINT

451 Perspective: A HOAP-ful future for hybrid operational and analytical processing

MARCH 23 2020

By James Curtis, Matt Aslett

In 2017, we introduced a new term, hybrid operational and analytical processing, or HOAP, that formally identified a trend of blending both operational transactions and analytics within a single system or platform. Since then, we continue to see enterprises gravitate to hybrid workloads, with an increasing number of vendors developing products and services to satisfy the hybrid processing need.

THIS REPORT, LICENSED TO MARKLOGIC, 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.



Introduction

Hybrid operational and analytical processing, or HOAP, was a term we introduced in an earlier report to formally identify and track the rising trend of blending both operational transaction and analytics within a single system or platform. Ongoing 451 Research suggests that database systems that are designed to support hybrid operational and analytic processing (HOAP) will continue to mature in the coming years, addressing a variety of new applications as well as existing workloads.

451 TAKE

Hybrid operational and analytical processing, or HOAP, continues to see broad adoption for many enterprises. Part of the appeal of systems that are capable of hybrid operational and analytic processing is more than just an efficiency strategy of fewer systems to maintain; it's also the ability to do analytics on incoming operational transactions. In an era where enterprises are expected to be more data-driven, systems that can enable analytics where they were previously not available is certainly welcome news – and certainly beneficial to organizations to enable real-time decision-making. But as organizations have embraced and continue to embrace HOAP, there has also been a growing list of vendors that are either targeting these workloads specifically or offer HOAP capabilities as a component of their systems. This makes for a rather competitive market. In nearly all cases, HOAP functionality is coming from vendors that peddle transactional systems and are adding analytics to them, although there's at least one vendor providing transactions on an analytic database. Research shows that HOAP is steadily taking a bite out of pure OLTP workloads and is expected to do so into the future.

Context

Hybrid operational and analytic processing – processing that can handle both transactional and analytical workloads – is an area that we have been watching and tracking closely for many years now. It is well known that the commonly accepted practice for handling both transactional, operational workloads and analytical workloads has been to keep them separated, each workload running in separate systems. The fact that one process may impede upon the other – long-running analytical queries affecting incoming transactions, for instance – is just one of the many reasons why it makes sense to separate out these two workloads.

But architecturally, it has made sense to separate these processes as well. Transactional processing often leverages row storage – storage that maintains the data in rows and columns. A table listing individuals would likely be represented with a different person in each row with each column denoting a different category: name, address, phone and so forth. For analytical-based systems – data warehouses, for instance – the data is often stored in a columnar format, where a column represents a particular area. How the data is stored offers certain benefits. Row storage, for instance, can write a lot of transactions quickly while also enabling consistency (normalization) within the database and can handle complex queries with joins. Conversely, columnar storage enables fast analytical queries because the data is stored in columns – indexed and sometimes compressed – so that they can be scanned. But the data may also be denormalized (multiple copies of the same data) and that would present challenges for processing transactions because data consistency would be required.

REPORT REPRINT

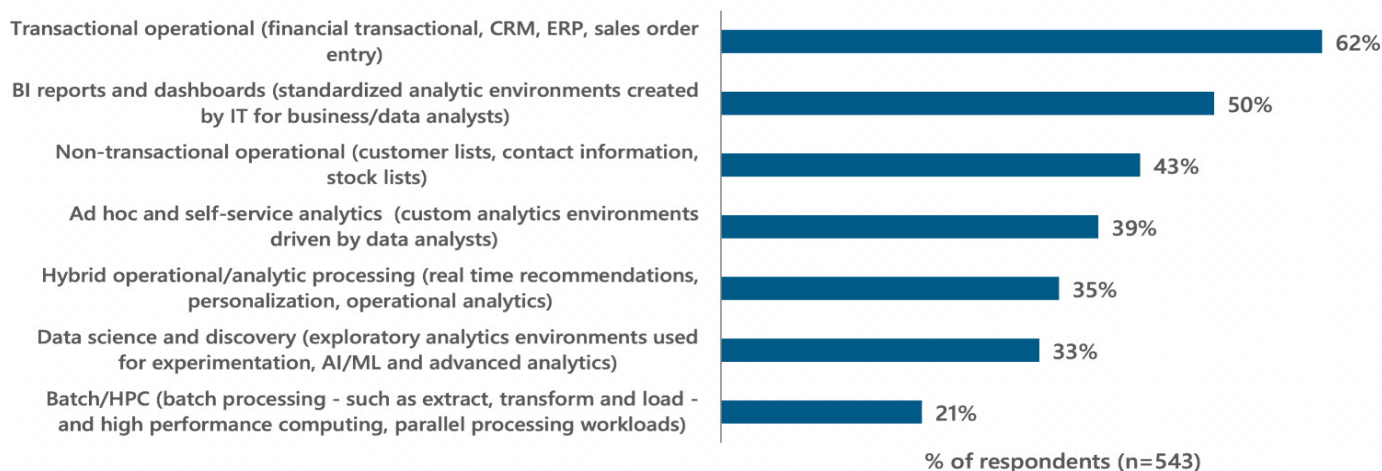
Background on not only how transactional and analytical processing occurs but also why makes the rise of hybrid database systems all the more remarkable. We noted in our earlier report how we have witnessed the emergence of a new breed of relational and non-relational database vendors that have developed specific technology to address HOAP workloads.

Current HOAP adoption and growth

In Figure 1, taken from 451 Research's *VotE: Data, AI, & Analytics 1H 2019* report, survey respondents were asked which type of workloads they were currently running on their data platform systems. Not surprisingly, transactional processing and analytical workloads were the top two responses, at 62% and 50%, respectively. But survey respondents also cited hybrid operational analytic processing at 35%, indicating that HOAP is not a niche workload but a known workload type for many enterprises.

Figure 1: Workloads for Data Platforms

Source: 451 Research, *VotE: Data and Analytics, Workloads and Adoption Patterns 2019*



Current adoption

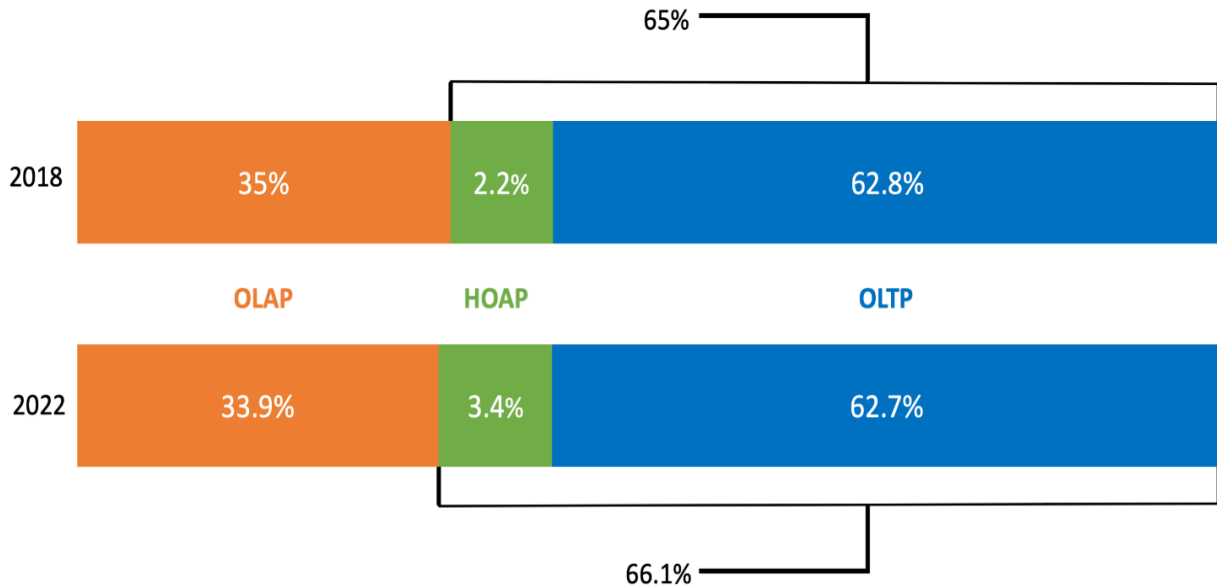
While 451 Research survey data indicates HOAP workloads are being adopted by many enterprises, these workloads still represent a small fraction of the total overall revenue generated from analytical processing (OLAP) and operational workloads (OLTP). Figure 2 shows data from 451 Research's *Data, AI and Analytics Market Monitor* service that supports this - indicating that HOAP workloads accounted for only 2.2% of total database revenue in 2018, compared with 62.8% for OLTP workloads and 35% for OLAP workloads.

Likewise, we see HOAP use growing and expect it to continue to do so at a greater rate than it has up to this point. Our *Data, AI and Analytics Market Monitor* service estimates suggest that HOAP workloads will grow to account for 3.4% of total database revenue in 2022, compared with 62.7% for OLTP workloads and 33.9% for OLAP workloads. However, it's important to note that while the HOAP percentages appear to be modest, our *Data, AI and Analytics Market Monitor* research points to the market growing by nearly \$20bn from 2018 to 2022, from \$58.6bn to \$78.5bn in 2022. A sizeable proportion of that total database revenue will be maintenance revenue for existing traditional transactional and analytic database workloads.

REPORT REPRINT

Figure 2: 2018-2022 Total Revenue

Source: 451 Research's Data, AI and Analytics Market Monitor



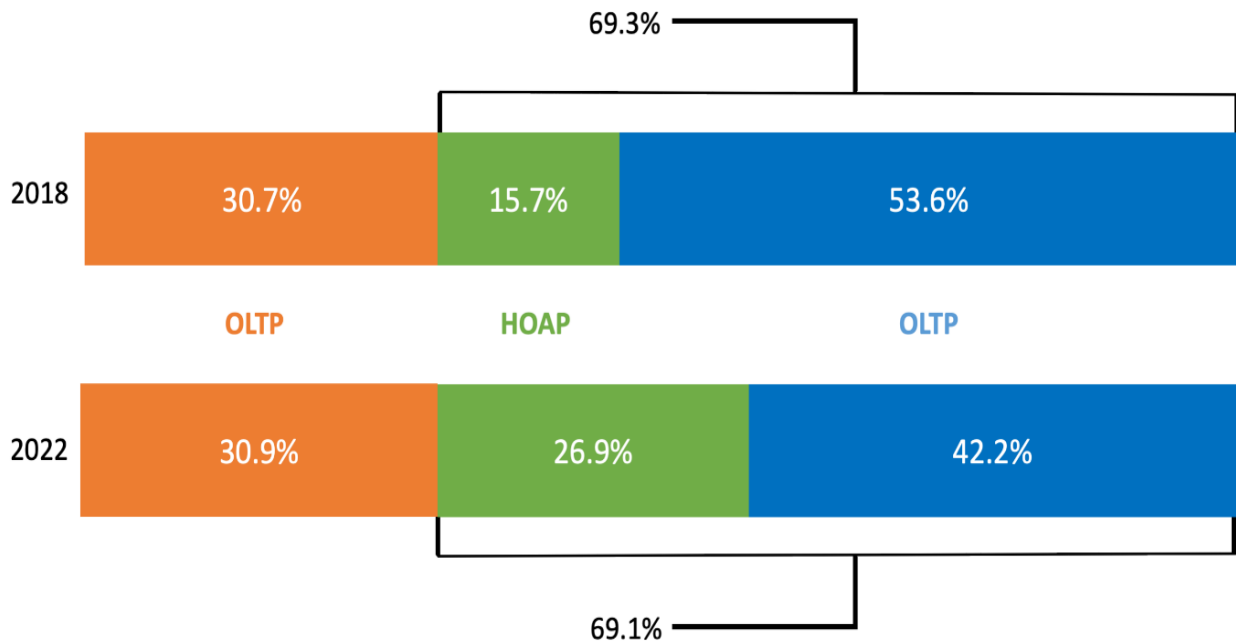
Growth opportunity

On that basis, we can get a greater insight of HOAP growth by looking at incremental revenue for new database workloads rather than total revenue, as shown in Figure 3. Our Data, AI and Analytics Market Monitor data shows that HOAP workloads accounted for 15.7% of incremental database revenue in 2018, compared with 53.6% for OLTP workloads and 30.7% for OLAP workloads. By 2022, we estimate that HOAP workloads will account for 26.9% of incremental database revenue, compared with 42.3% for OLTP workloads and 30.9% for OLAP workloads.

REPORT REPRINT

Figure 3: 2018-2022 Incremental Revenue

Source: 451 Research's Data, AI and Analytics Market Monitor



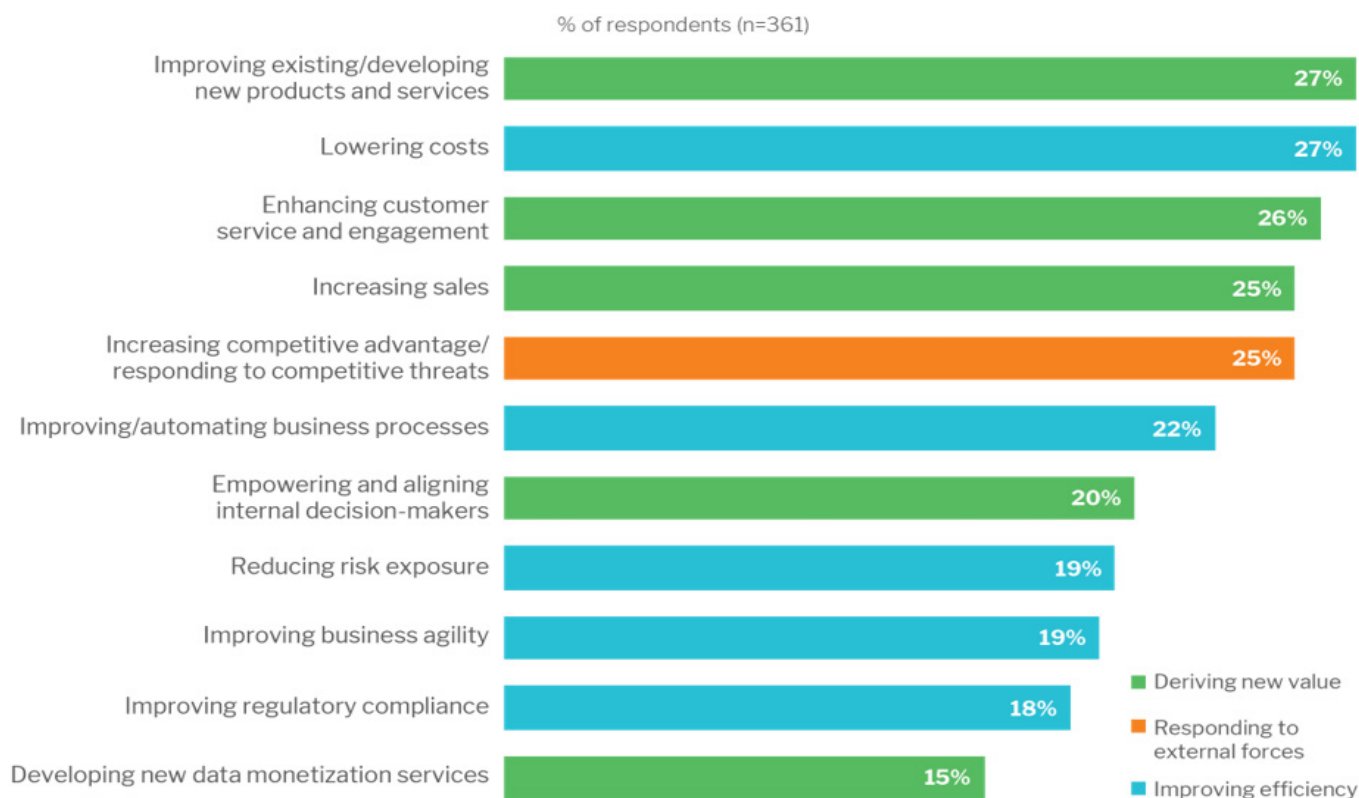
A few takeaways from the chart are worth noting. First, HOAP workloads are expected to be an even greater percentage of transactional, operational workloads, going from 22.6% in 2018 to 38.9% in 2022. Moreover, this increase will be reflected in real revenue dollars, where it is expected to double. That is, incremental HOAP revenue in 2018 was forecasted at \$735m and is expected to increase to \$1.4bn in 2022.

Drivers of adoption

What is driving this growth? One possible explanation for hybrid processing adoption can be found in 451 Research's *VotE: Data, AI, & Analytics 2019 2H* report regarding the use of data for decision-making, shown in Figure 4.

Figure 4: Benefits of Being More Data-Driven

Source: 451 Research's *VotE: Data, AI, & Analytics 2019 2H*



In Figure 4, survey respondents were asked about the benefits of being data driven. Except for increasing competitive advantage and responding to threats, which include responding to external needs, respondents found value in two core benefits: deriving new value and improving efficiency. We are not suggesting that HOAP is entirely the reason for enterprises becoming data driven, but it does play a role. For instance, combing transactional systems with analytics does reduce IT complexity. HOAP also can address customer engagement and possibly increase sales by applying recommendations to incoming transactions.

Another driver for growth includes recent technology innovations, including open source projects, which are aiding vendors in developing specific databases capable of handling hybrid workloads. For instance, there are many storage engines available, usually offered as open source, which have given rise to how data can be stored. There are multiple compute engine choices and hardware advances, particularly with memory, including the adoption of persistent memory.

Cloud computing has also taken off, particularly enabling the ability to split compute and storage for analytics, including the adoption of cloud object storage. Some vendors are also able to support hybrid IT systems blending on-premises with cloud infrastructure, taking advantage of deployment locations to optimize workloads and queries. These technologies have in a sense provided a sort of perfect storm and have thus aided vendors in developing HOAP systems.

451 Research's data illustrate that the database market is accelerating toward both database-as-a-service across hybrid IT systems and hybrid operational and analytic processing. Further, our research indicates that it will be those vendors that offer a consistent experience across the broadest portfolio of options that will likely benefit most from the changing market dynamics.

Vendors that provide HOAP

The HOAP vendor landscape has also evolved a good deal over the past few years. While not exhaustive, Figure 5 lists the primary vendors that either actively position their products as HOAP systems or market their systems or platforms as providing capabilities for HOAP. Noteworthy is the fact that the vendors providing HOAP hail from nearly every category of the database landscape, including relational databases, analytic databases, NoSQL and non-relational databases, distributed data processing frameworks (Hadoop/Spark) and distributed data grid/cache systems.

Figure 5: HOAP Vendors

Source: 451 Research

VENDOR	PRODUCT
ACTIAN	Actian X
AEROSPIKE	Aerospike Database
ALIBABA CLOUD	HybridDB for MySQL
AWS	Amazon Aurora
CLOUDERA	CDP Data Hub, Enterprise Data Hub
DATASTAX	DataStax Enterprise
ENTERPRISEDB	Enterprise Postgres
ESGYN	EsgynDB
FAIRCOM	c-treeACE
GIGASPACES	XAP
GRIDGAIN	GridGain Enterprise
IBM	IBM Db2
INTERSYSTEMS	InterSystems IRIS
MARIADB	MariaDB Platform X4, SkySQL
MARKLOGIC	MarkLogic Data Hub Platform
MEMSQL	MemSQL
MICROSOFT	Microsoft SQL Server 2019, Azure SQL Database
MONGODB	MongoDB Enterprise Server
NUODB	NuoDB
ORACLE	Oracle Database, Autonomous Transaction Processing

REPORT REPRINT

VENDOR	PRODUCT
PINGCAP	TiDB (TiFlash, TiKV)
REDIS LABS	Redis Enterprise
SAP	SAP HANA
SPLICE MACHINE	Splice Machine Data Platform
TIBCO	ComputeDB (SnappyData)
VMWARE	GemFire, Greenplum
VOLTDB	VoltDB
YUGABYTE	YugaByteDB