Bloor

Market Report

Market Report Paper by Bloor
Author **Philip Howard**
Publish date **December 2017**

SQL Engines on Hadoop

> **It is clear that Impala, LLAP, Hive, Spark and so on, perform significantly worse than products from vendors with a history in database technology.**

Author **Philip Howard**

Executive summary

Hadoop is used for a lot of different purposes and one major subset of the overall Hadoop market is to run SQL against Hadoop. This might seem contrary to Hadoop's NoSQL roots, but the truth is that there are lots of existing investments in SQL applications that companies want to preserve; all the leading business intelligence and analytics platforms run using SQL; and SQL skills, capabilities and developers are readily available, which is often not the case for other languages.
However, the market for SQL engines on Hadoop is not mono-cultural. There are multiple use cases for deploying SQL on Hadoop and there are more than twenty different SQL on Hadoop platforms. Mapping the latter to the former is not a trivial task, as different offerings are optimised for some purposes but not others.

The key differentiators between products are the use cases they support, their performance and the level of SQL they offer. While all of these are discussed in detail in this paper it is worth briefly explaining that SQL support has two aspects: the version supported (ANSI standard 1992, 1999, 2003, 2011 and so on) plus the robustness of the engine at supporting SQL queries running with multiple concurrent thread and at scale.

Figure 1 illustrates an abbreviated version of the results of our research. This shows various leading vendors, and our estimates of their product's positioning relative to performance and SQL support. Use cases are shown by the colour of each bubble but for practical reasons this means that no vendor/product is shown for more than two use cases, which is why we describe Figure 1 as abbreviated. Thus, for example, we are using "EDW" as shorthand for products that support both transactional lookups and complex analytics, which are otherwise individual use cases. Also, it excludes vendors targeting OLAP, as the leaders in this market – Jethro Data and Kyvos Insights – have distinct approaches that are not easily compared.

> **The key differentiators between products are the use cases they support, their performance and the level of SQL they offer.**

Figure 1 –
Use cases by performance and SQL support. Use cases include Hybrid Transactional and Analytic Processing (HTAP), a merger of the transactional look-ups and complex analytics (EDW: enterprise data warehouse), combined batch and real-time/streaming analytics (Lambda architectures), and machine learning (ML). OLAP and some other use cases are omitted.

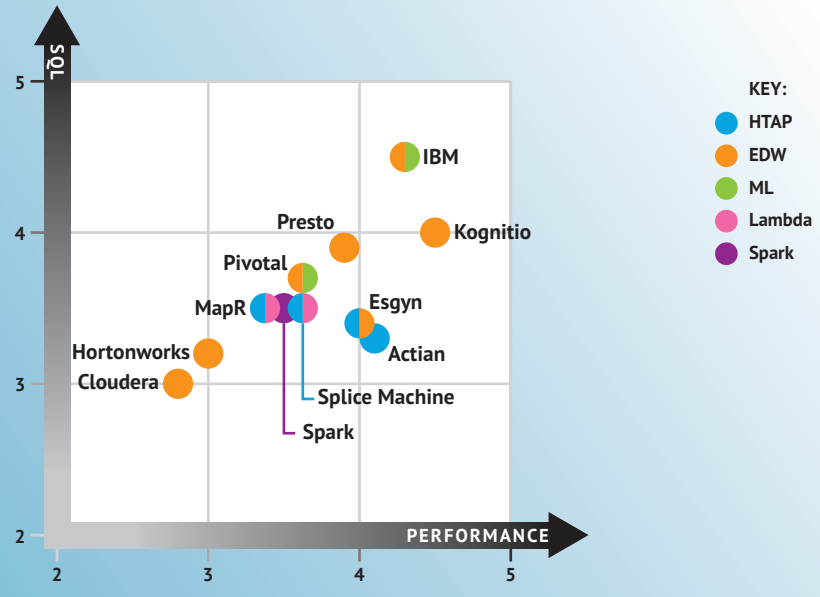

Use cases

We have identified six different use cases for SQL on Hadoop. Some of these overlap one another and there will also be instances where a user wants more than one of these use cases running on the same cluster. However, we believe that the examples detailed provide the bedrock for making decisions about potential solutions.

The main use cases we have identified, in no particular order, are:

1. Transactional look-ups. This will often be combined with other use cases.
2. Hybrid transactional analytic processing (HTAP).
3. Complex queries against large datasets. Typically involving many users. We might describe this as "traditional data warehousing" and, certainly, there are vendors aiming to replace enterprise data warehouses (EDW) via this use case. Often combined with transactional look-ups.
4. Online analytic processing (OLAP). May be either multi-dimensional OLAP (MOLAP) or relational OLAP (ROLAP).
5. To support machine (and deep) learning.
6. A "collapsed" lambda (or kappa) architecture designed to support both batch and real-time (streaming) analytics. Will often be combined with either or both of OLAP and machine learning,

There are several other uses cases where you might want to use SQL on Hadoop but, often enough, Hadoop on its own will be enough. These use cases include extract, load and transform (ELT) and archival, as well as (ad hoc) data preparation. The last of these was identified as a use case by one of the vendors, although none of the suppliers – including the identifier – we have spoken to, have claimed to target it. The same applies to data discovery and similar use cases where you would probably be better off to rely on an information/data catalogue running on your data lake. One vendor also suggested a use case as an operational data store.

"We have identified six different use cases for SQL on Hadoop. Some of these overlap one another and there will also be instances where a user wants more than one of these use cases running on the same cluster."

Offerings

Products in this market tend to fall into one of six categories and in the following lists we have highlighted those products we examine in more detail in this report. The groupings consist of:

- Pure-play open source projects. This category includes Hive, HBase, Tajo, Phoenix, Ignite and Spark. See also the OLAP-based projects below. All of these are Apache projects. Of the less well-known offerings Phoenix supports on-line transaction processing (OLTP) running against HBase; Ignite is an in-memory computing platform that is commercially supported (and was originally developed) by GridGain. It is typically used either as a Hadoop accelerator and/or to provide immediate consistency. Tajo is a big data warehouse. There have been no new releases of Tajo for 18 months, so we suspect that it is defunct.
- Vendor supported open-source projects. This group includes **Drill** (supported by MapR), **Presto** (Teradata/Starburst Data), HAWQ (**Pivotal**) and Trafodion (**Esgyn**). All of these, again, are Apache projects. Also in this category are Impala (Cloudera) and Hive + LLAP (Hortonworks – live long and process - previously known as Stinger). Note that Drill does not have to run on Hadoop.
- Traditional data warehousing products that have been used as the basis for SQL on Hadoop platforms. These include IBM Db2 (**Big SQL**), Oracle, Vertica, **Pivotal HDB** (HAWQ: effectively a port of Greenplum), **Kognitio** (which is free-to-use) and **Actian VectorH**. VectorH is the odd one out here because Actian Vector is a symmetric multi-processing (SMP) solution that has been developed into a massively parallel processing (MPP) environment. All the other products were MPP-based originally.
- Other MPP-based solutions. This category consists of Transwarp and **Esgyn**. The latter is is a descendant of Tandem NonStop, HP Neoview and other HPE-based warehousing developments.
- Specialist offerings. Mostly these are targeted at OLAP environments. In this category are Apache Kylin (MOLAP) and Apache Lens (ROLAP) as well as **Kyvos Insights** and **Jethro Data**. **Splice Machine** is also in this category but has rather broader capabilities (see later). AtScale will compete with products in this category but is a "BI on Hadoop" engine rather than a SQL on Hadoop platform: as such it is not discussed further here.
- Others that are often referred to as SQL on Hadoop engines, but which are not. Included in this category are Splout SQL, which is really about data serving, and Concurrent Lingual, which is used for application development. Druid, which started life as an MDX engine (and which now has limited SQL support) is another data serving product with OLAP capabilities. Apache Calcite is a general-purpose SQL optimiser but not an engine per se. None of the products in this group are discussed in this report.

In the vendor/product section of this report we include short descriptions of many, though not all, of the proprietary products (open source or otherwise), with the exception of Oracle, Vertica and Transwarp, none of which responded to our requests for information. While the omission of Oracle and Vertica is no great loss (a straight line can be drawn across from their traditional products), we would have liked to include details about Transwarp.

"Traditional data warehousing products have been used as the basis for SQL on Hadoop platforms. These include IBM Db2 (Big SQL), Oracle, Vertica, Pivotal HDB (HAWQ: effectively a port of Greenplum), Kognitio (which is free-to-use) and Actian VectorH."

Performance benchmarks

A great many vendors in this space have conducted and published benchmarks. Some of these have been validated by third parties, some of them have been conducted by third parties, but the majority have not involved any independent authorities. Although TPC (transaction processing council) tests have typically been the basis for these benchmarks, none of them have been authenticated by TPC. The individual product descriptions that follow outline the results of the various benchmarks that have been performed by different vendors. We will therefore confine ourselves here to general comments.

The first point that we would like to note is that TPC-DS (Decision Support) tests are not just an indicator of performance but also of SQL support. TPC-H, on the other hand, is based on SQL 92, which is hardly up-to-date. We are disappointed with Actian, therefore, that it is focused on TPC-H and not TPC-DS.

The second point is that many tests are done using relatively small datasets and a single processing thread, when what you are really want is multiple users running against large sets of data. IBM, for example, has demonstrated that while Spark is perfectly capable of running all TPC-DS queries at small scale it breaks down as you scale up.

Thirdly, some vendors, notably Hortonworks and Cloudera, both of which have been guilty of publishing partial results. For example, just selecting (no doubt the best ones) 15 of the 99 TPC-DS test to report on.

To conclude this section – while not all products have been benchmarked and some have been benchmarked against different standards – it is clear that Impala, LLAP, Hive, Spark and so on, perform significantly worse than products from vendors with a history in database technology. Moreover, it is much more likely that companies in the latter category will be able to support all of your queries and run them successfully: the level of SQL support from the pure-play, Cloudera or Hortonworks products, tends to be limited.

While on the subject of SQL support, it is worth commenting that the level of support for ANSI standard SQL varies widely. IBM – not just in Big SQL and Db2, but across its product range – is much the most advanced vendor in this respect. Conversely, there are a number of products whose ANSI support dates back to the last century.

“The level of support for ANSI standard SQL varies widely. IBM – not just in Big SQL and Db2, but across its product range – is much the most advanced vendor in this respect. There are a number of products whose ANSI support dates to the last century.”

Product suitability

While performance may be a major determinant in buying decisions, it is only relevant when comparing apples with apples, and the products covered in this paper constitute an entire fruit bowl. In this section we therefore match products to use cases.

1. Transactional look-ups. This will often be combined with use case 3 (see above). Various products, often in conjunction with HBase are suitable here. Notable contenders would be IBM Big SQL (with HBase), Splice Machine (which incorporates HBase into its Lambda architecture), Esgyn and Actian VectorH.
2. Hybrid transactional analytic processing (HTAP). This is a major focus area for both Esgyn and MapR Drill. Splice Machine is also a suitable contender here, though its emphasis is slightly different (more on leveraging transactional data for predictive analytics than embedding analytics into operational applications). The InterSystems IRIS Data Platform also competes here though it is not based on Hadoop (but is a clustered solution) as do others.
3. Complex queries against large datasets. Typically involving many users. We might describe this as "traditional data warehousing" and, certainly, there are vendors aiming to replace enterprise data warehouses (EDW) via this use case, often combined with transactional look-ups. All the "ported" data warehouses play in this space, as do Actian VectorH, Esgyn and Splice Machine. Kognitio comes out well in the benchmark studies we have investigated.
4. Online analytic processing (OLAP). May be either multi-dimensional OLAP (MOLAP) or relational OLAP (ROLAP). The vendor-based products in this area are much stronger than any open source offerings. Kyvos Insights, Jethro Data and Splice Machine are the vendors to consider.
5. To support machine (and deep) learning. Pivotal, IBM and Splice Machine are the companies most active in this area, but where IBM relies on Spark MLlib, Pivotal is a major contributor to the Apache MADLib project. Splice Machine ships with MLlib.
6. Lambda architectures. Both MapR and Splice Machine are in the business of "collapsing" lambda architectures to support batch and streaming analytics into a single platform. In the case of Splice Machine, a Spark processing engine is embedded into the platform. In this context it is worth commenting that independent benchmarking has found that tight integration with Spark results in a 11x performance improvement compared to simply connecting to Spark. We would expect an embedded engine to do even better.

There are two other use cases worth commenting on. Esgyn has identified operational data stores as a target use case. More interestingly, MapR Drill supports queries against semi-structured data such as JSON, as well as structured data. It has extended its SQL support to allow this. Competitors to MapR for this sort of functionality tend to come from other environments: the InterSystems IRIS data platform, for example, encompasses the same capabilities and extends to unstructured data.

"While performance may be a major determinant in buying decisions, it is only relevant when comparing apples with apples, and the products covered in this paper constitute an entire fruit bowl."

Conclusion

SQL on Hadoop is all about horses for courses and, in this paper, we have discussed both the horses and the courses. Table 1 highlights our results. Readers should recognise that you can do OLAP, for example, with any EDW product, but the likes of Kyvos and Jethro will typically provide better performance, hence our recommendations. We have also suggested some SQL but not Hadoop-based vendors that you might like to consider as alternatives to the SQL on Hadoop products, though these are not intended to represent an exhaustive list. Specifically, we have concentrated on scale-out clustered solutions and have omitted products such as IBM Informix or SAP HANA, both of which target HTAP (for example), because they employ architectures that are a long way removed from Hadoop.

Use case	Transactional	HTAP	Complex	OLAP	M/L	Lambda	Mixed data
Recommended	IBM Big SQL Actian VectorH	Esgyn MapR Drill Splice Machine	Kognitio IBM Big SQL Presto	Jethro Data Kyvos	Pivotal IBM Big SQL	Splice Machine MapR Drill	MapR Drill
Others	Esgyn		All EDW	Splice Machine All EDW	Splice Machine Some EDW		
Non-Hadoop		InterSystems		AtScale Druid			InterSystems

Actian VectorH

Actian VectorH (Vector for Hadoop) is Actian's Vector product ported to the Hadoop platform as a SQL engine. However, Vector uses a symmetric multi-processing (SMP) based architecture that scales up rather than out. VectorH, on the other hand is, by necessity, a massively parallel processing (MPP) solution that uses Hadoop for clustering purposes. So, this represents more than just a port from one MPP environment to another. It is typically implemented on either a Cloudera or Hortonworks platform with MapR supported upon request. While the principle components of the product are based on the Vector product, it also leverages the query planner and optimiser from what used to be known as Ingres but is now called Actian X. While there is access to Parquet and ORC files, which are treated as external tables, Actian has developed a proprietary storage mechanism on top of the Hadoop distributed file system (HDFS). Apache YARN is supported and available security is both row and role-based.

VectorH, like Vector is an in-memory, columnar database. As its names suggests, one of its major differentiators is the vectorised processing that it supports. It uses positional delta trees (in memory) that minimise the impact of updates while performing query operations. Both Spark and Scala are supported.

Actian has performed benchmark testing comparing the performance of VectorH with Hive, Impala, HAWQ and Spark. VectorH typically outperforms the others by an order of magnitude. However, there are a couple of caveats to be made about these figures. Firstly, the benchmarking was performed during the first half of 2016 and is therefore somewhat out of date. Secondly, these were TPC-H tests rather than the more common TPC-DS, which are usually regarded as more suitable for SQL on Hadoop engines. These are more onerous (there are 22 TPC-H queries but 99 in the TPC-DS standard) and do more to test the SQL syntax supported by the various engines.

Actian
2300 Geng Rd., Suite 150
Palo Alto, CA 94303

www.actian.com

Strengths

- While the performance figures quoted by Actian are out-of-date and not based on TPC-DS, the results are, nonetheless, impressive. This suggests that VectorH is a serious contender for traditional enterprise data warehousing (EDW) workloads. Indeed, drawing a line through other benchmarks – conducted by both vendors and independently – would suggest that VectorH has probably got the performance edge against all of its rivals.

Threats

- Performance isn't everything and at least some competitive vendors are offering specialised capabilities that go beyond supporting "traditional EDW workloads".
- Actian is targeting the general-purpose EDW market running on Hadoop. However, this is a crowded space, with many traditional EDW vendors also addressing this market. The company has not traditionally been known as a player in the high end of this market where solutions are MPP-based, so the company has a clear marketing issue in getting its presence known.

Recommended for

HTAP, transactional look-ups.

Also appropriate for

Mixed workload environments that include complex analytics.

EsgynDB

Esgyn Corporation, the developer of EsgynDB, is based in the United States and China. The company was founded by ex-HP staff that had worked on successive generations of that company's database technology, starting with Tandem NonStop SQL, via HP's enterprise data warehouses, to EsgynDB, which represents the 4th generation of development, and which is a SQL on Hadoop engine focused on HTAP (hybrid transactional and analytic processing), supporting mixed workloads and multi-structured datasets. It is a commercial implementation of the Apache Trafodion database, to which Esgyn is the prime contributor. The implication of supporting HTAP is that Esgyn is suitable for use across a range of processing environments including acting as an operational data store or as enterprise data warehouse, but its primary differentiator is its support for HTAP. The product is available both in the cloud and on premises. It is also, in conjunction with Supermicro, available as an appliance. EsgynDB has been deployed in a variety of areas, including Internet of Things (IoT), Banking and Insurance, Telecom, Manufacturing, Internet Security and Smart Cities.

EsgynDB is a massively parallel processing (MPP) based database that typically leverages Trafodion tables but which also supports Hive, HBase, Parquet, ORC and Avro formats. The product supports extensible user defined functions (UDFs) that allow external environments to be treated as tables. Currently supported external environments include Apache Solr and JanusGraph (which is a Linux Foundation open source graph database). In the latter case, you can embed Gremlin code into your UDF so that you can join graph data with data managed directly by EsgynDB. The product also integrates with and leverages technologies such as Hibernate, Spark and various streaming engines. In terms of other features, the product supports a wealth of capabilities, including secondary indexes, the ability to assign parallel resources as required (a part of its workload management), the sorts of high availability that you would expect from a product with a Tandem heritage, load balancing and so on. There is a distributed query manager to support transaction processing and the product is fully ACID compliant. SQL support is based on ANSI 2003 with extensions that support Oracle PL/SQL and Teradata functions, to encourage users to move off those platforms and onto EsgynDB.

From a performance perspective Esgyn has conducted benchmarks using both TPC-C and TPC-DS. In the case of the former, the product has demonstrated linear scalability for transaction processing between 6 and 13 node deployments, with the latter supporting just over 345,000 transactions per minute. In terms of query performance, EsgynDB runs all 99 TPC-DS queries (which is by no means always true of other SQL on Hadoop engines) with an average performance advantage of 2.7 times versus Impala version 2.2 and 5 times versus Hive 1.2.1. We should add by way of a caveat that benchmarks are moving feasts, as vendors introduce new capabilities, but of course this applies as much to Esgyn as to Cloudera or Hortonworks.

Esgyn
691 South Milpitas Boulevard, Suite 117
Milpitas, CA 95035

www.esgyn.com

Strengths

- Almost all proprietary vendors within the SQL on Hadoop space concentrate on one or another aspect of business intelligence or analytics. While they might have transaction support this is not an area they focus on. EsgynDB has a major advantage in its support for HTAP.
- Given its background and history we have no doubt that EsgynDB is robust and richly featured. Performance benchmarks are one thing but the true test is concurrency and here we would expect Esgyn to have a significant advantage. It is, for example, one of the few SQL on Hadoop engines to provide spill-to-disk capabilities when there is memory pressure.

Threats

- Esgyn is targeting other environments as well as HTAP. For transaction processing on Hadoop it has similar advantages as for HTAP. However, for purely analytic workloads it faces much more competition, from both open source and proprietary vendors. The company also claims benefits for use as an operational data store, though we are not sure that there is a significant market for this capability.

Recommended for

HTAP.

Also appropriate for

Mixed workload and other traditional environments.

IBM Big SQL

Big SQL is IBM's SQL on Hadoop engine. It used to be accompanied by BigInsights and other "Big" products, but the company's recently announced partnership with Hortonworks, means that these other products have now been, or are being, replaced. Big SQL runs in conjunction with Apache Hive and, optionally (if you want transactional look-ups), Apache HBase. In the latter case, the alternative would typically be to use Apache Phoenix along with HBase, but this would mean separate connections for Hive and HBase, whereas Big SQL does this with one connection as well as allowing transparent joins across Hive and HBase. Similarly, Big SQL provides a unified security architecture as opposed to two distinct security models. And finally, it means ANSI (2011) standard SQL support from IBM, as opposed to SQL that does not comply with any ANSI standard if you are using Phoenix.

Apart from that, Big SQL is based on Db2 and leverages Db2 capabilities such as the Db2 optimiser and the Db2 federation capabilities. The product does not just support processing via Spark – many platforms do that – but directly integrates with Spark, which will be useful if, for example, you want to extend your warehouse environment to support machine learning, or if your data lake is to support streaming analytics.

IBM has previously run benchmarks against Hive (on its own), Impala, HAWQ and Spark. These were based on TPC-DS both on a single stream of queries and for four concurrent queries, in both cases running against a 10 TB database. These were not official results (in fact, no one has posted official TPC-DS results) but "based on" TPC-DS. The most interesting outcome is a claimed 3.5 times performance improvement compared to Apache HAWQ. A more recent benchmark test comparing Big SQL and Spark 2.1, with a 100 TB database, found that Big SQL was 3.2 faster and used only a third of the resources required by Spark SQL. Moreover, although Spark can execute all TPC-DS queries at smaller scale, it was unable to do so at this level, completing only 83 of the 99 queries.

IBM
1 New Orchard Road, Armonk
New York 10504-1722, USA

www.ibm.com

Strengths

- Big SQL has the most advanced ANSI standard SQL implementation of any product in this market. Most other vendors, when asked, are vague about what they support or mutter about SQL 2003 (if you're lucky). There is also substantial support (over 95%) for Oracle PL/SQL: a capability inherited from Db2, as well as the SQL used by Netezza (IBM).
- Unlike many (not all) providers of SQL on Hadoop engines, IBM is quite clear about its products' strengths and it is firmly focused on complex analytics serving multiple users (which you might think of as traditional data warehousing), which may be extended to include transactional/fast look-up environments when Big SQL is deployed in conjunction with HBase.
- The Db2 heritage for Big SQL means that we would have no qualms about recommending it for high availability, resilience, security, load balancing and all those other capabilities that you would expect from a product with the maturity and longevity of Db2.
- Big SQL provides both row- and column-level security. This is unlike most competitors, which typically support one or the other but not both.

Threats

- In principle, there should be no particular threats to Big SQL. However, we have the impression that, outside the company's user base, companies do not really think of IBM as being in the Hadoop space. This also applies – despite ample evidence to the contrary – when it comes to open source projects. This is a perceptual issue rather than anything else, but it is real nevertheless.

Recommended for

Transactional look-ups, complex analytics, mixed workload environments, machine learning.

Jethro Data

Jethro is a SQL on Hadoop engine specialising in on what we might call "traditional business intelligence". That is to say, on online analytic processing (OLAP) with, necessarily, the ability to drill down through your OLAP cubes to fine level detail. And to do this at scale and with high performance. It is this last point – performance – that is most significant. To get the sort of performance you would like, Jethro offers two primary capabilities: auto-cubes and indexing, where the optimiser will use the former for aggregated queries and the latter for more granular queries or, where that is appropriate, both. Existing applications and queries built using MicroStrategy, Tableau, Qlik and so on, should run against Jethro without change.

In Jethro, all columns and measures (things like "price" are measures) are indexed. Further, there is an index of indexes. In order to ensure performance, indexes are never locked, and the company takes an append-only approach to avoid the need for locking when there is an update to the database. When you add new data, Jethro creates indexes on this incremental data which are then appended to the original indexes. There is a background index merge process as well as capabilities to identify and handle duplicates.

The way that auto-cubes work is that the first time that you get a query from Tableau (say), Jethro creates the relevant cube using indexes. Thereafter, Jethro serves that query from the auto-cube it has created. Moreover, it will also serve variants of that query from the auto-cube, which it does by treating any filters (for example, queries by country) as external to – but associated with – the auto-cube. There are similar processes involved for adding incremental data to a cube as for indexes.

We are not aware of any comparative benchmarks comparing Jethro with other SQL on Hadoop engines. However, the company does two live benchmarks that you can run for yourself at ***tableau.jethrodata.com*** (log in demo/demo) and ***http://jethrodata.qlik.com/*** (no login needed). These run in the cloud with 1TB of raw data, a fact table containing approximately 2.9 billion rows and six dimensions.

Jethro Data has recently partnered with Hortonworks to extend that company's Enterprise Data Warehouse solution.

Jethro
157 Columbus Avenue Suite 535
New York, NY 100235

www.jethro.io

Strengths

- The auto-cubing feature means that you do not have to define your OLAP cubes in advance. Moreover, you don't have to re-define them if they change.
- There are relatively few vendors offering OLAP on Hadoop. Most of the suppliers in this space are massively parallel implementations that focus on mixed workloads: in effect, an enterprise data warehouse on Hadoop. These products are unlikely to match the price/performance of Jethro. Of the few other products that do offer OLAP these tend not to offer the sort of indexing capabilities that Jethro does, meaning that drill-down performance will suffer. Also, cubes tend to get very large.
- The partnership with Hortonworks should help to drive adoption.

Threats

- OLAP is sort of old-fashioned and not up with the currently most hyped analytics technologies such as predictive analytics and machine learning. This can mean that doing the basic things well can get overlooked in favour of the latest flavour of the month. As a company that focuses on doing something traditional as well as it possibly can, Jethro Data is in danger of being overlooked for other technologies that may be perceived to be "sexier". We don't agree with this view, but we have to recognise that some people are swayed by what's fashionable.

Recommended for

OLAP.

Kognitio

Kognitio was founded in 1987 as a data warehousing vendor. At that time, the company was known as WhiteCross Systems. In the early 90s the company introduced what we would today call a database appliance, running on proprietary hardware. Key features are that, from the outset, the database employed a massively parallel architecture and that processing was in-memory. Indeed, given the recent hype about in-memory databases, Kognitio could reasonably claim to be the progenitor of this market.

The history of the Kognitio data warehousing product has been one of gradually moving away from its proprietary roots as the industry has caught up with its requirements. For example, in the mid-90s the company adopted standard industry chipsets, in the mid-2000s it moved to blade computing and away from its own hardware and, most recently (in 2016), the company ported Kognitio onto Hadoop (which involved changing the storage model so that it would work with Apache YARN), though the stand-alone version of the product is still available. And, moreover, the company has made the database free-to-use (with, optionally, paid-for support).

Kognitio targets complex queries against large datasets. For example, one of its clients has 10,000 Tableau dashboards running against a 9 PB database Hadoop cluster (all updating within seven seconds) and over a hundred individual analytic queries (some of which are complex). Needless to say, given the company's longevity, Kognitio has sophisticated optimiser, workload management, high availability, load balancing and so forth. It supports parallel processing for any supported language, such as R or Python, as well as SQL. Support is provided for ORC, JSON and Parquet. There is a query streaming capability that caters for situations where you don't have enough memory

From a performance perspective the company ran a series of benchmarks (validated by Enterprise Management Associates) comparing Kognitio 8.1.50, Impala 2.6.0 and Spark 2.0. In each case the same 12 node cluster was used, running the TPC-DS query set for both a single query stream and ten concurrent query streams. In both cases the data volume was set at 1TB. For the single query stream, Kognitio was the only product to complete all 99 queries and it was fastest on 92 of them (Impala fastest for six queries, Spark for one). For the ten query streams test, Kognitio was "long running" for four queries, Impala failed on more than a quarter of all the queries (mostly because it does not support the appropriate SQL syntax) and Spark did not complete fifteen of the queries. Of the 95 queries completed, Kognitio was fastest on all but eight of these. As an example of comparative performance, for single stream queries that were completed by all vendors (70), Kognitio completed in slightly more than 13 minutes, Impala took longer than 50 minutes and we gave up counting Spark when it had taken longer than Impala for just the first thirteen of the seventy queries.

Kognitio
3a Waterside Park, Cookham Road
Bracknell RG12 1RB, United Kingdom

www.kognitio.com

Strengths

- Because Kognitio has always been an in-memory database, it has been expensive. However, that is no longer the case, as memory prices have come down. Moreover, with the company adopting a free-to-use licensing model, Kognitio should be much better placed than it has been historically.

Threats

- Like many other (but not all) vendors in the enterprise data warehousing (EDW) space, Kognitio has ported its product onto Hadoop. This means that the company will be up against many of the same suppliers that it has historically competed against. Unfortunately, these vendors are all larger and better known than Kognitio and, while this shouldn't be a deciding factor, it often is.

Recommended for

Mixed workload environments that combine transactional look-ups with complex analytics.

Kyvos Insights

Kyvos Insights, now in version 4.0, is a Hadoop deployment that specialises in multi-dimensional on-line analytic processing (MOLAP) at scale. It supports both MDX and SQL and is available both on-premises and in the cloud. The product works with various business intelligence tools (in principle, any of them) and the company has partnerships with Tableau and Qlik.

The way that Kyvos Insights works is that it builds all possible dimensions for your cube and then distributes these across your Hadoop cluster. Data is pre-aggregated but this is based on generic principles rather than expected queries. We prefer this: it means that you don't have to rebuild your cubes if you start to get unexpected queries. The product supports incremental updates so that you don't need to rebuild your cubes, just update them. Similarly, the product supports slowly changing dimensions, as well as sliding windows (queries within a time frame). The engine can work in conjunction with Yarn and integrates with products such as Sentry and Ranger as well as LDAP, Kerberos and Active Directory. Active-active load balancing and high availability are provided.

From a performance perspective Kyvos has run benchmarks (based on TPC-DS) using a star schema (the product also supports other schemas, such as snowflakes) with 100 billion rows and a cardinality of 30 million. To make things more equitable the company turned off its caching capabilities, which retain data from previous queries to improve the performance of subsequent ones), so these tests were run "cold". Thirteen queries were run with 1, 10 and 50 concurrent user requests. To cut a long story short: Kyvos was typically two orders of magnitude faster than Impala for 1 and 10 requests. For 50 requests Impala timed out (more than 30 minutes) on nine of the thirteen queries. Performance was also tested against Spark – but this was Spark version 1.x – which, performed even worse than Impala.

Kyvos Insights
720 University Avenue, Suite 130
Los Gatos, CA 95032

www.kyvosinsights.com

Strengths

- Kyvos Insights specialises in OLAP. There are open source OLAP projects (Apache Kylan [MOLAP] and Apache Lens [ROLAP]) but, as far as we know, there is no commercial support for either of these engines. There are also companies that have built OLAP engines on top of Impala, but these aren't true SQL engines and, in any case, these are dependent on the performance of Impala which, as we have seen, is not impressive.
- The scalability offered by Kyvos is surprising. It has users (not all the same ones) ingesting two billion rows per day, with more than 200 billion fact rows, with more than 300 dimensions and attributes, with a cardinality of as much as 1.2 billion, and a maximum cube size of 20 terabytes. Those are impressive figures.

Threats

Providers of SQL on Hadoop engines fall into three categories: open source offerings (including those supported by HortonWorks, Cloudera and so forth), products that are essentially ported data warehouse offerings from big beasts like IBM and Oracle, and specialist providers such as Kyvos. These companies in the middle are in danger of getting squeezed on both sides, though Kyvos has the advantage that it has identified a niche that the major vendors are not targeting.

Recommended for

OLAP, especially when cubes have extreme attributes.

Bloor

Product Sheet

MapR Drill

Apache Drill is an open source SQL engine that will run on a variety of NoSQL platforms, including Hadoop. It will also run on top of document databases and natively supports JSON, so it is arguable that it is SQL+ engine, since it supports SQL queries against semi-structured and self-describing datasets (also including CSV and Parquet formats), data as well as structured data using either Apache HBase or MapR-DB tables. Its SQL support is ANSI 2003 compliant, which is more advanced than many of its open source rivals but not as up-to-date as some of its proprietary competitors. There are extensions to SQL that have been implemented to support semi-structured data.

Apache Drill is commercially supported by MapR. And, in that context, it runs on the MapR Converged Data Platform which includes MapR-FS (POSIX compliant file storage for high performance read/write), MapR-DB (a multi-model NoSQL database), and MapR Streams. This architecture has several important implications. Firstly, it effectively means that MapR can provide a "collapsed" lambda architecture. That is, it will support both batch and real-time (streaming) analytics on the same platform, in an integrated fashion. Secondly, the support for POSIX, combined with traditional analytic capabilities (not least, support for secondary indexes) means that the environment should be well-suited to hybrid transactional and analytic programming (HTAP). Thirdly, there is federation capability built in, so that you can query across from Drill to other environments.

The final result of this architectural approach is that Drill relies on the Converged Data Platform for a lot of its functionality. MapR stresses consistency, performance, resilience, scalability and so on. However, there is one downside to this, which is that there is no SQL optimiser in Drill. There is an optimiser in MapR-DB – relied on by Drill – but this is not a SQL optimiser. The implication of this is that things like join strategies may not be optimal.

For reference, it is worth noting that Apache Drill is based on Google Dremel.

MAPR

MAPR
350 Holger Way
San Jose, CA 95134

www.mapr.com

Strengths

- Of the big three providers of commercial Hadoop support, MapR has usually been – in our opinion – the most technically advanced. The company's converged data platform has been extended to support a wide variety of workloads.
- MapR Drill is the leading (though not the only) open source product in this market to be focused on HTAP.
- The support for semi-structured as well as structured (relational) data is a significant benefit.

Threats

- Like other open source SQL on Hadoop engines, MapR is facing increased competitive pressure from proprietary vendors, and the SQL optimisers that these suppliers can provide will mean that Drill will typically be out-performed when used with only structured data.

Recommended for

Mixed data analytics, HTAP, "collapsed" lambda architectures.

Pivotal HDB

Pivotal HDB is a SQL on Hadoop engine that started life as a massively parallel implementation of Pivotal Greenplum running on Hadoop. It is a part of the Pivotal Big Data Suite, which also includes the Greenplum data warehouse and the GemFire in-memory data grid. HDB has been open-sourced as Apache HAWQ and GemFire as Apache Geode.

In addition, and this is a major differentiator for Pivotal, the company also supports Apache MADlib (again, an in-house development – in conjunction with work done at several universities – which Pivotal has contributed to Apache). This provides machine learning capabilities using SQL and runs against PostgreSQL-based databases, which both Greenplum and HDB (HAWQ) are. MADlib processes are parallelised (where relevant) and run within the SQL engine. Around 40 different functions are currently available within MADlib and you can call these from R, Python and Java programs.

Major features of Pivotal HDB include dynamic pipelining, in-memory query processing, an HDFS metadata cache, automatic elastic query execution (where you spin up more resources, as required), dynamic cluster expansion/shrink capabilities when the product is deployed in the cloud, data federation capabilities that allow query processing across external environments, and the use of the Greenplum Orca query optimiser, amongst others.

The product integrates with HCatalog (for interoperability with Hive-based data), Apache Ranger (for security), and both Apache Ambari (for administration) and Apache YARN (for resource management). Pivotal is a significant contributor to the Apache Ambari project and it has extended YARN by providing "Resource "Qs" that provide multi-tenancy capabilities as well as workload prioritisation. The product's SQL support is for ANSI standard 2003.

Pivotal has not published any recent benchmark performance figures. There are some available from 2014 when it's results suggest that it was six times faster than Impala. However, we would not necessarily rely on these figures three years later. Perhaps more interesting was that Pivotal HDB was able to successfully run all TPC-DS queries even back in 2014, when none of the three competitors it benchmarked itself against could complete even a third of those queries. This ability to support all TPC-DS queries has been confirmed by third party benchmarks.

Pivotal
875 Howard Street, Fifth Floor
San Francisco, CA 94103, USA

www.pivotal.io

Strengths

- The MADlib capabilities offered by Pivotal are a major advantage if deploying Pivotal HDB for use by data scientists. While there are plenty of vendors that have – in effect – ported a proprietary data warehousing technology onto Hadoop, there are very few that can offer anything comparable to MADlib.
- Given that Greenplum was first introduced over a decade ago, we would expect HDB to have all the sorts of high availability, security, load balancing and so forth (including ACID compliance) that you would expect from a product with that sort of longevity.

Threats

- While MADlib gives Pivotal a distinct advantage when complex analytics (what we used to call data mining) are required, it is otherwise one of a number of vendors that have ported proprietary data warehousing products onto Hadoop. Nor is it the only supplier to have open-sourced its product. Its competitors (threats) are therefore the same as they have always been plus, of course, open-source developments not based on proprietary engines.

Recommended for

Mixed workload environments, machine learning.

Splice Machine

Splice Machine is a SQL on Hadoop engine that provides a "collapsed" lambda architecture. That is, it provides a single environment that combines support for traditional batch-based analytics with streaming analytics. Conventionally, a lambda architecture requires three different processing engines (batch, streaming and serving) but Splice Machine has "spliced" these together into a single environment so that, when implemented, each node in your cluster includes both Apache HBase and Apache Spark engines with underlying HDFS storage (both Apache Orc and Parquet are supported). A cost-based optimiser in each node routes queries appropriately, and this leverages the secondary indexes that Splice Machine supports. The intention is to support the complete gamut of analytic capability from transactional look-ups (the product is ACID compliant with multi-version concurrency control), through on-line analytic processing (OLAP) to complex and streaming analytics, and machine learning (the product ships with Apache MLlib). The product runs both on premises and in the cloud (Amazon S3) and, in the latter case, there are significant "database as a service" capabilities. Splice Machine supports Zeppelin notebooks and you can use traditional SQL analytic tools as well as Scala, Python, R and so on. One further significant feature is support for what the company calls "timeline tables", which effectively provide time series capabilities.

This last feature is important with respect to Splice Machine's target market, which the company describes as On-line Predictive Processing (OLPP). This can best be described as an inverted form of hybrid transactional analytic processing (HTAP) where HTAP basically does some OLAP type processing and then embeds that into operational data, while OLPP takes transactional data and embeds that into analytics – obviously, based on the name, with the primary intention of supporting processes such as predictive (and prescriptive) analytics.

SQL support is ANSI standard and there is also support for Oracle's PL/SQL extensions.

The company has not published any competitive benchmarks but there are some impressive customer performance figures.

Splice Machine
612 Howard Street, #300
San Francisco, CA 94105

www.splicemachine.com

Strengths

- There is increasing interest in both lambda and kappa architectures, but they are extremely complex to implement in a do-it-yourself fashion and require very significant skills. Being able to "collapse" these architectures into a single platform is a major advantage for Splice Machine.
- Having Spark internally within Splice Machine should provide significant performance advantages. One independent benchmark of Spark performance was that it performed eleven times better when tightly integrated with a host database rather than just connected to it. We would expect Splice Machine's approach to perform even better.

Threats

- Splice Machine has an impressive range of capabilities. However, given this breadth, there is a danger that the company could lose focus and take too much of a scatter gun approach. We therefore think that it is right for the company to focus on a particular market segment – in this case predictive processing – but it needs to retain that focus going forward.

Recommended for

"Collapsed" lambda architecture, HTAP, machine learning, predictive analytics.

Also appropriate for

OLAP.

Starburst Presto

Presto is an open source product available under an Apache license. Starburst Data – a spin-off from Teradata – provides commercial support for it and is a major contributor to the Presto project, as are companies such as Facebook, Twitter and Uber. It represents Teradata's – albeit now available via Starburst – SQL on Hadoop offering. This is a markedly different approach from rivals such as IBM, Oracle, HPE, Pivotal and so forth, all of which have ported their relational engines onto Hadoop.

Presto runs on both Cloudera and Hortonworks Hadoop distributions but not MapR (partly because MapR is promoting Drill and partly because of the money that MapR wants for certification). The architecture is based around a co-ordinator running with worker nodes, plus full in-memory capabilities and vectorised columnar processing. There are strong security capabilities, with support for LDAP and Kerberos. A major upcoming feature is that there will be a new cost-based optimiser that is the result of a collaboration between Teradata and Facebook. It has been designed specifically for Presto, as opposed to the Apache Calcite project, which is more of a generic optimiser. Another major feature that has been previously contributed by Teradata is spill-to-disk, which is designed to support query processing when you run out of memory. There are a number of other in-memory engines which grind to a halt if you run out of memory. Also planned for a future release, there will be new workload management capabilities with resource groups. There is also an intention to introduce more batch-based capabilities so that you can just run Presto and not require Presto running in conjunction with Hive.

Radiant Advisors ran some benchmarks for Teradata, the results of which were published in Q2 2016. That makes the figures rather out of date. At that time Impala was outperforming Presto more often than the other way around. We would expect, though we cannot confirm, that this situation is now reversed. What we can say is that at that time Presto was capable of running more of the TPC-DS queries than Impala. We understand that Presto will now run all TPC-DS and TPC-H benchmarks but tests run by other vendors suggest that this is still not the case for Impala.

Starburst
Boston
MA, USA

www.starburstdata.com

Strengths

- Presto forms a part of the Teradata ecosystem. In particular, it is supported by the Teradata QueryGrid, which means that queries can span not just Presto but also Teradata and Teradata Aster implementations. Moreover, Presto is provided free of charge as a part of the Teradata Appliance for Hadoop. It is also worth mentioning Teradata's Think Big acquisition, so the company has significant Hadoop implementation and consulting skills.
- Teradata – and therefore Starburst – knows what it is doing when it comes to databases that support analytics. We therefore have no qualms about features such as high availability, load balancing, workload management, query optimisation, and so on.

Threats

- Teradata (and now Starburst) has taken a different path compared to its major competitors: opting for an Apache-based open source approach, as opposed to porting a proprietary solution onto Hadoop. The downside of this is that some features of Presto, which you would expect in an enterprise data warehouse (EDW) are not there yet, though we are happy that they will be in the future. Moreover, at least some of its traditional competitors have made their products free to download and use.

Recommended for

Traditional data warehousing environments where there are mixed workloads, including both transactional look-ups and complex analytics.

About the author

PHILIP HOWARD
Research Director/Information Management

Philip started in the computer industry way back in 1973 and has variously worked as a systems analyst, programmer and salesperson, as well as in marketing and product management, for a variety of companies including GEC Marconi, GPT, Philips Data Systems, Raytheon and NCR.

After a quarter of a century of not being his own boss Philip set up his own company in 1992 and his first client was Bloor Research (then ButlerBloor), with Philip working for the company as an associate analyst. His relationship with Bloor Research has continued since that time and he is now Research Director, focused on Information Management.

Information management includes anything that refers to the management, movement, governance and storage of data, as well as access to and analysis of that data. It involves diverse technologies that include (but are not limited to) databases and data warehousing, data integration, data quality, master data management, data governance, data migration, metadata management, and data preparation and analytics.

In addition to the numerous reports Philip has written on behalf of Bloor Research, Philip also contributes regularly to ***IT-Director.com*** and ***IT-Analysis.com*** and was previously editor of both ***Application Development News*** and ***Operating System News*** on behalf of Cambridge Market Intelligence (CMI). He has also contributed to various magazines and written a number of reports published by companies such as CMI and The Financial Times. Philip speaks regularly at conferences and other events throughout Europe and North America.

Away from work, Philip's primary leisure activities are canal boats, skiing, playing Bridge (at which he is a Life Master), and dining out.

Bloor overview

Technology is enabling rapid business evolution. The opportunities are immense but if you do not adapt then you will not survive. So in the age of Mutable business Evolution is Essential to your success.

We'll show you the future and help you deliver it.

Bloor brings fresh technological thinking to help you navigate complex business situations, converting challenges into new opportunities for real growth, profitability and impact.

We provide actionable strategic insight through our innovative independent technology research, advisory and consulting services. We assist companies throughout their transformation journeys to stay relevant, bringing fresh thinking to complex business situations and turning challenges into new opportunities for real growth and profitability.

For over 25 years, Bloor has assisted companies to intelligently evolve: by embracing technology to adjust their strategies and achieve the best possible outcomes. At Bloor, we will help you challenge assumptions to consistently improve and succeed.

Copyright and disclaimer

Bloor Research International Ltd
20–22 Wenlock Road
LONDON N1 7GU
United Kingdom

Tel: **+44 (0)20 7043 9750**
Web: **www.Bloor.eu**
email: **info@Bloor.eu**