



Analytical Database Options

By David M. Raab

It has been clear for a long time that relational databases are ill-suited for analytical processing. The problem is fundamental: relational databases are designed to retrieve all elements from a few records, while analytical queries typically read all records but only a few elements. Until recently, database developers were able to overcome the problem with clever design and powerful hardware. Today, the largest databases - hundreds or thousands of terabytes - are too big for this to work at an acceptable price. Thus, the quest for alternatives is on.

But let's back up a bit. Relational databases are inefficient at analyzing all sizes of databases, not just the largest. Companies may have met their critical needs by throwing money at the problem, but they have also rejected less important applications that couldn't justify the cost. This suggests there is a significant hidden demand for analytical applications of all sizes if costs can be reduced.

This insight is important because many analytical technologies do not handle very large databases. They are sometimes dismissed for this reason. But this is a mistake: if alternative technologies are more cost-effective than conventional databases for smaller-scale projects, they still add value in those situations.

With that in mind, let's look at the major technologies available today for analytical systems.

Multidimensional databases. These systems read cubes of aggregated data. This makes them fundamentally different from the products listed in the rest of this column, which give direct access to the raw details. This is a critical distinction for analytical applications where the required queries are not known in advance. Products like BusinessObjects, Hyperion, Cognos 8 and Applix have long been the primary alternatives to relational databases for analytical purposes.

In-memory databases. These include Panoratio, PivotLink and QlikView. The technical details differ, but in general all these systems load the source data into memory in a compressed, non-relational format and query against it. Part of their value comes from the compression itself, which lets them handle more data by simply loading relational tables into memory. But compression varies greatly with both the technology and the data and is sometimes little better than one to one. At least as much value comes from the nonrelational structures, which can be more flexible and powerful for analytical queries than SQL-based systems. This makes the alternative systems cheaper to deploy because there is less fine-tuning of data structures. Still, conventional servers rarely run more than 32 or 64GB of memory. This means that, even with compression, a pure in-memory system would almost never hold more than 50 or 100GB of source data.

Simple columnar systems. These are "simple" in the

sense of running on conventional servers, not the massively parallel kind. Vendors include database marketing stalwarts Alterian and smartFOCUS. These products organize data by columns (i.e., data elements or fields) rather than rows. Because analytical queries typically read a few elements in all records, the columnar structure lets the database engine load only the information it needs. This reduces the amount of disk access and speeds the result. Like the in-memory systems, these tools have query languages that are more flexible than SQL, making them easier to work with. However, scalability is still limited, and these systems rarely exceed a terabyte of source data.

Massive columnar systems. This is the category receiving the most attention today. Vendors include Calpont, EXASOL AG, ParAccel and Vertica. Each is unique, but these systems share several basic characteristics: SQL compatibility, massively parallel "shared nothing" hardware and some flavor of a columnar data structure. These vendors generally supplement the columnar approach with other techniques, such as partitioning and storing multiple copies of the same information in different sort sequences. Some support in-memory data storage as well. These systems do scale, although it seems that most implementations are less than 10TB.

Database appliances. These are also massively parallel systems with a SQL-compatible database. Rather than a columnar approach, they focus on innovations to get the best performance from their hardware at volumes into the hundreds of terabytes. Competitors include DATAlegro, Greenplum, Kognitio and Netezza. The claim here is scalability at a much lower cost than massively parallel processor leader Teradata.

Extraction-based systems. For tasks like monitoring Web traffic, data volumes are so huge that it makes more sense to scan traffic and extract selected attributes than to store it all. Products from vendors such as Altosoft, Skytide and Visual IQ offer this capability.

Outliers. There are a few analytical systems that don't fit into any of the above categories. Infobright, illuminate Solutions and QD Technology work in the subterabyte range. All are SQL-compatible, disk-based rather than in memory, and run on standard servers. illuminate uses a very flexible database approach of its own devising. Essentially, it stores each value once and attaches index structures that show all the contexts in which that value appears. Infobright and QD Technology employ a conventional relational structure and focus on compression. Sybase IQ is a variation on a columnar database that reaches very high volumes (it cites a test with 155TB of input) without massively parallel hardware. 

David M. Raab is a Principal at Raab Associates Inc., a consultancy specializing in marketing technology and analytics. He can be reached via his blog at <http://customerexperiencematrix.blogspot.com/> or draab@raabassociates.com.

Copyright of DM Review is the property of SourceMedia, Inc. and its content may not be copied or emailed to multiple sites or posted to a listserv without the copyright holder's express written permission. However, users may print, download, or email articles for individual use.