

illuminate Solutions: a new addition to the data-warehousing crowd

Analyst: Krishna Roy

Sector: Enterprise Software

Data-warehousing vendor **illuminate Solutions** hails from Spain, where it has built up a customer base, and subsequently expanded into Latin America, the US and Europe. In the past six months or so, it opened sales and support offices in the Netherlands and France, as well as a sales office in Madrid, as part of ongoing expansion plans. The company's technology for data warehousing was funded and developed under a different incarnation, although it was subsequently rebuilt from scratch in 2005 when illuminate was established.

This technology, which illuminate dubs 'value-based storage,' is designed to optimize its database for analytics by effectively storing every value in a record only once, so that the database has a small physical footprint. The other fundamental element of the offering is the Explorer front end, which exploits the value-based nature of the database and specializes in particular on ad hoc and unpredictable queries. Illuminate also provides front-end tools for business users and Open Database Connectivity (ODBC) drivers to third-party business intelligence (BI) tools and third-party data mining tools, since the underlying database also has some basic algorithms for data mining built in as well.

The 451 Take

We've seen a number of architectures and techniques emerge to optimize a database for analytics and data warehousing, including in-memory and columnar approaches, but we've never seen illuminate's so-called 'value-based storage' approach before. This fact arguably gives it market differentiation, which we believe is vital in the warehousing market, given the number of new entrants that have emerged during the past 12 months. Illuminate also has a decent-sized customer base to act as reference accounts and underpin the claims of its technology, which should also help. That said, as history demonstrates, the advent of a new type of database, even if it's better than its predecessors, doesn't always herald automatic success. One only has to look to the object-oriented databases of the mid-1990s for proof.

Context

Illuminate was established in early 2005, although the genesis of its so-called value-based storage technology dates back to 1983, when illuminate's chief architect and cofounder, Joe Foley, started researching it. Foley brought the technology to market as a product in his previous startup, but it was a victim of the dot-com crash in 2001. Four years later, he teamed up with a partner in Spain, founder Juan Francisco García, and another founder, Anton Torrents, and illuminate was born with a product rebuilt from scratch.

Illuminate raised \$3m in funding from **Adara Venture Partners** in January 2007 to start the European rollout of its eponymous data-warehouse software, and will look for additional investment in the future, according to management. The company does both direct and reseller sales, and now claims roughly 25 direct and 11 indirect customers in Spain, and

about five direct and 20-plus indirect accounts in the Americas. None of these are customers from its original incarnation, according to management. About 90% of the 11 customers on the indirect side of its business in Spain have come via OEM **Grupo Castillo**. The indirect deals it has signed in the Americas are also mainly through resellers.

Insurance, followed by banking and retail, are the industries in which it has gained traction to date. Customer segmentation and profiling are the most frequent use case scenarios for its data-warehouse software, which is available in English and Spanish language versions.

Illuminate was reportedly profitable in 2005 but is now in growth mode, so is no longer in the black. It has, for example, opened sales and support offices in the Netherlands and France, as well as a sales office in Madrid in 2007, and is planning further investments in these regions. The company plans to achieve revenue of approximately \$1.6m for 2007 – up from \$991,000 in 2006.

Headquartered in Barcelona, illuminate also has an R&D center in Minneapolis, and sales and services operations in Latin America, where it also goes to market with partners, in addition to new aforementioned offices opened earlier this year. It has more than 30 employees in total.

Technology

The essence of illuminate's data-warehouse offering is its internally developed database that stores each value in a record only once – hence the term 'value-based storage,' which the startup has coined to describe this technology. Clearly many records contain the same value, so storing the value only once makes the physical database very compact. The original context of the data and its associated metadata is also maintained in the storage model.

Data from various sources including ERP, CRM applications, spreadsheets and other warehouses is broken down to the value level and all indexed separately and automatically during the loading process, which reportedly has comparable load rates to a relational database. The data is indexed through a number of techniques – principally bit maps, vectors, indices and B-trees – depending on the data type. ODBC drivers are also used, for example, to handle text-based data formats.

Compression is another element. The database reportedly has compression rates of about 6:1, which means the amount of data it can store is a great deal greater than it can read. The database can currently read up to 5TBs of data – although the game plan is to increase this amount in future. It also contains some internally developed algorithms for data mining.

Equally core to illuminate's offering is the query functionality dubbed Explorer that illuminate has built on top of the database to exploit its value-based nature. Explorer effectively enables queries to be asked that couldn't be asked by defining them in standard SQL, such as associative queries, partial queries, queries with no qualification and those that are unpredictable. The tool is also designed to provide fast results on any query – without the need for a query optimizer.

Products

Illuminate 3.3 is the current release, which sports a number of enhancements over its predecessor. One new feature, designed to make it easier to implement in existing environments, is modularization. Previously, the software came on one installation disk for the dashboards, the Explorer front-end and database. But with 3.3, customers can purchase these components separately. That said, management strongly recommends that customers

buy Explorer in addition to the database – even if they already have a third-party front-end BI tool from **Business Objects**, for example – since Explorer's query functionality maximizes the value-based nature of the database. The other main improvements in 3.3 are around metadata and further optimizations to the way the index is structured.

The illuminate offering can be broken down into a number of components. Illuminate Data Warehouse Server handles data storage and integration, while Explorer is the ad-hoc query and interactive analysis tool for data analysts – although it also has an API, so less-technical types can access its functionality via a third-party BI tool from **QlikTech, Cognos** or Business Objects, for example. Illuminate Dashboards and illuminate Maps are its internally developed BI tools for business users – available in client/server and Web-based releases. Illuminate Discoverer is its front-end data mining and pattern recognition tool designed to exploit the clustering algorithms for data mining built into the database. Discoverer is only meant for basic data mining. It can also connect to third-party data mining tools, including those from **SAS Institute** and **SPSS**, via ODBC drivers.

Pricing

An entry-level deployment costs around \$20,000 and is designed to be fast – typically up and running in a few weeks on average, according to management.

Competition

Illuminate reports that it is often brought into customer situations where the incumbent warehouse is **Oracle** or **IBM DB2**, and its own offering is bought for analysis of a subset of this data – and then subsequently expanded. Although this use case is a common scenario for a number of other data-warehouse players, including those that use a software/hardware combo under an appliance delivery model – **Netezza, Datallegro, Greenplum, Paracel, Dataupia** and now **Vertica Systems** – management reports that it has never competed directly against any of these vendors.

That's probably because, Dataupia aside, the rest are highly focused on query performance. In contrast, illuminate's principal focus is on query flexibility, i.e., the ability to run a number of different types of queries, many of which couldn't be written using standard SQL. Furthermore, illuminate is a software-only play and has no plans to deliver its warehouse appliance on optimized hardware, which could be another reason it hasn't come across Netezza and the rest of the data-warehouse appliance vendors in deals.

Illuminate also encounters **Microsoft's** Access and SQL Server databases in smaller customer sites as the incumbent warehouse when customers want to expand their warehouse, which makes sense given Microsoft's midmarket focus. Oracle and IBM are also pursuing the midmarket with data-warehouse appliances, which in Oracle's case are delivered in combination with partners including **Dell, EMC** and **Sun**. However, illuminate hasn't encountered IBM or Oracle's warehouse appliances.

Another common scenario is to come up against **MicroStrategy** when customers are unhappy with its inflexibility, according to illuminate. Again, that makes sense, since MicroStrategy is strong in the area of Relational Online Analytical Processing (ROLAP)-based analytics. That said, illuminate contends that MicroStrategy, alongside any other players with front-end BI tools, could be a partner. Even though illuminate has its own front-end tools for its database, the startup's primary focus is on the warehouse itself. Executives report that the most commonly used front-end tool with the illuminate warehouse is QlikTech's, although it also has customers that use Cognos, which is in the

throes of being acquired by IBM, and Business Objects, which is in the midst of being taken over by SAP.

Teradata, Sybase IQ, Infobright and Calpont, which will bring its data-warehouse offering to market in early 2008, are the other main players, none of which illuminate has competed against.

Strengths	Weaknesses
We haven't come across a so-called value-based storage approach to data warehousing from any other vendor to date, which gives illuminate market differentiation. It also has a reasonable-sized customer base, which is important for a vendor touting a new technology.	The startup needs to raise its profile since it's a relative unknown in a crowded sector. We're also unsure of how far its warehouse scales, given its sweet spot of up to 5TB of data.
Opportunities	Threats
Building on the strength of customer success in Spain and Latin America to push further into the all-important US market and Europe seems like the next logical step.	The number of different types of databases for analytics continues to expand with vendors touting in-memory, columnar and tokenized approach – and now a value-based storage approach from illuminate. But a new technology – even if works better – doesn't always flourish when up against database juggernauts Oracle, IBM and Microsoft.

About The 451 Group

The 451 Group is a technology industry analyst company focused on the business of enterprise IT innovation. The company's analysts provide critical and timely emerging-technology insight to clients at vendor, investor, services and end-user organizations – insight that aids both strategic and tactical decision making for competitive advantage.

The company's services include the 451 Market Insight Service, which delivers daily insight into emerging enterprise IT markets; 451 TechDealmaker, a weekly analysis service focused on forward-looking M&A within the enterprise IT business; 451 Special Reports, which are produced on a periodic basis to analyze key emerging enterprise IT markets in greater depth; and 451 Strategic Counsel, the company's analyst-inquiry program, which provides clients with direct access to 451 analysts. The company also produces via 451 Events periodic industry summits and investor conferences that address opportunities and obstacles facing emerging enterprise IT markets.

The 451 Group is headquartered in New York, with offices in key locations, including San Francisco, London and Boston. The company also operates Tier 1 Research – an independent division of The 451 Group, headquartered in Minneapolis – which analyzes the financial and industry implications of developments impacting public and private companies within the IT, communications and Internet sectors.

For additional information on the company or to apply for trial access to its services, go to: www.the451group.com