

KATE MITCHELL, CEO OF COPPEREYE, DISCUSSES OPPORTUNITIES FOR INNOVATION IN BUSINESS INTELLIGENCE.



For our business audience, what's the key message you would like to start our conversation with?

While it may seem a bit counter-intuitive, business leaders are the drivers behind many of the big changes in the technology landscape. They become dissatisfied with the status quo and demand something different from the technology establishment. While technology professionals are frequently accused of wanting the latest and greatest just because it's new, I believe the opposite is true – when pushed, information technology professionals will go with what they know rather than look to innovative solutions they perceive as higher risk, which in reality are simply unfamiliar to them. So it is critically important for business leaders to keep their eyes open for new and innovative technologies that can solve real business problems, and then get their technology teams to investigate and validate the new approach.

So that's why you're taking your argument for a new category of data management solution to the business person?

Exactly right. There's the old joke about "Why did the application developer put all the data in a database?" with the answer being "Because it was there." And there is way too much truth in this for comfort, particularly if you're the CFO writing big checks for database licenses and hardware simply because a database was the most convenient choice for solving a problem. It's important to note we are not talking about subtleties of technical architectures – but bold, common-sense things like "Why do I need a big, expensive database and the hardware to support it and the database administrators to manage it, for all this data that won't ever change and most of which I'll never need?"

What do you see as the biggest challenges of business intelligence?

I recently had an opportunity to present at a business intelligence conference in London, and was genuinely surprised to hear almost every speaker, customers as well as solution providers, acknowledge how business intelligence solutions were struggling to meet the expectations of business users. I believe a core challenge for business intelligence is grappling with the enormous volume of data being generated by business operations, and the need to scale to handle this increased volume at the same time the business is demanding more, and faster, access to the underlying data and the business intelligence generated from it. And while there's a lot of excitement around unstructured data such as email, documents and multimedia files, for business intelligence purposes we're really talking about structured transaction data such as point-of-sale transactions in retail, phone call records in telecommunications, or credit card transactions in the consumer finance markets.

So where's the opportunity for innovation?

Business needs are driving the requirement to eliminate overnight data loads so business intelligence can contribute to better decision making throughout the business day. This is key. This same requirement also demands access to detailed transaction data rather than simply providing summary reports of last week's or last month's operations. So this is where you will see a lot of attention in the business intelligence area over the next 18 months or so – getting inside the daily decision cycle and getting fast access to more detailed information to make better decisions.

You've been advocating a new category of data management solution for the past year or so – how does that relate to business intelligence?

There's an opportunity for a new category of data management solution specifically focused on the needs of business data that does not change after it is created. By calling out data that meets this definition as a new category, and managing it differently from other business data that can change, it's possible to dramatically reduce the cost of capturing, storing and quickly accessing this data. Most companies will tell you the volume of business data is more than doubling every year, requiring significant investment simply to keep up with the increased volumes. This makes it that much more challenging to meet new requirements for more timely business intelligence or provide access to more detailed information. To date the solution has been to buy ever-larger hardware and more database licenses which consume precious capital that could be better spent on innovation that drives business value.

With your background as an Oracle marketing vice president, it seems you'd be the last person to be taking a poke at the relational database?

Don't misunderstand – I have been and continue to be a big fan of relational databases. The relational database has a well-earned place in the architecture of virtually every enterprise. However, the characteristics of the data being generated by companies has changed dramatically in the last twenty years. For instance, all of the network systems like cellular phones, computerized point-of-sale systems, e-commerce, automated business processes and countless other applications of computer technology generate a far different mixture of data types than in the past. This growing volume of structured, unchanging business data does not need the power and sophistication of the relational database. By storing this large volume of data outside the relational database, the database is unburdened and becomes dramatically more efficient, delivers higher performance to end users, and requires significantly less hardware resources to achieve this improved performance.

I've heard you refer to "business event data" a couple of times – is this the new category you're advocating and if so, what makes it different?

Business event data is the name I've proposed for business data that cannot change once it is created. Again, this is not a subtle point but one that makes good sense once you think about it. Business Event Data are structured, unchanging, records of business events such as banking transactions, stock trades telephone calls or retail purchases. A point-of-sale retail transaction is a fact that can never change once it is created, even if there is a subsequent refund which creates another transaction. These "write-once" records are typically characterized by high transaction rates and long retention requirements. This results in databases of hundreds of terabytes that are quickly becoming slow to access and expensive to maintain using conventional means.

In your view, what makes dealing with enormous data volumes such a challenge?

There are three principal areas of compromise every data management solution must deal with, each with a significant impact on the business. First is latency, which is how soon after data is created it is available for use; shorter latency is more demanding, and longer latency is much easier. The second is the volume of data that is managed, with larger volumes being more difficult, as you would expect. And last is query performance, which is how quickly you can get data back once you submit a query, and there has been a lot of attention in this area recently. Taken individually, each of these three characteristics is desired – low latency, lots of stored data, and quick query response. The challenge is in the real world you can really only get two of these characteristics for a reasonable investment, with all three becoming exceedingly expensive. We need look no further than production applications in most companies that keep only six months or less of transactions online (they've compromised volume to get low latency and fast query response), or more appropriate to the business intelligence audience is a data warehouse application that relies on an overnight batch load to get data into the system (accepting high latency in the interest of large data volume and fast query response).

So is the argument for this new category all about saving money?

Yes and no. Yes, it's about saving money compared with using a relational database for all this business event data that is currently bogging down production databases. But no, because it's also about innovation and finally being able to do great new things that weren't ever considered before, because a database would have been completely unaffordable.

Big volumes of business data are nothing new – why is something different needed now?

Big volumes of business data may not be new – although I'd argue that they're a lot bigger now than anyone could have ever anticipated – but the thing that is really changing is the nature of what companies do with all that data. Five or ten years ago, a company with a data warehouse that was used to make quarterly planning decisions was considered a "data driven enterprise". Today, that level of solution is nothing more than a good starting point. Leading companies have moved their data warehouse and business intelligence solutions into mainstream support of the business as mission-critical systems. Business intelligence applications are also no longer the exclusive territory of big companies, but are also present in smaller firms without the big budgets or information technology staff required to run multi-terabyte database implementations. It all comes back to my argument in favor of a new category of data – business event data. Today's data warehouses and business intelligence solutions simply were never designed for dealing with the enormous volumes of unchangeable, structured transaction data currently being generated by business operations.

Thanks for answering our questions Kate – any parting comments?

I enjoyed the opportunity. I hope the readers have found my comments valuable. At a minimum, next time someone proposes using a database to solve a particular business problem they will pause and ask if a database is really the right solution, or just the most convenient choice.

About CopperEye

CopperEye is a provider of enterprise search software that allows companies to quickly retrieve exactly the records they need from months or years of history and billions of business transactions stored on low-cost file systems. Unlike other enterprise search solutions that focus on unstructured data such as web pages and word processing documents and return ambiguous results, CopperEye specifically addresses the needs of business transaction data that would otherwise require a costly and complex database solution to provide rapid retrieval of specific records. Using CopperEye, companies are able to implement more powerful applications that have dramatically enhanced access to historical data, while simultaneously realizing an order-of-magnitude reduction in data management infrastructure costs. CopperEye is an enterprise-scale solution that can be running in a full production environment in a matter of weeks. CopperEye has offices in Bristol, U.K.; Stamford, Conn. and San Mateo, Calif., U.S.A.

For further information, please visit <http://www.coppereye.com>.

Corporate Headquarters:

Suite 47, Aztec Centre
Almondsbury
Bristol, BS32 4TD
Tel. +44 (0) 1454 203610
Fax. +44 (0) 1454 203330

US Headquarters

CopperEye
101 Federal Street
Suite 1900, Boston
MA 02110, USA
Tel. +1-617-342-7173
Fax. +1-617-342-7080

www.coppereye.com