

Tokutek delivers ACID transaction support with TokuDB version 3.0

Analyst: Matt Aslett

Event summary

- Tokutek has released version 3.0 of its TokuDB for the MySQL storage engine, including support for ACID transactions.
- TokuDB is based on Fractal Tree indexing technology that reduces data-insertion times, improving the performance of MySQL for both read and write applications.
- The company is also working with two other database vendors that are experimenting with integrating the Fractal Tree indexing technology with their own database products.

The 451 take

The apparent ability of Fractal Tree indexing to enable querying of live operational data by improving data-insertion performance provides Tokutek with a truly differentiated story to take to market, one that would appear to be perfectly matched with its chosen sector of advanced Web applications. Targeting MySQL users makes a great deal of sense and we are also intrigued by the potential to partner with existing database providers to bring Fractal Tree indexing to a wider audience without cannibalizing Tokutek's direct business. Kayak.com is a significant customer win, although we would like to see a few more of that stature, especially now that the company can claim ACID-compliant transaction support.

Details

Tokutek was established in 2006 to build a business around the research undertaken by its founders into the potential of Fractal Tree indexing, a method of indexing data designed to overcome disk I/O limitations of traditional B-tree indexes. The company delivered the first generally available version of TokuDB (2.0) in May 2009 having decided that the best initial route to market for the technology was through a storage engine for the open source MySQL database. TokuDB for MySQL is targeted at advanced Web applications such as e-commerce, advertising, marketing, vertical search and social networking, and is designed to provide a way of overcoming MySQL's performance ceiling.

By improving data-insertion rates, Tokutek argues that it also disproves the theory that databases cannot perform transactional and analytical processes simultaneously. Despite the performance hit taken in providing support for ACID transactions, Tokutek claims a 10 times performance improvement compared to the MyISAM and InnoDB storage engines (it had claimed 10-50x for version 2.0 without ACID support). At that rate, the company

maintains, it becomes practical to run ad hoc queries against live data. Having delivered ACID support, Tokutek is now turning its attention to improving performance further with a focus on support for multithreading beyond four cores.

The company currently has nine employees, down from 12 in May 2009, and has less than half a dozen paying customers, up from one in May 2009. Customers include travel search site **Kayak.com**, which is using TokuDB for high-speed analysis of visitor traffic, and **Profile Technology**, which offers an advanced search capability for **Facebook** and reported a 58x insertion rate improvement by moving to TokuDB. Tokutek raised an undisclosed amount of seed funding from **General Catalyst Partners** and **Atlas Venture** in 2008 and an angel round of debt funding in 2009, and is investigating a first institutional funding round.

Competitive landscape

Tokutek's most direct competition is users trying to scale their MySQL Web applications themselves using techniques such as sharding, caching and partitioning. As such, other vendors that improve the performance of MySQL, including **ScaleDB** and **Codefutures**, must be considered indirectly competitive. While TokuDB is effectively an operational database technology, it does blur the lines between operations and analytics, and the company has also come up against **Netezza** for data-warehouse-centric applications and **Infobright** for data mining (Infobright also provides a storage engine for MySQL that improves its query performance for analytical workloads).

Given its focus on MySQL users, Tokutek does not see a lot of **Oracle** or **Microsoft SQL Server**, although they are clearly the long-term targets. Its technology might be more likely to vie with them via partners, however. Tokutek reports that two commercial database vendors are currently experimenting with integrating the Fractal Tree indexing technology with their own database products. Partnering with a larger database provider is an attractive option since it could boost Tokutek's revenue without the company having to go head to head with the likes of Oracle and Microsoft.

Reproduced by permission of The 451 Group, copyright 2010. This report was originally published within The 451 Group's Market Insight Service.

For additional information on The 451 Group or to apply for trial access, go to:
www.the451group.com