



Exclusive To: *LINUXINSIDER*

Open Source Databases: A Round-up

By Matt Benjamin

CTO and Co-Founder of The Linux Box

When the Internet and World Wide Web finally went mainstream in the mid-1990s, new Unix-like operating systems running on PCs went mainstream too.

Developers working on these systems had access to a wide range of development tools, but these platforms didn't initially offer anything like the mature, SQL-based and 4GL-based database frameworks common on traditional platforms of the time; commercial database vendors were in no hurry to remedy this situation. Ironically, several commercial RDBMS products were originally based on open-source code developed by Michael Stonebraker's Ingres project at Berkeley.

Open source databases were the inevitable result. Early systems like mSQL were primitive, but suitable for entry-level applications. They were also free for people to use and develop. A memory of mSQL survives today in MySQL, whose C-language API was based closely on that of mSQL to ease migration. The MySQL system was started by Swedish developers TCX, as an ISAM data store supporting business applications. Reborn as a multi-user SQL RDBMS with ISAM under the hood, MySQL would become nearly synonymous with open source databases for several years. Around the same time, intrepid developers were creating other fascinating systems, including Postgres95, the precursor to Postgresql, based on yet another research database from Stonebraker.

In 2010, open source databases have come of age and it can be a challenge to keep track of developments in the space. Client/server systems with SQL interfaces jockey for position against upstart NoSQL systems with intimidating (and exciting) new models for data representation, distribution, and consistency. In addition, more than a dozen embedded and special purpose databases have grown up to serve the needs of applications too small or too agile to require a full RDBMS. We'll survey some of these systems in their different categories.

Client/Server RDBMS (SQL)

The structured query language and its rich client/server support environment remain the industry standard for management of structured data. Open source is very well represented in this category. The best known open source RDBMS is of course MySQL, with many thousands of developers and applications relying on it since its appearance in the mid-1990s. While fully GPL-licensed, MySQL was primarily developed by its strong, core team of developers. Since the acquisition of MySQL AB by Oracle, a more significant development role has been played by several forks, such as MariaDB, a fork started by MySQL's original developer Monty Widenius. MySQL is distinctive in several ways in the current landscape. Initially noted for ease of use and integration with virtually every web application development framework, MySQL

has over the years become richer, more flexible and conformant to standards. Currently, MySQL supports several data management back ends (or "storage engines"), including the original MyISAM database, a newer InnoDB database with enhanced referential integrity and SQL procedures support, a clustered database framework, MySQL Cluster, based on both InnoDB and a distributed shared memory data store called NDB, as well as a federated database framework for operating on data distributed across an arbitrary number of database servers. The MySQL replication and clustering options are thought of as among the most mature and proven in the open source space.

Two other client/server RDBMS systems deserve specific mention. The Postgresql database is a very strong product with several distinguishing features. Postgresql is particularly valued for its strong referential integrity and virtual table (view) support, and especially for its rich procedure language support. Postgresql procedures can be authored in several general purpose languages (e.g., Perl, TCL), and also in a (partly) Oracle-compatible language - PL/PGSQL. The current 8.x release greatly simplifies management and tuning and the upcoming 9.x release introduces new XML-data management features. Postgresql is developed by a vibrant open-source community and also enjoys strong commercial support from several Postgresql-focused companies. A new clustered Postgresql technology contributed by EnterpriseDB promises significant increases in scalability and redundancy and an early preview code is already available. Finally, we should mention Ingres, which, after many years as a proprietary descendant of Berkeley Ingres, is now developed by Ingres Corp. as "Commercial Open Source." Ingres is most popular with organizations and developers familiar with original, closed-source Ingres. Nevertheless, Ingres is a rich product with a long history and rich feature set.

Compact/Embedded Databases

The flexibility and potentially low cost of open source/free software has ensured a prominent role for it in databases designed for bundling with larger software systems, as well as embedding in small form-factor systems such as electronic and mobile devices. The Berkeley Database (BerkeleyDB) from SleepyCat (now owned by Oracle) has successfully targeted this category for over a decade. BerkeleyDB is known for both compactness and performance and provides API's for C, C++, and Java applications, with bindings for a myriad of other languages. BerkeleyDB does not use a query language. More recently, the SQLite database has exploded onto the scene, offering the familiarity of SQL while at the same time providing excellent compactness and performance. SQLite is incorporated in thousands of free and commercial products thanks to both its features and its highly non-restrictive license. These are by no means the only open source databases designed for embedding. Firebird, for example, is a descendant of the Borland Interbase product and has a strong following. As well, there is an entire subcategory of databases implemented in pure Java and intended for embedding in Java applications. Products worth noting include Apache Derby, Hypersonic SQL, and McKoi SQL. These cover a range of requirements, from ultimate simplicity and compactness to full client/server capability.

New Architectures (NoSQL)

The open source database landscape has never been static. Some of the most talked-about databases don't use SQL and many aren't relational at all. The term "NoSQL" is used to describe a number of new databases that have recently become available that are not uniformly similar, but don't fit the RDBMS mold either. A number of the new databases are replicated while others are distributed. Among the latter group there are several such as Cassandra, Dynamite, and Voldemort, which was inspired by Amazon's Dynamo database. Replicated databases can increase durability as well as performance. CouchDB exemplifies a replicated, NoSQL database, with a rich but flexible data model. The more recent MongoDB has schema

flexibility comparable to CouchDB, but has a more traditional client-access protocol and implementation. The Dynano-like systems draw on peer-to-peer technology and relaxed consistency to achieve greatly increased transaction rates, while still replicating data for availability. Some of these systems (Cassandra, Voldemort, Dynamite) trade transactional semantics for performance. Other systems, like Scalarix, support transactional updates. Other systems use radically simplified data models and many are key-value stores. NoSQL databases are an important development. Their greater flexibility and simplicity can speed development and intrinsic distribution contribute to both availability and performance. That said, NoSQL is not for everyone. The power and sophistication of SQL databases are the foundation for a wide range of reporting and data access technologies, from application frameworks to report generation systems. It will be some time before NoSQL equivalents for such integration technologies can be put in place.

(Bio: Matt Benjamin, CTO and co-founder of The Linux Box, has a lengthy track record working with software programs, Unix computing systems and the web. The Linux Box Corporation is a software development organization specializing in open source technology and the Linux platform.)