Flying under the radar

In a SQL world, post-relational databases are holding their own

By Charles Barouch
IN NEWPORT BEACH, CALIF.

In the modern IT world, SQL and XML have become such key components in the database sector that it's difficult to remember a time when they weren't part of the architectural story.

However, in the late 1960s and early '70s, database architectures with flexible field sizes, nested tables and loose-data typing options were winning market share and mind share. While these earlier systems lacked the tagging that in many ways defines XML as XML, they behaved in a remarkably similar manner in all other respects.

Because these systems—the current terms to describe them are "post-relational" or "multivalue"—don't get a lot of press coverage, it is easy to write them off as one of the many noble experiments that have fallen by the wayside. The IT industry is littered with better mousetraps that eventually disappear.

However, post-relational technologies haven't fallen. There are many companies that still offer such database products, from such tech-sector mainstays as IBM to lesser-known companies such as InterSystems, Northgate and jBase. At the International Spectrum Conference here in late March, seven of the major database vendors in this arena presented their wares, touted new partnerships and detailed plans for expansion.

A quick look through the conference agenda reveals the same topics that can be seen at a relational database conference: how to develop robust Web-to-data integration, change control management, security, document management and all the other usual subjects.

In addition, just like their better-known counterparts in the SQL world, these databases have strongly partisan supporters.

In reviewing the information at the event and talking with participants at the conference, here's the summary of the partisan case for why a business might consider post-relational databases in addition to—or instead of—RDBMS (relational database management system) technology:

Attendees view IBM's display at the International Spectrum Conference.

Scale

SQL scales exceptionally well when scale means increasing the number of users without loss of speed. Post-relational systems scale exceptionally well when scale means increasing the complexity of the application.

The secret is in the data structure. Due to the XML-style nesting, data integrity is inherent in the model. This requires fewer computational resources to check and protect the completeness of the data. Additionally, nesting allows for dramatically fewer reads to achieve the equivalent volume of data retrieval. That means fewer read cycles and therefore longer MTBF (mean time between failures) relative to the volume of business. It also means that a programmer or analyst can see the primary relationships by looking directly at the data, without actually needing to see the schema.

IBM has documented excellent throughput with thousands of active users, so even user scaling works well in these environments.

Total cost of ownership

Many post-relational database environments in small and midsize businesses have a staff of one programmer. They tend to have one or two consultants who work less than 40 hours a month for them. For bigger companies, the staffing range usually caps out at one-third of the staff required for RDBMS, even when actively re-engineering.

Applications in the post-relational market are generally less expensive for equivalent functions. That, combined with smaller staff requirements to manage and update the applications once they are purchased,

mean post-relational systems offer real relief for businesses with financial constraints.

DBA

Other than occasional performance tuning or manual file sizing, post-relational systems have rarely needed a database administrator. A typical system can run for more than 10 years with no one performing traditional administration. While some SQL systems are just starting to move away from DBA as a full-time job category, these systems have been DBA-less from inception.

Modernization

With such technologies as encryption for data that is at rest—as opposed to being transmitted—and interfaces for .Net, Java and secure sockets, relational database vendors are continuing to keep their products up-to-date. For example, Revelation Software started offering inherent support for GUIs in 1994. OpenQM offers object programming support as part of the core product. Reality—currently owned by Northgate, previously owned by McDonnell Douglas—offers integrated Web tools.

Most of the players in this category work on Windows, Linux and Unix, so they can be integrated into whichever operating system strategy a company may follow.

Post-relational databases

The top vendors in this space include:

- **IBM** Big Blue's U2 Client Tools leverage the extended relational model supported by its UniData and UniVerse products
- **InterSystems** Its Caché high-performance object database runs SQL five times faster than relational databases
- **Northgate** Its Reality database offers a multitude of integrated interfaces
- **jBase** Its namesake offering is an open database product
- **Revelation Software** Its OpenInsight offers Windows-, Linux- and Java-based GUI tools
- **OpenQM** The open-source database got its start in 1993 embedded in software from Ladybridge Systems
- **TigerLogic** Its Pick databases include mvEnterprise and mvBase

Source: eWEEK reporting

Best of both worlds

All the major players offer extensive SQL support out of the box. This allows for use of both 1NF and NF2 data within the same database. The ability to support diverse data plays well for mashups, data warehousing and other aggregation processes.

For example, UniData and jBase offer full SQL compliance as an option.

Tools

While there are many more companies offering tools for RDBMS than for post-relational databases, all the RDBMS tools that require ADO (ActiveX Data Objects) or ODBC (Open Database Connectivity) capabilities can be used in both worlds. The tools specific to the post-relational products tend to be less expensive. With a strong variety of reporting options, XML conversion tools, 4GL programming language environments, IDEs (integrated development environments) and other core functionalities available with the traditional tools, post-relational vendors offer flexibility in supporting dynamic business models.

Reliability and breadth

Banks, colleges, insurance companies, hospitals and other traditionally conservative institutions routinely use post-relational technology. The technology can be found in everything from libraries to lumber yards.

Training

The learning curve for these systems is relatively small. Given the programming overlap between the compiled language embedded in the systems and newer languages, such as PHP, some business executives have said they can get a new employee up to a productive level in a few weeks.

A casual poll of attendees at the International Spectrum Conference indicated that many of them had extensive business knowledge but little IT training when they started. Many credited the easier syntax of post-relational queries as an asset in shortening the learning curve.

Query

Post-relational systems push their equivalent of the Join, called a Translate, into the schema, making it easier to write queries and more difficult to incorrectly associate data due to poor understanding of the schema. While the tools available to work with SQL are far more extensive and varied, the underlying query language in post-relational systems is easier to master, attendees said.

Charles Barouch has spent more than 25 years as a business and computer consultant. He splits his time between Key Ally, a consulting organization, and the New York Botanical Garden, where he does Web and data integration work. He can be reached at Results@KeyAlly.com.

Database news

Brian Prince keeps an eye on all the trends and breaking news in the database field.
At eweek.com