

How do you Stack Up?



Get business results with high quality data

Get the White Paper »

TRILLIUM SOFTWARE



information management

How Your Business Works

E-Mail Address:

Password:

LOGIN

Forgot your password?

HOME | NEWSLETTERS | eBooks | EVENTS | WHITE PAPERS | VENDOR SHOWCASE | RSS/XML

BLOGS | DM RADIO | WEB SEMINARS | SLIDE SHOW | INFO-MGMT TV | CLOUD COMPUTING EXCHANGE



Register Now for FREE Site Access and More

Enter Email Address

REGISTER

Will Triple Stores Replace Relational Databases?

Google, Microsoft and Yahoo! have found success with this flexible technology, though its short-term impact on RDBs is debatable

Information Management New letters, April 18, 2011

Jans Aasman

People ask me all the time, "Will triple stores replace relational databases in three or five years?" and I usually give two answers:

Like what you see? Click here to sign up for Information Management's daily newsletter to get the latest news, trends, commentary and more.

Answer 1: Yes, because triple stores provide 100 times more flexibility. For example, triple stores make it so much easier to add new predicates (think columns in relational databases) and write complicated ad hoc queries or perform inferencing and rule processing. Triple stores will soon be as robust, user-friendly and manageable as relational databases. Relational databases may continue to perform a bit better on simple joins, but triple stores already produce better performance when it comes to complicated queries, rule handling and inferencing. Given this robustness and usability – if the speed is roughly the same – many people will make the choice to switch to the more flexible solution.

Answer 2: No, triple stores will continue to be used in conjunction with relational databases for the near future. Many installed legacy systems took millions of dollars to implement, and it's impractical to replace these systems in the near term. In these cases, triple stores can enable smart integration of databases by adding intelligent metadata on top of databases. Many companies are already using triple stores as a "smart brain" on top of their legacy systems.

ADVERTISEMENT

Integrating Data: An Open Source Approach

Any technology deployed in the information system needs to interoperate with existing applications or databases. Learn through real-life scenarios how Open Source data integration solves the interoperability challenge. [Free White Paper.](#)

How Do Triple Stores Differ from RDBs?

In a recent conversation with a leading industry research firm, I explained how relational databases differ from triple stores by first showing a set of tables that describe a person with multiple properties and set of link tables to spouses, schools, professions and children.

I then showed the same information as a flat table of triples and explained that you:

1. Don't need to make schemas beforehand.
2. Don't need to link tables because you can do one-to-many relationships directly.
3. Can add new data attributes (predicates) on the fly that will be instantly available for querying because everything is automatically indexed.

The analyst stopped me and said, "Hold on – I've heard this argument before. This is just a completely denormalized database with just one table. There were all these database articles about knowledge databases in the seventies, but they never went anywhere. How is this different?"

Standards and Adoption

To make a leap forward in technology, standards are often a critical factor. There is a standard for the triple approach (RDF, RDFS, OWL, SPARQL, etc). If you visit the [W3C website](#), you can see that in just a few years it has become the most important development next to HTML5. The big initial success of SQL databases was that there was a standard. It allowed companies to train new people very easily on new technology, and the standard makes it easy to escape from one database company to the other. Competition between companies focused on features and performance but not on the basic access and storage technology.

Unstructured Data is Valuable

Businesspeople are realizing that in many cases, there is far more knowledge in unstructured data than in

ADVERTISEMENT

download your FREE White Paper #104 today!

Click here



MOST READ | MOST EMAILED

1. When Business Intelligence Pays Off
2. Customers In The Cloud: How Safe?
3. Interface: The Social Media Wave is Here
4. Three Types of Executives in Performance Management
5. Overall IT Job Market Improves

Information Management BuzzBox

What are people saying at SEMTECH 2011?



WorkingOntology #semtech a bit more about TopBraid support for #microdata <http://bit.ly/mIOYpk>
24 minutes ago · reply · retweet · favorite



Jenz514 learned at #semtech from @ivanherman that work is underway to reduce the time it takes to get wikipedia changes reflected in dbpedia
24 minutes ago · reply · retweet · favorite



CaptSolo what was the pwd for #semtech handouts?
25 minutes ago · reply · retweet · favorite



ConsultRamy Meet the #pragmatestech Team at #Semtech - @eddy2ma @sabawalid @johnnyazzi - booth 302 and try our live demo of the CTRL #semantic engine!

Join the conversation

MARKETPLACE

Find related products & services on Information Management Marketplace

CHANNEL RESOURCES

Articles from this Site

- CEO to Congress: Meaningful Use Helps EHR, Hampers IT Innovation
- Life Insurance Payouts, Silos and Data Management
- Social CRM's a Tough, Worthy Goal
- Data Management Gets Big
- Final PHI Protection Rule Won't Mandate Encryption

structured information. (Think emails, documents, drawings, spreadsheets, etc.) Everyone will agree that relational databases are not the perfect tool for unstructured data. For a while, people thought that XML might be the solution, but XML is not a flexible self-describing language that easily deals with graphs. (To some extent it can self-describe through XML schema and deal with graphs by using the ID tag but is cumbersome.) Most people now agree that RDF is a much better structure for that than XML.

Need for Standardized Metadata

Metadata has become increasingly important now that unstructured data is finally recognized to be as valuable as structured data. There are now metadata standards using RDF triples and companies are using triple stores in production not as a replacement of relational databases but as a complementary technology.

In his book, "Pull: The Power of The Semantic Web to Transform your Business," David Siegel explains how standardized metadata and meta tagging will change the business world. I'm not going to defend his book here, but he sketches in great detail the importance of metadata for taxes, the SEC, the world of health care, education, etc. If you read this book, you will understand why RDF (and thus triples) is the new language for metadata.

Real applications are using triple stores today. Most of these applications are programmed with a triplestore that is running next to a bunch of relational databases. At some point, triple stores will completely replace relational databases. But today they can also run in parallel and in cooperation.

For the longest time Google refused to use the word "semantics" in any of their communications. Their position was that keywords are enough to do whatever you need to do for information retrieval. However, Google bought Metaweb, which offers the biggest RDF-based encyclopedia on the Web and formally announced that it will use RDF embedded in Web pages to enhance the presentation of products and companies and events in their query results. Both Microsoft and Yahoo! have taken the same step, and we are now seeing the first success stories of companies that use this technology.

I'll leave you with one final thought: What if you need to buy a new car? The obvious choice is to buy another car just like the one you have in the garage. The type of car that has been around for 30 years, works great on the road, perfect mileage/fuel consumption and every auto shop knows how to repair it. But there is also a new car on the market that is just as stable and fast, similar cost, but much more flexible because it also can fly and run underwater. Which car would you buy?

Dr. Jans Aasman, CEO of Franz Inc., the leading supplier of Graph Database technology for the Semantic Web.

For more information on related topics, visit the following channels:

- [Analytics](#)
- [Data Integration](#)
- [Data Management](#)
- [Databases](#)
- [ETL](#)
- [Information Management](#)
- [Metadata Management](#)
- [Storage](#)
- [Unstructured Data](#)
- [Manufacturing/High Tech](#)
- [Retail/CPG/Logistics](#)

[More Articles from this Site](#)

White Papers

- [Selection Criteria for Database Archiving Products](#)
 - [Application consolidation and decommissioning projects: strategies that deliver ROI](#)
 - [Gaining the Performance Edge Using a Column-Oriented Database Management System](#)
 - [High Availability](#)
 - [Advanced Compression](#)
- [More White Papers](#)

Web Seminars

- [The Case for Database Conversion](#)
 - [Modeling Unstructured Data](#)
- [More Web Seminars](#)

Books

- [Oracle PL/SQL: Expert Techniques For Developers and Database Administrators](#)
 - [Information Modeling and Relational Databases](#)
 - [Implementing Enterprise Data Warehousing](#)
 - [SQL Queries for Mere Mortals](#)
 - [Physical Database Design](#)
- [More Books](#)

ADVERTISEMENT

Marketplace Message Center