

Franz's AllegroGraph® Sets New Record – 1 Trillion RDF Triples

AllegroGraph is the first NoSQL database to achieve the loading of over 1 Trillion RDF Triples – A major step forward in scalability for the Semantic Web.

OAKLAND, Calif. – August 16, 2011 – Franz Inc., a leading supplier of Graph Database technology, with critical support from Stillwater SuperComputing Inc. and Intel, today announced it has achieved its goal of being the first to load and query a NoSQL database with a trillion RDF statements. RDF (also known as triples or quads), the cornerstone of the Semantic Web, provides a more flexible way to represent data than relational database and is at the heart of the W3C push for the Semantic Web.

A trillion RDF Statements eclipses the current state of the art for the Semantic Web data management but is a primary interest for companies like Amdocs that use triples to represent real-time knowledge about telecom customers. Per-customer, Amdocs uses about 4,000 triples, so a large telecom like China Mobile would easily need 2 trillion triples to have detailed knowledge about each single customer.

Bill Guinn, CTO of Amdocs Product Enabler Group recently said, "We run the biggest Oracle installations in the world for telecom companies and we will keep using them in the foreseeable future. However, we just couldn't make our intelligent Amdocs Intelligent Decision Automation work in a relational database. We needed the flexibility of a triple store to make this new personalized predictive CRM application work. So now we have an application that works with literally tens of other relational databases and unstructured data

sources, which turns information into triples so we can raise the bar on customer experience.”

Dr. Aasman said, “Some people have asked, “Why not do this on a distributed cloud system with Hadoop?” The quick answer: NoSQL databases like Hadoop and Cassandra fail on joins. Big Enterprise, big web companies and big government intelligence organizations are all looking into big data to work with massive amounts of semi-unstructured data. They are finding that NoSQL databases are wonderful if one needs access to a single object in an ocean of billions of objects, however, they also find that the current NoSQL databases fall short if you need to run graph database operations that require many complicated joins. A typical example would be performing a social network analysis query on a large telecom call detail record database.”

“Relational databases are exactly the right way to deal with classic business structures and relationships – sales catalog items, warehouse inventory, shipping logistics, customers, orders, etc... And I’m not suggesting here that that’s likely to change. But relational databases are very difficult to press into service for ‘fuzzier’ applications, such as characterizing the linkages in a social network or predicting why a caller on a customer service line might be calling before the service representative picks up the phone.” stated Mitch Shults, Mission-Critical Segment Strategist, Intel Data Center Group, in a recent blog posting.

“Graph database technologies are delivering value to the Enterprise and we are seeing accelerating adoption,” added Dr. Aasman. “Our joint efforts with Stillwater and Intel to achieve this Industry record offers an unbeatable combination of software and hardware to further corporate adoption of scalable Graph database technologies, like what we see at Amdocs.”

“We found a great partner in Franz Inc. Our experience in

building high performance systems optimized for big data/deep analytics enabled Stillwater Supercomputing to help realize the industry's first Trillion Triple database. Stillwater's acceleration technology for Knowledge Processing will enable AllegroGraph to scale from the cloud, to plug-and-play hardware appliances, to tablets and smart phones and deliver the next generation of intelligent applications to our customers." said Dr. Theodore Omtzigt, Ph.D., CEO and Founder, Stillwater SuperComputing Inc.

"The history of computing is basically the story of finding creative ways to burn ever-less-expensive compute time in order to save ever-more-expensive programmer time. Semantic database techniques are the latest in a long list of innovations to advance generality and flexibility that are only practically possible thanks to Moore's Law. At Intel, our job is to keep Moore's Law on track so those innovations can keep happening." said Mr. Shults.

Dr. Aasman will be discussing ["When Relational Technology Can't do the Job"](#) as well as ["MongoGraph – Mongo DB Meets the Semantic Web"](#) on August 24th, at the 2011 NoSQLNow Conference in San Jose, California

About Franz Inc.

Franz's semantic technology solutions help bring [Web 3.0](#) ideas to reality. The company is the leading supplier of commercial, persistent and scalable Graph Database products. AllegroGraph is a high-performance database capable of storing and querying billions of RDF statements. The product provides solutions for customers to combine unstructured and structured data using W3C standard RDF for creating new Web 3.0 applications as well as identifying new opportunities for Business Intelligence in the Enterprise. AllegroGraph's Activity Recognition package provides a powerful means to aggregate and analyze data about individual and organizational behaviors, preferences, relationships, plus spatial and temporal linkages between

individuals and groups. Franz customers include Fortune 500 companies in the government, life sciences and telecommunications industries. For more information, visit franz.com.

All trademarks and registered trademarks in this document are the properties of their respective owners.