

What One Trillion Means for the Semantic Web

Mitchell Shults commented on the significance of Franz's recent success loading one trillion triplestores. Shults writes, "Triplestores are perfect for making sense out of extremely complex data. However, a triplestore is only useful if massive quantities of information can be loaded, updated and effectively queried in a reasonable amount of time. That is why Franz Technology's announcement is so interesting."

He continues, "Less than a month before the 6/7 announcement, Intel gave Franz access to one of our lab systems – a high performance server from IBM... Before Franz had an opportunity to work with this system, the largest triplestore they'd been able to assemble contained roughly 50 billion entries. Running on the 8-socket Xeon E7 system, Franz was able to load and efficiently query more than 320 billion triples, and the factor limiting scale wasn't memory or processors—it was the amount of disk space available. With some additional spindles and memory, Franz is confident that they can achieve the previously unthinkable result of a trillion triples."

Shults adds, "It's difficult for the human mind to grasp a trillion anything – dollars, stars, or triples. The important thing to understand here is that the amount of processing that goes into loading and querying a trillion triples is enormous. Unless you have a hardware platform that can deliver a corresponding amount of concentrated processing power at an affordable cost, it's all kind of pointless. What Franz demonstrated was that such a hardware platform exists, it performs even better than expected, and it delivers a level of capacity that allows customers to think about putting the full potential of the Semantic Web to use in important and creative ways."