Objectivity Case History

Customer Information

Customer:	Lawrence Berkeley Labs and Stanford Linear Accelerator
Industry:	High Energy Physics (HEP)
Application Domain:	Particle Analysis
Status:	Active customer with growing implementation
Platform:	Solaris and NT
Compiler:	C++ and Java
Other Tools:	

Customer Background

BaBar is a large "particle analysis" project spanning 500 people, spread across 80 different sites, in 10 countries. By colliding different sized particles at high rates of speed, Babar's physicists gain a better understanding of the physical nature of particles, and how they operate. That knowledge is used to understand why the universe is full of matter. After the big-bang, what happened to the anti-matter?

Utilizing the linear accelerator located at Stanford University, the Stanford Linear Accelerator (SLAC), BaBar produces huge volumes of scientific data. Upon entering its very last stage, that data is housed in a federated OBJY db. Getting to that last phase is quite a process.

The accelerator is used to collide different sized "electrons" with different sized "positrons". These two particles are, by their very nature, the antithesis (or anti-particle) of each other. First, the accelerator creates and "stores" the positrons. Then, a series of collisions are run where the positrons and electrons are collided at super-high speeds. Collisions are run with particles of various sized-weights, materials, and at various speeds. The intent is to see how these particles work, behave, and interact upon collision.

How many collisions are run? 25 million per second.

Special detectors "surround" the collisions at specific collision points, and these detectors are used to extract and record the relevant information about where the particles travel to, how they behave, how they decay (how long they take to do that), etc. It is estimated that only one of every 2,000 collisions actually produces some type of "useful" information.

Customer Environment

The data from this (first) stage of particle collisions is then fed into a bank of 300 specialized processors, running a VME based OS and connected via a high-speed "GigaBit Ethernet" network. That data is then crushed down and sent to a bank of 30

Solaris based machines, where the data is further analyzed and scrubbed. After a long collection period (creating about 300 TB's per year) it is sent to another station of 100 Solaris based processors with large disk arrays. That's where Objectivity comes in. The data is now "blown up" by a factor of four and shipped to a SUN Enterprise 10,000, which is augmented by HPSS, a high performance storage system designed as a cooperative venture between Lawrence Berkeley Labs and IBM. HPSS runs on an IBM AIX box and is outfitted with state of the art optical jukeboxes and other state of the art high performance storage hardware. OBJY clients running in 10 countries can access the particle data through the E-10000 and HPSS. BaBar clients are currently run on Sun/Solaris, HP/UX, DEC/OSF, and IBM/AIX machines.

Buying Criteria

Multi-Platform support: BaBar clients are many. Their languages are C++ and Java.

Performance: Performance with large amounts of data was important.

Replication: BaBar felt replication was a key requirement. It should be noted that BaBar will push the replication envelope. They will need OBJY to support replication across a wide area network (WAN). They are willing to work cooperatively with us on this issue, and are looking forward to making this happen.

Distribution: Our federated distributed architecture was of primary importance.

Competitive Landscape

BaBar purchased it five years ago. They looked at all the competitors back then. They selected OBJY for our distributed architecture and the concept of a federation. Scalability and performance were also important.

Contact Information

Objectivity Rep: Ray Gehring
Customer Contact: David Quarrie
Phone:
Email: