

## Bear Stearns Deploys Data Backbone

**N**EW YORK—Bear, Stearns & Co. is using a data backbone offering from GemStone Systems rather than relational database technology for some key trading and risk operations, says Buzz Moschetti, chief architecture officer for Bear Stearns.

The data management needs of the firm—particularly its global clearing and settlement risk analytics and program trading groups—are such that the traders must be able to read huge amounts of data quickly, Moschetti says. They also require low latency access to structured data such as objects, arrays and vectors. Thus, the high-performance data engine infrastructure that is needed to satisfy these demands has to be highly configurable and universally available as well as have “zero impact” on the performance of the firm’s incumbent systems, he says.

By overhauling data access and management, the firm hopes to improve transaction speeds and volumes, cut the time for risk computations, and get a more accurate view of its risk, pricing and portfolio valuations across multiple assets, officials say.

In general, conventional relational database offerings cannot handle this variety of highly complex data processing operations, Moschetti says. Among other impediments, if a relational database management system (RDBMS) is used, it will often require conversions, which “start to become a nuisance” in a distributed computing environment, he says. “You have to convert the data in the formats [of the database providers].”

After an evaluation and a proof-of-concept test, the firm chose GemFire Enterprise for managing in-memory data across its enterprise applications, officials say. The offering, to be implemented for at least three projects, provides the data access and support for key application program interfaces (APIs) that the firm needs, says Moschetti.

He declines to name the other vendors considered for the project.

Another capability of the data backbone is that it must spread cache across nodes “without creating more points of failure,” Moschetti says. It also has to be flexible as far as the rights to caching, persistence to disks and support of multiple APIs, including the C language and the Microsoft .Net framework. The major, overriding specifications for the project, though, are “the ‘get and put’ capabilities. It largely comes down to that,” he says. GemFire provides connections to databases, application stores, analytical tools, messaging systems, and mainframes, say officials from GemStone, based in Beaverton, Ore. (GemStone detailed the Bear Stearns implementation at the DWT 2005 New York conference and exhibition last week. See page 1.)

GemFire is being deployed in production for the firm’s program trading desk in New York, confirms Moschetti, who declines to provide further details. The program trading group was grappling with data latency issues, according to a GemStone press release.

The GemFire offering will also be a key part of a risk management project that is to debut in the summer, Moschetti says. The firm wants to cut the data latency of its risk management reporting. “We want to be able to pre-load and pre-calculate [via GemFire],” he says. Data from risk evaluations will be cached and made available to users within the firm via GemFire.

For yet another project, GemFire will be used to provide re-authentication and re-entitlement for applications that will be transitioned to a service-oriented architecture (SOA), Moschetti says. Once major applications are partitioned and changed into services, they require multiple re-authentication and re-entitlement procedures, which is “a huge stumbling block to SOAs. The answer is easy—you put authentication and entitlement into memory.” This will affect the trade processing and the market intelligence gathering capabilities of SOAs.

Eugene Grygo