



GEMFIRE BUSINESS BENEFITS

- Business continuity
- Elimination of complex architecture deployments
- Balanced distribution across all of the nodes in a distributed system



GemFire Enterprise: Resiliency Made Easy

WHY RESILIENCY SYSTEMS AND WHY NOW?

Today's electronic world has resulted in a major shift in terms of how organizations think about resiliency. Firms of all sizes are not only faced with the challenge of determining how resilient their mission-critical systems need to be, but also how they can efficiently and cost-effectively architect a "resilient" system.

Trading systems in equity and foreign exchange markets are no exception to this requirement. In fact, these systems demand continuous availability of systems at all times. With millions of dollars traded per minute in the trading systems, any downtime in systems' availability means:

- Loss of millions of dollars in revenue per minute
- Unable to calculate risks in ever changing global trading markets
- Payment of fines to exchanges and governments due to delayed execution of trades.
- Loss of customers to competitors.

THE TRADITIONAL APPROACH TO RESILIENCY

Traditionally, highly available systems have been built using N+1 clustering and storage area replication. A traditional solution has the following components:

- Relational databases in every region
- Local area messaging bus
- Message Oriented Middleware for WAN distribution
- Clustering solution for high availability
- Storage level replication for disaster recovery
- 3DNS or wide IP for IP failover

In this traditional approach, recovery from a failed node in the N+1 cluster would entail bringing the additional node online, starting the database on this node, fetching data from the database, recreating business objects from that relational data and reconnecting to exchanges and users. This procedure is not only susceptible to errors due to multiple recovery steps but also takes about 15 minutes of recovery time. The problem gets even worse if an entire cluster is lost to disaster. Not only does it take around 1-4 hours to follow the same tedious procedure to recover but also requires tremendous effort to replicate the steps for all nodes at the failed site. In a trading system, any downtime due to node, cluster or site failure converts to hundreds of millions of dollars in lost revenue.

WHAT IS THE BETTER SOLUTION?

The best way to provide nearly 100 percent uptime for data and deliver maximum resiliency is by using data management middleware to ensure there are multiple consistent copies of the active business objects in-memory at all times. As firms strive to get even closer to 100 percent uptime and ensure resiliency, distributed data caching becomes an ideal solution.

GemFire from GemStone Systems, offers a simple and powerful distributed caching solution that enables data replication across multiple nodes distributed across multiple geographical regions.

Presented as a simple HashMap API, GemFire delivers maximum value behind the scenes: You simply "put" your state into your local, in-process HashMap and under the covers the GemFire caching solution takes care of replicating this business object to multiple servers. As soon as an object is put into the

cache, it is replicated in-memory to at least one additional node. The data can be replicated to additional nodes either synchronously or asynchronously depending on sensitivity to latency and tolerance for data loss in the event of a catastrophic failure. Write-through to a database or other persistent store is done asynchronously as time permits. In essence, the distributed cache behaves much like RAID for the enterprise.

Additionally, GemFire solution eliminates the need of any fancy shared storage since the data is mirrored across multiple nodes in the distributed cache. This also provides the opportunity to distribute the workload across all of the nodes in the distributed system. Exchanges may be split between the two sites and clients will likely be distributed across the two sites, as everything except external connectivity is in a hot/hot configuration.


WHAT ABOUT RECOVERY TIME?

Application clusters deployed using GemFire are virtual, so the nodes do not need to be located close together within the datacenter -- they can be on separate subnets, using separate routers, power sources, etc. In case of a node failure, the recovery process includes detection of failure and reconnect-

tion of the clients. It's that simple. The total recovery time from point of detection takes less than 1 second. If the entire data center fails, the recovery process is as follows: detect the failure; reconnect the exchange at the alternate site; reconnect the clients; recovery is complete. The time to recover the data management layer from point of detection is around 1 second. That's a significant difference to the typical 1-4 hour disaster recovery time common in business today!

SUMMARY

As distributed computing deployments become the norm rather than the exception, resiliency will become one of the most critical issues facing global corporations. By using GemFire, firms can achieve nearly instantaneous recovery from outages -- real business continuity -- while simultaneously simplifying their architectures. This one product takes the place of an H/A solution, a shared-storage environment, storage-level replication and a wide-area data distribution product, providing a complete data resiliency architecture for mission-critical systems.



Corporate Headquarters:
1260 NW Waterhouse Ave., Suite 200 Beaverton, OR 97006 | Phone: 503.533.3000 | Fax: 503.629.8556 | info@gemstone.com | www.gemstone.com

Regional Sales Offices:
New York | 90 Park Avenue 17th Floor New York, NY 10016 | Phone: 212.786.7328
Washington D.C. | 3 Bethesda Metro Center Suite 778 Bethesda, MD 20814 | Phone: 301.664.8494
Santa Clara | 2880 Lakeside Drive Suite 331 Santa Clara, CA 95054 | Phone: 408.496.0242

Copyright© 2005 by GemStone Systems, Inc. All rights reserved. GemStone®, GemFire™, and the GemStone logo are trademarks or registered trademarks of GemStone Systems, Inc. Information in this document is subject to change without notice.