

## Industry Speak



### Using Data Fabric for Continuous Availability in Clustered Stateful Applications

*Jags Ramnarayan*  
*Chief Architect*  
*GemStone Systems*

High Availability (HA) refers to a system or component that is continuously operational for a reasonably long period. Generally, availability is measured relative to "100% operational" or "never failing." A widely held but difficult-to-achieve standard of availability for a system or product is known as "five 9s" (99.999 percent) availability.

A computer system or a network consists of several parts. Since all parts usually need to be present for a computer system or a network to be operational, meticulous planning is required for high availability centers in the areas of backup and failover processing, data storage and access.

HA is about the pursuit of continuous availability. To a certain degree, HA is addressed by most infrastructure software components - Web servers, application servers, messaging solutions and database software through clustering. But ensuring availability, especially for highly concurrent stateful applications with many moving parts remains a quest.

We present the notion of a "data fabric" (also referred to as a "data grid") that pools memory (and optionally disk) across the cluster to efficiently manage the state with guarantees of consistency even in the face of many conditions of failure. The data fabric also provides unique and reliable event distribution capabilities to execute application behavior close to where the data resides, enabling a new class of real-time event-driven applications.

High Availability Clusters (also known as HA Clusters or Failover Clusters) are computer clusters that provide a high availability of services. HA clusters are not a new phenomenon. Most application server platforms offer support for clustering to ensure availability of hosted services. Planning continuous availability not only requires dealing with network failures but also redundancy in application behavior execution mechanisms, redundancy in access to data or manipulation of shared data. Making infrastructure components highly available is not sufficient. They have to be extremely reliable, i.e. the software should perform exactly what it was designed to do. In other words, software should deliver predictable results, irrespective of the fluctuations in load conditions or exposure to conditions of failure.

Let us examine a data fabric (or a data grid) approach to continuous availability. The premise of a data fabric is that by keeping all application states available in a consistent form across the cluster and shielding the application from failures in the back-end data sources, the availability and reliability of the application can be significantly boosted. A data fabric is a cluster of nodes connected with each other like a mesh. It manages data redundantly in the middle-tier, using a "shared nothing" design, to achieve very high availability. A data fabric is a sophisticated distributed caching platform with distinctive architectural features:

- Critical data can be replicated to any number of nodes in the cluster synchronously (with support for optimistic to pessimistic concurrency control)

- When load increases, cluster nodes can be added dynamically and automatically to create more redundant copies of data, thereby increasing availability
- Data can be stored in memory (multiple copies). Alternatively, it can be stored on a disk. When stored on a disk, each copy on the disk is independent of each other. Synchronization is done only at the in-memory cache level, thereby avoiding a shared disk storage mechanism.
- When the data volume is high, it can be spread across nodes with a configurable number of redundant copies at all times
- When the load conditions change, the fabric can condition the load by moving client connections to maintain uniform load balancing
- Real-time disaster recovery requires replicating data using asynchronous communication to other clusters (arranged in a hub-spoke manner). The failure of the entire cluster will result in the clients move to other data clusters without any operator intervention.
- All application behavior, if triggered through events delivered through the data fabric, is guaranteed to be highly available. Any conditions of failure are detected and application behavior automatically triggered elsewhere. The application is automatically notified with a "possible duplicate" flag indicating that the behavior might have been executed partially elsewhere.

Ensuring high availability of stateful applications transcends simple clustering techniques that replay invocations on alternate nodes in the cluster during conditions of failure. You can address HA for stateful applications through the use of data fabric. A data fabric ensures data visibility across the entire cluster through main-memory replication and an architecture that eliminates all single points of failure.

## Reference

---

High availability clusters: [http://en.wikipedia.org/wiki/High-availability\\_cluster](http://en.wikipedia.org/wiki/High-availability_cluster)

## Bio

As the Chief Architect of GemStone Systems, Jags is responsible for technology leadership for its high performance distributed data grid platform. He has guided several customers in adopting and successfully deploying middle-tier data grids. With more than 18 years of experience, Jags helped GemStone architect its J2EE application server platform embedding an object database as a core cache in the 1990s.

As the Web Services Architect of BEA, Jags represented BEA in the W3C SOAP protocol specification, JAXM and other standards. Jags has also represented GemStone in the J2EE platform specification and the EJB expert group. He has presented at several conferences on high performance data management. Jags holds a Bachelor's Degree in Computer Science and a Master's Degree in Management of Science and Technology.