



IMPROVE RELIABILITY IN YOUR CLUSTERS WITH AUTOMATED STORAGE

Server farms, N+1 failover, and cluster file systems are server implementations that improve reliability through redundancy, particularly when configured with shared pools of storage.

Cluster file systems scale workloads by adding nodes to the cluster, and need intelligence and automation in SAN storage arrays to function smoothly without disruption or bottlenecks.

CONSOLIDATION EFFORTS RAISE RELIABILITY AND RECOVERY NEEDS

Because IT infrastructure is the foundation of most business operations, reliable and predictable service is critical. However, not all applications have the same business value. IT administrators constantly make decisions about service levels and the most economical methods to achieve them. Consolidation efforts (which involve providing greater IT services on fewer physical resources) raise reliability and recovery needs because of the potential impact of losing a valuable physical resource.

IMPROVING SERVER RELIABILITY THROUGH CLUSTERING

Reliability is commonly achieved through redundancy, which can mean dual controllers and RAID in your storage configuration or multiple physical servers for your application environment. Clustering describes a set of servers operating together to ensure reliable application services. Clusters come in several forms:

- Farms are collections of servers that process the same function – for example, web server farms. Farms are typically integrated with a front end that distributes work to each server and monitors performance, adjusting workloads when servers are down or operating slowly. This deployment is common with stateless applications – that is, applications that have little "unique" or modified data.
- Failover (or N+1) is generally used with small implementations of two to eight servers working together. The typical configuration includes (N) applications running on (N+1) servers; if a server fails, the extra server (+1) takes on the failed application. This deployment is common with "stateful" applications such as e-mail systems, databases, and file and print servers that have a lot of unique data.
- Shared cluster file systems are the most advanced form of clustering, not only supporting cluster farms and failover deployments, but also allowing all nodes in a cluster to operate directly on the same data at the same time.

CLUSTERS NEED SHARED STORAGE

Clusters require increasing levels of data sharing depending on the configuration. Web servers share mostly read-only data – this can be accomplished by shared storage or simply with local copies of data. Failover clusters share access to application data, but at any given time only a single server uses a given application's storage. Cluster file systems have all nodes accessing application storage at the same time.

HEAVY DATA LOADS REQUIRE AUTOMATED MANAGEMENT

There are many cluster file system products on the market: VMware® ESX Server™, PolyServe™ Matrix Server™, Symantec™ Storage Foundation, and Red Hat® Global File System (GFS) are popular examples. These products share common attributes: (1) they provide a common cluster file system across all the nodes, so all nodes can efficiently access the same data at the same time; (2) applications may be operating on multiple nodes in the cluster, accessing the same data at the same time; and (3) application performance can be scaled by adding cluster nodes.

This "all-active" clustering requires shared storage. Any- or all-node storage access puts new pressures on the storage configuration – access is more dynamic, and access patterns change quickly as the cluster workloads vary. With a non-clustered system, you might add an e-mail server and a database server making, specific slices of storage available to them. But with clustered file systems, applications can be actively modifying data and accessing the same volumes simultaneously. As you scale, more and more servers need to access the same data, using the same network links, cache, and controllers - you must have a storage array that can handle these dynamic data loads.

This type of heavy data access is best served by an intelligent, iSCSI SAN like the EqualLogic PS Series. These arrays offer key improvements such as automatic load balancing and scaling of network links, controllers, and drives. By automatically balancing loads across all available storage resources, PS Series arrays adjust automatically to multi-node data access. In addition, the storage virtualization and simple, online scalability dramatically improve management for the

dynamic workloads of cluster file systems. Importantly, as you add PS Series arrays, you add not only more disk spindles but more network connections – so your growing environment doesn't hit a connectivity bottleneck and you actually increase performance by adding capacity. And all of this is done while online, so you don't have to take down operations to shuffle volumes around.


Clustered systems provide improved application reliability. They require storage with greater reliability and ease of operations – therefore, they need a dynamic, easy-to-use, flexible SAN. The more intelligence and automation in your storage array, the more smoothly your cluster systems will function and scale.

To view other Coffee Break Bulletins or to learn more about EqualLogic, visit us at www.equallogic.com.


Cluster Types

A = application S = spare


Server Farm: sharing read-only data



Failover (N + 1): sharing access to application data



Cluster File System: accessing application storage from all nodes



110 Spit Brook Road, Building ZK02, Nashua, NH 03062
Tel 603.579.9762 / Fax 603.579.6910 / www.equallogic.com