

Lotus Domino Advanced Services: High Availability Powered By Notes

by Craig Smelser

[Editor's note: This article resides in "Notes Today", the technical Webzine located on the <http://www.notes.net> Web site produced by Iris Associates, the developers of Domino/Notes.]

In an online world, continuous, reliable access to networked information is indispensable for creating dynamic business advantages and competitive strength. While most large enterprises function within a networked global business market, more and more smaller companies are discovering the world of networked computing to expand their competitive edge beyond their customary borders.

As computing networks and networked applications become more critical, features like **high availability** become more important. Whether your company maintains a catalogue on the World Wide Web, or you are a stockbroker receiving constant updates from New York, London, and Tokyo, the computing resources that provide this information are critical to your work, and their availability is a non-negotiable requirement.

Lotus understands the world of high availability systems and the importance of providing continuously available data and servers. The new **Lotus Domino Advanced Services** are a set of advanced server features integrated with Lotus Notes to provide you with the availability and scalability you need. A fundamental component of Domino Advanced Services is the Domino server cluster(s) that provide high availability for Notes networked applications.

This paper describes high availability and its importance in delivering highly reliable access to networked information and introduces Lotus Domino Advanced Services.

What Is High Availability?

High availability is a method of computing that assures you and your colleagues access to critical information you need, anytime and from anywhere. High availability provides continuous, uninterrupted access to networked resources, in spite of serious failures that may occur within the network. Computer resources that are available "24x7" can greatly enhance workflow and productivity and together with IBM, Lotus remains focused on delivering systems that are highly available.

There are a number of different strategies to achieve continuous data access. Ten to 12 years ago, there were only a few methods, and the primary implementations were either through hardware or operating system software (or both). These initial methods introduced the fundamental concepts of synchronized data redundancy still in use today. However, more recent methods have improved on these early concepts and take advantage of new technological shifts in the computer industry. The result is that high availability can now be implemented in sophisticated networked applications using *standard* hardware and operating systems.

Clustering for High Availability

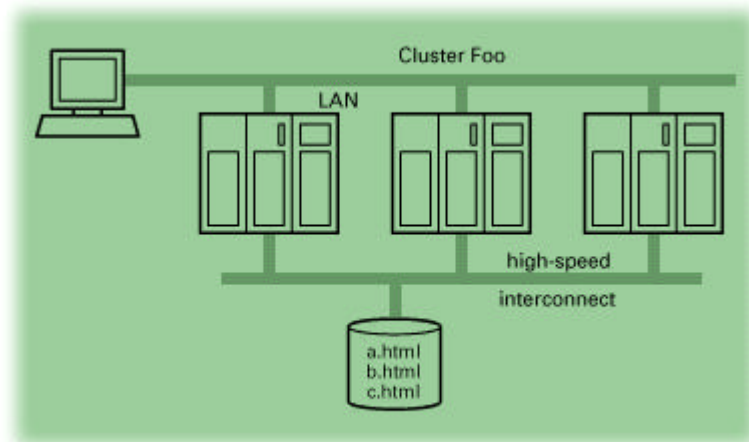
In networked computing systems, high availability is achieved by providing synchronized redundancy of critical system components through "clustering them". A cluster is a collection of similar hardware or software components that clients on the network access as a single, virtual resource that is highly available. Individual cluster components can be physically located in the same room or dispersed around the world and connected by a network that spans the enterprise. With clusters, key network elements are not only duplicated, but also connected to work together as a team. This enables the synchronized redundant components to automatically and transparently take over for failed components and maintain data availability.

Synchronization of the cluster components is the key to this style of high availability. For example, with Domino clusters, if one database is updated, and its replica remains unchanged, the two databases are

no longer duplicates. The databases must establish frequent communication over the cluster network to stay in synch. When databases on different servers are identical, they are said to be "synchronized."

For several years, computing researchers have been actively seeking the most efficient method for synchronizing data on highly available systems. Some of the strategies include:

- **Redundant Arrays of Independent Disks (RAID)**-- One of the earliest and most widely used implementations of hardware clusters today is RAID technology. RAID disk mirroring redundantly stores information on multiple disk drives, in "mirrors" of one another. Each RAID implementation requires a different method for achieving hardware redundancy. Each also offers its own unique combination of price, performance, maintenance, and recovery variables. Many don't consider RAID a cluster technique, but I list it here because it illustrates the properties that best define a cluster and because its benefits to high availability are well understood.
- **OS Clusters** -- While RAID links multiple, separate data sources, OS clusters provide highly available access to network and computing services. In OS clusters, multiple independent file servers can share a common view of the file system in use. Similarly, clustered file servers appear to network clients as a single resource, sometimes even given a single "alias" or name for client access. Some OS clusters can also perform load balancing. Load balancing distributes the work associated with the processing of a task or request among file servers in the cluster. Any OS cluster method you select to provide access to network resources will have a specific impact on overall system performance as well.



High Availability and Domino Advanced Services

When a clustered system focuses on the availability of hardware or operating systems, the value and importance of high availability may be obscured, even to users and companies highly dependent on their networks and servers. High availability however is not just about the systems and networks that connect users together. Enterprises that are serious about high availability acknowledge its critical contribution to their business and deliver it as an integral component of the applications on which that business depends.

Lotus Development Corporation, a subsidiary of IBM, has experience in and a strong commitment to assisting business enterprises in delivering information critical to their businesses. For Lotus, the challenge of providing high availability lies in both highly available servers and highly available networked applications and databases. To address this challenge, Lotus offers Lotus Domino Advanced Services, an enterprise-caliber solution.

Lotus Domino Advanced Services provides the ability to cluster Domino servers and manage Lotus Notes databases for high availability. Combined with the power of Lotus Notes, Domino clusters take advantage of Notes multi-platform environments and incorporate its many administration, tracking, and security functions.


Domino Clusters: The Next Generation of High Availability

Domino clusters are interconnected teams of Domino servers that provide uninterrupted access to Notes networked information resources, including Notes messaging, databases, and other service components. ***Domino servers, powered by the ever-expanding functionality of Notes, provides a multi-platform networked system that makes high availability, a feature once offered by large and expensive proprietary systems only, now accessible to large and small enterprises alike.***

Using Domino clusters you can connect up to six Domino servers in the same Notes domain with a local area network. There is no requirement that each Domino server run on the same operating system or hardware platform. Domino clusters include many OS cluster features but take advantage of Notes' excellence in database management to deliver much more.

- **Client and Server Failover** -- Domino clusters provide fail-over protection for business critical databases and servers, including pass-through server failover to other servers in a cluster. Domino servers failover clients by redirecting database requests to other servers in the cluster. This redirection is managed by a Cluster Manager responsible for tracking cluster membership and cluster server status. Except for changes in the database icon, failover is totally transparent to the user.
- **Workload Balancing** -- In addition to failover, Domino clusters can take advantage of Notes highly acclaimed replication technology to evenly distribute the workload of heavily used databases across multiple servers in the cluster. When a user (with the Notes client or a standard HTML browser) updates a database, Notes immediately synchronizes [replicates with] -- on a transaction-by-transaction basis -- all other replicas in the cluster. Servers in the cluster can be tuned based on their location, contents, or other performance parameters. Domino achieves this by creating a single, virtual address for the cluster and routing transmission requests among the various servers in the cluster to achieve greatest efficiency.
- **Synchronized Shared Access** -- Unlike OS clusters that might use a "scheduler" or "lock manager" to maintain synchronized shared data, Domino clusters replicate changes among databases as they occur. Domino cluster replication ensures that all changes, whether to a database or to the cluster membership itself, are immediately passed to other databases or servers in the cluster. This type of replication is known as "event-driven" replication and is performed by a special component called the Cluster Replicator. When the Cluster Replicator is unable to replicate database changes, it retains the information in memory and tries repeatedly to replicate the changes until it is successful. When the destination database becomes available, the Cluster Replicator pushes all changes to the database regardless of how many changes occurred on the source while the database was unavailable. In special cases, where event-driven replication may not be possible, an administrator can use "schedule-driven" replication to update databases. In these ways, Domino clusters actively maintain status on database and server availability and keep databases tightly synchronized to provide high availability.
- **Scalability** -- Domino clusters also support scalability of up to thousands of simultaneous Notes sessions, making information available to anyone who needs it. Special database attributes, set by a Notes database manager in the server NOTES.INI file, enable restricted access to specific databases and ensure security of data when necessary. As a Domino cluster grows and maintenance issues become more important, Notes further provides tools that enable an administrator to accurately analyze cluster events, monitor databases, and identify problems rapidly.
- **High Availability** -- Domino provides high availability in four ways:

1. Creating an appropriate number of database replicas on multiple servers
2. Managing client failover to available database replicas
3. Managing user and server access to data in the cluster. This includes event-driven replication of databases to ensure synchronized shared access of data within the cluster. All members of a Domino cluster share a common Cluster Database Directory that contains information about all databases and their replicas within the cluster. When one server updates its list of databases, these changes get replicated to the shared directory on other servers in the cluster. It is the task of the Cluster Database Directory Manager to keep this directory updated with the most current information. This information is used by the cluster to determine fail-over paths and to control access to a database.
4. Maintaining an accurate and up-to-date status on all servers in the cluster. This is the responsibility of the Cluster Manager, a software component installed on each Domino server when it is added to a cluster. The Cluster Manager tracks the state of all members in a cluster. It keeps a list of which servers are currently available and maintains information about active workloads for each server. Using this information, a Domino cluster is able to redirect client requests to servers that it knows are available and can handle the request. Each Cluster Manager maintains a memory cache of this information and updates it regularly by probing other cluster servers for their status.



The Future of High Availability and Lotus Domino

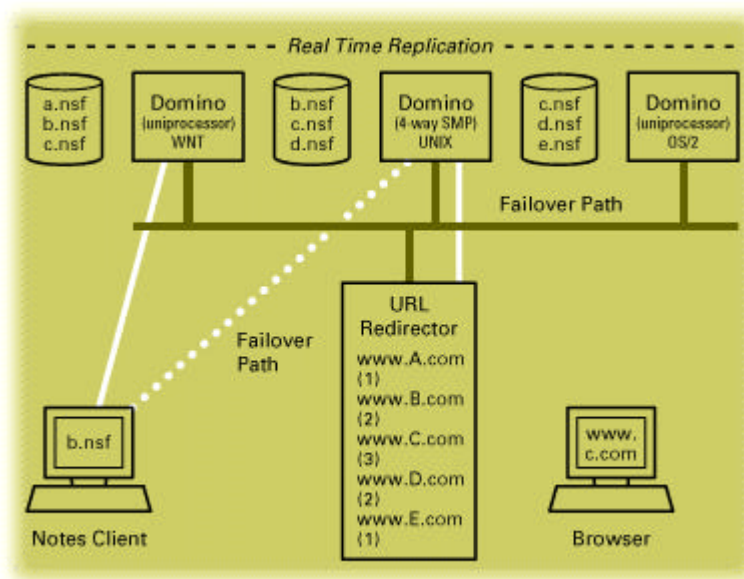
The explosive growth of the Internet, Intranets, and the World Wide Web has positioned browser software as a significant alternative to more traditional client software. The use of these technologies can potentially reduce the total cost of network ownership and management by leveraging the strengths of standards-based network protocols and tools. Lotus supports these recent developments and with the introduction of Domino provides support for popular Web browser features by significantly expanding their reach and accessibility with Domino Advanced Services. And, through Notes replication, it is fairly straightforward to deploy Domino clusters that are synchronized for load balance and failover.

The key to the future lies in also simultaneously supporting clustering for failover and load balancing with non-Notes clients (that is, "generic" non-Notes browsers such as Netscape Navigator) connected to tightly replicated Lotus Domino clusters. Traditional Web server clustering typically is achieved by disk sharing of the same physical HTML files or through manual administration to replicate or synchronize. In addition, load balancing through "URL redirectors," like "round-robin DNS" have been sub-optimal because they

attach to a particular server for an entire transaction session. To this point, these strategies have been acceptable because first-generation Web sites have typically consisted of HTML publishing -- or read-only information.

With the coming generations of true Web applications -- dynamic, content-rich transaction applications that involve filling out forms, adding data, or updating databases -- failover and load balancing take on increasing complexity. While Notes replication provides immediate clustering benefits, IBM's forthcoming Network Dispatcher takes this technology a giant step further. It will address some of these challenges through its ability to provide a "network packet" layer of control for how clients bind to servers during an application "session."

The Network Dispatcher will parcel out *individual TCP packets* among the clustered group of servers, providing a much finer granularity in the balancing of the processing load. Further, Web applications will be able to programmatically control the Network dispatcher, building intelligence into the failover process. The result will be highly available Web applications - just what the industry needs to move from "surfing" the Web to "Working the Web".



Summary

Clusters are a critical element of any system that provides high availability, regardless of whether the clusters are network hardware, operating systems, or applications. Domino Clusters improve on traditional OS cluster methods by providing:

- Higher availability of applications
- Higher degrees of scalability and growth flexibility
- More reliable data synchronization and access
- Well-known data security benefits through Notes
- Ability to take advantage of Web and Web-based intranets benefits.

In addition, Lotus values the advantages that OS clusters bring to networked systems and the role of OS clusters as important contributors to high availability. Lotus also believes that the technologies that underlay OS clusters may be able to help improve the overall availability of many Notes implementations

and intends to pursue the integration of Domino clusters with OS cluster systems as a major theme of the continued development and promotion of Domino Advanced Services.

ABOUT THE AUTHOR

Craig is a Domino server Product Manager at Lotus.

Copyright 1997 Iris Associates, Inc. All rights reserved.