



CaminoSoft

StandbyServer™ for NetWare

White Paper

August 1, 2003

CaminoSoft

**StandbyServer™
for NetWare**

Executive Summary

CaminoSoft is dedicated to providing products that ensure high server availability. CaminoSoft StandbyServer™ for NetWare is an automatic server recovery system that protects multiple NetWare servers (called primary servers) with a single standby machine. StandbyServer maintains a real-time mirror of the primary servers' data but will also automatically take over server operations for any one of the primary servers when that server experiences a hardware or software failure, or even when planned maintenance is being performed. When the standby machine is standing in for a failed primary server, it continues the mirroring process for the other protected primary servers, thus keeping their data synchronized and up to date.

StandbyServer provides the highest degree of system flexibility, ease of use, and open configuration of any fault tolerant or high server availability product. No proprietary disk systems, monitoring devices or special operating system software is required. It takes advantage of the economies and performance of industry-standard hardware to extend the configuration capabilities and flexibility of implementing automatic server failover to a standby machine.

The CaminoSoft high availability family of products provides market-leading solutions for automatic NetWare server recovery. StandbyServer is the latest multi-node offering in the family of high-availability solutions for NetWare.

StandbyServer for NetWare offers:

- Fully-automatic server failover
 - The standby machine detects and protects against both hardware and software faults
 - Automatic client reconnection to the failed over standby machine (when using the latest auto-reconnecting client software from Novell or Microsoft).
- Usage of Native OS
 - Does not replace or alter NetWare
 - Supports NetWare versions 4.x, 5.x, and 6
 - Uses native NetWare mirroring
- Support for IPX or IP protocols
- SMP (Symmetric Multi Processing) compliant
- SNMP (Simple Network Management Protocol) network management messaging capable
- Uses standard NetWare network adapter technology to mirror data between servers
 - Supports use of a dedicated inter-server network (recommended)
 - All mirroring traffic is directed over the dedicated network
 - No impact to the performance of the existing production network
 - Dedicated network uses industry-standard cables, cards and drivers (e.g., 100BaseT and IP or IPX)
 - Dedicated network is routable and bridgeable
 - Standby machine can be in a different room, building or locale
 - Data mirroring traffic can optionally be routed over the production network or backbone (not recommended for heavily-used networks)
- Affordability
 - A single standby machine can protect multiple primary servers
 - The standby machine need not be identical to the servers it protects
 - The standby machine can perform other services while monitoring the primary servers

- Benefits
 - Increased reliability of server operations
 - Increased availability of server data and services
 - Allows scheduling of server maintenance at a convenient time

© 2003 CaminoSoft Corporation (portions copyright CaminoSoft Corporation). CaminoSoft and the CaminoSoft logo are trademarks of CaminoSoft Corporation. StandbyServer, SnapShotServer, OFFSite Archive, and all other trademarks and trade names referenced herein are the property of their respective owners.

The software discussed in this paper is protected by one or more of the following patents and/or is the subject of pending U.S. and/or foreign patent applications:

US 5,649,152 US 5,812,748 US 5,835,953 AU 687,274

CaminoSoft Corporation
600 Hampshire Road, Suite 105
Westlake Village, CA 91361 USA

www.caminosoft.com

Phone 800-889-8248 • 805-370-3100
Fax 805-370-3200

Introduction

The primary objective in developing *StandbyServer* was to provide a solution to the problem of server downtime and a means to replicate mission critical data. System downtime is becoming progressively more expensive as organizations place higher demands on the availability and reliability of their system data and network services. Businesses are increasingly relying on their network data to make real-time decisions. Those systems must be available in the event of a hardware or software failure. With a server standing by with its own replica of the network data, continuous business operation is ensured in virtually any failure scenario.

To help reduce downtime some organizations keep a spare machine available to restore the functionality of a failed server by transferring the disk devices. Others have used only external disk devices, so that the transfer of system disks from a failed server to a spare server is easier. Either path requires a technician to make the system data available again. To make this manual approach to server fault tolerance work, full-time technicians must be on site or on call to perform the failover.

Moreover, the actual disk drives containing vital system data might have malfunctioned, forcing the system administrator to risk the loss of valuable system data by restoring from the latest tape backup. Using proprietary external storage is a disadvantage as well, since these proprietary systems might not be repairable or replaceable, and represent a single point of failure. These approaches are not satisfactory to businesses that rely on the high-availability of their system data and network services.


For years, *StandbyServer* for NetWare has been the leading solution in providing protection against server downtime. *StandbyServer* provides the following features and benefits:

- Downtime is essentially eliminated
- Comprehensive data reliability and high-availability are included in one solution
- Real-time mirroring provides the most complete and reliable solution to data availability
- Provides fault tolerance for the disk channel
- Provides high-availability for the entire server, including data and network services
- When a primary server fails, the standby machine uses the same startup process as the failed server to ensure that it functions (looks and feels) just like the server it replaced. The network clients see the same server, data and network services
- All server components are protected, not just the disk drives—complete and fully redundant server
- Automatic server failover—no manual intervention required
- Compatible with industry-standard hardware—uses standard communications adapters with IPX or IP protocols
- Does not require identical servers—*StandbyServer* allows mismatched machines to be used
- Using the utility server feature, the standby machine can perform other server functions such as a print server, backup (for all the primary servers), gateway, or router

StandbyServer provides all of the above features and benefits, and further, protects multiple primary servers with a single standby machine.

How StandbyServer for NetWare Works

StandbyServer for NetWare is a software solution providing a mechanism to connect several primary servers to a single standby machine. The standby machine performs real-time server mirroring for all the primary servers, and constantly monitors their health to ensure functionality. In the event of a primary server failure, the standby machine automatically stands in for the failed server. The standby machine boots up as the failed server and executes the original server's startup files.



When the standby machine boots up as the failed server, all necessary hardware drivers are loaded, the failed primary server's volumes are mounted, and the server takes on the same name, server addressing, bindery, and NDS information as the failed primary server. As a result of maintaining the server identity, client machines running 32-bit client software from Novell or Microsoft are automatically reconnected. No special operations are necessary to restore client-server connections. Non-compatible clients can simply re-login to the standby machine, using their same process (login names, scripts, etc.).

Normal Operation

Under normal operation, no clients are logged into the standby machine and therefore a second full user count of NetWare is not necessary. The standby machine uses a separate licensed copy of NetWare. When operating in standby mode, no user connections are needed by the standby machine. With earlier versions of NetWare (through 5.1), only a runtime license is required on the standby machine. For NetWare 6.x, a copy of NetWare capable of running StandbyServer in standby mode can be obtained for no charge via Novell's website. When operating in the primary role (during a failover), the standby machine uses the licensed user connections from the primary server. The standby machine loads a suite of CaminoSoft NetWare Loadable Modules (NLMs) that monitor the status of all primary servers and provide communications paths between the machines. The communication paths and CaminoSoft software allow the primary server to control the disk devices that are housed by or attached to the standby machine. Native NetWare mirroring is incorporated to mirror data (over standard IPX or IP protocols) from the primary machines to the standby machine. These NLMs also facilitate the flow of data between the machines and handle synchronization and data integrity issues.

The Failover Process

In the event that any of the primary servers fail, the standby machine downs and exits its version of NetWare. Then the standby machine boots up automatically as the failed primary server, using the same license of NetWare that was running on the failed primary server. Since all disk devices are mirrored, the same boot-up sequence, login scripts, bindery, NDS data, and primary server network services are restored. From the client perspective, the standby machine looks and feels like the primary server. The standby machine also continues mirroring the data from the other primary servers. This failover process only requires as much time as it takes to detect the failure of the primary server and reboot the standby machine. The volumes of the primary are automatically remounted so clients of the failed server can quickly regain access to data and applications.

After a primary server failure, the standby machine takes over the failed server's operations allowing the failed machine to be diagnosed and/or repaired. StandbyServer supports only one failover at a time. In order for StandbyServer to protect against other failures, the failed primary server must be repaired and placed back into service so that the standby machine can again function in its principal standby role. This restores the original configuration of the network and allows the standby machine to monitor and mirror all primary servers again.

Utility Server Operation

The standby machine has the ability to not only be a dedicated standby machine, but also to operate as a fully functioning NetWare server, capable of its own independent processes. This is called the utility server feature. Using this feature, the standby machine can function as an independent active server on the network, as well as a standby for the primary servers. The CaminoSoft SnapShotServer™ for NetWare backup enabler makes use of the utility server feature. Other applications such as fax server, print server, CD-ROM jukebox server, Internet services or other processes may be executed using this feature. The latest versions (v6 and later) of StandbyServer operate only in utility server mode.

Client Machine Behavior

When a primary server fails and the standby machine takes over, the client machines that were logged into the failed server will experience an out-of-service delay for a few minutes during the failover process. Clients attached to the other protected servers in the StandbyServer cluster continue to work without interruption. The actual failover time depends on the size and number of files contained in the volumes on the failed server. After the SYS volume is mounted on the standby machine, client machines can now access the server and server operation is resumed. After the failover has taken place, the standby machine acts as the failed primary server and continues to mirror data for the remaining primary servers.

Client machine behavior depends on the version of NetWare being used, the network protocol, and the client networking software. Older client machines will experience a server outage and will be disconnected from the network. Users must re-login after the failover takes place with the same user names, passwords, and login scripts as before. Any client application software will need to be restarted. Server performance and response time may vary depending on the differences in machine resources (memory, processor, etc.).

Client32 or 32-bit client machines will experience a short server outage during the failover. These client machines maintain a persistent connection and resume normal operation automatically after the failover is complete. All disk operations that were in progress when the failure occurred are reinstated and finished. Normally, no data loss occurs because the server will back out transactions using transaction tracking and the client will retransmit unconfirmed transmissions.

Machine Considerations

StandbyServer makes use of industry standard communication protocols, allowing it to be used in a wide variety of server platforms. Servers that do not incorporate standard busses can be used as either the primary or standby machines. If the machine can communicate over IPX or IP, it can be used in a StandbyServer configuration.

StandbyServer is flexible in that the standby machine does not need to be identical to any of the primary servers. Even the actual disk devices do not need to be identical between the primary servers and the standby machine. Only the actual disk partitions need to be sized similarly between the protected primary servers and the standby machine in order to allow NetWare to mirror the disparate devices.

While the standby machine does not have to be identical, it should be reliable and have as much memory configured as the protected server with the largest amount of memory. This requirement is necessary to ensure the standby machine is fully capable of handling the primary server role should a failover occur. The RAM requirements of NetWare should be followed for all machines. StandbyServer uses approximately 128 KB of RAM in each of the primary servers and approximately 1 MB of RAM per primary in the standby machine. See the section on *System Requirements* for further details.

The Dedicated Link

StandbyServer employs a dedicated link, which has several distinct advantages:

- Avoids a single point of failure by implementing multiple connection paths between servers
- The dedicated link is used for all mirroring data traffic between servers
- No data traffic is added to the production network
- Provides a second path to verify server health and operation
- Prevents inadvertent failovers
- Higher performance

The failover process does not occur until both the dedicated link and production network paths to a primary server have been checked to determine that a primary server has in fact failed. The dedicated link can be created using standard communications hardware such as, 100BaseT or Gigabit Ethernet, FDDI, Token Ring, or ATM adapters.

With StandbyServer, users are only limited by the constraints of the particular communication adapter used to implement the dedicated link(s). For example, FDDI boards and cabling can be used as the dedicated link, in which case the fiber cable can be up to 20 kilometers (single-mode fiber) or 2 kilometers (multi-mode fiber). The FDDI signals can be routed, bridged, or even placed on the same fiber backbone as production network traffic. Benefits and/or constraints of other media types also apply when used to implement the dedicated link between the primary server(s) and the standby machine.

Running without a dedicated link may be necessary if StandbyServer is implemented using a network backbone or when using long distance equipment between servers. In such cases care should be taken to ensure that the single link is highly available and secure; otherwise, the Autoswitch feature should be disabled. If a dedicated link is not used and some aspect of the production network fails, all communication between the primary servers and standby machine is stopped. This will cause the standby machine to detect a primary server failure and take over that server's operations. The ability of StandbyServer to run over existing network communication links makes it very flexible and open for use as a disaster-recovery solution. CaminoSoft's OFFSite Archive™ for NetWare functions with StandbyServer to enhance an organization's ability to perform wide-area disaster-recovery and remote data replication.

Server Placement

The primary servers and standby machine can be placed in different rooms, buildings, or in certain situations, even different cities. By taking advantage of modern networking communication topologies and wiring media, the system administrator can choose to implement a StandbyServer-based system in various configurations. Routers, bridges, and hubs may be used in conjunction with the dedicated links; however, the same rules that govern IPX and IP also apply to the dedicated link. For example, a slow IP or IPX channel will decrease system performance accordingly.

WAN Considerations

Some WAN topologies and media are very slow compared with LAN topologies and media. For example, a T1 bandwidth is capable of supporting 1.54 Mb/second—1/64 the speed of 100 Mb/second Ethernet. A T3 link is capable of 45 Mb/second and would make an ideal dedicated link in a WAN environment, but T3 cost would need to be weighed against its performance.

Under normal fully mirrored conditions, only incremental data changes are sent over the dedicated link, thus most WAN connections will support the amount of data being sent between the machines. But under heavier loads, when many data changes are taking place (like a heavily used database), common WAN connections are not fast enough to keep up with the data being sent over the dedicated link. This causes the primary server's dirty cache buffers to increase until the data can be sent over the WAN connection, thus requiring a larger amount of server memory and possibly reducing server performance.

Most LAN connections, particularly the recommended 100 Mb/second links, will not cause any performance degradation since they can keep up with most disk systems.

Benefits

CaminoSoft's *StandbyServer* can virtually eliminate server downtime in the event of a primary server failure. *StandbyServer* directly addresses high server availability in three areas: reliability, accessibility, and data replication.

- Protects multiple primary servers—up to 20 but the ideal range is between 2 and 6 servers.
- Automatic recovery from most common types of server downtime, including software, hardware and maintenance.
- High-availability of all server functions and network services even when a primary server crashes.
- Because the operating system and the data are mirrored on a separate machine, all server components are highly available.
- The machines need not be identical—the standby machine needs to be NetWare certified and have sufficient disk devices installed to mirror the disk storage of each primary server.
- Mirrored disk devices are in a separate environment. This allows the most critical and valuable part of the data system to be placed in a separate location in order to survive a data disaster. The expansion capabilities of a standby machine are also utilized—mirrored disks can be housed in a separate enclosure.
- The standby machine can function as an active server. Using the utility server feature, the standby machine can function as an active NetWare server while it protects each primary server.
- The automatic failover can be disabled in situations where the standby machine is being used as simply an off-site data repository or data vault.
- Uses native NetWare disk mirroring. There is minimal latency for writing data to the standby machine. All disk updates are performed on the standby machine at the same time they are executed on the primary servers.
- Disk mirroring provides fault tolerance for the entire disk channel. In the event of a hard drive or controller failure in the primary server, NetWare will automatically utilize the mirrored copy of the data on the standby machine—even if a total server failover does not take place.
- No files need to be tagged for mirroring—this process is completely automatic.
- Files or applications do not need to be shut down or closed.
- There is no data latency (file copy and replication products introduce data latency).
- With a dedicated link used between the primary servers and standby machine, no additional traffic is added to the production network. Redundant paths are implemented between the primary servers and the standby server allowing for redundant checks before switching over.
- Dedicated links are strongly recommended but not mandatory. For those installations that do not wish to add any additional hardware or are required to use the network backbone, *StandbyServer* will operate over the existing network.
- The dedicated links can use IPX or IP protocols, which can be fully routed or bridged. NetWare 5.x/6.x also include a proprietary, dedicated protocol, VIPX, which cannot be routed but adds some fault tolerance and security features.
- *StandbyServer* includes a remote disk read blocker that disables disk reads from the standby machine. This is advantageous when the primary servers and standby machines are separated geographically or when a slower link is used as the dedicated link. In a normal mirrored environment, both disk reads and writes are sent to all the devices in the mirrored set; the disk read blocker disables the read requests to the remote devices in the standby machine.
- *StandbyServer* supports a wide variety of third-party communication adapters. Network adapters that have an IPX or IP driver for use in NetWare systems can be used for the dedicated link. This makes third-party routers, bridges, and hubs useable as part of the dedicated link.

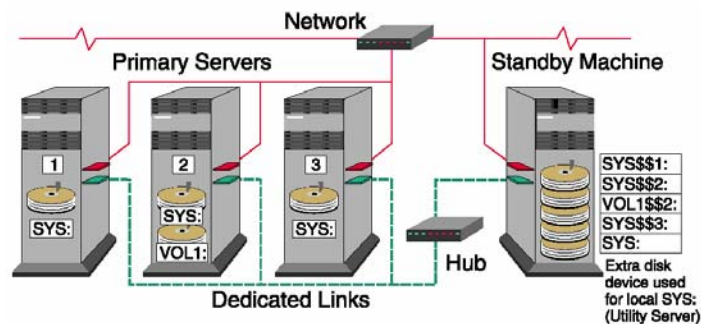
- It is possible to locate the primary servers and standby machine in different rooms, buildings, or in certain situations in different cities.
- StandbyServer is hardware independent, allowing it to function in non-standard or high performance equipment; for example, equipment that does not have an industry-standard bus.
- Server failures can be analyzed. Since the standby machine automatically fails over when a primary server fails, the failed primary can be debugged because it is left in its faulty state.
- Planned maintenance can be performed on any primary server without disrupting network services. The failover process to the standby machine can be invoked manually with little or no impact to network clients. This allows the system administrator to schedule primary server maintenance or upgrades during normal business hours.

StandbyServer for NetWare Implementations

The following diagrams show different configurations for implementing StandbyServer for NetWare:


Production Network With Dedicated Link

In production network implementations using a dedicated link, all mirrored data flows over the dedicated link. This creates two paths between the primary servers and the standby machine for redundant protection. Performance is enhanced by not using the production network for mirrored data transmission.



Wide Area Network With Dedicated Link

In a WAN implementation using a dedicated link, the read blocker option can be used to minimize WAN bandwidth usage.



Enhanced Features

StandbyServer includes several features that enhance the abilities of mirrored servers.

Utility Server

The utility server (required with StandbyServer 6.x) feature allows the standby machine to function as an independent file server while still maintaining a real-time mirror for the protected primary servers. StandbyServer can be configured with the standby machine dedicated to mirroring the disk devices of the primary servers and monitoring their status to ensure functionality. However, using the utility server feature, the standby machine contains its own SYS volume on a disk device that is not mirrored to the primary servers.

With its own SYS volume, the standby machine can have network clients of its own—or have processes that are independent of the primary servers. The only caution is that should any of the primary servers fail, the standby machine's clients are disconnected when it fails over. This allows the standby machine to share the load of network services by functioning as a backup server, print server, CD-ROM server, or some other necessary process.

The CaminoSoft SnapShotServer utility makes use of this feature to 'image' the data as it appeared at a specific moment in time, and keeps these point-in-time images on the standby machine. This allows backup engines to use the standby machine to perform backups at any time without encumbering the primary servers with the additional overhead normally associated with performing tape backups. In the case of StandbyServer, it also consolidates the backup of multiple servers in one location.

Remote Disk Read Blocking

Normally, when disk devices are mirrored, the operating system increases system performance by spreading read requests among the different devices in a mirrored set. However, when many primary servers are requesting reads from the standby machine, the standby machine can become overloaded and decrease system performance. By disabling remote mirrored disk reads, total system throughput and performance can be improved.

Throttling Mechanism

In cases where a WAN link can be used, or when using a shared data backbone for both the dedicated link and the production network, StandbyServer can implement a data throttling mechanism to control the amount of data that it sends over the network. That way it will not monopolize the available bandwidth.

Sharing Devices

StandbyServer allows the primary servers and standby machine to share devices. If the primary server has an external disk enclosure, such as a RAID subsystem, a tape changer, or a CD-ROM jukebox attached to it, the standby machine must have access to it after a failover. While some devices have a dual-connect capability and can be used automatically with StandbyServer, others require that special command files be executed before they can be accessed. StandbyServer can execute command files at failover to provide access to the external devices.

Automatic Disk Integrity Check

NetWare 4.x and later version servers automatically execute a disk integrity check, if needed, at volume mount time (volume or pool repair utilities). If the volumes do not require the integrity check, the process is skipped.

Simple Installation

The installation of StandbyServer is performed on the standby machine, which then automatically installs all code on the selected primary machines over the network.

Automatic Communication Configuration

StandbyServer automatically configures its server-to-server communication to accommodate the largest packet size available. A default packet size of 512 bytes is used to transfer data. However, to increase performance, this packet

size is automatically increased to make better use of high performance adapters and media types. For example, Ethernet can make use of 1492-byte packets; FDDI can make use 4096-byte packets or larger.

Servicing and Restoring a Failed Server

After a failover has occurred, the failed server should be repaired or replaced to restore the system to the original high-availability operation. *StandbyServer* allows the system administrator to schedule service or to debug a failed server while the standby machine is operating in its place. Although maintenance on a failed server can be performed at any time, it is important to restore high-availability functionality to the system.

When a primary server fails, the standby machine automatically takes its place. *StandbyServer* allows the standby machine to failover once—for any one of the primary servers. After a failed server has been serviced or replaced, the following steps need to be taken to restore high-availability operation:

Scenario 1 - Failed primary server, standby machine running as new primary; switching server roles back to original configuration:

1. Repair failed server
2. Resynchronize disk devices
3. Follow instructions in the user guide to restore server roles
4. Boot-up servers in their original roles

Scenario 2 - Failed standby machine, no failover takes place:

1. Repair standby machine
2. Boot-up the standby machine

System Requirements

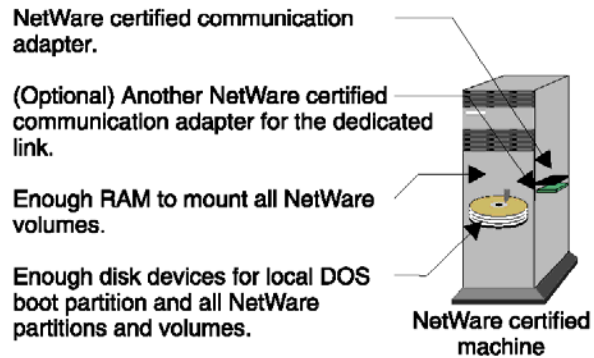
One of the advantages of *StandbyServer* is that the primary servers and standby machine do not need to be identical. They can have different processors, network adapters and drivers, bus structures, and disk devices. This section lists the minimum requirements for both primary and standby machines.

Processor	386 or better—NetWare-certified, NetWare 4; Pentium or better—NetWare-certified, NetWare 5, 6.
RAM	The standby machine must have the same or a greater amount of RAM as the largest primary server. The <i>StandbyServer</i> processes in the standby machine use 1 MB of memory per clustered primary server. Besides meeting the requirements for NetWare, <i>StandbyServer</i> requires that each primary server have approximately 128 KB of additional dedicated RAM.
Network adapters	NetWare-approved network interface cards are supported. An IPX or IP connection must exist between primary servers and the standby machine.
Dedicated link	If a dedicated link is used, it must meet the same requirements as the regular network connection. The dedicated link should consist of a high bandwidth, low latency type media; recommended minimum is 100 Mb Ethernet, but FDDI and other fiber-based adapters are acceptable as well.

If the dedicated link is to be used in a long distance or WAN environment, care should be taken in order to ensure low latency and maximize the link performance; routers and bridges should be used sparingly as they will induce latency and decrease link performance.


Primary Servers

A valid, licensed copy of NetWare 4.11, 4.2, 5.x or 6 must be installed and running on the primary servers.



Standby Machine

The standby machine should be configured much like the largest primary server, with as many disk devices as necessary to mirror all of the primary servers' disk devices. One dedicated link for each primary server can be used, or a single dedicated link adapter can be used with a hub to connect to each of the primary servers.



The standby machine must be bootable as a NetWare server, with NetWare 4.11, 4.2, 5.x, or 6 fully installed and tested. If different versions of NetWare are used on the primary servers, then the highest version used on the primary servers should be used on the standby machine. In the event of a primary failure, the standby boots up the version of NetWare previously running on the failed primary.

Other Products

CaminoSoft offers a full line of storage management and high availability solutions in addition to *StandbyServer*.

SnapShotServer for NetWare is available separately and is also bundled with *StandbyServer*. *SnapShotServer* provides the ability to capture live system data and keep this data in a frozen state, while still allowing the system to constantly change and access the data. *SnapShotServer* facilitates the backup of open files and live databases by allowing system administrators to perform backups on live data at any time. *SnapShotServer* was the recipient of LAN Magazine's product of the year award.

OFFSite Archive for NetWare replicates mission-critical data asynchronously to a remote data vault server in another location for backup, disaster recovery and permanent archiving. By utilizing asynchronous replication, slower links such as WANs and the Internet can be used to connect the source and target servers, allowing much greater distances between them.

CaminoSoft also offers Hierarchical Storage Management (HSM) solutions for Novell and Windows server environments.

For more information please contact an authorized CaminoSoft reseller or visit our website.

CaminoSoft Corporation

600 Hampshire Rd. Suite 105
Westlake Village, CA 91361

Phone: 800-889-8248 805-370-3100

Fax: 805-370-3200

www.caminosoft.com