

RAID Unbound: Storage Fault Tolerance in a Distributed Environment

Brian Ritchie

Alphatronix, Inc.

4022 Stirrup Creek Drive, P.O. Box 13978

Research Triangle Park, NC 27709

e-mail: sales@alphatronix.com

www.alphatronix.com

919-544-0001 or 800-849-2611

FAX: 919-544-4079

800-222

Introduction

Mirroring, data replication, backup, and more recently, RAID are all technologies utilized by corporate America to protect and ensure access to critical company data. IS managers have taken comfort in knowing that critical data was being copied and safely stored for future access in the event of an equipment failure, operator or user error, or even worse, a local disaster. If one of these events were to occur, this critical data could still be transparently accessed, or at least recovered, and operations would continue. Or would they?

A whole new set of problems have arisen as a corporation's data becomes more and more geographically distributed. Do conventional protection techniques — mirroring, data replication, RAID, backup etc. — truly provide the level of data protection and data accessibility needed under this changing environment? The answer to this question is — probably not. Each of these technologies provides important benefits; but each has failed to adapt fully to the realities of distributed computing. The key to data high availability and protection today is to take these technologies' strengths and “virtualize” them across a distributed network.

Traditional Backup and High Availability Methods

RAID and mirroring offer high data availability, while data replication and backup provide strong data protection. If we take these concepts at a very granular level (defining user, record, block, file, or directory types), and then liberate them from the physical subsystems with which they have traditionally been associated, we have the opportunity to create a highly-scaleable network-wide storage fault tolerance. The network becomes the virtual storage space in which the traditional concepts of data high availability and protection are implemented without their corresponding physical constraints. Let's look at the evolution of these technologies.

The concept of RAID has existed for several years, giving users the ability to copy and/or stripe data to an array of disks. Because of the redundant design, data remains accessible even if an individual disk should fail. But what if the server fails? Because the disk array is located in a single physical location, its data is vulnerable to a “single point of failure.”

Using RAID, no alternate “safe” locations from which to retrieve data exist, making it incomplete as a reliable, uninterrupted storage fault tolerant solution.

Disk mirroring, like RAID, enables IS managers to “duplicate” or “mirror” critical data to a second disk so that the data can later be retrieved in the event of a primary disk failure. This method also poses several drawbacks. For one, mirroring is not efficient; the entire contents of the disk must be duplicated. Neither RAID nor mirroring offers the level of granularity needed to define which users, records, blocks and/or files receive top-level protection. Additionally, mirroring, like RAID, is constrained by physical location and vulnerable to a single point of failure.

Backup and data replication technologies have been used for some time to protect mainframe and workgroup level data. Historically, they ensured that data was always available, but how do you back up a file system when it is hundreds of gigabytes to terabytes in size, or when you’re collecting hundreds of megabytes of data per day? The window of time available to perform these tasks is no longer enough. This magnitude of data can’t be backed up during regular business hours because of the already high level of network traffic versus the network pipe size available.

Emerging SFT Technology

Although there are still benefits to RAID, mirroring, data replication and backup, today’s storage needs demand what these technologies can’t provide — high data protection and availability across the entire enterprise. Administrators are looking to a new generation of software to take high data availability and protection concepts one step further. A new concept of network-wide storage fault tolerance (SFT) has emerged, which utilizes the entire network’s storage resources, giving administrators the ability to store multiple copies of information at multiple sites in the enterprise, even at remote storage vaults.

SFT technology evolved from a need to ensure that data was consistently and readily available to key users. If key users cannot get consistent and immediate access to their critical data, then individual productivity suffers, meaning loss of money to a company. SFT software sends key user data not to a single, same site location (as in RAID and mirroring), but to various storage devices located throughout a company’s computing environment, eliminating any single point of failure (Figure 1). In this way, it offers an enterprise-wide level of high data availability and data protection rather than traditional subsystem-specific security.

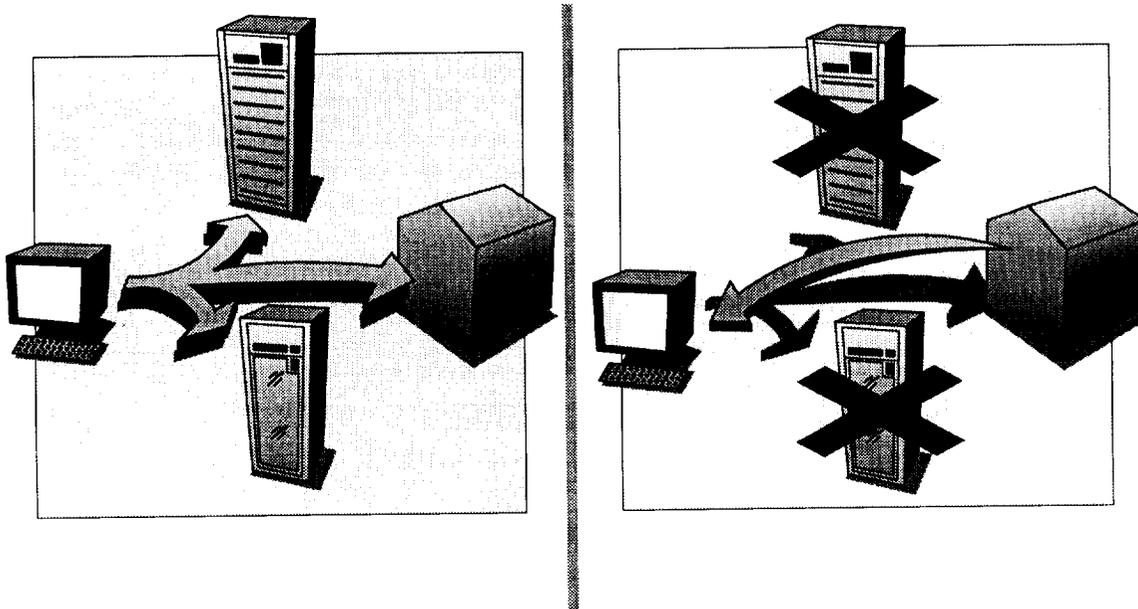


Figure 1. Data high availability and disaster recovery capability for key user data using storage fault tolerant software.

SFT software gives data control back to the administrator and saves time in the process. Possible storage destinations (such as RAID or optical and tape libraries) are defined by administrators based on various criteria including the type of data being stored. The most valuable data (for example scientific information or mathematical calculations) can be sent to two, three or more local storage centers or to a remote site that's considered "most secure." The software recognizes every storage site on the network and marks each site based upon its storage performance. This process ensures that whenever data must be accessed from a secondary source, it is retrieved automatically and immediately from the highest performance device available, regardless of where the device is located. When data is requested, the software automatically and transparently retrieves it without any "downtime" associated with restoring "save sets" to a drive. Because the software is managed centrally, the labor associated with individually managing distributed RAIDs or disk mirroring systems no longer exists. And because the data is always accessible, replication and/or backup processes may be eliminated, or performed much less frequently than before.

The "Virtual" Single Storage Environment

Unlike traditional technologies, Storage Fault Tolerant systems work to make a distributed environment look, act and feel, to the administrator, like a single logical environment (Figure 2).

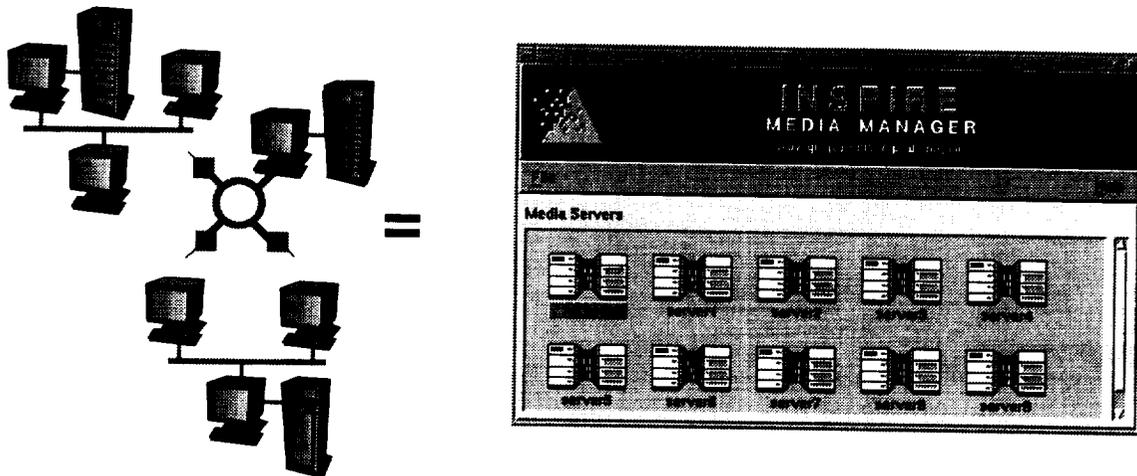


Figure 2. Creating a "virtual single storage environment" from multiple, distributed physical systems using storage fault tolerant software

This is important for several reasons. First, administrators are able to see every server and peripheral device connected to the network and can create a higher-level storage fault tolerant environment, across this network, than was possible before. Second, because administrators can select which users and what data get top protection priority, new users and data types can be easily defined and added from a central location. Third, all storage resources (across the enterprise) are visible through a single graphical user interface (GUI) for easy viewing, manipulation and management by the administrator. This is crucial in determining where, when, and to how many sites data should be sent. Fourth, administrators are able to balance their data load across a heterogeneous and distributed network from one common GUI.

SFT technology makes use of RAID or mirroring hardware already in place and offers a whole new set of capabilities and whole new level of control for systems administrators. It makes the need to purchase new hardware for mirroring or a new RAID box each time one fills up obsolete because the software makes every storage destination on the network available. Using SFT software affords administrators, sitting at one location, the ability to see each new server and every additional storage device, as they're connected, so that he/she can ensure the safe storage and availability of all crucial data across the enterprise; it enables a geographically distributed environment to look and act as a single logical environment, providing network-wide storage fault tolerance to key users' data.

Network-wide storage fault tolerant systems also have the potential to eliminate most, if not all, of the burden placed on traditional backup procedures. By migrating data to multiple, redundant locations, SFT systems are already, today, capable of eliminating the need to back up the bulk of a company's data. As SFT systems continue to evolve, they have the potential to entirely replace traditional backup procedures with enterprise-wide data high availability, scalable to the requirements of today's distributed computing environments.