

# **SANs for High Availability Systems**

**Whitepaper**

**Erik Ottem**



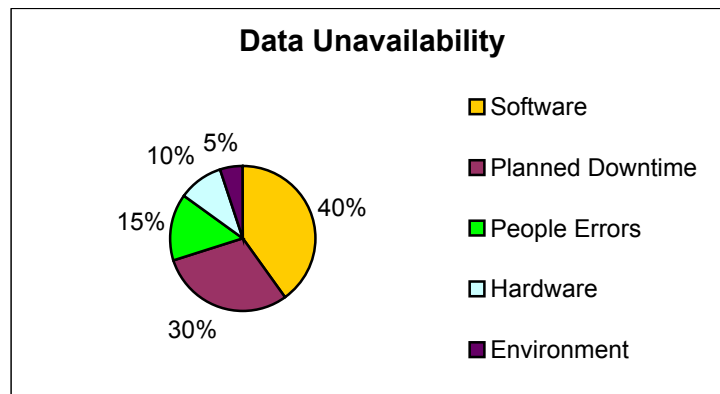
# SANs for High Availability Systems

As always, in today's business environment, time is money. However, it is no longer an eight-hour a day, five-day-a-week affair. Today's businesses run twenty-four hours a day, seven days a week, fifty-two weeks a year. Whenever data is not available, it costs businesses money in terms of lost transactions, lost management opportunities and lost customer relationships. While this has long been the case for large global enterprises, it is now also increasingly true for even small and medium sized organizations - lack of access to data at any time is painful, often disruptive and occasionally fatal to such organizations.

In this white paper, we examine the causes of data unavailability, the costs and solutions for High Availability (HA) systems and the applicability of the Storage Area Network (SAN) architecture to this issue.

## The Causes of Data Unavailability

There are several factors that affect downtime or unavailability of data as shown in Figure 1. The results are not surprising:



**Figure 1 - Causes of Data Unavailability**

While most data unavailability is due to unexpected systems outages, such as hardware and software problems, a significant amount of data unavailability is due to planned downtime. Altogether, systems outages of all kinds account for about 80% of the total. Environmental issues such as power outages, fires, floods, or people errors such as the man on a backhoe digging in the wrong place cause the remaining 20% of the downtime.

There are many ways to manage these problems. For example, data unavailability due to people errors and software problems can be managed with improved procedures and processes, but nevertheless will probably not be completely eliminated. Hardware problems can be managed with redundancy and other techniques, but likewise cannot be completely eliminated. Things will break, regardless.

Certainly, you can increase protection by increasing the safeguards, but such protection can be expensive. The question is, how much guaranteed data availability can you afford? The answer is "It depends."

## High Availability for Business Objectives

Given twenty-four hour, seven-day week utilization, 99% availability translates to data being unavailable the equivalent 3.9 days per year. Critical applications would generally not accept this amount of downtime, whereas some applications would find this acceptable. Think of your manufacturing management application not available for almost 4 days a year. Could your business afford it? Probably not.

Moving to 99.9% availability translates to a more acceptable 8.8 hours of downtime in a year. But it might also cost ten times as much for such a system. Is it worth it?

The reason that an application is important to a business is the economic value per hour it has to that business. For example, at a major financial center it might cost \$2 million per hour for such downtime, but non-critical applications may cost only \$10,000 per hour or less.


Thus, there is a gradient of trade-offs that must be assessed. This, in turn, has led to a gradient of High Availability options, each with its own cost/benefit ratio.

## High Availability Objectives

The traditional solution to data unavailability has been High Availability (HA) systems. This is a catchall phrase that covers many different approaches to increasing data availability. But generically, a HA system is designed to do several things:

- First, it must reduce downtime due to planned maintenance.
- Second, it must be resilient to unexpected hardware and software failure.
- Third, it must also deal with problems such as disaster recovery in case of fire, earthquake, flood, or even the errant backhoe operator.

Traditionally, the solution to a HA requirement has been the fault tolerant system, a term that is often applied to a hardware configuration that allows redundant or otherwise protected components to failover (or switch) to a new component or set of components so that downtime is minimized. A simplified example is given in Figure 2. Generally, these cost two to ten times as much as the equivalent non-fault tolerant system. But the potential savings from achieving improved data availability can far exceed this additional up-front cost.



**Figure 2 - A Simplified Fault Tolerant System**

A fault tolerant system consists of three layers:

1. Protection for the server with clustering and alternate pathing techniques that allow applications to migrate to functioning server platform in the event of server failure.
2. Redundant networking for robust connections and intelligent data and storage management.
3. Protection for the storage with RAID (redundant array of inexpensive disk) technology, data mirroring and/or replication.

The increasingly sophisticated development of server clustering for availability has created a requirement to assure that data can follow the application as it fails over to alternate sites. There is now a solution to this quandary. It is the Storage Area Network.

## The Storage Area Network

The Storage Area Network (SAN) was developed in response to the requirements of High Availability of data, scalable growth, and system performance. Not only does it increase the availability of data to the cluster but throughout the enterprise system.

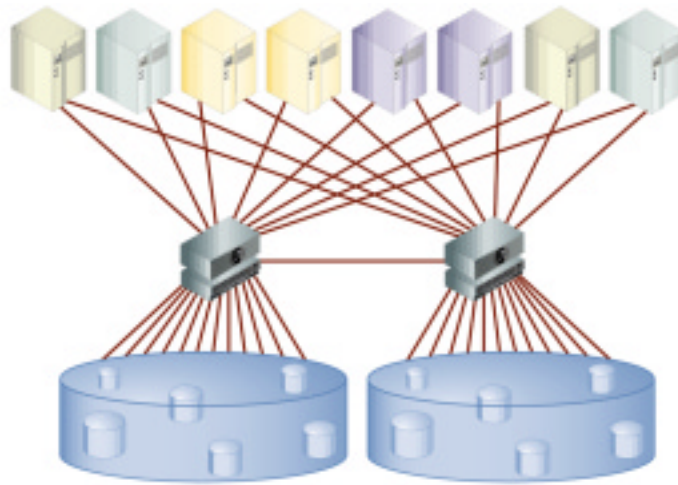
SAN employs high performance interconnects like Fibre Channel (FC) technology today to allow innovative ways to protect and distribute important data. Of the basic improvements of FC over the traditional SCSI connections the benefits of speed, distance and scalability are important in building fault tolerant configurations. FC can provide links of up to 120km to allow distance connections that appear to be local to the applications. And as we shall see, there is safety in distance. Redundancy of connection in the SAN can guard against the errant backhoe operator, fire and other local disasters. Protection can be provided by creating mirrors of data in two buildings on a campus connected by high speed, scalable, and long distance connections that appear local.

Even with all of the critical benefits of SANs, until recently SANs were considered difficult to implement, expensive and costly to manage. All that has changed with the advent of a new breed of SAN Appliances. SAN Appliances introduce a new concept called Network Storage Pools (NSP) that separate physical storage resources from logical volumes that servers and clients use. This storage “virtualization” provides the key enabling technology to truly provide host and platform independent HA configurations. In addition SAN Appliances provide key data protection technologies such as snapshot, remote and local data mirroring and massive scalability. Combined with component and network topology, high availability with no single point of failure is achieved in a properly configured Network Storage Pool SAN.

The SAN and NSP architecture can provide new options for reliability in a distributed environment. SAN can greatly lower the cost and improve the scalability of HA solutions, thus enabling a wider range of applications to benefit. We will review a series of possible topologies using SAN architecture to illustrate the SAN’s flexibility for HA designs, starting with a small company occupying a single building and proceed up to large corporations with buildings spread over a metropolitan area.

### The Single Building


The simplest scenario for a HA configuration is one that it is contained in a single building. This arrangement is typical of small companies. The SAN system shown in Figure 3 is a simple extension of the redundant system shown in Figure 2, except all the storage is now Fibre Channel based and thus available to all the systems on the SAN. In this case, both servers are now active, allowing a larger number of clients access.



**Figure 3 - HA SAN System with SAN Appliance and redundant SAN in one building**

This structure protects against server failure by clustering applications, protects against data loss with use of RAID technology and protects access to data by use of a redundant network topology in the SAN. Still, things can happen to impact operations. Fire, loss of power in the building or other local problems might cause considerable downtime unless a redundant remote stand-by site exists with mirrored data. Clearly, there is much to be said for a geographic distribution to improve availability.

However, even here, the individual components of a SAN HA configuration such as shown in Figure 3 can be spread out through that building, either on opposite ends of the same floor, or on different floors. However, the advent of FC based SANs provides additional options since their distance capability allows a reconsideration of what comprises a local connection



**Figure 4 - SANs in a campus environment**

## The Campus Cluster

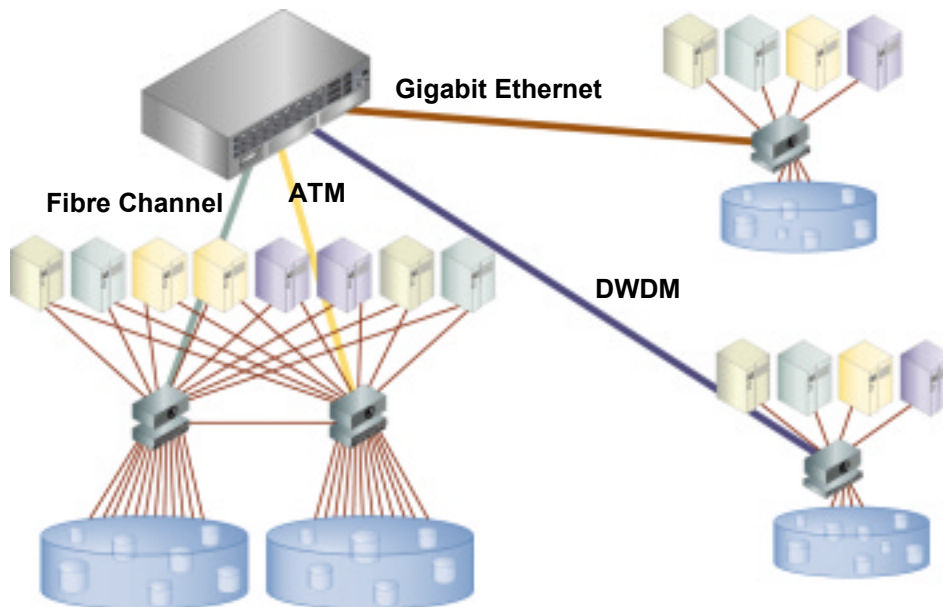
A campus cluster is usually found where multiple buildings exist in a private campus so that right of way to lay cable exists. A campus cluster usually consists of buildings in the same area, often no more than a few hundred meters apart so employees can walk among the buildings. This environment provides additional disaster planning options over single building configurations. For instance, fire loss is more easily contained.

An additional feature of the leading Fibre Channel switches like the Capellix is that the FC link can be up to 10 kilometers with single mode fiber optics and as much as 120 kilometers with special GBICs. This makes these switches not only a natural for the campus cluster but also for the metro cluster as well. Although the two buildings may be several hundred meters (or kilometers) apart, the storage and servers still look local to the applications due to the strengths of Fibre Channel connections. This allows the systems to mirror the data between buildings without interfering with application operation. The fact that both writes may occur without causing a time out for the application is known as synchronous operation- both writes occur in the normal operation of the system. This is possible due to the 10km distance capabilities of Fibre Channel allowing local devices to be more distant than previously possible.

## The Metro Cluster

While the campus cluster affords protection from fires, other natural hazards such as earthquake, tornado, and flood could adversely affect an entire campus. To gain even greater protection than even a campus environment, protecting select data in a metropolitan area provides greater protection. Although metro clusters increase the costs of management and inter-office travel, they also reduce risks to most natural hazards.

As noted, the SAN FC technology allows up to 10-kilometer distances, and with special additional hardware, up to 120 kilometers or more. When we go beyond FC distances a different kind of copy is required, that of asynchronous copy. Asynchronous copy assumes that the latency, or delay, of writing to two sites at a distance that the application will notice- and degrade performance- is undesirable. To address the latency of longer distances asynchronous copy makes one copy of the data and then releases the application to go to the next step and then writes the data to the remote location. This allows the data to be protected without impacting the performance of the application. There is a price, however. In the event of disaster at the main site, the remote site with the mirrored data will lose all data that was in the process of being copied to the mirrored site. This means that before the operations can be resumed in the remote site of the metro cluster a point in time will be selected and the application restarted. This, and the higher cost of the longer distance links, are why this technology is typically used only for selected information and applications.



**Figure 5 -** The Metro Area SAN



The failover capability of the Metro Area SAN can also be extended using different technologies to cover any two locations in the world. The longer the distances, the greater the loss of data during restarting the application in the remote site, but the greater the disaster protection potential. The technologies to connect over long distances would include Gigabit Ethernet, Dense Wave Division Multiplexing (DWDM), or Asynchronous Transfer Mode (ATM). In many cases a backbone of one of these technologies already exists and is used for the additional traffic of data and application mirroring between two existing sites of the enterprise.

## Data Availability Management

The typical way that data is mirrored requires identical configurations in the primary and mirrored site. With the addition of storage virtualization as found in a SAN Appliance that provides Network Storage Pools comes the ability to mirror on dissimilar devices. This means that as long as there is capacity in both sites the data can be mirrored, and the choice of storage devices is not dependent upon the physical storage since all of these are virtualized into a networked storage resource and is not dependent on brand or even kind of storage provided there is equal available capacity.

A further advantage of SAN Appliances as a way to manage storage is the consolidation of the storage resource and the storage networking so that from a single framework data can be allocated, replicated, and the network managed for greater personnel and hardware efficiency.

## Summary

In today's global economy, data availability is more important than ever. As enterprises and application grow, the ability to grow a data infrastructure with it is critical. The intelligent information infrastructure SAN products from Gadzoox Networks enable a wide variety of HA configurations that will grow with your enterprise and protect your investment in subsequent configurations. Data can be protected more easily than ever before, and the means to implement HA configurations are available today at very reasonable costs. Just decide how much protection is required and how much budget is available for the project, since the longer distance often translates into higher costs.

High availability increasingly depends on the intelligent infrastructure of a redundant Networked Storage Pool SAN configuration to assure that as clusters keep applications available, and mirroring and parity RAID protect data, that redundant intelligent SAN architecture will provide the structure to provide a highly scalable, high performance, and highly available system design.