



What is a SAN?

The Basics: What is a SAN?

It sounds simple enough: A SAN is a Storage Area Network. And that answer may get you a point on a test! But in the real world, many of us from the smaller markets would be lost if we had to come up with a shopping list to start building our first SAN.

This chapter will answer all your questions. We'll discuss some goals for data storage systems, storage alternatives (what's not a SAN?), introduce the SAN, and then dig into some details about building that first SAN. By the end of this chapter, you'll have no problem designing a basic storage area network and drawing up that shopping list.

We'll also discuss appropriate uses for the technology. After all, storage area networks are just another tool. When do you use this tool versus some other tool?

And, finally, we'll talk about the management software that makes everything work. Just as any other network needs a network operating system, SANs need some kind of central management software.

Putting Storage in Perspective

Before we get too far along, we need to look at our goals for a storage system. What do we want from a storage system? How much of that do we get without a SAN? Is there anything we *can't* get without a SAN?

We propose several goals for a disk-based storage system. Note that we're not including magnetic tape here. Tape systems are great for backups, but don't fill the bill for always-available data usage. Optical media are also not appropriate, due to both speed and capacity.

We believe you are looking for the following characteristics in an "ideal" disk-based storage system:

Reliability and Availability

There is nothing we want more from our disk storage than this: just work. For most of us, the only time we use the term "mean time between faults" is in relation to disk drives.

Fault Tolerance refers to the ability of a storage system to continue operating after one of the drives fails. For example, in a RAID 5 array, one drive can fail and the array will continue working. If the system has hot-swap drives, the bad drive can be replaced and the system will have zero downtime.

High Availability goes beyond fault tolerance within an array. With databases stored separately from servers, two or more servers can be configured as a cluster, with failover options. Thus, if either server fails, its service is assumed by another server. In such a case, an entire critical function (e.g., email) can have high availability.

Speed

As architectures change, we are constantly examining how to get the most speed from our physical media—disks. As we all know, data stored in memory (RAM) can be accessed very quickly. But hard drives have always created a bottleneck because they require us to physically get data to and from a device.

While at some level we make comparisons like "SCSI vs. SATA," the architecture of moving data from the storage device to the servers (or workstation) is sometimes the more important variable.

Effect on the network. If we use the local area network (LAN) to move data, we may put a significant strain on the network. This depends