

## Re: How does a SAN connect to a computer?

---

*Source:*

<http://coding.derkeiler.com/Archive/Delphi/borland.public.delphi.non-technical/2008-07/msg01993.html>

---

- *From:* "Brion L. Webster" <[brion.webster@xxxxxxxxxxxxxxxxxxxxxxxxxxxx](mailto:brion.webster@xxxxxxxxxxxxxxxxxxxxxxxxxxxx)>
  - *Date:* 22 Jul 2008 09:51:28 -0700
- 

Ian Boyd wrote:

i've heard of this mythical thing called a SAN (Storage Area Network). It's not a hard drive. It's not a RAID array. It's not a fileserver on the network. It's not a fast connection to a bunch of fast hard drives.

<snip>

Does a SAN connect to a pc using something like external SCSI (whether it be fibre or not), or external SATA? Does it require a custom PCI expansion card to do connection?

OK, lots of speculation. For those who don't have their own, and might be interested:

A Storage Area Network, or SAN, is a specialized chassis that holds a disk write controller and a bunch of hard drives. Kind of like a blade server requires a special chassis to hold the blades, and provides backplane communications between them, the SAN chassis holds 12-16 hard drives (usually), the write controllers (usually two for redundancy), power supplies (usually two for redundancy), and provides a backplane for communications between them.

The first piece of a SAN you purchase will be the most expensive piece. The initial chassis, with the write controllers, has most of the logic, the on board software, licensing, etc. Often, some number of the initial hard drives in the first chassis are reserved for the on board operating system. Sometimes you can use a portion of those disks for your own storage, but it's usually not recommended.

This is where the near-infinite configuration options come in. Most vendors offer either SAS or SATA drives. SAS drives are the current equivalent of SCSI drives we used to order for servers, while SATA drives are the what IDE drives used to be for workstations. SAS drives are fast, "enterprise class", and dual-ported. SATA drives are slower, single-ported, but offer vastly more storage. For primary file storage, you'd opt for SAS drives. For secondary disk-to-disk backups, or terabytes of video, order SATA drives.

If your first chassis can't hold enough disks to meet your anticipated storage needs, you backpack on another chassis. Say you ordered the first one as SAS drives. You might order the second one as SATA drives, for expanded storage, or second tier storage. You may decide to order four SAS chassis, and three SATA chassis. The chassis themselves typically connect with 2-4 GBPS FibreChannel connectors, depending on how current your SAN unit is. Again, the chassis backplane is responsible for providing enough I/O for all this traffic. Each expansion chassis will also have redundant write controllers, but without the total O/S software that came in the first chassis, so they're usually about half the cost of the initial unit.

## Re: How does a SAN connect to a computer?

This is what makes a SAN so tantalizing to server admins. How do you add more storage to a SAN? Order a new chassis, stuffed with disks, and snap it on. Voila – more room. Go into the management software, and either add those disks to an existing server's assigned storage pool, or tell the SAN to give those disks to a different server that needed more storage.

The other thing a SAN can do is automatically swap in a hot spare from anywhere in the SAN to any failed disk. Have 150 disks in 40 different RAID groups? 1 hot spare can be put into any one of those raid groups, no need for manual swapping around until you get to it.

How do you add more storage to a traditional server? Shut the server down, open the case, pray there are some free power connectors and data cables, add some more hard drives, power back on, double check the BIOS to make sure they're recognized, initialize the disks, format them, set up ACL permissions, shares, etc. With the SAN, you never have to crack your server, it stays running, the ACL stays in place, the shares are still there, the total capacity simply got bigger.

OK, back to the original question – how do you connect a SAN to a computer? There are currently two methods – FibreChannel or iSCSI.

FibreChannel uses a different card setup than most folks are used to. Again, FibreChannel offers 2 or 4 Gbps speed. Usually you have to order 2 cards per server, in addition to the ethernet nics they have. Somehow you get the FibreChannel connections into the SAN, I don't have direct experience with how you connect multiple servers to a SAN in this configuration.

With iSCSI, everything is done over regular TCP/IP ethernet. The big drawback is the speed – most vendors are stuck at 1 Gbps for data. Ideally, you'd have 4 Nics per server – 2 for "public" communications, and 2 for your iSCSI data connection. I'm getting by with legacy servers that only have two network cards, 1 for the public network, 1 for iSCSI. With a Gigabit switch, you set up a couple of Virtual LANs in the switch.

In both FibreChannel and iSCSI, it is vitally important that you reduce collisions between your public traffic and your SAN traffic, and between the two write controllers in your SAN. So – for my iSCSI solution:

Public NIC goes into a port on the switch that talks to the rest of my LAN. All client workstations talk to the server through that port. All clients access data shares over that port, that IP. That NIC is assigned a typical IP in our corporate addressing scheme.

iSCSI NIC Server 1 goes into a port on the switch that talks to one set of ports on my SAN. Server 2's iSCSI NIC goes into a port on the LAN with a different VLAN that talks to two different ports on my SAN. That way they avoid contention, I can expand this indefinitely. Each has a primary port on one write controller, and a secondary port on the other write controller. The iSCSI NIC's have manually assigned IP addresses that are non-routable, do not match the corporate standard, and due to the Virtual LAN on the switch, should never, ever, be able to be seen by anyone else on the corporate network.

I have chosen to leave operating system boot drives in my servers. It makes it a lot easier to troubleshoot things when they go wrong, and makes the SAN an optional extension of the servers. There are two versions of the Microsoft iSCSI initiator you can download – one is for booting off of an iSCSI SAN, one is for "normal" use.

You install the Microsoft iSCSI initiator on your Microsoft server. Depending on the vendor, you may have to install their drivers or management software as well. On your SAN, you carve up your available disk space, decide how you want the raid to work (RAID 1, 0, 10, 3, 5, 6, etc.), and assign it to the servers. Once it's

## Re: How does a SAN connect to a computer?

assigned to the servers, you go into the server, pull up Disk Management, scan for new drives, and go from there. At that point it looks like an internal hard drive to the server. The first time, you have to format it, and do all the NTFS ACL/Sharing work, but from there, you're done with internal server work.

Most of the time, you would not use a SAN for a workstation class computer. There's no reason you couldn't, but how often do you put a bigger hard drive in a workstation? You certainly wouldn't purchase a SAN for this kind of solution – just get an external hard drive.

Phew! Treatise done! ;-)

—

–Brion

There's no such thing as 'one, true way;'

– Mercedes Lackey

.