

## 9.1.5 SAN Facts

---

In a storage area network (SAN), multiple servers are configured to share a common storage device. In this lesson, you will learn about:

- How a SAN works
- SAN components
- SAN technologies
- Jumbo frames
- Clustering a SAN

### How SAN Works

The following are facts about how a SAN works:

- A SAN is a dedicated network separate from the production network.
- A SAN connects file servers with storage devices.
  - Multiple servers can connect through the SAN fabric to share the same storage device.
  - A single server can connect with multiple storage devices.
- Storage devices are called SAN targets.
- Servers are called SAN initiators and run initiator software to communicate with SAN targets.
- The storage device appears to be connected to a server like any other directly attached storage.

### SAN Components

There three main components that make up a SAN:

<b>SAN Component</b>	<b>Description</b>
Host	Hosts are not technically a SAN component. <ul style="list-style-type: none"><li>▪ Hosts connect to the SAN using a NIC or a host bus adapter (HBA).<ul style="list-style-type: none"><li>▪ An HBA is a circuit board that is inserted into a computer motherboard, just like a NIC.</li></ul></li></ul>

Storage	<p>Storage devices are typically:</p> <ul style="list-style-type: none"> <li>▪ Disk arrays</li> <li>▪ Tape libraries</li> <li>▪ Optical jukeboxes</li> </ul>
SAN Fabric	<p>The SAN fabric consists of:</p> <ul style="list-style-type: none"> <li>▪ SAN hubs</li> <li>▪ SAN switches</li> <li>▪ Network cables</li> </ul> <p>The term <i>fabric</i> is generally used identify multiple communication paths between servers and storage devices. This creates a redundant fabric rather than a single thread. Also, fabric serves as a reminder that the storage area network is different from the production network.</p>

## SAN Technologies

Two dominant SAN technologies are Fibre Channel (FC) and Internet Small Computer Systems Interface (iSCSI). Two more technologies, Fibre Channel over Ethernet (FCoE) and Infiniband, are being implemented more and more.

SAN Technology	Description
Fibre Channel (FC)	<p>A switched FC SAN uses fiber optic cabling, network adapters, and switches to build the SAN fabric. To build a redundant FC SAN:</p> <ul style="list-style-type: none"> <li>▪ Install two FC HBAs in each server.</li> <li>▪ Deploy two FC switches by connecting each server to each FC switch. <ul style="list-style-type: none"> <li>▪ Connect one HBA to the first FC switch.</li> <li>▪ Connect the other HBA to the second FC switch.</li> </ul> </li> <li>▪ Deploy shared storage devices. <ul style="list-style-type: none"> <li>▪ Use RAID devices containing multiple hard disks and two FC HBAs.</li> <li>▪ Connect one HBA to the first FC switch.</li> <li>▪ Connect the other HBA to the second FC switch.</li> </ul> </li> </ul>

When physically configured, FC uses SCSI protocols to access and manage the shared storage:

- The storage devices on the FC target appear to the operating system on the initiators as if they are locally attached SCSI hard disks.
- The initiators send SCSI commands over the SAN fabric to manage the remote storage on the target.
- Each server connected to the SAN is an initiator and can send disk I/O SCSI commands to the shared storage device.

Internet Small  
Computer  
Systems Interface  
(iSCSI)

An iSCSI SAN uses network protocols that encapsulates SCSI commands within IP packets and transmits them over a standard Ethernet network. To build a redundant iSCSI SAN:

- Install two dedicated Ethernet NICs in each server. These NICs can have either fiber or twisted pair connections.
- Deploy two dedicated Ethernet switches by connecting each server to each switch using the appropriate fiber optic or twisted pair cabling.
  - Connect one NIC to the first switch.
  - Connect the other NIC to the second switch.
- Deploy the shared storage devices.
  - Use RAID devices with two Ethernet NICs.
  - Connect one NIC to the first switch.
  - Connect the other NIC to the second switch.

As with FC, iSCSI storage devices are targets, while iSCSI servers are initiators.

- The iSCSI initiator connects to and communicates with iSCSI targets using port 3260 by default.
- The storage devices on the remote iSCSI target appear to the operating system on the iSCSI initiator as locally-attached hard disks.
- The iSCSI initiator sends SCSI commands within IP packets to the iSCSI target over the network.
- The iSCSI target redirects the SCSI commands to its locally-attached storage devices.

An iSCSI SAN should never be implemented on a

	<p>standard production network. The performance of the SAN will be heavily impacted. As a best practice, implement a dedicated network infrastructure (cabling, switches, and NICs) that is restricted to only iSCSI SAN traffic.</p>
<p>Fibre Channel over Ethernet (FCoE)</p>	<p>FCoE was developed to lower the cost of Fibre Channel solutions.</p> <ul style="list-style-type: none"> <li>▪ Like iSCSI, FCoE uses standard Ethernet switches and physical connections.</li> <li>▪ Unlike iSCSI, which can run on slower Ethernets, FCoE requires 10 gigabit Ethernet.</li> </ul>
<p>Infiniband</p>	<p>Infiniband was designed for high-performance supercomputers.</p> <ul style="list-style-type: none"> <li>▪ Infiniband features high throughput and very low latency.</li> <li>▪ Infiniband can be used in a SAN as well as between clustered computers.</li> </ul>

## Jumbo Frames

Since SANs move large amounts of data, enabling jumbo frames in an iSCSI or FCoE SAN can increase the performance of the SAN.

- Most Ethernet switches that support gigabit Ethernet are configured with a maximum transmission unit (MTU) size of 1500 bytes
  - The MTU specifies the maximum size of the payload in an Ethernet frame.
  - An MTU of 1500 is the maximum allowed by the IEEE 802.3 standard.
- Jumbo frames have payloads larger than 1500 bytes.
  - The size of jumbo frames varies by vendor and device.
  - Many network devices support payloads as large as 9000 bytes.
- To enable jumbo frames in a SAN, change the MTU setting on all servers, SAN fabric devices (SAN switches) and storage devices.

## Choosing a SAN Technology

Several factors may influence the choice of SAN technology:

- Cost wise, iSCSI and FCoE are cheaper, but Fibre Channel and Infiniband perform better.
- iSCSI and FCoE are easier to implement than Fibre Channel and Infiniband, which require specialized hardware and knowledge.

- iSCSI and FCoE are not as fast as Fibre Channel and Infiniband.
  - iSCSI and FCoE speeds will increase as Ethernet speeds increase.
- Infiniband has a short distance limitation of about 300 meters, and Fibre Channel has a distance limitation of 10 kilometers.
- iSCSI and FCoE, can route IP packets over many networks, accommodating longer distances.
  - Performance is impacted, as the longer distances increase latency.

This table compares SAN technologies.

<b>SAN Technology</b>	<b>iSCSI</b>	<b>FC</b>	<b>FCoE</b>	<b>Infiniband</b>
Cost	Least expensive	More expensive	Less expensive than FC	Most expensive
Specialization	No specialized hardware or knowledge	Requires specialized hardware and knowledge	No specialized hardware, but requires specialized knowledge	Requires specialized hardware and knowledge
Speed	Limited to Ethernet speeds			
Faster				
Limited to Ethernet speeds				
Fastest				
Distance	Accommodates longer distances, but introduces latency			

Distances up to 10 km
Accommodates longer distances, but introduces latency
Short distances less than 300 meters (Note: Infiniband devices exist that accommodate longer distances, but with increased latency.)

## Clustering a SAN

The servers in a SAN share common storage devices. They are commonly deployed in a clustered configuration.

- With clustering enabled, multiple SAN servers can be grouped together to provide a degree of fault tolerance.
- All of the cluster data exists on the shared storage; there is no need to replicate data between servers.
- To users on the network, the cluster appears as a single file server.
- If one of the servers in the cluster goes down, another server immediately takes over to provide access to the files on the shared storage device.

Clustered SAN servers can also be configured to load balance.

- A bottleneck can occur if one server in a cluster is overloaded with requests to access the SAN's shared storage.
- A load-balanced cluster can divide up the work and distribute it between multiple servers in the cluster.
- Load balancing can dramatically increase the performance of the SAN.