

ETHERNET VERSUS FIBER CHANNEL

Ethernet (and TCP/IP) is the number one used technology these days for communication between devices. But for storage, the dominant technology in a datacenter often is Fiber Channel (abbreviated to FC).

Fibre Channel

FC is a network standard to allow hosts (servers) to communicate with storage devices. By itself, it's completely separate from Ethernet. A storage network switch is not the same as an Ethernet network switch. There is one notable exception to this rule: the Cisco Nexus 5548UP and 5596UP have switchports that can be run in either Ethernet mode, or Fibre Channel mode, but not both modes at the same time. There's also no communication between both types of ports possible, as the protocols are incompatible.

One name you'll hear when talking about storage networking is Brocade: the most prominent vendor of storage networking hardware. Also, a bit of information about the name Fibre Channel: originally, FC's only transport medium was fiber, but these days twisted pair copper wire is also possible. That's the opposite of Ethernet, which originally ran only on copper wires and now can be used on fiber as well.

Reliability

Another important difference to remember between Ethernet and FC is the reliability: FC was designed with perfect reliability in mind. Not a single frame may be lost, and frames must be delivered in order, just like they would from a local attached storage device. FC switches even signal when they're congesting to other devices, so these devices stop sending frames, instead of dropping frames. This in contrast to Ethernet, which will just start dropping frames when congested, relying on upper layers (like TCP) to make sure everything keeps working.

SAN versus NAS

Some people think a Storage Area network, or SAN, is similar to a Network Attached Storage disk, or NAS. This is not true: a NAS provides access to files, a SAN provides access to raw storage. It also doesn't show up as a network drive in the operating system but as a local attached drive, and it is treated that way too.

Layers and command set

Wikipedia mentions that Fibre Channel does not follow the OSI layer. It's true but not completely: a FC frame be divided into layers. The biggest difference is layers 5 to 7 of the OSI

model are missing, as FC is raw storage data transport and not related to a particular application. I'll quote the layers from [Wikipedia](#):

- ◆ FC4: Protocol Mapping layer for protocols such as SCSI.
- ◆ FC3: Common Services layer, a thin layer for encryption or RAID redundancy algorithms.
- ◆ FC2: Network layer, consists of the core of Fibre Channel, and defines the main protocols.
- ◆ FC1: Data Link layer, which implements line coding of signals.
- ◆ FC0: PHY, includes cabling, connectors etc.

On FC4, SCSI or [Small Computer System Interface](#) is commonly used. SCSI is a command set to communicate with storage devices. It's the same command set used between a computer and a local attached SCSI drive (like a [SAS](#) drive). FC2 is the network layer and somewhat relates to OSI layer 2 and 3. A SAN is one flat network, best compared to a layer 2 subnet. There are discussions about whether FC is switching or routing, but it's a bit of both really. Personally, I use the term 'Fibre Channel switching' because it's a flat network. On the other hand, FSPF or Fibre Channel Shortest Path First, is commonly referred to as a routing protocol. Also, it doesn't use MAC addresses, but World Wide Names (WWNs) to identify source and destination nodes, which are hexadecimal numbers just like MACs.

Bandwidth

FC speeds aren't in multiples of 10 like Ethernet, but double with each implementation: there's 1GFC, 2GFC, 4GFC, 8GFC and 16GFC. The 'G' stands for Gigabit, as you need high bandwidth for storage. A FC adapter is not like an Ethernet NIC: it doesn't have an IP, and it will not be treated as a NIC by the operating system, but more like a storage device (which it is).

Fibre Channel over Ethernet

When data centers started to grow, this gave some scalability options when implementing redundancy. Redundancy meant two Ethernet NICs, but also two FC adapters for storage, giving a total of four connections per server. For this reason, Fibre Channel over Ethernet (or FCoE) was developed. FCoE uses Ethernet frames (up to OSI layer 2) and sets FC on top of that (from FC2 and up). The result is a converged network that can transport both device communications and storage blocks.

For this to work you'll need a Converged Network Adapter (CNA) and switches capable of FCoE. It's theoretically possible to use a normal NIC and let software calculate the FCoE frames, but few, if any, of these implementations exists. Also, I haven't found any sources claiming a standard Ethernet switch will or will not work. Most likely they'll work, but given the unreliable nature of Ethernet, you'll run into serious problems once congestion occurs, as SCSI does not recover well from lost or out-of-order-delivered frames (most likely your operating

system will crash or get corrupted). A FCoE enabled switch, like the Cisco Nexus series for example, provides lossless Ethernet techniques to handle this, and can use FC signalling to prevent congestion.

Fibre Channel over IP

So that's FCoE, but as this doesn't use IP, it's still a flat network. For WAN links, there are other standards too, that can span multiple hops and don't have distance limitations like native FC. It's possible to run FC on top of IP, using FCIP or iFCP (Internet Fibre Channel Protocol). Both don't seem to be commonly used.

iSCSI

One of the more widely used techniques for converged storage networking is iSCSI, which is running SCSI on top of TCP (using ports 860 and 3260). This doesn't really involve any FC formatting anymore in any part of the frame, so it's less overhead than FCIP and iFCP, which also run on TCP but then still require FC headers. TCP counters the unreliability of Ethernet, allowing for reliable frame delivery and sequence numbering to prevent out-of-order-delivery. iSCSI also doesn't require specialized networking gear, allowing for normal Ethernet network equipment. You can even implement QoS and basic firewalling matching on TCP port numbers.

Storage space

SCSI uses Logical Unit Numbers (or LUNs) to differentiate between different (virtual) partitions on a storage device. This means that you can have a large SAN server with several TB of storage, divided into many different LUNs, one for each server. Servers then communicate using SCSI (over any of the above technologies) using LUNs to address their part of the storage. This way, servers do not interfere with each other's storage. Most modern operating systems have support for iSCSI. VMWare's ESXi and vSphere even implement this on the hypervisor level, making the storage disks appear completely local to the virtual machines.

IP over Fibre Channel

Internet Protocol over Fibre Channel (IPFC) exists too, but it doesn't seem to be used a lot. Good documentation and drivers are hard to find, so why go through all the trouble? Most companies already have a working Ethernet infrastructure and Ethernet is usually less expensive. This is also another reason why iSCSI is popular: some claim that buying 10 Gigabit-Ethernet switches and NICs is less expensive than buying 8GFC switches and adapters, and the increased overhead of TCP and iSCSI is less than the speed gain from 8 Gbps to 10 Gbps. This was quite a write, but I hope to have cleared out the basic differences and similarities between these two technologies. Anything to add? Let me know in the comments.

Source : <http://reggle.wordpress.com/2011/10/15/ethernet-versus-fibre-channel-an-introduction/>