



Storage Area Networking Overview



<https://t.me/learningnets>

In This Course

- ▶ Storage Area Networking Hardware Components
 - Arrays, Switches, Host Adapters
- ▶ Storage Protocols
 - File vs. Block Storage, CIFS, NFS, Fibre Channel, FCoE, iSCSI
- ▶ Fibre Channel
 - FLOGI, PLOGI, FCNS, Domain IDs, FSPF, FCID, pWWN, nWWN, Zoning, Aliases
- ▶ Quality of Service
 - Buffer to Buffer Credits, DCBX, Priority Flow Control
- ▶ Storage High Availability
 - Storage Replication, SAN A / SAN B Design, Multipath I/O
- ▶ Unified Fabric
 - FCoE, Multihop FCoE
- ▶ Unified Computing System (UCS)
 - C Servers, B Servers, Boot From Fibre Channel & iSCSI SAN

Q&A

<https://t.me/learningnets>



Storage Area Networking Hardware Components



<https://t.me/learningnets>

Storage Area Network (SAN) Components

▷ Three Main Components:

- Storage Arrays
- Storage Switches
- Host Bus Adapters (HBAs)

Storage Arrays

- ▷ Physical disk enclosures
- ▷ Offers block or file level storage access to servers
- ▷ Two main disk configurations
 - Just a Bunch of Disks (JBOD)
 - Redundant Array of Independent Disks (RAID)



JBOD vs. RAID

- ▶ JBOD presents each disk separately to end host
 - E.g. Windows Server sees 3 disks and creates a stripe set out of them
- ▶ RAID presents logical volume as a physical disk to end host
 - E.g. Windows Server sees 1 disk
 - Array hides physical details of the disks, e.g. a RAID 5 set

SAN Array vs. DAS and NAS

▷ Direct Attached Storage (DAS)

- Server's local physical disks
- Can limit portability, e.g. VMware VMotion

▷ Network Attached Storage (NAS)

- NFS & CIFS/SMB file shares
- Can limit disk Input Output Per Second (IOPS) throughput

SAN Array Physical Connectivity

- ▷ SAN Arrays typically attach to network three ways
- ▷ Native Fibre Channel (FC)
 - 1/2/4/8/16 Gbps
- ▷ Fibre Channel over Ethernet (FCoE)
 - 10/40 GigE
- ▷ IP
 - iSCSI Block Level Storage
 - CIFS/NFS/etc. File Level Storage

Storage Switches

- ▶ Connect servers to storage
 - Traditionally FC to FC switches
 - Server is called the “initiator”
 - Storage is called the “target”
- ▶ Like IP routers, can do protocol conversions
 - Fibre Channel to FCoE
 - Fibre Channel to iSCSI
 - Fibre Channel to FCIP
- ▶ Directors vs. Switches
 - Directors are just bigger switches
 - Think Catalyst 6509 vs. Catalyst 2960
 - Higher port density
 - Higher redundancy (power, supervisors, line cards)



Cisco Storage Platforms

▷ Multilayer Director Switches (MDS)


- 9700 Directors
- 9300 / 9200 / 9100 Switches
- <http://www.cisco.com/go/mds>

▷ Nexus

- 7700/7000 FCoE Switches
- 5000/6000 Fibre Channel and FCoE Switches
- <http://www.cisco.com/go/nexus>

Host Bus Adapters (HBAs)

- ▷ NIC card for SAN
 - Connects servers and storage to SAN switches
- ▷ Typically three types of HBAs
 - Fibre Channel HBAs
 - SAN traffic only
 - 1/2/4/8/16 Gbps
 - iSCSI HBAs
 - 1/10/25/40 Gbps LAN plus iSCSI hardware offload
 - FCoE Converged Network Adapter (CNA)
 - E.g. “Unified Fabric”
 - 10/25/40 Gbps LAN plus FCoE hardware offload
 - Multiple generations of cards
 - More on this later...



Course Hardware Topology

- ▷ Nexus 5672UP
 - Native Fibre Channel, FCoE, & IP Storage (i.e. iSCSI)
- ▷ Nexus 7000 w/ F3 Linecard
 - FCoE only in Storage VDC
 - IP Storage in User VDC (i.e. iSCSI)
- ▷ UCS 6200 Fabric Interconnects
 - Native Fibre Channel, FCoE, & IP Storage (i.e. iSCSI)
- ▷ UCS B Series Blades
- ▷ UCS C Series Rack Mount Servers
- ▷ Dual Attached Storage Array

Q&A

<https://t.me/learningnets>



Fibre Channel Protocol



<https://t.me/learningnets>


What is Fibre Channel?

▷ From a high level, replaces SCSI disk cable with a network

▷ From:



▷ To:



What is Fibre Channel?

- ▶ Protocol stack primarily used to send SCSI commands over the SAN
 - Technically you could run IP over FC, but main application is SAN
 - RFC 2625 - IP and ARP over Fibre Channel
- ▶ Standard “T11” per InterNational Committee for Information Technology Standards (INCITS)
 - <http://www.t11.org/>
 - E.g. not IEEE or IETF
- ▶ FCoE is T11’s FC-BB-5 standard
 - <http://www.t11.org/fcoe>

Fibre Channel Topologies

- ▷ Supports three types of topologies
- ▷ Point-to-Point (FC-P2P)
 - Initiator (server) and Target (storage) directly connected
- ▷ Arbitrated Loop (FC-AL)
 - Logical ring topology, similar to Token Ring
 - Implies contention is required on the ring
- ▷ Switched Fabric (FC-SW)
 - Logical equivalent to a Switched Ethernet LAN
 - Switches manage the fabric allowing any-to-any communication without contention
 - Similar to how CSMA-CD is removed with Ethernet Switching

Fibre Channel Port Types

- ▷ FC has different port types to define a port's function
- ▷ N_port - Node Port
 - End host (target or initiator) in P2P or Switched Fabric
- ▷ NL_port - Node Loop Port
 - End host in an Arbitrated Loop
- ▷ F_port - Fabric Port
 - Switch's port that connects to a Node Port
- ▷ FL_port - Fabric Loop Port
 - Switch's port that connects to a Node Loop Port

Fibre Channel Port Types

▷ E_port - Expansion Port

- Inter Switch Link (ISL)

▷ TE_port - Trunking Expansion Port

- Extended ISL, analogous to an 802.1Q Trunk

▷ Other port types will be addressed as we progress

- E.g. NP (Proxy Node Port), VE (Virtual Expansion Port), ENode (Ethernet Node Port), etc.

Fibre Channel Addressing

- ▷ FC addressing is analogous to IP over Ethernet
 - IP addresses are logical and manually assigned
 - Ethernet MAC addresses are physical and burned in
- ▷ FC World Wide Names (WWNs)
 - 8 byte address burned in by manufacturer
- ▷ FC Identifier (FCID)
 - 3 byte logical address assigned by fabric

FC World Wide Names

- ▷ WWN is FC's 8 byte physical address
- ▷ WWN is subdivided further into...
- ▷ World Wide Node Name (WWNN or nWWN)
 - Switch, server, or disk's physical address
- ▷ World Wide Port Name (WWPN or pWWN)
 - Switch, server, or disk's port's physical address
 - E.g. a switches, HBAs, and arrays have more than one port
- ▷ WWN is not used for data plane switching

FC Identifiers

- ▷ FCID is FC's 3 byte logical address
- ▷ FCID is subdivided into three fields
 - Domain ID
 - Each switch gets a Domain ID
 - Area ID
 - Group of ports on a switch have an Area ID
 - Port ID
 - End station connected to switch gets a Port ID
- ▷ FCID is used for the actual data plane traffic switching

Fibre Channel Domain IDs

- ▷ Domain IDs identify a switch in the fabric
- ▷ Can be manually assigned, otherwise will be automatically assigned by the Principal Switch
- ▷ Principal Switch is analogous to the STP Root Bridge, and is chosen based on an election
- ▷ No configuration needed for PS election

Fibre Channel Routing

- ▷ FC does not use flooding to build topology like Ethernet does
- ▷ Fabric Shortest Path First (FSPF) is used to route traffic between switches
 - Same Dijkstra SPF as OSPF & IS-IS
 - Node ID in the SPT is the FCID's Domain ID
 - Traffic is routed via lowest cost path between Domain IDs
 - ECMP is supported for equal SPT branches
 - Unequal cost load distribution is not supported
- ▷ FSPF runs automatically as a Fabric Service
 - No configuration required
 - Can customize knobs, but typically not required

Fibre Channel Logins

- ▷ Ethernet networks are connectionless
 - Traffic in the data plane results in topology learning in the control plane
- ▷ Fibre Channel networks are connection oriented
 - All end stations must first register with the control plane of the fabric before sending any traffic
- ▷ Fabric registration has three parts
 - Fabric Login (FLOGI)
 - Port Login (PLOGI)
 - Process Login (PLRI)

FLOGI, PLOGI, & PLRI

▷ Fabric Login (FLOGI)

- Node Port (N_Port) tells switch's Fabric Port (F_Port) it wants to register
- Switch learns the WWNN and WWPN of Node
- Switch assigns FCID to Node

▷ Port Login (PLOGI)

- End-to-End login between Node Ports
- Initiator tells Target it wants to talk
- Used for applications such as end-to-end flow control

▷ Process Login (PLRI)

- Upper layer protocol login negotiation between Node Ports

Fibre Channel Name Server

- ▶ Fibre Channel Name Server (FCNS) is analogous to ARP Cache
- ▶ Used to resolve the WWN (physical address) to the FCID (logical address)
- ▶ Like Principal Switch and FSPF, FCNS is a distributed fabric service that requires no configuration

Zoning

- ▶ By default all initiators learn about all targets during the login process
 - FCNS maintains the mappings of everyone's WWPN to FCID
- ▶ Servers mounting the wrong volumes can corrupt data
 - E.g. Windows NTFS & MBR not compatible with Linux GPT
- ▶ Zoning prevents this by limiting which resources an initiator can use
 - Zoning is analogous to ACLs in Ethernet & IP world
 - Associates WWNs, FCIDs, aliases, etc. to control who can talk to who
- ▶ Like FCNS, Zoning is a distributed fabric service

Virtual SANs (VSANs)

- ▶ Traditionally multiple SANs were designed as physical separate networks
 - i.e. “SAN Islands”
 - Physical separation is costly in terms of equipment, power, space, cooling, management, etc.
- ▶ VSANs solve the isolation problem similar to how VLANs segment broadcast domains
 - Isolates the management and failure domain of the network
 - Separates FLOGI, FCNS, Zoning, Aliases, etc. per VSAN
- ▶ With VSANs, E Ports now become TE Ports
 - Similar to 802.1Q trunks in Ethernet

SAN Port Channeling

- ▶ Like Ethernet Port-Channeling, SAN PCs can be used to aggregate the bandwidth of physical links
- ▶ Supports Port Channeling Protocol (PCP) for negotiation of links
 - Similar to 802.3ad LACP in Ethernet

Q&A

<https://t.me/learningnets>



UCS C Series and B Series SAN Connectivity



<https://t.me/learningnets>

Q&A

<https://t.me/learningnets>



Unified Ports on Nexus 5000 and UCS B Series Fabric Interconnects



<https://t.me/learningnets>

Q&A

<https://t.me/learningnets>



FCoE over Enhanced vPC



<https://t.me/learningnets>

Enhanced vPC Overview

▷ Enhanced vPC is combination of...

- Fabric vPC from Parent Switches to dual-homed FEXes
- Host vPC from dual-homed FEXes to dual-homed host

▷ Enhanced vPC and FCoE

- SAN A / SAN B separation must be maintained through FEXes
- FEX 101 pins FCoE traffic to Parent A
- FEX 102 pins FCoE traffic to Parent B
- Result is that LAN traffic follows vPC but not FCoE traffic

FCoE over EvPC Configuration

- ▷ Final configuration can be built from two documentation examples:
 - [Nexus 5600 Series NX-OS Interfaces Configuration Guide - Configuring Enhanced Virtual Port Channels](#)
 - [Nexus 5600 Series NX-OS Fibre Channel over Ethernet Configuration Guide - Configuring FCoE VLANs and Virtual Interfaces](#)
- ▷ Configuration Caveats
 - Parent A must enable “fcoe” only under FEX 101
 - Parent B must enable “fcoe” only under FEX 102
 - VFC interfaces are bound to **physical** interface to host
 - vPC Consistency must match except for FCoE config

Q&A

<https://t.me/learningnets>



SAN Port-Channels



<https://t.me/learningnets>

Q&A

<https://t.me/learningnets>



FCoE on UCS B Series



<https://t.me/learningnets>

UCS B Series FCoE Connectivity Options

▷ UCS B Series FCoE Connectivity Options

- FCoE Uplinks
- FCoE Port-Channels
- Unified Uplinks
- Unified Port-Channels

▷ FCoE Uplinks vs. Unified Uplinks

- FCoE Uplinks only trunk VLANs mapped to VSANs (FCoE VLANs)
- Unified Uplinks trunk all VLANs by default
 - Can be modified with VLAN groups for disjoint layer 2

Unified Port-Channels

▷ Unified Port-Channels Order of Operations

- Create SAN FCoE Port-Channel and modify to Unified Port-Channel
- Create LAN Port-Channel and modify to Unified Port-Channel

▷ First order:

- Create FCoE Uplink
- Create FCoE Port-Channel number **X** and add member ports
- Create LAN Port-Channel number **X** and *add no member ports*
- Result is Unified Port-Channel

▷ Second order:

- Create LAN Uplink
- Create LAN Port-Channel number **X** and add member ports
- Create FCoE Port-Channel number **X** and *add no member ports*
- Result is Unified Port-Channel

Q&A

<https://t.me/learningnets>



Multihop FCoE



<https://t.me/learningnets>

Q&A

<https://t.me/learningnets>



Zoning & Device Aliases



<https://t.me/learningnets>

Q&A

<https://t.me/learningnets>



UCS SAN Pin Groups



<https://t.me/learningnets>

Q&A

<https://t.me/learningnets>