

Terminology

Direct Attached Storage (DAS) Storage that accessed directly from Host as block-level device with limited mobility and scalability

Network Attached Storage (NAS) file-level storage connected to network, serves files to other devices on the network using TCP & UDP i.e. (SMB & NFS)

Storage Area Network (SAN) high-speed network that provides block-level network access to storage. composed of hosts, switches, storage where fibre channel used for carrying SCSI or FICON traffic

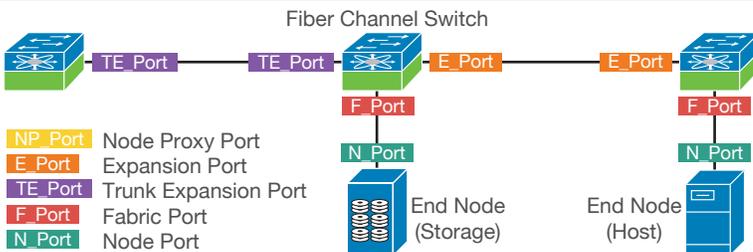
Fibre Channel (FC) high-speed network technology used to for storage area networking (SAN), running at 1, 2, 4, 8, 16, 32, & 128 Gbps

SCSI Standard that defines an interface between an initiator(host) and a target(storage) which used to carry block based storage commands

iSCSI carrying SCSI commands over IP, used to facilitate data over ethernet & managing connection between IP based storage and hosts

Fibre Channel over IP (FCIP) Protocol used to tunnel fibre channel packets over an IP infrastructure. Extend FC over long distance

Fibre Channel Port Types



N-Port Node port used to connect devices to Fibre Channel switch

F-Port Switch Ports used to connect to N-Port devices (Switch Side)

L-Port Loop port are used in arbitrated loop topology to build network without FC switches

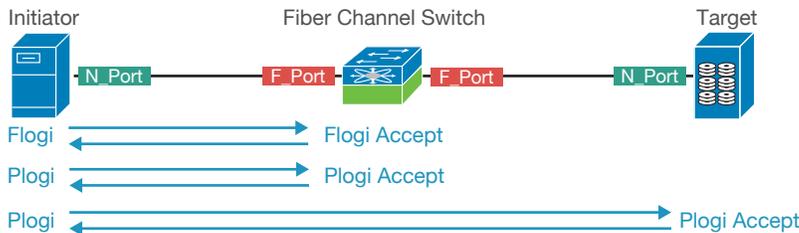
NL-Port Node port that connects to both loops and switches

FL-Port Switch port that connects to both loops and switches

E-Port Expansion port used to cascade and connect FC switches

TE-Port expand E ports functionality to support VSAN trunking, QOS, fctrace feature and create an Extended ISL (EISL) between switches

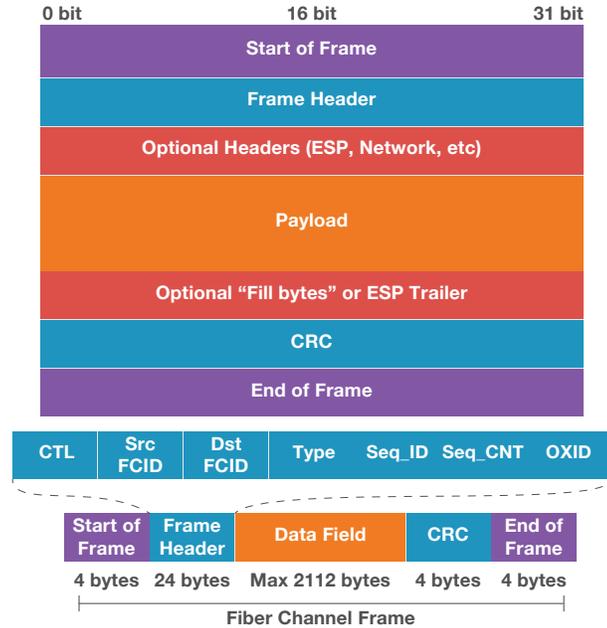
FLOGIs & PLOGIs



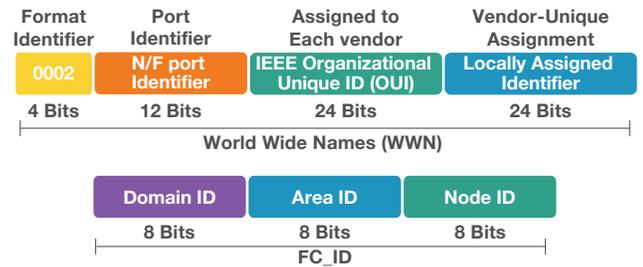
Fabric Login (FLOGI) N_Port determine presence of fabric and requests a unique 24-bit address (FCID/PID), class of service and Buffer to Buffer Credit from the Fabric Login Server

Port Login (PLOGI) End-to-End logging between node port, which require to communicate with each other

Fibre Channel Frame Format



Fibre Channel Addressing



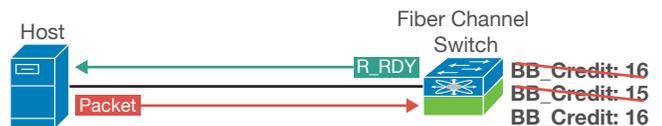
World Wide Names (WWN) 8 bytes burnt-in address assigned to switches, ports and node by manufacturer. These addresses are registered in fabric and mapped to an FC_ID. Unlike MAC address, the WWN is not used as the source or destination of a frame.

World Wide Node Names (WWNN) Every Switch and HBA card is assigned a WWNN at time of manufacture

World Wide Port Names (WWPN) Every Switch and HBA port will have different address. E.g. 2 port HBA card, include 1 WWNN and 2 WWPN at same time

FC_ID 3 bytes address used for FC data plane traffic switching & routing similar to IP address. The FCID is comprised of the switch's Domain ID, Area ID and Port ID

Buffer to Buffer Credits



B2B credits used to ensure FC transport is lossless, when the link is brought up number of credits negotiated between ports, this credits **decremented** with each packet placed on the wire (independent of packet size) and if number of credit = 0, transmission will stop

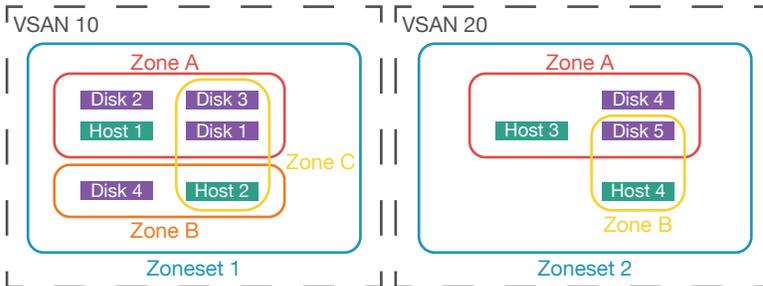
FC Fabric Services

VSAN Analogous to VLAN or VRF in Ethernet, provide cost-effective virtual fabric. Fabric events are isolated per VSAN which provide isolation for High Availability. Each VSAN provides separate Fabric Services (FSPF, Zones/Zoneset, DNS, RSCN)

Fabric Shortest Path First (FSPF) Analogous to OSPF in Ethernet. FSPF routes traffic based on destination domain ID found in destination FCID, each domain ID identifies a single switch for FSPF (domain ID is limited to 239), FSPF performs hop-by-hop routing use total cost as the metric to determine most efficient path and support equal cost load balancing across links

Fibre Channel Name Server (FCNS) Analogous to ARP cache in Ethernet. Used to resolve the WWN (physical address) to the FCID (logical address), like FSPF, FCNS is a distributed fabric service that requires no configuration and keep a mapping of FCID to WWPN inside its database. FCNS database will be populated by FLOGI database, and it will be shared across all the switches in fabric

Zoning



Zone are the basic form of data path security, Zone members can only talk to other members of the same zone. Devices can be members of more than one zone. Default zoning is “deny” similar to Access-List, zones belong to a zoneset and it should be “active” to enforce zoning. Only one active zoneset per VSAN is allowed

Basic Zoning Multiple administrators can make changes at the same time leading to potential conflicts and config over ride

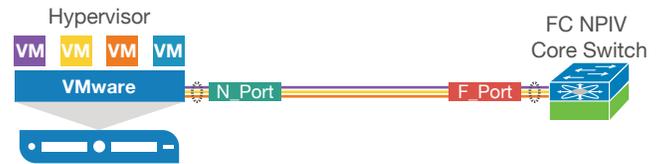
Enhanced Zoning Entering the zoning configuration mode locks the zone database across the fabric. Can be enabled per VSAN

Fibre Channel Services Configuration

```
# Interface and VSAN configuration
interface fc1/1
  switchport mode [E | F | NP | TE | TF | TNP | auto]
vsan database
  vsan 10 name "FABRIC_A"
  vsan 10 interface fc1/1
show vsan [usage | x | x membership | membership]

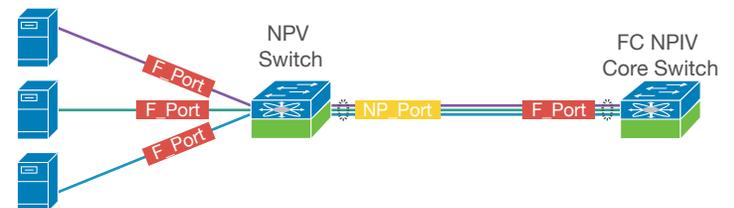
# Zoning
zone name ZONE-A vsan 10
  member pwwn 50:06:01:60:XX:XX:XX:XX
  member fcid 0xce00XX
zoneset name ZONESET1 vsan 10
  member ZONE-A
zoneset activate name ZONESET1 vsan 10
show [flogi database | zone | zoneset | zoneset active]
debug zone change [errors | events | packets] 10
```

NPIV



N-Port ID Virtualization (NPIV) NPIV switch allows multiple FLOGI and FCID assignment on its F_Port facing downstream or NPIV Virtualization provides a means to assign multiple FCIDs to a single N_Ports. Allows multiple application to share the same Fiber Channel adapter port. Usage applies to platform such as VMWare, MS HyperV, Citrix

NPV



N-Port Virtualization (NPV) NPV fixes the FCID (Domain ID) problem by removing the need for a switch to participate in Fabric Services (without NPV, storage fabric is limited to 239 domain ID/switch). Switches running NPV appear to the rest of the fabric as an end host (E.g. N_Port). Upstream facing link on the NPV switch is called the NP_Port. No local switching, routing and zoning on FC switch with NPV mode.

Port Trunking and Channeling

Port Channels is a logical bundling of identical links, to form port channel; link speed, port mode & state (auto, E, etc.) and VSAN membership must match. can bundle up to 16 links and treated as one logical ISL by upper layer protocols (FSPF). Provide load balancing base on src/dst FC_ID and OX_ID

Trunking VSAN trunking enables a link transmitting frames in the EISL format to carry traffic for multiple VSAN. When trunking is operational on an E port, that E port becomes a TE port and able to carry tagged from from multiple VSANs

F-Port Port Channel and F-Port Trunking

F-Port Port Channel with NPV Bundle multiple F-Port on enabled NPV switch into 1 logical link, Similar to ISL portchannels in FC and Etherchannels in Ethernet. Provide High-Availability, no disruption if cable port or line cards fail. Optimal bandwidth utilization & higher aggregate bandwidth with load balancing

F-Port Trunking with NPV Partition F-Port to carry multiple VSANs, extends VSAN benefits to blade server and separate management and isolation domain, it also provide QOS and Security features