Data Center Scalability Made Easy with Fibre Channel Services

Live Webcast
August 26, 2020
10:00 am PT

Today's Presenters

Kiran Ranabhor
Cisco Systems
Moderator

David Peterson
Broadcom Inc.
Presenter

Barry Maskas
HPE
Presenter

About the Fibre Channel Industry Association (FCIA)

25+ Years
Promoting Fibre Channel Technology

Industry Leading
Member Companies

142M+ FC Ports
Shipped Since 2001

Key Tenants of Fibre Channel

- Purpose-built as network fabric for storage and standardized in 1994, Fibre Channel (FC) is a complete networking solution, defining both the physical network infrastructure and the data transport protocols. Features include:
  - **Lossless, congestion free systems**—A credit-based flow control system ensures delivery of data as fast as the destination buffer can receive, without dropping frames or losing data.
  - **Multiple upper-layer protocols**—Fibre Channel is transparent and autonomous to the protocol mapped over it, including SCSI, TCP/IP, ESCON, and NVMe.
  - **Multiple topologies**—Fibre Channel supports point-to-point (2 ports) and switched fabric (224 ports) topologies.
  - **Multiple speeds**—Products are available supporting 8GFC, 16GFC, and 32GFC today.
  - **Security**—Communication can be protected with access controls (port binding, zoning, and LUN masking), authentication, and encryption.
  - **Resiliency**—Fibre Channel supports end-to-end and device-to-device flow control, multi-pathing, routing, and other features that provide load balancing, the ability to scale, self-healing, and rolling upgrades.

Agenda

- What are Fibre Channel Fabric Services?
- Overview of long-standing Fabric Services and what the newer Fabric Services provide
- What is FC-CT? And how does it relate to Fibre Channel Fabric Services?
- Fibre Channel Generic Services and Switch Fabric functionality
- Potential Fabric Services
- Q & A

But Before We Start, the Acronyms

- FC-GS-7 – Fibre Channel – Generic Services – 7
  - FC-CT - Common Transport
  - Directory Service, Management Service
- FC-SW-7 – Fibre Channel – Switch Fabric – 7
  - Distributed Services
  - Fabric Configuration – Principal Switch selection
  - Path Selection – FSPF
- FC-LS-4 – Fibre Channel – Link Services – 4
  - Fabric Login
  - State Change Notification
- FC-FS-5 – Fibre Channel – Framing and Signaling – 5
  - Clock Synchronization Server

What are Fibre Channel Fabric Services?

- Fabric Services are a set of functional entities that (as a whole) provide for the management and scalability of Fibre Channel Fabrics
  - Fabric Login Service (see FC-LS-4)
    - FLOGI ELS
  - State Change Service (see FC-LS-4 and FC-SW-7)
    - SCR and RSCN ELSs
  - Generic Services (see FC-GS-8)
    - Directory Service, Management Service, …
- Fabric Services are not base fabric operation functional entities such as…
  - Fabric Configuration i.e., Principal Switch selection and Domain_ID assignment
  - Path Selection i.e., the FSPF routing protocol
- Fabric Services are addressed via Well-Known Addresses (WKAs)
  - Similar concept as Telco model i.e., here is my number if you want to talk to me

Long-standing Fabric Services

- Fabric Login Service @ F_Port Controller WKA FFFFFEh
  - Fabric Login (FLOGI) ELS is the method by which an N_Port establishes its operating environment with a Fabric
    - N_Port_ID assignment
    - World Wide Names associated with the ports
    - Exchange Service Parameters
      - Buffer-to-buffer Credit, etc…

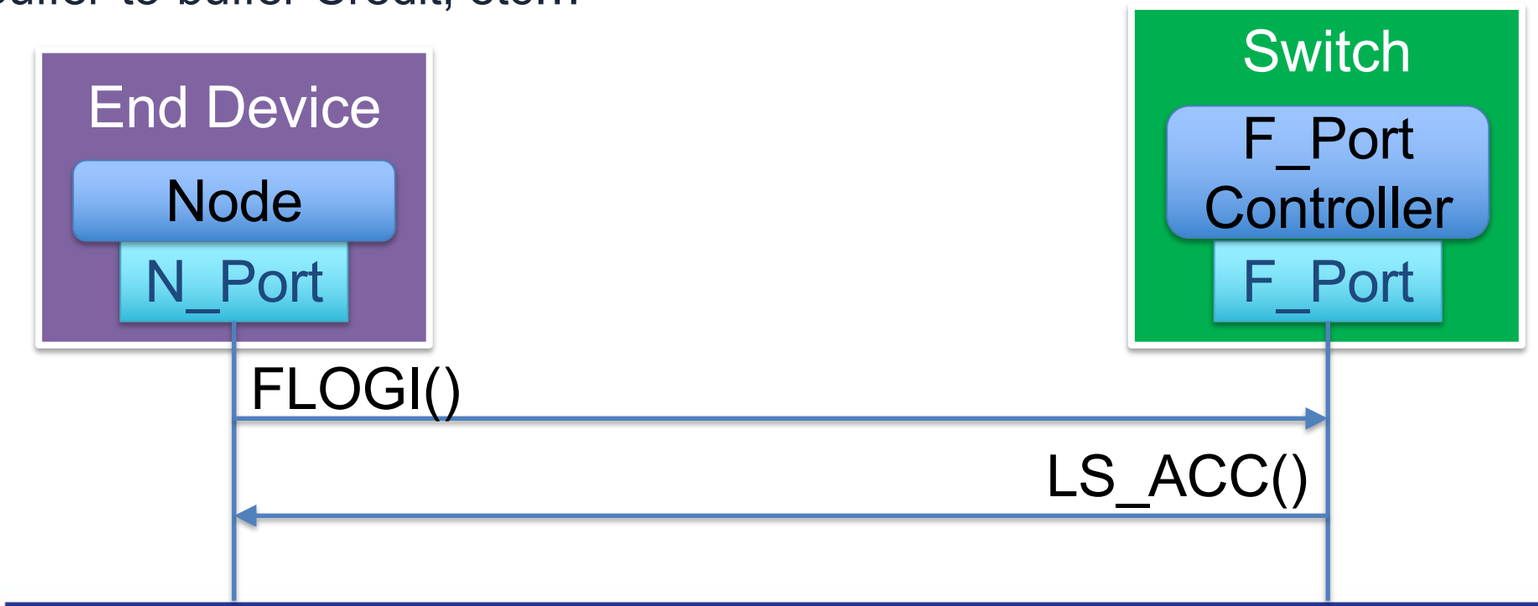

Long-standing Fabric Services

- Directory Service (Name Server subtype) @ WKA FFFFFCh
  - After Fabric Login (FLOGI) an end device sends Port Login (PLOGI) to the Name Server and registers FC-4 information such as Type and Features

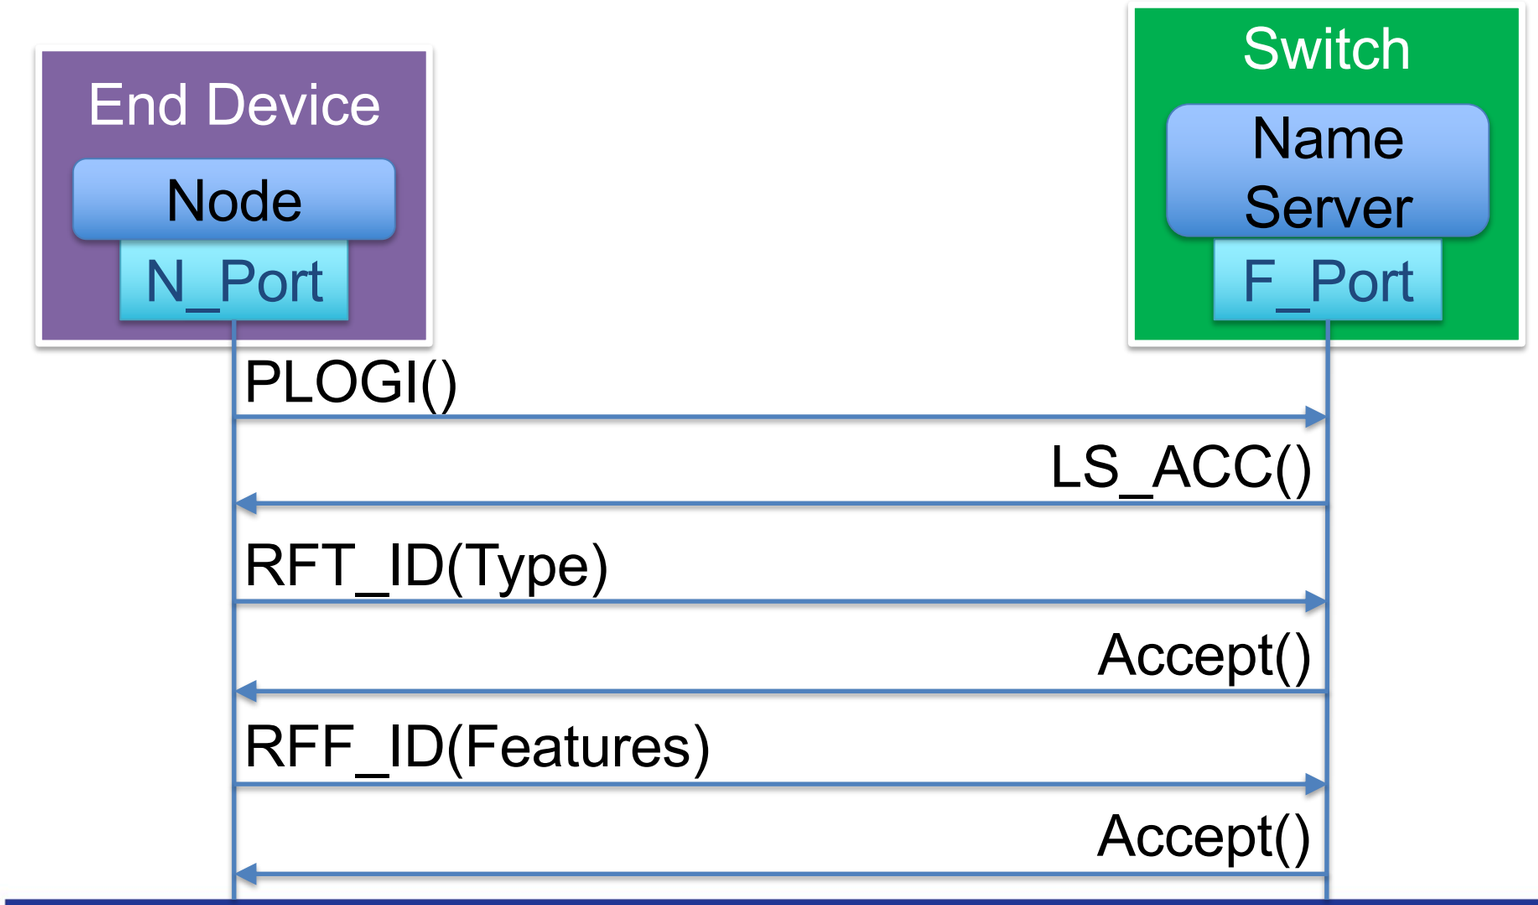

Long-standing Fabric Services

- State Change Service @ Fabric Controller WKA FFFFFDh
  - After registering with the Name Server, State Change Registration (SCR) ELS is sent to register to receive Registered State Change Notification (RSCN) ELSs from the Fabric Controller

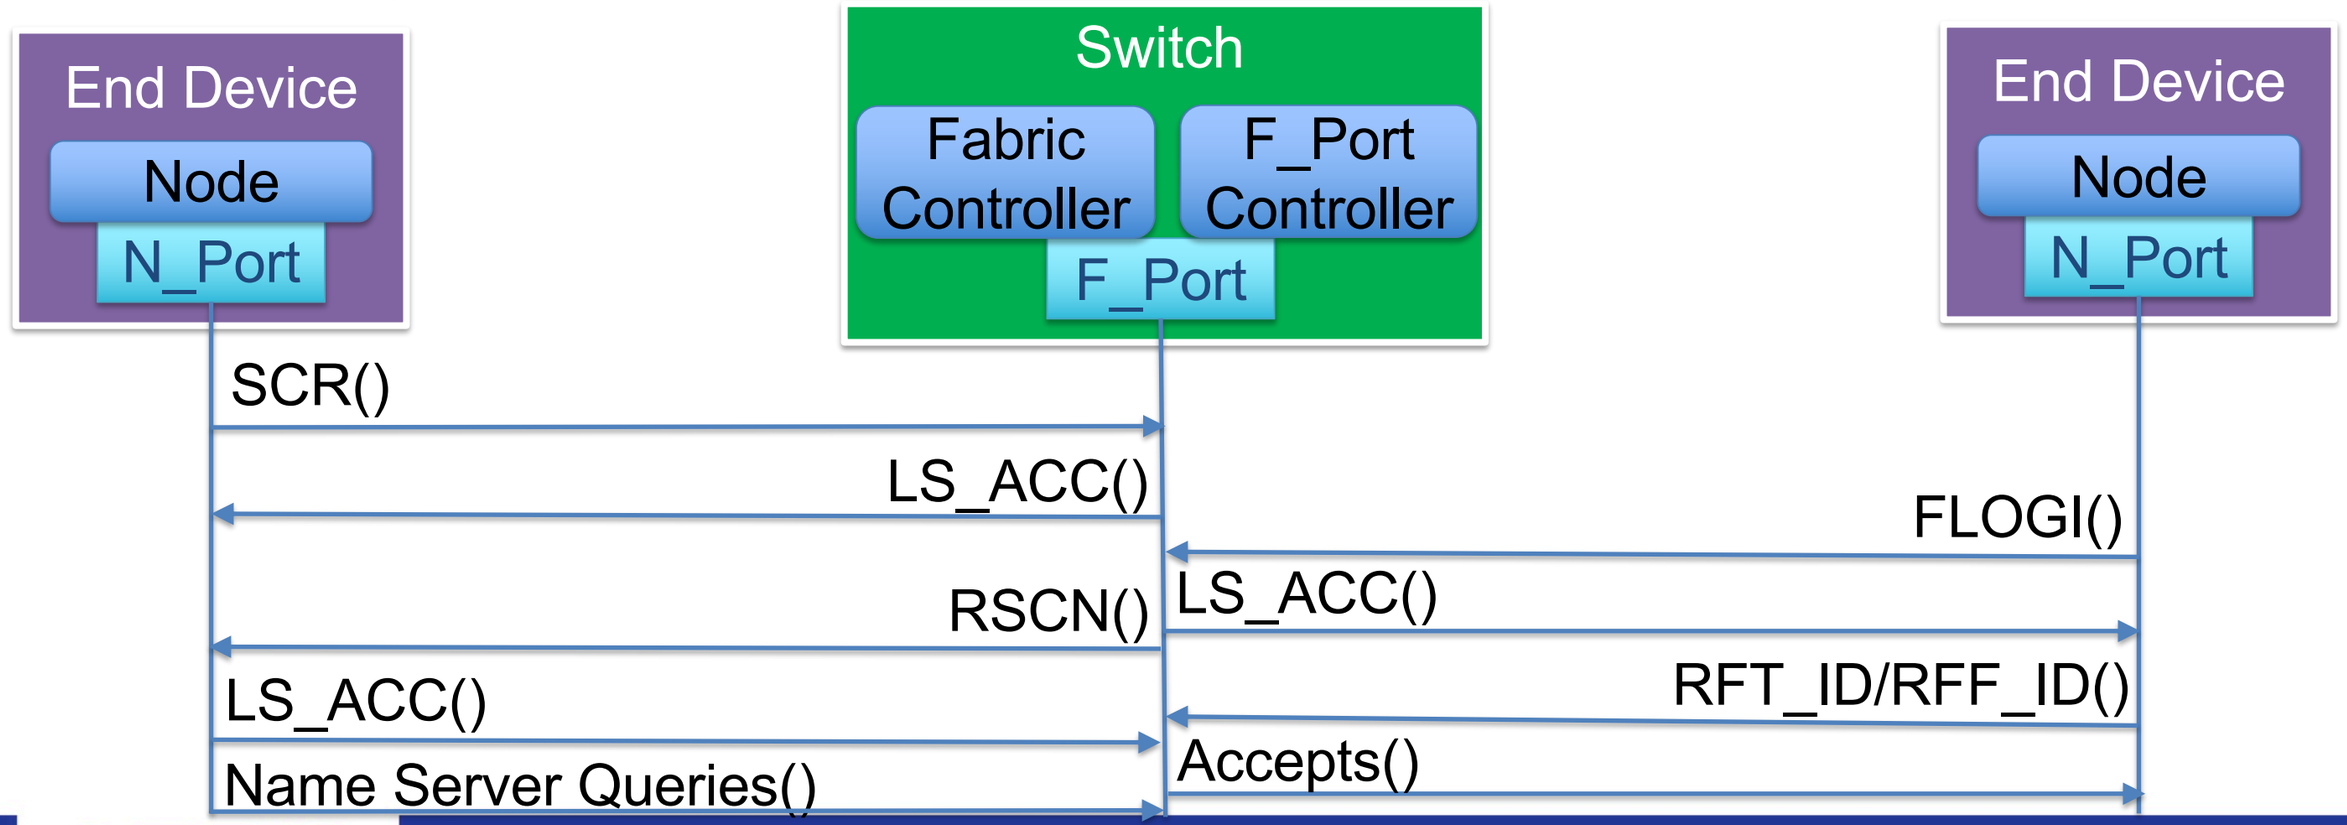

Long-standing Generic Fabric Services

- Directory Service FFFFFCh
  - Name Server
    - Provides a way for Nx_Ports to register and discover Fibre Channel attributes
  - VE Identification Server
    - Maintains the mappings between Global VE IDs and Fabric VE IDs used by the Priority Tagging mechanism (see FC-LS-4)
- Management Service FFFFFAh
  - Contains multiple management service subtypes
- Clock Synchronization Service FFFFF6h
  - Allow clocks located within nodes to be readily synchronized to microsecond accuracies
  - Used in Avionics environments e.g., FC-AE-1553 Technical Report
- Event Service FFFFF4h
  - This Service was developed with a goal of replacing RSCN functionality
  - I am not aware of any implementations

Management Service Subtypes

- Fabric Configuration Server
  - Provides a way for a management application to discover Fibre Channel Fabric topology and attributes
- Unzoned Name Server
  - Provides a management application access to the Name Server database without Zone constraints
- Fabric Zone Server
  - Provides a management application access to, and control of Fabric Zones

Management Service Subtypes

- Security Policy Server
  - Provides an in-band protocol for controlling and extracting security policy information
- Security Information Server
  - Provides an in-band protocol for observing and controlling operational security information
- Fabric Device Management Interface (FDMI)
  - Provides for the management of devices such as HBAs through the Fabric

(Newer) Management Service Subtypes

- Enhanced Fabric Configuration Server
  - Provides a way for a management application to discover Fibre Channel Fabric topology and attributes
  - Has additional queries along with a Fabric object and Transport Infrastructure object

- Application Server
  - Provides a way to manage application specific services
    - Virtual Machine (VM) identifiers
    - (ASCII) Attributes
      - Entity Name
      - Host Identifier
      - Symbolic Data

Newer Fabric Services

- Application Server – Virtual Machine Identifier flow example
  - Normal Fabric Login
  - N_Port Login
    - Enhanced for quick detection of supporting devices
  - FC-4 Type Registration
    - Unique type for Application Services
  - VM Registration
    - Allocation of tags for each VM
  - FC-4 Type Query
    - Identify VM peers
  - VM Tagged Flows
    - Fabric unique identifier

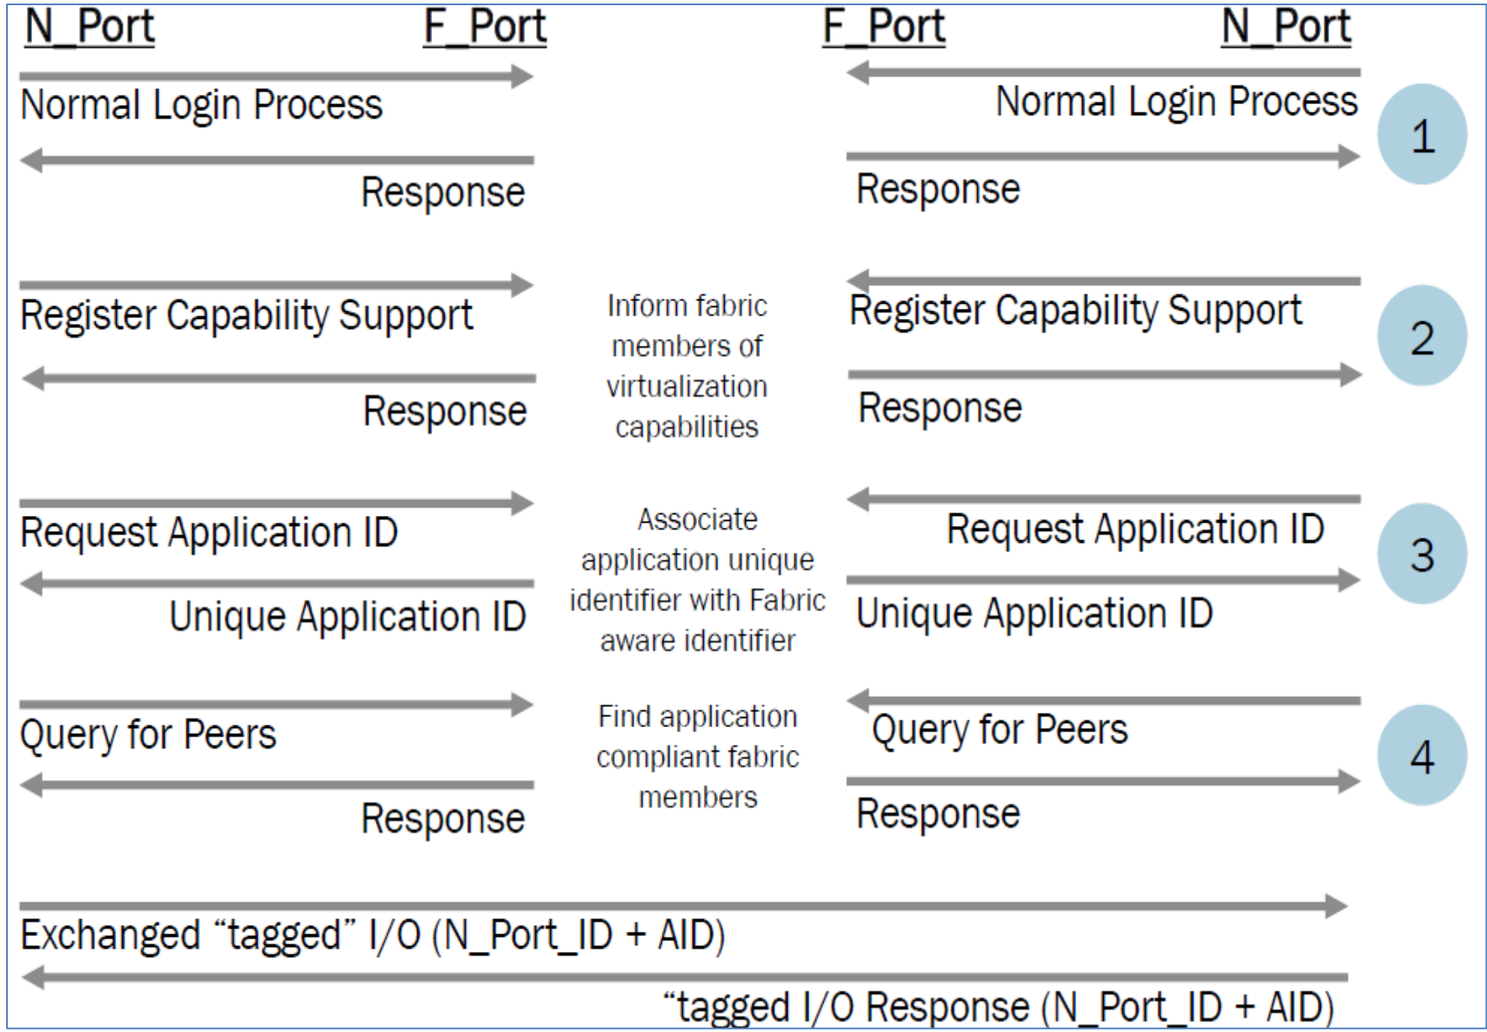

Obsolete Fabric Services

- Directory Service
  - X.500 Server
    - Obsoleted in FC-GS-2
  - IP Address Server
    - Obsoleted in FC-GS-4
- Management Service
  - Performance Server
    - Added to FC-GS-4 as an informative annex
    - Obsoleted in FC-GS-7

Obsolete Fabric Services

- Multicast Server FFFFF5h
  - Obsoleted in FC-FS-3
- (Security) Key Distribution Service FFFFF7h
  - Obsoleted in FC-GS-4
- Alias Service FFFFF8h
  - Obsoleted in FC-GS-7
- Time Service FFFFFBh
  - Obsoleted in FC-GS-7

FC-CT

Ensuring standardization and ease of deployment and scalability

FC services run on each switch, unlike Ethernet switches where services like DHCP and DNS normally run on a router, switch, or host

Building and Scaling up a FC SAN is low risk - you just plug devices in

FC Fabric services make building and scaling an FC SAN low risk and plug-and-play

- ✓ Plug-and-play in an FC SAN is a key attribute of Fabric Services

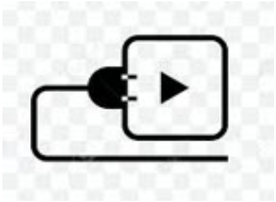

Domain_ID is dynamically assigned to a FC switch when it comes online

- Static IDs are also supported
- Before the switch can communicate with other switches, it will first configure itself to know what's attached
- The protocol determines what port mode is present, and determines the addresses of attached N_Ports and E_ports

Principal Switch (PS) selection begins followed by the Domain_ID Distribution process

- After completion of the PS election, a switch must begin the Domain_ID Distribution process.
- Even if the Domain_ID is manually configured, the distribution process occurs, because the PS needs to compile a list of Domain_IDs
- A switch assigns the FCID to each attached node
  - FCID is derived from the Domain_ID, Area_ID and WWN of the attached node

Configuring the Domain_ID is critical

- merging fabrics can be disruptive if conflicting Domain_IDs are present

N-ports utilize Fabric Services to communicate with other N-ports using the Directory Server Name Server services

Services share a Common Transport (CT) at the FC-4 level

Fibre Channel Model	Fibre Channel Implementation Standard
FC-4 (T11-3, T10, etc.) (protocol mappings)	FCP-5, FC-SB-7, FC-GS-8, FC-NVMe-2, FC-BB-6, FC-SP-2
FC-3, FC-2 (T11-3) (common services)	FC-LS-4, FC-FS-5, FC-SP-2
FC-2, FC-3 (T11-3) (network and framing)	FC-FS-5, FC-SW-7, FC-MI-2, FC-EE, FC-SP-2
FC-1 (T11-2, T11-3) (data link encoding/decoding)	FC-FS-5, FCoE mapping
FC-0 (T11-2) (physical layer)	FC-PI-7, FC-PI-7P

FC-DA-2 - Device Attach Technical Report

FC Generic Services

Services provided by FC-CT to WKA are used to support management and operation of a FC Fabric:

- device and topology discovery
- fabric zoning
- security
- address alias groups
- event reporting
- performance monitoring

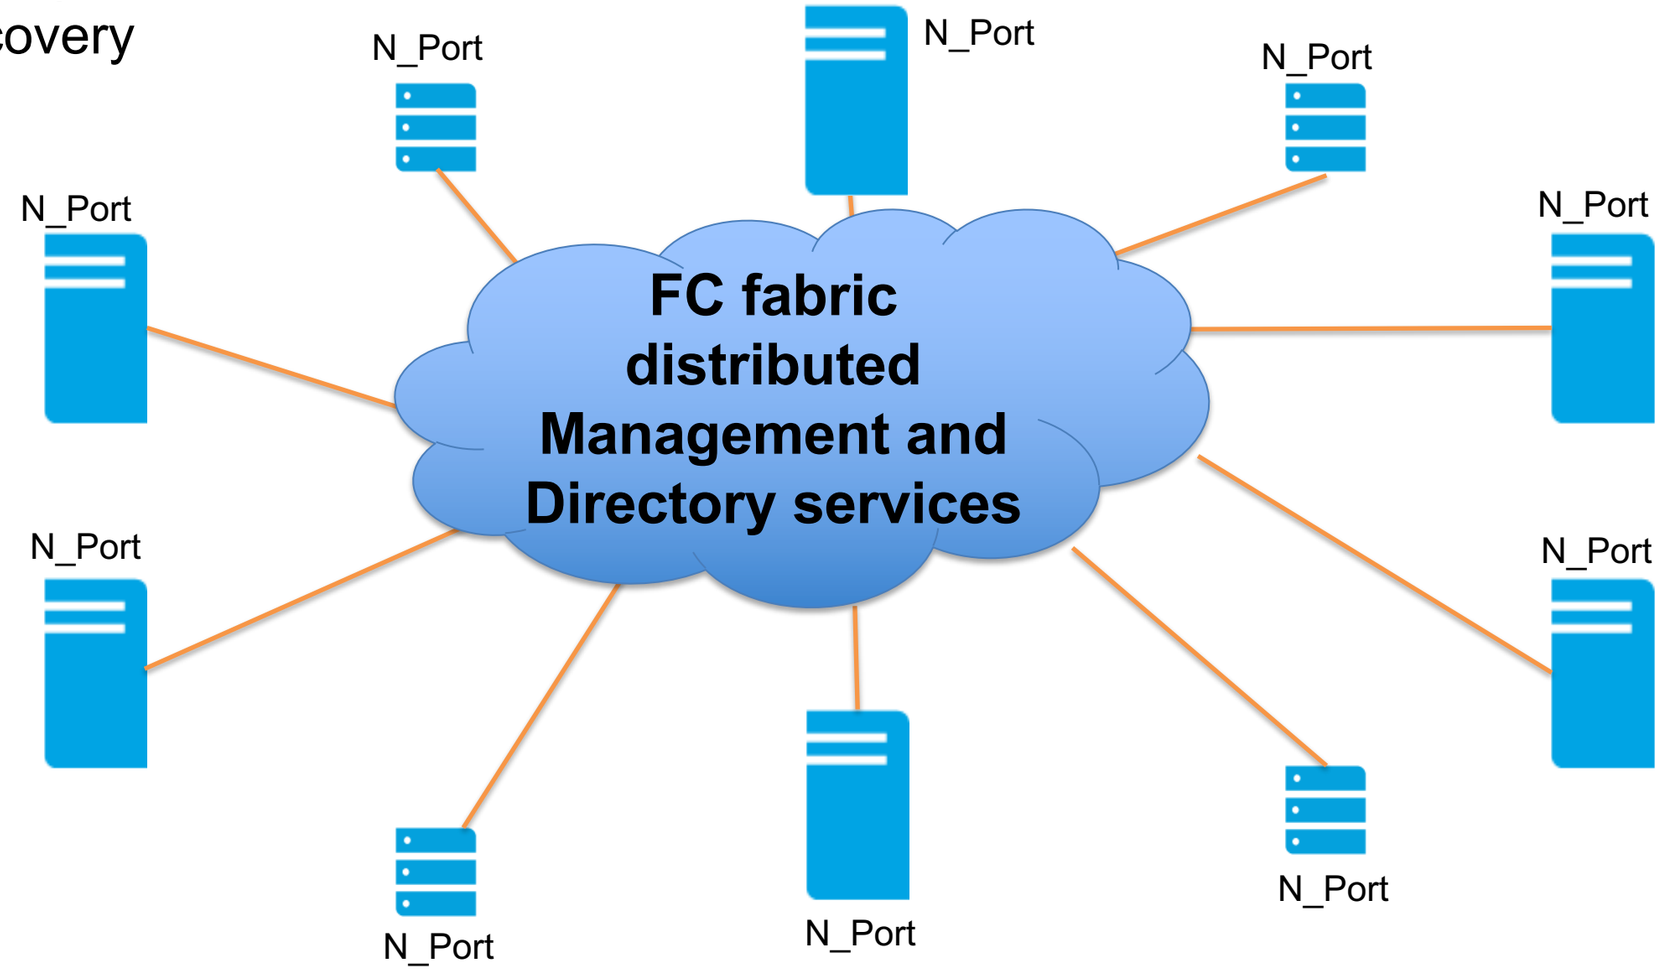

FC Generic services protocol

Fabric Services are embedded in each FC switch

- Fabric services provides the control plane for information transfer
- Technical name for FC Common Transport is FC-CT protocol
  - FC-CT protocol is a standard protocol for request and responses using well known addresses (WKA) of Distributed services
  - FC-CT provides access to a Service with a set of Service parameters
    - Is a level of multiplexing that simplifies the Server-to-Server communication for a distributed Service
  - There are Services and Operational Services – FC-CT is a service to obtain Services
    - FC-CT provides a common access method to services but does not operationally perform requested services
  - FC-CT does not participate in Principal switch election and Domain ID assignment
      - FSPF has its own protocol header and does not use the fabric configuration server
      - FSPF routing protocol is defined separately in the standard

FC Common Transport (FC-CT)

FC-CT protocol is specified by FC-GS-8

- Each service is a node port and clients must perform N-port Login (PLOGI) using the WKA to use the service
  - **Event Server** **WKA FF FF F4h (standard for enhanced rscn)**
  - **Management Server** **WKA FF FF FAh**
  - **Directory/Name Server** **WKA FF FF FCh**

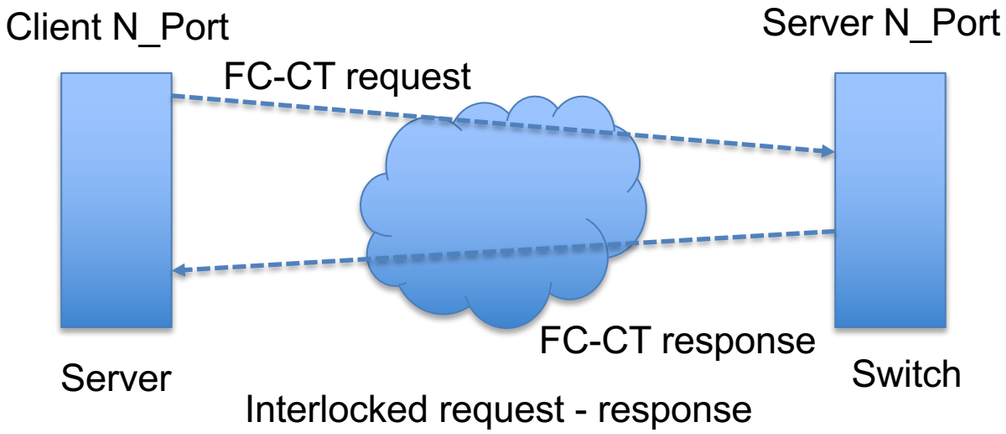

Available using Extended Link Services

- **Fabric Controller** **WKA FF FF FDh**
- **Fabric Port Login Server** **WKA FF FF FEh**
- **N_Port Controller** **WKA FF FF F0h - utilized for virtual fabrics**
- Clock Sync server WKA FF FF F6h

NOTE: WKA is not subject to Zone and Zone Set restrictions i.e., WKA is always accessible

FC Common Transport (FC-CT)

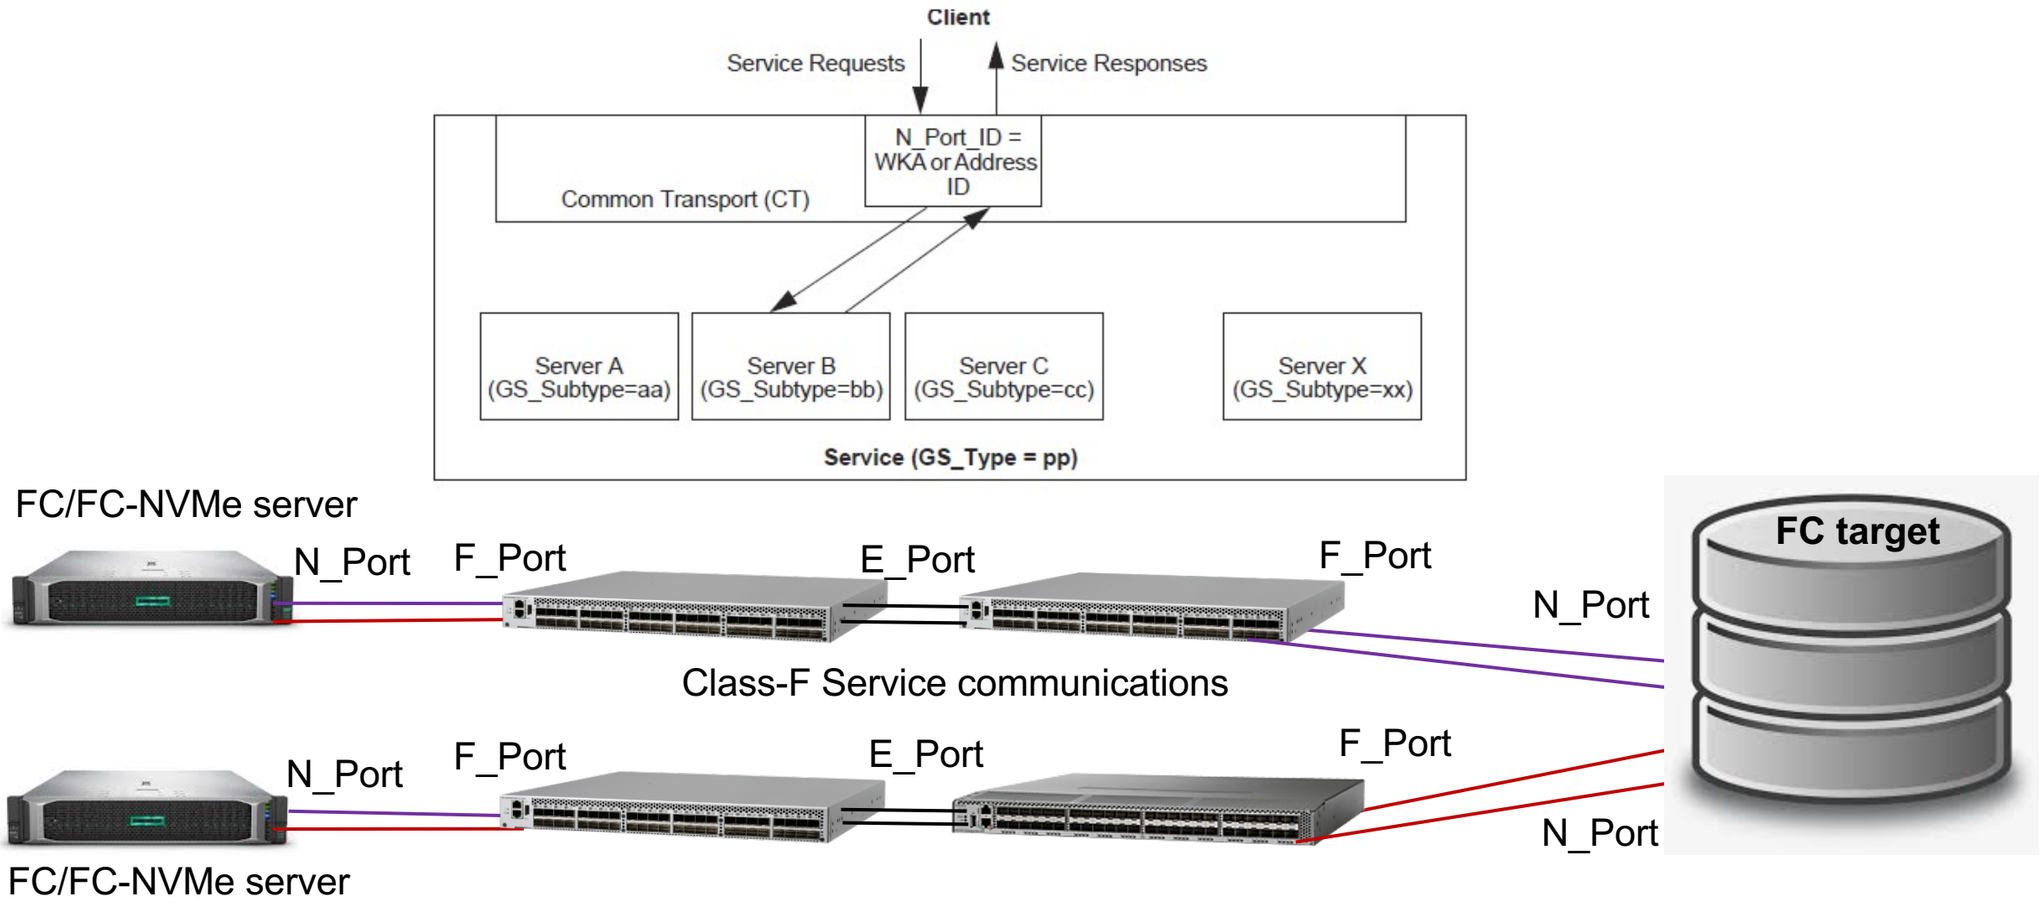

- Each switch provides a server service instance
- Collectively the switches provide services to N_Ports subscribed using WKA
- Server instances utilize Domain Controller addresses xFF FC xx for inter server service communication

FC Fabric Services

Management Server	Directory / Name Server	Fabric Controller	Fabric Login Server
Enables FC SAN management using fabric management software Located at WKA x**FFFFFAh**	Responsible for name registration and management of node ports. Located at WKA **xFFFFFCh**	Responsible for managing and distributing registered state change notifications (RSCNs) to attached node ports. Responsible for distributing SW-RSCNs to every other switch. SW-RSCNs keep the name server up-to-date on all switches. Located at WKA **xFFFFFDh**	Used during a node's fabric login process. FC address – domain ID, area ID and port ID assigned during login. Located at WKA **xFFFFFEh** Scale: 239 domain ID addresses 239 x 256 areas x 256 ports = 15663104 possible node ports

Not available for Ethernet

PLOGI Registration required by FC-CT service

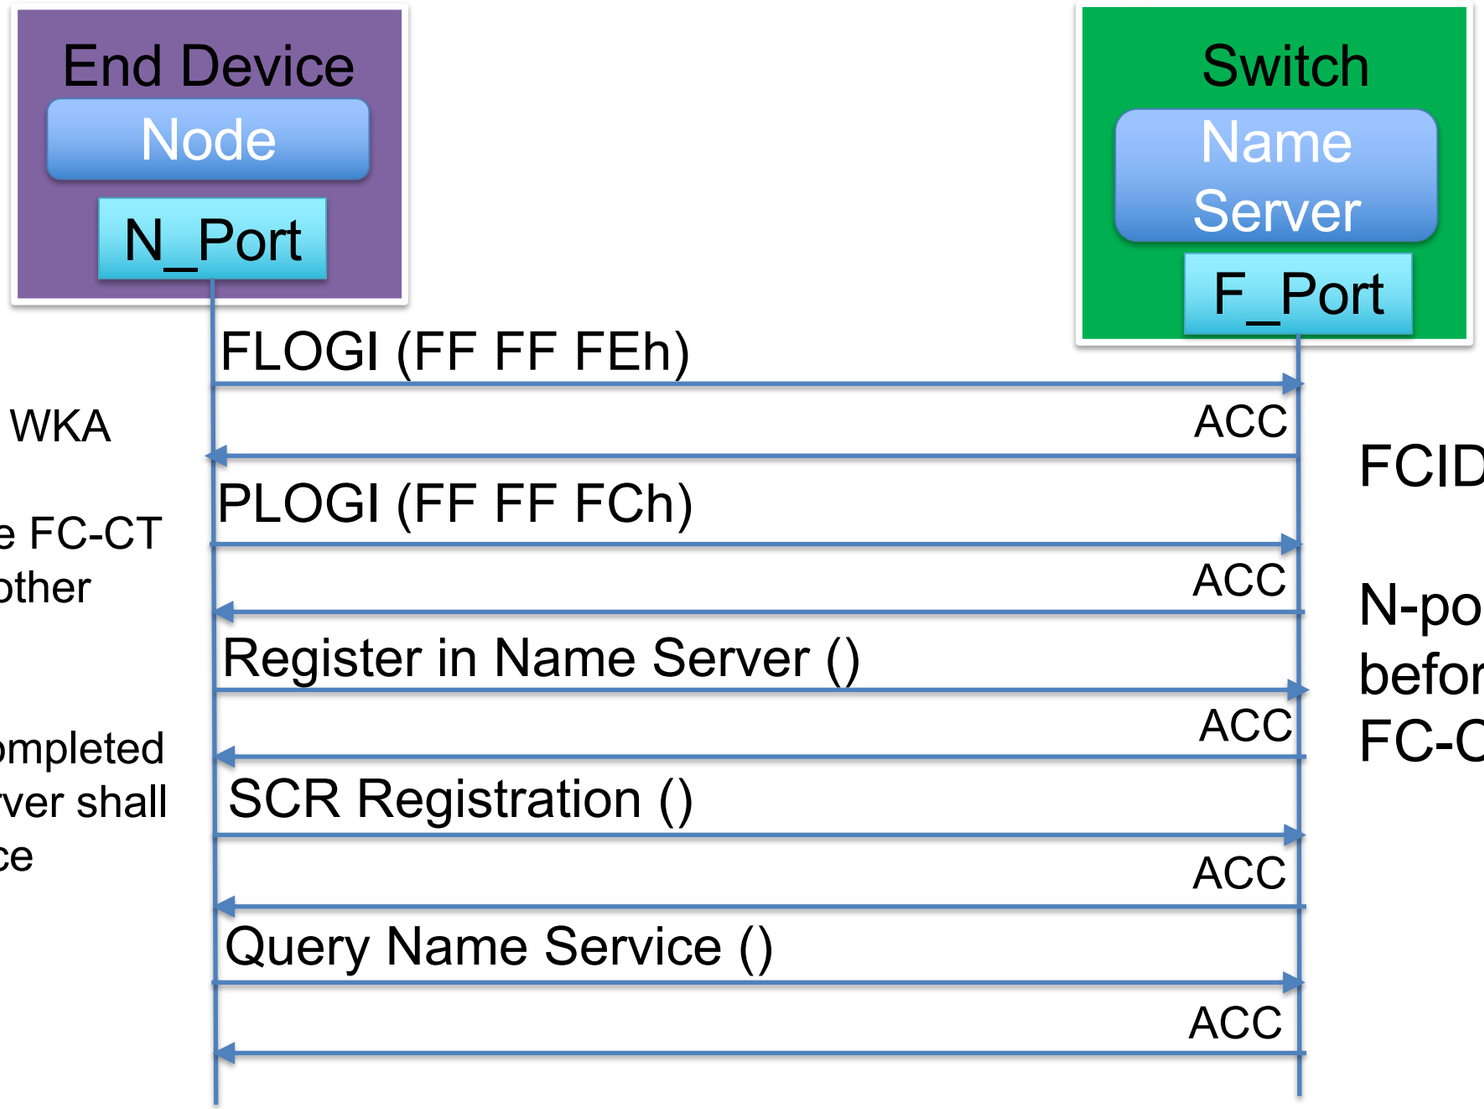

N_Port Login with the WKA

PLOGI required to use FC-CT to communicate with other devices

An N_Port that has completed registration with a Server shall Logout with the Service

FCID assigned

N-ports must register before requesting FC-CT services

PLOGI Registration required by FC-CT service

Port	Side A	Side B	Summary	Bytes	OX_Id	Source	Destination	RX_Id	F_Ctl
FC VCFC-bay5-11(1,3,1)	ExtLinkReq		ExtLinkReq; EX_LNK_SRV; FLOGI;	152	0377	000000	F_Port Controller	FFFF	290000
FC FF30-121(1,3,2)		FC4UCtl	FC4UCtl; FC-SW; ELP;	140	0003	Fabric Controller	Fabric Controller	FFFF	290000
FC VCFC-bay5-11(1,3,1)	ACK_1		ACK_1; Basic Link Service;	36	0003	Fabric Controller	Fabric Controller	037A	E90000
FC VCFC-bay5-11(1,3,1)	ExtLinkReq		ExtLinkReq; EX_LNK_SRV; FLOGI;	152	037B	000000	F_Port Controller	FFFF	290000
FC FF30-121(1,3,2)		ExtLinkRply	ExtLinkRply; EX_LNK_SRV; Accept;	152	037B	F_Port Controller	4A0000	000C	980000
FC VCFC-bay5-11(1,3,1)	ExtLinkReq		ExtLinkReq; EX_LNK_SRV; PLOGI;	152	037D	4A0000	Management Server	FFFF	290000
FC FF30-121(1,3,2)		ExtLinkRply	ExtLinkRply; EX_LNK_SRV; Accept;	152	037D	Management Server	4A0000	1000	993000
FC VCFC-bay5-11(1,3,1)	FC4UCtl		FC4UCtl; FC-GS; GFN;	60	037F	4A0000	Management Server	FFFF	290000
FC VCFC-bay5-11(1,3,1)	FC4UCtl		FC4UCtl; FC-GS; GMAL;	60	0380	4A0000	Management Server	FFFF	290000
FC VCFC-bay5-11(1,3,1)	FC4UCtl		FC4UCtl; FC-GS; GIELN;	60	0381	4A0000	Management Server	FFFF	290000
FC FF30-121(1,3,2)		FC4SCtl	FC4SCtl; FC-GS; Accept;	60	037F	Management Server	4A0000	1002	993000
FC VCFC-bay5-11(1,3,1)	FC4UCtl		FC4UCtl; FC-GS; GPPN;	60	0382	4A0000	Management Server	FFFF	290000
FC FF30-121(1,3,2)		FC4SCtl	FC4SCtl; FC-GS; Accept;	312	0380	Management Server	4A0000	1003	993000
FC FF30-121(1,3,2)		FC4SCtl	FC4SCtl; FC-GS; Accept;	308	0381	Management Server	4A0000	1004	993000
FC FF30-121(1,3,2)		FC4SCtl	FC4SCtl; FC-GS; Accept;	56	0382	Management Server	4A0000	1005	993000
FC VCFC-bay5-11(1,3,1)	FC4UCtl		FC4UCtl; FC-GS; RPL;	584	0387	4A0000	Management Server	FFFF	290000
FC FF30-121(1,3,2)		FC4SCtl	FC4SCtl; FC-GS; Reject; Invalid command code; No additional explanation;	52	0387	Management Server	4A0000	100F	993000
FC VCFC-bay5-11(1,3,1)	ExtLinkReq		ExtLinkReq; EX_LNK_SRV; FDISC;	152	038A	000000	F_Port Controller	FFFF	290000
FC VCFC-bay5-11(1,3,1)	ExtLinkReq		ExtLinkReq; EX_LNK_SRV; FDISC;	152	038F	000000	F_Port Controller	FFFF	290000
FC VCFC-bay5-11(1,3,1)	ExtLinkReq		ExtLinkReq; EX_LNK_SRV; FDISC;	152	0392	000000	F_Port Controller	FFFF	290000
FC VCFC-bay5-11(1,3,1)	ExtLinkReq		ExtLinkReq; EX_LNK_SRV; FDISC;	152	0394	000000	F_Port Controller	FFFF	290000
FC VCFC-bay5-11(1,3,1)	ExtLinkReq		ExtLinkReq; EX_LNK_SRV; FDISC;	152	0398	000000	F_Port Controller	FFFF	290000
FC VCFC-bay5-11(1,3,1)	ExtLinkReq		ExtLinkReq; EX_LNK_SRV; FDISC;	152	039A	000000	F_Port Controller	FFFF	290000
FC VCFC-bay5-11(1,3,1)	ExtLinkReq		ExtLinkReq; EX_LNK_SRV; FDISC;	152	039E	000000	F_Port Controller	FFFF	290000
FC VCFC-bay5-11(1,3,1)	ExtLinkReq		ExtLinkReq; EX_LNK_SRV; FDISC;	152	03A0	000000	F_Port Controller	FFFF	290000
FC VCFC-bay5-11(1,3,1)	ExtLinkReq		ExtLinkReq; EX_LNK_SRV; FDISC;	152	03A4	000000	F_Port Controller	FFFF	290000
FC VCFC-bay5-11(1,3,1)	ExtLinkReq		ExtLinkReq; EX_LNK_SRV; FDISC;	152	03A2	000000	F_Port Controller	FFFF	290000
FC VCFC-bay5-11(1,3,1)	ExtLinkReq		ExtLinkReq; EX_LNK_SRV; FDISC;	152	03A6	000000	F_Port Controller	FFFF	290000
FC VCFC-bay5-11(1,3,1)	FC4UCtl		FC4UCtl; FC-GS; GFN;	60	03A7	4A0000	Management Server	FFFF	290000
FC VCFC-bay5-11(1,3,1)	FC4UCtl		FC4UCtl; FC-GS; GIELN;	60	03A8	4A0000	Management Server	FFFF	290000
FC FF30-121(1,3,2)		FC4SCtl	FC4SCtl; FC-GS; Accept;	60	03A7	Management Server	4A0000	1154	993000
FC FF30-121(1,3,2)		FC4SCtl	FC4SCtl; FC-GS; Accept;	308	03A8	Management Server	4A0000	1155	993000
FC VCFC-bay5-11(1,3,1)	FC4UCtl		FC4UCtl; FC-GS; RPL;	584	03AD	4A0000	Management Server	FFFF	290000
FC FF30-121(1,3,2)		FC4SCtl	FC4SCtl; FC-GS; Reject; Invalid command code; No additional explanation;	52	03AD	Management Server	4A0000	1158	993000
FC VCFC-bay5-11(1,3,1)	ExtLinkReq		ExtLinkReq; EX_LNK_SRV; FDISC;	152	03B2	000000	F_Port Controller	FFFF	290000

CT services

A client originates an FC-CT exchange

- transmits the request sequence (transferring sequence initiative)
- waits for the response sequence
- sends the next request

- This provides an interlocked, half-duplex operation with the server because only one exchange and sequence is open at a time

A Common Transport IU (CT_IU) is the common Fibre Channel Sequence used to transfer all information between a Client and a Server

The first part of the CT_IU contains a preamble with information common to all CT_IUs

One or more of the following optional preambles may be used:

a) Extended CT_IU preamble
b) Vendor Specific preamble

FC-CT protocol

- Each generic service operation is an exchange consisting of one or more IUs
- Each IU is delivered using a Sequence of one or more frames
  - the actual number of frames depends on the size of the IU
- There are two IUs in the FC-CT protocol:
  - Request – a client sends a request to the generic service provider. The request consists of a single FC-CT IU containing the client request
  - Response – The generic service provider sends a response to the client. The response transfers a single FC-CT IU containing a response to the client request.
    - two Response CT_IUs - Accept CT_IU and Reject CT_IU

- The error policy is to abort a failed sequence and discard all subsequent sequences of that exchange (Abort, discard multiple sequences).

CT Confidentiality and Authentication

CT_Authentication may be used between a Client and a Server

- to authenticate CT_IU Requests and Responses
- optionally to encrypt CT_IU Requests and Responses

CT_Authentication is provided through the use of the extended CT_IU preamble

- this mechanism provides a means to validate that requests and responses are transferred without change
- distinguishes these from requests and responses that are corrupted by an intermediary agent or generated by a non-authorized CT (i.e., integrity and anti-replay protection)
- optionally this mechanism provides confidentiality to CT_IUs

CT Confidentiality

- The CT Confidentiality transform is a cryptographic transform of the CT_IU payload, used in combination with the extended CT_IU preamble

Directory Server

Services provided by the Directory server WKA **FF FF FCx**

- **Name Server is a function of the directory server and is available on any FC topology**
- The name server provides a rich set of commands to:
  - register information in the database
  - remove information from the database
  - query information in the database utilizing query parameters

- Directory server is a node port and requires N_Port Login (PLOGI) to gain directory services access
- The FC-CT protocol supports Directory server request and response exchange of IUs
- The Directory Service provides a means to discover information about fabric attached nodes and Nx_Ports

The consumer of a Directory Service is normally a “device driver” or other internal layer of an OS

The information provided by a Directory Service is operational, and therefore may be constrained by the operational environment (i.e., Zone) of the node

The Directory Service provides access to the distributed database in every switch/fabric

FC Common Transport (FC-CT)

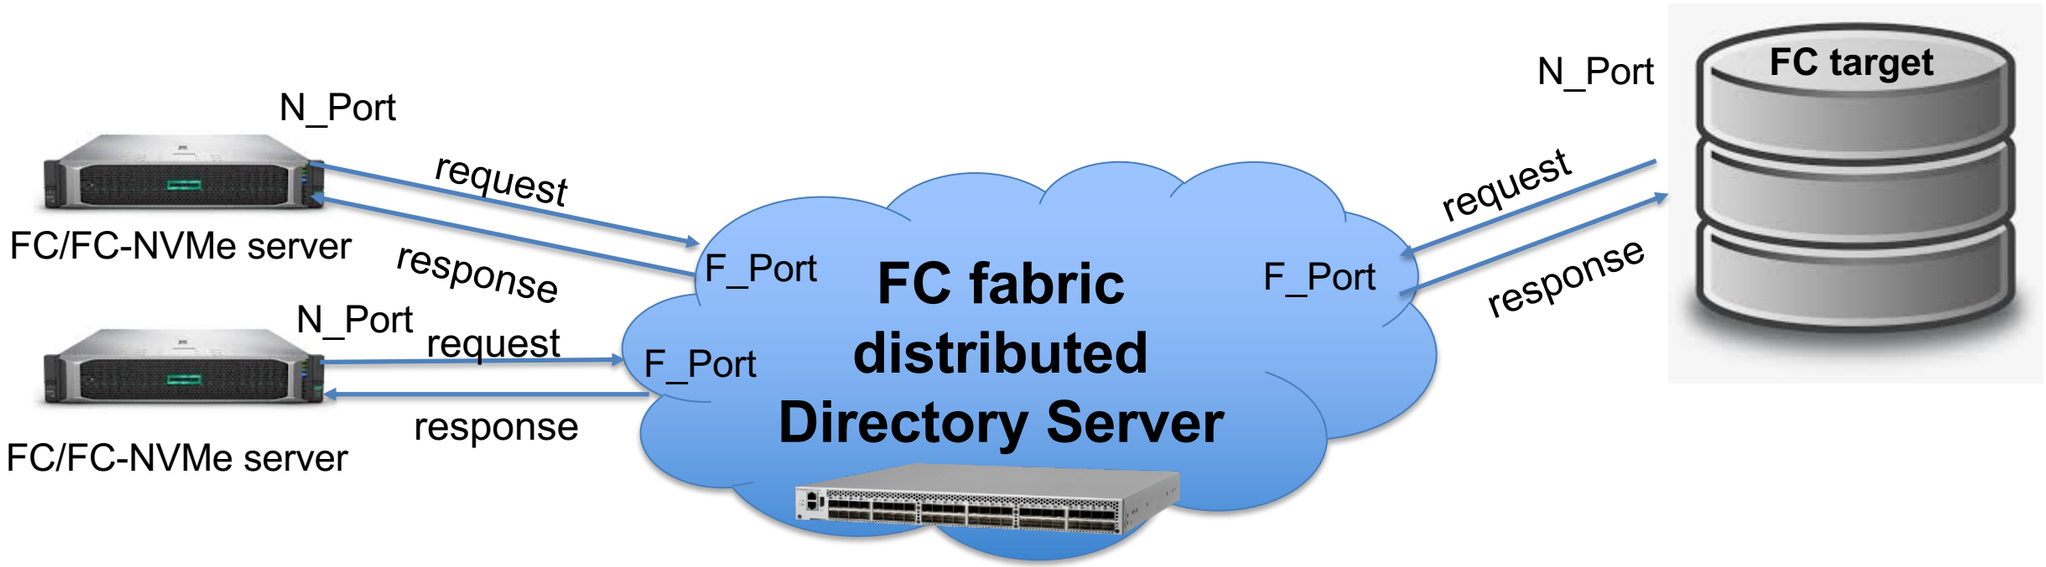

Node Port sends Requests to WKA FF FF FCh – which is a Node Port
Name Server Request is formatted as an FC-CT frame
Name Server – SubType = 02h
Typical requests are Register FC-4 type and Get IDs

- Each switch provides a server service instance
- Collectively the switches provide services to N_Ports subscribed using WKA
- Server instances utilize Domain Controller addresses xFF FC xx for inter server service communication

Name Server

Requests for the Name Server are carried over the FC-CT

The Name Server is a database of information about Nx_Ports and provides a way for Nx_Ports to register and discover Fibre Channel attributes

Note: **No Name Server scale limit is defined in the standard, it scales with the fabric**

Once registered, the attributes are made available to requestors using four types of object requests

1. Get
2. Register
3. Deregister
4. Delimiter

- Available Objects:

port identifier	port name	node name	class of service
FC-4 type	symbolic port and node name	port type	fabric port name
FC4 features	hard address	permanent port name	

The Name Server is intended to be distributed among Switches, making the Name Server immediately available to Nx_Ports once they have successfully completed Fabric Login

Name Server

Some information may be registered in the Name Server database when a node port logs in with the fabric using Fabric login (FLOGI).

Other information must be registered explicitly by the node port in order to appear in the database

Node ports query the Name Server to obtain information about other node ports

- For example, a SCSI initiator requests a list of all node ports supporting SCSI FCP

When zoning is in effect, responses to queries are filtered so that a requesting N_Port only receives information about ports within its zone(s)

Name Server protocol

- Name Server registration, deregistration, queries, and delimiter based actions are managed through protocols containing a set of request CT_IUs and response CT_IUs supported by the Name Server
- For a Name Server request, the Name Server payload is transported from the requestor to the Name Server using a request CT_IU
- The corresponding Name Server response is transported from the Name Server to the requestor, in the Exchange established by the requestor, using a response CT_IU
- If Zones exist within the Fabric:
  - the Name Server restricts access to information in the Name Server database based on the Zone configuration

Name Server data Example: nsshow

Type Pid COS PortName NodeName SCR
N 010000; 3;51:40:2e:c0:01:17:90:82;51:40:2e:c0:01:17:90:83; 0x00000003
SCR: Fabric-Detected Nx-Port-Detected
FC4s: **FCP FC-NVMe**
NodeSymb: [39] "SN1600Q FW:v8.08.230 DVR:v10.01.00.16-k"
Fabric Port Name: 20:00:c4:f5:7c:b5:b1:e1
Permanent Port Name: 51:40:2e:c0:01:17:90:82
Device type: Physical Initiator
Port Index: 0
Share Area: No
Redirect: No
Partial: No
LSAN: No
Slow Drain Device: No
Device link speed: 32G
Connected through AG: No
Real device behind AG: No
FCoE: No
FC4 Features [FCP]: Initiator

Type Pid COS PortName NodeName SCR
N 010300; 3;10:00:08:f1:ea:ba:c2:d3;20:00:08:f1:ea:ba:c2:d3; 0x00000003
SCR: Fabric-Detected Nx-Port-Detected
FC4s: **FCP FC-NVMe**
PortSymb: [34] "Emulex PPN-10:00:08:F1:EA:BA:C2:D3“
NodeSymb: [69] "Emulex SN1610E2P FV12.4.270.5 DV12.4.243.4 HN:R82-S09 OS:Windows 2016“
Fabric Port Name: 20:03:c4:f5:7c:b5:b1:e1
Permanent Port Name: 10:00:08:f1:ea:ba:c2:d3
Device type: Physical Initiator
Port Index: 3
Share Area: No
Redirect: No
Partial: No
LSAN: No
Slow Drain Device: No
Device link speed: 32G
Connected through AG: No
Real device behind AG: No
FCoE: No
FC4 Features [FCP]: Initiator
FC4 Features [FC-NVMe]: Initiator

Name Server dataExample: display fc name-service database

VSAN: 400 FCID: 0x4e002a
Port-WWN(vendor): 10:00:14:02:ec:ec:4d:aa
Node-WWN: 20:00:14:02:ec:ec:4d:aa
Class: 3
Node-IP-addr: 0.0.0.0
FC4-types(FC4_features): **SCSI-FCP:Initiator**
NVMeoFC:Initiator
Symbolic-port-name: Emulex PPN-10:00:14:02:EC:EC:4D:AA
Symbolic-node-name: Emulex SN1600E2P FV12.2.226.0 DV12.2.207.0 HN:RACK106-S01 OS:Windows 2012 R2
Port-type: 0x01(N)
Fabric-port-WWN: 28:01:d8:94:03:f2:b4:c7
Hard-addr: 0x000000

VSAN: 400 FCID: 0x4e002d
Port-WWN(vendor): 51:40:2e:c0:00:a1:78:d4
Node-WWN: 51:40:2e:c0:00:a1:78:d5
Class: 3
Node-IP-addr: 0.0.0.0
FC4-types(FC4_features): **SCSI-FCP:Initiator**
Symbolic-port-name: QLogic Port0 WWPN 51:40:2e:c0:00:a1:78:d4
Symbolic-node-name: SN1100Q FW:v8.08.230 DVR:v9.3.3.20 A1
Port-type: 0x01(N)
Fabric-port-WWN: 28:01:d8:94:03:f2:b4:cb
Hard-addr: 0x000000

Management Server

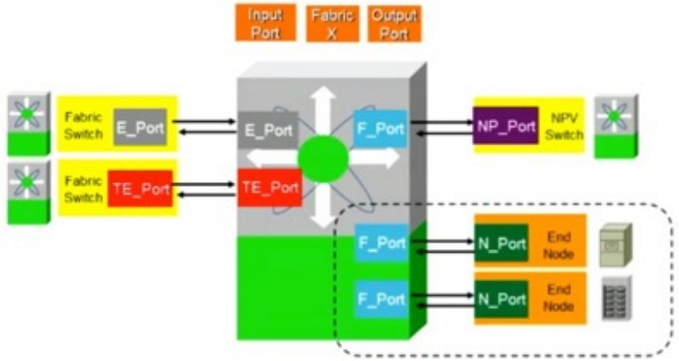

Services provided by the management server WKA **FF FF FAh**

- fabric configuration
- create and manage zones
- access the name server database without zone and zone set restrictions

- Management server is a node port and requires N_Port Login (PLOGI) to gain directory services access
- The FC-CT protocol supports Management server request and response exchange of IUs.
- The client management application uses the Management server to
  - determine the fabric configuration
  - create and manage zones within the fabric
  - access the Name Server database without zoning restrictions

The Fabric Configuration server provides configuration information to management clients

- This information identifies the interconnect elements (switches) within the fabric, platforms containing one or more nodes (each containing one or more node ports), and the connections between them.

Management Service

The Management Service provides a single management access point within the FC Fabric, the standard defines the model for requests and responses to access Management Service information.

- FC-GS-8 does not define the structure of this information.

Management Service covers the following areas:

1. Fabric Configuration Server - provides for the configuration management of the Fabric
2. Unzoned Name Server - provides access to Name Server information that is not subject to Zone constraints
3. Fabric Zone Server - provides access to and zone management and administration

✓ A management application can create or remove zones, add or remove members from the zone, group zones into zone sets, and activate a specific zone set

1. Security Policy Server - provides distribution of security policies
2. Security Information Server - provides access to operational security information
3. Fabric Device Management Interface (FDMI) - provides access to data associated with attached devices
4. Enhanced Fabric Configuration Server - provides for the configuration management of the Fabric
5. Application Server - provides for application specific management (e.g. application identifier management)

Basic and Enhanced Zoning

Fabric Zone Server is a control interface for Basic and Enhanced Zoning

If all Switches in a Fabric support the Enhanced request set, then it may be used.

- Otherwise only the Basic request set may be used

Two control management requests are defined to deal with these two distinct request sets:

a) GFEZ (Get Fabric Enhanced Zoning Support) to provide a management application with the ability to query the Fabric about its support for Enhanced Zoning

b) SFEZ (Set Fabric Enhanced Zoning Support) to provide a management application with the ability to change the Zoning operational mode of the Fabric from Basic to Enhanced.

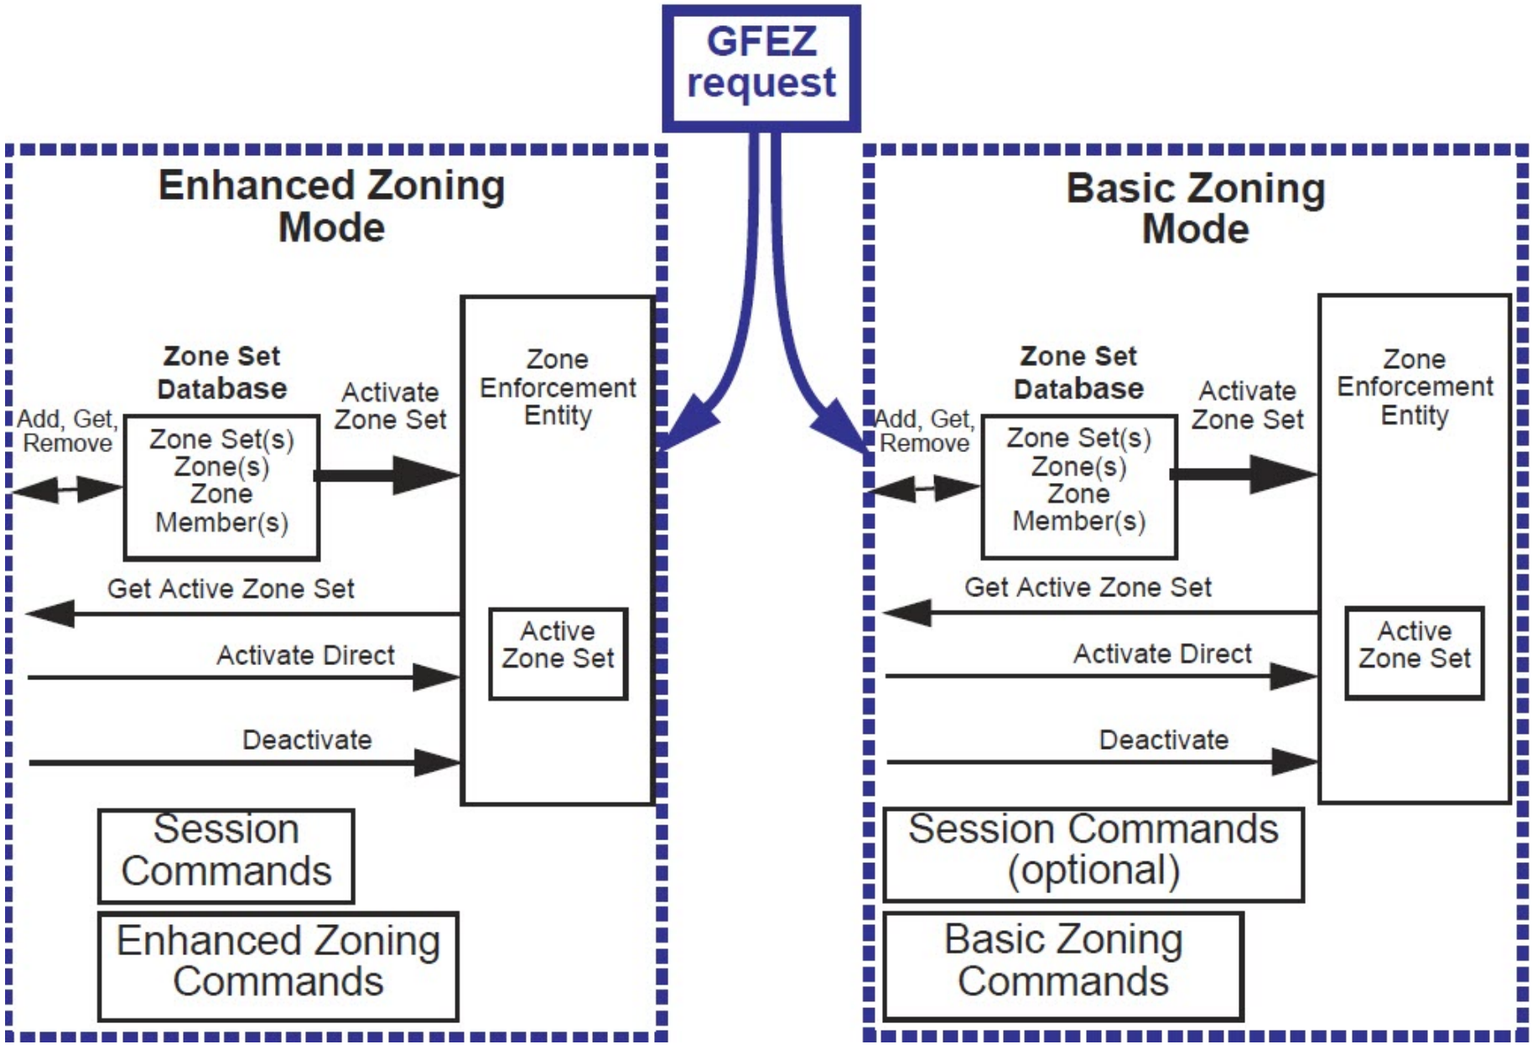

Fibre Channel Generic Services and Switch Fabric functionality

- A Distributed Services model is used to allow a Fabric to provide consistent services to all attached N_Ports
- A Distributed Server is “contained” in a Switch
  - In this case the term “contain” does not mean the entity is physically inside the Switch (i.e., it may be physically outside the Switch)
- Distributed Services are mapped to a common framework with goal to provide:
  - a consistent method for distributing services across Switches in a Fabric
  - a distribution method that is topology independent
  - a method that preserves processing facilities for existing frame formats

Distributed Services

- Features of a Distributed Server include:
  - A transport
  - Common characteristics
  - Work categories
  - Frame formats
- If Zoning is present in a Fabric, Distributed Services may be affected, with the following rules:
  - Switch-to-Switch communications are not Zoned
    - Only applies to the Class F CT Header based Distributed Services frames
  - Zoning is applied by the Entry Switch
    - If a particular Distributed Service is affected by Zoning, it is the responsibility of the Entry Switch to make sure that a requesting N_Port does not receive data for that Distributed Service that is outside of the N_Port's Zone

Distributed Services Transport

- Generic Service requests and responses are transported between Distributed Servers using the Common Transport (FC-CT) defined in FC-GS-8
- CT frames are sent using Class F service
  - D_ID is set to the Domain Controller Identifier of the destination Switch
  - S_ID is set to the Domain Controller Identifier of the source Switch
  - TYPE field is set to 20h, indicating Fibre Channel Fabric Generic Services

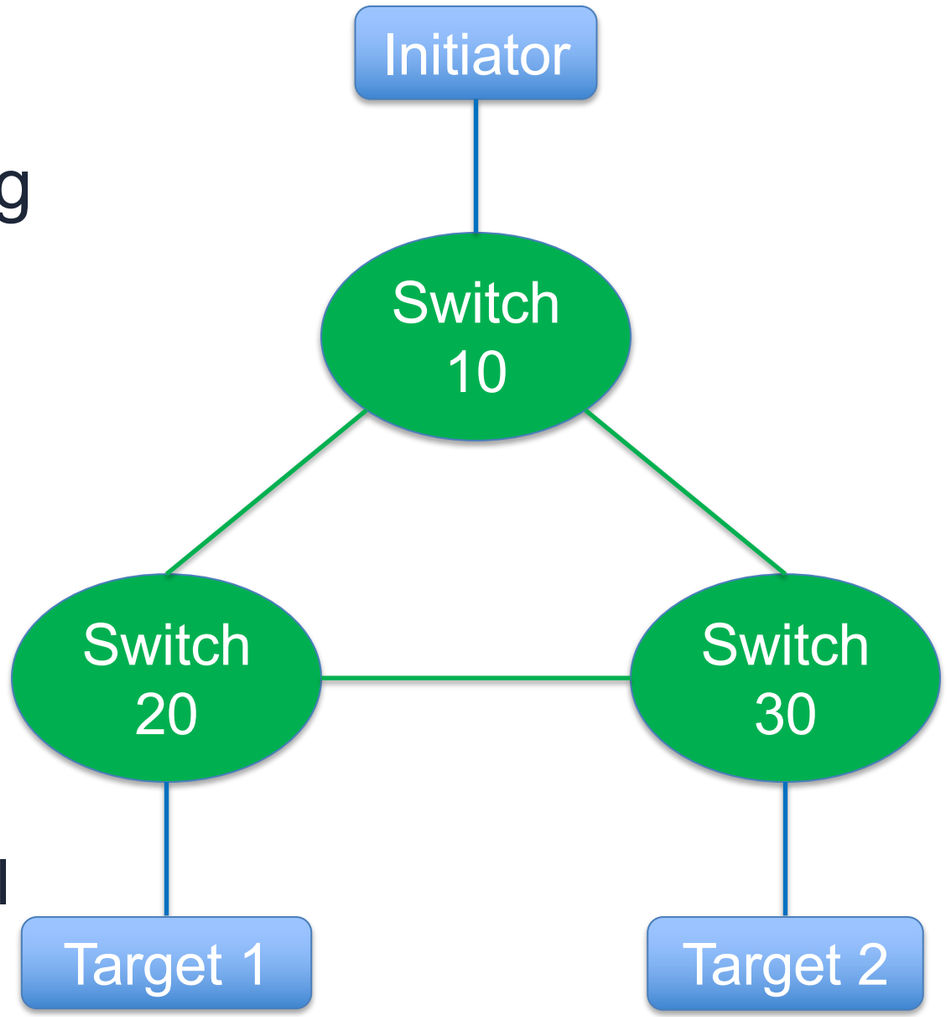

Distributed Services Transport

- For a Distributed Services request, a remote Switch does not send a response directly to the requesting N_Port
  - All responses are sent to the Entry Switch
  - Entry Switch is responsible sending the appropriate response to the requesting N_Port
- An N_Port communicates with a Distributed Service via the WKA of the Distributed Service
- N_Ports do send Distributed Service requests to Domain Controllers
  - Distributed Services requests/responses are sent between Switches, not between Switch and N_Port

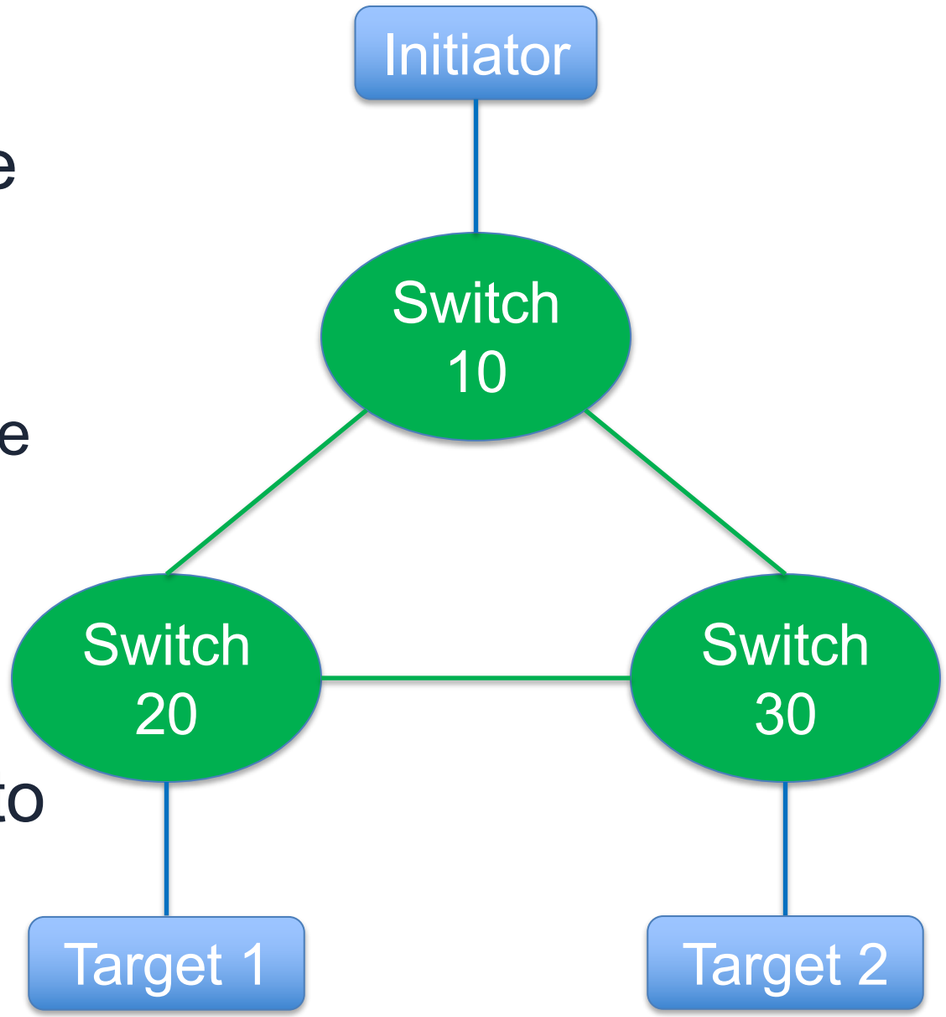

Distributed Services Common Characteristics

- Each Distributed Service shares a set of common characteristics
  - Timeouts
    - For requests between Switches, the timeout value is 5 seconds
  - Local data copies
    - Local data copies may be optionally allowed by a Distributed Service. If a Distributed Service allows local data copies it also specifies the method by which the integrity of the local copied data is maintained
  - Exchange management
    - Each request between Switches is mapped to a unique Exchange
    - Multiple outstanding requests are allowed between a pair of Switches up to the end-to-end credit resources specified by the receiving Switch

Distributed Services Common Characteristics

- Each Distributed Service shares a set of common characteristics
  - Responses
    - Each request sent receives a response
    - If the receiving Switch is unable to perform an operation, it responds with a Reject CT_IU specifying a reason code and reason code explanation
    - If a response is not received from all Switches to which a request was sent within the timeout period, then the request is considered partial and a response is be sent back to the N_Port as appropriate for the Service

Distributed Services Common Characteristics

- Each Distributed Service shares a set of common characteristics
  - Partial response
    - For many requests even a partial response to the requesting N_Port is useful
    - A partial response may occur for a number of reasons
      - one of the Switches a request is directed to is busy and unable to respond within the timeout period
      - one of the Switches a request is directed to does not support the service requested
    - A service may allow partial responses for a subset of its requests
    - If the response to a request is partial, the partial response bit is set to one in the CT Header of the response back to the N_Port

Distributed Services Common Characteristics

- Each Distributed Service shares a set of common characteristics
  - Data merge
    - Describes how data from multiple responses is consolidated
  - Error recovery
    - If an error on a Distributed Services frame is detected (e.g., No ACK, P_BSY), the frame may be retransmitted for a time interval up to 5 seconds

Distributed Services Work Categories

- Work categories are definitions that allow consistent mapping of services to Distributed Services
- These categories define how each Distributed Service maps its commands given the Distribution characteristic
- The work categories are:
  - Local
  - 1 to 1
  - 1 to Many
  - 1 to All

Distributed Services Work Categories - Local

- Local requests are those that may be handled entirely by the Entry Switch
- A request is local for the following reasons:
  - The data being requested is owned entirely by the Entry Switch
    - This situation is dependent on the type of request
  - The Entry Switch is maintaining a local copy of the data being requested
    - This situation may occur for any request depending on the local data copy rules of the Distributed Service to which the request belongs
- Any request that is determined to be local is processed as appropriate for the service as defined in FC-GS-8

Distributed Services Work Categories

- 1-to-1
  - A request that is unable to be handled entirely by the Entry Switch, and the Entry Switch has identified a single remote Switch that may handle the request
  - The local Switch sends the request frame directly to the Domain Controller of remote Switch

Distributed Services Work Categories

- 1-to-Many
  - A request that is unable to be handled entirely by the Entry Switch, and the Entry Switch has identified multiple remote Switches that may handle the request
  - The local Switch sends request frames directly to the Domain Controller of all remote Switches that it has identified to contain requested data

Distributed Services Work Categories

- 1-to-All
  - A request that is unable to be handled entirely by the Entry Switch, and the Entry Switch is unable to identify the set of remote Switches to query
  - The Entry Switch sends request frames directly to the Domain Controller of all Switches in the Fabric

Distributed Services Frame Formats

- Where possible Distributed Services use the same frame formats for Switch-to-Switch communications as are used for N_Port requests

Distributed Services
Distributed Name Server

- Each Switch contains its own resident Name Server, called a distributed Name Server (dNS)
- Each dNS within a Switch is responsible for the name entries associated with the Domain(s) assigned to the Switch
- Each dNS within a Switch only returns information associated with the Domain(s) for which the Switch is responsible
- A client N_Port sends its Name Service request to the Entry Switch via the Well-Known Address
- If the required information is not available locally, the dNS within the local Switch services the request by making any needed requests of other dNS's in other Switches

Distributed Services
Distributed Name Server

- A dNS may maintain local data copies
  - Integrity of locally copied data is maintained via SW_RSCN notification
  - This implies all Switches distribute SW_RSCN notifications throughout the Fabric whenever a change takes place in their local dNS database
- Communication between dNS's to acquire requested information is transparent to the original requesting client
- Partial responses to dNS queries are allowed
  - If an Entry Switch sends a partial response back to an N_Port it sets the partial response bit in the CT Header

Distributed Services
Distributed Management Server

- For each Management Server of the Management Service, each Switch contains its own instance of the Server. Generically, each Server instance is called a distributed Management Server (dMS)
- Each dMS within a Switch is responsible for the entries associated with the Domain(s) assigned to the Switch
- A client N_Port sends its Management Server request to the Entry Switch via the Well-Known Address and appropriate sub-type
- If the required information is not available locally, a dMS within the Entry Switch services the request by making any needed requests of other dMS instances contained by the other Switches

Distributed Services
Distributed Management Server

- A dMS may maintain local data copies, and a dMS notifies other dMS's that they should remove local data copies
- Communication between dMS's to acquire requested information is transparent to the original requesting client
- Partial responses for some dMS requests are allowed
- Responses returned to a client for some dMS Servers are not subject to Zoning
  - The impact of Zoning on the Security Information Server is separately specified for each Security Information request

Distributed Services
Distributed VE Identification Server

- Each Switch contains its own resident VE Identification Server, called a distributed VE Identification Server (dVEIS)
- Each dVEIS within a Switch is responsible for the VE mappings associated with the Domain(s) assigned to the Switch
- A client N_Port sends its VE Identification Server request to the Entry Switch via the Well-Known Address and appropriate sub-type
- If the required information is not available locally, a dVEIS within the Entry Switch services the request by making any needed requests of other dVEIS instances contained by the other Switches

Distributed Services
Distributed VE Identification Server

- A dVEIS may maintain local VE mapping copies, and a dVEIS notifies other dVEIS's of changes in the VE mappings
- Communication between each dVEIS to acquire requested information is transparent to the original requesting client

Potential Fabric Services

- A Congestion Management Service or Server may be developed in the future as a result of functionality related to Fabric performance that was recently approved
  - New Fabric Performance Impact Notification (FPIN) ELS was added to FC-LS-5 draft standard
    - Link Integrity Notification
    - Delivery Notification
    - Peer Congestion Notification
    - Congestion Notification
  - New Primitive Signals were added to FC-FS-6 draft standard
    - ARB(F1) – Warning Congestion Signal
    - ARB(F7) – Alarm Congestion Signal

Summary

Fibre Channel Fabric Services provide for a robust, scalable, and manageable Storage Area Network ☺

Q & A

After this Webcast

- Please rate this event – we value your feedback
- We will post a Q&A blog at http://fibrechannel.org/ with answers to the questions we received today
- Follow us on Twitter @FCIAnews for updates on future FCIA webcasts
- Visit our library of FCIA on-demand webcasts at http://fibrechannel.org/webcasts/ to learn about:
  - Fibre Channel Fundamentals
  - FC-NVMe
  - Long Distance Fibre Channel
  - Fibre Channel Speedmap
  - FCIP (Extension): Data Protection and Business Continuity
  - Fibre Channel Performance
  - FICON
  - Fibre Channel Cabling
  - 64GFC
  - FC Zoning Basics

Our Next FCIA Webcast:

What's New in FC-NVMe-2

October 15, 2020

Register at: https://www.brighttalk.com/webcast/14967/435209

Thank You