Inside a Modern Fibre Channel Architecture – Part 1

Live Webcast
September 15, 2021
10:00 AM PT/1:00 PM ET

FIBRE CHANNEL INDUSTRY ASSOCIATION

Today's Speakers

David Rodgers

FCIA BoD

Teledyne LeCroy

Patty Driever

FCIA BoD

Distinguished Engineer

Storage and Networking Solutions & Strategy

IBM

David Peterson

Principal Engineer

Broadcom

Craig Carlson

FCIA BoD

Senior Technologist

Marvell

About the FCIA

"The Fibre Channel Industry Association (FCIA) is a mutual benefit, non-profit, international organization of manufacturers, system integrators, developers, vendors, industry professionals, and end users."

About the Fibre Channel Industry Association (FCIA)

25+ Years
Promoting Fibre Channel Technology

Industry Leading
Member Companies

142M+ FC Ports
Shipped Since 2001

Agenda

- Overview
- Functional levels (FC-0, FC-1, FC-2, FC-3, FC-4)
- FC node architectural components
- Physical model
- Communication models (simplex, full-duplex, half-duplex)
- Interconnect topologies
- Classes of service
- Fabric model
- Generic Services

Overview

- Fibre Channel is a bi-directional, point-to-point, serial data communication channel, architected for high performance

- Fibre Channel may be implemented using any combination of the following three topologies:
  - a point-to-point link between two ports
  - a set of ports interconnected by a switching network called a Fabric
  - a set of ports interconnected with a loop topology
    - Loop topology is no longer is in wide use

Overview

- Fibre Channel provides a general transport for Upper Level Protocols (ULPs) such as:
  - Single-byte Common Command Set (SBCCS)
  - Small Computer System Interface (SCSI)
  - NVM Express (NVMe)
  - Internet Protocol (IP)
- Fibre Channel protocol provides many implementation possibilities…from minimum cost to maximum performance
- Transmission medium is isolated from control protocol so each implementation may use a technology best suited for the environment

Overview

- Effective transfer rate achieved by an FC configuration is determined by:
  - Physical variants
  - Communication model
  - Payload size
  - Speed of the Fibre link
  - Class of service
  - Overhead

Functional Levels

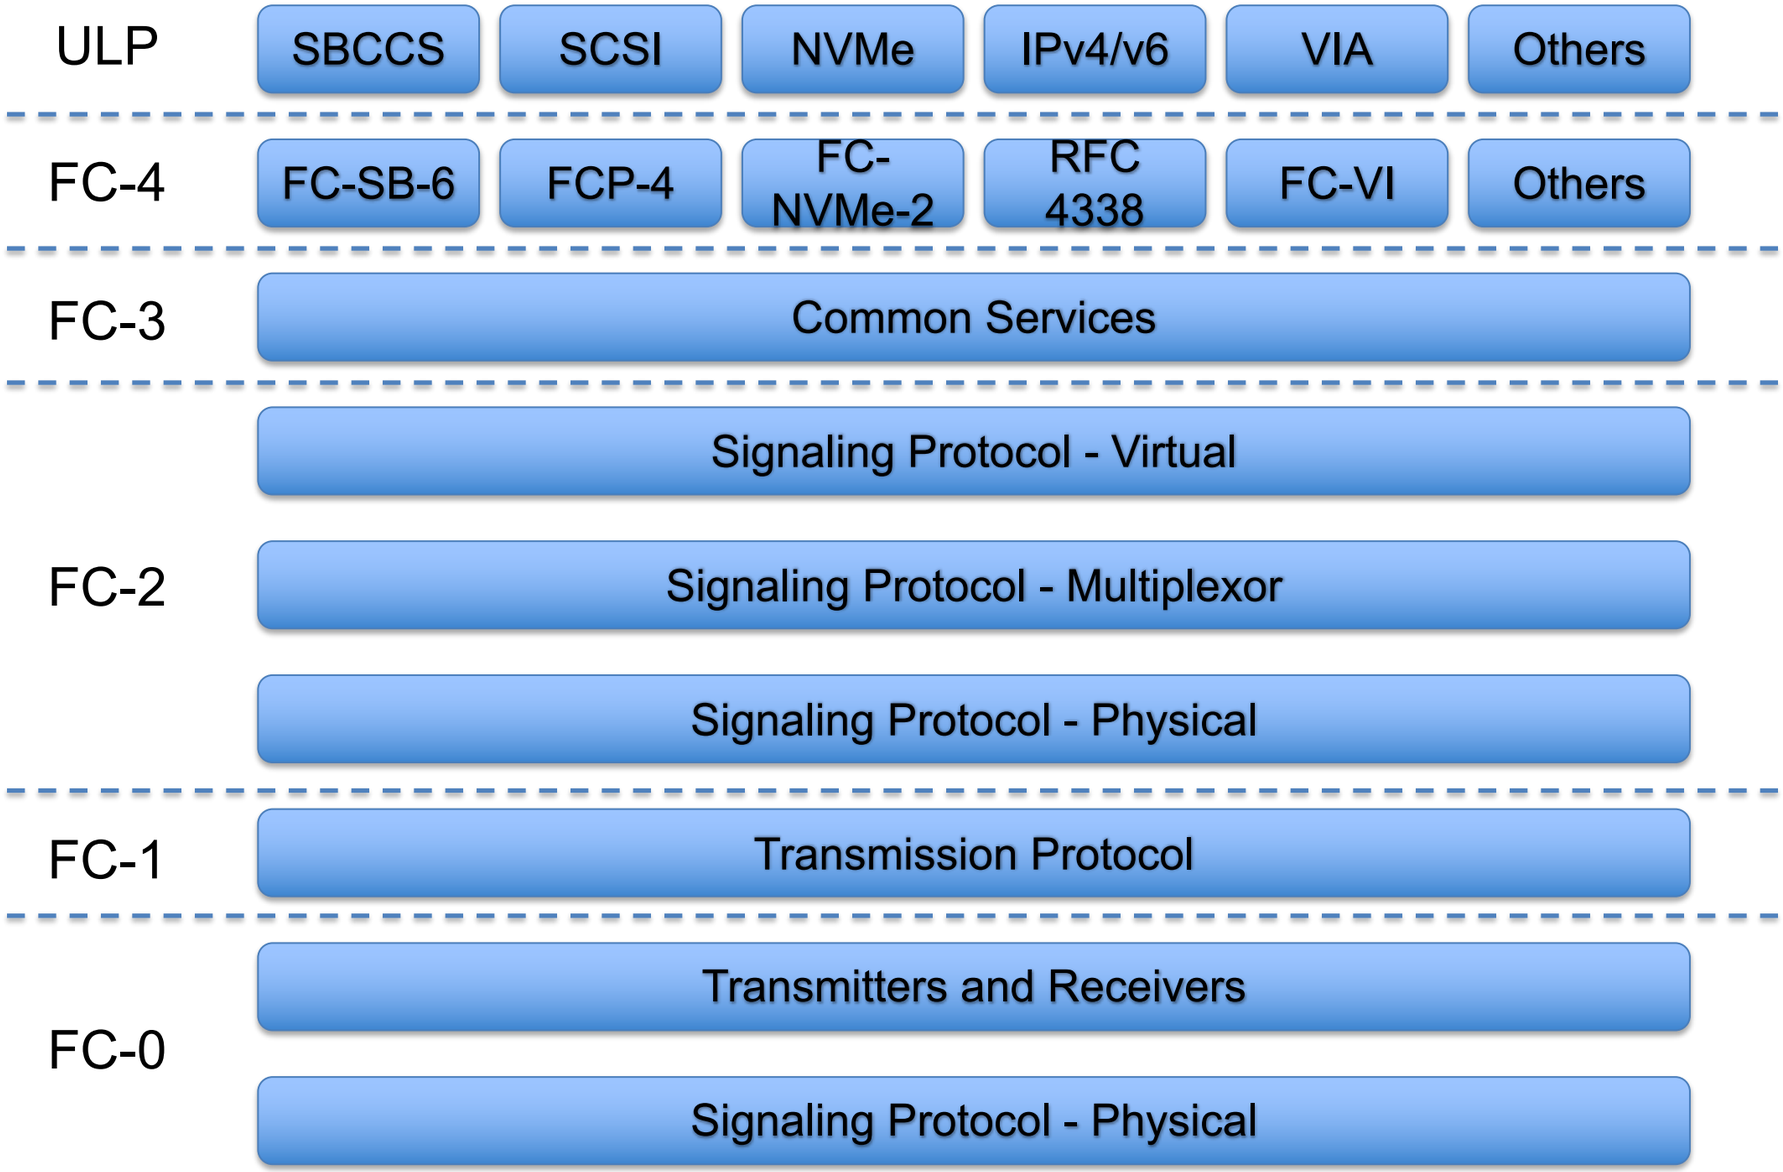

Functional Levels

- FC-0
  - the physical interface (FC-0) consists of transmission media, transmitters, and receivers and their interfaces
  - physical media, associated drivers and receivers capable of operating at various speeds are specified in the physical layer standards
    - FC-PI-x
    - FC-BaseT

Functional Levels

- FC-1
  - defines the transmission protocol that is used for FC-0 levels specified in FC-PI-x and FC-BaseT
  - includes the serial encoding, decoding, and error control
  - FC transmits information using either a 64B/66B transmission code or an adaptive 8B/10B transmission code
  - encoding process results in the generation of Transmission Words
  - specific encoded bit patterns, referred to as Ordered Sets, are designated to have special meaning
  - Ordered Sets are used by FC-2P sublevel to identify frame boundaries, transmit primitive function requests, and by FC-1 level to maintain proper link transmission characteristics during periods of inactivity

Functional Levels

- FC-1
  - transmitter and receiver behavior is specified via a set of states and their interrelationships
  - the states are divided into operational and not operational classes
  - error monitoring capabilities and special operational modes are also defined for operational receivers and transmitters

Functional Levels

- FC-2
  - serves as the transport mechanism for Fibre Channel
  - transported data is transparent to FC-2, and visible to FC-3 and above
  - three sublevels:
    - FC-2P - the FC-2 Physical sublevel
    - FC-2M - the FC-2 Multiplexer sublevel
    - FC-2V - the FC-2 Virtual sublevel

Functional Levels

- FC-2
  - FC-2P
    - specifies the rules and provides mechanisms that shall be used to transfer frames via a specific FC-1 level
    - functions include frame transmission and reception, buffer-to-buffer flow control, and clock synchronization using Primitive Signals
  - FC-2M
    - specifies addressing and functions used to route frames between a Link Control Facility (LCF) and a VN_Port
  - FC-2V
    - specifies facilities and functions that an Nx_Port may provide for FC-4 usage, regardless of the FC-1 that is used
    - functions include several classes of service, frame content construction and analysis, Sequence disassembly and reassembly, Exchange management, and Name_Identifiers

Functional Levels

- FC-3
  - provides a set of services that are common across multiple Nx_Ports of a node
  - includes protocols for Basic Link Services and Extended Link Services
  - Link Services represent a mandatory function required by FC-2

Functional Levels

- FC-4
  - the highest level in the Fibre Channel standards set
  - defines the mapping between the FC lower levels and an Upper Level Protocol
    - SCSI command set
    - SBCCS command set
    - NVMe command set
    - IPv4/IPv6
    - other Upper Level Protocols (ULPs)
      - FC provides a method for supporting many ULPs

FC Node

Architectural Components

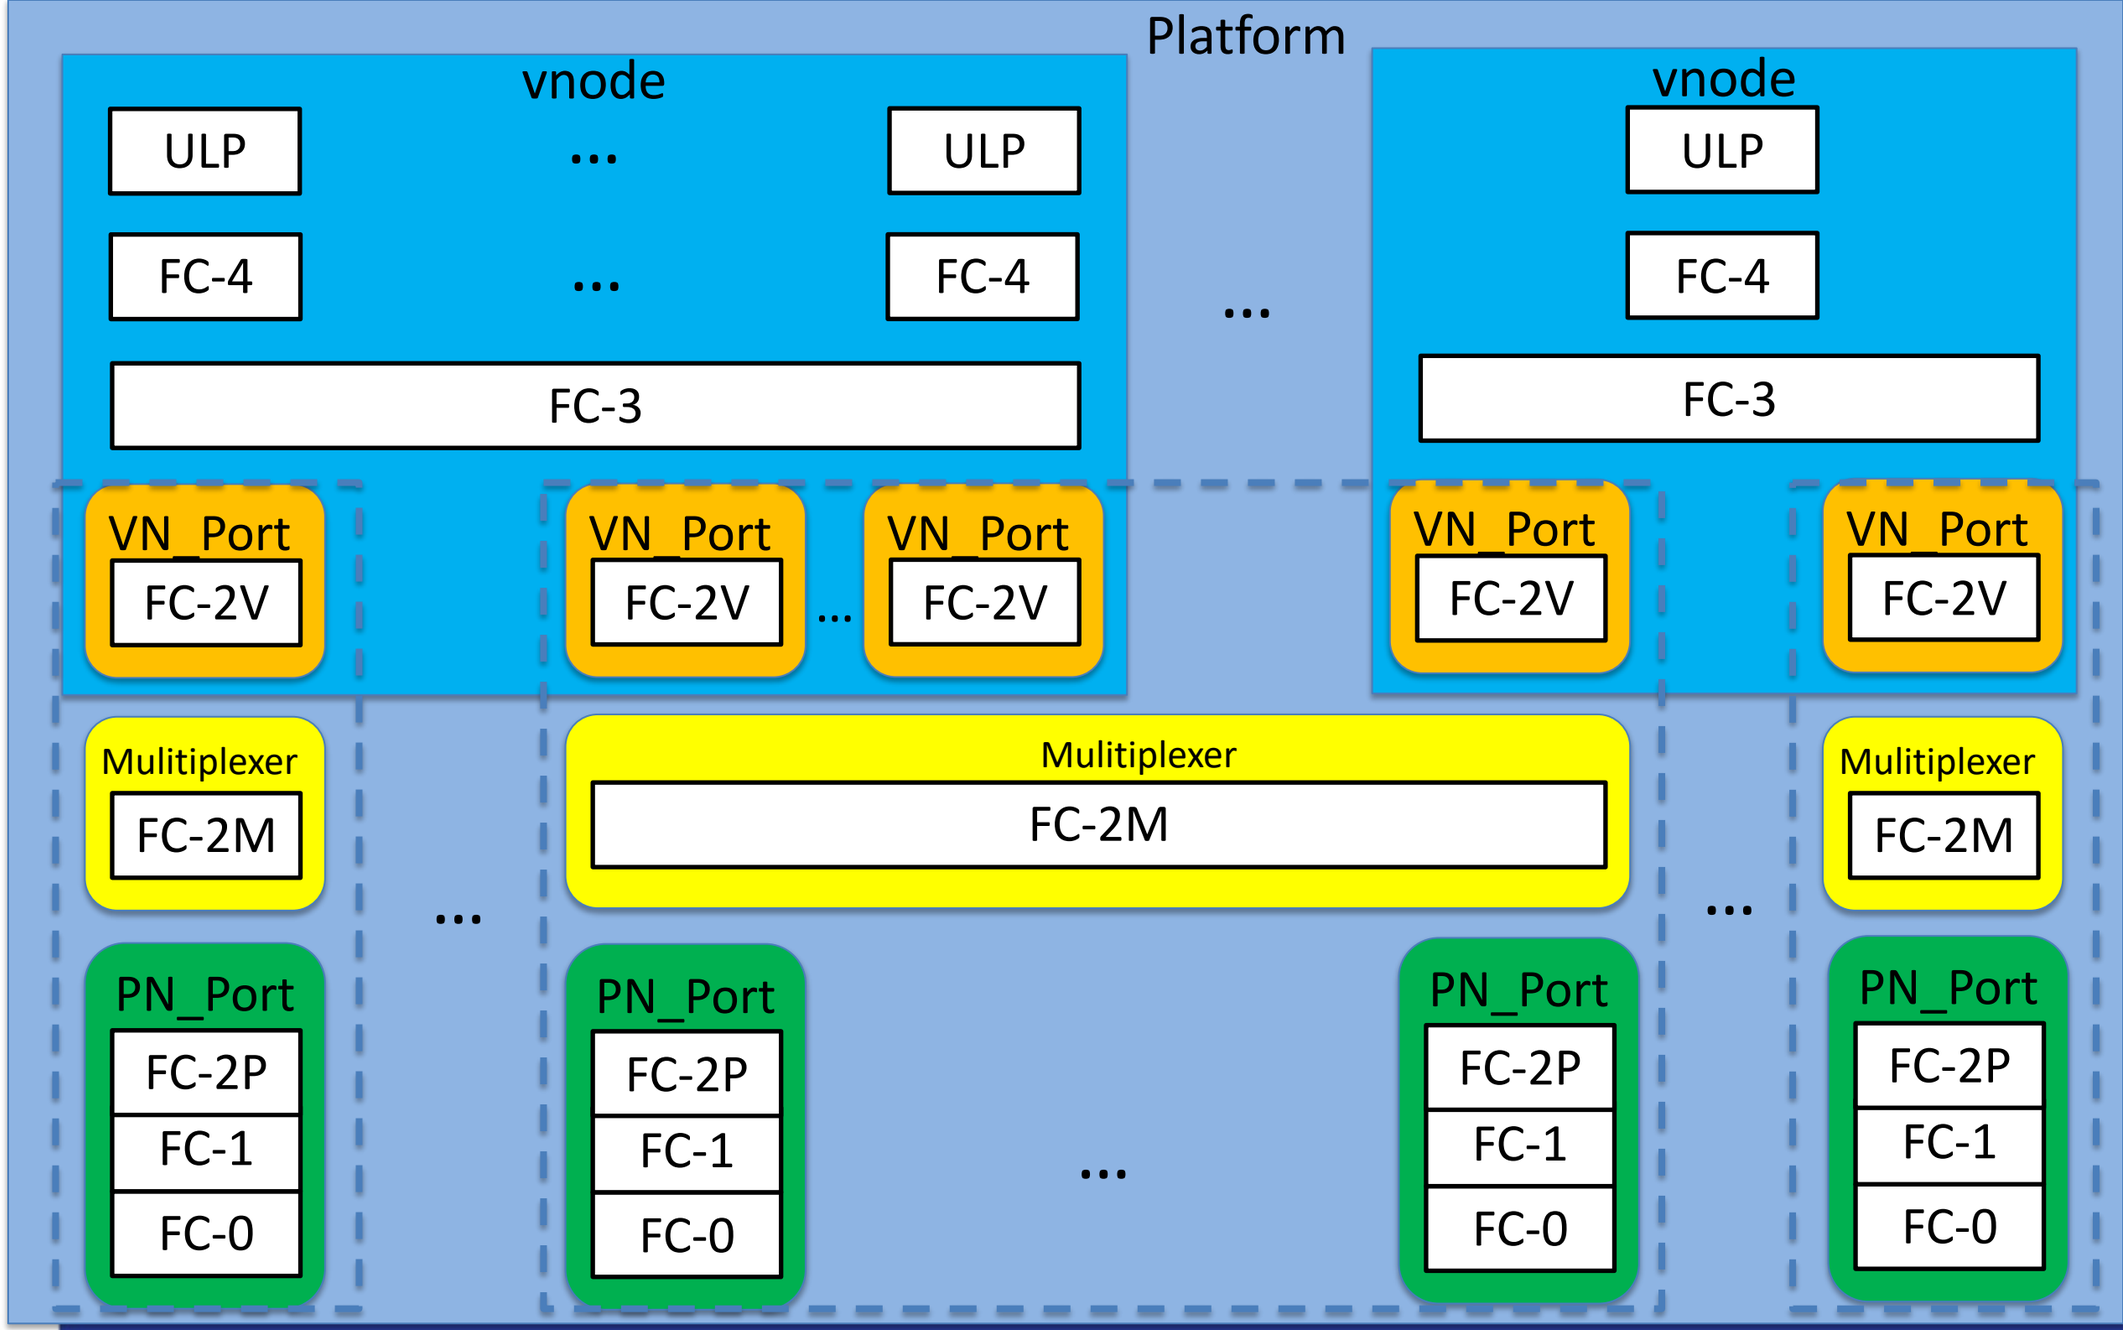

FC Node Architectural Components

- Relations among the architectural components and functional levels in an FC node
  - the term vnode is interchangeable with the term node
  - the term VN_Port is interchangeable with the terms
    - Nx_Port
    - N_Port in Fabric topologies
    - NL_Port in loop topologies

FC Node Architectural Components

- A node is an administratively defined group of ULPs and Nx_Ports within a physical entity (i.e., a Platform)
- Equivalent term vnode replaces the term node to emphasize multiple nodes may coexist within the same Platform
- Each node has a Name_Identifier that enables it to be referenced by certain functions of the FC environment
  - Name Server requests

Physical Model

- FC physically consists of a minimum of two PN_Ports, each associated with a Platform, interconnected by a pair of fibres - one outbound and the other inbound at each PN_Port
- This pair of unidirectional fibres transmitting in opposite directions, with their associated transmitters and receivers, is referred to as a link
- A link is used by the interconnected PN_Ports to perform data communication

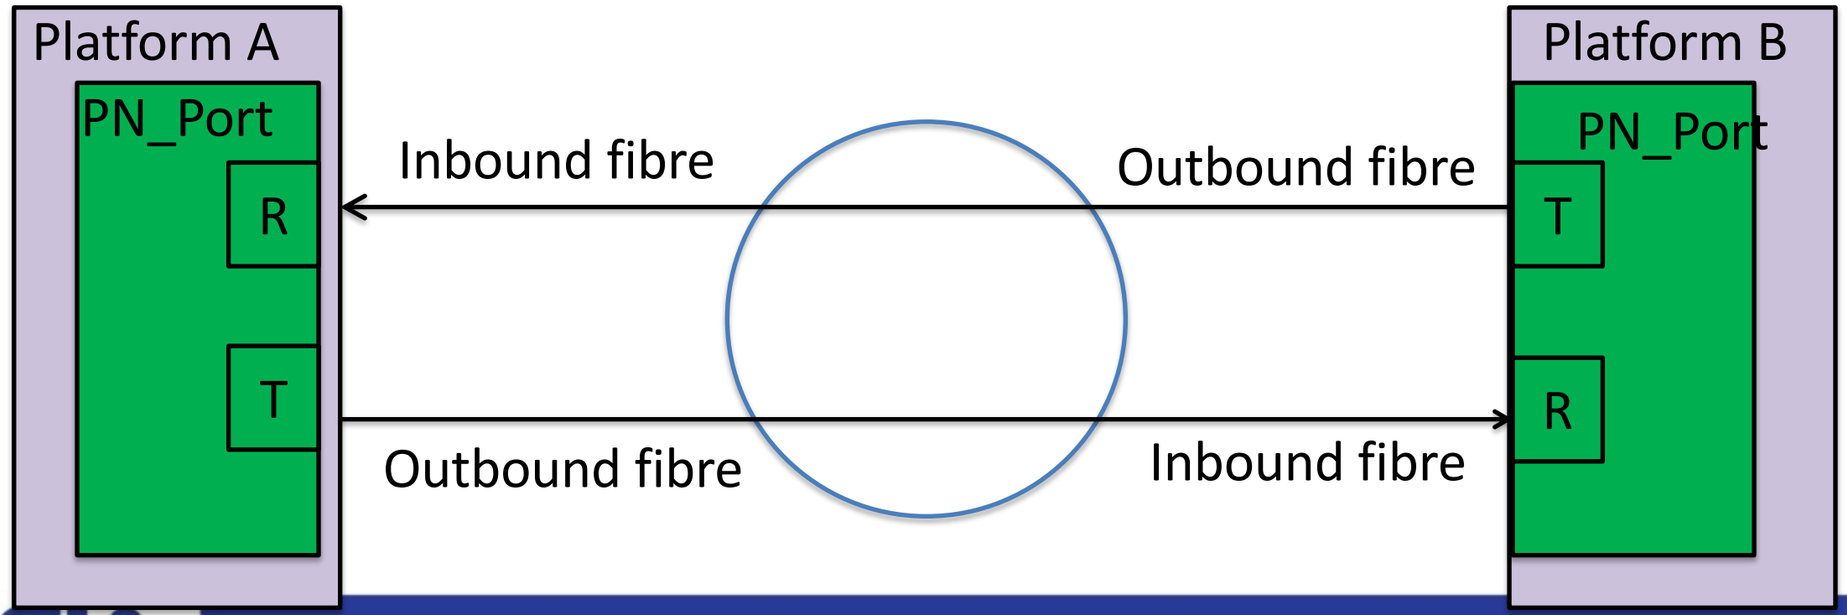

Physical Model

- One or more Platforms can be connected over Fibre Channel links
- Each Platform can contain one or more PN_Ports, each containing a transmitter and a receiver
- Inherently capable of simultaneous bi-directional flow
  - Transmitter sends data frames on outbound fibre
  - Receiver receives responses on inbound fibre

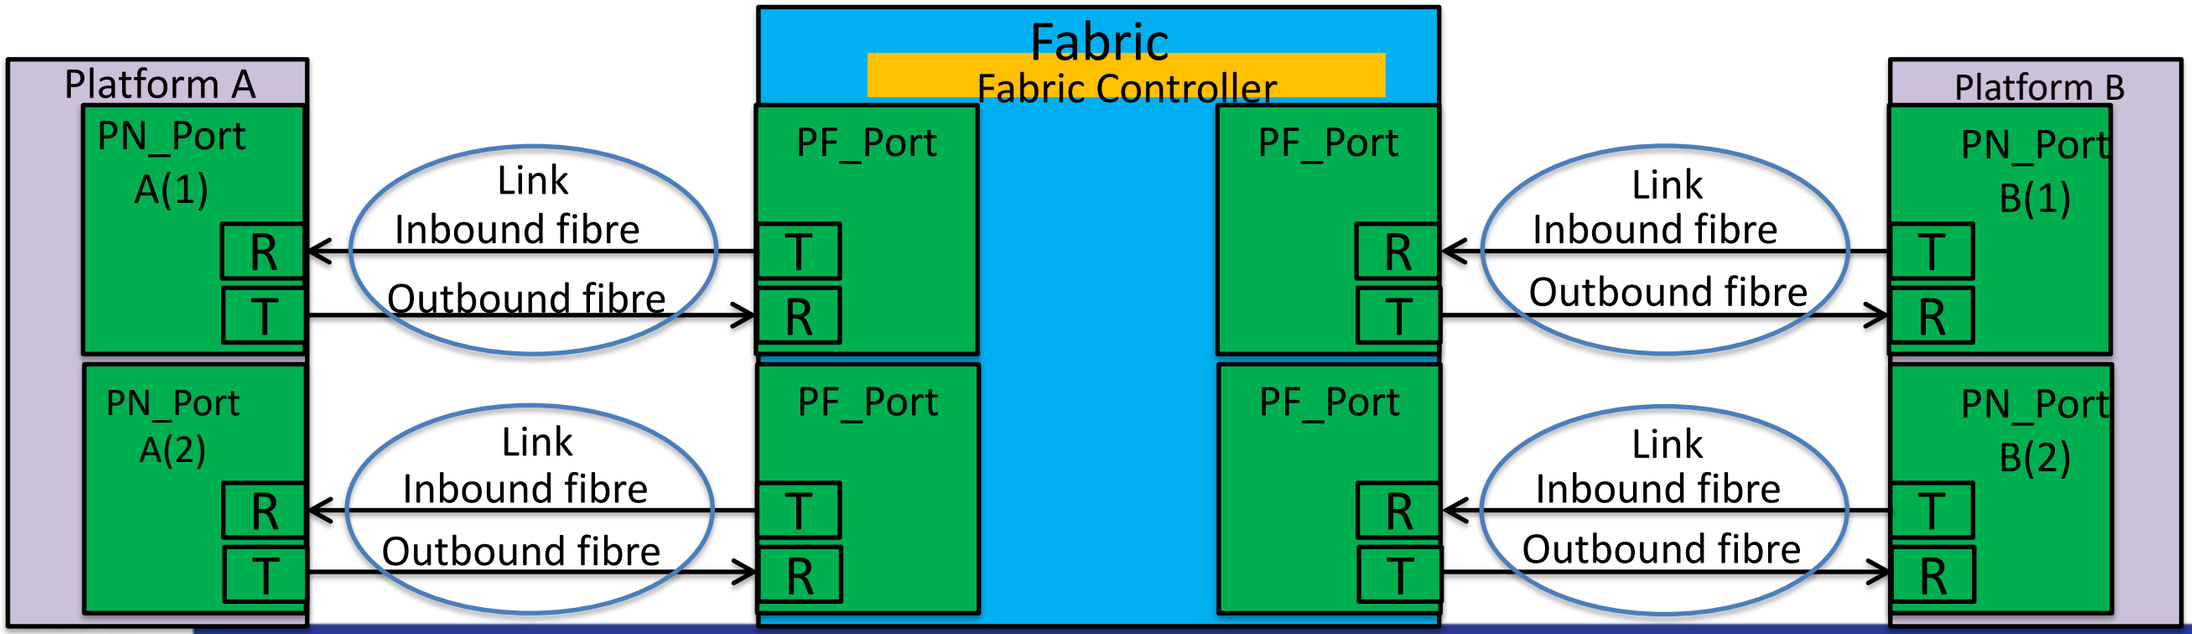

Physical Model

- Link Control facility (LCF) is a hardware facility that attaches to each end of a link and manages transmission and reception of data
  - in a node, an LCF is a PN_Port
  - in a Fabric, an LCF attached to a PN_Port is a PF_Port

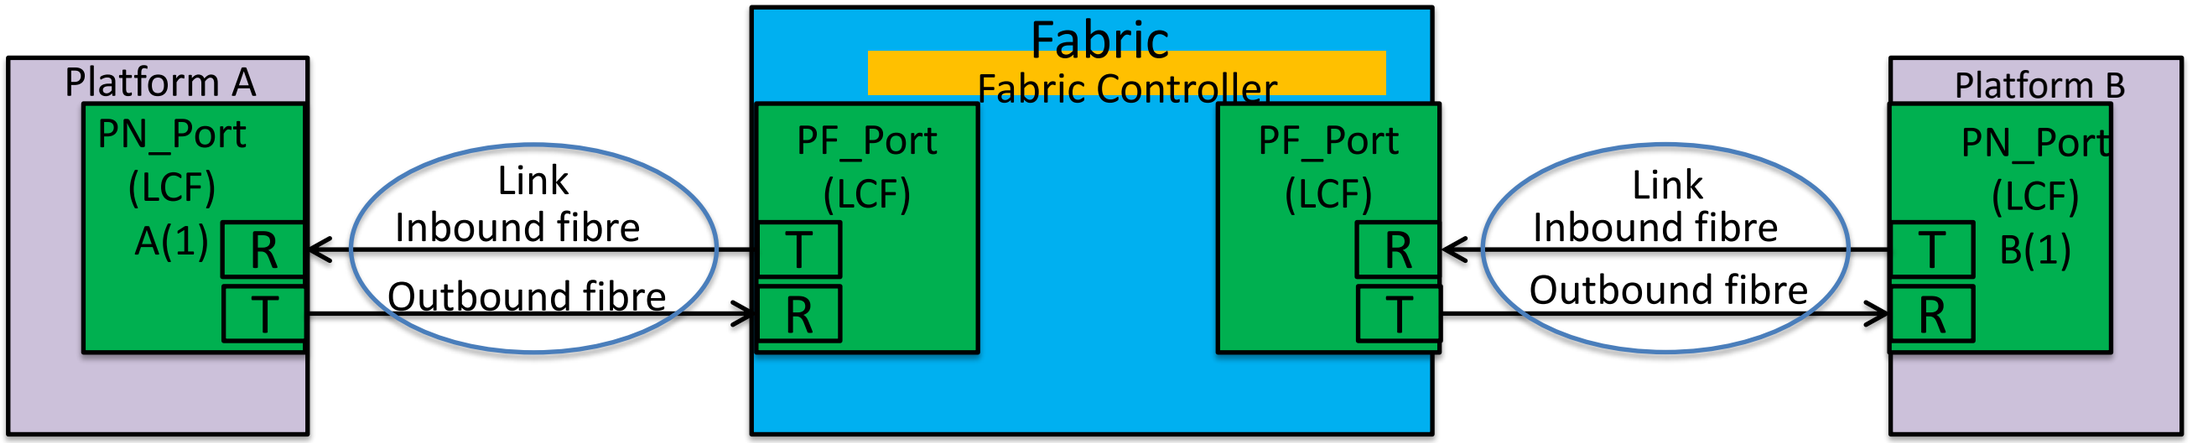

Communication Models

PN_Port:

- transmits Data frames per requests from upper level
- receives Link_Control responses for those Data frames
- receives Data frames from other PN_Ports
- transmits appropriate Link_Control responses for those frames to PN_Ports

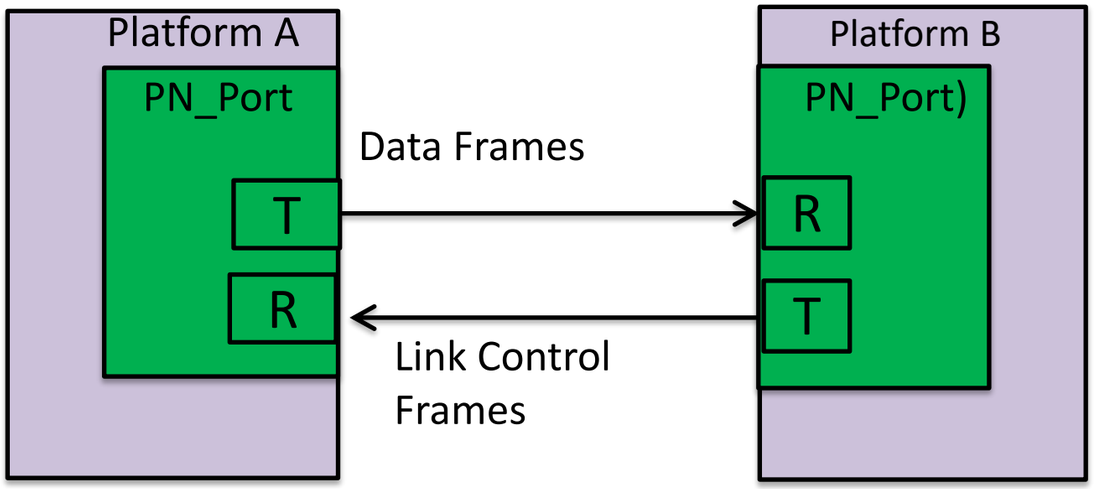

Communication Models

Simplex operation - PN_Port transfers Data frames in one direction only, with Link_Control frames flowing in the opposite direction

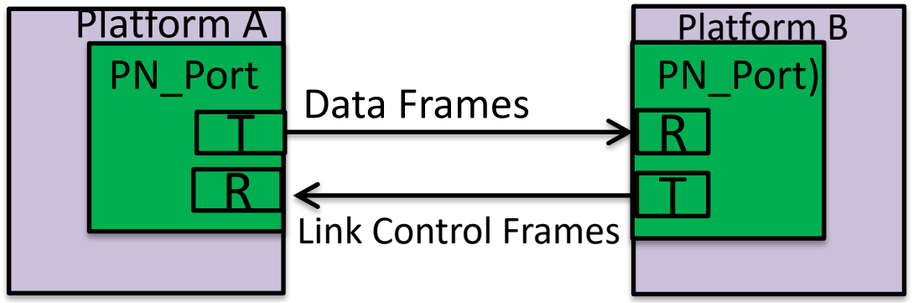

Full-duplex operation - PN_Port simultaneously transmits and receives Data frames, with Link_Control frames flowing in both directions as well

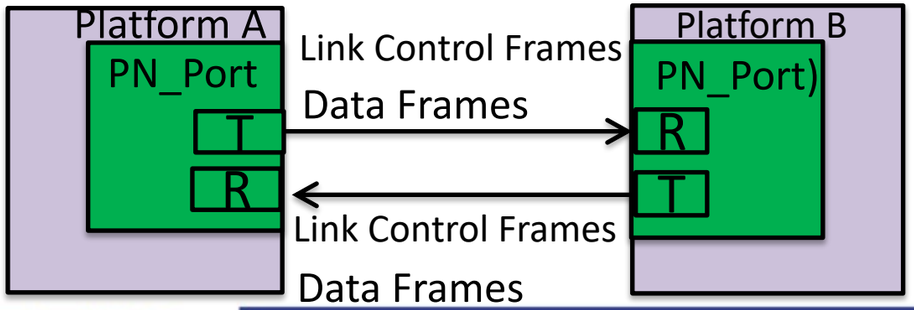

Half-duplex operation - PN_Port both transmitting and receiving data, but not simultaneously

- Data frames and Link_Control frames flow in both directions, but flow is limited to a single direction at a time

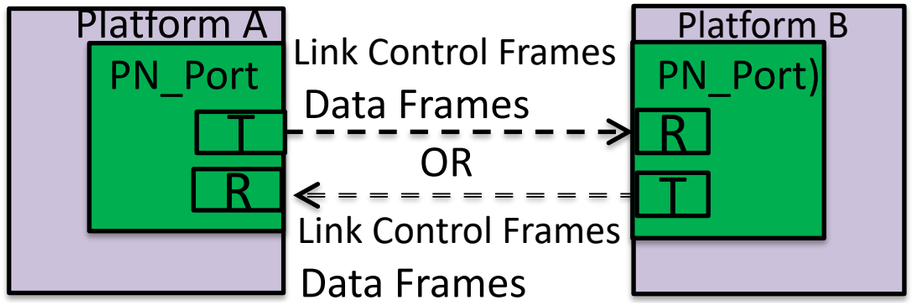

Interconnect Topologies

- Topologies are defined based on the capability and the presence or absence of Fabric between the PN_Ports
- Three basic topology types:
  - Point-to-point
  - Fabric
  - Arbitrated Loop
- The protocols in Fibre Channel are topology independent
  - However, attributes of the topology may restrict operation to certain communication models

Interconnect Topologies

- **Point-to-point** topology allows communication between PN_Ports without the use of a Fabric

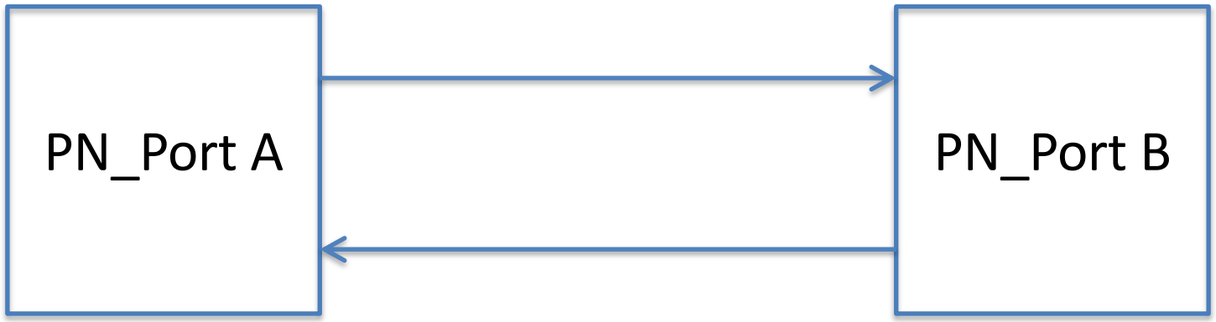

Interconnect Topologies

- **Fabric** topology uses the D_ID embedded in the Frame_Header to route frames through a Fabric to the desired destination PN_Port

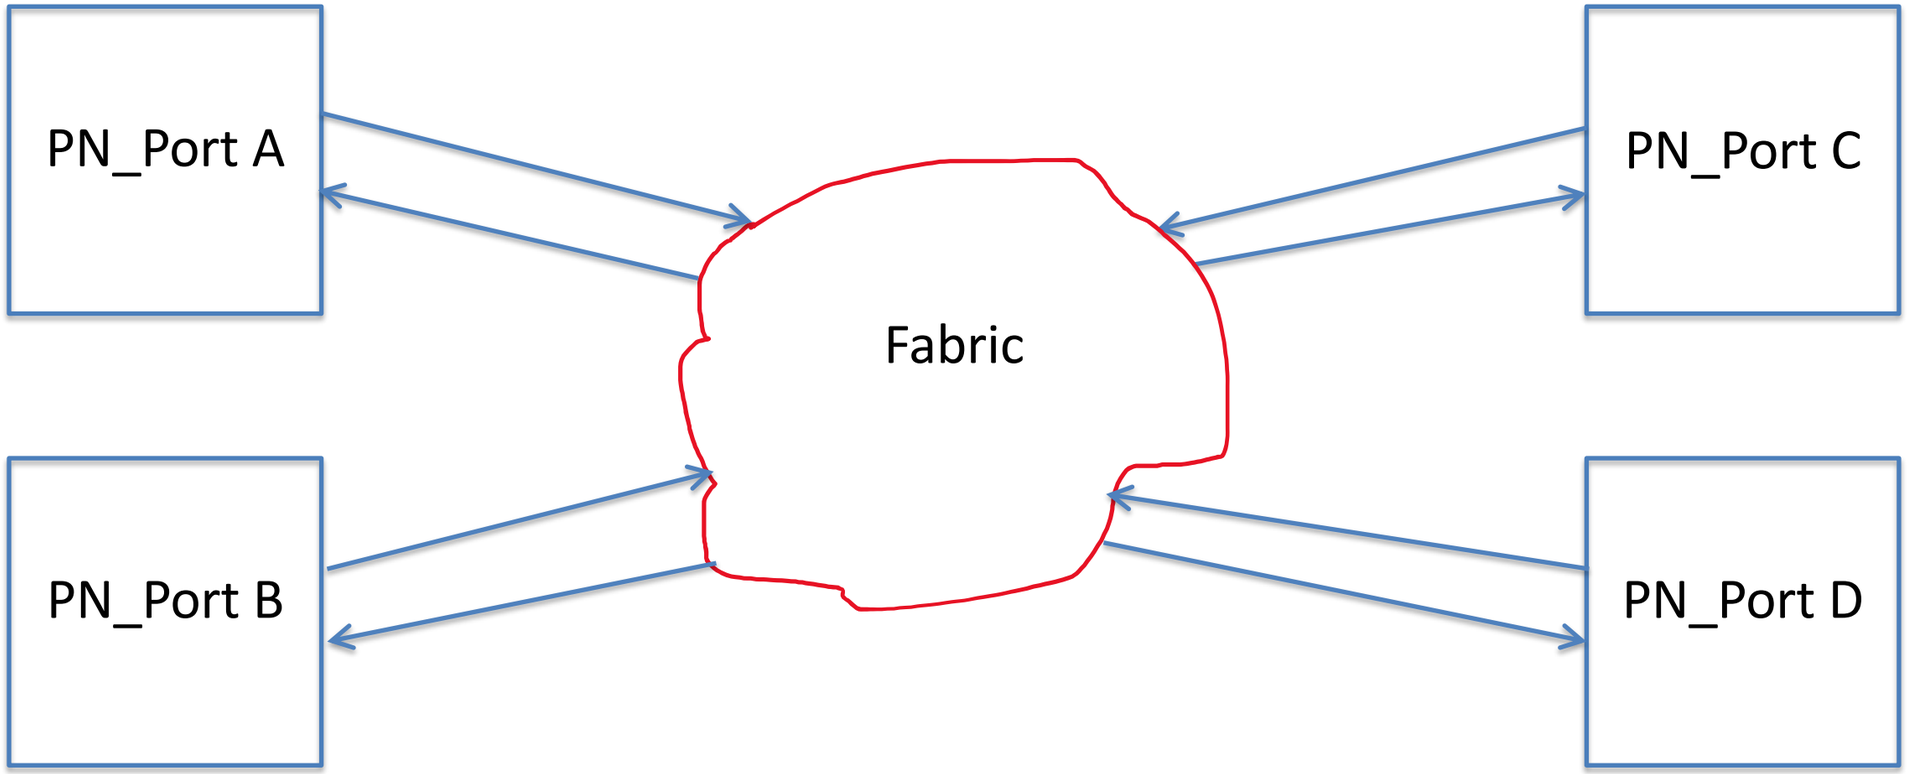

Interconnect Topologies

- **Arbitrated Loop** topology permits three or more L_Ports to communicate without the use of a Fabric
  - supports a maximum of one point-to-point circuit at a time
  - when two L_Ports are communicating, Arbitrated Loop topology supports simultaneous, symmetrical bi-directional flow

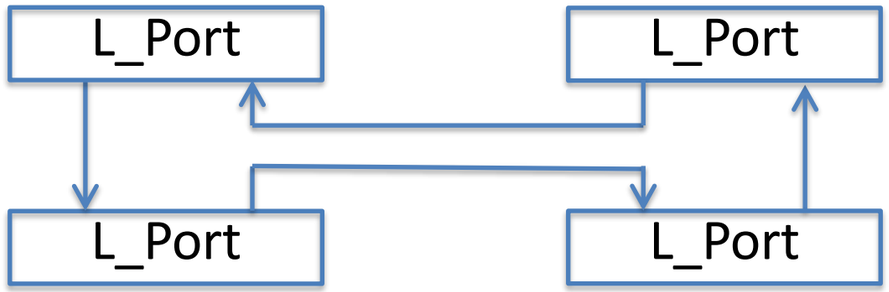

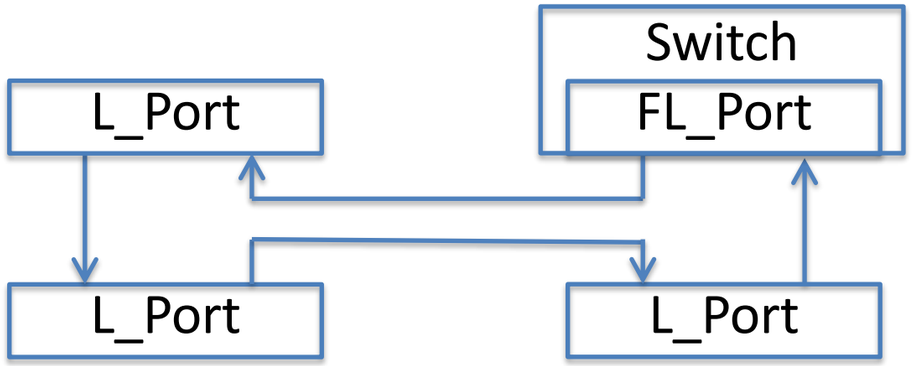

Classes of Service

- Based on the level of delivery integrity required for an application
  - Class 2
  - Class 3
  - Class F
- Topology independent
  - if a Fabric is not present, the class of service is provided as a special case of point-to-point
  - FC_Ports are not required to support all classes of service

Classes of Service

- Class 2 service – multiplex
  - frame delivery service multiplexing frames at frame boundaries ***with*** frame acknowledgement
  - transmitter transmits Class 2 Data frames in a sequential order within a given Sequence
    - however the Fabric may not guarantee the order of delivery and frames may be delivered out of order
  - Fabric or the destination Nx_Port guarantees notification of delivery in the absence of link errors
    - in case of link errors, notification is not guaranteed since the S_ID may not be error free

Classes of Service

- Class 3 service - datagram
  - frame delivery service with the Fabric multiplexing frames at frame boundaries ***without*** frame acknowledgement
  - supports only unacknowledged delivery where the destination Nx_Port does not send any confirmation of Link_Control frames on receipt of valid Data frames
    - any acknowledgement of Class 3 service is up to the upper levels
  - transmitter transmits Class 3 Data frames in sequential order within a given Sequence
    - however, the Fabric may not guarantee the order of delivery and frames may be delivered out of order
  - Fabric is expected to make a best effort to deliver the frame to the intended destination and does not issue a busy or reject frame to the source Nx_Port if unable to deliver the frame

Classes of Service

- Class F service - Fabric
  - frame delivery service used only for communication between switches in a Fabric

Fabric Model

- Primary function of the Fabric is to receive the frames from a source Nx_Port and route the frames to the destination Nx_Port whose address identifier is specified in the frames
- Each Nx_Port is physically attached through a link to the Fabric
- FC-2 specifies the protocol between the Fabric and the attached Nx_Ports
- A Fabric is characterized by a single address space where every Nx_Port has a unique N_Port_ID
- A Fabric specifies the classes of service it supports in its Service Parameters

Fabric Model

- Model is conceptual and may provide the following major functions
  - bi-directional Physical Fabric Ports (PF_Ports)
  - receive buffer
  - frame delivery service
  - receive buffer queue management

Fabric Model

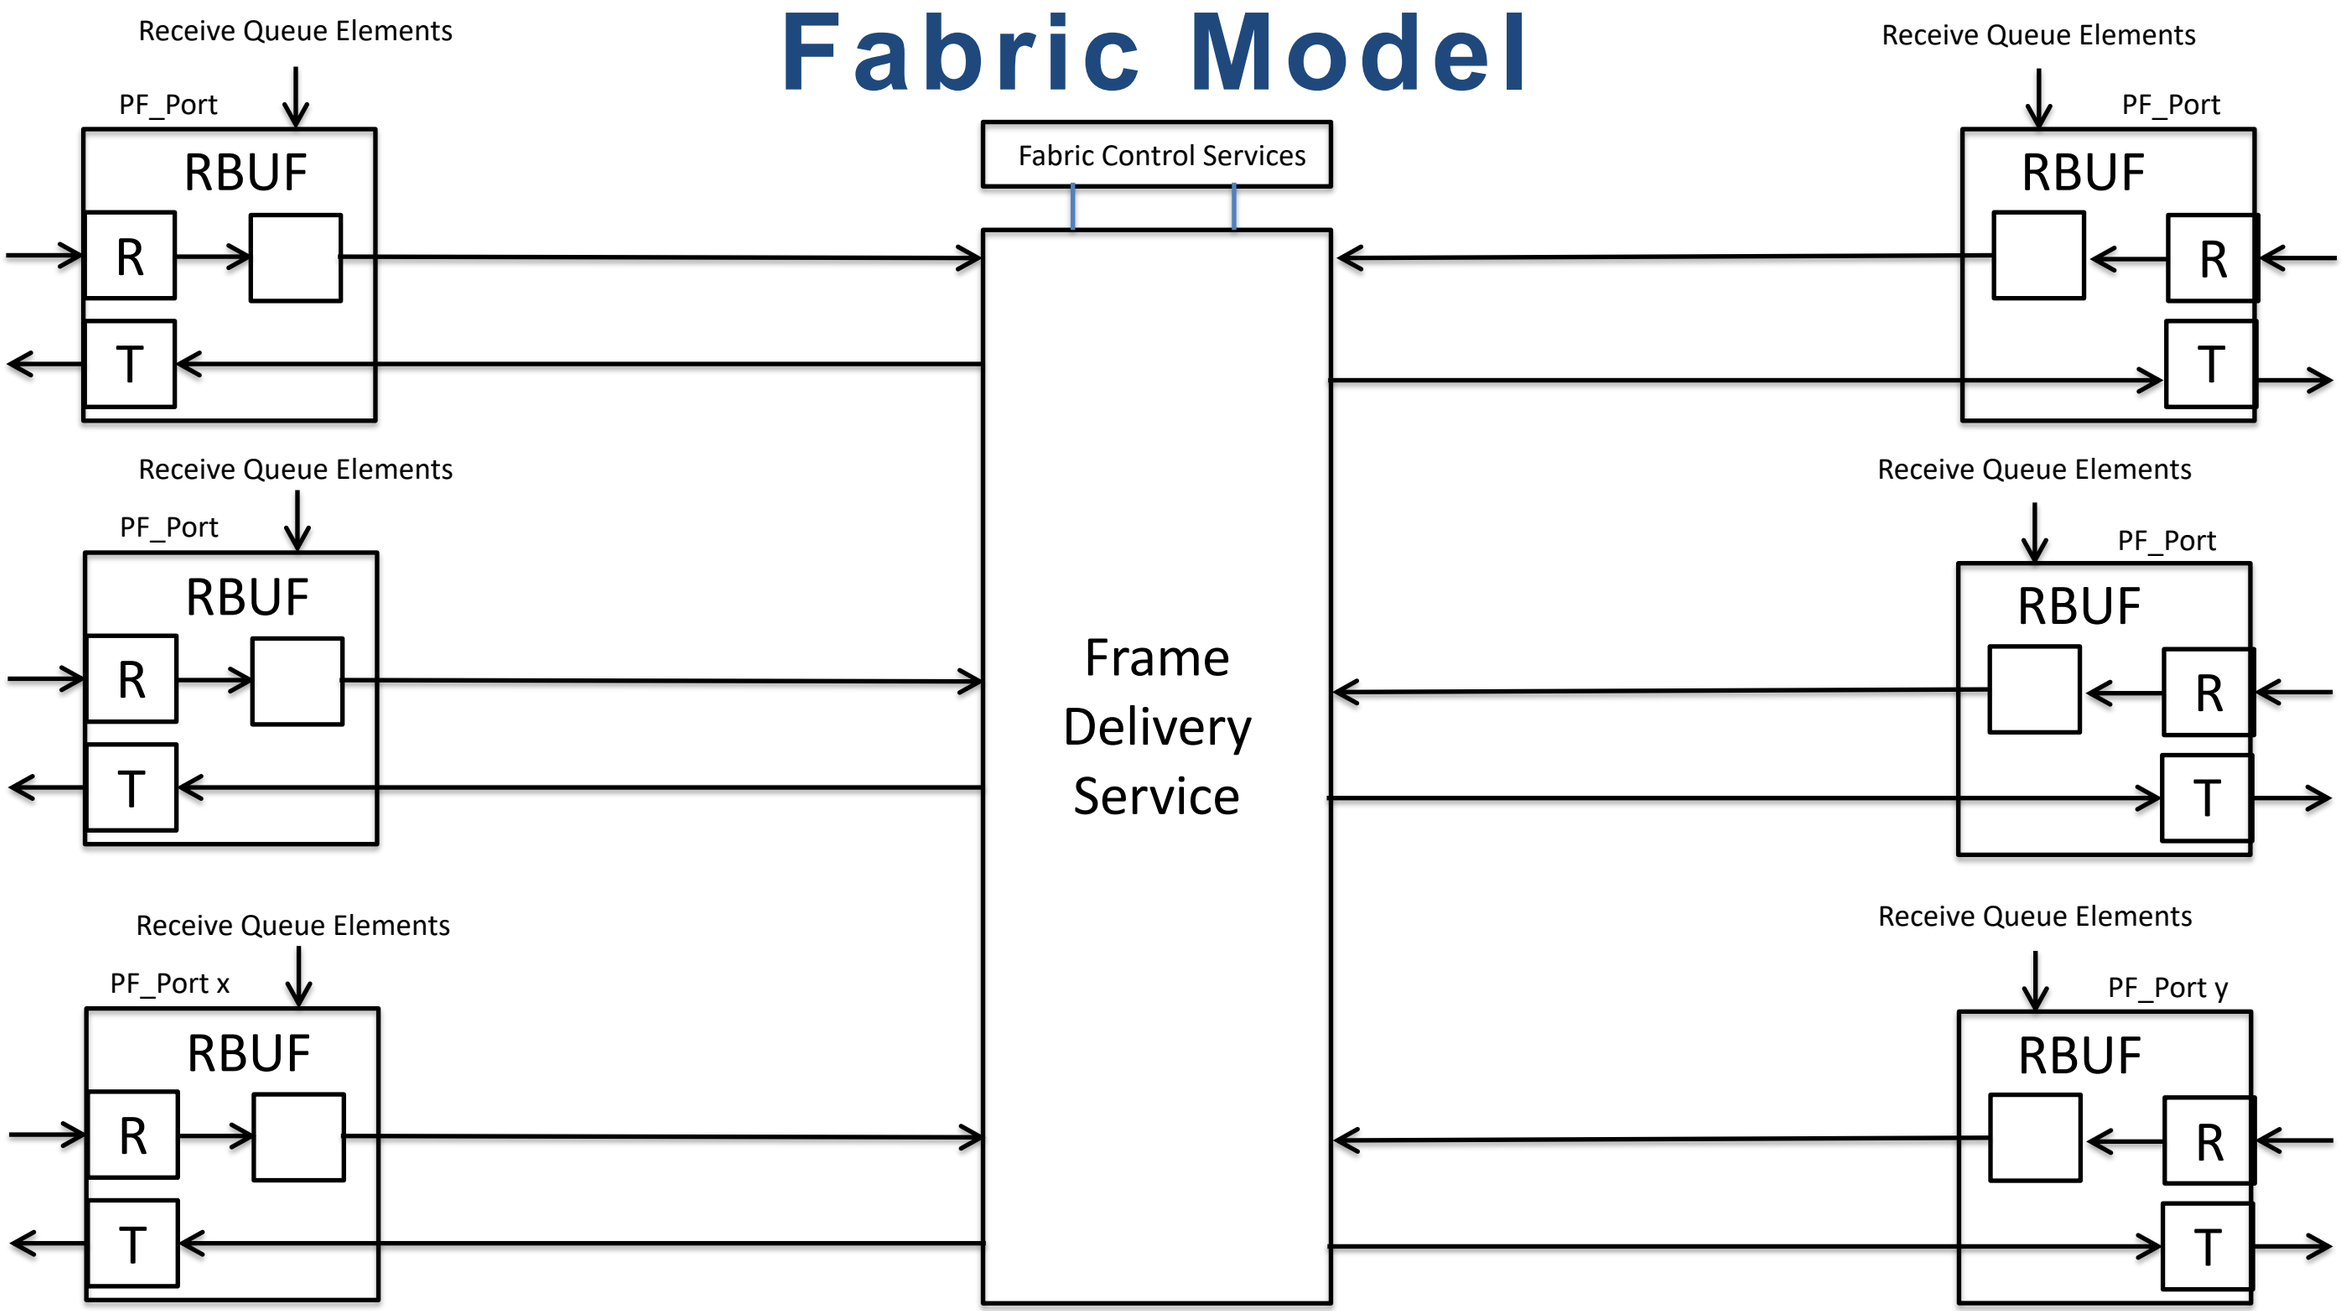

Fabric Model

- Fabric Ports (Fx_Ports)
  - Fabric model contains two or more Fx_Ports
  - Each Fx_Port is attached to one or more Nx_Ports at one or more PN_Ports through a link
  - Each Fx_Port is bi-directional and supports one or more communication models
  - Frames are routed to the Fx_Port attached to the destination Nx_Port
  - Receiving Fx_Port responds to sending Nx_Port according to FC-2 protocol
  - Fabric may verify the validity of the frame as it passes through the Fabric

Fabric Model

- Fabric Ports (Fx_Ports)
  - Fx_Port may contain receive buffers for the incoming frames
  - maximum Data_Field size that the Fabric is able to handle for frames is determined during Login
  - one of the Fabric Service Parameters indicates the maximum Data_Field size for the entire Fabric
  - Fabric routes the frame to the Fx_Port attached to the destination Nx_Port based on the value in the D_ID field embedded in the Frame_Header of the frame
  - routing mechanisms within the Fabric are transparent to Nx_Ports

Fabric Model

- Frame delivery service
  - multiplexes frames at frame boundaries
  - frame delivery service does not guarantee full link bandwidth between communicating Nx_Ports
  - Fabric notifies the transmitting Nx_Port with a reason code embedded in a Link_Response frame, if it is unable to deliver a Class 2 frame
  - for a Class 3 frame, the Fabric does not notify the transmitting Nx_Port if it is unable to deliver the frame

Fabric Model

- Frame delivery service
  - if frames from multiple Nx_Ports are targeted for the same destination Nx_Port in Class 2 or Class 3, congestion of frames may occur within the Fabric
  - management of this congestion is part of the frame delivery service and buffer-to-buffer flow control
  - if any buffer-to-buffer flow control error occurs, the Fabric logs the error and may discard the overflow frame without notification
    - error logging is vendor specific

Generic Services

- Generic Services are provided to meet the needs of the configuration
  - Directory Service
    - Name Server
    - VE Identification Server
  - Management Service
    - Fabric Configuration Server & Enhanced Fabric Configuration Server
    - Unzoned Name Server
    - Fabric Zone Server
    - Security Policy Server
    - Security Information Server
    - Fabric Device Management Interface
    - Application Server
  - Event Service

Generic Services

- Each of these services is addressed with an N_Port_ID for the Nx_Port providing the service or with a well-known address
- Well-known addresses are recognized and routed to by the Fabric
- Services may be centralized or distributed

Our Next FCIA Webcast:

Inside a Modern Fibre Channel Architecture - Part 2

October 27, 2021

Register at: https://bit.ly/FCArchPart2

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
- Our library of FCIA on-demand webcasts http://fibrechannel.org/webcasts/ includes:
  - 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

Thank You