

# FICON MultiHop

## Requirements and Supported Configurations

Dr. Steve Guendert

Brocade Communications

IBM Z Technologies CTO

[Stephen.Guendert@brocade.com](mailto:Stephen.Guendert@brocade.com)

@DrSteveGuendert

# Abstract

- This presentation will discuss the newly announced support by IBM for MultiHop FICON and the implications it has for performance, management, and business continuity architectures. We will go into detail on the Fabric Shortest Path First (FSPF) protocol mechanism, limitations, best practices, and specific configurations supported.

# Agenda

- Quick review of cascaded FICON basics
- Intro to FICON Multihop: why needed and what does it solve?
- FICON Multihop requirements
- Basics of fabric routing and ISL link mapping
- FICON Multihop supported configurations

# Legal Disclaimer

All or some of the products detailed in this presentation may still be under development and certain specifications, including but not limited to, release dates, prices, and product features, may change. The products may not function as intended and a production version of the products may never be released. Even if a production version is released, it may be materially different from the pre-release version discussed in this presentation.

Nothing in this presentation shall be deemed to create a warranty of any kind, either express or implied, statutory or otherwise, including but not limited to, any implied warranties of merchantability, fitness for a particular purpose, or non-infringement of third-party rights with respect to any products and services referenced herein.

Brocade, the B-wing symbol, and MyBrocade are registered trademarks of Brocade Communications Systems, Inc., in the United States and in other countries. Other brands, product names, or service names mentioned of Brocade Communications Systems, Inc. are listed at [www.brocade.com/en/legal/brocade-Legal-intellectual-property/brocade-legal-trademarks.html](http://www.brocade.com/en/legal/brocade-Legal-intellectual-property/brocade-legal-trademarks.html). Other marks may belong to third parties.

# Cascaded FICON Basics

A quick review

# FICON history review: 19 years of evolution

FC-SB-2, 3, 4, 5

- IBM originally announced FICON in May 1998 (9672 G5)
- FC-SB-2 did not support switched paths over multiple directors
- FC-SB-3 (January 2003) introduced support for cascaded FICON.
- Now on 9<sup>th</sup> generation of the FICON channel (FICON Express 16S+)
- See prior Computer Measurement Group (CMG) Proceedings for more on this history


# Cascaded FICON

- An implementation of FICON that allows two FICON SAN fabrics to be linked via connections between pairs of FICON Directors.
- Director to Director connections are known as interswitch links (ISLs)
- Permits greater flexibility, better utilization of links, and higher levels of availability in the enterprise.
- Facilitate designs for more robust business continuity/disaster recovery (BC/DR) architectures such as IBM GDPS.



# Non-Cascaded FICON Environment

## Geographically Dispersed Data Center Connections

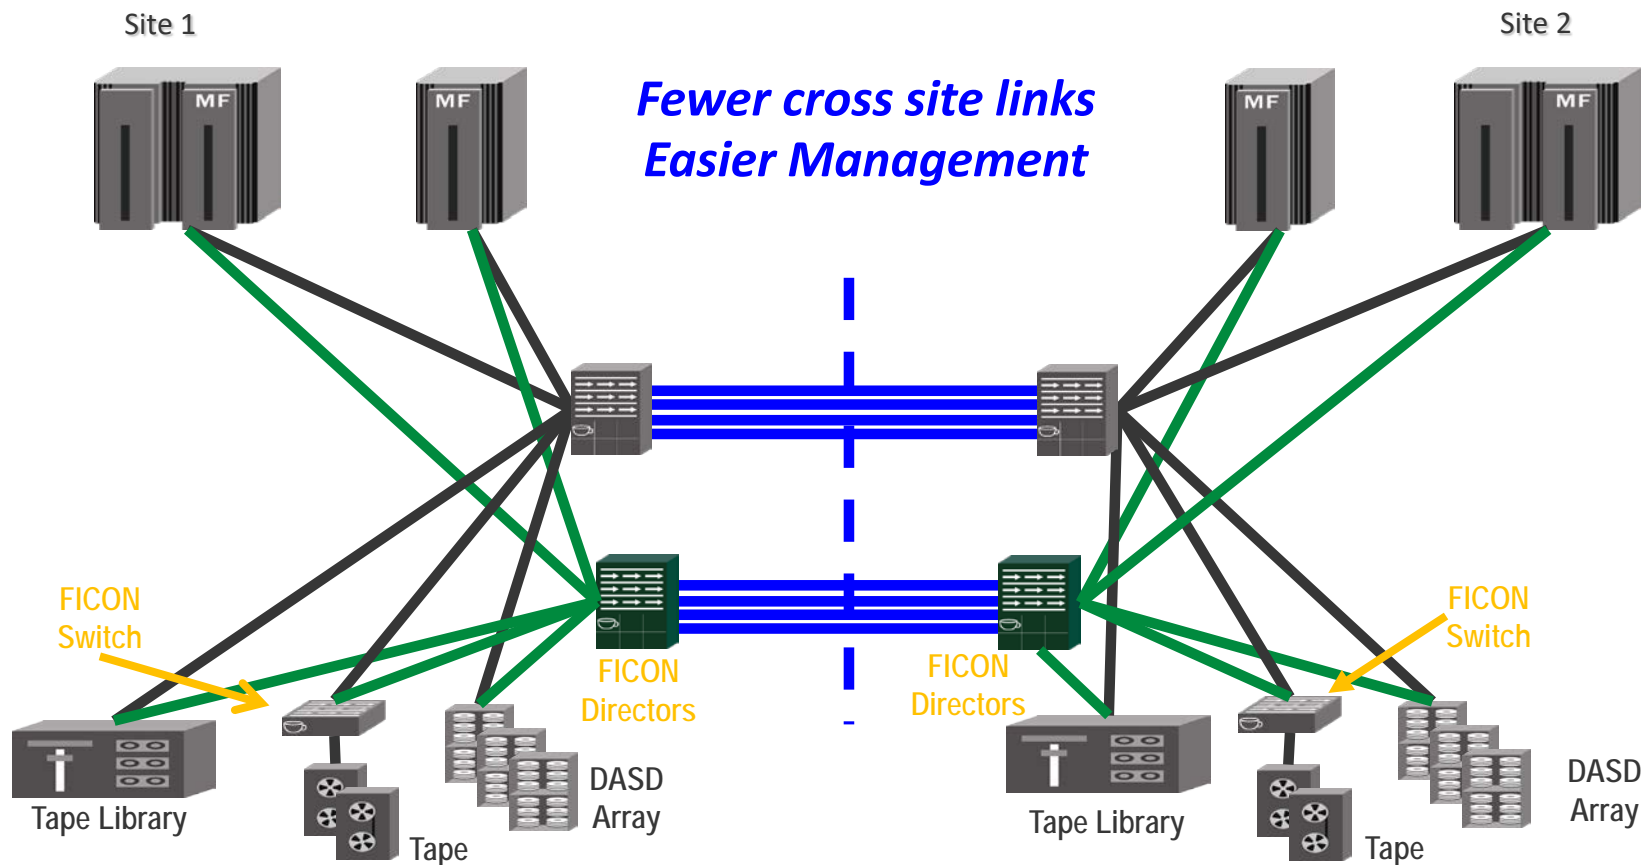






# Cascaded FICON Environment

## Geographically Dispersed Data Center Connections



# ISL Definition and Addressing

- HCD defines the relationship between channels and directors (via Switch ID).
- HCD also defines the relationship for the destination switch ID and specific switch ports to control units (CU) for the switch to CU side of the cascaded FICON link.
- **HCD does not define the ISL connections.**
- Management of the traffic over ISLs is controlled exclusively by the FICON Directors.

# ISL Definition and Addressing (2)

- Additional ISL bandwidth can be added to a cascaded FICON configuration without modifying existing HCD definitions.
- HCD assumes the links are present and requires 2 byte addressing be employed.
- During initialization the FICON switching devices identify their peers and create internal routing tables.

# Introduction to FICON Multihop

Why needed and what problem does it solve?

What is Required?

# Topology Expansion

## Multi-hop Support




- Challenge

- Fabric resiliency under failure is affected when “unsupported” routes must be avoided to meet support restrictions
- FICON fabrics are viewed as less reliable than FCP fabrics

- Solution

- Remove restrictions regarding hop counts between channel and control unit
- Verify FICON architecture over broader scope of configurations
- Establish best practice guidelines for FICON configurations



# Multi-Hop FICON

- FOS 8.1 qualification
- Support for up to 4 cascaded switches/ 3 hops
- Initially static routing only
  - Cannot be used in conjunction with FIDR
- IBM Technical White Paper co-authored by Brocade

# Hardware Requirements for FICON Multihop

IBM Z hardware requirements

- Z13/z13s/z14 running Driver 27L with bundle S08a or later
- FICON Express 16S+, FICON Express 16S or FICON Express 8S channels

# Hardware Requirements for FICON Multihop

## DASD/Storage hardware requirements

- DS8870/DS8880 running firmware release 7.5 or later
- 16Gbps or 8Gbps host adapter cards
- Other storage OEMs: check with your vendor to verify requirements



# SAN Hardware Requirements for FICON Multihop

## Hardware requirements

- All Brocade IBM Z qualified Gen 5 (16 Gbps) products
- All Brocade IBM Z qualified Gen 6 (32 Gbps) products
- FOS 7.4.0a and later
- Brocade Network Advisor 12.4.1 or later

# Network/DWDM and software requirements

- In general: devices qualified/supported in an IBM Z environment are supported for Multihop FICON
- Check [Resource Link](#)
- Still verify support between DWDM/network devices with FICON SAN hardware components
  - Verify specific support with ISL features
- FICON Multihop is supported on all operating systems that run on z13 and z14 that allow FICON channels to be defined.

# Basics of Fabric Routing and ISL Link Mapping

# Routing overview and ISL link mapping

- Internal routing tables
- Number of hops and differences between open systems SAN and FICON (the one hop rule).
- Metric (cost) based on number of hops, and bandwidth of ISL
- Best practice: same bandwidth for ISLs

# Fabric Shortest Path First (FSPF)

- Foundation for all fibre channel SAN routing mechanisms.
- FSPF is a link state path selection protocol that directs traffic along the shortest path between the source and destination based upon the link cost.
- FSPF detects link failures, determines the shortest route for traffic, updates the routing table, provides for fixed routing paths through a fabric, and maintains correct ordering of frames.
- FSPF is based on a replicated topology database present in every switching device in the FICON SAN.
- A Link State Record (LSR) describes the connectivity of a switch within the topology database.

# Link State Record (LSR)

- Describes the connectivity of a switch within the topology database
  - The topology database contains one LSR for each switch in the FICON SAN fabric.
- Each LSR consists of a link state record header, and one or more link descriptors.
- Each link descriptor describes an ISL associated with that switch
  - identifies an ISL by the Domain\_ID and output port index of the “owning” switch and the Domain\_ID and input port index of the “neighbor” switch.
  - This combination uniquely identifies an ISL within the fabric.
- LSRs are transmitted during fabric configuration to synchronize the topology databases in the attached switches.
- LSRs are also transmitted when the state of a link changes and on a periodic basis to refresh the topology database.

# Link Cost

- Associated with each ISL
- Reflects the desirability of routing frames via that ISL.
- The link cost is inversely proportional to the speed of the link
  - higher speed links are more desirable transit paths, therefore they have a lower cost.
- The topology database has entries for all ISLs in the fabric, enabling a switch to compute its least cost path to every other switching device in the FICON SAN fabric from the information contained in its copy of the database.

# FSPF's Three Main Functions

- Hello protocol
  - Establishes two-way communication with neighboring switch
  - Determines the connectivity of an ISL
- Initial topology database synchronization protocol
- Topology database maintenance protocol

The White Paper covers these in great detail



# Hello Protocol summary

- Used to determine when a neighbor switch is ready to begin routing frames on an ISL.
- Initially, a FICON switching device does not know the Domain\_ID of its neighbor(s) and the switch is then said to have “one-way” communications with the neighbor.
- Once a switch acquires a Domain\_ID, it begins transmitting periodic Hello messages, and indicates the source of the ISL by using the Domain\_ID and output port index of the sending switching device.
- When a FICON switching device receives a Hello on an ISL, it learns the Domain\_ID of the neighbor, and it uses its output port on this ISL to send its own Hello containing its own Domain\_ID and port index.
  - The switch is then aware of the neighbor switch’s Domain\_ID
  - At this point, the two are said to have “two-way” communications with each other and they can begin their initial topology database synchronization.

# Hello Protocol summary (2)


- Hello messages are transmitted on a periodic basis on each ISL even after two-way communication is established
  - In essence, the periodic Hello acts as a “heartbeat” between the switching devices
- The time interval between successive Hellos (in seconds) is defined by a timer called the Hello interval
  - The number of seconds that a switching device will wait for a Hello before reverting back to one-way communication is defined by the Dead interval.
  - Both the Hello and Dead intervals are communicated between switches as part of the Hello itself.
- If a switch fails to receive a Hello within the expected time it
  - assumes the neighbor switch is no longer operational and
  - removes the associated ISL from its topology database
  - the switch removes the Domain\_ID of the neighbor switch from the Hello messages and the Hello protocol reverts to the one-way communication state.
- If a link failure occurs on an ISL, two-way communication is lost. When the ISL is restored, the switching devices must re-establish two-way communication and synchronize their topology databases before the ISL may be again used for routing frames.

# FICON Multihop Supported Configurations

Native ISL Configurations


# Two-Hop Array

- Three switches connected linearly
- Must adhere to the FICON timeout limitations
  - Longest distance a FICON can traverse is 300km
  - the maximum distance of  $AB + BC \leq 300\text{km}$




# Triangle

- Three switches connected in a triangle
- Must observe the 300km max distance for FICON links.
  - Examine the worst case scenarios in the event of path loss



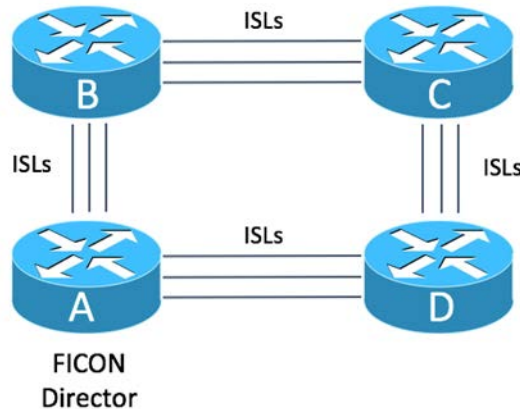
# Three-Hop, Four Array Configuration

- Three switches connected linearly
- 300 km distance limit still applies



# The Square

- Four switches connected in a square configuration
- 300 km distance limit still applies
  - Examine worst case scenario in event of path failure



# FICON Multihop Supported Configurations

FCIP Configurations




# FCIP configurations

- Primarily used for asynchronous replication applications
- Maximum distance supported over an FCIP connection for FICON=3000km

# Two-Hop array with FCIP Bridge


- Two switches connected with native ISLs (300km)
- Third switch connected over FCIP (3000km)



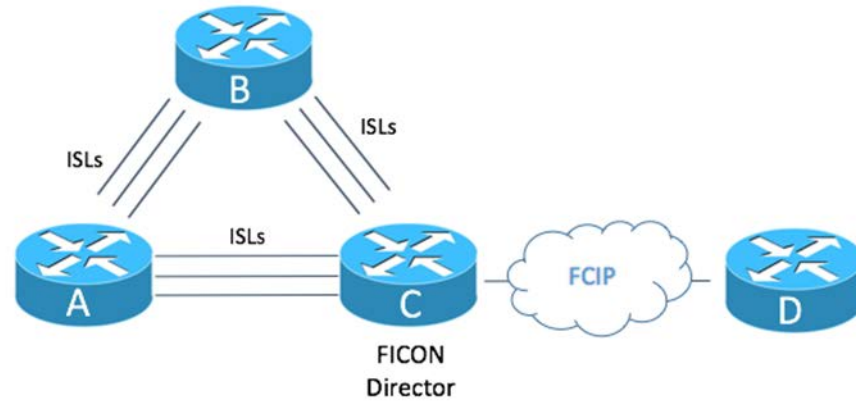
# Three Hop Four Switch with middle FCIP



# Three-Hop array with Leaf FCIP Bridge



# Triangle with Leaf FCIP Bridge



# FICON Multihop Design Considerations

# Multihop Design Considerations

- Bandwidth planning
  - Consider having additional bandwidth available between all routes
  - Perform a bandwidth sizing study
- Potential adverse performance impact to applications in the event or rerouting traffic due to failure/loss of connectivity
  - Increased latency traversing a longer path

# Bandwidth Sizing Study Considerations

- What type of traffic is going across the cascaded FICON links (ISLs)? DASD, tape, CTC, all of the above?
- Do I have replication traffic going across the ISLs? If so, what type (synchronous, asynchronous, both)?
- Do I wish to isolate a particular traffic type to its own set of ISLs (by OS, storage type, replication type)?
- Am I using trunking/port channels in conjunction with Multi-Hop ?
- Am I using virtual fabrics and how does that effect my ISL allocation?
- Do I have an SLA that must be met on a specific replication traffic?
- Does my environment include a GDPS or similar architecture? Will I be performing a HyperSwap?



# Summary

- FICON Multihop provides
  - Potential consolidation and simpler management
  - Better availability for 3 and 4 site architectures
- FICON Multihop does introduce more complexity into configurations
- Increased complexity means more detailed planning is required