

FCIP (Extension): Data Protection and Business Continuity

Live Webcast
December 7, 2017
10:00 AM PT

Today's Presenters



Mark Detrick
Principal Solutions Architect
Brocade Storage Networking
Broadcom



Rupin Mohan
Director R&D, Storage Networking
HPE


Agenda

FCIP (Extension): Data Protection and Business Continuity

- Why do replication?
- Why FCIP?
- What is FCIP?
- RPO/RTO and Sync/Async
- FCIP Design Considerations
- FCIP Deployments
- Summary and Q&A

Who Needs Data Protection?


- Who is immune to this?
- Whose business won't suffer severe consequences if massive data loss occurs?
- Who does not have to answer to shareholders, customers and/or regulatory demands?



Lightning & Fires



Hurricanes / Cyclones




Earthquakes




Tornadoes / Heavy Storms



Tsunami




Data Ransom and Breaches



Inaccessible Infrastructure

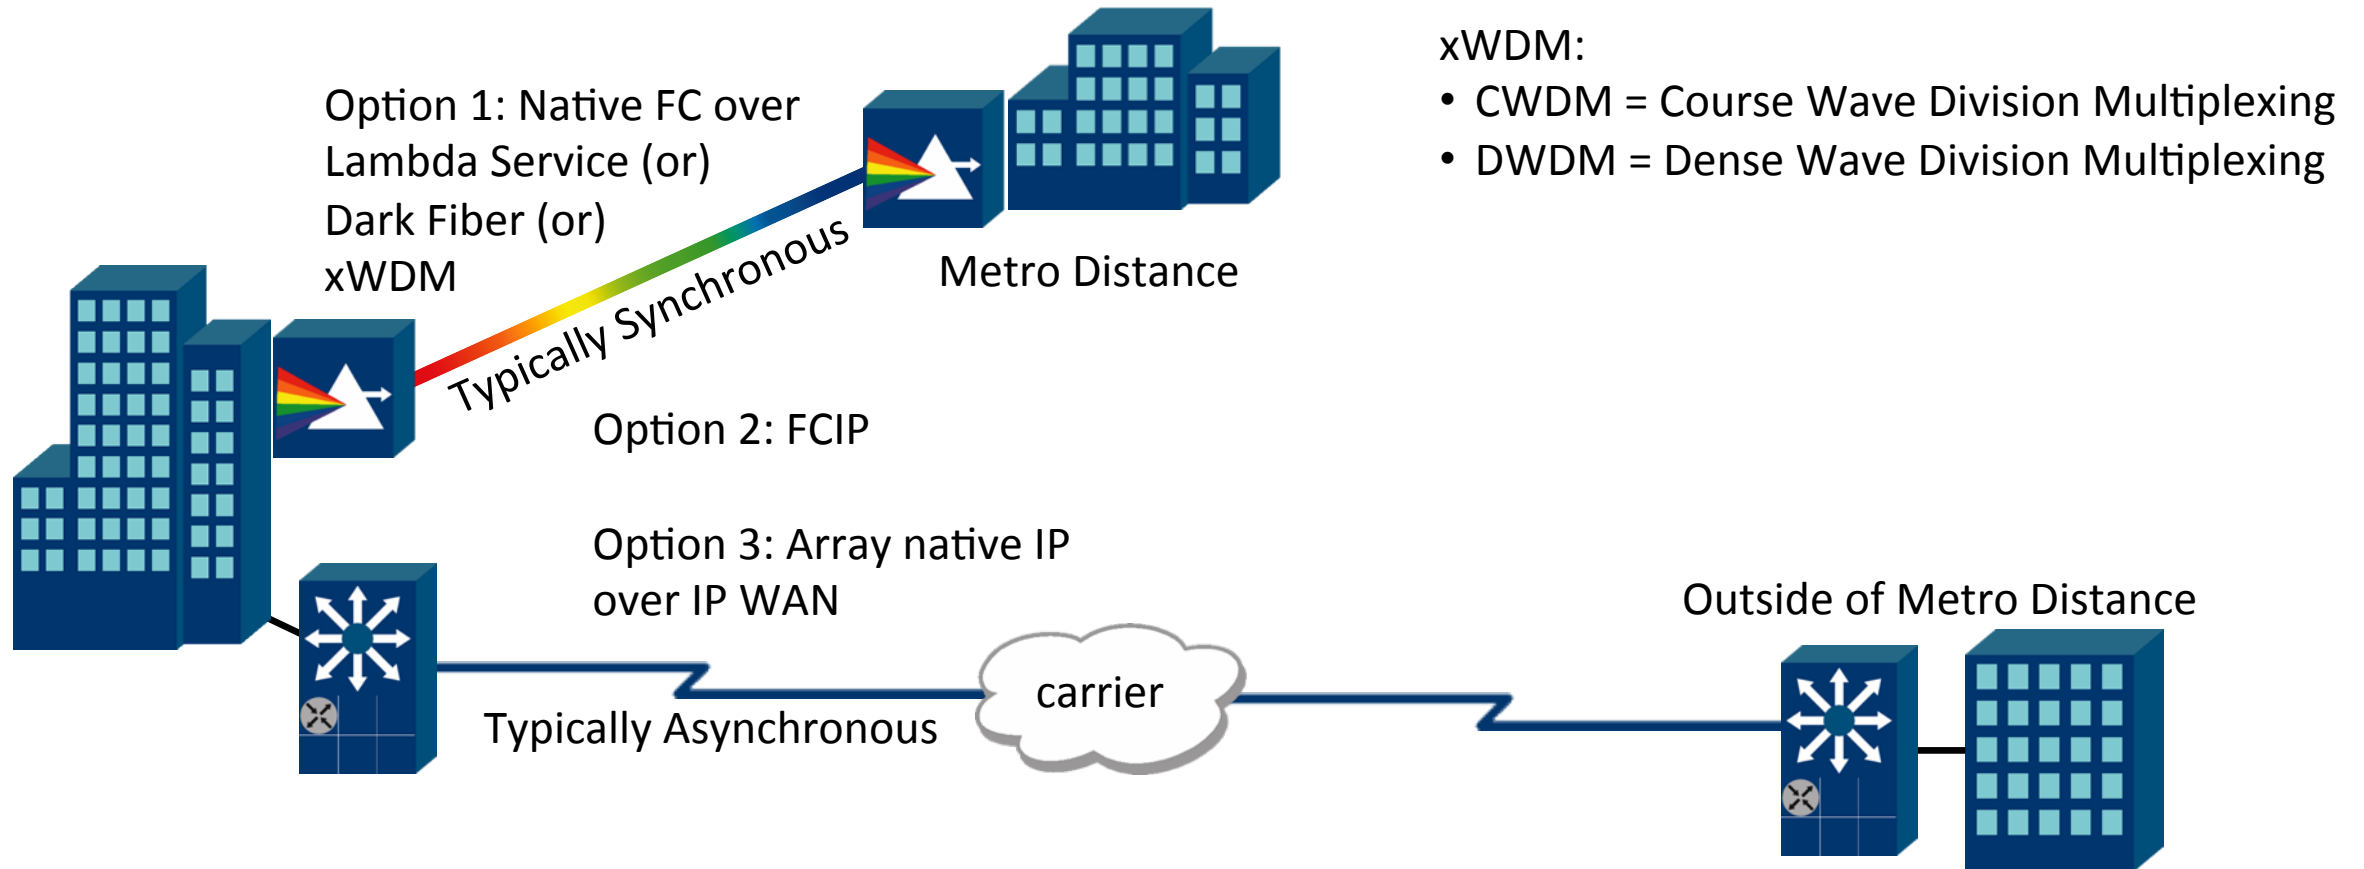


Cut Cables and Power



Terrorism

Long Distance Options



Solution Comparison

Capability	FCIP Async over IP WAN	Array Native IP Async / Sync	Native FC Sync over DWDM
Latency / Distance	1 to 200 ms RTT (100 km to global)	Async: 1 RTT to 150 ms global Sync: <1 ms (100 km or less)	<1 -3 ms RTT (100 km or less)
Cost	Middle	Lowest	Highest
High Availability	<ul style="list-style-type: none"> • Load balancing can be per exchange or frame • Robust and efficient • Frame based routing can be lossless failover 	<ul style="list-style-type: none"> • Load balancing across multiple paths • TCP is neither robust nor efficient 	<ul style="list-style-type: none"> • Load balancing can be per exchange or frame • Robust and efficient • Frame based routing can be lossless failover and failback
Compression	FCIP optimized compression	Typically, array can do compression	Some FC switches provide compression
Encryption	FCIP encryption	Require external device	FC switches provide encryption
Throughput and bandwidth efficiency	High throughput dependent on bandwidth of IP WAN – 10 Gbps IP WAN is common.	Array IP ports have limited throughput relative to FC ports	High throughput 10 Gbps interfaces/optics are common

FCIP: You can go longest distance with high throughput without any loss during failover


FC over IP (FCIP)

Traversing IP Links

- **FCIP is a TCP/IP tunneling protocol**
 - Connects distance separated FC devices via an IP WAN
 - Live mirroring (RDR = Remote Data Replication)
 - Tape backup (Open Systems and FICON tape)
- An FCIP link is a FC ISL transporting frames between domains
 - A VE_Port is a virtual E_Port facing an FCIP tunnel
 - VE_Port to VE_Port is an ISL (Inter-Switch Link) though FCIP
 - Control data (class F) passes across FCIP
- Tunneling
 - The IP network and associated protocols are invisible to the FC network
 - The FC network and associated protocol are invisible to the IP network

An FCIP ISL

Does it matter that an ISL is over IP?
No, it's still just an ISL



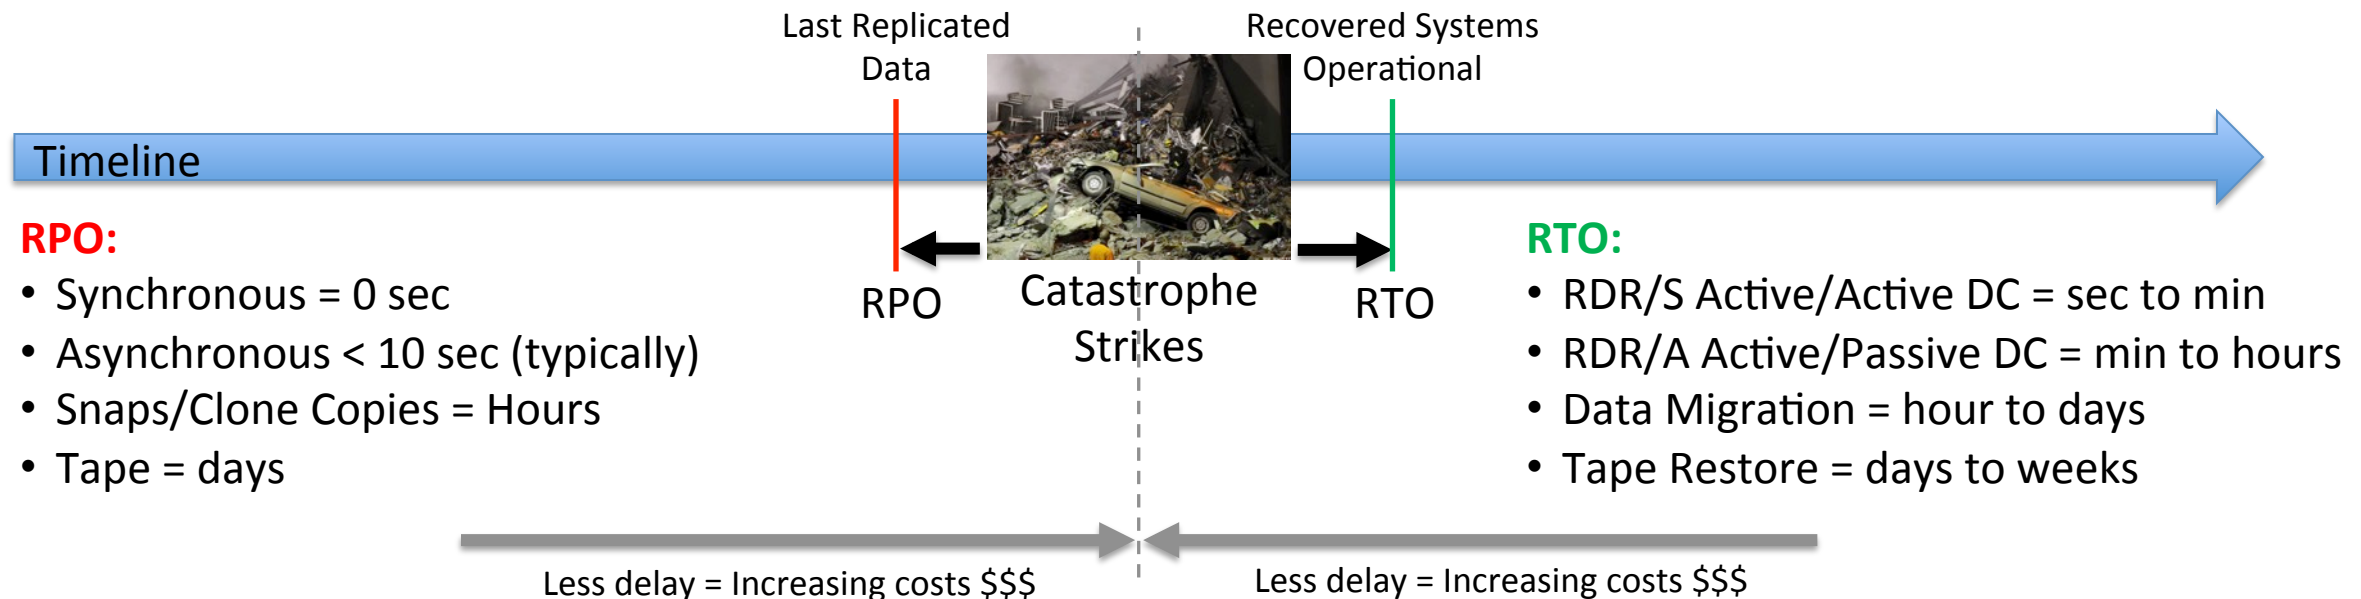
FCIP is TCP Based

- Reliable Transport
 - Every byte is acknowledged, otherwise, it is resent
- Connection Oriented
 - Both sides agree to data transmission
- Stream Oriented
 - TCP preserves exact byte sequence sent
- Unstructured Data
 - TCP does not denote discrete data fields in its transmission
- Full Duplex
 - Concurrent transfers in both directions
- BW Delay Product
 - Accommodates large pipes at long distances

RPO & RTO

DR Management Objectives

- Extension RPO/RTO are relatively short
- RPO = Recovery Point Objective
 - Time last bit of data was successfully saved outside disaster zone
- RTO = Recovery Time Objective
 - How quickly can systems be brought back online with the data?



Extending Storage over an IP WAN

Extension

- Advantages
 - Ubiquitous IP WAN Infrastructure
 - Cost effective IP WAN Infrastructure
 - Get outside of disaster area
 - Access data globally
 - Multiple paths/carriers
- Challenges
 - RPO: Propagation delay – sync vs. async
 - RTO: Remote data safe-keeping vs. Remote continued operations?
 - Security beyond the data center walls
 - WAN bandwidth, efficient utilization, and availability
 - Shared enterprise IP WAN links

Extension Attributes


- Performance
 - Maintain performance across global distances
- Flexibility
 - Wide variety of IP network connectivity types
- Resource optimization
 - WAN optimization techniques
 - Protocol acceleration techniques
 - Compression techniques
 - Security techniques
- Increased network integrity
 - High Availability techniques
 - Error recovery techniques
 - Fault, protocol, and fabric services isolation
 - Best practices

FCIP Technologies

- High Availability
- Compression
- Security/Encryption
- Storage replication mode
- IP WAN & Infrastructure
- Protocol, Encapsulation and Optimization
- QoS
- FC Fabrics and FC Routing (previous webinar on fc.org)

Terminology & Concepts

- RDR = Remote Data Replication
- Effective Throughput
 - Is the storage end-device throughput
 - Is not the WAN throughput
 - Effected by:
 - Compression
 - IPsec overhead
 - FCIP/TCP/IP overhead ($\approx 1.5\%$ with IPsec and without compression = worse case)
- Latency
 - Referred to as RTT (Round Trip Time) or Propagation Delay
 - Typically WAN is measured in milliseconds (ms)
 - RTT includes speed of light through cable medium + intermediate network equipment
 - Fiber propagation delay is $\approx 10 \mu\text{s}/\text{km}$ RTT
 - Each router/switch adds more latency
 - Network buffering/congestion = data sitting idle = increased latency
 - TCP retransmission of lost data adds a RTT to overall delay



Synchronous / Asynchronous

- Synchronous RPO $t=0$
 - All safely remote written data is acknowledged to host
 - Not tolerant of WAN latency
- Asynchronous RPO $t < 10$ seconds (typically)
 - Data written locally is acknowledged to host
 - Tolerant of WAN latency
 - Data may not be safely remotely written at time of catastrophe
- IP WAN per distance km, tends to be less costly
 - Dark Fiber, xWDM and Lambda services are expensive over long distance
 - Dark Fiber, xWDM and Lambda services may not exist over long distances
 - Asynchronous FCIP replication is most popular over long distance


Designing an RDR FCIP Network

Storage Subsystem Characteristics

- Synchronous, Asynchronous or Copy?
 - Sync & Async maintain write order and supports consistency groups
 - Copy doesn't maintain write order
 - Bandwidth requirements for each are different
 - WAN latency requirements are different
- How fast can the array deliver data?
 - Per LUN? How many LUNs?
 - Per FC port? How many ports?
 - Per Application?
- What is the data delta per period of time?
 - Do you have adequate cache and/or journal volumes?

Synchronous RDR (RDR/S)

Array via FC Direct Connect



- **As distance increases – so does I/O response time**
- R1 does not respond until write confirmation from R2
- All writes guaranteed to be safe at remote location (RPO = 0)
- FC switches used for adequate number of buffer credits

Primary Site



Metro distance with fiber infrastructure



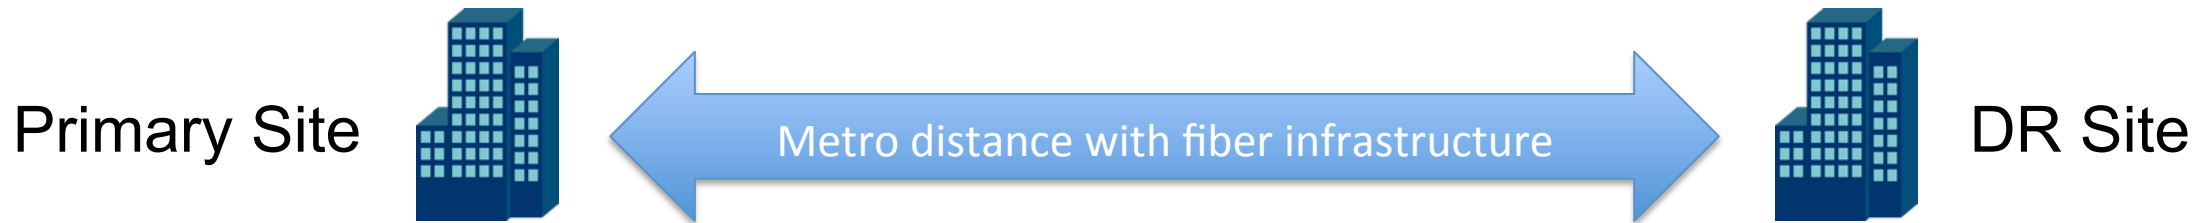
DR Site

Synchronous RDR

Extension via Direct Connect




- Extension can be used with RDR/S, will increase I/O response time
- Provides encryption
- Provides Trunking/Port Channel across multiple paths

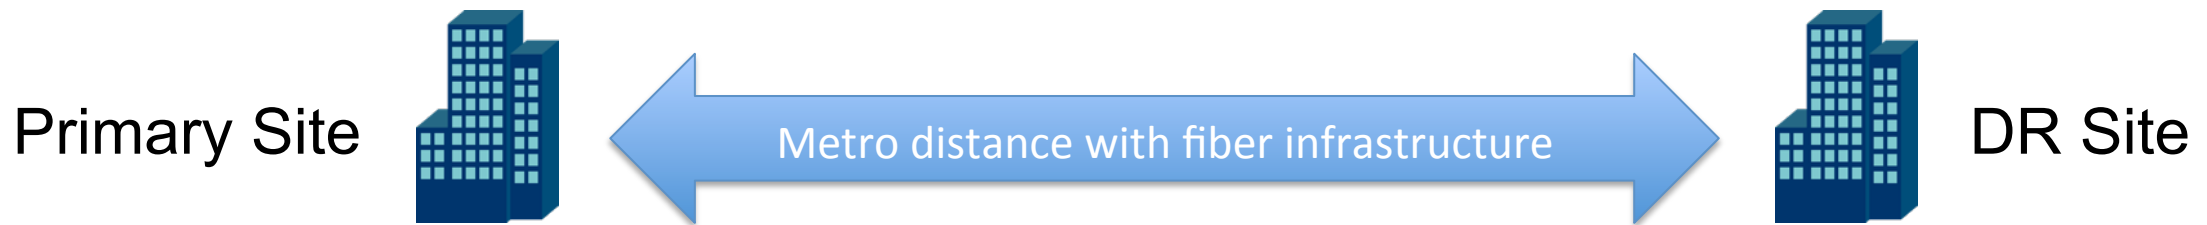


Synchronous RDR

Extension via IP Network

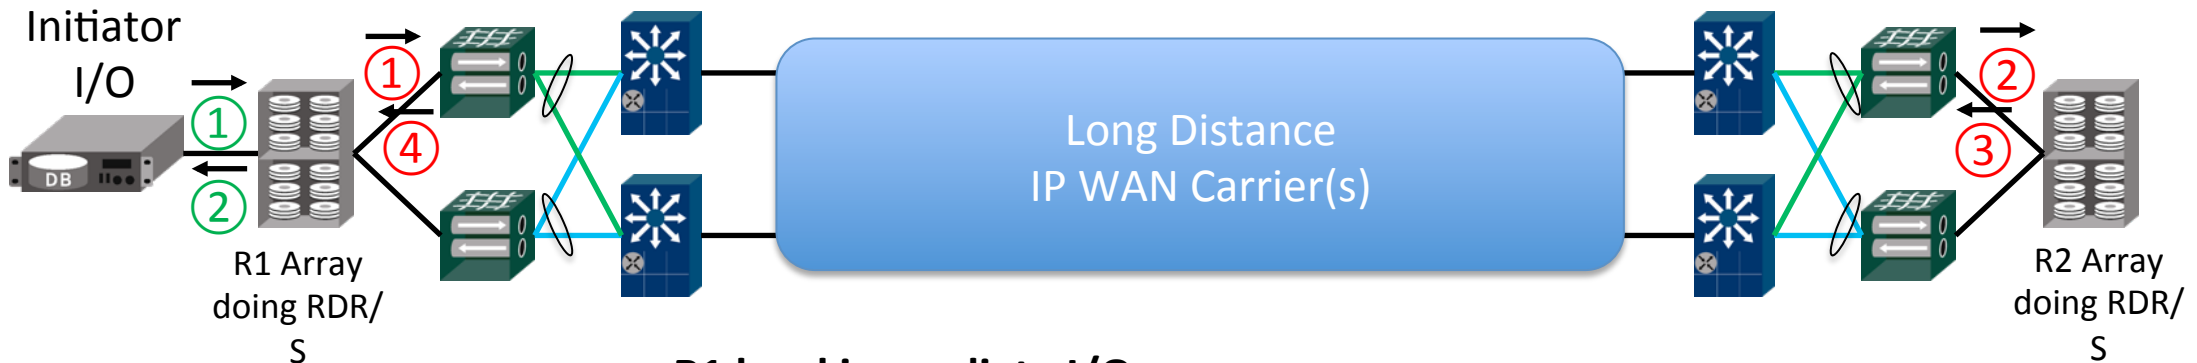


- If IP network is used, the IP network must accommodate synchronous demands
 - Paths not shared
 - Accommodates peak bandwidth
 - Not oversubscribed and no congestion possible (no buffering or retransmissions)



Asynchronous RDR

Extension via IP Network



- **R1 local immediate I/O response**
- Distance has **no effect** on I/O response time
- Added propagation delay of Extension equipment is irrelevant
- I/O is not yet written to remote safe location (**RPO = delta t**)

Primary Site



Long distance IP WAN infrastructure




DR Site

Bandwidth Requirements

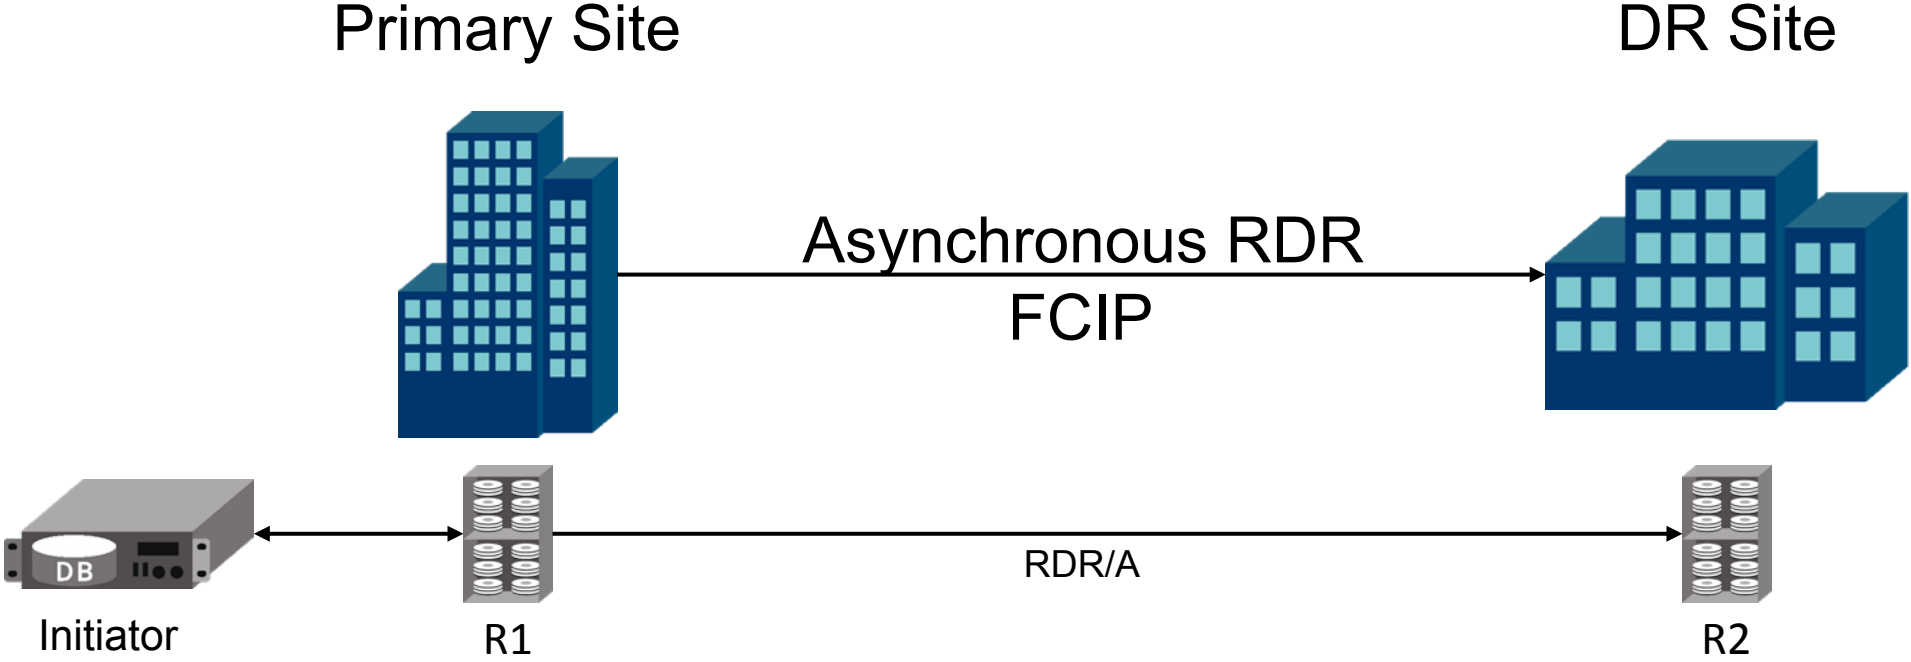
Synchronous vs. Asynchronous

- If RDR/S exceeds the available link capacity:
 - WAN bandwidth must accommodate peak load
 - Network buffered writes or lost packets result in poor response times
 - Result = Poor I/O response times
- If RDR/A or copy exceeds available link capacity:
 - WAN bandwidth should accommodate average load over finite period
 - Writes are cached/journaled in the array or buffered in the network
 - Result = No negative effect on I/O response times



2 Data Centers

Array to Array RDR/A




2 Data Centers

Cascaded RDR

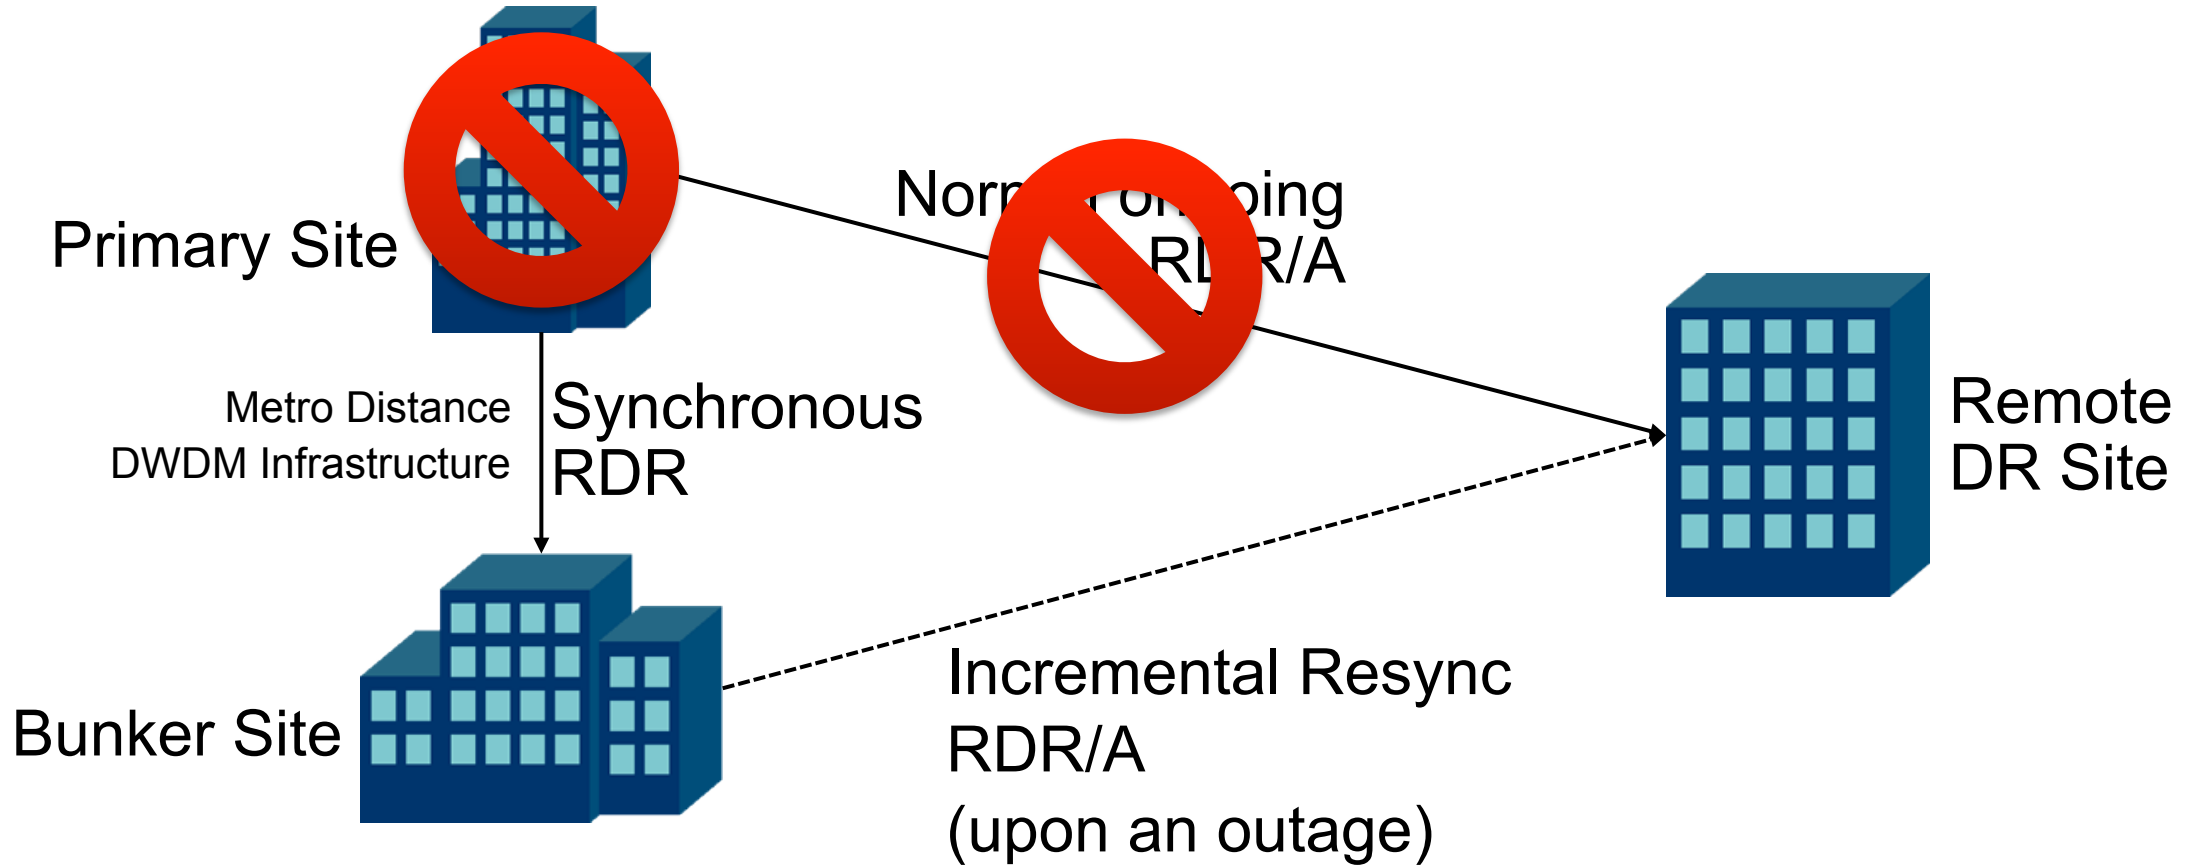
Primary Campus Site
or same DC R1 failure

DR Site



3 Data Centers

HPE 3PAR 3DC-PP, IBM MM/GM, HDS 3DC & EMC STAR



FCIP Interoperability

Brocade & Cisco FCIP Interoperability

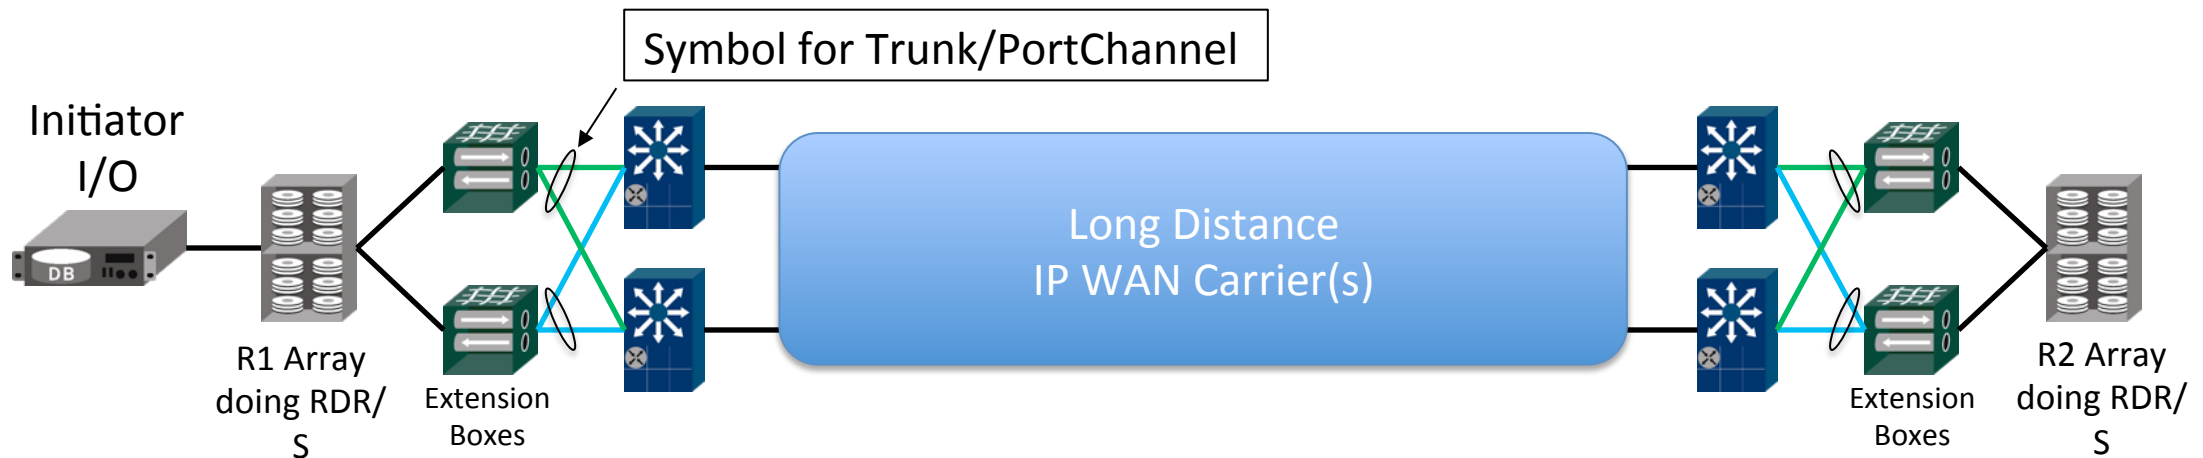
- Simply...
There's NO supported FCIP interoperability across vendors.

FCIP Features



Extension Trunking/Port Channels

- Bandwidth
 - Aggregated links
- Path Availability
 - Multiple disparate paths
- Load Sharing/Balancing
- Failover/Failback



FCIP Compression

- Multiple compression algorithms
 - Typically chosen based on amount of WAN BW
 - Optimized for ratio vs. throughput
 - Increases effective throughput (or)
 - Reduces required WAN bandwidth

Encryption

FCIP Data Inflight Security

- Is encryption a requirement (i.e., HIPPA) or just prudent?
- Encryption of data inflight only
 - No possibility of data loss
 - No key manager required
- Encryption for devices that:
 - Have a severe penalty when performing encryption
 - Cannot encrypt


Encryption

In the IP network?

- Best practice:
 - Encrypt from tunnel origin to destination
- Is doing encryption in the IP network better or faster?
 - Generally - No
 - More costly \$\$\$
 - Performance hit
 - Less control for storage admins (turn it on/off?)
 - Less secure, farther from the source/destination end-devices
 - Diagnostics & troubleshooting issues

Encryption

Prevent Eavesdropping and Attacks

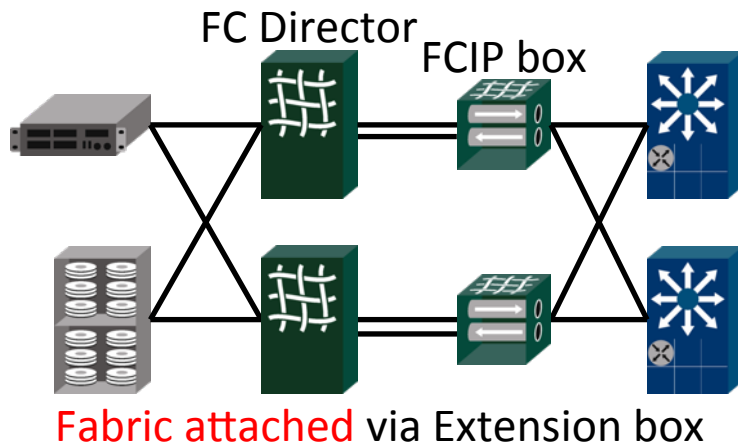


FCIP Deployment

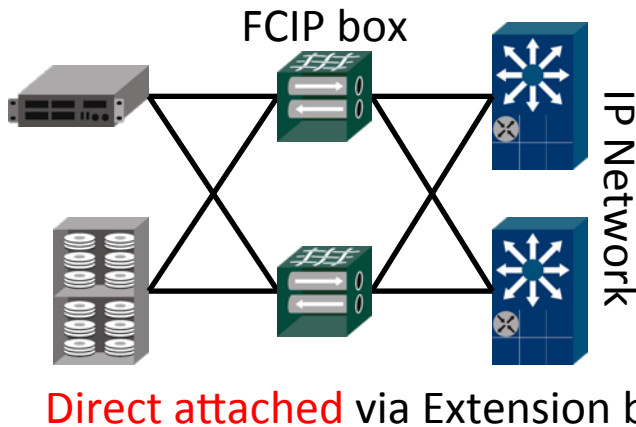


Designing an FCIP RDR Network

Direct or Fabric Attached?

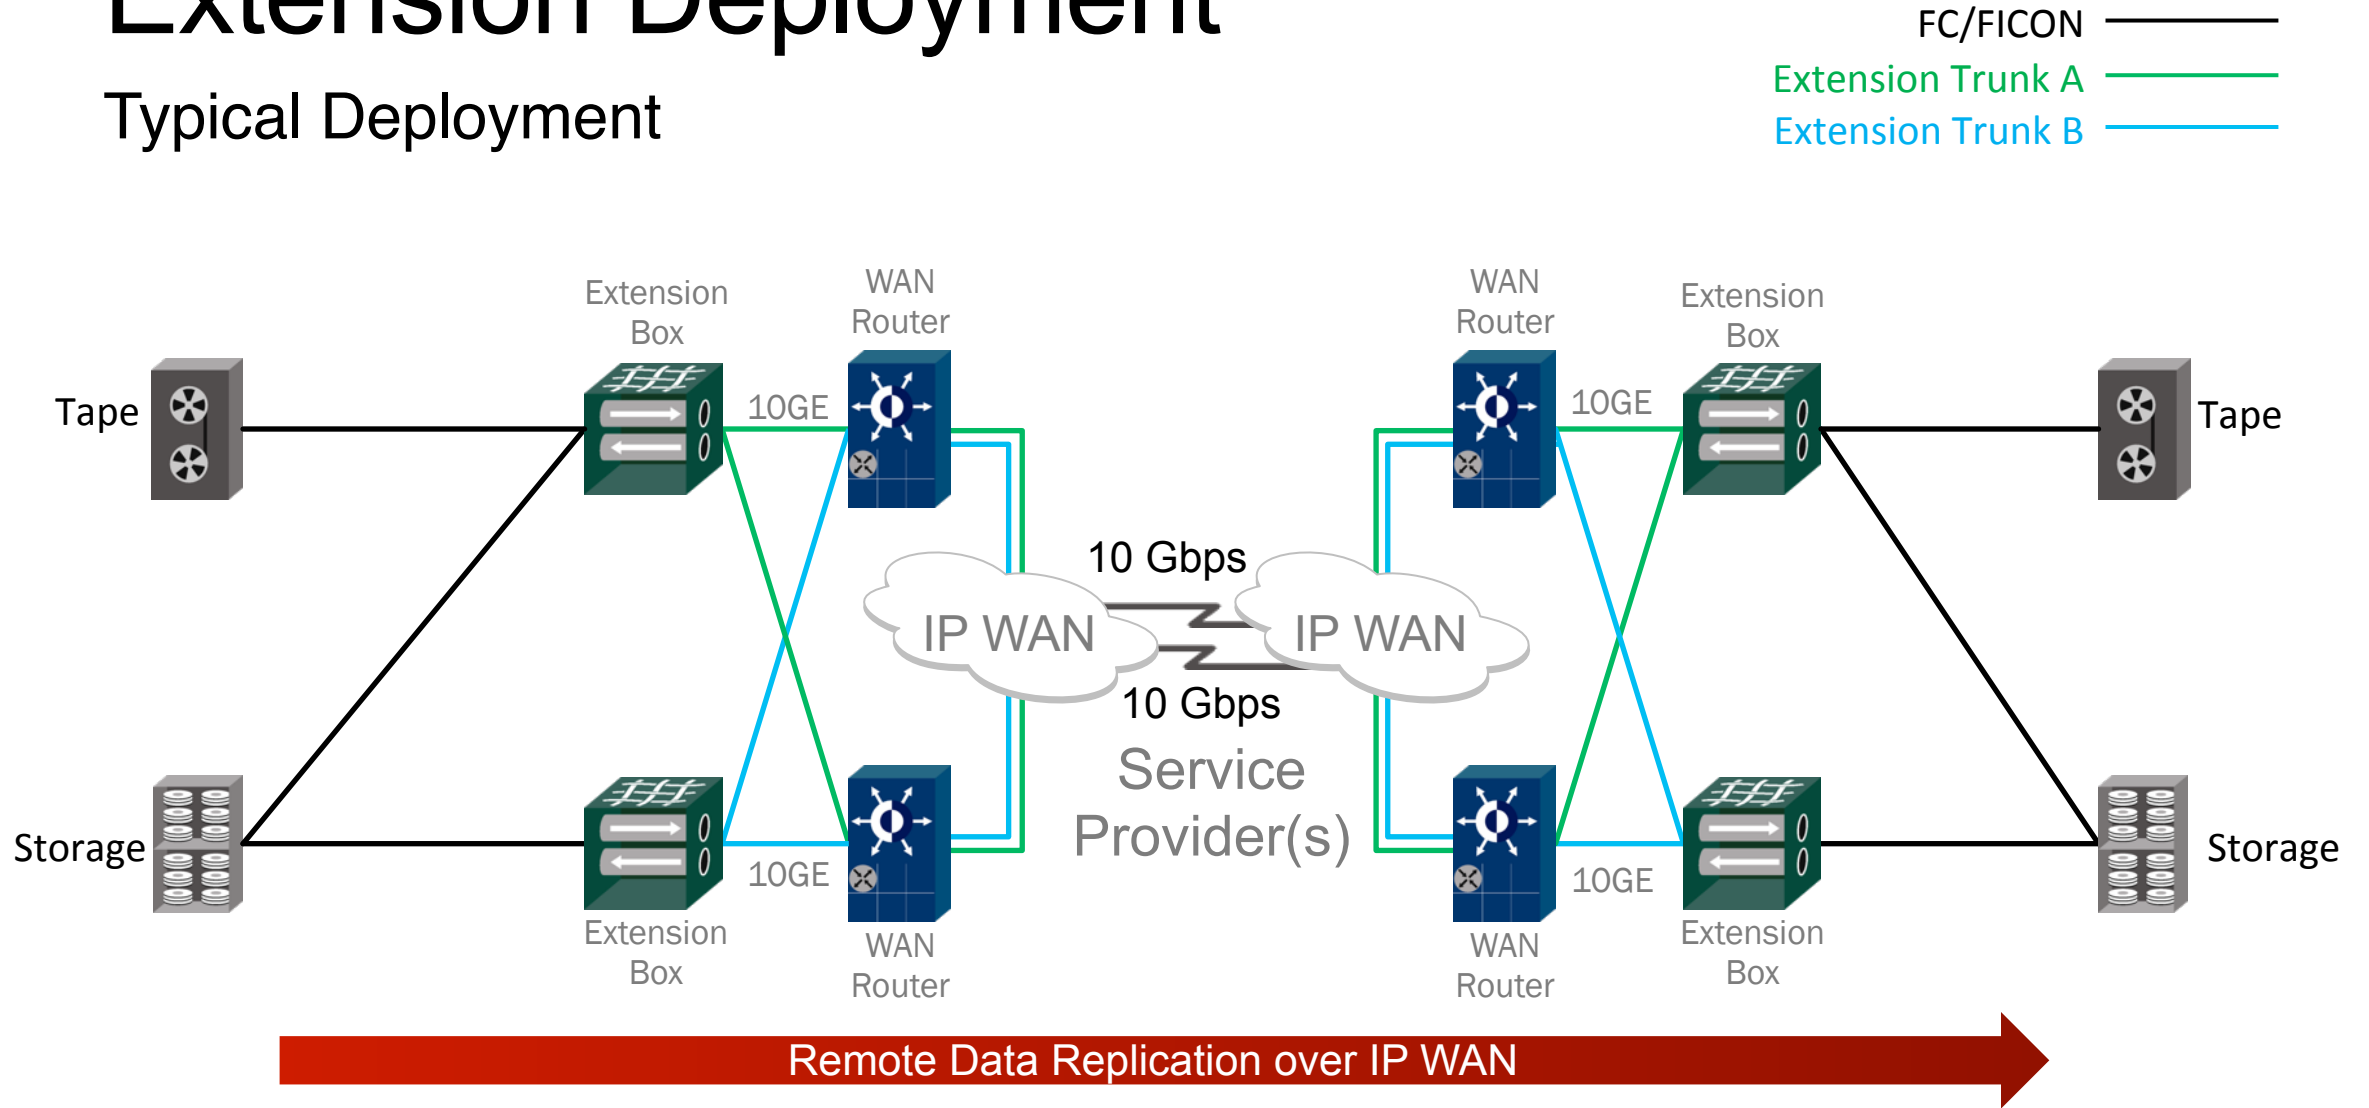


- Direct Attach or Fabric Attached?
 - Best practice is direct attached
 - Scalable solution is fabric attached
 - Don't preclude fabric attached, if it makes sense
- Why is direct attach best practice?
 - Less risk of unplanned disruption to production fabric
 - Firmware upgrades independent of production fabric
 - Single extension box at site doesn't compromise A & B fabrics
 - Less complexity (KISS)
 - Less path points of failure
 - Less chance for configuration error
 - Isolate RDR network from the production, just because you can't be too cautious!



Extension Deployment

Typical Deployment



Summary

FCIA FCIP Webinar

- Disasters could effect anyone
- Replication of data outside of the metro area is prudent
- Long distance replication usually requires asynchronous
- Asynchronous replication doesn't affect I/O response time
- FCIP is typically applied to asynchronous replication
- FCIP features HA, trunking, compression and encryption
- FCIP can be direct or fabric attached
- FCIP offers various architectures for redundancy

Our Next FCIA Webcast:

Fibre Channel Performance, Congestion, Slowdrain and Over Utilization

February 6, 2018

10:00 am PT

Register at:

<https://www.brighttalk.com/webcast/14967/295141>

After this Webcast

- Please rate this event – we value your feedback
- We will post a Q&A blog at <http://fibrenchannel.org/> with answers to all 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://fibrenchannel.org/webcasts/> to learn about:
 - Fibre Channel Fundamentals
 - FC-NVMe
 - Long Distance Fibre Channel
 - Fibre Channel Speedmap

Thank You

