

# Understanding Fibre Channel Scaling

Live Webcast  
November 6, 2019  
10 am PT



# SPEAKERS



Brandon Hoff  
Broadcom  
Software Architect



Mark Rogov  
Dell EMC  
Systems Engineer  
@rogovmark



Dr. J Metz  
Cisco  
R&D Engineer  
@drjmetz

# About the FCIA

- The Fibre Channel Industry Association (FCIA) is a mutual benefit, non-profit, international organization of manufacturers, system integrators, developers, vendors, and industry professionals, and end users
  - Promotes the advancement of Fibre Channel technologies and products that conform to the existing and emerging T11 standards
  - Maintains resources and supports activities to ensure multi-vendor interoperability for hardware, interconnection, and protocol solutions
  - Provides promotion and marketing of FC solutions, educational awareness campaigns, hosting public interoperability demonstrations, and fosters technology and standards conformance

<https://fibrechannel.org/>

3




# Agenda

- Survey
- Lossless via buffer credits
- Topology
- ISL oversubscription
- Mesh example
- Fabric Services
- Zoning
- Deterministic Performance

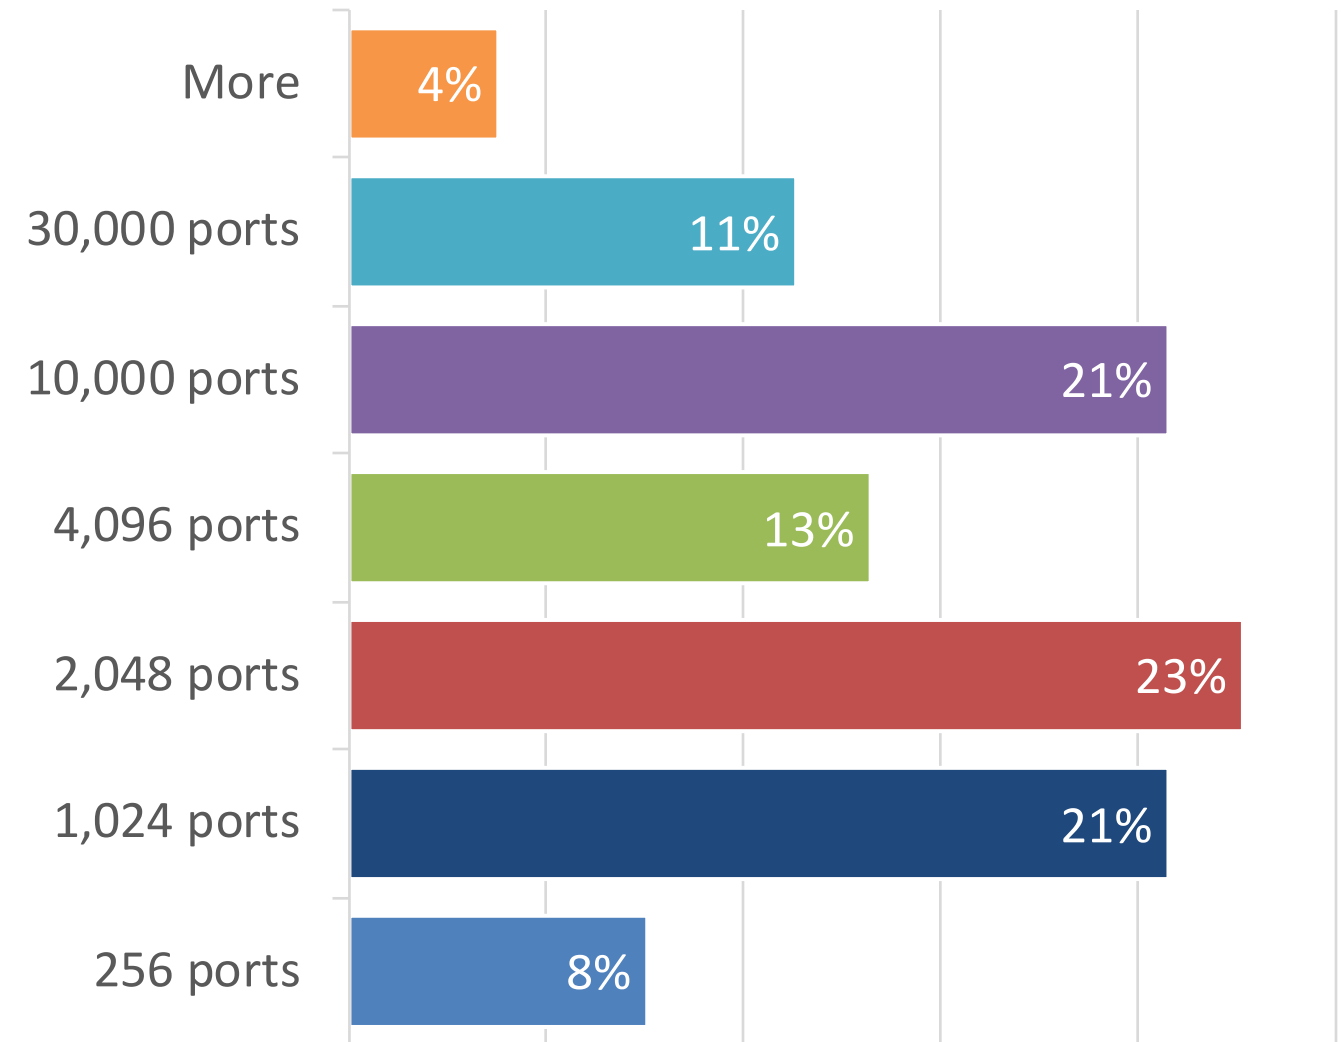
# Fibre Channel Flexibility

Storage vendor internal survey, 2019

## Typical Fabric



## Largest Fabrics



# Lossless & Buffer Credits

# Do you think your customers know what Buffer-to-Buffer credits are?

Yes

18%


No

82%

**FIBRE CHANNEL BUFFERS AND CREDITS  
ENABLE LOSSLESS COMMUNICATION**

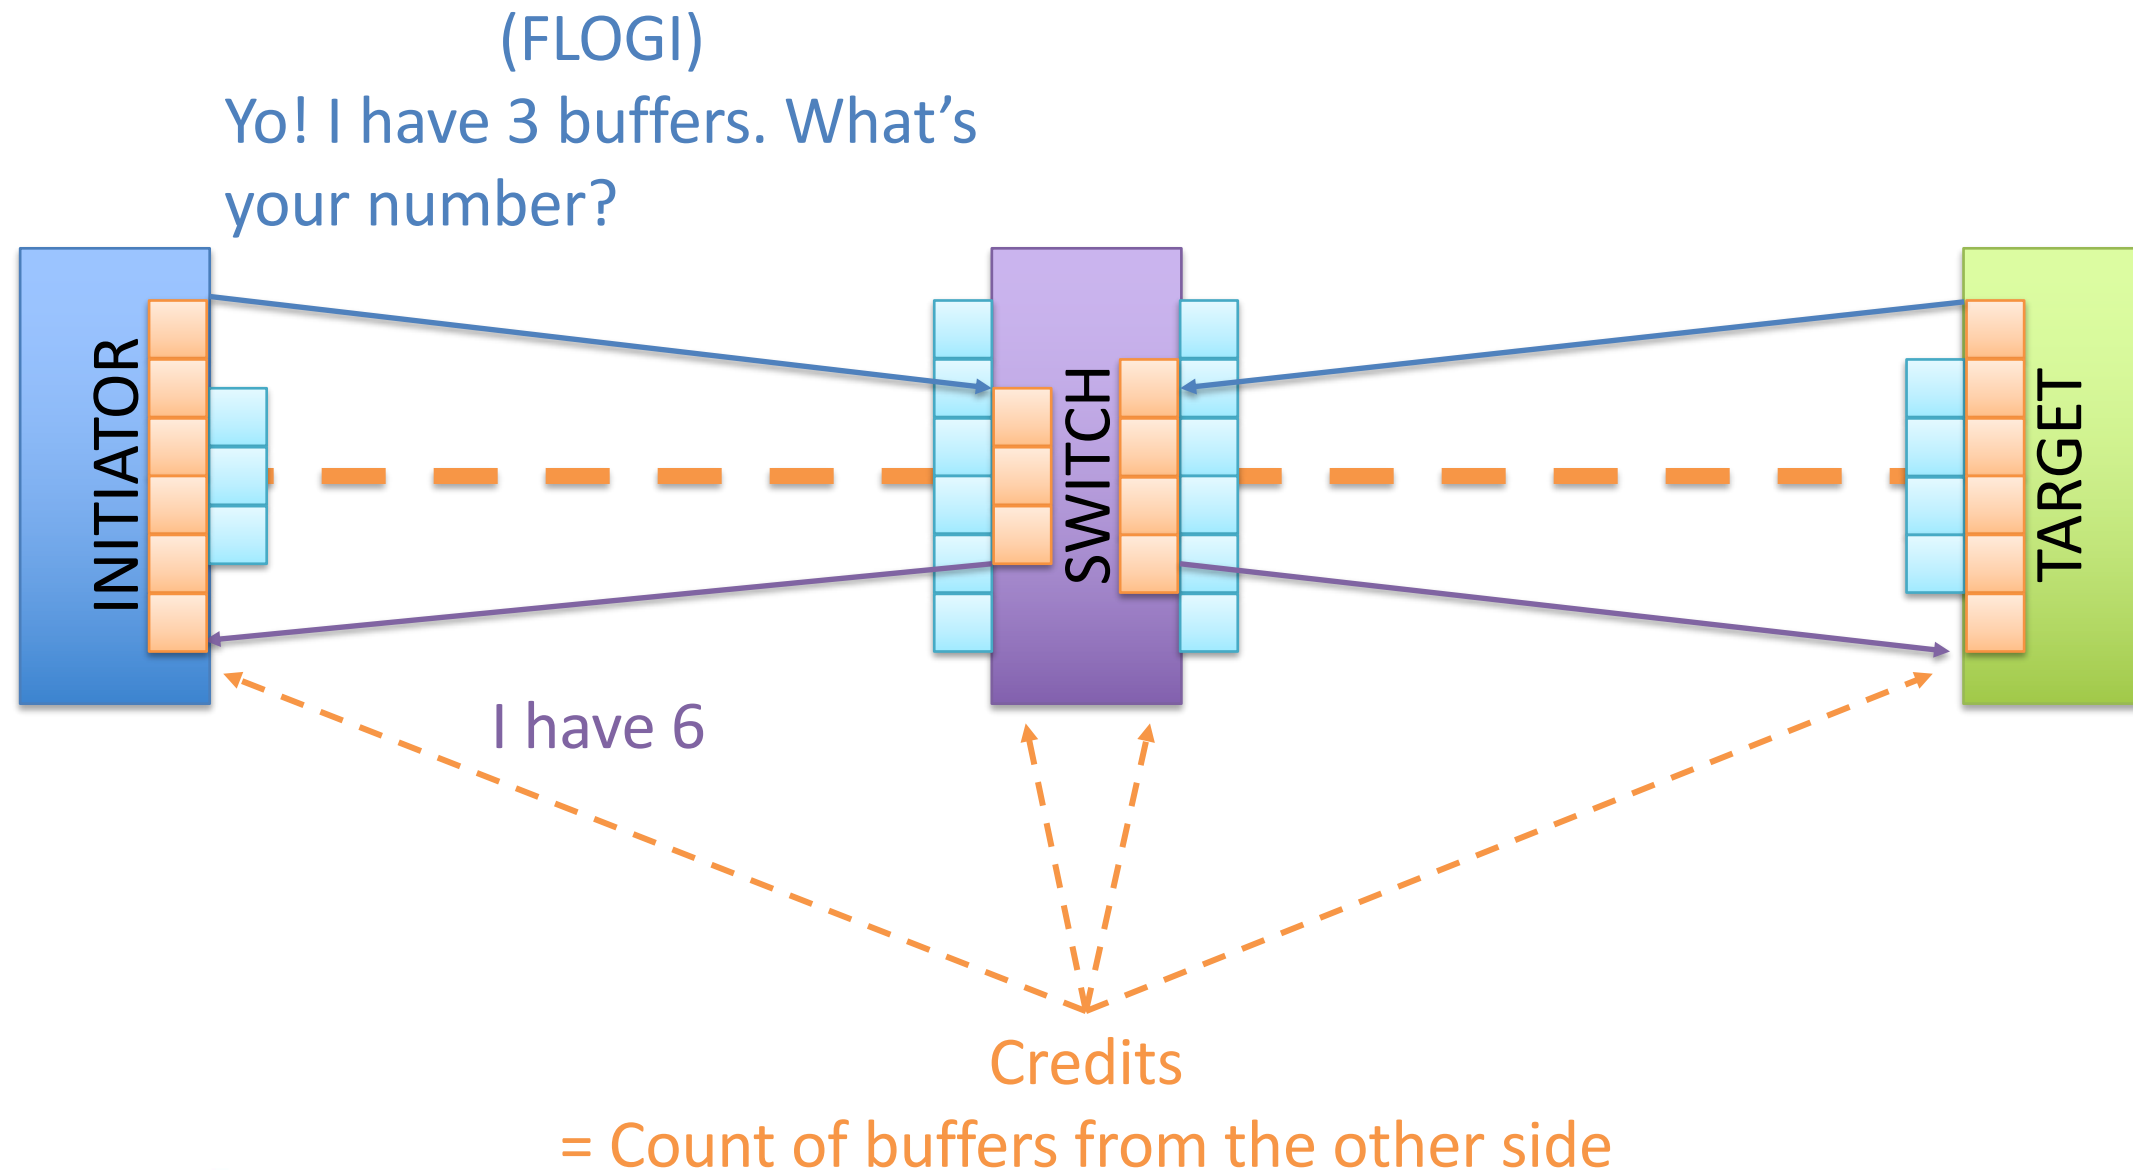
Source: Storage Vendor internal survey, 2019

# Buffers and Credits







# Buffers and Credits



# Buffers and Credits

Yo Target! Get me some data (aka READ)




Now, what happens next?  
All buffers and credits are used!

# Buffers and Credits

Oh, that was good!  
I have 2 buffers now!  
(R\_RDY)

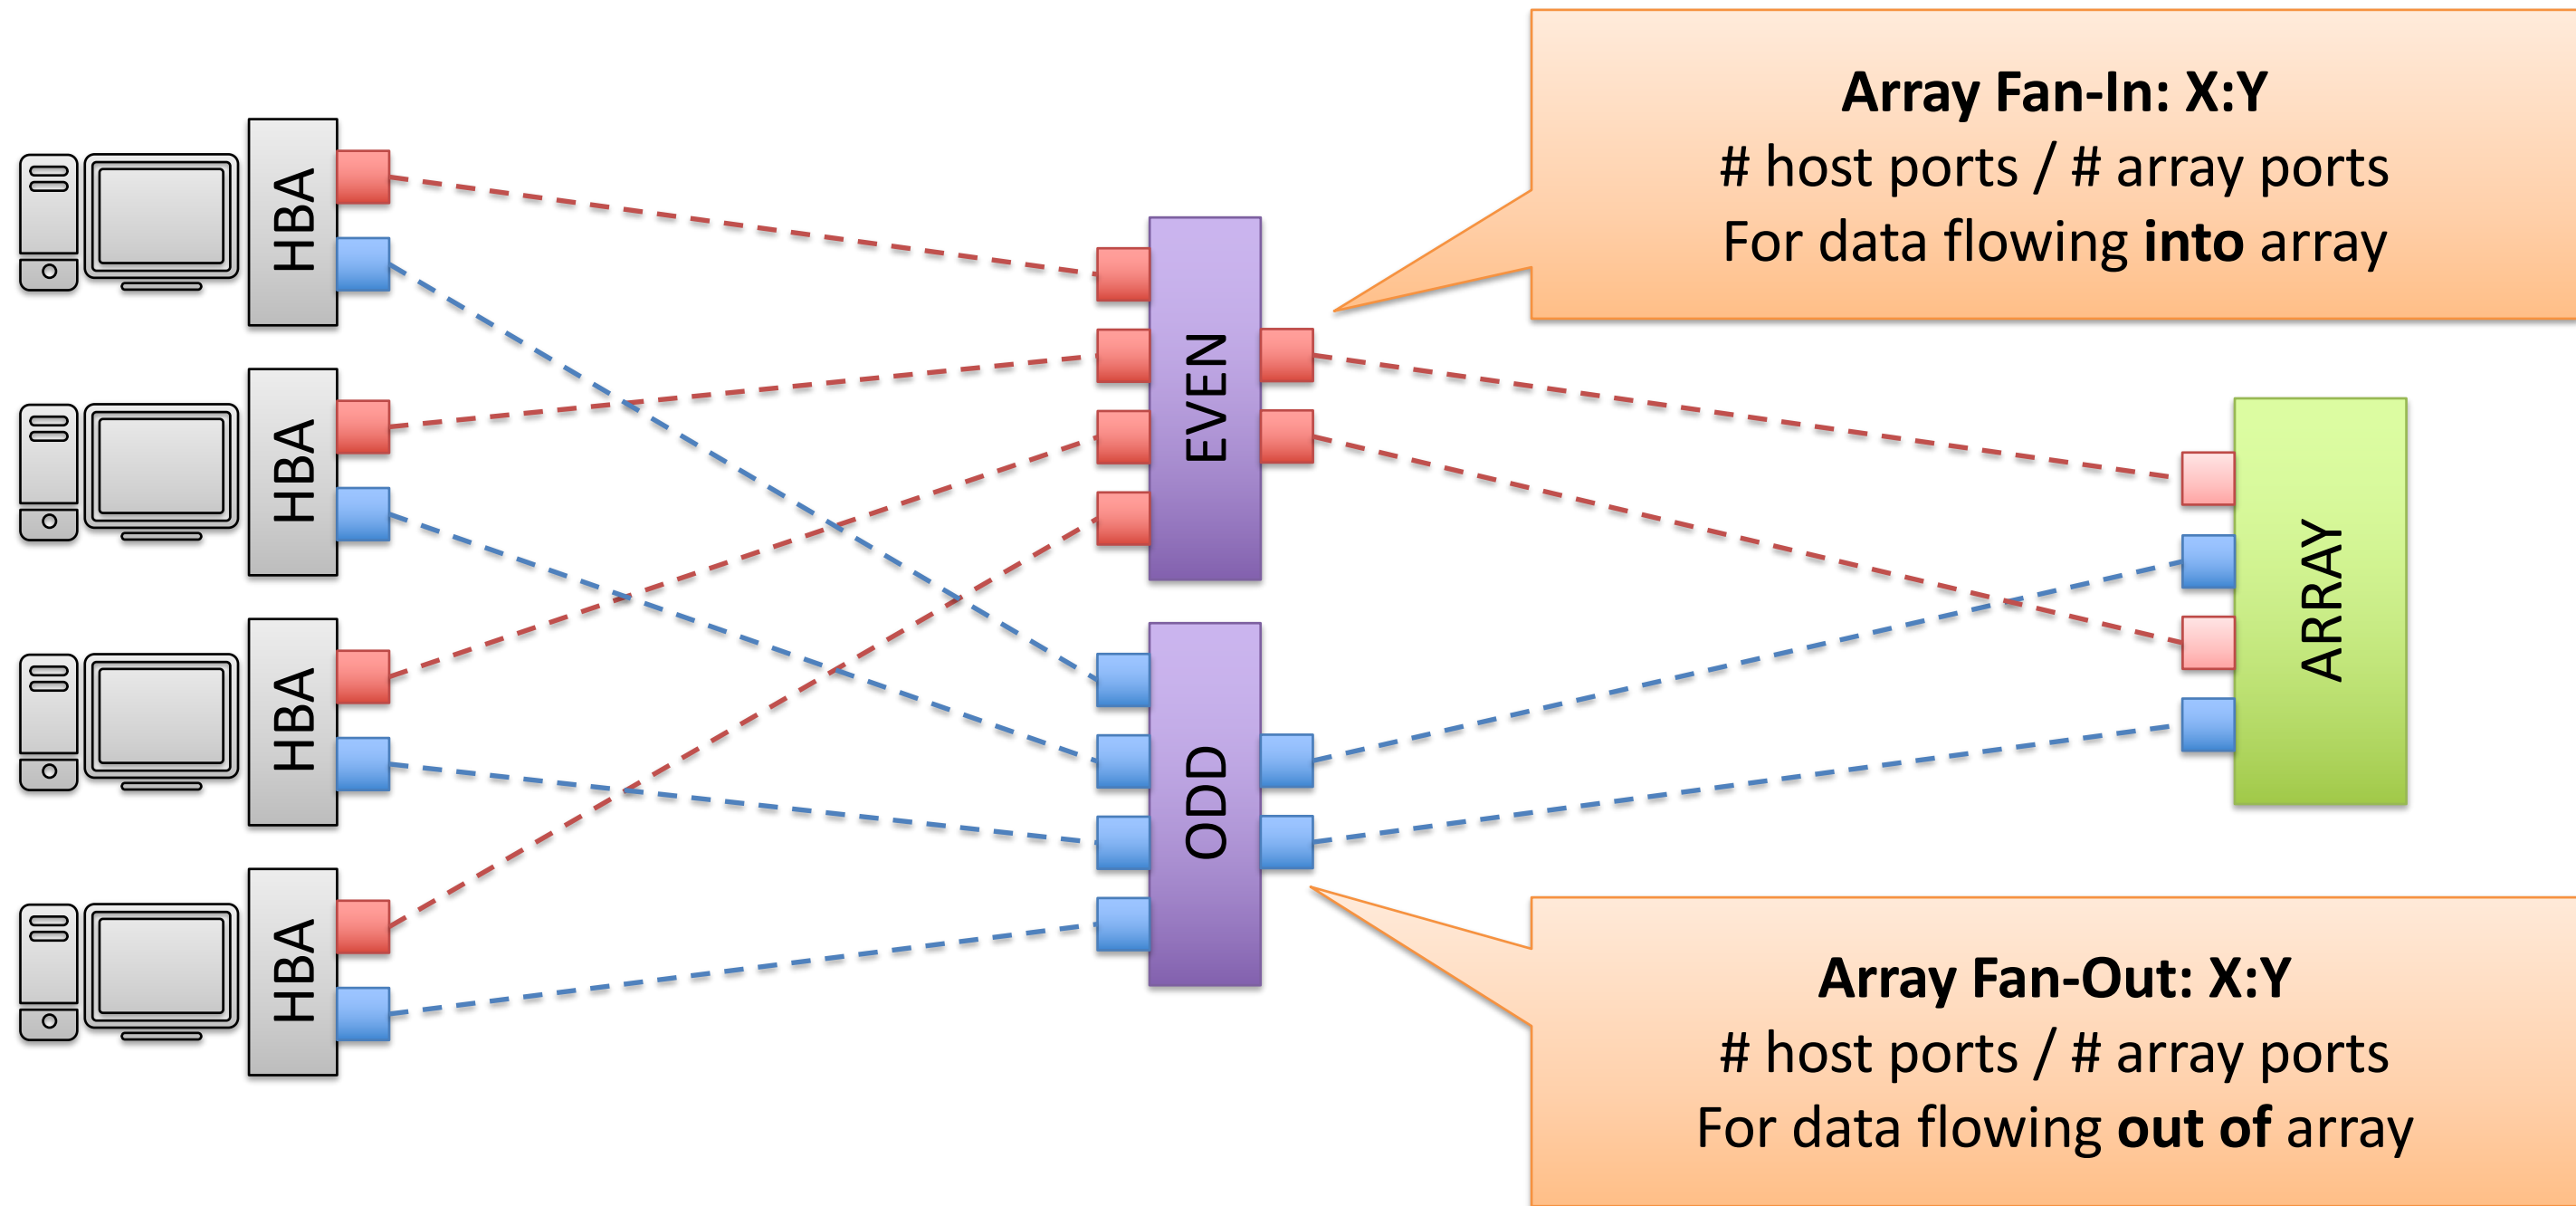
2 freed buffers (R\_RDY)



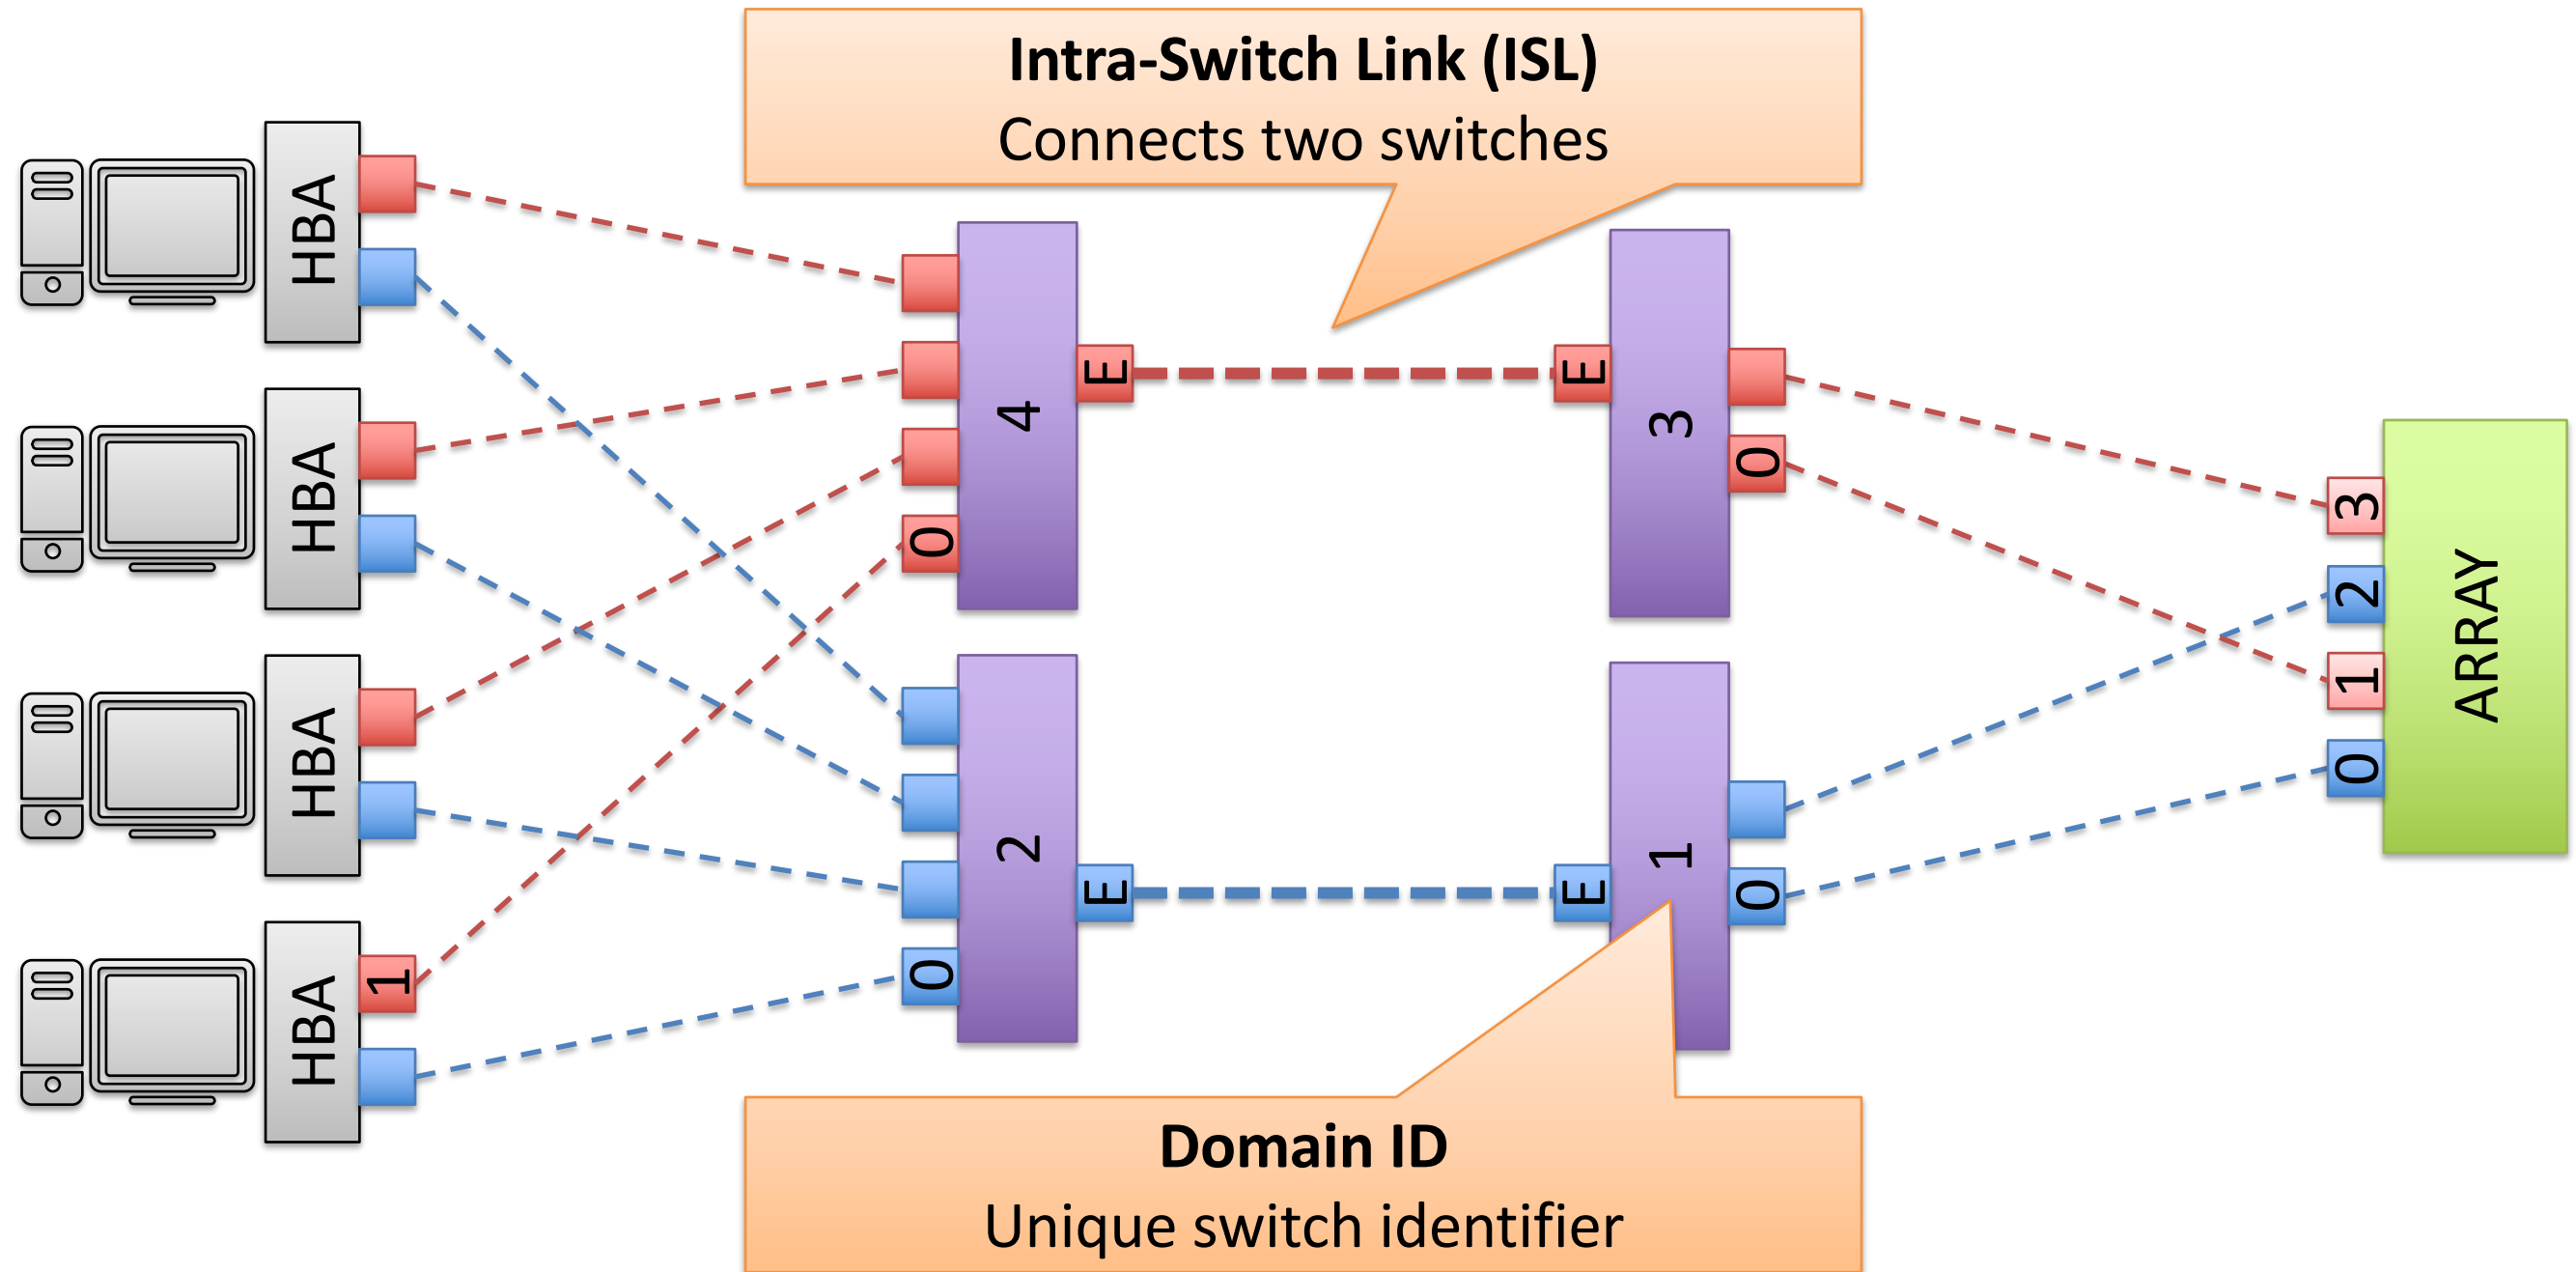
Host processed/received Two IOs  
Two more are waiting.  
Everyone knows the count  
**LOSSLESS COMMUNICATION**

# Topology

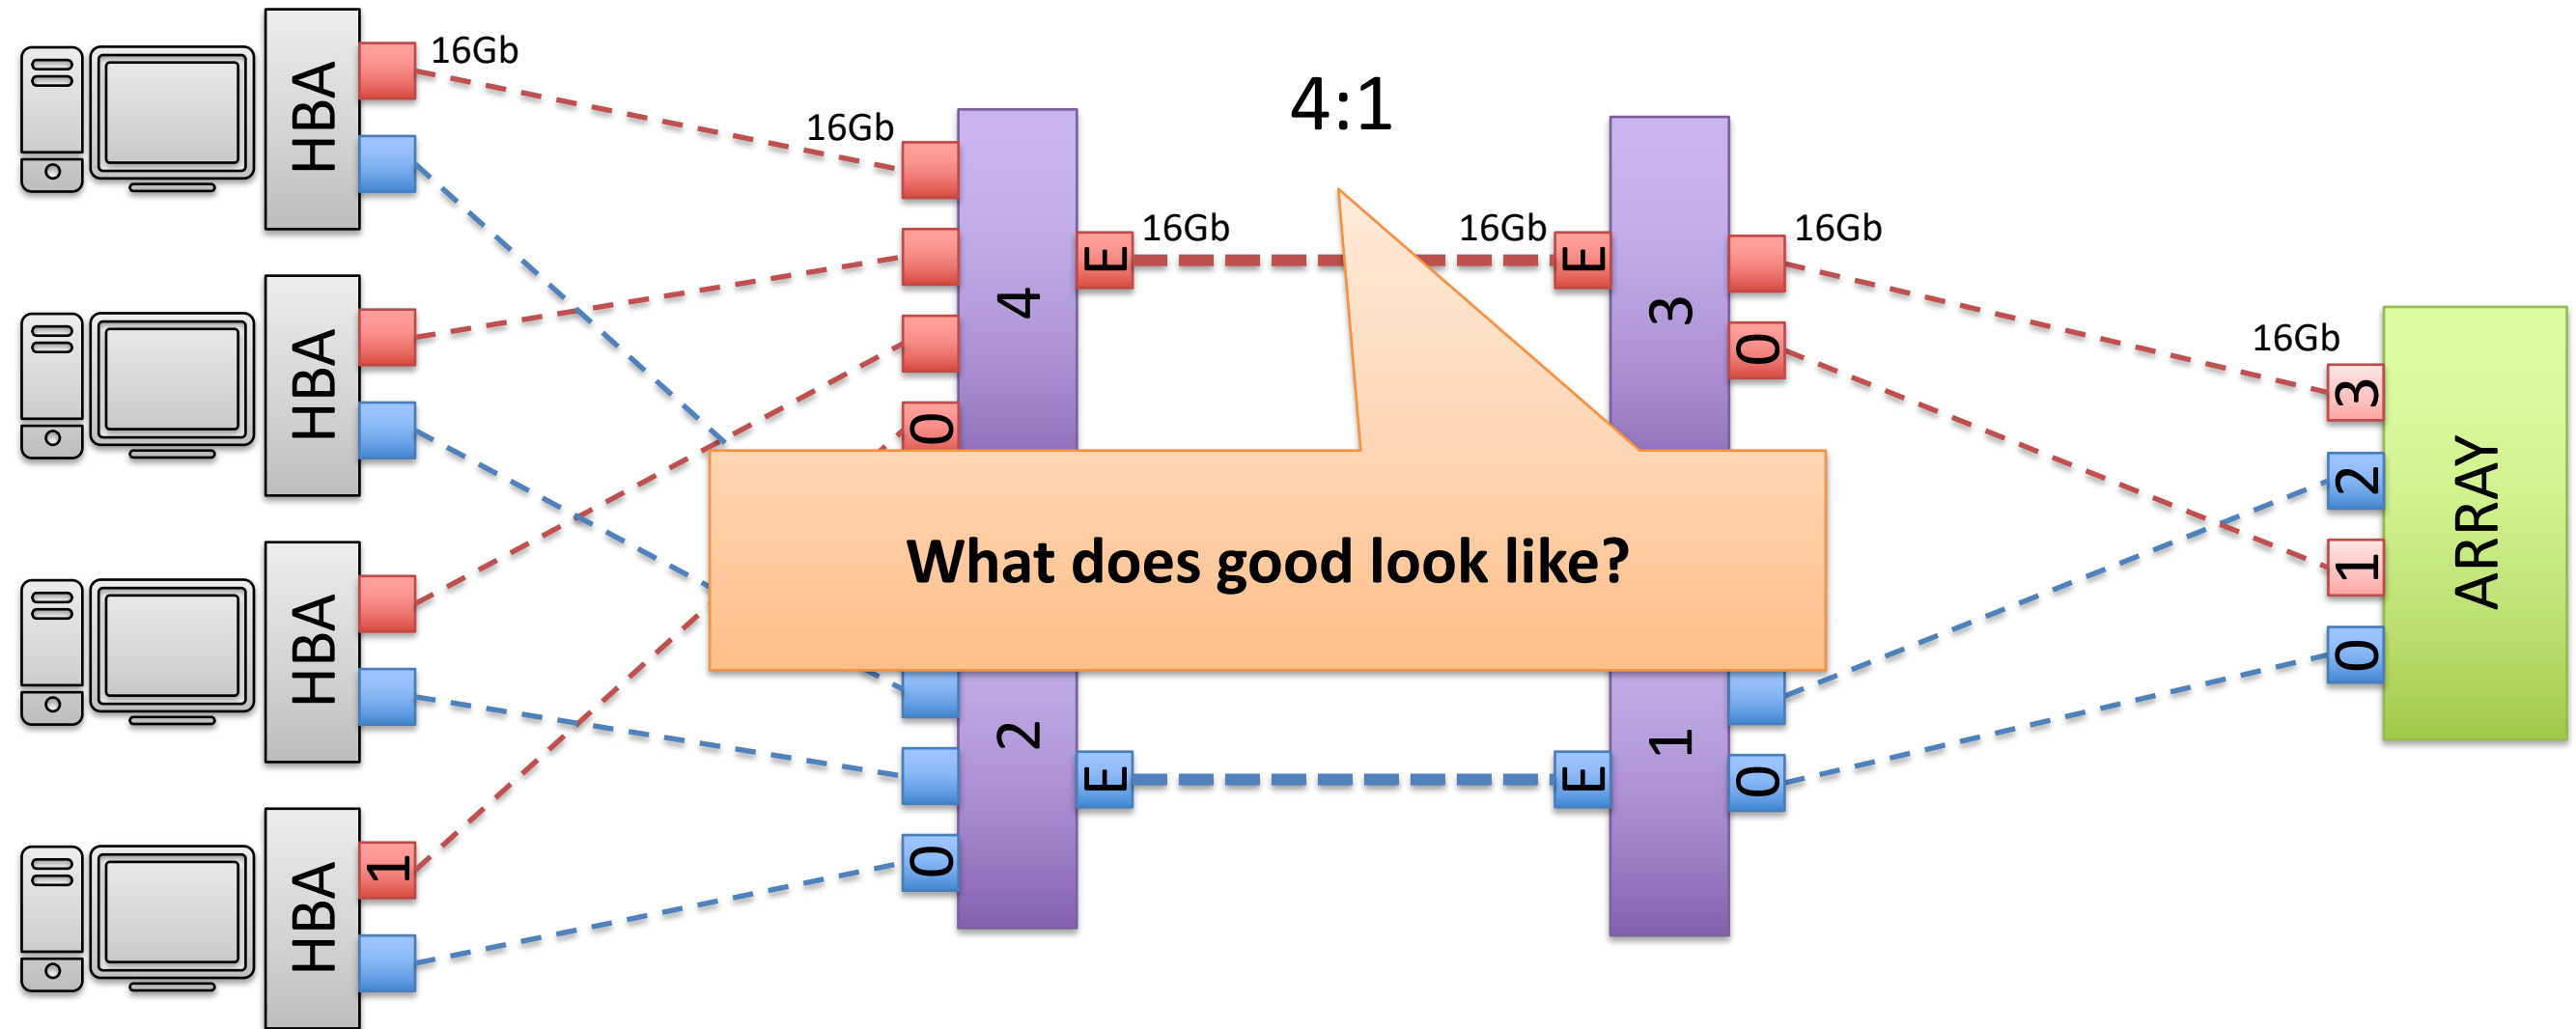
# Switched Fabric




# Domain & ISL



# ISL Oversubscription



# ISL Oversubscription



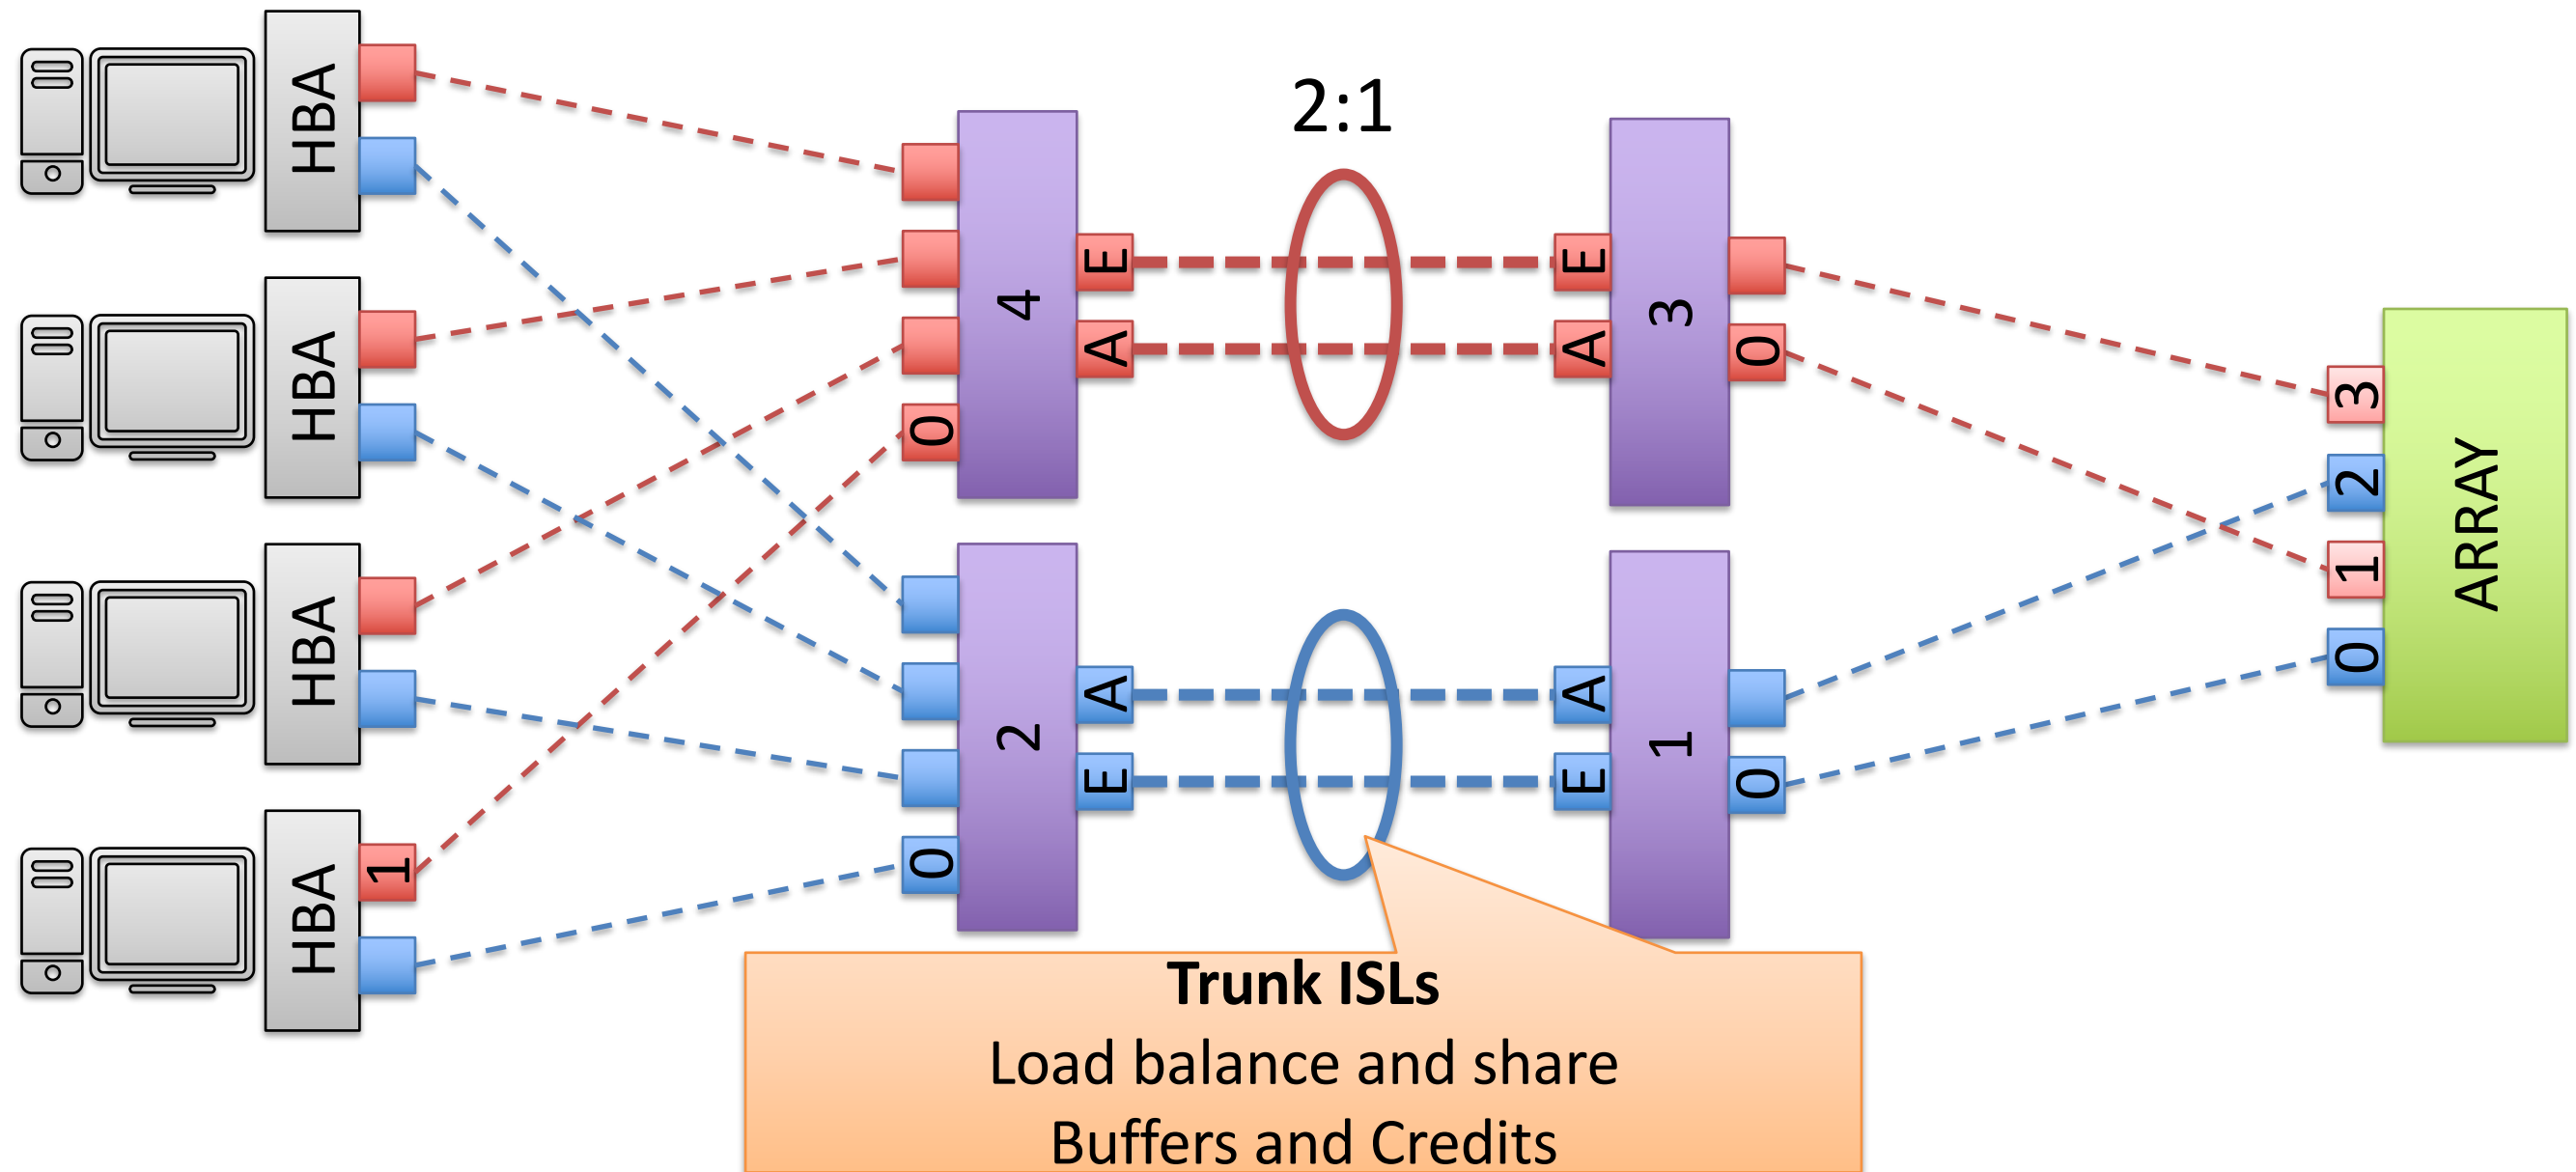
## Workload Considerations:

Think through congestion sensitivity  
and traffic direction

	Block Size	Read/Write	Random/Seq	ISL Oversubscription
OLTP (Database)	8KB	25R/75W	80R/20S	Lower
OLAP (Warehouse)	64KB	75R/25W	20R/80S	Higher
IoT	Mix (small)	5R/95W	5R/95S	Higher
VDI	4-16KB	50R/50W	100R	Mid




# Decrease ISL Oversubscription



# Scalable Topologies

## Core-Edge



Hosts


Edge Switches

Core Switches

Edge Switches

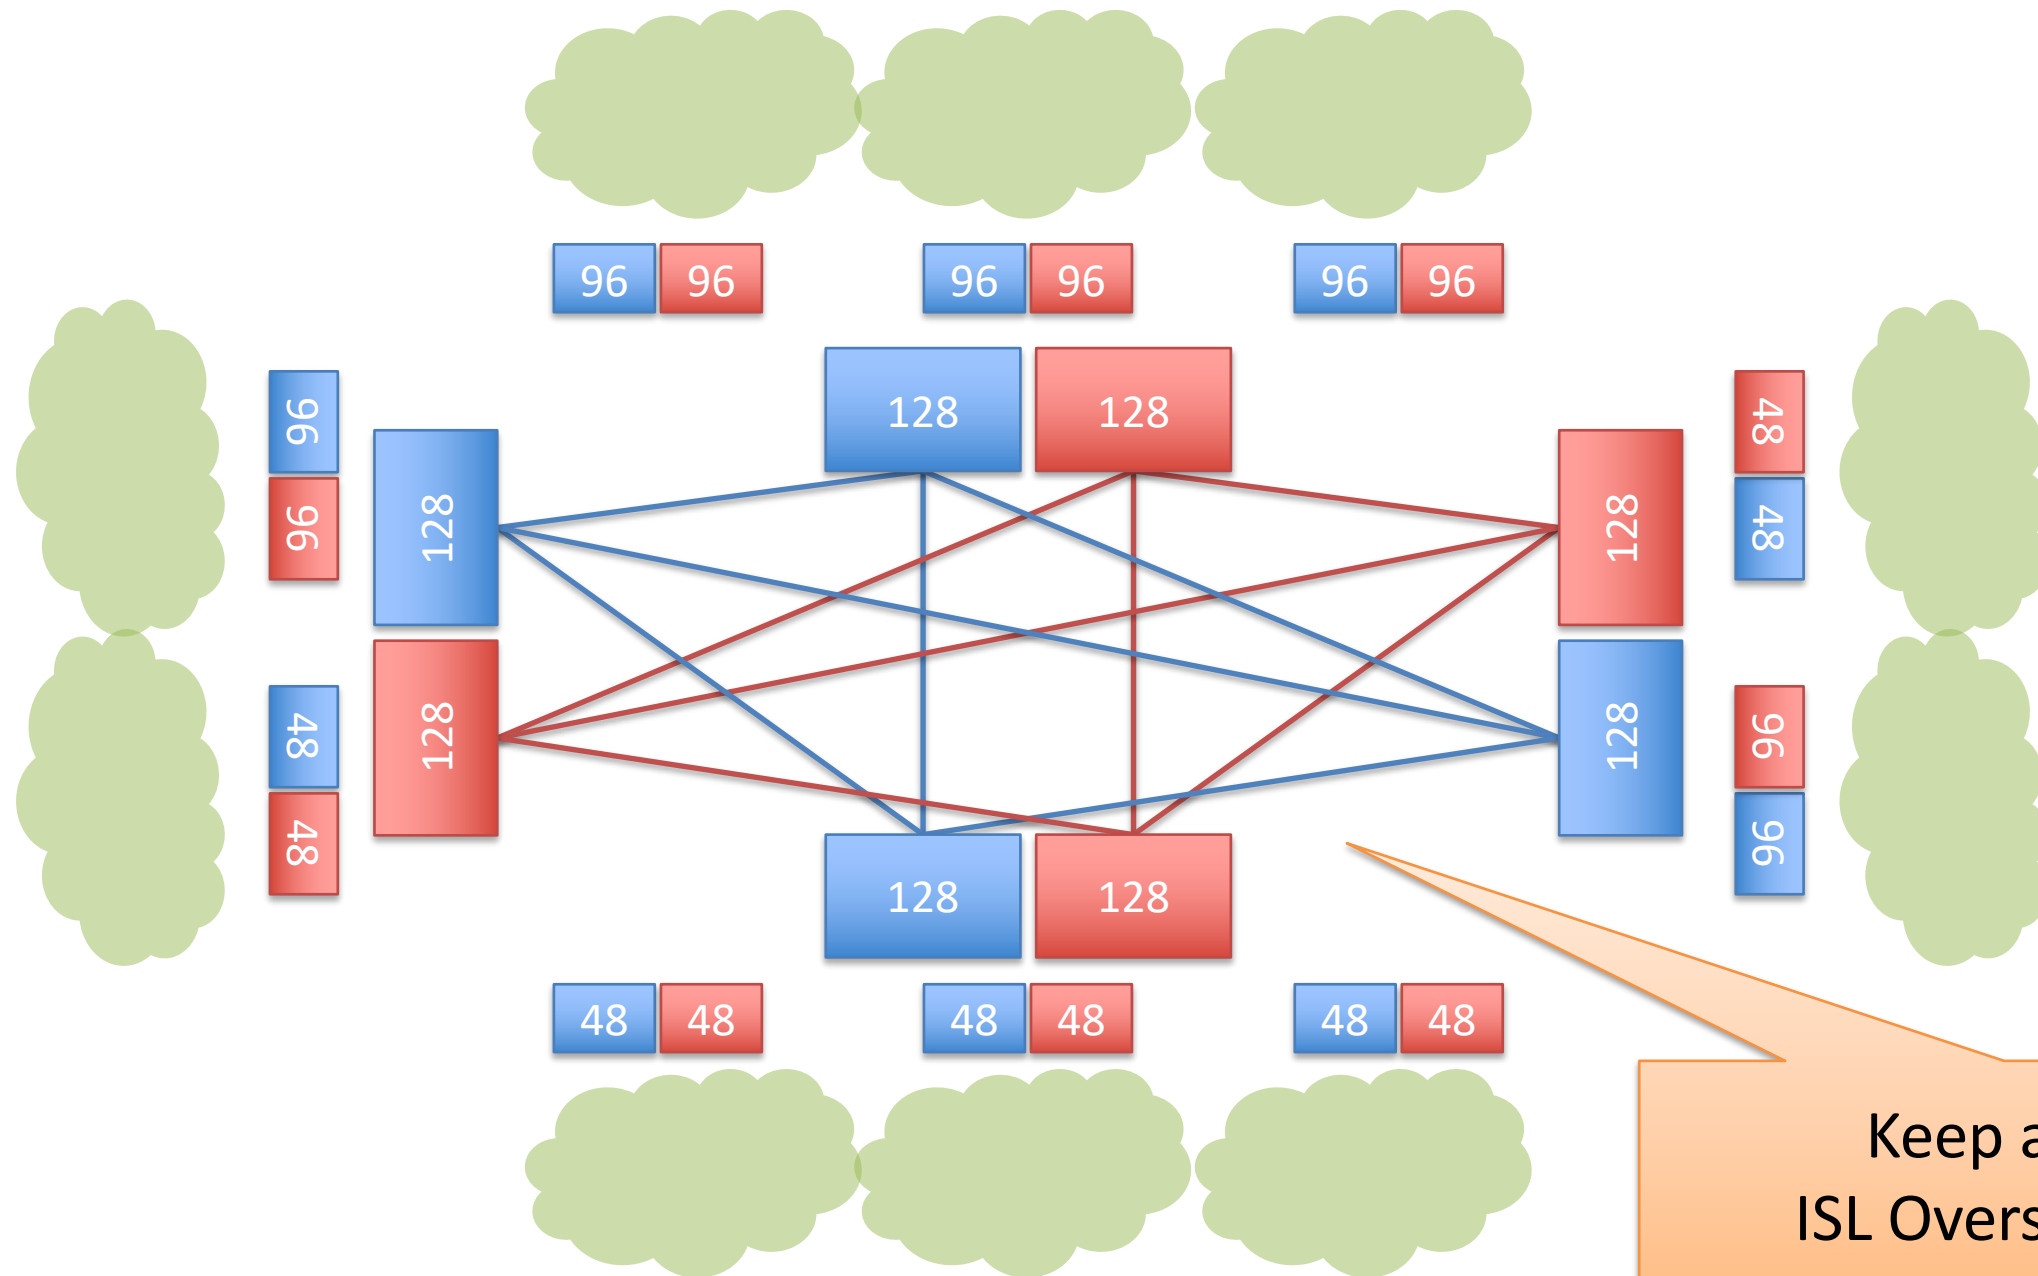
Storage Arrays

## Edge-Core-Edge



# Guaranteed Delivery, Performance, at Scale

Port Count:  
**2,208**




# Fabric Services & Zoning

# Fabric Services

- Discovery Services
  - Fibre Channel provides a name service that ports can discover all of the other ports they have access too.
- Fabric Shortest Path First routing (FSPF)
  - Fabric Shortest Path First (FSPF) is a routing protocol used in Fibre Channel networks. It calculates the best path between switches, establishes routes across the fabric and calculates alternate routes in event of a failure or topology change.
- Zoning
  - Fibre Channel zoning is the partitioning of a Fibre Channel fabric into smaller subsets to restrict interference, add security, and to simplify management.

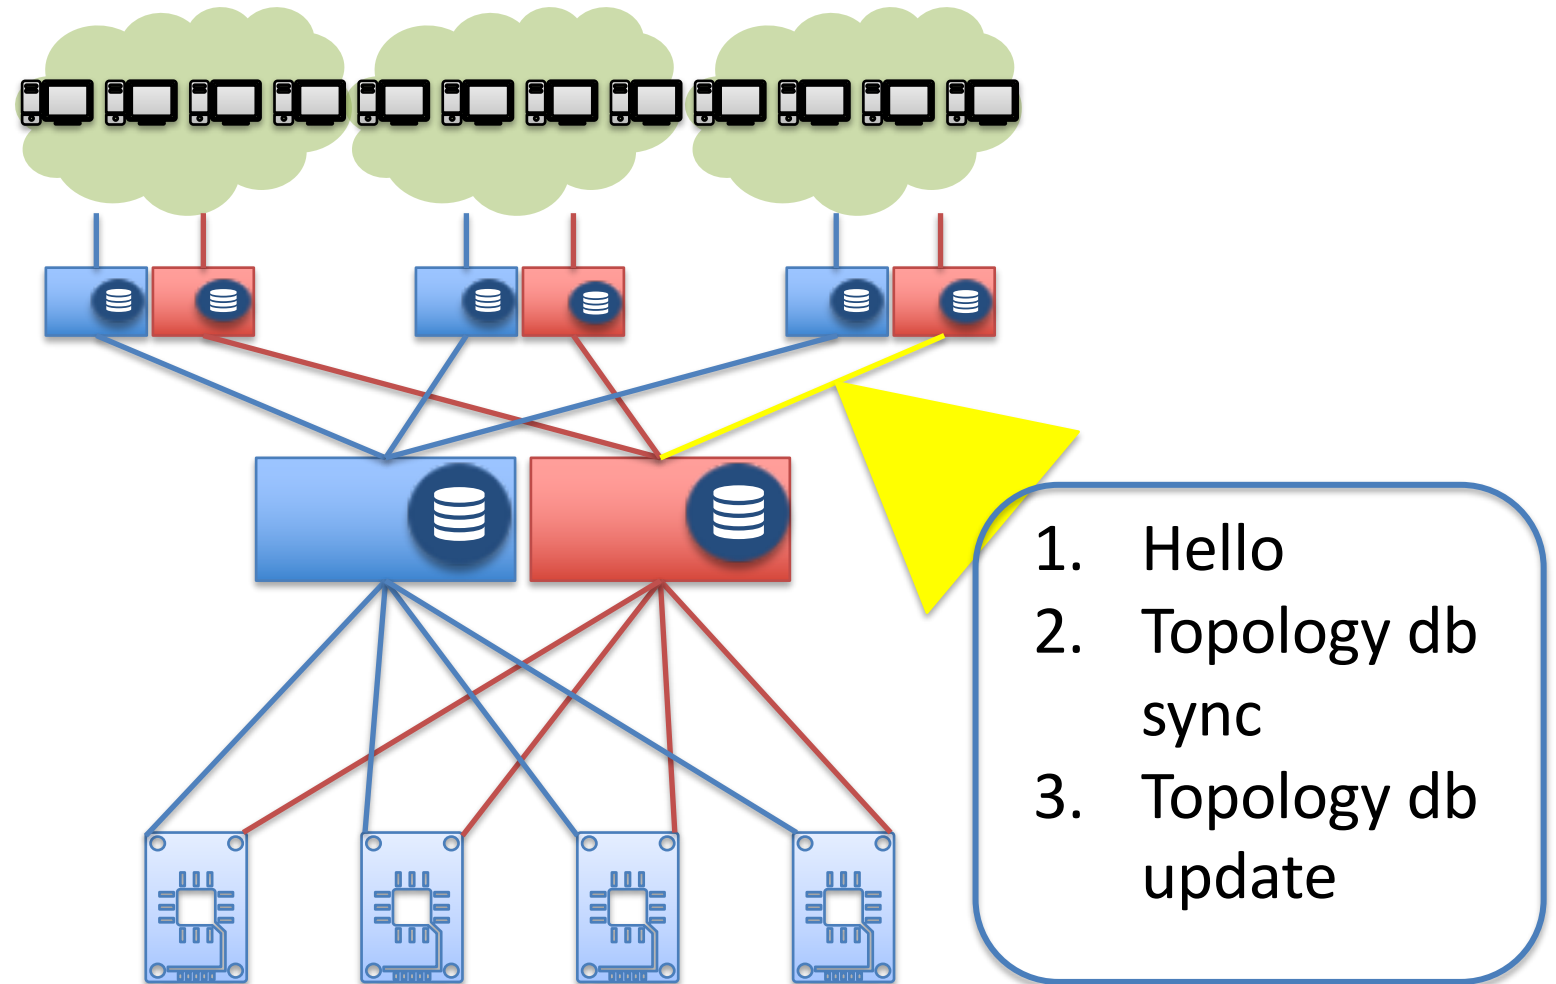
# Name/Discovery Services

- As each port logs in, the port name and location is stored in a database in the Fabric
- Ports can query the Fabric to discover what ports it can see
- Initiators then can scan the targets to discover available storage devices



# FSPF: Protocol


- A Hello protocol to establish two-way communication with a neighbor switch
- Initial topology database synchronization protocol
- Topology database maintenance protocol
- The topology database identifies available paths between ports





# FSPF: Path Selection


- Fibre Channel uses a least cost approach to determine paths to be used to forward frames
  - Weighted by ISL speed
- Frames are delivered in order
- Frames may be load balanced or trunked over ISLs






# FSPF: Path Selection

- Fibre Channel uses a least cost approach to determine paths to be used to forward frames
  - Weighted by ISL speed
- Frames are delivered in order
- Frames may be load balanced or trunked over ISLs
- Backup path also selected



# Zoning

- Zoning allows specific groups of devices to communicate with each other
- Individual zones limit communication between the devices to devices that “care” about each other
- In Fibre Channel switches, the “Fabric Zone Server” controls zoning



# Zoning Terminology

- Zone Set – A collection of zones
- Active Zone Set – The Zone Set currently enforced by the Fabric
- Zone – A “container” with members representing end devices
- Member
  - In a zoneset a **member** is a zone
  - In a zone a **member** represents an end device or group of end devices
- Zone Alias – A name that represents one or more FC devices
- Default Zone
  - Contains all devices not a member of any zone in the active zone set
  - This group of devices may be permitted to communicate or denied
- Basic Zoning Mode
  - Zoning changes done w/o fabric wide lock
  - Lock obtained once changes sent to fabric
  - Less efficient zone data structure
- Enhanced Zoning Mode
  - All zoning changes occur only after a fabric wide lock is obtained. Ends with a commit of changes.
  - More efficient zone data structure



See also: Fibre Channel Zoning Basics on-demand webcast: <http://bit.ly/FCZONING>

Fibre Channel Zoning webcast slides: <http://bit.ly/32nEFoc>

# Deterministic Performance

# Fibre Channel

Great performance in busy data center networks.



# Questions?

# More Great Resources

- **Fibre Channel Performance: Congestion, Slow Drain, and Over-Utilization, Oh My!**
  - Watch: <https://www.brighttalk.com/webcast/14967/295141>
  - Q&A Blog: <https://fibrechannel.org/fibre-channel-performance-qa/>
- **Intro to Incast, Head of Line Blocking and Congestion Management**
  - Watch: <https://www.brighttalk.com/webcast/663/356343>
  - Q&A Blog: <http://sniansfblog.org/storage-congestion-on-the-network-qa/>

# Our Next FCIA Webcast:

## **FC SAN Workloads**

Follow us @FCIANews  
for date and time



# 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 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
  - FCIP (Extension): Data Protection and Business Continuity
  - Fibre Channel Performance
  - FICON
  - Fibre Channel Cabling
  - 64GFC
  - FC Zoning Basics

# Thank You

