

# REALIZE

Dell EMC Integrated Offerings  
Connectrix Business Unit

DELL EMC / World

# Configuring your SAN to Support All Flash Arrays


Alan Rajapa  
Erik Smith

# Agenda

- What's new with storage connectivity
  - Protocol adoption
  - Connectrix Product Updates
- Congestion Spreading and its impact
  - Congestion Spreading terminology
  - Congestion Spreading & Innocent flows
- SAN Best practices for All Flash Arrays (AFAs)
- Introduction to NVMe
  - FC-NVMe

# How Customers Invest in Storage


- FC and IP continue as primary options



Source: IDC WW Quarterly Disk Storage Systems Forecast, 03/17

# How Customers Invest in Storage

- FC and IP continue as primary options



Source: IDC WW Quarterly Disk Storage Systems Forecast, 03/17

# Agenda

- What's new with storage connectivity
  - Protocol adoption
  - Connectrix Product Updates
- Congestion Spreading and its impact
  - Congestion Spreading terminology
  - Congestion Spreading & Innocent flows
- SAN Best practices for All Flash Arrays (AFAs)
- Introduction to NVMe
  - FC-NVMe

# Connectrix B-Series Updates

HW

48 SFP+ Port Blades (4,8,16,32G)  
32- 128Gb/s ICL ports  
FCIP Blade (32Gb/s FC, two 40GiE)

ED-DCX6-8/4B

24 SFP+ Ports (4,8,16,32G)

DS-6610B

Q2  
2016

Q3  
2016

Q1  
2017

Q2  
2017

FW

FOS 8.0.0  
FOS 8.0.1a

FOS 7.4.1d (T)  
FOS 8.0.1b

FOS 7.4.1e  
NOS 6.02c

FOS 8.1.0a  
FOS 8.0.2b

Mgmt

Connectrix  
Manager  
v14.0

Connectrix  
Manager  
v14.0.1

Connectrix  
Manager  
v14.2

# Connectrix B-Series Updates

HW

48 SFP+ Port Blades (4,8,16,32G)  
32- 128Gb/s ICL ports  
FCIP Blade (32Gb/s FC, two 40GiE)

ED-DCX6-8/4B

24 SFP+ Ports (4,8,16,32G)

DS-6610B

Q2  
2016

Q3  
2016

Q1  
2017

Q2  
2017

FW

FOS 8.0.0

FOS 8.0.1a

FOS 7.4.1d (T)

FOS 8.0.1b

FOS 7.4.1e

NOS 6.02c

FOS 8.1.0a

FOS 8.0.2b

Mgmt

Connectrix  
Manager  
v14.0

Connectrix  
Manager  
v14.0.1

Connectrix  
Manager  
v14.2



# Connectrix B-Series Updates

HW

48 SFP+ Port Blades (4,8,16,32G)  
32- 128Gb/s ICL ports  
FCIP Blade (32Gb/s FC, two 40GiE)

ED-DCX6-8/4B

24 SFP+ Ports (4,8,16,32G)

DS-6610B

Q2  
2016

Q3  
2016

Q1  
2017

Q2  
2017

FW

FOS 8.0.0  
FOS 8.0.1a

FOS 7.4.1d (T)  
FOS 8.0.1b

FOS 7.4.1e  
NOS 6.02c

FOS 8.1.0a  
FOS 8.0.2b

Mgmt

Connectrix  
Manager  
v14.0

Connectrix  
Manager  
v14.0.1

Connectrix  
Manager  
v14.2

# Connectrix MDS-Series Updates

HW

48 SFP+ Ports (4,8,16,32G)

DS-9648-1536k9

Q2  
2016

Q3  
2016

Q4  
2016

Q1  
2017

Q2  
2017

FW

NX-OS 6.2.(15)  
NX-OS 7.3(0)D1(1)

NX-OS 7.3(0)DY(1)

NX-OS 6.2.11e (FICON)

NX-OS 7.3(1)DY(1)

Mgmt

DCNM  
v7.2.3

DCNM  
V10.0.1

DCNM  
V10.1.1

DCNM  
V10.1.2

# Connectrix MDS-Series Updates

HW

48 SFP+ Ports (4,8,16,32G)

DS-9648-1536k9

Q2  
2016

Q3  
2016

Q4  
2016

Q1  
2017

Q2  
2017

FW

NX-OS 6.2.(15)  
NX-OS 7.3(0)D1(1)

NX-OS 7.3(0)DY(1)

NX-OS 6.2.11e (FICON)

NX-OS 7.3(1)DY(1)

Mgmt

DCNM  
v7.2.3

DCNM  
V10.0.1

DCNM  
V10.1.1

DCNM  
V10.1.2

Connectrix  
32 Gb/s  
Switches and Directors  
Now available!

# Agenda

- What's new with storage connectivity
  - Protocol adoption
  - Connectrix Product Updates
- **Congestion Spreading and its impact**
  - Congestion Spreading terminology
  - Congestion Spreading & Innocent flows
- SAN Best practices for All Flash Arrays (AFAs)
- Introduction to NVMe
  - FC-NVMe

# Congestion Spreading Terminology

# Congestion

- Oversubscription
  - Bandwidth mismatch
    - › Speed mismatch (e.g. 16G to 4G)
    - › Fan-in mismatch (e.g. 8-port to 1-port)

# Congestion

- Oversubscription
  - Bandwidth mismatch
    - › Speed mismatch (e.g. 16G to 4G)
    - › Fan-in mismatch (e.g. 8-port to 1-port)
- Misbehaving Devices (Slow Drain)
  - **Brocade**: Misbehaving Devices
    - › Credit stall (i.e. device induced credit latency)
  - **Cisco**: B2B credit starvation
    - › Slow Port
    - › Stuck Port

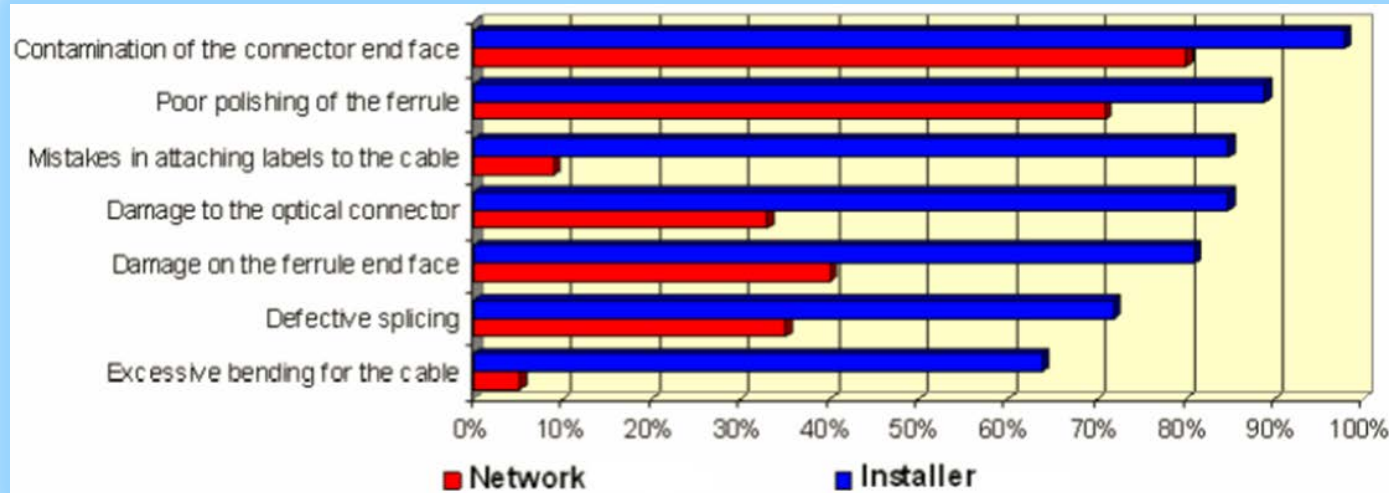


# Congestion

- Oversubscription
  - Bandwidth mismatch
    - › Speed mismatch (e.g. 16G to 4G)
    - › Fan-in mismatch (e.g. 8-port to 1-port)
- Misbehaving Devices (Slow Drain)
  - **Brocade**: Misbehaving Devices
    - › Credit stall (i.e. device induced credit latency)
  - **Cisco**: B2B credit starvation
    - › Slow Port
    - › Stuck Port
- Lost Credit
  - Physical errors

# Physical Errors

- 98% of network installation professionals reported connector/cable contamination was the greatest cause of network failure (Source: NTT-Advanced Technology)
- 83% of replaced optics were not actually faulty (Source: Dell EMC)



John Ford  
[john.ford@dell.com](mailto:john.ford@dell.com)

Tuesday at 4:10pm and Wednesday at  
12:50pm @ Brocade Booth

Wednesday at 4-5pm @ Customer  
Service Booth

# Metrics

- Credit Latency
  - Time at zero credit
  - Amount of time frames are waiting to be transmitted
    - › Credit Latency counters track the cumulative time frames wait for credit when ready for transmission
    - › Credit Latency monitors typically look for average delay times over varying periods to capture delay spikes, repetitive delays, or sustained delays

# Metrics

- Credit Latency
  - Time at zero credit
  - Amount of time frames are waiting to be transmitted
    - › Credit Latency counters track the cumulative time frames wait for credit when ready for transmission
    - › Credit Latency monitors typically look for average delay times over varying periods to capture delay spikes, repetitive delays, or sustained delays
- Brocade: Queue Latency
  - Transmit queue latency
  - Amount of time a frame is queued for transmission until it is transmitted
    - › i.e. the time a frames takes to move from the bottom to the top of the transmit queue plus transmission time

# Metrics

- Credit Latency
  - Time at zero credit
  - Amount of time frames are waiting to be transmitted
    - › Credit Latency counters track the cumulative time frames wait for credit when ready for transmission
    - › Credit Latency monitors typically look for average delay times over varying periods to capture delay spikes, repetitive delays, or sustained delays
- **Brocade:** Queue Latency
  - Transmit queue latency
  - Amount of time a frame is queued for transmission until it is transmitted
    - › i.e. the time a frames takes to move from the bottom to the top of the transmit queue plus transmission time
- Bandwidth Utilization
  - Percent utilization
  - Average amount of link capacity used

# Metrics

- Credit Latency
  - Time at zero credit
  - Amount of time frames are waiting to be transmitted
    - › Credit Latency counters track the cumulative time frames wait for credit when ready for transmission
    - › Credit Latency monitors typically look for average delay times over varying periods to capture delay spikes, repetitive delays, or sustained delays
- **Brocade:** Queue Latency
  - Transmit queue latency
  - Amount of time a frame is queued for transmission until it is transmitted
    - › i.e. the time a frames takes to move from the bottom to the top of the transmit queue plus transmission time
- Bandwidth Utilization
  - Percent utilization
  - Average amount of link capacity used
- Frame Loss
  - Discarded frame count
  - **Brocade:** Count of frame discarded due to transmission hold time timeout (220ms to 500ms)
  - **Cisco:** Frames discard due to congestion drop (500ms) or no-credit drop timeout

# Metrics

- Credit Latency
  - Time at zero credit
  - Amount of time frames are waiting to be transmitted
    - › Credit Latency counters track the cumulative time frames wait for credit when ready for transmission
    - › Credit Latency monitors typically look for average delay times over varying periods to capture delay spikes, repetitive delays, or sustained delays
- Brocade: Queue Latency
  - Transmit queue latency
  - Amount of time a frame is queued for transmission until it is transmitted
    - › i.e. the time a frames takes to move from the bottom to the top of the transmit queue plus transmission time
- Bandwidth Utilization
  - Percent utilization
  - Average amount of link capacity used
- Frame Loss
  - Discarded frame count
  - Brocade: Count of frame discarded due to transmission hold time timeout (220ms to 500ms)
  - Cisco: Frames discard due to congestion drop (500ms) or no-credit drop timeout
- Link Reset
  - Link reset count
  - Brocade: Number of times credit has not been returned for 2-seconds
  - Cisco: Number of times when credits were zero for 1 second on F\_port and 1.5 seconds on E\_port



# Brocade Severity

- Mild
  - Small credit delay
  - Brocade only Small queue latency (less than 10ms)
  - No frame loss (discards) or link resets

# Brocade Severity

- Mild
  - Small credit delay
  - Brocade only Small queue latency (less than 10ms)
  - No frame loss (discards) or link resets
- Moderate
  - Medium credit delay
  - Brocade only Medium queue latency (10ms – 80ms)
  - No frame loss (discards), but no link resets

# Brocade Severity

- Mild
  - Small credit delay
  - Brocade only Small queue latency (less than 10ms)
  - No frame loss (discards) or link resets
- Moderate
  - Medium credit delay
  - Brocade only Medium queue latency (10ms – 80ms)
  - No frame loss (discards), but no link resets
- Severe
  - Large credit delay
  - Brocade only Large queue latency (greater than 80ms)
  - Frame loss (discards) and some link resets

# Cisco Severity

- Level – 1 : Latency
  - Reduced number of remaining credits or small duration of credit unavailability
  - No discards, retransmission or link resets

\* Default configuration: congestion-drop timeout – 500ms, no-credit-drop timeout – off  
Configurable option: congestion-drop timeout – 100 - 500ms, no-credit-drop timeout – 1 – 500 ms  
Recommended configuration: congestion-drop timeout – 200ms, no-credit-drop timeout – 50 ms

# Cisco Severity

- Level – 1 : Latency
  - Reduced number of remaining credits or small duration of credit unavailability
  - No discards, retransmission or link resets
- Level – 2 : Retransmission
  - Longer duration of credit unavailability
  - Frames are discarded (but no link reset) due to congestion-drop timeout or no-credit-drop timeout\* leading to retransmission.

\* Default configuration: congestion-drop timeout – 500ms, no-credit-drop timeout – off  
Configurable option: congestion-drop timeout – 100 - 500ms, no-credit-drop timeout – 1 – 500 ms  
Recommended configuration: congestion-drop timeout – 200ms, no-credit-drop timeout – 50 ms

# Cisco Severity

- Level – 1 : Latency
  - Reduced number of remaining credits or small duration of credit unavailability
  - No discards, retransmission or link resets
- Level – 2 : Retransmission
  - Longer duration of credit unavailability
  - Frames are discarded (but no link reset) due to congestion-drop timeout or no-credit-drop timeout\* leading to retransmission.
- Level – 3 : Extreme delay
  - Prolonged duration of credit unavailability (1 sec for F-port, 1.5 sec for E-port)
  - Link resets or port flaps

\* Default configuration: congestion-drop timeout – 500ms, no-credit-drop timeout – off

Configurable option: congestion-drop timeout – 100 - 500ms, no-credit-drop timeout – 1 – 500 ms

Recommended configuration: congestion-drop timeout – 200ms, no-credit-drop timeout – 50 ms

# Brocade Terminology Cross Reference

Cause	Mild	Moderate	Severe
Oversubscription <sup>1</sup>	<ol style="list-style-type: none"> <li>1. High Bandwidth at the device port.</li> <li>2. Small Credit Latency at the ISL port.</li> <li>3. Less than 10ms Queue Latency at the ISL port.</li> <li>4. No Frame Loss or Link Resets.</li> </ol>	<ol style="list-style-type: none"> <li>1. High Bandwidth at the device port.</li> <li>2. Medium Credit Latency at the ISL port.</li> <li>3. Between 10m to 80ms Queue Latency at the ISL port.</li> <li>4. No Frame Loss or Link Resets.</li> </ol>	<ol style="list-style-type: none"> <li>1. High Bandwidth at the device port.</li> <li>2. Large Credit Latency at the ISL port.</li> <li>3. Greater than 80ms Queue Latency at the ISL port.</li> <li>4. Frame Loss at an upstream (ISL) port (indicates Queue Latency of 220ms-500ms).</li> <li>5. No Link Resets.</li> </ol>
Misbehaving Device	<ol style="list-style-type: none"> <li>1. Small Credit Latency at the device port and upstream ISL port.</li> <li>2. Less than 10ms Queue Latency at the device port and upstream ISL port.</li> <li>3. No Frame Loss or Link Resets.</li> </ol>	<ol style="list-style-type: none"> <li>1. Medium Credit Latency at the device port and upstream ISL port.</li> <li>2. Between 10ms to 80ms Queue Latency at the device port and upstream ISL port.</li> <li>3. No Frame Loss or Link Resets.</li> </ol>	<ol style="list-style-type: none"> <li>1. Large Credit Latency at the device port and upstream ISL port.</li> <li>2. Greater than 80ms Queue Latency at the device port and upstream ISL port.</li> <li>3. Frame Loss at device or upstream (ISL) port (indicates Queue Latency of 220ms-500ms).</li> <li>4. Link Reset at an ISL port (indicates credit stall for more than 2s).</li> </ol>
Lost Credit <sup>2</sup>	<ol style="list-style-type: none"> <li>1. Small Credit Latency at the port.</li> <li>2. Less than 10ms Queue Latency at the port or upstream from the port.</li> <li>3. No Frame Loss or Link Resets.</li> </ol>	<ol style="list-style-type: none"> <li>1. Medium Credit Latency at the port.</li> <li>2. Between 10ms to 80ms Queue Latency at the port or upstream from the port.</li> <li>3. No Frame Loss or Link Resets.</li> </ol>	<ol style="list-style-type: none"> <li>1. Large Credit Latency at the port.</li> <li>2. Greater than 80ms Queue Latency at the port or upstream from the port.</li> <li>3. Frame Loss at the port or upstream from the port (indicates credit stall for 220ms-500ms).</li> <li>4. Link Reset at the port or upstream from the port (indicates credit stall for more than 2s).</li> </ol>

<sup>1</sup> Severe congestion due to oversubscription is a rare to extremely rare occurrence.

<sup>2</sup> Causes for Lost Credit are typically transmission errors such as ITW, CRC, or other signal related problems.

# Cisco Terminology Cross Reference

Congestion type	Level – 1 : Latency	Level – 2 : Retransmission	Level – 3 : Extreme delay
Oversubscription	<ol style="list-style-type: none"> <li>1. High link utilization at the end-device port.</li> <li>2. No B2B credit starvation at the end-device port</li> <li>3. Congestion spreading towards the ISLs</li> <li>4. No Frame Loss or Link Resets.</li> </ol>	Retransmission or Extreme delay due to oversubscription is a rare to extremely rare occurrence.	
Credit Starvation	<ol style="list-style-type: none"> <li>1. Low link utilization at the end-device port</li> <li>2. Reduced number of remaining credits or small duration of credit unavailability</li> <li>3. Congestion spreading towards ISLs</li> <li>4. No discards, retransmission or link resets</li> </ol>	<ol style="list-style-type: none"> <li>1. Low link utilization at the end-device port</li> <li>2. Longer duration of credit unavailability</li> <li>3. Congestion spreading towards ISLs.</li> <li>4. Frames are discarded (but no link reset) due to congestion-drop timeout or no-credit-drop timeout* leading to retransmission.</li> </ol>	<ol style="list-style-type: none"> <li>1. No frames are transmitted to the end-device.</li> <li>2. Prolonged duration of credit unavailability (1 sec for F-port, 1.5 sec for E-port)</li> <li>3. Severe congestion towards ISLs</li> <li>4. Link resets or port flaps</li> </ol>

\* Default configuration: congestion-drop timeout – 500ms, no-credit-drop timeout – off  
 Configurable option: congestion-drop timeout – 100 - 500ms, no-credit-drop timeout – 1 – 500 ms  
 Recommended configuration: congestion-drop timeout – 200ms, no-credit-drop timeout – 50 ms




# Agenda


- What's new with storage connectivity
  - Protocol adoption
  - Connectrix Product Updates
- **Congestion Spreading and its impact**
  - Congestion Spreading terminology
  - Congestion Spreading & Innocent flows
- SAN Best practices for All Flash Arrays (AFAs)
- Introduction to NVMe
  - FC-NVMe

# Congestion spreading & Innocent flows


# Congestion Spreading




# Congestion Spreading




# Congestion Spreading



# Congestion Spreading




# Congestion Spreading



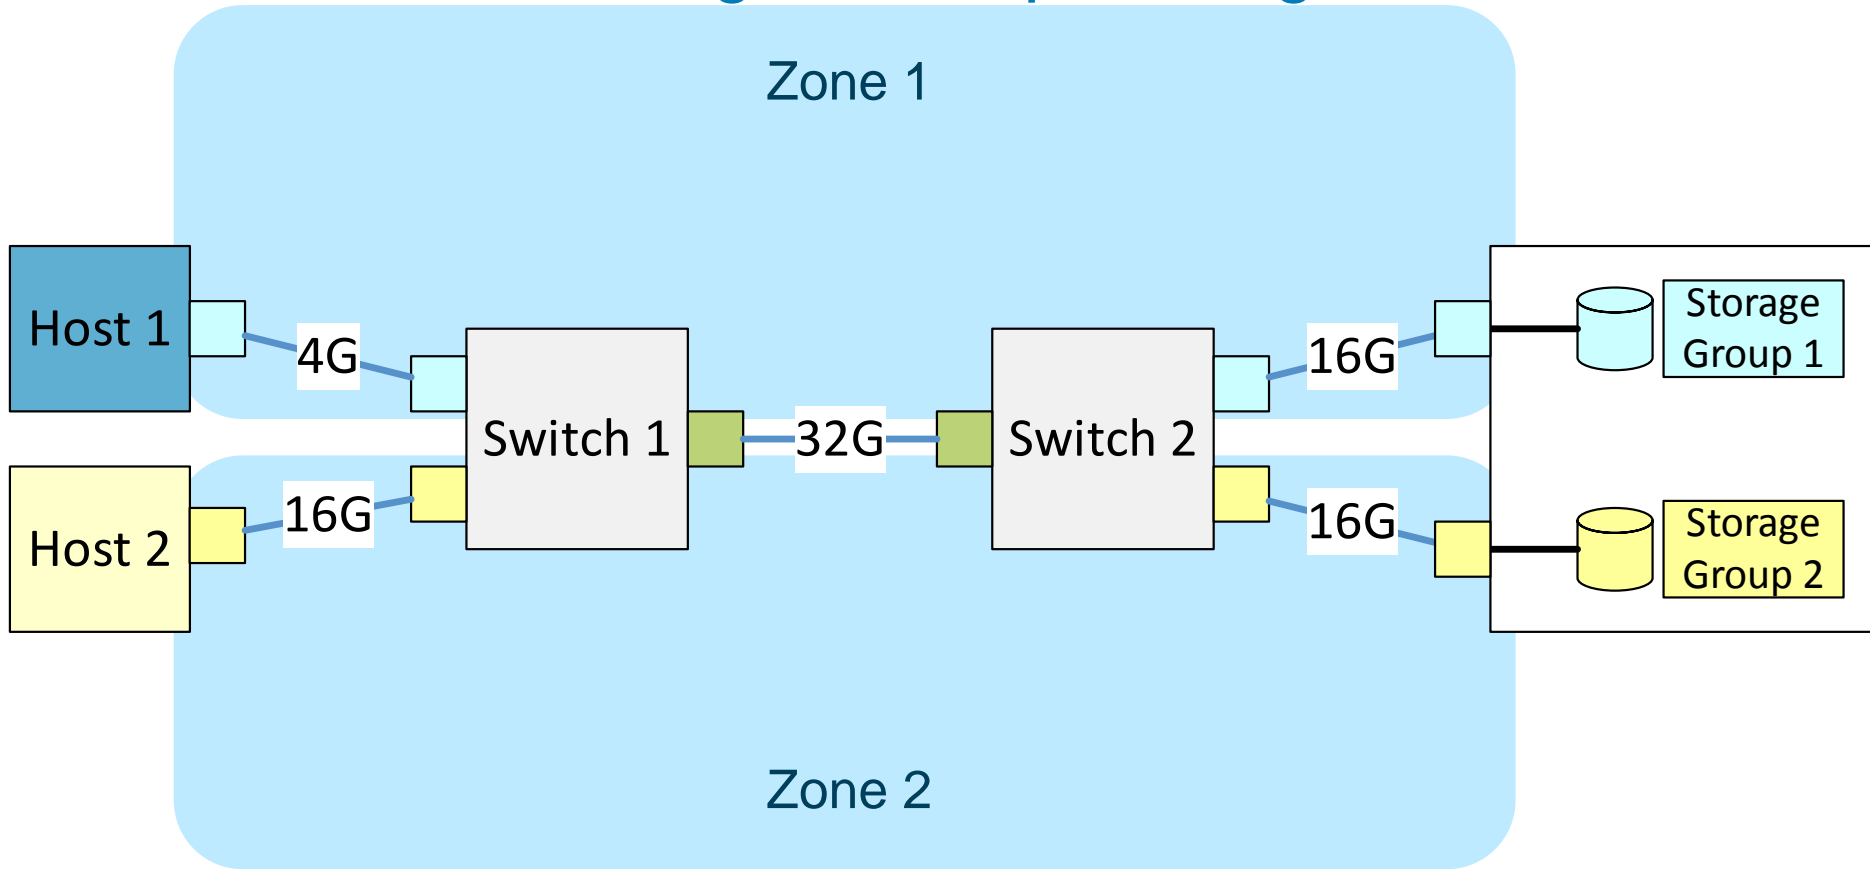
# Congestion Spreading

Zone 1

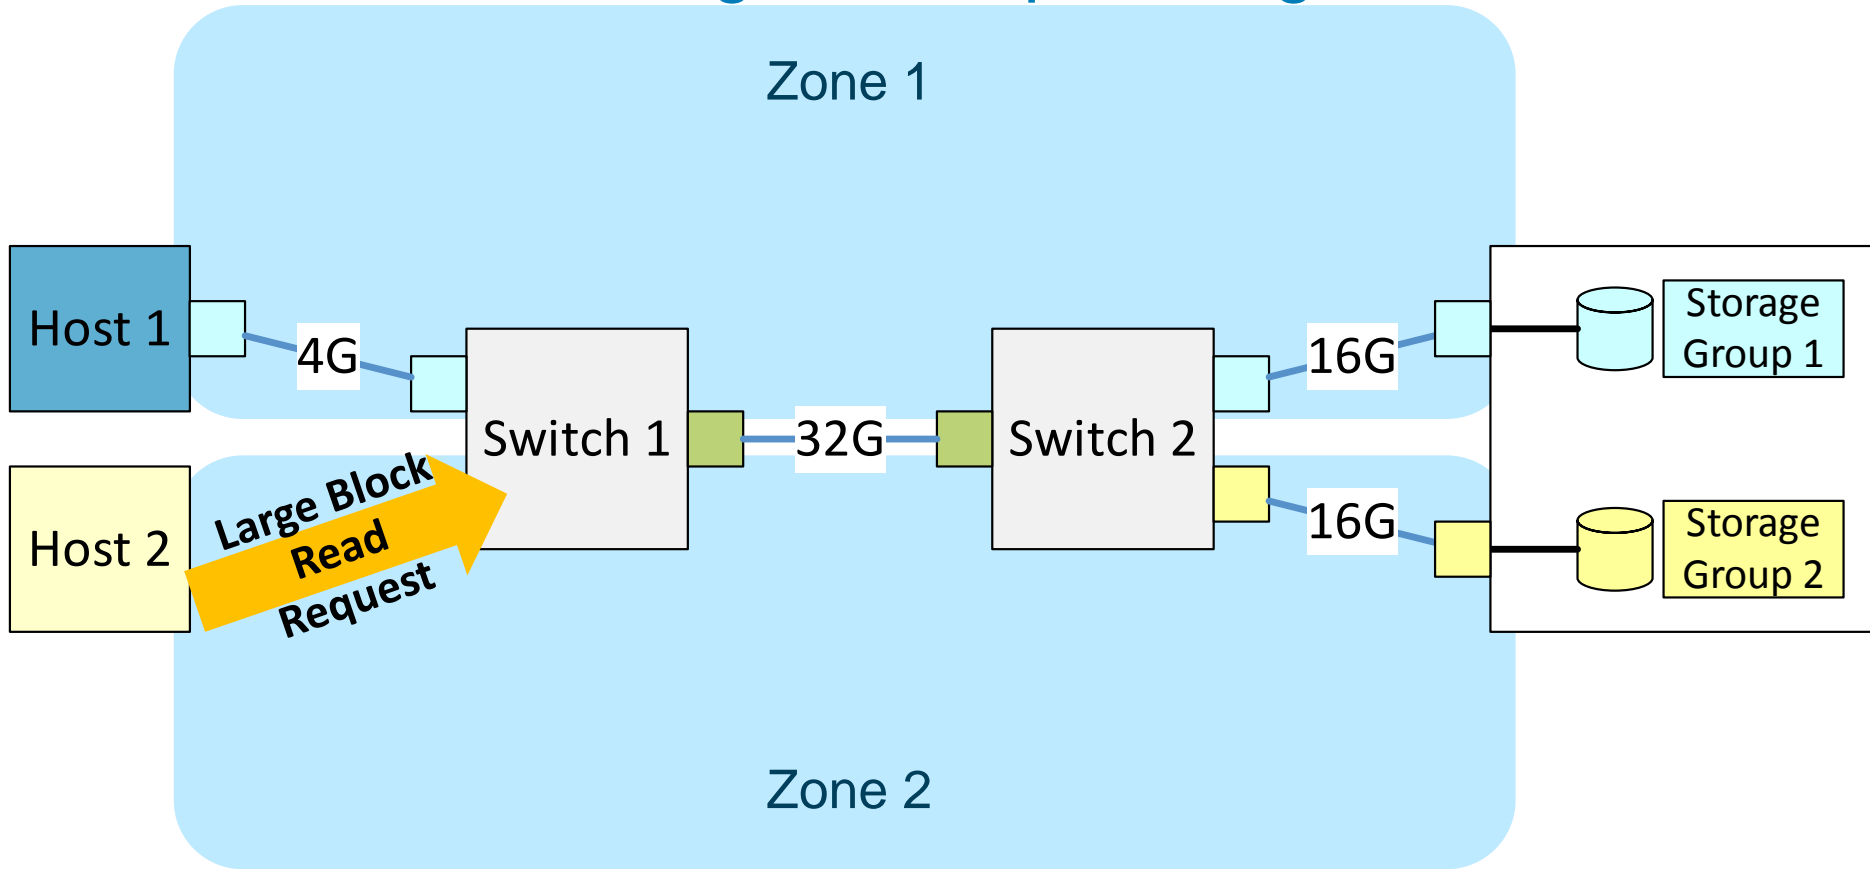




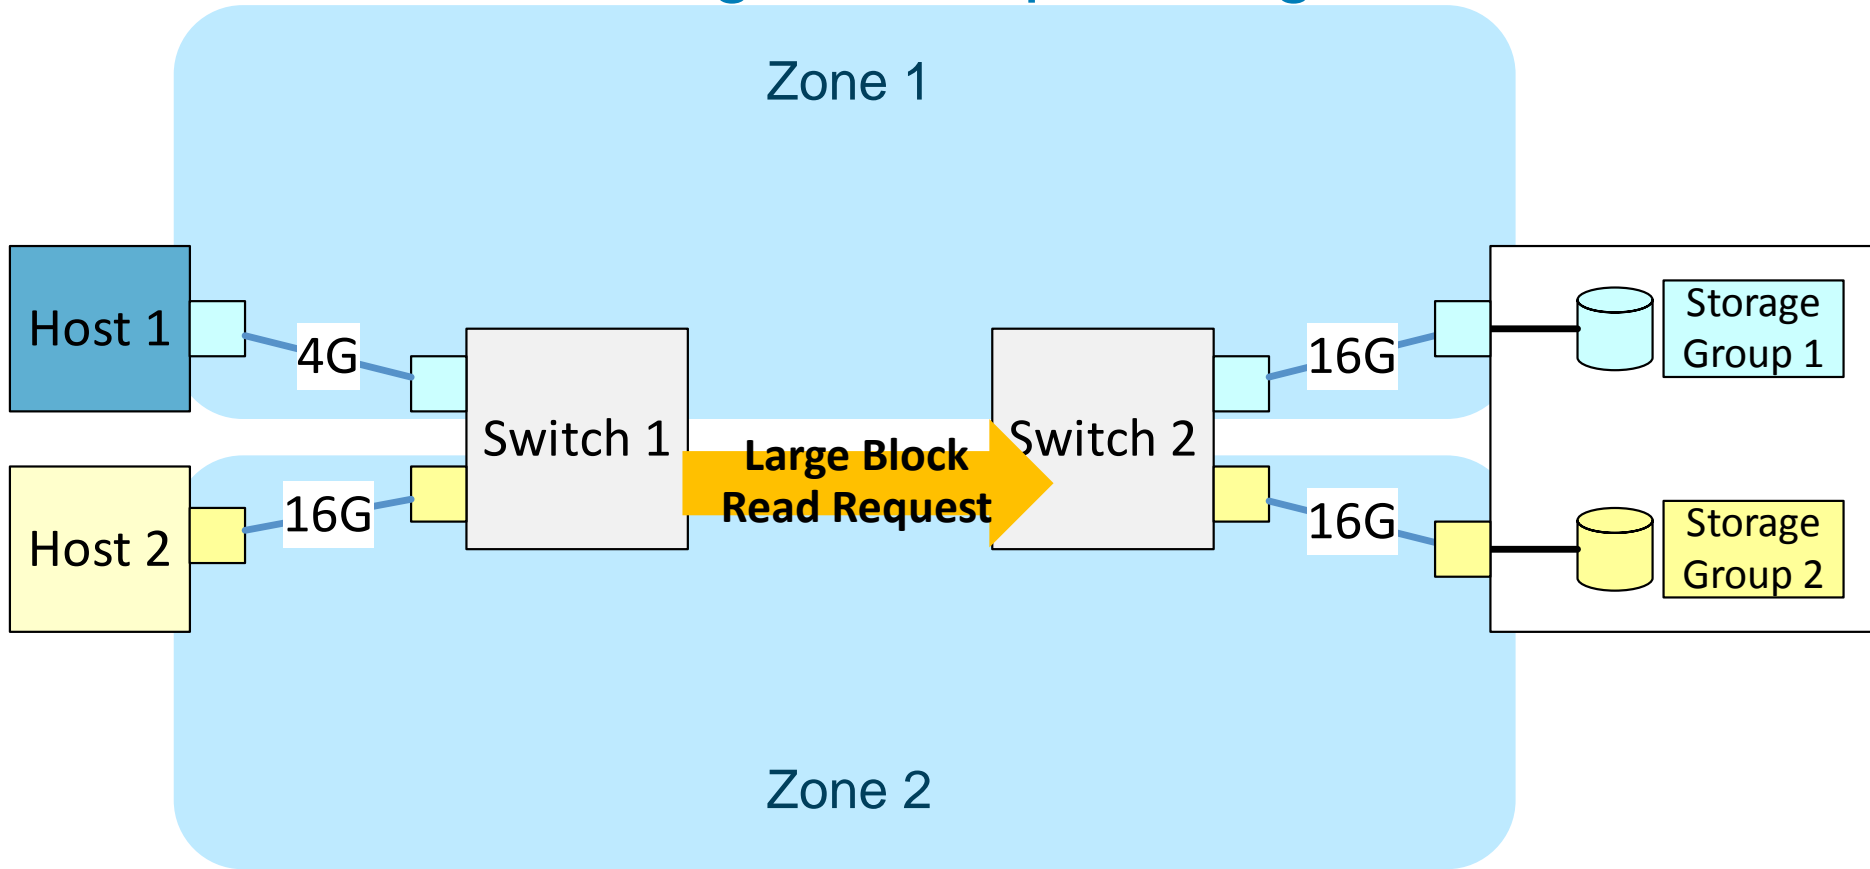
# Congestion Spreading



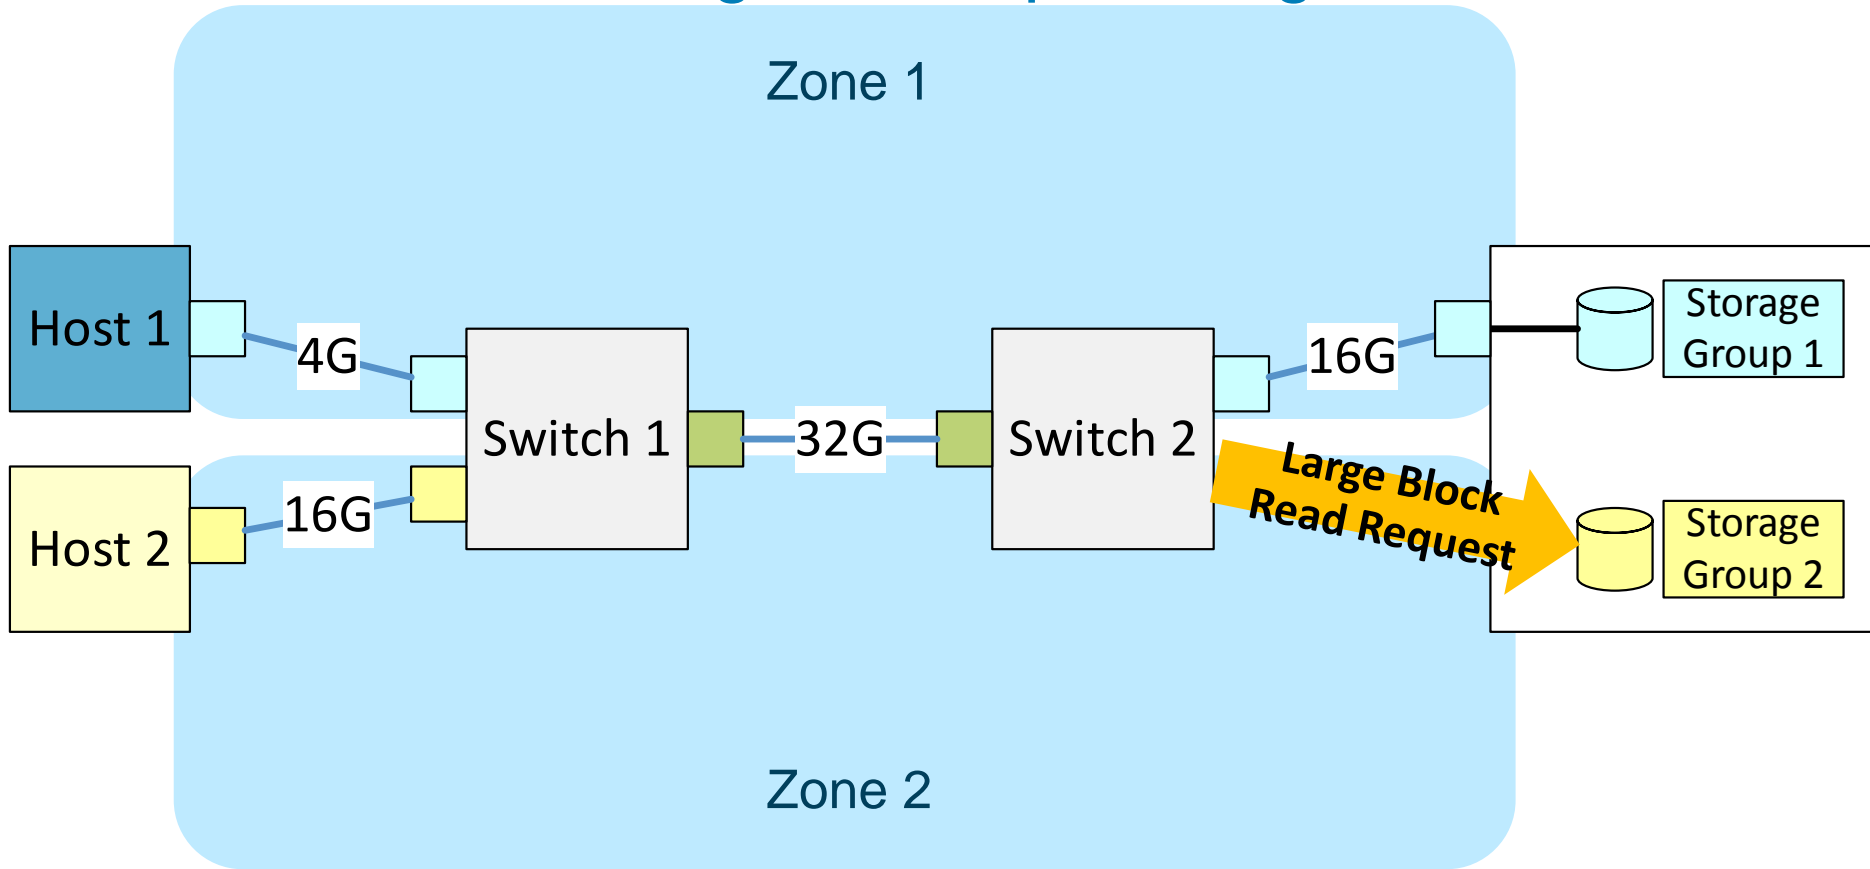
# Congestion Spreading



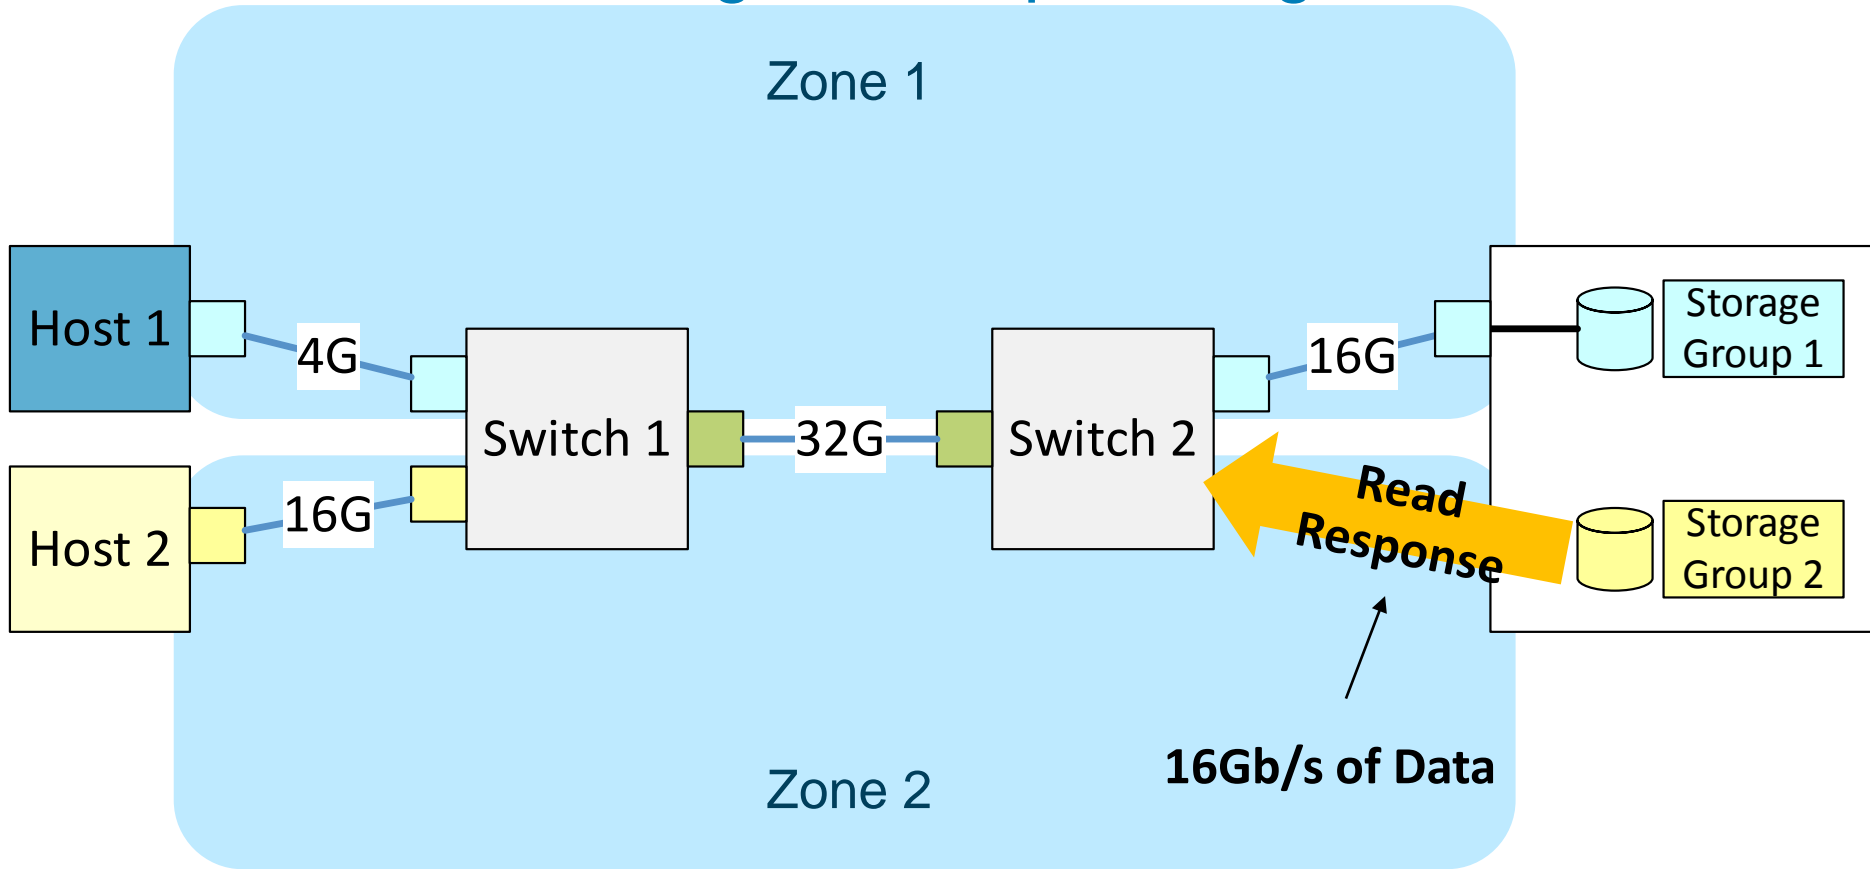
# Congestion Spreading



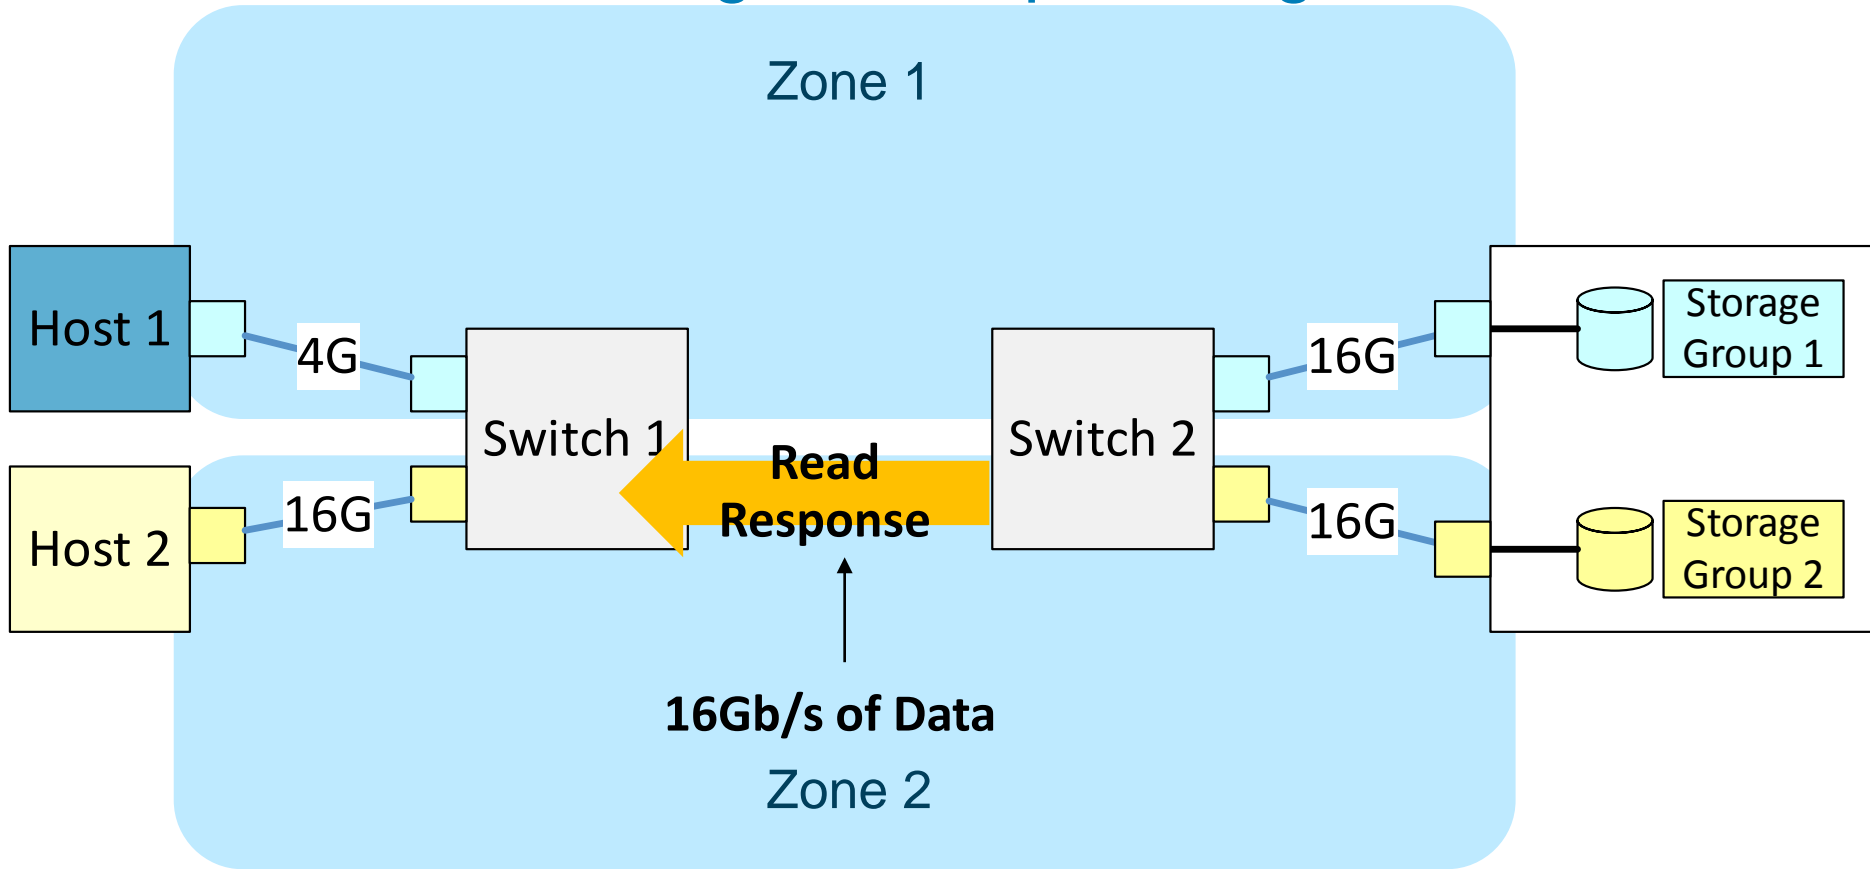
# Congestion Spreading



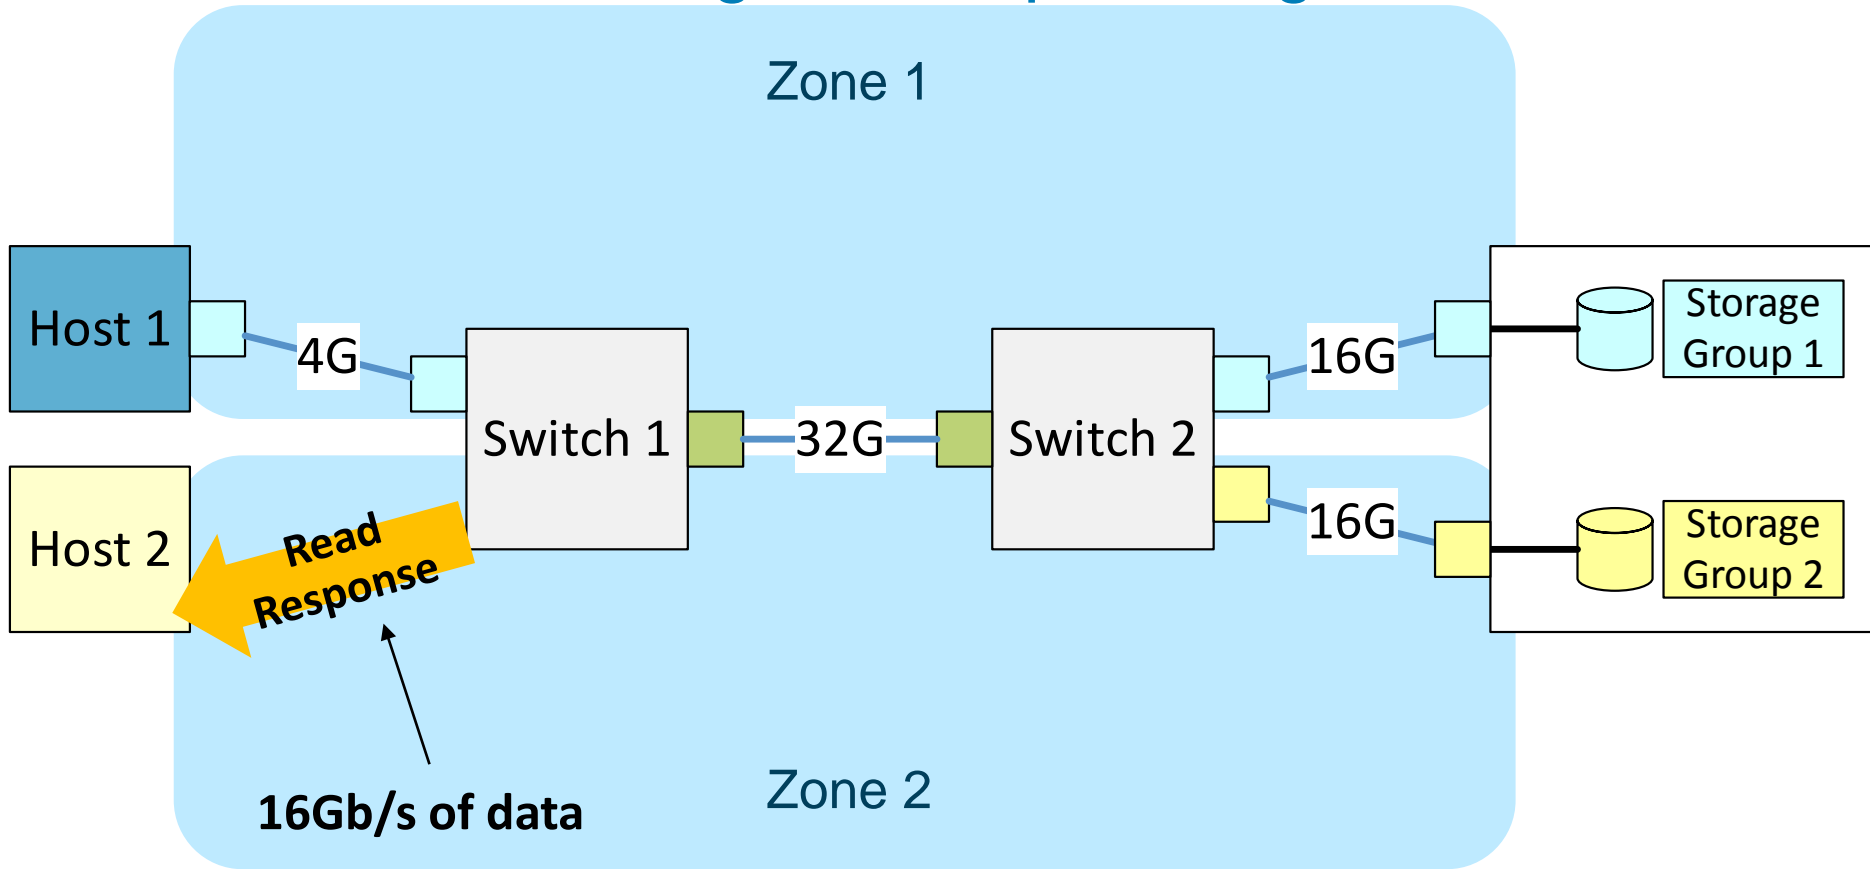
# Congestion Spreading



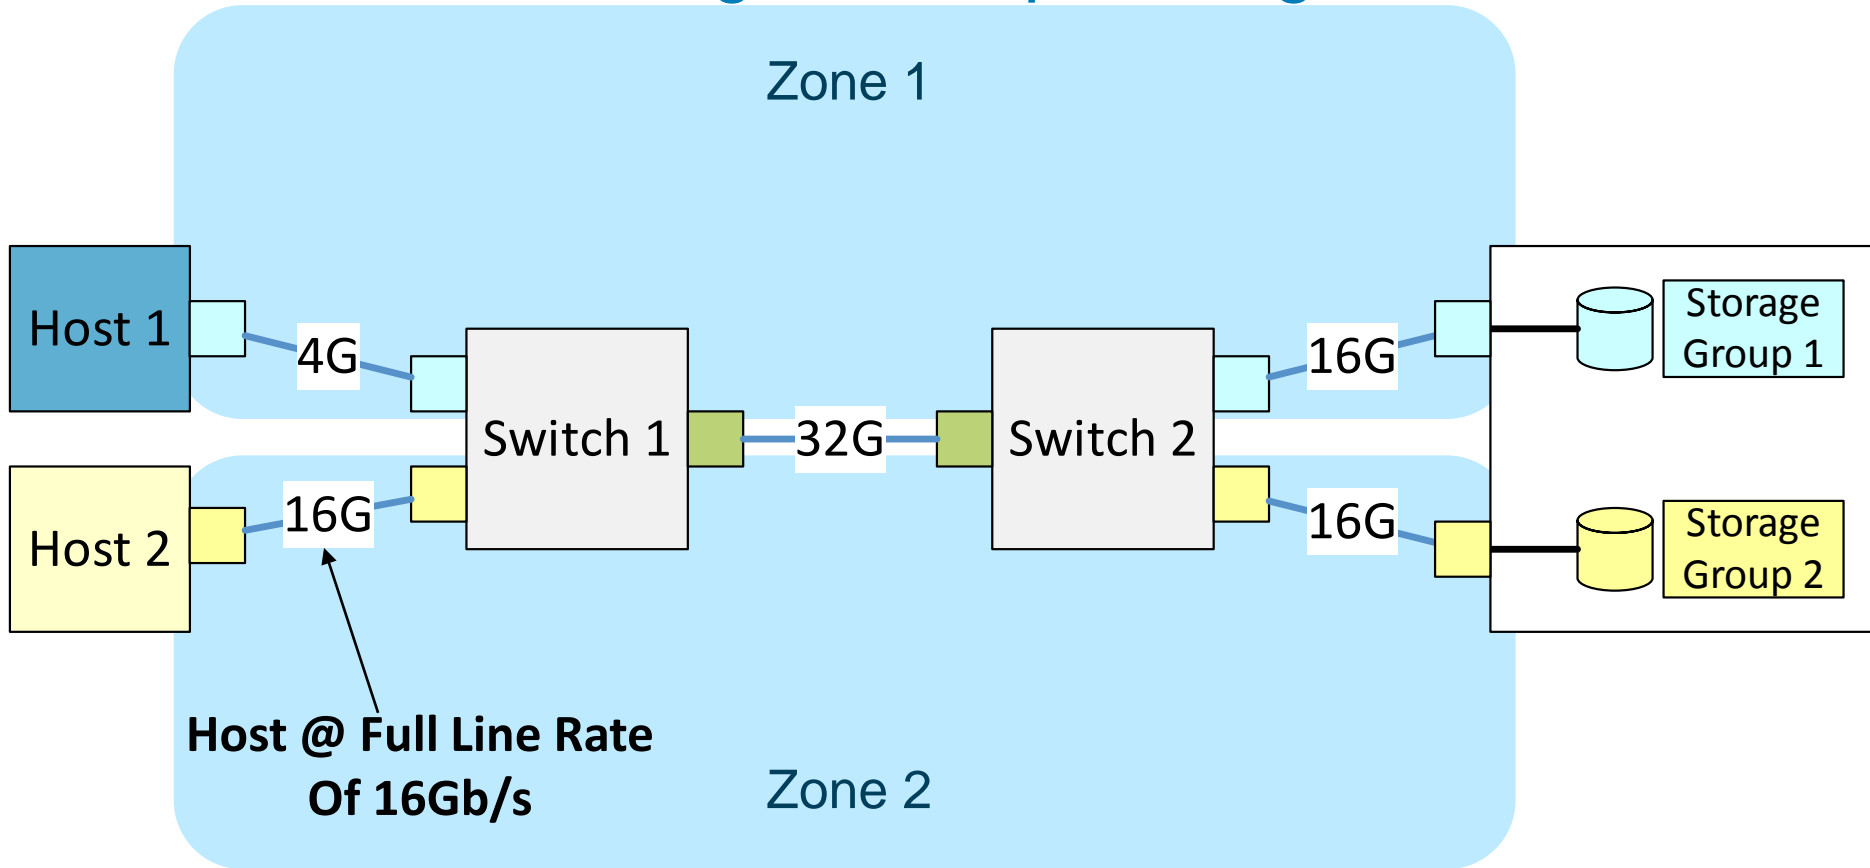
# Congestion Spreading



# Congestion Spreading

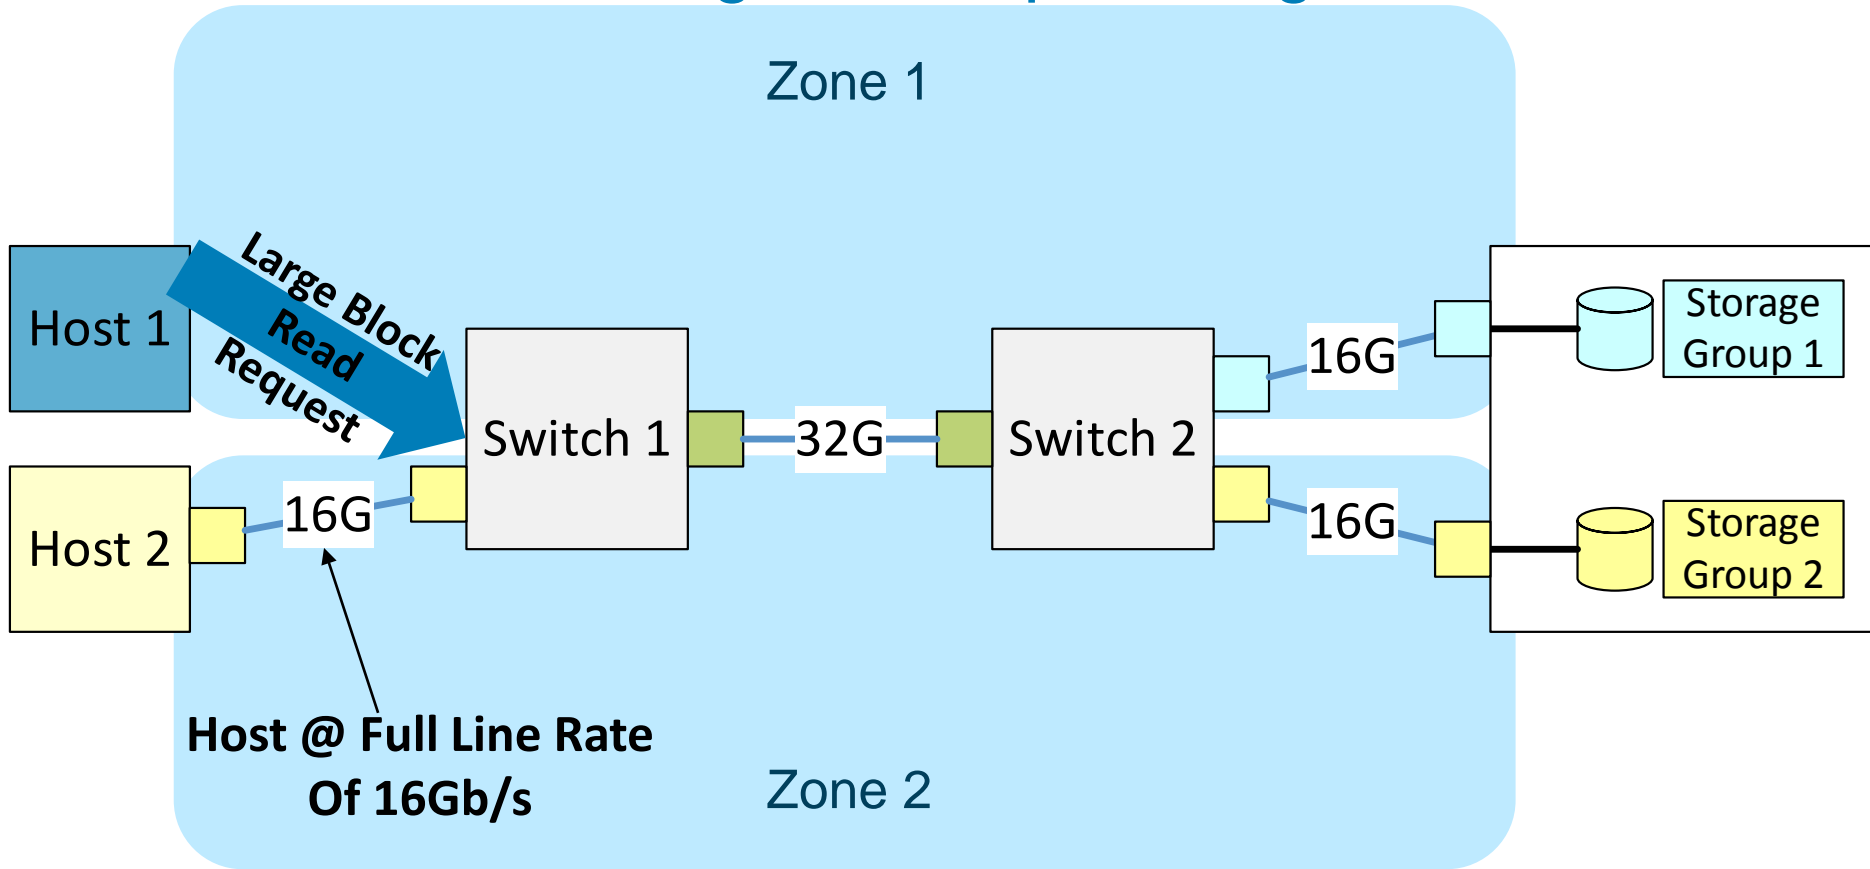


# Congestion Spreading

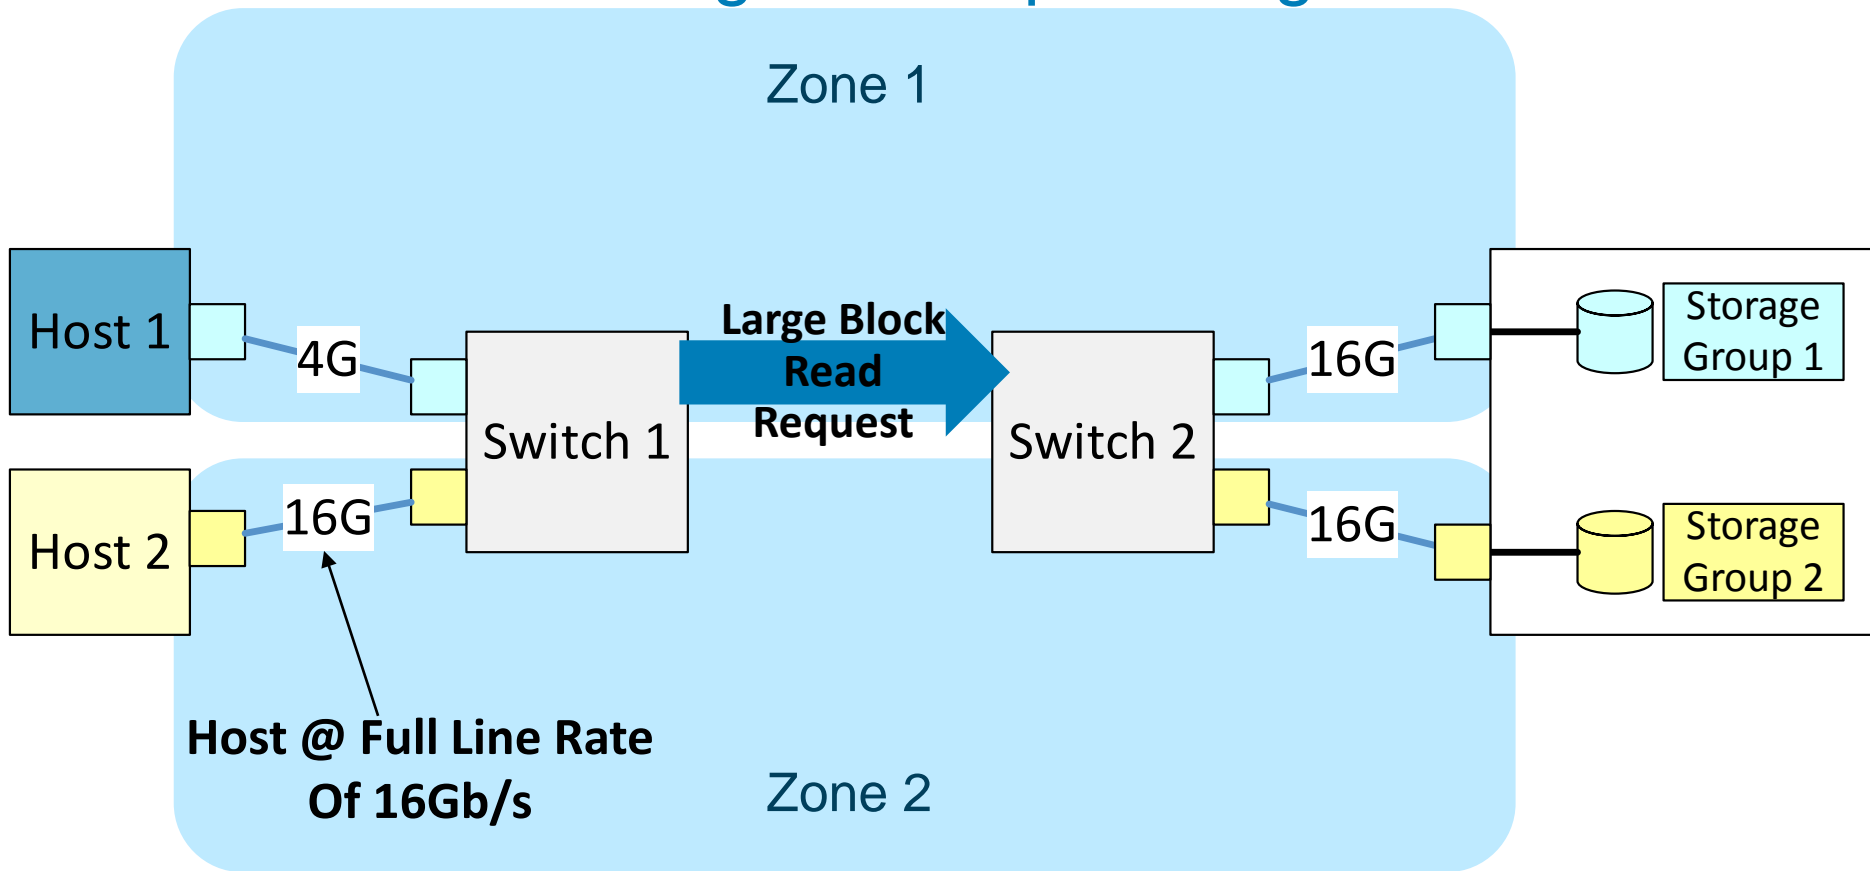




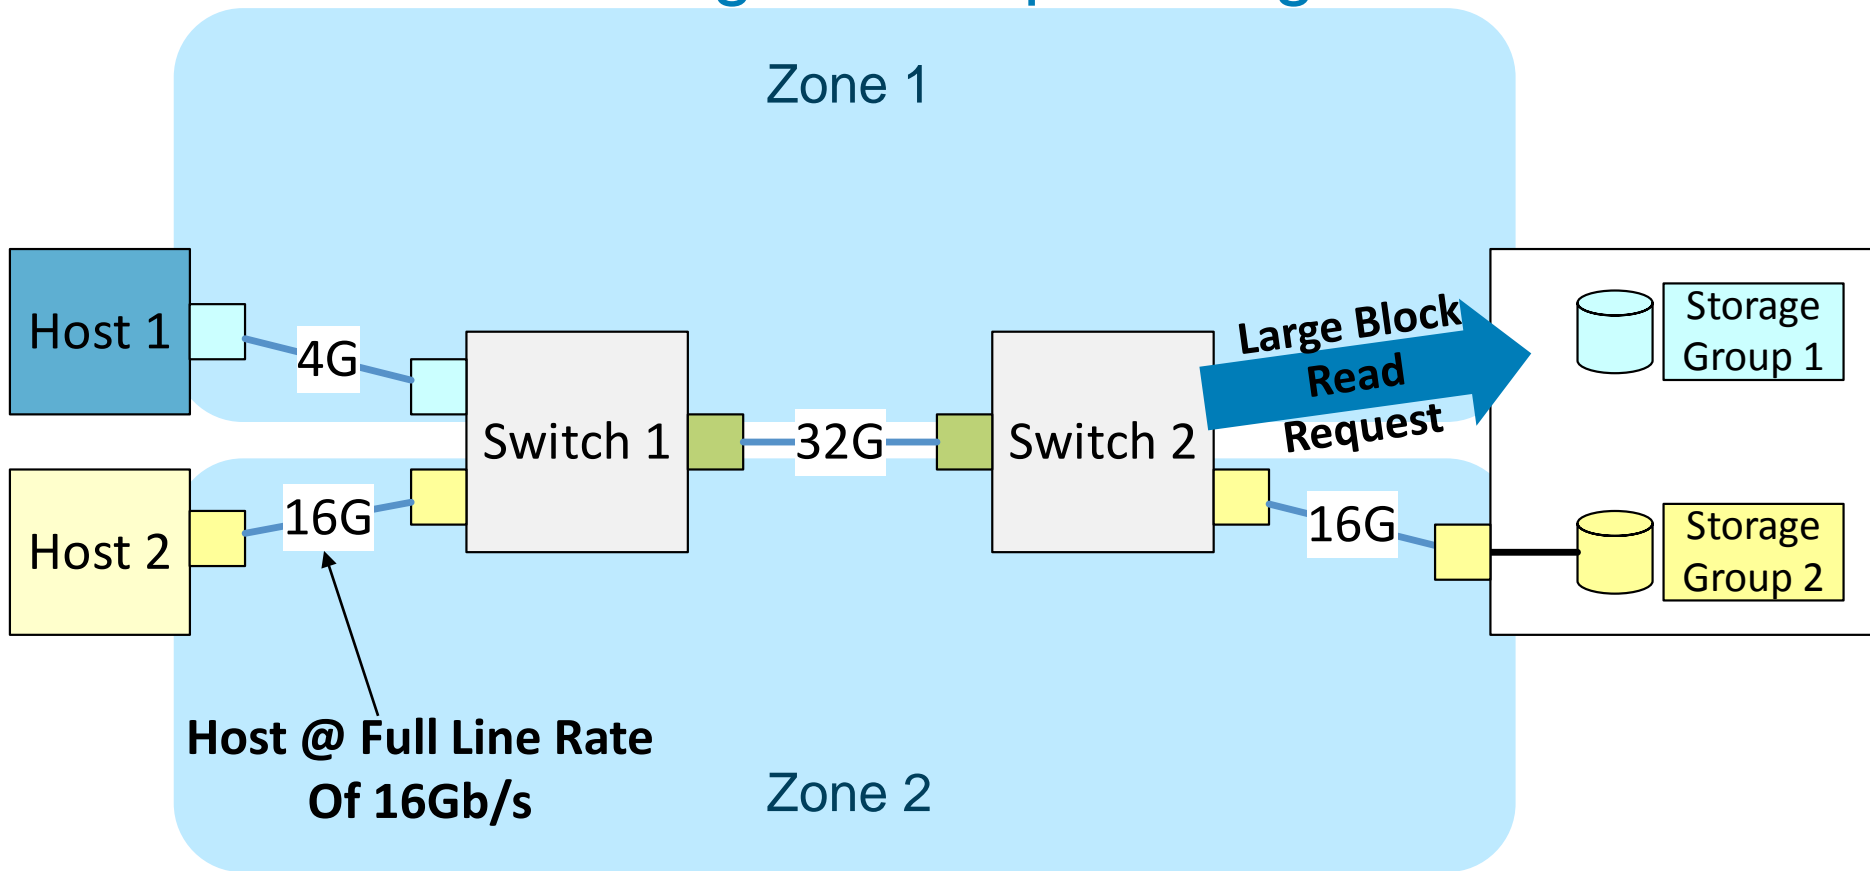
# Congestion Spreading



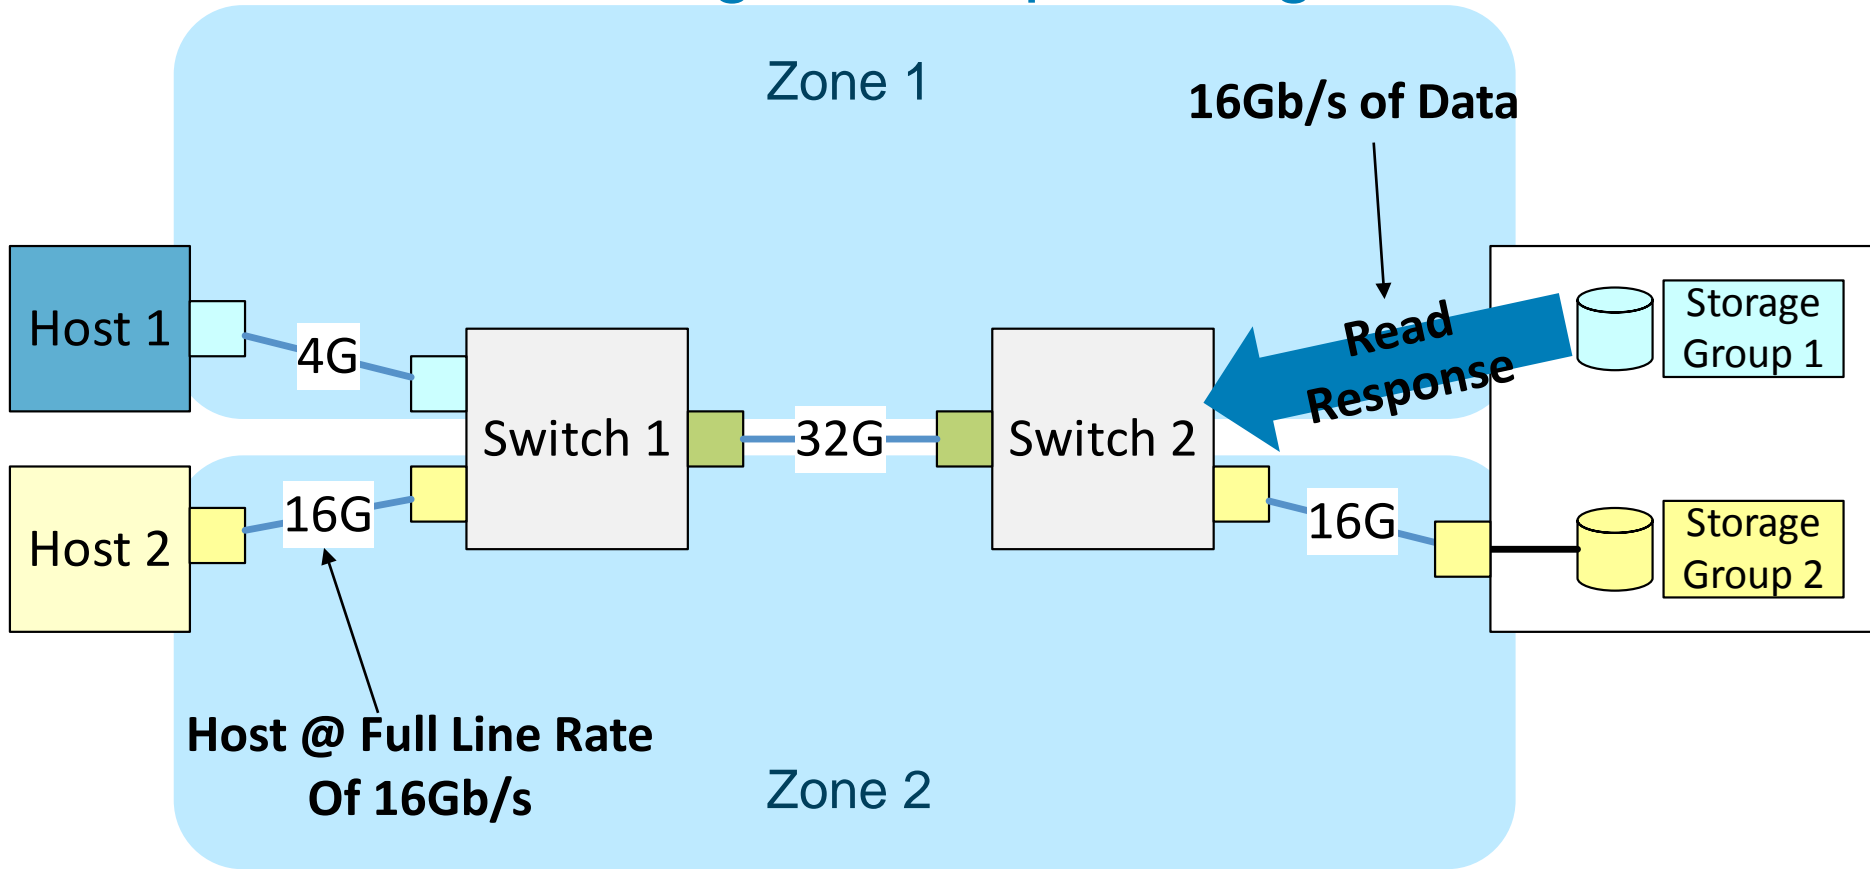
# Congestion Spreading



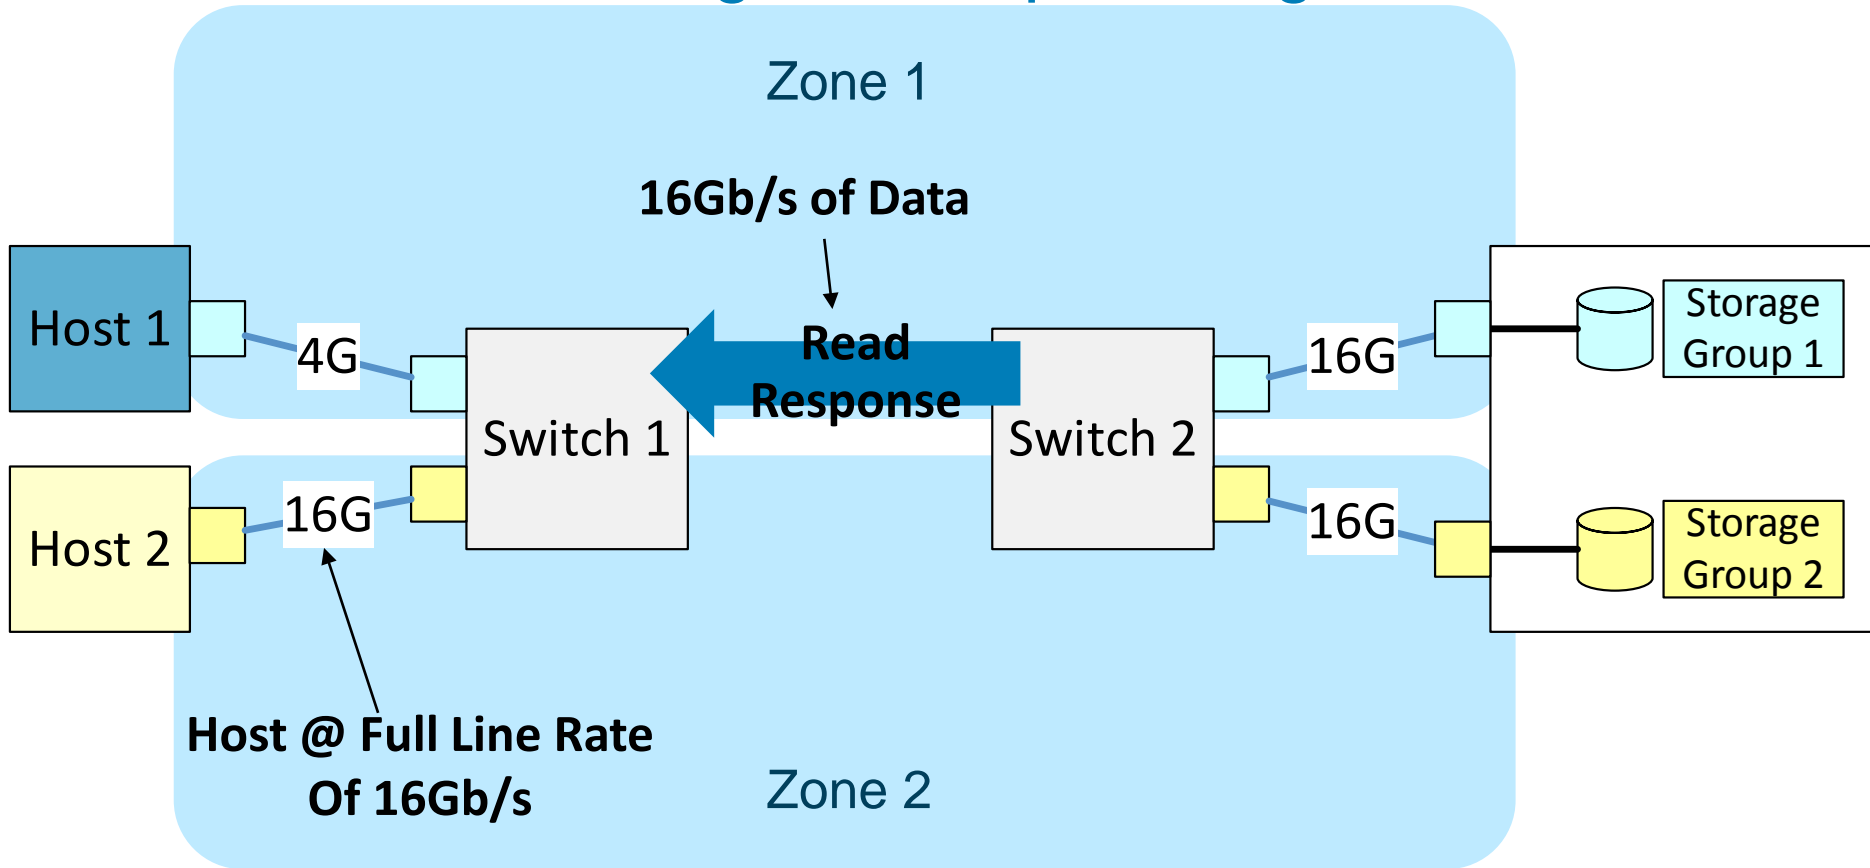
# Congestion Spreading



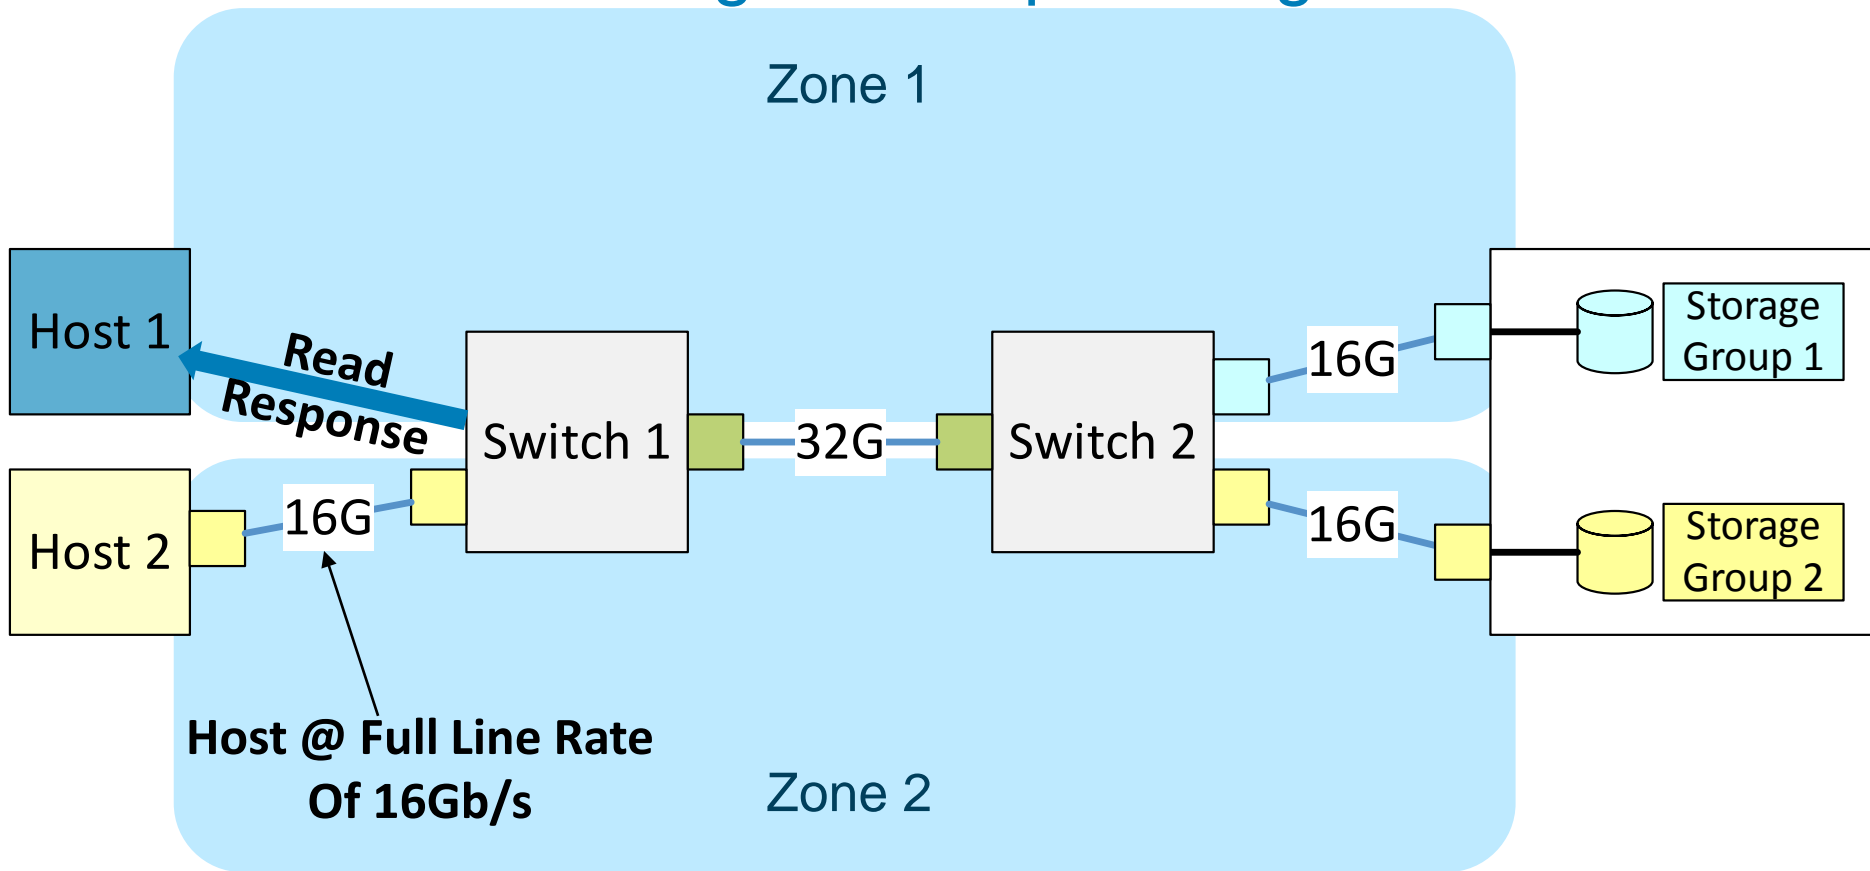
# Congestion Spreading



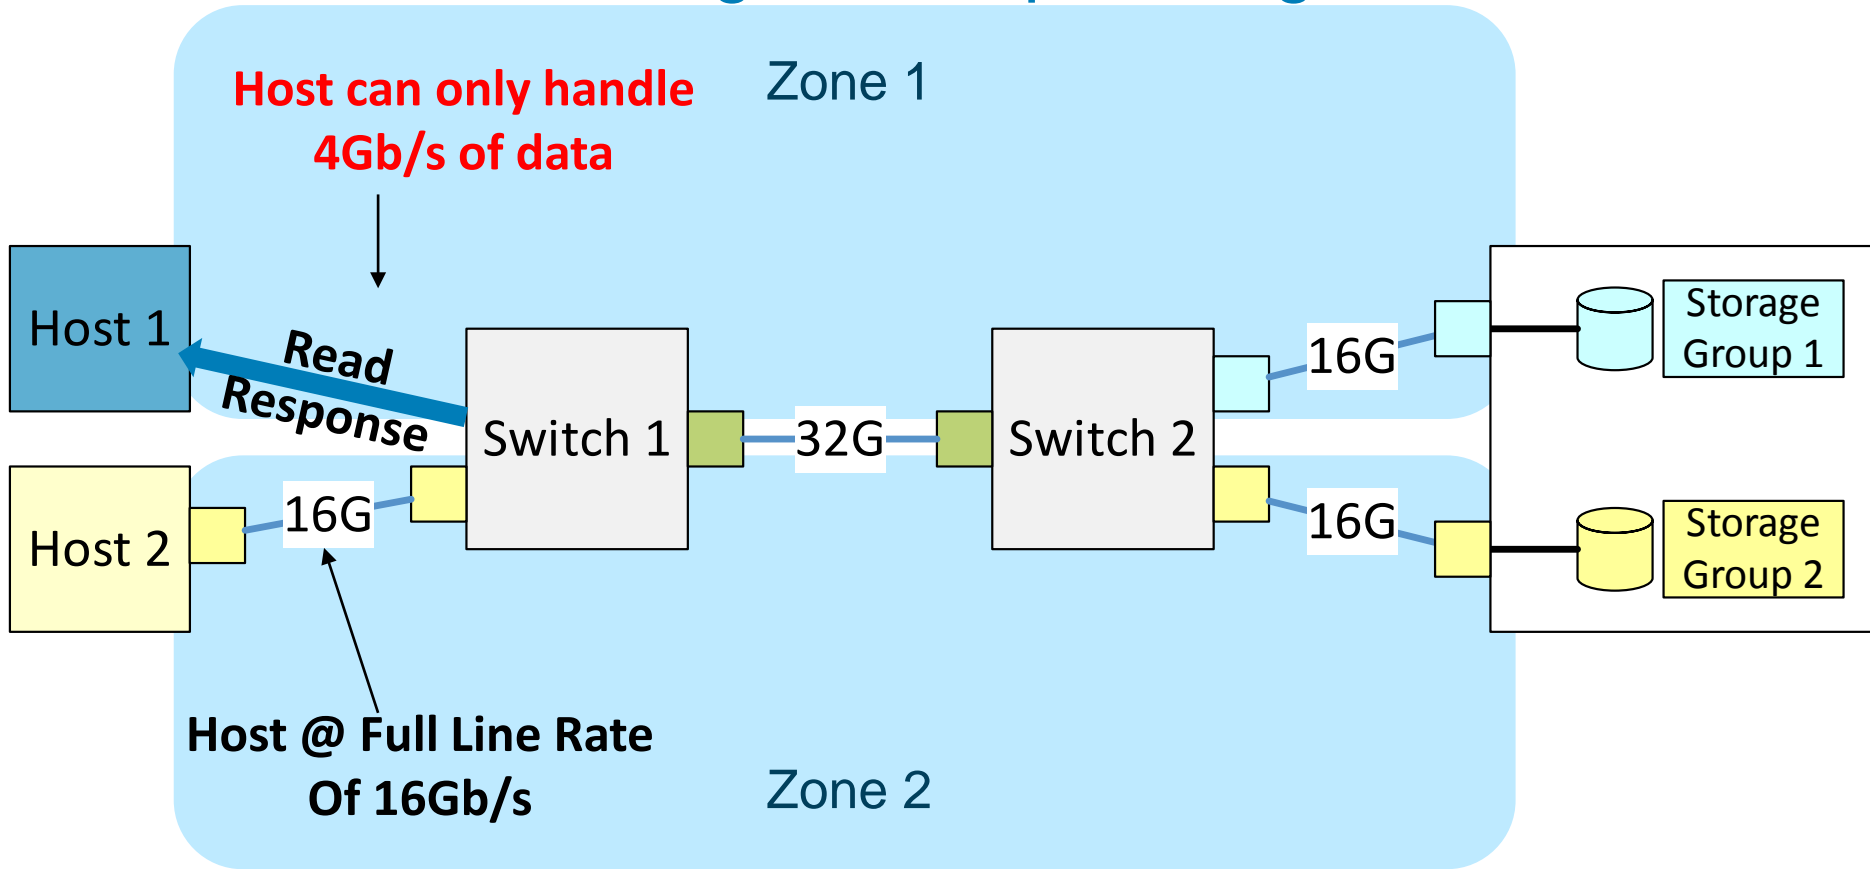
# Congestion Spreading



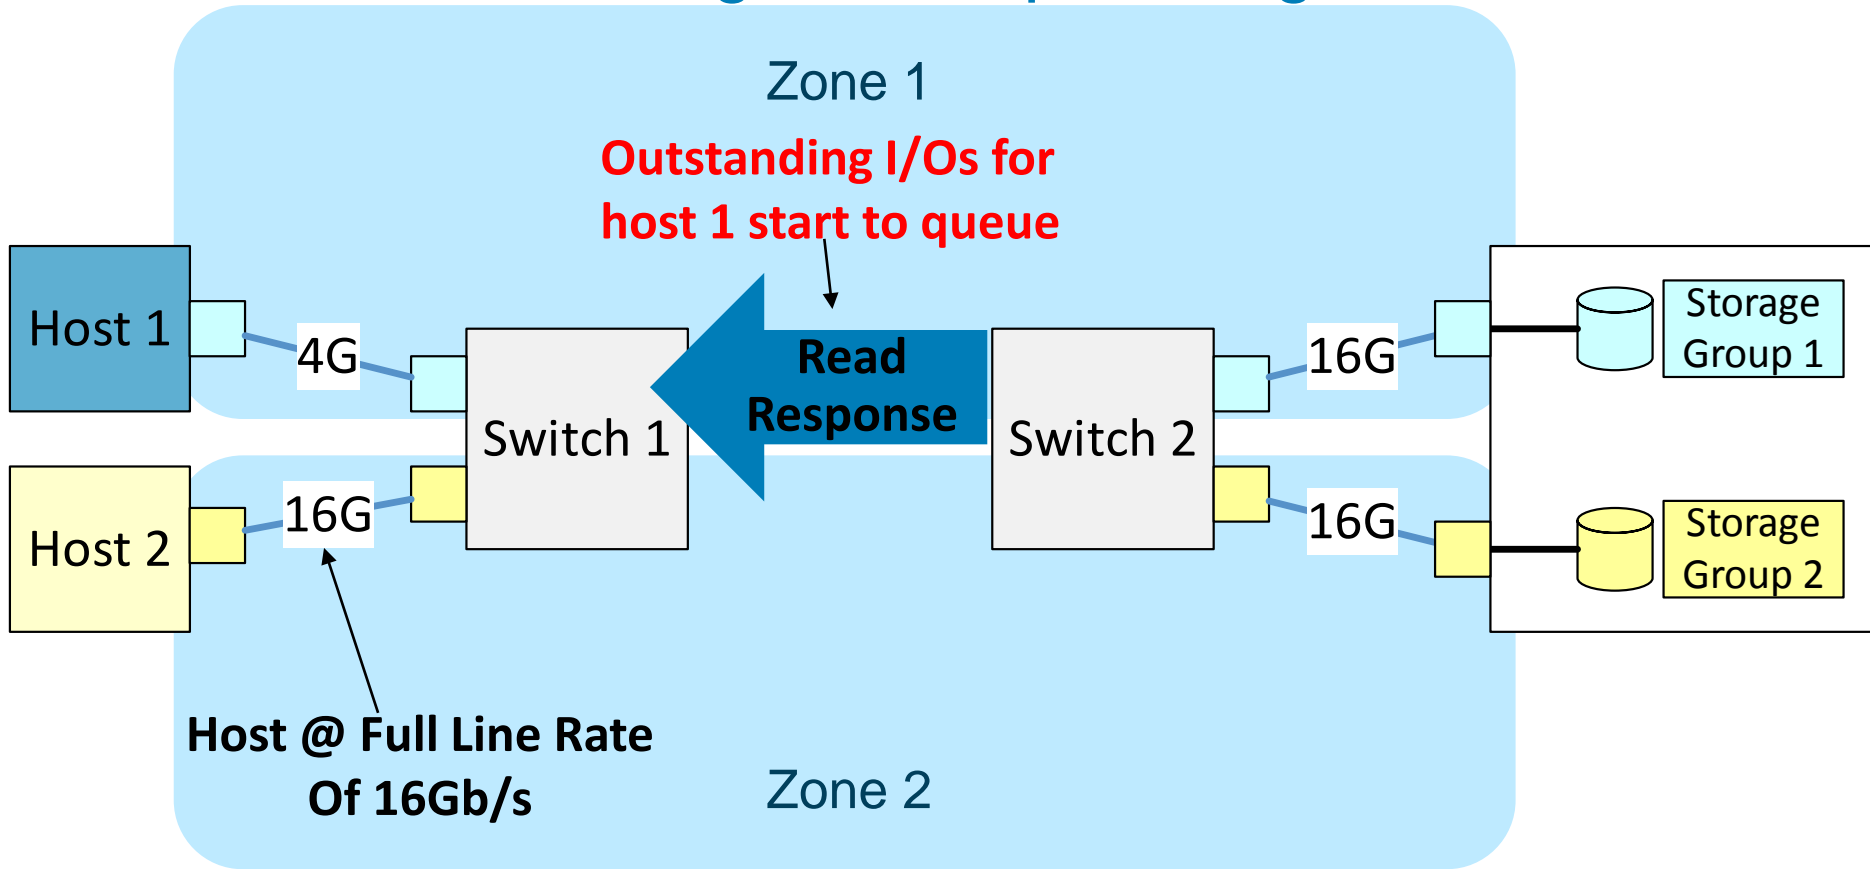
# Congestion Spreading



# Congestion Spreading

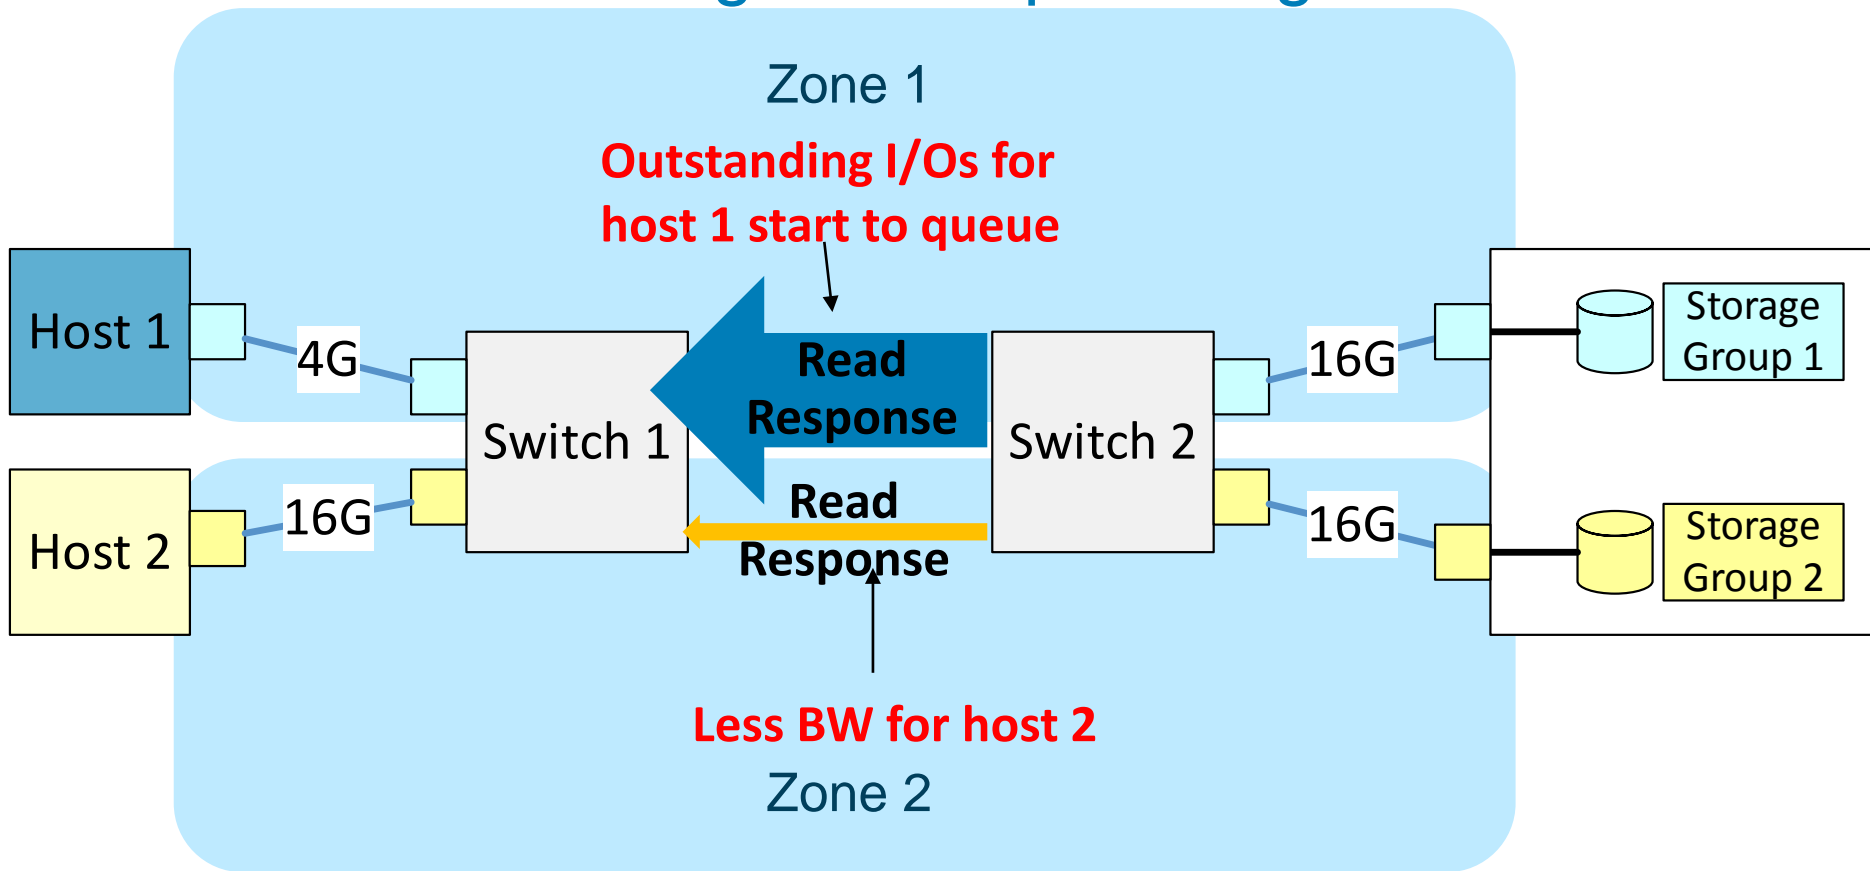


# Congestion Spreading

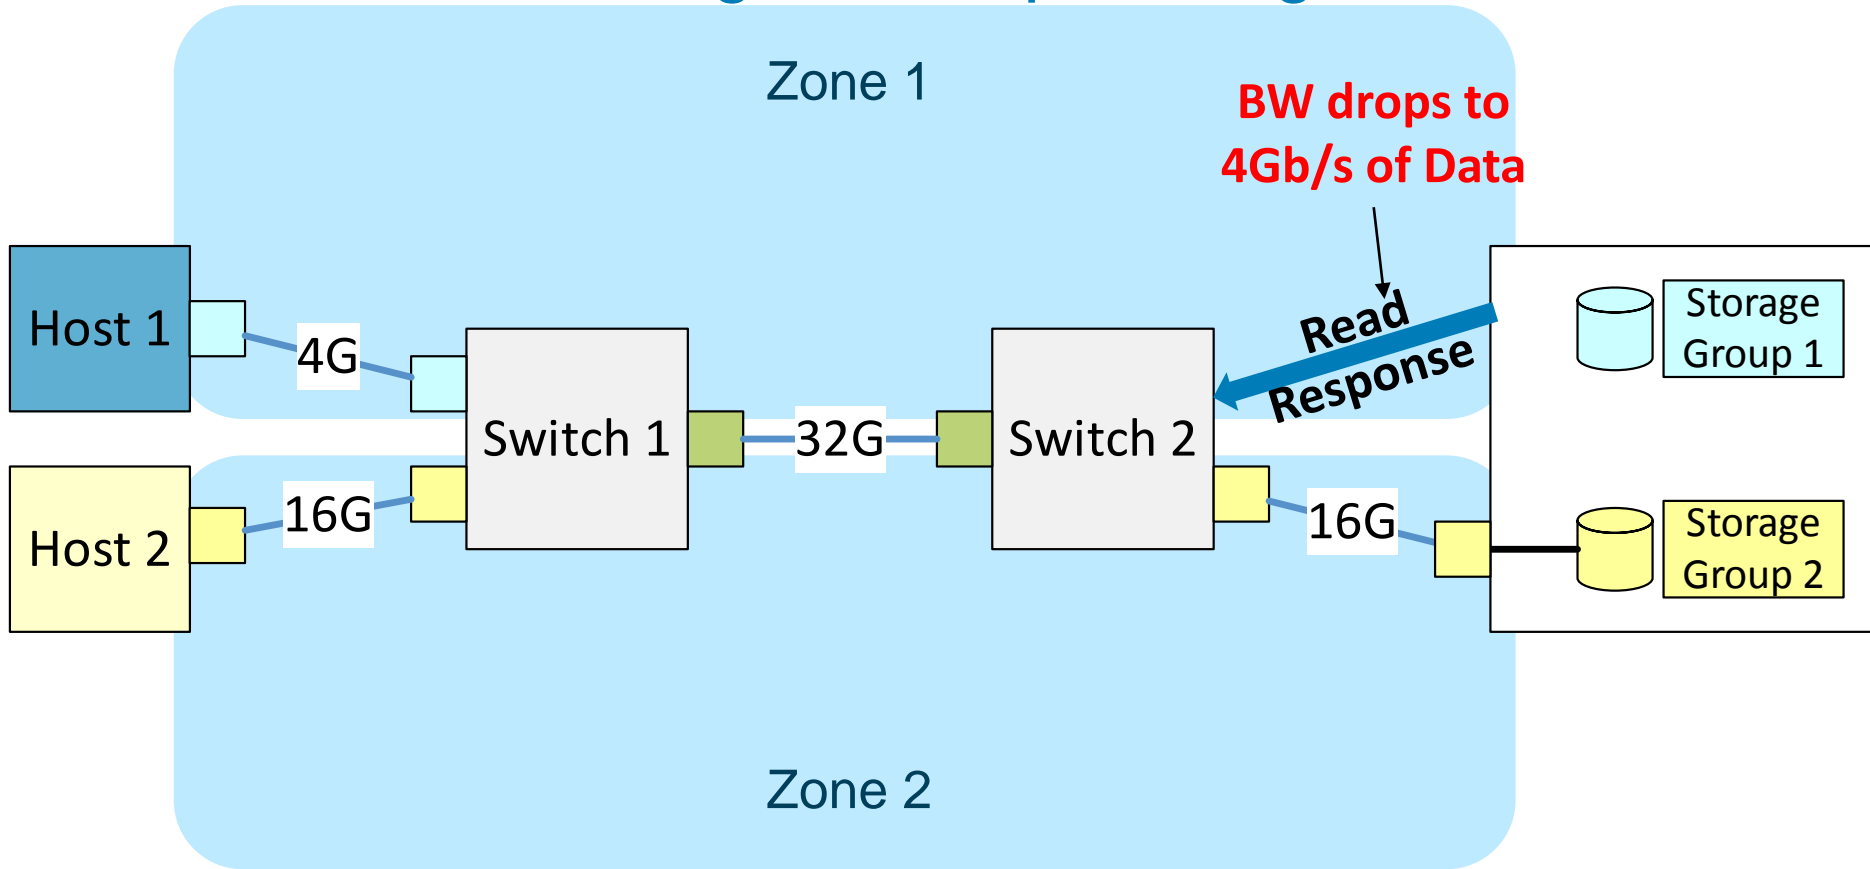




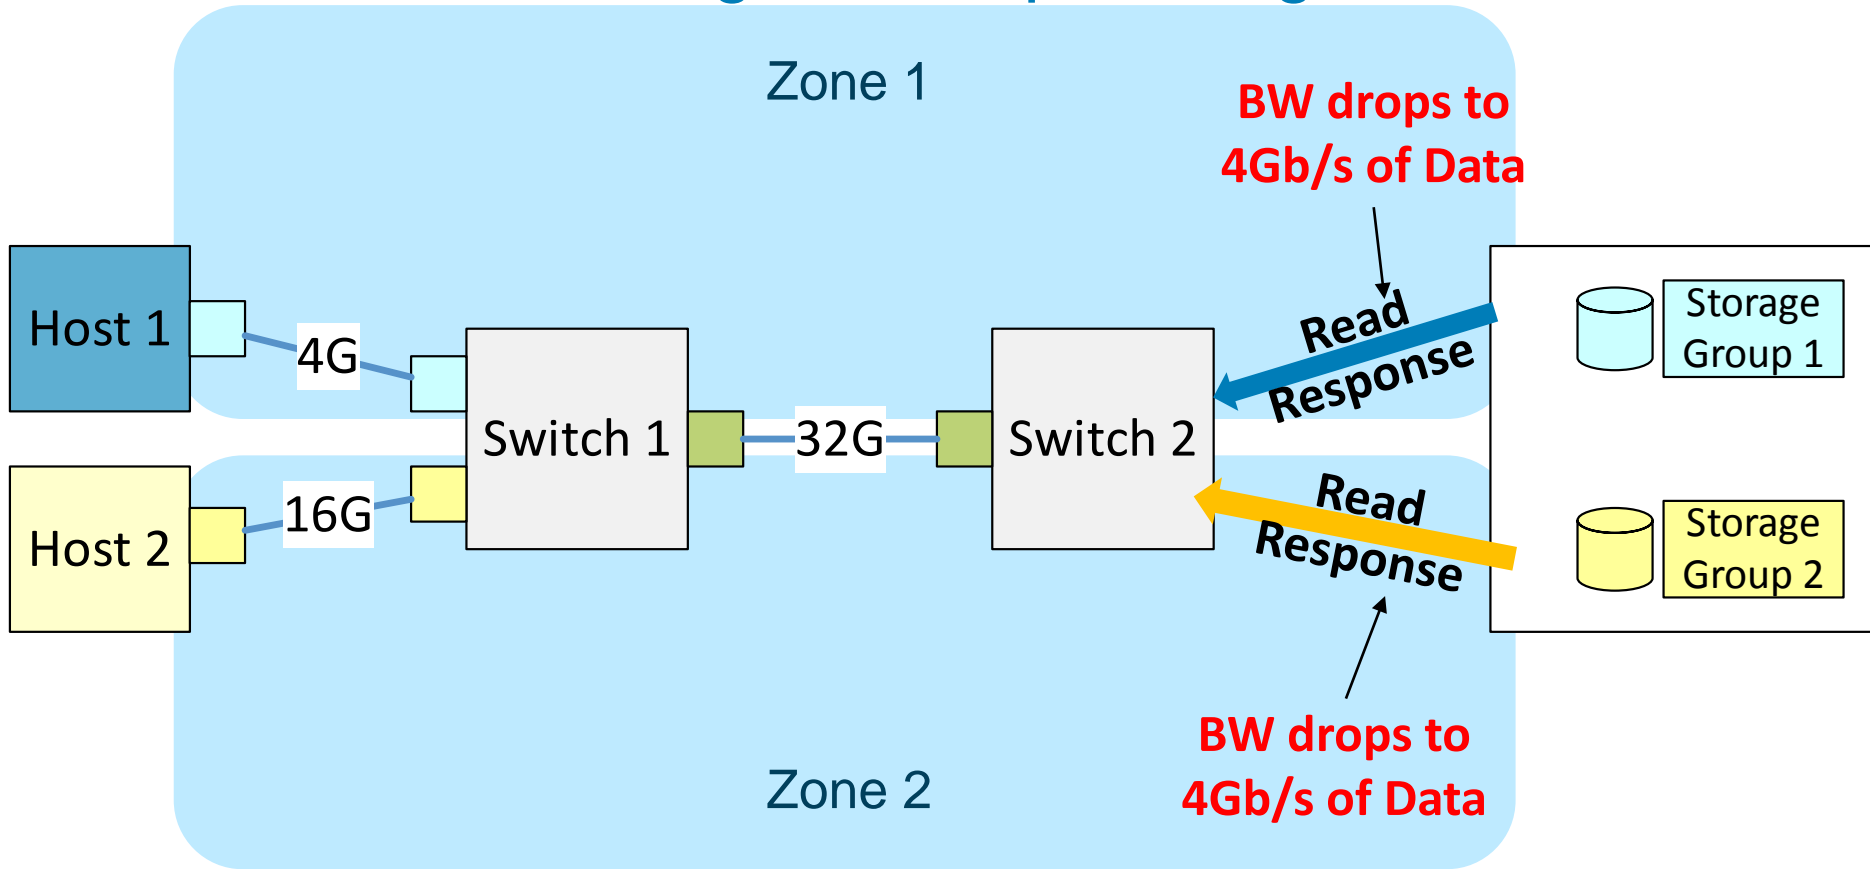
# Congestion Spreading



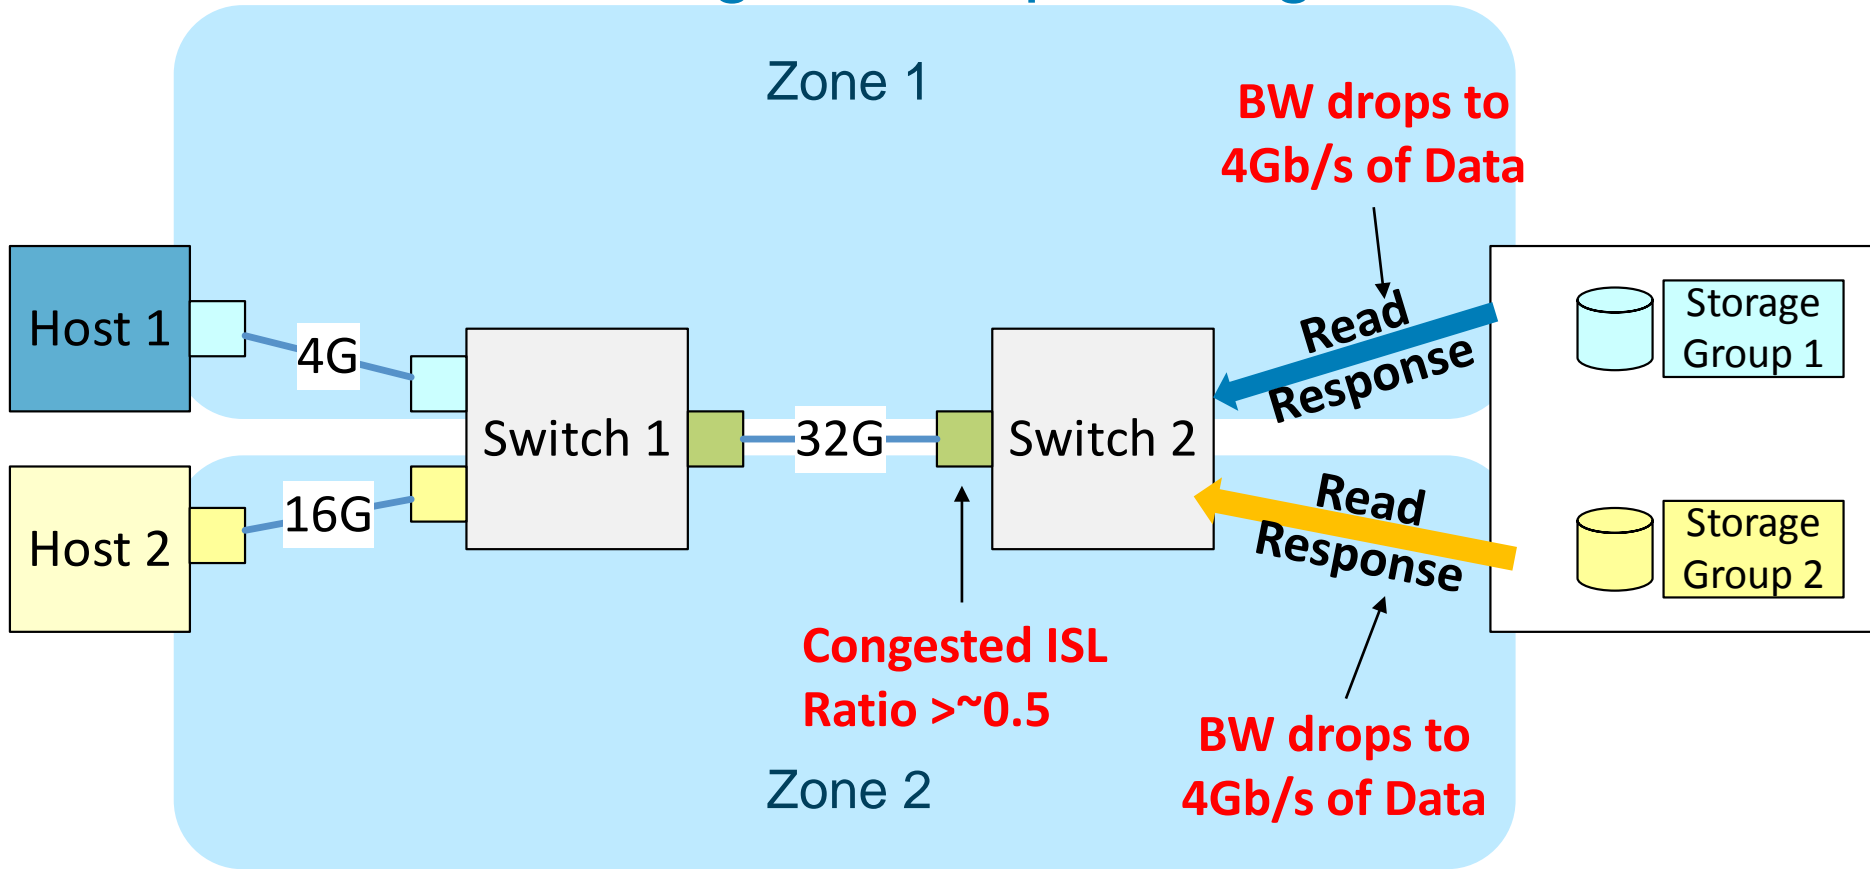
# Congestion Spreading




# Congestion Spreading



# Congestion Spreading




# Lost Credit Slow Drain




Thanks to Dell EMC E-lab for testing and results!

# Lost Credit Slow Drain




Thanks to Dell EMC E-lab for testing and results!

# Lost Credit Slow Drain



Thanks to Dell EMC E-lab for testing and results!

# Lost Credit Slow Drain




Thanks to Dell EMC E-lab for testing and results!



# Lost Credit Slow Drain

Performance impact  
when BB\_Credit is 5 or  
less.




Thanks to Dell EMC E-lab for testing and results!

# Lost Credit Slow Drain

$$\text{C Ratio} = \frac{\text{Time at zero BB\_Credit}}{\text{Frames transmitted}}$$

Ratio




Thanks to Dell EMC E-lab for testing and Results!

# Lost Credit Slow Drain

$$\text{C Ratio} = \frac{\text{Time at zero BB\_Credit}}{\text{Frames transmitted}}$$


Ratio



Thanks to Dell EMC E-lab for testing and Results!

# Lost Credit Slow Drain

$$\text{C Ratio} = \frac{\text{Time at zero BB\_Credit}}{\text{Frames transmitted}}$$




Thanks to Dell EMC E-lab for testing and Results!

# Lost Credit Slow Drain

$$\text{C Ratio} = \frac{\text{Time at zero BB\_Credit}}{\text{Frames transmitted}}$$


Ratio



Thanks to Dell EMC E-lab for testing and Results!

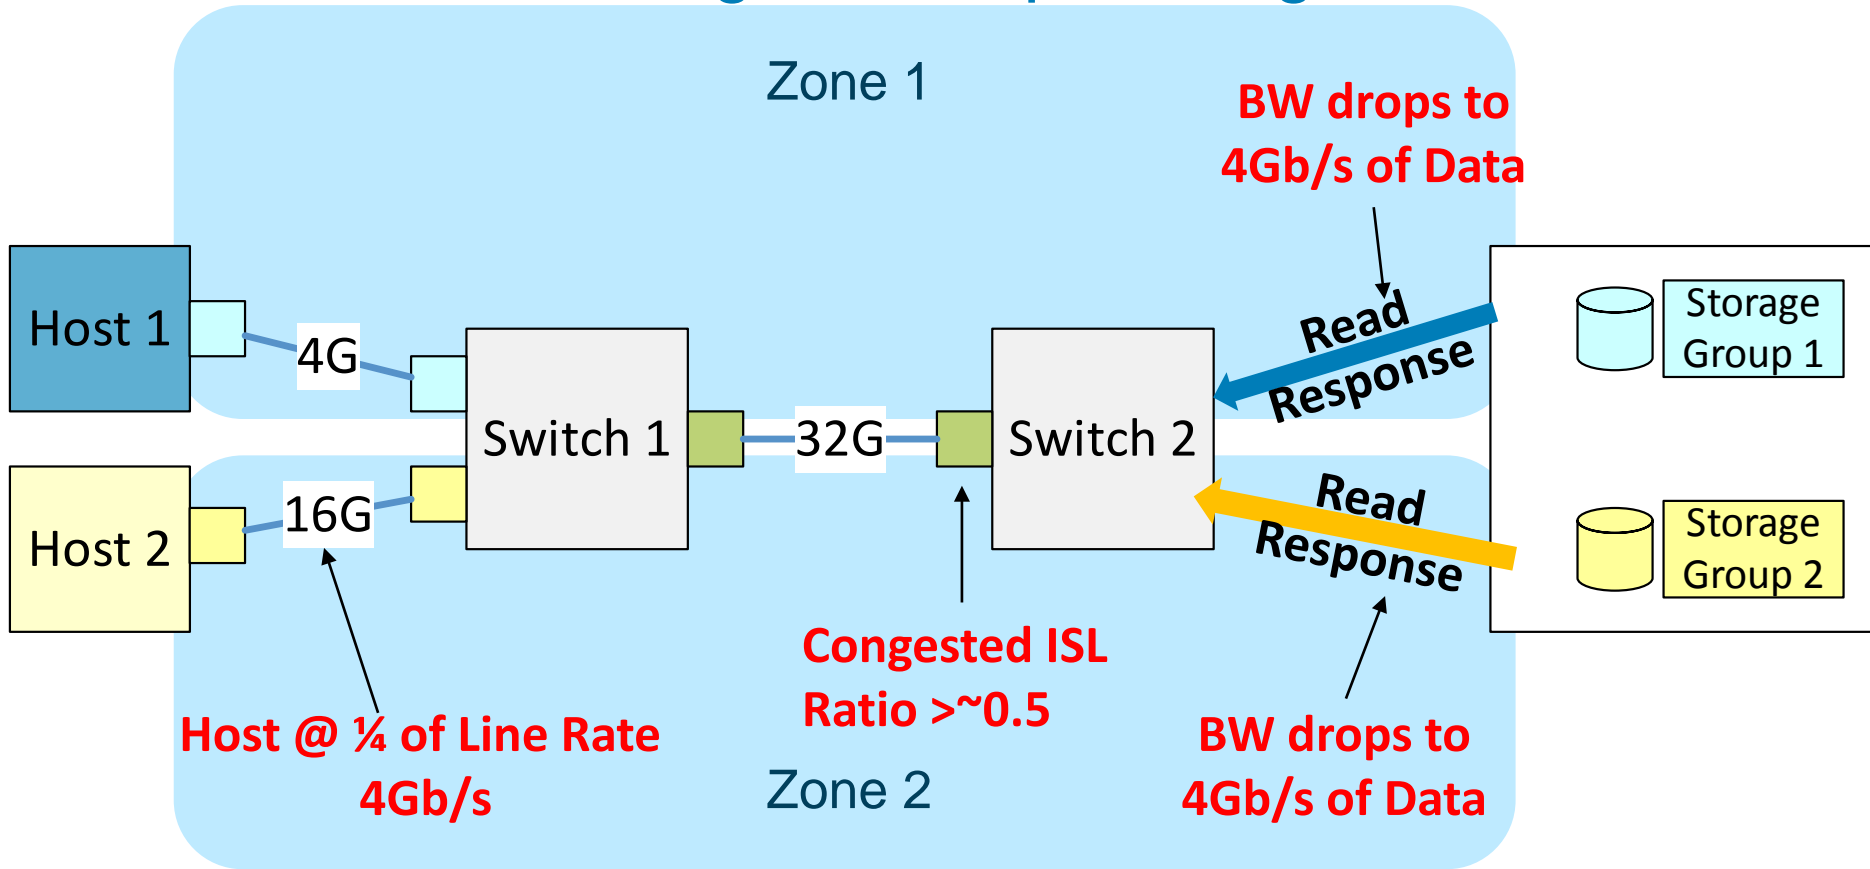
# Lost Credit Slow Drain

$$\text{C Ratio} = \frac{\text{Time at zero BB\_Credit}}{\text{Frames transmitted}}$$

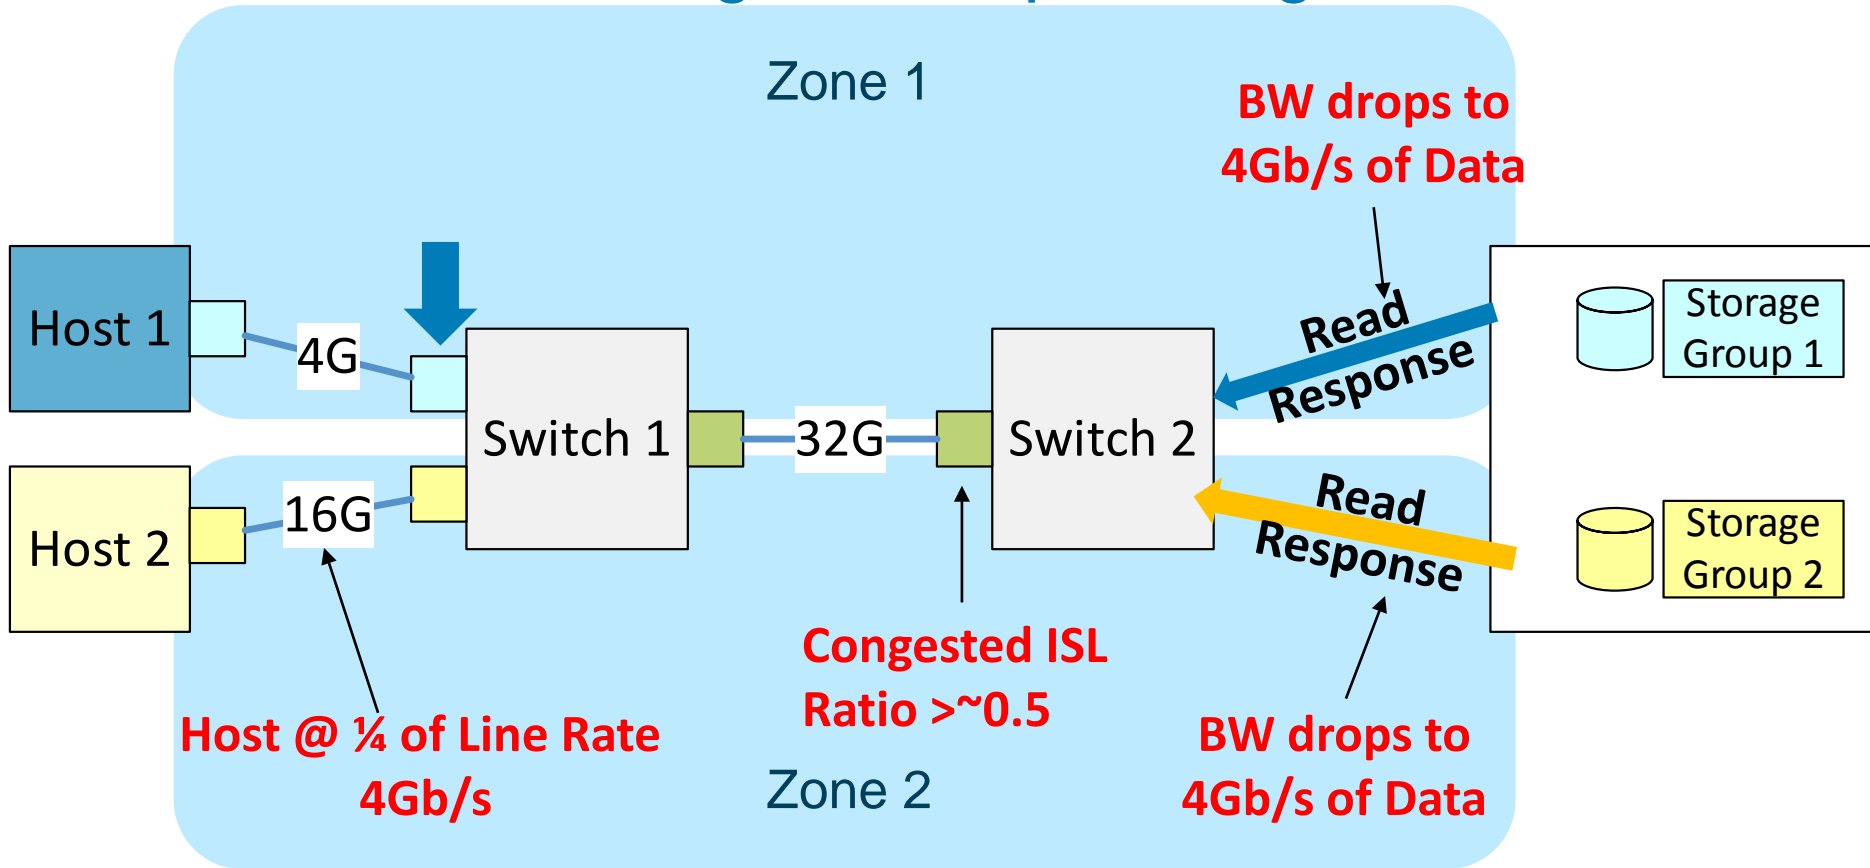


Thanks to Dell EMC E-lab for testing and Results!

# Congestion Spreading



# Congestion Spreading





# Agenda

- What's new with storage connectivity
  - Protocol adoption
  - Connectrix Product Updates
- Congestion Spreading and its impact
  - Congestion Spreading terminology
  - Congestion Spreading & Innocent flows
- **SAN Best practices for All Flash Arrays (AFAs)**
- Introduction to NVMe
  - FC-NVMe

# SAN Best Practices for AFAs

# Congestion Spreading Avoidance

# Congestion Spreading Avoidance

Short Term (today)			

# Congestion Spreading Avoidance

Short Term (today)			
Medium term			

# Congestion Spreading Avoidance

<b>Short Term (today)</b>			
<b>Medium term</b>			
<b>Long term</b>			

# Congestion Spreading Avoidance

	Prevention		
Short Term (today)			
Medium term			
Long term			

# Congestion Spreading Avoidance

	Prevention	Detection	
Short Term (today)			
Medium term			
Long term			




# Congestion Spreading Avoidance

	Prevention	Detection	Remediation
Short Term (today)			
Medium term			
Long term			

# Congestion Spreading Avoidance


	Prevention	Detection	Remediation
Short Term (today)	<ul style="list-style-type: none"><li>• BW Ratio</li></ul>		
Medium term			
Long term			

# BW Ratio (example 1)

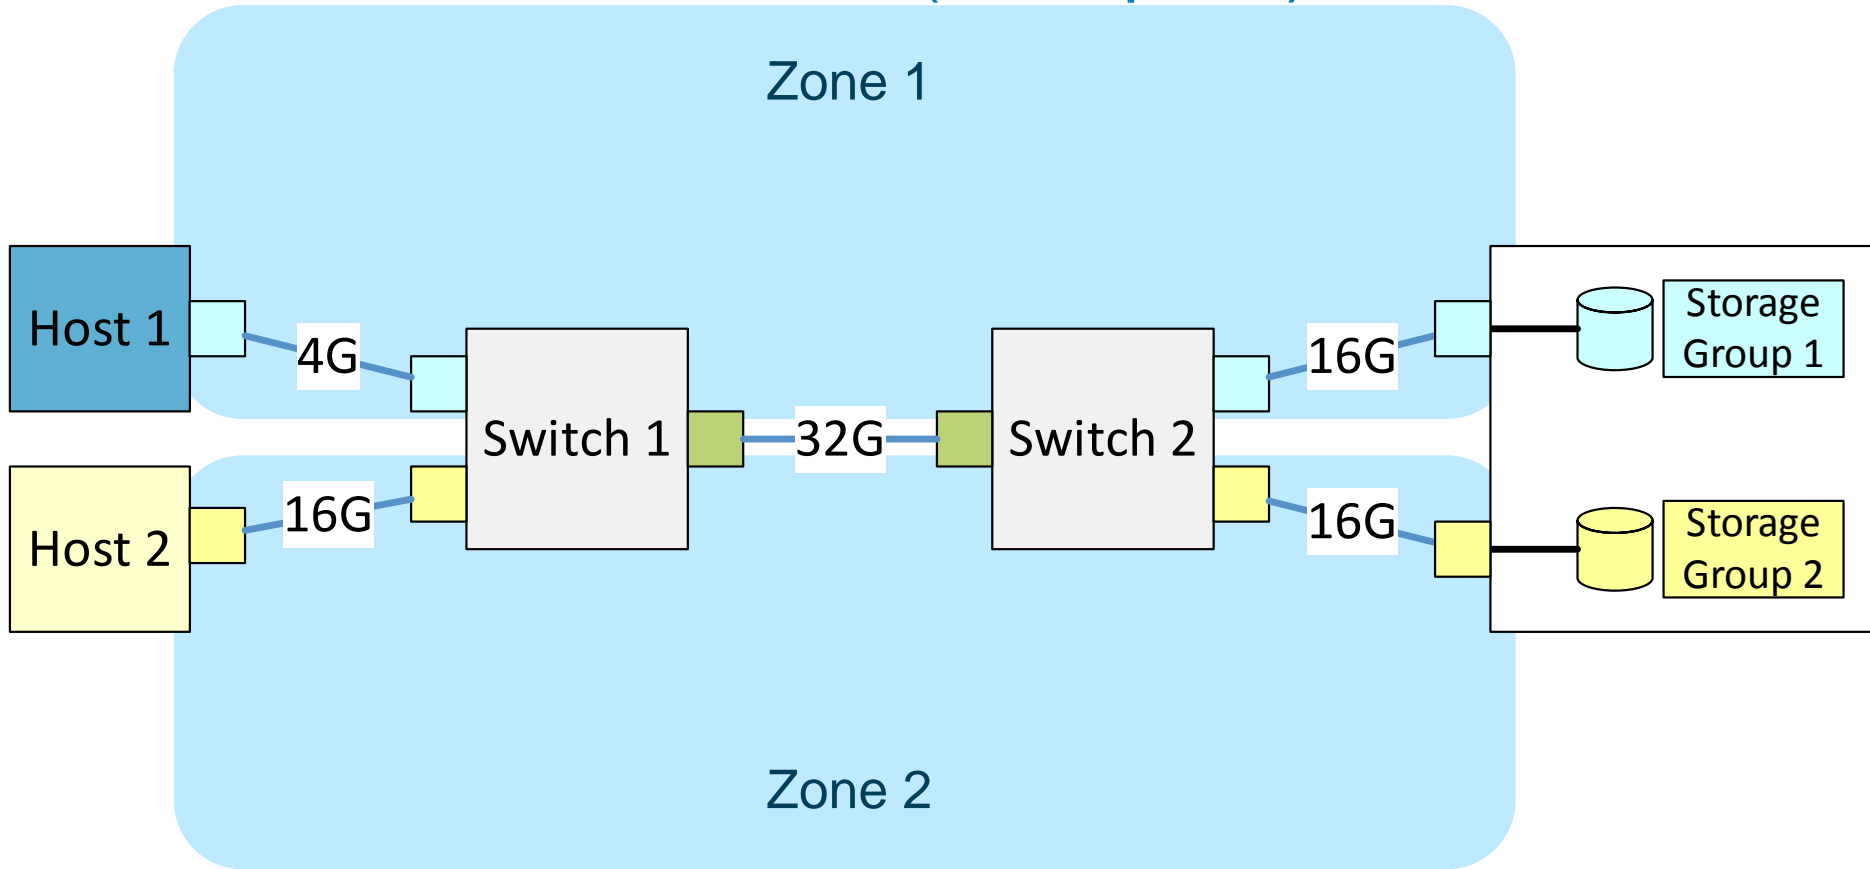


# BW Ratio (example 1)

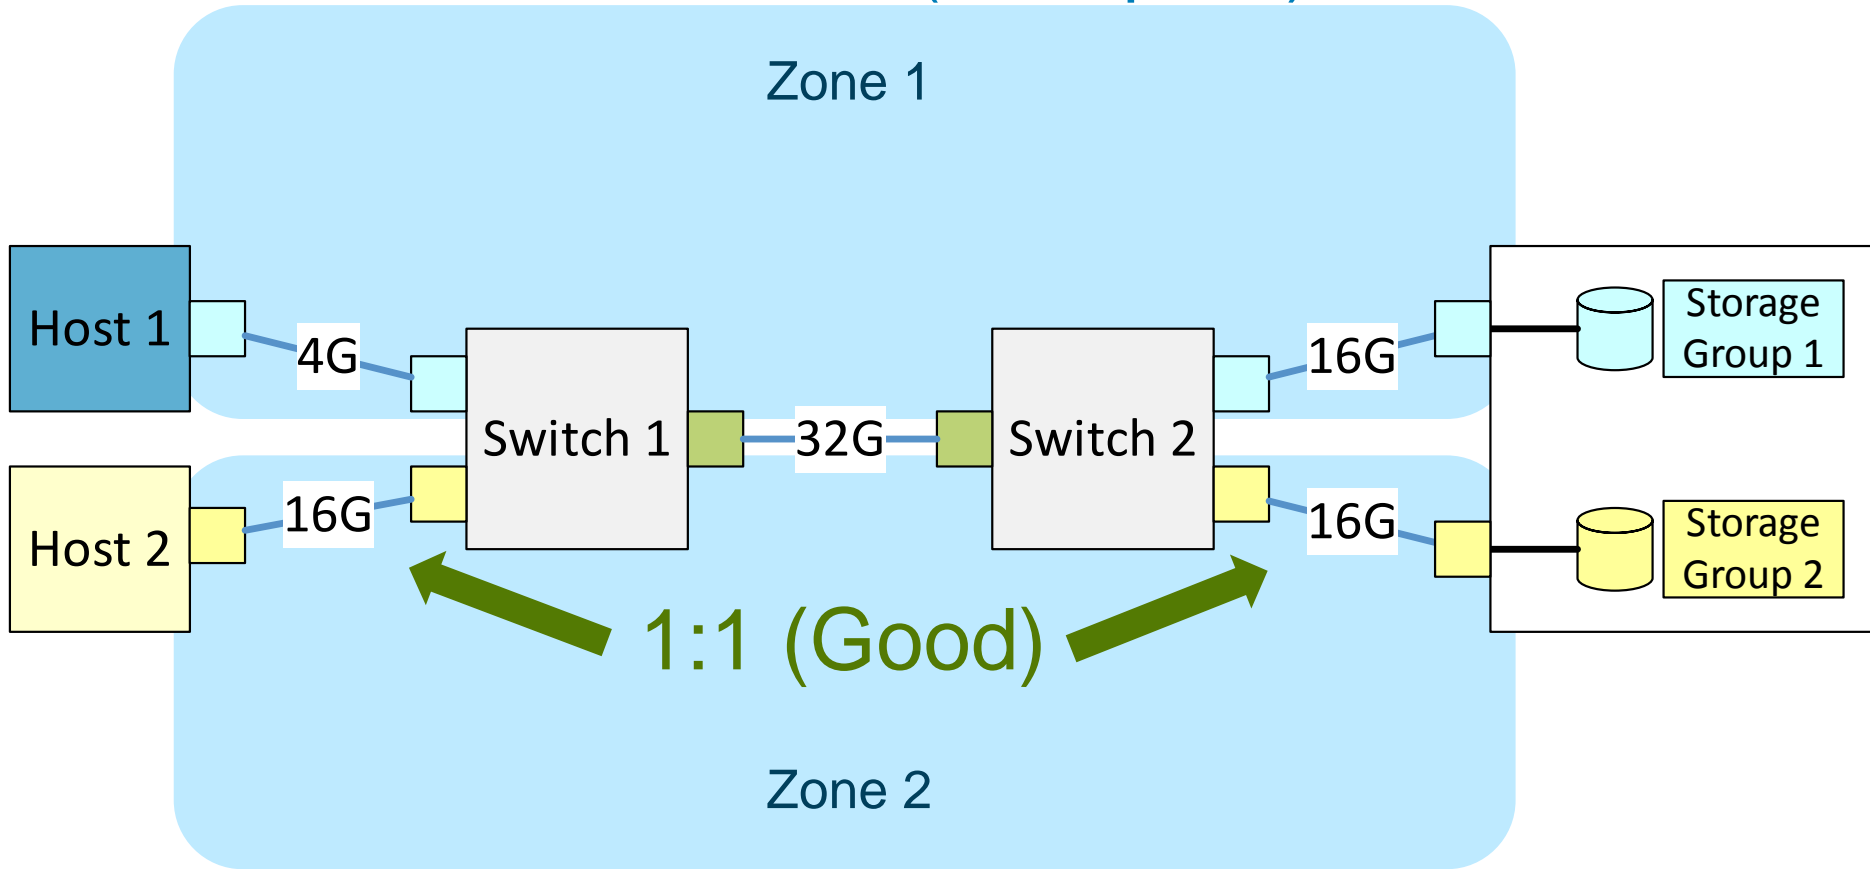
Zone 1



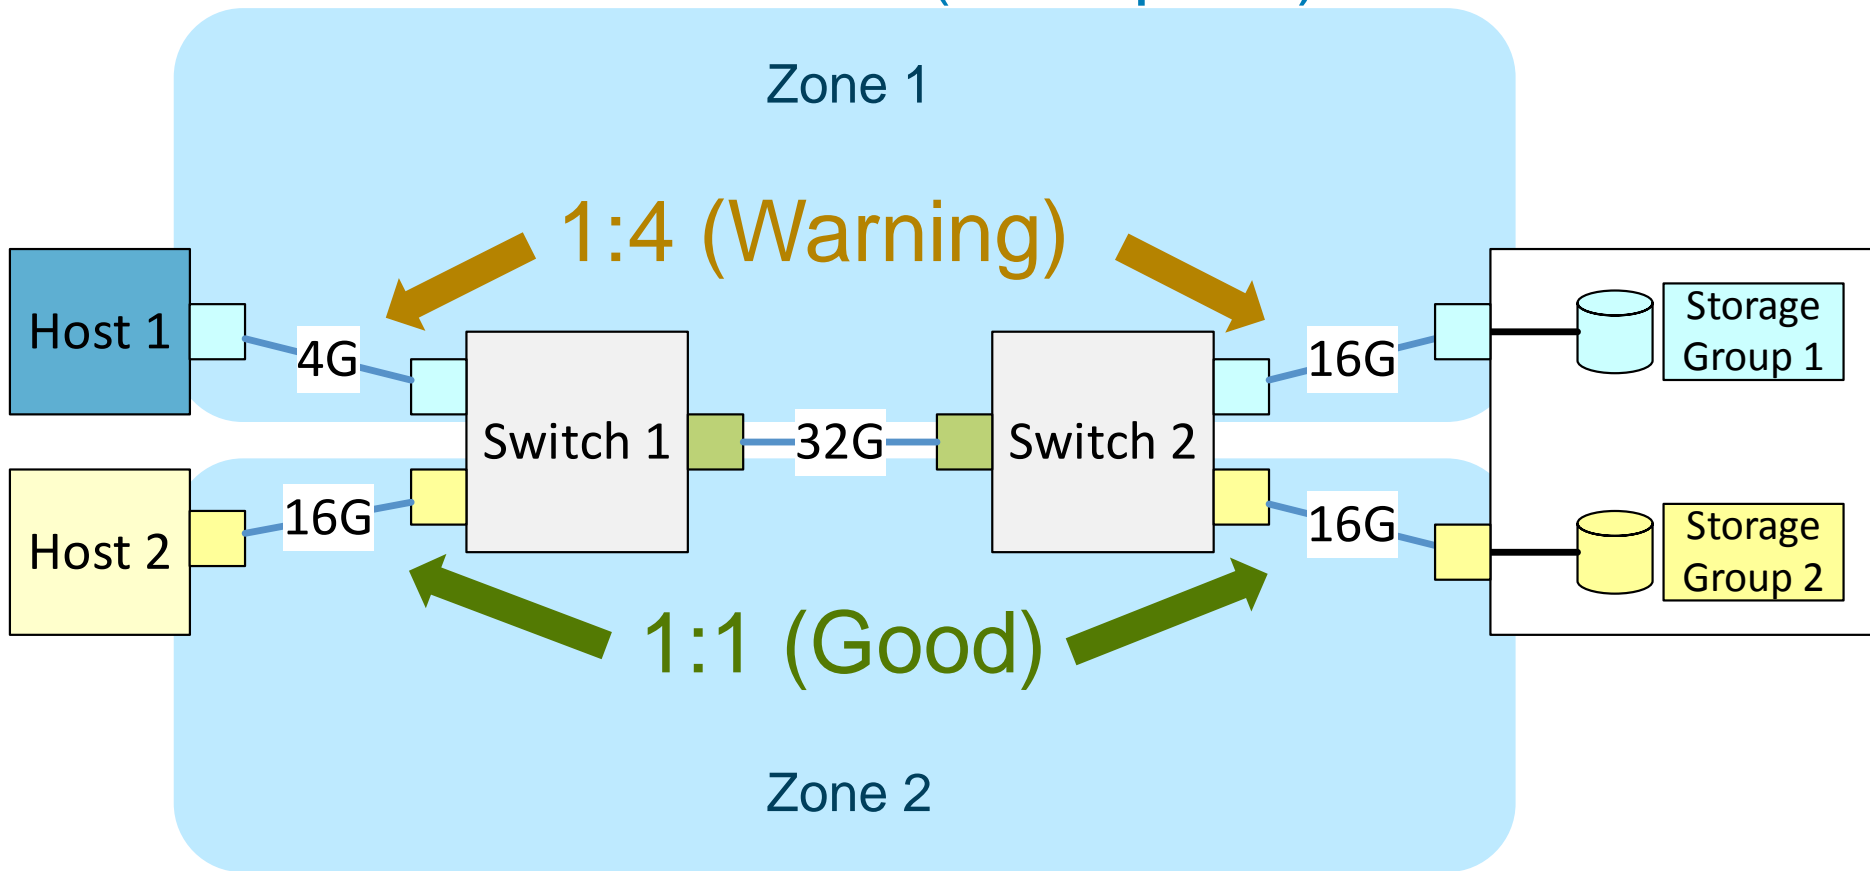
# BW Ratio (example 1)



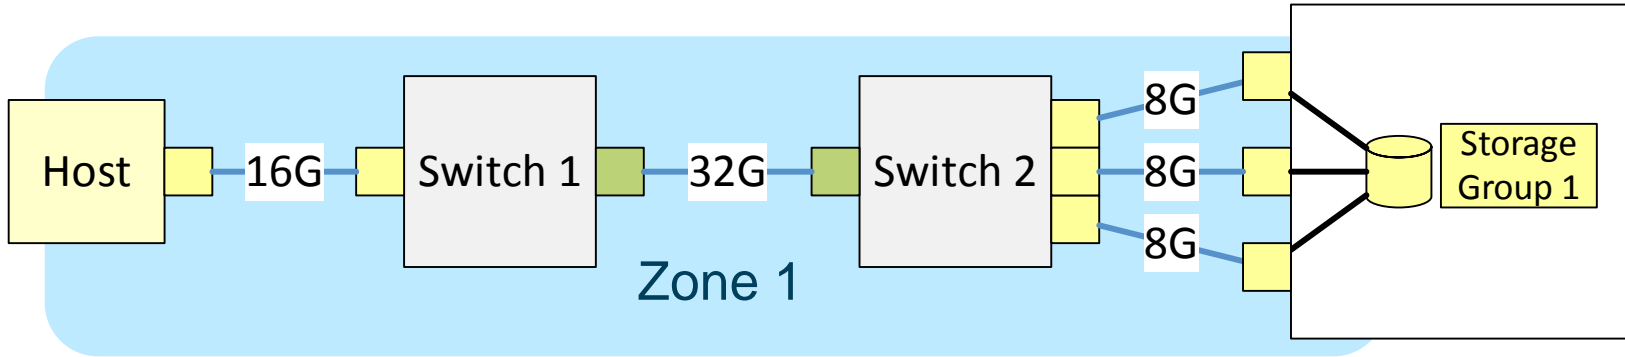
# BW Ratio (example 1)



# BW Ratio (example 1)




# BW Ratio (example 2)

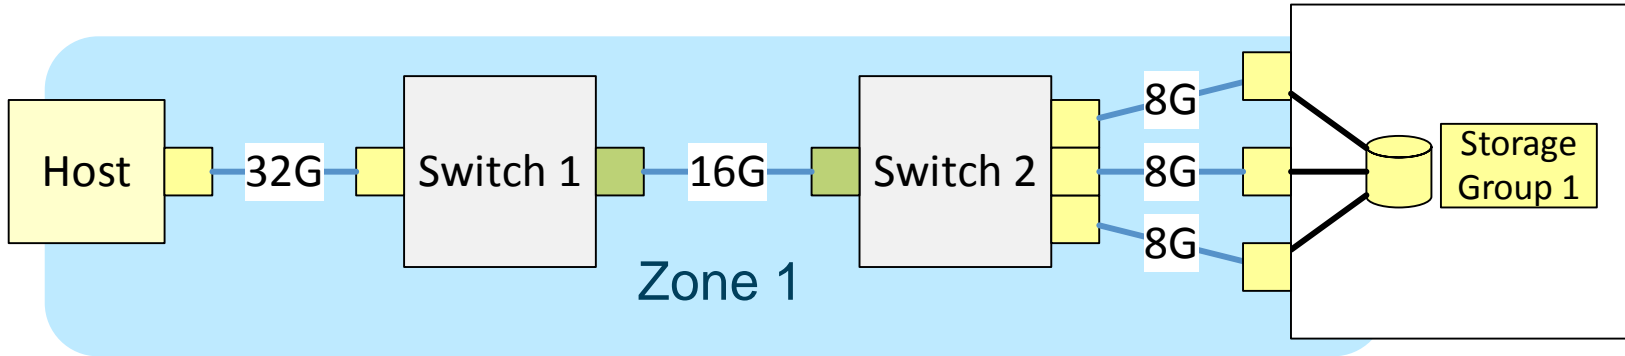





# BW Ratio (example 2)



# BW Ratio (example 3)



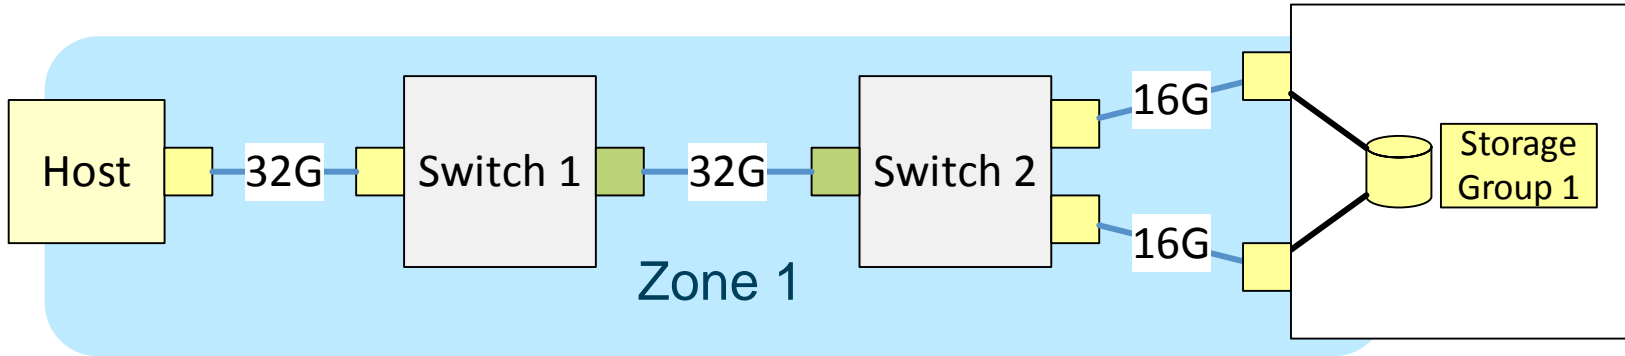
# BW Ratio (example 3)



# Congestion Spreading Avoidance

	Prevention	Detection	Remediation
Short Term (today)	<ul style="list-style-type: none"><li>• BW Ratio</li><li>• Modernize SAN</li></ul>		
Medium term			
Long term			

# Modernize




- Zero oversubscription (from end-to-end) is impractical in larger environments
- Focus on upgrading specific host, switch and storage connectivity

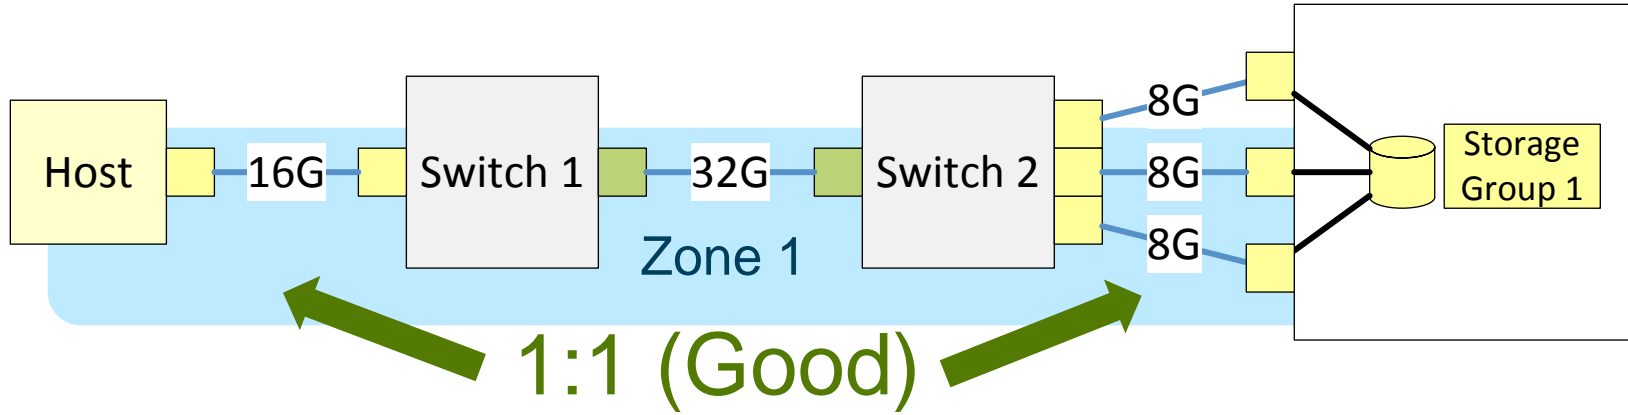
# Congestion Spreading Avoidance

	Prevention	Detection	Remediation
Short Term (today)	<ul style="list-style-type: none"><li>• BW Ratio</li><li>• Modernize SAN</li><li>• Fan-In/Out Ratio</li></ul>		
Medium term			
Long term			

# BW Ratio (example 2)



# BW Ratio (example 2)



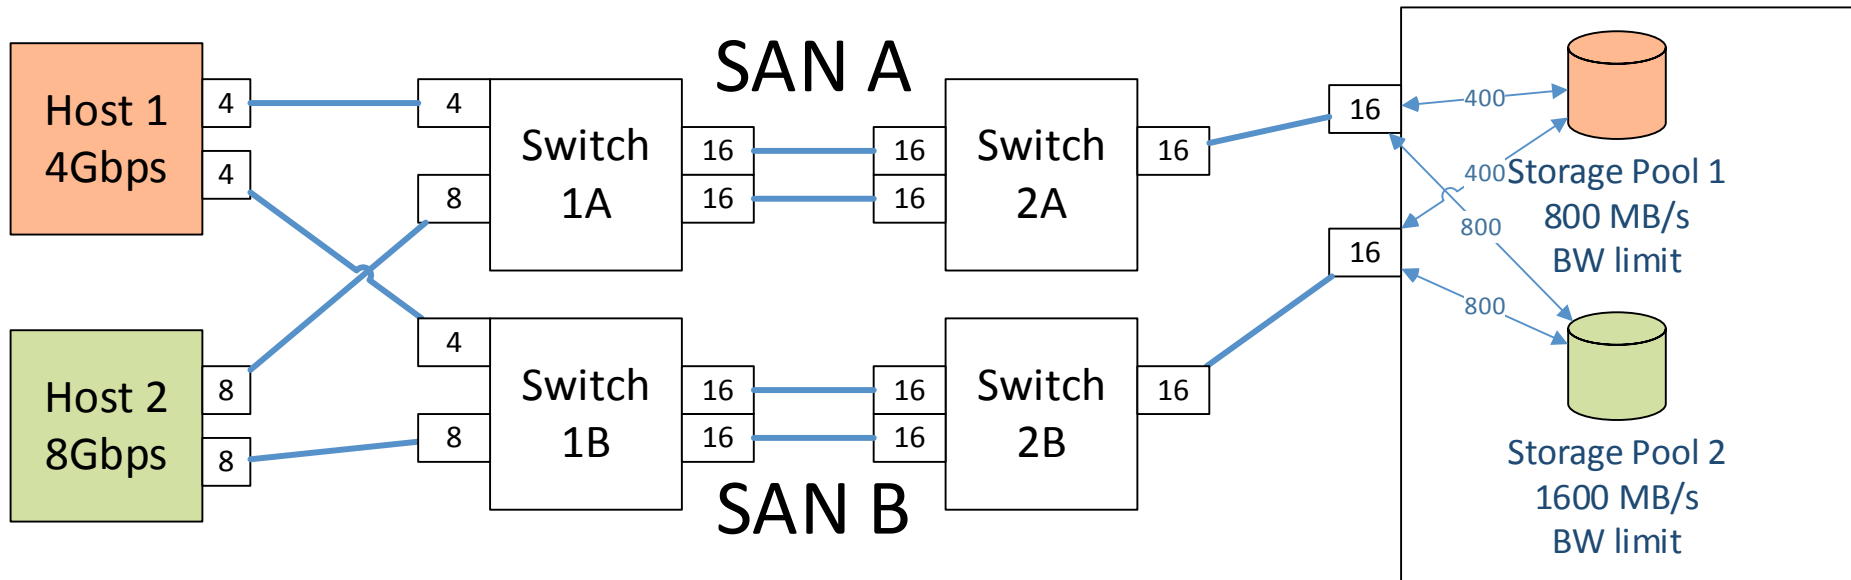


# Congestion Spreading Avoidance

	Prevention	Detection	Remediation
Short Term (today)	<ul style="list-style-type: none"><li>• BW Ratio</li><li>• Modernize SAN</li><li>• Fan-In/Out Ratio</li><li>• BW Limits</li></ul>		
Medium term			
Long term			

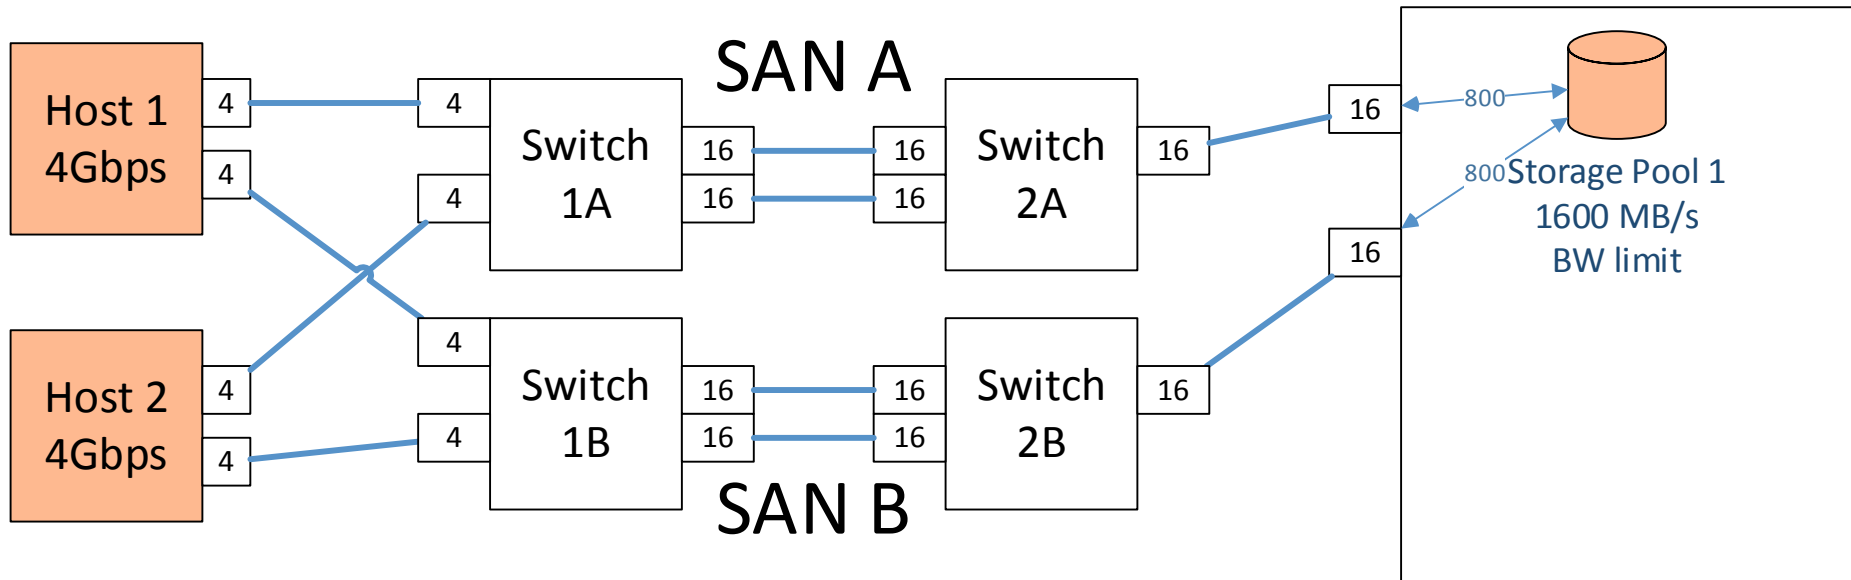
# BW Limit (example 1)

- Apply a BW Limit to a Storage Group or LUN
  - This approach works well when you need to limit BW to a single host.



# BW Limit (example 2)

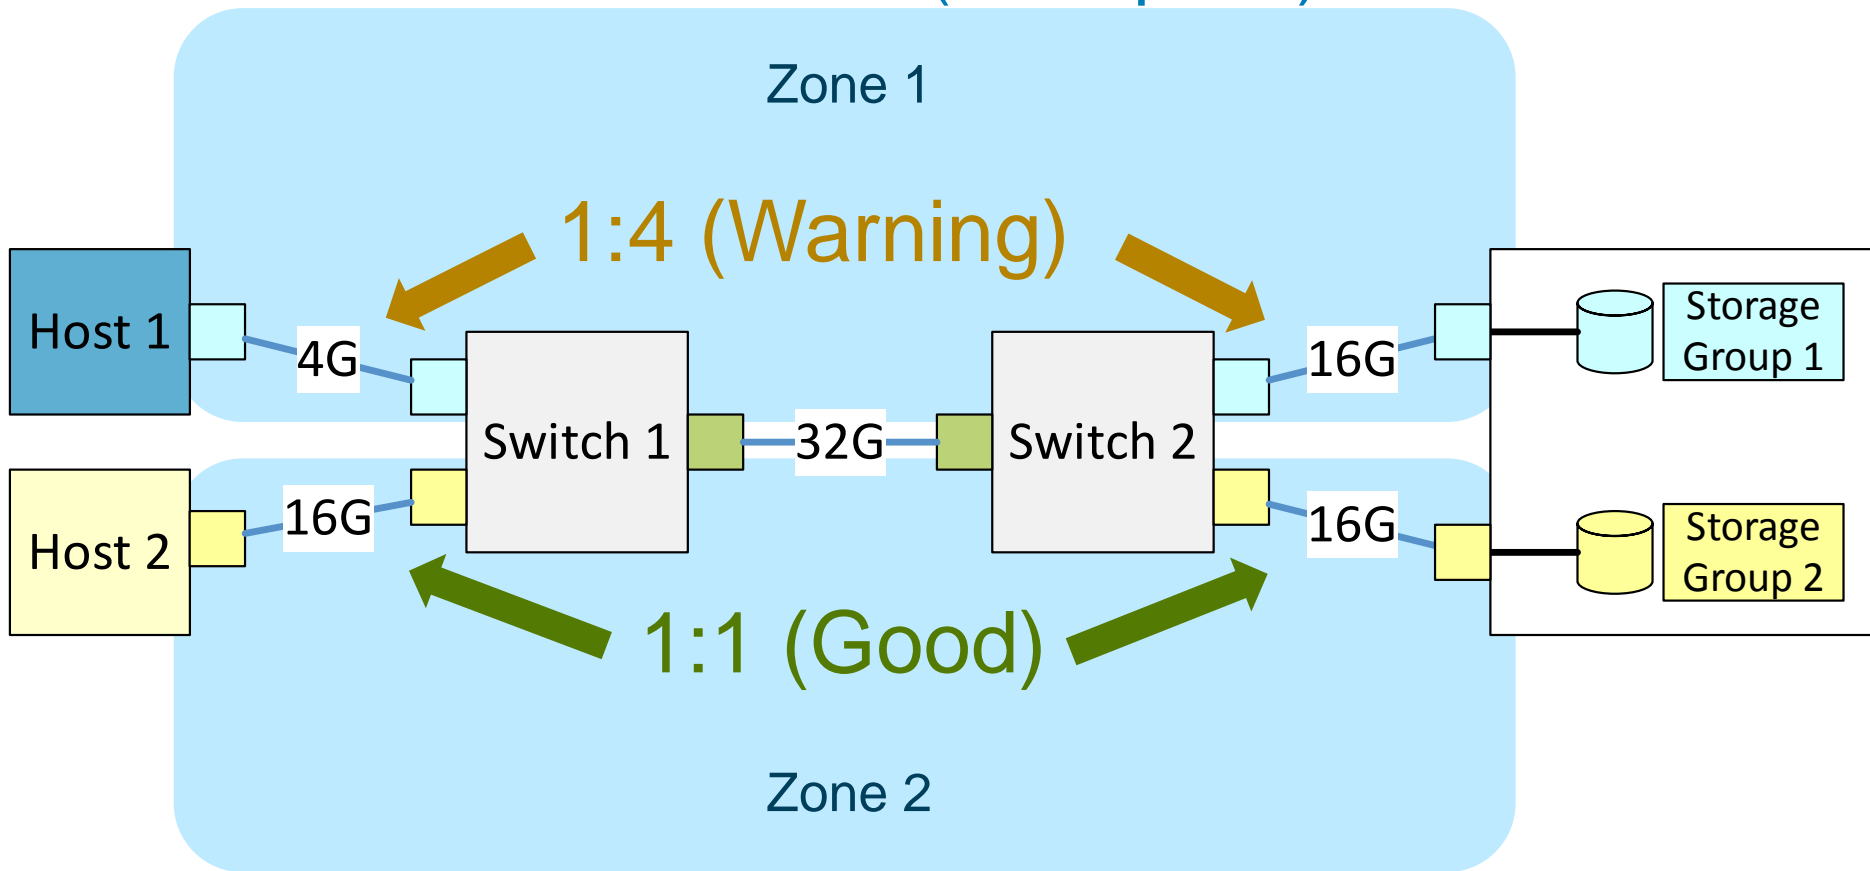
- But since the limit is applied to the SG or LUN, this becomes a problem when you have more than one host accessing it.
- The problem is, nothing prevents a single HBA from consuming all 800 MB/s



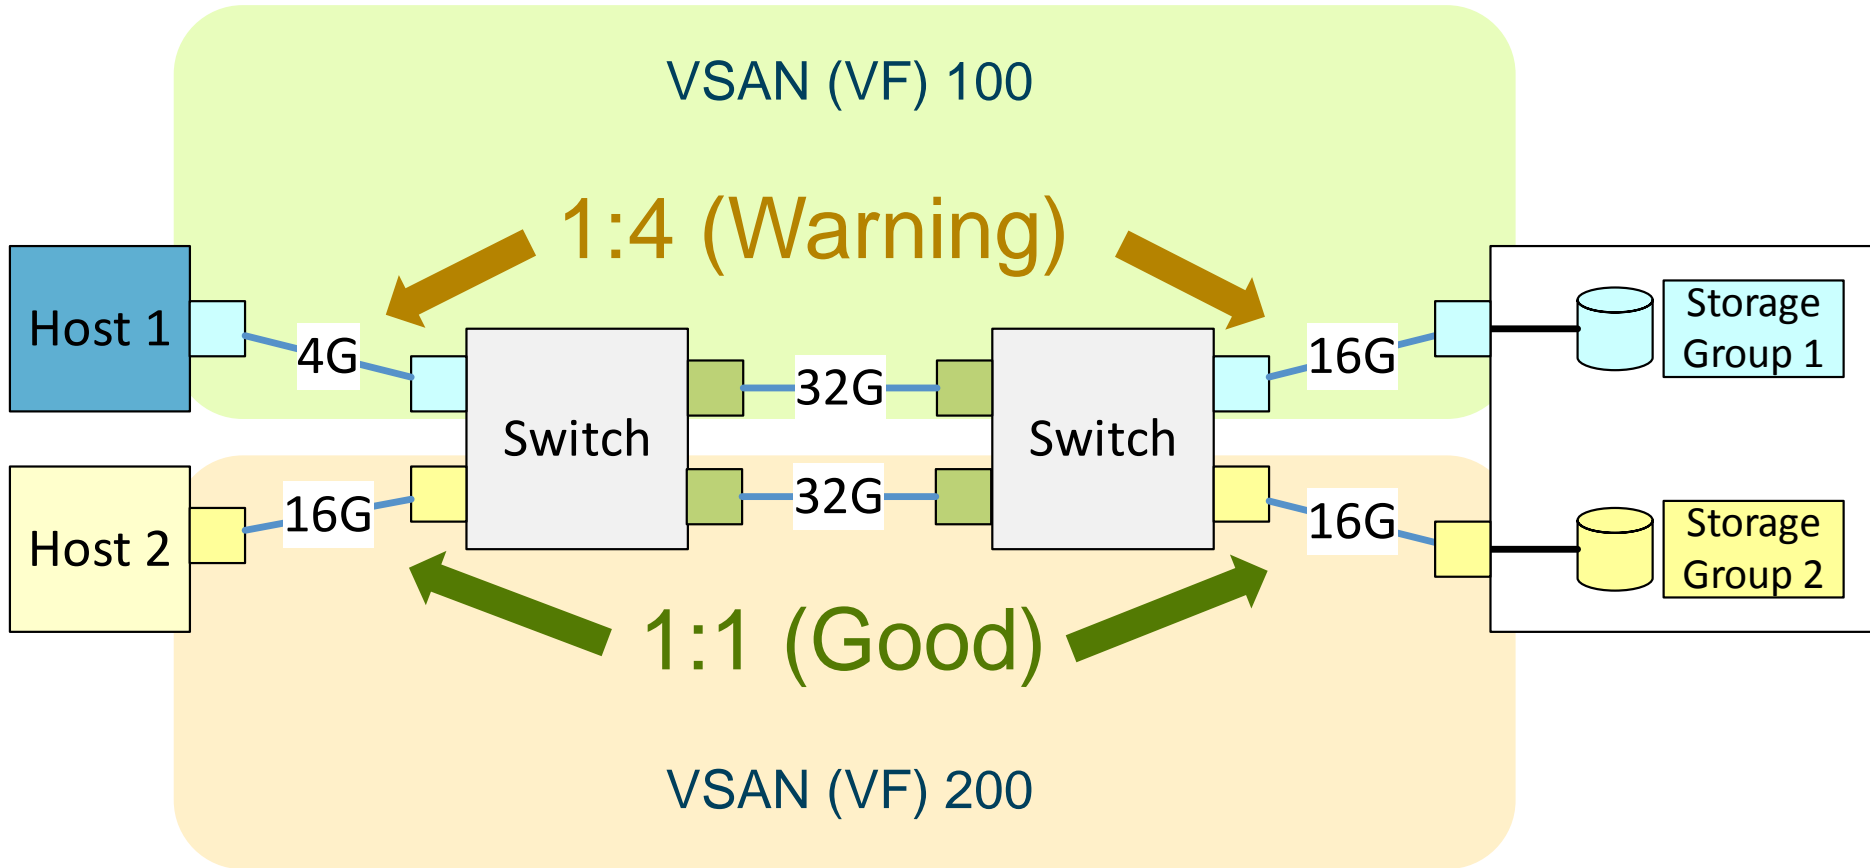
# Congestion Spreading Avoidance

	Prevention	Detection	Remediation
Short Term (today)	<ul style="list-style-type: none"><li>• BW Ratio</li><li>• Modernize SAN</li><li>• Fan-In/Out Ratio</li><li>• BW Limits</li><li>• Isolation</li></ul>		
Medium term			
Long term			

# BW Ratio (example 1)



# Isolation



# Congestion Spreading Avoidance

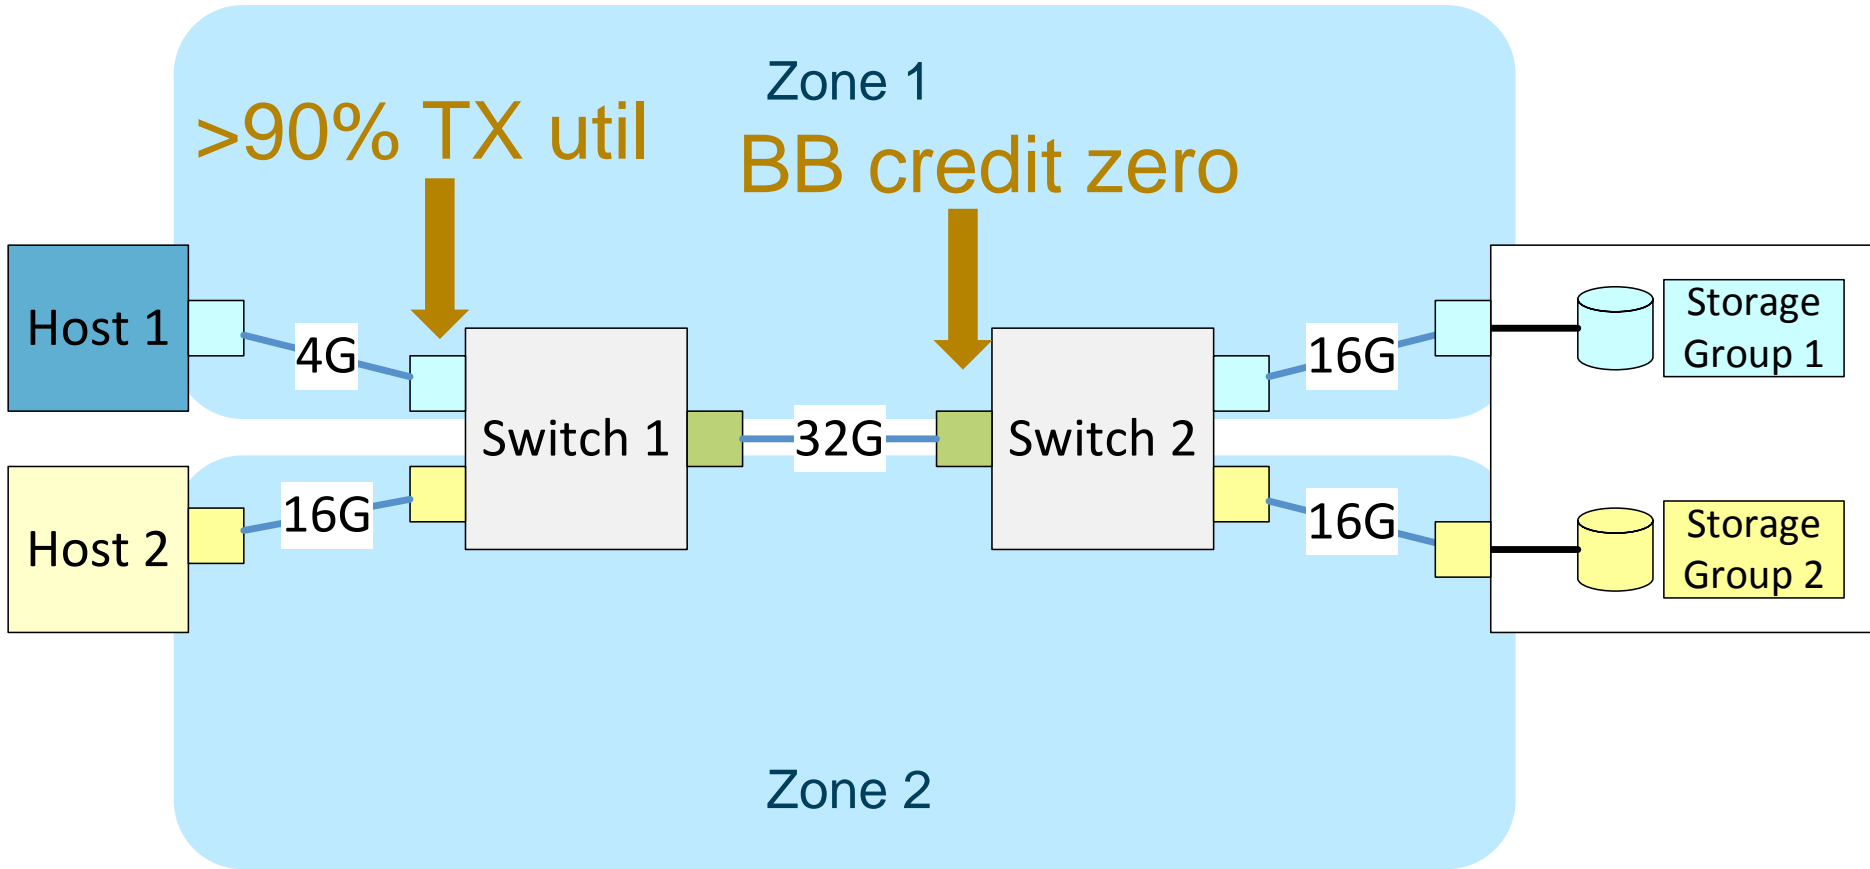
	Prevention	Detection	Remediation
Short Term (today)	<ul style="list-style-type: none"><li>• BW Ratio</li><li>• Modernize SAN</li><li>• Fan-In/Out Ratio</li><li>• BW Limits</li><li>• Isolation</li></ul>	<ul style="list-style-type: none"><li>• Ratio: bb_credit Zero(or TxWait) /tx frames</li></ul>	
Medium term			
Long term			

# Congestion Spreading Avoidance

	Prevention	Detection	Remediation
Short Term (today)	<ul style="list-style-type: none"><li>• BW Ratio</li><li>• Modernize SAN</li><li>• Fan-In/Out Ratio</li><li>• BW Limits</li><li>• Isolation</li></ul>	<ul style="list-style-type: none"><li>• Ratio: bb_credit Zero(or TxWait) /tx frames</li><li>• Brocade FPI/MAPS</li></ul>	
Medium term			
Long term			



# Brocade MAPS



## 419.927 - Top Port Traffic

Port	Min TX Traffic	TX Traffic	Max TX Traffic	Port
slot3 port47	418.26	419.927	421.111	sw
slot3 port44	400.015	401.607	402.971	sw
slot4 port4	203.864	206.821	209.896	sw
slot4 port5	202.238	205.23	208.395	sw
slot4 port7	201.845	204.829	207.714	sw
slot4 port6	201.751	204.743	207.513	sw
slot4 port4	0.103	0.104	0.106	sw
slot3 port46	0.053	0.054	0.054	sw
slot3 port47	0.051	0.051	0.051	sw
slot4 port5	0	0	0	sw

Refreshed- 1:22 PM

## 599200821 - BB Credit Zero

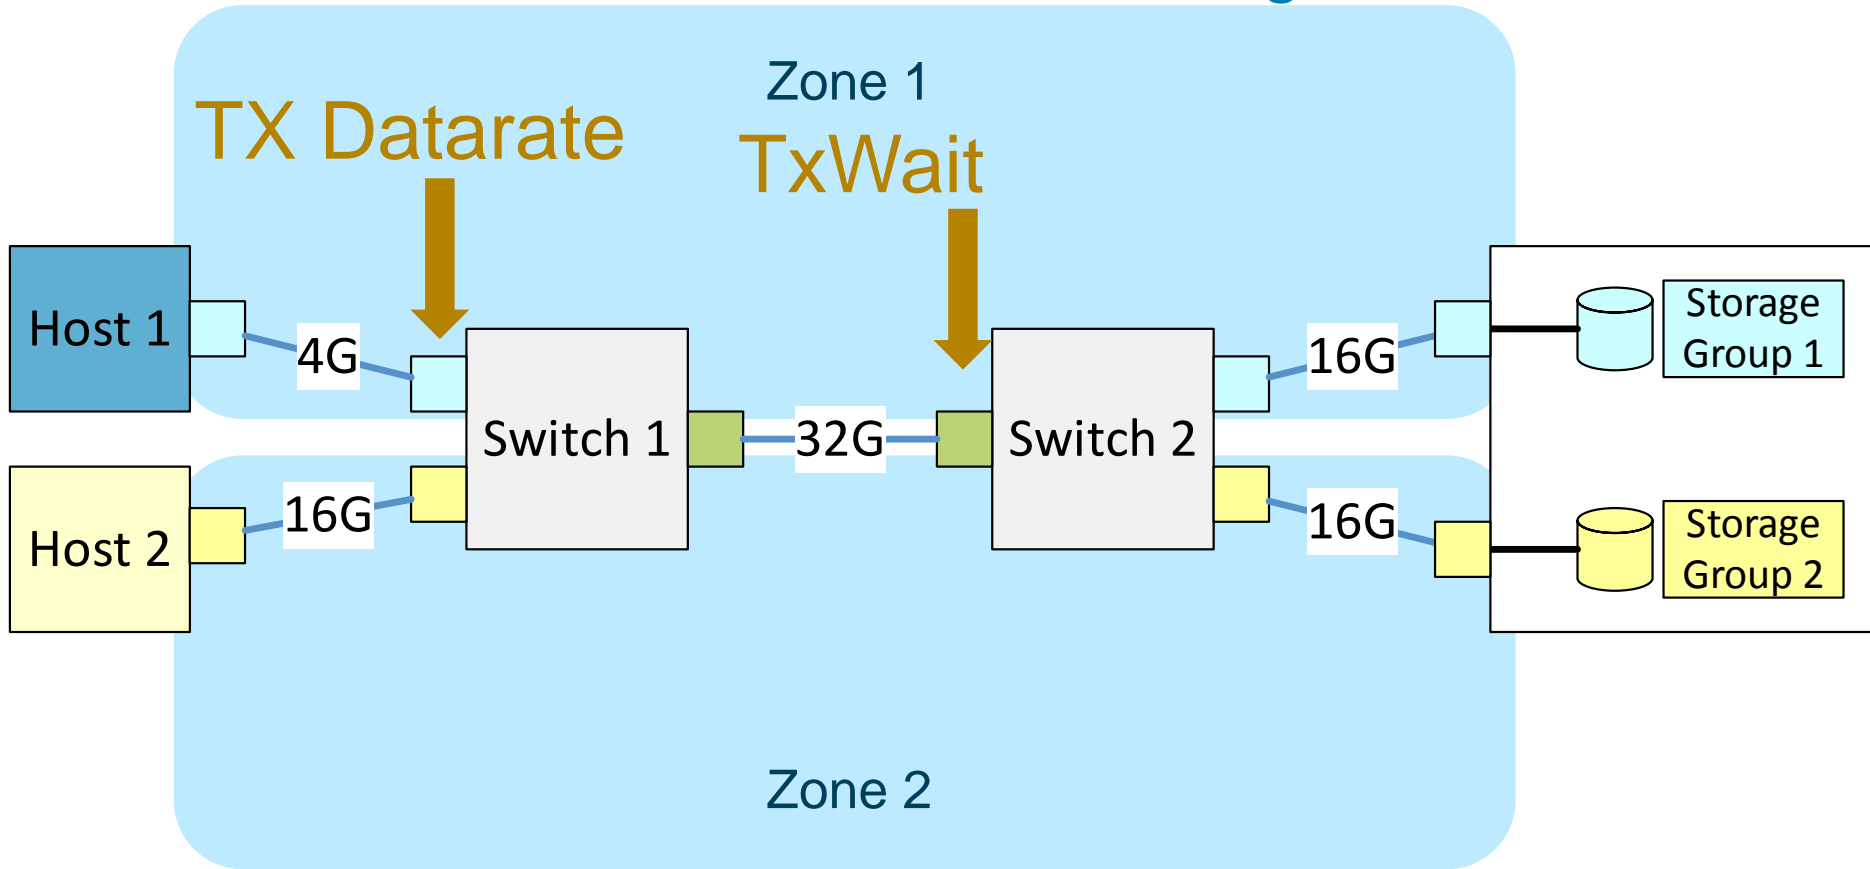
Port	Connected	BB Credit Zero	BB Credit Zero/sec	Product
slot4 port5	<a href="#">slot4 port5</a>	599200821	332889.345	switch
slot4 port6	<a href="#">slot4 port6</a>	599200821	332889.345	switch
slot4 port7	<a href="#">slot4 port7</a>	599200821	332889.345	switch
slot4 port4	<a href="#">slot4 port4</a>	598861625	332700.903	switch
slot3 port44	<a href="#">10:00:00:9...</a>	37470	20.817	switch

Refreshed- 1:26 PM

# Congestion Spreading Avoidance

	Prevention	Detection	Remediation
Short Term (today)	<ul style="list-style-type: none"><li>• BW Ratio</li><li>• Modernize SAN</li><li>• Fan-In/Out Ratio</li><li>• BW Limits</li><li>• Isolation</li></ul>	<ul style="list-style-type: none"><li>• Ratio: bb_credit Zero(or TxWait) /tx frames</li><li>• Brocade FPI/MAPS</li><li>• Cisco Port Monitor</li></ul>	
Medium term			
Long term			

# Cisco Port Monitoring





## Monitor / SAN / Slow Drain Analysis

## Collection Configuration

Scope: Fabric\_I0P054150

Duration: Once



Time Remaining: 00:00

☒ 10 Minutes ☐ 30 Minutes ☐ 1 Hour ☐ Other  Hours

Current jobs:



From: 2017-04-24 11:28:49 to: 2017-04-24 11:38:49

Zoom: 10 minutes 30 minutes MAX

## Slow Drain Details for Fabric\_I0P054150

Total 10

Show Non Zero data rows onl

	Type	Level 3			Level 2		TxWtAvg100...	RxB2Bto0	Level 1		
		TxCreditLoss	TxLinkReset	RxLinkReset	TxTimeoutD...	TxDiscard			TxB2Bto0	TxWaitInSec	TxWait(2.5us) ...
151 fc2/37	Switch	0	0	0	0	0	0	0	221654585	315.61	52.6012

# Congestion Spreading Avoidance

	Prevention	Detection	Remediation
Short Term (today)	<ul style="list-style-type: none"><li>• BW Ratio</li><li>• Modernize SAN</li><li>• Fan-In/Out Ratio</li><li>• BW Limits</li><li>• Isolation</li></ul>	<ul style="list-style-type: none"><li>• Ratio: bb_credit Zero(or TxWait) /tx frames</li><li>• Brocade MAPS</li><li>• Cisco Port Monitor</li></ul>	<ul style="list-style-type: none"><li>• Congestion Spreading KB</li></ul>
Medium term			
Long term			

# The Congestion Spreading Knowledge Map

## KB 464027



# Congestion Spreading Knowledge Map

- Knowledge Map – KB 464027
  - Summary of KB's and how they're related
- Congestion Spreading overview – KB 464206
  - Congestion Spreading types and how to detect them
- Known Issues – KB 464232
  - Emulex FW issue KB 464737:
    - › LPe11000/11002 (4Gb FC HBA): firmware version 2.82x14 or above
    - › LPe12000/12002 (8Gb FC HBA): firmware version 2.02a1 or above
    - › Not an issue with 16G+ HBAs
  - RecoverPoint FW issue: KB 453338
  - Congestion Spreading due to lost credit: KB 464246



# Congestion Spreading Avoidance

	Prevention	Detection	Remediation
Short Term (today)	<ul style="list-style-type: none"><li>• BW Ratio</li><li>• Modernize SAN</li><li>• Fan-In/Out Ratio</li><li>• BW Limits</li><li>• Isolation</li></ul>	<ul style="list-style-type: none"><li>• Ratio: <math>\text{bb\_credit} / \text{Zero(or TxWait) /tx frames}</math></li><li>• Brocade FPI/MAPS</li><li>• Cisco Port Monitor</li></ul>	<ul style="list-style-type: none"><li>• Congestion Spreading KB</li><li>• Brocade SDDQ</li><li>• Cisco Slow Drain Virtual Link Isolation</li></ul>
Medium term			
Long term			

# Congestion Spreading Avoidance

	Prevention	Detection	Remediation
Short Term (today)	<ul style="list-style-type: none"><li>• BW Ratio</li><li>• Modernize SAN</li><li>• Fan-In/Out Ratio</li><li>• BW Limits</li><li>• Isolation</li></ul>	<ul style="list-style-type: none"><li>• Ratio: <math>\text{bb\_credit} \text{ Zero(or TxWait) } / \text{tx frames}</math></li><li>• Brocade FPI/MAPS</li><li>• Cisco Port Monitor</li></ul>	<ul style="list-style-type: none"><li>• Congestion Spreading KB</li><li>• Brocade SDDQ</li><li>• Cisco Slow Drain Virtual Link Isolation</li><li>• Cisco no-credit-drop</li></ul>
Medium term			
Long term			

# Congestion Spreading Avoidance

	Prevention	Detection	Remediation
<b>Short Term (today)</b>	<ul style="list-style-type: none"><li>• BW Ratio</li><li>• Modernize SAN</li><li>• Fan-In/Out Ratio</li><li>• BW Limits</li><li>• Isolation</li></ul>	<ul style="list-style-type: none"><li>• Ratio: <math>\text{bb\_credit} / \text{Zero(or TxWait) /tx frames}</math></li><li>• Brocade FPI/MAPS</li><li>• Cisco Port Monitor</li></ul>	<ul style="list-style-type: none"><li>• Congestion Spreading KB</li><li>• Brocade SDDQ</li><li>• Cisco Slow Drain Virtual Link Isolation</li><li>• Cisco no-credit-drop</li><li>• Port fencing</li></ul>
<b>Medium term</b>			
<b>Long term</b>			

# Congestion Spreading Avoidance

	Prevention	Detection	Remediation
<b>Short Term (today)</b>	<ul style="list-style-type: none"><li>• BW Ratio</li><li>• Modernize SAN</li><li>• Fan-In/Out Ratio</li><li>• BW Limits</li><li>• Isolation</li></ul>	<ul style="list-style-type: none"><li>• Ratio: <math>\text{bb\_credit} / \text{Zero(or TxWait) /tx frames}</math></li><li>• Brocade FPI/MAPS</li><li>• Cisco Port Monitor</li></ul>	<ul style="list-style-type: none"><li>• Congestion Spreading KB</li><li>• Brocade SDDQ</li><li>• Cisco Slow Drain Virtual Link Isolation</li><li>• Cisco no-credit-drop</li><li>• Port fencing</li></ul>
<b>Medium term</b>	<ul style="list-style-type: none"><li>• BW Limit Enhancements</li></ul>		
<b>Long term</b>			

# Congestion Spreading Avoidance

	Prevention	Detection	Remediation
<b>Short Term (today)</b>	<ul style="list-style-type: none"><li>• BW Ratio</li><li>• Modernize SAN</li><li>• Fan-In/Out Ratio</li><li>• BW Limits</li><li>• Isolation</li></ul>	<ul style="list-style-type: none"><li>• Ratio: <math>\text{bb\_credit} / \text{Zero(or TxWait) /tx frames}</math></li><li>• Brocade FPI/MAPS</li><li>• Cisco Port Monitor</li></ul>	<ul style="list-style-type: none"><li>• Congestion Spreading KB</li><li>• Brocade SDDQ</li><li>• Cisco Slow Drain Virtual Link Isolation</li><li>• Cisco no-credit-drop</li><li>• Port fencing</li></ul>
<b>Medium term</b>	<ul style="list-style-type: none"><li>• BW Limit Enhancements</li><li>• ViPR-C</li></ul>		
<b>Long term</b>			

# Congestion Spreading Avoidance

	Prevention	Detection	Remediation
<b>Short Term (today)</b>	<ul style="list-style-type: none"><li>• BW Ratio</li><li>• Modernize SAN</li><li>• Fan-In/Out Ratio</li><li>• BW Limits</li><li>• Isolation</li></ul>	<ul style="list-style-type: none"><li>• Ratio: <math>\text{bb\_credit} / \text{Zero(or TxWait) /tx frames}</math></li><li>• Brocade FPI/MAPS</li><li>• Cisco Port Monitor</li></ul>	<ul style="list-style-type: none"><li>• Congestion Spreading KB</li><li>• Brocade SDDQ</li><li>• Cisco Slow Drain Virtual Link Isolation</li><li>• Cisco no-credit-drop</li><li>• Port fencing</li></ul>
<b>Medium term</b>	<ul style="list-style-type: none"><li>• BW Limit Enhancements</li><li>• ViPR-C</li></ul>	<ul style="list-style-type: none"><li>• ViPR-SRM</li></ul>	
<b>Long term</b>			

# Congestion Spreading Avoidance

	Prevention	Detection	Remediation
<b>Short Term (today)</b>	<ul style="list-style-type: none"><li>• BW Ratio</li><li>• Modernize SAN</li><li>• Fan-In/Out Ratio</li><li>• BW Limits</li><li>• Isolation</li></ul>	<ul style="list-style-type: none"><li>• Ratio: <math>\text{bb\_credit} / \text{Zero(or TxWait) /tx frames}</math></li><li>• Brocade FPI/MAPS</li><li>• Cisco Port Monitor</li></ul>	<ul style="list-style-type: none"><li>• Congestion Spreading KB</li><li>• Brocade SDDQ</li><li>• Cisco Slow Drain Virtual Link Isolation</li><li>• Cisco no-credit-drop</li><li>• Port fencing</li></ul>
<b>Medium term</b>	<ul style="list-style-type: none"><li>• BW Limit Enhancements</li><li>• ViPR-C</li></ul>	<ul style="list-style-type: none"><li>• ViPR-SRM</li></ul>	<ul style="list-style-type: none"><li>• ViPR-SRM &amp; ViPR-C</li></ul>
<b>Long term</b>			

# Congestion Spreading Avoidance

	Prevention	Detection	Remediation
<b>Short Term (today)</b>	<ul style="list-style-type: none"><li>• BW Ratio</li><li>• Modernize SAN</li><li>• Fan-In/Out Ratio</li><li>• BW Limits</li><li>• Isolation</li></ul>	<ul style="list-style-type: none"><li>• Ratio: bb_credit Zero(or TxWait) /tx frames</li><li>• Brocade FPI/MAPS</li><li>• Cisco Port Monitor</li></ul>	<ul style="list-style-type: none"><li>• Congestion Spreading KB</li><li>• Brocade SDDQ</li><li>• Cisco Slow Drain Virtual Link Isolation</li><li>• Cisco no-credit-drop</li><li>• Port fencing</li></ul>
<b>Medium term</b>	<ul style="list-style-type: none"><li>• BW Limit Enhancements</li><li>• ViPR-C</li></ul>	<ul style="list-style-type: none"><li>• ViPR-SRM</li></ul>	<ul style="list-style-type: none"><li>• ViPR-SRM &amp; ViPR-C</li></ul>
<b>Long term</b>	Future Fibre Channel Enhancements (TBD)		




# AGENDA

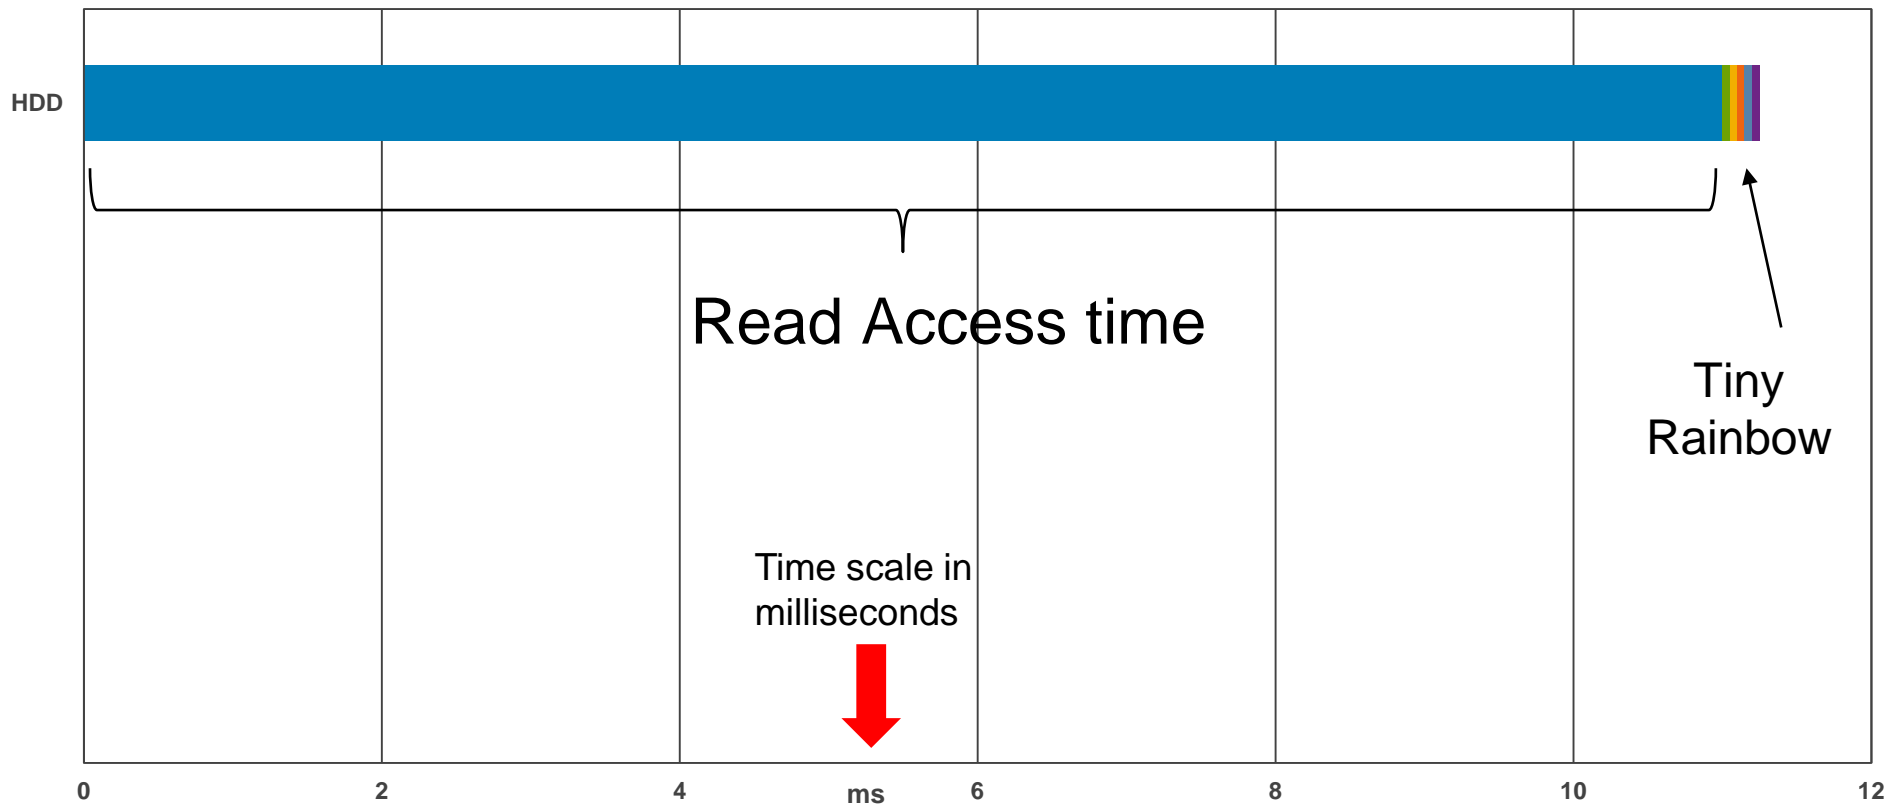
- What's new with storage connectivity
  - Protocol adoption
  - Connectrix Product Updates
- Congestion Spreading and its impact
  - Congestion Spreading terminology
  - Congestion Spreading & Innocent flows
- SAN Best practices for All Flash Arrays (AFAs)
- **Introduction to NVMe**
  - FC-NVMe

# Introduction to NVMe

# HDD vs SDD



# HDD vs SDD




■ Read Time ■ Transfer Time ■ Misc SSD ■ Link Transfer ■ Platform & Adapter ■ Software

Source: Intel Measurements

Source: Jim Handy, SDC 2015

[https://www.snia.org/sites/default/files/SDC15\\_presentations/gen\\_sessions/JimHandy\\_The\\_Long-Term\\_Future\\_PrintVersion.pdf](https://www.snia.org/sites/default/files/SDC15_presentations/gen_sessions/JimHandy_The_Long-Term_Future_PrintVersion.pdf)

# HDD vs SSD




■ Read Time ■ Transfer Time ■ Misc SSD ■ Link Transfer ■ Platform & Adapter ■ Software

Source: Intel Measurements

Source: Jim Handy, SDC 2015

[https://www.snia.org/sites/default/files/SDC15\\_presentations/gen\\_sessions/JimHandy\\_The\\_Long-Term\\_Future\\_PrintVersion.pdf](https://www.snia.org/sites/default/files/SDC15_presentations/gen_sessions/JimHandy_The_Long-Term_Future_PrintVersion.pdf)

# HDD vs SSD




■ Read Time ■ Transfer Time ■ Misc SSD ■ Link Transfer ■ Platform & Adapter ■ Software

Source: Intel Measurements

Source: Jim Handy, SDC 2015

[https://www.snia.org/sites/default/files/SDC15\\_presentations/gen\\_sessions/JimHandy\\_The\\_Long-Term\\_Future\\_PrintVersion.pdf](https://www.snia.org/sites/default/files/SDC15_presentations/gen_sessions/JimHandy_The_Long-Term_Future_PrintVersion.pdf)

# HDD vs SSD



■ Read Time ■ Transfer Time ■ Misc. SSD ■ Link Transfer ■ Platform & Adapter ■ Software


Source: Intel Measurements

Source: Jim Handy, SDC 2015

[https://www.snia.org/sites/default/files/SDC15\\_presentations/gen\\_sessions/JimHandy\\_The\\_Long-Term\\_Future\\_PrintVersion.pdf](https://www.snia.org/sites/default/files/SDC15_presentations/gen_sessions/JimHandy_The_Long-Term_Future_PrintVersion.pdf)

# SSD Access times

MLC NAND SATA 3 ONFi 2



■ Read Time ■ Transfer Time ■ Misc SSD ■ Link Transfer ■ Platform & Adapter ■ Software


Source: Intel Measurements

Source: Jim Handy, SDC 2015

[https://www.snia.org/sites/default/files/SDC15\\_presentations/gen\\_sessions/JimHandy\\_The\\_Long-Term\\_Future\\_PrintVersion.pdf](https://www.snia.org/sites/default/files/SDC15_presentations/gen_sessions/JimHandy_The_Long-Term_Future_PrintVersion.pdf)



## SSD Access times




■ Read Time ■ Transfer Time ■ Misc SSD ■ Link Transfer ■ Platform & Adapter ■ Software

Source: Intel Measurements

Source: Jim Handy, SDC 2015

[https://www.snia.org/sites/default/files/SDC15\\_presentations/gen\\_sessions/JimHandy\\_The\\_Long-Term\\_Future\\_PrintVersion.pdf](https://www.snia.org/sites/default/files/SDC15_presentations/gen_sessions/JimHandy_The_Long-Term_Future_PrintVersion.pdf)

## SSD Access times

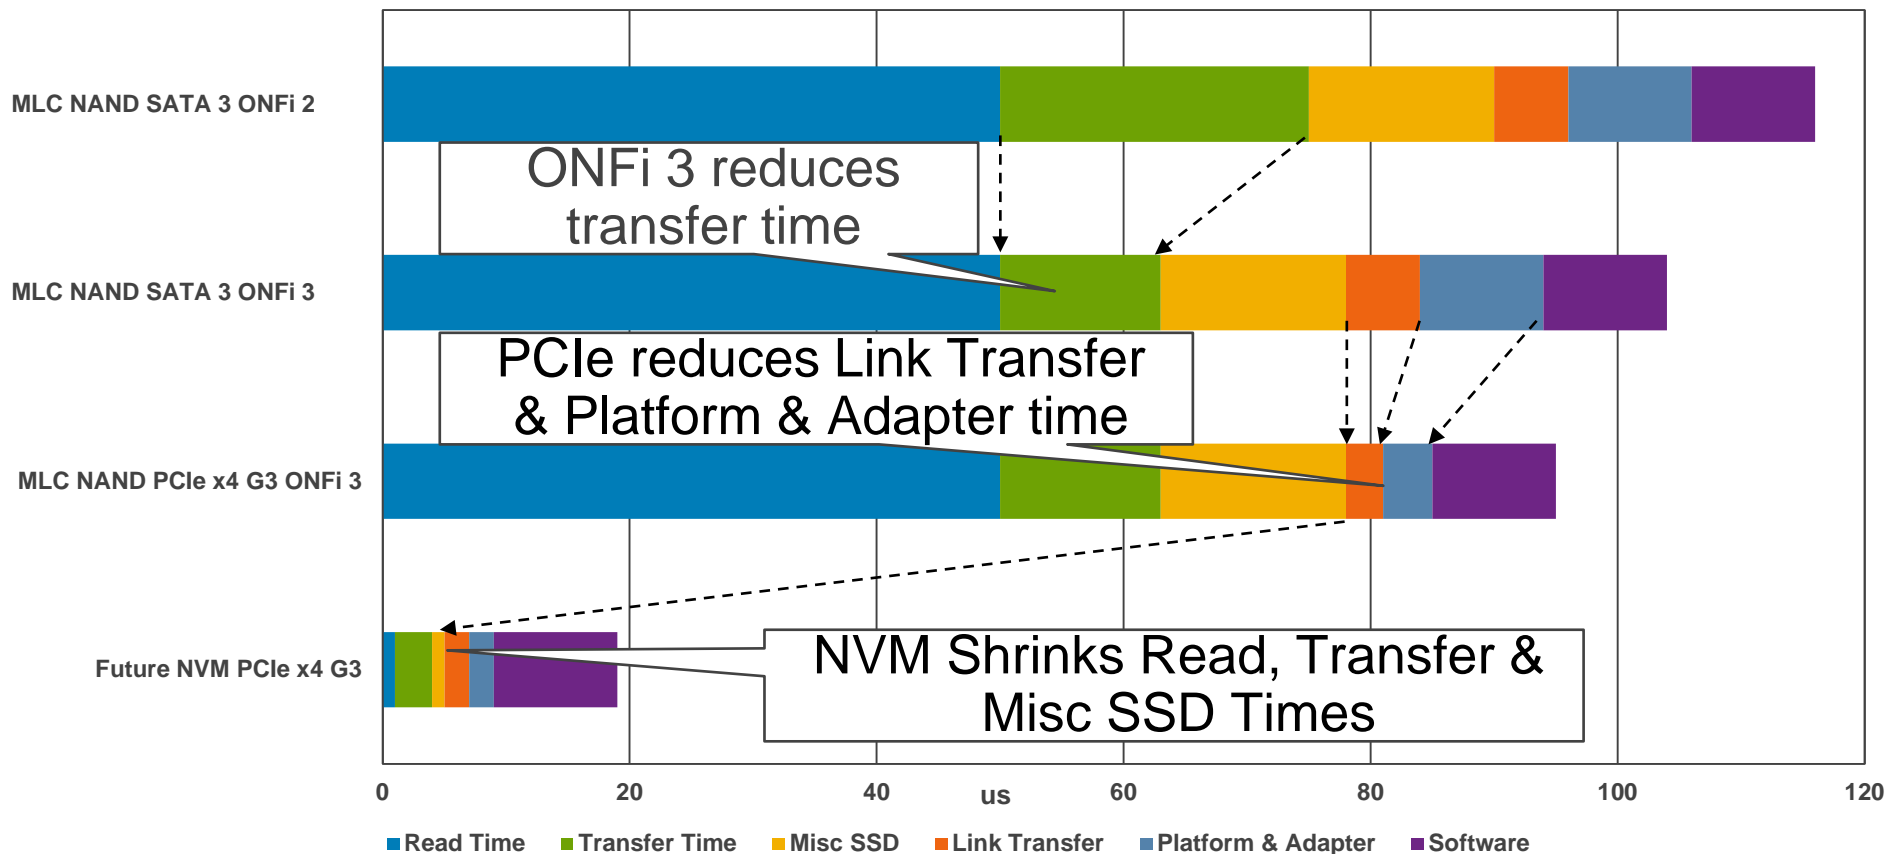


Source: Intel Measurements

Source: Jim Handy, SDC 2015

[https://www.snia.org/sites/default/files/SDC15\\_presentations/gen\\_sessions/JimHandy\\_The\\_Long-Term\\_Future\\_PrintVersion.pdf](https://www.snia.org/sites/default/files/SDC15_presentations/gen_sessions/JimHandy_The_Long-Term_Future_PrintVersion.pdf)

## SSD Access times



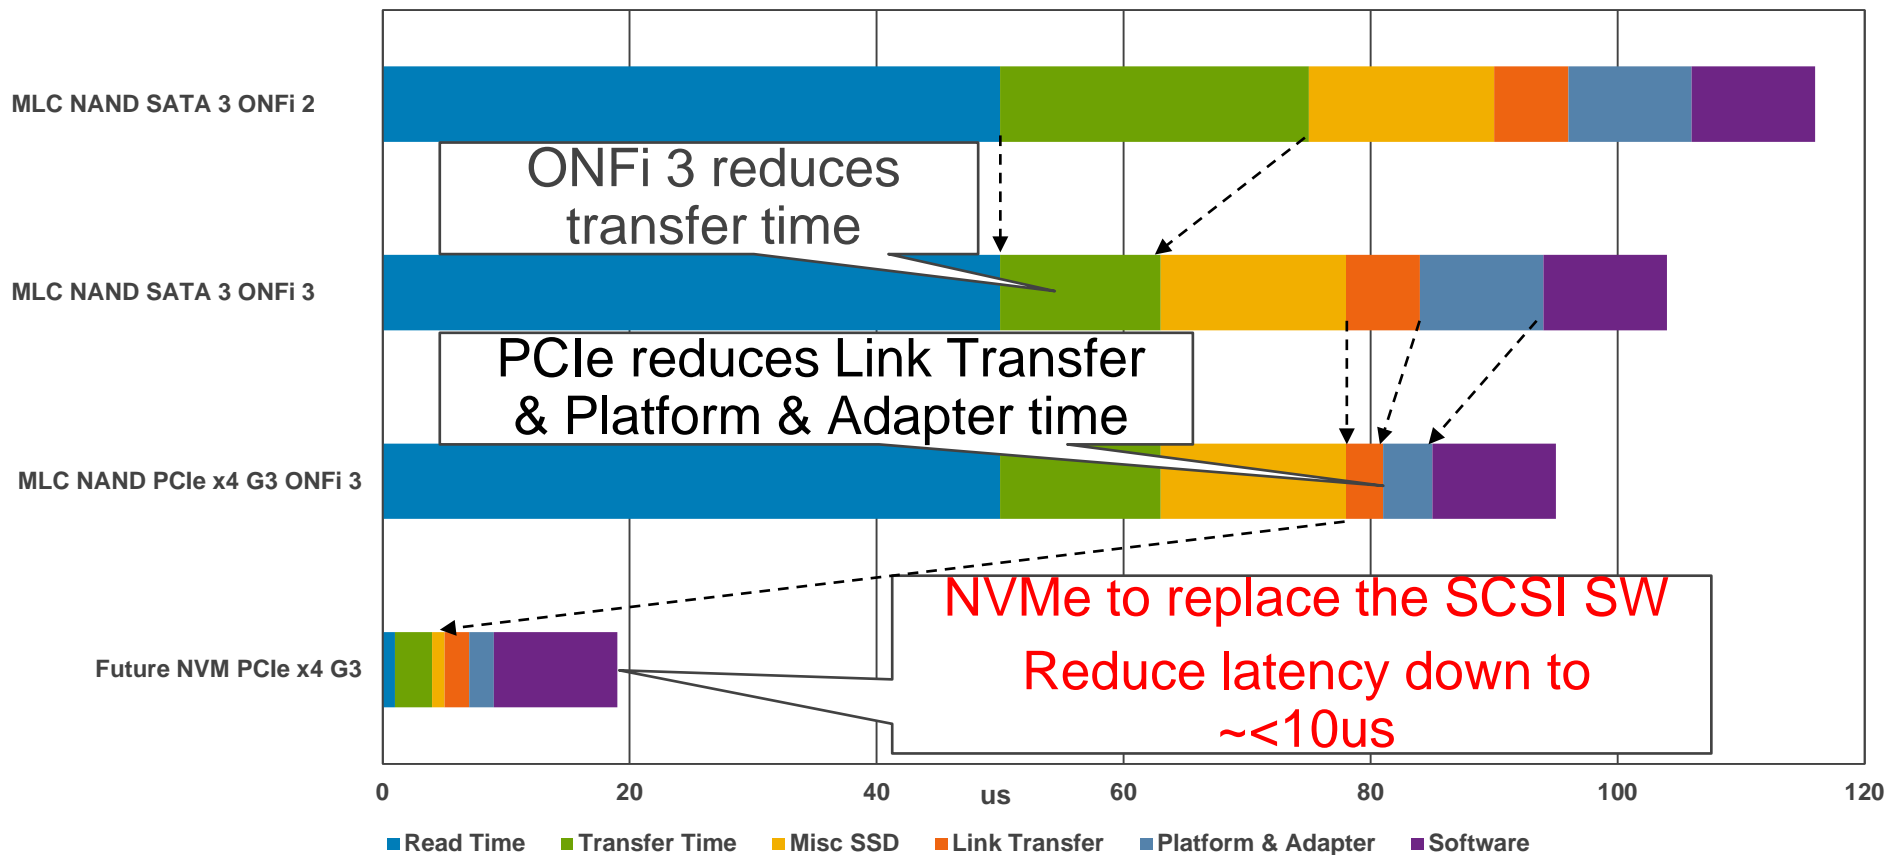
■ Read Time ■ Transfer Time ■ Misc SSD ■ Link Transfer ■ Platform & Adapter ■ Software

Source: Intel Measurements

Source: Jim Handy, SDC 2015

[https://www.snia.org/sites/default/files/SDC15\\_presentations/gen\\_sessions/JimHandy\\_The\\_Long-Term\\_Future\\_PrintVersion.pdf](https://www.snia.org/sites/default/files/SDC15_presentations/gen_sessions/JimHandy_The_Long-Term_Future_PrintVersion.pdf)

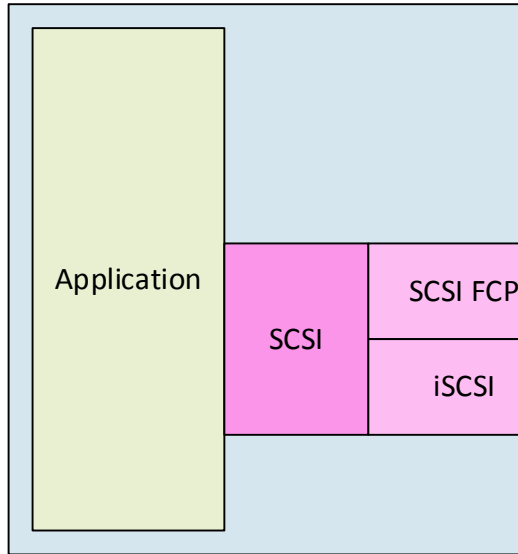
## SSD Access times



# Agenda

- What's new with storage connectivity
  - Protocol adoption
  - Connectrix Product Updates
- Congestion Spreading and its impact
  - Congestion Spreading terminology
  - Congestion Spreading & Innocent flows
- SAN Best practices for All Flash Arrays (AFAs)
- **Introduction to NVMe**
  - FC-NVMe

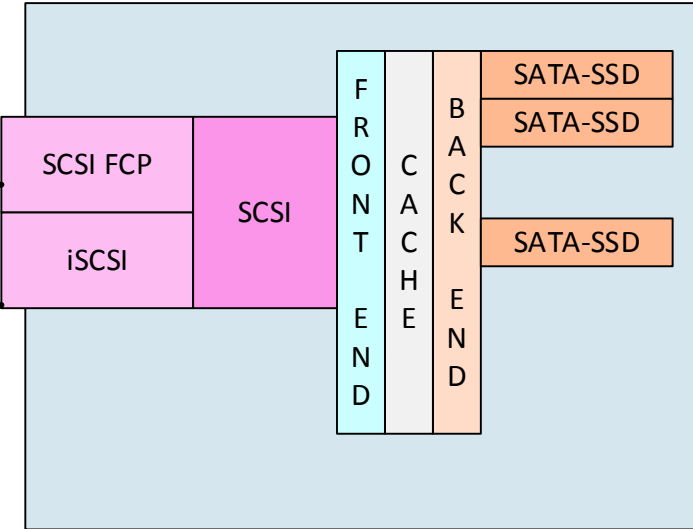
## Host



Fibre Channel

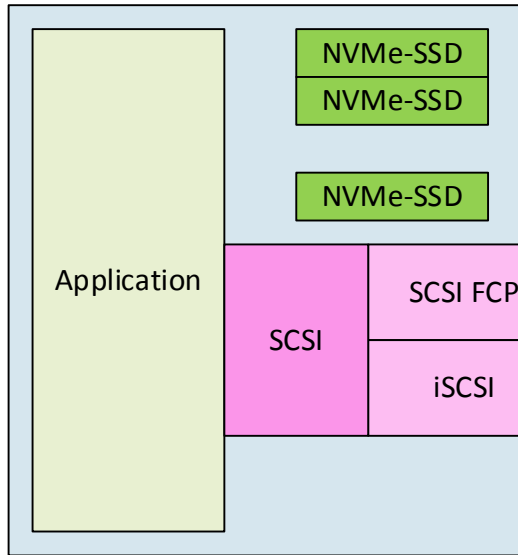
IP/Ethernet

## Storage Array





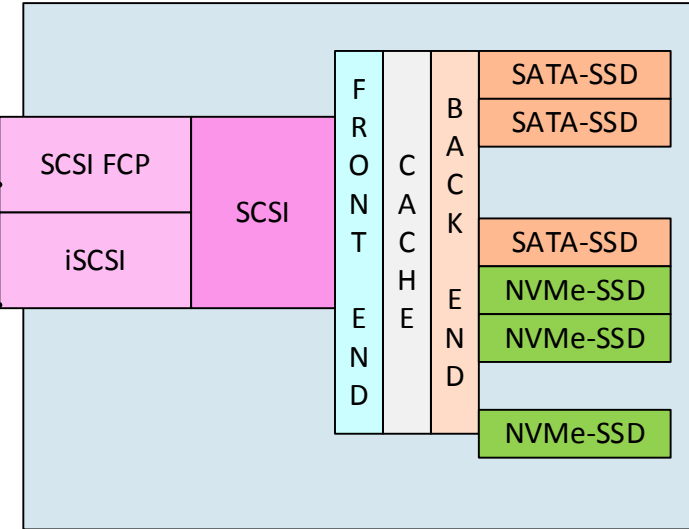
## Host



Fibre Channel

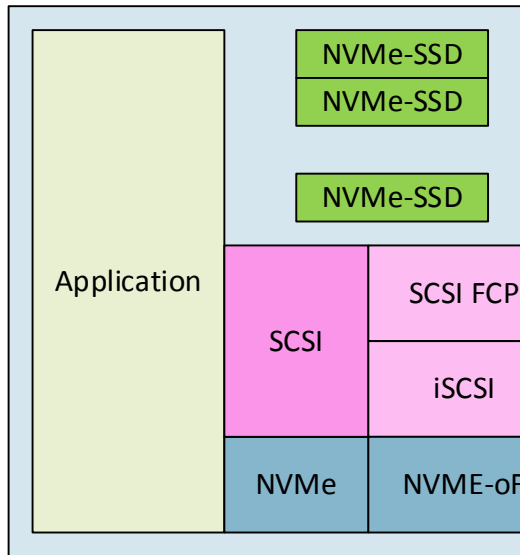
IP/Ethernet

## Storage Array

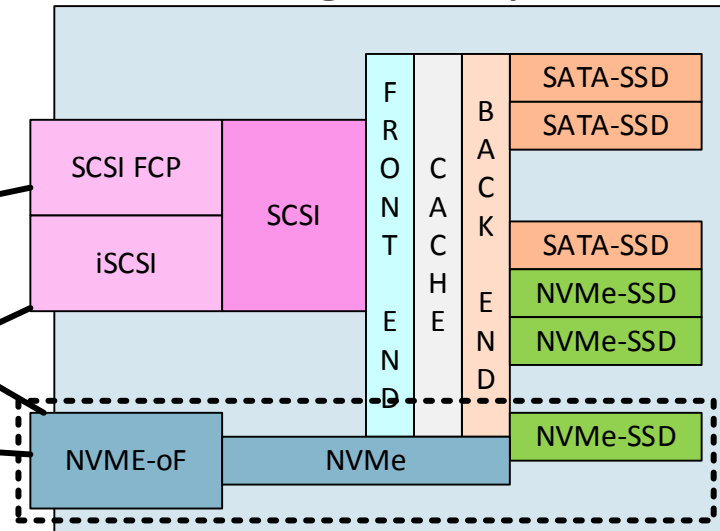




## Host



## Storage Array



Please Visit  
Brocade & Cisco  
Booths for  
FC-NVMe  
Demos

# Want to win a Levitating Death Star Speaker?

- Follow **@DellEMCStorage** while at Dell EMC World
- **2 Winners** will be chosen **daily** from Monday May 8 through Thursday May 11
- All winners will be **notified through Twitter** Direct Message



NO PURCHASE NECESSARY. Ends 05/11/2017. To enter and for Official Rules, visit <http://thecoreblog.emc.com/dell-emc-world-follow-win-sweepstakes-2017/>

## JOIN US!

- Birds of a feather (BOF)
  - Tuesday @ 1:30pm
- Second Session
  - Thursday @ 8:30am
- Drop by the Connectrix Booth
- Schedule your 1x1 personal sessions today!

# THANK YOU!

## Q & A

Erik Smith

[erik.smith@dell.com](mailto:erik.smith@dell.com)

@provandal

[www.brasstacksblog.typepad.com](http://www.brasstacksblog.typepad.com)

Alan Rajapa

[alan.rajapa@dell.com](mailto:alan.rajapa@dell.com)

# Sources/References

# Sources/References

- Erik Smith: Slow Drains are impacting your SAN
  - <http://brasstacksblog.typepad.com/brass-tacks/2016/02/slow-drains-are-impacting-your-san.html>
  - <https://www.brighttalk.com/webcast/13935/213779/slow-drains-are-impacting-your-san>
- Erik Smith: EMC World 2016 Presentation:
  - [https://www.emcworldonline.com/2016/connect/sessionDetail.ww?SESSION\\_ID=2539](https://www.emcworldonline.com/2016/connect/sessionDetail.ww?SESSION_ID=2539)
- Dell EMC Congestion Spreading knowledge Map
  - [https://emcservice--c.na55.visual.force.com/apex/KB\\_BreakFix\\_clone?id=kA2j0000000R59Q](https://emcservice--c.na55.visual.force.com/apex/KB_BreakFix_clone?id=kA2j0000000R59Q)
- Brocade MAPS Configuration Guide
  - <http://www.brocade.com/content/html/en/configuration-guide/fos-80x-maps/GUID-426E1CD4-3763-419D-9D54-91F824F463EB-homepage.html>
- Cisco Slow Drain Device Detection white paper
  - <http://www.cisco.com/c/dam/en/us/products/collateral/storage-networking/mds-9700-series-multilayer-directors/whitepaper-c11-737315.pdf>
- Dell EMC VMAX 3 and VMAX AFA I/O Limits
  - <https://www.emc.com/collateral/white-papers/h15498-dell-emc-vmax3-and-vmax-all-flash-quality-of-service-controls-for-multi-tenant-environments.pdf>
- Jim Handy- SDC 2015
  - [https://www.snia.org/sites/default/files/SDC15\\_presentations/gen\\_sessions/JimHandy\\_The\\_Long-Term\\_Future\\_PrintVersion.pdf](https://www.snia.org/sites/default/files/SDC15_presentations/gen_sessions/JimHandy_The_Long-Term_Future_PrintVersion.pdf)

# Connectrix Video Resources

- **Title:** Connectrix Optic Cable Cleaning Made Easy (2:31)

**Description:** In this video, you will learn how easy it is to clean your fiber optic cables to prevent connectivity issues in your Connectrix products

<https://www.youtube.com/watch?v=u1rgjzUQIDU&index=2&list=PL0CFE7365EB8D1577>



- **Title:** Connectrix SFP Optic Replacement Made Easy (1:51)

**Description:** In this video, you will learn how easy it is to replace your failed SFP optic in your Connectrix products

<https://www.youtube.com/watch?v=4kd6O0S7Fsg&index=1&list=PL0CFE7365EB8D1577>





# REALIZE

DELL EMC / World