

#SHAREorg



Under the Hood of Ficon



Dale Riedy
IBM®
riedy@us.ibm.com

7 August 2012
Session 11716



Legal Stuff

- Notice
 - IBM may have patents or pending patent applications covering subject matter described in this document. The furnishing of this document does not give you any license to these patents. You can send license inquiries, in writing to: *IBM Director of Licensing, IBM Corporation, North Castle Drive, Armonk, NY 10504-1785 U.S.A.*
 - Any references in this information to non-IBM Web sites are provided for convenience only and do not in any manner serve as an endorsement of those Web sites. The materials at those Web sites are not part of the materials for this IBM product and use of those Web sites is at your own risk.
- Trademarks
 - The following terms are trademarks of the International Business Machines Corporation in the United States, other countries, or both: FICON® IBM® Redbooks™ System z10™ z/OS® zSeries® z10™
 - Other Company, product, or service names may be trademarks or service marks of others.

Agenda



Frames, Sequences and Exchanges

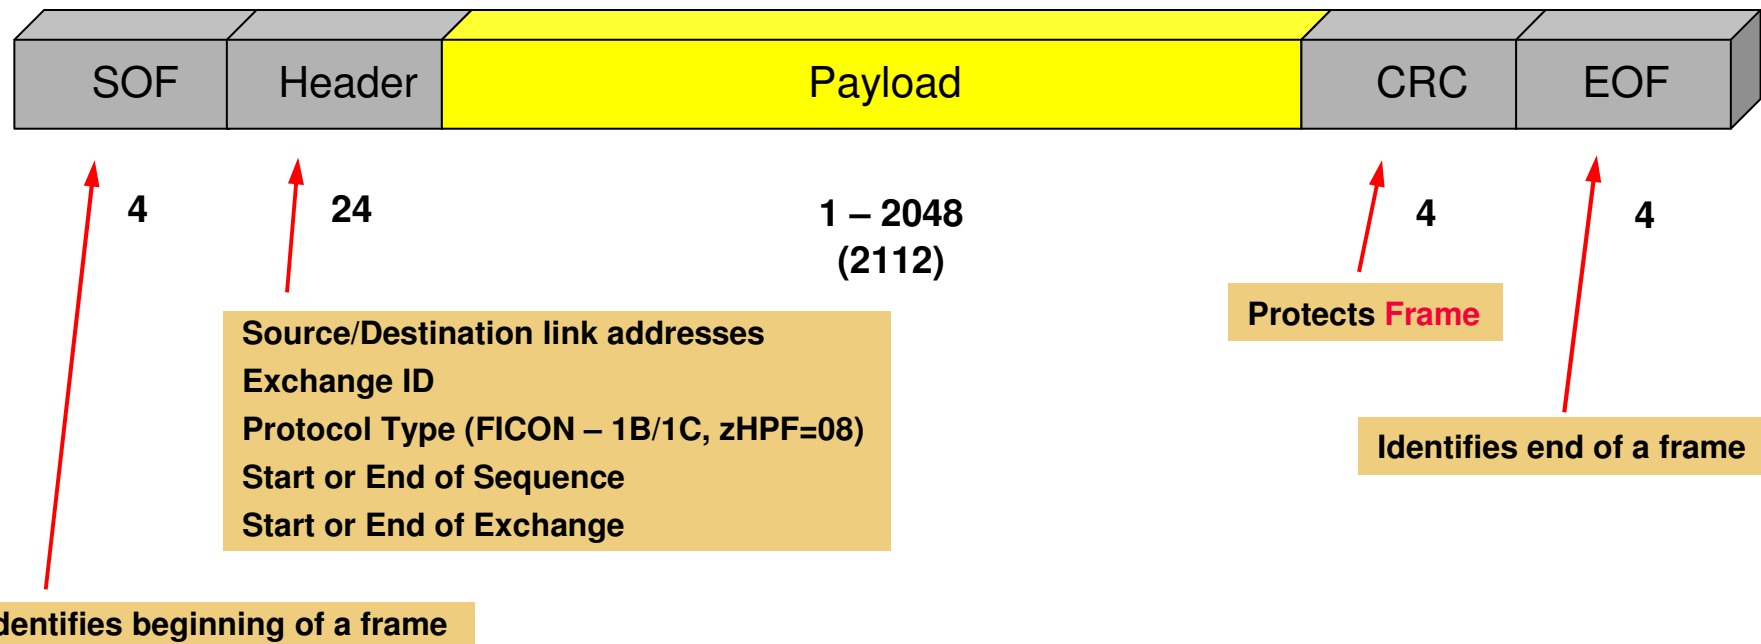
IDAW vs MIDAW

Buffer to Buffer Credit

Ficon Recovery

Fibre Channel Frame

The basic building block is the **FRAME**




Sequences and IUs

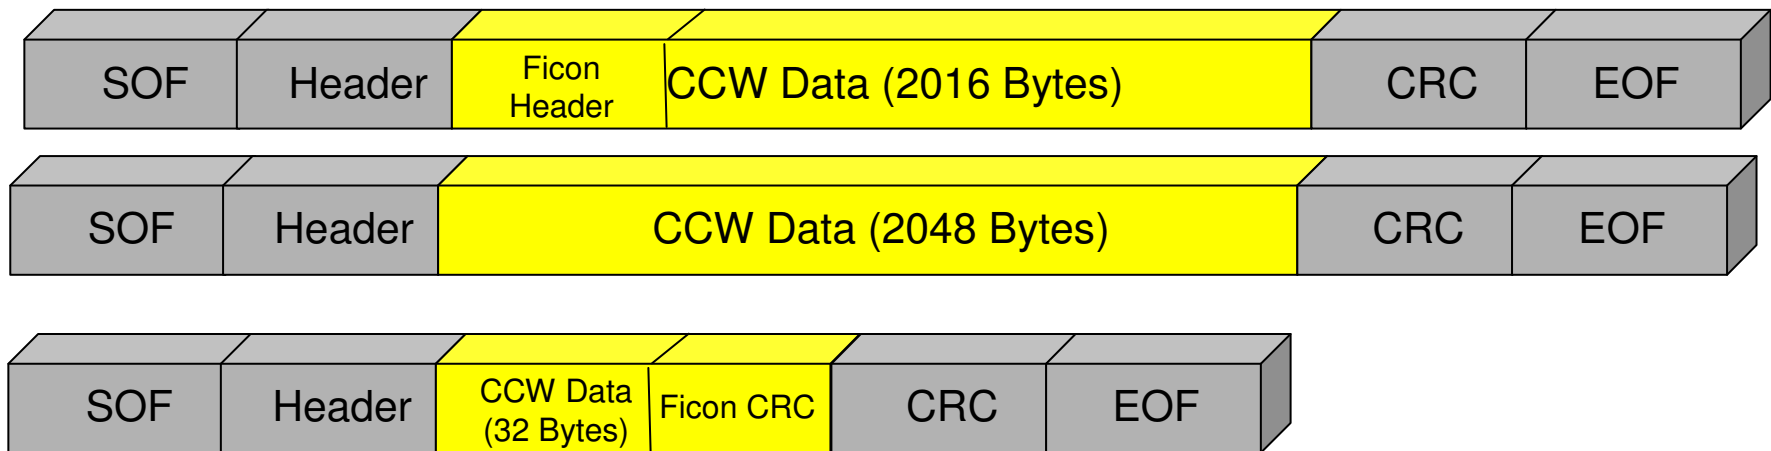
- Each Upper Layer Protocol (ULP) defines the contents and format of it's own **Information Units** (IUs)
 - Commands
 - Data
 - Status
 - Control
 - Etc
- Ficon IUs can be up to 8K (8192) in size
 - 8160 (8K-32) bytes of data
 - 32 bytes contain Ficon Header information
 - 4 frames are needed for the largest IU
- The collection of frame(s) that make up a IU are called a **Sequence**
 - A Sequence may be as small as a single Frame

Ficon IU Examples

1 Frame IU to transfer a Read CCW



3 Frame IU to transfer 4K of data




Exchanges

- Fibre Channel Architecture defines an **Exchange** as
 - “A mechanism for identifying and managing an operation between two ports“
- All IUs (a.k.a. Sequences) that make up a single I/O operation are part of an **Exchange**

Ficon Exchanges

- In Ficon, each concurrent I/O operation uses two Exchanges
 - One unidirectional Exchange for IUs from the Channel to the CU
 - A different unidirectional Exchange for IUs from the CU to the Channel
- The PAIR is commonly know as a “Ficon Exchange”



How many Exchanges do I need?

- Little's Law states:
 - *The number of “things” in a system can be determined by multiplying the average arrival rate of those “things” by the average time each “thing” stays in the system*
- Applied to Ficon:
 - The average number of Exchanges active at any given time = Average I/O rate * Average response time
 - Example: 5000 Ficon I/Os / Second on a given channel with 0.4ms service time¹ needs 2 Active Exchanges (pairs) at any given time

¹ The amount of time the I/O is active in the channel

Agenda

Frames, Sequences and Exchanges



IDAW vs MIDAW

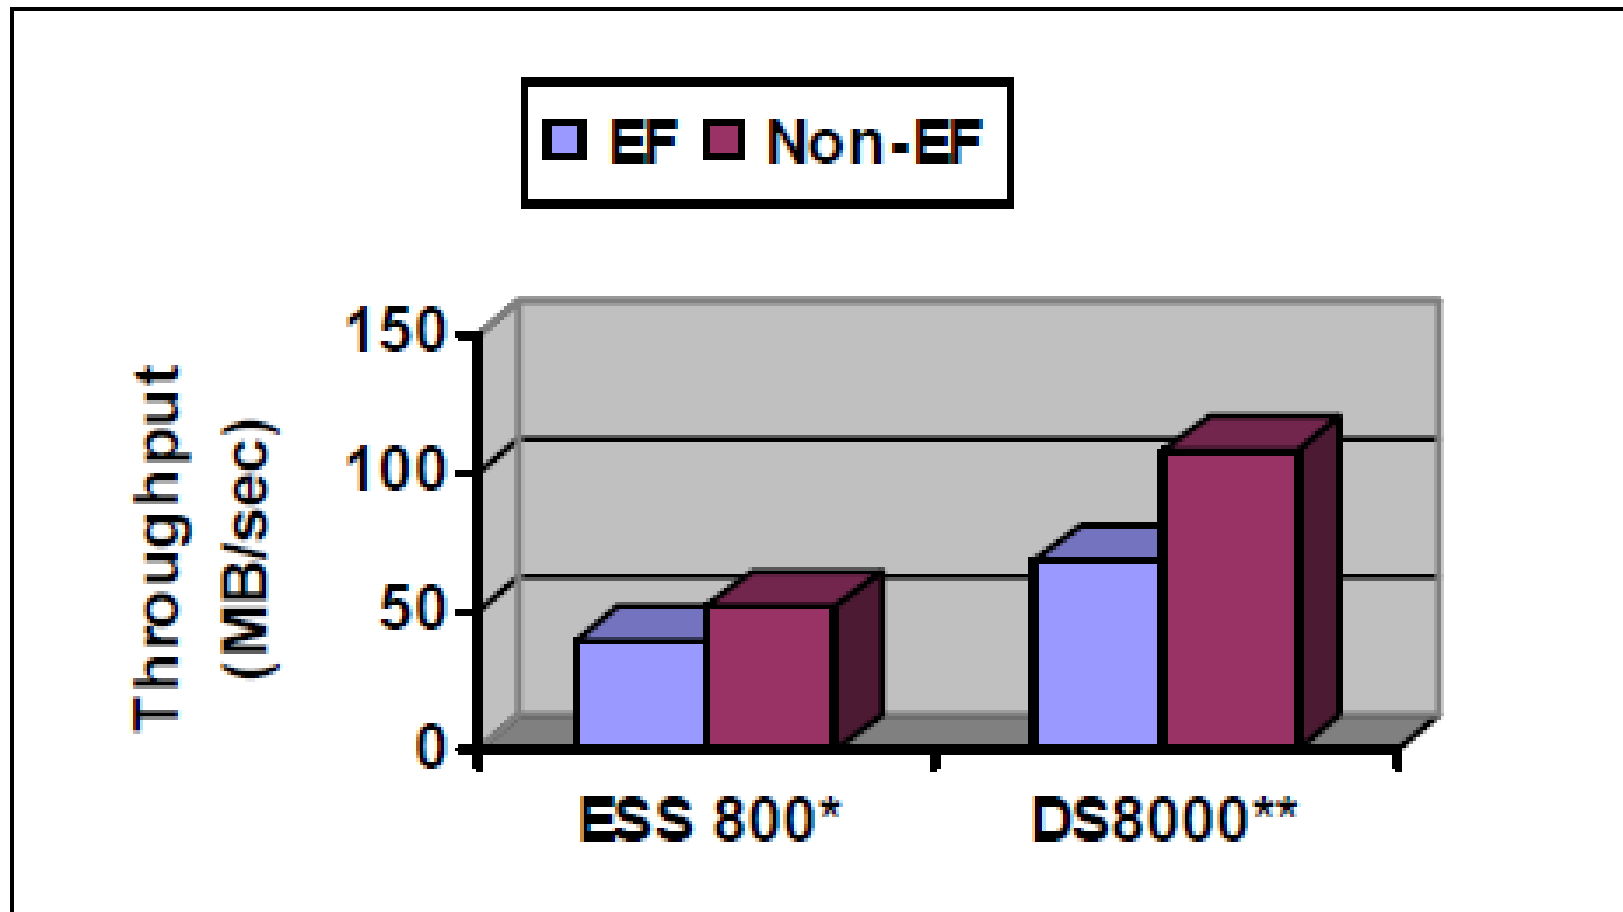
Buffer to Buffer Credit

Ficon Recovery

What Problem Does MIDAW Solve?


- Extended Format data sets have a small suffix appended to the data.
- Data and Suffix are in **DISCONTIGUOUS** virtual storage.
- Data and Suffix are in the same physical record on the DASD volume.
- This combination results in less than optimal channel, control unit, and link efficiencies.

Non-EF vs EF Performance before MIDAWs



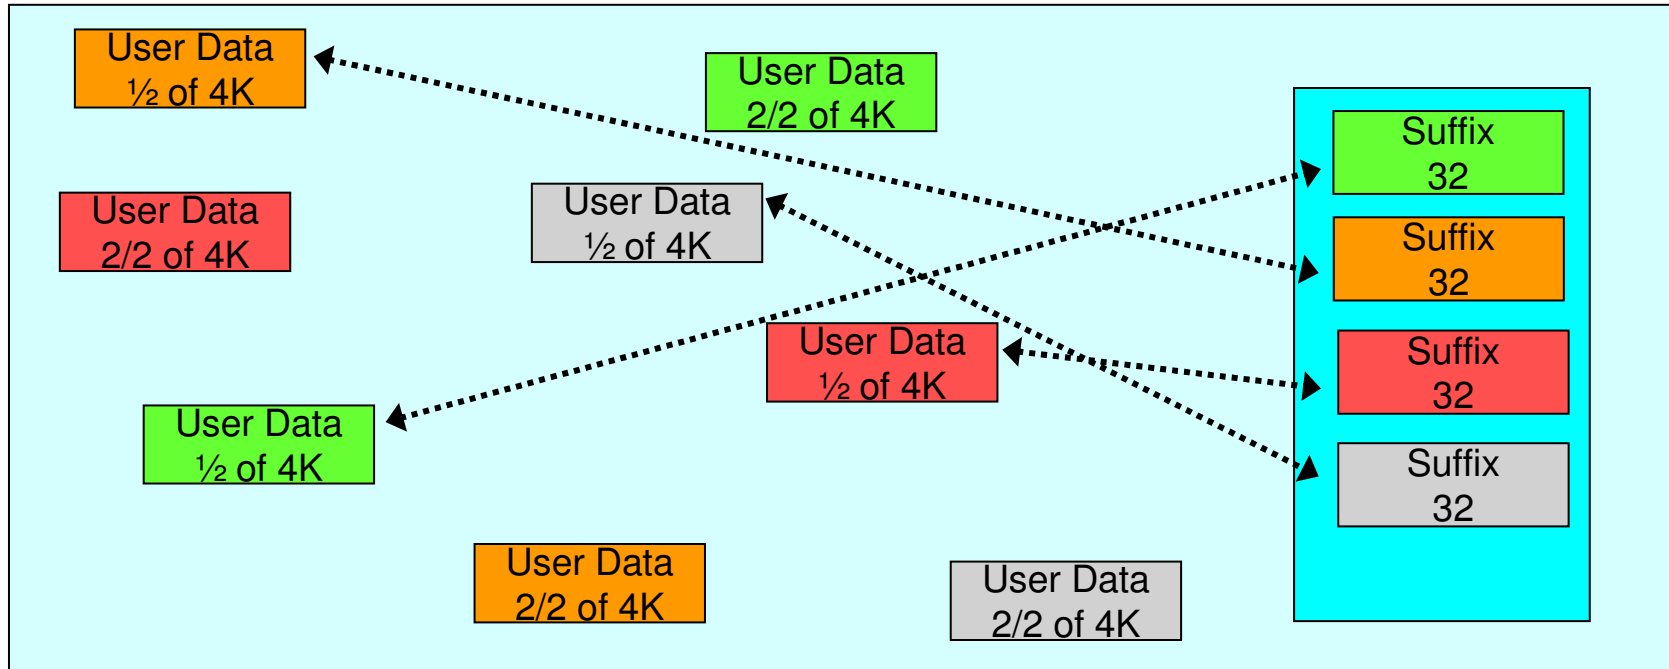
DB2 table scan with 4K pages prior to MIDAWs

EF Data in VIRTUAL Storage



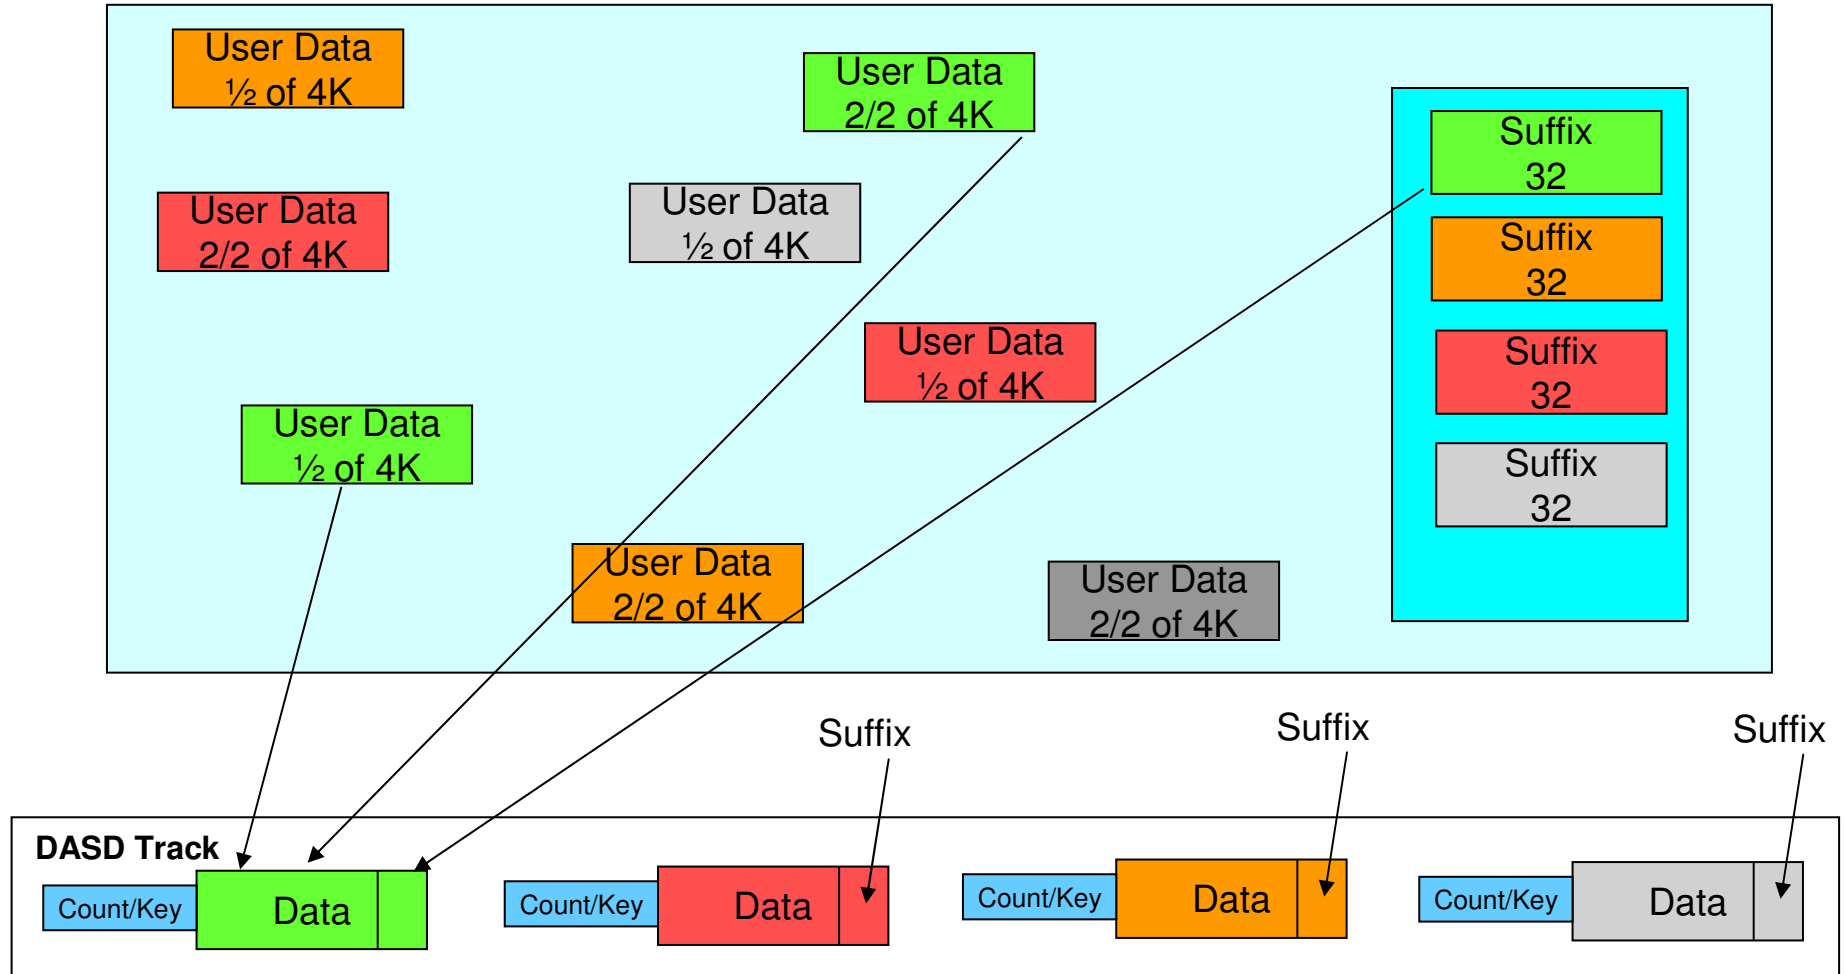
In virtual storage the User Data is contiguous

EF Data in REAL Storage

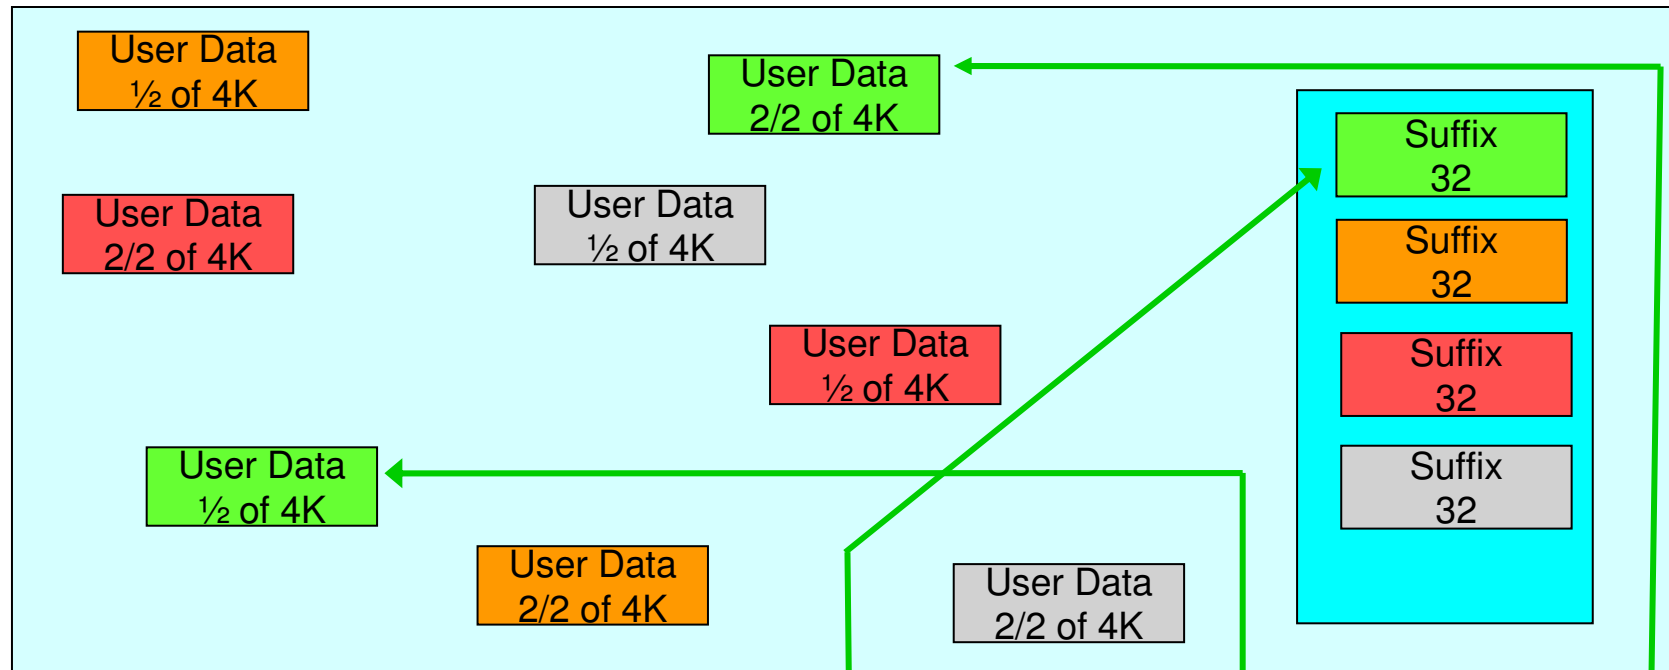


In REAL storage the User Data is scattered about

EF Data from REAL Storage to DASD volume

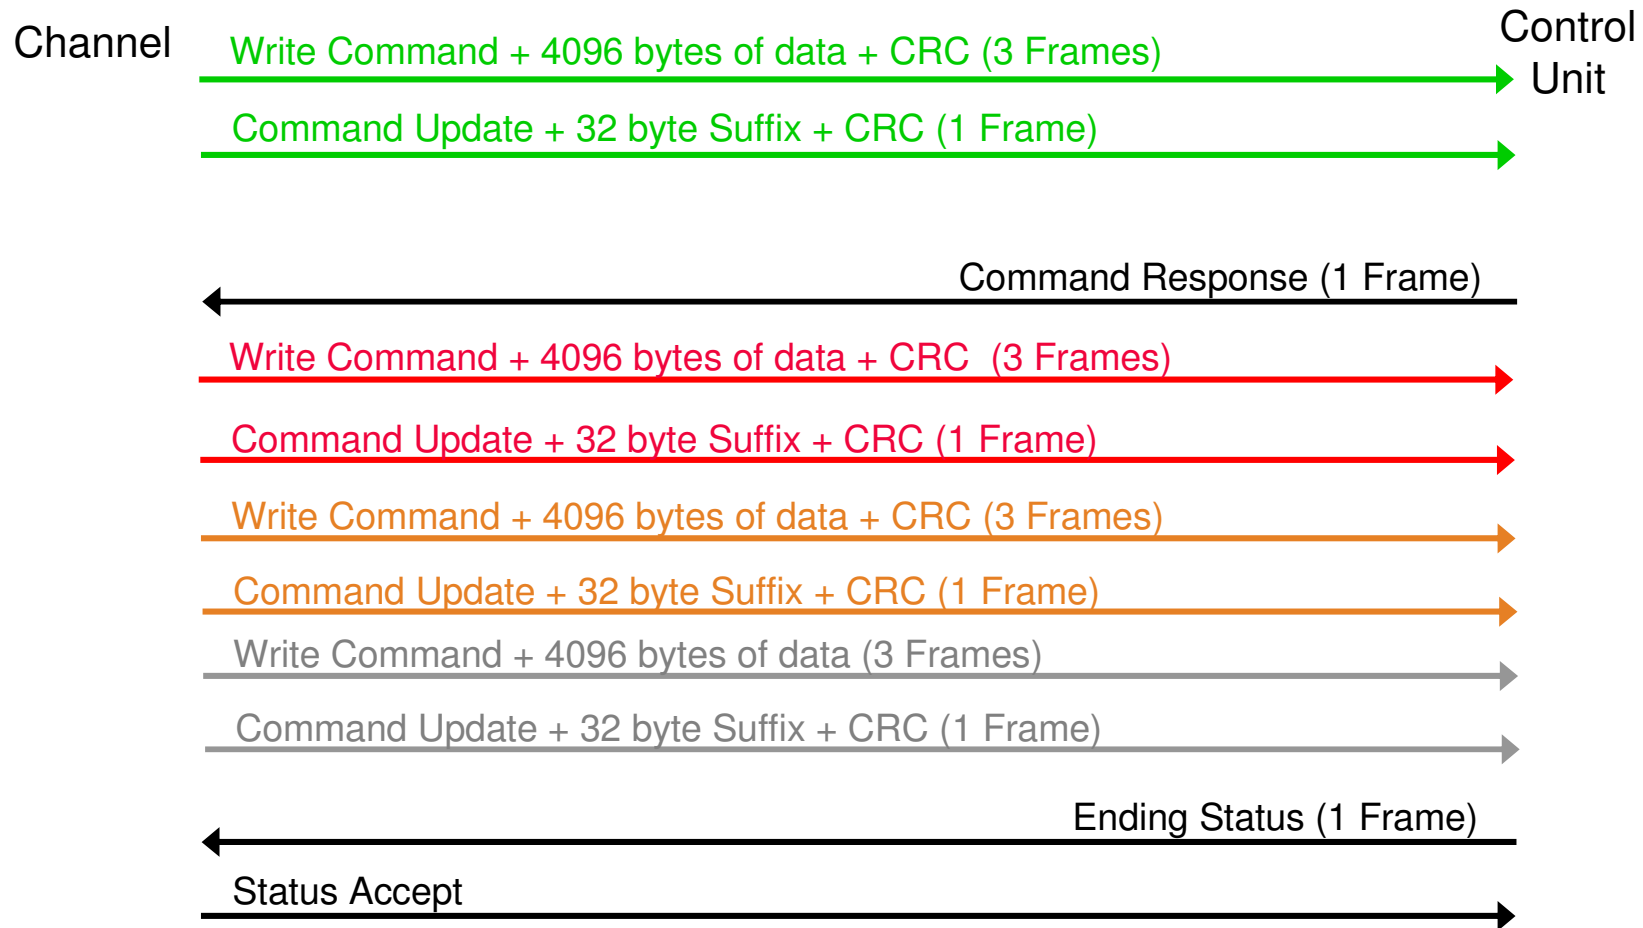


How It's Done With IDAW and Data Chaining



Write	CD,IDA	1000	00465E00	→	03EC4000
XX	CC	20	00175400	→	03C5B800
Write	CD,IDA	1000	00465E08	→	02FE9000
XX	CC	20	00175420	→	0C78D000
Write	CD,IDA	1000	00465E10	→	0E015800
XX	CC	20	00175440	→	097F8000
Write	CD,IDA	1000	00465E18	→	0D06E800
XX		20	00175460	→	0BF43000

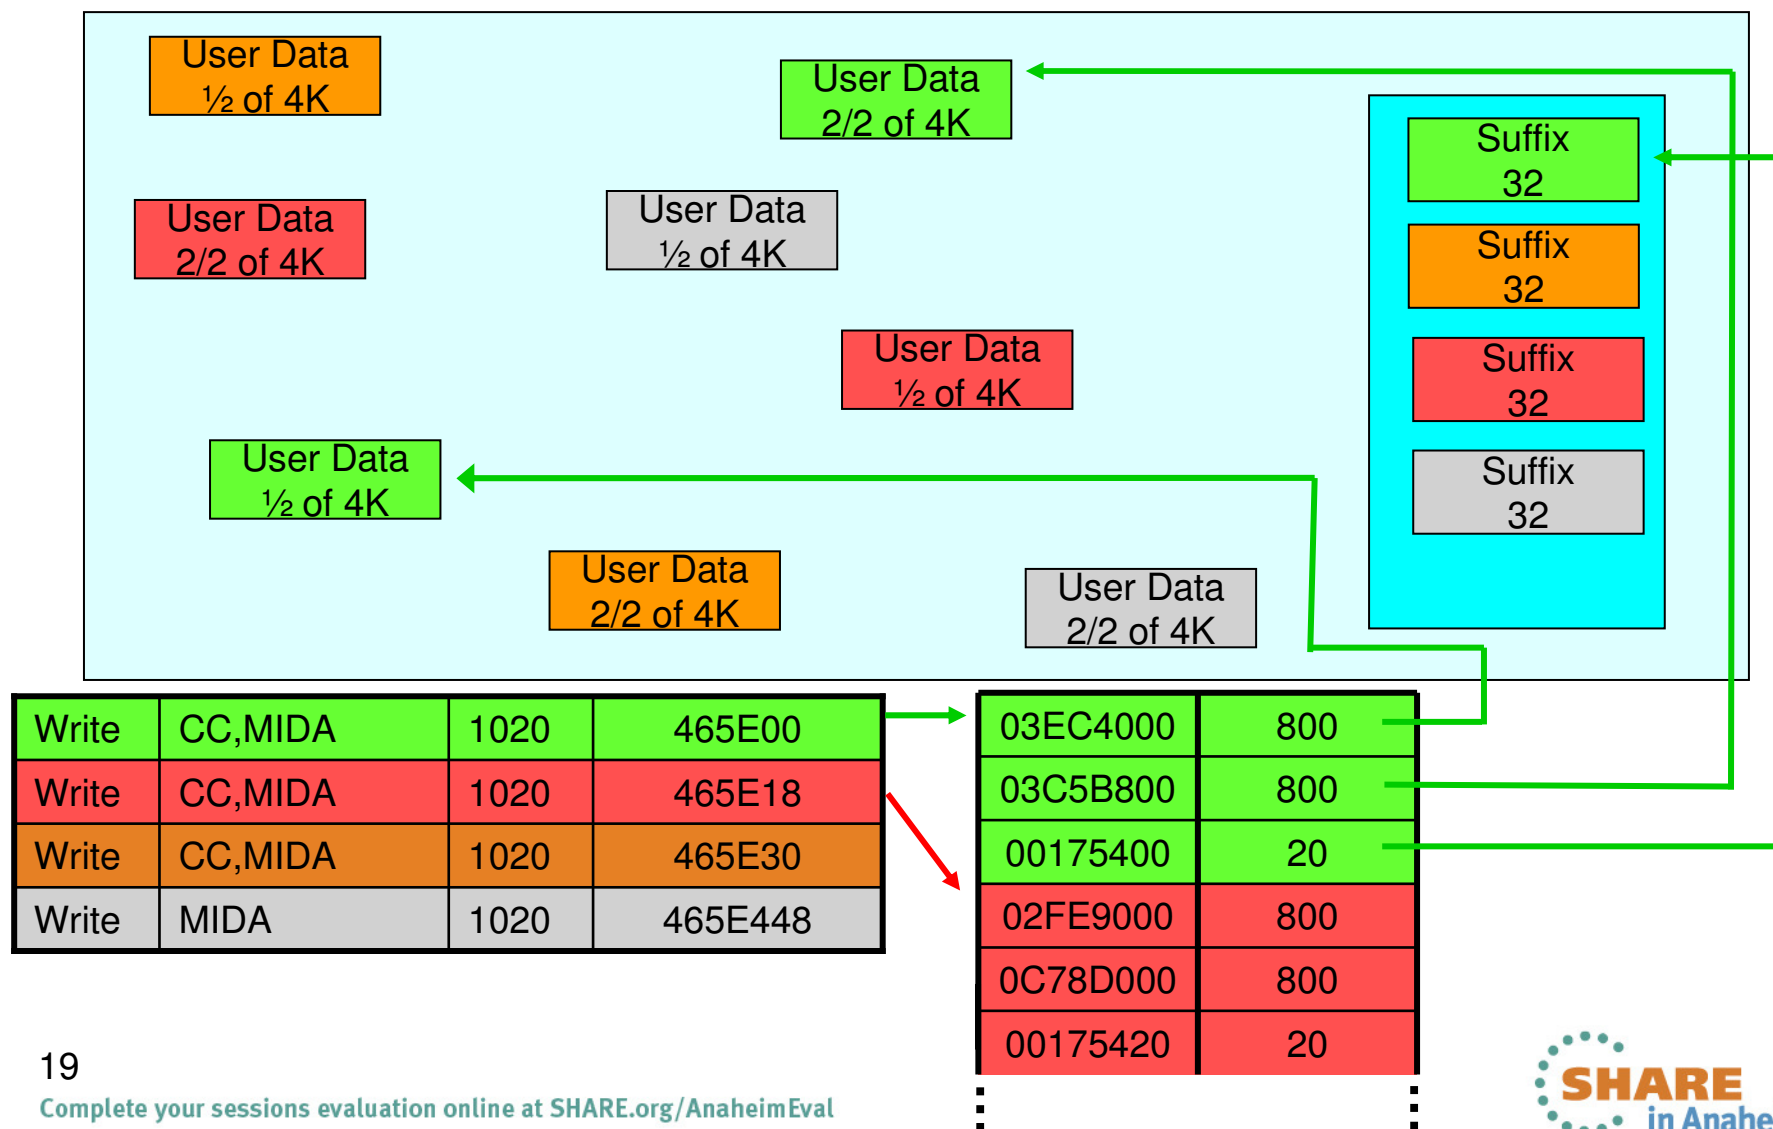
Sequence View Of the FC Link with IDAW



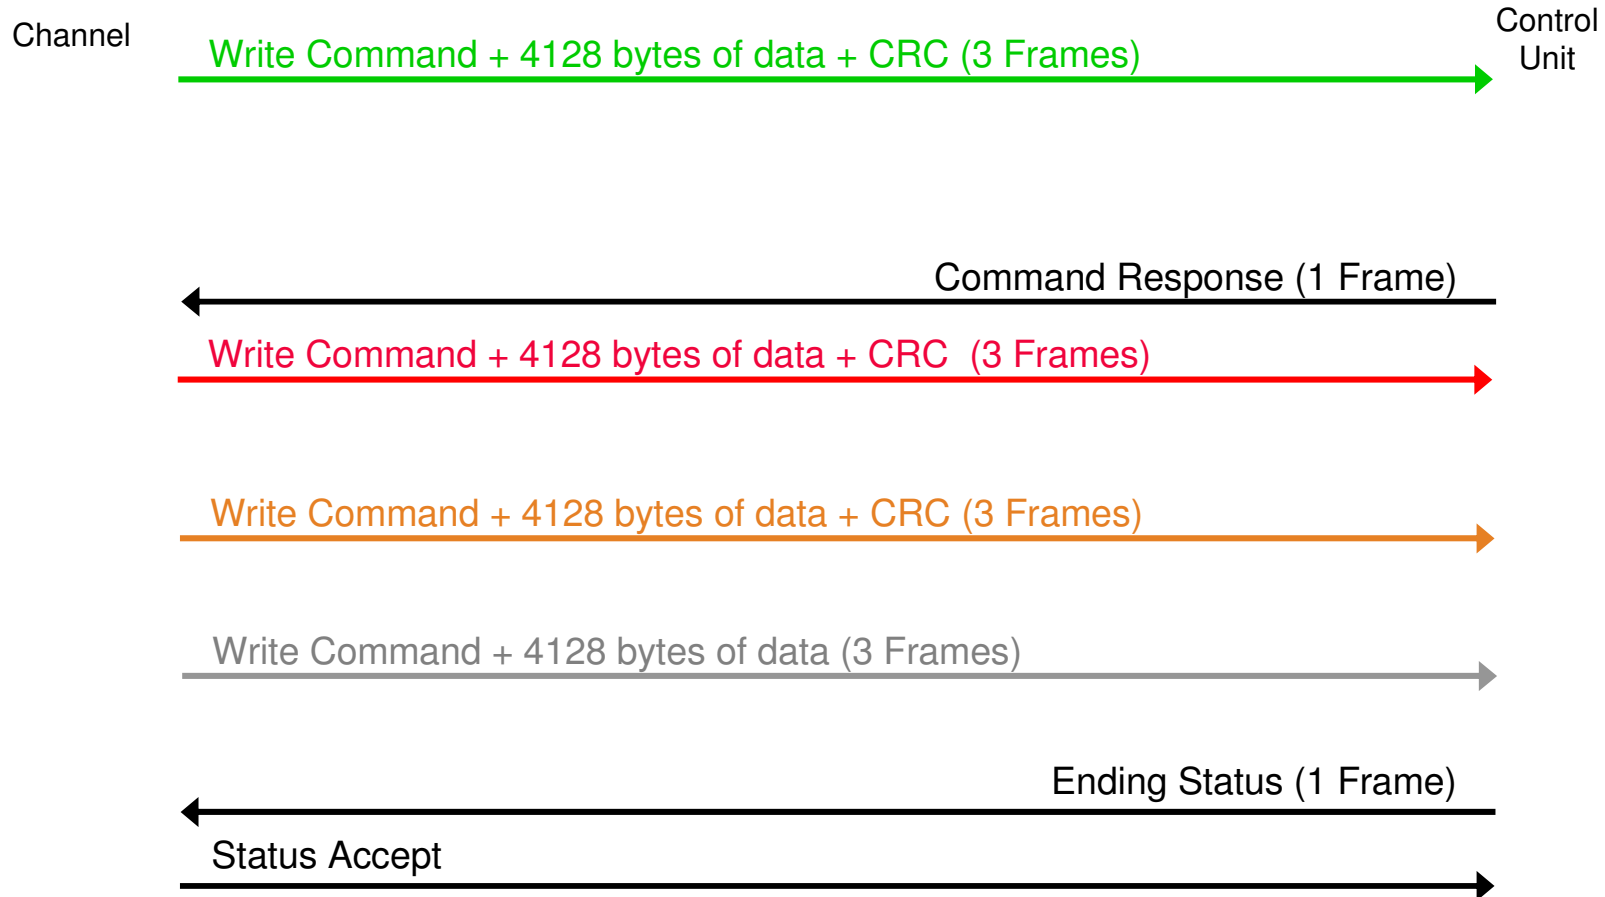
What is MIDAW?

- **M**odified **I**n**D**irect **A**ddress **W**ord
- A new method of gathering and scattering data into & from non-contiguous system z storage locations during an I/O operation.
- Designed to improve performance of certain applications
 - DB2 sequential workloads that use Media Manager to process small records with Extended Format data sets
- Reduces the number of CCWs in a channel program for these workloads

How It's Done With MIDAW



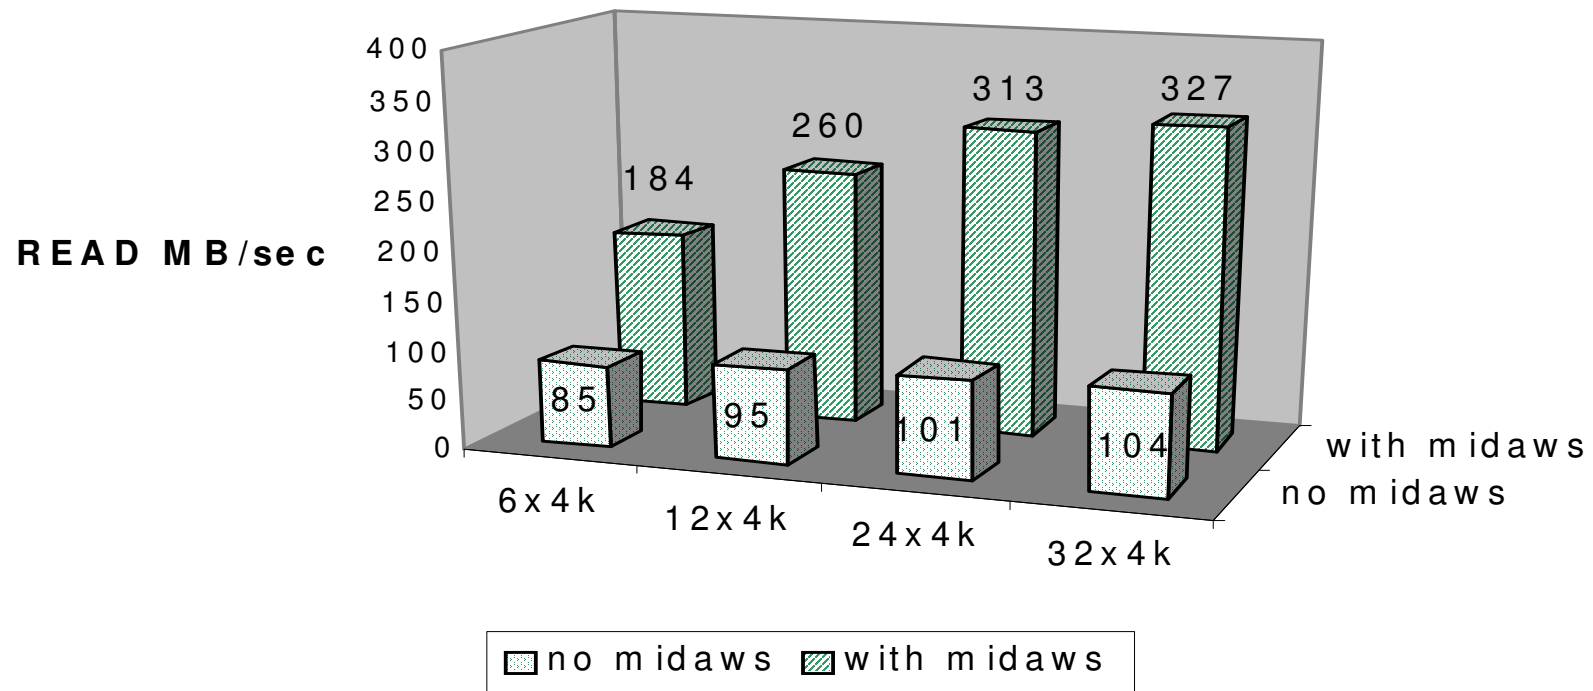
Sequence View Of the FC Link with MIDAW



IDAW vs MIDAW comparison

	Channel to CU		CU to Channel	
	IDAW	MIDAW	IDAW	MIDAW
Frames	17	13	2	2
Sequences	9	5	2	2

summary of FIC ON Express4 channel MIDAW s measurements



* This performance data was measured in a controlled environment running an I/O driver program under z/OS. The actual throughput or performance that any user will experience will vary depending upon considerations such as the amount of multiprogramming in the user's job stream, the I/O Configuration, the storage configuration, and the workload processed.

See also: <http://www.redbooks.ibm.com/redpapers/pdfs/redp4201.pdf>

Agenda

Frames, Sequences and Exchanges

IDAW vs MIDAW



Buffer to Buffer Credit

Ficon Recovery

What is Buffer-to-Buffer Credit?

- The greater the BB Credit....
 - A. The faster frames can be sent
 - B. The farther apart the two ports can be
 - C. The larger the frames can be
 - D. None of the above

What is Buffer-to-Buffer Credit?

- The greater the BB Credit....
 - A.
 - B. The farther apart the two ports can be
 - C.
 - D.

What is Buffer-to-Buffer Credit?

- BB Credit is the number of FRAME buffers a port provides for it's NEAREST neighbor

- BB Credit does NOT have to be symmetrical

What is Buffer-to-Buffer Credit?

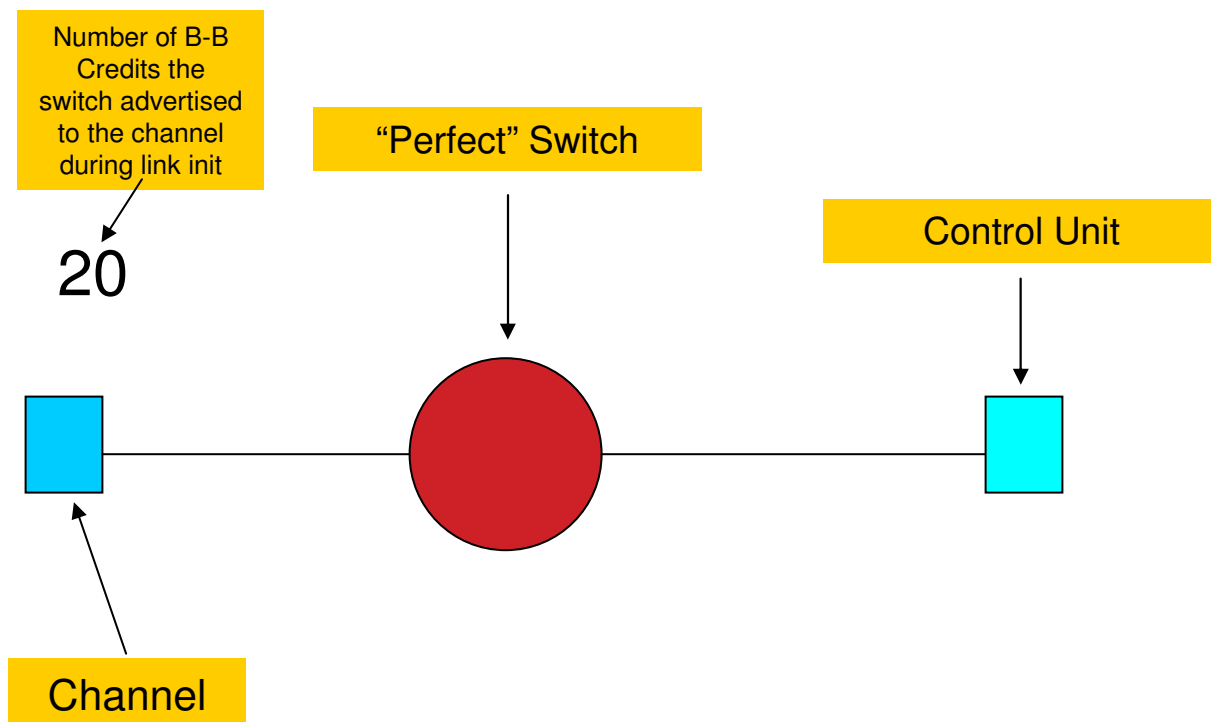
- BB credit value determines the DISTANCE two nodes can be apart and still maintain full link frame rate
- BB credit is on a FRAME basis, not frame SIZE basis
 - A 1 byte frame consumes 1 credit
 - A 2K byte frame consumes 1 credit
- Number of credits needed determined by:
 - Raw Link Speed
 - Distance between two adjacent nodes
 - Average frame size

What is Buffer-to-Buffer Credit?

- Each time a frame is sent, the sender decrements its available credit count by 1
- Each time a frame receiver clears a frame buffer it sends a “R_RDY”
 - Special 4 byte character – **NOT a frame** !
- Reception of a R_RDY causes the available credit count to be incremented by 1


BB-Credit Animations

Initial Conditions



NOTE: In these animations, both the frames and R_RDY’s are numbered. This is for illustrative purposes only. In reality, neither the frames nor the R_RDY’s are numbered. The arrival of an R_RDY only informs the receiver that **A** frame buffer has been freed, not **WHICH** frame buffer has been freed.


20



31

Complete your sessions evaluation online at SHARE.org/AnaheimEval


19



32


Complete your sessions evaluation online at SHARE.org/AnaheimEval

18




33

Complete your sessions evaluation online at SHARE.org/AnaheimEval



SHARE
Technology • Connections • Results

17




34

Complete your sessions evaluation online at SHARE.org/AnaheimEval



2012


16



35

Complete your sessions evaluation online at SHARE.org/AnaheimEval


15



36

Complete your sessions evaluation online at SHARE.org/AnaheimEval


14



37

Complete your sessions evaluation online at SHARE.org/AnaheimEval


13




38

Complete your sessions evaluation online at SHARE.org/AnaheimEval

12




11



40

Complete your sessions evaluation online at SHARE.org/AnaheimEval


10



41

Complete your sessions evaluation online at SHARE.org/AnaheimEval


9



42

Complete your sessions evaluation online at SHARE.org/AnaheimEval


8



43

Complete your sessions evaluation online at SHARE.org/AnaheimEval


7



44

Complete your sessions evaluation online at SHARE.org/AnaheimEval


6



45

Complete your sessions evaluation online at SHARE.org/AnaheimEval


5



46

Complete your sessions evaluation online at SHARE.org/AnaheimEval


4



47

Complete your sessions evaluation online at SHARE.org/AnaheimEval

4




48


Complete your sessions evaluation online at SHARE.org/AnaheimEval



4



4




50

Complete your sessions evaluation online at SHARE.org/AnaheimEval


This Page Intentionally Left Blank

Suppose the switch is too far away from the channel for the B-B credit it advertised to the channel

20




19



53


Complete your sessions evaluation online at SHARE.org/AnaheimEval

18




54

Complete your sessions evaluation online at SHARE.org/AnaheimEval



SHARE
Technology • Connections • Results

17




55

Complete your sessions evaluation online at SHARE.org/AnaheimEval



2012


16



56

Complete your sessions evaluation online at SHARE.org/AnaheimEval


15




57

Complete your sessions evaluation online at SHARE.org/AnaheimEval

14




13




59

Complete your sessions evaluation online at SHARE.org/AnaheimEval



SHARE
Technology • Connections • Results

12




60

Complete your sessions evaluation online at SHARE.org/AnaheimEval



2012


11




61

Complete your sessions evaluation online at SHARE.org/AnaheimEval


10




9



8




7




65

Complete your sessions evaluation online at SHARE.org/AnaheimEval

6




5




67

Complete your sessions evaluation online at SHARE.org/AnaheimEval


4




3




2




1



0




0




73

Complete your sessions evaluation online at SHARE.org/AnaheimEval

0




0




75

Complete your sessions evaluation online at SHARE.org/AnaheimEval


0




0




0



0




0




80

Complete your sessions evaluation online at SHARE.org/AnaheimEval


0




0




0




0




0





SHARE
Technology • Connections • Results


0




86

Complete your sessions evaluation online at SHARE.org/AnaheimEval


0




0



0




0




90

Complete your sessions evaluation online at SHARE.org/AnaheimEval


0




0




0






SHARE
Technology • Connections • Results


0





SHARE
Technology • Connections • Results

0




95


Complete your sessions evaluation online at SHARE.org/AnaheimEval




2012




0




0




0




0






SHARE
Technology • Connections • Results

0




100

Complete your sessions evaluation online at SHARE.org/AnaheimEval




SHARE
Technology • Connections • Results

0




101

Complete your sessions evaluation online at SHARE.org/AnaheimEval



SHARE
Technology • Connections • Results


0




102

Complete your sessions evaluation online at SHARE.org/AnaheimEval


0




0




0






SHARE
Technology • Connections • Results

0





0






0





SHARE
Technology • Connections • Results

0




109

Complete your sessions evaluation online at SHARE.org/AnaheimEval

This Page Intentionally Left Blank


Suppose there are multiple senders to one receiver
Each sender attempts to send at 100% link speed

20



20

19




19

112

Complete your sessions evaluation online at SHARE.org/AnaheimEval

18




18

113

Complete your sessions evaluation online at SHARE.org/AnaheimEval

17




17

114

Complete your sessions evaluation online at SHARE.org/AnaheimEval

16




16

115

Complete your sessions evaluation online at SHARE.org/AnaheimEval

15




15

116

Complete your sessions evaluation online at SHARE.org/AnaheimEval

14




14

117

Complete your sessions evaluation online at SHARE.org/AnaheimEval

13




13

118

Complete your sessions evaluation online at SHARE.org/AnaheimEval

12




12

119

Complete your sessions evaluation online at SHARE.org/AnaheimEval


11




11

120

Complete your sessions evaluation online at SHARE.org/AnaheimEval



SHARE
Technology • Connections • Results




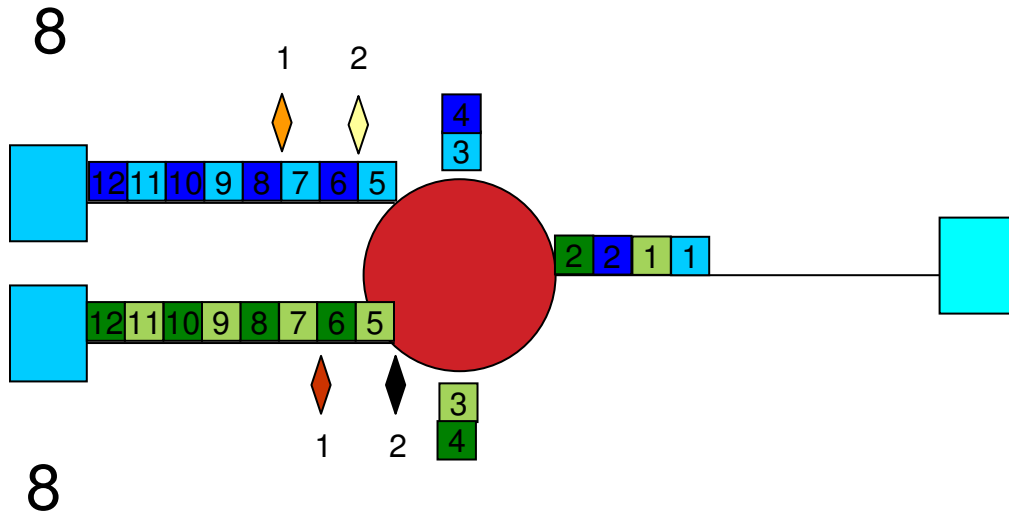
121


Complete your sessions evaluation online at SHARE.org/AnaheimEval





2012




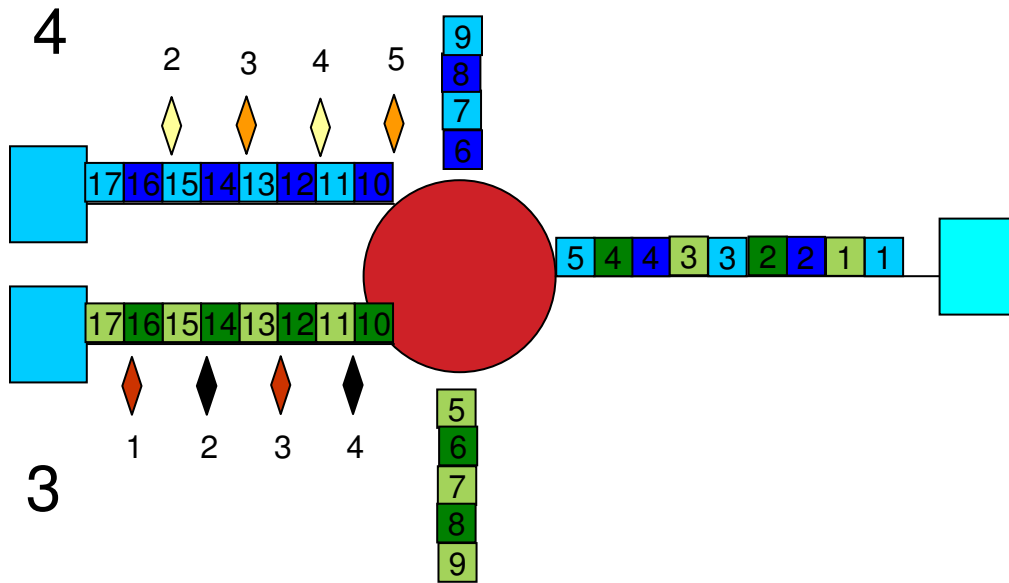






















SHARE
Technology • Connections • Results







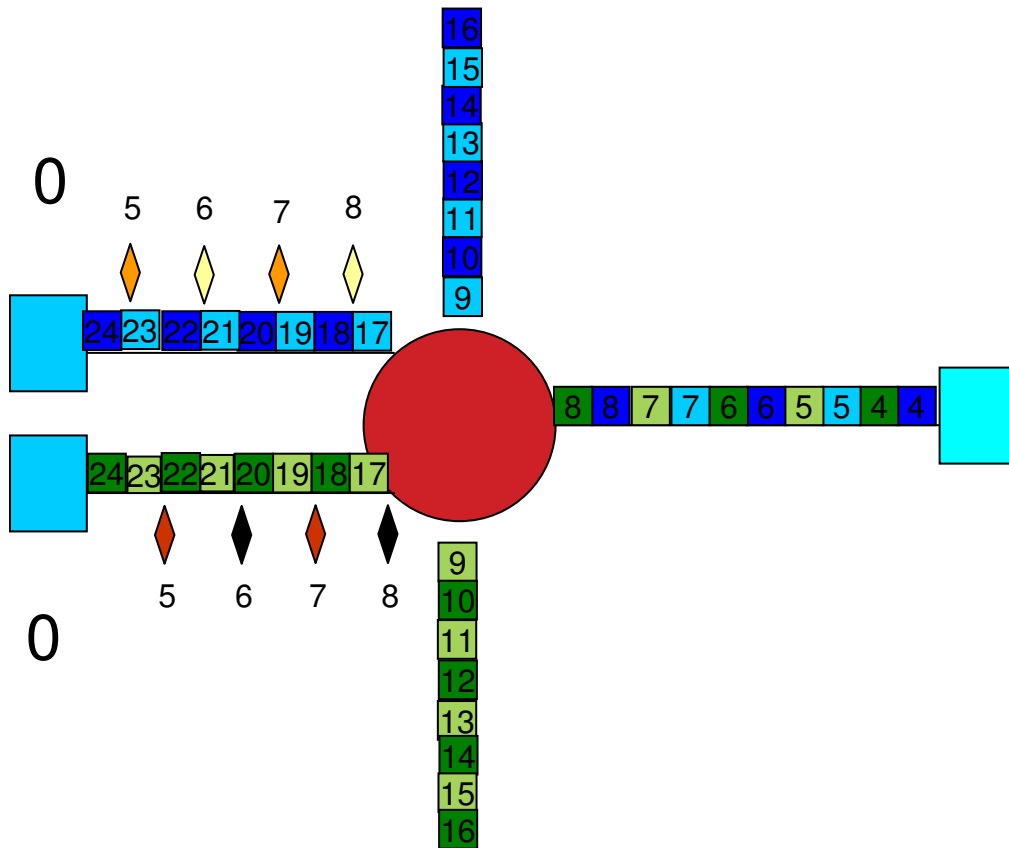


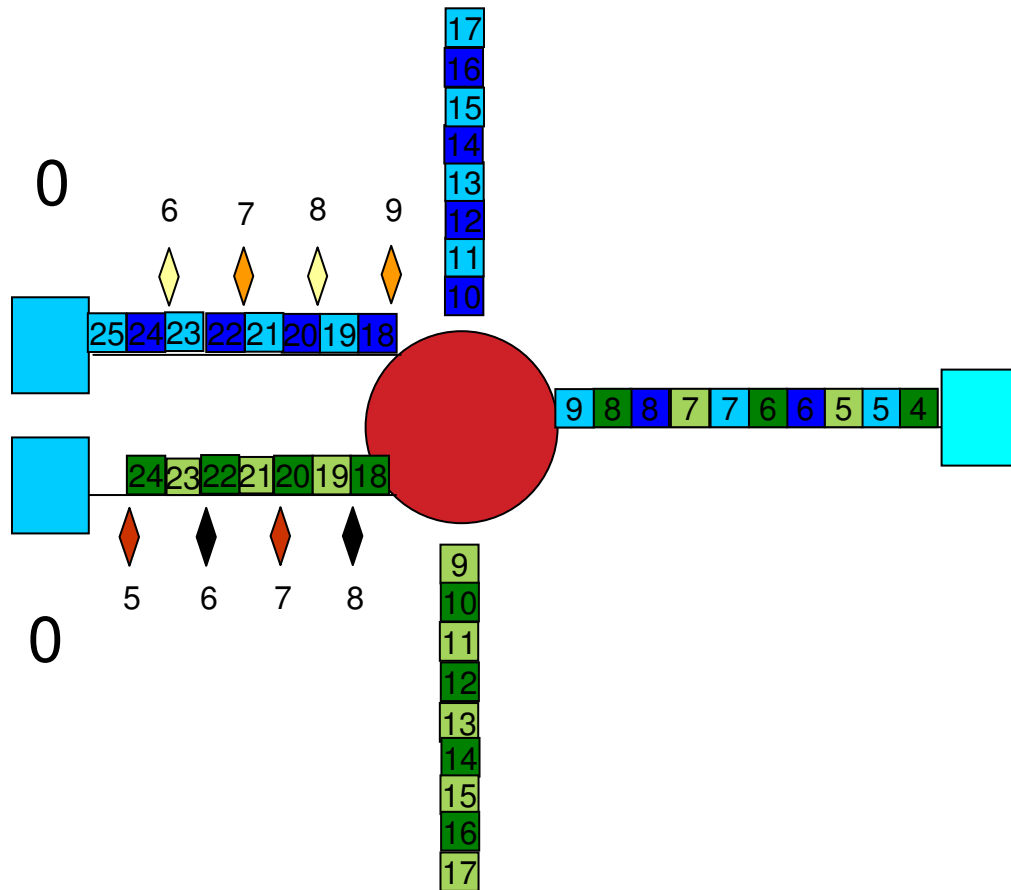


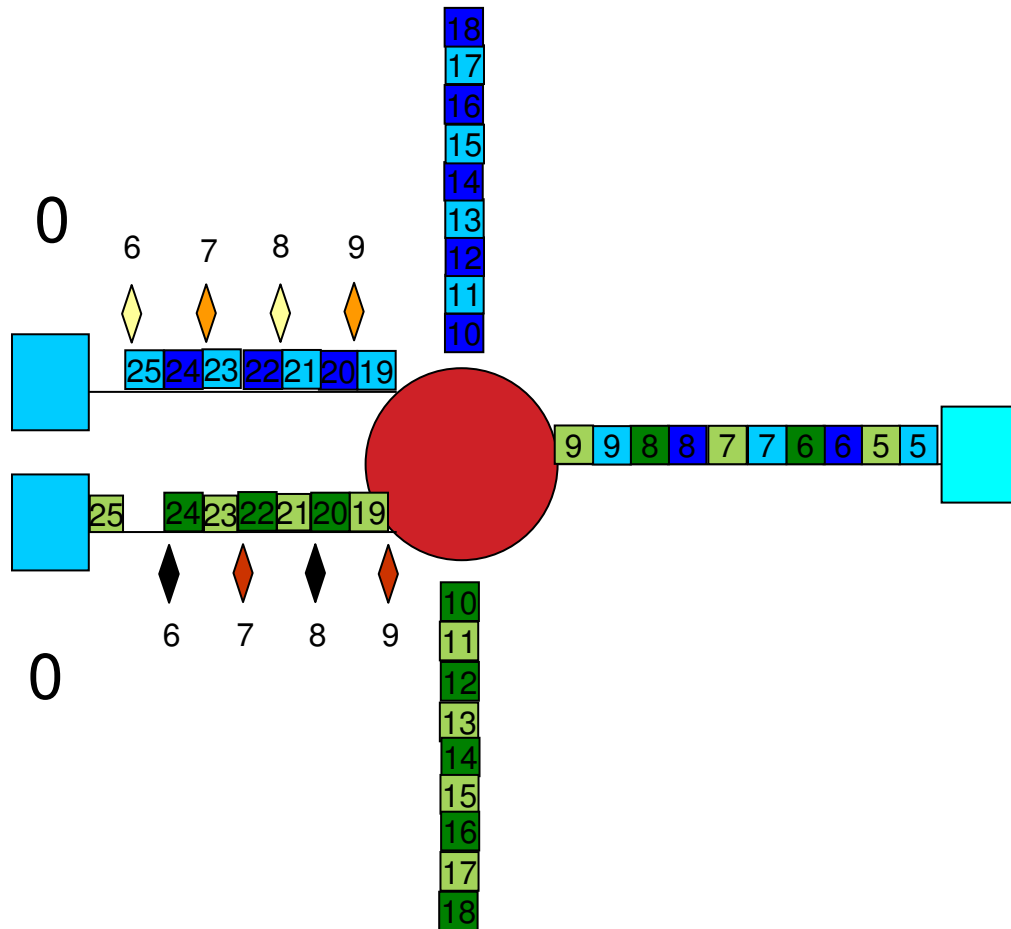
SHARE
Technology • Connections • Results




















SHARE
Technology • Connections • Results





SHARE
Technology • Connections • Results




140


Complete your sessions evaluation online at SHARE.org/AnaheimEval





2012





SHARE
Technology • Connections • Results















This Page Intentionally Left Blank


More real life example: Two senders sending at 30% - 50% link rate to one receiver

20



20

19




19

149

Complete your sessions evaluation online at SHARE.org/AnaheimEval

18




18

150

Complete your sessions evaluation online at SHARE.org/AnaheimEval


17



18


151

Complete your sessions evaluation online at SHARE.org/AnaheimEval



SHARE
Technology • Connections • Results

16




18

152

Complete your sessions evaluation online at SHARE.org/AnaheimEval

16




17

153

Complete your sessions evaluation online at SHARE.org/AnaheimEval


16



16


154

Complete your sessions evaluation online at SHARE.org/AnaheimEval



SHARE
Technology • Connections • Results

16



16


155

Complete your sessions evaluation online at SHARE.org/AnaheimEval



2012

15




16

156

Complete your sessions evaluation online at SHARE.org/AnaheimEval


15



16

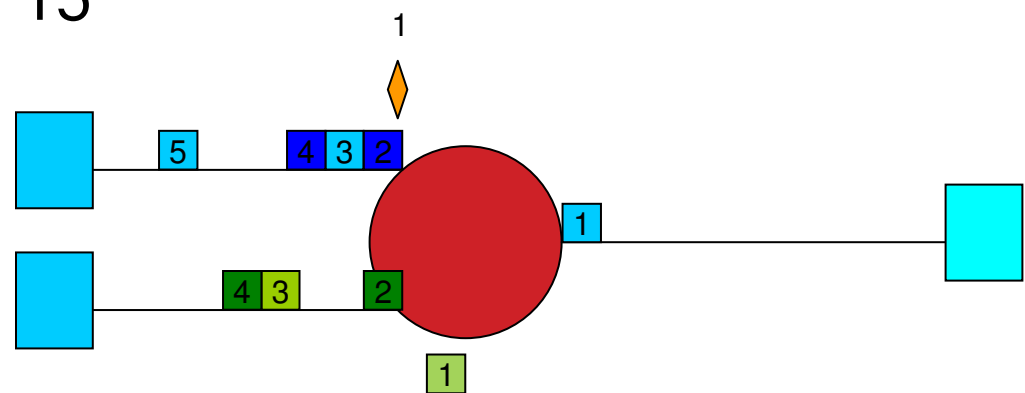
157

Complete your sessions evaluation online at SHARE.org/AnaheimEval



SHARE
Technology • Connections • Results

15



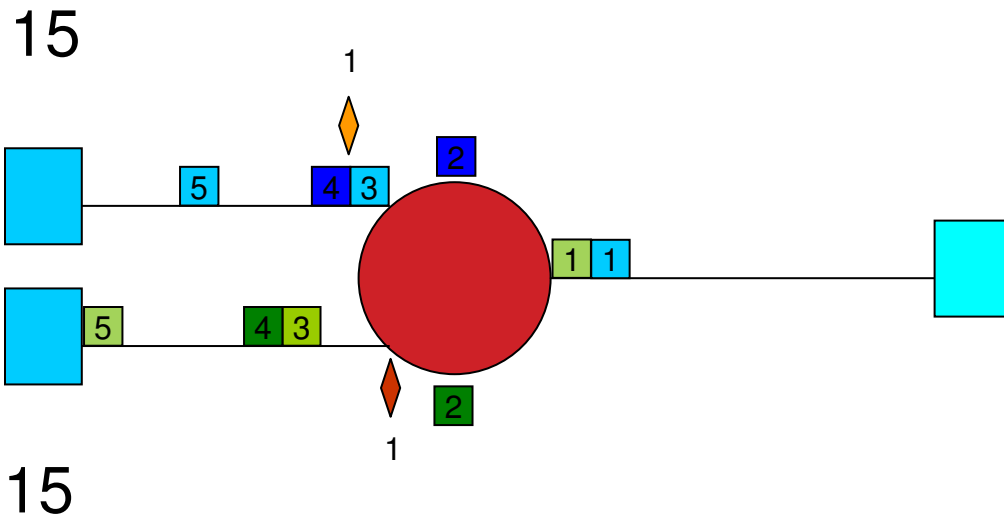
16


158

Complete your sessions evaluation online at SHARE.org/AnaheimEval




2012





SHARE
Technology • Connections • Results





160


Complete your sessions evaluation online at SHARE.org/AnaheimEval





2012


















SHARE
Technology • Connections • Results

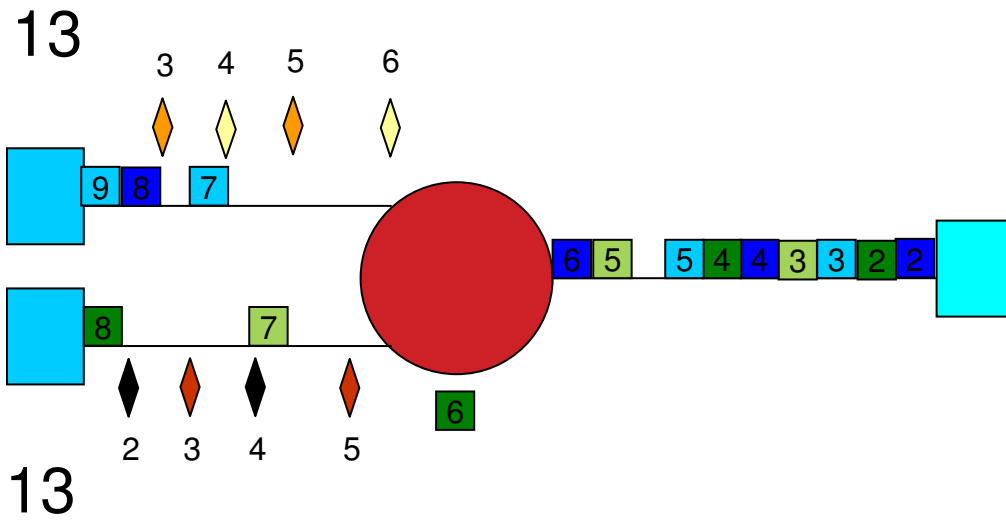








SHARE
Technology • Connections • Results







SHARE
Technology • Connections • Results




170

Complete your sessions evaluation online at SHARE.org/AnaheimEval




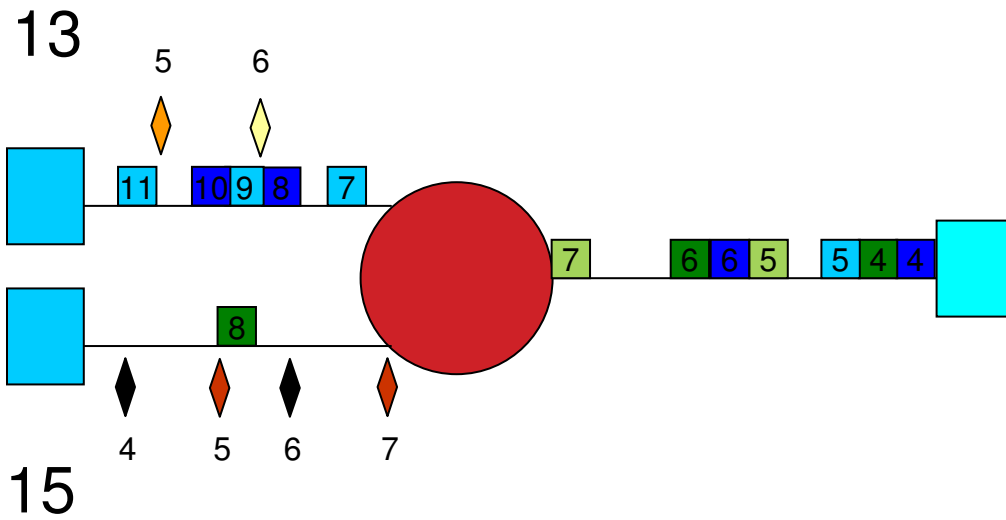
2012



171


Complete your sessions evaluation online at SHARE.org/AnaheimEval






173


Complete your sessions evaluation online at SHARE.org/AnaheimEval







175


Complete your sessions evaluation online at SHARE.org/AnaheimEval







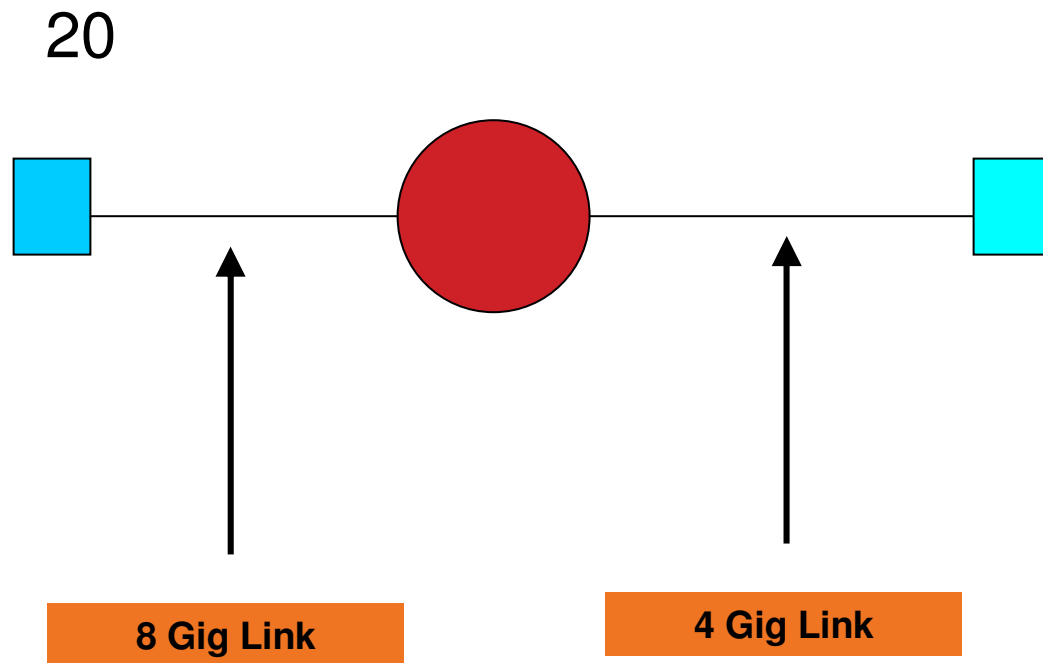
SHARE
Technology • Connections • Results



178
Complete your sessions evaluation online at SHARE.org/AnaheimEval

This Page Intentionally Left Blank


Fat Pipe / Skinny Pipe



180


Complete your sessions evaluation online at SHARE.org/AnaheimEval

19




181

Complete your sessions evaluation online at SHARE.org/AnaheimEval



SHARE
Technology • Connections • Results

18




182

Complete your sessions evaluation online at SHARE.org/AnaheimEval




2012



SHARE
Technology • Connections • Results

17




183

Complete your sessions evaluation online at SHARE.org/AnaheimEval



2012


16



184

Complete your sessions evaluation online at SHARE.org/AnaheimEval


15



185

Complete your sessions evaluation online at SHARE.org/AnaheimEval


14



186

Complete your sessions evaluation online at SHARE.org/AnaheimEval


13



187

Complete your sessions evaluation online at SHARE.org/AnaheimEval


12



188


Complete your sessions evaluation online at SHARE.org/AnaheimEval

11




189

Complete your sessions evaluation online at SHARE.org/AnaheimEval



SHARE
Technology • Connections • Results

10




190


Complete your sessions evaluation online at SHARE.org/AnaheimEval




2012





SHARE
Technology • Connections • Results







SHARE
Technology • Connections • Results







SHARE
Technology • Connections • Results







SHARE
Technology • Connections • Results

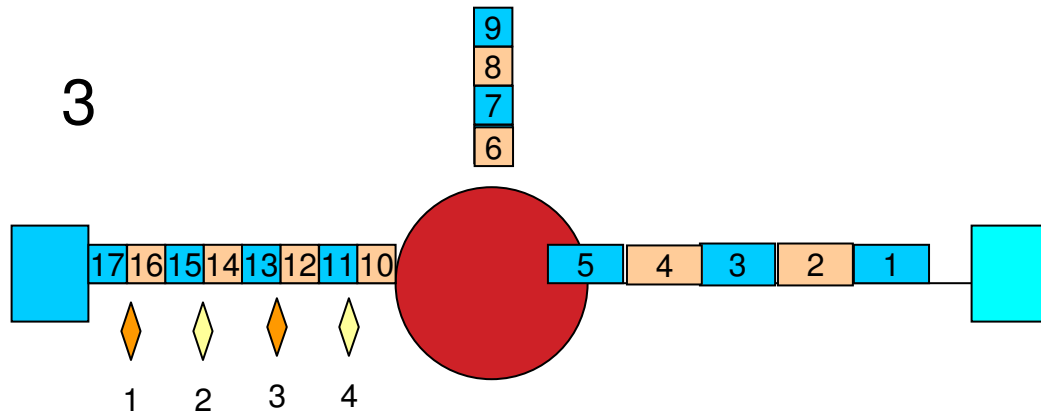








SHARE
Technology • Connections • Results












SHARE
Technology • Connections • Results





SHARE
Technology • Connections • Results





200


Complete your sessions evaluation online at SHARE.org/AnaheimEval




2012












SHARE
Technology • Connections • Results





SHARE
Technology • Connections • Results





205

Complete your sessions evaluation online at SHARE.org/AnaheimEval

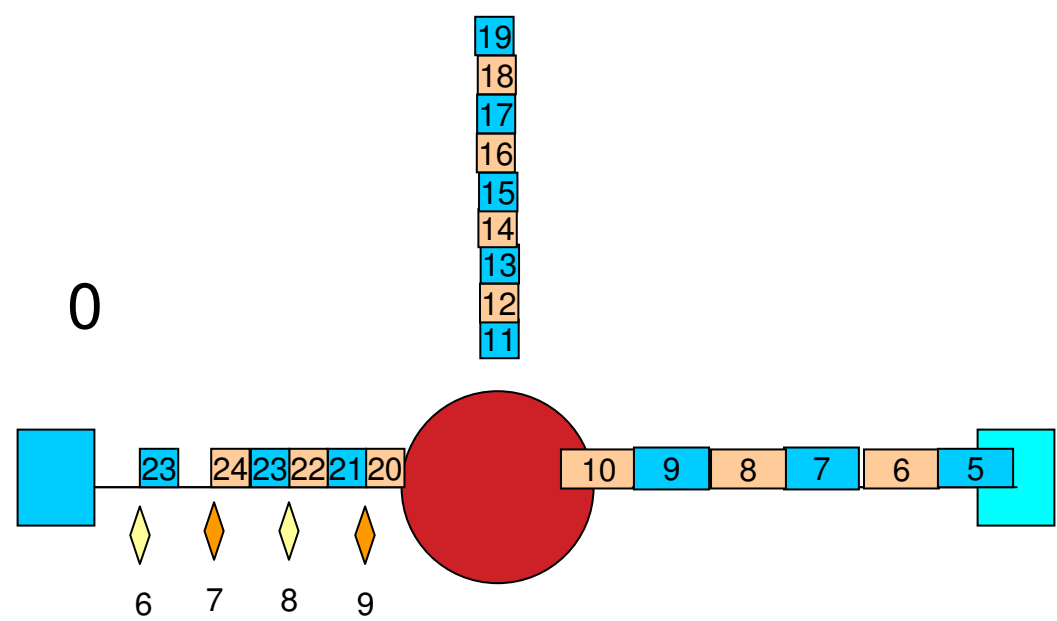



2012



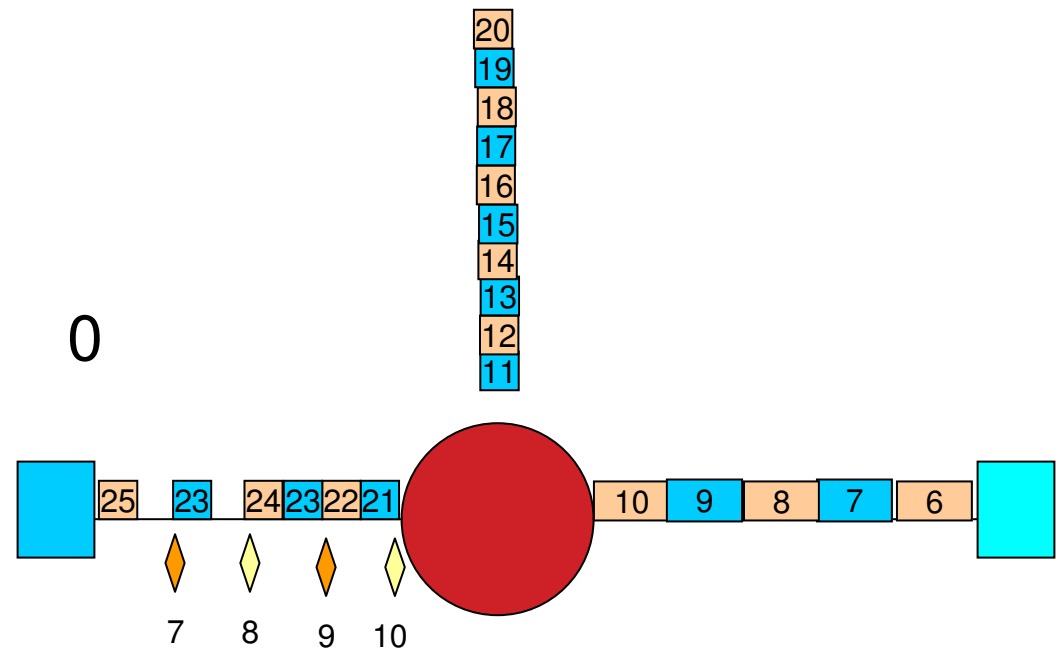




SHARE
Technology • Connections • Results







SHARE
Technology • Connections • Results







SHARE
Technology • Connections • Results




210


Complete your sessions evaluation online at SHARE.org/AnaheimEval



2012



SHARE
Technology • Connections • Results




211


Complete your sessions evaluation online at SHARE.org/AnaheimEval



2012



SHARE
Technology • Connections • Results





212

Complete your sessions evaluation online at SHARE.org/AnaheimEval




2012





SHARE
Technology • Connections • Results



214


Complete your sessions evaluation online at SHARE.org/AnaheimEval




2012




0





SHARE
Technology • Connections • Results

0



216

Complete your sessions evaluation online at SHARE.org/AnaheimEval



2012

How much credit do I need?

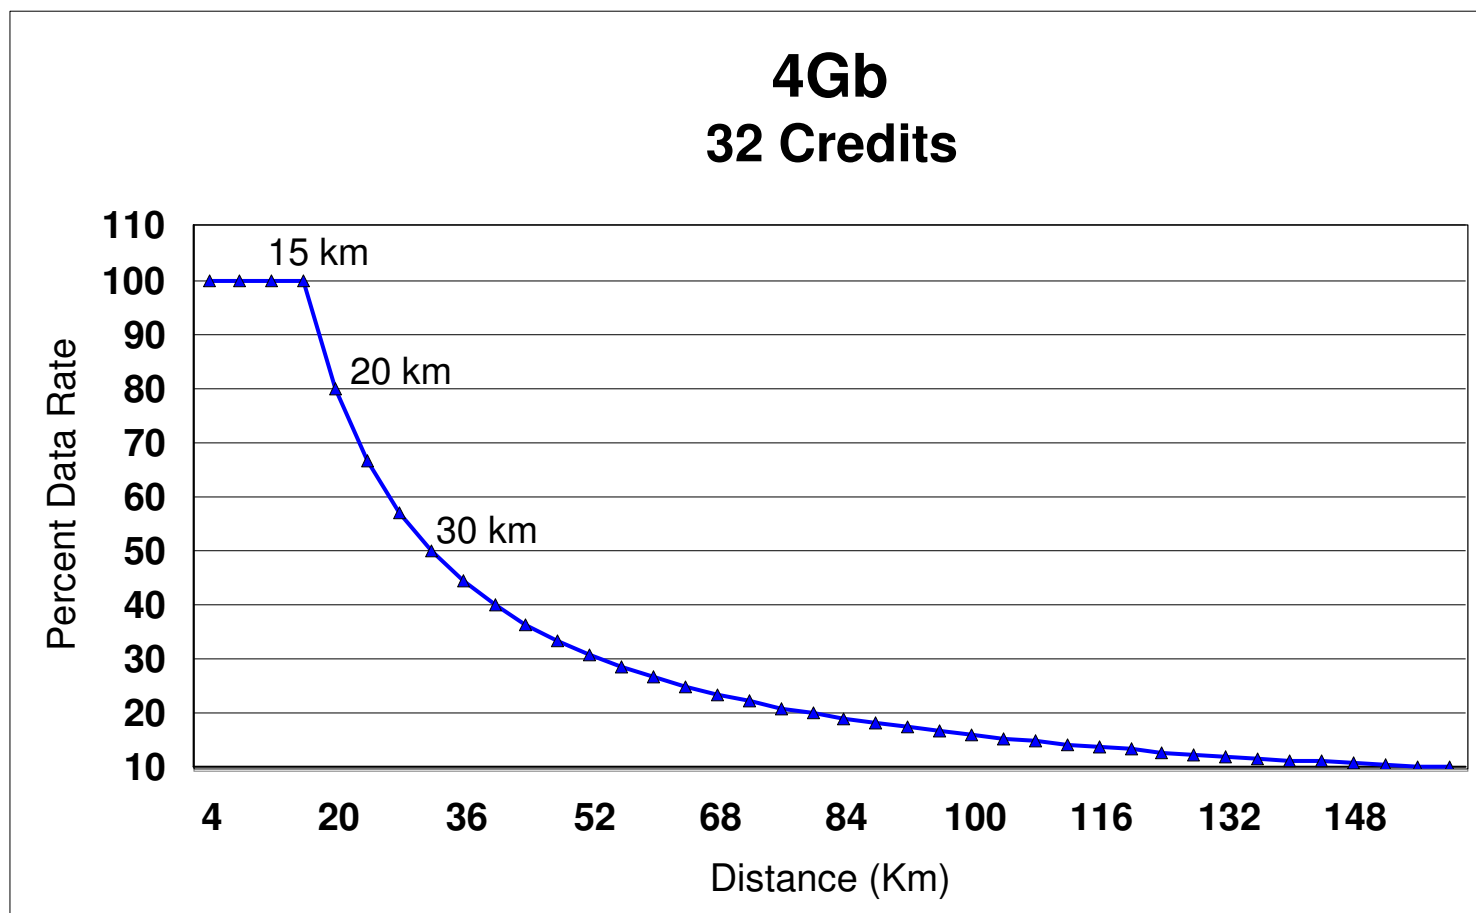
- Good “Rule of thumb”

Number of credits needed = $1 + \frac{\text{Link speed in Gb/s} * \text{Distance in Km}}{\text{Average Frame Size in Kb}}$


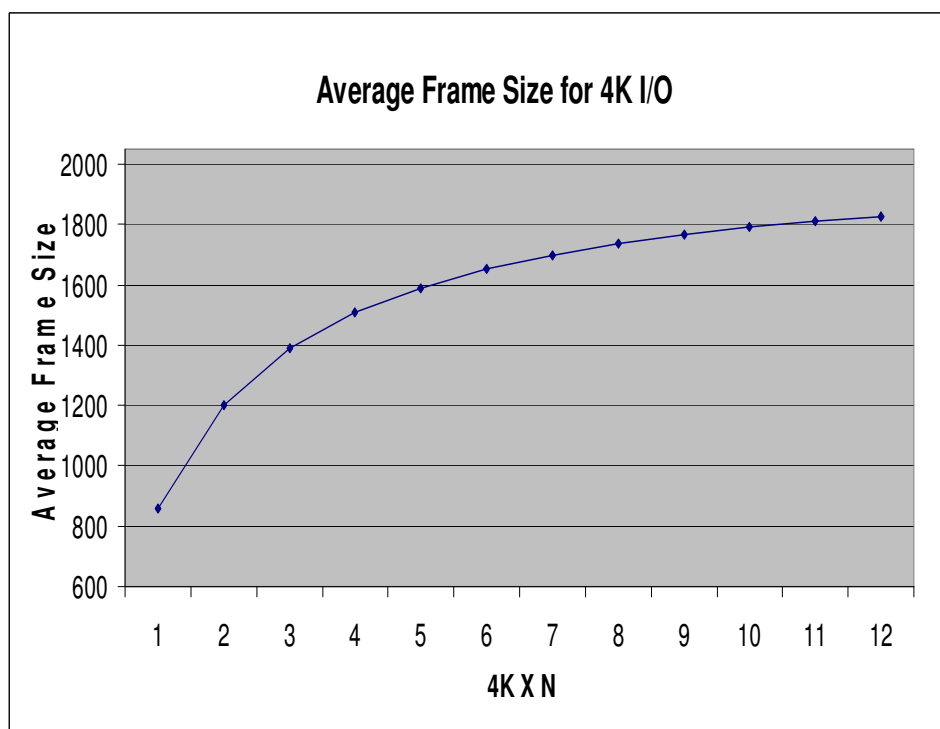
Example: 30 km at 4 Gb/s
 $1 + \frac{4 * 30}{2} = 61$

Example: 10 km at 8 Gb/s
 $1 + \frac{8 * 10}{2} = 41$

How does credit affect Bandwidth?



Average Frame Size vs Block Size



Agenda

Frames, Sequences and Exchanges

IDAW vs MIDAW

Buffer to Buffer Credit

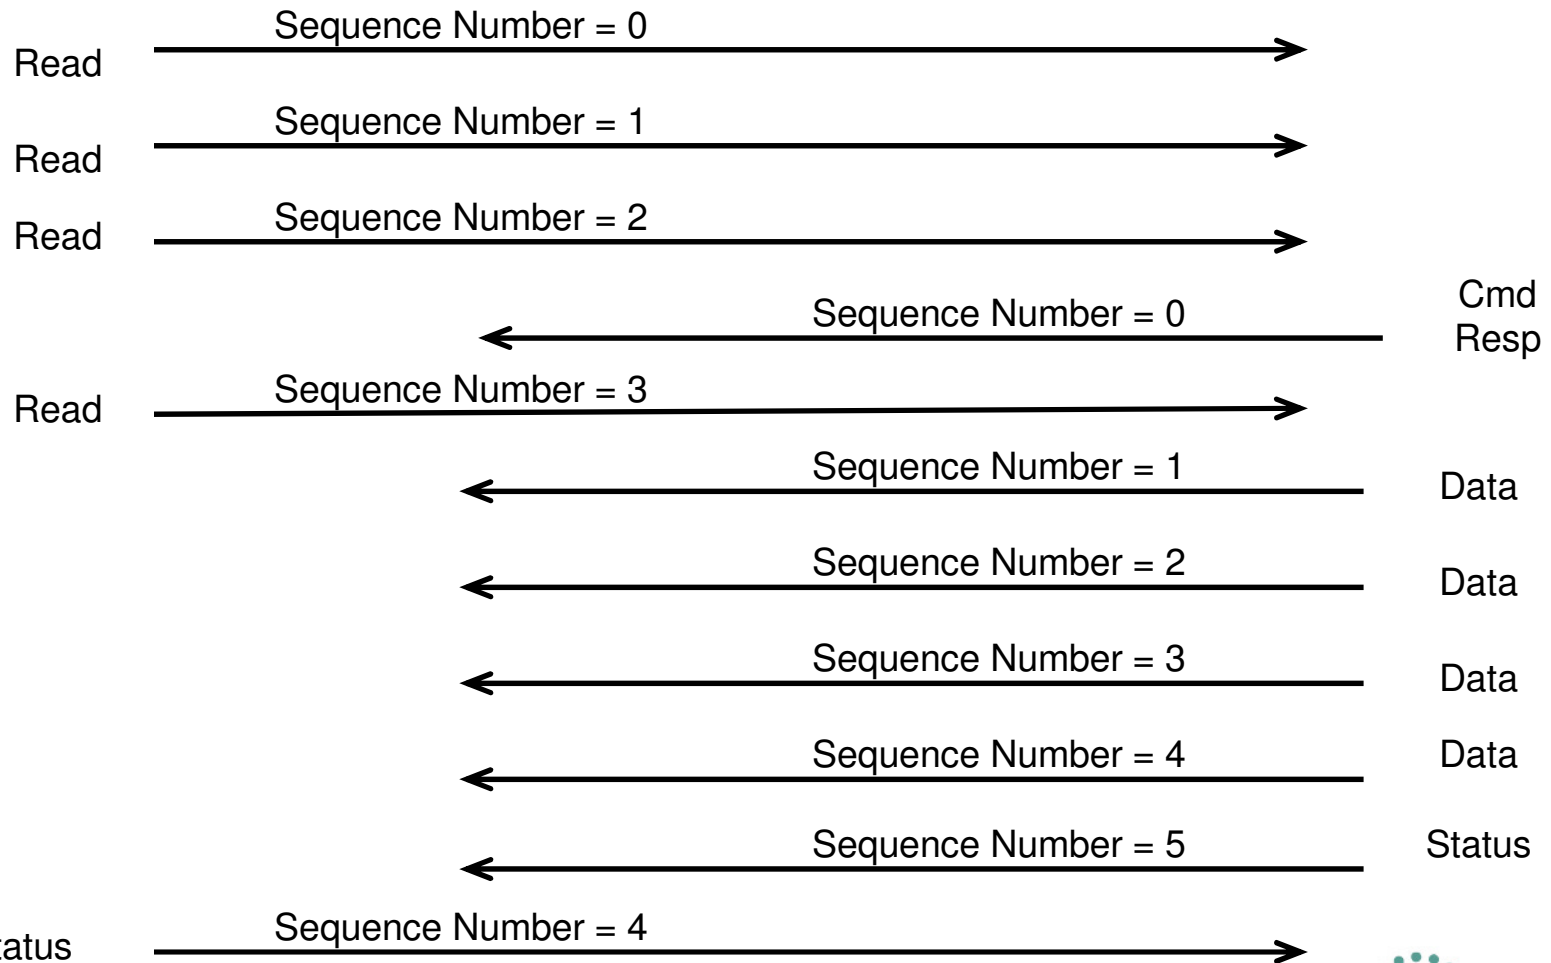


Ficon Recovery

Error Detection and Recovery

- Device Level
 - Error scope is one single I/O operation
- Link Level
 - Error scope is all active I/O on the affected link
 - Channel to Switch Link → All I/O active on the channel
 - Switch to CU Link → All I/O on that one destination link
- Channel Internal Errors


Device Level Errors



222 Accept

Complete your sessions evaluation online at SHARE.org/AnaheimEval

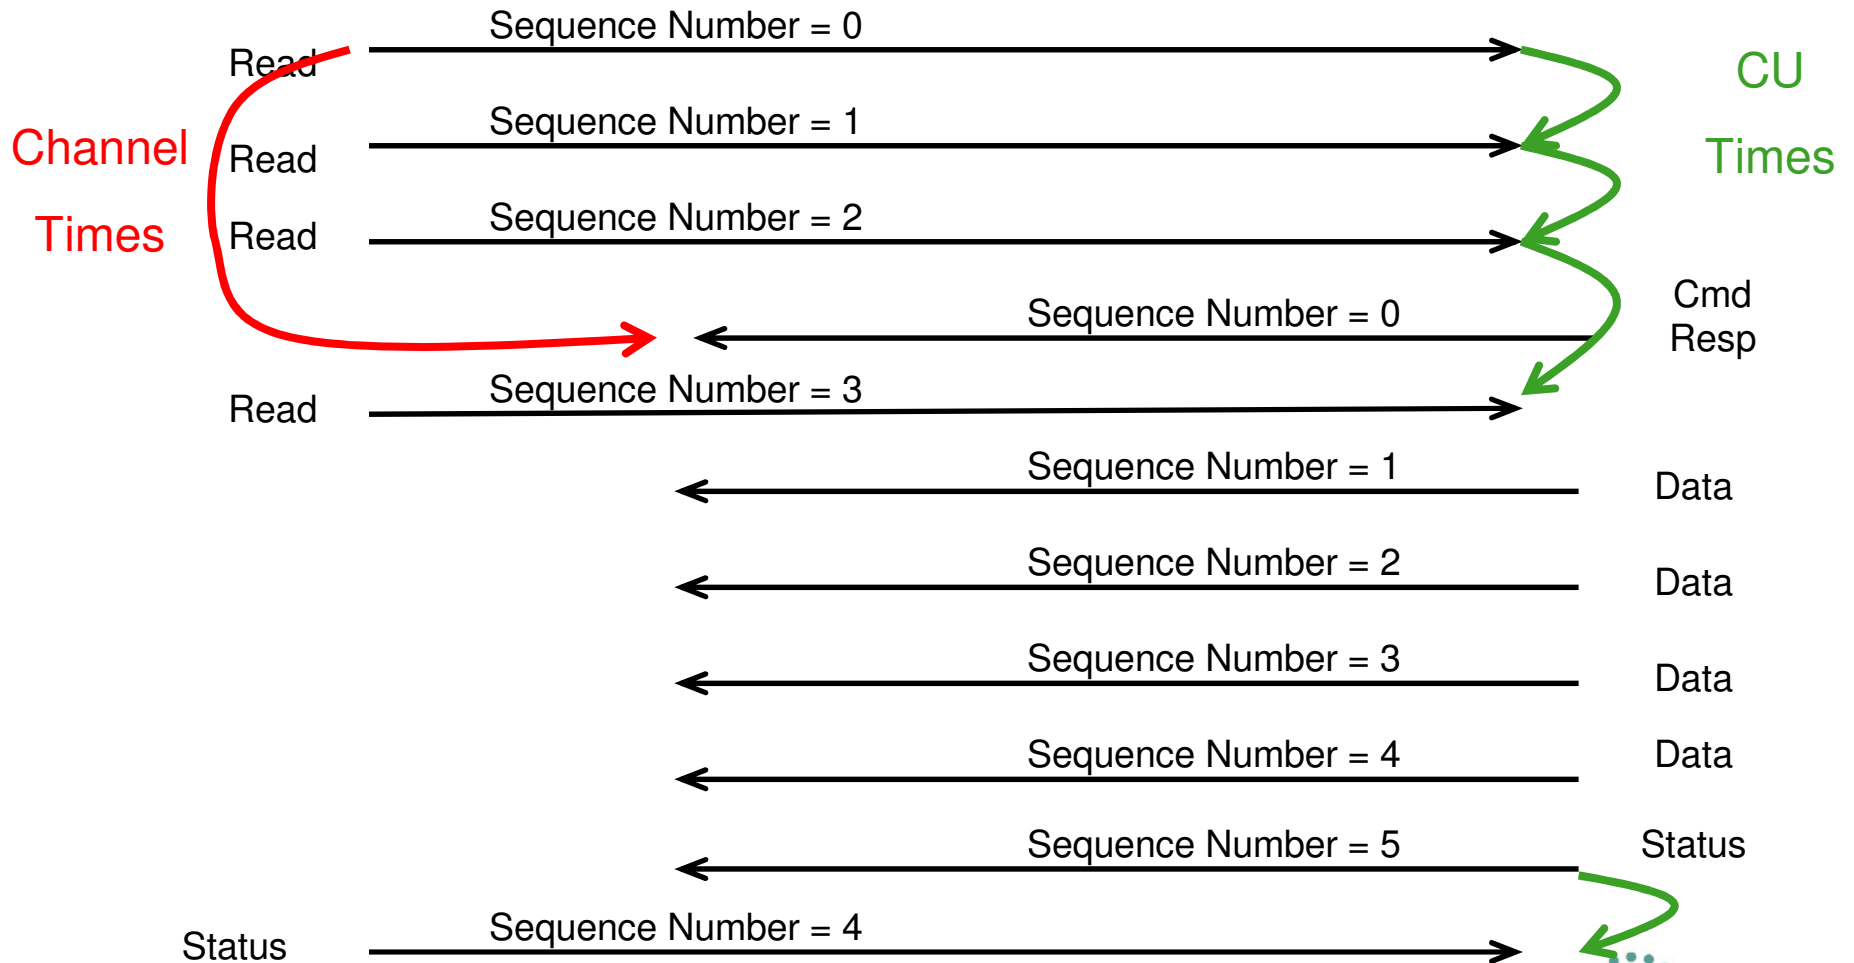
Device Level Errors



Status
223 Accept

Complete your sessions evaluation online at SHARE.org/AnaheimEval

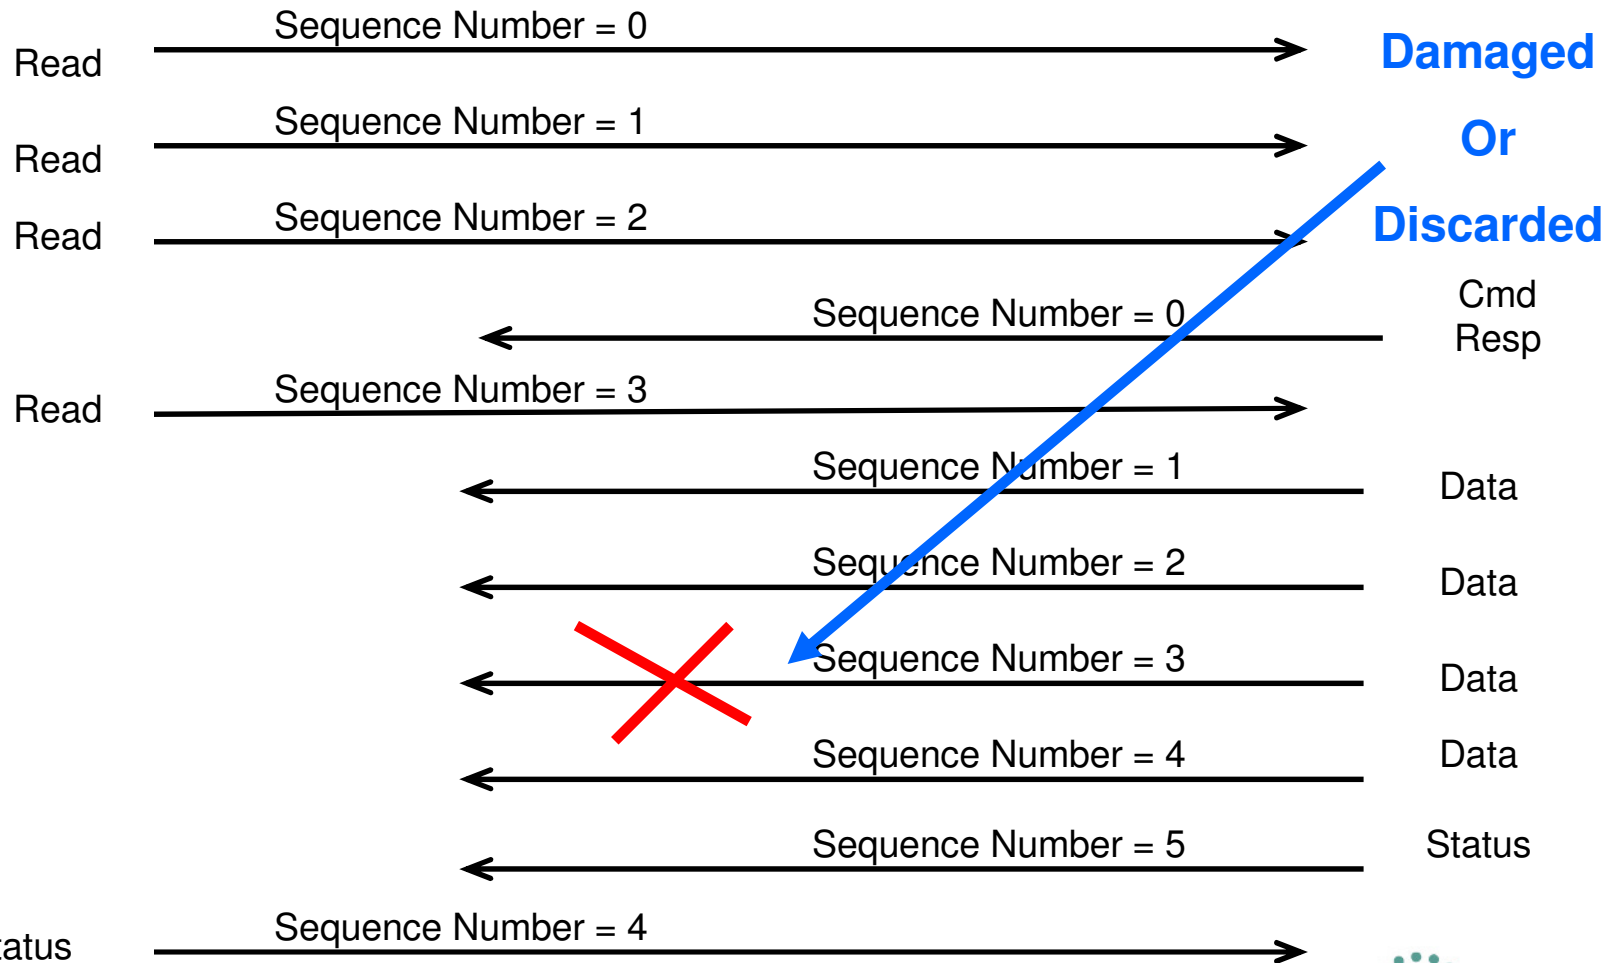
Device Level Errors



Status
224 Accept

Complete your sessions evaluation online at SHARE.org/AnaheimEval

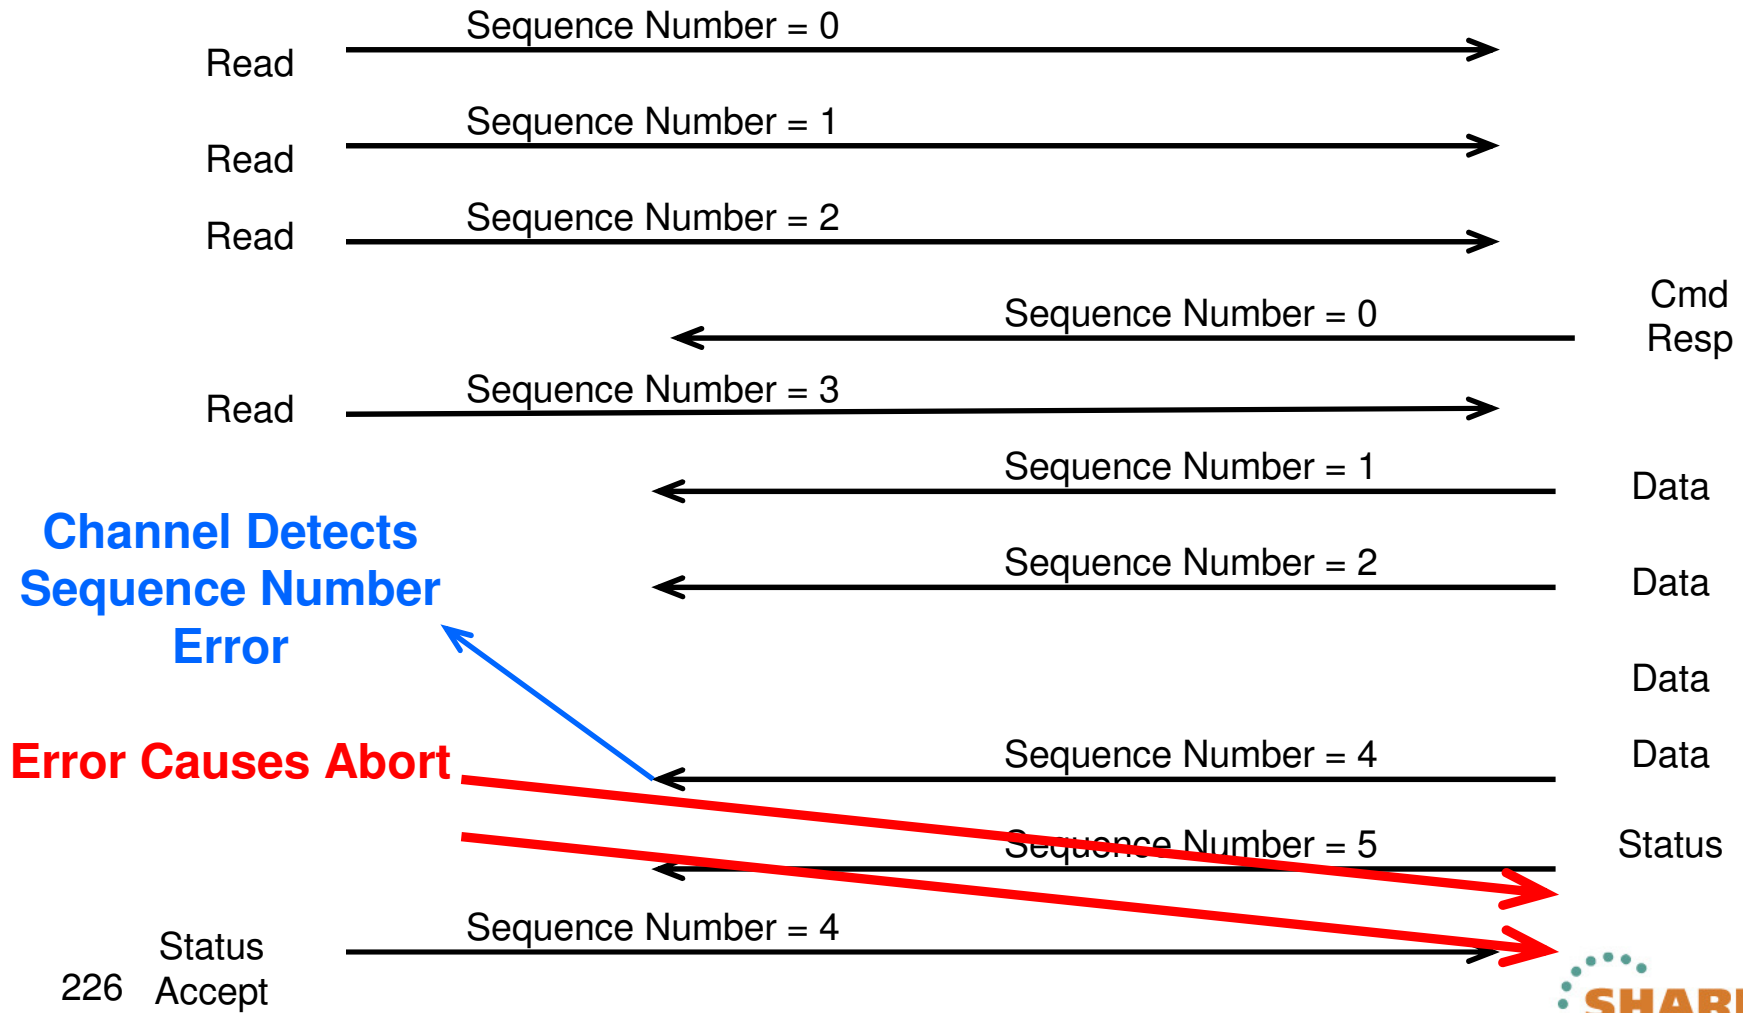
Device Level Errors



Status
225 Accept

Complete your sessions evaluation online at SHARE.org/AnaheimEval


Device Level Errors



226 Accept

Complete your sessions evaluation online at SHARE.org/AnaheimEval

Device Level Errors



* DACK will result in an IFCC (IOS050I) interrupt to software

Link Recovery

- Three initiators:
 - Link Between Channel and Switch 'Fails'
 - State Change from Switch
 - Timeout
- For Remote Links, Channel Attempts a 'Ping'
 - If channel gets a response – all is ok
 - If 'no one home' response, repeat Ping in 4 seconds
 - Still 'no one home' response – declare link dead
- 5 Link fails in 5 Minutes → Flapping Link Threshold

IOS001E dddd,INOPERATIVE PATHS pp pp pp

IOS2001I dddd,INOPERATIVE PATHS

STATUS FOR PATH(S) pp,pp,pp....

LOGICAL PATH IS REMOVED OR NOT ESTABLISHED (A0)

LINK RECOVERY THRESHOLD EXCEEDED FOR LOGICAL PATH (06)

Internal Errors

- Firmware detected
 - “Shouldn’t be able to get here from there”
- Hardware detected
 - Parity Error
 - Cross Check
- SAP Detected
 - Timeouts
 - Out of context messages

Internal Error Recovery

- Channel Hardware Reset
- Firmware Reloaded
- Channel is Re-Initialized
- Any in-progress Operations Terminated

Thank You For Your Time And Attention

