#SHAREorg

SHARE

I/O Synergy – The Whole is Greater than the Parts (I/O Performance)

Dale F. Riedy

IBM

riedy@us.ibm.com

March 13, 2014

Session 15081

SHARE in Anaheim

Introduction

- Synergy focus areas
- I/O Performance Related Topics
- For each topic:
  - Overview
  - Pointers/Tips
  - Triva/Interesting Facts

I/O and Storage Constraints

Focus Area	Description	Client Impact
I/O Rates and Bandwidth	Ability to provide a balanced system for MIPs and I/O requests within the existing footprint and power envelope	Inability to achieve high CPU utilization because of I/O delays and concurrency limits. Inability to consolidate more physical footprints. Inability to grow the workloads on z/OS.
Latency	Fit within the existing batch window, ability to handle workload spikes	Bottlenecks will force clients to split workloads and spend resources managing around the I/O limits
Scalability	Number of devices, channels, devices per channel, aliases, I/O slots for accelerators, addressable storage capacity	Inability to add additional work and data to the existing z/OS infrastructure
Resiliency	Ability to recover from a channel path, device, storage controller, or site failure with minimal impact to the system and applications.	Application delays and outages
Complexity	I/O configuration planning, tuning, problem determination, capacity planning	Added cost and risk. Frequency of manual effort increases chance for errors, degrades resilience, and adds cost.

Technology Evolution (I/O Performance)

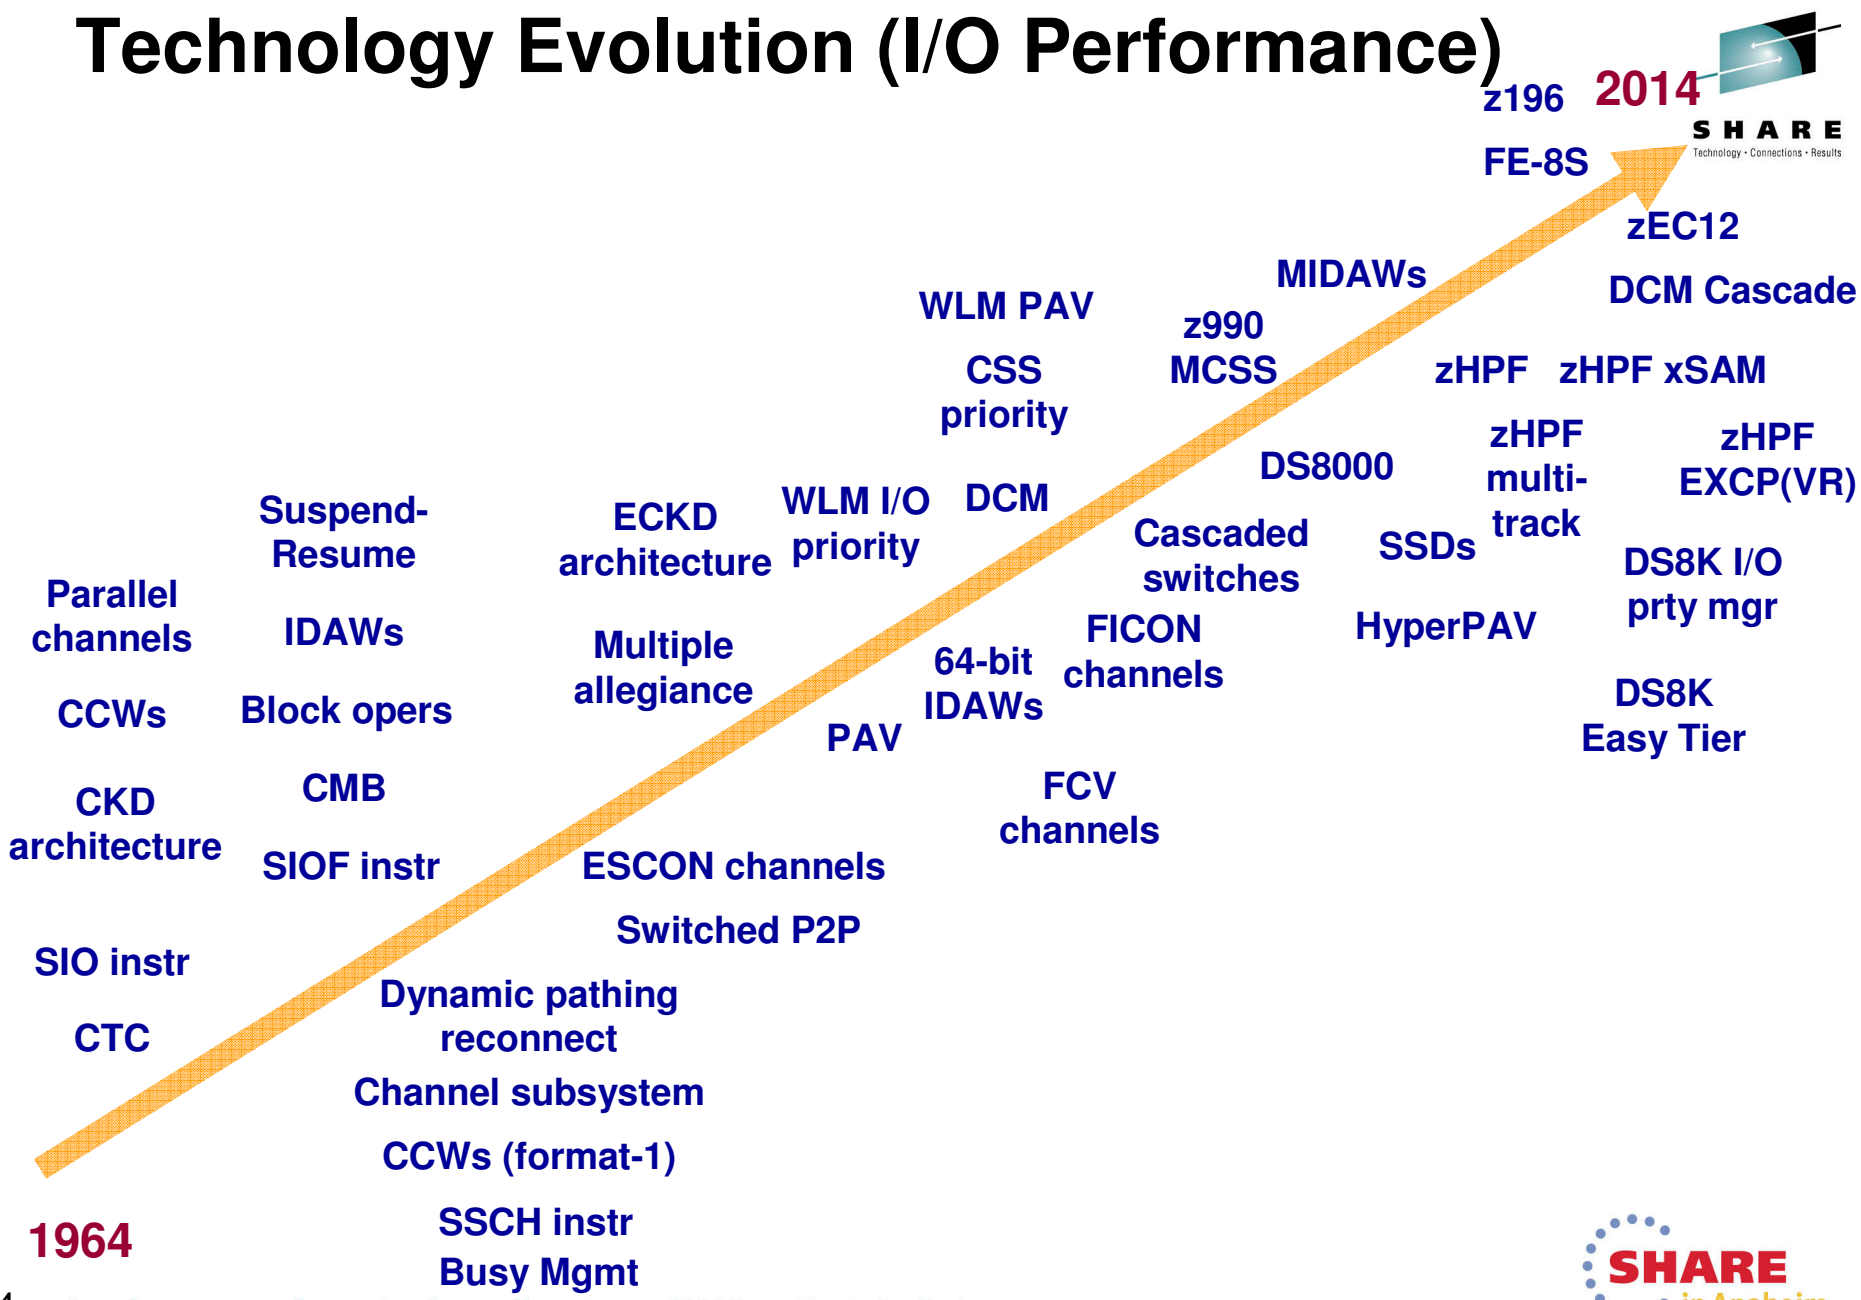

Technologies to Improve I/O Service Time

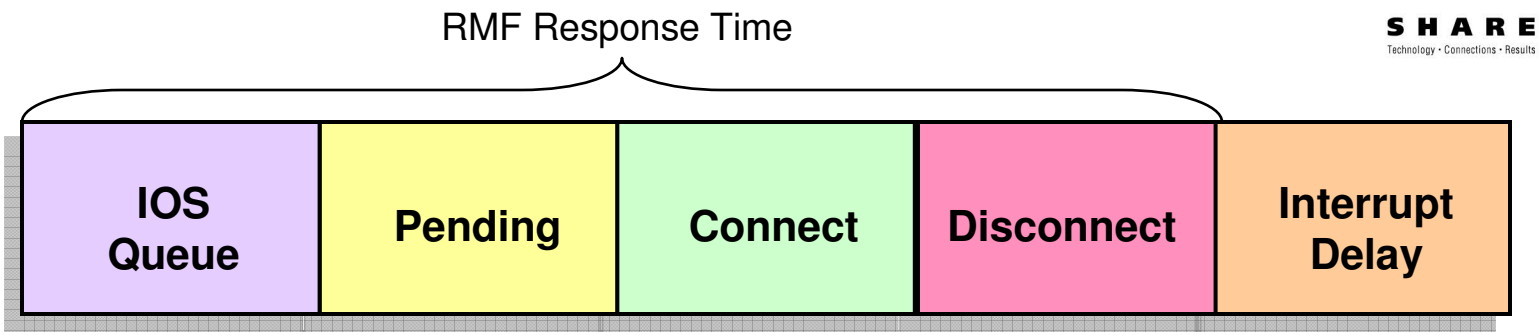

Response Time Components

Contributors

- High tran rates
- Long busy conditions
- Other I/O delays
- Device reserved
- CU recovery
- Fabric congestion
- Distance
- High port / HA utilization
- Extent conflict
- FICON channel multi-plexing
- Out of buffer credits
- Cache miss
- NVS destage
- PPRC delay
- Write pacing (XRC)
- LPAR weights
- Capping
- CPENABLE

Technologies

- HyperPAV
- WLM I/O priority
- zHPF
- CSS priority
- FICON DCM
- FICON express 8S
- zEC12 channel selection
- zHPF
- MIDAWs
- FICON express 8S
- CU cache alg
- Easy Tier
- SSDs
- PPRC FCP pre-deposit writes
- I/O prty mgr
- Hiper-dispatch

High Performance FICON for System z (zHPF)

zHPF is a critical element of the System z I/O strategy

How does zHPF improve I/O performance?

- Leverages optimizations made in the HBA for FCP
- Uses a protocol that has less chit-chat between the channel and device
- Packages entire command stream into one or two sequences
- Packages data more efficiently into frames (less buffer credits used)
- H/W data router on FICON express 8S provides additional acceleration

zHPF Evolution

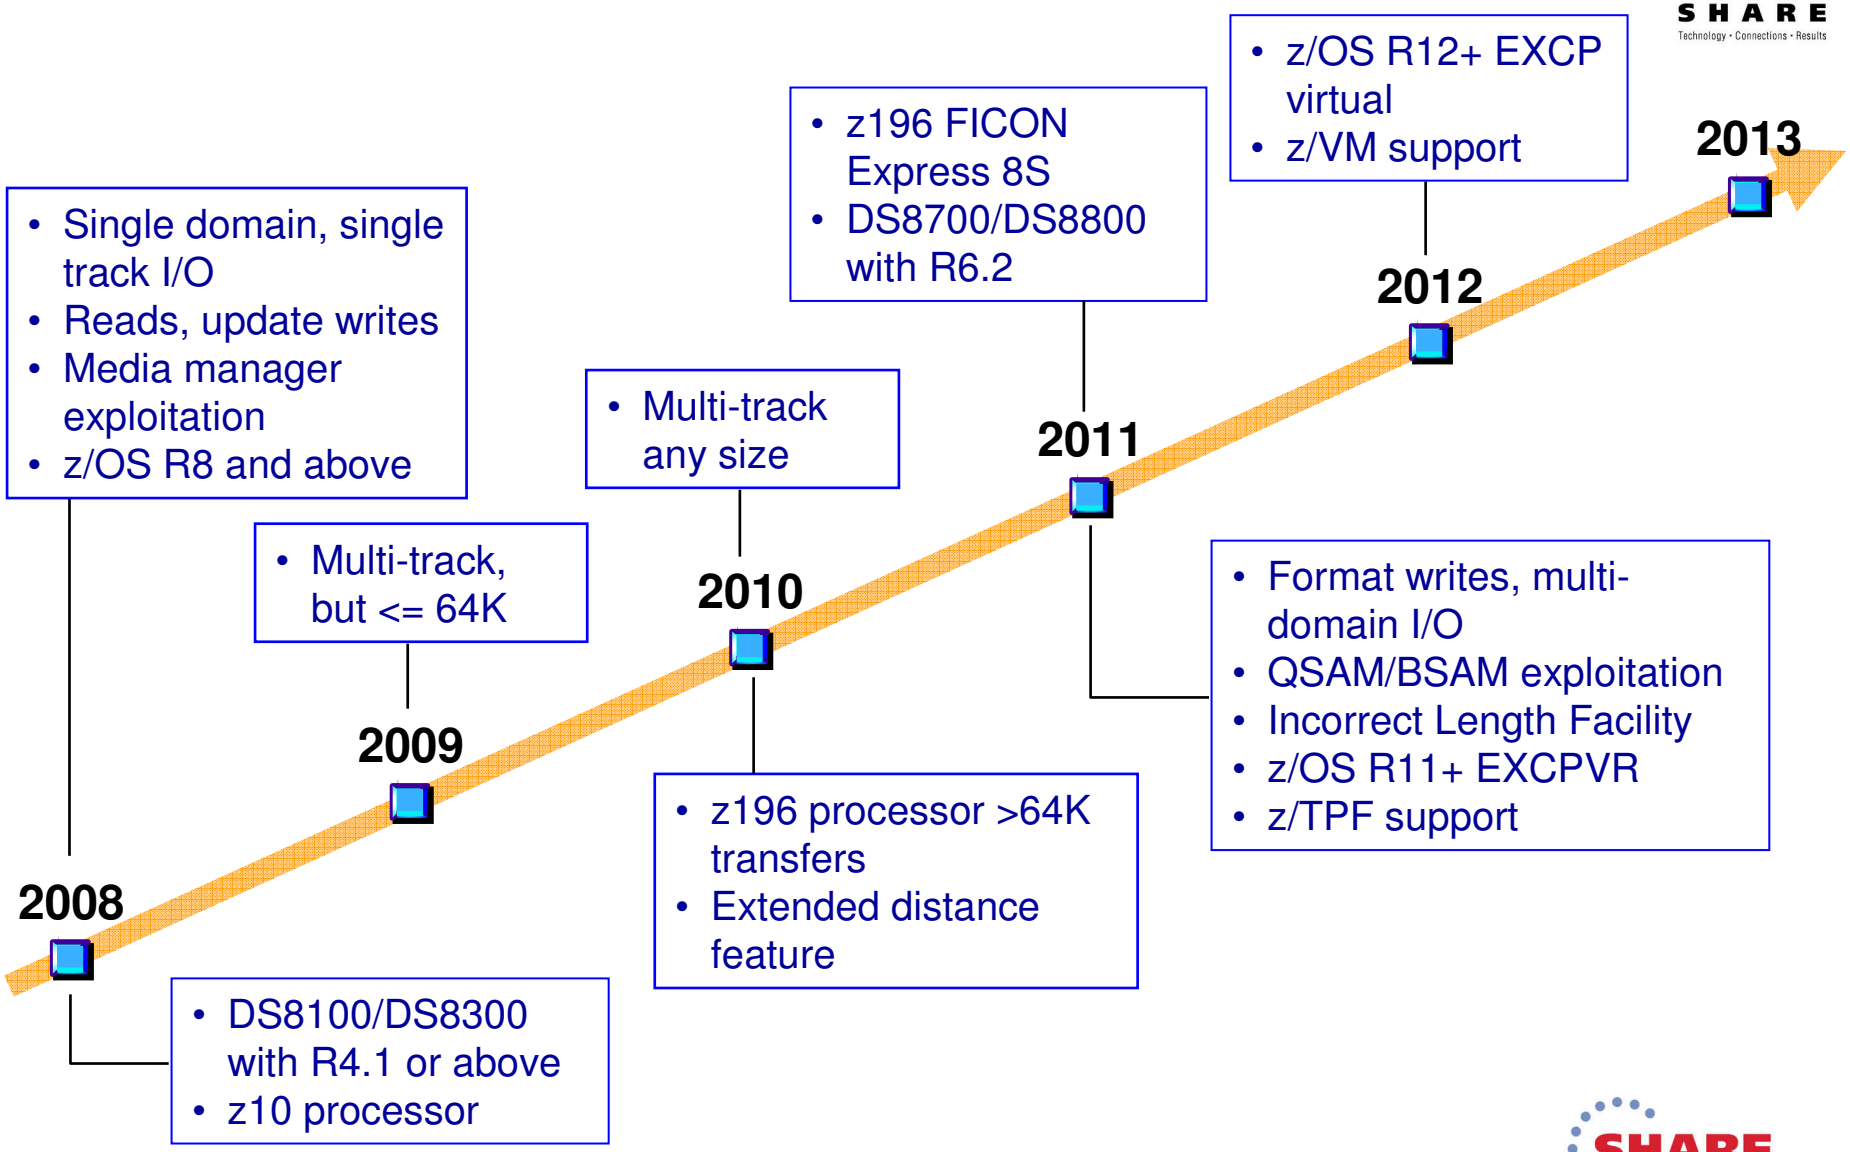

zHPF/FICON/ESCON Comparison

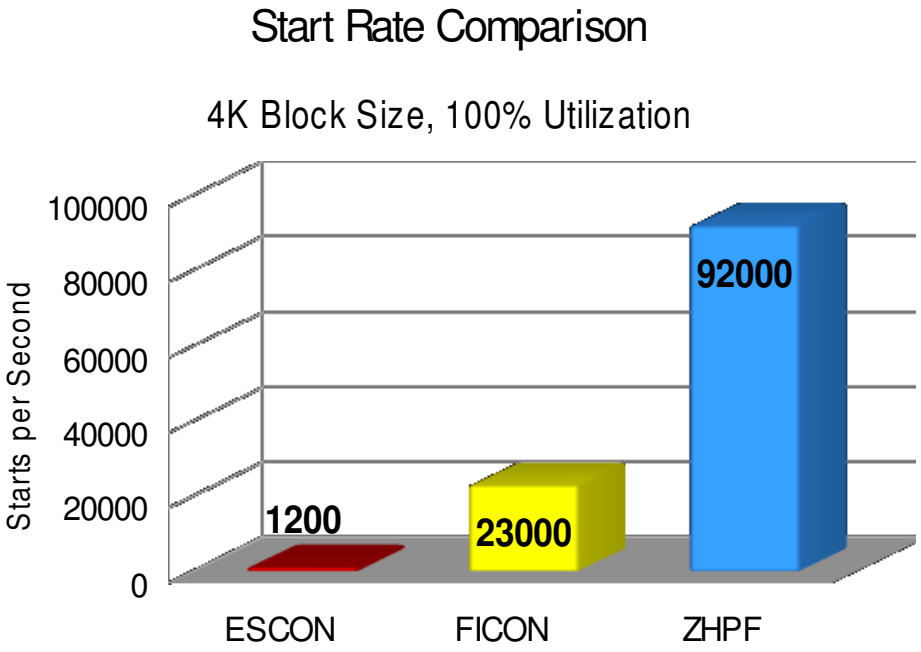

4X the start rate of FICON

2.5X the throughput of FICON

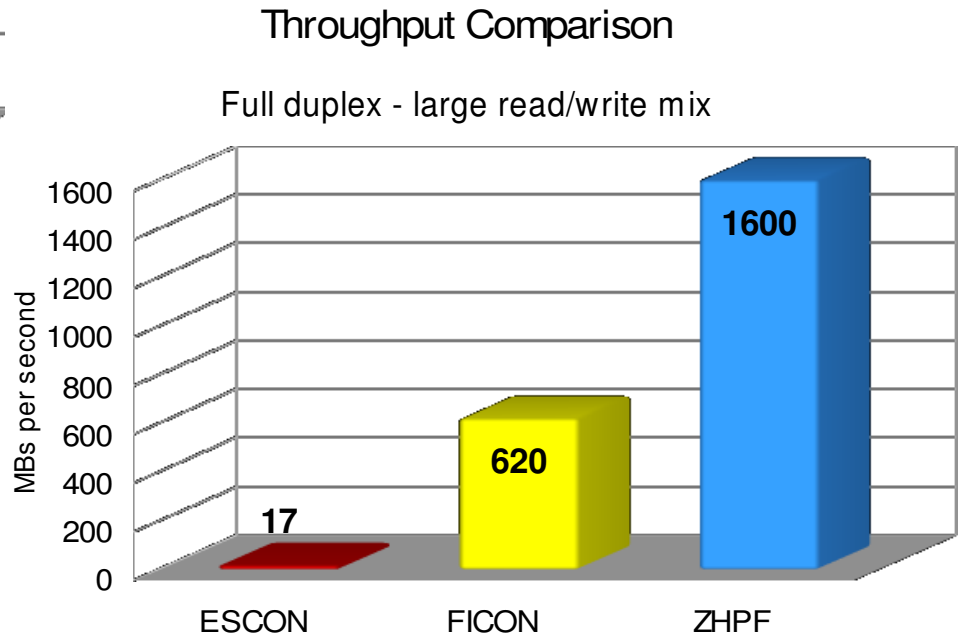

★ESCON only supports half duplex

zHPF Trivia

zHPF is the first time we have dramatically changed the format of a channel program since System 360. Changes have been made to basic CCW channel program (e.g., format-1 CCWs, MIDAWs), and major changes have been made to the underlying architecture (XA, ESCON, FICON), but the channel program format was not radically changed.

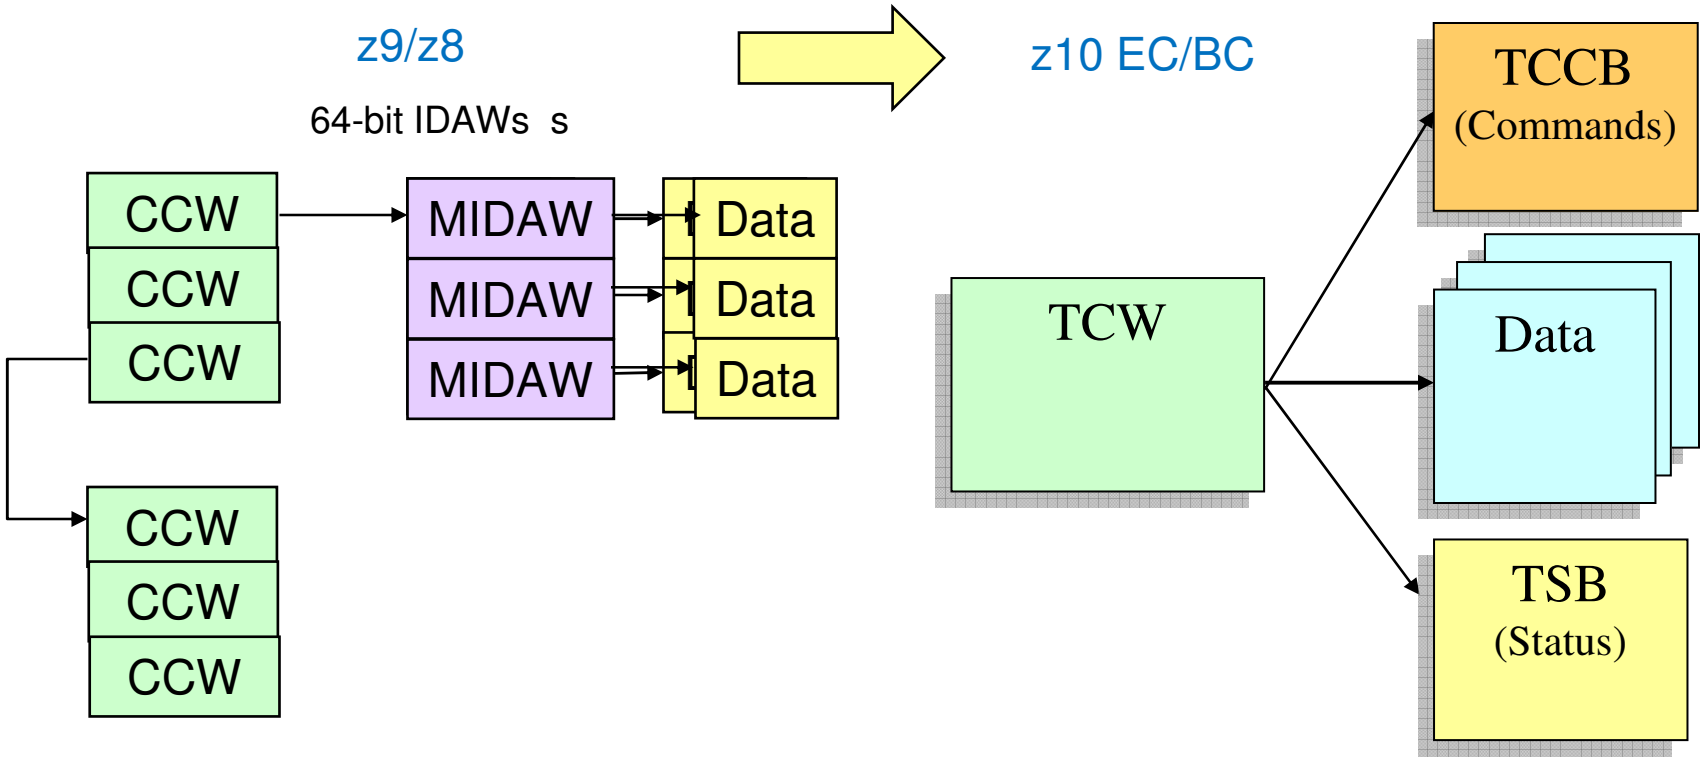

FC Link Protocol for FICON

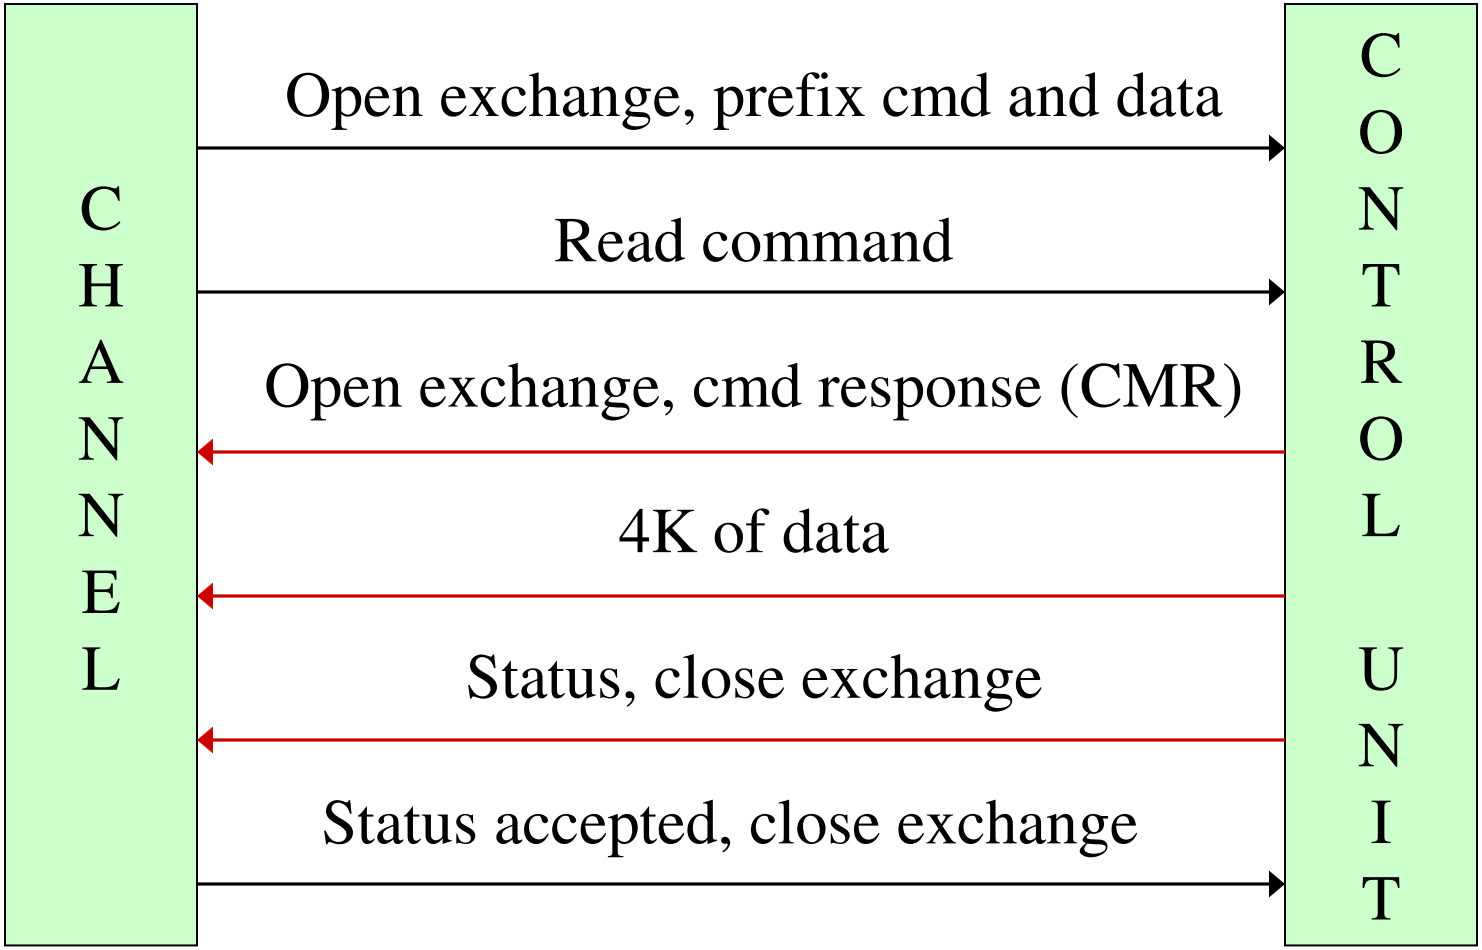

Example of a 4K read I/O request (6 sequences, 8 frames)

FC Link Protocol for zHPF (and FCP)

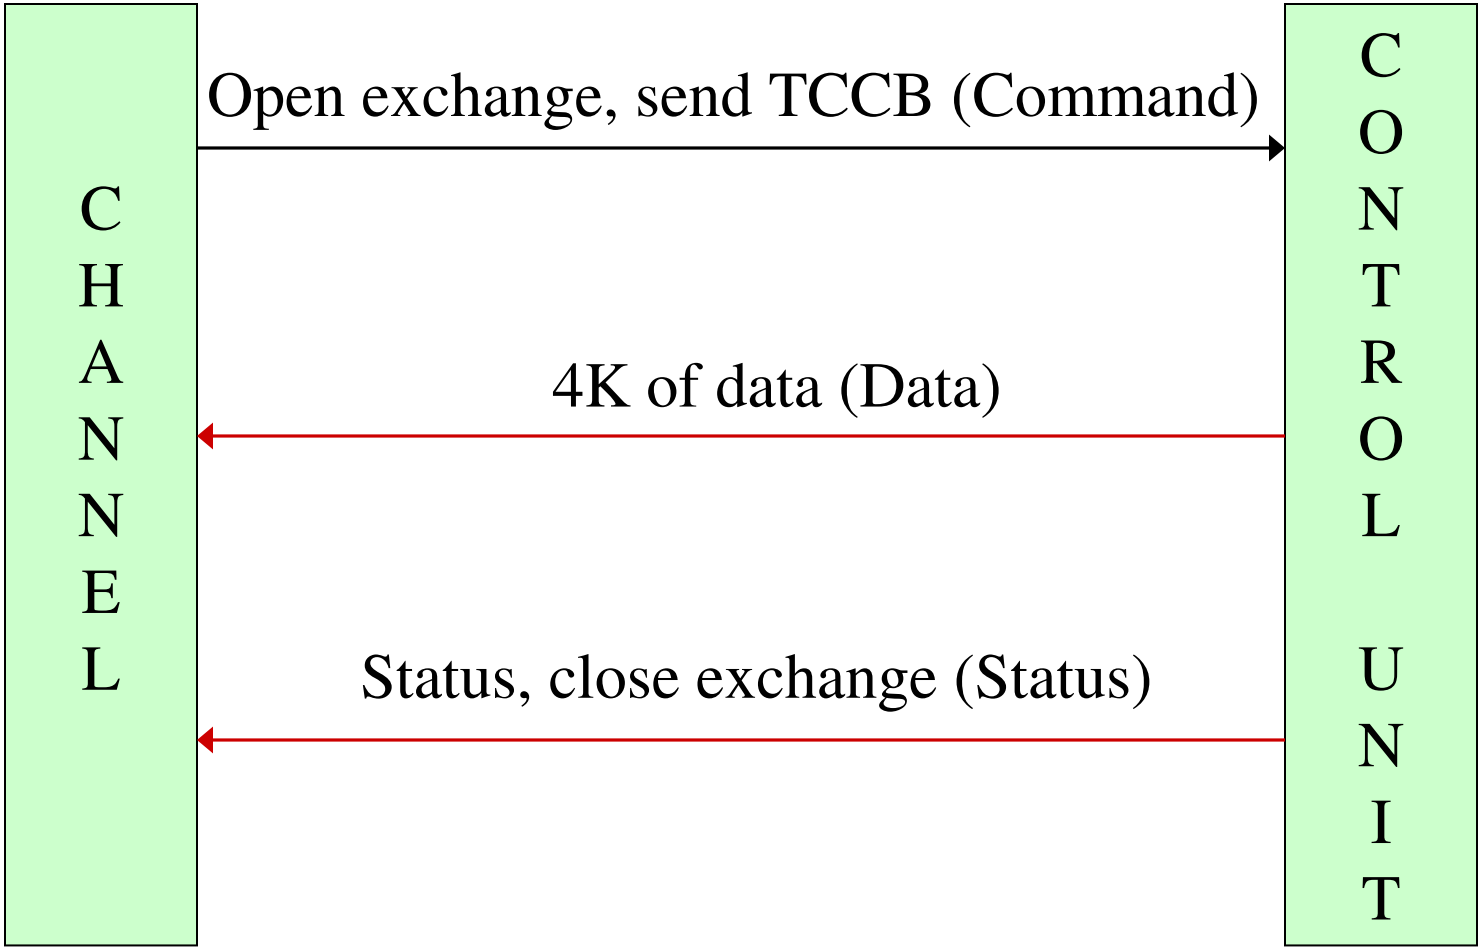

Example of a 4K read I/O request (3 sequences, 5 frames)

FICON – Read 3 Contiguous Tracks (12x4K Records)

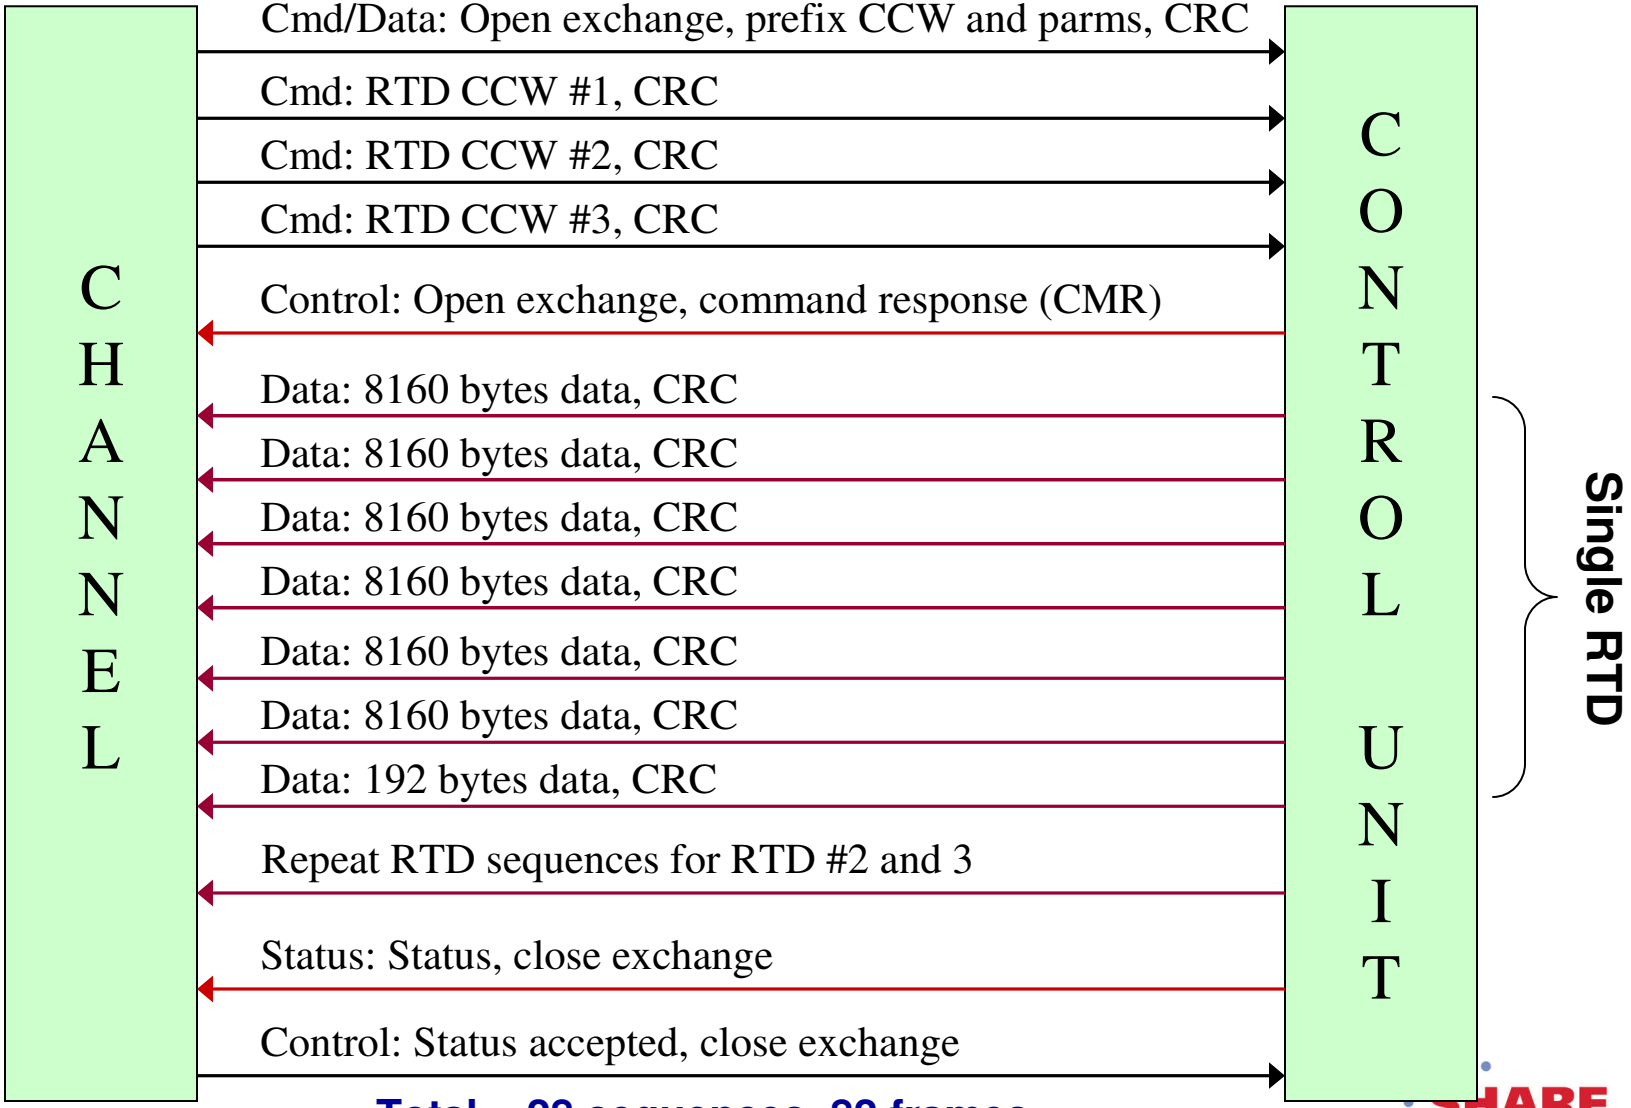

Total = 28 sequences, 82 frames

zHPF – Read 3 Contiguous Tracks (12x4K Records)

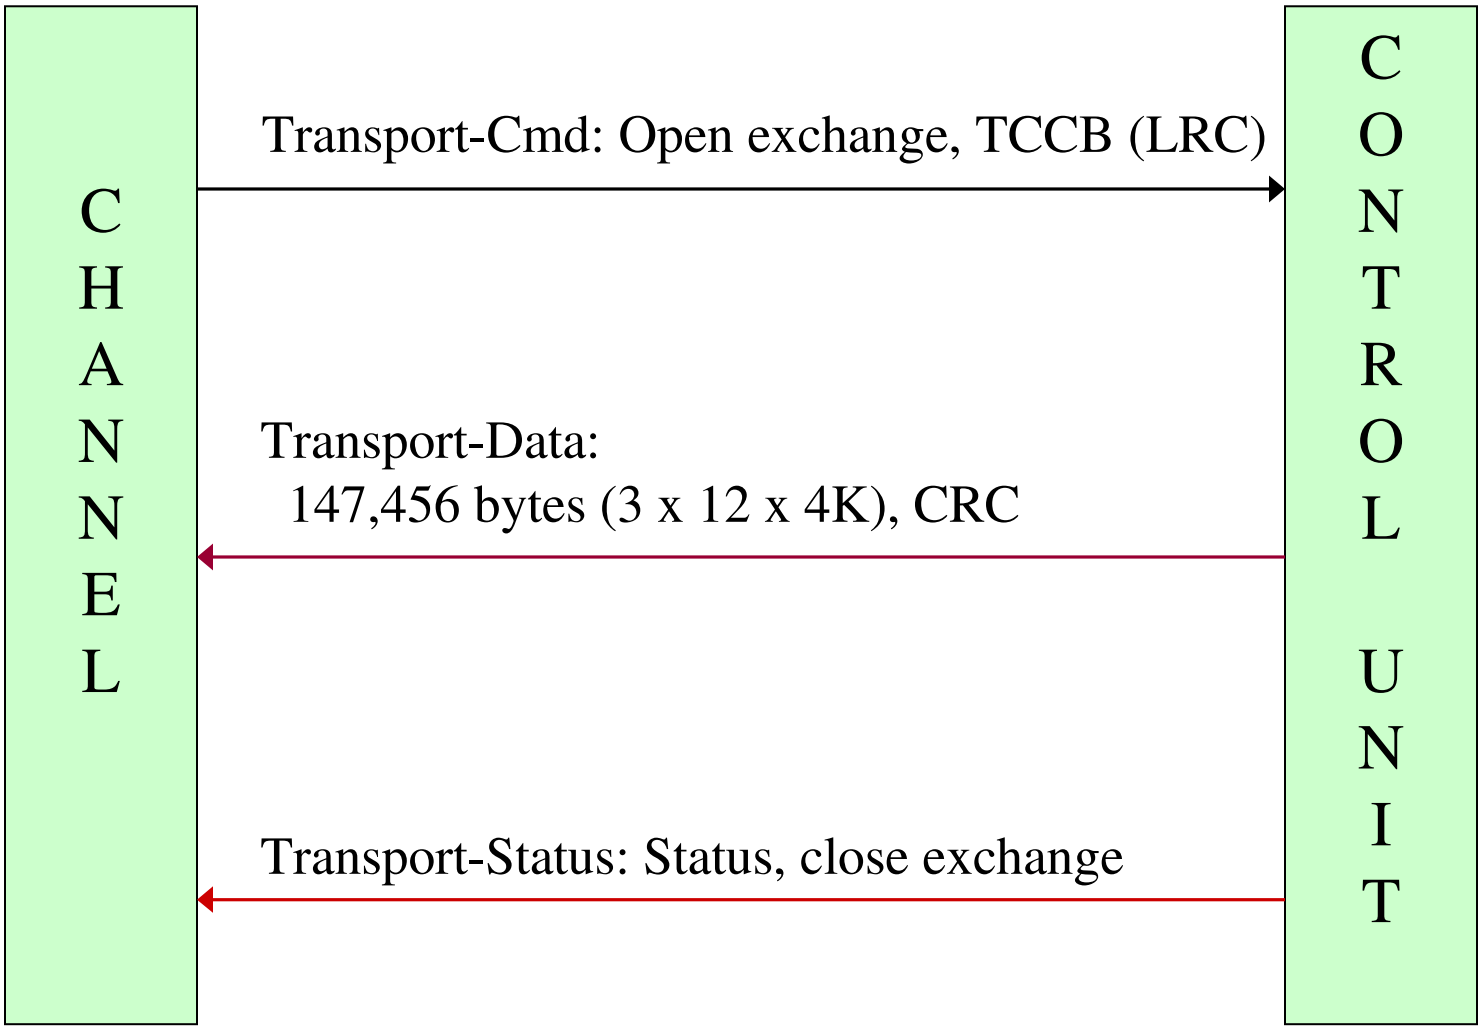

Total = 3 sequences, 75 frames

CCW Channel Programs – Virtually Unlimited in Size

A single CCW channel program is virtually unlimited in size since CCWs can be chained together implicitly or explicitly. The size of a channel program is only limited by the amount of 31-bit virtual and real storage available.

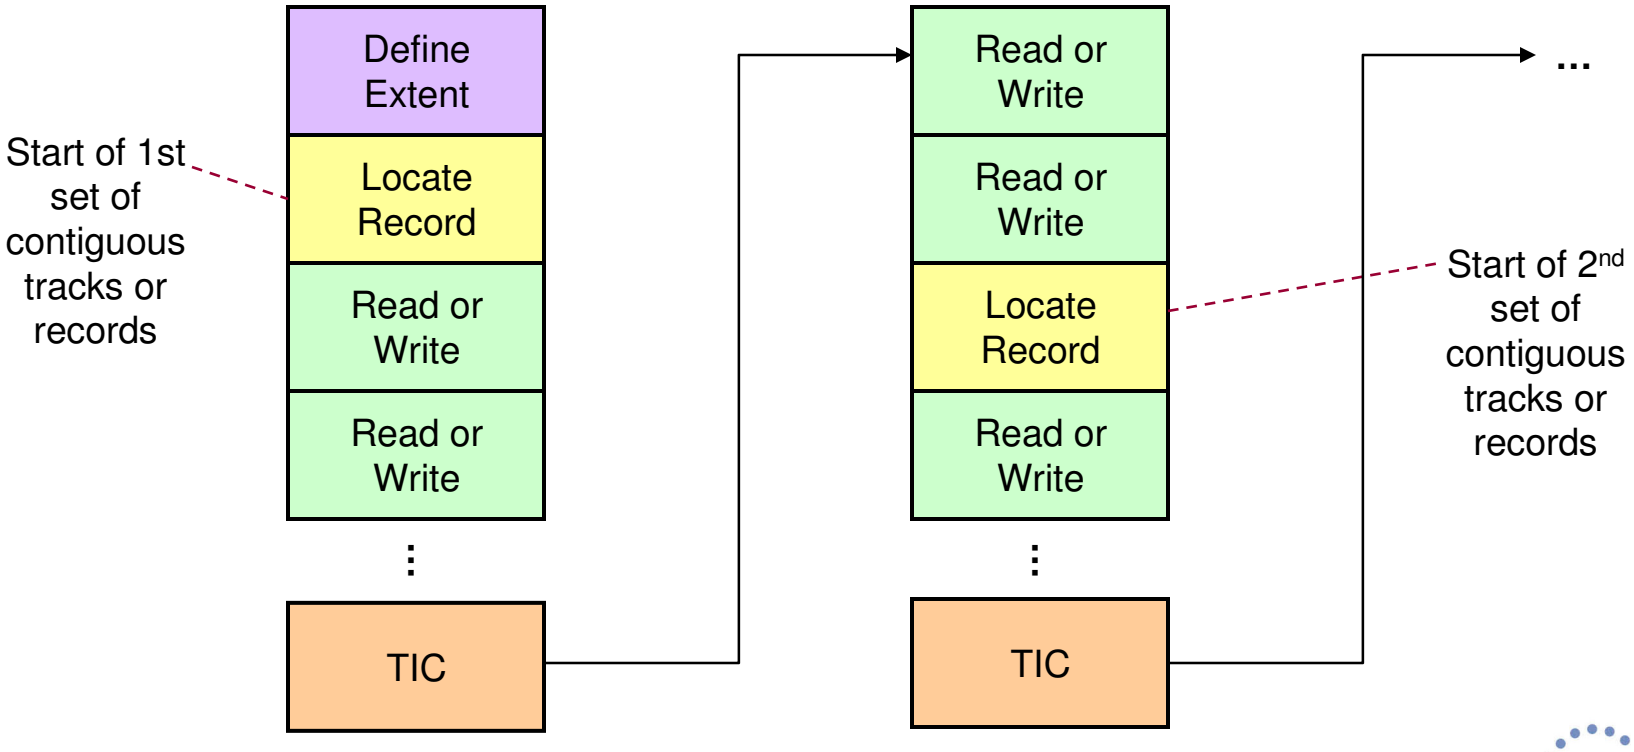

zHPF Channel Programs – Limited by Architecture

The command block for zHPF is basically a SCSI CDB, which limits the size to 240 bytes. How can we pack enough information into a channel program so that it can be used for things like DB2 list prefetch and QSAM/BSAM I/O?

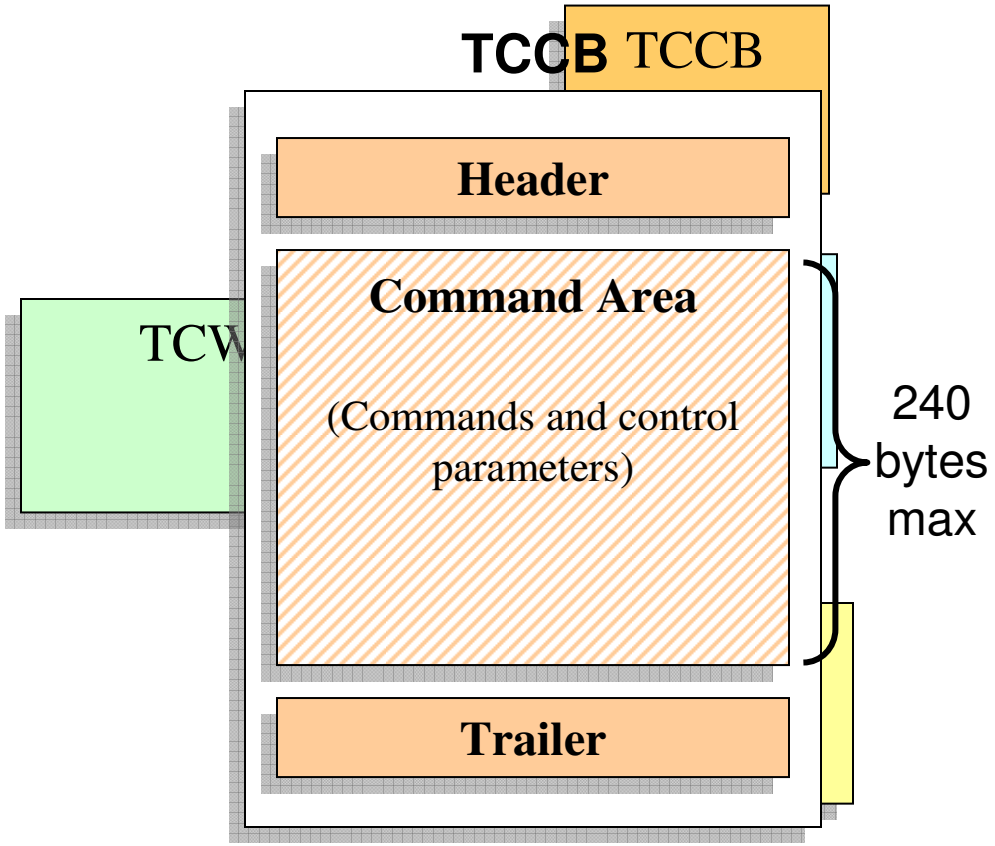

Step 1 – Collapse Lots of CCWs

For CCW channel programs, z/OS collapses the define extent and locate record CCWs into a single CCW. For zHPF, we can also collapse the read/write CCWs as well as the subsequent locate records into a smaller area.

★Up to 22 extents at 255 records or tracks per extent

★22 * 255 * 48K (per track) = 262 MB

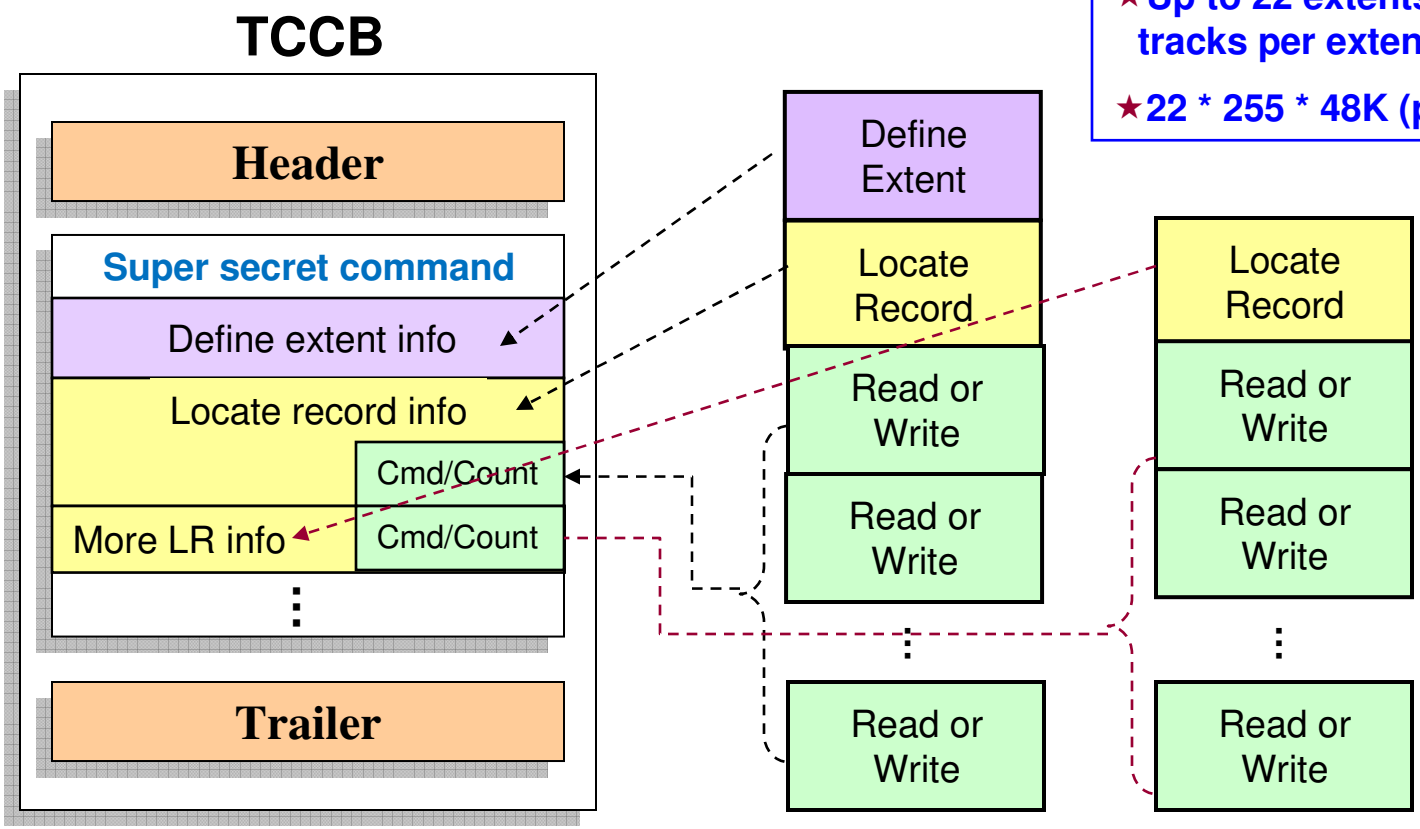

Step 2 – Extend Commands into Output Stream

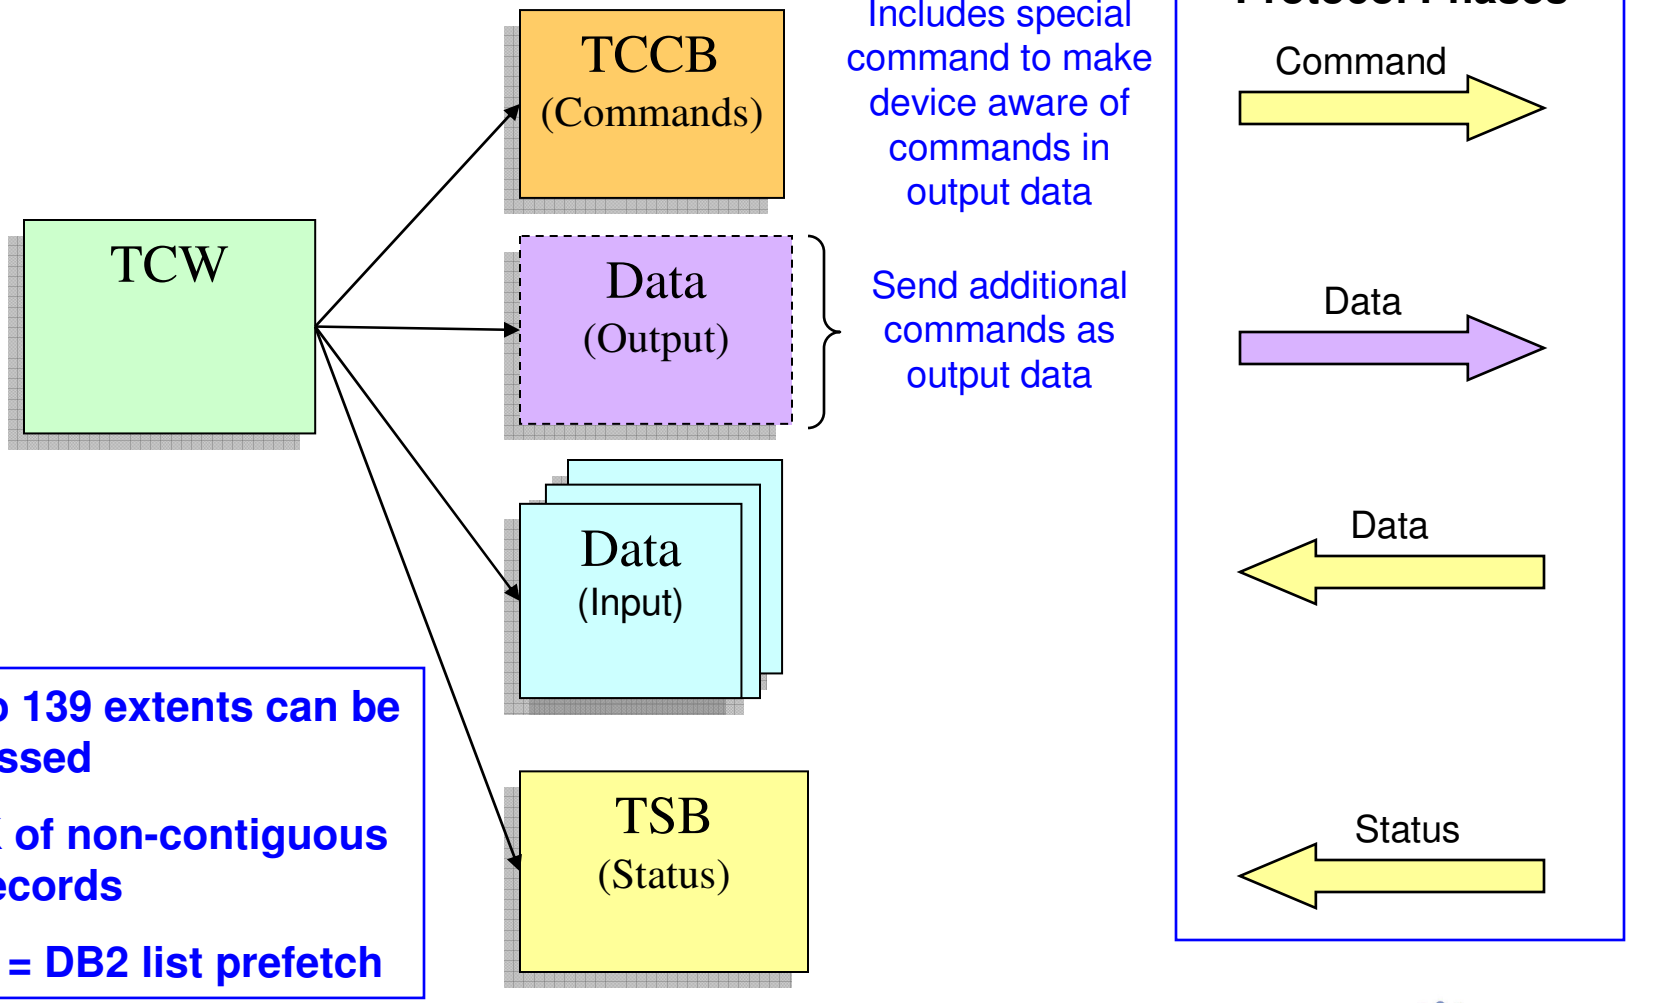

- ★ Up to 139 extents can be accessed
- ★ 556K of non-contiguous 4K records
- ★ User = DB2 list prefetch

zHPF Tips – How to Tell if zHPF is Used

- Channel path activity report shows FICON vs zHPF activity
  - This is the only place where zHPF usage is recorded

CHANNEL PATH			UTILIZATION(%)			READ(MB/SEC)		WRITE(MB/SEC)		FICON OPERATIONS			ZHPF OPERATIONS		
ID	TYPE	G SHR	PART	TOTAL	BUS	PART	TOTAL	PART	TOTAL	RATE	ACTIVE	DEFER	RATE	ACTIVE	DEFER
00	FC_S	5 Y	100.00	100.00	0.84	0.13	2.15	0.17	2.68	61.5	1.7	0.0	4.6	1.0	0.0
01	FC_S	5 Y	100.00	100.00	0.85	0.13	2.21	0.13	2.69	61.3	1.8	0.0	4.7	1.0	0.0
02	FC_S	4 Y	0.14	2.30	0.85	0.10	2.17	0.13	2.70	61.3	1.3	0.0	4.6	1.0	0.0
03	FC_S	4 Y	0.13	2.27	0.84	0.11	2.14	0.13	2.66	60.0	1.3	0.0	4.4	1.0	0.0
04	FC_S	5 Y	0.13	2.24	0.82	0.10	2.07	0.13	2.63	59.4	1.7	0.0	4.4	1.0	0.0
05	FC_S	5 Y	0.13	2.25	0.83	0.10	2.11	0.12	2.66	59.1	1.7	0.0	4.2	1.0	0.0
06	FC_S	4 Y	0.12	2.23	0.83	0.10	2.09	0.13	2.68	58.7	1.3	0.0	4.2	1.0	0.0

zHPF I/Os per second (physical channel)

Average number of active I/Os (open exchanges)

Average number of deferred I/Os (no available open exchanges)

zHPF Tips – Is zHPF Enabled?

```
D IOS,ZHPF
IOS630I 13.35.33 ZHPF FACILITY 742
HIGH PERFORMANCE FICON FACILITY IS ENABLED
```

```
D M=DEV(3100)
IEE174I 11.32.54 DISPLAY M 965
DEVICE 3100    STATUS=ONLINE
CHP                    A4    B4    D0
ENTRY LINK ADDRESS     22    2323  B3
DEST LINK ADDRESS      42    23B2  42
PATH ONLINE            Y     Y     Y
...
FUNCTIONS ENABLED = MIDAW, ZHPF
```

```
D M=CU(3100)
IEE174I 11.08.13 DISPLAY M 226
CONTROL UNIT 3100
CHP                    A4    B4    D0
ENTRY LINK ADDRESS     22    23    B3
...
ZHPF - CHPID           Y     Y     Y
ZHPF - CU INTERFACE    Y     Y     Y
...
FUNCTIONS ENABLED = ZHPF
```

- Enables zHPF at a system level, default is NO
- Device could be disabled for any number of reasons including CU LIC feature
- Certain I/O errors can disable zHPF at the device or CU level (these are usually a sign of a microcode problem)
- zHPF CHPID/CU interface should always be Y unless it is connected to an ESCON channel/interface or FICON express
- If zHPF doesn't appear in functions enabled, then either LIC feature not installed or microcode problem

zHPF Tips – Is zHPF Enabled?

```
D SMS,OPTIONS
IGD002I 13:36:10 DISPLAY SMS 750
ACDS      = SYS1SMS.DR5.ACDS
COMMDS    = SYS1SMS.DR5.COMMDS
ACDS LEVEL = UNAVAIL
SMS PARMLIB MEMBER NAME = IGDSMSDR
INTERVAL = 15  DINTERVAL = 150
...
OAMTASK =       PDSE_SYSEVENT_DONTSWAP = NO
DB2SSID =       SAM_USE_HPF = YES
...
```

- IGDSMSxx parmlib member indicates whether QSAM/BSAM should use zHPF.

```
DEVSERV,QDASD,7000,RDFEATS
IEE459I 11.57.32 DEVSERV QDASD 848
 UNIT VOLSER SCUTYPE DEVTYPE        CYL  SSID...
07000 ------ 2107951 2107900       1113  7000...
   FEATURE CODES AT V020A0AF8
00070006001F00A6 DDFFFDE164400000 FD80000000000000 C3D000BB00000000
E785C00000000000 AEC0000000000000
```

- zHPF capability/feature bits (+x'28'). At a minimum, bits 0 & 2 should be on. Other bits may also need to be on for other exploitation (e.g, bit 5 for QSAM/BSAM). Contact IBM support for specific questions.

zHPF Trivia

Although the channel supported a maximum transfer size of 2M in the initial zHPF support in the z10 processor, for performance reasons, the software limited the amount of data transferred to 68K. Where did this number come from?

Around 64K is where performance issues occurred.

512 bytes is the minimum VSAM control interval size. 128 of these fit into 64K (2 tracks are required).

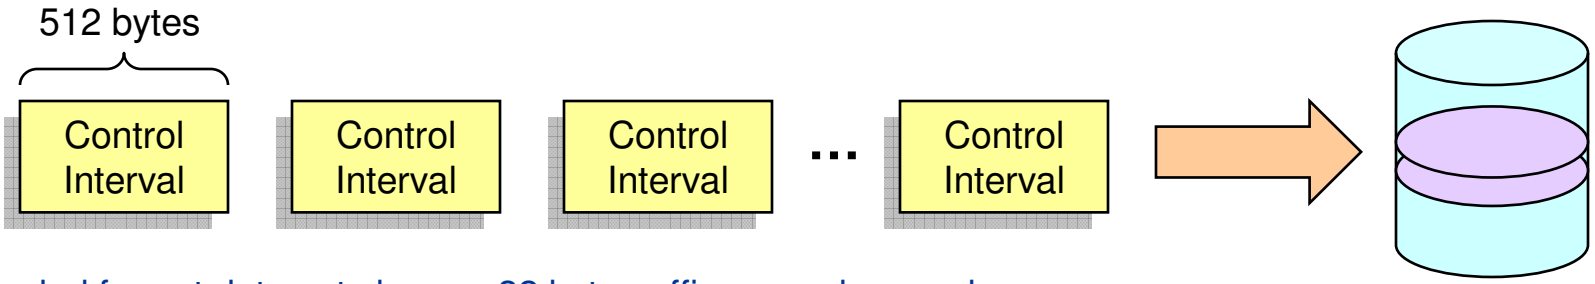

Extended format data sets have a 32 byte suffix on each record.

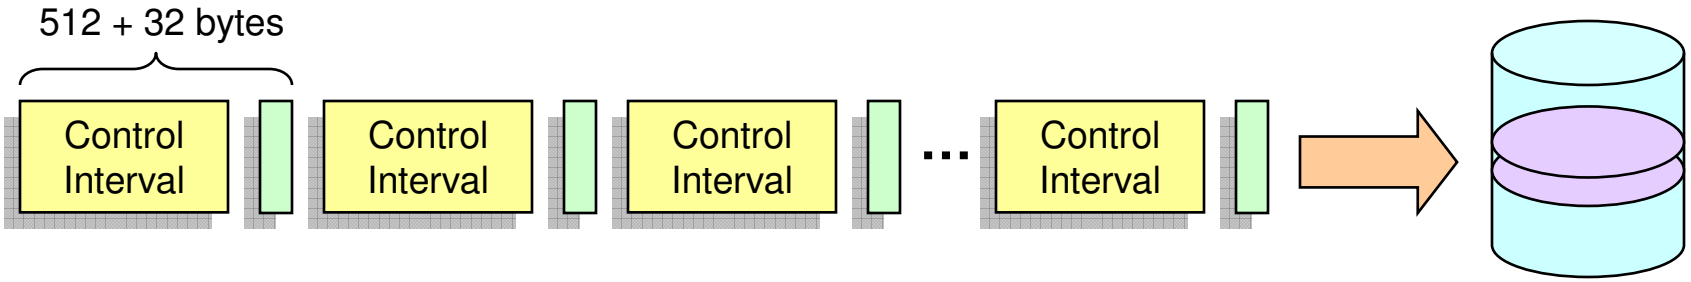

☺ Therefore, we need to add (64K / 512) * 32 = 4096 to limit to allow 64K transfers for extended format data sets to eligible for zHPF.

FICON Express 8S Channel

- zHPF protocol implemented in the hardware via a data router
- 76% higher start rates compared to zHPF on FICON express 8
- 107% higher throughput compared to zHPF on FICON express 8
- **100% of DB2 I/O can be converted to zHPF**

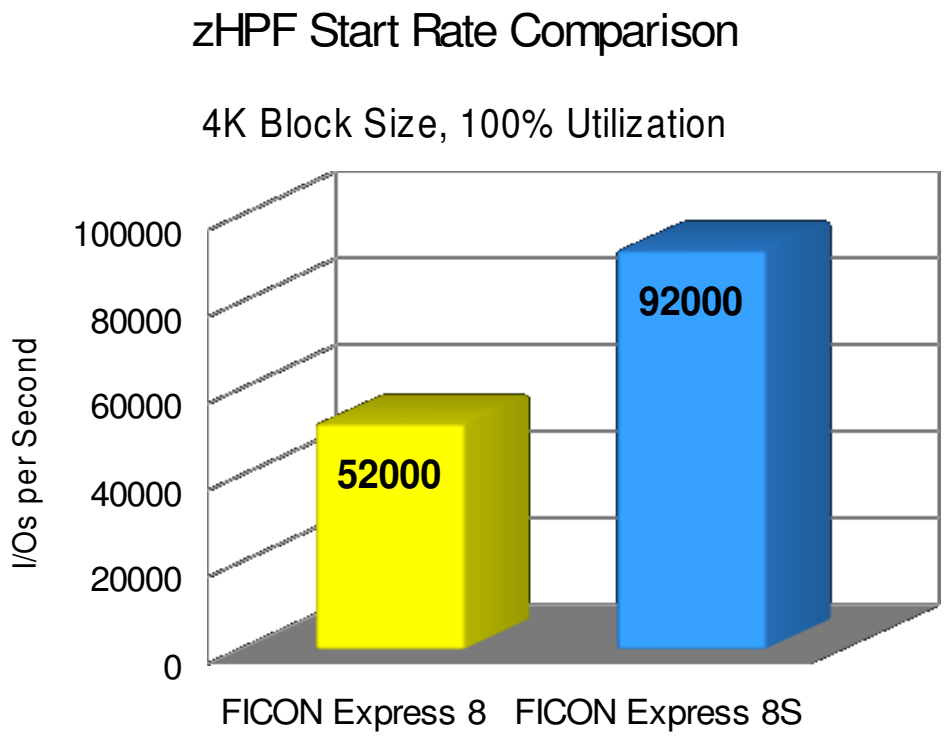

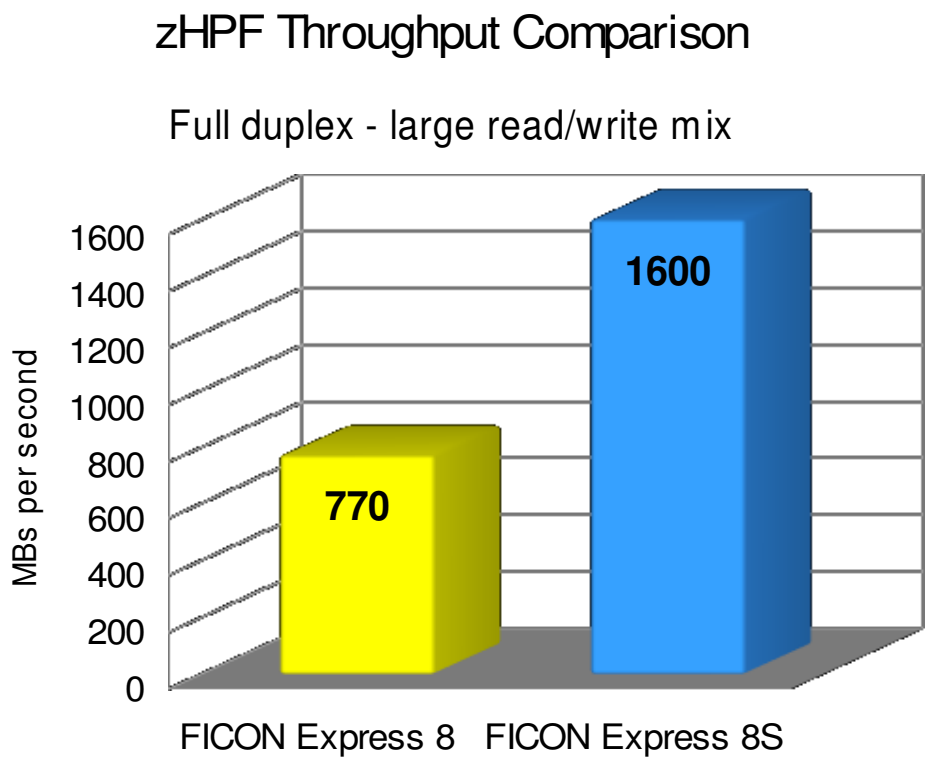

Older FICON Channels – Store and Forward

- In older channels, the HBA did not access System z memory directly. A store and forward technique was used where the channel firmware would fetch information from System z memory into channel memory, which would then be accessed by the HBA.
- A pipeline is established which allowed the channel to work in parallel with the HBA.

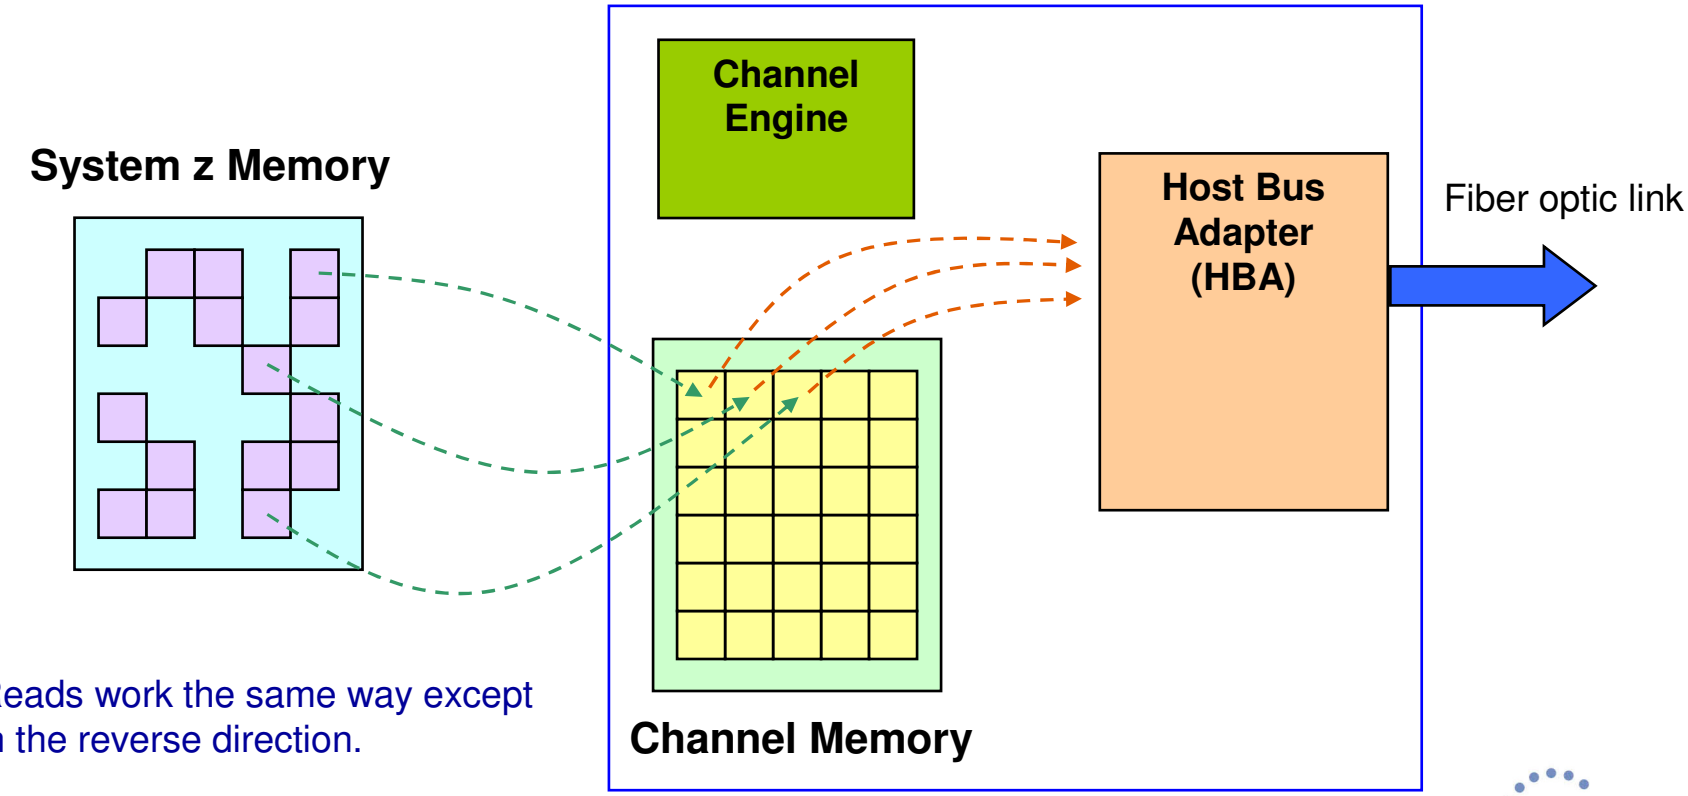

Reads work the same way except in the reverse direction.

FICON Express 8S – H/W Data Router

- With the FICON express 8S channel, the HBA accesses System z memory directly through and under the control of a H/W data router.
- The channel firmware is involved only in setting up the storage areas that need to be accessed in channel memory.

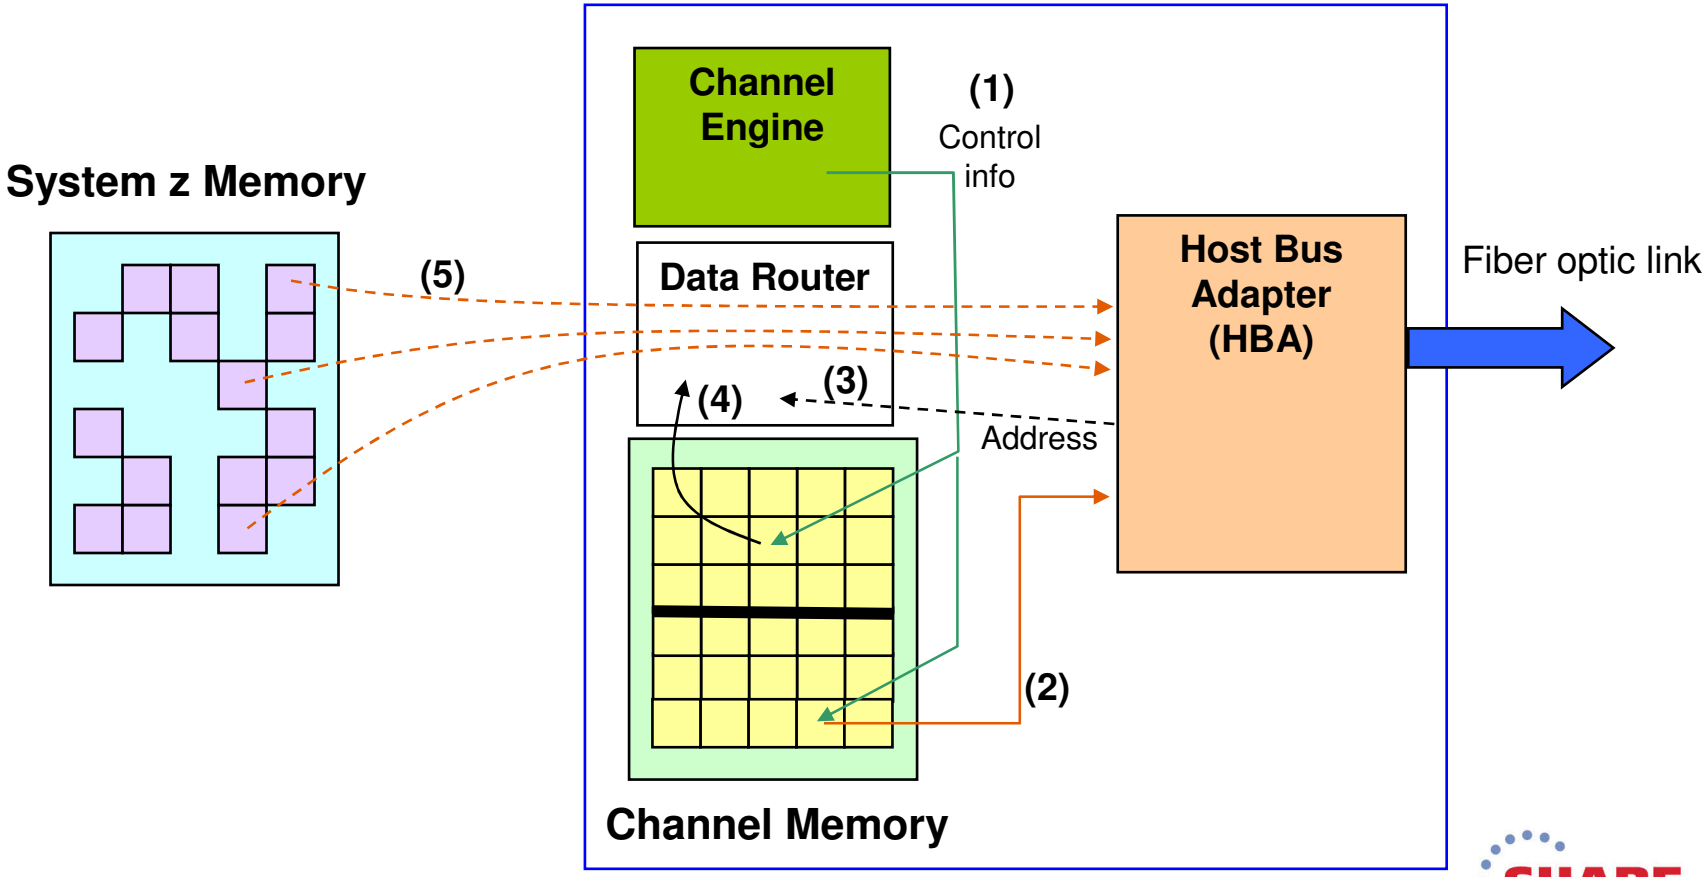

Parallel Access Volumes (PAV) – Before PAV

Although, storage controllers allow multiple I/Os to be issued to a single device from multiple systems (multiple allegiance), System z allows only one I/O to be issued to a device (subchannel in HSA) at a time for an LPAR.

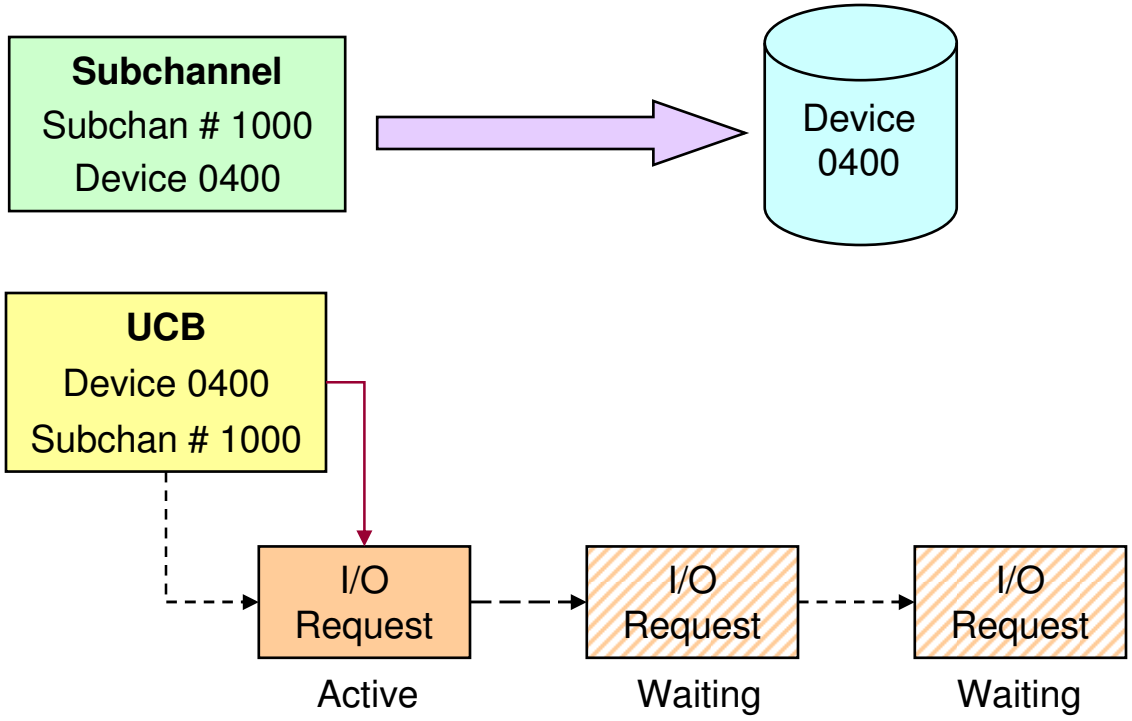

If multiple I/Os arrive for a device at the same time, the remaining I/O's will be queued internally in z/OS while waiting for the prior I/Os to complete (IOSQ time).

Static PAV

Static PAV allowed multiple *alias* devices to be defined for a *base* device. This allowed multiple I/Os to be done to a particular volume in parallel since a separate subchannel was assigned for each alias. However, the alias to base assignment was static, so this required careful planning to ensure that a base had adequate aliases defined.

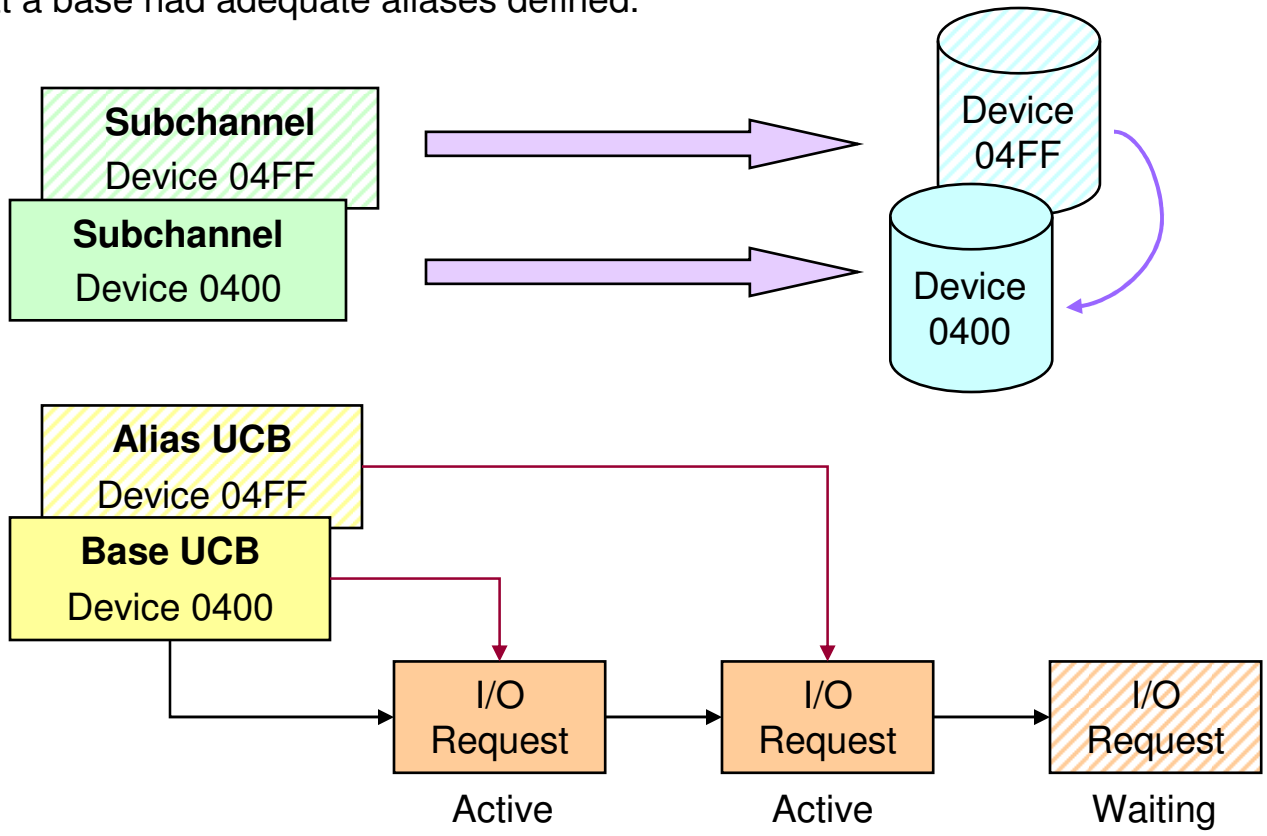

Dynamic (WLM Managed) PAV

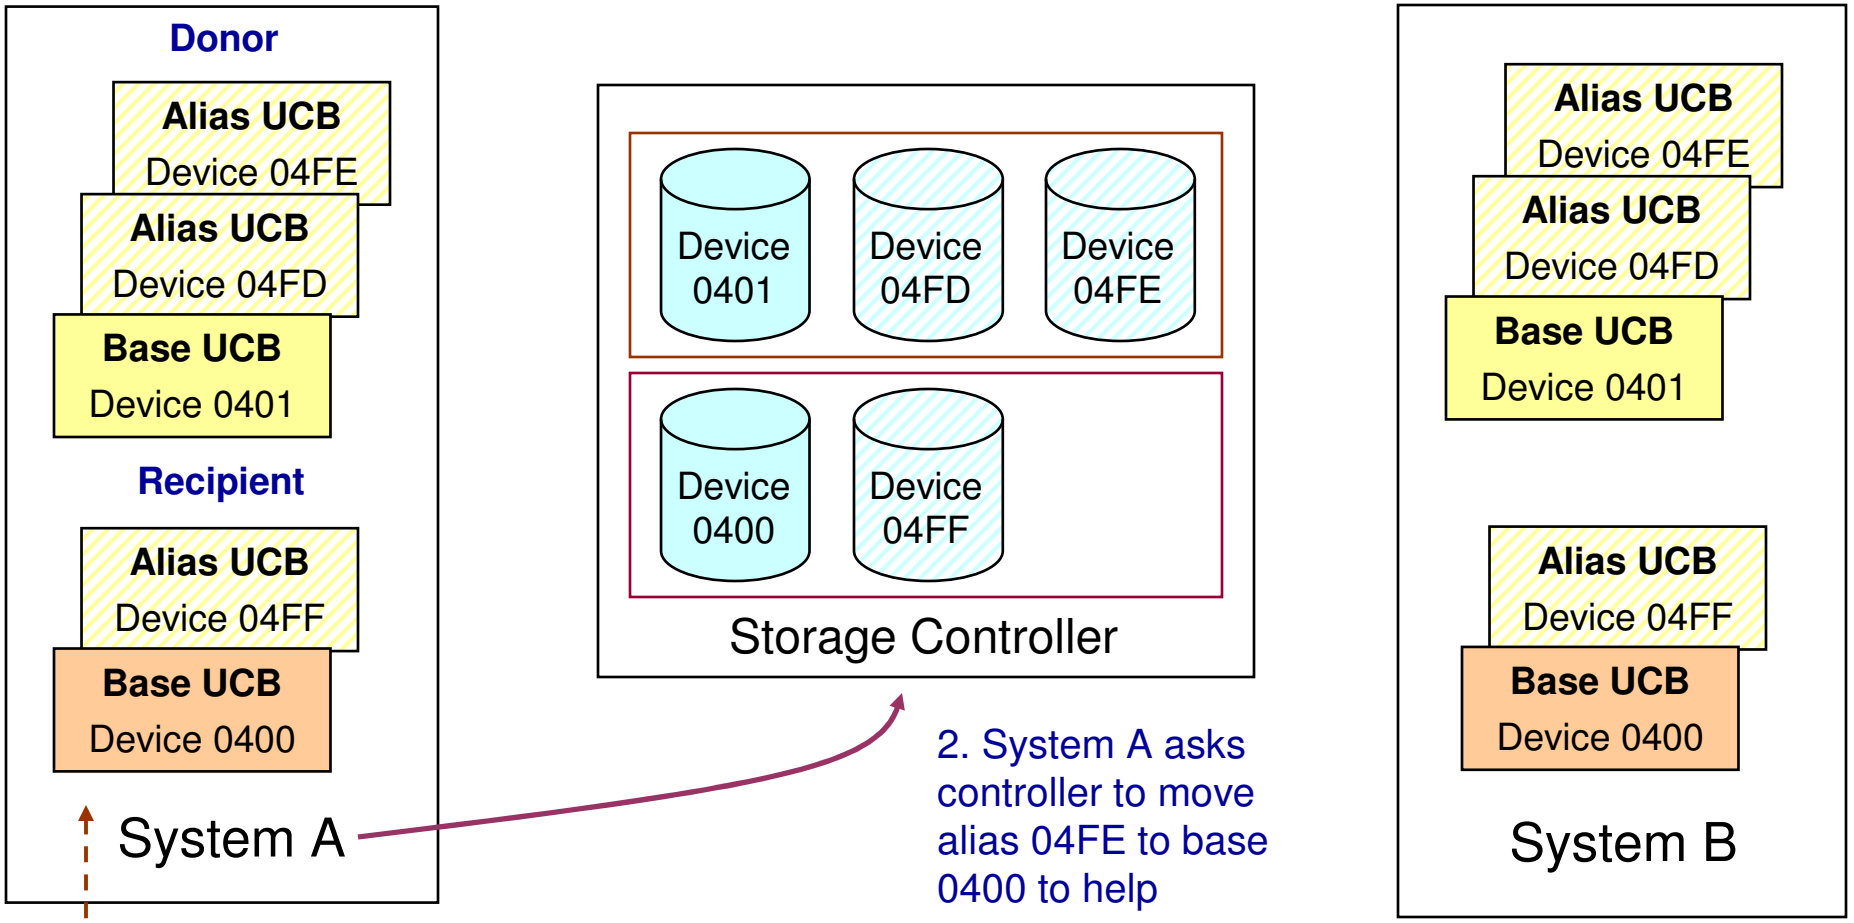

Dynamic (WLM Managed) PAV

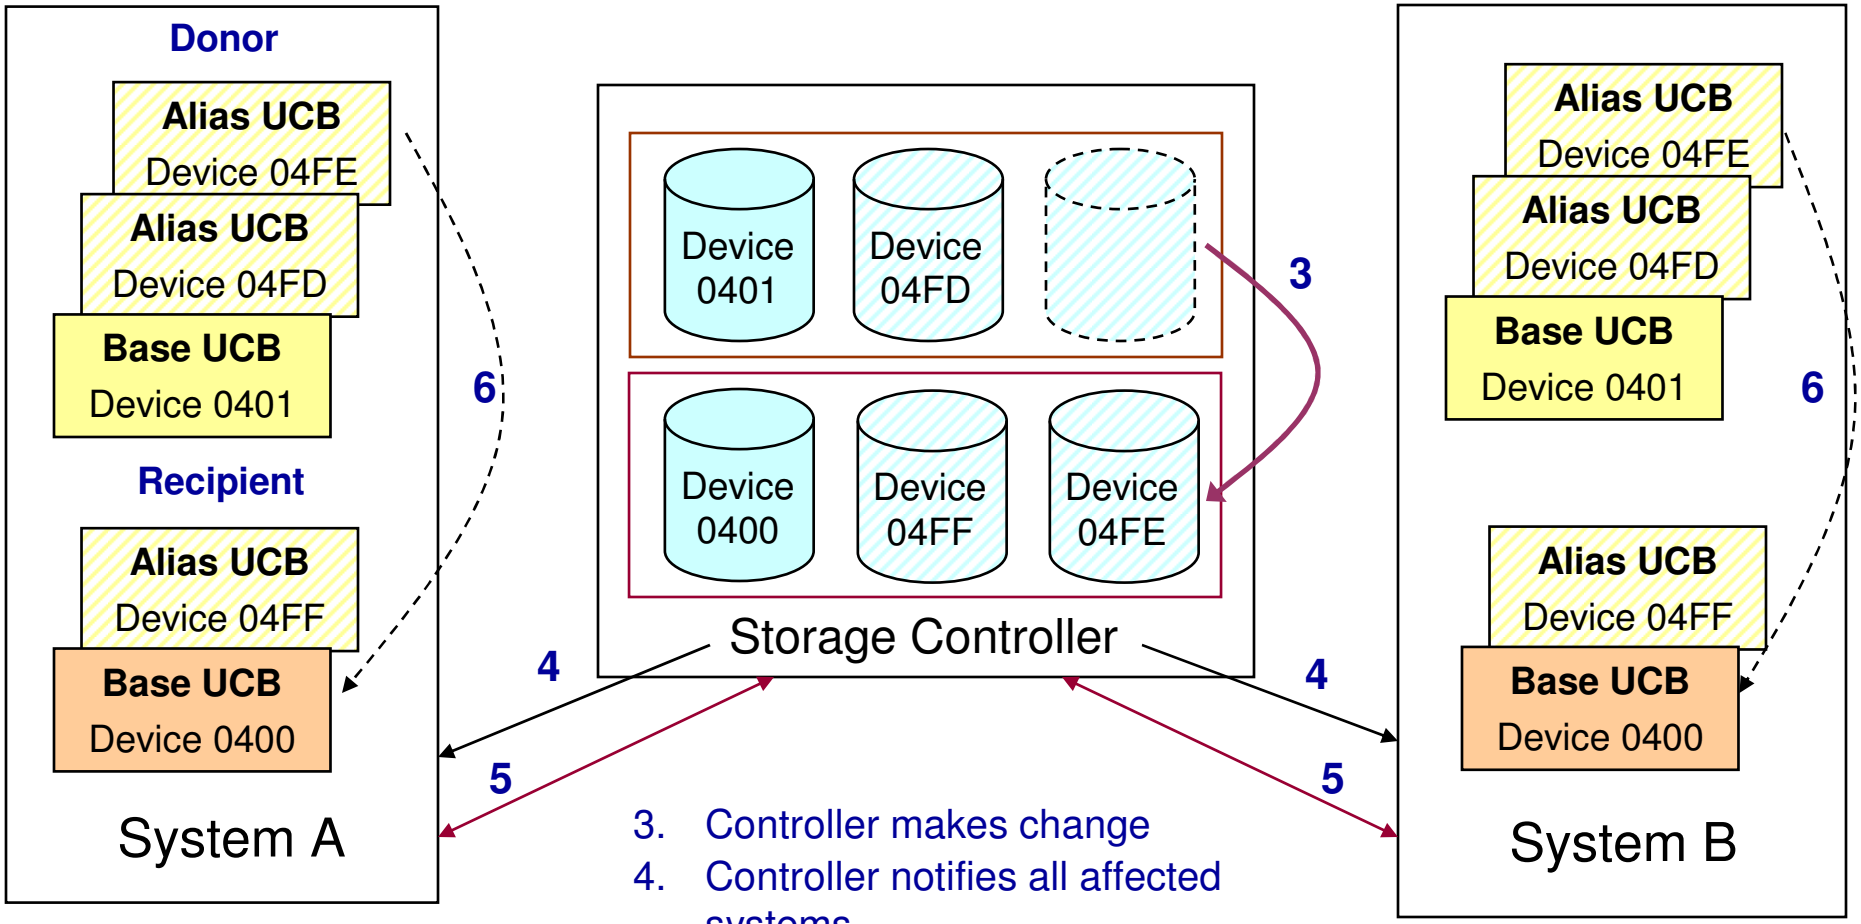

3. Controller makes change
4. Controller notifies all affected systems
5. Systems read new binding info
6. Systems adjust UCBs

HyperPAV

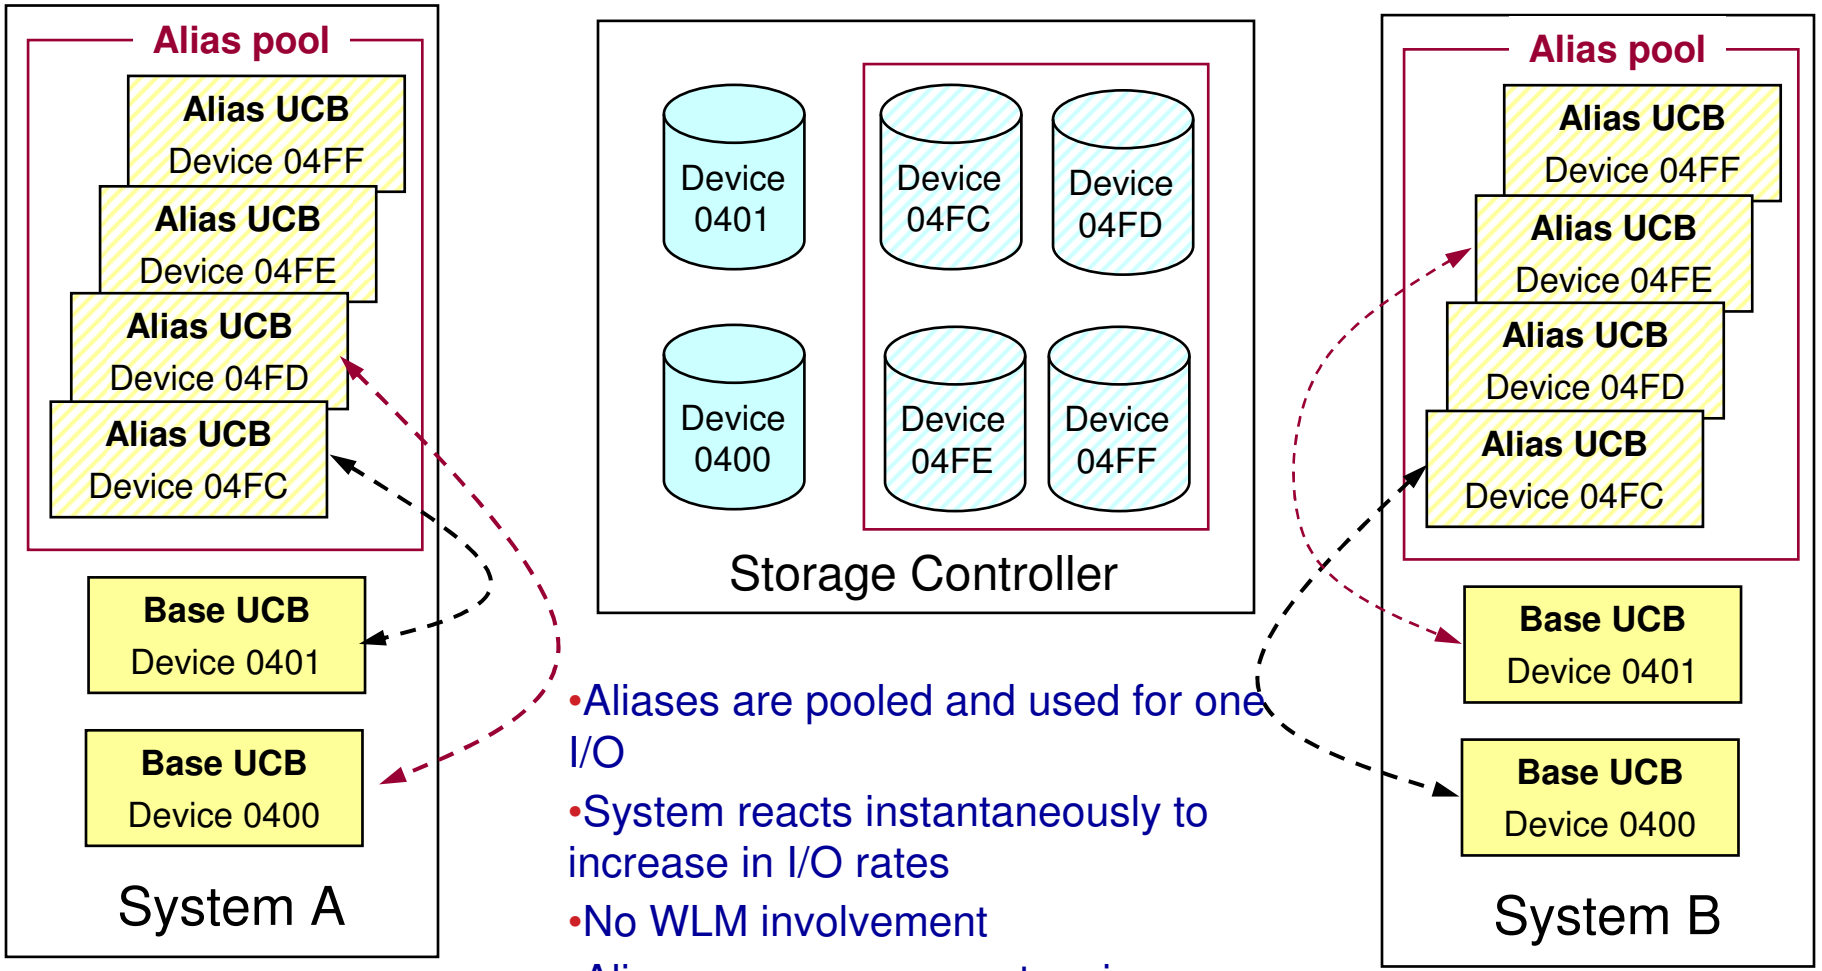

- Aliases are pooled and used for one I/O
- System reacts instantaneously to increase in I/O rates
- No WLM involvement
- Aliases use on one system is independent of others

HyperPAV Tips – How Many Aliases?

- Little's Law:
  - Number of active requests = I/O rate * I/O service time
- Applied to HyperPAV:
  - No. of aliases = Peak avg. I/O rate * service time * multiplier
- Multiplier added since peak average service time doesn't account for workload spikes
  - Multiplier should be at least 2 since peak concurrent I/Os will be higher than the average
  - Monitor for increases in IOSQ time or HPAV wait delay (see next page)

HyperPAV Tips - Monitoring

D I R E C T A C C E S S D E V I C E A C T I V I T Y

z/OS V1R13 SYSTEM ID TRX2 DATE 11/23/2011
CONVERTED TO z/OS V2R1 RMF TIME 12.30.00

STORAGE GROUP	DEV NUM	DEVICE TYPE	NUMBER OF CYL	VOLUME SERIAL	PAV	LCU	ACTIVITY RATE	RESP TIME	IOSQ TIME	CMR DLY	DB DLY	INT DLY	PEND TIME	DISC TIME	CONN TIME
XTEST	2209	33903	3339	TRXSXA	1	0032	24.901	6.54	2.27	.202	.000	.020	.212	2.96	1.10
	220A	33909	10017	TRXT01	1	0032	20.085	6.12	2.06	.207	.000	.019	.217	2.80	1.04
	220B	33909	10017	TRXT02	1	0032	30.702	7.18	2.28	.210	.000	.013	.220	3.73	0.95
	220C	33909	10017	TRXT03	1	0032	27.103	6.93	2.03	.201	.000	.016	.211	3.67	1.02
	220D	33909	10017	TRXT04	1	0032	31.040	6.30	2.25	.208	.000	.022	.218	2.85	0.98
	220E	33909	10017	TRXT05	1	0032	33.010	6.62	2.31	.203	.000	.017	.213	3.08	1.01

First sign that I may not have enough aliases defined for this LCU is high IOSQ time.

Note that there may be other performance problems that are causing high IOSQ queue. For example, high disconnect time due to staging data into cache. Adding aliases won't necessarily help in these cases.

HyperPAV Tips – Monitoring

```
                          I/O   Q U E U I N G   A C T I V I T Y

            z/OS V1R13                  SYSTEM ID TEST1            START ...
                                        RPT VERSION V1R13 RMF      END   ...
TOTAL SAMPLES =    901    IODF = 72     CR-DATE: 11/01/2013    CR-TIME:   ...
                                                                AVG    AVG       AVG
 LCU    CU      DCM GROUP    CHAN      CHPID     % DP    % CU   CUB    CMR  ... CSS      HPAV
              MIN MAX DEF   PATHS     TAKEN     BUSY    BUSY   DLY    DLY       DLY    WAIT  MAX
 00A0   2600                 0C        0.003     0.00    0.00   0.0    0.1
                             0F        0.007     0.00    0.00   0.0    0.0
                             42        0.007     0.00    0.00   0.0    0.2
                             45        0.007     0.00    0.00   0.0    0.6
                             00        0.007     0.00    0.00   0.0    1.2
                             03        0.007     0.00    0.00   0.0    0.5
                             06        0.003     0.00    0.00   0.0    1.3
                             09        0.003     0.00    0.00   0.0    2.2
                              *        0.043     0.00    0.00   0.0    0.6       0.1  0.000  10
```

Average number of I/O requests that had to wait because no aliases available

Maximum number of aliases in use during this interval. Be careful with this number because short spike may cause it to increase!

HyperPAV Tips - Monitoring

```
SY1  d m=dev(0710)
SY1  IEE174I 23.35.49 DISPLAY M 835
DEVICE 0710   STATUS=ONLINE
CHP                    10   20   30   40
DEST LINK ADDRESS      10   20   30   40
PATH ONLINE            Y    Y    Y    Y
CHP PHYSICALLY ONLINE  Y    Y    Y    Y
PATH OPERATIONAL       Y    Y    Y    Y
MANAGED                N    N    N    N
CU NUMBER              0700 0700 0700 0700
MAXIMUM MANAGED CHPID(S) ALLOWED:  0
DESTINATION CU LOGICAL ADDRESS = 07
SCP CU ND          = 002107.000.IBM.TC.03069A000007.00FF
SCP TOKEN NED      = 002107.900.IBM.TC.03069A000007.0700
SCP DEVICE NED     = 002107.900.IBM.TC.03069A000007.0710
HYPERPAV ALIASES IN POOL  32
```

Sometimes an I/O configuration or z/OS issue will cause aliases not to be available for use. You should compare this number with what you have defined for your LCU.

PAV Trivia

Why is HyperPAV supported only for FICON channels?

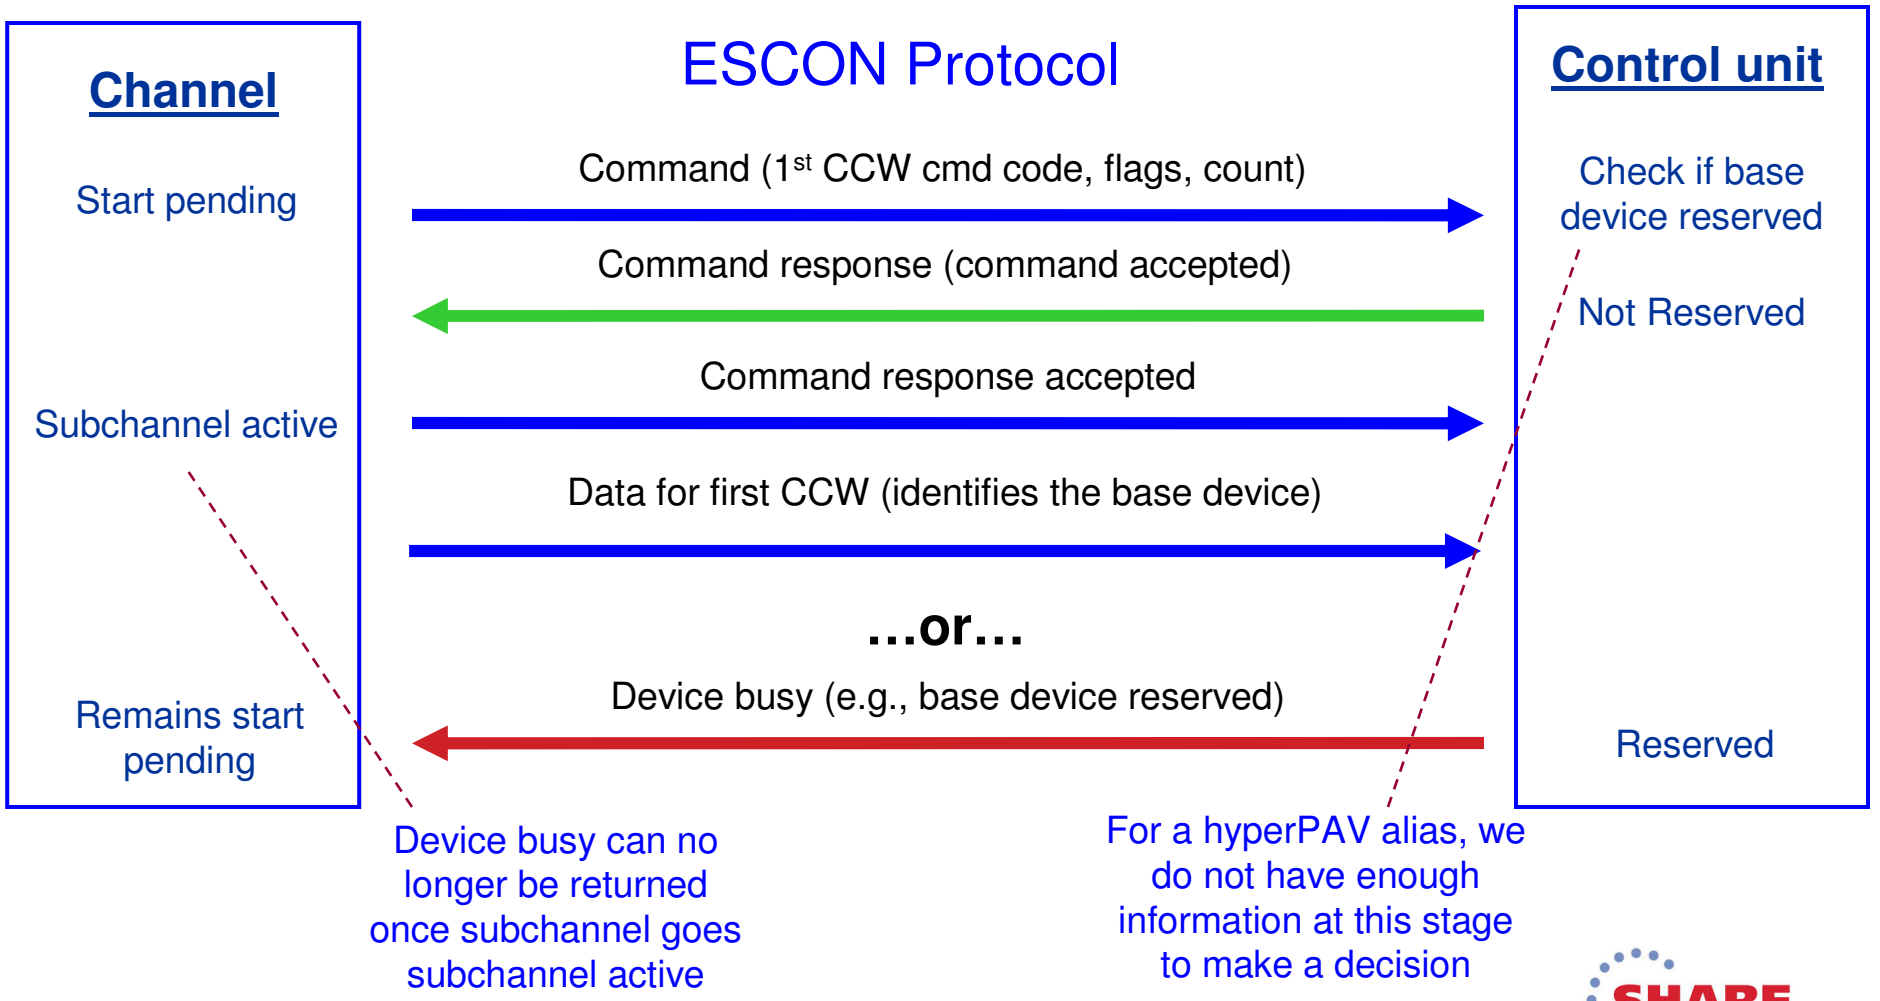

PAV Trivia

Why is HyperPAV supported only for FICON channels?

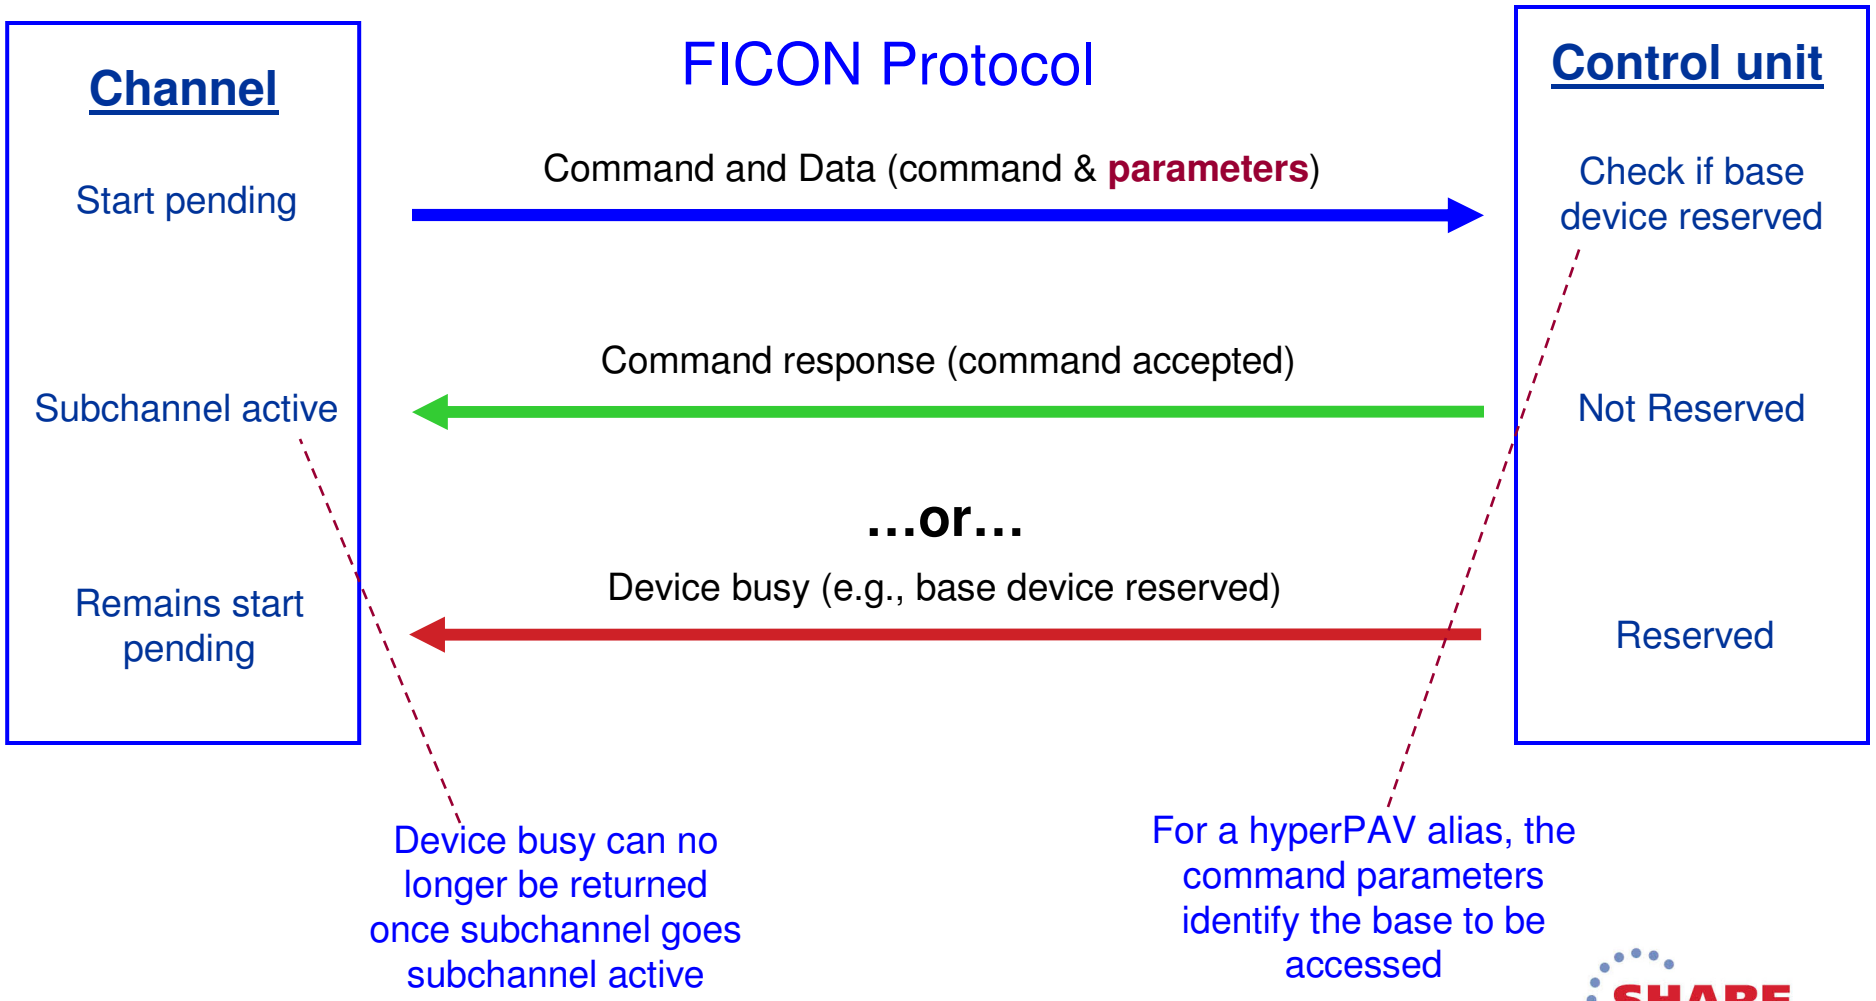

Dynamic Channel-Path Management (DCM)

- Goals
  - Simplify I/O configuration management
  - Improve performance management
  - Enhance RAS
- System manages a pool of channels that are dynamically added to control units as needed
  - Managed channel
  - Managed control unit
- System will add or move channels based on:
  - Channel utilization hot spots
  - Port utilization hot spots
  - Availability need – down to your last channel path

Dynamic Channel-Path Management (DCM)

Data Collection

- Collect performance data (e.g., channel and CU port utilization, connect time, bytes transferred)

Topology Discovery

- Internal representation of connections between:
  - Channels
  - Switches
  - Control Units
- Updated as changes to I/O topology occur

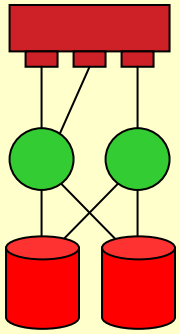

DCM Algorithms

- Check for hot spots and performance issues
- Determine what changes can be made
- Predict new channel utilizations
- Select the best change
- Change the I/O configuration via dynamic I/O
- Vary the path online to all systems

Possible changes

1. **Add A0 to 5000**
2. Add A1 to 5000
3. Add A0 to 4000
4. Add A1 to 4000

DCM Change

Dynamic I/O

- Add, replace or delete CHPID from control unit group

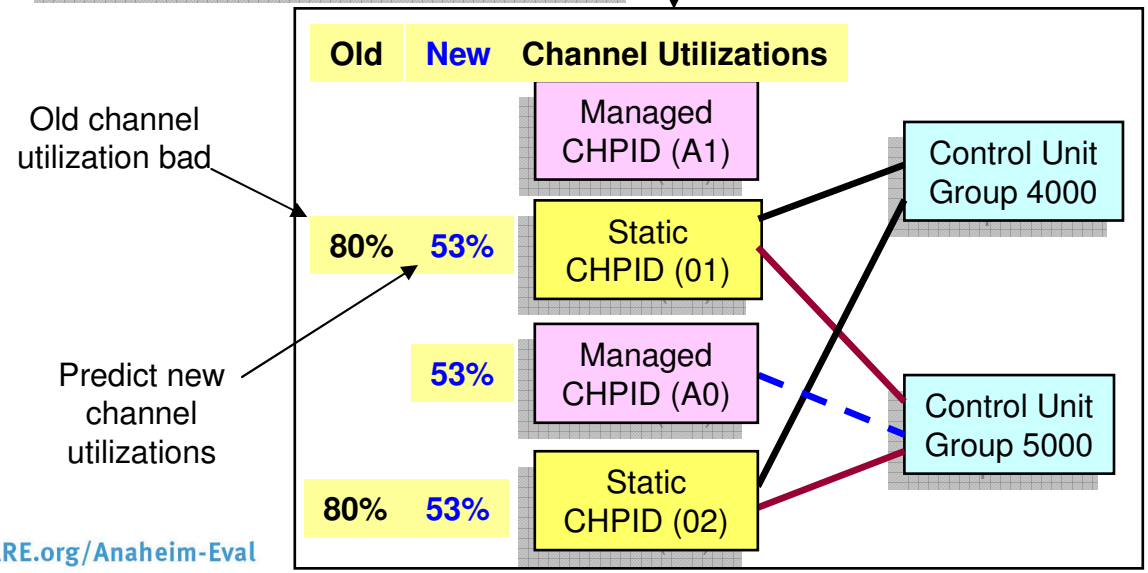

DCM Multi-Switch Cascade

Before z/OS 2.1

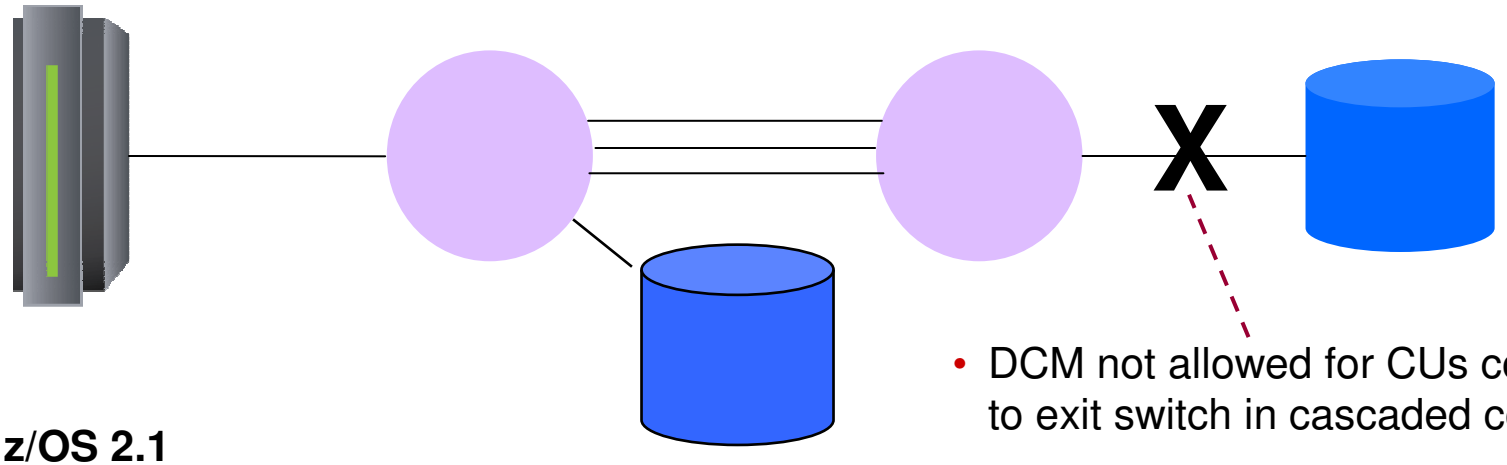

- DCM not allowed for CUs connected to exit switch in cascaded config

With z/OS 2.1

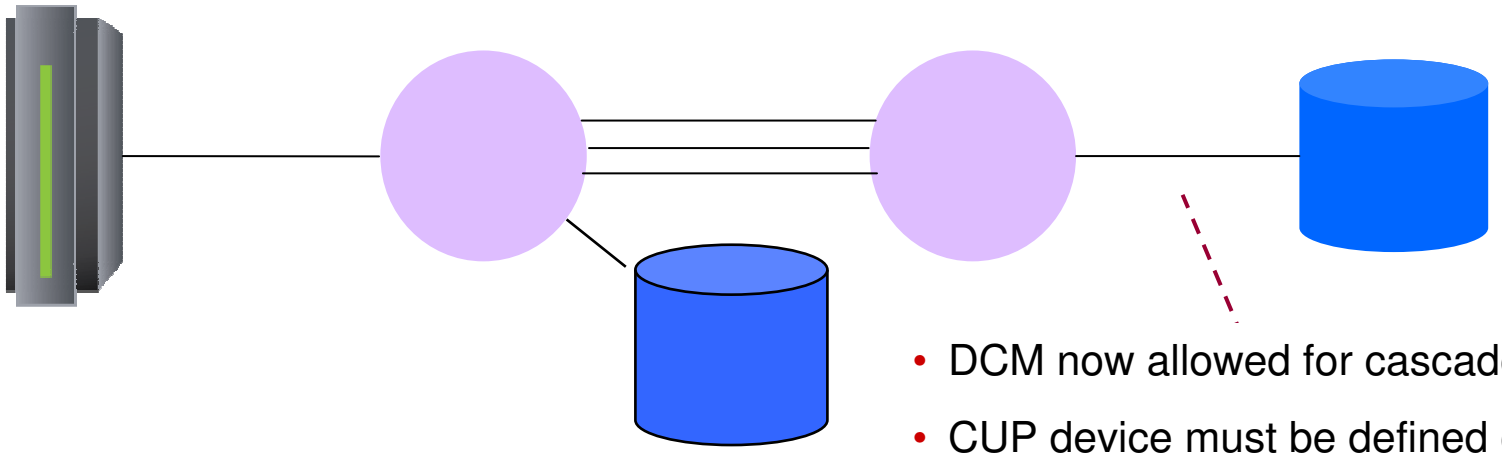

- DCM now allowed for cascaded
- CUP device must be defined on local and remote switches with 2 byte link addresses

zEC12 Enhanced Channel Path Selection

- Overall I/O performance can be affected by poor performance on one channel path
  - See RMF I/O queuing report
- There are many possible causes:
  - Multi-system workload spikes
  - SAN congestion or destination port congestion
  - Control unit host adapter or port congestion
  - Firmware failures in the SAN, channel extenders, control units
  - Hardware failures (link speeds did not initialize correctly)
  - Mis-configuration and cabling errors
- Channel subsystem will now balance paths based on initial command response time

What is Initial Command Response Time?

- Initial command response (CMR) time is the amount of time from when the channel sends the first command until it gets a response from the control unit
  - One round trip through the fabric
  - Originally created to distinguish between IFCCs that occur because of congestion vs lost frames
  - Good for detecting fabric congestion and other problems on a path

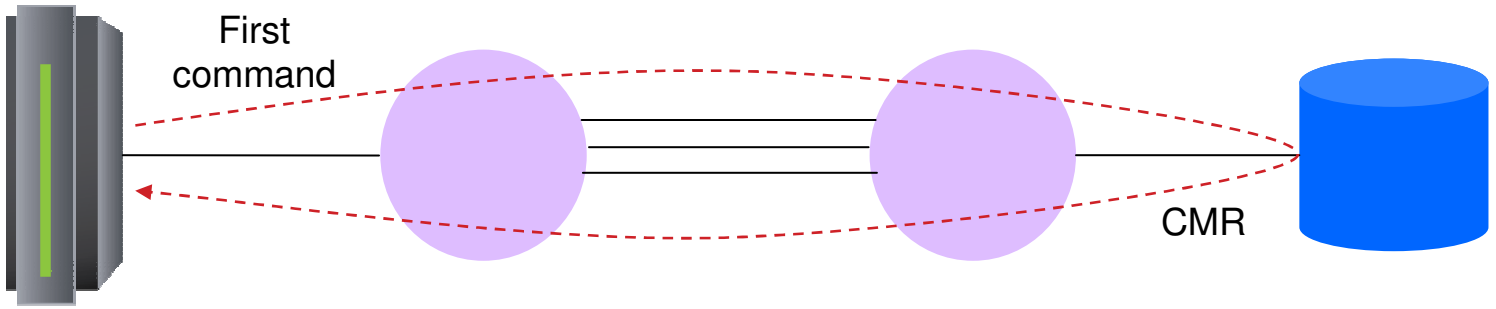

zEC12 Enhanced Channel Path Selection

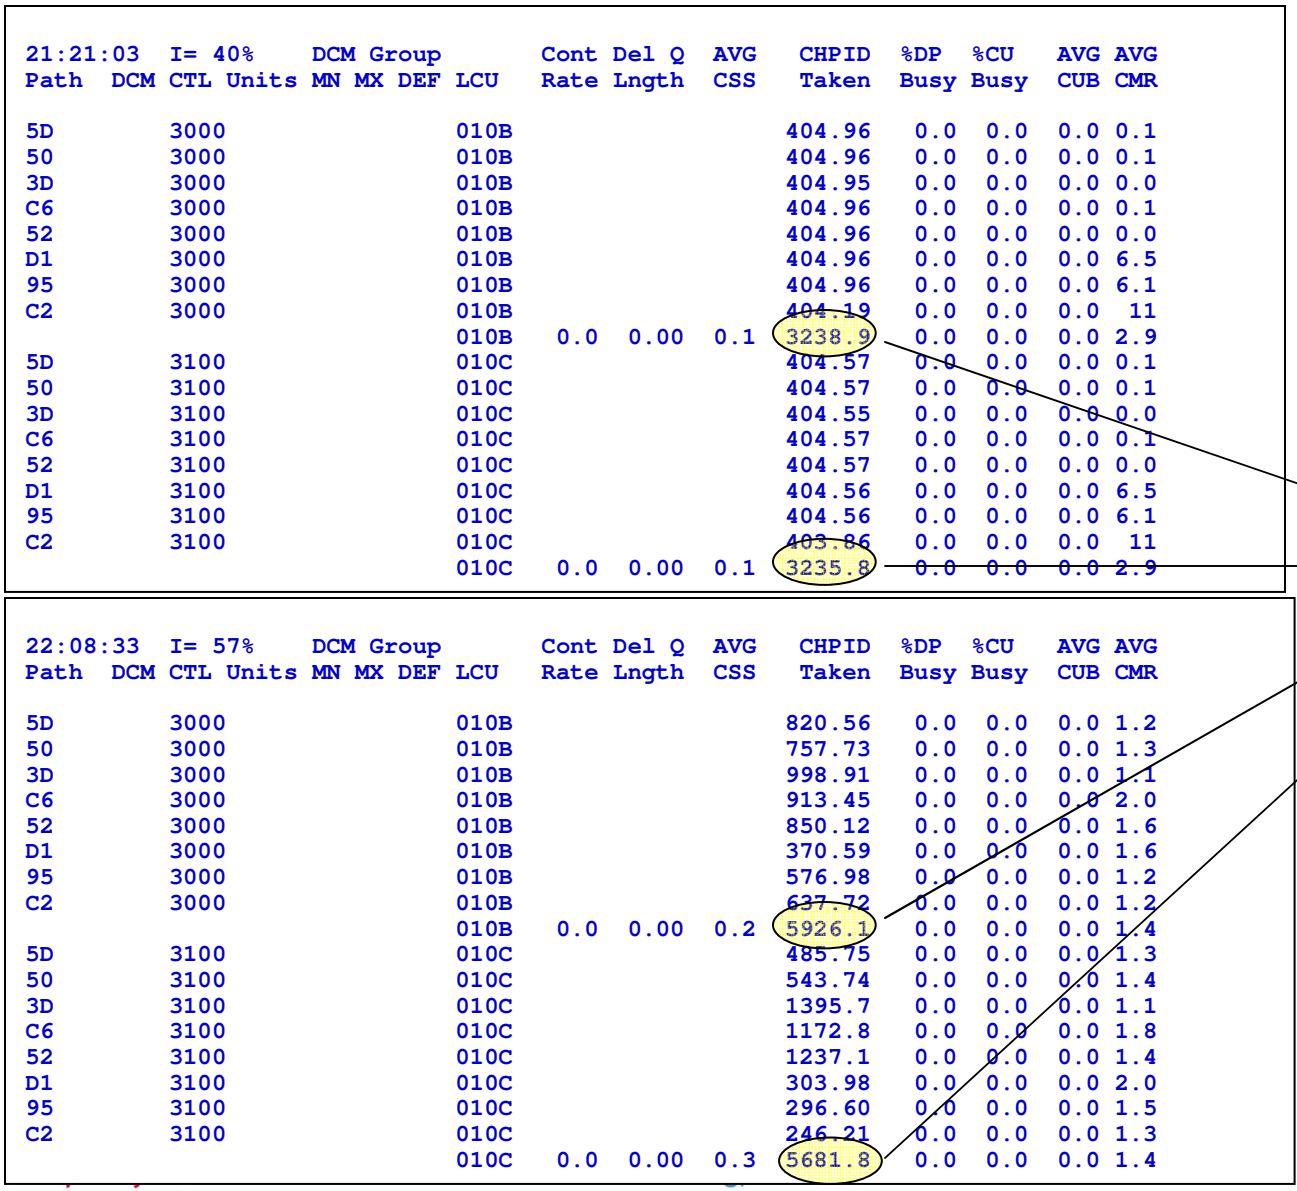

21:21:03 I= 40% DCM Group Cont Del Q AVG CHPID %DP %CU AVG AVG

Path	DCM	CTL Units	MN	MX	DEF	LCU	Rate	Lngth	CSS	Taken	Busy	Busy	CUB	CMR
5D		3000				010B				404.96	0.0	0.0	0.0	0.1
50		3000				010B				404.96	0.0	0.0	0.0	0.1
3D		3000				010B				404.95	0.0	0.0	0.0	0.0
C6		3000				010B				404.96	0.0	0.0	0.0	0.1
52		3000				010B				404.96	0.0	0.0	0.0	0.0
D1		3000				010B				404.96	0.0	0.0	0.0	6.5
95		3000				010B				404.96	0.0	0.0	0.0	6.1
C2		3000				010B				404.19	0.0	0.0	0.0	11
						010B	0.0	0.00	0.1	3238.9	0.0	0.0	0.0	2.9
5D		3100				010C				404.57	0.0	0.0	0.0	0.1
50		3100				010C				404.57	0.0	0.0	0.0	0.1
3D		3100				010C				404.55	0.0	0.0	0.0	0.0
C6		3100				010C				404.57	0.0	0.0	0.0	0.1
52		3100				010C				404.57	0.0	0.0	0.0	0.0
D1		3100				010C				404.56	0.0	0.0	0.0	6.5
95		3100				010C				404.56	0.0	0.0	0.0	6.1
C2		3100				010C				403.86	0.0	0.0	0.0	11
						010C	0.0	0.00	0.1	3235.8	0.0	0.0	0.0	2.9

22:08:33 I= 57% DCM Group Cont Del Q AVG CHPID %DP %CU AVG AVG

Path	DCM	CTL Units	MN	MX	DEF	LCU	Rate	Lngth	CSS	Taken	Busy	Busy	CUB	CMR
5D		3000				010B				820.56	0.0	0.0	0.0	1.2
50		3000				010B				757.73	0.0	0.0	0.0	1.3
3D		3000				010B				998.91	0.0	0.0	0.0	1.1
C6		3000				010B				913.45	0.0	0.0	0.0	2.0
52		3000				010B				850.12	0.0	0.0	0.0	1.6
D1		3000				010B				370.59	0.0	0.0	0.0	1.6
95		3000				010B				576.98	0.0	0.0	0.0	1.2
C2		3000				010B				637.72	0.0	0.0	0.0	1.2
						010B	0.0	0.00	0.2	5926.1	0.0	0.0	0.0	1.4
5D		3100				010C				485.75	0.0	0.0	0.0	1.3
50		3100				010C				543.74	0.0	0.0	0.0	1.4
3D		3100				010C				1395.7	0.0	0.0	0.0	1.1
C6		3100				010C				1172.8	0.0	0.0	0.0	1.8
52		3100				010C				1237.1	0.0	0.0	0.0	1.4
D1		3100				010C				303.98	0.0	0.0	0.0	2.0
95		3100				010C				296.60	0.0	0.0	0.0	1.5
C2		3100				010C				246.21	0.0	0.0	0.0	1.3
						010C	0.0	0.00	0.3	5681.8	0.0	0.0	0.0	1.4

I/O Rate Health Check (OA40508)

```
CHECK(IBMIOS,IOS_IORATE_MONITOR)
START TIME: 04/22/2013 08:44:14.271360
CHECK DATE: 20120430  CHECK SEVERITY: MEDIUM
CHECK PARM: THRESHOLD(100),RATIO(2),XTYPE(),XCU()

IOSHC133I I/O Rate Report

The following control units show inconsistent I/O rates based on these
parameters:
THRESHOLD = 100
RATIO = 2

I/O RATE EXCEPTION DETECTED AT: 04/22/2013 08:44:14.254730
CONTROL UNIT = 0500
ND = 002107.000.IBM.PK.000000000002

          ENTRY EXIT   CU     I/O      AVG      IOR
  CHPID   LINK  LINK   INTF   RATE     CMR      EXC
   14     B153  B177   0001   39.603   2.560     *
   44     B353  B375   0012   101.38   2.112
   16     B055  B277   0013   98.019   2.134
   46     B154  B376   0104   50.693   2.048

*Medium Severity Exception *

IOSHC132E Analysis of I/O rates detected one or more control units with
an exception.
```

I/Os per second to the CU must be at least this high

I/O rate ratio between the highest and lowest paths

This is the exception path

Exception message appears in system log

Additional Reading

- *Evolution of the System z Channel* – Presentation given at Share August 2011 by Patty Driever and Howard Johnson
- *IBM zEnterprise 196 and IBM zEnterprise 114 I/O and FICON Express8S Channel Performance*
- *FICON DCM for System Programmers* - Whitepaper on Techdocs
- *IBM zEnterprise Storage I/O Advancements* - IBM Journal of Research and Development, Vol 56

Thank you

Appendix

FICON Trivia

How long is a fiber channel frame?

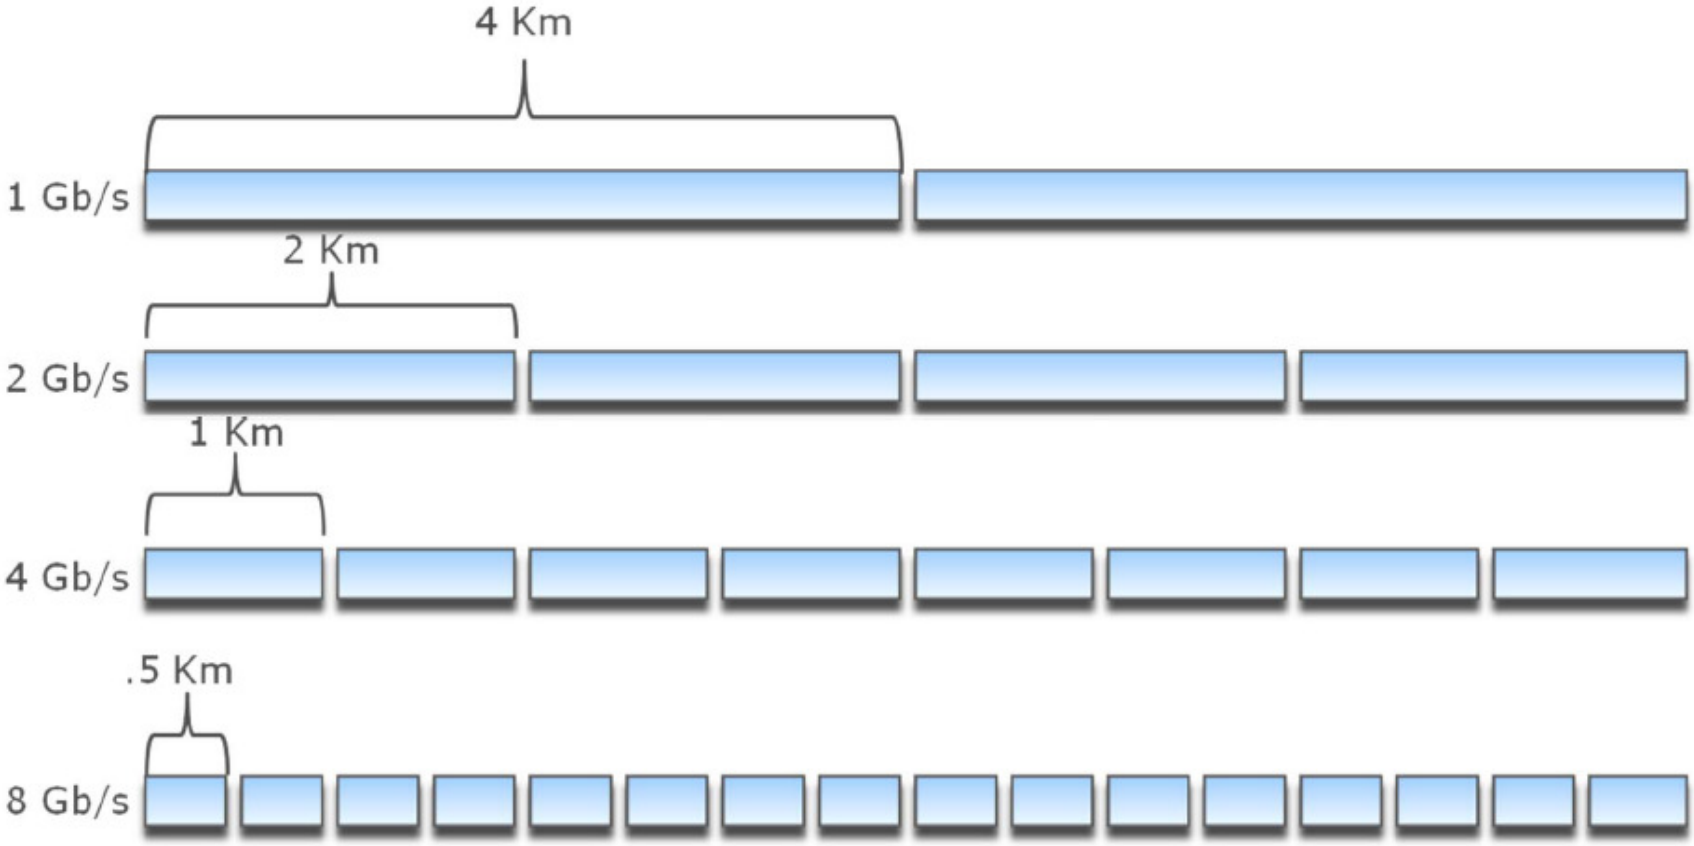

★ The fibre channel frame is the basic building block in FICON. A frame can have a payload up to 2112 bytes. Each byte is represented by 10 bits (8b/10b encoding).

Slide by Howard Johnson

zHPF - 4X More I/Os per Second than FICON I/O

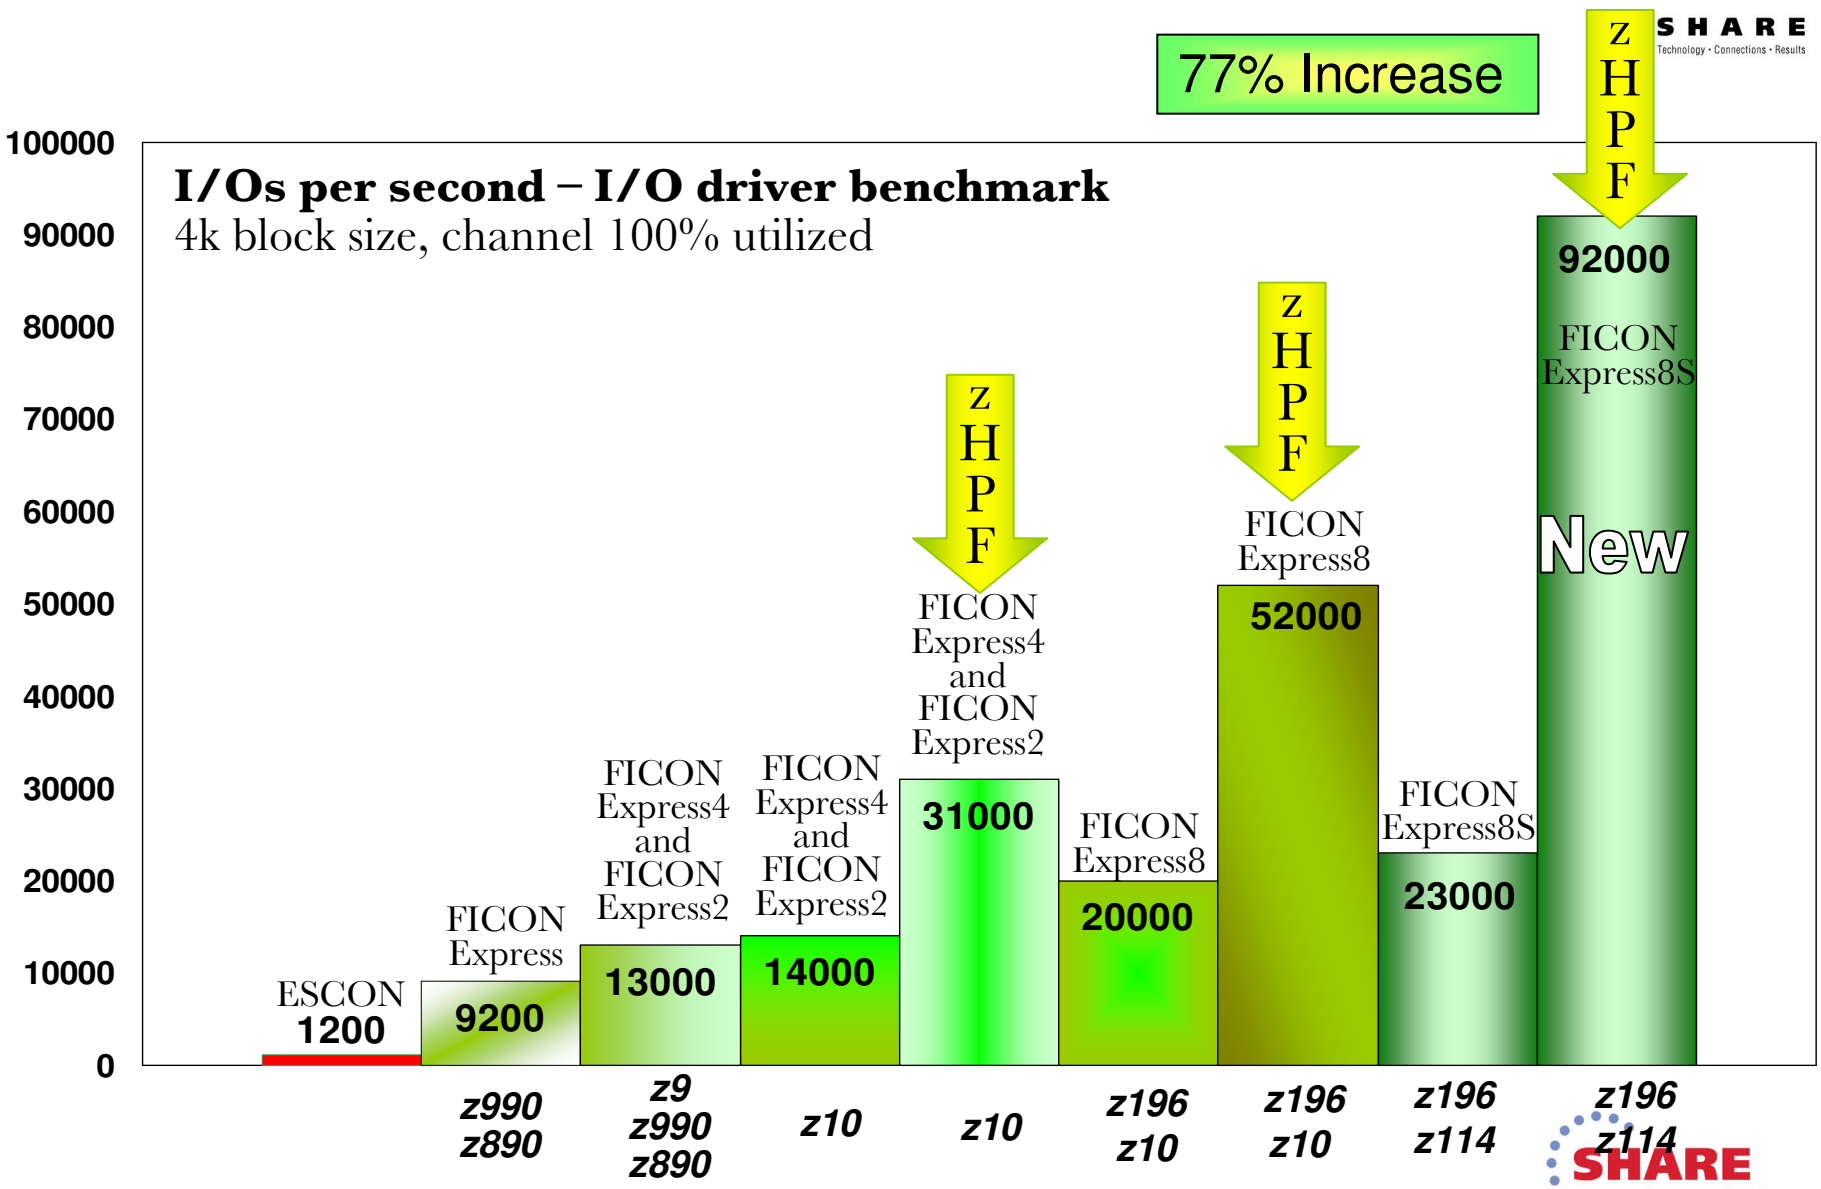

zHPF - More than 2X FICON Throughput

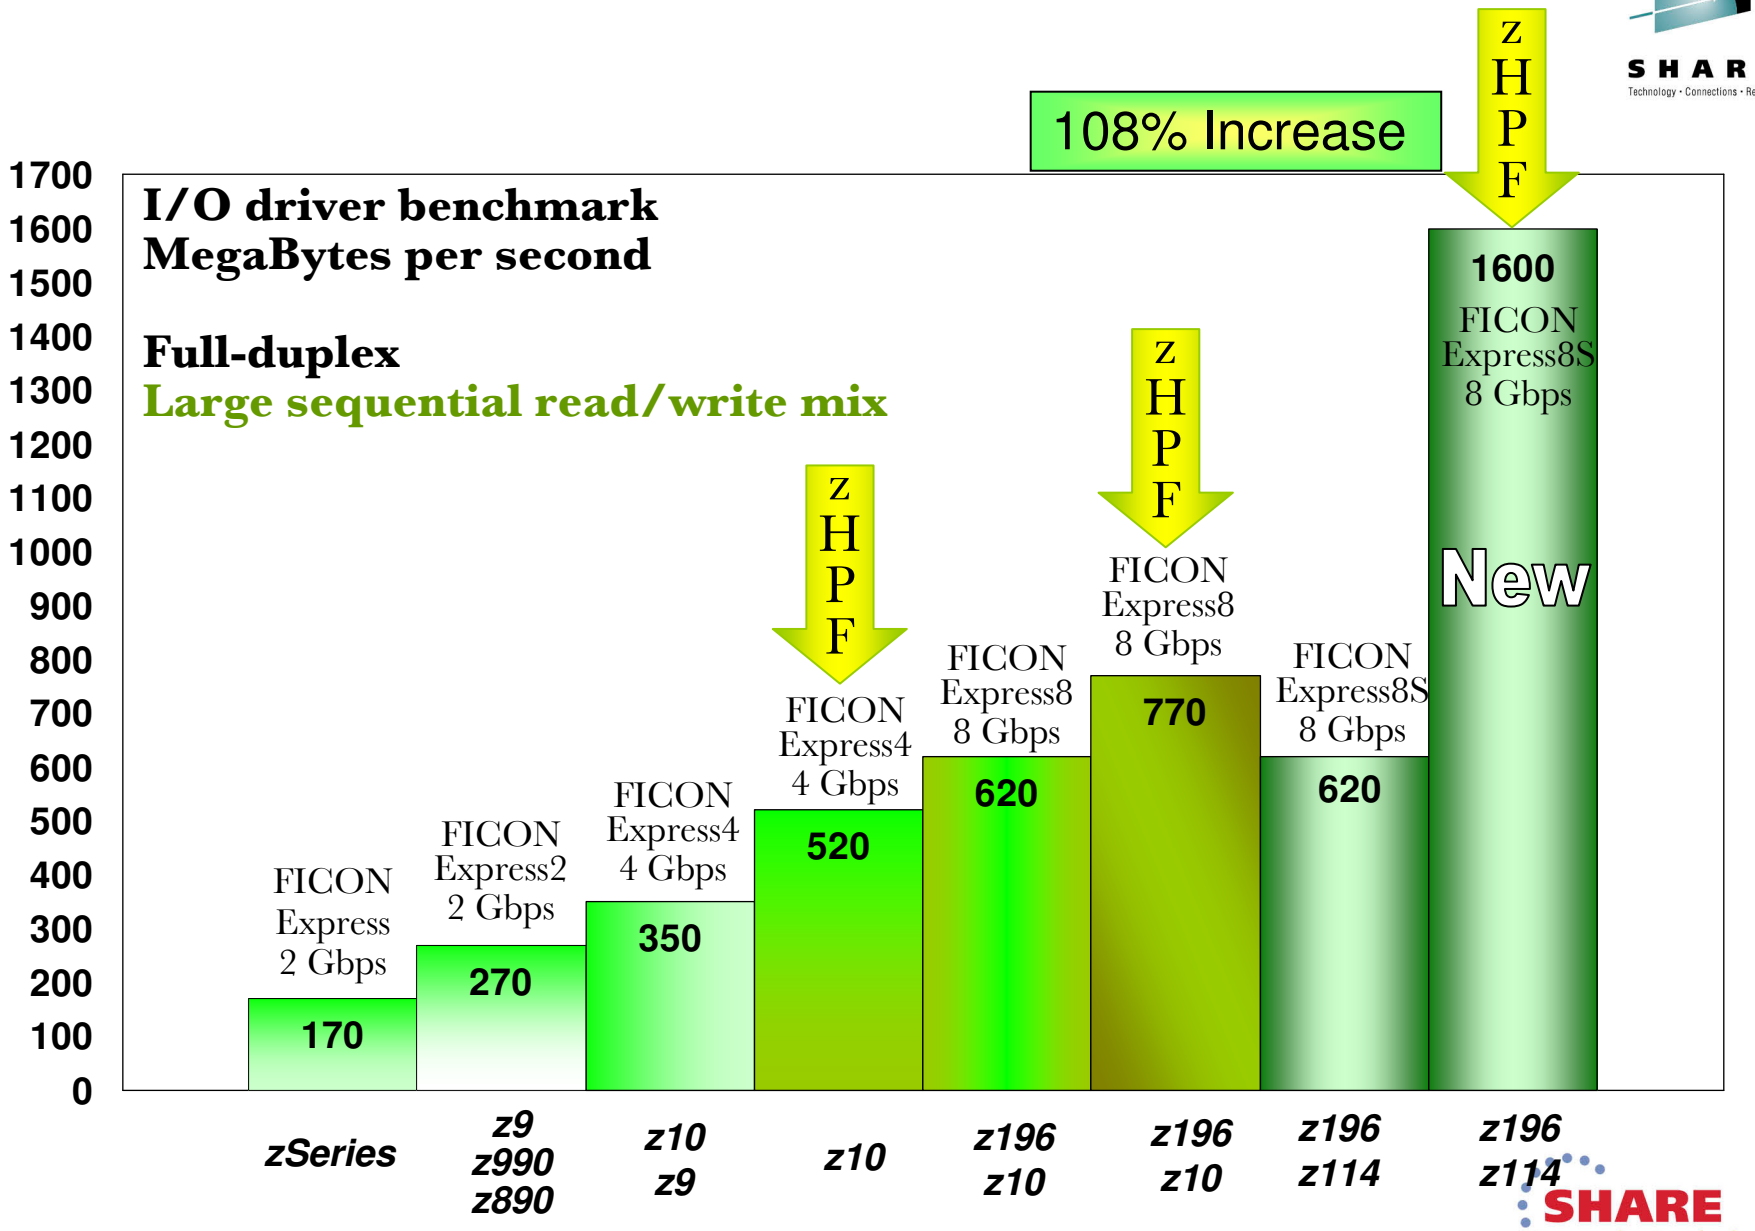

zHPF - Response Time Improves Too

Single FICON Express8S channel: zHPF vs FICON READ 4k bytes/IO
Total I/O Response Time vs IO/sec

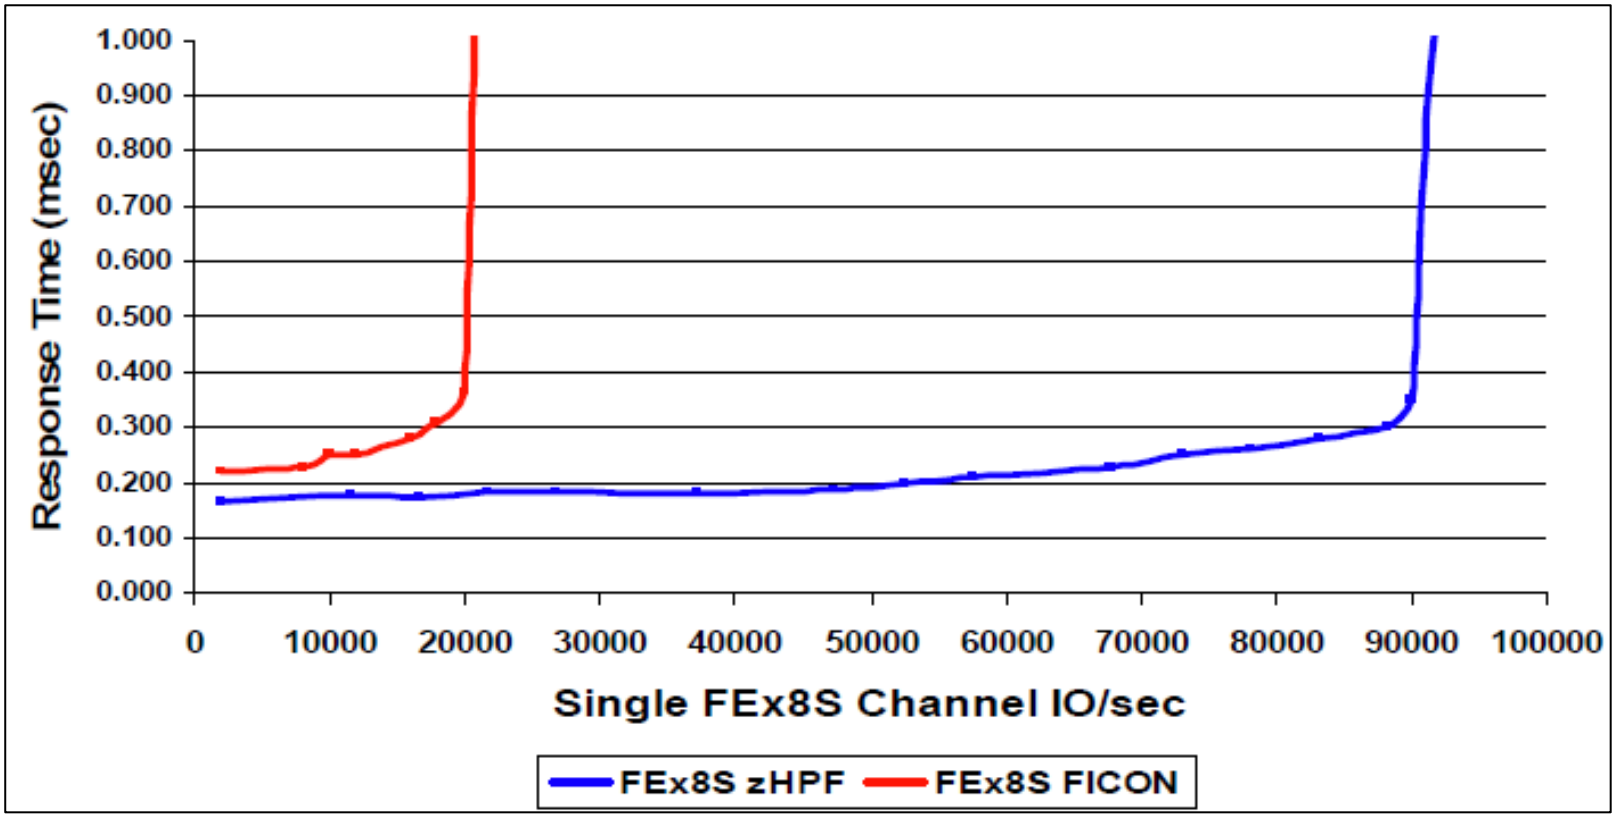

CMR Health Check Report Example

```
CHECK(IBMIOS,IOS_CMRTIME_MONITOR)
START TIME: 12/10/2011 16:34:03.455536
CHECK DATE: 20100501  CHECK SEVERITY: MEDIUM
CHECK PARM: THRESHOLD(3),RATIO(5),XTYPE(),XCU()

IOSHC113I Command Response Time Report

The following control units show inconsistent average command response
(CMR) time based on these parameters:
THRESHOLD = 3
RATIO = 5

CMR TIME EXCEPTION DETECTED AT: 12/10/2011 16:29:24.212239
CONTROL UNIT = 25C0
ND = 002107.941.IBM.75.0000000WH391

          ENTRY EXIT  CU    I/O     AVG
  CHPID   LINK  LINK  INTF  RATE    CMR
   81     2C51  2DC4  0030  72.330  9.21
   22     3C1B  3DC2  0031  71.651  9.47
   82     2C52  2DC0  0032  72.333  8.70
   84     2C54  2DCC  0100  71.810  1.92
   21     3C19  3DD2  0231  72.122  1.79
* Medium Severity Exception *

IOSHC112E Analysis of command response (CMR) time detected one or
more control units with an exception.
```

These are the exception paths

Exception message appears in system log