Improving Performance For FICON Infrastructures

Dr. Steve Guendert
Mainframe Global Solutions Architect
Brocade

Session 5614
Monday Feb. 25, 2008

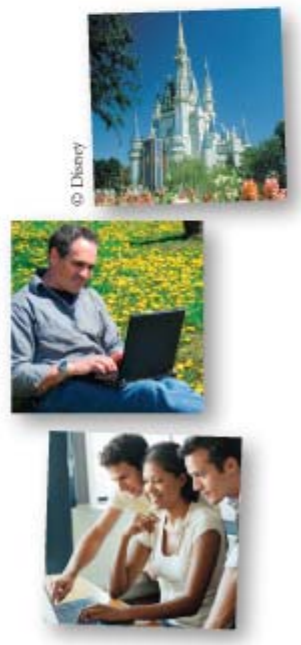

Abstract

- A performance optimized FICON environment will result in lower TCO. This session focuses on FICON performance and how you can diagnose problems and proactively take action to prevent performance issues.
- This session will emphasize maximizing the performance of your FICON environment today, so as to minimize your TCO, in preparation for the future enhancements to the FICON protocol. The presentation will be based on several recent zJournal articles written by the speaker. We'll discuss BB credit optimization, QoS, Modified Indirect Data Address Word (MIDAW), what is 4 Gbps really doing for you?, frame pacing delay, Open Exchange management and using CUP and the RMF 74-7 record (FICON director Activity report), as well as other SW for optimizing performance.
- **Session will not cover the basics/introductory material on discussed subtopics. It is suggested that you consult the references listed for further information.**

Key References

- Guendert, S. Buffer-to-Buffer Credits and Their Effect on FICON Performance, CMG Measure IT, March 2005.
- Artis, H.P. and Guendert, S. Designing and Managing FICON Interswitch Link infrastructures, Proceedings of the Computer Measurement Group, 2006.
- Guendert, S. FICON and Quality of Service (QoS): Making the Case for True, End-to-End, Host Managed QoS. zJournal February/March 2008.
- Guendert, S. and Lytle, D. Buffer to Buffer Credit Management: An Oxymoron? zJournal, June/July 2007
- Guendert, S. Understanding the Performance Implications of Buffer to Buffer Credit Starvation: Frame Pacing Delay, Proceedings of the Computer Measurement Group, 2007
- Allen, A.O., Probability, Statistics, and Queueing Theory, Academic Press, 1978.
- Guendert, S. Taking FICON To the Next Level: Cascaded High Performance FICON, Proceedings of the Computer Measurement Group, 2005.
- Cronin, C. Performance Considerations for Cascaded FICON Directors, www-1.ibm.com/servers/eserver/zseries/library/techpapers/gm130237.html, March 2003
- Artis, H.P. Managing Complex FICON Configurations. Performance Associates, Inc. 2005

Agenda

- Buffer-to-Buffer Credit Optimization
- Quality of Service (QoS)
- MIDAW
- What is 4 Gbit/sec FICON doing for me?
- Open Exchanges
- Software

SHARE
Technology · Connections · Results

Understanding the Performance Implications of Buffer to Buffer Credit Starvation In a FICON Environment: Frame Pacing Delay

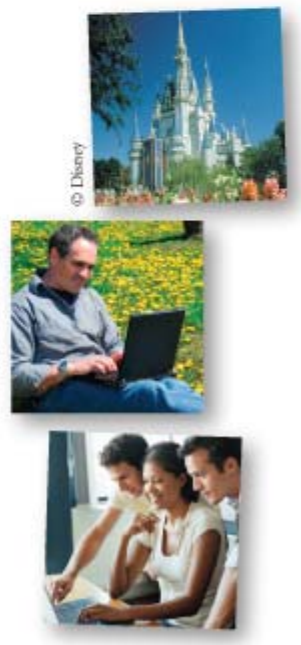

Wikipedia defines Oxymoron

- An oxymoron is a figure of speech that combines two normally contradictory terms. *Oxymoron* is from Greek *oxy* ("sharp") and *moros* ("dull"). Thus the word *oxymoron* is itself an oxymoron.
- Oxymorons are a proper subset of the expressions called contradictions in terms. What distinguishes oxymorons from other paradoxes and contradictions is that they are used intentionally, for rhetorical effect, and the contradiction is only apparent, as the combination of terms provides a novel expression of some concept, such as "cruel to be kind".

Buffer to Buffer Credit management: an oxymoron

- There is no way to actually report/track how many BB credits are being used.
- The RMF 74-7 record comes close, but names the field something else.
- Published rules of thumb mistakenly assume full frames
- Similar to dynamic PAVs and alias overkill, end users tend to overkill BB credit assignment
  - Can lead to director configuration issues which may cause an outage to fix
- Problems occur with BB credits for some underlying reason

ESCON Device Information Blocks (DIBs) -a brief historical perspective

- ESCON DIBs contain control, status, or user data.
- Storage director's (control unit) adapters define maximum DIB size employed by configuration.
  - 8,16,32,64,128,256,512, or 1024 bytes long
- A microprocessor on the ESCON adapter manages the adapter's data buffer.
  - 3990-3/6 ESCON adapter held (8) 256 byte DIBs.

ESCON Device Information Blocks (DIBs)- a brief historical perspective (2)

- During initialization, a subsystem's ESCON adapters notify the channel subsystem of DIB and data buffer sizes to be employed.
- The channel subsystem employs these values for all future communications with the subsystem.

- Assuming ESCON adapter is receiving data, it must successfully pass each DIB it receives to buffer areas in the storage director before it can request another DIB be sent by the channel to avoid potential of data overrun.
- While the channel may initially schedule a full data buffer of DIBs be sent to a device, further DIBs may only be transmitted after a request has been received from the ESCON adapter signaling that a prior DIB has been successfully passed along to the storage director.

Buffer-to-Buffer Flow Control

- Flow control between two optically adjacent ports in the I/O path
- Separate, independent pool of credits manages Buffer-to-Buffer flow control (BB Credits)
- Transmitter count non-zero: transmitter is free to continue sending data
- It takes light 5 nsec to propagate through 1 meter of optical fiber:
  - 50 µsec to travel 10 km
- Faster links, longer distances leads to a performance drag similar to ESCON droop:

Buffer Credit Concepts

- Define the maximum amount of data that can be sent prior to an acknowledgement
- Buffer credits are physical ASIC port or card memory resources and are finite in number as a function of cost
- Within a fabric, each port may have a different number of buffer credits
- The number of available buffer credits is communicated at fabric logon (FLOGI)

Buffer Credit Concepts (2)

- One buffer credit allows a device to send one 2112-byte frame of data (2K usable for z/OS data)
- Assuming that each credit is completely full, you need one credit for every 1 km of link length over a 2 Gbit/sec fiber
- **Unfortunately, z/OS disk workload rarely produce full credits. For a 4K transfer, the average frame size is 819 bytes.**
- **Hence, five credits would be required per km over a 2 Gbit/sec fiber**

Calculating the number of buffer credits

- What you must know to do this correctly is:
  - Link speed (1, 2, 4, 10Gbps) – easy to get
  - Actual fiber run distance that the frame must traverse – easy to get
  - The size of the frame – very hard to get
- Formula for assigning buffer credits (assumes 2148 frame size)
  - 1 Gbps
    - Distance (in km) / 2 + 20%
  - 2 Gbps
    - Distance (in km) + 20%
  - 4 Gbps
    - Distance (in km) x 2 + 20%
  - 10 Gbps
    - Distance (in km) x 6 + 20%
- But does this always work?

Buffer Credits Required

By Size of Frame and Link Speed

SHARE
Technology · Connections · Results

A distance of 20km with the link 100% utilized				2Gbps	4Gbps	8Gbps	10Gbps
SOF, Header, CRC, EOF	Payload	Total Frame Bytes	Smaller than full frame by x%	Buffer Credits Required 8b10b	Buffer Credits Required 8b10b	Buffer Credits Required 8b10b	Buffer Credits Required 64b66b
36	2112	2148	0.000%	20	40	80	117
36	2002	2038	5.138%	21	42	84	124
36	1902	1938	9.809%	22	44	88	130
36	1802	1838	14.481%	24	47	93	137
36	1702	1738	19.152%	25	49	98	145
36	1602	1638	23.823%	26	52	104	154
36	1502	1538	28.494%	28	56	111	164
36	1402	1438	33.165%	30	60	119	175
36	1302	1338	37.836%	32	64	128	188
36	1202	1238	42.507%	35	69	138	203
36	1102	1138	47.179%	38	75	150	221
36	1002	1038	51.850%	41	82	164	243
36	902	938	56.521%	46	91	182	268
36	819	855	60.398%	50	100	199	294
36	700	736	65.957%	58	116	232	342
36	600	636	70.628%	67	134	268	396
36	500	536	75.299%	80	159	318	469
36	400	436	79.970%	98	195	390	577
36	300	336	84.641%	127	254	507	748
36	200	236	89.312%	181	361	721	1065
36	100	136	93.984%	313	626	1251	1848
36	75	111	95.151%	383	766	1532	2264
36	50	86	96.319%	495	989	1978	2922

What Is the Optimal Number of BB Credits?

- Optimal number of credits is determined by:
  - Distance (frame delivery time)
  - Processing time at receiving port
  - Link signaling rate
  - **Size of frames being transmitted**
- Optimal # BB_Credit =

$$\frac{\text{(Round-trip receiving time+Receiving\_port processing time)}}{\text{Frame Transmission time}}$$

- As the link speed increases, the frame transmission time is reduced; therefore, as we get faster iterations of FICON such as FICON Express4 and Express8, the amount of credits need to be increased to obtain full link utilization, even in a short-distance environment!

Formula: Robert Kembel, *The Fibre Channel Consultant*

Why an Optimal Number?

- Analogous to DASD and cache sizing
- Law of diminishing marginal returns
- Exceeding the optimal number of BB_credits does nothing to increase performance, it merely increases your costs
- Optimal number of BB Credits allows for performance optimized distance solutions
- Optimal number allows for better capacity planning
- How does your director allocate BB Credits?

More bytes/frame→more efficient usage of buffer to buffer credits

SHARE
Technology · Connections · Results

Average frame size for FICON Express4 channel MIDAWs measurements

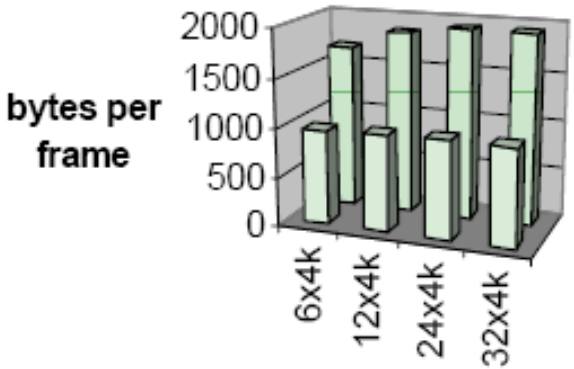

	6x4k	12x4k	24x4k	32x4k
no midaws	965	982	1000	997
with midaws	1692	1873	1943	1943

- Reference: Cathy Cronin “IBM System z9 and FICON Express 4 Performance Update. SHARE Tampa Proceedings, Feb 2006

How Do MIDAWs Effect ISLs?

- IDAW – Each Block Is A CCW
- MIDAW – Move Entire Chain As One I/O
- Example – 4K Block Extended Format

I/O	Data	BB Credits
Write	4096	2
EF Data	35	1

- Chain Of 16

  Total Data = (4096+35) * 16 = 66096
  IDAW, Total BB Credits = 3 * 16 = 48
  MIDAW, Total BB Credits = 66096/2048 = 33

MIDAW Uses 30% Fewer BB Credits So You Can Go 30% Farther

Data Droop for Over Distance at 2 Gbit/sec

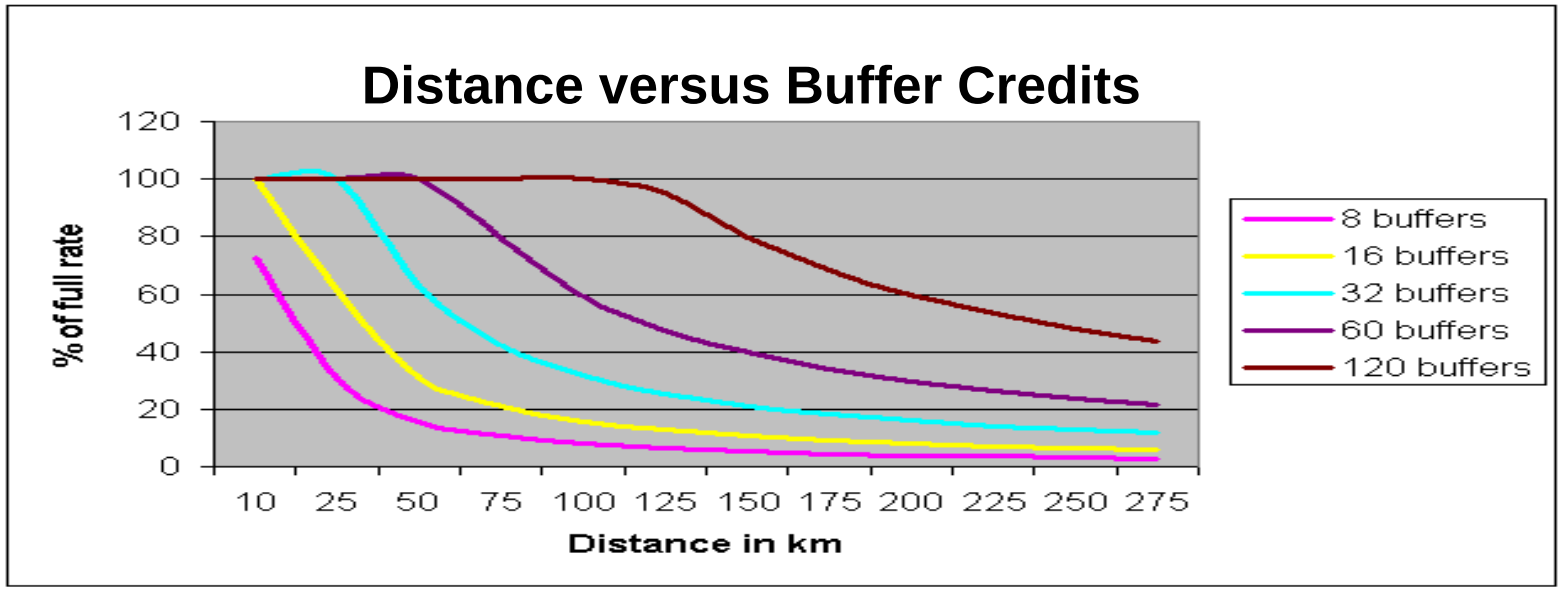

- Cronin, *Performance Considerations for Cascaded FICON Directors*, IBM, March 2003
- Hence, serious consideration must be give to the assignment of credits to ports on director architectures that share a pool of credits among the ports on a card. While relatively few credits (16) might be assigned to local devices, the bulk of the credits should be assigned to ISLs.
- *For data chaining OLTP workloads, assume a worst-case 512 byte average credit size to avoid any potential of droop*
- *MIDAWs and RTD/WTD substantially increase the average credit size*

What to Do?

- Mechanism for counting BB credit usage currently does not exist:
  - Data is there, how do we exploit it?
  - You find out you run out after the fact: let's be more proactive rather than being satisfied with being blind
- What we wind up with is BB credit assignment overkill, similar to how we overkill PAV aliases:
  - Just as you can run out of addresses, you can run out of BB credits
    (or will assign them poorly):
    - Do some analysis and plan things better! Save TCO $$ in the process!
- Director architectures have made BB Credits a capacity planning item unto themselves.

Frame Pacing Delay And Other RMF 74-7 Tidbits

The FICON Director Activity Report Is Your Friend

RMF 74 Subtype 7 Records

Global Data

Switch Data

Connector Data

Port Data

- Four data classes of data are reported by the 74 subtype 7
- Port data includes average read/write frame sizes, average bandwidth, error count, and pacing delays for each port. *Frame pacing occurs when a director port exhausts its available credits.* Frame pacing delays are measured in 2.5 micro-second units
- Data is collected for each RMF interval if FCD is specified in your *ERBRMFnn* parmlib member

Frame Pacing Delay

- AVG FRAME PACING:
  - Defined by RMF as the average time (microseconds) that a frame has to wait before it could be transmitted due to no buffer credits being available
- You always want to see a zero value in this field:
  - Reporting on this value was one of the primary reason that the RMF 74-7 record was developed – it was not needed for ESCON
  - A non-zero value in the AVG FRAME PACING field indicates that you have an issue with insufficient BB Credits
  - It is critical to use CUP in any FICON environment in which distance extension is being utilized
  - 4 Gbit/sec may create more Frame Pacing Delay issues than 2 Gbit/sec
- z/OS disk workloads rarely use a "full" 2148-byte credit:
  - For example, with a 4K block transfer, the average frame size for each 4K transfer is typically about 819 bytes

FICON Director Activity Report from RMF (Requires CUP Code)

F I C O N D I R E C T O R A C T I V I T Y

z/OS V1R4 SYSTEM ID KS01 START 10/03/2005-13.55.00 INTERVAL 000.05.00
RPT VERSION V1R2 RMF END 10/03/2005-14.00.00 CYCLE 1.000 SECONDS

IODF = 0C CR-DATE: 07/07/2005 CR-TIME: 18.47.41 ACT: POR
SWITCH DEVICE: 0001 SWITCH ID: ** TYPE: 006140 MODEL: 001 MAN: MCD PLANT: 01 SERIAL: 0000013A2

PORT ADDR	-CONNECTION- UNIT	ID	AVG FRAME PACING	AVG FRAME SIZE READ	WRITE	PORT BANDWIDTH (MB/SEC) -- READ --	-- WRITE --	ERROR COUNT
04	CHP-H	47	0	433	593	1.01	1.52	0
05	CU	----	0	830	1195	2.30	5.81	0
06	CU	----	0	140	69	0.00	0.00	0
07	CU	C052	0	591	85	0.00	0.00	0
	CU	C050						
08	CHP-H	45	0	400	577	0.98	1.55	0
09	CU	----	0	1355	374	8.42	1.00	0
0A	------	----	0	0	0	0.00	0.00	0
0B	------	----	0	0	0	0.00	0.00	0
0C	CHP-H	4B	0	431	601	1.02	1.56	0
0D	CU	----	0	776	374	3.57	1.41	0
0E	CU	----	0	0	0	0.00	0.00	0
0F	CU	C053	0	1773	78	0.08	0.00	0
	CU	C051						
10	CHP-H	46	0	366	716	1.06	2.32	0
11	CU	----	0	1099	393	4.62	0.97	0
12	------	----	0	0	0	0.00	0.00	0
13	------	----	0	0	0	0.00	0.00	0
14	CHP-H	50	0	533	832	0.32	0.58	0
15	CU	----	0	868	1223	2.25	5.55	0
16	CU	----	0	158	72	0.00	0.00	0
17	CU	C053	0	1761	77	0.09	0.00	0
	CU	C051						
18	CHP-H	49	0	378	745	1.04	2.33	0
19	CU	----	0	1118	399	4.83	0.99	0
1A	------	----	0	0	0	0.00	0.00	0
1B	------	----	0	0	0	0.00	0.00	0
1C	CHP-H	51	0	737	535	0.34	0.17	0
1D	CU	----	0	877	1230	3.22	7.91	0
1E	CU	----	0	0	0	0.00	0.00	0
1F	CU	C052	0	590	82	0.00	0.00	0
	CU	C050						
20	CHP-H	4A	0	374	756	1.04	2.40	0
21	CU	----	0	1472	413	3.51	0.36	0
22	------	----	0	0	0	0.00	0.00	0
23	------	----	0	0	0	0.00	0.00	0

Frame Pacing Delay Being Reported

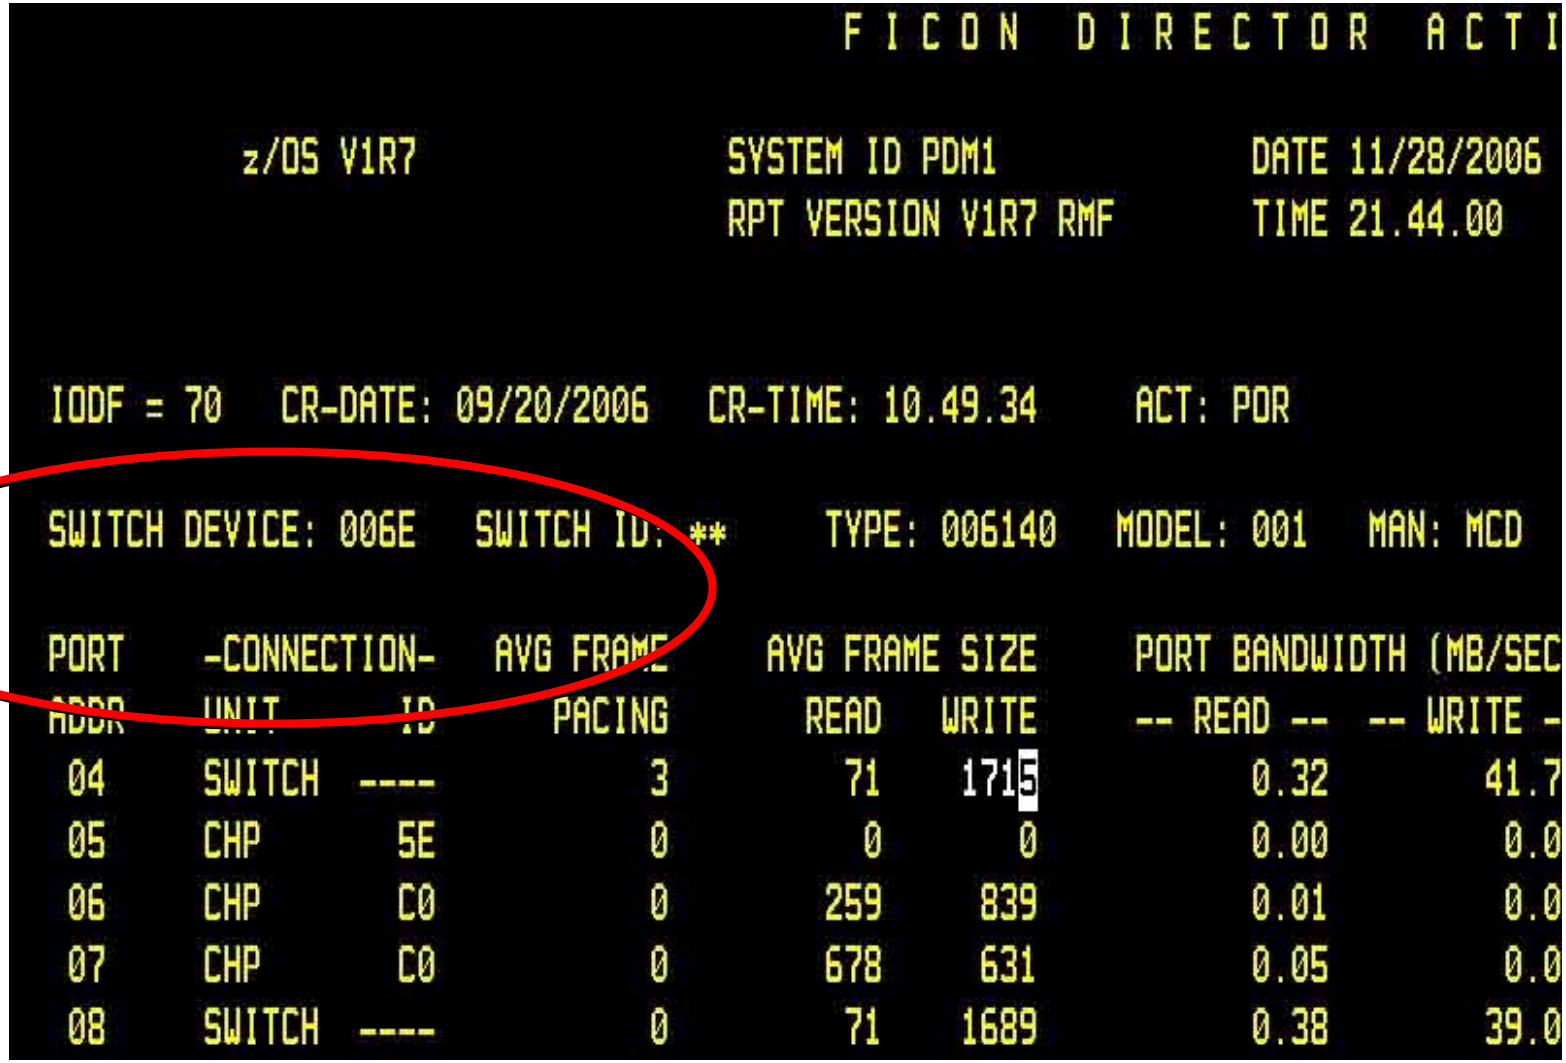

```
                                        F I C O N   D I R E C T O R   A C T I

          z/OS V1R7                     SYSTEM ID PDM1            DATE 11/28/2006
                                        RPT VERSION V1R7 RMF      TIME 21.44.00


IODF = 70   CR-DATE: 09/20/2006   CR-TIME: 10.49.34     ACT: POR

SWITCH DEVICE: 006E   SWITCH ID: **     TYPE: 006140   MODEL: 001   MAN: MCD

PORT    -CONNECTION-   AVG FRAME      AVG FRAME SIZE      PORT BANDWIDTH (MB/SEC
ADDR    UNIT      ID     PACING       READ   WRITE       -- READ --  -- WRITE -
 04     SWITCH  ----          3         71    1715            0.32        41.7
 05     CHP       5E          0          0       0            0.00         0.0
 06     CHP       C0          0        259     839            0.01         0.0
 07     CHP       C0          0        678     631            0.05         0.0
 08     SWITCH  ----          0         71    1689            0.38        39.0
```

Local Frame Pacing Delay

- How can you run out of buffer credits inside a data center?
  - Frame pacing delays occur when multiple, heavily used paths merge into a single FICON link
  - Frame pacing delays can contribute to PEND, DISC, and CONN time measurements

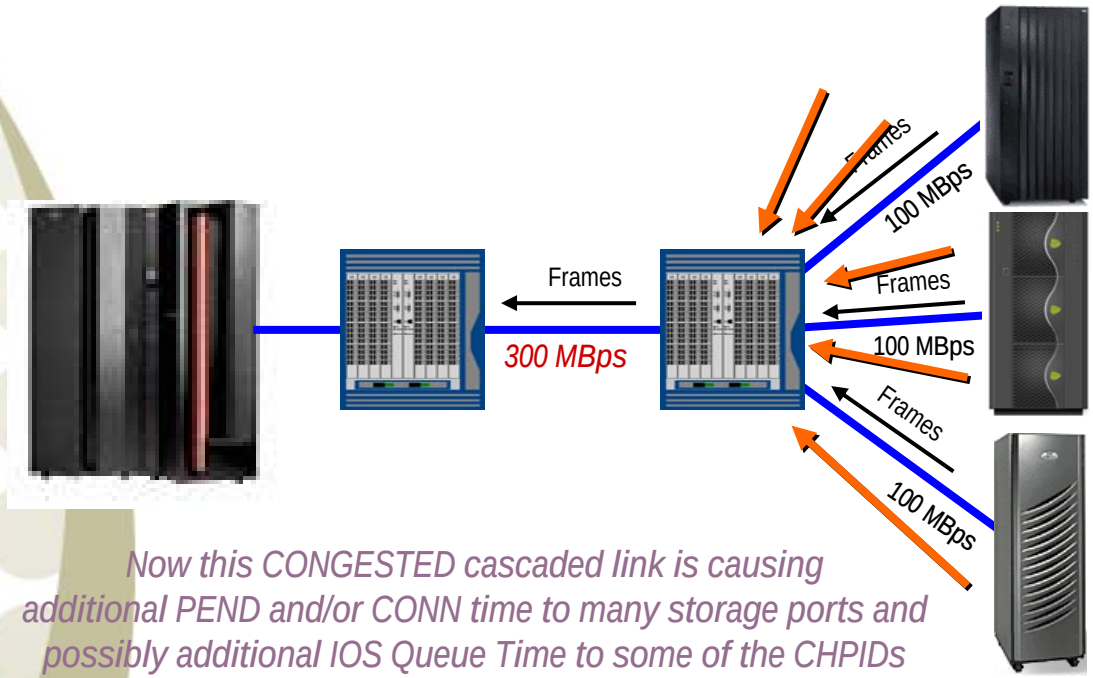

Now this CONGESTED cascaded link is causing additional PEND and/or CONN time to many storage ports and possibly additional IOS Queue Time to some of the CHPIDs

Frame Pacing Delay is caused by running out of buffer credits!

UCBs serviced by these storage ports are probably experiencing additional delays usually reported as PEND Time and CONN Time and sometimes as DISC time

Frame Pacing Delay came about with FC and FICON so it is not a factor in ESCON performance!

Suggestion

- Use the RMF 74-7 record as a way to help narrow down/troubleshoot performance problems in your environment.
- The RMF 74-7 record in and of itself will not allow you to precisely diagnose what is causing a performance problem, other than to indicate BB Credit issues.

QUALITY OF SERVICE (QoS)

Not all data is created equal, nor does all data require equal treatment….

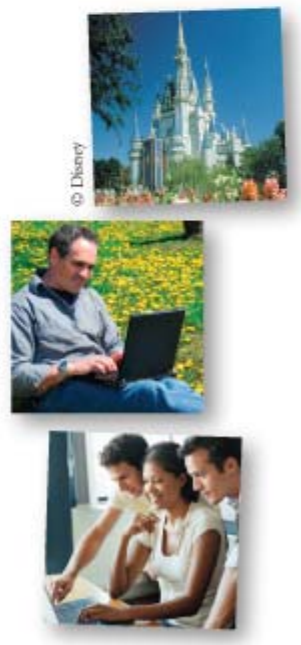

And now for a commercial..

- QoS topic has become a session unto itself
- Friday 8 am, Coronado K Session:
  - End to End Quality of Service (QoS) in FICON Environments
  - Dennis Ng (IBM) and myself co-presenting
- zJournal article in the brand new issue here at SHARE:
  - “FICON and Quality of Service (QoS): Making the Case for True, End-to-End, Host Managed QoS.”

WHAT IS 4 GBIT/SEC FICON REALLY DOING FOR ME?

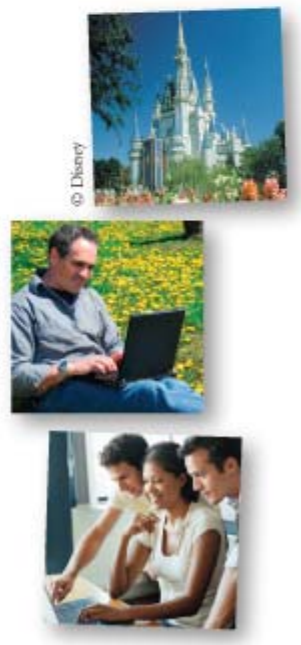

FICON Channel Implementations

Server	Channel Type	Microprocessor	Internal Bus To STI	Data Rate	STI
G5/G6	FICON	166 MHz Power PC 750	32-bit 33MHz	1 Gbit	1
z900	FICON	333 MHz Power PC 750 L2 Cache	32-bit 33 MHz	1 Gbit	1
z8xx/z9xx	FICON Express	333 MHz Power PC 750 L2 Cache	64-bit 66 MHz	2 Gbit	1
z890/z990 z9	FICON Express 2	448 MHz Power PC 750 L2 Cache	64-bit 112 MHz	2 Gbit	1(2 Z9)
z9	FICON Express 4	448 MHz Power PC 750 L2 Cache	64-bit 112 MHz	4 Gbit	2

Recent history of Mainframes
Keeping up with user demand

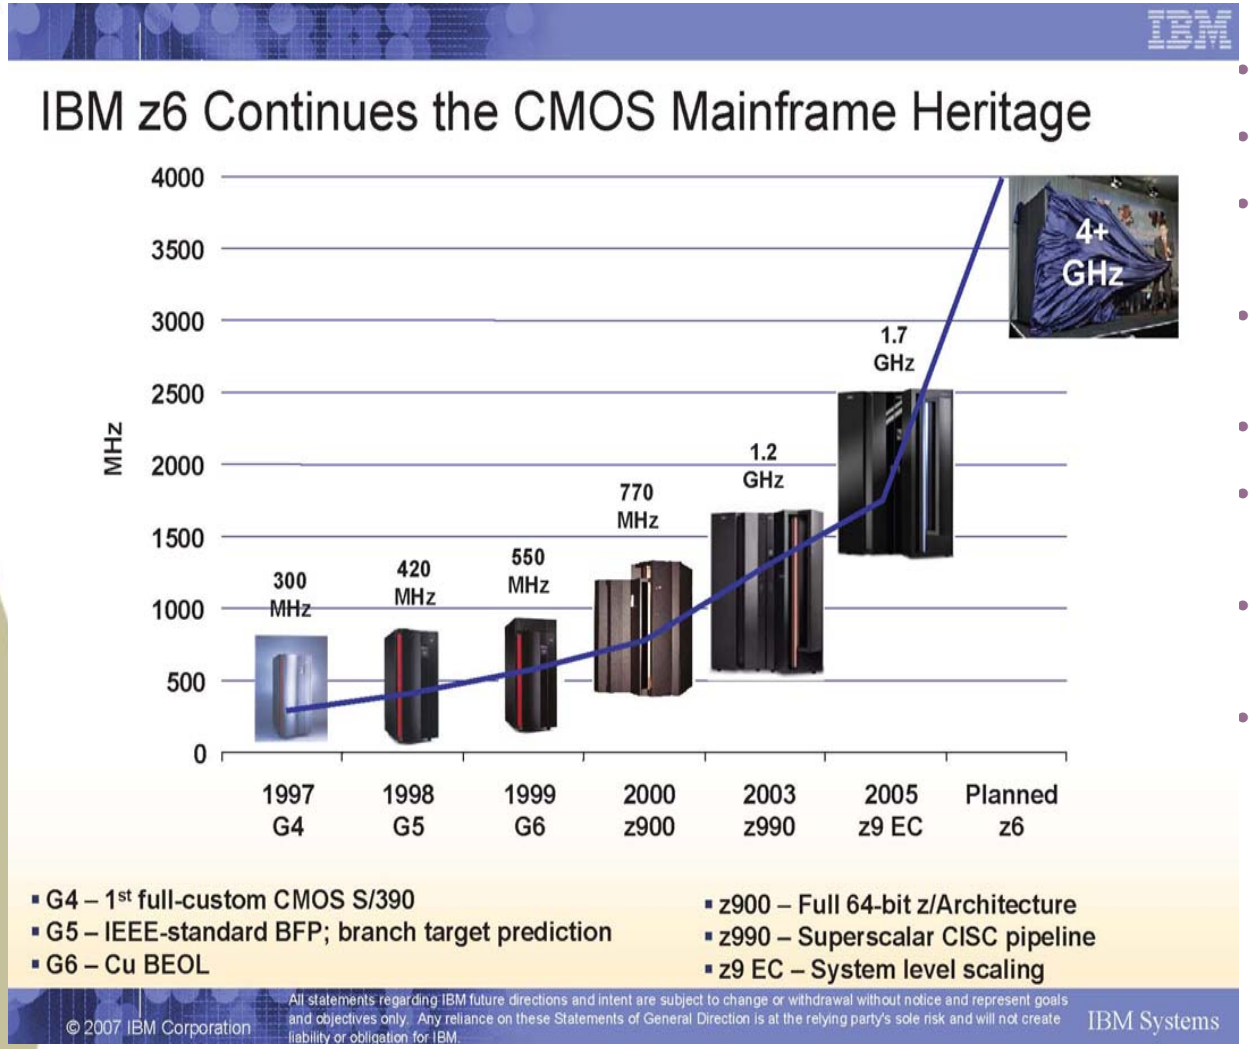

- 991m transistors
- 3 GHz Bus Speed
- 20,000 error checkers in a chip
- Full z/Architecture compatibility
- New performance features
- Enhanced virtualization capabilities
- Significant energy efficiency
- Industry leading RAS

FICON Channel Full-Duplex Read/Write Mix Data Transfers

Reference: IBM

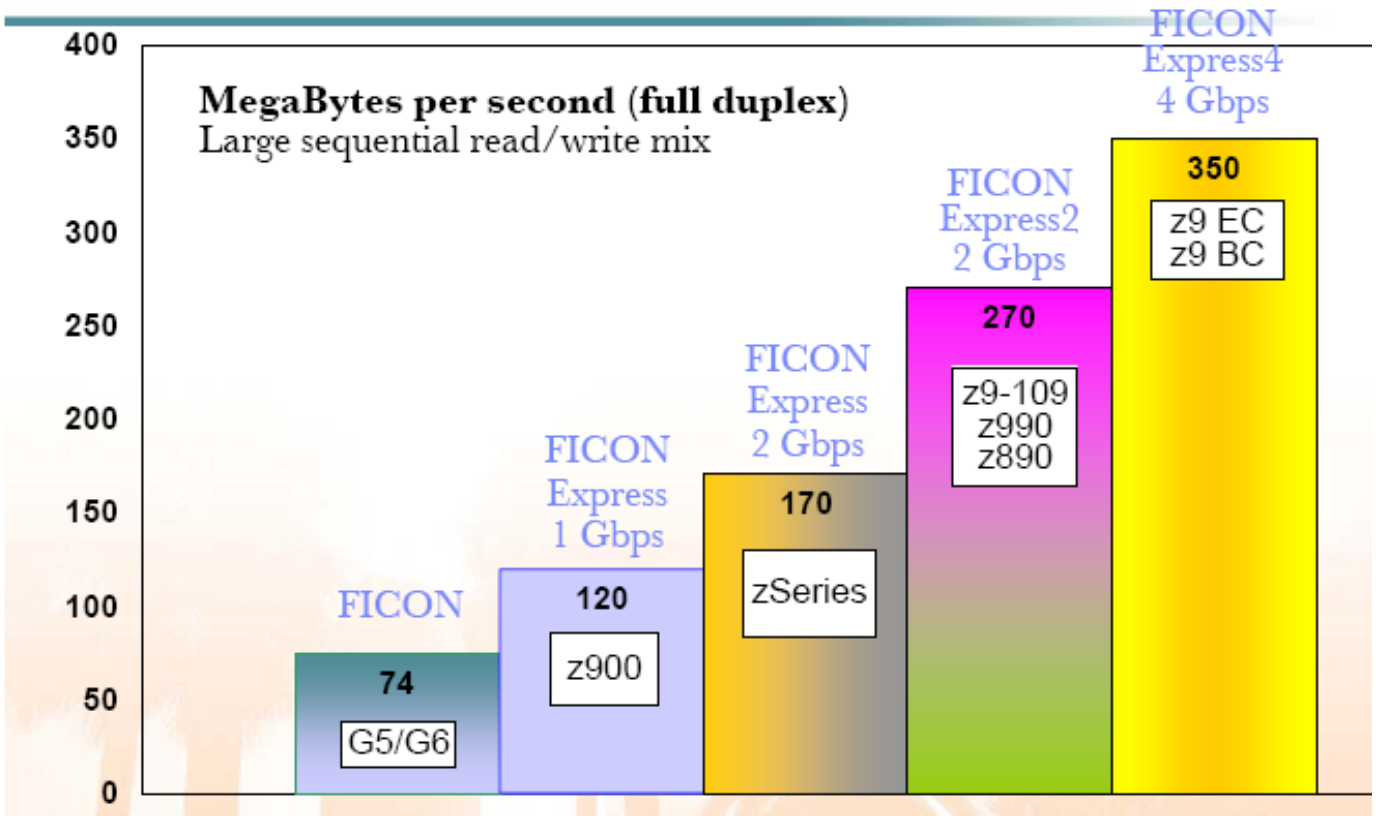

FICON Channel Performance-Start I/Os

Reference: IBM

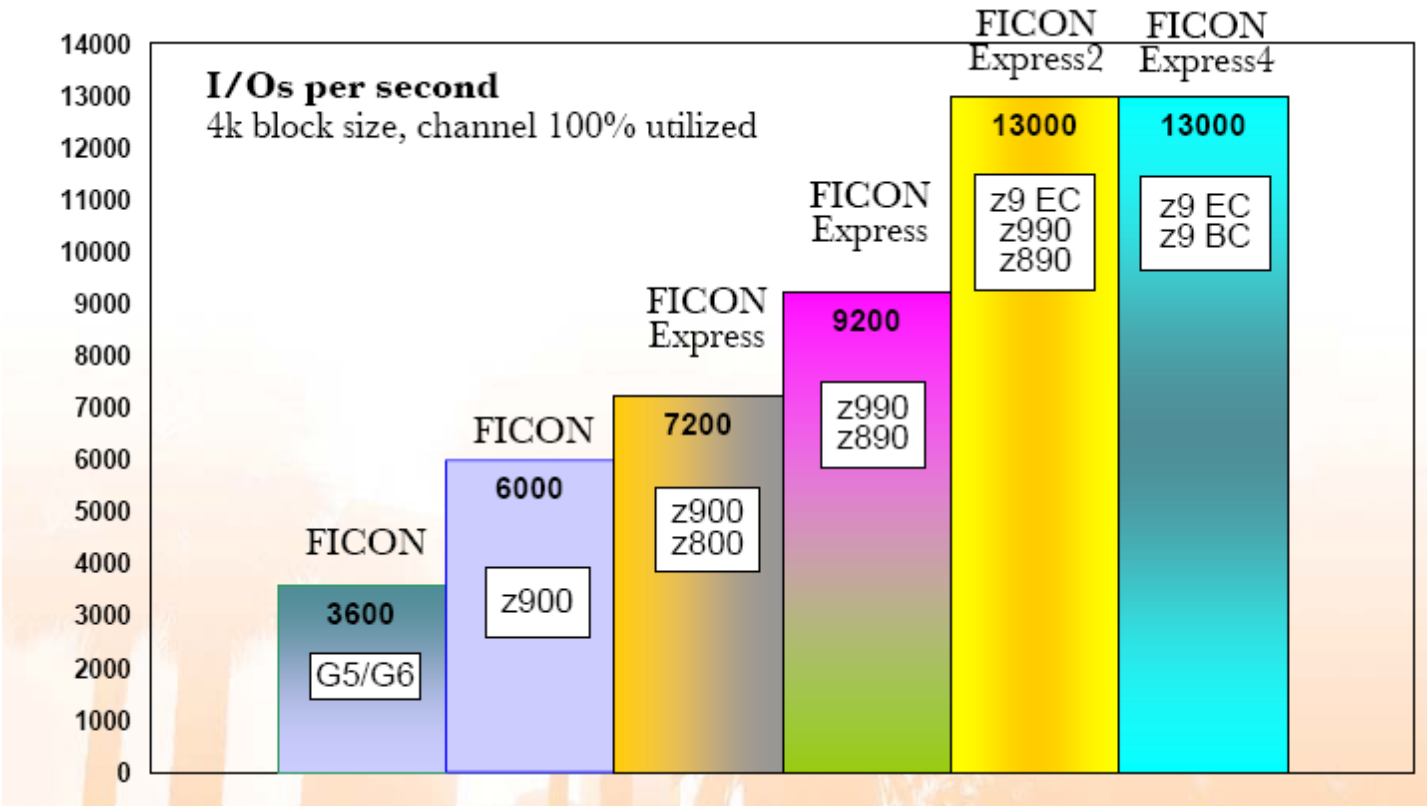

What about the host adapters on the DASD arrays- IBM example

- Full card aggregate
  - ESS800: 150 MB/sec
  - DS8000: 560 MB/sec
  - DS6800: 778 MB/sec
- Single port (reads)
  - ESS800: 145 MB/sec
  - DS8000 2 Gbit: 206 MB/sec
  - DS8000 4 Gbit: 307 MB/sec
  - DS6800: 206 MB/sec

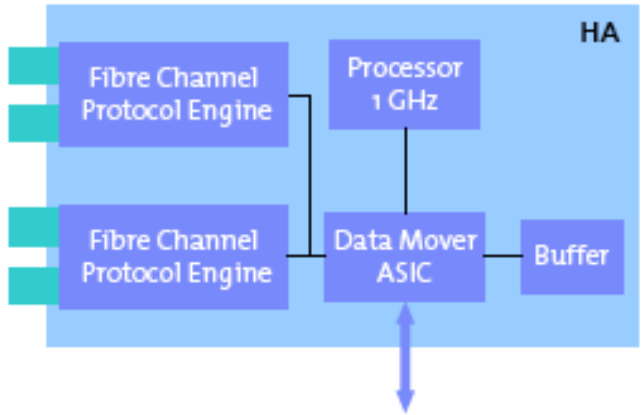

Simplified diagram based on IBM documentation

- Numbers above are for 2 Gbps FICON benchmarks

What to Do?

- Are you driving your 2 Gbit/sec channels hard enough to benefit?
  - Better bandwidth, but not increased I/O
- Are you running MIDAW?
  - If not, you'll probably see little benefit to running 4 Gbit/sec
- Have I modeled my environment with 4 Gbit/sec to see if I will gain anything?
  - Are the gains worth the cost?
- How much will it cost me to implement 4 Gbit/sec FICON?
  - Should I wait for the new 4 Gbit/sec channel cards with better I/O capability?
  - What about my storage?

SHARE
Technology · Connections · Results

Modified Indirect Data Address Words (MIDAWs)

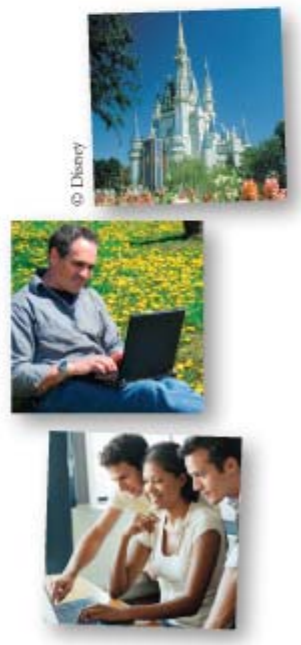

Intro

- MIDAW introduced with z9 to address:
  - Limitations in the existing IDAW implementation
  - Inefficiencies with FICON data chaining
- FICON channels historically have experienced persistent limitations for achievable data rates for small block size traffic
- FICON upper-level protocol requires that Information Units (IUs) contain a single CCW, a CCW and associated data, or only data:
  - Achieves far lower I/O rates for small blocks than open systems in exchange for better data integrity and recoverability
  - Small blocks require more CCWs per track than large blocks-→channels are less efficient for small blocks

Channel Efficiency

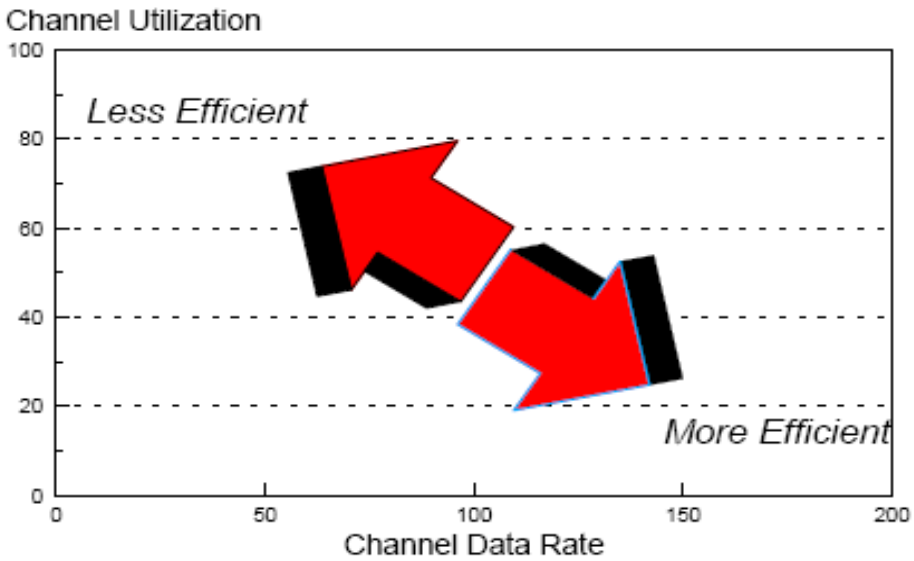

- Small block operations exhaust the microprocessor: channel busy metric
- Large data transfer operations saturate the bus: Bus busy metric
- Channel efficiency: ratio of bus busy to channel busy is an indication of the effective data transfer rate at saturation.
- Effective employment of the FICON channel bandwidth requires balancing channel and bus busy
- Reference: Dr. Pat Artis “Understanding the Performance Implications of MIDAWS”, 2006

What Is the MIDAW Facility?

- It is a system architecture and software exploitation
- New method of gathering and scattering data into and from non-contiguous System z9 storage locations during I/O operations
- By reducing the number of CCWs in some channel programs, FICON performance can be improved in certain applications such as DB2

Improvement in FICON Express 4 Channel Processor Utilization and Bandwidth

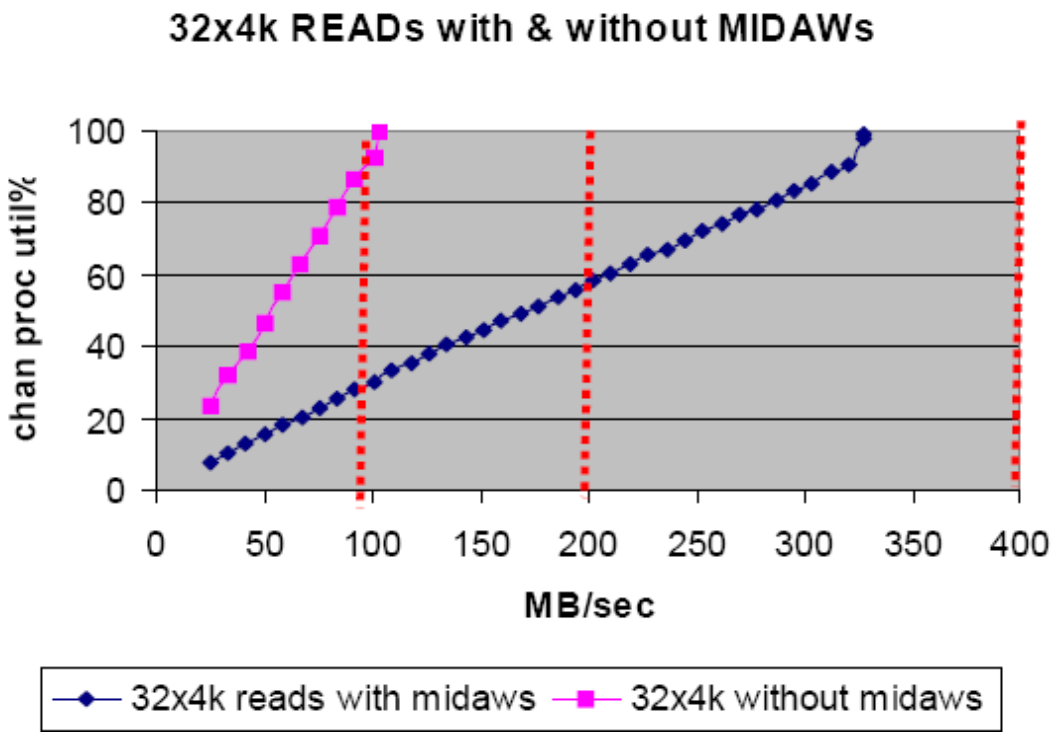

Reference: Cathy Cronin “IBM System z9 and FICON Express 4 Performance Update. SHARE Tampa Proceedings, Feb 2006

More Bytes/Frame→More Efficient Channel

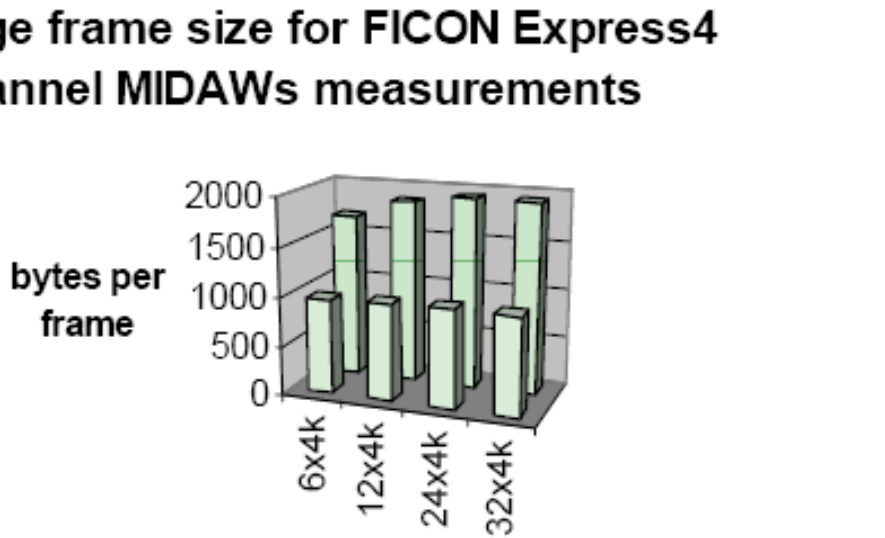

	6x4k	12x4k	24x4k	32x4k
no midaws	965	982	1000	997
with midaws	1692	1873	1943	1943

Reference: Cathy Cronin “IBM System z9 and FICON Express 4 Performance Update. SHARE Tampa Proceedings, Feb 2006

FICON Express 4 Experiment

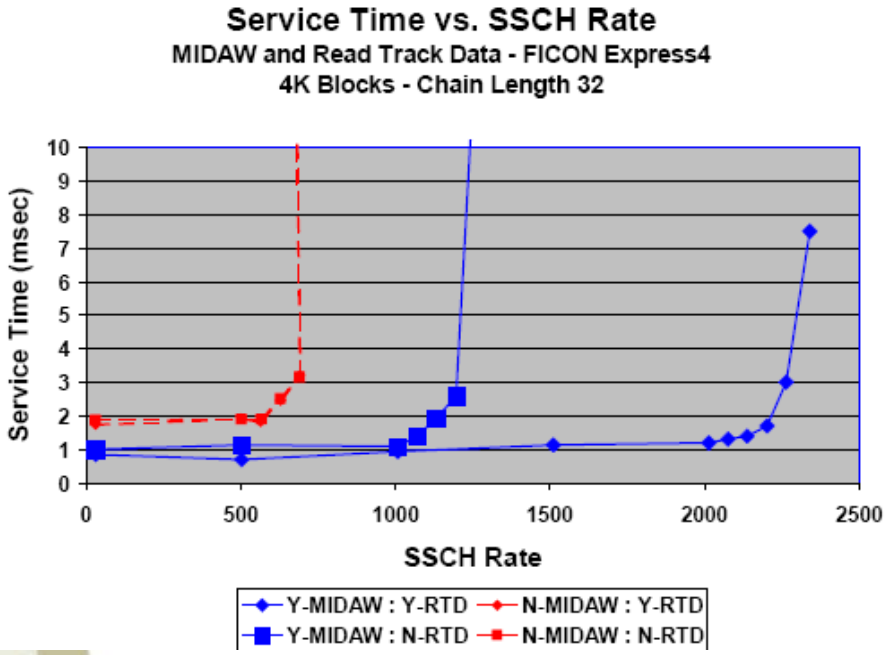

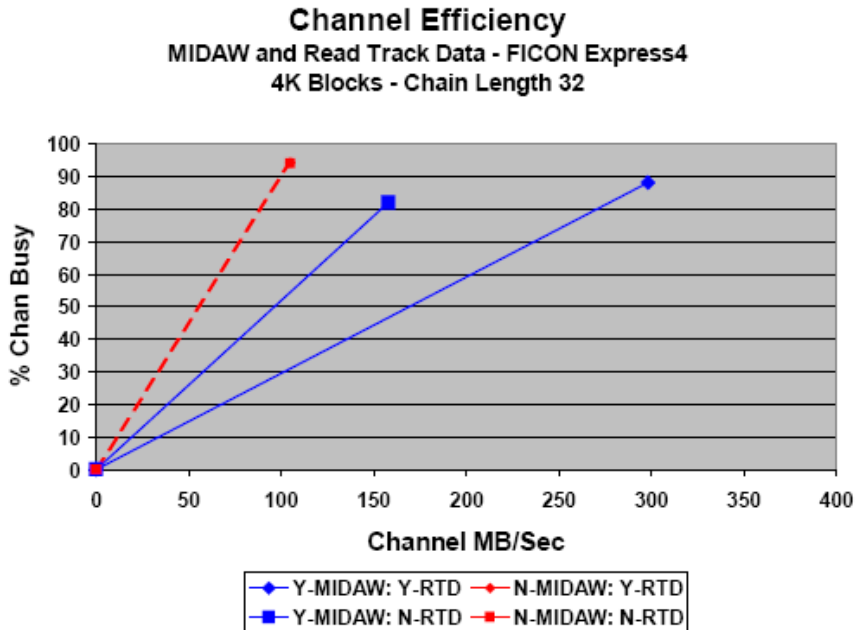

Reference: Dr. Pat Artis “Understanding the Performance Implications of MIDAWS”, 2006

MIDAW Conclusions

- MIDAW improves both the channel efficiency and service times for FICON Express 4 and FICON Express 2
- MIDAW dramatically increases the maximum SSCH rate achievable by subsystems
- MIDAW increases the efficiency of your channel links by transferring more data frames and less command frames
- Today, use of MIDAW is really the only way to get a benefit from moving to FICON Express 4

SHARE
Technology · Connections · Results

OPEN EXCHANGES

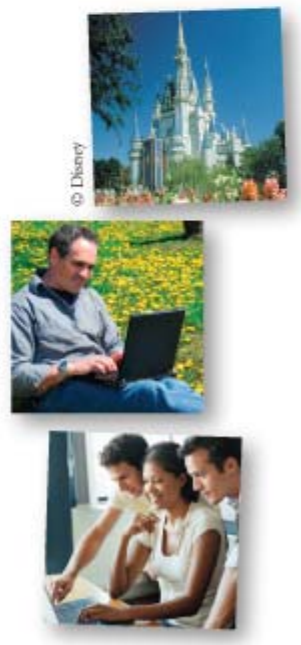

Open Exchanges

- An **Open Exchange (OE)** is a logical resource that represents an I/O being processed by a channel. An OE number is assigned to each SSCH request. FICON and FICON Express/EX2 channels are limited to 32 concurrent OEs for z8x0 and z9x0 processors. FICON Express 2 and FEX4 on the z9 is limited to 64 concurrent OEs.

Open Exchanges (2)

- An OE is a conversation number that is assigned by the Fibre Channel port. It is a unique hardware generated number that is assigned to every work unit related to a specific I/O. This value is contained within the frame header.
- Each work unit is also assigned a sequence number by the Fibre Channel port.
- The number of OEs for a channel can be estimated using Little's Law.

What's the big deal about Open Exchanges?

- Scenarios
  - Customers with channel utilizations<35%
  - PEND and CONN time skyrocketing beyond a certain number of SSCHs
  - Response times skyrocketing
  - What gives?

ESCON vs. FICON

- Traditional method of using channel utilization works well with ESCON
  - 1 I/O at a time
- FICON is very different
  - Multiple I/Os at a time
  - CONN elongation

Estimating the Number of Active OEs

- The number of Open Exchanges is NOT reported by RMF:
  - A calculation/estimation must be done to obtain this number
- Active_OEs = Chan_I/O_Rate * OE_Duration
  - *For the OE_Duration, we now use the sum of the average CMR, DISC, and CONN times for the I/Os performed by the channel:*

    Active_OEs = Channel_I/O_Rate * (CMR + DISC + CONN)
- MXG member ANALFIOE performs this calculation
- Intellimagic's RMF Magic SW performs this and graphs the trends

CMR=Command Response Time

Subsystem Interface Open Exchanges

- ANALFIOE adds a per-device OE value to the MXG Type 74 file
- Current DASD subsystem FICON interfaces are capable of supporting at least 48 concurrent OEs:
  - Manage them to values less than 20

SHARE
Technology · Connections · Results

SOFTWARE THAT CAN HELP

What is “Good Disk Performance”?

SHARE
Technology · Connections · Results

Traditional Questions

- Pending time 2 ms?
  - Good or bad?
- Disconnect time 2 ms?
  - Good or bad?
- Connect time 2 ms?
  - Good or bad?

Easy because at least we can measure these variables

But these questions no longer reveal disk subsystem health

Important Metrics

- Host Adapter busy
  - RMF does not report this
- Back-end Array Group busy
  - RMF does not report this
- Connect time elongation
  - RMF does not report this
- FICON open exchanges
- Effective FICON data rates

These metrics tell us what we need to know, but RMF does not report them!
(RMF Magic does!)

53

Example for Disconnect

Traditional Interpretation

Physical Device = Logical Device

- ‘overhead time when a device might disconnect to perform positioning functions such as SEEK/SET SECTOR, as well as any reconnection delay’

  [RMF z/OS 1.8 manual...]

New Rules

Physical Device =NOT Logical Device

- Disconnect time reasons include:
  - Read miss .. (database access)
  - Synchronous Copy
  - Microcode and back-end management delays (physical pend)
  - FW Bypass .. (FlashCopy or overcommitted write cache)
  - Pacing (asynchronous remote copy)

54

Understanding Disconnect Time

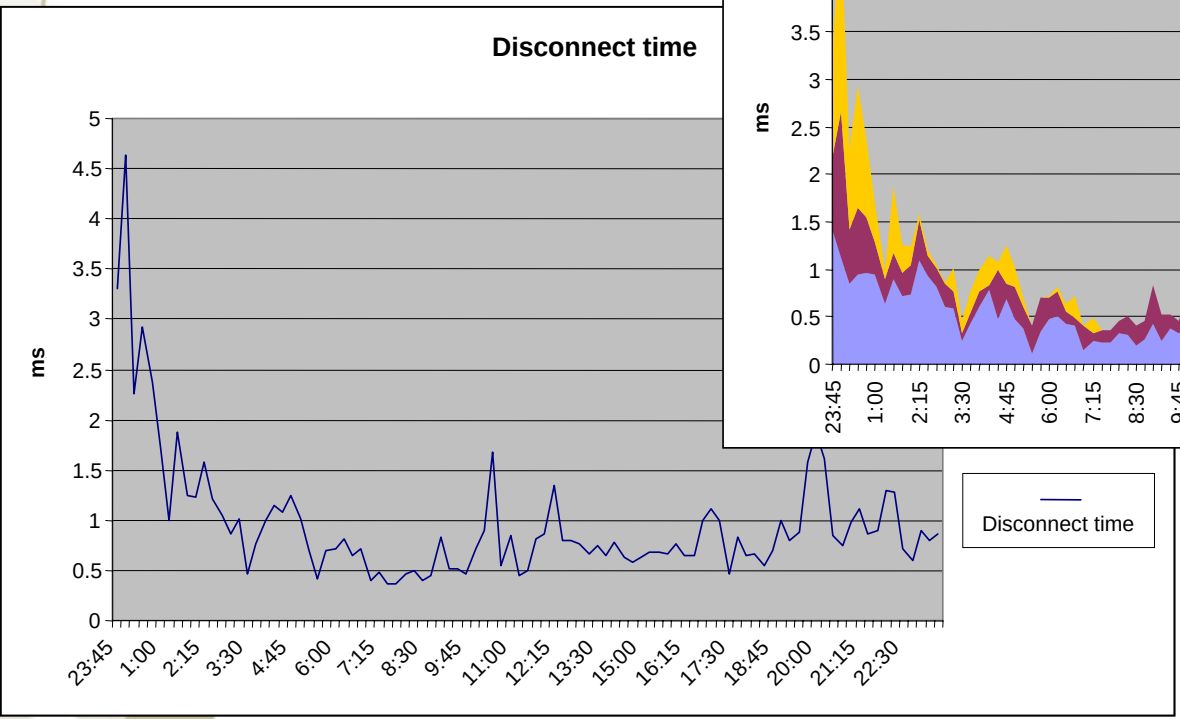

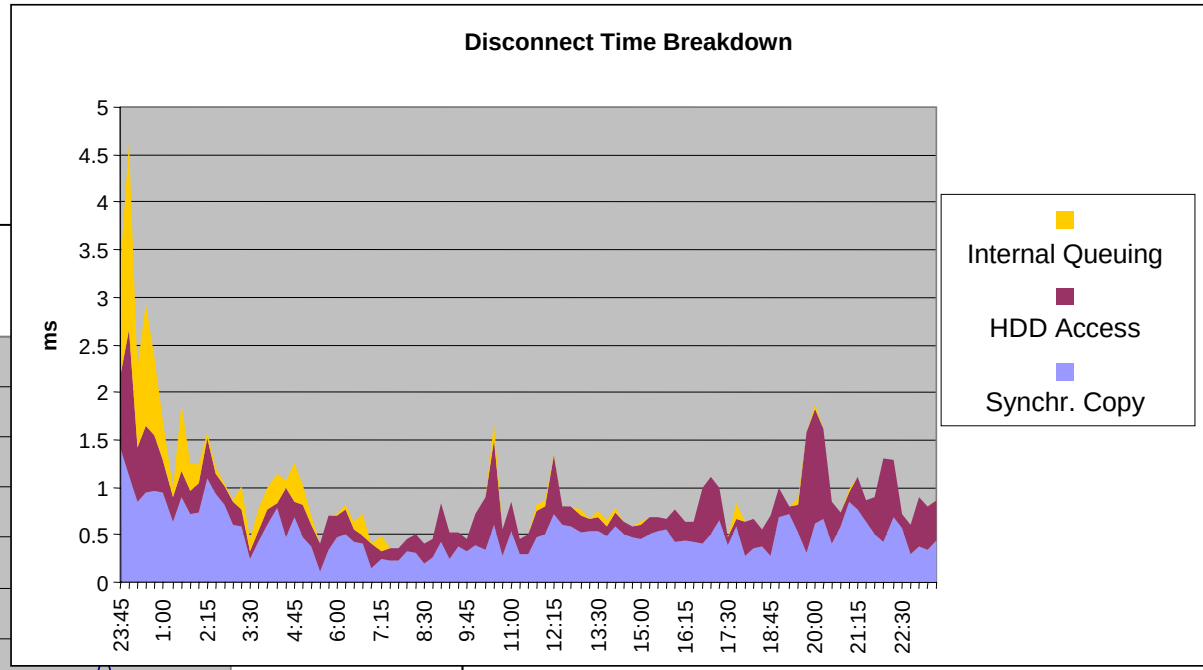

Which picture tells you more?

55

Channel example

- ESCON channels could run at relatively high utilization
- but utilization is misleading
- Should read ‘occupancy’ rather than utilization

- 2GB FICON is theoretically 10x faster
- But it is a multiplexing technology
- Parallel Access Volumes (PAV) and Multiple Allegiance (MA) compound the problem
- Visibility beyond the channel is necessary

56

Monitor FICON Health

(Summary Charts & Reports)

- Effective Data Rate shows data rate while connected
- Good indicator of channel overload (and technology)

Low value indicates that more connect time is needed to transfer a given amount of data; i.e. FICON elongation.

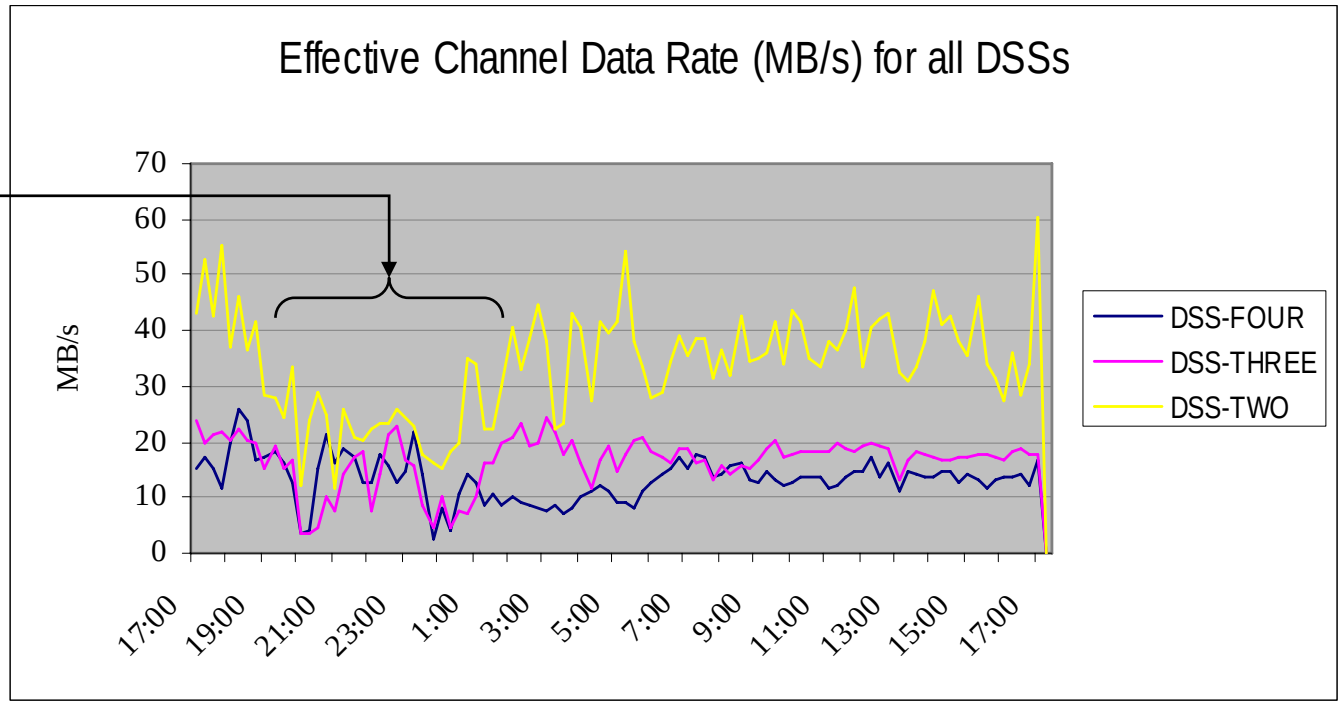

FICON Effective Data Rate

- Effective Data Rate shows data rate while connected
- Good indicator of channel overload (and technology)

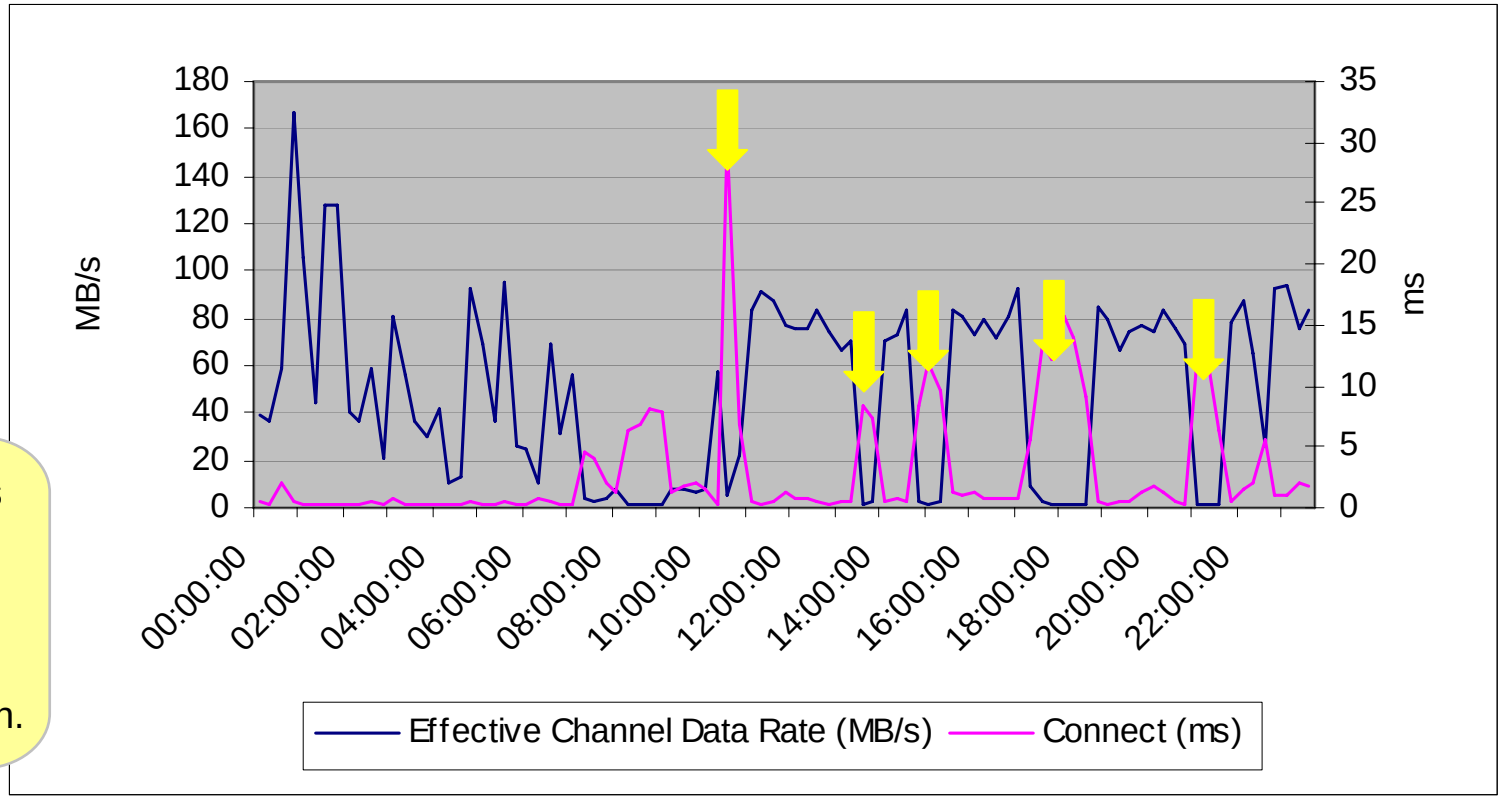

Low value indicates that more connect time is needed to transfer a given amount of data; i.e. FICON elongation.

58

Intellimagic RMF Magic
Effective Data Rate & Connect

- By plotting effective data rate and connect together it is easier to spot FICON elongation

Chart generated when requesting LSS-level response time charts (Breakdown charts)

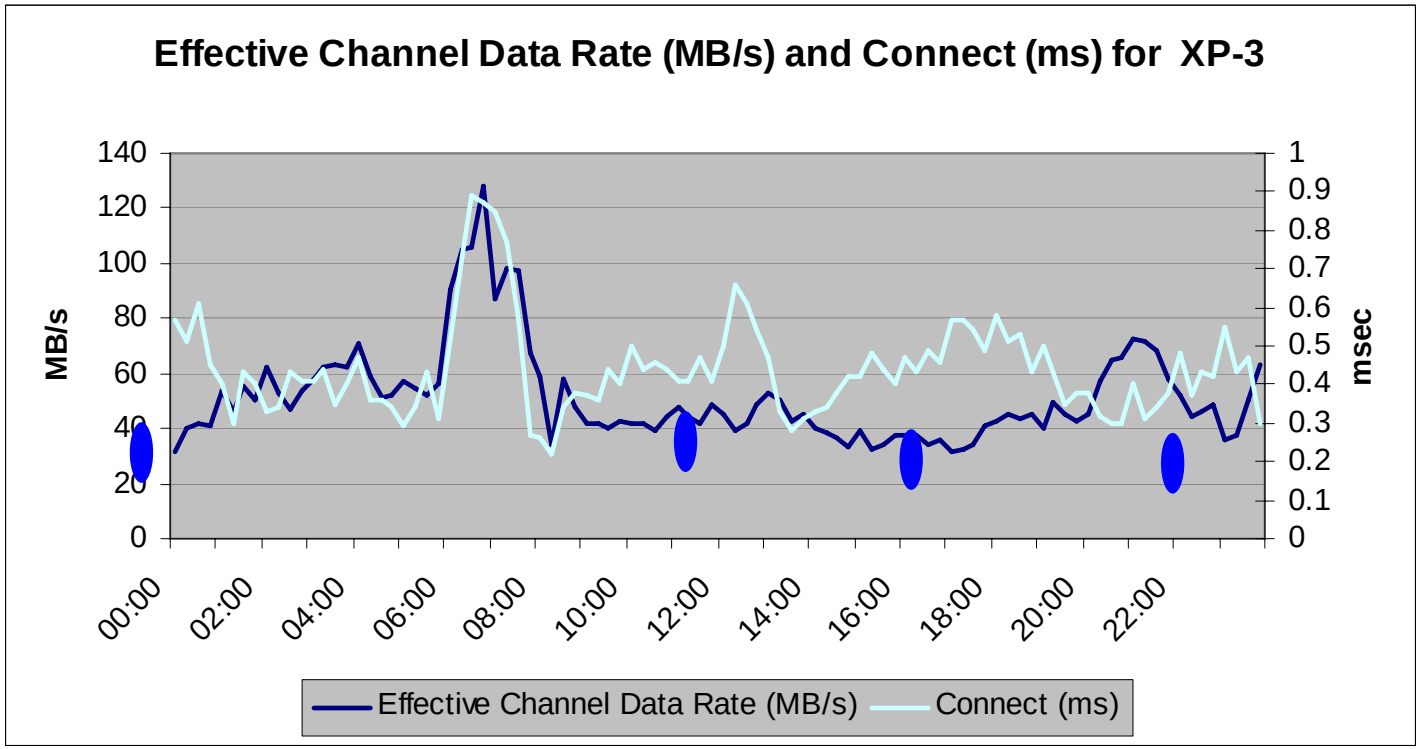

FICON Open Exchanges

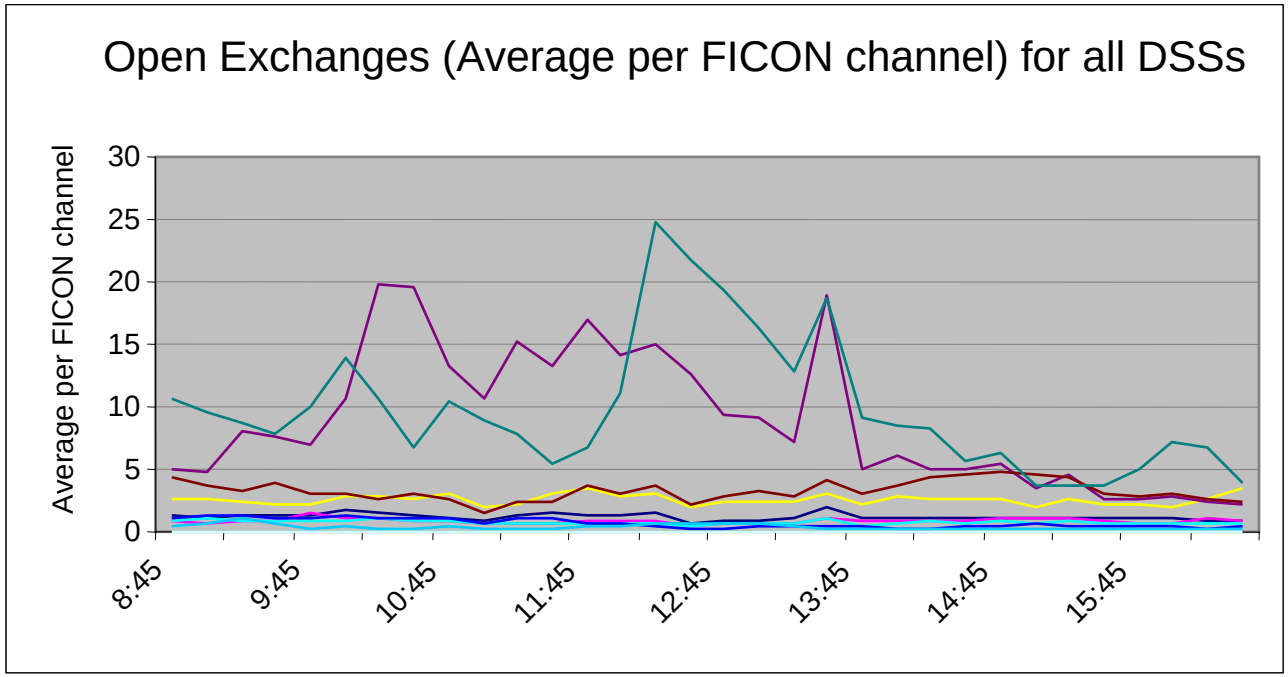

- Open Exchange is active I/O operation on FICON channel, maximum is 32 at any time, or 64 for newer FICON cards
- Chart shows average per Disk Subsystem link over time

Conclusion and observations

- FICON is much different than ESCON, hence how you manage it needs to be different.
- Performance management in the FICON space tends to be more of a react and panic damage control drill rather than a proactive planning exercise.
- 2 quotes:

Be Pro-active!

- “Setting the goal is not the main thing, it is deciding how you will go about achieving it and staying with that plan”.
  - Tom Landry

- “Some people want it to happen, some people wish it would happen, some people make it happen.”
  - Michael Jordan

SHARE
Technology · Connections · Results

Questions?

Thank You