



# High Performance FICON for System z (zHPF)

April 15, 2009

**Paul S Hagen**

**IBM** Washington Systems Center



**ON** DEMAND BUSINESS™

© 2009 IBM Corporation



# Agenda

- **Overview**

- What is zHPF
- Highlights of zHPF

- **zHPF hardware and software requirements**

- **Performance Analysis**

- Summary - z10 zHPF performance white paper Jan 28, 2009
- Summary - DS8000 Release 4.1 zHPF presentation, Dec. 2008

- **Tools and commands**

- Missing Interrupt Handler
- RMF Channel Activity Reports Enhancements
- z/OS Commands and Output

- **Migration considerations**

- **Summary**

- **References**



# Overview

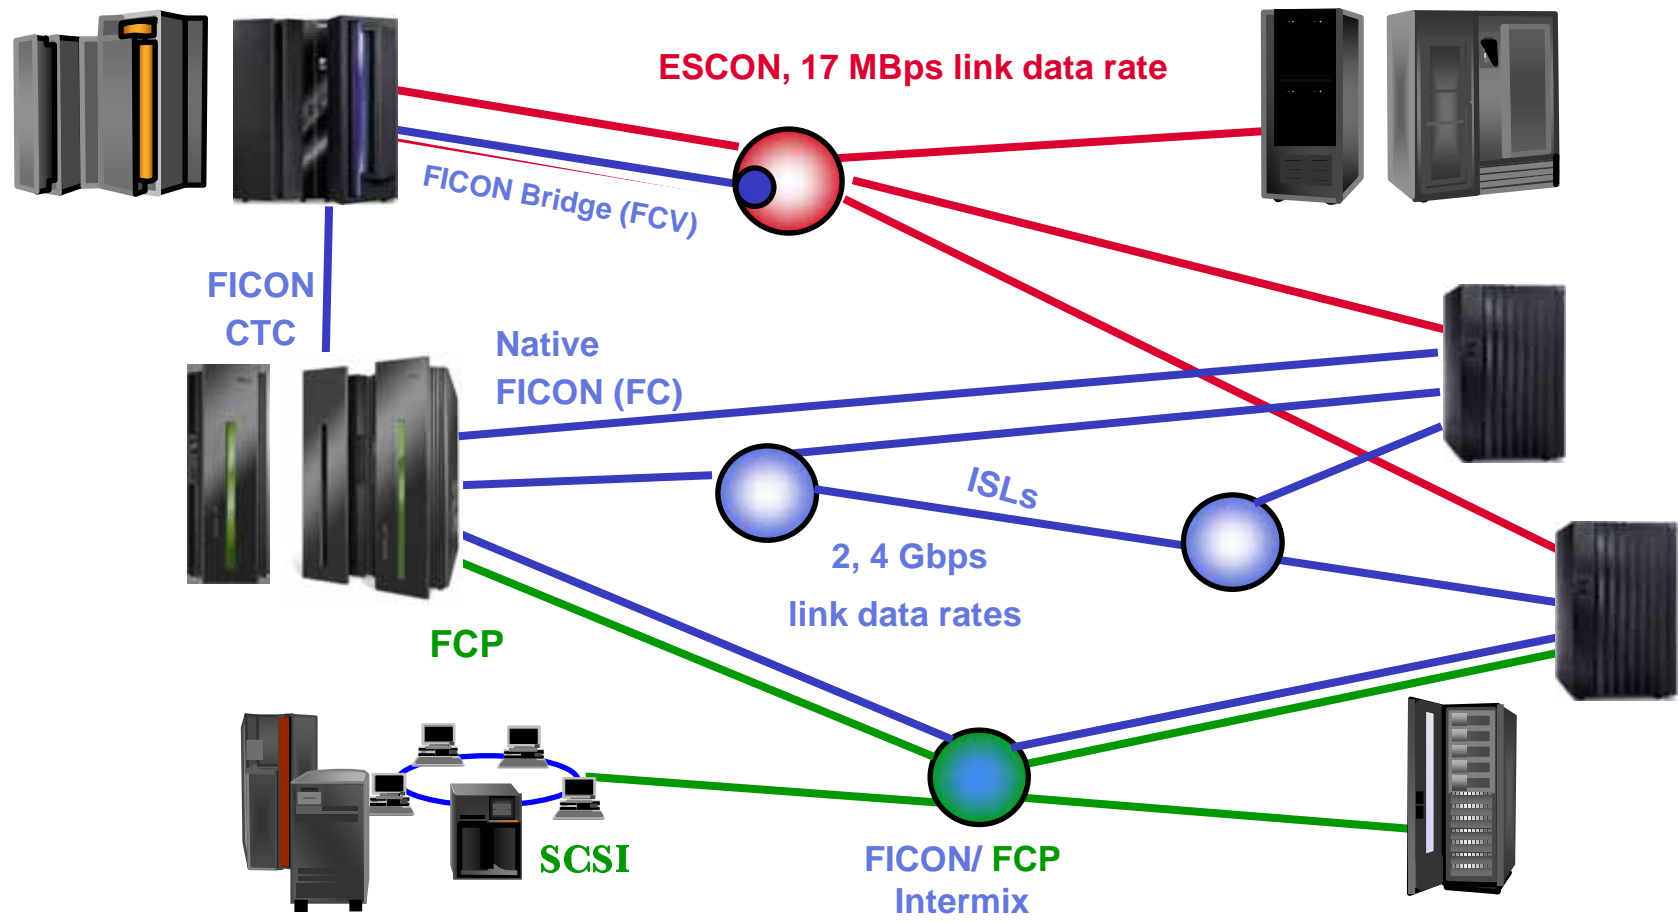


## What is High Performance FICON ?

- **High Performance FICON for System z (zHPF) is a new data transfer protocol that is optionally employed for accessing data from an IBM DS8000 storage and other subsystems.**
  - Data accessed by DB2, PDSE, VSAM, zFS and Extended Format SAM can benefit from the improved transfer technique
  - May help reduce the infrastructure costs for System z I/O by efficiently utilizing I/O resources so that fewer CHPIDs, fibers, switch ports and control unit ports are required
  - Complements the Extended Address Volumes for System z (EAV) strategy for growth by increasing the I/O rate capability as the volume sizes expand vertically

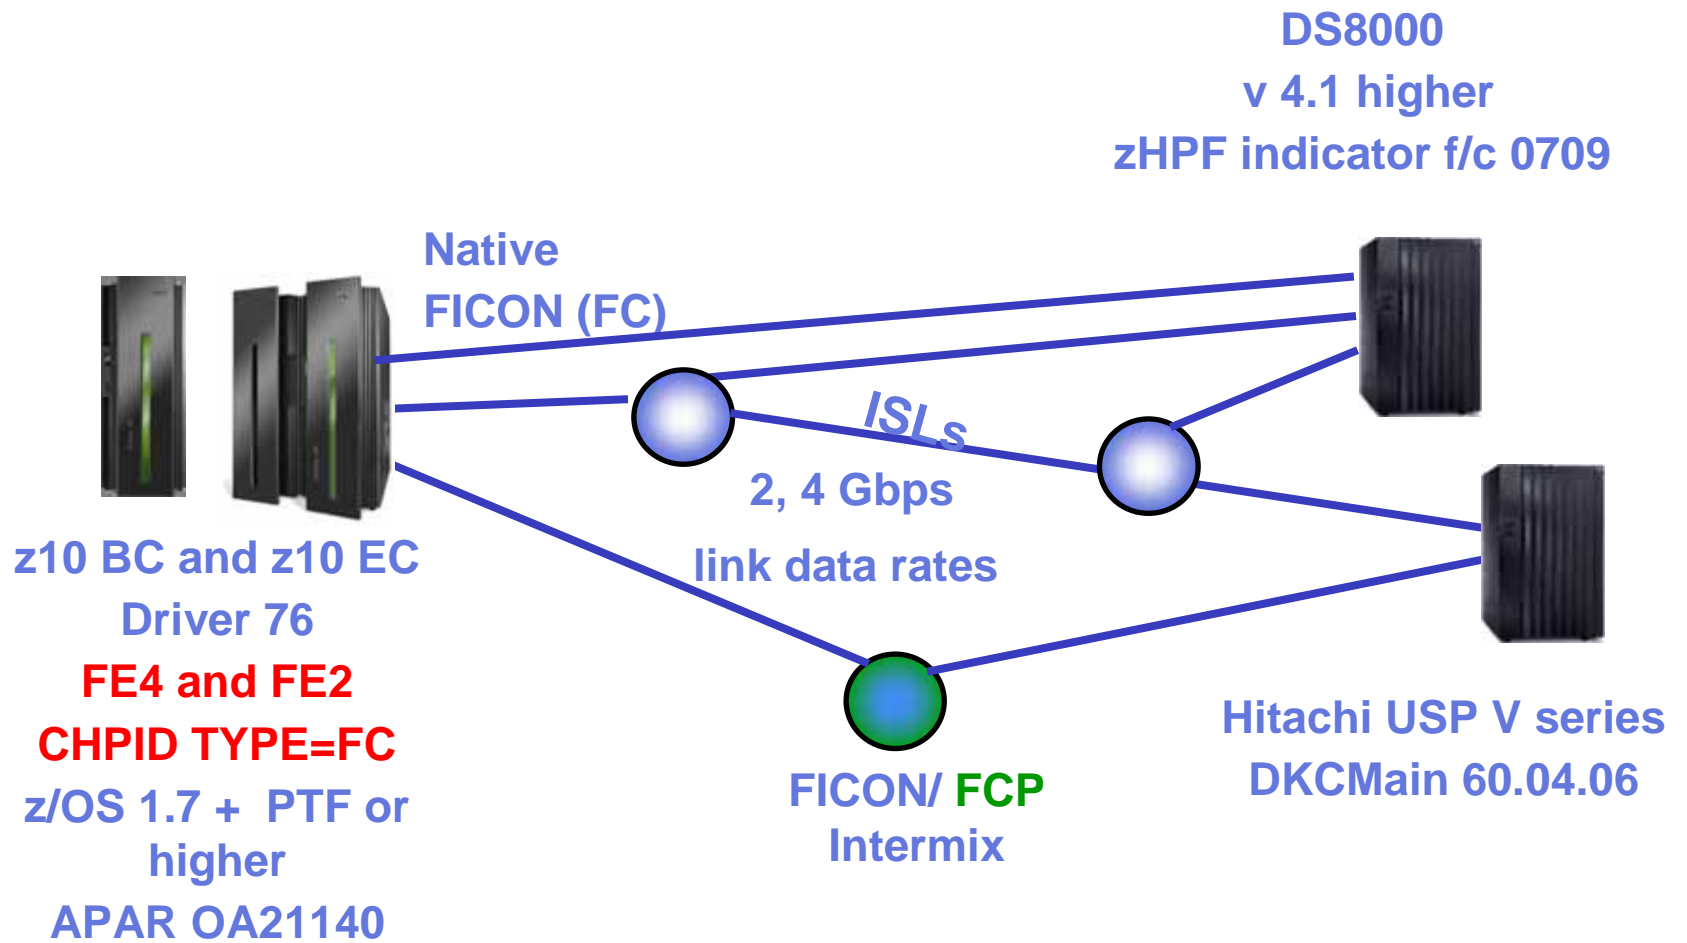


# System z Storage Area Network (SAN)





# System z Storage Area Network (SAN) with zHPF



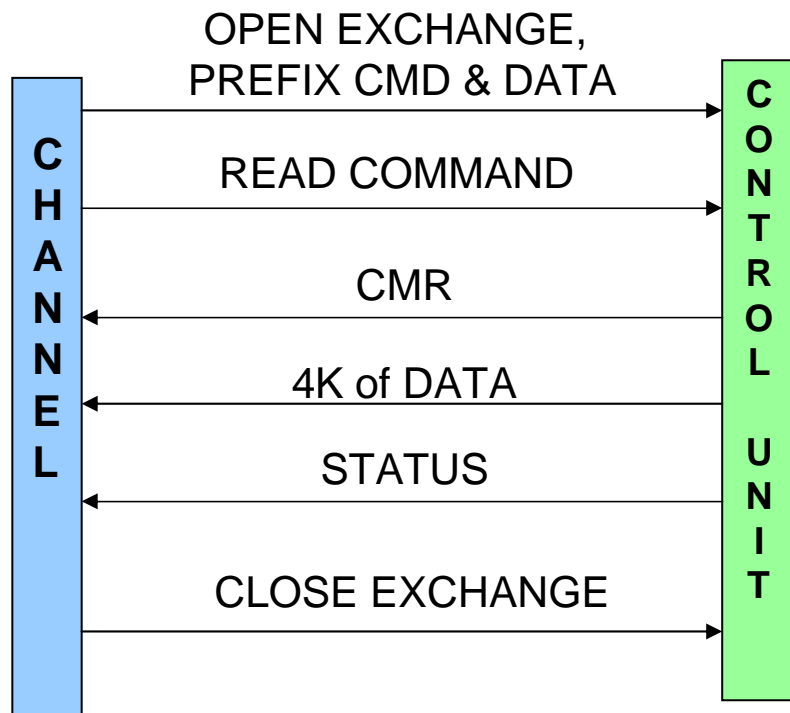
# High Performance FICON Highlights

- **Maximum Application Benefit for Typical OLTP Workloads**
  - Media Manager builds the Transport Control Word (TCW) versus the Command Control Word (CCW)
  - Only data accessed by DB2, PDSE, VSAM, HFS, zFS or Extended Format SAM benefit from the enhanced transfer technique
  - Maximum I/O rate for a channel with a simple 4KB read hit benchmark doubles with zHPF
  - Realistic production workloads with a mix of data transfer sizes may see up to 30% savings in channel utilization. compared to FICON
  - OLTP Workloads that exploit zHPF could see up to 30% improvement in DS8000 throughput.
- **DS8000 Code Structure Optimized for zHPF**
  - Streamlined Internal Communication Protocols
- **Improved RAS and Workload Management**
  - I/Os are queued in control unit when a device is reserved by another host
  - Additional channel and control unit diagnostics for MIH conditions
- **Compatibility Between Existing CCWs and New TCWs**
  - Bilingual Channel and Control Unit Ports
  - CCWs continue to use FICON protocols
  - TCWs use new Transport Mode Protocols
  - **CHPID TYPE = FC ... NO HCD changes required**

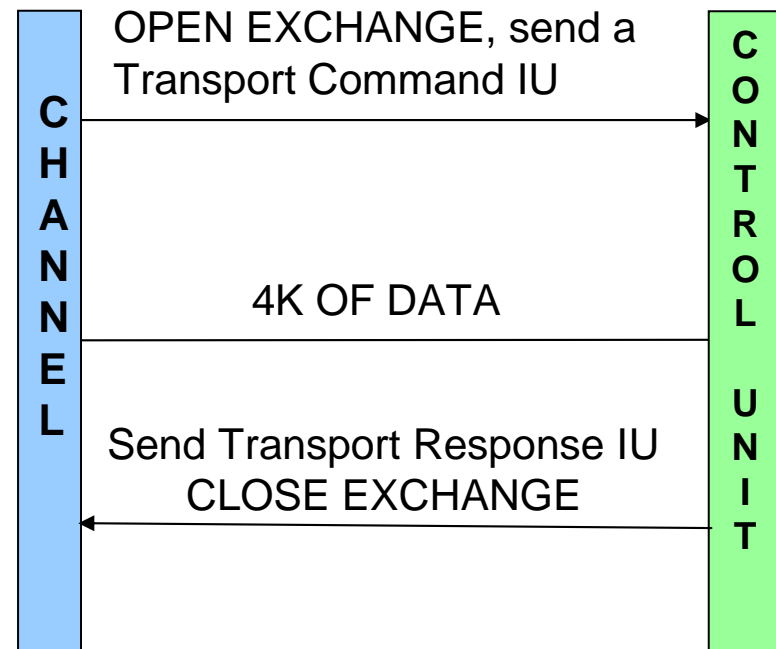


# Link Protocol Comparison for a 4KB READ

## FICON



## zHPF



zHPF provides a much simpler link protocol than FICON



## Maximum number of open exchanges increased from 64 with FICON to over 600 with zHPF

- **Number of open exchanges is a function of activity rates and service times**
  - An open exchange is an I/O that is active between the channel and the control unit and it includes I/Os that are cache hits as well as cache misses which experience a delay before the data can begin transferring
- **zHPF compared to FICON activity rates are higher now & expected to be higher in the future**
  - What is counted in service time for zHPF is potentially different from FICON for cases with **device reserves held**
    - i/o's waiting for devices reserved were queued at the SAP for FICON & not counted as an open exchange
    - Can be queued at the CU with zHPF and therefore counted as an open exchange



# Hardware and Software Requirements



# zHPF hardware and software requirements

## ■ Hardware

- IBM System z10 Server (z10 EC and z10 BC)
  - Driver level 76
- DS8000 version 4.1 or Higher
  - Code bundle R12q.9b081017b ( Bundle 64.1.16.0) for the R4.1 High Performance FICON indicator (feature code 0709 - one time charge)  
GA Date: Oct. 24<sup>th</sup>, 2008
- FICON connectivity (Express2 or Express4)
  - Native 4Gb end to end is recommended for maximum benefits
  - All online channels to the DASD Logical Control Unit must **MUST BE** FEx2 or FEx4 to support zHPF. The inclusion of any non-compliant features will result in the entire path group implementing FICON operation and not zHPF.

## ■ Software

- Minimum Level: z/OS 1.7 plus prerequisite PTFs
- RMF – Requires APAR OA21140 for new function support for High Performance FICON (zHPF)
  - For example, both FICON and zHPF operations are now reported in the RMF “Channel Activity” Report



# Performance Analysis



# Summary - z10 Performance Whitepaper




## CONNECT time and PEND time are calculated differently

- All of the current response time components reported for FICON (CONNECT time, PEND time, CMR time, DISC time and IOSQ time) will be reported for zHPF I/O operations also but the calculations for some of these response time components are different
- With zHPF, CONNECT time is accumulated by the control unit (CU) instead of the channel. This means the starting and ending points for the CONNECT time calculation are different
- Up to 50 microseconds of the time reported in PEND time may actually belong in CONNECT time
- **The total zHPF response times are correct and therefore show the response time benefits of zHPF**
- An update to the algorithms used to calculate zHPF PEND and CONNECT time is under development
- An updated zHPF performance paper will be released when new measurement data is available



# High Performance FICON for System z (zHPF)

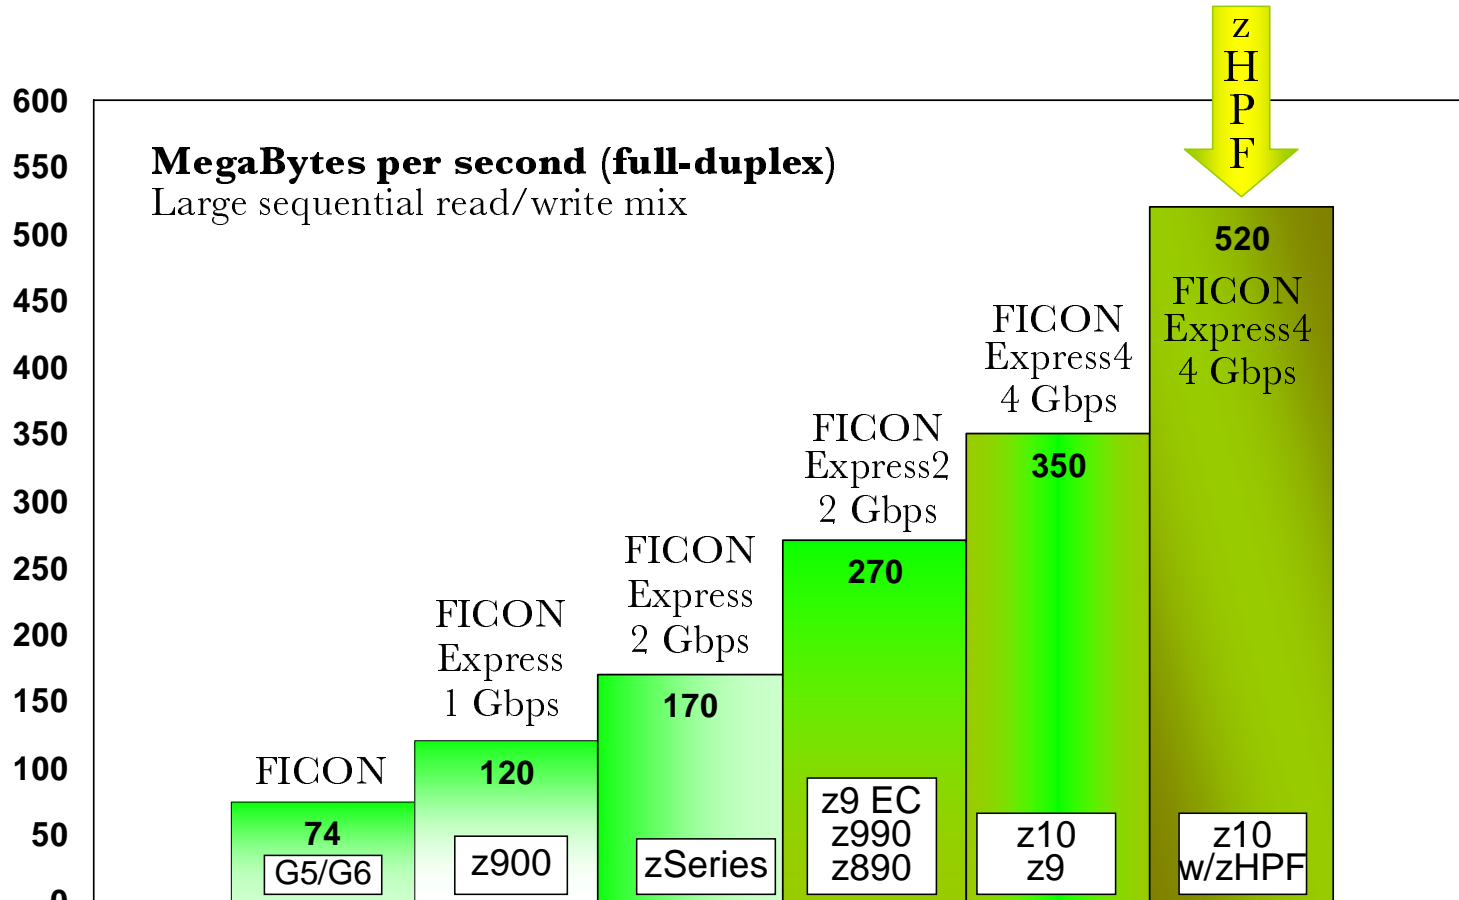
- I/O driver benchmark performance results for maximum I/Os per second





# High Performance FICON for System z (zHPF)

- I/O driver benchmark performance results for large sequential read/write mix

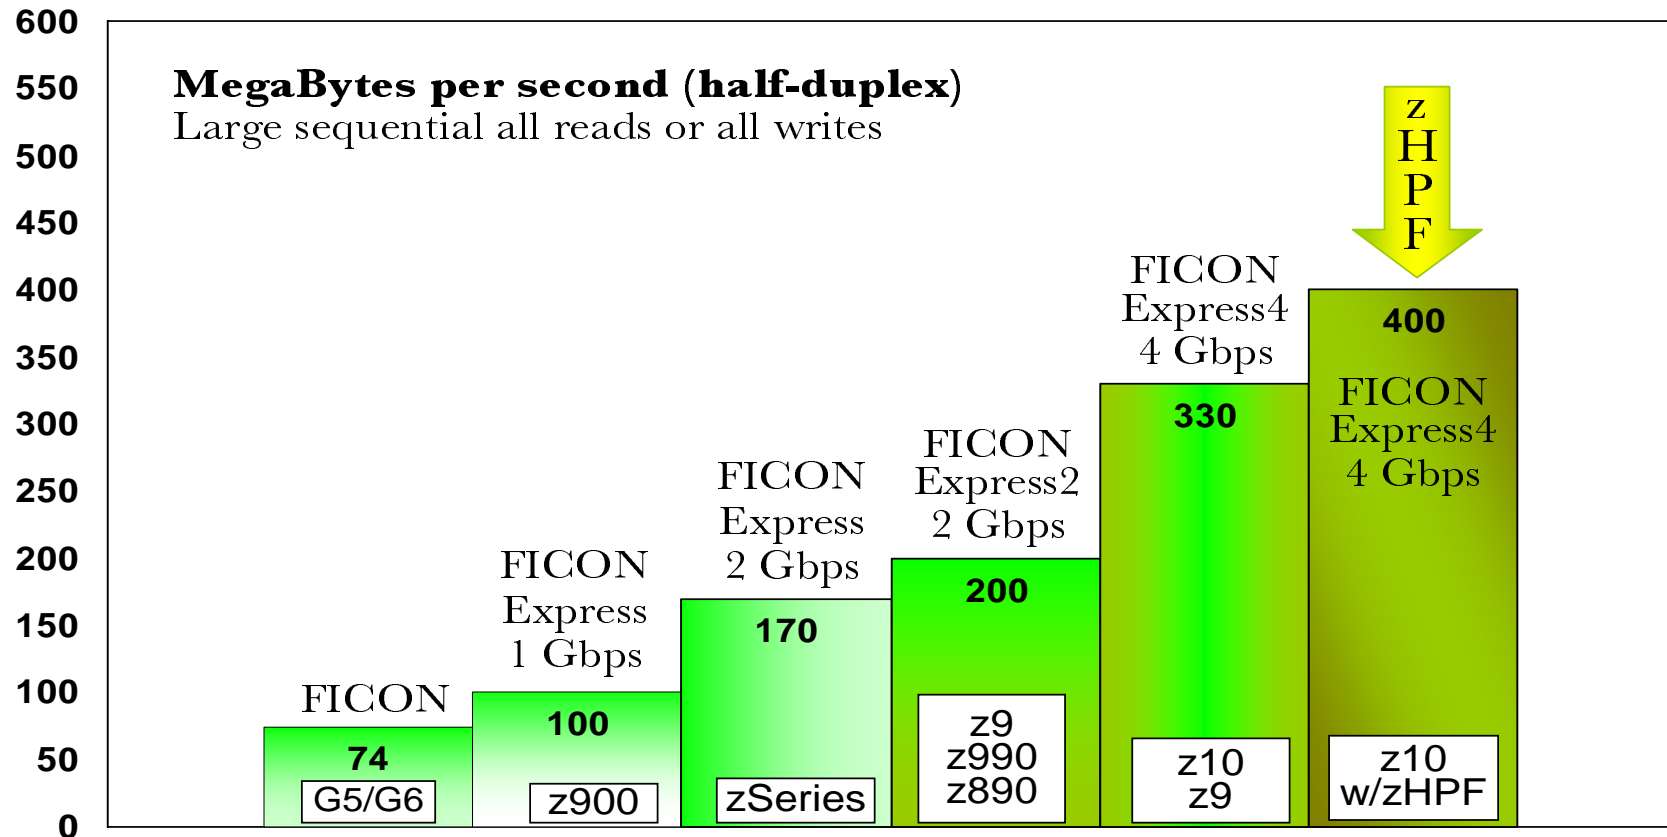


\*This performance data was measured in a controlled environment running an I/O driver program under z/OS. The actual throughput or performance that any user will experience will vary depending upon considerations such as the amount of multiprogramming in the user's job stream, the I/O configuration,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



# High Performance FICON for System z (zHPF)

- I/O driver benchmark performance results for large sequential reads or writes



\*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.




# Summary - DS8000 Performance Results



# DS8000 – Summary of zHPF vs FICON Cache Hits and Read Misses: Channel/HA Scalability


**DS8300 Turbo**  
**384 Disks**  
**73GB/15K RPM**





# DS8000 – Summary of zHPF vs FICON OLTP Workloads: Channel/HA Scalability

**DS8300 Turbo**  
384 Disks  
73GB/15K RPM





# Missing Interrupt Handler



## Missing Interrupt Handler (MIH)

- **Every second, the z/OS Missing Interrupt Handler (MIH) gets control to determine whether there are any I/O requests that are not completing**
- **For each of these I/O requests, a recovery action will be performed which includes one or more of the following:**
  - Issuing a message
  - Generating a logrec record for diagnostic purposes
  - Terminating the I/O request
  - Re-queuing the I/O request
- **New for zHPF I/O Requests only :**
  - z/OS has the ability to interrogate the status of the I/O request
  - Interrogate information returned by the control unit is surfaced in:
    - MIH status messages
    - MIH logrec records
  - If a device reserve is held, MIH will not terminate or re-queue the I/O request



## MIH Messages

IOS071I 031B,62,\*MASTER\*, START PENDING

**STATUS: DEVICE RESERVED BY ANOTHER SYSTEM**

IOS071I 0980,40,IOSAS, START PENDING

**STATUS: NO I/O OPERATION IS IN PROGRESS**

IOS071I 0410,F2,WHATEVER,START PENDING

**STATUS: I/O WAITING FOR EXTENT CONFLICT**

IOS071I 1029,\*\*,JES3,START PENDING

**STATUS: I/O OPERATION IS EXECUTING**

- **Status in red** appears only for zHPF I/O operations



## MIH Logrec Records

```
DEVICE NUMBER: 08810      REPORT: MIH EDIT      DAY YEAR      JOB IDENTITY: SMS
                        SCP: VS 2 REL. 3      DATE: 019 06
DEVICE TYPE: 3390
                        CPU MODEL: 2084      HH MM SS.TH
CHANNEL PATH ID: 0F      CPU ID: 0632CD      TIME: 16 01 32.78

MISSING INTERRUPT: 10 - START PENDING IN SUBCHANNEL      SUBCHANNEL ID NUMBER: 000169A9
                                                                VOLUME SERIAL: QGNQXB
                                                                UCB LEVEL BYTE: 08
                        HH MM SS.TH
TIME INTERVAL: 00 00 15.00

RECOVERY ACTIONS PERFORMED BYTE: 0C
  HALT OR CLEAR SUBCHANNEL 0
  SIMULATED INTERRUPT 0
  REDRIVE DEVICE 0
  REQUEUE I/O REQUEST 0
  ...

HEX DUMP OF SUBCHANNEL INFORMATION BLOCK
...
COMMAND CODE: E7      I/O DRIVER ID: 02
STATUS: DEVICE RESERVED BY ANOTHER SYSTEM

INTERROGATE INFORMATION:

FORMAT: 01  FLAGS: 40  CU STATE: 00  DEVICE STATE: 00  I/O STATE: 00

STATE DEPENDENT DATA:  xxxxxxxx xxxxxxxx xxxxxxxx
DEVICE LEVEL ID:      xxxxxxxx
DEVICE DEPENDENT DATA: xxxxxxxx xxxxxxxx xxxxxxxx xxxxxxxx
                        xxxxxxxx xxxxxxxx xxxxxxxx
```

- Status and Interrogate Information in red appears only for zHPF I/O operations



# RMF Channel Activity Reports



# DEFINITIONS: 6 new fields in the Channel Activity Report for RMF

## Both FICON and zHPF Operations (Physical Channel)

- **RATE** – The number of native FICON or zHPF operations per second at the physical channel level.
- **ACTIVE** – The average number of native FICON or zHPF operations that are concurrently active. Often referred to as the number of "open exchanges"
- **DEFER** – Number of deferred native FICON or zHPF operations per second. This is the number of operations that could not be initiated by the channel due to lack of available resources.

REQUIRES APAR OA21140: NEW FUNCTION - RMF SUPPORT FOR HIGH PERFORMANCE FICON FOR SYSTEM z (zHPF)



# Example of new FICON and zHPF fields on RMF Channel Activity report

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
2E	FC_S	5	Y	18.57	18.57	2.42	9.49	9.49	3.59	3.59	822.8	1.5	0	1463.7	1.5	0
2F	FC_S	5	Y	18.60	18.60	2.42	9.50	9.50	3.61	3.61	825.0	1.5	0	1463.5	1.5	0
30	FC_S	5	Y	18.59	18.59	2.42	9.50	9.50	3.57	3.57	823.8	1.5	0	1464.2	1.5	0
31	FC_S	5	Y	18.58	18.58	2.42	9.49	9.49	3.59	3.59	823.0	1.5	0	1464.1	1.5	0
32	FC_S	5	Y	18.58	18.58	2.42	9.50	9.50	3.58	3.58	823.4	1.5	0	1463.5	1.5	0
33	FC_S	5	Y	18.58	18.58	2.42	9.48	9.48	3.62	3.62	823.6	1.4	0	1463.2	1.5	0
34	FC_S	5	Y	28.32	28.32	4.28	9.60	9.60	12.39	12.39	1471.0	1.7	0	1479.7	1.5	0
35	FC_S	5	Y	28.33	28.33	4.28	9.61	9.61	12.39	12.39	1469.5	1.7	0	1480.7	1.5	0
36	FC_S	5	Y	28.34	28.34	4.27	9.60	9.60	12.38	12.38	1470.5	1.7	0	1479.6	1.5	0
37	FC_S	5	Y	28.35	28.35	4.28	9.61	9.61	12.37	12.37	1469.5	1.7	0	1481.4	1.5	0
38	FC_S	5	Y	28.35	28.35	4.27	9.60	9.60	12.38	12.38	1470.6	1.7	0	1480.3	1.5	0
39	FC_S	5	Y	28.33	28.33	4.27	9.61	9.61	12.36	12.36	1469.8	1.7	0	1480.5	1.5	0
3A	FC_S	5	Y	28.32	28.32	4.28	9.61	9.61	12.38	12.38	1469.9	1.7	0	1480.4	1.5	0



# Commands and Output



## SETIOS zHPF=YES|NO **Enables/disables zHPF for zOS**

### **SETIOS ZHPF=YES**

```
IOS090I SETIOS. zHPF UPDATE(S) COMPLETE
```

### **SETIOS ZHPF=NO**

```
IOS090I SETIOS. zHPF UPDATE(S) COMPLETE
```

## D IOS,zHPF

**Displays zHPF status for the zOS image**

### **D IOS,ZHPF**

```
RESPONSE=MVS8
```

```
IOS630I 18.23.27 zHPF FACILITY 751
```

```
FICON CHANNEL EXTENSIONS FACILITY IS ENABLED
```

### **D IOS,ZHPF**

```
RESPONSE=MVS8
```

```
IOS630I 18.23.27 zHPF FACILITY 751
```

```
FICON CHANNEL EXTENSIONS FACILITY IS DISABLED
```



## D M=DEV

### Displays zHPF status by device

#### D M=DEV(B020)

```
IEE174I 18.38.39 DISPLAY M 899
DEVICE B020 STATUS=ONLINE
CHP          A0    A1    A2    A3    A4    A5    A6    A7
ENTRY LINK ADDRESS  D0    D3    D1    D5    D2    D6    D4    D7
DEST LINK ADDRESS   88    89    8A    8B    8C    8D    8E    8F
PATH ONLINE        Y     Y     Y     Y     Y     Y     Y     Y
CHP PHYSICALLY ONLINE Y     Y     Y     Y     Y     Y     Y     Y
PATH OPERATIONAL   Y     Y     Y     Y     Y     Y     Y     Y
MANAGED            N     N     N     N     N     N     N     N
CU NUMBER          0A00 0A00 0A00 0A00 0A00 0A00 0A00 0A00
MAXIMUM MANAGED CHPID(S) ALLOWED:  0
DESTINATION CU LOGICAL ADDRESS = 00
SCP CU ND          = 002107.900.IBM.75.000000013941.0030
SCP TOKEN NED      = 002107.900.IBM.75.000000013941.0000
SCP DEVICE NED     = 002107.900.IBM.75.000000013941.0020
HYPERPAV ALIASES CONFIGURED = 64
FUNCTIONS ENABLED = MIDAW, zHPF
```



# D M=CHP

## Displays zHPF status by channel path

### D IOS, ZHPF

IOS630I 18.13.26 zHPF FACILITY 747  
 FICON CHANNEL EXTENSIONS FACILITY IS ENABLED

D M=CHP(A0)

IEE174I 18.20.45 DISPLAY M 749

CHPID A0: TYPE=1B, DESC=FICON SWITCHED, ONLINE

DEVICE STATUS FOR CHANNEL PATH A0

	0	1	2	3	4	5	6	7	8	9	A	B	C	D	E	F
0B00	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+
0B01	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+
0B02	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+
0B03	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+
0B04	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+
0B05	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+

SWITCH DEVICE NUMBER = NONE

ATTACHED ND = 010000.001.MCD.01.0000013M0628

PHYSICAL CHANNEL ID = 0121

FACILITIES SUPPORTED = zHPF

\*\*\*\*\* SYMBOL EXPLANATIONS \*\*\*\*\*

+ ONLINE @ PATH NOT VALIDATED - OFFLINE . DOES NOT EXIST

\* PHYSICALLY ONLINE \$ PATH NOT OPERATIONAL

BX DEVICE IS BOXED SN SUBCHANNEL NOT AVAILABLE

DN DEVICE NOT AVAILABLE PE SUBCHANNEL IN PERMANENT ERROR

AL DEVICE IS AN ALIAS UL DEVICE IS AN UNBOUND ALIAS



# Migration considerations



## z/OS High Performance FICON Analysis Tool

- **To evaluate the impact of zHPF on customer's current workload without having the prereqs for zHPF installed**
  - ATS Storage can help with the study if you provide the SMF data
  - Open a Service Request for this study to be performed by ATS-Storage
    - <http://dalnotes1.sl.dfw.ibm.com/atss/techxpress.nsf/request?OpenForm>
  - No cost service (Free) and takes about 5 days
  - For more information refer to :
  - <http://w3-03.ibm.com/support/techdocs/atmastr.nsf/WebIndex/FLASH10668>



## zHPF - Migration

- No HCD changes required
- **Can run without exploiting this feature if necessary using z/OS PARMLIB option**
  - SETIOS ZHPF=YES|NO on z/OS images that want to Enable / Disable the function
  - Default is ZHPF=NO
- **Recommendation: turn on zHPF, if server pre-req are satisfied.**
  - SETIOS ZHPF=YES



## Hitachi qualification on zHPF as of Dec. 4<sup>th</sup> 2008

- Hitachi Universal Storage Platform V (Hitachi USP V) series products
  - Servers IBM System z10 Enterprise Class (z10 EC™) and IBM System z10 Business Class™ (z10 BC™) at Driver 76D bundle 11
  - Operating System z/OS® V1.7 with the IBM Lifecycle Extension for z/OS V1.7 (5637-A01) with PTFs z/OS V1.8, V1.9, or V1.10 with PTFs
  - Storage Hitachi Universal Storage Platform V (Hitachi USP V) series at DKCMain 60.04.06 and zHPF and FICON protocols (CHPID type FC) on System z10 through supported directors, switches, as well as direct connection



# Summary



# Summary – Key learning points !!

## ■ Overview

- High Performance FICON for System z (zHPF) is a new data transfer protocol that is optionally employed for accessing data from an IBM DS8000 storage and other subsystems. **TCW versus CCW results in more open exchanges**

## ■ zHPF hardware and software requirements

### – Hardware

- IBM System z10 Server - z10 EC and z10 BC at Driver level 76 - **FICON Express2 or Express4**
- DS8000 version 4.1 or Higher at Code Bundle 64.1.16.0 (**feature code 0709 - one time charge**)

### – Software

- Minimum Level: z/OS 1.7 plus prerequisite PTFs
- RMF – Requires APAR OA21140

## ■ Performance Analysis

- z10 zHPF performance white paper Jan 28, 2009 – A simple 4KB read I/O benchmark doubles
- DS8000 Release 4.1 zHPF presentation, Dec. 2008 – Most benefit when cache friendly environment

## ■ Tools and commands

- Missing Interrupt Handler - If a device reserve is held, MIH will not terminate or re-queue the I/O request
- RMF Channel Activity Reports Enhancements – Incorporates both FICON and zHPF activity per channel
- z/OS Commands and Output - **Enabled with SETIOS command** – zHPF display commands

## ■ Migration considerations

- z/OS High Performance FICON Analysis Tool – Service provided by Storage ATS
- **No HCD changes required**



# References



## References

- **IBM Systems Mainframe Servers System z Hardware I/O connectivity - What's New ?**
  - <http://www-03.ibm.com/systems/z/hardware/connectivity/news.html>
- **Techdocs FQ127122 - High Performance FICON (zHPF) FAQ (Server)**
  - <http://w3-03.ibm.com/support/techdocs/atsmastr.nsf/WebIndex/FQ127122>
  - <http://www-03.ibm.com/support/techdocs/atsmastr.nsf/WebIndex/FQ127122>
- **IBM System z10 I/O and High Performance FICON for System z Channel Performance**
  - <ftp://ftp.software.ibm.com/common/ssi/sa/wh/n/zsw03059usen/ZSW03059USEN.PDF>
- **IBM System Storage DS8000 Turbo**
  - <http://www-03.ibm.com/systems/storage/disk/ds8000/index.html>
- **Techdocs FQ126612 – Accumulated questions and answers concerning High Performance FICON (Storage)**
  - <http://w3-03.ibm.com/support/techdocs/atsmastr.nsf/WebIndex/FQ126612>
  - <http://www-03.ibm.com/support/techdocs/atsmastr.nsf/WebIndex/FQ126612>
- **DS8000 and zHPF Teleconference 12/4/08 Presentations**
  - <http://w3-03.ibm.com/support/techdocs/atsmastr.nsf/WebIndex/PRS3441>