

# IBM System z – Hardware Workshop

[www.ibm.com/redbooks](http://www.ibm.com/redbooks)

*zHPF*

*High Performance FICON*



IBM System z10 Workshop

© 2009 IBM Corporation

## Notices

This information was developed for products and services offered in the U.S.A.

IBM may not offer the products, services, or features discussed in this document in other countries. Consult your local IBM representative for information on the products and services currently available in your area. Any reference to an IBM product, program, or service is not intended to state or imply that only that IBM product, program, or service may be used. Any functionally equivalent product, program, or service that does not infringe any IBM intellectual property right may be used instead. However, it is the user's responsibility to evaluate and verify the operation of any non-IBM product, program, or service.

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.

The following paragraph does not apply to the United Kingdom or any other country where such provisions are inconsistent with local law: INTERNATIONAL BUSINESS MACHINES CORPORATION PROVIDES THIS PUBLICATION "AS IS" WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESS OR IMPLIED, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF NON-INFRINGEMENT, MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE. Some states do not allow disclaimer of express or implied warranties in certain transactions, therefore, this statement may not apply to you.

This information could include technical inaccuracies or typographical errors. Changes are periodically made to the information herein; these changes will be incorporated in new editions of the publication. IBM may make improvements and/or changes in the product(s) and/or the program(s) described in this publication at any time without notice.

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.

IBM may use or distribute any of the information you supply in any way it believes appropriate without incurring any obligation to you.

Any performance data contained herein was determined in a controlled environment. Therefore, the results obtained in other operating environments may vary significantly. Some measurements may have been made on development-level systems and there is no guarantee that these measurements will be the same on generally available systems. Furthermore, some measurement may have been estimated through extrapolation. Actual results may vary. Users of this document should verify the applicable data for their specific environment.

Information concerning non-IBM products was obtained from the suppliers of those products, their published announcements or other publicly available sources. IBM has not tested those products and cannot confirm the accuracy of performance, compatibility or any other claims related to non-IBM products. Questions on the capabilities of non-IBM products should be addressed to the suppliers of those products.

This information contains examples of data and reports used in daily business operations. To illustrate them as completely as possible, the examples include the names of individuals, companies, brands, and products. All of these names are fictitious and any similarity to the names and addresses used by an actual business enterprise is entirely coincidental.

### COPYRIGHT LICENSE:

This information contains sample application programs in source language, which illustrate programming techniques on various operating platforms. You may copy, modify, and distribute these sample programs in any form without payment to IBM, for the purposes of developing, using, marketing or distributing application programs conforming to the application programming interface for the operating platform for which the sample programs are written. These examples have not been thoroughly tested under all conditions. IBM, therefore, cannot guarantee or imply reliability, serviceability, or function of these programs.

Note to U.S. Government Users Restricted Rights -- Use, duplication or disclosure restricted by GSA ADP Schedule Contract with IBM Corp.

yourdotcom  
IBM.COM

IBM

## Trademarks


**The following terms are trademarks of the International Business Machines Corporation in the United States, other countries, or both:**

IBM has two registered trademarks for the branding of ITSO publications. These registered marks are for the text word "IBM Redbooks" and the Redbooks logo. In a nutshell, the term Redbooks must always be used in the plural form (for both text and logo) since IBM only owns the registered mark for the plural form. Usage must follow the guidelines below:

**Using the term Redbooks in written text**  
 Redbooks are only to be referred to in the plural form, NEVER in the singular.  
 For the initial reference (first occurrence), you must use "IBM Redbooks®" and include "IBM" as well as the ®. For instances thereafter you may use "Redbooks" without "IBM" preceding the word or ® following it.

**Correct usage for written text:**  
 In this IBM Redbooks® publication we will explore.....(® symbol required for 1st usage)  
 This Redbooks publication will show you.....(2nd usage or later - no ® or "IBM" needed)


**Using the logo:**



Redbooks (logo)

**OTHER ITSO PUBLICATIONS - Marks not yet registered**  
 Trademark registration is a lengthy process and until we are officially registered, we cannot use the ® symbol. For those terms/logos in process, we will be using the ™ symbol. In contrast to the ® symbol (placed in the lower right hand corner), the ™ symbol is placed in the upper right hand corner. Please see examples below:

Redpaper ™  
 Redpapers ™  
 Redwiki ™  
 Redwikis ™



The following terms are trademarks of other companies:  
 Java and all Java-based trademarks are trademarks of Sun Microsystems, Inc. in the United States, other countries, or both.  
 Microsoft, Windows, Windows NT, and the Windows logo are trademarks of Microsoft Corporation in the United States, other countries, or both.  
 Intel, Intel logo, Intel Inside logo, Intel Centrino, Intel Centrino logo, Celeron, Intel Xeon, Intel SpeedStep, Itanium, and Pentium are trademarks or registered trademarks of Intel Corporation or its subsidiaries in the United States and other countries.  
 UNIX is a registered trademark of The Open Group in the United States and other countries.  
 Linux is a trademark of Linus Torvalds in the United States, other countries, or both.  
 Other company, product, or service names may be trademarks or service marks of others.

3

September 14, 2009

© 2009 IBM Corporation

Redbooks Workshops

yourdotcom  
IBM.COM

IBM



## System z10



- Features and Functions Part 2
- zHPF

4

September 14, 2009


© 2009 IBM Corporation



Redbooks Workshops

## z10 High Performance FICON for System z (zHPF)

- **Improve FICON efficiency, scale, value and RAS**
  - As data density behind a CU and device increases
    - scale the I/O rates and bandwidth to grow with the data
  - Reduce elapsed times
    - improved I/O rates and bandwidth
  - Reduce the number of channels, switch ports, control unit ports and optical cables required to balance CPU MIPS with I/O capacity
  - Reduce time and skill it takes to resolve MIH and IFCC conditions
    - Improve First Failure Data Capture
- **Solution**
  - Leverage optimizations in Host/Bus Adapter (HBA) to improve I/O performance, especially for small block I/O
    - Significant improvements in I/O rates (2 to 3x) for small block transfers
  - I/O architecture changes
    - FICON protocol changes between channel and CU (industry standard)
    - New System z architecture to build new channel programs to be exploited by z/OS
    - New ECKD commands for improved efficiency
  - Additional channel and CU diagnostics for timeout conditions


5
September 14, 2009
© 2009 IBM Corporation


## System z High Performance FICON (zHPF)

- FICON architecture for protocol simplification and efficiency designed to:
  - reduce the number of Information Units (IUs) processed
  - reduce FICON channel overhead and improve performance
- The maximum number of I/Os per second is designed to be improved up to 100%\* for small data transfers that can exploit the zHPF protocol
- Requires control unit exploitation
  - IBM System Storage DS8000 – require LIC level 4.1 and FC7092
    - ✓ traditional and HyperPAV supported
  - OEM Support
- z10 FICON Express8, Express4 and Express2 features, CHPID type FC
  - Implemented in System z10 Licensed Internal Code
  - FC channel supports multiple CUs using both FICON and zHPF protocols at the same time
  - zHPF will not be used unless **all** the CHPIDs / CU ports for the device support the new protocol
  - Channel and CU inform each other of capabilities at initialization time
    - ✓ Process Login (PRLI) added to Link Initialization
- Supported by z/OS 1.8 and higher releases with PTF
  - Media manager driver supports data sets
    - ✓ DB2, VSAM, HFS, zFS, PDSE, and extended format sequential data sets

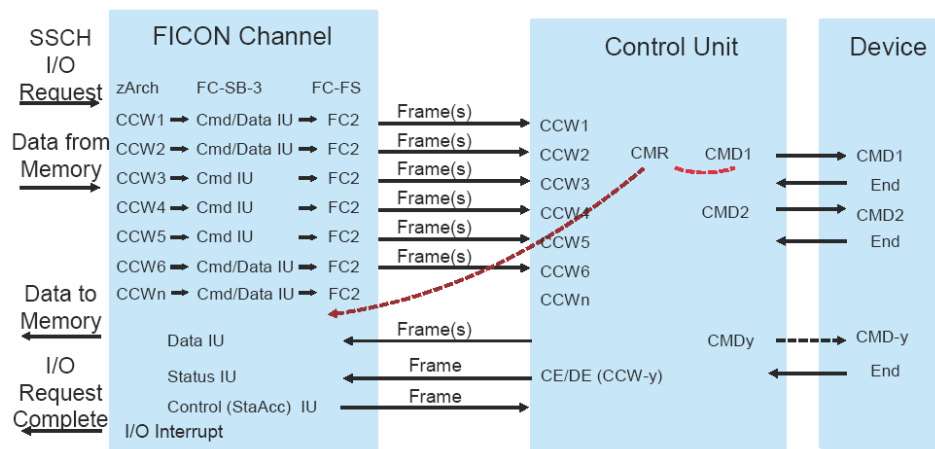
\*Note: Some specific channel programs may not be converted to zHPF protocol

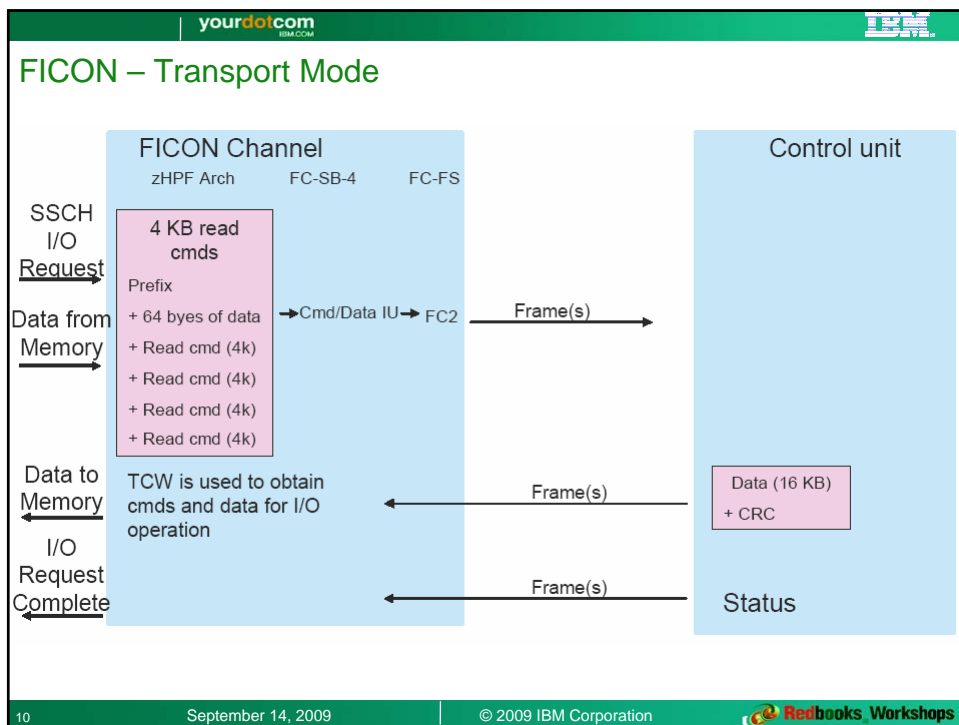
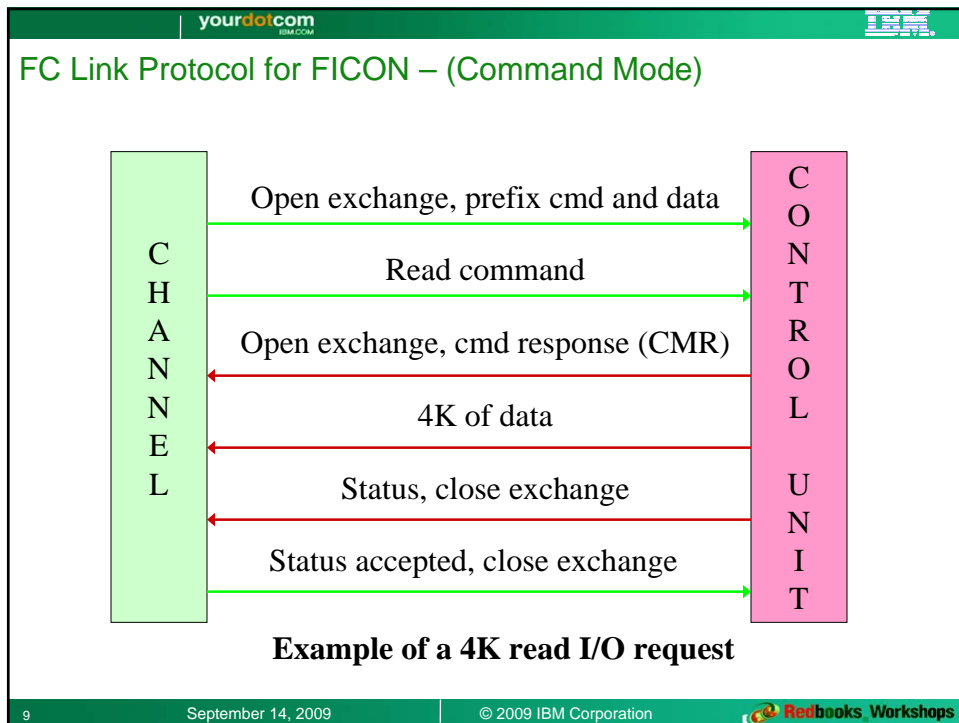
6
September 14, 2009
© 2009 IBM Corporation


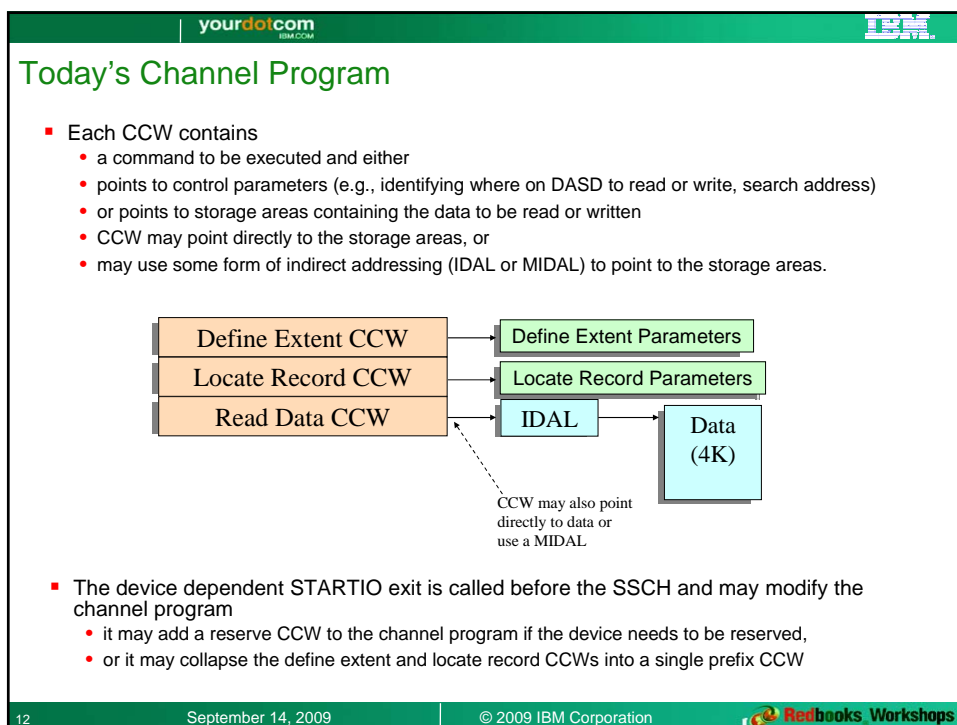
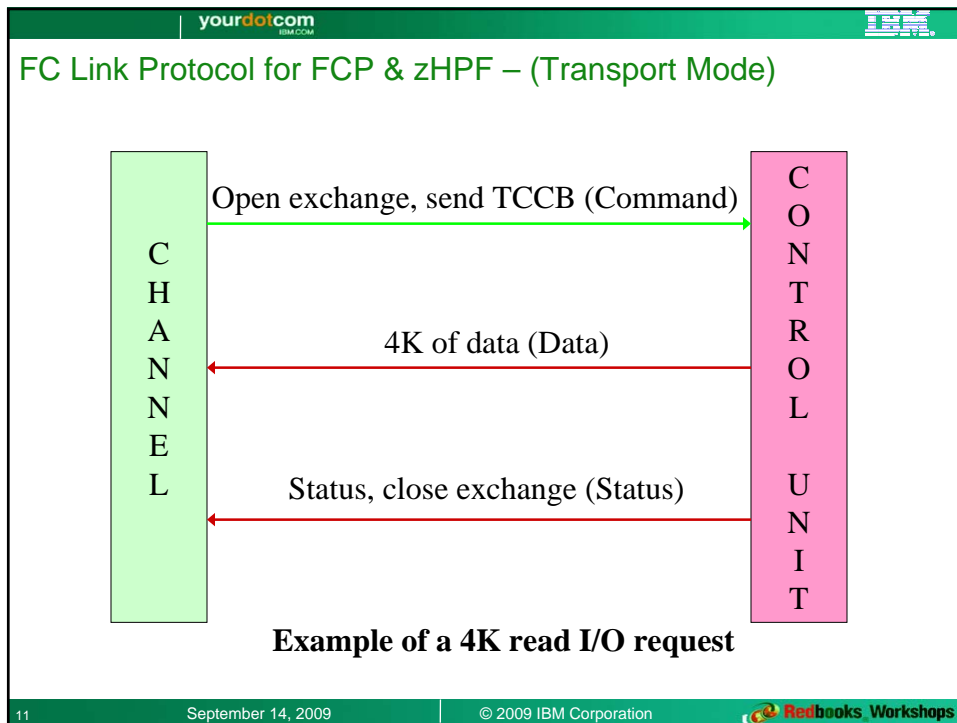
## zHPF vs FICON Architecture

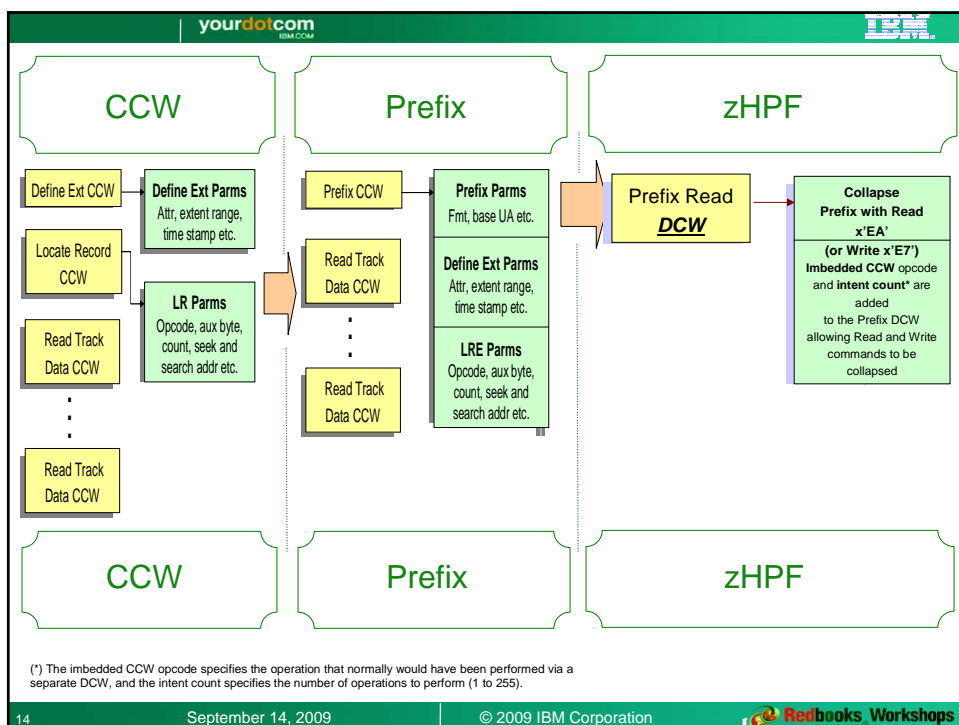
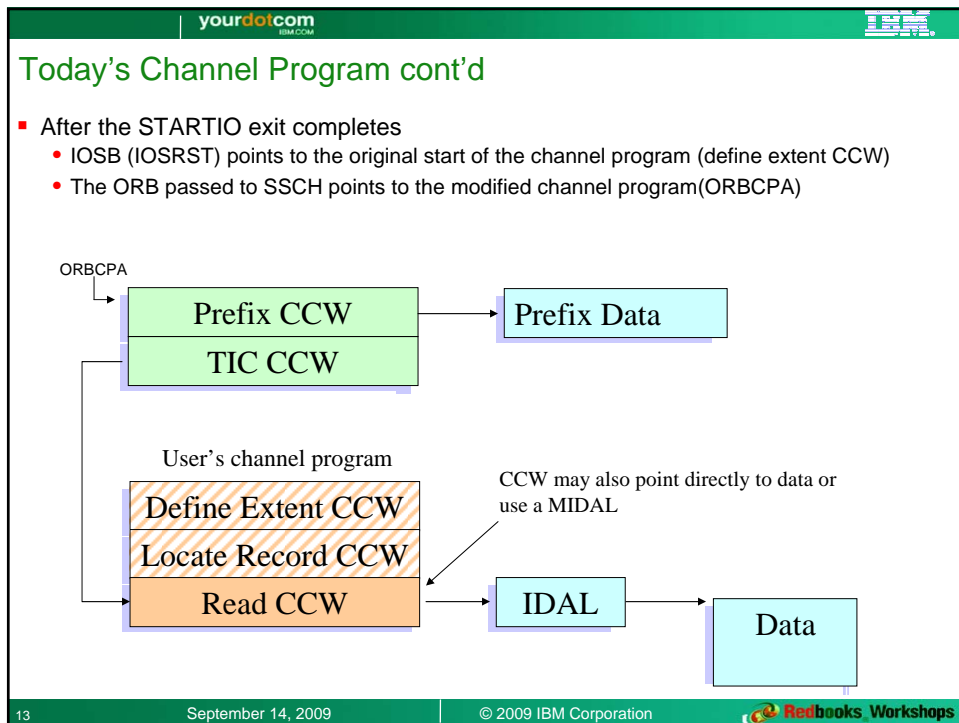
- zHPF architecture moves significant function out of channel firmware
  - Emulex – (the processor chip used by the CU host adapter)
  - Control Unit
  - Software
- Basic transport uses “FCP like”
  - The concept of CCWs “is gone”
  - At the channel there is only “ONE Command”
  - CU commands and control data are encapsulated in one “FCP” command descriptor block
- After some initial setup by firmware, Emulex controls the entire operation
- Data can move only ONE direction in any given Start Subchannel
  - All Reads OR
  - All Writes

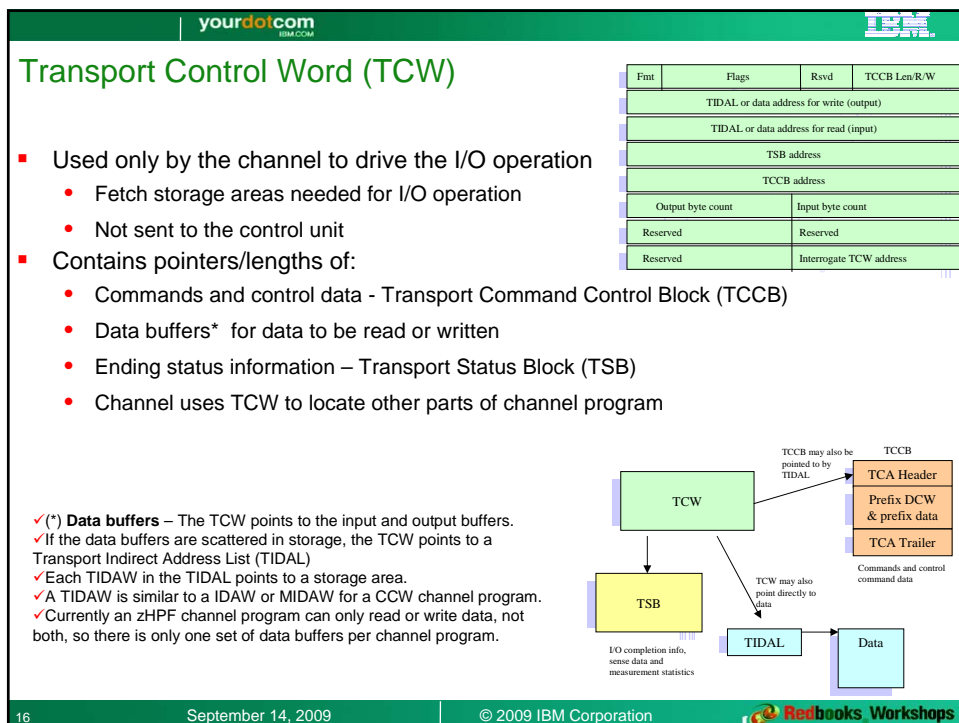
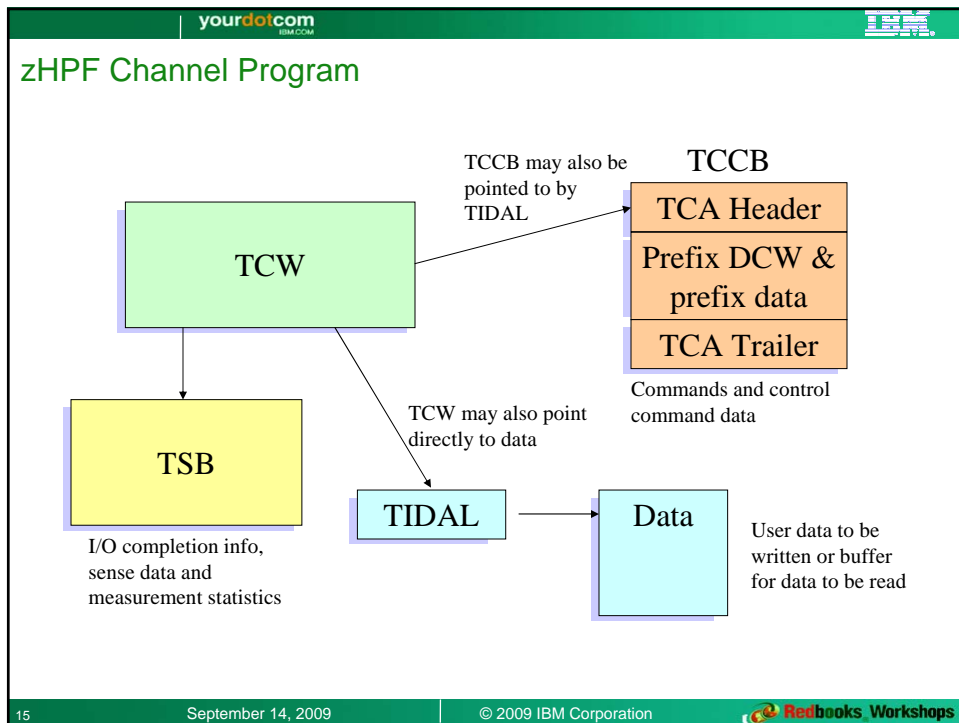
## FICON – Command Mode







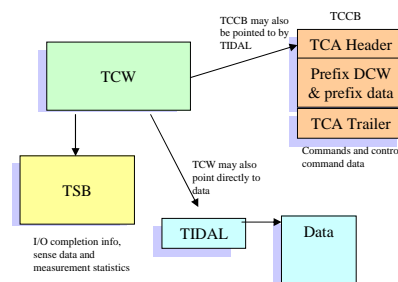




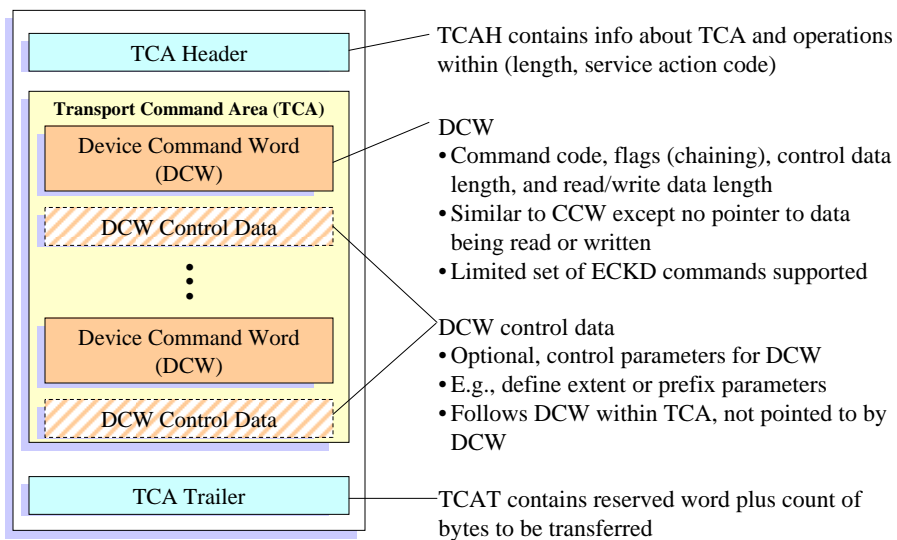


## Transport Command Control Block (TCCB)

- Contains the commands to be executed **by the control unit**
  - Like SCSI Command Descriptor Block (CDB)
- Pointed to by the TCW
- Channel does not look at the TCCB contents – simply packages it and sends it to the CU
  - Allows FCP transport protocols to be used instead of FICON



## Transport Command Control Block (TCCB)



## TSB – Transport Status Block

- Contains additional I/O completion status that is not presented in the IRB
  - 64 bytes in length
  - Stored by channel at I/O or interrogate completion time
  - It also contains response time statistics reported by the control unit
- I/O Status Format – TSB Format
  - Ending status for an I/O operation information not present in IRB
  - I/O measurement statistics used to update CMB/ECMB
  - Sense data for unit checks
- Device Detected Program Check – TSB Format
  - Errors detected by device that directs channel subsystem to present a program check
  - Describes reason for program check
- Interrogate Status – TSB Format
  - Status presented as a result of an interrogate request (MIH)
  - Describes state of I/O operation at time interrogate received

TCBW may also be pointed to by TIDAL

TCB

TCA Header

Prefix DCW & prefix data

TCA Trailer

Commands and control command data

TCBW may also point directly to data

TIDAL

Data

I/O completion info, sense data and measurement statistics

19
September 14, 2009
© 2009 IBM Corporation

## I/O Status TSB

Up to 64 bytes

0	Size	Flags and Fmt	Last DCW offset
1	DCW residual byte count of last DCW		
2	Reserved		
3	Device time		
4	Defer time (Device disconnect time)		
5	Queue time		
6	Device busy time		
7	Device active only time		
8-15	Sense Data (0 to 8 words)		

Replaces residual count in IRB

I/O statistics used to update IRB and CMB/ECMB

No longer need to issue separate sense command

TCBW may also be pointed to by TIDAL

TCB

TCA Header

Prefix DCW & prefix data

TCA Trailer

Commands and control command data

TCBW may also point directly to data

TIDAL

Data

I/O completion info, sense data and measurement statistics

User data to be written or buffer for data to be read

20
September 14, 2009
© 2009 IBM Corporation

## Up to 64 bytes

© 2009 IBM Corporation

Must be 1	Must be 64 bytes
-----------	------------------

© 2009 IBM Corporation

yourdotcom  
IBM.COM
IBM

## Transport Indirect Address Word (TIDAW)

- Used to indirectly address data and commands (TCCB)
  - TCCB TIDAL use is restricted to device support code

**MIDAW**

Reserved	Flags	Count
Data Address		

16 bytes in length,  
quadword aligned, list  
can't cross page boundary

**TIDAW**

Flags	Reserved	Count
Data Address		

\* 2 byte count = 64KB

\*\* 4 byte count = 4 GB

23
September 14, 2009
© 2009 IBM Corporation
Redbooks Workshops

yourdotcom  
IBM.COM
IBM

## Chaining TIDALs Together

- Used when TIDAL exceeds size of page
- Last TIDAW in list has T-TIC flag on and pointer to next TIDAL

```

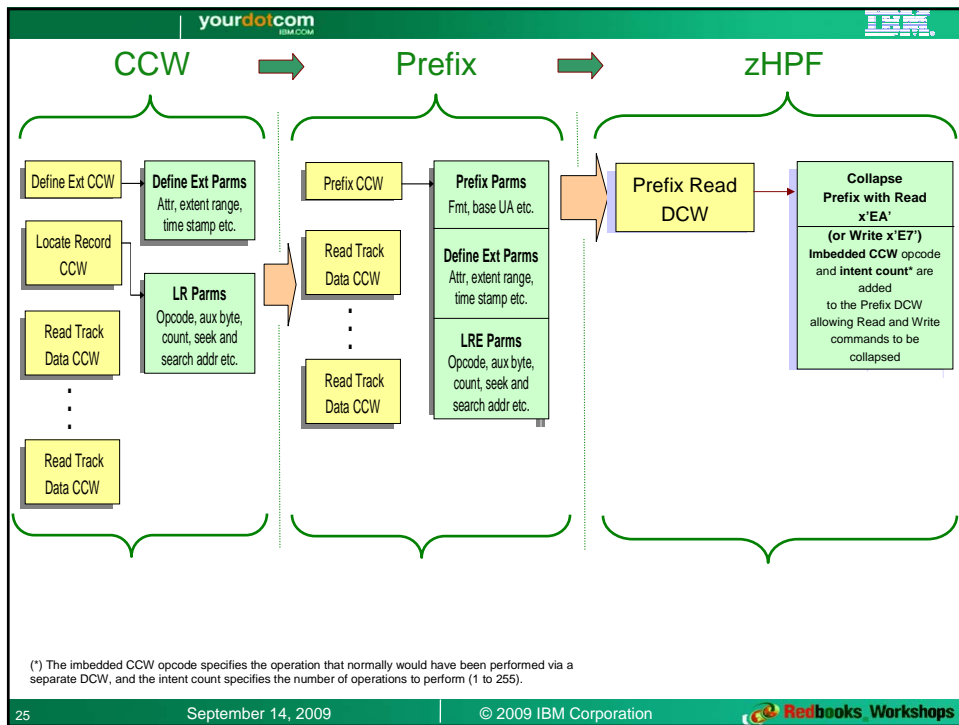
graph TD
    TCW[TCW] --> T1[TIDAW (Data)]
    T1 --> D1[Data]
    T2[TIDAW (Data)] --> D2[Data]
    T3[TIDAW (T-TIC)] --> T4[TIDAW (Data)]
    T4 --> D3[Data]
    T5[TIDAW (Data)] --> D4[Data]
    T6[TIDAW (Data)] --> D5[Data]
    T7[TIDAW (Data)] --> D6[Data]
  
```

24

September 14, 2009

© 2009 IBM Corporation

Redbooks Workshops



yourdotcom IBM.COM

## zHPF Interactions and Dependencies

- Hardware Dependencies
  - Processor
    - z10 EC GA2 / z10 BC
    - FICON channel cards may be carried over to z10, but zHPF is supported only on FICON Express2 and higher cards
  - Control Unit
    - DS8000 - Microcode level 4.1 - LIC feature – FC7092
    - OEM support
  - No HCD, switches or configuration changes required

26 September 14, 2009 © 2009 IBM Corporation Redbooks Workshops

## Interactions and Dependencies

### ■ Software Dependencies

- Supported on z/OS 1.7\*, 1.8\*, and 1.9 and later
- Support is enabled only when all H/W and S/W functions are present and zHPF feature/function is enabled

### ■ Exploiters

- Media manager is the current exploiter
  - DB2, VSAM, HFS, zFS, PDSE, and extended format sequential data sets
- Media manager will retry zHPF I/O operations using CCW channel programs if zHPF I/O fails
- No changes are needed to application programs to exploit zHPF once the function is enabled

\*some z/OS releases require Lifecycle Extension contract

## Customer Control of zHPF

- New IECIOSxx and SETIOS commands to enable and disable zHPF
  - Default: Disabled (1.7 and up)

```
IECIOSxx: ZHPF=YES | NO
SETIOS ZHPF=YES | NO
```

- New display IOS command to display status

```
D IOS,ZHPF
IOS630I hh.mm.ss XXX DATA
HIGH PERFORMANCE FICON FACILITY IS
ENABLED | DISABLED
```

## Determining if a Device Supports zHPF

D M=DEV(410)

```
IEE174I 11.00.11 DISPLAY M 258
DEVICE 0410 STATUS=ONLINE
CHP          A0    A1    A2    A3
DEST LINK ADDRESS A0    A1    A2    A3
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        0400 0400 0400 0400
MAXIMUM MANAGED CHPID(S) ALLOWED: 0
DESTINATION CU LOGICAL ADDRESS = 04
SCP CU ND        = 002107.000.IBM.TC.02069A00FF04.00FF
SCP TOKEN NED     = 002107.000.IBM.TC.02069A00FF04.0400
SCP DEVICE NED    = 002107.000.IBM.TC.02069A00FF04.0410
FUNCTIONS ENABLED = MIDAW, ZHPF
```

## Finding more info about zHPF

- <http://www14.software.ibm.com/webapp/set2/psp/srchBroker>

The screenshot shows the IBM Technical help database search interface. At the top, there's a navigation bar with links: Home, Products, Services & industry solutions, Support & downloads, and My IBM. Below this, a search bar is visible. The main heading is "Technical help database for mainframe Preventive Service Planning buckets". A descriptive paragraph follows, explaining that the database contains installation tips, high impact or pervasive problems, and service recommendations for the zSeries family of servers. Below the text, there's a section titled "Find the bucket by Type, Category and Release:". This section includes three dropdown menus: "Type:" (set to "Function"), "Category:" (set to "zHPF"), and "Release:" (set to "-- ALL RELEASES --"). There are also three radio buttons: "All Content" (selected), "Exclude extract files", and "Extract files only". A "Go" button is at the bottom left of this section.

yourdotcom  
IBM.COM

## DataBase Search Results

- <http://www14.software.ibm.com/webapp/set2/psp/srchBroker>

```

/* Preventive Service Planning */
/* UPGRADE: FUNCTION, SUBSET: zHPF */
/* Updates also available at ftp site: */
/* ftp site = ftp://ftp.software.ibm.com/s390/pspapartool/ */
/* ftp file = FUNCTION_zHPF.txt */
/* Last Extract: */
/* Time = 18:25:31 */
/* Date = 08/25/2009 */
APAR(AA004445) FMID(EER3500) FIX(U000849) UPG(FUNCTION) SUB(zHPF) .
APAR(AA18766) FMID(HBB7720) FIX(UA44235) UPG(FUNCTION) SUB(zHPF) .
APAR(AA19156) FMID(HBB7720) FIX(UA44249) UPG(FUNCTION) SUB(zHPF) .
APAR(AA19956) FMID(HBB7720) FIX(UA44245) UPG(FUNCTION) SUB(zHPF) .
APAR(AA22918) FMID(HBB7720) FIX(UA44122) UPG(FUNCTION) SUB(zHPF) .
APAR(AA18766) FMID(HBB7730) FIX(UA44236) UPG(FUNCTION) SUB(zHPF) .
APAR(AA19156) FMID(HBB7730) FIX(UA44250) UPG(FUNCTION) SUB(zHPF) .
APAR(AA19956) FMID(HBB7730) FIX(UA44246) UPG(FUNCTION) SUB(zHPF) .
APAR(AA22918) FMID(HBB7730) FIX(UA44127) UPG(FUNCTION) SUB(zHPF) .
APAR(AA27852) FMID(HBB7730) FIX(UA46183) UPG(FUNCTION) SUB(zHPF) .
APAR(AA18766) FMID(HBB7740) FIX(UA44237) UPG(FUNCTION) SUB(zHPF) .
APAR(AA19156) FMID(HBB7740) FIX(UA44251) UPG(FUNCTION) SUB(zHPF) .
APAR(AA19956) FMID(HBB7740) FIX(UA44247) UPG(FUNCTION) SUB(zHPF) .

APAR(AA21140) FMID(HRM7750) FIX(UA90470) UPG(FUNCTION) SUB(zHPF) .
APAR(AA22918) FMID(JBB772S) FIX(UA44129) UPG(FUNCTION) SUB(zHPF) .
APAR(AA21140) FMID(JRM773J) FIX(UA90475) UPG(FUNCTION) SUB(zHPF) .
APAR(AA21140) FMID(JRM774J) FIX(UA90473) UPG(FUNCTION) SUB(zHPF) .
APAR(AA21140) FMID(JRM775J) FIX(UA90471) UPG(FUNCTION) SUB(zHPF) .
/* end of PSP extract file */

```

31
September 14, 2009
© 2009 IBM Corporation

yourdotcom  
IBM.COM

## FICON and zHPF Comparison

Function	FICON	zHPF
Devices supported	DASD, tape, CTC, printers etc.	DASD
Channels supported	All	FICON (2.0+)
I/O drivers supported	All	Media manager
Parallel Access Volumes (includes HyperPAV)	Supported	Supported
Channel Measurement Block statistics	Supported	Supported
Dynamic Pathing <ul style="list-style-type: none"> <li>▪ Alternate path reconnect</li> <li>▪ Reserve allegiance</li> </ul>	<ul style="list-style-type: none"> <li>▪ Yes</li> <li>▪ Yes</li> </ul>	<ul style="list-style-type: none"> <li>▪ No (no physical disconnect)</li> <li>▪ Yes</li> </ul>

32
September 14, 2009
© 2009 IBM Corporation



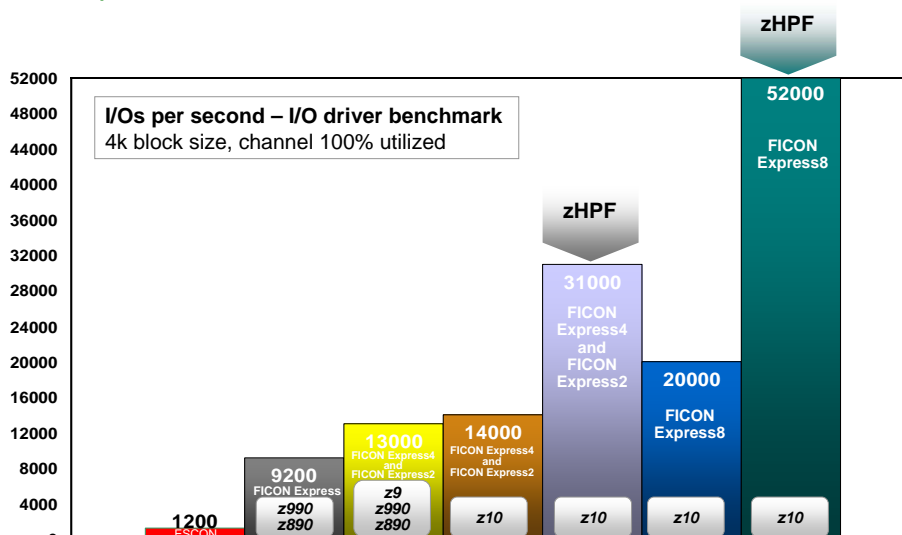
## When is a Device Enabled for zHPF ?

A device is enabled for zHPF if **all** of the following are true:

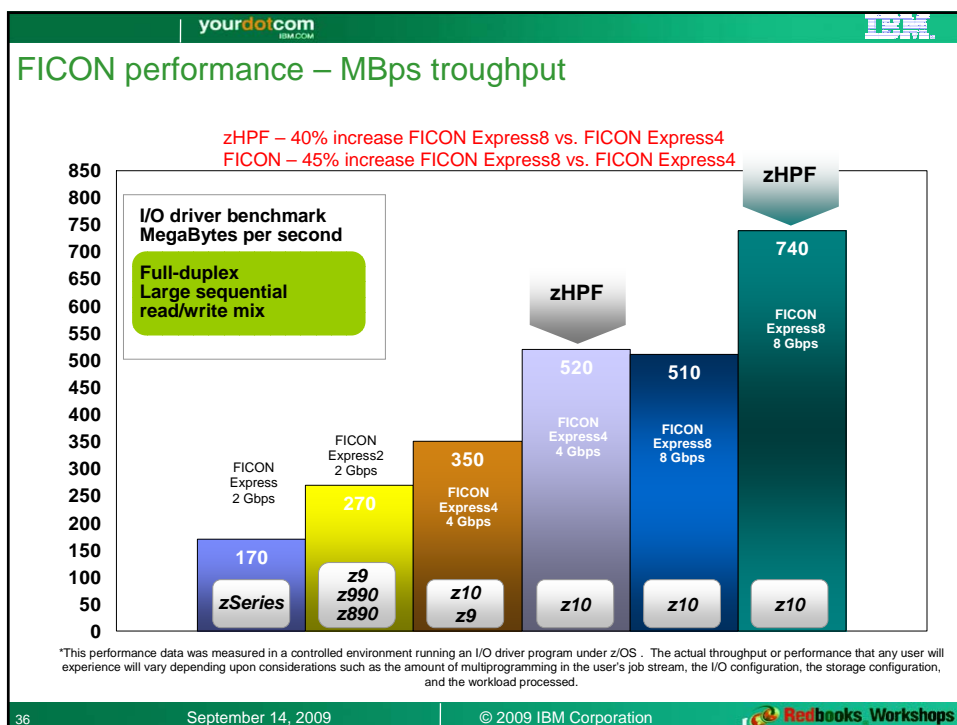

1. The processor supports zHPF
2. All of the online channels support zHPF
  - CONFIG CHP(xx) and VARY PATH may affect zHPF status
  - Loss of path may also affect zHPF status
  - No message issued when zHPF status changes
3. The control unit supports zHPF
  - Indicated via Read Feature Codes
4. The device support code supports zHPF
  - UIM indicates this – Only DASD for now
5. zHPF has not been disabled by the customer
  - zHPF=YES specified in IECIOSxx or SETIOS command

✓ New UCB bit (**UCBFCX**) is set if device supports FCX

## FICON performance – start I/Os



\*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.



yourdotcom  
IBM.COM
IBM

## Non zHPF GTF CCW trace

- SSCH CCW trace

I/O Driver id: EXCP=02,  
Media manager=03

Prefix CCW

CCW Addr

Read Track  
Data CCWs

Application I/O has IOS  
level and UCB level of  
01. Anything higher is  
IOS recovery I/O.

TIC CCW

Prefix parameters

No data shown in this area  
because this is a read command.  
(Data buffers are empty)

```

SSCH.... 0B42B  ASCB.... 00FCD600 CPUID... 0006  JOBN.... D5PDBM1 RST..... 0FF71AD0 VST..... 02CCAAD0
CC..... 00  SEEKA... 00000000 22000A00  GPMSK... 00  OPT..... 14
DVRID... 03  IOSLVL... 01  UCBVLVL... 01  UCBWGT... 00  BASE.... 0B42B
ORB..... 024171A8 03C2FF01 6202FAE8 0700FC00 00000000 00000000 00000000 00000000
GMT-07/09/2008 16:53:28.706838 LOC-07/09/2008 12:53:28.706838

CCW CHAIN  FORMAT 1  SSCH  DEV.... 0B42B  ASCB.... 00FCD600 CPU....
JOBN.... D5PDBM1
0F9DCAE8 E7600044 6202FAA4 01C0000C 00000000 00000000 40C40000  | .....D.. |
0000004C 0022000A 0022000D C2AAAE9  | .....B.WZ |
*** Back half of split data ***
80FC0000 00000000 0C800004 0022000A  | ..... |
0022000A 09964080 00000000 00000000  | .....O..... |

0F9DCAE0 08000000 0FF71AE0 02CCAAE0 A6401080 0FF71B00 02CCAAE8 A640C180 0FF71B80 02CCAAE0 A640C180 0FF71D00 02CCAAE8 A6004080 0FF71E80 02CCAAE0 A6004080 0FF71E80
IOSB 02CCA0F0

C0004800 0303015A 8131D950 057F14C0 024171A8 00000000 00000000 02CCA20C
00000000 00000000 00000000 00000000 00000000 8131B6F0 81445F38 8131E48
7B86FA00 01000000 08000000 00000000 00000000 00000000 22000A00 C9D6E2C2
00000000 00000162 C8000000 02CCA18C 00000000 00000000 00000000 00000000
  
```

37
September 14, 2009
© 2009 IBM Corporation
Redbooks Workshops

yourdotcom  
IBM.COM
IBM

## Non zHPF GTF CCW trace cont'd

- I/O Interrupt Trace

Define Extent  
CCW

MIDAW Data length  
(bytes 2-3)

64-bit Data address

Locate Record CCW

1st Read Track Data CCW (bytes = x'4080')

2nd Read Track Data CCW (bytes = x'C180')

Data length (again)

...other MIDAWs for this CCW...

```

IO..... 0B4FF  ASCB.... 00FCD600 CPUID... 0002  JOBN.... D5PDBM1 PSW..... 070C1000 B6800BCC
IRB..... 00C04007 55C89170 0C000000 00200008 00000000 TCB..... 007FF078 SENSE... N/A
OPT..... 14  DVRID... 03  IOSLVL... 01  UCBVLVL... 01  UCBWGT... 08
GMT-07/09/2008 16:53:28.707332 LOC-07/09/2008 12:53:28.707332

CCW CHAIN  FORMAT 1  IO  DEV.... 0B4FF  ASCB.... 00FCD600 CPU
JOBN.... D5PDBM1
1278C140 63400010 55C89570 40C40000 00000044 002C000A 002C000D  | .....D..... |
1278C148 47400010 55C89580 0C800004 002C000A 002C000A 09964080  | .....O..... |
1278C150 A6401080 55C89170 00000000 00000000 00000000 00000000  | ..... |
MIDAW 00000000 00001000 1000 00000000 00000000 00000000 00000000  | ..... |
00000001_18465000 00000000 00000000 00000000 00000000 00000000 00000000  | ..... |
*** Back half of split data ***
MIDAW 00000000 00000020 0020 00000000 00000000 00000000 00000000  | ..... |
00000000_55C895C0 00000000 00000000 00000000 00000000 00000000 00000000  | ..... |
...other MIDAWs for this CCW...
MIDAW 00000000 00800020 0020 00000000 00000000 00002000 00000000  | ..... |
00000000_55C89650 00000000 00000000 00000000 00000000 005A5AA5 00000000  | ..... |
.....

1278C158 A640C180 55C891F0 00000000 00000000 00000000 00000000  | ..... |
MIDAW 00000000 00001000 1000 00000000 00000000 00000000 00000000  | ..... |
00000001_18E27000 00000000 00000000 00000000 00000000 00000000 00000000  | ..... |
*** Back half of split data ***
MIDAW 00000000 00000020 0020 00000000 00000000 00000000 00000000  | ..... |
00000000_55C89680 00000000 00000000 00000000 00000000 005A5AA5 00000000  | ..... |
...other MIDAWs for this CCW...
  
```

38
September 14, 2009
© 2009 IBM Corporation
Redbooks Workshops

yourdotcom  
IBM.COM

## zHPF GTF CCW trace (SSCH)

```

SSCH.... 0B449   ASCB.... 00FCD600 CPUID... 0006   JOBN.... D5PDBM1 RST..... 7E28C680 VST..... 02CE6680
CC..... 00   SEEKA... 00000001 21000100   GPMSE... 00   OPT..... 14
DVRID... 03   IOSLVL... 01   UCBLVL... 01   UCWGT... 00   BASE.... 0B449
ORB..... 024184C8 0004FF01 6202FCC0 0700FC00 00000000 00000000 00000000 00000000
GMT-07/09/2008 16:53:28.713689 LOC-07/09/2008 12:53:28.713689

FORMAT 0          SSCH          TCW CHAIN          DEV..... 0B449   ASCB.... 00FCD600 CPU..... 0006
JOBN.... D5PDBM1

TCW at 0F9DCCC0 (6202FCC0)
Format..... 00          Flag1..... 06
Flag2..... 00          Flag3..... 00
TCCBL/R/W..... 4E
Output Address... 00000000 00000000
Input Address... 00000000 7E28C6D0
TSB Address..... 00000001 41373910
TCCB Address..... 00000001 413739C0
Output Count.... 00000000          Input Count..... 00001020
Interrogate TCW.. 00000000

TCCB TIDAW at 00000001_413739C0
Flags... 00 Count... 00000058 Addr... 00000001 41373950

TCA Header at 00000001_41373950
Format..... 7F          TCALen..... 54
Serv Act Code... F901          Priority..... FC

DCW at 00000001_41373960
Command.. EA Flags.. 00 CD Count.. 40 Count.. 00001020

DCW Control Data at 00000001_41373968
01C00000 00000000 00000000 40C01020 0000004E 01210001 01210001 C2AAA6E9 | ..... {.....*.....B
6D56486 2000000C 00000000 0C8A601 01210001 01210001 09001020 00000000 | .....Yw.....

TCCB TIDAW at 00000001_413739D0
Flags... 80 Count... 00000008 Addr... 00000000 7E28C778

TCA Trailer at 00000000_7E28C778
Transport Count... 00001024

```

**Real and virtual addresses of TCW**

**Total input bytes for channel program (TCCB)**

**TCCB TIDAW for part of channel pgm modified by device code**

**Prefix read**

**Control data (prefix parameter) length**

**Total data bytes transferred by prefix read**

**Prefix parameters**

**TCCB TIDAW for part of channel pgm not modified by device code**

39

September 14, 2009

© 2009 IBM Corporation

yourdotcom  
IBM.COM

## zHPF GTF CCW trace (I/O interrupt)

```

IO..... 0B449   ASCB.... 00FCD600 CPUID... 0003   JOBN.... D5PDBM1 PSW..... 071C6000 B6AEE826
IRB..... 00104007 6202FCC0 0C000100 00800000 00000000 TCB..... 007FF078 SENSE... N/A
OPT..... 14   DVRID... 03   IOSLVL... 01   UCBLVL... 01   UCWGT... 00
GMT-07/09/2008 16:53:28.714713 LOC-07/09/2008 12:53:28.714713

FORMAT 0          TCW CHAIN          IO          DEV..... 0B449   ASCB.... 00FCD600 CPU..... 0003
JOBN.... D5PDBM1

TCW at 02CE6680 (7E28C680)
Format..... 00          Flag1..... 04
Flag2..... 00          Flag3..... 00
TCCBL/R/W..... 4E
Output Address... 00000000 00000000
Input Address... 00000000 7E28C6D0
TSB Address..... 00000000 00000000
TCCB Address..... 00000000 7E28C720
Output Count.... 00000000          Input Count..... 00001020
Interrogate TCW.. 00000000

I/O Status TSB at 00000001_41373910
Length..... 20          Flags..... D1
DCW Offset..... 0000          Count..... 00000000
TotalDevTime... 000000248          DeferTime..... 00000000
CUQueueTime... 00000000          DevBusyTime... 00000000
DevActOnlyTime 00000000          00000000 00000000 00000000 00000000
Sense Data..... 00000000 00000000 00000000 00000000

TCA Header at 00000000_7E28C720
Format..... 7F          TCALen..... 54
Serv Act Code... F901          Priority..... 00

DCW at 00000000_7E28C730
Command.. EA Flags.. 00 CD Count.. 40 Count.. 00001020

DCW Control Data at 00000000_7E28C738
01800000 00000000 00000000 40C01020 0000004E 01210001 01210001 00000000 | ..... {.....
00000000 2000000C 00000000 0C8A601 01210001 01210001 09001020 00000000 | .....Yw.....

```

Size	Flags and Fmt	Last DCW offset
DCW residual byte count of last DCW executed		
Reserved		
Device time		
Defer time (Device disconnect time)		
Queue time		
Device busy time		
Device active only time		
Sense Data (0 to 8 words)		

**Read/Write indicator – low order two bits (Read=B'10', Write=B'01')**

**Residual count – should be zero if successful completion**

40

September 14, 2009

© 2009 IBM Corporation

yourdotcom  
IBM.COM
IBM

## zHPF GTF CCW trace (I/O interrupt) cont'd

First input buffer TIDAW (4K buffer)

```

Data TIDAW at 00000000_7E28C6D0
Flags... 00 Count... 00001000 Addr... 00000001 52652000
Data at 00000001_52652000
10000000 00000000 0000147C 00000000 00000000 002E4810 00570194 0DBC0000 | .....@.....m..
*** Back half of split data ***
024A0222 01FF01E1 01C801B4 01A5015F 011E00E7 00BA0092 006F0051 003800D5 | .ç.....H...v.^...X...k.?...
Data TIDAW at 00000000_7E28C6E0
Flags... 80 Count... 00000020 Addr... 00000000 7E28C798
Data at 00000000_7E28C798
10000000 00000000 00141000 00000000 00000000 00000000 00000000 005A5AA5 | .....!

TCA Trailer at 00000000_7E28C778
Transport Count... 00001024

IOSB at 02CCD0F0
FLA..... C0          FLB..... 00          FLC..... 40          PROC..... 00          DVRID.... 03          ...
ASID..... 015A        PGAD..... 8131D950    PKEY..... 05          COD..... 7F          OPT..... 14          ...
UCB..... 02417CE8     CCWAD..... 00000000    DSTAT.... 00          SSTAT.... 00          CSWRC.... 0000      ...
USB..... 02CCD020     IOPID..... 00000000    SCHC..... 0000      SMS..... 0000      IPIB..... 00000000  ...
ERP..... 00000000     PCI..... 00000000    NRM..... 8131B6F0    ABN..... 81445F38    DIE..... 81317E48  ...
VST..... 02CCD700     DSID..... 7B5D3C00    LEVEL.... 01          GPMSK.... 00          DCTI..... 0000      ...
CKEY..... 00          MDB..... 00          MDM..... 00          CTC..... 00000000    SKM..... 00          ...
SKCC..... 0011        SKH1..... 00          SKH2..... 03          SKR..... 00          XID..... IOSB      ...
XFLG1.... 00          XFLG2.... 00          XSSXA.... 00000000    XI0BE.... 02CCD820    XRCOD.... 00          ...
XASPR.... 015F        XFLG3.... C4          XIOD..... 02CCD18C    XMSC..... 00000000    00000000      ...
XRSVF.... 00000000    00000000

IOBE at 02CCD820
ID..... IOBE        VERS..... 01          FLG1..... 00          FLG2..... 00          ERPM..... 00          ...
IEDB..... 00000000    FLG3..... 00          SIOC..... 00          TIME..... 00          RESCOUNT. 00000000 ...
DDPC_RC.. 00          DDPC_RCQ.. 00

```

Amount of data formatted controlled by GTF CCW parameter

Second input buffer TIDAW (32 byte suffix)

41
September 14, 2009
© 2009 IBM Corporation
Redbooks Workshops

yourdotcom  
IBM.COM
IBM

# System z10

## ■ Questions ?

- **Luiz A. Fadel**  
— [fadel@br.ibm.com](mailto:fadel@br.ibm.com)
- **Ewerson Palacio**  
— [bird@br.ibm.com](mailto:bird@br.ibm.com)

42
September 14, 2009
© 2009 IBM Corporation
Redbooks Workshops

## Disclaimer

**THE INFORMATION CONTAINED IN THIS PRESENTATION IS PROVIDED FOR INFORMATIONAL PURPOSES ONLY. WHILE EFFORTS WERE MADE TO VERIFY THE COMPLETENESS AND ACCURACY OF THE INFORMATION CONTAINED IN THIS PRESENTATION, IT IS PROVIDED "AS IS" WITHOUT WARRANTY OF ANY KIND, EXPRESS OR IMPLIED. IN ADDITION, THIS INFORMATION IS BASED ON IBM'S CURRENT PRODUCT PLANS AND STRATEGY, WHICH ARE SUBJECT TO CHANGE BY IBM WITHOUT NOTICE. IBM SHALL NOT BE RESPONSIBLE FOR ANY DAMAGES ARISING OUT OF THE USE OF, OR OTHERWISE RELATED TO, THIS PRESENTATION OR ANY OTHER DOCUMENTATION. NOTHING CONTAINED IN THIS PRESENTATION IS INTENDED TO, NOR SHALL HAVE THE EFFECT OF, CREATING ANY WARRANTIES OR REPRESENTATIONS FROM IBM (OR ITS SUPPLIERS OR LICENSORS), OR ALTERING THE TERMS AND CONDITIONS OF ANY AGREEMENT OR LICENSE GOVERNING THE USE OF IBM PRODUCTS AND/OR SOFTWARE.**

IBM, the IBM logo, ibm.com, and DB2 for z/OS are trademarks or registered trademarks of International Business Machines Corporation in the United States, other countries, or both. If these and other IBM trademarked terms are marked on their first occurrence in this information with a trademark symbol (® or ™), these symbols indicate U.S. registered or common law trademarks owned by IBM at the time this information was published. Such trademarks may also be registered or common law trademarks in other countries. A current list of IBM trademarks is available on the Web at "Copyright and trademark information" at [www.ibm.com/legal/copytrade.shtml](http://www.ibm.com/legal/copytrade.shtml)

© Copyright IBM Corporation 2009. All rights reserved.  
U.S. Government Users Restricted Rights  
Use, duplication or disclosure restricted by GSA ADP Schedule Contract with IBM Corp.

System z10

Backup slides

yourdotcom  
IBM.COM
IBM

## System z FICON Enhancements over time

- **FICON Link Incident Reporting**
  - Designed to allow an OS image (without operator intervention) to register for any link incident reports  
This improves the ability to capture data for link incident analysis
  - z/OS will display this information on consoles and will also save in the system log and in LOGREC
  - Introduced with the z9 EC and requires z/OS 1.7 or higher
- **FICON open exchanges increased from 32 to 64**
  - The number of open exchanges was increased from 32 to 64 to improve the FICON throughput and to align with the longer FICON paths
- **Channel Link Recovery Thresholding (Flapping Links)**
  - A channel subsystem function designed to threshold link failure events when repeated failures that are subsequently recovered do not disrupt customer's workload
  - Requires z/OS 1.4.1 plus PTFs or later
- **RNID (Remote Note IDentifier)**
  - Enables the CU node id's on the downstream side of switches to be displayed at the SE and at the operating systems console using the D M=DEV command
  - Requires z/OS 1.4.1 plus PTFs or later.
- **MIDAW (Modified InDirect Address Word)**
  - The Modified Indirect Data Address Word (MIDAW) facility is a system architecture and software exploitation designed to improve FICON performance
  - Designed to reduce channel, director and control unit overhead by reducing the number of CCW's and frames that have to be processed

45
September 14, 2009
© 2009 IBM Corporation
Redbooks Workshops

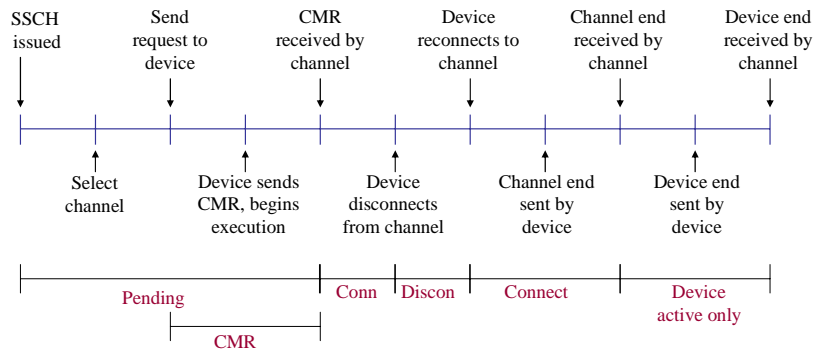
yourdotcom  
IBM.COM
IBM

## Other z10 FICON/FCP Enhancements

- **CTC Logical Paths**
  - CTC now "remembers" logical path information
  - Saved locally in the channel when it changes
  - Restored after a CCC
  - Eliminates lost logical path headaches after a CCC
- **RMF enhancements**
  - Channel now provides information about the number of exchanges in use and the number of SSCHs that could not be executed due to an "out of resource" condition
- **Engine utilization changes**
  - Out of resource now shows up as engine busy
  - Helps make SAD more accurate with respect to RMF and SAP
- **Name Server registration**
  - FICON will register with the Switch "Name Server"
  - The FICON channel now provides the same information to the fabric as is commonly provided by open systems, registering with the "Name Server" in the attached FICON directors
    - Enables quick and efficient management of storage area network (SAN) and perform problem determination and analysis
  - Platform registration is a standard service defined in the Fibre Channel - Generic Services 3 (FC-GS-3) standard (INCITS (ANSI) T11.3 group)
    - It allows a platform (storage subsystem, host, etc.) to register information about itself with the fabric (directors)
    - Applicable to all FICON Express4, FICON Express2, and FICON Express features - CHPID type FC

46
September 14, 2009
© 2009 IBM Corporation
Redbooks Workshops

## Channel Measurements



- All times shown above measured by channel
- Device may accumulate CU queuing and defer time during connect time
- CSS subtracts defer time from connect, adds to disconnect
- Operating system subtracts CU queuing time from connect, adds to pending
- Device busy (not shown) is also measured by channel when CU returns device busy to the initial selection sequence. It is part of pending time.

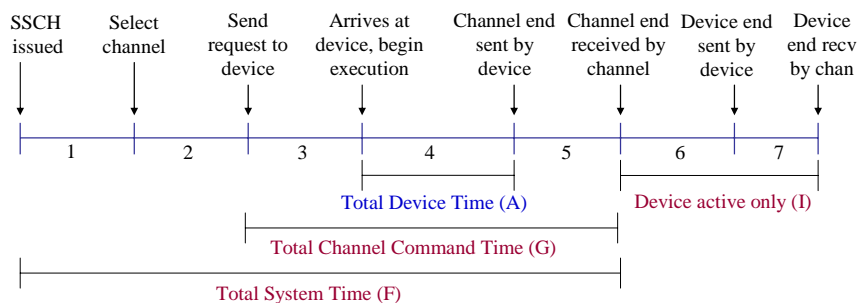
47

September 14, 2009

© 2009 IBM Corporation

Redbooks Workshops

## Channel Measurements cont'd



- CU will now collect 5 times: total device (A), CU defer (B), CU queuing (C), device busy (D), device active only (E). The last 4 times are a subset of the total device time.
- Connect = A - B - D - E (total device minus all CU times except CU queuing); time slot 4
- Pending = F - A (that portion of total system time that is not connect time); time slots 1, 2, 3, 5
- Disconnect = B (CU defer time)
- Device active only = E or I (time slots 6, 7)
- Device busy = D + what ever device busy was accumulated by the channel
- Initial command response = G - A (portion of channel time to traverse the link - both ways); time slots 3, 5

48

September 14, 2009

© 2009 IBM Corporation

Redbooks Workshops