More Data, Less Chatter : Improving Performance on z/OS with IBM zHPF

Holden O'Neal

Session 17576

August 14, 2015

Overview

1. What problem does zHPF solve
2. Overview of zHPF and History
3. How zHPF is used by zOS
4. SAS internal test results
5. How do you get zHPF

Lets plan Lunch!

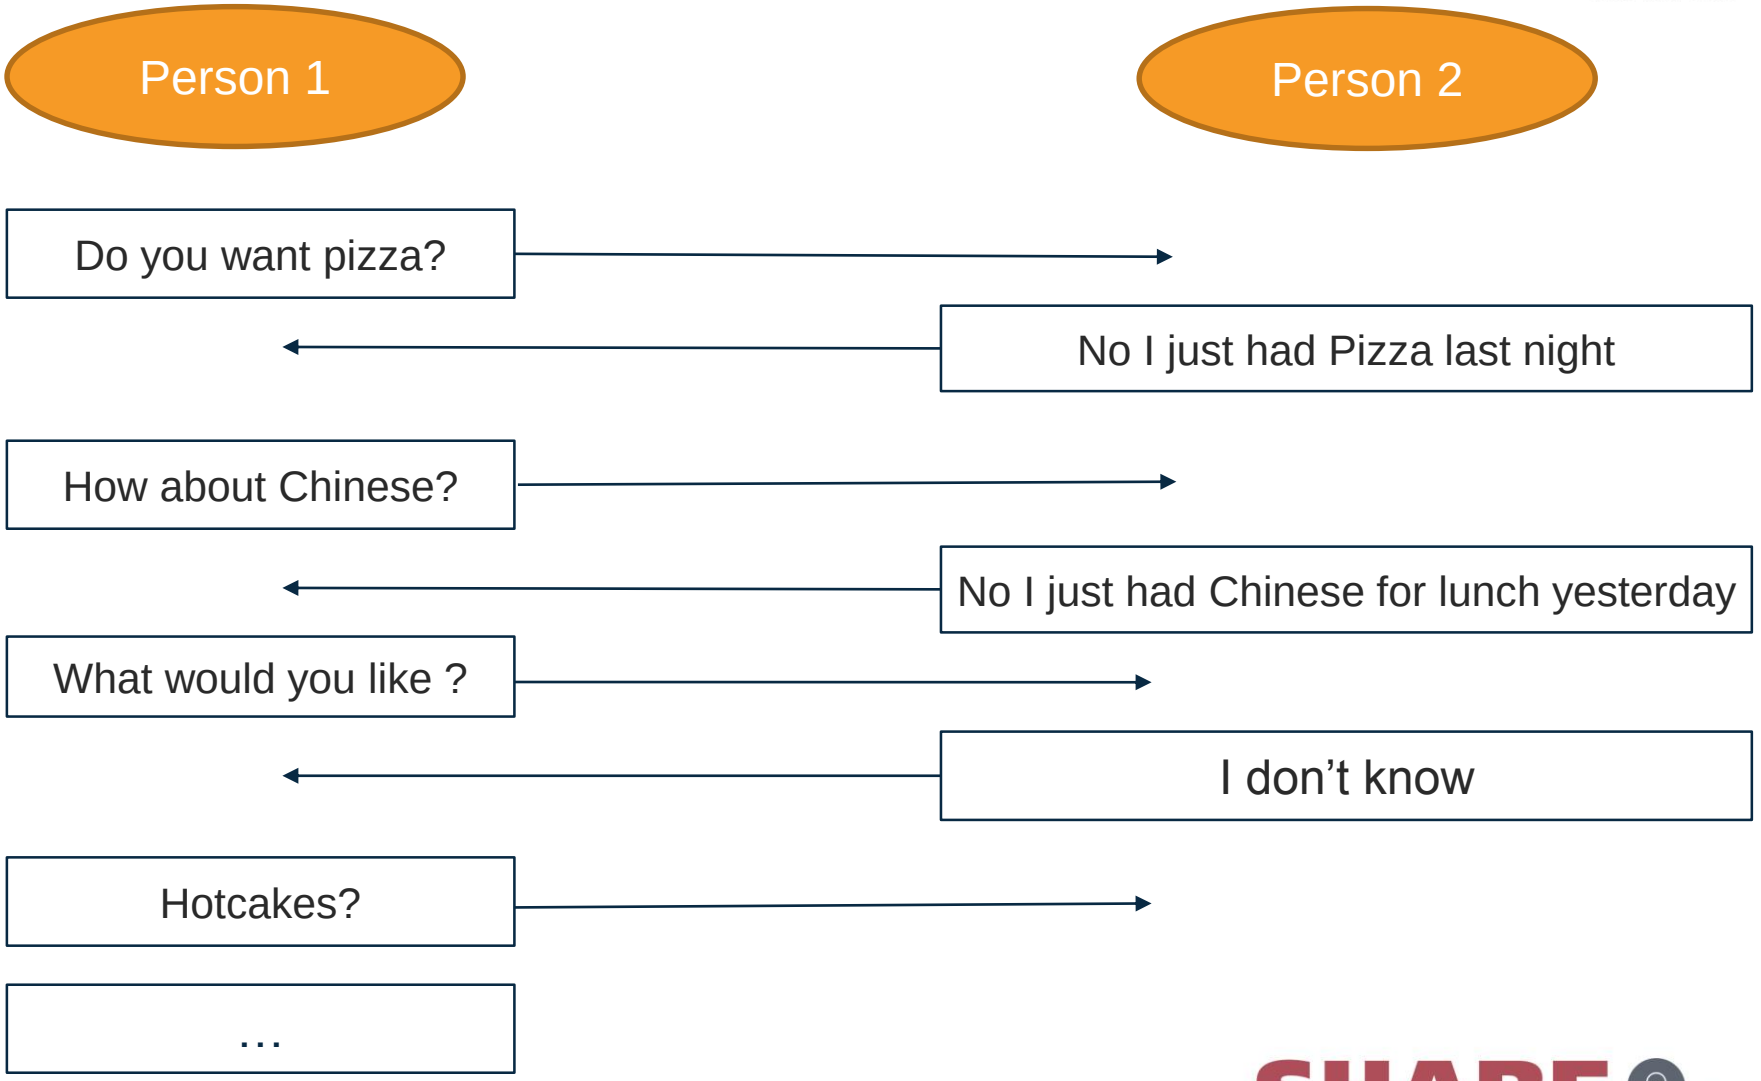

Lets try that again…

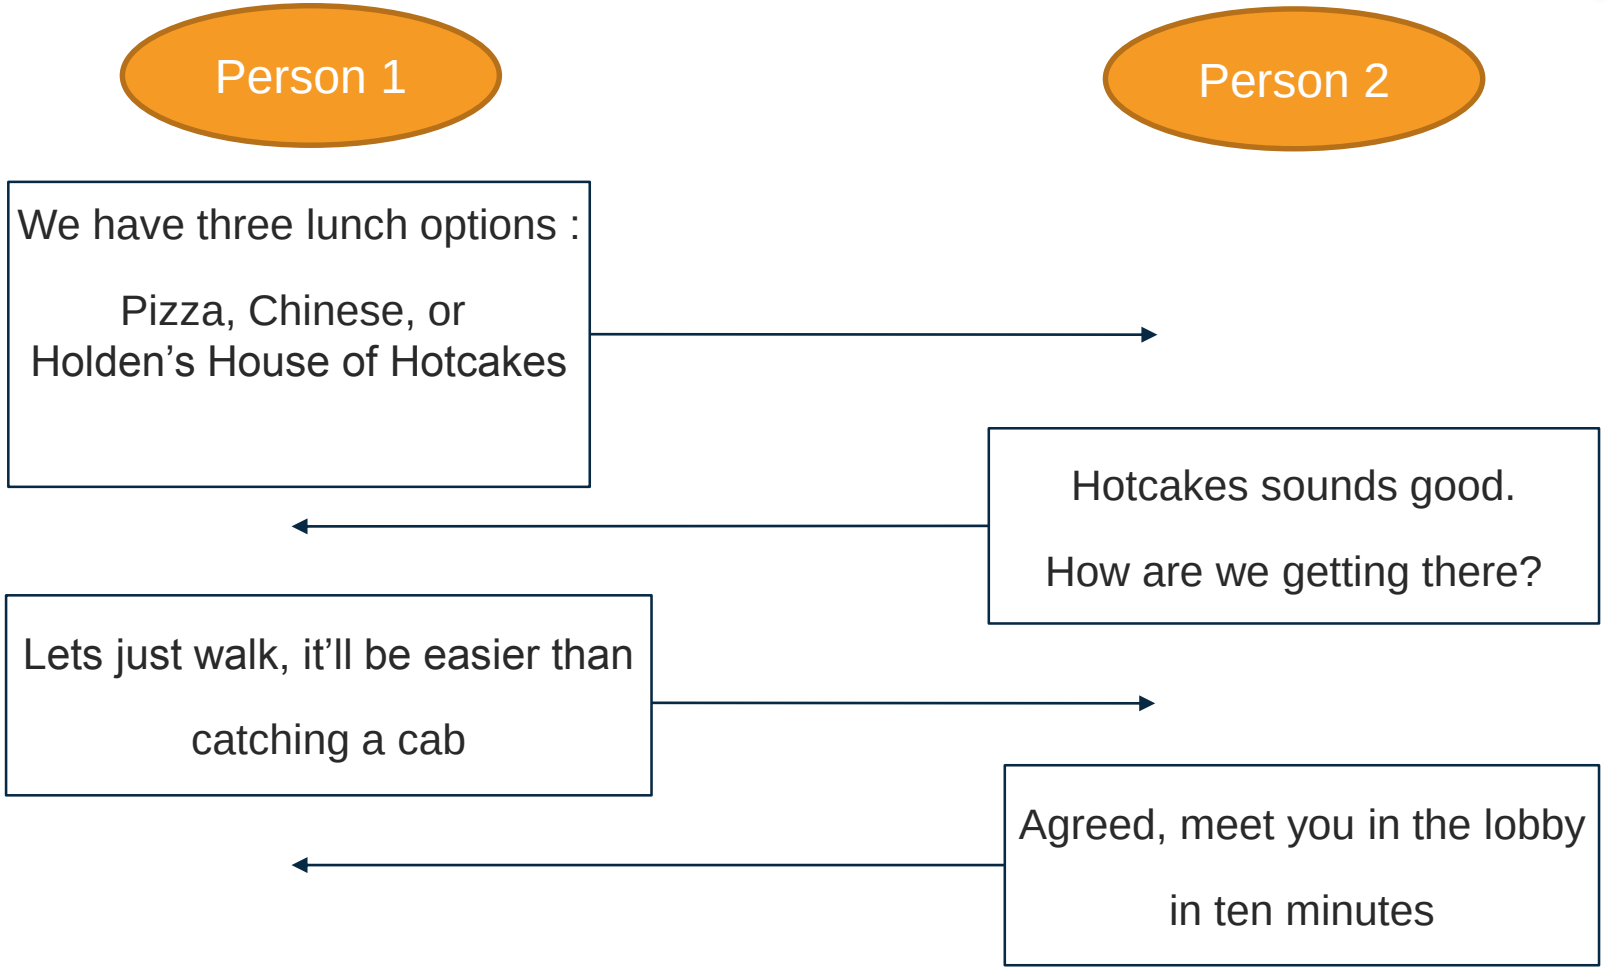

What is zHPF?

- Hiper Performance Ficon for SystemZ
- zHPF is an extension to FICON Architecture
- Improve Small Block I/O

History of zHPF and FICON

- Oct 2008 : Initial announcement of System z10 with FICON Express4 and FICON Express2 channels
  - Single track limit for zHPF data transfers
  - Max of 31k zHPF IOs/sec
- July 2009 : FICON Express8 channels on System z10
  - 64k byte limit for zHPF data transfers
  - Max of 52k zHPF IOs/sec
- July 2010 : Additional Support on zEnterprise 196
  - Extension to multi-tracks of zHPF data transfers
- July 2011 : Introduced FICON Express8S Channel
  - Hardware data router introduced
  - 100% DB2 I/O converted to zHPF

FICON and zHPF performance on z Systems

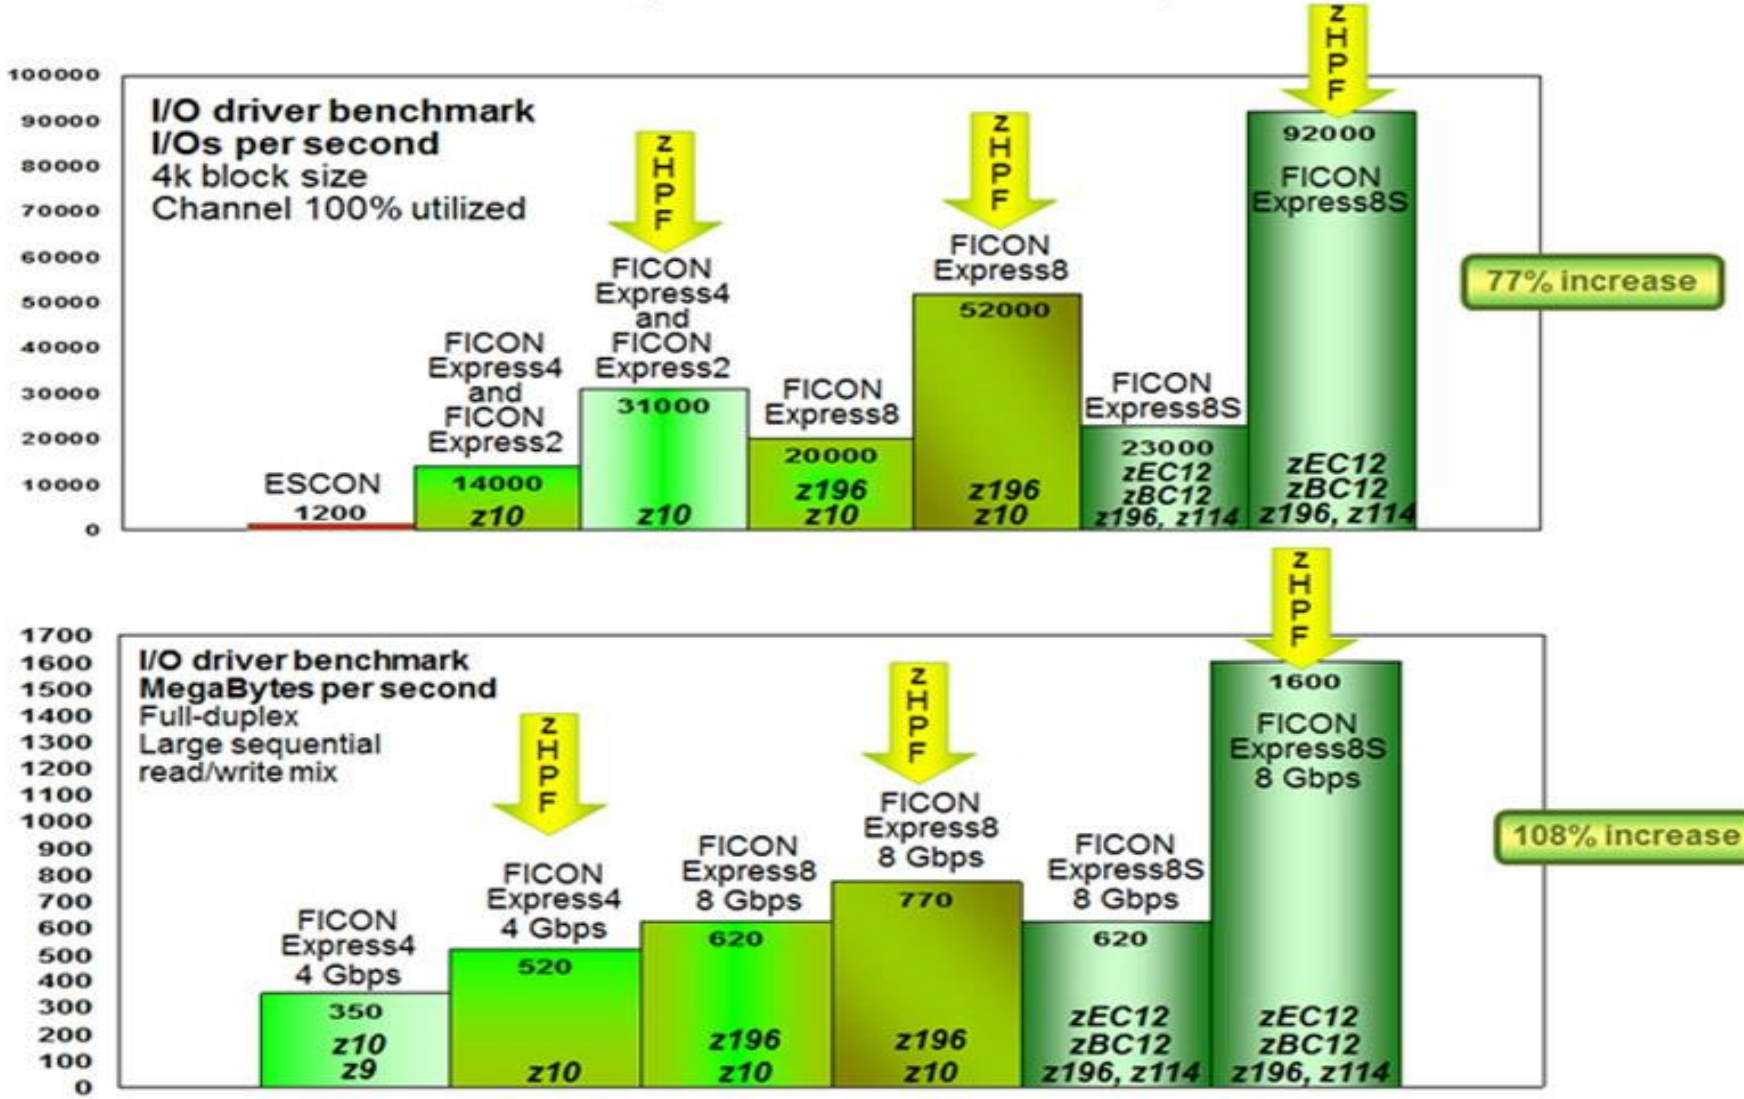

Photo and tests credited to IBM

ESCON

- Original System/360 disk drive controllers lacked on-board processors
- Supports only CCW
- Must wait for Channel End/Device End
- Jobs monopolize channel
- More channel paths required

FICON

- Support previous CCW & zHPF
- CCWs can be transferred without waiting for command response
- Multiple I/O operations concurrently on a channel
- Shares the pipe
- zHPF reduces connection time to I/O device

Questions?

Channel Architecture Overview

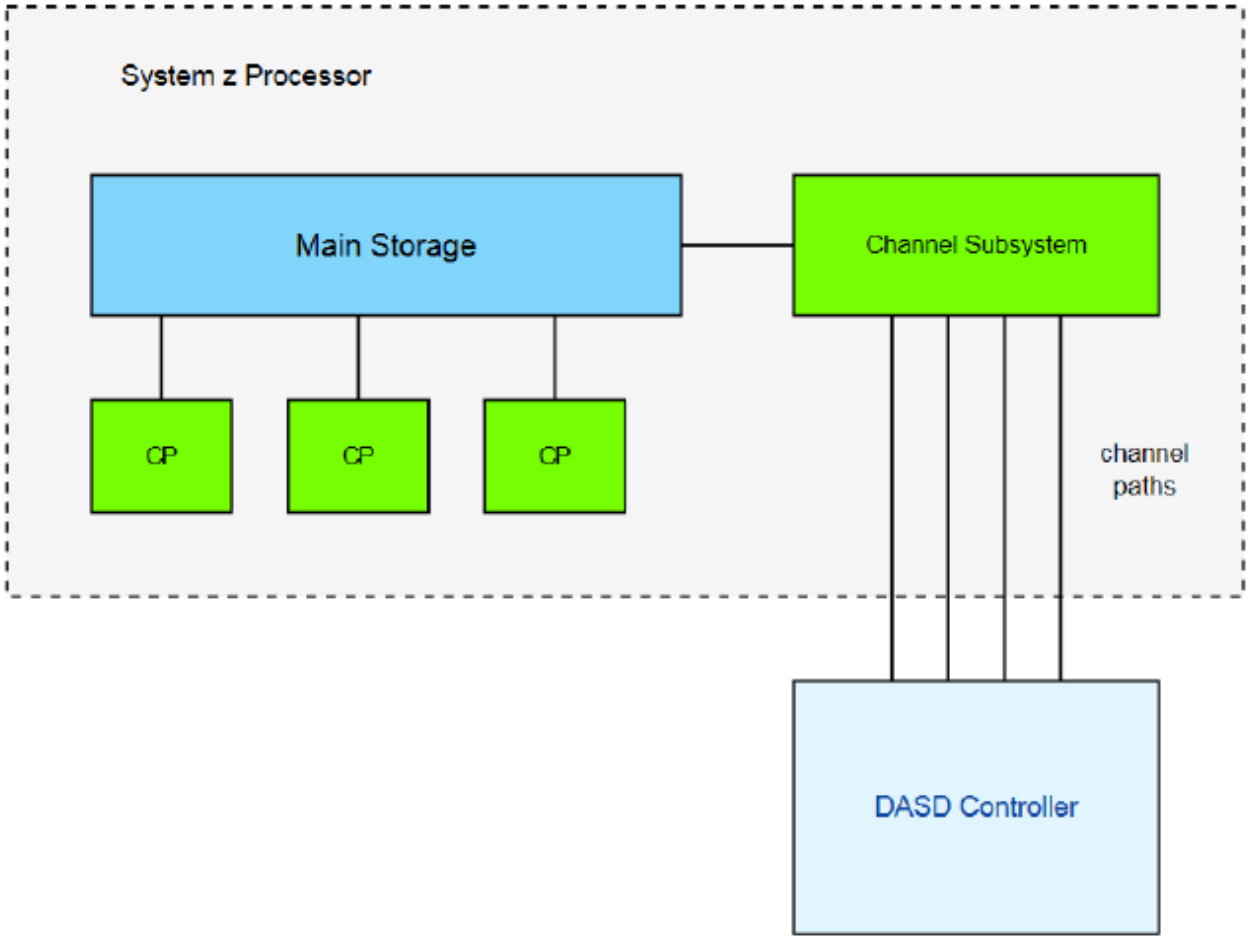

Figure 2 - Relationship of Channel Subsystem to Other Processor Components

Channel Command Word Program (CCP)

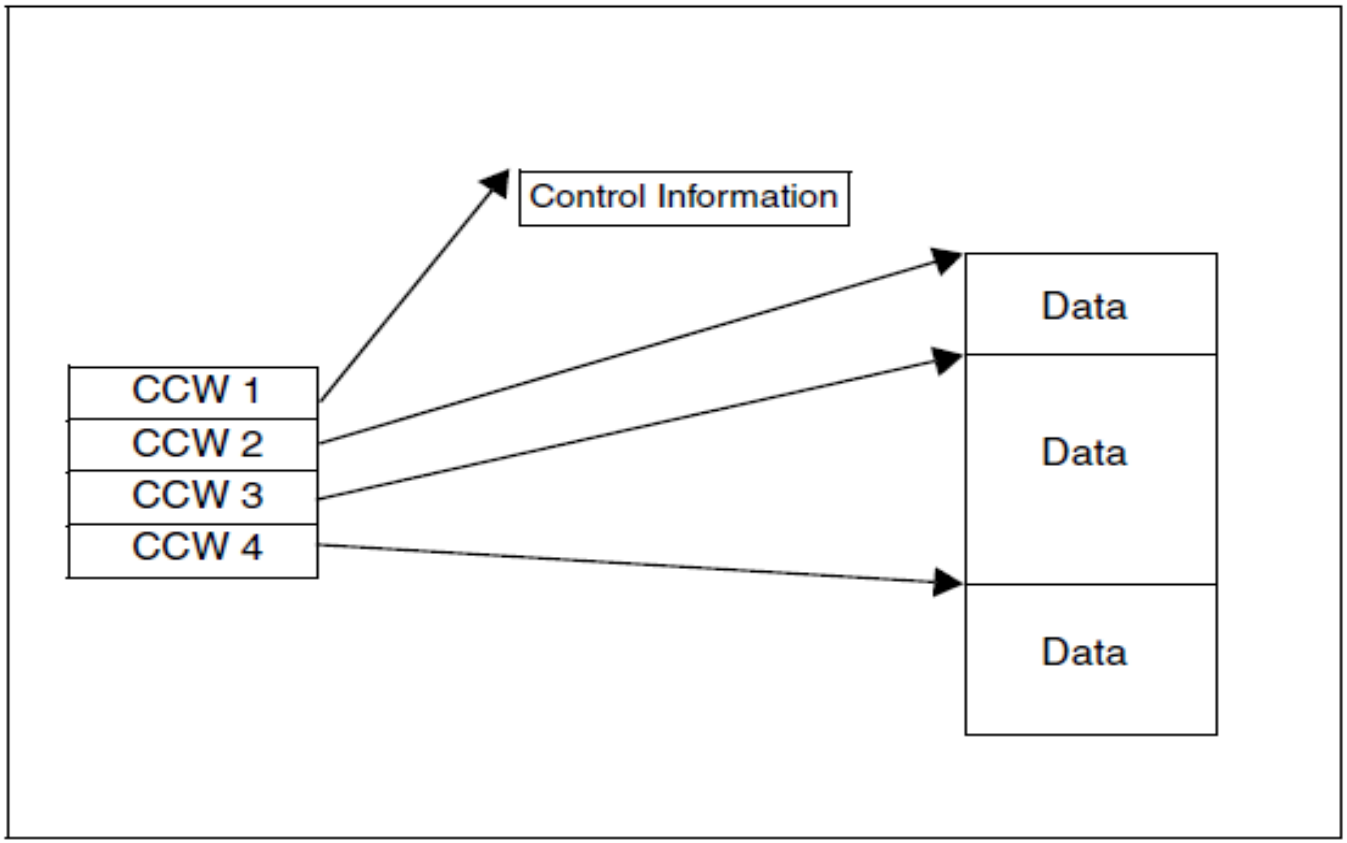

Figure 1. CCW Channel Program Example

Photo credited to IBM

Transport Command Word(TCW)

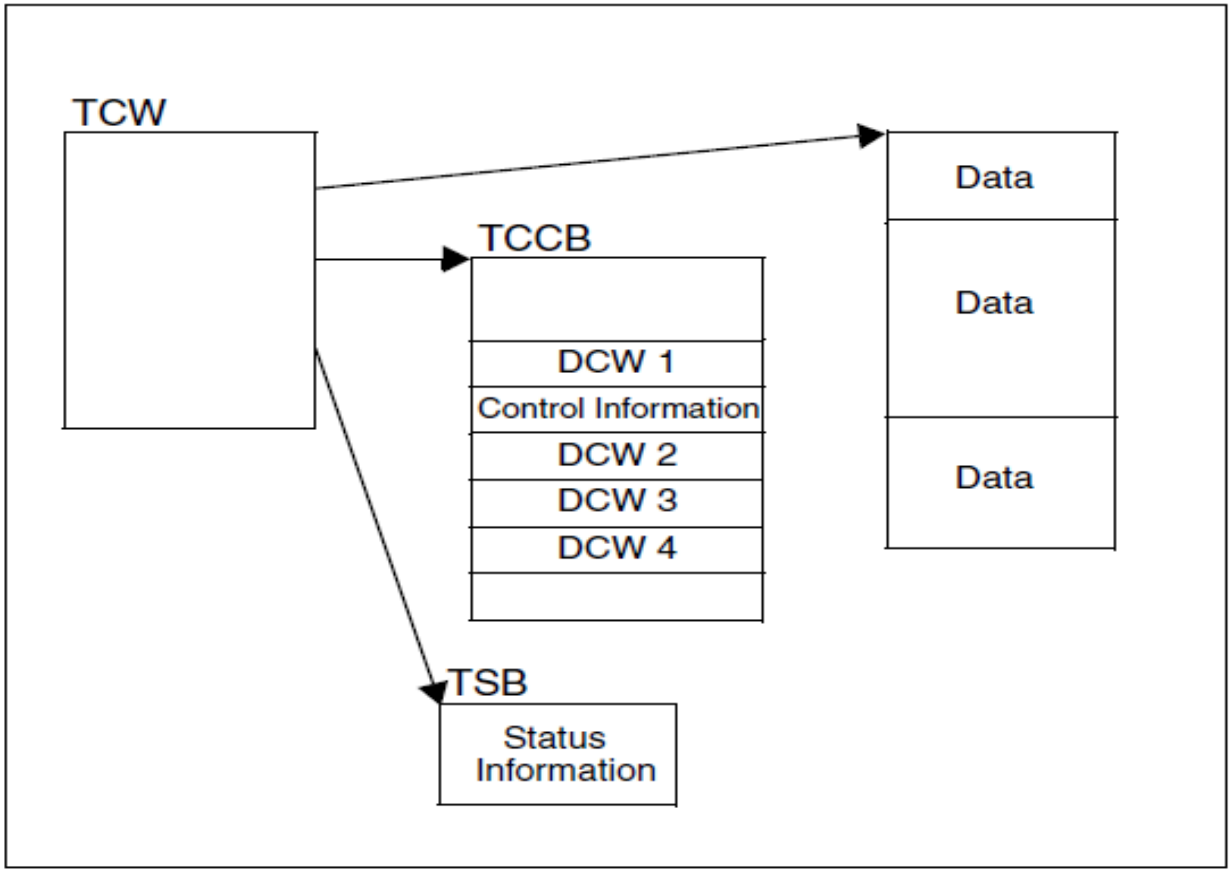

Figure 2. Example of TCW Channel Program Example Designating I/O

Photo credited to IBM

Questions?

zHPF vs FICON

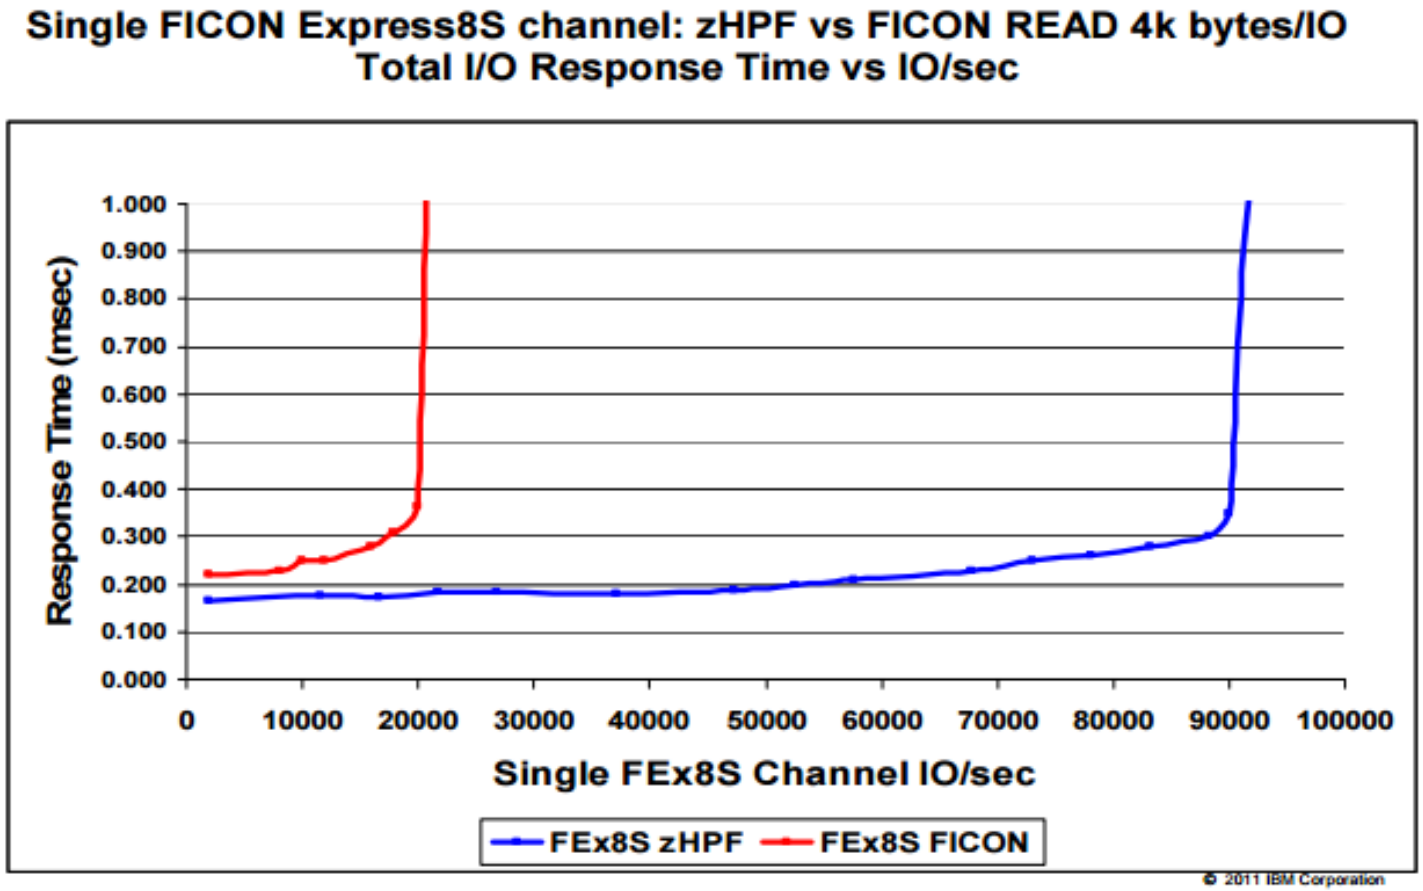

Figure 2

Photo credited to IBM

zHPF and CCP Mixtures

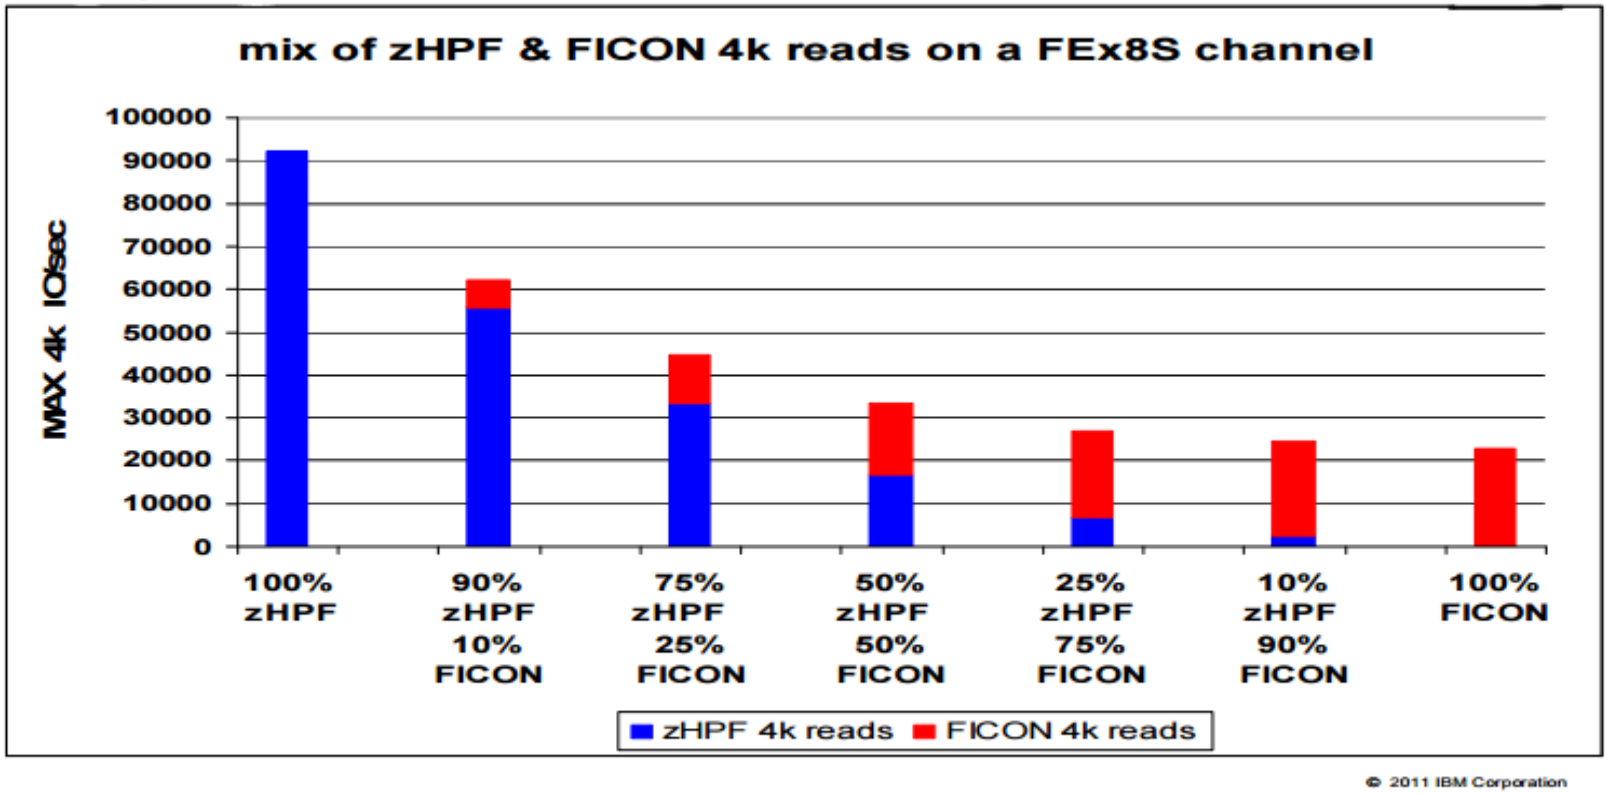

Photo credited to IBM

DB2 10

- Disorganized Indexes
- Sorted Row ID (RID) Lists
  - No more “death by table scan”
- RUNSTATS
  - Page sampling reduces CPU time and uses less buffers
  - Page sampling becomes faster than Row Sampling

QSAM Benefits

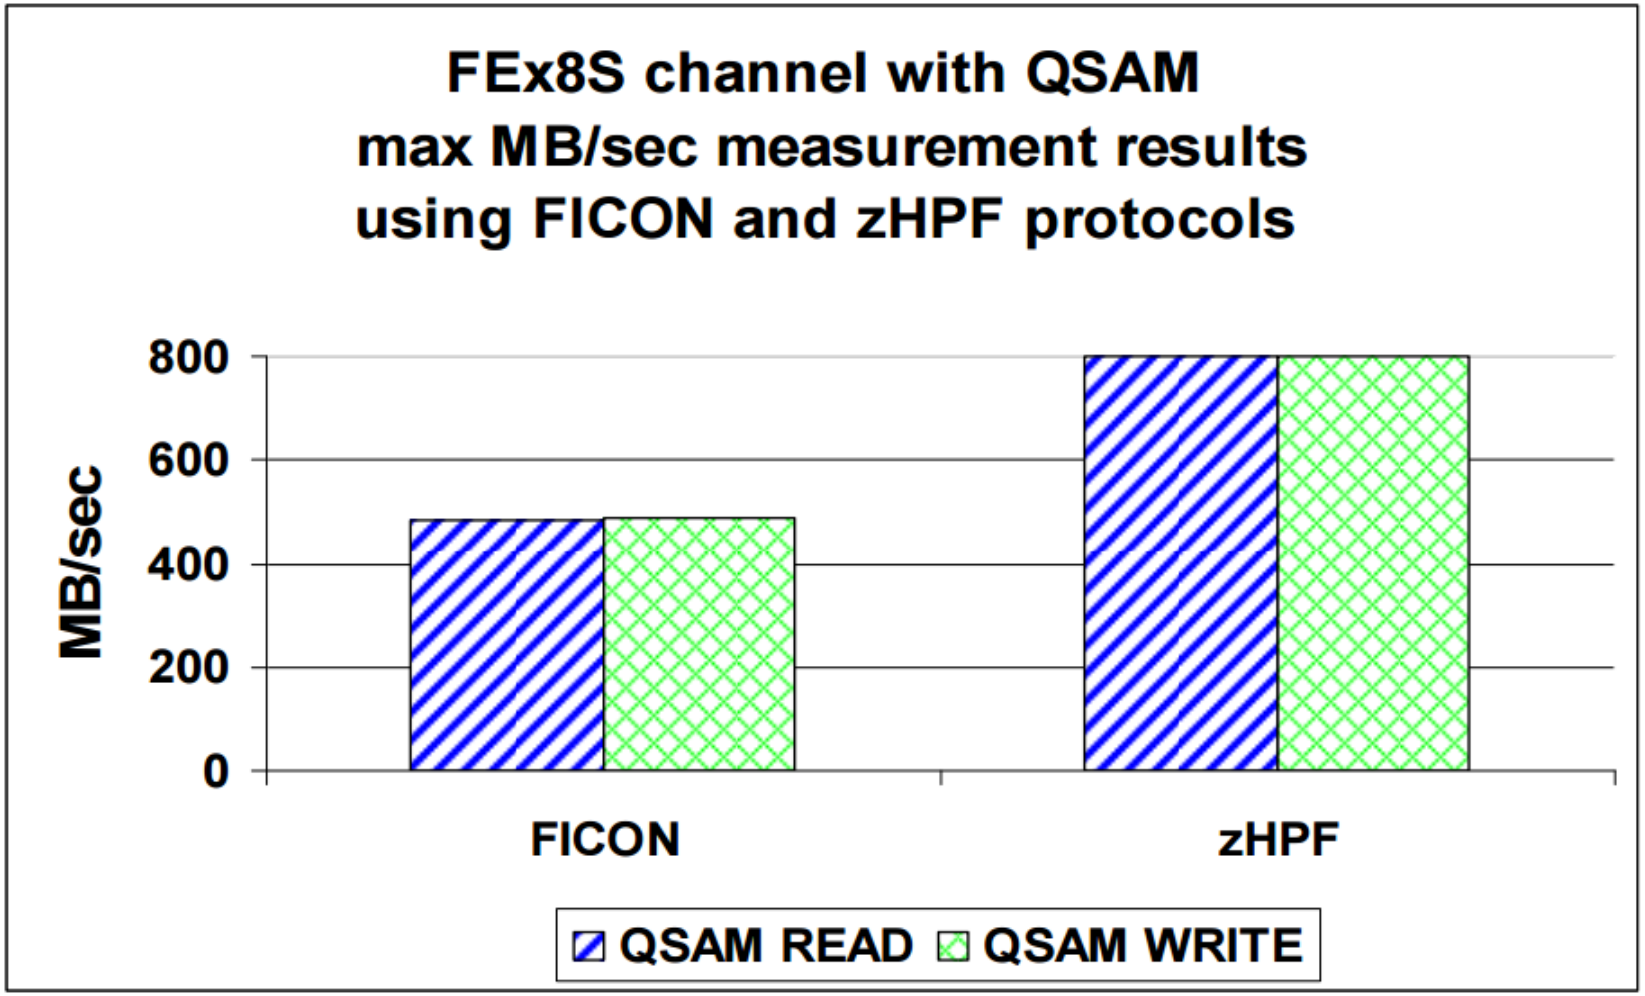

Photo credited to IBM

SAS Internal Testing Results

Parameters for Testing

- SAS Version 9.4 3rd Maintenance
- Little to no computational processing
- 27K Block Size
- SAS Direct Access Bound Libraries
- Tall and Wide

Member	Page Size	Number of Pages	Size in Mbytes
Tall	27648	240176	6333
Wide	55296	147061	7755

Increased I/O Throughput
BLKSIZE = 27K

Type of Test	Member	Elapsed Seconds CCW	Elapsed Seconds zHPF	Ratio
Random Read	Tall	32.7	22.2	0.68
Random Read	Wide	39.3	29.1	0.74
Random Write	Tall	129.5	89.7	0.69
Random Write	Wide	124.2	89.0	0.72
Sequential Read	Tall	57.7	48.0	0.83
Sequential Read	Wide	35.8	25.4	0.71

Median Elapsed Time for CCW versus zHPF Channel Programs – Library BLKSIZE = 27K

zHPF Improves I/O Rate
BLKSIZE = 27K

Type of Channel Program	Elapsed Seconds
CCW	419.2
zHPF	303.4
Ratio (zHPF/CCW)	0.72
Improvement	28%

- Using recommended Library Block Size of 27K
- 28% less elapsed time

Increased I/O Throughput
BLKSIZE = 6K

Type of Test	Member	Elapsed Seconds CCW	Elapsed Seconds zHPF	Ratio
Random Read	Tall	41.82	25.5	0.61
Random Read	Wide	48.06	32.41	0.67
Random Write	Tall	145.97	103.01	0.71
Random Write	Wide	171.1	102.04	0.60
Sequential Read	Tall	78.93	48.97	0.62
Sequential Read	Wide	61.79	30.65	0.50

zHPF Improves I/O Rate
BLKSIZE = 6K

Type of Channel Program	Elapsed Seconds
CCW	547.67
zHPF	342.58
Ratio (zHPF/CCW)	0.63
Improvement	37%

- Using Library Block Size of 6K
- 37% less elapsed time
- Strength of zHPF

zHPF Improves Channel Utilization – 8 Way

Type of Channel Program	Channel Utilization	Percentage of Bus Cycles Busy	Data Transfer Rate(MB/s)
CCW	67%	28%	440
zHPF	16%	44%	706

- 8 way Sequential Read
- Dedicated System and 2 Dedicated Channels
- zHPF had higher I/O rate but used less channel capacity.

zHPF Improves Channel Utilization – 64 Way

Type of Channel Program	Channel Utilization	Percentage of Bus Cycles Busy	Data Transfer Rate(MB/s)	Avg Connect Time
CCW	90%	25%	401	1.121
zHPF	41%	50%	797	0.135

- 64 way Random Read
- zHPF aggregate data rate almost 2xCCW
- zHPF has a significantly lower average connect time

IBM zHPF Support Requirements

Component	Requirements and Recommendations
z/OS	One of the following Levels of z/OS : • V1R12 (with APAR OA38185) • V1R13 (with APAR OA38185) • V2R1 - ZHPF = YES must be set in the IECIOSxx member of SYS1.PARMLIB(default for this parameter is 'NO') - z/OS must NOT be running as a guest under z/VM -We recommend that the fix be applied for APAR OA45589 – An incorrect zHPF Channel Program might be generated in certain retry situations
Processor	See paper reference "Which servers support zHPF?" in Burgess (2014)
DASD Controller	One of the following models in the IBM DS8000 series: • DS8700 with maintenance bundle 76.20.90.0 or higher • DS8800 with maintenance bundle 86.31.110.00 or higher • DS8870

How to use zHPF in SAS

- Implement the appropriate level of zHPF support for your System Z hardware
- Release SAS 9.4 Maintenance 2 :
  - Read operations
- Release SAS 9.4 Maintenance 3 :
  - Read and Write operations

Evaluating Performance – DLDEBUG=1X

```
libname mylib '.mylib.saslib' dldebug=1x;run;
libname mylib clear; run;
```

Channel Program Generation Statistics for Library HOONEA.MYLIB.SASLIB:

NumOperTotal	3
Method	zHPF
NumOperErrSchd	0
NumOperErrComp	0
NumOperRead	0
NumOperWrite	3

Evaluating Performance – LIBNAME LIST

- LIBNAME LIST statement display channel program type in 9.4 maintenance release 3

```
1          libname sashelp list;

NOTE :  Libref=      SASHELP
        ......
             Channel Program Type= zHPF
```

Evaluating Performance – DLDEBUG=20000x

- How to ensure a library is using CCW and not zHPF

```
libname x '.base.saslib' dldebug=20000x;
libname x list; run;
    libname x list;
NOTE: Libref=   X
Physical Name= HOONEA.BASE.SASLIB
      Unit= DISK
      Device= 3390
      Blocksize= 27648
      Blocks per Track= 2
      Total Library Blocks= 30
      Total Used Blocks= 12
      Percent Used Blocks= 40.0%
      Members= 0
      DSNTYPE= BASIC
      Data Representation= MVS_32
      Channel Program Type= CCW
```

Acknowledgements

- Lewis King, SAS (co-author)
- Fred Forst, SAS (co-author)
- Dale Riedy, IBM z/OS IOS Design and Development

Thank you for your time

- Holden O'Neal
- Email : Holden.O'Neal@sas.com
- Phone : 919-531-2312
- 100 SAS Campus Drive, Cary NC 27513

- Feel free to contact me if you have any questions or concerns