GUIDE SHARE EUROPE UK REGION

GSE UK Conference 2018
Better, stronger, faster; The Mainframe….. the Machine!

Storage Performance with zHyperLink

Nick Clayton
IBM

November 2018
Session DK

Business Value for Low Latency I/O

- Has anyone here not received a call when transactions have slowed down and lines of business are not meeting their service level agreements?
  - Why is this important?
    - Losing business?
    - Clients leaving stuff in their in-basket and leaving the site?
    - Restart times after an application or system failure?
    - Requirements to re-engineer applications to make them go faster?
    - Adding data sharing instances to get more log throughput?
- Is the corollary true?
  - If we can substantially improve transaction latency, does this add value? Open up new business opportunity? Reduce costs?

Latency comparisons for data access

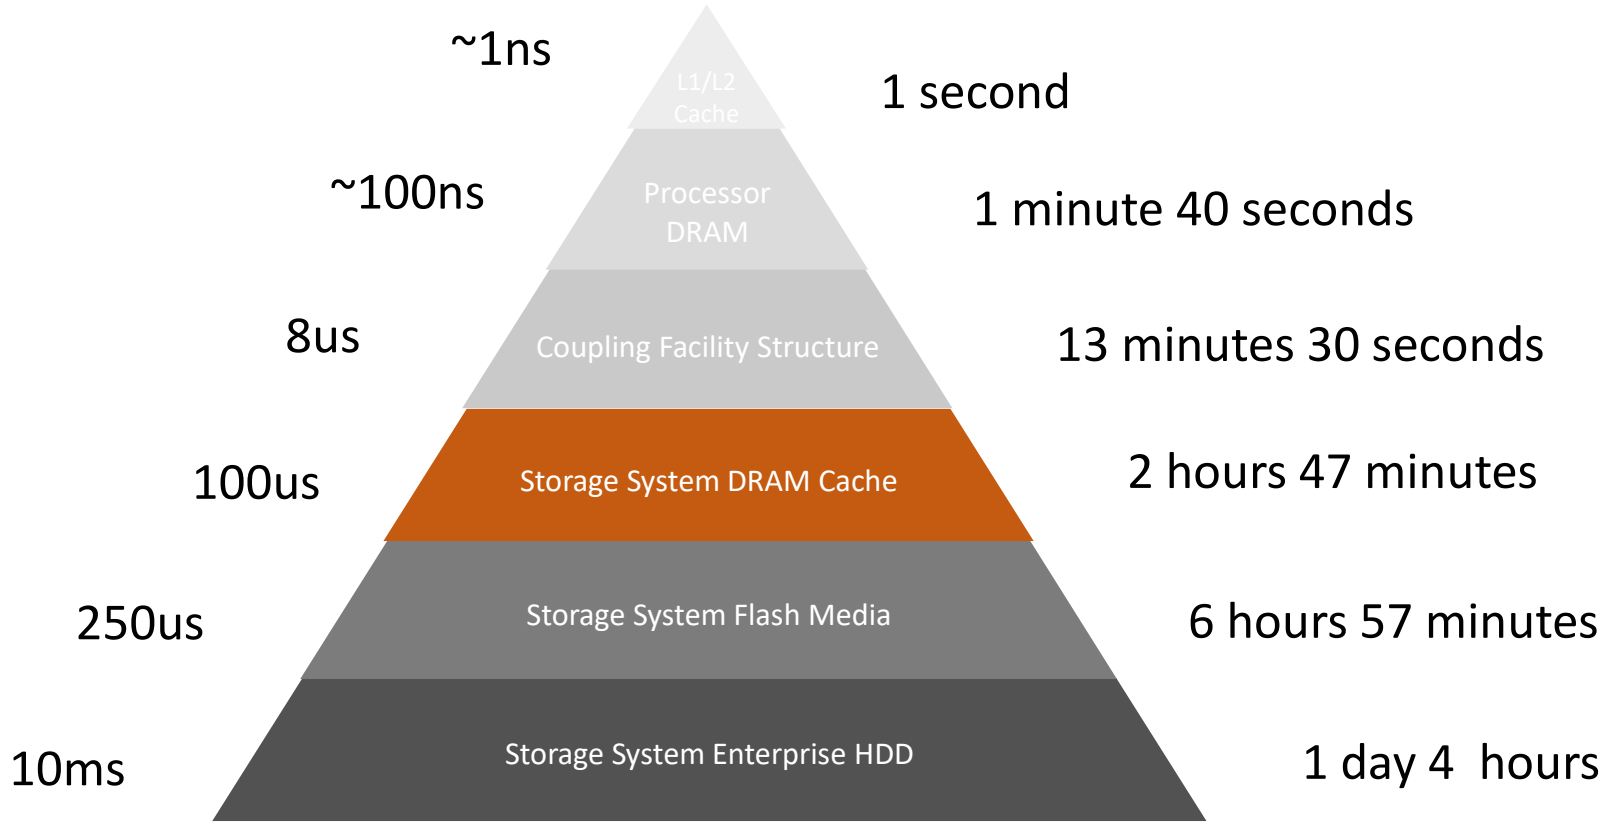

DS8880 Flash Storage for IBM Z

HDD response times become an inhibitor to realising the value of modern hardware even with high cache hit ratios seen for z/OS systems

Flash usage for IBM Z clients is continually increasing as clients realise the benefit in reduced transaction response times and faster batch processing.

While Hybrid Flash configurations are still the majority All- Flash usage is increasing rapidly.

The High Performance Flash Enclosure on the DS8880 uses Hardware RAID to provide the best latency and throughput leaving the POWER servers to perform advanced function such as replication

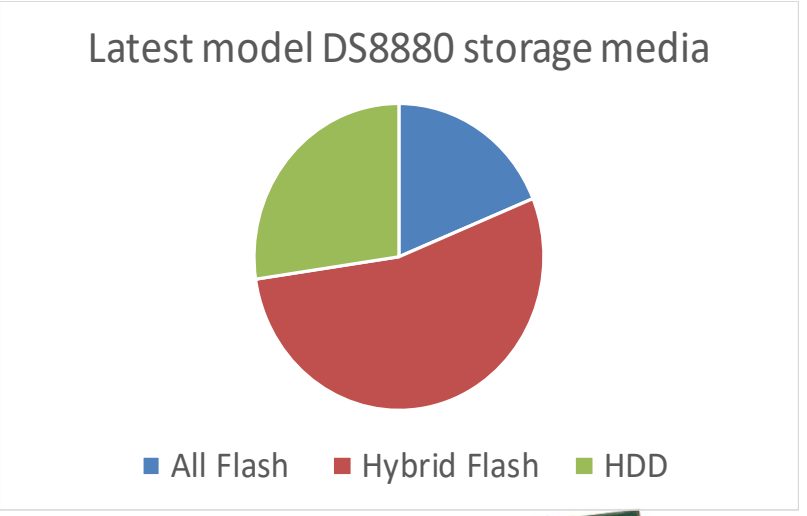

High Performance FICON (zHPF)

As hardware and storage media gets faster protocol overheads and limits become more significant

zHPF streamlines and reduces overheads compared to the FICON protocol and provides reduced latency and increased throughput. zHPF writes perform a single exchange significantly improving latency at distance but also more efficient locally

zHPF can be enabled/disabled in the OS so make sure it is enabled both overall and for Sequential IO

Not all software is exploiting zHPF so evaluate your usage using RMF and discuss future exploitation with your software vendor

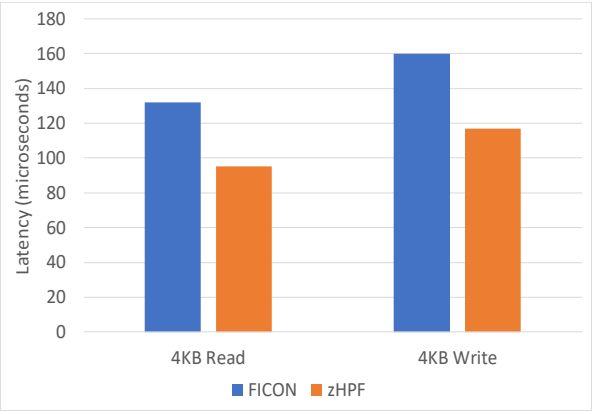

Performance in Adversity

Performance is not just about meeting SLAs in normal operation. If there are hardware or software problems then we need to minimize the impact to production performance

zHPF and FICON provides significantly faster retry times than SCSI avoiding waiting for software timeouts of many seconds (e.g. 30 seconds)

ASIC acceleration in Channels, Host Adapters and Flash RAID accelerates data transfer and enables data checking transparent to software without compromising on performance.

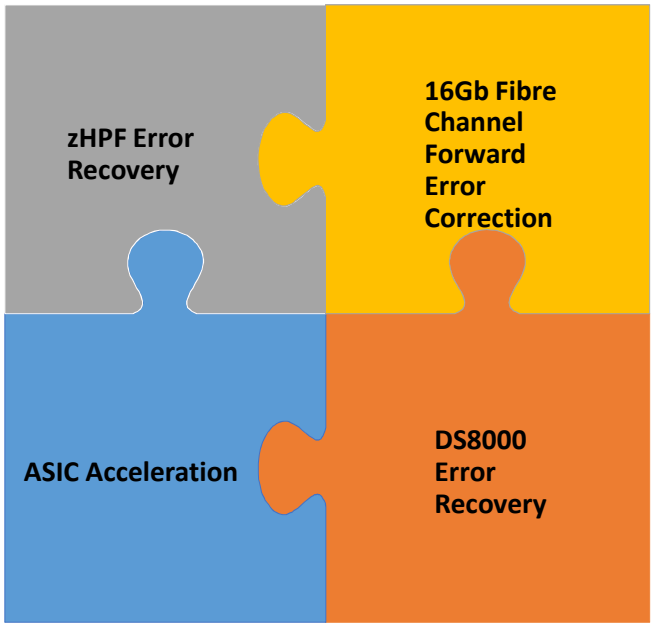

Forward Error Correction on 16Gb Fibre Channel avoids optical quality issues with faster speeds. Effective doubling the optical signal strength

DS8000 microcode provides faster error detection and recovery times. The design point is to complete ALL recoveries in 6 seconds or less – most are significantly less

zHPF Db2 list prefetch Optimizer

Db2 often requires to prefetch large amounts of data from the storage system into memory. This data could be contiguous on the storage or could be non-contiguous.

Db2 list prefetch enables the storage system to accurately prefetch non-contiguous data almost as fast as if the data was being read sequentially. This reduces the reliance on regular database reorg to maintain performance

The storage system is able to prefetch the data in parallel and so using Flash on the backend provides even greater throughput potential

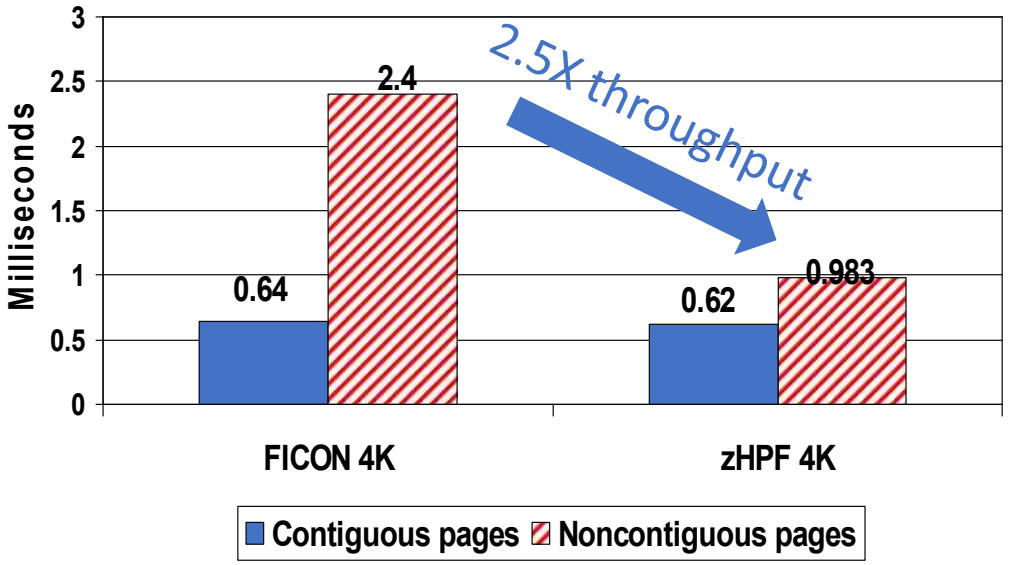

Allows users to consider reducing the frequency of database reorgs

zHyperWrite for DB2 and IMS log writes

- DB2 and IMS perform parallel Log writes with DS8880 Metro Mirror in a HyperSwap enabled environment
- Reduces latency overhead of storage based synchronous mirroring and reduces log write response time up to 43 percent providing improved log throughput
- Supported with GDPS and Copy Services Manager (CSM) for Metro Mirror and Multi-target Metro Mirror
- zHyperWrite is a pre-requisite to being able to use zHyperLink with Metro Mirror

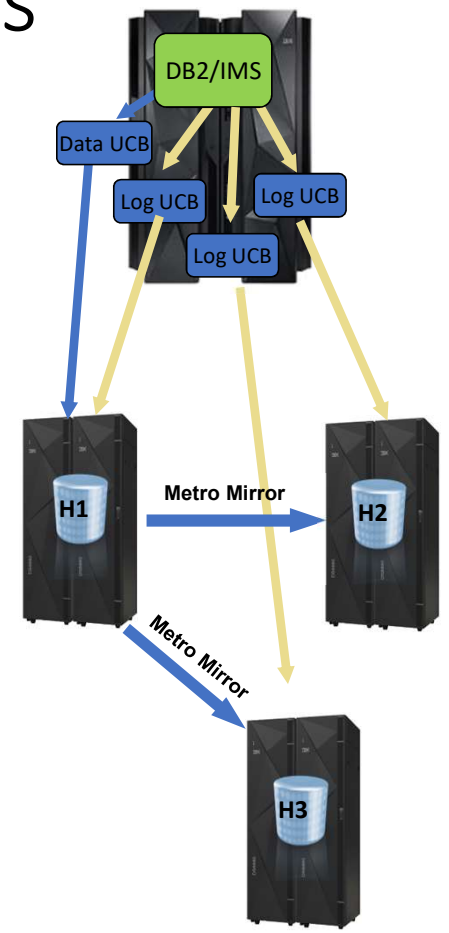

Db2 Castout Accelerator

- Db2 castout performs writes of pages from memory to storage. These are typically long chains of scattered writes to the Db2 tables.
- There are times when the performance of this process can be critical, including hardware failure events, Db2 buffer pool exhaustion and planned reconfiguration changes
- The Db2 castout accelerator changes way these writes are performed to increase throughput. For non-mirrored environments this could be up to 80% faster and in environments with Metro Mirror could be more than 5X faster.
- This enhancement also applies to other similar types of writes including zFS, PDSE and DB2 checkpoint processing.

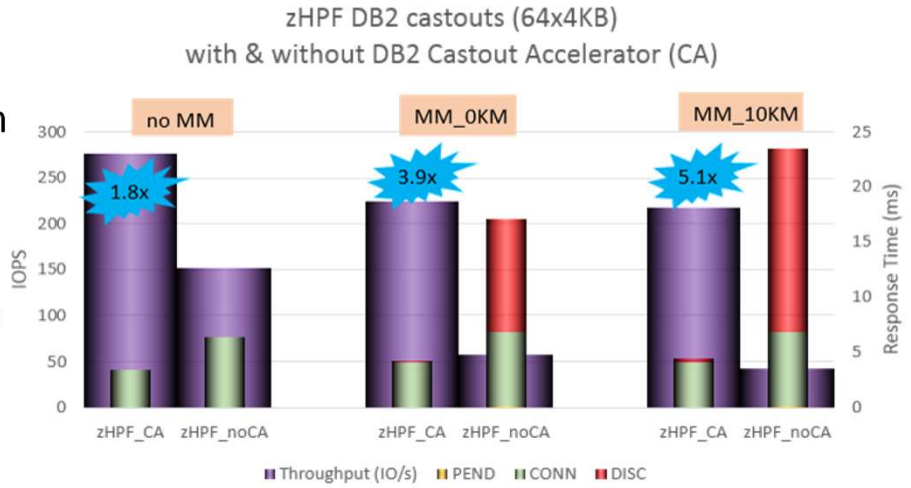

How can we fundamentally change the elapsed time needed to process a workload?

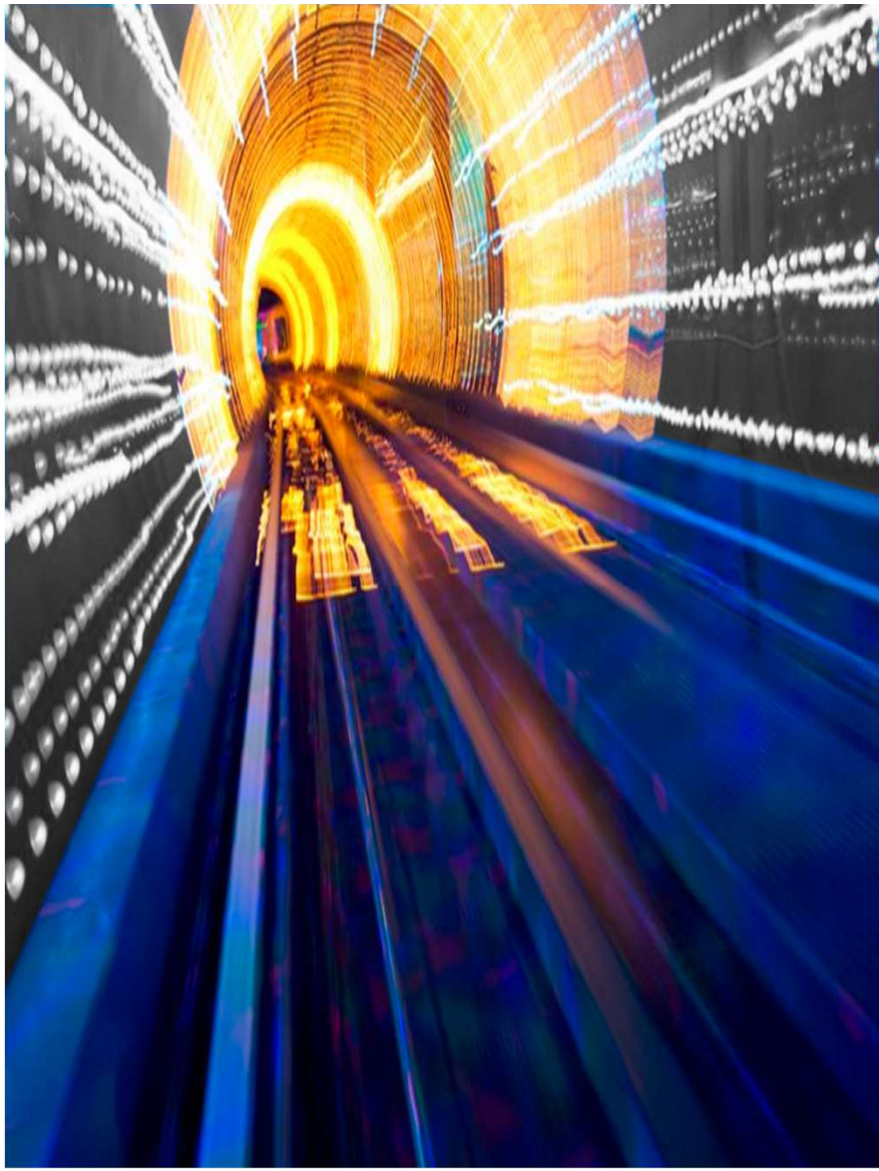

Response Time Breakdown for a Simple Transaction

- DB2 Server CPU time: 5%
- Lock/Latch + Page Latch: 2-4%
- I/O service time 60-65%
- Dispatcher (CPU) Latency: 20-25%
- Network (TCP/IP): 4-6%

- Time in BLUE is time spent Waiting

IBM zHyperLink™

- IBM zHyperLink™ is a new low latency link that connects IBM Z to storage
- IBM zHyperLink™ is the result of a cross brand project created to provide extreme low latency links between the mainframe and the storage
- IBM zHyperLink™ dramatically reduces latency by interconnecting the z System's Central Electronics Complexes (CECs) directly to the I/O Bay of the DS8880

Today's I/O

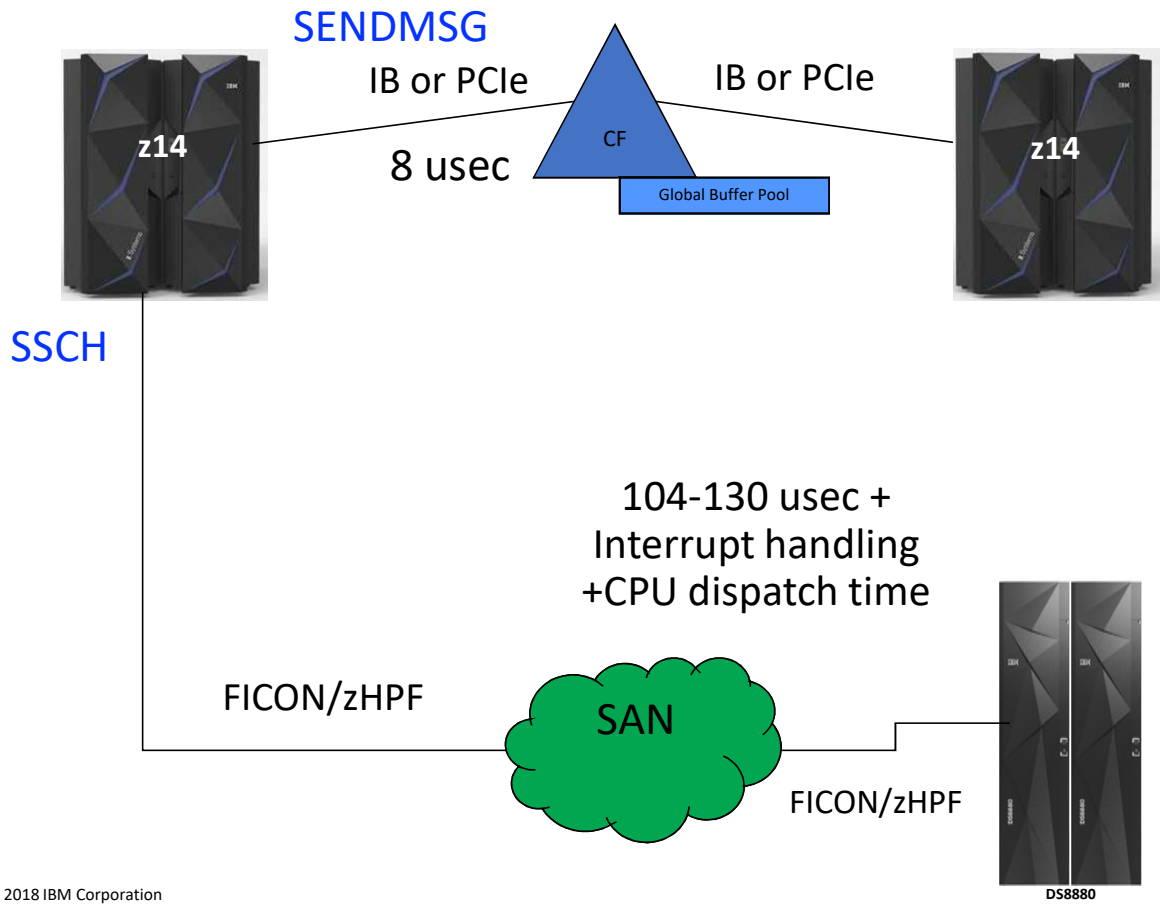

Typically over 80% cache hit ratio on random reads.

100% cache hit on writes.

How does IBM zHyperLink™ change the game?

- zHyperLink™ is FAST enough the CPU can just wait for the data
  - No un-dispatch of the running task
  - No CPU Queueing Delays to resume it
  - No host CPU cache disruption
  - Very small I/O service time
- Operating System and Middleware (e.g. DB2) are changed to keep running over an I/O
- Transparently Gives DB2 apps fundamentally better latency than apps on platforms without zHyperLink
  - Excluding 100% in memory databases

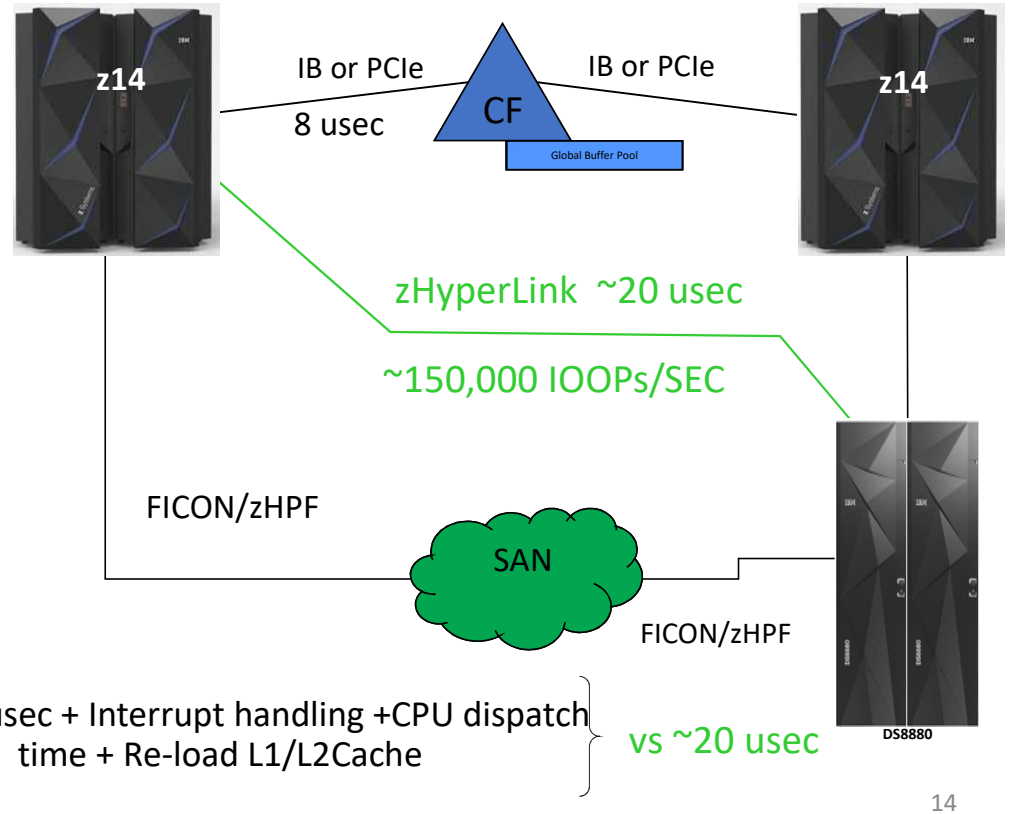

Db2 Online Transaction with zHyperLinks

(10GB buffer pools)

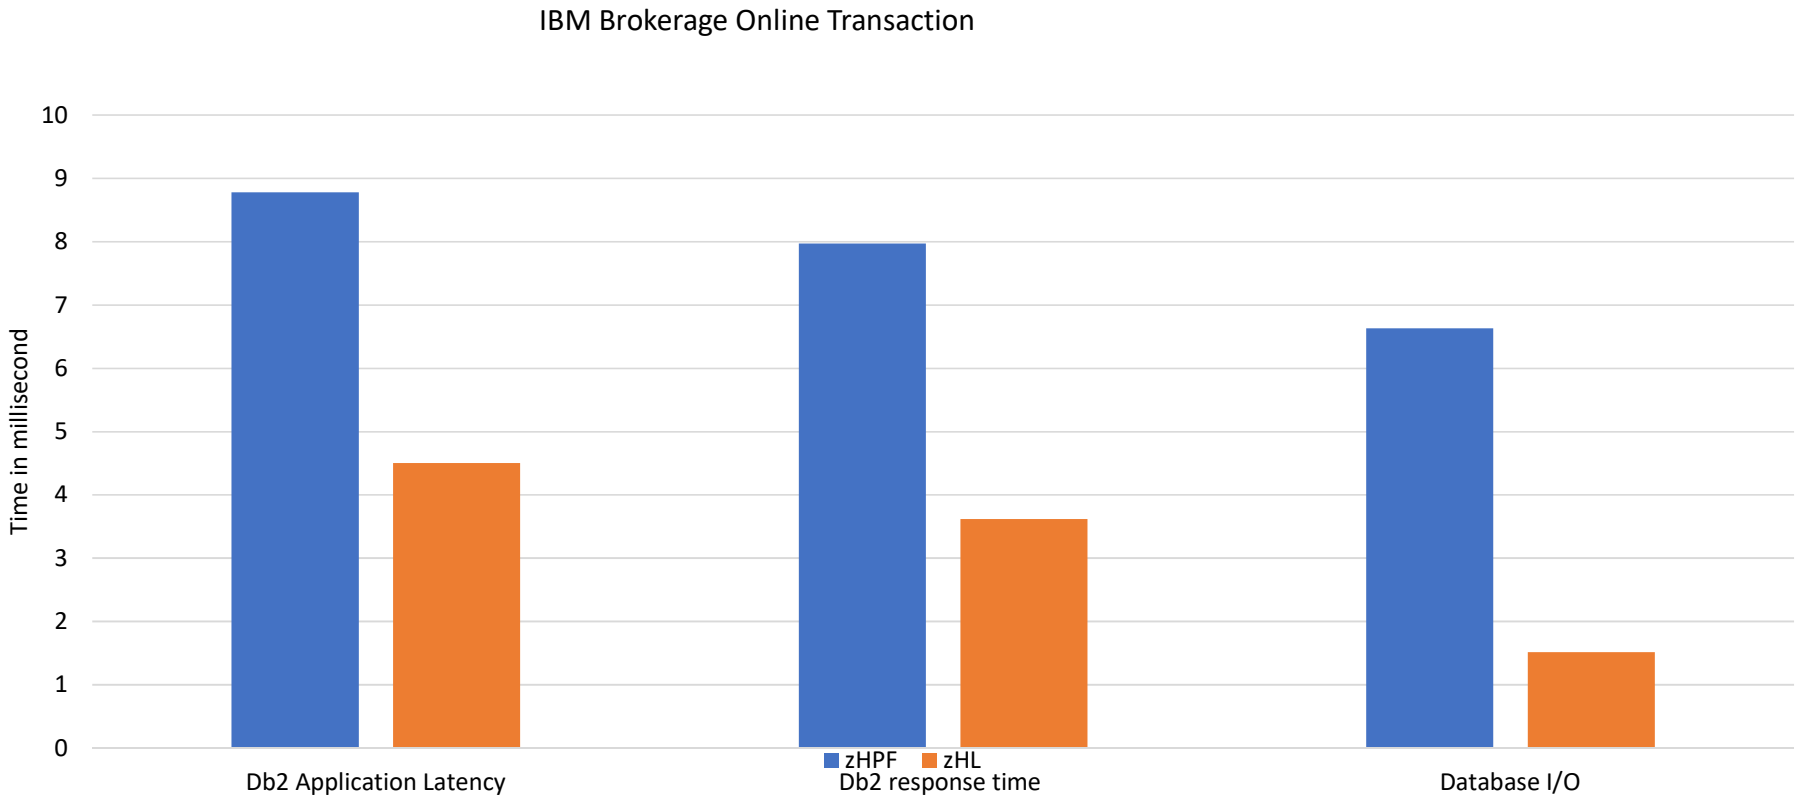

IBM DS8880 zHyperLink Ports

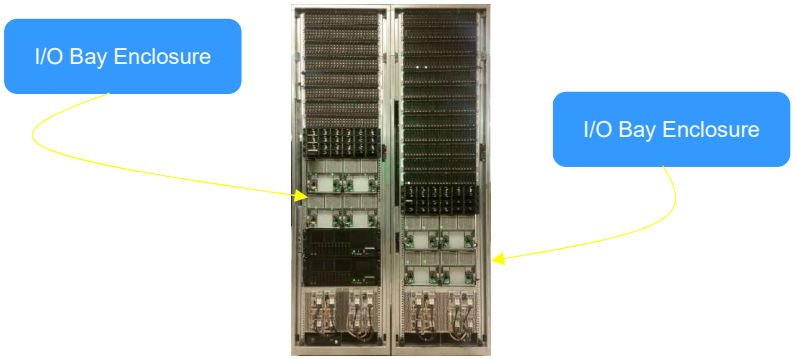

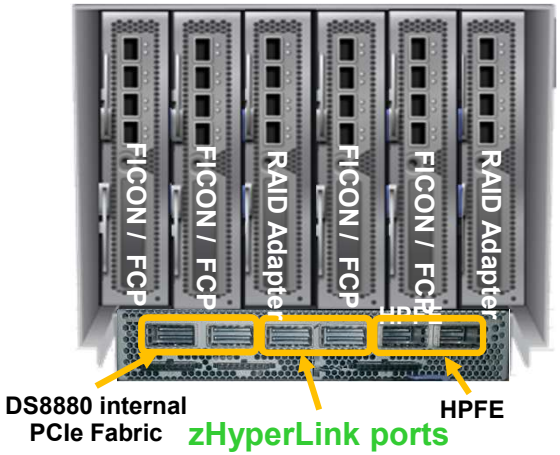

System	Cores per CEC	zHyperLink support	Max zHyperLink connections (increments of 4)
DS8882	6	No	None
DS8884 and DS8884F	6	No	None
	12	Yes	4
DS8886 and DS8886F	8	No	None
	16	Yes	8
	24	Yes	12
DS8888F	24	Yes	12
	48	Yes	16

Expansion of zHyperLink support requires R8.3.3

Investment Protection – DS8880 hardware shipping 4Q2016 (models 984, 985, 986 and 988), older DS8880's will be field upgradeable

Physical Connectivity

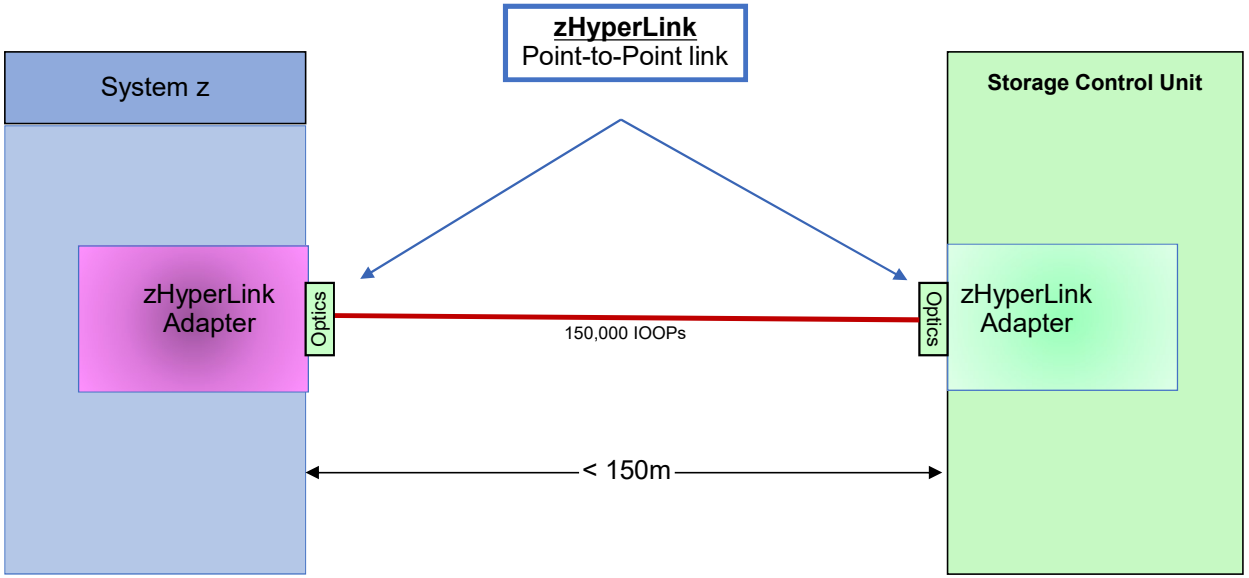

32 Synch I/O Ports supported on each z System CEC, 2 ports per card

Require 2 or more links to each CU

16 zHyperLink Ports supported on each Storage Subsystem

Require 2 or more links to each CEC

Local Primary/Remote Secondary

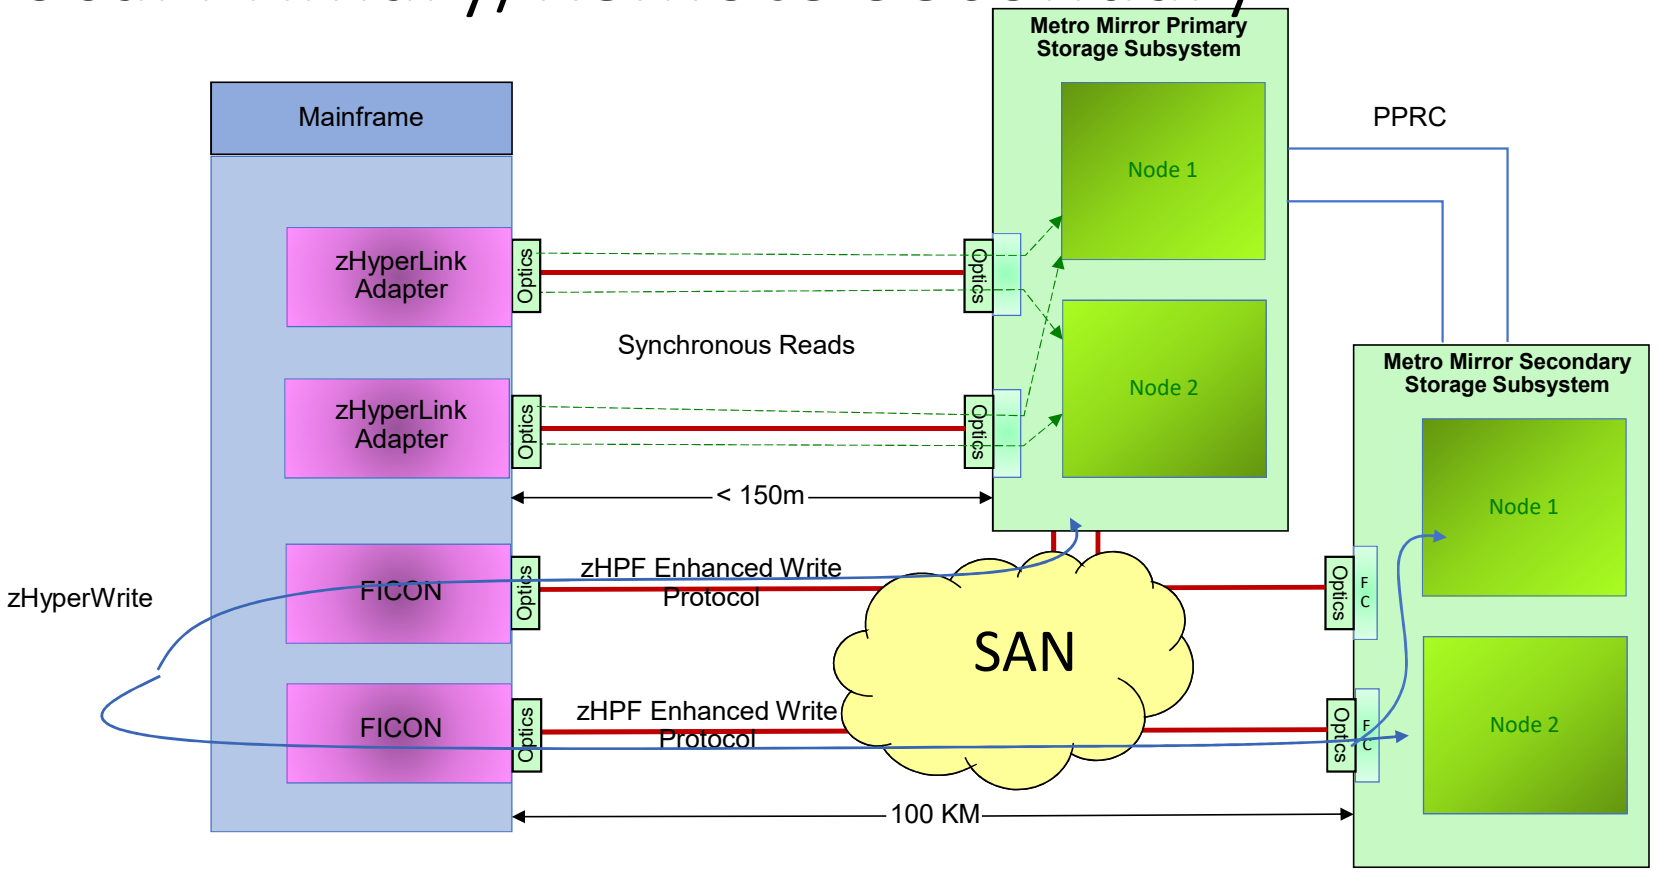

Local Primary uses zHyperLinks for reads, zHPF with enhanced write protocols and zHyperWrite for writes at distance

zHyperLink Performance – New RMF Report

```
1                              S Y N C H R O N O U S   I / O   D E V I C E   A C T I V I T Y
                                                                                                PAGE    1
            z/OS V2R2              SYSTEM ID MVS6            START 02/12/2018-17.50.00  INTERVAL 000.01.00
                                   RPT VERSION V2R2 RMF      END   02/12/2018-17.51.00  CYCLE 0.200 SECONDS
-
 TOTAL SAMPLES =    300   IODF = 71   CR-DATE: 02/02/2018   CR-TIME: 14.10.07     ACT: ACTIVATE
-                                     - DEVICE ACTIVITY RATE -  -- AVG RESP TIME --  AVG SYNCH I/O  %       %     %     %
 STORAGE  DEV   DEVICE   VOLUME  LCU  -- SYNCH I/O --   ASYNCH  -SYNCH I/O - ASYNCH  TRANSFER RATE  REQ     LINK  CACHE --REJECTS--
  GROUP   NUM   TYPE     SERIAL         READ   WRITE      I/O   READ  WRITE    I/O   READ   WRITE  SUCCESS  BUSY  MISS  READ WRITE
-         F000  33909    SRF000  0071 16083.7   0.000  4062.07  0.026 0.000  0.095  65.88  0.000    79.84 19.95  0.00  0.22  0.00
          F001  33909    SRF001  0071 15675.8   0.000  4140.55  0.026 0.000  0.095  64.21  0.000    79.11 20.68  0.00  0.21  0.00
0                         LCU    0071 31759.5   0.000  8202.61  0.026 0.000  0.095  130.1  0.000    79.47 20.31  0.00  0.21  0.00
0         F100  33909    SRF100  0072 13768.5   0.000  3322.40  0.026 0.000  0.095  56.40  0.000    80.56 19.25  0.00  0.19  0.00
          F101  33909    SRF101  0072 16067.9   0.000  3963.07  0.026 0.000  0.095  65.81  0.000    80.22 19.60  0.00  0.18  0.00
0                         LCU    0072 29836.4   0.000  7285.46  0.026 0.000  0.095  122.2  0.000    80.38 19.44  0.00  0.19  0.00
0         F200  33909    SRF200  0073 13884.8   0.000  4516.37  0.026 0.000  0.095  56.87  0.000    75.46 24.34  0.00  0.20  0.00
          F201  33909    SRF201  0073 15163.8   0.000  4171.43  0.026 0.000  0.095  62.11  0.000    78.43 21.37  0.00  0.20  0.00
0                         LCU    0073 29048.6   0.000  8687.81  0.026 0.000  0.095  119.0  0.000    76.98 22.82  0.00  0.20  0.00
0         F300  33909    SRF300  0074 13729.1   0.000  3320.81  0.026 0.000  0.095  56.23  0.000    80.52 19.29  0.00  0.19  0.00
          F301  33909    SRF301  0074 15724.3   0.000  4058.68  0.026 0.000  0.095  64.41  0.000    79.49 20.33  0.00  0.18  0.00
0                         LCU    0074 29453.3   0.000  7379.49  0.026 0.000  0.095  120.6  0.000    79.97 19.85  0.00  0.19  0.00
```

zHyperLink Performance – New RMF Report

```
1                                 E S S   S Y N C H R O N O U S   I / O   L I N K   S T A T I S T I C S
                                                                                                                         PAGE    24
             z/OS V2R2                  SYSTEM ID MVS6              START 02/23/2018-16.29.36  INTERVAL 000.00.23
                                        RPT VERSION V2R2 RMF        END   02/23/2018-16.29.59  CYCLE 0.200 SECONDS
0SERIAL NUMBER  00000RA321     TYPE-MODEL  002107-985     CDATE  02/23/2018     CTIME  16.29.37     CINT  00.00.23
0                              ----CACHE READ OPERATIONS----   ----CACHE WRITE OPERATIONS---   -----NVS WRITE OPERATIONS----
 SIID  ------LINK TYPE------      OPS    BYTES   RTIME  %SUCC      OPS    BYTES   RTIME  %SUCC      OPS    BYTES   RTIME  %SUCC
                                 /SEC      /OP     /OP             /SEC      /OP     /OP            /SEC      /OP     /OP
00080  Optical PCIe  GEN3  8   129723    4.1K     0.0   99.9      0.0    0.0      0.0    0.0      0.0    0.0      0.0    0.0
 0081  Optical PCIe  GEN1  0   NO DATA TO REPORT OR ZERO
 0180  Optical PCIe  GEN3  8   128421    4.1K     0.0   99.9      0.0    0.0      0.0    0.0      0.0    0.0      0.0    0.0
 0181  Optical PCIe  GEN1  0   NO DATA TO REPORT OR ZERO
 0280  Optical PCIe  GEN3  8   127040    4.1K     0.0   99.9      0.0    0.0      0.0    0.0      0.0    0.0      0.0    0.0
 0281  Optical PCIe  GEN1  0   NO DATA TO REPORT OR ZERO
 0380  Optical PCIe  GEN3  8   125507    4.1K     0.0   99.9      0.0    0.0      0.0    0.0      0.0    0.0      0.0    0.0
```

IBM zBNA Tool for zHyperLink

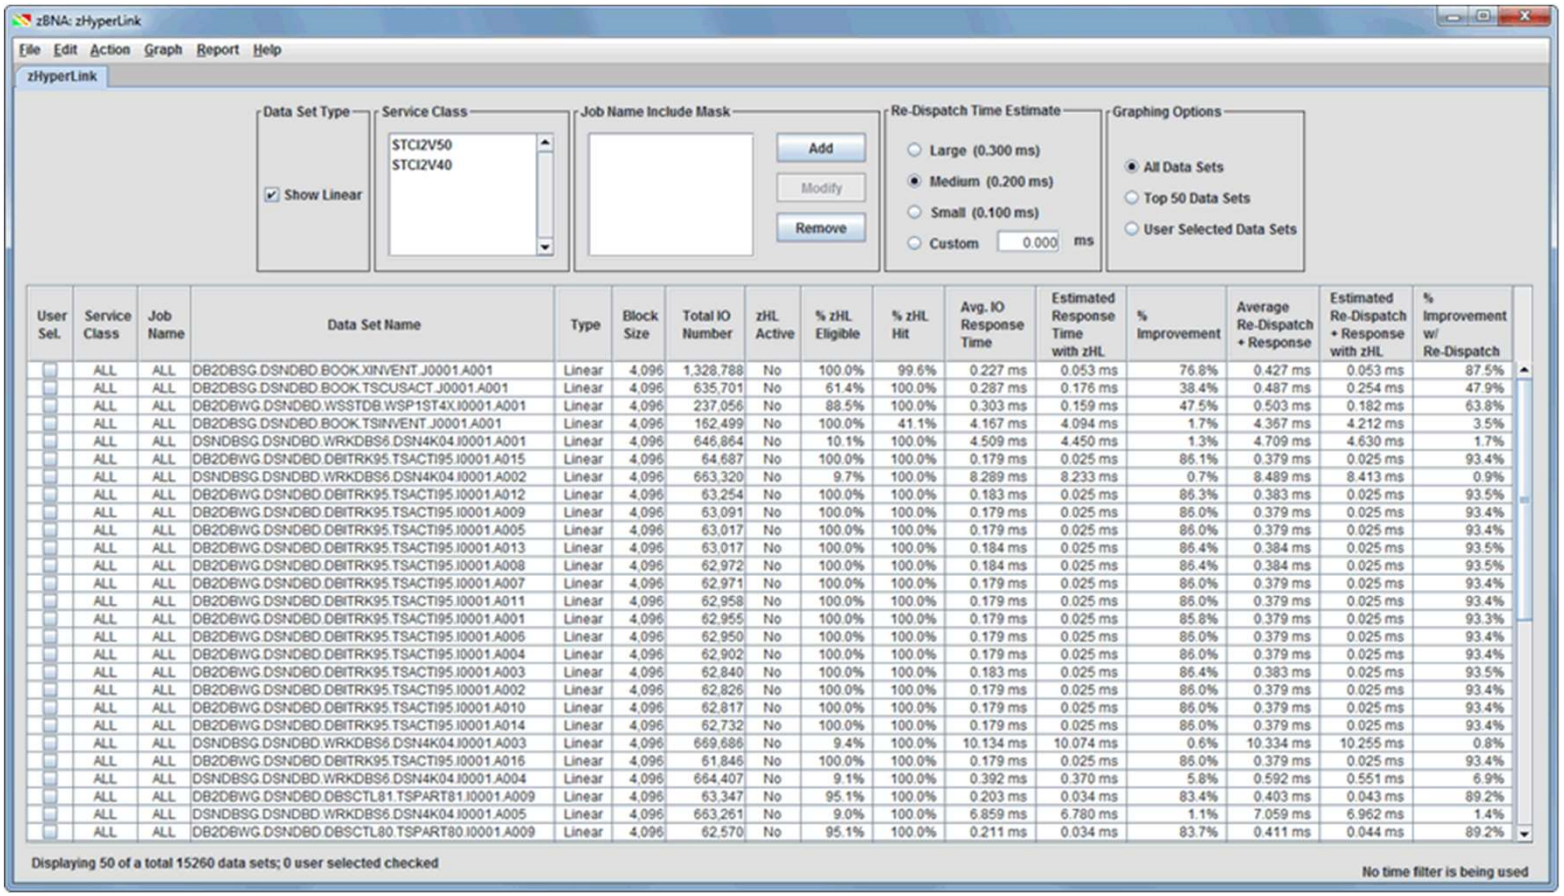

User Sel.	Service Class	Job Name	Data Set Name	Type	Block Size	Total IO Number	zHL Active	% zHL Eligible	% zHL Hit	Avg. IO Response Time	Estimated Response Time with zHL	% Improvement	Average Re-Dispatch + Response	Estimated Re-Dispatch + Response with zHL	% Improvement w/ Re-Dispatch
	ALL	ALL	DB2DBSG.DSNDBD.BOOK.XINVENT.J0001.A001	Linear	4,096	1,328,788	No	100.0%	99.6%	0.227 ms	0.053 ms	76.8%	0.427 ms	0.053 ms	87.5%
	ALL	ALL	DB2DBSG.DSNDBD.BOOK.TSCUSACT.J0001.A001	Linear	4,096	635,701	No	61.4%	100.0%	0.287 ms	0.176 ms	38.4%	0.487 ms	0.254 ms	47.9%
	ALL	ALL	DB2DBWG.DSNDBD.WSSTDB.WSP1ST4X.I0001.A001	Linear	4,096	237,056	No	88.5%	100.0%	0.303 ms	0.159 ms	47.5%	0.503 ms	0.182 ms	63.8%
	ALL	ALL	DB2DBSG.DSNDBD.BOOK.TSINVENT.J0001.A001	Linear	4,096	162,499	No	100.0%	41.1%	4.167 ms	4.094 ms	1.7%	4.367 ms	4.212 ms	3.5%
	ALL	ALL	DSNDBSG.DSNDBD.WRKDBS6.DSN4K04.I0001.A001	Linear	4,096	646,864	No	10.1%	100.0%	4.509 ms	4.450 ms	1.3%	4.709 ms	4.630 ms	1.7%
	ALL	ALL	DB2DBWG.DSNDBD.DBITRK95.TSACTI95.I0001.A015	Linear	4,096	64,687	No	100.0%	100.0%	0.179 ms	0.025 ms	86.1%	0.379 ms	0.025 ms	93.4%
	ALL	ALL	DSNDBSG.DSNDBD.WRKDBS6.DSN4K04.I0001.A002	Linear	4,096	663,320	No	9.7%	100.0%	8.289 ms	8.233 ms	0.7%	8.489 ms	8.413 ms	0.9%
	ALL	ALL	DB2DBWG.DSNDBD.DBITRK95.TSACTI95.I0001.A012	Linear	4,096	63,254	No	100.0%	100.0%	0.183 ms	0.025 ms	86.3%	0.383 ms	0.025 ms	93.5%
	ALL	ALL	DB2DBWG.DSNDBD.DBITRK95.TSACTI95.I0001.A009	Linear	4,096	63,091	No	100.0%	100.0%	0.179 ms	0.025 ms	86.0%	0.379 ms	0.025 ms	93.4%
	ALL	ALL	DB2DBWG.DSNDBD.DBITRK95.TSACTI95.I0001.A005	Linear	4,096	63,017	No	100.0%	100.0%	0.179 ms	0.025 ms	86.0%	0.379 ms	0.025 ms	93.4%
	ALL	ALL	DB2DBWG.DSNDBD.DBITRK95.TSACTI95.I0001.A013	Linear	4,096	63,017	No	100.0%	100.0%	0.184 ms	0.025 ms	86.4%	0.384 ms	0.025 ms	93.5%
	ALL	ALL	DB2DBWG.DSNDBD.DBITRK95.TSACTI95.I0001.A008	Linear	4,096	62,972	No	100.0%	100.0%	0.184 ms	0.025 ms	86.4%	0.384 ms	0.025 ms	93.5%
	ALL	ALL	DB2DBWG.DSNDBD.DBITRK95.TSACTI95.I0001.A007	Linear	4,096	62,971	No	100.0%	100.0%	0.179 ms	0.025 ms	86.0%	0.379 ms	0.025 ms	93.4%
	ALL	ALL	DB2DBWG.DSNDBD.DBITRK95.TSACTI95.I0001.A011	Linear	4,096	62,958	No	100.0%	100.0%	0.179 ms	0.025 ms	86.0%	0.379 ms	0.025 ms	93.4%
	ALL	ALL	DB2DBWG.DSNDBD.DBITRK95.TSACTI95.I0001.A001	Linear	4,096	62,955	No	100.0%	100.0%	0.179 ms	0.025 ms	85.8%	0.379 ms	0.025 ms	93.3%
	ALL	ALL	DB2DBWG.DSNDBD.DBITRK95.TSACTI95.I0001.A006	Linear	4,096	62,950	No	100.0%	100.0%	0.179 ms	0.025 ms	86.0%	0.379 ms	0.025 ms	93.4%
	ALL	ALL	DB2DBWG.DSNDBD.DBITRK95.TSACTI95.I0001.A004	Linear	4,096	62,902	No	100.0%	100.0%	0.179 ms	0.025 ms	86.0%	0.379 ms	0.025 ms	93.4%
	ALL	ALL	DB2DBWG.DSNDBD.DBITRK95.TSACTI95.I0001.A003	Linear	4,096	62,840	No	100.0%	100.0%	0.183 ms	0.025 ms	86.4%	0.383 ms	0.025 ms	93.5%
	ALL	ALL	DB2DBWG.DSNDBD.DBITRK95.TSACTI95.I0001.A002	Linear	4,096	62,826	No	100.0%	100.0%	0.179 ms	0.025 ms	86.0%	0.379 ms	0.025 ms	93.4%
	ALL	ALL	DB2DBWG.DSNDBD.DBITRK95.TSACTI95.I0001.A010	Linear	4,096	62,817	No	100.0%	100.0%	0.179 ms	0.025 ms	86.0%	0.379 ms	0.025 ms	93.4%
	ALL	ALL	DB2DBWG.DSNDBD.DBITRK95.TSACTI95.I0001.A014	Linear	4,096	62,732	No	100.0%	100.0%	0.179 ms	0.025 ms	86.0%	0.379 ms	0.025 ms	93.4%
	ALL	ALL	DSNDBSG.DSNDBD.WRKDBS6.DSN4K04.I0001.A003	Linear	4,096	669,686	No	9.4%	100.0%	10.134 ms	10.074 ms	0.6%	10.334 ms	10.255 ms	0.8%
	ALL	ALL	DB2DBWG.DSNDBD.DBITRK95.TSACTI95.I0001.A016	Linear	4,096	61,846	No	100.0%	100.0%	0.179 ms	0.025 ms	86.0%	0.379 ms	0.025 ms	93.4%
	ALL	ALL	DSNDBSG.DSNDBD.WRKDBS6.DSN4K04.I0001.A004	Linear	4,096	664,407	No	9.1%	100.0%	0.392 ms	0.370 ms	5.8%	0.592 ms	0.551 ms	6.9%
	ALL	ALL	DB2DBWG.DSNDBD.DBSCTL81.TSPART81.I0001.A009	Linear	4,096	63,347	No	95.1%	100.0%	0.203 ms	0.034 ms	83.4%	0.403 ms	0.043 ms	89.2%
	ALL	ALL	DSNDBSG.DSNDBD.WRKDBS6.DSN4K04.I0001.A005	Linear	4,096	663,261	No	9.0%	100.0%	6.859 ms	6.780 ms	1.1%	7.059 ms	6.962 ms	1.4%
	ALL	ALL	DB2DBWG.DSNDBD.DBSCTL80.TSPART80.I0001.A009	Linear	4,096	62,570	No	95.1%	100.0%	0.211 ms	0.034 ms	83.7%	0.411 ms	0.044 ms	89.2%

https://www.ibm.com/support/techdocs/atsmastr.nsf/WebIndex/PRS5132

or

Google:

IBM ZBNA DOWNLOAD

Disclaimer

The following information in this session is related to work still under development and is subject to change prior to GA.

zHyperLink Writes

- The first release of zHyperLink writes will focus on providing reduced latency for Db2 log IO to improve transaction response times

- zHyperLink write support is planned for 4Q 2018
  - DS8880 R8.5.1 microcode
  - Z14 GA2 microcode likely to be recommended
  - IOS, Media Manager, DFSMS and DB2 software will all require PTFs

zHyperLink writes and Copy Services

- zHyperLink writes will be supported on Metro Mirror primaries but only where zHyperWrite is used enabling software to write to primary and secondary

- zHyperLink writes will be supported for FlashCopy devices but significant Copy on Write activity may mean that writes at times revert to FICON

- zHyperLink writes will not be supported at GA on Global Mirror or XRC primary devices or for datasets in Concurrent Copy

zHyperLink Software exploitation

- Db2 V12 Read Support (Batch and Transactions)
  (IMS, CICS, etc. exploitation of zHyperLinks via Db2)

Future Software Support:

- Db2 V12 Writes support for Db2 logs to improve transaction response times – 4Q 18
- VSAM Read Support (Batch and Transactions)
  (IMS, CICS, etc. exploitation of zHyperLinks via VSAM)
- Other exploiters are under discussion

VSAM zHyperLink™ Exploitation

- When zHyperLink™ is enabled on a LPAR and for individual data sets:
  - All VSAM (physical) read i/o will request zHyperLink™ support.
  - Includes all record type VSAM data sets (KSDS, ESDS, RRDS, VRRDS, LDS).
  - Includes all types of VSAM read access (NSR,LSR,GSR,RLS), except NSR Sequential.
  - Currently only CISIZEs less than or equal to 4096 are supported.
  - User buffers must be on a quad-word boundary.
  - If the physical read i/o can be satisfied within 20 us, the task (TCB or SRB) issuing the i/o will not suspend or expect an interrupt.

Software Deliveries

Fix Category: IBM.Function.zHyperLink

DFSMS Exploitation for zHyperLink Express z/OS 2.1 and above.

The following are the z/OS 2.3 APAR numbers:

FMID	APAR	Comments
HDZ2230	OA53199	DFSMS (Media Mgr, Dev. Support)
	OA50681	DFSMS (Media Mgr, Dev. Support)
	OA53287	DFSMS (Catalog)
	OA53110	DFSMS (CMM)
	OA52329	DFSMS (LISTDATA)
	OA52876	**VSAM RLS zHyperlink Exp. (open)**
	OA52941	**VSAM zHyperlink Exp. (open)**
	OA52790	**SMS zHyperlink Exp. (open)**
	OA54822	SMS/CMM/BAM
	OA54824	Media Manager
	OA54825	ISMF
	OA54826	Naviquest
	OA54872	AMS

VSAM zHyperLink™ System Parameters

To enable VSAM zHyperLink™ Exploitation on each LPAR:

SYS1.PARMLIB(IGDSMSxx): **VSAM_zHyperLink(NO/YES)**

Where:
NO – Do not request zHyperLink™ exploitation for VSAM reads on this lpar (default).

YES – Request zHyperLink™ exploitation for VSAM reads on this lpar.

VSAM zHyperLink™ STORCLAS Enablement

To enable VSAM zHyperLink™ Exploitation for data sets via STORCLAS:

```
SCDS Name . . . . . : D50.RLS.SCDS
Storage Class Name  : SXPXXS02

To ALTER Storage Class, Specify:

  Guaranteed Space  . . . . . . . . . . . . .  Y        (Y or N)
  Guaranteed Synchronous Write  . . .  N       (Y or N)
  Multi-Tiered SG . . . . . . . . . . . . . . .        (Y, N, or blank)
  Parallel Access Volume Capability. .  N      (R, P, S, or N)
  CF Cache Set Name . . . . . . . . . . . .  CS2      (up to 8 chars or blank)
  CF Direct Weight  . . . . . . . . . . . . .  1       (1 to 11 or blank)
  CF Sequential Weight  . . . . . . . . . .  1        (1 to 11 or blank)
  CF Lock Set Name  . . . . . . . . . . . .            (up to 8 chars or blank)
  Disconnect Sphere at CLOSE  . . . .  Y       (Y or N)
  zHyperLink Eligible for Read  . . . . N      (Y or N)
  zHyperLink Eligible for Write . . . . . . N  (Y or N)

Use ENTER to Perform Verification; Use UP Command to View previous Page;
Use HELP Command for Help; Use END Command to Save and Exit; CANCEL to Exit.
```

VSAM zHyperLink™ Data Set Enablement

To enable VSAM zHyperLink™ Exploitation for data sets (intended for testing purposes):

VARY SMS,DSName(datasetname),{DELETE},
{,ZHLREAD={YES|NO|DEFAULT},
{,ZHLWRITE={YES|NO|DEFAULT}

Parameters:
DSN(datasetname) the data set name. For VSAM, this must be a component name or cluster name.

ZHLREAD=YES|NO|DEFAULT
Specifies whether the data set can use zHyperLink for read requests.
YES
This data set can use zHyperLink for read requests. It will override the Storage Class setting for zHyperLink read eligibility.
NO
This data set cannot use zHyperLink for read requests. It will override the Storage Class setting for zHyperLink read eligibility.
DEFAULT
This data set will use the Storage Class setting for zHyperLink read eligibility. When the data set is closed, zHyperLink read eligibility is governed by the Storage Class. If there is no Storage Class associated with the data set, the default of not using zHyperLink read will be active. If the data set is non-SMS, zHyperLink will not be used to read this data set.

VSAM zHyperLink™ SMF Support

The following new SMF 42-6 fields are available for zHyperlink reads:

- S42SNERD Eligible Synchronous I/O Reads
- S42SNERH Eligible Synchronous I/O Read Hits
- S42SNRDT No of sync_io read attempts.
- S42SNROS No of sync_io read successes.
- S42SNSEQ No of sync_io sequential operation
- S42SNCND No of sync_io cache candidates.
- S42SNHTS No of sync_io cache hits
- S42SNRMS No of sync_io read misses.
- S42SNMXR Max sync_io read response
- S42SNRDU Avg sync_io Read Time in usec
- S42SNCONC No of concurrent synch I/Os per req

The zBNA and CP3KEXTR programs can be used to analyze zHyperLink performance. The latest versions that must be used are zBNA version 1.8.2 and CP3KEXTR version 3.75.

Redbook: Getting Started with IBM zHyperLink for z/OS
(http://www.redbooks.ibm.com/redpapers/pdfs/redp5493.pdf)

zHyperLink™ VSAM NSR Time 82.4% Elapse Time Improvement

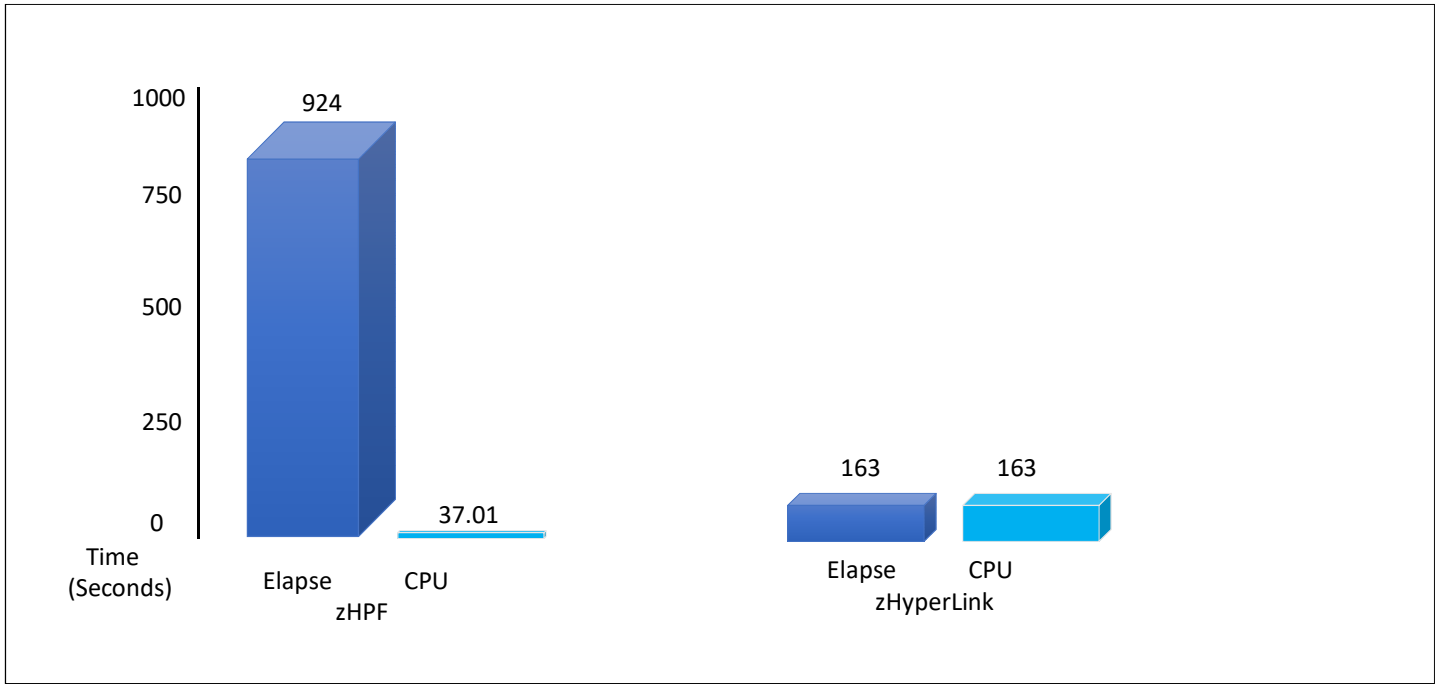

VSAM 2M NSR Direct Reads

4 KB Records / 4 KB CISIZE

DASD Cache primed

zHyperLink™ VSAM/RLS

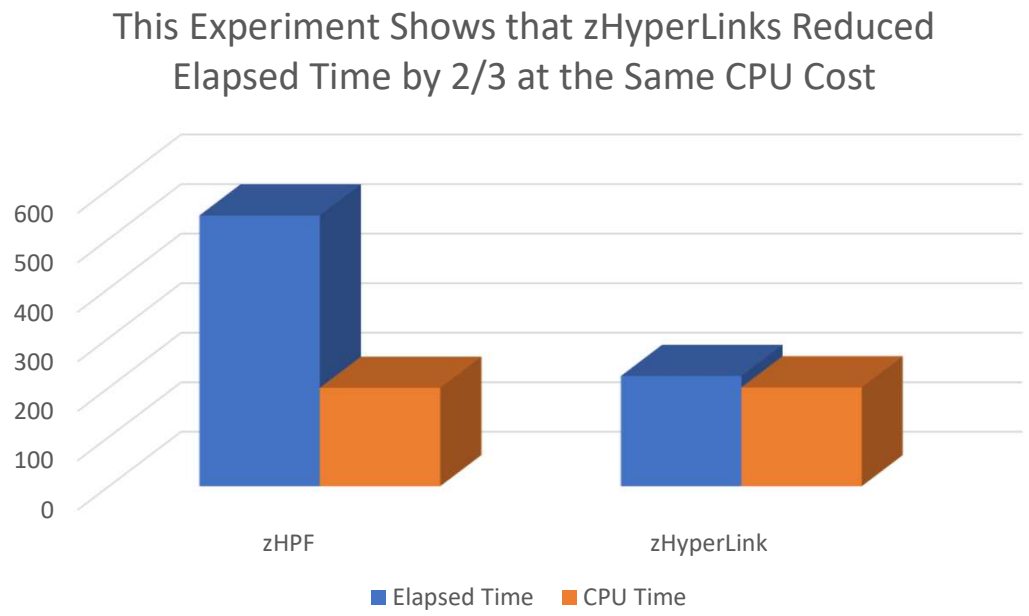

- VSAM RLS Sequential Reads, 4 KB Records
  - Ran job to load the DASD Cache
  - Started with no data in the local buffer pool
  - no CF cache usage

- High CPU contention from other LPARs sharing the processor (I/O Interrupt Delay time >40 usec)
- Very low CPU utilization from application LPAR, just one job.
- VSAM access method only touches dataset metadata, the extended format dataset suffix, requiring L1/L2 loading

We want your feedback!

- Please submit your feedback online at ….
  - http://conferences.gse.org.uk/2018/feedback/DK
- Paper feedback forms are also available from the Chair person
- This session is DK