IBM z15 FICON Express16SA Performance

January 2020

Brian Murphy
Yamil Rivera
Dinesh Kumar

IBM Systems

Table of Contents

Notices, Disclaimer and Trademarks

Acknowledgements

The authors would like to thank the following colleagues for their expertise and insight:

- Su KR Han – IBM Systems, Shanghai, China
- Fei Fei Li – IBM Systems, Shanghai, China
- Andrea Harris –IBM Systems, Poughkeepsie, NY
- Shelly Black –IBM Systems, Poughkeepsie, NY
- Sonny Williams – IBM Systems, Tucson, AZ

The authors would also like to thank all of the reviewers of this paper for their helpful comments.

Introduction

This whitepaper assumes a familiarity with the general concepts of Z Systems. Readers unfamiliar with these topics should consult the IBM Z I/O Connectivity website :
https://www.ibm.com/it-infrastructure/z/capabilities/connectivity

In this paper, "FICON Express16SA" and "FEx16SA" are used interchangeably unless otherwise noted in the context of the text. When zHPF is used by itself, it refers to the zHPF Protocol. When FICON is used by itself, it refers to the FICON Protocol. When FCP is used by itself, it refers to the FCP Protocol.

The z15 supports the following FICON features:

- FICON Express16SA (FEx16SA) - new build
- FICON Express16S+ (FEx16S+) - carry forward only
- FICON Express16S (FEx16S) - carry forward only
- FICON Express8S (FEx8S) - carry forward only

The FICON features conform to the following architecture:

- Fibre Connections (FICON)
- High Performance FICON on Z (zHPF)
- Fibre Channel Protocol (FCP)

The FICON features provide connectivity between any combination of servers, directors, switches, and devices (control units, disks, tapes, and printers) in a SAN.

The FEx16SA feature is encryption capable in support of Fibre Channel Endpoint Security (FC-ES).

Each FICON feature occupies one I/O slot in the PCIe I/O drawer. Each feature has two ports, with one PCHID and one CHPID associated with each port.

Summary

The z15 FEx16SA channel offers many benefits over previous generations of FICON channels.

The z15 supports the new form factor I/O cards which use PCIe Gen3 links with increased capacity, granularity, and infrastructure bandwidth.

Enabling Fibre Channel Endpoint Security protects your data in flight, offloading encryption processing to the FEx16SA channel with less than 2% impact on maximum zHPF throughput and less than 4% impact on maximum FCP throughput.

The increased capability of the FEx16SA channels is complemented by improved performance at many levels of the z15 I/O Subsystem.

- The z15 PCIe fanout card with two PCIe Interconnects with sixteen FEx16SA channel cards supports up to 50 GB/sec Read/Write (mix).
- The maximum I/O processing capability on a z15 with 22 SAPs is over 5.3 M IO/sec, a 23% increase compared to a z14 with 23 SAPs.
- Single zOS LPAR running on z15 saturates at about 950K IO/sec.

Additional zHPF, FCP, and FICON product information is available on the IBM Z websites:

https://www.ibm.com/it-infrastructure/z/capabilities/connectivity

https://www.redbooks.ibm.com/redpieces/pdfs/sg248851.pdf

Fibre Channel Endpoint Security

Fibre Channel Endpoint Security is an end-to-end solution which ensures all data flowing between the new encryption capable 16Gbps FICON Express SA channel and DS8900 storage is encrypted and protected. Fibre Channel Endpoint Security is an additional data security technology that contributes to the IBM Z approach of encryption everywhere, further minimizing the risk of a security breach by extending the value of pervasive encryption.

Fibre Channel Endpoint Security provides in-flight protection for all data over FICON and Fibre Channel (FCP) links, independent of the operating system, file system, or access method in use.

There are 3 key components of IBM Fibre Channel Endpoint Security:

1. IBM z15 with FICON Express16SA, the IBM z15 supports new encryption capable 16Gbps FICON cards for securing Channel-to-Storage Control Unit links.

2. IBM DS8900, the new IBM DS8900 supports new 32Gbps encryption capable FICON adapter cards to support secured links between IBM Z and the storage control unit.

3. IBM Security Key Lifecycle Manager, ISKLM version 3.0.1 External Key Manager (EKM) appliance is required to serve a shared key to the server and storage that is used to secure messages between the endpoints in which further key material is derived/exchanged. Key Management Interoperability Protocol (KMIP) is the standard key management protocol used for key manager communication. Channel and Storage System authenticate with key server, and Channel and Storage System CU query server for shared keys.

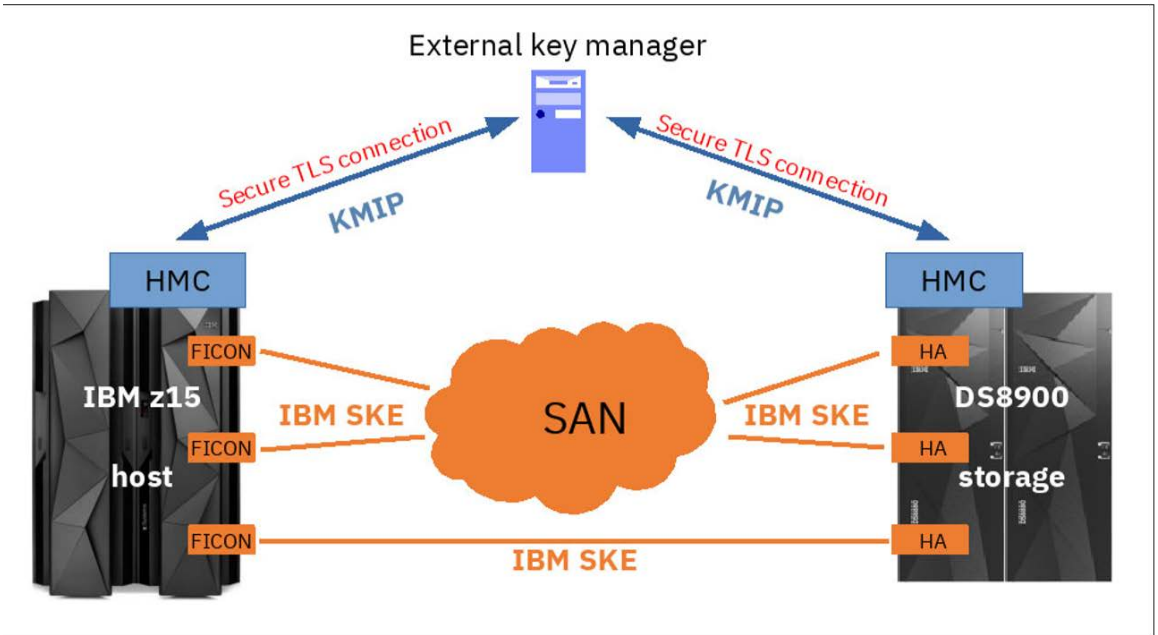

Figure 1 – Fibre Channel Endpoint Security

The High Performance FICON for IBM Z:

High Performance FICON for IBM Z (zHPF) is implemented for throughput and latency, which it optimizes by reducing the number of information units (IU) that are processed. Enhancements to the z/Architecture and the FICON protocol provide optimizations for online transaction processing (OLTP) workloads. Reflected in the bar charts in **Figure 2** and **Figure 3** below are the maximum single channel throughput capacities of each of the generations of FICON Express channels supported on z15. For each of the generations of FICON Express channels, displays the maximum capability using the zHPF protocol exclusively.

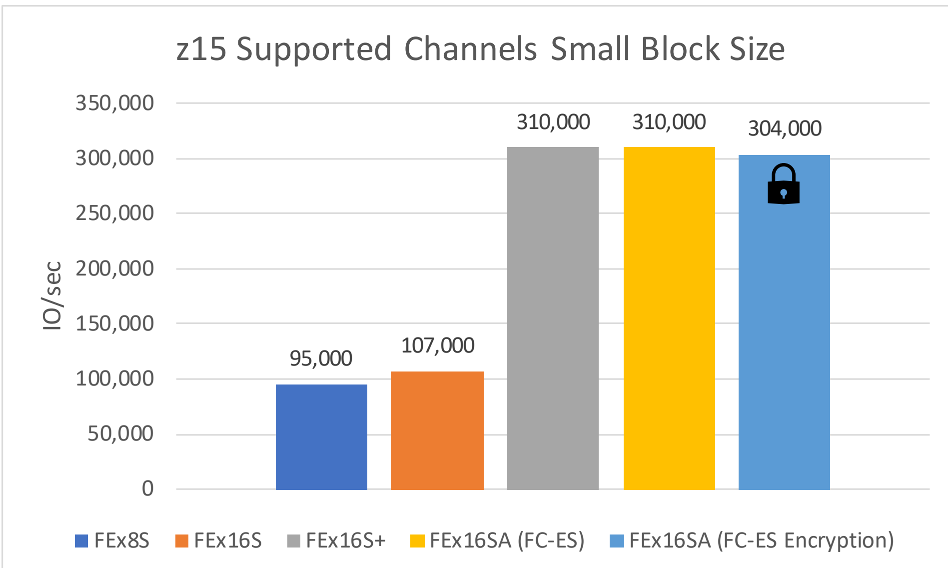

Figure 2 – FICON Express Channel Maximum IO/sec

The size of most online database transaction processing (OLTP) workload I/O operations is 4K bytes. Laboratory measurements, a FEx16SA channel, using the zHPF protocol and small data transfer I/O operations, achieved a maximum of 310K IO/sec. Also, using FEx16SA in a z15 using the zHPF protocol and small data transfer FC-ES Encryption enabled, achieved a maximum of 304K IO/sec.

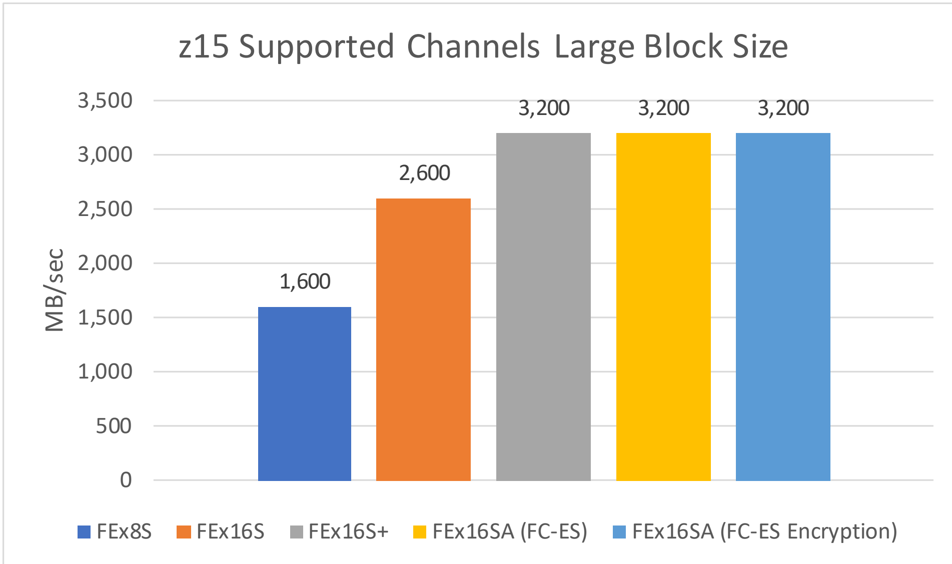

Figure 3 – FICON Express Channel Maximum MB/sec

The chart above displays the maximum READ/WRITE (mix) MB/sec for each channel. Laboratory measurements, using FEx16SA in a z15 with the zHPF protocol and a mix of large sequential read and update write data transfer I/O operations, achieved a maximum throughput of 3,200 READ/WRITE (mix) MB/sec. Also, using FEx16SA in a z15 with the zHPF protocol with FC-ES Encryption and a mix of large sequential read and update write data transfer I/O operations, achieved a maximum throughput of 3,200 READ/WRITE (mix) MB/sec.

The maximum zHPF 4K IO/sec measured on a FEx16SA channel approximately 12x increase of the maximum FICON protocol capability. Response time curves for these results will be described later in this document. As shown in **Figures 2 and 3**, with zHPF, it is possible to achieve an improvement in both small block IO/sec processing for OLTP workloads and large sequential READ/WRITE (mix) I/O processing.

High Performance FICON for System z (zHPF) architecture performance features

zHPF is an extension to the FICON architecture designed to improve the execution of small block I/O requests. zHPF streamlines the FICON architecture and reduces the overhead on the channel processors, control unit ports, switch ports, and links by improving the way channel programs are written and processed. To understand how zHPF improves upon FICON, one needs to review the relevant characteristics of FICON channel processing.

A FICON channel program consists of a series of Channel Command Words (CCWs) which form a chain. The command code indicates whether the I/O operation is going to be a read or an update write from disk, and the count field specifies the number of bytes to transfer. When the channel finishes processing one CCW, and either a command chaining or data chaining flag is turned on, it processes the next CCW, and the CCWs belonging to such a series are said to be chained. Each one of these CCWs is a FICON channel Information Unit (IU) which requires separate processing on the FICON channel processor and separate commands to be sent across the link from the channel to the control unit. The zHPF architecture defines a single command block to replace a series of FICON CCWs as illustrated in **Figure 4** below.

zHPF improves upon FICON by providing a Transport Control Word (TCW) that facilitates the processing of an I/O request by the channel and the control unit. The TCW has a capability that enables multiple channel commands to be sent to the control unit as a single entity instead of being sent as separate commands as is done with FICON CCWs. In addition, the channel is no longer expected to process and keep track of each individual channel command word. Instead, the channel forwards a chain of commands to the control unit to execute. The reduction of this overhead increases the maximum I/O rate possible on the channel and improves the utilization of the various sub-components along the path traversed by the I/O request.

zHPF provides a much simpler link protocol than FICON. **Figure 4** below shows an example of a 4K read FICON channel program, where three IUs are sent from the channel to the control unit plus three IUs from the control unit to the channel. In this example, zHPF reduces the total number of IUs sent by half, using one IU from the channel to the control unit and two IUs from the control unit to the channel.

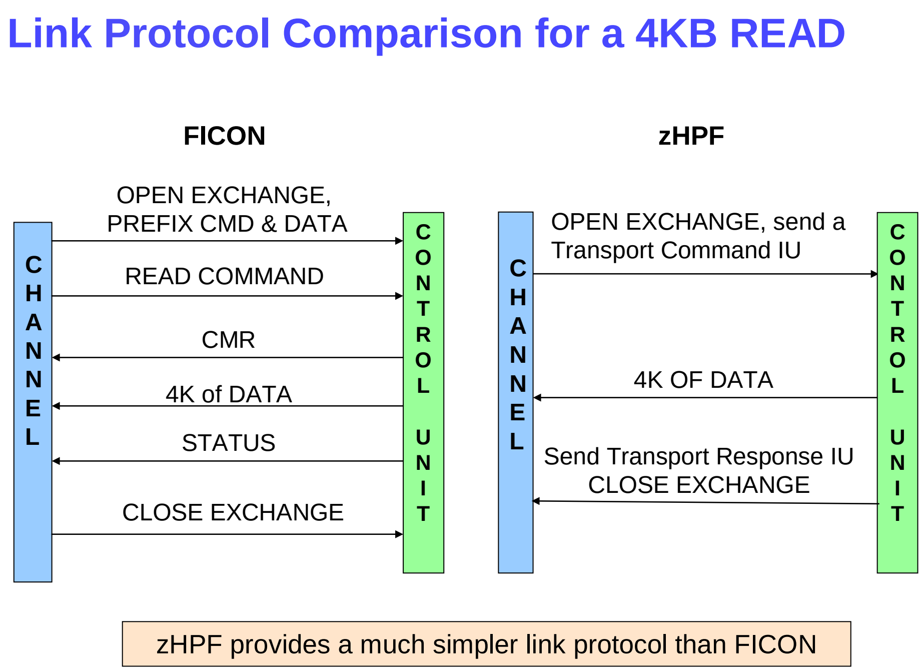

Figure 4 – Link Protocols

With zHPF, "well constructed" CCW strings are collapsed into a single new Control Word. Conceptually this is similar to the Modified Indirect Data Address Word (MIDAW) facility enhancement to FICON, which allowed a chain of data CCWs to be collapsed into one CCW. zHPF now allows the collapsing of both Command Chained as well as Data Chained CCW strings into one Control Word. zHPF-capable channels and devices support both FICON and zHPF protocols simultaneously.

The maximum number of open exchanges or the number of I/Os that can be simultaneously active on FEx16SA channels is designed to be significantly higher with zHPF compared to FICON. 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, which begin transferring data back to the channel immediately and those that are cache misses which might experience a delay of several milliseconds before the data can begin transferring back to the channel. Since higher I/O activity levels are both possible now and expected to increase in the future with zHPF, the maximum number of open exchanges allowed per channel has been increased with zHPF. More information on how to find the actual average number of open exchanges used in a production workload environment is provided in the **"zHPF fields on the RMF Channel Activity report"** section of this document.

FEx16SA channel performance using zHPF and FICON protocols

With the introduction of the FEx16SA channel, improvements can be seen in both response times and maximum throughput for IO/sec and MB/sec for workloads using the zHPF protocol on FEx16SA channels.

One primary performance improvement in a FEx16SA channel with zHPF compared to FICON is the maximum number of small block IO/sec (4K bytes per I/O) that can be processed. As displayed in **Figure 5** below, the maximum number of IO/sec that was measured on a FEx16SA channel running an I/O driver benchmark with a 4K bytes per I/O workload exploiting zHPF is 310K IO/sec, which is approximately twelve times what was measured with FICON. In this case, two separate experiments were conducted; one where the channel was executing all zHPF channel programs and another where the channel was executing all FICON channel programs.

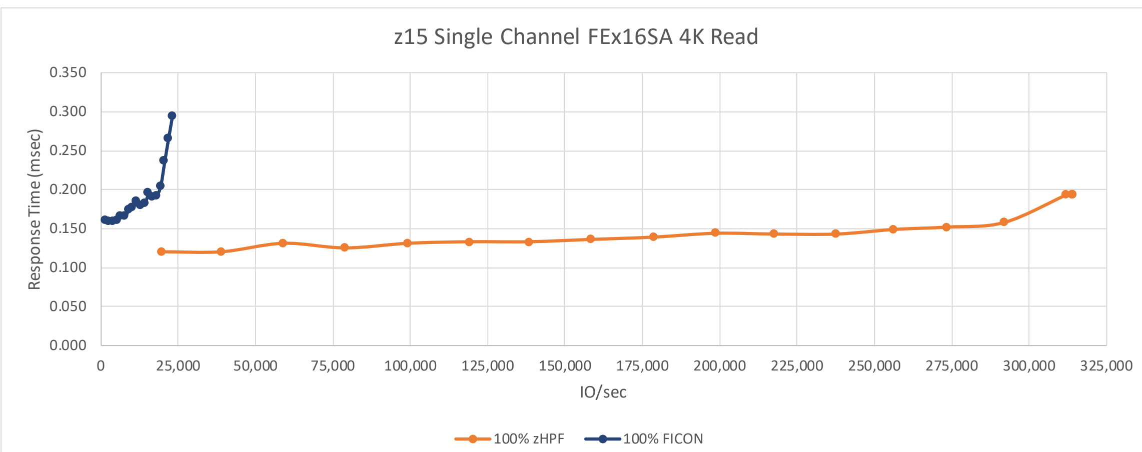

Figure 5 – FEx16SA Maximum IO/sec

Figure 5 shows that the total response time for a 4K read I/O operation is approximately 100 microseconds better for zHPF compared to FICON channel programs at I/O activity levels less than 20K IO/sec per channel. For FICON, response times increase gradually to approximately 0.2ms for activity levels up to 18K IO/sec per channel and then increase sharply at a level of activity between 20K and 23K IO/sec. In contrast, response times for zHPF continue at a moderate rate of increase up to approximately 0.19ms at 300K IO/sec, followed by the sharp increase indicative of reaching the maximum capability of the channel.

Please note that these measurements were done with a single z15 channel connected to multiple host adapter ports, each on an IBM System Storage DS8950. Contact an IBM representative for more information on the performance of a DS8900.

Since the difference in maximum capability with 100% zHPF compared to 100% FICON is so great on the FEx16SA channel, it is worthwhile to note how the maximum capability changes for I/O benchmarks that read 4K bytes of data using a mixture of zHPF and FICON. As shown in **Figure 8** below, with a mixture of 90% zHPF targeted devices and 10% FICON targeted devices, the maximum capability is 240K IO/sec. With a mix of 50% zHPF targeted devices and 50% FICON targeted devices, the maximum capability is 160K IO/sec on a FEx16SA channel. Higher performance improvements can be experienced with higher percent zHPF activity.

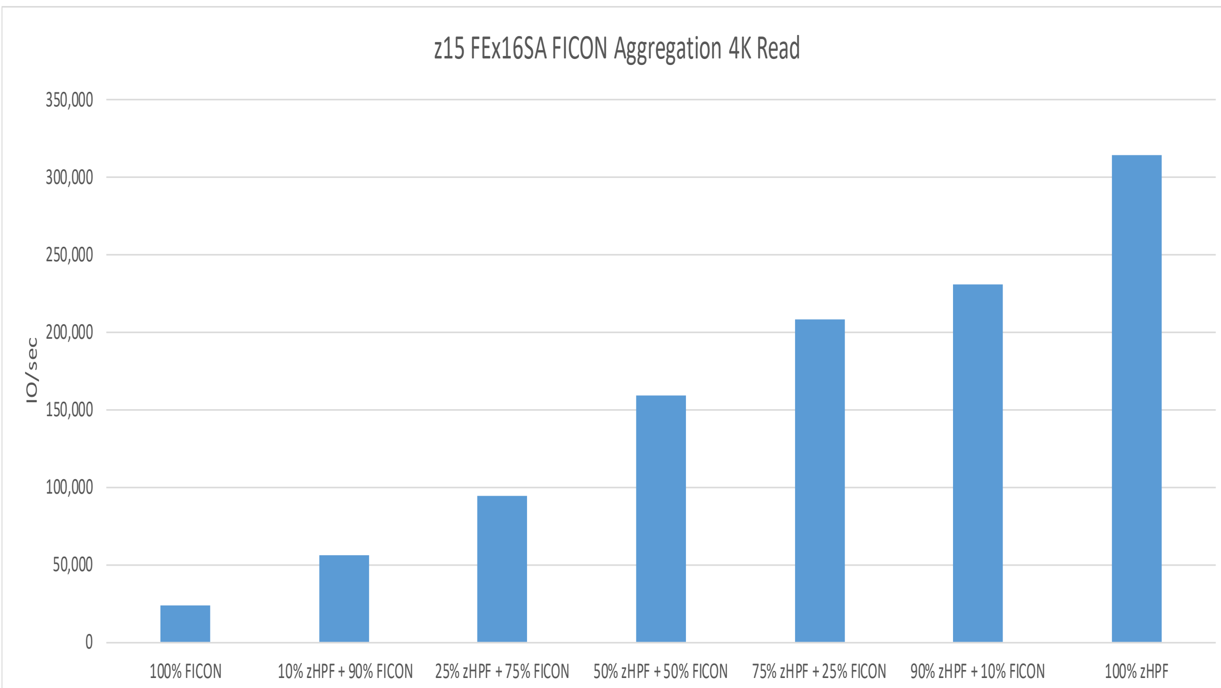

Figure 8 – FICON Aggregation

As shown in **Figure 9** below, using zHPF on a FEx16SA channel, a range from 1,200 MB/sec for a 1x4K block to over 3,200 READ/WRITE (mix) MB/sec when 16x4K (or larger) blocks are transferred per I/O operation was achieved. In contrast, using the FICON protocol with the MIDAW facility, a range from 100 MB/sec for a single 4K block to over 600 READ/WRITE (mix) MB/sec when 16 or more 4K blocks are transferred per I/O operation was measured on a FEx16SA channel. The maximum MB/sec possible with zHPF ranges from 4 times FICON for large data transfers (128K bytes) to 12 times FICON for smaller data transfers (8K bytes per I/O). The graph below demonstrates the efficiency of the zHPF protocol with the new FEx16SA channel, which increases the average bytes per frame and reduces the number of frames that is transferred per I/O.

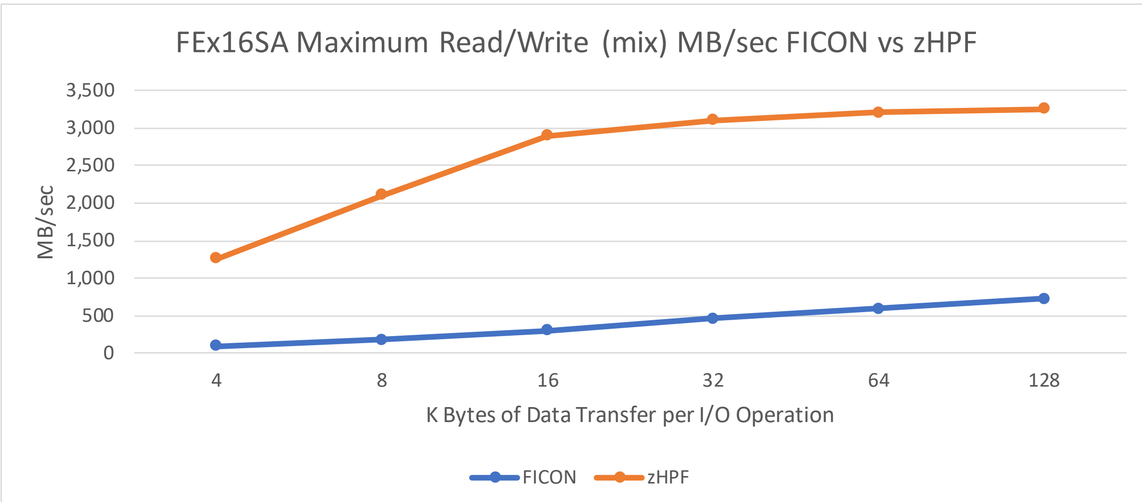

Figure 9 – FEx16SA Maximum MB/sec

Examples of the response time benefits of both the new FEx16SA channel and the efficiencies of the zHPF protocol are shown in **Figure 10** below. Simple benchmark experiments were conducted using a FEx16SA channel to READ or update WRITE 32 4K records or 128K bytes of data in total to a single device. PEND and CONN time response time components are displayed for zHPF compared to the FICON protocol. For I/O operations that READ 128K bytes of data, zHPF response times are 25% faster than FICON which is displayed in the CONN time component. For I/O operations that update WRITE 128K bytes of data, zHPF response times are 20% faster than FICON.

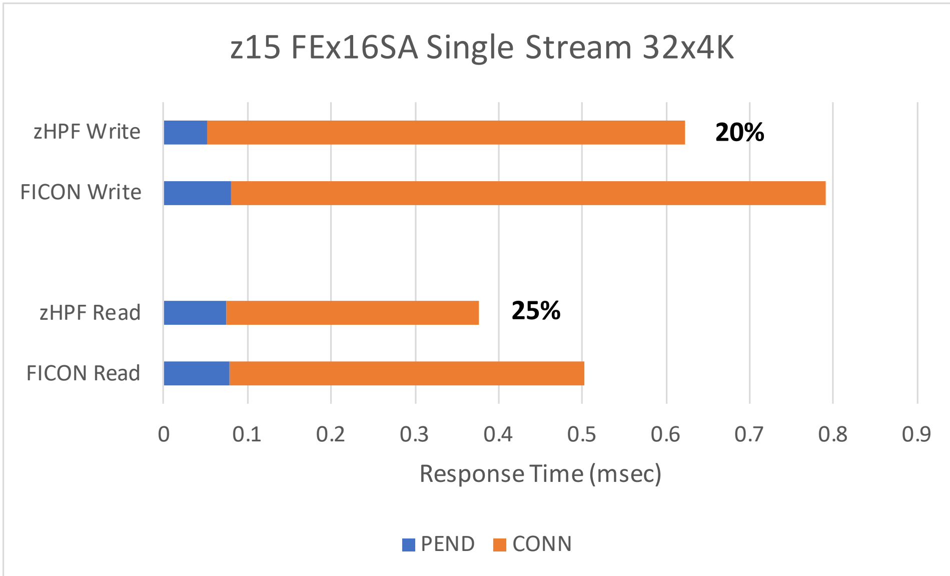

Figure 10 – FEx16SA Response Time

Benefits of both the new FEx16SA channel and the efficiencies of the zHPF protocol can also be observed in the channel processor utilization reported on the Resource Measurement Facility (RMF) Channel Activity Report. **Figure 11** below shows the FEx16SA channel processor utilization reported for both zHPF and FICON measurements for 32x4K READ I/O operations. Since the hardware data router is handling the movement of data for zHPF, the maximum FEx16SA channel processor utilization reported for these large data transfer zHPF I/O operations is less than 10%. In contrast, the FEx16SA channel microprocessor is involved in processing both the Channel Command Words (CCWs) and the transfer of each of the 8k byte data information units (IUs). The channel utilization reported for equivalent amounts of MB/sec processed is more than 20 times what is used for zHPF large data transfers. For large data transfer I/O operations using the zHPF protocol it is therefore more appropriate to observe either channel bus or link utilizations.

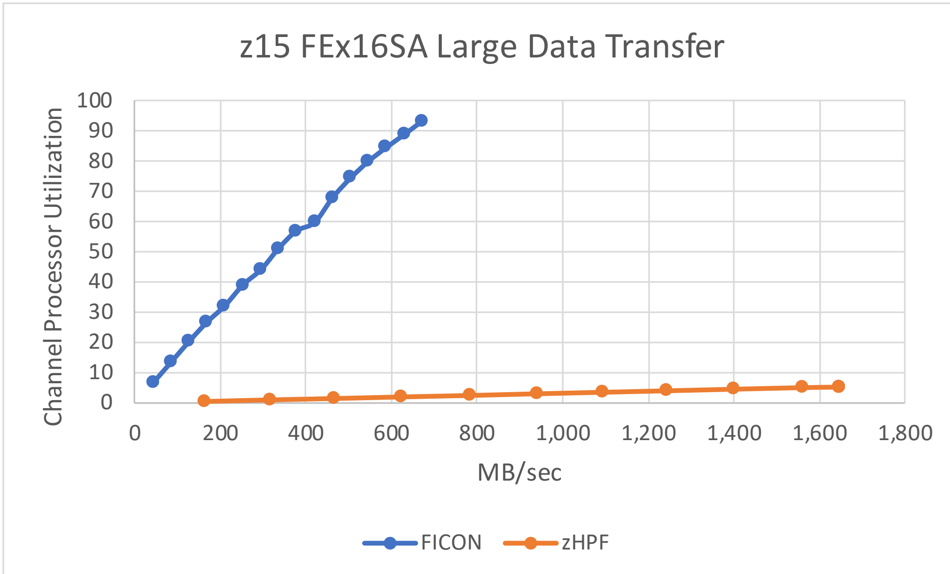

Figure 11 – FEx16SA Channel Processor Utilization

In general, IBM recommends the following guidelines to achieve good response times:

- keep Channel Processor Utilizations less than 50%
- keep Channel BUS Utilizations less than 70%

```
                                         C H A N N E L   P A T H   A C T I V I T Y
                                                                                                           PAGE    1
        z/OS V2R3            SYSTEM ID P34            DATE 08/15/2019            INTERVAL 00.59.999
                             RPT VERSION V2R3 RMF     TIME 09.39.44              CYCLE 1.000 SECONDS

IODF = CF   CR-DATE: 08/15/2019   CR-TIME: 08.55.18   ACT: POR        MODE: LPAR       CPMF: EXTENDED MODE          CSSID: 0
--------------------------------------------------------------------------------------------------------------------------
                                             DETAILS FOR ALL CHANNELS
--------------------------------------------------------------------------------------------------------------------------
   CHANNEL PATH          UTILIZATION(%)    READ(MB/SEC) WRITE(MB/SEC)      FICON OPERATIONS        ZHPF OPERATIONS
ID TYPE   G  SPEED SHR   PART  TOTAL   BUS  PART  TOTAL   PART  TOTAL    RATE  ACTIVE   DEFER    RATE ACTIVE  DEFER
42 FC_S  27    16G  Y    2.64   2.64 26.86 429.93 429.93 429.65 429.65    0.0     0.0     0.0  6558.1    3.9    0.0
4C FC_S  26     8G  Y    2.51   2.51 26.86 429.90 429.90 429.70 429.70    0.0     0.0     0.0  6558.2    4.2    0.0
61 FC_S  21    16G  Y    5.11   5.11 53.73 859.59 859.59 859.61 859.61    0.0     0.0     0.0   13117    6.8    0.0
6D FC_S  20     8G  Y    4.87   4.87 50.77 819.57 819.57 805.12 805.12    0.0     0.0     0.0   12395   31.5    0.0
94 FC_S  17    16G  Y   19.18  19.18 68.77 859.61 859.61 859.61 859.61    0.0     0.0     0.0   13117    7.3    0.0
95 FC_S  16     8G  Y   17.02  17.02 64.50 819.77 819.77 792.70 792.70    0.0     0.0     0.0   12302   31.1    0.0
A3 FC_S  13     8G  Y   10.94  10.94 53.72 429.48 429.48 430.12 430.12    0.0     0.0     0.0  6558.2    5.1    0.0
******************************************************** Bottom of Data ********************************************
```

Figure 12 – FEx16SA Channel Path Activity

When using either the FICON or the zHPF protocol, the FEx16SA channel processor utilization is the most important factor to observe for small data transfers. The link or bus

utilization is a more important factor to observe for large data transfers using the zHPF protocol. For large data transfers using the FICON protocol, both the channel processor and link utilization guidelines should be observed.

In summary, the FEx16SA channel offers performance improvements in both response times and maximum throughput for IO/sec and MB/sec for workloads using the zHPF protocol.

FEx16SA FCP Performance

FICON Express hardware supports the Fibre Channel Protocol (FCP) when defined as CHPID type FCP, allowing z/VM, z/VSE, KVM, and Linux on IBM Z LPARs to connect to industry-standard Small Computer System Interface (SCSI) storage controllers and devices.

The FEx16SA channel operates at the same 16Gbps line rate as the previous generation FEx16S+ channel. Therefore, the FEx16SA channel provides similar throughput performance when compared to the FEx16S+ channel for FCP protocol. With support for Fibre Channel Endpoint Security encryption enabled, the FEx16SA channel has the added benefit of Encryption of Data In Flight for FCP operations with less than 4% impact on maximum channel throughput.

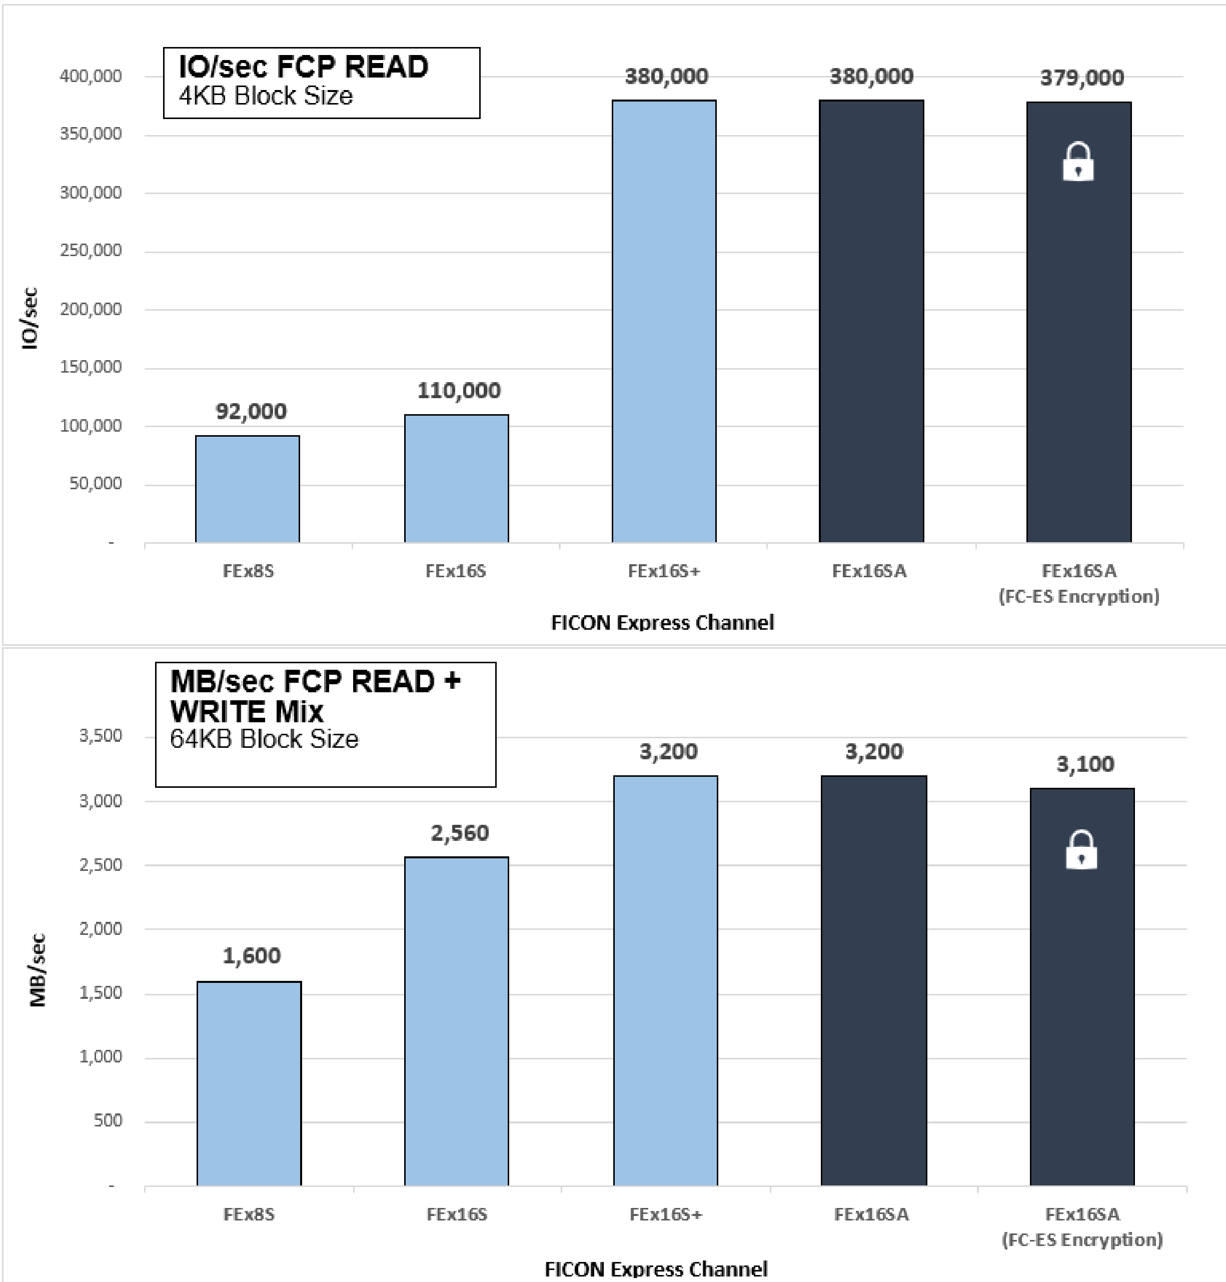

Figure 13 – Maximum FEx16SA Throughput

The results in **Figure 13** represent the maximum single channel throughput capacity for two benchmarks, executed across 4 generations of channel technology. These experiments were executed on a Linux on IBM Z LPAR using a single channel attached to multiple host adapter ports on a IBM System Storage DS8950.

FCP small block performance was measured by driving the channel with multiple concurrent 4KB READ operations, while large block throughput was determined using a mixture of 64KB READ + WRITE operations. These results show that a FICON Express16SA channel, operating with the FCP protocol and Fibre Channel Endpoint Security encryption, can execute up to 379K small block READ IO/sec and process up to a 3,100 MB/s mix of large block operations. Meanwhile, the same channel without Fibre Channel Endpoint Security encryption supports up to 380K IO/sec and up to 3,200 MB/s, comparable to the previous generation FEx16S+ channel. This represents less than 4% impact on maximum channel throughput for the added benefit of securing all FCP traffic between the FEx16SA channel and a Fibre Channel Endpoint Security enabled DS8900 Storage link.

Furthermore, since data protection is handled by channel hardware, the LPAR CPU or IFL is not utilized for encryption processing. **Figure 14** highlights the average IFL utilization of a 12 core LPAR, running a mixture of READ + WRITE FCP operations at various traffic load levels. Overall IFL busy was measured to be comparable when enabling and disabling Fibre Channel Endpoint Security.

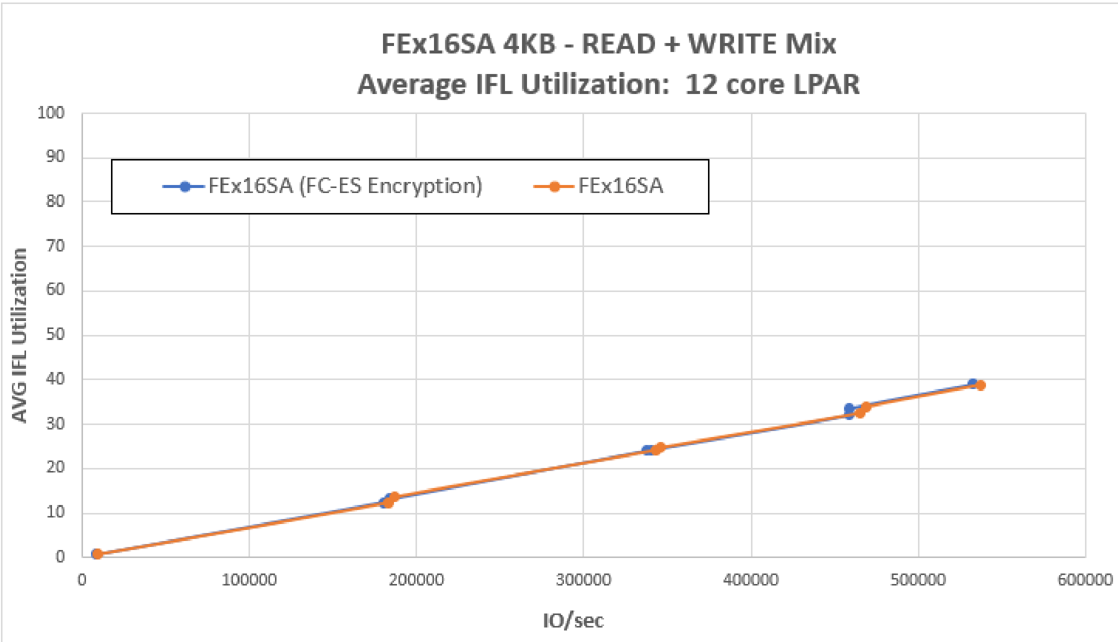

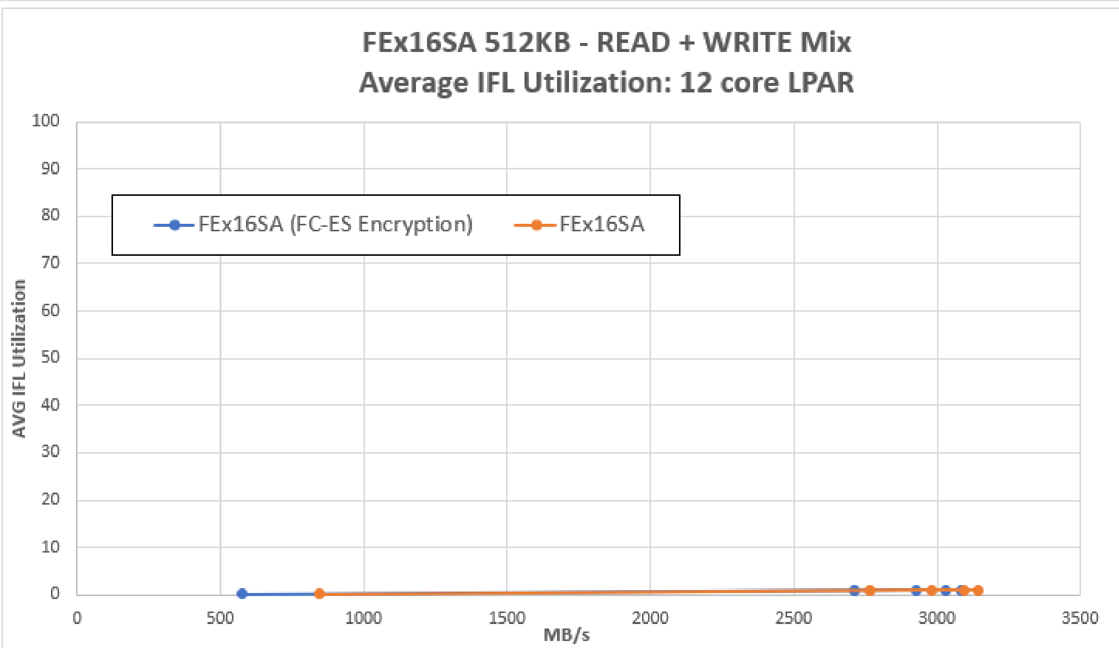

Figure 14 – Average IFL Utilization for 12 Core LPAR

Looking at the effects of Fibre Channel Endpoint Security encryption more closely in **Figure 15**, we see FEx16SA perform well when handling varying levels of FCP traffic. Using the same benchmark setup and increasing the number of FCP operations from a single I/O operation at a time up to 336 concurrent operations, the FEx16SA channel with FICON Express

Endpoint Security encryption was able to process an overall small block IO/sec load within 1% - 4% of the results seen without encryption enabled. This was consistent for every point tested along the curve with READ only, WRITE only and READ/WRITE (mix) operations.

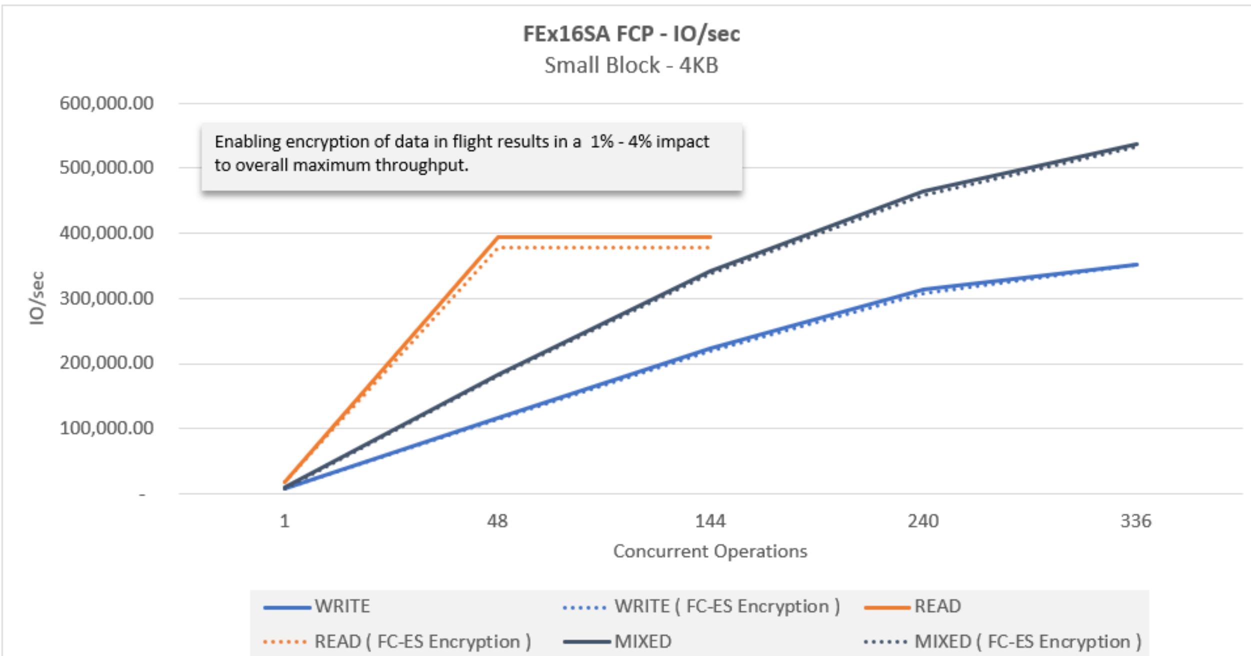

Figure 15 – FCP Small Block Performance Curve

On the other hand, as the size of the data block increases, some FCP I/O operations may see lower throughput when enabling encryption for workloads running with limited concurrency or parallelism. In the example included in **Figure 16**, produced using 64KB block sizes, the FEx16SA channel with Fibre Channel Endpoint Security encryption was able to write approx. 230 MB/s when the driver program issued a single FCP operation. Compared with the 300 MB/s measured when data protection is disabled and running the same WRITE operation, this represents a 23% drop in throughput. However, notice that as the number of concurrent operations increases, and subsequently the traffic load handled by the channel, the throughput delta is overcome, and overall channel throughput falls inline within 1% - 4% of the bandwidth seen when protection is disabled. READ operations in this case saw a smaller impact to single IO operations, only displaying a 2% impact on throughput.

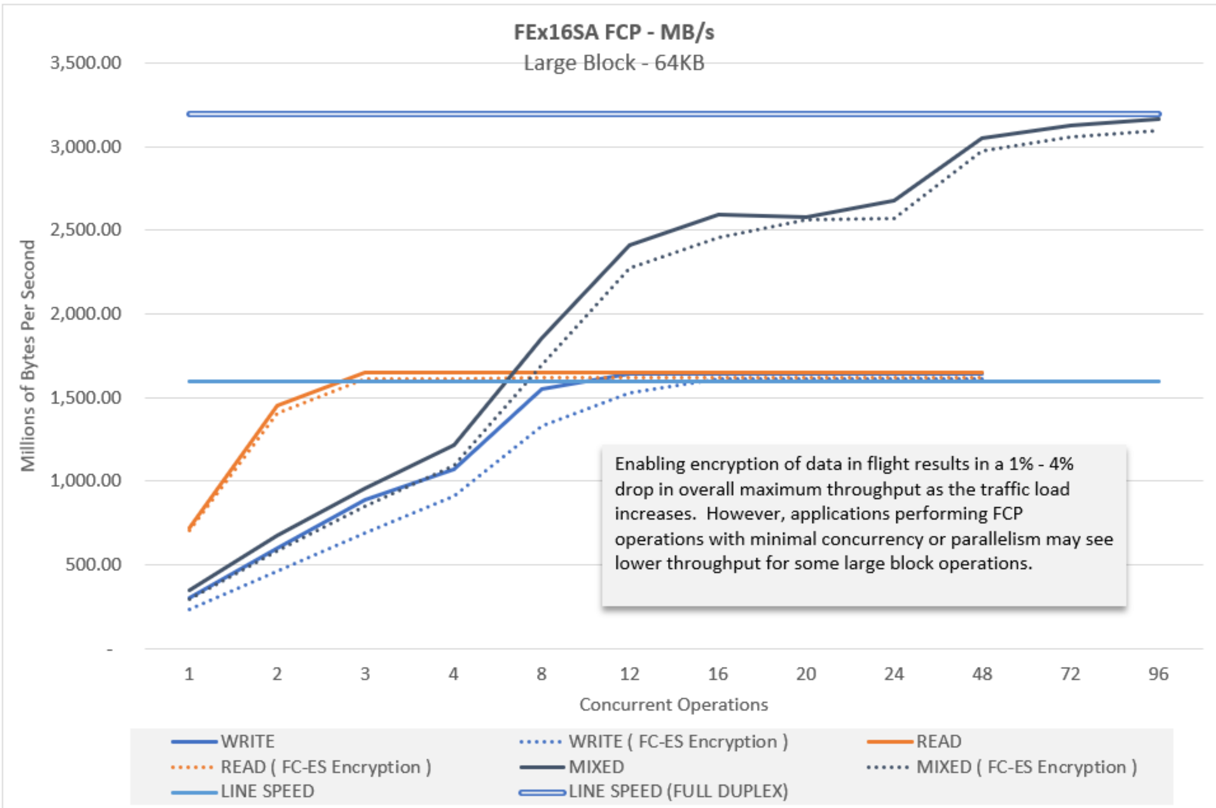

Figure 16 – FCP Large Block Performance Curve

FEx16SA channel performance using the FCP Hardware Data Router

The FEx16SA FCP channel continues to support the FCP Hardware Data Router introduced several generations ago with the FEx8S channel. This feature offers improved throughput and I/O processing capacity, in addition to reduced response times, by:

- providing a shorter path when transferring data from system memory
- reducing the channel's internal processing costs.

In fact, the FCP Hardware Data Router is required to achieve the maximum throughput numbers highlighted in the previous section.

As shown in **Figure 17**, FEx16SA small block performance greatly benefits from the use of the FCP Hardware Data Router, lowering the average I/O response times and overall IO/sec achieved with the channel. In this case, a curve was built by driving a FEx16SA with an increasing number of concurrent 4KB read operations:

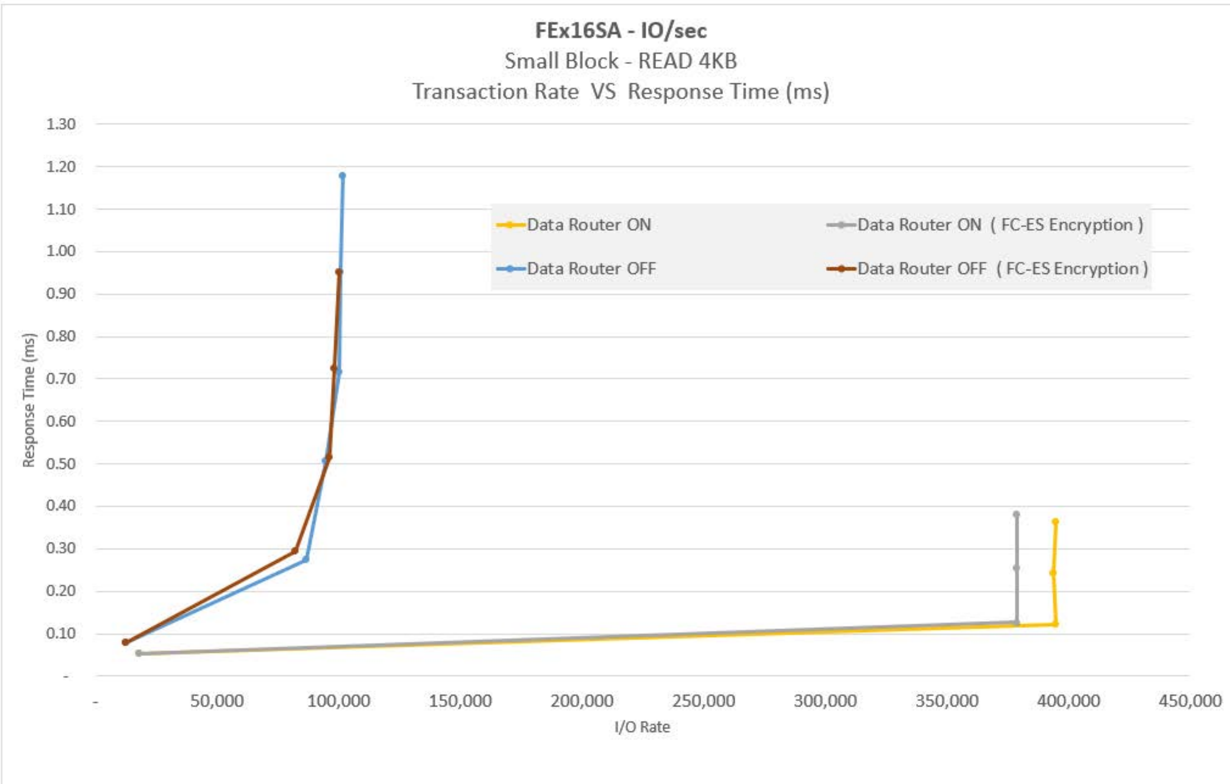

Figure 17 – FEx16SA Data Router Small Block I/O Rate VS Response Time

When the FCP Hardware Data Router is enabled, the FEx16SA channel can process nearly 3.8 times the maximum 4KB READ IO/sec that can be done without it. Latency is reduced as well, which can result in 30% lower response times when executing a single FCP 4KB READ operation or approximately up to a 25us reduction. Furthermore, as the amount of work increases in the channel, I/O response times increase at a lower rate with the Hardware Data Router, only showing a sharp increase once we reached the maximum capacity of the channel.

Additionally, the FCP Hardware Data Router benefits large block performance, increasing the overall bandwidth supported by the FEx16SA channel as shown on **Figure 18**. Running a mixture of 64KB FCP READ/WRITE (mix) operations results in a maximum full duplex throughput of 3,200 MB/sec, which is over twice the maximum achieved when the Hardware Data Router is not in use. Maximum throughput when executing READ only or WRITE only operations has a smaller gain of less than 10%, reaching over 1,600 MB/s (line speed).

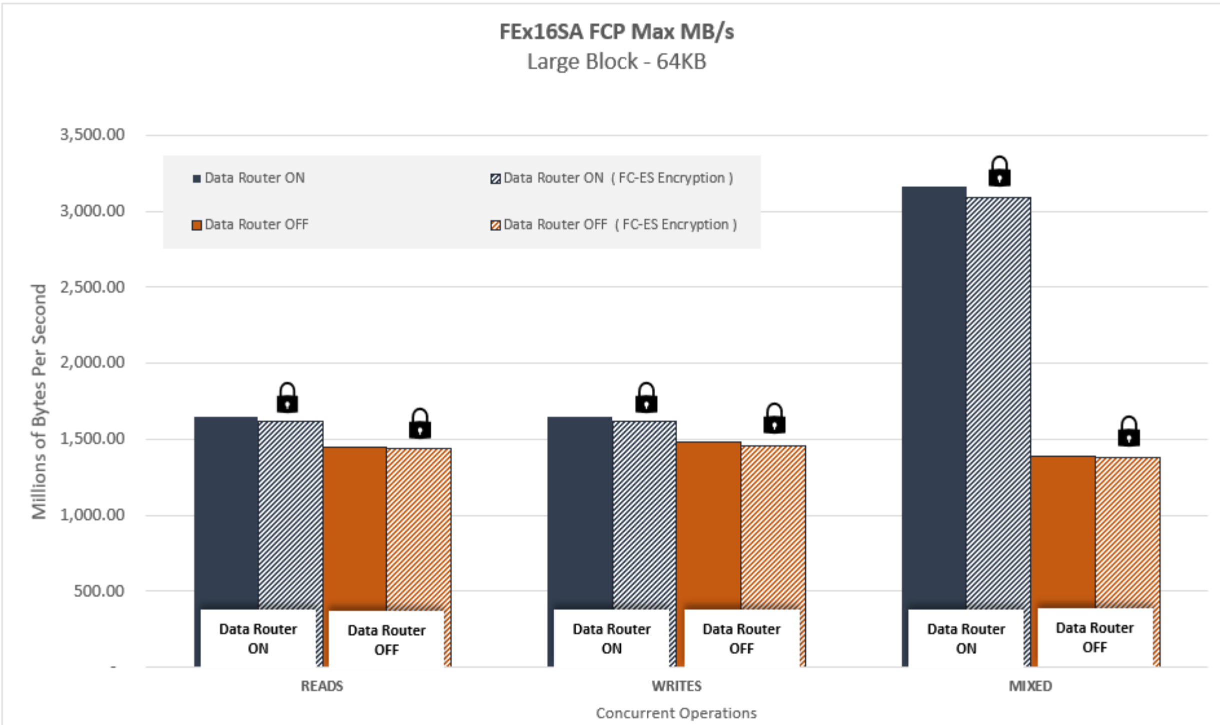

Figure 18 – FEx16SA Data Router Large Block Throughput

To make use of the Hardware Data Router, and start benefiting from this feature, the following requirements must be met:

FCP Hardware Data Router Minimum Requirements

- zEnterprise 196 GA2 and zEnterprise 114 with FEx8S.
- z/VM V6.3 for guest exploitation.
- Linux on IBM Z:
  – SLES 12 and SLES 11 SP3.
  – RHEL 7 and RHEL 6.4.
  – Ubuntu 16.04 LTS (or higher).
- KVM hypervisor which is offered with the following Linux distributions: SLES-12 SP2 or higher, and Ubuntu 16.04 LTS or higher. For minimal and recommended distribution levels refer to the IBM Z website.

Enabling the FCP Hardware Data Router

The use of the FCP hardware data router can be manually enabled / disabled by the user in a Linux on IBM Z:

- By modifying a kernel parameter:

  ```
  ON:   'zfcp.datarouter=1'
  OFF:  'zfcp.datarouter=0'
  ```

- Or by passing an additional parameter when loading the zfcp driver:

  ```
  ON:    modprobe zfcp datarouter=1
  OFF:   modprobe zfcp datarouter=0
  ```

To verify the status of the FCP data router, the following commands can be used:

```
# systool -m zfcp -v
Module = "zfcp"

  Attributes:
    initstate           = "live"
    refcnt              = "48"
    srcversion          = "52F4E65D11CC52E53FD2D0D"
    supported           = "Yes"

  Parameters:
    ...
    datarouter          = "Y"
    ...
```

Or

```
# cat /sys/module/zfcp/parameters/datarouter
Y
```

For additional information, see the '*Linux on IBM Z and LinuxONE – FCP Hardware Data Router support*' online document: https://www.ibm.com/support/knowledgecenter/en/linuxonibm/liaag/wecm/l0wecm00_fcp_hw_data_router_support.htm

Commercial Batch Workload Performance

To evaluate the effect of encryption enabled in the IBM Fibre Channel Endpoint solution we tested two different workload types. The first workload, known as 'Commercial Batch-Long' (CB-L), is a commercial batch job stream reflective of large batch jobs with fairly heavy CPU processing. The workload consists of a mix of Commercial Batch Jobs. The batch jobs include various combinations of C, COBOL, FORTRAN, and PL/I compile, link-edit, and execute steps. Sorting, DFSMS utilities (e.g. dump/restore and IEBCOPY), VSAM and DB2 utilities, SQL processing, GDDM® graphics, and FORTRAN engineering/scientific subroutine library processing are also included.

The results of running the CB-L workload are shown below in **Figure 19**. Execution of the CB-L workload with and without Fibre Channel Endpoint Security encryption enabled demonstrated that Fibre Channel Endpoint Security encryption had no significant impact on either the elapsed time, CPU consumption, nor the I/O completion rate.

Fibre Channel Endpoint Security	CPU%	I/O Interrupt Rate	Elapsed Time
Disabled	91.72	89,886.62	544.30
Enabled	91.74	89,817.79	544.16

Figure 19 – CB-L - Fibre Channel Endpoint Security

Commercial Batch-Short (CB-S) is an I/O intensive batch workload which utilizes BSAM and QSAM access methods and was the second workload used to measure the effects of Fibre Channel Endpoint Security encryption being enabled. CB-S consists of 18 jobs that repetitively copy data between input datasets and output datasets using various IBM standard utilities (e.g. IEBGENER, IEBDG, DFSORT, IFASMFDP, IDCAMS Repro and IEBCOMPR). For example, the DFSORT job splits one input dataset into three output datasets, and in the next step merges the three created datasets back into a single one. This workload has been used in the past to measure the effects of IBM Z asynchronous compression, so here it was employed to test the effects of Fibre Channel Endpoint Security encryption without compression enabled, and then the test was repeated to measure the effects of Fibre Channel Endpoint Security encryption in conjunction with IBM Z Integrated Accelerator for zEDC **Figure 20**.

		I/O Interrupt	
Fibre Channel Endpoint Secuity	**CPU%**	**Rate**	**Elapsed Time**
Disabled Without Compression	23.86	11,443	675.90
Enabled Without Compression	24.01	11,338	675.80
		I/O Interrupt	
Fibre Channel Endpoint Secuity	**CPU%**	**Rate**	**Elapsed Time**
Disabled with Integrated Accelerator for zEDC	23.54	6,480	555.50
Enabled with Integrated Accelerator for zEDC	23.22	6,369	565.52

Figure 20 – CB-S - Fibre Channel Endpoint Security

In both tests it was demonstrated that enabling Fibre Channel Endpoint Security encryption had negligible impact on CPU consumption, elapsed job execution time and I/O completion rate.

There is a new field within the DISPLAY M command to display the status of device indicating Connection Security for Fibre Channel Endpoint Security encryption.

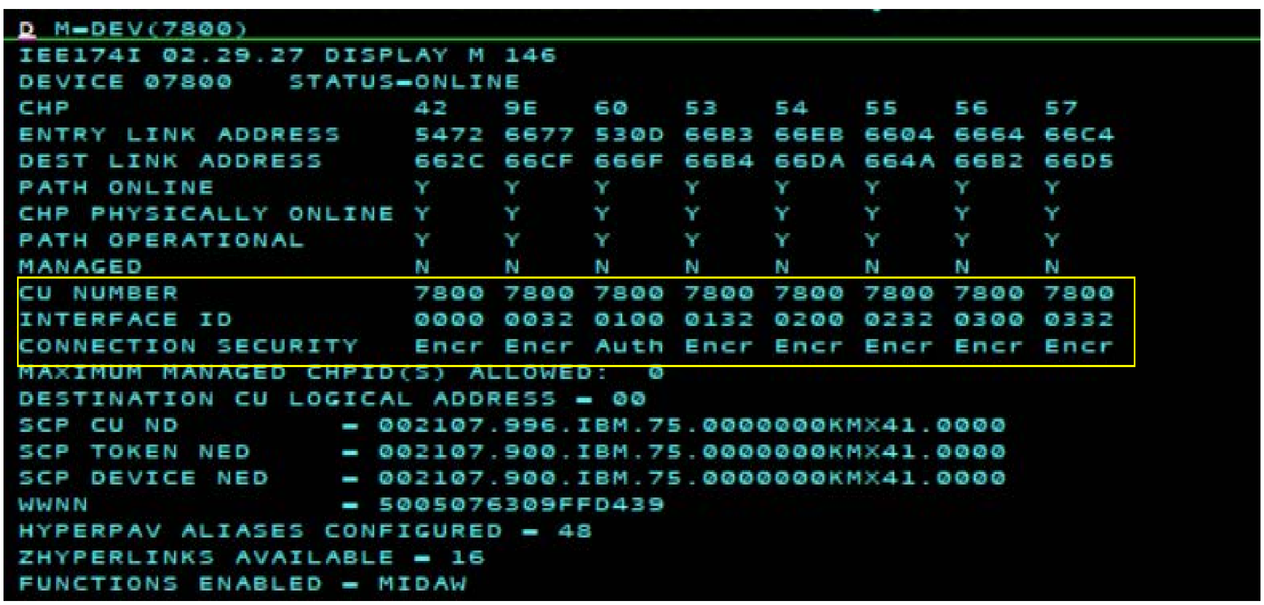

```
D M=DEV(7800)
IEE174I 02.29.27 DISPLAY M 146
DEVICE 07800   STATUS=ONLINE
CHP                      42   9E   60   53   54   55   56   57
ENTRY LINK ADDRESS       5472 6677 530D 66B3 66EB 6604 6664 66C4
DEST LINK ADDRESS        662C 66CF 666F 66B4 66DA 664A 66B2 66D5
PATH ONLINE              Y    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                7800 7800 7800 7800 7800 7800 7800 7800
INTERFACE ID             0000 0032 0100 0132 0200 0232 0300 0332
CONNECTION SECURITY      Encr Encr Auth Encr Encr Encr Encr Encr
MAXIMUM MANAGED CHPID(S) ALLOWED:  0
DESTINATION CU LOGICAL ADDRESS = 00
SCP CU ND          = 002107.996.IBM.75.0000000KMX41.0000
SCP TOKEN NED      = 002107.900.IBM.75.0000000KMX41.0000
SCP DEVICE NED     = 002107.900.IBM.75.0000000KMX41.0000
WWNN               = 5005076309FFD439
HYPERPAV ALIASES CONFIGURED = 48
ZHYPERLINKS AVAILABLE = 16
FUNCTIONS ENABLED = MIDAW
```

Figure 21 – Device Display for Fibre Channel Endpoint Security

The DISPLAY M command for channel path also indicates when Fibre Channel Endpoint Security encryption is enabled.

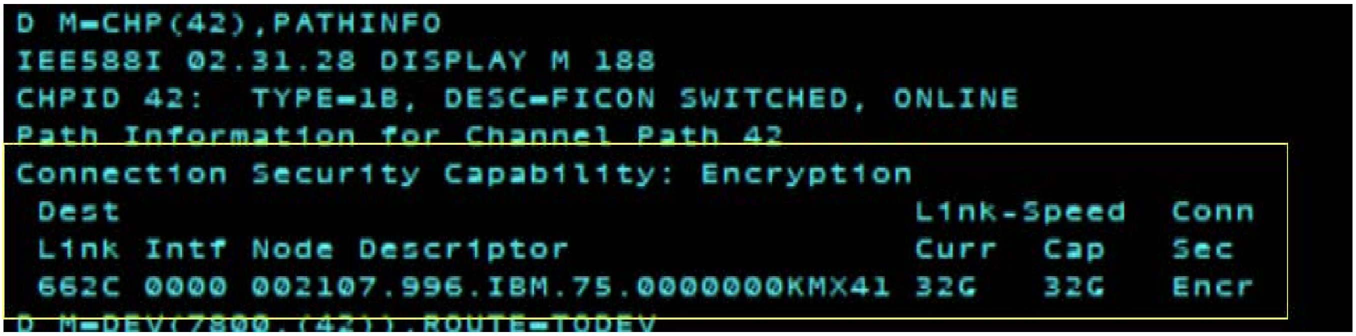

```
D M=CHP(42),PATHINFO
IEE588I 02.31.28 DISPLAY M 188
CHPID 42:  TYPE=1B, DESC=FICON SWITCHED, ONLINE
Path Information for Channel Path 42
Connection Security Capability: Encryption
 Dest                                        Link-Speed   Conn
 Link Intf Node Descriptor                   Curr   Cap   Sec
 662C 0000 002107.996.IBM.75.0000000KMX41    32G    32G   Encr
```

Figure 22 – CHPID Display for Fibre Channel Endpoint Security

FEx16SA Channel Performance using zHPF and FICON Protocols at up to 100km of Distance

The FEx16SA channel has 90 buffer credits for each individual channel port which is enough to support a 16 Gbps link speed at distances up to 10 km between each channel and its nearest neighbor, which could be a port on a director or tape or disk storage subsystem.

To evaluate FEx16SA channel performance using zHPF and FICON protocols at 100 km of distance, using Dense Wavelength Division Multiplexer (DWDM). A series of performance measurements were conducted with multiple FEx16SA channels on z15 connected to multiple storage subsystems using cascaded FICON directors with a 16 Gbps Inter Switch Link (ISL) between the two directors extended over a Dense Wavelength Division Multiplexer (DWDM) as illustrated in **Figure 23** below.

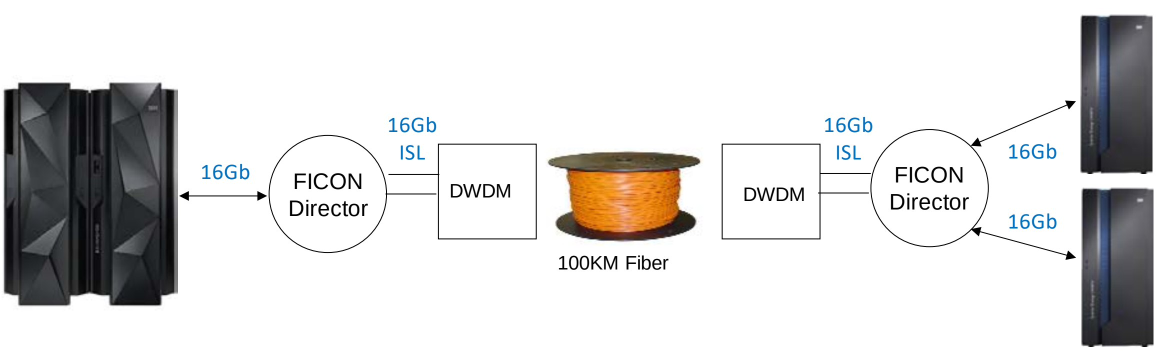

Figure 23 – Distance Diagram

The amount of throughput that can be achieved on an ISL at distances up to 100 km are dependent on a number of variables:

- Workload characteristics such as data transfer sizes and mix of READ and WRITE activity
- zHPF vs FICON protocols
- Link speed and the bit encoding scheme used on the ISL
- Buffer credits available on the ISL ports on each of the FICON directors
- Actual distance between the two ISL ports on each of the FICON directors

Figure 24 below displays the maximum MB/sec achieved on 16 Gbps ISLs between two FICON directors separated by 100 km DWDM with both 720 and 960 Buffer-to-Buffer (B2B) credits available on each of the ISL ports. Previously, the general recommendation was to use 0.5 buffers for each 1 km of distance and each 1 Gbps of link speed. This implies 50, 100, 200, 400 and 500 B2B credits for 1, 2, 4, 8 and 10 Gbps link speeds respectively. However, in contrast to the 1, 2, 4 and 8 Gbps FICON link speeds which use 8/10 bit encoding schemes, the 16 Gbps ISL link uses a 64/66 bit encoding scheme. This means that in contrast to the 1, 2, 4 and 8 Gbps FICON channel link speeds which are capable of a maximum of approximately 100, 200, 400 and 800 MB/sec, a 16 Gbps ISL link is capable of a maximum of approximately 1600 MB/sec in each direction. These measurements used 960 B2B credits to ensure the maximum was achieved.

For 128K byte READ or update WRITE I/O operations with 960 B2B credits, a maximum of 3,170 MB/sec was achieved using zHPF protocols compared to 2,850 MB/sec with 720 B2B credits. For a 50/50 mix of 128k byte READ and WRITE I/O operations with 960 B2B credits, a maximum of 5,920 MB/sec was achieved using zHPF protocols compared to 5,410 MB/sec with 720 B2B credits.

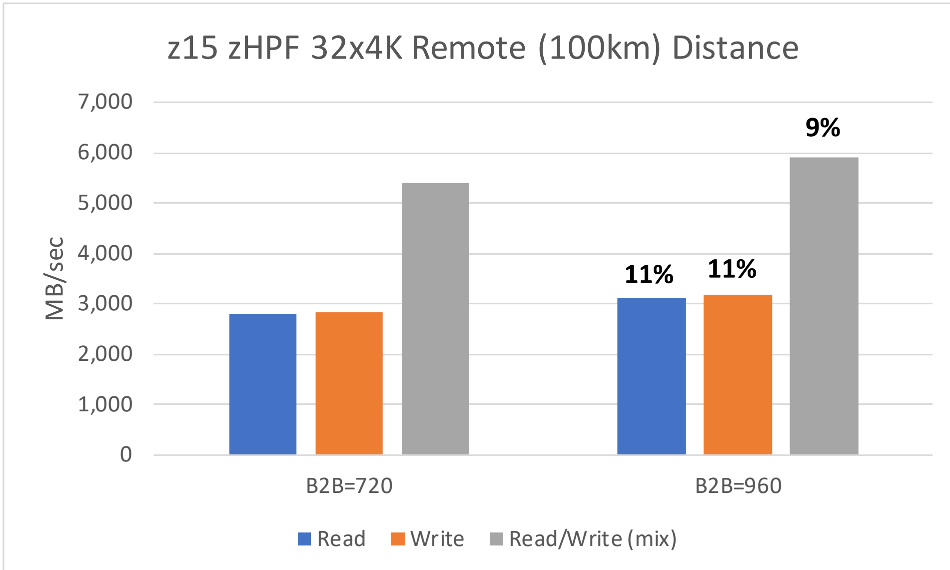

Figure 24 – Distance zHPF Large Data Transfer

Figure 25 below summarizes the Average Frame Size for zHPF for small, medium, and large data transfer sizes. For example, when the frames from a 50/50 mix of 4K byte READ and WRITE I/O operations are flowing across an ISL link, the average bytes per frame using zHPF protocols is 1080 bytes. For READ I/O operations, the average frame size shown is for the data frames flowing from the ISL port adjacent to the storage subsystem to the ISL port closest to the z15. For update WRITE I/O operations, the average frame size shown is for the data frames flowing to the ISL port adjacent to the storage subsystem from the ISL port closest to the z15.

Data Transfer	zHPF Average Frame Size
4K Read or Write	1083
4K Read/Write (mix)	886
32x4K Read or Write	2023
32x4K Read/Write (mix)	1995
128x4K Read or Write	2068
128x4K Read/Write (mix)	2061

Figure 25 - Average Frame Sizes

From a response time perspective, most workloads will experience a 1 millisecond (ms) adder at 100 km of distance compared to local response times due to the speed of light through a fiber. **Figure 26** below summarizes this 1ms adder to response times at reasonable levels of activity for READ I/O operations with small data transfers (4K bytes per I/O), large data transfers (32x4K or 128K bytes per I/O) and very large data transfers (128x4K or approximately 512K bytes per I/O) at 100 km compared to Local distance.

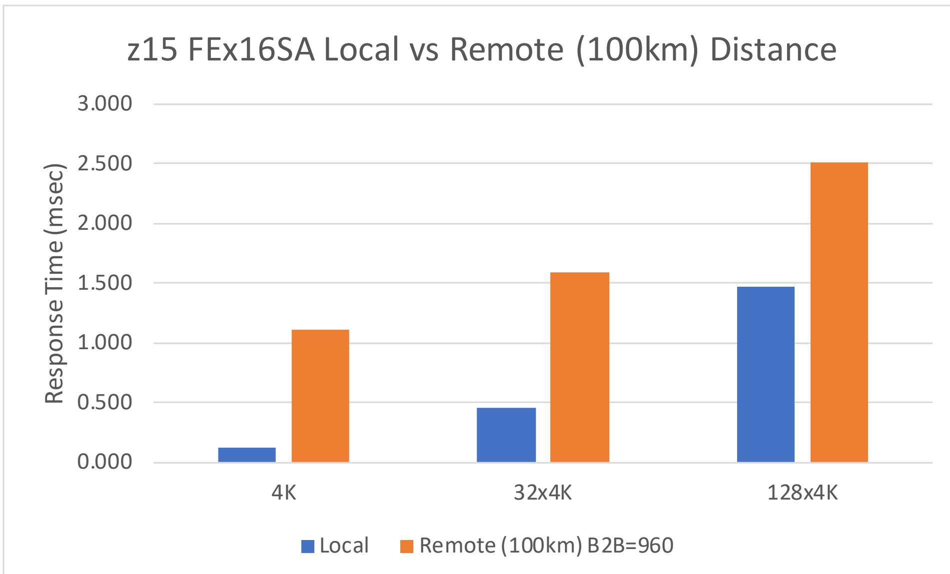

Figure 26 – Local vs Remote Distance

zHyperLink Express

zHyperLink is an ultra-low latency, synchronous I/O interface connecting the IBM System Z servers to the IBM System DS8000 storage series. The zHyperLink Express feature was first introduced on the IBM z14 system and the DS8880 storage family. zHyperLink support continues with the introduction of the new IBM z15 Server and the DS8900F storage family. This low latency interface provides a point-to-point connection, up to a maximum cable length of 150 meters, between the Z server and the DS8000 using PCIe Gen3 x8 infrastructure. To accommodate this new transport, a new Synchronous I/O command set has been defined to allow the operating system to synchronously read and write data. Finally, the optical MPO cables and transceivers are the same as those defined for the IBM Sysplex Coupling Facility links.

Traditional communication between servers and storage control units occur over a Storage Area Network requiring Host Bus Adapters (HBAs) with their associated device drivers, one or more switches and multiple layers of a protocol stack. Examples are Fibre Connectivity (FICON) and High Performance FICON (zHPF). These protocols with their associated overhead, use asynchronous I/O access methods such as the z/Architecture Start Subchannel instruction with its associated asynchronous completion events. The large variation in access times, and even the minimum access times, of SAN storage with protocols such as Fibre Channel, require asynchronous access to avoid unacceptable blocking of CPU processing. Alternatively, a synchronous I/O access method may be used and can provide significant performance advantages. Synchronous access causes a software thread to be blocked while waiting for the I/O to complete, but avoids the operating system overhead of scheduling, context switches, processor cache pollution and interrupts. zHyperLink, Express Synchronous I/O Paradigm is illustrated in **Figure 27** below.

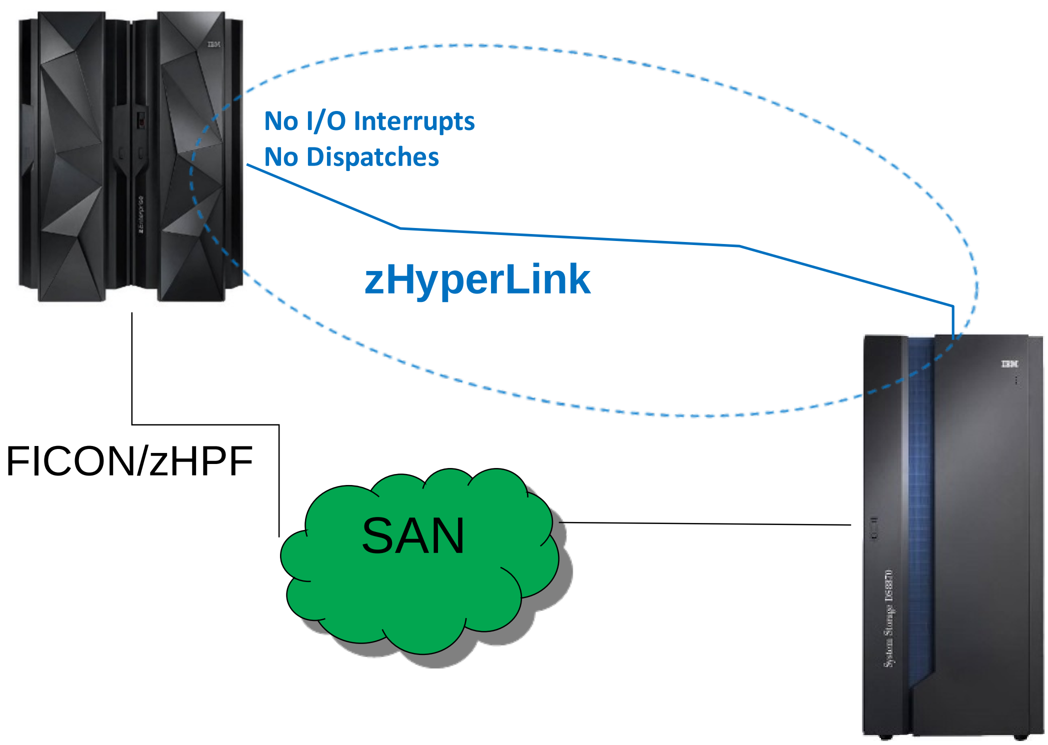

Figure 27 – zHyperLink Diagram

Latency is the most important performance metric for storage access. The zHyperlink low latency interface provides the opportunity for the operating system to read and write data synchronously, thus avoiding the scheduling and interrupt overhead associated with asynchronous IO operations. With 19 microseconds latency for IBM Z servers with zHyperLink, the DS8900F family provides ultra-low latency to help clients process huge volumes of transactions faster and deliver real-time insights to differentiate their products and gain a competitive advantage.

As illustrated in **Figure 28** using an IBM z15 zHyperLink Express attached to a DS8950F, Sigle-Stream 4 KB Read and Write link latency is reduced by up to 4X as compared to z15 FEx16SA attached to a DS8950F with a single 32 GFC Host Adapter.

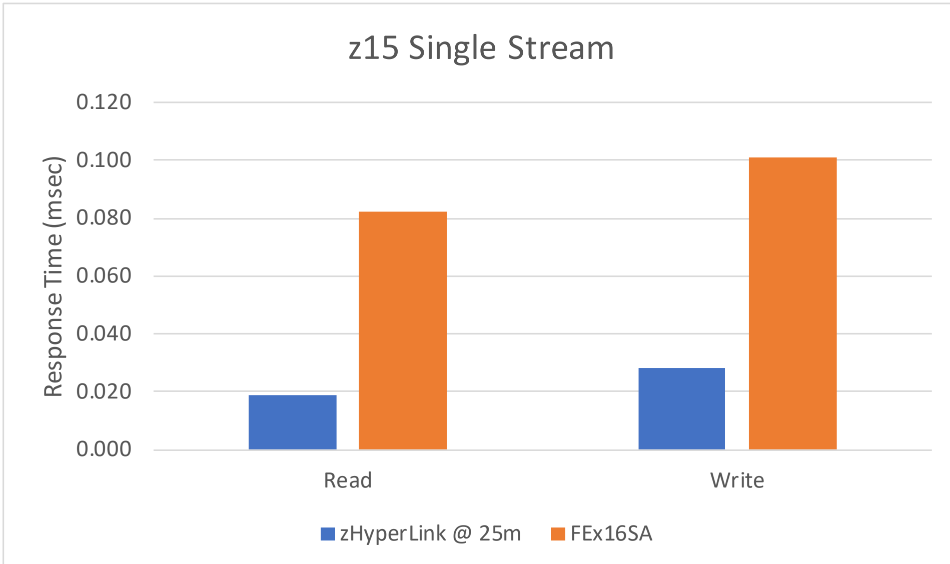

Figure 28 – FEx16SA vs zHyperLink Express

For additional information on zHyperLink Express please refer to the following guide: http://www.redbooks.ibm.com/Redbooks.nsf/RedbookAbstracts/redp5493.html?Open

RMF Channel Activity report

There are six fields on the RMF Channel Activity report that can be used to distinguish between FICON and zHPF traffic.

- The RATE field refers to the number of FICON or zHPF I/Os per second initiated at the total physical channel level (not by LPAR)
- The ACTIVE field refers to what we have previously called the "open exchanges", i.e. the number of I/Os that are simultaneously active within a channel.
- The DEFER field refers to the number of deferred FICON or zHPF IO/sec. This is the number of operations that could not be immediately initiated by the channel due to a temporary lack of resources.

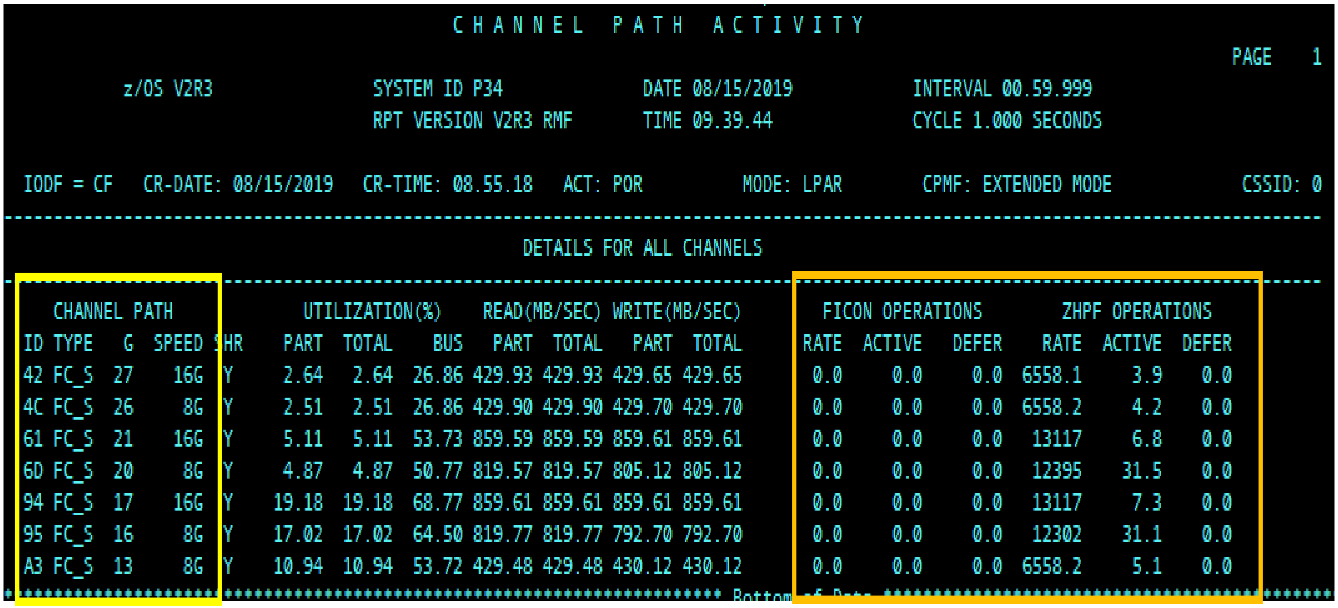

CHANNEL PATH ACTIVITY

PAGE 1

z/OS V2R3 SYSTEM ID P34 DATE 08/15/2019 INTERVAL 00.59.999

RPT VERSION V2R3 RMF TIME 09.39.44 CYCLE 1.000 SECONDS

IODF = CF CR-DATE: 08/15/2019 CR-TIME: 08.55.18 ACT: POR MODE: LPAR CPMF: EXTENDED MODE CSSID: 0

DETAILS FOR ALL CHANNELS

CHANNEL PATH					UTILIZATION(%)			READ(MB/SEC)		WRITE(MB/SEC)		FICON OPERATIONS			ZHPF OPERATIONS		
ID	TYPE	G	SPEED	SHR	PART	TOTAL	BUS	PART	TOTAL	PART	TOTAL	RATE	ACTIVE	DEFER	RATE	ACTIVE	DEFER
42	FC_S	27	16G	Y	2.64	2.64	26.86	429.93	429.93	429.65	429.65	0.0	0.0	0.0	6558.1	3.9	0.0
4C	FC_S	26	8G	Y	2.51	2.51	26.86	429.90	429.90	429.70	429.70	0.0	0.0	0.0	6558.2	4.2	0.0
61	FC_S	21	16G	Y	5.11	5.11	53.73	859.59	859.59	859.61	859.61	0.0	0.0	0.0	13117	6.8	0.0
6D	FC_S	20	8G	Y	4.87	4.87	50.77	819.57	819.57	805.12	805.12	0.0	0.0	0.0	12395	31.5	0.0
94	FC_S	17	16G	Y	19.18	19.18	68.77	859.61	859.61	859.61	859.61	0.0	0.0	0.0	13117	7.3	0.0
95	FC_S	16	8G	Y	17.02	17.02	64.50	819.77	819.77	792.70	792.70	0.0	0.0	0.0	12302	31.1	0.0
A3	FC_S	13	8G	Y	10.94	10.94	53.72	429.48	429.48	430.12	430.12	0.0	0.0	0.0	6558.2	5.1	0.0

Figure 29 – RMF Diagram

The above example above **Figure 29** shows a mix of FICON and zHPF I/O operations on each channel as shown in the "G" and "SPEED" fields.

- CHPID 42 (**FEx16SA**), the channel link is operating at 16 Gbps
- CHPID 4C (**FEx16SA**), the channel link is operating at 8 Gbps
- CHPID 61 (**FEx16S+**), the channel link is operating at 16Gbps
- CHPID 6D (**FEx16S+**), the channel link is operating at 8Gbps
- CHPID 94 (**FEx16S**), the channel link is operating at 16Gbps
- CHPID 95 (**FEx16S**), the channel link is operating at 8Gbps
- CHPID A3 (**FEx8S**), the channel link is operating at 8Gbps

The FEx16SA channels support the following generation field ("G" column) of the RMF Channel Activity:

- "26" the channel link is operating at a speed of 8 Gbps
- "27" the channel link is operating at a speed of 16 Gbps
- "28" the channel link is operating at a speed of 16Gbps with FEC (Forward Error Correction)

The FEx16S+ channels support the following generation field ("G" column) of the RMF Channel Activity:

- "1F" the channel link is operating at a speed of 4 Gbps
- "20" the channel link is operating at a speed of 8 Gbps
- "21" the channel link is operating at a speed of 16 Gbps
- "22" the channel link is operating at a speed of 16Gbps with FEC (Forward Error Correction)

The FEx16S channels support the following generation field ("G" column) of the RMF Channel Activity:

- "15" the channel link is operating at a speed of 4 Gbps
- "16" the channel link is operating at a speed of 8 Gbps
- "17" the channel link is operating at a speed of 16 Gbps
- "18" the channel link is operating at a speed of 16Gbps with FEC (Forward Error Correction)

The FEx8S channels support the following generation field ("G" column) of the RMF Channel Activity:

- "11" the channel link is operating at a speed of 2 Gbps
- "12" the channel link is operating at a speed of 4 Gbps
- "13" the channel link is operating at a speed of 8 Gbps

For additional descriptions please refer to RMF User's Guide:
https://www.ibm.com/support/knowledgecenter/en/SSLTBW_2.3.0/com.ibm.zos.v2r3.erbb500/chan3.htm

RMF Synchronous I/O Device Activity

There is a section of the RMF Report for Synchronous I/O Device Activity report which can be used to distinguish between Synchronous and Asynchronous Device Activity.

- DIRECT ACCESS DEVICE ACTIVITY
  - Device Activity Rate, if Synchronous I/O Activity will be signified by “S”.
- SYNCRONOUS I/O DEVICE ACTIVITY
  - Synchronous I/O Device Activity

For additional descriptions please refer to RMF User’s Guide: https://www.ibm.com/support/knowledgecenter/en/SSLTBW_2.3.0/com.ibm.zos.v2r3.erbb500/device.htm

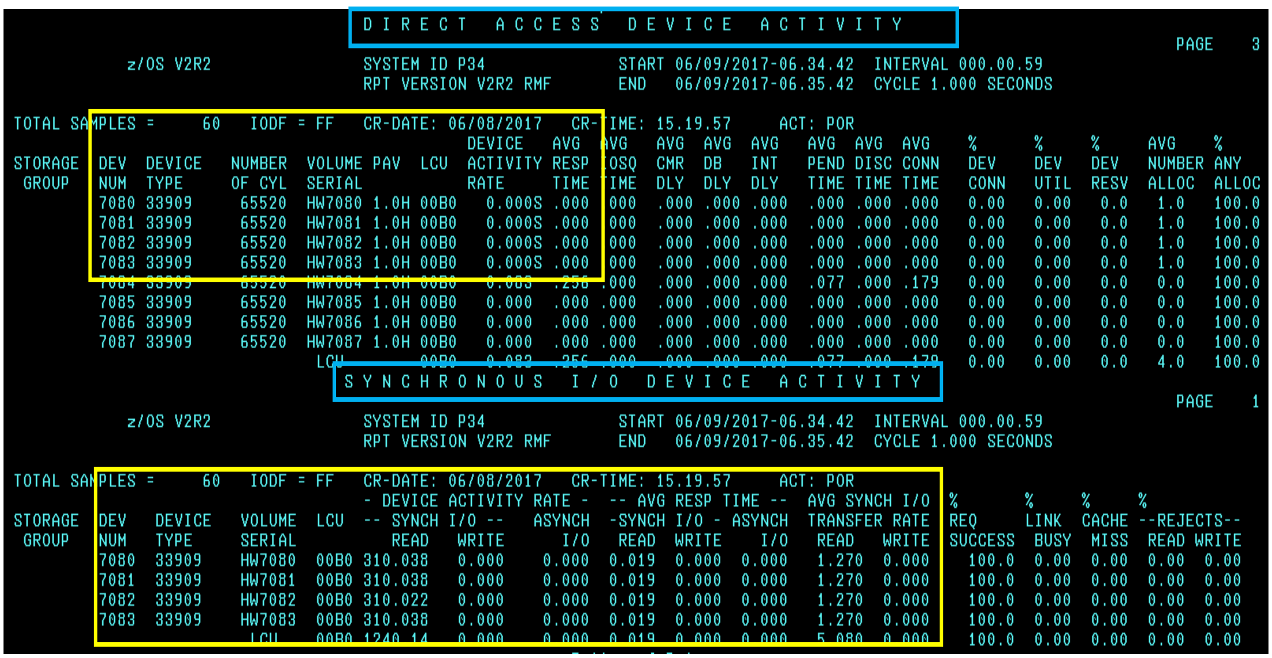

```
                                 D I R E C T   A C C E S S   D E V I C E   A C T I V I T Y
                                                                                                            PAGE    3
          z/OS V2R2              SYSTEM ID P34              START 06/09/2017-06.34.42  INTERVAL 000.00.59
                                 RPT VERSION V2R2 RMF       END   06/09/2017-06.35.42  CYCLE 1.000 SECONDS

TOTAL SAMPLES =     60   IODF = FF   CR-DATE: 06/08/2017   CR-TIME: 15.19.57     ACT: POR
                                             DEVICE   AVG  AVG   AVG  AVG  AVG   AVG  AVG  AVG    %      %      %    AVG    %
STORAGE DEV  DEVICE    NUMBER  VOLUME PAV  LCU   ACTIVITY RESP IOSQ  CMR  DB   INT   PEND DISC CONN   DEV    DEV    DEV  NUMBER ANY
 GROUP  NUM  TYPE      OF CYL  SERIAL            RATE     TIME TIME  DLY  DLY  DLY   TIME TIME TIME   CONN   UTIL   RESV ALLOC  ALLOC
        7080 33909      65520  HW7080 1.0H 00B0   0.000S .000 .000  .000 .000 .000  .000 .000 .000   0.00   0.00   0.0   1.0   100.0
        7081 33909      65520  HW7081 1.0H 00B0   0.000S .000 .000  .000 .000 .000  .000 .000 .000   0.00   0.00   0.0   1.0   100.0
        7082 33909      65520  HW7082 1.0H 00B0   0.000S .000 .000  .000 .000 .000  .000 .000 .000   0.00   0.00   0.0   1.0   100.0
        7083 33909      65520  HW7083 1.0H 00B0   0.000S .000 .000  .000 .000 .000  .000 .000 .000   0.00   0.00   0.0   1.0   100.0
        7084 33909      65520  HW7084 1.0H 00B0   0.083  .256 .000  .000 .000 .000  .077 .000 .179   0.00   0.00   0.0   0.0   100.0
        7085 33909      65520  HW7085 1.0H 00B0   0.000  .000 .000  .000 .000 .000  .000 .000 .000   0.00   0.00   0.0   0.0   100.0
        7086 33909      65520  HW7086 1.0H 00B0   0.000  .000 .000  .000 .000 .000  .000 .000 .000   0.00   0.00   0.0   0.0   100.0
        7087 33909      65520  HW7087 1.0H 00B0   0.000  .000 .000  .000 .000 .000  .000 .000 .000   0.00   0.00   0.0   0.0   100.0
                                    LCU    00B0   0.083  .256 .000  .000 .000 .000  .077 .000 .179   0.00   0.00   0.0   4.0   100.0
                              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 P34              START 06/09/2017-06.34.42  INTERVAL 000.00.59
                                 RPT VERSION V2R2 RMF       END   06/09/2017-06.35.42  CYCLE 1.000 SECONDS

TOTAL SAMPLES =     60   IODF = FF   CR-DATE: 06/08/2017   CR-TIME: 15.19.57     ACT: POR
                                 - 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
        7080  33909    HW7080 00B0 310.038   0.000   0.000 0.019 0.000  0.000    1.270  0.000    100.0  0.00  0.00  0.00  0.00
        7081  33909    HW7081 00B0 310.038   0.000   0.000 0.019 0.000  0.000    1.270  0.000    100.0  0.00  0.00  0.00  0.00
        7082  33909    HW7082 00B0 310.022   0.000   0.000 0.019 0.000  0.000    1.270  0.000    100.0  0.00  0.00  0.00  0.00
        7083  33909    HW7083 00B0 310.038   0.000   0.000 0.019 0.000  0.000    1.270  0.000    100.0  0.00  0.00  0.00  0.00
                          LCU  00B0 1240.14   0.000   0.000 0.019 0.000  0.000    5.080  0.000    100.0  0.00  0.00  0.00  0.00
```

Figure 30 – RMF Synchronous I/O

FEx16SA Card and I/O Subsystem (IOSS) performance with zHPF

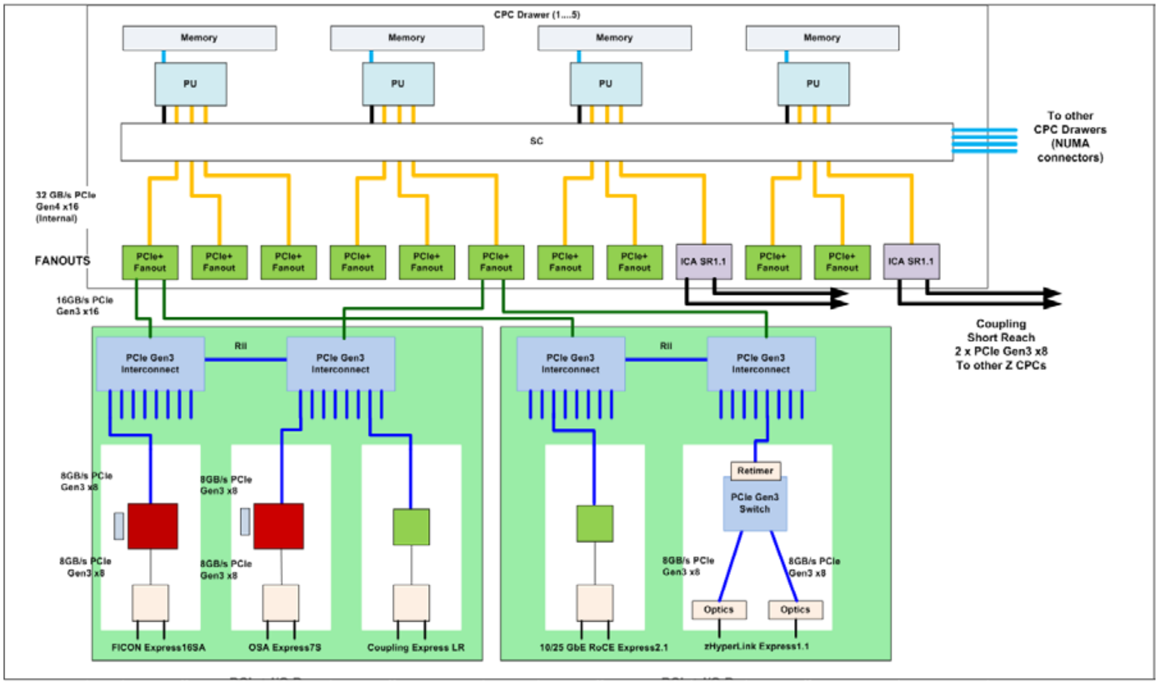

Figure 31 – z15 IOSS Diagram

The z15 supports the I/O drawer and form factor I/O cards which use Peripheral Component Interconnect Express Generation 3 (PCIe Gen3) links with increased capacity, granularity, and infrastructure bandwidth, as well as increased reliability, availability, and serviceability. The results of performance measurements done at each of the levels in this new I/O Infrastructure displayed in **Figure 31** are summarized in this section.

At the card level, the new FEx16SA channel card has two channels on the channel card.

Measurements were done with both FEx16SA channels active **Figure 32** on the same card to determine the maximum MB/sec capability of the FEx16SAcard with zHPF. When each I/O is transferring 128K bytes of data, the maximum MB/sec for two channels on a single card is approximately 3,200 MB/sec for READs or update WRITEs and over 6,400 MB/sec for a mix or READ and WRITE I/O operations. These card measurement results are about 2 times the maximum MB/sec of a single FEx16SA channel using the zHPF protocol.

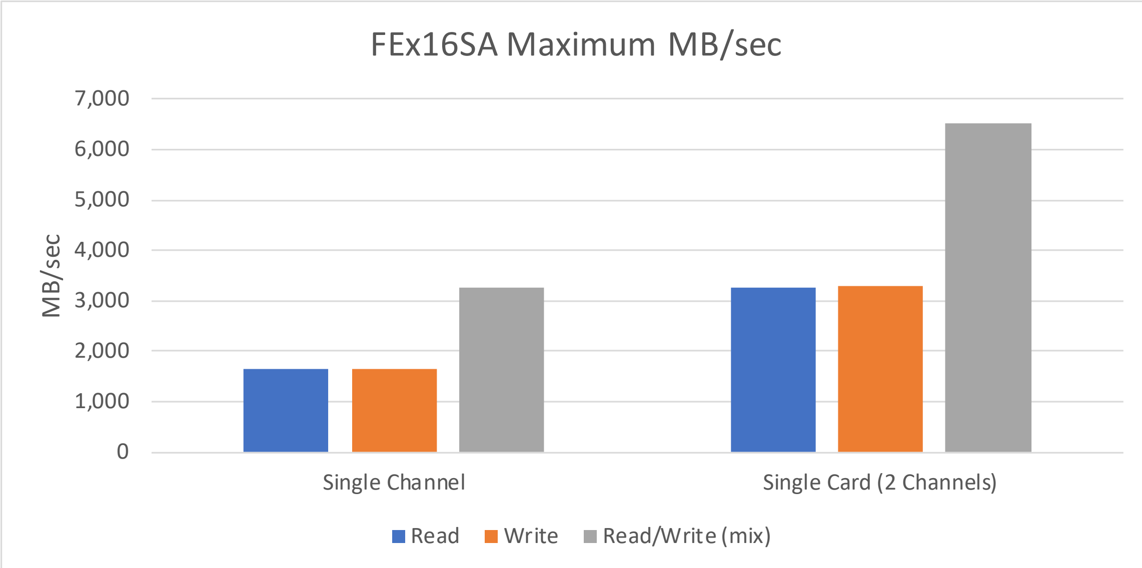

Figure 32 – FEx16SA Channel vs Card Bandwidth

Figure 33 summarizes the expected I/O bandwidth available at each level of the new z15 I/O Infrastructure including a single FEx16SA channel, a single FEx16SA card with two FEx16SA channels, single PCIe Interconnect domain with up to eight FEx16SA channel cards, and PCIe fanout card with two PCIe interconnects that allow connections to as many as 16 FEx16SA channel cards.

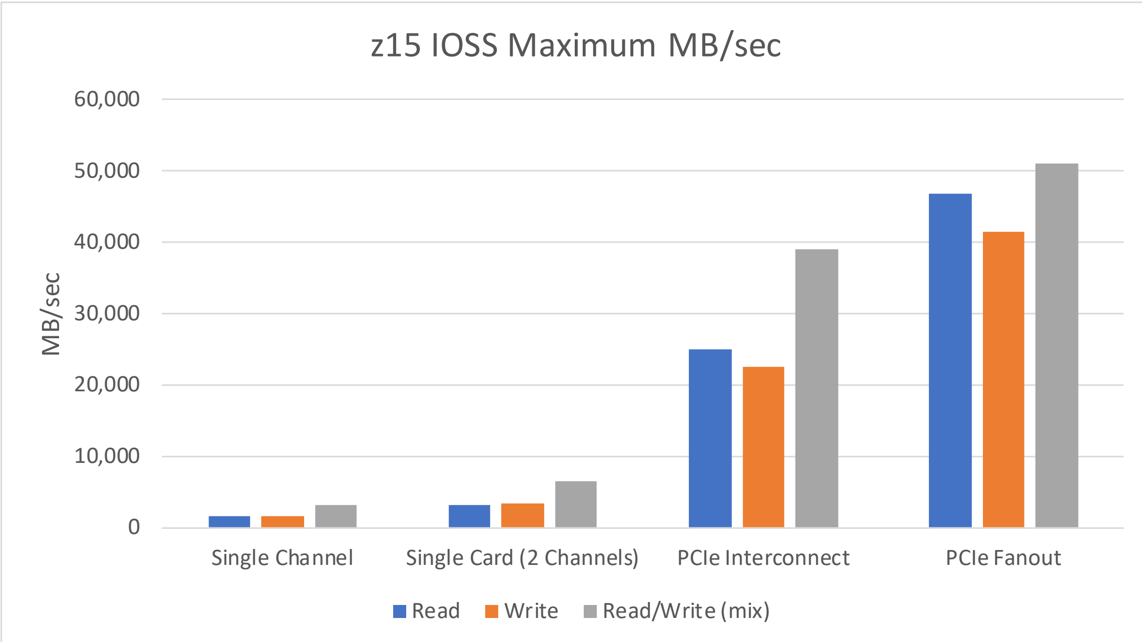

Figure 33 – z15 IOSS Bandwidth

The results achieved at the channel and card level of the I/O Infrastructure were each described in this document. For the PCIe interconnect level with eight FEx16SA channel cards, approximately 24 GB/sec was measured for READs or update WRITEs and over 38 GB/sec was measured for a mix or READ and WRITE I/O operations. For the PCIe fanout card with two PCIe Interconnects or sixteen FEx16SA channel cards, approximately 40 GB/sec was measured for READs or update WRITEs and over 50 GB/sec for a mix or READ and WRITE I/O operations.

The z15 supports the Redundant I/O Interconnect feature which is designed to help avoid unplanned outages by maintaining critical connections to I/O devices during path failures, or upgrades or repairs of a multi-drawer server. The z15 uses an alternate PCIe link which operates at the same speed as the primary PCIe link. The results achieved using two alternate path PCIe links into a single PCIe fanout card were similar to what was achieved using the primary PCIe links, namely 24 GB/sec READ or update WRITE.

This information about the maximum capabilities of various levels of the z15 IOSS is provided to help customers plan appropriately for high bandwidth demand workloads.

z15 SAP capacity

As opposed to 3 CP chips per cluster in the previous z14 architecture, the new z15 has a redesigned system architecture with 2 CP chips per cluster.. The bus architecture in z15 provides direct connectivity between any given drawer to any other drawer leading to a fully interconnected 5 drawer system as compared to a 4 drawer system in the previous z14. **Figure 34** shows this new 5 drawer system and the redesigned system architecture.

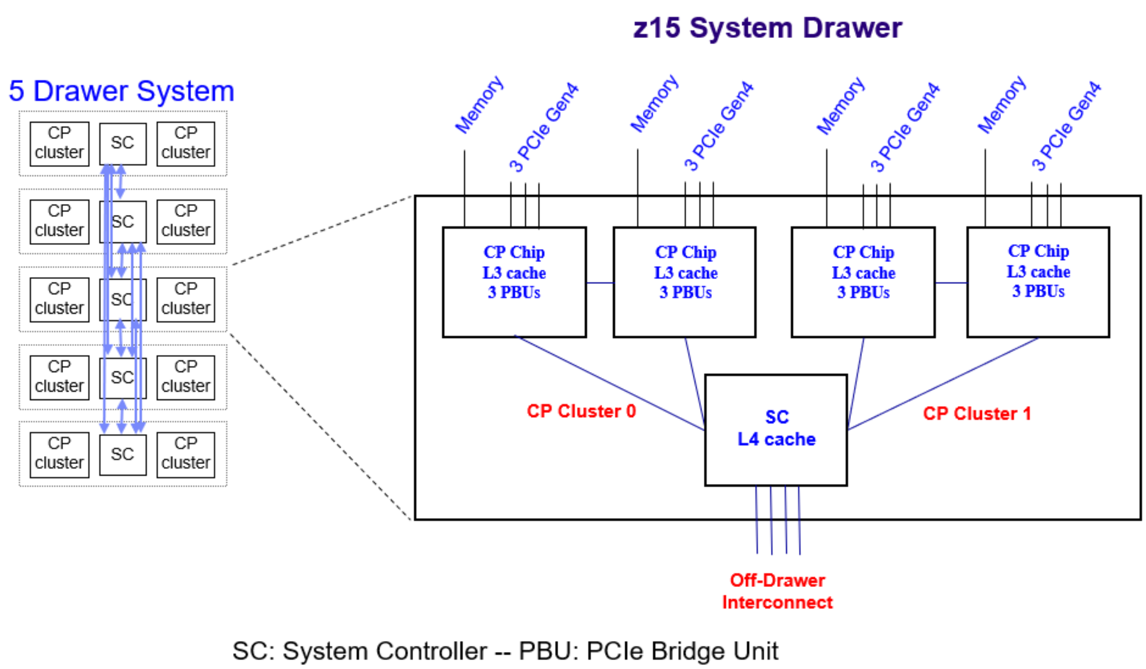

Figure 34 - Fully interconnected 5 drawer system

The System Assist Processors (SAPs) of z15 support the Simultaneous Multi-Threading (SMT) feature that was first introduced on z14. **Figure 35** shows a table of various configurations of z15 available to customers.

Model	Feature	CPs	IFLs	Unassigned IFLs	zIIP	ICFs	IFPs	Std SAPs	Add'l SAPs	IBM Spares
T01	Max34	0-34	0-34	0-33	0-22	0-34	1	4	0-8	2
T01	Max71	0-71	0-71	0-70	0-46	0-71	1	8	0-8	2
T01	Max108	0-108	0-108	0-107	0-70	0-108	1	12	0-8	2
T01	Max145	0-145	0-145	0-144	0-96	0-145	1	16	0-8	2
T01	Max190	0-190	0-190	0-189	0-126	0-190	1	22	0-8	2

Figure 35 – Available feature models (T01: Max34, Max71, Max108, Max145, Max190) of z15

As opposed to the separate High-End (HE) and Mid-Range (MR) models in z14, the new cloud and blockchain ready z15 is a one system offering called the T01 model. The various featured models of T01 (*Max34, Max71, Max108, Max145, Max190)* based on number of CPs are listed in column 'Feature' in **Figure 35**. The number of drawers in the Max34, Max71, Max108, Max145 and Max190 models are 1,2,3,4 and 5, respectively. Single processor capacity of z15 for equal n-way at common client configurations is approximately 14% greater than on z14, where as the largest IBM z15 is expected to provide approximately 25% more capacity than the largest z14 with some variation based on workload and configuration.

As before, any system configuration still has 1 XSAP core (which is not SMT enabled) and multiple other SMT enabled SAP cores. The Std SAPs listed in **Figure 35** include 1 XSAP for each of the featured models. For example, single drawer featured model (T01:Max34) will have 1 non-SMT XSAP core and 3 SMT enabled SAP cores resulting in a total of 4 standard SAPs. Additional SAPs can be purchased by customers for each drawer. Each SMT enabled SAP core will show as 2 IO Processor (IOP) threads in the RMF I/O Queueing Activity Report. The XSAP (non-SMT) will show as a single IOP thread and not as 2 threads. Thus, a model with n SAP cores will show $2n - 1$ threads in the RMF report. **Figure 36** shows a snapshot for the 4 SAP featured model T01:Max34 with 7 IOP threads.

I/O QUEUING ACTIVITY

PAGE

z/OS V2R3 SYSTEM ID P34 START 02/25/2019-10.44.10 INTERVAL 000.00.59

RPT VERSION V2R3 RMF END 02/25/2019-10.45.10 CYCLE 1.000 SECONDS

TOTAL SAMPLES = 59 IODF = F6 CR-DATE: 02/25/2019 CR-TIME: 02.54.12 ACT: POR

INPUT/OUTPUT PROCESSORS

	- INITIATIVE QUEUE -		------- IOP UTILIZATION -------			-- % I/O REQUESTS RETRIED --					-------- RETRIES / SSCH ---------				
IOP	ACTIVITY RATE	AVG Q LNGTH	% IOP BUSY	I/O START RATE	INTERRUPT RATE	ALL	CP BUSY	DP BUSY	CU BUSY	DV BUSY	ALL	CP BUSY	DP BUSY	CU BUSY	DV BUSY
00	14871.40	0.02	8.20	14871.40	22007.06	0.0	0.0	0.0	0.0	0.0	0.00	0.00	0.00	0.00	0.00
01	14871.34	0.02	8.05	14871.34	21759.60	0.0	0.0	0.0	0.0	0.0	0.00	0.00	0.00	0.00	0.00
02	14870.93	0.02	3.25	14870.93	896.536	0.0	0.0	0.0	0.0	0.0	0.00	0.00	0.00	0.00	0.00
03	14870.41	0.02	10.27	14870.41	28422.84	0.0	0.0	0.0	0.0	0.0	0.00	0.00	0.00	0.00	0.00
04	14870.83	0.02	3.38	14870.83	1285.487	0.0	0.0	0.0	0.0	0.0	0.00	0.00	0.00	0.00	0.00
05	14875.18	0.02	9.86	14875.18	28400.76	0.0	0.0	0.0	0.0	0.0	0.00	0.00	0.00	0.00	0.00
06	14871.20	0.02	3.34	14871.20	1445.266	0.0	0.0	0.0	0.0	0.0	0.00	0.00	0.00	0.00	0.00
SYS	104101.3	0.02	6.62	104101.3	104217.5	0.0	0.0	0.0	0.0	0.0	0.00	0.00	0.00	0.00	0.00

Figure 36 – RMF I/O Queueing Activity Report

With the PCIe+ Gen3 fanout cards and other enhancements mentioned in previous sub-section (including the SMT feature of SAPs), the maximum I/O processing capability on a z15™ machine with 22 standard System Assist Processors (SAPs) is now over 5.3M IO/sec, a significant 23% increase over the previous z14 with only 4.3M IO/sec.

These increases in IO capacity and the total number of CPs in a drawer make the new IBM z15™ an ideal offering for integration across the modern high-capacity hybrid cloud environments. This increase in IO capacity also enables agile deployment of IBM Z and greater flexibility in the cloud data center.

An activity rate of over 5.3M IO/sec was measured on the highest capacity z15 featured model: T01:Max190 running 10 z/OS V2.3 LPARs, having 22 System Assist Processors (SAPs) and a mix of FICON Express channels running at 16Gb speeds each executing zHPF small block (4K bytes/IO). This **maximum SAP capacity of 5.3 Million IO/sec** with 22 SAPs was achieved at 100% SAP utilization. However, for production environments, IBM recommends that SAP utilization levels be capped at 70% utilization. **Figure 37** below shows I/O capacity for the z15 and z14 models projected at 70% SAP utilization levels.

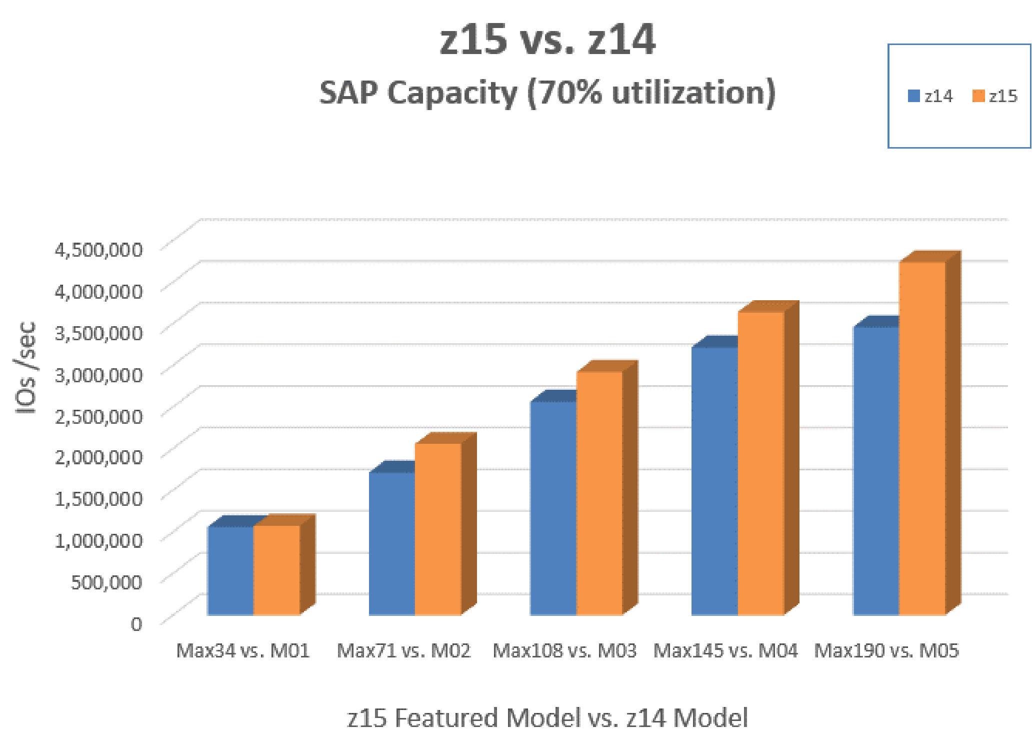

Figure 37 – SAP Capacity of z15 vs. z14

As shown in **Figure 37** above, at 70% SAP utilization levels:

- A z15 featured model T01:Max34 with 4 SAPs on 1 drawer achieved an IO capacity of 1.07M IO/sec with 2 LPARs. This is 1% higher than that of z14 M01 model with 33 CPs and 5 SAPs on 1 drawer.
- A z15 featured model T01:Max71 with 8 SAPs on 2 drawers achieved an IO capacity of 2.06M IO/sec with 4 LPARs. This is 20% higher than that of z14 M02 model with 69 CPs and 10 SAPs on 2 drawers.
- A z15 featured model T01:Max108 with 12 SAPs on 3 drawers achieved an IO capacity of 2.92M IO/sec with 6 LPARs. This is 14% higher than that of z14 M03 model with 105 CPs and 15 SAPs on 3 drawers.
- A z15 featured model T01:Max145 with 16 SAPs on 4 drawers achieved an IO capacity of 3.64M IO/sec with 8 LPARs. This is 13% higher than that of z14 M04 model with 141 CPs and 20 SAPs on 4 drawers.
- A z15 featured model T01:Max190 with 22 SAPs on 5 drawers achieved an IO capacity of 4.24M IO/sec with 10 LPARs. This is 23% higher than that of z14 M05 model with 170 CPs and 23 SAPs on 4 drawers.

Since IO intensity of customer workloads deployed on IBM Z has been growing rapidly in the past few years, it becomes critical to measure achievable IO intensity for an application running in a single LPAR instance. Thus, we measured the maximum IO capacity of a single LPAR instance separately. A single zOS LPAR executing zHPF small block (4K bytes) I/O operations on the T01:Max34 featured model with 4 SAPs achieved about 950K IO/sec.

IBM Corporation
New Orchard Rd.
Armonk, NY 10504
U.S.A

Produced in the United States of America
1/2020

....