* FICON Express channels were a new generation of FICON channels that offer improved performance capability over previous generations of ESCON channels.
  - IBM has made significant improvements to FICON channels since this product was initially shipped in 1999. NOTE: FICON® protocol first introduced in May of 1998; High Performance FICON for System z (zHPF) was introduced in October of 2008. Both FICON and zHPF are designed to support attachment to ECKD™ devices.

The High Performance FICON for IBM System z journey:

* First introduced in **October of 2008**: Initial zHPF announcement on IBM System z10® with FICON Express4 and FICON Express2 channels
  - Maximum of 31k zHPF IO/sec, 2.2 times the FICON protocol
  - Single track limit for zHPF data transfers
* In **July 2009** FICON Express8 channels were introduced on System z10:
  - Maximum of 52k zHPF IO/sec
  - 64k byte limit for zHPF data transfers
* In **July 2010** additional support on IBM zEnterprise 196 (z196) was introduced
  - Extension to multi-tracks of zHPF data transfers
* In **July 2011**, the new FICON Express8S channel:
  - Introduced a hardware data router for more efficient zHPF data transfers
  - FICON Express8S channel became the first channel with hardware specifically designed to support zHPF.

Note: With zHPF it is possible to achieve a significant improvement in both small block IO/sec processing for OLTP work loads and large sequential READ+WRITE I/O processing compared to previous FICON offerings.

zHPF performance: 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 the zHPF protocol improves upon the FICON protocol, 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 a 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.
    NOTE: The CCW is the original I/O operation used for communications with the channel subsystem.
  - 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.
* 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.
  - The channel is no longer expected to process and keep track of each individual channel command word.
  - 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.
  - The figure 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 in half, using one IU from the channel to the control unit and two IUs from the control unit to the channel.

‡ Information Units (IUs) consist of one to four Fibre Channel frames.

Link Protocol Comparison for a 4K READ

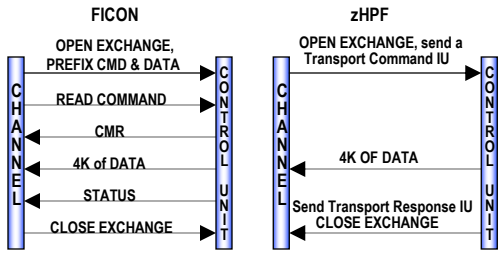

zHPFprovides a much simpler link protocol than FICON

zHPF compliments the Extended Address Volumes for System z (EAV) strategy for growth by increasing the I/O rate capability as the volume sizes increasing The I/O rate capability as The volume sizes expand vertically.

* 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.

NOTE: 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 FICON Express8S channels is designed to be significantly higher with the zHPF protocol compared to the FICON protocol.
  - 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 o channel has been increased with zHPF.

One primary performance improvement in a FICON Express8S channel with zHPF compared to FICON is the maximum number of small block I/Os per second (4k bytes per I/O) that can be processed.

zHPF exploitation - zHPF can be turned on or off

Note: For z/OS exploitation, there is a parameter in the IECIOSxx member of SYS1.PARMLIB (ZHPF=YES/NO) and on the SETIOS command to control whether zHPF is enabled or disabled. The default is ZHPF=NO.

i.e. *Console Command*
D M=DEV(B020)
:
FUNCTIONS ENABLED = MIDAW, zHPF

For zHPF exploitation, FICON Express8S (CHPID type FC) on a zEC12 requires at a minimum: z/OS V1.13, z/OS V1.12, z/OS V1.10 and V1.11 with PTFs, Linux on System z distributions: For single-track support: SLES 11 SP1, Red Hat RHEL 6. Multitrack support requires at a minimum SLES 11 SP2, Red Hat RHEL 6.1.

Reflected in the bar chart are the "best can do" capabilities of each of the last few generations of FICON channels that have been supported.

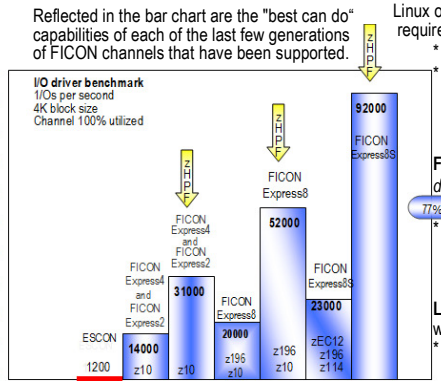

For an additional perspective, an ESCON® performance benchmark measured in the late 1990's is displayed in the I/Os per second bar chart. Each of the recent generations of FICON Express channels, has two bars, one which displays the maximum capability using the FICON protocol exclusively and another that shows the maximum capability using the zHPF protocol exclusively.

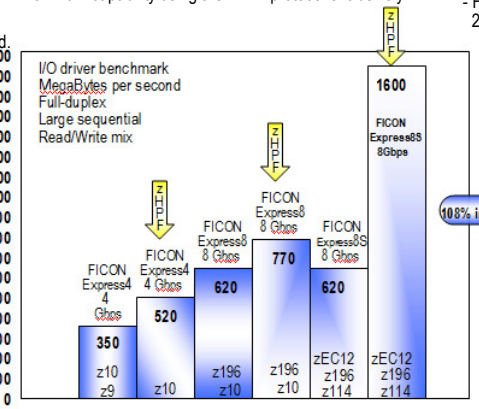

The chart above displays the maximum READ+WRITE MB/sec for each channel. In laboratory measurements, using FICON Express8S on a zEC12, z196, or z114 with the zHPF protocol and a mix of large sequential read and write data transfer I/O operations, FICON Express8S operating at 8 Gbps achieved a maximum throughput of 1600 READ+WRITE MB/sec compared to the maximum of 770 READ+WRITE MB/sec achieved with FICON Express8 operating at 8 Gbps. This represents approximately a 108% increase.

Express8S channel measurement results summarized here were collected using z/OS V1.13. For Linux® on System z measurements, Linux on System z distribution SUSE Linux Enterprise Server (SLES) 11 SP1 was used.

* zHPF-capable channels and devices support both FICON and zHPF protocols simultaneously.
* The Media Manager component of DFSMS™ detects whether the device supports zHPF or not and builds the appropriate channel programs.
  - Media Manager will build the zHPF Transport Mode channel programs for DB2® databases and access methods and file systems such as VSAM, HFS, z/FS, PDSE, and SAM extended format datasets.

FibreExpress8S channel performance using zHPF and FICON - With the introduction of the *hardware data router* on the FICON Express8S features, performance improvements occur using the zHPF protocol to process both small random access and large sequential I/O operations.

77% increase

* The hardware data router transfers data both to and from the channel to zEnterprise memory more efficiently than previous generation channels for I/O operations using the zHPF protocol.
  - As a result improvements can be seen in both response times and maximum throughput for IO/sec and MB/sec for workloads using the zHPF protocol on FICON Express8S channels.

Linux on System z operating system - Several FICON Express8S channel performance measurements were also completed using the Linux on System z operating system.

* The results achieved with Linux on System z are similar to those achieved with z/OS.
  - Using the zHPF protocol transferring 4k bytes per I/O operation, a maximum of 92,000 IO/sec was achieved or four times the maximum of 23,000 achieved using the FICON protocol.
  - Using the zHPF protocol transferring 48k bytes per I/O operation, a maximum of 800 READ MB/sec, 800 WRITE MB/sec and 1500 READ+WRITE MB/sec was measured or 1.5 to 2.7 times the maximum of 520 READ MB/sec, 480 WRITE MB/sec and 550 READ+WRITE MB/sec measured using FICON protocol.
* In general, IBM System z recommends the following guidelines to achieve good response times – keep channel processor utilizations less than 50% and keep channel link utilizations less than 70%.
  - For a 2, 4 or 8Gbps link, this means maximum READ or WRITE MB/sec should be kept below 140, 280 or 560MB/sec respectively.
  - When using either the FICON or the zHPF protocol, the FEx8S channel processor utilization is the most important factor to observe for small data transfers.
  - With the hardware data router on the FEx8S channel, 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.

108% increase

DB2 - One hundred percent of DB2 I/O is now convertable to zHPF for optimal I/O resource utilization, bandwidth and response time :

* Format write throughput increases up to 52% (4K pages)
* Applies to Load, Reorg writes to shadow / restores
* Preformatting throughput increases up to 100%
* Insert preformatting is asynchronous, except for when allocating a new extent
* Synch I/O cache hit response time decreases by up to 30%
* Sequential prefetch throughput increases up to 19%
* Dynamic prefetch throughput increases up to 23% (40% with SSD)
* Disorganized index scans
* DB2 10 throughput increases up to 111% (more with 8K pages)
* Together DB2 10 and zHPF is up to 11 times faster
* Skip sequential index-to-data access
* Cache misses are 3 to 4 times faster.

zHPF may help reduce the infrastructure costs for System z I/O by efficiently utilizing I/O resources so that fewer CHPIDs, fiber, switch ports and control unit ports are required.

Improve FICON Scale, Efficiency and RAS

* As the data density behind a CU and device increase, scale I/O rates and bandwidth to grow with the data
* Leverages HBA hardware optimizations done while preserving System z QOS
* Significant improvements in I/O rates (4-5x) for small block transfer
* Improved I/O bandwidth (ability to bi-directionally fill the link at 8 Gbps.
* New ECKD commands for improved efficiency
* Improved first failure data capture
* Additional channel and CU diagnostics for debugging MIH conditions.

Missing Interrupt Handler