


Prologue

- The IBM z13 was built from the casters up for **CAMSS workloads** where the I/O enhancements reduce transactional latency mitigating increases to transaction response times that might be introduced by adding cloud sourced data to the work flows.
- These performance improvements also improve the scalability of z/OS and middleware (such as DB2) to meet the demands of mobile applications.
- The increase in the volume of I/O drives the requirement for further improvements to the already industry-leading quality of service (QoS) capabilities of z Systems.
- These **new QoS enhancements** affect processing in the following ways:
 - Improve the quality of the Fibre Channel links to reduce error rates.
 - Extend the z/OS workload manager into the SAN fabric to manage the end-to-end work according to client policy.
 - Provide reduced cost for the physical infrastructure with enhanced virtualization to allow sharing more of the enterprise I/O traffic over shared ISLs.
 - Provide improved availability with additional flexibility and scale to the I/O configuration.

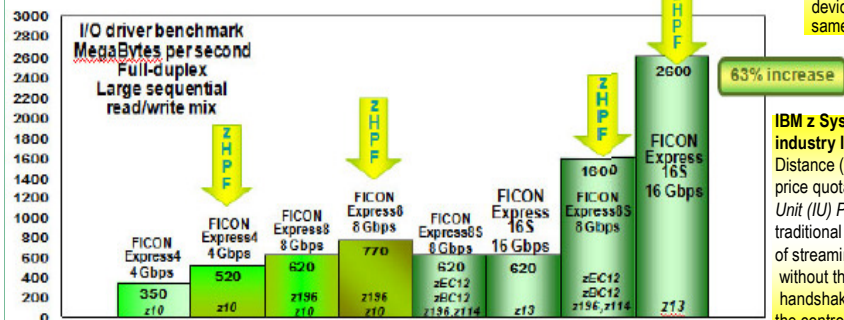
New value to I/O infrastructure

- IBM z13 is the latest generation of the IBM mainframe to deliver **substantial value to the I/O infrastructure**:
 - Fibre Connection (FICON) Dynamic Routing is a new feature that enables exploitation of SAN dynamic routing policies in the fabric to lower cost and improve performance for supporting I/O devices.
 - Mainframe SAN Fabric Priority, with exploiting storage products, extends the z/OS WorkLoad Manager (WLM) to the SAN infrastructure providing improved resilience and autonomic capabilities while enhancing the value of FICON Dynamic Routing.
 - FICON Express16S with the DS8870 can provide substantially improved DB2 transactional latency and up to 32% reduction in elapsed time for I/O bound batch jobs.
 - Forward Error Correction (FEC) with supporting storage capabilities to the Fibre Channel link protocol can operate at higher speeds, over longer distances, with reduced power and higher throughput, while retaining the same reliability and robustness that FICON has traditionally been known for.
 - zHPF Extended Distance II feature can reduce the impact of distance on I/O response times by 50% for large data writes, providing significant response time improvements for multi-site IBM Parallel Sysplex® environments.
 - Scales to six logical channel subsystems (LCS) allows for up to 85 client usable LPARs.
 - All z13 FICON features (FICON Express8, FICON Express8S, and FICON Express16S) can support up to 32K devices allowing further consolidation of more devices onto a set of channels.
 - A fourth subchannel set for each logical channel subsystem (LCSS) is provided to eliminate single points of failure for storage after a disk failure by facilitating the exploitation of IBM DS8870 multi-target Metro Mirror storage replication with IBM Geographically Dispersed Parallel Sysplex™ (IBM GDPS®) and IBM Tivoli Storage Productivity Center for Replication HyperSwap.

FICON Express16S



FICON I/O rates: The history of z Systems channel I/O rates. FICON Express16S is able to achieve 93,000 I/O operations per second.



FICON bandwidth: The history of z Systems channel bandwidth. FICON Express16S is able to achieve 2600 MBps bidirectional throughput on 16 gigabit per second (Gbps) Fibre Channel links



With z13, 16 Gb link speeds are supported by the IBM FICON Express16S features. The faster link speed provides reduced I/O latency for critical middleware I/O, such as writing the database log.

The SAN qualities of service (such as fabric priority) helps complement the sharing of the ISLs by extending the IBM z/OS® Work Load Manager (WLM) to manage I/O priority of the fabric.

With the FICON Express16S generation of features, IBM added Forward Error Correction (FEC) capabilities to the Fibre Channel link protocol, using the most advanced FEC coding in the industry.

FEC allows FICON channels to operate at higher speeds, over longer distances, with reduced power and higher throughput, while retaining the same reliability and robustness that FICON has traditionally been known for.

IBM System z I/O Exerciser (ESAIO) simplifies the chore of exercising the I/O connections in the I/O configuration before starting z/OS and running production work. This tool is intended to help identify possible cabling or definition errors by validating that all the paths defined to each device actually connect to the same physical device.

IBM z Systems have historically been the industry leader in I/O execution at long Distance (100 km, 200 km with a request price quotation (RPQ)). The Information Unit (IU) Pacing capability built into the traditional FICON protocol was capable of streaming write data to the control unit without the protocol requiring interlocked handshakes between the channel and the control unit.


- The right figure shows the I/O latency characteristics of a single 4K I/O request on both the z13 and zEC12 machines using 8 and 16 Gbps link technologies.
- The FICON Express16S channel with the DS8870 16 Gbps host bus adapter provides the best I/O latency result, up to 21% better than the zEC12.
- The FICON Express16S channel attached to a DS8870 with a 16 Gbps host bus adapter provides the **best I/O latency results**, as in the example.
- zHPF Read z13 with FICON Express16S Response Time is 21% lower (better) than IBM zEnterprise® EC12 (zEC12) with FICON Express8S (33 microseconds (usecs))
- zHPF Write z13 w/ FICON Express16S Response Time is 21% lower (better) than zEC12 w/ FICON Express8S (37 usecs)
- FICON Read z13 w/ FICON Express16S Response Time is 11% lower (better) than zEC12 w/ FICON Express8S
- FICON Write z13 with FICON Express16S Response Time is equivalent to zEC12 with FICON Express8S.

- The **primary value of moving to faster link technologies** is in the reduction of I/O latency for "large" read/write operations.
 - For 8 Gb Fibre Channel (FC) link speeds, 5.1 microseconds is required to transport every 4 KB block. For 16 Gb FC link speeds, 2.5 microseconds is required for every 4 KB block.
 - I/O operations that transfer much more data amortizes the overhead of starting the I/O across a very large data transfer leads to a much more substantial improvement, (i.e. a 128 KB write operation can theoretically have the elapsed time reduced by up to 83.2 microseconds (32 blocks at 2.6 usec/block).


- NOTE:** The actual latency improvements depend heavily on the HBA implementation and the channel's ability to move data from the host memory to the HBA.
- The time for **DB2 to write** to the log can be quite a substantial part of the DB2 transaction.
 - A reduction in log write latency improves DB2 transactional latency.
 - As described previously, a 128 KB log write using FICON Express16S for DB2 can result in a measurable I/O service time improvement, (i.e. for a well-tuned SAP banking workload, 60% of the transactional latency is spent waiting for synchronous I/O).
 - > This synchronous I/O is devoted to the write operation to the DB2 log .since a reduction to 60% of the latency can yield significant improvements to the DB2 transactional latency.
- The next two figures on right show the I/O latency times on **z13 versus zEC12** when running 256 KB reads and writes using the various link technologies available.
 - In the measurement, four I/O streams were active, half were reads and half were writes.
- The z13 with FICON Express16S and 16 Gbps DS8870 technology achieved the best possible I/O response times, **32% better** than with the zEC12 using 8 Gbps technology.
- The figure on lower right shows relative I/O latency of **z13 versus zEC12** with possible combinations of FICON channel and DS8870 link speeds.

- NOTE:** The FICON Express16S channel running at 16 Gbps to a DS8870 with 16 Gb links provides the **best performance** for a mixture of reads and writes at 256 KB.

- Relative I/O latency is improved** with the z13, as shown in this example:
 - Below 350 MBps z13 FICON Express8S 8Gb HBA response time is 17% lower (better) than zEC12 FICON Express8S 8Gb HBA.
 - Below 350 MBps z13 FICON Express16S 8Gb HBA response time is 22% lower (better) than zEC12 FICON Express8S 8Gb HBA
 - Below 350 MBps z13 FICON Express8S 16Gb HBA response time is 30% lower (better) than zEC12 FICON Express8S 8Gb HBA
 - Below 350 MBps z13 FICON Express16S 16Gb HBA response time is 32% lower (better) than zEC12 FICON Express8S 8Gb HBA

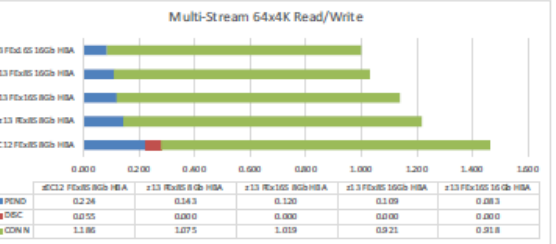


Single stream I/O latency of the z13 with FICON Express16S to a 16 Gb DS8870 versus zEC12 with FICON Express8S to 8 Gb DS8870.



Multi-Stream 64x4K Read/Write

Multi-Stream 64x4K Read/Write



Relative I/O latency of z13 versus zEC12 with possible combinations of FICON channel and DS8870 link speeds