Enhance your I/O performance with IBM System z

BY HARRY YUDENFRIEND

As you probably already know, IBM has made great strides in enhancing the I/O performance of IBM System z servers and storage. This article explores some of those enhancements, with a focus on FICON I/O operations and the IBM System Storage DS8000 solid state drive (SSD) technology.

IBM technologies to improve your I/O performance

Over the last few years, IBM System z has delivered a number of new I/O technologies to reduce I/O service times and more efficiently utilize the I/O resources in the System z I/O architecture. Figure 1 shows the individual components of the service time for a FICON I/O operation. System z I/O architecture is unique in that it allows the host to parse the individual components of I/O service time.

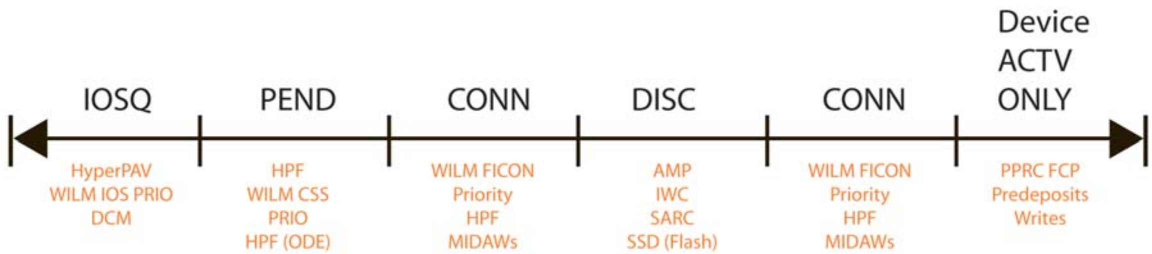

Figure 1: Recent IBM technologies to improve I/O performance

For each component of I/O service time, IBM has delivered innovations for improved I/O performance, as described in the sections that follow.

IOS queue time (IOSQ)

This is the amount of time that an I/O request remains queued in the operating system before being initiated for the device. HyperPAV technology virtually eliminates IOSQ time by assigning alias I/O devices on demand, as the application I/O requests need them. HyperPAV also works with z/OS Workload Manager (WLM) to ensure that available resources are managed to meet the goals specified by your installation.

HyperPAV also provides virtualization of System z I/O addressing to utilize more effectively the number of alias device addresses available for I/O among multiple sharing systems.

Pending time (PEND)

This is the amount of time that an I/O request remains queued in the channel subsystem, channel, or control unit before being started. WLM has the metrics and controls for effectively setting I/O priority for the channel subsystem and storage subsystem to meet the goals that your installation specifies.

High performance FICON (HPF) technology allows execution that is more efficient by allowing operations to remain queued in the storage subsystem instead of requiring re-transmission to the storage subsystem after a device busy/owed device end (ODE) sequence, which typically occurs when using reserve/release commands on the device.

Connect time (CONN)

This is the amount of elapsed time in which the I/O operation is actively executing through the channel subsystem and storage control unit. The WLM specified FICON I/O priority allows the control unit to throttle the link bandwidth to favor high priority requests at the expense of low priority requests.

Modified indirect addressing words (MIDAW) technology reduces the number of fiber channel sequences needed to execute I/O requests that, in the past, would have used data chaining protocols (for example, I/O to extended format data sets).

Furthermore, HPF technology provides a means for FICON operations to take advantage of hardware optimizations provided by the host bus adapters (HBAs) in System z, as well as any HBA in the storage subsystems.

Figure 2 shows how these high performance FICON optimizations improve channel efficiency to allow much higher I/O rates for online transaction processing (OLTP) workloads (in this case the DB2IO workload).

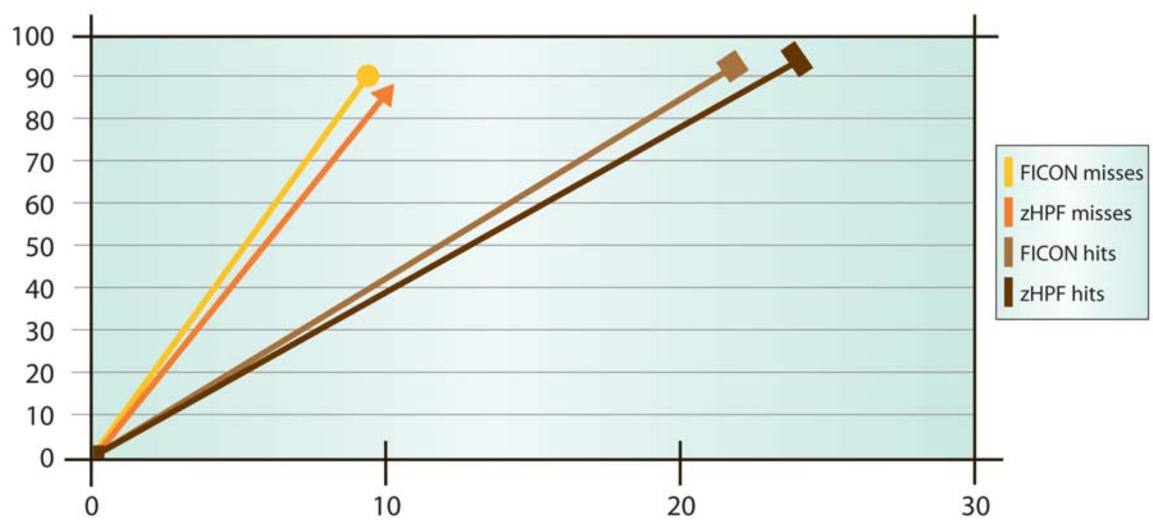

Figure 2: High Performance FICON improvements for higher I/O rates for OLTP workloads

Disconnect time (DISC)

This is the amount of time in which the I/O operation is not actively executing. For read operations, this includes the time needed to service a cache miss. For write operations, this includes the time needed for synchronous replication technologies.

System z and extended count key data (ECKD) architecture allow a single I/O operation to transfer data from multiple non-contiguous portions of the volume. The adaptive multi-stream prefetch (AMP) algorithm improves the efficiency of cache management and reduces the time required to satisfy cache misses. The algorithm for the DS8000 solid state drives (SSD) uses hints provided by DB2. The DS8000 must itself decide how much data to prefetch asynchronously from the disk. IBM introduced adaptive multi-stream prefetch in the 2.4G LIC release (July 2007) of the DS8000 to adjust the prefetch quantity, based on the needs of the application, thus avoiding over-utilizing the disks if the application did not require the resources.

DS8000 also has implemented an Intelligent Write Cache (IWC), which will order the write operations to the backend disks in order to get execution that is more efficient at the backend drives by minimizing head movement.

Device active only time (DAO)

This is the time it takes the synchronous replication technology to propagate the data to the secondary storage subsystem. The IBM DS8000 Metro Mirror feature exploits the ‘pre-deposit write’ option of the Fibre Channel Protocol (FCP) to eliminate a round trip signal needed by standard FCP operations.

The remainder of this article describes the benefits of SSD technology in reducing the time needed for resolving cache misses.

SSD technology

SSD technology can complement all of these technologies to provide optimal I/O performance for System z. With DFSMS policy-based storage management, System z, and z/OS instrumentation and tooling, your installation can manage the transition from traditional spinning storage devices to SSD. See “Stop spinning your storage wheels: z/OS Support for solid state drives in the DS8000 storage subsystem” in *z/OS Hot Topics Newsletter Issue 20*, March 2009, GA22-7501-16.

Figure 3 shows how this synergy between high performance FICON and SSD works for read response time (4K in microseconds).

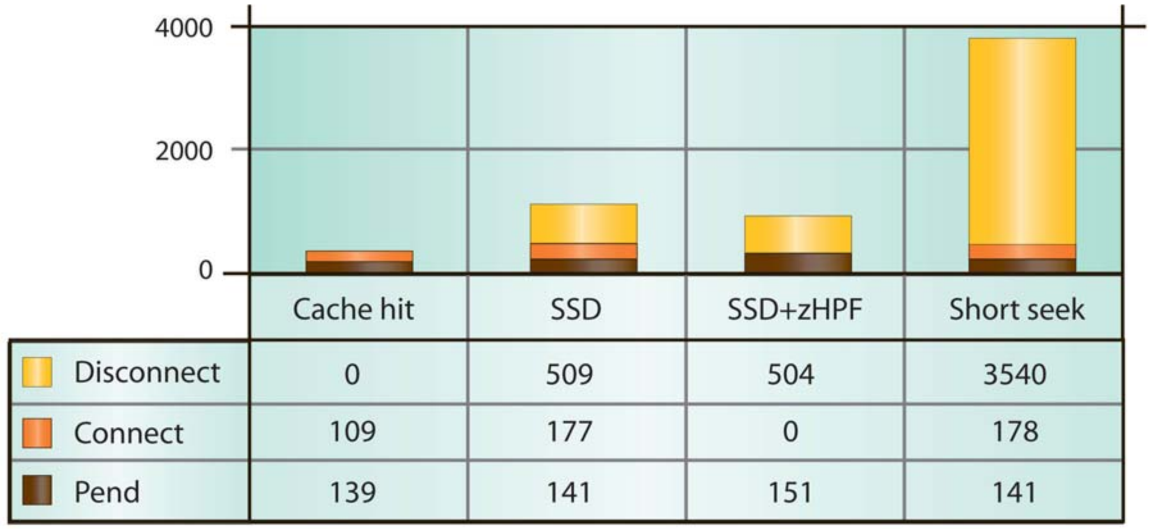

	Cache hit	SSD	SSD+zHPF	Short seek
Disconnect	0	509	504	3540
Connect	109	177	0	178
Pend	139	141	151	141

Figure 3: High Performance FICON and SSD technology to reduce I/O service times

Conclusion

System z continues to innovate in the I/O stack to enhance I/O performance for your installation. These innovations target various components of the I/O service time to provide differentiating value for the System z platform. The most recent of these I/O performance innovations complement the new extended address volume (EAV) feature of the DS8000 and z/OS V1R10 to provide enhanced scalability, while maintaining all of the z/OS classic qualities of service.