IBM TotalStorage Enterprise Storage Server Model 800 New Features in LIC Level 2.3.0 Performance White Paper

11/21/2003

Lee LaFrese
ESS Performance Evaluation

Storage Systems Group
International Business Machines Corporation

Notices and Disclaimer

Product data has been reviewed for accuracy as of the date of initial publication. Product data is subject to change without notice. This information may include technical inaccuracies or typographical errors. IBM may make improvements and/or changes in the product(s) and/or programs(s) at any time without notice. References in this document to IBM products, programs, or services does not imply that IBM intends to make such products, programs or services available in all countries in which IBM operates or does business.

THE INFORMATION PROVIDED IN THIS DOCUMENT IS DISTRIBUTED "AS IS" WITHOUT ANY WARRANTY, EITHER EXPRESS OR IMPLIED. IBM EXPRESSLY DISCLAIMS ANY WARRANTIES OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE OR NON-INFRINGEMENT. IBM shall have no responsibility to update this information. IBM products are warranted according to the terms and conditions of the agreements (*e.g.*, IBM Customer Agreement, Statement of Limited Warranty, International Program License Agreement, etc.) under which they are provided. IBM is not responsible for the performance or interoperability of any non-IBM products discussed herein.

The performance data contained herein was obtained in a controlled, isolated environment. Actual results that may be obtained in other operating environments may vary significantly. While IBM has reviewed each item for accuracy in a specific situation, there is no guarantee that the same or similar results will be obtained elsewhere.

The provision of the information contained herein is not intended to, and does not, grant any right or license under any IBM patents or copyrights. Inquiries regarding patent or copyright licenses should be made, in writing, to:

IBM Director of Licensing
IBM Corporation
North Castle Drive
Armonk, NY 10504-1785
U.S.A.

Introduction

IBM continues to deliver performance enhancements to their high end storage systems with the introduction of three new features for the TotalStorage Enterprise Storage Server (ESS) model 800. The following features are available beginning with Licensed Internal Code (LIC) Level 2.3.0:

- A faster turbo option (referred to as Turbo II)
- Peer-to-Peer Remote Copy over fibre channel protocol (referred to as PPRC over FCP)
- Arrays across loops (referred to as AAL)

This paper discusses and illustrates the performance of these features in both Open and z/OS environments. Some general configuration and usage guidelines are provided. Answers to several frequently asked questions are also included.

Workload Characteristics and Measurement Configurations

The information in this paper includes data from actual measurements completed in the IBM Tucson, IBM Mainz and IBM Gaithersburg performance laboratories. Based on experience with numerous customers, IBM has identified a number of workload profiles that we believe to be representative of typical application I/O workloads. High level descriptions of the benchmark workload characteristics and a summary of the various hardware configurations measured are listed below.

Open Workloads

- *70/30/50* -- 70% reads, 30% writes, 50% cache hits, this workload is considered comparable to typical online transaction processing in an Open systems environment.
  Read/Write Ratio = 2.33, Read Hit Ratio = 0.50, Destage Rate = 17.2%, Transfer = 4 KB
- *Sequential* -- these workloads are similar to typical batch processing
  100% Read or 100% Write, 64 KB transfers to/from disk, sequential access pattern

z/OSWorkloads

- *Cache Standard* -- this workload is considered comparable to typical online transaction processing in a z/OS environment.
  Read/Write Ratio = 3, Read Hit Ratio = 0.735, Destage Rate = 11.6%, Transfer = 4 KB
- *Cache Friendly* -- this workload displays high locality of reference and thus a high cache hit ratio.
  Read/Write Ratio = 4.9, Read Hit Ratio = 0.82, Destage Rate = 7.5%, Transfer = 4 KB
- *Cache Hostile* -- this workload is highly random and thus has a relatively low cache hit ratio.
  Read/Write Ratio = 2, Read Hit Ratio = 0.34, Destage Rate = 18.3%, Transfer = 4 KB
- *QSAM Sequential* -- these workloads are similar to typical batch processing.
  100% Read or 100% Write, 5 x 27 KB transfers to/from disk, sequential access pattern

Configurations

- ESS model 800 both with and without the Turbo or Turbo II options and 16 or 32 GB cache
- 256 or 128 logical volumes with 3 Parallel Access Volume (PAV) aliases per base address for z/OS
- 18.2, 36.4 and 72.8 GB disks with various quantities of HDDs, 15 K RPM, RAID 5
- 2 or 8 Fibre Channel 2 Gb Host Channels for Open measurements as indicated. 4, 8 or 16 FICON Express 2 Gb Host Channels for z/OS measurements as indicated.
- Direct attached Fibre Channel PPRC links for PPRC runs. ESCON links use 64-bit host adapters.
- pSeries 7040 or p690 AIX Server for Open measurements. z900 2064 Processor for z/OS measurements.

Performance of the Turbo II Option

The brains of the ESS is the Symmetric Multiprocessor (SMP) complex that runs the software that controls operation of the disk storage system. Each I/O processed by the ESS uses cycles within the SMP complex. With its dual cluster design, the ESS contains two dual-active SMPs. In the event of a cluster failure, one SMP is capable of taking over and running the entire system.

Before LIC level 2.3.0 became available, the ESS model 800 was available either as a non-turbo or with an optional Turbo feature. The non-turbo contains a pair of 600 MHz 4-way SMPs. The optional Turbo feature upgrades the SMP to a 668 MHz 6-way. With the new Turbo II option, the SMP is now a 750 MHz 6-way. This provides about 20% more raw processing power than the original Turbo feature. Compared to the non-turbo, the Turbo II has about twice the raw processing power. Although raw processing power does not directly translate into performance, the Turbo II feature does enable the ESS to handle more I/Os with better performance when compared to the non-turbo or the original Turbo feature. Of course, many real world workloads do not stress the SMPs in the ESS. Thus only certain types of workloads running at a high rate of activity will truly benefit from the Turbo II feature.

Figures 1 and 2 compare the performance of the ESS 800 non-turbo versus an ESS 800 with both the Turbo and Turbo II features. These charts show the maximum throughput for a number of standard Open and z/OS benchmark workloads. Workloads that are storage processor intensive such as read cache hits and on-line transaction processing (OLTP) oriented benchmarks such as Cache Standard show improvement with the Turbo II option. Other workloads such as sequential reads and writes are not processor constrained. There is no noticeable improvement in sequential throughput when using the Turbo II feature compared to the non-turbo ESS. Turbo II is recommended for workloads with high I/O rate demands. As a rule of thumb, Turbo II should be considered for peak on-line I/O rates over 20,000 IO/sec on a single ESS.

Channels are also a consideration with the Turbo II feature. You will not be able to fully exploit the capabilities of the Turbo II unless you have enough front end channel bandwidth. An ESS with the Turbo II option should be configured with eight or more 2 Gb FICON or fibre channel connections. There is little point to using the Turbo II option with a SCSI or ESCON attached ESS.

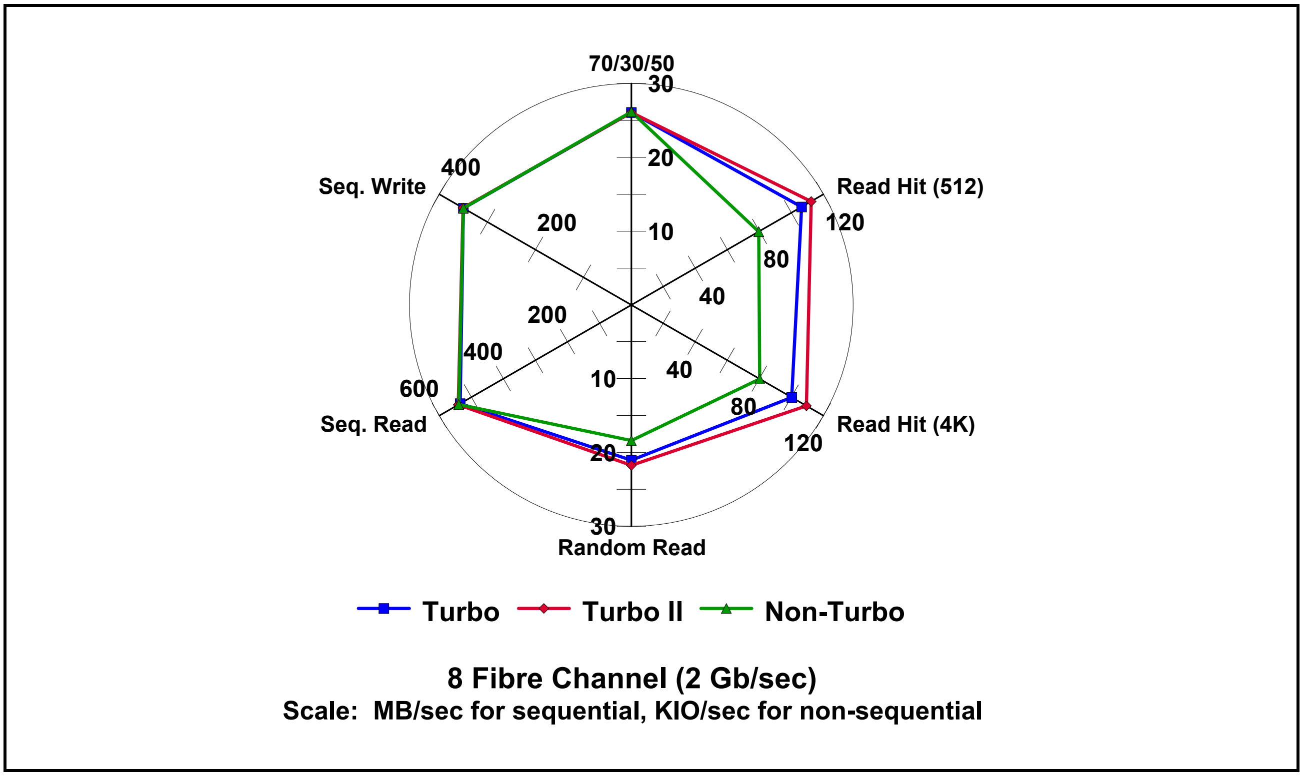

Figure 1 - ESS 800 Performance for Open Systems, NonTurbo vs. Turbo vs. Turbo II

-

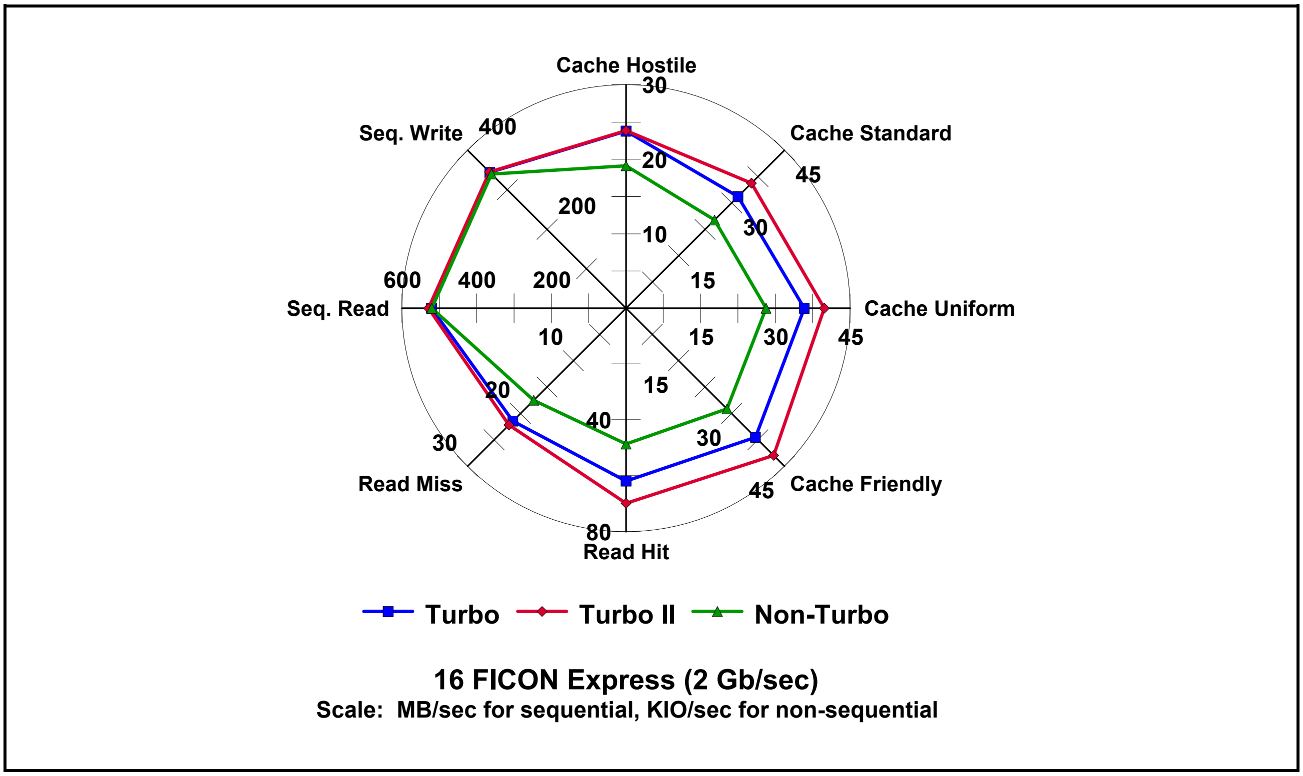

Figure 2 - ESS 800 Performance for z/OS, Non-Turbo vs. Turbo vs. Turbo II

The IBM Disk Magic tool can be used to compare modeled performance of an ESS non-turbo versus the Turbo II feature for a particular workload. It is recommended that modeling be done to explore whether or not the Turbo II feature is really necessary to meet your actual workload performance requirements.

The Turbo II feature may also improve performance when Copy Services are extensively used with ESS. PPRC, XRC and Flash Copy all place demands on the SMPs within the ESS. The Turbo II feature provides more SMP "horsepower" to handle these demands. It is not always possible to precisely predict the effect that a combination of copy services and customer workloads may place on the SMPs in an ESS. In these cases, deciding to use a Turbo II feature may simply be a matter of judgment. A combination of peak I/O rates over 15,000 IO/sec and regular use of one or more copy services would indicate that the Turbo II feature should be seriously considered.

Figure 3 compares the performance of an ESS 800 non-turbo to the Turbo and Turbo II features with PPRC active. The workload is z/OS Cache Standard and the PPRC distance is 0 km.

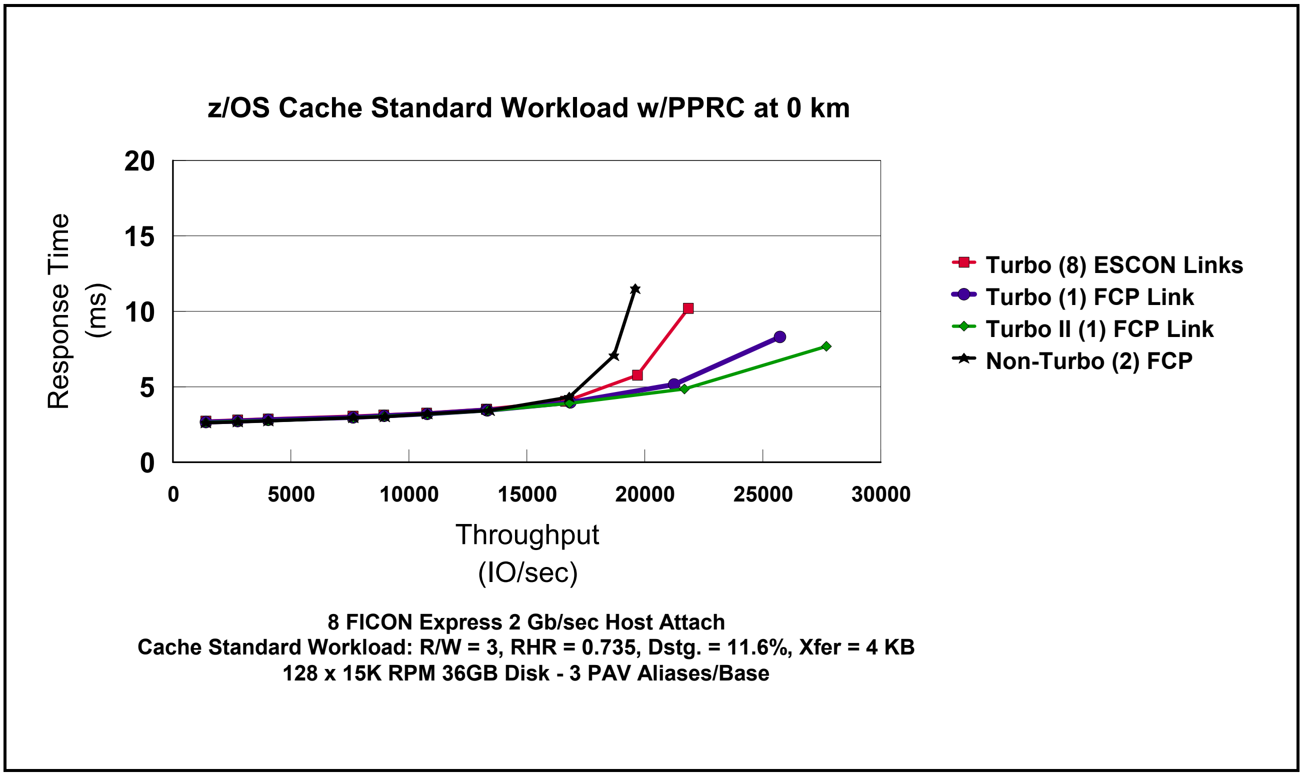

Figure 3 - ESS 800 Performance for z/OS with PPRC at 0 km, Non-Turbo vs. Turbo vs. Turbo II

Another area where the Turbo II feature may help is performance both during and immediately after a cluster failover. Cluster failover time may be up to 20% faster when comparing the non-turbo ESS and one with the Turbo II feature. The Turbo II feature will usually enable better performance and higher throughputs after the failover is complete. Performance during a concurrent microcode upgrade should also be somewhat faster with Turbo II for the same reasons. Customers that require robust failover performance are advised to use the Turbo II feature.

PPRC over Fibre Channel Performance

Disaster recovery (DR) is a fundamental business need for nearly any organization that has modern computing facilities. ESS has several advanced functions that can facilitate the implementation of DR solutions. One of these functions is Peer-to-Peer Remote Copy (PPRC). Although PPRC is not new, PPRC over FCP is now available for the ESS model 800 with LIC level 2.3.0. For a complete description of ESS PPRC performance with ESCON links, please see the papers listed in the reference section.

PPRC is a method of remotely replicating data. With synchronous PPRC, data written to a primary disk storage system is synchronously transmitted to a remote disk storage system. There is also a flavor of PPRC called PPRC-XD (XD for extended distance) which provides for asynchronous replication even at very long distances.

With the availability of LIC level 2.3.0, IBM added the capability to us fibre channel connections for PPRC links. Prior versions of PPRC utilized ESCON links. ESCON is a 200 Mb/sec serial channel which can achieve instantaneous data transfer rates of up to about 18 MB/sec. ESCON channels are single threaded -- they only allow one I/O to use the channel at a time. As a result additional PPRC write concurrency could only be achieved by using more ESCON links. For ESCON links on the ESS 800, distance can add about 20 microseconds per km of propagation delay to the PPRC write service time on a logical track basis. Because the ESCON link is single threaded, no other I/O can use the channel during the propagation delay. As a result, there is a substantial drop in throughput on an ESCON channel as the distance increases.

Fibre channel links are a significant technological improvement over ESCON. The ESS host adapter supports full duplex fibre channel operation at 2 Gb/sec. The instantaneous data transfer rate is up to about 200 MB/sec. Fibre channels are multi-threaded so many I/Os can operate concurrently on a single fibre channel link. The protocol used with fibre channel is more efficient than ESCON so only about 10 microseconds per km of propagation delay is added to the PPRC write service time at distance. Multi-threading helps prevent the aggregate throughput capability of a fibre channel link from degrading with distance. Theoretically, this applies to any distance; however IBM requires an RPQ for PPRC distances beyond 103 km so the anticipated I/O response time implications may be reviewed with the customer before implementing. Practically speaking, synchronous PPRC may provide acceptable performance for many users even at distances of up to 400 km when using fibre channel links with dark fiber or Dense Wave Division Multiplexing (DWDM). Please note that repeaters may be needed for fibre channel to function at extended distances.

The introduction of fibre channel links for ESS allows a substantial reduction in the infrastructure required for implementing a PPRC solution. This is because a single fibre channel link, depending on the workload, can handle between 4 and 12 times the throughput of an ESCON link at 0 km based on laboratory measurements. At distance, the improvement ratios can be even greater.

Establishing PPRC pairs is the first step needed to use PPRC. In some instances such as data migration or creating a remote point-in-time copy, it is the only step needed. Figure 4 compares ESCON and fibre channel link throughput for making initial PPRC copies of both Open (fixed block) and z/OS count key data (CKD) formatted volumes. In general, the fibre channel links are roughly 8 times faster than ESCON links for this type of activity. Figure 5 shows that even at 75 km distance the link throughput is roughly the same as at 0 km provided that at least 4 volumes are copied at the same time.

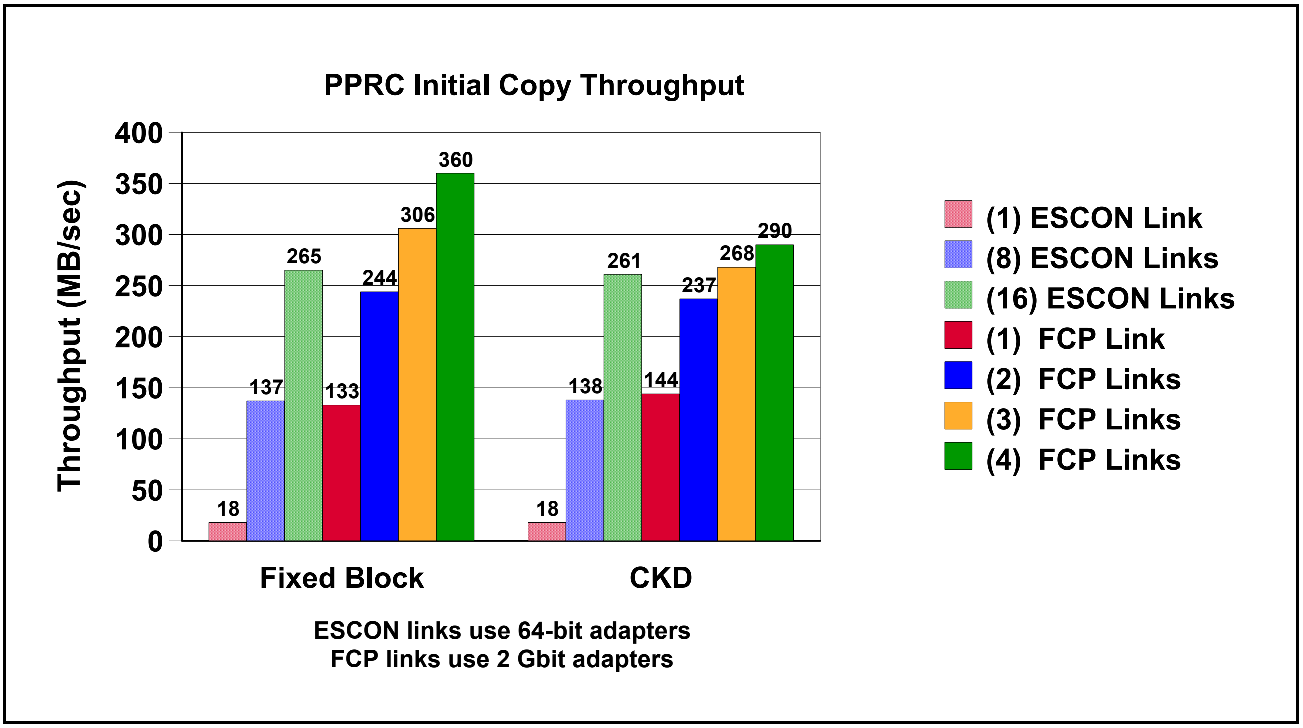

Figure 4 - PPRC Initial Copy Throughput Capability, ESCON Links vs. Fibre Channel Links

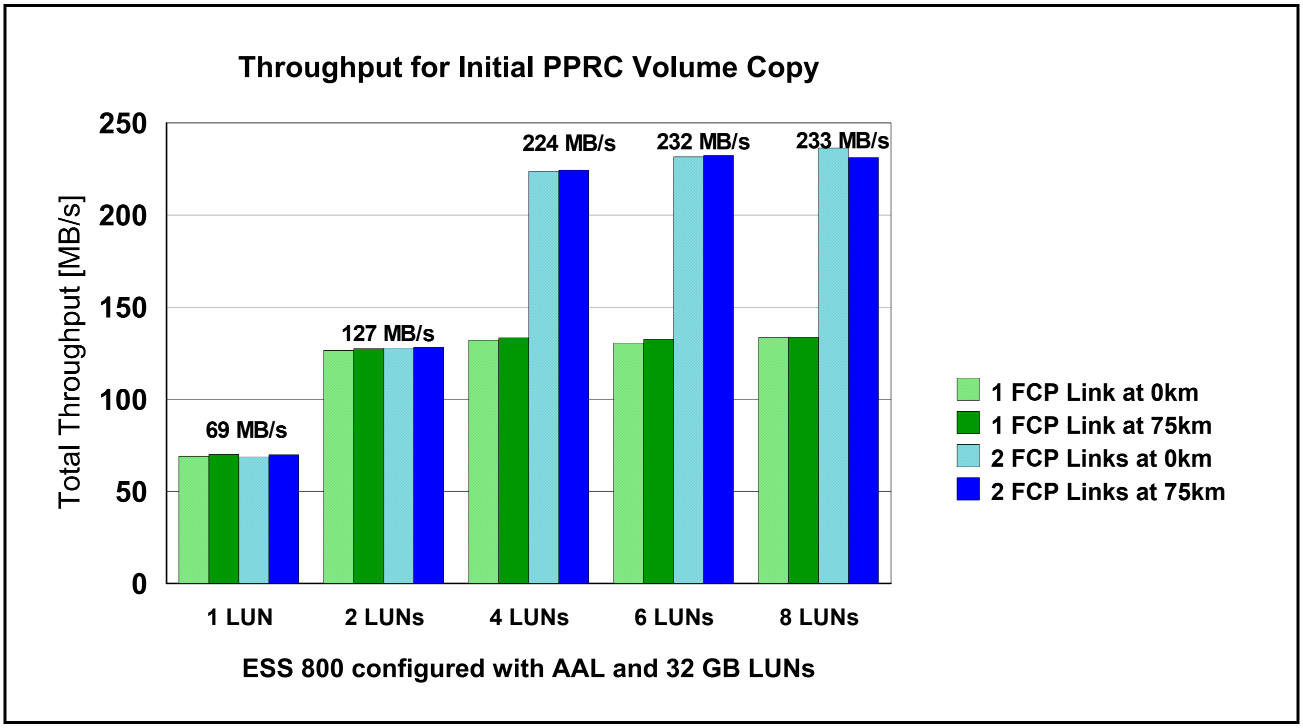

Figure 5 - PPRC Initial Copy Throughput Capability, Fibre Channel Links at 0 km vs. 75 km

For many applications, PPRC link throughput for normal operation is more important than initial copy throughput. Figure 6 compares ESCON and fibre channel link throughput for sequential writes with PPRC. Data for both fixed block and CKD environments are included. Sequential writes are usually a major component of batch processing. For this type of workload, the fibre channel links are roughly 4 times faster than ESCON links.

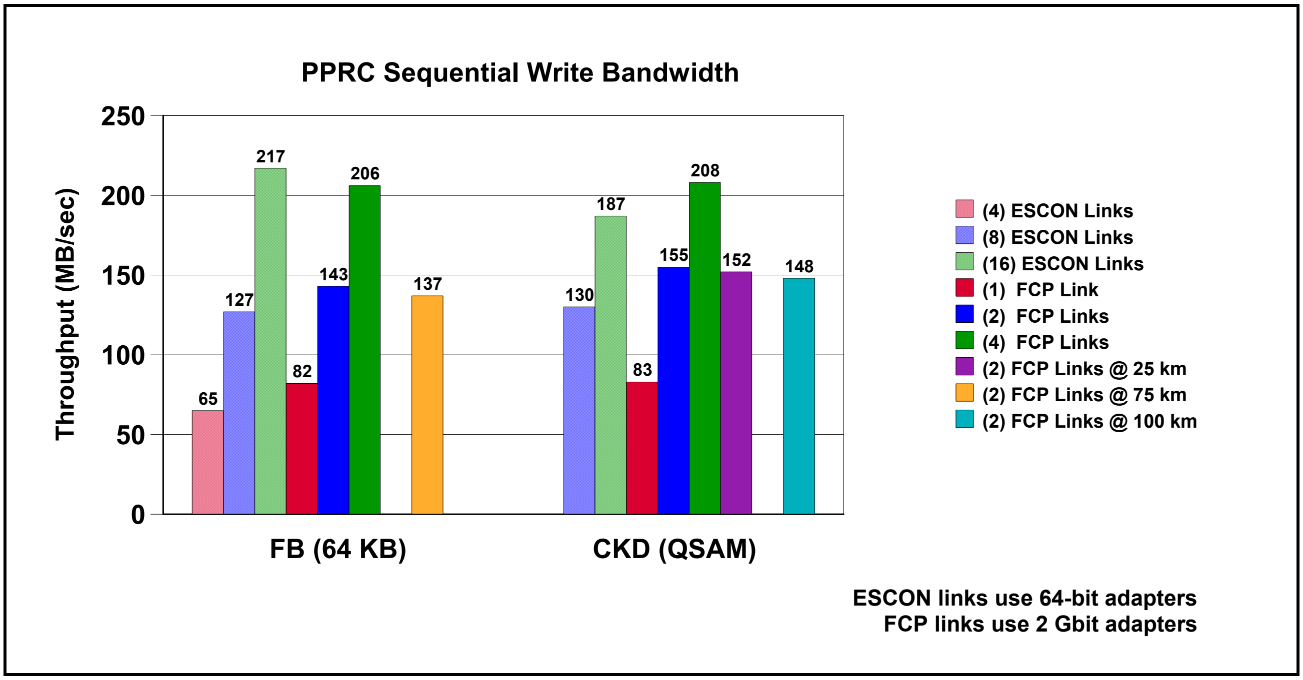

Figure 6 - PPRC Sequential Write Throughput Capability, ESCON Links vs. Fibre Channel Links

Figure 7 shows details of fixed block sequential throughput measurements for PPRC at both 0 km and 75 km. Distance does not introduce any noticeable decrease in aggregate link throughput although more concurrent streams are needed to saturate the links as distance increases.

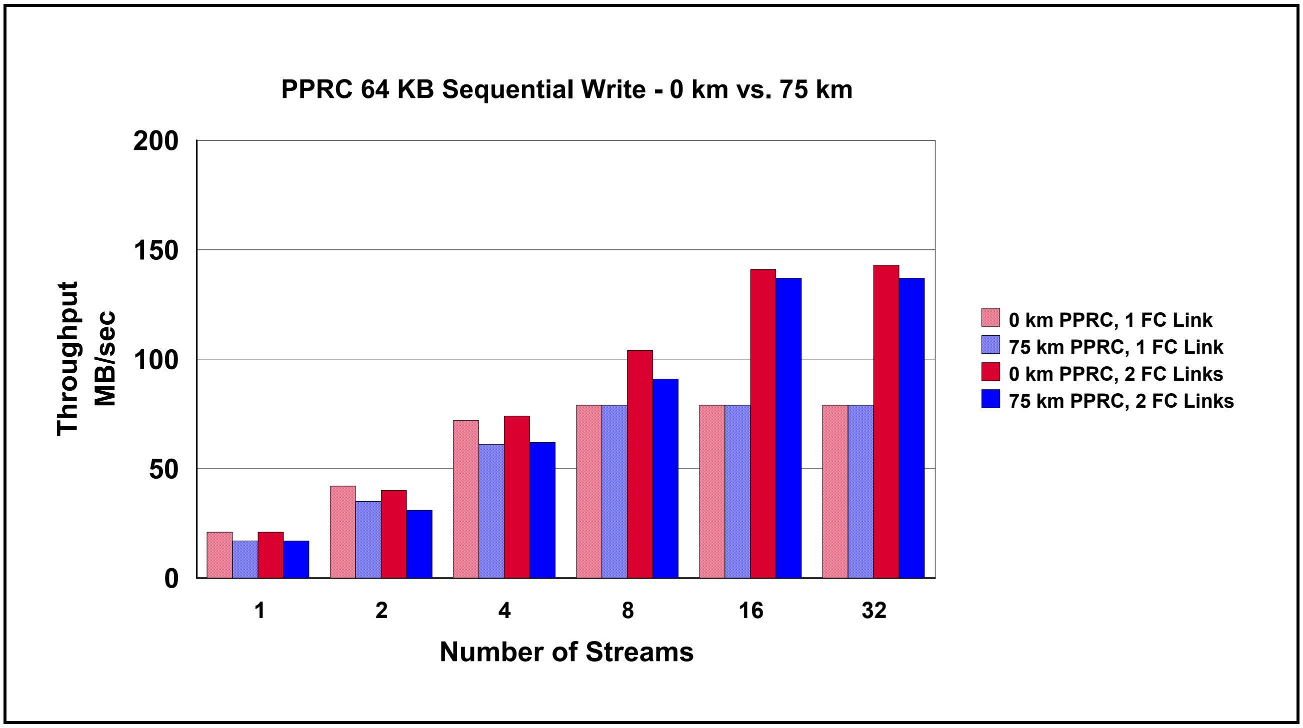

Figure 7 - PPRC Sequential Write Throughput vs. Number of Streams for Open Systems

Figure 8 shows a comparison of ESCON and fibre channel link throughput for 4 KB random write hits in an Open systems (fixed block) environment. Small block writes are often an important part of many on-line transaction processing (OLTP) workloads. Figure 9 has similar data for z/OS (CKD) workloads. For this type of workload, a single fibre channel link has nearly as much throughput as 16 ESCON links at 0 km distance. When more than one link is used, the throughput does not double. This is because the ESS SMPs become saturated before two links are fully utilized. The highest PPRC throughput for small block writes can be achieved by using the Turbo II feature on the ESS 800.

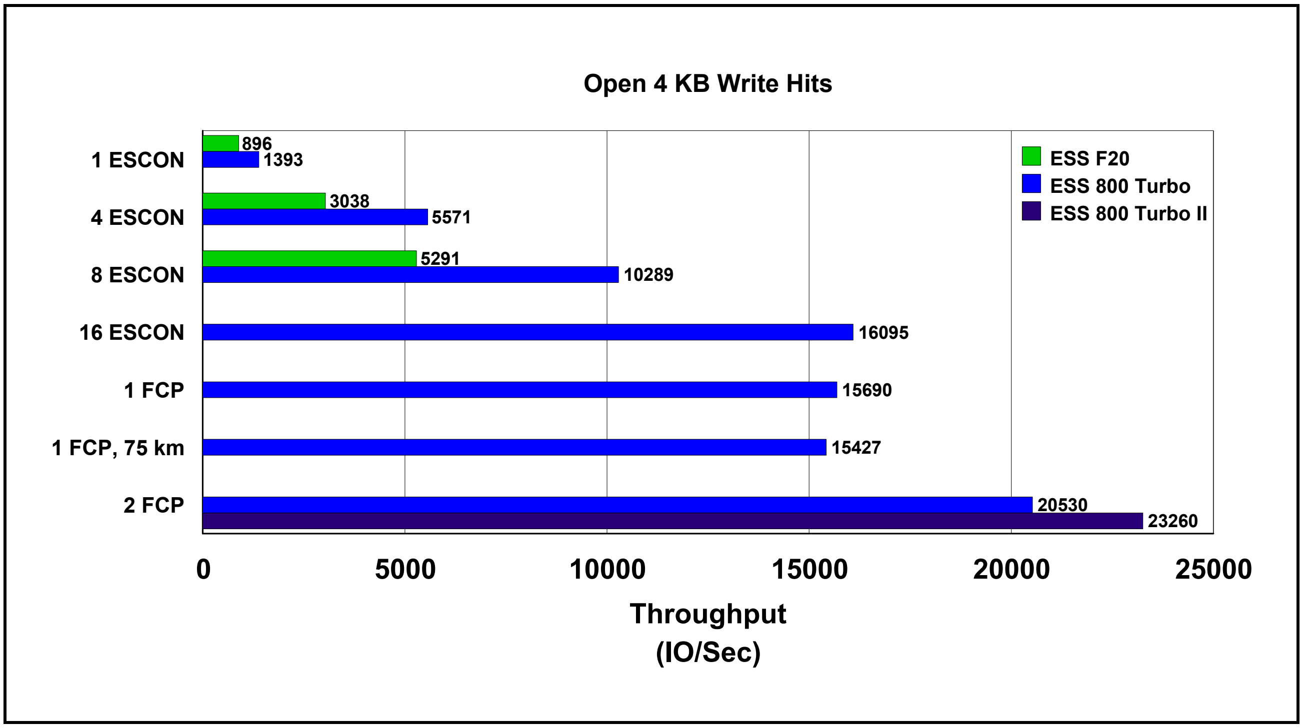

Figure 8 - PPRC 4 KB Write Hit Throughput for Open Systems

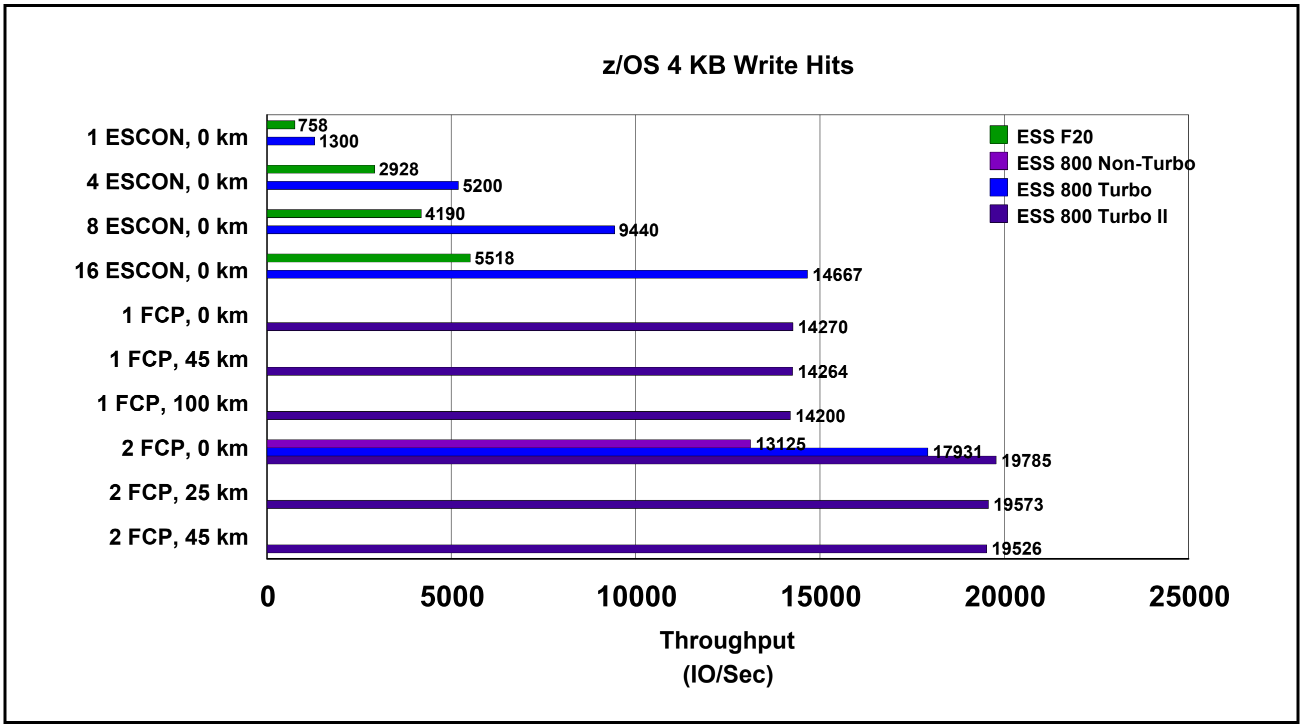

Figure 9 - PPRC 4 KB Write Hit Throughput for z/OS

It is interesting to note that the fixed block and CKD test results are very similar to each other for both initial copy, sequential writes and write hits. This is not at all unusual. In many cases fixed block performance measurements on the ESS are numerically similar to CKD performance measurements of comparable workloads. How the data is formatted is not terribly relevant to disk hardware performance.

Many real world user workloads will never exercise the maximum throughput capabilities of the ESS with fibre channel PPRC links. A single fibre channel link is often sufficient to meet performance requirements. However, it is suggested that two links always be configured to provide redundancy in the event of a link failure. Two fibre channel links per ESS will be more than adequate for the vast majority of PPRC implementations. With ESCON links, eight was the usual rule of thumb and sometimes even more were needed at longer distances Thus, fibre channel PPRC links enables at least a 4 to 1 reduction in fiber infrastructure needed for PPRC.

Another way to reduce infrastructure is by using switches and fibre channel Inter-Switch Links (ISLs). For sequential writes, a single link can be driven to about 83 MB/sec. The bottleneck is the ESS host adapter (HA), not the fibre channel itself. Figure 10 shows a suggested configuration that maps four ESS host adapters to two fibre channel ISLs. This supports a combination of efficiency and redundancy. By configuring this way, each ISL may reach throughputs up to 163 MB/sec. This strategy can be especially effective at long distances where the cost of dark fiber could be greater than the cost of the rest of the PPRC solution hardware and software.

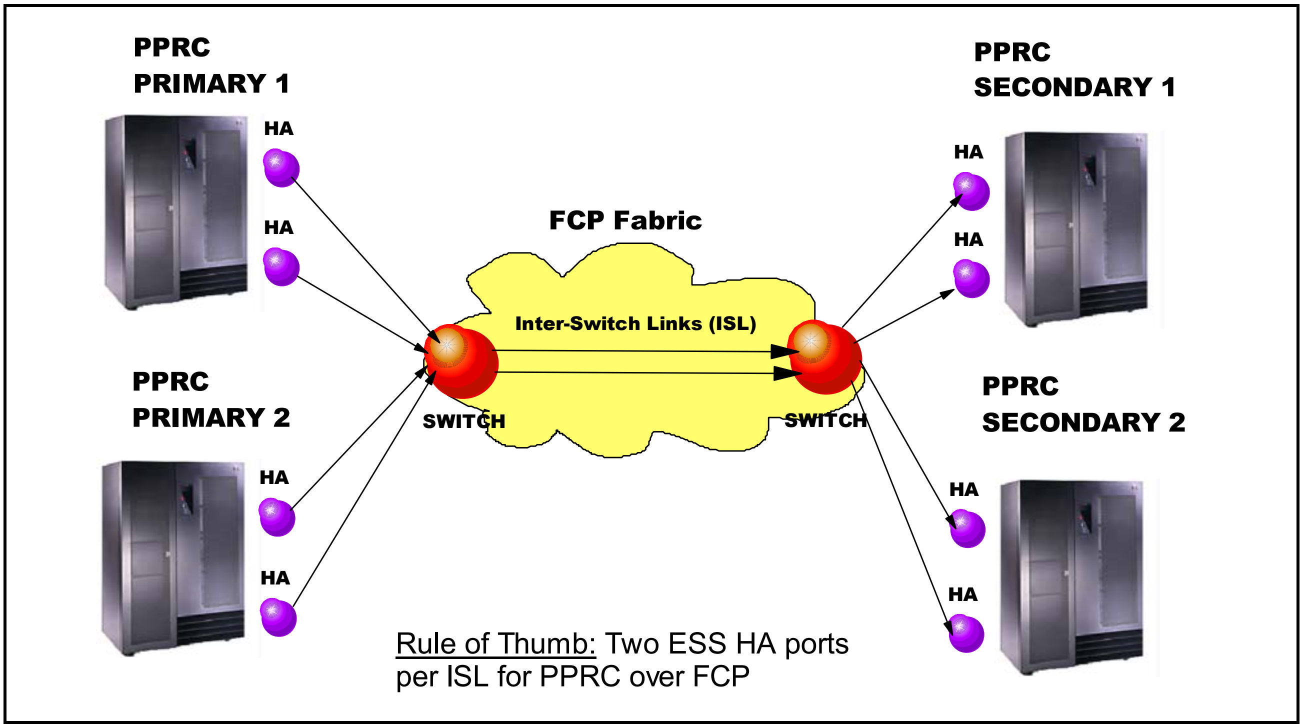

Figure 10 - Suggested PPRC Topology with Switches and ISLs

Note that initial copy can drive the ESS host adapters to much higher throughput levels (about 140 MB/sec) than sequential writes. During normal PPRC operation data must be written to both the cache and NVS on the secondary ESS to support data integrity. During initial copy, data only needs to be written to the cache on the secondary ESS. In the event of an unexpected cache failure during initial copy, the copy will fail and need to be restarted.

Besides throughput, disk response time is another important performance metric for PPRC. Using synchronous PPRC will always increase the response time of writes, even at 0 km distance. Reads are generally unaffected. There are four components to the response time increase: 1) Base overhead time for executing the PPRC write on the local and remote disk systems. 2) Data transfer time that depends on the instantaneous link datarate. 3) Propagation delay that is a function of the distance between the local and remote disk systems. 4) Queuing delays which are caused by contention for PPRC links or other components within the system. ESS PPRC is specifically designed to avoid volume contention. Multiple writes to the same ESS source volume may all proceed concurrently. Some competitive storage systems serialize synchronous remote copy writes on a volume basis.

The base overhead of PPRC on the ESS 800 is roughly the same for both ESCON and fibre channel links but the instantaneous data rate is much faster with fibre channel. For example, a 4 KB transfer takes about 200 us less time with fibre channel than with ESCON links. A 64 KB transfer may take up to 3 ms less time with a fibre channel link. At 100 km distance, PPRC over a fibre channel link will add 1 ms of latency to each PPRC write due to propagation delay. With ESCON links, the latency would be twice as long (2 ms). Because fibre channel is a multi-threaded protocol, there is less likelihood of contention for links which might otherwise cause queuing delays. As a result, the disk response time for PPRC writes are generally substantially lower with fibre channel links than it would be with ESCON links.

Figures 11 (Fixed Block) and 12 (CKD) compare response time measurements for base (no PPRC) as well as PPRC with both ESCON and fibre channel PPRC links at various distances. In all cases the response time with fibre channel links will be lower than with ESCON links. As the distance and transfer size increase, the response time improvement with fibre channel links becomes more dramatic.

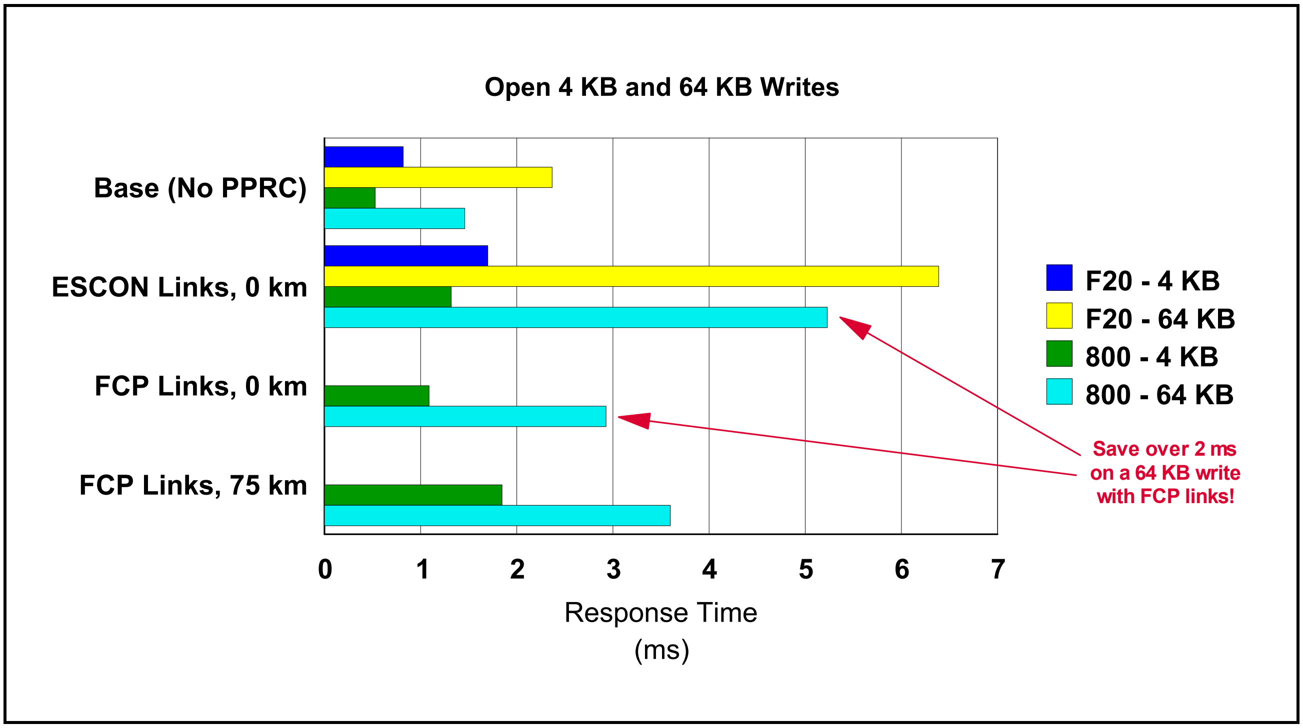

Figure 11 - PPRC Response Times for Open Systems

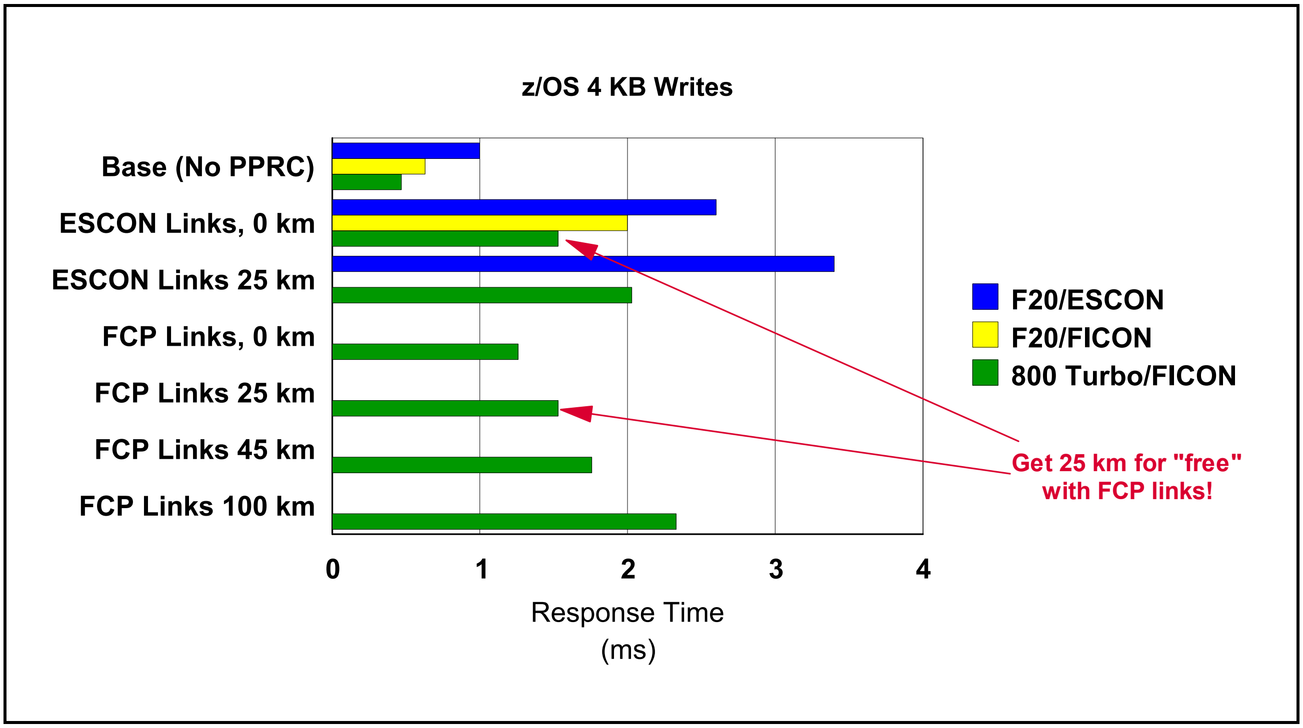

Figure 12 - PPRC Response Times for z/OS

Related to response time is single stream throughput. Aggregate throughput tells you how much total work can be done on a particular link. Single stream throughput tells you how fast a particular job or task may run. As response times rise, single stream throughput will decrease. Figures 13 (Fixed Block) and 14 (CKD) show single stream throughput measurements for ESS PPRC with both fibre channel and ESCON links. When ESCON links are used with either fibre channel or FICON host channels, there is a mismatch between the front end channel speed and the PPRC link speed. The single stream throughput is dominated by the slower ESCON link speed. With fibre channel PPRC links, the channels and links have the same raw datarates. This is reflected in the higher single stream datarates possible with fibre channel links. The single stream throughput does decrease with distance because the response time increases due to propagation delay.

Consider the result in figure 13 for ESS 800T (T for Turbo option) Base (no PPRC) of 42 MB/sec. Turning on PPRC at 0 km cuts the single stream rate in half. This makes sense since with PPRC you need to write the data twice so you should expect it to take about twice as long. When response time doubles, the single stream rate will be halved.

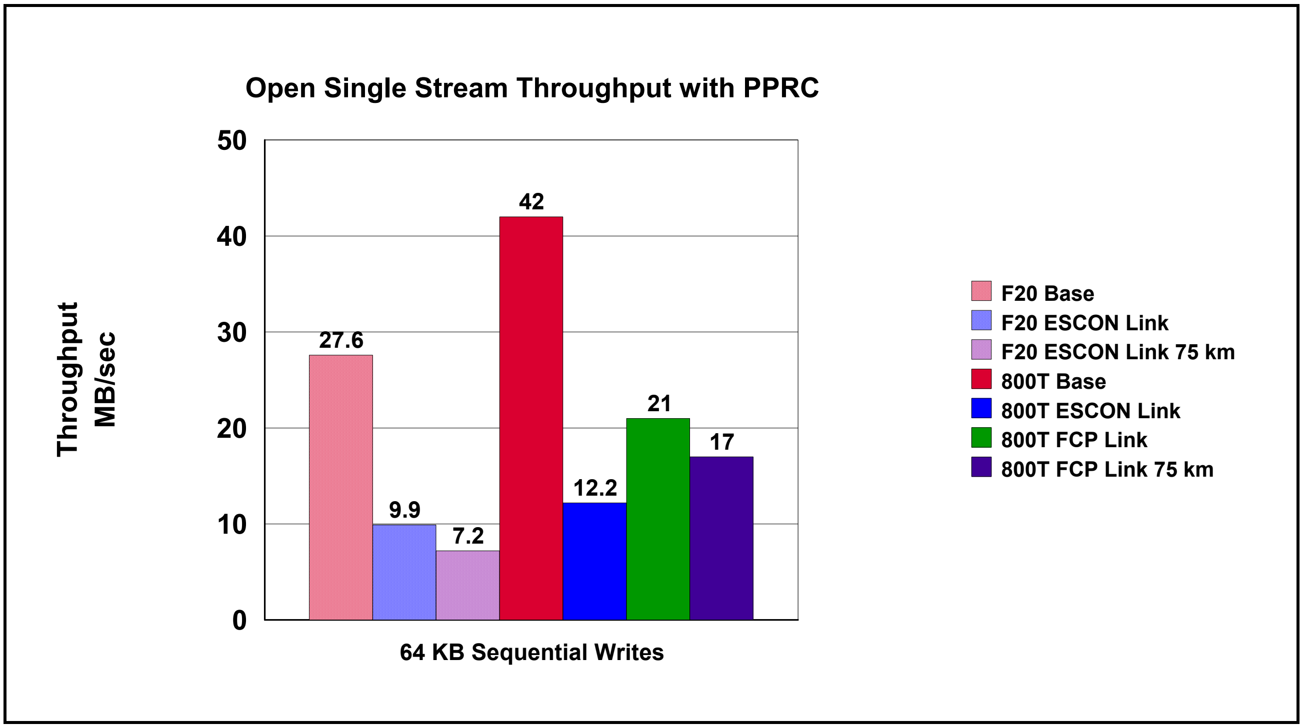

Figure 13 - PPRC Single Stream Throughput for Open Systems

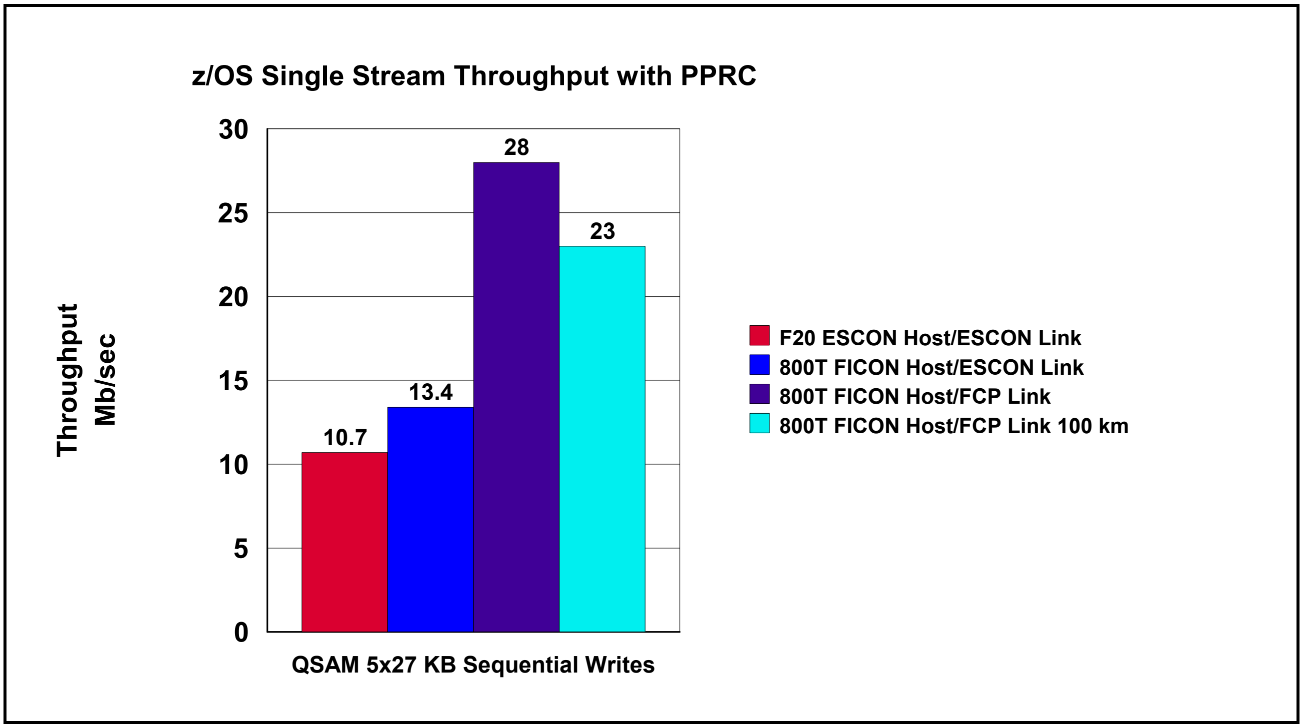

Figure 14 - PPRC Single Stream Throughput for z/OS

Although generic benchmark measurements may provide a good indication of PPRC performance, most users are more concerned with how PPRC may affect the performance of real applications. A series of test runs were made with Tivoli Storage Manager (TSM) doing a database restore and expiration processing. Two 32 GB logical volumes were used. Figure 15 shows the total run time for the series of TSM tests both with and without remote copy. The data shows that run times increase when using PPRC or AIX Logical Volume Mirroring (LVM). Compared to the no remote copy case, the run time more than doubles with LVM at 75 km distance but increases less than 50% using PPRC at the same distance. For this type of application, PPRC may be a better performance choice than LVM.

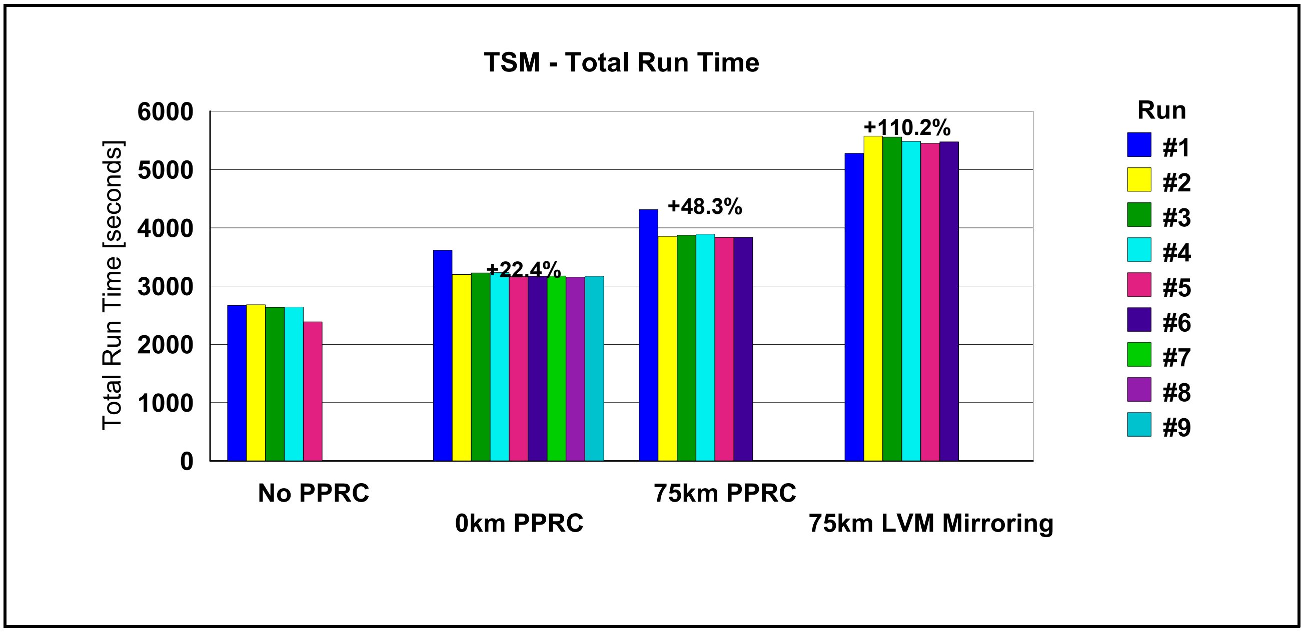

Figure 15 - TSM Application Tests with PPRC versus AIX LVM

Arrays Across Loops

Arrays across loops (AAL) is a new way of structuring RAID arrays to better exploit the available hardware in the ESS. Each device adapter on the ESS has two Serial Storage Architecture (SSA) loops attached. Without arrays across loops, each RAID array is connected to only one of these loops. By using arrays across loops, a single RAID array is spread across both SSA loops. Figure 16 compares a non-AAL array structure to AAL. The result is that single RAID arrays can achieve higher sequential bandwidths. The sequential bandwidth capability of a full ESS is unchanged.

It is recommended that new ESS 800 installations be configured with AAL whenever possible. AAL can help to provide the following benefits:

- Using AAL may help simplify performance management. With AAL, a single array can do more work. Thus if an array contains many volumes with high bandwidth demands, it is not usually necessary to move data to other arrays to balance the workload if AAL is in use.
- AAL may help reduce the need for implementing functions such as software striping. Software striping is sometimes used to get higher throughput for critical jobs. AAL supports higher throughput without the additional management effort.
- AAL can deliver bandwidth to demanding applications where and when it is needed. It supports autonomic management of array bandwidth and can help customers save time and effort. Real world workloads are often imbalanced. With AAL, there is less worry about manually balancing the workload across RAID arrays.

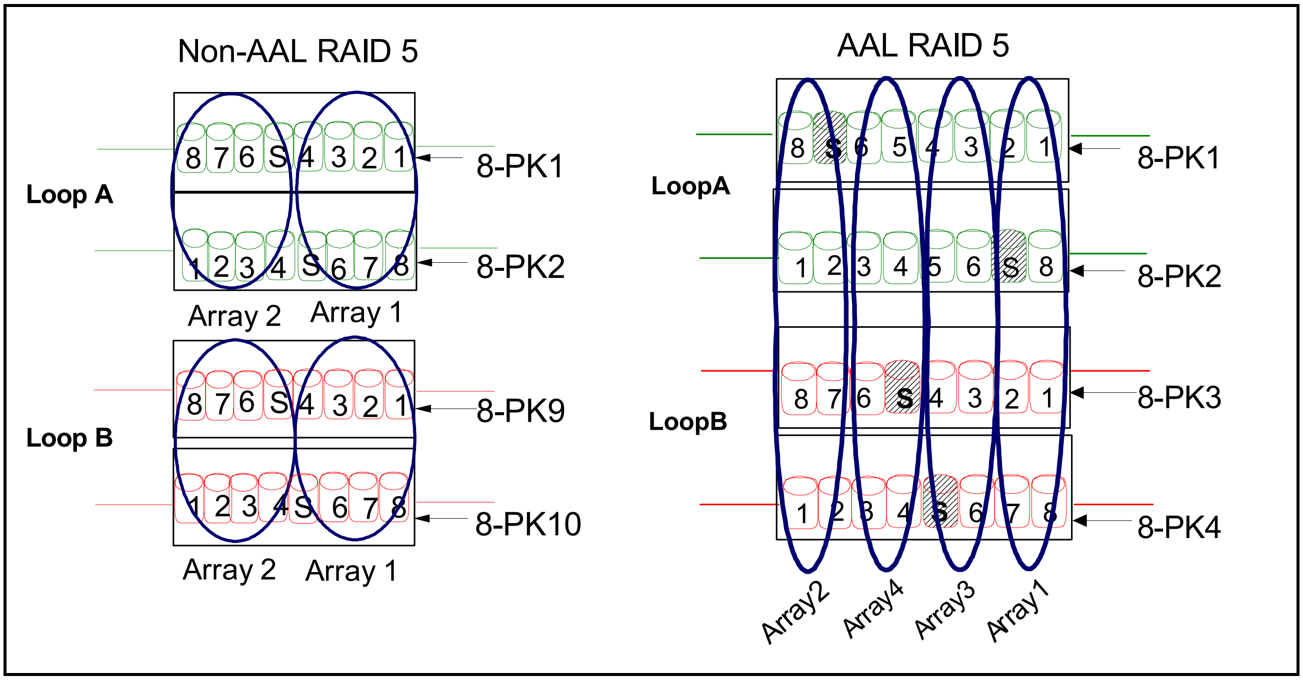

Figure 16 - Non-AAL versus AAL

Figure 17 compares single array sequential throughput measurements with and without AAL. Please note that once again the results for Open systems are very similar to z/OS. Improvements of about 60% for writes and 70% to 80% for reads may be achieved.

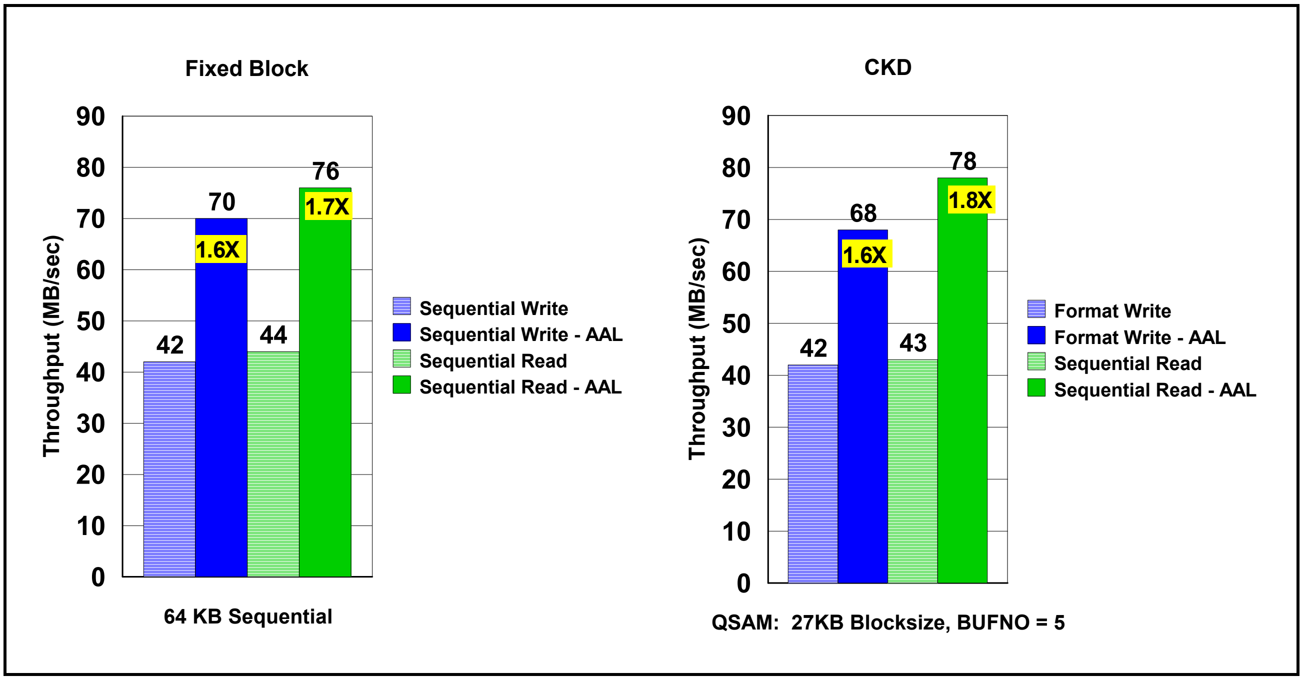

Figure 17 - ESS 800 Single Array Throughput with RAID 5

AAL also enables better performance for an individual sequential job. This is sometimes referred to as single stream throughput. A single job may be on the critical path for batch processing. Faster single stream throughput can make the difference between meeting a batch service level agreement or missing it. Figure 18 compares single stream read sequential throughput measurements with and without AAL. Again, results for Open systems and z/OS are very similar. Single stream throughput increases by about 40% with AAL.

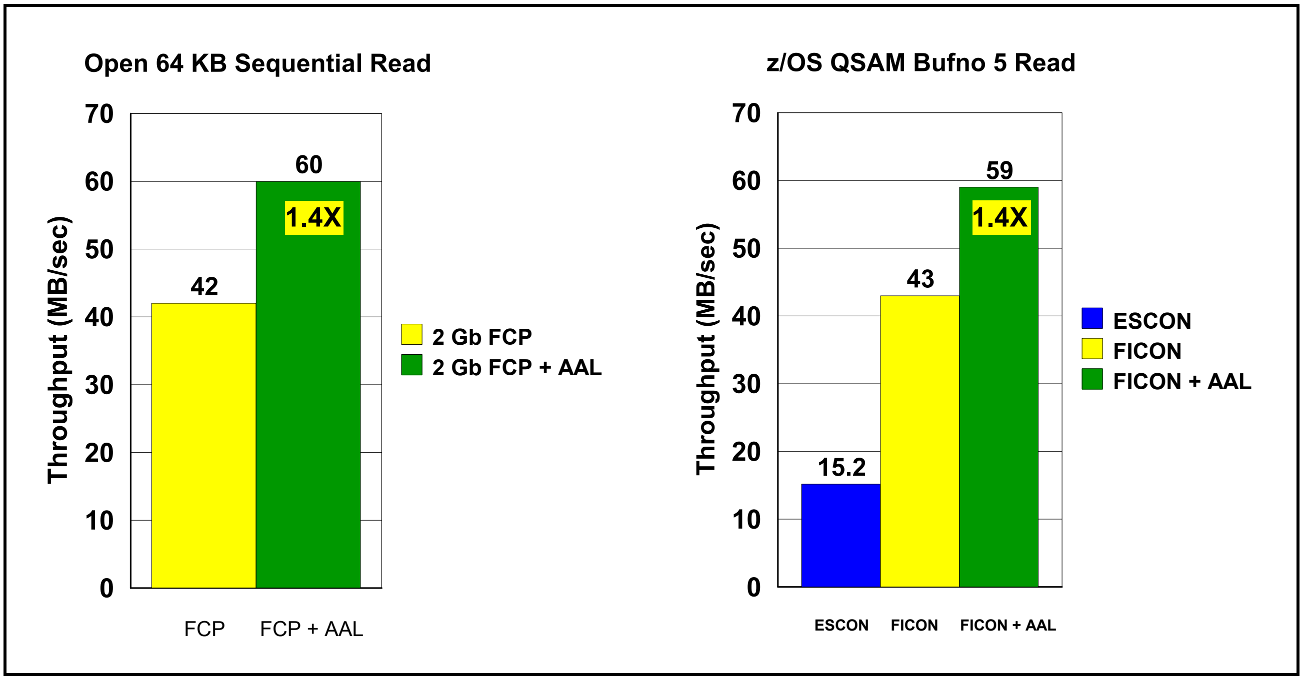

Figure 18 - ESS 800 Single Stream Sequential Throughput with RAID 5

A practical example of an application that may benefit from AAL is volume dumps. A single z/OS 3390-9 volume (8.514) can be dumped to 3590 tape from an ESS with AAL in about two minutes. Without AAL, the time will be nearly twice as long. Comparable improvements are expected when dumping Open (fixed block) LUNs.

With AAL, additional capacity must be added to the ESS in groups of four 8-packs instead of the two 8-pack increment that is required in non-AAL configurations. This may not allow the granularity that some customers require and should be taken into consideration before this option is chosen. Some customers may prefer not to configure with AAL because of this requirement.

AAL may not provide any performance benefit for very small numbers of disks and in some cases it will actually be sub-optimal. This is because some of the device adapters in the ESS may wind up being unused If an ESS configuration is never expected to grow beyond eight 8-packs (64 disks), it is recommended that AAL not be ordered. If capacity growth is anticipated, then AAL is usually the best choice.

AAL is a plant initial configuration option for new ESS 800 orders. It is not available as an MES upgrade, and once installed, it cannot be easily removed. Converting an existing ESS from AAL to non-AAL or vice versa is disruptive. It requires you to move all the data off the ESS, then reconfigure and move the data back, and would require a billable ITS services contract. For this reason, it is likely that most customers would choose not to implement AAL on an existing ESS or change an existing AAL configuration back to a non-AAL configuration. However, it is a performance enhancing option that should be strongly considered for nearly all new ESS 800 installations.

Disk Magic

Disk Magic is a performance modeling tool used by IBM which can help you estimate the expected performance of an ESS 800 with a specific configuration running a specific customer workload. The tool can be used to estimate performance characteristics for the Turbo II feature and for PPRC over fibre channel at various distances. It is not able to make projections for AAL since the tool is designed to model a full ESS as opposed to individual RAID arrays within an ESS.

As with any modeling tool, the quality of the output will depend on the accuracy of the input data describing the workload characteristics. Customers should contact their IBM Storage Sales Specialist or IBM Business Partner if they are considering the Turbo II feature or PPRC over fibre channel on ESS 800 and would like to have a performance modeling study done.

Frequently Asked Questions

Turbo II Feature

Q. Can I use a Turbo II feature with 145.6 GB disk instead of a non-turbo model with 72.8 GB disk and get the same performance?
A. Use Disk Magic to evaluate this. The answer depends on how busy the 146 GB disk will be. Remember that a Turbo II does nothing to make the underlying disk perform better.

PPRC over Fibre Channel Protocol

Q. If I use a fibre channel switch can I share host access and PPRC links on the same ESS host adapter?
A. This is a supported configuration. However, if PPRC is used extensively it is usually a good idea to have separate fibre channel connections for host attachment and links. This will help reduce configuration complexity. Sharing should only be considered if more ports are needed to support the desired host connectivity.

Arrays Across Loops

Q. Does the ESS Expert support AAL configurations?
A. No, not at this time. ESS Expert changes to support AAL configurations are planned and expected to be available in the first half of 2004.

Q. Can I configure a portion of an ESS for AAL and the remainder without AAL?
A. No. Although it may be technically possible to do this, it is unsupported and untested by IBM. An ESS must be configured either completely with or without AAL.

Conclusions

Features available in Enterprise Storage Server LIC Level 2.3.0 include three important performance enhancements. The Turbo II option provides more processing power to handle high throughput OLTP workloads and heavy copy services requirements. PPRC over fibre channel improves performance over long distances and reduces the infrastructure needed to support PPRC implementations. AAL is a new way to configure RAID arrays on the ESS which enables higher sequential throughput on a single array or single job stream. Using AAL does not increase the overall ESS sequential bandwidth but helps the ESS autonomically manage imbalanced sequential workloads.

Acknowledgments

Many thanks to Sonny Williams, Preston Rivers, Andrew Lin, Gero Schmidt, Peter Kimmel, Shona Robertson, Frank Krueger, Terese Souvignier, Rosemary McCutchen, and Craig Gordon for making performance measurements that are included in this paper. Thanks also to Brian Sherman, John Ponder, Steve West, Siebo Friesenborg, Hank Heinz, Alice Bird, and David Sacks for their comments, insights and assistance

References

[1] L. LaFrese, “IBM TotalStorage Enterprise Storage Server 800 145.6 GB Disk 8-packs and 64-bit ESCON Host Adapters Performance White Paper”, IBM Intranet, January 2003

[2] L. LaFrese, “IBM TotalStorage Enterprise Storage Server Peer to Peer Remote Copy Extended Distance Performance White Paper”, IBM Intranet, June 2002

[3] L. LaFrese, “IBM Enterprise Storage Server Peer to Peer Remote Copy Performance White Paper”, IBM Intranet, December 2000