



IBM® TS7760 and TS7760T
Release 4.0
Performance White Paper
Version 2.0

By Khanh Ly
Virtual Tape Performance
IBM Tucson

Table of Contents

Table of Contents2

Introduction.....3

TS7700 Performance Evolution5

TS7700 Copy Performance Evolution7

Hardware Configuration.....9

TS7700 Performance Overview10

TS7700 Basic Performance.....13

Additional Performance Metrics26

Performance Tools30

Conclusions31

Acknowledgements.....32

Product
Release
Enhancements

Introduction

This paper provides performance information for the IBM TS7760 and TS7760T, which are two current products in the TS7700 family of storage products. This paper is intended for use by IBM field personnel and their customers in designing virtual tape solutions for their applications.

This is an update to the previous TS7700 paper dated May 07, 2015 and reflects changes for release 4.0, which introduces the TS7760 and TS7760T. TS7760T supports the ability to connect a TS4500 or a TS3500 tape library to a TS7760. Like a TS7720T, TS7760T can perform the function of a TS7760 or TS7740 depending on the target partition specified in the customer's workload. Up to 8 target partitions could be defined in a TS7760T, namely CP0 through CP7. When workload targets CP0, the TS7760T behaves as a TS7760. When workload targets CPn (n = 1 through 7), the TS7760T behaves as a TS7740.

Unless specified otherwise, all runs to a TS7760T in this white paper target a 50TB CP1 tape managed partition with ten FC 5274.

Notes:

FC 5274 is the feature code used to manage the premigration mechanism in the TS7760T_{cp1-7}. Within a TS7760T, there is a global premigration queue for all tape partitions. The FC5274 limits the maximum amount of queued premigration content within a TS7760T. The features will be in 1 TB increment with a maximum of 10TB for all partitions. The priority and premigration throttle thresholds (PMPRIOR and PMTHLVL) can not exceed this limit.

The following are performance related changes in release 4.0:

- ❖ Server Refresh to new Power8 pSeries with:
 - Two 3.42GHz processors with 10 processor cores.
 - 32 GB of DDR3 CDIMM @ 1600 Mbps
- ❖ Disk Cache Refresh with larger capacity and better performance.
- ❖ TS4500 Tape Library Support
- ❖ Four 10 Gb Grid Ethernet Links
- ❖ 16 Gb Fibre Channel Adapter attached to backend disk cache
- ❖ 16 Gb Fibre Channel Switch attached to physical tape drives

TS7700 Release 4.0 Performance White Paper Version 2 adds data for the following configurations:

- ❖ Two-way hybrid grid VEC-T/CSA/26 Drawers/12 EH8/4x10Gb links.
- ❖ Three-way hybrid grid VEC-T/CSA/26 Drawers/12 EH8/4x10Gb links.
- ❖ Four-way hybrid grid VEC-T/CSA/26 Drawers/12 EH8/4x10Gb links.

The following conventions are used in this paper:

<i>Binary</i>			<i>Decimal</i>		
<i>Name</i>	<i>Symbol</i>	<i>Values in Bytes</i>	<i>Name</i>	<i>Symbol</i>	<i>Values in Bytes</i>
<i>kibibyte</i>	<i>KiB</i>	2^{10}	<i>kilobyte</i>	<i>KB</i>	10^3
<i>mebibyte</i>	<i>MiB</i>	2^{20}	<i>megabyte</i>	<i>MB</i>	10^6
<i>gibibyte</i>	<i>GiB</i>	2^{30}	<i>gigabyte</i>	<i>GB</i>	10^9
<i>tebibyte</i>	<i>TiB</i>	2^{40}	<i>terabyte</i>	<i>TB</i>	10^{12}
<i>pebibyte</i>	<i>PiB</i>	2^{50}	<i>petabyte</i>	<i>PB</i>	10^{15}

TS7700 Performance Evolution

The TS7700 architecture continues to provide a foundation for product growth in both performance and functionality. Figures 1 and 2 show the write and read performance improvement histories.




Figure 1. VTS/TS7700 Standalone Maximum Host Write Throughput. All runs were made with 128 concurrent jobs, using 32KiB blocks, and QSAM BUFNO = 20. The volume size is 2659 MiB (1000 MiB volumes @ 2.66:1 compression).

Notes:

- nDRs : number of cache drawers
- m x m Gb:
 - 4x4Gb -- four 4Gb FICON channels
 - 8x8Gb -- eight 8Gb FICON channels (dual ports per card)
 - 4x1x8Gb -- four 8Gb FICON channels (single port per card)
- CS9* -- the new higher-performance 3TB cache drive
- TS7720T_{cp1} -- cache partition 1 on the TS7720T
- TS7760T_{cp1} -- cache partition 1 on the TS7760T

Standalone
Performance
Evolution




Figure2. VTS/TS7700 Standalone Maximum Host Read Hit Throughput. All runs were made with 128 concurrent jobs, using 32KiB blocks, and QSAM BUFNO = 20. The volume size is 2659 MiB (1000 MiB volumes @ 2.66:1 compression).

See definition for Read Hit in the section “TS7700 Performance Overview”.

Notes:

- nDRs : number of cache drawers
- m x m Gb:
 - 4x4Gb -- four 4Gb FICON channels
 - 8x8Gb -- eight 8Gb FICON channels (dual ports per card)
 - 4x1x8Gb -- four 8Gb FICON channels (single port per card)
- CS9* -- the new higher-performance 3TB cache drive
- TS7720T_{cp1} -- cache partition 1 on the TS7720T
- TS7760T_{cp1} -- cache partition 1 on the TS7760T

TS7700 Copy Performance Evolution

Figures 3 through 4 display deferred copy rates. Data rate over the grid links are of compressed data. In each of the following runs, a deferred copy mode run was ended following several terabyte (TB) of data being written to the active cluster(s). In the subsequent hours, copies took place from the source cluster to the target cluster. There was no other TS7700 activity during the deferred copy except for premigration if the source or target cluster was a TS7740, a TS7720T, or a TS7760T. The premigration activity consumes resources and thus lower the copy performance on the TS7740, TS7720T, or TS7760T as compared to the TS7720 or TS7760.

The 8 Gb FICON requires an additional 16GB of memory (total 32GB) which accounts for the copy performance improvement with 8Gb FICON.

Copy
 Performance
 Evolution




Figure 3. Two-way TS7700 Single-directional Copy Bandwidth.




Figure 4. Two-way TS7700 Bi-directional Copy Bandwidth

Hardware Configuration

The following hardware was used in performance measurements. Performance workloads are driven from IBM System z196 or zEC12 host with eight 8 Gb FICON channels.

Standalone Hardware Setup

TS7700	Drawer count (cache size TB)	Tape lib	Tape Drives	IBM System z™ Host
TS7760 VEC 3956 CSA/XSA	10 (313.96 TB)	N/A	N/A	zEC12
	26 (816.49 TB)		N/A	
TS7760T VEC-T 3956 CSA/XSA	1 (31.11 TB)	TS4500	8 TS1150	z196
	3 (93.96 TB)			
	5 (156.82 TB)			
	7 (219.68 TB)			
	10 (313.96 TB)		12 TS1150	zEC12
	26 (816.49 TB)			

Grid Hardware Setup

TS7700	Drawer count (cache size TB)	Tape Lib	Tape Drives	Grid link (Gb)	IBM System z™ Host
TS7760 VEC 3956 CSA/XSA	10 (313.96 TB)	N/A	N/A	2x10	zEC12
	26 (816.49 TB)			4x10	
TS7760T VEC-T 3956 CSA/XSA	10 (313.96 TB)	TS4500	12 TS1150	2x10	
	26 (816.49 TB)			4x10	

TS7700 Performance Overview

Performance Workloads and Metrics

Performance shown in this paper has been derived from measurements that generally attempt to simulate common user environments, namely a large number of jobs writing and/or reading multiple tape volumes simultaneously. Unless otherwise noted, all of the measurements were made with 128 simultaneously active virtual tape jobs per active cluster. Each tape job was writing or reading 2659 MiB of uncompressed data using 32 KiB blocks and QSAM BUFNO=20 that compresses within the TS7700 at 2.66:1. Measurements were made with eight 8-gigabit (Gb) FICON channels on a z196 or zEC12 host. All runs begin with the virtual tape subsystem inactive.

Unless otherwise stated, all runs were made with default tuning values:

- ✓ DCOPYT=125,
- ✓ DCTAVGTD=100,
- ✓ PMPRIOR=1600, PMTHLVL=2000,
- ✓ ICOPYT=ENABLED,
- ✓ CPYPRIOR=DISABLED,
- ✓ Reclaim disabled,
- ✓ Number of premigration drives per pool=10.

Refer to the **IBM® TS7700 Series Best Practices - Understanding, Monitoring and Tuning the TS7700 Performance** white paper for detailed description of different tuning settings.

Metrics and Workloads

Types of Throughput

The TS7760 or TS7760T_{cp0} is a disk-cache only cluster, therefore read and write data rates have been found to be fairly consistent throughout a given workload.

The TS7760T_{cp1->7} contains physical tapes to which the cache data will be periodically written and read, and therefore it exhibits four basic throughput rates: peak write, sustained write, read-hit, and recall.

Peak and Sustained Write Throughput.

For all TS7760T_{cp1->7} measurements, any previous workloads have been allowed to quiesce with respect to pre-migration to backend tape and replication to other clusters in the grid. In other words, the test is started with the grid in an idle state. Starting with this initial idle state, data from the host is first written into the TS7760T_{cp1->7} disk cache with little if any premigration activity taking place. This allows for a higher initial data rate, and is termed the “peak” data rate. Once a pre-established threshold is reached of non-premigrated compressed data, the amount of premigration is increased, which can reduce the host write data rate. This threshold is called the premigration priority threshold (PMPRIOR), and has default value of 1600 gigabytes (GB). When a second threshold of non-premigrated compressed data is reached, the incoming host activity is actively

throttled to allow for increased premigration activity. This throttling mechanism operates to achieve a balance between the amount of data coming in from the host and the amount of data being copied to physical tape. The resulting data rate for this mode of behavior is called the “sustained” data rate, and could theoretically continue on forever, given a constant supply of logical and physical scratch tapes. This second threshold is called the premigration throttling threshold (PMTHLVL), and has a default value of 2000 gigabytes (GB). These two thresholds can be used in conjunction with the peak data rate to project the duration of the peak period. Note that both the priority and throttling thresholds can be increased or decreased via a host command line request.

Read-hit and Recall Throughput

Similar to write activity, there are two types of TS7760T_{cp1->7} read performance: “read-hit” (also referred to as “peak”) and “recall” (also referred to as “read-miss”). A read hit occurs when the data requested by the host is currently in the local disk cache. A recall occurs when the data requested is no longer in the disk cache and must be first read in from physical tape. Read-hit data rates are typically higher than recall data rates.

These two read performance metrics, along with peak and sustained write performance are sometimes referred to as the “four corners” of virtual tape performance. The following charts in this paper show three of these corners:

1. peak write
2. sustained write
3. read hit

Recall performance is dependent on several factors that can vary greatly from installation to installation, such as number of physical tape drives, spread of requested logical volumes over physical volumes, location of the logical volumes on the physical volumes, length of the physical media, and the logical volume size. Because these factors are hard to control in the laboratory environment, recall is not part of lab measurement.

Grid Considerations

Grid Considerations

Up to four TS7700 clusters can be linked together to form a Grid configuration. Five- and six-way grid configurations are available via iRPQ. The connection between these clusters is provided by two 1-Gb TCP/IP links (default). Four 1-Gb links, two 10-Gb links, or four 10-Gb links options are also available. Data written to one TS7700 cluster can be optionally copied to the one or more other clusters in the grid.

Data can be copied between the clusters in either RUN (also known as “Immediate”), deferred, or sync mode copy.

When using the RUN copy mode the rewind-unload response at job end is held up until the received data is copied to all peer clusters with a RUN copy consistency point.

In deferred copy mode data is queued for copying, but the copy does not have to occur prior to job end. Deferred copy mode allows for a temporarily higher host data rate than RUN copy mode because copies to the peer cluster(s) can be delayed, which can be useful for meeting peak workload demands. Care must be taken, however, to be certain that there is sufficient recovery time for deferred copy mode so that the deferred copies can be completed prior to the next peak demand. Whether delay occurs and by how much is configurable through the Library Request command.

In sync mode copy, data synchronization is up to implicit or explicit sync point granularity across two clusters within a grid configuration. In order to provide a redundant copy of these items with a zero recovery point objectives (RPO), the sync mode copy function will duplex the host record writes to two clusters simultaneously.

TS7700 Basic Performance

The following sets of graphs show basic TS7700 bandwidths. The graphs in Figures 5 and 6 show single cluster, standalone configurations. Unless otherwise stated, the performance metric shown in these and all other data rate charts in this paper is host-view (uncompressed) MB/sec.

TS7760 Standalone Performance

TS7760
 Standalone
 Maximum Host
 Throughput




Figure 5. TS7760 Standalone Maximum Host Throughput. All runs were made with 128 concurrent jobs, each job writing and/or reading 1000 MiB (with 1:1 compression) or 2659 MiB (with 2.66:1 compression), using 32KiB blocks, QSAM BUFNO = 20, and eight 8Gb (8x8Gb) FICON channels from a zEC12 LPAR.

Notes:

Mixed workload refers to a host pattern made up of 50% jobs which read hit and 50% jobs which write. The resulting read and write activity measured in the TS7760 varied and was rarely exactly 50/50

When workload targets CP0, the TS7760T behaves as a TS7760 – see 'Introduction' section for details.

TS7760T_{cp1} Standalone Performance

TS7760T_{cp1->7}
Standalone
Maximum Host
Throughput




Figure 6. TS7760T_{cp1} Standalone Maximum Host Throughput. All runs were made with 128 concurrent jobs, each job writing and/or reading 1000 MiB (with 1:1 compression) or 2659 MiB (with 2.66:1 compression), using 32KiB blocks, QSAM BUFNO = 20, and eight 8Gb (8x8Gb) FICON channels from a zEC12 (or a z196) LPAR.

Notes:

Mixed workload refers to a host pattern made up of 50% jobs which read hit and 50% jobs which write. The resulting read and write activity measured in the TS7760 varied and was rarely exactly 50/50

For workloads that target different cache partitions:

- If some workloads target a combination of TS7760_{cp0} and TS7760_{cp1->7}, sustained write data rate will be higher than the sustained rate shown in Figure 6 which targets TS7760_{cp1} and performance will depend on the TS7760_{cp0} and TS7760_{cp1->7} combination. Please contact IBM for additional information.
- If workloads target multiple tape managed partitions simultaneously, the combined data rate for all partitions (TS7760_{cp1->7}) will be the same as the data rate for TS7760_{cp1} (as shown in figure 6).

TS7700 Grid
Maximum Host
Throughput

TS7700 Grid Performance

Figures 7, 8, 10 through 12, 14, 16, 18, 20, 22, and 24 display the performance for TS7700 grid configurations.

For these charts “D” stands for deferred copy mode, “S” stands for sync mode copy and “R” stands for RUN (immediate) copy mode. For example, in Figure 8, RR represents RUN for cluster 0, and RUN for cluster 1. SS refers to synchronous copies for both clusters.

All measurements for these graphs were made at zero or near-zero distance between clusters.

Two-way TS7700 Grid with Single Active Cluster Performance

Two-way TS7700
Grid Single Active
Maximum Host
Throughput




Figure 7. Two-way TS7760 Single Active Maximum Host Throughput. Unless otherwise stated, all runs were made with 128 concurrent jobs. Each job writing or reading 2659 MiB (1000 MiB volumes @ 2.66:1 compression) using 32 KiB block size, QSAM BUFFNO = 20, and eight 8Gb FICON channels from a zEC12 LPAR. Clusters are located at zero or near zero distance to each other in laboratory setup. DCT=125.




Figure 8. Two-way TS7760T_{cp1} Single Active Maximum Host Throughput Unless otherwise stated, all runs were made with 128 concurrent jobs. Each job writing or reading 2659 MiB (1000 MiB volumes @ 2.66:1 compression) using 32 KiB block size, QSAM BUFFNO = 20, and eight 8Gb FICON channels from a zEC12 LPAR. Clusters are located at zero or near zero distance to each other in laboratory setup. DCT=125.




Figure 9. Two-way TS7700 Hybrid Grid H1





Figure 10. Two-way TS7700 Hybrid H1 Single Active Maximum Host Throughput. Unless otherwise stated, all runs were made with 128 concurrent jobs. Each job writing or reading 2659 MiB (1000 MiB volumes @ 2.66:1 compression) using 32 KiB block size, QSAM BUFFNO = 20, and eight 8 Gb FICON channels from a zEC12 LPAR. Clusters are located at zero or near zero distance to each other in laboratory setup. DCT=125.

Two-way TS7700 Grid with Dual Active Cluster Performance



Two-way TS7700
Grid Dual Active
Maximum Host
Throughput

Figure 11. Two-way TS7760 Dual Active Maximum Host Throughput. Unless otherwise stated, all runs were made with 256 concurrent jobs (128 jobs per active cluster). Each job writing or reading 2659 MiB (1000 MiB volumes @ 2.66:1 compression) using 32 KiB block size, QSAM BUFFNO = 20, and eight 8Gb FICON channels from a zEC12 LPAR. Read tests were driven from two zEC12 LPARs. Clusters are located at zero or near zero distance to each other in laboratory setup. DCT=125.




Figure 12. Two-way TS7760T_{cp1} Dual Active Maximum Host Throughput.
 Unless otherwise stated, all runs were made with 256 concurrent jobs (128 jobs per active cluster). Each job writing or reading 2659 MiB (1000 MiB volumes @ 2.66:1 compression) using 32 KiB block size, QSAM BUFFNO = 20, and eight 8Gb FICON channels from a zEC12 LPAR. Read tests were driven from two zEC12 LPARs. Clusters are located at zero or near zero distance to each other in laboratory setup. DCT=125.




Figure 13. Two-way TS7700 Hybrid Grid H2





Figure 14. Two-way TS7700 Hybrid H2 Dual Active Maximum Host Throughput. Unless otherwise stated, all runs were made with 256 concurrent jobs (128 jobs per active cluster). Each job writing or reading 2659 MiB (1000 MiB volumes @ 2.66:1 compression) using 32 KiB block size, QSAM BUFFNO = 20, and eight 8Gb FICON channels from a zEC12 LPAR. Read tests were driven from two zEC12 LPARs. Clusters are located at zero or near zero distance to each other in laboratory setup. DCT=125.

Three-way TS7700 Grid with Dual Active Cluster Performance



Three-way
 TS7700 Grid Dual
 Active Maximum
 Host Throughput

Figure 15. Three-way TS7700 Hybrid Grid H3




Figure 16. Three-way TS7700 Hybrid H3 Dual Active Maximum Host Throughput. Unless otherwise stated, all runs were made with 256 concurrent jobs (128 jobs per active cluster). Each job writing or reading 2659 MiB (1000 MiB volumes @ 2.66:1 compression) using 32 KiB block size, QSAM BUFFNO = 20, and eight 8Gb FICON channels from a zEC12 LPAR. Read tests were driven from two zEC12 LPARs. Clusters are located at zero or near zero distance to each other in laboratory setup. DCT=125.




Figure 17. Three-way TS7700 Hybrid Grid H4




Figure 18. Three-way TS7700 Hybrid H4 Dual Active Maximum Host Throughput. Unless otherwise stated, all runs were made with 256 concurrent jobs (128 jobs per active cluster). Each job writing or reading 2659 MiB (1000 MiB volumes @ 2.66:1 compression) using 32 KiB block size, QSAM BUFFNO = 20, and eight 8Gb FICON channels from a zEC12 LPAR. Read tests were driven from two zEC12 LPARs. Clusters are located at zero or near zero distance to each other in laboratory setup. DCT=125.

Four-way TS7700 Grid with Dual Active Cluster Performance




Figure 19. Four-way TS7700 Hybrid Grid H5

Four-way TS7700
 Hybrid Grid Dual
 Active Maximum
 Host Throughput




Figure 20. Four-way TS7700 Hybrid H5 Dual Active Maximum Host Throughput. Unless otherwise stated, all runs were made with 256 concurrent jobs (128 jobs per active cluster). Each job writing or reading 2659 MiB (1000 MiB volumes @ 2.66:1 compression) using 32 KiB block size, QSAM BUFFNO = 20, and eight 8Gb FICON channels from a zEC12 LPAR. Read tests were driven from two zEC12 LPARs. Clusters are located at zero or near zero distance to each other in laboratory setup. DCT=125.




Figure 21. Four-way TS7700 Hybrid Grid H6





Figure 22. Four-way TS7700 Hybrid H6 Dual Active Maximum Host Throughput. Unless otherwise stated, all runs were made with 256 concurrent jobs (128 jobs per active cluster). Each job writing or reading 2659 MiB (1000 MiB volumes @ 2.66:1 compression) using 32 KiB block size, QSAM BUFFNO = 20, and eight 8Gb FICON channels from a zEC12 LPAR. Read tests were driven from two zEC12 LPARs. Clusters are located at zero or near zero distance to each other in laboratory setup. DCT=125.

Four-way TS7700 Hybrid Grid H7



Four-way TS7700
Hybrid H7 Quadruple
Active Maximum
Host Throughput

Figure 23. Four-way TS7700 Hybrid Grid H7




Figure 24. Four-way Hybrid H7 Four Active Maximum Host Throughput. Unless otherwise stated, all runs were made with 512 concurrent jobs (128 jobs per active cluster). Each job writing or reading 2659 MiB (1000 MiB volumes @ 2.66:1 compression) using 32 KiB block size, QSAM BUFFNO = 20, and eight 8Gb FICON channels from a zEC12 LPAR. Read tests were driven from two zEC12 LPARs. Clusters are located at zero or near zero distance to each other in laboratory setup. DCT=125.

Additional Performance Metrics

TS7760 Performance vs. FICON Channel Configuration

The figure 25 shows that the number of FICON channels affects host throughput.

Host Data Rate vs.
FICON Channel
Configuration




Figure 25. TS7760 Standalone Maximum Host Throughput. All runs were made with 128 concurrent jobs, each job writing and/or reading 1000 MiB (with 1:1 compression) or 2659 MiB (with 2.66:1 compression), using 32KiB blocks, QSAM BUFNO = 20, and eight 8Gb (8x8Gb) FICON channels from a zEC12 LPAR.

TS7760T Sustained and Premigration Rates vs. Premigration Drives

TS7760T_{cp1->7} premigration rate, i.e. the rate at which cache-resident data is copied to physical tape, depend on the number of TS1150 tape drives reserved for premigration and the number of disk drawers installed. By default, the number of tape drives reserved for premigration is ten per pool.

TS7760T_{cp1->7} sustained write rate is the rate at which host write rate balanced with premigration to tape, also depends on the number of premigration tape drives.

Figure 26 shows how the number of premigration tape drives and the number of disk cache drawers affects premigration rate and sustained write rate.

Sustained and Premigration Data Rates vs. Premigration Tape Drives




Figure 26. Standalone TS7760T_{cp1} sustained write and Premigration rate vs. the number of TS1150 drives reserved for premigration. All runs were made with 128 concurrent jobs. Each job writing or reading 2659 MiB (1000 MiB volumes @ 2.66:1 compression) using 32 KiB block size, QSAM BUFFNO = 20, and eight 8 Gb FICON channels from a zEC12 LPAR.

Notes: VEC-T/CSA/10 drawer*: R 4.0 Pre-GA performance data.

TS7760T Premigration Rates vs. Drawer Counts

Figure 27 shows that the number of cache drawers affects premigration rate (with and without host activity). The TS7760T had 8 TS1150 installed. Seven TS1150 were used for premigration.

Premigration Data
Rate vs. Cache
Drawer Counts




Figure 27. Standalone TS7760T sustained write and Premigration rate vs. the number of cache drawers. All runs were made with 128 concurrent jobs. Each job writing 2659 MiB (1000 MiB volumes @ 2.66:1 compression) using 32 KiB block size, QSAM BUFFNO = 20, and eight 8 Gb FICON channels from a zEC12 LPAR

Performance vs. Block size and Number of Concurrent Jobs.

Figure 28 shows data rates on a standalone TS7760T_{cp0} VEC-T/CSA/26 drawers/8x8Gb FICON with different workload job counts driven from a zEC12 host using different channel block sizes. Significant performance improvement occurs using larger block size.

Host Date Rate vs.
Block Size and Job
Count




Figure 28. TS7760T_{cp0} Standalone Maximum Host Throughput. Each job writing 2659 MiB (1000 MiB volumes @ 2.66:1 compression) using different block sizes, QSAM BUFFNO = 20, using eight 8Gb FICON channels from a zEC12 LPAR

Performance Tools

Batch Magic

This tool is available to IBM representatives and Business Partners to analyze SMF data for an existing configuration and workload, to help projecting a suitable TS7700 configuration.

BVIRHIST plus VEHSTATS

BVIRHIST requests historical statistics from a TS7700, and VEHSTAT produces the reports. The TS7700 keeps the last 90 days of statistics. BVIRHIST allows users to save statistics for periods longer than 90 days.

Performance Aids

Performance Analysis Tools

A set of performance analysis tools is available on Techdocs that utilizes the data generated by VEHSTAT. Provided are spreadsheets, data collection requirements, and a 90 day trending evaluation guide to assist in the evaluation of the TS7700 performance. Spreadsheets for a 90 day, one week, and a 24 hour evaluation are provided.

<http://www-03.ibm.com/support/techdocs/atsmastr.nsf/WebIndex/PRS4717>

Also, on the Techdocs site is a webinar replay that teaches you how to use the performance analysis tools.

<http://www-03.ibm.com/support/techdocs/atsmastr.nsf/WebIndex/PRS4872>

BVIRPIT plus VEPSTATS

BVIRPIT requests point-in-time statistics from a TS7700, and VEPSTATS produces the reports. Point-in-time statistics cover the last 15 seconds of activity and give a snapshot of the current status of drives and volumes.

The above tools are available at one of the following web sites:

<ftp://public.dhe.ibm.com/storage/tapetool/>

Conclusions

Conclusions

The TS7700 provides significant performance and increased capacity. Release 4.0 introduces the TS7760T, which supports the ability to connect a TS4500 or a TS3500 tape library. TS7760T can perform as a TS7760 or TS7740 depending on the target partition specified in the customer's workload. Up to 8 partitions can be defined in a TS7760T, namely CP0 through CP7. When workload targets CP0, the TS7760T behaves as a TS7760. When workload targets CPn (n = 1 through 7), the TS7760T behaves as a TS7740. The TS7700 architecture provide a base for product growth in both performance and functionality.

Acknowledgements

Acknowledgements

The authors would like to thank Joseph Swingler, Randy Hensley, Louis Russo, Toy Phouybanhdyt, Toni Alexander, Katja Denefleh, Donald Denning, and Lawrence Fuss for their review comments and insight, and also to Eileen Maroney and Lawrence Fuss for distributing the paper.

The authors would like to thank Albert Veerland for Performance Driver Enhancement.

Finally, the authors would like to thank Louis Russo, Dennis Martinez, George Venech, Michael Frick, James Maudlin, Andrew Ruhl, James F. Tucker, and Kymberly Beeston for hardware/zOS/network support.

© International Business Machines Corporation 2016

IBM Systems

9000 South Rita Road

Tucson, AZ 85744

Printed in the United States of America

9-16

All Rights Reserved

IBM, the IBM logo, System Storage, System z, zOS, TotalStorage, , DFSMSdss, DFSMSHsm, ESCON and FICON are trademarks or registered trademarks of International Business Machines Corporation in the United States, other countries, or both.

Other company, product and service names may be trademarks or service marks of others.

Product data has been reviewed for accuracy as of the date of initial publication. Product data is subject to change without notice.

This information could 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.

Performance data for IBM and non-IBM products and services contained in this document were derived under specific operating and environmental conditions. The actual results obtained by any party implementing such products or services will depend on a large number of factors specific to such party's operating environment and may vary significantly. IBM makes no representation that these results can be expected or obtained in any implementation of any such products or services.

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 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.