



IBM® TS7770, TS7770T, and TS7770C
Release 5.1
Performance White Paper
Version 2.0

By Khanh Ly
Virtual Tape Performance
IBM Tucson

Table of Contents

Table of Contents	2
Introduction.....	3
Hardware Configuration.....	4
TS7700 Performance Overview	5
TS7770 Basic Performance.....	8
TS7770 Grid Performance	11
Additional Performance Metrics	22
Performance Tools	33
Conclusions	34
Acknowledgements.....	35

Introduction

This paper provides performance information for the IBM TS7770, TS7770T, and TS7770C, which are three 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 September 26, 2016 and reflects changes for release 5.1, which features the TS7770, TS7770T, and TS7770C. Some performance related features introduced since the last performance white paper (TS7700 R 4.0 Performance White Paper) will also be included.

Unless specified otherwise, in this white paper, all runs to a TS7770T target a 300TB CP1 tape managed partition with ten FC 5274. All runs to a TS7770C target a 300TB CP1 cloud managed partition with 4 FC5274.

Product
Release
Enhancements

The following are performance related changes since R 4.0:

- ❖ 16Gb FICON support – R 4.1
- ❖ Software compression (ZSTD and LZ4) – R 4.1.2
- ❖ TS7700 cloud support - R 4.2
- ❖ Server Refresh to new Power9 pSeries - R 5.0
- ❖ Disk Cache Refresh to V5300 – R 5.0
- ❖ TS7700C Grid-awareness support – R 5.1

TS7770 Release 5.1 Performance White Paper Version 2 adds data for the following configurations:

- ❖ Standalone TS7770T VED-T/CSB/n drawers (n=6, 8)
- ❖ TS7770T Premigration Rates vs Drawer Counts
- ❖ TS7770C Cloud Premigration Rate vs. Drawer Counts
- ❖ TS7770 Copy Performance Comparison
- ❖ Performance vs Block Size and Job counts
- ❖ Performance vs Compression Schemes and Job Counts
- ❖ Performance vs Compression Scheme and Block Sizes
- ❖ Virtual Mount Performance vs Configuration and Copy Modes

Hardware Configuration

The following hardware was used in performance measurements. Performance workloads are driven from IBM System zEC13 host with eight 16 Gb FICON channels.

Standalone Hardware Setup

TS7700	Drawer count	Tape Lib/ Tape Drives	Cloud	IBM System z™ Host
TS7770 VED 3956 CSB/XSB	2, 4, 6, 8, or 10	N/A	N/A	zEC13
TS7770T VED-T 3956 CSB/XSB	2, 4, or 10	TS4500/ 12 TS1150	N/A	
TS7770C VED-C 3956 CSB/XSB	2, 4, 6, 8, or 10	N/A	IBM COS 3403	

Grid Hardware Setup

TS7700	Drawer count	Tape Lib/ Tape Drives	Cloud	Grid links (Gb)	IBM System z™ Host
TS7770 VED 3956 CSB/XSB	10	N/A	N/A	2x10	zEC13
TS7770T VED-T 3956 CSB/XSB	10	TS4500/ 12 TS1150	N/A	2x10	
TS7770T VED-C 3956 CSB/XSB	10	N/A	IBM COS 3403	4x10	

IBM COS 3403

Maneger	Accesser	Slicestor	Gridlinks/Access	Gridlinks/Slicestore
M01	2 A00	9 S01	2x10Gb	2x10Gb

The following conventions are used in this paper:

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

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 10.7 GiB of uncompressed data using 32 KiB blocks and QSAM BUFNO=20 that compresses within the TS7770 at 5.351 using ZSTD compression. Measurements were made with eight 16-gigabit (Gb) FICON channels on a zEC13 host. All runs begin with the virtual tape subsystem inactive.

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

- ✓ DCOPYT=125,
- ✓ DCTAVGTD=100,
- ✓ ICOPYT=ENABLED,
- ✓ LINKSPEED=1000
- ✓ CPYPRIOR=DISABLED,

For TS7770T, there are additional settings:

- ✓ PMPRIOR=3600, PMTHLVL=4000,
- ✓ Reclaim disabled,
- ✓ Number of premigration drives per pool=10,

For TS7770C, there are additional settings:

- ✓ CPMCNTL=0
- ✓ CPMCNTN=60
- ✓ CDELDNT=16
- ✓ CLDPRIOR=3600 with only 4 FC5274 installed (notes: I only installed 4 FC5274 with CLDPRIOR=3600 so that I would have reasonable peak and sustained periods in a 6-hour run).

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

Types of Throughput

The TS7770 or TS7770T_{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 TS7770T_{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.

The TS7770C_{cp1->7} connects to the cloud to which the cache data will be periodically written and read, and therefore it also exhibits four basic throughput rates: peak write, sustained write, read-hit, and recall.

Metrics and
Workloads

Peak and Sustained Write Throughput.

For all TS7770T_{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 TS7770T_{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. For all the run in this white paper, PMPRIOR and PMTHLVL were set to 3600 and 4000 respective to achieve a longer peak duration.

For all TS7770C_{cp1->7} measurements, CLDPRIOR was set to 3600 to establish the cloud premigration priority threshold. The premigration throttle threshold was determined by the number the FC5274 installed. I only installed 4 FC5274 so that there were reasonable peak and sustained periods in a 6-hour run).

Read-hit and Recall Throughput

Similar to write activity, there are two types of TS7770T_{cp1->7} (or TS7770C_{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 (or from the cloud). Read-hit data rates are typically higher than recall data rates.

TS7770T_{cp1->7} 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.

TS7770C_{cp1->7} recall was not included in this white paper.

Grid Considerations

Grid Considerations

Up to five TS7700 clusters can be linked together to form a grid configuration. Six-, seven-, and eight-way grid configurations are available via iRPQ. The connection between these clusters is provided by two 1-Gb, four 1-Gb links, two 10-Gb links, or four 10-Gb TCP/IP links. 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 deferred, RUN (also known as “Immediate”), 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 if DCT is set to zero (default). 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.

TS7770 Basic Performance

The following sets of graphs show basic TS7770 bandwidths. The graphs in Figures 1, 2, and 3 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.

TS7770 Standalone Performance

TS7770
Standalone
Maximum Host
Throughput




Figure 1. TS7770 Standalone Maximum Host Throughput. All runs were made with 128 concurrent jobs, each job writing and/or reading 2000 MiB (with 1:1 compression) or 10.7 GiB (with 5.35:1 compression), using 32KiB blocks, QSAM BUFNO = 20, using eight 16Gb (8x16Gb) FICON channels from a zEC13 LPAR.

Notes:

Mixed 1:1 workload refers to a host pattern made up of 50% jobs which read and 50% jobs which write. The resulting read and write activity measured in the TS7770 varied and was rarely exactly 50/50

Mixed 5:35:1 workload refers to host pattern made up of 25% read jobs and 75% write jobs.

TS7770T_{cp1} Standalone Performance

TS7770T_{cp1->7}
 Standalone
 Maximum Host
 Throughput




Figure 2. TS7770T_{cp1} Standalone Maximum Host Throughput. All runs were made with 128 concurrent jobs, each job writing and/or reading 2000 MiB (with 1:1 compression) or 10.7 GiB (with 5.35:1 compression), using 32KiB blocks, QSAM BUFNO = 20, using eight 16Gb (8x16Gb) FICON channels from a zEC13 LPAR.

Notes:

Mixed 1:1 workload refers to a host pattern made up of 50% jobs which read and 50% jobs which write. The resulting read and write activity measured in the TS7770T varied and was rarely exactly 50/50

Mixed 5:35:1 workload refers to host pattern made up of 25% read jobs and 75% write jobs.

TS7770C_{cp1} Standalone Performance

TS7770C_{cp1->7}
 Standalone
 Maximum Host
 Throughput

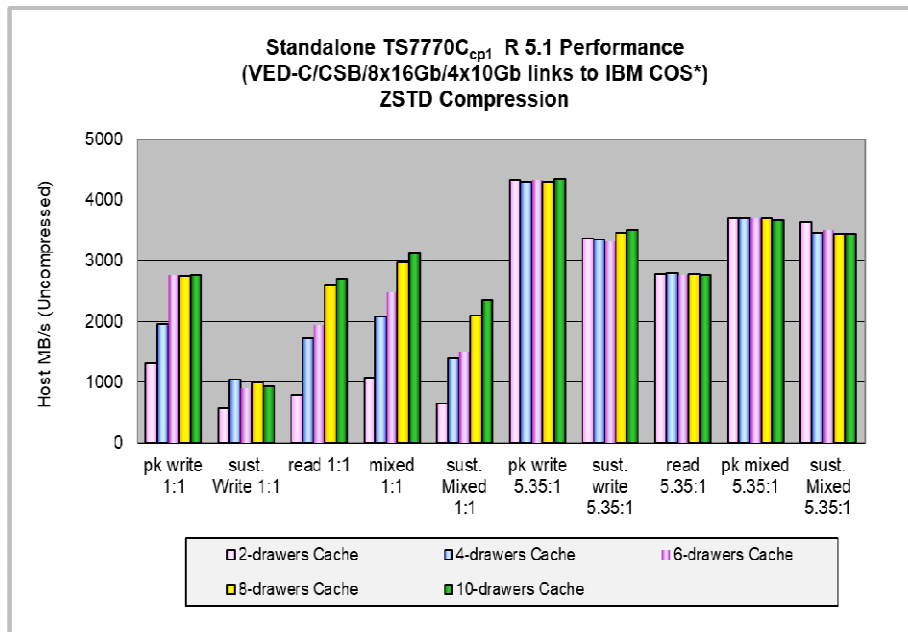


Figure 3. TS7770C_{cp1} Standalone Maximum Host Throughput. All runs were made with 128 concurrent jobs, each job writing and/or reading 2000 MiB (with 1:1 compression) or 10.7 GiB (with 5.35:1 compression), using 32KiB blocks, QSAM BUFNO = 20, using eight 16Gb (8x16Gb) FICON channels from a zEC13 LPAR.

Notes:

Mixed 1:1 workload refers to a host pattern made up of 50% jobs which read and 50% jobs which write. The resulting read and write activity measured in the TS7770C varied and was rarely exactly 50/50

Mixed 5:35:1 workload refers to host pattern made up of 25% read jobs and 75% write jobs.

TS7700 Grid
Maximum Host
Throughput

TS7700 Grid Performance

Figures 4 through 6, 8 through 11, 13, 15, 17, and 19 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 4, 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 4. Two-way TS7700 Single Active Maximum Host Throughput. Unless otherwise stated, all runs were made with 128 concurrent jobs. Each job writing 10.7 GiB (2000 MiB volumes @ 5.35:1 compression) using 32 KiB block size, QSAM BUFFNO = 20, using eight 16Gb FICON channels from a zEC13 LPAR. Clusters are located at zero or near zero distance to each other in laboratory setup. DCT=125.

Notes:

SDT/AES-256 (Security Data Transfer with TLS 1.2 AES256): Encrypted user data for grid replication.

Two-way
TS7770T_{cp1} Grid
Single Active
Maximum Host
Throughput




Figure 5. Two-way TS7770T_{cp1} Single Active Maximum Host Throughput. Unless otherwise stated, all runs were made with 128 concurrent jobs. Each job writing 10.7 GiB (2000 MiB volumes @ 5.35:1 compression) using 32 KiB block size, QSAM BUFFNO = 20, using eight 16Gb FICON channels from a zEC13 LPAR. Clusters are located at zero or near zero distance to each other in laboratory setup. DCT=125.

Notes:

SDT/AES-256 (Security Data Transfer with TLS 1.2 AES256): Encrypted user data for grid replication.

Two-way
TS7770C_{cp1} Grid
Single Active
Maximum Host
Throughput




Figure 6. Two-way TS7770C_{cp1} Single Active Maximum Host Throughput. Unless otherwise stated, all runs were made with 128 concurrent jobs. Each job writing 10.7 GiB (2000 MiB volumes @ 5.35:1 compression) using 32 KiB block size, QSAM BUFFNO = 20, using eight 16Gb FICON channels from a zEC13 LPAR. Clusters are located at zero or near zero distance to each other in laboratory setup. DCT=125.

Notes:

* HTTPS: Communication protocol between TS7770C and cloud.

* SDT/AES-256 (Security Data Transfer with TLS 1.2 AES256): Encrypted user data for grid replication.




Figure 7. Two-way TS7700 Hybrid Grid H1

Two-way TS7700
Hybrid Grid H1
Single Active
Maximum Host
Throughput




Figure 8. Two-way TS7700-TS7700T_{CP1} Hybrid H1 Single Active Maximum Host Throughput. Unless otherwise stated, all runs were made with 128 concurrent jobs. Each job writing 10.7 GiB (2000 MiB volumes @ 5.35:1 compression) using 32 KiB block size, QSAM BUFFNO = 20, using eight 16Gb FICON channels from a zEC13 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 Clusters Performance

Two-way TS7700
Grid Dual Active
Maximum Host
Throughput




Figure 9. Two-way TS7770 Dual Active Maximum Host Throughput. Unless otherwise stated, all runs were made with 256 concurrent jobs (128 jobs per active cluster). Each job writing 10.7 GiB (2000 MiB volumes @ 5.35:1 compression) using 32 KiB block size, QSAM BUFFNO = 20, using eight 16Gb FICON channels from a zEC13 LPAR. Clusters are located at zero or near zero distance to each other in laboratory setup. DCT=125.

Notes:

SDT/AES-256 (Security Data Transfer with TLS 1.2 AES256): Encrypted user data for grid replication.

Two-way
TS7700T_{CP1} Grid
Dual Active
Maximum Host
Throughput




Figure 10. Two-way TS7700T_{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 10.7 GiB (2000 MiB volumes @ 5.35:1 compression) using 32 KiB block size, QSAM BUFFNO = 20, using eight 16Gb FICON channels from a zEC13 LPAR. Clusters are located at zero or near zero distance to each other in laboratory setup. DCT=125.

Notes:

SDT/AES-256 (Security Data Transfer with TLS 1.2 AES256): Encrypted user data for grid replication.

Two-way
TS7700_{CP1} Grid
Dual Active
Maximum Host
Throughput




Figure 11. Two-way TS7700_{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 10.7 GiB (2000 MiB volumes @ 5.35:1 compression) using 32 KiB block size, QSAM BUFFNO = 20, using eight 16Gb FICON channels from a zEC13 LPAR. Clusters are located at zero or near zero distance to each other in laboratory setup. DCT=125.

Notes:

* HTTPS: Communication protocol between TS7700C and cloud.

* SDT/AES-256 (Security Data Transfer with TLS 1.2 AES256): Encrypted user data for grid replication.




Figure 12. Two-way TS7700 Hybrid Grid H2

Two-way TS7770
Hybrid Grid H2
Dual Active
Maximum Host
Throughput




Figure 13. Two-way TS7770-TS7770T_{CP1} 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 10.7 GiB (2000 MiB volumes @ 5.35:1 compression) using 32 KiB block size, QSAM BUFFNO = 20, using eight 16Gb FICON channels from a zEC13 LPAR. Clusters are located at zero or near zero distance to each other in laboratory setup. DCT=125.

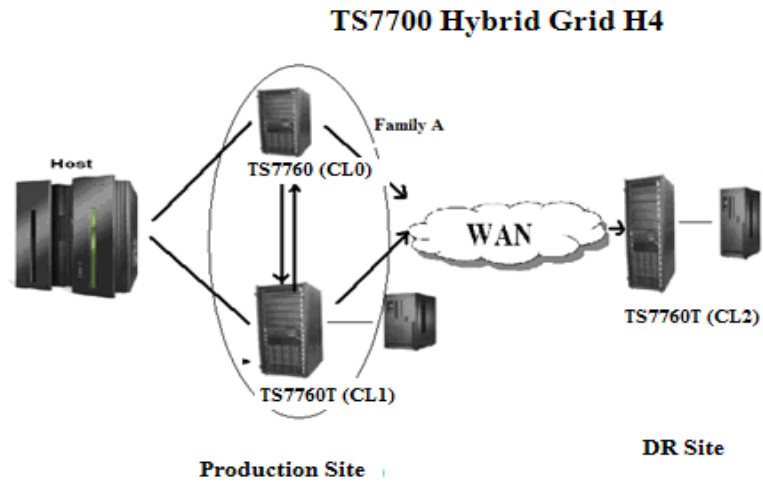


Figure 14. Three-way TS7700 Hybrid Grid H4

Three-way TS7770
 Hybrid Grid H4
 Dual Active
 Maximum Host
 Throughput




Figure 15. Three-way TS7770-TS7770T_{CP1} 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 10.7 GiB (2000 MiB volumes @ 5.35:1 compression) using 32 KiB block size, QSAM BUFFNO = 20, using eight 16Gb FICON channels from a zEC13 LPAR. Clusters are located at zero or near zero distance to each other in laboratory setup. DCT=125.

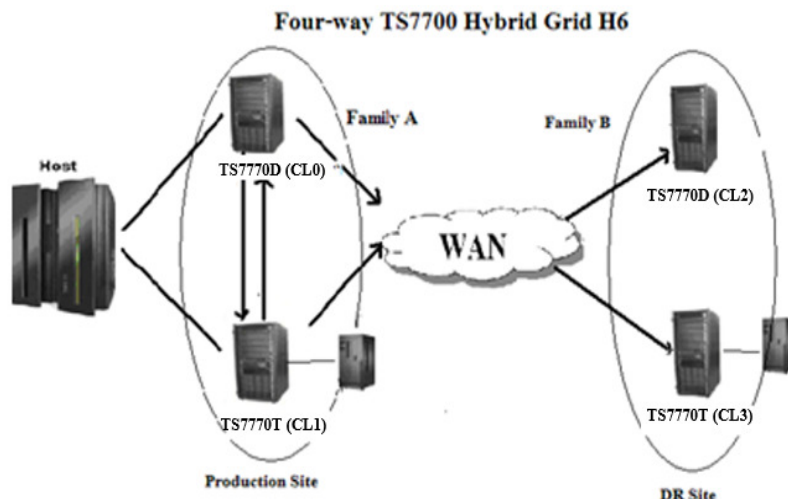


Figure 16. Four-way TS7700 Hybrid Grid H6

Four-way TS7770
Hybrid Grid H6
Dual Active
Maximum Host
Throughput

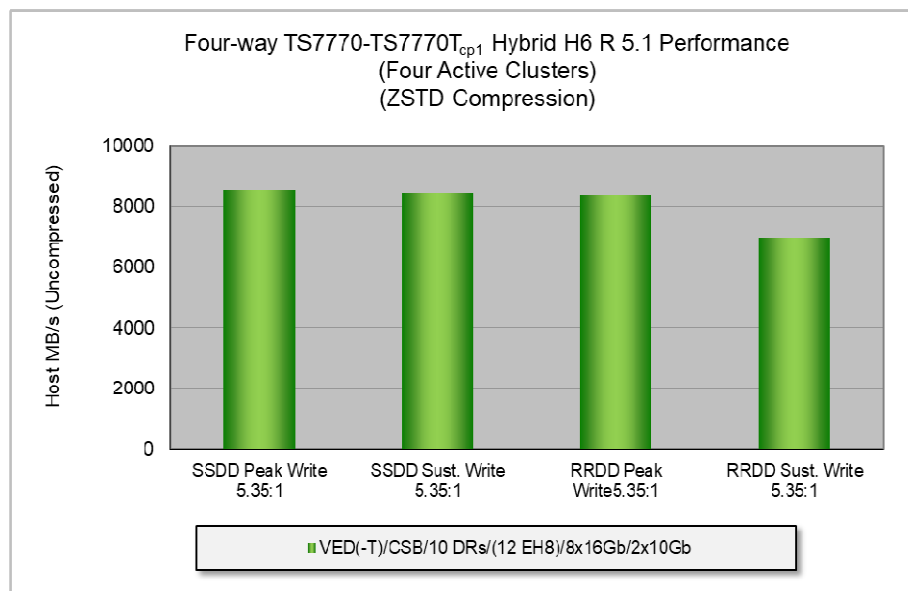


Figure 17. Four-way TS7770-TS7770T_{cp1} Hybrid H6 Dual Active Maximum Host Throughput. Unless otherwise stated, all runs were made with 256 concurrent jobs (128 jobs per active cluster). Each job writing 10.7 GiB (2000 MiB volumes @ 5.35:1 compression) using 32 KiB block size, QSAM BUFFNO = 20, using eight 16Gb FICON channels from a zEC13 LPAR. Clusters are located at zero or near zero distance to each other in laboratory setup. DCT=125.

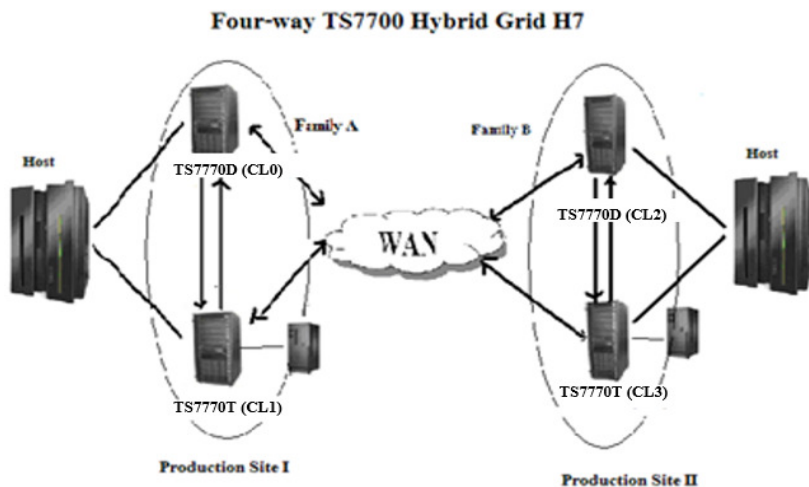


Figure 18. Four-way TS7700 Hybrid Grid H7

Four-way TS7700
Hybrid H7
Quadruple Active
Maximum Host
Throughput




Figure 19. Four-way TS7770-TS7770T_{cp1} Hybrid H7 Four Active Maximum Host Throughput. Unless otherwise stated, all runs were made with 256 concurrent jobs (128 jobs per active cluster). Each job writing 10.7 GiB (2000 MiB volumes @ 5.35:1 compression) using 32 KiB block size, QSAM BUFFNO = 20, using eight 16Gb FICON channels from a zEC13 LPAR. 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 20 shows how the number and/or configuration of the FICON channels affects host throughput.

Host Data Rate vs.
FICON Channel
Configuration




Figure 20. TS7770 Standalone Maximum Host Throughput. All runs were made with 128 concurrent jobs, each job writing 10.7 GiB (with 5.35:1 compression), using 32KiB blocks, QSAM BUFNO = 20, using eight 16Gb (8x16Gb) FICON channels from a zEC13 LPAR.

TS7770T Sustained and Premigration Rates vs. Premigration Drives

TS7770T_{cp1->7} premigration rates, i.e. the rates at which cache-resident data is copied to physical tapes, 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.

TS7770T_{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.

The figure 21 shows how the number of premigration tape drives affects premigration rate and sustained write rate.

Sustained Host rate and Premigration rates vs. Premigration Tape Drives




Figure 21. Standalone TS7770T_{cp1} sustained write rate and tape premigration rate vs. the number of TS1150 drives reserved for premigration. All runs were made with 128 concurrent jobs. each job writing 10.7 GiB (with 5.35:1 compression), using 32KiB blocks, QSAM BUFNO = 20, using eight 16Gb (8x16Gb) FICON channels from a zEC13 LPAR.

TS7770T Premigration Rates vs. Drawer Counts

The figure 22 shows that the number of cache drawers affects premigration rate (with and without host activity).

Sustained Host rate and Premigration rates vs. Cache Drawer Counts

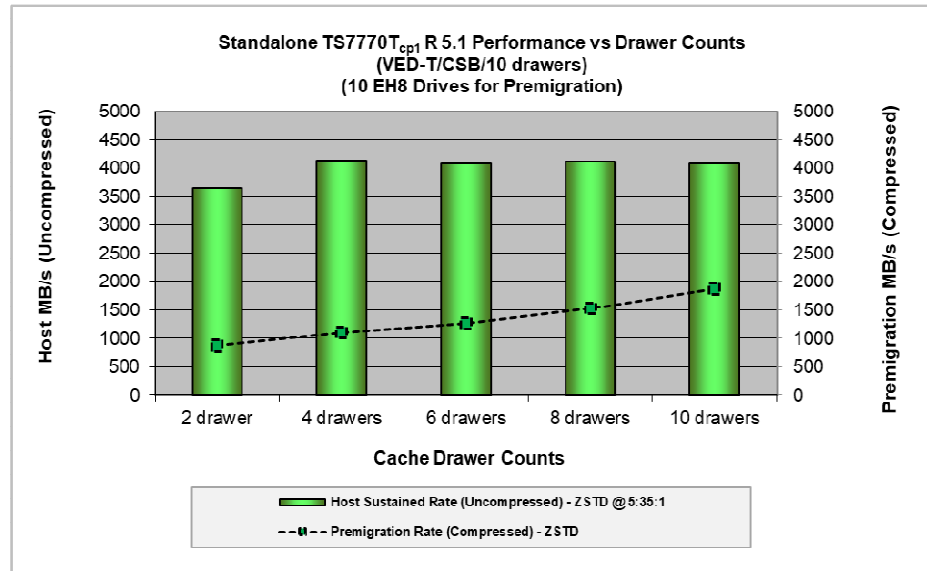


Figure 22. Standalone TS7770T_{cp1} sustained write rate and tape premigration rate vs. the number of cache drawers. All runs were made with 128 concurrent jobs. each job writing 10.7 GiB (with 5.35:1 compression), using 32KiB blocks, QSAM BUFNO = 20, using eight 16Gb (8x16Gb) FICON channels from a zEC13 LPAR.

TS7770C Cloud Premigration Rates vs. Drawer Counts

The figure 23 shows that the number of cache drawers affects cloud premigration rate (with and without host activity).

Sustained Host Rate and Cloud Premigration Rates vs. Cache Drawer Counts




Figure 23. Standalone TS7770C_{cp1} sustained write rate and cloud premigration rate vs. the number of cache drawers. All runs were made with 128 concurrent jobs. each job writing 10.7 GiB (with 5.35:1 compression), using 32KiB blocks, QSAM BUFNO = 20, using eight 16Gb (8x16Gb) FICON channels from a zEC13 LPAR.

TS7700 Copy Performance Comparison

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.

TS7700 Copy Performance Comparison




Figure 24. Two-way TS7700 Maximum Copy Throughput. Clusters are located at zero or near zero distance to each other in laboratory setup.

Notes:

SDT/AES-256 (Security Data Transfer with TLS 1.2 AES256): Encrypted user data for grid replication.

For 2-way TS7700T_{cp1}, the premigration activity on the source/target cluster consumes resources and thus lower the copy performance on the TS7700T as compared to the TS7700.

TS7700T
Sustained Copy
Performance
Evolution




Figure 25. Two-way TS7700T_{cp1} Maximum Copy Throughput. Clusters are located at zero or near zero distance to each other in laboratory setup.

Notes:

SDT/AES-256 (Security Data Transfer with TLS 1.2 AES256): Encrypted user data for grid replication.

Performance vs. Block size and Number of Concurrent Jobs.

Figure 26 shows data rates on a standalone TS7770 VED/CSB/10 drawers/8x16Gb FICON with different job counts driven from a zEC13 host using different channel block sizes. Significant performance improvement occurs using 64KB and 128KB block size with 64 concurrent jobs.

Standalone
TS7770
Performance vs
Block Sizes and
Job Counts




Figure 26. TS7770 Standalone Maximum Host Throughput. All runs were made with 128 concurrent jobs. each job writing 10.7 GiB (with 5.35:1 compression), with QSAM BUFNO = 20, using eight 16Gb (8x16Gb) FICON channels from a zEC13 LPAR.

Performance vs. Compression Schemes and Job Counts

Figures 27 and 28 show data rates on a standalone TS7770 VED/CSB/10 drawers/8x16Gb FICON with different workload job counts driven from a zEC13 host using different TS7770 compression schemes (FICON, LZ4, and ZSTD). The difference between figure 27 and figure 28 was the block size used (32KB vs 256KB).

Standalone TS7770
Performance vs
Compression
Schemes and Job
Counts
(32KB Block Size)




Figure 27. TS7770 Standalone Maximum Host Throughput. All runs were made with 128 concurrent jobs. each job writing 10.7 GiB (with 5.35:1 compression), with 32KB block size and QSAM BUFNO = 20, using eight 16Gb (8x16Gb) FICON channels from a zEC13 LPAR.

Standalone TS7770
Performance vs
Compression
Schemes and Job
Counts
(256KB Block Size)




Figure 28. TS7770 Standalone Maximum Host Throughput. All runs were made with 128 concurrent jobs. each job writing 10.7 GiB (with 5.35:1 compression), with 256KB block size and QSAM BUFNO = 20, using eight 16Gb (8x16Gb) FICON channels from a zEC13 LPAR.

Performance vs. Compression Schemes and Block Sizes

Figure 29 shows how data compression ratio varies depending on the compression scheme and block size used.

Standalone TS7770
Performance vs
Compression
Schemes and Block
Sizes

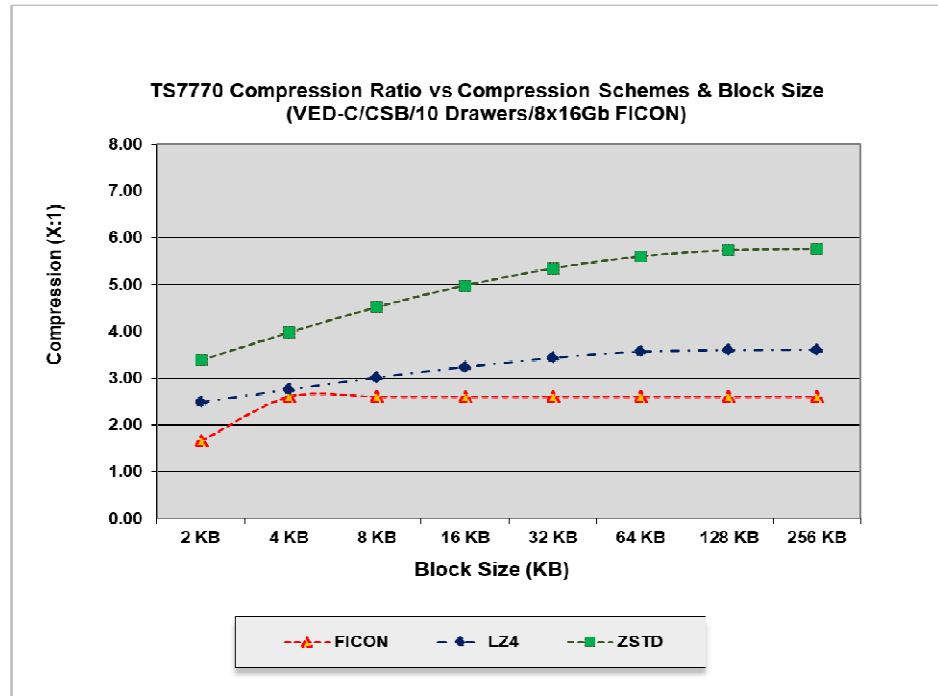


Figure 29. TS7770 Compression Ratio. All runs were made with 128 concurrent jobs. each job writing 10.7 GiB (with 5.35:1 compression), with QSAM BUFNO = 20, using eight 16Gb (8x16Gb) FICON channels from a zEC13 LPAR.

Virtual Mount Performance vs. Configurations and Copy Modes

In each of the following mount intensive runs, very small volumes were used to obtain the highest number of virtual mounts.

TS7700 Maximum Virtual Mount vs Hardware Configurations and Copy Modes




Figure 30. TS7770 Maximum Mount Rate. All runs were made with 128 concurrent jobs, each job writing 100 MiB (with 5.35:1 compression), with 32KB block size and QSAM BUFNO = 20, using eight 16Gb (8x16Gb) FICON channels from a zEC13 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, and project a suitable TS7700 configuration.

BVIRHIST plus VEHSTATS

BVIRHIST requests historical statistics from a TS7700, and VEHSTATS 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/atmastr.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/atmastr.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 improvement, increased capacity, and new functionality over the years. Release 4.1 introduced 16Gb FICON channel support which increased the maximum channel performance from 2500 MB/s to over 4000 MB/s. Release R 4.1.2 introduced software compression LZ4 and ZSTD which increase the compression ratio very significantly as compared to the traditional hardware compression at the FICON adapter level (FICON compression). Release 4.2 introduced cloud support. Release 5.0 introduces new TS7770, TS7770T, and TS7770C models with new Power 9 server and V5000 cache. The TS7700 architecture provides a base for product growth in both performance and functionality.

Acknowledgements

Acknowledgements

The author would like to thank Joseph Swingler, and Donald Denning for their review comments and insight, and also to Gary Anna for publishing the paper.

The author would like to thank Albert Veerland for Performance Driver support.

Finally, the author would like to thank Donald Denning, James F. Tucker, Kymberly Beeston, Dennis Martinez, George Venech, and Jeffrey Watson 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

12-20

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.