Technical Validation

IBM TS7700 Series Virtual Tape Solutions

Boosting the Performance and Resilience of IBM z/TPF Workloads

By Vinny Choinski, Senior Validation Analyst and Christophe Bertrand, Senior Analyst
September 2020

Contents

ESG Technical Validations

The goal of ESG Technical Validations is to educate IT professionals about information technology solutions for companies of all types and sizes. ESG Technical Validations are not meant to replace the evaluation process that should be conducted before making purchasing decisions, but rather to provide insight into these emerging technologies. Our objectives are to explore some of the more valuable features and functions of IT solutions, show how they can be used to solve real customer problems, and identify any areas needing improvement. The ESG Validation Team's expert third-party perspective is based on our own hands-on testing as well as on interviews with customers who use these products in production environments.

Introduction

This technical validation report documents ESG testing of the latest capabilities of the TS7700 and auditing of the performance of IBM TS7700 Series virtual tape solutions located in an IBM facility in Tucson, Arizona. Validation testing was designed to demonstrate z/TPF integration and management, the capacity-on-demand feature, overall system performance of the latest model, and compression performance, in particular.

Background

It has been said that the world runs on the mainframe.[1] Because so many businesses depend on mainframes to deliver their critical applications, solutions that improve processes and performance in these environments can play a critical role in helping IT organizations deliver reliable continuous access to business-critical assets. A recent ESG research survey illustrates that the considerations most important to IT decision makers for justifying IT investments to their organizations' business management teams over the next 12 months include improving data security and business risk management, improving business process, and enabling digital transformation (see Figure 1).[2]

Figure 1. Top Five Most Important Considerations for Justifying IT Investment

Which of the following considerations do you believe will be most important in justifying IT investments to your organization's business management team over the next 12 months? (Percent of respondents, N=658, five responses accepted)

Consideration	Percent
Improves security/risk management	36%
Increases employee productivity	34%
Improves business process	32%
Enables digital transformation	30%
Improves customer experience	29%

Source: Enterprise Strategy Group

IBM TS7700 virtual tape solutions help organizations address multiple concerns most important to IT decision makers. They help improve business processes by increasing system performance and offering flexible options for addressing exploding data volumes; they improve data security and business risk management through their grid architecture and support of Z pervasive encryption, among other attributes; and through technologies such as Transparent Cloud Tiering (TCT), which helps enable the multi-cloud architectures that are driving many digital transformation initiatives.

---

[1] Source: Forbes, *Running These Workloads? You Should Take A Look at The IBM Z15*, March 2020.
[2] Source: ESG Master Survey Results, *2020 Technology Spending Intentions Survey*, January 2020.

IBM TS7700 Series

IBM TS7700 is a family of mainframe virtual tape library (VTL) solutions that optimize data protection and business continuance in IBM Z mainframe environments. Tape virtualization can help improve reliability and multi-regional resilience; reduce the time needed for critical batch windows; decrease services downtime caused by physical tape drive and library outages; and reduce data processing cost, time, and complexity by moving primary workloads to virtual tape.

Through the use of virtualization and disk cache, the TS7700 family operates at disk speeds while maintaining compatibility with existing tape operations. Its fully integrated, tiered storage hierarchy takes advantage of both disk and tape technologies to deliver performance for active data and the best economics for inactive and archive data. A web-based graphical user interface (GUI), based on the interface used in several other IBM storage solutions, is provided to configure and monitor TS7700 systems. TS7700 writes data by policy to physical tape through attachment to high-capacity, high-performance IBM TS1150 and IBM TS1140 tape drives that are installed in IBM TS4500 or TS3500 tape libraries.

IBM TS7700 is now in its seventh generation, providing over 22 years of IBM Z virtual tape support. The IBM TS7770 for IBM Z is the latest model of the TS7700 family. R5.0 is the microcode level associated with the new IBM TS7770. R5.0 can be installed on the IBM TS7770 and IBM TS7760 models only. The IBM TS7770 mimics the previous generation IBM TS7760 with more options and functions, including object storage support.

The new TS7770 VTL solutions are designed to improve storage economics and data security in mission-critical hybrid multi-cloud environments. Increased disk density helps reduce IT infrastructure costs, new POWER9 controllers substantially reduce processor energy consumption, and Transparent Cloud Tiering technology enables up to ten times faster cloud object storage. Figure 2 shows IBM Z processors connected to a TS7770 VTL system over a FICON SAN. The TS7770 solution also includes an IBM TS4500 tape library for policy-driven, cost-efficient storage of inactive data. The TS7770 and its various components, known as a cluster, are connected by WAN to other TS7700 family clusters in a geographically dispersed grid. The new TS7770 can now support up to eight clusters in a single grid.

Figure 2 illustrates a typical TS7770 implementation and highlights many of the new features:

- Enhancements to the TS7700 **Capacity-on-Demand** feature enable users who deploy the new TS7770 model to pay only for the storage actually used. The Capacity-on-Demand feature allows users to enable disk capacity concurrently from just 20 TB up to 780 TB on the base frame, 2.36 PB with one expansion frame, and 3.94 PB with two expansion frames.

- TS7700 systems are IBM Z intelligent—no additional software is required to support them in z/OS, z/TPF, and also z/VM and z/VSE environments, while IBM Z enjoys full access to all IBM proprietary tape library command sets. IBM Z platforms see the entire TS7700 grid, instead of a series of independent tape libraries. Advanced policy management enables cache management for volume retention and deletion and is tightly integrated with DFSMS policy management. Volume pooling allows the grouping of logical volumes on physically separate cartridges or cartridge pools, and cross-site replication creates copies of logical volumes at different sites.

- As part of the overall support for IBM Z pervasive encryption, new TS7770 systems provide **full AES256 encryption** for data in flight and at rest. The TS7700 also now provides full AES256-strength encryption for all logical volume content in flight between peers in a grid network and TLS1.2 RSA level encryption for all external SKLM communications. The inherent air gap isn't the only data protection feature; TS7770 solutions also offer enhanced key management and write once/read many times (WORM) technologies.

- The IBM POWER9 servers used in TS7770 systems plus enhancements in microcode enable **twice the bandwidth** between hosts and TS7700 storage to increase application performance. The POWER9 servers also offer more powerful CPUs, more RAM, and support for higher capacity SSDs.

- Using **TCT technology**, TS7700 grids can now offload data volumes to object storage targets in the cloud. A new Automated Grid Cloud Failover capability helps ensure data is available from anywhere at any time, providing nearly zero seconds failover across up to 8 grid-linked TS7770 systems, 20 GB/sec of throughput, and nearly 12 PB compressed capacity, allowing mainframes to send more data faster while reducing the CPU utilization associated with data management.

Figure 2. IBM TS7770 for Enterprise-class IBM z/TPF Workloads
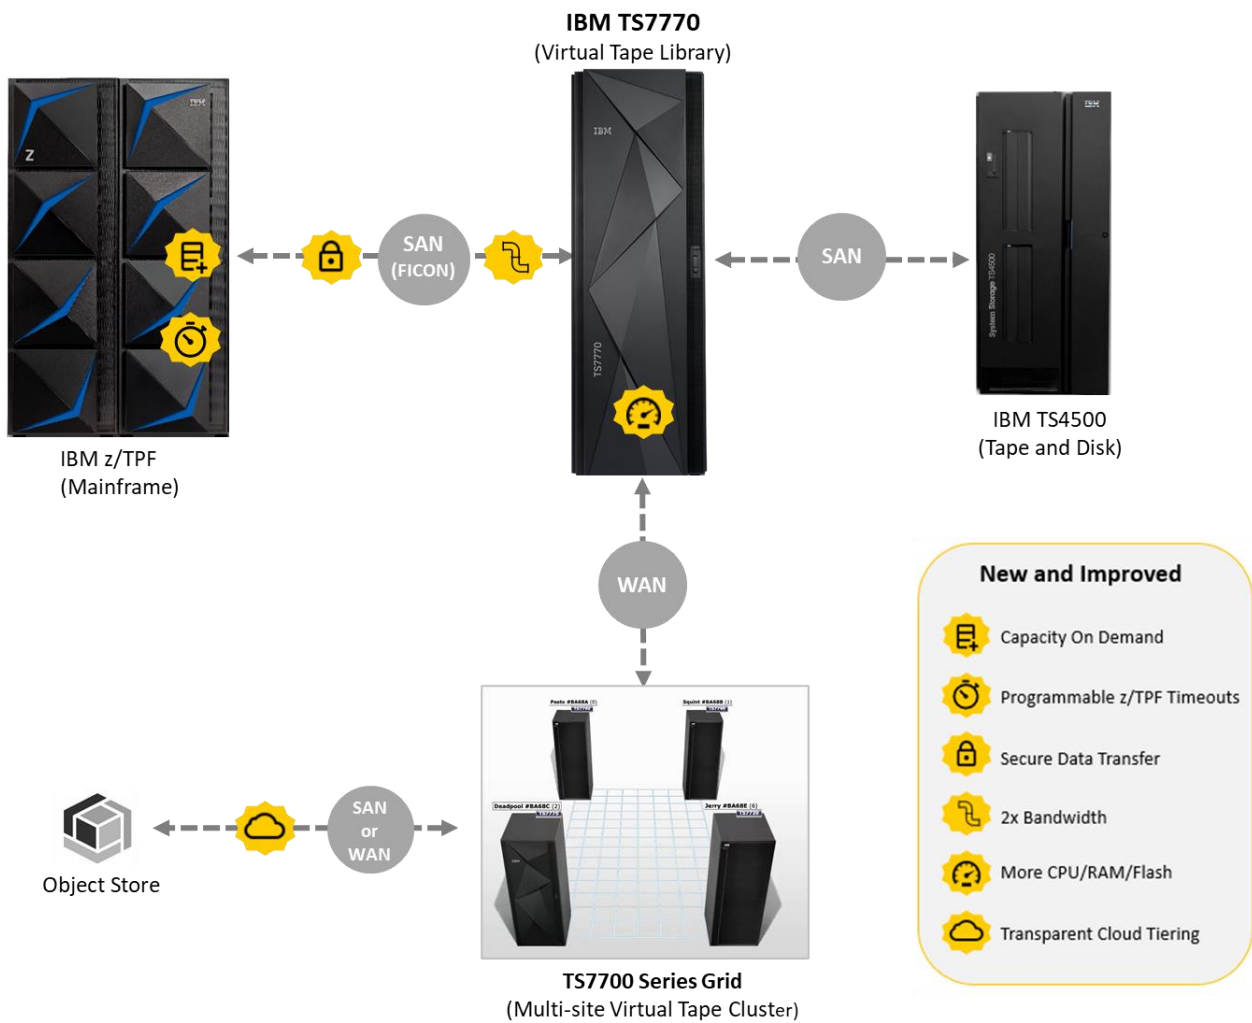
Source: Enterprise Strategy Group

ESG Technical Validation

ESG conducted validation testing of TS7700 series of virtual tape solutions that were located in an IBM facility in Tucson, Arizona. Testing was designed to demonstrate z/TPF integration and management, Capacity on Demand, overall system performance, and data compression performance.

IBM z/TPF Integration and Management

IBM z/TPF is a high-volume, high-throughput transaction processor that can handle large, continuous loads of essentially simple transactions across large, geographically dispersed networks. But it is more than just a transaction processor. z/TPF is also an operating system and a unique database, all designed to work together as one system.

Figure 3 illustrates a common IBM TS7700 grid deployment supporting TPF data processing. z/TPF makes data write requests to a target TS7700 across a FICON SAN. The target system is WAN-connected to four other TS7700 clusters that together comprise a geographically distributed TS7700 grid. The grid, functioning as a single storage resource, is connected to an IBM TS4500 tape library where inactive data is moved for cost-efficient storage. TPF is often deployed in always-on 24/365 mission-critical environments.

**Figure 3. IBM z/TPF Integration and Management**

Source: Enterprise Strategy Group

Figure 4 shows the results of ESG z/TPF – TS7700 system testing. The test objective was to exercise TPF data transfer from a z/TPF host through the TS7700 test system to a back-end TS4500 physical tape library in order to investigate the buildup of write requests or write queue depth during simulated workloads. The images in Figure 4 show the GUI where a planned cluster failover has been configured at a queue depth threshold of 50,000. The simulation involved up to three tape

devices running at about 100 MB/sec each. The zdtap cmd command was used to monitor tape queue depths during the ESG simulation.

During the entire test sequence, ESG never recorded an actual queue depth above 1,000. Figure 4 is a screen shot showing a queue depth of 134 with 3 tape devices running a 300 MB/sec write workload. Queue depth during the ESG TPF demo usually remained under 1,000 and never got anywhere near the tape switch threshold of 50,000.

Figure 4. System Monitoring

300+ MB/sec
(3 **TS7760T** TPF volumes)

Low z/TPF tape queue
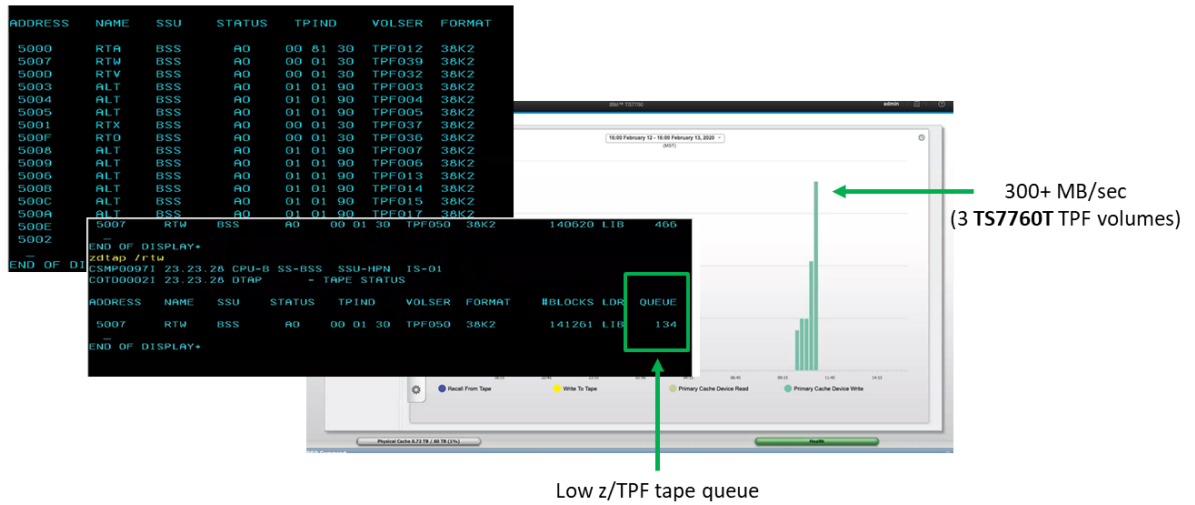
ⓘ Why This Matters

z/TPF is a foundation of modern global business. The performance of transactional applications running over z/TPF matters to anyone who uses a credit card or makes a travel reservation, among many other commerce activities.

ESG has confirmed that IBM TS7700 VTL solutions are more than capable of handling typical transaction volumes in IBM Z environments. The performance of TS7700 means that enterprises generating very large transactional workloads can still build storage solutions for mainframes that optimize cost-efficiency while enabling leading-edge multi-cloud architectures for high availability, cyber resilience, and disaster recovery purposes.

Capacity on Demand

With the release of the latest model, IBM has enhanced the Capacity-on-Demand options available with TS7700 solutions. As noted above, the Capacity-on-Demand feature allows users to pay only for storage capacity currently in use, and then add to available capacity as needed in 20 TB or 100 TB increments.

ESG validated the enhanced TS7770 Capacity-on-Demand capabilities by simulating the addition of storage capacity through the TS7700 UI. Figure 5 illustrates the simple process involved. The image and exploded "Cluster Settings" UI dialog box in the upper right of the figure shows the "Feature Licenses" and "Current Available Resources" of TS7700 cluster #BA871. In the middle dialog box, a 20 TB storage addition is being enabled. The final image and dialog box confirm that 20 TB of storage capacity has been added to the cluster.

Figure 5. Capacity on Demand
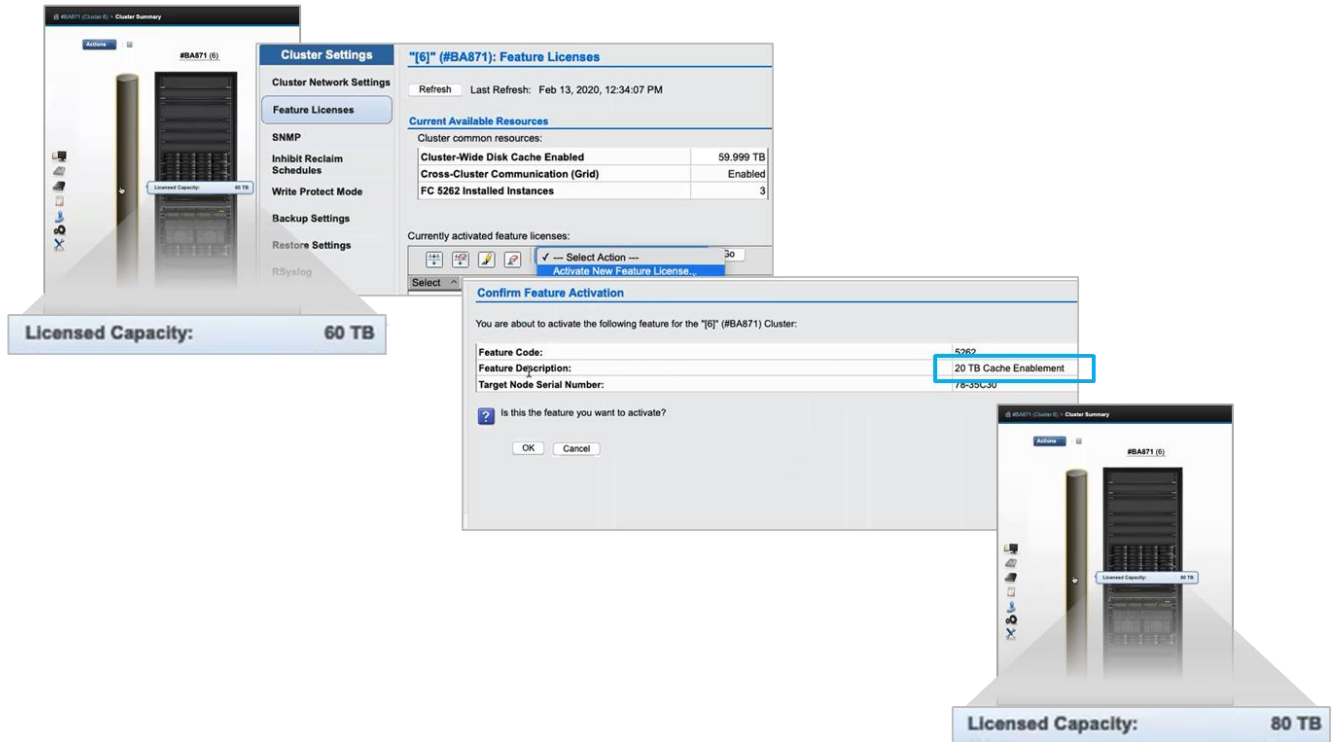
Source: Enterprise Strategy Group

ℹ Why This Matters

Over the past decade, the pay-as-you-go consumption model has gained significant traction in the IT world. The rise of cloud computing provides the most powerful example. ESG research confirms that today, pay-per-use is the preferred storage resources payment model—more than three-quarters (76%) of IT organizations are procuring at least some on-premises storage capacity in this manner at present.[3]

In the modern era of exploding data volumes and flat IT budgets, flexible consumption models matter to everyone, from the storage administrator, up to the CIO and beyond, to the business executive responsible for overall profit performance.

Performance

One of the main drivers of virtual tape solution deployments is the need for higher system performance. When mainframe applications run faster, the entire world of commerce benefits.

The new TS7770 model comes with several new features and capabilities that drive higher system performance. IBM POWER9 servers replace the previous generation POWER8 technology, leveraging two 10-core, 3.8GHz processors and 128GB of DDR4 memory to provide 12% faster processing at 20% less energy consumption. Also, TS7770 supports 16Gb FICON connectivity, enabling IBM customers to exploit their most current FICON infrastructure and maximize FICON

---

[3] Source: ESG Research Report, *Data Storage Trends in an Increasingly Hybrid Cloud World*, March 2020.

throughput to IBM Z. 20 GB/sec of throughput and nearly 12 PB of compression-enhanced storage capacity enable mainframes to send more data faster while reducing the CPU utilization associated with data management.
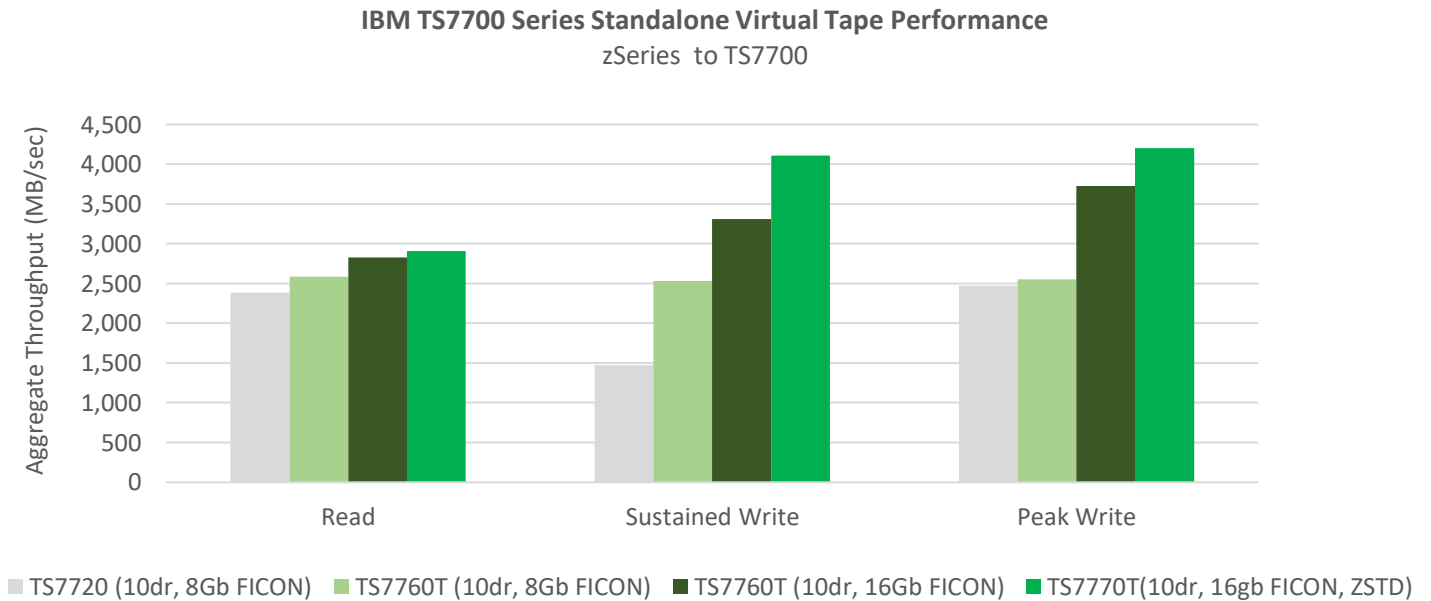
Getting Started

The TS7700 performance results shown in this paper have been derived from measurements that generally attempt to simulate common user environments, namely many jobs writing and/or reading multiple tape volumes simultaneously. Unless otherwise noted, all measurements were made with 128 simultaneously active virtual tape jobs per active cluster. Each tape job was writing or reading 2,659 MB 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 16Gb FICON channels from an IBM z13 LPAR. All runs began with the virtual tape subsystem inactive. Unless otherwise stated, all runs were made with default tuning values.

Standalone Performance

As shown in Figure 6, ESG audited throughput performance for common z/TPF host operations. "Sustained Write" measured host to FICON attached disk (e.g., virtual tape) performance with a concurrent FC-SAN attached tape copy. "Peak Write" measured host to FICON attached disk performance (e.g., no tape copy), and the "Read" results demonstrated FICON disk to host read performance.

Figure 6. Standalone Performance

IBM TS7700 Series Standalone Virtual Tape Performance
zSeries to TS7700

Legend: ■ TS7720 (10dr, 8Gb FICON)   ■ TS7760T (10dr, 8Gb FICON)   ■ TS7760T (10dr, 16Gb FICON)   ■ TS7770T(10dr, 16gb FICON, ZSTD)

Source: Enterprise Strategy Group

During the ESG performance test session, approximately 100 MB/sec throughput was measured to one tape device, and approximately 300 MB/sec was observed from three devices running in parallel.

For all TS7700 cluster measurements, any previous workloads were quiesced with respect to premigration to back-end tape and replication to other clusters in the grid. In other words, tests were started with the grid in an idle state. Then data from the host was first written into the TS7700 disk cache with little if any premigration activity taking place. This allowed for a higher initial data rate and is termed the "peak" data rate.

Once a preestablished threshold was reached, the amount of premigration was increased, which can reduce the host write data rate. This threshold is called the premigration priority threshold (PMPRIOR) and has a default value of 1,600 GB. When a second threshold was reached, the incoming host activity was actively throttled to allow for increased premigration activity. This throttling mechanism operated 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 forever, given a constant supply of logical and physical scratch tapes.

Table 1. Standalone Performance

| Model	Read (MB/s)	Sustained Write (MB/s)	Peak Write (MB/s)
TS7720 (10dr, 8Gb FICON)	2386	2588	2827
TS7760T (10dr, 8Gb FICON)	1473	2530	3309
TS7760T (10dr, 16Gb FICON)	2471	2554	3730
TS7770T (10dr, 16GB FICON, ZSTD)	2910	4112	4204

Source: Enterprise Strategy Group

Copy Performance

For most workloads, tape copy performance into and/or through a grid doesn't matter much, other than the ability of copy performance to shrink RPOs for disaster recovery. But in z/TPF environments, copy performance is crucial. TPF logs everything for transaction-level recoverability. If the copy to tape process can't keep pace, then transaction queues start building up in the host.

Up to eight TS7700 clusters can be linked together to form a grid configuration. The connection between these clusters is provided by two or four 10-Gb TCP/IP links. Data written to one TS7700 cluster can be optionally copied to other clusters in the grid. Data can be copied between the clusters in either RUN (also known as "Immediate"), deferred, sync mode copy, or no 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 delays occur 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 RPO, the sync mode copy function will duplex the host record writes to two clusters simultaneously.

Figure 7 shows the audited copy performance results from ESG testing. These tests measured throughput when creating copies between TS7700 nodes in deferred copy mode. The nodes were connected via 10 Gigabit Ethernet. Testing was first run on a TS7720 and then repeated for the TS7760.[4]

---

[4] Configuration: 26 drawer TS7720t version 3.2 and a 26 drawer TS7760t version 4.0 with 2.6:1 data compression.

Figure 7. Copy Performance

IBM TS7700 Series Two-way Single Direction Copy Performance
TS7700 Series to TS7700 Grid
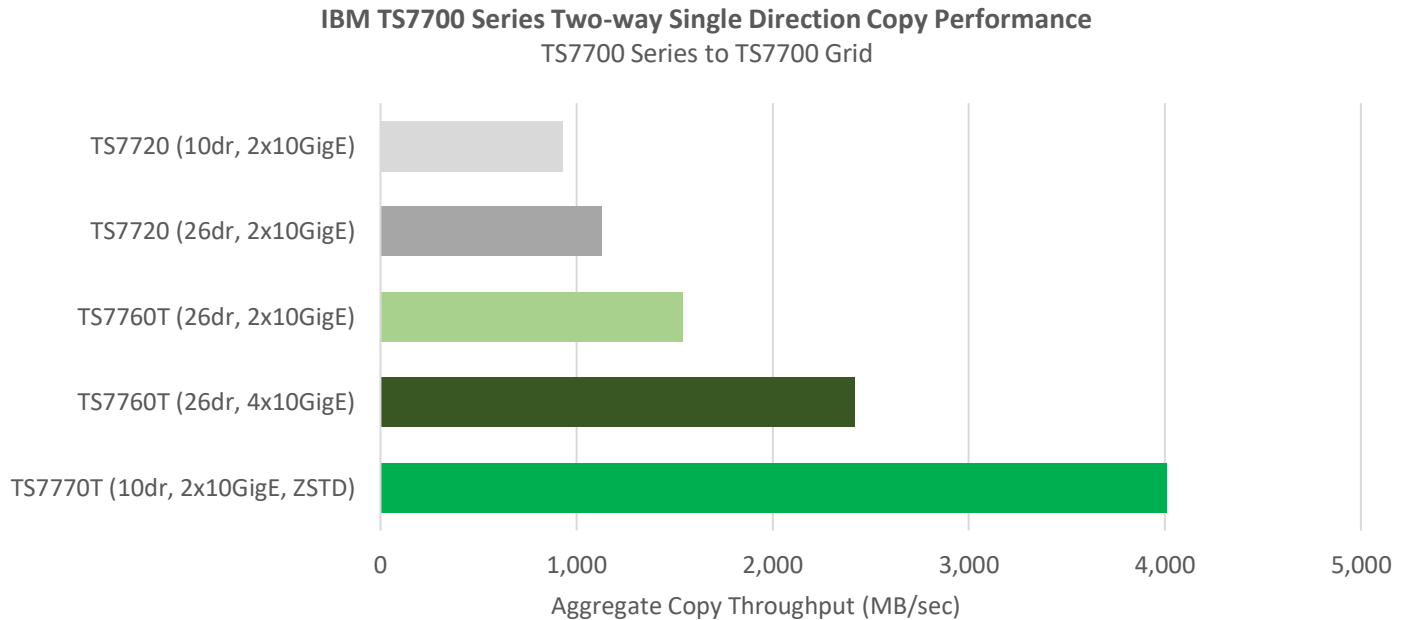
Source: Enterprise Strategy Group

Table 2. Copy Performance

	TS7720 (10dr, 2x10GigE)	TS7720 (26dr, 2x10GigE)	TS7760 (26dr, 2x10GigE)	TS7760 (26dr, 4x10GigE)	TS7770T (10dr, 2x10GigE, ZSTD)
Throughput (MB/s)	933	1129	1545	2420	4010

Source: Enterprise Strategy Group

The data rates shown in Figure 7 and Table 2 are produced by compressed data over TS7700 grid links. In each ESG test, a deferred copy mode run was ended following several TBs of data being written to the active cluster. In 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 appropriate premigration activity. The premigration activity consumed resources and thus lowered the copy performance.

ⓘ **Why This Matters**

The continuously growing amount of data that needs to be managed in modern data center environments puts pressure on IT professionals as they design and implement today's storage solutions. It's not easy building a solution with the agility to deliver high-performance and scalable capacity as storage demands grow and environments evolve.

ESG has validated that the TS7700 series from IBM can help organizations meet these data storage challenges. With cluster transfer rates exceeding 2,400 MB/s and storage capacity up to 100PB leveraging the TS4500 high-density tape library option, ESG confirmed that TS7700 solutions can be scaled to meet just about any storage need.

Compression

Data compression is an important element of any modern data management solution. It helps lower costs while increasing network performance and overall system efficiency, among other benefits. But not all compression solutions are the same; each makes its own compromise between effectiveness and throughput. TS7700 has supported a form of ALDC compression in the FICON adapters since its first release. The FICON adapter compression is an older algorithm that produces lower than average compression results. IBM Z began offering zEDC compression at the host in the z13/z14 models. Recently, with release R4.1, IBM has introduced two new compression options within the TS7700 itself (the LZ4 and ZSTD algorithms) that offer choice in balancing performance demands with storage requirements. And the FICON adapter compression is still supported. Figure 8 illustrates the current data compression options within the IBM Z / TS7700 ecosystem.

Figure 8. IBM Z / TS7700 Compression Options

IBM z/TPF
(Mainframe)

IBM TS7760 and IBM TS7770
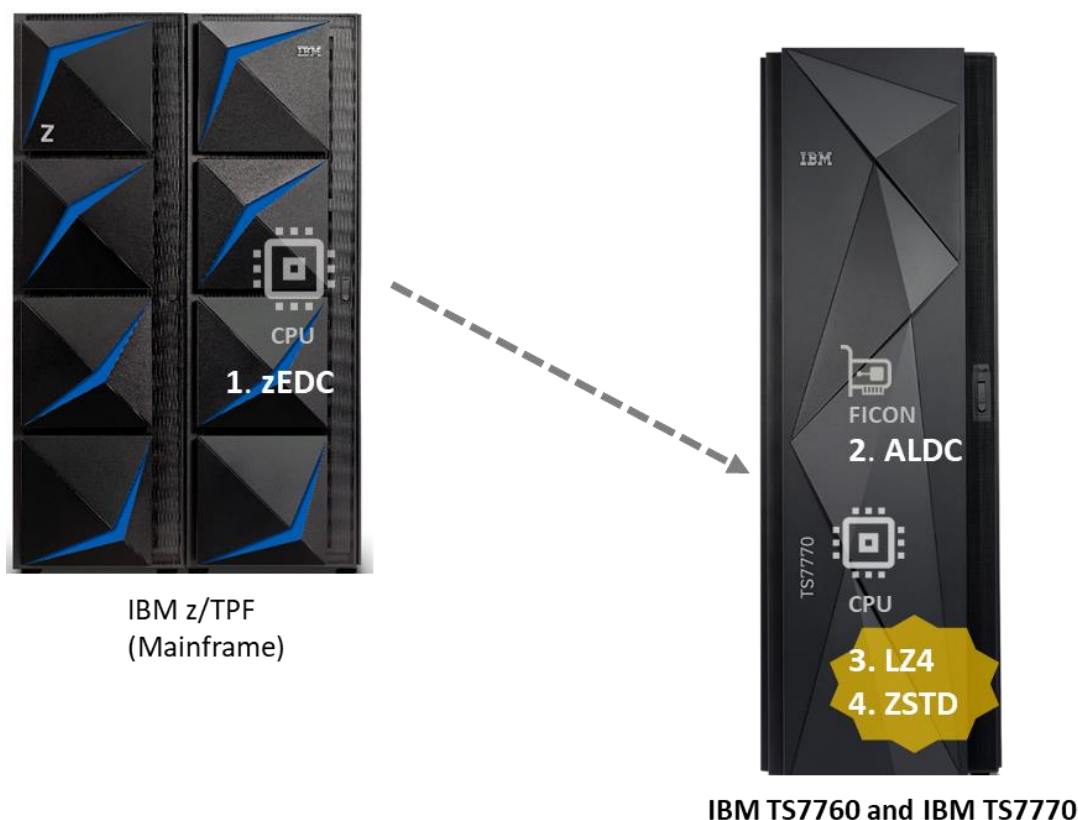
Source: Enterprise Strategy Group

Figure 9 shows the results of ESG testing of the three compression solutions available beyond the host IBM Z. zEDC was not tested because it doesn't bear directly on TS7700 performance. Plus, when zEDC is used, the host may inform TS7700 not to attempt any compression because it would not have value and would cost either storage or MIPS. ESG utilized smaller pattern files termed RECS4. All runs were made with 128 concurrent jobs.[5] As expected, the older LZ1/ALDC FICON-based

---

[5] Each job wrote/read a volume (2GB after compression at the TS7700) using a 32K block size; eight 16Gb FICON channels from an IBM Z LPAR; and tuning values: DCT=125, PMPIOR=5600, PMTHLVL=6000, ICOPYT=ENABLED, Reclaim=disabled, LINKSPEED=1000, number of premigration drives per pool=10. The TS7700 cluster used for testing was located at nearly zero distance from the IBM Z in the validation setup.

algorithm produced lower compression rates. With Release R5.0, IBM is recommending that clients choose the new software based LZ4 and ZSTD algorithms for almost all workloads.

**Figure 9. TS7700 Series Compression Data Reduction Testing**

TS7700 Compression Ratios
(TS7770, 128 jobs, RECS4 data pattern, 32KB block size)
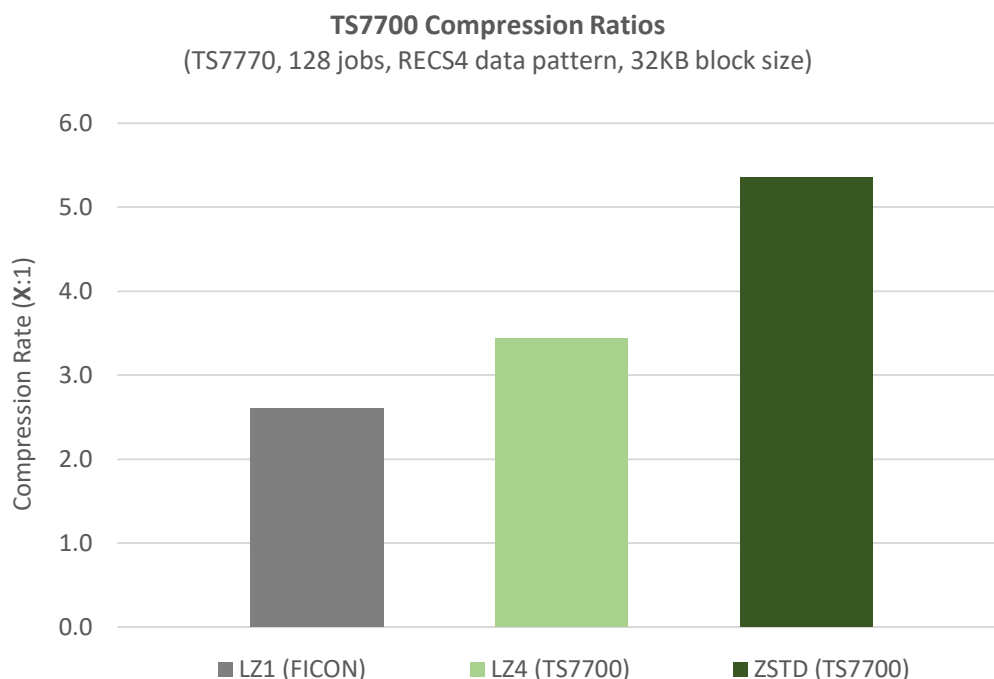
Source: Enterprise Strategy Group

The ZSTD algorithm produces better compression results but to do so it consumes more CPU resources than LZ4 and the FICON-based ALDC. This difference is revealed by the testing at different block sizes in Figure 10, which shows the performance of all three compression algorithms at different block sizes. As noted previously, ZSTD consumes more CPU resources than LZ4 and ALDC. As a result, system performance can change based on compression ratio and the CPU resources required to compress the data. At small block sizes, all three algorithms performed about the same, but as block size increased, the different compromises made by LZ4 (higher speed / lower compression strength) and ZSTD (higher compression strength but at greater CPU cost) were revealed.

At smaller job counts, ZSTD was slower in ESG testing. This is a direct result of the overhead the ZSTD algorithm incurs. Assuming the solution is not disk-cache-bound, a single stream LZ4/FICON workload can run up to two times faster than a single stream ZSTD workload. But if enough concurrent jobs are running, the cumulative performance of ZSTD equals LZ4 because ZSTD produces better overall compression. Better compression means less overhead for data movement through the TS7700. For batch periods where very few active devices run in parallel, going with LZ4 would be ideal if performance is most important. If compression is more important for these lower active device periods, then ZSTD would be a better choice. You can also mix workloads so some use LZ4 while others use ZSTD. If many workloads are expected to run in parallel, then ZSTD is usually the best choice for both cumulative performance and compression. If some of these many jobs require higher performance, they can choose LZ4 so that they can run independently faster, but with less compression. Very small drawer configurations could also produce results where lower job counts would actually run faster with ZSTD because the improved compression put less strain on the limited disk cache.

Figure 10. TS7700 Series Compression Performance Testing

TS7700 Compressed Backup Performance
(TS7760, 128 Jobs, RECS4 data pattern)
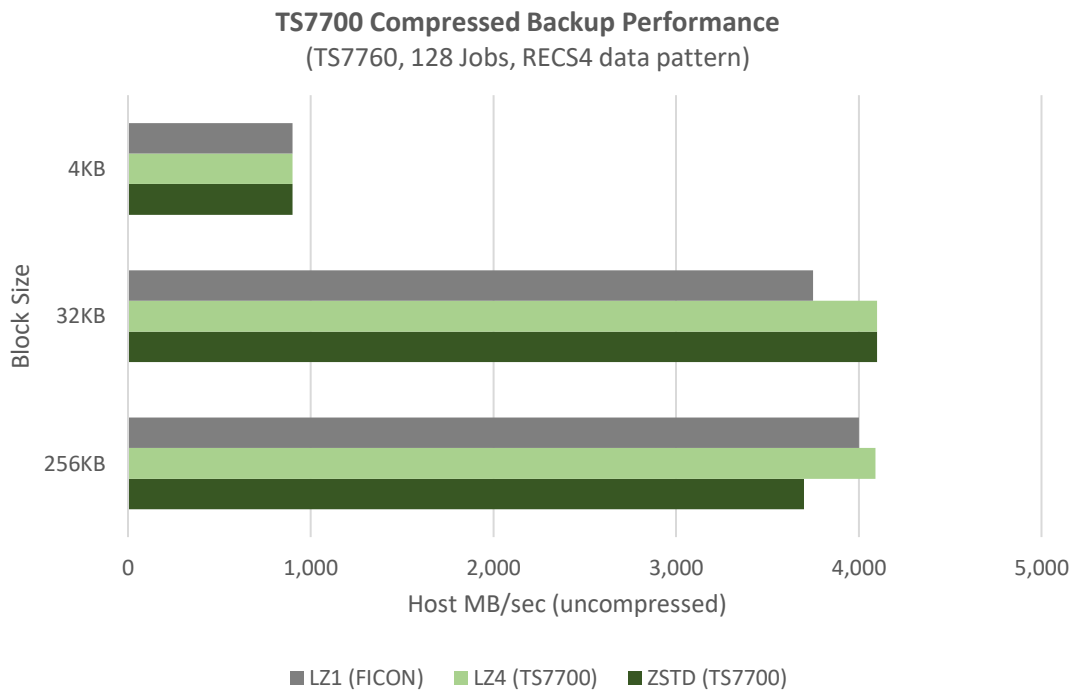
Source: Enterprise Strategy Group

ℹ **Why This Matters**

Compression is a valuable tool in leveling the playing field between exploding data volumes and stagnant IT budgets. But not all compression algorithms perform the same. As revealed in ESG testing, certain tradeoffs need to be considered when choosing a compression method—compression strength, CPU consumption, and speed, to name a few.

IBM TS7700 allows users to choose between multiple compression engines in order to match a variety of factors such as use case, block size, and performance requirements to the unique traits of each compression algorithm. And in fact, TS7700 offers users the option of mixing compression methods within the same workload. The multiple choices and flexibility of applications enable TS7700 solutions to provide effective data compression solutions across a wide range of workloads and use cases, while enhancing the cost-efficiency of the systems. These advantages will matter to every enterprise hoping to optimize cost and performance.

The Bigger Truth

The TPF operating system and mainframes in general share some important character traits. They both have roots stretching back over half a century. And neither of them would be around today if it weren't for very active R&D programs producing ongoing innovation. But the dynamic evolution of TPF and IBM Z technologies places a burden on supporting systems to keep pace. This is where IBM TS7700 stands out. As rapidly as transaction processing and mainframe servers have changed, so have IBM virtual tape solutions.

The latest iterations of IBM TS7700 support leading-edge mainframe technologies such as pervasive encryption, multi-cloud support, containerization, and object storage, to name only a few. But some things haven't changed—the requirements for performance and cost-efficiency.

ESG testing demonstrated that TS7700 solutions have not only kept up with the vibrant pace of IBM Z and z/TPF evolution, but they can more than accommodate modern performance demands while offering fresh choices to enhance cost-efficiency.

When you are looking for a virtualized storage solution that matches your mainframe requirements, you might want to look to a team that has been doing it for a while. If your organization relies on IBM Z mainframes running z/TPF as a platform for business-critical transaction processing, ESG recommends that you take a good look at the TS7700 series of virtual tape solutions from IBM. And don't overlook its ability to support other mainframe workloads like your next cloud, mobile, or big data initiative.