November 2016

Tape goes high speed

Tests confirm advantages of combining tape, flash and software-defined storage in a single, low-cost solution

Highlights

- Gain high performance, high capacity and low-cost storage all in one solution
- Leverage the cost and capacity advantages of tape to build powerful analytics systems
- Lower storage costs without losing capabilities or performance

Forward-thinking IT architects have extolled for years the benefits of building an overall storage solution using only flash and tape storage media. Place active data on flash to accelerate storage performance and efficiency; move less-active data to tape to lower overall storage costs. This storage architecture has informally been named "flape." But until recently, implementing effective flape solutions has proved problematic.

Software-defined storage (SDS) technologies, however, have changed the flape equation. SDS capabilities such as those provided by members of the IBM® Spectrum Storage™ family open up exciting new possibilities, enabling the deployment of easily managed high-performance, low-cost storage solutions using SDS components integrated with IBM tape and flash systems.

A proof of concept of this innovative new flash, SDS and tape-based storage architecture was recently performed in the Ennovar Solution Reference Architecture laboratory at Wichita State University. Using IBM Spectrum Scale™ and IBM Spectrum Archive™ SDS components, an IBM DCS3860 solid-state drive (SSD) array and IBM tape systems with analytics application workloads provided by the National Institute for Aviation Research (NIAR), the Ennovar team demonstrated the advantages of the flape architecture while increasing analytics performance by an average of more than three times.

Figure 1. "Flape" storage combines IBM software-defined storage technologies and both solid-state and tape-based storage media.

Back to the future

Next-generation analytics, cognitive and cloud workloads are driving the need for storage speed—but at the same time, managing large amounts of unstructured data places significant pressure on storage infrastructure to provide high capacity at very low cost. The Ennovar team leveraged IBM Spectrum Scale to manage the movement of data between flash and tape in a true flape architecture, and also a solution using disk for comparison. The flape architecture used an IBM System Storage® DCS3860 enclosure with SSDs as a high-performance scratch space for analytics processing. The comparison solution used a conventional spinning disk array for the scratch space. Both solutions moved inactive research data to low-cost IBM tape storage leveraging IBM Spectrum Archive.

Ennovar

The Ennovar laboratory at Wichita State University provides hardware logistics, test plan configurations, testing, data acquisition and verification. The lab's technical deliverables include best practices, technical reports, reference architectures, and sales collateral and sales enablement for the solution in multiple languages. Ennovar employs more than 50 engineering, communications, graphic design and graduate-level business students, along with full-time professionals, to lead collaborative projects with industry partners.

NIAR

In order to provide flape testing results with real world significance, the Ennovar and IBM team used an advanced general-purpose multi-physics simulation software package known as LS-DYNA developed by the Livermore Software Technology Corporation. LS-DYNA performs "transient dynamic" analysis of high-speed, short-duration events where inertial forces are important. Such events can include:

- Crash simulations (deformation of chassis or structure, airbag inflation, seatbelt tensioning)
- Explosions (underwater naval mines, shaped charges)
- Manufacturing (sheet metal stamping)

The LS-DYNA application and data set were provided by NIAR at Wichita State University. Since its inception in 1985, the institute has made a name for itself as the most capable university-based aviation research center in the United States, providing research, design, testing, certification and training

to the aviation manufacturing industry, government agencies, educational entities and other clients that can benefit from their services.

Flape project objectives

The objectives of the joint Ennovar/IBM flape project were to verify the benefits and performance of a software-defined tiered storage solution using flash as a high-performance scratch space for analytics processing while placing inactive research data on low-cost IBM tape. More specifically, the testing hoped to demonstrate the following findings regarding flape:

- Cost efficiency
  - Placing inactive data on tape significantly reduces cost.
  - The total cost of ownership (TCO) for tape on average is more than six times lower than disk.[1]
- Performance
  - Leveraging the flash tier provides for high-performance input/output (I/O).
- Flexibility
  - SSD and tape modular building blocks provide scale-out capacity and performance.
- Intelligent management
  - Leveraging the IBM Spectrum Scale ILM (Information Lifecycle Manager) policy engine to place data in the right tier at the right time—transparent to applications—can improve and simplify management.

Test bed configuration

The Ennovar/IBM team designed a flape test regime that would compare the performance of configurations using only nearline-SAS disk storage for the active-data scratch space, only flash storage for the scratch space, and a combination of both flash and disk. To provide further value, the tests also employed both Gen 6 and Gen 7 IBM tape systems. This offered a means to compare the performance improvements between the tape technology generations. Figure 2 provides details of the three test configurations.

	Balanced-workload nearline-SAS building block	Balanced-workload SSD building block	Balanced-workload mixed building block
Compute workload servers	6 IBM System x3690 X5 Dual Intel Xeon CPU E6540: 6 cores each @ 2.00 GHz “Beckton”	6 System x3690 X5 Dual Intel Xeon CPU E6540: 6 cores each @ 2.00 GHz “Beckton”	6 System x3690 X5 Dual Intel Xeon CPU E6540: 6 cores each @ 2.00 GHz “Beckton”
Controllers/ NSD servers	2 Lenovo System x3650 M5 Dual E5-2650 v4: 12 cores each @ 2.2 GHz “Broadwell”	2 System x3650 M5 Dual E5-2650 v4: 12 cores each @ 2.2 GHz “Broadwell”	2 System x3650 M5 Dual E5-2650 v4: 12 cores each @ 2.2 GHz “Broadwell”
Storage subsystem	1 IBM System Storage DCS3860 with Gen2 Controllers 60 – 6 TB nearline-SAS per system 12 Gb SAS connections Max. 60 disks, 262 TB usable capacity	1 DCS3860 with Gen2 Controllers 20 – 800 GB SSD per system 12 Gb SAS connections Max. 20 disks, 12 TB usable capacity	2 DCS3860 with Gen2 Controllers 20 – 800 GB SSD per system 100 – 6 TB nearline-SAS per system 12 Gb SAS connections Max. 100 disks, 436 TB usable capacity Max. 20 disks, TB usable capacity
IBM Spectrum Scale tape library	IBM TS4500 Tape Library, tape drives 4 LTO Ultrium 7 (300 MB/s) and 4 LTO Ultrium 6 (160 MB/s)	TS4500 Tape Library, tape drives 4 LTO Ultrium 7 (300 MB/s) and 4 LTO Ultrium 6 (160 MB/s)	TS4500 Tape Library, tape drives 4 LTO Ultrium 7 (300 MB/s) and 4 LTO Ultrium 6 (160 MB/s)

Figure 2. Flape test configuration details

For the application workload, the testing regime used 320 GB of LS-DYNA output data from an actual Boeing 737-800 crash landing simulation. This simulation represents a 30-foot drop of a single-aisle commercial aircraft full of passengers, and has more than 10 million elements. The file size grows to 335 GB during the analytics processing and requires more than 46 hours of computational process time. The LS-DYNA simulation ran simultaneously while the flape infrastructure migrated data from flash to tape and recalled data from tape to flash in the background.

Tape, the next generation of storage

The streaming I/O performance from the various flape test runs was measured on the IBM Spectrum Scale file system with public domain software known as "IOR HPC benchmark." This software is designed for measuring I/O performance specifically on parallel file systems and was tuned for sequential writes/random reads.

Figure 3 provides the test results comparing flash to disk. Essentially, 20 flash drives provided better performance than 60 nearline-SAS drives. "Cluster" refers to load balancing only on data reads, while "scatter" balances both writes and reads.

	IBM Spectrum Scale block allocation type: Cluster				IBM Spectrum Scale block allocation type: Scatter			
Drives	Write (MB/s)	Normalized (MB/s)	Read (MB/s)	Normalized (MB/s)	Write (MB/s)	Normalized (MB/s)	Read (MB/s)	Normalized (MB/s)
20 SSD	3,579	179.0	9,426	471.3	2,557	127.9	9,124	456.2
20 nearline-SAS	1,963	98.15	2,756	137.8	930	46.5	1,480	74
60 nearline-SAS	2,221	37.02	5,789	96.5	2,271	37.85	4,731	78.9
Normalized SSD: nearline-SAS (20 drives)		1.82		3.42		2.75		6.16
The average (20-drive) normalized performance of SSD was 3.54 times faster than nearline-SAS over all tests.								

Figure 3. Tests results comparing flash and disk performance

Figure 4 provides the results when comparing the performance of Linear Tape-Open (LTO) Ultrium 6 tape drives to LTO Ultrium 7 drives. With disk storage, the improvements offered by the Gen 7 tape system were dramatic. And even when flash storage was used as the analytics scratch space, the LTO Ultrium 7 drives offered significant benefits compared to the previous-generation drives.

	6 TB 7200 RPM nearline-SAS		800 GB SSD	
	Migrate (sec)	Recall (sec)	Migrate (sec)	Recall (sec)
LTO Ultrium 6(4 drives)	594	703	405	576
LTO Ultrium 7(4 drives)	385	559	311	491
Performance increase from LTO Ultrium 6 to LTO Ultrium 7	67.3%	38.66%	30.23%	17.31%

Figure 4. IBM tape system Gen 6 vs. Gen 7 performance comparison

Perhaps most important, the tests showed a negligible performance impact on data movement or application processing when the powerful SDS combination of IBM Spectrum Scale and IBM Spectrum Archive were migrating and recalling data to and from the flash and tape storage systems in parallel with the LS-DYNA simulation. Plus, the Ennovar team learned that it was very easy to scale up tape performance by simply adding additional drives.

Speaking on how valuable the flape architecture could be for enterprises running big-data analytics, Shawn Ehrstein, Director of the NIAR CAD/CAM Lab and Virtual Reality Center, noted: "*The benefit of using this type of technology would help us drive efficiencies in how we currently handle our data today. Being able to migrate and recall data sets while running our simulation applications provides a practical solution to having this in one location.*"

Figure 5. The LTO Ultrium 7 drive that goes into IBM TS4500 Tape Library.

Tape, the storage foundation

Tape storage provides the foundation of any successful flape solution. Although in the past many industry analysts have predicted tape's demise, in fact tape is enjoying a deployment renaissance thanks to the explosion of data volumes and tape's ultra-low storage costs. Some analysts believe that cloud storage might replace tape as the inexpensive foundational layer of 21st-century storage architectures, but tape actually offers many advantages to cloud service providers (CSPs), thanks in part to the rapid evolution of both tape technologies and complementary SDS capabilities. In fact, today many CSPs are deploying tape in their infrastructures to address challenges of capacity growth, scalability and cost.

In the flape testing conducted in the Ennovar lab, the performance of IBM TS1060 Tape Drives, LTO generation 6, was compared to the performance of the newest generation-7 IBM TS1070 Tape Drives. Incorporating the latest generation of industry-leading LTO technology, the TS1070 model is designed to handle backup, save and restore, and archival data storage needs with higher capacity and higher data transfer rates than the previous generation. Both LTO Ultrium 7 and 6 technologies are designed to support media partitioning and IBM Spectrum Archive technology. They also support data encryption and write-once-read-many (WORM) media.

Figure 6 shows the consistent performance and capacity improvements realized between each successive IBM tape system generation, with the metrics from the LTO generation 7 and 6 models highlighted because of their use in the Ennovar testing. TS1070 (LTO Ultrium 7) data transfer performance has increased over the previous generation with a transfer rate of up to 300 MBps with 8 Gbps Fibre Channel interface connectivity. The drive offers two Fibre Channel ports and one Ethernet port per drive to improve availability.

LTO generation to generation	Capacity increase	Performance increase
4 to 5	87.5%	16.7%
5 to 6	66.7%	14.3%
6 to 7	140%	87.5%

Figure 6. Comparison of capacity and performance increases between IBM tape drive generations

LTO Ultrium 7 and 6 tape drives are available in both full-height and half-height form factors. They are IBM Spectrum Archive-enabled for file system-based storage of data and incorporate IBM Linear Tape File System™ (LTFS) format. LTFS technology allows files stored on magnetic tape to be accessed in a similar fashion to those on disk or removable flash drives, bringing tape storage into the same range of use cases as disk. LTO Ultrium 7 drives provide 2.4 times more capacity than the previous generation of drives, with a single cartridge supporting up to 15 TB (with 2.5:1 compression) of effective storage capacity. This means that when fully loaded, IBM TS4500 Tape Library in which TS1070 drives are used can support up to 526.5 PB of data storage. TS1070 drives provide 87 percent higher performance with 300 MB/s bandwidth per drive for up to 96 GBps of throughput in a single TS4500 library.

Figure 7. IBM TS4500 Tape Library

TS4500 itself is a high-density enterprise tape library that integrates seamlessly into the IBM Spectrum Scale SDS topology. The library reduces setup and configuration time with IBM Spectrum Archive and offers automated tiering for fast data migration, plus a "single-pane-of-glass" management console and easy-to-use GUI.

Flape, the new favorite

Until the advent of powerful, deeply integrated SDS solutions such as IBM Spectrum Scale and IBM Spectrum Archive, the practicality and benefits of deploying a flash-and-tape-only storage architecture were unclear. Now, testing in the Ennovar lab at Wichita State University has proven that flash and tape storage can effectively work together to provide storage solutions with excellent performance, high capacity, and very competitive, low costs. Application workloads and business use cases that require fast transactional processing of data that will not be modified later—such as the IBM Blockchain service on IBM Bluemix®, the Internet of Things and many types of financial/banking applications—can be especially good candidates for flape solutions. In fact, almost any type of workload can benefit from this revolutionary new storage architecture. Soon, "flape" may be a word heard around the world.

For more information

To learn more about IBM IBM Spectrum Scale and IBM Spectrum Archive, please contact your IBM representative or IBM Business Partner, or visit: **ibm.com**/systems/storage/spectrum

Additionally, IBM Global Financing provides numerous payment options to help you acquire the technology you need to grow your business. We provide full lifecycle management of IT products and services, from acquisition to disposition. For more information, visit: **ibm.com**/financing

IBM Systems
Route 100
Somers, NY 10589

Produced in the United States of America
November 2016

Actual available storage capacity may be reported for both uncompressed and compressed data and will vary and may be less than stated.

[1] Mike Kahn and David Reine, "The Impact of LTO-7 on The TCO of Long-Term Storage," *The Clipper Group*, September 2015. http://www.clipper.com/research/TCG2015008.pdf

Please Recycle

TSW03515-USEN-00