



**Scaling the
PERFORMANCE
Wall**

*Spectra Raises the Tape
Performance Bar*

**HORISON**
Information Strategies

Fred Moore
President
Horison Information
Strategies
Horison.com

Introduction

I attended the Spectra® Large Tape Conference (LTC) which was held May 3-4, 2018, in Broomfield, CO. bringing together users from the large-scale tape community for hands-on product demonstrations, in-depth technology discussions, and many networking opportunities. This year's LTC topic, "Scaling the Performance Wall," highlighted Spectra's emphasis on improving tape performance for large-scale storage environments and revealed an array of new tape library performance improvements.

Increasing tape cartridge capacity to lower the price per gigabyte has long been the primary focus of the tape industry. As native cartridge capacities have presently soared to 15TB, and with roadmaps pointing to capacities of hundreds of terabytes per cartridge in the not-too-distant future, both the number and size of files stored on a single tape cartridge will steadily increase. Thus, the probability for concurrent file accesses on the same cartridge will also increase signaling the need for faster tape file access to reduce contention.

Fortunately, in the past few years, the tape industry has stepped up its focus on improving performance by delivering new capabilities including the Active Archive, RAIT (Redundant Arrays of Independent Tape), and RAO (Recommended Access Order) while providing much faster data rates than any other storage products. With tape capacities projected to soar for the foreseeable future, it's time for the tape industry to increase its focus on initial access time (time to first byte), robotic access time (mount time), and overall throughput. Presentations at the LTC revealed that effort is well underway at Spectra.

Rational Robotics Improves ExaScale Library Performance

Spectra directly addressed tape library performance by introducing a series of improvements called Rational Robotics for the richly featured [TFinity® ExaScale](#) Tape Library which is currently the industry's highest capacity (up to 1.6EB compressed) library. Rational Robotics is a combination of intelligent

hardware and software tape library features that minimize the time required for a robotic library to locate a tape cartridge and access data on that cartridge while improving library and drive reliability.

Rational Robotics Software Optimizes Robotic Movement

Spectra has developed new software-enhanced management controls for greater productivity that helps the library's high-performance transporter (HPT) perform even faster with more efficient movements with Media IQ™ and Slot IQ™ (see below). The Media IQ software feature queues requests from all connected hosts and then assigns them to movers based on their optimal proximity between media and drives. This allows the library to assign the robot that is best positioned to service each particular request in the least amount of time, providing the most efficient robotic movement for reduced cartridge access times.

The Slot IQ software "move" algorithm will virtualize the slot locations or physical inventory inside the library and take advantage of the empty spaces (slots) nearest to tape drives. Slot IQ allows the robots to physically move shorter distances or less often as they take advantage of the available storage "holes" closest to the drive bay, eliminating unneeded robotic movement and improving performance. This reduces time when inserting new cartridges into the library. Media IQ and Slot IQ are Spectra-exclusive time-saving and intelligent streamlining features.

HPT (High Performance Transporter) Redesign

The TFinity ExaScale high-performance transporter is a redesign of the robotic hand used to manipulate media and represents the next step in library performance and reliability. The HPT uses new robotics and firmware and is designed with four goals: better performance, better reliability, mixed media support, and better sensing.



The new HPT reduces tape mount time (faster access time/performance) and has increased the MTBF (mean time between failure) for better reliability. New sensors and features, including temperature and humidity readings, provide additional reliability of Spectra's HPT. The HPT is the central technological improvement that now drives the industry's fastest library.

TAOS (Time-based Access Order System) Reduces File Access Time

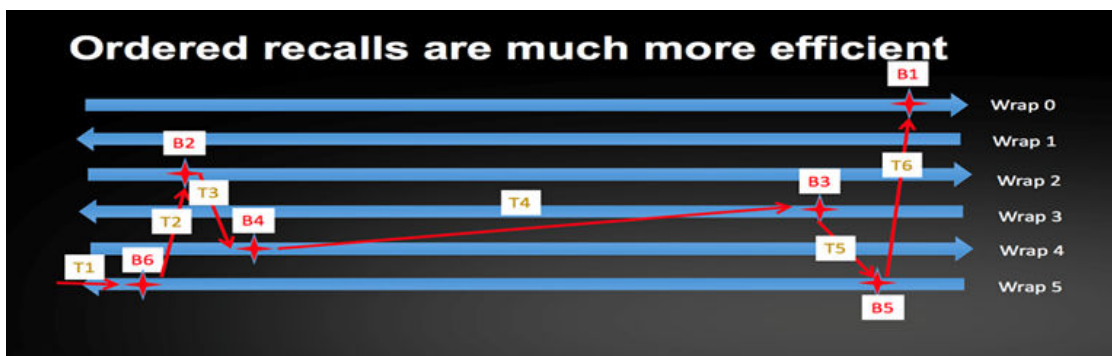
In addition to several advanced robotic performance improvements for the TFinity ExaScale library, Spectra also addressed improving the file access time (time to first byte) with a new capability called TAOS. TAOS is a unique – and long overdue - advancement for LTO that provides up to a 4 times improvement in overall access speed on reads and up to a 13 times reduction in physical tape movement across the drive head. Less tape movement results in reduced tape media and drive wear providing the additional benefit of improving overall tape system reliability and reducing cost.

Tape files are written on tape in sequential order but are most often accessed (reading data) in random order. Files are laid out in wraps that follow a serpentine pattern starting at the physical beginning of the tape, traversing all the way to the end of the tape, then back to the beginning. An LTO-7 cartridge has 112 wraps and LTO-8 has 208 wraps. As tape capacities and the number of files on a cartridge continue to grow, file access times will increase as contention mounts. Retrieving non-consecutive files from a tape can result in inefficient seek times between file reads. With tape media length of nearly 1,000 meters, making inefficient seeks to files can become very costly. The **red arrows** below illustrate the total distance traversed to access 6 files (T1-T6) in standard linear mode without TAOS.



TAOS Orders the Read Requests

The TAOS algorithm runs on the Spectra library processor. This implementation allows direct access to tape information upon load and provides a simple interface mechanism to the higher level HSM software. With an estimated amount of time to go between each file on tape, TAOS uses a “Nearest Neighbor” algorithm to develop the list of files in order to create an optimally ordered list. Additional optimization techniques can be applied to this ordered list to further reduce the overall time needed to recall all the files requested. The **red arrows** show the reduced distance travelled using TAOS to access 6 files (T1-T6).



When a list of multiple files to recall are provided, every file’s start and end position are translated to a physical position on the tape and an estimate for how long in milliseconds between each file’s end position and every other file’s beginning locate time position can be calculated. An optimized best access order list provides the least amount of time that is needed to locate and read all concurrently requested files or data sets on a tape. Spectra indicated that read requests for files less than 100MB in size typically result in a 4 times improvement in overall file access time making smaller files more appealing for tape. TAOS also improves the performance of restore operations making the recovery process from tape much faster.

TAOS represents a major and long overdue breakthrough for LTO file access times. TAOS will take on greater importance going forward as tape cartridge capacities and, therefore, the number of files on a cartridge, continue to increase. TAOS support is presently available for LTO-7, LTO-8, and future LTO drives. TAOS is similar to the RAO (Recommended Access Order) feature for IBM's enterprise tape drives. One of the most appealing features yet available for LTO, TAOS can significantly reduce access time.

Summary

Spectra is scaling the performance wall with the new HPT, the Rational Robotics features Media IQ and Slot IQ, and TAOS. These new features have given the tape industry a huge overdue boost providing much faster initial access times and throughput levels. Large amounts of unstructured object data, including text and binary data, archives storing big data, images, videos, surveillance, audio, and scientific data are ideally suited to be stored and retrieved most cost-effectively on modern tape, and these applications will place additional demands on future tape performance. The impact of the IoT on tape demand is unknown. These emerging trends would present an enormous problem for large-scale tape storage systems, and possibly a potential impediment to future storage growth, if not for Spectra's focus on performance.

TFinity Exascale Library Performance Features	
Feature	Summary
High Performance Transporter	Faster Robotic Moves, Multi-Media Support, Improved Reliability
Media IQ	Sorting Move Commands and Optimizing Move Sequence-Based on Robot Location
Slot IQ	Take Advantage of Empty Library Slots Nearest Tape Drives to Improve Performance
TAOS	Reduce File Access Time by Optimizing Tape Movement

I left the Spectra LTC refreshed and energized, realizing that many aspects of tape performance were seriously being addressed. The message and benefits were evident – by ‘Scaling the Performance Wall,’ Spectra has clearly raised the tape performance bar.