

TAPE ACCORDING TO SPECTRA LOGIC

GET READY FOR A NEW RIDE

OCTOBER 2013



I attended the Spectra Summit for the first time earlier this month in Broomfield, CO. I figured it would be another analyst day held annually by larger vendors to bring the analyst community up to date on their strategy and ensure the vendor gets adequate mindshare from the analyst community. In some ways, Spectra Summit 2013 was similar and in another way, entirely different.

Spectra Logic has been consistently gaining market share in the mid-range and enterprise tape library market for several years in a row. Granted the market has been shrinking steadily at a slow rate of decline given the movement towards disk-based technologies over the past decade. Yet in this declining market, Spectra Logic has stayed true to their calling and maintained a laser-like focus on tape-based technologies. In so doing, they have continued to extract market share from Oracle (StorageTek), IBM, Quantum and HP.

You would think they would be happy with that. You would be wrong.

Under the Covers

Spectra Logic has been working under the covers to find ways to bring tape back into the mainstream of enterprise IT and Cloud. They have recognized (and insisted that the market listen) that tape has certain inherent advantages over disk (even the lowest cost disk) and these include removability and cost. But that logic has fallen on deaf ears in the recent past and disk has been eroding tape revenues for a while.

Enter Big Data, Cloud, increased desire on the part of the business units to extract meaning from mounds of data, and increased pressures from governmental and other agencies to preserve data for long periods of time for compliance reasons, and everything about tape use changes. In spite of data deduplication technologies that have made disk superior to tape in backup environments, tape still enjoys serious cost advantages when it comes to massive amounts of data that does not need immediate access. And this is what Spectra Logic has been concentrating on.

The Strategy

At the Spectra 2013 conference they exposed their long term strategy. They first defined what they call “deep storage” where large amounts of data need to be kept for decades (maybe forever) at very low costs, using very little power, requiring little cooling, yet maintaining excellent data integrity and delivering “reasonably fast” online access. This storage also needs to be restored at file or object level, rather than at an entire volume level. And it needs to be accessed without knowing the path to where it is located. Add to that the requirement that scaling must be seamless, i.e. storage systems cannot have boundaries that the application (or the user) need to be aware of.

When you marry tape technology with an object interface, interesting things start to become possible. And this is exactly what Spectra Logic has done. They have developed BlackPearl, a

“caching” box so to speak, that has a RESTful interface (plus 10GbE) on the front side and FC or SAS on the backend that would connect to the tape libraries. The object-based API is DS3, an extension of Amazon Web Services’ S3 interface. LTFS is supported on the backend to enable “random access out of a sequential device.” In addition, BlackPearl maintains a catalog and all the metadata associated with the data located in the backend tape libraries. Tape Library management, data integrity verification, data security integration are some of the other capabilities of BlackPearl. SSDs inside BlackPearl act as a cache rather than a tier and make more frequently accessed data readily available without invoking the tape library. The object interface on the frontend presents a flat address space and allows access using web-based technologies, rather than CIFS/NFS and other traditional storage technologies. The application only needs to know what it is looking for, rather than where it is located in a sea of data.

DS3 API was needed because S3 was designed for making many short transfers rather than a small number of large transfers. Its HDD and flash orientation is right for using AWS’ cloud for backup and DR purposes but not for bulk transfers of billions of objects that may be located in a deep archive. DS3 adds “Bulk Put” and “Bulk Get”, along with a few other commands to S3.

Three configurations of BlackPearl were announced: one version sits in front of a Spectra T950 tape library with TS 1140 drives and this configuration enables up to 6.4PB of raw data (before compression) to be stored at 9c/GB. The second configuration associates BlackPearl with a T950 tape library using LTO drives and this enables storing 2.4PB of raw data, at 10c/GB. A third configuration uses T380 tape library to store up to 1.9PB at 14c/GB.

The implications of BlackPearl are huge, not simply for Spectra Logic but to the industry. First and foremost, it brings tape back into the game. For most storage admins, tape had become something to run away from. If Spectra Logic is successful, and I am betting they will be, we will see tape back into the conversation at enterprise IT. Look at it this way. You can send your data for long term storage to AWS Glacier and pay 1c/GB/month, or you can buy and own BlackPearl and a T950 tape library for 9c/GB. In other words, you can create your own private cloud and in nine months, reach a point of equivalence with the most inexpensive long term storage available in a public cloud today. Granted you have to maintain the system yourself but, at this pricing, and the relentless effort on the part of Spectra Logic to make day-to-day management of tape libraries “disappear”, it is hard to ignore the strategy.

I expect media and entertainment, life sciences, genomics, oil and gas, several branches of DoD, healthcare and web services companies will show a lot of interest in this methodology. NASCAR and Yahoo! were at hand to lend support to the idea. Spectra Logic has already written a DS3 client module for Hadoop and this will allow Hadoop applications to directly dump to and extract data from tape libraries, without doing backup. It is hoping that many ISPs will write client code using DS3 API for other applications that can then interface directly to Black Pearl. The game has only just begun.

At Taneja Group we have realized for a long time that only ignorant people think tape is dead. However, we have also realized for a decade or so that tape revenues were declining in light of low-cost SATA HDDs and brilliant data deduplication technologies. For a while the pendulum was shifting rather hastily towards disk. LTFS pumped some new life into tape but still suffered from the classic file system limitations of hierarchy, limited number of objects it could handle, and the requirement for knowing the path to the data one wanted. DS3 and Black Pearl inject new life into tape in a major way.

But more importantly, it makes you re-think tape technology for deep archive. Spectra Logic showed its prowess by maintaining a steady growth over the years in mid-range and enterprise tape libraries in a declining market. This move shows its creativity goes beyond making solid tape libraries. It shows thought leadership and ability to think out of box. We believe many application

vendors will write client code for DS3 and direct dump to tape. Let's not forget application vendors don't get involved in religious wars between disk and tape; they only care that their applications are used to interact with data, whatever media it is stored on. Bravo, Spectra!

NOTICE: The information and product recommendations made by Taneja Group are based upon public information and sources and may also include personal opinions both of Taneja Group and others, all of which we believe to be accurate and reliable. However, as market conditions change and not within our control, the information and recommendations are made without warranty of any kind. All product names used and mentioned herein are the trademarks of their respective owners. Taneja Group, Inc. assumes no responsibility or liability for any damages whatsoever (including incidental, consequential or otherwise), caused by your use of, or reliance upon, the information and recommendations presented herein, nor for any inadvertent errors that may appear in this document.