

STORAGE SWITZERLAND

A SOLUTION TO THE PROBLEM OF 'FOREVER DATA'



Eric Slack, Senior Analyst

Data sets are growing, but so are the periods of time that they're being saved. Once primarily driven by regulatory compliance, companies are now finding that there are other factors pushing data retention to seemingly unlimited duration. The repurposing of digital content, the explosion of data analytics and the longevity of cloud storage are all causing storage problems for companies today. A new "Deep Storage" interface (DS3) and tape appliance called BlackPearl promises to bring a solution to this problem of storing more data sets, essentially forever.

The 'Forever Data' Problem

Digital assets, like old movies, are being saved without expiration dates to support the potential for future reuse or 'monetization', such as anniversary releases of Disney classics. Data sets used in time-based analytics are also being kept on an open-ended basis in order to provide longer historical comparisons. Even cloud providers are having retention problems as they find many customers have little motivation to delete old data sets. This is data the provider has promised to store forever.

These growing requirements to keep data for extended periods are causing companies to re-evaluate their long term-storage infrastructures and look for some new alternatives to traditional disk storage and archive-based

systems. "[Deep Storage](#)" is a term introduced by Spectra Logic to describe a new storage architecture designed to address the problem of forever data in terms of accessibility, capacity, cost, longevity, etc.

Deep Storage

Deep Storage is a tape storage tier that's been made accessible to applications and workflows using a REST-based interface called "[Deep Simple Storage Services](#)" (DS3). DS3 is created by adding 'bulk' data movement commands to the S3 protocol to support sequential access tape operations in addition to the random access disk operations traditionally supported. By connecting this high capacity, long-term storage area to an object storage system or directly to an application, deep storage brings along tape's benefits (economics, density, longevity, security and transportability) to any environment.

What is BlackPearl?

[BlackPearl](#) is a 2U rack-mounted appliance that provides SSD caching and a REST interface between applications or storage systems and tape, using the DS3 interface. On the front end it supports multiple connectivity options, including 10Gb Ethernet and InfiniBand.

On the back end it handles the tape interface protocol, the library management and the data buffering required to support tape's higher single stream throughput. It also manages the catalog of input data objects and the conversion to LTF5-compatible output that's written to each tape cartridge. And, BlackPearl handles security and long-term data integrity in conjunction with Spectra Logic's BlueScale library management software.



What BlackPearl Does

The BlackPearl appliance takes objects as input from the client using the DS3 interface, consolidates them into larger objects called “buckets” and writes each bucket to a tape drive. BlackPearl manages the distribution of these objects in each bucket so that the appropriate tape cartridges are read to recall the objects requested by a client. Currently, LTO5, LTO6, and TS1140 tape generations are supported and data is stored using the LTF5 format.

As a library interface that's Internet enabled, BlackPearl allows tape libraries to be located remotely from the application servers. In the past, with tape drives connected via Fibre Channel or SCSI-based protocols, most users would locate libraries in the same data center as the application servers. This fact alone makes deep storage a viable option for cloud environments, without complicated data handling and replication, etc.

To be clear, BlackPearl is not an object storage system; it's an appliance that sits between tape libraries and any REST-enabled server or storage device (see diagram above). This makes the BlackPearl appliance an ideal way to add a deep storage tier on the back end of an existing object storage infrastructure.

What BlackPearl Doesn't Do

BlackPearl does not create data objects; this must occur in either an object-based storage system or a REST-enabled client driver that's integrated with a software application. As in all traditional, RESTful architectures data management is done by the Client and storage management is done by the Server. BlackPearl is the Server component that handles, stores and retrieves data objects via the DS3 interface.

BlackPearl also doesn't read or modify the metadata that's part of each object. It does manage the placement of objects within data buckets but doesn't know anything about the contents of the data it's storing below the object level.

While it's true that any S3-compatible storage can connect to BlackPearl, a modification into the DS3 interface would be needed in order to fully leverage the advantages that tape brings. For most implementations the primary method for integrating deep storage would be through a DS3 client that's written into an application.

DS3 Client Development

Spectra Logic has created the first DS3 client, one that's designed to bring Deep Storage to Hadoop implementations. But the plan moving forward is for end users, complimentary hardware vendors and ISVs that can benefit from Deep Storage connectivity to develop their own Client applications. To support this activity Spectra has created a DS3 Developer community with a Software Developer Kit and a Developer Forum to facilitate the exchange of ideas within the community. Resources available include APIs, client code examples, a simulator download for testing that code, plus documentation.

LTFS

With the BlackPearl interface appliance, the benefits of LTFS can be leveraged in the deep storage system. As an open-source, self-describing, file-aware tape format, LTFS enables LTO and TS1140 tapes to be interchanged between platforms that support it. This means deep storage data sets can be transported and handled efficiently, in bulk, even between different software applications or vendors in the workflow. BlackPearl also allows for the ingest of LTFS tapes in traditional, file-based formats.

Storage Swiss Take

Tape has always been able to provide benefits for storing data sets that are too large or kept for too long to make disk storage feasible. But accessibility has been a problem for many environments because tape interfaces have been application-specific or they required file-based archive systems in order to communicate with workflows. With the advent of object storage, workflows can now directly access data via standard web protocols.

About Storage Switzerland

Storage Switzerland is an analyst firm focused on the virtualization and storage marketplaces. For more information please visit our web site: <http://www.storage-switzerland.com>

Copyright © 2013 Storage Switzerland, Inc. - All rights reserved