STORAGE SWITZERLAND

HOW BACKUP VIRTUALIZATION CAN SAVE VTL

Storage Switzerland, LLC

George Crump, Senior Analyst

Many data center managers, especially those responsible for open systems, have a low opinion of Virtual Tape Libraries (VTL) and as a result, VTL adoption within the data center has been equally low. There are however, data protection challenges that data centers do need solutions for. Backup virtualization technology may provide the solutions to these problems without the challenges caused by VTL.

What is Backup Virtualization?

Backup virtualization is similar to server virtualization in that it places a layer of abstraction between two components of technology. In server virtualization the abstraction is between the server hardware and the operating system. With backup virtualization the abstraction is between the backup application / server and the backup or archive devices that the server reads and writes data to. Like server virtualization they are typically open in nature, supporting almost any application as a source and almost any device as a target.

In comparison the VTL is typically a 'sealed' system: essentially a disk array with software intelligence to make the disk act like a tape library. The supplying vendor almost always forces the user to purchase disk and tape from a single source, and in most cases doesn't support mixed-target devices.

Problem 1: Integration of Tape

Despite the popularity of disk to disk backup and cloud based storage, many, if not the overwhelming number of large data centers, still count on tape as their primary backup and archive storage mechanism. This isn't a disk vs. tape discussion, but a current reality vs. future possibility issue. There may be an economic justification to stay with tape or just a familiarity with the tape process. In either case, one of the key problems that many VTL systems have is that they either relegate tape to a secondary status (as a weaker technology) or worse, they replace tape all together and have no clean way to move data from disk to tape without going through the backup server again. Customers that are not swayed to replace tape are forced to develop convoluted workarounds to maintain data in both the VTL and the tape environments.

Backup virtualization solutions like Tributary Systems' Storage Director, are able to provide a tight integration of tape to the backup virtualization platform, leveraging disk to make tape *better*. For example, disk can be used to cache data as it's being backed up, data which can then be queued up to send to tape quickly. Doing so keeps tape drives spinning at full speed without slowing down for repositioning. This reduces the number of tape drives needed, improves the life span of those drives, and reduces the number of tape mounts that the library needs to handle.

Problem 2: Reduction of Backup Window

The second problem area that data center managers are looking for VTLs to solve is the reduction of the backup window. The first challenge is that any product, VTL or backup virtualization can't overcome the speed of the network. In the best case, they can only receive data as fast as it can be sent. However, many of these larger data centers have invested in infrastructures like Fibre Channel SANs that can deliver data very quickly. Assuming that's the case, VTL still has a specific problem when trying to reduce the backup window - moving the data to tape. Many data centers do not declare a backup complete until the data is either on tape or on tape and off-site.

Many VTLs on the market cannot stream data to tape while receiving and storing data to disk. They have to wait until all disk based backup jobs complete before sending to tape. Worse, the move to tape often is not integrated at the back end of the device, meaning data has to go back through the backup server to get to the tape device, which is slow and places a load on the backup server.

Backup virtualization, on the other hand, is designed to cache the inbound data to a disk buffer before sending it to the tape device. The data center manager has the option to either have that data removed from disk as it's being written to tape, or it can be held on disk until space is needed. For organizations that don't declare data protected until it is on tape, this is an ideal situation, one that also keeps data local for as long as possible to help with the recovery process.

Most importantly this data is written to tape directly from the backup virtualization appliance. It doesn't need to go back through the backup server to be processed and copied to the tape device. This allows for extremely high speed copies to tape without impacting backup application performance, including its ability to continue to write data to the backup virtualization appliance.

Problem 3: Consolidation Of Devices

This problem has actually occurred primarily as a result of the failure of VTL to be the focal point for data protection in the data center. Since VTLs do not resolve the above issues, many alternate solutions were tried. Environments today have often accumulated a variety of tape and disk solutions in repeated attempts to reduce the backup window or integrate tape. Many of these devices have capabilities like deduplication or power efficient MAID that can be appealing in certain environments, but can be problematic in others. For example, MAID may be a good way to store the quarterly full on, but not as good for daily incremental backup jobs.

Backup virtualization allows all of these devices to be consolidated behind a single appliance, allowing them to be used at the same time across applications and data sets as it makes sense. Devices of similar or complimentary capabilities can be grouped, allowing the backup application to communicate with a single device. For example, full backups can be written to what the application thinks is a single tape library, but the backup virtualization device may actually write that data to high speed disk first, then make a copy to tape for off-site data movement, then move the disk copy to a deduplicated disk for longer term on-site retention. Backup virtualization also allows for easy migration to new tape or disk technology when those become available without having to redirect backup jobs or start over from scratch.

VTLs have only achieved limited adoption because they haven't solved the problems that data center managers faced, like reducing backup windows, nor do they compliment the technology organizations want to use, like tape. Backup virtualization not only addresses these challenges, it actually makes tape better, as well as solving new problems such as consolidating desperate devices.