

# STORAGE SWITZERLAND REPORT

## BACKUP OPTIMIZATION - GETTING THE MOST OUT OF YOUR BACKUP INVESTMENT



George Crump, Senior Analyst

In life there are three 'givens': you can never be too thin, you can never be too rich, and your backups can never be completed too soon. This last given has led to a tremendous amount of growth in the data protection marketplace. There are solutions that protect in real time, recover almost instantly, optimize backup data so it can be stored on disk for a long period of time, and those that can replicate backup jobs. The problem is how do backup managers make sure they are spending their precious budget dollars on solutions that will provide the most improvement for their investment?

Backup managers, just like storage managers, need tools such as [Tek-Tools Backup Profiler](#) to understand the nature of their backup environments and make sure that the solutions they invest in will fix, or at least ease the problem. Similar to the storage manager, backup managers have classes of service when it comes to data protection. There's data that is protected in real time, data that's protected to disk, and data that's protected to tape. No single protection class is appropriate for all data. To maximize their investment the backup manager wants to send the right data set to the right destination.

Backup optimization is often an endless ride where backup managers try to fix individual segments of the process to shorten the backup window. The first stop is often the backup infrastructure itself. It would seem after all, that building a faster highway is a sure way to get data from point A to point B more rapidly. The first step however, prior to making any additional infrastructure investment, is to understand the capabilities of the existing environment. If the server to be protected can only send 4MB/s of data at a time, then putting that server on a 1GbE segment is going to provide limited, if any additional performance. A backup analyzer software solution can show the transfer rates from each client that's being protected and determine if a faster network is warranted.

This limitation may not be the fault of the client. While the client itself may have plenty of horsepower the application that's creating the backup stream may also be limiting performance. Database applications and email systems are notorious for having slow APIs that simply can't generate enough backup data, especially during online backups. Another culprit is file servers with very high file counts. The backup application can spend an inordinate amount of time inspecting all these files to determine if they need to be backed up.

A backup analysis tool would be able to determine the cause of the problem by allowing you to compare different backup streams. For example, in an SQL backup you could do a dump of the data from within the application across the network and compare that to the backup time from the backup application that the backup analysis tool would provide you. If both show slow results then you may have under-powered application hardware. If one shows significantly better results than the other you may have a backup software issue. If both are performing at the maximum of what the current network can deliver, then a network hardware upgrade may be called for.

In addition, a backup analysis tool like Tek-Tools Backup Profiler can also provide file information that will help in file server backup. The time from a backup job's start to the point where data begins to move across the network should be examined. If there is a significant lag time then there may be an issue with the number of files being backed up. The file server could then be analyzed to determine how many files it has on it. Ideally, older files should be archived and moved out of the backup path, a function that again, can be provided by the proper analysis tool.

The next step is often focused on the backup target. If the network is fast enough then a faster backup target may help complete backup jobs faster. It's not just fast clients that can see a benefit from faster targets. The type of target may be a better match for certain types of backup jobs. For example, a long backup job that's due to a slow network or slow backup application could benefit from a random access device like disk that can spin at a constant speed no matter how fast or slow the backup data is coming in from the client. Tapes have to slow down and spin back up based upon the amount of data being fed to them, often called "shoe shining", something disk does not. A slow client with either a lot of jobs or a lot of files may benefit from a faster target. These types of clients are often the more frequent sources for restoration. Again, a backup analyzer tool will help in this determination.

Tape drives are often left out of the backup optimization discussion. When looked at in terms of performance per

capacity, nothing is more affordable than tape. And in many cases, if the tape can be kept streaming at full speed then there's a good chance that a modern tape drive can provide significantly better performance than disk. However, tape as a target, of course, brings a concern about recovery speeds, because the data has to be searched sequentially on the media.

An ideal scenario for tape would be as a device either locally attached or attached via a high speed network connection that also doesn't get frequent restore requests. Application servers typically fit this bill perfectly. In addition to the backup process, they're often protected via snapshots or replication, they're on high speed servers and networks and finally, are often in some sort of highly available mode - via a cluster or something similar. All that protection means that the backup would be the point of last resort for most application recoveries. After all, why waste expensive disk backup resources on this?

This scenario has a lot of moving parts and separate processes providing protection of data. This is the perfect situation for a backup and storage analysis tool. It allows confirmation that snapshots and/or replication copies are being made frequently and it confirms that backup jobs are moving to tape in a rapid fashion. In many cases, the protection can be seen from the point of view of the application, the storage or the backup. Each of the stakeholders in the protection of that data can see from their perspectives that everything is working according to plan.

The result is less data going to expensive backup disk, more data going to high speed tape, yet better leverage of the other data protection processes, saving time and money.

Backup optimization is not only accomplished by upgrading and improving the components of the backup process but also leveraging the capabilities of the surrounding environment. This can be archiving to move old unchanging data out of the backup process, or correctly targeting servers being backed up based on their backup and recovery needs.

### **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>.**