

# A Growing Need for Storage Optimization

Posted by Arthur Cole on October 3, 2008 at 5:40 pm

It makes no sense to put non-crucial or little-used data on your most expensive storage tier. But how many of you actually know what percentage this static data occupies on your primary storage systems?

Apparently, not many, or else there would be a massive industry shift to more effective storage systems.

According to a recent survey of Fortune 1000 companies by Copan Systems, more than 70 percent report storing what Copan calls “persistent data” on their primary tiers or on a combination of primary and archive tiers. And while the actual percentage of unused data residing on primary tiers varies from enterprise to enterprise, the company estimates that about 70 percent of total enterprise data is static, meaning there’s likely a lot of useless information clogging up your system right now.

The key to solving this problem is not simply adding an archival system, however. What’s needed are more effective means of storage optimization, intelligent software that can identify static information and move it off the primary tier, and then seamlessly move it back when someone goes looking for it.

According to George Crump, analyst and founder of Storage Switzerland, the habit of archiving data after a year or two of inactivity will have to go. Future migrations will take place after several days, turning primary storage into essentially a large cache. The rise of solid-state disks, with their extremely fast I/O performance, should fit into this new environment nicely, accompanied by new migration and utility software, as well as new lines of intelligent storage controllers.

One of the newest software approaches is from a company called [Double-Take Software](#) of Southborough, Mass. The company is taking aim at unstructured data with its Double-Take Cargo that leaves a small “stub file” in the primary server whenever an inactive file is removed. The stub file provides a transparent link to the full file should anyone request it.

Still another approach is a new generation of appliances that can accommodate not only optimization, but disk-saving techniques such as compression and deduplication.

The storage farm is the perfect example of the old adage “More is not necessarily better.” Virtualization and advanced networking will produce tremendous pressure to add more storage, but it will not be money well spent unless you have the right tools to optimize your overall data infrastructure.