

Corporate IT Managers:

Are you losing control of your data?

Use these eight steps to immediately slash storage requirements and costs

=====

Cheap storage is the crack cocaine of corporate IT, and the addiction has allowed unstructured data to explode. As analysts estimate annual data growth at 50 percent, the day of reckoning can't be far off.

It's time to kick the seductive – but wasteful – habit of buying excess storage rather than establishing sound storage policies. And the sooner the better, because every month you wait means more data to deal with. Meanwhile, belt-tightening execs might even now be turning their attention to reducing storage costs.

Fortunately, savvy IT managers have access to powerful solutions. By taking action now, you can gain control before it's too late.

Eight steps to optimize your storage system

Sounds straightforward. Minimize the capacity required to house your information, and access it more efficiently. That reduces the overall data management cost per GB. But it takes careful planning and specialized tools to help you do it.

Job one – develop a tiered storage system

A well-planned, tiered storage system has always been a sound move. Today, it's absolutely essential, and it should be the first step in your storage optimization strategy. By storing less important or active information on lower tiers, you make the primary tier more efficient, and you reduce your average storage costs.

To begin with, your system can be as simple as two tiers, with the enterprise archive on the lower tier. Add intermediate tiers as you determine the need.

The challenge is deciding what goes where. It's not a job for users, and you need to automate it as much as possible. Tools which analyze the global storage and reveal which files are dormant or haven't been accessed in a long time (among other parameters) are one step. And archive packages can crawl file systems to identify which files should be moved, moving them when system activity is low.

Build a true enterprise archive system

Often, numerous backups end up becoming your "archives", meaning duplication runs rampant, and finding information becomes onerous.

A true enterprise archive system is smaller and more accessible. It significantly reduces backup storage and the time required, since it's not included in regular backups. And it's easier to weed out valueless data from a single archive rather than numerous backups.

Fortunately, MAID-based storage systems have made lower tiers more reliable and efficient, which allows you to archive even more data, freeing up more tier one space.

Dedupe and cull

To minimize capacity requirements, you need to eliminate duplication and stop wasting space.

Backups are the biggest culprits in duplicate data. While it seems prudent, backing up data sets that are only slightly different from last week's eats up a vast amount of storage space. And while it can be tough to determine what data is obsolete, and therefore of no value to the organization, eliminating this information can both free up storage space and improve accessibility – particularly if it, too, has been duplicated.

Investigate systems designed to cut down on duplication and wasted space. These systems dedupe and compress data across the enterprise, and they do it in real time – so they may be used

for primary storage as well as backups. They can also reduce capacity needs since they don't need dedicated space to cache data while they're deduping.

Reduce file size

An enterprise archive will reduce your primary tier storage requirements significantly. By reducing file size for popular software, you can reduce it even further. Some packages claim to reduce popular files by as much as 90 percent – without the need to decompress them before use – and further reduce already reduced files, such as compressed image formats.

Consider cloud storage

Cloud storage can provide faster scalability – at a safe off-site location – and as a "sexy" technology, you'll likely be called on to explore its relevance to your organization. Cloud storage can be used for active data, but more often you'll start with backup or archived data.

If you're considering using it for archiving, you need to have periodic, ongoing access to the data. So find out how the provider enables you to move data in and out of the cloud. Some providers are already allowing for API sets, which allow users to control migration and retention parameters, taking some of the load off of you.

Virtualize your files

File virtualization systems can help automate cloud archiving, and make access to data stored there more intuitive.

These work like file analysis tools, identifying archive candidates by size, access activity, and other parameters, but they also move the data to the cloud gateway and set up transparent links.

They also help users access files by simply knowing a name and not the location, like accessing a website by its domain name rather than its IP address.

And determine how granular you need to get. Some systems are less specific than others – allowing parameters to be applied only to a directory level rather than a file level.

Employ image backup

Think of image backup as a snapshot of your system at a point in time – a snapshot that includes system state, operating system, application configurations and data – all in a single file, which can be used for anything from full restores to individual file restores.

The process uses server images, such as those created by VMware, to provide a comprehensive representation of the system. Protecting the image is faster and easier than traditional file-based processes, which copy and manage thousands of files.

Know your storage needs

Greater efficiency and storage optimization reduces the need for constantly increasing storage space. But you still need to be able to project your storage requirements – not just a year down the road, but next week.

Fortunately, a number of tools now monitor capacity planning and storage utilization, and can warn you of capacity problems that may crop up. The analysis helps you determine whether you need to buy new storage, how much you need, and where to deploy it.

IT managers ignore unfettered data growth at their peril. By taking charge of data optimization now, you'll save yourself and your company massive headaches down the road. And you'll anticipate and manage an issue that, if it isn't already, will soon be on upper management's radar.