



Stretching Storage by Shrinking Files

Older Data Gets the Squeeze in
Solution-Soft Compression Strategy

Article by Christopher Bucholtz

While many storage companies are encouraging their customers to bring in bigger and better boxes, another strategy for coping with the exploding data storage demand calls for making storage smarter.

One example of this strategy is e-Space, a storage management solution from Santa Clara, Calif.-based SolutionSoft Systems, Inc. The e-Space solution combines compression technology with the time-based controls of hierarchical storage management to automatically compress unused, but still important, files to get the most out of storage systems without requiring human intervention.

In the e-Space strategy, data can go through a cycle of compression and decompression to more efficiently use disk space.

Solution-Soft is targeting the software-at-large, Fortune 1000 companies and industries like telecommunications, health care, and insurance that require data-intensive applications. Another major target area is the e-service community of application service providers (ASPs), portals, and e-mail services.



"There is an explosion of data today, and all that data is important for things like trend analysis and understanding customer behavior," said Paul Wang, Solution-Soft's CEO. "With the advent of e-business, the requirements to have this data available online have increased greatly. Because of this, the total cost of ownership has gone up even as the cost of hardware has gone down. For every dollar spent on hardware, a company must spend \$5 to \$7 to manage that software, and that's a recurring cost you must pay every year."

Accessible Archives

For that reason, the number of online storage devices being managed needs to be reduced at the same time that the amount of readily available online data is increasing. The e-Space approach uses compression to give storage devices the ability to double as archival storage, making management easier while keeping older data accessible.

The system allows users to establish parameters for when files are compressed, much as they would in a traditional hierarchical storage management solution. The default setting is 30 days, but users can assign different thresholds according to the type or size of the file.

"This allows the users to juggle storage space by compressing more files with more aggressive storage policies," said Wang.

When the assigned threshold is reached, the software automatically chooses the best compression algorithm for each file and compresses the file by an average of 70 percent, storing it on the same device it originally resided on. Compression occurs at a rate of 1 to 3 megabytes per second. The space saved by compression is then released back to the storage system to

accommodate new data. When the data is needed again, the file is decompressed at a rate of 2 to 6 megabytes per second and remains in its decompressed form until it again reaches the time threshold set for compression.

Show Me

The software also includes a feature that allows managers to "audition" proposed changes to compression policies by previewing them and analyzing their impact over time. "This allows managers to fine-tune their software to get the biggest bang from their equipment," Wang said.

Part of that bang is the standardization of storage devices. Without the requirement for archival storage, storage management can be simplified and hardware failures can be reduced through the use of fewer devices. Out-of-disk-space errors can be minimized, and management costs can be reduced.

The compressed files also take less time to back up, said Wang, an important consideration as volumes of data in readily accessible formats increase exponentially. Compressed files also will make disaster recovery faster.

The software will be supported by "all major UNIX and Windows NT platforms," said Wang. Support is currently available for HP MPE, HP-UX, and Sun Solaris, with IBM AIX and Windows NT versions in beta testing. The software will work with "any hardware the operating system can recognize," Wang said, and it can be integrated with HP OpenView. Pricing for the e-Space solution starts at \$249 and is based on the user's server configuration.

Files Keep All Attributes

Compression technology has been employed in other solutions, notably as part of the POSIX interface included with MPE/iX 5.0. The difference the e-Space solution provides is the ability to compress files "in place," meaning that the compressed file retains all the file and security attributes of the uncompressed file, including the file name and the last access date. Applications that periodically access older files are able to call up these compressed files and utilize them without requiring modifications, and users will be able to locate compressed files just by searching for the original name.

Originally appeared in HP World Magazine, February 2000 issue