Data Reduction Technologies Provide a Highly Effective Backup and Recovery Strategy

By Pat Hanavan

Businesses are confronting the growth of data with a new face as they strive to effectively protect and manage their critical information. Added challenges around storage capacity, power and cooling are direct effects from data's rapid growth within organizations. Storing the surplus backup data is not only time-consuming and costly, but also a highly ineffective use of valuable and scarce storage resources.

Data deduplication directly addresses these issues, reducing total backup storage needs by 20x to 50x compared to traditional tape-based backup. Rather than storing everything with every backup, data deduplication stores a unique file shuck or records once and creates a pointer to that information, reducing the amount of storage needed for backups and overall data center energy use.

However, not all organizations have data deduplication capabilities included in their backup methodologies. In these cases, many organizations can leverage their backup tool's data reduction technologies to enjoy similar benefits. The difference? Data deduplication achieves high reduction rates compared to traditional full/incremental backup methodologies which generate masses of redundant backup data. Data reduction technologies reduce the amount of data being backed up in the beginning, decreasing the amount of backup data being stored in the end. For data-rich application environments such as Microsoft Exchange, data reduction technologies can cut backup storage in half.

Database Backup

Email, the preferred form of communication for most business transactions, has become mission-critical to organizations of all sizes. Consequently, the amount of data generated through email is growing exponentially. IT departments must be able to back up and store this data and then quickly recover it as needed. Traditional backup approaches entails running at least two separate backups of identical information, resulting in redundant backup storage.

However, database backup is mandatory since the only way to retrieve all Exchange Server data in the event of a disaster is to restore the databases. But administrators must also be able to recover at a more granular level, including whole mailboxes, a single email message, a calendar item, a contact, or a public or private folder. If administrators want to quickly recover individual emails, folders or mailboxes, they typically have to

run a separate backup of the specific parts of Exchange which they may wish to recover later. These granular backups, or “mailbox” or “brick-level” backups as they are commonly called, are dramatically slower to back up but easier to restore than database backups. Due to the sluggish backup rates, IT administrators seldom backup all of the individual Exchange mailboxes and folders.

Granular Recovery

Data reduction technologies enable IT to completely eliminate Exchange mailbox backups and still restore individual messages or mailboxes in seconds. More specifically, granular recovery technology provides the power and flexibility to recover any part of Exchange from the database backups themselves, making it one of the most effective data reduction technologies.

With granular recovery technology, the backup tool assembles additional information during the backup and places it in its catalogs. This extra information enables the recovery of single mail messages or folders from Exchange database backups sent to disk and/or tape without requiring separate mailbox level backups.

Granular recovery technology-enabled backups can recover Exchange mailboxes, messages, folders, individual calendar items or tasks since all of this information is contained within the Exchange databases.

Backup tools that include this technology feature make enabling granular recovery simple. An IT administrator simply checks an option to enable granular recovery and then chooses which Exchange storage groups to back up. By eliminating the redundant mailbox backups, granular recovery technology is virtually an effortless strategy for capitalizing on precious backup time and existing storage resources.

Continuous Data Protection

Another component of various backup tools is continuous data protection technology, which also provides benefits of granular recovery technology. With this component, granular recovery-enabled backups run frequently to ensure that Exchange recoveries are possible at any time.

Continuous data protection technology captures all Exchange server updates as they occur. A full backup is completed once a week or once a month, and between those full backups the continuous data protection technology keeps Exchange transaction logs continuously backed up. These backup logs are consolidated into easily managed recovery points so that Exchange databases are protected up to the latest complete transaction log.

Using continuous data protection technology, IT administrators can enable their backup tool to make recovery points at intervals they specify. Recovery points create backup sets that can be browsed, allowing for easy recovery of individual messages or folders from a point in time when either a full backup or recovery point was run. Each time a recovery point is created, it truncates the Exchange transaction log so that log growth is controlled.

In addition, recovery points are virtualized to save the disk space typically required by new complete backups of Exchange when a new recovery point is created. The actual disk space consumed by an individual recovery point is only the size of the transaction logs that have been continuously protected, combined with the associated last full backup.

The combination of data reduction technologies, such as granular recovery technology and continuous data protection, can significantly reduce storage considered necessary for backing up Exchange databases and mailboxes. What's more, these technologies can be used to reduce backup storage and energy consumption across other challenging platforms in the data center, including Microsoft Active Directory and Microsoft SharePoint environments. Granular recovery technology can also be applied to database backups of Active Directory and SharePoint.

Data reduction technologies provide an intelligent, flexible solution for addressing today's proximate and pressing needs. They provide solutions for maximizing existing resources, eliminating redundancies, and reducing energy consumption. Having emerged as a critical component of an efficient and effective backup and recovery strategy, data reduction technologies will help organizations as they face business challenges that lie ahead.

Pat Hanavan is the vice president of product management for the Backup Exec product family within Symantec's Data Protection Group. In this role, Hanavan is responsible for the overall product direction, strategies and business performance for Backup Exec and Backup Exec System Recovery. These market leading products ensure successful backup and recovery of information and systems for organizations worldwide.

For more information, visit www.symantec.com.