


# A LOOK INSIDE MICROSOFT SYSTEM CENTER DATA PROTECTION MANAGER 2007

Microsoft® System Center Data Protection Manager (DPM) 2007 is designed to combine the advantages of traditional tape backup and disk-based continuous data protection processes in a unified platform. This article delves into the key functionality of DPM 2007 and how it can help streamline and enhance data protection processes.



By Jason Buffington  
Sanjeet Singh

## Related Categories:

Business continuity

Data security

Dell PowerEdge servers

Dell PowerVault storage

Microsoft System Center

Storage

Visit [DELL.COM/PowerSolutions](http://DELL.COM/PowerSolutions)  
for the complete category index.

For years, the basic minimum standard in traditional data protection has been nightly tape backup. The standard approach has changed very little: once a week, production server data is streamed to tape media during nonpeak hours, a process referred to as a full backup. Then, each night, backup agents identify files or other data objects that have changed and send those to tape as well, in the form of either incremental backups (collective changes in the data since the last full backup) or differential backups (changes in the data each day). Incremental backups typically allow faster recoveries than differential backups, while differential backups typically allow faster nightly backup processes than incremental backups.

Although tape hardware technology has advanced, the disadvantages of nightly tape backup have remained the same, including the potential to lose a full day's worth of data, lengthy recovery times, and poor performance when restoring individual files or other data objects—which can constitute a significant percentage of recovery operations. Although disk-based approaches such as virtual tape libraries (VTLs), real-time replication (RTR), and continuous data protection (CDP) are designed to avoid or

mitigate some of these problems, they also come with their own disadvantages. VTLs use the same approach as traditional tape backup, simply substituting disk media for tape media. RTR maintains a near-current copy of production data but does not typically provide multiple recovery points, and therefore is more suitable for high availability than data protection. And CDP is often designed specifically for individual applications, requiring enterprises to purchase and deploy multiple systems that provide differing capabilities and may not be compatible with one another.

Microsoft System Center Data Protection Manager (DPM) 2007 is designed to combine the advantages of traditional tape backup and disk-based CDP processes using an integrated disk-to-disk-to-tape system that is optimized for Microsoft workloads such as Microsoft Exchange, Microsoft SQL Server™, Microsoft Office SharePoint® Server, and Microsoft Virtual Server software, as well as file services in the Microsoft Windows Server® 2003, Windows Server 2008 (code-named “Longhorn”), Windows® XP, and Windows Vista® operating systems. Utilizing key DPM 2007 functionality can help administrators streamline and enhance their data protection processes while increasing flexibility and simplifying backup management.<sup>1</sup>

<sup>1</sup> For a more detailed discussion of traditional tape backup; disk-based VTL, RTR, and CDP approaches; and how Dell and Microsoft are working together to create DPM 2007–based data protection systems, see “Safeguarding Data with Dell PowerVault Data Protection Solutions,” by Sanjeet Singh and Jason Buffington, in *Dell Power Solutions*, February 2008, [DELL.COM/Downloads/Global/Power/ps1q08-20070484-Singh.pdf](http://DELL.COM/Downloads/Global/Power/ps1q08-20070484-Singh.pdf).

## UNDERSTANDING MICROSOFT SYSTEM CENTER DATA PROTECTION MANAGER 2007

DPM 2007 incorporates two key data protection features—express full backups and transaction replication. The express full feature provides block-level synchronization utilizing Microsoft Volume Shadow Copy Service (VSS) along with the built-in VSS writers for supported Microsoft server applications and the VSS writer in the native file systems for Windows Server 2003, Windows Server 2008, Windows XP, and Windows Vista. The transaction replication feature uses application-aware agents that can protect Microsoft Exchange and SQL Server data through their transaction logs.

### Express full backups

The express full feature creates a bitmask of the physical volumes that are being protected and is designed to work across different data sources. When administrators select a data source to protect, such as a SQL Server database, DPM 2007 identifies the files that make up the database and which disk blocks compose those files. The DPM bitmask table essentially contains a 0 or 1 for each block that composes a file or other data object, with 0 denoting a block that has not changed and 1 denoting a block that has changed. As data changes during the course of a day or week and the blocks are updated, the DPM agent changes the bits for those blocks from 0 to 1.

When a scheduled express full task executes, the DPM agent invokes a VSS snapshot of the data area, including the Windows file system and any source-specific VSS writers (such as Exchange, SQL Server, SharePoint, or Virtual Server). Within each application that uses VSS, the VSS writer helps ensure that the data is consistent and ready to be backed up. In addition, VSS provides the ability to have a “frozen” set of disk blocks (a shadow copy) while the production disk continues to service active I/O. The DPM agent then fetches only the changed

blocks from within the shadow copy, while the production processes continue to function.

Rather than using approaches such as RTR (which is usually application agnostic and may not ensure write-order integrity) or CDP (which is often designed for a specific application), DPM uses the underlying VSS capabilities in Windows file systems and Microsoft applications to provide data-consistent protection capabilities across different data sources. After invoking the VSS snapshot, DPM can transmit the changed blocks while the production file system continues to operate. When DPM has synchronized all the marked blocks with the DPM server, the snapshot is released.

Microsoft terms these *express full backups*. Administrators should note that this option is not the same as a traditional full tape backup. When the backup process is complete, DPM does have a near-exact copy of the production data as it was when the process began—effectively a full backup. However, unlike a traditional full tape backup, DPM tracks and sends only changed blocks, ignoring unchanged blocks and stagnant files. The advantage of this approach is significantly reduced bandwidth requirements, which helps make DPM particularly useful for centralized backup of data from branch offices and other sources across a wide area network (WAN).

The result is a network-optimized synchronization of production data to the DPM server. DPM can then facilitate snapshots of the replicated area to provide up to 512 recovery points by storing only the block-level changes between each

iteration. DPM express full backups would typically occur as a nightly scheduled task, but can occur as often as every half hour or as seldom as once per week.

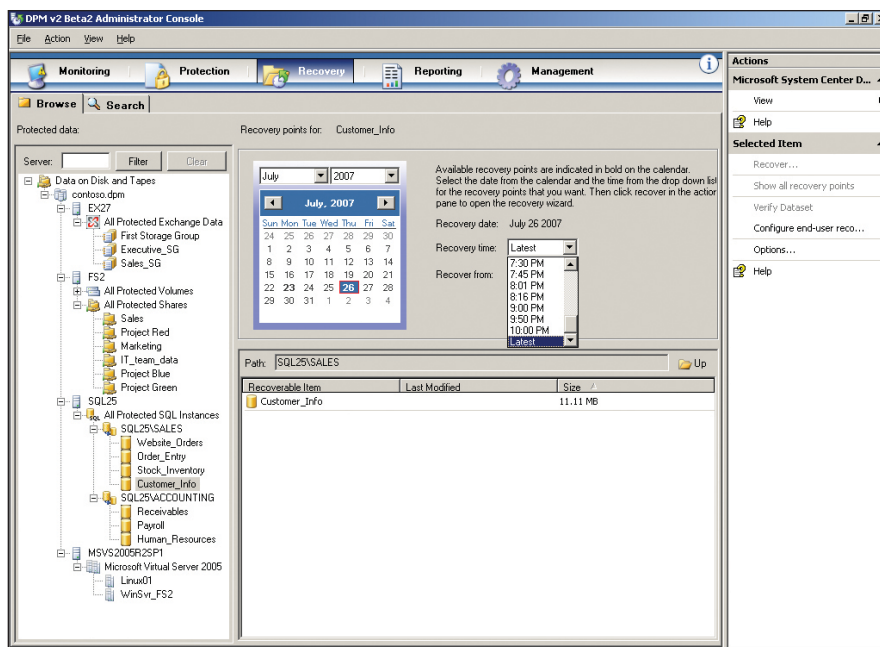
### Transaction replication

Transaction replication uses DPM agents to secure any application transaction logs that have closed since the last time window (as often as every 15 minutes) and transmit them to the DPM server. This feature sounds similar to log shipping, in which the secondary server applies logs automatically and repeatedly to maintain a near-current copy of the data, with the goal of increasing availability. However, because DPM is designed to increase data protection rather than availability, the DPM server retains each set of transaction logs as a differential backup. And just as a traditional tape backup might use its latest full backup plus some set of differentials to restore data to a previous point in time, DPM can use the live replica data plus the differentials to restore a production server to any previous point in time that the DPM server has transaction logs for.

## INTEGRATING REPLICATION FEATURES

Combining the express full and transaction replication features can provide significant administrative flexibility. For example, because a DPM 2007 server can maintain up to 512 VSS snapshots, performing one express full backup each week in addition to synchronizing transactions every 15 minutes could result in an Exchange or SQL Server system having more than 344,000 total recovery

**“DPM 2007 is designed to combine the advantages of traditional tape backup and disk-based CDP processes using an integrated disk-to-disk-to-tape system optimized for Microsoft workloads.”**



**Figure 1.** Simplified recovery in the Microsoft System Center Data Protection Manager 2007 Administrator Console

points. Performing one express full backup each day in addition to capturing transaction replication differentials every 15 minutes could result in nearly 50,000 recovery points, and potentially allow administrators to recover data to any 15-minute point in time over the past year and a half. Realistically, organizations would normally have 30–90 days of data on disk-based backup, including 15-minute transactional differentials, and then offload this data to tape for long-term retention.

Although this potentially large number of recovery points may seem imposing, the DPM Administrator Console is designed to simplify data recovery (see Figure 1). For example, if a database server's hard drive fails at 11:58, the DPM server could potentially recover the data to 11:45, 11:30, 11:15, and so on. However, the DPM Administrator Console also allows administrators to simply choose "Latest" for the latest available recovery point. DPM could then first recover the server to 11:45. Then—provided administrators have followed the best practice of maintaining the databases on one volume and the transaction logs on a different volume—it can automatically play

forward any surviving transaction logs from 11:46 up through 11:58. When this process is complete, the database could be within a single transaction of its state when the hard drive failed.

Using the DPM 2007 express full and transaction replication features together offers multiple advantages. Instead of needing one tool for nightly tape backup and then one or more "better than nightly" disk-to-disk technologies, DPM enables a single platform to first protect data on disk and then offload that data to tape. This approach helps simplify management by using a single agent on the production server, a unified replication engine from production server to protection server, and a consistent interface regardless of media type. In addition, rather than requiring different application-specific agents and protection features, DPM provides a unified system to protect Microsoft Exchange, SQL Server, SharePoint, and Virtual Server data—meaning that if administrators repurpose a SQL Server system to run Virtual Server, they would not have to waste a SQL Server-specific backup agent or purchase an additional Virtual Server-specific backup agent.

## PROTECTING CRITICAL DATA FOR KEY MICROSOFT WORKLOADS

Dell and Microsoft are working together to create unified disk-to-disk-to-tape Data Protection Solutions based on Dell™ PowerVault™ hardware and Microsoft System Center Data Protection Manager 2007 software. Designed to be cost-effective and easy to deploy and manage, these systems can provide flexible, simplified data protection in environments built on Dell hardware and Microsoft software. 

**Jason Buffington** is the senior technical product manager for Data Protection Manager and other storage solutions at Microsoft. He has more than 16 years of experience in data protection and business continuity, and studied Computer Science at Texas A&M University.

**Sanjeet Singh** is a global product marketing manager in the Dell Enterprise Storage Group. He has eight years of experience in developing and delivering business-critical technologies, including databases and data protection. Sanjeet has an M.S. in Computer Engineering from Purdue University and an M.B.A. from the University of Texas.

**MORE  
ONLINE**  
DELL.COM/PowerSolutions

### QUICK LINKS

Singh, Sanjeet, and Jason Buffington. "Safeguarding Data with Dell PowerVault Data Protection Solutions." *Dell Power Solutions*, February 2008. [DELL.COM/Downloads/Global/Power/ps1q08-20070484-Singh.pdf](http://DELL.COM/Downloads/Global/Power/ps1q08-20070484-Singh.pdf)

**Dell storage:**  
[DELL.COM/Storage](http://DELL.COM/Storage)

**Microsoft System Center Data Protection Manager 2007:**  
[www.microsoft.com/dpm](http://www.microsoft.com/dpm)