Planning a System Center Data Protection Manager 2007 Deployment

Microsoft Corporation

Published: Sep 2007

Abstract

This content explains how DPM works and provides guidance for planning a DPM deployment.

The information contained in this document represents the current view of Microsoft Corporation on the issues discussed as of the date of publication. Because Microsoft must respond to changing market conditions, it should not be interpreted to be a commitment on the part of Microsoft, and Microsoft cannot guarantee the accuracy of any information presented after the date of publication.

This White Paper is for informational purposes only. MICROSOFT MAKES NO WARRANTIES, EXPRESS, IMPLIED OR STATUTORY, AS TO THE INFORMATION IN THIS DOCUMENT.

Complying with all applicable copyright laws is the responsibility of the user. Without limiting the rights under copyright, no part of this document may be reproduced, stored in or introduced into a retrieval system, or transmitted in any form or by any means (electronic, mechanical, photocopying, recording, or otherwise), or for any purpose, without the express written permission of Microsoft Corporation.

Microsoft may have patents, patent applications, trademarks, copyrights, or other intellectual property rights covering subject matter in this document. Except as expressly provided in any written license agreement from Microsoft, the furnishing of this document does not give you any license to these patents, trademarks, copyrights, or other intellectual property.

Contents

Planning a DPM 2007 Deployment

This content explains how DPM works and provides guidance for planning a DPM deployment.

In This Section

Introducing Data Protection Manager 2007

Planning Protection Groups

Planning for DPM Deployment

Deployment Plan Checklist and Roadmap

Introducing Data Protection Manager 2007

Microsoft System Center Data Protection Manager (DPM) 2007 is a key member of the Microsoft System Center family of management products, designed to help IT professionals manage their Windows environment. DPM is the new standard for Windows backup and recovery—delivering seamless data protection for Microsoft application and file servers by using integrated disk and tape media.

In This Section

DPM Features

How DPM Works

System Requirements

DPM Licensing

DPM Features

Data protection is essential to a business or organization, and Microsoft System Center Data Protection Manager (DPM) 2007 is an effective solution for providing that protection. DPM provides the following benefits for your organization:

- Disk-based data protection and recovery.
- Tape-based backup and archive solutions.
- Disaster recovery solutions.

You can back up the DPM database to tape, or you can use a second DPM server in a geographically separated location to protect the primary DPM server.

If you use a second DPM server, you can restore data to protected computers directly from the secondary DPM server. The secondary DPM server can also protect computers until the primary DPM server is brought back online.

- DPM provides protection of the following items:
  - File data from volumes, shares, and folders.
  - Application data, such as Microsoft Exchange Server storage groups, Microsoft SQL Server databases, Windows SharePoint Services farms, and Microsoft Virtual Server and its virtual machines.
  - Files for workstations running Windows XP Professional SP2 and all Windows Vista editions except Home.
  - Files and application data on clustered servers.
  - System state for protected file and application servers.

In This Section

Backup Solutions Combining Disk and Tape

Protection for Multiple Data Types

Protection for Clustered Servers

Management Tools

See Also

How DPM Works

Backup Solutions Combining Disk and Tape

With DPM data protection, you can use disk-based storage, tape-based storage, or both.

Disk-based storage, also called *D2D,* for "disk-to-disk," is a type of backup in which data from one computer is stored on the hard disk of another computer. This contrasts with the more traditional method of backing up data from one computer to a storage media such as tape, also called *D2T,* for "disk-to-tape." For extra protection, the two methods can be combined in a disk-to-disk-to-tape (*D2D2T*) configuration that provides the rapid recovery benefits of disk-based storage in the short term and tape-based, archive storage for critical data in the long term. The following illustration shows the three storage methods.

Data Storage Methods

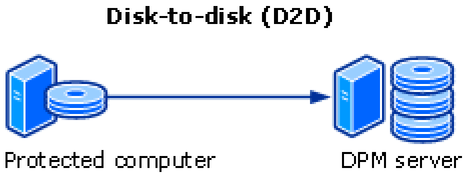

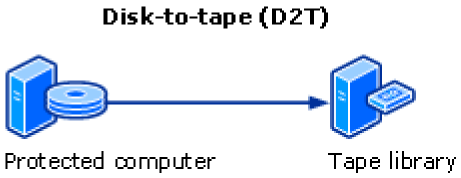

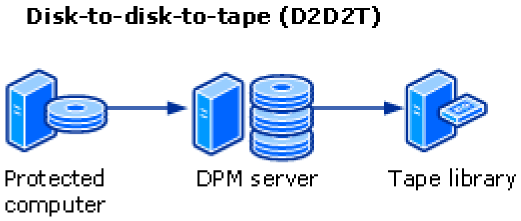

To determine which storage method to use, you must consider the relative importance of your organization's protection requirements.

- **How much data your organization can afford to lose.** Realistically, not all data is equally valuable. Organizations must weigh the impact of loss against the costs of protection.
- **How quickly recovered data must be available.** Recovery of data that is critical to ongoing operations is typically more urgent than routine data. On the other hand, organizations should identify servers providing essential services during working hours that must not be disrupted by recovery operations.
- **How long your organization must maintain data.** Long-term storage might be necessary for business operations, depending on the type and contents of the data. An organization might also be subject to legal requirements for data retention, such as the Sarbanes-Oxley Act and the Data Retention Directive.
- **How much your organization can spend on data protection.** When considering how much to invest in data protection, organizations must include the cost of hardware and media, as well as the personnel costs for administration, management, and support.

You can use DPM to back up data to both disk and tape, giving you the flexibility to create focused, detailed backup strategies that result in efficient and economic data protection. When you need to restore a single file or an entire server, recovery is fast and simple: you identify the data, and DPM locates the data and retrieves it (although your assistance might be needed if the tape has been removed from the library).

Disk-Based Protection and Recovery

One advantage of disk-based data protection is the potential time savings. Disk-based data protection requires none of the preparation time that tape-based protection does—locating the specific tape required for a job, loading the tape, positioning the tape to the correct starting point. The ease of using a disk encourages sending incremental data more frequently, which reduces the impact on the computer being protected and on network resources.

Data recovery with disk-based data protection is more reliable than that of tape-based systems. Disk drives typically have a much greater mean time between failure (MTBF) rating than tapes.

Recovery of data from disk is quicker and easier than recovery from tape. Recovering data from disk is a simple matter of browsing through previous versions of the data on the DPM server and copying selected versions directly to the protected computer. A typical file recovery from tape takes hours and can be costly, and administrators in a medium-size data center can usually expect to perform 10 to 20 or more of these recoveries each month.

Using DPM and disk-based data protection, data can be synchronized as frequently as every 15 minutes and maintained as long as 448 days.

Tape-Based Backup and Archive

Magnetic tape and similar storage media offer an inexpensive and portable form of data protection that is particularly useful for long-term storage.

In DPM, you can back up data from a computer directly to tape (D2T). You can also back up data from the disk-based replica (D2D2T). The advantage of creating your long-term backup on tape from the disk-based replica is that the backup operation can occur at any time with no impact on the computer being protected.

Additionally, a thorough disaster recovery plan includes offsite storage of critical information—you want to be able to recover your organization's data, should your facility be damaged or destroyed. Tape is a popular and convenient medium for offsite storage.

Using DPM, data can be backed up to tape as frequently as daily for short-term protection, and it can be maintained as long as 99 years for long-term protection.

By using software solutions from DPM partners, you can use removable media such as a USB hard drive in place of tape. For more information, see Data Protection Manager Partners (http://go.microsoft.com/fwlink/?LinkId=98869).

See Also

Management Tools

Protection for Clustered Servers

Protection for Multiple Data Types

Protection for Multiple Data Types

The following table lists the types of data that DPM can protect and the level of data that you can recover by using DPM.

Note

For information about the specific software requirements for protected computers, see DPM System Requirements (http://go.microsoft.com/fwlink/?LinkId=66731).

Protectable and Recoverable Data

Product	Protectable Data	Recoverable Data
Exchange Server 2003 Exchange Server 2007	• Storage group	• Storage group • Database • Mailbox
SQL Server 2000 SQL Server 2005	• Database	• Database
Microsoft Office SharePoint Server 2007 Windows SharePoint Services 3.0	• Farm	• Farm • Database • Site • File or list
Windows Server 2003 Windows Storage Server 2003	• Volume • Share • Folder	• Volume • Share • Folder • File
Microsoft Virtual Server 2005 R2 SP1	• Virtual server host configuration • Virtual machines • Data for applications running on virtual machines[1]	• Virtual server host configuration • Virtual machines • Data for applications running on virtual machines[1]
All computers that can be protected by DPM 2007 except computers running Windows Vista or Windows Server 2008	• System state	• System state
Workstations running Windows XP Professional SP2 and all Windows Vista editions	• File data	• File data

Product	Protectable Data	Recoverable Data
except Home (must be member of a domain)		

[1] Data for applications running in virtual machines must be protected and recovered as an application data source, not as a component of a protected virtual machine. For example, to protect and recover data for an instance of SQL Server running on a virtual machine, you install the DPM protection agent on the virtual machine and select the data source as a SQL Server database. When you install the protection agent on the virtual host and protect a virtual machine on the host, application data is also protected but can be recovered only by recovering the virtual machine itself.

See Also

Managing Protected File Servers and Workstations

Managing Protected Servers Running Exchange

Managing Protected Servers Running SQL Server

Managing Protected Servers Running Windows SharePoint Services

Managing Protected Virtual Servers

Protection for Clustered Servers

DPM 2007 supports shared disk clusters for file servers, Exchange Server 2003, SQL Server 2000, and SQL Server 2005. DPM 2007 supports both non-shared disk clusters and shared disk clusters for Exchange Server 2007.

For DPM protection agent installation, when you select a server that is a cluster node, DPM notifies you so that you can choose to install the protection agent on other nodes in the cluster as well.

End-user recovery is available for both clustered and nonclustered resources on clustered file servers.

On planned failover, DPM continues protection. On unplanned failover, DPM issues an alert that a consistency check is required.

See Also

Protection for Multiple Data Types

Management Tools

To facilitate the performance of key management tasks, DPM 2007 provides the following tools and capabilities for IT administrators:

- DPM Administrator Console
- Reports and notifications
- DPM Management Packs
- Windows PowerShell integration
- Remote administration
- End-user recovery

DPM Administrator Console

DPM Administrator Console uses a task-based administration model that automates common tasks, enabling the administrator to get the job done with the fewest number of steps.

To simplify the management of data protection activities, DPM builds on Microsoft Management Console (MMC) functionality to provide a familiar, intuitive environment for performing configuration, management, and monitoring tasks.

DPM Administrator Console organizes tasks into five easily accessible task areas: monitoring, protection, recovery, reporting, and management. Wizards guide the administrator through basic configuration tasks such as adding disks, installing agents, and creating protection groups. Search and browse features are provided in the **Recovery** task area to assist in finding and recovering previous versions of files.

DPM Administrator Console provides both a **Jobs** tab and an **Alerts** tab for monitoring data protection activity. The **Jobs** tab provides the status and operational details for each scheduled, completed, running, canceled, or failed job. The **Alerts** tab aggregates informational alerts and error conditions to provide a summary view of activity for the entire system and provides recommended actions for each error.

For details about using DPM Administrator Console, see Appendix A: DPM Administrator Console (http://go.microsoft.com/fwlink/?LinkId=98871) in *Deploying DPM 2007*.

Reports and Notifications

DPM provides a comprehensive set of reports that provide data about protection success and failures, recovery success and failures, and disk and tape utilization. You can also identify common errors and manage circulation of tapes. Summary reports aggregate information for all protected computers and protection groups. Detailed reports provide information about individual computers or protection groups. An administrator can use these reports to fine-tune protection after the initial DPM deployment.

DPM notifications provide a convenient way to stay informed when critical, warning, or informational alerts are generated. You choose the severity of alert that you want to be notified about; for example, you can choose to receive only critical alerts. You can also choose to receive notifications of the status of recovery jobs, and you can have scheduled DPM reports delivered as e-mail attachments so that you can monitor data protection trends and analyze data protection statistics at your convenience. You can also use the DPM Management Pack for System Center Operations Manager 2007 to provide customized notifications.

For details about the reports available in DPM 2007, see Managing DPM Servers (http://go.microsoft.com/fwlink/?LinkId=91853). For instructions on subscribing to notifications, see DPM 2007 Help.

DPM Management Packs

Management Packs for Microsoft Operations Manager 2005 (MOM) and System Center Operations Manager 2007 will be available for DPM 2007. As part of your data management strategy, you can use the DPM Management Pack to centrally monitor data protection, state, health, and performance of multiple DPM servers, and the servers that they protect. From the Operations Manager Operations Console, an administrator can monitor DPM and network infrastructure simultaneously, analyzing issues with data protection in the context of other factors in system and network performance. The administrator also can monitor other mission-critical applications, such as SQL Server.

To download the DPM Management Packs, see the Management Pack Catalog (http://go.microsoft.com/fwlink/?LinkId=47215).

Windows PowerShell Integration

Windows PowerShell is an interactive command-line technology that also supports task-based scripting.

DPM provides its own set of Windows PowerShell commands that can be used for performing data protection management tasks. You access the DPM cmdlets through DPM Management Shell.

A DPM administrator can use DPM cmdlets to perform all the administrative tasks that can be performed in the console, including sets of cmdlets designed to be used for the following tasks:

- To configure DPM
- To manage tapes and disks
- To manage protection groups
- To protect and recover data

In addition, DPM cmdlets enable administrators to perform the following tasks, which cannot be performed by using DPM Administrator Console:

- To remove recovery points

- To customize the start time for library maintenance jobs, such as detailed inventory and cleaning
- To specify the local area network (LAN) configuration to be used for a backup job

Remote Administration

You can establish a Remote Desktop connection to a DPM server to manage DPM operations remotely.

DPM Management Shell can be installed on computers other than the DPM server, enabling you to administer multiple DPM servers remotely. You can even install DPM Management Shell on desktop computers running Windows XP or Windows Vista.

End-User Recovery

In addition to administrator-provided data recovery, DPM enables users to independently retrieve previous versions of their files by using the familiar Windows Explorer interface or any of the Microsoft Office 2007 applications. End-user recovery is not available for application data.

See Also

Protection for Clustered Servers

Protection for Multiple Data Types

How DPM Works

The method that Data Protection Manager uses to protect data varies according to the type of data being protected and the method of protection you select.

In This Section

Disk-Based Protection Process

Tape-Based Protection Process

Recovery Process

Protection Policy

Auto Discovery Process

DPM Directory Structure

Disk-Based Protection Process

To provide disk-based data protection, the DPM server creates and maintains a *replica*, or copy, of the data that is on protected servers. The replicas are stored in the *storage pool* which consists of a set of disks on the DPM server, or on a custom volume. The following illustration shows the basic relationship between a protected volume and its replica.

Replica Creation

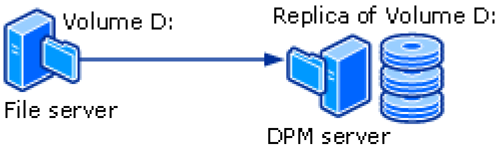

Whether you are protecting file data or application data, protection begins with the creation of the replica of the data source.

The replica is *synchronized*, or updated, at regular intervals according to the settings that you configure. The method that DPM uses to synchronize the replica depends on the type of data being protected. For more information, see The File Data Synchronization Process and The Application Data Synchronization Process. If a replica is identified as being inconsistent, DPM performs a consistency check, which is a block-by-block verification of the replica against the data source.

A simple example of a protection configuration consists of a DPM server and a protected computer. The computer is protected when you install a DPM *protection agent* on the computer and add its data to a *protection group*.

Protection agents track changes to protected data and transfer the changes to the DPM server. The protection agent also identifies data on a computer that can be protected and is involved in the recovery process. You must install a protection agent on each computer that you want to protect by using DPM. Protection agents can be installed by DPM or you can install protection agents manually using applications such as Systems Management Server (SMS).

Protection groups are used to manage the protection of data sources on computers. A protection group is a collection of data sources that share the same protection configuration. The protection configuration is the collection of settings that are common to a protection group, such as the protection group name, protection policy, disk allocations, and replica creation method.

DPM stores a separate replica for each *protection group member* in the storage pool. A protection group member can be any of the following data sources:

- A volume, share, or folder on a desktop computer, file server, or server cluster
- A storage group on an Exchange server or server cluster
- A database of an instance of SQL Server or server cluster

See Also

The Application Data Synchronization Process

[The Difference Between File Data and Application Data](#)

[The File Data Synchronization Process](#)

The File Data Synchronization Process

In DPM 2007, for a file volume or share on a server, the protection agent uses a volume filter and the change journal to determine which files have changed and then performs a checksum procedure for these files to synchronize only the changed blocks. During synchronization, these changes are transferred to the DPM server and then applied to the replica to synchronize the replica with the data source. The following figure illustrates the file synchronization process.

File Synchronization Process

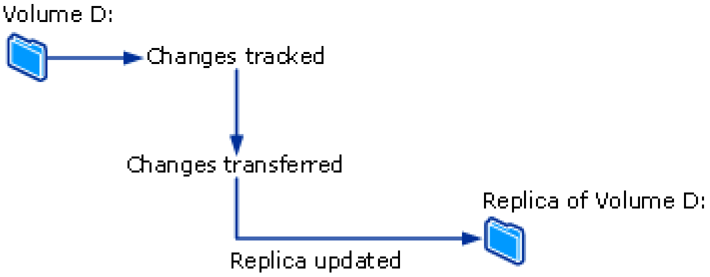

If a replica becomes inconsistent with its data source, DPM generates an alert that specifies which computer and which data sources are affected. To resolve the problem, the administrator repairs the replica by initiating a *synchronization with consistency check*, also known as simply a *consistency check*, on the replica. During a consistency check, DPM performs a block-by-block verification and repairs the replica to bring it back into consistency with the data source.

You can schedule a daily consistency check for protection groups or initiate a consistency check manually.

At regular intervals that you can configure, DPM creates a *recovery point* for the protection group member. A recovery point is a version of the data from which data can be recovered. For files, a recovery point consists of a shadow copy of the replica, which is created by using the Volume Shadow Copy Service (VSS) functionality of the operating system on the DPM server.

See Also

[The Application Data Synchronization Process](#)

[The Difference Between File Data and Application Data](#)

[Disk-Based Protection Process](#)

The Application Data Synchronization Process

For application data, after the replica is created by DPM, changes to volume blocks that belong to application files are tracked by the volume filter.

How changes are transferred to the DPM server depends on the application and the type of synchronization. The operation that is labeled *synchronization* in DPM Administrator Console is analogous to an incremental backup, and it creates an accurate reflection of the application data when combined with the replica.

During the type of synchronization that is labeled *express full backup* in DPM Administrator Console, a full Volume Shadow Copy Service (VSS) snapshot is created but only changed blocks are transferred to the DPM server.

Each express full backup creates a recovery point for application data. If the application supports incremental backups, each synchronization also creates a recovery point. The synchronization type supported by each type of application data is summarized as follows:

- For protected Exchange data, synchronization transfers an incremental VSS snapshot using the Exchange VSS writer. Recovery points are created for each synchronization and express full backup.
- SQL Server databases that are log-shipped, in read-only mode, or that use the simple recovery model do not support incremental backup. Recovery points are created for each express full backup only. For all other SQL Server databases, synchronization transfers a transaction log backup, and recovery points are created for each incremental synchronization and express full backup. The transaction log is a serial record of all the transactions that have been performed against the database since the transaction log was last backed up.
- Windows SharePoint Services and Microsoft Virtual Server do not support incremental backup. Recovery points are created for each express full backup only.

Incremental synchronizations require less time than performing an express full backup. However, the time required to recover data increases as the number of synchronizations increases. This is because DPM must restore the last full backup and then restore and apply all the incremental synchronizations up to the point in time selected for recovery.

To enable faster recovery time, DPM regularly performs an express full backup, a type of synchronization that updates the replica to include the changed blocks.

During the express full backup, DPM takes a snapshot of the replica before updating the replica with the changed blocks. To enable more frequent recovery point objectives, as well as to reduce the data loss window, DPM also performs incremental synchronizations in the time between two express full backups.

As with the protection of file data, if a replica becomes inconsistent with its data source, DPM generates an alert that specifies which server and which data source are affected. To resolve the problem, the administrator repairs the replica by initiating a synchronization with consistency check on the replica. During a consistency check, DPM performs a block-by-block verification and repairs the replica to bring it back into consistency with the data sources.

You can schedule a daily consistency check for protection groups or initiate a consistency check manually.

See Also

The Difference Between File Data and Application Data

Disk-Based Protection Process

The File Data Synchronization Process

The Difference Between File Data and Application Data

Data that exists on a file server and which needs to be protected as a flat file qualifies as file data, such as Microsoft Office files, text files, batch files, and so forth.

Data that exists on an application server and which requires DPM to be aware of the application qualifies as application data, such as Exchange storage groups, SQL Server databases, Windows SharePoint Services farms, and Virtual Server.

Each data source is presented in DPM Administrator Console according to the type of protection that you can select for that data source. For example, in the Create New Protection Group Wizard, when you expand a server that contains files and is also running Virtual Server and an instance of SQL Server, the data sources are treated as follows:

- If you expand **All Shares** or **All Volumes**, DPM displays the shares and volumes on that server and will protect any data source selected in either of those nodes as file data.
- If you expand **All SQL Servers**, DPM displays the instances of SQL Server on that server and will protect any data source selected in that node as application data.
- If you expand **Microsoft Virtual Server**, DPM displays the host database and virtual machines on that server and will protect any data source selected in that node as application data.

See Also

The Application Data Synchronization Process

Disk-Based Protection Process

The File Data Synchronization Process

Tape-Based Protection Process

When you use short-term disk-based protection and long-term tape-based protection, DPM can back up data from the replica volume to tape so that there is no impact on the protected computer. When you use tape-based protection only, DPM backs up the data directly from the protected computer to tape.

DPM protects data on tape through a combination of full and incremental backups from either the protected data source (for short-term protection on tape or for long-term protection on tape when DPM does not protect the data on disk) or from the DPM replica (for long-term protection on tape when short-term protection is on disk).

Note

> If a file was open when the replica was last synchronized, the backup of that file from the replica will be in a *crash consistent state*. A crash consistent state of the file will contain all data of the file that was persisted to disk at the time of last synchronization. This applies only to file system backups. Application backups will always be consistent with the application state.

For specific backup types and schedules, see Planning Protection Groups.

See Also

How DPM Works

Disk-Based Protection Process

Recovery Process

The method of data protection, disk-based or tape-based, makes no difference to the recovery task. You select the recovery point of data that you want to recover, and DPM recovers the data to the protected computer.

DPM can store a maximum of 64 recovery points for each file member of a protection group. For application data sources, DPM can store up to 448 express full backups and up to 96 incremental backups for each express full backup. When storage area limits have been reached and the retention range for the existing recovery points is not met yet, protection jobs will fail.

Note

> To support end-user recovery, the recovery points for files are limited to 64 by Volume Shadow Copy Service (VSS).

As explained in The File Data Synchronization Process and The Application Data Synchronization Process, the process for creating recovery points differs between file data and application data. DPM creates recovery points for file data by taking a shadow copy of the replica on a schedule that you configure. For application data, each synchronization and express full backup creates a recovery point.

The following illustration shows how each protection group member is associated with its own replica volume and recovery point volume.

Protection Group Members, Replicas, and Recovery Points

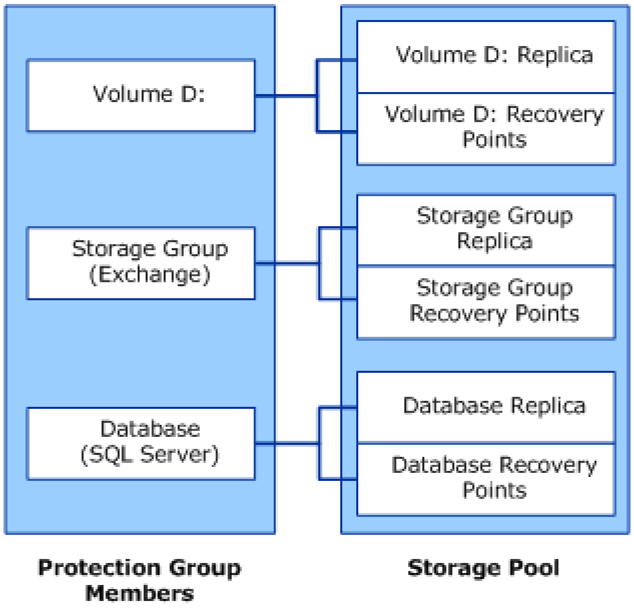

Administrators recover data from available recovery points by using the Recovery Wizard in DPM Administrator Console. When you select a data source and point in time from which to recover, DPM notifies you if the data is on tape, whether the tape is online or offline, and which tapes are needed to complete the recovery.

Users can recover previous versions of protected files. Because recovery points retain the folder and file structure of protected data sources, users browse through familiar volumes, folders, and shares to recover the data they want. End-user recovery is not available for application data such as an Exchange mailbox. Also, the versions of file data that are available for end-user recovery are those stored on disk in the DPM storage pool; file data that has been archived to tape can be recovered only by an administrator.

End users recover protected files by using a client computer that is running the shadow copy client software. Users can recover previous versions through shares on file servers, through Distributed File System (DFS) Namespaces, or by using a command on the **Tools** menu for Microsoft Office applications.

See Also

The Application Data Synchronization Process

The File Data Synchronization Process

Protection Policy

DPM configures the *protection policy*, or schedule of jobs, for each protection group based on the recovery goals that you specify for that protection group. Examples of recovery goals are as follows:

- “Lose no more than 1 hour of production data”
- “Provide me with a retention range of 30 days”
- “Make data available for recovery for 7 years”

Your *recovery goals* quantify your organization's data protection requirements. In DPM, the recovery goals are defined by retention range, data loss tolerance, recovery point schedule, and, for database applications, the express full backup schedule.

The *retention range* is how long you need the backed-up data available. For example, do you need data from today to be available a week from now? Two weeks from now? A year from now?

Data loss tolerance is the maximum amount of data loss, measured in time, that is acceptable to business requirements, and it will determine how often DPM should synchronize with the protected server by collecting data changes from the protected server. You can change the synchronization frequency to any interval between 15 minutes and 24 hours. You can also select to synchronize just before a recovery point is created, rather than on a specified time schedule.

The *recovery point schedule* establishes how many recovery points of this protection group should be created. For file protection, you select the days and times for which you want recovery points created. For data protection of applications that support incremental backups, the synchronization frequency determines the recovery point schedule. For data protection of applications that do not support incremental backups, the express full backup schedule determines the recovery point schedule.

Note

When you create a protection group, DPM identifies the type of data being protected and offers only the protection options available for the data.

See Also

How DPM Works

Auto Discovery Process

Auto discovery is the daily process by which DPM automatically detects new or removed computers on the network. Once a day, at a time that you can schedule, DPM sends a small packet (less than 10 kilobytes) to the closest domain controller. The domain controller responds to the LDAP request with the computers in that domain, and DPM identifies new and removed computers. The network traffic created by the auto discovery process is minimal.

Auto discovery does not discover new and removed computers in other domains. To install a protection agent on a computer in another domain, you must identify the computer by using its fully qualified domain name.

See Also

How DPM Works

DPM Directory Structure

When you begin protecting data with DPM, you will notice that the installation path of DPM contains three folders in the Volumes directory:

- \Microsoft DPM\DPM\Volumes\DiffArea
- \Microsoft DPM\DPM\Volumes\Replica
- \Microsoft DPM\DPM\Volumes\ShadowCopy

The DiffArea folder contains mounted shadow copy volumes that store the recovery points for a data source.

The Replica folder contains mounted replica volumes.

The ShadowCopy folder contains local backup copies of the DPM database. In addition, when you use DPMBackup.exe to create backup shadow copies of the replicas for archive by third-party backup software, the backup shadow copies are stored in the ShadowCopy folder.

See Also

How DPM Works

System Requirements

For DPM and protected computer hardware and software requirements, see System Requirements (http://go.microsoft.com/fwlink/?LinkId=66731).

DPM Licensing

You use a single license for each computer protected by DPM. License type correlates to the type of data being protected.

DPM has two license types: standard and enterprise. The standard license entitles you to protect volumes, shares, and folders, as well as computer system state. The enterprise license entitles you to protect application data, such as mailboxes and databases on an Exchange Server, in addition to files. On a server cluster, DPM installs an agent on each node of the cluster. A license is used for each server node.

The following table lists the license applied for each data type.

DPM Licenses Used for Data Types

Type of protected data	License used
Files only.	Standard
Files on a single node of a server cluster.	Standard
System state.	Standard
SQL Server. (A DPM protection agent on a computer running SQL Server entitles you to protect databases for all SQL instances on that computer.)	Enterprise
Exchange Server.	Enterprise
Windows SharePoint Services. (On a Windows SharePoint Services farm, a license is used for each back-end server and one license is used for the front-end Web server.)	Enterprise
Virtual Server. (On a computer running Virtual Server, a single protection agent installed on the computer enables you to protect any number of virtual machines, or guests, on the host computer. To protect specific application data within a virtual machine, such as to protect the databases for an instance of SQL Server running on a virtual machine, you must install a protection agent directly to the virtual machine. When you protect data on a virtual machine that has a protection agent installed, the appropriate license is used for the type of data being protected.)	Enterprise
Another DPM server.	Enterprise
Data for bare metal recovery using DPM System Recovery Tool.	Enterprise

You do not use a license when you install a protection agent on a computer. The license is applied only when data on a computer is added to a protection group. When you are no longer protecting any data on a specific computer, you can reuse that license on another computer.

When the type of data being protected changes, DPM automatically updates the license usage. For example, you are protecting an Exchange storage group and files on a single server, so you have used an enterprise license to protect that server. Later, you stop protection of the Exchange storage group. Because DPM is now protecting file data only on that server, your license use will change to a standard license.

In a situation where you have only enterprise licenses available and you need to protect file data on a new computer, an enterprise license can be used. For example, you have three standard licenses and three enterprise licenses. You are protecting file data on three computers. You add file data from a fourth computer to a protection group. Because all standard licenses have been used already, DPM will apply an enterprise license.

During DPM installation, you enter the number of licenses that you have purchased. After installation, to update the license information, in the **Protection** task area of DPM Administrator Console, in the **Actions** pane, click **View DPM licenses**, and then change the number of purchased licenses as appropriate.

You can purchase additional DPM licenses through the Microsoft Partner program (http://go.microsoft.com/fwlink/?LinkId=71663).

Planning Protection Groups

To create an effective plan for deploying Microsoft System Center Data Protection Manager (DPM) 2007, you must carefully consider your organization's requirements for data protection and recovery and weigh those requirements against the capabilities of DPM.

This section presents the information you require to plan the membership and configuration of your protection groups. As part of the protection group configuration, you will define your recovery goals for the data being protected.

In the context of the Microsoft Operations Framework (MOF), this section assumes that the change—incorporating DPM in your organization to provide data protection and recovery—has been approved and that your task is planning how to implement the change.

For more information about change management in MOF, see Service Management Functions: Change Management (http://go.microsoft.com/fwlink/?LinkId=68729).

This section also assumes that you are adding DPM to an existing disaster recovery strategy for your business. For more information about planning a disaster recovery strategy, see Introduction to Backup and Recovery Services (http://go.microsoft.com/fwlink/?LinkId=71721).

In This Section

What Do You Want to Protect?

To begin planning for DPM deployment, you should first decide which data you want to protect. DPM 2007 offers protection for the following types of data, which are explained in more detail in subsequent topics:

- File data, at the level of volumes, folders, and shares, on file servers running Microsoft Windows Server 2003 or the Windows Server 2008 operating system
- File data on workstations running Microsoft Windows XP Professional SP2 and all editions of the Windows Vista operating system except Home
- Microsoft Exchange Server 2003 SP2 and Exchange Server 2007 data, at the level of storage groups
- Microsoft SQL Server 2000 SP4, SQL Server 2005 SP1, and SQL Server 2005 SP2 data, at the level of databases
- Windows SharePoint Services 3.0 and Microsoft Office SharePoint Server 2007, at the level of farms
- Microsoft Virtual Server 2005 R2 SP1 host and guest configurations
- System state

See Also

Application Data

Clustered Resources

File Data on Servers and Workstations

System State

File Data on Servers and Workstations

You can protect volumes that can be accessed through either drive letters or mount points, folders, and shares.

The simplest approach to selecting data for protection is to select all file data that you include in your current backups. Alternatively, you can select only specific subsets of your data for protection.

The principal factor to consider when selecting data is your need to quickly recover point-in-time copies of the data if data is lost or corrupted. Key candidates for protection are files that change frequently. Other good candidates are files that are frequently accessed, regardless of how often they change.

Important

Although volumes on file servers are typically formatted as NTFS, which is required for DPM protection, it is not uncommon for volumes on workstations to be formatted as FAT or FAT32. To protect these volumes, you must convert them to NTFS. For instructions, see How to Convert FAT Disks to NTFS (http://go.microsoft.com/fwlink/?LinkId=83022).

See Also

File and Folder Exclusion

Protecting Data in DFS Namespaces

Unsupported Data Types

What Do You Want to Protect?

File and Folder Exclusion

You can configure data protection to exclude specified folders and also file types by file name extension.

When you select a volume or share for protection, you automatically select all protectable child items in that volume or share, as shown in the following illustration.

All Child Items Automatically Selected

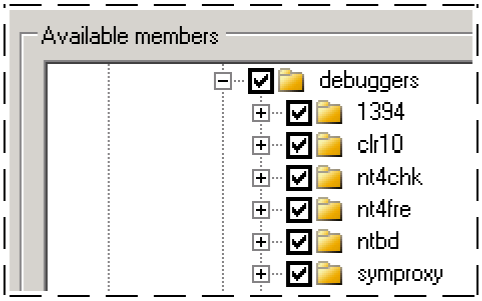

Available members

- ☑ debuggers
  - ☑ 1394
  - ☑ clr10
  - ☑ nt4chk
  - ☑ nt4fre
  - ☑ ntbd
  - ☑ symproxy

To exclude a folder from protection, you ensure that the parent of the folder that you do not want protected is selected, and then you clear the check box of the folder that you do not want protected, as shown in the following illustration.

Folder Excluded from Protection

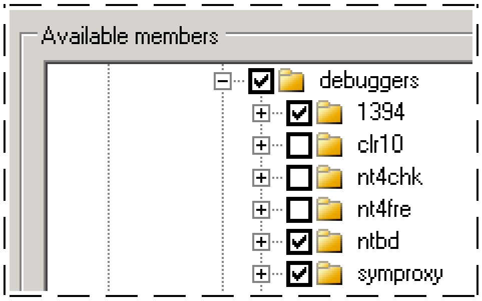

Available members

- ☑ debuggers
  - ☑ 1394
  - □ clr10
  - □ nt4chk
  - □ nt4fre
  - ☑ ntbd
  - ☑ symproxy

When you finish selecting the members for your protection group, you can view the excluded folders, as shown in the following illustration.

View Excluded Folders

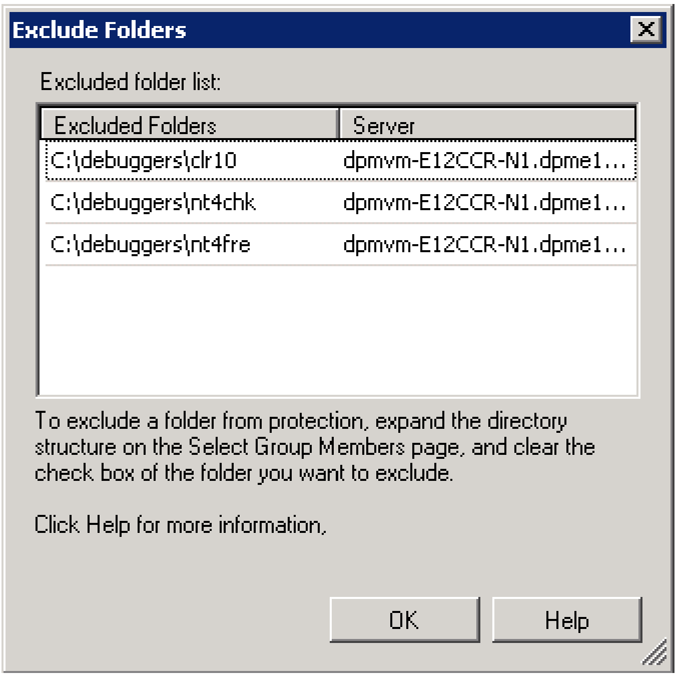

Exclude Folders

Excluded folder list:

Excluded Folders	Server
C:\debuggers\clr10	dpmvm-E12CCR-N1.dpme1...
C:\debuggers\nt4chk	dpmvm-E12CCR-N1.dpme1...
C:\debuggers\nt4fre	dpmvm-E12CCR-N1.dpme1...

To exclude a folder from protection, expand the directory structure on the Select Group Members page, and clear the check box of the folder you want to exclude.

Click Help for more information.

You can also specify file name extensions to exclude from protection at the protection group level. For example, a file server might contain music files or personal files that the business does not want to use disk space or network bandwidth protecting. Exclusion by file name extension applies to all members of the protection group.

The following illustration shows how to exclude files from protection by file name extension.

Exclusion by File Name Extension

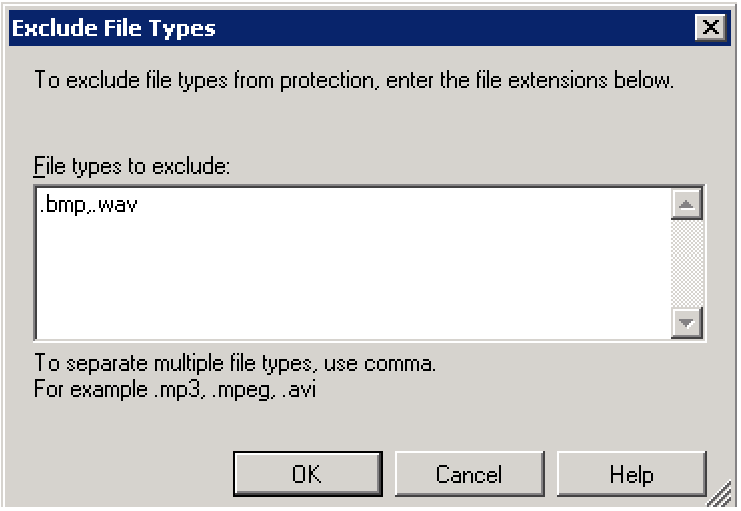

Exclude File Types

To exclude file types from protection, enter the file extensions below.

File types to exclude:

.bmp,.wav

To separate multiple file types, use comma.
For example .mp3, .mpeg, .avi

See Also

Protecting Data in DFS Namespaces

Unsupported Data Types

Protecting Data in DFS Namespaces

You can protect data that is part of a Distributed File System (DFS) Namespaces hierarchy. However, you cannot select shares for protection through the DFS Namespaces hierarchy. Instead, you can select shares for protection only by their target paths.

If your namespace includes roots or links that have multiple targets with the same data, we recommend that you protect only one of the targets. Protecting multiple targets with the same data is redundant.

The following illustration shows DPM protection of a DFS Namespaces target.

Protecting a DFS Namespaces Target by Using DPM

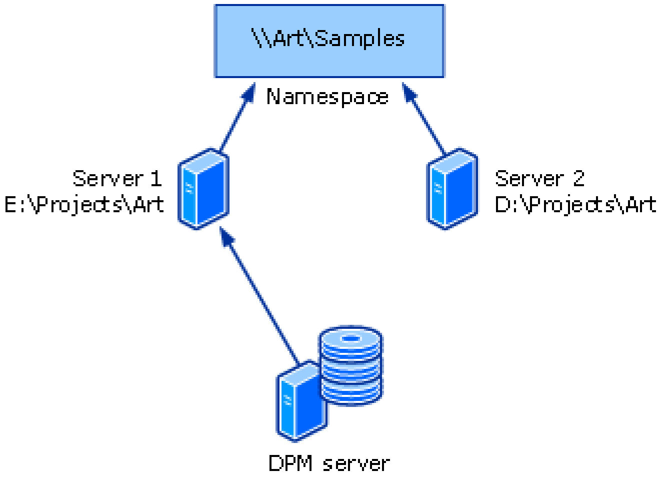

When end-user recovery is enabled for a protected target, users can access previous versions of files through the DFS Namespaces hierarchy. When end users attempt to access previous versions of files on a share that has multiple targets, DPM transparently directs them to the protected target.

See Also

File and Folder Exclusion

Unsupported Data Types

Unsupported Data Types

If a protected data source contains an unsupported data type, DPM continues to protect the supported data types in the affected data source, but it does not protect the unsupported data.

If DPM detects any of the following unsupported data types in a protected data source, the affected data is not protected:

- Hard links
- Reparse points, including DFS links and junction points

  Important

  A protection group can contain data with mount points. When mount points are included in a protection group, DPM protects the mounted volume that is the target of the mount point, but it does not protect the mount point metadata. When you recover data that contains mount points, you must manually re-create your mount point hierarchy. DPM does not support protection of mounted volumes within mounted volumes.
- Recycle Bin

- Paging files
- System Volume Information folder

  Note

  The System Volume Information folder cannot be protected as a file data source. To protect system information for a computer, you must select the computer's system state as the protection group member in the Create New Protection Group Wizard.

- Volumes that are not formatted with NTFS

If a file contains hard links or symbolic links from Windows Vista, DPM cannot replicate or recover the files.

DPM cannot protect files that have any of the following combinations of file attributes:

- Encryption and reparse
- Encryption and Single Instance Storage (SIS)
- Encryption and case sensitivity
- Encryption and sparse
- Case sensitivity and SIS
- Sparse and reparse
- Compression and SIS

See Also

File and Folder Exclusion

Protecting Data in DFS Namespaces

Application Data

You can use DPM to protect the following types of application data:

- **Exchange Server Storage Groups.** DPM can protect storage groups for Microsoft Exchange Server 2003 SP2 and Exchange Server 2007.
  - You cannot exclude from protection any database in the selected storage group.
  - All storage groups on a computer running Exchange Server 2003 must be members of the same protection group or protection of these storage groups will fail.
  - You should disable circular logging for protected storage groups.
- **SQL Server Databases.** DPM can protect databases for Microsoft SQL Server 2000 SP4, SQL Server 2005 SP1, and SQL Server 2005 SP2.
  - Each database in an instance of SQL Server can belong to the same or a different protection group.
  - You cannot exclude from protection any data in the selected database.

- DPM does not support incremental backups for the following databases:
  - SQL Server 2000 and SQL Server 2005 master databases
  - SQL Server 2000 msdb database
  - SQL Server 2000 model database
- **Windows SharePoint Services Data.** DPM can protect server farms for servers running Windows SharePoint Services 3.0 or Office SharePoint Server 2007.
  - You cannot exclude from protection any data in the selected farm.
- **Virtual Server and Virtual Machines.** DPM can protect a Virtual Server host (a computer running Virtual Server 2005 R2 SP1) and the *guests*, or virtual machines, running in the context of that host.

In addition, DPM can protect the data of applications running in the guest. However, data for applications running on virtual machines must be protected and recovered as an application data source, not as a component of a protected virtual machine. For example, to protect and recover data for an instance of SQL Server running on a virtual machine, you select the data source as a SQL Server database. When you protect a virtual machine, application data is also protected, but it can be recovered only by recovering the virtual machine itself.

See Also

Clustered Resources

File Data on Servers and Workstations

System State

Clustered Resources

DPM can protect shared disk clusters for the following:

- File servers
- SQL Server 2000 with Service Pack 4 (SP4)
- SQL Server 2005 with Service Pack 1 (SP1)
- Exchange Server 2003 with Service Pack 2 (SP2)
- Exchange Server 2007

DPM can protect non-shared disk clusters for Exchange Server 2007 (cluster continuous replication). DPM can also protect Exchange Server 2007 configured for local continuous replication.

See Also

Application Data

File Data on Servers and Workstations

System State

System State

DPM can protect the system state for any computer on which a DPM protection agent can be installed, except computers running Windows Vista or Windows Server 2008.

Workstation and Member Server System State

When DPM backs up the system state of a workstation or member server, the following components are protected:

- The boot files
- The COM+ class registration database
- The registry
- System files that are under Windows File Protection

Domain Controller System State

When DPM backs up the system state of a domain controller, the following components are protected:

- Active Directory Domain Services (NTDS)
- The boot files
- The COM+ class registration database
- The registry
- The system volume (SYSVOL)

Certificate Services System State

When DPM backs up the system state of a member server or domain controller with Certificate Services installed, Certificate Services is protected in addition to the member server or domain controller system state components.

Cluster Server System State

When DPM backs up the system state of a cluster server, the cluster service metadata is protected in addition to the member server system state components.

See Also

Application Data

Clustered Resources

File Data on Servers and Workstations

What Are Your Goals for Recovery?

In planning for data protection, you must set realistic recovery goals for each data source that you will protect. Not all information or data maintained on your company's computers requires equal protection, nor does all of it merit the same investment in protection. Your deployment plan should establish recovery goals for each data source according to your business needs for protection of that data.

In DPM, you set your recovery goals in terms of *synchronization frequency*, *recovery point schedule*, and *retention range*, as follows:

- Synchronization frequency should be selected based on your data loss tolerance*,* or how much data you can lose. You can specify the synchronization for a protection group to occur as frequently as every 15 minutes. You can also specify less frequent synchronizations. At a minimum, DPM must synchronize the replicas for a protection group at least once between recovery points.
- The recovery point schedule establishes how many recovery points of this data should be created and when. A recovery point is the date and time of a version of a data source that is available for recovery from media that is managed by DPM.
- The retention range is how long you need the backed-up data available. To determine your retention range needs, consider the pattern of recovery requests you experience in your enterprise. If requests are concentrated within two weeks of data loss, 10 days might be an appropriate retention range for you. If requests are concentrated at a later time, you might need a longer retention range.

For example, your recovery goals for a specific Exchange Server database could be that the most recent data is never more than 30 minutes old, that you can select from versions created at 30-minute intervals, that it will be available for recovery from disk for 14 days, and that it will be available for recovery from tape for 3 years.

See Also

Planning Protection Configurations

Recovery Goals for Disk-Based Protection

Recovery Goals for Tape-Based Protection

What Do You Want to Protect?

Recovery Goals for Disk-Based Protection

Although all members of a protection group share the same synchronization frequency, the synchronization process and the resulting recovery point schedule differ based on the type of data being protected. For more information, see How DPM Works.

Synchronization and Recovery Points for Files

For a file volume or share, the protection agent on the protected computer tracks changed blocks in the change journal that is part of the operating system. During synchronization, these changes are transferred to the DPM server and then applied to the replica to synchronize the replica with the data source.

You can select a synchronization frequency interval of anywhere from 15 minutes to 24 hours. The default is 15 minutes. You can also select to synchronize only before a recovery point is created.

Recovery points, which are shadow copies of the replica for file data, are created from the synchronized replica on a configurable schedule. Each file synchronization does not result in a recovery point unless you synchronize only before each recovery point; however, you can manually create a recovery point from the most recent file synchronization.

For example, a volume is synchronized hourly and a recovery point for the volume is created at 8:00 A.M., 12:00 P.M., and 6 P.M. A user makes changes to a file on the volume at 1:30 P.M.; however, when another user makes changes an hour later, the file is inadvertently corrupted, and you are asked to recover the file with the first user's changes. Because the changes at 1:30 P.M. were made after the most recent recovery point was created at 12:00 P.M., you cannot recover the file from the most recent recovery point. However, you can manually create a recovery point from the appropriate synchronization of that replica and then recover the file from that new recovery point.

The default schedule creates recovery points at 8:00 A.M., 12:00 P.M., and 6:00 P.M. daily. You can modify both the times and the specific days. You cannot specify different times for different days. For example, you can schedule recovery points for 2:00 A.M. and 2:00 P.M. on weekdays only; however, you cannot schedule recovery points for 2:00 A.M. on weekdays and at 12:00 P.M. on weekends.

Retention Range for Files

Retention range is the duration of time for which the data should be available for recovery. When the retention range for a recovery point expires, the recovery point is deleted.

You can select a retention range between 1 and 448 days for short-term disk-based protection, up to 12 weeks for short-term tape-based protection, and up to 99 years for long-term tape-based protection. DPM can store a maximum of 64 recovery points for each file member of a protection group.

For example, if you select to synchronize before each recovery point and you schedule 6 recovery points daily, and you set a retention range of 10 days, recovery points for the files in that protection group never exceed 64. However, if you choose a combination of settings that exceeds the limit of 64 recovery points, DPM warns you during the configuration process so that you can modify your selections; you cannot configure a protection configuration for files that exceeds the limit of 64 recovery points.

Synchronization and Recovery Points for Application Data

For application data, changes to volume blocks belonging to application files are tracked by the volume filter. Synchronization of application data is analogous to an incremental backup and creates an accurate reflection of the application data when combined with the replica.

You can select a synchronization frequency interval of anywhere from 15 minutes to 24 hours. The default is 15 minutes. You can also select to synchronize only before a recovery point is created. If you select to synchronize only before a recovery point is created, DPM performs express full backup to synchronize the replica according to the recovery point schedule.

For applications that support incremental backups, the default schedule results in recovery points for each synchronization (every 15 minutes) and for the express full backup at 8:00 P.M. daily. For applications that do not support incremental backups, the default schedule results in a recovery point for the express full backup at 8:00 P.M. daily.

You can modify both the times and the specific days. You cannot specify different times for different days. For example, you can schedule recovery points for 2 A.M. and 2 P.M. on weekdays only; however, you cannot schedule recovery points for 2 A.M. on weekdays and at 12:00 P.M. on weekends.

Exception for Some SQL Server Databases

Transaction log backups, which DPM uses for incremental synchronization of application data, cannot be performed for a SQL Server database that is read-only, configured for log shipping, or configured to use the Simple Recovery Model. For those SQL Server databases, recovery points correspond to each express full backup.

Comparing Synchronization and Express Full Backup

To enable faster recovery time, DPM will regularly perform an express full backup in place of incremental synchronization. An express full backup is a type of synchronization that updates the replica to include the changed blocks.

Note

> You can modify the express full backup schedule for any protection group that contains application data by using the **Optimize performance** action in the **Protection** task area or by using the Modify Group Wizard.

Retention Range for Application Data

You can select a retention range between 1 and 448 days for short-term disk-based protection, up to 12 weeks for short-term tape-based protection, and up to 99 years for long-term tape-based protection.

For example, if you select to synchronize every 15 minutes and you set a retention range of 10 days, those recovery goals result in a protection plan that maintains 960 recovery points for application data in that protection group after the initial 10 days of data protection.

See Also

Recovery Goals for Tape-Based Protection

Recovery Goals for Tape-Based Protection

DPM protects data on tape through a combination of full and incremental backups from either the protected data source (for short-term protection on tape or for long-term protection on tape when DPM does not protect the data on disk) or from the DPM replica (for long-term protection on tape when short-term protection is on disk).

The choices for retention range, frequency of backups, and recovery options are different for short-term and long-term protection.

Note

You can select disk or tape for short-term protection, but not both.

Short-Term Protection on Tape

For short-term data protection on tape, you can select a retention range of 1–12 weeks. DPM provides management support of your tapes through alerts and reports, and it uses the specified retention range to establish the expiration date for each tape.

Your options for backup frequency are daily, weekly, or biweekly, depending on the retention range.

If you select short-term protection on tape using both incremental and full backups, the retention range will be longer than the one you specified (up to a maximum of 1 week longer) because of a dependency between full and incremental backups. Tapes containing full backup are recycled only after all dependent incremental tapes are recycled. Because full backup happen once a week and the incrementals daily, the weekly full backup tape must wait for the six daily incremental backup tapes to be recycled before the full backup tape is recycled. If an incremental backup fails and there is no incremental tape to recycle, the full backup tape will be recycled earlier.

Long-Term Protection on Tape

For long-term data protection, also known as tape archive, you can select a retention range between 1 week and 99 years. DPM provides management support of your tape archives through alerts and reports, and it uses the specified retention range to establish the expiration date for each tape.

The frequency of backup is based on the specified retention range, as shown in the following list:

- When the retention range is 1–99 years, you can select backups to occur daily, weekly, biweekly, monthly, quarterly, half-yearly, or yearly.
- When the retention range is 1–11 months, you can select backups to occur daily, weekly, biweekly, or monthly.
- When the retention range is 1–4 weeks, you can select backups to occur daily or weekly.

See Also

Recovery Goals for Disk-Based Protection

Planning Protection Configurations

After you identify the data sources that you need to protect and determine your recovery goals, your next step is to analyze the information that you gathered so that you can organize the data sources into protection groups.

A *protection group* is a collection of data sources that share the same protection configuration. The *protection configuration* consists of the protection group name, and settings for disk allocations, replica creation method, and on-the-wire compression.

To plan a protection group, you must make the following decisions:

- Which data sources will belong to the protection group?
- Which protection method (disk-based, tape-based, or both) will you use for the protection group?
- What are your recovery goals for the members of the protection group?
- How much storage space will be needed to protect the selected data?
- Which tape and library should be used?
- What method will you use to create the replica for the members of the protection group?

The topics in this section provide guidelines for making the decisions involved in creating a protection group.

In This Section

Selecting Protection Group Members

Selecting a Data Protection Method

Defining Recovery Goals

Allocating Space for Protection Groups

Specifying Tape and Library Details

Choosing a Replica Creation Method

See Also

What Are Your Goals for Recovery?

What Do You Want to Protect?

Selecting Protection Group Members

With Data Protection Manager (DPM) 2007, there are several approaches you can take to organize data sources into protection groups, including the following:

- **By computer**, with all data sources for a computer belonging to the same protection group.
  - An advantage of this approach is that with all data from a computer in the same protection group, you have a single point of adjustment for performance loads.
  - A constraint of this approach is that all data sources of a type on that computer must be assigned the same recovery goals.
- **By data type**, separating files and each application data type into different protection groups.
  - An advantage of this approach is that you can manage data types as a group.
  - A constraint of this approach is that recovering a server can require multiple tapes from several protection groups.

By definition, all members of a protection group share recovery goals—that is, all data sources of a type in a protection group must have the same retention range and data loss tolerance.

If you have only a single stand-alone tape, use a single protection group to minimize the effort to change tapes. Multiple protection groups require a separate tape for each protection group.

Guidelines for Protection Groups

As you design the structure of your protection groups, keep the following guidelines and restrictions in mind:

- Data sources on a computer must be protected by the same DPM server. In DPM, a data source is a volume, share, database, or storage group that is a member of a protection group.
- You can include data sources from more than one computer in a protection group.
- When you select a parent folder or share, its subfolders are automatically selected. You can designate subfolders for exclusion and also exclude file types by extension.

- Verify that you do not have more than 100 protectable data sources on a single volume. If you do, distribute your data sources across more volumes if possible.
- All protection group members of the same type (file or application data) will have the same recovery goals. However, within the same protection group, files can have different recovery goals than application data.

  Exception: If a SQL Server database is configured to use the Simple Recovery Model or is the primary database in a log shipping pair, the recovery goals for that database will be configured separately from the recovery goals for all other application data.
- All storage groups on a computer running Exchange Server 2003 must be members of the same protection group.
- When you select a data source that contains a reparse point (mount points and junction points are data sources that contain reparse points), DPM prompts you to specify whether you want to include the target of the reparse point in the protection group. The reparse point itself is not replicated; you must manually re-create the reparse point when you recover the data.

Special Considerations for Protecting Data on Workstations

Your recovery goals for data on user workstations might differ from the recovery goals for data on file servers. You should consider placing file servers and workstations in different protection groups so that you can adjust the synchronization schedules separately. For example, if you synchronize data on file servers every 15 minutes, any workstations that belong to the same protection group as the file servers are also synchronized every 15 minutes.

Special Considerations for Protecting Data Over a WAN

Network bandwidth usage throttling and on-the-wire compression are performance optimization features that are particularly important for deployments in which a DPM server protects data over a wide area network (WAN) or other slow network.

On-the-wire compression is configured at the protection-group level.

Network bandwidth usage throttling is configured at the protected-computer level. In addition, you can specify different network bandwidth usage throttling rates for work hours, non-work hours, and weekends, and you define the times for each of those categories.

When protecting application data such as Exchange storage groups or SQL Server databases over a WAN, consider reducing the schedule for express full backups.

How Important Is the Protection Group Membership Decision?

Protection group members cannot be moved between protection groups. If you decide later that a protection group member needs to be in a different protection group, you must remove the member from its protection group and then add it to a different protection group.

If you determine that the members of a protection group no longer require protection, you can stop protection of the protection group. When you stop protection, your options are to retain protected data or to delete protected data.

- **Retain protected data option**: Retains the replica on disk with associated recovery points and tapes for the specified retention range.
- **Delete protected data option**: Deletes the replica on disk and expires data on the tapes.

See Also

Planning Protection Configurations

Selecting a Data Protection Method

Data Protection Manager (DPM) 2007 offers the following data protection methods: disk-based (D2D), tape-based (D2T), or a combination of disk-based and tape-based protection (D2D2T).

The data protection method is configured at the protection-group level. If you want to use different methods to protect two data sources, the data sources cannot belong to the same protection group.

The following table compares the advantages and disadvantages of each method.

Comparison of Data Protection Methods

Method	Advantages	Disadvantages	When to use
Disk-based protection only	• Speed of data recovery. • Speed of data backup. • Backups are less likely to have errors. • Ability to have redundancy to handle failure using technologies such as RAID.	• Disks are not a simple solution for archive needs, because of the cost of disks and the inconvenience of storing offsite.	• When you have a limited data loss tolerance. • When you need faster recovery times.

Method	Advantages	Disadvantages	When to use
	• Less manual intervention, such as changing tapes.		
Tape-based protection only	• Can be stored offsite for security and as a contingency for disaster recovery. • Easy to increase capacity by adding more tapes.	• Slower and more cumbersome recovery process. • Prone to errors.	• When data loss tolerance is more generous. • When recovery time objective is generous. • For data that does not change frequently and does not require backup as frequently. • For lengthy retention period.
Both disk-based and tape-based protection	• Combined advantages of above, while balancing out each method's disadvantages. • A single point of management.		

See Also

Planning Protection Configurations

Defining Recovery Goals

After you select the members of a DPM protection group and the methods to use for data protection, you define the recovery goals for the file data and application data in that protection group.

The recovery goals are defined by the configuration of retention range, synchronization frequency, and recovery point schedule. DPM provides default settings for the recovery goals; however, you can modify each or all of the settings.

At least one synchronization must be scheduled to occur between scheduled recovery points. For example, you specify a synchronization frequency of every 45 minutes. You cannot then configure recovery points to be created at 1:00 P.M. and 1:30 P.M. because there is no intervening synchronization between the recovery points.

When a SQL server is configured to use the Simple Recovery Model or is the primary server in a log shipping pair, the recovery points for any protected databases on that server are created according to the schedule for express full backups.

The following topics in this section provide detailed information to help you plan your recovery goals:

- Recovery Goal Options for Each Protection Method
- Recovery Point Schedules for Long-Term Protection
- Scheduling Options for Long-Term Protection
- Customizing Recovery Goals for Long-Term Protection

See Also

Planning Protection Configurations

Recovery Goal Options for Each Protection Method

The following table lists the recovery goal options for each DPM protection method.

Recovery Goal Options for Protection Methods

Protection method	**Retention range**	**Synchronization frequency or backup schedule**	**Recovery points**
Short-term on disk	1–448 days	Select a frequency between 15 minutes and 24 hours, or select **Just before a recovery point**.	When a specific synchronization frequency is selected: • Recovery points for files are created according to the schedule you configure.

Protection method	Retention range	Synchronization frequency or backup schedule	Recovery points
			- Recovery points for application data are created after each synchronization. When **Just before a recovery point** is selected, recovery points for all protection group members are created according to the schedule you configure.
Short-term on tape	1–12 weeks	Select to back up: - Every day - Every week - Every two weeks	Instead of recovery points, you configure one of the following backup types: - Full and incremental backups - Only full backup When you select weekly or every two weeks, only full backup is available. You specify the day and time. When you select daily full backups, you specify the time. When you select daily full and incremental, you specify the day and time for the full backup and for the incremental backup.
Long-term on tape	Minimum: 1 week Maximum: 99 years	Select to back up: - Daily - Weekly - Biweekly - Monthly	See Recovery Point Schedules for Long-Term Protection and Customizing Recovery Goals for Long-Term Protection.

Protection method	Retention range	Synchronization frequency or backup schedule	Recovery points
		• Quarterly • Half-yearly • Yearly	

See Also

Defining Recovery Goals

Recovery Point Schedules for Long-Term Protection

The following table lists the DPM recovery point schedule for the different long-term protection combinations.

Recovery Point Schedules for Long-Term Protection

Backup frequency and retention range	Recovery point schedule
Daily, 1–4 weeks	Full backup daily
Daily, 1–11 months	1 full backup each day for 4 weeks 1 full backup each month after the initial 4 weeks
Daily, 1–99 years	1 full backup each day for 4 weeks 1 full backup each month after the initial 4 weeks, until the 12th month 1 full backup each year after the initial 11 months
Weekly, 1–4 weeks	Full backup weekly
Weekly, 1–11 months	1 full backup each week for 4 weeks 1 full backup each month after the initial 4 weeks
Weekly, 1–99 years	1 full backup each week for 4 weeks 1 full backup each month after the initial 4 weeks, until the 12th month

Backup frequency and retention range	Recovery point schedule
	1 full backup each year after the initial 11 months
Bi-weekly, 1–11 months	1 full backup every 2 weeks for 4 weeks 1 full backup each month after the initial 4 weeks
Bi-weekly, 1–99 years	1 full backup every 2 weeks for 4 weeks 1 full backup each month after the initial 4 weeks, until the 12th month 1 full backup each year after the initial 11 months
Monthly, 1–11 months	Full backup monthly
Monthly, 1–99 years	1 full backup each month, until the 12th month 1 full backup each year after the initial 11 months
Quarterly, 1–99 years	1 full backup every 3 months until the 12th month 1 full backup each year after the initial 11 months
Half-yearly, 1–99 years	1 full backup every 6 months until the 12th month 1 full backup each year after the initial 11 months
Yearly, 1–99 years	Full backup yearly

See Also

Defining Recovery Goals

Scheduling Options for Long-Term Protection

The following table lists the scheduling options you can modify for long-term protection with DPM.

Scheduling Options for Long-Term Protection

For this backup frequency	Depending on retention range, you can configure
Daily	• Time for daily backup • Specific day or day of week and time for monthly backup • Specific day or date and time for yearly backup
Weekly	• Time and day of week for weekly backup • Specific day or day of week and time for monthly backup • Specific day or date and time for yearly backup
Biweekly	• Time and day of week for biweekly backup • Specific day or day of week and time for monthly backup • Specific day or date and time for yearly backup
Monthly	• Specific day or day of week and time for monthly backup • Specific day or date and time for yearly backup
Quarterly	• Time and date for quarterly backup (Quarterly backups are performed in January, April, July, and October on the specified day of the month.) • Specific day or date and time for yearly backup
Half-yearly	• Time, specific day or date, and months for half-yearly backup • Specific day or date and time for yearly backup
Yearly	• Specific day or date and time for yearly backup

See Also

Defining Recovery Goals

Customizing Recovery Goals for Long-Term Protection

When you specify a retention range and backup frequency, DPM generates a schedule of backup jobs. (For more information, see Recovery Point Schedules for Long-Term Protection.) You can also customize the schedule of backup jobs for your recovery goals, to replace the default schedule.

When you customize the schedule of backup jobs for a protection group, you specify a recovery goal for each backup interval. Your interval choices for backup frequency are as follows:

- Daily
- Weekly
- Monthly
- Yearly

You can specify a recovery goal for up to three backup frequency intervals. For each backup frequency interval, you specify the retention range for the tape, the number of copies of the tape that should be made, and the tape label.

For example, by customizing the recovery goals for a protection group, you could configure backups to occur according to the following schedule:

- One copy of weekly backups, retained for two weeks
- Two copies of monthly backups, retained for six months
- One copy of the yearly backup, retained for five years

See Also

Planning Protection Configurations

Allocating Space for Protection Groups

When you create a protection group and select disk-based protection, you must allocate space on the storage pool for the replicas and recovery points for each data source that you have selected for membership in the group, and you must allocate space on protected file servers or workstations for the change journal.

DPM provides default space allocations for the members of the protection group. The following table shows how DPM calculates the default allocations.

How DPM Calculates Default Space Allocations

Component	Default Allocation	Location
Replica volume	For files: • (Data source size x 3) / 2 For Exchange data: • Data source size x (1 + log change) / (alert threshold - .05) For SQL Server data: • Data source size x (1 + log change) / (alert threshold - .05) For Windows SharePoint Services data: • Total size of all databases/ (alert threshold - .05) For Virtual Server data: • Data source size x 1.5 For system state: • (Data source size x 3) / 2	DPM storage pool or custom volume
Recovery point volume	For files: • (Data source size x retention range in days x 2) / 100 + 1600 MB For Exchange data: • 4.0 x retention range in days x log change x data source size + 1600 MB For SQL Server data: • 2.5 x retention range in days x log change x data source size + 1600 MB For Windows SharePoint Services data: • 1.5 x retention range in days x log change x total size of all databases + 1600 MB For Virtual Server data:	DPM storage pool or custom volume

Component	Default Allocation	Location
	• (Data source size x retention range in days x 0.02) + 1600 MB For system state: • (Data source size x retention range in days x 2) / 100 + 1600 MB	
Change journal (for file protection only)	300 MB	Protected volume on the file server or workstation

The values used in the preceding table are defined as follows:

- **Alert%—**Threshold for the alert associated with replica growth; typically 90%.
- **Log change—**This is the change rate on the database or storage group in question. This varies widely, but for the purposes of the default recommendation in DPM, it is assumed to be 6% for Exchange and SQL Server data and 10% for Windows SharePoint Services data.
- **Retention range (RR)—**This is the number of recovery points stored; it is assumed to be 5 for purposes of the DPM default recommendation.
- **System state data source size—**The data source size is assumed to be 1 GB.

When you create a protection group, in the **Modify Disk Allocation** dialog box, the **Data Size** column for each data source displays a **Calculate** link. For the initial disk allocation, DPM applies the default formulas to the size of the volume on which the data source is located. To apply the formula to the actual size of the selected data source, click the **Calculate** link. DPM will determine the size of the data source and recalculate the disk allocation for the recovery point and replica volumes for that data source. This operation can take several minutes to perform.

We recommend that you accept the default space allocations unless you are certain that they do not meet your needs. Overriding the default allocations can result in allocation of too little or too much space.

Allocation of too little space for the recovery points can prevent DPM from storing enough recovery points to meet your retention range objectives. Allocation of too much space wastes disk capacity.

If, after you have created a protection group, you discover that you have allocated too little space for a data source in the protection group, you can increase the allocations for the replica and recovery point volumes for each data source.

If you discover that you have allocated too much space for the protection group, the only way to decrease allocations for a data source is to remove the data source from the protection group, delete the replica, and then add the data source back to the protection group with smaller allocations.

To help you estimate your storage space needs, download the DPM storage calculator (http://go.microsoft.com/fwlink/?LinkId=104370).

See Also

Planning Protection Configurations

Specifying Tape and Library Details

If you select protection using tape, you must specify the number of copies of each tape that DPM should create and the configuration options for the backup tapes. You can choose one of the following options:

- **Compress data**

  If you select this option, DPM compresses the data as it is written to the tape, which reduces the space needed on the tape and increases the number of backup jobs that can be stored on the same tape. Compression does not significantly increase the time required to complete the backup job. The rate of compression varies according to the type of data.

- **Encrypt data**

  If you select this option, DPM encrypts the data as it is written to the tape, which increases the security for archived data. Encryption does not significantly increase the time required to complete the backup job.

  > **Note**
  >
  > To enable encryption, a valid encryption certificate must be available on the DPM server. For instructions, see "How to Encrypt Data in a Protection Group" in DPM Help.

See Also

Planning Protection Configurations

Choosing a Replica Creation Method

When you create a protection group, you must choose a method for creating the replicas for the volumes included in the group. Replica creation involves copying all the data selected for protection to the DPM server and then running synchronization with consistency check for each of the replicas.

DPM can create the replicas automatically over the network, or you can create the replicas manually by restoring the data from removable media such as tape. Automatic replica creation is easier, but, depending on the size of the protected data and the speed of the network, manual replica creation can be faster.

To help you choose a replica creation method, the followingtable provides estimates for how long DPM takes to create a replica automatically over the network given different protected data sizes and network speeds. The estimates assume that the network is running at full speed and that other workloads are not competing for bandwidth. Times are shown in hours.

Hours to Complete Automatic Replica Creation at Different Network Speeds

Size of Protected Data	512 Kbps	2 Mbps	8 Mbps	32 Mbps	100 Mbps
1 GB	6	1.5	< 1	< 1	< 1
50 GB	284	71	18	5	1.5
200 GB	1137	284	71	18	6
500 GB	2844	711	178	45	15

Important

If you are deploying DPM to protect data over a WAN and your protection group includes more than 5 GB of data, we recommend that you choose the manual method for creating the replicas.

Automatic Replica Creation

For large replica creation jobs, you might want to schedule the job to run only during periods of light network traffic.

Manual Replica Creation

If you choose manual replica creation, DPM specifies the precise locations on the DPM server where you must create the replicas. Typically, you create the replicas by restoring your most recent backup of the data source from removable media such as tape. After you restore the data, you complete the process by running synchronization with consistency check for each of the replicas.

It is crucial that when you restore the data to the DPM server to create the replica, you retain the original directory structure and properties of the data source, such as time stamps and security permissions. The more discrepancies that exist between the replicas and the protected data source, the longer the consistency checking part of the process takes. If you do not preserve the original directory structure and properties, manual replica creation can take as long as automatic replica creation.

See Also

Planning Protection Configurations

Planning for DPM Deployment

When you create your deployment plan for Microsoft System Center Data Protection Manager (DPM) 2007, you should plan your protection groups first because the needs of the protection groups—size, rate of data change, location, recovery goals—will inform your decisions for creating and locating DPM servers and tape libraries.

After you plan your protection groups, you can complete the deployment plan by determining the configurations of DPM servers necessary to protect your data most efficiently. The topics in this section include security and management considerations that might influence your deployment plan.

In This Section

Planning the DPM Server Configurations

End-User Recovery Considerations

Security Considerations

See Also

Planning Protection Groups

Planning the DPM Server Configurations

Your deployment plan should specify the number of DPM servers necessary to protect your data and where you plan to locate each DPM server on your network.

Your deployment plan should also specify which instance of Microsoft SQL Server each DPM server will use. DPM requires an instance of SQL Server for the DPM and reporting databases. DPM will install SQL Server during installation on the DPM server, or you can use an existing instance of SQL Server on a remote computer.

A critical component of your DPM server configuration is the *storage pool*, a set of disks that store replicas and recovery points for protected data. The capacity of the storage pool and any custom volumes that you assign to DPM must be sufficient to provide disk-based protection of the selected data sources.

If your deployment plan requires tape-based protection for any data sources, you will need to attach a tape library or stand-alone tape drive to the DPM server.

If you plan to protect a large Windows SharePoint Services farm, you should install DPM on a volume that has sufficient disk space for the DPM database. The DPM database requires about 1 GB for every million items that exist in the farm. For example, if you protect a farm with 5 million items, you would plan about 5 GB storage in the DPM database to hold the catalog for such a farm. This space requirement is in addition to the storage space that DPM requires for the tape backup catalogs, job logs, and so forth.

In This Section

Selecting the Number of DPM Servers

Locating the DPM Servers

Selecting the Instance of SQL Server

Planning the Storage Pool

Planning the Tape Libraries Configuration

See Also

End-User Recovery Considerations

Security Considerations

Selecting the Number of DPM Servers

As you consider the number of DPM servers that your organization requires, keep in mind that there is no precise formula for determining the number of DPM servers. In practice, the number of servers and amount of data that a single DPM server can protect will vary based on the following factors:

- Change rate of the data sources to be protected
- The amount of space available in the storage pool
- How often the data will be synchronized
- Available bandwidth at each protected computer
- Aggregate bandwidth on the DPM server

To get an estimate of your data change rate, you can review an incremental backup for a recent, average day. The percentage of your data included in an incremental backup is usually indicative of your data change rate. For example, if you have a total of 100 GB of data and your incremental backup includes 10 GB, your data change rate is likely to be approximately 10 percent each day.

However, because the method that DPM uses to record changes to data is different from that of most backup software, incremental backup size is not always a precise indicator of data change rate. To refine your estimate of your data change rate, consider the characteristics of the data you want to protect.

For example, while most backup software records data changes at the file level, DPM records changes at the byte level. Depending on the type of data that you want to protect, this can translate to a data change rate that is lower than the incremental backup might suggest.

The following table lists the data source limits that a DPM server that meets the minimum hardware requirements can protect and the recommended disk space required per DPM server.

Platform	Data Source Limit	Recommended Disk Space
32-bit computers	150 data sources. We recommend approximately 30 to 40 servers fanning into a single DPM server.	10 TB **Note** There is a Volume Shadow Copy Service (VSS) non-paged pool limitation on x86 32-bit operating systems. If you are protecting data using a secondary DPM server, the recommended disk space is only 6 TB.
64-bit computers	300 data sources Data sources are typically spread across 50 to 75 physical servers.	40 TB

Snapshot Limit

A DPM server can store up to 9,000 disk-based snapshots, including those retained when you stop protection of a data source.The snapshot limit applies to express full backups and file recovery points, but not to incremental synchronizations.

The snapshot limit applies per DPM server, regardless of storage pool size. When you configure protection groups, the DPM server is provisioned for the number of snapshots to accommodate the protection group configuration. You can use the following cmdlet in DPM Management Shell to identify the number of snapshots the server is provisioned for:

$server=Connect-DPMServer –
DPMServerName *Name* **$server.CurrentShadowCopyProvision**

When planning your DPM deployment, you need to consider the snapshot limit as part of the DPM server capacity. The following table lists examples of the number of snapshots that result from different protection policies.

Protection policy	**Snapshots**
Exchange storage group: daily express full backup and 15-minute incremental synchronization with a retention range of 5 days	5
Volume on a file server: 3 daily recovery points with a retention range of 21 days	63
SQL database: 2 express full backups daily with a retention range of 14 days	28
Total:	**96**

See Also

Planning the DPM Server Configurations

Locating the DPM Servers

DPM requires a Windows Server 2003 Active Directory Domain Services directory services structure to support its protection and recovery operations.

DPM can protect servers and workstations that are located in the same domain as the DPM server or in a domain that has a two-way trust relationship with the domain that the DPM server is located in.

When deciding where to locate your DPM server, consider the network bandwidth between the DPM server and the protected computers.

DPM supports teamed network interface cards (NICs). Teamed NICs are multiple physical NICs that are configured to be treated as a single NIC by the operating system. Teamed NICs provide increased bandwidth by combining the bandwidth available using each NIC and failover to the remaining NIC or NICs when a NIC fails. DPM can use the increased bandwidth achieved by using teamed NICs on the DPM server.

Another consideration for the location of your DPM servers is the need to manage tapes and tape libraries manually, such as adding new tapes to the library or removing tapes for offsite archive.

See Also

Planning the DPM Server Configurations

Selecting the Instance of SQL Server

A typical DPM installation includes an instance of SQL Server that is installed by DPM Setup. The instance of SQL Server that is installed by DPM Setup is included in the DPM software and does not require a separate SQL Server license.

However, when you install DPM 2007, you can specify a remote instance of SQL Server to be used by DPM instead of the default instance of SQL Server that is included with DPM.

To use a remote instance of SQL Server, the server running SQL Server and the DPM server should be located in the same domain. A specific instance of SQL Server can be used by only one DPM server. The remote instance of SQL Server cannot be on a computer that is running as a domain controller.

Note

> If the remote instance of SQL Server is running as a domain account, you should enable the named pipes protocol for communication with the DPM server. For instructions on configuring the named pipes protocol, see Configuring Client Network Protocols (http://go.microsoft.com/fwlink/?LinkId=87976).

The remote instance of SQL Server must be running Internet Information Services (IIS) and SQL Server 2005 Standard or Enterprise Edition with SP2, including the following components:

- SQL Server Database Engine
- Reporting Services

We recommend you use the following settings on the remote instance of SQL Server:

- Use the default failure audit setting.
- Use the default Windows Authentication mode.
- Assign a strong password to the sa account.
- Enable password policy checking.
- Install only the SQL Server Database Engine and Reporting Services components.
- A remote instance of SQL Server should not run as Local System.
- Run SQL Server by using a low-privileged domain user account.

See Also

Planning the DPM Server Configurations

Planning the Storage Pool

The storage pool is a set of disks on which the DPM server stores the replicas and recovery points for the protected data. Planning the storage pool involves calculating capacity requirements and planning the configuration of the disks.

You can also substitute custom volumes that you define in Disk Management for volumes in the storage pool.

DPM can use any of the following for the storage pool:

- Direct attached storage (DAS)
- Fibre Channel storage area network (SAN)
- iSCSI storage device or SAN

The storage pool supports most disk types, including Integrated Drive Electronics (IDE), Serial Advanced Technology Attachment (SATA), and SCSI, and it supports both the master boot record (MBR) and GUID partition table (GPT) partition styles.

If you use a SAN for the storage pool, we recommend that you create a separate zone for the disk and tape used on DPM. Do not mix the devices in a single zone.

You cannot add USB/1394 disks to the DPM storage pool.

We recommend that you use disks with capacity of no more than 1.5 terabytes. Because a dynamic volume can span up to 32 disks, if you use 1.5-terabyte disks, DPM can create replica volumes of up to 48 terabytes in size.

Important

> Some original equipment manufacturers (OEMs) include a diagnostic partition that is installed from media that they provide. The diagnostic partition might also be named the OEM partition, or the EISA partition. EISA partitions must be removed from disks before you can add the disk to the DPM storage pool.

In This Section

Calculating Capacity Requirements

Planning the Disk Configuration

Defining Custom Volumes

See Also

Planning the DPM Server Configurations

Calculating Capacity Requirements

Capacity requirements for the DPM storage pool are variable and depend primarily on the size of the protected data, the daily recovery point size, expected volume data growth rate, and retention range objectives.

Daily recovery point size refers to the total size of changes made to protected data during a single day. It is roughly equivalent to the size of an incremental backup. Retention range refers to the number of days for which you want to store recovery points of protected data on disk. For files, DPM can store a maximum of 64 recovery points for each volume included in a protection group, and it can create a maximum of 8 scheduled recovery points for each protection group each day.

Note

> The limit of 64 recovery points for files is a result of the limitations of the Volume Shadow Copy Service (VSS), which is necessary for the end-user recovery functionality of DPM. The recovery point limit does not apply to application data.

In general, we recommend making the storage pool two times the size of the protected data for protection of files. This recommendation is based on an assumed daily recovery point size of approximately 10 percent of the protected data size and a retention range of 10 days (two weeks, excluding weekends).

If your daily recovery point size is larger or smaller than 10 percent of your protected data size, or if your retention range objectives are longer or shorter than 10 days, you can adjust the capacity requirements for your storage pool accordingly.

Regardless of how much capacity you decide to allow for the storage pool in your initial deployment, we recommend that you use extensible hardware so that you have the option of adding capacity should the need arise.

The sections that follow provide guidelines for determining your daily recovery point size and retention range objectives.

Estimating Daily Recovery Point Size

Our recommendation to make the storage pool two times the size of the protected data assumes a daily recovery point size of 10 percent of the protected data size. Daily recovery point size is related to data change rate and refers to the total size of all recovery points created during a single day. To get an estimate of the daily recovery point size for your protected data, you can review an incremental backup for a recent, average day. The size of the incremental backup is usually indicative of the daily recovery point size. For example, if the incremental backup for 100 GB of data includes 10 GB of data, your daily recovery point size will probably be approximately 10 GB.

Determining Retention Range Objectives

Our recommendation to make the storage pool two times the size of the protected data assumes a retention range objective of 10 days (two weeks, excluding weekends). For the typical enterprise, requests for recovery of data are concentrated within two to four weeks after data loss events. A retention range of 10 days provides for recovery of data up to two weeks after a data loss event.

The longer your retention range objective, the fewer recovery points you can create each day. For example, if your retention range objective is 64 days, you can create just one recovery point each day. If your retention range objective is eight days, you can create eight recovery points each day. With a retention range objective of 10 days, you can create approximately six recovery points each day.

See Also

Defining Custom Volumes

Planning the Disk Configuration

Planning the DPM Server Configurations

Planning the Disk Configuration

If you are using direct-attached storage for the DPM storage pool, you can use any hardware-based configuration of redundant array of independent disks (RAID), or you can use a "just a bunch of disks" (JBOD) configuration. Do not create a software-based RAID configuration on disks that you will add to the storage pool.

To decide on the configuration for the disks, consider the relative importance of capacity, cost, reliability, and performance in your environment. For example, because JBOD does not consume disk space for storing parity data, a JBOD configuration makes maximum use of storage capacity. For the same reason, the reliability of JBOD configurations is poor; a single disk failure inevitably results in data loss.

For the typical DPM deployment, a RAID 5 configuration offers an effective compromise between capacity, cost, reliability, and performance. However, because the DPM server workload is composed primarily of write operations, RAID 5 is likely to degrade the performance of a DPM server more markedly than it would in the case of a file server. This degradation in performance can in turn affect the scalability of DPM. The ability of DPM to effectively protect data degrades as performance degrades.

To help you evaluate options for configuring the disks in your storage pool, the following table compares the trade-offs between JBOD and the various levels of RAID, on a scale from 4 (very good) to 1 (acceptable).

Comparison of Configuration Options for Storage Pool Disks

Disk Configuration	Capacity	Cost	Reliability	Performance and Scalability
JBOD	4	4	1	4
RAID 0	4	4	1	4
RAID 1	1	1	4	3
RAID 5	3	3	3	2
RAID 10	1	1	4	4

For more information about RAID, see Achieving Fault Tolerance by Using RAID (http://go.microsoft.com/fwlink/?LinkId=46086).

See Also

Calculating Capacity Requirements

Defining Custom Volumes

Planning the DPM Server Configurations

Defining Custom Volumes

In DPM 2007, you can assign a *custom volume* to a protection group member, in place of the DPM storage pool. A custom volume is a volume that is not in the DPM storage pool and is specified to store the replica and recovery points for a protection group member.

Although the DPM-managed storage pool is sufficient for most business needs, you might want a greater amount of control over storage for specific data sources. For example, you have critical data that you want to store using a high-performance logical unit number (LUN) on a storage area network.

Any volume that is attached to the DPM server can be selected as a custom volume in the Create New Protection Group Wizard, except the volume that contains the system and program files. To use custom volumes for a protection group member, two custom volumes must be available: one volume to store the replica and one volume to store the recovery points.

DPM cannot manage the space in custom volumes. If DPM alerts you that a custom replica volume or recovery point volume is running out of space, you must manually change the size of the custom volume by using Disk Management.

You cannot change the selection of storage pool or custom volume for a protection group member after the group is created. If you must change the storage location for a data source's replica or recovery points, you can do so only by removing the data source from protection and then adding it to a protection group as a new protection group member.

See Also

Calculating Capacity Requirements

Planning the Disk Configuration

Planning the DPM Server Configurations

Planning the Tape Libraries Configuration

You can add tape libraries and stand-alone tape drives to DPM to enable short-term and long-term data protection on tape. The tape libraries and stand-alone tape drives must be physically attached to the DPM server.

Note

> The term *tape libraries* refers to both multi-drive tape hardware and stand-alone tape drives.

Consider the number of tape backup jobs and the size of the protected data when planning the capacity of your tape library. You must also consider the hardware features: a tape library without an autoloader requires manual tape rotations when jobs are being performed.

To plan for the number of tapes you will need for each protection group, multiply the backup frequency by the retention range.

The tape labels for tapes used for long-term protection are assigned when you create a protection group. DPM will assign a default tape label in this format: **DPM - <ProtectionGroupName> - long-term tape <number>**. Before you begin creating protection groups, you should plan your tape naming scheme if you do not want to use the default scheme.

For more information, see Managing Tape Libraries (http://go.microsoft.com/fwlink/?LinkId=91964).

See Also

Planning the DPM Server Configurations

End-User Recovery Considerations

Your deployment plan should specify the data for which end-user recovery will be enabled and the DPM servers that must be configured in Active Directory Domain Services to provide end-user recovery.

End-user recovery enables end users to independently recover data by recovering previous versions of their files. End users can recover previous versions through shares on file servers, through DFS Namespaces, or by using a command on the **Tools** menu of Microsoft Office 2003 applications.

If you currently have Shadow Copies of Shared Folders enabled on a computer that you protect with DPM, you can disable that feature and regain the disk space that it uses. End-users and administrators will be able to recover files from the recovery points on the DPM server.

Enabling end-user recovery requires configuring the schema of Active Directory Domain Services, enabling the end-user recovery feature on the DPM server, and installing the recovery point client software on the client computers.

Configuring Active Directory Domain Services

Configuring Active Directory Domain Services to support end-user recovery involves four operations:

1. Extending the schema

2. Creating a container
3. Granting the DPM server permissions to change the contents of the container
4. Adding mappings between source shares and shares on the replicas

The schema is extended only once; however, you must configure the Active Directory schema extension for each DPM server. When you enable end-user recovery for additional DPM servers in the domain, the process performs steps 3 and 4 for each additional server. DPM will update the share mapping (step 4) after each synchronization, if needed.

DPM administrators who are both schema and domain administrators in the Active Directory Domain Services domain can complete these steps with a single click in DPM Administrator Console. DPM administrators who are not schema and domain administrators can complete these steps by directing a schema and domain administrator to run the DPMADSchemaExtension tool.

The DPMADSchemaExtension tool is stored on the DPM server in the folder Microsoft Data Protection Manager\2006\End User Recovery. A user who is both a schema and domain administrator can run the tool on any computer running Windows Server 2003 that is a member of the domain in which the DPM server is deployed. The administrator must specify the name of the DPM server when running the tool.

If you use the DPMADSchemaExtension tool to enable end-user recovery, you must run it once for each DPM server.

Installing the Shadow Copy Client Software

Before end users can begin independently recovering previous versions of their files, the DPM recovery point client software must be installed on their computers. If a client for Shadow Copies of Shared Folders is present on the computer, the client software must be updated to support DPM.

The recovery point client software can be installed on computers running the Windows XP operating system with Service Pack 2 (SP2) or later and Windows Server 2003 with or without SP1.

See Also

Planning the DPM Server Configurations

Security Considerations

Security Considerations

DPM operates as a high-privileged server on the network. To help ensure the security of the DPM server, the DPM security architecture relies on the security features of Windows Server 2003 and Active Directory Domain Services, SQL Server 2005, and SQL Server Reporting Services.

To maintain the DPM security architecture:

- Accept all default security settings.
- Do not install unnecessary software on the DPM server.
- Do not change security settings after DPM is deployed. In particular, do not change SQL Server 2005 settings, Internet Information Services (IIS) settings, DCOM settings, or settings for the local users and groups that DPM creates during product installation.
- A remote instance of SQL Server should not run as Local System.

Installing unnecessary software and changing default security settings can seriously compromise DPM security.

In This Section

Configuring Antivirus Software

Configuring Firewalls

Security Considerations for End-User Recovery

Granting Appropriate User Privileges

See Also

End-User Recovery Considerations

Planning the DPM Server Configurations

Configuring Antivirus Software

DPM is compatible with most popular antivirus software products. However, antivirus products can affect DPM performance, and, if they are not configured properly, they can cause data corruption of replicas and recovery points. This section provides instructions for mitigating such problems.

Configuring Real-Time Monitoring for Viruses

To minimize performance degradation on the DPM server, disable antivirus real-time monitoring of replicas for all protected data sources by disabling real-time monitoring of the DPM process msDPMprotectionagent.exe, which is located in the folder Microsoft Data Protection Manager\DPM\bin. Real-time monitoring of replicas degrades performance because it causes the antivirus software to scan all affected files each time DPM applies changes to the replicas.

Additionally, if you experience degraded performance while using DPM Administrator Console, disable real-time monitoring of the csc.exe process, which is located in the folder Windows\Microsoft.net\Framework\v2.0.50727. The csc.exe process is the C# compiler. Real-time monitoring of the csc.exe process can degrade performance because it causes the antivirus software to scan files that the csc.exe process emits when generating XML messages.

For instructions for configuring real-time monitoring for individual processes, see your antivirus product documentation.

Setting Options for Infected Files

To prevent data corruption of replicas and recovery points, configure the antivirus software on the DPM server to delete infected files rather than automatically cleaning or quarantining them. Automatic cleaning and quarantining can result in data corruption because these processes cause the antivirus software to modify files with changes that DPM cannot detect. Any time that DPM attempts to synchronize a replica that has been modified by another program, data corruption of the replica and recovery points can result. Configuring the antivirus software to delete infected files avoids this problem. Note, however, that you must run manual synchronization with consistency check each time that the antivirus software deletes files from a replica. For instructions for configuring your antivirus software to delete infected files, see the product documentation.

See Also

Security Considerations

Configuring Firewalls

If the computers you want to protect reside behind a firewall, you must configure the firewall to allow communication between the DPM server, the computers it protects, and the domain controllers.

Protocols and Ports

Depending on your network configuration, you might need to perform firewall configuration to enable communication between DPM, the protected servers, and the domain controllers. To help with firewall configuration, the following table provides details about the protocols and ports used by DPM.

Protocols and Ports Used by DPM

Protocol	Port	Details
DCOM	135/TCP Dynamic	The DPM control protocol uses DCOM. DPM issues commands to the protection agent by invoking DCOM calls on the agent. The protection agent responds by invoking DCOM

Protocol	Port	Details
		calls on the DPM server. TCP port 135 is the DCE endpoint resolution point used by DCOM. By default, DCOM assigns ports dynamically from the TCP port range of 1024 through 65535. However, you can configure this range by using Component Services. For more information, see Using Distributed COM with Firewalls (http://go.microsoft.com/fwlink/?LinkId=46088).
TCP	5718/TCP 5719/TCP	The DPM data channel is based on TCP. Both DPM and the protected computer initiate connections to enable DPM operations such as synchronization and recovery. DPM communicates with the agent coordinator on port 5718 and with the protection agent on port 5719.
DNS	53/UDP	Used between DPM and the domain controller, and between the protected computer and the domain controller, for host name resolution.
Kerberos	88/UDP 88/TCP	Used between DPM and the domain controller, and between the protected computer and the domain controller, for authentication of the connection endpoint.
LDAP	389/TCP 389/UDP	Used between DPM and the domain controller for queries.
NetBIOS	137/UDP 138/UDP 139/TCP 445/TCP	Used between DPM and the protected computer, between DPM and the domain controller, and between the protected computer and the domain controller, for miscellaneous operations. Used for SMB directly hosted on TCP/IP for DPM functions.

Windows Firewall

Windows Firewall is included with Windows Server 2003 SP1. If you enable Windows Firewall on the DPM server before you install DPM, DPM Setup will properly configure the firewall for DPM. If you enable Windows Firewall on the DPM server after you install DPM, you must configure the firewall manually to permit communication between the DPM server and protected computers. Configure Windows Firewall on a DPM server by opening port 135 to TCP traffic and specifying the DPM service (Microsoft Data Protection Manager/DPM/bin/MsDPM.exe) and the protection agent (Microsoft Data Protection Manager/DPM/bin/Dpmra.exe) as exceptions to the Windows Firewall policy.

For instructions for configuring Windows Firewall, search on "Windows Firewall" in Windows Help and Support for Windows Server 2003.

See Also

Security Considerations

Security Considerations for End-User Recovery

You can enable end-user recovery for file data, but not for application data. Use only domain-based security groups for permissions to files and folders on which you plan to enable end-user recovery. DPM cannot guarantee consistency between end-user access to data on protected computers and end-user access to recovery points of that data on the DPM server if you rely on local security groups.

For example, if the set of users included in the protected computer's local Users group differs from the set of users included in the DPM server’s local users group, different sets of users will have access to the data on the protected computer and to the recovery points of that data.

See Also

Security Considerations

Granting Appropriate User Privileges

Before you begin a DPM deployment, verify that appropriate users have been granted required privileges for performing the various tasks. The following table shows the user privileges that are required to perform the major tasks associated with DPM.

User Privileges Required to Perform DPM Tasks

Task	Required Privileges
Adding a DPM server to an Active Directory domain	Domain administrator account, or user right to add a workstation to a domain

Task	Required Privileges
Installing DPM	Administrator account on the DPM server
Installing the DPM protection agent on a computer	Domain account that is a member of the local administrators group on the computer
Opening DPM Administrator Console	Administrator account on the DPM server
Extending the Active Directory Domain Services schema to enable end-user recovery	Schema administrator privileges in the domain
Creating an Active Directory Domain Services container to enable end-user recovery	Domain administrator privileges in the domain
Granting a DPM server permissions to change the contents of the container	Domain administrator privileges in the domain
Enabling end-user recovery feature on a DPM server	Administrator account on the DPM server
Installing recovery point client software on a client computer	Administrator account on the client computer
Accessing previous versions of protected data from a client computer	User account with access to the protected share
Recovering Windows SharePoint Services data	Windows SharePoint Services farm administrator account that is also an administrator account on the front-end Web server that the protection agent is installed on

See Also

Security Considerations

Deployment Plan Checklist and Roadmap

This checklist includes the planning tasks necessary to prepare to deploy Data Protection Manager (DPM) 2007.

Task	Reference
Identify each data source to be protected, including the following information: • Data source type (file, Microsoft Exchange,	What Do You Want to Protect?

Task	Reference
Microsoft SQL Server, Microsoft Windows SharePoint Services, Microsoft Virtual Server, system state) - Data source size - Any folders or file name extensions to be excluded from protection - Fully qualified domain name (FQDN) of computer - Cluster name (if applicable)	
Identify one of the following methods for each protection group: - Short-term disk-based protection - Short-term tape-based protection - Long-term tape-based protection - Short-term disk-based protection and long-term tape-based protection - Short-term tape-based protection and long-term tape-based protection	Selecting a Data Protection Method
For each data source, determine the recovery goals for each data protection method that you will use. For short-term disk-based protection, identify the following information: - Retention range - Synchronization frequency - Number of recovery points For short-term tape-based protection, identify the following information: - Retention range - Backup schedule - Type of backup - Number of backup copies - Tape labeling scheme For long-term tape-based protection, identify the following information: - Retention range	What Are Your Goals for Recovery? Defining Recovery Goals

Task	Reference
- Backup schedule and scheduling options - Number of backup copies - Tape labeling scheme	
Organize the data sources into protection groups.	Selecting Protection Group Members
Determine your storage needs, based on your information about the protected data sources and recovery goals.	Allocating Space for Protection Groups
If you are using tape-based protection, decide if you want to compress or encrypt the data on tapes.	Specifying Tape and Library Details
Decide which method of replica creation you will use for each protection group.	Choosing a Replica Creation Method
Identify the DPM server configurations necessary, including the following information: - The number of DPM servers - Where to locate each DPM server - Which instance of SQL Server each DPM server will use	Planning the DPM Server Configurations
Determine the disk configurations each DPM server will require to meet the storage needs of the protection groups. Include any custom volumes that specific data sources will use.	Planning the Storage Pool
Identify the DPM servers that require tape libraries and the capacity of each library.	Planning the Tape Libraries Configuration
Identify the DPM servers for which end-user recovery will be enabled and which clients will require installation of the recovery point client software.	End-User Recovery Considerations

See Also

Introducing Data Protection Manager 2007

Planning for DPM Deployment

Planning Protection Groups