

Chapter 2 Backup Concepts

Overview

This chapter describes:

- Automatic jobs, including rotation and non-rotation day operations
- Custom jobs
- File rules
- Media scheduling
- Retaining backup copies
- Concurrent backup operations and jobs

Contents

Introduction

What Happens On Rotation Day?

What Happens on Non-Rotation Day?

Custom Jobs

Backup Operations

How Rules Affect Backups

Media Rotation

Tower of Hanoi

Weekly versus Daily Rotation

Grandfather-Father-Son

Other Backup Issues

Retaining Backup Copies

When You Use an Unexpected Media Set

Early Rotation

Deferred Rotation

Moving Daily Backups Off-Site

Concurrency in Backup Director

DOS Workstations

Running Concurrent Backup Operations

Running Concurrent Jobs

Introduction

Backup Director offers data protection capabilities that can make run-time decisions for you such as:

- Which media should be on-site or off-site
- Which files can safely be removed to free disk space
- Which files can be skipped because they are already protected

Backup Director submits operations in one of two ways, as an automatic job or as a custom job.

Automatic Jobs

Automatic jobs are the foundation of your data protection strategy and make Backup Director unique. **Automatic** means that Backup Director automatically determines:

- Which operation to perform—Automatic jobs include database maintenance and backup operations. Automatic jobs perform different operations on rotation and non-rotation days.
- Which resources should be included in the automatic operations
- Which media to use
- When to perform the backup
- When to move media sets on-site and off-site. You must decide whether to implement the recommendations.

The program's decisions are based on the Media Scheduling parameters in Configuration Manager. See Chapter 4, “Customizing Your Installation,” for information about defining automatic operations, including rotation day and non-rotation day operations.

What Happens On Rotation Day?

Rotation day is the day Backup Director asks for a different media set to perform the next scheduled automatic job. When you perform your first automatic backup, Backup Director will perform the operations it normally performs on a rotation day.

Rotation day is the only day Backup Director will automatically perform certain database maintenance operations, such as updating the File History Database to reflect any files that have been deleted from the disk.

On rotation day, Backup Director:

- Updates media records and verifies the integrity of the databases (referred to as **database maintenance**).
- Copies to media all eligible files on all protected resources (**full backup**). As a result, the program overwrites the sessions written prior to the current rotation.

What Happens on Non-Rotation Day?

Non-rotation day is a day on which the program asks for a media from the same media set as the previous day's automatic job. On a non-rotation day, Backup Director performs:

- An incremental (or differential or full backup) operation on volume resources.
- A full backup operation on non-volume resources (for example, the Bindery files).

Automatic jobs on these days can be significantly faster than those on rotation days if incremental or differential backups are used. These backup operations copy only changed files to media.

See Chapter 3, "Getting Started," for information about submitting automatic jobs.

Custom Jobs

Custom jobs are requests to perform a single operation on selected items (servers, resources, directories, or files). Custom jobs can supplement protection by automatic jobs in special circumstances, or they can provide an alternate backup schedule. Unlike automatic jobs, custom jobs are not required to record session information in the database. These jobs can write to managed or non-managed media (media labeled by the user). Usually, you create a custom job for a special purpose such as:

- Taking snapshots of your system (possibly to store at an off-site location) at regular intervals (for example, weekly or monthly)
- Transferring data to another office (for example, you have to send a snapshot of a resource or directory each week to another location)
- Backing up an entire machine or resource prior to replacing it

The table below summarizes the differences between automatic and custom jobs.

Automatic Jobs...

Custom Jobs...

Automatically label and use managed media

Use managed or non-managed media.

Include the entire Protected Resource List

Can include the entire Protected Resource List or specific resources or files.

Can perform multiple operations automatically

Can perform a single operation (backup, restore, or utility operation).

Perform the operation at each scheduled interval

Perform the operation once, periodically, or following another scheduled job.

Backup Operations

There are three types of backup operations.

- **Full backup**—Makes a backup copy of all eligible files on all resources and is the default backup operation that occurs on a rotation day. Files eligible for a backup operation are those files that are not excluded by file rules. Full backup operations always occur on rotation day.
- **Incremental backup**—The default backup operation on non-rotation days. An incremental backup operation writes copies of all files that have changed since the previous full or incremental backup operation.
- **Differential backup**—Makes a backup copy of every file that has changed since the previous full backup. Because this operation does not clear the archive bit, this can result in the same version of files being copied even if they haven't changed since the last full backup.

Incremental backup operations tend to complete more quickly than differential backups. However if the program uses differential backups, Backup Director can restore an entire resource more quickly because a greater number of the most recently changed files are located on the same media. The program only needs to open one session on each required media set.

How Rules Affect Backups

While users would prefer to have all of their data backed up continuously, this strategy is rarely feasible given the burden to the network and inconvenience to users. Most data rarely warrants this type of protection. As an administrator you need to find the best balance of data protection, convenience, and network resources. Backup Director's file rules can help you define the best protection for the different kinds of data on your network.

File rules are applied to each resource that Backup Director protects. Because network requirements vary depending on the environment, corporate policy, government specifications, etc., users may find it necessary to customize their system to meet defined requirements.

- **Include**—Files with this backup rule are always available for backup operations.
- **Exclude**—Files with this backup rule are never available for backup operations.

By default, most of your files are covered by the "*" filename pattern. This pattern is set to **Include**. This system rule is sufficient for most files.

Rules affect file protection and media space. A more specific filename pattern can supersede a system rule. The filename pattern "*.WK" is more specific than "*". For example, if you add rule to exclude all files with a BAK extension, you will prevent the program from backing up these files and wasting media space. For a complete list of Backup Director's other system rules, see Chapter 9, "Managing Files."

Media Rotation

The purpose of media rotation is to protect media set copies at an off-site location when they are not in use. If a disaster should occur, you can recover data from media stored at the off-site location. By using several media sets that correspond to different time periods, you achieve two data protection goals:

- Copies on different media sets help you to safeguard at least one copy at any time. Extra media sets can be stored off-site.
- Copies corresponding to different dates exist on different media sets. This provides users with the flexibility of restoring from any of several file versions.

Backup Director provides two configurable rotation patterns: Tower of Hanoi (TOH) and Grandfather-Father-Son (GFS). Of the two, TOH is considered to offer the greatest depth of file versions. However, GFS is the more traditional pattern.

Tower of Hanoi

TOH rotation pattern is based on a sophisticated algorithm that guarantees an organized, systematic approach for scheduling media. One distinctive feature of this pattern is the number of unique file versions that are kept across multiple media sets.

Each time a media set is introduced, the rotation interval doubles for the new media set. For example, using weekly rotation, the program rotates the A set every two weeks, the B set every four weeks, the C set every eight weeks, etc. The frequently used media sets have the most recent copies of files, while the infrequently used sets have older versions available. Therefore a rich assortment of file versions are available across different media sets.

With a weekly rotation schedule, five media sets maintain various file versions spanning a two- to four-month period. This provides a variety of file versions—equal to or better than that provided by a 12-media set Grandfather-Father-Son system.

Weekly versus Daily Rotation

If you implement TOH rotation with weekly rotation, the program requires a different media set each week. If you implement TOH rotation with daily rotation, rotation day operations occur every day because the program requires a different media set every day. Which rotation schedule is better? It depends on your individual needs.

Weekly rotation has several advantages:

- Automatic jobs take less time overall since rotation day operations occur only once a week, not every day as in the daily rotation alternative.
- More recent backup versions of evolving files.
- Fewer backup media and media changes are required.
- If users do not add and delete a lot of data every day, a weekly rotation schedule is sufficient.
- A greater variety of file versions spanning a longer time period.

Daily rotation also has advantages:

- Media will not reach capacity as quickly since only one day's worth of backups are copied between rotation.
- If the media becomes damaged, only copies of that day's data are lost.
- If users add, modify, and delete data often and/or your resources are near capacity, you should use a daily rotation scheme.

Grandfather-Father-Son

Grandfather-Father-Son (GFS) is the traditional rotation scheme in the PC LAN environment. Backup Director's implementation includes five pre-defined media set time periods: daily, weekly, monthly, quarterly, and yearly. GFS allows up to 32 media sets, seven of which are automatically designated for your daily media. You can change the number of weekly and monthly media sets to follow a round-robin sequence rather than the calendar. For example, you can set the number of monthly sets to eight rather than twelve media sets.

When using the GFS rotation pattern, the program performs a full backup operation at every automatic job.

The diagram above illustrates the number of backup sessions that may be available on managed media during a three-month window.

Note that the program may not actually build the number of media sets you have configured. For example, if Friday is the rotation day used for weekly media sets, the program will never create a FRIDAY daily media set.

Other Backup Issues

Retaining Backup Copies

Backup sessions written to managed media exist on a media set until the next time the media is used for the rotation day operations. The program overwrites the previous backup sessions and begins writing a new cycle of backup sessions. With backup sessions written to non-managed media, the program does not automatically overwrite previous backup sessions; you have the choice of overwriting the previous session or appending to it.

You can ensure that data located on the session is available for a limited period of time or indefinitely by:

- Preserving backup sessions—Increase the value of the **Preserve Backups** parameter to set the minimum number of days the file version remains on the media. For example, if your installation uses the GFS rotation pattern, the program automatically “preserves” backup copies written to a “MONDAY” media set until the media is rotated the following Monday. If you need to extend this time period, you can set the **Preserve Backups** parameter to seven days or more. Once all of the sessions on a media are older than the value you set for the parameter, the program is able to overwrite the backup sessions.
- Writing to non-managed media—Since, by definition, the program does not rotate non-managed media, the program will not automatically overwrite session on this media. Take this media off-site to prevent other users from overwriting or appending to this media.
- Retiring managed media—You can de-activate, or “retire,” managed media so that the program no longer schedules it for rotation. As a result, the sessions on that media all become permanent. You can still restore data from retired media.

When You Use an Unexpected Media Set

Despite the automation of Backup Director's media rotation schedule, the media itself (or the person implementing the schedule) can still fail. Backup Director performs the automatic backup on any managed media set that is loaded when the program begins servicing the job. Whether you need to substitute the current set for a damaged set or forgot to retrieve the required media set in time for rotation, the program can adapt to whatever managed or blank media you load.

Early Rotation

If you change media sets prior to rotation day, you are performing an **early rotation**. The program performs the usual rotation and non-rotation day operations using the current media set. At the next rotation day, the program rotates to the media set that normally follows the substitute media set.

In GFS rotation, the program attempts to resume the schedule at the next rotation by asking

for the expected media set. For example, if you load PAL:WEEK3:1 instead of the scheduled PAL:WEEK1:1 media, the program will request PAL:WEEK2:1 at the next rotation.

Deferred Rotation

Similarly, you may decide to continue using the current media set rather than rotate to the required media on rotation day. The program continues to append backup sessions to the current media. In TOH rotation, the program will again attempt to rotate to the media set that normally follows the current set. In GFS rotation, the program will again refer to the original schedule and attempt to resume that schedule.

Moving Daily Backups Off-Site

If you are required to move your backup media off-site on a daily basis, Backup Director provides two options to assist you: copying media and the **Daily Media Change Within Set** option. Both options require extra backup media that Backup Director would normally not demand.

The *Copy* operation in Media Manager allows you to duplicate media. You can store the duplicate set of media off-site for security, while keeping the original media set on-site for convenience.

To use Media Manager's *Copy* operation, you must have two backup devices configured. These do not have to be the same media type, nor do they have to be on the same SCSI bus. For example, you could have 4mm and 8mm tape drives or a tape drive and optical drive.

After running your normal backups, keep the media in the drive and perform the media copy operation as described in Chapter 10, "Managing Media."

When **Daily Media Change within Set** option is turned on, the program requests a different media for each day. This option is only available if you use the TOH rotation pattern. You must determine when to remove and retrieve individual media. Because the Off-Site Media Advisor window displays storage recommendations for media sets only, you must determine the appropriate schedule for moving daily media off-site.

See Chapter 4, "Customizing Your Installation," for more information about this option.

Concurrency in Backup Director

Concurrency refers to Backup Director's ability to simultaneously use multiple devices to:

- Perform backup operations on multiple resources simultaneously
- Perform backup jobs and restore and/or utility jobs simultaneously

To take advantage of concurrency, your installation must have at least two devices configured.

If you are going to use concurrent operations, be sure:

- You have eligible media loaded in each device.
- You have at least 2 MB RAM available on your server for each concurrent process (for example, if you have two backup devices you need at least 4 MB of RAM).

DOS Workstations

Concurrency does not apply to multiple local drives on DOS workstations. TSASMS.COM (on the workstation) can only communicate with one drive volume at a time. The solution is to separate these local drives on the Protected Resource List so that the program can back up each drive concurrently with resources that do not belong to the same workstation. See the *"Re-arranging Resources"* section in Chapter 8, "Managing Resources."

Running Concurrent Backup Operations

Backup Director allows you to simultaneously back up multiple volumes to multiple devices during a single backup job or automatic job. The backup engine is the only engine that can simultaneously process resources using multiple devices from a single job request.

Concurrent backup operations are limited by the number of backup devices. For example, if you have two backup devices and submit an backup job for three volumes, the program cannot process the third volume until one of the operations performed on the first two volumes has completed. However, users with a single device can benefit from concurrent database maintenance operations.

Depending on the number of backup devices you are using, concurrent backup operations may be affected if another job is using a device. For example, you have two backup devices and are performing a restore operation, only one device is available for backup operations. If the restore job completes while the backup job is running, you can run the remainder of the backup operations concurrently by loading the second device with the appropriate media.

Running Concurrent Jobs

Running jobs concurrently allows your installation to load up to three different engines (backup, restore, and utility) at once. Each engine can perform only one job at a time. While a backup job may run multiple backup operations simultaneously, Backup Director can process only one backup job in the queue at a time. If each of the three engines are loaded at the same time and there are three backup jobs in the job queue, the program can only process the first backup job submitted to the queue. The remaining backup jobs must wait until the previous backup job has completed.
