Retrospect 10 for Mac

Reviewer's Guide

About this Reviewer's Guide

This document provides a concise guide to understanding Retrospect 10 for Mac OS X. While it is not designed to replace the Retrospect User's Guide, it will provide the reader with…

- descriptions of Retrospect's technologies and terminology,
- knowledge on how Retrospect's components work together to provide reliable backup and recovery, and
- examples and methods for protecting laptops that come and go from the network, as well how to employ offsite (cloud) backups for added protection.

Understanding how Retrospect works

This section introduces certain terminology that is described in the Glossary of Terms found in the Retrospect User's Guide. Such terms are highlighted with **bold** emphasis.

Retrospect is *file-based* backup software, as opposed to drive imaging software that copies blocks or sectors on a hard disk drive. While an image- or block-based tool, such as Mac OS X's Disk Utility application, can be very fast at recovering an entire disk partition or providing a point-in-time view into a partition's contents, a file-based backup provides several distinct advantages. For example, it's possible to quickly search multiple backup archives and recover multiple versions of a document, while still being able to provide a complete restore of an entire disk. It's also trivial to restore files backed up from a Windows PC's NTFS volume to a Mac, because the volume format doesn't matter for file-based recovery.

*Retrospect employs several technologies to ensure that it only and always backs up the minimum number of files necessary to restore whatever **volume** it is protecting.*

That statement says a lot about Retrospect's design. Understanding this basic principle provides a good deal of insight into why Retrospect works the way it does and what we feel is the most important aspect of a backup system: reliability.

Retrospect's key components and terminology

The following components are the building blocks of every **backup**, **copy**, or **restore** operation in Retrospect:

- **Source** – a **volume** attached to the Retrospect server or a computer on the network running the Retrospect **Client** software that contains files and folders to be protected; can also be a network share
- **Destination** – the target for the files being copied or backed up; can be a **Media Set** (see below) for backups or another volume for a copy operation
- **Rule** – a built-in or user-defined set of conditions used to filter what gets copied during an operation
- **Tag** – a label applied to a source for the purpose of logical or physical grouping and abstraction; examples of tags are: *Laptops*, *Accounting Dept.*, *2nd Floor*
- **Schedule** – the time or times during which Retrospect will run its operations
- **Script** – a saved procedure that defines the settings for an operation; scripts can be run manually or automated with one or more schedules

Two important components work together to store and track backed-up data:

- **Media Set** – a logical container made up of one or more data storage mediums (including hard disks, network shares, digital tapes, USB thumb drives, etc.) that hold backed up files
- **Catalog** – a database that tracks all of the files stored in a Media Set, as well as point-in-time information about each of the sources that have been backed up; each Media Set has its own Catalog

The following example ties it all together: *A backup script runs nightly at 8:00 PM to protect all sources tagged as "Desktops," and copies any files that match the "User Files And Settings" rule to a destination Media Set comprised of volumes on a Thunderbolt RAID.*

Smart Incremental backup technology

Retrospect doesn't use the traditional concepts of *full*, *incremental*, and *differential* backups. These are outdated modes of backup that have significant drawbacks with regard to performance, restore precision, or both. Traditional full backups are incredibly time consuming but offer precise restores. Incremental backups save time during the backup process, yet they restore unwanted files that were renamed, moved, or deleted since the last full backup.

Retrospect's Smart Incremental technology works by matching the files on a source volume with the files that are already stored on the destination Media Set. By doing so, Retrospect only needs to back up those files that have changed or been newly created since files were last written to the destination. This provides the same performance advantage of a traditional incremental backup.

In addition to copying just the files not already present on the destination, the Smart Incremental process saves a complete listing of all the files and folders that were present on the source at the time of the backup. This provides a point-in-time **snapshot** of the exact state of a volume, which Retrospect can later use as a guide to select the proper files for a restore.

Instant Scan technology

In order to build the list of new and changed files for any particular volume being backed up, previous versions of Retrospect had to spend time scanning all the files present to determine what had changed since the last backup to that Media Set. For a volume containing a million files, this scanning process could take 10-15 minutes—often far longer than the time needed to actually back up the new bits. When multiplied by 20 computers being backed up, that's several hours of time spent calculating what exactly to backup.

Retrospect now employs Instant Scan technology to significantly reduce the overall time spent determining what files need to be copied during a backup (and for certain types of restores). By using FSEvents on HFS+ volumes and the USN change journal for NTFS volumes, Retrospect and the Retrospect Client software are able to pre-scan disks and folders on Mac and Windows PCs, so that the list of new and changed files is ready to go when the backup starts. For a typical network environment where computers are backed up on a regular basis, *Instant Scan technology cuts overall backup times in half*.

Instant Scan technology is a major step forward in managing the ever increasing amount of data present on an organization's network and allows users to reduce costs and/or increase their level of protection:

- More computers (or data) can be protected during the backup window by each Retrospect host server.
- Critical data can be protected more often.

Data deduplication

Retrospect's Smart Incremental technology provides another benefit to users: data deduplication that saves time and storage space needed for backups.

Because Retrospect only backs up files that aren't already contained in the target Media Set, it doesn't bother storing multiple copies of files that are duplicated around the network. For example, if Retrospect encounters a Keynote presentation on Computer B that's exactly the same (in terms of name, size, creation and modification dates, etc.) as one it just backed up from Computer A, it doesn't need to copy that file again. Likewise, if Computer A and Computer B share most of the same 60 GB iTunes music library, Retrospect only needs to copy the matching files one time. The more data that is duplicated around a network, the more time and storage space Retrospect saves.

Using rules to further refine backups and restores

Retrospect allows the user to define **rules** that can be used to filter unimportant files, or to specifically select files that meet certain criteria. Several rules come pre-defined, such as the "All Files Except Cache Files" rule, which tells Retrospect to ignore temporary cache files like those created by Web browsers. Not backing up such files can save significant storage space, since it's typically unnecessary to restore them.

Rules are flexible and powerful. They can be used to prevent operating system files from being backed up, or to restore all Microsoft Excel files larger than 2 MB, which were modified in the past 60 days, and contain the word "Forecast" in their names.

Smart Restores

Retrospect's modular design allows the Smart Incremental technology to be used in reverse during a restore operation. By using the snapshot saved with the backup as a guide, Retrospect ensures that only those files not already on the volume being restored need to be written, while (depending on the restore options selected) files that don't belong are deleted. This gives Retrospect the restore precision of a traditional full backup, only Retrospect doesn't have to unnecessarily re-copy matching files that are already present on the destination.

Retrospect offers three types of restores:

- Restore an entire volume to a previous point in time
- Browse and restore selected files and folders from a specific point in time
- Search for and restore one or more versions of one or more files from any backup

These options cover a multitude of cases and provide incredible restore flexibility.

Multiple backups improve reliability

We've seen that Retrospect only needs to copy a file once to be able to restore it to any destination. But what happens if a Media Set is destroyed in a fire along with the computers for which it was storing backups? Not good!

To protect against data loss due to events such as theft, fire, and flood that can damage locally-stored backups just as easily as the original data, it's critical that more than one backup be created, and optimally, at least one copy be stored offsite for safe keeping.

By design, Retrospect's Smart Incremental backups select files to copy by comparing only against those already on the target destination. By simply targeting a different destination, it's possible to affect what files will be selected for backup. Targeting an empty Media Set will result in all files being backed up, even if they've already been backed up to a different Media Set.

Retrospect tracks each Media Set independently.

The following table shows what files will be copied for each destination:

Media Set Name	Last Written To	What Gets Copied
Backup A	Yesterday	Unique files changed or added since yesterday
Backup B	1 Week Ago	Unique files changed or added in the past week
Backup C	Never (empty)	All unique files

Unlike other backup software that chooses files based solely on dates or archive flags, Retrospect ensures that each destination is kept complete. By doing so, only one set is needed to perform a restore, and individual Media Sets may be retired as archives without impacting other Media Sets.

This is a great example of how Retrospect is designed to do the right thing. Give it an empty destination, and Retrospect will copy everything to it. Give it a set that already contains data, and Retrospect will copy just what's needed. The user doesn't have to change the type of backup to ensure that the right files get copied.

Data Verification

Retrospect's modular design and focus on reliability is readily apparent in how it verifies that backed up files can be restored successfully.

After Retrospect completes backing up a source volume, it reads back the files it copied and compares them to the originals. But the way Retrospect does this is unique. To read back the files it copied, Retrospect actually uses its restore module. Then, instead of overwriting the original files with the backed-up copies, it simply compares them. This method has the benefit of testing the actual path that files will take when being restored, so in essence, file restorability is verified along with file integrity.

A note on performance testing: Retrospect defaults to thorough file verification, which can almost double the time that a backup operation takes. Retrospect's Media Verification option uses MD5 checksum digests generated during the backup and written to the media to quickly verify data integrity. This method doesn't test the entire restore path during a verification, but it is several times faster. Many other backup applications default to no verification, sacrificing reliability for performance.

Common Uses for Retrospect

This section provides guidance for using Retrospect in a couple of scenarios that illustrate how the software meets common usage needs. In addition to its focus on reliability, Retrospect is designed to be flexible for a variety of needs not touched on here.

A note about scripts: A single Retrospect script can control the backups of an entire network of computers, with multiple sources, multiple destinations, and multiple schedules. However, each script's options, such as which rule to use, whether to use encryption, and on what schedule to run, will apply to all of that script's sources. If you want to back up desktops only at night and laptops 24 hours a day, you will need two scripts to do that. If you want to encrypt the backups for the accounting system and customer database, you will need to use a separate script than for non-encrypted backups of less sensitive data.

Nightly onsite backups with weekly offsite copies

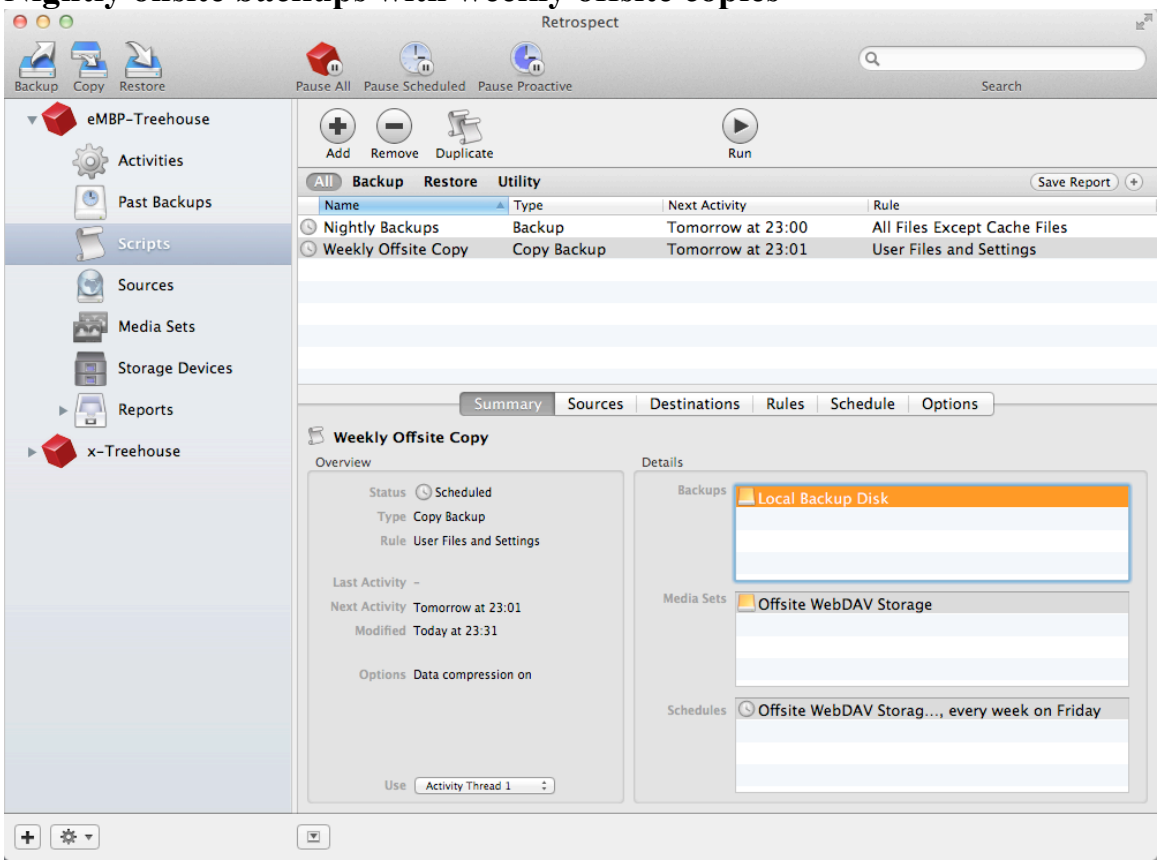

With just two scripts, Retrospect can protect an entire network of computers using both onsite and cloud-based storage. This exercise also provides a good example of how powerful Retrospect's rules can be.

Script #1 – Nightly backup to local disk

1. Create a Backup script using Scripts→Add.

2. Select the sources to be protected.

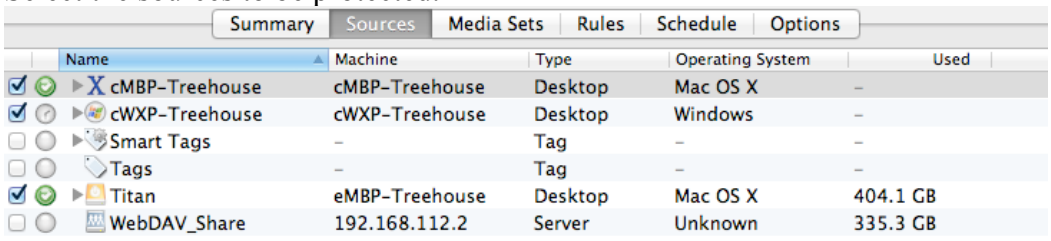

3. Select the local disk destination (define a Media Set if you have not already done so by using Media Sets→Add).

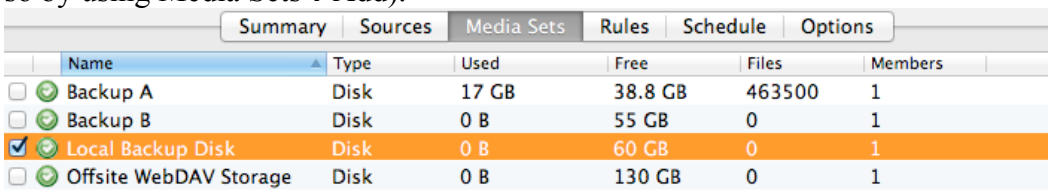

4. Select the All Files Except Cache Files rule.
5. Add a daily schedule, selecting the local Media Set as the destination.

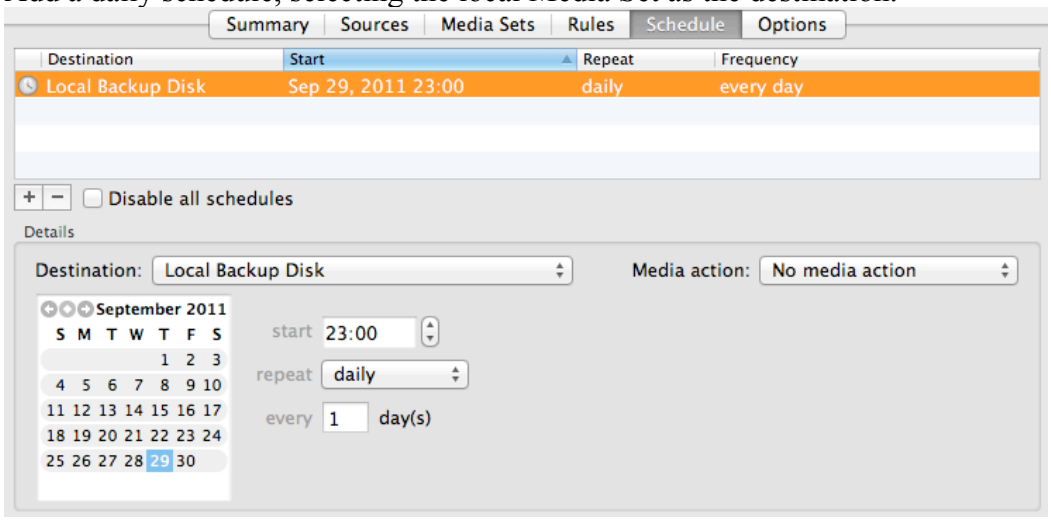

6. Return to the Summary tab and use the Activity Thread pop-up menu at the bottom of the Overview section to select Activity Thread 1. (This ensures the weekly offsite backup will run immediately after this one by setting it to run one

minute later on the same Activity Thread.)

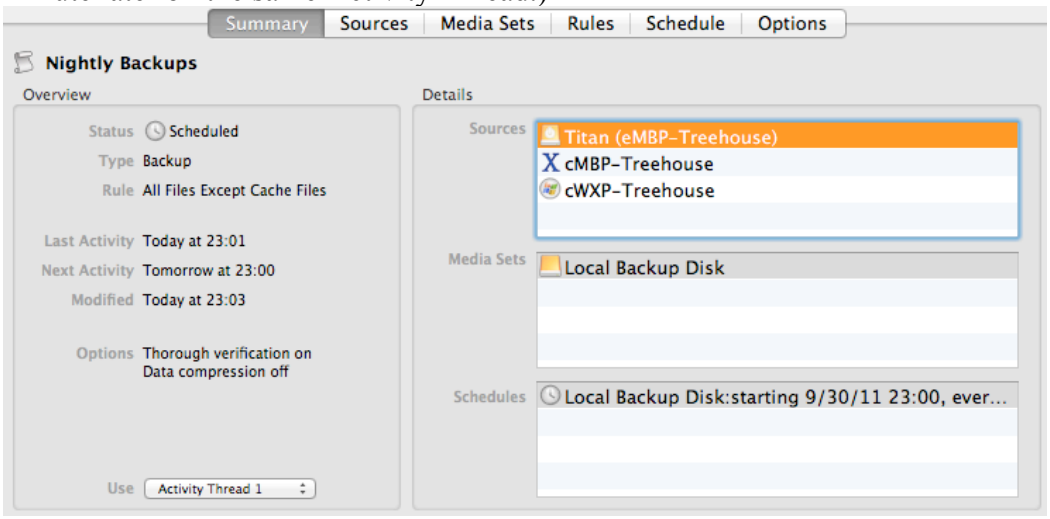

Script #2 – Weekly copy to cloud-based storage

1. Create a Copy Backup script using Scripts→Add. (A note about Copy Backup scripts: The source for a Copy Backup script is another Media Set. This allows Retrospect to transfer data from one Media Set to another, a useful method of copying files that have already been backed up without having to touch the original source volumes again.)
2. On the Sources tab, select "Copy most recent backups for each source" from the pop-up menu above the list of Media Sets, and select the Media Set that's used to store the nightly backups from Script #1. This tells Retrospect not to bother transferring files that may have been deleted from the source volumes earlier in the week.

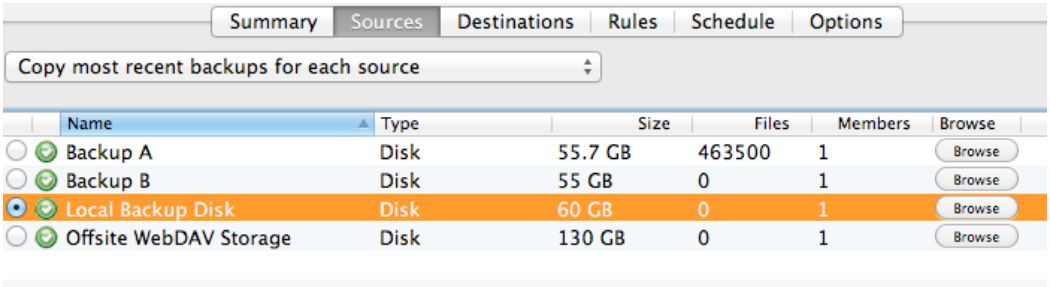

3. Select the offsite Media Set for the destination. (Define a Media Set using WebDAV or other offsite storage if you have not already done so by using Media Sets→Add.)
4. On the Rules tab, select the User Files And Settings rule. This tells Retrospect to only add files from the Users directories on Mac OS, Windows, and Linux systems, thereby ignoring operating system files and applications. This will keep our offsite backups lean and mean.

5. Add a weekly schedule to backup to the offsite Media Set on Fridays, and set the time to one minute after the Backup Script we created for Script #1.

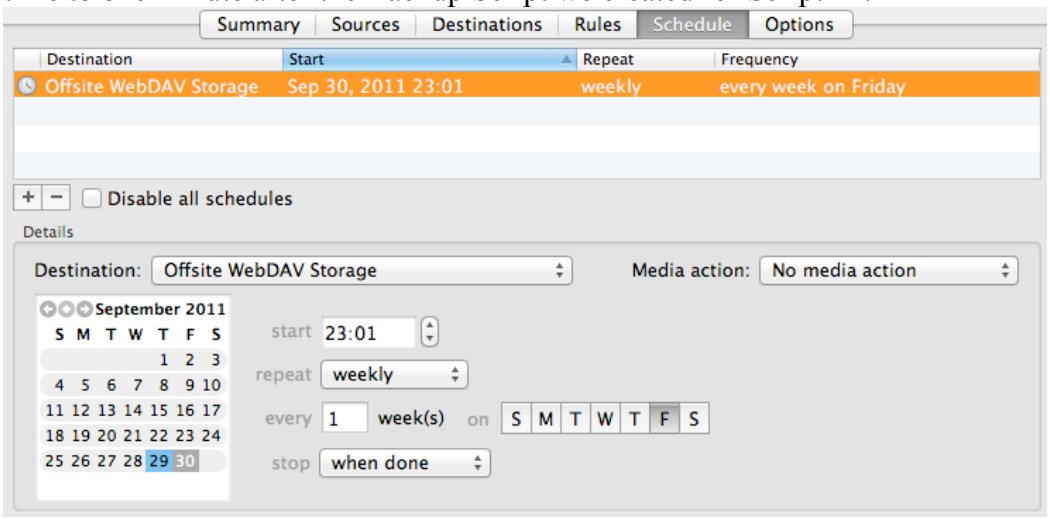

6. On the Options tab, turn on Data Compression. This will save time when Retrospect transfers data to the cloud storage.
7. Return to the Summary tab and use the Activity Thread pop-up menu at the bottom of the Overview section to select Activity Thread 1.

That's all there is to it! Retrospect will now back up nightly to local storage. Then on Friday, as soon as the nightly backups are done, Retrospect will transfer the latest backups to offsite storage.

Using Proactive Backup to protect laptops

Laptops represent a particular challenge to backup administrators. They often come and go from the network, which makes it difficult to protect them on a predictable schedule. Retrospect's answer to this problem is the **Proactive Backup** script, which uses dynamic prioritization, scheduling, and resource management to ensure that systems are effectively protected when they become available.

Part 1 – Defining the laptops

1. Go to the Sources view and click on the Tags tab.
2. Create a new tag by clicking on the + button at the bottom of the Tags list. Name the tag "Laptops."
3. Highlight one or more laptop clients that you have logged in and check the box next to the Laptops tag. (If you don't have any Retrospect Client systems to test with, you may define any source volume as a laptop for this exercise.) Later, any additional systems you tag as a Laptop will be automatically added to your Proactive Backup script for protecting laptops.

Part 2 – Create the Proactive Backup script

1. Create a Proactive Backup script using Scripts→Add.

2. Click on the Sources tab and check the Laptops tag.

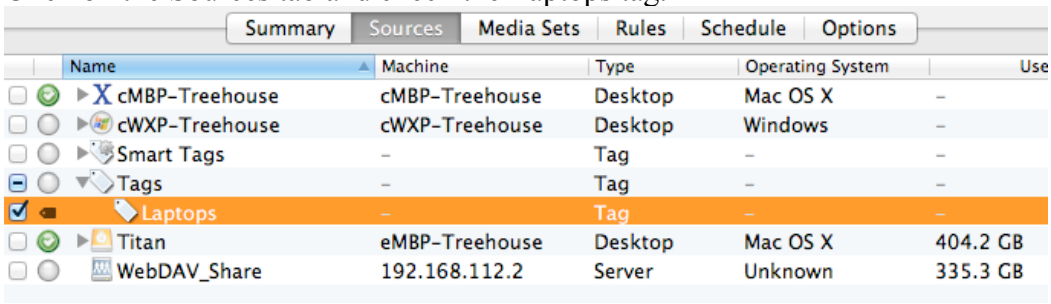

3. Skip the Media Sets tab for now, because Proactive Backup will start running as soon as it has a valid source, destination Media Set, and schedule.
4. Select the All Files Except Cache Files rule.
5. On the Schedule tab, add a schedule that backs up sources every day and runs all day. (Unlike a regular schedule that starts at a specific time and runs through the list of sources once, Proactive Backup will keep looking for valid sources as long as it has an active schedule. If we were protecting desktops, we might instead set one schedule to run during the week from 8:00 PM to 7:00 AM and another to run all weekend long.)

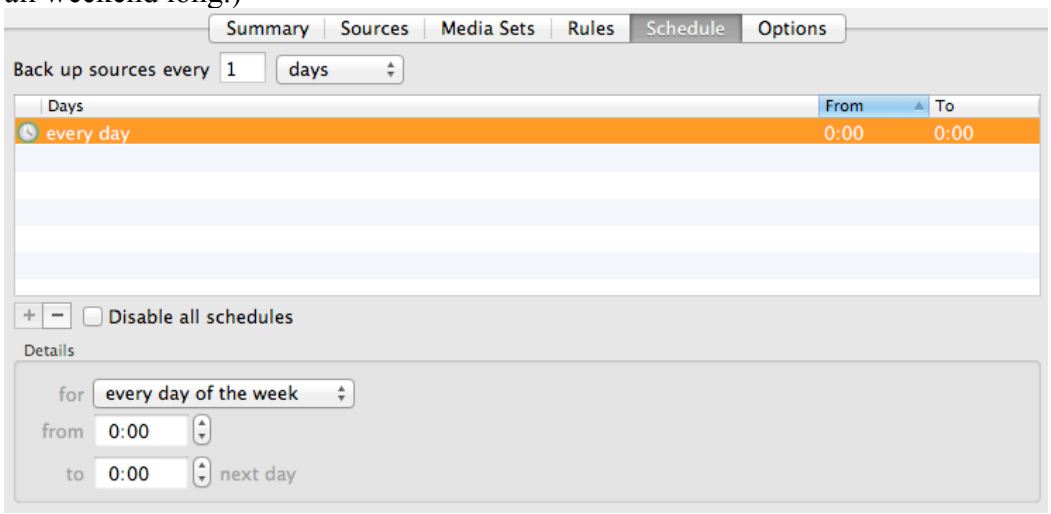

6. Finally, select one or more Media Sets to be targets for our laptop backups, and Proactive Backup will begin running, looking for laptops to back up.

If multiple destinations are selected in step 6, Proactive Backup will automatically alternate between them. If one's not available, it will keep running to the available Media Set(s). This provides flexibility to swap out the Media Sets (e.g., to take one offsite for safekeeping and protection from fire, flood, theft) without needing to adhere to a rigid schedule.

Where to go for additional information

The Retrospect User's Guide and Retrospect User's Guide addendum contain useful information for beginners and Retrospect veterans alike. The first two chapters of the Retrospect User's Guide cover getting started and how Retrospect works in more detail, while the Retrospect User's Guide Addendum covers the features new to version 10.

Additional resources, including user forums, a searchable knowledgebase, and support contact information, are available online at www.retrospect.com.