



Data Rescue User's Guide

PROSOFT
engineering, inc.

Licensing and Copyright Information

©2002-2007 Prosoft Engineering, Inc. All rights reserved.

303 Ray St., Pleasanton, CA 94566

This document is protected by copyright. No part of this document may be reproduced in any form by any means without prior written authorization of Prosoft and its licensors, if any.

The information described in this document may be protected by one or more U.S. patents, foreign patents, or pending applications.

TRADEMARKS

Apple®, AppleShare®, AppleTalk®, Macintosh®, iMac®, iBook®, and Mac® are registered trademarks of Apple Computer, Inc., registered in the U.S. and other countries.

Iomega®, Zip®, and Jaz® are registered trademarks of Iomega Corporation, Inc., registered in the U.S. and other countries.

Word®, Excel®, and PowerPoint® are registered trademarks of Microsoft Corporation, Inc.

All product names mentioned herein are the trademarks of their respective owners.

THIS DOCUMENT IS PROVIDED “AS IS” WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESS OR IMPLIED, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE, OR NON-INFRINGEMENT.

THIS DOCUMENT COULD INCLUDE TECHNICAL INACCURACIES OR TYPOGRAPHICAL ERRORS. CHANGES ARE PERIODICALLY ADDED TO THE INFORMATION HEREIN; THESE CHANGES WILL BE INCORPORATED IN NEW EDITIONS OF THE DOCUMENT. PROSOFT ENGINEERING, INC. MAY MAKE IMPROVEMENTS AND/OR CHANGES IN THE PRODUCT(S) AND/OR THE PROGRAM(S) DESCRIBED IN THIS DOCUMENT AT ANY TIME.

8500-0005D

Table of Contents

1	Introduction.....	8
1.1	Welcome	8
1.2	What Makes Data Rescue Different?.....	8
1.3	Latest version of the software.....	8
1.4	Program Setup	8
1.5	Uninstalling the Software	8
1.6	Launching and Activating the Software.....	9
1.7	Technical support.....	9
1.8	Contact Prosoft Engineering	9
1.9	Data Rescue Demo – Risk-Free Trial.....	9
2	About This User Guide.....	10
2.1	User Guide Organization.....	10
2.2	Conventions	10
2.3	User Experience Level.....	11
2.4	System Requirements	11
2.4.1	Hardware	11
2.4.2	Software	11
2.5	General Usage	12
2.6	CAUTION: Drives with Suspected Hardware Problems	12
3	Installing and Starting Data Rescue.....	13
3.1	Install / Uninstall Data Rescue	13
3.1.1	To install Data Rescue.....	13
3.2	To uninstall Data Rescue	13
3.3	Before you start Data Rescue.....	13
3.3.1	Note on Data Rescue temporary file space	13
3.4	Starting Data Rescue	14
4	Using Assistant Mode.....	16
4.1	Selecting a drive or volume	16
4.2	Selecting the working (temporary) volume	17
4.3	Selecting the scan method.....	18
4.4	Scanning the media	19
4.5	Scan results.....	21
4.6	Recovering files and folders.....	23
4.6.1	Before doing a large recovery	24
4.7	Double-click opening of individual files	24
4.8	Searching the scan results	24
5	Using Expert Mode	27
5.1	Selecting the source media and scan mode	27
5.2	Scan results.....	28
5.3	Analyzing Media.....	29

5.4	Cloning Media	32
6	Preferences	33
6.1	User Interface Preferences	34
6.1.1	Report level	34
6.1.2	Open report on any error	34
6.2	Scan Engine Preferences	34
6.2.1	Analyze preferences	35
6.2.2	Scanning preferences	35
6.2.3	Recovery preferences	35
6.2.4	File Module preferences	36
7	After the Thorough Scan	39
7.1	Saving a Scan File	39
7.2	Allocation Blocks Layout Dialog	40
7.3	Recovering Files from Digital Cameras and Music Players	41
8	Concepts	41
8.1	Selecting a Volume or Drive to Scan	42
8.1.1	Drives versus Volumes	42
8.2	File Systems and Catalogs	43
8.3	Catalog Scans	45
8.4	Content Scan	46
8.5	Deleted Files Scan	47
8.6	Scan File	47
8.7	Allocation Blocks and Media Blocks	48
8.8	Locating Your File After a Scan, Revisited	49
8.8.1	Locating a file found by a Catalog scan	49
8.8.2	Locating a file found by a Content scan	50
9	Using Data Rescue – More Details	50
9.1	Scan Time	50
9.2	Block Level Tools	51
10	Troubleshooting and Frequently Asked Questions	53
10.1	Data Rescue seems to freeze. What should I do?	53
10.2	Why doesn't my drive or volume show up?	53
10.3	Why are some files listed twice?	53
10.4	Why do my found files total more than the size of the original drive?	54
10.5	Why is there a delay when I click the mark checkboxes?	54
10.6	Does Data Rescue handle RAID volumes?	54
10.7	I recovered files, but they won't open. What's wrong?	55
10.7.1	Files recovered by catalog (i.e. not from the CBR folder)	55
10.7.2	Files recovered by content (i.e. from the CBR folder)	55
10.8	Can I recover files that I deleted accidentally?	55
10.9	What are Data Rescue's general content scan capabilities?	55
10.10	Specifically what types of files can Data Rescue II find by content?	56

11	Glossary	58
-----------	-----------------------	-----------

Table of Figures

Figure 1: Data Rescue Authorization Screen.....	14
Figure 2: Activation Window	15
Figure 3: Assistant Mode Introduction Screen	16
Figure 4: Media Selection Screen	16
Figure 5. Working Volume Selection Screen	17
Figure 6. Scan Method Selection Screen.....	18
Figure 7. Scan Confirmation Screen	19
Figure 8. Scan Progress	20
Figure 9. Scan Results Summary	21
Figure 10. File Browser screen	22
Figure 11. Selecting things for recovery	23
Figure 12. Choosing a folder to recover files in	24
Figure 13. Searching for files or folders by name	25
Figure 14. The Find dialog	25
Figure 15. Search results	26
Figure 16. Expert Mode Screen	27
Figure 17. Selecting a location for working storage	28
Figure 18. Scan Results	29
Figure 19. Analyzing a Volume or Drive.....	30
Figure 20. Analyze stride and read buffer sizes.....	31
Figure 21. Analyze Results.....	31
Figure 22. A drive with read problems	32
Figure 23. User Interface Preferences	34
Figure 24. Scan Engine Preferences.....	35
Figure 25. Top level file module preferences.....	36
Figure 26. Non-editable preferences	37
Figure 27. The Text file modules are disabled by default.....	38
Figure 28. Particular file module preferences.....	39
Figure 29. Changing an editable inherited value.	39
Figure 30. Allocation Block Layout Dialog	41
Figure 31. Drives vs Volumes	42

Figure 32. Media selection	43
Figure 33: File System Data Storage and Catalog Structures	44
Figure 34: Media and Allocation Blocks	48
Figure 35: Block Content Window	52

1 Introduction

1.1 Welcome

Thank you for purchasing Data Rescue.

Prosoft Engineering's Data Rescue is an easy-to-use utility to help you recover files and folders from crashed or corrupted hard disks, floppy disks, or removable cartridges. Recovered data is saved to another medium, leaving the original disk untouched.

1.2 What Makes Data Rescue Different?

The main objective of Data Rescue is that the strategy for finding lost files is different and better from the strategies used by other tools. Not only will it find your data in situations where the other tools fail, but also it will recover them in a much better condition, restoring icons, dates and even folder hierarchy.

Data Rescue is the only utility that focuses on data recovery as opposed to hard drive repair. Your entire hard drive is examined for file content, ignoring the possibly corrupt directory entries. This is one of the reasons that Data Rescue is more successful than "repair" utilities on the market. The data on the drive is painstakingly evaluated and files are meticulously re-assembled and stored in a safe location. Either with intense manual evaluation, or with the help of powerful software tools, this is the process that data recovery professionals routinely use to restore their client's data.

Why? Because this is the only reliable way to safely recover your data.

1.3 Latest version of the software

Check our web site "<http://www.prosofteng.com>" to find out if you have the latest version of this software.

If you do not have the latest version of Data Rescue, download it from the web site using your serial number.

Note: serial numbers are compatible with all versions having the same main version number: a serial number of version 11.0 is useable with versions 11.1, 11.2, etc. Data Rescue II will not accept serial numbers from Data Rescue X or earlier.

1.4 Program Setup

Data Rescue installation is very easy. Simply drag the application folder <Data Rescue Folder> from the CD-ROM (or from the mounted disk image if you downloaded it from the web site) and drop in your Applications folder. The latest version of Data Rescue may be downloaded from our web site <http://www.prosofteng.com>.

1.5 Uninstalling the Software

To remove this software from your hard disk:

- Quit Data Rescue
- Move the application folder <Data Rescue Folder> to the Trash.

1.6 Launching and Activating the Software

Double-click the application icon to run Data Rescue. When first launching, you will be prompted to enter your serial number (sent to you by email, or on a sticker on your CD sleeve).

This serial number will activate all of the features of Data Rescue. Please, keep this number in a safe place (for upgrades and technical support).

1.7 Technical support

Our technical support is free and can be accessed by e-mail, fax or phone. Please have your serial number and the software version number ready before calling.

1.8 Contact Prosoft Engineering

If you have comments about this product, problems or questions about this user guide or with our web site - or if you are interested in a site license - please contact PROSOFT ENGINEERING (specify the version and serial numbers of your copy):

- by e-mail address: support@prosofteng.com
- by phone or fax: Tel.: 925-426-6306
 Fax: 925-426-6309
- by mail:
 PROSOFT ENGINEERING
 303 Ray St.
 Pleasanton, CA 94566

1.9 Data Rescue Demo – Risk-Free Trial

You can use an unregistered copy (demo mode) of Data Rescue to do a full scan. This full scan works exactly as the fully licensed version would, which will allow you to see all the files that are available for recovery. However, the demo mode only allows you to recover one single file, no larger than 5MB. Once you've confirmed that Data Rescue can indeed see the files that you wish to recover you can simply purchase a copy of Data Rescue from www.prosofteng.com and a serial number will be immediately sent to you via email. Because Data Rescue will not try and modify your disk, you can test out the software without the worry of "making things worse" like you can with other disk repair utilities.

2 About This User Guide

This guide contains both instructions on how to use Data Rescue and reference information on the concepts one needs to know about files and how the files are recovered. It is a “how to” user guide, so the reference material is for the most part in chapters separate from instructional chapters. If you are trying to look up a concept, use the glossary and the index.

In addition, it is important to understand that Data Rescue is a file recovery utility, not a file repair utility. This guide endeavors to help distinguish the difference between recovery and repair.

2.1 User Guide Organization

This guide is divided into the following chapters:

Overview: Provides a brief introduction to Data Rescue, describes system requirements, and explains conventions used in this guide.

Installing and Starting Data Rescue: This chapter provides instructions for quickly installing and starting Data Rescue for the first time.

Using Data Rescue: This chapter walks the user through a simplified procedure for volume selection, scan selection, and file recovery with a minimum of technical information. This procedure should suffice for most users. If this does not result in recovery of the desired files, the user can then refer to later chapters for a more in-depth presentation. Also included is a short special section for users wishing to recover images and music files from plug-in digital media cards.

Concepts: This chapter explains the concepts needed to understand and use the more advanced Data Rescue capabilities and options.

Using Advanced Features: Provides instructions on using Data Rescue options, preferences, and other capabilities not covered in the earlier simplified procedure.

Troubleshooting: Provides important troubleshooting information.

Glossary: Provides short definitions of important or common Data Rescue terms.

2.2 Conventions

The following conventions are used in this guide:

<i>Convention</i>	<i>Description</i>
Text in bold type	Text for GUI-based commands that are often part of a menu or dialog, is printed in bold type: Click Start Rescue to begin.
Text in <code>monospace</code> type	Text used to represent names of

	directories, files, and paths is printed in a monospaced font.
Monospace italic text in angle brackets (<>)	Text used as a placeholder for variables is printed in an italicized monospaced font and place in angle brackets: <pre><logicalblock> = (<allocationblock> x <allocationblocksize>) + <filespaceoffset></pre>
Text in <i>italic</i> type.	Text of all words that can be found defined in the Glossary is printed in italic type: <i>Allocation block</i>
Text in bold type, offset by a “►”	Text used to introduce a set of step-by-step procedures is printed in bold, and offset by an arrow (►): ► ► To install Data Rescue 1. Double-click the

2.3 User Experience Level

This guide assumes that the reader is familiar with the Macintosh computer, including how to use the mouse, select menu items, and so on. If you are not familiar with how to perform basic functions like these on a Macintosh, please read the documentation that came with your computer.

2.4 System Requirements

2.4.1 Hardware

Rescue II is designed to run on Macintosh computers capable of running Mac OS X 10.2.4 and later. Compatible Hardware includes all computers developed in the last three years from the iMac to the G3, G4, G5, iBook, and G3/G4 PowerBooks. If you have a machine that is not capable, you can contact Prosoft Engineering for an older version of Data Rescue that will work with your system. Data Rescue needs at least 128 MB of RAM to run but at least 256 MB is recommended.

2.4.2 Software

Data Rescue II requires Mac OS X version 10.2.4 or later to and includes support for Mac OS X system with HFS or HFS+ partitions. It also includes limited support for file recovery from non-HFS/HFS+ disks and media (see the later section on Content Based Recovery).

2.5 General Usage

Data Rescue was designed to be used in a wide range of situations. Specifically, Data Rescue can handle:

- Any type of media which appears as a disk device with 512 byte sectors (floppy disk, hard disk, zip, jazz, CF camera cards, etc.)
- Volumes which can't be mounted, even if the driver is damaged.
- Disks with a damaged partition map.
- SCSI, IDE (also called ATA disks) FireWire and USB disks.
- HFS and HFS+ ("Mac OS Extended").
- Limited support for non-HFS/HFS+ filesystems
- Large volumes (4 GB and larger).
- Password-protected volumes (if data is not scrambled).
- Non-Roman-script (e.g. Japanese) file names.
- Recovery of fragmented files.

However, you should be aware that Data Rescue will not be useful in all situations. A demo version of Data Rescue may be downloaded from <http://www.prosofteng.com> to test whether it can find your files.

Also note that you need disk space on another device for Data Rescue's temporary working storage, and to recover your data. This must be a separate drive from the drive you are trying to recover from, in order to avoid altering that drive. Data Rescue offers the most extensive, safe recovery of your important data. Because your corrupt drive is not stable, Data Rescue will not allow you to try and save files to that same drive, nor will Data Rescue try and "repair" that corrupt drive (doing so could actually cause more corruption and data loss).

2.6 CAUTION: Drives with Suspected Hardware Problems

A word of caution before you use Data Rescue on a drive where you suspect a hardware problem: Even though Data Rescue exercises caution to never damage your drive's data by writing to it, there are certain types of drive hardware failures which by their nature, can cause additional damage to the media and data just by continuing to use the drive. One example of this kind of damage is a crashed disk head, which could potentially cause additional media scratching and data loss with continued use. One clue to this type of damage may be unusual mechanical sounds emanating from the drive. If you suspect this type of damage, and the data on the disk is extremely critical, you may wish to send your drive to a professional drive recovery service with clean-room facilities.

That being said, the cost of professional recovery services can run into the thousands of dollars. If you suspect hardware damage, but need to avoid the expense of the professional services, you can minimize the chances of further damage to your drive by first cloning it to a spare drive, then do the recovery on the clone. See the section later on using the Clone function in Expert mode for details.

3 Installing and Starting Data Rescue

This chapter provides basic instructions on installing and starting Data Rescue. If you are using Data Rescue for the first time, you should read this chapter in detail.

3.1 Install / Uninstall Data Rescue

If Data Rescue was shipped to you on a bootable CD, it is not necessary to install Data Rescue. If you plan on installing Data Rescue to a disk, you should NOT install it on the disk you are trying to recover files from. If you are trying to recover files from a disk that you normally use as your boot drive, you will need to either boot from a Data Rescue bootable CD, or else set up a different drive to boot from, and install Data Rescue on that drive, and boot from it. The following instructions are for users with a downloaded version of Data Rescue. You should always keep a copy of Data Rescue on an emergency disk (a system disk if possible).

3.1.1 To install Data Rescue

From the location to which you downloaded Data Rescue, double-click on the **Data Rescue** image file to mount it on your computer.

1. Drag the contents of the Data Rescue drive (the mounted disk image) to your applications folder.

A Data Rescue folder is created in your applications folder. The Data Rescue folder contains the Data Rescue Application, Release Notes, and User's Guide.

3.2 To uninstall Data Rescue

If at any time Data Rescue or any of its parts needs to be uninstalled, simply drag its folder to the trash.

3.3 Before you start Data Rescue

After you have installed Data Rescue (if a downloaded version) or inserted the Data Rescue CD into your system's CD-ROM drive (if on a bootable CD), you are ready to begin using Data Rescue.

Data Rescue requires access to a safe location for its temporary files and to store your recovered files. In the interests of performance, your best option is to plug in an external, freshly formatted hard drive through the SCSI or FireWire connection at the back of your computer, or use a separate internal drive. Make sure the capacity of this drive is great enough to contain all the files you plan to recover from your damaged hard drive, plus space for Data Rescue's temporary files. It is not recommended to try and use a slow drive such as a network drive, USB 1.0 drive, etc. for the temporary storage because doing so will have a severe performance impact on Data Rescue.

3.3.1 Note on Data Rescue temporary file space

As Data Rescue scans and processes your volume, it needs access to temporary storage for its internal temporary working data in a directory named **DataRescueTemp** at the top level of the drive you specified for temporary storage. (Note: these are internal data,

unrelated to the files that you will recover from your volume.) The space required for these temporary files will vary depending on the size of the drive you are scanning, and how many files it has. To give you a general idea: a 60GB drive approximately 40% full, with approximately 600000 files on it required about 200Mbytes of free space for the temporary files. Free space equal to 1% of the size of the volume or drive you are scanning will often be plenty for Data Rescue's temporary storage needs. This directory is automatically removed when you quit Data Rescue.

3.4 Starting Data Rescue

Launch the application by double-clicking the Data Rescue icon in the Data Rescue folder. If Data Rescue has not been run before, an authorization window will come up first, as shown in Figure 1. Because Data Rescue needs unrestricted access to your disk devices, it must run with special privileges. This step authorizes Data Rescue to run with these privileges. Unlike older Data Rescue versions, Data Rescue II retains its privileges, so this step only needs to be done once. Enter the user name and password of an Administrative user. (If you are the only user of your system, then you are most likely the administrative user, and you can use your own name and password.) If you do not have an administration password for your computer, please contact your system administrator to login for you.




Figure 1: Data Rescue Authorization Screen

After clicking on the OK button, if Data Rescue has not yet been activated, an activation window will appear, as shown in Figure 2. Fill in your name and company name, if applicable. If you have a serial number, enter it as shown. If you are just trying Data Rescue before buying it, you may leave the serial number field blank. Click on the Read Agreement button to display the license agreement, and check the box indicating your agreement to its terms. Finally, click the Activate button (if you have a serial number), or the Demo button (if you don't). In Demo mode, Data Rescue will run, but you will only be able to recover one file. To do a full recovery, please purchase a licensed copy from www.prosofteng.com and a serial number will be immediately emailed to you.




Figure 2: Activation Window

Note: When Data Rescue is installed on a hard disk and activated with the serial number, this information is recorded in the user's preference files, and the activation will not need to be repeated unless the preference files are removed. When launched from a CD ROM, Data Rescue has no place to record the serial number, so activation with the serial number will be required on each launch.

Following this, the main window will appear in Assistant mode, as shown in Figure 3. Data Rescue II has a simplified Assistant mode, which will be described first, and an Expert mode with more features, which will be described later.

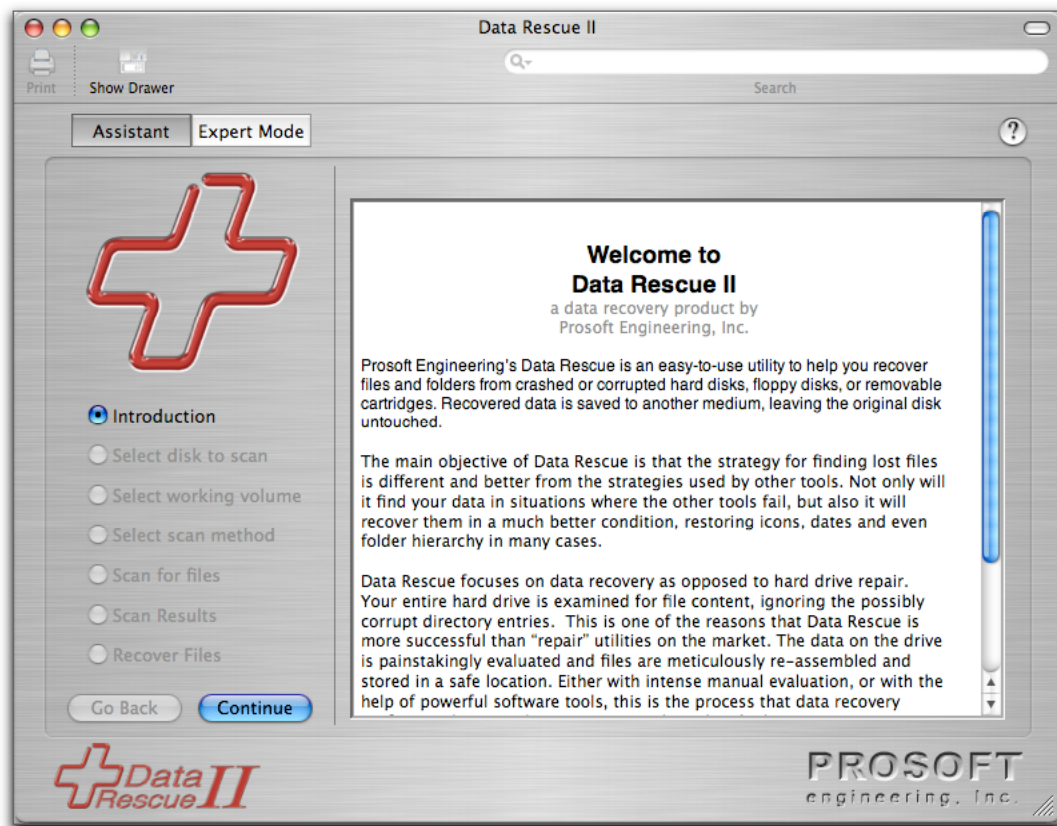


Figure 3: Assistant Mode Introduction Screen

4 Using Assistant Mode

Assistant mode is designed to make it as simple as possible to run Data Rescue and recover files by guiding the user through the steps. Following the Introduction screen comes the screen to select the media (drive or volume) to be scanned for files (Figure 4).

4.1 Selecting a drive or volume

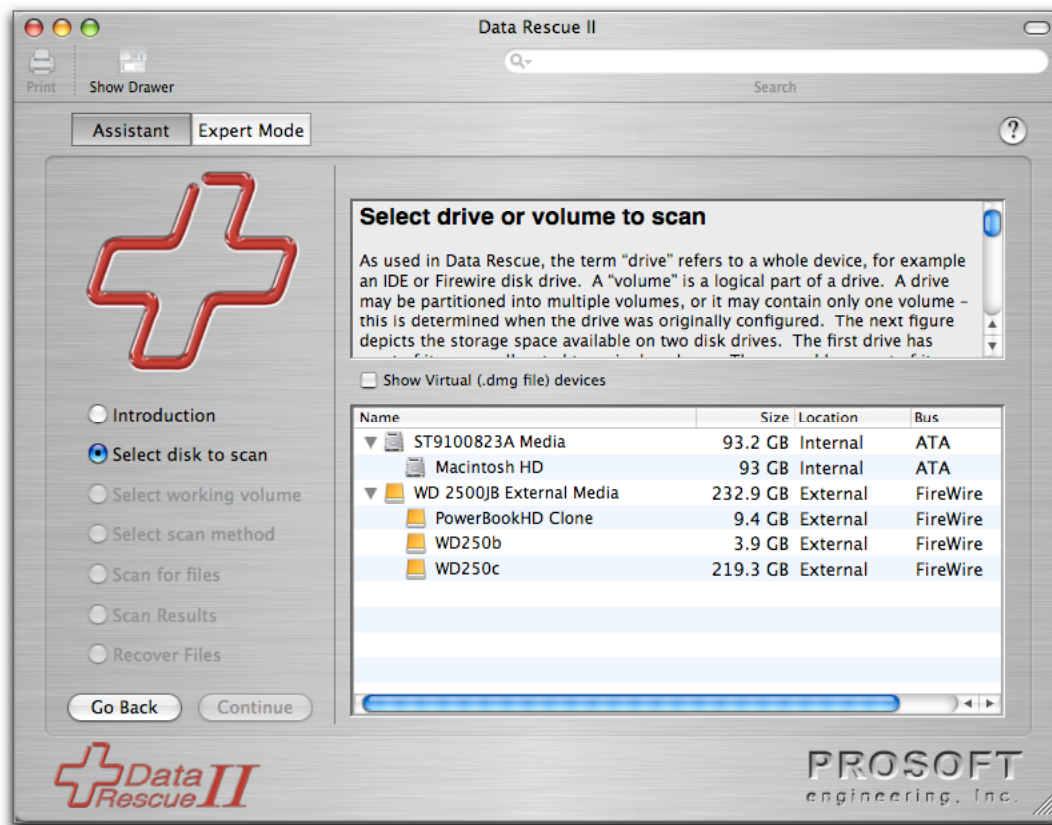


Figure 4: Media Selection Screen

This screen shows a two-level list of all the disk drives and volumes that Data Rescue can see on the system. The first level shows the hardware names of whole drives, for example “WD 2500JB External Media” in the above figure. The second level shows any volumes which were found on the drive, for example “PowerBookHD Clone” in the above figure. You may choose either a volume or a whole drive. We suggest that if the volume you want to scan is listed, that you choose it. If not, choose the whole drive that the volume is located on. If you can’t tell by the hardware name which disk is which, you can usually tell by looking at the size, bus, and location fields.

Before proceeding, if you have chosen a volume, check that the reported size seems about right. If a volume is damaged, it will sometimes show an incorrect size. In that case, you are better off scanning the whole drive. After choosing the source scan drive, click Continue to the next step.

4.2 Selecting the working (temporary) volume

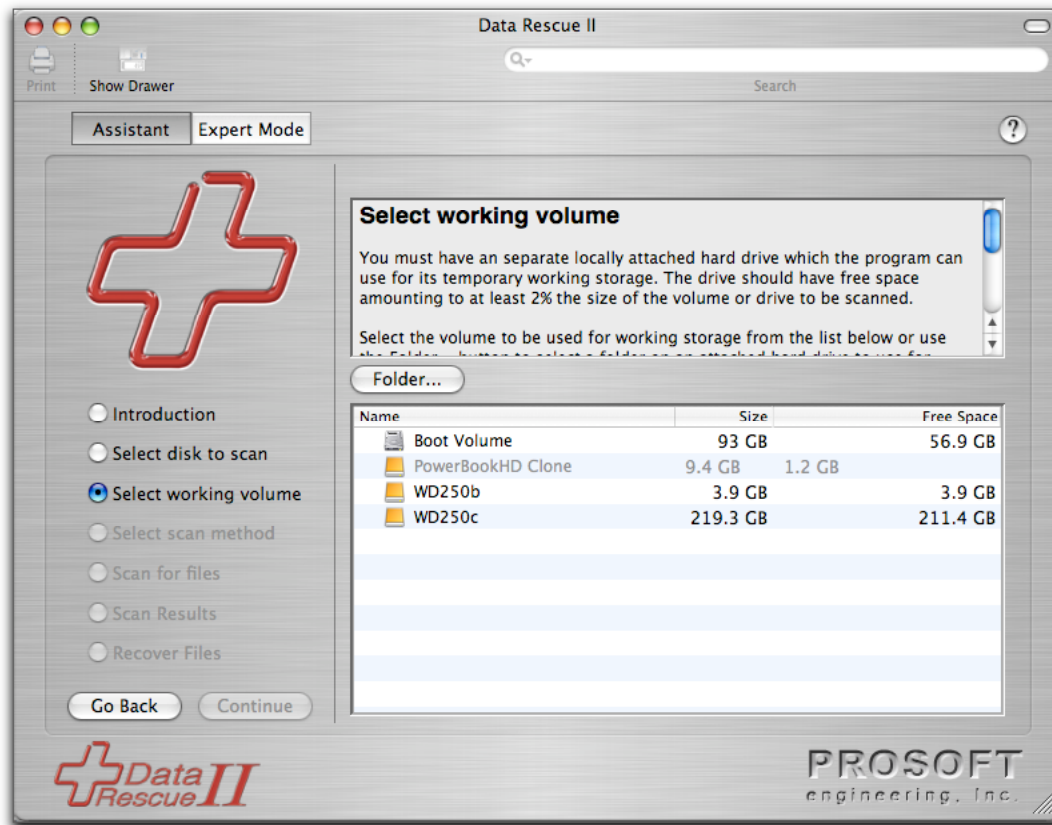


Figure 5. Working Volume Selection Screen

Data Rescue requires a temporary storage location to keep its working files while scanning. These can grow quite large when scanning a large drive or volume. It is not possible to calculate exactly in advance the temporary storage required, but normally 1% of the size of the media being scanned will be more than enough. It should rarely require more than 2%. The volume chosen for the temporary storage may not be the same as the volume being scanned. We recommend that the temporary volume be located on a different drive from the one being scanned, especially if you suspect any hardware problems with the scanned drive.

During scanning, Data Rescue will create a folder at the top level of the volume chosen as the working volume named DataRescueTemp. (On the boot volume, it is placed in /tmp rather in the root of the volume). This folder and its contents is normally removed automatically when you quit Data Rescue. If you prefer that this temporary folder go

someplace other than at the top level of the volume, you can use the Folder... button to choose a folder for it instead.

After choosing the working volume, click Continue to go to the next step.

4.3 Selecting the scan method

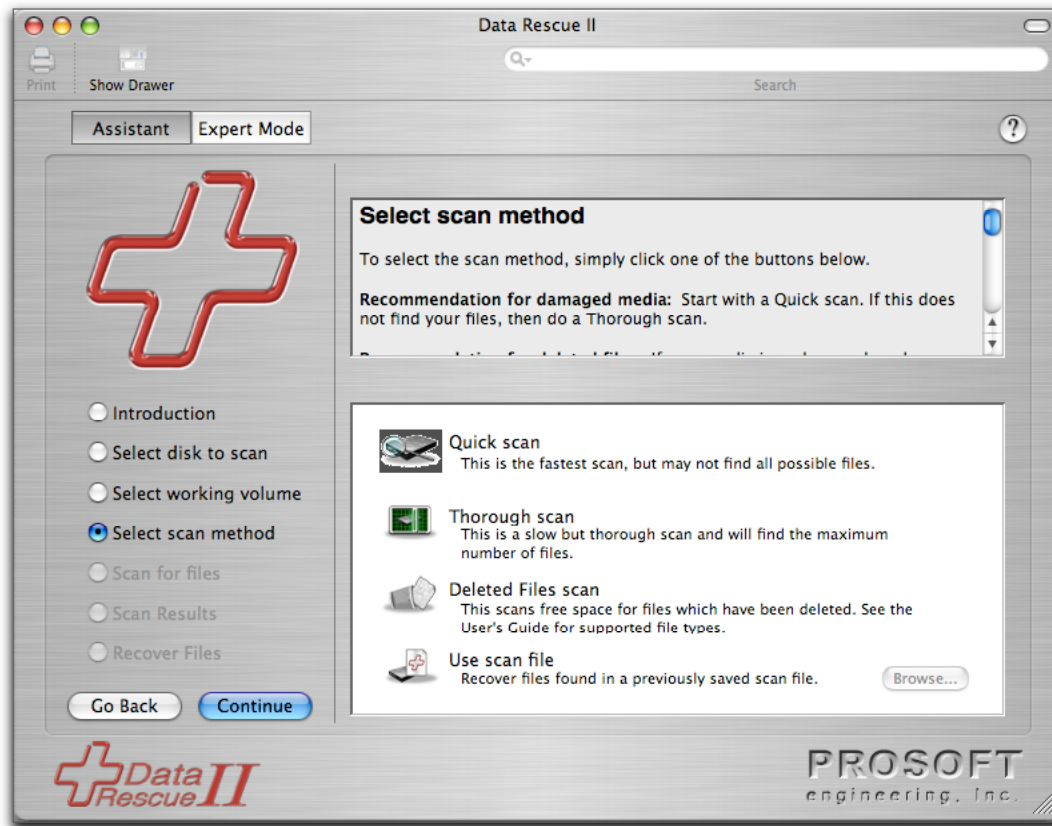


Figure 6. Scan Method Selection Screen

The next step is to select the scan method (Figure 6). The type of scan you need to do depends on your situation. The two most common situations are: 1. Finding accidentally deleted files, and 2. Finding files in a damaged file system.

If your file system and disk are thought to be in good shape, and you are trying to recover files which were accidentally deleted, you should choose the Deleted Files scan. You should be aware that in Mac OS X, when files are deleted, their names and folder locations are lost, so files found by this scan method will have a program-generated name, and you will have to open the recovered files to find the ones you want.

If you believe that your file system is damaged, then you should first try a Quick scan and see if that locates the files you want. If not, then go back and try a Thorough scan.

The fourth choice, Use Scan File, may be used if you have already done a scan and saved a scan file during a previous session with Data Rescue. Loading from the scan file is

primarily done to avoid the need to repeat the previous scan, and is most useful with a scan file saved from a previous Thorough scan.

After choosing the scan method, click Continue. This will display the confirmation screen (Figure 7).

4.4 Scanning the media

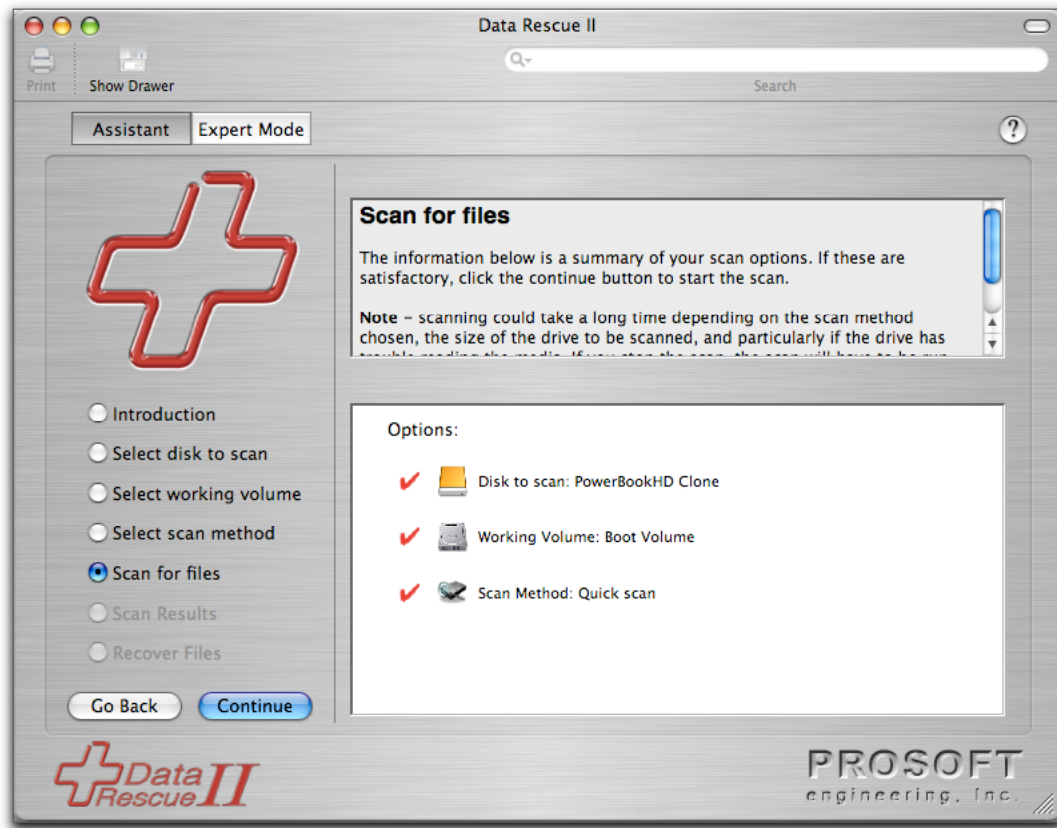


Figure 7. Scan Confirmation Screen

If the scan disk or volume, the working volume, and scan method are all correct, click Continue to start the scan. This will bring up a progress display, showing the various steps of processing (Figure 8).

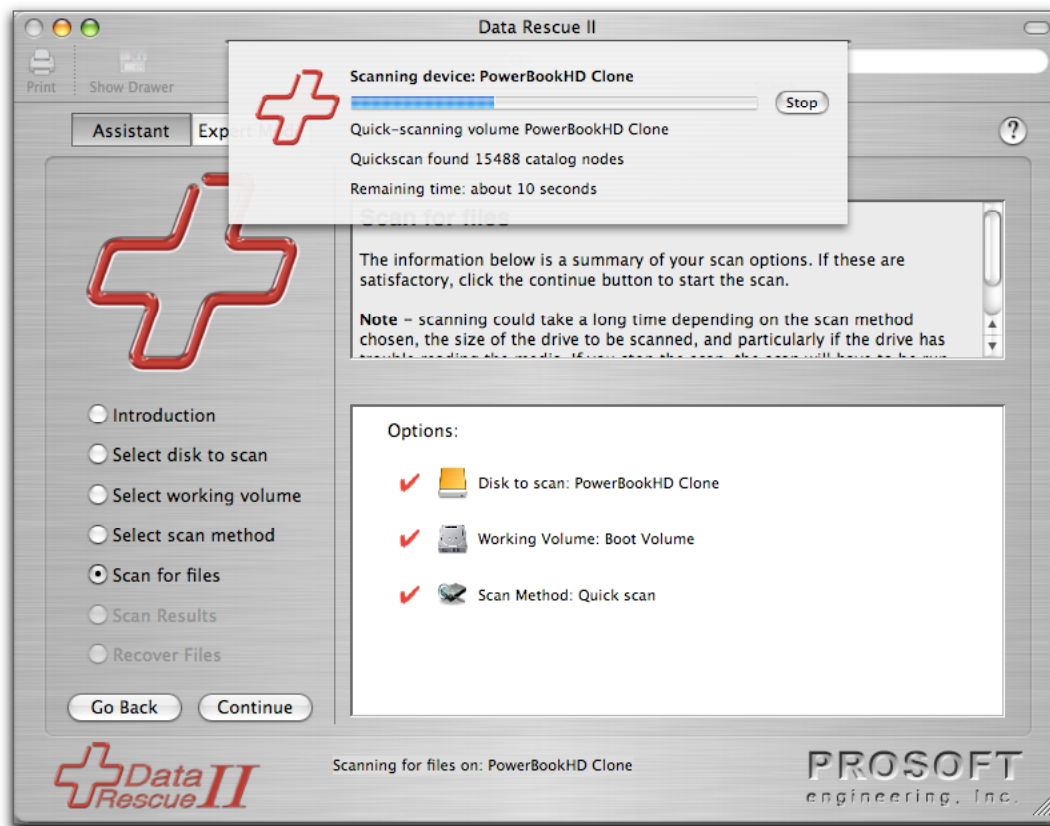


Figure 8. Scan Progress

When the scanning and processing is done, the progress sheet will disappear, and a Scan Results Summary screen will be shown (Figure 9).

4.5 Scan results

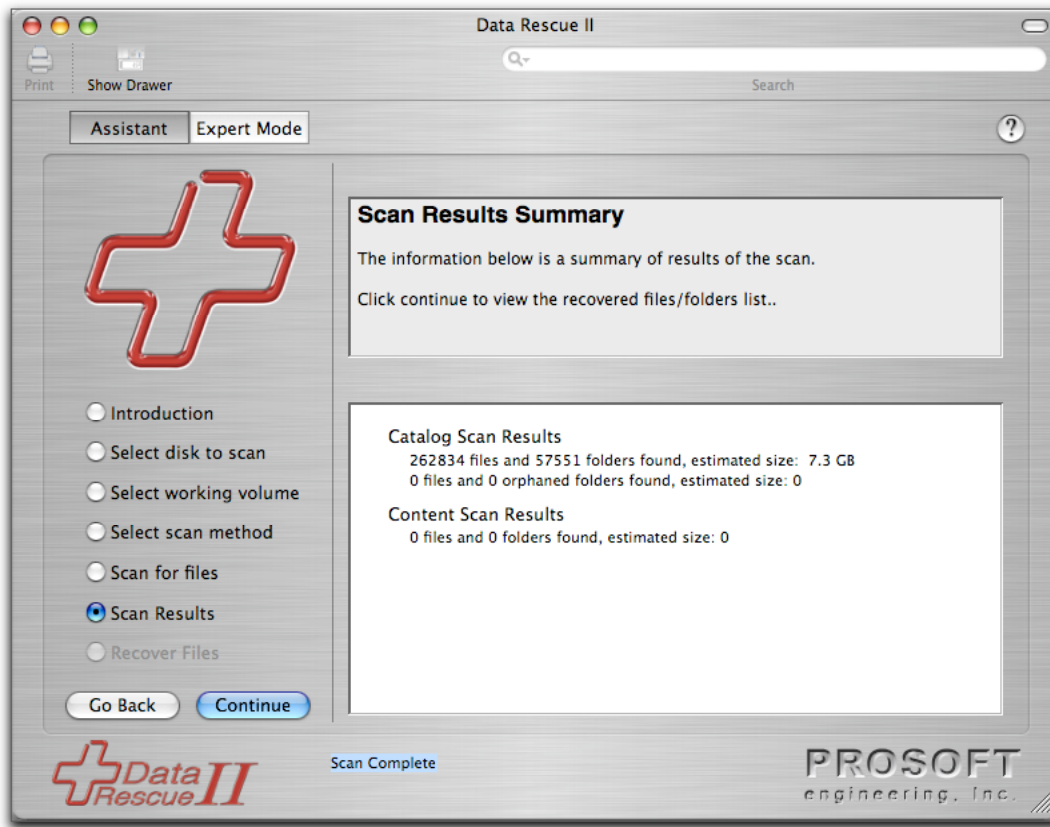


Figure 9. Scan Results Summary

This summarizes how many files and folders were found during the scan. Click Continue to see the found files (Figure 10).

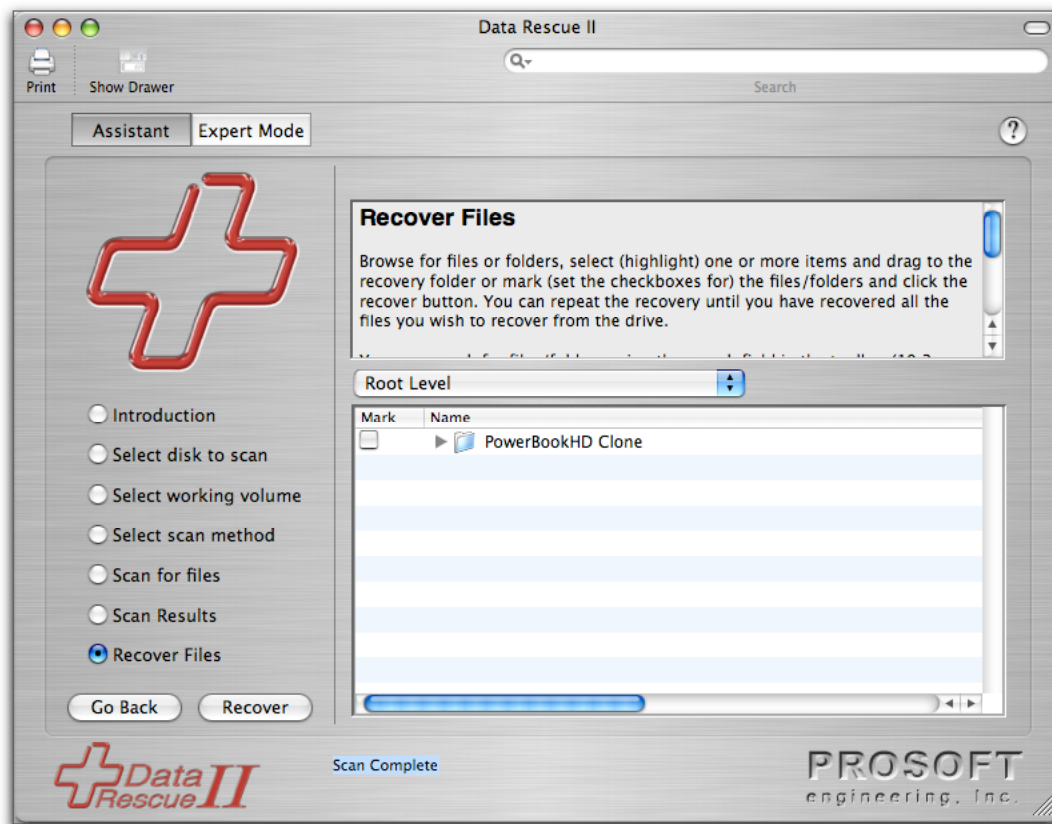


Figure 10. File Browser screen

You can use the disclosure triangle buttons in the window to open folders, and navigate to find things you want to recover, and click the checkboxes for those things in the Recover column (Figure 11).

4.6 Recovering files and folders

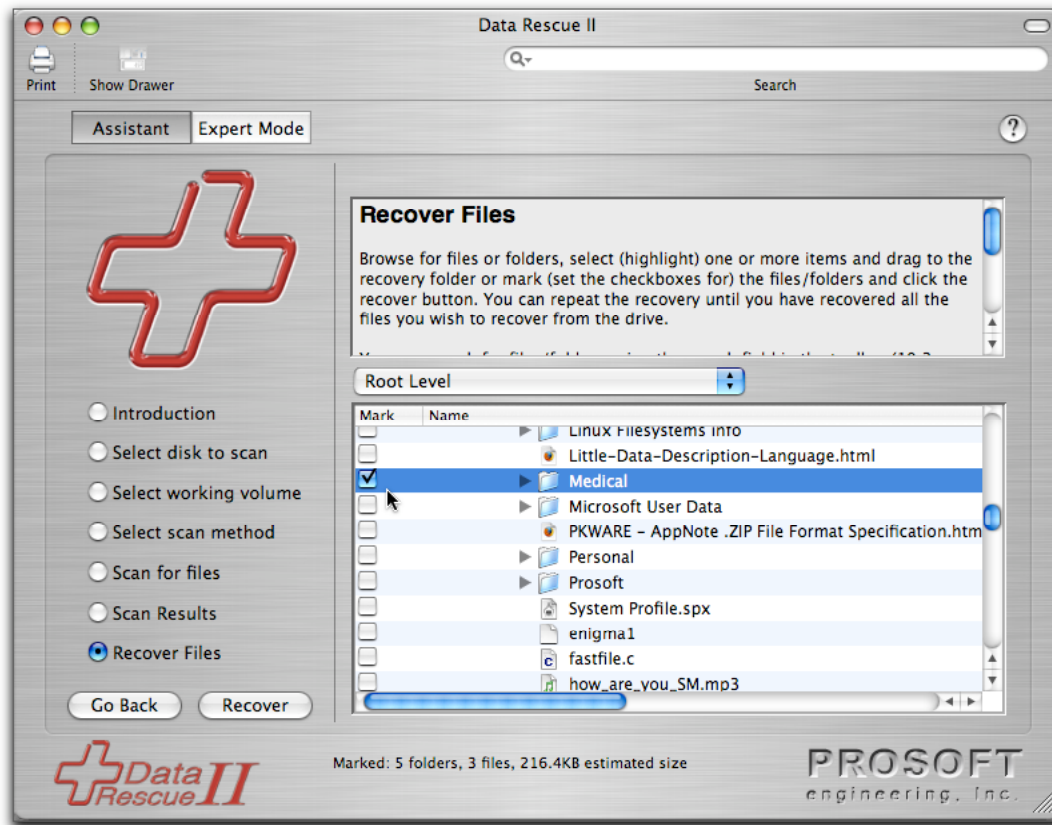


Figure 11. Selecting things for recovery

Once you have everything selected that you want to recover, click the Recover button. You will be prompted for a location (Figure 12).




Figure 12. Choosing a folder to recover files in

Files recovered this way will have the folder structure, starting with the root node (in this case, a folder named “PowerBookHD Clone”) and going down to all the checkmarked items created in the directory you chose. So in this example, a folder named “PowerBookHD Clone” will be created in the existing folder “My Recovered Files”, and all the folder hierarchy between there and the Medical folder, and all the files and folders underneath Medical.

If you don’t want the whole folder path back to the root to be created, then instead of using the checkboxes, you can simply select (highlight) the file or folder you want to recover, and drag it with the mouse to some location in a Finder window. Only the item(s) highlighted and whatever is contained in them (in this example the Medical folder and everything it contains) will be recreated in the folder you dragged it to.

4.6.1 Before doing a large recovery

After doing a Thorough scan, before recovering a large number of files, you should first try recovering and opening a few of your important files to make sure they contain good data. The reason is that with a thorough scan, there can be some uncertainty about the proper ABL choice. (See the later section entitled “After the Through Scan”.) If this setting is incorrect, the files you recover will contain garbage, and you will have wasted your time. Note: These comments about the ABL setting do not apply to files listed under the CBR folder – the ABL setting does not apply to those.

4.7 Double-click opening of individual files

As a convenience to the user, double-clicking on a file in the Recover Files view will cause this file to be automatically recovered to a fixed location (it is put inside the Data Rescue temporary storage folder), and the system is asked to open the file with the appropriate application. So for example, double-clicking on a JPEG graphics file would recover it and open it with the default image viewer, e.g. Preview. Note: This feature is disabled in Demo mode. Also, it will not work when booted from the Data Rescue CD ROM because it lacks the normal system applications for opening files.

4.8 Searching the scan results

You can also search for files and folders by name. Just type a string which the name(s) would include (case-insensitive) into the search field (Figure 13) and press the Return key.

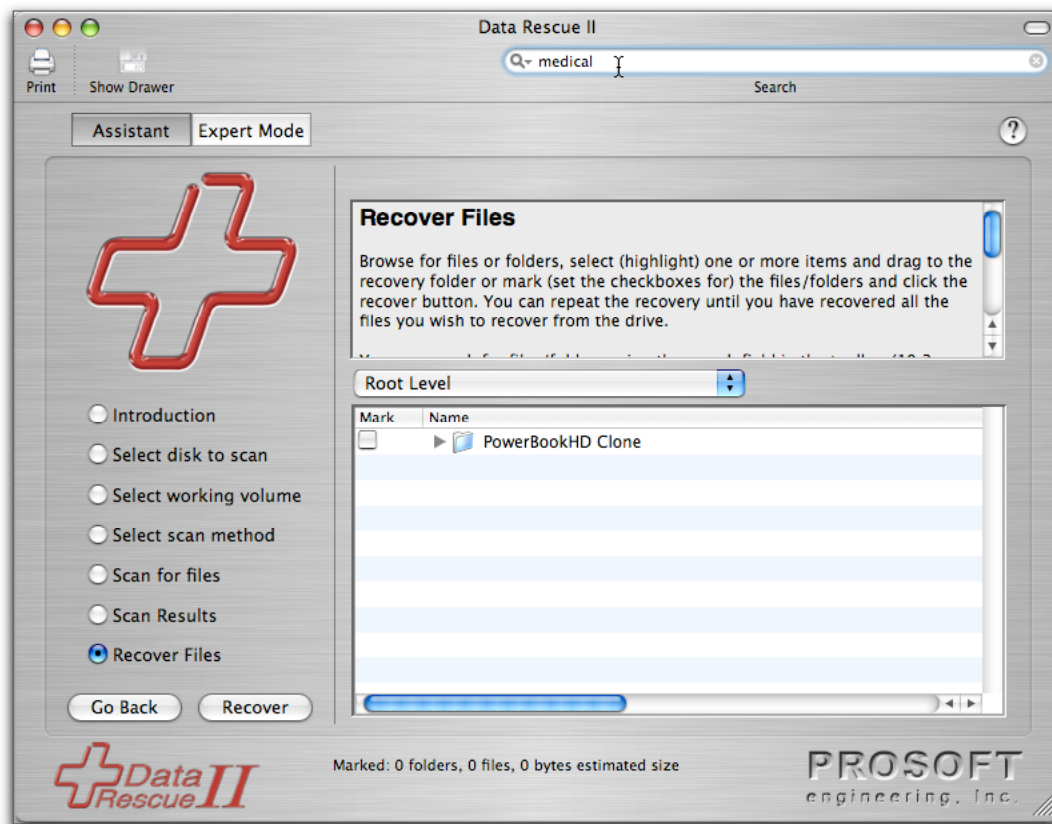


Figure 13. Searching for files or folders by name

Alternatively, you may go to the Edit → Find menu item to bring up a Find dialog (Figure 14). With this dialog, you can do more specific searches, including searches by size and date, as well as by name.




Figure 14. The Find dialog

A list of matches found will be displayed in the upper part of the search results area (Figure 15).

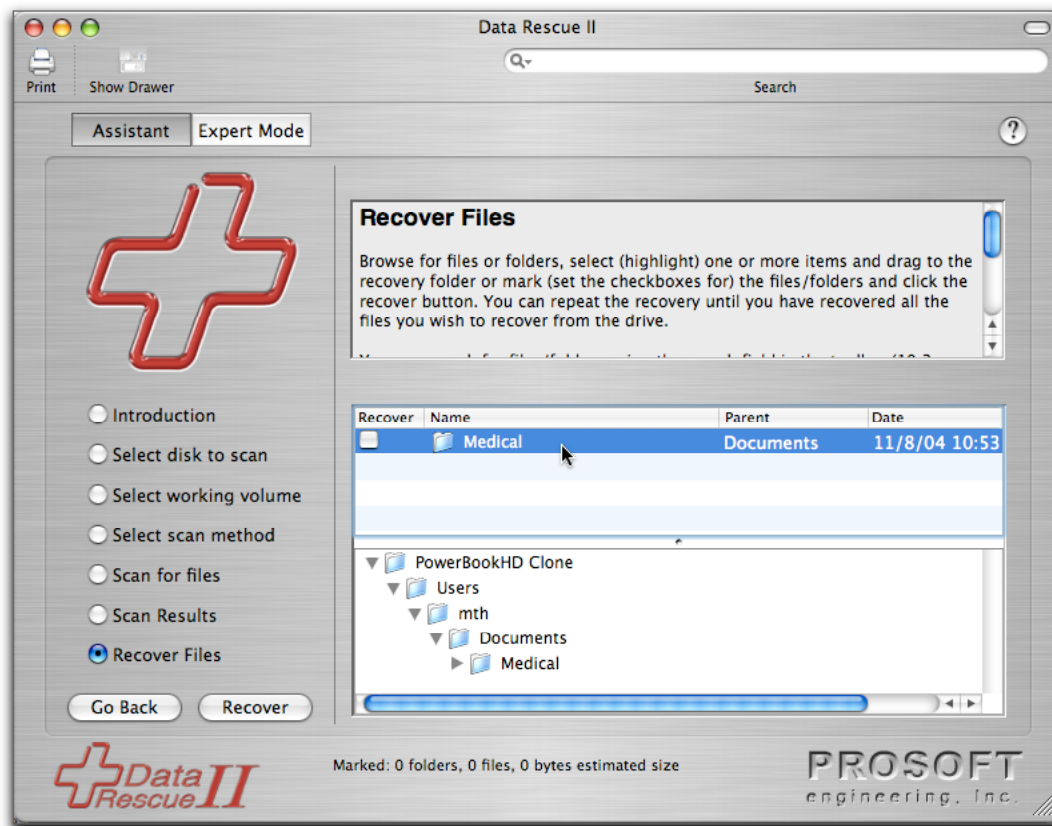


Figure 15. Search results

In this case, only one match was found. To recover, highlight or mark (checkmark) the item from the upper list and click Recover. The distinction between highlighting and marking is the same as before: using the marks will cause the folder hierarchy to be rebuilt all the way up to the root (“PowerBookHD Clone” in this example), whereas highlighting only will not rebuild the folders above the one selected. Note: if some items are marked and some are highlighted, the marked items take precedence. If any items are marked, only marked items will be recovered.

Whichever item you have highlighted in the upper list will also be shown in the lower list, the purpose of which is to show you the location of the upper item within the directory hierarchy.

To return to the previous recovery browser view, select the View → Show Scan Results menu item.

Note: Sometimes, and especially with a Thorough scan, an item that you are looking for will appear more than once and/or be found in the wrong folder or in one of the orphan folders. For this reason, if you do not find what you’re looking for, or if you do find it but it doesn’t recover properly, you should use the search by name and see if there are other copies elsewhere. Some of the copies that you find may be the remnants of older versions of your file. Data Rescue has no way to determine which version you’re looking for, so it must show all the ones it finds.

5 Using Expert Mode

Expert mode has some additional options “Analyze” and “Clone”, which are explained later. The other scan modes, Quick, Thorough, Deleted, and Use Scan File are identical in operation whether run from Assistant or Expert mode.

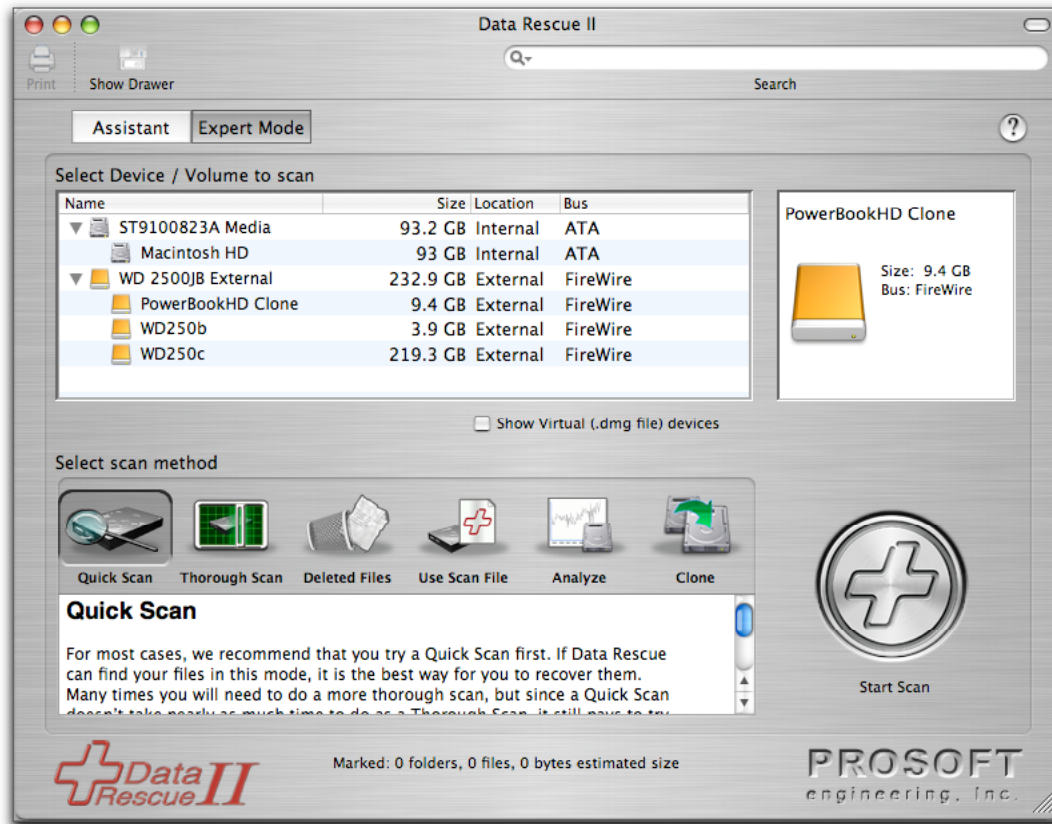


Figure 16. Expert Mode Screen

5.1 Selecting the source media and scan mode

To start, select a source media (device or volume) and a scan mode, in either order, and click the Start button. This will bring up the dialog to select the temporary (working) storage. Choose a volume with adequate free space, typically at least 1% of the size of the media to be scanned, and preferably not on the same device as the scanned media.

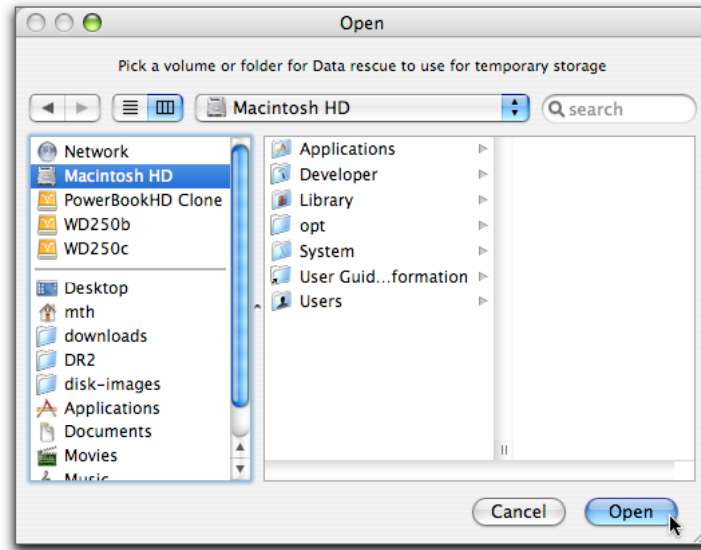


Figure 17. Selecting a location for working storage

You may, if you wish, choose some place other than the top of the volume. Wherever you choose, a folder named DataRescueTemp will be created to hold Data Rescue's temporary storage. (Note: choosing the boot volume will create the folder in /tmp.) Since the scan engine runs as root, it can create the folder anywhere you wish. It automatically changes the owner to the user running the program, to make it easier to remove manually later, if that should be necessary. Note: Data Rescue will normally automatically remove this folder and its contents when it quits.

5.2 Scan results

In Expert mode, the scan results are shown in a separate window (Figure 18).

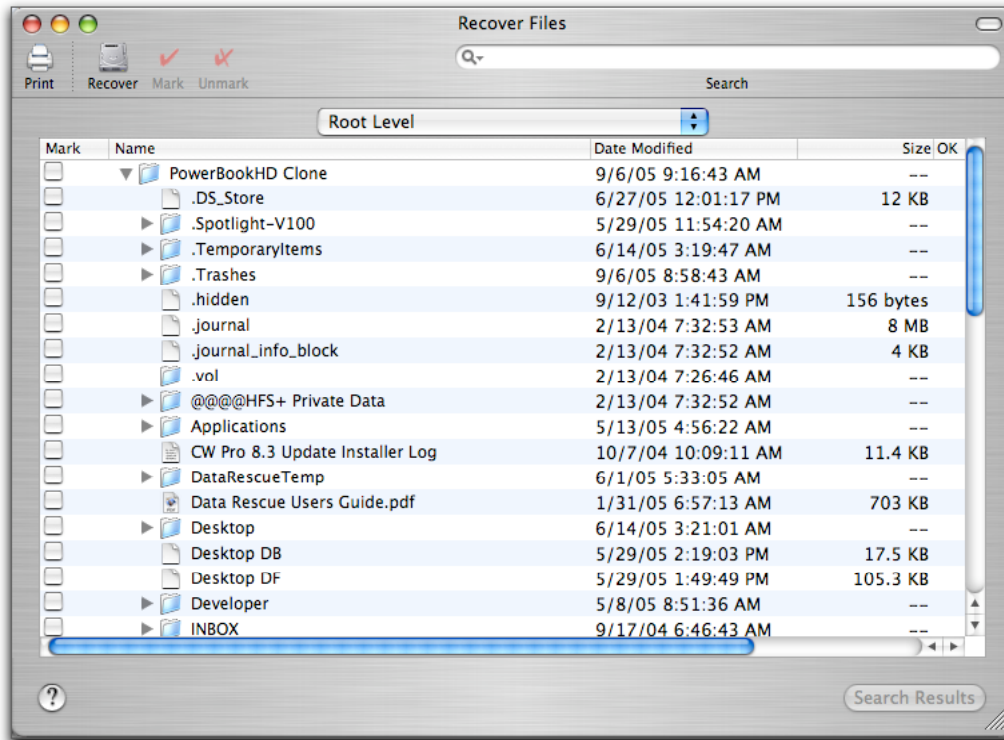


Figure 18. Scan Results

Navigation and selecting files and folders is the same as for Assistant mode. You can drag individual files and folders from this window to a Finder window to recover, or you can highlight or mark files and folders and click the Recover button. As with Assistant mode, if the marks (checkboxes) are used, the parent folders will be rebuilt back to the root, whereas with highlighting, only the files and folders highlighted and those included in the highlighted folders will be rebuilt. Double-click opening also works.

Searching works the same way as for Assistant mode. After doing a Thorough scan, before recovering a large number of files, you should first try recovering and opening a few of your important files to make sure they contain good data. The reason is that with a thorough scan, there can be some uncertainty about the proper ABL choice. (See the later section entitled “After the Thorough Scan”.) If the setting is incorrect, the files you recover will contain garbage, and you will have wasted your time.

5.3 Analyzing Media

Expert mode supports an Analyze function, which can provide a quick check of a drive or volume’s ability to be read by sampling its read times throughout the media. To do an analysis, select the media to be analyzed and the Analyze function button. Select the granularity with the popup menu on the panel (Quick, Medium, or Thorough Analysis) then click the Start Analyze to start it (Figure 19).

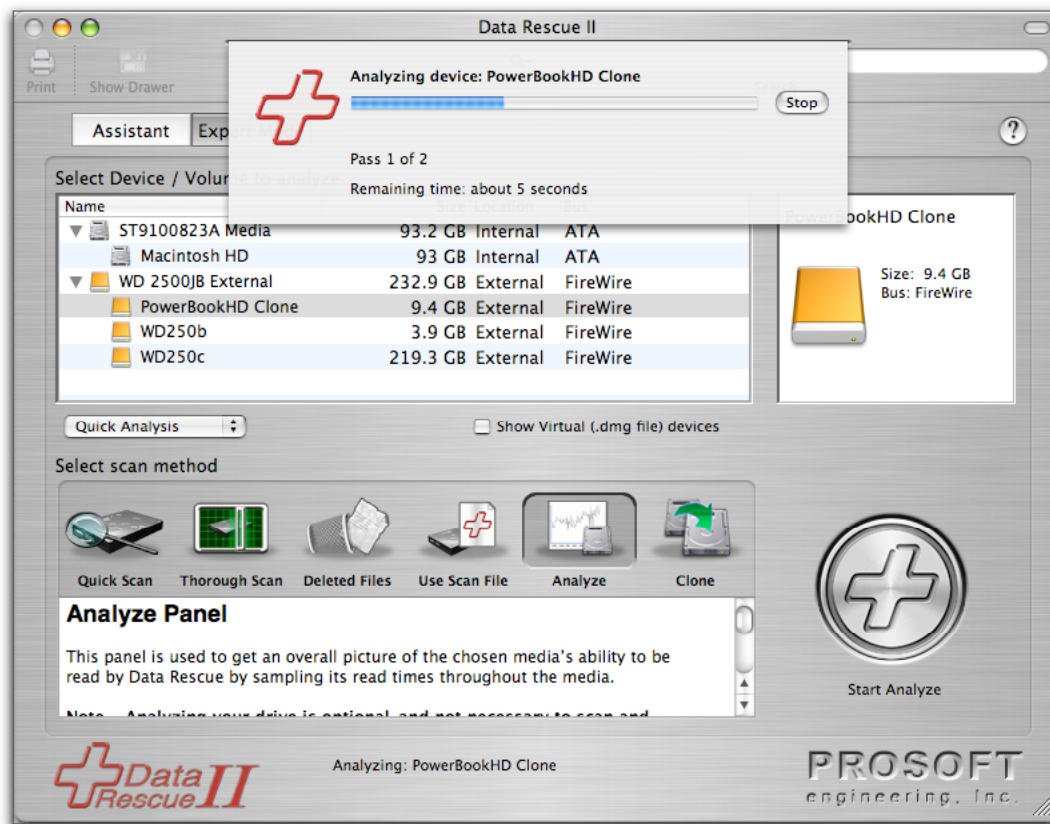


Figure 19. Analyzing a Volume or Drive

During the analysis, the 2MB blocks are read from the disk, starting at the beginning, and skipping to the next read location, etc. until the end of media. The Quick, Medium, and Thorough analyses differ only in their stride (media spacing between one read and the next). The Quick Analysis has the largest stride, and the Thorough the smallest. The default stride sizes are 128MB (Quick), 64MB (Medium), 8MB (Thorough). The stride sizes and read buffer size are adjustable in the Scan Engine preferences (Figure 20). See the later section on preferences.

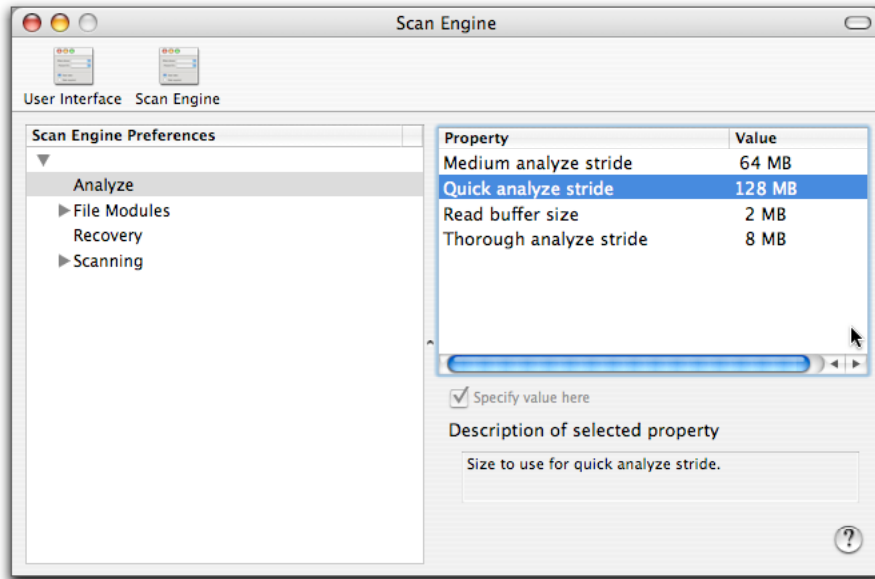


Figure 20. Analyze stride and read buffer sizes

The time taken to do each read during the scan is measured and kept. Two passes are made, and the lowest read time of the two for each read position is plotted (Figure 21).




Figure 21. Analyze Results

Note: This function is intended to give a reasonably quick indication of drive functioning, and does not necessarily indicate a very accurate read time for the drive, because other

Mac OS system interrupts and events can cause some inaccuracy. For the most consistent results, you should not be using the system for other things while analyzing.

If a disk is having trouble reading in a particular area, that will be reflected by read times in those areas which are significantly above average. A red line is plotted on the graph as a visual aid to seeing where read times are significantly above average. If your drive shows some read times above the red line, it doesn't necessarily mean that anything is wrong, because the system might have become busy just about then, artificially increasing the measured read time. However, if you repeat the analysis additional times and the long read times persist in the same part of the media each time, then your disk is having trouble reading those areas. Figure 22 shows such a case.




Figure 22. A drive with read problems

The drive of Figure 22 has persistent read problems in two areas near the end of the disk. Notice that the read times in these areas exceed one second. If Data Rescue were doing a thorough scan of this disk, when it got to these areas, progress would appear to stop or freeze, because of the greatly increased read time for blocks in these areas, although if allowed to proceed, it would eventually read through these areas successfully. This particular disk does not have any read errors. Read errors are reported as a fixed value of 2.000 seconds.

Another thing to notice is the gradual increase in read times toward the end of the disk. This is normal for many drives, and does not indicate a malfunction.

5.4 Cloning Media

The intended use of the Clone function is to make an exact copy of the drive you are having problems with, so that you can do the remaining scans and recovery operations on

the clone. Obviously, doing this would require another spare drive with the same or larger capacity as the original. Cloning the original source is not required to use Data Rescue, so why would you want to do it? There are a couple of reasons that you might want to.

If your original drive appears to be having some hardware problems, then you want to minimize the use of it, because certain types of drive hardware problems can become worse with use, and further damage data. At this point, it is prudent to reiterate the warning that was given at the beginning of this guide: *If you suspect hardware damage to your drive, and your data is extremely critical, you may wish to consider sending the drive to a professional drive recovery service.* However, if you cannot afford the considerable expense of these services, and you suspect hardware damage, cloning provides a way to reduce the use of your source drive to an absolute minimum. The clone function makes a single pass straight through the original drive and copies as much of the raw data as can be read to the clone drive. Once you have done that, you may put the original away for safekeeping, and do all the recovery operations on the clone.

If the original drive is intermittent, cloning may provide the best chance of getting a good snapshot of the data on it. We have had at least a couple of instances in which a drive was good when cold, but after running for an hour or two, started getting read errors. Since cloning takes considerably less time to complete than scanning and recovering, we were able to get the raw data off the drive while it was still working properly.

In another case, a malfunctioning drive seemed to be able to handle reading data when accessing it sequentially with a straight-through read, but couldn't handle reading when lots of seeking was required. In this case, the cloning operation itself was all that was required, and the clone disk functioned normally. (This was a rare case.)

Finally, although this is not its intended purpose, you can use the clone function to just duplicate a drive, if you need another identical copy for some reason.

You can select and clone volumes rather than whole drives if you want, although we recommend not doing that unless Prosoft technical support advises you to, or you are knowledgeable about disk partitions and volumes, and know what you are doing.

6 Preferences

You can bring up the preferences window by choosing the Preferences item of the Data Rescue menu. The preferences are divided into two groups – the ones which affect the user interface process, and the ones which affect the scan engine process.

The preference settings are persistent – that is, they are recorded in the user's Data Rescue preference files, and will be remembered for the next time Data Rescue is launched. Note: The preferences are not persistent when booted from the Data Rescue CD ROM.

The user interface preferences are recorded in the user's Library/Preferences/DataRescueII.plist file, and the scan engine preferences are recorded in the user's Library/Preferences/DataRescueIIBE.plist file. You can restore preferences

to the default values by deleting these files. Note: If you do this, you will need to reenter the activation serial number the next time you launch Data Rescue.

6.1 User Interface Preferences

The preferences open initially to the user interface group (Figure 23).



Figure 23. User Interface Preferences

6.1.1 Report level

The report level sets the level of information included in the Data Rescue Report. This report is accessible through the View → Report menu item.

6.1.2 Open report on any error

If an error occurs, it is recorded in the report. If this option is checked, the report window is automatically opened when any error occurs.

6.2 Scan Engine Preferences

Clicking on the Scan Engine icon in the toolbar area will bring up the scan engine preferences (Figure 24).

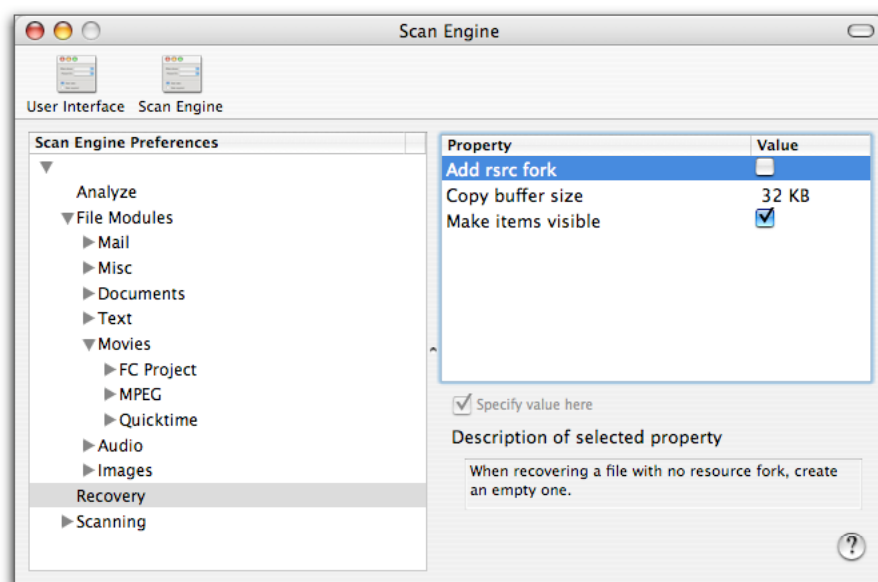


Figure 24. Scan Engine Preferences

There are quite a few scan engine preferences, and these are organized hierarchically and displayed as a tree. Note: Although there are a large number of preferences, the default values should be fine for most situations and users, so normally you don't have to worry about any of them. They are mostly there to provide extra flexibility for expert users and in unusual circumstances, where the default settings don't do the job for you, or you have some special requirement.

Each part of the tree may have properties and values. Highlighting a portion of the tree in the left-hand window will cause any associated properties and values to appear in the right-hand window. If you highlight a property, a description is displayed in the small text window in the lower-right. Properties with Boolean (on/off) values, are represented with a checkbox. Properties with numerical or text values are shown with a text field. If a property is not editable, it will be shown as gray (dimmed). Otherwise, you can change the value. To edit a numerical or text value, double-click the current value, then type in the new value.

For properties which represent sizes of things, you can usually input a new value using KB, MB, or GB size shortcuts. Thus, to specify a size of 2 kilobytes, you can input either 2048 or 2KB, etc.

The various scan engine preferences are discussed in the sections following.

6.2.1 Analyze preferences

In these preferences, you can customize the size of the strides and the amount of data read on each stride used for the Expert mode Analyze function. So for example, if the read buffer size is 2MB and the Quick analyze stride size is 128MB, then during a quick analyze, Data Rescue will read 2MB at media location 0, then another 2MB at media location 128MB, etc.

6.2.2 Scanning preferences

The Maximum files/folder preference is used during scanning when Data Rescue has to create its own folders for catalog files, such as in the Orphan's folder. When such a folder would exceed this number of files, Data Rescue will create a new folder. The default is 1000 files per folder.

The Scanning/HFS Enabled preference can be used to disable scanning for HFS catalog structures during a thorough scan. You might want to do this if you are using Data Rescue to scan a drive which is known to have some file system other than HFS, in order to speed things up a little. Normally, leave this enabled.

6.2.3 Recovery preferences

The Add rsrc fork preference controls whether files that are found without a resource fork should have an empty resource fork created for them. A few Mac applications may complain or refuse to open a file without a resource fork. If you experience this, you may turn this flag on. The default is off.

The Copy buffer size preference specifies the size of the buffer to use for copy operations, such as when recovering files. A larger buffer size will speed the recovery of files, up to a point, but will of course use more memory.

The Make items visible preference controls whether files found which are invisible are set to visible when recovered.

6.2.4 File Module preferences

Data Rescue uses two strategies to find files, catalog and content. The content scan strategy is divided into file modules. Each file module is responsible for detecting one kind of file. Most file modules have a number of preference properties in common, such as enable, maximum file size, etc. Since there are over 100 file modules, and many of the properties may have the same value for groups of file modules, for convenience, some of these properties may be specified at a higher level of the tree, allowing the lower level parts of the tree to use that value, without needing to specify the same value for every individual module. Individual file modules may still override these higher level values on an individual basis. Figure 25 shows the properties specified at the File Modules node of the tree.

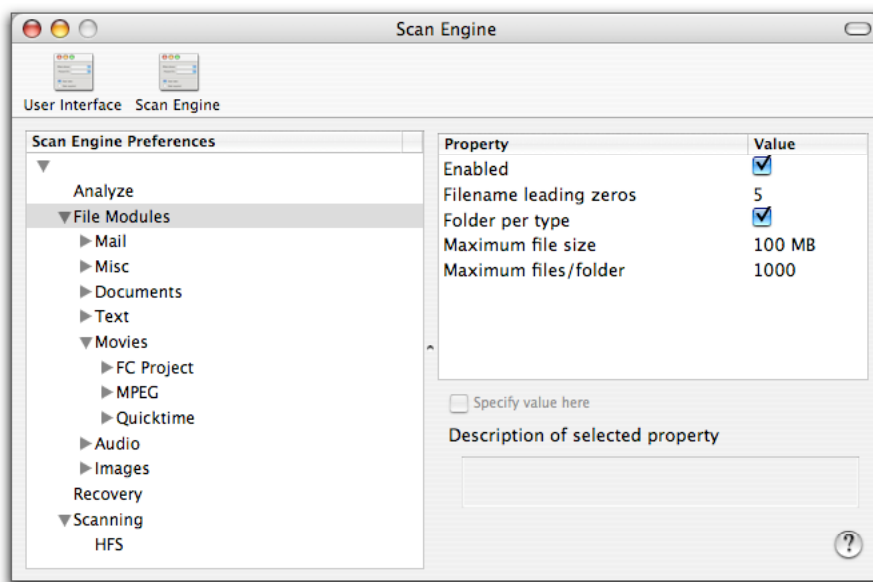


Figure 25. Top level file module preferences

6.2.4.1 Top level file modules preferences

“Enabled” specifies whether a module is enabled to look for files during a scan. The default enabled value for this node is On, so all the file modules underneath this point will be enabled, unless overridden at a lower level.

“Filename leading zeros” controls how many numerical digits are used when synthesizing names as the file modules find new files. So for example, with the default setting of 5, filenames like abc-00001, abc-00002, etc. will be generated. With a setting of 3, names like abc-001, abc-002, etc. will be generated. The purpose of the leading

zeros is so that Finder will list these in numerical order. Setting this to 0 means that no leading zeros will be added.

“Folder per type” controls whether the files found by the file modules will be listed in their own per-type directory or not. So for example, with this option On (the default), a raw CRW image type would appear under CBR/Images/CRW/, a JPEG image under CBR/Images/JPEG/, etc. With the option off, these would both appear under CBR/Images/.

“Maximum file size” controls the maximum size that certain files can be. Not all file modules implement this, but for the ones which do, if a file would exceed this number, it is ignored.

“Maximum files/folder” sets the maximum number of files per folder for file modules (i.e. the ones under the content (CBR) folder).

6.2.4.2 General file type preferences.

These consist of the first level names under the content folder (CBR), such as Mail, Documents, etc. Most of these don’t specify their own preference properties. The preferences that are inherited from the top-level file modules preferences are still shown so you can see their values, but are not editable (grayed out), as the one for Mail (Figure 26).

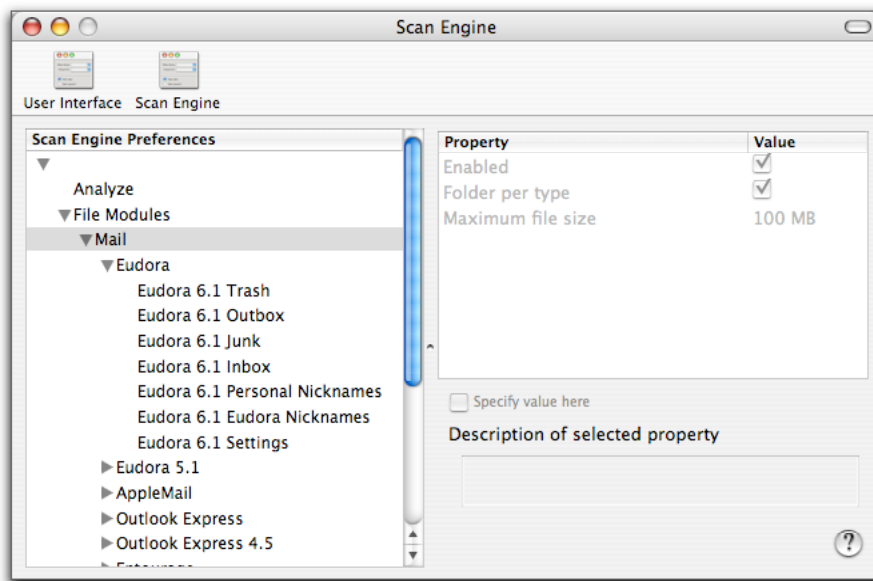


Figure 26. Non-editable preferences

A few types have their own properties at this level – the Text file modules for example, which by default are disabled because of their rather severe impact on performance (Figure 27). You may want to turn this on if you want to scan for a particularly important text document that is not recovered by any of the other file modules.

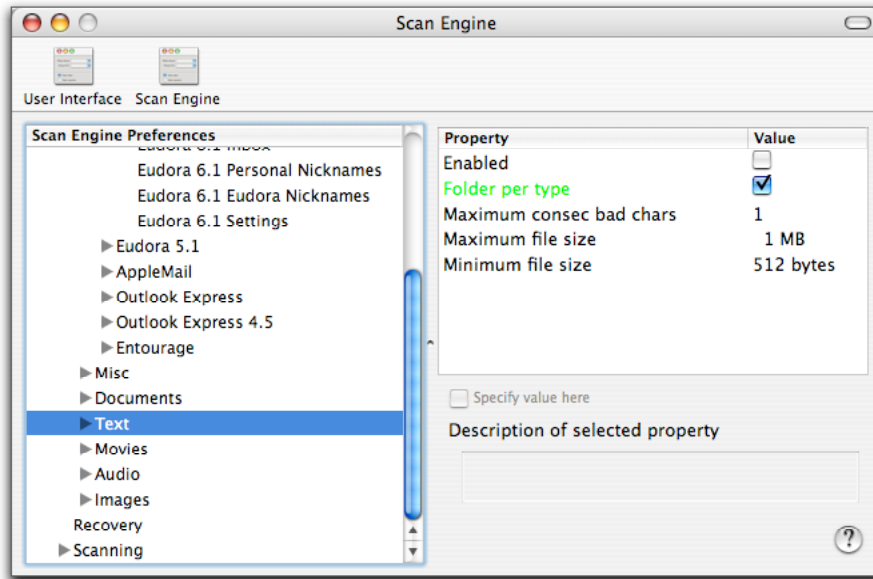


Figure 27. The Text file modules are disabled by default

One of the properties, “Maximum consec bad chars” is unique to the Text file modules, and so that is specified here, as opposed to the top level.

You may also notice that the “Maximum file size” preference has been specified here, rather than using the larger one specified at the top level.

6.2.4.3 Particular file type preferences

In Figure 28, we have navigated all the way down to a particular file type preference screen.

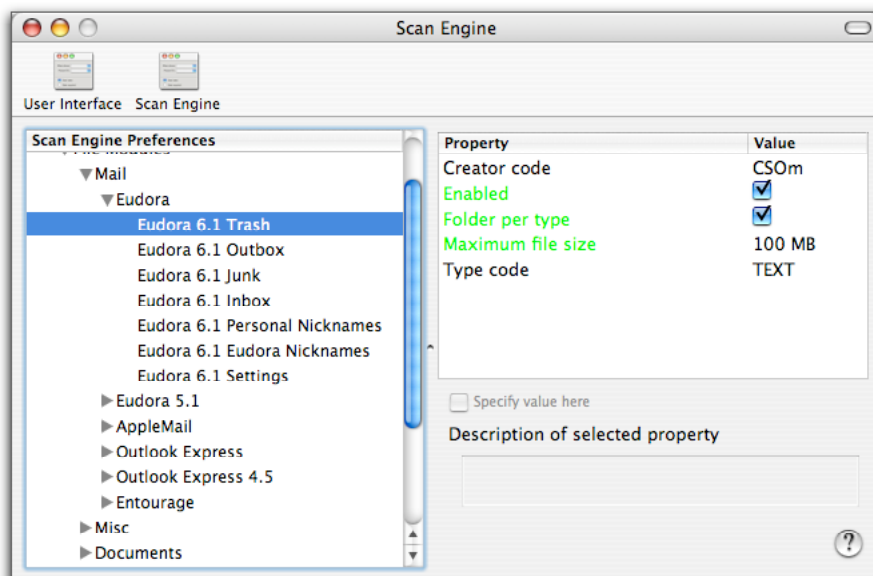


Figure 28. Particular file module preferences

The properties which are listed in green are ones which are specifiable for this file type, but currently are using a value from a higher place in the tree. In this example, the Maximum file size property is not specified here, but inherits the value of 100MB from higher in the tree. If you wanted to specify a different maximum file size for Eudora 6.1 Trash files, you could do so here. For example, we could set it to 1GB by double-clicking the value field, and entering 1GB, as we did here:

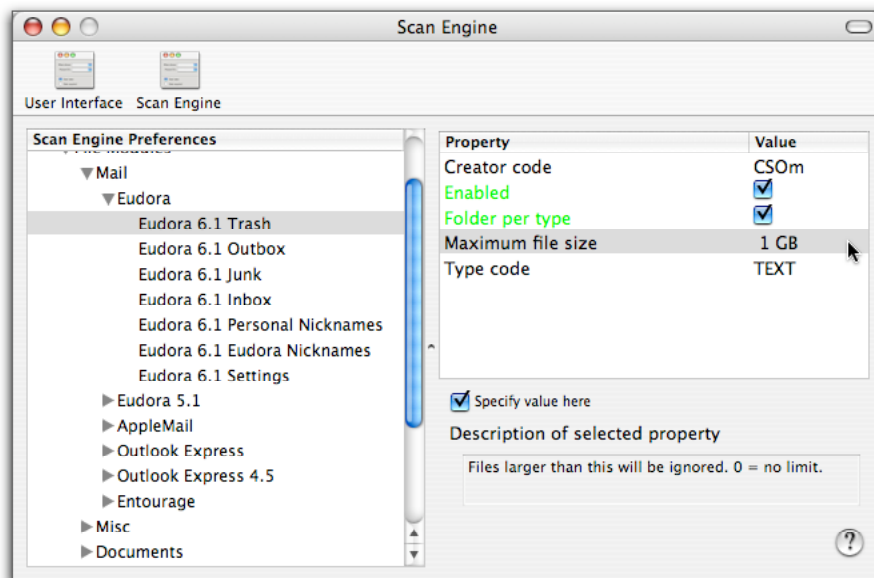


Figure 29. Changing an editable inherited value.

Notice that the property is no longer shown in green, and the “Specify value here” checkbox is set. To unspecify the value and return to the value specified by the higher part of the tree, uncheck that checkbox. You should see the property change to green (inherited) once again, and take on its former value.

7 After the Thorough Scan

7.1 Saving a Scan File

Following a Thorough scan, there are a couple of things that you should do before proceeding to recover a large number of files. The first thing is that you should probably do after a thorough scan is to save a scan file, which you can do with the File → Save Scan File menu item. This will bring up a dialog and allow you to save the information collected during the scan to a file of your choice. Later, if you wish to recover other files from this same volume, you can use the scan file rather than repeating all of the Thorough scan, thus saving you time. To use a previously saved scan file, choose the

Use Scan File as your scan method. You will be prompted for the name and location of your scan file.

Note: The scan file only saves indexes to the files found on the scanned drive, not the file data itself. To recover files later using the scan file, you will need to retain the scanned drive in its unmodified state.

7.2 Allocation Blocks Layout Dialog

The second thing you should do after your Thorough scan is to check several of your important files in the catalog scan folders (the ones whose top-level folder name is the name of the drive or volume you scanned). Either recover a few of these, then open them, or you can double-click them, which will automatically recover and open them. If none of them open, or they all contain garbage, then chances are that the allocation block layout (ABL) is wrong. In this case, go to the Expert → Allocation Blocks Layout menu item. A dialog will open:

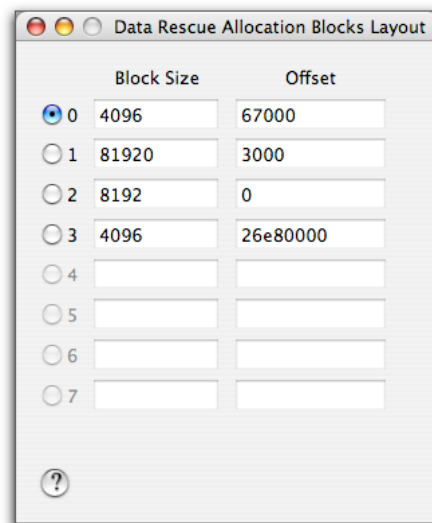


Figure 30. Allocation Block Layout Dialog

During its scan, Data Rescue determines some possibilities for the proper ABL setup, and adds them to this window. The Block Size is the number of bytes, in decimal, of the allocation block. The Offset is the number of bytes, in hexadecimal, from the start of the media to the start of the file system.

The choice Data Rescue considers most likely is put first, but if that doesn't result in good files, try each of the others. Hopefully one will work. In some cases, in particular if the drive you scanned had more than one partition, or if there are remnants of older file systems still present on the disk, you will find that some files require one of the ABLs and other files will require another. The only way to tell at the current time is to try them. Note: The ABL setting has no effect on the files under the CBR folder.

Allocation Blocks are explained in more detail in a later section entitled "Allocation Blocks and Media Blocks".

7.3 Recovering Files from Digital Cameras and Music Players

If your digital camera or music player uses removable flash media (e.g. CompactFlash cards, etc.) and you have a card reader device attached to your Mac, or if the camera or player can attach directly to your Mac and emulate a disk device, then you can use Data Rescue's Thorough scan to recover pictures or songs from it. If your digital camera uses JPEG format for still images, and Quicktime or MPEG-1 or AVI for its movie clips, then you should have success recovering these. Some popular RAW camera image formats are also supported. Many digital music players use MP3 or AAC/M4A formats for the songs, and you should be able to recover these also.

If you discover that your digital device uses a format not currently supported by Data Rescue, let us know about it. We will use this type of input to guide our decisions about what type of Content formats to add support for in the future. You can mail us at support@prosofteng.com.

8 Concepts

The next chapters are for users wanting a more detailed understanding of Data Rescue options, or who were unable to recover the files they wanted using Chapter 4's simplified instructions. This chapter explains some concepts needed to understand how to set the Data Rescue options not covered earlier and how to use Data Rescue to its fullest. The next chapter goes through the specific options, and what they do and when and why you might use them.

The concept of volumes versus drives was discussed earlier in Chapter 4. If you don't understand the difference, please read that section first.

8.1 Selecting a Volume or Drive to Scan.

8.1.1 Drives versus Volumes

As used in this document, the term “drive” refers to a whole device, for example an IDE or Firewire disk drive. A “volume” is a logical part of a drive. A drive may be partitioned into multiple volumes, or it may contain only one volume – this is determined when the drive was originally configured. The next figure depicts the storage space available on two disk drives. The first drive has most of its space allocated to a single volume. The second has most of its space split between two volumes.




Figure 31. Drives vs Volumes

Each volume resides in a partition of the drive. Stored within a volume are that volume’s name and all its files, folders, and catalog structures. The files and folders on a volume are all defined relative to the beginning of that volume on the drive.

In normal Macintosh operation, the user will see on his desktop an icon for each volume, labeled with the volume’s name. These are names that may be specified or changed by the user; names like “Macintosh HD”, or “My backup disk”, etc. The drive itself also has a name, defined by the drive manufacturer; names like “IBM-DTLA-307045 Media”, “DMI MAXTOR 6Y060L0 Media”, etc. In normal Macintosh use the user doesn’t see drive names – only the volume names. However you may see both when you use Data Rescue, so you need to be aware of them.

When Data Rescue starts, it makes a list of drive and/or volume names and puts them in the media selection part of the main window.

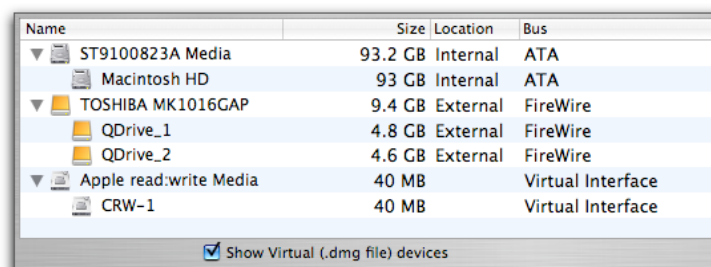


Figure 32. Media selection

The 1st drive has a hardware drive name of “ST9100823A Media” which has a single volume on it named “MacintoshHD”. The 2nd drive (“TOSHIBA MK1016GAP”) is divided approximately in half, and contains two volumes, with names QDrive_1 and QDrive_2.

For drives whose space is mostly taken up with a single volume, it usually won’t matter whether the drive or the volume is chosen. Unlike Data Rescue X, Data Rescue II can run either Quick or Thorough scan on either a drive or a volume. One difference is that if a volume is chosen, the limits of the scan are determined by the recorded size and starting location of the volume. These parameters come from structures on the drive that are subject to being overwritten if the drive gets corrupted. If you suspect that the volume is showing an incorrect size, then you should choose the drive to scan instead. The limits of the scan for a drive are gotten from the value reported by the drive interface, and these are not subject to being overwritten in the same way that volume limits are.

If your drive is known to be partitioned into more than one volume, then you should choose the volume you wish to scan, rather than the drive, as long as the volume appears to be the correct size. If the whole drive is scanned, not only will it take longer (because of scanning the whole drive rather than just the volume portion), but also the results from more than one volume will be intermixed in the output, and it will complicate the recovery, because the files belonging to one volume must be recovered using a different ABL setting than the files belonging to another volume.

A future enhancement to Data Rescue II may allow automatic separation of such results, making recovery in this situation easier, but that is not in place currently.

8.2 File Systems and Catalogs

Each disk volume is formatted as some sort of file system. Examples of file system formats are HFS, HFS+, FAT16, FAT32, UFS, NTFS, etc. Macintosh computers typically use HFS+. Windows PCs typically use FAT32 or NTFS. The details of these formats differ considerably, but they have some things in common: they all support file and folder (directory) objects, and they all have some form of bookkeeping information, which we will call “catalog” information. Figure 33 shows part of a simplified generic file system to help illustrate the concepts.

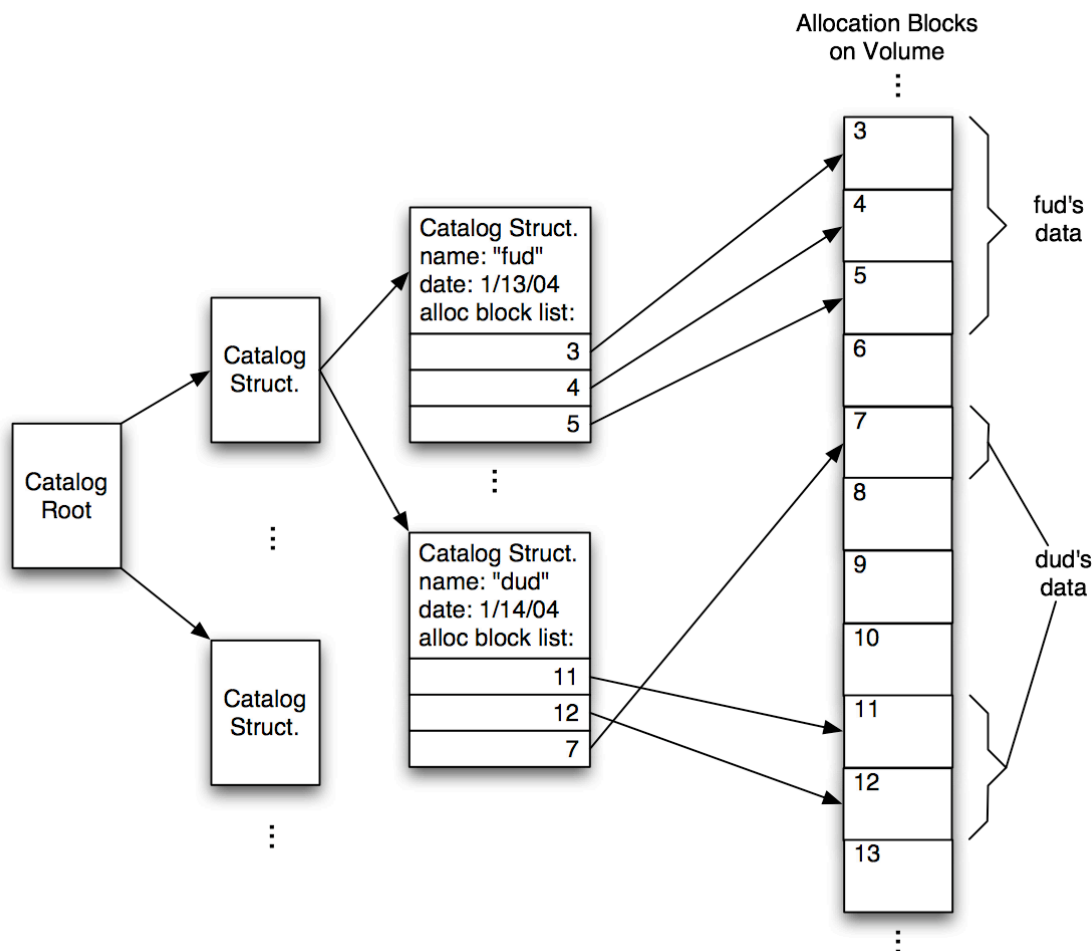


Figure 33: File System Data Storage and Catalog Structures

The files are used to store the content, such as the bytes making up an image, song, etc. Each file has its own format (JPEG, MP3, etc.) which is independent of the filesystem format. Folders (directories) are special files which are used to organize files and other folders into a logical hierarchy called a “directory structure”.

The catalog information keeps track of files and folders, what their names are, and where they are stored within the volume.

Figure 33 illustrates a file system with two files named “fud” and “dud”. Each file has a catalog structure. A catalog structure stores all the information about the file except the data itself, which for file fud is stored in 3 consecutive allocation blocks numbered 3, 4, and 5. The operating system uses the information in the catalog structure to get the location of the file’s allocation blocks on the media (the disk). Other catalog structures are in turn used to get the location of those catalog structures, etc. It all forms a big branching tree-like structure. At the base of the tree is a special catalog structure called the catalog root which the operating system normally knows how to find directly. From

the catalog root, all the other catalog structures, and file allocation blocks can be found by the OS.

Figure 21 also shows another file “dud” which is similar to file “fud”. “dud” also has its data stored in 3 allocation blocks, but the blocks are not consecutive. The first 2/3 of dud’s data are stored in blocks 11 and 12, but the last part is stored in allocation block 7. File dud is therefore not stored consecutively and in order, and is said to be “fragmented”. Since the location and ordering of the blocks is maintained in dud’s catalog structure, the operating system is able to use this to maintain the illusion of continuity, making it appear to applications as though dud’s data were all one continuous piece.

With the preceding discussion as background, we are now in a better position to explain the Data Rescue II scan types. Note: Data Rescue X supported separate Thorough Catalog and Content scans. In Data Rescue II, these two scan types have been integrated together in the Thorough scan. Since both types of scanning are still occurring in Data Rescue II (just concurrently) we will still discuss them as separate concepts below.

8.3 Catalog Scans

A catalog scan looks for and relies on catalog structures to find files. To the extent that the files and directory structure are still referenced by undamaged portions of the catalog information that is found, these files along with their original names and directories can be recovered. File recovery using a catalog scan is independent of the file type or format; in other words, files of all types are equally well recovered by catalog scans.

There are two types of catalog scans supported in Data Rescue: Quick and Thorough. The Quick Catalog scan traverses the catalog tree much the same as the operating system would, and will find files and folders that are part of the catalog tree. The Thorough Catalog scan doesn’t rely on the higher level catalog tree structures. Instead, it searches the media directly for catalog structures that point to files and folders.

The Thorough scan is generally more useful, because it can find catalog structures that are not referenced by a (possibly damaged) catalog tree. (In addition, the Thorough scan incorporates a Content scan, discussed below.) The Quick catalog scan assumes that the most important catalog structures are valid and intact, and only looks for catalog information in the places they are supposed to be. For this reason, the Quick catalog scan is often not very useful for recovering files on file systems that have sustained more than minor damage. Still, there are cases in which the operating system can’t see the files but the Quick catalog scan can find them. Since the Quick catalog scan is very fast compared to the Thorough catalog scan, it doesn’t hurt to give it a try first. If it doesn’t find the files you wanted, you can then go ahead with a Thorough scan.

The catalog scans will only be effective for file system formats which Data Rescue understands. Currently this includes HFS and HFS+ file systems, though other types may be added in the future.

8.4 Content Scan

The Content scan (included in the Data Rescue II Thorough scan) doesn't depend on any catalog structures, but instead searches the media directly for the file data. It does this by recognizing the format of the file's contents. As part of the Thorough scan, the Content scan checks the entire volume or drive. Because it searches for the file contents directly, rather than the catalog structures, it cannot recover a file's name or folder (directory). On the other hand, it is able to recover file data that is no longer referenced by any catalog structure, which catalog scans cannot do by themselves.

A Content scan can only recover files whose format is recognized and supported by Data Rescue. Current this list includes the following types, but more will be added in the future:

Archives:

Backup, DMG, GZ Files, SIT, ZIP

Audio:

AAC, AIFF, Loop, MP3, STMP, WAV, iTunes Library, M4a, M4b, m4p

Documents:

AppleWorks, ArchiCAD, Bookmarks, ClarisWorks, Excel, FileMaker Pro, FreeHand, HTML, InDesign, Office, OpenOffice, PDF, PageMaker, PowerPoint, Programming, Quark, Quark EPS, Quickbooks, QuickCeph, Quicken, RTF-Mac Files, Web Archive, Word Document, iChat, iTunes Library

Images:

CRW, CaptureOne, Comic Life, EPS, GIF, Illustrator, Illustrator EPS, JPEG, NIB, ORF, PNG, Photoshop, Photoshop BMP, Photoshop EPS, Photoshop JPEG, Photoshop PCX, Photoshop TIFF, TIFF

Mail:

AppleMail, Entourage, Entourage v11, Eudora, Eudora 5, Outlook Express, Outlook Express 4.5

Misc:

ETC, LPG, MISC, Norton, Property Lists, RIFF, RSRC

Movies:

AVI, After Effects, DV, FC Project, Flash, MPEG, QTZ, Quicktime, iMovie

Text:

HTML, PLIST, Postscript, Rich Text, TXT, XML

In the current version of Data Rescue, only non-fragmented files (i.e. those which are stored sequentially on consecutive disk blocks) can be successfully recovered with the Content scan. (Catalog scans have no such limitation.) This limitation may be addressed in a future enhancement to Data Rescue. Fortunately, many filesystems including HFS/HFS+, FAT, NTFS store most of their files in a non-fragmented manner.

8.5 Deleted Files Scan

A Deleted Files scan is similar to a Thorough scan for a volume, except that only the free space of the volume is searched. Since there are no catalog entries for deleted files, this search does not look for them. In other words, only a content scan is done for the free space. As a result, all files found with this method are found by the content scan method, and as explained above in the Content Scan section, only the file contents will be available. The original filename and folder information is not available, so the file contents are given a generated name and folder by file type. The files found by this type of scan will be a subset (often, a very much smaller subset) of files found by a Thorough scan, and therein lies one of its advantages: the user will have many fewer results to look through to find the files he is after. In addition, because the free space will often be only a fraction of the total volume space, the search time will be only a fraction of that of a Thorough scan.

Because this scan method depends on an accurate map of free space on the volume, it is only suitable for volumes which are in good shape (undamaged). If your volume is damaged, you will most likely have to use Thorough scan to find your files.

8.6 Scan File

Regardless of the method chosen, as Data Rescue scans your media, it builds an internal list of information about files and folders it has found. When a scan is complete, this information may be saved as a file called a “scan file”. In this way, the results of a lengthy scan can be saved, then loaded in later without having to repeat the scan. We recommend that you always save a scan file immediately following a lengthy scan.

Note: The scan file only saves indexes to the files found on the scanned drive, not the file data itself. To recover files later using the scan file, you will need to retain the scanned drive in its unmodified state.

To utilize a previously saved scan file, choose the “Use Scan File” option on the main window rather than a catalog or content scan. When you click the Start button, you will be presented with a dialog to choose the previously saved scan file.

Since a scan file represents a snapshot of information collected from a volume in a particular state, it should be clear that if the volume is allowed to change (for example if the volume is put or remains in service) following the saving of the scan file, the scan file will be stale with respect to the volume, and will no longer accurately represent its contents. The more stale the scan file is, the more inaccurate it will be.

8.7 Allocation Blocks and Media Blocks

This section will explain media blocks and allocation blocks, how they relate to each other and why their relationship is important. Normally you don't need to know about them, unless Data Rescue has trouble determining the proper relationship between the two.

We use “media” as a generic term to mean any type of device that holds volumes with files and folders. Digital media is usually divided into fixed sized chunks which we will call Media Blocks, or MB for short. (An MB is typically 512 bytes, but the size is not important for this discussion.) Figure 34 illustrates the media broken into a string of consecutively numbered MBs, starting at 0 at the beginning of the device, all the way to the end. Data for a file always starts at the beginning of an MB, so the start of any file may be accessed by its MB number (MB#). The software that talks directly to the storage device does so using the MB numbers.

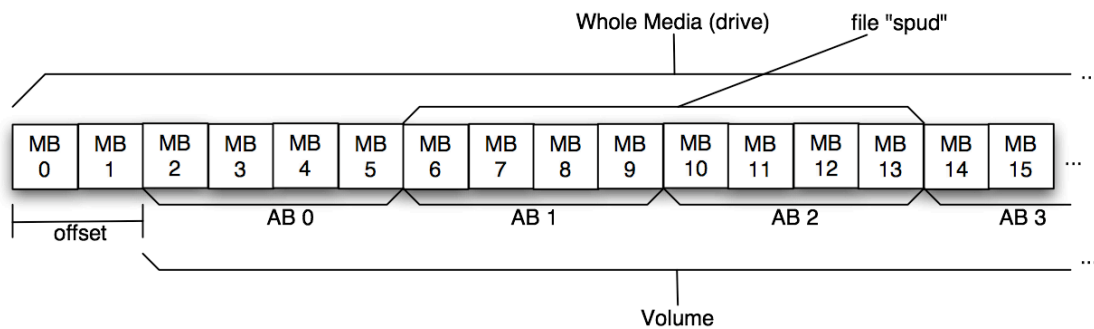


Figure 34: Media and Allocation Blocks

Earlier section 8.1, with the help of Figure 33, explained that a file's data is broken into allocation blocks which are stored in a volume. We call these “allocation blocks” (Abs for short) because whenever the operating system needs to allocate space for a file in the volume, it always allocates a whole number of these blocks. Figure 34 illustrates how an example volume might lay out on the media. Allocation blocks are always a fixed whole number of MBs. Volumes always consist of a whole number of ABs. In this example, there are 4 MBs to every AB, and the volume starts at MB#2. The distance from the start of the media and the start of the volume is called the “offset”, and it is not necessarily a multiple of the AB size. The example illustrates one file “spud”. Files always start at the beginning of an AB. In this case, spud starts at MB#6 and AB#1.

The catalog structures (discussed earlier) contain a list of AB numbers making up a file's data. Therefore, the catalog structure for spud would contain “1” as the starting location of spud, because it starts at AB#1 relative to the start of the volume.

In order for the operating system software to read the first part of the data from spud though, it must ask the storage device for MB#6, because that is where spud starts in terms of MB numbers. Therefore, the AB# in the catalog structure must be translated into an MB# in order to access the data. To do this, you need to know two numbers:

- What is the offset?

- What is the size of an AB in terms of MBs?

You can see that for our example, offset = 2 because the volume starts at MB#2, and size = 4, because each AB is comprised of 4 MBs. The formula to do the translation in general is quite simple:

$$\text{MB\#} = \text{AB\#} \times \text{size} + \text{offset}$$

These two important numbers, size and offset, are together referred to as the “allocation block layout”, or ABL for short. (Just a fancy name for a simple concept.)

Normally the ABL numbers for a volume are stored in a known location so the operating system can find them. For HFS and HFS+ file systems, they are stored in a volume header structure called the MDB/VH. For HFS+ systems, a copy of this structure is also stored, called the Alternate MDB/VH.

The significance of the ABL numbers should now be obvious – if a Data Rescue scan has found a catalog structure for a file, it will still need to know the proper ABL numbers in order to read the data from the media when recovering the file. For example, in the previous example, if the offset were accidentally changed to 1 rather than 2, when you tried to recover spud, the starting point would be calculated as MB#5 rather than MB#6, which might be part of some other file, or just complete gibberish. In fact, the wrong ABL would mean that *every* file on the volume would be read and recovered incorrectly.

8.8 Locating Your File After a Scan, Revisited

Files located by their content (i.e. those listed under the CBR folder) have several significant differences from those found by catalog scan, that the user should understand. This section endeavors to explain those differences and the recommended way to deal with them. We will use the following abbreviations for this section:

catalog-file = a file found by Data Rescue using a Catalog scan

content-file = a file found by Data Rescue using a Content scan

8.8.1 Locating a file found by a Catalog scan

A catalog-file will usually (but not always) show up in the recovery window with its original filename, and in its proper place in the directory structure tree (the one whose topmost node is the name of the volume or drive that was scanned). In this case, the user may locate the file in the recovery window by navigating through the directory structure to the place where he remembers the file to have been, or by searching for the file by name.

In some cases, a catalog-file’s folder will be unknown. This can happen for example, if the catalog information for the file’s directory is invalid. In these cases, Data Rescue synthesizes a new folder name, and puts it in the top level directory named “Orphans Folder”. If you have done a scan and haven’t found a file you are looking for in the normal directory structure, and an Orphan Folder is present in the top level of the recovery window, it is possible that your missing file might be among those orphans. It is also possible in some cases that the file’s original name may have been changed to

something else, such as “inode123456” by the operating system. These cases too may end up in the Orphans Folder. The recommended way to deal with these “inode” files is to recover them all, then use the Finder to see their type. If one has a type matching the type you are looking for, you can try to open it with your application to see if it is the one you want or not.

8.8.2 Locating a file found by a Content scan

As stated earlier, a Content scan finds only the file contents, not its name or folder. Therefore, content-files always have a synthesized name and folder. One thing that is known about the content-file is its type, and this type is used to generate a meaningful filename extension for the synthesized filename. In addition, files of a like type are grouped together in a synthesized directory named for that type. So for example, MP3 files found by a Content scan will be found in an MP3 subfolder, and the filename will be something like “M123.mp3”. In some cases, it is possible for Data Rescue to extract some additional information from the file contents to make the synthesized filename more meaningful. For example, a JPEG content-file will have a name like “J1200x768-123.jpg”. This not only identifies it as a JPEG file with the .jpg extension, but includes the pixel dimensions of the image (1200 by 768) as part of the name.

Data Rescue is currently not able to display previews of files in its recovery window, so the recommended procedure to browse content-files is to go ahead and recover all of them, then use the Finder to browse and preview them; or alternatively, double-click a file, which will automatically recover that file and ask the system to open it with the appropriate application.

It is a current limitation of Data Rescue that it cannot properly recover fragmented content-files. A fragmented content-file may appear like any other content-file in the recovery window, but upon recovering and trying to use it, you may discover that the contents are truncated, or contain unrelated garbage part way through the file, or (if the application is particular) that the file won’t open, or in rare cases, that the application will crash or the Finder will restart when you attempt to open or preview that file. In that case, you should just go ahead and delete that file, since it isn’t any good anyway. Data Rescue can’t always tell whether a content-file is in good shape or not. You can look at the right-most column of the recovery window, labeled “OK”. If that contains “bad”, then it’s almost certain that the corresponding content-file is bad. (We leave these content-files in the list because in most cases the first part of the file is good, and in some cases this may be better than nothing.) If it contains “good”, then there is a good chance that the content-file is good (but it still might be bad). If it contains a blank, then the condition of the content-file is unknown. (Note: catalog-files always have a blank in this column.)

9 Using Data Rescue – More Details

9.1 Scan Time

The scan process is mostly I/O bound, though a faster CPU will usually result in somewhat faster scans. The time needed to do a Thorough Catalog on a drive is

primarily determined by the speed of the drive, the size of the volume, and whether the disk has many bad blocks, or not and to a lesser extent, by the speed of the drive used for temporary storage. To give you an idea, the Thorough Scan time for one machine with a 40GB internal IDE volume with no bad blocks, and a Firewire external drive for temporary storage was about an hour. If the disk has many bad blocks, this time can be lengthened considerably.

9.2 Block Level Tools

Data Rescue features block-level tools to display raw disk blocks. Most users won't need this, but if you are an advanced user, you can use it to view absolute drive locations, such as viewing the partition table at the beginning of a drive, or viewing the data at the start of a file.

To display a block, select **New Window With Block#...**, and then enter the number of the logical block you want to display, relative to the start of the drive or volume you selected earlier as your source.

You can also read the first block of the data or resource fork of a file by highlighting this file in the recovery window and selecting **New Window With This File's DATA** (or **RSRC**). In that case, the logical block read depends on the allocation blocks layout. If the displayed data appears incorrect, try one of the other ABL choices in the Allocation Blocks Layout dialog. The ABL window automatically updates any open block display windows, so you can immediately see the effect of choosing a different ABL.

Note: Normally the block window you open in this way remains open showing the start of the file that was selected (highlighted) when you opened it. You can open a special “inspector” version of the block window if you hold down the Option key while clicking on the **Blocks → New Window With Block#...** menu item. When opened in this way, the block display window will track whichever single file is currently selected (highlighted) in the recovery list. This allows you to quickly look at the starts of a number of files without having to bring up a separate window for each one.

The next figure shows an example block window.

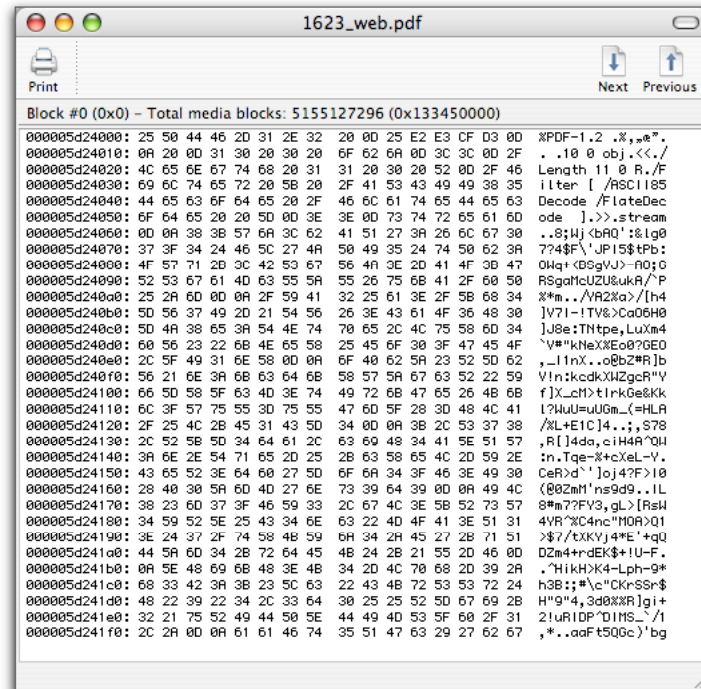


Figure 35: Block Content Window

The first number in a row is the media address of the data (in hex). Then follow the block's data in hex and its character representation. You can view the next or previous block clicking the appropriate icon in the toolbar. These commands have keyboard equivalents Command + and Command - for convenience. You can also jump to a particular block by selecting **Go To Block#...** in the same menu.

Note that you can't edit a block in the block display.

10 Troubleshooting and Frequently Asked Questions

The following chapter contains some commonly asked and troubleshooting questions and their solutions or workarounds. If problems still persist, contact Prosoft Engineering Technical Support via our website at <http://www.prosofteng.com>

10.1 Data Rescue seems to freeze. What should I do?

Most cases of apparent freezing are actually slow reads due to disk problems. If Data Rescue runs into a disk area where the disk is having trouble reading, it may appear to freeze, but actually just be very slow. If you think it has frozen, make a note of what is being displayed on the progress display, for example, the block number, then check again in half an hour or so and see if it has advanced at all. If not, it is probably just slow due to disk read problems. If so, then contact Prosoft technical support for assistance. Here are a couple more tips:

- Make sure cables connecting to disk drives are properly connected.
- Disconnect or turn off unnecessary devices.

10.2 Why doesn't my drive or volume show up?

If you have attached a drive with a damaged volume since you started Data Rescue, Data Rescue may not automatically notice it. Try the Expert → Refresh Volume List menu item. If the volume you are looking for still does not appear, it may be because Data Rescue is unable to find the correct name for it, in which case it may show some other name such as “unknown” for the volume name. Or it may be that Data Rescue can see the device, but is not able to recognize a volume structure on the device. If that is the case, you will need to choose the device name for scanning rather than the volume.

Finally, it could be possible that your drive is malfunctioning to the point where your computer is not able to talk to it at all, in which case even the device name will not appear. In this case, Data Rescue will not be able to scan your device. In this latter case, you might try some or all of the following things to see if your device can be made to appear: Double-check the drive cables and power source; remove/reattach the drive and/or power cycle it; power down and restart your computer. You can also double check that the problem isn't with the Data Rescue software itself by looking for your volume/drive using Apple's Disk Utility. If Disk Utility can see it, Data Rescue should be able to also; if Disk Utility can't, then Data Rescue won't be able to either.

10.3 Why are some files listed twice?

Sometimes, your volume contains obsolete catalog entries. If Data Rescue finds two catalog entries with identical names, it first checks if the two entries refer to the same data. If so, then the items are called “true duplicates”, and only one of these will be retained, and the other automatically removed. If the two entries refer to different data, both of them are retained, even if they have identical names and appear to be in the same directory, because there is no way to know which one is the correct one. To find which is the correct entry, check dates and file contents.

10.4 Why do my found files total more than the size of the original drive?

There could be two reasons for this, as explained below.

When you do a Thorough scan, Data Rescue uses two different algorithms to locate files. (See the explanations of catalog and content scans in section 8.) These methods will often locate many of the same files twice – once under a catalog folder and again under the CBR folder. (Data Rescue does not currently have the capability to automatically cross correlate these sets of files, but may do so in a future version.) So if you elect to recover everything, it will often be the case that the total space required exceeds the original media size.

The second reason why the found files may total more than expected is the possibility of anomalously (and incorrectly) large files. In the course of scanning the media, Data Rescue will often come across bad files and catalog entries. Data Rescue is able to filter out the vast majority of these bad entries, but not all of them. Occasionally a few of these may show up in the recovery list with incorrect and large sizes. If you suspect this may be the case, you can easily find these large files by searching for files greater than a certain size using the Edit → Find menu item. A useful technique to eliminate these from the recovery is: first mark everything by clicking the checkbox for the top level folders, then search for and uncheck the large files which appear to be bogus.

10.5 Why is there a delay when I click the mark checkboxes?

After scanning a large drive, there may be hundreds of thousands or millions of files and folders to deal with. When you click on a mark checkbox in the recovery list to mark or unmark files for recovery, especially the top-level folders, Data Rescue has to walk recursively down through most of these items to individually mark or unmark them. The time this takes is proportional to the number of found files.

10.6 Does Data Rescue handle RAID volumes?

Yes and no. With RAID, several component drives are set up to act logically as if they were a single composite drive. If your RAID drives or file system is in good enough shape that the system can still make them appear as a single logical drive, then Data Rescue should be able to scan that composite drive and recover files from it.

If the situation is such that the system cannot present the component drives as a single composite drive, then the answer depends on what kind of RAID your drives are set up to do. If the drives are set up to do mirroring (i.e. each file is stored completely on each mirror drive), then Data Rescue should be able to Thorough-scan any of these component drives and recover files. If the RAID setup is striped, so that a single file's data is stored spread out over more than one drive, then Data Rescue will not be able to recover most files from the component drives. In other words, Data Rescue does not have the ability to mimic the behavior of a striped RAID controller in order to scan the individual component drives on its own. Such a capability may be added to a future release.

10.7 I recovered files, but they won't open. What's wrong?

The answer to this question depends on whether you're talking about files recovered from under the "CBR" folder, or other ones.

10.7.1 Files recovered by catalog (i.e. not from the CBR folder)

It's normal for a small percentage of recovered catalog files to be bad. But if you have checked a number of them and none of them are good, then chances are that they were recovered with the wrong ABL setting. Pick out a few files to use as test files, and go back and select a different ABL choice (Expert → Allocation Blocks Layout... menu) and recover and test those few. If you find an ABL setting that gives good results on those files, re-do your big recovery with that setting. For more information, read the material on Allocation Block Layout in sections 7 and 8.

10.7.2 Files recovered by content (i.e. from the CBR folder)

It's also normal for a small percentage of these files to be bad. If the original file was fragmented (not stored on disk in consecutive blocks), then it can not be properly recovered by content. In most normal situations, most files on a user's disk will not be fragmented. Typical fragmentation rates for files tend to be just a few percent.

Note: The ABL setting plays no role in recovering files under the CBR folder.

10.8 Can I recover files that I deleted accidentally?

In OS X, when a file is deleted, for example by dragging it to the trash and emptying the trash, the file's name and folder information is usually erased by the system, and is therefore irretrievably lost. However in most cases, the most important part of the file – its contents – are still present on the disk. Finding the file by its contents alone is in general a difficult task, requiring algorithms to recognize those contents among all the billions of bytes of data on the disk. This task requires algorithms that are tailored to each specific file type. Data Rescue X made a start at this by providing support for about ten important file types that can be recovered by content. Data Rescue II increases this content scan capability by an order of magnitude, supporting more than one hundred separate file types. (See the following questions.) If the file types you are looking for are among the supported types, there is a good chance that Data Rescue will be able to find it.

Beyond the ten-fold increase in number of content scan file types supported, Data Rescue II enhances your ability to find deleted files in another way: the Deleted Files Scan scans only the free space portion of your drive, greatly reducing the number of files that you must look for to find particular ones.

10.9 What are Data Rescue's general content scan capabilities?

In generic technical terms, Data Rescue's content scan is potentially suitable for any file type that consists of a single fork. For files that have both a data and a resource fork, the forks may be recoverable separately by content scan, but Data Rescue has no means to connect the two pieces back together. Some "files" are in reality a collection of separate files which the Finder treats like a single file. This is called a "bundle". The individual

components of a bundle may be recoverable by content scan, but again Data Rescue has no automated way to associate these components back together into a bundle.

An additional requirement for a successful recovery by content is that the file's data not be fragmented. In other words, it must be stored from beginning to end in consecutive media locations. Fortunately, most files get stored on disk that way. OS X tries to store files in a non-fragmented way when it can. Still it is typical for a small percentage of files to be fragmented, and these will not recover properly when found by content.

Important: Do not mistakenly use a defragmentation utility (or any other program which will alter your disk) after losing your files, and prior to scanning with Data Rescue. Doing so will only make it less likely that you will be able to recover your files.

Note: The above limitations apply only to the files found by content. For files that are found by their catalog entries, the original bundles and forks may be properly recovered, and fragmentation is not an issue.

10.10 Specifically what types of files can Data Rescue II find by content?

The actual list of file types recoverable by content can be seen from within Data Rescue by going to the Data Rescue II Preferences menu and browsing through the File Modules tree. We expect to be adding to this list on an ongoing basis. The following was current as of the time of this writing.

Archives:

Backup, DMG, GZ Files, SIT, ZIP

Audio:

AAC, AIFF, Loop, MP3, STMP, WAV, iTunes Library, M4a, M4b, m4p

Documents:

AppleWorks, ArchiCAD, Bookmarks, ClarisWorks, Excel, FileMaker Pro, FreeHand, HTM, InDesign, Office, OpenOffice, PDF, PageMaker, PowerPoint, Programming, Quark, Quark EPS, Quickbooks, QuickCeph, Quicken, RTF-Mac Files, Web Archive, Word Document, iChat, iTunes Library

Images:

CRW, CaptureOne, Comic Life, EPS, GIF, Illustrator, Illustrator EPS, JPEG, NIB, ORF, PNG, Photoshop, Photoshop BMP, Photoshop EPS, Photoshop JPEG, Photoshop PCX, Photoshop TIFF, TIFF

Mail:

AppleMail, Entourage, Entourage v11, Eudora, Eudora 5, Outlook Express, Outlook Express 4.5

Misc:

ETC, LPG, MISC, Norton, Property Lists, RIFF, RSRC

Movies:

AVI, After Effects, DV, FC Project, Flash, MPEG, QTZ, Quicktime, iMovie

Text:

HTML, PLIST, Postscript, Rich Text, TXT, XML

11 Glossary

Allocation block: this is a block of file data within the file space of a volume. The size of allocation blocks is the same within a volume, but different volume can have different allocation block sizes. This size is always a multiple of 512 bytes so allocation blocks boundaries are always aligned with logical blocks.

Creator type: Associated with every file on a Macintosh is a four letter code that specifies the file's creator. The creator is typically the application that created the file to begin with. This code is hidden from the user, but is used by the file system in a number of ways. For one, along with the *File type*, it helps determine the icon to display for that file. In addition, when you double click on a document, it is the application that has the same creator code as the document that will be launched to open that document.

File type: Associated with every file on a Macintosh is a four letter code that specifies the type of file it is (e.g. 'TEXT' denotes a text file, 'APPL' denotes an application). On many other operating systems, the *file type* is part of the name; it is usually a three-letter code following a period (e.g. myfile.txt or letter.doc). On the Mac, this code is hidden from the user but serves the same purpose: to identify the kind of file. It uses the *file type*, along with the *Creator type* (see below) to determine the icon to display for a given file.

HFS, HFS+: This stands for "Hierarchical File System" and "Hierarchical File System – Plus". These are the names of Apple's proprietary File system layout for Mac OS computers. The newer, HFS+ system affords greater file and disk capacities, as well as performance improvements over that of the older HFS architecture. HFS+ was released with Apple's 8.1 operating system and is designed to supersede HFS. All Mac OS versions from 8.1 Classic through OS X 10.x support both HFS and HFS+.

Orphan: An *orphaned file or folder* is a file or folder which Data Rescue has discovered a catalog entry for, but no parent catalog folder can be found. Items like this are presented under the Orphans Folder in the Data Rescue recovery window.

Physical block: this is a 512-byte block of information on a physical device. Physical blocks start at zero so physical block #0 is usually the first block of the partition map of a physical device.

Remote volume: A *remote volume* is a volume that is not directly connected to your computer. Typically a *remote volume* is a file server. If your Macintosh is not connected to a computer network (e.g. at your company), you probably don't have access to any *remote volumes*.

- Activation, 14
- allocation block
 - defined, 58
- Allocation blocks, 48
- Allocation blocks layout, 40
- Analyzing drive performance, 29
- Assistant mode, 15, 16
- Authorization, 14
- block-level tools, 51
- Catalog*, 53
- Catalog scans, 45
- Checking before a large recovery, 24
- Cloning a drive, 32
- Concepts, 41
- Content scans, 46
- creator type
 - defined, 58
- Data Rescue
 - and RAID volumes, 54
 - general uses for, 12
 - installing and starting, 13
- Deleted files, 47, 55
- Deleted Files scan, 18
- Demo mode, 9
- dictionary. *See* Glossary
- Differences from other tools, 8
- drives vs volumes, 42
- Drives with hardware problems, 12
- Expert mode, 27
- File systems and catalogs, 43
- file type
 - defined, 58
- glossary, 10, 11, 58
- index, 10
- Install, 13
- Installation and setup, 8
- Introduction, 10
- Latest version, 8
- Launching and Activating, 9
- Locating files after a scan, 49
- orphaned file
 - defined, 58

- Overview, 8, 10
- physical block
 - defined, 58
- Preferences, analyze, 35
- Preferences, recovery, 35
- Preferences, scan engine, 34
- Preferences, scanning, 35
- Preferences, user interface, 34
- Preferences, file modules, 36
- Previewing files, 24
- Quick scan, 18
- Recovering files and folders, 23
- remote volume
 - defined, 58
- Saving a scan file, 39
- Scan drive or volume, selecting, 16
- Scan file, 47
- Scan method, selecting, 18
- Scan results summary, 21
- Scan results, expert mode, 28
- Searching for files and folders, 24
- selecting a volume or drive, 42
- Starting, 14
- Supported file types
 - content scan, 56
- Supported file types, content scan, 46
- System requirements, 11
- table of figures, 6
- Technical support, 9
- temporary file space, 13
- Temporary storage, selecting, 17
- Thorough scan, 18
- Troubleshooting, 10, 53
- Uninstall, 13
- Uninstalling, 8
- Use Scan File, 18
- User Guide, 10, 11
 - about, 10
 - conventions used in, 10
 - user level, 11
- viewing preferences, 33