Data Rescue User's Guide


PROSOFT
engineering, inc.

Licensing and Copyright Information

4725 First Street, Suite 270, Pleasanton, CA 94566

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 Registering the Software 9
1.7 Technical support 9
1.8 Contact Prosoft Engineering 9
1.9 Data Rescue Demo 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
2.5 General Usage 12
3 Installing and Starting Data Rescue 12
3.1 Install / Uninstall Data Rescue 12
3.2 Before you start Data Rescue 13
3.3 Starting Data Rescue 13
4 Using Data Rescue – Quick Start 15
4.1 Viewing Preferences 16
4.2 Selecting a Volume or Drive to Scan 17
4.3 Choosing the Scan Mode 18
4.4 Starting the Scan 20
4.5 Locating your Files after a scan 24
4.6 Recovering Files from Digital Cameras and Music Players 27
5 Concepts 28
5.1 File Systems and Catalogs 28
5.2 Catalog Scans 30
5.3 Content Scan 30
5.4 Scan File 31
5.5 Allocation Blocks and Media Blocks 32
5.6 Locating Your File After a Scan, Revisited 33
6 Using Data Rescue – More Details 35
6.1 Scan Time 35
6.2 Content Scan Control Settings 35
6.3 Verifying the Correct ABL before Recovery 39

6.4 Examining Files with the Block Content Window 39
6.5 Recovering Files 41
6.6 Save Scan Option 42
6.7 Setting Data Rescue Preferences 43
6.8 Viewing the Data Rescue Report 45
6.9 Allocation Block Layout Window 45
6.10 Block Level Tools 51
7 Troubleshooting 53
7.1 Data Rescue freezes/crashes. What should I do? 53
7.2 Why are some files listed twice? 53
7.3 Does Data Rescue handle RAID volumes? 53
8 Glossary 54

Table of Figures

Figure 1: Data Rescue Start Screen 14
Figure 2: Authentication Window 14
Figure 3: Register Data Rescue Window 15
Figure 4: MainWindow 15
Figure 5: Preference Window Settings 16
Figure 6: Drives and Volumes 17
Figure 7: Example Popup with Ignore Partitioning off. 18
Figure 8: Example Popup with Ignore Partitioning on 18
Figure 9: Content Based Recovery control for image files 19
Figure 10: Content Based Recovery control for text files 19
Figure 11: Thorough Catalog Scan mode selected 20
Figure 12: Choosing a Temporary Storage Volume 21
Figure 13: During the Scan 22
Figure 14: Allocation Blocks Layout alerts 23
Figure 15: Allocation Blocks Layout Window 23
Figure 16: Recovery Window 24
Figure 17: Browsing Catalog Scan Results 25
Figure 18: Browsing Content Scan Results 26
Figure 19: Selecting a Destination for Recovery 27
Figure 20: File Recovery Progress Window 27
Figure 21: File System Data Storage and Catalog Structures 29
Figure 22: Media and Allocation Blocks 32
Figure 23: Content Scan Control, Images Tab 35
Figure 24: Content Scan Control, Audio Tab 36
Figure 25: Content Scan Control, Movies Tab 37
Figure 26: Content Scan Control, Text Tab 37
Figure 27: Content Scan Control, Docs Tab 38
Figure 28: Content Scan Control, General Tab 38
Figure 29: Recover Files Window 40
Figure 30: Block Content Window 41
Figure 31: Use Scan File 43

Figure 32: Preferences Dialog Box 44
Figure 33: Data Rescue Report Window 45
Figure 34: Allocation Block Layout (ABL) Window 46
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 10.0.0 is useable with versions 10.0.1, 10.0.2, etc.

1.4 Program Setup

On the CD-ROM, you will find the application folder <Data Rescue Folder>. Drag this folder into your hard disk (you can also download Data Rescue 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 Registering 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
  4725 First Street, Suite 270
  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, allowing 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 500k. 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:

Convention	***Description***
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 `monospace` 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: `<logicalblock> = (<allocationblock> x <allocationblocksize>) + <filespaceoffset>`
Text in *italic* type.	Text of all words that can be found defined in the Glossary is printed in italic type: *Allocation block*
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

Data Rescue X is optimized to run on Macintosh computers capable of running Mac OS X. 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 a version of Data Rescue that will work with your system. Data Rescue needs at least 128 MB of RAM to run.

2.4.2 Software

Data Rescue X requires Mac OS X version 10.1 or later to run (10.1.5 or later is recommended) 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).

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

1. From the location to which you downloaded Data Rescue, double-click on the **Data Rescue** image file to mount it on your computer.
2. 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.

3.1.2 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.

3.2 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.

3.2.1 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.

3.3 Starting Data Rescue

Launch the application by double-clicking the Data Rescue icon in the Data Rescue folder. The Data Rescue start screen opens, as shown in Figure 1.

Figure 1: Data Rescue Start Screen

Click the Red Plus Sign to continue.

Because Data Rescue X needs unrestricted access to your entire file system, it must run with special privileges. After clicking on the Red Plus Sign, a login window will appear, as shown in Figure 2. 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 2: Authentication Window

Now, you can register Data Rescue in order to unlock its complete functionality.

If this is the first time Data Rescue has been run from the hard drive, or if you are booted from the Data Rescue X bootable CD, a registration window will appear, as shown in Figure 3.

Figure 3: Register Data Rescue Window

If you have purchased a copy of Data Rescue, enter your name, organization and registration code in the fields provided, then click **Register**. If you only want to test Data Rescue, click **Demo**. 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.

An example main window is shown in Figure 4.

Figure 4: MainWindow

You are now ready to use Data Rescue.

4 Using Data Rescue – Quick Start

After starting Data Rescue, using it in general involves these steps:

1. Set/change preferences.
2. Select a volume or drive to scan.
3. Select a scan mode.
4. Start the scan and wait for it to finish.
5. Verify/select the proper allocation block layout, if needed.
6. Locate the file(s) you wish to restore.
7. Restore the file.
8. Repeat, as needed.

The following sections aim to guide you through these steps, and make some recommendations for the choices you will need to make. In the interest of making the process as easy as possible for the most users, these steps are presented as a simplified procedure in a sort of cookbook fashion without delving into the reasons and details. The recommendations supplied should suffice for most users. However the number of ways that a drive or volume can be corrupt are endless, and not all can be successfully dealt with this way. If you experience difficulty finding or restoring your files, don't give up yet -- you may just need to depart from the simplified procedure. Then you will need to refer to the later sections which explain more about the different Data Rescue options and how to use them. If you get stuck or need help with something, you can contact our customer support for help.

4.1 Viewing Preferences

You can bring up the preferences window by choosing the Preferences item of the Data Rescue menu. Data Rescue remembers these settings from the last time it was run, so it's worth taking a quick look at the settings before starting a scan to make sure they're set the way you expect. Figure 5 shows the recommended (and default) preference settings. Except for the "Ignore Partitioning" item (see the following section), you should normally not need to change the default settings.

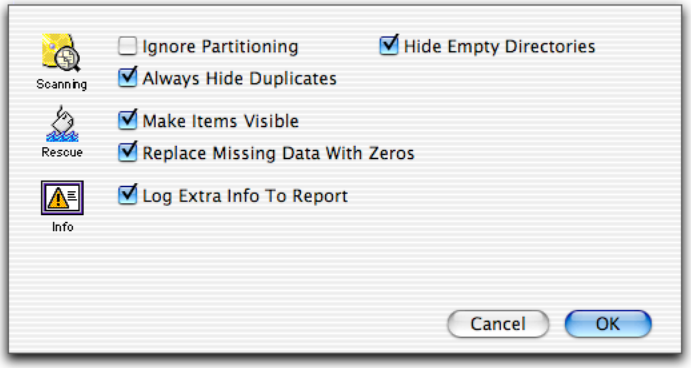

Figure 5: Preference Window Settings

For an explanation of the other preference items, see the later section on preference details.

4.2 Selecting a Volume or Drive to Scan.

4.2.1 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.

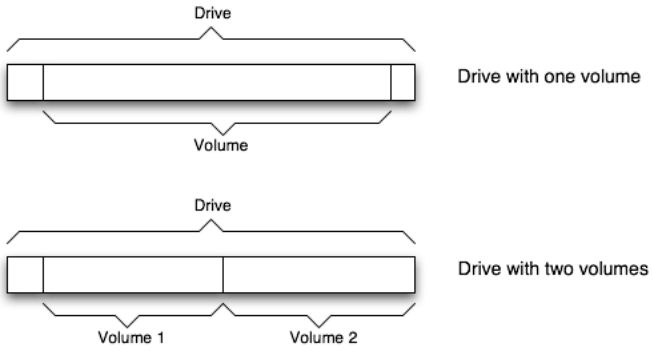

Figure 6: Drives and Volumes

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 popup menu of the main window. This list is affected by the setting of the “Ignore Partitioning” checkbox in the Data Rescue Preferences window. If this box is checked, Data Rescue will list only drive names in the popup menu; if unchecked, Data Rescue will list volume names for those drives whose partitioning it understands. The following figures show what the popup looks like on an example system with 7 attached drives for the different Ignore Partitioning settings.

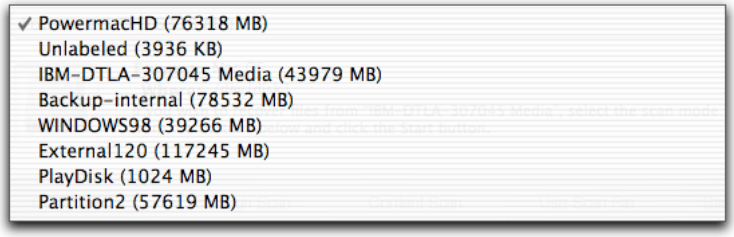

Figure 7: Example Popup with Ignore Partitioning off.

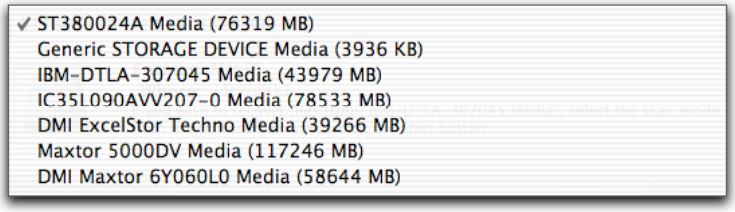

Figure 8: Example Popup with Ignore Partitioning on.

The 1st drive has a hardware drive name of "ST380024A Media" which has a single volume on it named "PowermacHD". The 3rd item ("IBM-DTLA-307045 Media") displays the drive name in both cases, because Data Rescue doesn't recognize the partitioning on this drive. The 7th drive on the system ("DMI Maxtor 6Y060L0 Media") is partitioned into two volumes: a small one named "PlayDisk" and a large one named "Partition2". That is why the first figure shows 8 items instead of 7.

4.2.2 Choosing the "Ignore Partitioning" preference.

Recommendation: Start with "Ignore Partitioning" off (i.e. uncheck the box in the preferences window). After doing this, see if you can find the name of your volume (the one you want to recover from) in the main window's volume selection popup menu (on the Main Window, under the text "Select Drive to Recover").

Under some circumstances, for example if the drive's partition table is corrupt or not understood by Data Rescue, you may not be able to find your volume in the popup menu. In that case, try the Refresh Volume List menu item (in the Expert menu).

If you still can't see it, set the "Ignore Partitioning" option. Then look for a drive name in the popup menu which corresponds to the drive you want to recover from. It's not always easy to tell which drive name is the one you want, so Data Rescue also includes the drive size in the popup menu as an additional clue.

4.3 Choosing the Scan Mode

After finding and selecting your volume, next choose a scan mode. To select the scan mode, simply click the corresponding main window button(s).

Recommendation for HFS/HFS+ volumes: Start with a Quick Catalog scan, if you can. If this does not find your files, or if it cannot be done (Quick scan is only allowed in

certain cases), and your volume is HFS or HFS+ format, then do a Thorough Catalog scan. If the Thorough Catalog scan can't find your files, and you are looking for specific file types that are supported by the Content scan, try that. You can select Thorough Catalog and Content scans concurrently if you wish. The combination of Thorough Catalog and Content scans gives the most thorough scanning that Data Rescue has to offer.

Recommendation for non-HFS/HFS+ volumes: Currently, the only option available in Data Rescue for recovering files from non-HFS/HFS+ volumes is the Content scan.

Before initiating a Content scan, bring up the "Content Based Recovery…" dialog from the Expert menu, and make sure that the file types you want to recover are enabled. Figure 9 shows the dialog with the Images tab visible.

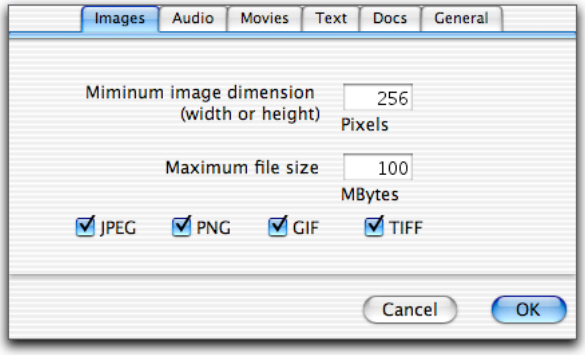

Figure 9: Content Based Recovery control for image files

Figure 10 shows the dialog with Text tab visible.

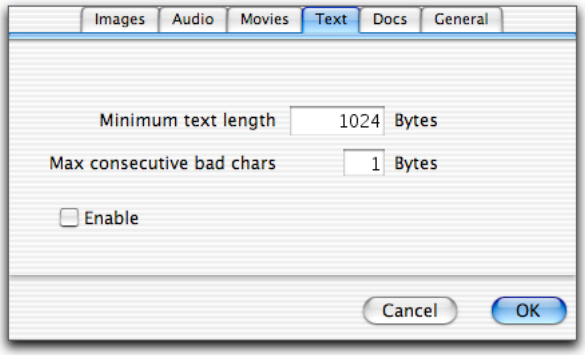

Figure 10: Content Based Recovery control for text files

Recommendation: For best performance in a scan, leave the Text files disabled unless you are particularly looking for a text file.

If you do a Content scan, be aware that only certain file types are recovered with this method, and these will not have their original file and folder names – see the section on locating your recovered Content scan files.

Figure 11 shows both the Through Catalog and Content scans selected and ready to start scanning the selected HFS+ volume “PlayDisk”.

Figure 11: Thorough Catalog Scan mode selected

4.4 Starting the Scan

To start the scan, just click the “Start Now” button. You will be presented with some dialog boxes similar to the following. Answer “Yes”, then choose the volume you want to use for temporary storage. Note: Don’t choose the volume you want to scan. It is also best not to choose a volume on the same drive that is being scanned.

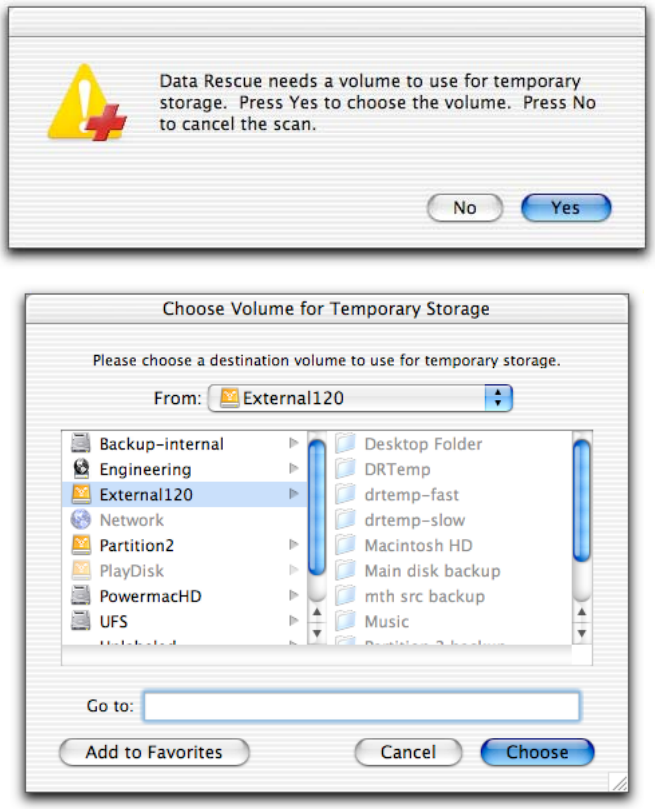

Figure 12: Choosing a Temporary Storage Volume

After you choose the temporary storage volume, the scan will begin. Below are some example pictures of Data Rescue during different phases of the scan. The whole process can take anywhere from a few minutes to several hours, depending on the size of your volume or drive and other factors. The rightmost numbers above the progress bar give an estimate of the minutes and seconds remaining in the current phase of the scan. In the first picture, the numbers in the middle (e.g. “75+693”) indicate how many potential file and folder items Data Rescue has identified for later analysis. The first number is generated by the Catalog scan and the second by the Content scan.

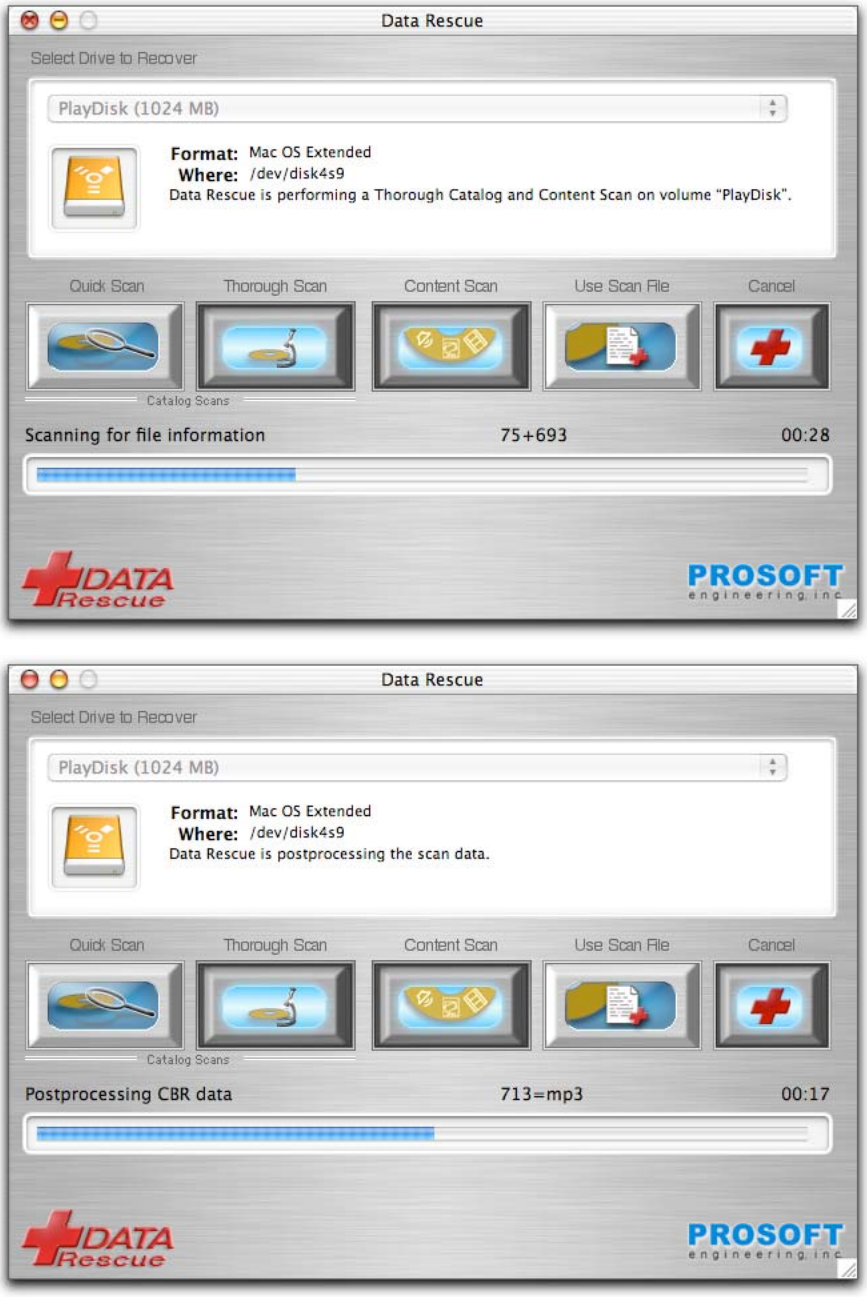

Figure 13: During the Scan

In the second picture, Data Rescue is busy analyzing items found during the Content scan. The example "713=mp3" indicates that it is working on the 713th item, which is an MP3 file.

You can cancel Data Rescue at any time during scan processing by clicking the Cancel button. Note: A canceled scan cannot be resumed. If you quit the scan before it is complete, you will not be able to recover or save any files. To go back to the original

window after a scan has finished or been canceled (e.g. to choose a different volume), use the File/Recover New Volume menu item.

4.4.1 If you get an "allocation blocks layout" alert

If at the end of the scan process you get either of these alert boxes, then Data Rescue needs help determining the allocation blocks layout:

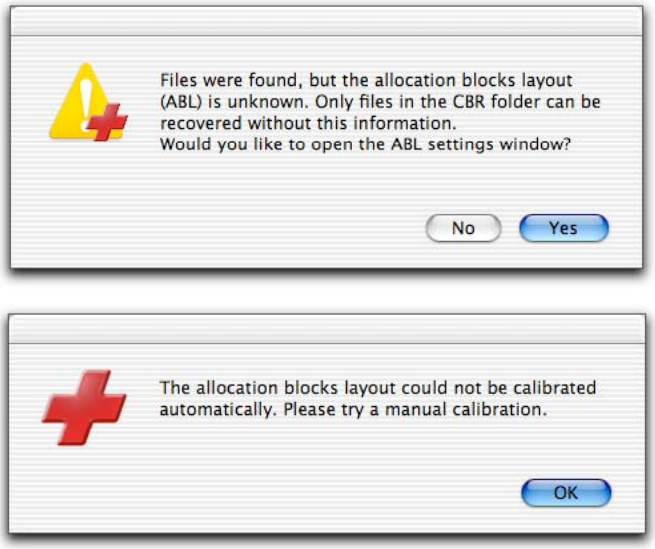

Figure 14: Allocation Blocks Layout alerts

Answer "Yes" or "OK" to these. This will cause the ABL window to come up, as shown in the next figure.

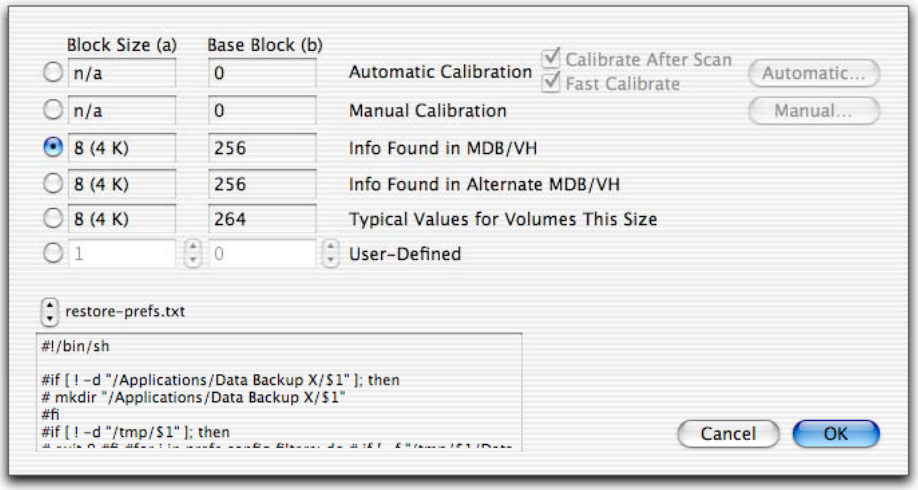

Figure 15: Allocation Blocks Layout Window

Recommendation: Look at the numbers in the Block Size and Base Block columns for the 3rd radio button (labeled "Info Found in MDB/VH"). Compare these to those just

below (labeled “Info Found in Alternate MDB/VH”). If these agree, and are not shown as “n/a”, then the chances are good that they are the correct values. Choose either the 3rd or 4th radio button, and click “OK”.

If the numbers for MDB/VH and Alternate MDB/VH don’t agree or are “n/a”, you’ll need to refer to the later detailed section on Allocation Block Layout.

4.5 Locating your Files after a scan

Following the scan, the Recovery Window is displayed, as in the following figure.

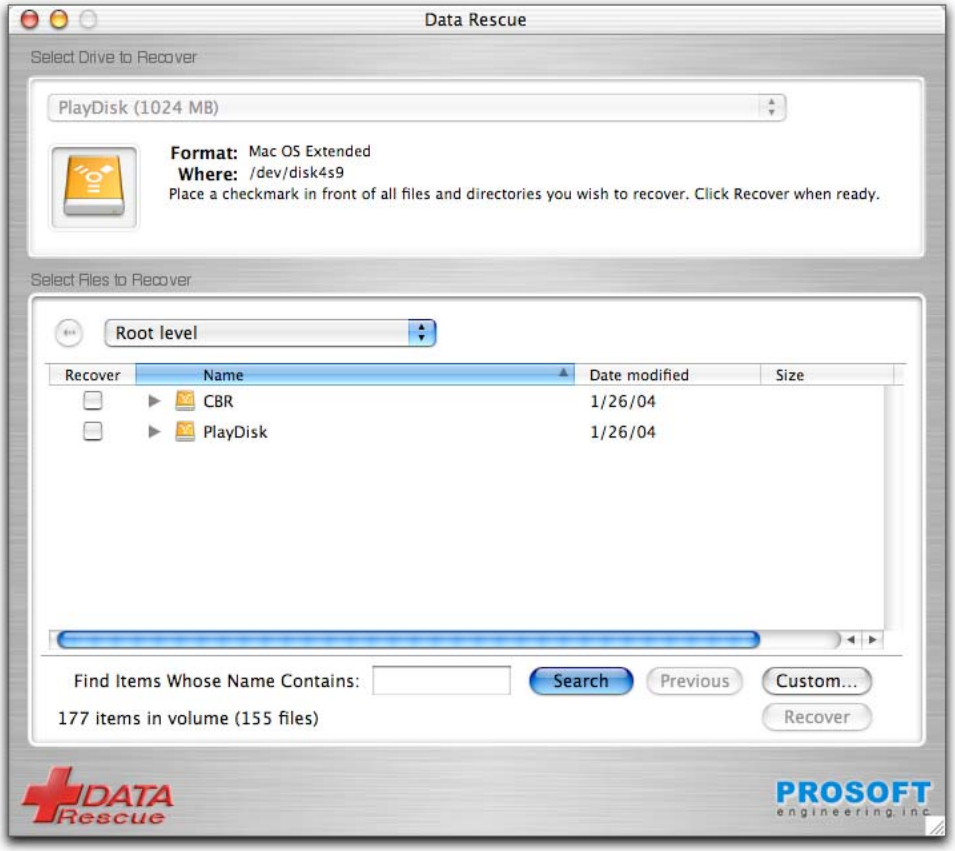

Figure 16: Recovery Window

Recommendation: If your scan took a long time to complete, we recommend that before proceeding further, you save a copy of the scan (File/Save Scan File… menu item). This will save you from needing to repeat the scan if you decide after quitting Data Rescue that you need to view the scan results again. *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.*

The Recovery Window allows you to browse and search for files that were found during the scan. When initially displayed, the window shows only top-level folders. Files found by a Content scan will be placed in subfolders under the top-level folder named “CBR”.

Files and folders found by a Catalog scan will be placed under a folder whose name is the volume or drive name that was scanned.

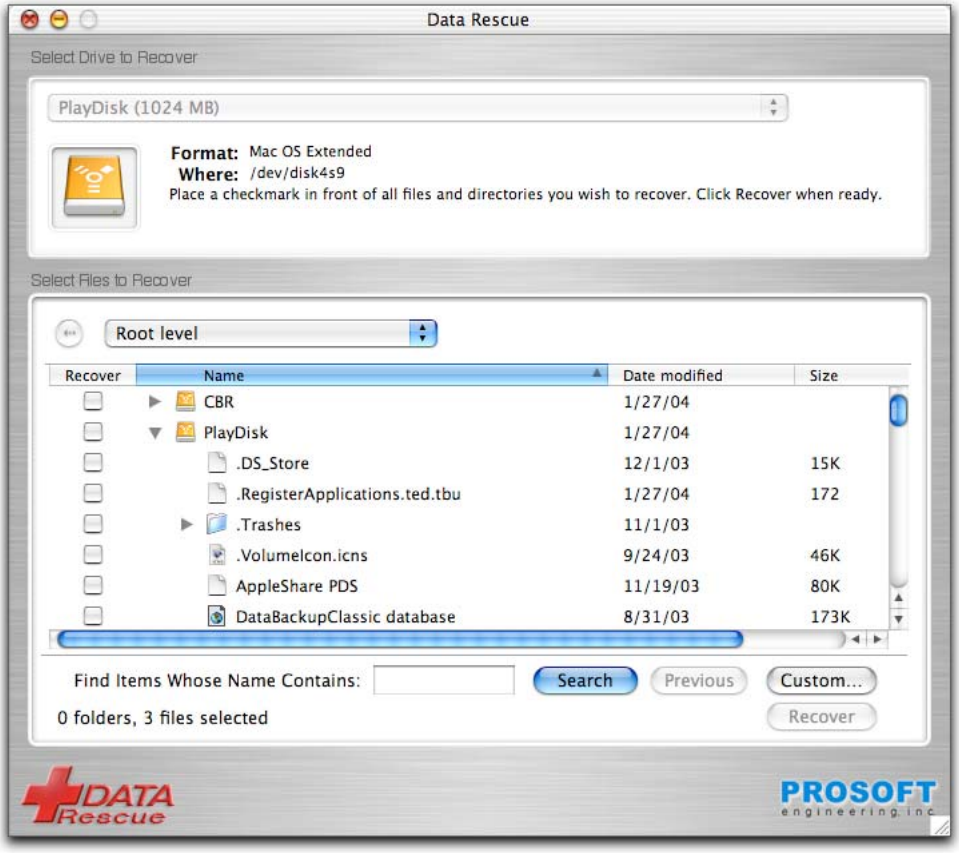

Figure 17: Browsing Catalog Scan Results

The preceding figure shows some files found by the Catalog portion of the scan, under the volume name “PlayDisk”. To select files and folders for recovery, click their Recover checkboxes. Note: If a folder is checked, all the files and folders under it are also checked.

The next figure shows some image files found by the Content portion of the scan, under the top-level name “CBR”. The file and folder names and dates under CBR are all synthesized (created) by Data Rescue.

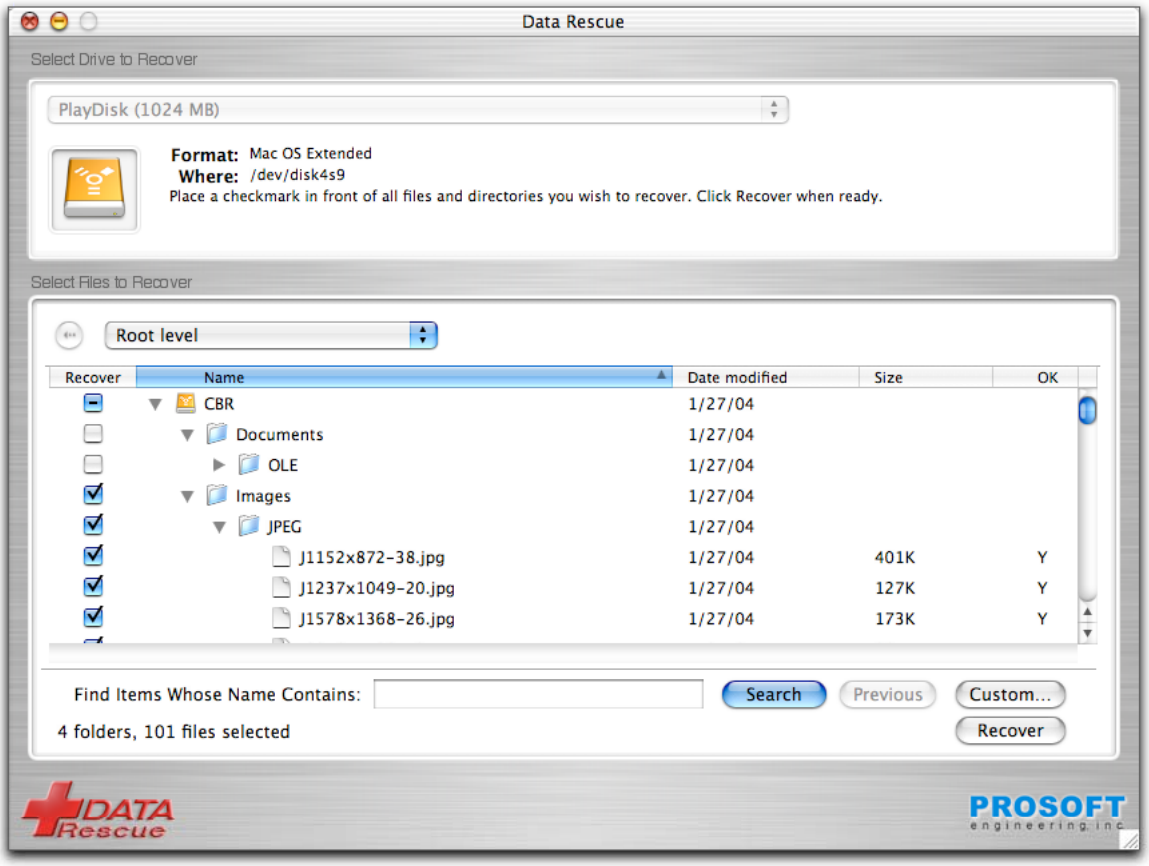

Figure 18: Browsing Content Scan Results

For image files, the pixel dimensions (e.g. 1152 by 872) are included in the synthesized names. The actual file size is also shown.

Recommendation: To locate the file you want in the CBR directory, it is best to recover whole groups of these files at once, then browse the results using the Finder or other appropriate Macintosh application. For example, to locate a particular image, recover the entire CBR/Images directory to disk, then use Finder or iPhoto to browse the files.

After you have set Recovery checkmarks on everything you wish to recover, click the Recover button to recover it all to disk. You will get a “Choose Destination” dialog similar to this:

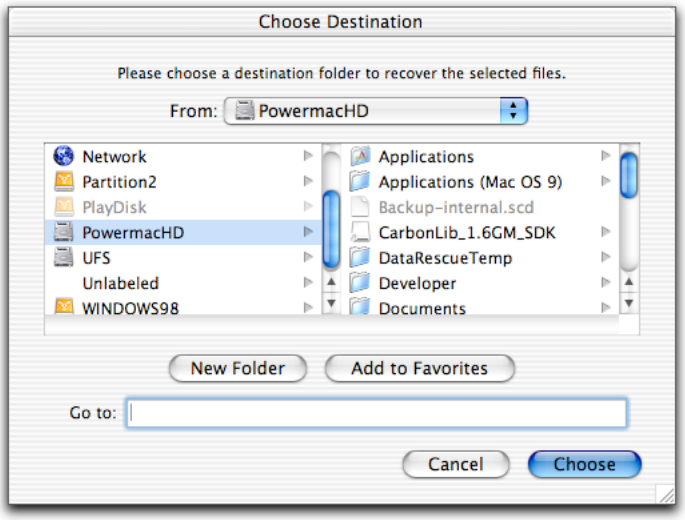

Figure 19: Selecting a Destination for Recovery

Choose a destination folder somewhere (not on the drive you scanned!), and click Choose. Note: If you have already done some recoveries earlier, it is okay to choose the same destination. You should see a progress bar window, as in the following figure:

Figure 20: File Recovery Progress Window

After the recovery you can go check the files you recovered (you do not have to quit Data Rescue for this). They will be in subfolders of your chosen destination folder.

4.6 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 Content scan to recover pictures or songs from it. Make sure that the image or audio types your player uses are enabled in the Expert / Content-based recovery dialog, and then do a Content scan. If your digital camera uses JPEG format for still images, and Quicktime or MPEG-1 for its movie clips (most do), then you should have success recovering these. 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.

5 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 Chaper 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.

5.1 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 21 shows part of a simplified generic file system to help illustrate the concepts.

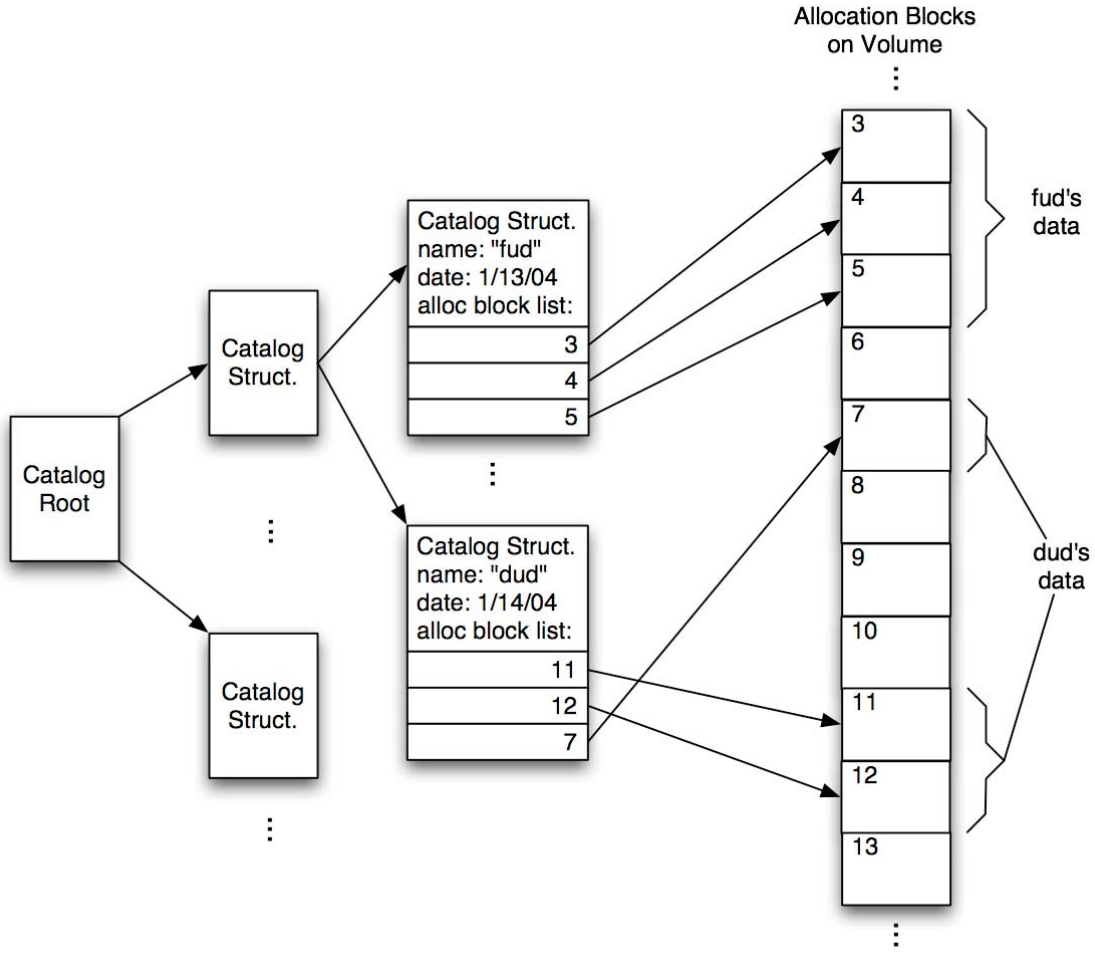

Figure 21: File System Data Storage and Catalog Structures

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 21 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 three Data Rescue scan types.

5.2 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 which 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 which point to files and folders.

The Thorough catalog scan is generally more useful, because it can find catalog structures that are not referenced by a (possibly damaged) catalog tree. The Quick catalog scan assumes that the catalog structures are valid and intact, and only looks for catalog information in the places they are supposed to be, i.e. starting at the catalog root, much like the operating system itself would. For this reason, the Quick catalog scan is often not useful for recovering files on file systems that have sustained more than very minor damage. Still, there are a few 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 usually 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 Through catalog 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.

5.3 Content Scan

The Content 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. Like the Thorough catalog 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 which is no longer referenced by any catalog structure, which catalog scans cannot do.

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:

IMAGES:	JPEG, TIFF, PNG, GIFF
MOVIES:	Quicktime, MPEG-1
AUDIO:	MP3, AAC/M4A
TEXT:	Generic ASCII text, RTF, XML, PLIST, Postscript (non-binary)

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.

5.4 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 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.

Note: In addition to scan files which you save yourself, Data Rescue also has the ability to load scan files that were produced by Data Recycler, another Prosoft Engineering Inc. application. Data Recycler saves its scan file in a folder named ".DataRecyclerCache" which is in a folder named ".DataRecycler Folder", which is in the root of the drive it scanned. The name of the file is ".DataRecyclerScanFile".

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. This could be an

issue if you are trying to use a scan file produced by Data Recycler. If you are trying to recover a file by using a Data Recycler scan file, and can't locate the files you want to recover, you should go ahead and do a traditional Data Rescue scan of the volume.

5.5 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 22 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.

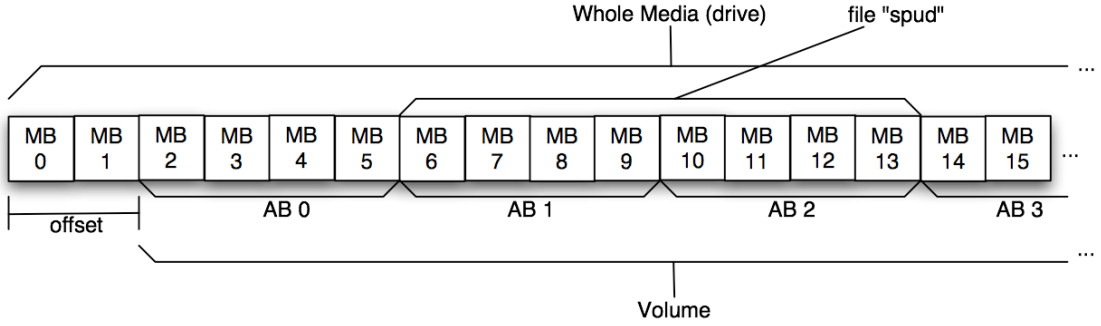

Figure 22: Media and Allocation Blocks

Earlier section 5.1, with the help of Figure 21, 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 22 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:

$$MB\# = AB\# \times size + offset$$

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.

5.6 Locating Your File After a Scan, Revisited

Content based scanning and recovery works in a different way from the traditional Data Rescue Catalog based scanning. As a result, there are several significant differences in the results 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

5.6.1 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 Catalog 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.

5.6.2 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 recovery all of them, then use the Finder to browse and preview them.

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 an "N", 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 a "Y", 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.)

6 Using Data Rescue – More Details

Chapter 4 presented a straightforward procedure for running Data Rescue which will suffice for most users. If you haven't read that, please do so before proceeding here. This chapter covers some options and details that were intentionally left out of that chapter.

6.1 Scan Time

The scan process is I/O bound. The time needed to do a Thorough Catalog or Content Scan 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 40 minutes. If the disk has many bad blocks, this time can be lengthened considerably.

6.2 Content Scan Control Settings

A Content Scan (see section 5.3) may be selected from the Main Window scan mode buttons as an option, either by itself, or concurrently with a Thorough Catalog scan. This section explains the options for controlling a Content Scan. You bring up the dialog via the "Expert/Content-Scan Control…" menu item. Most of the Content Scan Control tabs have two types of settings. The first type is an enable checkbox for each different file type. If the box is checked, then the Content Scan process will try to find that type of file; otherwise it won't. The second type of common setting is a numerical limit, usually to limit some size or length property of the file, in order to limit the files found to those of interest.

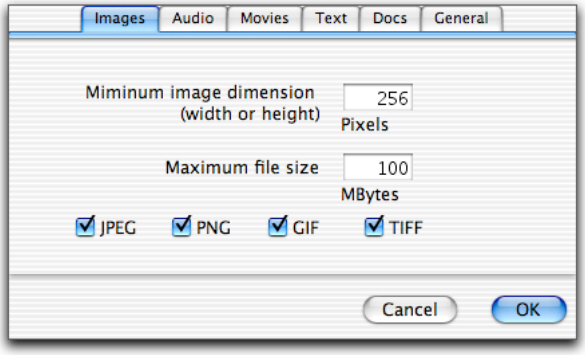

Figure 23: Content Scan Control, Images Tab

6.2.1 Images Tab

Figure 23 shows the Images control tab, which is displayed by default. The currently supported image types are shown in the figure. Check the boxes for the ones you are interested in recovering. JPEG format is particularly popular in digital cameras and web pages. If you don't know which one you want, just check them all (the default setting).

The Minimum image dimension limits the images found to those which have a horizontal and vertical pixel dimension that are both ≥ the number you put in this box. If you know the approximate smallest dimension of the images you are interested in, put a number a little less than that in this box. That way you won't have to wade through tons of smaller images when browsing the recovered results, and it will make the scan process faster too.

The Maximum file size is a safety feature to limit the maximum file size in case the end of file mark isn't found for certain image types (JPEG in particular). Set this box to a number which you're sure all the image files you're interested in will have a size less than.

6.2.2 Audio Tab

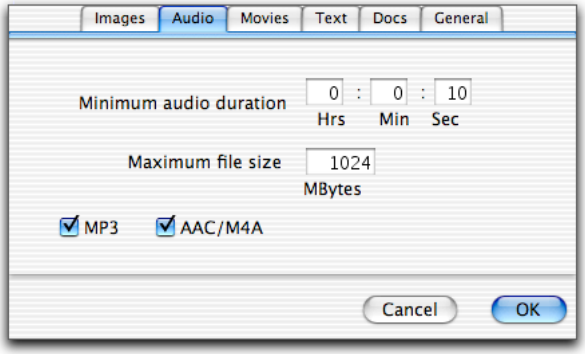

Figure 24: Content Scan Control, Audio Tab

The Audio Tab controls scanning for the supported audio types. The default setting is to have all the audio types enabled for scanning. MP3 format is popular for portable music players.

The Minimum audio duration number filters out music segments which are shorter than the specified time length. (Currently this applies only to MP3 format.)

The Maximum file size performs the same safety function for audio files as that for the Images Tab (see section 6.2.1). Set it to a number that you're sure is larger than the largest legitimate audio file you're looking for.

6.2.3 Movies Tab

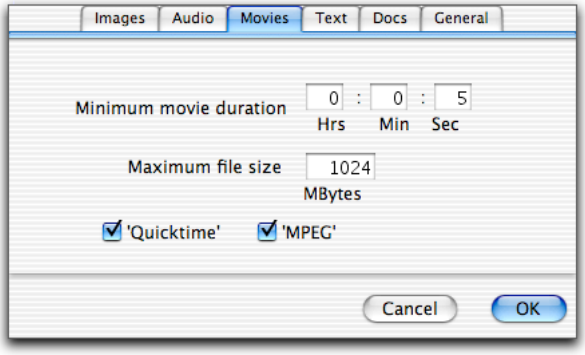

Figure 25: Content Scan Control, Movies Tab

The Movies Tab controls scanning for the supported movie types. The default setting is all enabled.

The Minimum movie duration and Maximum file size perform the same functions for movie files as those for the Audio Tab. You'll probably want to set the Maximum file size for movies larger than you did for audio, since movie files are typically much larger.

6.2.4 Text Tab

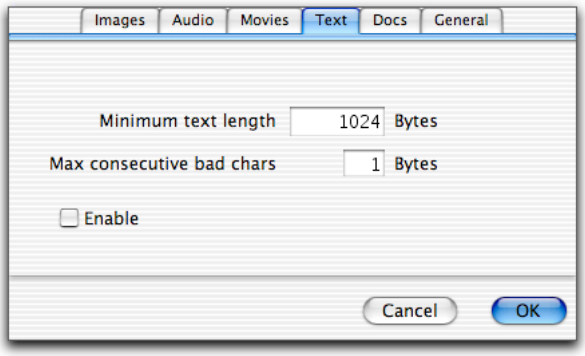

Figure 26: Content Scan Control, Text Tab

The Text Tab controls scanning for files which consist mostly of "good" ASCII text characters (letters, numbers, common punctuation marks, etc.). It looks for consecutive runs of these characters and considers such a run to be a file if the run length is ≥ the Minimum text length number. If more than "Max consecutive bad chars" characters are found consecutively, that is considered the end of the run also.

The Data Rescue recognizes certain patterns that occur at the beginning of a few common text file types, such as the pattern common at the start of HTML files. If one of those patterns is detected, then the file will be shown in the recovery window with a filename extension appropriate for that type, for example ".htm" is used for HTML text files. Data Rescue currently recognizes the following variants of text file: HTML, XML, PLIST, and RTF. Others will be added in the future.

If Data Rescue doesn't recognize a pattern in the text run, it is considered a generic text file, and will be shown with a ".txt" extension.

The recommended setting for this tab is disabled, unless you are particularly looking for a text file. This is because a typical large disk has many tens or hundreds of thousands of text runs on it, which slows scan processing considerably.

6.2.5 Docs Tab

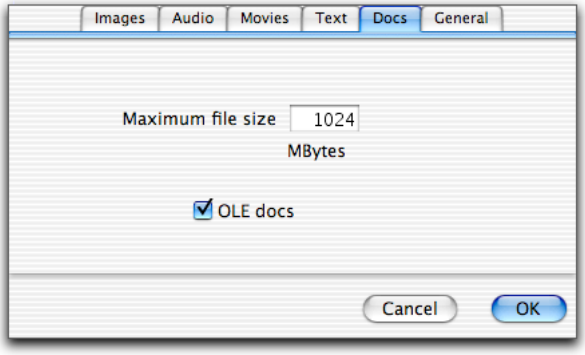

Figure 27: Content Scan Control, Docs Tab

The Docs Tab controls content scanning for structured documents. Currently Data Rescue supports only OLE/POIFS documents. This includes Microsoft Office documents that use OLE, such as Word, PowerPoint, and Excel. More document types may be added in the future. The default setting is all types enabled.

The Maximum file size field has the same function as for the other tabs.

6.2.6 General Tab

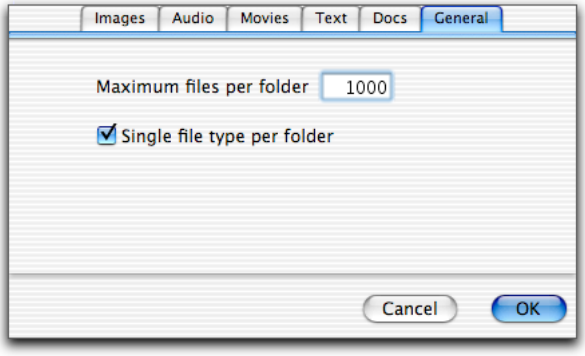

Figure 28: Content Scan Control, General Tab

The General Tab contains settings that apply to all the file types generally. It currently contains two settings, though more will probably be added in the future.

The Maximum files per folder setting controls how many files which are found by the Content Scan are put into a single directory. The default setting is 1000. If too many files are put into the same folder, browsing the folders becomes extremely slow. (This

comment applies both to browsing the recovery window from within Data Rescue, and also to browsing folders by the Finder and other applications after recovery to disk.) When the Content Scan would exceed this number, it automatically creates a new folder with a number appended to its name, and starts putting files in the new folder, etc.

The Content Scan puts files it finds under generic directory names that correspond to the tab name. For example, images all go under an "Images" folder, etc., either directly or in subfolders of the generic name. The Single file type per folder setting controls this choice.

If this option is unchecked, then all the file types of one tab will go directly into the folder for that tab; for example, different image types (JPEGs, GIFs, etc.) will all be stored together in the Images folder.

If the option is checked, then each file type is put under its own file type subfolder. For example, JPEG files are put in Images/JPEG, GIFs are put in Images/GIF, etc.

6.3 Verifying the Correct ABL before Recovery

This section only applies if you have done a Thorough Catalog scan (i.e., doesn't apply to Quick Catalog or Content scans).

After the scan is complete, Data Rescue will normally (by default, unless you have disabled it) run an automatic calibration routine to detect the proper ABL. If you are planning to recover many files, we recommend that you verify the ABL before doing the big recovery. Basically, if you can verify that a few files are correct, that verifies the ABL, since a bad ABL affects retrieval of data for all files in a volume.

There are two main ways to do this: Examine a few files in Data Rescue with the Block Content window (see section 6.4); and recover a few files and try them out (see section 6.5). The reason for checking more than one file is that even if the ABL is correct, some files found in the scan may simply be corrupt. Once you have verified some correct files, you may proceed with the recovery of all the files you need.

6.4 Examining Files with the Block Content Window

This method works for files which you know should be composed of (or at least start with) readable text, or files for which you can recognize the start of by looking at the hexadecimal data.

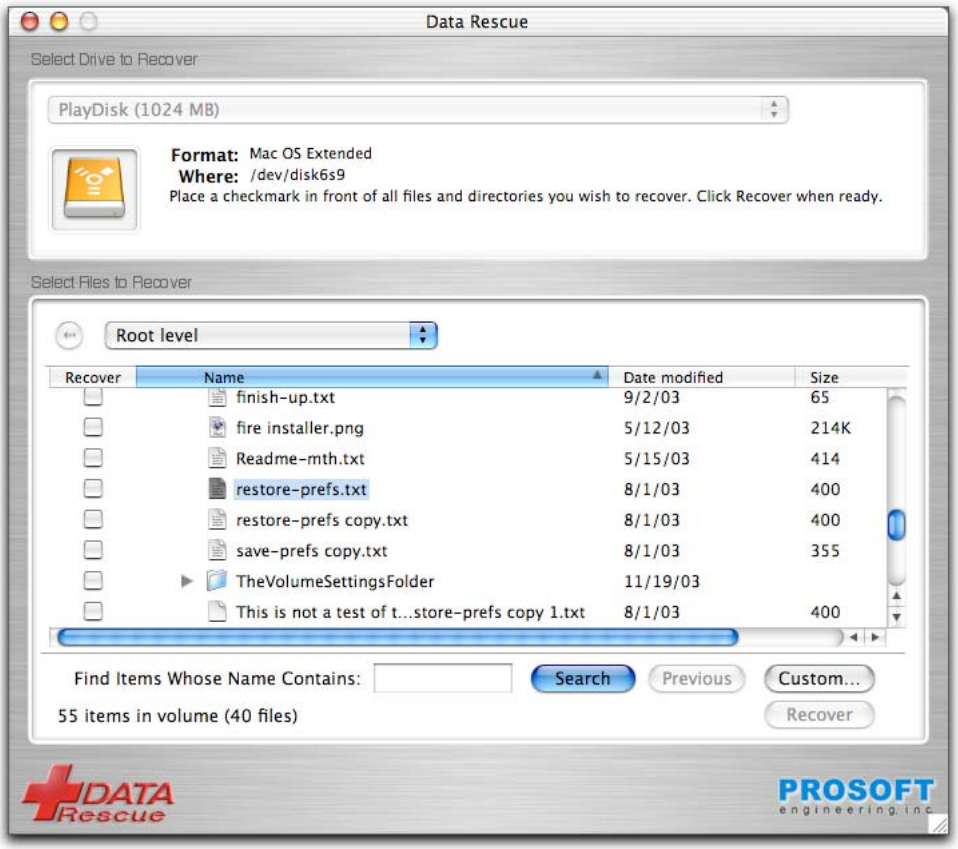

Figure 29: Recover Files Window

To use this method, select a file from the recovery window, as in Figure 29. (Note: The file to be examined is the one whose name is highlighted, not the one with the checkbox checked). From the **Blocks** menu, select **New Window With This File's DATA**, and check that the text in the rightmost column is what you expect.

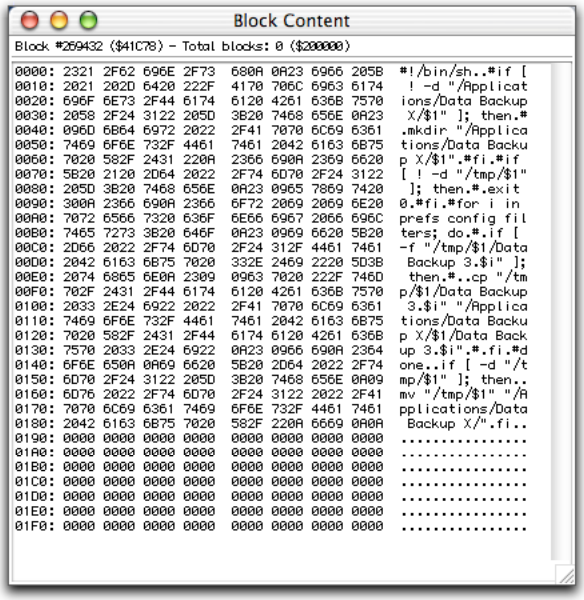

Figure 30: Block Content Window

Alternatively, you can perform these steps by finding a file that contains resources (e.g. an application, type "APPL"), then selecting **New Window With This File's RSRC**. The name of the file is usually written in the fourth row (rightmost column). Verify that this is correct.

You can also save blocks as a file, and examine arbitrary blocks on the device, not just those in a file. See section 6.10 for more details.

6.5 Recovering Files

Files found by Data Rescue are presented in a list with their names, size, and modification date. You can scroll up and down this list to look for your files. As in the standard Macintosh **Open** window, you can view the contents of folders by double-clicking them, and back out of the folder by using the hierarchy pop-up menu above the list. Alternatively, you can open a folder by clicking on its "disclosure triangle".

> *Note: there may be a folder named **Orphans Folder** present. This folder contains one or more subfolders, named Folder 001, etc. The items in these folders are files and folders which were found on the disk, but whose place in the normal hierarchy could not be found. In some cases, the items no longer have their original names, having instead names like "inode01237856", etc.*

Additionally, you can use the **Search** button to locate files having a particular name, size, date. By default, when the **Search** button is pressed, Data Rescue searches for file names that contain the text you type into the text box next to the button. This default behavior can be modified by using the **Custom…** button. For instance, if you are looking for files

that were modified since August 1, 1998, click the **Custom…** button and then set the **Find** menu to **Date ≥** and enter `8/1/98` in the **Find** box. Then click **OK** and then **Search** in Data Rescue's main window. Data Rescue will jump to the folder containing the matching items. If several folders contain matching items, use the **Next** and **Previous** buttons to move to other folders.

> *Performance tip: If you know the exact, full name of the file you are searching for, setting the Custom name search criteria to "Name is" (rather than the default "Name contains") will speed up the name search considerably.*

When you have located the file(s) or folder(s) that you want to recover, select them by placing a checkmark in the section box adjacent to the file names. When all your desired files are check marked, simply click the **Recover** button. You will be asked for a place to save your files to.

Enter the location that you want to save your files to. You cannot select the volume that you are recovering from.

A folder with the name of the volume you are recovering from will be created and your folders and files will be saved to it. You should try to open a few files to check that they were properly recovered. Additionally, you should also check the report created; see the section "Viewing the Data Rescue Report" for more information.

> *Note: Always check that your data is fully recovered by opening your documents before you reformat (or initialize) your original disk. You may think that your data is recovered when it is not. Data Rescue never alters the original disk, so you can try other methods to recover your files if Data Rescue fails. After you reformat or initialize your disk, it is very likely that Data Rescue won't be able to recover data anymore. In general, always try non-destructive methods like Data Rescue before you attempt to fix your disk.*

6.6 Save Scan Option

As Data Rescue does a Quick or Thorough scan of a volume, it gathers information ("scan data") about the files and folders it finds. You can save this scan data in a file using the "Save Scan File..." choice in the File menu. Later, after quitting Data Rescue, if you decide you need to recover additional files from the volume, you can restart Data Rescue and choose the "Use scan data" scan method in the main window, and the subsequent dialog this brings up, to find and load the scan file you saved earlier. In this way, you can avoid the need for a possibly lengthy rescan.

Note: It is not a good idea to write to a damaged volume, so if you know or suspect the volume you scanned is damaged, we recommend that you save the scan file to a different volume.

Figure 31: Use Scan File

6.7 Setting Data Rescue Preferences

You can set Data Rescue preferences for how scans are performed. Most preferences are meaningful only at scan-time, which is why they are presented here. Once set, options are saved in the **Preferences** folder of your system folder. The default options are suitable for most situations. (Note: Preferences are **not** permanently saved when booted from the Data Rescue boot CD, and must be reentered each time the system is booted from the CD.)

From the **Data Rescue X** menu, click **Preferences**. The Preferences dialog opens, as shown in Figure 32.

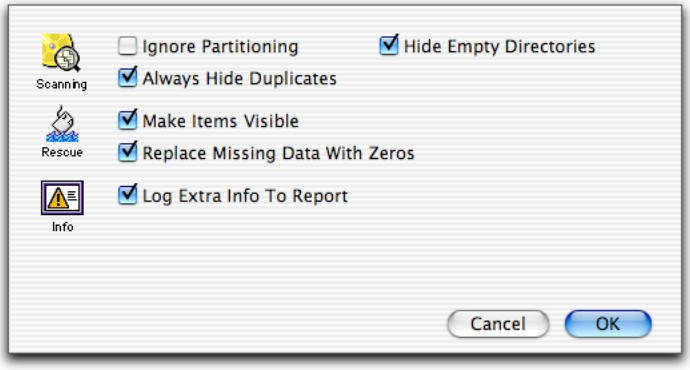

Figure 32: Preferences Dialog Box

Set desired Data Rescue preferences according to the following descriptions:

- **Ignore Partitioning**: If you check this box, partitions will be ignored. This means that your whole disk (including its partition map and drivers) will be seen as a volume. Don't use this option if you know that there is more than one volume on your disk. Note that if this option is set, some volumes may appear twice in the list (once as the volume, once as the whole disk). You should check this box only if you feel that the partition map is damaged. In any event, Data Rescue may ignore partitioning automatically for disks with questionable partition maps. You can see that partitioning is ignored for a particular disk because the name of the disk maker appears instead of that of the volume. The recommended setting for this option is "off" (not selected). See additional discussion of this item in section 4.2.2
- **Always Hide Duplicates**: If Data Rescue determines that two files are duplicates of each other, the duplicate file will be hidden. This option is recommended.
- **Hide Empty Directories**: Directories (folders) which are empty will be hidden if this box is checked. Note: In the current implementation, folders which contain any items, even if they are only hidden or empty directories, are not considered empty. This option is recommended.
- **Make Items Visible**: Data Rescue will change invisible files to visible, so that you can see them in the finder (Data Rescue will not change the names of files, so files that start with "." will remain invisible on Mac OS X.
- **Replace Missing Data With Zeros**: This option is meaningful only at recover-time. If a read error occurs while a file is recovered (possibly because of a bad medium), and if this option is set, then the missing data is replaced with zeros and more data is read if possible. If this option is not set, no more data is read for this file, and the recovery goes on to the next file.
- **Log Extra Info To Report**: Data Rescue will out put more than just the basic information about your recovery to the report window. The can help with recovering highly corrupted drives. (Using Data Rescue reports is covered in Chapter 3, "Using Advanced Features"). This option is recommended.

Click **OK** to save preferences settings.

6.8 Viewing the Data Rescue Report

Data Rescue reports comments and errors in the **Data Rescue Report** window, which you can view by selecting **Show Report** from the **Report** menu. This window is read-only and, as such, cannot be edited, but you can save its content to a text file by selecting **Save Report As…** in the **Report** menu. You can also log the report continuously to a file by selecting **Log Report To…** in the **Report** menu. These files can be edited or printed with a basic text editor like TextEdit.

The following are options you can use for data rescue reporting:

- To get useful information about your computer's configuration, select **Report Configuration** from the **Report** menu. This includes Data Rescue configuration, hardware configuration, system software configuration and a list of active system extensions and control panels.

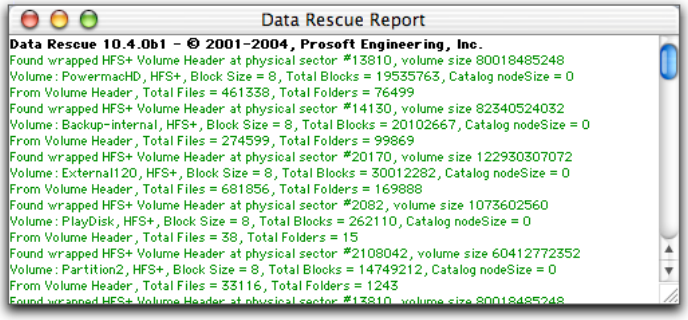

Figure 33: Data Rescue Report Window

- To list all files and folders in the current folder, select **List Files/Folders** from the **Report** menu.
- To get more information about the selected file(s) or folder(s), select **Report File/Folder Info** from the **Report** menu. This includes file/folder ID, type & creator, data and resource forks sizes, creation & modification dates, and allocation block numbers.
- To clear the report window, select **Clear Report** from the **Report** menu.

The report window contains 500 lines at most. When it reaches 500 lines, the first lines are removed as new lines are added. However, you can record longer reports by logging them to a file with the **Log Report To…** command.

6.9 Allocation Block Layout Window

This section explains the ABL window settings and use, and assumes you already understand the ABL concept (see section 5.5 for an explanation of the concepts). Note: The "size" and "offset" numbers explained in the concepts section correspond to the

"Block Size (a)" and "Base Block (b)" fields, respectively, of the ABL window (Figure 34).

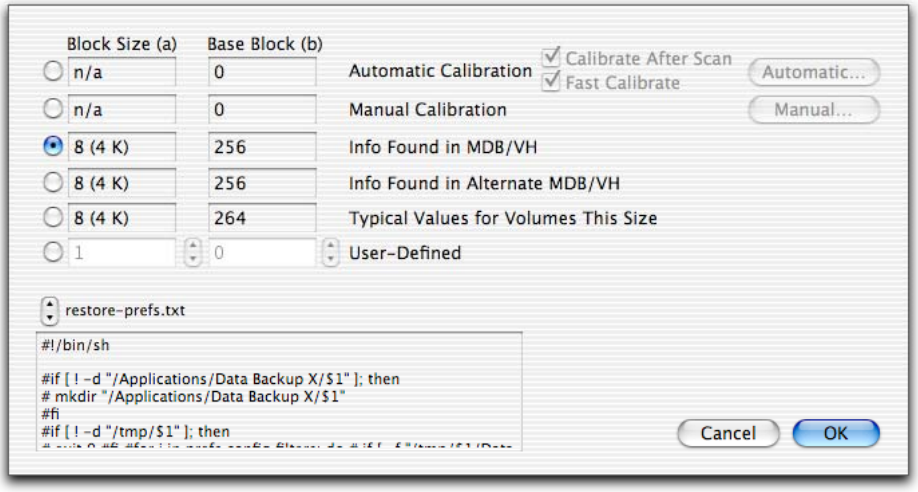

Figure 34: Allocation Block Layout (ABL) Window

By default, Data Rescue does an ABL calibration automatically after a Thorough Catalog scan. To detect the allocation blocks layout automatically, Data Rescue picks 200 files in the volume, and finds which combination of `size` and `offset` matches the most files respecting the above formula. It can take a while to find the 200 files, so Data Rescue includes a "Fast calibrate" option to pick only 10 files. If a fast calibrate yields questionable values, you should restart with a normal detection. Fortunately, the automatic detection of the allocation blocks layout rarely fails.

The following is additional information about the variables `a` and `b`:

- In HFS, the value of *a* is often around the volume size in KB divided by 32768.
- In HFS+, *a* can only be 8, 16, 32 or 64 and is typically 8.
- In HFS, the value of *b* is always greater than zero and is typically 19 or larger.
- In HFS+, *b* is theoretically zero, but most HFS+ volumes have a wrapper volume (a fake HFS volume that contains the real HFS+ volume), which offsets the file space. In practice, the value of *b* is the offset of the HFS+ volume within the wrapper volume.

6.9.1 Determining the Allocation Blocks Layout

The ABL is determined by various methods selected and configured by you in the Allocation Blocks Layout dialog (shown in Figure 34).

Below are descriptions of the methods that can be selected for use by Data Rescue:

6.9.1.1 Automatic Calibration

Auto-detect allocation blocks layout. If this option is set (default), Data Rescue will attempt to detect these parameters automatically. Otherwise, it will use the parameters

you provide. If the “Fast Calibrate” box is checked, Data Rescue will use 10 files to make its calculations; otherwise, it will use 200 files. Note that the **Manual** allocation blocks layout values are not stored in the preferences file.

The recommended setting is **Automatic Calibration** on and **Fast Calibrate** off.

For this method to work, the volume under recovery must contain files with resource forks. Prior to Automatic Calibration, the volume must have been scanned for files and files with resource forks must have been found.

Data Rescue looks for the (a, b) pair that matches the greatest number of equations and displays that information within the “Automatic Calibration” section of the ABL window. This ABL is later used for file recovery if the “Automatic Calibration” method is selected with the corresponding radio button in the ABL window.

The drawback of this method is that there must be resource files with unique sizes on the volume, so this may not work on some volumes (for instance, a volume containing only a library of resource-less JPEGs cannot be calibrated).

If the **Calibrate After Scan** box of the ABL window is checked, calibration is automatically attempted just after the disk is scanned for catalog information. This option, and the Automatic Calibration method, are the defaults.

6.9.1.2 Manual Calibration

Manual Calibration is the same as Automatic Calibration except that the volume is not required to contain files with resource forks. Instead, it uses backup files provided by the user.

The phases of Manual Calibration are:

1 - The user provides backup files through an external volume. Each of these files must:

- Be a copy of a file that exists on the volume under recovery and whose catalog information was found by Data Rescue,
- Have a unique name (it is useless if there are several files with the same name on the volume under recovery), and
- Have at least 512 bytes of data and/or resources.

2 - Data Rescue scans the volume under recovery for the original files. It uses the first block of data and/or resource to identify the original files.

If Data Rescue can find enough unique matches between the external volume files and files on the scanned disk, it can determine the correct ABL.

As in Automatic Calibration, Data Rescue looks for the (*a*,*b*) pair that matches the greatest number of equations and displays that information within the "Manual Calibration" section of the ABL window. This ABL is later used for file recovery if the "Manual Calibration" method is selected with the corresponding radio button in the ABL window.

The drawback of this method is that the user must have a backup copy of some files on a second volume.

6.9.1.3 Info Found in MDB/VH

This is the most natural way of determining the ABL. It yields the values that are stored by the file system in the volume's headers. If both the MDB/VH and the alternate MDB/VH contain the same nonzero values, then these values are highly likely to be the correct ones.

For HFS volumes, *a* and *b* are stored in the MDB (Master Directory Block), at logical block #2 of the volume.

For HFS+ volumes, *a* is stored in the Volume Header, at logical block #2 of the (inner) HFS+ volume. The location of the inner volume is stored in the MDB of the wrapper volume. *b* is the offset of the inner volume.

For more information on the MDB and the Volume Header, see "Inside Macintosh: Files" and TN 1150: "HFS Plus Volume Format".

If a valid MDB or Volume Header can be found in the volume, Data Rescue displays the corresponding information within the "Info found in MDB/VH" section of the ABL window. This ABL is later used for file recovery if the "Info found in MDB/VH" method is selected with the corresponding radio button in the ABL window.

The drawbacks of this method are as follows:

- If the boundaries of the volume are unknown (e.g. damaged partition map), the MDB cannot be located and the ABL cannot be found

- If the volume was reinitialized or altered by a disk utility, the ABL may be invalid even if the MDB looks valid.

6.9.1.4 Info Found in Alternate MDB/VH

- For HFS volumes, there is a backup copy of the MDB in the previous to last block of the volume.
- For HFS+ volumes, there is a backup copy of the Volume Header in the previous to last block of the inner volume.

If a valid alternate MDB or Volume Header can be found in the volume, Data Rescue displays the corresponding information within the "Info found in alternate MDB/VH" section of the ABL window. This ABL is later used for file recovery if the "Info found in alternate MDB/VH" method is selected with the corresponding radio button in the ABL window.

The drawbacks of this method are the same as that of (3). Unfortunately, if the MDB (or Volume Header) is unavailable or corrupted, then most probably the backup copy also is.

6.9.1.5 Typical Values for Volumes this Size

The ABL is chosen and applied at initialization time, essentially according to the volume size. This method simulates the default behavior of Apple's initialization tools.

- For HFS volumes, it uses the formula given in "Inside Macintosh" for calculating the default allocation block size. It then calculates the room needed for the volume bitmap to get b.
- For HFS+ volumes, it calculates the ABL for the volume wrapper, then calculates the room needed for the extents overflow file, catalog, Desktop DB, Read-Me and localized Read-Me files of the volume wrapper to get the offset of the inner volume, which is b. a depends on the size of the volume according to the following Apple-documented table:
    - up to 256 MB: $a = 1$
    - up to 512 MB: $a = 2$
    - up to 1 GB: $a = 4$
    - above: $a = 8$

If the size of the volume is known, Data Rescue displays the corresponding ABL within the "Typical Values for Volumes This Size" section of the ABL window. This ABL is later used for file recovery if the "Typical Values" method is selected with the corresponding radio button in the ABL window.

The drawbacks of this method are:

- If partitioning is ignored (e.g. bad partition map), the size of the volume is unknown and the ABL cannot be calculated.

- A tool that doesn't use these formulas (even Apple's tools seem not to always use the same formulas) may have initialized the volume. In this case, the displayed ABL will be wrong.

6.9.1.6 User-Defined

The user has the possibility to enter their own values in the ABL window. Here are a few techniques for determining the ABL by hand:

- The user may look in the MDB of a similar volume (a volume that has the same size and was initialized with the same tools using the same parameters).
- For HFS+ volumes: check the **Log Extra Info** box in the **Settings** window. Scan the volume. Whenever Data Rescue finds a block with an "H+" signature, indicating a possible Volume Header, the logical number of this block is logged to the report. "*b*" can be calculated by subtracting 2 from this value. In HFS+, only values of 1, 2, 4, 8, 16, 32 or 64 are allowed for "*a*" (the value is typically 8), so this method usually yields only a few values to try. It fails if the original Volume Header was overwritten.
- For all volumes: expert users can perform a calibration similar to that of Data Rescue using a disk utility such as SEdit. For instance, the user spots a well-known text file in Data Rescue's list. He or she gets *<allocationblock>* for this file with the **Report File/Folder Info** command. He or she then searches for the first logical block of this file with SEdit's text search feature; this yields *<logicalblock>*. The user must repeat this for another file. He or she then has a system of 2 equations with 2 unknown values. Solving this system yields the ABL.

  <logicalblock1> = (*<allocationblock1>* x *a*) + *b*

  <logicalblock2> = (*<allocationblock2>* x *a*) + *b*

  yield:

  a = (*<logicalblock1>* - *<logicalblock2>*) / (*<allocationblock1>* - *<allocationblock2>*)

  b = *<logicalblock1>* - (*<allocationblock1>* x *a*)

  Note: The user must access the volume the same way in Data Rescue and SEdit: if partitioning is ignored in Data Rescue, the volume must be accessed as a physical device in SEdit. Otherwise, there may be an offset issue with "b".

6.9.2 Using the Allocation Blocks Layout Dialog

Whatever method is selected in the ABL window, if this yields valid values for *a* and *b*, sample clips from found text files are shown in the text window at the bottom of the dialog as they would be recovered if this ABL was used.

You should initially try Automatic Calibration (selected by default). If Automatic Calibration fails, select other calibration choices and check whether the sample text files

appear as they should. If the sample text files appear as they should, the files can then be recovered using that calibration choice.

If the allocation blocks layout cannot be detected, Data Rescue shows an alert box. This may happen when not enough files have been found for calibration. If so, or if the values found by Data Rescue are incorrect, you will have to set *a* and *b* manually in the options window, where *a* is called **Block Size**, and *b* is called **Base Block**.

To set ***a*** and ***b*** manually

- Open the options window by selecting **Allocation Blocks Layout…** from the **Expert** menu.
- Click the **Use-Defined:** button, and enter *a* and *b* in the lower and upper boxes, respectively. Once you have set values manually, check the contents of a few files, as explained in the previous chapter, and change values again if the contents do not match.

You should start from the values given in the Data Rescue Report Window if you don't know what to set as *a* and *b*. Also try other values around these. Sometimes, *a* is exactly twice what Data Rescue guessed.

In HFS volumes, *b* is often about 19 (*a* depends on the size of the volume).

In HFS+ volumes, *a* is a power of 2 (often 8). Some tools let you choose *a* on initialization. A block size of 512 bytes (1/2 KB) means *a*=1, a block size of 1 KB means *a*=2, etc.

> *Note: Please note that if Data Rescue did not find the file(s) you want to recover, there is no point in determining the allocation blocks layout, as your files won't be recovered anyway. However, before concluding that Data Rescue did not find your file, be sure and check the contents of any files named "inodeNNNNNN" etc. in the Orphans Folder subfolders, in case it's one of those.*

6.10 Block Level Tools

Data Rescue features block-level tools to display and save raw disk blocks.

To display a block, select **New Window With Block#…**, and then enter the number of the logical block you want to display. The block is read in the currently selected volume.

You can also read the first block of the data fork (resp. resource fork) of a file by selecting this file and selecting **New Window With This File's DATA** (resp. RSRC). In that case, the logical block read depends on the allocation blocks layout.

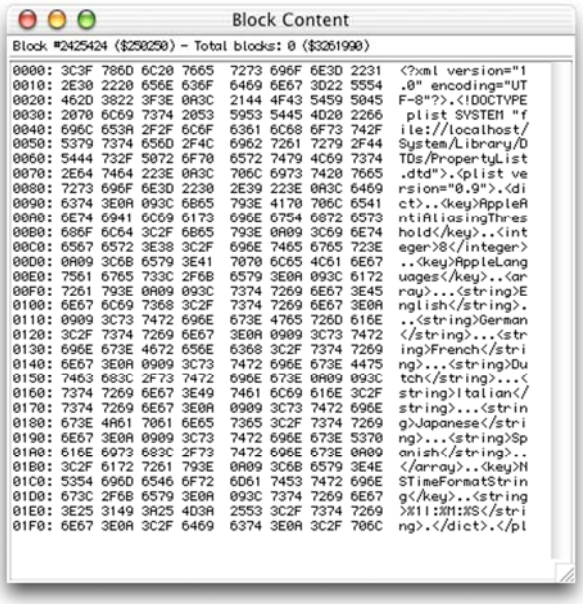

Figure 35: Block Content Window

The first number in a row is the offset in the block (in hex). Then follow the block's data in hex and ASCII. You can view the next (resp. previous) block by selecting **Next Block** (resp. **Previous Block**) in the **Blocks** menu. 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.

If you wish to view your disk's partition map, you must ignore partitioning, since the partition map does not belong to a volume. Also note that you can't edit a block in the block display.

You can save raw blocks to another medium by selecting **Save Blocks As…** in the **Blocks** menu.

Select a location to save your blocks and enter the number of the first logical block to save and the number of logical blocks to save. You can also set a particular type or creator for the file. Click **Save** to save the blocks.

7 Troubleshooting

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

7.1 Data Rescue freezes/crashes. What should I do?

- Check SCSI connections, IDs and terminations.
- Disconnect unnecessary devices.
- Bad SCSI cables can cause the Macintosh to freeze. Try with better cables.

7.2 Why are some files listed twice?

- Sometimes, your volume contains obsolete Catalog entries. If Data Rescue finds more than one Catalog entry for the same file in your volume, it displays them all. You should check which is the correct entry and hide all others; otherwise Data Rescue will recover the first entry in the list and generate “file already exists” errors for the others. To find which is the correct entry, check dates and file contents.
- Special protection or optimization tools such as TechTool Pro can leave partial or full Catalog copies on your volume (incidentally or on purpose). If so many duplicates will appear in Data Rescue’s list. You can get rid of these duplicates by selecting **Hide Duplicates** from the **Expert** menu.

7.3 Does Data Rescue handle RAID volumes?

- Data Rescue X currently does not handle RAID volumes. Robust RAID support will be added in future releases.

8 Glossary

Alias: An *alias* is a reference to a file. When organizing your disk, you may have a very important file or application that you want to get at on a very regular basis, and thus want to have it easily accessible – say on the desktop. At the same time, that file may "belong" inside, say, your Documents folder for a particular project. You don't want two copies of the file on your disk: for one, they take up more space. And secondly, the two might get out of sync. The way around this problem is the *alias* – the *alias* is a really small file that just points to the original. So, it takes up almost no space and yet can be put in a different folder organizationally for easy access. Opening the *alias* will, in reality, open the original file, wherever it is. You can also have *aliases* of folders and even whole volumes, whether they're local (attached to your computer) or across the network.

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.

Compressed file: A *compressed file* (e.g. a StuffIt file or a self-extracting archive) is a file that has been literally compressed - encoded in a way to save space. Files are usually compressed either for archival purposes or to transfer files over the internet, as smaller files will transfer more quickly than bigger ones. Encoded files (e.g. a MacBinary, uu-encoded, or BinHex file) are files that have been encoded to travel over the internet more easily (there are certain rules about what type of data will travel best). Encoded files usually take up a little more space than their unencoded version. Often files are first compressed and then encoded for travel over the Internet.

Fat application: Not long after the Macintosh was introduced, Macs were made using the Motorola 68000 line of CPUs (68000, 68020, 68030, and 68040). Since then, Macs have been made using the PowerPC line of processors (601, 603, 604, G3, G4), also from Motorola. Applications compiled to run on the 68K (as they're called) line of processors do not run by default on the PowerPC processors, and vice-versa. To avoid compatibility problems, Apple created an emulator for the PowerPC processors that allowed them to emulate a 68K processor and thus support applications developed for the 68K line. Alas, emulation means more work, so applications that run in "emulation mode" are slower than "PowerPC native" applications. To give old users (people with 68K Macs) and new users the best of both worlds, application developers have and still do provide "*Fat applications*" – fat because they contain both sets of code. This is good for speed but bad for hard disk space: a *fat application* is often over 50% larger than one that has only one version of the code.

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 releases with Apple's 8.1 operating system and is designed to supercede HFS. All Mac OS versions from 8.1 Classic through OS X 10.x support both HFS and HFS+.

Logical block: this is a 512-byte block of information on a volume. Logical blocks start at zero within a volume. Since a physical device usually contains additional information before the volume (partition map, drivers, other partitions…), the logical block number is usually a number smaller than the *physical block* number for the same block. In Data Rescue, the whole physical device is considered a volume if partitioning information is ignored (e.g. bad partition table) - in that case logical blocks are the same as physical blocks.

Orphaned alias: An *orphaned alias* is an *alias* (see above) that has lost track of the file it originally pointed to. The most common reason this happens is because the file it pointed to was deleted, but sometimes the file just gets lost. Apple's *aliases* do a pretty good job of staying linked, even when you move the original file around. Occasionally, though, Apple's system will fail and you will have to "reattach" the *alias* by telling the Mac OS where the original file now resides.

Orphaned file: An *orphaned file* is a file with a *creator type* (see above) that does not match any of the applications currently available to your machine. A file can be an *alias* if it was created by an application you no longer have on your machine or if it was given to you from elsewhere, and was created with an application you do not have. *Orphaned files* can still be used because many types of documents can be read by multiple applications, so as long as you have an application that can read it, you can look at it. Double-clicking on an *orphaned file* will cause the Mac OS to ask you for suggestions about which application to use to view the document since the Mac doesn't know the document's preference.

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.

Remap: *Remap*ping is the process by which you assign a file from one application to another. Setting the file's creator type (see above) from the old application to the new one does this. The net result of this is that the icon displayed for the file will now by associated with the new application, and the new application that will be launched when you double-click on the file.

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*.