

User Guide for DDRescue- GUI v1.7.1 and higher

generated by DAPS

User Guide for DDRescue-GUI v1.7.1 and higher: generated by DAPS

Version 1.0

Revision: 5

Copyright (c) 2017-2018 Hamish McIntyre-Bhatty.

Permission is granted to copy, distribute and/or modify this document under the terms of the GNU Free Documentation License, Version 1.3 or any later version published by the Free Software Foundation; with no Invariant Sections, no Front-Cover Texts, and no Back-Cover Texts. A copy of the license is included in the section entitled "GNU Free Documentation License".

Link To GNU Free Documentation License [<http://www.gnu.org/licenses/fdl-1.3-standalone.html>]

Source code is available at <https://github.com/hamishmb/docs>

Table of Contents

Preface	vi
1. What is DDRescue-GUI?	1
What is GNU ddrescue?	1
What is DDRescue-GUI?	2
Why should I use DDRescue-GUI?	3
System Requirements	4
Where can I download DDRescue-GUI?	5
2. Data Recovery Concepts	6
Device Names	6
Partitions and File Systems	6
"Mounting" and "Unmounting" Devices	7
Bad Sectors	7
Data Corruption	7
"Disk"	7
Formatting	7
Deleting and Undeleting	8
Image Files	8
The Disk Info Window	9
3. Getting Started With DDRescue-GUI	10
The Main Window	10
Setting up for your recovery.	14
Start your recovery	16
Finishing your recovery	17
4. Advanced Functionality	20
Why should I use a mapfile?	20
The advanced options in the settings window	20
Direct Disk Access	20
Reading backwards	21
Preallocating disk space	21
Soft run	21
Overwrite output file	22
Number of times to retry bad sectors	22
Maximum number of errors before exiting	22
Number of clusters to copy at a time	22
Device-Device recoveries	22
Recovering individual files with DDRescue-GUI	25
Mounting Output Files	26
When recovering from an entire device	26
When recovering from a partition	28
5. Frequently Asked Questions	30
Setup-related Questions	30
Runtime Questions	31

User Guide for DDRes-
cue-GUI v1.7.1 and higher

Post-recovery Questions 32
Troubleshooting Questions 33
A. License For The Stylesheets used to generate this documentation. 35
B. License For This Documentation 37

List of Figures

1.1. GNU ddrescue in action	1
1.2. DDRescue-GUI in action	2
1.3. Terminal output in DDRescue-GUI	4
2.1. DDRescue-GUI's Disk Information Window (Linux)	9
2.2. DDRescue-GUI's Disk Information Window (Mac)	9
3.1. DDRescue-GUI's Main Window (Linux)	10
3.2. DDRescue-GUI's Main Window (Mac)	10
3.3. Detailed info in DDRescue-GUI	12
3.4. Detailed info in action in DDRescue-GUI	12
3.5. Disk info in DDRescue-GUI (Linux)	13
3.6. Disk info in DDRescue-GUI (Mac)	13
3.7. DDRescue-GUI's Settings (Linux)	14
3.8. DDRescue-GUI's Settings (Mac)	15
3.9. DDRescue-GUI's Settings Pre-sets	15
3.10. DDRescue-GUI in action	16
3.11. Recovery Success (Linux)	17
3.12. Recovery Success (Mac)	17
3.13. Recovery Failure/Aborted (Linux)	18
3.14. Recovery Failure/Aborted (Mac)	18
3.15. Reverse read, and preset.	19
4.1. Manually selecting an output device (Linux-only)	23
4.2. Navigating to /dev on Linux	24
4.3. Confirming your destination (Linux)	25
4.4. Confirming your destination (Mac)	25
4.5. The mount button (Linux)	26
4.6. The mount button (Mac)	27
4.7. Selecting a partition (Linux)	27
4.8. Selecting a partition (Mac)	27
4.9. Successfully mounting a partition in your device (Linux)	28
4.10. Successfully mounting a partition in your device (Mac)	28
4.11. Successfully mounting your partition (Linux)	29
4.12. Successfully mounting your partition (Mac)	29

Preface

Abstract

DDRescue-GUI is a Graphical User Interface for GNU ddrescue, a data recovery tool. GNU ddrescue is a command-line (text-only) program, so using this interface with it brings you all of the power ddrescue has to offer, combined with the ease of use of a simple, cross-platform GUI. This document will teach you how to use both GNU ddrescue and DDRescue-GUI, and serves as a general-purpose troubleshooting reference.

Why read this document?

Well, for a lot of reasons. Firstly, data recovery is a complex and delicate subject, and even when using tools as powerful as GNU ddrescue, things can go wrong, especially if you're not experienced. That might sound quite negative, but in truth it's exactly why I wrote DDRescue-GUI; I've made a fair few mistakes when doing data recovery. Some of them have simply wasted a bit of time, but some of them have meant I've lost data. I'm simply fortunate that the data I lost wasn't important, but if a GUI like this had existed when I first used ddrescue, I could have avoided making some of those mistakes. I wrote this guide so you can avoid making some of the mistakes I've made.

The other reason is that while (and, perhaps, because) GNU ddrescue is extraordinarily powerful, it can be quite confusing to use. It has a lot of options, some of which are generally useful, and some of which apply more to corner-cases, so I hope this document will be helpful when you're deciding which options to use. When writing the GUI, I've tried to help you make the best decisions by using reasonable defaults, but here I will explain the options in detail, which should help with some of the confusion some of my users have experienced.

What is covered in this document?

Hopefully, it covers anything you might want to know about DDRescue-GUI, ddrescue, and data recovery in general. I know that's a large undertaking for a small guide, so at various points I may direct you to read other materials for more detail in particular areas.

To be specific, this document covers:

Basic Information About
DDRescue-GUI

Why you should use it, what systems it works on, and what GNU ddrescue is.

Data Recovery Concepts

Explains basic data recovery concepts in plain language, and other important things like device names.

DDRescue-GUI's Basic Functionality

Just enough to get you on your feet with using DDRescue-GUI. Covers general use of the GUI and some of the more basic features.

Advanced Functionality

Advanced functionality of ddrescue and DDRescue-GUI, including use of map files, mounting output files, and explanations of advanced options and when/why to use them.

Frequently Asked Questions and Troubleshooting

Some of the questions I get asked a lot, and also some extras to help you if you encounter problems.

Conventions

I have used a few conventions in this book, but I've tried to keep it simple where possible. DDRescue-GUI runs on Linux and macOS, and I will provide screenshots for both platforms where necessary. At any point where you need to be careful about something, this icon is used with some information:



Caution

Some general advice about how to do something, or what you should check.

Alternately, if there's some important information you need to read, I'll use this:



Important

Some important information which may apply in a few scenarios.

If there's a little snippet of useful information that might apply to you, I'll put it in a note:



Note

These are used to point out little pieces of information that may serve as shortcuts or as platform-specific features.

If there's a general-usage tip for you, it might look like this:



Tip

These might suggest something extra or optional you can do in particular circumstances.

If you need to follow instructions carefully, I might use this:



Warning

These are used to emphasise important part of instructions, or instructions that may have severe consequences if followed improperly.

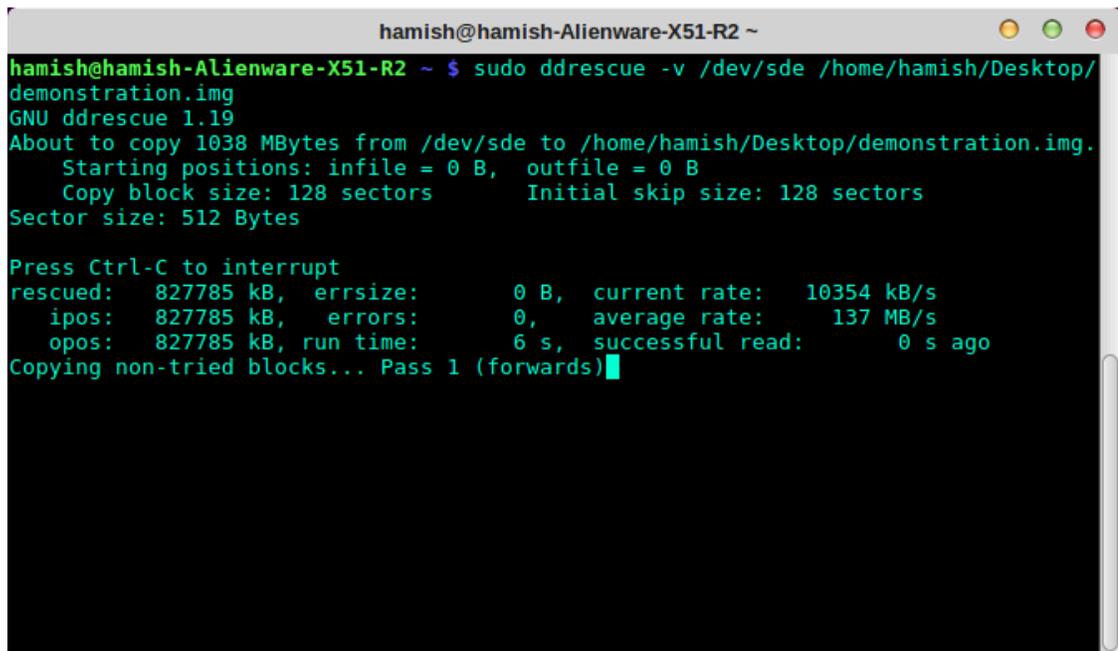
Chapter 1. What is DDRescue-GUI?

Abstract

DDRescue-GUI is an open source GUI front-end for GNU ddrescue. It is designed to make ddrescue easier to use, as it is a command-line application. DDRescue-GUI brings ddrescue to a wider audience than was previously possible because "normal" users can use it. This chapter provides basic information about ddrescue and DDRescue-GUI. I'm going to start at the start, to make sure everything is very clear. If you already know about ddrescue and DDRescue-GUI, you can probably skip to the next chapter.

What is GNU ddrescue?

Figure 1.1. GNU ddrescue in action



```
hamish@hamish-Alienware-X51-R2 ~ $ sudo ddrescue -v /dev/sde /home/hamish/Desktop/demonstration.img
GNU ddrescue 1.19
About to copy 1038 MBytes from /dev/sde to /home/hamish/Desktop/demonstration.img.
  Starting positions: infile = 0 B,  outfile = 0 B
  Copy block size: 128 sectors      Initial skip size: 128 sectors
Sector size: 512 Bytes

Press Ctrl-C to interrupt
rescued: 827785 kB, errsize: 0 B, current rate: 10354 kB/s
 ipos: 827785 kB, errors: 0, average rate: 137 MB/s
 opos: 827785 kB, run time: 6 s, successful read: 0 s ago
Copying non-tried blocks... Pass 1 (forwards)
```

GNU ddrescue is a data recovery program, written by Antonio Diaz Diaz. It's amazingly powerful, and has a lot of features other tools don't have, such as the ability to use a recovery mapfile (previously called a log file), which means that if your system crashes during a recovery, you can start again where you left off.

Other features include multiple passes over drives to try and get the easy-to-read data first, direct disk access (complicated so I will explain it later), and disk-to-disk recovery. It was originally written for Linux, ddrescue can actually run just fine on macOS and even Windows, which makes it all the more impressive. Note that DDRescue-GUI doesn't currently run on Windows, but may do so in a later release.

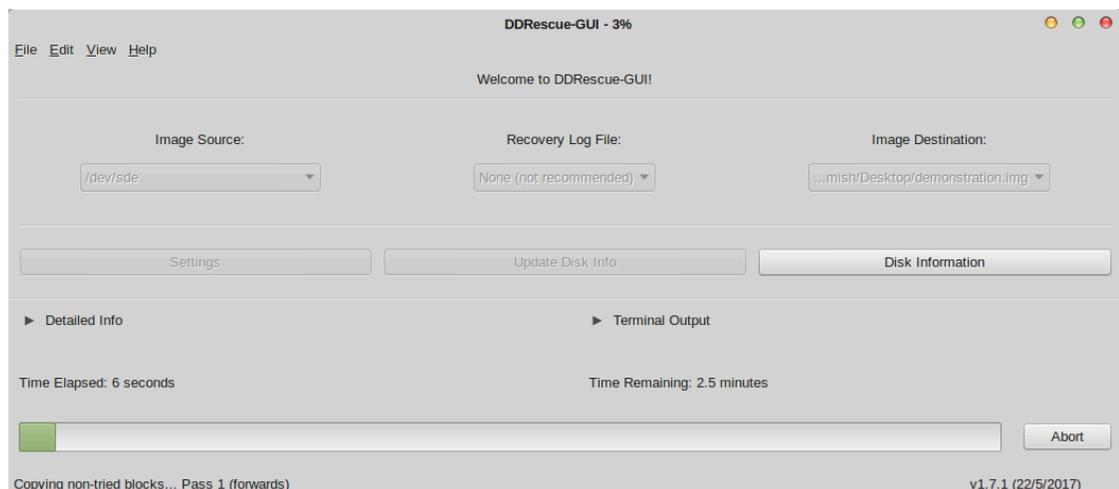
That's where the difficulty comes in, though. Firstly, ddrescue is a command-line tool, which means only people who are used to using the text-only command-line are going to be able to use it. It's possible to figure it out anyway, but if you're not confident that you're doing it right you could cause loss of data. Despite this, ddrescue is probably still the best data recovery tool out there.

The second problem is that while ddrescue will run on macOS, you have to compile it, which is often slow, difficult, and confusing, especially if you aren't used to doing this sort of thing. DDRescue-GUI provides a graphical interface for ddrescue, and solves some of these other problems as well.

Nevertheless, remember that GNU ddrescue is an excellent tool. For more information about ddrescue, see <https://www.gnu.org/software/ddrescue/>.

What is DDRescue-GUI?

Figure 1.2. DDRescue-GUI in action



DDRescue-GUI aims to bring GNU ddrescue to a wider audience by providing a graphical interface. DDRescue-GUI is cross platform, running on macOS and Linux. On macOS, DDRescue-GUI comes with a pre-compiled copy of ddrescue, so you can just install it and get to work. The GUI also provides other features not present in ddrescue, such as disk information, to help make sure you select the right disks for your recovery, auto-unmounting of devices (I will explain this later), and useful tips and suggestions to help you get set up.

DDRescue-GUI lets you use most of ddrescue's features, as well as guiding you. On macOS, some of these features aren't available if you're just using ddrescue by itself, but in a few cases DDRescue-GUI can work around the problem and provide a way for you to use them.

Why should I use DDRescue-GUI?

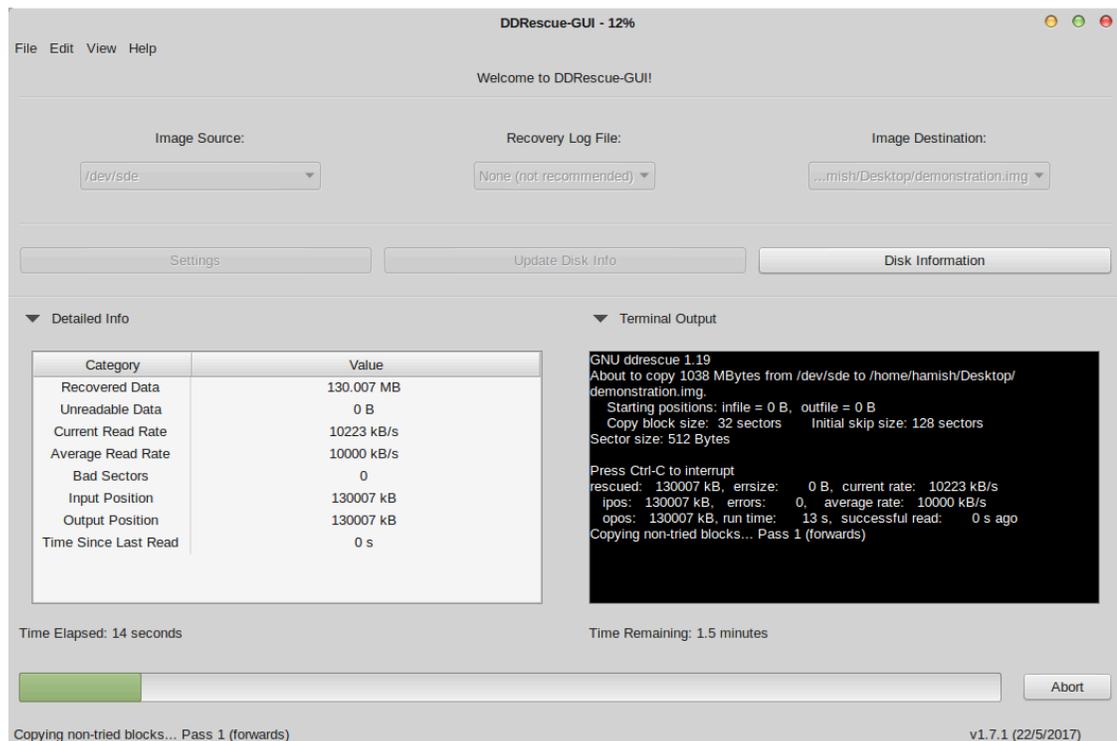
If the features above didn't convince you, then this might :) If you're a normal user and wondering why you should use this program with ddrescue instead of one of the other popular tools, you should know that ddrescue is smarter than quite a lot of the other tools:

Let's imagine that you have an old hard drive with precious data on it, say, family photos. GNU ddrescue will try to read the easily-readable parts of the disk first, which might grab, say, 70% of your photos quickly. This is important because drives often fail quite soon after showing warning signs, and it might let you get a fair amount of data before your drive completely gives up. After this, ddrescue will make several more passes to try and get the rest of the data.

Some other tools, such as SpinRite, tend to read the drive in sequence, and refuse to give up when they find an unreadable part of the drive. Combined with other techniques, this can theoretically get the data off your drive and even fix your drive. However, it could also break your drive before you get much data, especially if there are a load of bad areas at the start of the drive, but most of the rest of it is readable. Note that SpinRite is a well-respected tool, and has helped a lot of people, so I mean no disrespect to its users or authors. I would worry about using a disk again after it failed though, even if it had been fixed.

That aside, if you're an administrator, you might enjoy looking at ddrescue's output so you can see exactly what going on, rather than use a styled GUI. I certainly do. In this case, you're in luck: DDRescue-GUI makes it easy for you to see the output as well, if you want to:

Figure 1.3. Terminal output in DDRescue-GUI



System Requirements

DDRescue-GUI is designed to be lightweight, so it will run on most systems.

Linux Users: You need to be running a modern Linux distribution. Anything made after about 2014 should be fine. More specifically:

- Fedora 26 and Higher, and all derivatives/spinoffs.
- Ubuntu 14.04 and Higher, and all derivatives/spinoffs.
- Any recent version of Parted Magic.
- Theoretically, any other similar Linux distributions, like Slackware, Debian and Red Hat Linux.

Mac Users: Shouldn't be a problem for any even slightly modern mac. Specifically:

- macOS 10.9 (Mavericks) and Higher.
- 64-bit Intel macs only.

If you really need to run DDRescue-GUI on an older mac, run v1.6.1. This is an old version, and no longer updated, but it will work on:

- macOS 10.6 (Snow Leopard) and Higher.
- 32-bit and 64-bit Intel macs (choose at download time).

Note that there is absolutely no support whatsoever for macOS versions earlier than 10.6, and PowerPC macs.

Where can I download DDRescue-GUI?

Both Mac and Linux users can download from my website at www.hamishmb.com [https://www.hamishmb.com/html/downloads.php?program_name=ddrescue-gui]. This is the primary site for downloads, but you can also download the GUI from launchpad.net, from www.launchpad.net/ddrescue-gui [<https://www.launchpad.net/ddrescue-gui>].

Chapter 2. Data Recovery Concepts

Abstract

In order to be able to use any data recovery tools, you need to understand some of the concepts behind data recovery. This chapter will explain the basics, and it will also show you how to get help when you're using the GUI. You can safely skip this chapter if you already understand the concepts, because I've gone into fairly basic detail here.

Device Names

Device names are a bit of a weird concept to a lot of people, particularly those who use Windows, where drive letters are used instead, like C:, D: and so on. On Linux and macOS, however, devices (like hard drives, USB drives, DVD drives) have names. These names are used to access the drives, so for example, to read directly from the DVD drive, you might read from a device called `/dev/sr0` (Linux), or `/dev/disk1` (macOS), which you can think of as being similar to the D: drive in Windows.

Linux and macOS tend to use different names for devices. On Linux, it's quite common to have names like `/dev/sda` (first hard drive), `/dev/sdb` (second hard drive), and `/dev/sr0` (first optical drive). You might also see names like `/dev/sda1`, which represents the first partition on the first hard drive. macOS, however, tends to call disks `/dev/disk0` (first disk in the system, usually the boot drive), `/dev/disk1` (could be a hard drive, USB drive, or optical drive), and `/dev/disk1s2` (second partition on the second drive). It's a bit easier to understand the Linux names. If you're finding this confusing, don't worry, because I will explain it more clearly later along with some of the other concepts. It's quite complicated.

Now for the really confusing bit. While you CAN read directly from these devices, you won't be able to understand what you're reading. This is where it becomes different from Windows. On Linux and macOS, pretty much everything is accessed as a file, including devices. So, essentially, when you read from, say, `/dev/disk0`, you're actually reading data directly from the hard drive as if it were a file. In order to be able to use the drive in any meaningful way, it needs to have partitions and filesystems.

Partitions and File Systems

Okay. Device names can be a bit confusing, but it'll start to make more sense soon. A partition is, literally, just a section of a hard drive or other device. Usually, you only need partitions on hard drives, and other devices like DVDs often don't have them at all. You can have many partitions on one drive, often up to 4, but sometimes many more than that.

Right, so we have devices and partitions, but to actually store data, the partitions need a file system. This organises the files and keeps track of where they are and how large they are. A file system is what you see in Windows as, say, C:, or E:. Remember that some devices like DVDs don't need partitions, and the device has the filesystem itself.

"Mounting" and "Unmounting" Devices

You might have heard these terms before if you use Linux or macOS. If you have a partition with a file system, you need to "mount" it to make it usable. You could think of it like opening a book, I guess: All the information is there in the book, but you need to open it before you can read it. In Windows, mounted devices get drive letters like C:. In Linux and macOS, mounted devices are given a folder where they can be accessed, like `"/media/data"`, for example. When you're finished with a drive, you should unmount it, which is usually done when you shut down your computer.

Bad Sectors

A sector is a very small area of a disk. The exact size can vary, but nevertheless, it's very small. When a sector becomes unreadable, it is called a bad sector. Bad sectors can happen because of corrosion, scratches, shock damage and so on.

Data Corruption

This often happens because of bad sectors, but can also be due to power loss while writing data and many other things. Basically, it's data that can be read, but it isn't completely intact so it might not be meaningful or usable. If the sector is damaged the data may have been corrupted or misread. If you're unlucky, you may find this happened to you. In that case, you'll likely need to use some more tools to fix your output file from DDRescue-GUI before you can use it.

"Disk"

This is quite a vague term that can mean a lot of things. In this user guide, and in the program, a disk is a partition or device that contains a file system.

Formatting

If you format a disk, you are giving it a file system. This can also mean erasing the previous filesystem and everything it contains. If you say a disk is formatted as FAT32, it means the disk has a FAT32 filesystem on it. Often when a disk

is formatted, the data on it isn't actually overwritten, and it can sometimes be recovered. The same goes for deletion.

Deleting and Undeleting

You almost certainly know what deleting is. What you might not know, though, is that when you delete a file it isn't actually "gone" most of the time. Often, the reference that tells your computer the file exists is deleted, and the space the file was using is marked as free, but the actual file is still there, until it gets overwritten by something else later.

Because the files aren't often overwritten until later, if a file was deleted recently, it can often be "undeleted". This is where a special program reads through all the empty space on your disk and looks for file fragments that haven't been overwritten yet. You can often recover a surprising amount of data by doing this.

NOTE: GNU ddrescue and DdRescue-GUI aren't undelete tools, but you can often use them alongside undelete tools to get more data back.

Image Files

These hold a filesystem, much like a partition. This is a weird concept, but if you think about it, if everything in macOS and Linux is represented as a file, including partitions, it makes sense that a file can hold all the data from a partition or device. Image files are usually created by literally copying the file that represents a partition or a disk. Image files can have partitions inside too, just like a hard drive.

You might be wondering why these are helpful. Often they're used for backups, but also for data recovery. You can recover data from a bad DVD, say, and save the data into a file. This is good, because often after doing a recovery, some of the data you got turns out to be corrupted, and it needs to be analysed with another tool. It's much easier to do this if you recover to a file.

The Disk Info Window

Figure 2.1. DDRescue-GUI's Disk Information Window (Linux)

Name	Type	Vendor	Product	Size	Description
/dev/cdrom	Device	HL-DT-ST	DVD+-RW GA50N	N/A	DVD-RAM writer
/dev/sda	Device	MSDOS...	Samsung SSD 850	500 GB	ATA Disk
/dev/sda1	Partition	MSDOS...	Host Device: Samsung SSD 850	520 MB	Windows FAT volume
/dev/sda2	Partition	MSDOS...	Host Device: Samsung SSD 850	16 MB	Windows FAT volume
/dev/sda3	Partition	Windows	Host Device: Samsung SSD 850	134 MB	reserved partition
/dev/sda4	Partition	Windows	Host Device: Samsung SSD 850	9 GB	Windows NTFS volume
/dev/sda5	Partition	Windows	Host Device: Samsung SSD 850	87 GB	Windows NTFS volume
/dev/sda6	Partition	Windows	Host Device: Samsung SSD 850	447 MB	Windows NTFS volume
/dev/sda7	Partition	Windows	Host Device: Samsung SSD 850	447 MB	Windows NTFS volume

Figure 2.2. DDRescue-GUI's Disk Information Window (Mac)

Name	Type	Vendor	Product	Size	Description
/dev/disk0	Device	VBOX	HARDDISK	42 GB	Internal Hard Disk Driv.
/dev/disk0s1	Partition	VBOX	HARDDISK	209 MB	Internal Hard Disk Driv.
/dev/disk0s2	Partition	VBOX	HARDDISK	42 GB	Internal Hard Disk Driv.
/dev/disk0s3	Partition	VBOX	HARDDISK	650 MB	Internal Hard Disk Driv.

A lot of these terms are confusing to people who haven't heard them before. To make it easier to use the GUI, I have included a disk information window. This window shows you device names, but also other information about disks such as size, label, the model, and the description. All of that will help you figure out what device names you want to use, and help prevent data loss.

Chapter 3. Getting Started With DDRescue-GUI

Abstract

This chapter covers the basic functionality you'll need to use in DDRescue-GUI. It's intended to be just enough for you to get started using DDRescue-GUI.

The Main Window

That was quite a lot of introduction material there. Now we'll get started with using the GUI. Firstly, here's the main window:

Figure 3.1. DDRescue-GUI's Main Window (Linux)

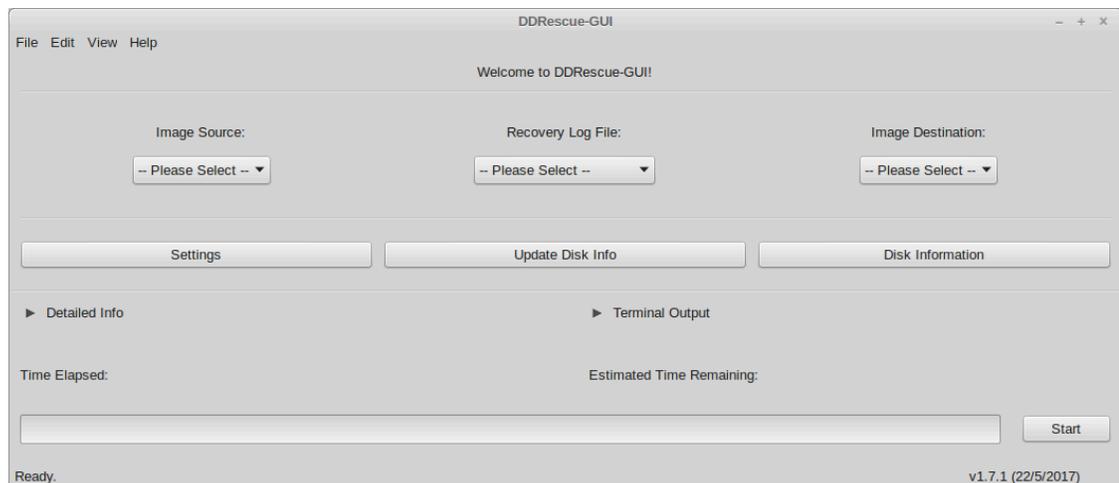
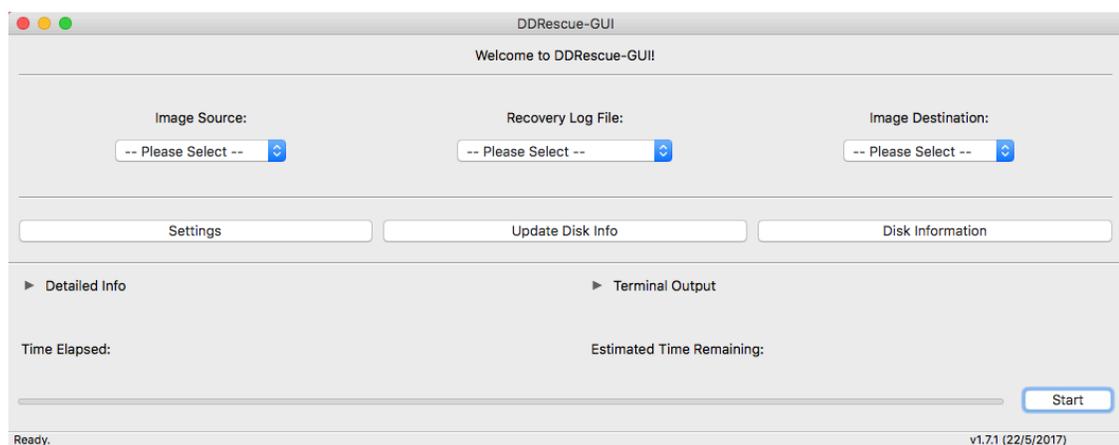


Figure 3.2. DDRescue-GUI's Main Window (Mac)





Caution

If you're using a version of ddrescue I don't support, you'll get a warning at this point. If you see the warning, check you're running the latest version of DDRRescue-GUI. The GUI may still run, but you may also encounter errors or formatting issues.

This window is displayed after DDRRescue-GUI has finished starting up. There are 3 main things you need to set here. The "Image Source" is the damaged drive. The map file is optional, and used to resume where you left off if you don't want to recover all the data at once. If you set the map file, make sure you put it on a third drive (perhaps a USB stick or another hard drive) for safety.



Tip

It's highly recommended that you use a map file. It only takes a few seconds to set up, and it means that you can stop and start the recovery as you please. It also allows you to resume your recovery if, say, there's a power outage, or your computer crashes. If that happened and you didn't use a map file you would probably have to start from the start!

Finally, you have the "Image Destination", which is where you want to recover your data to. This can be another device, or an image file. It's a good idea to go with an image file like "recovery.img". Put it in a convenient place so you don't lose it, and make sure you have enough disk space first.

You might have noticed the "Detailed Info" pull-down. That's use to show you more information about what's happening during your recovery. If you click on it, you'll see this:

Figure 3.3. Detailed info in DDRescue-GUI

Category	Value
----------	-------

When you're doing a recovery this section will have lots of information, like this:

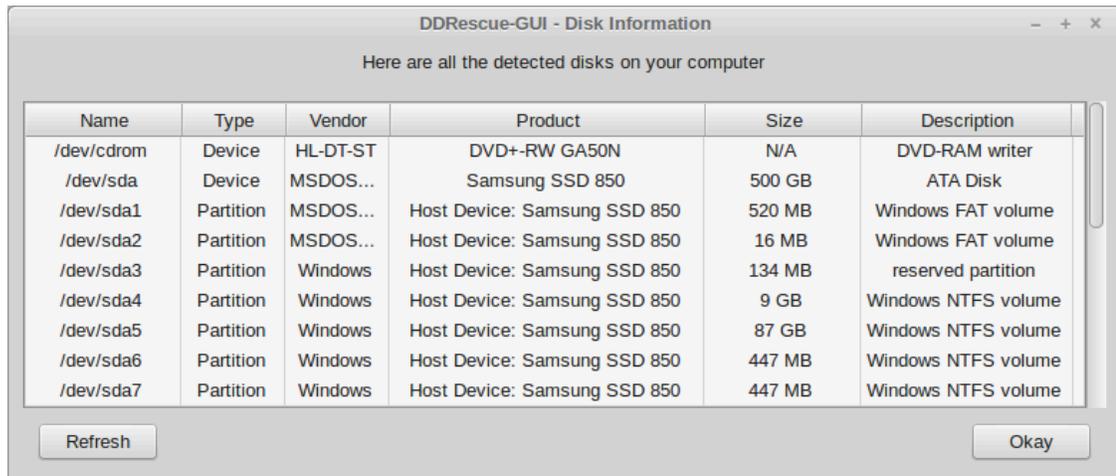
Figure 3.4. Detailed info in action in DDRescue-GUI

Category	Value
Recovered Data	104.579 MB
Unreadable Data	0 B
Current Read Rate	10485 kB/s
Average Read Rate	9507 kB/s
Bad Sectors	0
Input Position	104579 kB
Output Position	104579 kB
Time Since Last Read	0 s

The next helpful thing on this window is the button called "Disk Information". This contains all the the disk information I mentioned earlier, and it's very helpful when you're trying to match device names to disks. Here are two examples of how to use it, one with Linux and one with macOS:

Getting Started With DDRRescue-GUI

Figure 3.5. Disk info in DDRRescue-GUI (Linux)

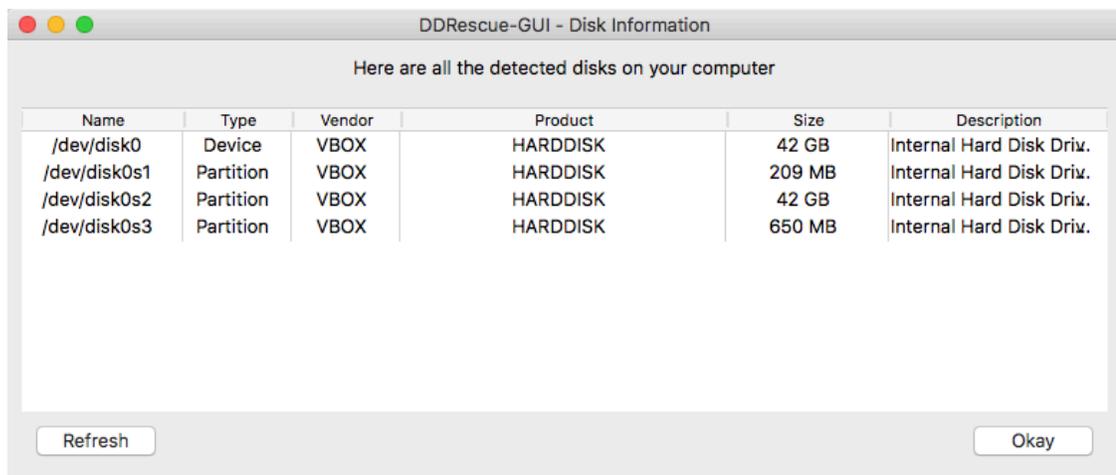


Name	Type	Vendor	Product	Size	Description
/dev/cdrom	Device	HL-DT-ST	DVD+-RW GA50N	N/A	DVD-RAM writer
/dev/sda	Device	MSDOS...	Samsung SSD 850	500 GB	ATA Disk
/dev/sda1	Partition	MSDOS...	Host Device: Samsung SSD 850	520 MB	Windows FAT volume
/dev/sda2	Partition	MSDOS...	Host Device: Samsung SSD 850	16 MB	Windows FAT volume
/dev/sda3	Partition	Windows	Host Device: Samsung SSD 850	134 MB	reserved partition
/dev/sda4	Partition	Windows	Host Device: Samsung SSD 850	9 GB	Windows NTFS volume
/dev/sda5	Partition	Windows	Host Device: Samsung SSD 850	87 GB	Windows NTFS volume
/dev/sda6	Partition	Windows	Host Device: Samsung SSD 850	447 MB	Windows NTFS volume
/dev/sda7	Partition	Windows	Host Device: Samsung SSD 850	447 MB	Windows NTFS volume

In this example, you can see in the screen-shot that I have a device whose model is "Samsung SSD 850", and because I have that in a table with device names, I now know that my SSD's name is `/dev/sda`.

And now for the Mac example:

Figure 3.6. Disk info in DDRRescue-GUI (Mac)



Name	Type	Vendor	Product	Size	Description
/dev/disk0	Device	VBOX	HARDDISK	42 GB	Internal Hard Disk Driv.
/dev/disk0s1	Partition	VBOX	HARDDISK	209 MB	Internal Hard Disk Driv.
/dev/disk0s2	Partition	VBOX	HARDDISK	42 GB	Internal Hard Disk Driv.
/dev/disk0s3	Partition	VBOX	HARDDISK	650 MB	Internal Hard Disk Driv.

In this example, you can see that there's a virtualbox hard disk with name `/dev/disk0`, and it has three partitions: `/dev/disk0s1`, `/dev/disk0s2`, and `/dev/disk0s3`. Note that the description field shows that it's an internal hard disk drive as well.

So if you were starting a recovery, you'd select the image source, map file, and image destination. You also need to pick some more settings before you can start your recovery though, and these are on the settings window.



Tip

If you accidentally click "Start" now, don't worry, because the GUI will prompt you to take a look at the settings first. This is designed to stop you from making mistakes and losing data.

Setting up for your recovery.

You're pretty close to being ready to go now, but first you need to check the settings:

Figure 3.7. DDRRescue-GUI's Settings (Linux)

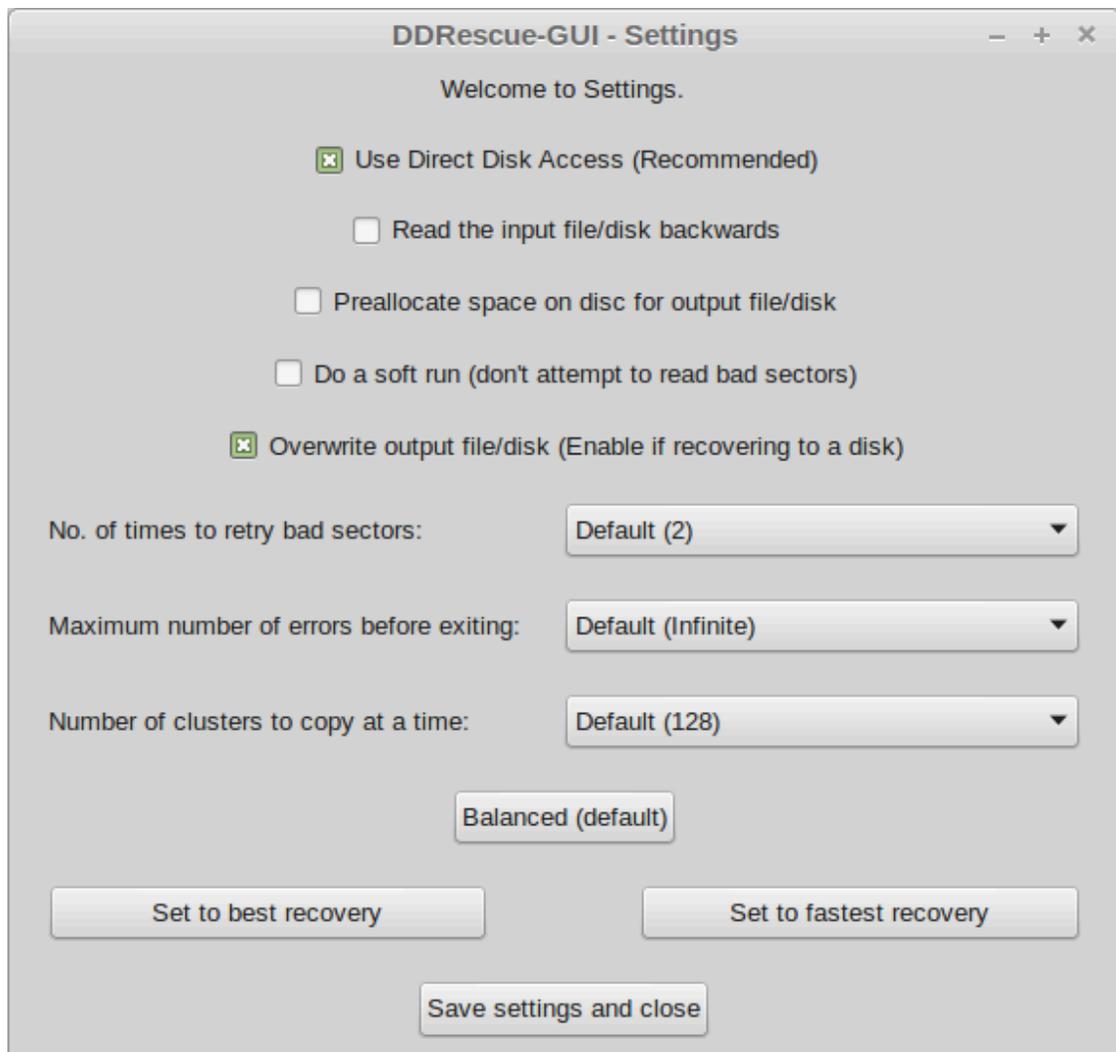
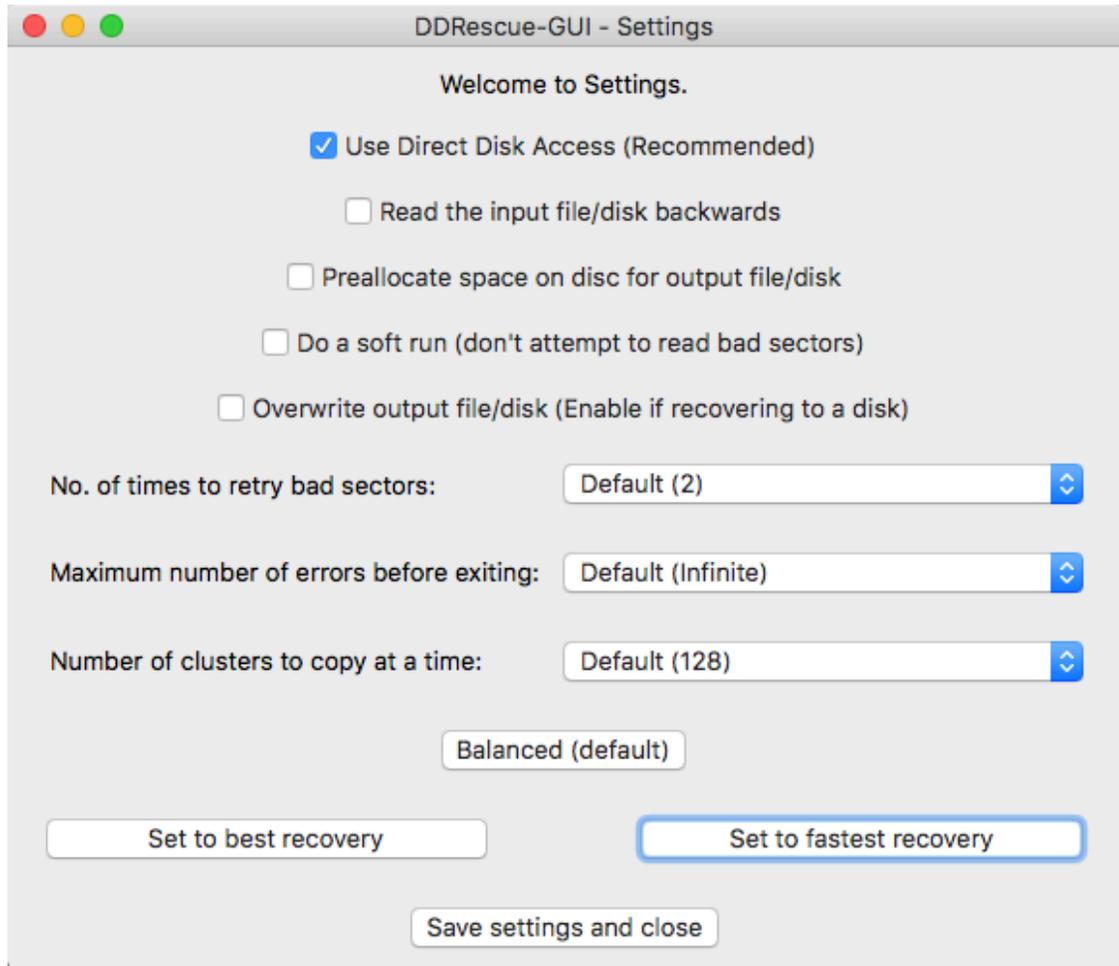
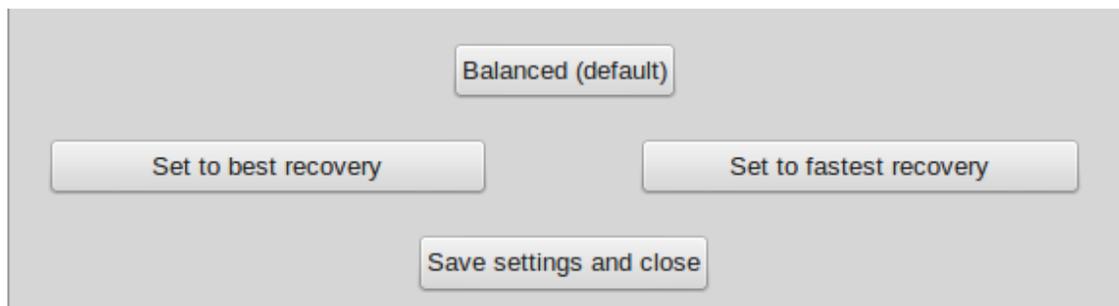


Figure 3.8. DDRRescue-GUI's Settings (Mac)



You might think this looks like a lot of options, and you'd be right. To help you get started quickly, there are three pre-sets at the bottom of the window:

Figure 3.9. DDRRescue-GUI's Settings Pre-sets



If you're in a hurry to get started, just click one of these and then save the settings, and it should be fine. If you think your disk is mostly readable, go with "Fastest".

If you aren't sure, "Balanced" is a good all-around choice. If you know that your disk is heavily damaged, pick "Best".



Note

These pre-sets set up DDRescue-GUI quickly so you can get started. If you pick the wrong option, it's okay, but it might just mean your recovery takes a bit longer. Nothing bad will happen though. If it's taking a long time, and you used a map file, feel free to stop the recovery, and then start again with a different preset. **But**, make sure you pick the same options on the main window!

There are other options to discuss here, but this is the basic functionality section, so I'm mostly focused on getting you running at the moment. For now, save the settings and close the window.

Start your recovery

You should now be ready to start recovering data. Go ahead and click the start button.

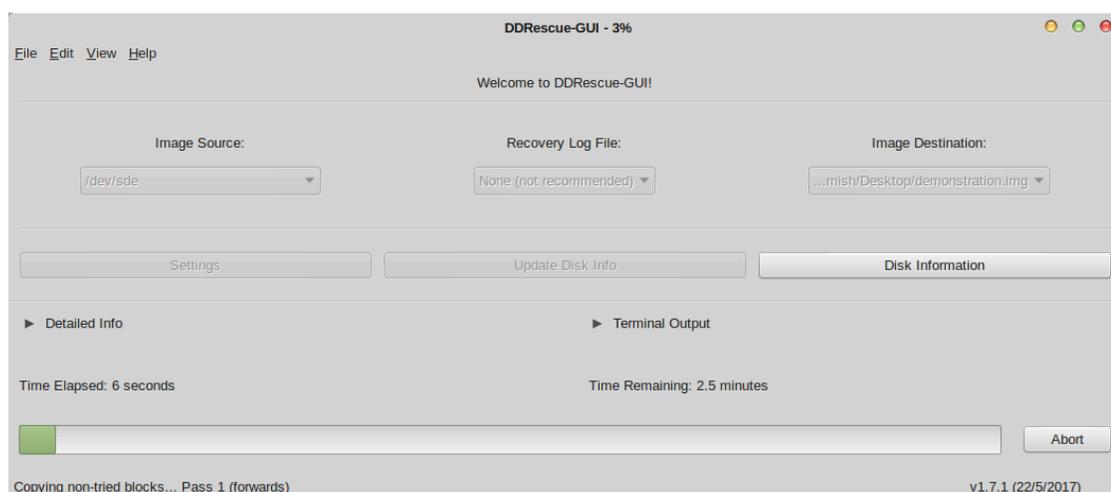


Tip

If anything is missing or set wrong, you'll be notified at this point. If you see any warnings or errors, go back and check that you set everything up right.

Hopefully, you're seeing something like this:

Figure 3.10. DDRescue-GUI in action



If you want to, click on the "Detailed Info" drop-down I mentioned before to see what's going on in more detail. At this point, you probably just need to let your

system get on with it. Data recovery can take hours, days, or even weeks. Hopefully you will have a speedy recovery, but it all depends on how badly your disk was damaged, how big it is, and what kind of damage it has taken.

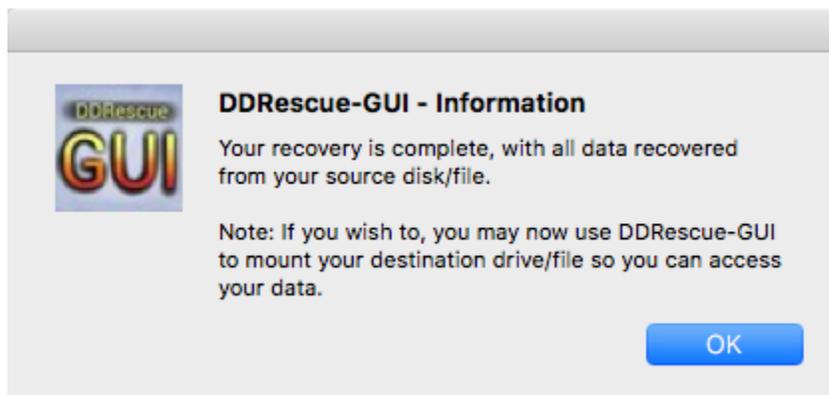
Finishing your recovery

When your recovery is finished, you should see something like this:

Figure 3.11. Recovery Success (Linux)



Figure 3.12. Recovery Success (Mac)



Hopefully, you have all of your data, in which case you can now mount your output file/device to read it if you like. To read about how to do that, go to the section called “Mounting Output Files”.

If you're unlucky, or if you stopped your recovery, you might see something more like this:

Figure 3.13. Recovery Failure/Aborted (Linux)

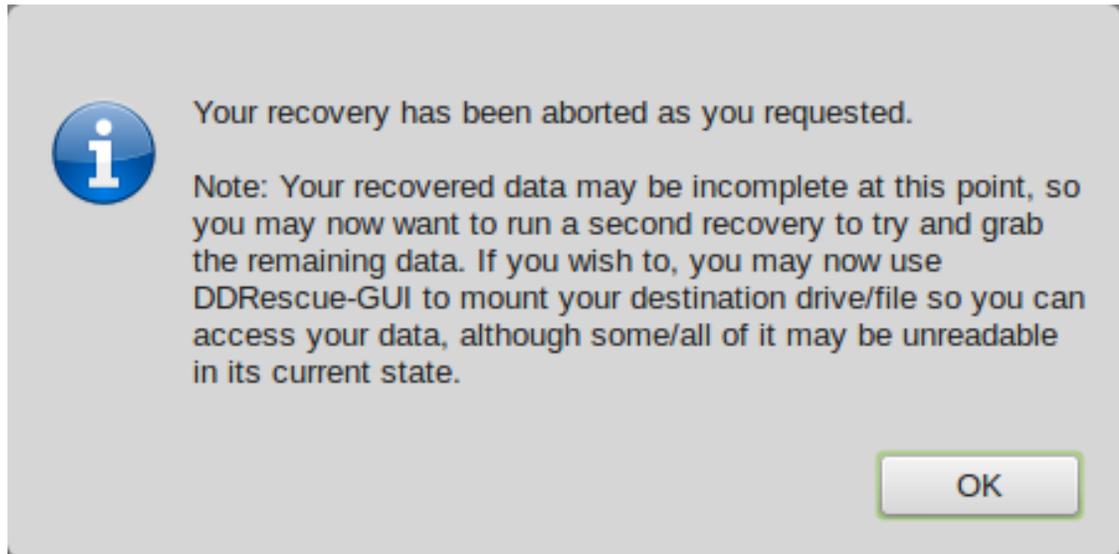
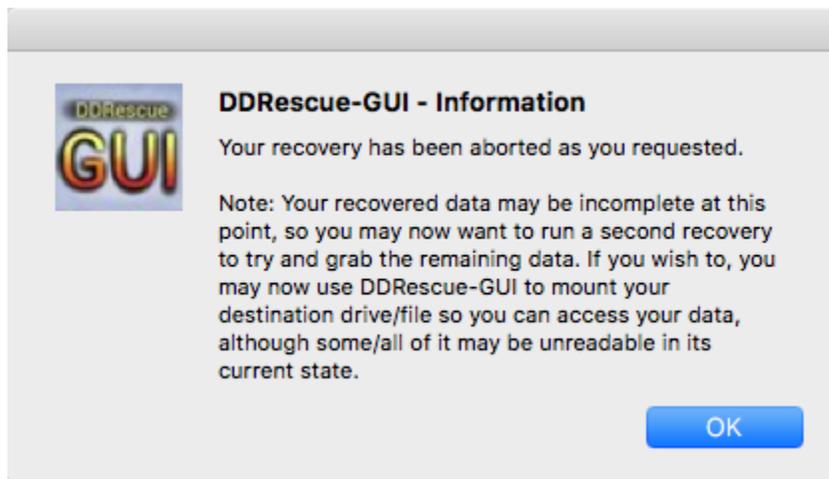
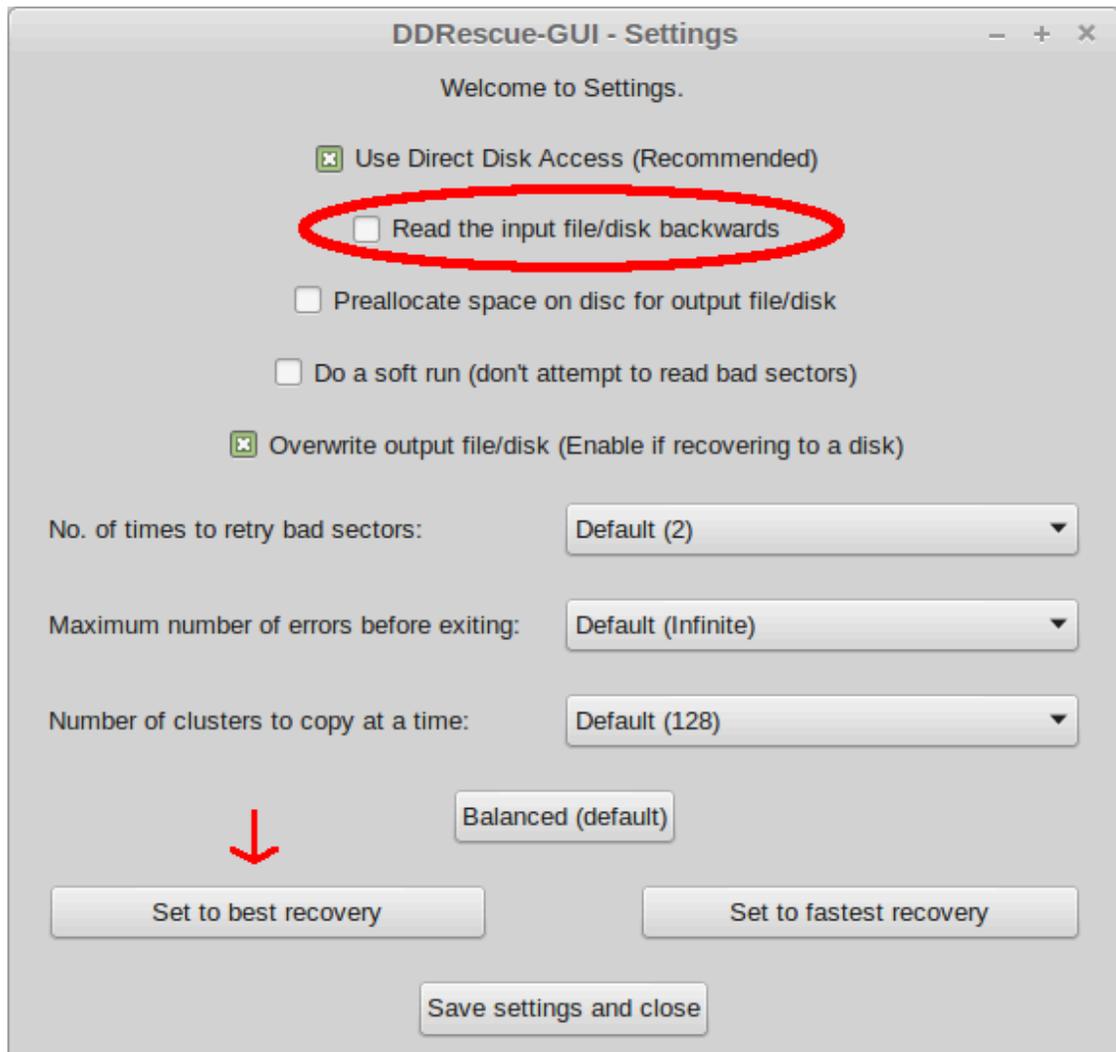


Figure 3.14. Recovery Failure/Aborted (Mac)



If so, don't worry; you can probably still get more data back, if you specified a map file. You can simply start the recovery again, with the same settings on the main window, but try reading the input file backwards, and/or set to "Best Recovery" preset. Those settings are here:

Figure 3.15. Reverse read, and preset.



Caution

Make sure you use the same settings on the main window, including the map file. Otherwise you might lose data or inadvertently start from the start!

Chapter 4. Advanced Functionality

Abstract

The last chapter covered just enough material to get you started, in case you're in a rush and something bad just happened to one of your drives. This one covers all the bases and will give you a complete understanding of how to use all the features of this software.

Why should I use a mapfile?

Using a map file is always a good idea for many reasons:

- It allows you to stop and start the recovery as and when you please.
- It allows you to restart your recovery with different settings in Settings Window if you want/need to.
- It means you can resume your recovery if your system crashes for the recovery stops for any other reason.

Without a map file, you are imposing serious limitations on yourself and the program, and also risking your data in the case of a system crash or power outage. When you set a map file, it's a good idea to put it on a drive different to your destination drive. Put it somewhere else, like on a USB stick or a third hard drive. This means that if you somehow overwrite the destination drive, you still have a map file.



Warning

If you really want to, you can opt to not use a map file. This is a really bad idea. Only do this if you've thought everything through, and are really sure you won't make use of ANY of the features in the list above.

The advanced options in the settings window

There are a lot of options here that we didn't cover in Chapter 3, so let's get started.

Direct Disk Access

This is one of GNU ddrescue's best features. Surely, though, if we're recovering from the file that represents the device, we're reading directly from it anyway? The answer is no, for a variety of reasons.

The Linux Kernel (or Mach Kernel for Mac users) buffers input and output to devices. This is good:

- It means we can write to devices faster (up to a point).
- Reading data can also be faster if we've already read that bit of data since booting.

Basically, it makes our devices seem faster. When you read from a device file like `"/dev/sda"`, the kernel manages all the reads through its buffer. The problem is that the buffer holds data in sections, of, say, 4 KB. This is normally fine, but if you can only read 3KB from a damaged disk before getting an error, you won't get ANY of that data from the kernel, because it can't fill the buffer.

Direct disk access solves this by bypassing the kernel's buffer entirely. Note that it can make recoveries slower, but you normally want to leave it enabled anyway.

Reading backwards

GNU ddrescue will normally start by reading from the start of the drive to the end of the drive ("forwards"). After that, it does several forwards and backwards passes to get as much data as possible.

However, it is possible to start by reading from the end of the disk to the start ("backwards"). Using this option also reverses all of the other passes. Reading backwards is often slower, but if you have a lot of bad sectors at the start of your drive, this might help you.

Preallocating disk space



Note

This feature is only available on Linux.

This is only really meaningful if you're using an image file. Say, you're recovering an 8GB USB drive. What it will do is make an 8GB disk image and then fill it with data, rather than expanding the disk image as new data arrives. This is useful if you're not sure you have enough disk space for the data. It's also possible it might speed the recovery up.

Soft run

As the name suggests, this option will skip over bad sectors. This option can be helpful if you want to make a really quick first pass before trying the difficult parts of the disk in some situations. Normally, it's not very helpful though, as all it does is disable the last phase, "Retrying bad sectors".

Overwrite output file

You need to enable this option if you're doing a device to device recovery, rather than recovering to an image file. This is just a safety feature.



Tip

DDRescue-GUI will automatically set this option for you when it is needed. If it's set wrong by the GUI you can override it though.

Number of times to retry bad sectors

When it's finished getting all of the good data off of the drive, GNU ddrescue will retry all of the bad sectors it's found. You can use this option to skip this step, or run it as many times as you want. Values range from 0 - 5, and then forever. The default of 2 is normally fine.

Maximum number of errors before exiting

You can tell ddrescue to exit after it's encountered a certain number of bad areas on the drive. This could be useful if you're just reading through the drive to see if there are any bad sectors, or if you're impatient and want to try different settings when the threshold is reached. Values range from infinite, and 1000 - 10. You normally just want to leave this as it is (infinite).

Number of clusters to copy at a time

On the first pass, you can set how many clusters you want ddrescue to copy at a time. It's ignored on later passes, but it can make the recovery a lot faster or slower, depending on how you set it, and what kind of disk you're recovering from.

Values range from 32 - 256. Mostly, the default of 128 is fine, but if you find it to be slow, you can adjust it. Bigger values are often faster.

Device-Device recoveries

Device to device recoveries are a little bit more difficult than device to file recoveries. The main process is the same, but:



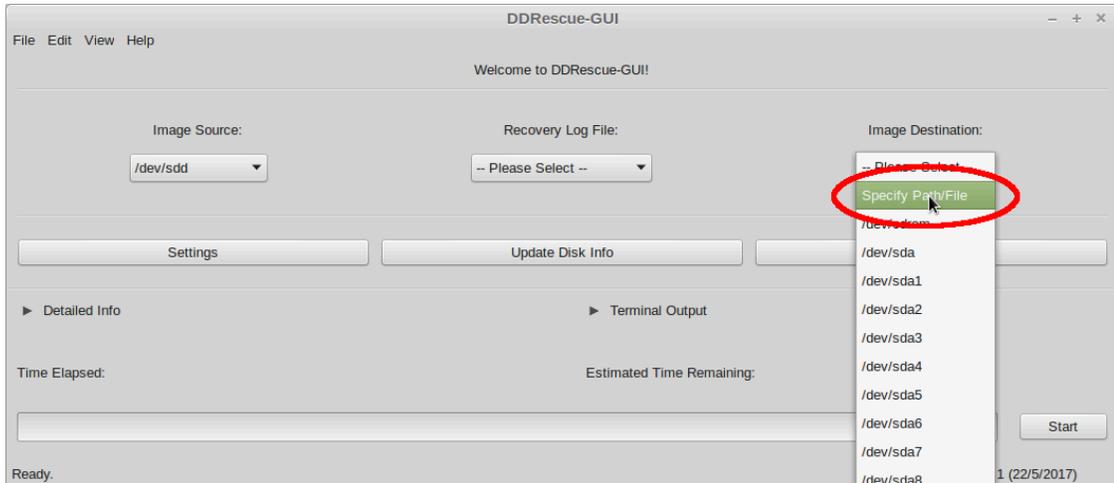
Warning

You must be *absolutely* sure that the device you're recovering to is really the device you intend to recover to!

Sadly, it often happens that people make a mistake at this point, and overwrite important data. Make sure that you use the Disk Information Window to double and triple check that the device you're recovering to is the right one. Do NOT continue until you're absolutely sure.

If you need to select a device that DDRescue-GUI hasn't detected, you can select your own destination like this on Linux:

Figure 4.1. Manually selecting an output device (Linux-only)

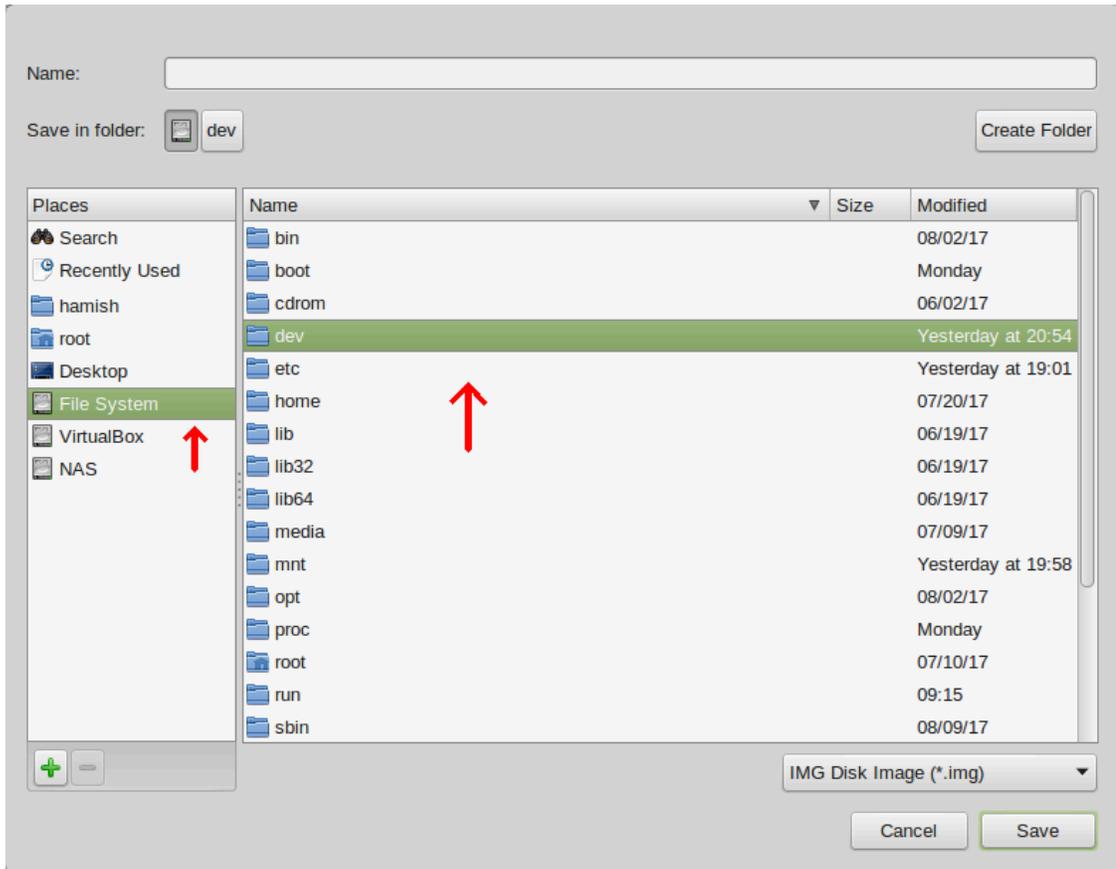


Note

On Macs, you can't yet select devices not shown in the list because they are hidden from you by macOS. In a future release of DDRescue-GUI, I will probably provide a workaround for this.

Continue, and navigate to /dev.

Figure 4.2. Navigating to /dev on Linux



Note

You may need to select "Show All Files/Devices" in the corner of the file dialog to find what you're looking for.

If you do select a device or an existing file as your destination, the GUI will warn you like this:

Figure 4.3. Confirming your destination (Linux)

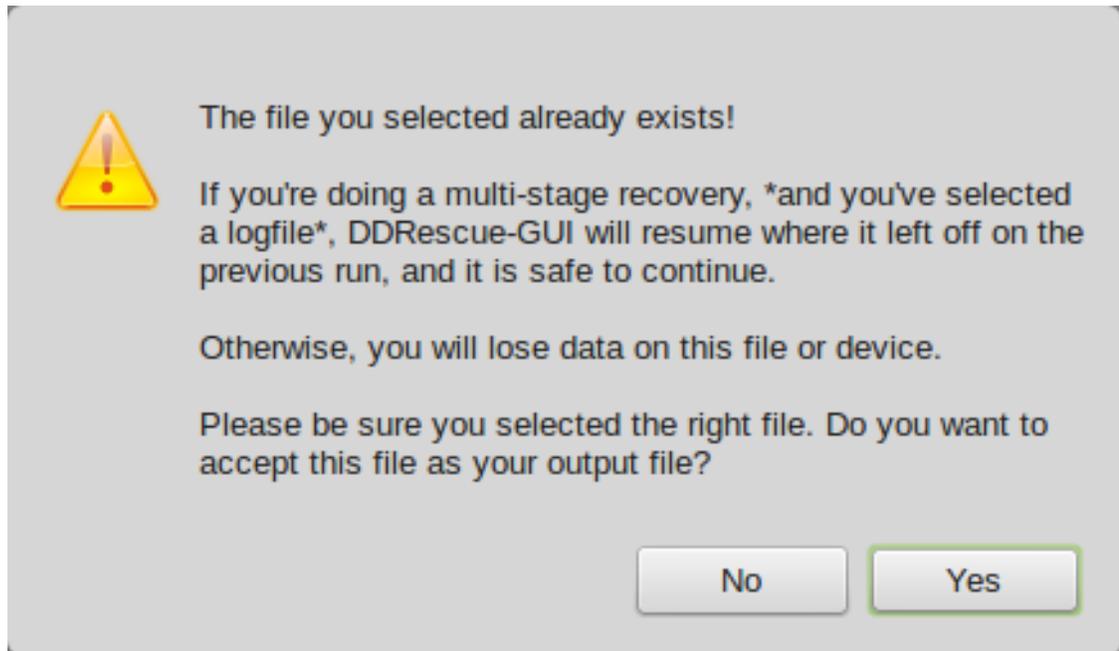
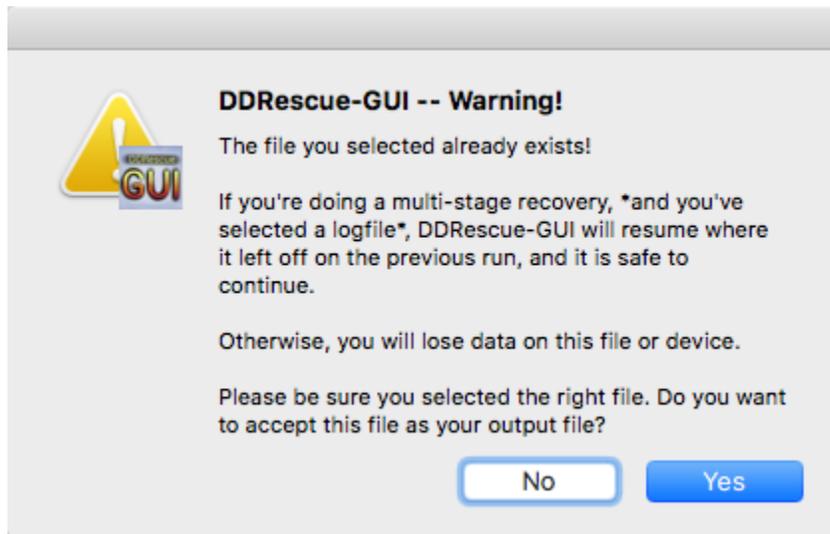


Figure 4.4. Confirming your destination (Mac)



After this point, follow the instructions as usual, making sure the "Overwrite output file/disk" option is left enabled. This leads us quite nicely on to the next topic.

Recovering individual files with DDRescue-GUI

Seeing as you can recover from device files, it stands to reason that you can recover individual files as well. Just select them under the source option, and create a new file under the destination, and all should be fine. Again, you may

need to select "View All Files/Devices" in the file dialog to see most files. This might be difficult to do on Macs.



Important

Make sure that you haven't got "Direct Disk Access" checked in the settings window, as it may cause strange behaviour.



Note

If you're wondering why this is helpful, it might be if you've only got a few files you care about on a failing drive, and you've checked that you can't read/copy them the normal way. This should work especially well if they're large files.

Mounting Output Files

This is one of the most helpful features of DDRescue-GUI, especially to people who are new to data recovery or are not technically-minded. It allows you to view all the data you recovered easily. The procedure is a bit different depending on whether you recovered a partition or a device, but it is very similar.

When recovering from an entire device

When your recovery is finished, click "Mount Image/Disk".

Figure 4.5. The mount button (Linux)

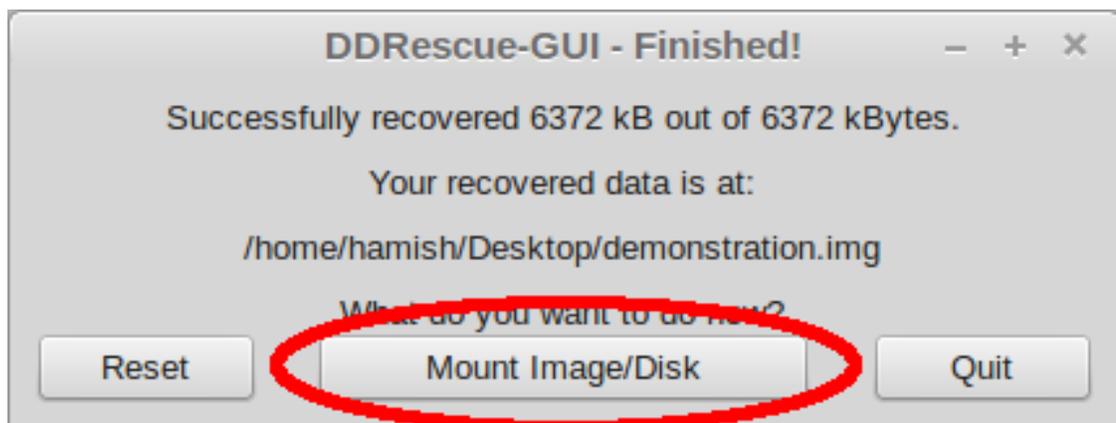
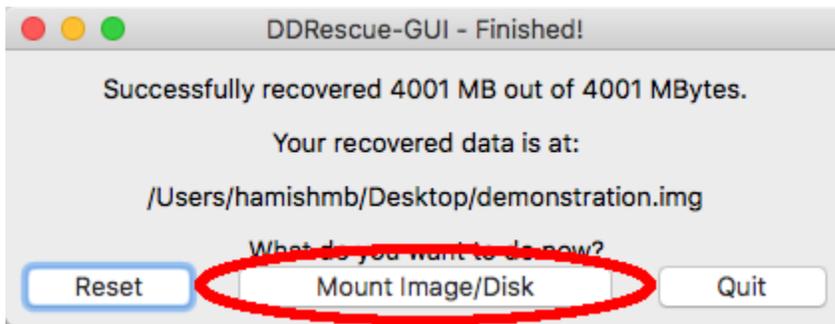


Figure 4.6. The mount button (Mac)



You will be presented with a window that looks like this:

Figure 4.7. Selecting a partition (Linux)

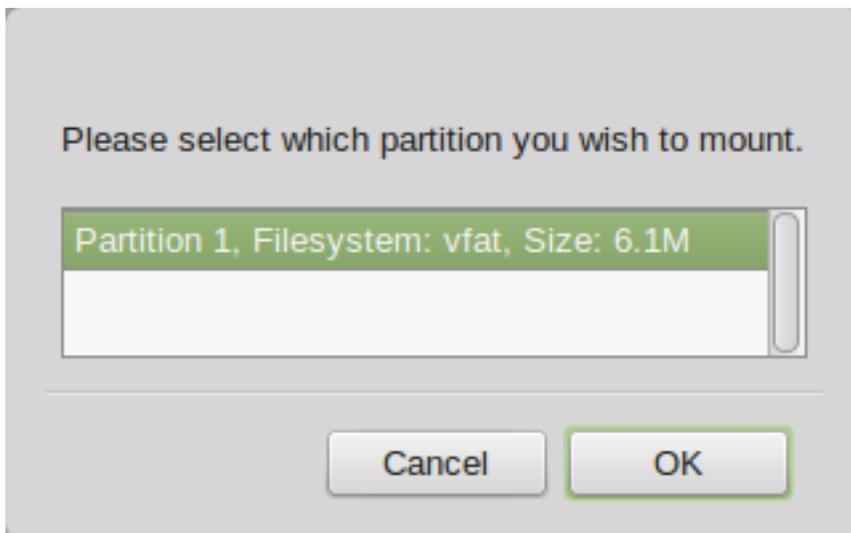
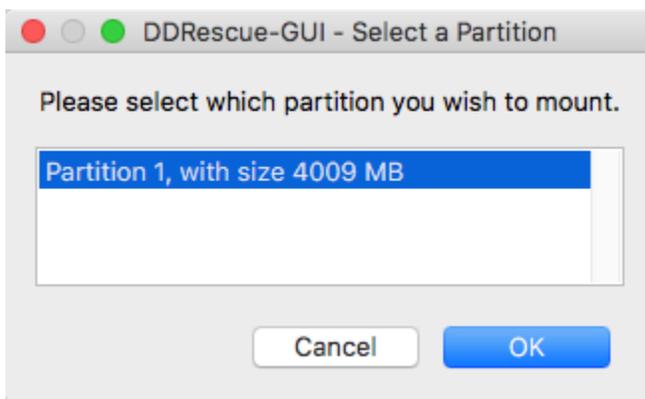


Figure 4.8. Selecting a partition (Mac)



The information here is like the Disk Information Window, but more basic. Select a partition you want to mount, and click okay. After a few seconds, you should see this:

Figure 4.9. Successfully mounting a partition in your device (Linux)

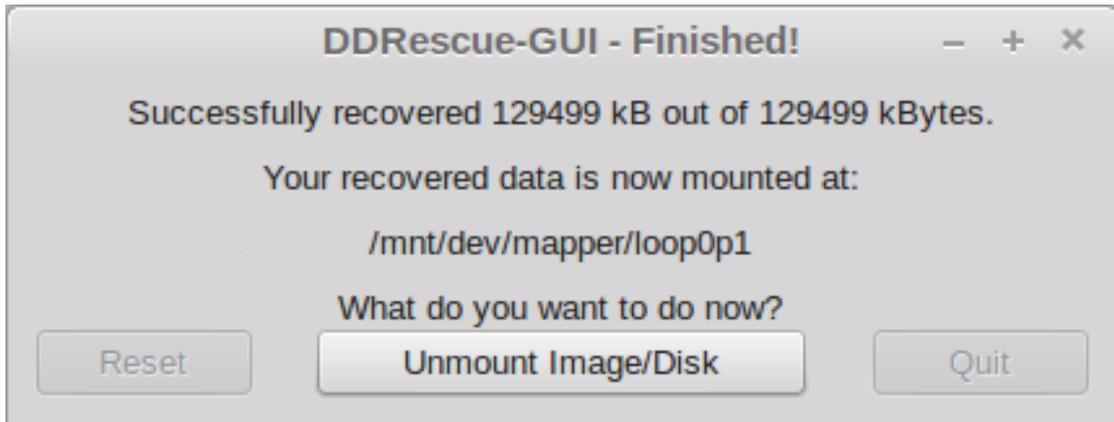
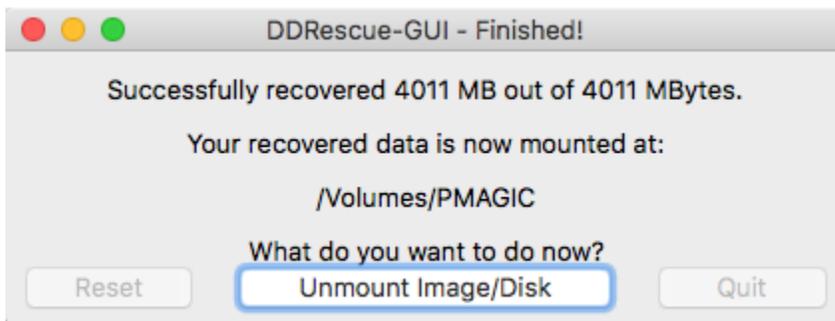


Figure 4.10. Successfully mounting a partition in your device (Mac)



Tip

Don't worry if you select the wrong partition; you can just try again. Often, there will be a few partitions that are unmountable. These don't hold any meaningful data, so don't worry about it, just pick a different partition and try again until you find your data.

And simply navigate to where your files are, and view or copy as desired. When you're finished, click "Unmount", and you'll be returned to the previous window.

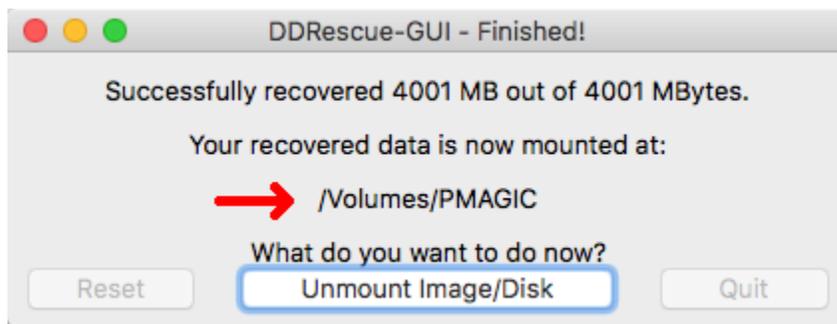
When recovering from a partition

This is a bit simpler; just click "Mount Image-Disk", and it'll do it for you. After a few seconds, you'll see:

Figure 4.11. Successfully mounting your partition (Linux)



Figure 4.12. Successfully mounting your partition (Mac)



And then navigate to where your files are. When you're finished, click "Unmount", and you'll be returned to the previous window.

Chapter 5. Frequently Asked Questions

Abstract

The final chapter in this user guide has frequently asked questions, which will hopefully help resolve any issues you have. I've also added a few that I thought might be helpful.

Setup-related Questions

Q: Can I run DDRescue-GUI on Windows?

A: No, but I might add support in a future release if there's demand.

Q: Why do I need to use a mapfile?

A: There's are lots of reasons, and it really is very important. See the section called "Why should I use a mapfile?".

Q: What is "Direct Disk Access"?

A: Mostly, you should just leave this on, because it usually makes it easier to read damaged drives and helps you recover more data. If you want to know how it works, see the section called "Direct Disk Access".

Q: Why is the read backwards option helpful?

A: If the very start of your disk is difficult to read, it may enable you to get most of the data more quickly. If you think your drive is about to stop working completely, this could help you a lot.

Q: Can I recover to optical (CD/DVD/Bluray) drives?

A: Unfortunately, no. This would be difficult to implement, and also writing to these drives is slow; if you think your source drive might be about to die, the last thing you want is to hold it up writing an optical disk.

You could recover to an image file, then use tools to split that file into CD/DVD sized chunks and write them, but that may not be helpful anyway; sometimes you need to repair filesystems and so on before you can read data. Reading disks like that would be difficult as well, and more suitable for a backup program to do.

Q: Can I run more than one instance of DDRescue-GUI at once?

A: At the time of writing, no. While there's nothing stopping you from doing this, the GUI's log file would be open twice, and would probably corrupt. So, you probably CAN do it, but I haven't tested that scenario, and your GUI

logfile would probably be corrupted and unusable if anything went wrong. Note that the recovery map file would probably be fine though.

Note that I might add support for this in a future release.

Q: Can I restart DDRescue-GUI after aborting my recovery?

A: Absolutely, but only if you used a map file. Just make sure you have exactly the same input, output, and map file settings, and you should be good to go. You can safely change the settings in Settings Window, which allows you to try different approaches to get your data more quickly.

Q: What preset/settings should I use?

A: See the section called "Setting up for your recovery." for the basics. If you want more details, have a look at the section called "The advanced options in the settings window".

Q: MacOS: Where are the /dev/rdisk* devices?

A: They are intentionally hidden. If you want to use direct disk access, pick your device the usual way, and tick "Use Direct Disk Access" in the settings window.

Q: Can I use a domain mapfile to speed my recovery up with DDRescue-GUI?

A: At the time of writing, you can't. This is because DDRescue-GUI supports old versions of ddrescue that don't have that feature, and the process of making a domain mapfile is different for different filesystems; it would probably become complicated, and it would be hard to make a simple interface for this.

Nevertheless, I might implement this in a future release if there is demand for it.

Q: Can I use DDRescue-GUI to recover from encrypted volumes?

A: Possibly. On Linux, LUKS encryption doesn't work with DDRescue-GUI at the moment, but other encryption methods may work, but I haven't tested the GUI with them. The same goes for MacOS encryption, so your mileage may vary; it may work, but is untested at the time of writing.

Runtime Questions

Q: The remaining time estimate doesn't seem very accurate, why is that?

A: My algorithm for calculating this is very simple, so if you have a large bad area of a disk, and then lots of easily-readable data, the time estimate

might take a while to adjust. New versions of GNU ddrescue have a better estimate, so you can show terminal output to see that if you like. If you can't see it, you'll just have to make do with the GUI's estimate, which I may improve in the future.

Q: Recovery is very slow, why?

A: This is quite normal, sadly. It's much harder to read damaged disks than good ones, so transfer rates can be only a few KB/s or worse. You can try different settings in the settings window to see if that improves the situation.

Note that on certain types of drive with flash memory, like USB sticks and SD cards, using the "Best Recovery" preset can slow you down. Instead try "Balanced", or try turning direct disk access off.

Post-recovery Questions

Q: Should I mount my output file/device?

A: It really depends on what you want to do. If you want to check what kind of state your data is in, it might be helpful. You might even be able to start copying data from, say, an image file to a USB stick if it is intact.

However, you may find that you can't mount your output file. This is especially likely to happen if you couldn't get all the data off your source drive, but don't worry! You can use tools like photorec and testdisk to recover your files from the output file/device, and you'll probably get a lot of them back still.

Q: Why can't I mount my output file?

A: It could be because you didn't recover enough data from your source drive, the data is damaged, or that your operating system doesn't understand the filesystem. On MacOS, you probably won't be able to read Linux-style partitions, and on Linux you may struggle to read exFAT, and HFS+ partitions depending on your setup.

Don't panic if this happens to you. If there's more data to recover, try doing that first. If it still doesn't work, make sure the filesystem is supported on your OS, and then in turn try using testdisk, photorec or similar tools to read the output file/device.

Q: Can I restart DDRescue-GUI and try again?

A: Yes, but only if you have a map file. Just make sure you use the same options on the main window. Restarting with different settings often allows

you to pull more data off your drive. This work especially if it's say, a CD, and you have multiple CD drives you can try to read it with. In that case, you may need to change the input device, but make sure you get it right, and use the Disk Info window to help you.

Q: How can I mount my output file/disk after closing DDRescue-GUI?

A: If you are recovering to a disk, just use your operating system's tools like you would with, for example, a USB drive.

If you are recovering to a file, the situation is somewhat different. It's easier to just start the recovery again with the same settings and abort immediately. In future, I will make this easier to do.

Troubleshooting Questions

Q: Linux: I get a warning that my ddrescue version is unsupported on startup.

A: First, make sure you're running the newest version of DDRescue-GUI. If you still have the error, you may be using a newer version of ddrescue than I support, in which case please try to downgrade. If you can't downgrade, you may be able to use the GUI anyway. Alternatively, if your ddrescue version is too old, the GUI will most likely not work for you, unless you can upgrade.

Q: MacOS: The GUI doesn't start.

A: Please make sure you're using a supported version of MacOS. I support versions from 10.9 (Mavericks) onwards, and I test with the betas for the new versions of macOS before they come out. If you need support for an older release like 10.6, 10.7 or 10.8, try DDRescue-GUI version 1.6.1 or 1.6. Note that these older versions look slightly different, and have fewer features, so the screenshots in the guide may be slightly wrong for you. There is absolutely no support for versions of MacOS prior to 10.6, or for PowerPC Macs.

Q: It's taking a long time "Getting new Disk Information" when I open the GUI, why?

A: This is normal if you have a lot of disks, or an older computer. If you think the GUI has crashed, try rebooting your computer and starting DDRescue-GUI again.

If you're still having trouble, you may have disks the GUI doesn't support, so remove any external disks and try again. Note that encrypted disks are not supported and may cause problems.

Q: The information/formatting is wrong on the GUI

A: Linux Users: Please check that you're using a supported version of ddrescue. You can check in the About window. If you don't have a supported version of GNU ddrescue, you will see a warning when you start the GUI.

Mac Users: Check that you're using a supported version of MacOS (10.9 or higher). If you are, please report this bug to me by contacting me at www.launchpad.net/~hamishmb [<http://www.launchpad.net/~hamishmb>].

Q: The GUI crashed while recovering data! Help!

A: It's very unlikely that this would happen, but nevertheless:

Don't panic. Your data is almost definitely fine. Reboot your computer, select all the same settings, and try again. Note that this issue is more likely to happen if you're using an unsupported version of ddrescue.

Q: I can't mount my output file!

A: See [Why can't I mount my output file?](#).

Appendix A. License For The Stylesheets used to generate this documentation.

Note that I have slightly modified the PDF stylesheet so that it uses a sans-serif font instead of a serif font for ease of reading. Other than that, the styleheets used to generate this documentation are identical to the standard DocBook XSL Stylesheets.

Copyright

Copyright (C) 1999-2007 Norman Walsh

Copyright (C) 2003 JiÅ™Ã Kosek

Copyright (C) 2004-2007 Steve Ball

Copyright (C) 2005-2014 The DocBook Project

Copyright (C) 2011-2012 O'Reilly Media

Permission is hereby granted, free of charge, to any person obtaining a copy of this software and associated documentation files (the ``Software"), to deal in the Software without restriction, including without limitation the rights to use, copy, modify, merge, publish, distribute, sublicense, and/or sell copies of the Software, and to permit persons to whom the Software is furnished to do so, subject to the following conditions:

The above copyright notice and this permission notice shall be included in all copies or substantial portions of the Software.

Except as contained in this notice, the names of individuals credited with contribution to this software shall not be used in advertising or otherwise to promote the sale, use or other dealings in this Software without prior written authorization from the individuals in question.

Any stylesheet derived from this Software that is publically distributed will be identified with a different name and the version strings in any derived Software will be changed so that no possibility of confusion between the derived package and this Software will exist.

License For The Stylesheets
used to generate
this documentation.

Warranty

THE SOFTWARE IS PROVIDED "AS IS", WITHOUT WARRANTY OF ANY KIND, EXPRESS OR IMPLIED, INCLUDING BUT NOT LIMITED TO THE WARRANTIES OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE AND NONINFRINGEMENT. IN NO EVENT SHALL NORMAN WALSH OR ANY OTHER CONTRIBUTOR BE LIABLE FOR ANY CLAIM, DAMAGES OR OTHER LIABILITY, WHETHER IN AN ACTION OF CONTRACT, TORT OR OTHERWISE, ARISING FROM, OUT OF OR IN CONNECTION WITH THE SOFTWARE OR THE USE OR OTHER DEALINGS IN THE SOFTWARE.

Contacting the Author

The DocBook XSL stylesheets are maintained by Norman Walsh, <ndw@nwalsh.com>, and members of the DocBook Project, <docbook-developers@sf.net>

Appendix B. License For This Documentation

GNU Free Documentation License
Version 1.3, 3 November 2008

Copyright (C) 2000, 2001, 2002, 2007, 2008 Free Software Foundation, Inc.
<<http://fsf.org/>>

Everyone is permitted to copy and distribute verbatim copies of this license document, but changing it is not allowed.

0. PREAMBLE

The purpose of this License is to make a manual, textbook, or other functional and useful document "free" in the sense of freedom: to assure everyone the effective freedom to copy and redistribute it, with or without modifying it, either commercially or noncommercially. Secondly, this License preserves for the author and publisher a way to get credit for their work, while not being considered responsible for modifications made by others.

This License is a kind of "copyleft", which means that derivative works of the document must themselves be free in the same sense. It complements the GNU General Public License, which is a copyleft license designed for free software.

We have designed this License in order to use it for manuals for free software, because free software needs free documentation: a free program should come with manuals providing the same freedoms that the software does. But this License is not limited to software manuals; it can be used for any textual work, regardless of subject matter or whether it is published as a printed book. We recommend this License principally for works whose purpose is instruction or reference.

1. APPLICABILITY AND DEFINITIONS

This License applies to any manual or other work, in any medium, that contains a notice placed by the copyright holder saying it can be distributed under the terms of this License. Such a notice grants a

License For This Documentation

world-wide, royalty-free license, unlimited in duration, to use that work under the conditions stated herein. The "Document", below, refers to any such manual or work. Any member of the public is a licensee, and is addressed as "you". You accept the license if you copy, modify or distribute the work in a way requiring permission under copyright law.

A "Modified Version" of the Document means any work containing the Document or a portion of it, either copied verbatim, or with modifications and/or translated into another language.

A "Secondary Section" is a named appendix or a front-matter section of the Document that deals exclusively with the relationship of the publishers or authors of the Document to the Document's overall subject (or to related matters) and contains nothing that could fall directly within that overall subject. (Thus, if the Document is in part a textbook of mathematics, a Secondary Section may not explain any mathematics.) The relationship could be a matter of historical connection with the subject or with related matters, or of legal, commercial, philosophical, ethical or political position regarding them.

The "Invariant Sections" are certain Secondary Sections whose titles are designated, as being those of Invariant Sections, in the notice that says that the Document is released under this License. If a section does not fit the above definition of Secondary then it is not allowed to be designated as Invariant. The Document may contain zero Invariant Sections. If the Document does not identify any Invariant Sections then there are none.

The "Cover Texts" are certain short passages of text that are listed, as Front-Cover Texts or Back-Cover Texts, in the notice that says that the Document is released under this License. A Front-Cover Text may be at most 5 words, and a Back-Cover Text may be at most 25 words.

A "Transparent" copy of the Document means a machine-readable copy, represented in a format whose specification is available to the general public, that is suitable for revising the document straightforwardly with generic text editors or (for images composed of pixels) generic paint programs or (for drawings) some widely available drawing editor, and that is suitable for input to text formatters or for automatic translation to a variety of formats suitable for input to text formatters. A copy made in an otherwise Transparent file format whose markup, or absence of markup, has been arranged to thwart or discourage subsequent modification by readers is not Transparent.

An image format is not Transparent if used for any substantial amount of text. A copy that is not "Transparent" is called "Opaque".

Examples of suitable formats for Transparent copies include plain ASCII without markup, Texinfo input format, LaTeX input format, SGML or XML using a publicly available DTD, and standard-conforming simple HTML, PostScript or PDF designed for human modification. Examples of transparent image formats include PNG, XCF and JPG. Opaque formats include proprietary formats that can be read and edited only by proprietary word processors, SGML or XML for which the DTD and/or processing tools are not generally available, and the machine-generated HTML, PostScript or PDF produced by some word processors for output purposes only.

The "Title Page" means, for a printed book, the title page itself, plus such following pages as are needed to hold, legibly, the material this License requires to appear in the title page. For works in formats which do not have any title page as such, "Title Page" means the text near the most prominent appearance of the work's title, preceding the beginning of the body of the text.

The "publisher" means any person or entity that distributes copies of the Document to the public.

A section "Entitled XYZ" means a named subunit of the Document whose title either is precisely XYZ or contains XYZ in parentheses following text that translates XYZ in another language. (Here XYZ stands for a specific section name mentioned below, such as "Acknowledgements", "Dedications", "Endorsements", or "History".) To "Preserve the Title" of such a section when you modify the Document means that it remains a section "Entitled XYZ" according to this definition.

The Document may include Warranty Disclaimers next to the notice which states that this License applies to the Document. These Warranty Disclaimers are considered to be included by reference in this License, but only as regards disclaiming warranties: any other implication that these Warranty Disclaimers may have is void and has no effect on the meaning of this License.

2. VERBATIM COPYING

You may copy and distribute the Document in any medium, either commercially or noncommercially, provided that this License, the copyright notices, and the license notice saying this License applies to the Document are reproduced in all copies, and that you add no

other conditions whatsoever to those of this License. You may not use technical measures to obstruct or control the reading or further copying of the copies you make or distribute. However, you may accept compensation in exchange for copies. If you distribute a large enough number of copies you must also follow the conditions in section 3.

You may also lend copies, under the same conditions stated above, and you may publicly display copies.

3. COPYING IN QUANTITY

If you publish printed copies (or copies in media that commonly have printed covers) of the Document, numbering more than 100, and the Document's license notice requires Cover Texts, you must enclose the copies in covers that carry, clearly and legibly, all these Cover Texts: Front-Cover Texts on the front cover, and Back-Cover Texts on the back cover. Both covers must also clearly and legibly identify you as the publisher of these copies. The front cover must present the full title with all words of the title equally prominent and visible. You may add other material on the covers in addition. Copying with changes limited to the covers, as long as they preserve the title of the Document and satisfy these conditions, can be treated as verbatim copying in other respects.

If the required texts for either cover are too voluminous to fit legibly, you should put the first ones listed (as many as fit reasonably) on the actual cover, and continue the rest onto adjacent pages.

If you publish or distribute Opaque copies of the Document numbering more than 100, you must either include a machine-readable Transparent copy along with each Opaque copy, or state in or with each Opaque copy a computer-network location from which the general network-using public has access to download using public-standard network protocols a complete Transparent copy of the Document, free of added material. If you use the latter option, you must take reasonably prudent steps, when you begin distribution of Opaque copies in quantity, to ensure that this Transparent copy will remain thus accessible at the stated location until at least one year after the last time you distribute an Opaque copy (directly or through your agents or retailers) of that edition to the public.

It is requested, but not required, that you contact the authors of the Document well before redistributing any large number of copies, to

give them a chance to provide you with an updated version of the Document.

4. MODIFICATIONS

You may copy and distribute a Modified Version of the Document under the conditions of sections 2 and 3 above, provided that you release the Modified Version under precisely this License, with the Modified Version filling the role of the Document, thus licensing distribution and modification of the Modified Version to whoever possesses a copy of it. In addition, you must do these things in the Modified Version:

- A. Use in the Title Page (and on the covers, if any) a title distinct from that of the Document, and from those of previous versions (which should, if there were any, be listed in the History section of the Document). You may use the same title as a previous version if the original publisher of that version gives permission.
- B. List on the Title Page, as authors, one or more persons or entities responsible for authorship of the modifications in the Modified Version, together with at least five of the principal authors of the Document (all of its principal authors, if it has fewer than five), unless they release you from this requirement.
- C. State on the Title page the name of the publisher of the Modified Version, as the publisher.
- D. Preserve all the copyright notices of the Document.
- E. Add an appropriate copyright notice for your modifications adjacent to the other copyright notices.
- F. Include, immediately after the copyright notices, a license notice giving the public permission to use the Modified Version under the terms of this License, in the form shown in the Addendum below.
- G. Preserve in that license notice the full lists of Invariant Sections and required Cover Texts given in the Document's license notice.
- H. Include an unaltered copy of this License.
- I. Preserve the section Entitled "History", Preserve its Title, and add to it an item stating at least the title, year, new authors, and publisher of the Modified Version as given on the Title Page. If there is no section Entitled "History" in the Document, create one stating the title, year, authors, and publisher of the Document as given on its Title Page, then add an item describing the Modified Version as stated in the previous sentence.
- J. Preserve the network location, if any, given in the Document for public access to a Transparent copy of the Document, and likewise the network locations given in the Document for previous versions it was based on. These may be placed in the "History" section.

License For This Documentation

- You may omit a network location for a work that was published at least four years before the Document itself, or if the original publisher of the version it refers to gives permission.
- K. For any section Entitled "Acknowledgements" or "Dedications", Preserve the Title of the section, and preserve in the section all the substance and tone of each of the contributor acknowledgements and/or dedications given therein.
 - L. Preserve all the Invariant Sections of the Document, unaltered in their text and in their titles. Section numbers or the equivalent are not considered part of the section titles.
 - M. Delete any section Entitled "Endorsements". Such a section may not be included in the Modified Version.
 - N. Do not retitle any existing section to be Entitled "Endorsements" or to conflict in title with any Invariant Section.
 - O. Preserve any Warranty Disclaimers.

If the Modified Version includes new front-matter sections or appendices that qualify as Secondary Sections and contain no material copied from the Document, you may at your option designate some or all of these sections as invariant. To do this, add their titles to the list of Invariant Sections in the Modified Version's license notice. These titles must be distinct from any other section titles.

You may add a section Entitled "Endorsements", provided it contains nothing but endorsements of your Modified Version by various parties--for example, statements of peer review or that the text has been approved by an organization as the authoritative definition of a standard.

You may add a passage of up to five words as a Front-Cover Text, and a passage of up to 25 words as a Back-Cover Text, to the end of the list of Cover Texts in the Modified Version. Only one passage of Front-Cover Text and one of Back-Cover Text may be added by (or through arrangements made by) any one entity. If the Document already includes a cover text for the same cover, previously added by you or by arrangement made by the same entity you are acting on behalf of, you may not add another; but you may replace the old one, on explicit permission from the previous publisher that added the old one.

The author(s) and publisher(s) of the Document do not by this License give permission to use their names for publicity for or to assert or imply endorsement of any Modified Version.

5. COMBINING DOCUMENTS

You may combine the Document with other documents released under this License, under the terms defined in section 4 above for modified versions, provided that you include in the combination all of the Invariant Sections of all of the original documents, unmodified, and list them all as Invariant Sections of your combined work in its license notice, and that you preserve all their Warranty Disclaimers.

The combined work need only contain one copy of this License, and multiple identical Invariant Sections may be replaced with a single copy. If there are multiple Invariant Sections with the same name but different contents, make the title of each such section unique by adding at the end of it, in parentheses, the name of the original author or publisher of that section if known, or else a unique number. Make the same adjustment to the section titles in the list of Invariant Sections in the license notice of the combined work.

In the combination, you must combine any sections Entitled "History" in the various original documents, forming one section Entitled "History"; likewise combine any sections Entitled "Acknowledgements", and any sections Entitled "Dedications". You must delete all sections Entitled "Endorsements".

6. COLLECTIONS OF DOCUMENTS

You may make a collection consisting of the Document and other documents released under this License, and replace the individual copies of this License in the various documents with a single copy that is included in the collection, provided that you follow the rules of this License for verbatim copying of each of the documents in all other respects.

You may extract a single document from such a collection, and distribute it individually under this License, provided you insert a copy of this License into the extracted document, and follow this License in all other respects regarding verbatim copying of that document.

7. AGGREGATION WITH INDEPENDENT WORKS

A compilation of the Document or its derivatives with other separate and independent documents or works, in or on a volume of a storage or distribution medium, is called an "aggregate" if the copyright

resulting from the compilation is not used to limit the legal rights of the compilation's users beyond what the individual works permit. When the Document is included in an aggregate, this License does not apply to the other works in the aggregate which are not themselves derivative works of the Document.

If the Cover Text requirement of section 3 is applicable to these copies of the Document, then if the Document is less than one half of the entire aggregate, the Document's Cover Texts may be placed on covers that bracket the Document within the aggregate, or the electronic equivalent of covers if the Document is in electronic form. Otherwise they must appear on printed covers that bracket the whole aggregate.

8. TRANSLATION

Translation is considered a kind of modification, so you may distribute translations of the Document under the terms of section 4. Replacing Invariant Sections with translations requires special permission from their copyright holders, but you may include translations of some or all Invariant Sections in addition to the original versions of these Invariant Sections. You may include a translation of this License, and all the license notices in the Document, and any Warranty Disclaimers, provided that you also include the original English version of this License and the original versions of those notices and disclaimers. In case of a disagreement between the translation and the original version of this License or a notice or disclaimer, the original version will prevail.

If a section in the Document is Entitled "Acknowledgements", "Dedications", or "History", the requirement (section 4) to Preserve its Title (section 1) will typically require changing the actual title.

9. TERMINATION

You may not copy, modify, sublicense, or distribute the Document except as expressly provided under this License. Any attempt otherwise to copy, modify, sublicense, or distribute it is void, and will automatically terminate your rights under this License.

However, if you cease all violation of this License, then your license from a particular copyright holder is reinstated (a) provisionally,

unless and until the copyright holder explicitly and finally terminates your license, and (b) permanently, if the copyright holder fails to notify you of the violation by some reasonable means prior to 60 days after the cessation.

Moreover, your license from a particular copyright holder is reinstated permanently if the copyright holder notifies you of the violation by some reasonable means, this is the first time you have received notice of violation of this License (for any work) from that copyright holder, and you cure the violation prior to 30 days after your receipt of the notice.

Termination of your rights under this section does not terminate the licenses of parties who have received copies or rights from you under this License. If your rights have been terminated and not permanently reinstated, receipt of a copy of some or all of the same material does not give you any rights to use it.

10. FUTURE REVISIONS OF THIS LICENSE

The Free Software Foundation may publish new, revised versions of the GNU Free Documentation License from time to time. Such new versions will be similar in spirit to the present version, but may differ in detail to address new problems or concerns. See <http://www.gnu.org/copyleft/>.

Each version of the License is given a distinguishing version number. If the Document specifies that a particular numbered version of this License "or any later version" applies to it, you have the option of following the terms and conditions either of that specified version or of any later version that has been published (not as a draft) by the Free Software Foundation. If the Document does not specify a version number of this License, you may choose any version ever published (not as a draft) by the Free Software Foundation. If the Document specifies that a proxy can decide which future versions of this License can be used, that proxy's public statement of acceptance of a version permanently authorizes you to choose that version for the Document.

11. RELICENSING

"Massive Multiauthor Collaboration Site" (or "MMC Site") means any World Wide Web server that publishes copyrightable works and also provides prominent facilities for anybody to edit those works. A

License For This Documentation

public wiki that anybody can edit is an example of such a server. A "Massive Multiauthor Collaboration" (or "MMC") contained in the site means any set of copyrightable works thus published on the MMC site.

"CC-BY-SA" means the Creative Commons Attribution-Share Alike 3.0 license published by Creative Commons Corporation, a not-for-profit corporation with a principal place of business in San Francisco, California, as well as future copyleft versions of that license published by that same organization.

"Incorporate" means to publish or republish a Document, in whole or in part, as part of another Document.

An MMC is "eligible for relicensing" if it is licensed under this License, and if all works that were first published under this License somewhere other than this MMC, and subsequently incorporated in whole or in part into the MMC, (1) had no cover texts or invariant sections, and (2) were thus incorporated prior to November 1, 2008.

The operator of an MMC Site may republish an MMC contained in the site under CC-BY-SA on the same site at any time before August 1, 2009, provided the MMC is eligible for relicensing.

ADDENDUM: How to use this License for your documents

To use this License in a document you have written, include a copy of the License in the document and put the following copyright and license notices just after the title page:

Copyright (c) YEAR YOUR NAME.

Permission is granted to copy, distribute and/or modify this document under the terms of the GNU Free Documentation License, Version 1.3 or any later version published by the Free Software Foundation; with no Invariant Sections, no Front-Cover Texts, and no Back-Cover Texts. A copy of the license is included in the section entitled "GNU Free Documentation License".

If you have Invariant Sections, Front-Cover Texts and Back-Cover Texts, replace the "with...Texts." line with this:

with the Invariant Sections being LIST THEIR TITLES, with the Front-Cover Texts being LIST, and with the Back-Cover Texts being LIST.

If you have Invariant Sections without Cover Texts, or some other

License For This Documentation

combination of the three, merge those two alternatives to suit the situation.

If your document contains nontrivial examples of program code, we recommend releasing these examples in parallel under your choice of free software license, such as the GNU General Public License, to permit their use in free software.