Author:  Jonathan Yaeger
jon@datasaversllc.com

# DATA RECOVERY: BEST PRACTICES FOR SERVICE PROVIDERS

## OVERVIEW

Computer service centers encounter failed hard drives.  Data recovery can be a revenue source, but it must be wisely and carefully done. Well-intentioned but ill-informed efforts and practices can reduce or even ruin the prospects for successful recovery, and they can even expose the provider to legal liability.

As in medicine, the main precept of data recovery is “first, do no harm.”  The purpose of this article is to share basic data recovery practices intended to minimize the chances of harming a drive during the initial diagnostic and imaging (or copying) phases of the data recovery process.

The article will present general principles as well as specific examples.

## THE LEARNING CURVE

Gaining expertise in data recovery is neither a quick nor easy process.  There are many different drive makes and models; what works for one brand may not apply to another.  Models from the same manufacturer can be markedly different in construction and operation. A huge amount of technical information and details must be absorbed and mastered to become even modestly adept.

Hard drive manufacturers are frugal about releasing technical details of their products because they want to safeguard their intellectual property and trade secrets.

Good drive diagnostic and firmware tools are expensive.  Comprehensive training on how to use them is also generally expensive, when available. Most third-party data recovery training classes provide only a cursory introduction; there is no substitute for hours of hands-on experience.

Although much of what data recovery technicians learn and use comes from trial and error, **it is important not to get one’s education at the client’s expense**. We receive a lot of drives that have been worked on by other technicians. A fair number of those have been damaged or ruined – along with the chances of recovery – by mistake. Incorporating a set of best practices for hard drive data recovery can improve the outcomes and be beneficial to all concerned.

## GENERAL PRINCIPLES

1. **Recognize your own limitations**. *A customer's best interests should be top priority*. Resist the temptation to experiment or gamble with a client's data – you can make things much worse! If you don't have the knowledge, equipment, experience and confidence to tackle a particular recovery job, then simply stop. It is better to pass on a service opportunity than to ruin forever a client's chances of getting back data.

   Note that recovery pros sometimes sub out difficult jobs to others with specialized expertise, if it will ultimately benefit the client.

2. **Proper diagnosis is essential.** You can't fix the problem unless you are sure what it is. However, you can make things worse by randomly trying things without understanding the nature of the failure.

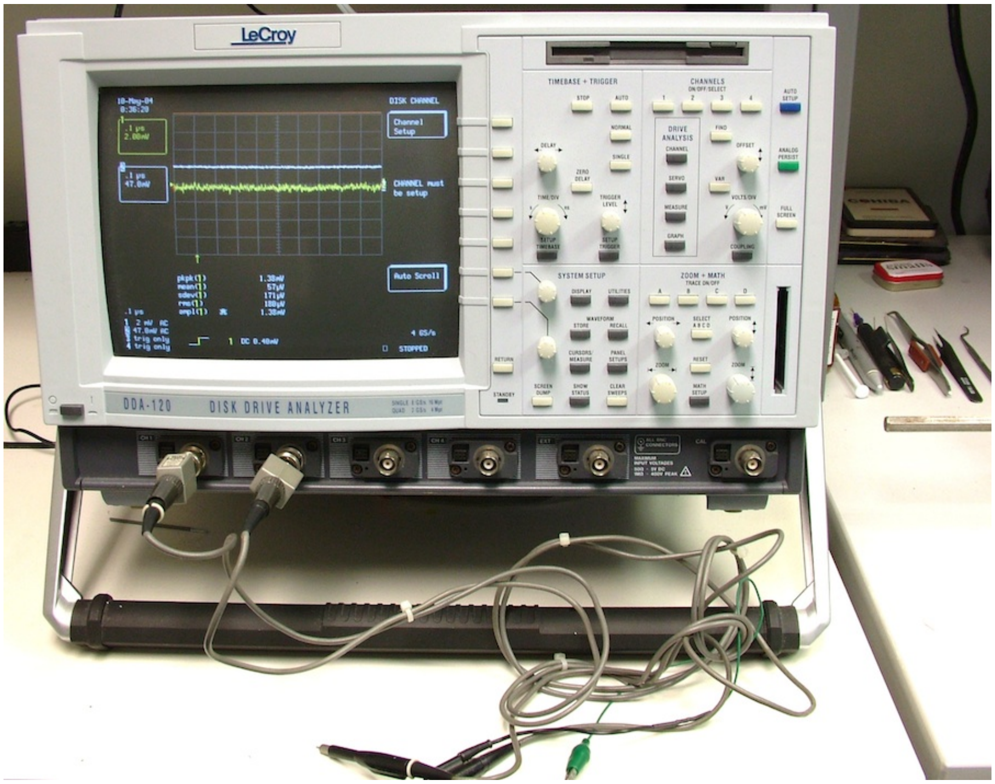

Fig. 1 **Data recovery requires the proper tools.**

3. **If you don’t have the proper tools, don’t try to improvise**. If you must open the sealed chamber of a hard drive, do it only in a certified clean environment, which isn't a kitchen, bathroom, or garage.

4. **Don’t do something that can’t be undone later on**. Keep a path open to the previous step so you can restore the drive to the same condition as you received it, if possible.

5. **Document and label everything**. Record every step that you take in a recovery effort. It will help you or the next technician later. If you take good notes, you can build a library of “how to” steps for the next drive with similar issues.

   Put a label on each item of clients' property, including cables, to make sure that they receive all of their property back and that it is not commingled with someone else’s. Label the drive’s PCB to be sure that the original is restored if swapped.

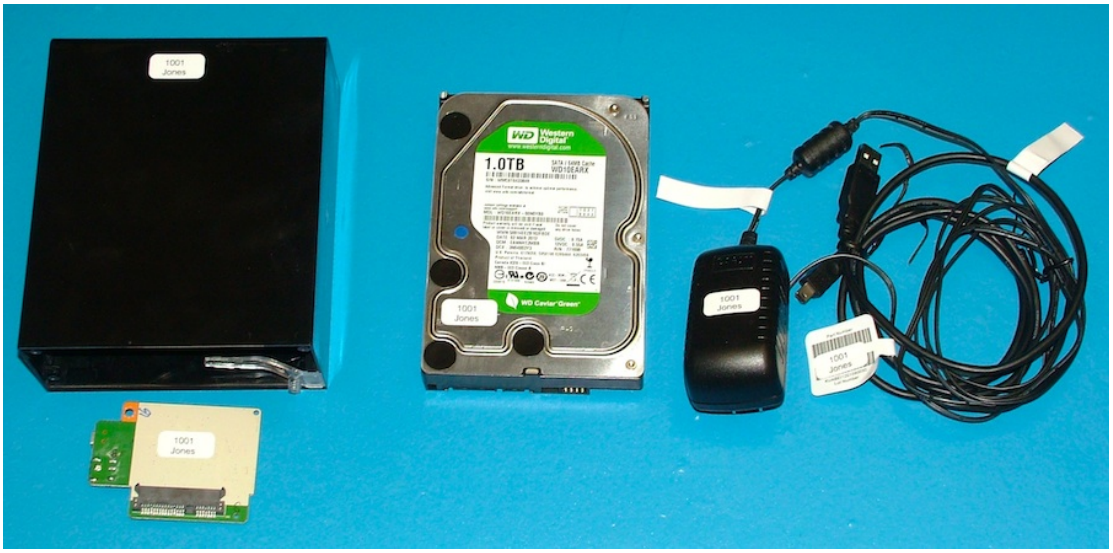

Fig. 2 **Label everything.**

6. **Try a new procedure on a test drive first**. Sometimes a new process or procedure will ruin a drive. Don’t let it be a client’s drive.

7. **You can't use a software tool to compensate for a hardware problem**. If a drive is failing, trying to copy the data without fixing the underlying hardware issue is a recipe for disaster; the drive may fail and become unrecoverable. Learn to recognize hardware failure patterns, and when to stop what you are doing.

8. **Look under the streetlight**. Use your five senses to examine a drive prior to working on it. A problem might be obvious if you take the time to look. A misdiagnosis not only wastes time, but also can put data at risk.

9. **Treat all data as confidential data**. The client's data is valuable. Do not expose it to security risk or loss. After a recovery is complete, securely wipe any media used in the process and confirm the erasure.

10. **Be honest and ethical**. Of all the "best practices," this is by far the most important. The "Golden Rule" certainly applies.

    Doctors have learned that the best policy with patients is to be candid and forthright, and so it is with data recovery.

## INITIAL EVALUATION: EXTERNAL APPEARANCES

The first step in the recovery process is to try to get a history from the client, i.e., what happened when the drive failed? Was it dropped? Was there a power outage? Clients may not have additional information, but when they do, it can be invaluable.

What do your senses tell you? Is the lid bent at a corner? Are there other signs of physical damage that suggest the drive had been dropped or abused?

The practice of "looking under the streetlight" means that you should take time to study a drive before doing anything. You may find obvious as well as hidden problems that will affect the diagnosis and dictate the course of the recovery. Overlooking a problem can result in wasted time and, sometimes, in loss of data.

## EXAMINING THE PC BOARD

Examine the PC board area, especially with laptop hard drives. Is there any sign of fluid contamination or corrosion on the board? Does the drive smell like coffee or wine? Is it sticky? If so, there is contamination and likely PCB damage.

A drive that has been under water should be opened in a clean room as well, to check for fluid contamination on the inside. All contamination must be removed before the recovery process can commence.

Are there burned areas on the circuit board?

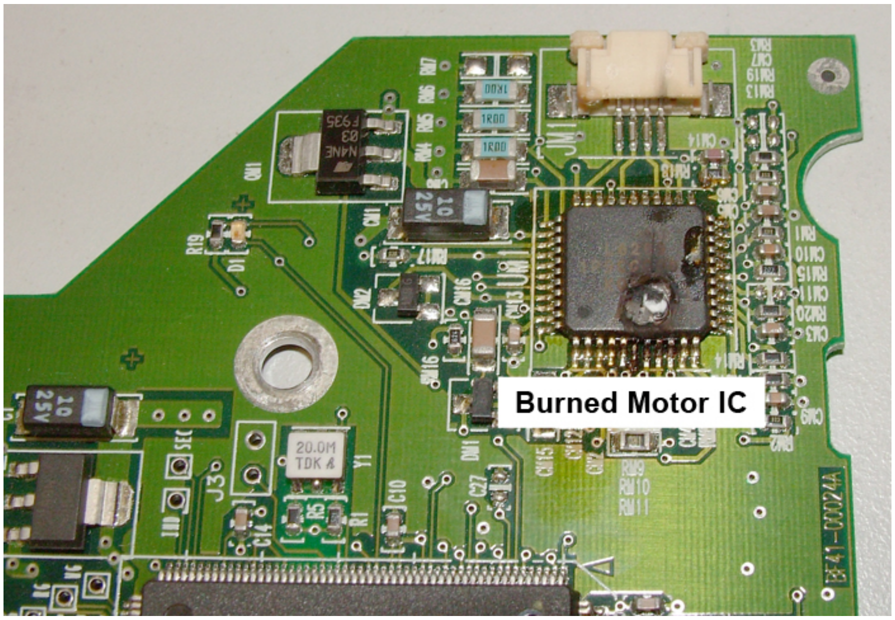

Fig. 3 A **burned motor integrated circuit.**

Figure 3 shows a burned motor chip on a PCB of a drive that – no surprise – did not spin up. Other PCB-related problems besides corrosion and contamination include damaged or defective parts.

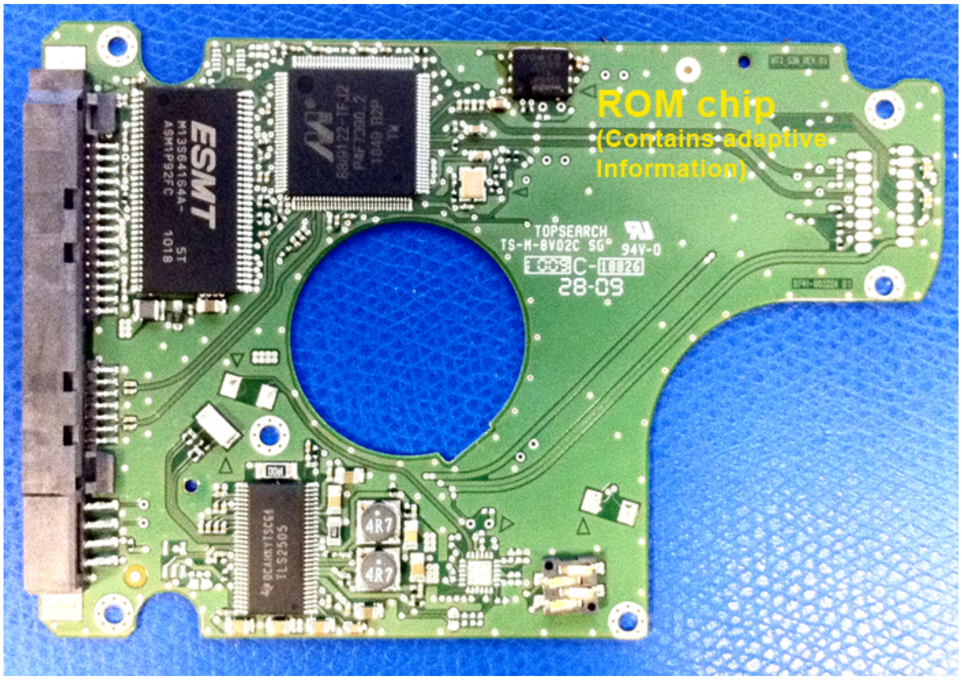

Fig. 4 A **PCB showing a ROM chip.**

Most hard drive circuit boards contain unique information about the particular drive, either as part of a **ROM** (usually an **EEPROM**, or **E**lectrically **E**rasable and **P**rogrammable **R**ead-**O**nly **M**emory circuit or “chip”), or embedded into the main processor.  This unique information is called *adaptive information* or *firmware*.

Figure 4 shows a laptop PCB and the placement of the ROM on the board.

> **RULE:** *You can’t just replace a damaged PCB without transferring the adaptive information to the replacement PCB, a process that requires special equipment.  The days of simply swapping PCBs are largely over.*

Sometimes service centers will swap a PCB and not tell the customer, or sometimes they discard the bad PCB.  This makes future recovery efforts much more difficult — and sometimes even impossible.

**Please note that recovery of Toshiba laptop drives and the newer Seagate drives in the F3 series may not be possible without their original PCB adaptive information.**

> **RULE:** *Best practices dictate returning the bad drive with the original PCB attached, even if damaged, if the data cannot be fully recovered.*

## BEFORE APPLYING POWER TO A DRIVE

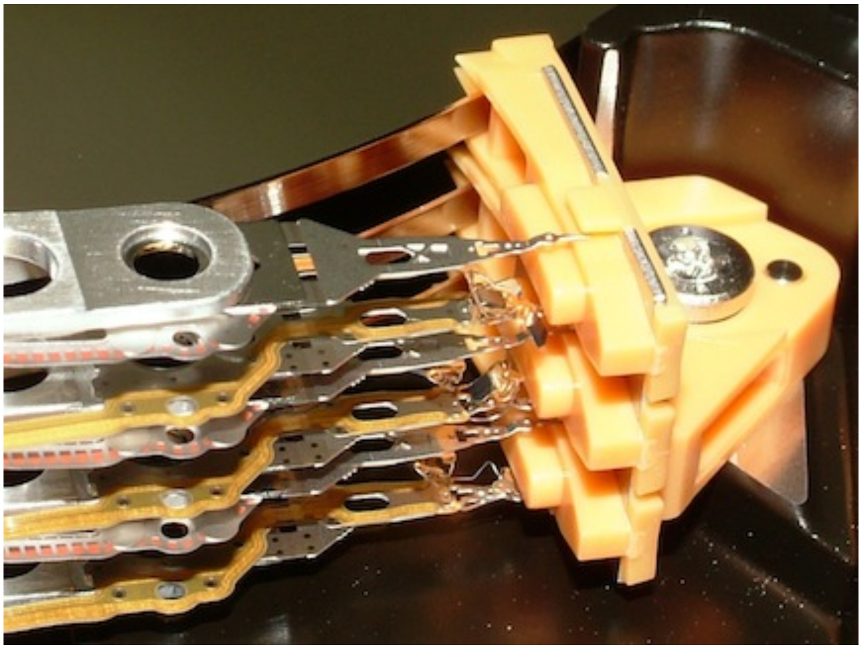

Fig. 5 **Damaged heads in a parking ramp.**

> **RULE:** *<u>Do not</u> apply power to a drive that you know or suspect has been dropped or subjected to shock! A drive that has been immersed in liquids or otherwise contaminated also shouldn't be powered. Read on to learn why.*

Hard drives are amazing collections of high-tech mechanics, electronics, and internal software programs (also called *firmware*), but they are also very delicate machines.

During operation, tiny read-write heads fly over the data platters. There is a very tiny gap between the heads and the platters. If a drive receives a physical shock, the heads may crash into the platters, resulting in damage to both the heads and the platters. Sometimes damage may occur to the delicate heads if the drive is dropped when powered off. Figure 5 shows a set of mangled heads sitting in the parking ramp.

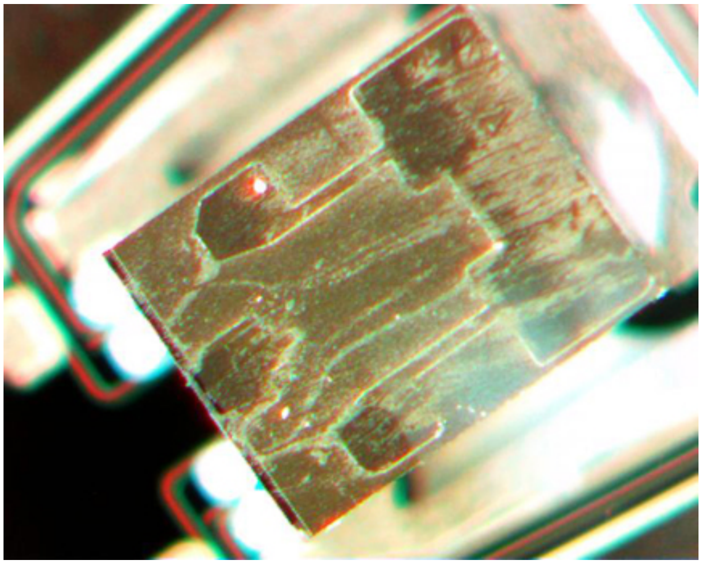

Fig. 6 **A scratched and contaminated head (magnified).**

Figure 6 shows a close-up of a head damaged by a head crash. The head has scratched the media, and some of the drive's magnetic media is now tiny particles of metallic dust. Damaged platter areas are unrecoverable.

As soon as the drive is turned on, the deformed heads will move onto the platter and act like sandpaper or tiny knives, scraping data away along with the chance of recovery.

The debris kicked up by the process acts like large boulders randomly dropped onto a highway, meaning more crashes will follow. Particles larger than 0.25 microns cannot pass through the air cushion gap between the head and the platter.

For perspective, the diameter of a human hair is on average about 80 microns, which is 320 times larger.

Figure 7 shows what the drive will look like if it is operated with crashed heads. In this particular case, the data is gone for good; the client's data is "in the filter."

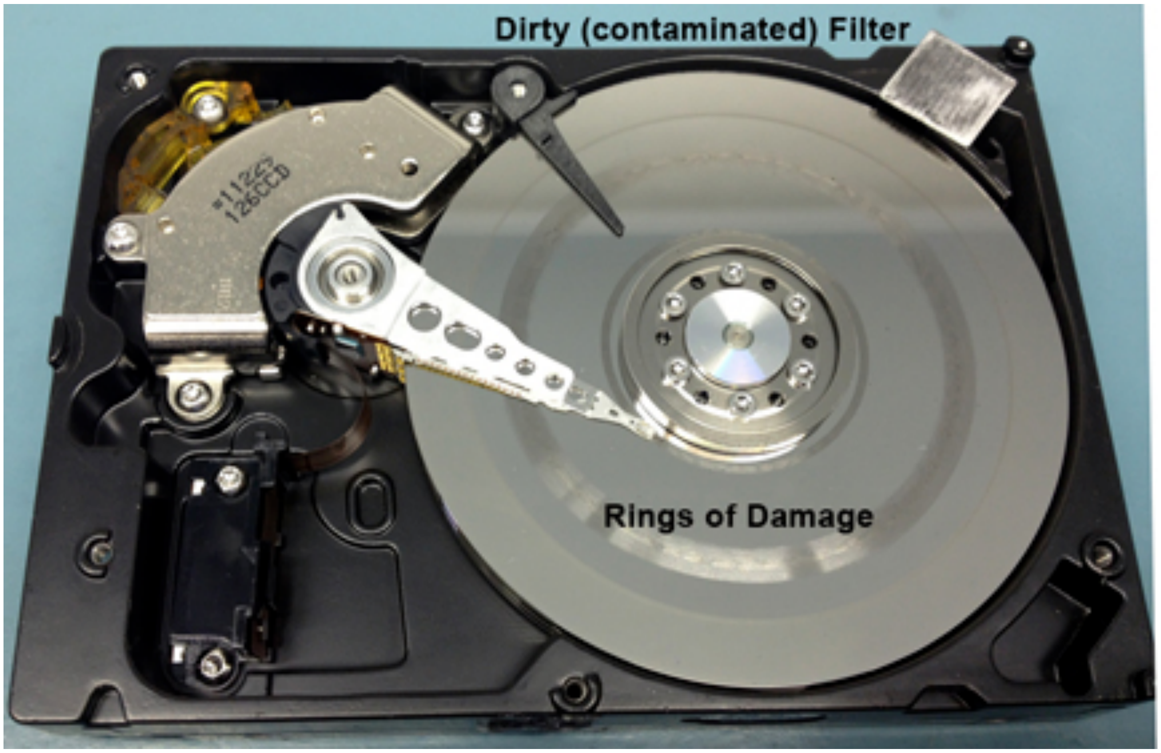

Fig. 7 **A hard drive with a head crash and media damage.**

Drives with slight media damage often can be recovered, though sometimes only partial recovery is possible; it depends on the extent of the damage. Working with media-damaged drives is a specialty, and some data recovery companies produce better results than others.

At this point, it should be clear that platter and head inspection are an important diagnostic part of the data recovery process, **but only in a true clean room environment.**

## OPENING THE SEALED CHAMBER OF A HARD DRIVE

**It is never a good idea to open up a hard drive outside of a clean room environment**. Why? Because, as previously noted, it only takes a small particle to cause a head crash, and if you open a drive outside of a clean room environment, you will contaminate the drive and increase the chances of premature failure or no recovery.

Opening a crashed drive in a clean room can contaminate the work area, too.

Sometimes removing the lid of a perfectly good drive can also contaminate the sealed chamber; debris often accumulates around the lid's seal.

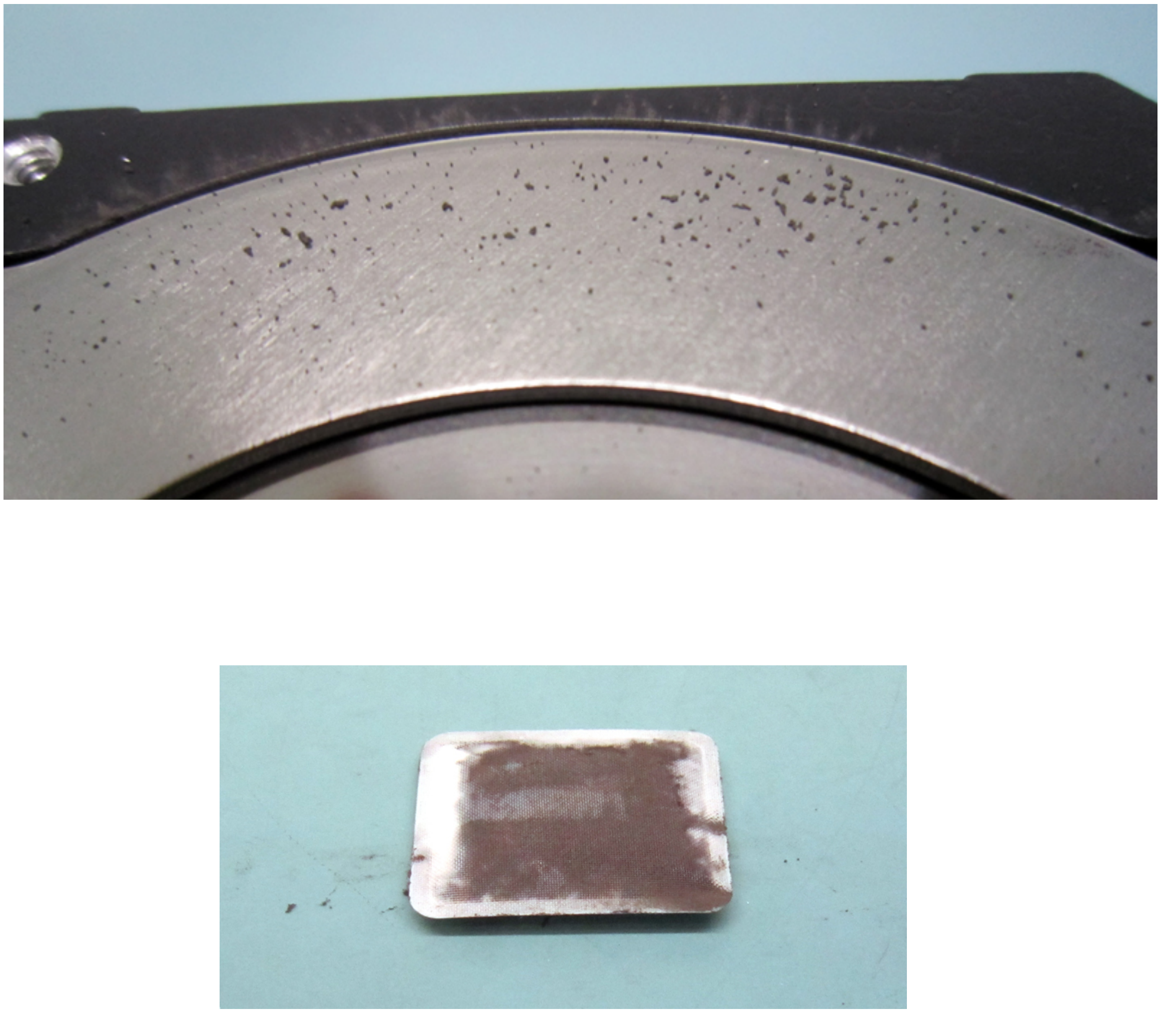

Fig. 8 **Close-up views of drive contamination.**

Another reason to not break the seal of a hard drive? If you are not authorized by the manufacturer to do so, you will *void the warranty*.

Resist the urge to "peek." Send the drive to a professional if clean room service is indicated. Many professional data recovery firms charge little or nothing for inspecting a drive prior to recovery.

## APPLYING POWER TO A DRIVE

The temptation is usually unavoidable, but a dropped drive should not be powered up without a clean room inspection of the heads and media.

Ask your client if the drive was dropped. Unless you have a clean room and the experience and expertise, it's a good idea to refer dropped drives to qualified data recovery labs.

You may safely assume that the client has attempted to access the drive until a point of defeat was acknowledged. <u>Sometimes it is just a matter of a few seconds</u> of operation of a mechanically failed drive that later results in a difficult or impossible recovery job.

Once that you are satisfied that a drive is safe to power up, listen carefully to it. Drives that make unusual noises, especially scraping and high-pitched sounds, should be shut off immediately. Unusual drive noises often indicate damage in progress. **Whenever there is uncertainty, always err on the side of caution**.

## IMAGING A DRIVE

A fundamental best practice of data recovery is to make a copy, or *image,* of the defective drive, then do the file recovery from the image. There are good reasons for this practice:

1. A failing hard drive may have a limited number of hours of operation left before it fails completely. It's best to try to get the data off a drive as quickly as possible, and with minimal damage to the original drive.

2. Sometimes more than one file recovery algorithm is required for optimal results. An image allows multiple recovery passes, whereas sometimes you get one shot with a failing drive.

3. The drive image is an archive of the original data and is often the only useful one.

It is important to understand how hard drives process read errors and perform other background tasks during normal operation.

Hard drives incorporate *housekeeping routines* into their firmware*.* These run in the background and are invisible to the end user. One standard routine is *defect management*, in which data in weak or failing sectors are moved to another region of the drive. The bad sectors are "marked out" and added to the *grown defect list*, also called the *G-List*.

Another housekeeping function is regularly updating the drive's *S.M.A.R.T. (Self-Monitoring, Analysis and Reporting Technology)* logs.

These logs keep track of *attributes* of a drive's performance; examples include the number of times the drive is started up, the inability to read sectors, and updates to error logs. These attributes are expressed in terms of "threshold exceeded" and "threshold not exceeded" values, which can provide an early warning of a drive's impending failure (Figure 9).

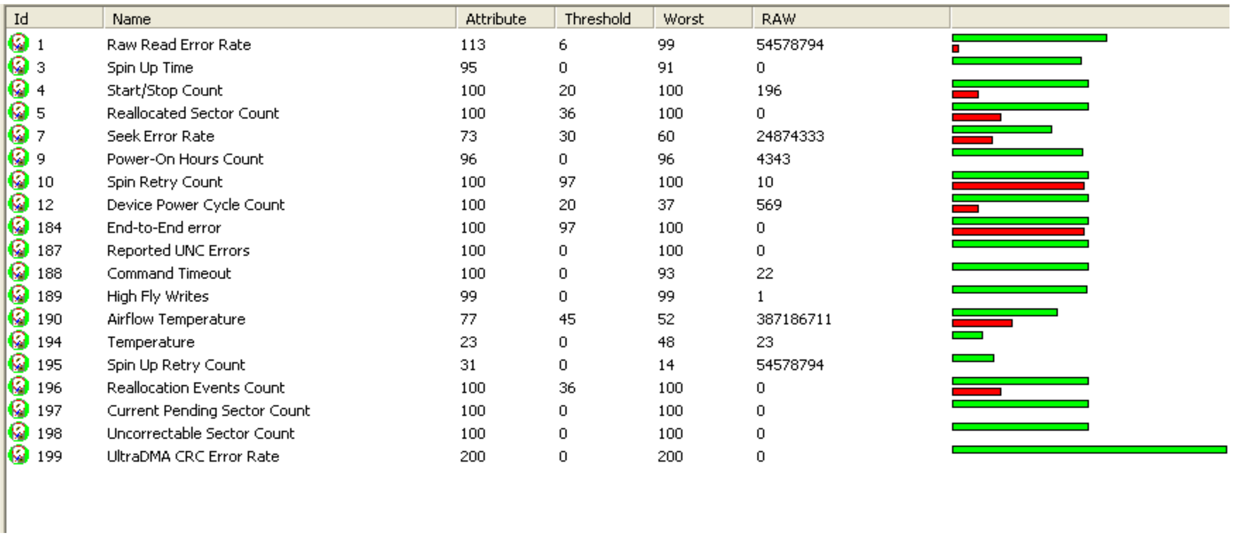

| Id | Name | Attribute | Threshold | Worst | RAW | |
|---|---|---|---|---|---|---|
| 1 | Raw Read Error Rate | 113 | 6 | 99 | 54578794 | 🟩<br>🟥 |
| 3 | Spin Up Time | 95 | 0 | 91 | 0 | 🟩 |
| 4 | Start/Stop Count | 100 | 20 | 100 | 196 | 🟩<br>🟥 |
| 5 | Reallocated Sector Count | 100 | 36 | 100 | 0 | 🟩<br>🟥 |
| 7 | Seek Error Rate | 73 | 30 | 60 | 24874333 | 🟩<br>🟥 |
| 9 | Power-On Hours Count | 96 | 0 | 96 | 4343 | 🟩 |
| 10 | Spin Retry Count | 100 | 97 | 100 | 10 | 🟩<br>🟥 |
| 12 | Device Power Cycle Count | 100 | 20 | 37 | 569 | 🟩<br>🟥 |
| 184 | End-to-End error | 100 | 97 | 100 | 0 | 🟩<br>🟥 |
| 187 | Reported UNC Errors | 100 | 0 | 100 | 0 | 🟩 |
| 188 | Command Timeout | 100 | 0 | 93 | 22 | 🟩 |
| 189 | High Fly Writes | 99 | 0 | 99 | 1 | 🟩 |
| 190 | Airflow Temperature | 77 | 45 | 52 | 387186711 | 🟩<br>🟥 |
| 194 | Temperature | 23 | 0 | 48 | 23 | 🟩 |
| 195 | Spin Up Retry Count | 31 | 0 | 14 | 54578794 | 🟩 |
| 196 | Reallocation Events Count | 100 | 36 | 100 | 0 | 🟩<br>🟥 |
| 197 | Current Pending Sector Count | 100 | 0 | 100 | 0 | 🟩 |
| 198 | Uncorrectable Sector Count | 100 | 0 | 100 | 0 | 🟩 |
| 199 | UltraDMA CRC Error Rate | 200 | 0 | 200 | 0 | 🟩 |

Fig. 9 **S.M.A.R.T. attributes display.**

Hard drives often fail in a cascading mode. For example, a drive's head might become dirty and fail to read or write data. Good sectors will be reported as bad, and the drive will try to move the data to another sector (i.e., through reallocation). If it cannot successfully write to the new location and verify the written data, it will report the new sector as bad, too. Each time it moves an allegedly bad sector, it will update both the G-List and S.M.A.R.T. logs. <u>The drive suffers a vicious cycle that interferes with imaging and eventually can lead to complete drive failure. Therefore, if the bad sector and S.M.A.R.T. management processes can be turned off during imaging, the chances of a successful and more complete image are greater.</u>

## SOFTWARE VERSUS "HARDWARE" IMAGING

Hard drives may be duplicated (or imaged) by computers using software programs for that purpose. Drives can also be copied by *hardware imaging*, using specially designed circuit boards or dedicated machines.

Some hardware imagers have features that greatly surpass software-only capabilities. They can:

- Disable defect management and S.M.A.R.T. attribute processing;
- Permit multiple recovery passes on a failing hard drive;
- Adjust the speed and "depth" of the reading process by adjusting the types of resets, number of attempts to read each bad sector, etc.;
- Build a heads map and image by head;
- Turn off imaging or power, jump to a different head, skip a pre-assigned number of sectors, etc. upon programmed conditions;
- Have dedicated imagers that make faster imaging (versus PCs running software);
- Support file imaging by MFT or other directory structures.

According to Serge Shirobokov of DeepSpar, Inc., *"Regarding the software vs hardware imaging section, the biggest difference between the two is probably the fact that hardware can automatically reset drives on timeout and software just has to wait . . . If a read command falls on a bad area, hardware tools can limit the damage from that by cutting the processing time down to a few hundred milliseconds (then resetting the drive), while software must wait the full few seconds for a drive to come back with an error. This alone makes hardware imaging multiple times less damaging."*

Because most drives presented for data recovery are unhealthy by definition, using a hardware-based imaging solution is generally considered a best practice.

When software imaging, it is a good idea to use a write blocker with software imaging tools, because a PC may try to initialize a drive and can overwrite important boot or partition information. Write blockers permit reading from a drive, but block attempts to write or change data. Most are hardware devices, but some use software. Note that a write blocker usually does not prevent a drive from updating the S.M.A.R.T. or the G-List.

Unix's "dd" command-line instruction is sometimes suitable for making a drive image because it can copy sectors that are not part of the active directory, including "erased" files.

On the other hand, DOS command-line instructions such as XCOPY are only good for copying files that are listed in a directory. If the directory is corrupted, or files have been erased, XCOPY will skip those files.

There are a number of "off the shelf" utility programs for disk imaging. Software imaging solutions work well with healthy hard drives. However, they are limited, and possibly harmful, when working with problem drives such as those with many bad sectors, firmware issues, or failing heads.

As noted, the best way to work with a defective drive is to turn off the drive's defect processing (called *sector reallocation*) and to disable the S.M.A.R.T. logging function. This reduces the chances that the drive will self-destruct during the imaging process.

However, turning off sector reallocation requires a vendor command, i.e., code specifically written for a particular drive, model, or series. Some of these commands are proprietary and are not generally published. A viable software solution would have to incorporate the vendor commands and have a means of matching the commands to the drive to be imaged.

For the tech-savvy: The BIOS of most PCs offers a menu choice to turn off S.M.A.R.T. notifications of impending drive failure, but that feature does not affect the internal S.M.A.R.T. operations of the hard drive.

A standard ATA command to disable a drive's S.M.A.R.T. operations is described in section 7.52 of the T13 ATA standard, which states: "This command disables all S.M.A.R.T. capabilities within the device including any and all timer and event count functions related exclusively to this feature."[6]

The command is *B0h with a Feature register value of D9h*

## WHEN TO STOP IMAGING

Remembering that a main precept of data recovery is "first, do no harm," it's critical to recognize when to stop, or when continuing is likely to do more harm than good.

Stop imaging when:

1. The drive makes clicking or other "unhealthy" noises;

2. The drive fails to come ready, or aborts;

3. The drive exhibits signs of cascading failure, i.e., one or more heads that were working when the process was started begin to fail.

A “gray” area is when there is a high unread or bad sector count. Sometimes it is hard to tell if this is due to media degradation (bad sectors), logic board failure, or another cause.

A head may spontaneously fail during the image process.  The drive may be able to finish imaging the data on the remaining heads, or it may “crater” and cause a ring of damage and widespread contamination.  At this point, having a clean room and the appropriate tools might save the recovery and a client’s data.

## RAIDS

A RAID originally stood for a “redundant array of inexpensive disks,” which later became “an independent array of independent disks.”  Nonetheless, a RAID is a storage system that has more than one drive. RAIDs have level classifications, such as RAID 0, RAID 1, RAID 5 or RAID 10. The Wikipedia entry for RAID gives a good overview: http://en.wikipedia.org/wiki/RAID.

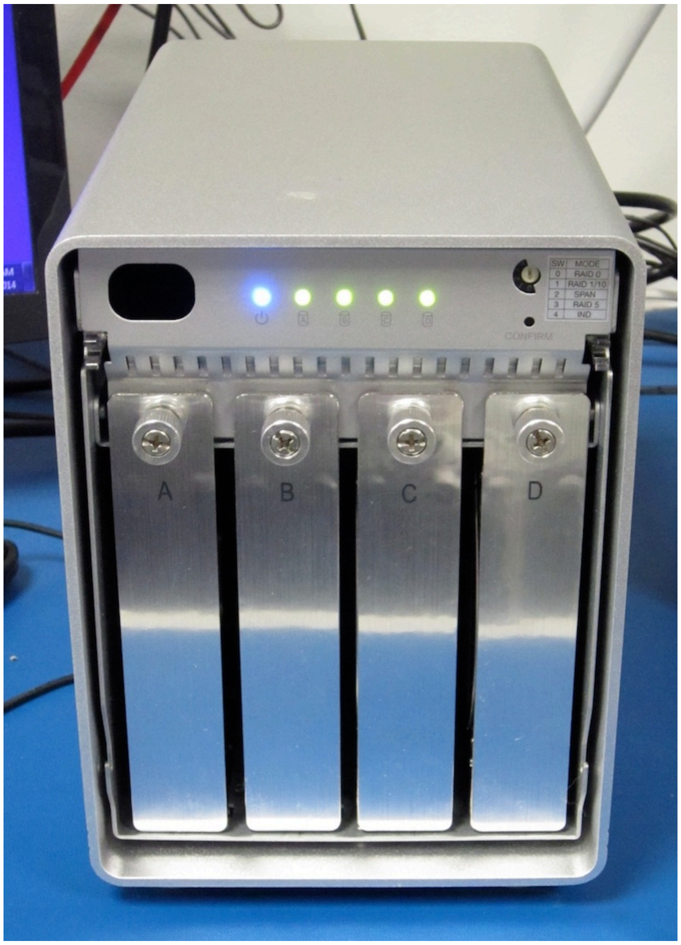

Fig. 10   **Configurable four-drive RAID array with front panel removed.**

ARAID 0 consisting of two drives isn’t redundant at all and, as such, is misnamed.  Both hard drives have to be recovered to get useful data from a RAID 0.

The redundancy feature of true RAIDs permits one or more drives of the array to fail and still retain functionality.  The RAID senses the failure and adjusts the reading and writing of data accordingly.

For example, a single drive from a three-drive RAID 5 can fail.  If a technician replaces the defective drive with another, the RAID “knows” – or prompts – the user to rebuild the RAID using the new drive.  However, if two drives fail in the RAID 5 configuration, data recovery is required; the RAID can’t rebuild from only one good drive.

The capability of automatic or semi-automatic RAID recovery after a drive failure is a great feature for the client, but if it's not completely understood by the technician, “experimenting” with a failed RAID array can compromise or even defeat subsequent data recovery attempts.  It’s easy to get into trouble if you don’t know what you are doing.

<u>Best practice rules for RAIDs</u>:

1. If you don’t know *precisely* what you are doing, <u>stop</u> what you are doing.

2. Drives should <u>never</u> be swapped around and placed in different positions in a RAID.

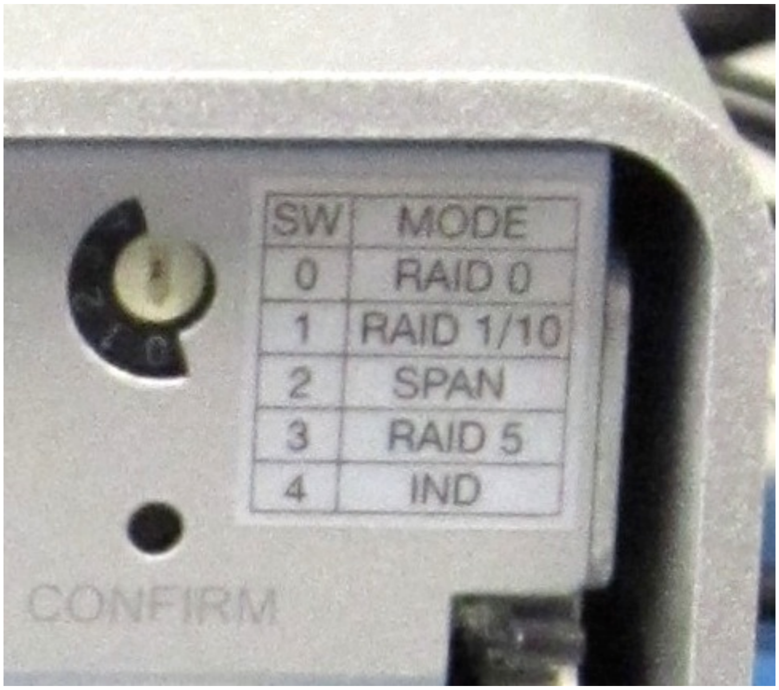

| SW | MODE |
|---|---|
| 0 | RAID 0 |
| 1 | RAID 1/10 |
| 2 | SPAN |
| 3 | RAID 5 |
| 4 | IND |

Fig. 11   A **RAID configurable switch.**

3. Figure 11 shows the configuration selector for the above RAID, which allows it to be set up in different RAID levels or modes. Never change the mode on the fly. This will cause severe data corruption.

4. The best practice when working with a problematic RAID is to make or clone copy of every drive, then work on the copy only. This will preserve the original RAID in "as received" condition, using the general principles of "do no harm" and "don't do anything you can't undo later."

## POST-IMAGE PROCESSING

Generally, data recovery software (to rebuild files and folders) is run on the image to try to recover lost files, rebuild directories, etc. Another reason for running recovery software, versus simply copying data, is that most programs fix permissions issues with the recovered data set. There are several data recovery programs, and each seems to have a particular set of strengths and weaknesses. There is not a single program that is consistently the "best."

## GOING FORWARD

Try to incorporate the General Principles (listed at the beginning) underlying best practices of data recovery.

Strive to always do what is in the best interest of the customer, whenever there is a choice, and know your limits.

Data recovery is an art as well as a science, but there is no substitute for knowledge acquired through experience.