Data Recovery

A Reference Manual — May 2026 Edition

By Mhessan Kouassi
Senior Data Recovery Expert at DAFOTEC

DAFOTEC Laboratory • French laboratory, ISO 5 cleanroom • Since 2004

Anatomy of storage media, professional methods, open source and commercial tools, prevention. All quantitative claims are sourced or explicitly labeled as orders of magnitude. No fabricated cases: the historical incidents cited are public and referenced; the DAFOTEC lab reports come from real anonymized cases published on dafotec.fr.

A reference document for technicians, forensics students, advanced individuals and IT decision-makers.

Preface

Why this manual exists — and what it reacts against

I have practiced data recovery since 2004. In the DAFOTEC laboratory in Roubaix, France, I have opened, diagnosed and imaged more than 120,000 storage devices — failed hard drives, silent SSDs, RAID arrays degraded by power surges, crushed smartphones, NAS units encrypted by ransomware. This practice, accumulated over 22 years, has taught me one simple thing: **in data recovery, the first wrong decision is almost always more costly than the failure itself**.

This manual was not written to promise recovery everywhere; it was written to avoid the mistakes that make recovery impossible. That distinction sounds modest, but it draws the outline of the entire discipline. When a client arrives at the laboratory with a clicking drive and has already run three software tools "just to see", we no longer have the same problem to solve as at the start of the incident. When a CIO restarts a RAID rebuild on a degraded array without imaging first, the gap between what was recoverable and what still is can sometimes be measured in years of archives.

The industry suffers from chronic misinformation. Consumer software promises to recover "100%" of files from any medium. YouTube tutorials recommend the freezer for hard drives and rice for waterlogged phones. Online comparisons announce success rates with decimal places and no published methodology. This manual is explicitly built against that literature: every figure is sourced or qualified as an order of magnitude, every method is documented with its limits, every promise is qualified by what it assumes.

The other reason for this manual is more practical. Modern media — SSDs with aggressive TRIM, Macs with Secure Enclave, NAS units encrypted by ransomware, SMR RAIDs — have made obsolete some of the reflexes inherited from the golden age of magnetic disks. On an HDD from the 2000s, you had days or even weeks to recover a deleted file. On a modern NVMe SSD in good health, the window sometimes measures in seconds after TRIM. The discipline has changed; the reflexes must follow.

I distill here four principles that summarize 22 years of practice:

- **Field experience shows that a well-intentioned wrong attempt destroys more data than a simple logical failure.**
- **In data recovery, the first ten minutes often decide the next ten years of a company's archives.**
- **At DAFOTEC, diagnosis always precedes the tool: we don't launch a scan to see, we first qualify the risk.**
- **The role of a serious professional is not to promise a miracle, but to say honestly what is still recoverable, what is not, and what must absolutely not be done next.**

You may not resell it, nor publish a modified version. This free distribution is a deliberate choice: we would rather have this manual circulate widely, because each avoided wrong first decision — for an individual, a small business, a hospital's IT department — is worth more than what selling a few hundred copies would bring.

Enjoy the reading. And good backups.

Mhessan Kouassi
Senior Data Recovery Expert • DAFOTEC, Roubaix
May 2026

Disclaimer and methodology

Three rules governed the writing of this manual:

- Every percentage is sourced or explicitly labeled as an order of magnitude when no public statistic justifies it.
- No case is fabricated. The historical incidents cited (NotPetya/Maersk 2017, Code Spaces 2014) are public and referenced. The "DAFOTEC lab reports" are real interventions, anonymized, published by DAFOTEC on dafotec.fr.
- The DAFOTEC proprietary methods mentioned in the text — VeriFiles, SEAD, HSA transplant, Spider Web, Reverse FTL, CPU Swap — are documented on the official website; this manual presents them in their technical context.

Who this book is for

IT technicians who have to manage data loss incidents; digital forensics students looking for a pedagogical overview; advanced individuals who want to understand what is happening under the hood before acting or sending their device to a laboratory; lawyers and judicial experts who need a technical view to evaluate the admissibility of digital evidence; executives and CIOs who want to evaluate their prevention posture.

Who it is not for

Anyone with an active emergency who doesn't have time to read. In that case, the rule is simple: **unplug the device, do not write anything to it, and either consult chapter 5 (diagnosis) to find your bearings, or contact a laboratory directly**. Diagnosis is free at DAFOTEC; you can therefore get a professional opinion before any financial decision.

Pricing note

All pricing mentioned in this book (DAFOTEC flat fees in euros) reflects the French market. Pricing varies significantly between countries, between laboratories within the same country, and between consumer and enterprise tiers. Use these figures as orders of magnitude, not as international references.

Legal notice

Recovering data from media that does not belong to you, or that contains data protected by professional secrecy or by privacy law (GDPR in Europe, HIPAA in the US, etc.), is regulated. The techniques described here are for educational and professional purposes only. For any serious matter — medical, legal, intellectual property — going through a laboratory under a non-disclosure agreement remains the only acceptable path. DAFOTEC operates under systematic NDA for professional media, in compliance with GDPR and ISO 27001.

Contents

INTRODUCTION

Chapter 0

Data recovery in 2026

A discipline at the crossroads of several trades

Data recovery is the set of techniques that allow lost or inaccessible information to be retrieved from a storage device. The discipline involves four families of skills: materials physics (magnetism for hard drives, floating-gate or charge-trap electronics for NAND memory), file system algorithms, reverse engineering of proprietary controllers, and forensic procedural rigor.

It is carefully distinguished from **backup restoration**, which falls under prevention. Recovery intervenes after the loss, on a device that has no usable copy. It is, by construction, an emergency discipline where one works with what remains.

Two worlds, never to be confused

Logical recovery: the device is physically intact and detected by the machine. The problem is in the software — corrupted file system, deleted partition, erased files, ransomware encryption. The raw data is almost always still there; one just needs to know how to read it.

Physical recovery: hardware failure. The device is no longer detected, or emits abnormal sounds, or its controller no longer responds. The intervention requires a specialized environment: cleanroom for mechanical drives, micro-soldering station for SSDs.

All recovery work begins with a diagnosis that decides between these two worlds. Choosing the wrong world costs data. Chapter 5 details this step.

The evolution since 1990

For three decades, recovery happened on magnetic hard drives. The principle was stable: as long as the platters were not physically rewritten, the data remained in place. Deletion, quick format, logical corruption: all of these affected only the allocation table, not the contents. It was the golden age of consumer software like Norton Utilities or, later, TestDisk.

The massive arrival of SSDs from 2010 onwards shifted the discipline. NAND memory does not behave like a magnetic disk: you cannot rewrite in place, and the controller must continuously consolidate and erase entire blocks. With the **TRIM** command (introduced in Windows 7, macOS 10.6.8, and the Linux kernel 2.6.33), deleting a file triggers a *physical* erasure of the relevant cells, often within seconds to minutes. On a modern healthy SSD, the logical recovery window shrinks dramatically.

In parallel, two other developments reshaped the landscape: the spread of **hardware encryption** (SED, TCG Opal, BitLocker with TPM, FileVault with Secure Enclave) which renders the physical data unusable without the key, and the explosion of **ransomware** attacks that now explicitly target backups. Verizon DBIR 2025 quantifies this latter trend: ransomware was involved in 44% of documented data breaches in 2024, up 37% year over year.

On statistics — No public consolidated statistics give the global success rate of data recovery. Professional laboratories sometimes publish figures in their marketing materials. DAFOTEC publishes its own publicly on dafotec.fr — by failure type, on the 120,000+ cases handled since 2004: 95% on logical HDD failures, 88% on electronic HDD failures, 78% on mechanical HDD failures, 82% on SSD firmware failures, 61% on SSD NAND failures, 91% on degraded RAID 5, 69% on smartphones. These figures apply only to DAFOTEC, on a perimeter of cases that were not aggravated by previous attempts. They reflect historical performance, not a guarantee on any individual case.

How this book is organized

Seven parts, nineteen chapters, five appendices. The progression is deliberate: start by understanding what a storage device physically is (part I), then how to diagnose a failure (part II), then which methods to apply (part III), then how to handle special cases (part IV), then which tools to choose in practice (part V), then how to make all of this unnecessary through good prevention (part VI), and finally where the discipline is heading in the coming years (part VII).

Four layout conventions:

- **Blue** boxes are pedagogical notes.
- **Orange** boxes are operational warnings — to read before acting.
- **Green** boxes are publicly sourced case studies (historical incidents publicly documented).
- **Beige** boxes marked "LAB REPORT — DAFOTEC" are real interventions, anonymized, published by DAFOTEC on its official website.

Part I

Physical foundations

Before any method, one must understand *physically* how data is written, read, and deleted. Three chapters: a mechanical hard drive (HDD), a NAND memory (SSD, eMMC, UFS, SD card), and a file system (the logical layer that organizes blocks into a tree). Without this foundation, the methods of the following chapters are just magic.

PART I — PHYSICAL FOUNDATIONS

Chapter 1

Anatomy of a mechanical hard drive

1.1 Overview

A hard disk drive (HDD) is a miniaturized precision mechanism. Under its hermetic enclosure are: one or more rigid **platters** coated with a magnetic layer, a **spindle motor** that rotates them at constant speed, **read/write heads** carried by an **actuator arm** that sweeps the surface, and an external **printed circuit board** (PCB) containing the controller, the firmware ROM, and the SATA or SAS interface.

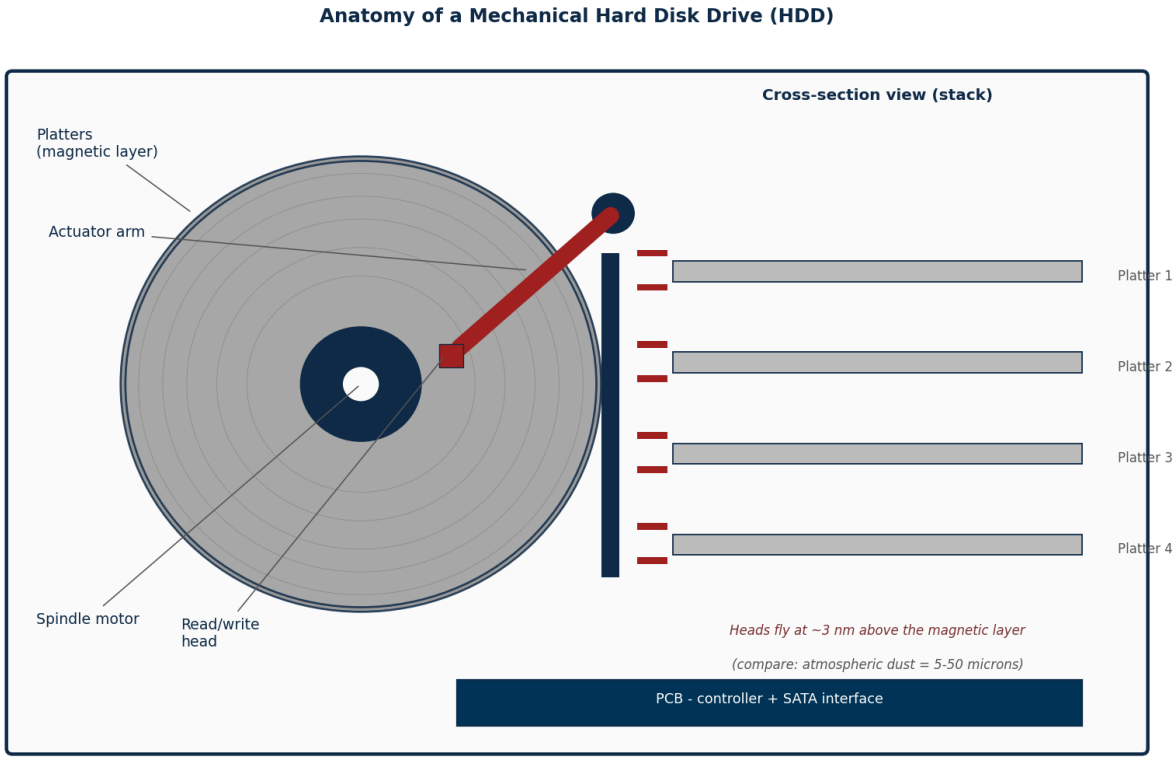

Schematic anatomy of a 3.5" HDD: platters, arm, heads in nanometric flight, PCB.

Typical orders of magnitude for a 3.5" consumer HDD in 2026:

Parameter	Typical value
Rotational speed	5,400 or 7,200 RPM (consumer), up to 15,000 (enterprise SAS)
Capacity per platter	2 to 4 TB
Number of platters	1 to 10 depending on total capacity
Linear density	Over 1 million bits per track inch
Head fly height	A few nanometers above the platter
Sequential throughput	150 to 300 MB/s

Parameter	Typical value
Random access latency	5 to 15 ms

1.2 How data is written

The magnetic layer of the platter is divided into billions of small magnetic domains. The write head, by generating a very strong and very brief local magnetic field, orients the polarization of a domain in one direction (bit 1) or the other (bit 0). The transition between two opposite polarizations is what the read head later detects, by induction or more recently by tunnel magnetoresistive effect (TMR).

As long as nothing overwrites the area, these magnetic polarizations are **stable over decades**. This is the fundamental property that makes HDDs so recoverable: erasing a file at the file-system level does not touch the magnetic domains themselves.

1.3 PMR, CMR, SMR, HAMR: the writing technologies

Four technologies coexist or succeed each other:

- **PMR** (Perpendicular Magnetic Recording): since 2005, magnetic domains are oriented perpendicular to the surface rather than parallel. This is the basis of all modern HDDs.
- **CMR** (Conventional Magnetic Recording): modern term that designates "classic" PMR with non-overlapping tracks. Allows any track to be rewritten without touching neighbors.
- **SMR** (Shingled Magnetic Recording): since 2013, tracks partially overlap like roof shingles. Gain of 20 to 25% in density, but each modification of a track requires rewriting the entire band of adjacent tracks. Firmware must manage a cache zone and garbage collection comparable to an SSD.
- **HAMR** (Heat-Assisted Magnetic Recording): technology that briefly heats the domain by laser during writing to reduce the size of stable domains. Commercialized from 2024 on very high-capacity enterprise drives (30 TB+).

SMR and recovery — SMR is significantly more complex to recover in case of firmware failure: the translation between logical addresses (LBA) and physical locations on the band is managed by the controller, and corrupted firmware can render content unreadable even on intact platters. Laboratories use specialized modules (PC-3000 has published SMR modules since 2020) that must manage both the main translator and the cache translator. Source: Rossmann Group, *CMR vs SMR: How Recording Technology Affects Recovery*, 2026.

1.4 Typical mechanical failure causes

1. **Head crash.** A head comes into contact with the platter surface — shock, vibration, fly-height defect. The result is often a progressive scratch that physically destroys the magnetic layer. Classic symptom: *repeated clicks* of the head trying in vain to position itself.
2. **Stiction.** Heads remain stuck to the platter instead of parking properly at shutdown. The motor can no longer start rotation. Symptom: *short buzzing noise then silence*.
3. **Spindle motor failure.** Bearings wear out, motor seizes. Symptom: *platter no longer turning at all, or turning in jerks*.

4. **PCB burnout.** A power surge destroys components of the printed circuit, often the TVS (protection transistor). The drive is no longer detected at all; sometimes it literally smokes.
5. **Firmware corruption.** The PCB ROM or a system area of the platters (service area, detailed below) becomes unreadable. The drive spins but does not mount, or mounts with an absurd capacity (0 GB, 8 MB, inconsistent value).

> **Caution —** A clicking HDD must be powered off immediately. Each additional platter rotation while the heads touch the surface extends the scratched area — every second, data is lost. This is one of the rare true emergencies in recovery.

1.5 Why an HDD remains highly recoverable

When a file system deletes a file, it modifies only its own internal tables (MFT for NTFS, inodes for ext4, FAT table for FAT32/exFAT). The physical sectors that contained the file are neither erased nor demagnetized. They will be only when a new file comes to write over them.

This is why on an HDD:

- A deleted file is recoverable as long as its sectors have not been reused.
- A quick format only reinitializes the basic structures of the FS; the sectors remain intact.
- A full format (which rewrites everything) does effectively destroy the data, but takes hours and is almost never done by accident.

On an unoverwritten HDD, logical recovery tools (TestDisk, R-Studio, UFS Explorer, PhotoRec as a last resort) retrieve in practice the overwhelming majority of data.

1.6 The average reliability of an HDD in 2026

Backblaze, a cloud hosting provider, has been publishing failure statistics for its hard drive fleet since 2013. The 2025 annual report (published in February 2026) counted 344,196 drives spread across 30 models. Three key figures:

- Annual AFR 2025: **1.36%** (down from 1.55% in 2024).
- Lifetime AFR across drive lifespan: **1.30%**, stable quarter over quarter.
- Q4 2025 showed a quarterly AFR of 1.13%, the lowest since 2022.

In other words: on a massive sample, about 1.4% of drives fail each year. For an individual with a single drive, this doesn't say much individually — your personal drive will fail or not, it's binary. But it reminds us that across a fleet, failure is statistically certain.

Source: Backblaze, Drive Stats for 2025, annual report published February 12, 2026.

1.7 Service area, firmware modules, translator

A crucial part of the intelligence of an HDD is found neither on the PCB nor in the controller, but in a particular area of the platters themselves: the **Service Area** (SA), invisible to the operating system. It contains over a hundred firmware modules indispensable to the operation of the drive. Three families deserve to be named:

- **P-List** (Primary defect list, module 0A on WD ROYL). List of defective sectors identified *at the factory*, at the end of the manufacturing line. These sectors are remapped from the start and never

accessible via LBAs.

- **G-List** (Grown defect list, module 0B). Sectors that became defective *during the life* of the drive and that firmware automatically reallocated to reserve sectors.
- **Translator** (module 028 on WD ROYL, equivalents on Seagate, Toshiba, etc.). Central module that translates the logical addresses (LBA) exposed to the system into physical addresses on the platters (cylinder, head, sector). Its corruption renders a perfectly mechanically healthy drive *completely unreadable* by the OS.
- **Adaptives** (modules 102 to 109 on WD). Head calibration parameters, specific to each drive specimen, adjusted at manufacturing. This is what makes a simple PCB swap impossible: these parameters must be transferred.

On an SMR drive, the translator is even more critical: it must also manage the mapping between CMR cache zones and shingled bands. A sudden power outage during a garbage collection operation can corrupt this mapping and render terabytes of data unaddressable. Recovery then consists of **reconstructing or repairing the translator** using specialized tools like PC-3000 — in practice, this is done daily in equipped laboratories. Sources: ACE Lab (PC-3000 public documentation), Rossmann Group, ISA Group *HDD Service Area Modules Reference*.

Diagnosing a SA failure — Typical symptom of a Service Area corruption: the drive is detected by the BIOS but displays an absurd capacity (0 GB, 8 MB, plausible but wrong value), or is detected under a manufacturer generic name. The drive spins normally, emits no abnormal noise — it is purely logical on the internal firmware side. Diagnosis reserved for the laboratory: the TTL serial terminal on the PCB is necessary to access the SA.

LAB REPORT — DAFOTEC • WD Blue 2 TB hard drive, heads failed (case #1)

Hardware: WD Blue 2 TB 3.5" • **Turnaround:** 72 hours • **Flat fee:** 650 EUR (French market)

Symptom. Repeated clicks every 3 seconds, drive not detected, after an 80 cm fall onto tile flooring of the external enclosure.

Intervention. Opening in ISO 5 cleanroom, replacement of the head stack assembly (HSA) by an identical donor drive of matching firmware revision, with sub-micron alignment under 0.3 microns. Sector-by-sector cloning with reduced timeout to preserve the new heads/platters combination.

Outcome. 1.94 TB recovered out of 2 TB. 3 partially corrupted 4K videos could not be fully restored due to a slightly scratched platter area caused by the fall.

Professional photographer, Paris area - 4 years of archives saved. Case published on dafotec.fr.

Part II

Diagnosis

Before acting, understand. Two chapters: a quantified overview of the causes of data loss in 2025-2026, and a triage method to decide whether you are facing a logical or physical case, and whether you can intervene yourself or if laboratory shipping is required.

PART I — PHYSICAL FOUNDATIONS

Chapter 2

Anatomy of an SSD: NAND, controller, FTL

2.1 A paradigm shift

An SSD (Solid State Drive) has no mechanical parts. All the complexity is in the electronics. This sounds like a simplification, but for recovery it is the opposite: NAND memory imposes physical constraints that force the controller to actively erase deleted data. On an HDD, deleting destroys nothing; on a modern SSD, deleting *destroys* within the following minutes.

2.2 NAND cells: floating-gate and charge-trap

The fundamental unit is the NAND cell. Historically, it was a **floating-gate** transistor (floating-gate MOSFET): an electric charge trapped between two insulating layers modifies the threshold voltage of the transistor. By measuring this voltage, the stored value is read.

Since the 3D NAND generations (2013 then massively from 2017), the technology has shifted to **charge-trap flash** (CTF). Instead of a conductive gate, an insulator that traps electrons is used. Advantages: better resistance to wear, simpler 3D manufacturing, fewer leaks between neighboring cells. Almost all consumer SSDs in 2026 (176 to 232-layer BiCS6, Samsung V8, etc.) use charge-trap.

Depending on the number of voltage levels distinguished in a single cell, more or fewer bits are stored. This is the core trade-off between density, endurance and reliability:

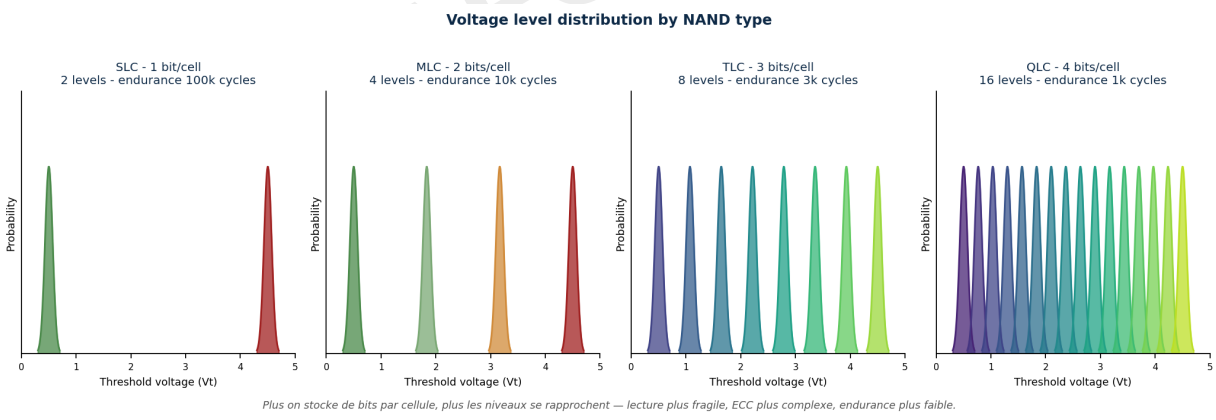

The more bits stored per cell, the closer the levels (SLC: 2 well-separated levels; QLC: 16 levels separated by less than 200 mV).

Type	Bits/cell	Levels	Endurance (P/E cycles)	Usage
SLC	1	2	50,000 to 100,000	Industrial, critical enterprise
MLC	2	4	3,000 to 10,000	Enterprise (declining)
TLC	3	8	1,000 to 3,000	Common consumer SSD 2026
QLC	4	16	150 to 1,000	High capacity budget
PLC	5	32	less than 150 (est.)	In development, rarely shipped

Sources: Kingston Technology, Lexar Enterprise, TechTarget, OSCOO. Endurance orders of magnitude vary by exact model; these values represent the high end of the typical range in 2024-2026.

The more bits stored per cell, the narrower the margin between levels (in QLC, less than 200 mV between two adjacent states), the more frequent the read errors, the more the controller must apply error-correcting codes (ECC), and the faster the cell wears with each write/erase cycle.

2.3 ECC, LDPC and adaptive read-retry

As margins have tightened, error-correcting codes have evolved. The first MLC SSDs used **BCH** codes (Bose-Chaudhuri-Hocquenghem), simple and deterministic. Modern TLC/QLC SSDs almost all use **LDPC** codes (Low-Density Parity-Check), more powerful but more complex: the decoder operates over several probabilistic iterations.

When the raw read produces too many errors for ECC alone to correct (heavily worn cells, disturbance from neighboring reads, charge loss over time), the controller triggers the **read-retry**: it changes the voltage reference threshold and re-reads the same page, sometimes five, ten or fifteen times with different thresholds. Combined with *soft* LDPC decoding (which uses probabilistic rather than binary information), these retries recover data that the initial read would have abandoned.

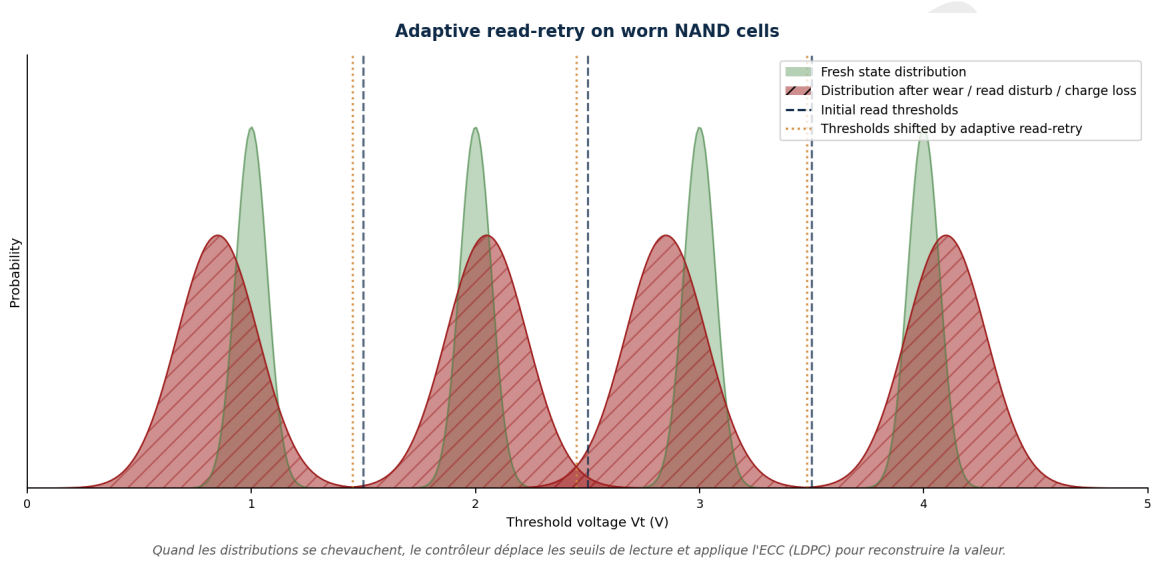

Read-retry shifts thresholds to distinguish distributions that overlap after wear or read disturb.

ECC and chip-off — For chip-off recovery (chapter 10), all this constitutes a major difficulty: reading raw NAND outside its original controller means obtaining data *after scrambling and before ECC decoding*. One must know how to reverse engineer both the scrambler and the ECC pipeline of the original controller. This is one of the reasons serious laboratories (DAFOTEC, ACE Lab clients) maintain profile databases by controller and firmware.

2.4 The fundamental constraint: write at page, erase at block

This is **the** detail that explains everything else. NAND memory can be read and written at the **page** level (typically 4, 8 or 16 KB), but it can only be erased at the **block** level (256 to 512 pages, typically 1 to 8 MB). Rewriting in place is impossible: the entire block must be erased first.

To remain performant, the controller never does this naively. When a file is modified, it:

1. Writes the new version to a free page, elsewhere on NAND.
2. Updates its internal mapping table (**FTL — Flash Translation Layer**) so the logical LBA now points to the new page.
3. Marks the old page as "invalid" but does not erase it immediately.
4. Later, in the background, the **garbage collector** consolidates the remaining valid pages of partially used blocks and applies the erase voltage on empty blocks.

2.5 Wear leveling, over-provisioning, hardware encryption

Since each cell has limited endurance, the controller applies **wear leveling**: it spreads writes across all available cells, so none wears faster than others. An **over-provisioning** zone (7% minimum, often 14% to 28% on enterprise SSDs) is invisible to the user but usable by the controller for reorganization operations and to replace cells that eventually die.

Often-ignored critical element: on the majority of modern SSDs (Phison E18, Silicon Motion SM2264, Samsung Pablo, etc.), **all NAND content is hardware-encrypted with AES-256**, even when the user has set no password. The master key is derived from a unique controller identifier (UID) and stored in a protected area. When the controller dies, both the FTL table and the key are lost. Chip-off then yields encrypted data whose key is permanently lost.

2.6 TRIM: the command that changes everything

Without TRIM, the controller has no idea which pages are still in use at the file-system level. When the OS deletes a file, it modifies its own tables but does not tell the SSD. The controller therefore considers these pages as containing valid data, and its garbage collector wastes time unnecessarily moving them around.

TRIM (ATA DATA SET MANAGEMENT command with the Trim attribute, or its NVMe equivalent DEALLOCATE) solves this problem by informing the controller of LBAs freed on the OS side. The controller can then:

- Update its mapping table immediately.
- Schedule these blocks for physical erasure at the next garbage collection cycle.

On most modern SSDs implementing DRAT (*Deterministic Read After Trim*) or DZAT (*Deterministic Zero After Trim*), any subsequent read of trimmed LBAs returns either a deterministic unspecified value or zeros. Logical recovery tools see nothing useful anymore.

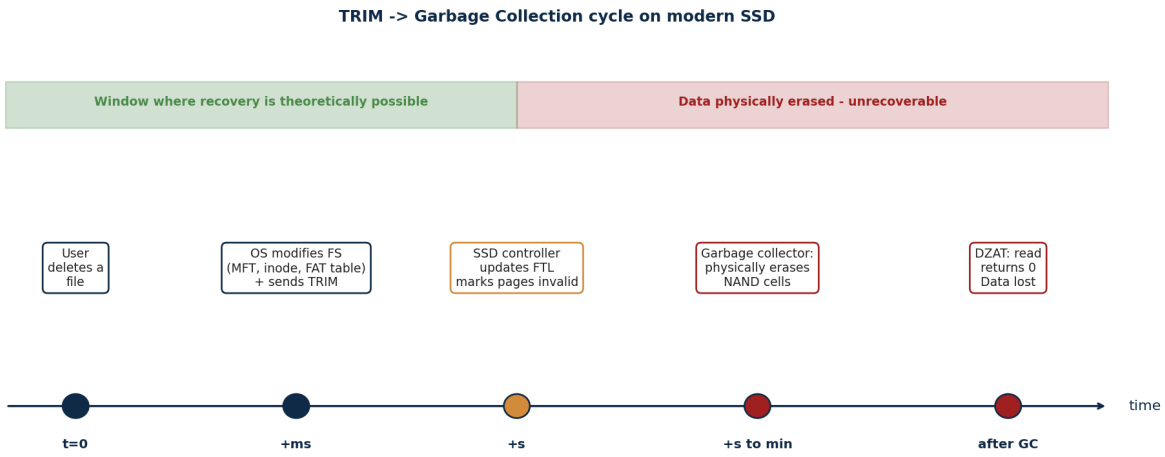

TRIM then garbage collection cycle: the recovery window closes within seconds to minutes.

2.7 The recovery window

How long before data is physically lost? On a modern NVMe SSD with TRIM active:

- On deletion of a file from an internal SSD mounted as NTFS/APFS/ext4 under recent Windows/macOS/Linux: TRIM is sent immediately (milliseconds).
- The garbage collector can activate at the next idle period, so within seconds to minutes at most.
- Once the block is physically erased, no chip-off recovers anything. The cells are in their neutral state.

> **Caution —** If you have deleted an important file from an SSD: **unplug the device immediately**, do not plug it back into the same machine, and send it for analysis. Every second under power reduces the chances. And even acting quickly, consider loss to be probable — not certain, but probable.

2.8 When TRIM does not work

SSD recovery remains possible in several configurations where TRIM is short-circuited, documented notably by Belkasoft (*Recovering Evidence from SSD Drives*, Forensic Focus). The main ones:

- **Hardware RAID.** Most RAID controllers do not pass TRIM to the underlying disks.
- **External SSDs via old USB-SATA bridges.** Older JMicron JMS539 and ASMedia ASM1051 do not pass TRIM. Recent UASP bridges (JMS578, ASM235CM, RTL9210B) do.
- **NAS.** Depending on firmware and volume configuration, TRIM may be absent or delayed.
- **Low-end pseudo-SSDs.** Some USB sticks and SD cards marketed as SSDs do not implement TRIM.
- **Buggy firmware.** Several models (notably some Crucial M4, OCZ Vertex, Intel 320) shipped with broken TRIM.
- **Small files stored internally in NTFS MFT.** Files of about 700 bytes or less are stored directly in the MFT entry (resident $DATA attribute) and are never affected by TRIM.

- **Fragments in still partially used blocks.** As long as a NAND block contains at least one valid page, it cannot be erased entirely — the invalid pages of the same block survive until the next GC.

2.9 How to check if TRIM is active

```
Windows: fsutil behavior query DisableDeleteNotify
(DisableDeleteNotify=0 -> TRIM enabled,
DisableDeleteNotify=1 -> TRIM disabled)

Linux: cat /sys/block/sdX/queue/discard_max_bytes
(0 = no TRIM support,
non-zero = TRIM available)

To see if fstrim runs automatically:
systemctl status fstrim.timer

macOS: system_profiler SPSerialATADataType | grep -i 'TRIM Support'
```


PART I — PHYSICAL FOUNDATIONS

Chapter 3

File systems: the treasure map

3.1 Why the file system decides

The file system (FS) is the software layer that transforms a raw block device (a sequence of sectors) into a tree of named files and folders. It maintains **internal structures** that link file names, metadata (size, dates, permissions) and physical location of data on the device.

When a file is deleted, the FS typically modifies two or three of these internal structures. The content of the file itself is not touched. It is this asymmetry that makes logical recovery possible: **the raw data survives the deletion of its entry in the "map"**.

But each FS manages this "map" differently. Some keep many traces (NTFS, with its $LogFile journal, is very verbose), others few (ext4 releases inodes and extents quite aggressively). This translates into significantly different recovery chances.

3.2 FAT32 and exFAT: simplicity

These are the simplest FS still massively used, mainly on removable media (USB sticks, SD cards, cameras, dashcams). FAT32 is limited to 4 GB per file; exFAT (2006, Microsoft) lifts this limit and has become the cross-platform standard for mass storage.

Structure: a **boot sector** at the start, one or two **FAT tables** describing the cluster chain of each file, and the rest of the volume as data area. Each file in a directory is described by a 32-byte entry containing short name, attributes and first cluster.

On deletion:

- The first character of the directory entry is replaced by 0xE5, marking the entry as deleted.
- The clusters in the FAT table are marked free (reset to zero).
- The content of the clusters themselves remains intact until rewritten.

Recovery is generally very effective on FAT/exFAT, especially for contiguous files (low fragmentation, common on SD cards used in sequential mode). Main limit: the first letter of the name is lost. Tools often put a generic character (_) in its place.

3.3 NTFS: the forensic richness

NTFS (*New Technology File System*, Microsoft, 1993) is the standard Windows FS since Windows NT. It is technically the most generous FS in residual metadata, which makes it the most recoverable from a logical point of view.

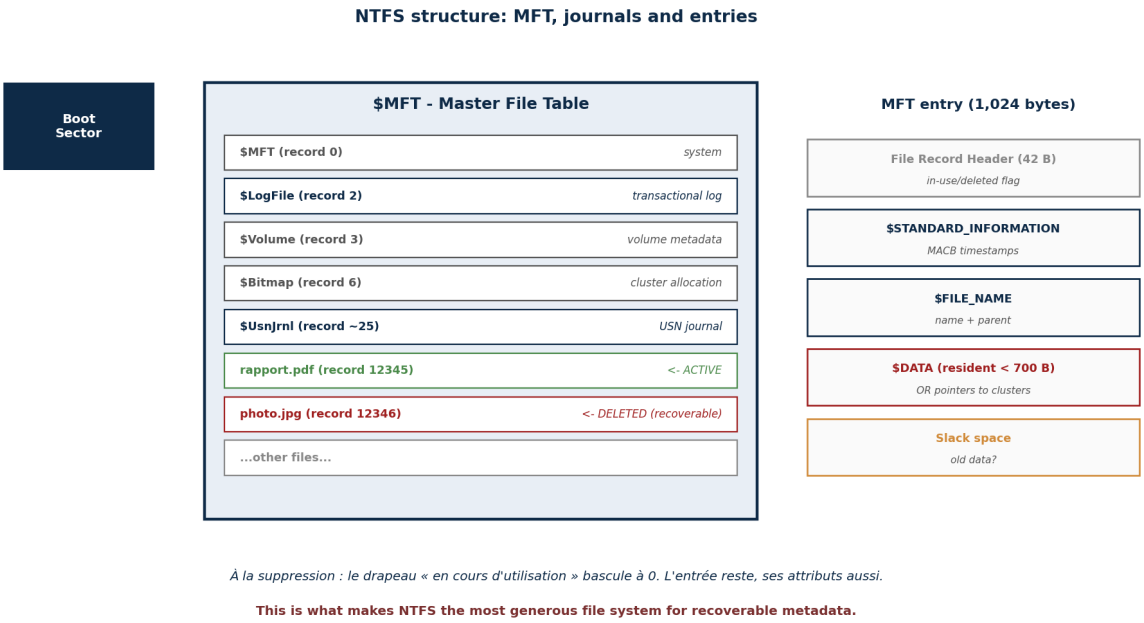

NTFS structure: Boot Sector, $MFT and 1,024-byte entries containing header + attributes.

Its central structure is the **Master File Table** ($MFT): a special file containing a 1,024-byte entry for each file and directory on the volume. This entry includes:

- A header (first 42 bytes) with a flag indicating whether the entry is in use or deleted.
- A $STANDARD_INFORMATION attribute with the four MACB timestamps (Modified, Accessed, Created, Birth/MFT change).
- A $FILE_NAME attribute with the name and a reference to the parent directory.
- A $DATA attribute containing either the file content directly (if smaller than ~700 bytes, called a *resident* attribute), or a list of *runs* pointing to data clusters.

3.4 Multi-source temporal correlation on NTFS

The true power of NTFS in forensics is revealed when multiple metadata sources are correlated:

- The **$MFT** entries and their two sets of timestamps (the $STANDARD_INFORMATION attribute, modifiable by the system, and the $FILE_NAME attribute, modified only at creation and rename). The divergence between the two is a classic marker of timestamp manipulation.
- The $LogFile transactional journal that retains the latest operations on the volume.
- The $UsnJrnl:$J USN journal that lists each creation, modification, rename and deletion with timestamp and reason code.
- The **Windows event logs** (.evtx) that track logins, executions, accesses — correlable at the second level.
- The **registry** (UserAssist, ShellBags, MUICache) that records executions and Explorer navigations.

Typical workflow: export the $MFT to CSV with MFTECmd (Eric Zimmerman), the $LogFile with LogFileParser, the $UsnJrnl with UsnJrnl2Csv, merge these streams in a timeline tool (Timeline Explorer, Timesketch, or Plaso/log2timeline on the open source side) and reconstruct the sequence

of events. This multi-source approach is indispensable as soon as there is a legal stake — a single stream can be manipulated, several concordant streams are much more difficult to falsify.

NTFS slack space — For deleted-file recovery (without legal stake), NTFS remains the most permissive FS: the MFT entry stays after deletion, with all its attributes and often its pointers to data clusters. Tools like MFTECmd can also recover fragments in the **MFT slack space** — unused space inside MFT entries, which often contains remnants of old entries (reference: Sygnia, *The Forensic Value of MFT Slack Space*, 2025).

3.5 ext4: the Linux specificity

ext4 (2008) is the default FS of most Linux distributions. Three key structural elements:

- The **superblock**, which contains the global FS parameters (total number of inodes, block size, etc.).
- The **inode table**, which contains a structure per file/folder with its metadata.
- The **extents**: ext4 does not describe data blocks file by file (as ext3 did with its indirect pointers), but by contiguous ranges (extents). More efficient for large files, but more destructive on deletion.

On deletion of a file on ext4:

1. The inode is marked free in the inode bitmap.
2. The extents are freed in the block bitmap.
3. In the inode itself, ext4 partially erases the pointers to blocks (unlike ext3 which preserved them). This is what makes recovery more difficult on ext4 than on ext3.

The reference tool is **extundelete** (open source), which exploits the ext4 journal to retrieve old versions of deleted inodes before the journal itself rewrites them. When that fails, **debugfs** (included in e2fsprogs) allows the structure to be examined manually:

```
$ sudo debugfs /dev/sda1
debugfs: lsdel # list recently deleted inodes
debugfs: stat <12345> # details of inode 12345
debugfs: dump <12345> /path/file.bin
```

Caution — On ext4, if you have just deleted something important, immediately remount the partition read-only before any other action: sudo mount -o remount,ro /dev/sdXY. Any write, even a simple system log, can reuse freed inodes or blocks.

3.6 APFS: copy-on-write and snapshots

APFS (*Apple File System*, 2017) replaced HFS+ on macOS, iOS and iPadOS. It is a modern FS based on two principles:

- **Copy-on-write.** Any modification writes elsewhere and updates pointers; old versions are never directly overwritten. This guarantees consistency in case of power loss.
- **Snapshots.** APFS can preserve instant images of previous volume state, at marginal cost (since copy-on-write already preserves old blocks). Time Machine on macOS relies intensively on this mechanism.

Consequence for recovery: on an unencrypted APFS, files deleted weeks ago may be present in a local snapshot. Tools like R-Studio, UFS Explorer and Disk Drill exploit these snapshots.

The wall is FileVault. Enabled by default on modern Macs with Apple Silicon, FileVault encrypts the entire volume with AES-256, and the key is protected by the user password and the *Secure Enclave*. Without the password, the physical data on the device is only a pseudo-random stream. See chapter 13 for methods specific to modern Macs.

3.7 Btrfs and ZFS: maximum robustness

Btrfs (Oracle, integrated into the Linux kernel since 2009) and ZFS (Sun Microsystems, 2006, now OpenZFS) are two FS in the family of "copy-on-write + checksums + snapshots", designed for resilience at scale. They are found mainly on NAS (Synology, QNAP), Linux servers, and storage appliances (TrueNAS).

- **Checksums.** Each block is protected by a checksum. Silent corruption detection (*bit rot*) is built in.
- **Self-healing.** On a mirror or RAID-Z volume, ZFS can automatically rewrite a corrupted block from a healthy copy.
- **Snapshots.** Like APFS, at near-zero cost, and exploited by modern NAS units to offer users point-in-time backups. Critical in post-ransomware: most NAS ransomware encrypts visible files but does not touch the Btrfs or ZFS snapshots (see chapter 12).

For recovery: these FS are in practice very little vulnerable to loss by simple logical corruption; they are by contrast to **extreme fragmentation** (copy-on-write naturally fragments content over modifications) and to the complexity of their internal structures, which makes classic data carving very inefficient. The most productive approach on these FS is almost always **via snapshots**, not via carving.

3.8 Summary table

FS	Platforms	Behavior on deletion	Logical recovery quality
FAT32	USB, SD, older systems	Entry marked 0xE5, FAT reset	Very good (loss of 1st letter)
exFAT	USB, high-capacity SD	Same as FAT32, extended capacity	Very good
NTFS	Windows	MFT entry marked deleted, $LogFile and $UsnJrnl retain traces	Excellent — the richest
ext4	Linux	Inode freed, extents erased aggressively	Medium — extundelete within a short window
APFS	macOS, iOS	Copy-on-write, snapshots retained per policy	Excellent via snapshots; nil if FileVault without key
Btrfs / ZFS	NAS, Linux servers	Copy-on-write, snapshots, checksums	Excellent via snapshots

PART II — DIAGNOSIS

Chapter 4

Causes of data loss: 2025-2026 figures

4.1 Four families of causes

Losses are classified into four categories that call for very different responses:

- **Human and logical**: accidental deletion, format, wrong manipulation, configuration error. The device is healthy; the data is logically inaccessible but physically present.
- **Cyber**: ransomware, destructive malware, wipers, malicious deletion. Ransomware adds an encryption layer; wipers (NotPetya for example) actually destroy.
- **Hardware**: mechanical failure (HDD), electronic failure (PCB burnout), NAND wear (SSD).
- **Environmental**: fire, flood, power surge, theft, physical destruction.

4.2 Ransomware, the new norm

The Verizon DBIR 2025 report is the statistical reference on data breaches. Over the period covered (November 2023 to October 2024), Verizon analyzed over 22,000 incidents and 12,195 confirmed breaches.

Indicator	2024 value (DBIR 2025)	Trend
Ransomware in breaches	44%	+37% vs DBIR 2024
Ransomware in SMB breaches	88%	Worsening
Ransomware in large enterprises	39%	Stable
Stolen credentials (initial vector)	22%	Still #1
Exploited vulnerabilities	20%	+34%
Third-party involvement (supply chain)	30%	Doubled
Median ransom paid	$115,000	Decreasing
Refused to pay ransom	64% of victims	+14 pts in 2 years

Source: Verizon Business, 2025 Data Breach Investigations Report, published April 23, 2025.

Two opposite readings of the same report. The pessimistic: ransomware has become a dominant modus operandi, particularly devastating for SMBs which are victims in nine out of ten cases. The optimistic: the median ransom is decreasing, two out of three victims now refuse to pay.

4.3 Human error, still the majority

The human element (in the broad sense: error, privilege abuse, social engineering) remains involved in a dominant share of breaches. The DBIR 2025 quantifies this at 60% of all studied breaches. Definitions vary from one report to another — some reach 95% by including any human action upstream of the incident — but the order of magnitude is constant.

For data recovery specifically, the human causes most frequently encountered by laboratories are accidental deletion of files or directories, mistaken formatting (often during OS installation), file overwriting by bad manipulation, mass deletion by script or command, and loss of password or encryption key.

4.4 Average hardware reliability

For HDD hardware failures, Backblaze remains the public reference with its *Drive Stats 2025* report. Over 337,192 production drives end of 2025, global annual AFR of 1.36%. Cumulative AFR over drive lifespan (lifetime): 1.30%. For SSDs, no comparable large-scale public report exists; the theoretical endurance of modern SSDs (TBW guaranteed by manufacturers) normally covers 5 to 10 years of consumer use; actual failure is often due to a failing controller or corrupted firmware, not cell wear.

4.5 Summary

1. **Human error** remains the number one cause in volume, particularly for individuals and SMBs.
2. **Ransomware** has exploded to become the first cause in terms of financial impact — it is now present in nearly one in two enterprise breaches.
3. **Hardware failures** are slowly decreasing on HDDs (better quality, AFR under 1.5%); on SSDs, they are less frequent but often more catastrophic.
4. **Environmental incidents** represent a few percent of cases, but their cost can be massive.

PART II — DIAGNOSIS

Chapter 5

Diagnosis and triage

5.1 Why diagnosis is the critical step

Diagnosis decides everything: the choice between software or physical intervention, between personal attempt or laboratory shipping, the predictable cost and time, and often the outcome itself. A wrong diagnosis leads to inappropriate actions that make the situation worse.

> **Caution —** As long as the diagnosis is not made, do not touch anything. Above all, do not run any recovery tool "to see": many of them write to the source device as soon as they are installed or on the first scan.

5.2 Five-step triage procedure

1. **Observe.** Is the device powered? Is it making noise? What noise? Is it hot? Does the BIOS/UEFI detect it? Does the OS display it in standard tools (Windows Disk Management, macOS Disk Utility, lsblk Linux)? Note everything without modifying anything.
2. **Classify.** From the observations, classify as: *physical* (not detected, abnormal noises, overheating), *logical* (detected but inaccessible, missing files, corrupted FS), or *hybrid* (intermittent detection, partial reads).
3. **Assess the stakes.** Is the data irreplaceable? Does it have legal, medical, professional value? Does it exist elsewhere (backup, cloud)? Is the urgency real?
4. **Decide the path.** Depending on the combination (type / stakes / available skills), choose: personal intervention, intervention by a generalist IT technician, or shipment to a specialized laboratory.
5. **Document.** If the stake may become legal, photograph the state of the device, note serial numbers, trace each manipulation. This is the chain of custody (see chapter 14).

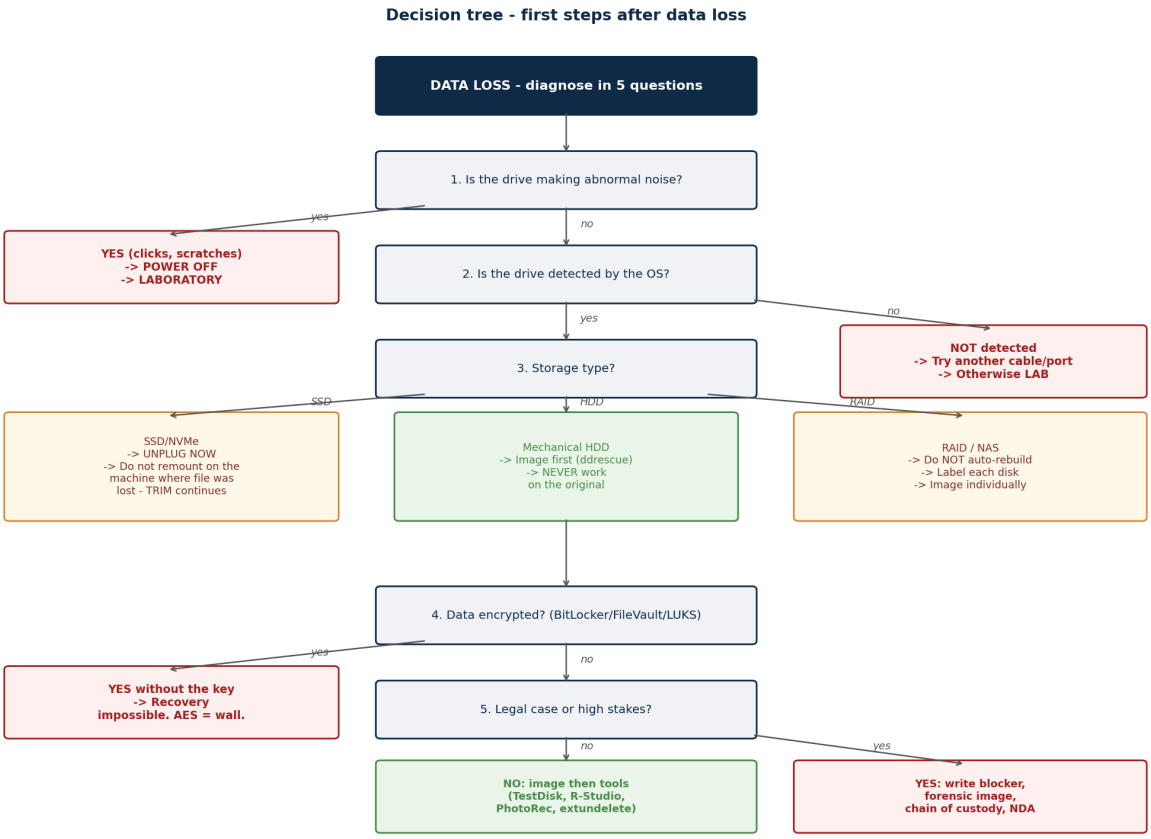

Quick decision tree after data loss: the five first questions that guide the next steps.

5.3 Typical symptoms and their meaning

Observed symptom	Probable diagnosis	Immediate action
Repeated clicks on HDD	Failed heads - physical	Power off immediately, do not reconnect
Silent HDD, not detected	PCB or motor failed - physical	Power off, ship to lab
HDD detected with absurd capacity (0, 8 MB)	Corrupted Service Area	Ship to lab (PC-3000)
HDD detected, major slowness	Bad sectors - evolving physical	Emergency imaging with ddrescue
SSD not detected at all	Controller failed - physical	Power off, ship to lab (JTAG/chip-off)
SSD detected, weird capacity (8 MB, 0 MB)	Corrupted firmware - physical	Ship to lab
Missing partition, corrupted table	Logical	ddrescue image then TestDisk
Files accidentally deleted on HDD	Logical - wide window	Unplug, image, R-Studio
Files deleted on SSD	Logical - short window (TRIM)	UNPLUG IMMEDIATELY
RAW volume (NTFS/exFAT corrupted)	Logical	Image then R-Studio/UFS Explorer
Files encrypted with unknown extensions	Ransomware	See chapter 12
Password prompt on entire disk	BitLocker/FileVault/LUKS	Find the recovery key

5.4 When to go to a laboratory

Three criteria, of which at least one is enough:

- The device is not detected, or emits abnormal sounds → physical → lab.
- The data is irreplaceable and critical (medical, legal, high-stake professional) → lab, even if the case seems simple.
- You have already tried something that made the situation worse → stop, and let a pro assess what is still recoverable.

5.5 Choosing a laboratory

Criteria of seriousness to check with a laboratory:

- Certified ISO 5 cleanroom (ask for the ISO 14644-1 certificate).
- Free diagnosis and payment on success — market standard in 2026.
- List of recoverable files communicated before payment (DAFOTEC calls this service "VeriFiles": the client explicitly validates the list of files from the diagnosis before the flat fee is invoiced).
- Written confidentiality (NDA), especially for professional cases.
- Verifiable reputation (independent reviews like Ekomi/Trustpilot, technical publications, mentions in specialized press). For international references: Ontrack, Kroll, DriveSavers, SalvageData, Gillware, Secure Data Recovery. For France: DAFOTEC in Roubaix (ISO 5 lab since 2004, public institutional clients: French National Police, Tourcoing University Hospital, CNRS, INSERM, French universities).

Signs of a dubious laboratory — If you are asked for full payment before diagnosis, or promised a guaranteed success rate, leave. No serious laboratory promises a recovery before opening the device.

Part III

Methods

Six chapters that constitute the operational heart of the book: how to image a device (the foundation of any other operation), how to analyze the logical layer, how to sculpt files from raw data (carving), how to physically intervene on a hard drive and an SSD, and how to handle RAID configurations.

Three golden rules — Three transversal rules apply to all these chapters and will not be repeated on every page: (1) **image first**, then work on the image; (2) **never write to the source device**, do not install software on it, do not recover files to it; (3) on SSDs with TRIM, **every second under power reduces the window** — unplugging early is better than diagnosing late.

PART III — METHODS

Chapter 6

Safe imaging

6.1 The principle

Imaging consists of creating a bit-by-bit copy of the source device to an image file, then working exclusively on this image. The original is set aside, safe.

Three reasons:

1. **Save what can be saved.** A failing device often fails more in the following hours. The first read may be the last one that succeeds.
2. **Work calmly.** On the image, you can make as many attempts as you want. If a tool corrupts something, another copy is made.
3. **Preserve evidence.** In forensics, the original is sealed with its reference hash; all work is done on a copy also hashed. This is what makes the procedure admissible in court (ISO 27037 standard).

6.2 ddrescue: the reference tool

GNU ddrescue (package gddrescue on Debian, Ubuntu, and most Linux distributions) is the open source reference tool. Its superiority over standard dd rests on three elements:

- A mapfile that precisely records the areas already copied, to be re-copied, slow or failed. It allows resuming exactly where it left off, and programming targeted passes.
- A multi-pass reading strategy: fast first (slow areas are skipped), aggressive next (the difficult areas are revisited).
- Fine handling of error sectors, with configurable maximum retry count.

Typical workflow

Identify the device precisely before anything else. A letter mistake, and a healthy disk is overwritten.

```
$ lsblk -o NAME,SIZE,MODEL,SERIAL,TRAN
NAME SIZE MODEL SERIAL TRAN
sda 500G ST500LM030-2E717D WCC6Y0L1234 sata
sdb 2T WDC_WD20EZRZ WD-WCC4M1234 usb
...

# The serial number confirms we target the right drive.
```

First pass, fast: copy what reads easily, skip the rest.

```
$ sudo ddrescue -f -n -d /dev/sdX /path/image.img /path/image.map

-f : allows output to a device
-n : 'no-scrape': does not dwell on difficult areas
-d : direct access to the device (bypass OS cache)
```

Second pass, targeted on difficult areas with a few retries:

```
$ sudo ddrescue -f -d -r3 /dev/sdX /path/image.img /path/image.map
```



```
-r3 : up to 3 retries per error sector
```

On a clicking drive (failed heads), we avoid going above -r3: each additional attempt is one more chance to damage the surface. On a drive whose problem is purely electronic, we can go up to -r10 without physical risk.

Useful option on particularly damaged cases: reverse-read pass, which helps when the order of sector presentation matters (difficult positioning on certain cylinders):

```
$ sudo ddrescue -f -d -R -r3 /dev/sdX /path/image.img /path/image.map

-R : 'reverse': reads from end to start
```

Once the image is obtained

```
$ sha256sum /path/image.img > /path/image.img.sha256

# To explore the image read-only:
$ sudo losetup --read-only --find --show /path/image.img
/dev/loop0

# On ext4: noload prevents journal replay (which would write)
$ sudo mount -o ro,noload /dev/loop0 /mnt/recovery

# On NTFS via ntfs-3g:
$ sudo mount -t ntfs-3g -o ro,norecover /dev/loop0 /mnt/recovery

# Unmount and free the loop when done:
$ sudo umount /mnt/recovery
$ sudo losetup -d /dev/loop0
```

Caution — Store the mapfile and image on a separate device from the source. If the image is on the same USB stick as the failing device, and that stick fails, everything is lost. Always use a healthy, fast disk with enough space.

6.3 Alternatives to ddrescue

FTK Imager (AccessData/Exterro, free). Very popular Windows tool in forensics. Creates images in E01 (EnCase) or raw DD format, automatically computes MD5/SHA1/SHA256, supports hardware write blockers.

Guymager (Linux, GUI). Graphical interface for forensic image creation. More practical than ddrescue for occasional users but less flexible on difficult cases.

dc3dd (DoD Cyber Crime Center). Variant of dd enhanced for forensics: on-the-fly hashing, logging, validation.

Hardware tools: PC-3000 Disk Imager (ACE Lab), DeepSpar Disk Imager, Atola Insight Forensic. These hardware solutions handle severely damaged drives better than purely software tools (direct controller control, fine reset management, masking of failed heads on HDD). Reserved for laboratories due to cost.

6.4 Write blockers

For forensic cases, between the source device and the analysis machine, a hardware **write blocker** is interposed: a device that lets reads pass but physically blocks any writes. Reference brands: Tableau (OpenText), WiebeTech (CRU).

For non-forensic but cautious use, the software write blocker via blockdev --setro on Linux is sufficient in practice.

PART III — METHODS

Chapter 7

Logical analysis and FS repair

7.1 The principle

Once the image is obtained, logical analysis seeks to **repair or interpret the file system structures** present on the image, rather than directly searching for files in the raw data. This is almost always more effective than carving: we recover not only the content, but also names, dates, the tree structure.

7.2 TestDisk: the partition table

TestDisk (Christophe Grenier, CGSecurity) is the reference tool for repairing MBR and GPT partition tables, and for recovering deleted partitions. Cross-platform.

Workflow on an image that has lost its table:

1. Launch testdisk /path/image.img.
2. Choose "None" for log creation.
3. Select the disk/image, then the table type.
4. Run "Analyse" then "Quick Search".
5. If Quick Search is insufficient, run "Deeper Search".
6. Verify the partitions found, then write the table ("Write").

7.3 NTFS: MFT, $LogFile, $UsnJrnl

- **R-Studio** and **UFS Explorer** (commercial) are the references for NTFS reconstruction.
- **MFTECmd** (Eric Zimmerman, free) parses the $MFT file into exploitable CSV.
- **LogFileParser** and **UsnJrnl2Csv** (Eric Zimmerman) exploit the journals.
- **The Sleuth Kit** (TSK) and its **Autopsy** interface offer a complete open source framework.

```
# With The Sleuth Kit: list all files, including deleted
$ fls -r -p /path/image.img > files.txt
# Deleted files are marked with '*' at the start of the line

# Recover a deleted file by its inode
$ icat /path/image.img 12345 > recovered_file.bin
```


7.4 ext4: extundelete, debugfs, ext4magic

```
# Restore a specific file whose name is known
$ sudo extundelete --restore-file 'home/user/important.pdf' /dev/sdb1

# Restore everything recoverable
$ sudo extundelete --restore-all /dev/sdb1

# With debugfs, list deleted inodes
$ sudo debugfs /dev/sdb1
debugfs: lsdel
Inode Owner Mode Size Blocks Time deleted
1234 1000 100644 524288 128 Sun May 4 14:32:11 2026
debugfs: dump <1234> /tmp/recovered
debugfs: quit
```

7.5 APFS: exploiting snapshots

```
# List APFS snapshots on a live Mac
$ tmutil listlocalsnapshots /

# List snapshots on a volume mounted read-only
$ diskutil apfs listSnapshots /Volumes/data
```

R-Studio and UFS Explorer Professional can exploit APFS snapshots — including from a raw image.

7.6 Btrfs and ZFS: snapshots and recovery

```
# List Btrfs snapshots
$ sudo btrfs subvolume list /volume1

# Restore a file from a snapshot
$ cp /volume1/.snapshots/123/file.pdf /volume1/file.pdf

# On ZFS, even more direct:
$ sudo zfs list -t snapshot
$ ls /pool/dataset/.zfs/snapshot/snapshot_name/
```

On modern Synology DSM 7 / QNAP QTS 5 NAS units, Btrfs snapshots are enabled by default. This is why ransomware attacks of the eCh0raix, QlockerBunny or DeadBolt type can often be circumvented without paying: the ransomware encrypts visible files but does not touch read-only snapshots (chapter 12).

PART III — METHODS

Chapter 8

Deep data carving

8.1 When to carve

Data carving reconstitutes files by searching for their binary signatures in raw data, **without using FS structures**. Last-resort operation when the logical layer is too damaged.

Typical cases: full format (FS structure rewritten), raw image from chip-off, OS reinstallation that wrote its new system over the old, device so corrupted that logical tools find nothing.

8.2 Four levels of sophistication

Level 1 — Simple signature (header / footer)

The historical approach. We scan the device looking for a *magic number* characteristic of the start of the format, then read until the end magic or until a fixed maximum size.

Format	Header (hex)	Footer / end
JPEG (JFIF)	FF D8 FF E0	FF D9
JPEG (Exif)	FF D8 FF E1	FF D9
PNG	89 50 4E 47 0D 0A 1A 0A	49 45 4E 44 AE 42 60 82
PDF	25 50 44 46 ("%PDF")	25 25 45 4F 46 ("%%EOF")
ZIP / DOCX / XLSX	50 4B 03 04	Variable
MP4 / MOV	(offset 4) 66 74 79 70	No fixed footer
RAR (v5+)	52 61 72 21 1A 07 01 00	Variable
SQLite	53 51 4C 69 74 65 20 66 6F 72 6D 61 74 20 33 00	

Very effective on **non-fragmented** files. Fails as soon as a file is fragmented. Tools: **PhotoRec**, **Foremost**, **Scalpel**.

Level 2 — Semantic carving (structure-aware)

We exploit the internal structure of the format. For a PDF, the xref table. For a ZIP, the central directory. This makes it possible to reconstitute fragmented files.

Level 3 — Entropy analysis

Calculating Shannon entropy per block distinguishes text / compressed data / encrypted data and targets carving.

Level 4 — Machine learning

Recent approaches (2023-2026) use neural networks for classification and reassembly. Belkasoft X and Magnet AXIOM announced ML modules in 2024-2025. Promising results on some formats, in practice limited.

8.3 PhotoRec: the open source standard

```
# Launch PhotoRec on an image
$ sudo photorec /path/image.img
```

Steps: choose the device, table type, partition, original FS, **File Opt** to select only the desired types, output directory. PhotoRec is slow: count several hours for 1 TB.

Caution — PhotoRec **never** recovers original names — it renames them f0000001.jpg, etc. For serious forensic use, prefer an FS-aware approach (ch. 7) whenever possible.

8.4 Hard limits of carving

- **Fragmentation.** On highly fragmented FS (old ext4, ZFS/Btrfs copy-on-write), files are split. Simple signature carving recovers the first fragment then noise.
- **Encryption.** Encrypted content has maximal entropy and resembles no known format.
- **Compression.** A ZIP, MP3 or JPEG file that has lost its first blocks is unrecoverable even if the rest is intact.
- **False positives.** On 1 TB of raw NAND from chip-off, carving can produce millions of files, of which 99% are noise.

PART III — METHODS

Chapter 9

Physical intervention on HDDs

9.1 The cleanroom: why it's non-negotiable

The heads of an HDD fly a few nanometers above the magnetic layer. Atmospheric dust (typically 5 to 50 microns) trapped under a moving head is, at scale, the equivalent of a car driving at 200 km/h hitting a concrete wall. The magnetic layer is scratched, often irrecoverably.

To open a hard drive without destroying the platters, an environment is required where particle concentration is drastically reduced. ISO 14644-1 classifies cleanrooms by maximum particle concentration:

ISO Class	Particles >=0.5 micron / m<super>3</super>	Typical use
ISO 1	10	Ultra-precise semiconductor manufacturing
ISO 2	100	Cutting-edge research
ISO 3	1,000	Advanced cleanroom
ISO 4	10,000	Standard semiconductor manufacturing
ISO 5	100,000	Professional HDD recovery
ISO 6	1,000,000	Insufficient for HDDs
Ambient air	~35,000,000 (dusty to clean)	Out of scope

The standard for HDD recovery is ISO 5. This is what serious international laboratories have (DriveSavers, SalvageData, Gillware, Secure Data Recovery, Ontrack) and in France at DAFOTEC in Roubaix.

Caution — Opening an HDD in an ordinary room, even a clean one, is almost certain additional destruction. YouTube videos showing amateurs doing a head swap in their garage are irresponsible: they ignore the environment of 35 million particles per m^3 and overvalue the rare cases where it works anyway.

9.2 Head swap (head transplant)

When the heads fail (from wear or shock), they are replaced with those of a physically identical **donor** drive. Procedure:

1. Diagnosis confirming that the platter mechanics are intact and that we are facing a head problem (typical clicks, abnormal read-signal amplitude at startup).
2. Sourcing of a strictly identical donor drive. Laboratories keep stocks of hundreds of models, or buy them from specialized wholesalers.
3. Opening of both drives under laminar flow, dismantling of head/arm assemblies with dedicated tools (specific wrenches, head separators to prevent them from touching when outside the drive).
4. Transfer of the head/arm assembly from donor to patient.

5. Closing of the patient, powering on.
6. Immediate imaging with a hardware imager (PC-3000, DeepSpar), before the new combination wears out in turn.

The donor must be identical down to the firmware revision. An apparently identical donor model but of another firmware revision can yield a drive that spins but reads nothing: the head/firmware calibration is frozen at the manufacturer.

9.3 PCB swap and ROM reprogramming

If the failure is on the external printed circuit (power surge, blown TVS), the PCB can be replaced with that of a donor. But beware: on most modern drives, calibration parameters specific to the physical device (bad sector map, head parameters, etc.) are stored in a ROM on the PCB. Without transfer of this ROM (by desoldering and resoldering, or via PC-3000), the patient drive with a donor PCB will at best produce unreadable data.

PC-3000 also allows the ROM to be reprogrammed directly via a COM/UART connection to the test points of the PCB. Fast and non-destructive.

9.4 Special cases

9.4.1 SMR drives

Drives in SMR (see 1.3) pose specific challenges. When firmware or the translation zone is corrupted, the platters cannot simply be read: the LBA → physical location mapping must be reconstructed taking into account the cache and the bands. PC-3000 published dedicated SMR modules from 2020. The success rate remains lower than equivalent CMRs.

9.4.2 Stiction

When heads remain stuck to the platter at startup, one can try to manually "unstick" them in the cleanroom by rotating the platters by hand while briefly applying power. Very delicate — there is a risk of ripping off the magnetic layer under the heads.

9.4.3 Scratched platters

If the scratch is superficial and localized, what is outside the scratched area can often be recovered (with lost sectors). If the scratch is deep or extensive, it is terminal — the magnetic layer has been torn off, the data is no longer there.

Some extremely specialized laboratories practice **platter swap**: transfer of platters from a patient drive to a donor's mechanical enclosure. Technically the most delicate (moving a platter without desynchronizing its relative position to the heads is almost impossible) and the success rate is low. Reserved for extreme cases with very high stakes.

9.5 Professional hardware platforms

Three platforms dominate:

- **PC-3000** (ACE Lab, Russia): global de facto standard. Modules for HDD, SSD, Flash, RAID, mobile. Knowledge base by controller and firmware kept up to date.
- **DeepSpar Disk Imager** (Canada): very effective alternative for HDD imaging with low-level control (head bypass, precise timeout configuration).

- **Atola Insight Forensic**: advanced forensics with multi-pass imaging and automatic detection of failing heads.

Complete laboratory equipment (PC-3000 HDD + Flash + Express + adapters + micro-soldering tools + ISO 5 cleanroom) represents an investment of several tens of thousands of euros. This is the main reason physical recovery is invoiced at several hundred euros minimum in any serious laboratory.

PART III — METHODS

Chapter 10

Physical intervention on SSDs

10.1 Why it's harder than an HDD

On an HDD, the data is magnetic and persistent. If we manage to spin the platters and read them (by head swap, PCB swap, etc.), we access the bits as they were written.

On an SSD, the data is:

- Electric (charge trapped in cells), therefore potentially volatile over time if not refreshed.
- Scrambled by the controller to balance electrical charges — descrambling must be known.
- Coded with an ECC proper to the controller — decoding must be known.
- Logically mapped by a **Flash Translation Layer** (FTL) that only the original controller knows perfectly.
- Often hardware-encrypted (SED, TCG Opal) by a key only the original controller can unlock.

Three intervention techniques, in order of preference (least destructive first):

10.2 JTAG / ISP: non-destructive

Modern SSD controllers often expose test points corresponding to an internal debug protocol: **JTAG** (Joint Test Action Group) or **ISP** (In-System Programming). By temporarily soldering thin wires onto these points, and connecting to a dedicated programmer, one can:

- Read the controller firmware.
- Inject a recovery *loader* that short-circuits the dead firmware and accesses the NAND directly via the controller's capabilities.
- On certain configurations, make the controller speak as if it worked normally, and image via SATA or NVMe.

Advantages: the NAND is read by its original controller, so descrambling and ECC are handled automatically, and the hardware encryption remains decrypted if the key is present. The physical device is not destroyed.

Disadvantages: requires specialized tools (PC-3000 Flash with JTAG adapters, or third-party solutions like RTPro, Medusa Pro), good precision micro-soldering skills, and knowledge of test points for each controller family. PC-3000 maintains a schematic database by controller, accessible to labs under contract.

10.3 Chip-off: the last-resort technique

If JTAG fails or is not applicable (dead controller, missing test points), the NAND chips are physically desoldered to be read independently of the controller.

10.3.1 Physical removal

Modern NAND chips are in BGA (Ball Grid Array) packaging with tens to hundreds of solder balls under the component. Desoldering it requires:

- A calibrated hot-air rework station with precise thermal profile (controlled heating ramp to neither crack the chip nor destroy the PCB).
- Thermal protection of the rest of the PCB (kapton, heat shields).
- A binocular microscope to reposition the chips on reading sockets.

10.3.2 The read

Once the chip is desoldered and cleaned (balls remade if needed), it is read on a universal NAND programmer: PC-3000 Flash, Soft-Center Flash, FlashExtractor. These programmers read raw pages as they are physically stored.

10.3.3 Logical reconstruction (the real challenge)

At this stage, we have a raw dump of each NAND chip. To extract usable data from it, we must:

1. **Descramble** the pages. The original controller applied an XOR with a pseudo-random sequence dependent on the address to balance the charges. Without knowing the LFSR polynomial (Linear Feedback Shift Register) used, impossible to invert.
2. **Decode the ECC**. Modern controllers use LDPC with several hundred parity bits per page. Without reproducing this pipeline, read errors are not corrected and a lot of data is lost.
3. **Reassemble the pages**. On multi-chip SSDs, the logical pages are distributed between chips according to an interleaving scheme proper to the controller. This order must be reconstructed.
4. **Reconstruct the FTL**. From the metadata embedded in each page (spare area), reconstitute the LBA → physical location mapping table. This is the most complex work.
5. **If the SSD was hardware-encrypted**, and the controller key could not be recovered — it's over. The readable content is encrypted.

This reconstruction is what makes a complete chip-off on a modern SSD take weeks in the laboratory, and the cost is high (laboratories like DAFOTEC display flat fees up to 950 EUR for complex failures).

10.4 Apple Silicon case: maximum difficulty

Macs with T2 chip (2018+) or Apple Silicon M1/M2/M3/M4 integrate the SSD controller into the CPU/SoC itself. Concretely:

- The NAND is soldered directly to the motherboard (no removable SSD).
- The SSD controller is in the SoC.
- Hardware encryption is bound to a unique identifier of the SoC's **Secure Enclave**.
- FileVault is enabled by default on Apple Silicon.

For recovery without the user password, it's impossible: the key is protected by the Secure Enclave. If the password is known but the motherboard is defective, specialized laboratories can as a last resort transplant the critical components (CPU/SoC containing the Secure Enclave) onto a blank donor motherboard to preserve cryptographic coherence, then read the NANDs. DAFOTEC publicly documents this intervention under the name "Mobile CPU Swap" for smartphones, and applies

comparable techniques on Macs with soldered SSDs.

PART III — METHODS

Chapter 11

RAID and advanced storage

11.1 Configuration recap

RAID (Redundant Array of Independent Disks) combines several disks for performance, resilience, or both. The common levels in 2026:

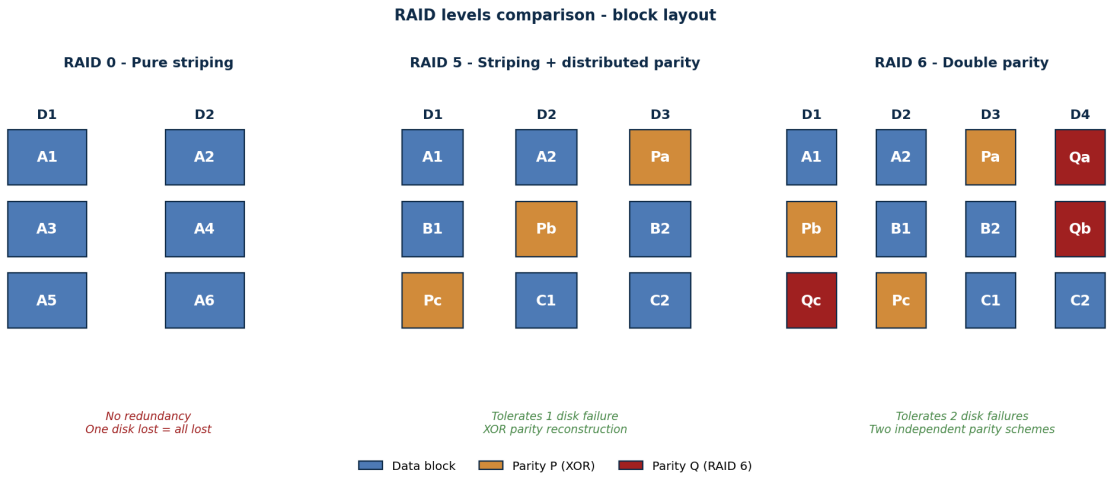

Comparison of RAID 0 / RAID 5 / RAID 6: layout of data blocks and parity.

Level	Description	Fault tolerance	Usage
RAID 0	Pure striping, performance	None	Caches, scratch — NEVER for data
RAID 1	Mirroring	1 disk	Small servers, simplicity
RAID 5	Striping + distributed parity	1 disk	NAS, SMB servers — common
RAID 6	Striping + double parity	2 disks	High-capacity storage
RAID 10	Mirror of stripes	Up to N/2	Performance + redundancy
SHR	Synology Hybrid RAID, variable size	1 (SHR) or 2 (SHR-2)	Synology NAS

11.2 The RAID golden rule: image BEFORE everything

The mortal trap in RAID recovery is the **automatic rebuild** on a degraded array. When a disk fails, many hardware RAID controllers propose (or automatically launch) a rebuild onto a replacement disk. The problem: the rebuild **writes massively** on the remaining disks. If another disk was silently dying, the additional load from the rebuild can tip it over.

> **Caution —** On a degraded RAID containing important data: **stop the server, physically label each disk with its bay position, and image each disk individually before any other operation**. Rebuilding will be done later, on the images, in an isolated environment.

11.3 Software reconstruction

Once each disk is imaged read-only, the RAID is reconstructed virtually with tools like R-Studio Network, UFS Explorer RAID, ReclaiMe Pro, or (on the open source side) mdadm on Linux. The challenge is to find the exact parameters:

- The **disk order** in the stripe.
- The **stripe size** (chunk size: 64 KB, 128 KB, 256 KB are common values).
- The **parity scheme** (left-symmetric, left-asymmetric, right-symmetric, right-asymmetric).
- The initial **offset** (the first MBs of each disk are often reserved for the RAID controller).
- The **parity delay** on certain Dell PERC and HP Smart Array configurations.

On proprietary hardware RAID controllers (Dell PERC, HP Smart Array, LSI/Avago/Broadcom), tools like R-Studio and UFS Explorer Professional automatically recognize configuration metadata. When this metadata is corrupted, parameters must be guessed by entropy analysis.

LAB REPORT — DAFOTEC • RAID 5 Dell PowerEdge, 3 disks in "Foreign" state (case #2)

Hardware: RAID 5 / 6 disks of 4 TB • **Turnaround:** 48 hours (on-call) • **Flat fee:** 2,700 EUR (French market)

Symptom. Power outage and surge. On restart, the PERC H730 controller flags 3 disks as "Foreign" and refuses to mount the array. Auto-rebuild had fortunately been disabled by the administrator.

Intervention. Physical labeling of each disk in its bay. Forensic cloning of each of the 6 disks with ddrescue and write blocker, onto identical destination disks. Virtual reconstruction of XOR parity without any writes to the originals. Identification of disk order and parity delay specific to the PERC H730 through entropy analysis.

Outcome. 22 TB of accounting data recovered intact. Firm activity resumed within 48 hours.

Accounting firm (38 staff), Lyon area - business interruption avoided. Case published on dafotec.fr.

11.4 Synology / QNAP NAS

Consumer and SMB NAS units often use modern FS (Btrfs on Synology, ZFS on some QNAP) with their own configuration metadata (SHR, SHR-2 at Synology). When a NAS fails:

- Never reset the NAS — the RAID configuration is in the metadata on the disks.
- Extract the disks noting their original order.
- On Linux, mdadm --examine on each disk reveals the configuration.
- UFS Explorer Professional and ReclaiMe Pro automatically recognize Synology and QNAP configurations.

NAS and ransomware — On modern Synology DSM 7 NAS units, the Btrfs file system enables snapshots by default. In case of a ransomware attack, the typical procedure is: 1) forensic cloning of the disks, 2) read-only assembly with mdadm --assemble --readonly, 3) btrfs -o ro,recovery mount, 4) btrfs subvolume list -s to identify snapshots prior to the attack, 5) extraction of clean data from these snapshots. On a large number of eCh0raix, QlockerBunny and DeadBolt attacks, this procedure avoids any ransom payment.

11.5 Recovery from chip-off on NAS

When the disks of a NAS are themselves physically failed (multi-disk failure after power surge for example), the previous techniques must be combined: recover each disk individually (head swap, PCB swap, emergency imaging if possible), then reconstruct the RAID virtually from the partial images obtained. The success rate depends on the combination.

11.6 Recovery on monolith devices (microSD, all-in-one USB)

Monolith devices are components where the NAND, controller and interface are embedded in the same epoxy resin: microSD, all-in-one USB 3.0 sticks, some CFast cards. When the component is physically damaged (oxidized contacts, cracked resin, crushed device), neither the chips nor the controller can be accessed by classical means.

The recovery technique consists of:

1. Abrading the epoxy resin by laser or chemically to expose the internal traces of the component.
2. Identifying under microscope the contact points corresponding to the NAND data pins.
3. Micro-soldering extremely fine copper wires (0.02 mm) onto these points.
4. Connecting to a universal NAND programmer and performing a raw dump as in a classic chip-off.
5. Logically reconstructing (descrambling, ECC, FTL) as in 10.3.3.

This highly technical procedure is documented publicly by DAFOTEC under the name *Spider Web*, referring to the web of fine copper wires that results under the microscope. Very few laboratories in France master this technique.

LAB REPORT — DAFOTEC • microSD card from DJI Mavic 3 drone (case #6)

Hardware: microSD 256 GB UHS-II (monolith device) • **Turnaround:** 4 days • **Flat fee:** 320 EUR (French market)

Symptom. Drone crashed from about 40 meters height. microSD card not recognized, contacts visibly oxidized, component resin cracked along the upper edge.

Intervention. Spider Web technique: laser abrasion of the epoxy resin at the NAND pins, identification of contact points under microscope, micro-soldering of 28 copper wires of 0.02 mm on the data pins. Raw dump on NAND programmer, then logical reconstruction (descrambling, ECC, page reassembly).

Outcome. 240 GB recovered out of 256 GB - 4 hours 20 minutes of 4K architectural scouting footage. The missing 16 GB corresponded to physically crushed sectors at impact.

Professional videographer, Nice area - commercial shoot saved. Case published on dafotec.fr.

Part IV

Special cases

Three domains where general techniques run into specific constraints: encryption (which can turn a trivial case into an impossible one), mobile devices (smartphones, modern Macs), and the judicial context (where procedure matters as much as technique).

PART IV — SPECIAL CASES

Chapter 12

Encryption and recovery

12.1 A paradigm shift

When a device is unencrypted, the data is readable by anyone who accesses the physical storage. When a device is encrypted, the data is no longer information; it is a sequence of bits indistinguishable from random noise. Recovery becomes a cryptanalysis problem, which in practice means: **impossible without the key**.

12.2 BitLocker (Windows)

BitLocker encrypts Windows volumes with AES-128 or AES-256. The master key (FVEK) is protected by one or more **protectors**:

- TPM (Trusted Platform Module): the key is stored in the motherboard module, released at boot if the system state is compliant.
- User password.
- 48-digit recovery key (saved in Microsoft account / Azure AD, or printed).
- USB key, smart card.

Without any protector, correctly implemented AES-128 or AES-256 resists all known attacks. Give up.

Good systematic lead: the recovery key saved in the Microsoft account. Retrievable via account.microsoft.com or via the AD/Azure administrator in enterprise.

12.3 FileVault (macOS)

FileVault 2 encrypts the entire APFS volume with AES-XTS. On Apple Silicon Macs (M1 to M4), it is enabled by default and managed by the Secure Enclave.

Recovery paths:

- User password.
- 24-character recovery key (saved in iCloud on activation, or noted by the user).
- Institutional key for enterprise Macs (configured by IT department before deployment).

Without any of these paths, chip-off has no interest: only encrypted data is obtained.

12.4 LUKS (Linux)

LUKS (Linux Unified Key Setup) is the Linux standard. Implemented by cryptsetup. The LUKS header contains up to 8 passphrase slots; each protects an encrypted copy of the master key.

- With known passphrase: cryptsetup open.
- Corrupted header but backup (cryptsetup luksHeaderBackup): restore.

- Without anything: Argon2/PBKDF2 + high cost factor make brute force impractical on modern passphrase.

Prevention tip — Backing up the LUKS header at installation is a good practice. Corruption of the first sectors without backup renders the volume permanently unreadable.

12.5 Self-Encrypting Drives (SED, TCG Opal)

Hardware encryption built into the SSD controller is increasingly the norm. It is:

- Always active (the controller decrypts with a default key if no password is defined).
- Activatable via BIOS/UEFI (ATA password) or TCG Opal Manager software.
- Instantly erasable by *crypto erase* — this is what makes SSD disposal secure in a few seconds.

Several SSDs have had faulty Opal implementations (Crucial MX100/MX200, Samsung 840/850 EVO before firmware EMT02B6Q) that allowed bypassing — but counting on it is playing the lottery.

12.6 Ransomware: what's really possible

When a workstation or server is hit by modern ransomware, user files are encrypted with a key (typically AES) itself encrypted by the attacker's public key. Without the corresponding private key, decryption is mathematically impossible.

Paths to systematically examine before giving in to despair:

1. **Public decryptor?** The *No More Ransom* project (nomoreransom.org), carried by Europol, publishes free tools for variants whose keys have been seized or vulnerabilities discovered. Over 200 tools in 2026.
2. **Machine still on?** If the attack is recent, the AES key in clear may still be in memory. RAM analysis with Volatility or DumpIt.
3. **Shadow Copies?** On Windows, vssadmin list shadows may reveal Shadow Copies that ransomware did not manage to delete.
4. **Btrfs/ZFS snapshots on NAS?** Most NAS ransomware encrypts visible files but forgets read-only snapshots. Typical procedure in chapter 11.
5. **Immutable or air-gapped backups?** Primary question (chapter 17).
6. **Should you pay?** Complex question. 64% of victims refused in 2024 (DBIR 2025). Paying offers no guarantee of decryption.

Caution — Keep the encrypted files even if you cannot decrypt them today. The keys of historical ransomware are regularly seized by law enforcement (LockBit, Hive, REvil) and published months or years later. Store the encrypted files on an offline disk; one day, you may be able to decrypt them.

LAB REPORT — DAFOTEC • Synology DS920+ NAS encrypted by eCh0raix ransomware (case #4)

Hardware: NAS RAID 5 / 4 disks of 8 TB • **Turnaround:** 7 days • **Flat fee:** 1,800 EUR (French market)

Symptom. eCh0raix attack on an Internet-exposed NAS. All user files encrypted with extension .encrypt. Two disks also showing errors on the RAID controller.

Intervention. Forensic cloning of all 4 disks with ddrescue. Read-only assembly of the SHR Btrfs volume via mdadm. Analysis of Btrfs snapshots: identification of the most recent snapshot predating the attack, dated the previous night. Extraction of the clean subvolume via btrfs send/receive to clean storage.

Outcome. 24 TB recovered via the Btrfs snapshot dated D-1. No ransom paid.

E-commerce SMB (12 employees), Lille area - product database intact. Case published on dafotec.fr.

12.7 Case study — Maersk and NotPetya (June 2017)

■ Maersk / NotPetya: backup saved by a power outage

On June 27, 2017, the shipping giant Maersk was hit by **NotPetya**, destructive malware disguised as ransomware, spread via a booby-trapped update of the Ukrainian accounting software M.E.Doc. In 7 minutes, the malware spread across the entire Maersk network: 45,000 to 49,000 workstations, 4,000 servers destroyed, including all ~150 **Active Directory domain controllers**. NotPetya is not reversible: the machines are dead.

Maersk has backups of individual servers (3 to 7 days old), but no backups of domain controllers — the architecture assumed the 150 controllers backed up each other through replication. But they were all destroyed simultaneously. Without AD, nothing can be restored.

Salvation: a domain controller in Ghana was **offline** at the time of the attack, due to a local power outage. It had survived. Maersk had it physically shipped (the network being destroyed) to London where the recovery center had been set up. This controller served as the basis to rebuild the entire infrastructure.

Bottom line: 10 days of total paralysis, estimated loss of $250-300 million for Maersk. Globally, NotPetya cost about $10 billion across Merck, FedEx/TNT, Mondelez, Saint-Gobain. Sources: Wired *The Untold Story of NotPetya* (2018); Control Engineering *Throwback Attack* (2025).

Lesson: online and synchronously interconnected backups do not protect against an attack that destroys them all at once. Survival came from luck. An air-gapped or immutable backup would have avoided the luck.

PART IV — SPECIAL CASES

Chapter 13

Mobile and Apple Silicon devices

13.1 The context

A modern smartphone contains typically more personal data than a desktop computer. But it is also one of the most difficult devices to recover:

- Default encryption (iPhone since 2010; Android 6 since 2015 in practice).
- Monolithic eMMC or UFS storage (NAND and controller in the same chip).
- Strong link with a cloud account.

13.2 iPhone and iOS

Since iPhone 5s (2013), the **Secure Enclave** stores keys and applies a strict policy on decryption. Any data on the NAND is encrypted by a key bound to the passcode and a hardware UID burned into the Secure Enclave.

- Chip-off yields uninterpretable encrypted data.
- Professional forensic tools (Cellebrite UFED, GrayKey Grayshift, Magnet GrayKey) exploit undisclosed vulnerabilities in certain iOS versions. Windows that close at each update.
- For individuals: iTunes/Finder backup or iCloud backup.

13.3 Android

More heterogeneous. Default encryption since Android 6, based on the passcode and a hardware keystore (**TEE**). Android 10+ uses file-based encryption.

Tools:

- **Professional forensic suites**: Cellebrite UFED, MSAB XRY, Oxygen Forensic Detective, Magnet AXIOM.
- **ADB**: on an unlocked live phone with USB debugging enabled.
- **Download / EDL** mode on some Qualcomm chips — now signed by manufacturer.

For a lost locked Android, the most practical path remains the **Google backup** (Photos, Drive, Contacts, Google One backups).

13.4 Physically destroyed smartphones

When the smartphone is physically destroyed (fall, crushing, motherboard broken in two), advanced techniques exist but require a laboratory:

- **Board-level micro-soldering**: repair of cut traces, replacement of defective components (PMIC, U2).
- **Mobile CPU swap**: transplant of the processor (which contains the Secure Enclave / TEE and therefore the encryption keys) onto a blank donor motherboard, preserving the cryptographic

chain. Standardized procedure in lab-level laboratories, documented publicly by DAFOTEC.

- **Direct NAND reading** on the desoldered eMMC or UFS chip — but without the Secure Enclave, only encrypted data is obtained.

LAB REPORT — DAFOTEC • iPhone 15 Pro crushed by a tractor (case #5)

Hardware: iPhone 15 Pro • **Turnaround:** 10 days • **Flat fee:** 380 EUR (French market)

Symptom. iPhone run over by the wheels of an agricultural tractor. Chassis bent at 90 degrees, motherboard cracked transversally in two, screen pulverized. No response to power.

Intervention. Forensic micro-soldering of the motherboard under microscope: reconnection of cut traces on either side of the crack. Transplant of the Apple A17 Pro processor (mobile CPU swap, containing the Secure Enclave and therefore the device's FileVault keys) onto a blank donor motherboard at BGA station. Preservation of the cryptographic chain. Extraction of the UFS image after boot in DFU mode.

Outcome. 100% of contacts, SMS, photos, calendar and app data (WhatsApp, banking) recovered.

Farmer, Normandy - professional contact book saved. Case published on dafotec.fr.

13.5 Apple Silicon Macs: T2 and M1-M4

On Apple Silicon Macs (M1, M2, M3, M4) and older Intel Macs with T2 chip (2018+), the SSD is **soldered to the motherboard** and its controller is integrated into the SoC. Three consequences:

- No removable SSD — any intervention requires opening and working on the motherboard itself.
- FileVault active by default, key bound to the SoC's Secure Enclave.
- A board-level power failure (PPBUS_G3H, PP3V3_S5, PP1V8_NAND rails) is enough to make the entire device non-bootable despite intact NAND.

Typical procedure to recover an Apple Silicon Mac whose motherboard is defective but whose user knows the FileVault password:

1. Schematics and boardview analysis to identify dead power rails.
2. Component-level repair of blown MOSFETs and PMICs.
3. If the motherboard is too damaged to be repaired, transplant of critical components (SoC, soldered SSD controller, NAND) onto a donor motherboard of identical model.
4. Read the internal SSD via hardware DFU interface.
5. Decrypt with the user password.
6. Extract the image onto a new device.

LAB REPORT — DAFOTEC • MacBook Pro M2, soldered SSD, macOS update blocked (case #3)

Hardware: MacBook Pro 14" M2 (2023) • **Turnaround:** 5 days • **Flat fee:** 480 EUR (French market)

Symptom. MacBook stuck at the Apple logo after a macOS 14.4 update. Recovery mode startup impossible. No standard DFU read. User has the FileVault password.

Intervention. Extraction of the motherboard, boardview analysis. Identification of a defective power rail on the soldered SSD side. Direct read of Apple NAND chips via proprietary hardware DFU interface after board-level repair. Decryption via the FileVault password provided by the client.

Outcome. 890 GB recovered out of 1 TB - complete Lightroom photos, intact Final Cut Pro projects, professional documents.

Independent filmmaker, Bordeaux area - 3 years of rushes saved. Case published on dafotec.fr.

13.6 The role of cloud backups

For mobile devices, this is almost always the most productive path. Recover from iCloud (Photos, contacts, mail, Drive, iOS backups), Google (Photos, contacts, calendar, Drive, Android backups), WhatsApp/Signal/Telegram. The majority of individuals losing mobile data are unaware of what they have automatically backed up.

PART IV — SPECIAL CASES

Chapter 14

Judicial forensics

14.1 The fundamental difference

Classical recovery has a simple goal: recover what can be recovered. Judicial forensics adds a second: **produce a result admissible in court**. Procedure becomes as important as technique.

- Continuous documentation (*chain of custody*).
- Integrity verification by hash at each step.
- Use of write blockers.
- Reproducibility: another expert must be able to redo the analysis and obtain the same results.
- Recognized and validated tools.

14.2 ISO 27037: the reference standard

ISO/IEC 27037:2012 defines the international framework in four phases:

1. **Identification**: locate evidence-bearing devices.
2. **Collection**: take physical possession in a documented manner.
3. **Acquisition**: create a verified forensic copy.
4. **Preservation**: maintain integrity over time, with documented chain of custody.

In France, judicial computer experts operate within the Penal Code framework and must be registered on the list of experts of a Court of Appeal. Seizures are performed by the Judicial Police (cybercrime sub-directorate) with expert support. In the US, FBI, USSS, and ICAC task forces operate under federal rules (Federal Rules of Evidence, Rule 902(14) on digital evidence authentication added in 2017).

14.3 Chain of custody step by step

Here is a complete and concrete procedure, as it can be applied to keep forensic work admissible:

1. **Receipt documentation.** High-resolution photos of the device from all angles, under uniform lighting. Notation of manufacturer, model, serial number, declared capacity. Weighing if relevant. General physical condition. All these elements in a reception report signed by the bearer and receiver, dated to the minute.
2. **Labeling and sealing.** Application of a numbered sealed bag containing the device. Number traced in a master register. Cellophane or tamper-evident bag preferred.
3. **Storage.** Storage in a locked safe or cabinet with traced access (badge, paper register, or both).
4. **Acquisition preparation.** Seal removal: verification that the seal is intact, photo, notation in the master register of the operator, date, time, and reason.
5. **Connection to the analysis workstation.** Always via a hardware write blocker (Tableau Forensic Universal Bridge, WiebeTech Forensic UltraDock, Atola Insight). Notation of model and

serial number of the write blocker.

6. **Acquisition.** Image creation in E01 or raw DD format with FTK Imager, dc3dd, or Atola. Automatic computation of MD5 + SHA1 (or SHA256). Hashes noted in the report.
7. **Cross-verification.** Compute a hash directly on the source via the write blocker (read-only). Compare with the image hash. Both must match bit-for-bit.
8. **Re-sealing.** Immediate return of the source to its bag, re-sealing, photo, register.
9. **Analysis.** All work is done exclusively on copies of the image (not on the original E01 image). Each tool used is noted with its exact version. Each manipulation is logged.
10. **Report.** Description of the device, complete procedure, hashes, tools, methodology, conclusions, list of produced artifacts. Expert signature and date.

Why it's crucial — A faulty chain of custody (undocumented broken seal, missing or divergent hash, inadvertent write to the source, unidentified tool) can be enough to render all conclusions inadmissible — regardless of the technical quality of the downstream analysis. This is the dimension most often underestimated by non-specialists.

14.4 Reference forensic tools

- **EnCase Forensic** (OpenText): historical suite.
- **FTK** (AccessData / Exterro). FTK Imager alone is free.
- **X-Ways Forensics**: European reference, lightweight and fast.
- **Magnet AXIOM**: focus on cloud, mobile and browser artifacts.
- **Belkasoft X**: very good on mobile and messaging.
- **The Sleuth Kit + Autopsy**: open source, free, largely sufficient for many cases.

14.5 Windows artifacts to analyze

- **Registry** (NTUSER.DAT, SOFTWARE, SYSTEM, SAM): config, USB connected, programs executed (UserAssist, ShellBags, MUICache).
- **Prefetch**: trace of executed programs.
- **ShellBags**: folders viewed in Explorer.
- **RecycleBin**: files deleted via the recycle bin.
- **$LogFile and $UsnJrnl**: NTFS journal.
- **Event Logs** (.evtx).
- **Browsers**: history, cookies, cache, downloads.

Caution — *Never* boot the target machine on its original OS. Any boot modifies hundreds of files (timestamps, logs, registry), which can be enough to invalidate the evidence. Always remove the disk or boot on a read-only forensic live distribution (CAINE, DEFT, Tsurugi Linux).

Part V

Practice

Two chapters to move from theory to concrete choices: which tools to use in which situation, and which pitfalls to avoid — with four step-by-step scenarios that cover the most common cases.

PART V — PRACTICE

Chapter 15

Tools in 2026: a realistic overview

15.1 Method

This chapter lists tools without giving a "score out of 5" or quantified "recovery rates". These rankings are almost always: published by affiliate sites of the publishers, based on non-reproducible tests, or copied year after year. What follows is a functional description and honest positioning.

15.2 Open source tools

ddrescue (GNU)

The essential imaging tool. See chapter 6.

TestDisk (CGSecurity)

Partition table repair, deleted partition recovery. For many common logical cases, it's all you need.

PhotoRec (CGSecurity)

Signature-based carving. Over 480 formats recognized. Cross-platform. Open source reference for carving.

The Sleuth Kit + Autopsy

Complete forensic suite. fls, icat, tsk_recover in CLI, Autopsy in web interface.

extundelete, ext4magic, debugfs

The Linux trio for ext4.

Eric Zimmerman's tools

Free Windows suite: MFTECmd, RECmd, LECmd, JLECmd, PECmd, Timeline Explorer. Now references for Windows DFIR.

15.3 Consumer and SMB commercial tools

Disk Drill (CleverFiles)

Windows and macOS. Accessible interface. Free version limited to 500 MB.

EaseUS Data Recovery Wizard

Very clean, multi-FS, free version up to 2 GB.

Recuva (CCleaner)

Free, simple. Suitable for recent accidental deletions on Windows. Limited for complex cases.

Stellar Data Recovery

Complete range. Good reputation on Office and multimedia.

15.4 Professional tools

R-Studio (R-Tools Technology)

Reference for IT technicians and small laboratories. Very good on NTFS, ext4, APFS, RAID.

UFS Explorer (SysDev Laboratories)

Excellent on advanced formats (APFS, ZFS, Btrfs, Synology/QNAP NAS, complex RAID). Very complete Professional edition. Used in many European labs.

ReclaiMe and ReclaiMe Pro

Specialized in RAID reconstruction and proprietary configurations (SHR, ZFS, Storage Spaces).

15.5 Judicial forensic tools

EnCase, FTK, X-Ways, Magnet AXIOM, Belkasoft X, Cellebrite UFED, Oxygen Forensic, MSAB XRY. Reserved for institutions and specialized enterprises.

15.6 Laboratory hardware platforms

- **PC-3000** (ACE Lab): de facto standard. Modules for HDD, SSD, Flash, RAID, mobile.
- **DeepSpar / Atola**: targeted alternatives.

Complete laboratory equipment represents several tens of thousands of euros.

15.7 Decision matrix

Situation	First tools to try
Recent accidental deletion on HDD	TestDisk + PhotoRec (free) or Disk Drill / EaseUS
Deletion on SSD (TRIM probable)	Try Recuva/EaseUS without much hope; SSD unplugged -> lab
RAW volume (NTFS/exFAT corrupted)	TestDisk for partition, R-Studio or UFS Explorer for FS
Synology/QNAP NAS failed	Pull disks, UFS Explorer Professional or ReclaiMe Pro
Degraded RAID 5	Image first, then R-Studio Network or UFS Explorer RAID
Mac with FileVault, known key	Boot in target disk mode, copy; or R-Studio for Mac
Mac with FileVault, unknown key	Check iCloud (recovery key); otherwise give up
Clicking HDD	Power off immediately -> lab (never yourself)
SSD not detected	Lab (JTAG / chip-off)
Apple Silicon Mac non-bootable, FileVault known	Lab (board-level + DFU read)
Judicial case (evidence to produce)	FTK Imager (free) + Autopsy, or pro suite (X-Ways, EnCase, AXIOM)
Locked Android phone	Google backup; for forensics: Cellebrite, MSAB
Locked iPhone	iCloud backup; for forensics: GrayKey (vulnerable models)
Synology NAS encrypted by ransomware	Clone + Btrfs snapshot analysis (chap. 11)

PART V — PRACTICE

Chapter 16

Fatal pitfalls and step-by-step scenarios

16.1 The seven mistakes that kill data

1. **Installing recovery software on the source device.** Installation writes where deleted files are.
2. **Recovering files onto the source device.** Variant just as catastrophic.
3. **Leaving an SSD powered on after the incident.** TRIM and GC continue.
4. **Accepting Windows automatic repair.** chkdsk /f, autorepair on a corrupted FS: Windows actively writes, moves, deletes.
5. **Opening an HDD outside a cleanroom.** See chapter 9.
6. **Desoldering a NAND chip with a soldering iron.** Without a calibrated rework station, the chip is destroyed.
7. **Keeping backups in the same environment as production.** Code Spaces case (below). Modern ransomware targets accessible backups.

16.2 Case study — Code Spaces (June 2014)

■ Code Spaces: 12 hours from company to gone

Code Spaces was a code hosting platform (Subversion and Git) with 7 years of history, based in the UK, fully hosted on AWS.

On June 17, 2014, the company suffered a DDoS attack followed by a message extorting payment, left directly in the EC2 console. The attacker had obtained access to the AWS panel — through credential compromise, with no MFA enabled.

Code Spaces refused to pay and tried to regain control. But the attacker had already created several accounts in the background. Realizing Code Spaces was trying to take back control, he launched a methodical deletion: **EBS snapshots, S3 buckets, AMIs, EC2 instances, storage instances**.

The critical point: *the backups were in the same AWS account as production*. Once access was obtained, the attacker could delete everything simultaneously.

On June 18, 2014 — 12 hours after the start of the attack — Code Spaces announced the permanent cessation of activity. Sources: Threatpost (June 2014); InfoWorld *Murder in the Amazon cloud* (2014); Wiz analysis *breaches.cloud* (2023).

Lessons: (1) never store backups in the same account/domain as production; (2) MFA mandatory on all cloud management accounts; (3) principle of least privilege; (4) have a tested incident response plan.

16.3 Four step-by-step scenarios

Scenario A — exFAT USB stick mistakenly formatted

Symptoms: the stick appears empty after a quick format triggered by mistake. User realizes immediately.

Diagnosis: logical. Quick format rewrote the FAT table and boot sector, but not the data clusters. The files are there, their directory entries too (with possibly lost first letter).

Recommended path:

1. Unplug the stick immediately.
2. On another machine, image the stick: sudo ddrescue -f -n -d /dev/sdX stick.img stick.map.
3. Work on the image. Run TestDisk on it to recover the table: testdisk stick.img.
4. If TestDisk does not rebuild the FS, run PhotoRec on the image for signature carving: photorec stick.img.
5. Recover into a destination folder **on another device**.

Typical outcome: full recovery, with often intact names.

Scenario B — NVMe SSD, critical file deleted 30 minutes ago

Symptoms: an entire directory was deleted on the internal SSD. TRIM is active by default (Windows 10+). The machine is still on and being used.

Diagnosis: logical on SSD with TRIM. Window almost closed — every minute counts.

Recommended path:

1. **Power off the machine immediately** (physical button, not "shut down" which may trigger finalization operations).
2. Remove the SSD. If it is a soldered SSD (ultrabook laptop, Mac), do not power back on — go to lab.
3. For a removable SSD: connect on another machine via a SATA-USB or NVMe-USB adapter **without mounting the file system in write mode**. On Linux: sudo blockdev --setro /dev/sdX before anything.
4. Image: sudo ddrescue -f -n -d /dev/sdX ssd.img ssd.map.
5. On the image, attempt logical recovery with R-Studio or TestDisk. Chances depend strongly on the SSD model and what GC did during the interval.

Typical outcome: random. If TRIM was transmitted and GC has run, little or nothing. If the machine did little work after deletion and GC did not activate, a significant portion can be recovered. On healthy NVMe Samsung 980 Pro or WD SN850 SSDs, do not be optimistic.

Scenario C — RAID 5 degraded after bad handling

Symptoms: on a Dell PowerEdge server with 5 disks of 4 TB in RAID 5, one disk was marked "failed" last week. The admin replaced a disk (perhaps the wrong one) and launched a rebuild. The rebuild crashed midway and the controller now shows two "foreign" disks.

Diagnosis: RAID 5 in maximum danger. A bad intervention may have partially overwritten parity or data. This is typically the lab case.

Recommended path:

1. **Stop the server immediately**. Do not launch any more rebuilds.
2. Physically label each disk: bay position (1, 2, 3, 4, 5), model, serial number, date.
3. Pull the disks and image them **individually** on a separate workstation with write blockers: sudo ddrescue on each, to files named disk1.img, disk2.img, etc.
4. Work on the images only.
5. Either in-house with R-Studio Network / UFS Explorer Professional, or by calling a specialized lab. For this kind of case, professional intervention is generally preferable.

Typical outcome: if a second rebuild was not launched on top, recoverable in the majority of cases. If data was overwritten by a partial rebuild, harder.

Scenario D — Synology NAS encrypted by ransomware

Symptoms: all NAS files appear with a foreign extension (.encrypt, .lockbit, etc.). A README or HOWTODECRYPT file is present at the root. The NAS works but files are unusable.

Diagnosis: NAS ransomware (eCh0raix, QlockerBunny, DeadBolt and their variants). Btrfs snapshots are probably intact if the NAS was running DSM 7+.

Recommended path:

1. **Isolate the NAS from the network immediately** (unplug Ethernet cable). Do not power it off abruptly either, do a normal shutdown.
2. Pull the disks noting the order (bay number).
3. Connect the disks on a dedicated Linux workstation, isolated from the Internet.
4. Assemble the RAID read-only: sudo mdadm --assemble --readonly /dev/md0 /dev/sd[abcd]3.
5. Mount Btrfs in recovery: sudo mount -t btrfs -o ro,recovery /dev/md0 /mnt/nas.
6. List the snapshots: sudo btrfs subvolume list -s /mnt/nas.
7. Identify the most recent snapshot **preceding** the attack (generally D-1 or D-2).
8. Extract the files from this snapshot to new storage.
9. Restore the NAS on a new OS and change all passwords.

Typical outcome: often 100% recovery, without paying ransom, in the majority of cases where snapshots were configured. Otherwise, check nomoreransom.org.

16.4 When to know to stop

Four signs: physically damaged device; multiple unsuccessful attempts; stake greater than pro cost; legal stake. In these cases, hand off.

Part VI

Prevention

The best recovery is the one you never have to do. Two chapters to close the circle: modern backup strategies, and an honest assessment of what remains permanently out of reach of recovery in 2026.

PART VI — PREVENTION

Chapter 17

Modern backup strategies

17.1 The 3-2-1 rule and its extension

The **3-2-1 rule** was formulated in 2005 by Peter Krogh, photographer, in *The DAM Book*. It boils down to: **3** copies, **2** different media, **1** offsite.

Twenty years later, the omnipresence of ransomware has led Veeam to propose an extension **3-2-1-1-0**: **+1** immutable or air-gapped copy, **+0** restoration errors (tested regularly).

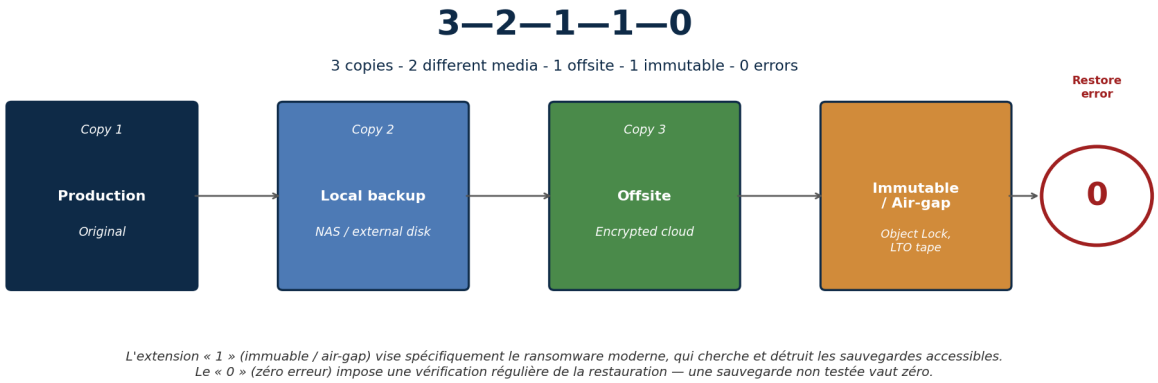

The 3-2-1-1-0 rule — four copies, one immutable, zero restoration errors.

On the origin — The 3-2-1-1-0 rule is Veeam marketing, not an ANSI or NIST standard. That does not detract from the substance: immutability and air-gap have become essential against modern ransomware. The principles are widely adopted (Object First, Wasabi, Backblaze B2 with Object Lock, Azure Blob immutability, Synology SnapLock).

17.2 Implementing immutability

- **S3 Object Lock** (AWS, Backblaze B2, Wasabi, MinIO). *Compliance* or *governance* mode.
- **Azure Blob immutability** and Google Cloud Storage retention policies.
- **Hardened Linux Repository** with chattr +i attribute and SSH disabled for root.
- **WORM** on tapes or optical disks.
- **NAS snapshots** immutable: Synology SnapLock, QNAP WORM.

17.3 Air-gap

An **air-gapped** device is physically disconnected from the network most of the time. Variants:

- **LTO tape** removed from the robot and stored in a safe.
- **USB drive** connected only for backup, then stored.
- **Rotation of several external drives** with rotating offsite storage.

17.4 Restoration verification

The **0** of 3-2-1-1-0 is the most neglected. Many organizations have backups that have never been tested. Common causes of restoration failure:

- Silent corruption (undetected bit rot).
- Broken incremental chain.
- Agent silently crashed months ago.
- Application that does not start from backup (dependency, version, inconsistent database).

Good practice: **quarterly** minimum restoration test, **monthly** for critical systems. Document each test.

17.5 For individuals

1. External drive at 50-100 EUR for Time Machine / File History.
2. Personal cloud: iCloud, Google One, Dropbox, OneDrive, Backblaze Personal Backup.
3. For truly irreplaceable files, third copy on USB stick stored at a relative's place.
4. Annual test: try to restore a random file from each of the three copies.

17.6 For SMBs

- Daily automatic backup (Veeam, Acronis, Datto, Synology Active Backup).
- At least one copy in immutable cloud.
- At least one weekly air-gapped copy.
- MFA mandatory on all admin accounts.
- Documented monthly restoration test.
- Written disaster recovery plan: who does what, in what order, with what contacts.
- Annual exercise: fictional incident, see if the plan holds.

PART VI — PREVENTION

Chapter 18

Current limits in 2026

18.1 What is definitively lost

1. **SSD with TRIM passed and GC completed.** Cells are in their neutral state. No technique recovers.
2. **Data encrypted by AES-256 without the key.** Mathematically infeasible with classical computers. A quantum computer could break AES-128 by Grover but not AES-256, and does not exist at useful scale in 2026.
3. **HDD platters with magnetic layer torn off.** The information was in the metal.
4. **Files encrypted by modern ransomware without the key and without implementation error.**
5. **RAID 5 with more than one disk failed, RAID 6 with more than two.** Parity no longer enough.
6. **Data overwritten by complete rewriting.** The myth of magnetic remanence is debunked for modern drives: a single pass of zeros on a modern HDD renders the data unrecoverable.

18.2 What becomes difficult

- **Mobile recovery**: Secure Enclave, TEE, default encryption close the doors. Cellebrite and GrayKey exploit vulnerabilities that close with each update.
- **Classic SSDs**: reliable TRIM, generalized hardware encryption, shortened window.
- **Cloud**: provider-side definitive deletion is increasingly rigid.

18.3 The final message of this chapter

Data recovery is a real, technically complex discipline that has made enormous progress but runs into increasingly tight physical and mathematical limits. The methods of chapters 6 to 11 work in a majority of cases, but depend almost always on the time elapsed between the incident and the first good decision.

The best strategy remains, unsurprisingly, not to have to recover: prevention, real tested backup, operational discipline when something goes wrong. Chapter 17 is, in practice, the most useful in the book.

Part VII

Horizon

A forward-looking chapter. Where is the discipline heading in the next five to ten years? Which storage technologies will reshape it? Which emerging approaches — AI, post-quantum — will change the landscape? Without promises, without peremptory predictions: trends already taking shape.

PART VII — HORIZON

Chapter 19

Horizon 2030: where the discipline is heading

19.1 On the storage side: density, complexity

PLC NAND and beyond

The first **PLC** memories (5 bits per cell, 32 voltage levels) were announced by Solidigm and Kioxia between 2023 and 2025. Progressive commercialization over 2026-2028 for very high-capacity SSDs (32 TB+ consumer). Consequences for recovery:

- Margins between voltage levels even tighter — under 100 mV.
- Theoretical endurance in free fall (probably less than 150 P/E cycles per cell). Compensated by very sophisticated controllers (fine temperature management, retention monitoring, automatic cell refresh).
- Raw read by chip-off more difficult: requires an ultra-precise VNR, multiple reads, and reconstruction of an even heavier LDPC ECC pipeline.
- Practical consequence: chip-off on PLC will very likely be reserved for top-tier laboratories; the JTAG/ISP path via a live controller will remain the main option.

HAMR at scale

Seagate began commercializing HAMR on its Mozaic 3+ (30 TB and more) in 2024. Toshiba and Western Digital will follow. By 2030, consumer 24 and 30 TB drives will be HAMR. Implications:

- Even higher linear density (above 1.4 TB/platter).
- More complex heads (integrated laser, punctual writing at 450 degrees C).
- Service Area even more critical — laser/heads/platters calibration is sensitive.
- For recovery: no fundamental change in method (head swap, PCB swap remain applicable), but technical margins tighten. Laboratories will have to update their donor stock and HAMR-specific procedures.

DNA, holographic, 5D optical storage

Three *archival* storage technologies have been in R&D; since the mid-2010s. None has reached general consumer commercial stage in 2026, but the promises are real:

- **DNA storage.** Microsoft, Twist Bioscience, Catalog. Encodes data in synthetic DNA sequences. Theoretical density on the order of petabyte per gram, lifespan of several thousand years. Latency and cost still prohibitive for anything other than very long-term archival.
- **Holographic storage** (HVD). Several times announced as ready for commercialization since 2008. Still not generalized in 2026.
- **5D optical / Superman memory crystal** (University of Southampton). Glass nanostructured by femtosecond laser. Experimental capacity of 360 TB on a disk, estimated lifespan of 13.8 billion years.

For the data recovery specialist: these technologies will pose, when they arrive, radically new problems. Reading DNA data in 50 years without the original sequencer? It is no longer the same trade. At this stage, it is too early for these subjects to be anything other than foresight.

19.2 On methods: AI and automation

AI-assisted carving

Several forensic publishers have announced machine learning modules between 2023 and 2025:

- **Magnet AXIOM AI** (since 2024): automatic classification of potentially illegal images, face detection, audio transcription.
- **Belkasoft X AI** (since 2024): reassembly modules for carving image fragments, detection of AI-generated content.
- **Cellebrite AI Reveal** (announced 2025): automatic conversation summary in smartphone extractions.

Realistic promises: yes, these approaches accelerate post-extraction classification on large volumes. Real limits: none works miracles at the low-level carving stage (recovering fragmented files from raw NAND remains a hard combinatorial problem). AI helps downstream, not upstream.

ML-assisted FTL reconstruction

Some academic publications (USENIX, FAST 2023-2025) explore the use of neural networks to identify the scrambling polynomial and ECC pipeline of an SSD controller from signatures on raw NAND. Still experimental; to watch over 3-5 years.

19.3 On cryptography: post-quantum

NIST finalized in 2024 the first post-quantum standardizations (CRYSTALS-Kyber for key exchange, CRYSTALS-Dilithium for signature, SPHINCS+ for hashing). Large-scale deployment in storage will take years. Implications for recovery:

- AES-256 remains robust against Grover (reduces effective security from 256 to 128 bits, which remains unsolvable). AES-128 will become theoretically breakable but not in practice for a long time.
- Post-quantum replacements for hardware encryption (TCG Opal v2.x with post-quantum algo) will arrive probably from 2028-2030.
- Implementation errors in these new algorithms — as happened with some early Opal implementations — will temporarily reopen unexpected recovery windows. Watch CVE/CERT advisories.

19.4 On legal and regulatory aspects

European framework

- **NIS2** (Network and Information Security Directive 2, EU) imposes from October 2024 enhanced obligations for incident management, backup and business continuity. Recovery becomes a regulatory point.
- **DORA** (Digital Operational Resilience Act) for the European financial sector, applicable since January 2025. Strict requirements on disaster recovery tests.

- **GDPR**: case law on retention of recovery copies is refining. Serious laboratories have contractualized destruction of temporary copies after restitution to the client.

US framework

- **HIPAA** (Health Insurance Portability and Accountability Act): mandates breach notification, encryption of PHI, and audit trails. Data recovery on medical devices requires Business Associate Agreement (BAA) with the lab.
- **SOX** (Sarbanes-Oxley Act): requires data integrity for financial reporting in publicly traded companies. Section 404 audits scrutinize backup and recovery procedures.
- **CCPA / CPRA** (California Consumer Privacy Act, amended by California Privacy Rights Act): individual rights to deletion/portability create tension with long-retention recovery copies.
- **FedRAMP** (Federal Risk and Authorization Management Program): authorization framework for cloud services used by US federal agencies. FedRAMP High imposes strict requirements on data location and chain of custody.
- **State breach notification laws**: all 50 states have their own laws, with varying deadlines (30 to 90 days typically). Laboratories handling US data must navigate this patchwork.

19.5 On economics

Three trends:

- **Public pricing** becomes the norm — fifteen years ago, almost no laboratory published its prices. In 2026, DAFOTEC publishes a complete price list, and several French laboratories follow. Commercial opacity pressure recedes.
- **Payment on results** becomes standard. Market acceptance that a client does not pay if nothing was recovered.
- **Validation before payment** (VeriFiles at DAFOTEC, equivalents at other labs) becomes an increasingly explicit client expectation.

19.6 On training

Data recovery remains a trade without a diploma curriculum. Training is done by apprenticeship in laboratories and publisher certifications (ACE Lab Certified, EnCase Certified Examiner, GCFE, CFCE). By 2030, dedicated modules could emerge in cybersecurity masters — some French universities have experimented with this since 2023 (Lille, Rennes, Compiègne, Lyon Master 2 Cyber programs). In the US, SANS DFIR programs and the GCFE certification already cover much of the discipline.

19.7 The trade does not disappear

A recurring fear: with cloud, generalized backups and default encryption, is data recovery a future-proof trade or a dying one?

Three elements converge to say the trade is not disappearing:

1. The total volume of storage produced worldwide grows exponentially. Even if a growing share is encrypted or backed up, the residual unprotected volume increases in absolute value.
2. Hardware failures remain inevitable (HDD AFR ~1.3% per year, finite SSD lifespan). Volumes grow, so do recovery needs.

3. Regulatory requirements (NIS2, DORA, ISO 27001, GDPR, HIPAA, SOX) impose increasingly rigorous forensic analyses and documented recoveries. The trade moves up in procedural demands.

What changes: the trade becomes more technical (more physics, more advanced electronics, more reverse engineering), more standardized (chain of custody, compliance), and more concentrated (laboratories capable of handling modern cases through to the end are few). The discipline does not disappear. It becomes more professional.

Appendices

Quick reference

Five appendices: command reference and useful Python scripts, glossary of technical terms, complete bibliography of public sources consulted, "About DAFOTEC" page, and thematic index.

APPENDICES

Chapter A

Command and script reference

A.1 Device identification

```
# Linux
$ lsblk -o NAME,SIZE,MODEL,SERIAL,TRAN
$ sudo hdparm -I /dev/sdX
$ sudo smartctl -a /dev/sdX

# macOS
$ diskutil list
$ diskutil info /dev/diskN

# Windows (PowerShell)
PS> Get-PhysicalDisk
PS> Get-Disk | Format-List
```

A.2 ddrescue imaging

```
# First fast pass
$ sudo ddrescue -f -n -d /dev/sdX image.img image.map

# Second pass with retry
$ sudo ddrescue -f -d -r3 /dev/sdX image.img image.map

# Reverse pass for severe cases
$ sudo ddrescue -f -d -R -r3 /dev/sdX image.img image.map

# Statistics
$ ddrescuelog -t image.map
```

A.3 Mount an image read-only

```
$ sudo losetup --read-only --find --show image.img
/dev/loop0
$ sudo blkid /dev/loop0

# If partition table:
$ sudo partx --show /dev/loop0
$ sudo partx --add /dev/loop0

# Mount according to FS
$ sudo mount -o ro,noload /dev/loop0p1 /mnt/recovery # ext4
$ sudo mount -t ntfs-3g -o ro,norecover /dev/loop0p1 /mnt/recovery
$ sudo mount -o ro /dev/loop0p1 /mnt/recovery # FAT/exFAT
```

A.4 Integrity hashes

```
$ sha256sum image.img > image.img.sha256
$ md5sum image.img > image.img.md5
$ sha256sum -c image.img.sha256

# On very large files, BLAKE3 is faster:
$ b3sum image.img
```


A.5 NTFS analysis

```
# Sleuth Kit
$ fls -r -p image.img
$ icat image.img 12345 > out.bin

# Zimmerman tools (Windows)
PS> MFTECmd.exe -f C:\Image\$MFT --csv .\out --csvf mft.csv
PS> LogFileParser.exe -f C:\Image\$LogFile -o logfile.csv
PS> UsnJrnl2Csv.exe -f C:\Image\$J -o usnjrnl.csv
```

A.6 ext4 analysis

```
$ sudo extundelete --restore-file 'path/file' /dev/sdb1
$ sudo extundelete --restore-all /dev/sdb1

$ sudo debugfs /dev/sdb1
debugfs: lsdel
debugfs: stat <12345>
debugfs: dump <12345> /tmp/out
```

A.7 Linux RAID (mdadm)

```
$ sudo mdadm --examine /dev/sd[a-d]1
$ sudo mdadm --assemble /dev/md0 /dev/sd[a-d]1
$ sudo mdadm --assemble --force --run /dev/md0 /dev/sd[a-d]1
$ sudo mdadm --detail /dev/md0
$ cat /proc/mdstat
```

A.8 Btrfs and snapshots (NAS)

```
$ sudo mount -t btrfs -o ro,recovery /dev/md0 /mnt/nas
$ sudo btrfs subvolume list -s /mnt/nas
$ sudo btrfs send /mnt/nas/.snapshots/123 \
| sudo btrfs receive /restoration
```

A.9 TRIM verification

```
# Windows
C:\> fsutil behavior query DisableDeleteNotify

# Linux
$ cat /sys/block/sdX/queue/discard_max_bytes
$ systemctl status fstrim.timer

# macOS
$ system_profiler SPSerialATADataType | grep -i 'TRIM Support'
```

A.10 RAM capture (live forensics)

```
# Windows: DumpIt (MoonSols/Comae)
C:\> DumpIt.exe

# Linux: LiME or AVML
$ sudo insmod lime.ko 'path=/tmp/mem.lime format=lime'

# Volatility 3 analysis
$ vol -f memory.dump windows.info
$ vol -f memory.dump windows.pslist
$ vol -f memory.dump windows.netscan
```


A.11 Python script: ddrescue mapfile parser

Small utility that parses the mapfile of a ddrescue session and displays statistics by block category (rescued, non-tried, non-trimmed, non-scraped, slow, bad-sector):

```
#!/usr/bin/env python3
"""parse_ddrescue_mapfile.py - Statistics for a ddrescue mapfile."""
import sys
from collections import Counter

STATUS_LABELS = {
'+': 'rescued',
'?': 'non-tried',
'*': 'non-trimmed',
'/': 'non-scraped',
'-': 'bad-sector',
'F': 'finished',
'L': 'slow',
}

def parse(path):
totals = Counter()
with open(path) as f:
for line in f:
line = line.strip()
if not line or line.startswith('#'):
continue
parts = line.split()
if len(parts) < 3 or parts[2] not in STATUS_LABELS:
continue
size = int(parts[1], 16)
totals[parts[2]] += size
return totals

if __name__ == '__main__':
totals = parse(sys.argv[1])
grand = sum(totals.values())
print(f'Total: {grand/1e9:.2f} GB')
for status, size in totals.most_common():
label = STATUS_LABELS.get(status, status)
pct = size / grand * 100 if grand else 0
print(f' {label:<14} {size/1e9:>8.2f} GB ({pct:>5.2f} %)')
```

Usage: python3 parse_ddrescue_mapfile.py image.map. Useful to decide whether to launch an additional pass (if many *non-scraped*) or if the device is hopeless (many *bad-sector*).

A.12 Python script: extract deleted MFT entries

Small pedagogical parser that reads an extracted $MFT file and lists entries marked as deleted with their filename from the $FILE_NAME attribute:

```
#!/usr/bin/env python3
"""parse_mft_deleted.py - List deleted MFT entries."""
import sys, struct

RECORD_SIZE = 1024
MFT_RECORD_IN_USE = 0x01

def parse_mft(path):
deleted = []
with open(path, 'rb') as f:
record_idx = 0
```



```
while True:
data = f.read(RECORD_SIZE)
if len(data) < RECORD_SIZE:
break
if data[:4] != b'FILE':
record_idx += 1
continue
flags = struct.unpack('<H', data[22:24])[0]
if not (flags & MFT_RECORD_IN_USE):
# Deleted - look for FILE_NAME attribute (type 0x30)
name = find_file_name(data)
if name:
deleted.append((record_idx, name))
record_idx += 1
return deleted

def find_file_name(record):
# Offset of first attribute
attr_offset = struct.unpack('<H', record[20:22])[0]
while attr_offset < len(record) - 8:
attr_type = struct.unpack('<I', record[attr_offset:attr_offset+4])[0]
if attr_type == 0xFFFFFFFF:
return None
attr_len = struct.unpack('<I', record[attr_offset+4:attr_offset+8])[0]
if attr_len == 0:
return None
if attr_type == 0x30: # FILE_NAME
content_offset = struct.unpack('<H', record[attr_offset+20:attr_offset+22])[0]
base = attr_offset + content_offset
name_len = record[base + 64]
try:
return record[base+66:base+66+name_len*2].decode('utf-16-le')
except UnicodeDecodeError:
return None
attr_offset += attr_len
return None

if __name__ == '__main__':
for idx, name in parse_mft(sys.argv[1]):
print(f'{idx:>8} {name}')
```

Disclaimer — These scripts are pedagogical and minimal. For serious use, prefer MFTECmd (Zimmerman) on the NTFS side and the Sleuth Kit suite on the generic side. The scripts above are for understanding the structure, not for replacing battle-tested tools.

APPENDICES

Chapter B

Glossary

AES — Advanced Encryption Standard. Symmetric encryption algorithm standardized by NIST in 2001. AES-128 and AES-256 resist all known attacks with current classical means.

AFR — Annualized Failure Rate of a drive fleet.

Air-gap — Physical disconnection of a system or device from the network.

APFS — Apple File System (2017). Copy-on-write, snapshots, native FileVault support.

BGA — Ball Grid Array. Package with a grid of solder balls under the component, common for NAND chips.

Btrfs — Linux B-tree file system. Copy-on-write with snapshots and checksums.

Carving — Reconstruction of files from binary signatures in raw data, without using FS structures.

Chain of custody — Continuous documentation of the handling of digital evidence.

Charge trap (CTF) — Modern NAND cell architecture where charge is trapped in an insulator, more robust than floating gate. Adopted massively since 2017.

Chip-off — Physical desoldering of a NAND chip to read it independently of its controller.

CMR — Conventional Magnetic Recording. HDD mode with non-overlapping tracks.

Copy-on-write — Any modification writes elsewhere and updates pointers. Used by APFS, Btrfs, ZFS.

CPU Swap (DAFOTEC method) — Transplant of a smartphone processor onto a donor motherboard, preserving the Secure Enclave / TEE. Documented publicly by DAFOTEC.

DBIR — Data Breach Investigations Report. Annual Verizon report since 2008.

ddrescue — GNU ddrescue. Bit-by-bit imaging tool tolerant to errors.

DRAT / DZAT — Deterministic Read After Trim / Deterministic Zero After Trim. SSD guarantees.

ECC — Error-Correcting Code. Correction codes applied by the SSD controller to each NAND page.

eMMC — Embedded MultiMediaCard. Flash memory with integrated controller, compact format for low/mid range smartphones.

ext4 — Default Linux file system since 2008.

FileVault — macOS volume encryption since Mac OS X 10.3, then FileVault 2 since 10.7. Managed by the Secure Enclave on Apple Silicon Macs.

FTL — Flash Translation Layer. Layer in the SSD controller that maps logical LBAs to physical NAND pages.

Garbage collection (GC) — Background process of the SSD controller that physically erases blocks marked free.

HSA transplant — Head Stack Assembly transplant. Transplant of the complete head/arm assembly from a twin donor hard drive to a recipient drive, with sub-micron alignment under 0.3 microns.

HAMR — Heat-Assisted Magnetic Recording. HDD technology that heats the magnetic layer by laser during writing. Commercialized since 2024 on very large drives.

HDD — Hard Disk Drive. Mechanical magnetic disk.

Imaging — Creation of a bit-by-bit copy of a device to an image file.

ISO 14644-1 — International standard for cleanroom classes by particle concentration.

ISO 27037 — International standard for the collection and preservation of digital evidence.

JTAG / ISP — Internal debug protocol exposed by SSD controllers. Exploited in non-destructive recovery.

LBA — Logical Block Address exposed by the SATA/NVMe interface.

LDPC — Low-Density Parity-Check. Error-correcting code used by modern SSD controllers, succeeding classical BCH codes.

LUKS — Linux Unified Key Setup. Standard Linux volume encryption.

Mapfile — Map file used by ddrescue.

MFT — Master File Table. Central structure of NTFS.

NAND flash — Floating-gate or charge-trap memory, technology under all SSDs, eMMC, UFS, microSD, USB sticks.

NTFS — New Technology File System. Windows FS since NT.

Over-provisioning — NAND space reserved by the SSD controller, invisible to the user.

PCB — Printed Circuit Board. External board of an HDD or SSD.

PC-3000 — ACE Lab hardware platform, de facto standard for professional recovery.

PLC — Penta-Level Cell. 5 bits per cell NAND memory, 32 voltage levels. Announced 2023-2025, progressive commercialization.

PMR — Perpendicular Magnetic Recording. HDD mode generalized since 2005.

RAID — Redundant Array of Independent Disks. Combination of several disks for performance or resilience.

Ransomware — Malicious software that encrypts files and demands a ransom.

Read-retry — NAND read technique consisting of shifting voltage thresholds to recover a page that the initial read could not decode.

Reverse FTL — Software reconstruction of the LBA-NAND mapping table of an SSD with a failed controller, from NAND page metadata.

Cleanroom — Controlled-atmosphere room defined by ISO 14644-1. For HDD recovery: ISO Class 5.

Secure Enclave — Apple security coprocessor (iPhone 5s+, Mac T2, Apple Silicon) that manages cryptographic keys.

SED — Self-Encrypting Drive. SSD that automatically encrypts via its controller, TCG Opal standard.

Service Area (SA) — Reserved area on HDD platters, invisible to the OS, containing over a hundred firmware modules (P-list, G-list, translator, adaptives).

SMR — Shingled Magnetic Recording. HDD mode with partially overlapping tracks.

Snapshot — Instant image of a volume at near-zero cost on copy-on-write FS.

Spider Web (DAFOTEC method) — NAND extraction procedure on monolith devices (microSD, all-in-one USB sticks). Laser abrasion of the resin then micro-soldering of 15 to 30 fine copper wires. Documented publicly by DAFOTEC.

TCG Opal — Trusted Computing Group standard for SSD hardware encryption.

TEE — Trusted Execution Environment. Android secure coprocessor (equivalent to Secure Enclave).

Translator — HDD firmware module that converts LBA to physical coordinates (cylinder, head, sector). Module 028 on WD ROYL.

TRIM — ATA command (DATA SET MANAGEMENT) or NVMe (DEALLOCATE) that informs the SSD controller of LBAs freed on the OS side.

UFS — Universal Flash Storage. eMMC successor for high-end smartphones.

VeriFiles (DAFOTEC service) — Service: the client reviews the complete list of recoverable files before any payment.

Wear leveling — SSD strategy that spreads writes across all NAND cells.

WORM — Write Once Read Many. Immutable storage.

Write blocker — Hardware or software device that blocks any writes to a device, for forensic integrity.

ZFS — Sun Microsystems file system (2006), now OpenZFS. Copy-on-write, checksums, self-healing.

APPENDICES

Chapter C

Bibliography

Public sources actually consulted for this manual (May 2026). Simplified URLs (https:// and www. prefixes omitted).

Industry reports

Verizon Business	2025 Data Breach Investigations Report. April 2025.	verizon.com/about/news/2025-data-breach-investigations-report
Backblaze	Drive Stats for 2025. February 2026.	backblaze.com/blog/backblaze-drive-stats-for-2025/
IBM Security	Cost of a Data Breach Report 2024.	ibm.com/reports/data-breach

Standards

ISO	ISO 14644-1:2015 — Cleanrooms classification.	iso.org/standard/53394.html
ISO	ISO/IEC 27037:2012 — Digital evidence guidelines.	iso.org/standard/44381.html
Trusted Computing Group	TCG Storage Opal 2.0.	trustedcomputinggroup.org/resource/storage-work-group-storage-security-subsystem-class-opal/
US Courts	Federal Rules of Evidence, Rule 902(14).	uscourts.gov/rules-policies/current-rules-practice-procedure/federal-rules-evidence

Cleanroom and HDD recovery

DriveSavers	Certified ISO Class 5 Cleanroom.	drivesaversdatarecovery.com/why-us/certified-iso-class-5-cleanroom/
Rossmann Group	CMR vs SMR: How Recording Technology Affects Recovery.	rossmanngroup.com/technical-reference/cmr-vs-smr-hard-drives
ACE Lab	PC-3000 public documentation HDD/SSD modules.	acelab.eu.com/

SSD, NAND, TRIM, FTL

Rossmann Group	What TRIM Does and Why It Destroys Data.	rossmanngroup.com/technical-reference/what-trim-does-and-why-it-destroys-data
Seagate	What Are SSD TRIM and Garbage Collection?	seagate.com/blog/what-are-ssd-trim-and-garbage-collection/
Kingston	The Importance of Garbage Collection and TRIM.	kingston.com/en/blog/pc-performance/ssd-garbage-collection-trim-explained
Belkasoft / Forensic Focus	Recovering Evidence from SSD Drives.	forensicfocus.com/articles/recovering-evidence-from-ssd-drives-in-2014/

File systems

Sygnia	Forensic Value of MFT Slack Space. 2025.	sygnia.co/blog/the-forensic-value-of-mft-slack-space/
Brian Carrier	File System Forensic Analysis (reference book).	sciencedirect.com/topics/computer-science/master-file-table

Tools

GNU	GNU ddrescue manual.	gnu.org/software/ddrescue/
CGSecurity	TestDisk and PhotoRec.	cgsecurity.org/wiki/TestDisk
Sleuth Kit	TSK and Autopsy.	sleuthkit.org/
Eric Zimmerman	Forensic tools.	ericzimmerman.github.io/
SANS Institute	DFIR papers and posters.	sans.org/white-papers/

Backup

Veeam	3-2-1 Backup Rule Explained.	veeam.com/blog/321-backup-rule.html
Object First	3-2-1-1-0 Backup Rule.	objectfirst.com/blog/how-object-first-and-veeam-bring-3-2-1-1-0-to-life/

US regulatory

HHS	HIPAA Security Rule.	hhs.gov/hipaa/for-professionals/security/
SEC	Sarbanes-Oxley Act compliance.	sec.gov/about/laws/soa2002.pdf
California OAG	CCPA / CPRA full text.	oag.ca.gov/privacy/ccpa
FedRAMP	Federal Risk and Authorization Management Program.	fedramp.gov/

Historical case studies

Wired	The Untold Story of NotPetya. 2018.	wired.com/story/notpetya-cyberattack-ukraine-russia-code-crashed-the-world/
Threatpost	Code Spaces Out of Business. June 2014.	threatpost.com/hacker-puts-hosting-service-code-spaces-out-of-business/106761/
InfoWorld	Murder in the Amazon cloud. 2014.	infoworld.com/article/2179073/murder-in-the-amazon-cloud.html
breaches.cloud / Wiz	Codespaces (2014).	breaches.cloud/incidents/codespaces/

DAFOTEC sources

DAFOTEC	Official site: pricing, technical expertise, real case studies, institutional clients.	dafotec.fr
DAFOTEC Belgium	Belgian version of the laboratory.	dafotec.be

APPENDICES

Chapter D

About DAFOTEC

DAFOTEC is a French data recovery laboratory founded in Roubaix in 2004. This manual was written by Mhessan Kouassi, senior expert at DAFOTEC since the lab's founding, drawing on 22 years of field practice and more than 120,000 cases handled.

Identity

- **Laboratory**: 59 Bis rue du Curoir, CS 40082, 59052 Roubaix Cedex (France).
- **Sites**: dafotec.fr (France) and dafotec.be (Belgium).
- **Age**: 22 years (2004 — 2026).
- **Volume handled**: over 120,000 cases since 2004.
- **Coverage**: 36 drop-off centers in metropolitan France, national pickup service, single analysis laboratory in Roubaix.

Certifications and compliance

- **ISO 5 cleanroom** (ISO 14644-1 standard) in Roubaix.
- **GDPR** compliance for personal data handling.
- **ISO 27001** compliance for information security.
- Handling of sensitive devices on a physically Internet-isolated network (**air gapped**).
- For US clients: NDA available, BAA on request for HIPAA-covered devices.

Institutional clients

DAFOTEC is regularly retained by French public organizations and scientific institutions for handling failed devices, under non-disclosure agreements (NDA). Clients cited with their authorization:

- French National Police (Gendarmerie Nationale)
- Tourcoing University Hospital
- CNRS (French National Center for Scientific Research)
- INSERM (French National Institute of Health and Medical Research)
- French universities

Business model

- **Free diagnosis** within 24 hours at the laboratory.
- **Payment on results**: the recovery flat fee is invoiced only in case of success. In case of failure, only a 25 EUR reconditioning and return fee is due.
- **VeriFiles service**: the complete list of recoverable files is provided to the client *before any payment*. The client validates this list then decides to accept or refuse the quote without charge.

- **Public pricing**: 300 EUR for a hard drive, 400 EUR for an SSD, 300 EUR for a smartphone, 550 EUR per disk for a NAS, 700 EUR per disk for a RAID. Up to 950 EUR for the most complex failures. Pricing reflects French market; US recovery labs typically charge 2-5x these amounts.

Independent reviews

4.9/5 on 797 verified Ekomi reviews as of publication of this manual (May 2026).

Why this manual is free

This book is distributed free of charge under the Creative Commons BY-NC-ND 4.0 license, downloadable without registration from dafotec.fr and dafotec.be.

The reason is consistent with the editorial mission of the manual: reduce misinformation in the data recovery sector; save individuals, SMBs and administrations from the wrong first decisions that make recovery impossible; provide technicians and forensics students with a solid and sourced reference.

DAFOTEC is a commercial laboratory, but the effort of popularization and technical transparency has value in itself. If this manual prevents one bad manipulation, its writing is justified. If it guides a few readers toward the right reflexes (tested backup, diagnosis before tool, lab in case of doubt), it is even better.

Distribution license

End of manual. May 2026.
DAFOTEC — Roubaix — since 2004.

APPENDICES

Chapter E

Thematic index

Alphabetical index of the main technical terms covered in the manual, with their page occurrences. For precise definitions, see also Appendix B (Glossary).

F

G

H

I

J

L

M

N

P

R

S

T

U

V

W

Z

