

# Squinting over the tape moat

An exceedingly brief look at magnetic tape storage  
from a security perspective

---

Richard Bachmann

08.04.2025

Why are we talking about magnetic tape?

## Why are we talking about magnetic tape?



**Department of Government Efficiency**  @DOGE · Apr 4



The @USGSA IT team just saved \$1M per year by converting 14,000 magnetic tapes (70 yr old technology for information storage) to permanent modern digital records.

 1.2K

 4.4K

 28K

 1.5M



## Why are we talking about magnetic tape?




- Your institution and/or cloud provider are using it
- Archived data is super interesting
  - Important data to keep
  - Personal information
  - Backup data
- Tape is a solid last line of defense against ransomware attacks

## Tape Media / Cartridge


- Present generation holds 18–50TB each
- Capacity doubles every couple of years
- Sequential access on ~1 cm wide but ~1km long surface area
  - Time to first byte requires seeking but then the data starts to flow fast
- Format may be self describing (LTFS) or not
  - At scale often the latter

CARTRIDGE



LINEAR SERPENTINE LAYOUT

TAPE LENGTH




BOUSTROPHEDON

FILE 31352682: /VAR/LOG/NGINX/ERROR.LOG

# Tape Drive

- Reads/writes tape, up to 400MB/s
- Hardware for a Tape Server/Data Mover
- Capable of performing compression and encryption




TAPE SERVER




TAPE DRIVE

# Tape Library

- Physical container for tape media and drives
- Inactive tapes are stored in a shelf  $\Rightarrow$  air gap
- One (maybe two) robots inside, moving tapes between slots and drive
  - Limits how fast tapes can go to active state
- May be partitioned into Logical Libraries



# Putting the pieces together



Security!

SECURITY!

---

## Media – The data you deleted is probably still there

- Overwriting tape takes time
- Delete file != Data gone
  - In most systems it is just a metadata operation

## Media – The data you deleted is probably still there

- Overwriting tape takes time
- Delete file != Data gone
  - In most systems it is just a metadata operation

### Remedy

- Encrypt your data by default
- Use Lethe [1]
- Be sure that retired tapes are wiped/destroyed

## Media – Overwriting the beginning means data is gone

- Overwriting the Beginning of Data (BoD) = all data gone
- Relatively quick operation,  
 $T_{queue} + T_{mount} + T_{write\_a\_few\_bits}$
- Only the vendor, or your local friendly 3-letter agency, has the means to recover data after the End of Data (EoD) mark

## Media – Overwriting the beginning means data is gone


- Overwriting the Beginning of Data (BoD) = all data gone
- Relatively quick operation,  
 $T_{queue} + T_{mount} + T_{write\_a\_few\_bits}$
- Only the vendor, or your local friendly 3-letter agency, has the means to recover data after the End of Data (EoD) mark

### Remedy

- Restrict access, especially to full tapes
- Monitor access patterns
- Read-only logical libraries for full tapes

# Media – Example: Read-only logical library

## PHYSICAL LIBRARY



## Encryption – Hostile encryption key switch

Patient attacker's key switch  
attack:

1. Compromise backup machines
2. Switch your key(s) for one of their own just as you write or do a backup
3. Switch back right after the process starts
4. Repeat until retention policy is exceeded (maybe 6 months?)

⇒ Backup data unreadable, except for the attacker

## Encryption – Hostile encryption key switch

Patient attacker's key switch attack:

1. Compromise backup machines
2. Switch your key(s) for one of their own just as you write or do a backup
3. Switch back right after the process starts
4. Repeat until retention policy is exceeded (maybe 6 months?)

⇒ Backup data unreadable, except for the attacker

### Remedy

- Outsource key management to users, make them practice restores from backup
- Monitor key source (DB, file, etc.)
- Take samples, try to read them
- Use distributed architectures
  - Make it hard to change everywhere
  - Make it break if discrepancies are found

## The DB – Your only metadata source?

- Tape data is often not self-describing. We may only know where files begin and end.
- A database is used to keep track of what is where
  - It would be a shame if something were to happen to it...

## The DB – Your only metadata source?

- Tape data is often not self-describing. We may only know where files begin and end.
- A database is used to keep track of what is where
  - It would be a shame if something were to happen to it...

### Remedy

- Make someone keep DB backups
- Make sure you aren't keeping the DB backups
  - Or if you are, put them on self-describing protected/Safeguarded tapes

## Summary and Q&A

- Tape infrastructure is likely to carry important data
- Tape storage is different to disk due to physical properties
  - And that happens to make it good for secure storage
- Patient malicious actors can and will figure out tape, so you should too!

### Acknowledgements:

Special thanks to Vladimír Bahyl for profound tape insight and patience in sharing it!

# References

[1] Eugene Chou et al. *Lethe: Secure Deletion by Addition*. 2023.

## Images:

- Slide 3 tape drive: LTO Consortium <https://www.lto.org/wp-content/uploads/2023/01/Technology-and-Sustainability-%E2%80%93-LTO-Saving-the-Planet.pdf>
- Slide 4 tape library 1: CERN, <https://cta.web.cern.ch/cta/>, 2023
- Slide 4 tape library 2: Hackers, 1995