CS 798: Digital Forensics and Incident Response

Lecture 18 - Residue-free Computing

Diogo Barradas

Winter 2025

University of Waterloo

What is this, and why do we care?

- Criminals may attempt to use anti-forensic solutions to better hide their activities in computer systems
- A loose definition of **residue-free computing**: operate any existing application installed on a computer in a way that prevents data from being recorded to disk, thus frustrating the forensic process and enabling privacy-preserving computing
- We will also cover **plausibly-deniable storage**, which aims to obscure the existence of entire filesystems

Outline

1. Disk Encryption
   - Full disk encryption
   - Plausibly-deniable storage
2. Residue-free computing
   - Live CDs and VMs
   - System-wide incognito mode

Disk Encryption

Disk Encryption

Full disk encryption

Full disk encryption

- **Full disk encryption (FDE):** encrypts users' and OS data in a disk to prevent unauthorized access
- There is a growing use (and offer) of FDE software
  - Integrated with OS, e.g., BitLocker, FileVault
  - Commercial tools, e.g., McAfee's Safeboot, GuardianEdge
- Also available for portable storage media

How can FDE hamper digital investigations?

- Full disk encryption can significantly hamper digital investigations, potentially preventing access to digital evidence
  - It will be difficult to recover evidence if the computer is shut down by first responders
  - It will waste investigators' time during early forensic triages
  - Investigators may lose the opportunity to search for passphrases written down on physical notes (typically used as FDE password recovery mechanism)

How can investigators tackle FDE?

- Preserving the live memory (RAM) of the machine becomes fundamental to any digital forensics investigation
  - Since FDE implementations typically decrypt data on the fly, it may be possible to recover FDE encryption keys from memory
- Enact more fitting legal approaches
  - Provide investigators with the ability to "surprise" the targets of a warrant (e.g., search warrants that permit surprise entry)
- In general, investigation procedures must be revised to reduce opportunities for anyone to power off the target system

Disk Encryption

Plausibly-deniable storage

XCKD 538

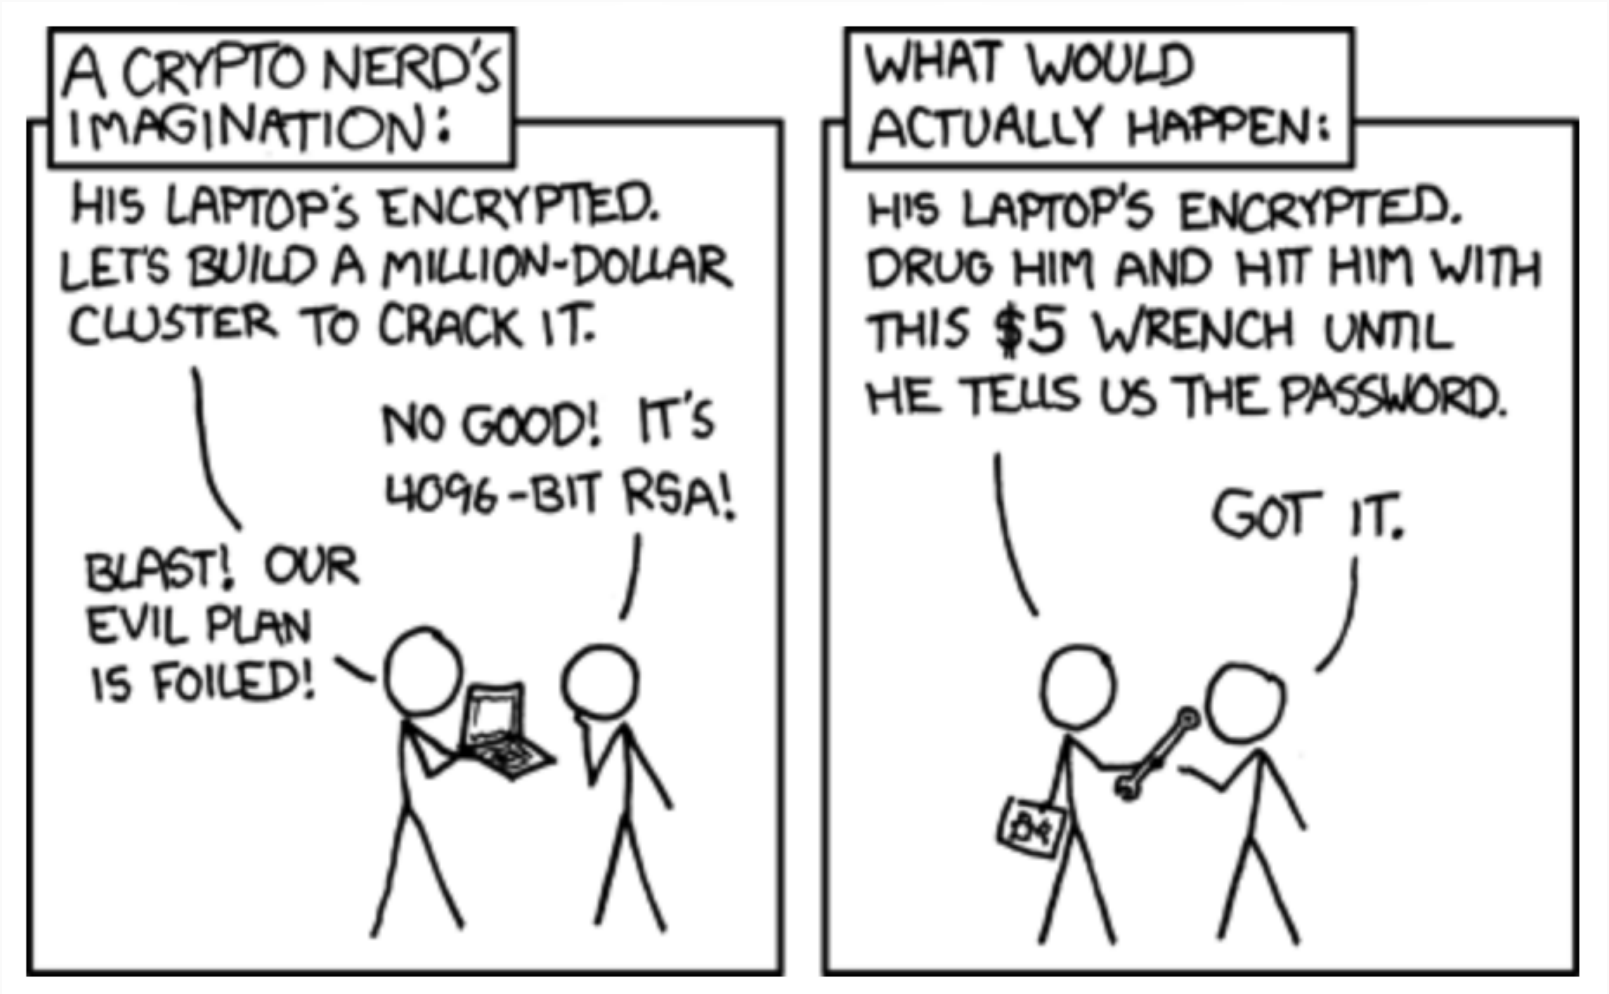
A CRYPTO NERD'S IMAGINATION:
HIS LAPTOP'S ENCRYPTED. LET'S BUILD A MILLION-DOLLAR CLUSTER TO CRACK IT.
NO GOOD! IT'S 4096-BIT RSA!
BLAST! OUR EVIL PLAN IS FOILED!
WHAT WOULD ACTUALLY HAPPEN:
HIS LAPTOP'S ENCRYPTED. DRUG HIM AND HIT HIM WITH THIS $5 WRENCH UNTIL HE TELLS US THE PASSWORD.
GOT IT.

What if we can be more selective on which data to encrypt?

- FDE systems are all-or-nothing: they either encrypt the whole disk or none of it
- Other software solutions allow for the creation of virtual filesystems that can remain hidden
  - These are known as steganographic file systems or plausibly-deniable storage
- **Steganographic file systems**: File systems that have multiple layers of files, with a given layer being exposed by providing its corresponding key while all other layers remain hidden
  - e.g., TrueCrypt, VeraCrypt, ShuffleCake

`https://www.veracrypt.fr/en/Home.html`

VeraCrypt volume creation

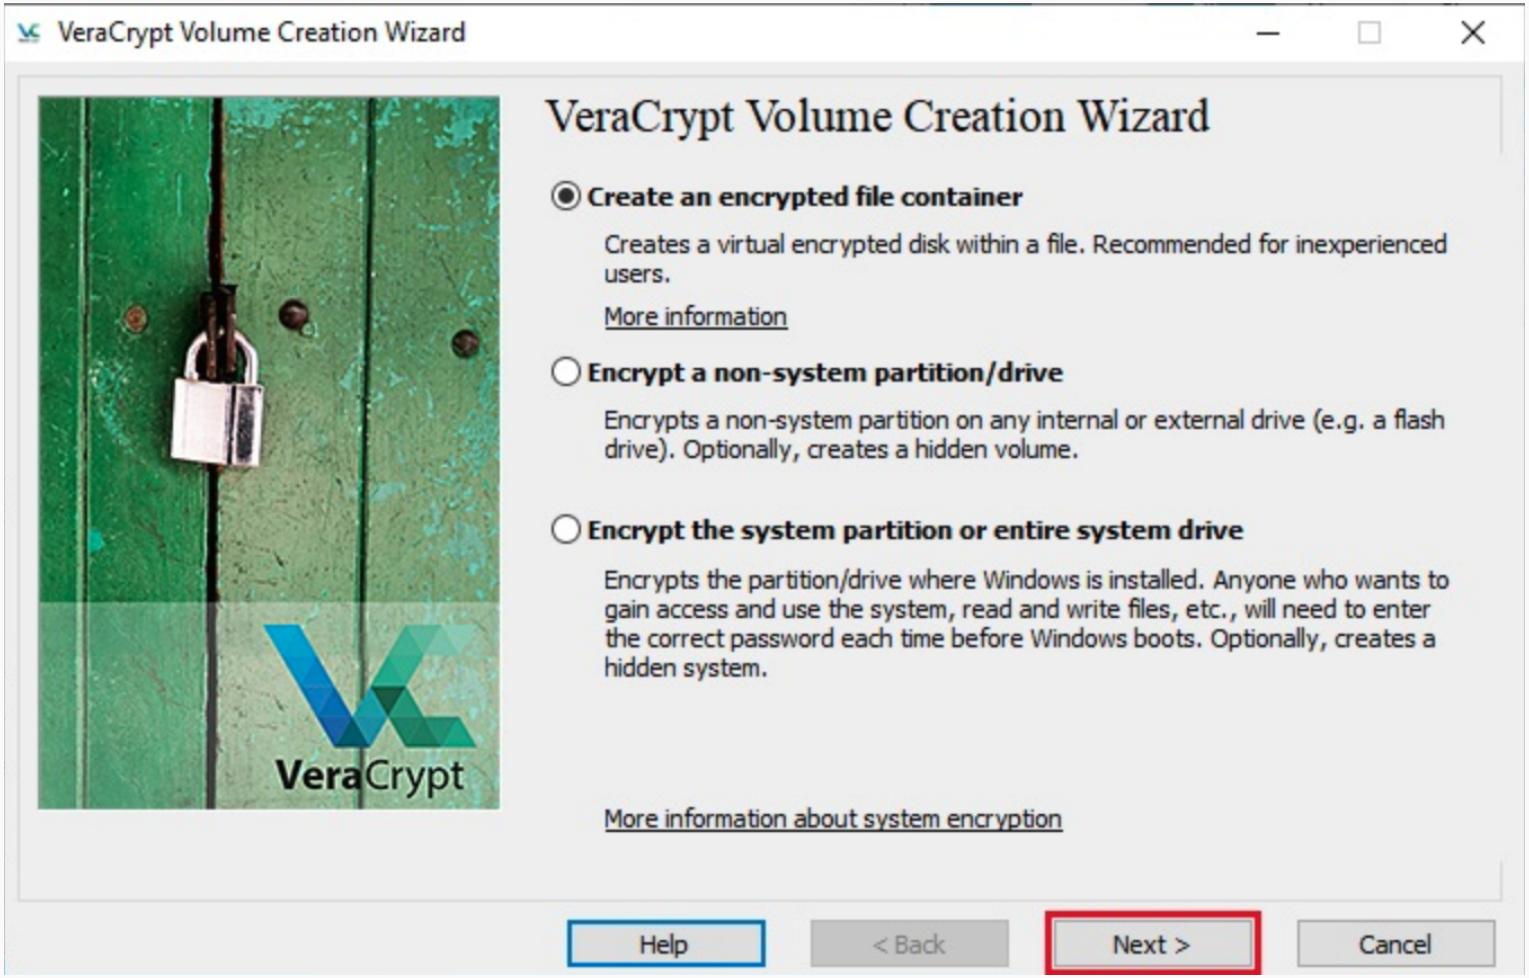

VeraCrypt volume creation (2)

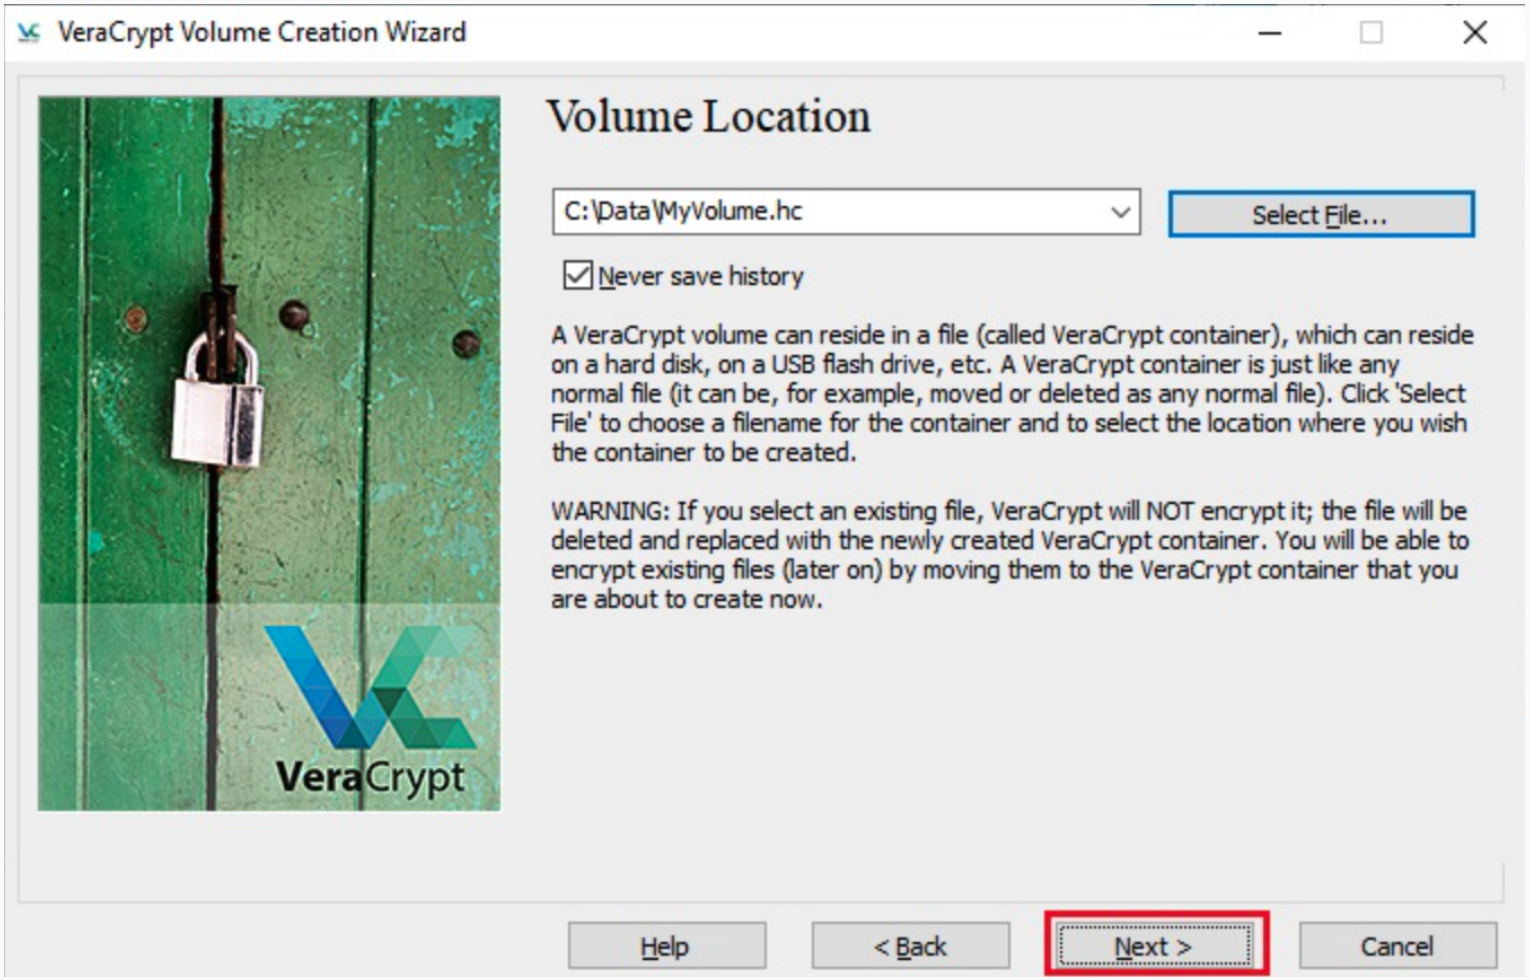

VeraCrypt volume creation (3)

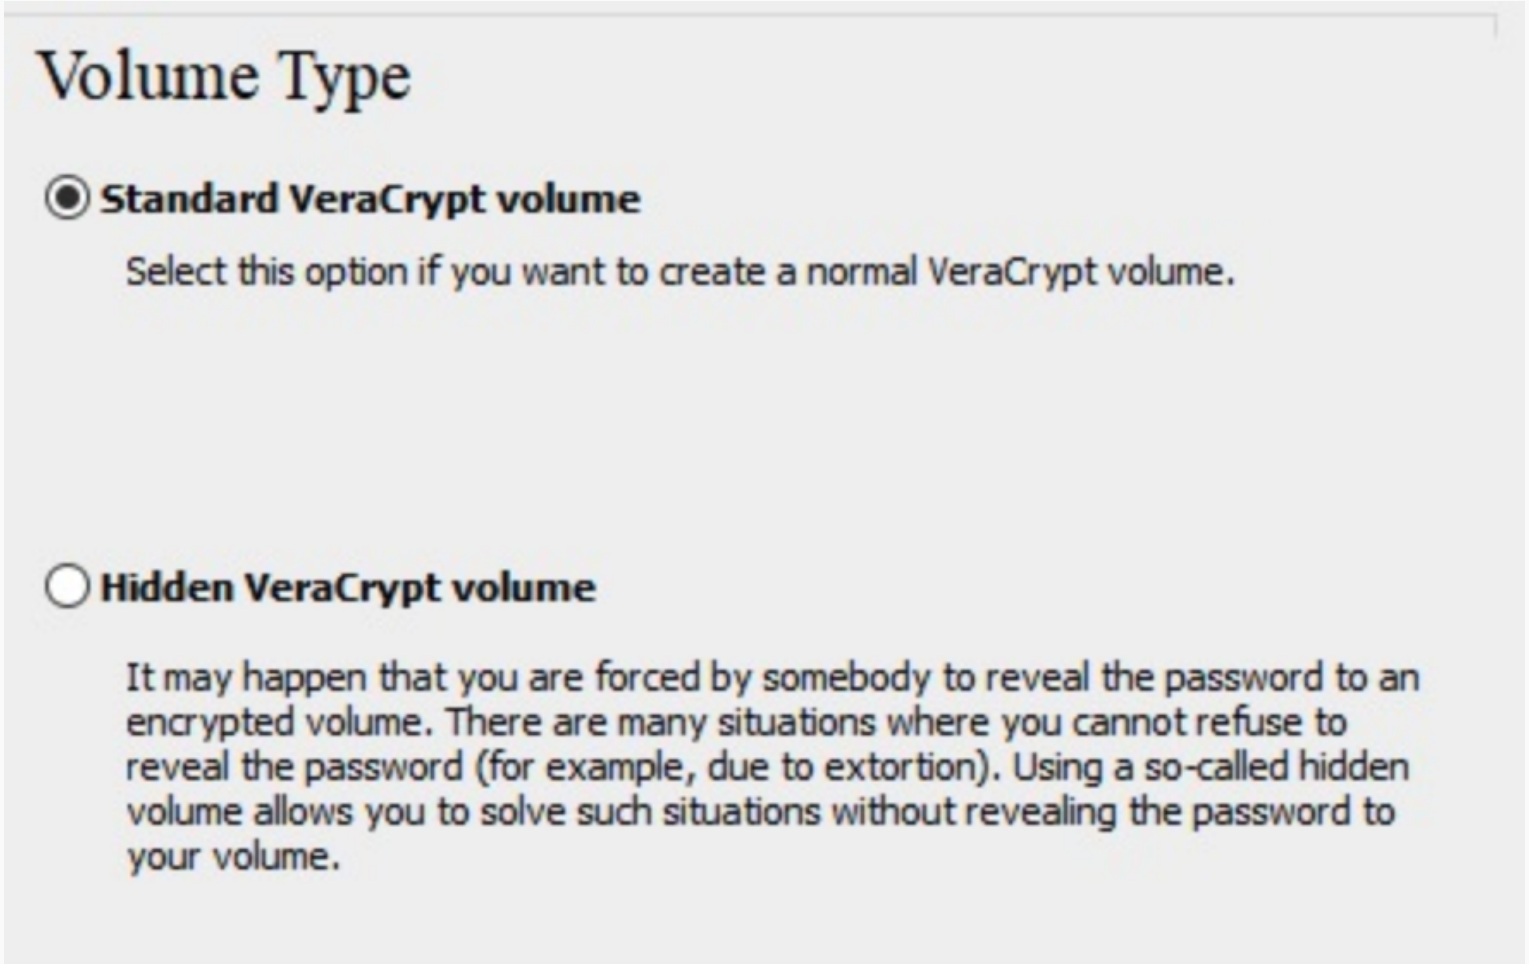

A look into VeraCrypt's plausibly-deniable storage

- Until decrypted, a VeraCrypt partition appears to consist of nothing more than random data
  - It turns out that wiping (i.e, securely erasing) your disk often fills it with random data
- The idea then is that it should be impossible to prove that a given partition is a VeraCrypt volume

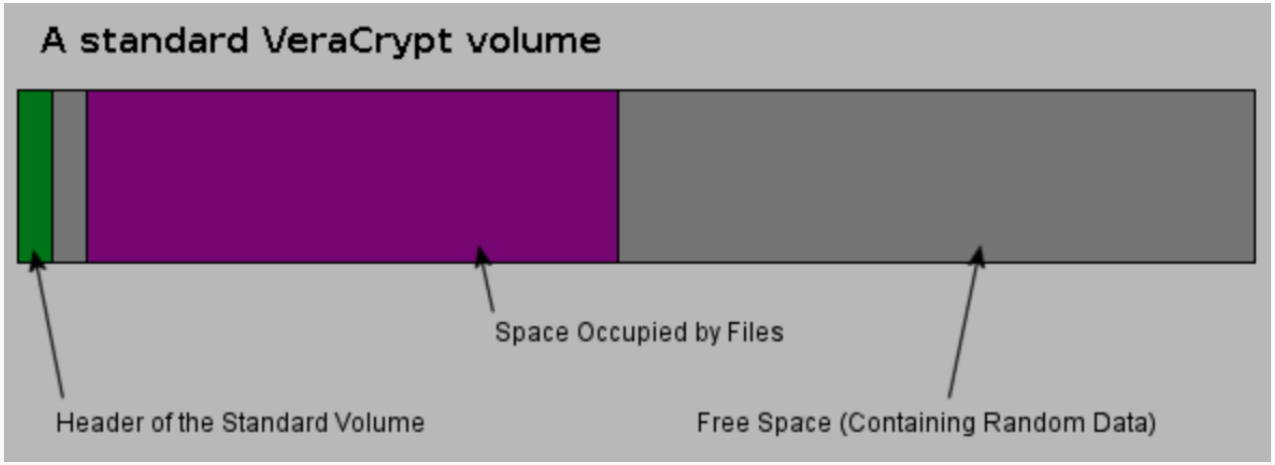

Coercion resistance

- VeraCrypt also supports hidden volumes (and OSes)
- **Idea**: Create a VeraCrypt volume within (the free space of) another VeraCrypt volume
- It should be impossible to prove whether there is a hidden volume not, because free space on any VeraCrypt volume is always filled with random data
- Under coercion (or to seem cooperative), a target user could simply reveal the password for the outer volume
    - Investigators have no way to know the truth

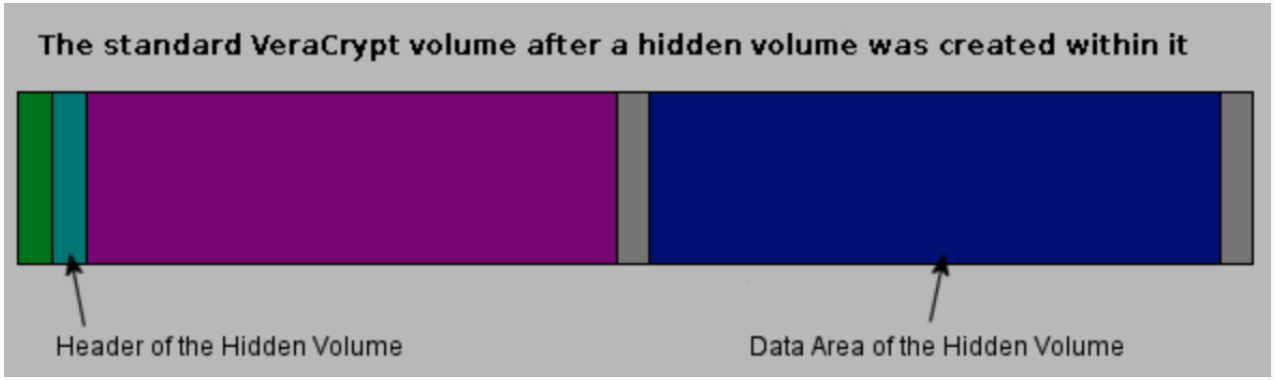

Coercion resistance (Hidden OSes)

- An unencrypted copy of the VeraCrypt Boot Loader has to be stored on the system drive
- The existence of the decoy OS is not secret
- You should use the decoy operating system as frequently as you use your computer
  - if you revealed the password for the decoy operating system to an adversary, he could find out that the system is not used very often, which might indicate the existence of a hidden operating system on your computer

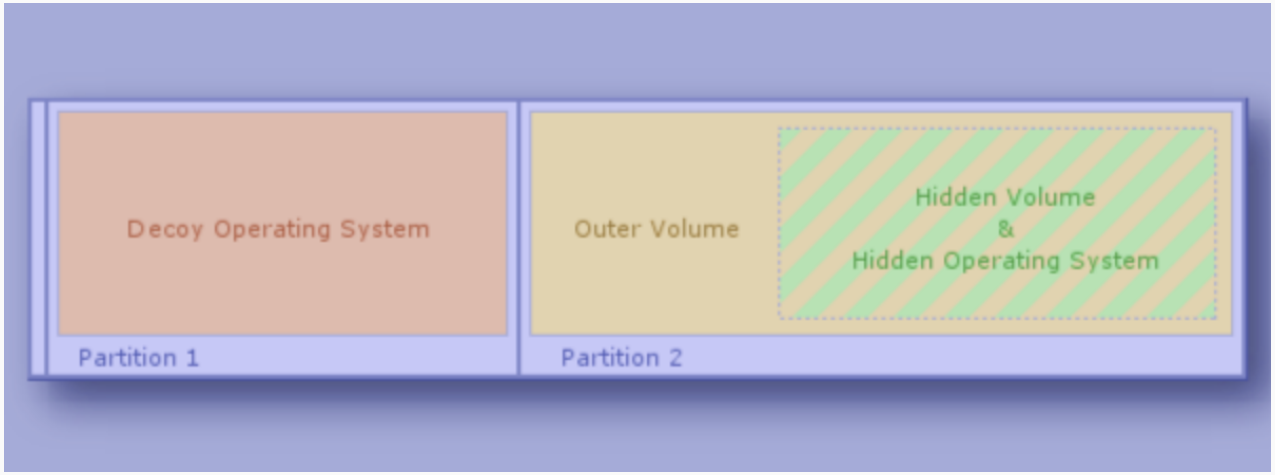

Multi-snapshot security

- The previous mechanism only works if investigators can inspect the device a single time (a.k.a., single-snapshot security)
- In multi-snapshot security), an investigator may access a VeraCrypt volume at several points over time
  - If a user changed the contents of a hidden volume, the contents of sectors (ciphertext) in the hidden volume area will be different across observations
  - After being given the password to the outer volume, investigators might wonder why these sectors changed.

ORAM-based plausible-deniable storage

- Multi-snapshot security requires mechanisms that hide users access patterns to hidden data
- **Oblivious RAM (ORAM)** techniques ensure that a database server cannot determine which database entries are accessed by a client
  - ORAMs can hide both the locations and contents of each access to a storage volume via high-overhead randomization
- Examples:
  - HIVE
  - DataLair
  - MobiCeal

A quick look into HIVE

- HIVE divides the storage into a public volume and a hidden volume, and both are accessed using an ORAM mechanism
- For every access to a volume (either public or hidden), the system also executes dummy accesses to the other volume.
- Since ORAM accesses are indistinguishable from each other (whether dummy or not), investigators cannot tell the difference between:
  1. accesses to the public volume, and
  2. accesses to the hidden volume
     - Essentially, a user could argue they only ever write to the public volume

Plausible-deniable storage based on canonical forms

- Randomization is not entirely necessary to obtain plausible deniability
- Other solutions involve the use of canonical forms like writing data sequentially in a circular buffer fashion, e.g., such as in a log-structured file system.
- These are enough to decouple a user's logical from physical access patterns (and retain data locality)
- Examples:
  - PD-DM
  - ECD

A quick look into PD-DM

- PD-DM leverages an append-only filesystem
- In PD-DM, whenever a public data record is written on the device, an additional random string (i.e., the "payload") is written in the immediately adjacent next block
- To store hidden data, PD-DM will first encrypt it and then write it as the payload of some public data write

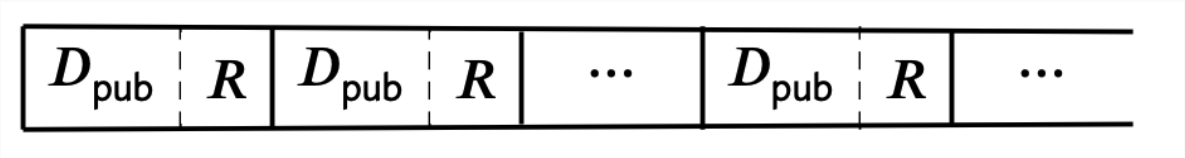

Residue-free computing

What if we could avoid writing to disk altogether?

- Apart from specific user files, a computer system might store a large amount of data to disk, without a user's knowledge
  - Application data, logs, file access info, etc.
- **Residue-free computing** techniques can avoid the writing of data to disk altogether, by running applications and entire OSes in the live memory of a system
- We have three main "footprint minimization" techniques to accomplish this:
  - Live CDs
  - Virtual machines
  - Systems like ResidueFree

Residue-free computing

Live CDs and VMs

What is a Live CD?

- A **live CD** typically refers to operating systems that boot from a removable device
- Privacy-focused live CDs, like Tails, avoid any writes to non-volatile storage and route all traffic through the Tor anonymity network

`https://tails.net/index.en.html`

How to limit data storage with a VM

- Virtual machines allow users to potentially destroy traces of past activities by reverting to an earlier snapshot
- Lightweight VMs (like Docker containers) can provide application isolation and ensure that filesystem modifications are not directly applied to the system's underlying filesystem

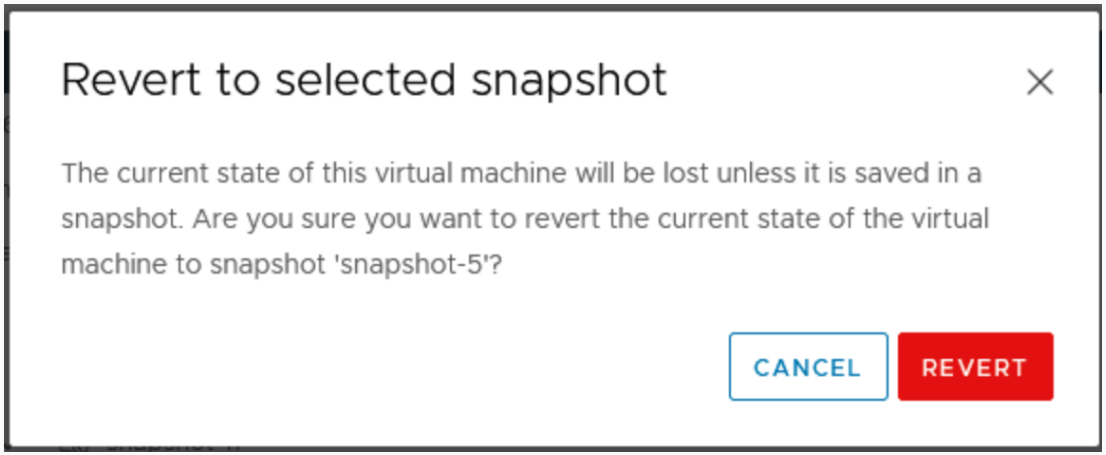

Drawbacks of Live CDs and VMs

- Erasure of past state of a VM relies on the secure erasure of the unsaved machine state; in practice it may be risky
- Containers can provide a level of indirection, but a containerized filesystem is still lodged on the main filesystem
- Live CDs and VMs require users to keep two computing environments: ordinary use and private computing
  - Switching between one and the other brings a significant usability burden

Residue-free computing

System-wide incognito mode

The residue-free computing vision

- **Idea**: to create an anti-forensics mode of operation for operating systems without requiring kernel modifications
- Users have access to their normal OS and existing files, and can run any installed program in residue-free mode
- The application has read access to all existing data on the filesystem, but file modifications (including deletions) are made to a volatile filesystem stored in memory (i.e., RAM).
    - Upon exiting the application, the filesystem modifications are permanently erased

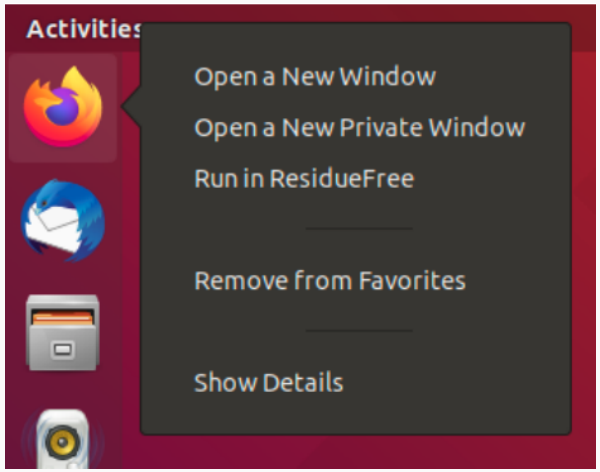

ResidueFree's architecture

- Three main components: a volatile RAM disk, a union filesystem, and an isolated execution environment.
- The union filesystem takes the volatile RAM disk and merges it with the user's filesystem
  - allows the application running in residue-free mode to appear to have read-write access to the user's regular filesystem
  - it can only read from the regular FS, but can write to the RAM FS

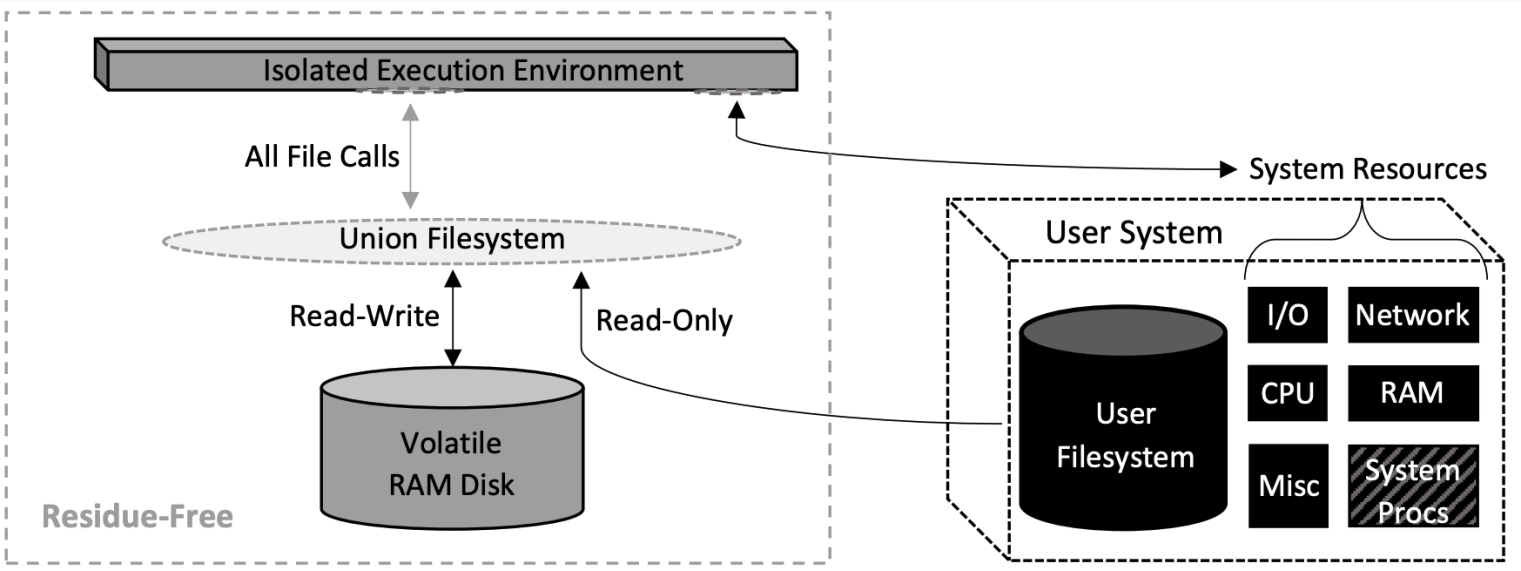

ResidueFree's limitations

- Focuses on leaving no residues on the filesystem
  - Data can still linger in RAM
  - RAM contents may be paged to disk
- Leaves system logs about ResidueFree's executions
- Lacks support for the execution of server applications

Takeaways

- The increasing use of disk encryption has detrimental implications for digital forensic investigations
- Plausibly deniable storage can make it hard for investigators to ascertain whether a target user is hiding the existence of data within a machine's physical storage. Yet, building secure and efficient plausibly deniable storage systems is a hard task.
- Residue-free computing techniques can avoid the writing of data to disk altogether

Pointers

- **Literature:**
  - The growing impact of full disk encryption on digital forensics
  - SoK: Plausibly Deniable Storage
  - Residue-Free Computing