



What happened to my file?

Tim Standing

Vice President Software Engineering - Mac
Other World Computing, Inc.







1. Data Corruption
2. Volume Corruption
3. Apple's T2 Chip



1. Data Corruption



**"...an average of one
single-bit-error every 14
to 40 hours per Gigabit
of DRAM."**



**"An older CERN study
marked their average as one
silent error in every 10^{16}
bits."**




"With an unrecoverable read error rate of 10^{-14} , you could expect a failed block read once every 12.5 TB or so."




**"..a four drive RAID5 - and
2TB drives - you would have
around a 40% chance of a
rebuild failure."**

Where can data corruption occur?



Where can data corruption occur?



"DRAM Errors in the Wild: A Large-Scale Field Study"



**Schoeder - University of Toronto
Pinheiro & Weber - Google
2009**

- **Average error rate of 2,000 - 6,000 per GB per year**
- **Over 8% of DIMMs exhibit errors**
- **Error rate highly dependent on motherboard design**



Testing for Data Corruption in RAM



Testing for Data Corruption in RAM



Testing for Data Corruption in RAM



Test Results:

- 12 Macs running for over 4 months
(iMac Pro, 2013 Mac Pro, 2018 Mac mini, 2011, 2013,
2105 and 2017 MacBook Pros)**
- 28 DIMMs total in all test computers**
- 15 PB data read, 15 PB data written**
- 0 Bits corrupted**

"Data Integrity"

Panzer-Steindel - CERN
2007

- 492 Servers with 1.5 PT of storage total
- Parity checked once per week
- 300 Incorrect parity blocks in 4 weeks



Testing for Data Corruption on Disk



Testing for Data Corruption on Disk



- **2 RAID 5 volumes with 15 Disks each**
- **First volume is active, constant reads and writes**
- **Second volume is passive, no reads or writes**
- **Every 10 days, verified that the parity information is correct**
- **Repeat for 6 months (18 parity verify operations)**
- **Only incorrect parity was on active volume and was result of power loss while writing to the volume**

"Data Integrity"

**Panzer-Steindel - CERN
2007**

- **Write 2 GB file to each of 3,000 servers**
- **Write file and verify every 2 hours**
- **360 TB total transferred in 5 weeks**
- **500 data corruption events observed**

Testing for Data Corruption While Reading or Writing



- **RAID 5 volume with 15 Disks each**
- **Write 1 GB random data files to fill volume, then read back and verify data**
- **4 Threads running continuously**
- **Continuous operation for 10 days at a time (6 months total)**
- **7.7 PB Transferred in 6 months (almost 1 billion i/os)**
- **Only incorrect data read was from power loss event**

"Why RAID 5 stops working in 2009"



Robin Harris, ZDNet (2007)

Uses disk specs to state that RAID 5 volumes will no longer be reliable with 2 TB disk

"...a 7 drive RAID 5 with 1 TB disks has a 50% chance of a rebuild failure. RAID 5 is reaching the end of its useful life."



"Why RAID 5 still works, usually"



Robin Harris, ZDNet (2016)


**Makes a case for RAID 5 volumes being
unreliable during rebuilds due to disk
Uncorrectable Error Rates.**





2. Volume Corruption








Apollo 13

Accident



NASA

Service Module





Figure 4-4.- Arrangement of fuel cells and cryogenic systems in bay 4.

Oxygen Tank

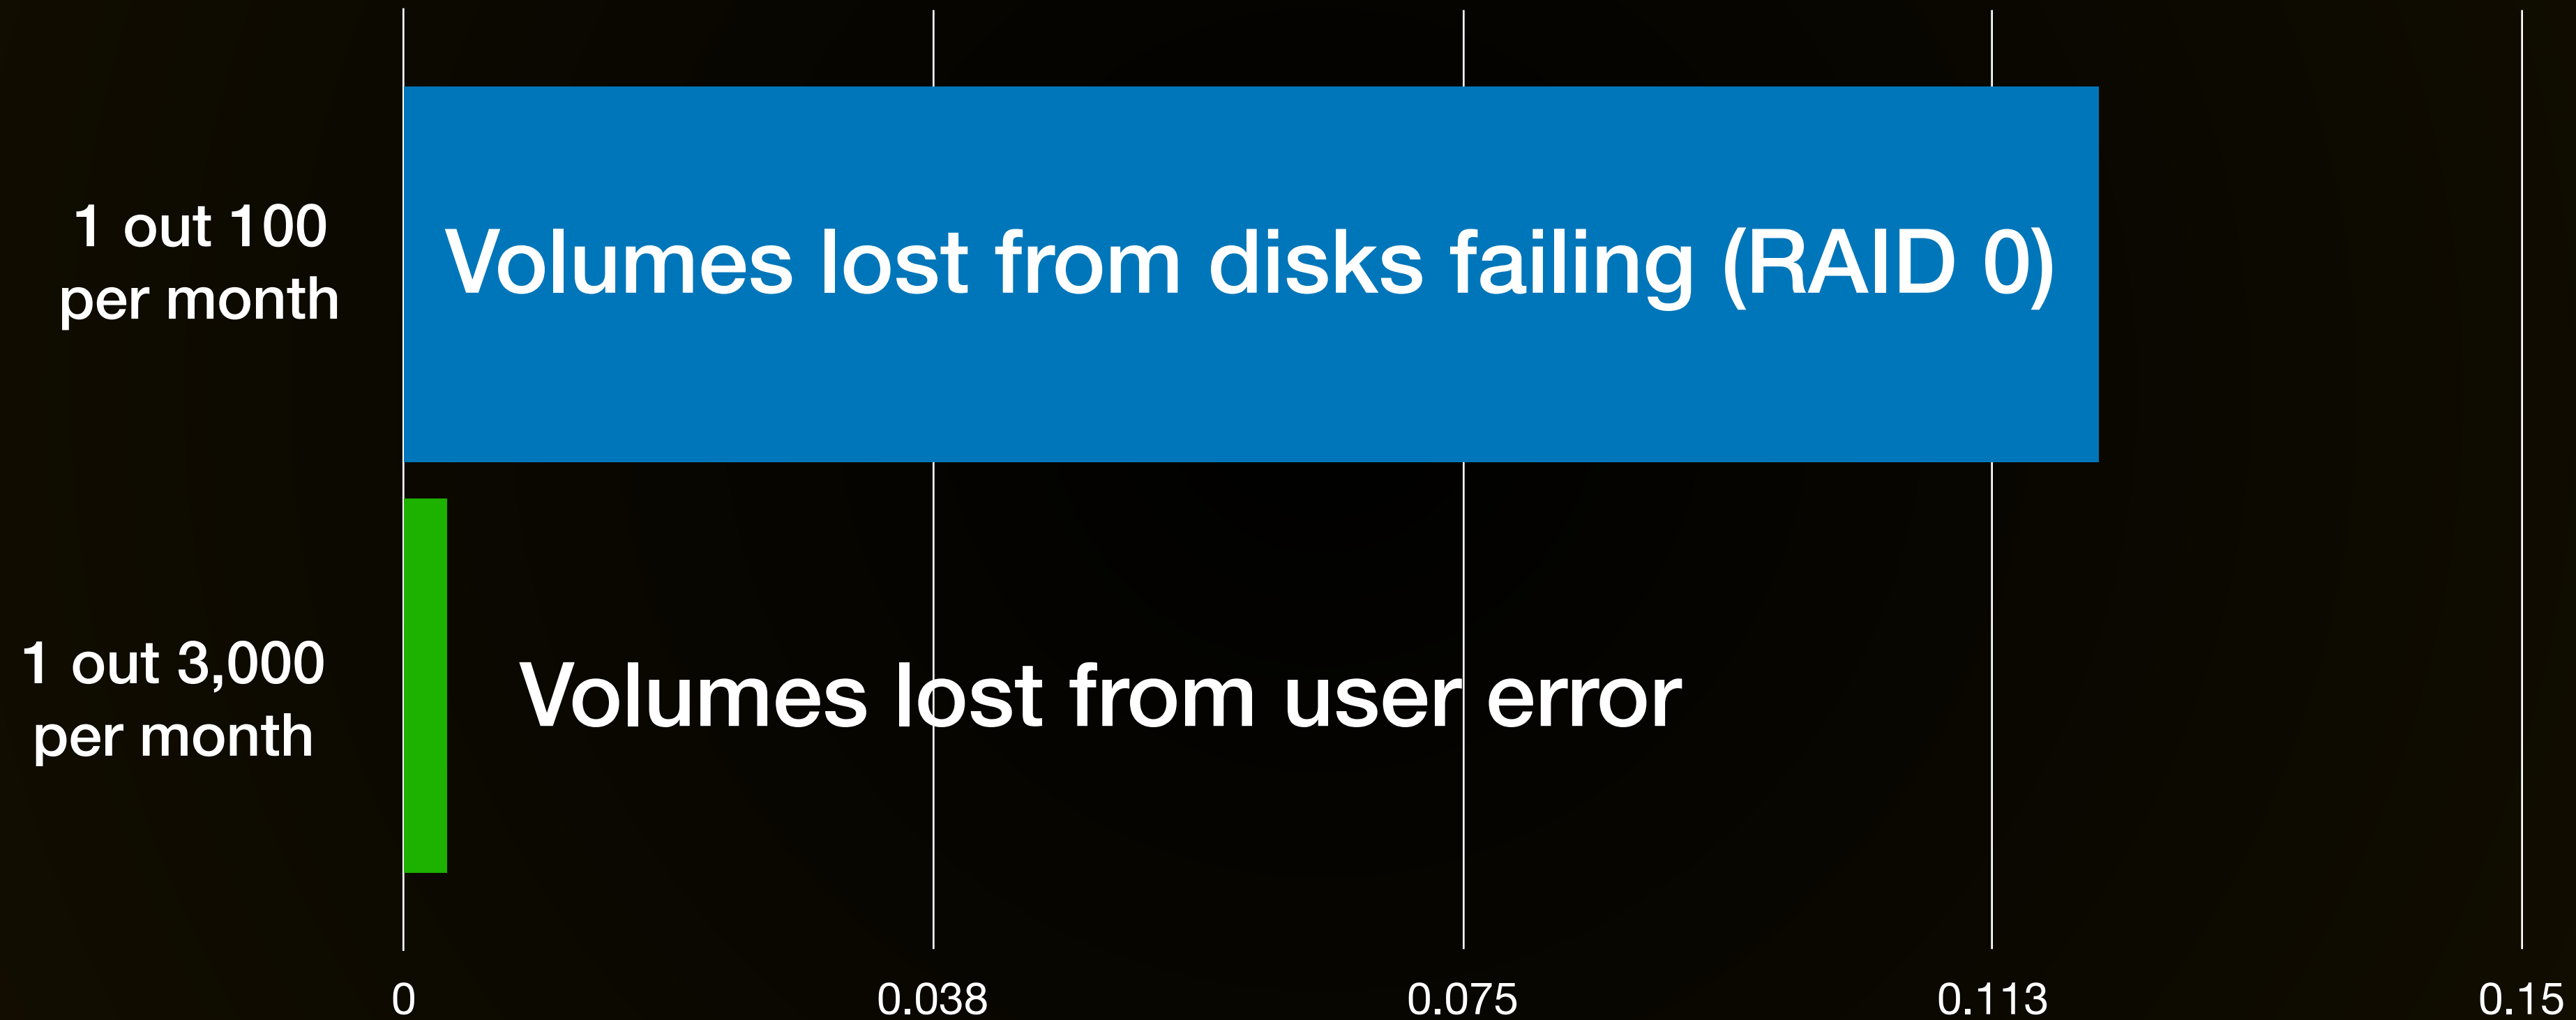






Volume Failure Rates

During the year 2008



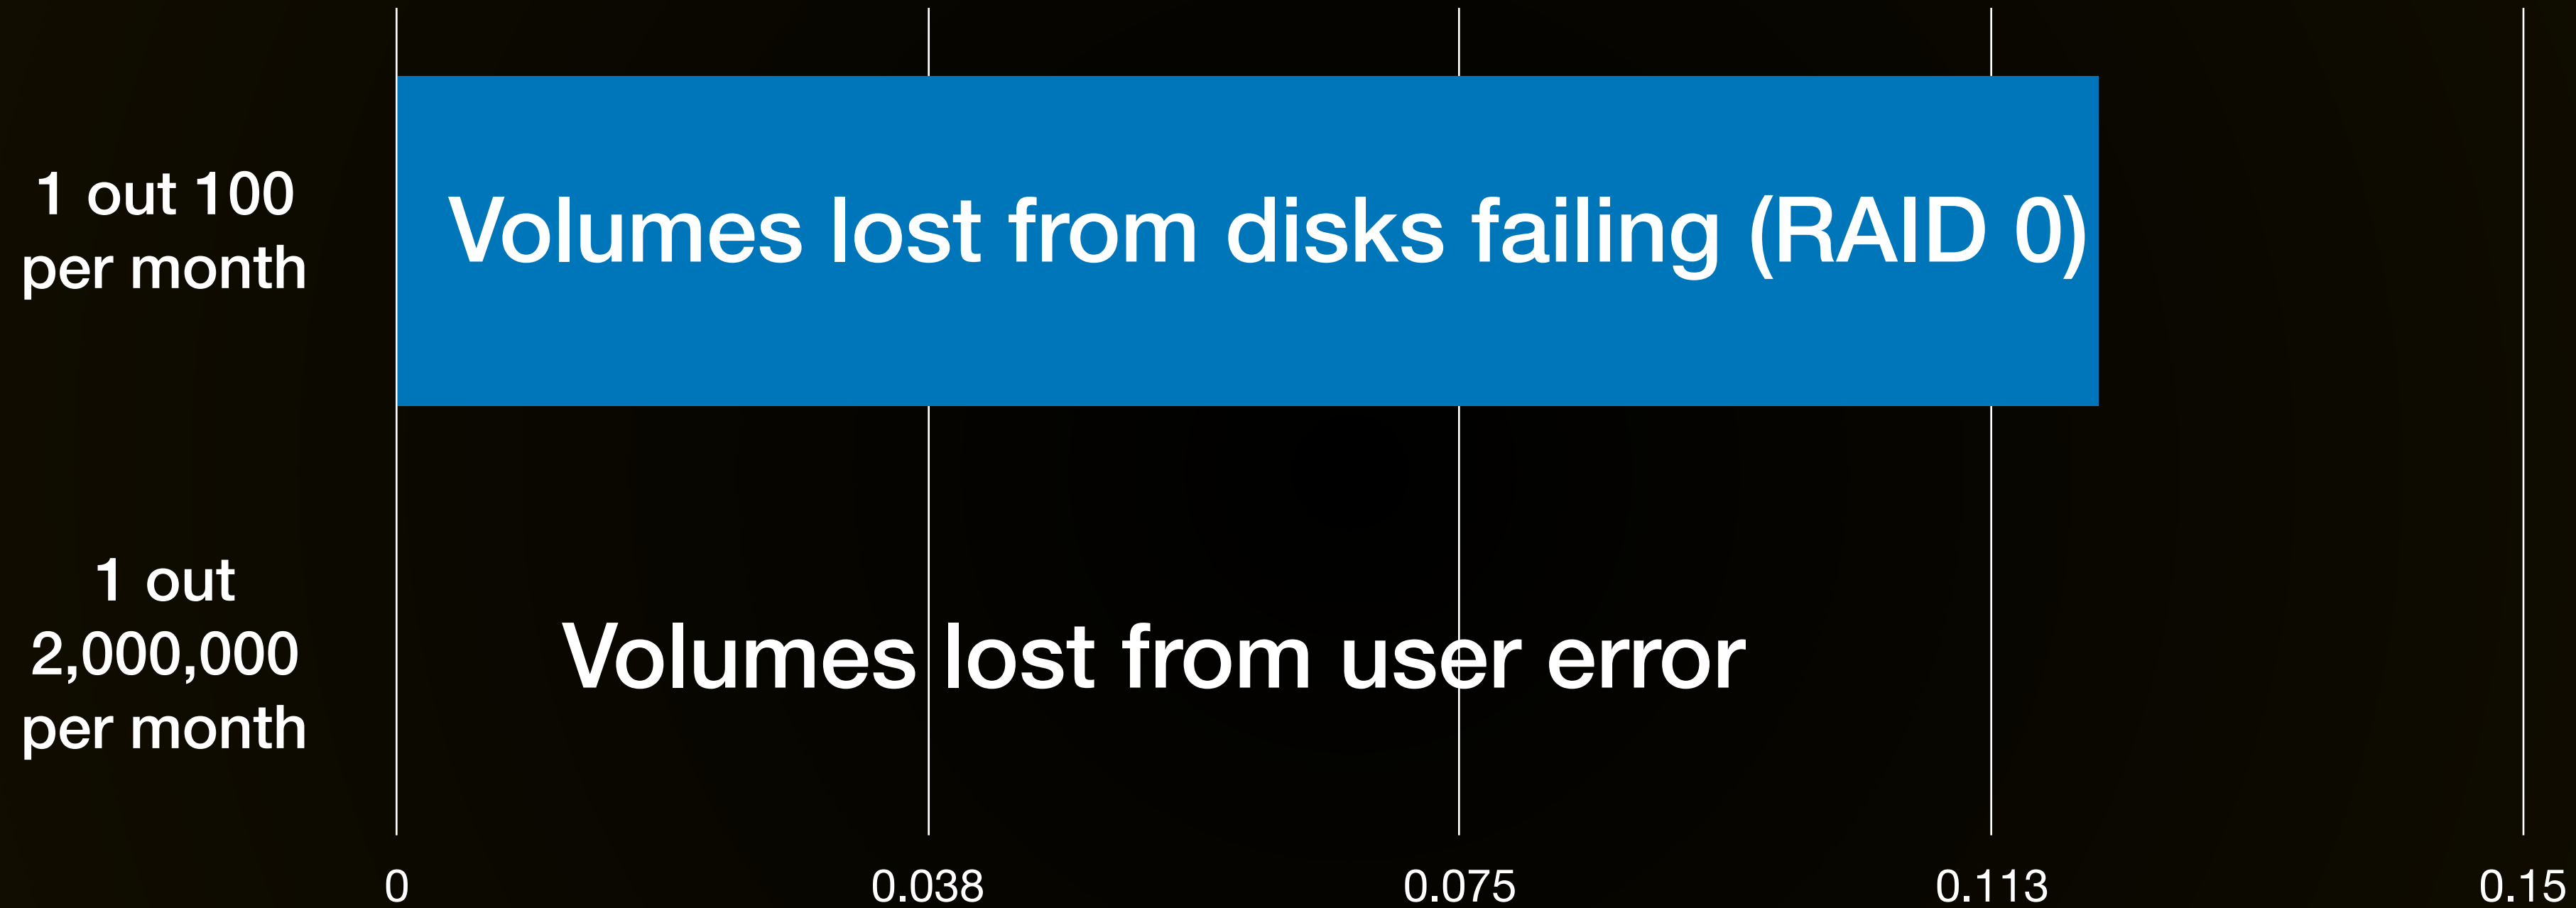
Annual Failure Rates

(all numbers approximate)



Volume Failure Rates

During the year 2010



Annual Failure Rates

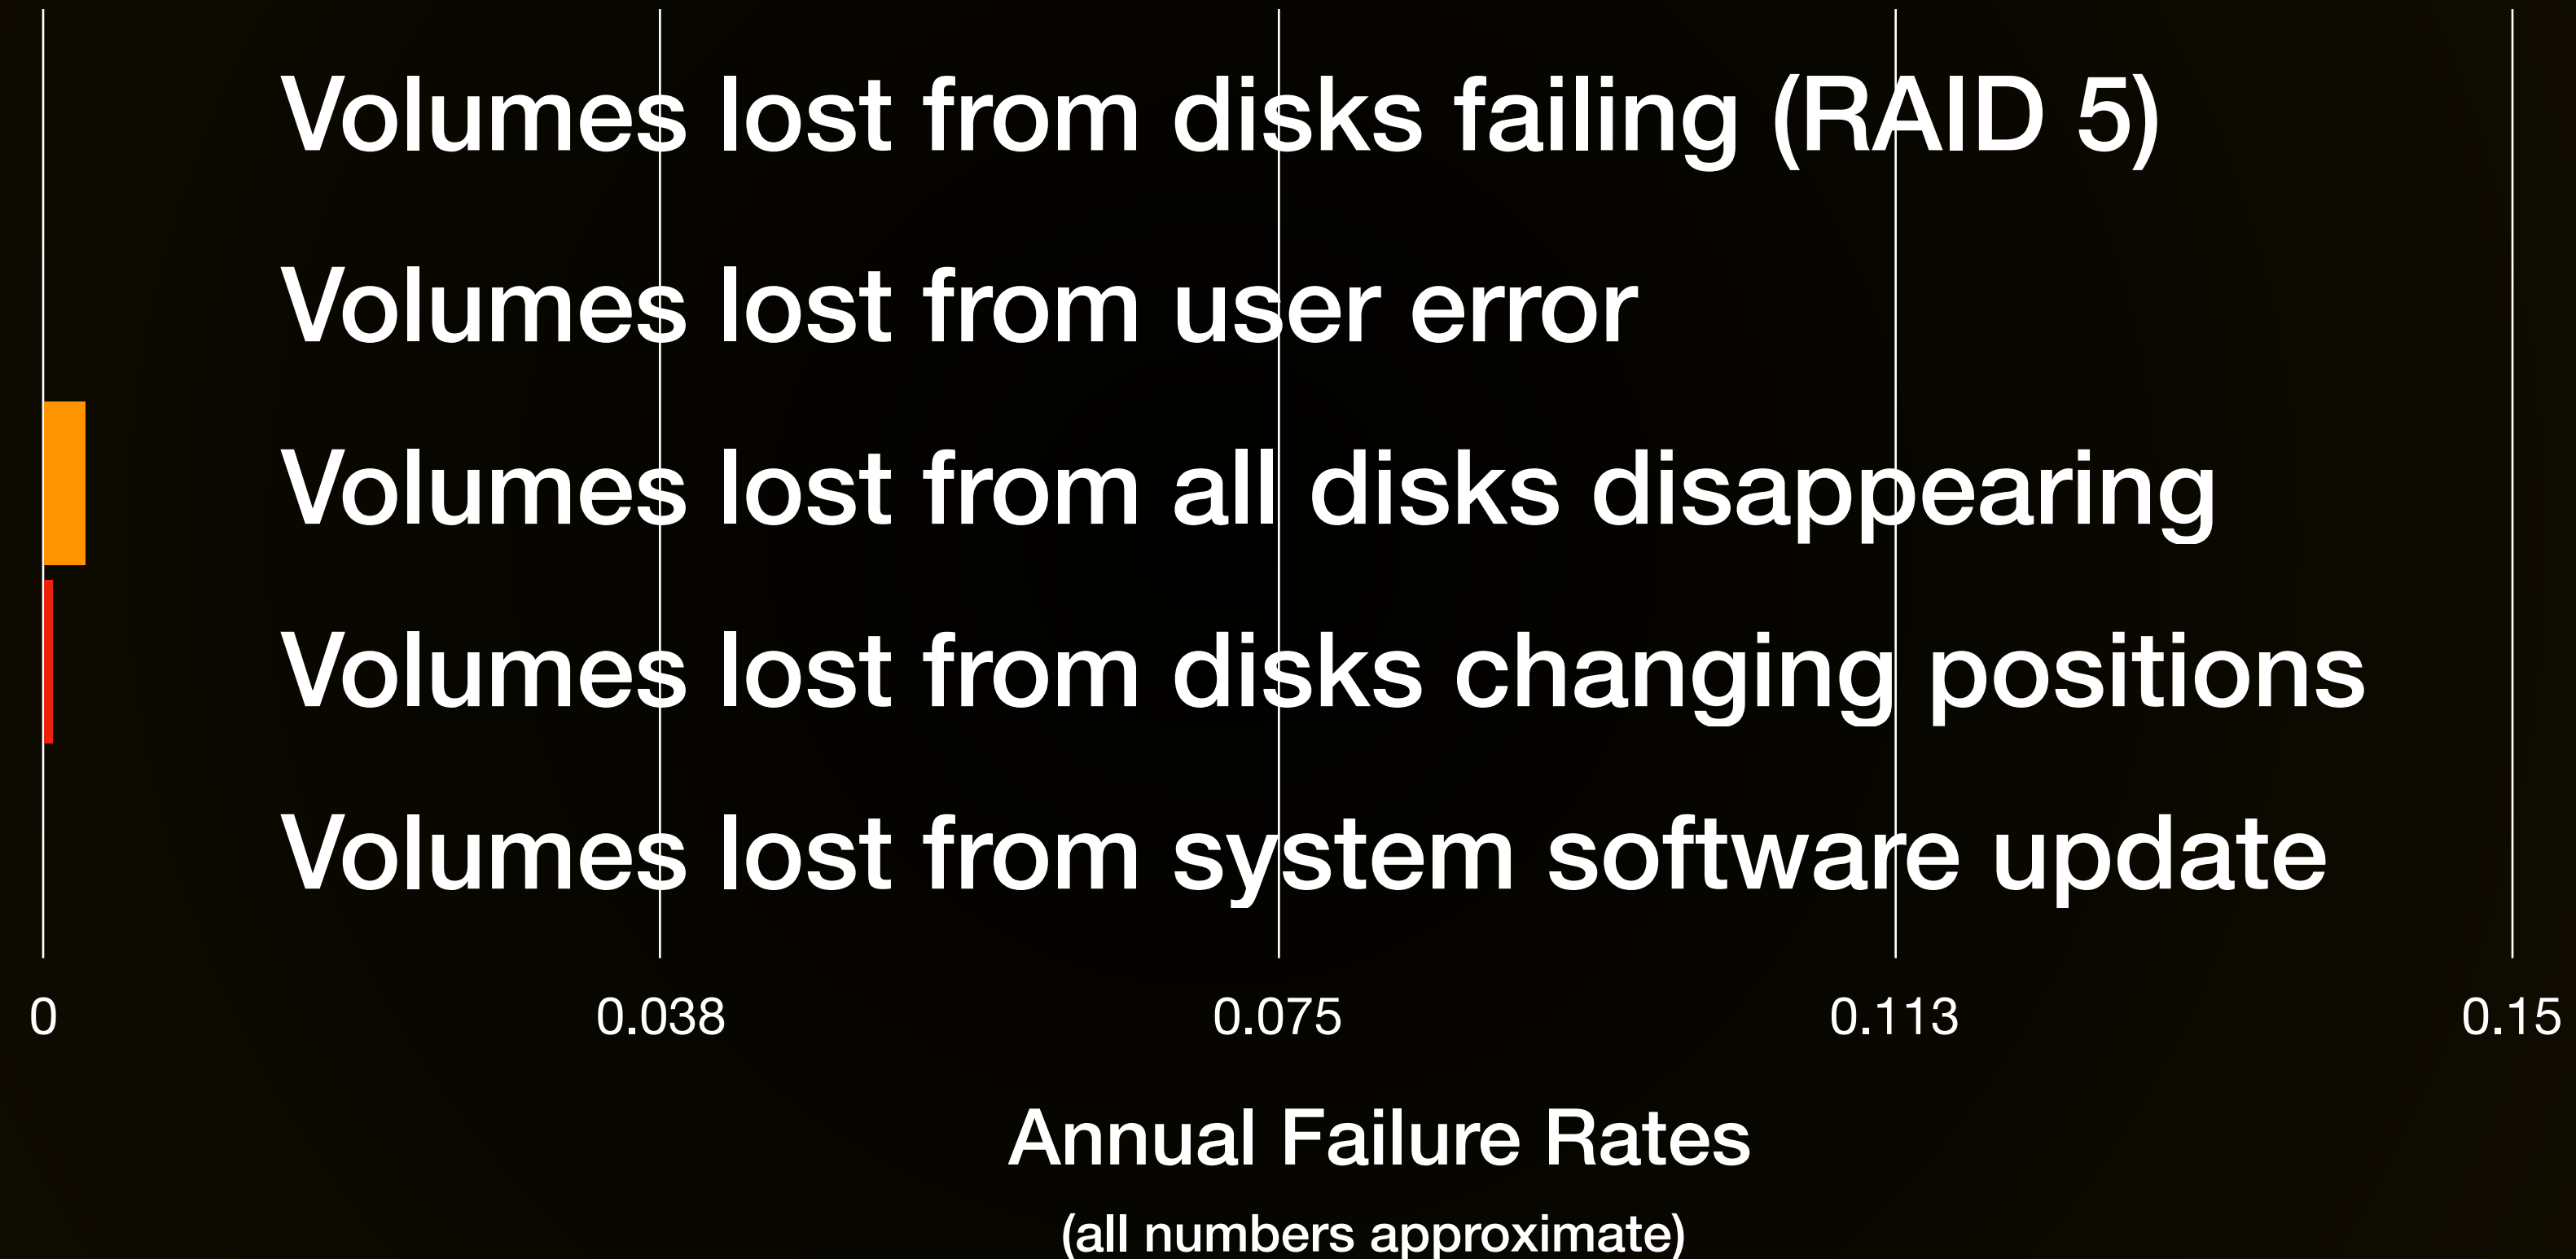
(all numbers approximate)

SoftRAID v.4 added Volume Safeguards = protect from user error



Volume Failure Rates

During the year 2014

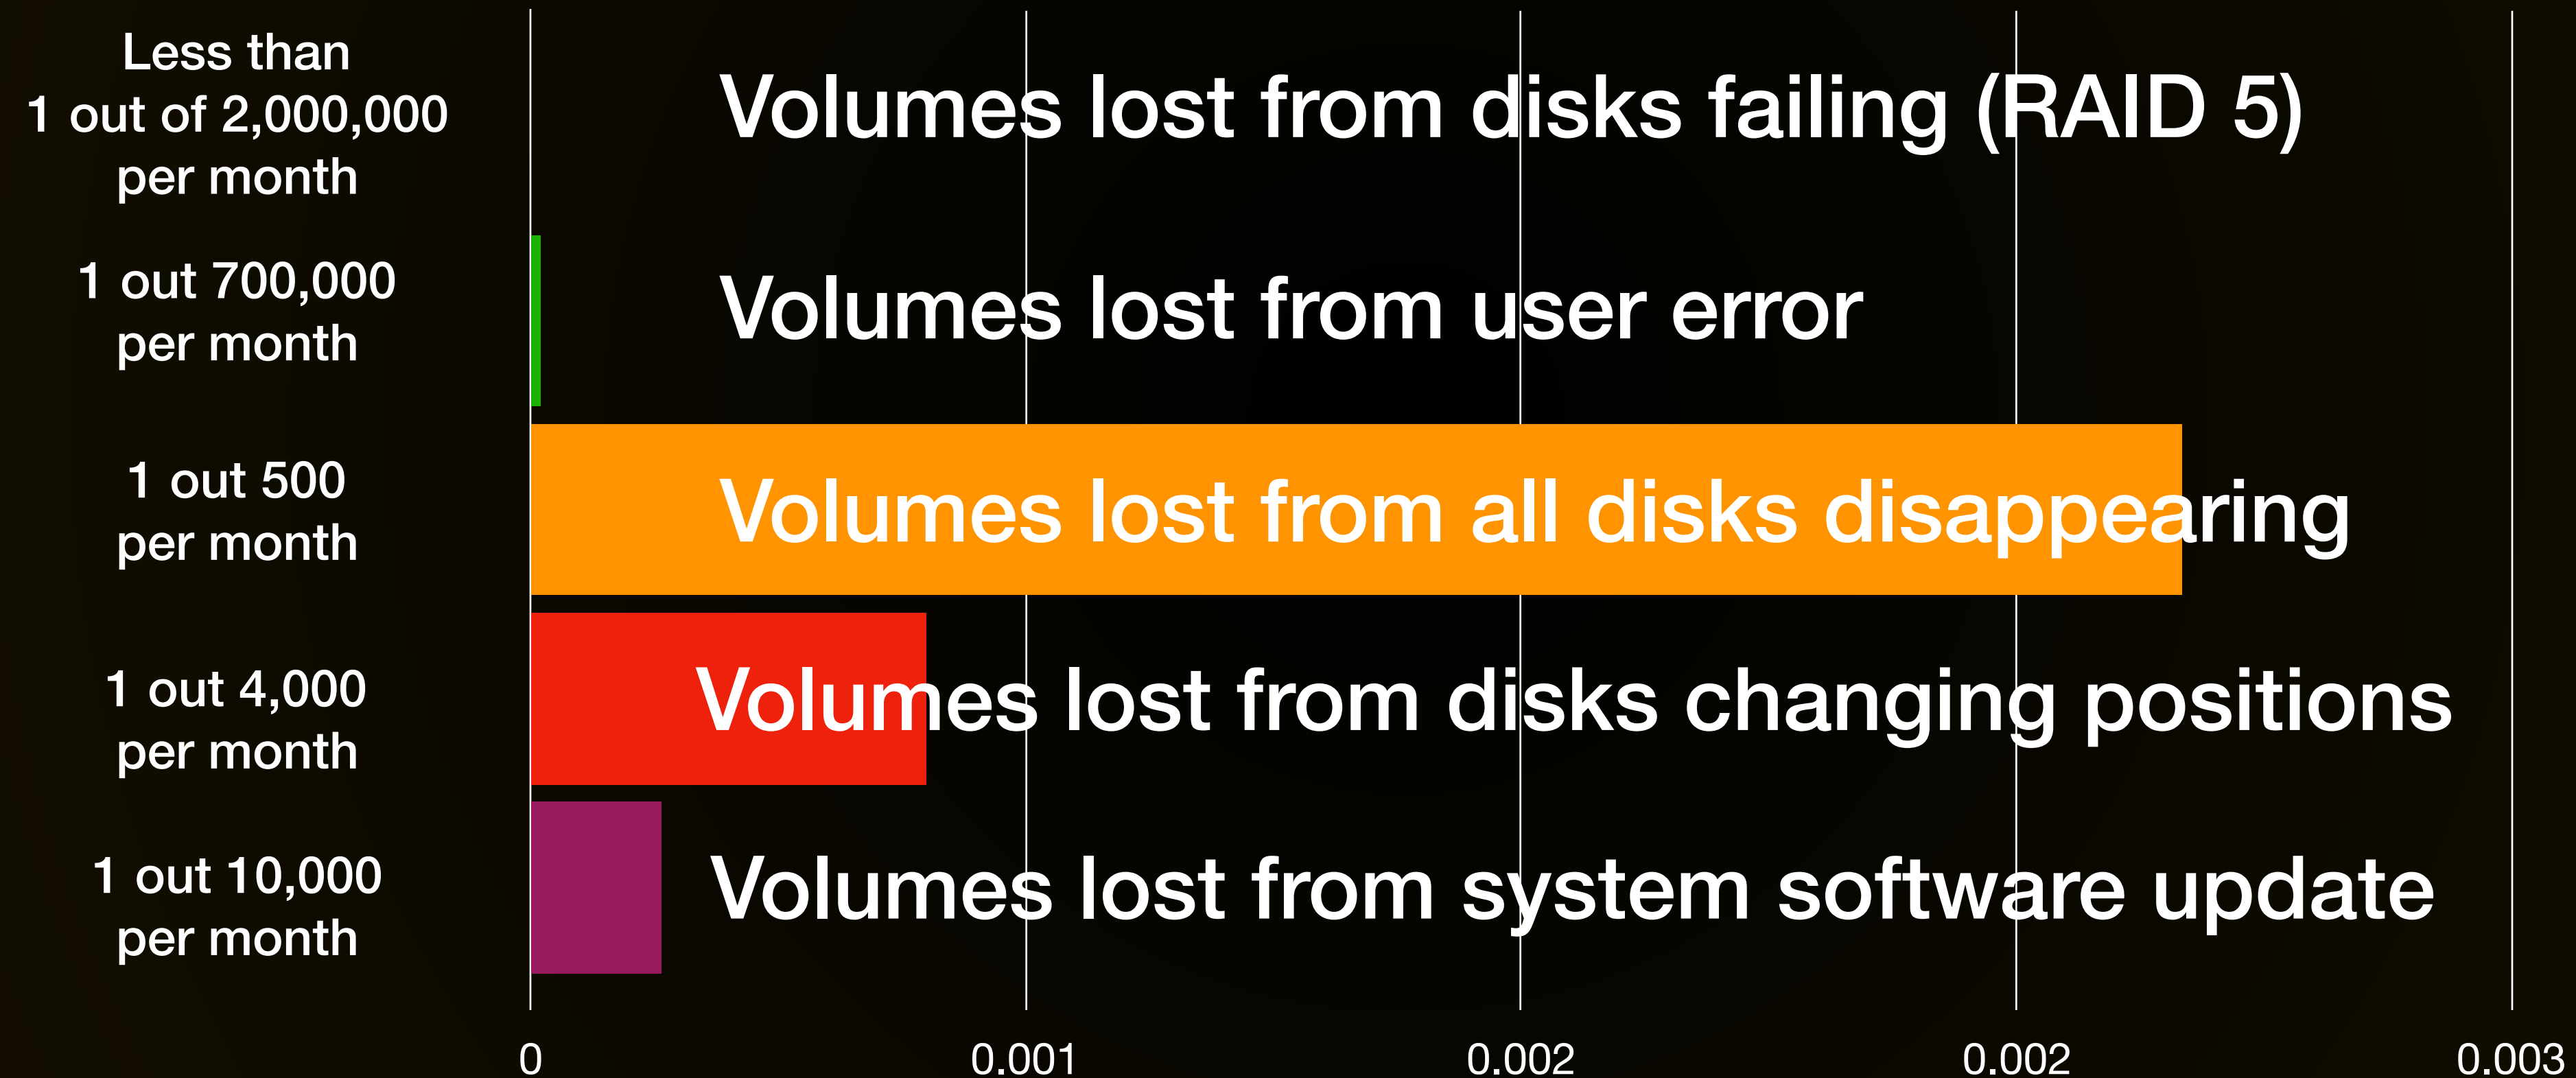


SoftRAID v. 5 added RAID 4 and 5 volumes



Volume Failure Rates

During the year 2014



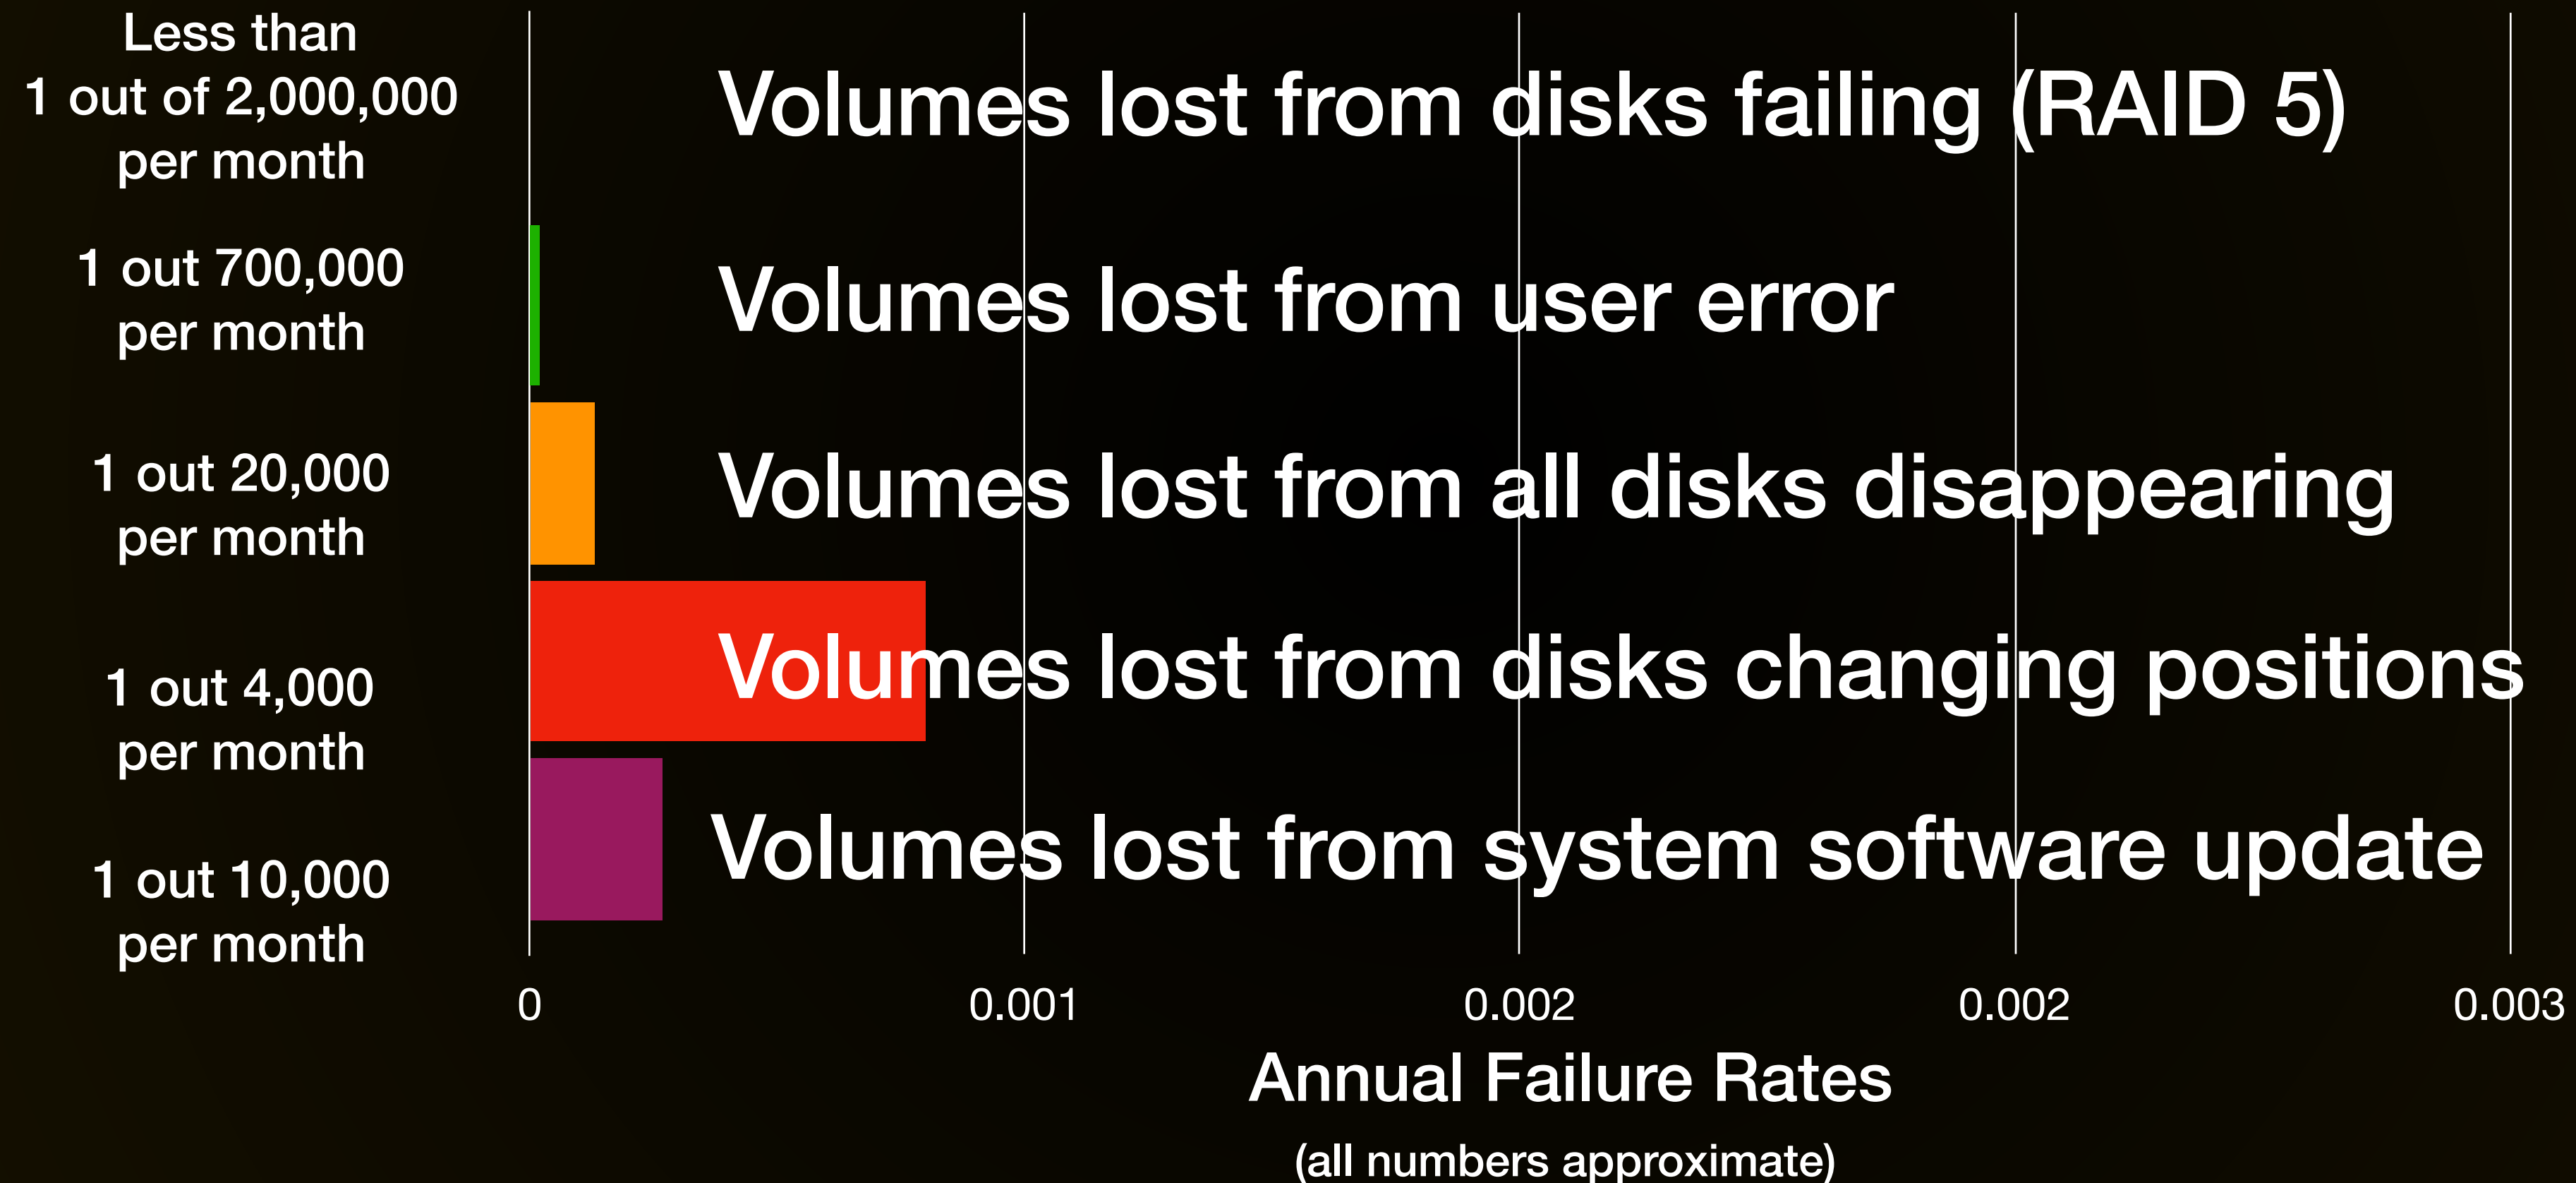
Annual Failure Rates

(all numbers approximate)



Volume Failure Rates

During the year 2017

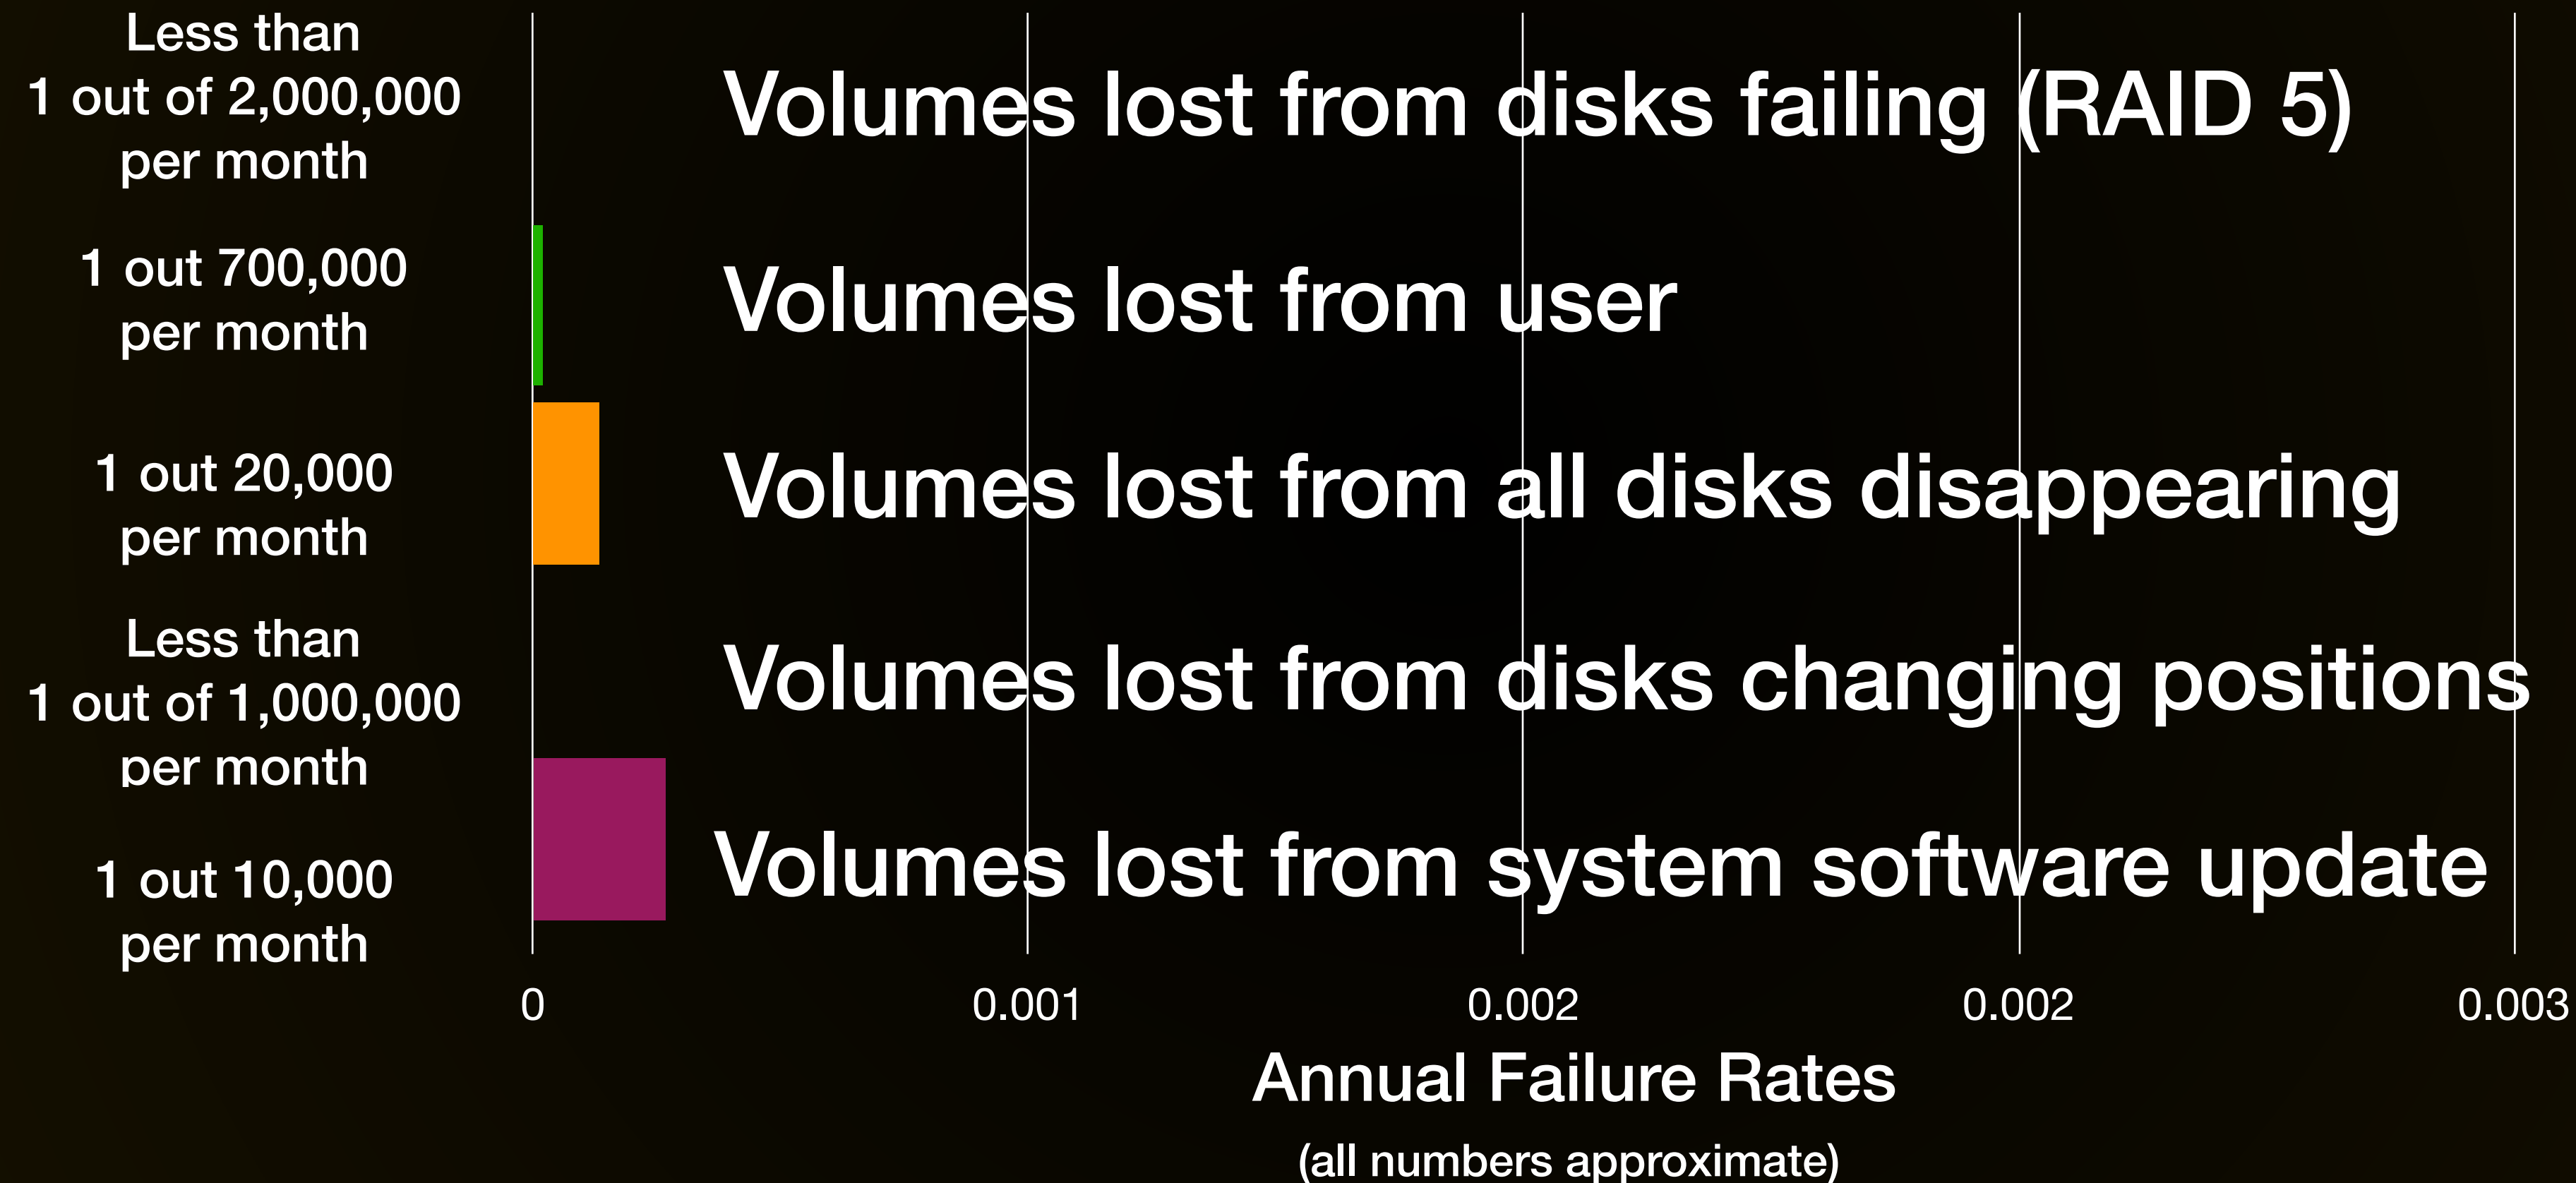


Mac OS 10.12.1 fixes kernel bug which causes all disks to



Volume Failure Rates

During the year 2019



SoftRAID 5.7 detects disks changing position due to kernel bug

Volume Failure Rates



Things we have not observed:

- **Loss of volumes during RAID 5 rebuild**
- **Bit rot on SSDs or HDDs**














ADRESS

**GIBRALTARGATAN 20
411 32 GÖTEBORG**



TELEFON

070-825 95 61











3. Apple's T2 Chip

How does the T2 chip protect your volumes?

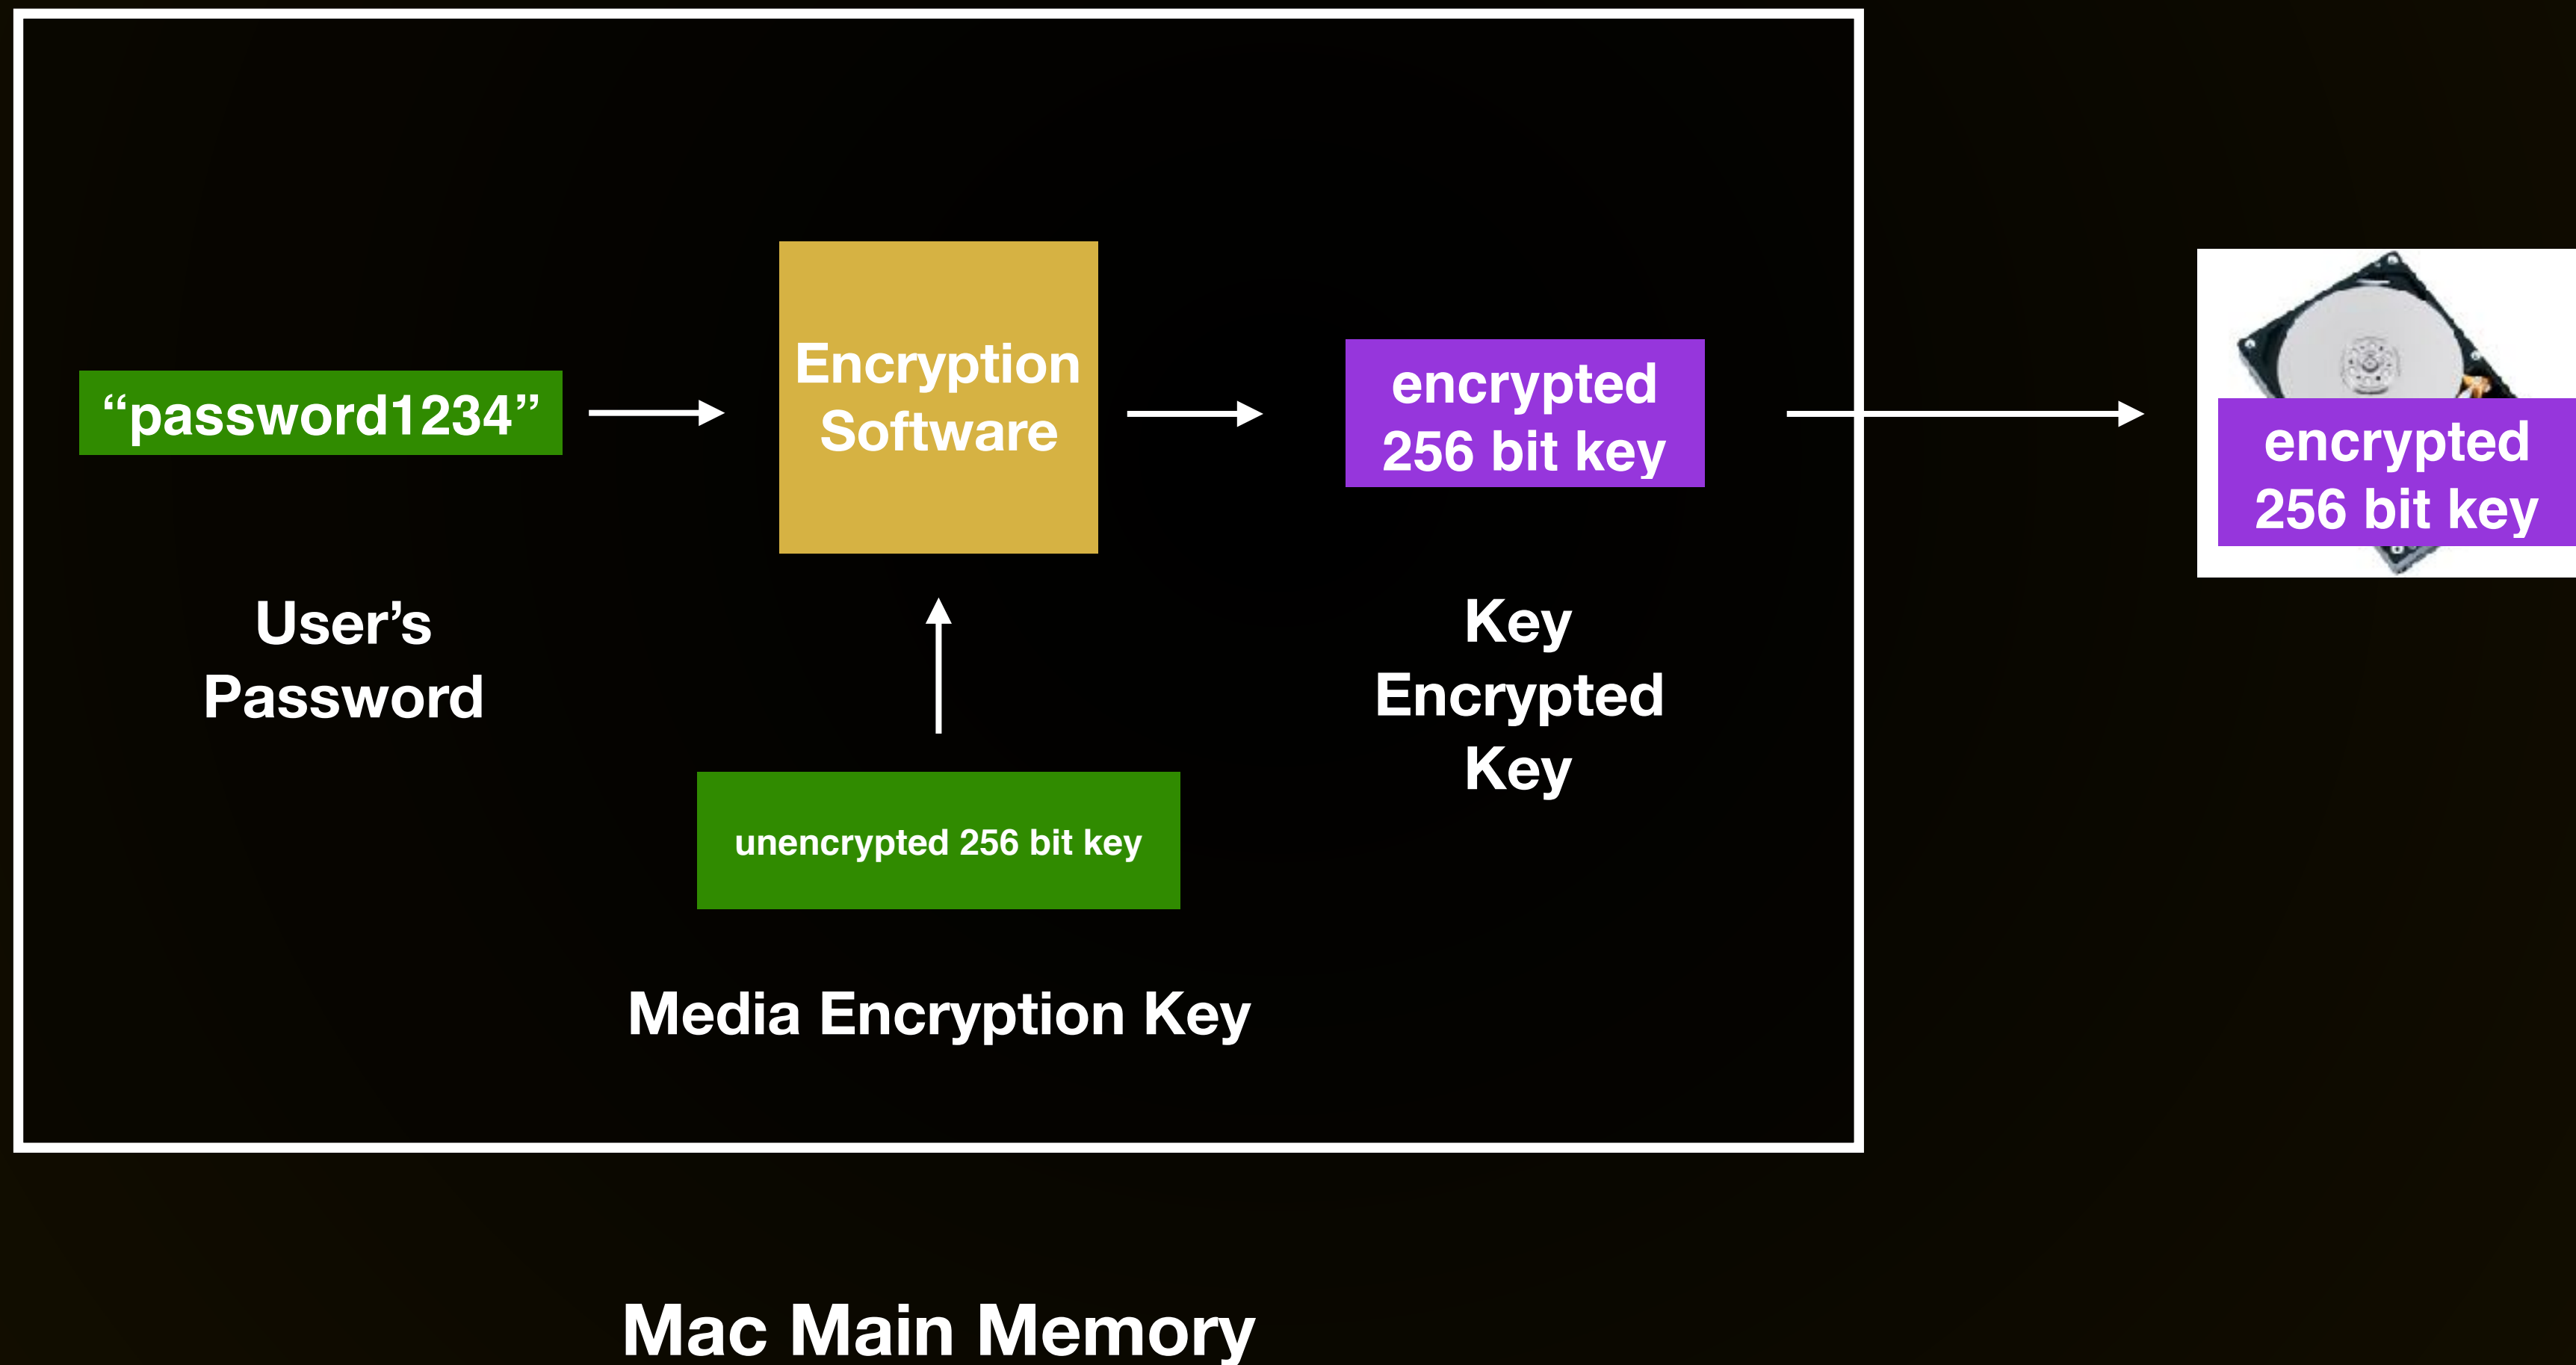


Review How FileVault Volumes Work



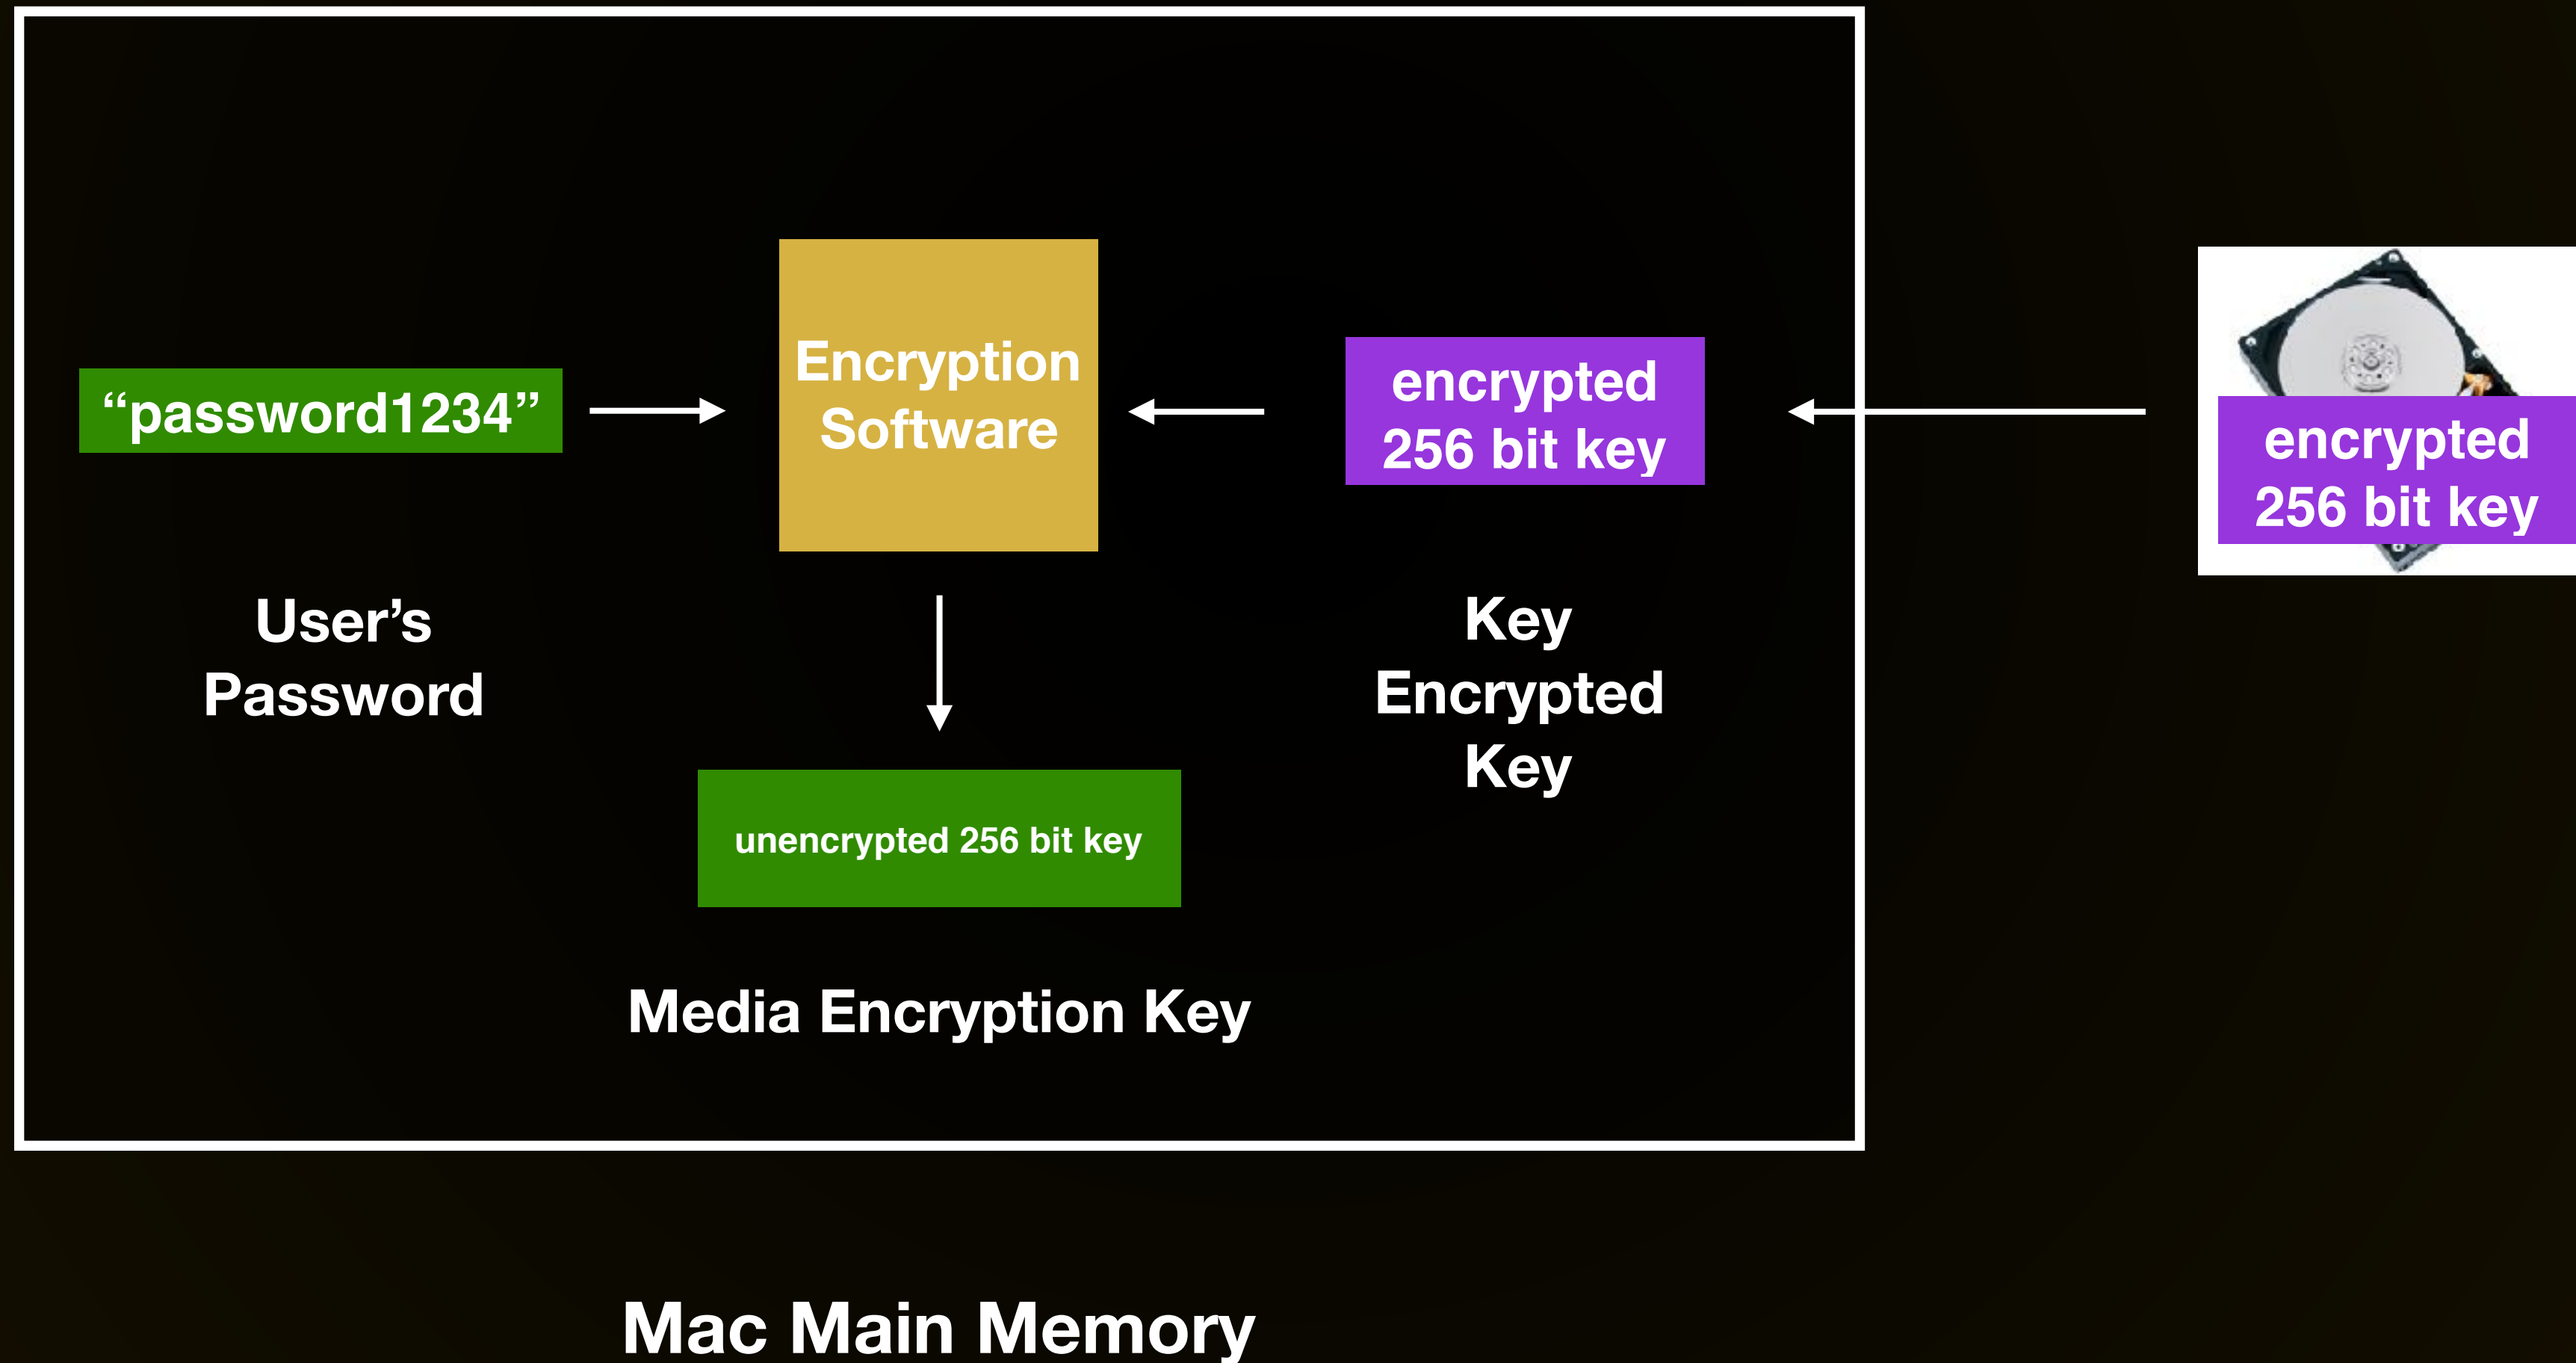
FileVault Volume

Creating a Volume



FileVault Volume

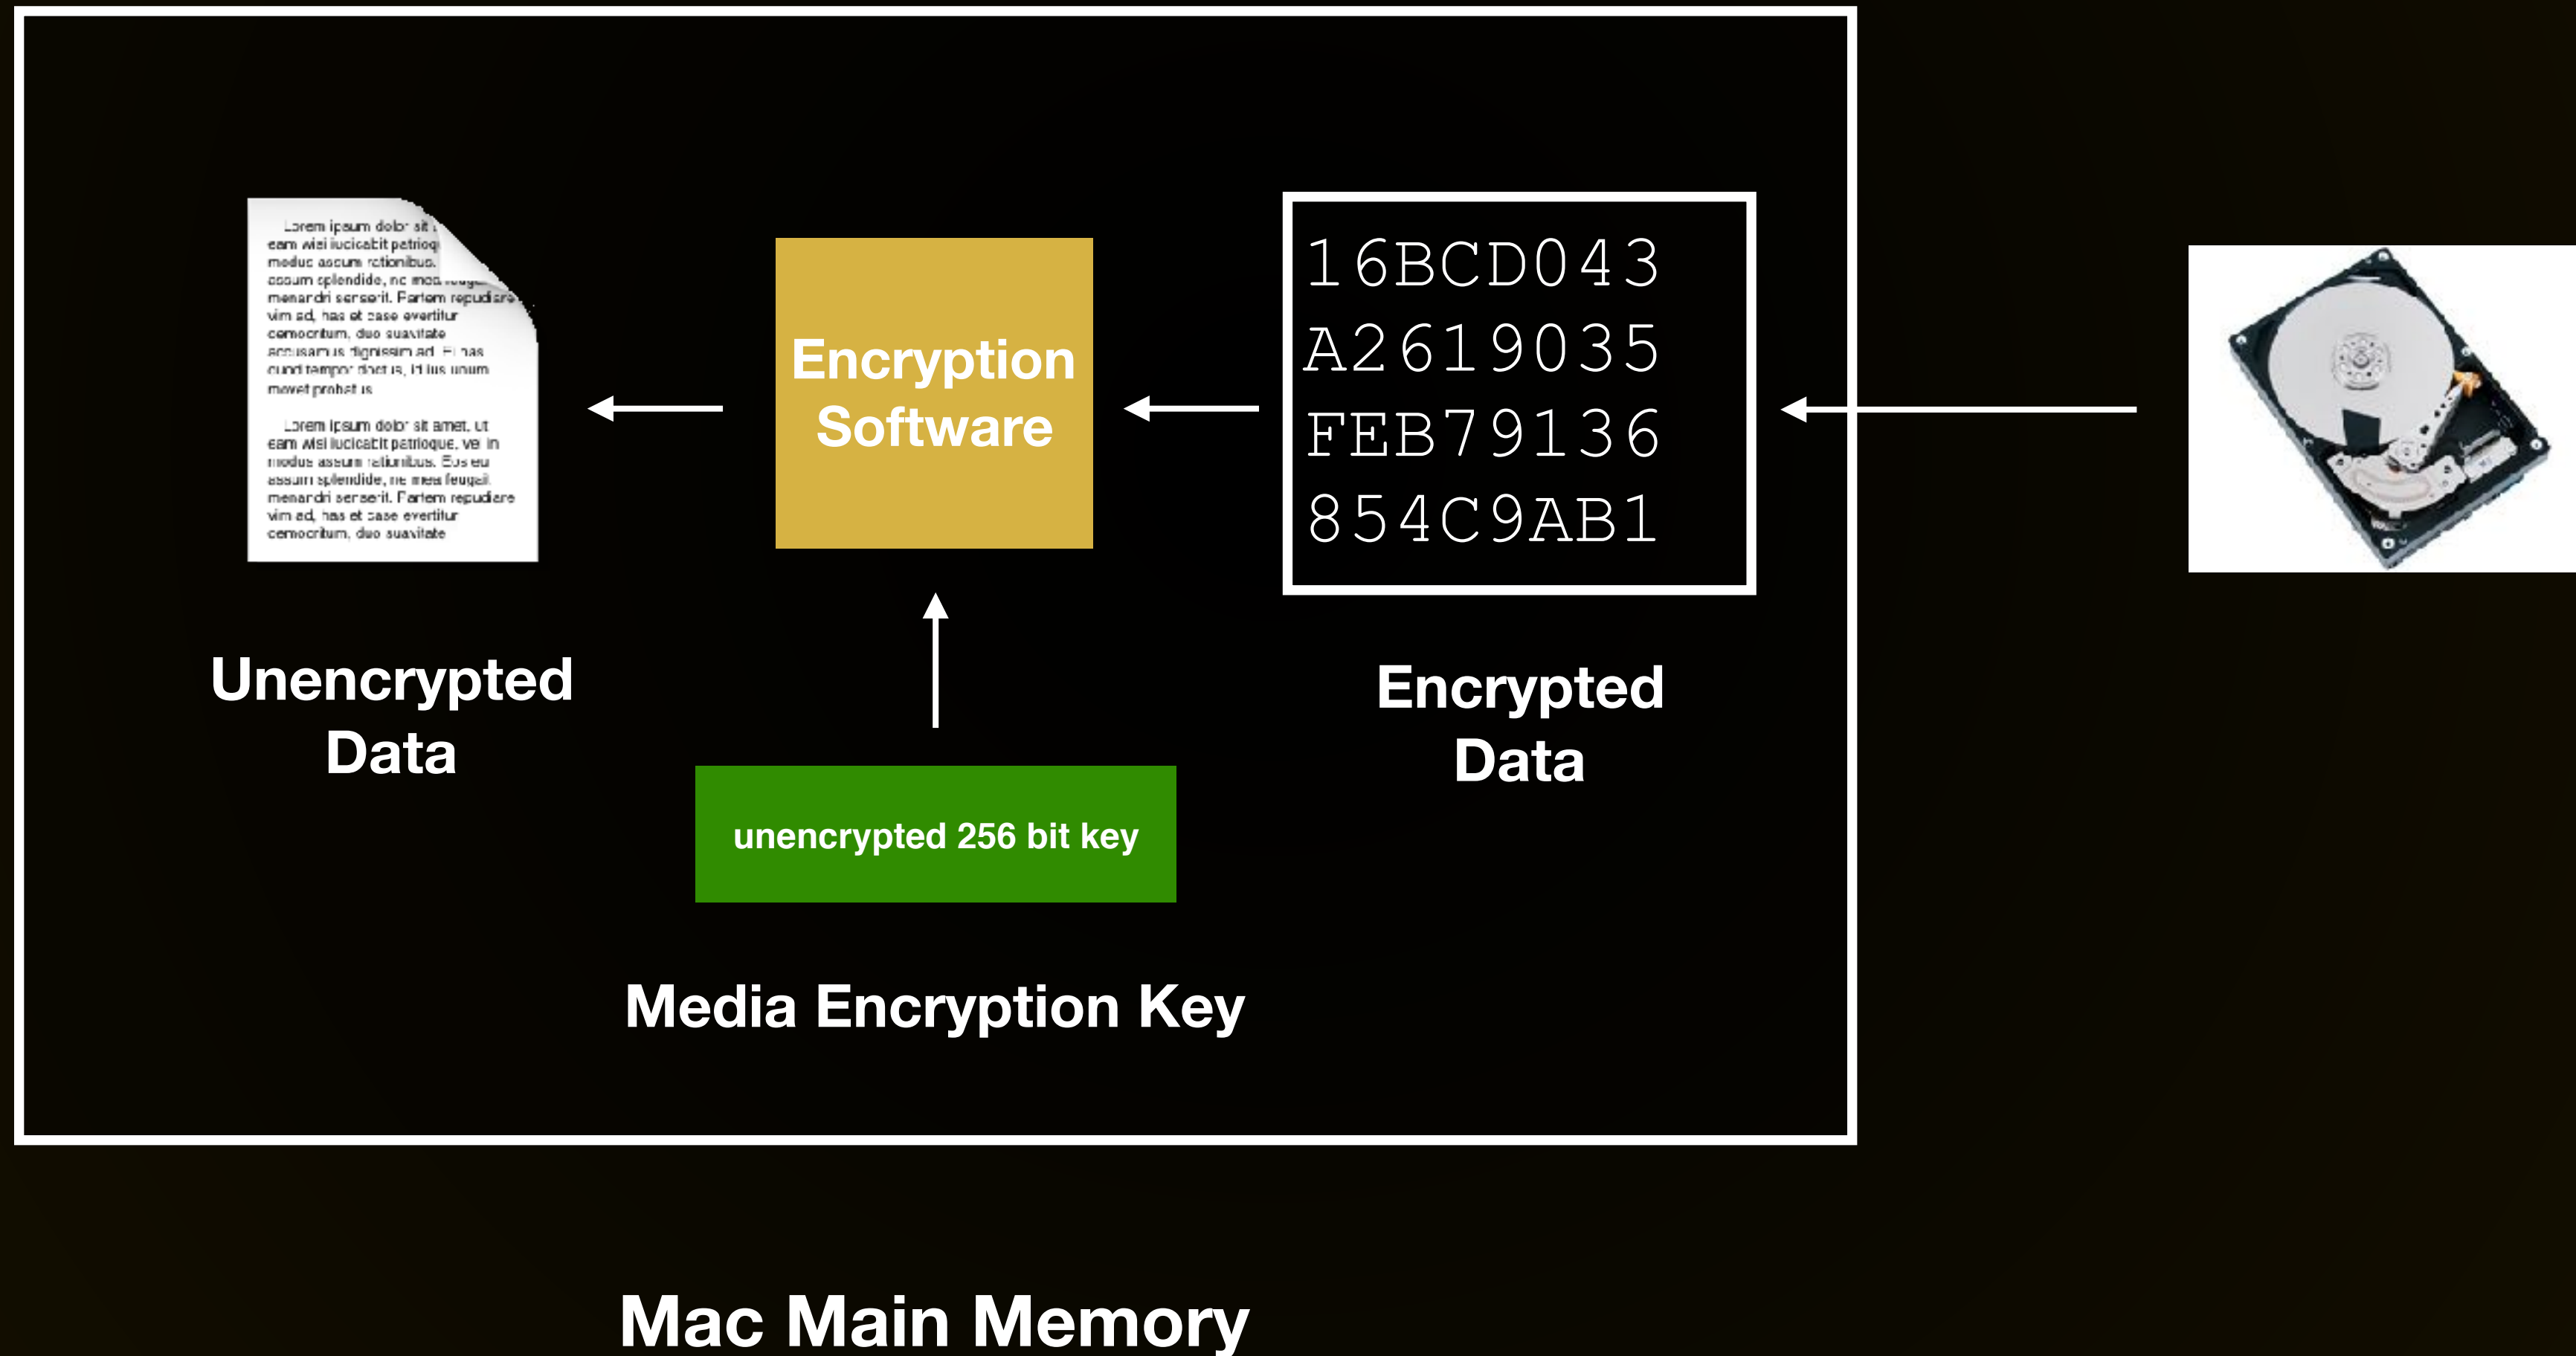
Unlocking a Volume





FileVault Volume

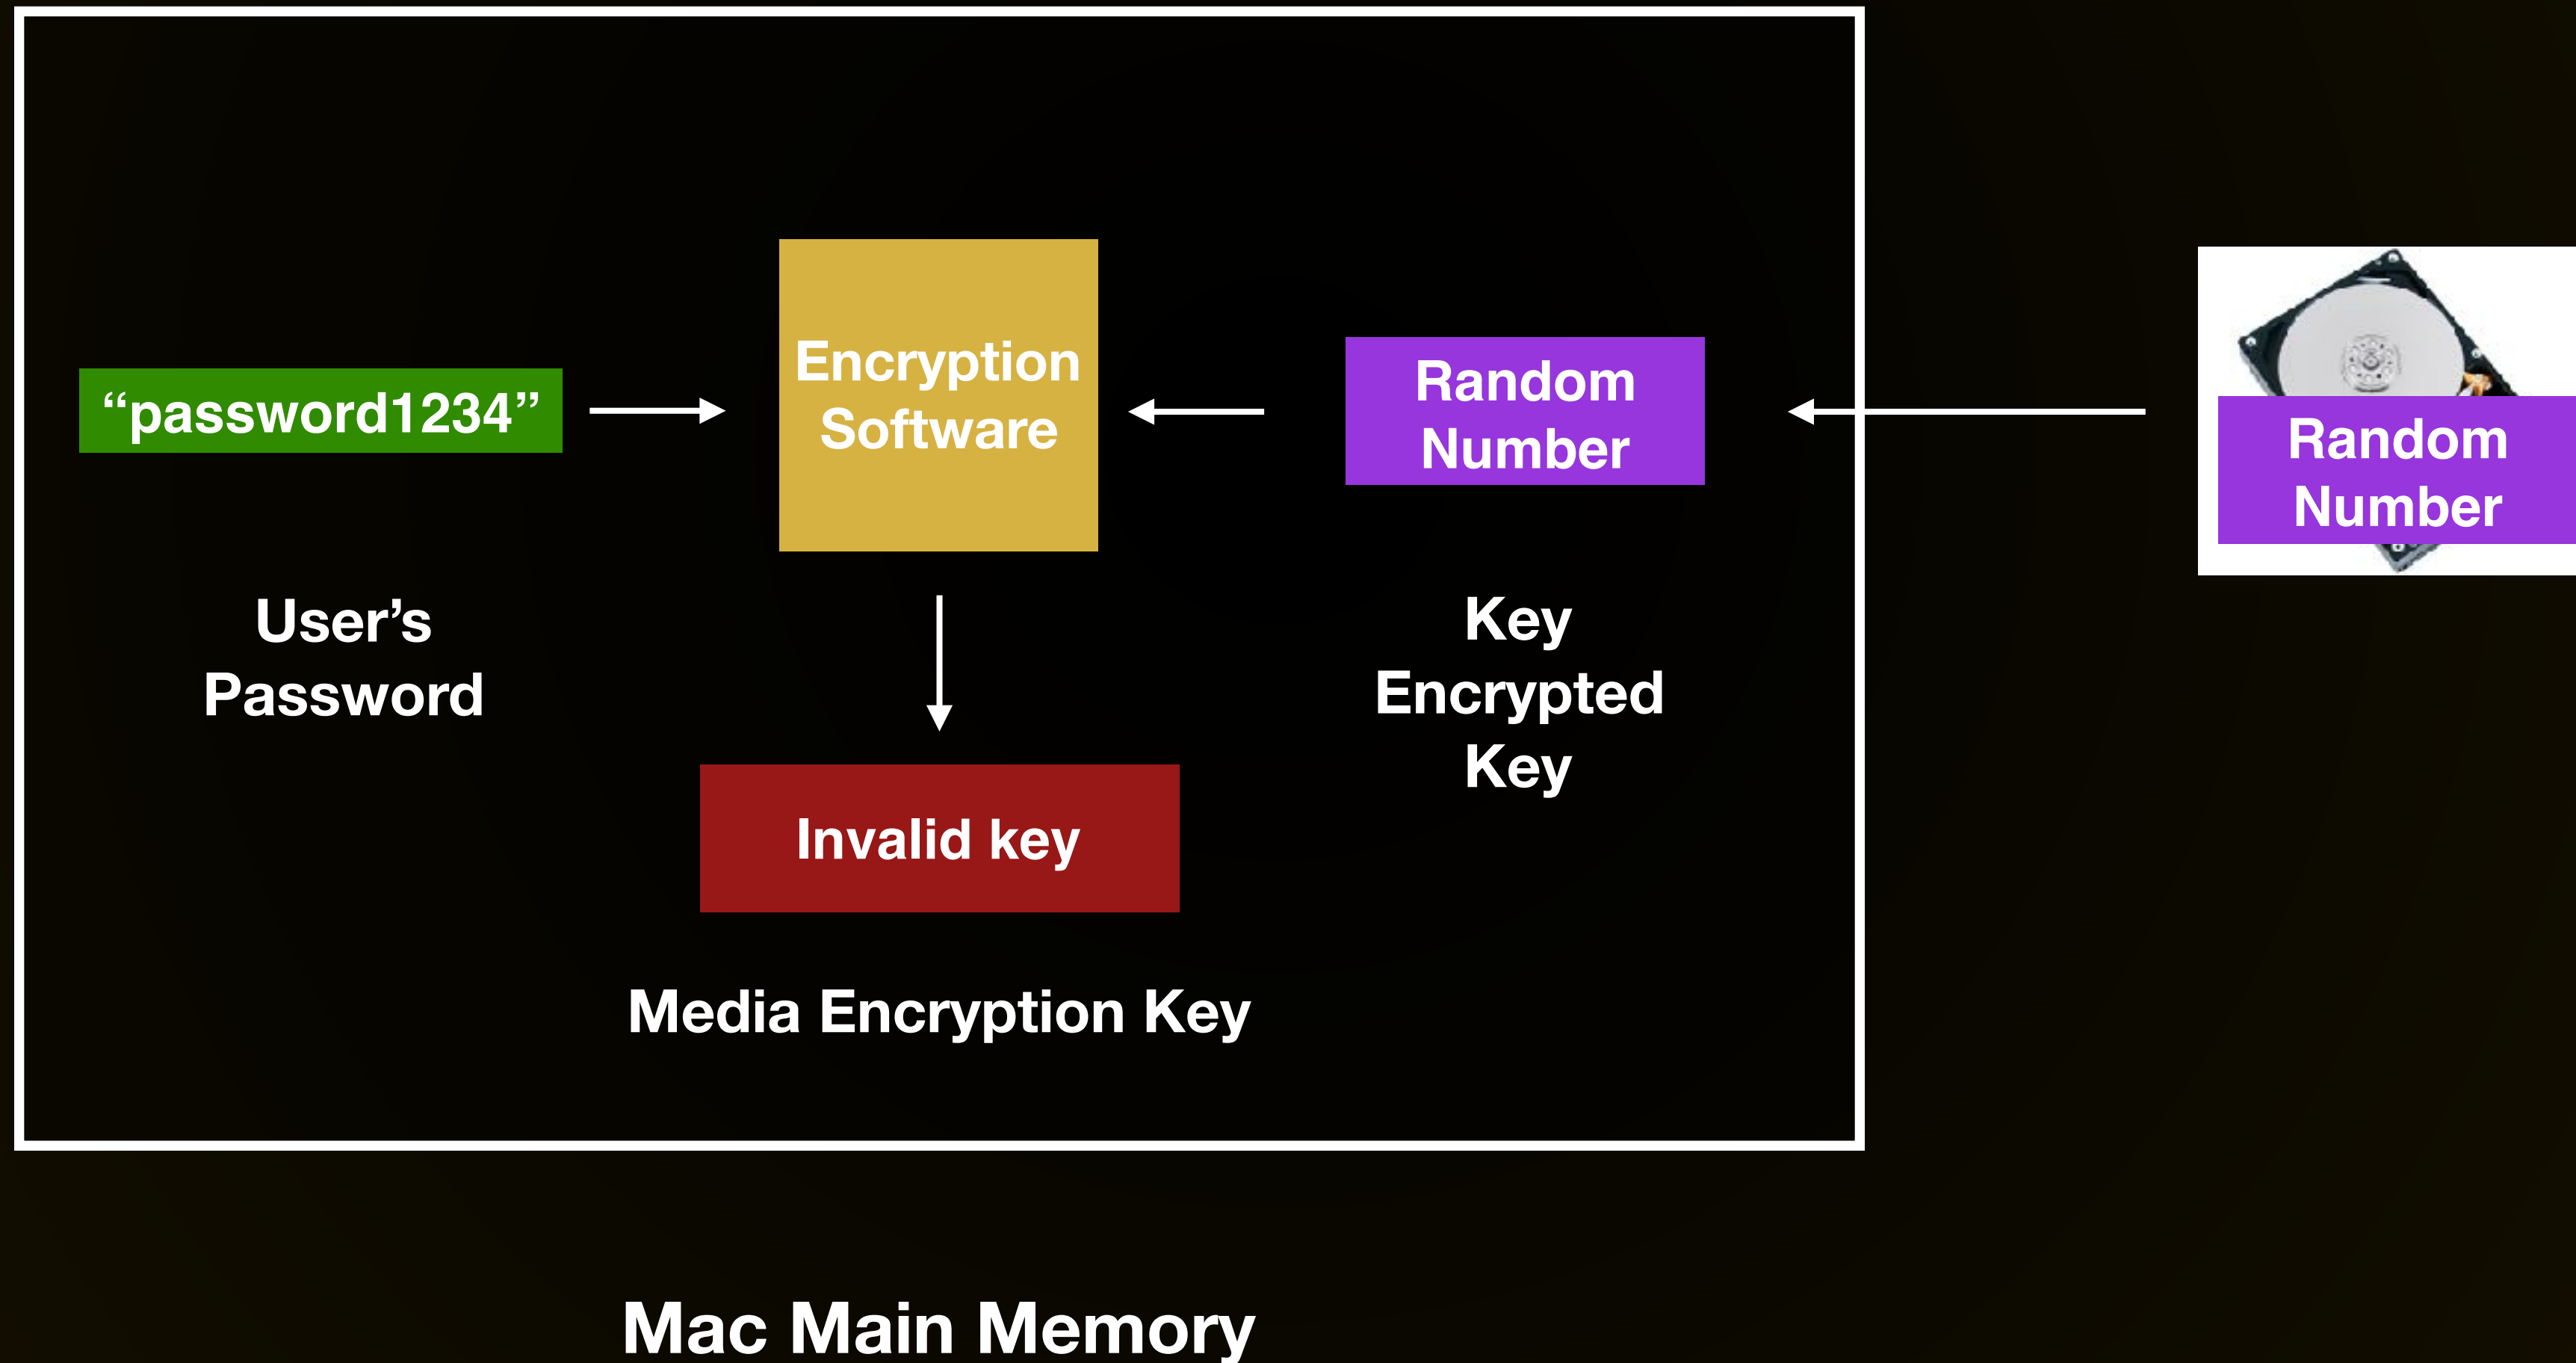
Reading Encrypted Files





FileVault Volume

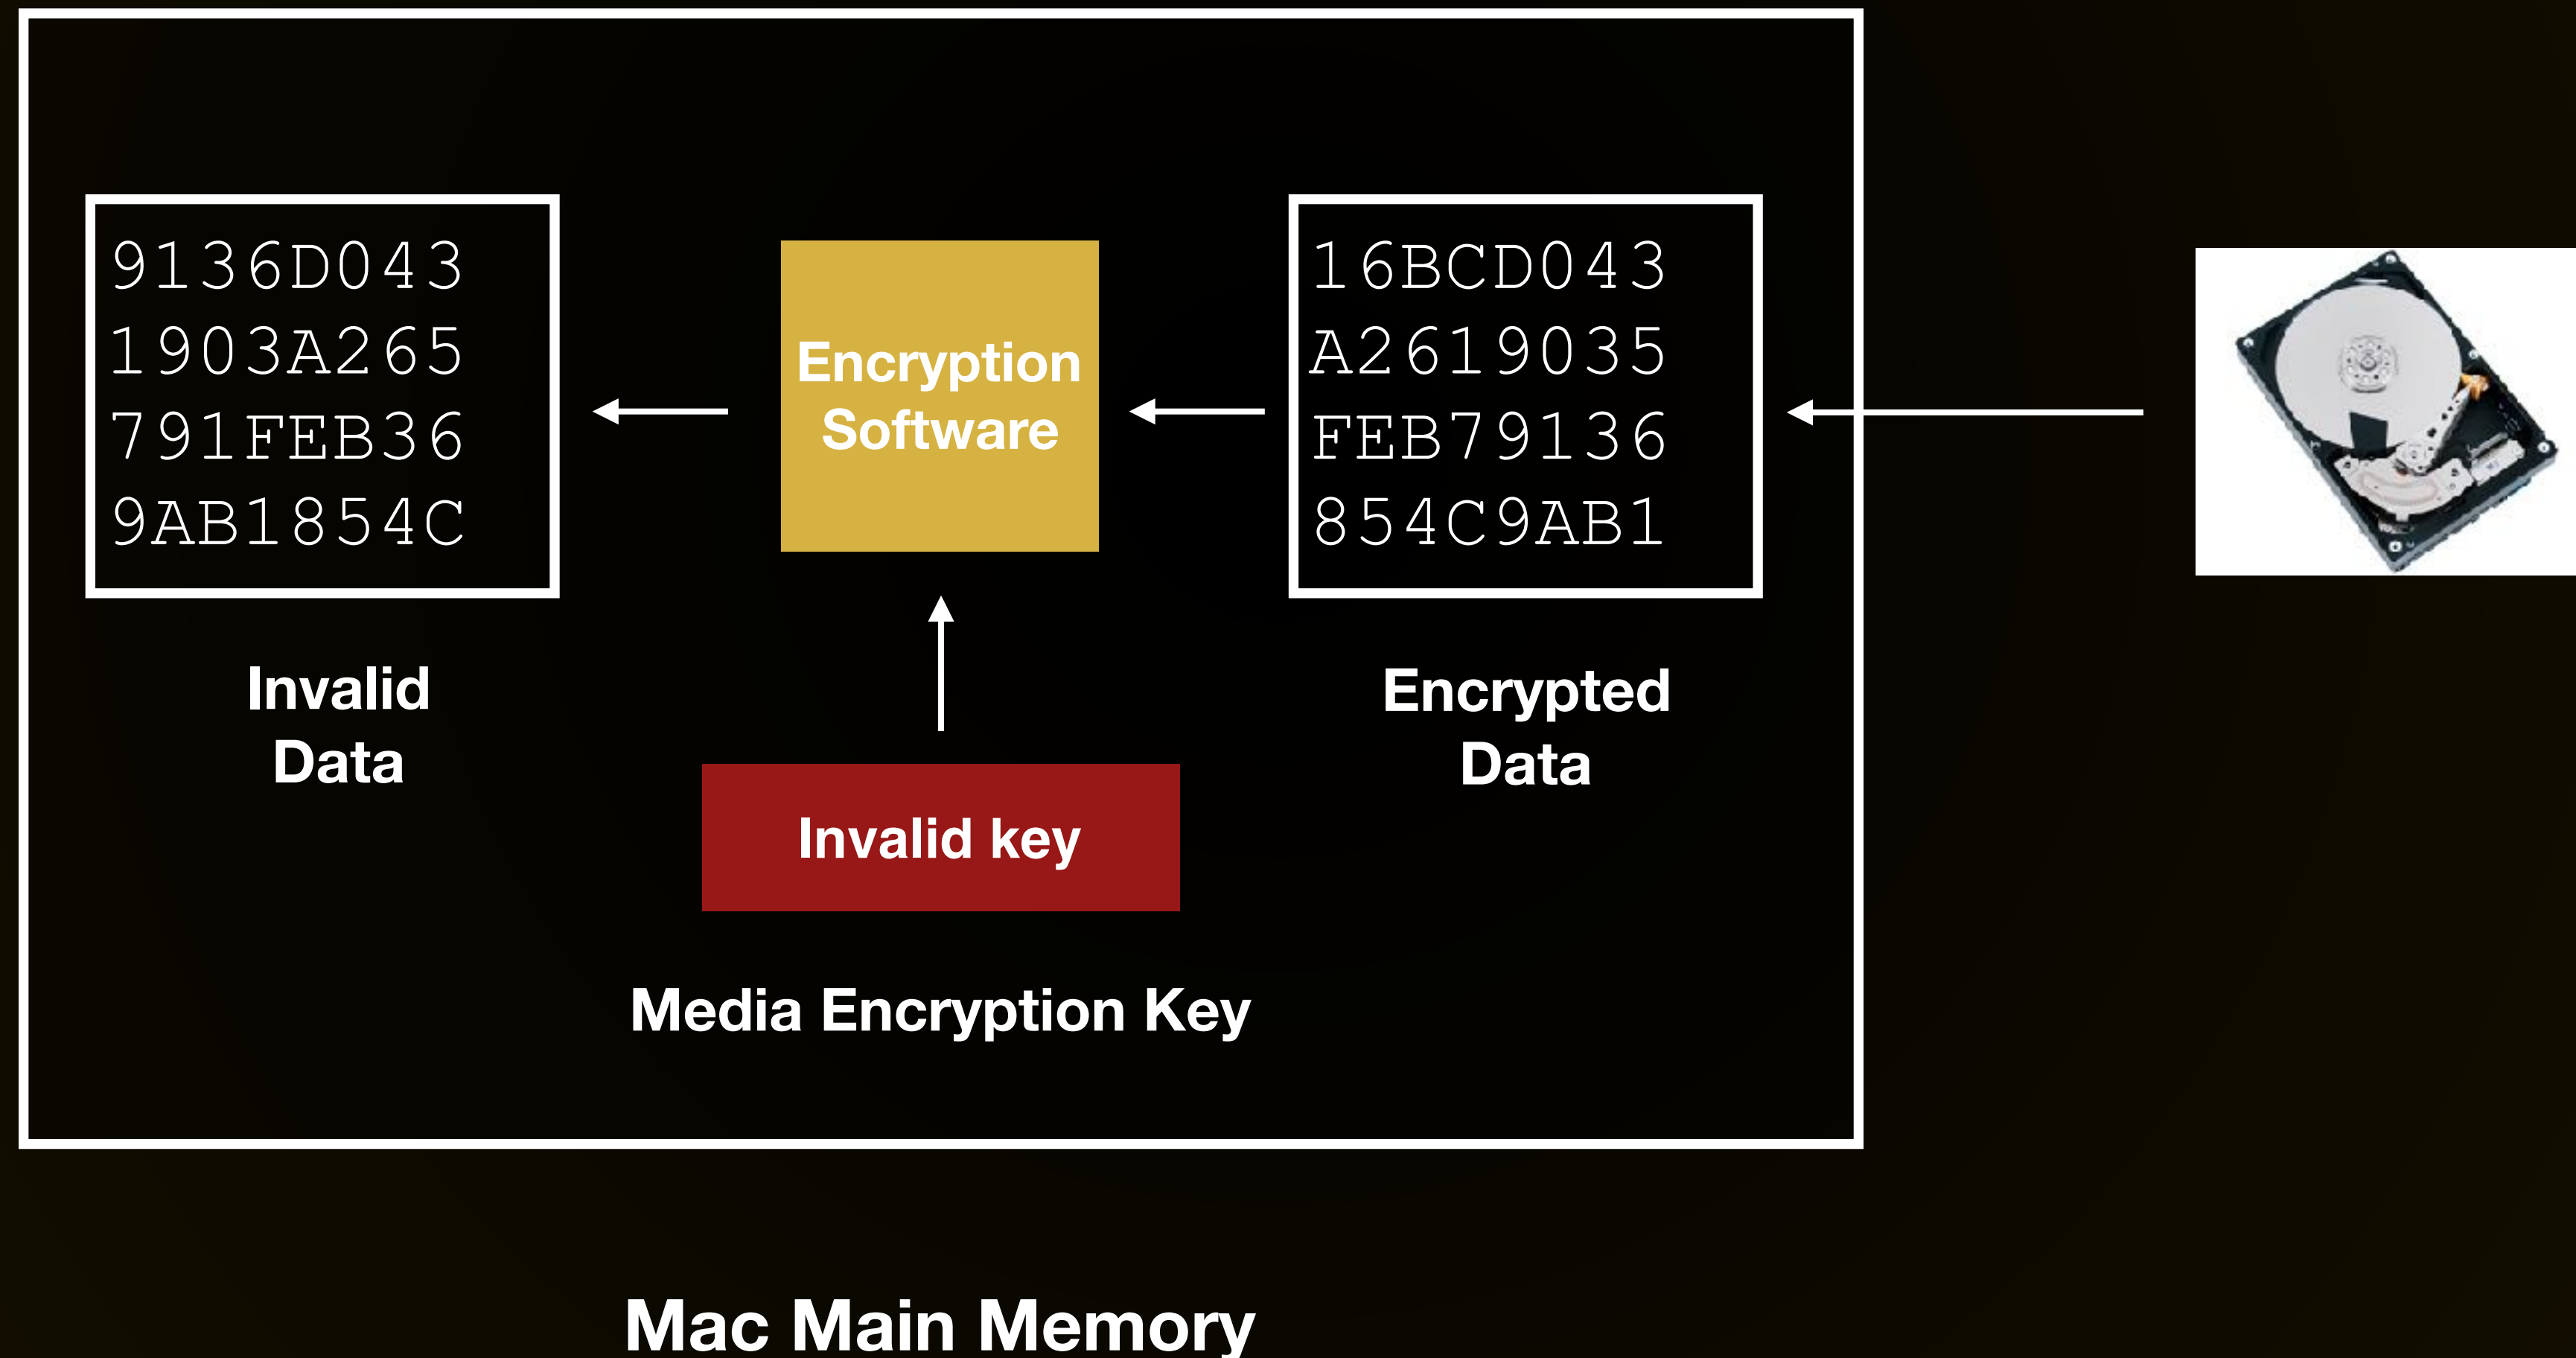
Invalid Key Encryption Key on Disk





FileVault Volume

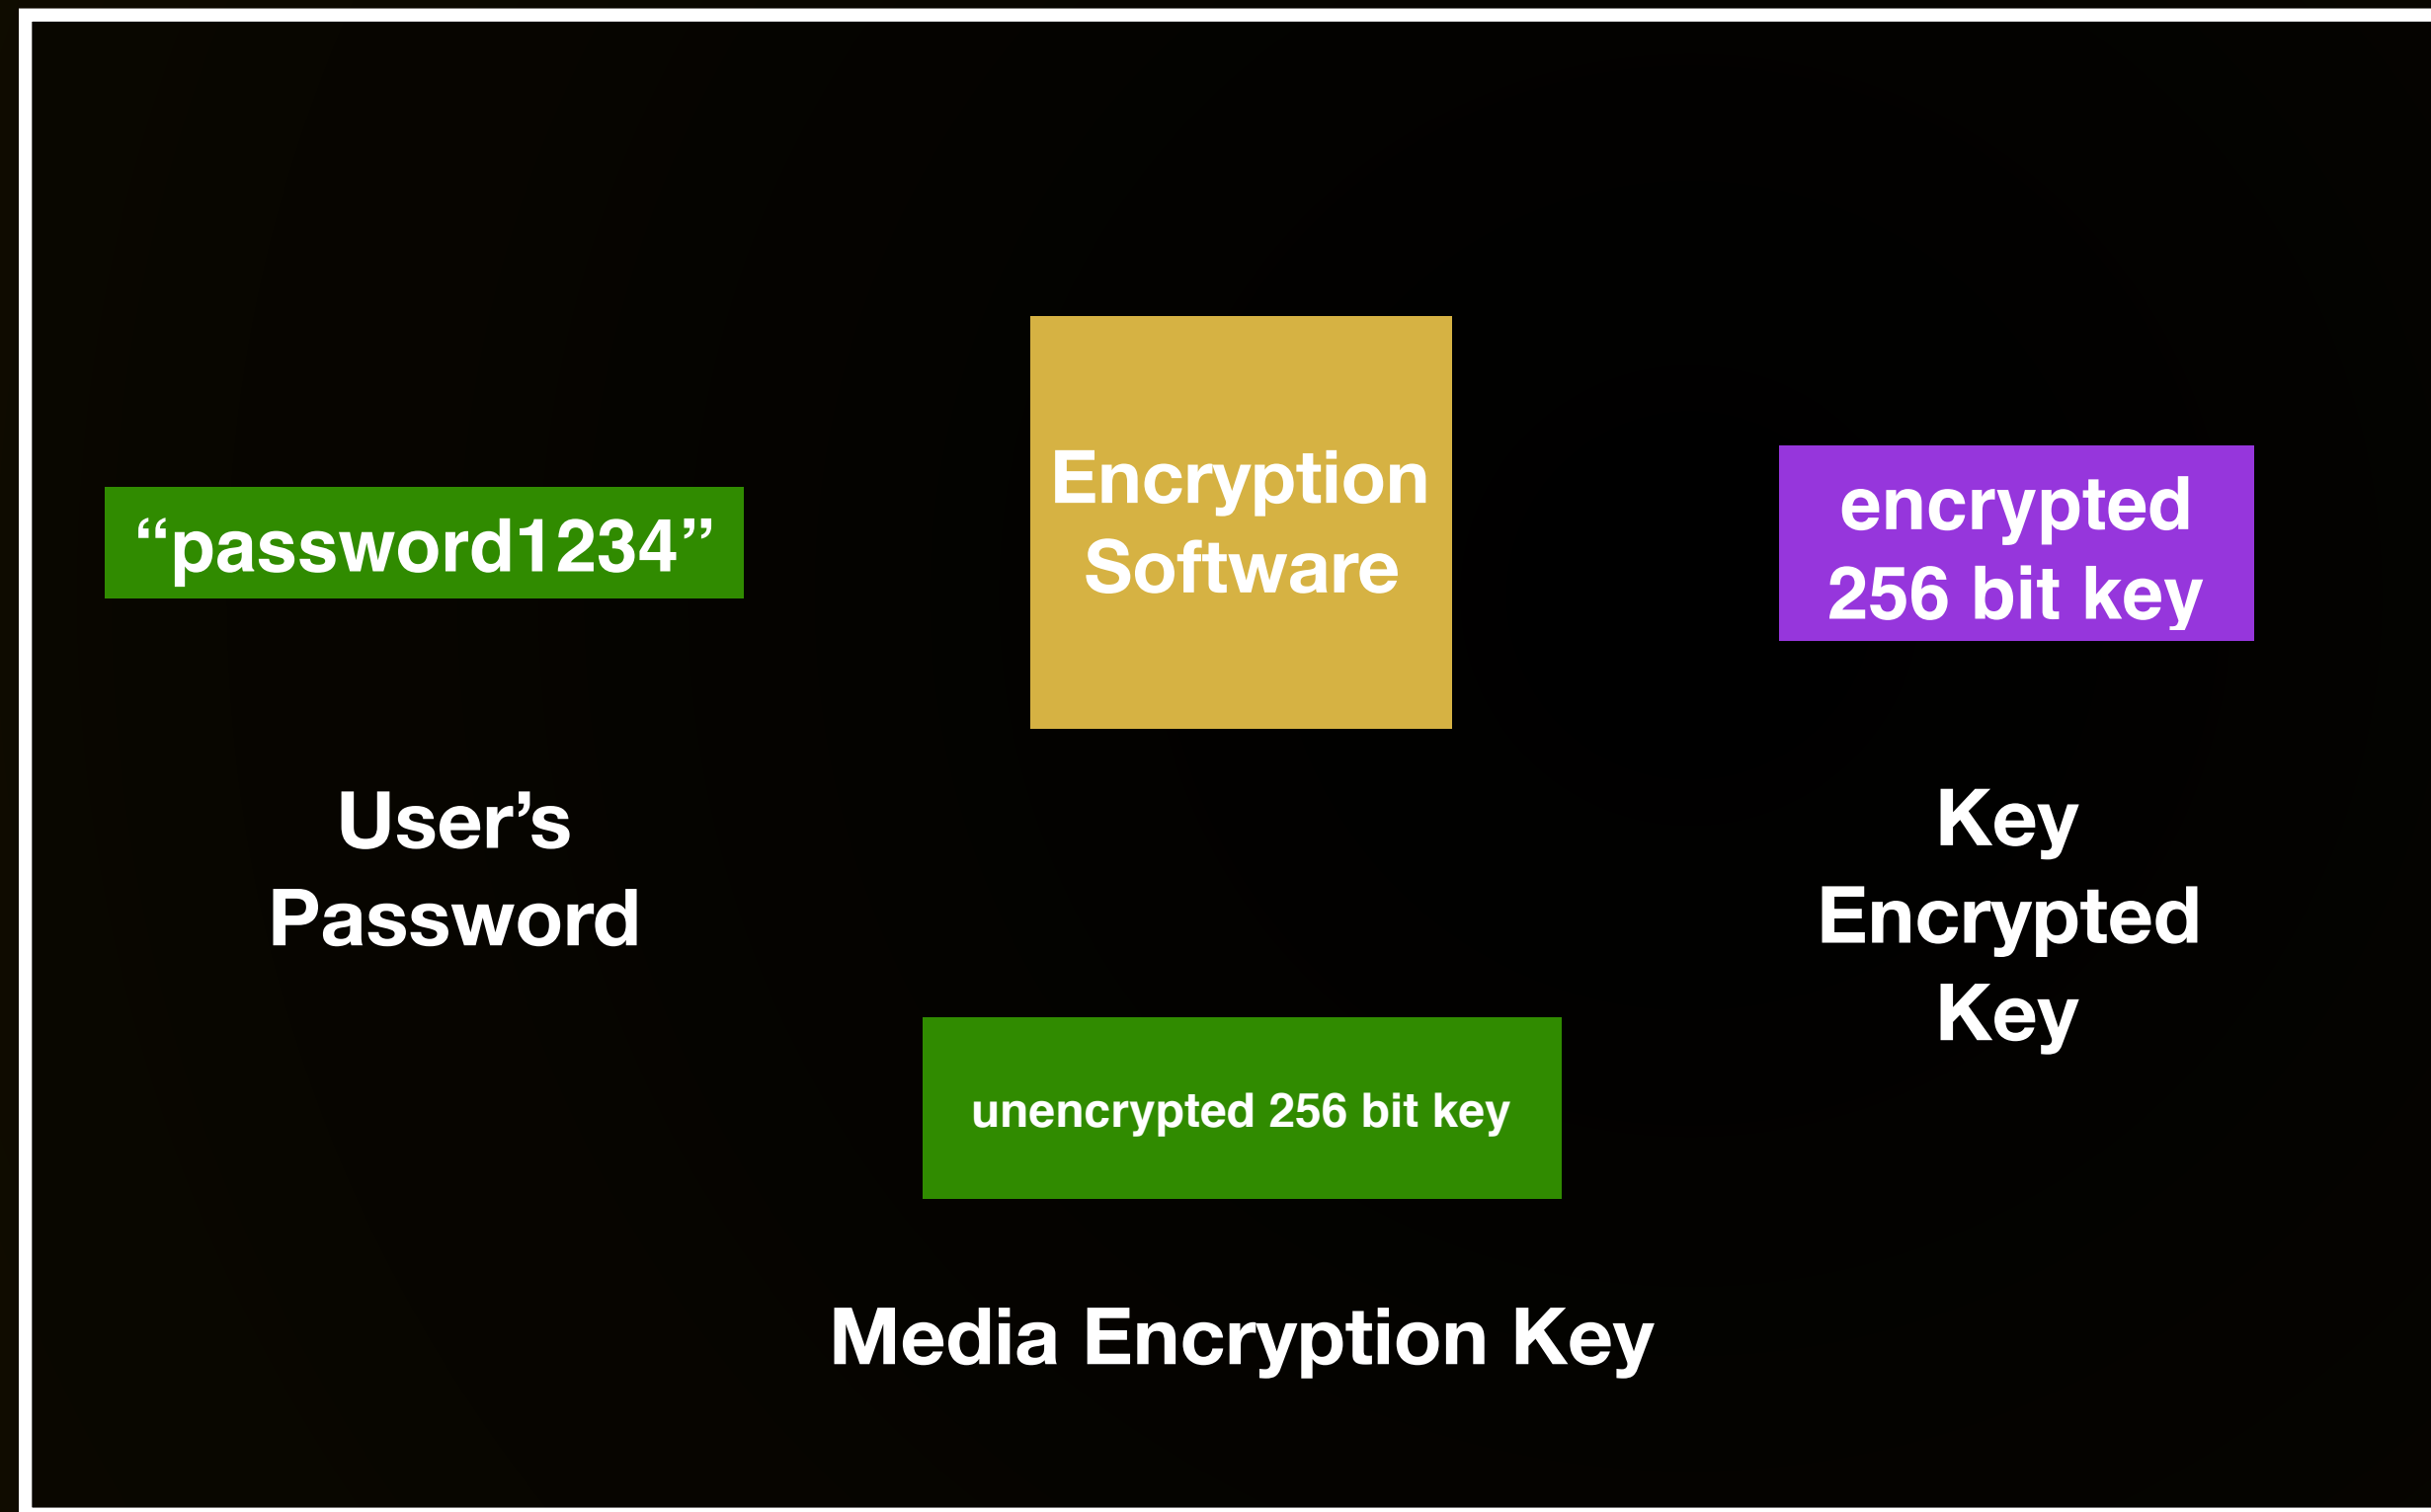
Reading Files with an Invalid Key Encryption Key





FileVault Volume

Key Encrypted Key is Stored on Disk



Mac Main Memory



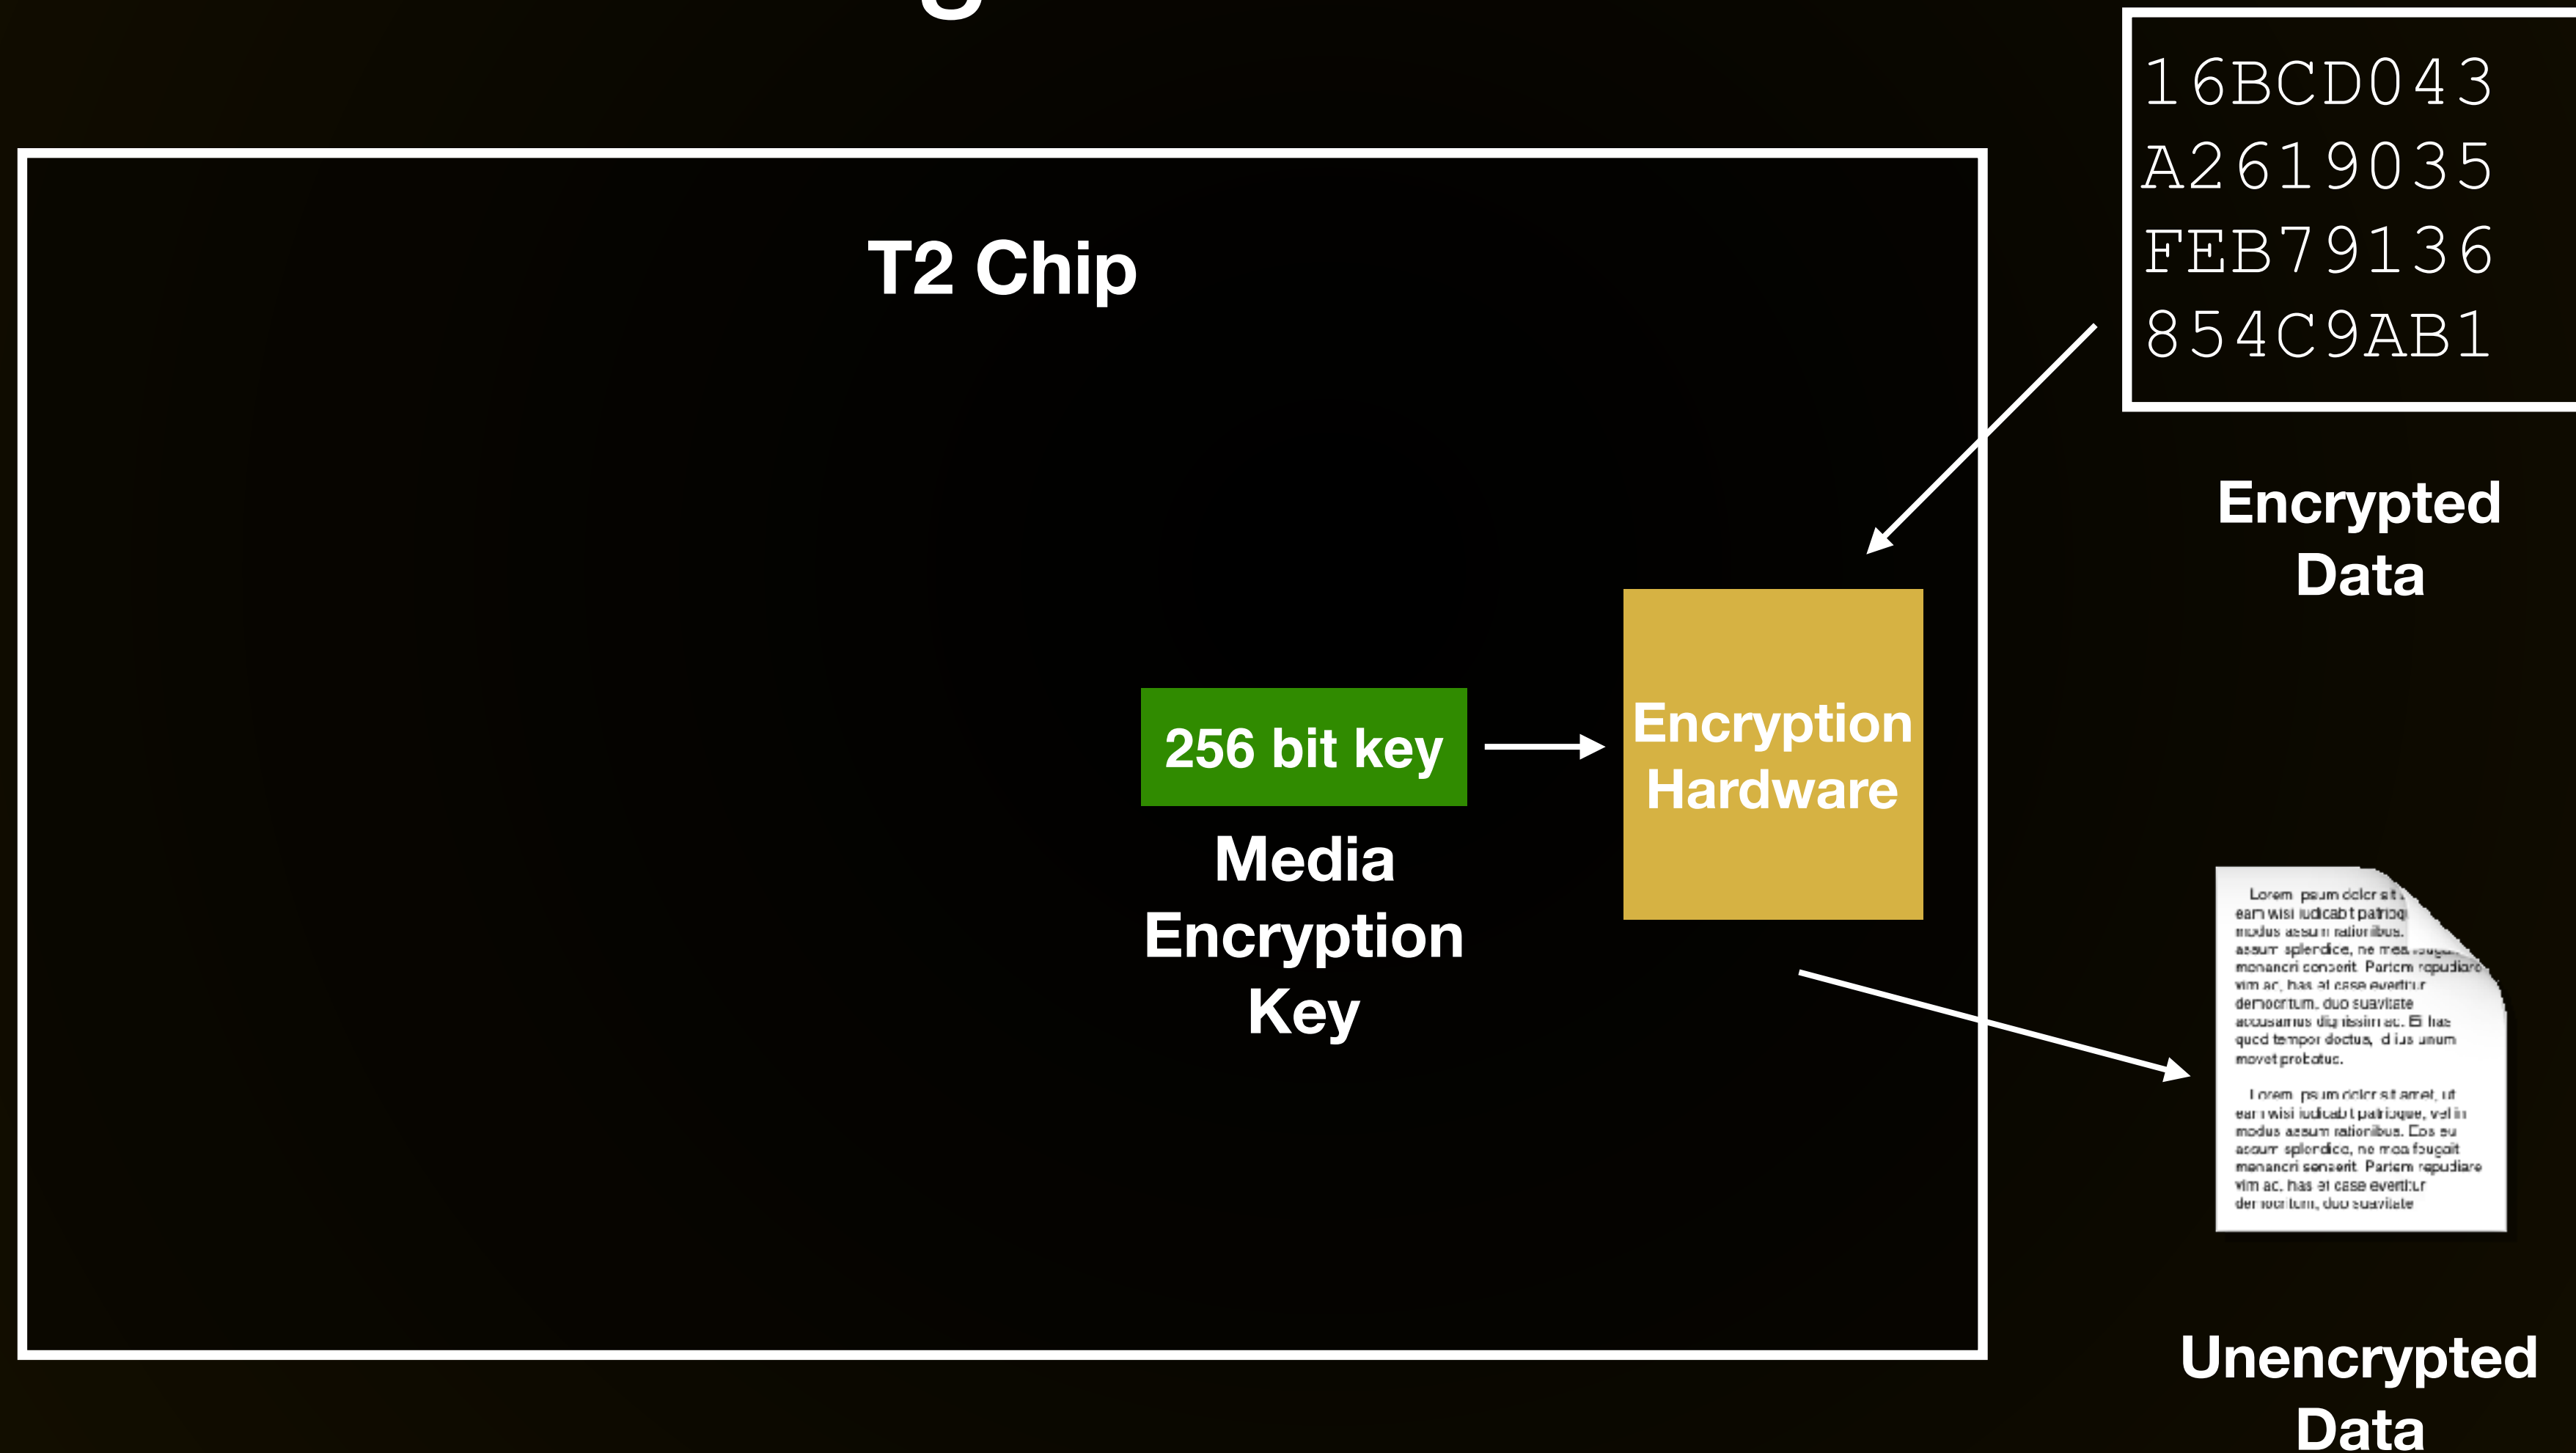
How is the T2 chip used for APFS



APFS Encryption With The T2 Chip



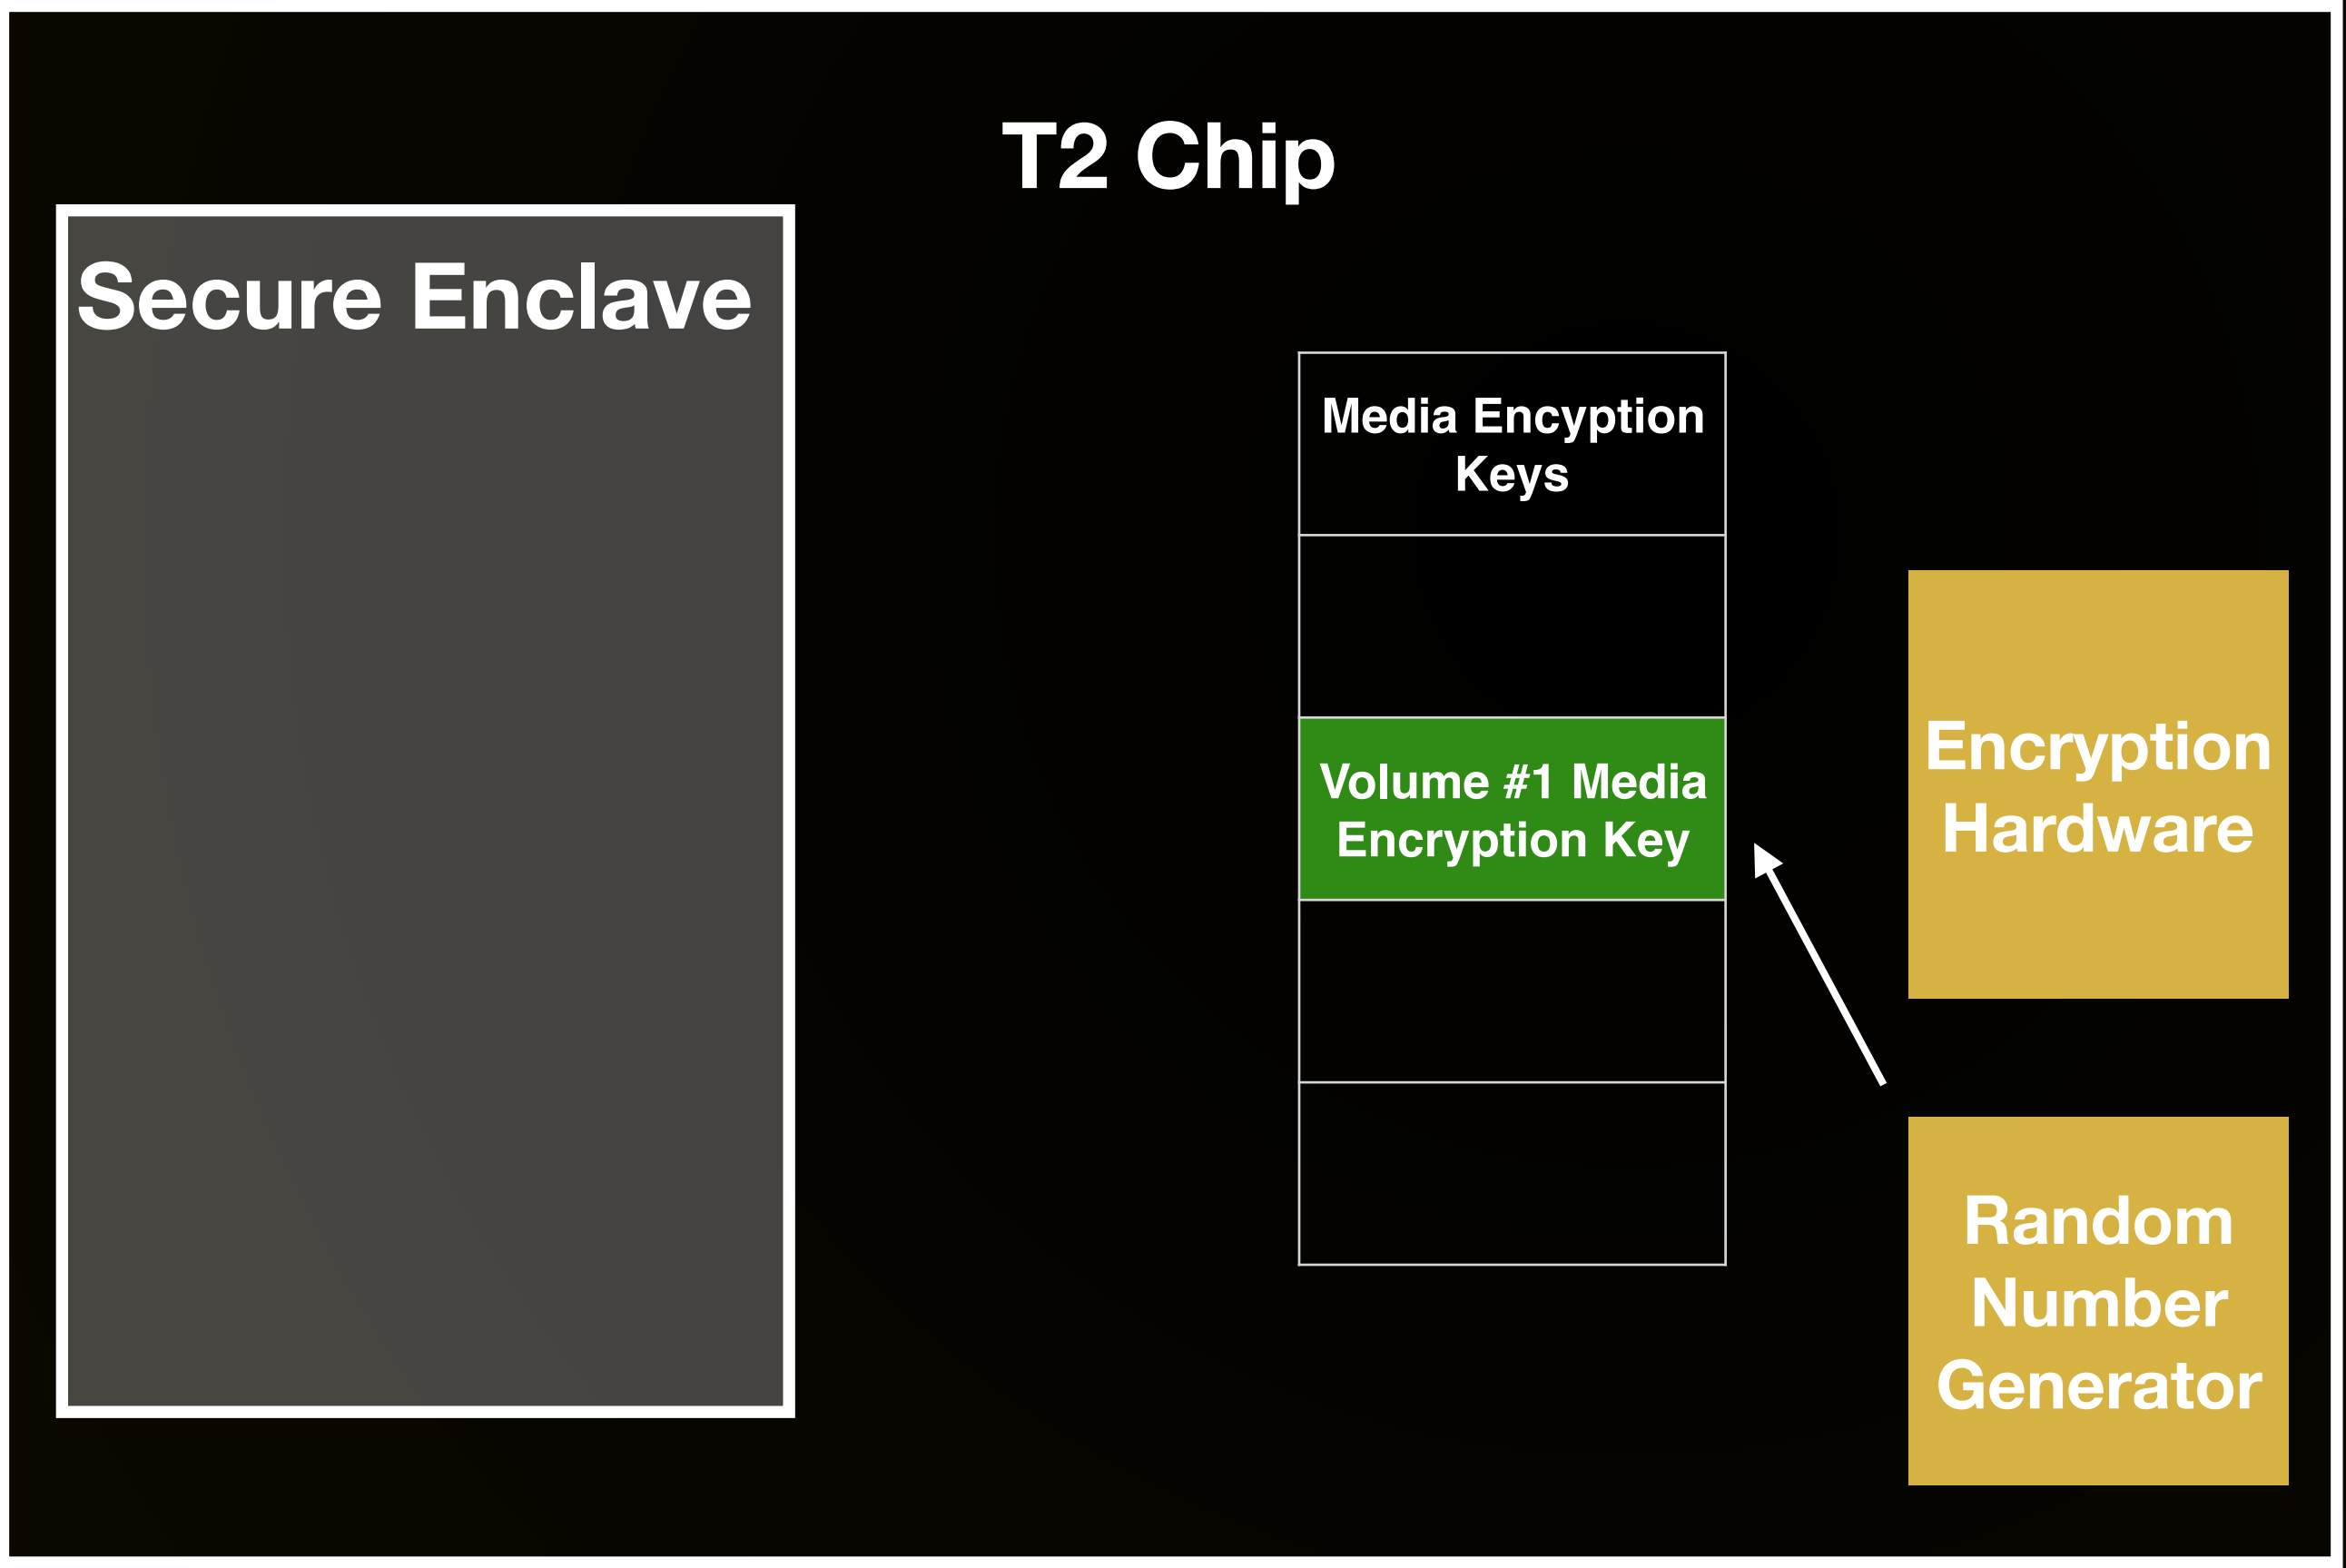
Reading from a Volume



APFS Encryption With The T2 Chip



Creating a new APFS volume



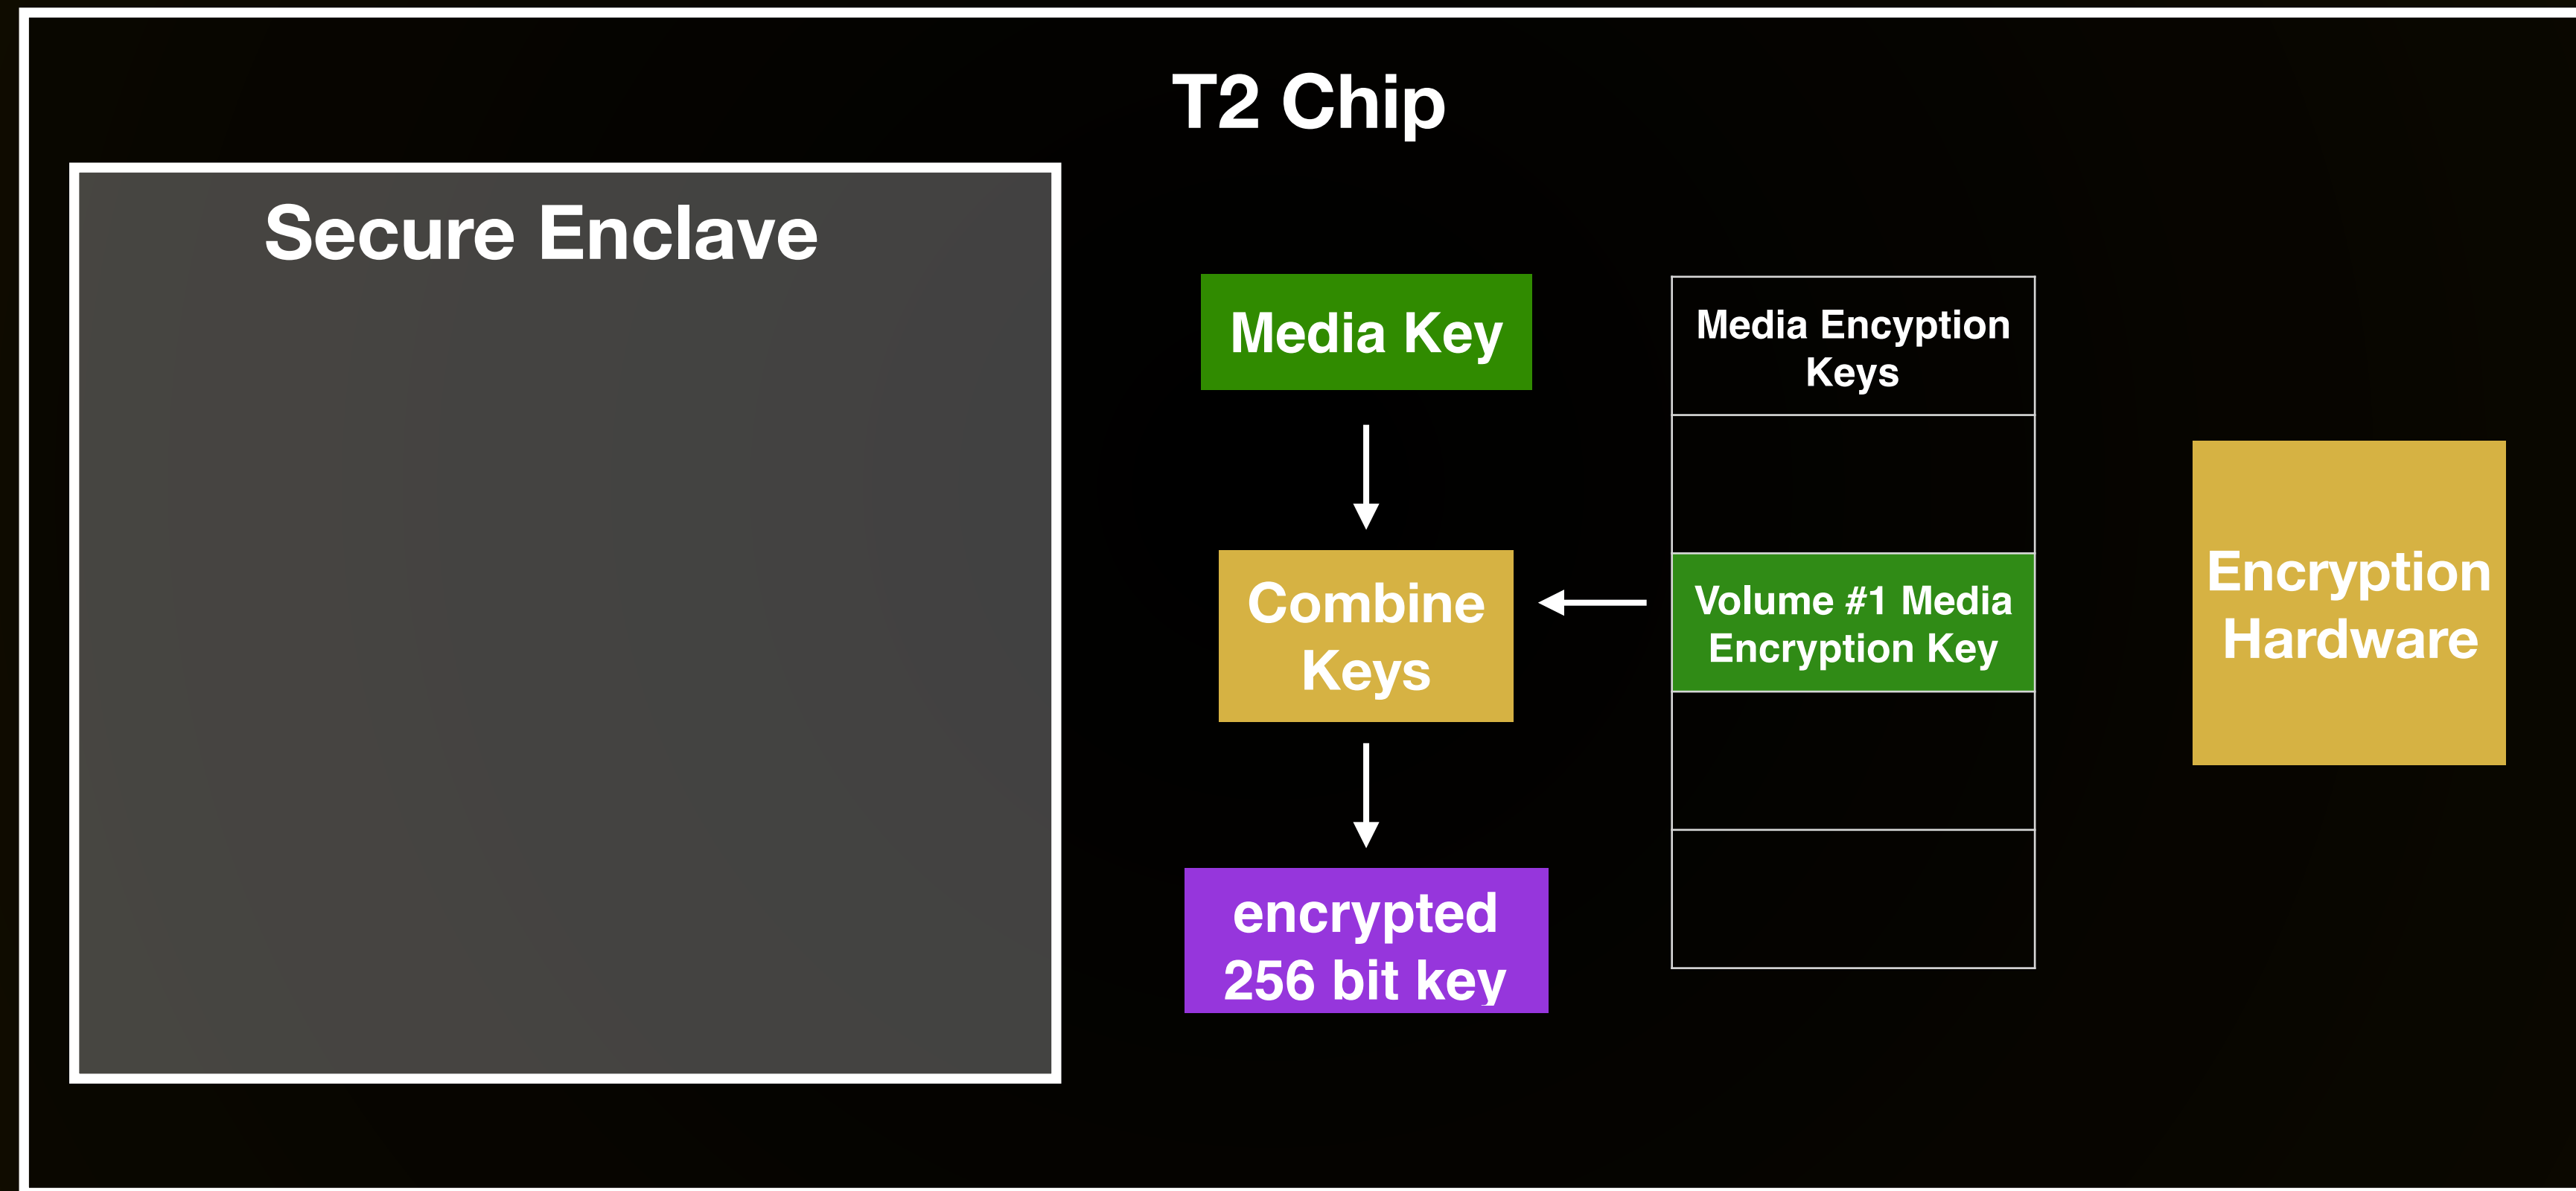
Object	Object ID
Container	401
Volume #1	526
Volume #2	123
Volume #3	996

AFP
Volume
Info

APFS Encryption With The T2 Chip



Creating a new APFS volume

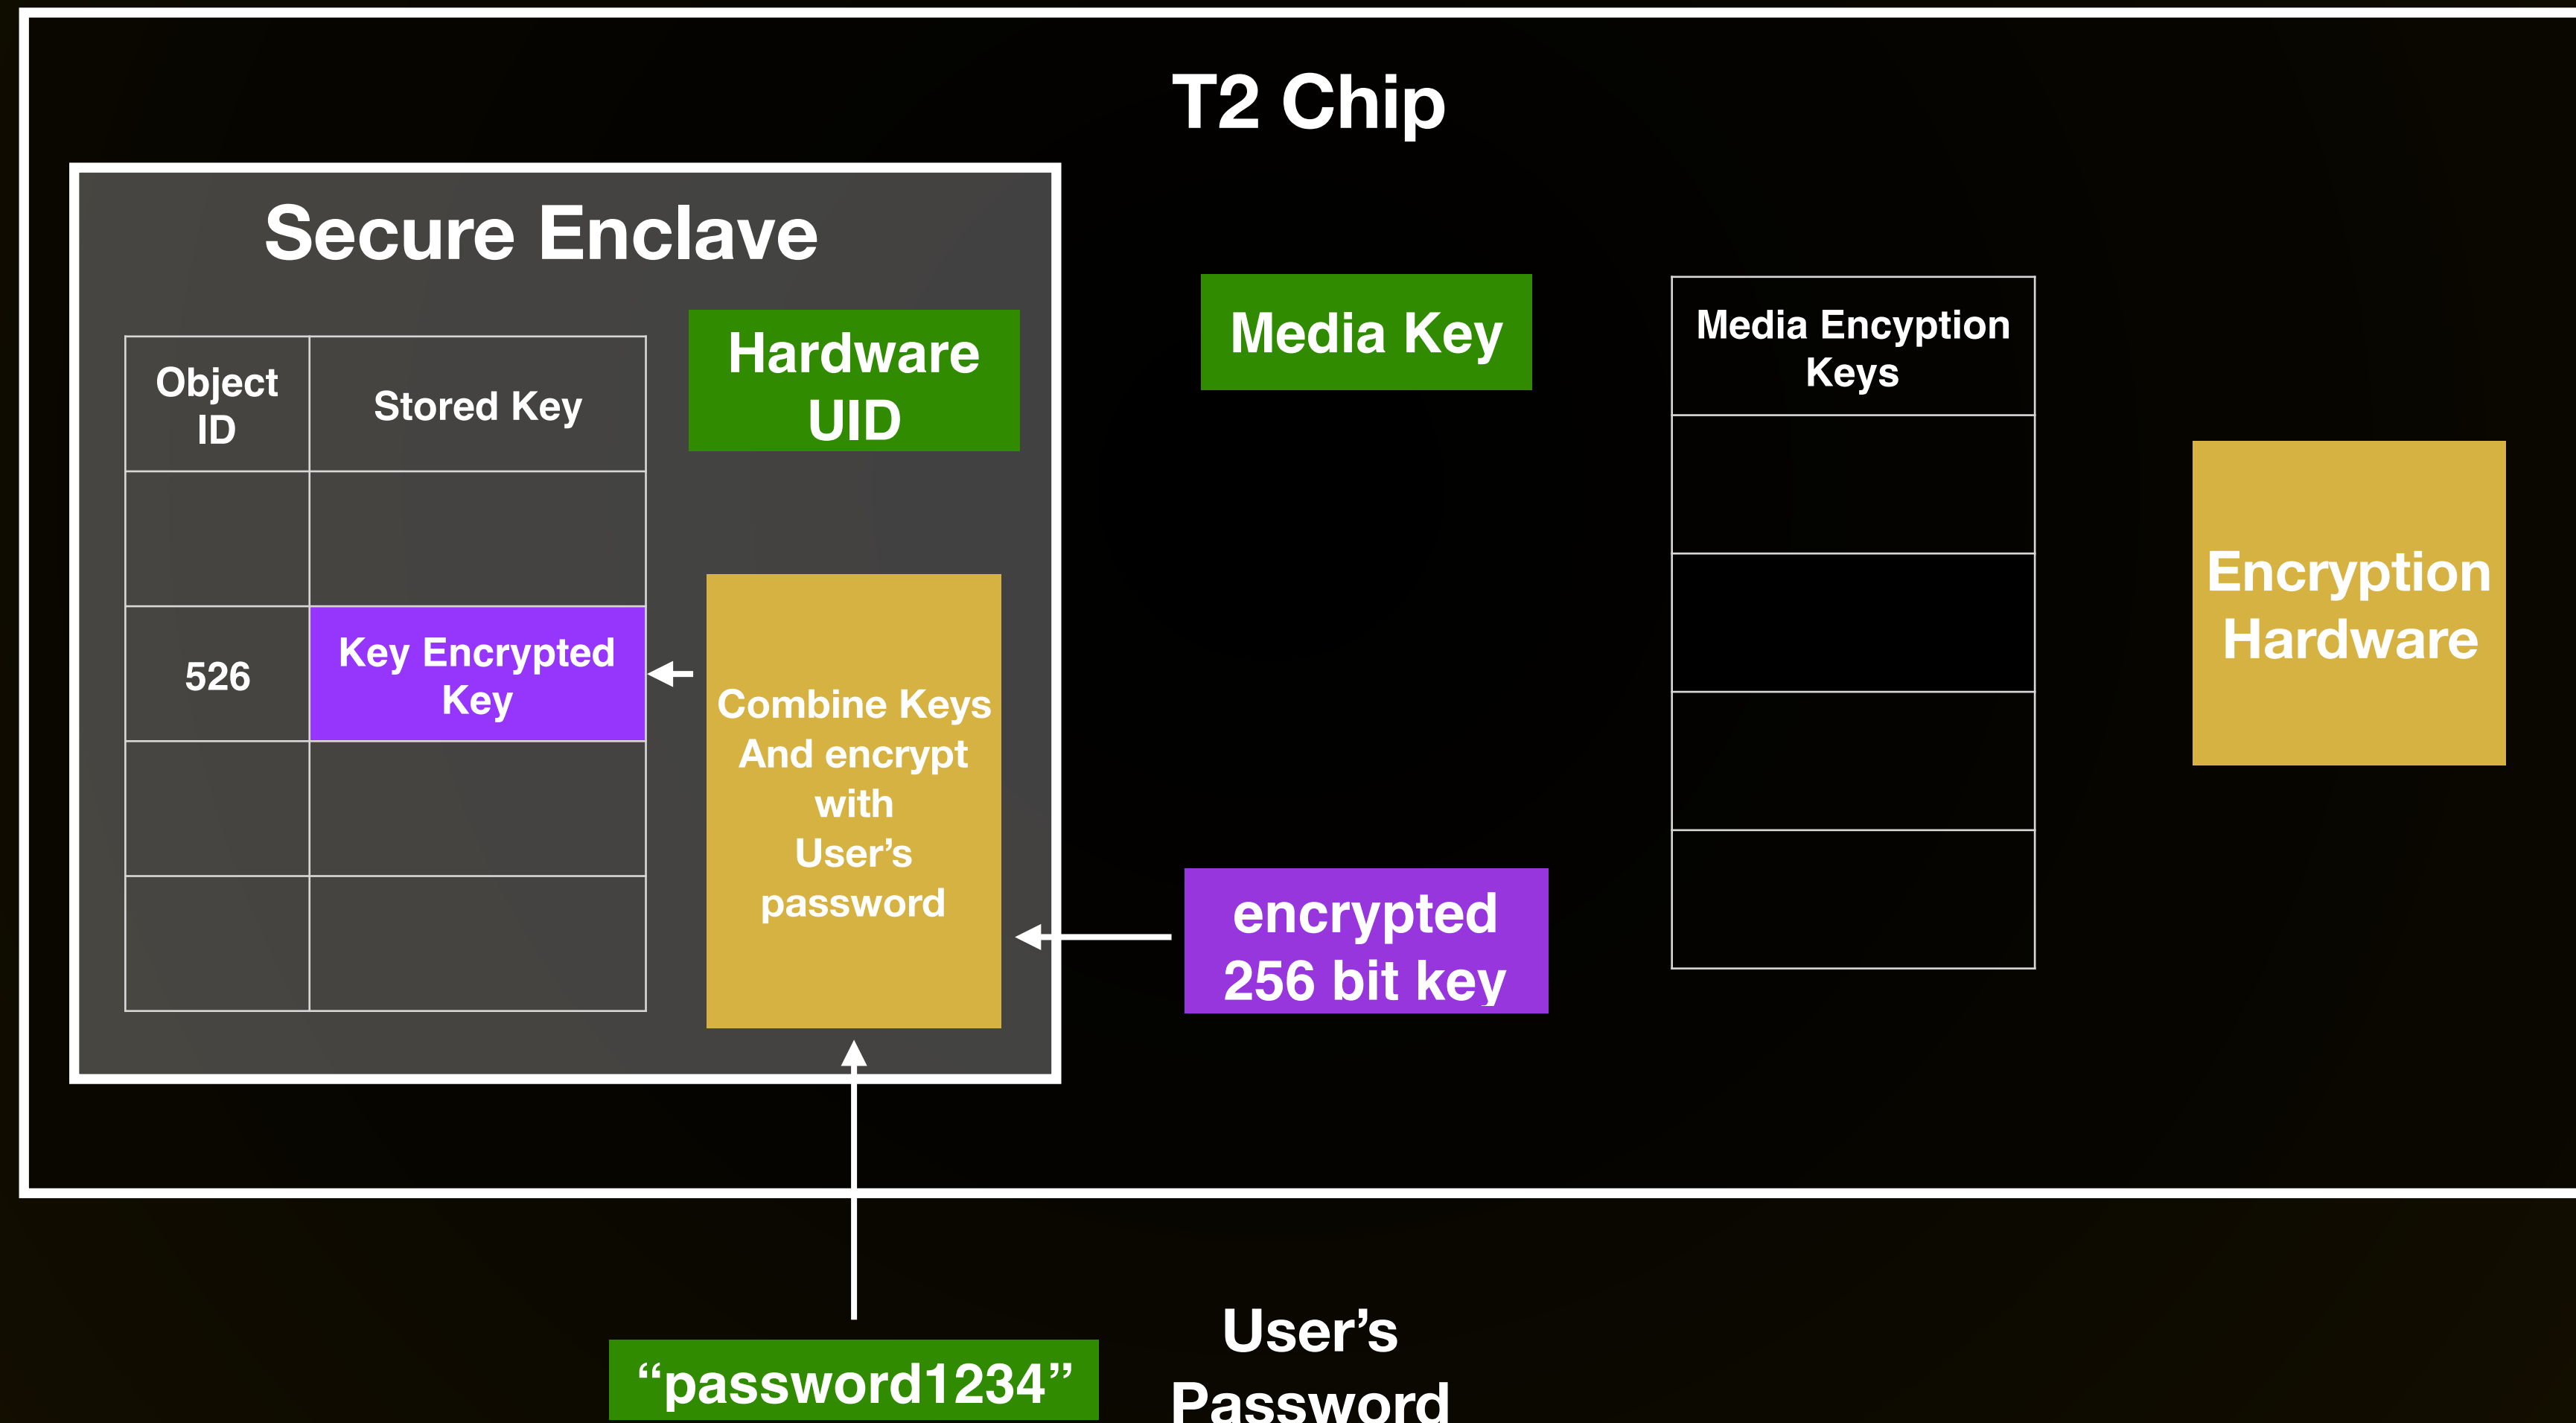


Media Key used for Remote Wipe feature in Find My Mac

APFS Encryption With The T2 Chip



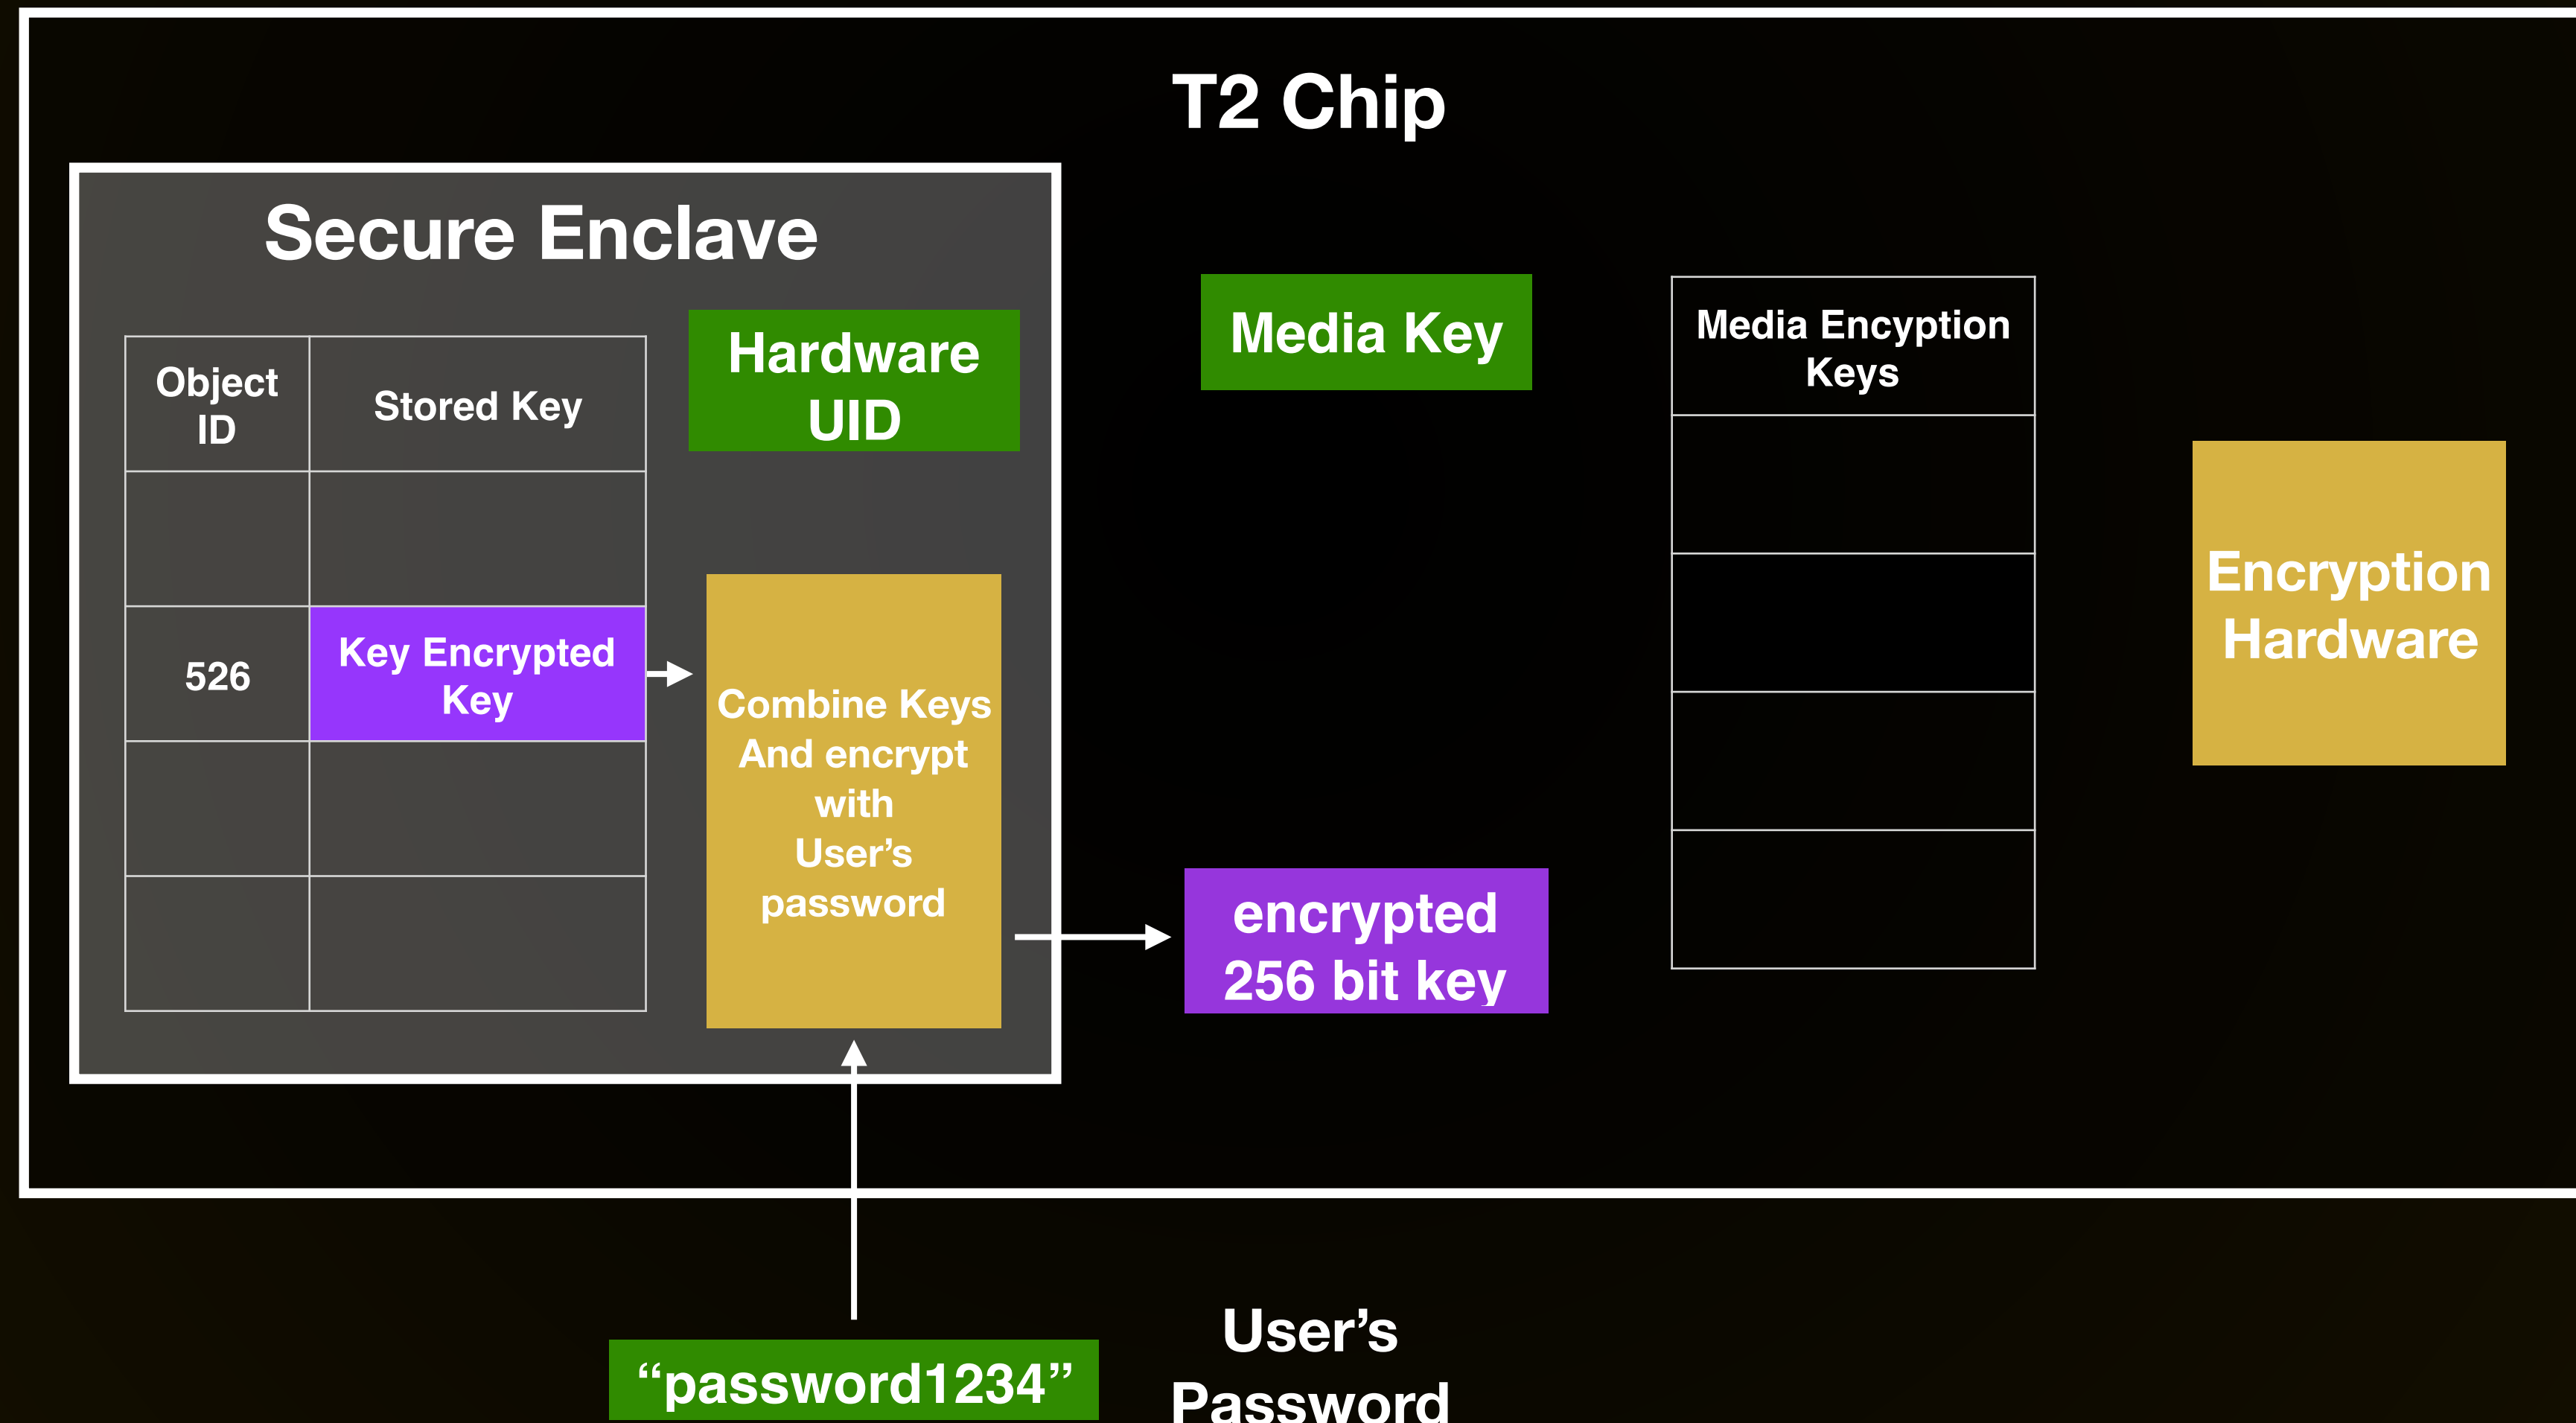
Creating a new APFS volume



APFS Encryption With The T2 Chip



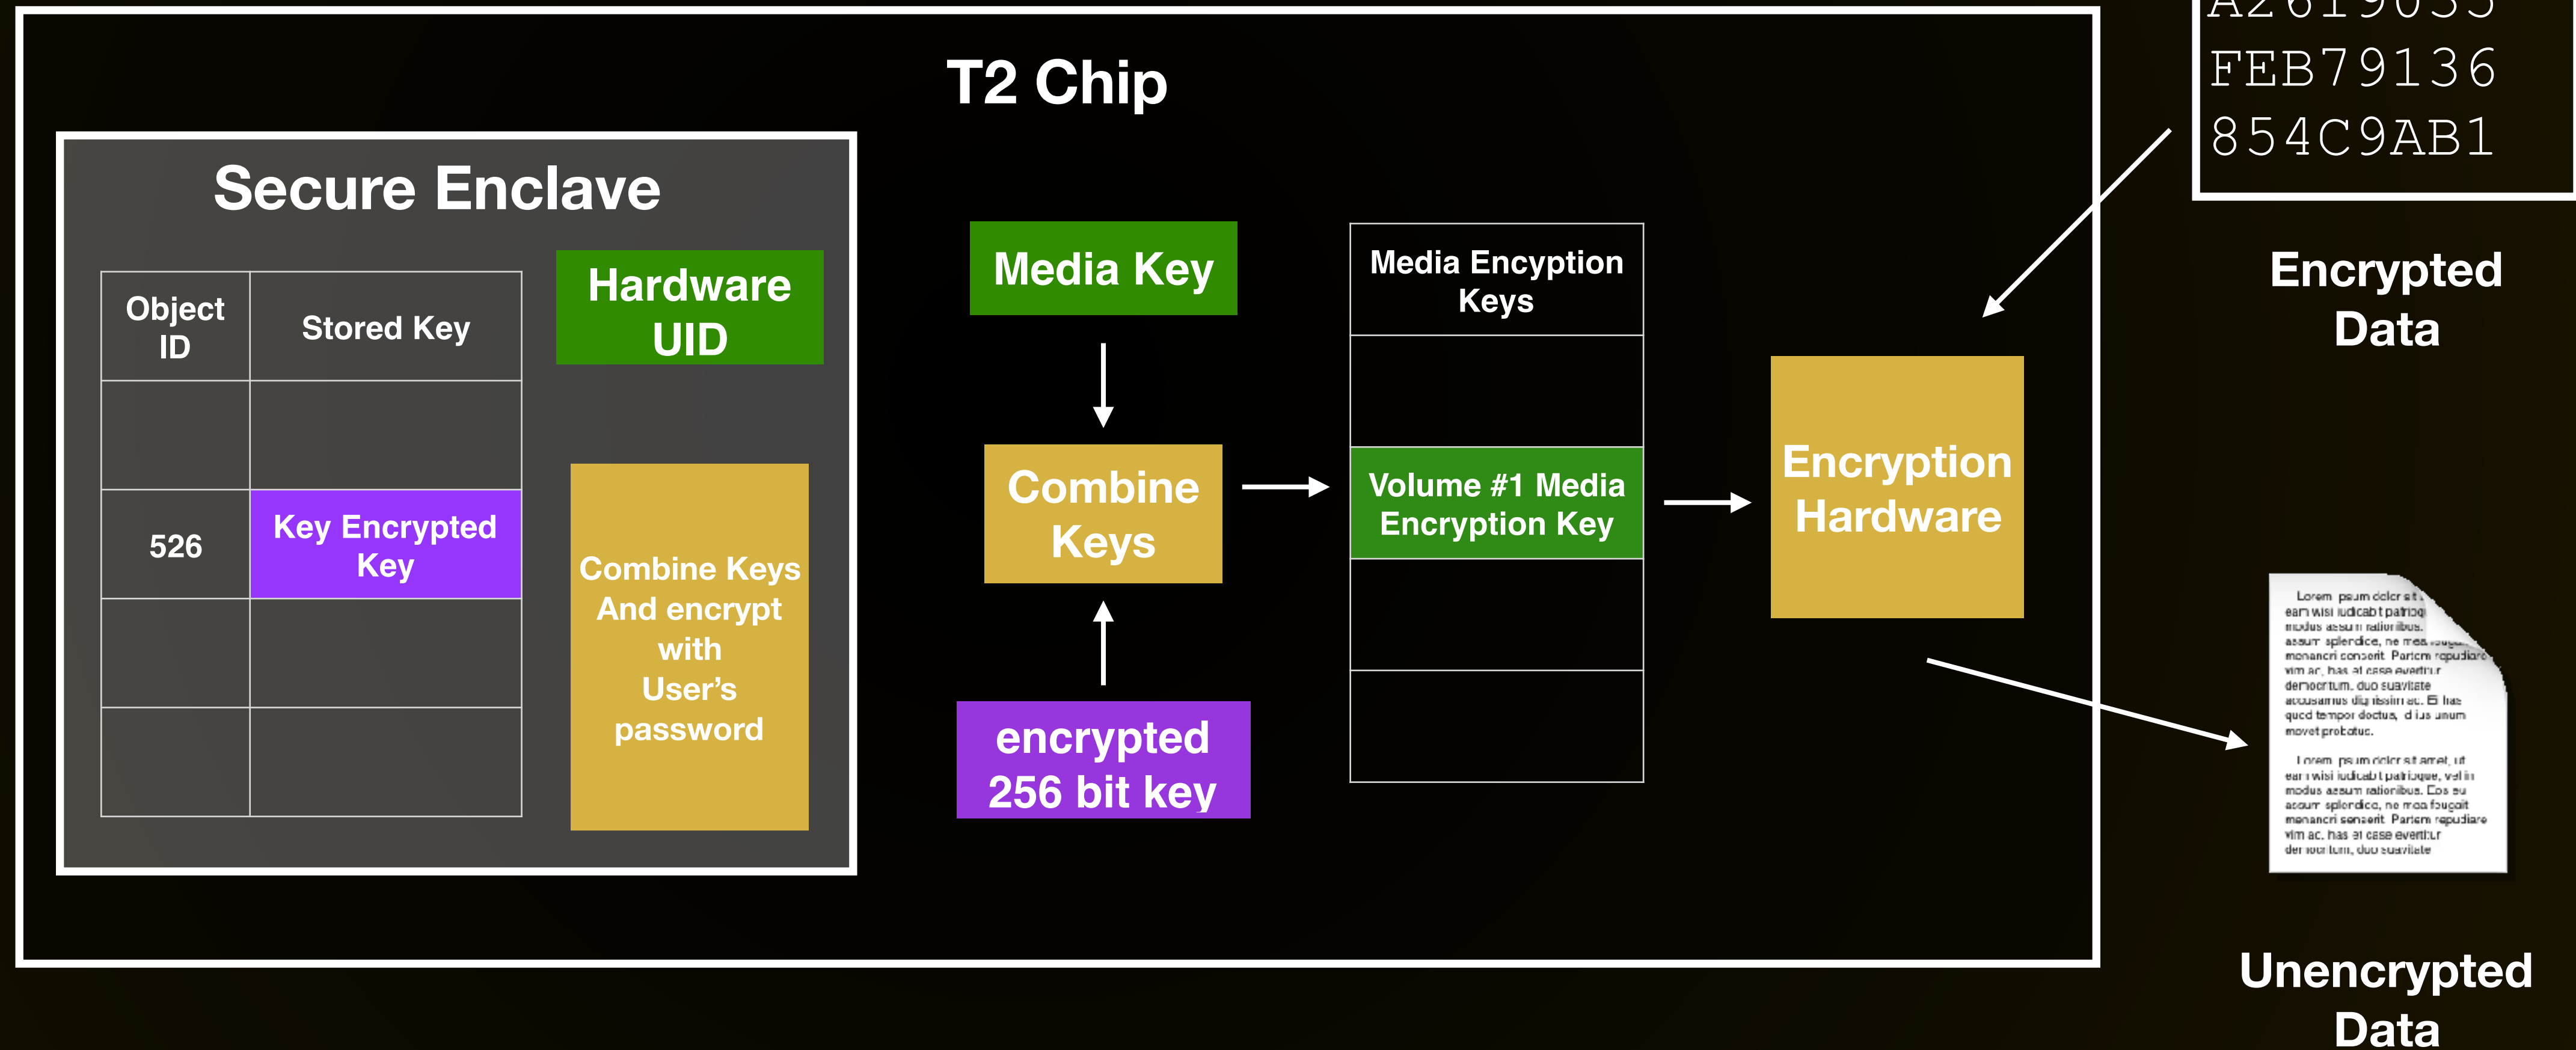
Mounting an Encrypted APFS volume



APFS Encryption With The T2 Chip



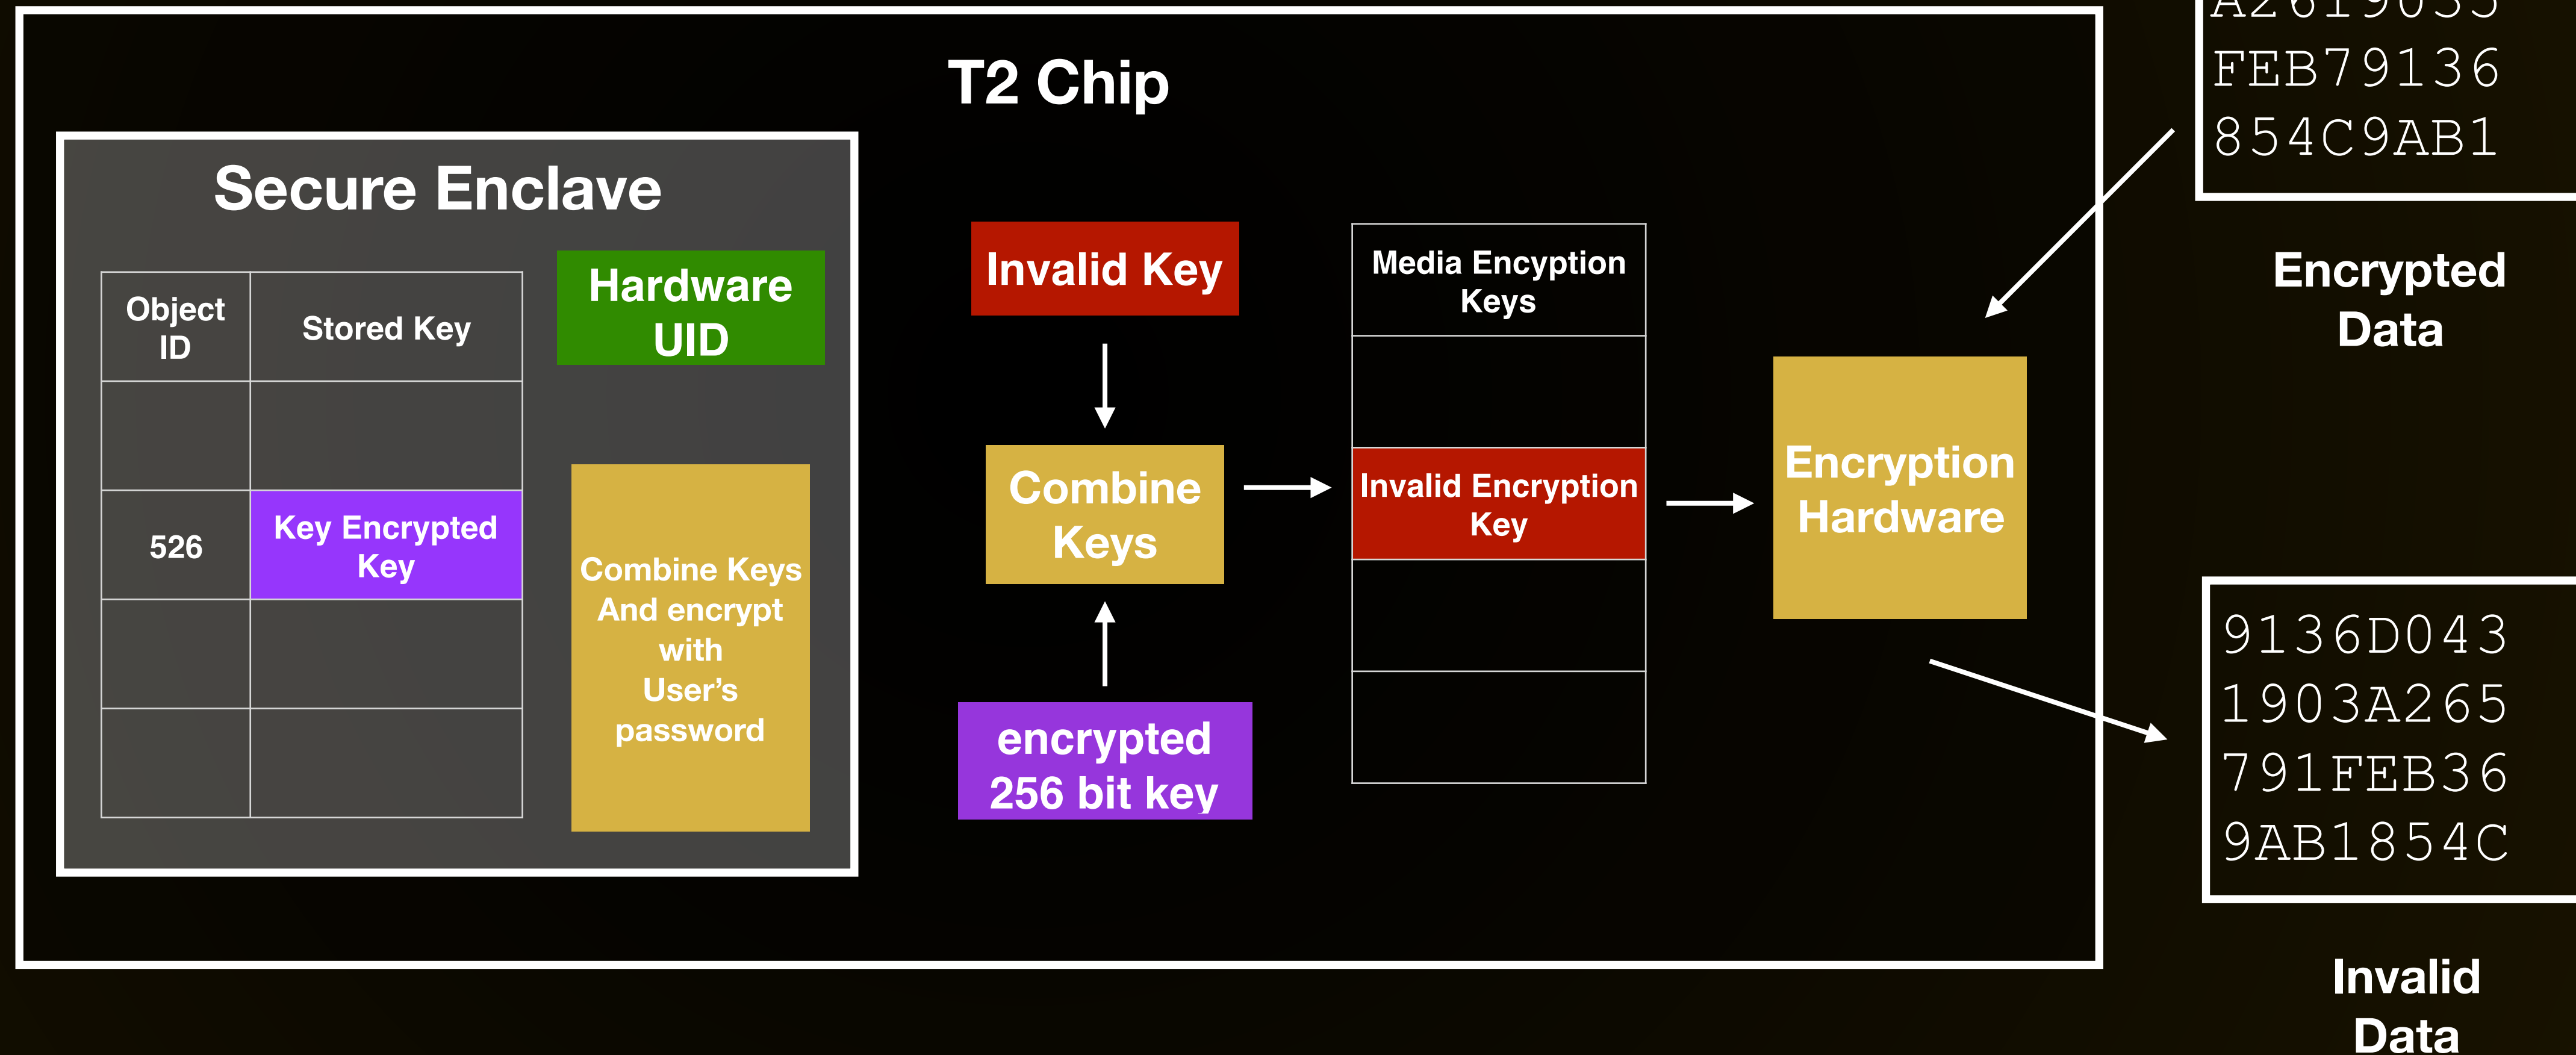
Mounting an Encrypted APFS volume



APFS Encryption With The T2 Chip



Accessing a Volume after a Remote Wipe



APFS Encryption With The T2 Chip



Advantages:

- **Media Encryption Key never in Mac memory**
- **Remote Wipe harder to circumvent**
- **Key Encrypted Key is stored in Secure Enclave not on disk**



OWC SHOP

Advanced Search

Categories

- OWC Internal SSD
- OWC External Storage
- OWC Memory
- HD / SSD Mountings
- Enclosures
- Batteries & Chargers
- Drive Docking Solutions
- iPhone / iPad / iPod
- Mounts
- Accessories
- Cables / Adapters
- Expansion Cards

[Catalog](#)[What's New?](#)[Shopping Cart](#)[Sign in](#)[My Account](#)

OWC Envoy Pro EX — USB-C

The fastest USB-C SSD ever built.

[BUY NOW >](#)

Up to 980MB/s



Up to 2TB



Mac & PC



<https://owcshop.eu>



Q & A