



The Magic Which Makes macos Volumes Fast

Tim Standing

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





The Magic Which Makes macos Volumes Fast and Reliable



Fast

- Storage Speed Increases
- PCIe Gen4
- Thunderbolt 5



Storage Speed Increases Year After Year



Why SoftRAID?



- Used in all OWC enclosures with 4 or more disks
- No hardware dependencies
- Much better performance
- Better disk health tracking
- Just plug in and go (with macOS 13.3 and later)



Testing Speed



- **AJA System Test Settings**
 - 64 GB Test File
 - 16 bit RGBA
 - 8192x4320 8K
 - DMA Mode: Dual Engines



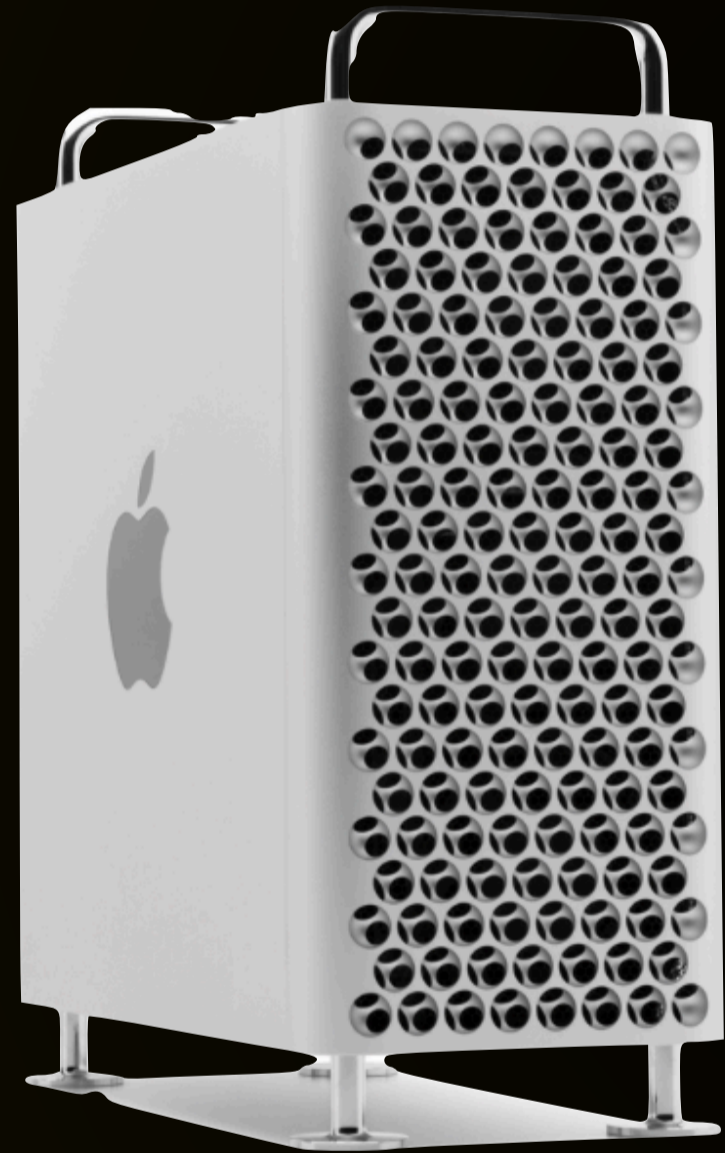
2018



2 ThunderBlades + 2018 MacBook Pro, RAID 0 = 4.0 GB/sec



2019



Accelsior 4M2 + 2019 Mac Pro, RAID 0 = 6 GB/sec



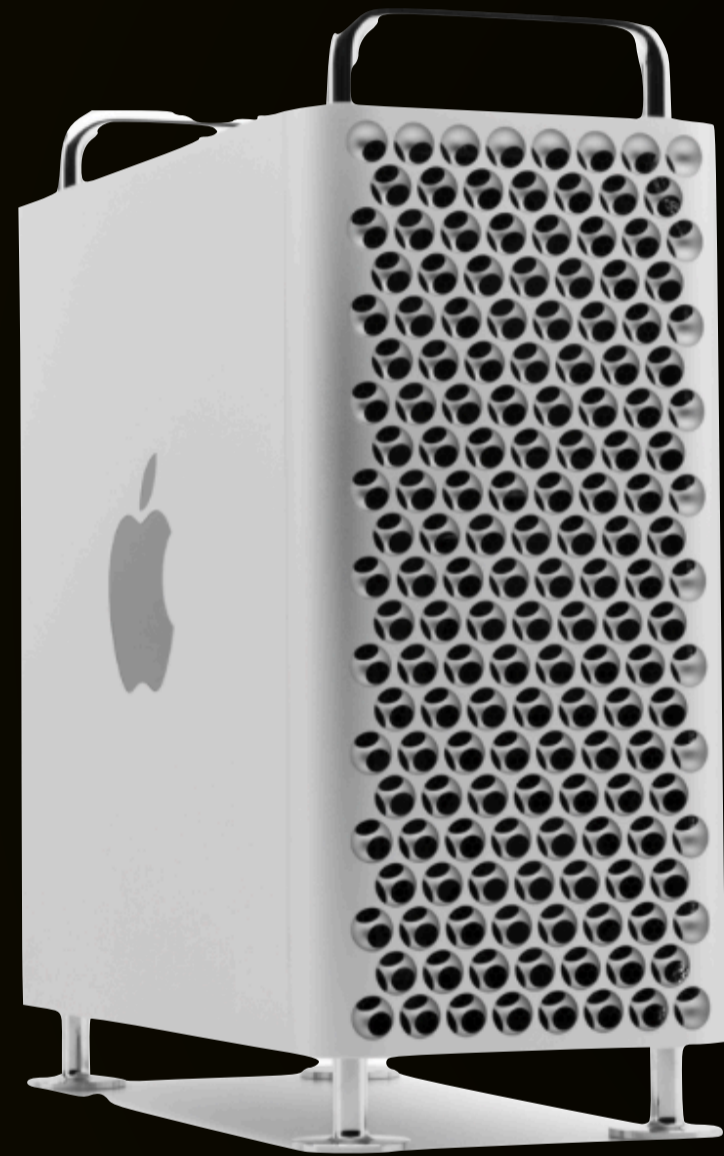
2022



4 Thunderblades + Mac Studio, RAID 0 = 10 GB/sec



2022



Accelsior 8M2 + 2019 Mac Pro, RAID 0 = 12 GB/sec



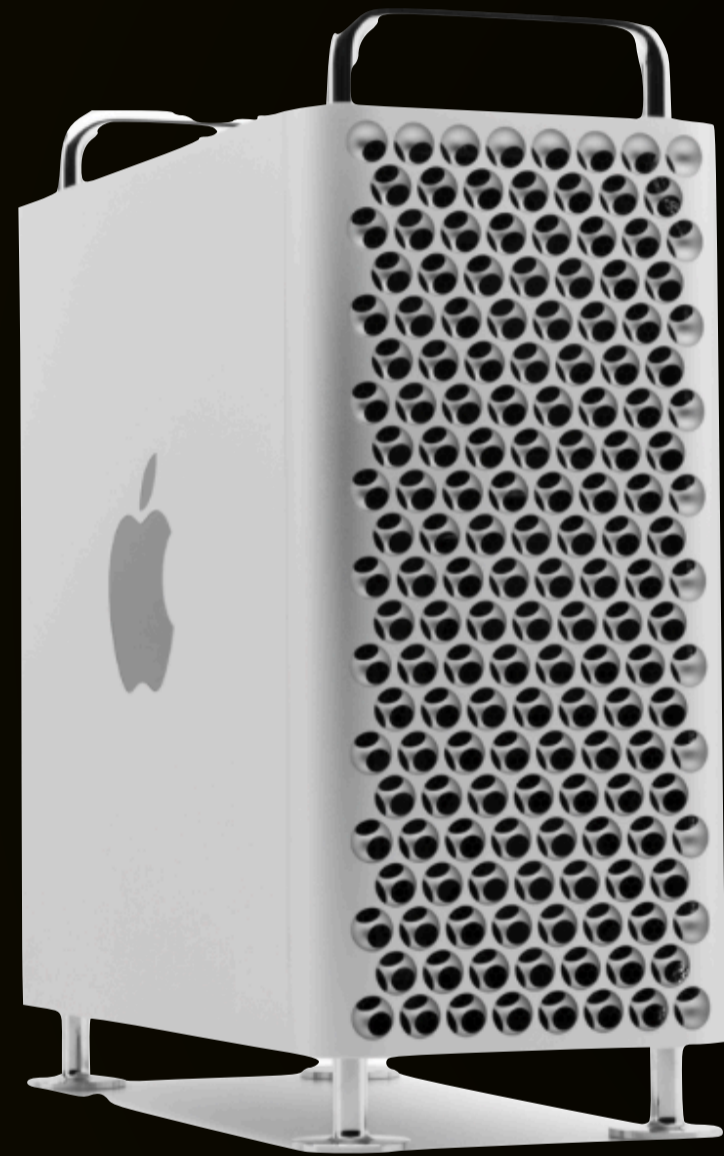
Accelsior 8M2



- 16 Lane - PCIe Gen 4
- Sizes up to 64 TB
- 8 NVMe blades



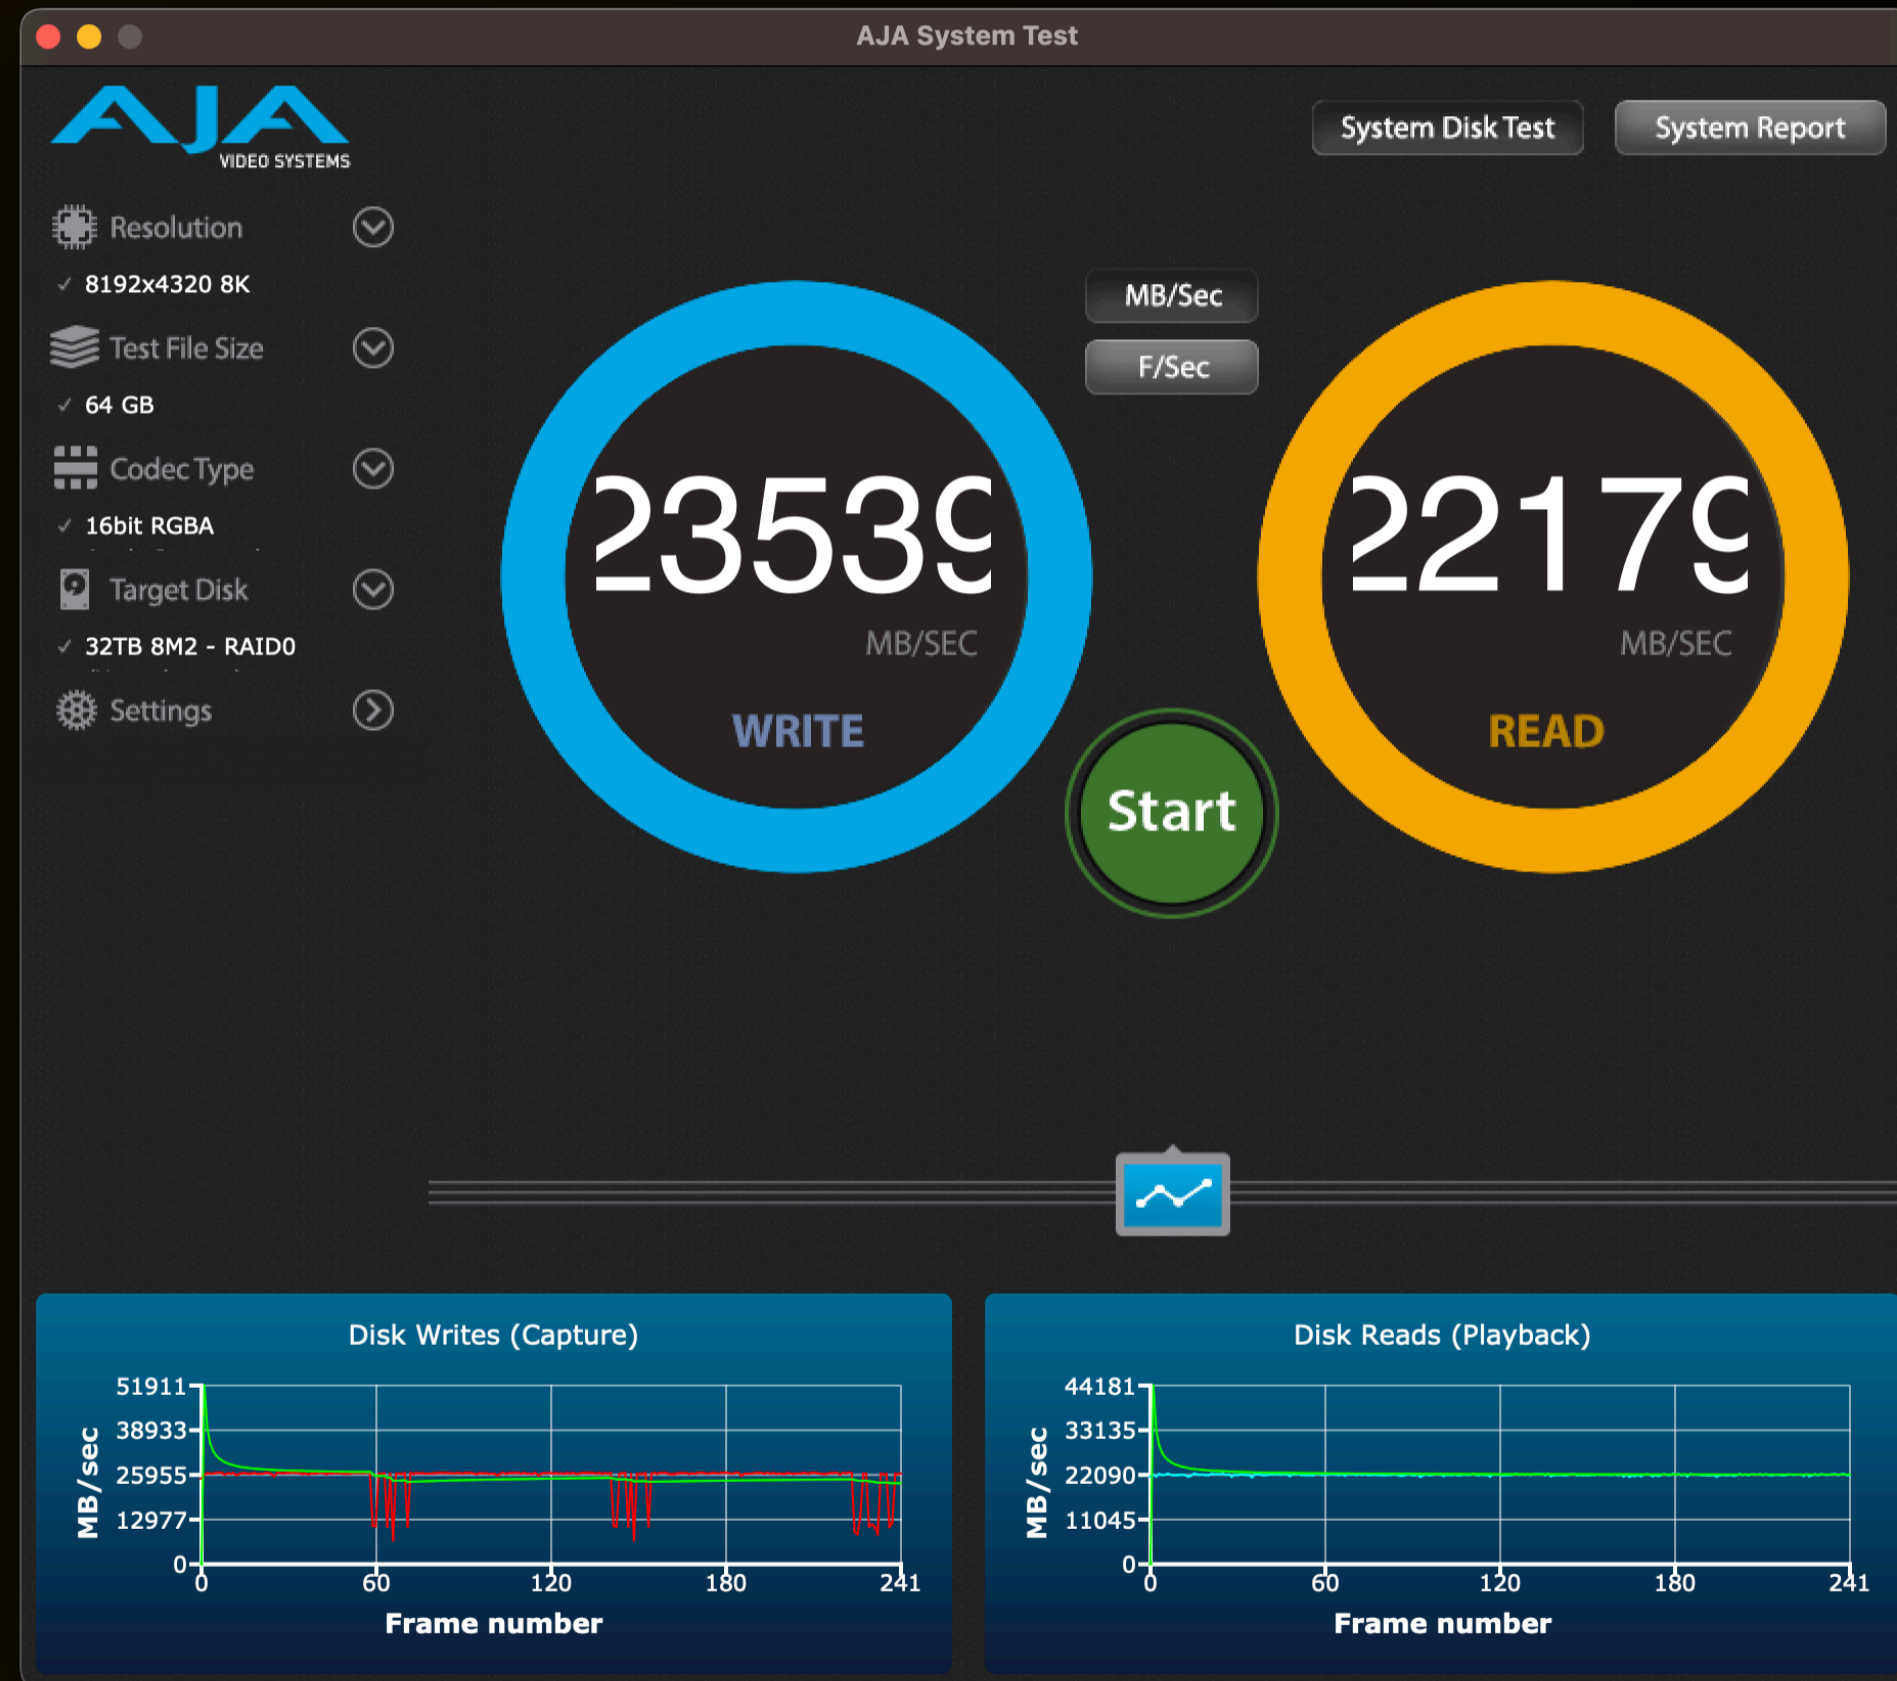
2023



Accelsior 8M2 + 2023 Mac Pro, RAID 0 = 24 GB/sec



Accelsior 8M2



RAID 0



RAID 5



PCI Gen 4



PCI Gen 4

- Double speed of PCI Gen 3
- 2 GB/sec/lane
- More power & more heat



2023 Mac Pro PCIe Lanes

2 Pools:
-16 lanes
-8 Lanes

Expansion Slot Utility

Automatic bandwidth configuration allows your Mac to dynamically assign bandwidth between pools for best performance.

☒ Automatic Bandwidth Configuration

		A	B
	7	x4 <input type="radio"/>	<input checked="" type="radio"/>
	6	x8 <input type="radio"/>	<input type="radio"/>
	5	x8 <input type="radio"/>	<input type="radio"/>
	4	x8 <input type="radio"/>	<input type="radio"/>
	3	x8 <input type="radio"/>	<input type="radio"/>
	2	x16 <input type="radio"/>	<input type="radio"/>
	1	x16 <input checked="" type="radio"/>	<input type="radio"/>

Pool A Allocation

100%

Pool B Allocation

88%



Thunderbolt 5



Thunderbolt

Thunderbolt Version	1	2	3	4	5
Gb/s	10 Gb/s	20 Gb/s	40 Gb/s	40 Gb/s	80 Gb/s
Data Rate GB/s	0.7 GB/s	1.4GB/s	2.8 GB/s	2.8 GB/s	5.6 GB/s
PCIe	Gen 2 2 lanes	Gen 2 4 lanes	Gen 3 4 lanes	Gen 3 4 lanes	Gen 4 4 lanes



Thunderbolt 5 Signalling

PAM-3

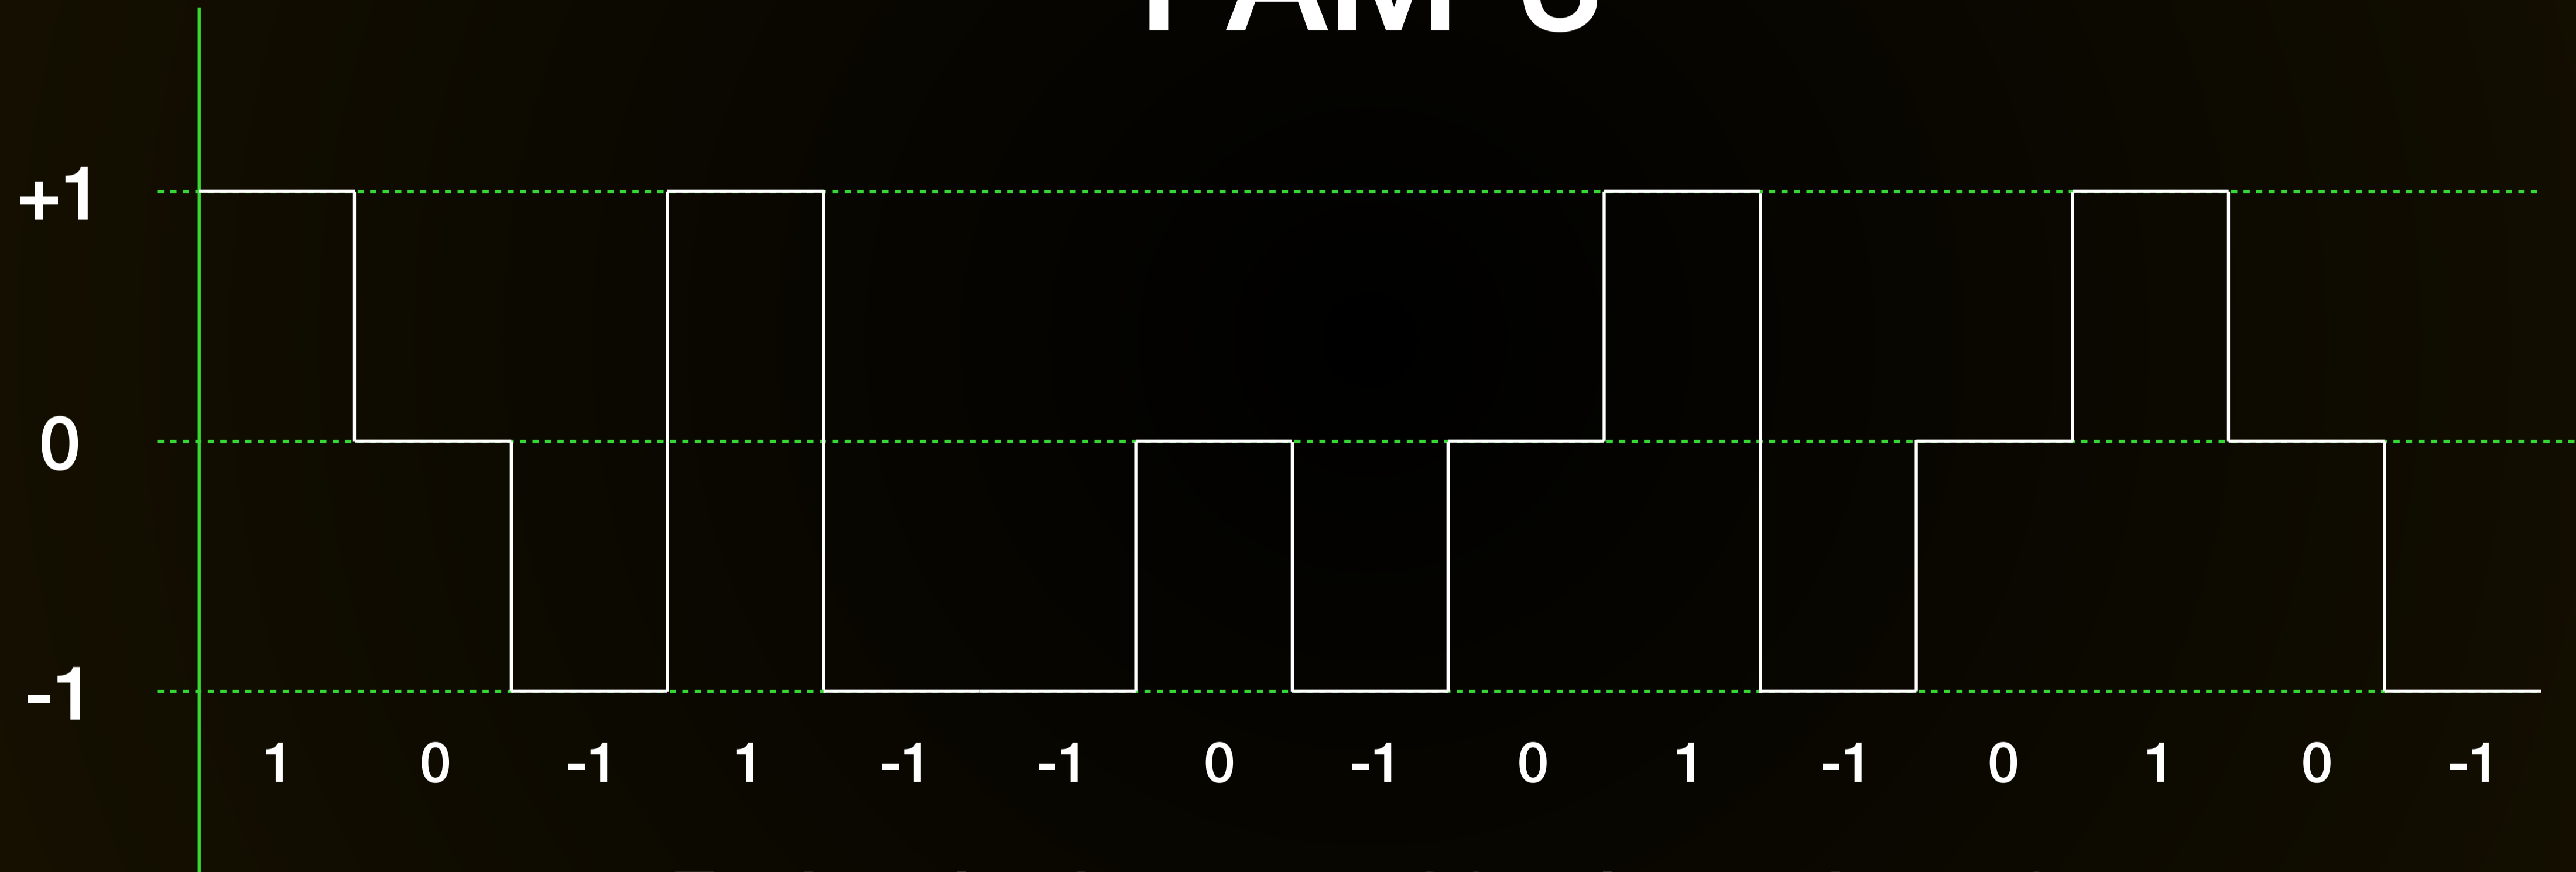


Pulse Amplitude Modulation
3 Levels



Thunderbolt 5 Signalling

PAM-3

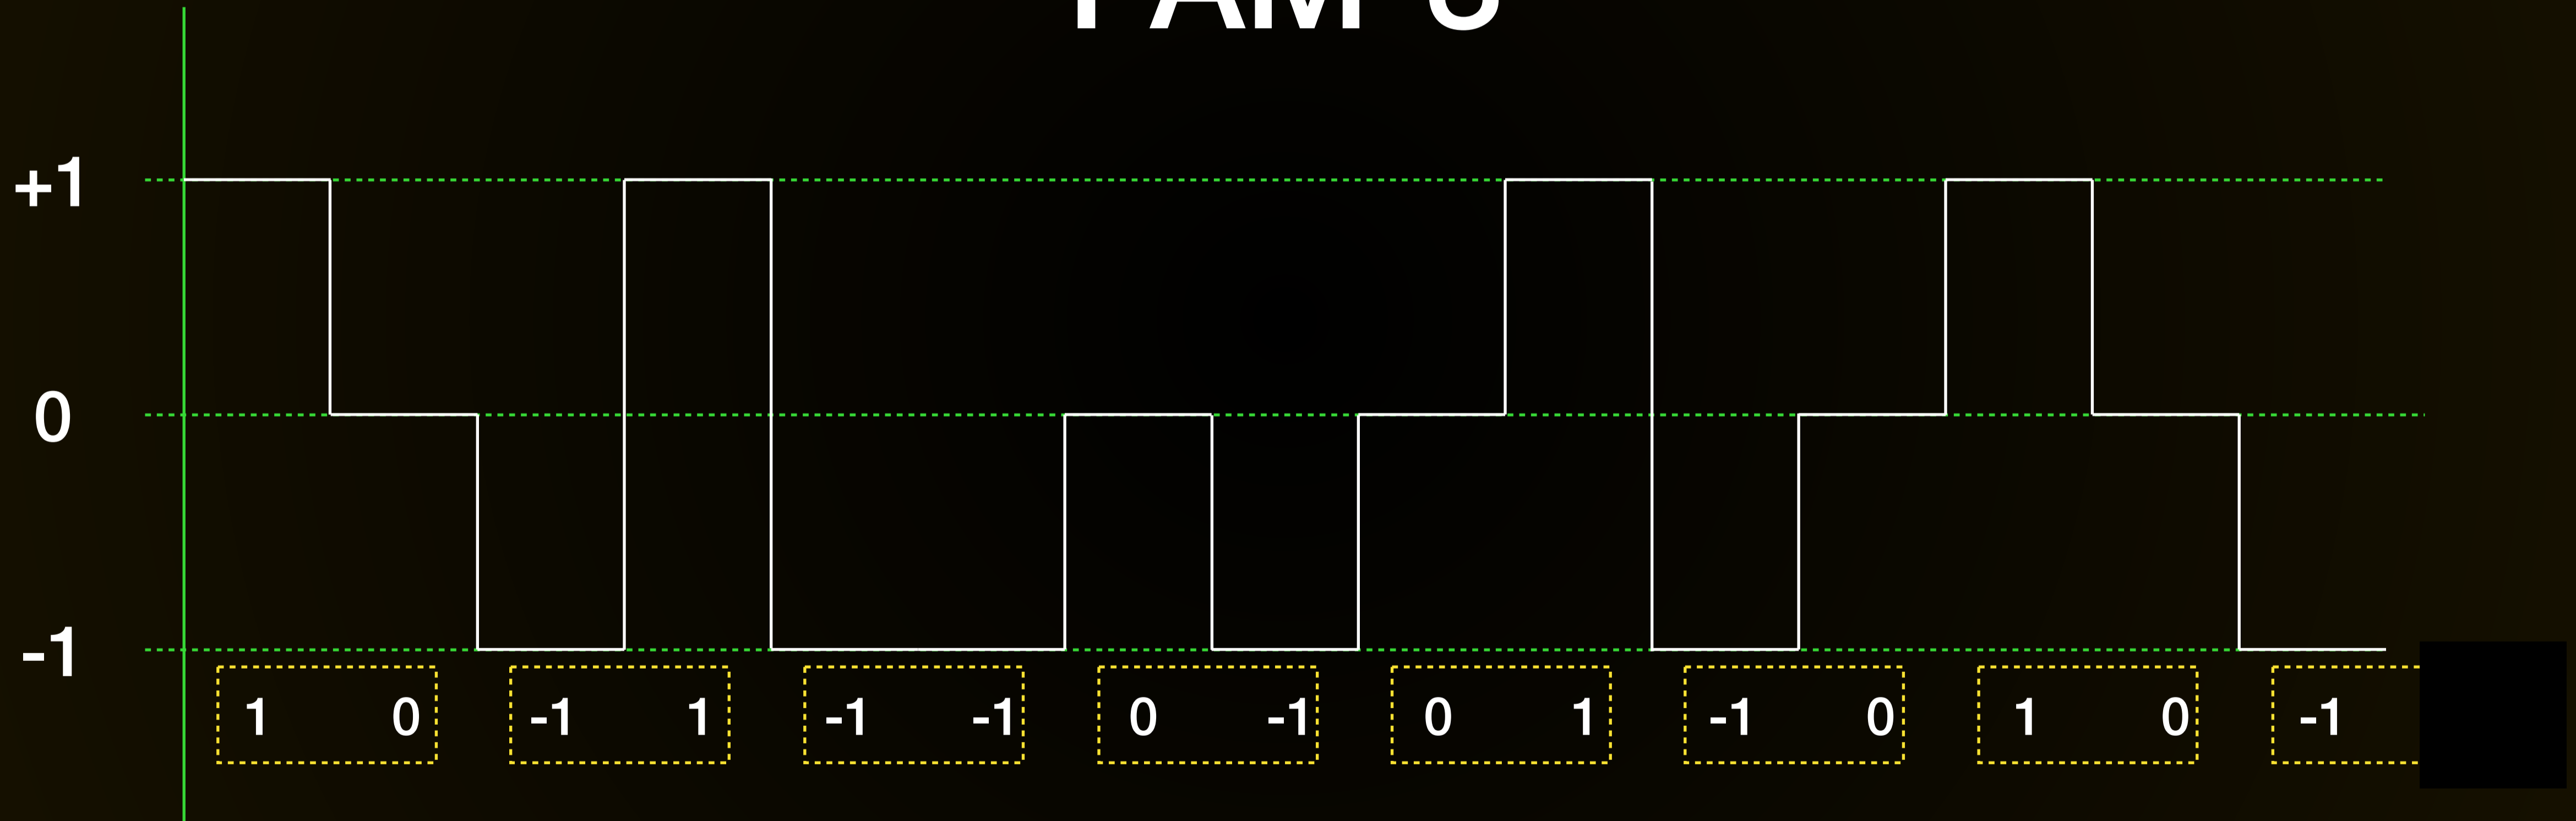


Each pulse has one of 3 values = base 3
2 pulses = 3^2 possible values



Thunderbolt 5 Signalling

PAM-3



2 pulses used to encode 3 binary bits



Thunderbolt Cables

USB



Don't Use

Thunderbolt
Passive



Works for
Thunderbolt 3 - 5

Thunderbolt 3
Active



Don't Use

Thunderbolt 4
Active



Use for
Thunderbolt 3 & 4

Thunderbolt 5
Active



Use for
Thunderbolt 5



Thunderbolt Cables

USB



- Might work for Thunderbolt
- Slower
- Less Reliable
- Not Certified

Don't Use



Thunderbolt Cables

Thunderbolt
Passive



Works for
Thunderbolt 3 - 5

- Works for all Thunderbolt levels
- Avoid 2 m cables (= Thunderbolt 2 speed)



Thunderbolt Cables

Thunderbolt 3
Active



USB devices run at USB 2.0 speed (40 MB/sec)

Don't Use



Reliability

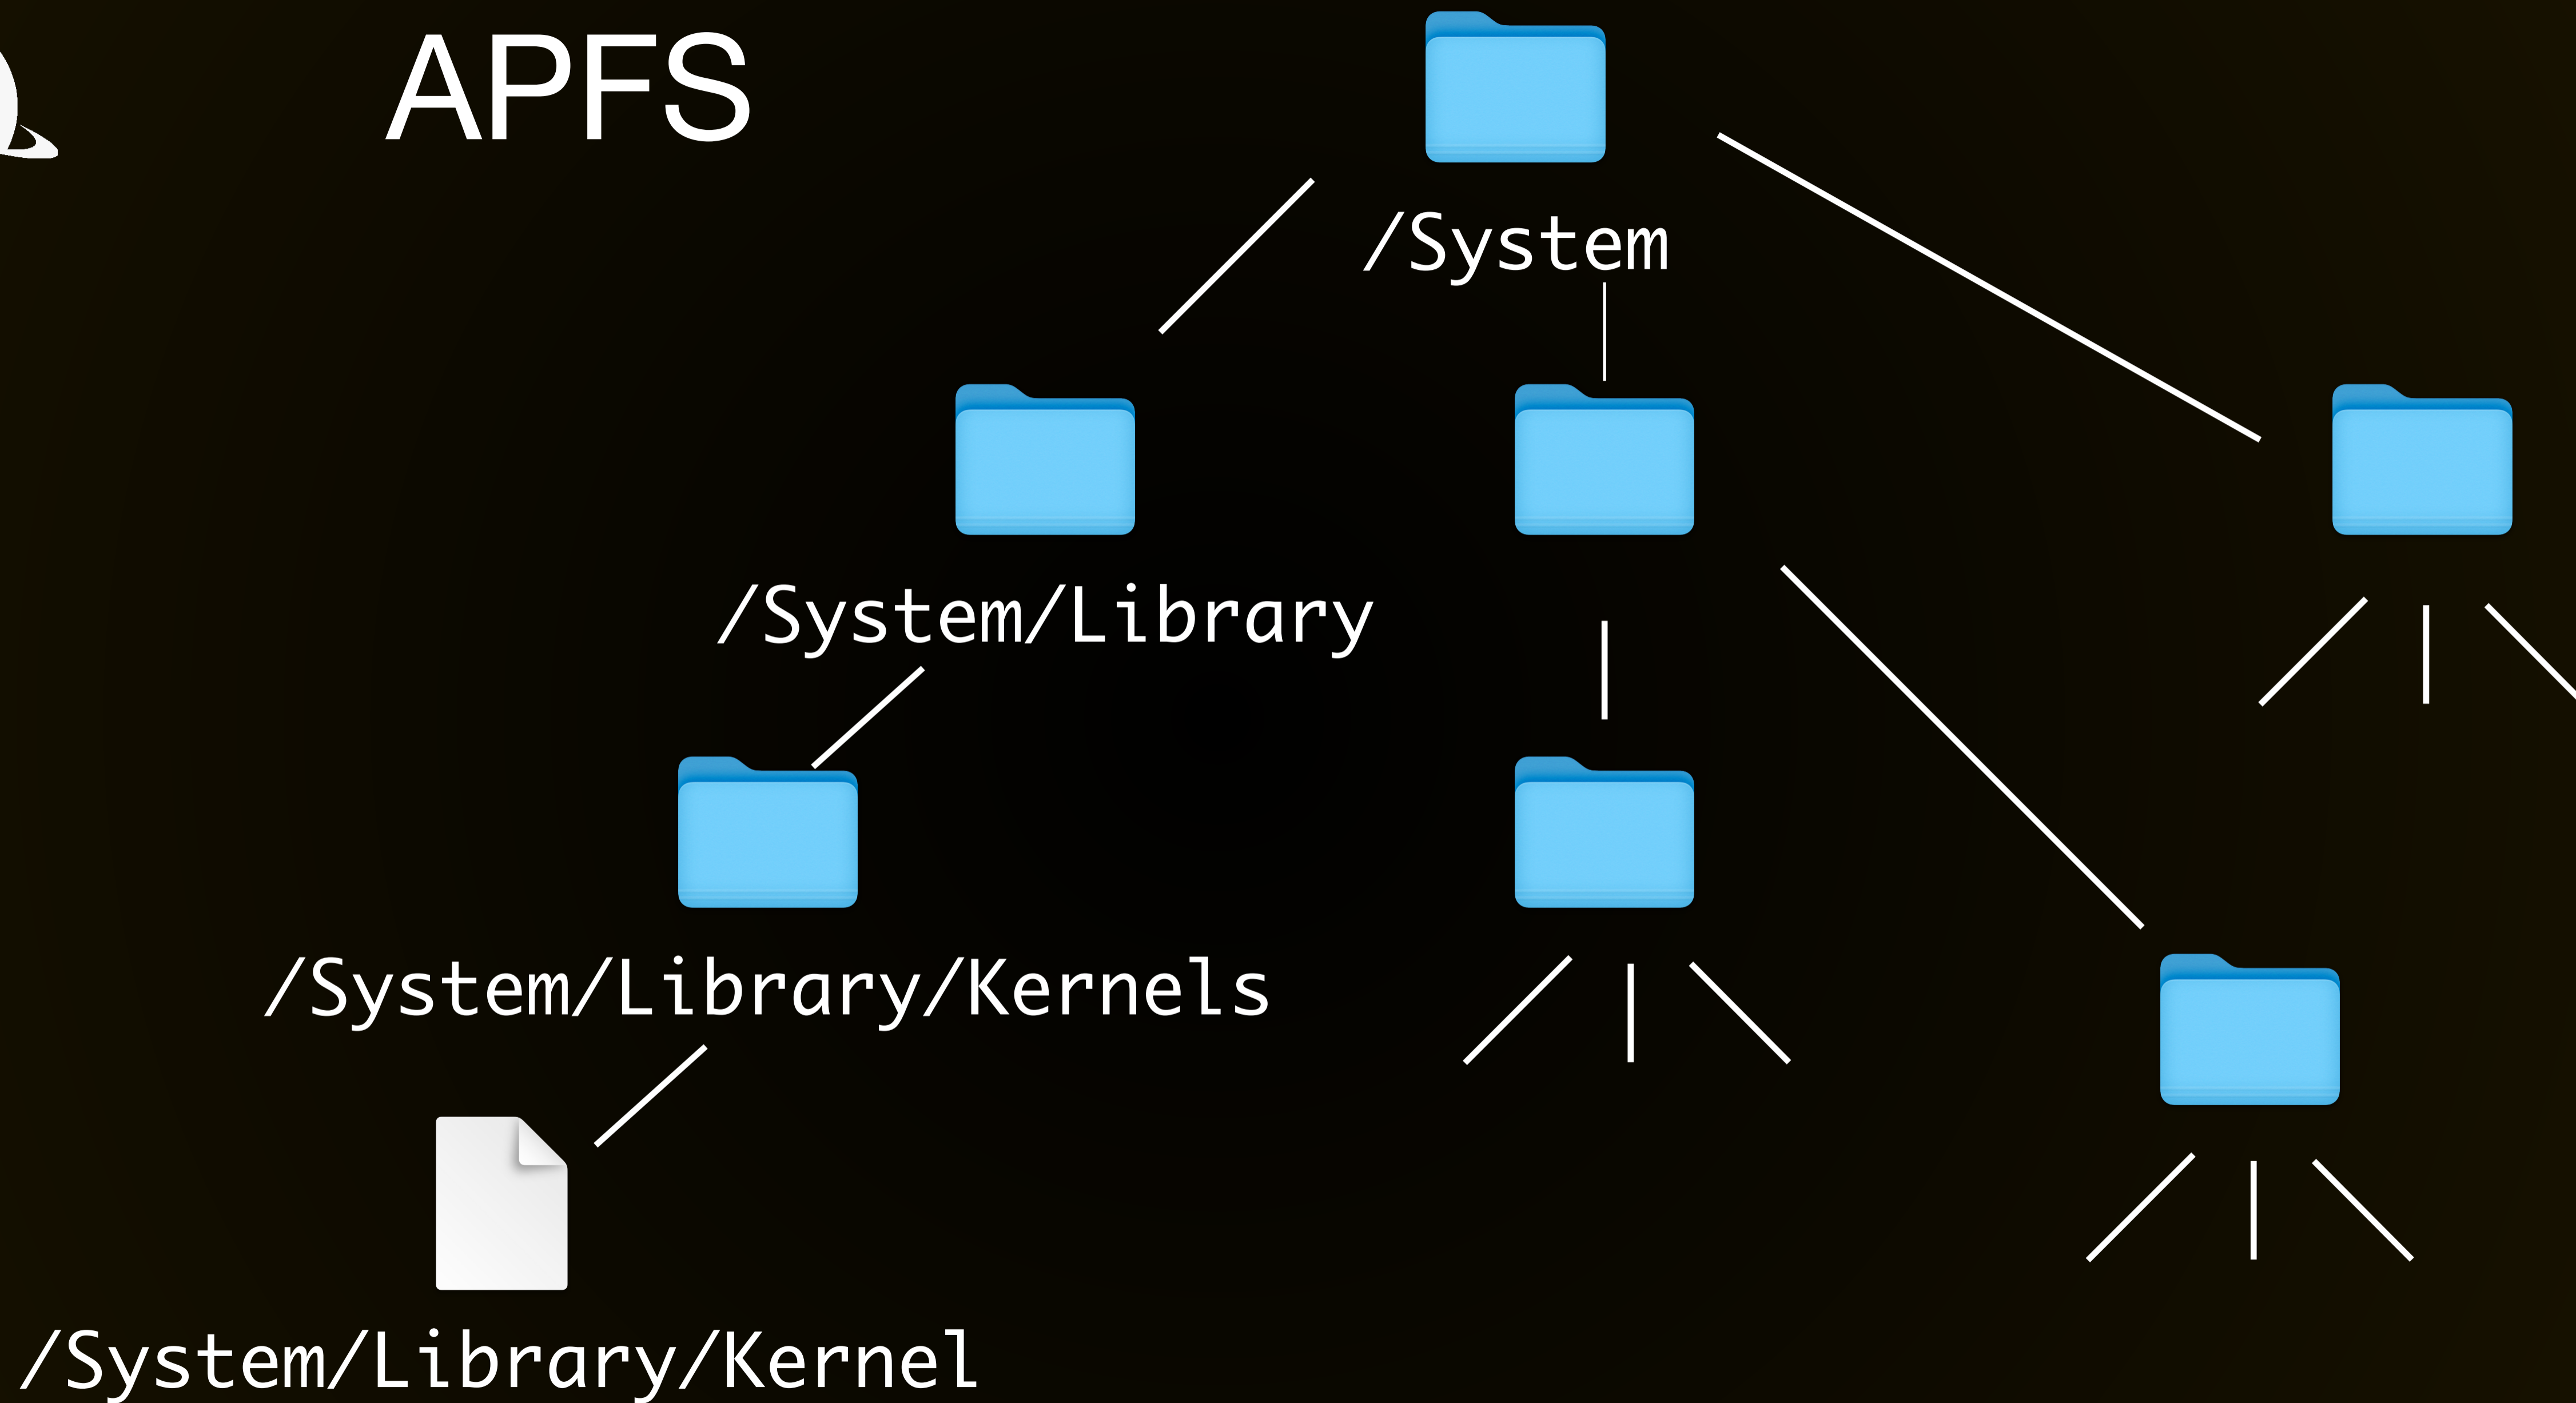
- **Protecting the System Volume**
- **Are kexts going to disappear?**
- **What is Bit-Rot**
- **History of Volume Loss**
- **An Example of Volume Loss**




Protecting the System Volume




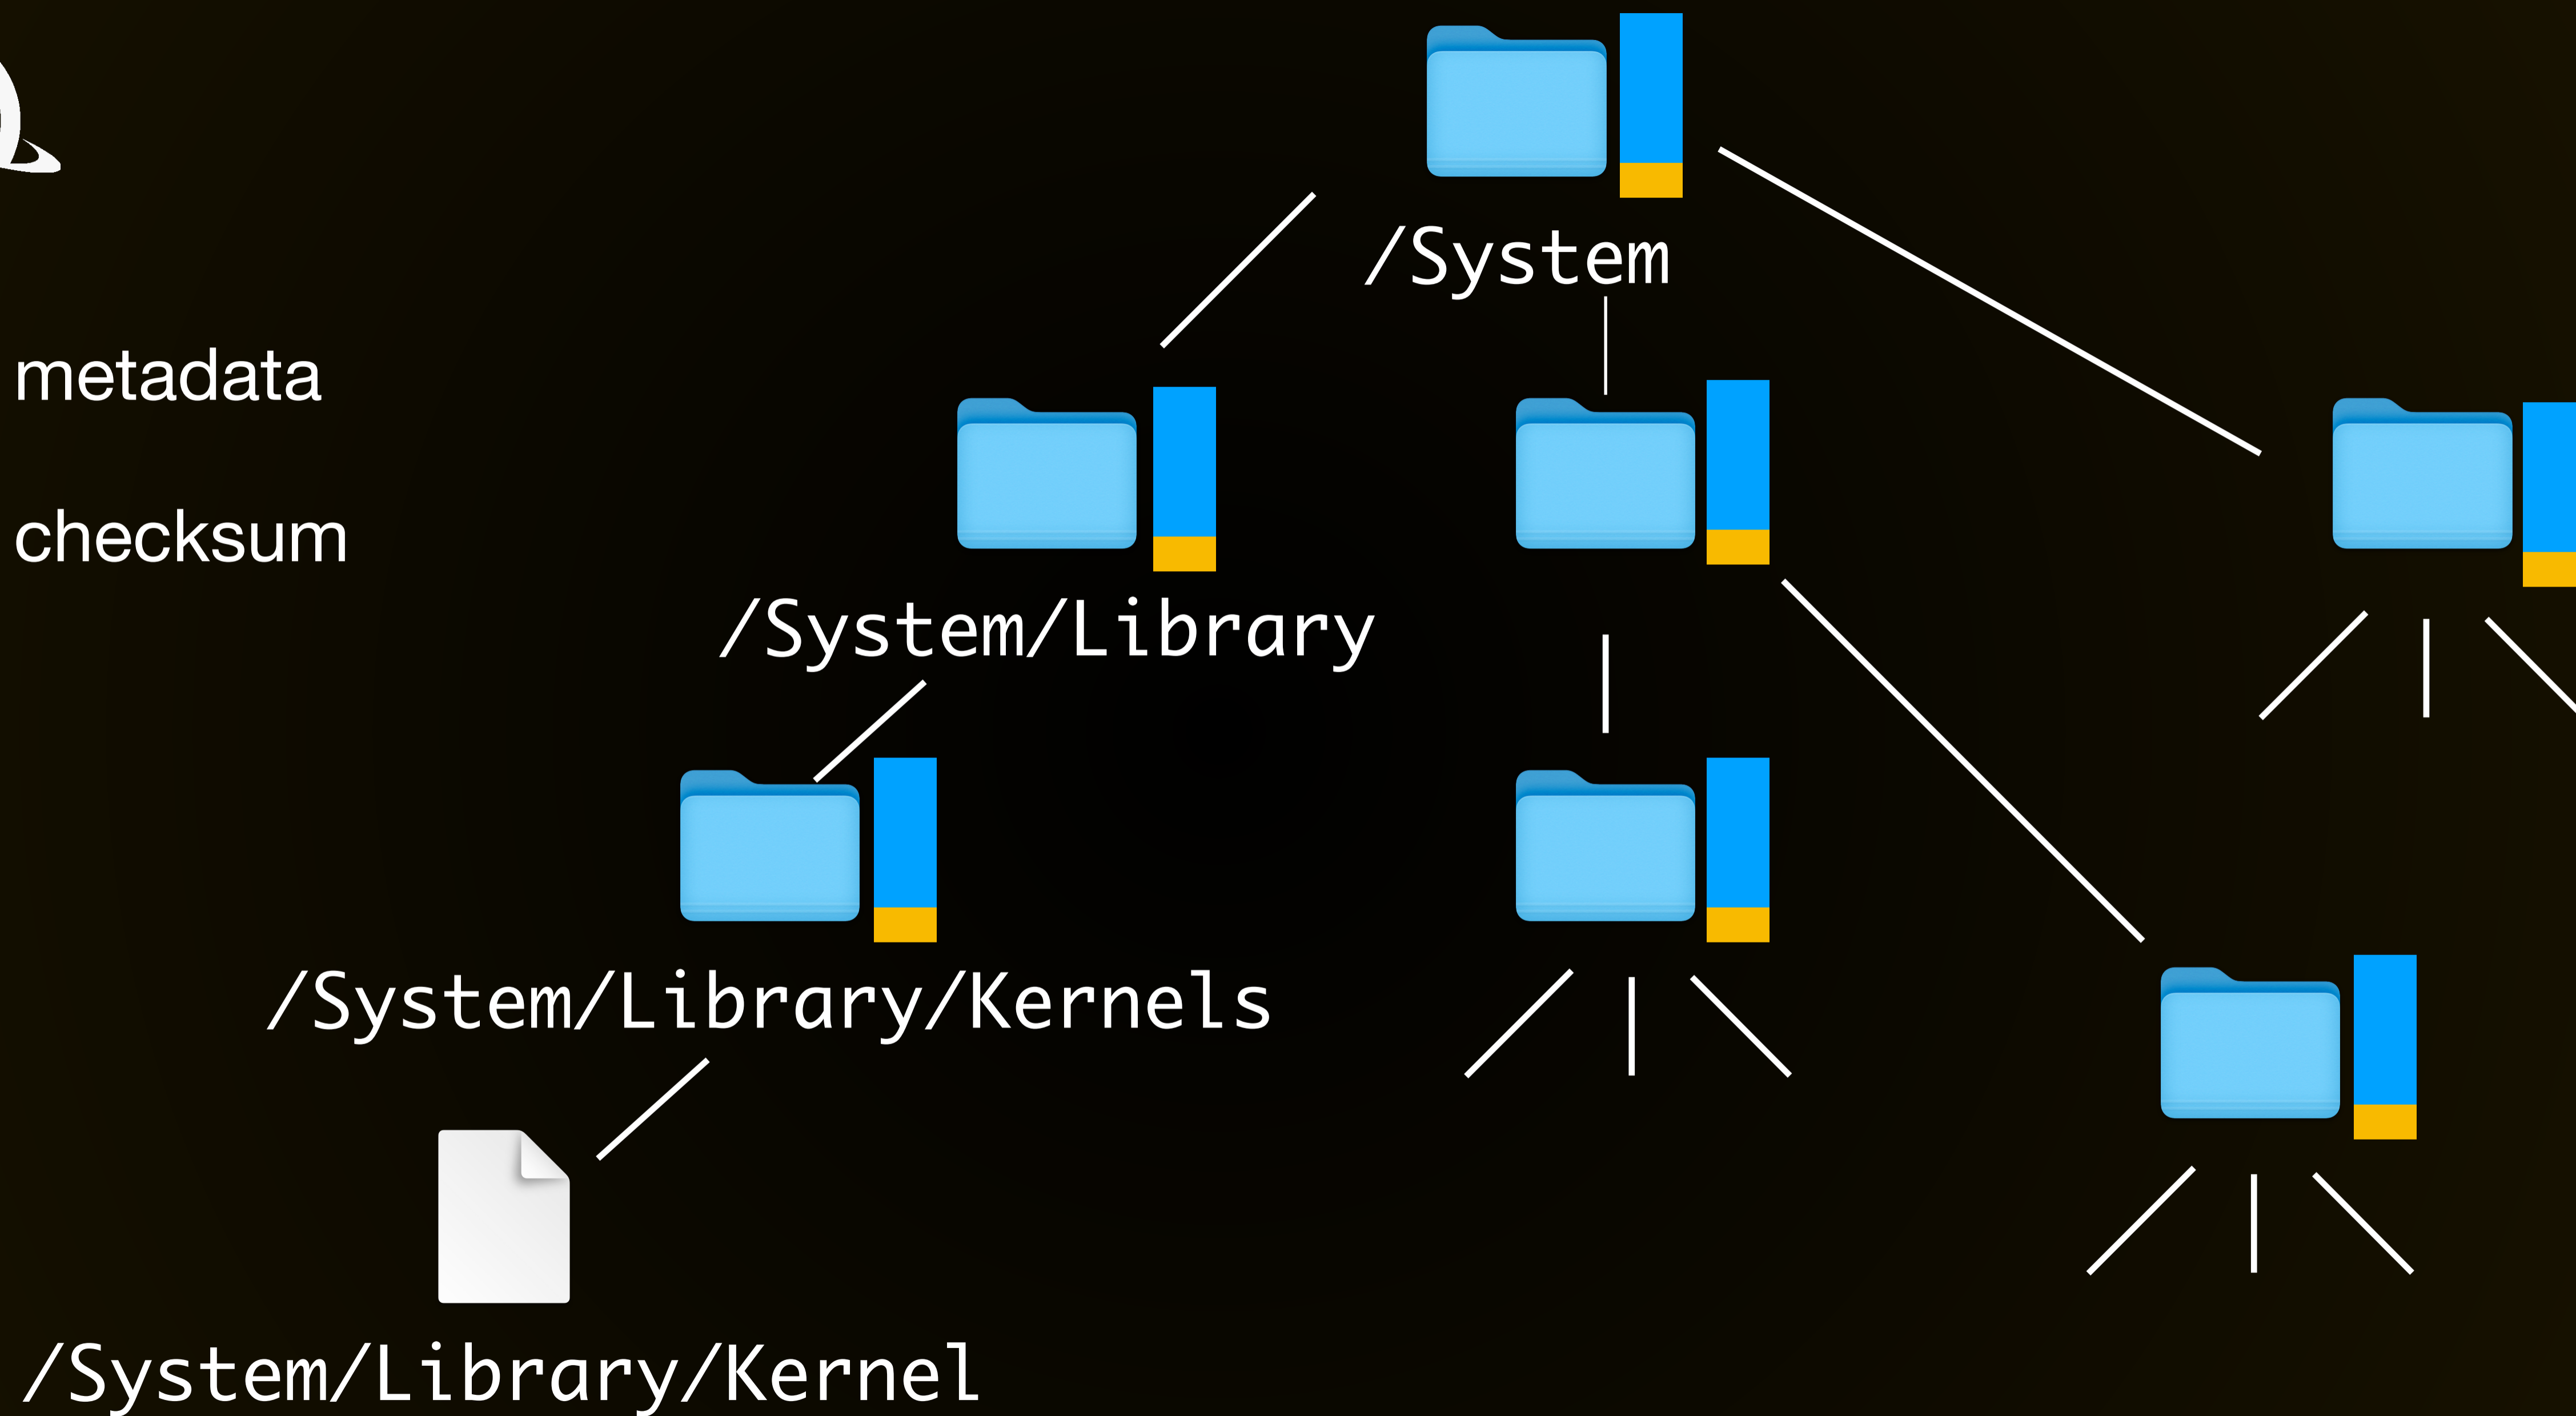
APFS





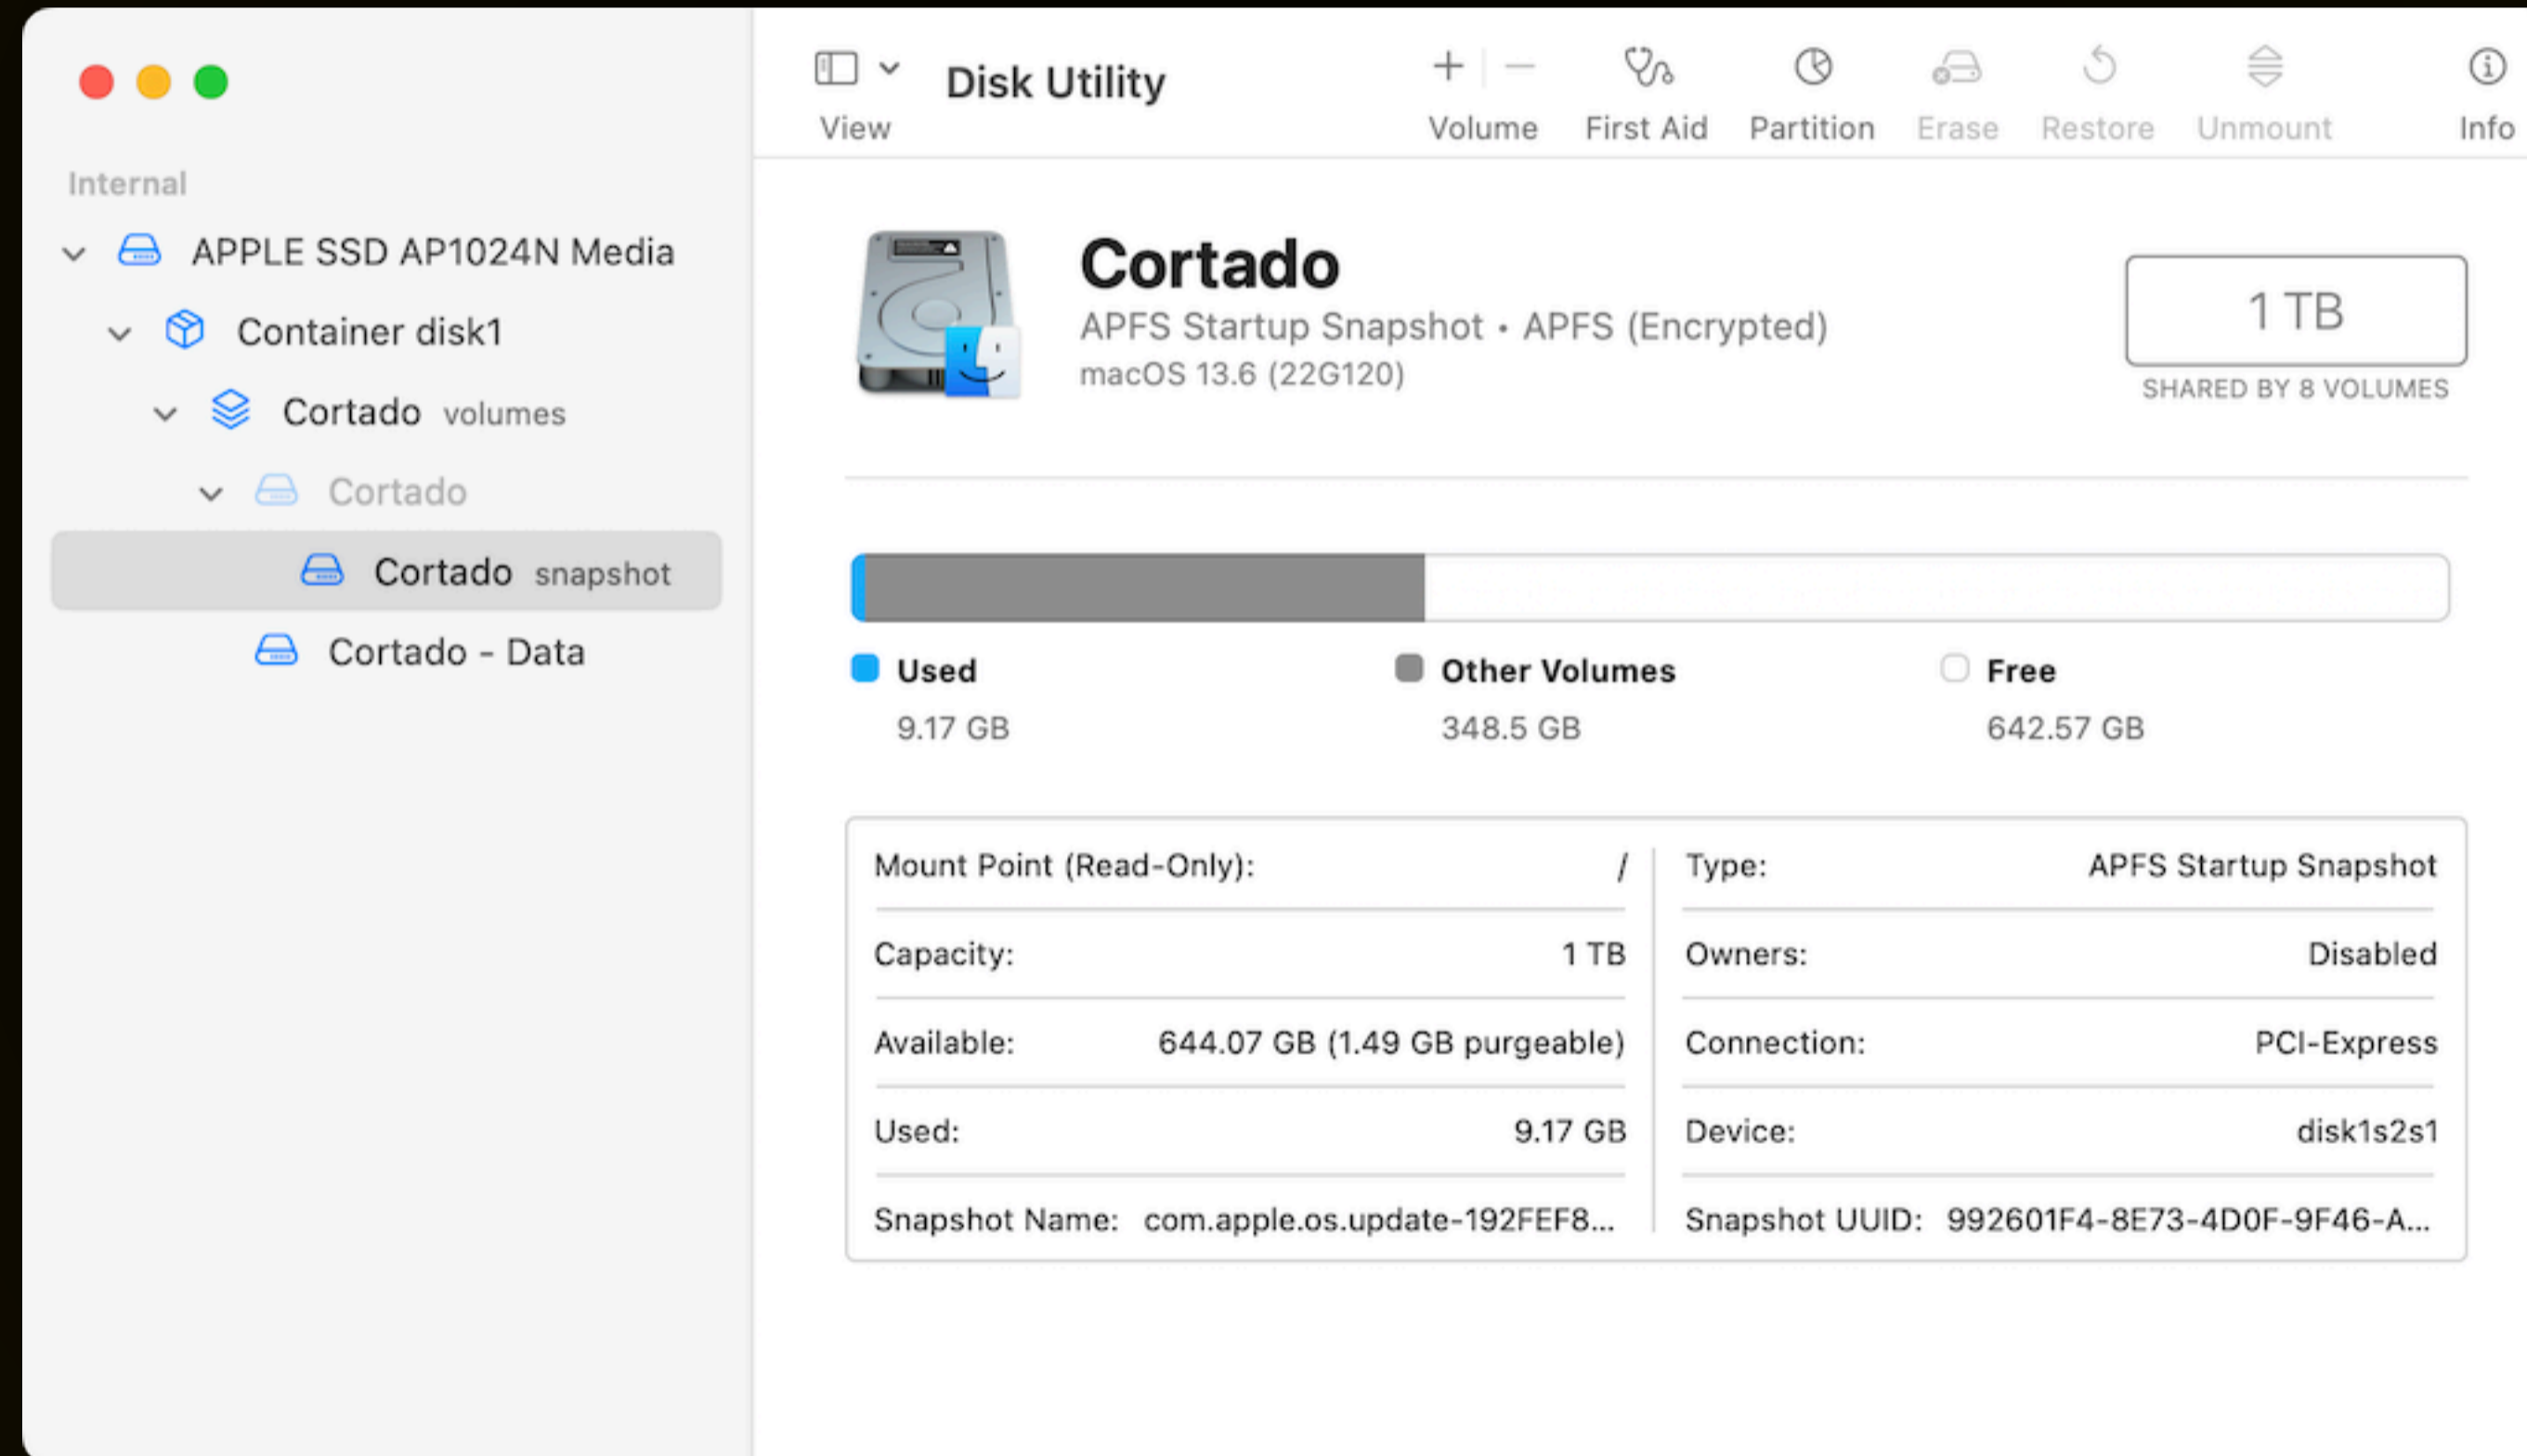
 = metadata

 = checksum





Signed System Volumes





System Volume



```
standing — -zsh — 77x38

standing@Yu ~ % diskutil apfs list disk1
|
+-- Container disk1 6EF02048-2C05-4EA1-B78F-CEED3C791B7D
=====
APFS Container Reference:      disk1
Size (Capacity Ceiling):      1000240963584 B (1.0 TB)
Capacity In Use By Volumes:    357651013632 B (357.7 GB) (35.8% used)
Capacity Not Allocated:        642589949952 B (642.6 GB) (64.2% free)
|
+--< Physical Store disk0s2 E8F4A4ED-AEE8-444E-AE04-276585F07163
-----
APFS Physical Store Disk:     disk0s2
Size:                          1000240963584 B (1.0 TB)
|
+--> Volume disk1s1 90953322-8CE8-4CC3-8252-31D47B320E31
-----
APFS Volume Disk (Role):      disk1s1 (Data)
Name:                          Cortado - Data (Case-insensitive)
Mount Point:                   /System/Volumes/Data
Capacity Consumed:             169620041728 B (169.6 GB)
Sealed:                         No
FileVault:                      Yes (Unlocked)
|
+--> Volume disk1s2 1AEB3E44-DE4C-49CB-BB61-DD6822E31F56
-----
APFS Volume Disk (Role):      disk1s2 (System)
Name:                          Cortado (Case-insensitive)
Mount Point:                   Not Mounted
Capacity Consumed:             9167335424 B (9.2 GB)
Sealed:                         Yes
FileVault:                      Yes (Unlocked)
Encrypted:                      No
|
Snapshot:                      992601F4-8E73-4D0F-9F46-A3073FE99140
Snapshot Disk:                 disk1s2s1
Snapshot Mount Point:          /
Snapshot Sealed:               Yes
```



Snapshot of System Volume →

```
standing — -zsh — 77x38

standing@Yu ~ % diskutil apfs list disk1
|
+-- Container disk1 6EF02048-2C05-4EA1-B78F-CEED3C791B7D
=====
APFS Container Reference:    disk1
Size (Capacity Ceiling):    1000240963584 B (1.0 TB)
Capacity In Use By Volumes: 357651013632 B (357.7 GB) (35.8% used)
Capacity Not Allocated:    642589949952 B (642.6 GB) (64.2% free)
|
+--< Physical Store disk0s2 E8F4A4ED-AEE8-444E-AE04-276585F07163
-----
APFS Physical Store Disk:    disk0s2
Size:                        1000240963584 B (1.0 TB)
|
+--> Volume disk1s1 90953322-8CE8-4CC3-8252-31D47B320E31
-----
APFS Volume Disk (Role):    disk1s1 (Data)
Name:                       Cortado - Data (Case-insensitive)
Mount Point:                /System/Volumes/Data
Capacity Consumed:          169620041728 B (169.6 GB)
Sealed:                     No
FileVault:                  Yes (Unlocked)
|
+--> Volume disk1s2 1AEB3E44-DE4C-49CB-BB61-DD6822E31F56
-----
APFS Volume Disk (Role):    disk1s2 (System)
Name:                       Cortado (Case-insensitive)
Mount Point:                Not Mounted
Capacity Consumed:          9167335424 B (9.2 GB)
Sealed:                     Yes
FileVault:                  Yes (Unlocked)
Encrypted:                  No
|
Snapshot:                   992601F4-8E73-4D0F-9F46-A3073FE99140
Snapshot Disk:              disk1s2s1
Snapshot Mount Point:       /
Snapshot Sealed:            Yes
```

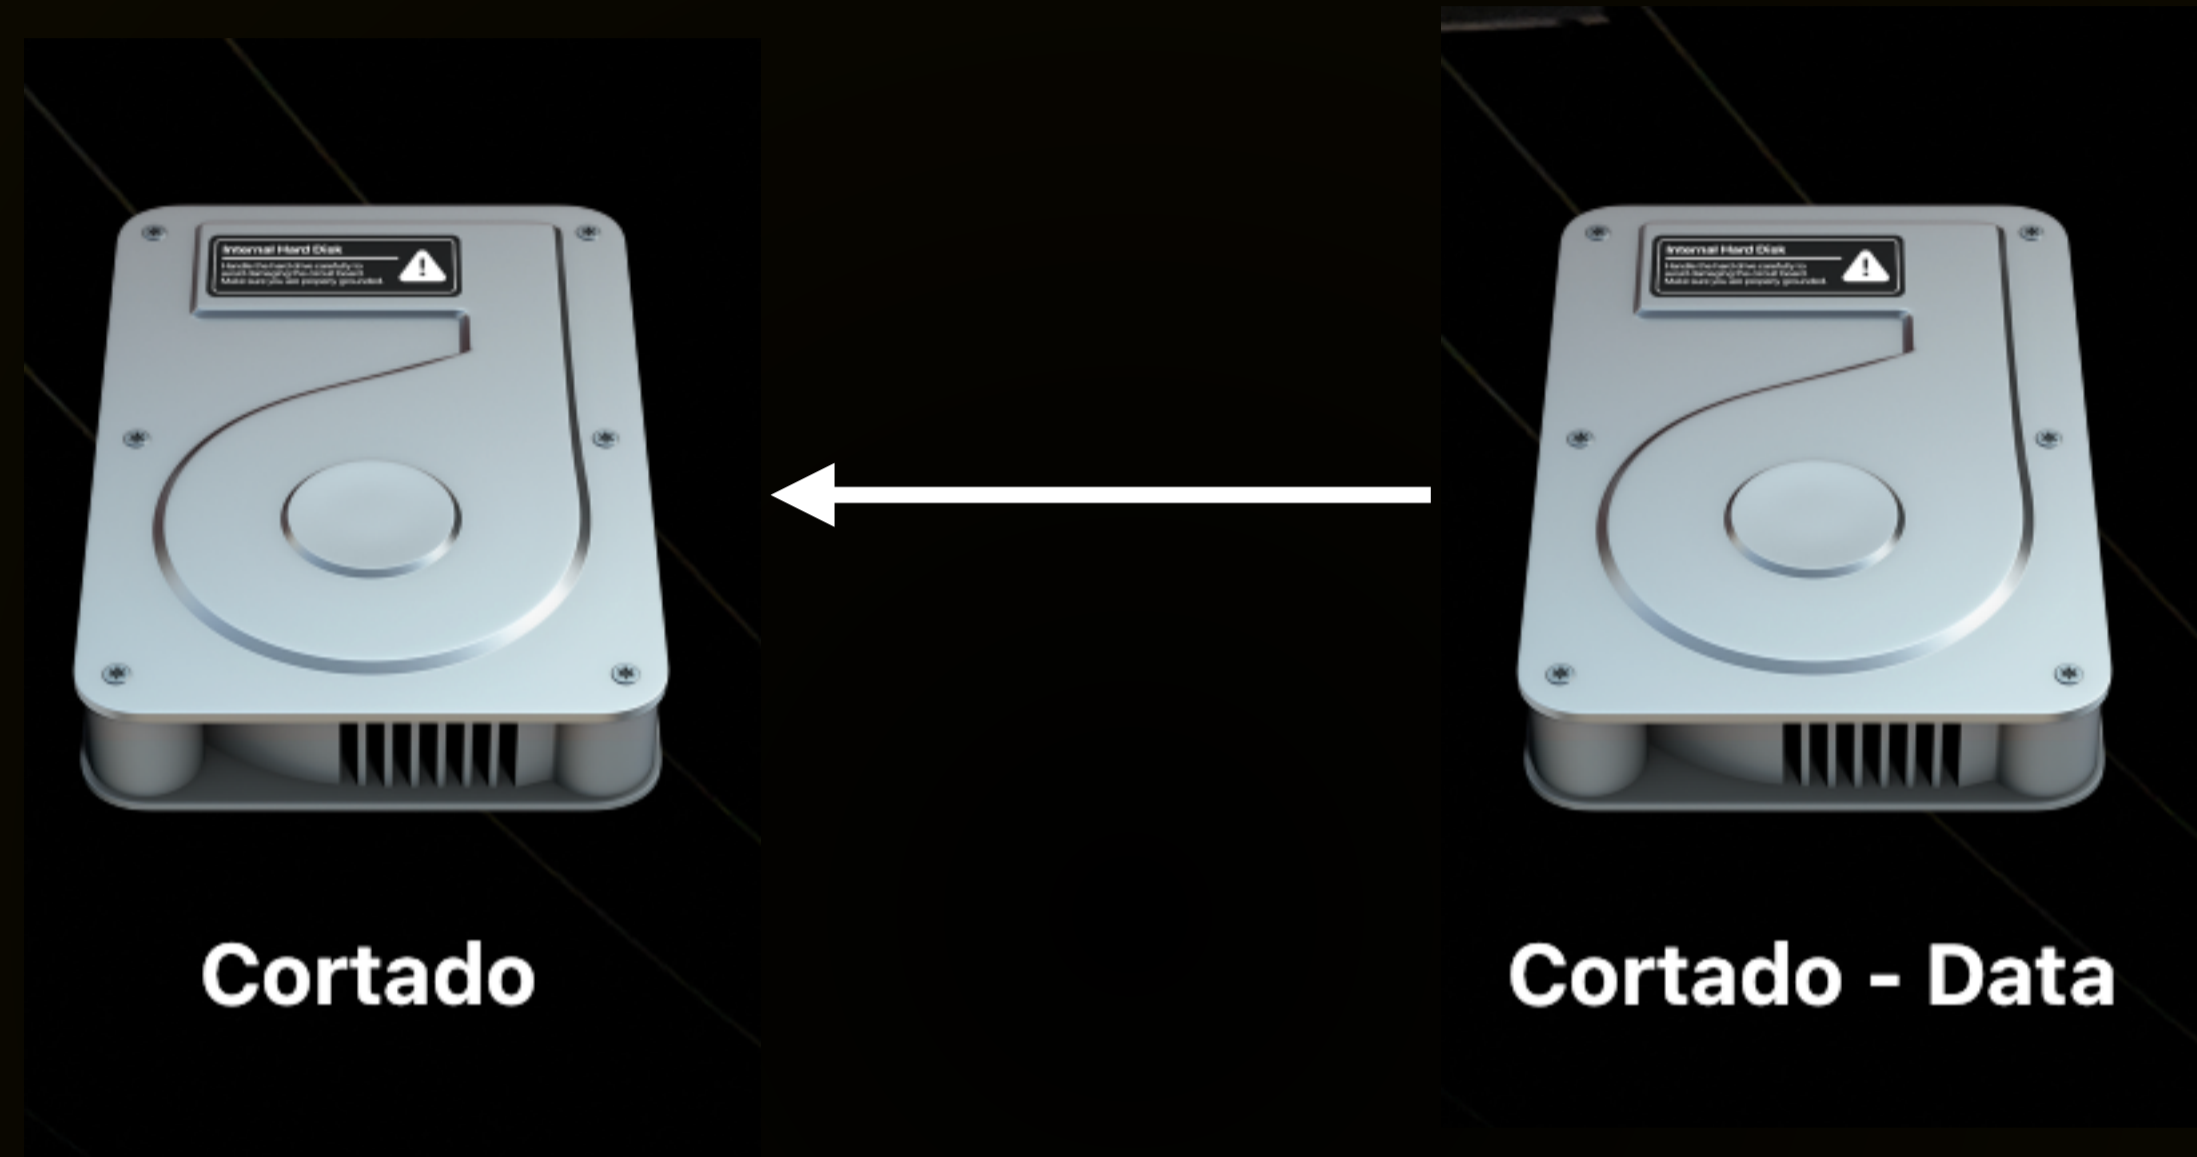


Data Volume



```
standing — -zsh — 77x38

standing@Yu ~ % diskutil apfs list disk1
|
+-- Container disk1 6EF02048-2C05-4EA1-B78F-CEED3C791B7D
=====
APFS Container Reference:    disk1
Size (Capacity Ceiling):    1000240963584 B (1.0 TB)
Capacity In Use By Volumes: 357651013632 B (357.7 GB) (35.8% used)
Capacity Not Allocated:    642589949952 B (642.6 GB) (64.2% free)
|
+--< Physical Store disk0s2 E8F4A4ED-AEE8-444E-AE04-276585F07163
-----
APFS Physical Store Disk:   disk0s2
Size:                       1000240963584 B (1.0 TB)
|
+--> Volume disk1s1 90953322-8CE8-4CC3-8252-31D47B320E31
-----
APFS Volume Disk (Role):    disk1s1 (Data)
Name:                       Cortado - Data (Case-insensitive)
Mount Point:                /System/Volumes/Data
Capacity Consumed:          169620041728 B (169.6 GB)
Sealed:                     No
FileVault:                  Yes (Unlocked)
|
+--> Volume disk1s2 1AEB3E44-DE4C-49CB-BB61-DD6822E31F56
-----
APFS Volume Disk (Role):    disk1s2 (System)
Name:                       Cortado (Case-insensitive)
Mount Point:                Not Mounted
Capacity Consumed:          9167335424 B (9.2 GB)
Sealed:                     Yes
FileVault:                  Yes (Unlocked)
Encrypted:                  No
|
Snapshot:                   992601F4-8E73-4D0F-9F46-A3073FE99140
Snapshot Disk:              disk1s2s1
Snapshot Mount Point:       /
Snapshot Sealed:            Yes
```



Data volume mounted at:
`/System/Volumes/Data`



Firm Links

```
standing — -zsh — 83x20
standing@Yu ~ % cat /usr/share/firmlinks
/Applications  Applications
/Library       Library
/System/Library/Caches  System/Library/Caches
/System/Library/Assets  System/Library/Assets
/System/Library/PreinstalledAssets  System/Library/PreinstalledAssets
/System/Library/AssetsV2  System/Library/AssetsV2
/System/Library/PreinstalledAssetsV2  System/Library/PreinstalledAssetsV2
/System/Library/CoreServices/CoreTypes.bundle/Contents/Library  System/Library/Core
Services/CoreTypes.bundle/Contents/Library
/System/Library/Speech  System/Library/Speech
/Users  Users
/Volumes  Volumes
/cores  cores
/opt  opt
/private  private
/usr/local  usr/local
/usr/libexec/cups  usr/libexec/cups
/usr/share/snmp  usr/share/snmp
```

/usr/share/firmlinks



Where are your applications?

```
standing — -zsh — 78x24
[standing@Yu ~ % ls -l /System/Applications
total 0
drwxr-xr-x  3 root  wheel  96 Sep 16 09:48 App Store.app
drwxr-xr-x  3 root  wheel  96 Sep 16 09:48 Automator.app
drwxr-xr-x  3 root  wheel  96 Sep 16 09:48 Books.app
drwxr-xr-x  3 root  wheel  96 Sep 16 09:48 Calculator.app
drwxr-xr-x  3 root  wheel  96 Sep 16 09:48 Calendar.app
drwxr-xr-x  3 root  wheel  96 Sep 16 09:48 Chess.app
drwxr-xr-x  3 root  wheel  96 Sep 16 09:48 Clock.app
drwxr-xr-x  3 root  wheel  96 Sep 16 09:48 Contacts.app
drwxr-xr-x  3 root  wheel  96 Sep 16 09:48 Dictionary.app
drwxr-xr-x  3 root  wheel  96 Sep 16 09:48 FaceTime.app
drwxr-xr-x  3 root  wheel  96 Sep 16 09:48 FindMy.app
drwxr-xr-x  3 root  wheel  96 Sep 16 09:48 Font Book.app
drwxr-xr-x  3 root  wheel  96 Sep 16 09:48 Freeform.app
drwxr-xr-x  3 root  wheel  96 Sep 16 09:48 Home.app
drwxr-xr-x  3 root  wheel  96 Sep 16 09:48 Image Capture.app
drwxr-xr-x  3 root  wheel  96 Sep 16 09:48 Launchpad.app
drwxr-xr-x  3 root  wheel  96 Sep 16 09:48 Mail.app
drwxr-xr-x  3 root  wheel  96 Sep 16 09:48 Maps.app
drwxr-xr-x  3 root  wheel  96 Sep 16 09:48 Messages.app
drwxr-xr-x  3 root  wheel  96 Sep 16 09:48 Mission Control.app
drwxr-xr-x  3 root  wheel  96 Sep 16 09:48 Music.app
drwxr-xr-x  3 root  wheel  96 Sep 16 09:48 News.app
```

/System/Applications

```
standing — -zsh — 78x24
[standing@Yu ~ % ls -l /System/Volumes/Data/Applications
total 0
drwxr-xr-x@ 3 standing  staff  96 Sep 13 20:17 1Password.app
drwx-----@ 3 standing  staff  96 Sep 27  2022 AJA System Test 16.2.25.app
drwxr-xr-x@ 3 root      wheel  96 Mar 29  2023 AJA System Test Lite.app
drwxr-xr-x@ 3 root      wheel  96 Jan 10  2023 Apple Configurator.app
drwxr-xr-x@ 3 standing  staff  96 May 24 23:08 Carbon Copy Cloner.app
drwxrwxr-x  3 standing  staff  96 Mar 25  2022 Cisco Webex Meetings.app
drwxr-xr-x@ 3 standing  staff  96 Jul 10 18:44 Copy That.app
drwxr-xr-x@ 3 standing  admin  96 Sep 30  2020 DataVault Password Manager.app
drwxr-xr-x@ 3 root      wheel  96 Jun  5 18:09 Developer.app
drwxr-xr-x@ 3 root      wheel  96 Jun  9 21:34 Focusrite Control.app
drwxr-xr-x  3 root      wheel  96 Sep 26 20:48 Install macOS Sonoma.app
drwxr-xr-x@ 3 standing  staff  96 Feb  8  2022 Kaleidoscope.app
drwxr-xr-x@ 3 root      wheel  96 Sep 26 20:24 Keynote.app
drwxr-xr-x@ 3 standing  admin  96 May  8 11:42 LastPass.app
drwxr-xr-x@ 3 root      wheel  96 Sep 13 19:53 Microsoft Teams.app
drwxr-xr-x@ 3 root      wheel  96 Sep 26 20:31 Numbers.app
drwxr-xr-x@ 3 root      wheel  96 Sep 26 20:30 OmniFocus.app
drwxr-xr-x@ 3 standing  admin  96 Sep 20 06:48 OmniOutliner.app
drwxr-xr-x@ 3 root      wheel  96 Mar 29  2023 PDFScanner.app
drwxr-xr-x@ 3 root      wheel  96 Sep 26 20:26 Pages.app
drwxr-xr-x@ 3 root      wheel  96 Sep 26 20:26 Pixen.app
drwxr-xr-x@ 3 standing  admin  96 Nov 28  2022 Qobuz.app
```

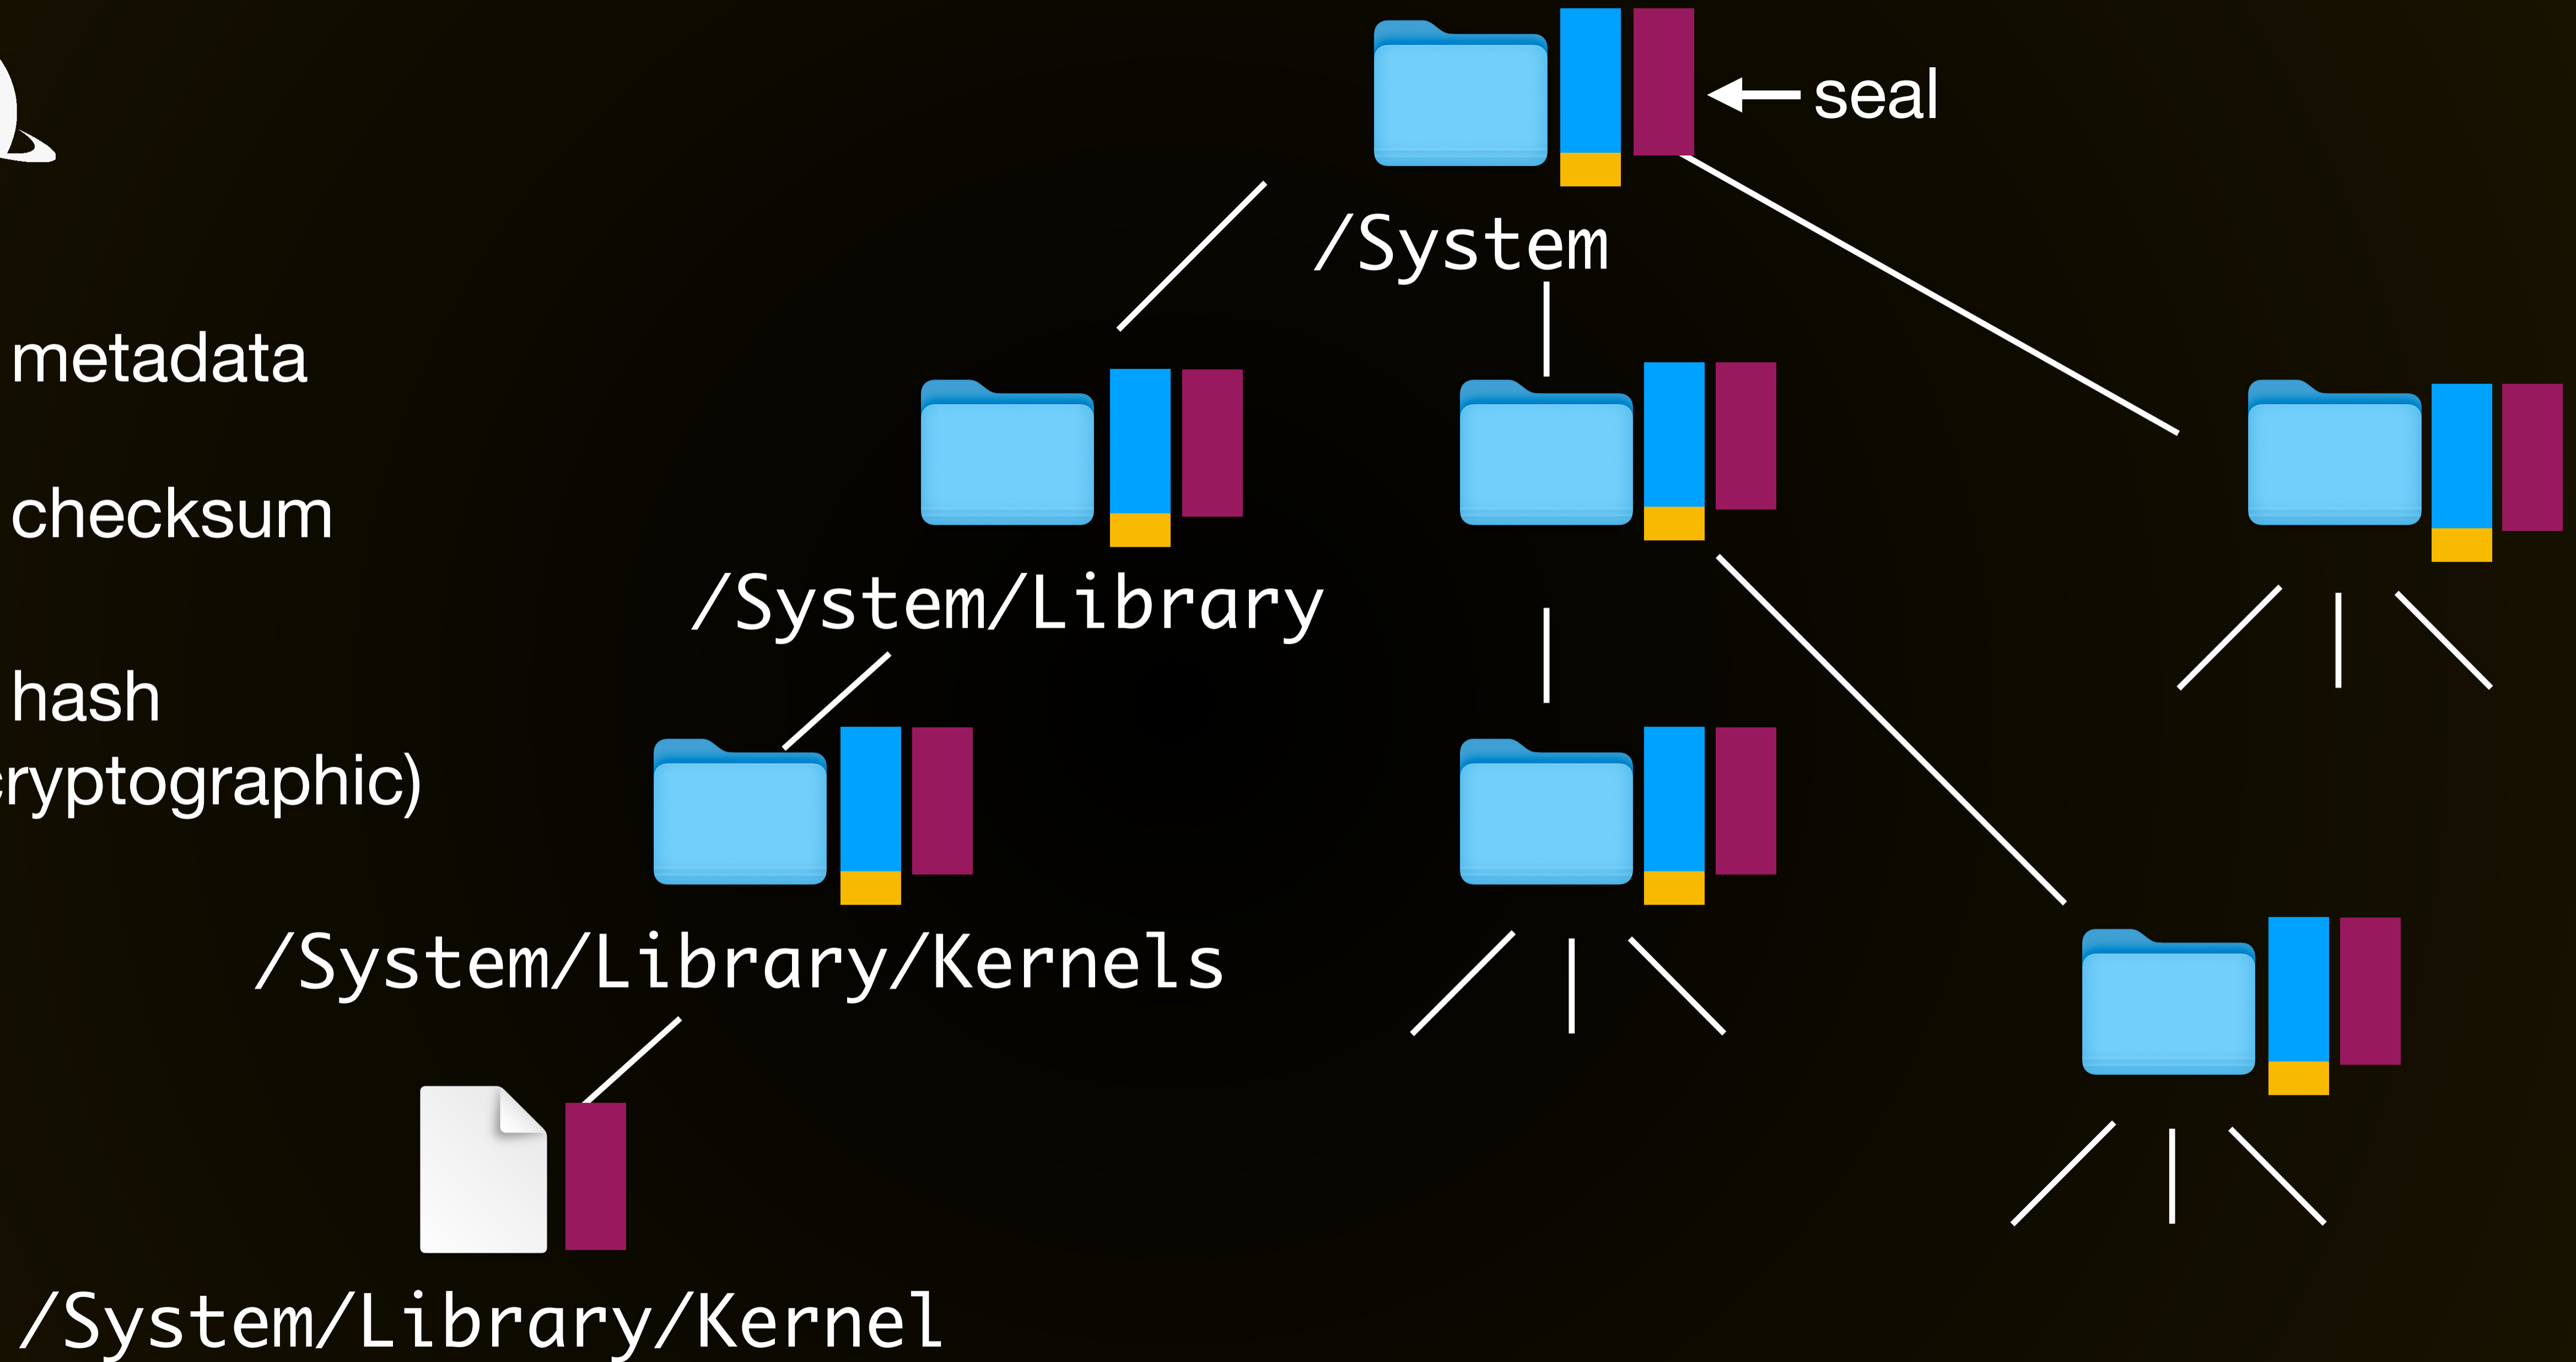
/System/Volumes/Data/Applications



 = metadata

 = checksum

 = hash
(cryptographic)





Creating a bootable volume

- 1. Install or update files**
- 2. Calculate cryptographic hash for all files and directories using file data and directory metadata**
- 3. Save B-tree in volume metadata**
- 4. Compare hash of root with value from Apple**
- 5. Take snapshot of volume**



During startup

- 1. Verify that seal is a valid value before mounting system volume**
- 2. Verify cryptographic hash for every file as it is read**



Places to Visit in Göteborg

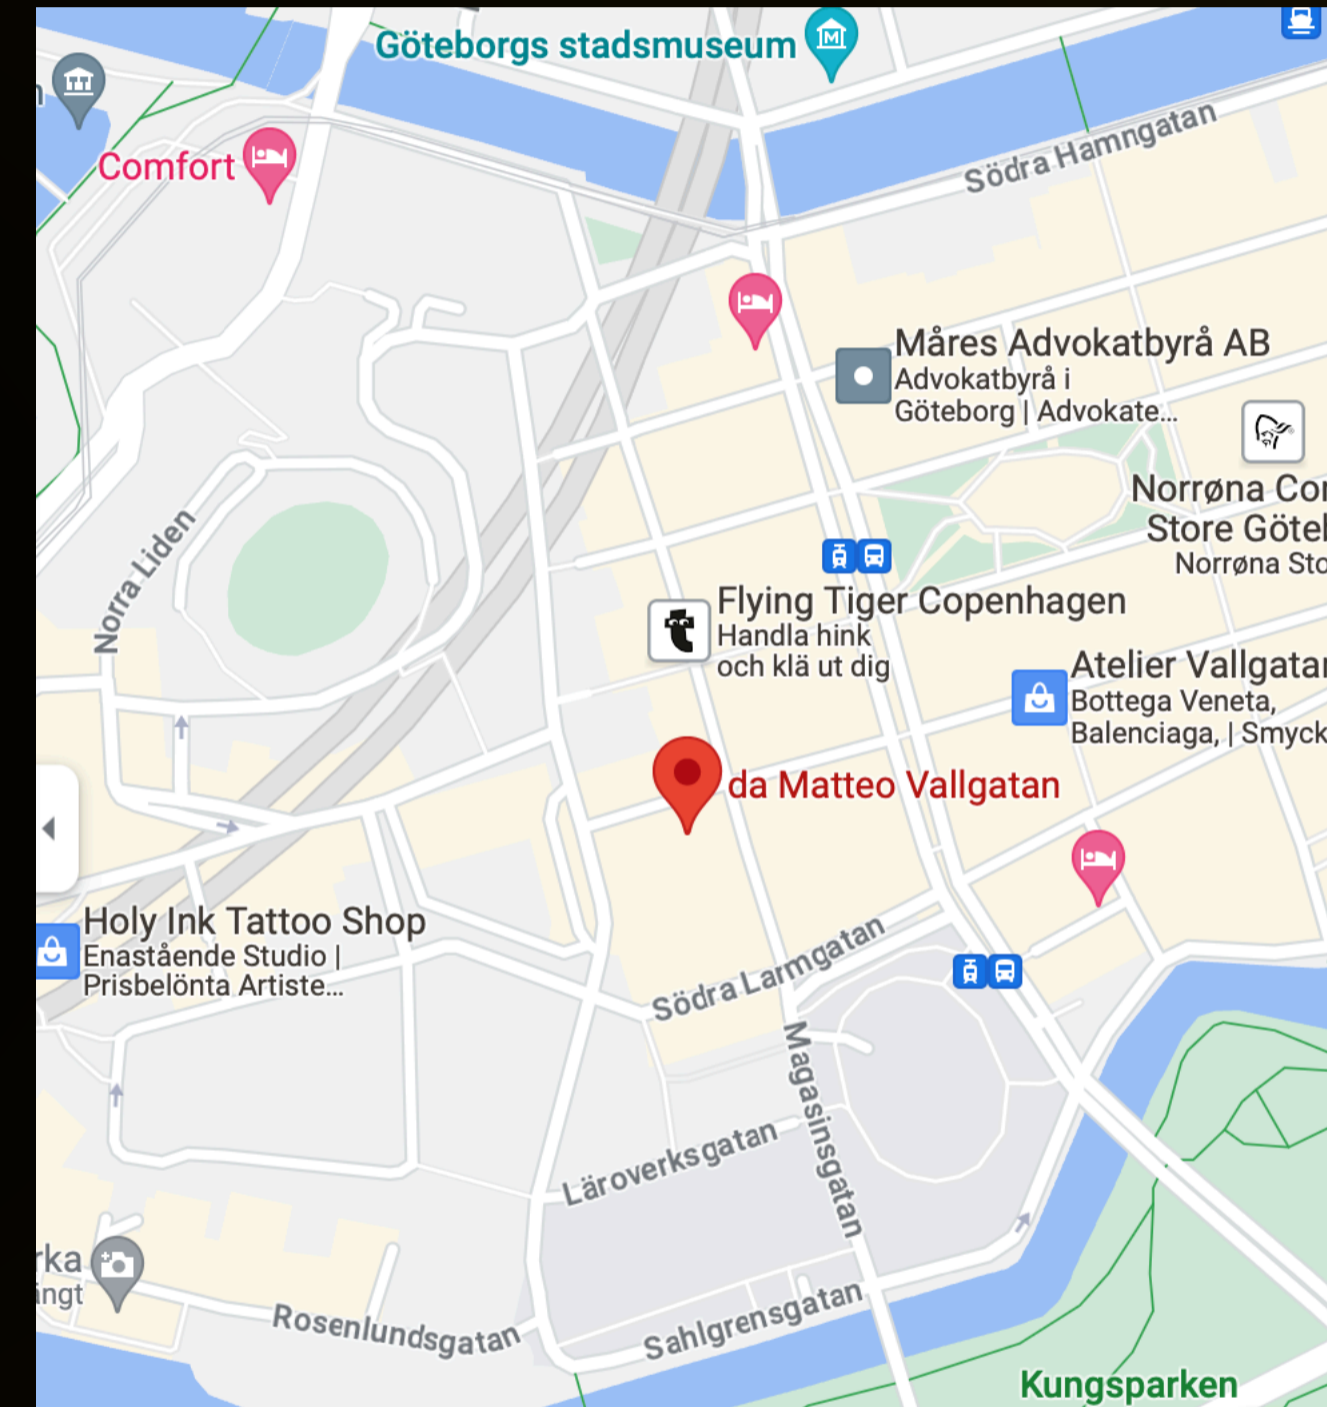




Da Matteo

Valgatan 5

Göteborg

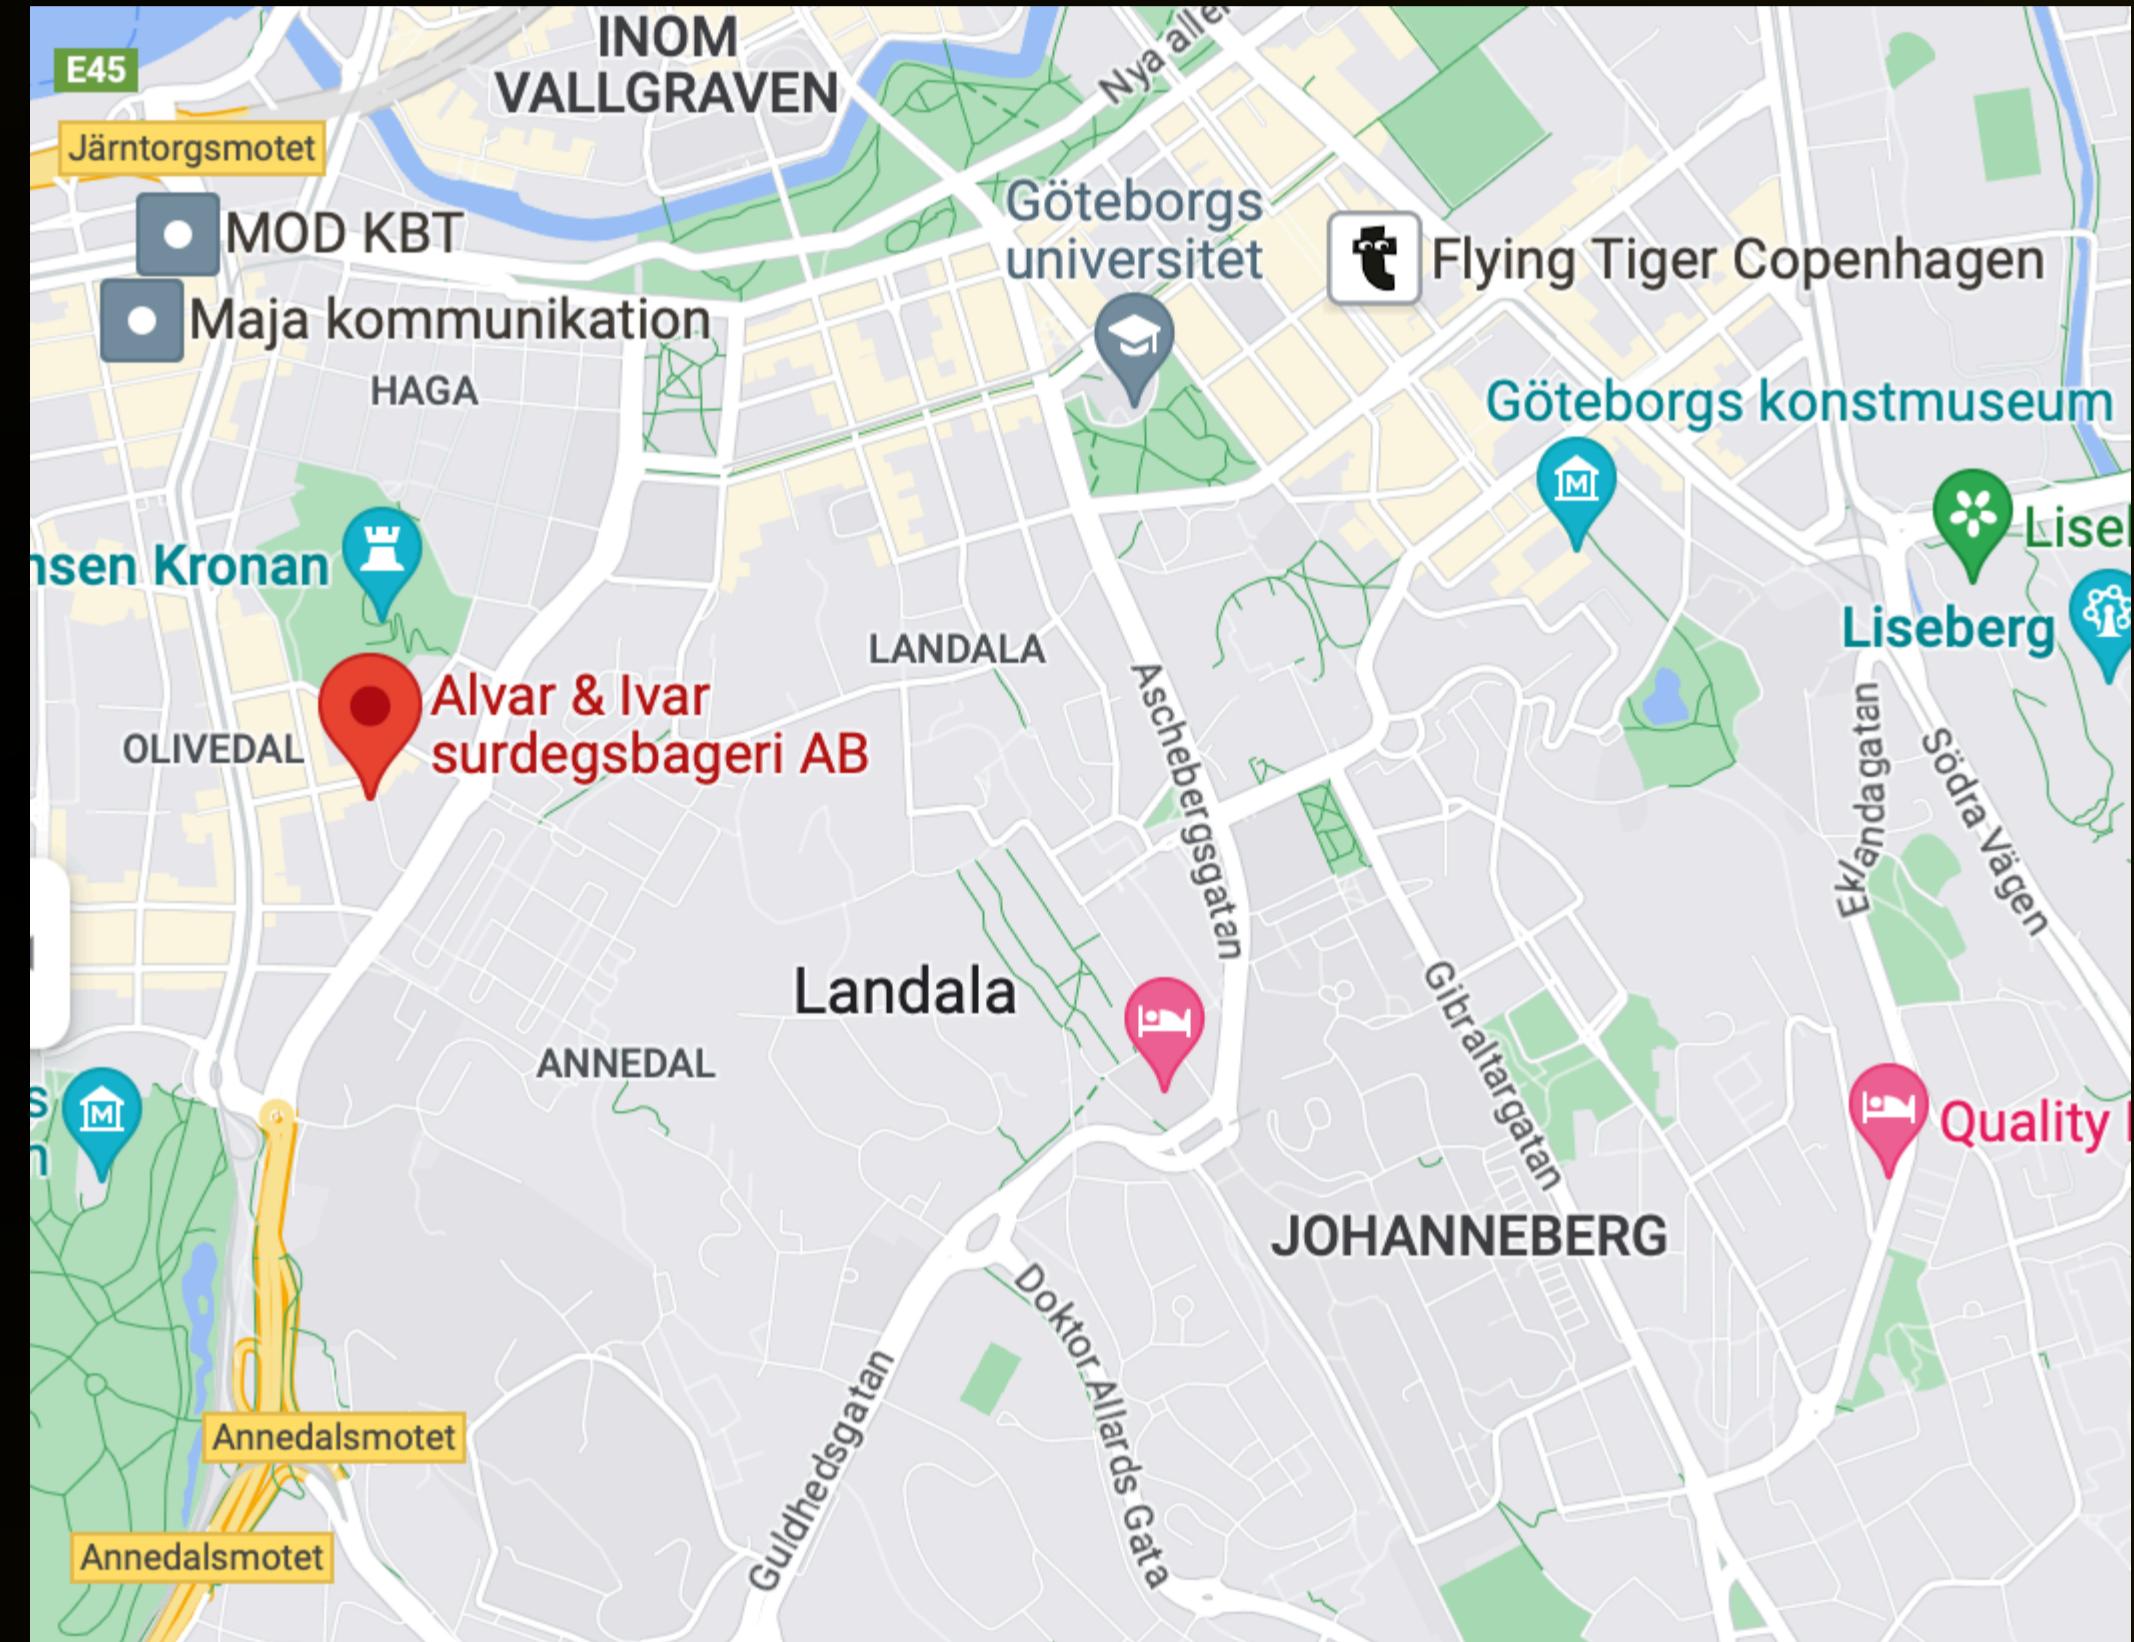








Alvar & Ivar Kastellgatan 11 Göteborg











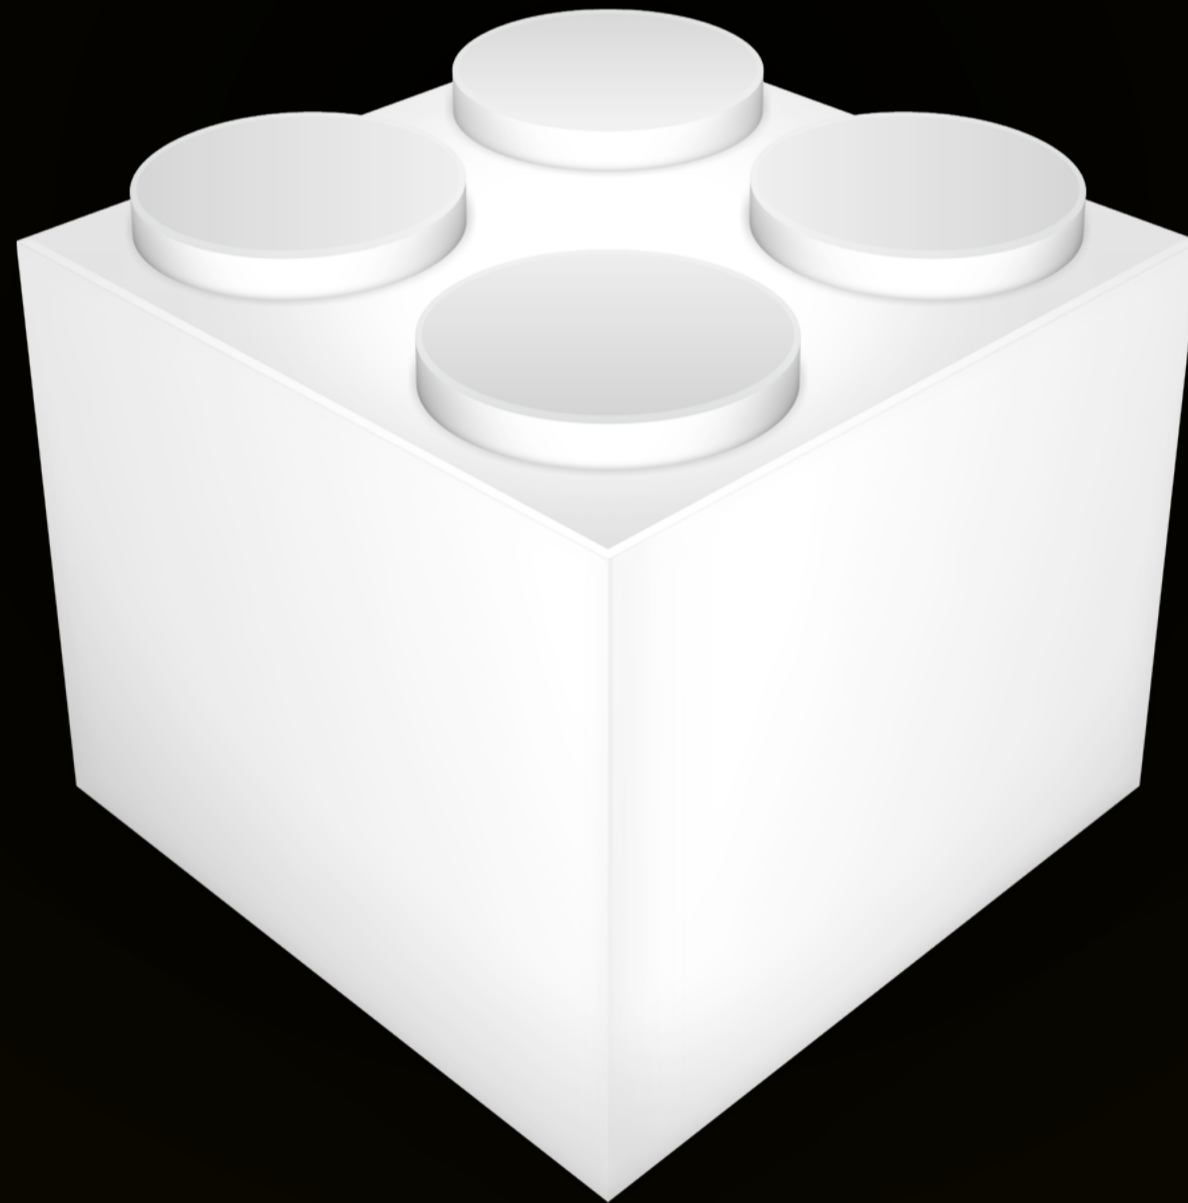
Da Storm & Bille Chokladfabrik

Sveagatan 19
Göteborg



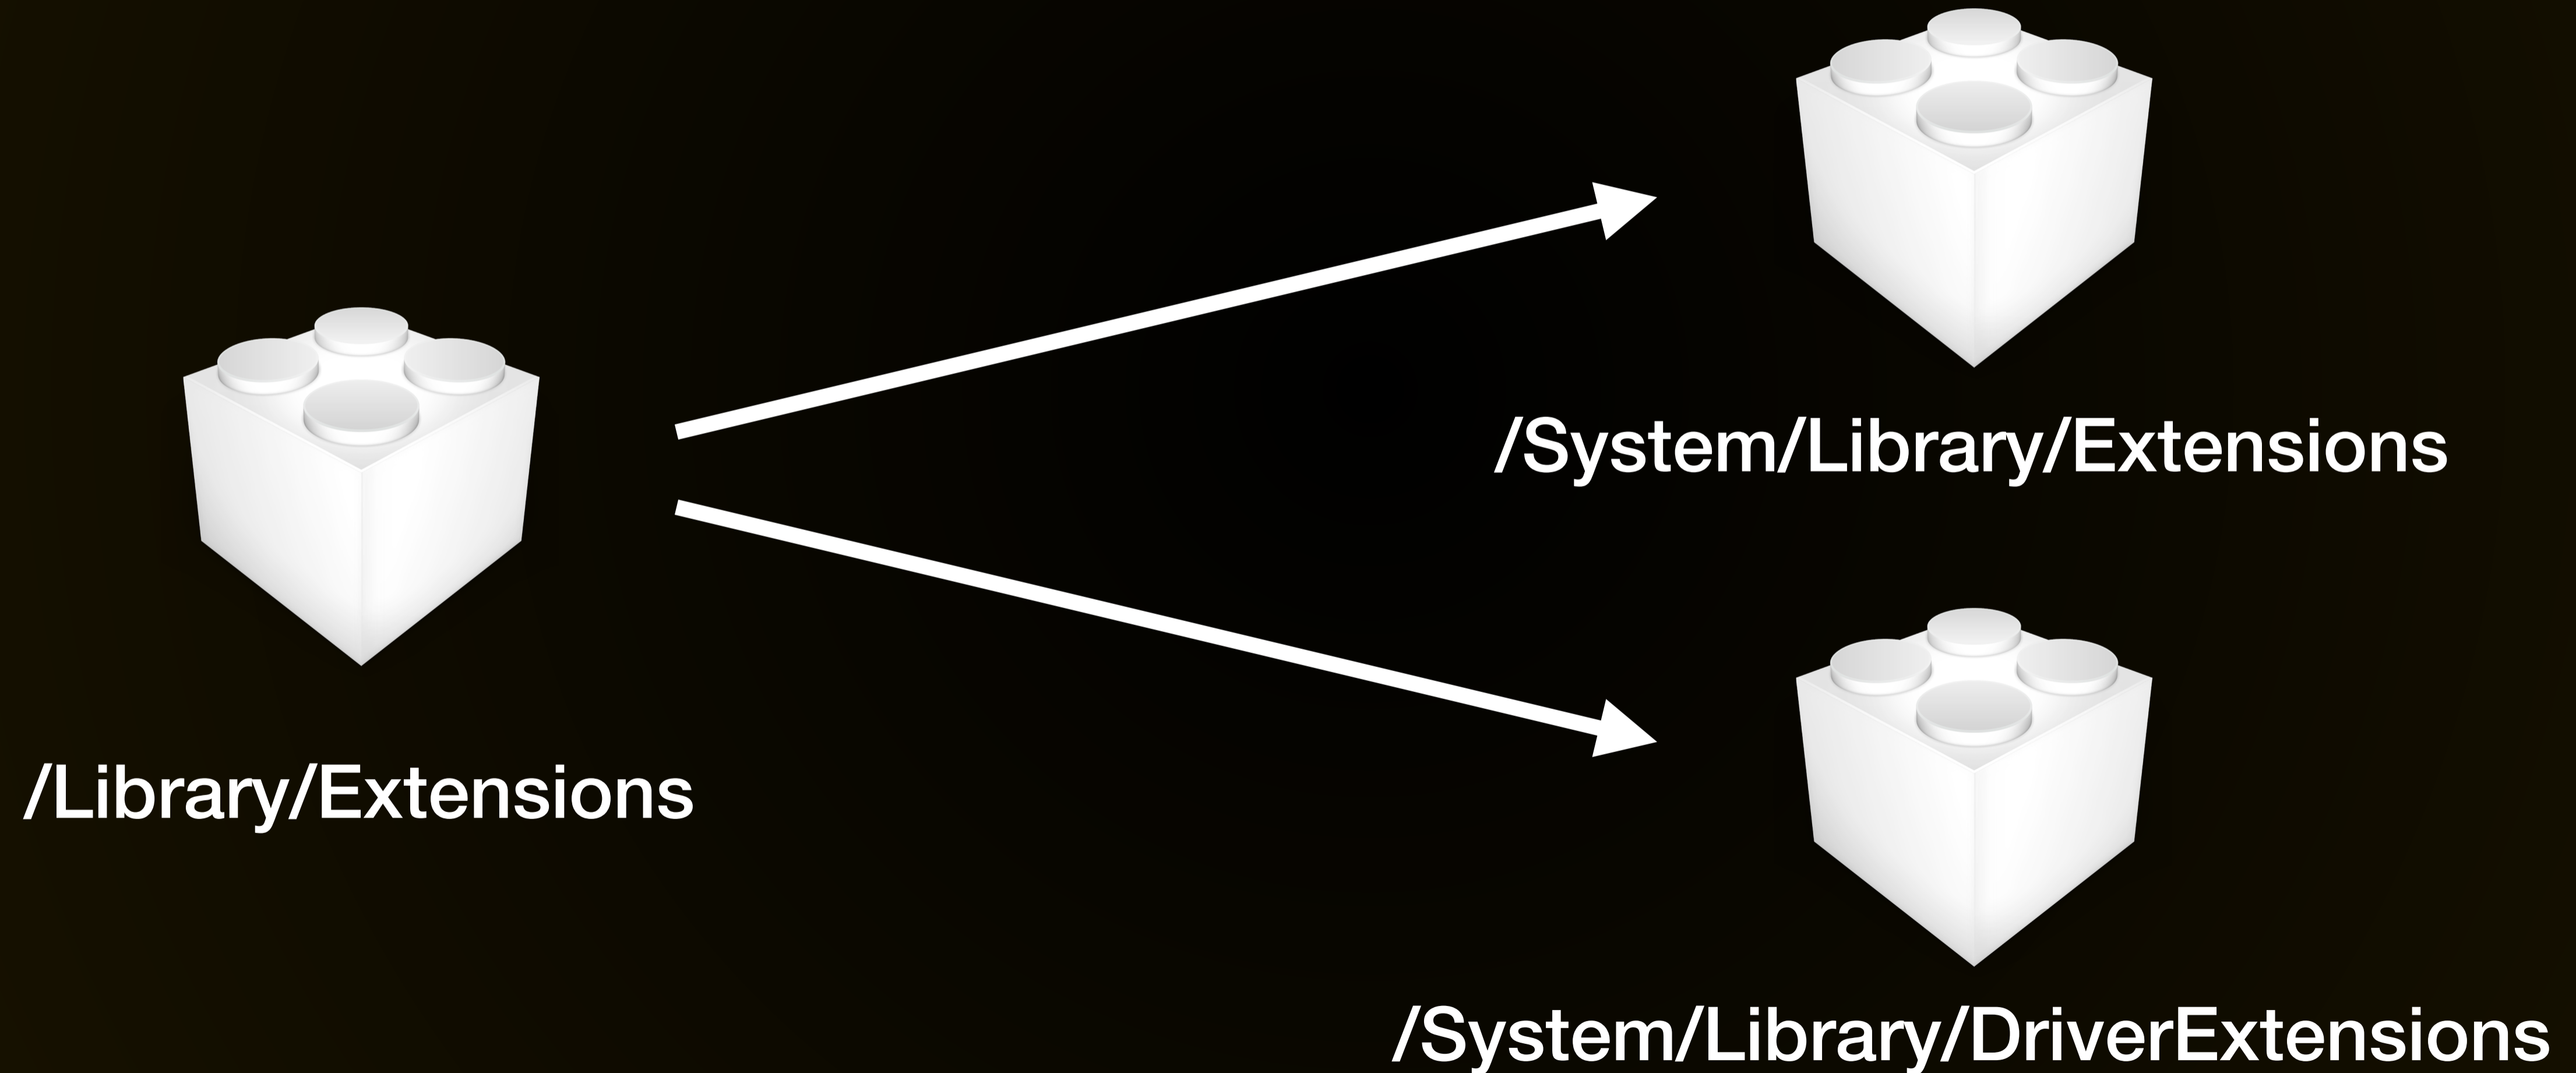


Are kexts going to disappear?





The Future of Kexts





/System/Library/Extensions



/Library/Extensions

Protected by SSV

Yes

No

Requires Reduced Security

No

Yes

Requires Admin Approval

No

Yes

User Loadable

No

Yes

Codes Signed

No

Yes

Notarized

No

Yes



SoftRAID kext

Before macOS 13.3

/Library/Extensions/SoftRAID.kext

Requires Reduced Security

Requires Admin Approval

**Requires New Driver Install
on Apple silicon**

macOS 13.3 and later

/System/Library/Extensions/SoftRAID.kext

No Reduced Security

No Admin Approval

No Driver Install Possible



SoftRAID kext

macOS 13.3 - 13.6

SoftRAID kext version 7.5

macOS 14.0 and later

SoftRAID kext version 7.6



MacBook Pro

Extension N...	Version	Last Modified	Notarized	Loaded	Obtained from
smbfs	5.0	9/16/23, 9:48 AM	Unknown	No	Not Signed
SMCMotionSensor	3.0.4	9/16/23, 9:48 AM	Unknown	No	Not Signed
SoftRAID	7.5	9/16/23, 9:48 AM	Unknown	No	Not Signed
SonyXDCAMDriver	554	9/16/23, 9:48 AM	Yes	No	Apple
StorageLynx	554	9/16/23, 9:48 AM	Unknown	No	Not Signed
System	22.6.0	9/16/23, 9:48 AM	Yes	No	Apple
tmpfs	1.0	9/16/23, 9:48 AM	Unknown	No	Not Signed
triggers	3.0	9/16/23, 9:48 AM	Unknown	No	Not Signed

Bundle ID:

com.apple.filesystems.smbfs

Notarized:

Unknown

Loaded:

No

Obtained from:

Not Signed

Location:

/System/Library/Extensions/smbfs.kext

Kext Version:

5.0

Loadable:

No

Validity Errors:

Validation Failures:

Kext has a CFBundleExecutable property but the executable can't be found: smbfs

Signature Validation Errors:

Not Signed

Dependencies:

Incomplete

Signed by:

Not Signed

Yu > Software > Extensions > smbfs



MacBook Pro

Extension N...	Version	Last Modified	Notarized	Loaded	Obtained from
smbfs	5.0	9/16/23, 9:48 AM	Unknown	No	Not Signed
SMCMotionSensor	3.0.4	9/16/23, 9:48 AM	Unknown	No	Not Signed
SoftRAID	7.5	9/16/23, 9:48 AM	Unknown	No	Not Signed
SonyXDCAMDriver	554	9/16/23, 9:48 AM	Yes	No	Apple
StorageLynx	554	9/16/23, 9:48 AM	Unknown	No	Not Signed
System	22.6.0	9/16/23, 9:48 AM	Yes	No	Apple
tmpfs	1.0	9/16/23, 9:48 AM	Unknown	No	Not Signed
triggers	3.0	9/16/23, 9:48 AM	Unknown	No	Not Signed

Bundle ID:

com.apple.driver.SoftRAID

Notarized:

Unknown

Loaded:

No

Get Info String:

SoftRAID version 7.5, Copyright © 2002-23 Other World Computing, Inc. All rights reserved.

Obtained from:

Not Signed

Location:

/System/Library/Extensions/SoftRAID.kext

Kext Version:

7.5

Loadable:

No

Validity Errors:

Validation Failures:

Kext has a CFBundleExecutable property but the executable can't be found: SoftRAID

Signature Validation Errors:

Not Signed

Dependencies:

Incomplete

Signed by:

Not Signed

Yu > Software > Extensions > SoftRAID



What is Bit-Rot?



"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 on Disk





Testing for Data Corruption on Disk

- **2 - RAID 5 volumes, each with 15 disks**
- **First volume is active, constant reads and writes**
- **Second volume is passive, contains data, no reads or writes**
- **Every 10 days, verified that the parity information is correct**
- **Repeat for 18 months (54 parity verify operations)**

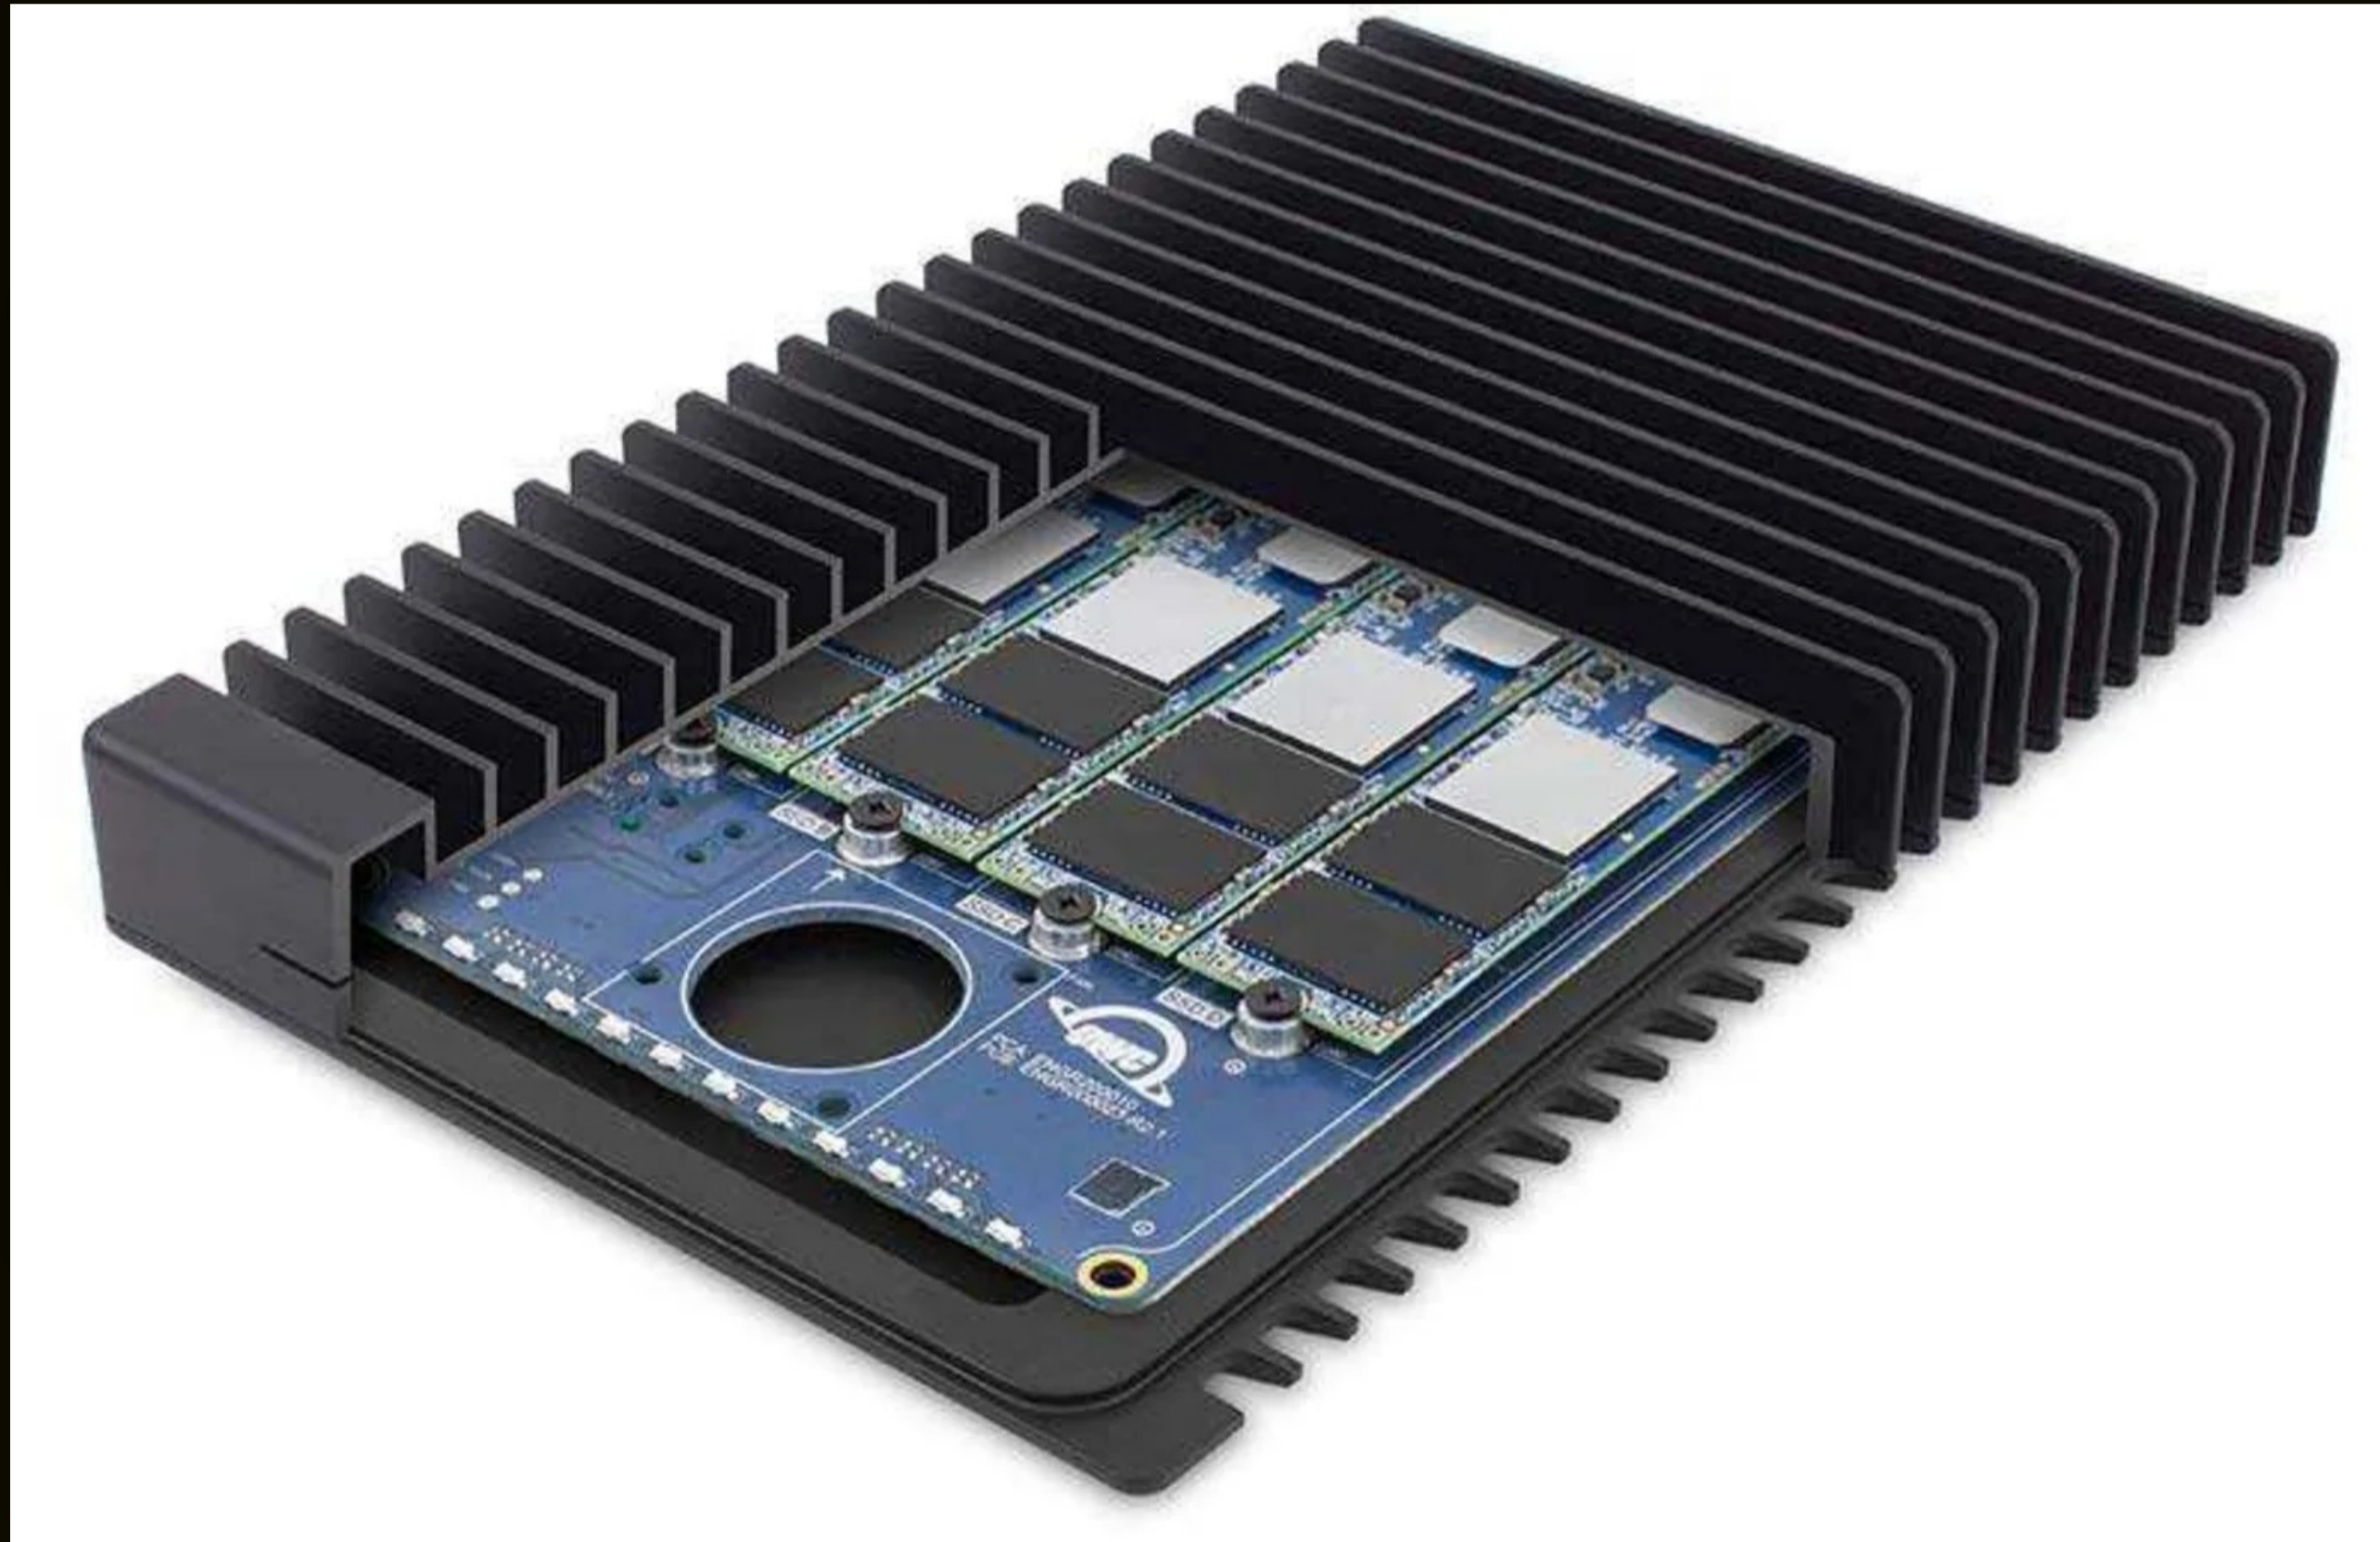


Testing for Data Corruption on Disk

- **21 PB transferred, 3 billion i/os**
- **Only incorrect parity was on active volume and was result of power loss while writing to the volume**



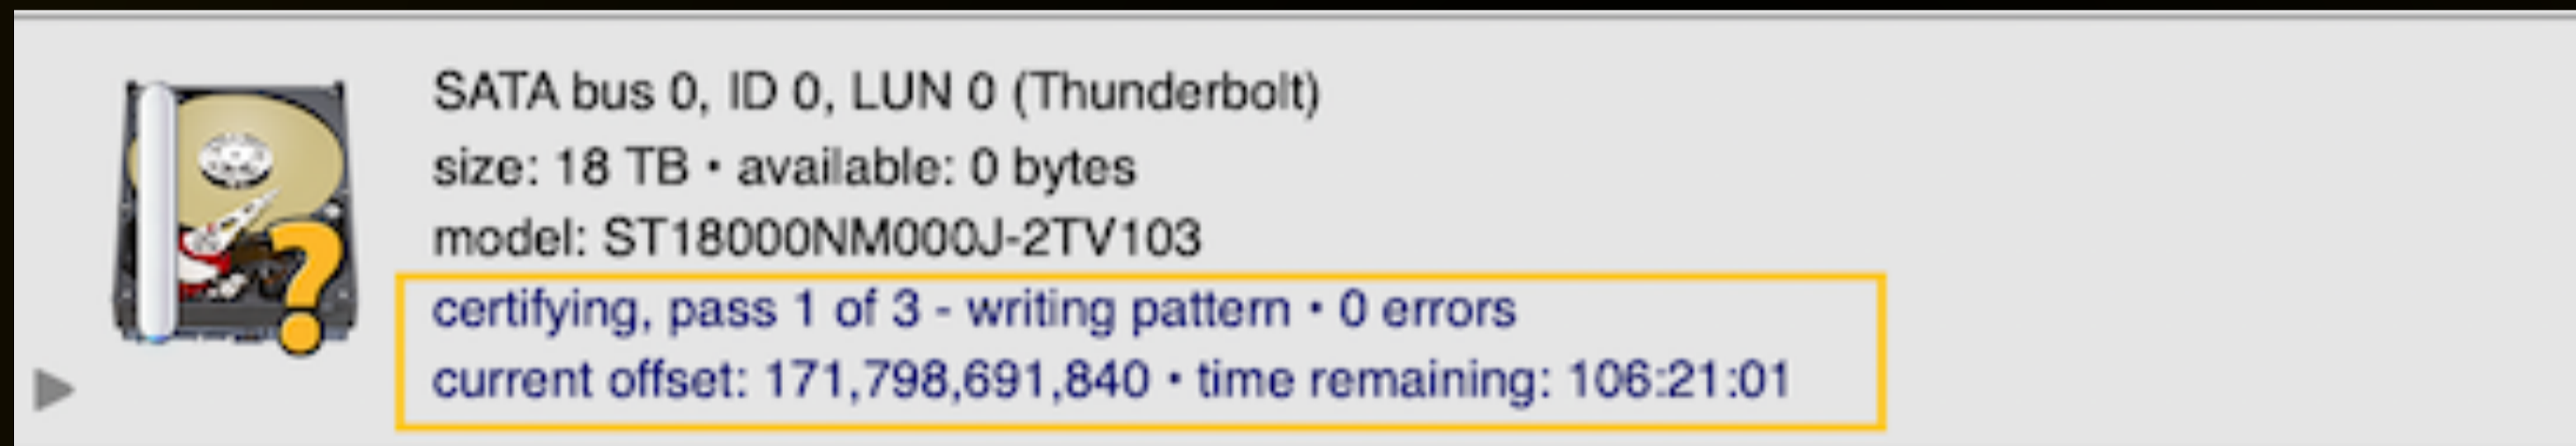
Bit-Rot Observed





Detecting Bit-Rot

1) SoftRAID Certify Function



2) Copy and verify 200 GB file

```
cp ~/Desktop/200GB.zip /Volumes/TestVolume/
```

```
cmp -l ~/Desktop/200GB.zip /Volumes/TestVolume/200GB.zip
```

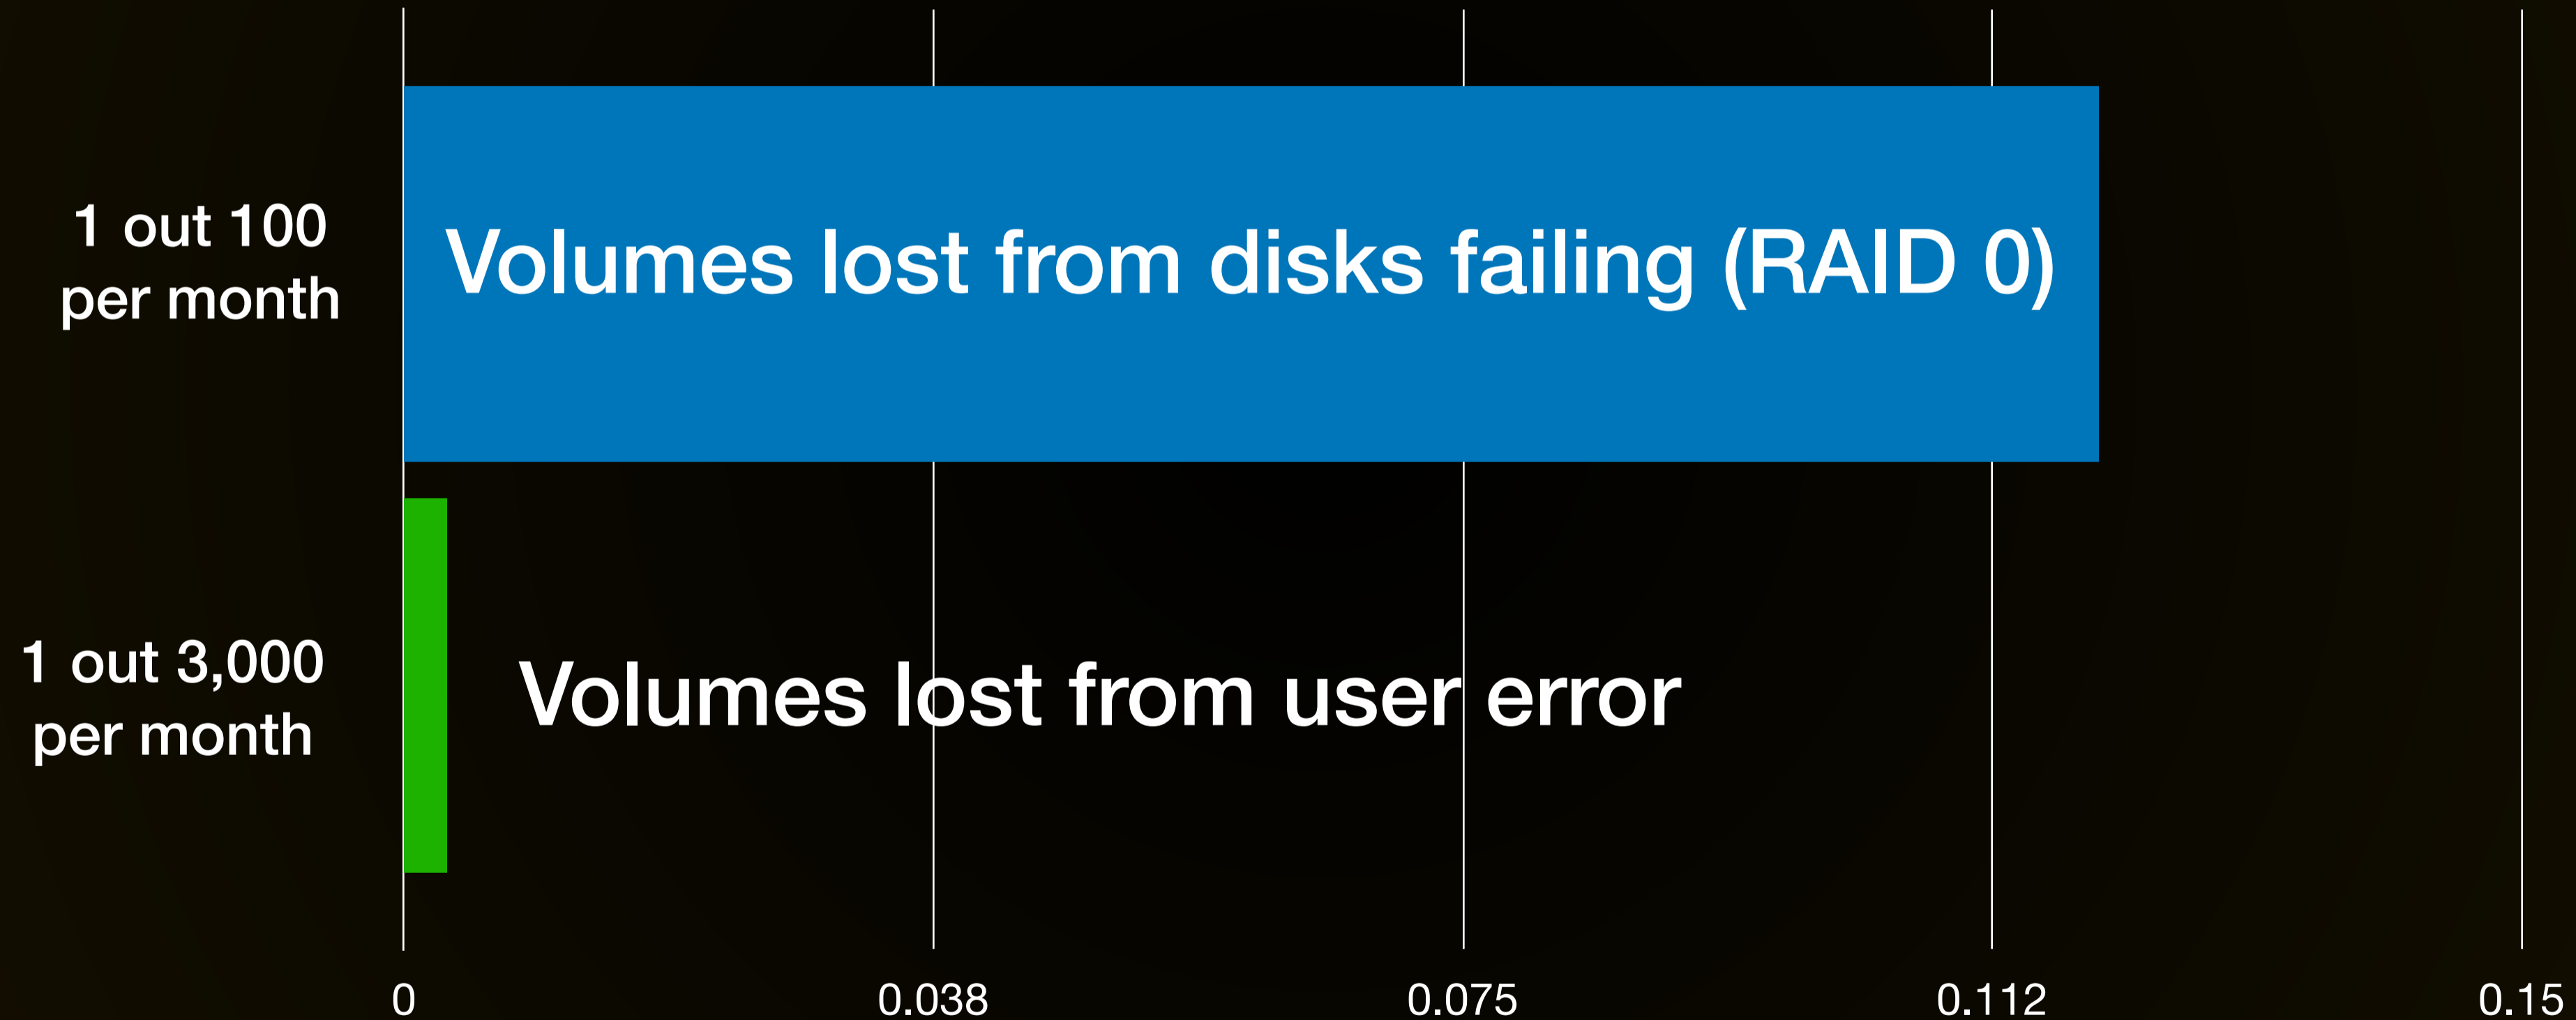


History of Volume Loss



Volume Failure Rates

During the year 2008



Annual Failure Rates

(all numbers approximate)



Volume Failure Rates

During the year 2010

1 out 100
per month

Volumes lost from disks failing (RAID 0)

1 out
2,000,000
per month

Volumes lost from user error

0

0.038

0.075

0.112

0.15

Annual Failure Rates

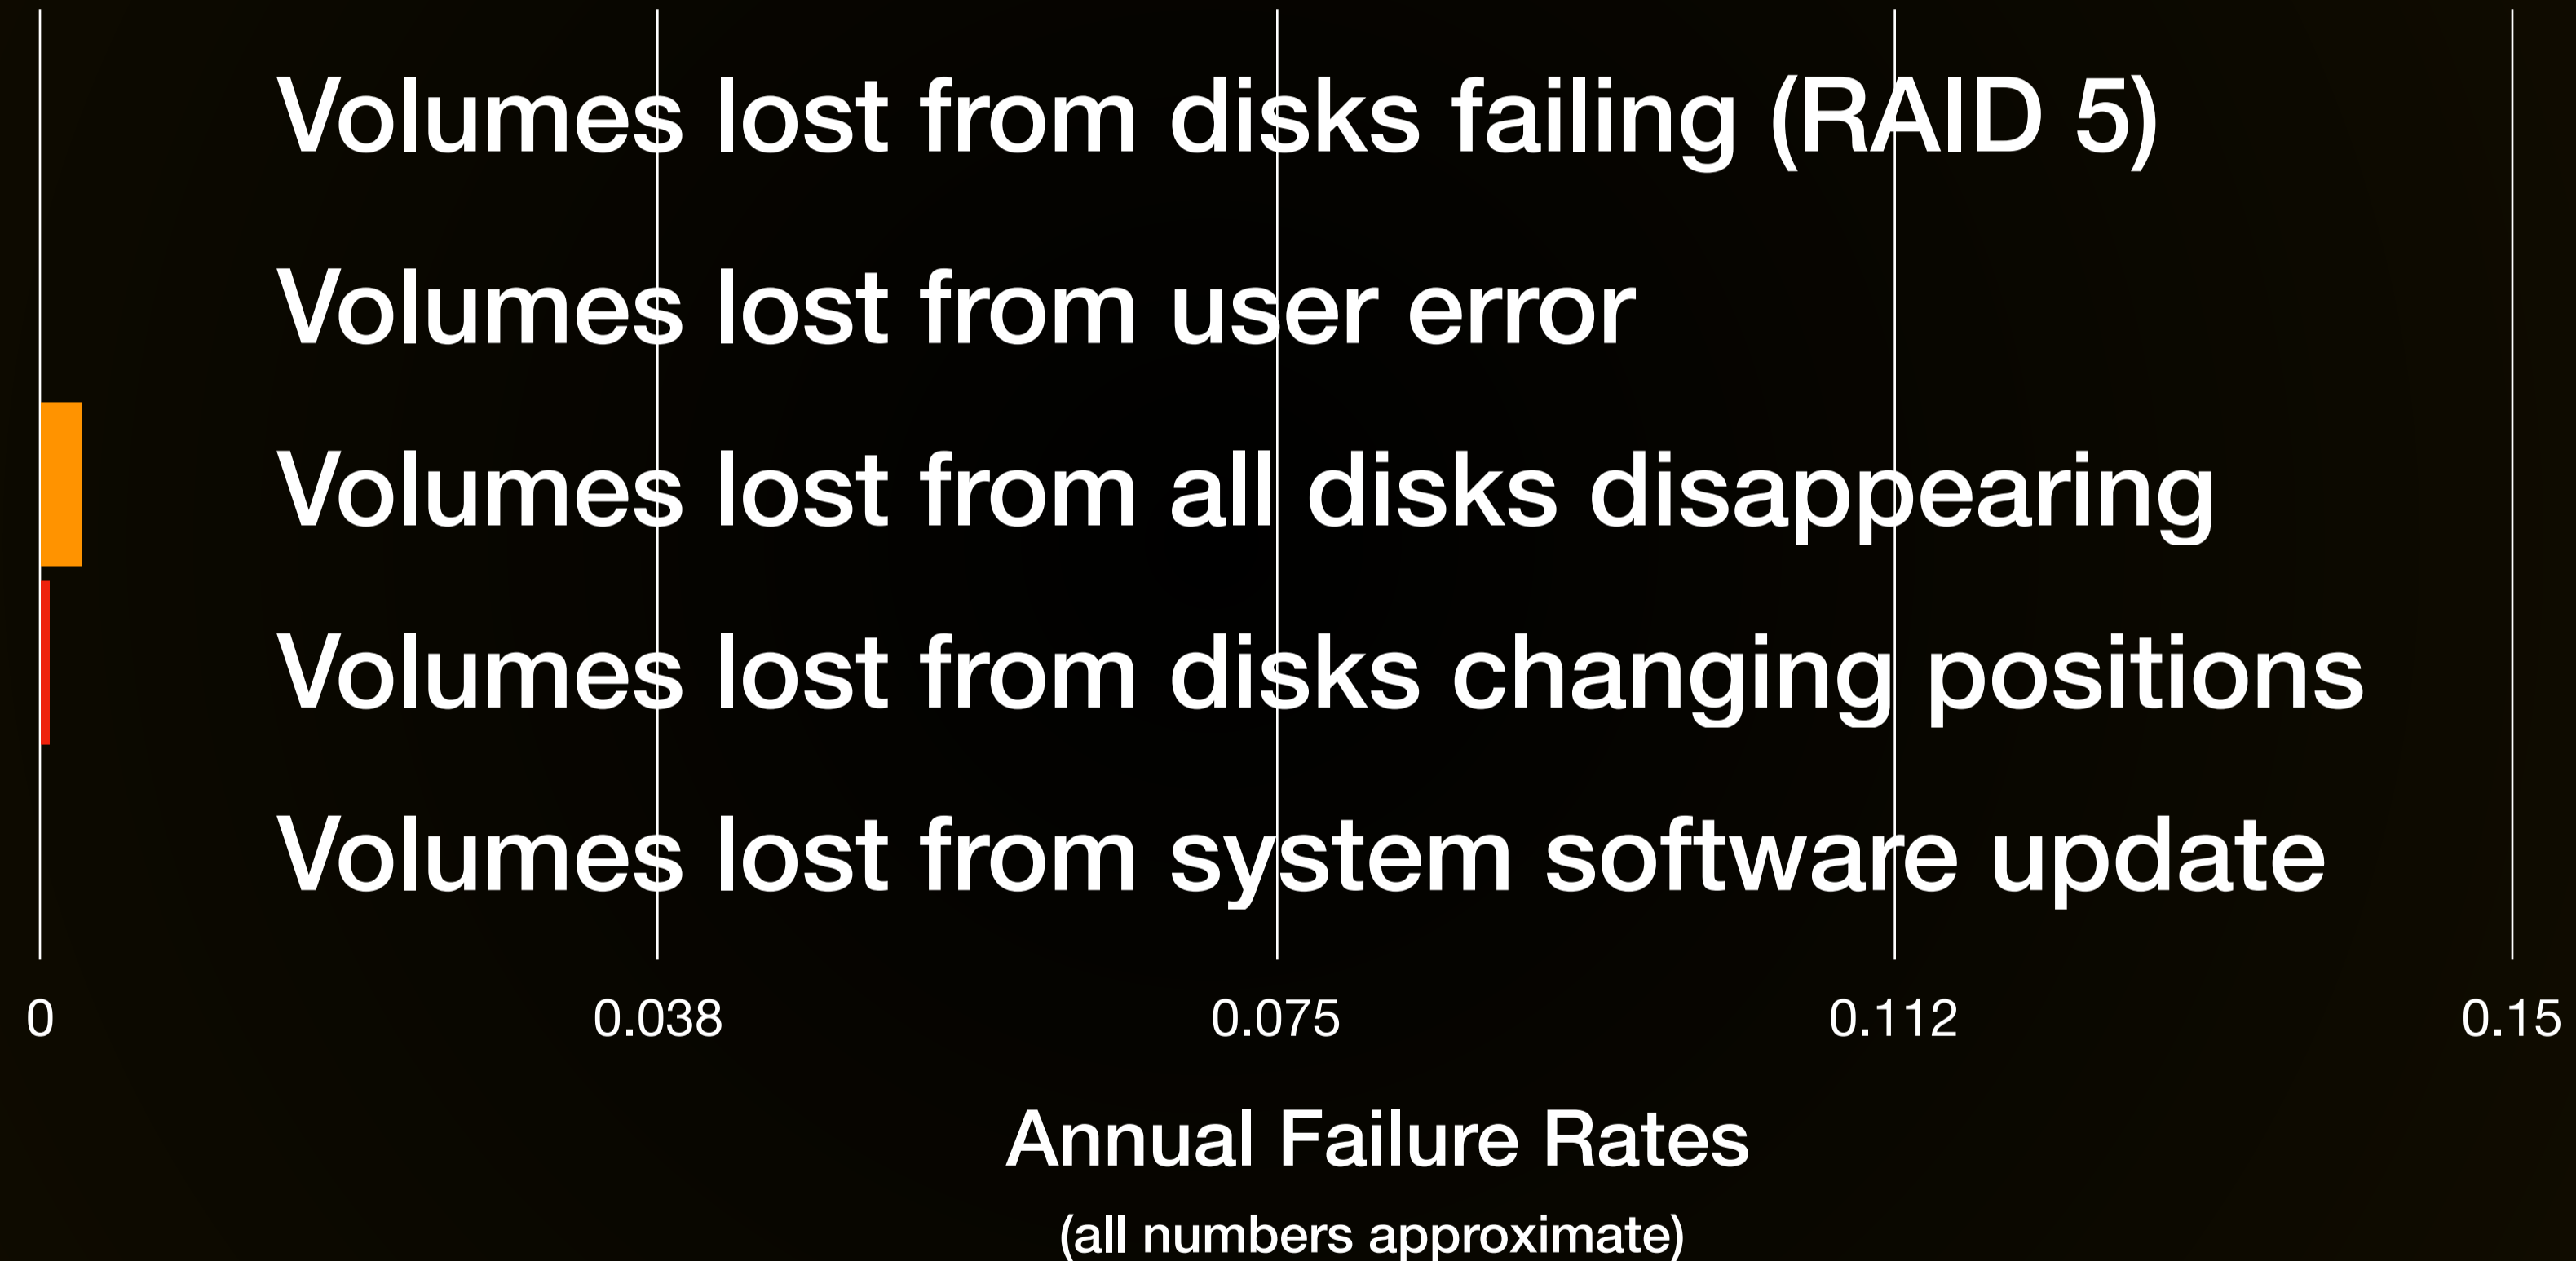
(all numbers approximate)

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



Volume Failure Rates

During the year 2014

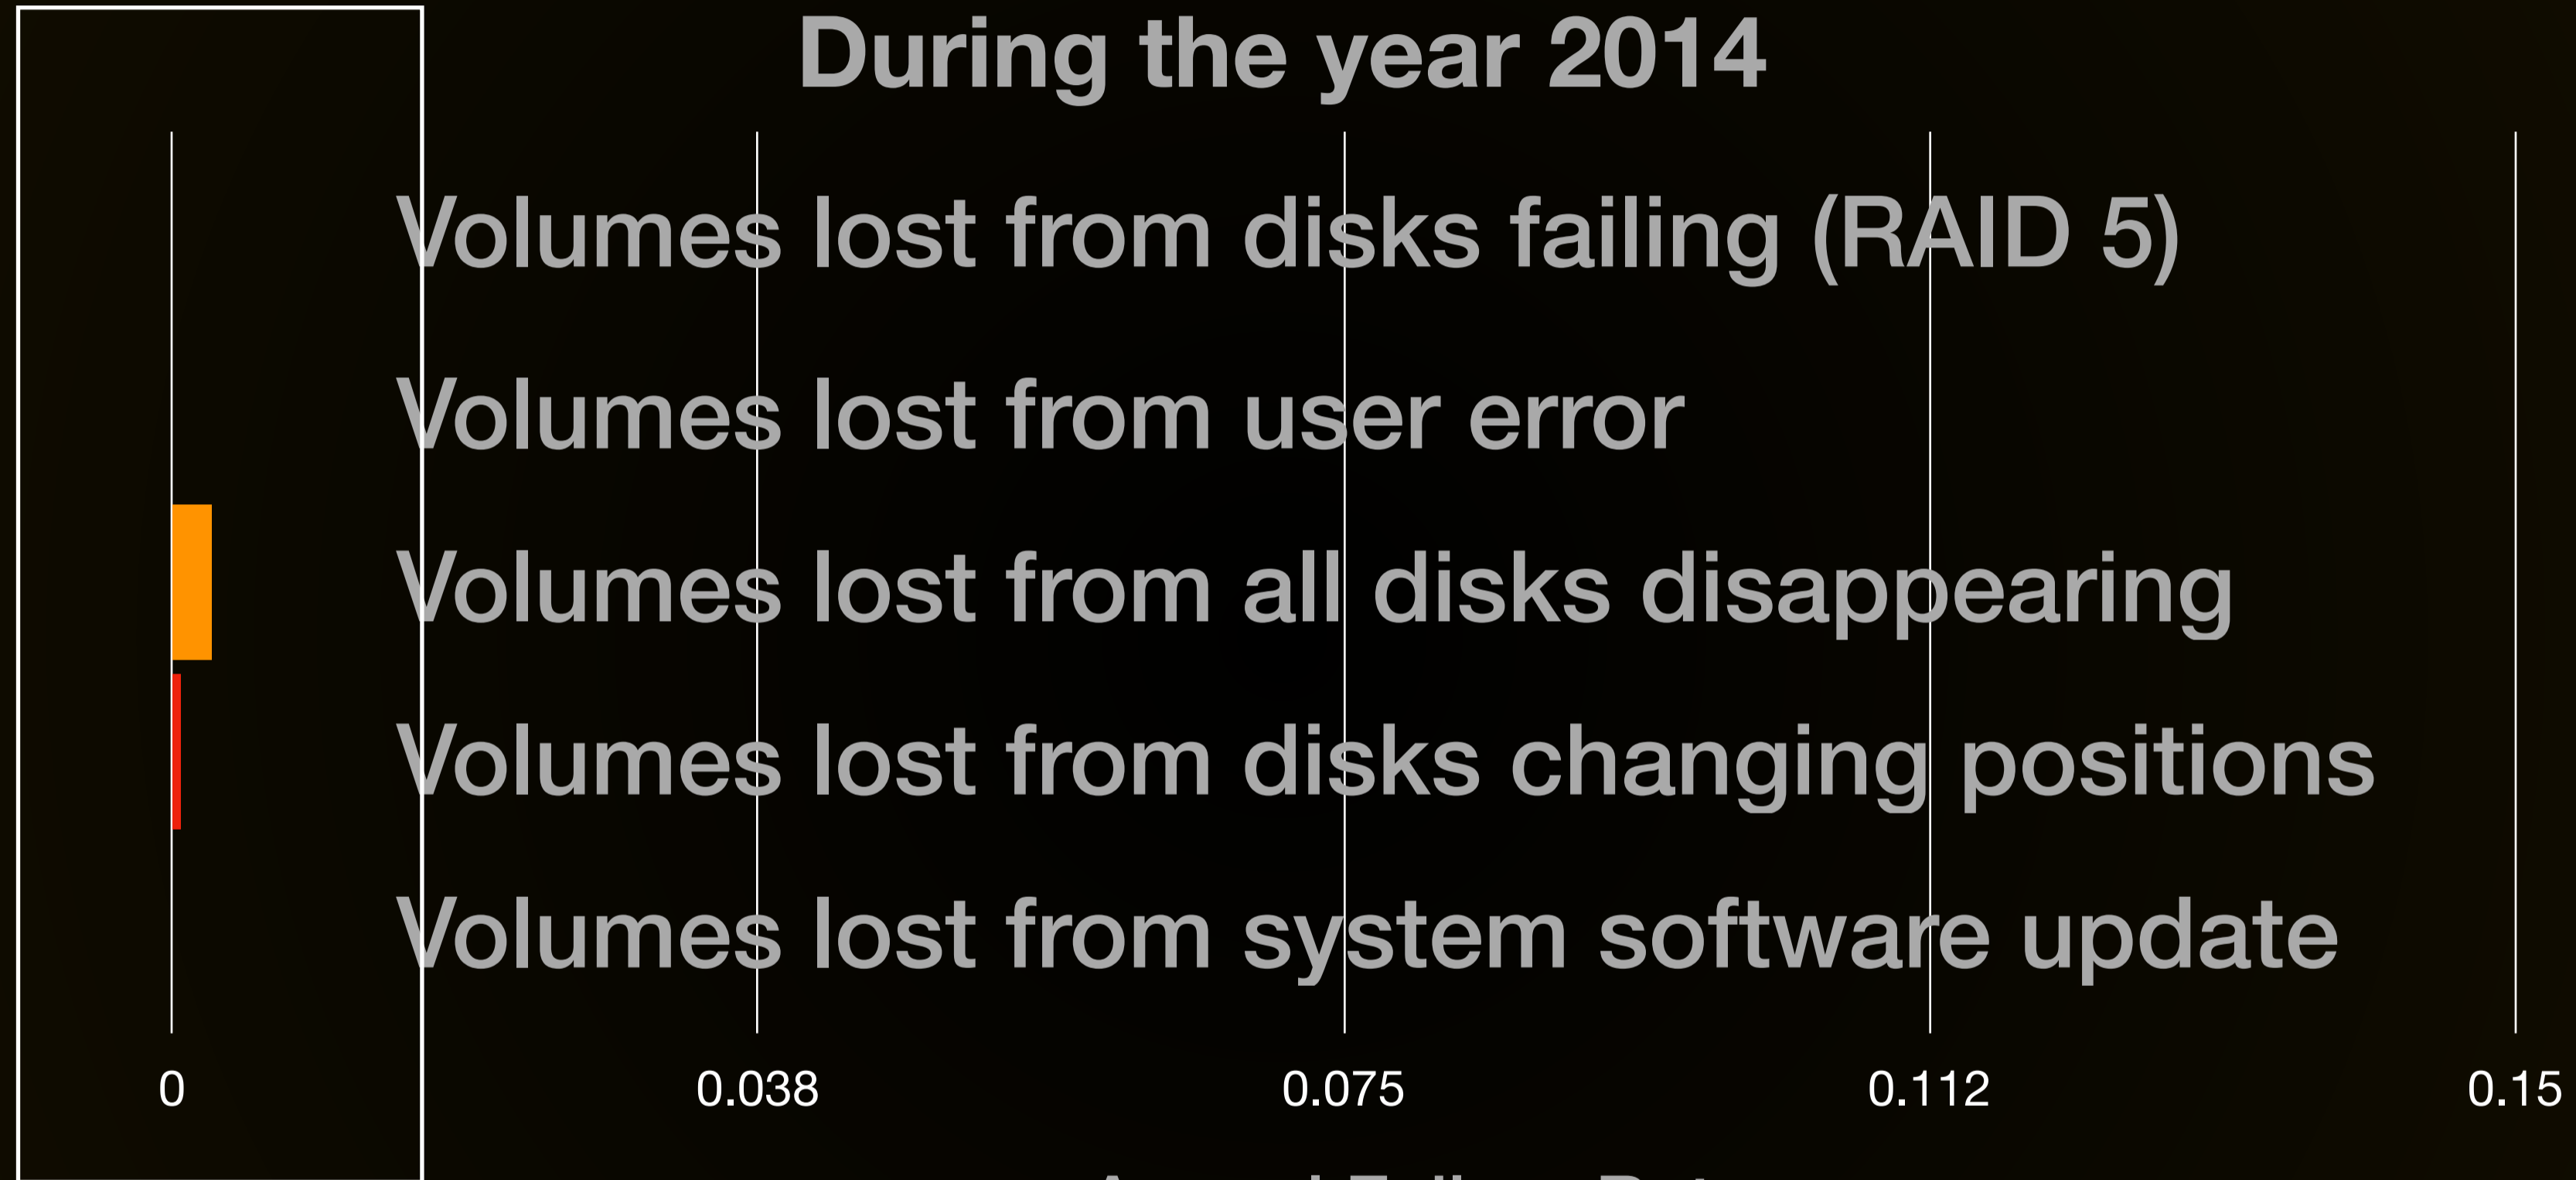


SoftRAID v. 5 added RAID 4 and 5 volumes



Volume Failure Rates

During the year 2014



Annual Failure Rates

(all numbers approximate)

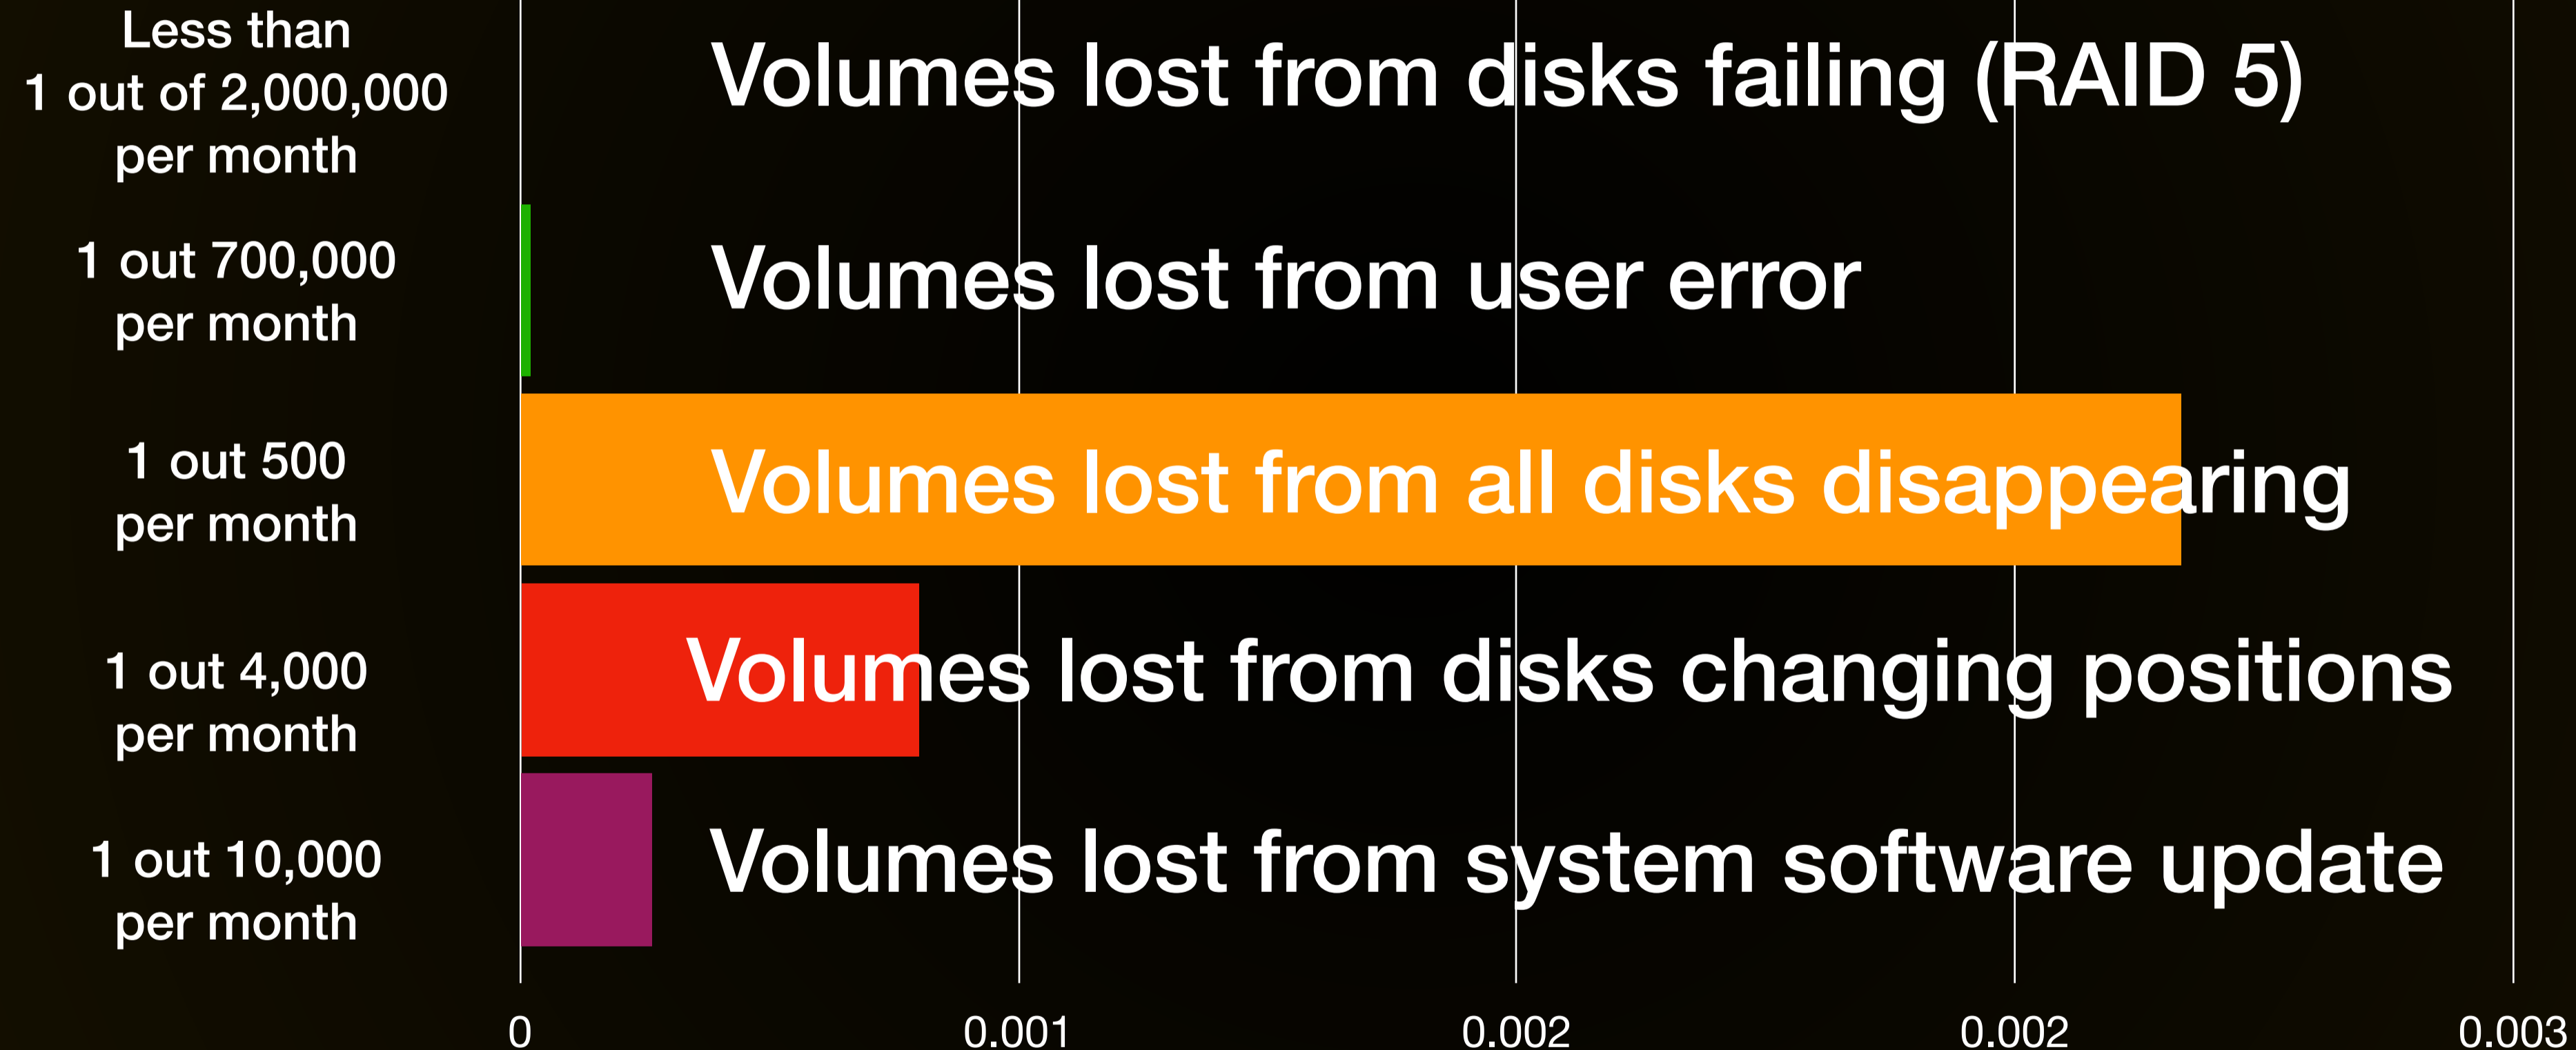


Zoom in 50x



Volume Failure Rates

During the year 2014



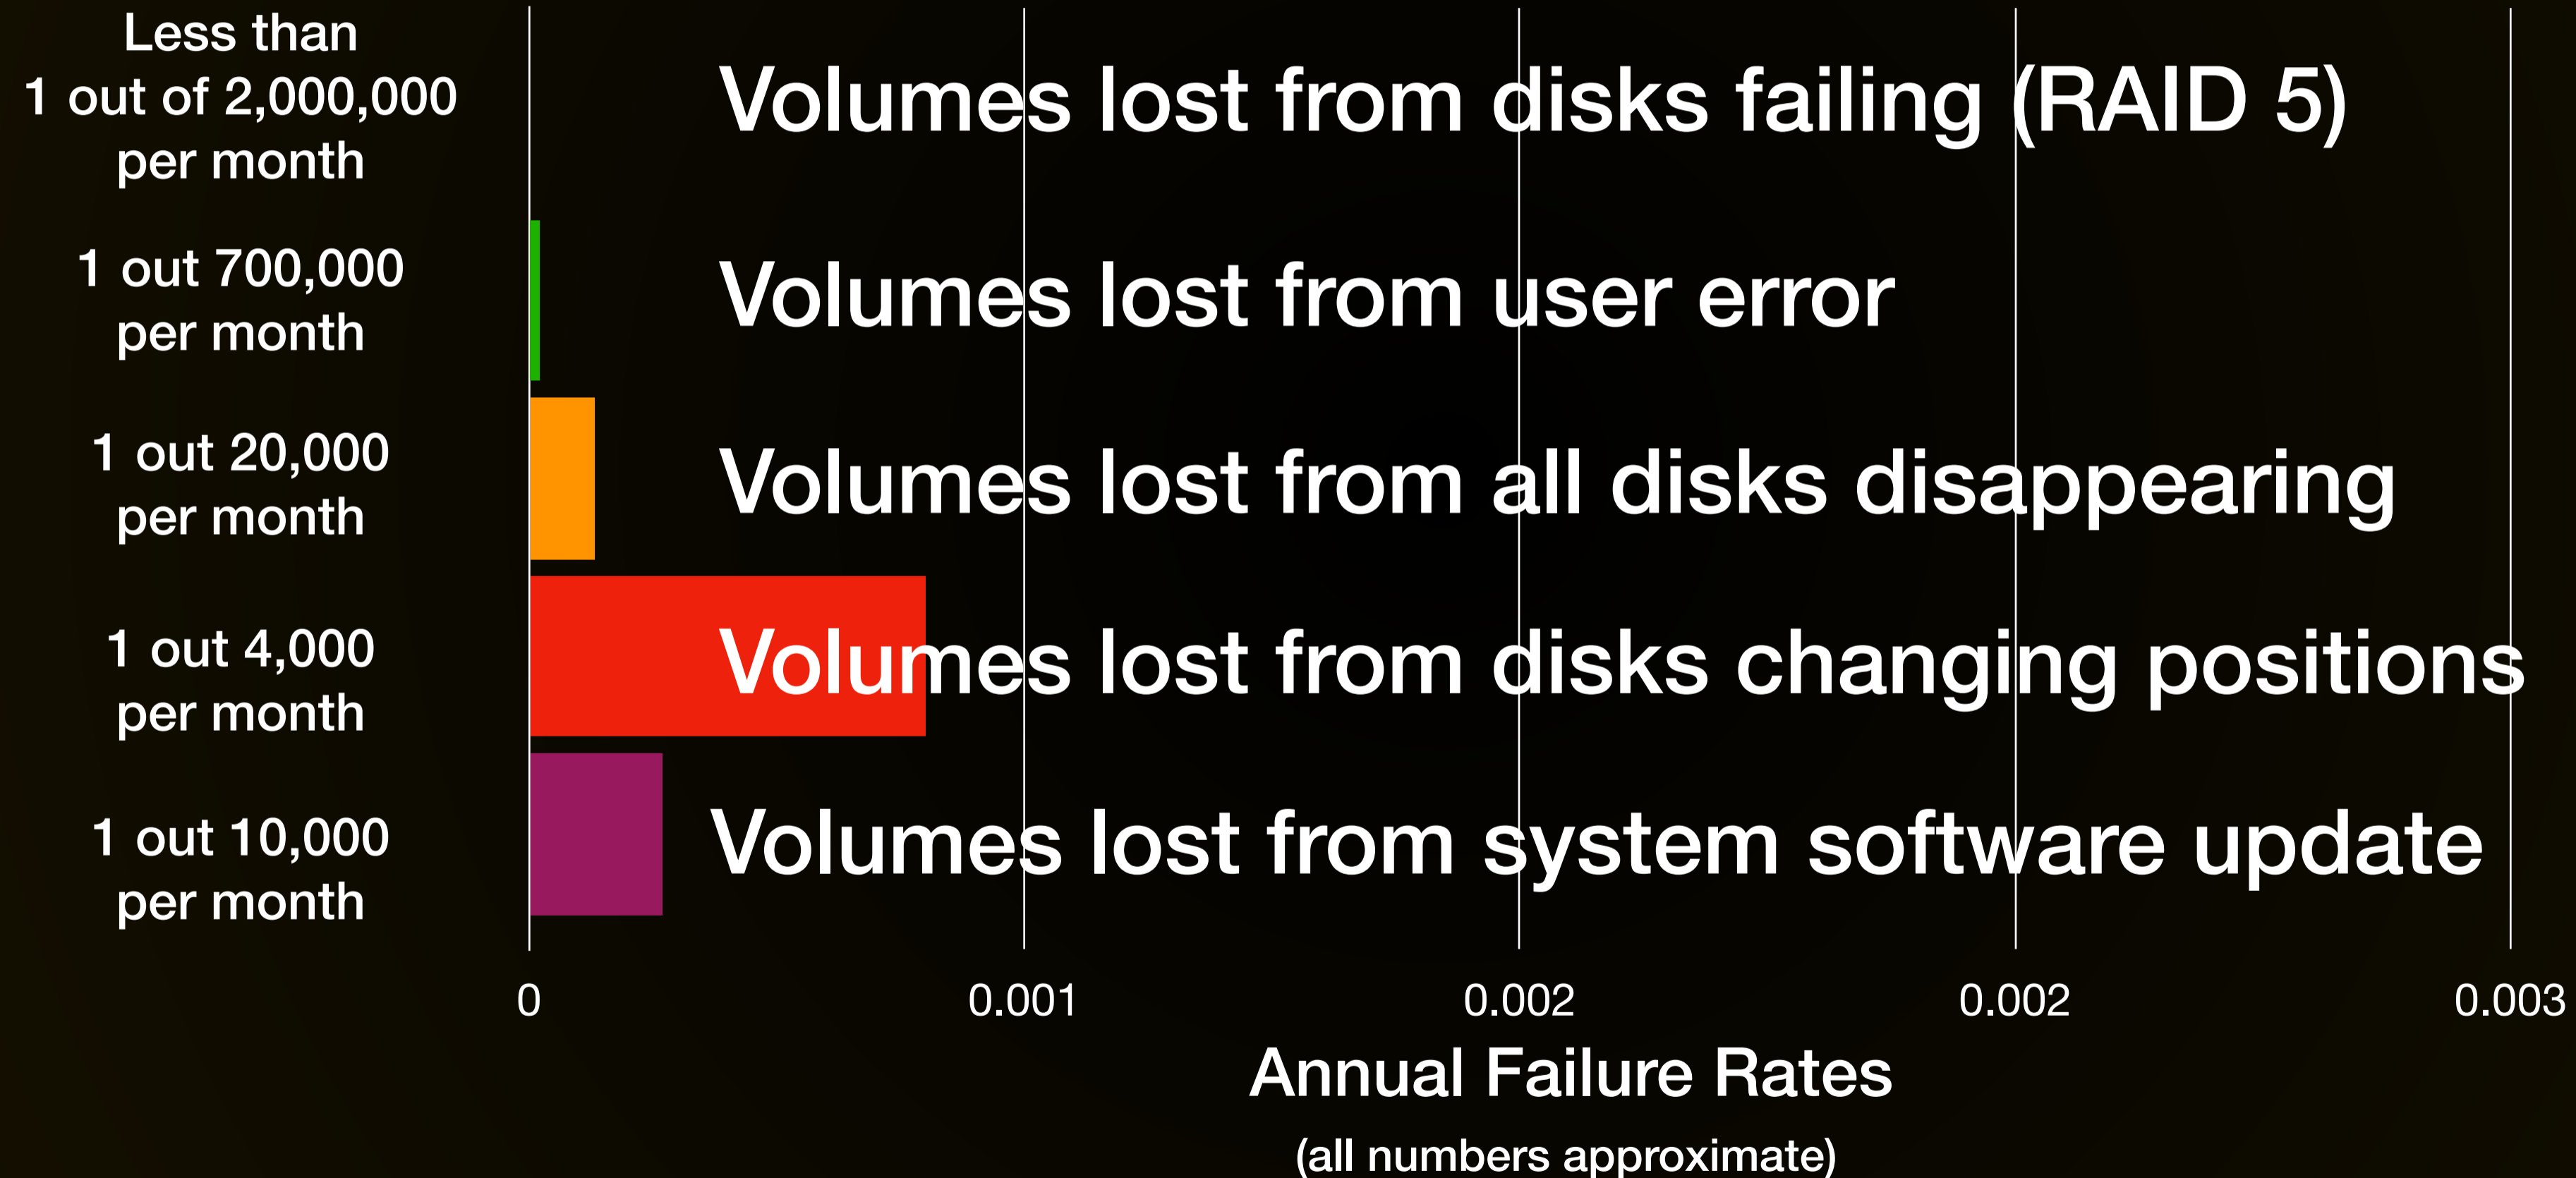
Annual Failure Rates

(all numbers approximate)



Volume Failure Rates

During the year 2017

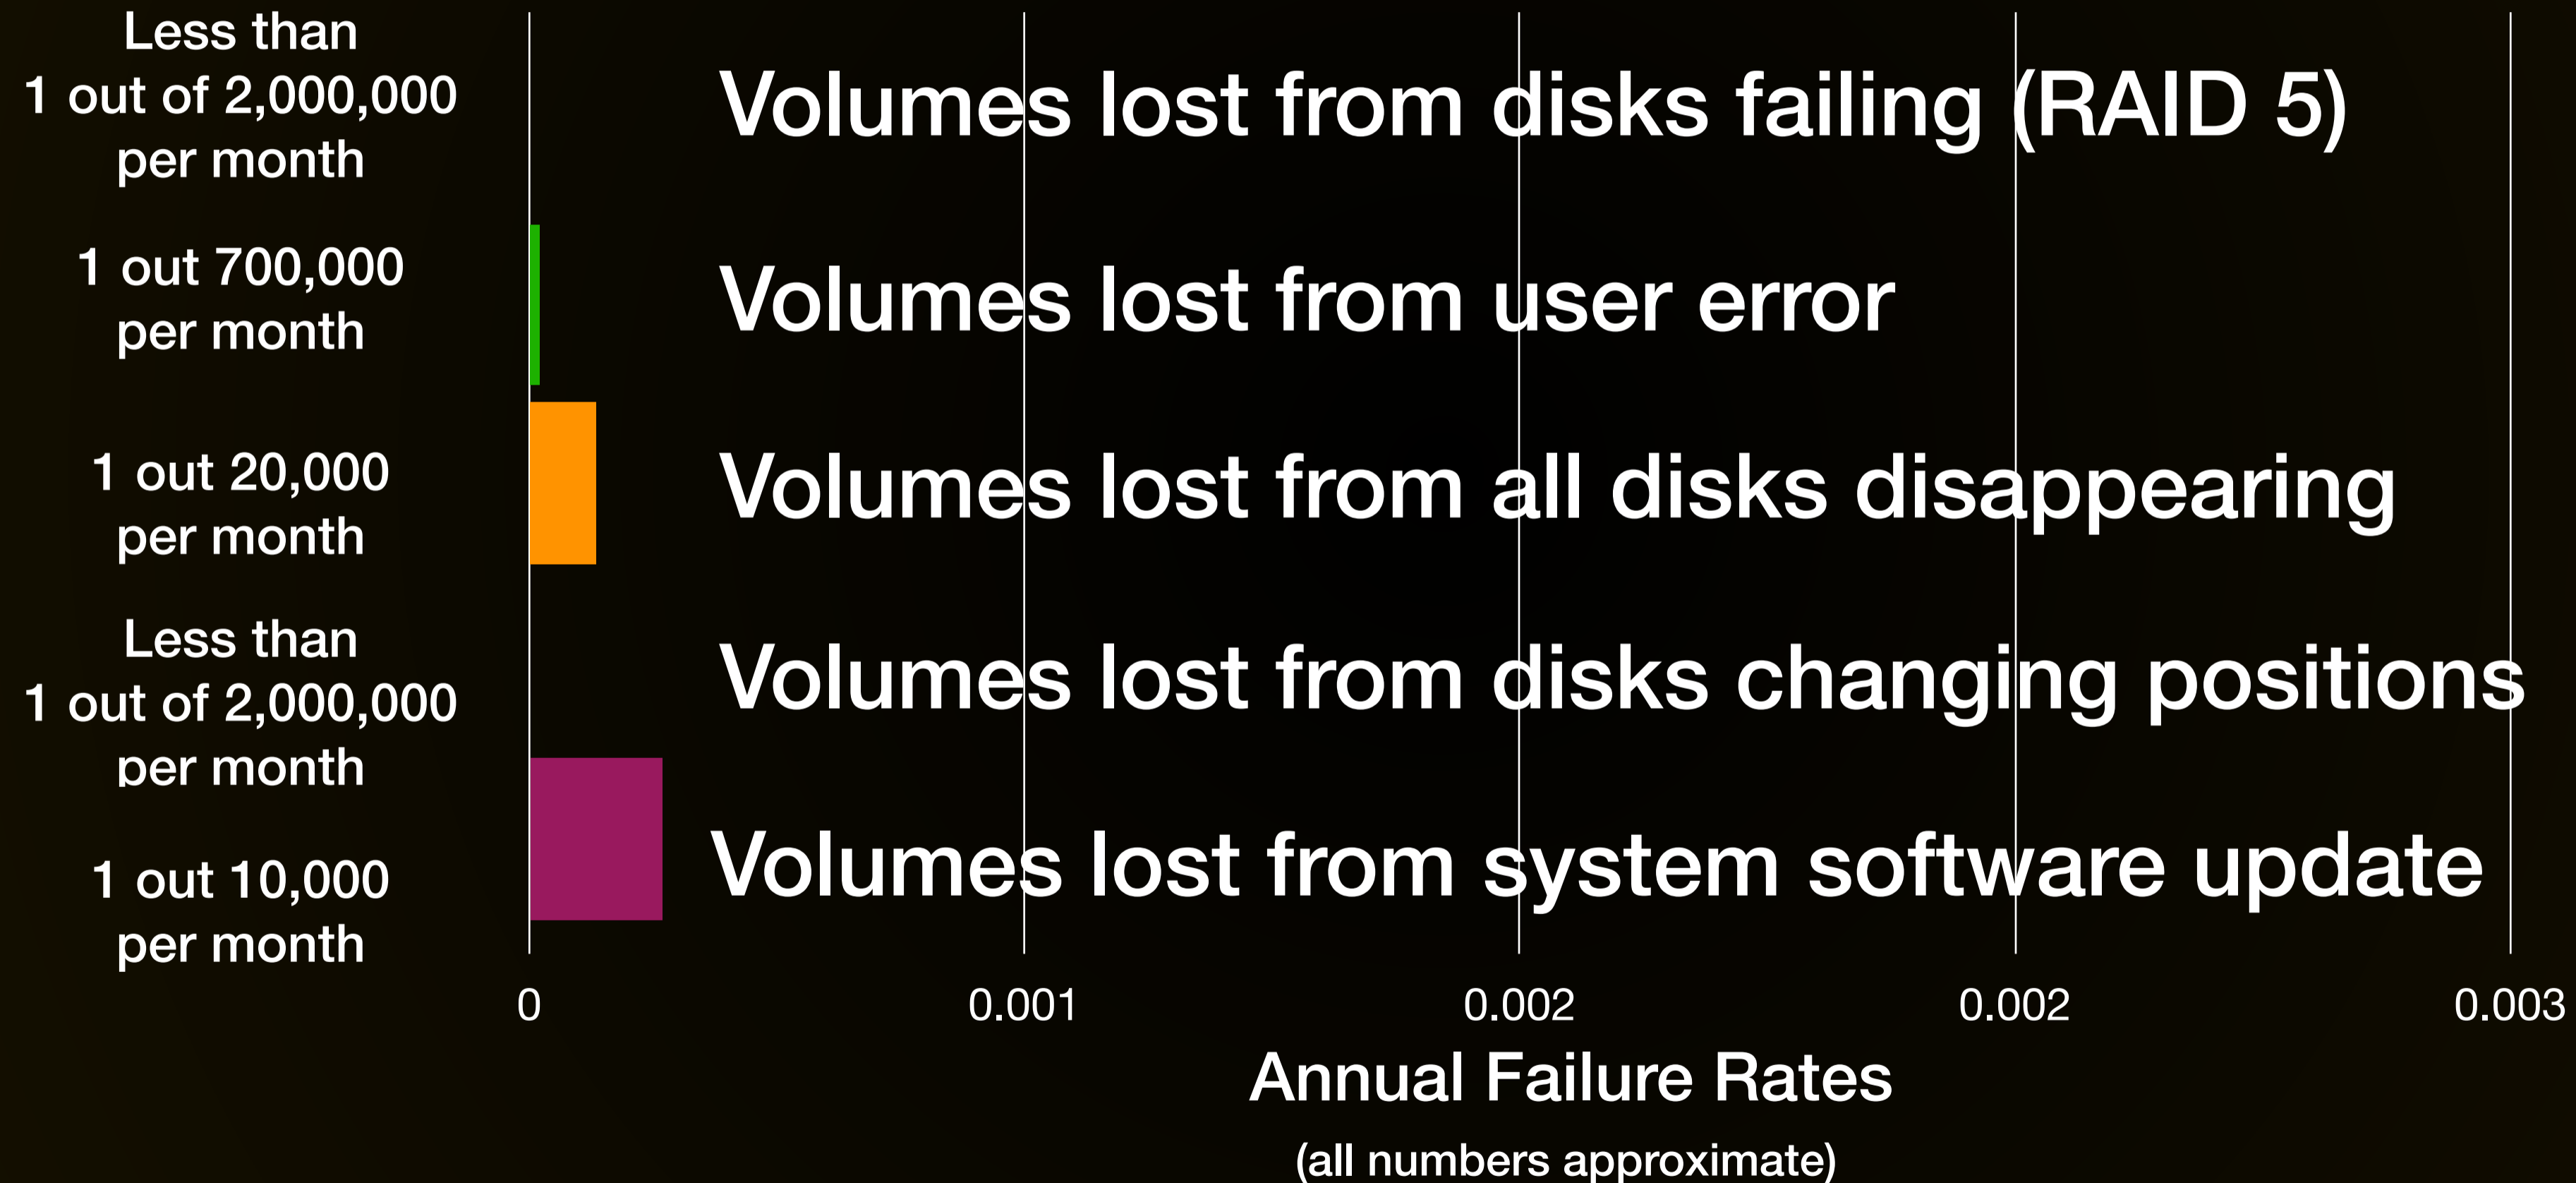


Mac OS 10.12.1 fixes kernel bug which causes all disks to disappear



Volume Failure Rates

During the year 2019

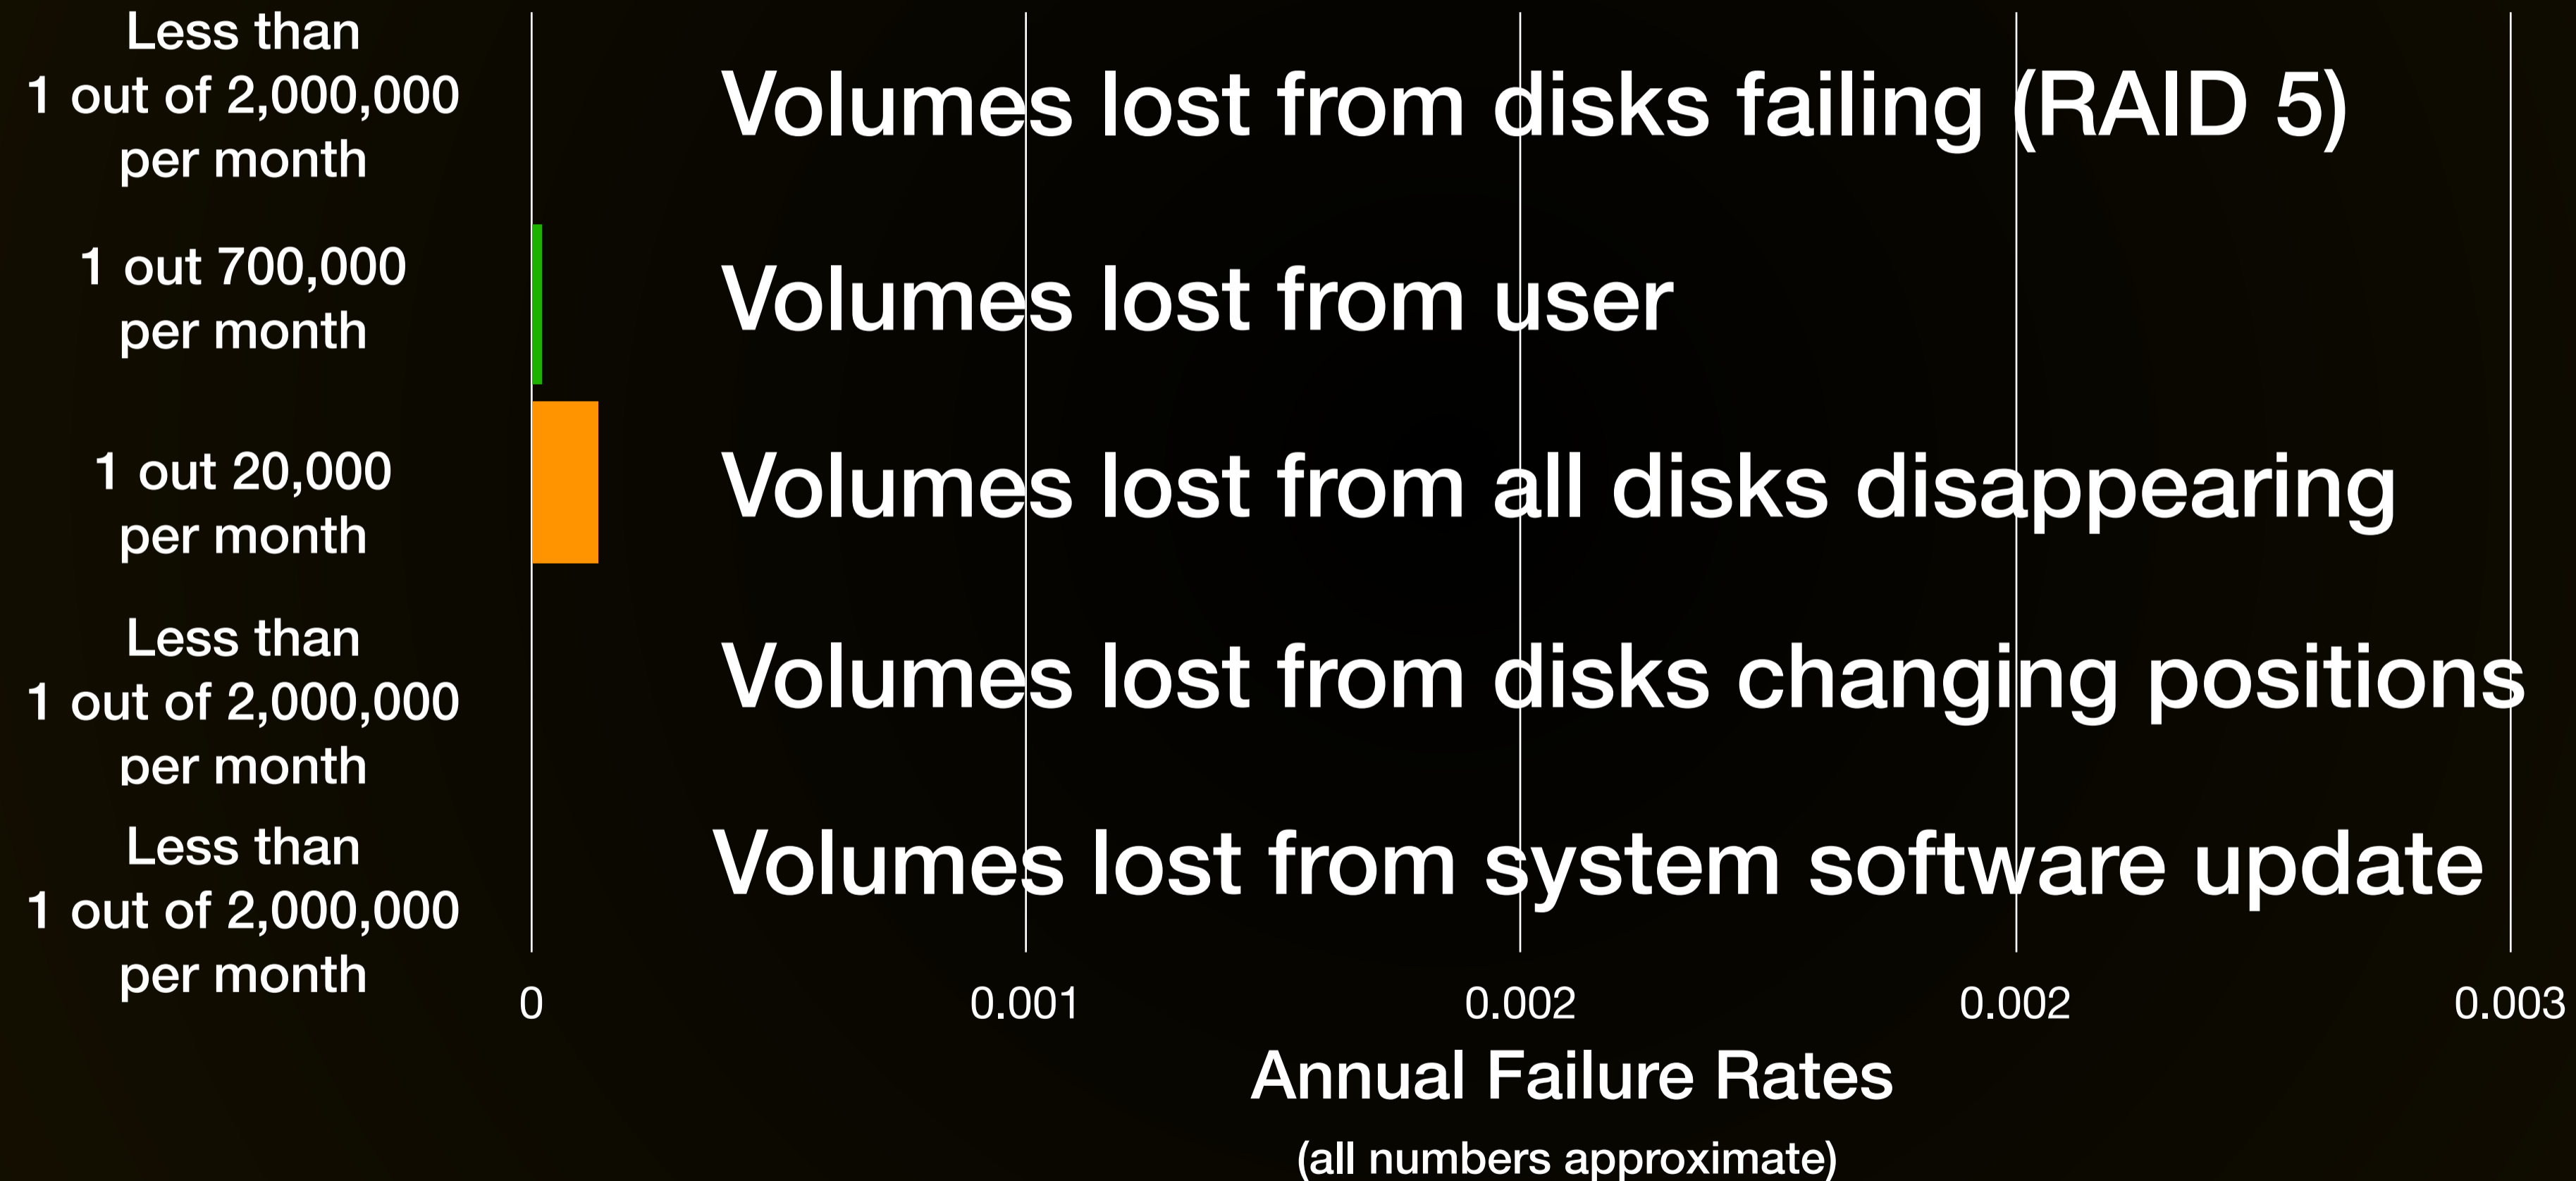


SoftRAID 5.7 detects disks changing position due to kernel bug



Volume Failure Rates

During the year 2019



macOS 12 fixes volume loss on system software update



An Example of Volume Loss



OWC Gemini
(Hardware RAID 0 - 2 HDDs - 28 TB total storage)



Customer's Mac: M1 MacBook Pro



Do you ever see this?



Disk Not Ejected Properly

Eject "Backup" before disconnecting or turning it off.

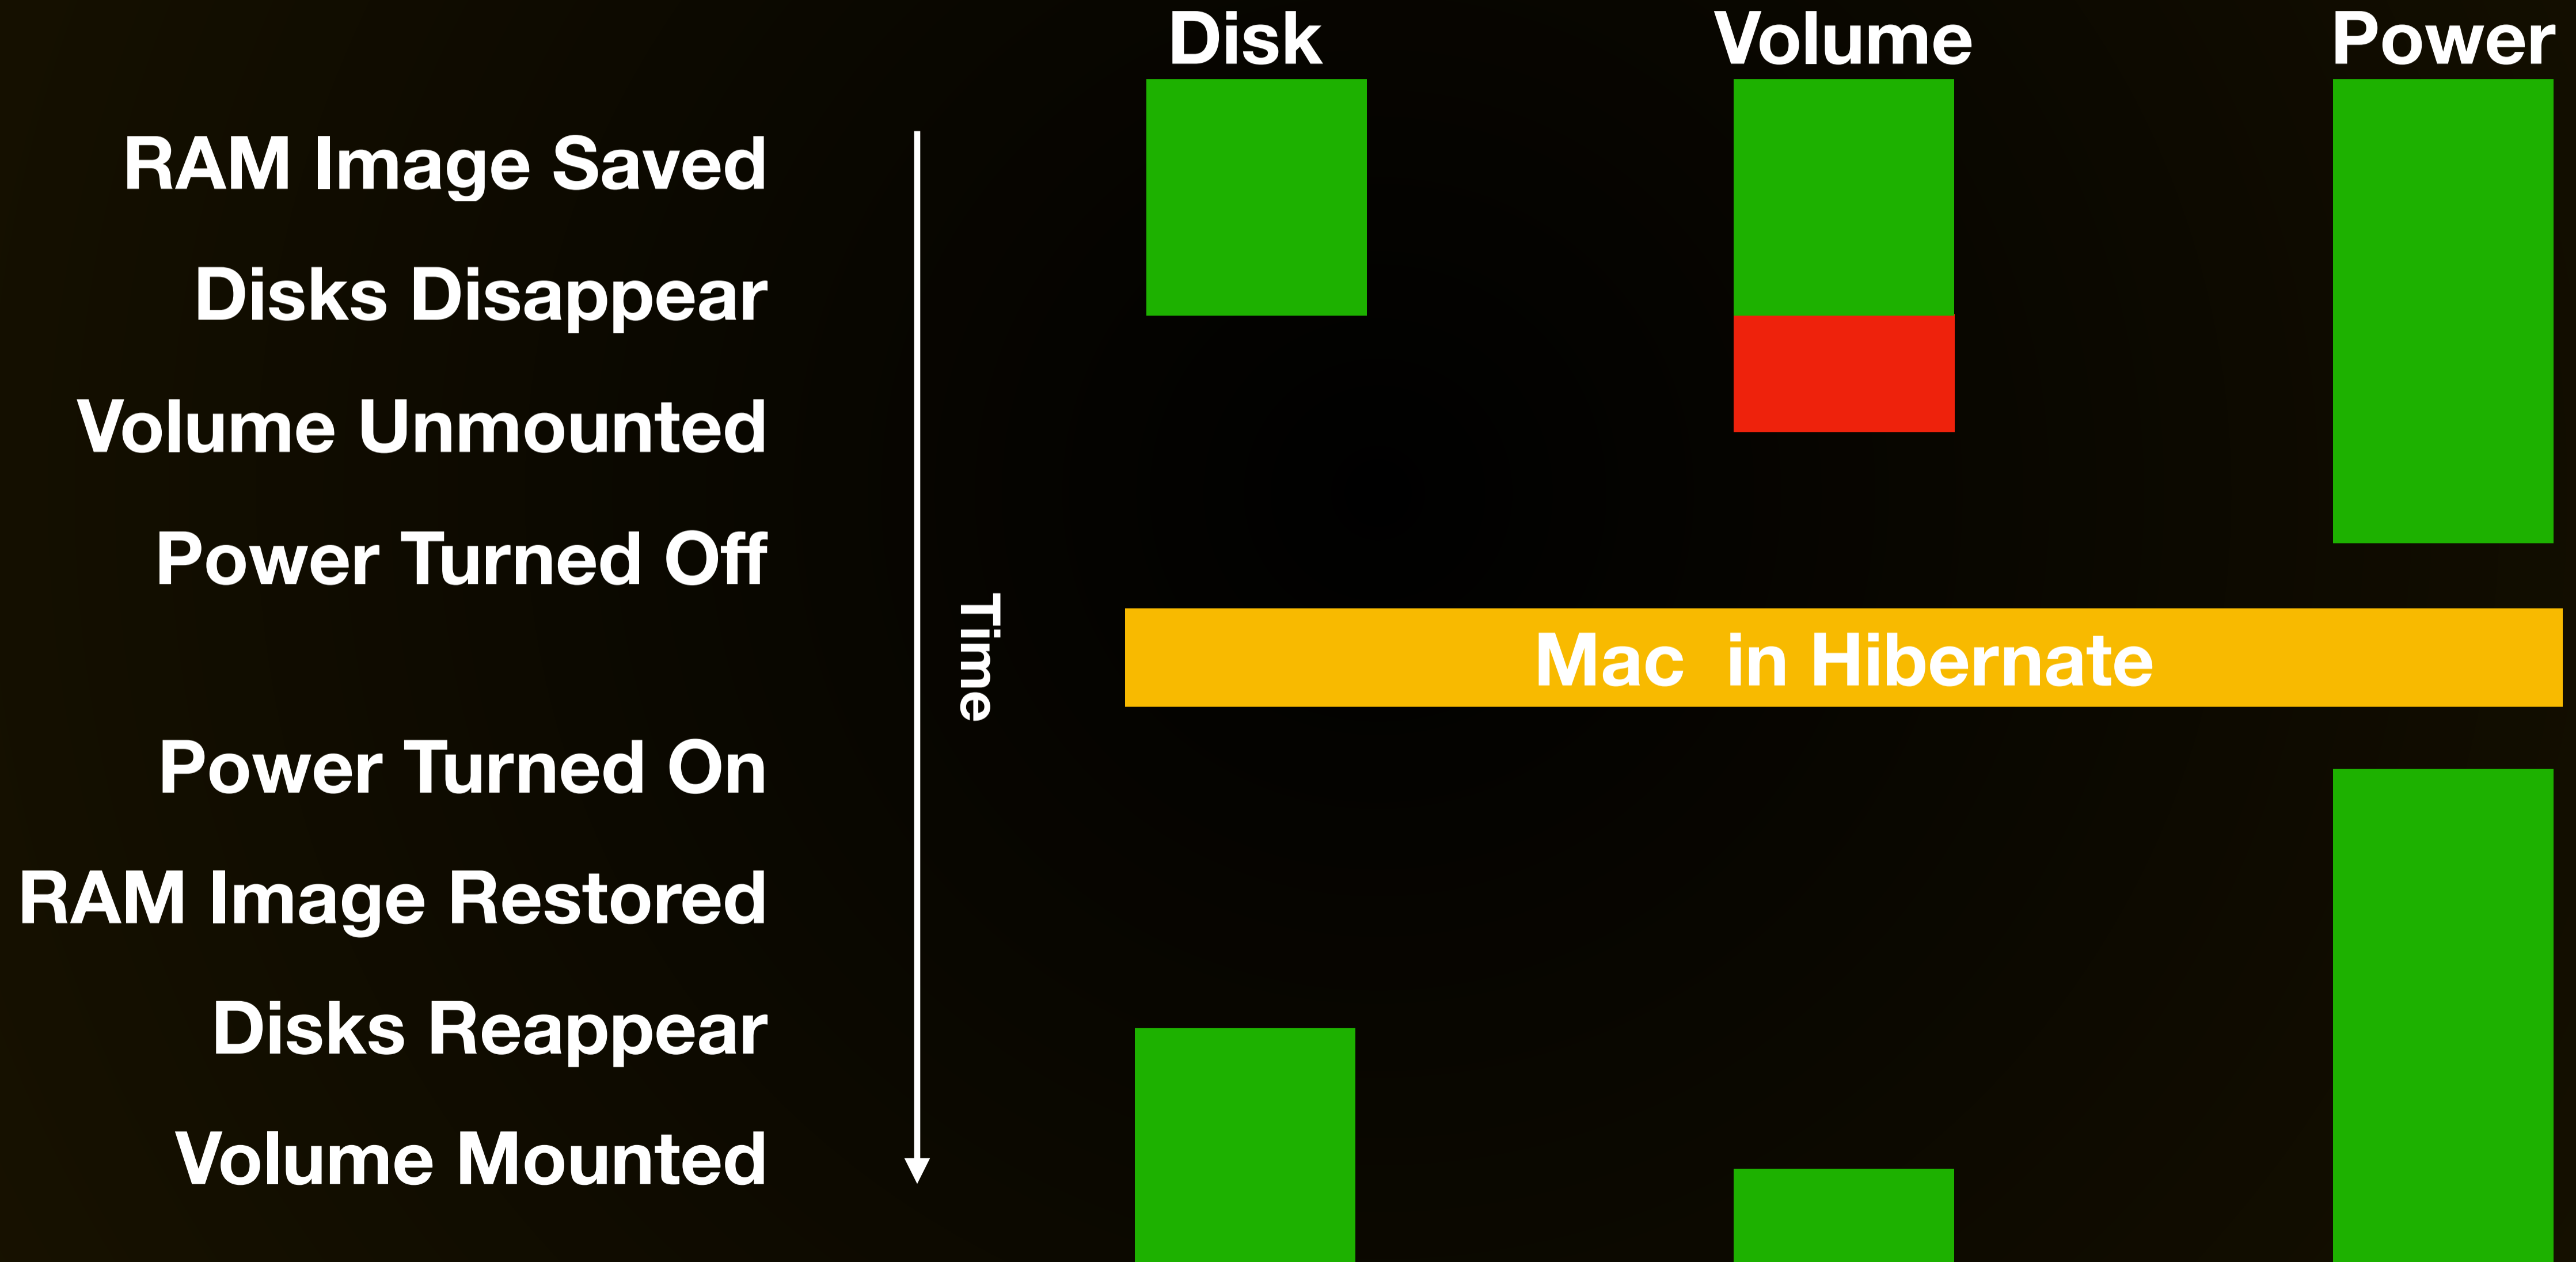


Why Hibernate?

- Laptop with no power source
- Prevents work loss when battery goes dead
- Unmount external volumes

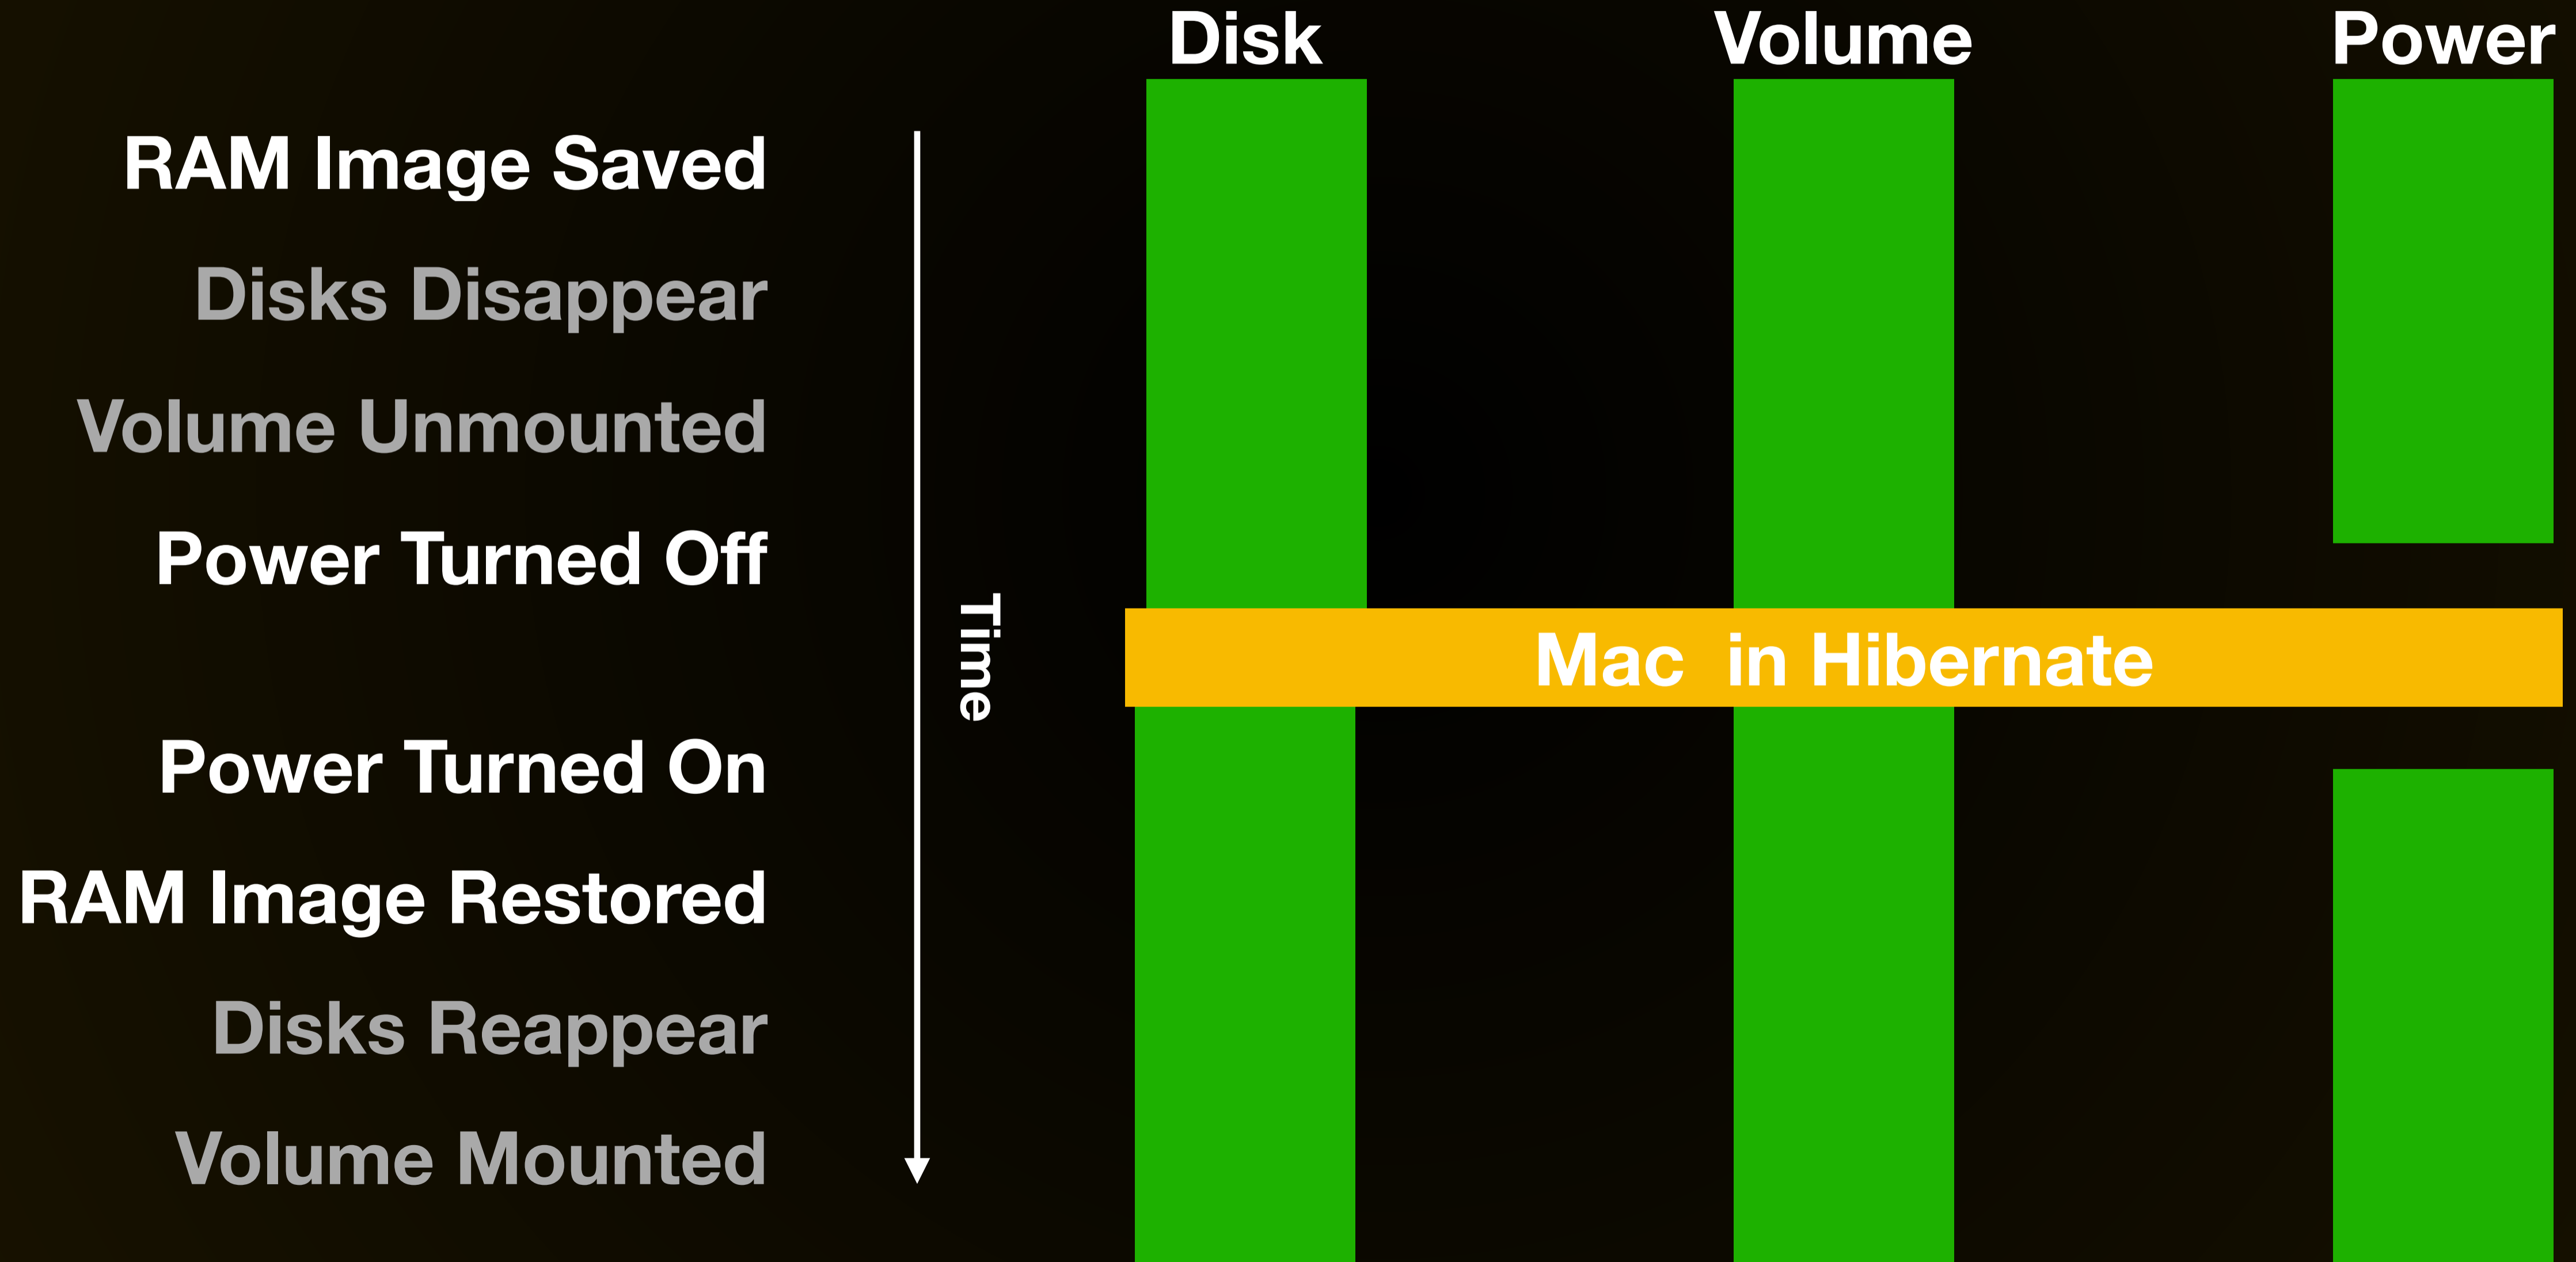


Hibernate on SATA Volumes





Hibernate on NVMe Volumes





Q&A