SNIA CMSI | COMPUTE, MEMORY, AND STORAGE

Is the Data Really Gone?

A Primer on the Sanitization of Storage Devices

A SNIA CMSI Webcast

Presented by:
Jonmichael Hands, Chia Network
John Geldman, KIOXIA
Jim Hatfield, Seagate

Today’s Speakers

Jonmichael Hands
Chia Network
Member,
SNIA Computational Storage
Special Interest Group

Jim Hatfield
Seagate Technology LLC
Member,
SNIA Security
Technical Work Group

John Geldman
KIOXIA
Member,
SNIA Board of Directors

SNIA Legal Notice

- The material contained in this presentation is copyrighted by the SNIA unless otherwise noted.
- Member companies and individual members may use this material in presentations and literature under the following conditions:
    - Any slide or slides used must be reproduced in their entirety without modification
    - The SNIA must be acknowledged as the source of any material used in the body of any document containing material from these presentations.
- This presentation is a project of the SNIA.
- Neither the author nor the presenter is an attorney and nothing in this presentation is intended to be, or should be, construed as legal advice or an opinion of counsel. If you need legal advice or a legal opinion please contact your attorney.
- The information presented herein represents the author's personal opinion and current understanding of the relevant issues involved. The author, the presenter, and the SNIA do not assume any responsibility or liability for damages arising out of any reliance on or use of this information.

  NO WARRANTIES, EXPRESS OR IMPLIED. USE AT YOUR OWN RISK.

What Does SNIA Do?

- SNIA is a non-profit global organization dedicated to developing standards and education programs to advance storage and information technology.

Industry Leading Organizations	Active Contributing Members	IT End Users & Storage Pros Worldwide
180	2,500	50,000

Who is CMSI?

- Part of SNIA, the SNIA Compute, Memory, and Storage Initiative is a community of storage professionals and technical experts who support:
  - The industry drive to combine processing with memory and storage,
  - The creation of new compute architectures and software to analyze and exploit the explosion of data creation over the next decade.
- CMSI's three Special Interest Groups – Computational Storage, Persistent Memory, and Solid State Drives – evangelize and educate on these technologies to the industry

www.snia.org/cmsi

Agenda

- Sustainability and media sanitization
- What is sanitization?
- Standards – IEEE P2883
- Verification of sanitization
- Sanitize in NVMe with some examples

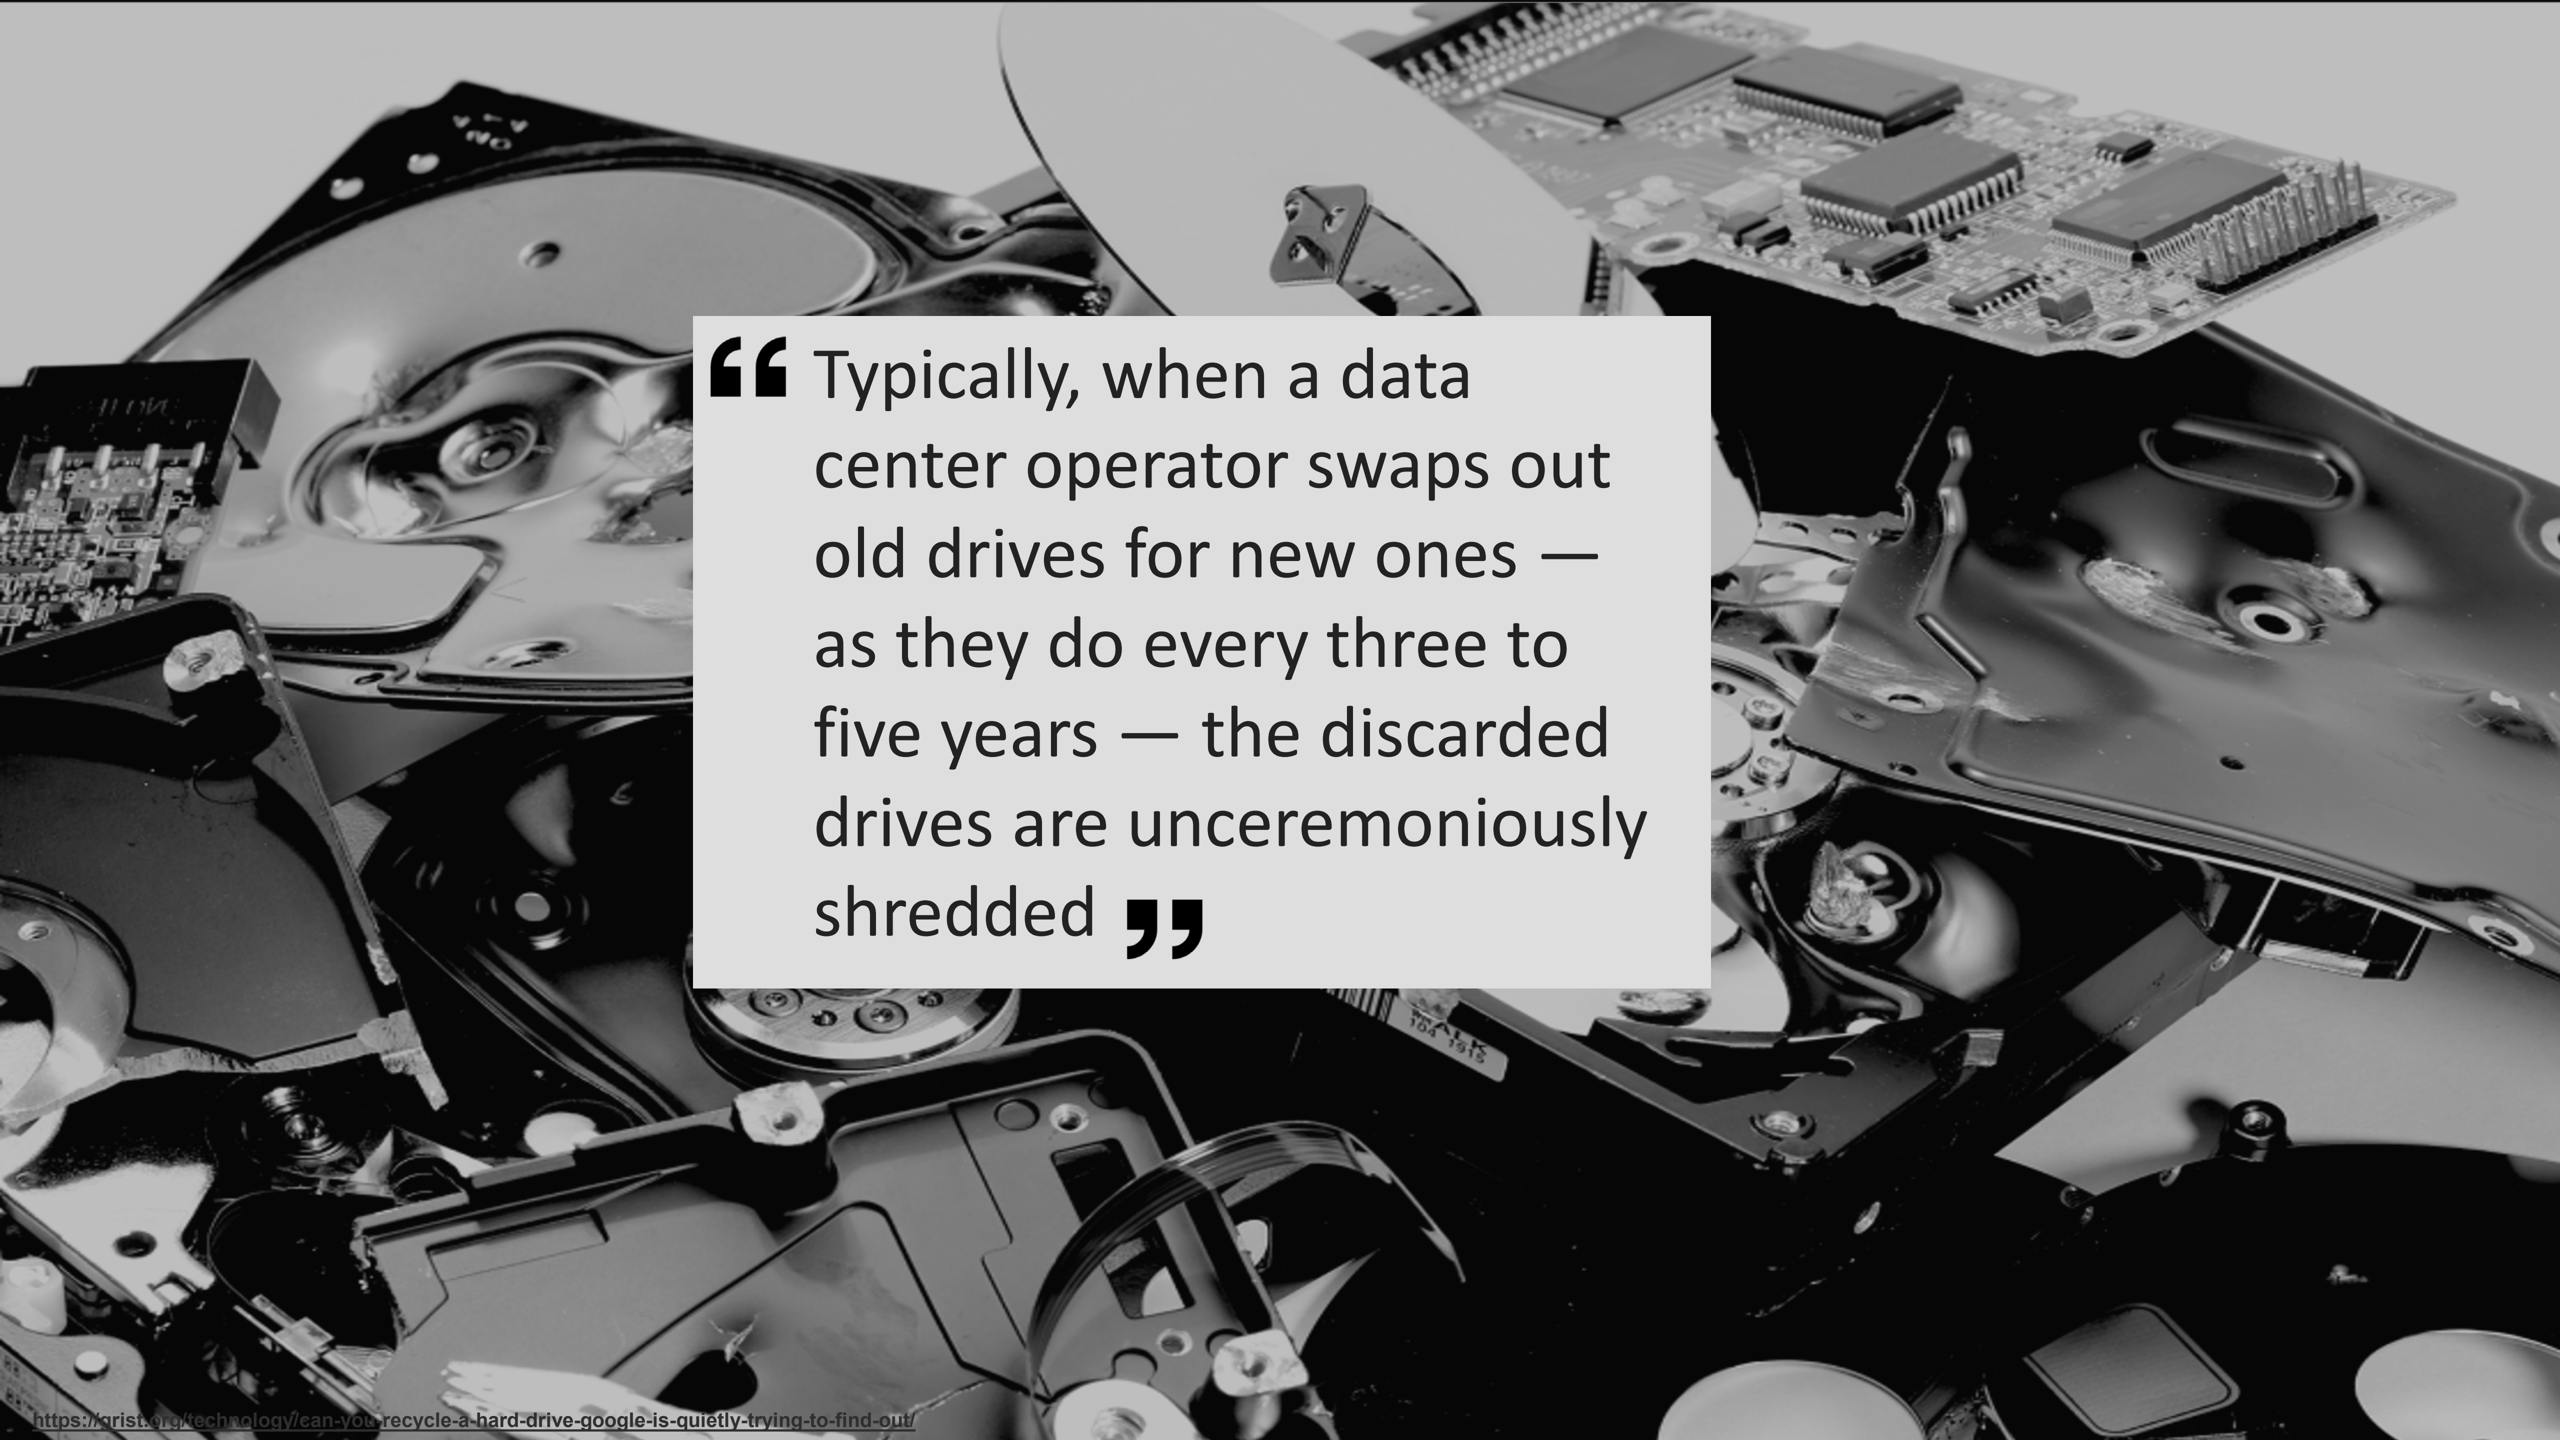

> “Typically, when a data center operator swaps out old drives for new ones — as they do every three to five years — the discarded drives are unceremoniously shredded”

https://grist.org/technology/can-you-recycle-a-hard-drive-google-is-quietly-trying-to-find-out/

Circular Economy for Storage

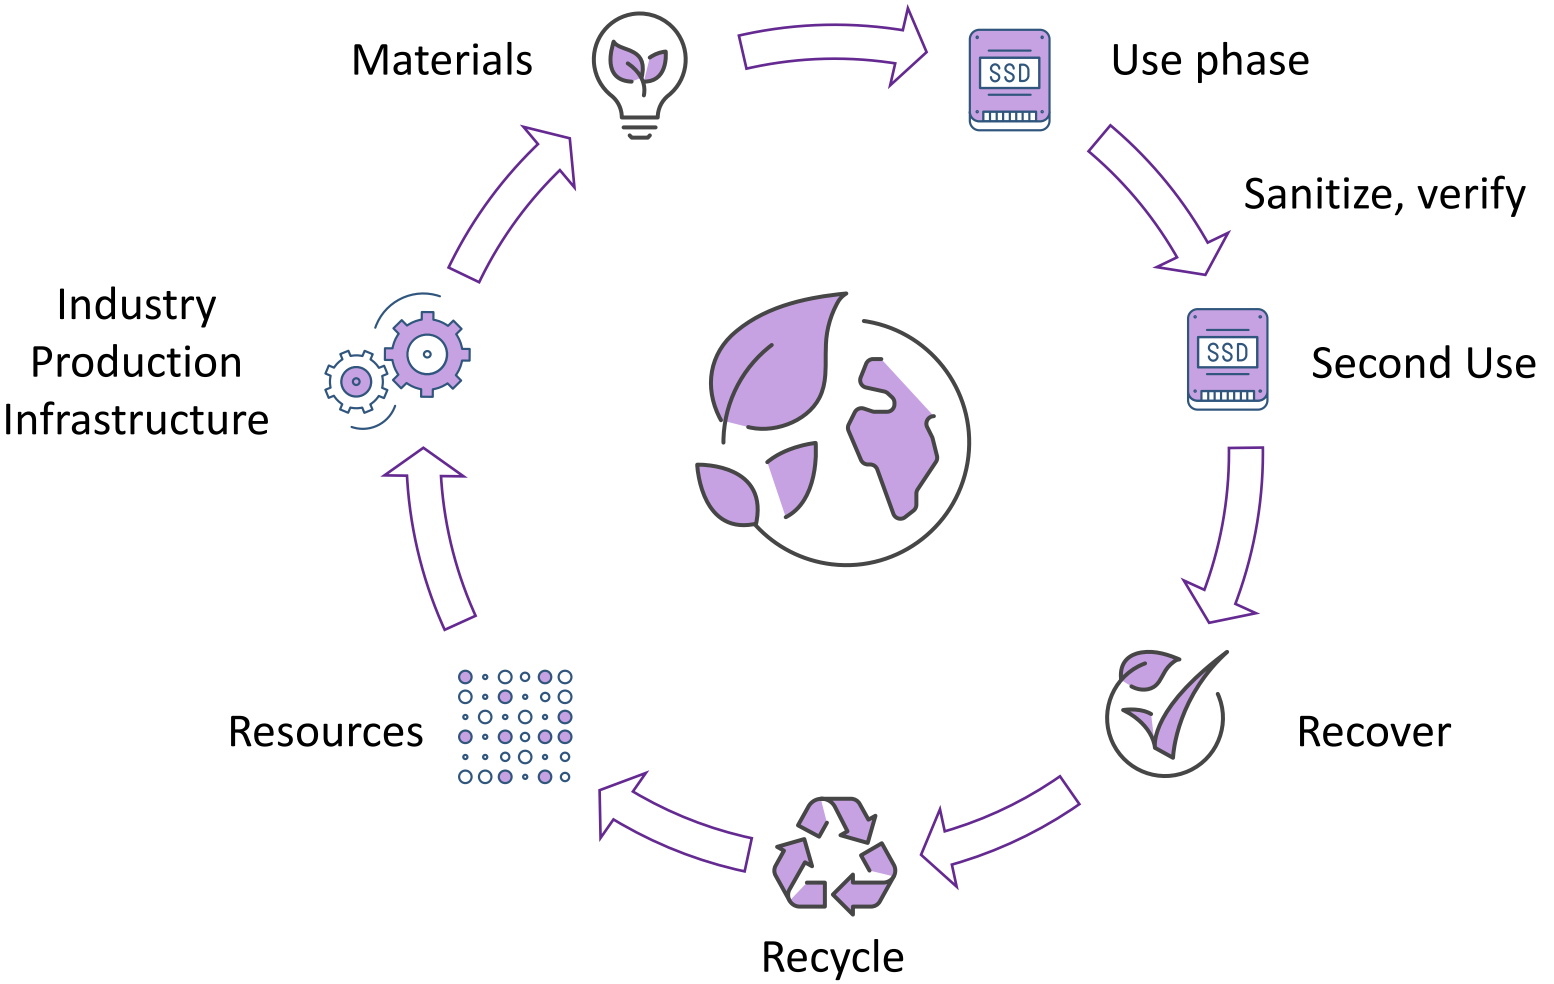

Sanitization is not…

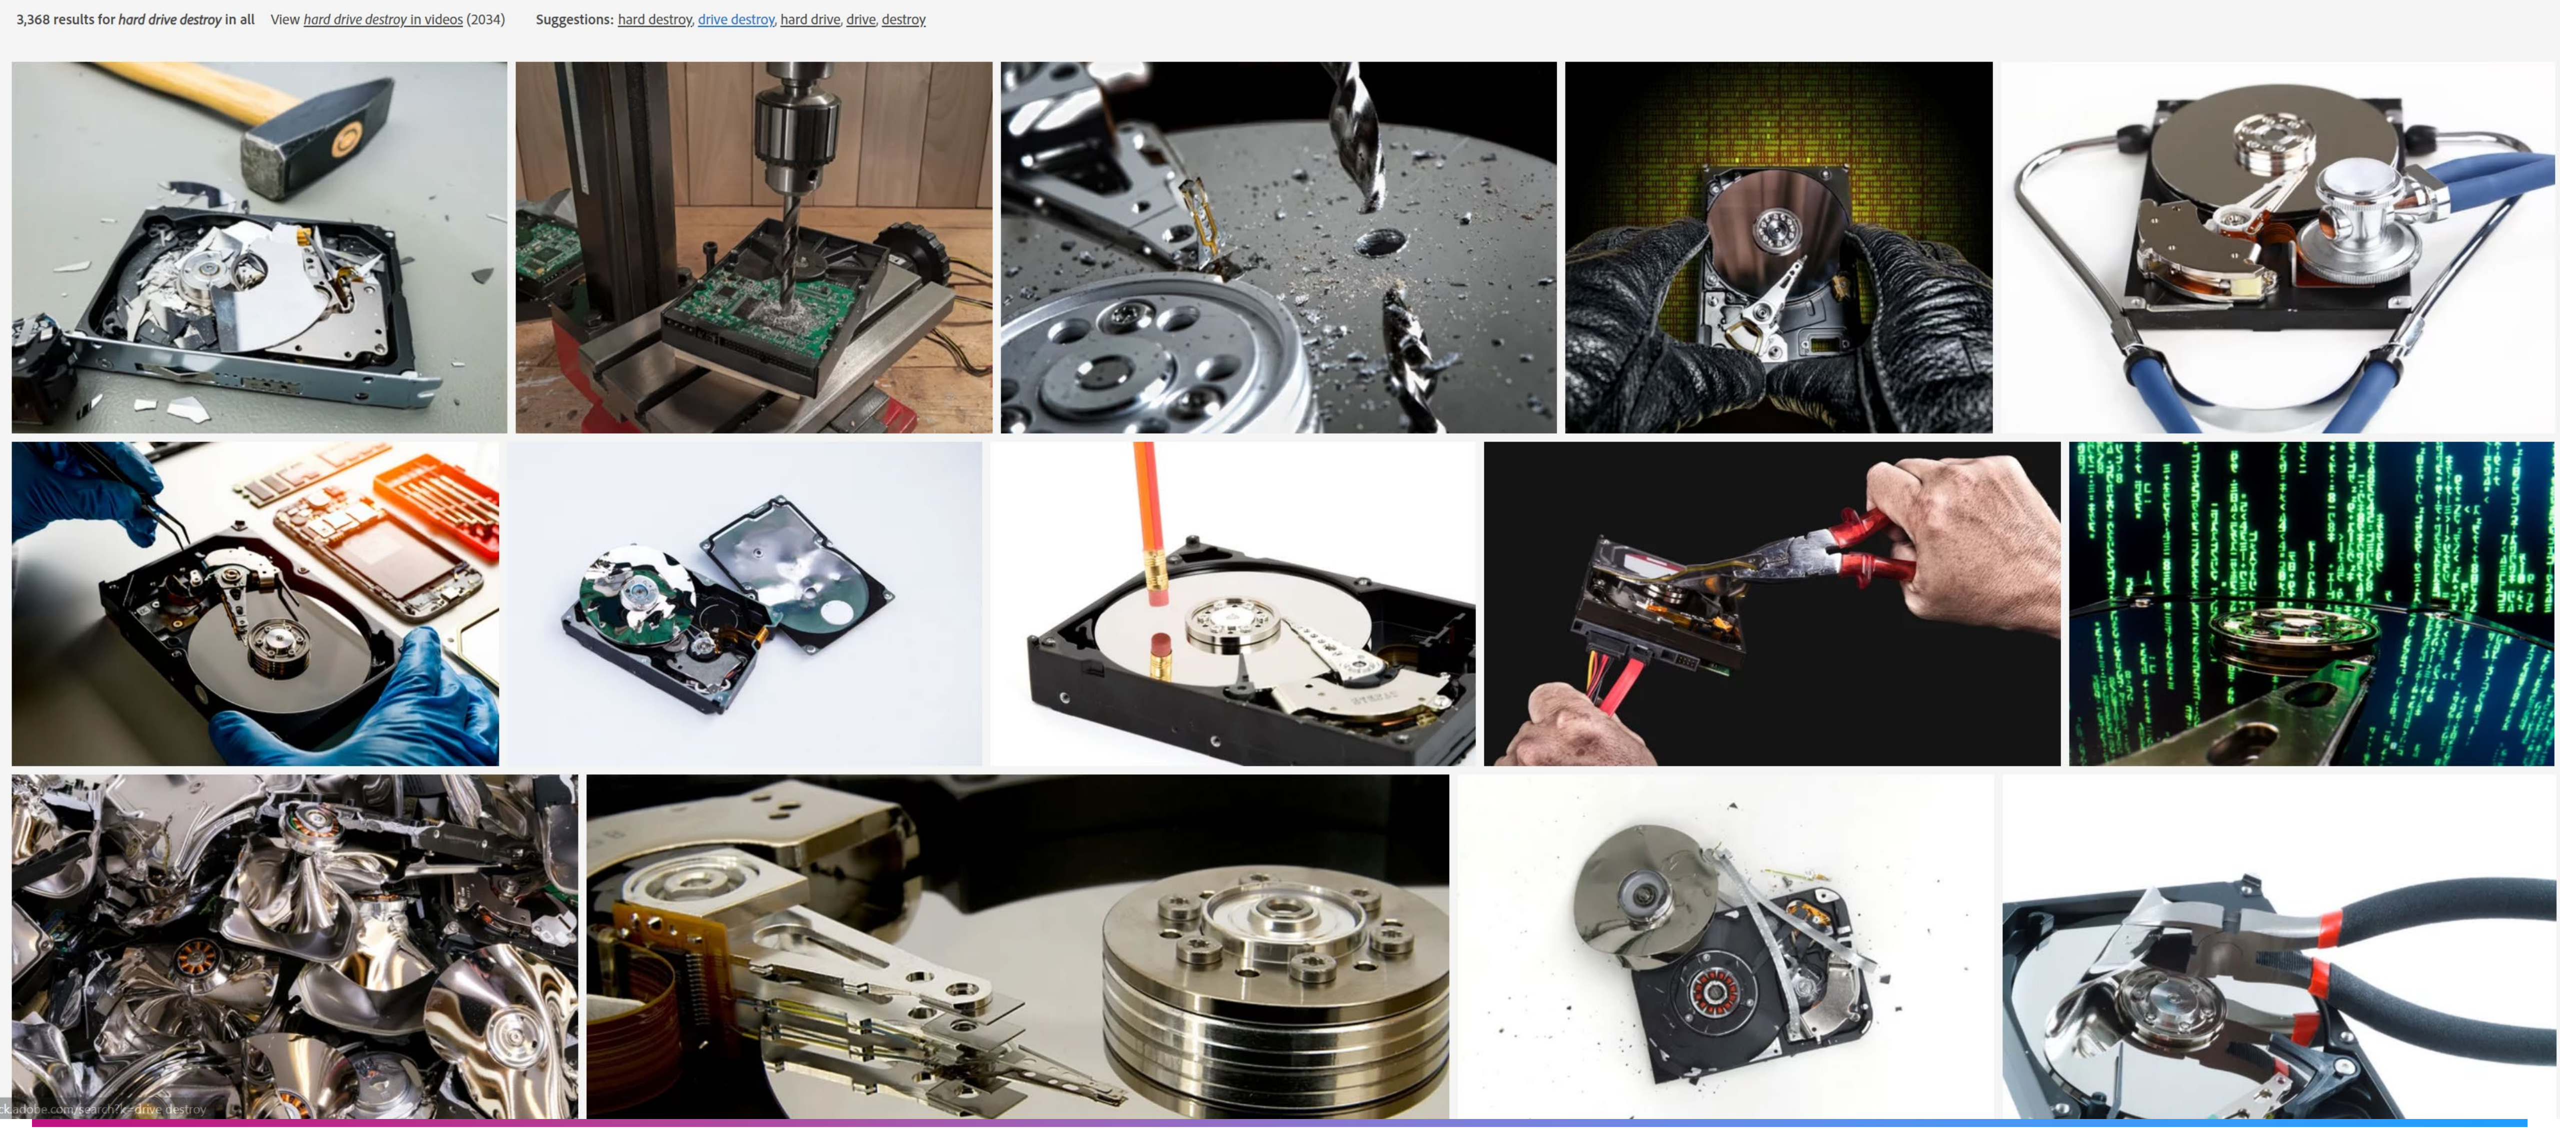

3,368 results for *hard drive destroy* in all   View *hard drive destroy* in videos (2034)   **Suggestions:** hard destroy, drive destroy, hard drive, drive, destroy

ck.adobe.com/search?k=drive destroy

What is Sanitization? And what is it not?

- Sanitization is NOT:
  - Related to FIPS or Common Criteria compliance
  - Simply deleting files
  - Trim/Unmap/Deallocation of storage
  - Fuzzy industry terms without a clear definition, and do not ensure the elimination of data
    - Secure data deletion, data clearing, data erasure, data destruction, data wiping, data overwriting, Data shredding
  - Blindly degaussing, shredding, pulverizing, disassembling a storage device
  - Scratching the media with a screwdriver (or) shooting it with a shotgun (these really do happen!)
  - Cutting a floppy disk with scissors (but it can be recovered by taping the pieces together)
  - Powering off a device
  - Exposure to radiation
  - And many more things….

What is Sanitization

A process or method to render access to target data on storage media infeasible for a given level of effort.

IEEE P2883 Draft Standard for Sanitizing Storage

- Sanitization - A process or method to render access to target data on storage media infeasible for a given level of effort.
- Defines Sanitization Methods and Techniques for specific media type (HDD, SSD, optical, removable, etc.)
- Specifies interface specific techniques (SATA, SAS, NVMe)
- Align industry on terminology and modern techniques for media sanitization
- Target all logical and physical locations for data – including user data, old data, metadata, overprovisioning , etc.

Defines purge method of sanitization that is secure, fast, and enables device reuse!

Sanitization Methods

Clear

Uses logical techniques to remove data on all addressable storage

Prevent against simple non-invasive data recovery

Format, deallocate

Purge

Uses logical or physical techniques to remove all data

Infeasible data recovery with state of the art techniques

Block erase, crypto erase, overwrite

Destruct

Infeasible data recovery with state of the art techniques

Leaves device in unusable state

Scope of Sanitization (Purge and Destruct)

All physical and logical locations that:

- currently contain user data
- used to contain user data (e.g., deallocated data, data reallocated because of media errors)
- could contain user data (e.g., overprovisioning, unused capacity, spare pools)
- are able to contain data that discloses information about user data (e.g., data that is useable to direct forensic analysis)

State of the art laboratory techniques

This includes such things as

- Disassembly, and mounting a different circuit board to an HDD spindle
- Reading raw signal from an HDD platter on a spin stand
- Electron microscopy
- X-ray probing
- And many more things that a well funded adversary or a nation state has at its disposal

The sanitization process

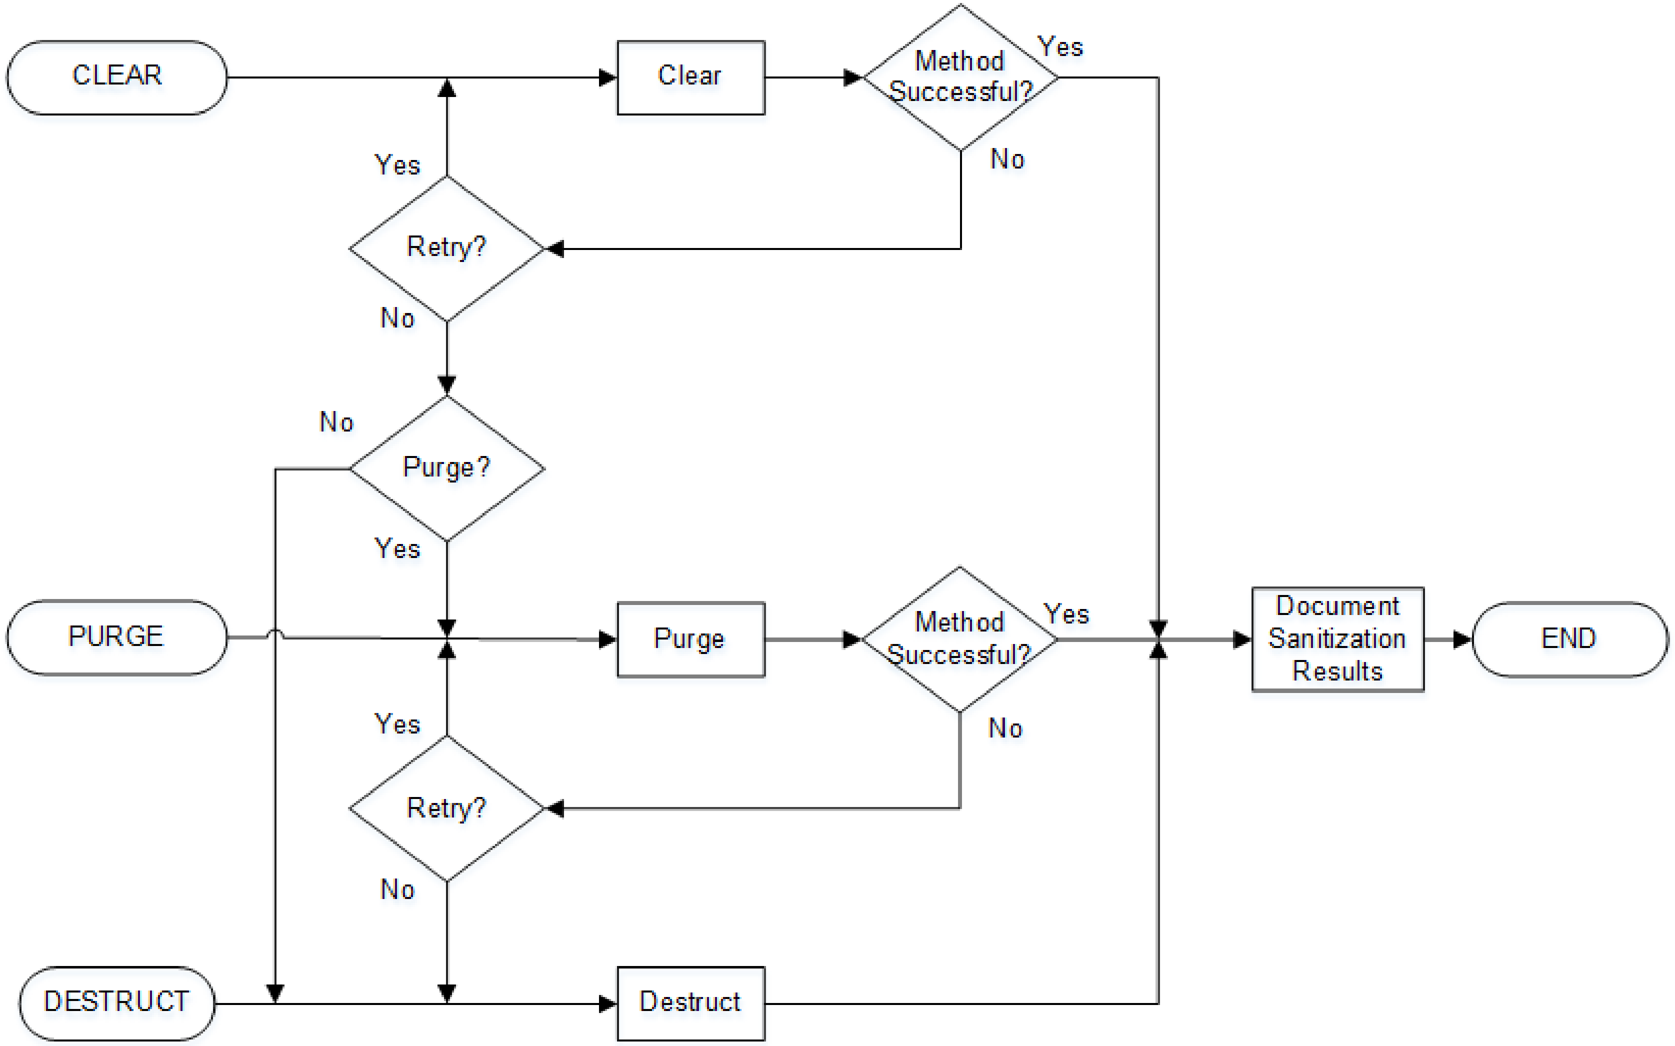

Sanitization Techniques

Factors Affecting the Ability to Sanitize

- The storage media is not identifiable.
  - For example, tape cartridges are usually are labeled with the technology and generation, but some may not be labeled.
  - A storage device may be in an enclosure that is sealed (e.g., laptop, USB enclosure, mobile device)
  - A storage device may be embedded in a chip that is not accessible
- The organization lacks the expertise to sanitize the storage media (while leaving it usable) or to verify that sanitization was successful.
- The equipment is not working or is anticipated to not be working soon.
- The equipment or software needed to perform the operations is not available.
  - Examples include a storage device to access removable storage media, an interface for the storage device, a degausser with sufficient strength to erase newer magnetic storage media, etc.

Sanitization Techniques - Overwrite

- The term ‘overwrite’ has multiple meanings. A distinction has been made between

  Simple overwrite (clear): writing a pattern or deallocating only logical locations

    - Write commands
    - SECURITY ERASE UNIT (ATA), FORMAT UNIT (SCSI), Format NVM (NVMe)
    - UNMAP (SCSI), DATASET MANAGEMENT (ATA), Dataset Management (NVMe)
    - Note: there could be physical copies of data that are not erased

  Sanitize overwrite (purge): writing a pattern to all logical and physical locations within the scope of sanitize

    - Note: SCSI, ATA, and NVMe all have a special ‘Sanitize’ command that accomplishes this

- Simple overwrite is not appropriate for
  - Non-magnetic media (paper, optical media, etc.)
  - NAND flash: overwriting negatively affects write amplification and the endurance of the device

Sanitization Techniques – Block Erase

- Block Erase
  - Allows a relatively large region of storage (e.g., an erase block) to be erased in a single operation
  - Is not appropriate for magnetic media
  - Is useful for types of memory devices (e.g., NAND) that support it
  - The media may or may not be readable without errors after block erasure (because CRCs also be erased and may not recomputed)
  - The media may be all binary 0's or all binary 1's after block erasure (depending on the media vendor)
  - Reduces the negative impact on write amplification that the Overwrite technique has

Sanitization Techniques – Cryptographic Erase

- Cryptographic Erase
  - "Method of sanitization in which the encryption key for the encrypted target data is sanitized, making recovery of the decrypted target data infeasible using state of the art laboratory techniques"
  - media based cryptographic erase: Method of cryptographic erase in which the encryption key is only resident on the storage device.
  - Without the encryption key used to encrypt the target data, the data are unrecoverable
  - ISO/IEC 27040 pre-conditions for cryptographic erase:
    - encryption of all data intended for cryptographic erase prior to recording on the storage;
    - the strength of the cryptographic algorithm (including mode of operation) used to encrypt the target data is at least 128 bits;
    - the level of entropy of the encryption key used to encrypt the target data is at least 128 bits; and
    - all copies of the encryption keys used to encrypt the target data are sanitized; if the target data's encryption keys are, themselves, encrypted with one or more wrapping keys, it is acceptable to perform cryptographic erase by sanitizing a corresponding wrapping key.
  - Only media-based storage sanitization is supported in IEEE 2883

Sanitization Techniques – Cryptographic Erase

- The level of effort needed to decrypt this data without the encryption key is the lesser of:
  - the strength of the cryptographic algorithm used to encrypt the data (including mode of operation);
  - the level of entropy of the target data's encryption.
- Sanitization may be performed with high assurance much faster than with other sanitization techniques. Cryptographic erase can be executed in seconds.
- Cryptographic erase can also be used as a supplement or in addition to other sanitization approaches.
- Some organizations perform an additional, but unneeded, sanitization using a clear method to reduce the attack surface by preventing access to the ciphertext.

Sanitization Techniques - Destruct

- Melting
  - “Destruct by changing storage media from a solid to a liquid state, generally by the application of heat”
- Incineration
  - “Destruct by burning a storage device completely”
- Both of these techniques are not ‘green’
  - environmental risks associated with disposing of potentially hazardous materials (e.g., plastics, lead, heavy metals)
  - no possibility of recovering valuable materials (e.g., gold, rare earth elements)
  - require large amounts of energy to perform

Sanitization Techniques - Destruct

- Shred
  - “An obsolete form of Destruct that cuts or tears a storage device or storage media into small particles”
- Pulverize
  - “An obsolete form of Destruct that grinds a storage device to a powder or appropriately small particles”
- With the increased density of data in all types of media, shredding and pulverizing can leave significant amounts of information on the remaining particles.

Verification of Sanitization

Oh what a tangled web we weave when at first we start to store

Let’s set expectations: Proof of Sanitize correctness is very difficult

- Issues:
  - The physical interface does not provide observability to all sanitized functionality
  - Many devices have internal integrity checks that prevent bad data from leaving the SSD
  - To avoid wasted write cycles, many devices deallocate all storage before leaving the sanitize operation
  - Use of media for non-data purposes (firmware, statistics, vendor specific device ‘state’)
- Let’s call correct implementation testing: Validation
  - This needs to be proven with third parties that can both hold design reviews and disassemble sacrificial devices
- Let’s call practical device testing: Verification
  - This is what responsible storage providers can and should perform

Let’s set expectations:
Proof of Sanitize correctness is very difficult

- The physical interface does not provide observability to all sanitized functionality
  - Overprovisioned areas are not accessible
    - Overprovisioning is extra storage used to increase endurance and improve performance
  - Caches and buffers in the data path are not accessible
- Many devices have internal integrity checks that prevent bad data from leaving the SSD
  - Crypto Erase and Block Erase sanitize operations affect
    - Zero/one balancing encoding
    - Integrity data structures
  - Overwrite sanitize operations include the normal encoding and creation of the integrity data structures

Let’s set expectations:
Proof of Sanitize correctness is very difficult

- To avoid wasted write cycles, many devices deallocate all storage before leaving the sanitize operation
  - This avoids a writing the entire device’s data in way that creates fake data that needs to be erased before storing user data (wasted endurance)
- Internal complexity:
  - Not all media storage is user data: firmware, logs, non-user information are not sanitized
  - Media does not contain simple readable text, data is typically binary and encoded for reliability
  - Without firmware context, it is difficult to identify what is user data and what is device data

NVM Express has some tools

- In NVM Express Base Revision 1.3, the no-deallocation bit was added to the Sanitize command
  - Support for this bit was very difficult and expensive
  - Requires an additional full write of observable space which has both time and endurance costs
- In NVM Express Base Revision 2.0, the support for that additional full write is reportable, and vendors can report whether or not the cost of such operations is supported
  - A device can report if it does not support no-deallocation
    - The wonderfully named: No-Deallocation Inhibit bit

All is not lost: Sanitize Verification

- Sanitize verification determines the adequacy or effectiveness of the storage sanitization
- For clear or purge, verification is a check of the *sanitization* outcomes
  - Much (not all) of the results of sanitization can be checked
- For destruct, physical inspection is used to check the *sanitization* outcomes

Verification 101

- Documentation: Keeping records by storage vendors of what was done (who, when, what) may help prove responsible practices
- For clear methods, if data verification is required, then representative sampling should be performed to evaluate the remaining data
- For purge methods, if data verification is required, then full verification should be performed to evaluate the remaining data
  - On the other hand, for cryptographic erase, as there is no predictable data, it may not be possible to perform verification of expected results
- For destruct, physical inspection is the only option as the device has been made unusable

Result Verification methods

- Representative sampling:
  - Select random locations totaling at least 5% of the addressable space
  - Select locations dispersed over the entire addressable space
    - For example, divide the addressable space into equal size areas and select more than one random block within each area
- Full verification - full reading of all areas to be verified and comparing the read data to the expected value
- Caveats:
  - There is no expected value for cryptographic erase, use simple checks for the lack of expected (pre-sanitize) values
  - Crypto erase and block erase methods may destroy internal encoding and integrity data, resulting in de-allocated data pattern returns (rather than unformatted raw media)

Example: NVMe Sanitize

NVMe Sanitize

- Make user data unrecoverable. All user data in NVM, PMR, CMB, cache, metadata, unallocated, or overprovisioned space
- Background operation with log and status
- Types of sanitize supported, read identify controller
  - Block erase – low level media specific block erase (e.g. NAND erase block)
  - Crypto erase – change media encryption key
  - Overwrite- overwrite with a fixed pattern
- Sanitize log page for status and estimated times for each method
- Send async notification upon completion

NVMe Sanitize

- Very well documented in NVMe specification, download at nvmexpress.org!

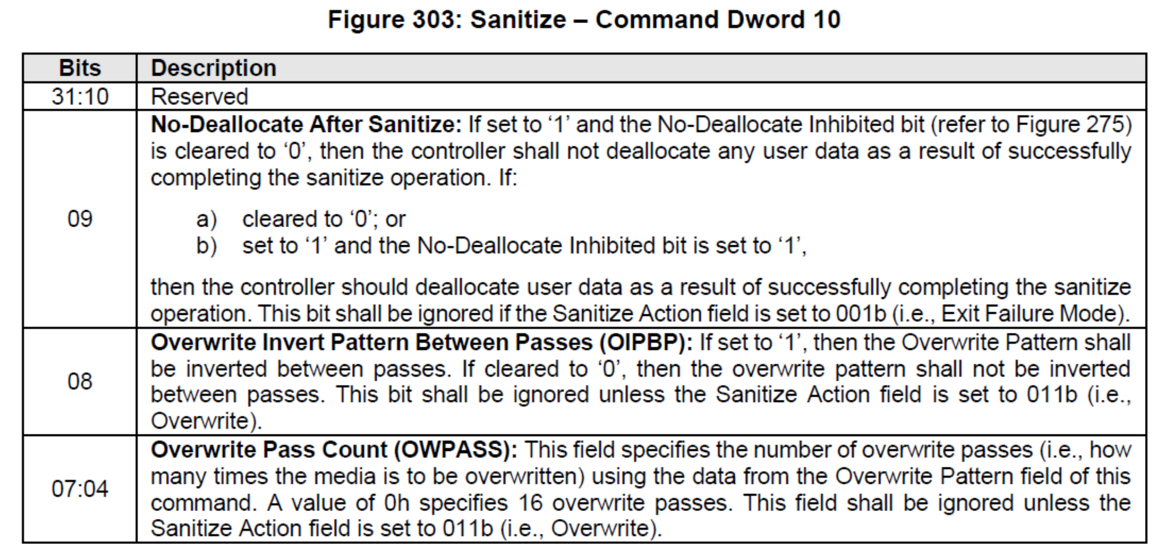

Figure 303: Sanitize – Command Dword 10

Bits	Description
31:10	Reserved
09	**No-Deallocate After Sanitize:** If set to '1' and the No-Deallocate Inhibited bit (refer to Figure 275) is cleared to '0', then the controller shall not deallocate any user data as a result of successfully completing the sanitize operation. If: a) cleared to '0'; or b) set to '1' and the No-Deallocate Inhibited bit is set to '1', then the controller should deallocate user data as a result of successfully completing the sanitize operation. This bit shall be ignored if the Sanitize Action field is set to 001b (i.e., Exit Failure Mode).
08	**Overwrite Invert Pattern Between Passes (OIPBP):** If set to '1', then the Overwrite Pattern shall be inverted between passes. If cleared to '0', then the overwrite pattern shall not be inverted between passes. This bit shall be ignored unless the Sanitize Action field is set to 011b (i.e., Overwrite).
07:04	**Overwrite Pass Count (OWPASS):** This field specifies the number of overwrite passes (i.e., how many times the media is to be overwritten) using the data from the Overwrite Pattern field of this command. A value of 0h specifies 16 overwrite passes. This field shall be ignored unless the Sanitize Action field is set to 011b (i.e., Overwrite).

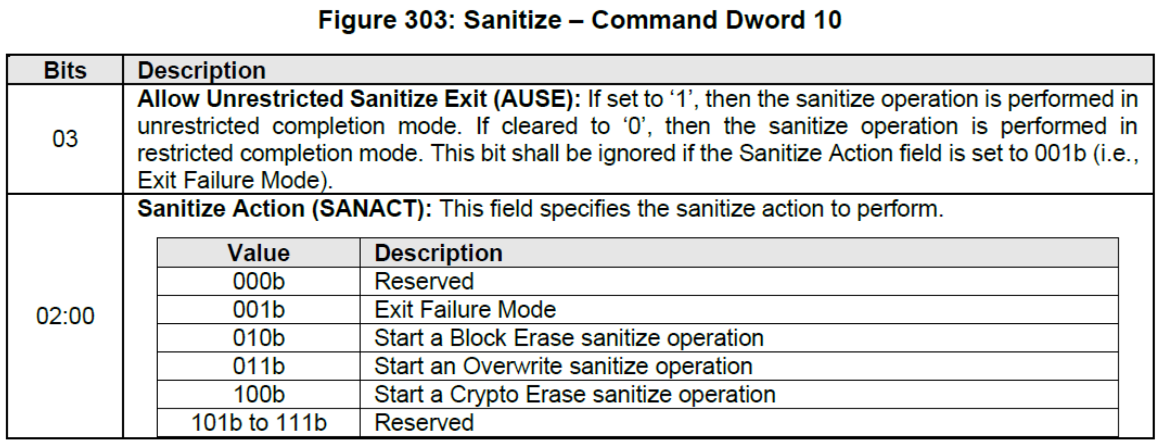

Figure 303: Sanitize – Command Dword 10

Bits	Description
03	**Allow Unrestricted Sanitize Exit (AUSE):** If set to '1', then the sanitize operation is performed in unrestricted completion mode. If cleared to '0', then the sanitize operation is performed in restricted completion mode. This bit shall be ignored if the Sanitize Action field is set to 001b (i.e., Exit Failure Mode).
02:00	**Sanitize Action (SANACT):** This field specifies the sanitize action to perform.

Value	Description
000b	Reserved
001b	Exit Failure Mode
010b	Start a Block Erase sanitize operation
011b	Start an Overwrite sanitize operation
100b	Start a Crypto Erase sanitize operation
101b to 111b	Reserved

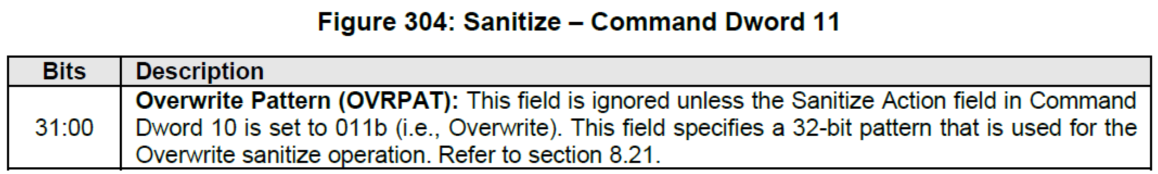

Figure 304: Sanitize – Command Dword 11

Bits	Description
31:00	**Overwrite Pattern (OVRPAT):** This field is ignored unless the Sanitize Action field in Command Dword 10 is set to 011b (i.e., Overwrite). This field specifies a 32-bit pattern that is used for the Overwrite sanitize operation. Refer to section 8.21.

Example with NVMe SSD (thanks Kioxia!)

1. Find a NVMe SSD’s sanitize capabilities through Identify Controller command
2. Send Sanitize command with action -2, block erase
3. Loop: Monitor Sanitize Status with Sanitize Log
4. Sanitize completes

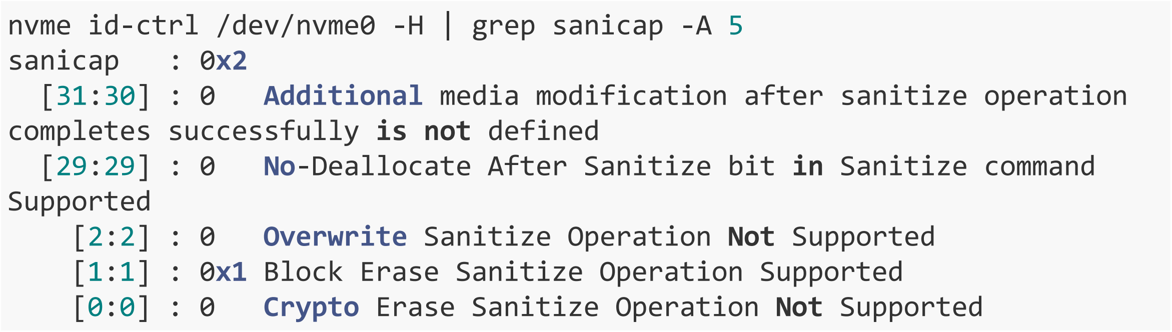

```
nvme id-ctrl /dev/nvme0 -H | grep sanicap -A 5
sanicap   : 0x2
  [31:30] : 0   Additional media modification after sanitize operation
completes successfully is not defined
  [29:29] : 0   No-Deallocate After Sanitize bit in Sanitize command
Supported
    [2:2] : 0   Overwrite Sanitize Operation Not Supported
    [1:1] : 0x1 Block Erase Sanitize Operation Supported
    [0:0] : 0   Crypto Erase Sanitize Operation Not Supported
```

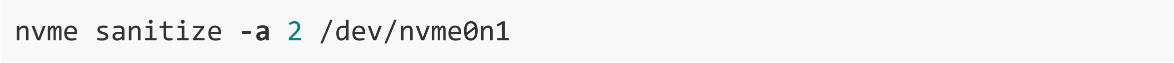

```
nvme sanitize -a 2 /dev/nvme0n1
```

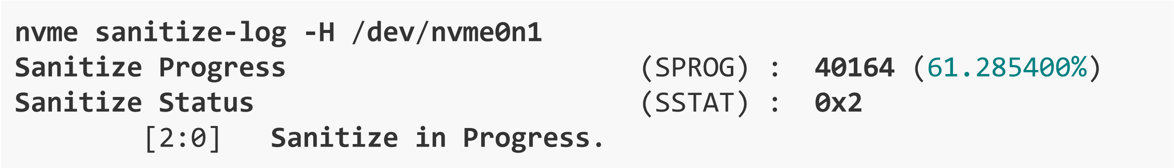

```
nvme sanitize-log -H /dev/nvme0n1
Sanitize Progress                       (SPROG) :  40164 (61.285400%)
Sanitize Status                         (SSTAT) :  0x2
        [2:0]   Sanitize in Progress.
```

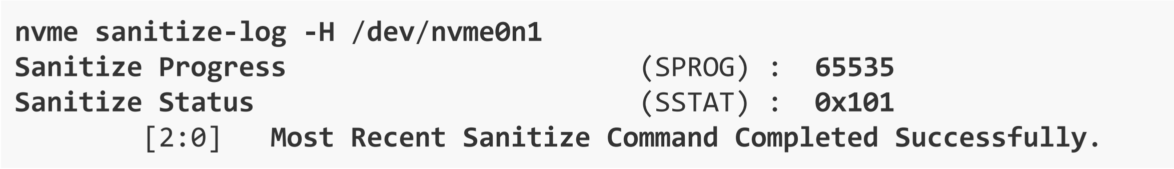

```
nvme sanitize-log -H /dev/nvme0n1
Sanitize Progress                       (SPROG) :  65535
Sanitize Status                         (SSTAT) :  0x101
        [2:0]   Most Recent Sanitize Command Completed Successfully.
```

Questions?

Thanks for Watching Our Webcast

- Please rate this webcast and provide us with feedback
- A link to this webcast and the PDF of the slides are posted to the SNIA Compute Memory and Storage Initiative website at https://www.snia.org/forums/cmsi/knowledge/articles-presentations
- You can also find this webcast and many other videos and presentations on today's topics in the SNIA Educational Library
- A Q&A from this webcast will be posted to the SNIA Compute, Memory, and Storage Blog
- Learn more about Data Governance and Security at https://www.snia.org/technology-focus/data-governance-security

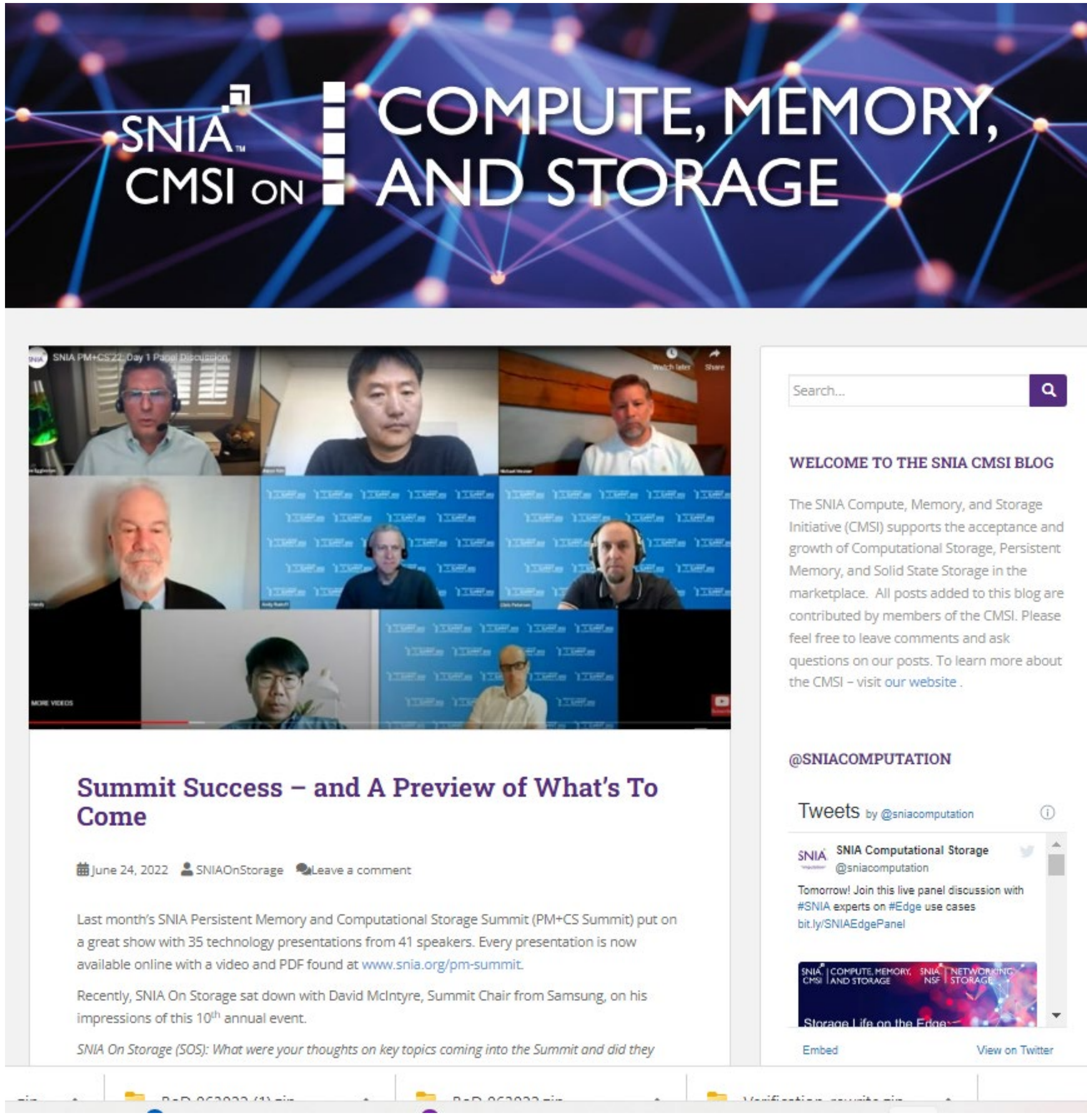
SNIA CMSI ON
COMPUTE, MEMORY, AND STORAGE
SNIA PM+CS'22: Day 1 Panel Discussion
MORE VIDEOS

Summit Success – and A Preview of What's To Come

June 24, 2022 SNIAOnStorage Leave a comment

Last month's SNIA Persistent Memory and Computational Storage Summit (PM+CS Summit) put on a great show with 35 technology presentations from 41 speakers. Every presentation is now available online with a video and PDF found at www.snia.org/pm-summit.

Recently, SNIA On Storage sat down with David McIntyre, Summit Chair from Samsung, on his impressions of this 10th annual event.

SNIA On Storage (SOS): What were your thoughts on key topics coming into the Summit and did they

Search...

WELCOME TO THE SNIA CMSI BLOG

The SNIA Compute, Memory, and Storage Initiative (CMSI) supports the acceptance and growth of Computational Storage, Persistent Memory, and Solid State Storage in the marketplace. All posts added to this blog are contributed by members of the CMSI. Please feel free to leave comments and ask questions on our posts. To learn more about the CMSI – visit our website .

@SNIACOMPUTATION

Tweets by @sniacomputation

SNIA Computational Storage @sniacomputation

Tomorrow! Join this live panel discussion with #SNIA experts on #Edge use cases bit.ly/SNIAEdgePanel

SNIA CMSI COMPUTE, MEMORY, AND STORAGE
SNIA NSF NETWORKING STORAGE
Storage Life on the Edge

Embed View on Twitter