STORAGE DEVELOPER CONFERENCE

SDC 23

BY Developers FOR Developers

Storage Sanitization: Why, When, How

Paul Suhler

KIOXIA Corporation

Abstract

Operators of data storage systems are legally obligated to protect customer data and can be subject to significant penalties for data breaches. This presentation will explore existing and upcoming standards to show the best practices for sanitizing customer data.

Agenda

- Data breaches
- Sanitization of storage devices
- Customer concerns
- Circularity and reuse
- The standards environment
- New directions for sanitization

The Players

- Vendor: The manufacturer of a storage device.
- Organization: The operator (and usually owner) of a storage device.
- User: The entity associated with the data stored on a storage device.
  - May be the organization, storing their corporate data.
  - May be a customer of the organization, renting compute and storage from the organization.
  - “User” can be recursive: A user may handle data private to customers of their own.

Avoiding Data Breaches

- Organizations must ensure that user data does not escape their control.
- Data breach: User data is accessible to an unauthorized entity.
  - Device stolen or disposed of without removing user data.
  - Attacker who has gained entry to the organization's system.
  - An authorized user of the system who accesses another user's data.
- Devices must be sanitized before being repurposed or discarded.

What is Sanitization?

- Sanitization: Eradication of all user data from a storage device.
- Recovery of user data must be infeasible.
  - Different methods of sanitization are resistant to different levels of attacks (See below.)
- Devices implement commands to sanitize user data.

Sanitization Process

- Take device out of customer's production environment.
  - Must ensure that applications do not write it.
  - If the device will be placed back in use, then it can be taken logically offline and sanitized *in situ*.
  - If the device is being moved to a different system or discarded, then it may be moved to a workstation dedicated to sanitization.
- Issue sanitize command.
  - NVMe Sanitize; SCSI SANITIZE; ATA has multiple commands.
- Confirm successful completion of command.
- If not successful, then try again.
- Verify that data was eradicated.
  - Read some or all of device to confirm removal of known pattern.
- If sanitization failed or verification failed, then the device may be destroyed.
- Document sanitization of the device.

Sanitization Methods

- Different methods provide different levels of protection.

- **Clear**: Device remains usable, and user data cannot be read from the device.
- **Purge**: Device remains usable, but user data cannot be recovered from media – even if the device were to be disassembled and the media read at a low level.
- **Destruct**: Device is destroyed and data cannot be recovered from the remains of the media.

- Source: IEEE 2883™–2022.

Techniques for Clear and Purge

- **Overwrite**: Sanitize command writes a specified pattern to all accessible media. Slow process for large capacity devices.
- **Block Erase**: Applies to solid-state media. One “erase block” is erased in a fixed time. Faster then Overwrite, but time can grow with the number of erase blocks.
  - Command standard may specify data to be returned.
- **Cryptographic Erase:** All data is encrypted by the device as it is written. Erase is performed in constant time by changing the encryption key.
  - Command standard may specify data to be returned.

Techniques for Destruct

- **Disintegrate**: Decompose media or break into small pieces.
- **Incinerate**: Burn media until reduced to ashes.
- **Melt:** Liquefy media.

Verifying Results of Sanitization

- Verification: Read device to ensure that user data does not remain.
- Block Erase and Crypto Erase can leave media error correction codes invalid.
- Read commands will fail (“media error”) until new data is written.
- Workarounds:
  - Device front end fabricates read data and does not access media.
  - Devices perform “additional media modification” to make media readable. Slow process.

- Help is coming: NVM Express is defining a mechanism for reading media without reporting media errors.

Sanitization Behavior Varies among Device Types

- There is no single standard describing how sanitize commands and the destruct method work on different device types.
- Different standards organizations – NVM Express, T10/SCSI, T13/ATA Storage Interfaces – have the same set of participants who try to align behaviors.
- IEEE Std 2883™–2022 describes use of different devices' commands.

- Help is coming: Proposed IEEE Project P3406 (Standard for a Purge and Destruct Sanitization Framework).

Customer Concerns

- Liability for a data breach can be in the tens of millions of dollars.
- Liability can exist in perpetuity.
- Is the storage device sanitization firmware buggy?
- Has an attacker compromised the firmware?
- Are there bugs in the software tool that issues the sanitization command?
- Did the technician correctly use the software tool?

- Without confidence in the entire process, the customer may decide to destroy the device, rather than risk a breach.

Circularity and Reuse

- Why *not* destroy the device?
- Destruction of devices is wasteful and has an environmental impact.
  - Device cannot be reused; a replacement must be purchased.
  - Destruction costs money (power, labor, facilities, etc.).
  - Devices contain various metals and other chemicals that should not end up in landfills.
- Nevertheless, if the customer decides that the device must be destroyed, then disassemble it first.
  - Component materials can be fed into different recycling streams.
  - New standards are in development that will provide guidance.

The Standards Environment

- IEEE Security in Storage Working Group (SISWG)
  - IEEE Std 2883™-2022 (IEEE Standard for Sanitizing Storage)
- ISO/IEC 27040 (Storage security)
  - 2nd edition is nearing publication.
  - Has requirements and guidance for technologies and practices.
  - Covers both logical and media-based sanitization.
  - Defers to IEEE 2883 for specific sanitization techniques.
- Trusted Computing Group Storage Working Group (TCG SWG)
  - Opal Subsystem Class and other SSCs.

The Standards Environment

- NIST – National Institute of Science and Technology
  - Cryptographic Module Verification Program (FIPS 140-3).
    - Testing is performed by certified testing labs.
  - Special publications – various aspects of cryptography and security.
    - New research includes algorithms resistant to attacks by quantum computers.
- EU Regulation 2019/424 (“Lot 9”)
  - Refers to “secure data deletion” standards; 27040 and 2883 would fit this category.

New Directions – Verification

- The Crypto Erase and Block Erase techniques invalidate error correction codes in the media, causing reads to fail due to a media error.
- Existing devices can perform “additional media modification” to place readable data on the media – this is very slow.
- New media verification mechanisms skip the media modification and allow reading the media without reporting media errors.
- Repeated reads of the same location must be allowed to return different data, otherwise proprietary media reliability characteristics can be inferred.

New Directions – Sanitize Subcomponents

- Sanitization of subcomponents (e.g., NVMe namespaces).
  - One storage device may be shared by multiple VMs (users), each of which has a different namespace in the same storage device.
  - Swapping a user out requires that their namespace must be sanitized.
  - Other namespaces may continue to be written and read.
  - Some other parts of the storage device must not be sanitized.
    - A controller memory buffer (CMB) may contain a data buffer used for I/Os to other namespaces.

New Directions – Encryption at a Fine Granularity

- Example: A file with data for one person is written with a unique encryption key. This may be a small part of a namespace.
  - NVM Express Key Per I/O functionality allows each Write to use a different key.
  - TCG SWG has defined a standard interface for injecting keys into devices.
- Encryption keys are stored in key management appliances, but are ephemeral in the device.
  - The Key Management Interface Protocol (KMIP) defines the interface between the host and the appliance.
- If the customer is ordered to forget that person's data, then all copies of that key are deleted from the appliance.
- Complexities:
  - How to prove that all copies of the key have been deleted?
  - Purging the entire device may require a second-level key that can be changed.

New Directions – Guidance

- Customers need guidance on appropriate sanitization methods.
  - What are the risks of exposure of different data?
  - What are the appropriate sanitization methods?
  - What are the effectiveness, economics, and environmental consequences?
- IEEE SISWG is developing new standards:
  - IEEE P2883.1 Recommended Practice for Use of Storage Sanitization Methods
    - How to use sanitization to meet your organization's needs.
    - Analysis of value of data and risks from data breaches.
  - IEEE P2883.2 Recommended Practice for Virtualized and Cloud Storage Sanitization
    - How to implement sanitization for virtualized and cloud storage systems.
    - Will address the concerns for storage at scale.

New Directions – Guidance

- IEEE SISWG is developing new standards:
  - IEEE P3406 (Standard for a Purge and Destruct Sanitization Framework) – pending approval of project.
    - Will provide requirements for standards organizations defining purge and destruct techniques.
    - Especially important for new storage technologies (e.g., DNA or crystal storage).
    - Need to make data recovery “infeasible using state of the art laboratory techniques”.
    - Some techniques will need to be deprecated.
      - E.g., if AES were to be broken, then Crypto Erase implementations that rely on it would be ineffective.

Compliance Testing and Device Certification

- Private testing companies are usually engaged by buyers of storage devices.
  - Most testing involves directly reading the media (HDD spin stands, NAND raw interface.
  - Device vendors may or may not help by showing where user data is stored.
- Will device vendors pay for certification?
  - Cost must be passed to customers; will it be overall higher?
  - NIST FIPS 140-3 compliance testing has been paid for by vendors.
- IEEE SISWG is exploring using the IEEE Conformity Assessment Program to establish a media sanitization certification program.

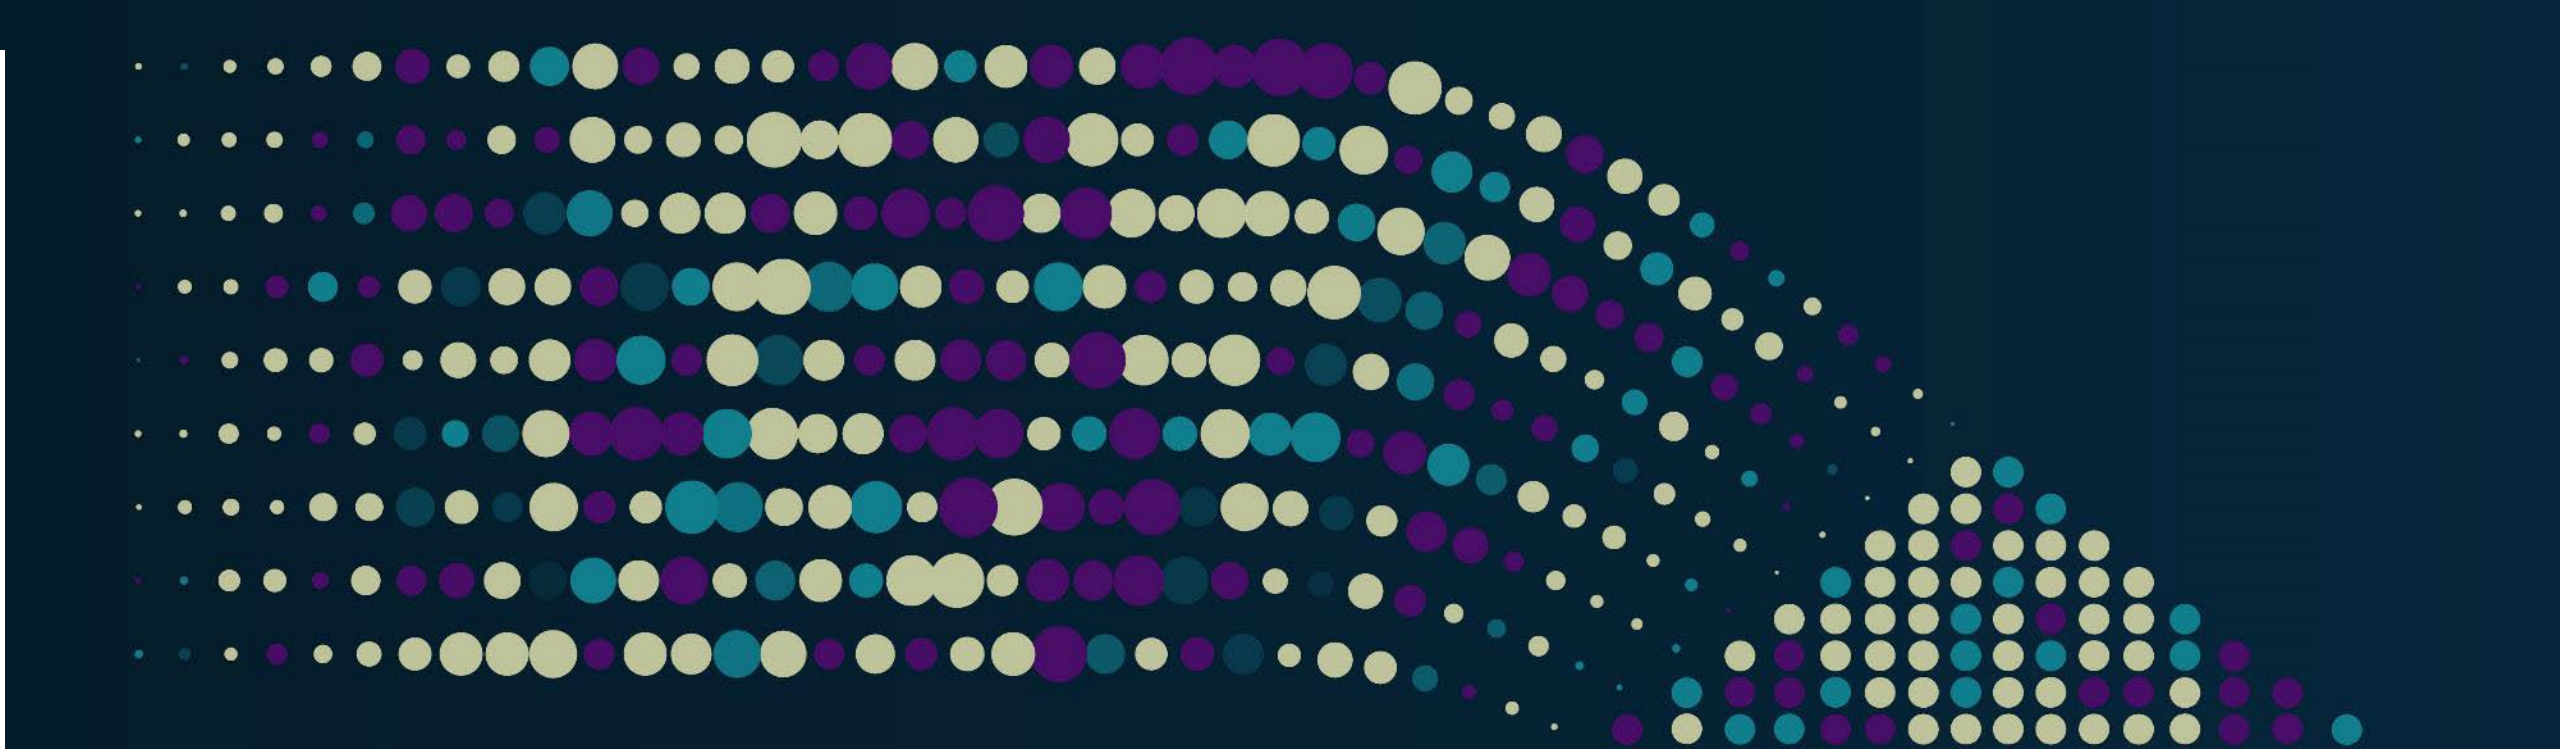

Please take a moment to rate this session.

Your feedback is important to us.

