SNIA STORAGE SECURITY SUMMIT
Wednesday, May 11, 2022 • Virtual

Key Per IO - Fine Grain Encryption For Storage

With TCG's Per-I/O Encryption Key Selection

Frederick Knight, NetApp

A SNIA Event

SNIA Legal Notice

- The material contained in this tutorial 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 containin 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.

Agenda

- Data at Rest Protection - Background
- Key Per I/O Overview
- Key Per I/O SSC And I/O Architectures Interactions

Existing Data At Rest Protection vs. Key Per IO

Background On Data At Rest Protection

Data At Rest Protection

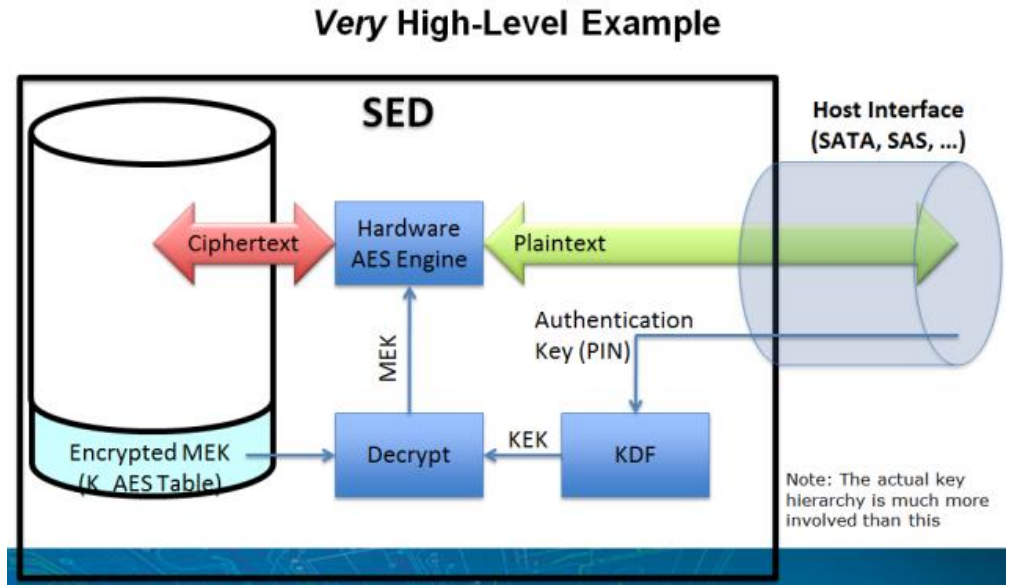

Properties

- Encrypt all user accessible data all the time, at interface speeds
- Keys generated & stored in NVM by the storage device
- Media Encryption Key (MEK) associated with contiguous LBA ranges or Namespaces
- Opal/Enterprise SSC* deliver passwords to drive in the clear (when not using Trusted Computing Group (TCG)* - Secure Messaging)

*Other names and brands may be claimed as the property of others.

Can we do better?

Desired properties:

1. Select an encryption key for each I/O to a Storage Device?
   - Associate encryption domains with higher-level objects (abstractions) than drives or volumes.
   - Crypto erase individual higher-level objects
   - Easier to support European Union's General Data Protection Regulations' "Right to be forgotten"
2. Externally manage Media Encryption keys?
   - Centralized key management infrastructure, consistent key policies
   - High assurance key generation and control, e.g., master keys in HSM (Hardware Security Module)
3. Ensure that a Storage Device with no power has no encryption keys?
   - Shorter physical drive loss/theft discussion with security auditor
   - Easier decommissioning process

Key Per I/O Overview

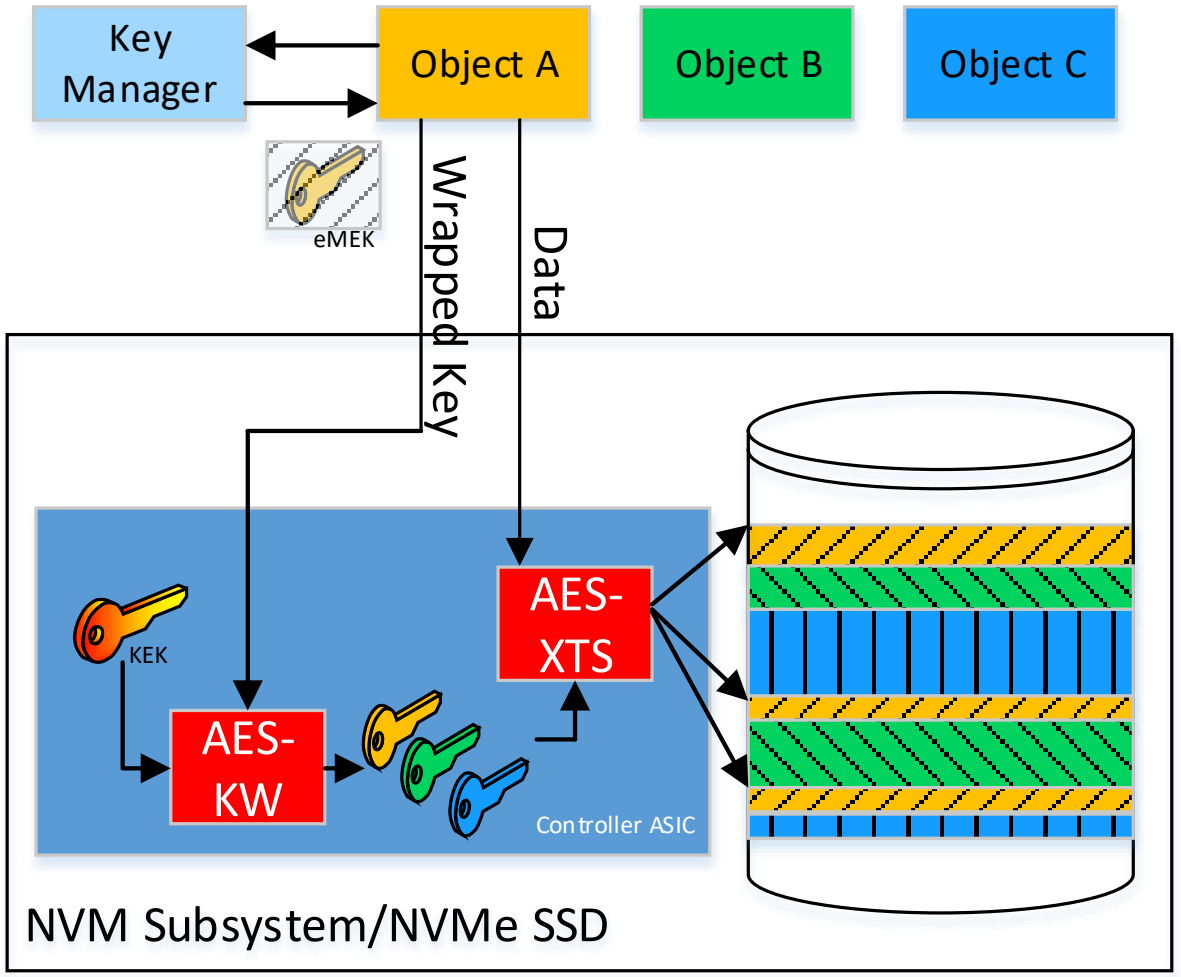

- Encrypted Media Encryption Keys are injected into Self Encrypting Drive key cache and assigned a “Key Tag” by SW
- Subsequent I/O can use the “Key Tag” to encrypt/decrypt data to/from the storage device in a non-contiguous fashion
- Media Encryption Keys (MEKs) are encrypted (wrapped) by a Key Encryption Key (KEK)
- Media Encryption Keys (MEKs) are not stored in the NVM of the drive and are lost on power loss
- Crypto erase accomplished by deleting the MEK from the Key Manager and the SSD or by sanitizing entire SSD

Key Per I/O Architectural Elements

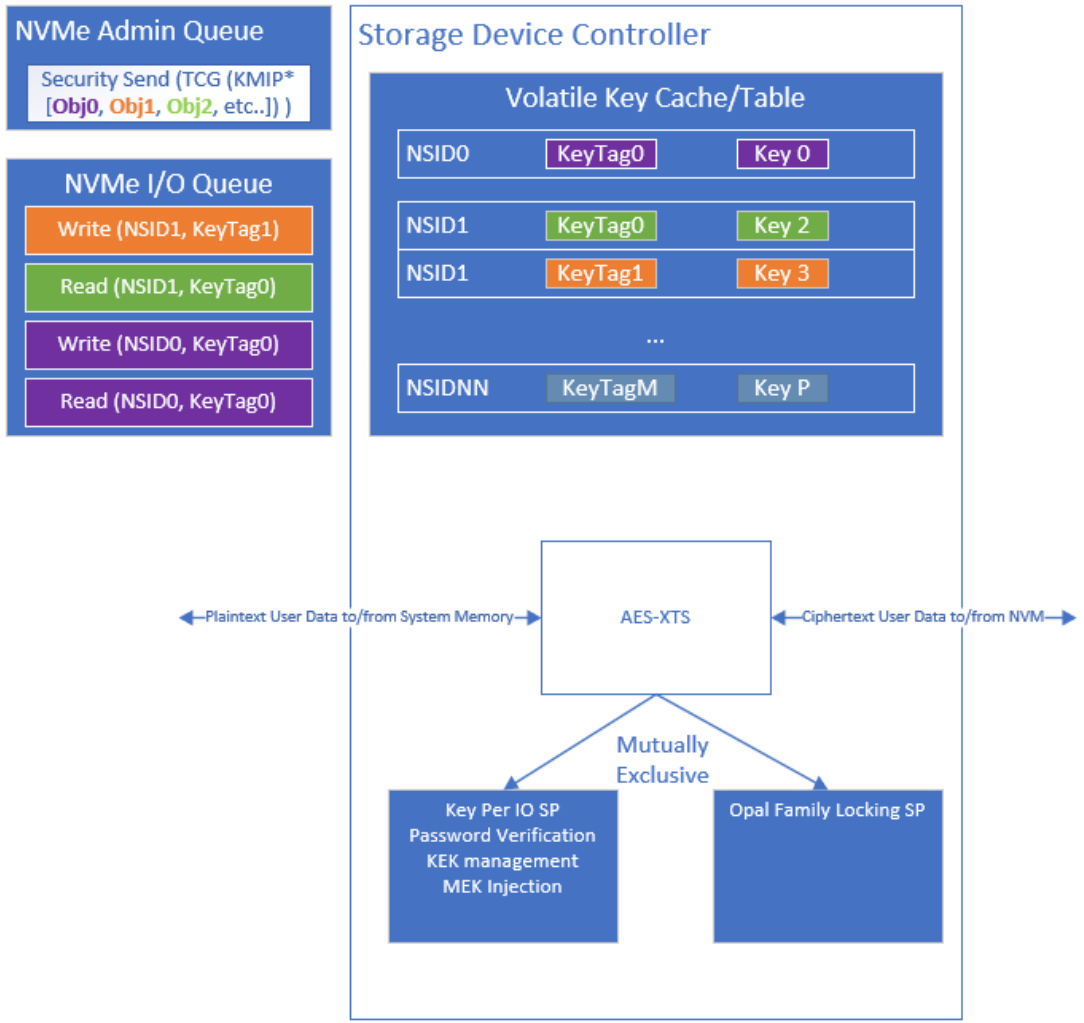

- Encrypted Media Encryption Keys (eMEKs) and their wrapping Key Encryption Keys (KEKs) are injected into the storage device via the Security Send & Receive
  - Specification in progress within the TCG SWG*
- OASIS KMIP* V2.1 for specifying Key data and its transportation over Security Send & Receive
  - Specification engagement in progress between TCG SWG & OASIS KMIP*
- Subsequent I/O can then use the “Key Tag”, a newly defined field in I/O commands, to specify the key that the device uses to encrypt/decrypt data to/from the storage device
  - Specification work in progress within NVMe*

*Other names and brands may be claimed as the property of others.

Data At Rest Tenant Isolation with Key Per I/O

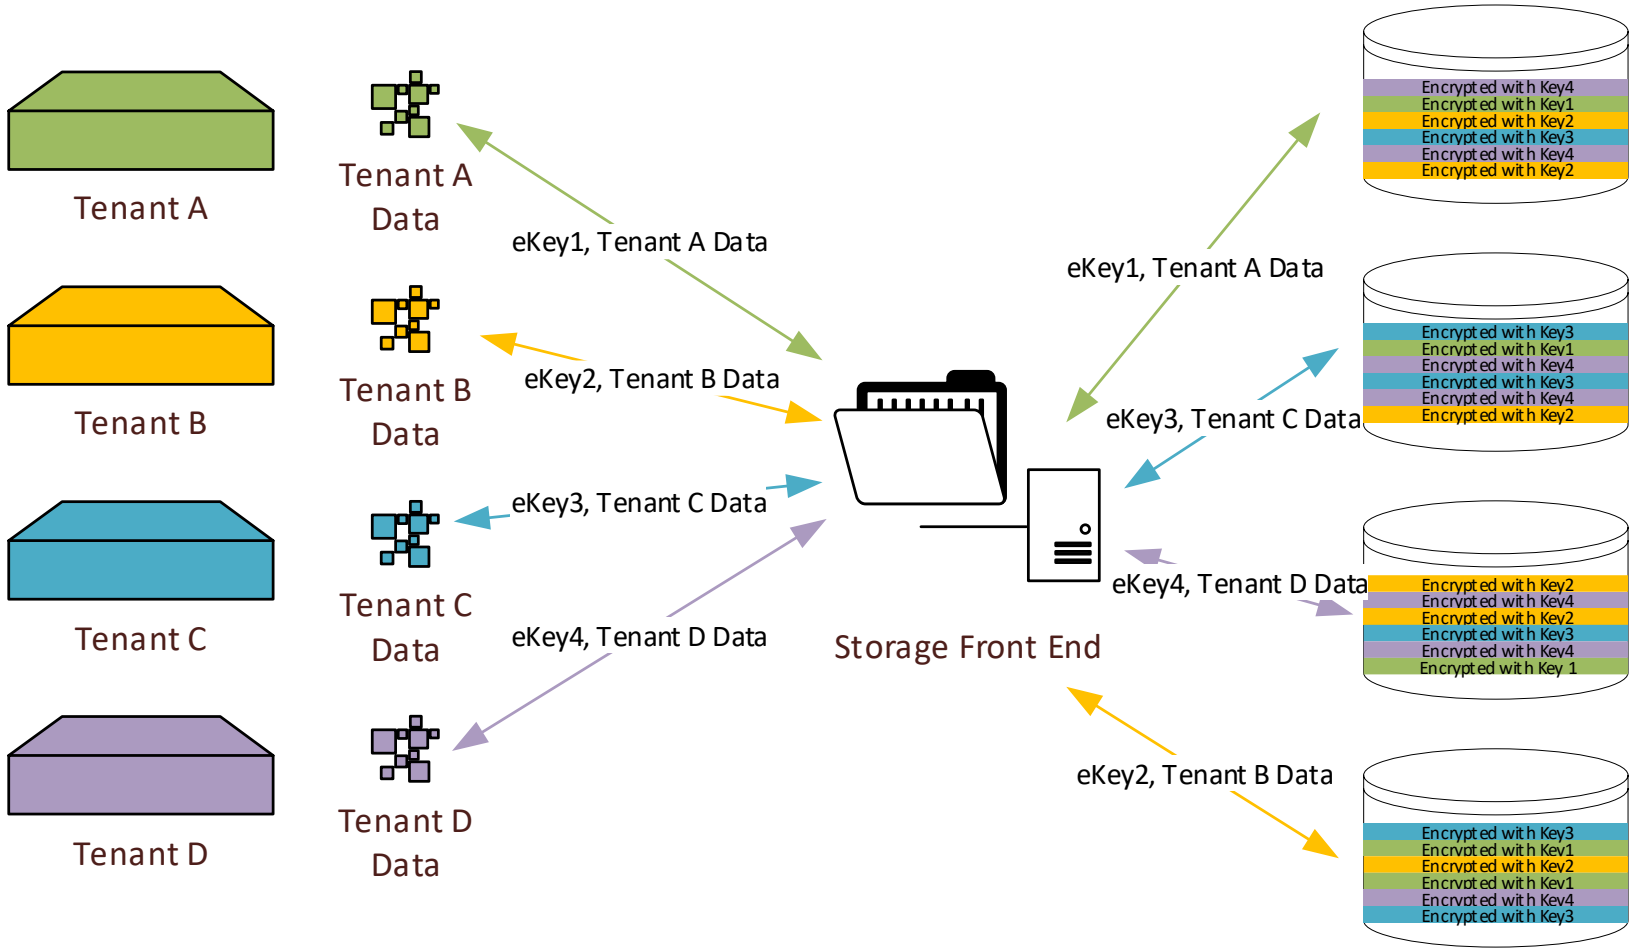

Key Per I/O SSC and I/O Architecture Interactions

KPIO Discovery

Host Detection of KPIO

- Number of Key Tags supported
- Granularity and alignment of operations
- NVMe Identify command
  - Per namespace
- TCG Discovery (Security Send and Security Receive)
  - Authenticate
  - Security Receive (Level 0 Discovery)
  - Discovery security characteristics

KPIO Configuration

Establish KEKs (Key Encryption Keys)

- Agreement between host and device for secure transmission of the KEKs (secure manufacturing (pre-shared keys), public/private key pairs (PKI), certificate, etc)
- Host obtains MEK (Media Encryption Keys) from a key management database (e.g., KMIP)
- MEKs are “wrapped” with KEKs and sent to the device

Wrapped MEKs sent from the host to the device

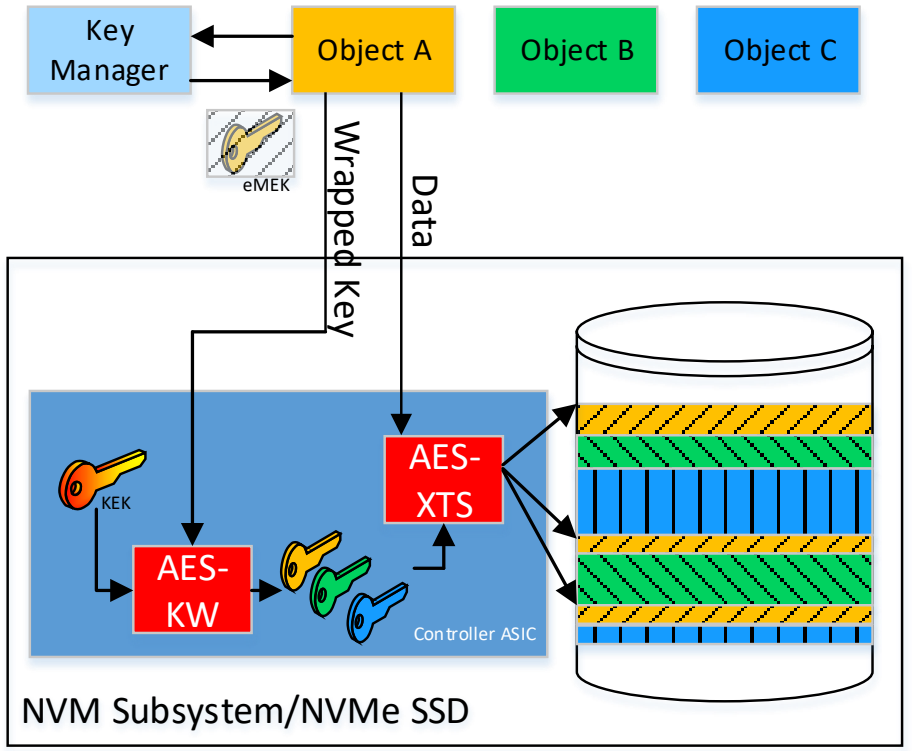

KPIO Configuration

MEKs are Loaded the Device Key Cache

- Associate each MEK injected with a Key Tag
  - Per namespace – loaded using Security Send command

Key Tag	MEK example (256 bit)
1	0x1234567890ABC…90ABCDEF
2	0x1234567890ABC…90ABCDE0
100	0x1234567890ABC…90ABCDE1
101	0x1234567890ABC…90ABCDE2
103	0x1234567890ABC…90ABCDE3
200	0x1234567890ABC…90ABCDE4
217	0x1234567890ABC…90ABCDE5

KPIO Configuration

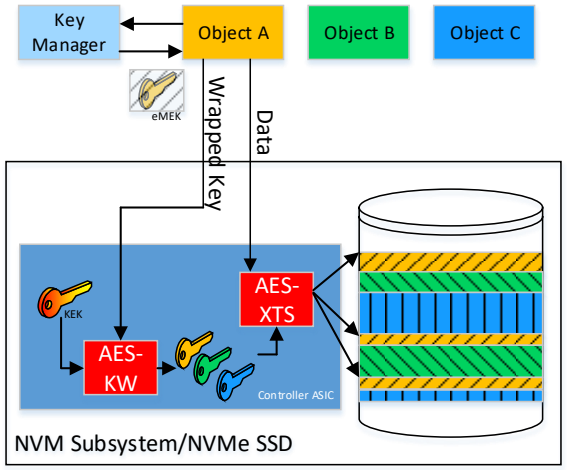

Load New Keys

- Associate a Key Tag with a different MEK
  - Per namespace – loaded using Security Send command

Key Tag	MEK example (256 bit)
1	0x1234567890ABC…90ABCDEF
2	0x1234567890ABC…90ABCDE0
100	0x1234567890ABC…90ABCDE6
109	0x1234567890ABC…90ABCDE7
110	0x1234567890ABC…90ABCDE8
111	0x1234567890ABC…90ABCDE9
220	0x1234567890ABC…90ABCDEA

KPIO Usage

I/O Command Usage

- Compare
- Copy
- Verify
- Read
- Write
- Write Zeroes
- Zone Append

- A field in each command to specify the Key Tag value to use for that individual I/O
- An indicator that a Key Tag is present

Key Tag	MEK example (256 bit)
1	0x1234567890ABC…90ABCDEF
2	0x1234567890ABC…90ABCDE0
100	0x1234567890ABC…90ABCDE6
109	0x1234567890ABC…90ABCDE7
110	0x1234567890ABC…90ABCDE8
111	0x1234567890ABC…90ABCDE9
220	0x1234567890ABC…90ABCDEA

KPIO Example Commands

- WRITE (LBA=100, LEN=8, flag=1, keytag=1)

MEK = 0x1234567890ABC…90ABCDEF

- WRITE (LBA=200, LEN=16, flag=1, keytag=100)

MEK = 0x1234567890ABC…90ABCDE6

- READ (LBA=100, LEN=8, flag=1, keytag=1)
  - Gets your data back
- READ (LBA=200, LEN=16, flag=1, keytag=1)
  - Gets error or bogus data
- READ (LBA=200, LEN=16, flag=1, keytag=100)
  - Gets your data back

flag = 1 means the keytag is present

Key Tag	MEK example (256 bit)
1	0x1234567890ABC…90ABCDEF
2	0x1234567890ABC…90ABCDE0
100	0x1234567890ABC…90ABCDE6
109	0x1234567890ABC…90ABCDE7
110	0x1234567890ABC…90ABCDE8
111	0x1234567890ABC…90ABCDE9
220	0x1234567890ABC…90ABCDEA

KPIO Impact in TCG

Security Send / Security Receive Commands

- Uses new TCG protocol ID
- Authentication
- Discovery
- Key Injection method (Establish Key Tag to MEK association)
- Key Clear method (Remove Key Tag to MEK association)
- Key Replacement method (Replace MEK for a Key Tag)
- Securely Purge Key Cache
- Define encryption / decryption algorithms that can be supported (e.g., XTS-AES-256)

KPIO Impact On Hosts

Host Responsibilities to use KPIO

- Hosts must manage the full life cycle of the Keys
  - Including secure purging of the keys
- Host is responsible for the correctness of the MEK injection / key tag association and use of the correct key tag for each I/O command
- Host is responsible for preventing incorrect key tag use
  - Key tag associations may change during operation – such as key tag cache size smaller than key tag needed usage
  - Using the key tag associated with the correct MEK
- Host must handle errors for improper use of key tags
  - Invalid key tag value (out of range), or a key tag with no associated MEK
  - Trying to use a key tag before injection is complete or after removal

KPIO Project Status

Current Key Topics in Progress

- Details of the MEK / KEK secure injection process
  - KMIP based methods
- Incorrect MEK detection optional capability
  - Incorrect MEK should not look like a Media Error
  - Does incorrect MEK just return “bogus” data
  - UUID association
- Testing use cases
- Still a work in Process

- Work at NVMe is nearly complete
- TCG work is continuing
- Come join us at TCG Storage Work Group to continue the discussions!

https://trustedcomputinggroup.org/

or

admin@trustedcomputtinggroup.org

KPIO For Other IO Architectures

What about SCSI and/or SATA

- The same TCG architecture is used by SCSI and SATA
  - Security Send / Security Protocol Out / Security Receive / Security Protocol In
- But completely new I/O commands would be required
  - Such as 32-byte CDBs for SCSI (to carry the Key Tag value)

- NO interest being shown to undertake such an effort

KPIO Key Takeaways

- The KPIO SSC is being defined such that an SD that claims TCG Opal SSC compatibility could be a KPIO SSC.
- Intended to protect confidentiality of data at rest from unauthorized access once it leaves the owner's control.
- Creating a fine-grained approach to enhance SED technology to better support multi-tenancy usage models.
- Standards based designs for multi-vendor interoperability.

Please take a moment to rate this session.

Your feedback is important to us.

Organizations

- Trusted Computing Group (TCG) is a not-for-profit organization formed to enable secure computing through open standards and specifications. Benefits of TCG technologies include protection of business-critical data and systems, secure authentication and strong protection of user identities, and the establishment of strong machine identity and network integrity. Trusted hardware and applications reduce enterprise total cost of ownership and support regulatory compliance. Through its member-driven work groups, TCG enables the benefits of trust in computing devices from mobile to embedded systems, as well as networks, storage, infrastructure, and cloud security. Almost all enterprise PCs, many servers, and embedded systems include the TPM; while networking equipment, drives, and other devices and systems deploy other TCG specifications, including self-encrypting drives and network security specifications.
- The NVM Express® (NVMe®) family of specifications define how host software communicates with non-volatile memory across multiple transports like PCI Express® (PCIe®), RDMA, TCP and more. It is the industry standard for solid state drives (SSDs) in all form factors (U.2, M.2, AIC, EDSFF). NVM Express is the non-profit consortium of tech industry leaders defining, managing and marketing NVMe technology. The latest NVMe 2.0 library of specifications consists of multiple documents, including the NVMe Base specification, Command Set specifications (NVM Command Set specification, ZNS Command Set specification, KV Command Set specification), Transport specifications (PCIe Transport specification, Fibre Channel Transport specification, RDMA Transport specification and TCP Transport specification) and the NVMe Management Interface specification.
- OASIS: KMIP Technical Committee: The KMIP Technical Committee will develop specification(s) for the interoperability of key management services with key management clients. The specifications will address anticipated customer requirements for key lifecycle management (generation, refresh, distribution, tracking of use, life-cycle policies including states, archive, and destruction), key sharing, and long-term availability of cryptographic objects of all types (public/private keys and certificates, symmetric keys, and other forms of "shared secrets") and related areas.

