# Security for NVMe

**August 11, 2015**

**Jason Cox**
**Security Architect, Intel Corporation**
**Co-Chair, TCG Storage Work Group**

# Objectives

- Background on Trusted Computing Group Storage specifications
- Why Opal:  Details on TCG Opal “Family” specifications and their value as security management interface for NVMe client and enterprise storage devices
    - Opal overview
    - SED overview
    - The Opal “Family”
- Ongoing TCG, NVMe engagement
- Comparing Opal to alternative security management mechanisms

# Trusted Computing Group

## Trusted Computing Group (TCG)

- Cross-industry organization formed to develop, define, and promote standards
  - Work Groups focused on TPM, Storage, Networking, Mobile, and more
  - **Booth #550**
- TCG Storage Work Group
  - Defines specifications related to Storage Device-based security features

www.trustedcomputinggroup.org

# TCG Storage Specifications

### Core Specification (Core Spec)

- Overall architecture – a description of the underlying constructs to be used in the device specifications.

### Storage Interface Interactions Specification (SIIS)

- Describes the interactions of the TCG SWG specifications with the underlying storage interface protocols, such as ATA, SCSI, USB, etc.

### Security Subsystem Class (SSC)

- Device specifications, consist primarily of a subset of the functionality contained in the Core Spec.
- Opal, Opalite, Pyrite, Enterprise

### Feature Sets

- These are documents that define extensions to the basic functionality of SSCs.
  - Created to allow for simple extensions to be added to the SSC at a faster pace.
  - Additionally, it allows for features that only appeal to a subset of the market to be standardized.
  - Generally "Optional", may be "Mandatory" by spec (e.g., PSID)

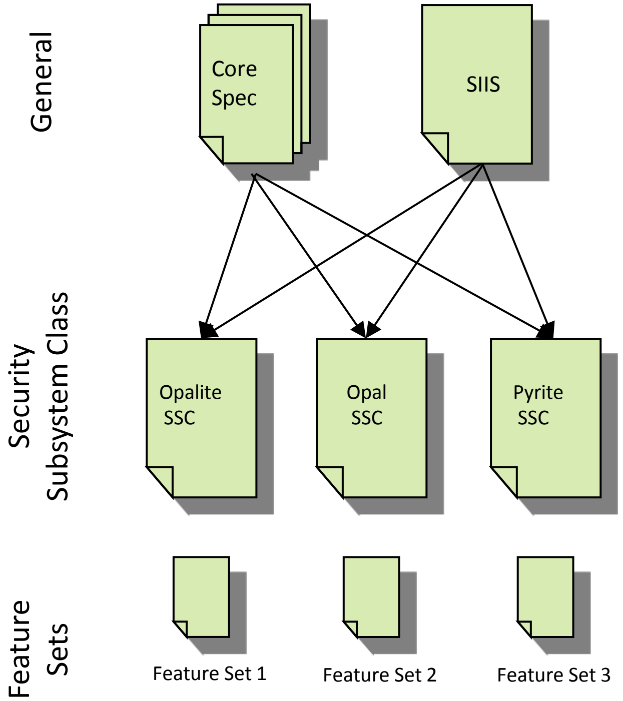

TCG Storage Specifications can be downloaded here:
http://www.trustedcomputinggroup.org/developers/storage

# Self-Encrypting Drive (SED)

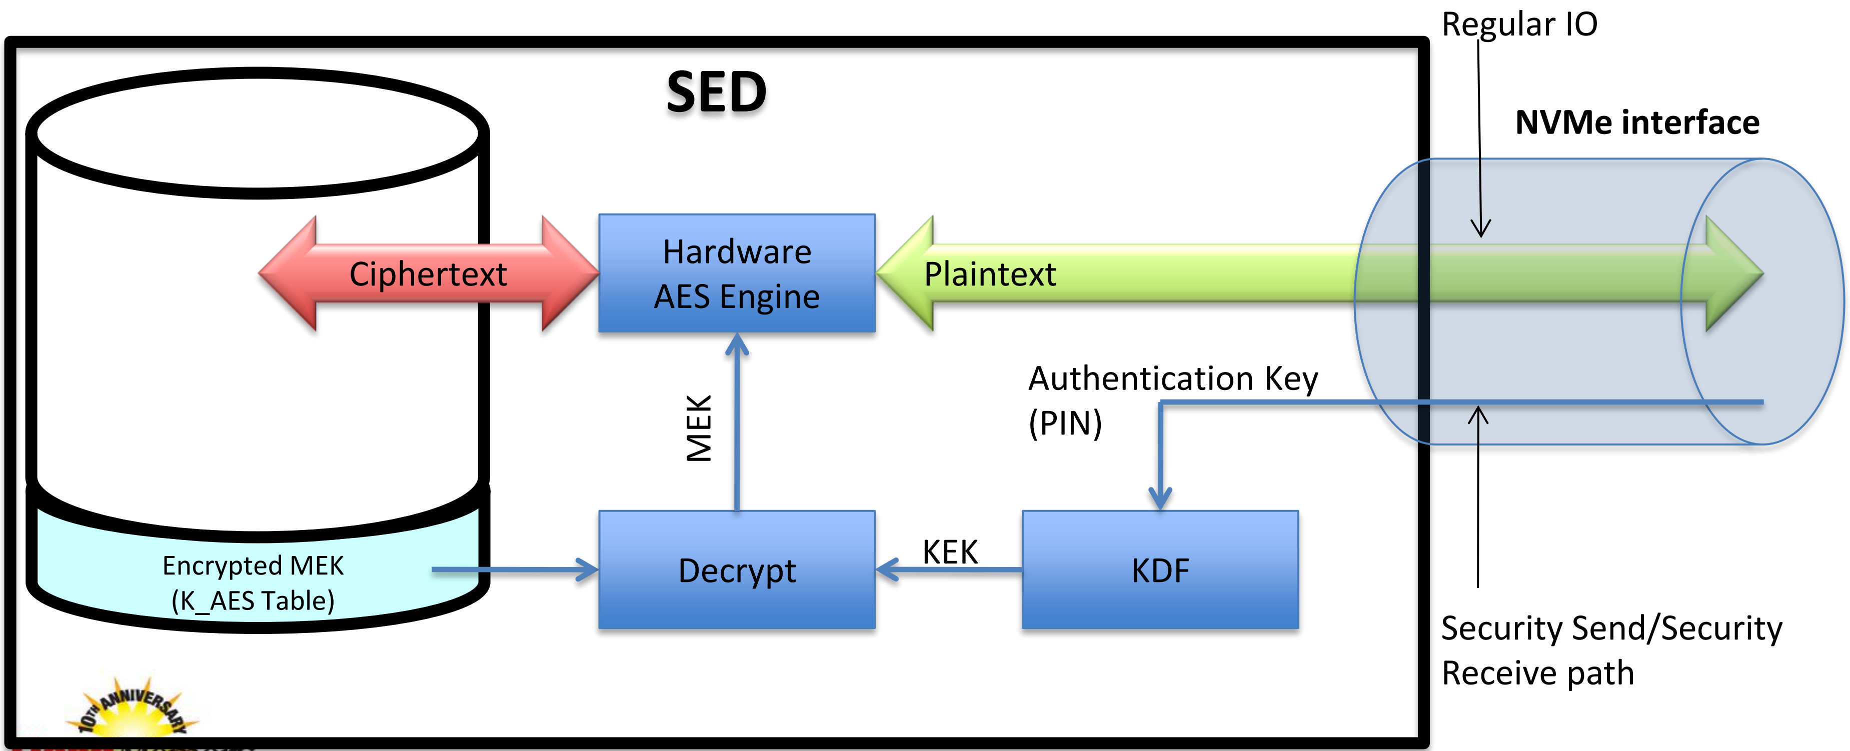

**High-Level Example**

# Opal in One Slide

**Opal SSC:**

- Defines the full-featured interface for managing security features in a storage device, including device encryption.
- **Threat model: protect confidentiality of stored user data against unauthorized access once it leaves the owner's control (when drive and system are powered off)**

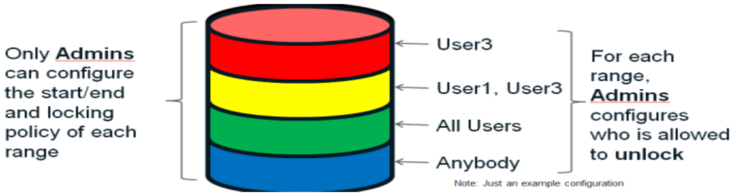

Only **Admins** can configure the start/end and locking policy of each range

User3

User1, User3

All Users

Anybody

For each range, **Admins** configures who is allowed to unlock

Note: Just an example configuration

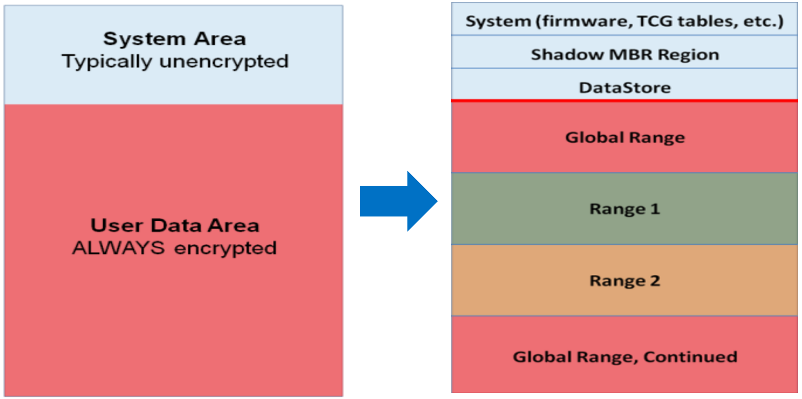

**System Area**
Typically unencrypted

**User Data Area**
ALWAYS encrypted

**System (firmware, TCG tables, etc.)**

**Shadow MBR Region**

**DataStore**

**Global Range**

**Range 1**

**Range 2**

**Global Range, Continued**

## Important Points:

- Each LBA Locking Range has its own media encryption key.
- Locking Ranges are locked after a storage device power cycle.
- Admin assigns access to unlock Ranges to 0 or more Users.
- Each Locking Range can be independently cryptographically erased.
- The Shadow MBR region stores ISV SW "Pre Boot Environment" to capture unlock password and unlock Ranges to allow OS boot.

# The Opal “Family”

The Opal “Family” – defined by request of NVMe to scale across the needs of NVMe in Client and Enterprise solutions

- Opalite – subset of Opal
  - Supports only a single “Global” range
  - Supports fewer User credentials
- Pyrite – “non-encrypting” version of Opalite
  - Does not specify capabilities for cryptographic protection of user data at rest

Opal, Opalite, and Pyrite:

- Common communications protocol, data structures, and commands

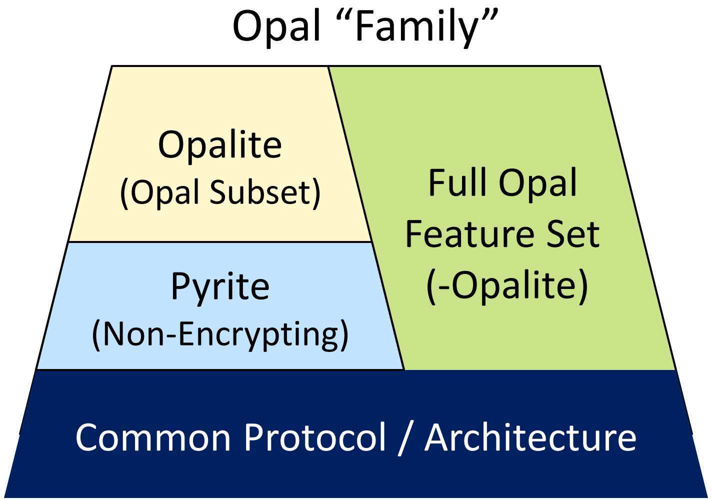

# Opal, Opalite, Pyrite Comparison picture

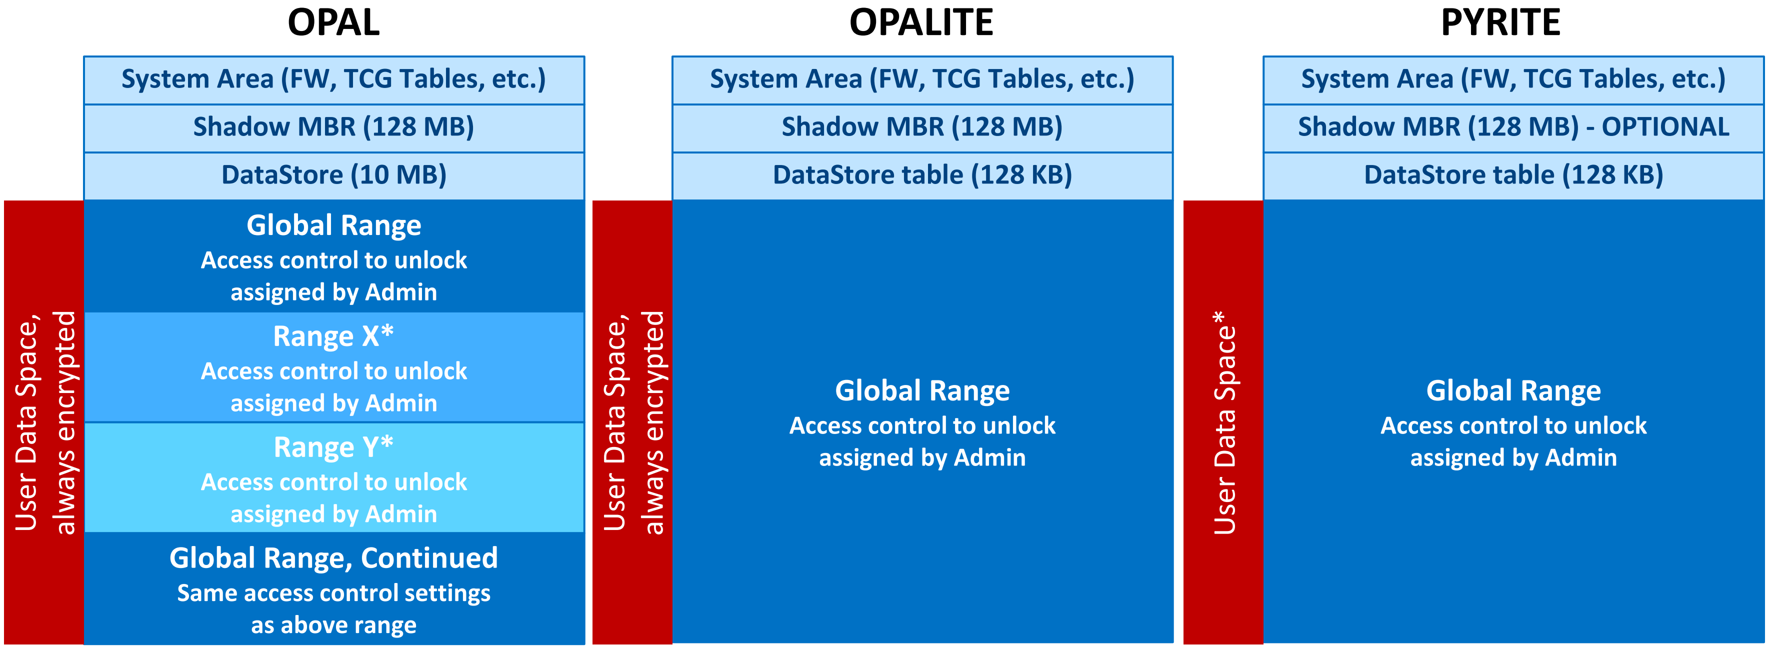

## OPAL

System Area (FW, TCG Tables, etc.)

Shadow MBR (128 MB)

DataStore (10 MB)

User Data Space, always encrypted

**Global Range**
Access control to unlock assigned by Admin

**Range X***
Access control to unlock assigned by Admin

**Range Y***
Access control to unlock assigned by Admin

**Global Range, Continued**
Same access control settings as above range

## OPALITE

System Area (FW, TCG Tables, etc.)

Shadow MBR (128 MB)

DataStore table (128 KB)

User Data Space, always encrypted

**Global Range**
Access control to unlock assigned by Admin

## PYRITE

System Area (FW, TCG Tables, etc.)

Shadow MBR (128 MB) - OPTIONAL

DataStore table (128 KB)

User Data Space*

**Global Range**
Access control to unlock assigned by Admin

*Opal 2.00 supports Global Range plus at least 8 configurable ranges

*Pyrite SSC does not specify encryption of user data

# Opal Family - Compared

| Feature | Opal V2.00 SSC | Opalite SSC<br>(Opal 2.00 subset) | Pyrite SSC (Non-encrypting version of Opalite) |
|---|---|---|---|
| Core Spec Version Supported | V2.00 | V2.00 | V2.00 |
| Activation and Life Cycle | Yes | Yes | Yes |
| Number of Admins/Users | 4 Admin,<br>8 User | 1 Admin,<br>2 User | 1 Admin,<br>2 User |
| Min Number of Required LBA Ranges | Global Range + 8 | Global Range only | Global Range only |
| Min DataStore Size (General Purpose Storage) | 10MB | 128KB | 128KB |
| Min MBR Table Size | 128MB | 128MB | 128MB (Optional) |
| Configurable Access Control | Yes | Yes | Yes |
| PSID | Optional (Required in v2.01) | Required | Not Required (recommended as Prohibited due to lack of integrated data sanitization) |
| Media Encryption | Required | Required | Not Specified |
| Crypto Erase | Revert, RevertSP, GenKey methods for device and locking range level erase granularity | Revert, RevertSP, GenKey methods for device and locking range level erase granularity | No user data erase supported – relies on native interface erase capability |

# WIP: Namespace Interactions

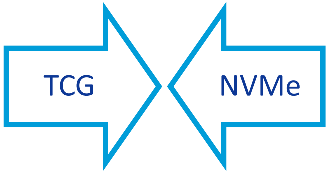

## TCG Storage Interface Interactions

- Updates to Namespace Interactions in progress (targets SIIS v1.05)

## Specifies required support for 2 scenarios:

- Multiple namespaces can be supported with all mapped to the Opal Global Range
- A single namespace can be supported with multiple Opal “Locking ranges” all mapped within the 1 namespace

### Multiple Namespaces

| Opalite | |
|---|---|
| Range | Namespace |
| Global | NS1 |
| | NS2 |
| | ...NSN |
| **Pyrite** | |
| Range | Namespace |
| Global | NS1 |
| | NS2 |
| | ...NSN |
| **Opal** | |
| Range | Namespace |
| Global | NS1 |
| | NS2 |
| | ...NSN |
| Range1 | “Blocked” |
| Range2 | “Blocked” |
| Range3 | “Blocked” |
| Range4 | “Blocked” |
| Range5 | “Blocked” |
| Range6 | “Blocked” |
| Range7 | “Blocked” |
| Range8 | “Blocked” |

If multiple namespaces are created, then locking of all are controlled together.

### Multiple Locking Ranges

| Opalite | |
|---|---|
| Range | Namespace |
| Global | NS1 |
| **Pyrite** | |
| Range | Namespace |
| Global | NS1 |
| **Opal** | |
| Range | Namespace |
| Global | NS1 |
| Range1 | NS1 |
| Range2 | NS1 |
| Range3 | NS1 |
| Range4 | NS1 |
| Range5 | NS1 |
| Range6 | NS1 |
| Range7 | NS1 |
| Range8 | NS1 |

If multiple Locking ranges are configured, then they all are within a single namespace, and additional namespaces cannot be created.

# Namespace Interactions

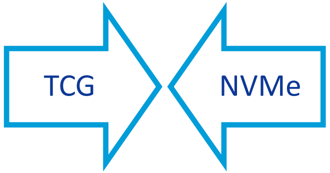

Architecture of enhanced configurability is in process as well.

- When namespaces are created, the Global Range settings apply.
- Namespaces can be associated with one or more Locking objects, to enable separate locking of that namespace or LBA ranges within that namespace.

TCG SWG is seeking input on use cases.

| Range | Namespace |
|---|---|
| Global | NS1<br>NS3<br>NS7 |
| Range1 | NS2 |
| Range2 | NS4 |
| Range3 | NS4 |
| Range4 | NS5 |
| Range5 | NS6 |
| Range6 | NS6 |
| Range7 | NS8 |
| Range8 | NS9 |

One or more locking ranges associated with “configured” namespaces, allowing these namespaces to be unlocked separately, with differently configurable access controls.

# IEEE 1667 and NVMe

IEEE 1667 TCG Transport Silo is a requirement for "eDrive" support

- eDrive in 30 seconds:
  - Starting with Windows 8, MS BitLocker is able to manage SEDs that implement Opal 2.00, Single User Mode Feature Set, and the IEEE 1667 TCG Transport Silo

IEEE 1667 has begun working on a IEEE 1667 transport technical proposal for NVMe

- Enables general access to IEEE 1667 silos over NVMe, including 1667 TCG Transport Silo
  - TCG Transport Silo – alternate transport for TCG Opal commands
- Enables management of Windows eDrive for NVMe Opal SEDs which use Opal 2.00

See www.ieee1667.com for more information on IEEE 1667

# Opal and Assurance

- Opal SSC Test Cases Specification
  - Baseline for Opal Certification
    - Covers Opal 1.00, 2.00, and 2.01
  - ***Currently in pre-publication review:***
    - http://www.trustedcomputinggroup.org/resources/specifications_in_public_review
      - http://www.trustedcomputinggroup.org/files/resource_files/99188CB2-1A4B-B294-D0DB1CF3A7136274/Opal_SSC_Certification_Test_Cases_v2_00_r1_85_Public%20Review.pdf
- Common Criteria Encryption Engine and Authorization Acquisition cPPs (Feb 2015)
  - Specifies security evaluation for Self-Encrypting Drives (SED) and SED management software
  - http://www.commoncriteriaportal.org/pps/?cpp=1a

Opal compliance and assurance are high priority OEM/customer requests.

# Opal, Enterprise Comparison

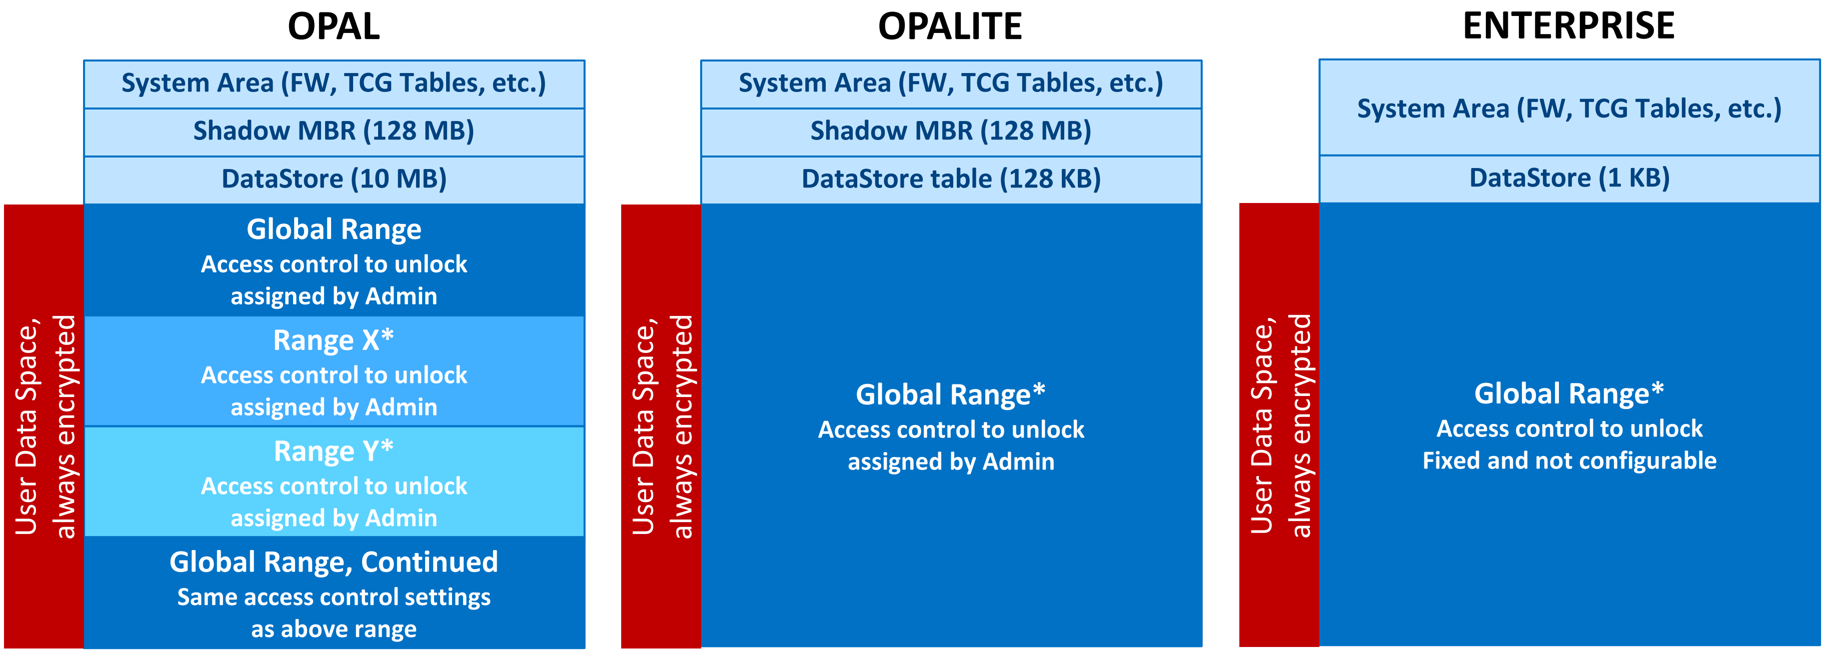

**OPAL**

User Data Space, always encrypted

System Area (FW, TCG Tables, etc.)

Shadow MBR (128 MB)

DataStore (10 MB)

Global Range
Access control to unlock assigned by Admin

Range X*
Access control to unlock assigned by Admin

Range Y*
Access control to unlock assigned by Admin

Global Range, Continued
Same access control settings as above range

**OPALITE**

User Data Space, always encrypted

System Area (FW, TCG Tables, etc.)

Shadow MBR (128 MB)

DataStore table (128 KB)

Global Range*
Access control to unlock assigned by Admin

**ENTERPRISE**

User Data Space, always encrypted

System Area (FW, TCG Tables, etc.)

DataStore (1 KB)

Global Range*
Access control to unlock Fixed and not configurable

*Opal 2.00 supports Global Range plus at least 8 configurable ranges and 8 Users

*Opalite requires only Global Range support plus 2 Users

*Enterprise SSC requires only Global Range support

The Opal family serves straightforward use cases while scaling for enhanced configurability.

# Opal Family and Enterprise SSC Features

| Feature | Opal V2.00 SSC | Opalite SSC<br>(Opal 2.00 subset) | Pyrite SSC (Non-encrypting version of Opalite) | Enterprise SSC |
|---|---|---|---|---|
| Core Spec Version Supported | V2.00 | V2.00 | V2.00 | V1.00 r0.9 (DRAFT) |
| Activation and Life Cycle | Yes | Yes | Yes | No |
| Number of Admins/Users | 4 Admin,<br>8 User | 1 Admin,<br>2 User | 1 Admin,<br>2 User | 1 “Bandmaster”,<br>1 “Erasemaster” (No Admin supported) |
| Min Number of Required LBA Ranges | Global Range + 8 | Global Range only | Global Range only | Global Range only |
| Min DataStore Size (General Purpose Storage) | 10MB | 1MB | 1MB | 1KB |
| Min MBR Table Size | 128MB | 128MB | 128MB (Optional) | 0 MB<br>(no pre-boot authentication support) |
| Configurable Access Control | Yes | Yes | Yes | No |
| PSID | Optional (Required in v2.01) | Required | Not Required (recommended as Prohibited due to lack of integrated data sanitization) | Not Supported |
| Media Encryption | Required | Required | Prohibited | Required |
| Crypto Erase | Revert, RevertSP, GenKey methods for device and locking range level erase granularity | Revert, RevertSP, GenKey methods for device and locking range level erase granularity | No user data erase supported – relies on native interface erase capability | Erase method |

Aligning on Opal across NVMe use cases, form factors, etc. enables a single configurable, scalable solution to address the widest variety of use cases in a common way.

# Alternatives – ATA Security

| Capability | ATA Security | Opal |
|---|---|---|
| Simple access control using a User password | ✓ | ✓ |
| Specified to require industry grade AES cipher for data-at-rest protection | X | ✓ |
| Remote management | X | ✓ |
| Extensibility to other security usage models | X | ✓ |
| Specified support for Crypto Erase | X | ✓ |
| “Purge” level erase as specified by NIST SP 800-88 | X | ✓ |

ATA Security – the “hard drive password”

- Not specified with support for media encryption
- BIOS management only by design (i.e., no OS component)
- Limited extensibility to address additional threats/usages

# Summary

The Opal family of specifications provide an established means of enabling security functionality, scalable across market segments and form factors.

TCG Storage WG is committed to engaging with NVM Express to support interactions with new features, and to meet necessary requirements; and to continuing to grow the TCG Storage specifications to expand the current set of use cases.

# References

Trusted Computing Group:

- http://www.trustedcomputinggroup.org/

White Paper:  TCG Storage, Opal, and NVMe

- http://www.trustedcomputinggroup.org/resources/tcg_storage_opal_and_nvme

TCG Storage Specifications:

- http://www.trustedcomputinggroup.org/developers/storage/specifications

1667:

- http://www.ieee1667.com/

# Thank You

nvm EXPRESS ™

*Architected for Performance*