TRUSTED® COMPUTING GROUP

SPECIFICATION

TCG Storage Opal Family Feature Set: Shadow MBR for Multiple Namespaces

Version 1.00
Revision 1.21
May 12, 2020

Contact: admin@trustedcomputinggroup.org

PUBLIC REVIEW

Work in Progress

This document is an intermediate draft for comment only and is subject to change without notice. Readers should not design products based on this document.

DISCLAIMERS, NOTICES, AND LICENSE TERMS

CHANGE HISTORY

REVISION	DATE	DESCRIPTION
Version 1.00 Draft 1.00	January 22, 2019	• Initial Release of Version 1.00.
Version 1.00 Draft 1.01	February 28, 2019	• Incorporated comments at January F2F meeting.
Version 1.00 Draft 1.02	April 22, 2019	• Incorporated comments at March F2F meeting.
Version 1.00 Draft 1.03	April 22, 2019	• Incorporated editorial comments at April F2F meeting.
Version 1.00 Draft 1.04	June 5, 2019	• Incorporated editorial changes received at WG vote.
Version 1.00 Draft 1.05	June 11, 2019	• Incorporated editorial comments from TC review.
Version 1.00 Draft 1.06	June 26, 2019	• Add language in Section 4.1.1.1.1 to incorporate comments from TC review.
Version 1.00 Draft 1.07	September 3, 2019	• Add informative content in Section 4.1.2.1.
Version 1.00 Draft 1.08	October 8, 2019	• Incorporated technical comments at September F2F meeting.
Version 1.00 Draft 1.09	October 16, 2019	• Incorporated technical comments at October F2F meeting.
Version 1.00 Draft 1.10	October 30, 2019	• Incorporated editorial changes discussed at October F2F meeting.
Version 1.00 Draft 1.11	December 2, 2019	• Incorporated changes discussed at Nov conference call.
Version 1.00 Draft 1.12	December 3, 2019	• Incorporated changes discussed at Dec F2F meeting.
Version 1.00 Draft 1.13	March 2, 2020	• Incorporated comments from TC review.
Version 1.00 Draft 1.14	March 17, 2020	• Incorporated Jim's editorial update proposals.
Version 1.00 Draft 1.15	March 17, 2020	• Incorporated Artem's comments.
Version 1.00 Draft 1.16	March 18, 2020	• Incorporated Artem's comments.
Version 1.00 Draft 1.17	March 25, 2020	• Editorial clean up.
Version 1.00 Draft 1.18	March 31, 2020	• Reformatting based on the format TCG Admin provided when r1.12 public review.
Version 1.00 Draft 1.19	April 14, 2020	• Incorporated TC's comments.
Version 1.00 Draft 1.20	April 15, 2020	• Incorporated changes discussed at the conference call on Apr 15.
Version 1.00 Draft 1.21	May 12, 2020	• Incorporated TC's comments.

CONTENTS

1 Introduction

1.1 Document Purpose and Scope

The Storage Workgroup specifications provide a comprehensive architecture for putting Storage Devices under policy control as determined by the trusted platform host, the capabilities of the Storage Device to conform with the policies of the trusted platform, and the lifecycle state of the Storage Device as a Trusted Peripheral.

1.2 Intended Audience

This specification defines the Shadow MBR for Multiple Namespaces feature set for the Opal Family Security Subsystem Classes (SSCs). Any Storage Device that claims Opal Family SSCs Shadow MBR for Multiple Namespaces feature set compatibility SHALL conform to this specification.

The intended audience for this specification is both trusted Storage Device manufacturers and developers that want to use these Storage Devices in their systems.

This document assumes familiarity and working knowledge of [2] [3] [4] [5] [6] [7] [8] [9] [10] [11].

1.3 Document References

[1] Internet Engineering Task Force (IETF), "Key words for use in RFCs to Indicate Requirement Levels" (RFC 2119)

[2] TCG Storage Architecture Core Specification, Version 2.01

[3] TCG Storage Interface Interactions Specification, Version 1.08

[4] TCG Storage Security Subsystem Class: Opal, Version 1.00

[5] TCG Storage Security Subsystem Class: Opal, Version 2.00

[6] TCG Storage Security Subsystem Class: Opal, Version 2.01

[7] TCG Storage Security Subsystem Class: Opalite, Version 1.00

[8] TCG Storage Security Subsystem Class: Pyrite, Version 1.00

[9] TCG Storage Security Subsystem Class: Pyrite, Version 2.00

[10] TCG Storage Security Subsystem Class: Ruby, Version 1.00

[11] NVM Express, Inc., “NVM Express”, Revision 1.3

1.4 Key Words

Key words are used to signify SSC requirements.

The Key Words “**SHALL**”, “**SHALL NOT**”, “**SHOULD**,” and “**MAY**” are used in this document. These words are a subset of the RFC 2119 key words used by TCG, and have been chosen since they map to key words used in T10/T13 specifications. These key words are to be interpreted as described in [1].

In addition to the above key words, the following are also used in this document to describe the requirements of particular features, including tables, methods, and usages thereof.

- **Mandatory (M):** When a feature is Mandatory, the feature SHALL be implemented. A Compliance test SHALL validate that the feature is operational.
- **Optional (O):** When a feature is Optional, the feature MAY be implemented. If implemented, a Compliance test SHALL validate that the feature is operational.
- **Excluded (X):** When a feature is Excluded, the feature SHALL NOT be implemented. A Compliance test SHALL validate that the feature is not operational.

- **Not Required (N)** When a feature is Not Required, the feature MAY be implemented. No Compliance test is required.

1.5 Conventions

1.5.1 Informative Text

Informative text is used to provide background and context. Informative text does not define requirements. Informative text is formatted as follows:

Begin Informative Text

Hello World!

End Informative Text

1.5.2 Precedence

The order of precedence to resolve conflicts between text, tables, or figures is text, then tables, then figures.

1.5.3 Lists

If the item in a list is not a complete sentence, the first word in the item is not capitalized. If the item in a list is a complete sentence, the first word in the item is capitalized.

Each item in a list ends with a semicolon, except the last item, which ends in a period. The next to the last entry in the list ends with a semicolon followed by an “and” or an “or” (i.e., “…; and”, or “…; or”). The “and” is used if all the items in the list are required. The “or” is used if only one or more items in the list are required.

Lists sequenced by letters show no ordering among the listed items. The leftmost level uses lower case letters and the next level uses capital letters. The following list shows no ordering among the named items:

a) oak;
b) maple; and
c) soft wood:
    A) pine; or
    B) cedar.

List sequenced by numbers show an ordering relationship among the listed items. All levels use Arabic numerals. The following list shows an ordered relationship among the named items:

1) hydrogen;
2) helium; and
3) lithium:
    1) lithium-6; and
    2) lithium-7.

1.5.4 Table Legend

The following legend defines SP table cell coloring coding, with the RGB values for the shading of each cell indicated in parentheses. This color coding is informative only. The table cell content is normative.

Table 1 SP Table Legend

Table Cell Legend	R-W	Value	Access Control	Comment
Arial-Narrow (230, 230, 230)	Read-only	Configurable Namespace Locking Feature Set specified	Fixed	• Cell content is Read-Only. • Access control is fixed. • Value is specified by the Configurable Namespace Locking Feature Set
Arial Narrow bold-under (230, 230, 230)	Read-only	VU	Fixed	• Cell content is Read-Only. • Access Control is fixed. • Values are Vendor Unique (VU). A minimum or maximum value may be specified.
Arial-Narrow (0, 0, 0)	Not Defined	(N)	Not Defined	• Cell content is (N). • Access control is not defined. • Any text in table cell is informative only. • A `Get` MAY omit this column from the method response.
Arial Narrow bold-under (179, 179, 179)	Write	Preconfigured, user personalizable	Preconfigured, user personalizable	• Cell content is writable. • Access control is personalizable • `Get` Access Control is not described by this color coding
Arial-Narrow (179, 179, 179)	Write	Preconfigured, user personalizable	Fixed	• Cell content is writable. • Access control is fixed. • `Get` Access Control is not described by this color coding

1.5.5 Fonts

Names of methods and SP tables are in Courier New font (e.g., the `Set` method, the `Locking` table). This convention does not apply to method and table names appearing in headings or captions.

1.6 Document Precedence

In the event of conflicting information in this specification and other documents, the precedence for requirements is:

1. This specification;
2. TCG Storage Architecture Core Specification [2];

3. TCG Storage Interface Interactions Specification [3];
4. TCG Storage Security Subsystem Class: [4] [5] [6], [7], [8], [9], [10]; and
5. NVM Express 1.3 [11]

1.7 Dependencies on Other Feature Sets

This document has no dependencies on other feature sets.

1.8 Interactions with Other Feature Sets

This document has no interactions with other feature sets.

1.9 Terminology

This document does not define any new terms.

2 Overview

2.1 Overview

When MBR shadowing (see [2]) is supported by the TPer and there is more than one namespace exists in the NVM subsystem, there are two use cases:

a) it is shared by all namespaces and controllers within the NVM subsystem; and

b) it is applied to one namespace and controller within the NVM subsystem only.

SIIS (see [3]) covers the first use case. The purpose of this feature set is to cover the second use case and to allow the host to specify to which namespace MBR Shadow is applicable.

3 Feature Set Requirements

This section defines the Mandatory (M) and Optional (O) requirements for the Shadow MBR for Multiple Namespaces feature set, when it is implemented in an Opal-compliant device.

3.1 Level 0 Discovery

A Storage Device implementing the Shadow MBR for Multiple Namespaces feature set SHALL:

a) return the Namespace Feature Descriptor as defined in section 3.1.1; and
b) support the Level 0 Discovery response requirements defined in [4], [5], [6], [7], [8], [9] or [10].

3.1.1 Shadow MBR for Multiple Namespaces Feature Descriptor (Feature Code = 0x0407) (M)

This feature descriptor SHALL be returned when the Storage Device supports the Shadow MBR for Multiple Namespaces feature set. The contents of the feature descriptor are defined in Table 2.

Table 2 Level 0 Discovery – Shadow MBR for Multiple Namespaces Feature Descriptor

Byte \ Bit	7	6	5	4	3	2	1	0
0	(MSB)			Feature Code (0x0407)				
1								(LSB)
2	Version				Reserved			
3	Length							
4	Reserved							ANS_C
5-15	Reserved							

3.1.1.1 Feature Code

0x0407

3.1.1.2 Version

This field indicates 0x1 or any version that supports the features described in this specification.

3.1.1.3 Length

This field indicates the number of bytes in the descriptor following byte 3. The value SHALL be set to 0x0C.

3.1.1.4 ANS_C

The All Namespace Capable (ANS_C) field is set to one to indicate that the Storage Device supports a value of 0xFFFF_FFFF for the NamespaceID column value of the `MBRControl` table. The ANS_C field is set to zero to indicate that the Storage Device does not support a column value of 0xFFFF_FFFF for the Namespace ID column value of the `MBRControl` table.

4 SSC Specific Functionality

This section specifies the additional SSC-specific functionality (not contained in [4], [5], [6], [7], [8], [9] or [10]).

4.1 Tables

This section defines new tables and modifications to existing tables required for this feature set.

4.1.1 Modified Tables

This feature set modifies the following tables:

a) `MBRControl`.

b) `ACE`.

4.1.1.1 MBRControl

This feature set modifies the `MBRControl` table by adding the following columns, in addition to those defined in [2]:

Table 3 LockingSP – MBRControl Table Column

Column Number	Column Name	IsUnique	Column Type
0x04	NamespaceID		bytes_4

4.1.1.1.1 NamespaceID (M)

This column value identifies the namespace to which the MBR Shadow is applied. The initial NamespaceID column value SHALL be either 0x0000_0000 or the Namespace Identifier of the existing namespace.

Support for the value of 0xFFFF_FFFF in the Namespace ID column is Optional. If the Storage Device reports a value of one in the ANS_C field of the Shadow MBR for Multiple Namespaces Feature Descriptor (see Section 3.1.1), then the initial value of the NamespaceID column MAY also be set to 0xFFFF_FFFF.

If this column value is equal to 0xFFFF_FFFF, then the `MBR` and the `MBRControl` tables in the Locking SP are shared by all namespaces and controllers within the NVM subsystem as defined in Section 5.6.1.4.6 in [3].

If this column value is equal to the value of the Namespace Identifier of an existing namespace, the `MBR` and the `MBRControl` tables in the Locking SP are applied to that namespace only.

4.1.1.2 ACE

This feature set modifies ACE table in the Locking SP as follows:

Table 4 Locking SP – ACE Table Preconfiguration

Table Association -Informative Column	UID	Name	CommonName	BooleanExpr	Columns
MBRControl					

Table Association -Informative Column	UID	Name	CommonName	BooleanExpr	Columns
	00 00 00 08 00 03 F8 00	"ACE_MBRControl_Admin_Set"		Admins	Enable, Done, DoneOnReset, NamespaceID

4.1.2 Modified Method

4.1.2.1 Set

If the `Set` method is invoked on the NamespaceID column of the `MBRControl` table and its value is equal to the Namespace Identifier of the non-existing namespace except when the value of the Namespace Identifier is 0x0000_0000, then the `Set` method SHALL fail with a status of INVALID_PARAMETER.

If the `Set` method is invoked on the NamespaceID column of the `MBRControl` table and the value of the Enabled column of the `MBRControl` table is TRUE, then the `Set` method SHALL fail with a status of INVALID_PARAMETER.

If the Storage Device reports a value of zero in the ANS_C field of the Shadow MBR for the Multiple Namespaces Feature Descriptor (see Section 3.1.1), and the `Set` method is invoked on the NamespaceID column of the `MBRControl` table and its value is 0xFFFF_FFFF, then the `Set` method SHALL fail with a status of INVALID_PARAMETER.

If the NamespaceID column value of the `MBRControl` table is 0x0000_0000, and the `Set` method is invoked on Enable column of the `MBRControl` table and its value is TRUE, then the `Set` method SHALL fail with a status of INVALID_PARAMETER.

If the Set method is invoked with the Enabled column value set to TRUE on the MBRControlObj and the LBA Format of the namespace corresponding to the value of NamespaceID column of `MBRControl` table is incompatible with the content of `MBR` table, then the Set method MAY fail with a status of INCOMPATIBLE_MBR_FORMAT.

5 Interaction with the Namespace Management Command and the Format NVM command

SIIS (see [3]) specifies that the `MBR` and the `MBRControl` tables in the Locking SP are shared by all namespaces and controllers within the NVM subsystem. It also defines the interactions with the Namespace Management Command and the Format NVM command.

This feature set modifies interactions with the Namespace Management command and the Format NVM command as follows:

If:

a) the Select (SEL) field of a Namespace Management command is Delete; and

b) the Namespace Identifier (NSID) field of that command is equal to the value of the NamespaceID column of the `MBRControl` table,

then that command SHALL fail with a status of Operation Denied.

If:

a) the Enabled column value of the MBRControlObj is TRUE; and

b) a Format NVM command specifies an LBA Format (see [11]) of the namespace corresponding to the value of the NamespaceID column of the `MBRControl` table that is different from the original LBA Format of that namespace,

then that command SHALL fail with a status of Invalid Security State.

If:

a) the NamespaceID column value of the MBRControlObj is 0xFFFF_FFFF;

b) the Enabled column value of the MBRControlObj is TRUE; and

c) the Select (SEL) field of a Namespace Management command is Create and specifies an LBA Format (see [11]) which is different from any existing namespace,

then the Namespace Management command SHALL fail with a status of Operation Denied.

6 Modifications to Core Specification

The Core Specification defines the Storage Device response for attempts by the host to read or write user data (see [2]). This feature set overwrites the response for attempts by the host to read or write user data as follows:

If the value of the NamespaceID column in the `MBRControl` table is not equal to 0xFFFFFFFF, then the device response for all cases when the host attempts to read user data is specified in Table 5.

If the value of the NamespaceID column in the `MBRControl` table is not equal to 0xFFFFFFFF, then the device response for all cases when the host attempts to write user data is specified in Table 6.

Table 5 Interface Read Command Access

MBRControl Enable	MBRControl Done	LBA belonging to Namespace equal to the value of NamespaceID column in MBRControl table	Starting LBA Within MBR	Ending LBA within MBR	ReadLockEnabled for Requested LBA range	ReadLocked for Requested LBA Range	Required Behavior
True	False	True	True	True	N/A	N/A	Return Data from MBR table
True	False	True	True	False	N/A	N/A	Transfer no data to the host and terminate the command with a "Data Protection Error" (see [3])
True	False	True	False	False	False	N/A	Return user data
True	False	True	False	False	True	False	Return user data
True	False	True	False	False	True	True	Return all zeroes.
True	False	True	False	False	True	Mixed (when crossing range boundaries)	Transfer no data to the host and terminate the command with a "Data Protection Error" (see [3])
True	False	False	N/A	N/A	False	N/A	Return user data
True	False	False	N/A	N/A	True	False	Return user data
True	False	False	N/A	N/A	True	True	Transfer no data to the host and terminate the command with a "Data Protection Error" (see [3])
True	False	False	N/A	N/A	True	Mixed (when crossing range boundaries)	Transfer no data to the host and terminate the command with a "Data Protection Error" (see [3])
True	True	N/A	N/A	N/A	False	N/A	Return user data
True	True	N/A	N/A	N/A	True	False	Return user data
True	True	N/A	N/A	N/A	True	True	Transfer no data to the host and terminate the command with a "Data Protection Error" (see [3])
True	True	N/A	N/A	N/A	True	Mixed (when crossing range boundaries)	Transfer no data to the host and terminate the command with a "Data Protection Error" (see [3])
False	N/A	N/A	N/A	N/A	False	N/A	Return user data
False	N/A	N/A	N/A	N/A	True	False	Return user data

MBRControl Enable	MBRControl Done	LBA belonging to Namespace equal to the value of NamespaceID column in MBRControl table	Starting LBA Within MBR	Ending LBA within MBR	ReadLockEnabled for Requested LBA range	ReadLocked for Requested LBA Range	Required Behavior
False	N/A	N/A	N/A	N/A	True	True	Transfer no data to the host and terminate the command with a "Data Protection Error" (see [3])
False	N/A	N/A	N/A	N/A	True	Mixed (when crossing range boundaries)	Transfer no data to the host and terminate the command with a "Data Protection Error" (see [3])

Table 6 Interface Write Command Access

MBRControl Enable	MBRControl Done	LBA belonging to Namespace equal to the value of NamespaceID column in MBRControl table	Starting LBA Within MBR	Ending LBA within MBR	WriteLockEnabled for Requested LBA range	WriteLocked for Requested LBA Range	Required Behavior
True	False	True	True	N/A	N/A	N/A	Transfer no data from the host and terminate the command with a "Data Protection Error" (see [3])
True	False	False	N/A	N/A	False	N/A	Write user data
True	False	False	N/A	N/A	True	False	Write user data
True	False	False	N/A	N/A	True	True	Transfer no data to the host and terminate the command with a "Data Protection Error" (see [3])
True	False	False	N/A	N/A	True	Mixed (when crossing range boundaries)	Transfer no data to the host and terminate the command with a "Data Protection Error" (see [3])
True	False	True	False	False	False	N/A	Write user data
True	False	True	False	False	True	False	Write user data
True	False	True	False	False	True	True	Transfer no data from the host and terminate the command with a "Data Protection Error" (see [3])
True	False	True	False	False	True	Mixed (when crossing range boundaries)	Transfer no data from the host and terminate the command with a "Data Protection Error" (see [3])

MBRControl Enable	MBRControl Done	LBA belonging to Namespace equal to the value of NamespaceID column in MBRControl table	Starting LBA Within MBR	Ending LBA within MBR	WriteLockEnabled for Requested LBA range	WriteLocked for Requested LBA Range	Required Behavior
True	True	N/A	N/A	N/A	False	N/A	Write user data
True	True	N/A	N/A	N/A	True	False	Write user data
True	True	N/A	N/A	N/A	True	True	Transfer no data from the host and terminate the command with a "Data Protection Error" (see [3])
True	True	N/A	N/A	N/A	True	Mixed (when crossing range boundaries)	Transfer no data from the host and terminate the command with a "Data Protection Error" (see [3])
False	N/A	N/A	N/A	N/A	False	N/A	Write user data
False	N/A	N/A	N/A	N/A	True	False	Write user data
False	N/A	N/A	N/A	N/A	True	True	Transfer no data from the host and terminate the command with a "Data Protection Error" (see [3])
False	N/A	N/A	N/A	N/A	True	Mixed (when crossing range boundaries)	Transfer no data and terminate the command with a "Data Protection Error" (see [3])