

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

Version 1.00
Revision 1.06
June 26, 2019

Contact: admin@trustedcomputinggroup.org

DRAFT

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

THIS SPECIFICATION IS PROVIDED “AS IS” WITH NO WARRANTIES WHATSOEVER, INCLUDING ANY WARRANTY OF MERCHANTABILITY, NONINFRINGEMENT, FITNESS FOR ANY PARTICULAR PURPOSE, OR ANY WARRANTY OTHERWISE ARISING OUT OF ANY PROPOSAL, SPECIFICATION OR SAMPLE.

Without limitation, TCG disclaims all liability, including liability for infringement of any proprietary rights, relating to use of information in this specification and to the implementation of this specification, and TCG disclaims all liability for cost of procurement of substitute goods or services, lost profits, loss of use, loss of data or any incidental, consequential, direct, indirect, or special damages, whether under contract, tort, warranty or otherwise, arising in any way out of use or reliance upon this specification or any information herein.

This document is copyrighted by Trusted Computing Group (TCG), and no license, express or implied, is granted herein other than as follows: You may not copy or reproduce the document or distribute it to others without written permission from TCG, except that you may freely do so for the purposes of (a) examining or implementing TCG specifications or (b) developing, testing, or promoting information technology standards and best practices, so long as you distribute the document with these disclaimers, notices, and license terms.

Contact the Trusted Computing Group at www.trustedcomputinggroup.org for information on specification licensing through membership agreements.

Any marks and brands contained herein are the property of their respective owners.

DRAFT

CHANGE HISTORY

REVISION	DATE	DESCRIPTION
Version 1.00 Draft 1.00	January 22, 2019	<ul style="list-style-type: none">Initial Release of Version 1.00.
Version 1.00 Draft 1.01	February 28, 2019	<ul style="list-style-type: none">Incorporated comments at January F2F meeting.
Version 1.00 Draft 1.02	April 22, 2019	<ul style="list-style-type: none">Incorporated comments at March F2F meeting.
Version 1.00 Draft 1.03	April 22, 2019	<ul style="list-style-type: none">Incorporated editorial comments at April F2F meeting.
Version 1.00 Draft 1.04	June 5, 2019	<ul style="list-style-type: none">Incorporated editorial changes received at WG vote.
Version 1.00 Draft 1.05	June 11, 2019	<ul style="list-style-type: none">Incorporated editorial comments from TC review.
Version 1.00 Draft 1.06	June 26, 2019	<ul style="list-style-type: none">Add languages in Section 4.1.1.1.1 to incorporated comments from TC review.

DRAFT

CONTENTS

DISCLAIMERS, NOTICES, AND LICENSE TERMS	1
CHANGE HISTORY	2
1 Introduction	5
1.1 Document Purpose and Scope.....	5
1.2 Intended Audience	5
1.3 Document References	5
1.4 Key Words.....	5
1.5 Document Precedence	6
1.6 Dependencies on Other Feature Sets	6
1.7 Interactions with Other Feature Sets	6
1.8 Terminology	6
2 Overview.....	7
2.1 Overview	7
3 Feature Set Requirements	8
3.1 Level 0 Discovery.....	8
3.1.1 Shadow MBR for Multiple Namespace Feature Descriptor (Feature Code = 0x0407) (M).....	8
4 SSC Specific Functionality	9
4.1 Tables	9
4.1.1 Modified Tables	9
4.1.2 Modified Mothod.....	10
5 Interaction with the Namespace Management Command.....	11
6 Modifications to Core Specification	12

Table 1 Level 0 Discovery – Shadow MBR for Multiple Namespace Feature Descriptor 8

Table 2 LockingSP – MBRControl Table Column..... 9

Table 3 Locking SP – ACE Table Preconfiguration 9

Table 4 Interface Read Command Access 12

Table 5 Interface Write Command Access..... 13

DRAFT

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 Namespace Feature Set for the Opal Family Security Subsystem Classes (SSCs). Any Storage Device that claims Opal Family SSCs Shadow MBR for Multiple Namespace 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] TCG Storage Feature Set: Configurable Namespace Locking, Version 1.00
- [12] 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 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 [12]

1.6 Dependencies on Other Feature Sets

This feature set has no dependencies on other feature sets.

1.7 Interactions with Other Feature Sets

This feature set has interactions with Configurable Namespace Locking feature set.

1.8 Terminology

This feature does not define any new terms.

DRAFT

2 Overview

2.1 Overview

When MBR shadowing (see [2]) is supported by the TPer and 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.
- b) It is applied to one namespace and controller within the NVM subsystem only.

SIS (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.

DRAFT

3 Feature Set Requirements

This section defines the Mandatory (M) and Optional (O) requirements for the Shadow MBR for Multiple Namespace 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 Namespace 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 Namespace Feature Descriptor (Feature Code = 0x0407) (M)

This feature descriptor SHALL be returned when the Storage Device supports the Shadow MBR for Multiple Namespace feature set. The contents of the feature descriptor are defined in Table 1.

Table 1 Level 0 Discovery – Shadow MBR for Multiple Namespace Feature Descriptor

Byte	Bit	7	6	5	4	3	2	1	0
0	(MSB)	Feature Code							
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 column value of 0xFFFF_FFFF for the Namespace ID 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 [2], [4], [5], [6], [7], [8] or [9]).

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 2 LockingSP – MBRControl Table Column

Column Number	Column Name	IsUnique	Column Type
0x04	NamespaceID		bytes_4

4.1.1.1.1 Namespace ID (M)

This column value identifies Namespace to which MBR Shadow belongs. 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` for the Namespace ID column is Optional. See 4.1.2.1 for the expected behavior when `0xFFFF_FFFF` is specified for this column using the Set method. If the Storage Device reports a value of one in the `ANS_C` field of the Shadow MBR for Multiple Namespace Feature Descriptor (see 3.1.1), then the initial `NamespaceID` column value MAY also be set to `0xFFFF_FFFF`.

If this column value is equal to `0xFFFF_FFFF`, then the `MBR` and `MBRControl` tables in the Locking SP are shared by all namespaces and controllers within the NVM subsystem as defined in Section 5.5.1.4.6 in [3].

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

4.1.1.2 ACE

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

Table 3 Locking SP – ACE Table Preconfiguration

Table Association -Informative Column	UID	Name	CommonName	BooleanExpr	Columns
<i>MBRControl</i>					

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 `NamespaceID` column of `MBRControl` table and its value is equal to the Namespace Identifier of the non-existing Namespace except `0x0000_0000`, 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 Multiple Namespace Feature Descriptor (see 3.1.1), and the `Set` method is invoked on `NamespaceID` column of `MBRControl` table and its value is `0xFFFF_FFFF`, then the `Set` method SHALL fail with a status of `INVALID_PARAMETER`.

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

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

DRAFT

5 Interaction with the Namespace Management Command

SIIS (see [3]) specifies that the MBR and 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.

In addition to the requirements of interactions with Namespace Management Command defined in SIIS, this feature set modifies interactions with the Namespace Management Command as follows;

If:

- a) the Select (SEL) field of the command is Delete;
- b) the Namespace Identifier (NSID) field of the command is equal to the value of NamespaceID column of MBRControl table

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

DRAFT

6 Modifications to Core Specification

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 NamespaceID column in 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 4.

If the value of NamespaceID column in 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 5.

Table 4 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 5 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])