

# ENDL TEXAS

Date: 27 March 2008  
 To: T10 Technical Committee  
 From: Ralph O. Weber  
 Subject: Fixes for seven OSD-2 bugs

## Revision History

- r0 Initial revision
- r1 Modified change 2 to swap highest and lowest as requested by the SNIA OSD TWG.
- r2 Add fixing a boot epoch bug (see change 7) which Rob Elliott identified last November.

Differences between r1 and r2 are marked with change bars.

Unless otherwise indicated additions are shown in [blue](#), deletions in ~~red-strikethrough~~, and comments in [green](#).

## Proposed Changes

### Change 1 – In what user objects can attributes be accessed?

#### Description

The following statement no longer describes reality:

An OSD command may only retrieve or store attributes in the attributes pages associated with the OSD object addressed by the command.

Most multi-object commands as well as the LIST command and LIST COLLECTION command provide mechanisms for accessing the attributes of user objects for which the mode of object addressing might most charitably be described as indirectly.

#### Proposed changes in OSD-2 r03

##### 4.7.1 Overview

...

The LIST command (see 6.15) and LIST COLLECTION command (see 6.16) allow application clients to retrieve attributes in the attributes pages associated with the objects being listed. Commands that use collections to affect multiple user objects (see 4.6.6.2) allow application clients to retrieve and/or store attributes in the attributes pages associated each object in the collection to which the command is addressed. Otherwise, an ~~An~~ OSD command may only retrieve or store attributes in the attributes pages associated with the OSD object addressed by the command.

## Change 2 – Conflicting requirements

### Description

The definition of the CREATE command contains conflicting requirements about which User\_Object\_ID is to be returned when multiple user objects are created. Also, the incorrect attribute name is referenced.

### Proposed changes in OSD-2 r03

## 6.4 CREATE

...

The NUMBER OF USER OBJECTS field specifies the number of user objects to be created. If the NUMBER OF USER OBJECTS field contains zero or one, one user object shall be created. Otherwise:

- a) The number of user objects created shall equal the value in the NUMBER OF USER OBJECTS field;
- b) The user objects created shall be assigned consecutive valued User\_Object\_IDs; and
- c) The lowest valued User\_Object\_ID shall be placed in the ~~created User\_Object\_ID~~ Collection\_Object\_ID or User\_Object\_ID attribute of the Current Command attributes page (see 7.1.2.29).

If the NUMBER OF USER OBJECTS field contains a value that is greater than one and the requested User\_Object\_ID is not zero, the command shall be terminated with a CHECK CONDITION status, with the sense key set to ILLEGAL REQUEST and the additional sense code set to INVALID FIELD IN CDB.

The command shall be terminated with a CHECK CONDITION status, with the sense key set to ILLEGAL REQUEST and the additional sense code set to INVALID FIELD IN CDB if:

- a) If the NUMBER OF USER OBJECT field contains a value that is greater than one;
- b) The GET/SET CDBFMT field contains 10b; and
- c) The GET ATTRIBUTES PAGE field (see 5.2.4.3) contains a value other than FFFF FFEh (i.e., the Current Command attributes page).

...

{{This text is needed to cover the case where the number of user objects field contains zero or one.}}

The ~~highest~~ lowest valued assigned User\_Object\_ID shall be placed in the Collection\_Object\_ID or User\_Object\_ID attribute in the Current Command attributes page (see 7.1.2.29).

### Change 3 – The CREATE PARTITION command does not initialize secret keys

#### Description

Clarification is needed to show that the CREATE PARTITION command does not initialize the secret keys needed by security methods other than NOSEC.

#### Proposed changes in OSD-2 r03

{{Add the following somewhere near the end of 6.7 (CREATE PARTITION).}}

The CREATE PARTITION command does not initialize the partition key or the working keys (see 4.12.9.1) for the new partition. Proper operation of any security method other than NOSEC (see 4.12.4) requires that the following commands be processed without errors before other commands are addressed to the partition:

- a) A SET KEY command (see 6.29) that establishes the partition key; and
- b) One or more SET KEY commands that establish one or more working keys for the partition.

A CREATE PARTITION command whose capability (see 4.11.2.2) has the SET\_ATTR bit set to one and POL/SEC bit set to one is allowed to avoid the need for SET KEY commands by setting the default security method attribute to NOSEC in the Partition Policy/Security attributes page (see 7.1.2.22) for the created partition.

## Change 4 – The interactions between multi-object commands and timestamps are not specified

### Description

How the processing of a multi-object command affects timestamps is not specified. The following specification is proposed.

### Proposed changes in OSD-2 r03

#### 7.1.2.17 Collection Timestamps attributes page

...

The data accessed time attribute (number 4h) shall contain the value of the clock attribute in the Root Information attributes page at the completion of the most recent LIST COLLECTION command (see 6.16) that transferred a list of user objects in the collection to the application client.

The data modified time attribute (number 5h) shall contain the value of the clock attribute in the Root Information attributes page at the completion of the most ~~recent: recent set attributes command function to a user object Collections attributes page (see 7.1.2.19) that added or removed a member from the collection.~~

- a) Set attributes command function to a user object Collections attributes page (see 7.1.2.19) that added or removed a member from the collection; or
- b) Multi-object command (4.6.6.2) that removed a user object from the collection.

...

#### 7.1.2.18 User Object Timestamps attributes page

...

The attributes accessed time attribute (number 2h) shall contain the value of the clock attribute in the Root Information attributes page at the completion of the most ~~recent: recent command whose CDB get and set attributes parameters (see 5.2.4) transferred any attributes pages or values associated with the user object to the application client.~~

- a) Command whose CDB get and set attributes parameters (see 5.2.4) transferred any attributes pages or values associated with the user object to the application client; or
- b) QUERY command (see 6.21) that evaluated any attributes values associated with the user object to the application client during attempts to match the values specified in the query list.

{{All other cases for multi-object commands appear to be covered by existing text because their access to and modifications of user object attributes are described in the CDB get and set attributes parameters.}}

## Change 5 – The force collection removal definition does not account for type 1 collections

### Description

Type 1 collections do not have 'back links' like Type 0 collections have. The definition of the FCR bit in the REMOVE COLLECTION command needs to be updated to account for this difference.

### Proposed changes in OSD-2 r03

## 6.23 REMOVE COLLECTION

...

The FCR (force collection removal) bit specifies the actions to be taken if the collection contains user objects. If the FCR bit is set to zero and the collection contains user objects, the command shall be terminated with a CHECK CONDITION status, the sense key shall be set to ILLEGAL REQUEST, and the additional sense code shall be set to PARTITION OR COLLECTION CONTAINS USER OBJECTS. If the FCR bit is set to one, the collection shall be removed as follows even if it contains user objects:

- 1) The user objects in the collection shall be processed as follows based on the contents of collection type attribute in the Collection Information attributes page (see 7.1.2.19):
  - A) If the collection type attribute is set to zero, each ~~Each~~ user object in the collection shall be modified to indicate that the user object no longer is a member of the collection; or
  - B) If the collection type attribute is set to a value other than zero, the user objects in the collection shall not be modified;
 and
- 2) The collection shall be removed.

## Change 6 – Multi-object processing accounts for replaced user objects but not for removed ones

### Proposed changes in OSD-2 r03

#### 4.6.6.2 Commands that use collections to affect multiple user objects

...

Each user object in the specified collection shall be processed as follows:

- 1) If the user object has been removed, the specified operation shall not be performed on the non-existent user object. This shall not be considered to be an error;
- 2) If user object has not been removed and the creation time attribute in the User Object Timestamps attributes page (see 7.1.2.18) is earlier than or equal to (i.e., less than or equal to) the creation time attribute in the Collection Timestamps attributes page (see 7.1.2.17) (i.e., if the user object has not been replaced), then the quotas (see 4.10) that apply to the specified operation shall be evaluated and processing of the operation shall be handled as follows:
  - A) If a quota error condition is detected, the multi-object command shall be terminated as described in this subclause; or
  - B) If no quota error condition is detected, the specified operation shall be performed on the user object and whether or not an error is detected shall be noted;
- 3) If the creation time attribute in the User Object Timestamps attributes page is later than (i.e., greater than) the creation time attribute in the Collection Timestamps attributes page, then the specified operation shall not be performed on the user object. This shall not be considered to be an error; and
- 4) If no error has been detected, the user object shall be removed from the specified collection.

## Change 7 – Boot epoch is incremented too easily

### Description

The current boot epoch reference to SAM-4 device conditions includes all of the following: power on, hard reset, logical unit reset, I\_T nexus loss, and power loss expected.

Per comments from Rob Elliott:

- "I\_T nexus loss is problematic; I don't think incrementing the value every time any I\_T nexus is lost would be correct. One I\_T nexus's permissions are probably not related to the status of other I\_T nexuses. The boot epoch attribute is not an I\_T nexus specific value (I hope). I\_T nexus loss should be ignored."
- "Power loss expected is problematic. It is a very lightweight condition, causing no more than a bunch of CLEAR TASK SETs. SAM-4 letter ballot comment resolution is booting it out of the "specificity" list in table 36. Power loss expected should be ignored. (if a power loss does happen, then "power on" will increment the boot epoch later)"

The boot epoch attribute definition needs to be restricted to power on, hard reset, and logical unit reset.

### Proposed changes in OSD-2 r03

#### 7.1.2.21 Root Policy/Security attributes page

...

When the OSD device is manufactured, the boot epoch attribute (number Ah) should be set to a non-zero value. The processing of **any a power on SCSI device condition, hard reset SCSI device condition, or logical unit reset SCSI device condition** ~~(e.g., logical unit reset)~~ established in response to an event (see SAM-4) shall cause the boot epoch attribute to be updated as follows:

- a) If the boot epoch attribute is not set to FFFFh, then one shall be added to the boot epoch value; or
- b) If the boot epoch attribute is set to FFFFh, then the boot epoch value shall be set to one.