

ENDL TEXAS

Date: 26 June 2004
To: T10 Technical Committee and SNIA OSD TWG
From: Ralph O. Weber
Subject: A FLUSH for every object type

OSD Letter Ballot comments HP 76) and HP 77) note that FLUSH OBJECT is the only command for which there is not a different command for each object type.

This proposal attempts to correct that by defining FLUSH xxx commands for all object types.

The contents of this proposal represent a replacement for 6.7 (FLUSH OBJECT) in OSD r09.

Revision History

- r0 Original revision
- r1 Cleanup some inconsistencies in the definitions of the various flush scope codes, assign service action codes to the new commands, and alphabetize FLUSH OSD and FLUSH PARTITION in clause 6
- r2 Correct some field names that are not in small caps.

Differences between r0 and this revision are marked with change bars.

Detailed OSD r09 Changes

For completeness, accepted changes from the following OSD Letter Ballot comments (see T10/04-108) are included in this document: HP 34), HP 78), HP 79), Lingua 28), and Veritas 64).

Text already appearing in OSD r09 is shown in black. Text appearing in OSD r09 that is to be removed is shown in red strike through. Text to be added is shown in blue.

4.9.4.4 Credentials and commands allowed

...

Table 19 — Commands allowed by specific capability field values

Commands allowed and CDB fields whose contents are restricted by credential field contents	Capability Field values that allow a command		
	Object Type Name	Permission Bits That Are Set To One	Object Descriptor Name
...			
A FLUSH OBJECT command with the CDB PARTITION_ID field containing a value that matches the contents of the credential PARTITION_ID field, the CDB USER_OBJECT_ID field containing a value that matches the contents of the object descriptor SINGLE OBJECT_ID field, and the USER_OBJECT_ID field not containing a Collection_Object_ID.	USER	OBJ_MGMT	1OBJECT
A FLUSH OBJECT COLLECTION command with the CDB PARTITION_ID field containing a value that matches the contents of the credential PARTITION_ID field, the CDB USER_OBJECT_ID COLLECTION_OBJECT_ID field containing a value that matches the contents of the object descriptor SINGLE OBJECT_ID field, and the USER_OBJECT_ID field containing a Collection_Object_ID.	COLLECTION	OBJ_MGMT	1OBJECT
A FLUSH OBJECT PARTITION command with the CDB PARTITION_ID field containing a value that matches the contents of the object descriptor SINGLE OBJECT_ID field and the CDB-USER_OBJECT_ID field equal to zero.	PARTITION	OBJ_MGMT	1OBJECT
A FLUSH OBJECT OSD command with the CDB PARTITION_ID field and the CDB USER_OBJECT_ID field both equal to zero or a FORMAT OSD command.	ROOT	OBJ_MGMT	1OBJECT ^a
A FORMAT OSD command.	ROOT	OBJ_MGMT and GLOBAL	1OBJECT ^a
...			
Combinations of OBJECT TYPE field, PERMISSION BITS field, and OBJECT DESCRIPTOR TYPE field values not shown in this table and table 20 are reserved. The credential and capability fields not shown in this table may place additional limits on the objects that are allowed to be accessed.			
^a The object descriptor SINGLE OBJECT_ID field shall contain zero.			

4.9.5.3 Reconstructing the credential

The device server reconstructs a credential from a CDB capability by:

- 1) Reconstructing the credential fields as follows:
 - A) Copy the value in the OSD system ID attribute in the Root Information attributes page (see 7.1.2.8) to the OSD SYSTEM ID field of the reconstructed credential; and
 - B) Reconstruct the credential PARTITION_ID field as follows:
 - a) If the command is CREATE PARTITION, FLUSH PARTITION, FLUSH OSD, FORMAT OSD, LIST, PERFORM SCSI COMMAND, REMOVE PARTITION, or SET KEY, then place zero in PARTITION_ID field of the reconstructed credential;
 - b) If the CDB USER_OBJECT_ID field contains zero and the command is ~~FLUSH OBJECT~~ or PERFORM TASK MANAGEMENT, then place zero in PARTITION_ID field of the reconstructed credential; or
 - c) Otherwise, copy the contents of the CDB PARTITION_ID field to the PARTITION_ID field of the reconstructed credential;
- and
- 2) Copying the capability from the CDB to the reconstructed credential.

4.15 Reservations

...

Table 30 — OSD commands that are allowed in the presence of various reservations

OSD Command	Addressed LU has this type of persistent reservation held by another I_T Nnexus				
	From any I_T Nnexus		From registered I_T Nnexus (RR all types)	From I_T Nnexus not registered	
	Write Excl	Excl Access		Write Excl RR	Excl Access – RR
...	...				
FLUSH OBJECT	Conflict	Conflict	Allowed	Conflict	Conflict
FLUSH COLLECTION	Conflict	Conflict	Allowed	Conflict	Conflict
FLUSH OSD	Conflict	Conflict	Allowed	Conflict	Conflict
FLUSH PARTITION	Conflict	Conflict	Allowed	Conflict	Conflict
...	...				
Key: LU=Logical Unit, Excl=Exclusive, RR=Registrants Only or All Registrants					

6.1 Summary of commands for OSD type devices

...

Table 39 — Commands for OSD type devices

Command name	Operation code	Service action ^a	Type	Reference
...	...			
FLUSH OBJECT	7Fh	8808h	M	6.w
FLUSH COLLECTION	7Fh	881Ah	M	6.x
FLUSH OSD	7Fh	881Ch	M	6.y
FLUSH PARTITION	7Fh	881Bh	M	6.z
...	...			

6.w FLUSH **OBJECT**

The FLUSH **OBJECT** command (see table 45) ensures that the specified data and attribute bytes for the specified user object are written to stable storage (see 4.10).

Table 45 — FLUSH **OBJECT command**

Bit Byte	7	6	5	4	3	2	1	0
8	(MSB) _____							
9	SERVICE ACTION (8808h)							(LSB)
10	Reserved						FLUSH SCOPE	
11	Reserved		GET/SET CDBFMT		Reserved			
12	TIMESTAMPS CONTROL							
13	Reserved							
15	Reserved							
16	(MSB) _____							
23	PARTITION_ID							(LSB)
24	(MSB) _____							
31	USER_OBJECT_ID							(LSB)
32	Reserved							
51	Reserved							
52	Get and set attributes parameters (see 5.2.1)							
79	Reserved							
80	Security parameters (see 5.2.5)							
173	Reserved							

The FLUSH SCOPE field (see table 46) specifies the scope of the data and attribute bytes that are written to stable storage.

Table 46 — User object flush scope values

Value	Scope of data and attributes that shall be written to stable storage
00b	User object data and attributes
01b	User object attributes only
10b - 11b	Reserved

Table 46 — Flush scope values

Value	Scope of data and attributes that shall be written to stable storage	
	Contents of USER_OBJECT_ID field	
	Not Zero (i.e., user object)	Zero (i.e., partition)
00b	User object data and attributes	List of OSD objects contained in the partition ^a
01b	User object attributes only	Partition attributes only
10b	Reserved	a) List OSD objects contained in the partition; ^a b) Partition attributes; and c) User object data and attributes for listed OSD objects ^b
11b	Reserved	Reserved

^a For all partitions except partition zero, the listed OSD objects are user objects and collections. For partition zero, the listed OSD objects are partitions.

^b A FLUSH OBJECT command with a flush scope of 10b that is addressed to partition zero, flushes the entire OSD logical unit.

The GET/SET CDBFMT field specifies the format of the get and set attributes parameters as described in 5.2.1.

The contents of the TIMESTAMPS CONTROL field are defined in 5.2.7.

The contents of the PARTITION_ID field are defined in 5.2.4.

The contents of the USER_OBJECT_ID field are defined in 5.2.8.

The get and set attributes parameters are defined in 5.2.1. The format of the Data-In Buffer and Data-Out Buffer when attributes are being retrieved or set is described in 4.11.

The security parameters are defined in 5.2.5.

6.x FLUSH COLLECTION

The FLUSH COLLECTION command (see table 47) ensures that the attribute bytes for the specified collection are written to stable storage (see 4.10).

Table 47 — FLUSH COLLECTION command

Bit Byte	7	6	5	4	3	2	1	0
8	(MSB) _____							
9	SERVICE ACTION (881Ah) _____ (LSB)							
10	Reserved						FLUSH SCOPE	
11	Reserved		GET/SET CDBFMT		Reserved			
12	TIMESTAMPS CONTROL							
13	_____							
15	Reserved _____							
16	(MSB) _____							
23	PARTITION_ID _____ (LSB)							
24	(MSB) _____							
31	COLLECTION_OBJECT_ID _____ (LSB)							
32	_____							
51	Reserved _____							
52	_____							
79	Get and set attributes parameters (see 5.2.1) _____							
80	_____							
173	Security parameters (see 5.2.5) _____							

The FLUSH SCOPE field (see table 48) specifies the scope of the data and attribute bytes that are written to stable storage.

Table 48 — Collection flush scope values

Value	Scope of data and attributes that shall be written to stable storage
00b	List of user objects contained in the collection
01b	Collection attributes only
10b	a) List of user objects contained in the collection; and b) Collection attributes
11b	Reserved

Table 49 — Flush scope values

Scope of data and attributes that shall be written to stable storage		
Contents of USER_OBJECT_ID field		
Value	Not Zero (i.e., user object)	Zero (i.e., partition)
00b	User object data and attributes	List of OSD objects contained in the partition ^a
01b	User object attributes only	Partition attributes only
10b	Reserved	a) List OSD objects contained in the partition; ^a b) Partition attributes; and e) User object data and attributes for listed OSD objects ^b
11b	Reserved	Reserved

^a For all partitions except partition zero, the listed OSD objects are user objects and collections. For partition zero, the listed OSD objects are partitions.

^b A FLUSH OBJECT command with a flush scope of 10b that is addressed to partition zero, flushes the entire OSD logical unit.

The GET/SET CDBFMT field specifies the format of the get and set attributes parameters as described in 5.2.1.

The contents of the TIMESTAMPS CONTROL field are defined in 5.2.7.

The contents of the PARTITION_ID field are defined in 5.2.4.

The COLLECTION_OBJECT_ID field specifies Collection_Object_ID (see 4.6.6). If the collection identified by the COLLECTION_OBJECT_ID field does not exist, 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 get and set attributes parameters are defined in 5.2.1. The format of the Data-In Buffer and Data-Out Buffer when attributes are being retrieved or set is described in 4.11.

The security parameters are defined in 5.2.5.

6.y FLUSH OSD

The FLUSH OSD command (see table 49) ensures that the specified data and attribute bytes for the OSD logical unit are written to stable storage (see 4.10).

Table 49 — FLUSH OSD command

Bit Byte	7	6	5	4	3	2	1	0
8	(MSB) _____							
9	SERVICE ACTION (881Ch)							(LSB)
10	Reserved						FLUSH SCOPE	
11	Reserved		GET/SET CDBFMT		Reserved			
12	TIMESTAMPS CONTROL							
13	Reserved							
51	_____							
52	_____							
79	Get and set attributes parameters (see 5.2.1)							
80	_____							
173	Security parameters (see 5.2.5)							

The FLUSH SCOPE field (see table 50) specifies the scope of the data and attribute bytes that are written to stable storage.

Table 50 — Root object flush scope values

Value	Scope of data and attributes that shall be written to stable storage
00b	List of partitions contained in the OSD logical unit
01b	Root object attributes only
10b	a) List of partitions contained in the OSD logical unit; b) Root object attributes; c) Lists of user objects and collections contained in the every partition in the OSD logical unit; d) Partition attributes for every partition in the OSD logical unit; e) User object data for every user object in the OSD logical unit; f) User object attributes for every user object in the OSD logical unit; g) List of user objects contained in every the collection in the OSD logical unit; and h) Collection attributes for every collection in the OSD logical unit
11b	Reserved

Table 49 — Flush scope values

Scope of data and attributes that shall be written to stable storage		
Contents of USER_OBJECT_ID field		
Value	Not Zero (i.e., user object)	Zero (i.e., partition)
00b	User object data and attributes	List of OSD objects contained in the partition ^a
01b	User object attributes only	Partition attributes only
10b	Reserved	a) List OSD objects contained in the partition; ^a b) Partition attributes; and c) User object data and attributes for listed OSD objects ^b
11b	Reserved	Reserved

^a For all partitions except partition zero, the listed OSD objects are user objects and collections. For partition zero, the listed OSD objects are partitions.

^b A FLUSH OBJECT command with a flush scope of 10b that is addressed to partition zero, flushes the entire OSD logical unit.

The GET/SET CDBFMT field specifies the format of the get and set attributes parameters as described in 5.2.1.

The contents of the TIMESTAMPS CONTROL field are defined in 5.2.7.

The get and set attributes parameters are defined in 5.2.1. The format of the Data-In Buffer and Data-Out Buffer when attributes are being retrieved or set is described in 4.11.

The security parameters are defined in 5.2.5.

6.2 FLUSH PARTITION

The FLUSH PARTITION command (see table 51) ensures that the specified data and attribute bytes for the specified partition are written to stable storage (see 4.10).

Table 51 — FLUSH PARTITION command

Bit Byte	7	6	5	4	3	2	1	0
8	(MSB) _____							
9	SERVICE ACTION (881Bh) _____ (LSB)							
10	Reserved						FLUSH SCOPE	
11	Reserved		GET/SET CDBFMT		Reserved			
12	TIMESTAMPS CONTROL							
13	_____							
15	Reserved _____							
16	(MSB) _____							
23	PARTITION_ID _____ (LSB)							
24	_____							
51	Reserved _____							
52	_____							
79	Get and set attributes parameters (see 5.2.1) _____							
80	_____							
173	Security parameters (see 5.2.5) _____							

The FLUSH SCOPE field (see table 52) specifies the scope of the data and attribute bytes that are written to stable storage.

Table 52 — Partition flush scope values

Value	Scope of data and attributes that shall be written to stable storage
00b	List of user objects and collections contained in the partition
01b	Partition attributes only
10b	a) List of user objects and collections contained in the partition; b) Partition attributes; c) User object data for every user object in the partition; d) User object attributes for every user object in the partition; e) List of user objects contained in every the collection in the partition; and f) Collection attributes for every collection in the partition
11b	Reserved

Table 49 — Flush scope values

Scope of data and attributes that shall be written to stable storage		
Contents of USER_OBJECT_ID field		
Value	Not Zero (i.e., user object)	Zero (i.e., partition)
00b	User object data and attributes	List of OSD objects contained in the partition ^a
01b	User object attributes only	Partition attributes only
10b	Reserved	a) List OSD objects contained in the partition; ^a b) Partition attributes; and e) User object data and attributes for listed OSD objects ^b
11b	Reserved	Reserved

^a For all partitions except partition zero, the listed OSD objects are user objects and collections. For partition zero, the listed OSD objects are partitions.

^b A FLUSH OBJECT command with a flush scope of 10b that is addressed to partition zero, flushes the entire OSD logical unit.

The GET/SET CDBFMT field specifies the format of the get and set attributes parameters as described in 5.2.1.

The contents of the TIMESTAMPS CONTROL field are defined in 5.2.7.

The contents of the PARTITION_ID field are defined in 5.2.4. If the PARTITION_ID field contains 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 get and set attributes parameters are defined in 5.2.1. The format of the Data-In Buffer and Data-Out Buffer when attributes are being retrieved or set is described in 4.11.

The security parameters are defined in 5.2.5.

B.1 Service action codes

...

Table B.1 — Numerical order OSD service action codes

Service Action	Command
...	...
8808h	FLUSH OBJECT
...	...
881Ah	FLUSH COLLECTION
881Bh	FLUSH PARTITION
881Ch	FLUSH OSD
...	...