ENDL
TEXAS

Date: 17 September 2008
To: T10 Technical Committee & SNIA OSD TWG
From: Ralph O. Weber
Subject: Response to T10 Letter Ballot comments on OSD-2

Introduction

This document reviews and responds to the comments included with votes submitted during the forwarding letter ballot on OSD-2 r04 (project T10/1729-D).

Every effort has been exerted to make the comment numbers in this document match those in the Letter Ballot report document T10/08-319r1.

Note: For each revision of this document, an equivalent FDF file is provided which can be imported into OSD-2 r04 to track the comment resolution process using Acrobat markups.

Revision History

r0 Initial revision with limited features, intended for preliminary review by the SNIA OSD TWG

Resolution Summary

The lists of comments on the following pages may be used to locate comments with specific types of resolutions and each entry is a PDF hot link to the comment and resolution text. The PDF bookmarks may be used to locate comments based on their source company.

The following table summarizes numbers of comments with specific types of resolutions by source company.

	Technical		Editorial						
Company	As Is	Changed	As Is	Changed	Rejected	Deferred or No Action Taken	Open	TBD	Total
Dell, Inc.								13	13
ENDL Texas								26	26
Hewlett Packard Co.							1	45	46
IBM Corp.							2	17	19
LSI Corp.								5	5
Seagate Technology						1	8	57	66
Symantec							13	8	21
Late Comments			1						1
Total			1			1	24	171	197

18484 Preston Road, Suite 102 PMB 178, Dallas, TX 75252 214-912-1373 972-596-2775 Fx ENDL_TX@computer.org

Unresolved Comments List

Rejected Comments List

Comments Deferred to the Next Version List

Substantive Comments Accepted As Proposed

Substantive Comments Accepted With Noted Changes

Accepted As Proposed Non-Substantive Comments List

Accepted With Noted Changes Non-Substantive Comments List

No Action Requested, No Action Taken Comments List

Unprocessed Comments List

Unprocessed Comments List (continued)

Unprocessed Comments List (continued)

Unprocessed Comments List (continued)

1. Dell, Inc.

Kevin Marks of Dell, Inc. submitted the following comments on a Yes vote.

Dell 1) Unprocessed Comment (Unprocessed) [1]
PDF pg 1, pg i, Global

I see both SPC-3 and SPC-4 referenced though out. Which is it?

Dell 2) Unprocessed Comment (Unprocessed) [2]
PDF pg 25, pg 1, Clause 1, a,b,c list on this page

devices;
s/b
devices; and

Dell 3) Unprocessed Comment (Unprocessed) [3]
PDF pg 26, pg 2, Clause 1, a,b,c list on this page

commands;
s/b
commands; and

Dell 4) Unprocessed Comment (Unprocessed) [4]
PDF pg 46, pg 22, 4.6.6.6, 1st ?,2,3 list in subclause

a)
s/b
1)

I think

Dell 5) Unprocessed Comment (Unprocessed) [5]
PDF pg 59, pg 35, 4.11.1, list entry B)

(see 4.11.3.1);
s/b
(see 4.11.3.1); and

Dell 6) Unprocessed Comment (Unprocessed) [6]
PDF pg 81, pg 57, 4.12.1, 1st a,b,c list on pg

Delete extra "and" in a,b,c list

Dell 7) Unprocessed Comment (Unprocessed) [7]
PDF pg 178, pg 154, 6.12, 2nd to last a,b,c list on pg

undefined;
s/b
undefined; and

Dell 8) Unprocessed Comment (Unprocessed) [8]
PDF pg 216, pg 192, 6.26.1, 1st a,b,c list on pg

0000h.
s/b
0000h;

Dell 9) Unprocessed Comment (Unprocessed) [9]
PDF pg 260, pg 236, 7.1.3.4, 1st a,b,c list in subclause

represents;

needs and/or?

Dell 10) Unprocessed Comment (Unprocessed) [10]
PDF pg 261, pg 237, 7.1.3.5, 1st a,b,c list in subclause

represents;

needs and/or?

Dell 11) Unprocessed Comment (Unprocessed) [11]
PDF pg 262, pg 238, 7.1.3.6, 1st a,b,c list in subclause

represents;

needs and/or?

Dell 12) Unprocessed Comment (Unprocessed) [12]
PDF pg 263, pg 239, 7.1.3.7, 1st a,b,c list in subclause

represents;

needs and/or?

Dell 13) Unprocessed Comment (Unprocessed) [13]
PDF pg 321, pg 297, 7.1.3.30, 2nd a,b,c list after table 222

,
s/b
;

2. ENDL Texas

Ralph O. Weber of ENDL Texas submitted the following comments on a Yes vote.

ENDL 1) Unprocessed Comment (Unprocessed) [14]
PDF pg 25, pg 1, Global

SPC-3 s/b SPC-4

ENDL 2) Unprocessed Comment (Unprocessed) [15]
PDF pg 28, pg 4, 2.5

SAM-4 is close enough to published to move it to normative references

ENDL 3) Unprocessed Comment (Unprocessed) [16]
PDF pg 29, pg 5, Global in 3.1

Glossary entries that use 'object' in the SAM-3 sense are no longer correct. Reword them.

ENDL 4) Unprocessed Comment (Unprocessed) [17]
PDF pg 29, pg 5, Global

SAM-3 s/b SAM-4

ENDL 5) Unprocessed Comment (Unprocessed) [18]
PDF pg 29, pg 5, 3.1.15

SCSI tasks seems unlikely to appear in the SAM-4 definition of a device server.

ENDL 6) Unprocessed Comment (Unprocessed) [19]
PDF pg 30, pg 6, 3.1.26

The definition of object needs to be updated to be consistent with the SAM-4 switch to UML.

ENDL 7) Unprocessed Comment (Unprocessed) [20]
PDF pg 31, pg 7, 3.1.50

What SAM-3 calls a task, SAM-4 calls a command. The 'task' definition needs to be restructured to coincide with SAM-4.

ENDL 8) Unprocessed Comment (Unprocessed) [21]
PDF pg 38, pg 14, 4.1

Task Priority s/b Command Priority [per SAM-4]

ENDL 9) Unprocessed Comment (Unprocessed) [22]
PDF pg 48, pg 24, 4.7

After write access to an object has been denied, there appears to be no way to use a command to restore that access.

ENDL 10) Unprocessed Comment (Unprocessed) [23]
PDF pg 102, pg 78, 4.12.10, p3, s1

QUERY TASK s/b QUERY TASK, QUERY TASK SET, and QUERY ASYNCHRONOUS EVENT

ENDL 11) Unprocessed Comment (Unprocessed) [24]
PDF pg 112, pg 88, 4.14, p 4, s 2

reported as deferred errors (see SPC-3)
s/b
handled as described in 4.11.3.

ENDL 12) Unprocessed Comment (Unprocessed) [25]
PDF pg 117, pg 93, 4.16.1, p 1, s 1

SPC-4 subclause 4.5.1 defines exceptions that make this "all errors" statement invalid. Resolve the conflict between OSD-2 and SPC-4.

ENDL 13) Unprocessed Comment (Unprocessed) [26]
PDF pg 117, pg 93, 4.16.1, last p

I believe that 4.11.3 defines the error recovery mechanism for errors that occur after a command has completed, and the mechanism does not involve deferred error reporting.

ENDL 14) Unprocessed Comment (Unprocessed) [27]
PDF pg 122, pg 98, 4.17, 2nd p before table 58

before the task containing that command enters the enabled task state
s/b
before the command enters the enabled command state [per SAM-4]

ENDL 15) Unprocessed Comment (Unprocessed) [28]
PDF pg 124, pg 100, 5.1, p 3

The structure of the operation code field is described in SPC-4, not SAM-4

ENDL 16) Unprocessed Comment (Unprocessed) [29]
PDF pg 127, pg 103, 5.2.6.2, table 62

byte 47 s/b byte 51 [to match table 60]

ENDL 17) Unprocessed Comment (Unprocessed) [30]
PDF pg 128, pg 104, 5.2.6.2, end of 2nd p on pg

[insert]If the attribute length field is set to a value that is less than 18, the unused attribute value bytes shall be placed at the highest offsets in the attribute value field.

ENDL 18) Unprocessed Comment (Unprocessed) [31]
PDF pg 128, pg 104, 5.2.6.3, table 63

byte 47 s/b byte 51 [to match table 60]

ENDL 19) Unprocessed Comment (Unprocessed) [32]
PDF pg 128, pg 104, 5.2.6.3, table 63

table 63

The attribute length field s/b 2 bytes not 4 bytes. For compatibility, insert two reserved bytes before the field.

ENDL 20) Unprocessed Comment (Unprocessed) [33]
PDF pg 130, pg 106, 5.2.6.4, table 64

byte 47 s/b byte 51 [to match table 60]

ENDL 21) Unprocessed Comment (Unprocessed) [34]
PDF pg 131, pg 107, 5.2.7, last p in subclause

bit being set s/b bit being set to one

ENDL 22) Unprocessed Comment (Unprocessed) [35]
PDF pg 210, pg 186, 6.24, table 120 & text following it

[field name] TASK TAG s/b COMMAND IDENTIFIER [per SAM-4]

ENDL 23) Unprocessed Comment (Unprocessed) [36]
PDF pg 211, pg 187, 6.24, table 121

Task Tag Specified s/b Command Identifier Specified

ENDL 24) Unprocessed Comment (Unprocessed) [37]
PDF pg 211, pg 187, 6.24, table 121

QUERY UNIT ATTENTION s/b QUERY ASYNCHRONOUS EVENT

ENDL 25) Unprocessed Comment (Unprocessed) [38]
PDF pg 328, pg 304, 7.1.4.3, table 229

table 229

The attribute value field is not 8-byte aligned. Six bytes must be added before the field begins. Suggest adding these bytes before the attribute length field.

ENDL 26) Unprocessed Comment (Unprocessed) [39]
PDF pg 48, pg 24, 4.7, p 1, s 1

p1 s1

Because of this, it is not possible to make a set of objects writeable once the write prohibition has been set. If writing is prohibited, then writing of the object accessibility attribute is prohibited too.

3. Hewlett Packard Co.

Rob Elliott of Hewlett Packard Co. submitted the following comments on a Yes vote.

HP 1) Unprocessed Comment (Unprocessed) [40]
PDF pg 27, pg 3, 2.2

(global)

ISO/IEC 14776-413, SCSI Architecture Model - 3 (SAM-3) [ANSI INCITS 402-2005]

Upgrade all SAM-3 references to SAM-4, since SAM-4 is published.

HP 2) Unprocessed Comment (Unprocessed) [41]
PDF pg 28, pg 4, 2.5

T10/1731-D
s/b
ANSI INCITS xxx-2008

since SPC-4 is published

HP 3) Unprocessed Comment (Unprocessed) [42]
PDF pg 28, pg 4, 2.5

T10/1683-D
s/b
ANSI INCITS xxx-2008

HP 4) Unprocessed Comment (Unprocessed) [43]
PDF pg 38, pg 14, 4.1

Task Priority
s/b
Command Priority

to match final version of sam4

HP 5) Unprocessed Comment (Unprocessed) [44]
PDF pg 38, pg 14, 4.2

response to an INQUIRY command
s/b
the standard INQUIRY data

HP 6) Unprocessed Comment (Unprocessed) [45]
no location specified

Delete
and

HP 7) Unprocessed Comment (Unprocessed) [46]
PDF pg 40, pg 16, 4.4, 1 p after figure 3

disc drives
s/b
disk drives

Only Seagate spells it "disc"

HP 8) Unprocessed Comment (Unprocessed) [47]
PDF pg 40, pg 16, 4.4, 2nd p after figure 3, s 1

share directly access

does not parse

HP 9) Unprocessed Comment (Unprocessed) [48]
PDF pg 40, pg 16, 4.4, 2nd p after figure 3, s 2

possibly
s/b
possibly the

HP 10) Unprocessed Comment (Unprocessed) [49]
PDF pg 40, pg 16, 4.4, 1st a,b,c list

initiator devices
s/b
SCSI initiator devices

also in figure 3

HP 11) Unprocessed Comment (Unprocessed) [50]
PDF pg 40, pg 16, 4.4, 3rd p after figure 3, s 1

application clients place
s/b
that application clients place

HP 12) Unprocessed Comment (Unprocessed) [51]
PDF pg 41, pg 17, 4.5, p 1, s 1

OBSD (see 3.1.27) logical unit
s/b
OSD logical unit

That's the phrase used in the next sentence, among other places

HP 13) Unprocessed Comment (Unprocessed) [52]
PDF pg 41, pg 17, 4.6.1, b) Partition

This OSD object
s/b
This type of OSD object

since there may be more than one. As written, it sounds like there is just one.

HP 14) Unprocessed Comment (Unprocessed) [53]
PDF pg 41, pg 17, 4.6.1, c) Collection

This OSD object
s/b
This type of OSD object

since there may be more than one. As written, it sounds like there is just one.

HP 15) Unprocessed Comment (Unprocessed) [54]
PDF pg 41, pg 17, 4.6.1, d) User object

This OSD object
s/b
This type of OSD object

since there may be more than one. As written, it sounds like there is just one.

HP 16) Unprocessed Comment (Unprocessed) [55]
PDF pg 41, pg 17, 4.6.1, d) User object

it
s/b
the user object

HP 17) Unprocessed Comment (Unprocessed) [56]
PDF pg 42, pg 18, 4.6.2, table 3

assigned by
s/b
are assigned by

HP 18) Unprocessed Comment (Unprocessed) [57]
PDF pg 42, pg 18, 4.6.2, table 3

assigned by
s/b
are assigned by

HP 19) Unprocessed Comment (Unprocessed) [58]
PDF pg 42, pg 18, 4.6.2, table 3

Footnote a also applies to "well known collections and "collection or user object", since those rows have multiple Partition_ID values to choose from as well.

HP 20) Unprocessed Comment (Unprocessed) [59]
PDF pg 44, pg 20, 4.6.6.3, table 4

The "A CREATE USER" and "A REMOVE" cells should be left justified to match the column headers

HP 21) Unprocessed Comment (Unprocessed) [60]
PDF pg 45, pg 21, 4.6.6.5.2, p 1, s 1

which
s/b
that

HP 22) Unprocessed Comment (Unprocessed) [61]
PDF pg 45, pg 21, 4.6.6.5.3, 1st a,b,c list

REFERESH
s/b
REFRESH

HP 23) Unprocessed Comment (Unprocessed) [62]
PDF pg 46, pg 22, 4.6.6.5.3, 1st a,b,c list on pg

REFERESH
s/b
REFRESH

HP 24) Unprocessed Comment (Unprocessed) [63]
PDF pg 47, pg 23, 4.6.6.6, 2nd p after note 2, s 1

device sever
s/b
device server

HP 25) Unprocessed Comment (Unprocessed) [64]
PDF pg 59, pg 35, 4.11.1, list entry C)

conditions(
add space

HP 26) Data sharing and error recovery (Unresolved) [65]

PDF pg 77, pg 53, 4.11.3.1

With snapshots allowing partitions to share parts of there data, it is more likely that an error will affect multiple partitions. For c) A) c), it may be worth reporting all affected Partition_IDs rather than just one, or the extent of damage will be unclear.

HP 27) Unprocessed Comment (Unprocessed) [66]
PDF pg 108, pg 84, 4.13.2.4, p 1

SHAPSHOT
s/b
SNAPSHOT

HP 28) Unprocessed Comment (Unprocessed) [67]
PDF pg 110, pg 86, 4.13.3, table 43, DO NOT CARE row

device sever
s/b
device server

HP 29) Unprocessed Comment (Unprocessed) [68]
PDF pg 110, pg 86, 4.13.3, 2nd p after table 43

Delete
the

HP 30) Unprocessed Comment (Unprocessed) [69]
PDF pg 112, pg 88, 4.14

Should OSD add an FUA_NV bit to distinguish between these data locations:

FUA: data is safe on the medium. (if the OSD controller is using disk drives as its medium, the disk drives can be moved to another OSD controller)

FUA_NV: data is safe either in NV cache or on the medium. (disk drives cannot be moved without also moving the NV cache)

HP 31) Unprocessed Comment (Unprocessed) [70]
PDF pg 117, pg 93, 4.16.1, p 1, s 1

OSD logical units shall use descriptor format sense data (see SPC-3) to report all errors.

While the intent is admirable, this is a bit too aggressive (like SAT-2's attempt to mandate descriptor format for the ATA PASS-THROUGH commands).

1. OSD ought to honor the D_SENSE bit in the Control mode page. If fixed format is requested, then return very abbreviated (almost useless) fixed-format data rather than return another format.

2. SAM-4/SPC-4 require that fixed format be used to report reset unit attention conditions and the codes needed to run MODE SELECT to turn on fixed-format, so old software just doesn't get confused

HP 32) Unprocessed Comment (Unprocessed) [71]
PDF pg 124, pg 100, 5.1, 1st p after table 59, s 2

216
s/b
228

to match the value in table 59.

May be better to say
"the value specified in table 59" to avoid this problem.

HP 33) Unprocessed Comment (Unprocessed) [72]
PDF pg 126, pg 102, 5.2.3, p 1

The definition of DPO in 4.14 and 5.2.3 is a bit misleading - it sounds like a weak "should" version of FUA, advising the device server to not put data in the cache.

Really, DPO signals that the application client does not expect to read the data again, so the device server need not retain the data in the cache in hopes of a cache hit on a read.

To maintain reasonable write performance, however, it is important that the device server temporarily put the data in its cache, performing the write to medium when convenient. DPO should not cause a drop to the FUA level of performance.

HP 34) Unprocessed Comment (Unprocessed) [73]
PDF pg 126, pg 102, 5.2.2, p 2, s 2

ADDITIONAL LENGTH
s/b
smallcaps

HP 35) Unprocessed Comment (Unprocessed) [74]
PDF pg 134, pg 110, 5.2.13, table 67, 0h row

0h
s/b
00h
to match width of other values in the table

HP 36) Unprocessed Comment (Unprocessed) [75]
PDF pg 146, pg 122, 6.1, table 80, footnote a

Service action codes values between…
s/b
For the operation code 7Fh, service action codes values between…

and join with the next sentence to share that "For…"

There are some non-7Fh opcodes in the table, and this comment does not apply to their service action values.

HP 37) Unprocessed Comment (Unprocessed) [76]
PDF pg 205, pg 181, 6.22.2 & 6.22.3

REBUILD IN PROGRESS

was intended for SCC (RAID volumes) to report that they are rebuilding a RAID volume. OSD performing an object check seems a bit different. I think a new additional sense code is worth adding for this reason. The progress indicator will make more sense if the operation in progress is clearly identified. The time for OBJECT STRUCTURE CHECK to complete is based on the amount of metadata; the time for a RAID rebuild is based on the amount of data.

Same comment applies to 6.22.3.

HP 38) Unprocessed Comment (Unprocessed) [77]
PDF pg 206, pg 182, 6.22.3

With snapshots, several partitions might share the same data. It seems like OBJECT STRUCTURE CHECK might need to ensure that all commands to all such partitions are terminated in 6.22.3.

HP 39) Unprocessed Comment (Unprocessed) [78]
PDF pg 211, pg 187, 6.24, table 121

QUERY UNIT ATTENTION
s/b
QUERY ASYNCHRONOUS EVENT

to match final SAM-4 terminology

HP 40) Unprocessed Comment (Unprocessed) [79]
PDF pg 228, pg 204, 6.30.1, table 136 title

REFERESH
s/b
REFRESH

HP 41) Unprocessed Comment (Unprocessed) [80]
PDF pg 240, pg 216, 6.35.1, p 1

partition(

add space

HP 42) Unprocessed Comment (Unprocessed) [81]
PDF pg 249, pg 225, 6.38.2, 1st p on pg

device sever
s/b
device server

HP 43) Unprocessed Comment (Unprocessed) [82]
PDF pg 264, pg 240, 7.1.3.8, table 160, Maximum CDB continuation length row

The field that the Maximum CDB continuation length field establishes the upper limit for is a 4-byte field. Should the length of this attribute be 4 bytes to match?

HP 44) Unprocessed Comment (Unprocessed) [83]
PDF pg 332, pg 308, 7.5.2.1, table 234, last row

F1
s/b
F1h

HP 45) Unprocessed Comment (Unprocessed) [84]
PDF pg 335, pg 311, B.1, table B.1, footnotes

Consider adding a footnote explaining how all the OSD-1 opcodes were made obsolete and replaced in OSD-2. Item d) on page 2 already explains that, but this table listing opcodes would be a good place to highlight it.

HP 46) Unprocessed Comment (Unprocessed) [85]
PDF pg 335, pg 311, B.1, table B.1 title

Add "(part n of 2)" to table B.1 header

4. IBM Corp.

Kevin Butt of IBM Corp. submitted the following comments on a No vote.

IBM 1) Unprocessed Comment (Unprocessed) [86]
PDF pg 29, pg 5, 3.1.11

Here you talk about a single capability while in 3.1.16 you say "not the first". In credential you may want to say also a capability or several capabilities

IBM 2) Unprocessed Comment (Unprocessed) [87]
PDF pg 36, pg 12, 3.7.2

the use of term object is somewhat confusing here (as we talk about storage objects). You may want to consider "modules" or something else

IBM 3) Unprocessed Comment (Unprocessed) [88]
PDF pg 41, pg 17, 4.5, p 1, s 1

LBAs are not defined here. You may want to remove the statement that reffers to LBAs (it contains no information)

IBM 4) Unprocessed Comment (Unprocessed) [89]
PDF pg 41, pg 17, 4.6.1, ???

This is hard to parse (can be parsed ambiguously)

IBM 5) Unprocessed Comment (Unprocessed) [90]
PDF pg 43, pg 19, 4.6.5, p 1

has to clarify uniqueness within a partition. The "assigned by OSD LU" is a bit confusing (perhaps explain)

IBM 6) Unprocessed Comment (Unprocessed) [91]
PDF pg 50, pg 26, 4.8.1, 4th p on pg, s 2

with

IBM 7) Unprocessed Comment (Unprocessed) [92]
PDF pg 61, pg 37, 4.11.2.2.1, 1st p on pg

I think that this formulation is problematic. I would like to see a clause that says that at least the effected changes are reflected in results (sense?)

IBM 8) Unprocessed Comment (Unprocessed) [93]
PDF pg 80, pg 56, 4.12.1, last p on pg

that statement is not completely correct. I would state it that"as a result of the communications between SM and client the client should be able to build a capability and should have the cap-key. The rest of the text calls a combination of those two a credential and uses for it a data-structure similar to those used in communications with the data server for ease of illustration"

IBM 9) Unprocessed Comment (Unprocessed) [94]
PDF pg 81, pg 57, 4.12.1, 2nd p on pg, s 2

mising "is"

IBM 10) Unprocessed Comment (Unprocessed) [95]
PDF pg 82, pg 58, 4.12.2, table 28

should mention that only cap-keys have to be confidential

IBM 11) Unprocessed Comment (Unprocessed) [96]
PDF pg 98, pg 74, 4.12.7.2, last p in subclause

A flow diagram of the checks might help a lot the reader and perhaps replace the text

IBM 12) Unprocessed Comment (Unprocessed) [97]
PDF pg 112, pg 88, 4.13.5, 1st p on pg

should not say something about "no space" as opposed to quota exhaustion.

IBM 13) Unprocessed Comment (Unprocessed) [98]
PDF pg 117, pg 93, 4.16.1, last p

are all references consistently to SPC3 or SPC4

IBM 14) Unprocessed Comment (Unprocessed) [99]
PDF pg 126, pg 102, 5.2.2, p 2, s 2

mention table at 172?

IBM 15) Unprocessed Comment (Unprocessed) [100]
PDF pg 132, pg 108, 5.2.8

In 5.28 and others can't the error report be more specific - NOT supported xxx - or have the form major, minor with even more detail

IBM 16) Unprocessed Comment (Unprocessed) [101]
PDF pg 134, pg 110, 5.2.13, table 67, 0h row

be updated

IBM 17) Unprocessed Comment (Unprocessed) [102]
PDF pg 144, pg 120, 4.5.4, 2nd to last p in subclause

that is better but I would add UNSUPPORTED to all field that are illegal due to lack of a feature

IBM 18) Remove COPY USER OBJECTS (Unresolved) [103]
PDF pg 153, pg 129, 6.4

I would prefer this command removed

IBM 19) CLEAR interactions w/ READ MAP (Unresolved) [104]

PDF pg 221, pg 197, 6.28.1

It should be stated explicitly that an area of an object cleared with a clear command or having as content the default value of 0 in all bytes may be reported as a data hole!

5. LSI Corp.

John Lohmeyer of LSI Corp. submitted the following comments on a Yes vote.

LSI 1) Unprocessed Comment (Unprocessed) [105]
PDF pg 3, pg iii, Revision History

This needs be be removed before public review.

LSI 2) Unprocessed Comment (Unprocessed) [106]
PDF pg 25, pg 1, Clause 1, 1st a,b,c list, item c)

This should be << devices; and >>

LSI 3) Unprocessed Comment (Unprocessed) [107]
PDF pg 26, pg 2, Clause 1, 1st a,b,c, list on pg, item c)

This should be << commands; and >>

LSI 4) Unprocessed Comment (Unprocessed) [108]
PDF pg 41, pg 17, Global

Global - Change all references to SAM-3 to SAM-4.

LSI 5) Unprocessed Comment (Unprocessed) [109]
PDF pg 41, pg 17, 4.5, p 2

Change this to << In addition to the sets of objects (i.e, classes) defined in SAM-4, this >>

6. Seagate Technology

Gerry Houlder of Seagate Technology submitted the following comments on a No vote.

Seagate 1) Unprocessed Comment (Unprocessed) [110]
PDF pg 22, pg xxii, Foreword

Should be 2008.

Seagate 2) Unprocessed Comment (Unprocessed) [111]
PDF pg 22, pg xxii, Foreword

Grammar. Either drop "of" or reword the sentence.

Seagate 3) Unprocessed Comment (Unprocessed) [112]
PDF pg 42, pg 18, 4.6.3, p 2

The description of the root object lists several data manipulation commands that are not allowed (READ, WRITE, APPEND). There are now other data manipulation commands (CLEAR and PUNCH). Should these commands be added to the list? Or, maybe the description should be generalized, such as "The device server shall terminate all commands that are defined to manipulate user data (e.g., READ, WRITE, APPEND, …) that are sent to the root object with CHECK CONDITION status, …"

Seagate 4) Unprocessed Comment (Unprocessed) [113]
PDF pg 43, pg 19, 4.6.4, last p in subclause

There are now other data manipulation commands (CLEAR and PUNCH). Should these commands be added to the list? Or, maybe the description should be generalized, such as "The device server shall terminate all commands that are defined to manipulate user data (e.g., READ, WRITE, APPEND, …) that are sent to the partition object with CHECK CONDITION status, …"

Seagate 5) Unprocessed Comment (Unprocessed) [114]
no location specified

There are now other data manipulation commands (CLEAR and PUNCH). Should these commands be added to the list? Or, maybe the description should be generalized, such as "The device server shall terminate all commands that are defined to manipulate user data (e.g., READ, WRITE, APPEND, …) that are sent to the collection object with CHECK CONDITION status, …"

Seagate 6) Unprocessed Comment (Unprocessed) [115]
PDF pg 43, pg 19, 4.6.6.1, p 6

There are now other data manipulation commands (CLEAR and PUNCH). Should these commands be added to the list? Or, maybe the description should be generalized, such as "The device server shall terminate all commands that are defined to manipulate user data (e.g., READ, WRITE, APPEND, …) that are sent to the collection object with CHECK CONDITION status, …"

Seagate 7) Unprocessed Comment (Unprocessed) [116]
PDF pg 45, pg 21, 4.6.6.5.3, 1st a,b,c list

Spelling error in "REFERESH"

Seagate 8) Unprocessed Comment (Unprocessed) [117]
no location specified

Spelling error in "REFERESH"

Seagate 9) Unprocessed Comment (Unprocessed) [118]
PDF pg 46, pg 22, 4.6.6.5.3, 1st a,b,c list on pg

Spelling error in "REFERESH"

Seagate 10) Unprocessed Comment (Unprocessed) [119]
PDF pg 46, pg 22, 4.6.6.6, p 2, s 2

Should be "other than TRACKING"

Seagate 11) Unprocessed Comment (Unprocessed) [120]
PDF pg 46, pg 22, 4.6.6.6, p 3, s 1

Should be "an object"

Seagate 12) Unprocessed Comment (Unprocessed) [121]
PDF pg 47, pg 23, 4.6.6.6, 2nd p after note 2, s 1

Should be "server"

Seagate 13) Unprocessed Comment (Unprocessed) [122]
PDF pg 47, pg 23, 4.6.6.6, 1st a,b,c list after note 2

Wrong reference, should be 5.2.7, not 5.2.5

Seagate 14) Unprocessed Comment (Unprocessed) [123]
PDF pg 48, pg 24, 4.6.6.6, table 6

This notation is not defined until section 4.8.5

Seagate 15) Unprocessed Comment (Unprocessed) [124]
PDF pg 49, pg 25, 4.7, 2nd p before table 7

States that the denial of write access to an object with members means denial of the ability to create new members in that object. It should also deny the ability to remove member objects.

Seagate 16) Unprocessed Comment (Unprocessed) [125]
PDF pg 49, pg 25, 4.7, 1st p before table 7

Once the object accessibility attribute is set to 1, can it be set to zero? Seems like this write protects the object AND ATTRIBUTES, preventing it from being set back to 0. That is, once an object is made read-only, it cannot be reverted back to read/write mode.

Seagate 17) Unprocessed Comment (Unprocessed) [126]
PDF pg 50, pg 26, 4.8.1, 4th p on pg, s 2

It is inconsistent that this paragraph indicates that multi-object commands can retrieve or store attributes to multiple objects, while table 6 indicates that only the GET and SET MEMBER OBJECTS command can do this.

Seagate 18) Unprocessed Comment (Unprocessed) [127]
PDF pg 50, pg 26, 4.8.1, 4th p on pg, s 2

Should be "associated with each"

Seagate 19) Unprocessed Comment (Unprocessed) [128]
PDF pg 54, pg 30, 4.8.6, p 3 w/ its a,b,c list

The error conditions described in subparagraphs "a" and "b" refer to an "invalid attribute length". This is not relevant to the error being described. The error is the attribute number.

Seagate 20) Unprocessed Comment (Unprocessed) [129]
PDF pg 68, pg 44, 4.11.2.2.4, p 1, s 1 & s 2

Leave a space between sentences.

Seagate 21) Unprocessed Comment (Unprocessed) [130]
PDF pg 70, pg 46, 4.11.2.3, table 25, CREATE CLONE row

Shouldn't this require CREATE and WRITE permission bits for destination?

Seagate 22) Create permission for CREATE CLONE (Unresolved) [131]
PDF pg 70, pg 46, 4.11.2.3, table 25, CREATE CLONE row

Shouldn't this require CREATE and WRITE permission bits for destination?

Seagate 23) Create permission for CREATE SNAPSHOT (Unresolved) [132]
PDF pg 70, pg 46, 4.11.2.3, table 25, CREATE SNAPSHOT row

Shouldn't this require CREATE and WRITE permission bits for destination?

Seagate 24) Why Append permission for REFRESH (Unresolved) [133]
PDF pg 72, pg 48, 4.11.2.3, table 25, REFRESH SNAPSHOT OR CLONE row

Why is this APPEND and not WRITE?

Seagate 25) Unprocessed Comment (Unprocessed) [134]
PDF pg 73, pg 49, 4.11.2.3, table 25, RESTORE PARTITION FROM SNAPSHOT row

Looks like READ and WRITE are swapped here. Main partition should be WRITE and snapshot partition should be READ for the RESTORE PARTITION command.

Seagate 26) Current Command attributes page permission bits requirements (Unresolved) [135]
PDF pg 74, pg 50, 4.11.2.3, table 26, USER retrieve from Current Command page row

The first rows indicate that the GET_ATTR bit is required to get attribute from the Current Command Page. The description of GET_ATTR in 4.11.2.2.1 specifically says that GET_ATTR is NOT required to access the Current Command Page attributes (see page 39, 3rd paragraph).

Seagate 27) Current Command attributes page permission bits requirements (Unresolved) [136]
PDF pg 74, pg 50, 4.11.2.3, table 26, COLLECTION retrieve from Current Command page row

The first rows indicate that the GET_ATTR bit is required to get attribute from the Current Command Page. The description of GET_ATTR in 4.11.2.2.1 specifically says that GET_ATTR is NOT required to access the Current Command Page attributes (see page 39, 3rd paragraph).

Seagate 28) Current Command attributes page permission bits requirements (Unresolved) [137]
PDF pg 74, pg 50, 4.11.2.3, table 26, PARTITION/ROOT retrieve from Current Command page row

The first rows indicate that the GET_ATTR bit is required to get attribute from the Current Command Page. The description of GET_ATTR in 4.11.2.2.1 specifically says that GET_ATTR is NOT required to access the Current Command Page attributes (see page 39, 3rd paragraph).

Seagate 29) Unprocessed Comment (Unprocessed) [138]
PDF pg 77, pg 53, 4.11.3.1, 1st B) on pg

Bad grammar: "established affected…".

Seagate 30) Unprocessed Comment (Unprocessed) [139]
PDF pg 79, pg 55, 4.11.3.3, 1st a,b,c list on pg

The OBJECT STRUCTURE CHECK command should be included in the list of allowed commands.

Seagate 31) Unprocessed Comment (Unprocessed) [140]
PDF pg 81, pg 57, 4.12.1, 2nd p on pg, s 2

Bad grammar, missing word(s).

Seagate 32) Unprocessed Comment (Unprocessed) [141]
PDF pg 85, pg 61, 4.12.4.1, table 30, footnote d

This should be "authorized", not "unauthorized".

Seagate 33) Unprocessed Comment (Unprocessed) [142]
PDF pg 88, pg 64, 4.12.4.5, 1st p on pg

Grammar: drop the word "contains"

Seagate 34) Unprocessed Comment (Unprocessed) [143]
PDF pg 96, pg 72, 4.12.6.3.2, 2nd B) on pg

Grammar: repeated words "in the in the".

Seagate 35) Unprocessed Comment (Unprocessed) [144]
PDF pg 104, pg 80, 4.13.2.2, 1st p on pg

Replace the word "show" with "shown"

Seagate 36) Unprocessed Comment (Unprocessed) [145]
PDF pg 117, pg 93, 4.16.1, last p

This is no longer true. OSD-2 (section 4.11.3) defines a new exception management mechanism that should take care of deferred errors.

Seagate 37) Acrobat Overhead Comment (No Action Taken) [146]
PDF pg 117, pg 93, 4.16.1, last p

Marked set by IrenS [Acrobat overhead from FDF file]

Seagate 38) Unprocessed Comment (Unprocessed) [147]
PDF pg 124, pg 100, 5.1, 1st p after table 59, s 2

States that Additional CDB Length should be 216. The table says it should be 228.

Seagate 39) Unprocessed Comment (Unprocessed) [148]
PDF pg 126, pg 102, 5.2.4, p 1, s 1

Can we use a defined verb (shall/should) instead of "is"?

Seagate 40) Unprocessed Comment (Unprocessed) [149]
PDF pg 127, pg 103, 5.2.6.2, table 62

Byte offset is incorrect. Should start at 52.

Seagate 41) Unprocessed Comment (Unprocessed) [150]
PDF pg 128, pg 104, 5.2.6.3, table 63

Byte offset is incorrect. Should start at 52.

Seagate 42) Unprocessed Comment (Unprocessed) [151]
PDF pg 130, pg 106, 5.2.6.4, table 64

Byte offset is incorrect. Should start at 52.

Seagate 43) Unprocessed Comment (Unprocessed) [152]
PDF pg 142, pg 118, 5.4.4, table 77

The CDB Continuation Descriptor Length is a constant (16). It currently refers to "n", which is undefined in this table. The paragraph that describes the descriptor length can be specific - it must be 16.

Seagate 44) Unprocessed Comment (Unprocessed) [153]
PDF pg 143, pg 119, 5.4.5, table 78

The CDB Continuation Descriptor Length is a constant (20). It currently refers to "n", which is undefined in this table. The paragraph that describes the descriptor length can be specific - it must be 20.

Seagate 45) Unprocessed Comment (Unprocessed) [154]
PDF pg 143, pg 119, 5.4.5, table 78

The CDB Continuation Descriptor Length is a constant (20). It currently refers to "n", which is undefined in this table. The paragraph that describes the descriptor length can be specific - it must be 20.

Seagate 46) Unprocessed Comment (Unprocessed) [155]
PDF pg 143, pg 119, 5.4.5

We suggest the removal of COPY USER OBJECTS command from this specification as agreed in the OSD twg. This feature will be revisited for future versions of OSD.

Seagate 47) Unprocessed Comment (Unprocessed) [156]
PDF pg 147, pg 123, 6.1, table 80, REMOVE COLLECTION row

REMOVE COLLECTION should have footnote "b".

Seagate 48) Remove COPY USER OBJECTS (Unresolved) [157]

PDF pg 153, pg 129, 6.4

We suggest the removal of COPY USER OBJECTS command from this specification as agreed in the OSD twg. This feature will be revisited for future versions of OSD.

Seagate 49) Unprocessed Comment (Unprocessed) [158]
PDF pg 156, pg 132, 6.5, last p on pg

Indicates that attribute list type Fh should be used. The referenced section says that list type Fh is obsolete. It should be type Eh.

Seagate 50) Unprocessed Comment (Unprocessed) [159]
PDF pg 160, pg 136, 6.7.1, last p on pg

Wrong reference. Should be 5.2.7, not 5.2.5. Check all other places as well.

Seagate 51) Unprocessed Comment (Unprocessed) [160]
PDF pg 169, pg 145, 6.10.1, 2nd p after table 89

Wrong reference. Should be 5.2.7, not 5.2.5

Seagate 52) Unprocessed Comment (Unprocessed) [161]
PDF pg 194, pg 170, 6.20.1, table 104, 2nd description row

This only really applies to the get attributes parameters. "Set" should be removed.

Seagate 53) Unprocessed Comment (Unprocessed) [162]
PDF pg 200, pg 176, 6.21, table 113, 2nd description row

This only really applies to the get attributes parameters. "Set" should be removed.

Seagate 54) Unprocessed Comment (Unprocessed) [163]
PDF pg 200, pg 176, 6.21, 5th p after table 113

What does Flash have anything to do with this command?

Seagate 55) Unprocessed Comment (Unprocessed) [164]
PDF pg 215, pg 191, 6.26.1, 2nd p after table 125, s 1

Wrong reference. Should be 5.2.7, not 5.2.5

Seagate 56) Unprocessed Comment (Unprocessed) [165]
PDF pg 229, pg 205, 6.30.1, 1st p on pg

Wrong reference. Should be 5.2.7, not 5.2.5

Seagate 57) Unprocessed Comment (Unprocessed) [166]
PDF pg 236, pg 212, 6.32, 5th from last p on pg

Not clear on what this paragraph is trying to address. The command tracking attributes page is removed by this command, which certainly qualifies as modifying them.

Seagate 58) Unprocessed Comment (Unprocessed) [167]
PDF pg 237, pg 213, 6.33, 3rd p after table 141

Wrong reference. Should be 5.2.7, not 5.2.5

Seagate 59) Unprocessed Comment (Unprocessed) [168]
PDF pg 240, pg 216, 6.35.1, last p on pg

Wrong reference. Should be 5.2.7, not 5.2.5

Seagate 60) Unprocessed Comment (Unprocessed) [169]
PDF pg 261, pg 237, 7.1.3.5, table 157

Missing the following entry: "P+7h"

Seagate 61) Unprocessed Comment (Unprocessed) [170]
PDF pg 267, pg 243, 7.1.3.8, 1st p on pg

Default isolation method attribute should be 110h, not 111h.

Seagate 62) Unprocessed Comment (Unprocessed) [171]
PDF pg 267, pg 243, 7.1.3.8, 2nd p after table 162

Supported isolation methods attribute should be 111h, not 112h.

Seagate 63) Unprocessed Comment (Unprocessed) [172]
PDF pg 272, pg 248, 7.1.3.9, table 169, 10th row

missing "C0h reserved".

Seagate 64) Unprocessed Comment (Unprocessed) [173]
PDF pg 272, pg 248, 7.1.3.9, table 169, 19th row

Should be 2FFh, not 1FFh.

Seagate 65) Unprocessed Comment (Unprocessed) [174]
PDF pg 278, pg 254, 7.1.3.11, 1st full p on pg

Grammar: missing "is" between "object" and "the"

Seagate 66) Current Command space consumption attribute (Unresolved) [175]

PDF pg 323, pg 299, 7.1.3.31

On behalf of Panasas:

For purposes of capacity accounting, it is helpful to know how much space a given command has consumed. can we add (if we don't already have) a specific attribute on the "current command" page or some such thing which returns "capacity delta" for the given command [type = int64]? i realize this is a little late in the game, but this was pointed out to me pretty recently at panasas, and should be pretty simple and non-intrusive. i believe there are some commands which already return this as part of the standard response code, so having a common dropping area for this would be cleaner as well.

7. Symantec

Roger Cummings of Symantec submitted the following comments on a No vote.

Symantec 1) Sequential read/write implication (Unresolved) [176]

PDF pg 55, pg 31, 4.9.2, p 2

i don't like the marked paragraph since it seems to imply that data must be read or written by the device sequentially, and that's not what higher performance disks do today.

for example, if you send a disk a full track write, it will typically start writing with whatever sector happens to passing under the head (once the head reaches the appropriate track). it doesn't wait for the first sector of the write to appear.

bad block re-vectoring has a similar effect. if the 3rd sector in a write has been re-vectored, it will be read or written after the other sectors in the read or write. so an i/o error may not be detected until all sectors other than the failed one have been transferred.

we can say that the atomicity guarantees affect the boundaries of how much data is affected by a media error detected when an operation is in progress (presumably the media error will occur at an atomic i/o boundary), but that's about it.

note that neither of the last two paragraphs really requires any particular behavior.

Symantec 2) Unprocessed Comment (Unprocessed) [177]

PDF pg 126, pg 102, 5.2.4, p 1, s 1

"design" seems like a poor choice of word. i suggested "required", but according to Ralph Weber that's a bad idea:
> RE Q1: 'required' is a standards word that will cause T10 readers
> to expect to see exactly how the security method other than NOSEC
> returns an error … inline … in the same paragraph as the word
> 'required'. Sorry, but 'designed' is the best word I can think of
> to put in the cited text.
perhaps "is design to" could be changed to "will"?

if Ralph (or anyone else) continues to object to making a change, then i can live with "designed". but please fix the tense.

Symantec 3) Unprocessed Comment (Unprocessed) [178]

PDF pg 273, pg 249, 7.1.3.9, 1st p after first a,b,c list on pg

i think there's an "is" missing between "partition" and "the". probably "object duplications", in the same phrase, should be singular rather than plural.

Symantec 4) Remove COPY USER OBJECTS (Unresolved) [179]

PDF pg 153, pg 129, 6.4

there are a number of questions that i raise, below, about the COPY USER OBJECTS command. i think the consensus of the group was that we would simply drop the command from the OSD-2 spec. regardless, i raise the individual issues here …

4.1) in section _6.4 COPY USER OBJECTS_ the command requires that the target object cannot previously exist (i.e. it's like create and write).

given that the command can copy multiple objects to the target object, and can pick out multiple ranges within each object, this seems like an odd restriction. it's not just a snapshot command, it's more of a general purpose copy command.

in it's current form my company doesn't have any particular use for the Copy User Objects command, regardless of whether or not it creates the target object, but if we work out the issues with OSD to OSD communication such that a Copy User Objects command can copy from one OSD to another, then i would very much like a Copy User Objects command that worked with pre-existing objects.

do we need a Copy User Objects command that works with existing objects? (my inclination would be to change this to requre an existing object, but we could also add another copy command that works with existing objects.)

4.2) in section _6.4 COPY USER OBJECTS_, the command allows the user to specify the "duplication method". while i can imagine how a copy-on-write implementation of copying an entire object would work, i have a more difficult time imagining copy-on-write for a duplicated object that's composed of a bunch of pieces from a bunch of other source objects.

assuming a given OSD implementation can support a copy-on-write copy of one object to another, but cannot support copy-on-write for an object composed of pieces of 27 other objects, what should the implementation report for the "Supported object duplication method attributes" attribute on the root attributes page (Section 7.1.2.8)?

should it report that it only supports byte-by-byte copies for "Copy User Objects", or should it report that it supports many space efficient types, but then fail the "Copy User Objects" command if any kind of range based-copying is requested?

one possible solution is to have a "clone user object" that only makes a writable clone of an entire existing user object, and another command for "Copy Data", that simply copies data from one or more ranges in one or more user objects to a target object.

4.3) in section _5.4.5 Copy user object source_, the specification of the CPY_ATTR bit (set to copy user attributes), says:

… If the CPY_ATTR bit is set to one, all application client settable attributes (see 7.2.1) are copied from this source user object to the destination user object. …

in the context of COPY USER OBJECTS, this is confusing. "Object Logical Length" is a client settable attribute. how is it treated? presumably, in the case where only part of the original object is copied, or multiple objects copied, we won't change the target object logical length to be the same as one of the source objects.

it seems to me that Object Logical Length should not be copied even if CPY_ATTR bit is set. if there's a single source object that's copied in it's entirety then it should be the same length as the source object anyway.

the draft does describe what happens for the "Reserved Data Space" attribute (it is not copied, it is added to what already exists for the duplicated object).

Symantec 5) Time of Duplication DO NOT CARE (Unresolved) [180]

PDF pg 111, pg 87, 4.13.4.2, table 44, DO NOT CARE row

it's not clear to me how loose this is intended to be. is there a requirement that the duplicated object be a snapshot or not? in my view, a "snapshot" would represent the state of the source object at *some* point in the past, but i don't think this statement makes that a requirement. for example, on UNIX/LINUX systems if you copy a file that's being written to in a random fashion (say start of file, end of file, start of file, end of file, …) using the cp command:
cp busyfile busyfile_copy

the copy will consist of data that was in the source object at some point in time, but not all of the data will have been present at the same point in time (you'll get some of the writes to the end of the file, but not the corresponding writes to the beginning of the file).

i suggest that we clarify this and (possibly) add another Time of Duplication method such that one is "snapshot corresponds to the contents of the source object at some point in time while the copy operation was in progress", and DO NOT CARE is "snapshot may not correspond to any point in time".

Symantec 6) Time of Duplication DO NOT CARE (Unresolved) [181]

PDF pg 111, pg 87, 4.13.4.2, table 44, DO NOT CARE row

in the case where the object is duplicated using one of the above [see comment Symantec 5] methods (either DO NOT CARE or the new one that i suggested), is it reasonable to expect that the timestamps in the duplicated object will indicate the exact time that the copy occurred?

we should probably specify the behavior one way or the other. my preference, is for the object time stamp to represent the exact time that the copy occurred (assuming the copy is consistent).

Symantec 7) Unprocessed Comment (Unprocessed) [182]

PDF pg 103, pg 79, 4.13.2.1, table 36, 2nd to last row

[insert the following at the end of this text]

For snapshots, this indicates a primary or clone partition, and for clones this indicates a snapshot partition.
…

clones may have a primary or snapshot partion as their source and snapshots may have a primary or clone partion. i've added words to indicate that.

Symantec 8) Unprocessed Comment (Unprocessed) [183]

PDF pg 160, pg 136, 6.7.1, last p on pg

the text indicates that The IMMED_TR bit is described in 5.2.5, but it is not. it it described in 5.2.7.

Symantec 9) Unprocessed Comment (Unprocessed) [184]

PDF pg 161, pg 137, 6.7.1, 1st p after 1st a,b,c list on pg, s 2

CREATE SNAPSHOT
s/b
CREATE CLONE

Symantec 10) Unprocessed Comment (Unprocessed) [185]

PDF pg 164, pg 140, 6.7.3, p 2, s 1

be maintained to restarting
s/b
be maintained so as to facilitate restarting
…
the same wording exists in the CREATE SNAPSHOT section

Symantec 11) Unprocessed Comment (Unprocessed) [186]
PDF pg 172, pg 148, 6.10.2, 1st p after 1st a,b,c list on pg

change s/b chain
… check for similar errors globally

Symantec 12) Snapshots commands interlocking (Unresolved) [187]

PDF pg 178, pg 154, 6.12, first c) on pg

the DETACH CLONE, REFRESH SNAPSHOT OR CLONE, and RESTORE PARTITION FROM SNAPSHOT command have checks against create completion time and referesh completion time. i assume the reason for these checks is to try and interlock the operations (so a clone can't be detached while it's being created, refreshed, or restored).

unfortunately, the checks don't really accomplish this. for example, the DETACH CLONE command specifies that the command will fail if …

c) The create completion time attribute is undefined (see 3.1.51) and the refresh completion time attribute is undefined

this serves to protect partition while it's being created (because both timestamps are undefined), but does nothing to interlock with a refresh operation (since the create completion time attribute remains unchanged throughout the operation).

for the most part, these checks would accomplish what i assume they're intended to do if we made the following changes:

1) set Refresh Completion Time and Restore Completion Time attributes to the creation time when a partition is created

2) continue the current behavior of making Refresh Completion Time undefined while the the refresh is in progress

3) make restore completion time undefined while a restore is in progress

3) change the checks for "refresh completion time is undefined *and* creation completion time attribute is undefined" to or instead, and add checks against the refresh completion time attribute

this would also require that primary partitions have a restore completion time attribute that is set when they are created, which may not be possible.

as a suggestion, instead of using the sundry completion times for the purpose of interlocking, it would probably be simpler to simply check if a tracking well known collection exists in the partition being detached, refreshed, or created and fail the operation if it exists.

or we could even leave the mechanism out entirely, and simply specify that if one of the other operations is in progress, the DEATCH CLONE operation will fail (no specification of how this is determined).

there is a generic mechanism described in 4.6.6.6 for to insure only one multi-object command is running on a partition (or is it collection), but the snapshot and clone commands are not multi-object commands.

Symantec 13) CREATE SNAPSHOT interlocking w/ other commands (Unresolved) [188]

PDF pg 169, pg 145, 6.10

like the issues raised in Symantec 12, the CREATE SNAPSHOT command doesn't seem to have any interlocking with the REFRESH CLONE and RESTORE PARTITION FROM SNAPSHOT.

i think it should be an error to create a snapshot of a partition that is currently being modified by one of those other command.

Symantec 14) REMOVE PARTITION race condition (Unresolved) [189]

PDF pg 238, pg 214, 6.34

there may be a race condition similar to the one described in Symantec 12 with REMOVE PARTITION and RESTORE PARTITION FROM SNAPSHOT. what happens (i.e. what error is returned) if a REMOVE PARTITION is executed on a partition while RESTORE PARTITION FROM SNAPSHOT is running with the same partition as the target?

Symantec 15) Duplication Method s/b Snapshot attribute (Unresolved) [190]

PDF pg 320, pg 296, 7.1.3.30

duplication method does not appear as an attribute in the Snapshots Information attributes page. i think it should. adding this is also required in some later comments.

Symantec 16) REFRESH/RESTORE Duplication Methods (Unresolved) [191]

PDF pg 228, pg 204, 6.30

in 6.30 "REFRESH SNAPSHOT OR CLONE" and 6.35 "RESTORE PARTITION FROM SNAPSHOT", the duplication method can be defined, which means that it can be different than the duplication method used originally to create the partition.

for refresh, we should require that the same duplication method be used to refresh the partition as was used to create. for restore, we should require that the same duplication method used to create the snapshot be used to restore the source partition.

(if someone believes that the REFRESH SNAPSHOT OR CLONE should also support changing the duplication method, then i might withdraw this objection. however, i think it would be better to have a separate command for that purpose since this one will both change the duplication method and, potentially, change the contents).

Symantec 17) Snapshot commands restart flaws (Unresolved) [192]

PDF pg 231, pg 207, 6.30.2, last p in subclause

in 6.30 "REFRESH SNAPSHOT OR CLONE", the mechansim whereby an incomplete command can be completed by the REFRESH SNAPSHOT OR CLONE command seems to have a flaw or two.

While the snapshot/clone tracking well known collection is, roughly speaking, supposed to contain the set of objects that still need processing, the highlighted text doesn't make sense if we're completing an operation that was already in progress. in particular, adding all of the objects back to the collection seems to preclude processing only the objects that weren't completed.

Symantec 18) Overlapping RESTOREs (Unresolved) [193]

PDF pg 240, pg 216, 6.35

in 6.35 "RESTORE PARTITION FROM SNAPSHOT", i don't see a mechanism to prevent a partition from being restored from two different snapshots at the same time. we should create such a mechanism, or at least define the error to be returned if this is attempted.

Symantec 19) REFRESH/RESTORE and deleted object removal (Unresolved) [194]

PDF pg 240, pg 216, 6.35 & 6.30

in 6.35 "RESTORE PARTITION FROM SNAPSHOT" and 6.30 REFRESH SNAPSHOT OR CLONE, i don't see any clauses which require objects be removed from the target partition. i would assume that after the command runs, the target partition is an exact duplicate of the original partition, which should include removing any objects that were created after the snapshot was created.

Symantec 20) Unprocessed Comment (Unprocessed) [195]

PDF pg 242, pg 218, 6.35.4, last p on pg

processing complete
s/b
processing is complete

Symantec 21) Too many attributes controlling REFRESH (Unresolved) [196]

PDF pg 268, pg 244, 7.1.3.8, 4th p on pg

[see highlighted text]…however, there is a separate attribute, 311h Support for snapshot refreshing that is supposed to control whether or not snapshots can be refreshed (see text above Table 167):

"If it is defined (see 3.1.14), the support for snapshot refreshing attribute (number 311h) (see table 167) shall indicate how the REFRESH SNAPSHOT command (see 6.30) is supported. If the support for snapshot refreshing attribute is undefined (see 3.1.51), then the REFRESH SNAPSHOT command is not supported."

there are two obvious ways to fix this. because i think a profliferation of optional features is a bad idea, i'd suggest that we change the above text to require that the attribute be defined.

however, if we feel that this feature needs to be optional (separate from CREATE SNAPSHOT), then RESTORE PARTITION FROM SNAPSHOT should also be optional even when CREATE SNAPSHOT is defined. but we should lump REFRESH SNAPSHOT and RESTORE PARTITION FROM SNAPSHOT together (i.e. support both or neither).

8. Late Comments

The following issues were discovered during the letter ballot comments resolution process.

Other 1) Remove expected keyword (Accepted, Editorial) [197]
PDF pg 33, pg 9, 3.3.1

The expected keyword must be removed, per the decision of the September CAP working group (minutes in 08-354).