

ENDL TEXAS

Date: 23 June 2004
 To: T10 Technical Committee
 From: Ralph O. Weber
 Subject: Response to T10 Letter Ballot comments on OSD

This document contains the responses to the T10 Letter Ballot comments on forwarding OSD to first public review. The summary of the T10 Letter Ballot results can be found in document T10/04-064.

All references to OSD pages are based on osd-r09.pdf.

Revision History

- r0 All comments from T10/04-064r0 included as unprocessed comments.
- r1 Add 32 late comment from Brocade Communications as unprocessed comments.
- r2 All comments were reviewed and those comments needing discussion in the SNIA OSD TWG were marked as Unresolved. During the review, a few comments were processed to the point of having their resolutions defined and those resolutions were noted.
- r3 Proposed resolutions for all but 18 security-related comments.

Notes to Reviewers

Remember:

- The splitting of Credential and Capability validation is not in this revision
- Most of the other Security rewrites are not in this revision
- If you think a security feature is missing, review all two pages of comment IBM 141)

Navigation Aids:

- All comment numbers are hot links to the comment
 - Comment numbers in the lists of comments
 - Comment numbers in the comments resolution text
- Point and whack

To review the comments:

- Use the Bookmarks to find the kind of comments you want to review
 - some favorite are:
 - Substantive Comments Accepted As Proposed
 - Substantive Comments Accepted With Noted Changes
 - Rejected Comments List
 - Comments from a given company
 - The Resolution Summary Bookmark goes to a table counting the comments by company and resolution
 - The numbers in square brackets are running counters of all the comments in this document.

Warning

The responses to comments that are "Accepted With Noted Changes" can include anything short of totally reversing the intent of the comment (that would be rejecting the comment, right).

Review such comments carefully, especially if they are yours.

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.

Company	Unre-solved	Technical		Editorial		Rejected	Deferred or No Action Taken	Total
		As Is	Changed	As Is	Changed			
Agilent Technologies			1		6	2	1	10
AMCC					2	8	1	11
Brocade			2	3	9	14	4	32
EMC Corp.	9				2	2		13
ENDL Texas	1		2					3
Hewlett Packard Co.	2	1	16	38	41	18	6	122
IBM Corp.	3		18	25	76	26	5	153
Intel Corp.				4	3	9	5	21
Lingua Data	1		2	18	11	10	2	44
LSI Logic Corp.	1		3	5	5			14
Panasas	1		2	2	1		1	7
Seagate Technology		3	14	4	17	12	1	51
Veritas Software		2	5	46	27	31	10	121
Late Comments		3		7				10
Total	18	9	65	152	200	132	36	612

Unresolved Comments List

EMC 2) Nonce Processing May Depend on Clock Synchronization	38
EMC 3) Specify security manager behavior	38
EMC 4) Protection Against Unauthorized Usage of a Credential.	38
EMC 5) CAPKEY Equivalent In Strength To Password-In-The-Clear	39
EMC 7) CMDRSP and ALLDATA Require Nonce Verification.	39
EMC 8) Credential Validation Algorithm Not Correct.	40
EMC 9) Timestamp Checks Should Not Be Optional.	40
EMC 10) Far-In-The-Future Nonces Invite Record-And-Replay Attacks	40
EMC 13) Key Algorithms Do Not Provide Forward Secrecy	41
ENDL 2) Separate Credentials From Capabilities & Enable Fencing	42
HP 33) Channel ID issue	49
HP 35) Channel ID issue	49
IBM 14) Remove SPC-2 acronym	69
IBM 141) Capability Validation is NOT a Security Function	97
IBM 149) ABORT TASK and QUERY TASK for Non-Root Objects.	101
Lingua 19) Channel ID issue	113
LSI 2) Insert SNIA OSD TWG membership	120
Panasas 5) Per-Partition Security Is Wrong	124

Rejected Comments List

Agilent 1) Several data field termination and padding comments	19
Agilent 5) No partition zero	20
AMCC 1) Provide explicit non-volatile cache management	23
AMCC 2) Make data-in usage consistent	23
AMCC 3) Allow for smaller sized segments in the Data-In and Data-Out Buffers	23
AMCC 5) Require sense key specific data	24
AMCC 6) Add pointer to the table key	24
AMCC 7) Assign a new operation code group code to OSD commands.	24
AMCC 8) Correct CDB additional length field	24
AMCC 9) Clarify REPORT LUNS behavior	25
Brocade 5) Blocking access via quota changes	27
Brocade 6) Security ladder diagrams in informative annex	28
Brocade 7) Make Credential format match the one defined in IKE	28
Brocade 9) OBSD to protect against multiple concurrent commands	29
Brocade 12) Specify order of bidirectional transfers	30
Brocade 18) Root and Partition Directory pages should include attribute 0h (the label)	32
Brocade 20) Errors for unimplemented attributes pages	33
Brocade 22) Root Directory page is read-only	33
Brocade 23) Make previously specified changes in 7.1.2.5.	34
Brocade 24) Make previously specified changes in 7.1.2.6.	34
Brocade 25) Make previously specified changes in 7.1.2.7.	34
Brocade 30) Truncating a user object should be an error	36
Brocade 31) Are Security attributes pages should not be required	37
Brocade 32) Root Partition security dominates	37
EMC 6) Define Security Tokens Precisely	39
EMC 11) Capability key new for each Capability, not Credential.	40
HP 3) Dedicate OSD to Gene Milligan.	43
HP 5) Delete list of standards in SCSI family.	43
HP 9) Several data field termination and padding issues	44
HP 12) Unnecessary consistency in error response descriptions	45
HP 15) Unnecessary consistency in error response descriptions	45
HP 16) Unnecessary consistency in error response descriptions	45
HP 20) P, C, and R are difficult to search for use something else.	46
HP 22) needed [s/b] used	46
HP 24) Pepper standard with references to 4.7.2	46
HP 25) One attribute page number space across an entire logical unit.	47
HP 41) The Scope of a Credential is One Object	50
HP 44) Drop "objects" for consistency.	50
HP 45) Drop "object" for uniformity	50
HP 46) Describe the Future of the Object Descriptor.	51
HP 53) Report Maximum Number of Far-In-The-Future Nonces.	52
HP 97) Put 'quota' in all quota attribute names	60
HP 103) Assign different attribute names in different attributes pages	61
HP 108) Duplicate quota testing statement	62
IBM 6) Delete list of standards in SCSI family	67
IBM 7) Remove SPC-2 because there are no references to it.	68
IBM 9) Zero or more is wrong in security method definition.	68
IBM 11) Sense data is point of reference for sense key, Not!	68
IBM 13) 4.6.1 contains details about User_Object_IDs, Not!	68
IBM 16) Use an a,b,c list	69
IBM 19) Enhance list of commands that cannot be directed to a partition.	69

Rejected Comments List (continued)

IBM 21) Enhance list of commands that cannot be directed to a collection.	70
IBM 23) Eliminate '(see table x in x.x.x)' construct.	70
IBM 38) Send Credential To Device	74
IBM 45) Data includes meta data.	75
IBM 51) Remove 'in most cases'	77
IBM 54) Eliminate description of the threat addressed	77
IBM 68) Specify behavior if total capacity is exceeded	80
IBM 76) Add discussion of quota violation for root capacity	82
IBM 79) Provide single definition for allocation length	82
IBM 80) Provide single definition for additional length	83
IBM 81) Make common definition for LIST and LIST COLLECTION parameter data	83
IBM 82) Change PERFORM SCSI COMMAND name.	83
IBM 85) Change PERFORM TASK MANAGEMENT FUNCTION name.	83
IBM 92) Define key identifier and seed just once.	86
IBM 98) (number xh) [s/b] just (xh)	87
IBM 120) LIST COLLECTION affects Partition Timestamps	93
IBM 144) Add Object Type to CDB	100
IBM 150) Add Capacity Quota for OBSD.	101
IBM 151) Key Invalidation Changes.	102
Intel 2) OSD not in figure 1	104
Intel 5) contain [s/b] specify	104
Intel 8) Wordy partition capacity.	105
Intel 12) Get attributes suggestions	106
Intel 15) set [s/b] collection	107
Intel 16) Example sounds like a requirement.	107
Intel 18) 'considered' is not very precise	107
Intel 19) reduce [s/b] limit.	108
Intel 20) reduce [s/b] limit.	108
Lingua 7) Should SAM-2 be referenced?.	109
Lingua 9) Add User_Object_ID to user object definition	110
Lingua 13) Remove policy/storage manager from model	111
Lingua 15) Add Current Command attributes page to table 3	112
Lingua 18) Put Security Method values in the Model clause	112
Lingua 29) Increase Granularity of Permissions for Commands Accessed via PERFORM SCSI COMMAND.	115
Lingua 33) Use small caps for field contents	116
Lingua 34) Add Partition_ID Field in FORMAT OSD Command	117
Lingua 35) Task Management Functions Are Not Directed To User Objects	117
Lingua 42) Eliminate Page Formats for Attributes	118
Seagate 4) Enhance list of commands that cannot be directed to a collection	126
Seagate 16) Add example regarding changing security version key values	129
Seagate 17) Require Command Terminations if working keys restrict far-in-the-future nonces	130
Seagate 25) Missing Collection_Object_ID parameter description	132
Seagate 28) FLUSH OBJECT command should include collections.	134
Seagate 29) Reference subclass for COLLECTION_OBJECT_ID description	134
Seagate 31) Describe intended use of LIST IDENTIFIER field.	134
Seagate 32) Explain native command usage in PERFORM SCSI COMMAND definition.	135
Seagate 36) Add "lists format only" note	136
Seagate 40) "Default" not used consistently	137
Seagate 41) "Quota" not used consistently	137
Seagate 42) "Quota" not used consistently	137
Veritas 3) Align attributes definition with OSD command names.	141

Rejected Comments List (continued)

Veritas 5) object-based storage device definition is too redundant	142
Veritas 12) Use UTF-8 instead ASCII	143
Veritas 13) Encryption (Privacy) Of Credential Transmission To Application Clients	144
Veritas 14) Define what is outside the scope of this standard	144
Veritas 15) or as [s/b] or exist as	144
Veritas 17) OSD can assign user object IDs	145
Veritas 25) Clients Can Assign Partition/User Object IDs	148
Veritas 27) Clients Can Assign Partition IDs	148
Veritas 28) partition OSD object [s/b] OSD partition object	149
Veritas 30) partition OSD object [s/b] OSD partition object	149
Veritas 31) User objects contain Partition_IDs	149
Veritas 32) Make user object membership in one partition an enforceable requirement.	150
Veritas 33) Clients Can Assign User Object IDs	150
Veritas 34) Clients Can Assign Collection_Object_IDs	150
Veritas 38) Add Current Command attributes page to overview	151
Veritas 40) command functions [s/b] command function	152
Veritas 41) GET ATTRIBUTES Processing Order.	152
Veritas 46) meaning that [s/b] consequently	153
Veritas 48) Add a sentence break	154
Veritas 54) should [s/b] could & add cross reference	155
Veritas 60) attributes [s/b] attribute	156
Veritas 80) When both DPO and FUA are set	160
Veritas 81) PARTITION_ID field and CREATE PARTITION	161
Veritas 83) USER_OBJECT_ID field and CREATE and CREATE AND WRITE	161
Veritas 89) to the [s/b] when	163
Veritas 92) to the [s/b] when	163
Veritas 94) to the [s/b] when	164
Veritas 101) to the [s/b] when	166
Veritas 104) Clarify attribute access limitations	166
Veritas 115) get and set [s/b] get.	169

Comments With Implementation Deferred to OSD-2

Brocade 2) Define failure states. 26
Brocade 8) Add detail on deferred write errors 29
IBM 143) Add Create Attributes Permission 99

Substantive Comments Accepted As Proposed

HP 69) Eliminate Discussion of Encrypted CDBs 54
Seagate 12) Add Permission Bit for APPEND 129
Seagate 37) Partition Directory lists only partition attributes pages 136
Seagate 51) LIST COLLECTION listing partition contents is a partition function 140
Veritas 65) Capability Validation is NOT a Security Function 158
Veritas 106) Replace 'them' with specific wording 167
Other 6) Define values of the x+0h attributes in the directory attributes pages 173
Other 8) Increasing a user object's logical length can produce a quota error 174
Other 10) ACA requirements 174

Substantive Comments Accepted With Noted Changes

Agilent 7) What happens to the Credential Partition_ID when the Object Descriptor identifies a Partition?	20
Brocade 4) Test all quotas before the media is modified.	27
Brocade 26) Clock drift	34
ENDL 1) Finish Error Reporting Definition	42
ENDL 3) Persistent Reservations Should Be Optional	42
HP 26) Description of attribute number FFFF FFFFh not clear	47
HP 27) application clients may not [s/b] application clients shall not	47
HP 56) Master key belongs to the logical unit	52
HP 62) Fencing	53
HP 65) Remove Sense Data Implementation Requirements from SPC-3.	54
HP 70) Clarify fua bit on reads.	55
HP 72) Specify Current Command attributes page usage for created objects	55
HP 74) When to Test Quotas	56
HP 76) Define FLUSH COLLECTION Command	56
HP 77) Define FLUSH PARTITION Command	57
HP 81) REQUEST SENSE returns parameter data with GOOD status.	57
HP 90) Clarify the length and padding of directory attribute values.	59
HP 95) Username is binary	60
HP 110) Incorporate new timestamp bypass controls	63
HP 113) Setting a Collections attribute can cause a quota error	63
HP 120) Use OBSD consistently as target device	65
IBM 27) Clarify REMOVE processing	71
IBM 32) Attributes associated with any object type do not appear in any directory pages	73
IBM 36) Clarify quota errors sense data descriptor	74
IBM 42) NOSEC definition change	74
IBM 50) Root attributes access includes partition zero attributes access	76
IBM 72) Add discussion of quota violation for partition capacity	81
IBM 75) Add discussion of quota violation for partition capacity	81
IBM 77) What Does FLUSH COLLECTION do?	82
IBM 78) No progress indication in OSD, maybe in OSD-2.	82
IBM 91) SET KEY Cannot Get A Quota Error	85
IBM 108) CREATE and CREATE COLLECTION cause partition capacity quota violations	89
IBM 118) Updates of "access" timestamps should involve the application client.	91
IBM 128) Disallow sparse lists of supported integrity check value algorithms.	94
IBM 142) Security Hole in Data-Out Integrity Check	99
IBM 145) Provide Reporting for Collections Support	100
IBM 146) Provide Reporting for Volatile Storage Support	100
IBM 147) Extended The Non-Standard Communications Between the Security Manager and Application Clients	101
IBM 153) define pre SET MASTER KEY master key identifier	103
Lingua 14) Non-Support for the FORMAT OSD Command.	111
Lingua 30) CREATE and CREATE AND WRITE can retrieve attributes from any created user object.	115
LSI 5) Remove requirement for GOOD status	120
LSI 6) Confusing FORMAT OSD requirement to be removed	120
LSI 7) Remove potentially troublesome error recovery statement.	121
Panasas 2) Reporting credential and capability errors is critical	123
Panasas 4) When to Test Quotas	124
Seagate 2) Confusing FORMAT OSD requirement to be removed.	126
Seagate 7) Add Security Method Field to CDB	127
Seagate 8) Capability Validation is NOT a Security Function	127
Seagate 9) Add Security Method to Capability	128

Substantive Comments Accepted With Noted Changes (continued)

Seagate 11) Add Permission Bit for APPEND	128
Seagate 18) SHA1 shall not be truncated	130
Seagate 21) Working Key Generation Key Values Are Not Saved	131
Seagate 22) Prohibiting changes in unused Data-In Buffer bytes	131
Seagate 26) Update Timestamp Bypass Mechanism	132
Seagate 33) Removing a non-existent collection should not be an error.	135
Seagate 39) Copy new partition username from partition 0, not from the root	137
Seagate 44) Require Zero User Object Length in Partition Zero	138
Seagate 45) Require Zero Collections Per User Object in Partition Zero	138
Seagate 46) Require Zero Object Count in Partition Zero	138
Veritas 64) FORMAT OSD should require global permission?	157
Veritas 76) Prohibiting changes in unused Data-In Buffer bytes	160
Veritas 84) Persistent Reservations Should Be Optional.	161
Veritas 105) Use Length FFFF FFFFh for Undefined Attributes	167
Veritas 116) Does Changing the User Object Logical Length Update the Data Modified Timestamp?	170

Accepted As Proposed Non-Substantive Comments List

Brocade 1) store [s/b] stored	26
Brocade 14) Simplify the association between attributes and objects	31
Brocade 16) only one [s/b] one	31
HP 2) SNIA, not SINA	43
HP 4) Get the peripheral device type field name right	43
HP 7) Format URLs consistently	43
HP 14) Incorrect cross reference, (see 6.5) [s/b] (see 6.6)	45
HP 17) Processing [s/b] Process	45
HP 18) store [s/b] stored	46
HP 21) use [s/b] used	46
HP 23) Processing [s/b] Process	46
HP 30) Add missing semicolon on a,b,c list	48
HP 36) semicolon [s/b] colon	49
HP 37) semicolon [s/b] colon	49
HP 38) semicolon [s/b] colon	49
HP 39) semicolon [s/b] colon	49
HP 40) Add missing period at end of sentence	49
HP 54) of [s/b] set to	52
HP 55) of [s/b] set to	52
HP 61) lost [s/b] loss	53
HP 67) of [s/b] set to	54
HP 68) of [s/b] set to	54
HP 75) object count quota [s/b] partition count	56
HP 78) List [s/b] List of	57
HP 79) listed [s/b] all of the listed	57
HP 82) Add (MSB) and (LSB) to seed field in table 63	57
HP 83) Add (MSB) and (LSB) to seed field in table 65	58
HP 87) INCITS [s/b] the ASCII characters "INCITS"	58
HP 88) undefined [s/b] unidentified	58
HP 91) Put field names in small caps	59
HP 93) Change date format	59
HP 94) Serial number [s/b] Product serial number	60
HP 98) 0001 [s/b] 1	61
HP 99) 0001 [s/b] 1	61
HP 100) 0002 [s/b] 2	61
HP 112) Fix table 87 border	63
HP 114) Add pro forma attribute page format field definitions	64
HP 115) Change date format	64
HP 116) in the in the [s/b] in the	64
HP 121) SHA is Secure Hash Algorithm	66
HP 122) Both FIPS 180-1 and FIPS 198 are needed to define HMAC-SHA	66
IBM 1) Remove revision history	67
IBM 5) Pilot implementations are not relevant	67
IBM 8) Use commas, not parentheses	68
IBM 10) SPC-3 is the document of reference for the sense data definition	68
IBM 12) Reflect prohibition on linked commands in the glossary	68
IBM 15) Clarify that no other bytes can contain a null	69
IBM 18) many [s/b] any number of	69
IBM 22) Missing 'or' in a,b,c list	70
IBM 26) 'and' is on the wrong 1,2,3 list entry	71
IBM 39) a integrity [s/b] an integrity	74

Accepted As Proposed Non-Substantive Comments List (continued)

IBM 40) New sentence, not semi-colon	74
IBM 41) performed [s/b] processed	74
IBM 52) Eliminate Undetectable Alterations?	77
IBM 53) security version key [s/b] security version tag	77
IBM 56) meaning that [s/b] as a result	78
IBM 59) performed [s/b] processed	78
IBM 61) and [s/b] and/or	79
IBM 64) Make note 5 normative	79
IBM 66) Make table last CDB byte match text	79
IBM 83) performed [s/b] processed	83
IBM 84) performed [s/b] processed	83
IBM 86) performed [s/b] processed	84
IBM 87) performed [s/b] processed	84
IBM 88) taSk [s/b] Task	84
IBM 139) Eliminate field treatments	96
Intel 1) SNIA, not SINA	104
Intel 7) between [s/b] among	105
Intel 9) Insert missing comma	105
Intel 13) store [s/b] stored	106
Lingua 2) NCITS [s/b] INCITS	109
Lingua 3) NCITS [s/b] INCITS	109
Lingua 4) SNIA, not SINA	109
Lingua 5) an SCSI [s/b] a SCSI	109
Lingua 6) Devices [s/b] Device	109
Lingua 8) lost [s/b] loss	110
Lingua 10) Typo in Volatile Cache definition	110
Lingua 12) Add cross reference to user objects	111
Lingua 20) the same [s/b] the same protections	113
Lingua 23) Delete redundant paragraph	114
Lingua 26) OSO [s/b] OSD	114
Lingua 32) Incorrect cross reference	116
Lingua 37) An integrity check values [s/b] Integrity check values	117
Lingua 38) rest [s/b] reset	117
Lingua 39) lost [s/b] loss	118
Lingua 40) is [s/b] of.	118
Lingua 41) zero [s/b] NOSEC	118
Lingua 43) taSk [s/b] Task	118
LSI 1) Remove revision history	120
LSI 3) rest [s/b] reset	120
LSI 4) e.g. [s/b] i.e.	120
LSI 11) by command [s/b] by a command	121
LSI 13) rest [s/b] reset	122
Panasas 1) user object [s/b] partition.	123
Panasas 3) partition [s/b] collection	123
Seagate 1) partition [s/b] collection	126
Seagate 5) store [s/b] stored	127
Seagate 13) Incorrect cross reference.	129
Seagate 14) is [s/b] in	129
Veritas 1) Peer-to-peer does not align with common usage	141
Veritas 2) Peer-to-peer does not align with common usage	141
Veritas 4) from single [s/b] from a single	142

Accepted As Proposed Non-Substantive Comments List (continued)

Veritas 6) lost [s/b] loss	142
Veritas 7) greenwich mean time [s/b] Greenwich Mean Time	142
Veritas 9) Devices [s/b] Device	143
Veritas 11) same writing [s/b] same as writing	143
Veritas 18) An OSD contains [s/b] An OBSD contains	145
Veritas 24) Multiple Object Operations in this standard	147
Veritas 26) Insert missing comma	148
Veritas 29) user object [s/b] partition	149
Veritas 37) store [s/b] stored	151
Veritas 43) insert 'in order'	153
Veritas 45) OSD Device [s/b] OBSD	153
Veritas 47) does allow [s/b] allows	153
Veritas 49) Application Client Trust Assumption Rewording	154
Veritas 50) send [s/b] sent	154
Veritas 51) and the [s/b] and by the	154
Veritas 53) Incorrect cross reference	155
Veritas 57) to [s/b] into	156
Veritas 61) of OSD object [s/b] of the OSD object	157
Veritas 66) is [s/b] in	158
Veritas 67) is [s/b] in	158
Veritas 68) An integrity [s/b] Integrity	158
Veritas 69) is has [s/b] has	158
Veritas 70) Remove redundant 'defining the'	159
Veritas 71) As [s/b] As a	159
Veritas 72) insert 'a'	159
Veritas 73) rest [s/b] reset	159
Veritas 74) lost [s/b] loss	159
Veritas 75) is [s/b] of	159
Veritas 77) zero [s/b] NOSEC	160
Veritas 78) containing service action [s/b] containing a service action	160
Veritas 79) Remove redundant 'there is'	160
Veritas 82) updated or not updated [s/b] established	161
Veritas 85) Give that sentence a verb	162
Veritas 86) Give that sentence a verb	162
Veritas 91) of partition [s/b] of the partition	163
Veritas 100) object [s/b] objects	165
Veritas 102) CREATE AND WRITE [s/b] WRITE	166
Veritas 109) number partitions [s/b] number of partitions	168
Veritas 114) the allow [s/b] allowed	169
Veritas 118) an request [s/b] a request	170
Veritas 119) an request [s/b] a request	170
Veritas 120) an request [s/b] a request	171
Veritas 121) in the in the [s/b] in the	171
Other 1) Clarify lack of data area in the root object, partitions, and collections	172
Other 2) partition for [s/b] partition that represents	172
Other 3) Correct object count testing wording	172
Other 4) Correct partition count testing wording	173
Other 5) Use consistent notation in table B.1	173
Other 7) Remove 'successfully' globally	174
Other 9) Remove comma in table 39 column heading	174

Accepted With Noted Changes Non-Substantive Comments List

Agilent 2) An ASCII data field may be terminated by a null character	19
Agilent 3) Delete redundant paragraph	19
Agilent 6) Is there a partition zero?	20
Agilent 8) Some Permissions definitions overlap	21
Agilent 9) Make table 20 more closely match table 19	21
Agilent 10) The NONE object descriptor is not allowed for root objects	22
AMCC 4) Byte references are always relative to the user object.	24
AMCC 10) taSk [s/b] Task	25
Brocade 3) FrameMaker cross reference marker not on the correct paragraph	26
Brocade 10) Show the format of common OSD CDB fields.	30
Brocade 11) shall terminated [s/b] shall be terminated	30
Brocade 13) not themselves to encoded [s/b] not themselves to be encoded.	31
Brocade 15) them [s/b] attributes.	31
Brocade 17) Clarify title of table 67	31
Brocade 19) Attributes pages are never undefined	32
Brocade 21) Attribute number 0h example needed	33
Brocade 29) Support retrieving/setting attributes on any command is mandatory.	36
EMC 1) Capability keys must be kept secret	38
EMC 12) Remove discussion of very short capability key life times	41
HP 1) Peer-to-peer not appropriate in a SCSI abstract	43
HP 6) Format URLs consistently	43
HP 8) greenwich mean time [s/b] Greenwich Mean Time	44
HP 11) applications clients [s/b] application clients	44
HP 13) Insert missing comma	45
HP 19) 1,2,3 list 'and' is on the wrong list entry	46
HP 28) FrameMaker cross reference marker not on the correct paragraph	47
HP 29) Clarify Quota Revocation.	48
HP 31) Glossary has better description of Credential than Model clause	48
HP 32) Explicit Statement About No Authentication	48
HP 34) Nexus [s/b] nexus	49
HP 42) Root partition not defined.	50
HP 43) Delete redundant paragraph	50
HP 47) is [s/b] in.	51
HP 48) is [s/b] in.	51
HP 49) Incorrect cross reference	51
HP 50) Change date format	51
HP 51) Change 'drive key' to 'root key'	51
HP 57) As [s/b] As a.	52
HP 58) Remove discussion of very short capability key life times	53
HP 63) All OSD commands assume bidirectional data transfers.	53
HP 64) set attributes offset field missing from table 28	54
HP 66) Clarify forced reservation clearing	54
HP 71) FrameMaker cross reference marker not on the correct paragraph	55
HP 73) FrameMaker cross reference marker not on the correct paragraph	56
HP 80) Define Interaction of FORMAT OSD and CDB Set Attributes Parameters	57
HP 84) CREATE AND WRITE [s/b] WRITE & FrameMaker cross-reference bug.	58
HP 85) Indicate which attributes pages have a page format defined.	58
HP 86) Insert "null-terminated".	58
HP 89) The Collection Directory, User Object Directory and Null attributes pages do not have an attribute number 0h.	59
HP 92) Improper use of 'may not'.	59

Accepted With Noted Changes Non-Substantive Comments List (continued)

HP 101) Root Quotas page format incorrect	61
HP 102) Attribute number 0h is not a quota	61
HP 104) Root Quotas page format incorrect	62
HP 105) FrameMaker cross reference marker not on the correct paragraph	62
HP 106) FrameMaker cross reference marker not on the correct paragraph	62
HP 107) Attribute number 0h is not a quota	62
HP 109) Attribute number 0h is not a quota	63
HP 117) that [s/b] the	64
HP 118) all attributes [s/b] each attribute	64
HP 119) Clarify list format attribute number definition	64
IBM 2) Input/output logical units are the same as logical units	67
IBM 3) the [s/b] this	67
IBM 4) 200X [s/b] 2004	67
IBM 20) Three corrections in 4.6.6 (Partitions)	70
IBM 24) Written document is documentation	71
IBM 25) performed [s/b] processed	71
IBM 28) performed [s/b] processed	72
IBM 29) performed [s/b] processed	72
IBM 30) performed [s/b] processed	72
IBM 31) Publicly available vendor specific attribute page definitions enhance this standard	72
IBM 33) FrameMaker cross reference marker not on the correct paragraph	73
IBM 34) FrameMaker cross reference marker not on the correct paragraph	73
IBM 35) Eliminate gratuitous use of a,b,c list	73
IBM 37) 'appropriate capability' not specific enough	74
IBM 43) Clarify CAPKEY trust discussion	75
IBM 44) Make suggestion an example	75
IBM 46) Make suggestion an example	75
IBM 47) root partition [s/b] partition zero	76
IBM 48) The object creation time field is explained twice	76
IBM 49) Note 3 is normative	76
IBM 55) Add cross reference to far-in-the-future nonce processing	77
IBM 57) Remove 'meaning that'	78
IBM 58) Remove 'but'	78
IBM 62) What are offset fields?	79
IBM 63) What sense data descriptor?	79
IBM 65) Make table additional CDB length match text	79
IBM 67) semi-colon [s/b] and	80
IBM 69) Eliminate 'perform'	80
IBM 70) FrameMaker cross reference marker not on the correct paragraph	80
IBM 71) FrameMaker cross reference marker not on the correct paragraph	80
IBM 73) FrameMaker cross reference marker not on the correct paragraph	81
IBM 74) FrameMaker cross reference marker not on the correct paragraph	81
IBM 89) Use standard notation for CHECK CONDITION definition	84
IBM 90) Use standard notation for CHECK CONDITION definition	84
IBM 93) CREATE AND WRITE [s/b] WRITE	86
IBM 94) FrameMaker cross reference marker not on the correct paragraph	86
IBM 95) them [s/b] attributes	86
IBM 96) preferred [s/b] recommended	87
IBM 97) Improper keyword use in column heading	87
IBM 99) Match comment IBM 97) changes in text	87
IBM 100) Improper keyword use in column heading	88

Accepted With Noted Changes Non-Substantive Comments List (continued)

IBM 101) Match comment IBM 100) changes in text	88
IBM 102) Improper keyword use in column heading	88
IBM 103) Match comment IBM 102) changes in text	88
IBM 104) Improper keyword use in column heading	88
IBM 105) Match comment IBM 104) changes in text	89
IBM 106) Improper keyword use in column heading	89
IBM 107) Match comment IBM 106) changes in text	89
IBM 109) Improper keyword use in column heading	89
IBM 110) FrameMaker cross reference marker not on the correct paragraph.	90
IBM 111) FrameMaker cross reference marker not on the correct paragraph.	90
IBM 112) Match comment IBM 109) changes in text	90
IBM 113) Improper keyword use in column heading	90
IBM 114) the allow [s/b] the allowed	90
IBM 115) Match comment IBM 113) changes in text	91
IBM 116) Improper keyword use in column heading	91
IBM 117) Match comment IBM 116) changes in text	91
IBM 119) Improper keyword use in column heading	92
IBM 121) Match comment IBM 119) changes in text	93
IBM 122) Improper keyword use in column heading	93
IBM 123) Match comment IBM 122) changes in text	93
IBM 124) Improper keyword use in column heading	93
IBM 125) Match comment IBM 124) changes in text	94
IBM 126) Improper keyword use in column heading	94
IBM 127) Match comment IBM 126) changes in text	94
IBM 129) Improper keyword use in column heading	94
IBM 130) Match comment IBM 129) changes in text	95
IBM 131) 16 or sixteen, not both	95
IBM 132) Improper keyword use in column headings	95
IBM 133) Newest valid nonce, not oldest.	95
IBM 134) Match comment IBM 132) changes in text	95
IBM 136) Match comment IBM 132) changes in text	96
IBM 137) Match comment IBM 132) changes in text	96
IBM 138) Match comment IBM 132) changes in text	96
IBM 140) is [s/b] are.	97
IBM 152) user object [s/b] partition	102
Intel 4) 'Inherited' Is The Wrong Word	104
Intel 6) Abstracted?	105
Intel 21) 'drive key' is un-SCSI.	108
Lingua 1) Make ISO part number correct.	109
Lingua 11) Clarify character encoding description.	110
Lingua 16) FrameMaker cross reference marker not on the correct paragraph	112
Lingua 17) FrameMaker cross reference marker not on the correct paragraph	112
Lingua 22) CMDRSP protects more than just command and CDB parameters	113
Lingua 24) field says 'creation' but attribute says 'created'	114
Lingua 25) field says 'creation' but attribute says 'created'	114
Lingua 27) Clarify capabilities tables 19 and 20 introduction.	114
Lingua 28) At least one row per command in table 19.	115
Lingua 31) Add commas to clarify security attributes page exceptions	116
Lingua 36) Credential integrity check value is not used in a reconstructed credential.	117
LSI 8) 'and' is on the wrong 1,2,3 list entry	121
LSI 9) All OSD object types is confusing	121

Accepted With Noted Changes Non-Substantive Comments List (continued)

LSI 10) FrameMaker cross reference marker not on the correct paragraph	121
LSI 12) far in the future nonce(s) [s/b] far-in-the-future nonce(s).	122
LSI 14) Admit that no attribute page numbers are assigned by other standards.	122
Panasas 6) FrameMaker cross reference marker not on the correct paragraph	124
Seagate 3) Insert missing comma	126
Seagate 6) FrameMaker cross reference marker not on the correct paragraph	127
Seagate 10) Delete redundant paragraph	128
Seagate 15) Working key selection wording is confusing	129
Seagate 19) is has [s/b] has	130
Seagate 20) Integrity check values for SET KEY and SET MASTER KEY not clear.	130
Seagate 23) Inconsistencies in Data-In Buffer integrity check value segment	131
Seagate 24) Inconsistencies in Data-Out Buffer integrity check value segment	131
Seagate 27) FrameMaker cross reference marker not on the correct paragraph	133
Seagate 30) FORMAT OSD should not change security attributes.	134
Seagate 34) CREATE AND WRITE [s/b] WRITE & FrameMaker cross-ref bug	135
Seagate 35) Attribute number 0h exceptions list should include Collection Directory.	136
Seagate 43) FrameMaker cross reference marker not on the correct paragraph	138
Seagate 47) the allow [s/b] allowed	138
Seagate 48) shall [s/b] shall be	139
Seagate 49) the each [s/b] each	139
Seagate 50) SET ATTRIBUTES service action is 880Fh	139
Veritas 8) Two types in Volatile Cache definition.	142
Veritas 10) to encoded [s/b] to be encoded	143
Veritas 16) "Abstracted subsets" not clear.	145
Veritas 19) Clarify OSD logical unit	146
Veritas 20) Eliminate Root Object Data	146
Veritas 21) Eliminate Partition Data	146
Veritas 23) Eliminate Collection Data	147
Veritas 35) SBC-based systems is unclear	151
Veritas 36) All OSD object types is confusing	151
Veritas 42) Publicly available vendor specific attribute page definitions enhance this standard	152
Veritas 44) FrameMaker cross reference marker not on the correct paragraph	153
Veritas 52) security mode [s/b] security method	154
Veritas 58) as follows [s/b] by	156
Veritas 59) as follows [s/b] by	156
Veritas 62) Add a sentence break in place of 'and'	157
Veritas 63) three editorial changes in one sentence	157
Veritas 87) FrameMaker cross reference marker not on the correct paragraph	162
Veritas 88) Insert a comma	162
Veritas 90) FrameMaker cross reference marker not on the correct paragraph	163
Veritas 93) shall terminated [s/b] shall be terminated	164
Veritas 95) for returned list [s/b] for the returned list	164
Veritas 96) that command [s/b] that new command.	164
Veritas 97) that command [s/b] that new command.	165
Veritas 107) Clarify title of table 67	167
Veritas 108) Indicate which attributes pages have a page format defined	168
Veritas 112) FrameMaker cross reference marker not on the correct paragraph	169
Veritas 117) Clarify adjustable clock attribute definition.	170

No Action Requested, No Action Taken Comments List

Agilent 4) What permissions bit combinations are errors?	20
AMCC 11) On what basis are tags unique?	25
Brocade 27) How is the length of each attribute determined?	35
Brocade 28) All attributes pages specified by this standard are mandatory	35
HP 10) Some field termination and padding formats are not used	44
HP 52) shall not be accepted a second time with the same nonce value?	52
HP 59) In addition to the key in question or instead of?	53
HP 60) This is undefined	53
HP 96) Questions about sparse user objects	60
HP 111) Future timestamps	63
IBM 17) Text is not syntactically clear but no changes proposed	69
IBM 60) Should explicit non-volatile cache support be provided?	78
IBM 135) 16 or sixteen, not both	96
IBM 148) Redefined Buffer Offset Values	101
Intel 3) OSD is in figure 1	104
Intel 10) How does REMOVE COLLECTION work?	105
Intel 11) Questions about collections	106
Intel 14) Why GET ATTRIBUTES and SET ATTRIBUTES?	106
Intel 17) Is 'propriety' the right word?	107
Lingua 21) Correct security methods in SPC-3	113
Lingua 44) Require Task Management Functions be Sent To the Root Object.	119
Panasas 7) Collection_Object_IDs in same name space as User_Object_IDs	125
Seagate 38) Logical unit device identifier attribute disappeared from Root Information attributes page	136
Veritas 22) Is the root object a member of partition 0?	147
Veritas 39) an command [s/b] a command	152
Veritas 55) Why denial of service?	155
Veritas 56) Why denial of service?	155
Veritas 98) Make support for PERFORM SCSI COMMAND mandatory	165
Veritas 99) Make support for PERFORM TASK MANAGEMENT FUNCTION mandatory	165
Veritas 103) Who defines page format?	166
Veritas 110) OSD Initialized Attributes Should be 'Yes' In OSD Provided.	168
Veritas 111) OSD Initialized Attributes Should be 'Yes' In OSD Provided.	168
Veritas 113) OSD Initialized Attributes Should be 'Yes' In OSD Provided.	169

1. Agilent Technologies

Pat Thaler from Agilent Technologies submitted the following comments on a Yes vote.

Agilent 1) Several data field termination and padding comments (Rejected) [1] 3.7.2

Perhaps I'm being excessively picky, but some of the statements here seem to be contradictory:

"A data field that is described as being null-terminated shall have one byte containing an ASCII null. and all other bytes in the field shall not contain the ASCII null character. " but farther down a data field that is described as being both null-terminated and null-padded can have multiple null characters at the end. A "null- terminated and null-padded field" is a kind of "null-terminated field" so it shouldn't violate the null-terminated field rule. The first sentence could be modified to "being null-terminated but that is not null-padded".

Also Note 1 seems to be wrong - The two sentences above say there is a difference between pad byte contents - in one case they contain space and in the other case null. This sentence says the null-padded and zero-padded fields have the same pad byte contents and there is some difference in the other bytes of the field. Delete or correct the note.

Reason for Rejection: Regarding the 'null-terminated but that is not null-padded' issue (second paragraph of comment): This subclause more closely resembles a glossary than a set-theory formulation. The phrasing is, "If you read this, it means that." Of course it is necessary to read all the keyword phrases described in the subclause before determining the meaning of a statement found elsewhere in the standard, but that is the way standard generally work.

Regarding note 1 (third paragraph of the comment): The two sentences above the note describe null-padding and zero-padding. The discussion of space-padding is three sentences above the note. The note seems to be ideally located to shed additional light on the two paragraphs that immediately precede it.

Agilent 2) An ASCII data field may be terminated by a null character (Accepted, Editorial) [2] 3.7.1

As long as I'm being picky - according to the next section an ASCII data field can also contain the null character. Add "and the null code (00h).

The null code shall only be used where required for null-termination and null-padding.

Editor's Note: The cited sentence will be modified as follows:

ASCII data fields shall contain only ASCII graphic codes (i.e., code values 20h through 7Eh) **and may be terminated with one or more ASCII null (00h) characters.**

Agilent 3) Delete redundant paragraph (Accepted, Editorial) [3] 4.8.4.3 At the bottom of page 34

the paragraph referencing table 13 plus the table appear to duplicate the information in the previous paragraph (the two paragraphs starting "The OBJECT CREATION TIME field").

Editor's Note: The paragraph that does not reference table 13 will be removed.

Agilent 4) What permissions bit combinations are errors? (No Action Taken) [4]

4.8.4.3

Permissions bit mask, WRITE bit - Presumably the WRITE bit is either ignored for the root, partition and collection object types - or setting the bit to 1 is an error for these types. Which is it?

Editor's Note: This issue is covered by the first table note in table 19, table 20 as well as the definition of reserved, to whit:

Combinations of OBJECT TYPE field, PERMISSION BITS field, and OBJECT DESCRIPTOR TYPE field values not shown in this table and table 20 are reserved.

3.3.9 reserved: A keyword referring to bits, bytes, words, fields and code values that are set aside for future standardization. A reserved bit, byte, word or field shall be set to zero, or in accordance with a future extension to this standard. Recipients are not required to check reserved bits, bytes, words or fields for zero values.

[Receipt of reserved code values in defined fields shall be reported as error.](#)

Agilent 5) No partition zero (Rejected) [5]

4.8.4.3 NOTE 3

NOTE 3 looks incorrect. From 4.6.2 there is no partition associated with Partition_ID 0h. Table 2 says that Partitions have Partition_IDs from 10000 h to all. Suggest deleting the NOTE 3.

Reason for Rejection: The first row in table 2 shows the combination of Partition_ID 0h and User_Object_ID 0h belonging to the root object. Glossary entry 3.1.31 further clarifies as follows:

3.1.31 partition zero: The partition with the Partition_ID (see 3.1.30) zero. The partition for the root object (see 3.1.33).

Agilent 6) Is there a partition zero? (Accepted, Editorial) [6]

Comment not listed in 04-064r0

General

Is there a "root partition"/"partition zero" or is there not"? 4.6.2 defines a "root object" but not a "root partition" and it does not allow 0h as a Partition_ID. A number of other places refer to partition zero, partition 0h or the root partition: e.g. 4.9.4.1, table 18, 6.8, 4.9.5.1, 4.9.5.4, Table 22, 4.9.8.2, Table 46, 6.8. Please make consistent.

Editor's Note: This comment will be resolved as described in the response to comment IBM 20).

Agilent 7) What happens to the Credential Partition_ID when the Object Descriptor identifies a Partition? (Accepted, Substantive) [7]

Comment not listed in 04-064r0

4.9.4.2

It is possible that there could be cases where the Object type is PARTITION and the OBJECT DESCRIPTOR field contains a different value than the Credential Format's PARTITION_ID field. It appears that for the cases in Table 19 where the Object Type Name is PARTITION that only the SINGLE_OBJECT_ID field is checked against the CDB PARTITION_ID. Is the Credential Format PARTITION_ID field to be ignored when Object Type is PARTITION or is a mismatch to be detected as an error?

Editor's Note: As described in the response to comment IBM 141), the Partition_ID will be taken out of the Credential and placed in the Capability Object descriptor. Details of how this comment is being resolved can be found in 04-193.

Agilent 8) Some Permissions definitions overlap (Accepted, Editorial) [8]

Comment not listed in 04-064r0

Table 19

There are two conflicting entries that overlap for some commands (near the top of page 40) "A PERFORM TASK MANAGEMNT command with function code of ABORT TASK or QUERY TASK" and " Any PERFORM TASK MANAGEMNT command" both applied to Object Type ROOT an and 1OBJECT with SINGLE OBJECT_ID of zero.

An ABORT TASK or QUERY TASK command for the root object fits in both entries but the first only requires DEV_MGMT and the second requires DEV_MGMT and GLOBAL.

Suggest adding "except those with function code of ABORT TASK or QUERY TASK" to the second entry to eliminate the overlap.

A similar overlap occurs in Table 20 between each item with SET_ATTR and the equivalent items with SET_ATTR and SECURITY in the Permission Bits.

Alternatively, one might add a statement to 4.9.4.4 that a command or attributes operation is allowed if at least one row in the table allows it column even if another row also applies and doesn't allow it.

I think it is clearer to remove the overlap.

Editor's Note: The first paragraph of 4.9.4.4 (Credentials and commands allowed) will be modified as follows:

... the presence of specific combinations of values in capability fields as shown in table 19. **A command function is allowed if at least one row in table 19 allows it, even if a different row that applies does not allow it.**
 ... capability fields that allow those functions are shown in table 20. **Retrieving or setting attributes is allowed if at least one row in table 20 allows it, even if a different row that applies does not allow it.**

The first paragraph of 4.9.4.4 will be broken into 3 paragraphs to better distinguish the ideas being presented.

Agilent 9) Make table 20 more closely match table 19 (Accepted, Editorial) [9]

Comment not listed in 04-064r0

Tables 19 and 20

The format differences between tables 19 and 10 tables are odd. The capability field columns have a joint heading in Table 19 but the same columns in Table 20 don't. Position of those columns is on the right in Table 19 and on the left in Table 20. It would be nice if these were made to match.

Editor's Note: The capability fields header will be added to table 20 and the capability fields will be moved to the right-hand side of the table.

Agilent 10) The NONE object descriptor is not allowed for root objects (Accepted, Editorial) [10]

Comment not listed in 04-064r0

Table 20

For the items where Object Type Name is ROOT table 20 allows Object Descriptor Name to be NONE or 1OBJECT and places no restrictions on the 1OBJECT value. Table 19 requires the Object Descriptor Name to be 1OBJECT and the 1OBJECT value be zero when Object Type Name is ROOT. There doesn't seem to be any reason for the difference and it would cause odd behavior - a credential that works for getting/setting attributes but not for commands despite the permissions bits. Suggest putting the same constraints on Table 20 ROOT items as on Table 19.

Editor's Note: Good catch. Until very recently, the object type ROOT was allowed with a NONE object descriptor in table 19. That has been removed in table 19 and the change needs to be propagated to table 20.

Note that nowhere in table 20 are the contents of the 1OBJECT object descriptor specified. Making a special case for the root object is not appropriate.

2. AMCC

Neil Wanamaker from AMCC submitted the following comments on a Yes vote.

AMCC 1) Provide explicit non-volatile cache management (Rejected) [11]

The author marked this comment as technical.

pg 53, 4.10

Does not match 03-388r1. Provide for nonvolatile cache as in 03-388r1.

Reason for Rejection: This standard explicitly includes non-volatile cache as part of stable storage, as described in the cited subclause, to wit:

Individual OBSD (see 3.1.26) implementations may use whatever technologies they choose to implement stable storage (e.g., [an OBSD may implement stable storage as a combination of non-volatile random access memory](#) and disk devices).

Considering the resistance to explicit non-volatile storage management from RAID vendors during recent CAP meetings, one could easily argue that this standard has gotten it right.

AMCC 2) Make data-in usage consistent (Rejected) [12]

The author marked this comment as technical.

pg 55, 4.11.2

4.7.2 suggests that data-in segments could include get/set attributes. Make consistent.

Reason for Rejection: The words 'data-in' never appear in 4.7.2. Perhaps the confusing arises from the notion that a READ command might also set attributes. That is accomplished using both the Data-In Buffer and the Data-Out Buffer (i.e., all OSD commands are bidirectional). See the response to comment HP 63).

AMCC 3) Allow for smaller sized segments in the Data-In and Data-Out Buffers (Rejected) [13]

The author marked this comment as technical.

pg 57, 4.11.4

Minimum size for get/set segments is 256 bytes. Allow for smaller increments.

Reason for Rejection: The size was chosen to allow for all known virtual memory management mechanisms. Real world initiators cannot manage memory boundaries smaller than their virtual memory page size. The editor is aware of operating systems with 512 byte virtual memory pages and it has been confirmed that Windows uses 4,096 byte virtual memory pages.

AMCC 4) Byte references are always relative to the user object (Accepted, Editorial) [14]

The author marked this comment as technical.
pg 58, 4.13, par 3

It is unspecified whether the byte is within the compass of the user object or the physical device; the former requires massaging before becoming useful for maintainability. Specify frame of reference.

Editor's Note: Change the last sentence in the paragraph as noted:

If it is possible to identify a specific byte or range of bytes within a user object as being associated with an error, the information sense data descriptor (see SPC-3) shall be included in the sense data with the INFORMATION field set to the byte **within the user object** associated with the error or the first byte in the range of bytes **within the user object** associated with the error.

AMCC 5) Require sense key specific data (Rejected) [15]

The author marked this comment as technical.
pg 58, 4.13, par 7

should -> shall.

Reason for Rejection: Sense key specific data is helpful only in resolving problems where the application client sends an incorrectly formatted command. Essentially, sense key specific data is a debugging aid that the device server provides to poorly coded application clients. It is not the practice of SCSI standards to require that device servers assist in the debugging of application clients. If those who purchase OBSDs wish the debugging aid, they can require it as a sales condition.

AMCC 6) Add pointer to the table key (Rejected) [16]

pg 60, Table 30

RR not readily understood. Add note reference or ?? to draw attention to key at bottom of diagram

Reason for Rejection: Everybody knows how to find the M and O in the table key. Finding RR requires exactly the same skills.

AMCC 7) Assign a new operation code group code to OSD commands (Rejected) [17]

The author marked this comment as technical.
pg 61, 5.1

"Since the CDB length is fixed, why not assign a group code to these, and treat them as fixed?"

Reason for Rejection: Saving a byte or two in an OSD CDB is not worth the cost to the SCSI family of standards that would result from assigning a group code to OSD. Besides, OSD gets more usable operation codes (aka service actions) this way, and how knows, maybe OSD-4 will change the CDB size.

AMCC 8) Correct CDB additional length field (Rejected) [18]

The author marked this comment as technical.
pg 61, Table 31

"Either last entry sb n-1, or additional length = n-8" make consistent

Reason for Rejection: If there were only one additional byte in the CDB, the byte number on that byte would be 8 and the additional length would be 1 or 8-7. The cited values are correct as written. See also comments IBM 65) and IBM 66).

AMCC 9) Clarify REPORT LUNS behavior (Rejected) [19]

The author marked this comment as technical.
pg 90, Table 55, and following paragraph

The status of REPORT LUNS is somewhat ambiguous here - is support required for both LUN 0 and the well-known LUN? Neither? Clarify

Reason for Rejection: The behavior of the REPORT LUNS command should be clarified in SPC-3 not in this standard.

AMCC 10) taSk [s/b] Task (Accepted, Editorial) [20]

pg 92, Table 57

taSk

Editor's Note: taSk will be changed to Task.

AMCC 11) On what basis are tags unique? (No Action Taken) [21]

The author marked this comment as technical.
pg 92, 6.13

On what basis are tags unique? I_T_L? I_T_L_Object?

Editor's Note: OSD is not modifying the tag uniqueness requirements stated in SAM-3, nor should it since tag uniqueness is a SCSI transport protocol issue.

3. Brocade Communications

Robert Snively from Brocade Communications submitted the following comments on a Yes vote.

Brocade 1) store [s/b] stored (Accepted, Editorial) [22]

Page 20, 4.7.1, paragraph 3

"retrieved and store" should be "retrieved and stored"

Brocade 2) Define failure states (Deferred to OSD-2) [23]

The author marked this comment as technical.

Page 20, 4.7.2

There is significant work in specifying an order of behavior for each of the types of commands. However, I have not located anything that shows what the state of each of these should be when a failure is encountered during one of the steps. Clause 14.3 addresses that, but ending states are not defined there either.

Proposed Solution:

I believe that there needs to be either a normative annex or a normative statement in the text specifying the expected behavior and final state for failures in each of the steps specified in 4.7.2. Some of that may need to be included in particular commands, especially reads and writes that are also doing get attributes.

Editor's Note: The OSD failure states are almost as well defined as the SBC-2 failure states and sufficiently well defined to allow implementations consistent with the level of definition for other OSD features (e.g., the data location and data access pattern features defined in this standard are nil). The OSD-2 project proposal (see 04-173) includes improved error handling.

The the NOT INITIATED COMMAND FUNCTIONS field and COMPLETED COMMAND FUNCTIONS field defined in the OSD error identification sense data descriptor (see T10/04-095) may partially address this comment.

Brocade 3) FrameMaker cross reference marker not on the correct paragraph (Accepted, Editorial) [24]

Page 23, 4.8.1

There is an incompleting reference in the 5th text line of the clause. The reference should be completed or removed.

Editor's Note: The problem is a technical gaff in the FrameMaker source file. The cross reference marker used to generate the text in entry c of the a, b, c list is located on the paragraph before the header for the User Object Quotas attributes page subclause. The correct location for the cross reference marker is on the header. The cross reference marker will be moved and the cross references will be regenerated to correct this problem.

Brocade 4) Test all quotas before the media is modified (Accepted, Substantive) [25]

The author marked this comment as technical.
Page 24, 4.8.3

Quota testing errors may leave the state of the OSD logical unit in an unknown state. The state should be either exactly as before the command processing or exactly as it should be after the command processing. There should be no unannounced case of unknown state.

Proposed Solution:

In 4.8.3, "The processing ... has been detected." should be replaced with: "Quota error detection should be performed before media is modified. The device server shall be sure that the state of the OSD logical unit is unchanged if a quota error is detected."

Editor's Note: The following new paragraph will be added to 4.8.2 (Quota errors):

~~The device server shall not terminate a command for quota errors after any user data or attributes have been modified.~~

The following paragraph will be removed from 4.8.3 (Quota testing):

~~Tests for quota errors may be made at any time during the processing of a command. The processing of a command may be partially completed at the time a quota error is detected. The device server is not required to restore the state of the OSD logical unit to the state that was present before processing was begun for the command in which the quota error has been detected.~~

Brocade 5) Blocking access via quota changes (Rejected) [26]

The author marked this comment as technical.
Page 24, 4.8.4

The text indicates that changing quotas to values less than the resources already consumed is not an error and shall not result in the truncation or removal of any information. It would seem to me that necessarily results in the unavailability of those resources, and therefore of the data represented in those resources.

Proposed Solution:

In the appropriate section, the text should indicate that setting a quota to a value less than the resources consumed shall be an error. The resources must first be freed before the quota can be modified.

Reason for Rejection: The SNIA OSD TWG has discussed at length the effects of setting quotas lower than the resources already consumed. The behavior specified is the behavior desired. It allows access to be temporarily blocked by change quota values.

Brocade 6) Security ladder diagrams in informative annex (Rejected) [27]

The author marked this comment as technical.
Page 24, 4.9

The security model is a new concept to SCSI. I believe that the model information, particularly with respect to a tutorial overview, is perhaps less helpful to me than I would have hoped. Additional overview text is desirable. This is particularly true of the different security methods.

Proposed Solution:

As a starting point, example security ladder diagrams should be provided for each of the security methods, perhaps in an informational annex referenced in the text.

Reason for Rejection: Without seeing some sample ladder diagrams and understanding how they might be helpful, neither the OSD editor nor the SNIA OSD TWG is willing to invest the time to construct a complete set of ladder diagrams.

Brocade 7) Make Credential format match the one defined in IKE (Rejected) [28]

The author marked this comment as technical.
Page 32, 4.9.4

This credential structure does not appear to be the same as the standard credentials provided by various available credential services. As a result, the system is required to have a new and unique credential service.

Proposed Solution:

Modify the text to use a standard IKE credential, allowing the use of standard security services.

Reason for Rejection: OSD credentials are unlike credentials available from any existing credential service, so a new and unique OSD-specific credential service is required no matter what credential format is used. Standard IKE credentials cannot be used, as they are not functionally compatible with OSD (e.g., a common IKE credential is an X.509 certificate).

The IKE credential format is a tag, length, value format. Although such a format is delightfully extensible, it is not an efficient consumer of bytes. The number of bytes in a SCSI CDB is limited to 260 and the capability part of the credential must fit in those bytes. There are not enough spare bytes in the OSD CDB to allow the luxury of using the IKE credential format.

Brocade 8) Add detail on deferred write errors (Deferred to OSD-2) [29]

The author marked this comment as technical.
Page 53, 4.1

The use of a volatile cache and deferred errors does not provide an indication of what state the OSD shall establish for the data that failed to be written. Since the OSD has accepted the data, it has accepted responsibility for the data at a "file system" level. Data that was not successfully transferred to nonvolatile store must be identified as such in subsequent read operations, as well as through a deferred error.

Proposed Solution:

The paragraph "The device server... (see SPC-3)." should be rewritten to:

The device server may transfer data from the volatile cache to stable storage after status has been returned for the command that placed the data in the volatile cache. Errors that occur during such data transfer operations shall be reported as deferred errors as specified in SPC-3. Such data shall be flagged such that subsequent read operations will indicate that the data is not valid.

Error conditions should be established in the READ type commands for this.

Note that the loss or unavailability of meta-data under the same circumstances is unacceptable.

Editor's Note: The OSD-2 project proposal (see 04-173) includes improved error handling, which includes errors detected after a command completes.

Brocade 9) OBSD to protect against multiple concurrent commands (Rejected) [30]

The author marked this comment as technical.
Page 57, 4.12

There is a constraint specified that application clients shall not request two or more commands that may adversely affect the return of future data. I assume such commands include the assignment of object parameters by two commands, or possibly two separate write operations. I believe that the target should provide most required protections for this kind of event. In particular, get/set attributes for a command should be atomic with respect to the influence of other commands. In addition, writes by one command should be atomic with respect to writes by another command.

Proposed Solution:

Specify the rare conditions that may allow disruptive interactions among commands. Require that the target protect the application client against interactions among commands other than those caused by variations in the timing of command entry to the execution phase.

Reason for Rejection: In approving T10/03-257r0, the T10 technical committee when out of its way to allow SBC-2 devices not to protect application clients from the ill effects of concurrent commands. The OSD editor and the SNIA OSD TWG see no reason why OSD devices need to suffer the performance penalties necessitated by such protections any more than SBC-2 devices.

Brocade 10) Show the format of common OSD CDB fields (Accepted, Editorial) [31]

Page 62, 5.2

See also comments:

This text explains common format values in the CDB. However, the CDB overview has not yet been specified, so the placing of these values within the command requires unreferenced reading forward to particular commands. I used table 40 as the example of this.

Proposed Solution:

I propose expanding Table 31 to include a more complete and generic example of the 174 byte command structure. Things like the placement of the time stamps and other fields are invariant and can be defined in such a generic command space.

Editor's Note: This comment will be resolved as follows:

- The table 31 will be modified as follows: "~~Typical~~ Basic OSD CDB".
- The description of the last bytes in table 31 will be modified as follows: "Service action specific fields (see 5.2.1)".
- A new subclause will be added under 5.2 (Fields commonly used in OSD commands), "5.2.1 Overview".
- The new subclause will contain a table showing the commonly used CDB fields and their placement in OSD CDBs.

The OSD specific fields will not be added to table 31 because table 31 aligns with the common variable length CDB format shown in SPC-3.

Brocade 11) shall terminated [s/b] shall be terminated (Accepted, Editorial) [32]

Page 75, 6.4

"command shall terminated" should be changed to "command shall terminate".

Editor's Note: As with comment Veritas 93), the cited sentence will be modified as follows:

If the REQUESTED_USER_OBJECT_ID field contains any value other than zero and the device server is unable to assign the requested User_Object_ID to the created user object, the user object shall not be created and 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.

Brocade 12) Specify order of bidirectional transfers (Rejected) [33]

The author marked this comment as technical.

Page 73, 6.3 & Global

I first noticed this in page 73, clause 6.3. There is a missing tutorial piece in the model describing clearly how these commands are executed. It is my understanding from reading and re-reading that every command is CMD, then then DATA OUT, then DATA IN, then RSP. That means that any CREATE is required to return a set of parameters indicating, among other things, the ID of the created object. This is never clearly pointed out in the model, though it is implied in 4.7.2. I believe additional text, together with appropriate examples using ladder diagrams or equivalent explanatory diagrams would be very valuable. The examples may go in a referenced informational annex.

Reason for Rejection: As described in FCP-2 Annex F, the order of Data-In and Data-Out operations is chosen by the device sever, with no constraints on which is chosen first or which follows which, or how many of each there are.

That is the way this standard works.

Brocade 13) not themselves to encoded [s/b] not themselves to be encoded (Accepted, Editorial) [34]

Page 11, 3.6.1

"not themselves to encoded." should be changed to "not themselves encoded."

Editor's Note: The word 'be' will be added to the cited text, as shown in the comment title.

Brocade 14) Simplify the association between attributes and objects (Accepted, Editorial) [35]

Page 14, 4.2

"All stored data objects (see 4.6) have attributes (see 4.7) associated with them." should be changed to "All stored data objects (see 4.6) have associated attributes (see 4.7)."

Brocade 15) them [s/b] attributes (Accepted, Editorial) [36]

Page 104, 7.1.1

"Page format parameter data allows retrieval of them in formatted pages where only the attribute values appear in the parameter data." should be changed to "Page format parameter data allows retrieval of the attributes in formatted pages where only the values of the attributes appear in the parameter data."

Editor's Note: The cited sentence will be changed as follows:

Page format parameter data allows retrieval of ~~them~~ attributes in formatted pages where only the attribute values appear in the parameter data.

Brocade 16) only one [s/b] one (Accepted, Editorial) [37]

The author marked this comment as technical.

Page 18, 4.6.3

"There is only one root object" should be "There is one root object"

Brocade 17) Clarify title of table 67 (Accepted, Editorial) [38]

The author marked this comment as technical.

Page 105, 7.1.2.1

If I understand the attributes page definitions correctly, there are a bunch more reserved values not mentioned here. Table 3 and 4 may be intended to do that. The tables are not referenced from 7.1.2.1, creating the impression of an incomplete set.

Proposed Solution:

Add reference near table 67 to tables 3 and 4 for other usage of attribute page numbers.

Editor's Note: The title for table 67 is misleading. It will be modified as follows: "Attributes pages **defined by this standard**"

Brocade 18) Root and Partition Directory pages should include attribute 0h (the label) (Rejected) [39]

The author marked this comment as technical.
Page 106, 7.1.2.2

Why are the Root and Directory Partition pages exceptions to the rule about labeling the pages? I believe that the exception specified in the first sentence and second sentence of the clause should be removed.

Proposed Solution:

Remove the exceptions in the first paragraph of clause 7.1.2.2, requiring all pages to be labeled.

Reason for Rejection: Making the requested change would cause confusion in the definitions for the cited Directory attributes pages. For example, the Root Directory attributes page would look like this:

Table 69 — Example Root Directory attributes page contents

Attribute Number	Attribute Value (ASCII characters)
0h	"INCITS T10 Root Directory"
R+0h	"INCITS T10 Root Directory"
R+1h	"INCITS T10 Root Information"
R+2h	"INCITS T10 Root Quotas"
R+3h	"INCITS T10 Root Timestamps"
R+5h	"INCITS T10 Root Security"

The result would be Letter Ballot comments requesting explanations for the duplicated attributes values.

Brocade 19) Attributes pages are never undefined (Accepted, Editorial) [40]

The author marked this comment as technical.
Page 106, 7.1.2.3

Incorrect dynamic generation of a page does not make the page undefined. It either makes the page incorrect/invalid, or it makes the page unavailable. The page is still identifiable by the system that is trying to create it. I believe that there should be a value of "unavailable" that is used instead of the "undefined" that is suggested.

Proposed Solution: unimplemented

Change "undefined" to "unavailable" in the first paragraph of 7.1.2.3. Change the text string to "attributes information unavailable" in the second paragraph of 7.1.2.3.

Editor's Note: All three instances of "undefined" in the subclause will be changed to "unidentified", making them match the subclause title. This includes the text string cited by the comment. Also, "undefined" will be changed to "unidentified" in each of the four subclauses that reference 7.1.2.3, specifically 7.1.2.4, 7.1.2.5, 7.1.2.6, and 7.1.2.7.

Brocade 20) Errors for unimplemented attributes pages (Rejected) [41]

The author marked this comment as technical.
Page 104, 7.1.1

What is the behavior if an attributes page has an established definition but is not implemented?

Proposed Solution:

Provide error codes for unimplemented pages.

Reason for Rejection: The design goal of OSD is that all attributes pages and all attribute values may be set and retrieved by application clients unless otherwise specified. OSD 09 implements that goal. The goal is not stated in OSD r09 because standards are not tutorial. There is no need to define error codes.

Brocade 21) Attribute number 0h example needed (Accepted, Editorial) [42]

The author marked this comment as technical.
Page 107, 7.1.2.4

An expanded example of the root attributes page in table 69 would be helpful. It is not necessarily very clear that the "Attribute Value" in the title page is a character string inserted in bytes 8-39 of table 68. I believe there may be some naming problems. As an example, the right title of Table 69 s/b "Value placed in Attributes Page Identification field" or something like that.

Proposed Solution:

Clarify Table 69, providing a fully expanded example, at least in informative text in an annex. Correlate the names between the different values more accurately.

Editor's Note: The attribute value shown in table 69 is bytes 0-39 of the bytes shown in table 68. Taking the attribute number R+0h as an example, "INCITS " is bytes 0-7 and the null-padded "T10 Root Directory" is bytes 8-39. As far as the editor can tell, all the names are correlated precisely.

The concepts are not sufficiently complex to require an informative annex. The most that seems necessary is adding the following note to the end of 7.1.2.2:

Note: Using the User Object Directory attributes page as an example, the VENDOR IDENTIFICATION field contains the ASCII characters "INCITS" and the ATTRIBUTES PAGE IDENTIFICATION field contains the ASCII characters "T10 User Object Information". The attribute number 0h attribute value is "INCITS T10 User Object Directory".

Brocade 22) Root Directory page is read-only (Rejected) [43]

The author marked this comment as technical.
Page 107, 7.1.2.4

The last paragraph on the page appears to be a very complicated way of saying that the Root Directory attributes page is "read-only".

Proposed Solution:

Provide a tutorial sentence explaining that, because the page cannot be set, the following error conditions will occur when an attempt is made to set the values.

Reason for Rejection: The interpretation of the cited paragraph is correct. Standards are not tutorial. Therefore, no changes are necessary.

Brocade 23) Make previously specified changes in 7.1.2.5 (Rejected) [44]

The author marked this comment as technical.
Page 108, 7.1.2.5

Include the necessary corrections from Brocade 020, 021, and 022.

Reason for Rejection: The changes made in response to comment Brocade 21) affect only subclause 7.1.2.2. Comments Brocade 20) and Brocade 22) were rejected.

Brocade 24) Make previously specified changes in 7.1.2.6 (Rejected) [45]

The author marked this comment as technical.
Page 109, 7.1.2.6

Include the necessary corrections from Brocade 020, 021, and 022.

Reason for Rejection: The changes made in response to comment Brocade 21) affect only subclause 7.1.2.2. Comments Brocade 20) and Brocade 22) were rejected.

Brocade 25) Make previously specified changes in 7.1.2.7 (Rejected) [46]

The author marked this comment as technical.
Page 110, 7.1.2.7

Include the necessary corrections from Brocade 020, 021, and 022.

Reason for Rejection: The changes made in response to comment Brocade 21) affect only subclause 7.1.2.2. Comments Brocade 20) and Brocade 22) were rejected.

Brocade 26) Clock drift (Accepted, Substantive) [47]

The author marked this comment as technical.
Page 111, 7.1.2.8

How is the clock set?

Proposed Solution:

Clause 4.9.5.2 says this is outside the scope of the standard. While a nice evasion, it leaves a non-functional technology, since the clock values must be known to some accuracy for some of the security and authentication to work. I propose that clause 4.9.5.2 specify a value of accuracy for the clock, to be set by a mechanism outside the standard. Some protocols are hard pressed to get the number under 1 second, and many clocks drift on the order of several seconds per year as clock processing interacts with the hardware counters.

Editor's Note: The following text will be added at the end of the paragraph describing the clock attribute in 7.1.2.8 (Root Information attributes page):

The clock attribute value should not gain or lose more than one second in any 24-hour interval.

The following text will be added at the end of the paragraph describing the adjustable clock attribute in 7.1.2.20 (Root Security attributes page):

The adjustable clock attribute value should not gain or lose more than one second in any 24-hour interval.

Brocade 27) How is the length of each attribute determined? (No Action Taken) [48]

The author marked this comment as technical.
Page 111, 7.1.2.8

How is the length of each attribute determined? As an example, the Unit Serial Number VPD page has a length count, but the Product Serial Number itself is defined by that length count. I have not located a comparable length field in the attributes explanation to this point.

Proposed Solution:

Either there are some length fields missing that prevent valid parsing of the attributes pages, or the explanation and tutorial information is not adequate to explain these values. It is possible that all of these values are exactly 32 bytes (after padding), but at least two are allowed to be longer.

Editor's Note: The length of an attribute value is determined by the standard, product specification, set attributes command function, etc. that defines or stores it. The lengths of attributes defined by this standard appear in the tables that define the attributes. How defining attribute lengths will be handled elsewhere is outside the scope of this standard.

As for the length of the serial number attribute, note that table 73 specifies the length as variable. Since there is no page format for the Root Information attributes page, the attributes in the Root Information attributes page can only be retrieved using the list attributes format, which is tag, length, value encoded.

Brocade 28) All attributes pages specified by this standard are mandatory (No Action Taken) [49]

The author marked this comment as technical.
Page "All", All

I believe that there is a fundamental problem with interoperability among devices implementing this standard. While "T10" attribute pages are defined, there is an explicit assumption that these are only a small subset of the attribute pages expected to be defined. That assumption is made in 7.1.2.2. That assumption, together with the absence of any referenced profiles or operational descriptions, has the possibility of creating interoperability nightmares.

Proposed Solution:

Specify that all T10 attribute pages and attributes supporting an implemented function are mandatory. Of course the choice of not implementing optional functions may decrease the actual number of attribute pages supported. Further specify that all OSD devices will operate correctly with only the T10 attributes specified. Vendor specific extensions shall not be required for T10 compliant operation of a device. Devices with vendor specific attributes set shall be accessible and operable by devices using only the T10 subset of defined attributes. This may be specified somewhere, but I sure did not see it. The obvious place to document this would be clause 4.7.1.

Editor's Note: The definition of the various Directory attributes pages seems to state that all attributes pages defined by this standard are mandatory to implement.

The remainder of the proposed changes are effectively statements about vendor specific behaviors and attributes, all of which are clearly beyond the scope of the standard.

Note that the response to comment IBM 145) makes the Collections attributes page optional to implement and shows how an optional to implement attributes page is represented in this standard.

Brocade 29) Support retrieving/setting attributes on any command is mandatory (Accepted, Editorial) [50]

The author marked this comment as technical.

Page 20, 4.7.2

The text says: "OSD commands provide the application client with the optional ability to get and set attributes as part of processing the command (e.g., a WRITE command may also retrieve the user object logical length attribute)." It appears to me that the capability is mandatory for the target to implement, since it must support functions like CREATE AND WRITE. While the application client can, at first glance, get away with executing only commands that do not require a simultaneous inbound/outbound operation and do not necessarily require treating attributes at the same time operations are going on, that falls into the "invocable" class typically used by profiles to identify mandatory behavior.

Proposed Solution:

There are three choices here:

- a) Confess that the ability to set and get attributes as part of the processing is a mandatory behavior of HBAs supporting OSD devices. Modify the text in 4.7.2 accordingly.
- b) Identify certain commands and operation combinations as optional for both the device server and the application client. Provide appropriate VPD information to specify which combinations are supported by the device server and make it mandatory that the application client not exercise those combinations. Commands invoking those combinations would be considered as errors.
- c) Change the text in 4.7.2 to indicate that the ability is not optional, but rather may be invoked at will by the application client. This would also require a short explanation of how things could be done without invoking those functions. In particular, the CREATE AND WRITE function would have to be singled out as prohibited for such application clients. This may also effect those security actions that require atomic behavior.

Editor's Note: Option a) is the intent of OSD. The word "optional" will be removed from the first sentence in 4.7.2.

Brocade 30) Truncating a user object should be an error (Rejected) [51]

The author marked this comment as technical.

Page 115, 7.1.2.11

The text indicates that changing the user object logical length to a value shorter than the logical length known to the OSD device server shall cause truncation. This is a data integrity exposure and large warning bells should go off when such an event occurs.

The text also conflicts with some of the definitions associated with quotas, which are related to object logical length. Particularly see 4.8.4.

Proposed Solution:

In 7.1.2.11, add text to the explanation that indicates that:

- a) the occurrence of a truncation shall be accompanied by an appropriate check condition, or
- b) the attempt to execute a command that creates a truncation shall be prohibited and a check condition provided. Shortening an object would require copying data to a new shorter object and removing the old object.

Note that a) would require corresponding text changes in 4.8.4.

Reason for Rejection: Truncating an object by setting its User Object Logical Length attribute is an intended feature of OSD. It is not an error.

The User Object Logical Length attribute is in the User Object Information attributes page, not the User Object Quotas attributes page, meaning that decreasing the value in the User Object Logical Length attribute cannot produce a quota error, meaning that 4.8.4 does not apply and need not be modified.

Brocade 31) Are Security attributes pages should not be required (Rejected) [52]

The author marked this comment as technical.
Page 132, 7.1.2.20

If the supported security methods attribute indicates that NOSEC is the supported security attribute, what is the status of all the other security pages.

Proposed Solution:

I believe that several of the security pages are optional and need only be supported if supported security methods other than NOSEC are supported. Those pages should be identified accordingly in their own sections. It is possible this is buried in the security model somewhere and I missed it. If it is, a reference would be sufficient in this section.

Reason for Rejection: The Root Security attributes page is required in all cases, to indicate that only the NOSEC security method is supported if nothing else. The Collection and User Object Security attributes pages will transition to Policy/Storage Management attributes pages as part of the splitting of Credentials and Capabilities described in comment IBM 141), since they only contain a Version Tag attribute and that attribute will stop being a Security attribute and become a Policy/Storage Management attribute. That leaves only the Partition Security attributes page, hardly worth overhead of making it optional.

Besides, standards with fewer options are better standards.

Brocade 32) Root Partition security dominates (Rejected) [53]

The author marked this comment as technical.
Page 135, 7.1.2.21

In this section and perhaps others, security methods are defined for the Partition. I would have expected that security methods other than those indicated as supported in the Root Partition would be considered as errors, though I do not find anywhere that is explicitly stated.

Proposed Solution:

I believe that this should be specified each place table 99 or its equivalent is referenced. It is possible this is buried in the security model somewhere and I missed it. If it is, a reference would be sufficient near each table.

Reason for Rejection: The design model for OSD is different partitions for different operating systems. With this in mind, the information in the root partition is only the default values to be used in a partition until the operating system that is using that partition establishes its own security modus operandi.

4. EMC Corp.

David Black from EMC Corp. submitted the following comments on a No vote.

EMC votes "No" due primarily to the following four comments:

EMC 3) is a significant functional omission from the standard.
EMC 5), EMC 9), and EMC 13) appear to be significant weaknesses or omissions in the specified security.

EMC 1) Capability keys must be kept secret (Accepted, Editorial) [54]

Section 4.9.2 Trust Assumptions

A crucial assumption has been omitted: the capability keys must be kept secret. All of the security properties depend on the attacker not being able to obtain the capability key that goes with a credential. This requires that the unspecified Application Client to Security Manager protocol be protected against eavesdropping (e.g., via encryption). In addition all participants (OBSD, security manager, and application) must protect capability keys from being disclosed to unauthorized parties.

Editor's Note: Details of how this comment is being resolved can be found in 04-193.

EMC 2) Nonce Processing May Depend on Clock Synchronization (Unresolved) [55]

Section 4.9.2 Trust Assumptions

The nonce processing described in Section 4.9.6.2 may depend on synchronized clocks. This assumption needs to be stated, along with the need to protect clocks from being changed by an attacker. There is some text about this in 4.9.5.2, but it is applied only to credential expiry, which may be less sensitive to clock synchronization than nonce validation.

Editor's Note: Details of how this comment is being resolved can be found in 04-193.

EMC 3) Specify security manager behavior (Unresolved) [56]

The author marked this comment as being a principle source of the EMC "No" vote.

Section 4.9.2 Trust Assumptions

There is no specification of the security functions of the security manager. This needs to be provided to explain how the server generates credentials and capability keys. This comment is tagged as being against Section 4.9.2, but is actually about an omission in the entire security section.

Editor's Note: Details of how this comment is being resolved can be found in 04-193.

EMC 4) Protection Against Unauthorized Usage of a Credential (Unresolved) [57]

4.9.3.1 Introduction

The table seems to define "forgery of credential" as the ability of an attacker to construct a new valid credential from scratch. Based on the comment on CAPKEY (see comment EMC 5), another line appears to be needed to capture the "unauthorized use of credential" threat. In its current form, CAPKEY does not protect against this threat.

Editor's Note: Details of how this comment is being resolved can be found in 04-193.

EMC 5) CAPKEY Equivalent In Strength To Password-In-The-Clear (Unresolved) [58]

The author marked this comment as being a principle source of the EMC "No" vote.

4.9.3.3 CAPKEY

See also comments EMC 6), HP 33), HP 34), HP 35), Lingua 19), and Lingua 20)

CAPKEY is equivalent in strength to password-in-the-clear in that an attacker who can observe a credential and integrity check value can make use of both via impersonating the initiator, as the integrity check value is calculated over only the initiator and target identities. In essence, the CAPKEY <credential, integrity check> pair functions as a limited use password for the initiator with the target - this is weaker than the CHAP authentication in iSCSI and Fibre Channel. CAPKEY's integrity check should be redefined to prevent reuse of the <credential, integrity check> value pair in a different I_T Nexus. One possibility would be to have the target generate a nonce for each I_T nexus, and include that value in all integrity check calculations on that I_T nexus. The passive observer who does not have the capability key is then thwarted by the fact that this nonce changes when she creates a new I_T nexus to impersonate the initiator.

Editor's Note: Details of how this comment is being resolved can be found in 04-193.

EMC 6) Define Security Tokens Precisely (Rejected) [59]**4.9.3.3 CAPKEY**

The specification of the security token is insufficiently precise about format and padding of the initiator and target identifiers. This is sufficiently important that appealing to the SCSI transport specs for the definition of these identifiers is not sufficient - they should be defined here with precise instructions on how to construct the security token.

Reason for Rejection: The requested change is incompatible with SAM layering. The definition of initiator and target identifier formats is to be found in the applicable transport protocol standard, not in a command set standard such as OSD. Having command set standards place requirements on transport protocol standards would restrict the options open to the developers of future transport protocols.

It would be appropriate for SCSI transport protocol standards wish to define their security token formats. Although it is something of a stretch, SAM-4 might be able to define a security token format that fits all protocols. However, OSD already says all it can say on the subject.

The current security token definition generates no additional avenues of attack. Both the initiator and the target must have the same security token format in order for the capability key algorithms to work. Incompatible security token formats only causes failures of credential validation process, most probably leading to both ends deciding to use CMDRSP security method instead.

EMC 7) CMDRSP and ALLDATA Require Nonce Verification (Unresolved) [60]**4.9.3.4 and 4.9.3.5 CMDRSP and ALLDATA**

These sections need to state that nonce verification is required as part of these security methods.

Editor's Note: Details of how this comment is being resolved can be found in 04-193.

EMC 8) Credential Validation Algorithm Not Correct (Unresolved) [61]

4.9.5.3 Reconstructing the Credential

Reconstruction steps 1) A), 1) B) a), and 1) B) b) modify the credential received in the CDB in a fashion that will cause device server integrity verification of a credential to fail if the modification causes any change. These steps appear to implement permission enforcement (e.g., CREATE PARTITION requires a credential for the root partition). At a minimum, this needs to be stated explicitly, but the specification would be improved by

- incorporating the A) check into the credential validation algorithm.
- incorporating the B a) and B) b) checks into Table 19 (e.g., the partition ID in the credential MUST be zero for a CREATE PARTITION command)

Editor's Note: Details of how this comment is being resolved can be found in 04-193.

EMC 9) Timestamp Checks Should Not Be Optional (Unresolved) [62]

The author marked this comment as being a principle source of the EMC "No" vote.

4.9.6.2 Device server validation of request nonces

The fact that timestamp checks are optional allows record-and-replay attacks on device servers that don't use timestamp checks. The specific attack of concern is:

- Stimulate client to produce commands to be replayed and record them
- Corrupt the commands in flight (e.g. overwrite nonces) so the server doesn't record their nonces being used
- Replay the commands at some later point in time

The root of this problem is that a device server cannot detect replay of a nonce that it has never seen - this is not a problem for sequence-based and time-based mechanisms. The timestamp checks should be made mandatory including recommendations for delta values to avoid this exposure.

Editor's Note: Details of how this comment is being resolved can be found in 04-193.

EMC 10) Far-In-The-Future Nonces Invite Record-And-Replay Attacks (Unresolved) [63]

4.9.6.3 Far-in-the-future nonces

This appears to be an invitation to record-and-replay attacks even in the presence of timestamp checks. It needs to be removed from the standard or heavily justified.

Editor's Note: Details of how this comment is being resolved can be found in 04-193.

EMC 11) Capability key new for each Capability, not Credential (Rejected) [64]

4.9.8.1 (Secret keys) Introduction

The Capability key row in Table 22 should say that the key is new with each Capability, not Credential.

Reason for Rejection: 4.9.5.4 (Computing the credential integrity check value) states that the computation of the credential integrity check value (aka, capability key, per 4.9.4.2 – Capability key) involves all the fields in the Credential except the CREDENTIAL INTEGRITY CHECK VALUE field. Since a change in any Credential field necessarily results in a change in the credential integrity check value (aka, capability key), it seems reasonable to say that the Capability key is new for each Credential.

EMC 12) Remove discussion of very short capability key life times (Accepted, Editorial) [65]

4.9.8.1 (Secret keys) Introduction

Footnote d refers to a short time interval during which the Capability key is valid. The fact that the time interval must be short (and what "short" means) is not specified. This is probably part of the larger problem specification of security manager functions has been omitted.

Editor's Note: The cited sentence contains no useful information and will be removed.

EMC 13) Key Algorithms Do Not Provide Forward Secrecy (Unresolved) [66]

The author marked this comment as being a principle source of the EMC "No" vote.

4.9.8.3 Computing updated generation keys and new authentication keys

This functionality cannot provide forward secrecy. If any master key is ever compromised, all keys based on it (including subsequent master keys) are potentially compromised. This can be disastrous.

Versions of SET KEY and SET MASTER KEY are needed that perform some sort of cryptographic key exchange so that forward secrecy can be provided if desired. A particularly disastrous situation is that when the OBSD owner changes, if the old owner (who knows the previous master generation key value) can observe the seed value in the new owner's SET MASTER KEY command, she can determine the new master keys (generation and authentication). This is undesirable for obvious reasons.

Editor's Note: Details of how this comment is being resolved can be found in 04-193.

5. ENDL Texas

Ralph O. Weber from ENDL Texas submitted the following comments on a No vote.

ENDL 1) Finish Error Reporting Definition (Accepted, Substantive) [67]

The author marked this comment as technical.

PDF pg 1, pg i, Global

See also comments: AMCC 4), IBM 35), IBM 36), and IBM 63)

Incorporate 04-095 to define additional quota enforcement features, new sense data descriptors, and additional error reporting features.

Editor's Note: No additional quota enforcement features (other than the new sense data descriptor used for quota errors) appear in 04-095.

ENDL 2) Separate Credentials From Capabilities & Enable Fencing (Unresolved) [68]

The author marked this comment as technical.

PDF pg 1, pg i, Global

related comments: HP 62), IBM 141)

Incorporate 04-100 to separate capability checking from security, the object version tag from security, and define basic object fencing features.

Editor's Note: Details of how this comment is being resolved can be found in 04-193.

ENDL 3) Persistent Reservations Should Be Optional (Accepted, Substantive) [69]

The author marked this comment as technical.

PDF pg 87, pg 68, 6.1, table 39

see also: comment Veritas 84)

PERSISTENT RESERVE IN and PERSISTENT RESERVE OUT should be optional implement, despite the statements made in approved document T10/02-260r1 and disregarding the written request from Ed Gardner that T10/02-260r1 be observed in OSD. These two commands make virtually no sense in an OSD device because of the vastly more versatile security capabilities feature. Requiring their implementation places an unreasonable burden on OSD products.

Editor's Note: The M's in table 39 will be changed to O's.

6. Hewlett Packard Co.

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

HP 1) Peer-to-peer not appropriate in a SCSI abstract (Accepted, Editorial) [70]

Page vii, Abstract

What is "peer-to-peer" about OSD? It follows the usual SCSI model of client/server =initiator/target.

Editor's Note: As described in comment Veritas 1), 'peer-to-peer' will be removed from the abstract.

HP 2) SNIA, not SINA (Accepted, Editorial) [71]

Page xviii, Page xviii

Change "Storage Industry Network Association" to "Storage Networking Industry Association (see <http://www.snia.org>)"

HP 3) Dedicate OSD to Gene Milligan (Rejected) [72]

Page xix, Introduction

Add a Gene Milligan dedication page since he edited the original drafts of this standard. See SBC-2 for picture and text.

Reason for Rejection: Dedicating too many standards to an individual reduces the significance of all dedications. Beyond all doubt, Gene Milligan contributed significantly to the success of T10 and SCSI. However, at least one version of SPI and one version of SBC have already been dedicated to Gene. That is sufficient.

HP 4) Get the PERIPHERAL DEVICE TYPE field name right (Accepted, Editorial) [73]

Page 1, 1 Scope

Change "device type field" to "PERIPHERAL DEVICE TYPE field" with small caps

HP 5) Delete list of standards in SCSI family (Rejected) [74]

Page 2, 1 Scope

Delete "At the time this standard was generated, examples of the SCSI general structure included:" and the list that follows on pages 2-3. It's always out of date and generally irrelevant to OSD itself.

Reason for Rejection: I have to maintain this list in conjunction with the SPC-3 version descriptors. As long I have to deal with the version descriptors, this list is appearing in my standards.

HP 6) Format URLs consistently (Accepted, Editorial) [75]

Page 5, 2.5 Ref under def

Format "(www.t10.org)" like other web page references

Editor's Note: In addition to the cited URL, www.ietf.org will be changed to <http://www.ietf.org> in subclause 2.4.

HP 7) Format URLs consistently (Accepted, Editorial) [76]

Page 5, 2.5 Ref under def

Format "(www.incits.org)" like other web page references

HP 8) greenwich mean time [s/b] Greenwich Mean Time (Accepted, Editorial) [77]

Page 8, 3.1.47 greenwich

Capitalize the city name Greenwich (and possible Mean and Time too). According to the reference and other pages, UTC is probably more correct than UT.

Editor's Note: 'greenwich mean time' will be changed to 'Greenwich Mean Time'.

UT will not be changed to UTC. It takes <http://aa.usno.navy.mil/faq/docs/UT.html> two large paragraphs of explanation to introduce the UTC acronym and the explanation appears to be necessary since UTC apparently stands for Coordinated Universal Time. Furthermore, the first sentence of the reference clearly compares UT (not UTC) to Greenwich Mean Time, to wit:

The times of various events, particularly astronomical and weather phenomena, are often given in "Universal Time" (abbreviated UT) which is sometimes referred to, now colloquially, as "Greenwich Mean Time" (abbreviated GMT).

HP 9) Several data field termination and padding issues (Rejected) [78]

Page 13, 3.7.2 Data field termination and padding requirements

Delete "and all other bytes in the field shall not contain the ASCII null character." and/or phrase the condition as "as being null-terminated but not null-padded" Null padding means other bytes are set to null. There are current no uses of "null-terminated" in the standard except in 3.7.2. Perhaps the attribute value strings (see 7.1.2.4 for the first example) should be - or are they null-terminated null-padded? If so, then just "null-padded" should suffice.

Reason for Rejection: Null-terminated means terminated with a single null character. The cited subclause defines it that way. Null-padded is defined later in the same subclause and is clearly a different concept.

Because each attribute has a length there is no global need for attribute values to be either null-terminated or null-padded.

HP 10) Some field termination and padding formats are not used (No Action Taken) [79]

Page 13, 3.7.2 Data field termination and padding requirements null-terminated and null-padded

There is only one use of "null-padded" in the standard, and no uses of null-terminated on its own.

Editor's Note: To avoid having to add one or more of the set of field termination and padding definitions when future OSD changes employ them, all the known cases have been included in the dpANS. Doing so reduces the potential for future oversights in this regard. The situation is not all that dissimilar from the inclusion of the keyword obsolete (see 3.3.7) even though the obsolete keyword cannot be used in this standard since there is no previous version of the standard from which to make a feature obsolete.

HP 11) applications clients [s/b] application clients (Accepted, Editorial) [80]

Page 16, 4.4 Elements of the example configuration

Delete clients (or make application singular)

Editor's Note: The cited sentence will be modified as follows:

The security manager may reside in the OBSD, in ~~applications~~ application clients, or as a separate entity, but the security requirements on the communications mechanism shall not change based on the location of the security manager.

HP 12) Unnecessary consistency in error response descriptions (Rejected) [81]

Page 18, 4.6.3 Root object

Change "status, the sense key shall be set to ILLEGAL REQUEST and the additional sense code shall be set to INVALID FIELD IN CDB." to "status status with the sense key set to ILLEGAL REQUEST and the additional sense code set to INVALID FIELD IN CDB."

Reason for Rejection: The cited description correctly describes the required behavior. The grammatical structure of the description is consistent with similar descriptions in other SCSI standards. There is no reason to change the text.

HP 13) Insert missing comma (Accepted, Editorial) [82]

Page 18, 4.6.2 Identifying OSD objects

Add comma between object and each

Editor's Note: A second comma will be added between "collection" and "and", to whit:

The combination of Partition_ID and User_Object_ID uniquely identifies the root object, each partition, each collection, and each user object. Partition_ID and User_Object_ID values are assigned as shown in table 2.

HP 14) Incorrect cross reference, (see 6.5) [s/b] (see 6.6) (Accepted, Editorial) [83]

Page 18, 4.6.4 Partitions

Broken cross-reference to CREATE COLLECTIONS; 6.5 should be 6.6

HP 15) Unnecessary consistency in error response descriptions (Rejected) [84]

Page 19, 4.6.4 Partition object

Change "status, the sense key shall be set to ILLEGAL REQUEST and the additional sense code shall be set to INVALID FIELD IN CDB." to "status status with the sense key set to ILLEGAL REQUEST and the additional sense code set to INVALID FIELD IN CDB."

Reason for Rejection: The cited description correctly describes the required behavior. The grammatical structure of the description is consistent with similar descriptions in other SCSI standards. There is no reason to change the text.

HP 16) Unnecessary consistency in error response descriptions (Rejected) [85]

Page 19, 4.6.6 Collections

Change "status, the sense key shall be set to ILLEGAL REQUEST and the additional sense code shall be set to INVALID FIELD IN CDB." to "status status with the sense key set to ILLEGAL REQUEST and the additional sense code set to INVALID FIELD IN CDB."

Reason for Rejection: The cited description correctly describes the required behavior. The grammatical structure of the description is consistent with similar descriptions in other SCSI standards. There is no reason to change the text.

HP 17) Processing [s/b] Process (Accepted, Editorial) [86]

Page 20, 4.7.2

Change Processing to Process

HP 18) store [s/b] stored (Accepted, Editorial) [87]

Page 20, 4.7.1 Overview

Change store to stored (or set)

HP 19) 1,2,3 list 'and' is on the wrong list entry (Accepted, Editorial) [88]

Page 20, 4.7.2

Move and to row 4) from row 3)

Editor's Note: The 'and' will be moved from row 2) to row 3).**HP 20) P, C, and R are difficult to search for use something else (Rejected) [89]**

Page 21, 4.7.3 Attributes pages (global)

P, C, and R are difficult to search for. How about slightly longer acronyms like P_BASE, C_BASE, and R_BASE?

Reason for Rejection: Search for P+, C+, or R+. The purpose of the shorthand is to be short, not long.**HP 21) use [s/b] used (Accepted, Editorial) [90]**

Page 21, 4.7.3 Attributes pages, Table 3 - Attributes page numbers

Change use to used

HP 22) needed [s/b] used (Rejected) [91]

Page 21, 4.7.3 Attributes pages

Change needed to used

Reason for Rejection: Changing to 'used' would eliminate the point of the cited sentence. There is no need to define a constant for attribute page numbers for user object because the value of such a constant would be zero.**HP 23) Processing [s/b] Process (Accepted, Editorial) [92]**

Page 21, 4.7.2

Change Processing to Process

HP 24) Pepper standard with references to 4.7.2 (Rejected) [93]

Page 21, 4.7.2 Command function ordering

Each command section that has unusual processing orders (GET ATTRIBUTES, SET ATTRIBUTES, REMOVE, REMOVE PARTITION, and REMOVE COLLECTION) should cross-reference back to here. If you get the Root Information page number of partitions value, it reflects the value before not after removing a partition.

Reason for Rejection: Subclause 4.7.2 describes the command ordering for every OSD command. The processing of every command defined in this standard is affect by the rules in this subclause in the same way that the processing of every command defined in the standard is affect by the information in clause 7. The requested cross references attempt to segregate commands in an unrealistic manner.

HP 25) One attribute page number space across an entire logical unit (Rejected) [94]

Page 21, 4.7.3 Attributes pages

It is not clear if the attributes page numbers form one numbering space across the entire logical unit. Please state the scope explicitly.

Reason for Rejection: It is not clear because it is not true. For example, each user object within a logical unit has its own set of attributes pages and the attribute page numbers that go with them.

HP 26) Description of attribute number FFFF FFFFh not clear (Accepted, Substantive) [95]

Page 22, 4.7.4 Attributes

Not sure I follow this, "all attribute values" attribute is represented in what format? And how is this useful beyond a boolean value (i.e. beyond whether or not attributes are present)? Or is just a boolean value meant here?

Editor's Note: The cited sentence will be modified as follows:

The attribute number FFFF FFFFh may be used to request the retrieval of all the attributes in a page having a non-zero attribute length (see 7.1.3). If an attributes list specifies the setting of attribute number FFFF FFFFh, 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 INVALID FIELD IN PARAMETER LIST. ~~shall represent all attribute values in the page, not an attribute value.~~

HP 27) application clients may not [s/b] application clients shall not (Accepted, Substantive) [96]

Page 23, 4.7.5 Attributes directories

Change may not to shall not

Editor's Note: The cited sentence will be removed:

~~Application clients may not modify the contents of an attributes directory page in any other way.~~

The comment correctly notes that there is a requirement here. However, the requirement is on the device server to prevent improper modification of the attributes directory, not on the application client to refrain from attempting to do so.

The necessary requirements on the device server are stated in 7.1.2.4, 7.1.2.5, 7.1.2.6, and 7.1.2.7, so they need not be restated here.

HP 28) FrameMaker cross reference marker not on the correct paragraph (Accepted, Editorial) [97]

Page 23, 4.8.1 [Quotas] introduction

This entire line is a messed up cross reference. "c) The COLLECTIONS PER USER OBJECT field contains the value of the collections per user object attribute. (see)." It should just be User Object Quotes attributes page (see 7.1.2.14)."

Editor's Note: The problem is a technical gaff in the FrameMaker source file. The cross reference marker used to generate the text in entry c of the a, b, c list is located on the paragraph before the header for the User Object Quotas attributes page subclause. The correct location for the cross reference marker is on the header. The cross reference marker will be moved and the cross references will be regenerated to correct this problem.

HP 29) Clarify Quota Revocation (Accepted, Editorial) [98]

Page 24, 4.8.4 Changing quotas

"As long as the quota value remains set to a value that is less than the applicable resources already consumed, all commands that attempt to consume the applicable resource shall be terminated with a quota error."

If:

- 1) the quota is set to a high value Z
- 2) objects are creating consuming that amount
- 3) the quota is reduced to a lower value Y
- 4) a command is run which reduces the number of objects but not to a value below Y...

Is the command required to be accepted or might it run into a quota problem?

This wording could be interpreted as letting such commands be terminated, because they end up still above the quota. It should allow any commands which reduce the number of objects even if the current quota is exceeded.

Editor's Note: The cited sentence will be modified as follows:

As long as the quota value remains set to a value that is less than the applicable resources already consumed, all commands that attempt to consume ~~the additional~~ applicable resource shall be terminated with a quota error.

HP 30) Add missing semicolon on a,b,c list (Accepted, Editorial) [99]

Page 24, 4.9.1

add; after row a)

HP 31) Glossary has better description of Credential than Model clause (Accepted, Editorial) [100]

Page 24, 4.9.1 Basic security model

This description oversimplifies the fact. Definitions 3.1.4 and 3.1.11 are better descriptions, suggest adopting text from them here.

Editor's Note: The changes resulting from comment IBM 141) affect the cited paragraph because of the clear split between credentials (security) and capabilities (access controls). However, the point made by the comment is a good one and will be accounted for in the final comment IBM 141) changes.

HP 32) Explicit Statement About No Authentication (Accepted, Editorial) [101]

Page 26, 4.9.1 Basic security model

I think it helps to state that the OSD device server does not authenticate via an explicit protocol. An indirect authentication takes place via the checking of the integrity checksums.

Editor's Note: The last paragraph in the cited subclause will be modified as follows (including changes requested by comment Veritas 50):

The security manager may authenticate the application client, but the OSD device server does not ~~need to~~ authenticate the application client. It is sufficient for the OSD device server to verify the capabilities and integrity check values ~~send sent~~ by the application client.

HP 33) Channel ID issue (Unresolved) [102]

Page 27, 4.9.3.3 The CAPKEY security method
See comment EMC 5)

These terms are obsolete in SAM-3: "1) Initiator identifier; and 2) Target identifier." (used 3 times each in this section) Use "Initiator port identifier" and "Target port identifier" if that is what is intended.

Editor's Note: Details of how this comment is being resolved can be found in 04-193.

HP 34) Nexus [s/b] nexus (Accepted, Editorial) [103]

Page 27, 4.9.3.3 The CAPKEY security method

Change Nexus to nexus (global except when referring to a Nexus argument in a function call)

Editor's Note: In addition to the cited paragraph, this change will be made in table 30. Also, Nexus Loss will be changed to nexus loss in the cited paragraph.

HP 35) Channel ID issue (Unresolved) [104]

Page 27, 4.9.3.1 Introduction
See comment EMC 5)

I think the following is implied in this text. Please make it explicit if it's truly implied.

The application client shall recompute the security token on receiving the I_T Nexus Loss event notification.

Editor's Note: Details of how this comment is being resolved can be found in 04-193.

HP 36) semicolon [s/b] colon (Accepted, Editorial) [105]

Page 29, 4.9.3.5.

Change ; to : in "segments (see 4.11.3);"

HP 37) semicolon [s/b] colon (Accepted, Editorial) [106]

Page 30, 4.9.3.5

Change ; to : in "Buffer;"

HP 38) semicolon [s/b] colon (Accepted, Editorial) [107]

Page 31, 4.9.3.5

Change ; to : in "segments (see 4.11.2);"

HP 39) semicolon [s/b] colon (Accepted, Editorial) [108]

Page 31, 4.9.3.5

Change ; to : in "Buffer;"

HP 40) Add missing period at end of sentence (Accepted, Editorial) [109]

Page 32, 4.9.4.1

Add . after "in 4.9.4.3"

HP 41) The Scope of a Credential is One Object (Rejected) [110]

Page 32, 4.9.4.1 Credential format

Somewhere around here, it needs to be stated that the scope of a credential is one object on the OSD logical unit. OBJECT CREATION TIME that shows up later is the first hint to this effect.

Reason for Rejection: Adding such a statement would be more confusing than helpful. In addition, the intention is that eventually the scope of one Capability (or one Credential) will be larger than a single object. There is little point in adding a statement now that will have to be found and removed during the OSD-2 Letter Ballot.

HP 42) Root partition not defined (Accepted, Editorial) [111]

Page 32, 4.9.4.1 Credential format

It would help if it is clarified whether the credential on root "partition" is qualitatively different from that on any other object (I think the answer is no). This sentence without any additional clarification would make the reader wonder if there's something here that's different.

Editor's Note: This comment will be resolved as described in the response to comment IBM 47), eliminating the term 'root partition' in favor of the glossary defined term 'partition zero'.

HP 43) Delete redundant paragraph (Accepted, Editorial) [112]

Page 34, 4.9.4.3 Capability format

This entire paragraph is completely redundant with Table 13. Suggest deleting.

Editor's Note: Actually, the cited paragraph is redundant with the paragraph that introduces table 13 and table 13. The specific text to be removed is shown in the response to comment Seagate 10).

HP 44) Drop "objects" for consistency (Rejected) [113]

Page 35, 4.9.4.3 Capability format, Table 14

Drop "objects" for consistency.

Reason for Rejection: Consistency is in the eye of the beholder. The glossary defines root object (3.1.33), partition (3.1.29), collection (3.1.6), and user object (3.1.48). Table 14 is consistent with the glossary and that is the consistency that is important.

HP 45) Drop "object" for uniformity (Rejected) [114]

Page 35, 4.9.4.3 Capability format, Table 14

Drop "object" for uniformity.

Reason for Rejection: Uniformity is in the eye of the beholder. The glossary defines root object (3.1.33), partition (3.1.29), collection (3.1.6), and user object (3.1.48). Table 14 is consistent with the glossary and that is the uniformity that is important.

HP 46) Describe the Future of the Object Descriptor (Rejected) [115]

Page 36, 4.9.4.3 Capability format, Table 14

This elaborate protocol structure with OBJECT DESCRIPTOR TYPE is simply confusing. Seeing that 1OBJECT is the only allowed type, I can only surmise that this structure is intended to allow future expansions that allow a single set of permission bits (and capability) to be shared across multiple objects. If this true, it is good to state that explicitly. As of now, the Object Descriptor Type is a simple valid/invalid bit.

Reason for Rejection: T10 standards do not speculate about the future. Too many things can happen to make such speculations wrong.

HP 47) is [s/b] in (Accepted, Editorial) [116]

Page 45, 4.9.5.2 Credential and capability validation

is should be in or delete and change use to used

Editor's Note: 'is' will be changed to 'in', per comments Seagate 14) and Veritas 66).

HP 48) is [s/b] in (Accepted, Editorial) [117]

Page 45, 4.9.5.2 Credential and capability validation

is should be in or delete and change use to used

Editor's Note: 'is' will be changed to 'in', per comment Veritas 67).

HP 49) Incorrect cross reference (Accepted, Editorial) [118]

Page 45, 4.9.5.2 Credential and capability validation

should be 4.9.3.3 or not here at all

Editor's Note: The cross reference will be corrected.

HP 50) Change date format (Accepted, Editorial) [119]Page 47, 4.9.6.1 Request nonce format
see also comment HP 93) and comment HP 115)

Change "January 1, 1970" to "1 January 1920"

Editor's Note: The cited text will be changed from "~~January 1, 1970~~" to "**1 January 1970**". The date will not be changed from 1970 to 1920.

HP 51) Change 'drive key' to 'root key' (Accepted, Editorial) [120]Separate comment that was not numbered by author.
Global, 4.9.5.4 Computing the credential integrity check value (and global)
See comment Intel 21)

What is a "drive"? Change "drive key" to "root key"

Editor's Note: Must also fix "device key" which is the same thing. Details of how this comment is being resolved can be found in 04-193.

HP 52) shall not be accepted a second time with the same nonce value? (No Action Taken) [121]

Page 48, 4.9.6.2 Device server validation of request nonces

shall not be accepted a second time with the same nonce value?

Editor's Note: The requirement is correctly stated. The purpose of the nonce in this context is to prevent replay attacks. A command cannot be replayed by an attacker if every command is required to have a unique nonce value.

HP 53) Report Maximum Number of Far-In-The-Future Nonces (Rejected) [122]

Page 48, 4.9.6.2 Device server validation of request nonces

The maximum number of far in the future nonces might be useful to know too.

Reason for Rejection: Reporting the maximum number would simplify attempts to attack the OSD by sending too many far-in-the-future nonces.

HP 54) of [s/b] set to (Accepted, Editorial) [123]

Page 49, 4.9.6.3.2

Change "a sense key of ILLEGAL REQUEST, and an additional sense code of SECURITY AUDIT VALUE FROZEN." to "a sense key set to ILLEGAL REQUEST, and an additional sense code set to SECURITY AUDIT VALUE FROZEN."

HP 55) of [s/b] set to (Accepted, Editorial) [124]

Page 49, 4.9.6.3.3

Change "a sense key of ILLEGAL REQUEST, and an additional sense code of SECURITY WORKING KEY FROZEN;" to "a sense key set to ILLEGAL REQUEST, and an additional sense code set to SECURITY WORKING KEY FROZEN;"

HP 56) Master key belongs to the logical unit (Accepted, Substantive) [125]

Page 50, 4.9.8.1 Introduction

What is the granularity of ownership implied here? 3.1.26 for OBSD says that it's a SCSI device. I tend to think that a SCSI LU is implied here for SET MASTER KEY.

Editor's Note: The cited text in table 22 will be modified as follows:

Change of ~~OBSD (see 3.1.26)~~ logical unit owner

HP 57) As [s/b] As a (Accepted, Editorial) [126]

Page 50, 4.9.8.1 Introduction

Change As to "As a" or delete entirely

Editor's Note: The cited text will be changed to 'As a'

HP 58) Remove discussion of very short capability key life times (Accepted, Editorial) [127]

Page 50, 4.9.8.1 Introduction
See comment EMC 12)

Please give a sense of what "very short" is - of the order of minutes or microseconds? If the answer is "it depends", then the same can be attributed to the "is very short" description.

Editor's Note: The cited sentence contains no useful information and will be removed.

HP 59) In addition to the key in question or instead of? (No Action Taken) [128]

Page 51, 4.9.8.1 Introduction

In addition to the key in question or instead of?

Editor's Note: The cited text contains neither the word 'additional' nor the word 'instead'

HP 60) This is undefined (No Action Taken) [129]

Page 51, 4.9.8.1 Introduction

Is this the same as the working key in the case of working keys?

Editor's Note: With five paragraphs of text to consider, it is not possible to find 'this'.

HP 61) lost [s/b] loss (Accepted, Editorial) [130]

Page 53, 4.10 Data persistence model

change lost to loss

HP 62) Fencing (Accepted, Substantive) [131]

Page 53, 4.10 Data persistence model
see comment ENDL 2)

Somewhere around this clause, a discussion and mandatory features of fencing need to be added in. The discussion on the SNIA OSD TWG is at best inconclusive and at least one company believes that lack of fencing support in OSD could lead to undetected data corruptions. The OSD TWG however did not arrive at any conclusions on this opinion as of 3/23/04. Before the Letter Ballot process is closed, OSD TWG (and perhaps T10 at large) needs to drive this issue to closure.

Editor's Note: This comment will be resolved as described in the response to comment ENDL 2).

HP 63) All OSD commands assume bidirectional data transfers (Accepted, Editorial) [132]

Page 54, 4.11.1 OSD meta data

I presume this is meant to address the set/get attr operations, but this brings up another question. It would be good to clarify on bidirectional commands.

Editor's Note: The following new subclause will be added:

4.11.1 Bidirectional data transfers

All commands defined by this standard use both the Data-In Buffer and Data-Out Buffer.

HP 64) set attributes offset field missing from table 28 (Accepted, Editorial) [133]

Page 56, 4.11.3 OSD Data-Out Buffer format

and SET ATTRIBUTES OFFSET?

Editor's Note: A row will be added to table 28 for the SET ATTRIBUTES OFFSET field.

HP 65) Remove Sense Data Implementation Requirements from SPC-3 (Accepted, Substantive) [134]

Page 58, 4.13 Error reporting

This implies that the caption in Table 10 of SPC-3 r16 needs changing. S/optional for all others/OSD/

Editor's Note: As agreed by the May CAP working group, the "Implementation Requirements" column in SPC-3 r16 Table 10 has been removed in SPC-3 r19.

HP 66) Clarify forced reservation clearing (Accepted, Editorial) [135]

Page 59, 4.15 Reservations

Only one "I" can be associated with an "I_T" nexus. So, is "every initiator port associated with the released registrations" meant?

Editor's Note: The cited sentence will be modified as follows:

A unit attention condition (see SAM-3) shall be established for **every** the initiator port associated with **a every** registered I_T nexus.

HP 67) of [s/b] set to (Accepted, Editorial) [136]

Page 61, 5.1 OSD CDB format

Change "additional sense code of" to additional sense code set to"

HP 68) of [s/b] set to (Accepted, Editorial) [137]

Page 61, 5.1 OSD CDB format

Change "additional sense code of" to additional sense code set to"

HP 69) Eliminate Discussion of Encrypted CDBs (Accepted, Substantive) [138]

Page 61, 5.1 OSD CDB format

Why would you think they would be encrypted? Delete the sentence

HP 70) Clarify FUA bit on reads (Accepted, Substantive) [139]

Page 65, 5.2.3 Options byte

The READ command has an options byte, but FUA only means results of a command being written. Is the intent to force the read to be sourced from the stable storage, as it does in SBC-2?

Editor's Note: If the FUA bit is set to one, the intent is that the data be located on stable storage at the time status is returned, as per the first sentence in the definition of the FUA bit:

The FUA (force unit access) bit controls whether or not the results of a command shall be written to stable storage (see 4.10) before status is returned to the application client.

If the data is in volatile cache at the beginning of processing for a READ command with the FUA bit set to one, before status is returned the data is required to be written to stable storage. Where the data is transferred from is outside the scope of this standard.

The sentences describing operation with the FUA bit will be modified as follows to clarify this:

If the FUA bit is set to zero, the device server may return status as soon as the data transferred by this command ~~has been placed~~ is in the volatile cache. If the FUA bit is set to one, the device server shall not return status until the data transferred by this command (i.e., either read data or write data) has been written to stable storage.

HP 71) FrameMaker cross reference marker not on the correct paragraph (Accepted, Editorial) [140]

Page 71, 6.2 APPEND

Fix "in the The COLLECTIONS PER USER OBJECT field contains the value of the collections per user object attribute. (see), then"

Editor's Note: The problem is a technical gaff in the FrameMaker source file. The cross reference marker used to generate the text the User Object Quotas attributes page subclause cross reference for the maximum user object length attribute description is located on the wrong paragraph. The correct location for the cross reference marker is on the header. The cross reference marker will be moved and the cross references will be regenerated to correct this problem.

HP 72) Specify Current Command attributes page usage for created objects (Accepted, Substantive) [141]

Page 72, 6.3 CREATE

I think this should also be followed with a sentence like the following (the same comment applies to CREATE and WRITE).

The assigned User_Object_ID shall be placed in the created User_Object_ID attribute of the Current Command attributes page (see 7.1.2.24).

Editor's Note: The following additions will be made, each as a new paragraph in the cited subclause:

- In 6.3 (CREATE command): "The highest 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 7.1.2.24)."
- In 6.4 (CREATE AND WRITE command): "The assigned User_Object_ID shall be placed in the Collection_Object_ID or User_Object_ID attribute in the Current Command attributes page (see 7.1.2.24)."
- In 6.4 (CREATE COLLECTION command): "The assigned Collection_Object_ID shall be placed in the Collection_Object_ID or User_Object_ID attribute in the Current Command attributes page (see 7.1.2.24)."

- In 6.4 (CREATE PARTITION command): "The assigned Partition_ID shall be placed in the Partition_ID attribute in the Current Command attributes page (see 7.1.2.24). The Collection_Object_ID or User_Object_ID attribute in the Current Command attributes page shall be set to zero."

HP 73) FrameMaker cross reference marker not on the correct paragraph (Accepted, Editorial) [142]

Page 75, 6.4 CREATE AND WRITE

Fix "in the The COLLECTIONS PER USER OBJECT field contains the value of the collections per user object attribute. (see), then"

Editor's Note: The problem is a technical gaff in the FrameMaker source file. The cross reference marker used to generate the text the User Object Quotas attributes page subclause cross reference for the maximum user object length attribute description is located on the wrong paragraph. The correct location for the cross reference marker is on the header. The cross reference marker will be moved and the cross references will be regenerated to correct this problem.

HP 74) When to Test Quotas (Accepted, Substantive) [143]

Page 75, 6.4 CREATE AND WRITE

If CREATE AND WRITE encounters a quota error, did the CREATE work or not?

Quota errors can be checked at any time during the command according to 4.8.3.

The same question applies to any CREATE command with a set attributes operation that fails for any reason (quota, invalid attribute number, etc.).

It would be easiest for software if the CREATE operation completed before the WRITE operation or attribute modification begins, so the object is assumed created if a CHECK CONDITION for one of those reasons occurs.

Editor's Note: The following new paragraph will be added to 4.8.2 (Quota errors):

The device server shall not terminate a command for quota errors after any user data or attributes have been modified.

The following paragraph will be removed from 4.8.3 (Quota testing):

~~Tests for quota errors may be made at any time during the processing of a command. The processing of a command may be partially completed at the time a quota error is detected. The device server is not required to restore the state of the OSD logical unit to the state that was present before processing was begun for the command in which the quota error has been detected.~~

HP 75) object count quota [s/b] partition count (Accepted, Editorial) [144]

Page 78, 6.6 CREATE PARTITION

Change "object count quota" to "partition count". This paragraph was probably copied from CREATE, CREATE AND WRITE, and CREATE COLLECTION, but those deal with the "object count quota" field.

HP 76) Define FLUSH COLLECTION Command (Accepted, Substantive) [145]

Page 78, 6.7 FLUSH OBJECT

There should be a FLUSH COLLECTION command to flush a collection object to stable storage, perhaps flushing all the user objects that are part of that collection.

Editor's Note: The FLUSH COLLECTION will be defined as described in T10/04-189.

HP 77) Define FLUSH PARTITION Command (Accepted, Substantive) [146]

Page 78, 6.7 FLUSH OBJECT

Other commands like CREATE and REMOVE have separate versions for user data objects and partitions. Why is this one combined? What does this have OBJECT in its title? I suggest splitting it into FLUSH, FLUSH COLLECTION, and FLUSH PARTITION.

Editor's Note: The FLUSH, FLUSH COLLECTION, FLUSH PARTITION, and FLUSH OSD commands will be defined as described in T10/04-189.

HP 78) List [s/b] List of (Accepted, Editorial) [147]

Page 79, 6.7 FLUSH OBJECT, Table 46

Change List to List of

HP 79) listed [s/b] all of the listed (Accepted, Editorial) [148]

Page 79, 6.7 FLUSH OBJECT, Table 46

Change "listed" to "all of the listed"

HP 80) Define Interaction of FORMAT OSD and CDB Set Attributes Parameters (Accepted, Editorial) [149]

Page 81, 6.8 FORMAT OSD

It might be helpful to include a note here mentioning that, because a set attributes can be included with the FORMAT OSD CDB, happen the command can "complete" at the SCSI level with some attributes in the R range set.

Editor's Note: The cited text will be modified as follows:

- e) Root object attributes and partition **zero** attributes as defined by this standard;
- f) Vendor specific additional root object attributes and partition **zero** attributes;
- g) **Root object attributes and partition zero attributes updated as specified by the CDB parameters;**

HP 81) REQUEST SENSE returns parameter data with GOOD status (Accepted, Substantive) [150]

Page 81, 6.8 FORMAT OSD

Change "a sense key of NOT READY with the additional sense code set to LOGICAL UNIT NOT READY FORMAT IN PROGRESS," to "GOOD status with parameter data containing a sense key set to NOT READY and an additional sense code set to LOGICAL UNIT NOT READY FORMAT IN PROGRESS"

Editor's Note: Including the changed described in comment IBM 78), the cited text will be modified as follows:

In response to the REQUEST SENSE command, unless an error has occurred, the device server shall return **GOOD status with parameter data containing the a sense key of set to NOT READY with and** the additional sense code set to LOGICAL UNIT NOT READY FORMAT IN PROGRESS, ~~with the sense key specific bytes set for progress indication (as described in SPC-3).~~

HP 82) Add (MSB) and (LSB) to SEED field in table 63 (Accepted, Editorial) [151]

Page 99, 6.19 SET KEY

add (MSB) and (LSB) to SEED. The text below actually refers to the "least significant bit" of the field.

HP 83) Add (MSB) and (LSB) to SEED field in table 65 (Accepted, Editorial) [152]

Page 101, 6.20 SET MASTER KEY

Add (MSB) and (LSB) to SEED field. The text below actual refers to the "least significant bit".

HP 84) CREATE AND WRITE [s/b] WRITE & FrameMaker cross-reference bug (Accepted, Editorial) [153]

Page 103, 6.21 WRITE

Delete all this text that only belongs in the CREATE AND WRITE section. "If a CREATE AND WRITE command causes the value in the user object logical length attribute in the User Object Information attributes page (see 7.1.2.11) to exceed the value in the maximum user object length attribute in the The COLLECTIONS PER USER OBJECT field contains the value of the collections per user object attribute., then a quota error shall be generated (see 4.8.2). The quota testing principles described in 4.8.3 apply to the testing of the maximum user object length quota. If a CREATE AND WRITE command causes the value in the used capacity attribute in the Partition Information attributes page (see 7.1.2.9) to exceed the value in the capacity quota attribute in the Partition Quotas attributes page (see 7.1.2.13), then a quota error shall be generated (see 4.8.2). The quota testing principles described in 4.8.3 apply to the testing of the capacity quota."

Editor's Note: All of the cited text applies to the WRITE command (just as much as it applies to the APPEND command or the CREATE AND WRITE command). Clearly, the CREATE AND WRITE should be WRITE.

In addition the cited text suffers from a technical gaff in the FrameMaker source file. The cross reference marker used to generate the text the User Object Quotas attributes page subclause cross reference for the maximum user object length attribute description is located on the wrong paragraph. The correct location for the cross reference marker is on the header. The cross reference marker will be moved and the cross references will be regenerated to correct the problem.

HP 85) Indicate which attributes pages have a page format defined (Accepted, Editorial) [154]

Page 105, 7.1.2.1 Attributes pages overview, Table 67

See also: comment Seagate 36) and comment Veritas 108)

Table 67 should indicate which pages have defined page formats and which use the list format

Editor's Note: Since all pages are accessible via the list format, adding a column containing that information would be pointless. However, a column will be added indicating which pages have a defined page format.

HP 86) Insert "null-terminated" (Accepted, Editorial) [155]

Page 106, 7.1.2.2

Add "null-terminated" before "null-padded" or delete "null-terminated" from the 3.7.2

Editor's Note: 'null-terminated,' will be inserted.

HP 87) INCITS [s/b] the ASCII characters "INCITS" (Accepted, Editorial) [156]

Page 106, 7.1.2.2 Attribute number 0h in all attributes pages

Change INCITS to: the ASCII characters "INCITS"

HP 88) undefined [s/b] unidentified (Accepted, Editorial) [157]

Page 106, 7.1.2.3

Change unidentified to undefined

HP 89) The Collection Directory, User Object Directory and Null attributes pages do not have an attribute number 0h (Accepted, Editorial) [158]
Page 106, 7.1.2.2

"With the exception of the Root Directory and Partition Directory attributes pages, all attribute pages ... shall have attribute number 0h." The Collection Directory, User Object Directory, and Null pages should be included in this exclusion list. The Collection Directory attributes are all C+ page numbers; it doesn't mention if attribute 0h is supported too. The User Object Directory attributes are all 0+ page numbers, where attribute 0h contains just a string for the User Object Directory, not the vendor ID and page identification string indicated by table 68. The Null page isn't really a page, and it doesn't have an attribute 0h describing it compliant with table 68.

Editor's Note: The Collection Directory attributes page will be added to the list in the cited text (see comment Seagate 35). Table 72 clearly shows the User Object Directory attributes page containing a value having the format shown in table 68. So, the User Object Directory attributes page will not be added to the list. The Null attributes page is a place holder and as such is outside the scope of this discussion.

HP 90) Clarify the length and padding of directory attribute values (Accepted, Substantive) [159]
Page 108, 7.1.2.5
See also comment Other 6)

Nowhere is it defined that these Attribute Values are null-terminated. See 3.1.7.2

Editor's Note: The following statements will be added to the cited subclauses:

- 7.1.2.4 (Root Directory attributes page) "Attribute values in the Root Directory attributes page have the format described in 7.1.2.2."
- 7.1.2.5 (Partition Directory attributes page) "Attribute values in the Partition Directory attributes page have the format described in 7.1.2.2."
- 7.1.2.6 (Collection Directory attributes page) "Attribute values in the Collection Directory attributes page have the format described in 7.1.2.2."
- 7.1.2.7 (User Object Directory attributes page) "Attribute values in the User Object Directory attributes page have the format described in 7.1.2.2."

HP 91) Put field names in small caps (Accepted, Editorial) [160]
Page 111, 7.1.2.8 Root Information attributes page

use small caps for field names like code set (several times in this page)

HP 92) Improper use of 'may not' (Accepted, Editorial) [161]
Page 112, 7.1.2.8 and global

Change "states may not" to "states shall not". The table has a row for "may: yes or no". It's providing a choice of "may or may not" vs. "shall not". Apply to all uses of "states may not" in the document.

Editor's Note: This comment will be resolved as described in the response to comment IBM 97). With respect to the 'global' marking on this comment, a large number of IBM comments specifically cite and resolve all occurrences of the cited phrasing.

HP 93) Change date format (Accepted, Editorial) [162]
Page 112, 7.1.2.8 Root Information attributes page
see also comment HP 50) and comment HP 115)

Change January 1, 1970 to 1 January 1970

HP 94) Serial number [s/b] Product serial number (Accepted, Editorial) [163]

Separate comment that was not numbered by author.
Page 111, 7.1.2.8 Root Information attributes page

Change serial number to product serial number to match the SPC-3 name and the 3 preceding attribute names.

HP 95) Username is binary (Accepted, Substantive) [164]

Page 113, 7.1.2.9 Partition Information attributes page

Are there any requirements that the username attribute be an ASCII or UTF-8 string, or is treated as an arbitrary binary field? This question also applies to the Collection Information (7.1.2.10) and User Object Information (7.1.2.11) attributes pages.

Editor's Note: In 7.1.2.9 (Partition Information attributes page) the definition of the username attribute will be modified as follows:

The username attribute (number 9h) shall contain ~~an~~ a binary valued identification of the user of the partition specified by the application client.

In 7.1.2.10 (Collection Information attributes page) the definition of the username attribute will be modified as follows:

The username attribute (number 9h) shall contain ~~an~~ a binary valued identification of the user for the collection specified by the application client.

In 7.1.2.11 (User Object Information attributes page) the definition of the username attribute will be modified as follows:

The username attribute (number 9h) shall contain ~~an~~ a binary valued identification of the user for the user object specified by the application client.

HP 96) Questions about sparse user objects (No Action Taken) [165]

Page 115, 7.1.2.11

What is an "unwritten byte"? It'd be safest for software if bytes of 00h could be assumed to be written, rather than vendor-specific data (probably previous object contents). The tradeoff is slower performance.

Editor's Note: OSD user objects are allowed to be allocated sparsely. This is a critical OSD feature. While OBSDs are not required to write sparsely allocated bytes as zero or anything else, they are required to read sparsely allocated bytes as zeros (see 6.14 READ, the first paragraph below the second a,b,c list on page 94). This allows a variety of OSD implementations while maintaining fully interoperable results from the perspective of the application client.

HP 97) Put 'quota' in all quota attribute names (Rejected) [166]

Page 116, 7.1.2.12 Root Quotas attributes page

Given that "All attributes in the Root Quotas attributes page are quotas" why not include "quota" or "maximum" in all their names (or in none of their names)? "Partition capacity quota" and "Default maximum user object length" are clearly named, but the others are not. "Partition object count" and "Partition count" sound like current values not maximums. "Maximum partition object count" and "Maximum partition count" would be better. "Partition collections per user object" could pass as a limit; "Maximum partition collections per user object" would be better.

Reason for Rejection: If the attribute appears in a quotas attributes page, then the attribute is a quota. Since all references to attributes include both the attribute name and the attribute page name, there will never be any lack of clarity. There is no reason to make the attribute names longer (and harder to read) than they already are.

HP 98) 0001 [s/b] 1 (Accepted, Editorial) [167]

Page 116, 7.1.2.12 Root Quotas attributes page

Change 0001 to 1

HP 99) 0001 [s/b] 1 (Accepted, Editorial) [168]

Page 116, 7.1.2.12 Root Quotas attributes page

Change 0001 to 1

HP 100) 0002 [s/b] 2 (Accepted, Editorial) [169]

Page 116, 7.1.2.12 Root Quotas attributes page

Change 0002 to 2

HP 101) Root Quotas page format incorrect (Accepted, Editorial) [170]

Page 116, 7.1.2.12 Root Quotas attributes page

In table 78 on the next page, the partition count field is shown as 4 bytes, not 8 bytes.

Editor's Note: This comment will be resolved as described in comment HP 104).

HP 102) Attribute number 0h is not a quota (Accepted, Editorial) [171]

Page 116, 7.1.2.12 Root Quotas attributes page

"All attributes in the Root Quotas attributes page are quotas" is not true for number 0h.

Editor's Note: The cited sentence will be modified as follows:

A# Except for attribute number 0h, all attributes in the Root Quotas attributes page are quotas (see 4.8).

HP 103) Assign different attribute names in different attributes pages (Rejected) [172]

Page 116, 7.1.2.12 Root Quotas attributes page

There are two attributes called "default maximum user object length attribute", one in the Root Quotas page and the other in the Partition Quotas page. Since the partition values can be changed after created, they should use different names. Following the convention of other fields, call the root level attribute the "Partition default maximum user object length" attribute.

Reason for Rejection: Curious how this was not a problem with username attribute or the used capacity attribute both of which appear in the Root Information, Partition Information, Collection Information, and User Object Information attributes pages. Since all references to attributes include both the attribute name and the attribute page name, there will never be any lack of clarity. There is no reason to make the attribute names longer (and harder to read) than they already are.

HP 104) Root Quotas page format incorrect (Accepted, Editorial) [173]

Page 117, 7.1.2.12 Root Quotas attributes page

In table 77, partition count is an 8 bit value, implying this should be 43 and the page length should be 24h. (or, change table 77 to make the partition count 4 bytes)

Editor's Note: Actually, partition count is an 8 byte value (see comment HP 101). However, the proposed changes are correct and will be made.

HP 105) FrameMaker cross reference marker not on the correct paragraph (Accepted, Editorial) [174]

Page 118, 7.1.2.13 Partition Quotas attributes page

Change "in the" to "in the User Object Quotas attribute page."

Editor's Note: The problem is a technical gaff in the FrameMaker source file. The cross reference marker used to generate the text the User Object Quotas attributes page subclause cross reference for the maximum user object length attribute description is located on the wrong paragraph. The correct location for the cross reference marker is on the header. The cross reference marker will be moved and the cross references will be regenerated to correct this problem.

HP 106) FrameMaker cross reference marker not on the correct paragraph (Accepted, Editorial) [175]

Page 118, 7.1.2.13 Partition Quotas attributes page

Delete "The COLLECTIONS PER USER OBJECT field contains the value of the collections per user object attribute. (see) for each user object, when it is created." which seems like a cut-and-paste error.

Editor's Note: The problem is a technical gaff in the FrameMaker source file. The cross reference marker used to generate the text the User Object Quotas attributes page subclause cross reference for the maximum user object length attribute description is located on the wrong paragraph. The correct location for the cross reference marker is on the header. The cross reference marker will be moved and the cross references will be regenerated to correct this problem.

HP 107) Attribute number 0h is not a quota (Accepted, Editorial) [176]

Page 118, 7.1.2.13 Partition Quotas attributes page

"All attributes in the Partition Quotas attributes page are quotas" is not true for number 0h.

Editor's Note: The cited sentence will be modified as follows:

A# Except for attribute number 0h, all attributes in the Partition Quotas attributes page are quotas (see 4.8).

HP 108) Duplicate quota testing statement (Rejected) [177]

Separate comment that was not numbered by author.

Page 118?, 7.1.2.13 Partition Quotas attributes page

Add "If setting an attribute causes the capacity quota to be exceeded, a quota error shall be generated." See last paragraphs of 5.2.1.2 and 5.2.1.3

Editor's Note: The quota testing rules apply to all the attributes in a quota attributes page. That is the reason for the 'all the attributes in ...' statement cited in comment HP 107). The quota testing requirements are specified as part of the command definitions and that is why the cited 5.2.x.y subclauses contain the text that they do. It is not appropriate to state the quota testing requirements twice.

HP 109) Attribute number 0h is not a quota (Accepted, Editorial) [178]

Page 119, 7.1.2.14 User Object Quotas attributes page

"All attributes in the User Object Quotas attributes page are quotas" is not true for number 0h.

Editor's Note: The cited sentence will be modified as follows:

All Except for attribute number 0h, all attributes in the User Object Quotas attributes page are quotas (see 4.8).

HP 110) Incorporate new timestamp bypass controls (Accepted, Substantive) [179]

Page 121, 7.1.2.15 Root Timestamps attributes page

This section needs to incorporate the formula hashed out on the SNIA OSD mailing list - that a FFh value would leave the timestamp control to the CDB timestamp bypass value.

Editor's Note: This comment will be resolved as described in comment Seagate 26).

HP 111) Future timestamps (No Action Taken) [180]

Page 123, 7.1.2.16 Partition Timestamps attributes page

What if you get the data modified attribute with a CREATE or REMOVE CDB? Do you get the future time where that command is estimated to complete, or the time the last CREATE or REMOVE command ran, or the time the create or remove operation part of the command completed? This applies to all the fields that return a timestamp which can be affected by the command carrying the get attributes itself (either by the core command operation or by a set attributes that was also specified).

Editor's Note: The answers appear to be different for CREATE and REMOVE (as described in 4.7.2 – Command function ordering for commands that get and/or set attributes).

For CREATE, the state of the attributes after processing the command function and after processing implicit attributes updates from the command functions determines the attribute values returned. Therefore, the timestamps returned reflect the time at which the create command function completed.

For REMOVE, the state of the attributes before processing the command function determines the attribute values returned. Therefore, the timestamps returned reflected the values in place before the command was received (including if applicable, the timestamps from the last remove command function).

HP 112) Fix table 87 border (Accepted, Editorial) [181]

Page 124, 7.1.2.16 Partition Timestamps attributes page, Table 87

Fix border

HP 113) Setting a Collections attribute can cause a quota error (Accepted, Substantive) [182]

Separate comment that was not numbered by author.

Page 129?, 7.1.2.19 Collections attributes page

Add mention of the collections per user object quota check that is performed when setting this page

Editor's Note: The following new paragraph will be inserted:

If setting a collection pointer attribute causes the number of collection pointer attributes with non-zero attribute lengths to exceed the value in the collections per user object attribute in the Partition Quotas attributes page (see 7.1.2.13), then a quota error shall be generated (see 4.8.2). The quota testing principles described in 4.8.3 apply to the testing of the object count quota.

HP 114) Add pro forma attribute page format field definitions (Accepted, Editorial) [183]

Page 130, 7.1.2.19 Collections attributes page

Add:

The PAGE NUMBER field contains the attributes page number of the Collection attributes page. The PAGE LENGTH field contains the number of additional bytes in the page format of the Collection attributes page.

HP 115) Change date format (Accepted, Editorial) [184]Page 132, 7.1.2.20 Root Security attributes page
see also comment HP 93) and comment HP 115)

Change January 1, 1970 to 1 January 1970

HP 116) in the in the [s/b] in the (Accepted, Editorial) [185]

Page 146, 7.1.3.3

Delete extra "in the"

HP 117) that [s/b] the (Accepted, Editorial) [186]

Page 146, 7.1.3.2 List entry format for retrieving attributes for this OSD object

Change that to the

Editor's Note: Both instances of this mistake on page 146 will be corrected as well as an instance on page 147.**HP 118) all attributes [s/b] each attribute (Accepted, Editorial) [187]**

Page 146, 7.1.3.2 List entry format for retrieving attributes for this OSD object

Change all to each (or change having a non-zero attribute length" to "having non-zero attribute lengths")

Editor's Note: '~~all attributes~~' will be changed to '~~each attribute~~'**HP 119) Clarify list format attribute number definition (Accepted, Editorial) [188]**

Page 146, 7.1.3.2 List entry format for retrieving attributes for this OSD object "one attribute"

If attribute number is FFFFFFFFh, more than one attribute is going to be returned. Use of "one" is incorrect

Editor's Note: Including the change described in the resolution to comment HP 118), the cited text will be modified as follows:

The ATTRIBUTE NUMBER field specifies:

- a) ~~the~~ The attribute number within the attributes page specified by the ATTRIBUTES PAGE field of the one attribute value to be returned; or
- b) The value FFFF FFFFh to request the return of ~~all attributes~~ each attribute in the attributes page ~~specified by the ATTRIBUTES PAGE field~~ having a non-zero attribute length.

HP 120) Use OBSD consistently as target device (Accepted, Substantive) [189]

Global

In 4.2 (OSD type devices), the first paragraph will be modified as follows:

An OBSD (see 3.1.26) ~~is a~~ contains one or more logical units ~~unit~~ that returns the OSD peripheral device type value in response to an INQUIRY command (see SPC-3). From the perspective of the application client, an OBSD logical unit contains OSD objects (see 3.1.27), not logical blocks (see 4.5). All stored data objects (see 4.6) have attributes (see 4.7) associated with them.

In 4.5 (Description of the OSD), the first sentence will be modified as follows:

Data is stored in abstracted subsets by the OBSD (see 3.1.26) logical unit.

In 4.7.3 (Attributes pages), table 4 footnote a will be modified as follows:

... The attributes in these pages may be set repeatedly and the ~~OBSD~~ OSD logical unit shall maintain the most recently set values for retrieval using the CDB get attributes parameters. The ~~OBSD~~ OSD logical unit shall not modify attribute values in these pages except in response to information provided in the set attributes parameters in a CDB.

In 4.9.8.1 (Secret Keys ... Introduction), top of page 51, the following modifications will be made (additional changes noted in comments HP 51), Intel 21), and Seagate 21):

~~The OBSD shall save two~~ Two secret key values shall be saved for each master, drive, partition, and working secret key as follows:

In 4.9.8.1 (Secret Keys ... Introduction), the fourth paragraph on page 51 will be modified as follows:

When an OBSD is manufactured, both the master authentication key and master generation key values for each logical unit shall be provided, but the two values may be identical.

The 6.8 (FORMAT OSD) command definition will be modified as follows:

The FORMATTED CAPACITY field specifies the total capacity of the formatted OSD logical unit in bytes. If the FORMATTED CAPACITY field is set to zero, the entire logical unit ~~OBSD (see 3.1.26)~~ is formatted as one OSD logical unit and the logical unit capacity established accordingly. ...

The 7.1.2.4 (Root Directory attributes page) definition will be modified as follows:

The contents of the Root Directory attributes page shall be maintained by the OSD logical unit ~~OBSD (see 3.1.26)~~.

The 7.1.2.4 (Partition Directory attributes page) definition will be modified as follows:

The contents of the Partition Directory attributes page shall be maintained by the OSD logical unit ~~OBSD (see 3.1.26)~~.

The 7.1.2.4 (Collection Directory attributes page) definition will be modified as follows:

The contents of the Collection Directory attributes page shall be maintained by the OSD logical unit ~~OBSD (see 3.1.26)~~.

The 7.1.2.7 (User Object Directory attributes page) definition will be modified as follows:

The contents of the User Object Directory attributes page shall be maintained by the **OSD logical unit ~~OSD~~** ~~(see 3.1.26)~~.

HP 121) SHA is Secure Hash Algorithm (Accepted, Editorial) [190]
Page 9, HMAC-SHA1 acronym definition

Modify as follows:

HMAC-SHA1 Keyed-Hash Message Authentication Code - Secure Hash ~~Standard~~ **Algorithm** 1 (see 2.3)

HP 122) Both FIPS 180-1 and FIPS 198 are needed to define HMAC-SHA (Accepted, Editorial) [191]
Page 133, 7.1.2.20 (Root Security attributes page), table 96

Modify row 2 as follows:

01h HMAC-SHA1 FIPS 180-1 (1995) **and FIPS 198 (2002)**

7. IBM Corp.

George O. Penokie from IBM Corp. submitted the following comments on a No vote.

IBM 1) Remove revision history (Accepted, Editorial) [192]
PDF pg 3, pg iii, Revision Information

The revision information has to be removed before this goes to public review.

IBM 2) Input/output logical units are the same as logical units (Accepted, Editorial) [193]
PDF pg 17, pg xvii, Foreword, 2nd paragraph

This is no need to the << input/output >> terminology in front of logical units.

Editor's Note: The cited text will be modified as follows:

This SCSI command set provides multiple operating systems concurrent control over one or more ~~input/output~~ logical units.

IBM 3) the [s/b] this (Accepted, Editorial) [194]
PDF pg 17, pg xvii, Foreword, 2nd paragraph

The statement << scope of the SCSI command set. >> should be << scope of this SCSI command set. >>

Editor's Note: The cited text will be modified as follows:

This SCSI standard provides commands that assist with coordination between multiple operating systems. However, details of the coordination are beyond the scope of ~~the this~~ SCSI command set.

IBM 4) 200X [s/b] 2004 (Accepted, Editorial) [195]
PDF pg 17, pg xvii, Foreword, 6th paragraph

The statement << 200X. >> should be << 2004.>> as I assume this will get done this year.

Editor's Note: In any case, the cooperation part of the effort more or less ends with this Letter Ballot. The cited text will be modified as follows:

This standard was developed by T10 in cooperation with industry groups during 1999 through ~~200X~~ 2004.

IBM 5) Pilot implementations are not relevant (Accepted, Editorial) [196]
PDF pg 17, pg xvii, Foreword, 6th paragraph

The statement << Most of its features have been tested in pilot products implementing these concepts in conjunction with standard transport protocols. >> is marketing hype and should be removed.

IBM 6) Delete list of standards in SCSI family (Rejected) [197]
PDF pg 21, pg 2, 1 Scope

There is no point in listing all the SCSI standards. The entire list should be deleted.

Reason for Rejection: I have to maintain this list in conjunction with the SPC-3 version descriptors. As long I have to deal with the version descriptors, this list is appearing in my standards.

IBM 7) Remove SPC-2 because there are no references to it (Rejected) [198]

PDF pg 23, pg 4, 2.2 Approved ISO references
See also comment IBM 14)

The standards reference << ISO/IEC 14776-452, SCSI Primary Commands - 2 (SPC-2) [ANSI NCITS.351-2001] >> should be deleted as the only place SPC-2 is used is in the acronym list.

Reason for Rejection: SPC-2 is the only approved reference. If SPC-2 goes, then why not the whole subclause, with all its ANSI contact points?

IBM 8) Use commas, not parentheses (Accepted, Editorial) [199]

PDF pg 25, pg 6, 3.1.3 attributes:

This << Data (sometimes called meta data) that >> should be << Data, sometimes called meta data, that >>.

IBM 9) Zero or more is wrong in security method definition (Rejected) [200]

PDF pg 27, pg 8, 3.1.39 security method:

The statement << zero or more >> seems pointless and should be deleted.

Reason for Rejection: No. The NOSEC security method enables zero security features and thwarts (surprise) zero security threats.

IBM 10) SPC-3 is the document of reference for the sense data definition (Accepted, Editorial) [201]

PDF pg 27, pg 8, 3.1.41 sense data:

The statement <<server delivers to an application client as described in SPC-3. >> should be << server delivers to an application client. See SPC-3. >>

IBM 11) Sense data is point of reference for sense key, Not! (Rejected) [202]

PDF pg 27, pg 8, 3.1.42 sense key:

The statement << field in the sense data (see 3.1.41). >> should be << field in the sense data. See 3.1.41. >>.

Reason for Rejection: No. One looks to 3.1.41 to understand sense data, not sense key. If 3.1.41 defined sense key in detail, then the proposed change would be correct. But 3.1.41 defines sense data, making the current construct correct.

IBM 12) Reflect prohibition on linked commands in the glossary (Accepted, Editorial) [203]

PDF pg 27, pg 8, 3.1.46 task:

The statement << or a group of linked commands >> should be deleted as linked commands are not allowed for this command set.

IBM 13) 4.6.1 contains details about User_Object_IDs, Not! (Rejected) [204]

PDF pg 27, pg 8, 3.1.49 User_Object_ID:

The statement << user object (see 4.6.1). >> should be << user object. See 4.6.1. >>

Reason for Rejection: No. One looks to 4.6.1 to understand a user object, not a User_Object_ID. If 4.6.1 defined User_Object_ID in detail, then the proposed change would be correct. But 4.6.1 defines user object, making the current construct correct.

IBM 14) Remove SPC-2 acronym (Unresolved) [205]

PDF pg 28, pg 9, 3.2 Acronyms
See also comment IBM 7)

The SPC-2 acronym should be deleted as it is not used in this standard.

IBM 15) Clarify that no other bytes can contain a null (Accepted, Editorial) [206]

PDF pg 32, pg 13, 3.7.2 Data field termination and padding requirements, 1st paragraph

The statement << and all other bytes in the field shall not contain the ASCII null character. >> would read better as << and no other bytes in the field shall contain the ASCII null character. >>

IBM 16) Use an a,b,c list (Rejected) [207]

PDF pg 33, pg 14, 4.1 The request-response model, 2nd paragraph

This should be an ab,c list << CDB, Data-Out Buffer, Data-Out Buffer Size, Data-In Buffer, Data-In Buffer Size, and Sense Data. >>

Editor's Note: The list is not sufficiently complex to require the a,b,c format. The constituents of the list are already presented in a manner that makes them stand out in the procedure call above the cited text. Reformatting this information as an a,b,c list would serve only to increase the white space on the page, communicating no useful information, providing the reader no insight that is not already conveyed, generally wasting time and space.

IBM 17) Text is not syntactically clear but no changes proposed (No Action Taken) [208]

PDF pg 37, pg 18, 4.6.2 Identifying OSD objects, first sentence in 4.6.2.

'identifies the root object, each partition,...' is not clear syntactically. (IBM)

Editor's Note: Comment HP 13), comment Intel 9), comment Seagate 3), and comment Veritas 26) all requested the addition of a comma after the word "object", as showed in the text citation but not in the cited text. The requested comma will be added.

None of the other four comments requested any other changes or described a lack of clarity in the cited text. The editor finds no syntactic problems with the cited text. In the absence of any proposed additional changes, no changes will be made.

IBM 18) many [s/b] any number of (Accepted, Editorial) [209]

PDF pg 37, pg 18, 4.6.4 Partitions, 1st paragraph

The statement << be many partitions, >> should be << be any number of partitions, >>

IBM 19) Enhance list of commands that cannot be directed to a partition (Rejected) [210]

PDF pg 37, pg 18, 4.6.4 Partitions

See also comment IBM 21), comment Seagate 4), and comment Other 1)

Device shall terminate these commands applied to a partition object: append, read, remove, write. (IBM)

Reason for Rejection: The propose list is not complete and cannot be maintained completely over time (e.g., what about REMOVE COLLECTION, SET KEY, soon to be FLUSH COLLECTION, and ...?).

The intent of the cited paragraph is to state that a partition does not include data. With that in mind, comment Other 1) describes the only changes that will be made in the cited paragraph.

IBM 20) Three corrections in 4.6.6 (Partitions) (Accepted, Editorial) [211]

PDF pg 37, pg 18, 4.6.4 Partitions

- a) Second paragraph should be 'when the partition is created'.
- b) Last sentence, should be (see 6.6)
- c) Contrast partition zero with all other partitions (IBM)

Editor's Note: The requested changes will resolved by modifying the cited paragraphs as follows:

A Partition_ID uniquely identifies each partition. Partitions have a User_Object_ID of zero and a Partition_ID (see 4.6.2) that is assigned by the OSD logical unit when the ~~user-object~~ partition is created. **The partition with Partition_ID zero represents the root object and is called partition zero.**

When a partition is created using the CREATE PARTITION command (~~see 6.5~~) (see 6.6), a partition OSD object shall be created to provide navigation between user objects in the partition.

IBM 21) Enhance list of commands that cannot be directed to a collection (Rejected) [212]

PDF pg 38, pg 19, 4.6.6 Collections

See also comment IBM 19), comment Seagate 4), and comment Other 1)

device shall terminate these commands applied to a collection object: append, read, remove, write. (IBM)

Reason for Rejection: The propose list is not complete and cannot be maintained completely over time (e.g., what about REMOVE PARTITION, SET KEY, soon to be FLUSH PARTITION, and ...?).

The intent of the cited paragraph is to state that a partition does not include data. With that in mind, comment Other 1) describes the only changes that will be made in the cited paragraph.

IBM 22) Missing 'or' in a,b,c list (Accepted, Editorial) [213]

PDF pg 39, pg 20, 4.7.1 Overview, a.b.c list

Item d should be << d) A user object; or >>

IBM 23) Eliminate '(see table x in x.x.x)' construct (Rejected) [214]

PDF pg 39, pg 20, Global

The use of the format << (see table 3 in x.x.x), >> should not be used as there is no obvious rule as to when the x.x.x is there and when it is not. There is no way you can say that if the table or figure is on the same page then there is no x.x.x because there is no way to know when a table will move to a different page as the document is modified. So either remove the << in x.x.x >> and make it consistent with all the other standards (which is my choice) or make all figures and tables have the << in x.x.x >> format.

Reason for Rejection: There is a clear rule for when the '(see table x in x.x.x)' construct is used. It is used only when the table is in a subclause other than the subclause in which the reference appears. The usage does not depend on the relative pages of the table and the reference and so is not subject to the vagaries described in the comment. Regardless of page location, tables stay in the same subclauses.

The '(see table x in x.x.x)' construct is used to aid readers using printed copies of the standard that lack the hot links present in PDF files. This editor views the construct as a critical feature of the standards he edits that is fully consistent with the T10 and ISO style guides.

IBM 24) Written document is documentation (Accepted, Editorial) [215]

PDF pg 39, pg 20, 4.7.1 Overview, Last paragraph

I would like to see the statement << and by other written documentation. >> deleted. But if not deleted then at least delete the term << written >>.

Editor's Note: The list would not be as inclusive as the standard intends if the entire cited phrase were removed. However, 'written' will be removed.

IBM 25) performed [s/b] processed (Accepted, Editorial) [216]

PDF pg 39, pg 20, 4.7.2 Command function ordering for commands that get and/or set attributes, 2nd paragraph

The statement << shall be performed in the following order: >> should be << shall be processed in the following order: >>

Editor's Note: The cited text will be modified as follows:

~~The processing of commands~~ **Commands** other than GET ATTRIBUTES, SET ATTRIBUTES, REMOVE, REMOVE PARTITION, and REMOVE COLLECTION that include getting or setting attributes shall be ~~performed~~ **processed** in the following order:

IBM 26) 'and' is on the wrong 1,2,3 list entry (Accepted, Editorial) [217]

PDF pg 39, pg 20, 4.7.2 Command function ordering for commands that get and/or set attributes, 1st 1,2,3 list

The << and >> should be moved from item 2 to item 3.

IBM 27) Clarify REMOVE processing (Accepted, Substantive) [218]

PDF pg 39, pg 20, 4.7.2 Command function ordering for commands that get and/or set attributes

Remove may implicitly change some attributes (e.g. capacity). Therefore the order of these changes with respect to other operations should be explicitly state like in other commands in this section. (IBM)

Editor's Note: Including changes made in response to comment IBM 30) and to comment HP 23), the cited text will be modified as follows:

~~The processing of a~~ **A** REMOVE command, a REMOVE PARTITION command, ~~and or~~ a REMOVE COLLECTION command that ~~include includes~~ getting or setting attributes shall be ~~performed~~ **processed** in the following order:

- 1) Process any set attributes command functions specified in the CDB;
- 2) Process any get attributes command functions specified in the CDB; ~~and~~
- 3) ~~Processing Process~~ those command functions not related to attributes; ~~and~~
- 4) ~~Process any set attributes command functions resulting from the processing of the command (e.g., updating timestamps).~~

IBM 28) performed [s/b] processed (Accepted, Editorial) [219]

PDF pg 40, pg 21, 4.7.2 Command function ordering for commands that get and/or set attributes, 3rd paragraph

The statement << shall be performed in the following order: >> should be << shall be processed in the following order: >>

Editor's Note: The cited text will be modified as follows:

~~The processing of a~~ A GET ATTRIBUTES command shall be ~~performed processed~~ in the following order:

IBM 29) performed [s/b] processed (Accepted, Editorial) [220]

PDF pg 40, pg 21, 4.7.2 Command function ordering for commands that get and/or set attributes, 4th paragraph

The statement << shall be performed in the following order: >> should be << shall be processed in the following order: >>

Editor's Note: The cited text will be modified as follows:

~~The processing of a~~ A SET ATTRIBUTES command shall be ~~performed processed~~ in the following order:

IBM 30) performed [s/b] processed (Accepted, Editorial) [221]

PDF pg 40, pg 21, 4.7.2 Command function ordering for commands that get and/or set attributes, 5th paragraph

The statement << shall be performed in the following order: >> should be << shall be processed in the following order: >>:

Editor's Note: The cited text will be modified as follows:

~~The processing of a~~ A REMOVE command, a REMOVE PARTITION command, ~~and or~~ a REMOVE COLLECTION command that ~~include includes~~ getting or setting attributes shall be ~~performed processed~~ in the following order:

IBM 31) Publicly available vendor specific attribute page definitions enhance this standard (Accepted, Editorial) [222]

PDF pg 41, pg 22, 4.7.3 Attributes pages, Table 4

The distinction between 'Defined by OBSD manufacturer product specifications' and 'Vendor Specific' is not clear to us. (IBM)

Editor's Note: Any publicly available documentation of attributes pages enhances the value of this standard. Encouraging such publication of attributes pages definitions is the purpose of the cited attribute page number value range.

The cited text will be changed to:

Defined by ~~publicly available OBSD (see 3.1.26)~~ manufacturer ~~product~~ specifications

IBM 32) Attributes associated with any object type do not appear in any directory pages (Accepted, Substantive) [223]
PDF pg 42, pg 23, 4.7.5 Attributes directories

Does attribute page FFFFFFFE (Current Command Page) appear in the directory-page of specific-object-type ? it doesn't appear in the examples in section 7. (IBM)

Editor's Note: The following paragraph will be added at the end of the subclause:

Attributes pages that are not associated with any object type (i.e., attributes pages with page numbers between F000 0000h and FFFF FFFFh inclusive) do not appear in any attributes directory.

IBM 33) FrameMaker cross reference marker not on the correct paragraph (Accepted, Editorial) [224]
PDF pg 42, pg 23, 4.8.1 Introduction

Missing quota attribute page for user-object. (IBM)

Editor's Note: The problem is a technical gaff in the FrameMaker source file. The cross reference marker used to generate the text in entry c of the a, b, c list is located on the paragraph before the header for the User Object Quotas attributes page subclause. The correct location for the cross reference marker is on the header. The cross reference marker will be moved and the cross references will be regenerated to correct this problem.

IBM 34) FrameMaker cross reference marker not on the correct paragraph (Accepted, Editorial) [225]
PDF pg 42, pg 23, 4.8.1 Introduction, Item c

There is a missing reference in item c.

Editor's Note: The problem is a technical gaff in the FrameMaker source file. The cross reference marker used to generate the text in entry c of the a, b, c list is located on the paragraph before the header for the User Object Quotas attributes page subclause. The correct location for the cross reference marker is on the header. The cross reference marker will be moved and the cross references will be regenerated to correct this problem.

IBM 35) Eliminate gratuitous use of a,b,c list (Accepted, Editorial) [226]
PDF pg 43, pg 24, 4.8.2 Quota errors, a,b,c list

You should combine items a), b), and c) into an item a

Editor's Note: So that the rejection of comment IBM 16) does not look fickle, the cited subclause will be modified as follows:

4.8.2 Quota errors

If one of the quota error conditions described in 5.2.1 and clause 6 occurs, ~~processing of~~ the command shall be terminated with a CHECK CONDITION status, with the sense key set to DATA PROTECT and the additional sense code set to QUOTA ERROR. ~~a quota error shall be reported as follows:~~

~~a) The status shall be CHECK CONDITION;~~

~~b) The sense key shall be DATA PROTECT;~~

~~c) The additional sense code shall be QUOTA ERROR; and~~

~~d) The sense data shall include the OSD attribute identification sense data descriptor (see SPC-3) with one or more attribute descriptors identifying the quota attribute or attributes that have been exceeded.~~

Note that the a,b,c list has been completely eliminated.

IBM 36) Clarify quota errors sense data descriptor (Accepted, Substantive) [227]

PDF pg 43, pg 24, 4.8.2 Quota errors
See also comment ENDL 1)

What type of sense-data-descriptor should be used to specify violated quota attribute? Can this be explicitly explained in the document? (IBM)

Editor's Note: The changes described on 04-095 address this issue.

IBM 37) 'appropriate capability' not specific enough (Accepted, Editorial) [228]

PDF pg 43, pg 24, 4.8.4 Changing quotas, 1st paragraph

So in the statement << command with an appropriate capability. >> what exactly is the << appropriate capability. >> supposed to be? That kind of wording is not appropriate for a standard. Either be more specific or delete it.

Editor's Note: The cited text will be modified to read, "appropriate capability (see 4.9.4.3)".

IBM 38) Send Credential To Device (Rejected) [229]

PDF pg 44, pg 25, 4.9.1 Basic security model, Figure 4

should be 'Send credential to the device' (IBM)

Reason for Rejection: The design center for Credentials is that the device need not receive or cache Credentials because it can reconstruct them from knowledge that it has available; specifically the CDB, the Capability in the CDB, information it knows about the object being accessed, and the secret keys (knowledge of which it shares with the Security Manager).

IBM 39) a integrity [s/b] an integrity (Accepted, Editorial) [230]

PDF pg 44, pg 25, 4.9.1 Basic security model, 1st paragraph under figure 4

The statement << a integrity check >> should be << an integrity check >>.

IBM 40) New sentence, not semi-colon (Accepted, Editorial) [231]

PDF pg 44, pg 25, 4.9.1 Basic security model, 2nd paragraph under figure 4

The statement << by this standard; however, the >> should be << by this standard. However, the >>

IBM 41) performed [s/b] processed (Accepted, Editorial) [232]

PDF pg 44, pg 25, 4.9.1 Basic security model, item B

The statement << are to be performed. >> should be << are to be processed. >>

IBM 42) NOSEC definition change (Accepted, Substantive) [233]

PDF pg 45, pg 26, 4.9.3.1 Introduction, Table 6

NOSEC Description should be 'No security, access control only' (IBM)

Editor's Note: This comment will be resolved as described in the response to comment IBM 141).

IBM 43) Clarify CAPKEY trust discussion (Accepted, Editorial) [234]

PDF pg 47, pg 28, 4.9.3.3 The CAPKEY security method, Last sentence in 4.9.3.3

'from forging, modifying a credential or replaying it on a different secure channel' (IBM)

Editor's Note: The cited sentence will be modified as follows:

Even when communications are secured by such means, it is necessary to prevent the untrusted application client (see 4.9.2) from forging or ~~otherwise~~ modifying a credential, and from ~~replaying a credential on a different secure channel~~.

IBM 44) Make suggestion an example (Accepted, Editorial) [235]

The author marked this comment as technical.

PDF pg 48, pg 29, 4.9.3.4 The CMDRSP security method, 2nd to last paragraph

The statement << One possible action is to request a new credential from the security manager and retry the command. >> should not be here as the previous sentence states the action is not defined by this standard. If you reject this recommendation then at a minimum the statement needs to be stated as a example in an (e.g., ...).

Editor's Note: The cited sentence will be made an e.g. example.

IBM 45) Data includes meta data (Rejected) [236]

PDF pg 48, pg 29, 4.9.3.5 The ALLDATA security method, first sentence

Add 'validates the integrity of all data, mete data included, in transit'. (IBM)

Reason for Rejection: The requested change is unnecessarily tutorial. The details of ALLDATA operation follow immediately in the same subclause and the details clearly include so call meta data. In addition, it is impossible to maintain a specific list of data covered correctly as the standard evolves.

IBM 46) Make suggestion an example (Accepted, Editorial) [237]

The author marked this comment as technical.

PDF pg 50, pg 31, 4.9.3.5 The ALLDATA security method, 2nd to last paragraph

The statement << One possible action is to request a new credential from the security manager and retry the command. >> should not be here as the previous sentence states the action is not defined by this standard. If you reject this recommendation then at a minimum the statement needs to be stated as a example in an (e.g., ...).

Editor's Note: The cited sentence will be made an e.g. example.

IBM 47) root partition [s/b] partition zero (Accepted, Editorial) [238]

PDF pg 51, pg 32, 4.9.4.1 Credential format
See also comment HP 42)

Change 'root partition' to 'partition zero' (IBM)

Editor's Note: The following changes will be made:

a) The cited sentence will be modified as follows:

If the PARTITION_ID field contains zero, the credential applies to ~~root partition~~ partition zero (see 3.1.31).

b) In the description of the FORMAT OSD, the description of the OSD's contents following a successful format operation will be modified as follows:

b) One partition OSD object for partition zero (~~i.e., the root partition~~) (see 3.1.31);

IBM 48) The OBJECT CREATION TIME field is explained twice (Accepted, Editorial) [239]

PDF pg 53, pg 34, 4.9.4.3 Capability format, 2 paragraphs above table 13
See also comment Seagate 10)

OBJECT CREATION TIME is explained twice. (IBM)

Editor's Note: The first explanation will be removed. The specific text to be removed is shown in the response to comment Seagate 10).

IBM 49) Note 3 is normative (Accepted, Editorial) [240]

PDF pg 55, pg 36, 4.9.4.3 Capability format, NOTE 3

This looks more like normative text than a note << NOTE 3 A Partition_ID of zero specifies that access is allowed to both the partition numbered zero and the root object. >> and as such should be included as the main line text.

Editor's Note: The cited note will be replaced with the following normative text:

If the OBJECT TYPE field contains 01h (i.e., ROOT) a Partition_ID of zero specifies that access is allowed to the root object. If the OBJECT TYPE field contains 02h (i.e., PARTITION) a Partition_ID of zero specifies that access is allowed to partition zero (see 3.1.31).

IBM 50) Root attributes access includes partition zero attributes access (Accepted, Substantive) [241]

PDF pg 61, pg 42, 4.9.4.4 Credentials and commands allowed, Table 20

Applies to all ROOT entries: do attribute pages of partition. (IBM)

Editor's Note: The three table 20 rows with ROOT in the first column will be modified as follows:

- [ROOT | GET_ATTR] The retrieval of attributes from any attributes page associated with the root object **or in any attributes page associated with partition zero.**
- [ROOT | SET_ATTR] The setting of attributes in any attributes page other than Root Security attributes page associated with the root object **or in any attributes page other than Partition Security attributes page associated with partition zero.**
- [ROOT | SET_ATTR and SECURITY] The setting of attributes in any attributes page associated with the root object **or in any attributes page associated with partition zero.**

IBM 51) Remove 'in most cases' (Rejected) [242]

The author marked this comment as technical.

PDF pg 64, pg 45, 4.9.5.2 Credential and capability validation, 4th paragraph

The statement << In most cases the order in which the >> is just a tease to the reader as there is no indication as to cases in which the order is important. I recommend deleting << In most cases >>.

Reason for Rejection: There is an indication as to which case places importance on the order of processing. It appears after the word 'but' in the cited sentence, to wit:

In most cases the order in which the validation steps are performed is not critical, [but the validation of the request nonce should be performed as early in the process as possible](#) to increase the chances of detecting and rejecting duplicate nonces.

IBM 52) Eliminate Undetectable Alterations? (Accepted, Editorial) [243]

The author marked this comment as technical.

PDF pg 65, pg 46, 4.9.5.2 Credential and capability validation, last paragraph

The statement << shall not be altered in any detectable way. >> implies there is an undetectable way to alter the information. Why have that in a standard? I suggest deleting << in any detectable way >>.

IBM 53) security version key [s/b] security version tag (Accepted, Editorial) [244]

PDF pg 66, pg 47, 4.9.5.5 Invalidating credentials

See comment IBM 141)

Should be 'security version tag attribute (and not key)

IBM 54) Eliminate description of the threat addressed (Rejected) [245]

PDF pg 66, pg 47, 4.9.6.1 Request nonce format, 1st paragraph

The statement << to thwart attempts to capture OSD commands (e.g., FORMAT OSD) and replay them. >> is not needed in the standard and should be deleted.

Reason for Rejection: Those who read the security subclauses of OSD for the purpose of determining whether or not the security features are truly secure require descriptions of the threats to be addressed by various features in order to compare the feature to the threat.

IBM 55) Add cross reference to far-in-the-future nonce processing (Accepted, Editorial) [246]

PDF pg 67, pg 48, 4.9.6.2 Device server validation of request nonces, 6th paragraph

In the statement << Commands containing request nonces with timestamps that are greater than the contents of the clock attribute in the Root Information attributes page plus a delta value may be terminated with... >> the may implies a may not but there is no description of what happens if the choice is to not error. This needs to be fixed. It looks like that may be described below. If that is the case then just add a reference to were it is defined.

Editor's Note: The cited sentence will be modified as follows:

Commands containing request nonces with timestamps that are greater than the contents of the clock attribute in the Root Information attributes page plus a delta value [may be processed as described in 4.9.6.3](#) or they may 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 NONCE TIMESTAMP OUT OF RANGE.

IBM 56) meaning that [s/b] as a result (Accepted, Editorial) [247]

PDF pg 69, pg 50, 4.9.6.2 Device server validation of request nonces, 1st paragraph

The statement << bytes for integrity check values, meaning that the HMAC-SHA1 function output of 160 bits is truncated into 96 bits. >> should be << bytes for integrity check values, as a result the HMAC-SHA1 function output of 160 bits is truncated into 96 bits. >>

Editor's Note: This comment would have been accepted, except that comment Seagate 18) specifies that the entire cited paragraph be removed.

IBM 57) Remove 'meaning that' (Accepted, Editorial) [248]

PDF pg 69, pg 50, 4.9.8.1 Introduction, Table 22 footnote d

The statement << meaning that, even though the security manager computes it, the computation is based on values beyond the security manager's control >> should be << (i.e., even though the security manager computes it, the computation is based on values beyond the security manager's control) >>

Editor's Note: The cited sentence will be modified as follows:

The capability key is the credential integrity check value, ~~meaning that, even.~~ Even though the security manager computes it, the computation is based on values beyond the security manager's control (e.g., the user object to which the credential allows access).

IBM 58) Remove 'but' (Accepted, Editorial) [249]

PDF pg 70, pg 51, 4.9.8.1 Introduction, 3rd paragraph from end

The statement << provided, but the two values may be identical. >> should be << provided, and the two values may be identical. >>

Editor's Note: The cited sentence will be modified as follows:

When an OBSD is manufactured, both the master authentication key and master generation key values shall be provided, ~~but the.~~ The two values may be identical.

IBM 59) performed [s/b] processed (Accepted, Editorial) [250]

PDF pg 71, pg 52, 4.9.9 OSD security interactions with SPC-3 commands and

SAM-3 task management functions, 3rd paragraph in 2 places This << performed. >> should be << processed >>

IBM 60) Should explicit non-volatile cache support be provided? (No Action Taken) [251]

PDF pg 72, pg 53, 4.10 Data persistence model

Should the new FUA_NV bit be used in these commands?

Editor's Note: No. This standard explicitly includes non-volatile cache as part of stable storage, as described in the cited subclause, to wit:

Individual OBSD (see 3.1.26) implementations may use whatever technologies they choose to implement stable storage (e.g., [an OBSD may implement stable storage as a combination of non-volatile random access memory](#) and disk devices).

Considering the resistance to explicit non-volatile storage controls from RAID vendors during recent CAP meetings, one could easily argue that this standard has gotten it right.

IBM 61) and [s/b] and/or (Accepted, Editorial) [252]

PDF pg 73, pg 54, 4.11.1 OSD meta data, a,b,c list

This looks like it should be an << and/or >> list rather than an << and >> list.

IBM 62) What are offset fields? (Accepted, Editorial) [253]

PDF pg 76, pg 57, 4.11.4 Data-In and Data-Out buffer offsets, 4th and 5th paragraphs

The term << offset field >> is used but I see no field titled offset nor is offset is small caps. As a result I have no idea what the offset field is supposed to be. This needs to be fixed.

Editor's Note: The cited sentence will be modified as follows:

Offset fields (see table 29) in the CDB (e.g., the RETRIEVED ATTRIBUTES OFFSET field described in 4.11.2 and the SET ATTRIBUTES LIST OFFSET field described in 4.11.3) specify the starting byte of segments of the Data-In Buffer or Data-Out Buffer other than the command data or parameter data segment.

Note that the question of 'offset field' without small caps has not been addressed because no comments have been made regarding similar usage in 4.11.2 and 4.11.3, to wit:

The CDB [offset fields](#) that assist in locating the Data-In Buffer segments are shown in table 26.

IBM 63) What sense data descriptor? (Accepted, Editorial) [254]PDF pg 77, pg 58, 4.13 Error reporting
See also comment ENDL 1)

What type of sense data descriptor should be used for OSD object identification? (IBM)

Editor's Note: The changes described on 04-095 address this issue.

IBM 64) Make note 5 normative (Accepted, Editorial) [255]

PDF pg 77, pg 58, 4.13 Error reporting, NOTE 5

This note looks like normative text to me. I think it should be placed in main line text.

IBM 65) Make table additional CDB length match text (Accepted, Editorial) [256]

PDF pg 80, pg 61, 5.1 OSD CDB format, table 31

The << (n-7) >> should be << 166 >> as that is the requirement stated below the table.

IBM 66) Make table last CDB byte match text (Accepted, Editorial) [257]

PDF pg 80, pg 61, 5.1 OSD CDB format, Table 31

The byte number labeled << n >> should be << 173 >> as that is the requirement stated below the table.

IBM 67) semi-colon [s/b] and (Accepted, Editorial) [258]

PDF pg 82, pg 63, 5.2.1.2 Get an attributes page and set an attribute, 3rd paragraph under table 33

The statement << allocation length; this shall not be >> should be << allocation length and shall not be >>

Editor's Note: The cited text will be modified as follows:

If the get attributes allocation length is not sufficient to accommodate all bytes in the specified attributes page, the transfer of attributes data shall be truncated at the specified get attributes allocation length; **and** this shall not be considered to be an error.

IBM 68) Specify behavior if total capacity is exceeded (Rejected) [259]

PDF pg 87, pg 68, 6 Commands for OSD type devices

A general comment for section 6: all write/create commands violating total device capacity should also be aborted. (IBM)

Reason for Rejection: SCSI standard do not specify every error that a device server may report. Many errors are sufficiently clear that no specification of their reporting is necessary. Commands that attempt to exceed the capacity of the OSD device fall in this category.

IBM 69) Eliminate 'perform' (Accepted, Editorial) [260]

PDF pg 87, pg 68, 6.1 Summary of commands for OSD type devices, 1st paragraph

The term << perform. >> should be << process. >>

Editor's Note: The cited sentence will be modified as follows:

The commands **for OSD type devices** in table 39 ~~are command functions that an OSD logical unit may perform.~~

IBM 70) FrameMaker cross reference marker not on the correct paragraph (Accepted, Editorial) [261]

PDF pg 90, pg 71, 6.2 APPEND, 9th paragraph

The statement << in the maximum user object length attribute in the The COLLECTIONS PER USER OBJECT field contains the value of >> makes no sense and needs to be fixed.

Editor's Note: The problem is a technical gaff in the FrameMaker source file. The cross reference marker used to generate the text the User Object Quotas attributes page subclause cross reference for the maximum user object length attribute description is located on the wrong paragraph. The correct location for the cross reference marker is on the header. The cross reference marker will be moved and the cross references will be regenerated to correct this problem.

IBM 71) FrameMaker cross reference marker not on the correct paragraph (Accepted, Editorial) [262]

PDF pg 90, pg 71, 6.2 APPEND, 9th paragraph

There is a missing cross reference << (see), >>.

Editor's Note: The problem is a technical gaff in the FrameMaker source file. The cross reference marker used to generate the text the User Object Quotas attributes page subclause cross reference for the maximum user object length attribute description is located on the wrong paragraph. The correct location for the cross reference marker is on the header. The cross reference marker will be moved and the cross references will be regenerated to correct this problem.

IBM 72) Add discussion of quota violation for partition capacity (Accepted, Substantive) [263]

PDF pg 91, pg 72, 6.3 CREATE

See also comment IBM 75) and comment IBM 108)

Quota violation for partition capacity:

Creating a user object can also cause quota violation of the partition capacity. (IBM)

Editor's Note: The following paragraph will be added as the last paragraph in the cited subclause:

If a CREATE command causes the value in the used capacity attribute in the Partition Information attributes page (see 7.1.2.9) to exceed the value in the capacity quota attribute in the Partition Quotas attributes page (see 7.1.2.13), then a quota error shall be generated (see 4.8.2). The quota testing principles described in 4.8.3 apply to the testing of the capacity quota.

IBM 73) FrameMaker cross reference marker not on the correct paragraph (Accepted, Editorial) [264]

PDF pg 94, pg 75, 6.4 CREATE AND WRITE, 2nd to last paragraph

The statement << to exceed the value in the maximum user object length attribute in the The COLLECTIONS PER USER OBJECT field contains the value of >> makes no sense and needs to be fixed.

Editor's Note: The problem is a technical gaff in the FrameMaker source file. The cross reference marker used to generate the text the User Object Quotas attributes page subclause cross reference for the maximum user object length attribute description is located on the wrong paragraph. The correct location for the cross reference marker is on the header. The cross reference marker will be moved and the cross references will be regenerated to correct this problem.

IBM 74) FrameMaker cross reference marker not on the correct paragraph (Accepted, Editorial) [265]

PDF pg 94, pg 75, 6.4 CREATE AND WRITE, 2nd to last paragraph

There is a missing cross reference << attribute. (see), then >>.

Editor's Note: The problem is a technical gaff in the FrameMaker source file. The cross reference marker used to generate the text the User Object Quotas attributes page subclause cross reference for the maximum user object length attribute description is located on the wrong paragraph. The correct location for the cross reference marker is on the header. The cross reference marker will be moved and the cross references will be regenerated to correct this problem.

IBM 75) Add discussion of quota violation for partition capacity (Accepted, Substantive) [266]

PDF pg 95, pg 76, 6.5 CREATE COLLECTION

See also comment IBM 72) and comment IBM 108)

Quota violation for partition capacity:

Creating a collection object can also cause quota violation of the partition capacity. (IBM)

Editor's Note: The following paragraph will be added as the last paragraph in the cited subclause:

If a CREATE COLLECTION command causes the value in the used capacity attribute in the Partition Information attributes page (see 7.1.2.9) to exceed the value in the capacity quota attribute in the Partition Quotas attributes page (see 7.1.2.13), then a quota error shall be generated (see 4.8.2). The quota testing principles described in 4.8.3 apply to the testing of the capacity quota.

IBM 76) Add discussion of quota violation for root capacity (Rejected) [267]

PDF pg 96, pg 77, 6.6 CREATE PARTITION

Part of this comment has split to form comment IBM 152) because the comment describes half editorial and half substantive changes. The blue strikeout text is covered in comment IBM 152).

~~User object should be replace with partition object in a few instances and~~ creating a partition may cause a violation of the root capacity (IBM)

Reason for Rejection: The Root Quotas attributes page does not contain a capacity quota and the IBM comments contain no request to define one. So there is nothing to enforce against in the Root Quotas attributes page.

Attempting to enforce against the capacity quota attribute in the Partition Quotas attributes page for partition zero is meaningless because the comparison value would be the used capacity attribute in the Partition Information attributes page (i.e., the comparison would have to be with "the number of bytes used by partition zero and all other partitions, all collections, and all user objects within all partitions including attributes bytes"). The result would be a quota restriction that is equivalent to reducing the total capacity of the OBSD. Why not just reduce the total capacity and be done with it?

The intent appears to be that the partition count attribute in the Root Quotas attributes page is the only quota restriction on the creation of partitions.

IBM 77) What Does FLUSH COLLECTION do? (Accepted, Substantive) [268]

PDF pg 97, pg 78, 6.7 FLUSH OBJECT

What exactly should be performed when flush is done on collection-object, other than flushing attributes associated with it? (IBM)

Editor's Note: FLUSH COLLECTION (see comment HP 76) may flush the list of collection members, the collection attributes, or both. This comment will be resolved as described in T10/04-189.

IBM 78) No progress indication in OSD, maybe in OSD-2 (Accepted, Substantive) [269]

PDF pg 100, pg 81, 6.8 FORMAT OSD, 2nd to last a,b,c list item c

The statement << progress indication (as described in SPC-3). >> should be << progress indication as described in SPC-3. >> or << progress indication (see SPC-3). >>

Editor's Note: The SNIA OSD Technical Working Group agreed that returning progress indication should deferred from OSD (see 04-004r3). Therefore, the cited sentence will be modified as follows:

In response to the REQUEST SENSE command, unless an error has occurred, the device server shall return a sense key of NOT READY with the additional sense code set to LOGICAL UNIT NOT READY FORMAT IN PROGRESS, ~~with the sense key specific bytes set for progress indication (as described in SPC-3).~~

IBM 79) Provide single definition for allocation length (Rejected) [270]

PDF pg 106, pg 87, 6.11 LIST COLLECTION

The allocation length field definition is repeated in several commands. It should only be defined one time then referenced from that point on.

Reason for Rejection: The ALLOCATION LENGTH field appears in just two commands, LIST and LIST COLLECTION. Several fields appear in just two command definitions and yet do not appear in these comments as needing a common definition (e.g., the LIST IDENTIFIER field and the INITIAL OBJECT_ID field). The definition for the ALLOCATION LENGTH field will remain with the majority, defined in the two places where it is used.

IBM 80) Provide single definition for additional length (Rejected) [271]

PDF pg 106, pg 87, 6.11 LIST COLLECTION

The additional length h field definition is repeated in several commands. It should only be defined one time then referenced from that point on.

Reason for Rejection: The ADDITIONAL LENGTH field appears in just two commands, LIST and LIST COLLECTION. Several fields appear in just two command definitions and yet do not appear in these comments as needing a common definition (e.g., the LIST IDENTIFIER field and the INITIAL OBJECT_ID field). The definition for the ADDITIONAL LENGTH field will remain with the majority, defined in the two places where it is used.

IBM 81) Make common definition for LIST and LIST COLLECTION parameter data (Rejected) [272]

PDF pg 106, pg 87, 6.11 LIST COLLECTION

In fact the whole parameter data seems to be a duplicate of the list commands. A single reference to that would prevent duplicate definitions.

Reason for Rejection: There are subtle differences between the LIST and LIST COLLECTION parameter data. The former may list partition IDs and the latter may list collection IDs. Combining the two would create confusion. Also, changes being considered for OSD-2 are likely to increase the differences between LIST and LIST COLLECTION parameter data.

IBM 82) Change PERFORM SCSI COMMAND name (Rejected) [273]

PDF pg 108, pg 89, 6.12 PERFORM SCSI COMMAND

The name of the PERFORM SCSI COMMAND should be changed to PROCESS SCSI command.

Reason for Rejection: Although 'process' (to handle through an established set of procedures) is probably the correct verb when describing the actions of a device server, 'process' is not the correct word to use in formulating a request and SCSI commands are requests that an application client makes to a device server. 'PERFORM' (i.e., do, or more specifically do in a prescribed manner) is the correct word for the name of a SCSI command.

IBM 83) performed [s/b] processed (Accepted, Editorial) [274]

PDF pg 108, pg 89, 6.12 PERFORM SCSI COMMAND, 1st paragraph in 2 places

The term << performed >> should be << processed >>.

IBM 84) performed [s/b] processed (Accepted, Editorial) [275]

PDF pg 109, pg 90, 6.12 PERFORM SCSI COMMAND, 5th paragraph after table 54

The term << performed >> should be << processed >>.

IBM 85) Change PERFORM TASK MANAGEMENT FUNCTION name (Rejected) [276]

PDF pg 110, pg 91, 6.13 PERFORM TASK MANAGEMENT FUNCTION

The PERFORM TASK MANAGEMENT FUNCTION should be renamed to PROCESS TASK MANAGEMENT FUNCTION.

Reason for Rejection: Although 'process' (to handle through an established set of procedures) is probably the correct verb when describing the actions of a device server, 'process' is not the correct word to use in formulating a request and SCSI commands are requests that an application client makes to a device server. 'PERFORM' (i.e., do, or more specifically do in a prescribed manner) is the correct word for the name of a SCSI command.

IBM 86) performed [s/b] processed (Accepted, Editorial) [277]

PDF pg 110, pg 91, 6.13 PERFORM TASK MANAGEMENT FUNCTION, 1st paragraph in

2 places The term << performed >> should be << processed >>.

IBM 87) performed [s/b] processed (Accepted, Editorial) [278]

PDF pg 111, pg 92, 6.13 PERFORM TASK MANAGEMENT FUNCTION, 5th paragraph after table 56

The term << performed >> should be << processed >>.

IBM 88) taSk [s/b] Task (Accepted, Editorial) [279]

PDF pg 111, pg 92, 6.13 PERFORM TASK MANAGEMENT FUNCTION, Table 57

The term << SAM-3 taSk >> should be << SAM-3 Task >>. (IBM)

IBM 89) Use standard notation for CHECK CONDITION definition (Accepted, Editorial) [280]

PDF pg 113, pg 94, 6.14 READ, 1st abc list

Items b,c and d should be combined into an item b) to make the wording more consistent with other error words in the standard.

Editor's Note: The cited text will be modified as follows:

If the STARTING BYTE ADDRESS field specifies a byte that is beyond the user object logical length attribute value in the User Object Information attributes page (see 7.1.2.11), then:

- a) No bytes shall be transferred; **and**
- b) 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.**
- ~~c) The sense key shall be set to ILLEGAL REQUEST; and~~
- ~~d) The additional sense code shall be set to INVALID FIELD IN CDB.~~

IBM 90) Use standard notation for CHECK CONDITION definition (Accepted, Editorial) [281]

PDF pg 113, pg 94, 6.14 READ, 2nd abc list

Items b,c and d should be combined into an item b) to make the wording more consistent with other error words in the standard.

Editor's Note: The cited text will be modified as follows:

If the values in the LENGTH field and STARTING BYTE ADDRESS field result an attempt to read a byte that is beyond the user object logical length attribute value in the User Object Information attributes page, then:

- a) The bytes between the starting byte address and the user object logical length shall be transferred;
- b) The command shall be terminated with a CHECK CONDITION status, **with the sense key shall be set to RECOVERED ERROR and the additional sense code set to READ PAST END OF USER OBJECT;**
- ~~c) The sense key shall be set to RECOVERED ERROR;~~
- ~~d) The additional sense code shall be set to READ PAST END OF USER OBJECT;~~
- e) The command-specific information sense data descriptor (see SPC-3) shall be included in the sense data; and
- f) The COMMAND-SPECIFIC INFORMATION field shall contain the number of bytes transferred.

IBM 91) SET KEY Cannot Get A Quota Error (Accepted, Substantive) [282]

PDF pg 118, pg 99, 6.19 SET KEY

It should be impossible for a SET KEY command to violate quota. (IBM)

Editor's Note: In 7.1.2.20 (Root Security attributes page), the following modifications will be made (including changes from the responses to comments HP 51) and Intel 21):

The **drive root** key identifier attribute (number 7FFEh) contains the key identifier value from the most recent successful SET KEY command (see 6.19) with the KEY TO SET field set to 01b (i.e., update **drive root** key). If the **drive root** key is invalid (i.e., never set or invalidated by a SET MASTER KEY command), the **drive root** key identifier attribute length shall be zero. **Regardless of the root key identifier attribute length, the used capacity attribute in the Partition Information attributes page (see 7.1.2.9) for partition zero (see 3.1.31) shall reflect an attribute length of seven (i.e., it shall not be possible for a SET KEY command to cause the partition zero used capacity attribute value to exceed the capacity quota attribute in the Partition Quotas attributes page (see 7.1.2.13) for partition zero and generate a quote error).**

7.1.2.21 (Partition Security attributes page), the following modifications will be made:

The partition key identifier attribute (number 7FFFh) contains the key identifier value from the most recent successful SET KEY command (see 6.19) with the KEY TO SET field set to 10b (i.e., update partition key). If the partition key is invalid (i.e., never set, invalidated by a SET MASTER KEY command (see 6.20), or invalidated by a SET KEY command), the partition key identifier attribute length shall be zero. **Regardless of the partition key identifier attribute length, the used capacity attribute in the Partition Information attributes page (see 7.1.2.9) shall reflect an attribute length of seven (i.e., it shall not be possible for a SET KEY command to cause the partition's used capacity attribute value to exceed the capacity quota attribute in the Partition Quotas attributes page (see 7.1.2.13) and generate a quote error).**

The working key identifier attributes (numbers 8000h to 800Fh) contain the key identifier value from the most recent successful SET KEY command with:

- a) The KEY TO SET field set to 11b (i.e., update working key); and
- b) The KEY VERSION field set to the attribute number minus 8000h (e.g., a version key of three sets attribute 8003h and a version key of eight sets attribute 8008h).

If a working key is invalid (i.e., never set, invalidated by a SET MASTER KEY command, or invalidated by a SET KEY command), the working key identifier attribute length for the associated working key shall be zero. **Regardless of the lengths of any of the working key identifier attributes, the used capacity attribute in the Partition Information attributes page shall reflect an attribute length of seven for all sixteen working key identifier attributes (i.e., it shall not be possible for a SET KEY command to cause the partition's used capacity attribute value to exceed the capacity quota attribute in the Partition Quotas attributes page and generate a quote error).**

IBM 92) Define key identifier and seed just once (Rejected) [283]

PDF pg 120, pg 101, 6.20 SET MASTER KEY

The key identifier and seed fields are defined in another command and therefore there should only be a reference to those definition at this point.

Reason for Rejection: The key identifier has a different definition for SET KEY and SET MASTER KEY. While the seed definition is the same for both, forcing readers to switch from one page to the next just to save one line of text seems like putting the pain in the wrong place. Changes in one seed definition might or might not affect the other seed definition and there are numerous other blocks of repeated text that require coordinated changes (e.g., quota testing requirements) so nothing substantial is saved in that department either.

The SEED field appears in just two commands, SET KEY and SET MASTER KEY. Several fields appear in just two command definitions and yet do not appear in these comments as needing a common definition (e.g., the LIST IDENTIFIER field and the INITIAL OBJECT_ID field). The definitions for the SEED field will remain with the majority, defined in the two places where it is used.

IBM 93) CREATE AND WRITE [s/b] WRITE (Accepted, Editorial) [284]

The author marked this comment as technical.

PDF pg 122, pg 103, 6.21 WRITE, 2nd to last paragraph and last paragraph

What have the << CREATE AND WRITE command >> description have to do with the WRITE command description??

Editor's Note: The cited text contains a cut and paste from the CREATE AND WRITE description. CREATE AND WRITE should be changed to WRITE.

IBM 94) FrameMaker cross reference marker not on the correct paragraph (Accepted, Editorial) [285]

PDF pg 122, pg 103, 6.21 WRITE, 2nd to last paragraph

The statement << exceed the value in the maximum user object length attribute in the The COLLECTIONS PER USER OBJECT field contains the value of the collections per user object attribute., then a quota error shall be generated (see 4.8.2). >> makes no sense and needs to be fixed.

Editor's Note: The problem is a technical gaff in the FrameMaker source file. The cross reference marker used to generate the text the User Object Quotas attributes page subclause cross reference for the maximum user object length attribute description is located on the wrong paragraph. The correct location for the cross reference marker is on the header. The cross reference marker will be moved and the cross references will be regenerated to correct this problem.

IBM 95) them [s/b] attributes (Accepted, Editorial) [286]

PDF pg 123, pg 104, 7.1.1 Attributes parameter formats, 2nd paragraph

The term << them >> is not specific enough. What them is being referred to?

Editor's Note: The cited sentence will be modified as follows:

Page format parameter data allows retrieval of ~~them~~ attributes in formatted pages where only the attribute values appear in the parameter data.

IBM 96) preferred [s/b] recommended (Accepted, Editorial) [287]

PDF pg 125, pg 106, 7.1.2.3 Attribute number 0h for unidentified attributes pages, 1st paragraph

The statement << page as preferred by this standard. >> should be << page as defined by this standard. >>

Editor's Note: The word 'defined' seems in appropriate because the cited statement is referring to a 'should' in the second sentence in 7.1.2.2. The cited sentence will be modified as follows:

Certain attributes pages may be created dynamically making them subject to programming errors that fail to define an attribute number 0h for the attributes page as ~~preferred recommended~~ by this standard.

IBM 97) Improper keyword use in column heading (Accepted, Editorial) [288]

The author marked this comment as technical.

PDF pg 130, pg 111, 7.1.2.8 Root Information attributes page, Table 73

What is the point in the << May Be Set No >> column? If give no useful information as it could just as easily be read as << May or may not >> be set. Either change it to a should or a shall or remove it.

Editor's Note: The cited column heading will be changed to "~~Application Client Settable~~".

IBM 98) (number xh) [s/b] just (xh) (Rejected) [289]

PDF pg 131, pg 112, 7.1.2.8 Root Information attributes page, Global

What the heck is << (number xh) >> supposed to mean? I see no point in calling a hex value a number what else could it be? Why is it not just << (xh) >>? I recommend all the << (number xh) notations be changed to << (xh) >>.

Reason for Rejection: The word 'number' stands for 'attribute number'. In the usage context '(xh)' might be read as either the attribute number or incorrectly as the attribute value. The additional word 'number' is present to avoid this. The word 'attribute' is unnecessary because of the usage context, specifically "The page identification attribute (number 0h) ...".

The most formal phrasing would be '(i.e., attribute number xh)'. However, use of the completely formal format would cause less important information to clutter what are commonly complex definitions.

The editor can find no prohibition regarding '(number xh)' in either the T10 style guide or the ISO style guide. The '(number xh)' usage varies only slightly from the '(xh)' usage that frequents T10 standards.

IBM 99) Match comment IBM 97) changes in text (Accepted, Editorial) [290]

The author marked this comment as technical.

PDF pg 131, pg 112, 7.1.2.8 Root Information attributes page, last paragraph
see comment HP 92)

This statement << states may not be set, the >> looks like you are taking a may and making it into a shall. This needs to be fixed.

Editor's Note: The cited text refers to the column heading described in comment IBM 97). In keeping with the changes made in response to comment IBM 97), the cited text will be modified as follows: "states ~~may not be set~~ is not application client settable, the".

IBM 100) Improper keyword use in column heading (Accepted, Editorial) [291]

The author marked this comment as technical.

PDF pg 132, pg 113, 7.1.2.9 Partition Information attributes page, Table 74

What is the point in the << May Be Set No >> column? If give no useful information as it could just as easily be read as << May or may not >> be set. Either change it to a should or a shall or remove it.

Editor's Note: The cited column heading will be changed to "**Application Client Settable**".

IBM 101) Match comment IBM 100) changes in text (Accepted, Editorial) [292]

The author marked this comment as technical.

PDF pg 132, pg 113, 7.1.2.9 Partition Information attributes page, last paragraph

This statement << states may not be set, the >> looks like you are taking a may and making it into a shall. This needs to be fixed.

Editor's Note: The cited text refers to the column heading described in comment IBM 100). In keeping with the changes made in response to comment IBM 100), the cited text will be modified as follows: "states **may not be set is not application client settable**, the".

IBM 102) Improper keyword use in column heading (Accepted, Editorial) [293]

The author marked this comment as technical.

PDF pg 133, pg 114, 7.1.2.10 Collection Information attributes page, Table 75

What is the point in the << May Be Set No >> column? If give no useful information as it could just as easily be read as << May or may not >> be set. Either change it to a should or a shall or remove it.

Editor's Note: The cited column heading will be changed to "**Application Client Settable**".

IBM 103) Match comment IBM 102) changes in text (Accepted, Editorial) [294]

The author marked this comment as technical.

PDF pg 133, pg 114, 7.1.2.10 Collection Information attributes page, last paragraph

This statement << states may not be set, the >> looks like you are taking a may and making it into a shall. This needs to be fixed.

Editor's Note: The cited text refers to the column heading described in comment IBM 102). In keeping with the changes made in response to comment IBM 102), the cited text will be modified as follows: "states **may not be set is not application client settable**, the".

IBM 104) Improper keyword use in column heading (Accepted, Editorial) [295]

The author marked this comment as technical.

PDF pg 134, pg 115, 7.1.2.11 User Object Information attributes page, Table 76

What is the point in the << May Be Set No >> column? If give no useful information as it could just as easily be read as << May or may not >> be set. Either change it to a should or a shall or remove it.

Editor's Note: The cited column heading will be changed to "**Application Client Settable**".

IBM 105) Match comment IBM 104) changes in text (Accepted, Editorial) [296]

The author marked this comment as technical.

PDF pg 134, pg 115, 7.1.2.11 User Object Information attributes page, last paragraph

This statement << states may not be set, the >> looks like you are taking a may and making it into a shall. This needs to be fixed.

Editor's Note: The cited text refers to the column heading described in comment IBM 104). In keeping with the changes made in response to comment IBM 104), the cited text will be modified as follows: "states **may not be set is not application client settable**, the".

IBM 106) Improper keyword use in column heading (Accepted, Editorial) [297]

The author marked this comment as technical.

PDF pg 135, pg 116, 7.1.2.12 Root Quotas attributes page, Table 77

What is the point in the << May Be Set No >> column? If give no useful information as it could just as easily be read as << May or may not >> be set. Either change it to a should or a shall or remove it.

Editor's Note: The cited column heading will be changed to "**Application Client Settable**".

IBM 107) Match comment IBM 106) changes in text (Accepted, Editorial) [298]

The author marked this comment as technical.

PDF pg 136, pg 117, 7.1.2.12 Root Quotas attributes page, 2nd paragraph above table 78 paragraph

This statement << states may not be set, the >> looks like you are taking a may and making it into a shall. This needs to be fixed.

Editor's Note: The cited text refers to the column heading described in comment IBM 106). In keeping with the changes made in response to comment IBM 106), the cited text will be modified as follows: "states **may not be set is not application client settable**, the".

IBM 108) CREATE and CREATE COLLECTION cause partition capacity quota violations (Accepted, Substantive) [299]

PDF pg 137, pg 118, 7.1.2.13 Partition Quotas attributes page

See also comment IBM 72) and comment IBM 75)

Quota violation for partition capacity:

Reference to CREATE and CREATE COLLECTION commands for the capacity quota (IBM)

Editor's Note: The cited sentence will be modified as follows:

If the setting of an attribute value (see 5.2.1), an APPEND command (see 6.2), a **CREATE command (see 6.3)**, a CREATE AND WRITE command (see 6.4), a **CREATE COLLECTION command (see 6.5)**, or a WRITE command (see 6.21) attempts to exceed the capacity quota, a quota error (see 4.8.2) shall be generated.

IBM 109) Improper keyword use in column heading (Accepted, Editorial) [300]

The author marked this comment as technical.

PDF pg 137, pg 118, 7.1.2.13 Partition Quotas attributes page, Table 79

What is the point in the << May Be Set No >> column? If give no useful information as it could just as easily be read as << May or may not >> be set. Either change it to a should or a shall or remove it.

Editor's Note: The cited column heading will be changed to "**Application Client Settable**".

IBM 110) FrameMaker cross reference marker not on the correct paragraph (Accepted, Editorial) [301]

PDF pg 137, pg 118, 7.1.2.13 Partition Quotas attributes page, 2nd paragraph after table 79

The statement << ...user object length attribute in the The COLLECTIONS PER USER OBJECT field contains the value of the collections per user object attribute. >> does not make any sense. This needs to be fixed.

Editor's Note: The problem is a technical gaff in the FrameMaker source file. The cross reference marker used to generate the text the User Object Quotas attributes page subclause cross reference for the maximum user object length attribute description is located on the wrong paragraph. The correct location for the cross reference marker is on the header. The cross reference marker will be moved and the cross references will be regenerated to correct this problem.

IBM 111) FrameMaker cross reference marker not on the correct paragraph (Accepted, Editorial) [302]

PDF pg 137, pg 118, 7.1.2.13 Partition Quotas attributes page, 2nd paragraph after table 79

There is a missing cross reference << (see) >>

Editor's Note: The problem is a technical gaff in the FrameMaker source file. The cross reference marker used to generate the text the User Object Quotas attributes page subclause cross reference for the maximum user object length attribute description is located on the wrong paragraph. The correct location for the cross reference marker is on the header. The cross reference marker will be moved and the cross references will be regenerated to correct this problem.

IBM 112) Match comment IBM 109) changes in text (Accepted, Editorial) [303]

The author marked this comment as technical.

PDF pg 137, pg 118, 7.1.2.13 Partition Quotas attributes page, 2nd paragraph above table 80 paragraph

This statement << states may not be set, the >> looks like you are taking a may and making it into a shall. This needs to be fixed.

Editor's Note: The cited text refers to the column heading described in comment IBM 109). In keeping with the changes made in response to comment IBM 109), the cited text will be modified as follows: "states ~~may not be set~~ is not application client settable, the".

IBM 113) Improper keyword use in column heading (Accepted, Editorial) [304]

The author marked this comment as technical.

PDF pg 138, pg 119, 7.1.2.14 User Object Quotas attributes page, Table 81

What is the point in the << May Be Set No >> column? If give no useful information as it could just as easily be read as << May or may not >> be set. Either change it to a should or a shall or remove it.

Editor's Note: The cited column heading will be changed to "~~Application Client Settable~~".

IBM 114) the allow [s/b] the allowed (Accepted, Editorial) [305]

PDF pg 139, pg 120, 7.1.2.14 User Object Quotas attributes page, 2nd paragraph under table 81

The statement << value the allow in >> should be << value the allowed in >> (IBM)

Editor's Note: Actually, the word 'the' needs to be removed in addition to the change requested. This comment will be resolved as shown in the response to comment Seagate 47).

IBM 115) Match comment IBM 113) changes in text (Accepted, Editorial) [306]

The author marked this comment as technical.

PDF pg 139, pg 120, 7.1.2.14 User Object Quotas attributes page, 2nd paragraph above table 82 paragraph

This statement << states may not be set, the >> looks like you are taking a may and making it into a shall. This needs to be fixed.

Editor's Note: The cited text refers to the column heading described in comment IBM 113). In keeping with the changes made in response to comment IBM 113), the cited text will be modified as follows: "states ~~may not be set~~ is not application client settable, the".

IBM 116) Improper keyword use in column heading (Accepted, Editorial) [307]

The author marked this comment as technical.

PDF pg 140, pg 121, 7.1.2.15 Root Timestamps attributes page, Table 83

What is the point in the << May Be Set No >> column? If give no useful information as it could just as easily be read as << May or may not >> be set. Either change it to a should or a shall or remove it.

Editor's Note: The cited column heading will be changed to "~~Application Client Settable~~".

IBM 117) Match comment IBM 116) changes in text (Accepted, Editorial) [308]

The author marked this comment as technical.

PDF pg 140, pg 121, 7.1.2.15 Root Timestamps attributes page, 2nd paragraph above table 85 paragraph

This statement << states may not be set, the >> looks like you are taking a may and making it into a shall. This needs to be fixed.

Editor's Note: The cited text refers to the column heading described in comment IBM 116). In keeping with the changes made in response to comment IBM 116), the cited text will be modified as follows: "states ~~may not be set~~ is not application client settable, the".

IBM 118) Updates of "access" timestamps should involve the application client (Accepted, Substantive) [309]

PDF pg 142, pg 123, 7.1.2.16 Partition Timestamps attributes page

Sections 7.1.2.16 - 7.1.2.18, what if the operation on the data (e.g. READ) failed. Should the data accessed time attribute be modified? (IBM)

Editor's Note: Timestamps for "access" actions should involve the application client. This is a legitimate mirror of the timestamps for "modify" actions that involve the device server.

In 7.1.2.15 (Root Timestamps attributes page) the following changes will be made:

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 command whose CDB get and set attributes parameters (see 5.2.1) successfully ~~transferred retrieved~~ any attributes pages or values associated with the root object ~~to the application client~~.

In 7.1.2.16 (Partition Timestamps attributes page) the following changes will be made:

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 command whose CDB get and set attributes parameters (see 5.2.1) successfully ~~transferred retrieved~~ any attributes pages or values associated with the partition ~~to the application client~~.

For all partitions except partition zero, 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:

- a) LIST command (see 6.10) that ~~listed the transferred a list of~~ user objects in the partition ~~to the application client~~; or
- b) LIST COLLECTION command (see 6.11) that transferred a list of collections in the partition to the application client.

[move this sentence to the beginning of the paragraph and modify it]

For partition zero, the data accessed time attribute shall contain the value of the clock attribute in the Root Information attributes page at the completion of the most recent LIST command ~~listed the transferred a list of~~ partitions in the root object ~~to the application client~~.

In 7.1.2.17 (Collection Timestamps attributes page) the following changes will be made:

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 command whose CDB get and set attributes parameters (see 5.2.1) successfully ~~transferred retrieved~~ any attributes pages or values associated with the collection ~~to the application client~~.

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 6.11) that ~~transferred a list of user objects in the collection to the application client~~. ~~accessed the collection~~.

In 7.1.2.18 (User Object Timestamps attributes page) the following changes will be made:

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 command whose CDB get and set attributes parameters (see 5.2.1) successfully ~~transferred retrieved~~ any attributes pages or values associated with the user object ~~to the application client~~.

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 READ command (see 6.14) that ~~accessed transferred data from~~ the user object ~~to the application client~~.

IBM 119) Improper keyword use in column heading (Accepted, Editorial) [310]

The author marked this comment as technical.

PDF pg 142, pg 123, 7.1.2.16 Partition Timestamps attributes page, Table 86

What is the point in the << May Be Set No >> column? If give no useful information as it could just as easily be read as << May or may not >> be set. Either change it to a should or a shall or remove it.

Editor's Note: The cited column heading will be changed to "~~Application Client Settable~~".

IBM 120) LIST COLLECTION affects Partition Timestamps (Rejected) [311]

PDF pg 142, pg 123, 7.1.2.16 Partition Timestamps attributes page, 5th paragraph under table 86

LIST COOLECTION may also affects the data accessed time attribute. (IBM)

Reason for Rejection: Collections are specialized user objects, with their own timestamps attributes page. Listing a collection does not access partition information and thus should not affect any Partition Timestamps.

IBM 121) Match comment IBM 119) changes in text (Accepted, Editorial) [312]

The author marked this comment as technical.

PDF pg 143, pg 124, 7.1.2.16 Partition Timestamps attributes page, 2nd paragraph above table 87 paragraph

This statement << states may not be set, the >> looks like you are taking a may and making it into a shall. This needs to be fixed.

Editor's Note: The cited text refers to the column heading described in comment IBM 119). In keeping with the changes made in response to comment IBM 119), the cited text will be modified as follows: "states ~~may not be set is not application client settable~~, the".

IBM 122) Improper keyword use in column heading (Accepted, Editorial) [313]

The author marked this comment as technical.

PDF pg 144, pg 125, 7.1.2.17 Collection Timestamps attributes page, Table 88

What is the point in the << May Be Set No >> column? If give no useful information as it could just as easily be read as << May or may not >> be set. Either change it to a should or a shall or remove it.

Editor's Note: The cited column heading will be changed to "~~Application Client Settable~~".

IBM 123) Match comment IBM 122) changes in text (Accepted, Editorial) [314]

The author marked this comment as technical.

PDF pg 145, pg 126, 7.1.2.17 Collection Timestamps attributes page, 2nd paragraph above table 89

This statement << states may not be set, the >> looks like you are taking a may and making it into a shall. This needs to be fixed.

Editor's Note: The cited text refers to the column heading described in comment IBM 122). In keeping with the changes made in response to comment IBM 122), the cited text will be modified as follows: "states ~~may not be set is not application client settable~~, the".

IBM 124) Improper keyword use in column heading (Accepted, Editorial) [315]

The author marked this comment as technical.

PDF pg 146, pg 127, 7.1.2.18 User Object Timestamps attributes page, Table 90

What is the point in the << May Be Set No >> column? If give no useful information as it could just as easily be read as << May or may not >> be set. Either change it to a should or a shall or remove it.

Editor's Note: The cited column heading will be changed to "~~Application Client Settable~~".

IBM 125) Match comment IBM 124) changes in text (Accepted, Editorial) [316]

The author marked this comment as technical.

PDF pg 146, pg 127, 7.1.2.18 User Object Timestamps attributes page, 2nd paragraph above table 91

This statement << states may not be set, the >> looks like you are taking a may and making it into a shall. This needs to be fixed.

Editor's Note: The cited text refers to the column heading described in comment IBM 124). In keeping with the changes made in response to comment IBM 124), the cited text will be modified as follows: "states ~~may not be set is not application client settable~~, the".

IBM 126) Improper keyword use in column heading (Accepted, Editorial) [317]

The author marked this comment as technical.

PDF pg 147, pg 128, 7.1.2.19 Collections attributes page, Table 92

What is the point in the << May Be Set No >> column? If give no useful information as it could just as easily be read as << May or may not >> be set. Either change it to a should or a shall or remove it.

Editor's Note: The cited column heading will be changed to "~~Application Client Settable~~".

IBM 127) Match comment IBM 126) changes in text (Accepted, Editorial) [318]

The author marked this comment as technical.

PDF pg 148, pg 129, 7.1.2.19 Collections attributes page, 2nd paragraph above table 93

This statement << states may not be set, the >> looks like you are taking a may and making it into a shall. This needs to be fixed.

Editor's Note: The cited text refers to the column heading described in comment IBM 126). In keeping with the changes made in response to comment IBM 126), the cited text will be modified as follows: "states ~~may not be set is not application client settable~~, the".

IBM 128) Disallow sparse lists of supported integrity check value algorithms (Accepted, Substantive) [319]

PDF pg 150, pg 133, 7.1.2.20 Root Security attributes page, p 1 on pg

We suggest to clarify that the supported (non zero valued) algorithms should appear consecutively. (IBM)

Editor's Note: The following sentence will be added to the cited paragraph:

~~If a supported integrity check value algorithm attribute contains zero, then all supported integrity check value algorithm attributes with higher valued attribute numbers also shall contain zero.~~

IBM 129) Improper keyword use in column heading (Accepted, Editorial) [320]

The author marked this comment as technical.

PDF pg 150, pg 131, 7.1.2.20 Root Security attributes page, Table 94

What is the point in the << May Be Set No >> column? If give no useful information as it could just as easily be read as << May or may not >> be set. Either change it to a should or a shall or remove it.

Editor's Note: The cited column heading will be changed to "~~Application Client Settable~~".

IBM 130) Match comment IBM 129) changes in text (Accepted, Editorial) [321]

The author marked this comment as technical.

PDF pg 152, pg 133, 7.1.2.20 Root Security attributes page, 2nd paragraph above table 97

This statement << states may not be set, the >> looks like you are taking a may and making it into a shall. This needs to be fixed.

Editor's Note: The cited text refers to the column heading described in comment IBM 129). In keeping with the changes made in response to comment IBM 129), the cited text will be modified as follows: "states ~~may not be set~~ is not application client settable, the".

IBM 131) 16 or sixteen, not both (Accepted, Editorial) [322]

PDF pg 154, pg 135, 7.1.2.20 Root Security attributes page, Last paragraph

In the statement <<The 16 SUPPORTED INTEGRITY CHECK VALUE ALGORITHM fields >> the term << 16 >> is used but in other places the term << sixteen >> is used. Pick one and use it consistently.

Editor's Note: ~~16~~ will be change to ~~sixteen~~.

IBM 132) Improper keyword use in column headings (Accepted, Editorial) [323]

The author marked this comment as technical.

PDF pg 154, pg 135, 7.1.2.21 Partition Security attributes page, Table 98

What is the point in the << May Be Set No >> column? If give no useful information as it could just as easily be read as << May or may not >> be set. Either change it to a should or a shall or remove it.

Editor's Note: The cited column heading will be changed to "~~Application Client Settable~~". The same change will be made in the following tables for which specific comments have not been lodged: table 102, table 104, and table 106.

IBM 133) Newest valid nonce, not oldest (Accepted, Editorial) [324]

PDF pg 155, pg 136, 7.1.2.21 Partition Security attributes page, 3rd paragraph above table 100

Text describing newest valid nonce is a copy of the text describing oldest valid nonce. This should be edited. (IBM)

Editor's Note: The cited sentence will be modified as follows:

The newest valid nonce attribute (number 3h) shall contain the minimum number of milliseconds ~~prior to later than~~ the value in the clock attribute in the Root Information attributes page to which the device server constrains the contents of the TIMESTAMP field in a request nonce (see 4.9.6) received in a command addressed to the partition, a collection in the partition, or a user object in the partition.

IBM 134) Match comment IBM 132) changes in text (Accepted, Editorial) [325]

The author marked this comment as technical.

PDF pg 156, pg 137, 7.1.2.21 Partition Security attributes page, 2nd paragraph above table 101

This statement << states may not be set, the >> looks like you are taking a may and making it into a shall. This needs to be fixed.

Editor's Note: The cited text refers to the column heading described in comment IBM 132). In keeping with the changes made in response to comment IBM 132), the cited text will be modified as follows: "states ~~may not be set~~ is not application client settable, the".

IBM 135) 16 or sixteen, not both (No Action Taken) [326]

PDF pg 158, pg 139, 7.1.2.21 Partition Security attributes page, Last paragraph

In the statement <<The sixteen WORKING KEY IDENTIFIER fields contain the >> the term << sixteen >> is used but in other places the term << 16 >> is used. Pick one and use it consistently.

Editor's Note: As noted in the response to comment IBM 131), sixteen will be used. So, no changes are need in the cited text.

IBM 136) Match comment IBM 132) changes in text (Accepted, Editorial) [327]

The author marked this comment as technical.

PDF pg 159, pg 140, 7.1.2.22 Collection Security attributes page, 2nd paragraph above table 103

This statement << states may not be set, the >> looks like you are taking a may and making it into a shall. This needs to be fixed.

Editor's Note: The cited text refers to the table 102 column heading that will be changed as part of the response to comment IBM 132). In keeping with the table 102 changes made in response to comment IBM 132), the cited text will be modified as follows: "states ~~may not be set~~ is not application client settable, the".

IBM 137) Match comment IBM 132) changes in text (Accepted, Editorial) [328]

The author marked this comment as technical.

PDF pg 160, pg 141, 7.1.2.23 User Object Security attributes page, 2nd paragraph above table 105

This statement << states may not be set, the >> looks like you are taking a may and making it into a shall. This needs to be fixed.

Editor's Note: The cited text refers to the table 104 column heading that will be changed as part of the response to comment IBM 132). In keeping with the table 104 changes made in response to comment IBM 132), the cited text will be modified as follows: "states ~~may not be set~~ is not application client settable, the".

IBM 138) Match comment IBM 132) changes in text (Accepted, Editorial) [329]

The author marked this comment as technical.

PDF pg 161, pg 142, 7.1.2.24 Current Command attributes page, 2nd paragraph above table 107

This statement << states may not be set, the >> looks like you are taking a may and making it into a shall. This needs to be fixed.

Editor's Note: The cited text refers to the table 106 column heading that will be changed as part of the response to comment IBM 132). In keeping with the table 106 changes made in response to comment IBM 132), the cited text will be modified as follows: "states ~~may not be set~~ is not application client settable, the".

IBM 139) Eliminate field treatments (Accepted, Editorial) [330]

PDF pg 167, pg 148, 7.4 Mode parameters, 3rd paragraph

The statement << OSD type devices shall treat the following mode parameter header fields (see SPC-3) as reserved: >> should be << OSD type devices shall reserve the following mode parameter header fields (see SPC-3): >>.

IBM 140) is [s/b] are (Accepted, Editorial) [331]

PDF pg 171, pg 152, A.1 Attributes page numbers assigned by other standards, 1st paragraph

The statement << publication is shown in >> should be << publication are shown in >>

Editor's Note: The resolution for this comment has been incorporated in the resolution for comment LSI 14).

IBM 141) Capability Validation is NOT a Security Function (Unresolved) [332]

The author marked this comment as technical.

Global

Changing the semantics of the NOSEC security mode.

Currently NOSEC security method supports any command, ignoring the capability.

We would like to change the semantics as follows:

- a) Under the NOSEC security method, all commands should carry a capability, like in other security methods. The NOSEC method should enforce the same rules regarding Permission bits, security version tag, expiration time, audit etc. However, no integrity check computations will be done or enforced, and as a result no key maintenance is required.
- b) We suggest splitting section 4.9.1 into two sections: one on the 'Basic Access Control model' and the other on the 'Basic security model'. The first section will specify all the relevant items from current 4.9.1 that have to do with access control only and not integrity enforcement. For example, only item c) on page 25. The access control model should also be accompanied with a Figure similar to Figure 4 but without the capability key components.
- c) Section 4.9.2. should indicate the trust assumptions under NOSEC - all parties are supposed to 'behave' according to the protocol. Specifically, the client is trusted to copy the capability exactly as obtained from the security manager without any modification into the CDB. Moreover, if the capability was requested for partition p then the client must use it in a CDB with Partition_ID p (as there is no way to verify this). Even under these trust assumptions, other unintentional errors/mistakes that may occur at the client can be detected by NOSEC.
- d) Page 8, Definition 3.1.39 - remove 'zero or more'
- e) Revise 4.9.3.2
- f) Persistent reservations should have the same status in NOSEC as otherwise, that is it should be performed as a PERFORM SCSI COMMAND with a device management permission bit.
- g) Table 10 - indicate that the CREDENTIAL INTEGRITY CHECK VALUE is zero for the NOSEC method.
- h) Section 4.9.4.3 needs references to the NOSEC model.
- i) Suggest to change the name 'Security Version Tag' to 'Access Version Tag', as it is relevant now to the NOSEC as well; clarify that it is used to deny access.
- j) SET KET and SET MASTER KEY should not be serviced if the corresponding security model of the partition or the device is NOSEC.

Editor's Note: The resolution for this comment is not complete in this revision. Expect further changes in future revisions of this document (e.g., Capability preparation probably will become a function of the Policy/Storage Manager).

Comments IBM 42), Seagate 8), request essentially the same changes as this comment but in far less detail.

The responses to several other comments need to be reflected by the changes made in response to this comment:

- Agilent 7) – This comment requests clarification of behaviors resulting from the presence of a Partition_ID in both the Credential and the Capability. Agilent 7) can be resolved by removing the Partition_ID from the Credential as described below.
- EMC 1) – This comment requests that Trust Assumptions subclause describe the requirement for "privacy" (with "confidentiality" being the preferred term) covering the message exchanges between the security manager and the application client. Currently, that information is in the Elements of the example configuration subclause.
- IBM 147) and Seagate 9) – These comments describe the addition of a security method field to the Capability. Note that the device server must be required to compare the contents of the new field to the appropriate attribute in the Partition Security attributes page and terminate the command if the two values are not equal.
- Panasas 2) – This comment raises concerns about properly specifying the order in which the various Credential and Capability validation functions are performed. Whether the agreed order will match the one requested by comment Panasas 2) is to be determined.
- Seagate 18) – This comment changes the number of bytes in all integrity check values from 12 to 20.

Per discussions on the SNIA OSD TWG reflector, the following additional changes are required:

- Move the Partition_ID out of the Credential and in to the Capability Object Descriptor
- Create a Whole Partition Object Descriptor containing just the Partition ID to be used (at least) for CREATE when the device server assigns the User Object ID
- Add attributes list format fields describing the types of attributes in the get/set lists. This allows the lists to be parsed just once, at the time they are actually processed. The security validation of the lists will use the new fields and the list parsing will validate that nothing in the list conflicts with the new fields. Size of the new fields and the bit definitions in them are to be determined.

Details of how this comment is being resolved can be found in 04-193.

IBM 142) Security Hole in Data-Out Integrity Check (Accepted, Substantive) [333]

The author marked this comment as technical.

The model of the Data-out integrity information format (section 4.9.3.5) seems to open a potential security hole. Since it is the client that indicates the NUMBER OF bytes it includes in the integrity check computation, the client may actually compute the integrity check on less bytes than it actually transfers. To close this hole, we suggest adding in the validation process at the server (p. 30) that only these bytes (on which the integrity check is computed) will affect the command. Or maybe there is another way to verify that the check was done on the entire data, without relying on the application to specify this?

Need to reflect this on page 32 also.

Editor's Note: The cited text in 4.9.3.5 (The ALLDATA security method) will be modified as follows:

The NUMBER OF COMMAND OR PARAMETER BYTES field specifies the number of bytes from the command data or parameter data segment that are included in the data-out integrity check value. If the value in the CDB LENGTH field, if any, or the value in the CDB PARAMETER LIST LENGTH field, if any, is larger than the value in the NUMBER OF COMMAND OR PARAMETER BYTES field, 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 NUMBER OF SET ATTRIBUTES BYTES field specifies the number of bytes from the set attributes segment that are included in the data-out integrity check value. If the value in the CDB SET ATTRIBUTE LENGTH field, if any, or the value in the CDB SET ATTRIBUTES LIST LENGTH field, if any, is larger than the value in the NUMBER OF SET ATTRIBUTES BYTES field, 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 NUMBER OF GET ATTRIBUTES BYTES field specifies the number of bytes from the get attributes segment that are included in the data-out integrity check value. If the value in the CDB GET ATTRIBUTES LIST LENGTH field, if any, is larger than the value in the NUMBER OF GET ATTRIBUTES BYTES field, 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.

No text needing modification can be found on page 32.

IBM 143) Add Create Attributes Permission (Deferred to OSD-2) [334]

The author marked this comment as technical.

No page reference specified

We suggest adding another bit to the capability permission bits called 'Set_new_attr'. This bit, if set to one, allows creating a new attribute that has not been defined before on the OSD. This will avoid creation of such attributes 'by mistake' - note that the mechanism for creating new attributes is trivial - simply set_att with a page number and attribute name that hasn't been seen before.

Editor's Note: OSD-2 will do better.

IBM 144) Add Object Type to CDB (Rejected) [335]

The author marked this comment as technical.
No page reference specified

We suggest adding an Object Type Value (ROOT, PARTITION, COLLECTION, USER) as specified in Table 14 to the CDB in addition to the capability; alternatively, can use 4 out of the 64 bits of the Object_ID and Partition_ID values for that.

Reason for Rejection: With the addition of separate FLUSH commands for every object type (see comment HP 76) and comment HP 77), the command code (i.e., service action value) identifies the object type being processed.

IBM 145) Provide Reporting for Collections Support (Accepted, Substantive) [336]

The author marked this comment as technical.
pg 19, 4.6.6 (Collections), p 1, s 1

Support for Collections is optional. The device should have a mean to notify whether it supports this feature, either via INQUIRY or through an attribute in the Root page.

Editor's Note: The cited text will be modified as follows:

Support for collections is optional. **If collections are not supported:**
a) The length of attribute number 4h in the User Object Directory attributes page (see 7.1.2.7) shall be zero for every user object (i.e., no Collections attributes pages identified); and
b) Zero shall be returned as the length of attribute number 0h in every Collections attributes page (see 7.1.2.19).

In 7.1.2.7 (User Object Directory attributes page), the introduction to table 72 will be modified as follows:

Table 72 shows the attributes in the User Object Directory attributes page when only the attributes pages defined in this standard are accessible via the logical unit **and collections are supported**. **If collections are not supported, attribute number 4h shall have a length of zero.**

In 7.1.2.19 (Collections attributes page), the definition of attribute number 0h will be modified as follows:

If collections are supported, the The page identification attribute (number 0h) shall have the format described in 7.1.2.2 with the VENDOR IDENTIFICATION field containing the ASCII characters "INCITS" and the ATTRIBUTES PAGE IDENTIFICATION field containing the ASCII characters "T10 Collections". **If collections are no supported, the length of the page identification attribute shall be zero.**

In table 92 in 7.1.2.19 (Collections attributes page), the length of attribute number 0h will be changed from "**40**" to "**0 or 40**".

IBM 146) Provide Reporting for Volatile Storage Support (Accepted, Substantive) [337]

The author marked this comment as technical.
pg 53, 4.10 Data persistence model, 2nd p after a,b list

Same notification requirement applies to the model of volatile storage at the device. The device should be able to notify whether it supports it.

Editor's Note: The cited text will be modified as follows:

Implementation of a volatile cache is optional. **Support for volatile cache, including support for the FUA bit and the DPO bit, may be indicated by setting the v_SUP bit to one in the Extended INQUIRY Data VPD page (see SPC-3).**

IBM 147) Extended The Non-Standard Communications Between the Security Manager and Application Clients (Accepted, Substantive) [338]

The author marked this comment as technical.
No page reference specified

As an optimization, suggest that when the security manager grants a credential to the client, it also indicates the security method for this partition. Since the security manager is expected to be aware of the security method currently employed by every partition, it can save a new client from going thru the discovery process of this attribute, a process that requires a few commands.

Editor's Note: The requested change has been included in the response to comment IBM 141). Details of how this comment is being resolved can be found in 04-193.

IBM 148) Redefined Buffer Offset Values (No Action Taken) [339]

The author marked this comment as technical.
No page reference specified

The offset field format in Table 29 section 4.11.4 may waste buffer space for small data segments. One possible solution is to have an exception whereby if the exponent equals 0, then byte offset = mantissa * (2^{exponent}).

Editor's Note: This comment was withdrawn by its authors.

IBM 149) ABORT TASK and QUERY TASK for Non-Root Objects (Unresolved) [340]

The author marked this comment as technical.
No page reference specified

Is it meaningful to consider an ABORT TASK and QUERY TASK for a USER, COLLECTION or PARTITION object?

Editor's Note: They are useful if the desire is for a user application to be able to abort its own commands.

From SNIA OSD TWG: IBM 149 - resolved

- add a short paragraph to 6.13 describing that it is meaningful to consider the above suggestion. If the command is targeted to root or partition, it can operate on any object.

IBM 150) Add Capacity Quota for OBSD (Rejected) [341]

The author marked this comment as technical.
No page reference specified

We suggest to maintain the capacity of the entire device as follows:

- a. Capacity quota in the Partition Quotas page of partition zero should imply a quota on the capacity used by all partitions (entire device, excluding space used for root attribute pages).
- b. Used capacity in the Partition Information page of partition zero should not exceed the Capacity quota in the Partition Quotas page of partition zero.

Reason for Rejection: Maintaining a capacity quote for the entire logical unit adds overhead to the processing of almost every OSD command and the information has almost no practical usefulness.

IBM 151) Key Invalidation Changes (Rejected) [342]

The author marked this comment as technical.

No page reference specified

The c part of this comment was split into comment IBM 153) so that the rejected proposals were clearly separated from the accepted one.

Invalidation of keys:

a. We suggest that a valid key shall have a key identifier with a non-zero length. Therefore, a SET KEY command shall not use an empty key identifier. An invalid key shall have a zero length key identifier.

b. Upon changing the security method attribute of a partition (including partition zero and root), the working keys for the relevant partition should be invalidated by the OBSD, before changing the security method.

~~c. On page 132, 2nd paragraph before last, since a device always comes with an (initial) Master Key, the master key identifier attribute length shall never be zero. It is possible to define a standard name for this initial master key.~~

Reason for Rejection: a) Is already handled by the operation of the SET KEY command (see 6.9) because the SET KEY command always sets the attribute length to 7 and by text in 7.1.2.21, to wit:

If the partition key is invalid (i.e., never set, invalidated by a SET MASTER KEY command (see 6.20), or invalidated by a SET KEY command), the partition key identifier attribute length shall be zero.

and

If a working key is invalid (i.e., never set, invalidated by a SET MASTER KEY command, or invalidated by a SET KEY command), the working key identifier attribute length for the associated working key shall be zero.

b) Can be accomplished by 1) changing the security level and 2) issuing SET KEY commands.

IBM 152) user object [s/b] partition (Accepted, Editorial) [343]

PDF pg 96, pg 77, 6.6 CREATE PARTITION

This comment was split from comment IBM 76) because the original comment described half editorial and half substantive changes. The blue strikethrough text is covered in comment IBM 76).

User object should be replace with partition object in a few instances ~~and creating a partition may cause a violation of the root capacity~~ (IBM)

Editor's Note: The cited sentence will be modified as follows:

If the REQUESTED PARTITION_ID field contains any value other than zero and the device server is unable to assign the requested Partition_ID to the created ~~user-object~~ partition, the ~~user-object~~ partition shall not be created and the command shall 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.

IBM 153) define pre SET MASTER KEY master key identifier (Accepted, Substantive) [344]

PDF pg 151, pg 132, 7.1.2.20 Root Security attributes page

This comment was split from comment IBM 151) so that the accepted proposal could be easily identified.

c. On page 132, 2nd paragraph before last, since a device always comes with an (initial) Master Key, the master key identifier attribute length shall never be zero. It is possible to define a standard name for this initial master key.

Editor's Note: The cited text will be modified as follows:

If a SET MASTER KEY command has never been processed, the master key identifier attribute length shall be ~~zero seven~~ and the master key identifier attribute value shall be the ASCII characters "1st key".

8. Intel Corp.

Robert Sheffield from Intel Corp. submitted the following comments on a Yes vote.

Intel 1) SNIA, not SINA (Accepted, Editorial) [345]
PDF Page 18

Storage Industry Network Association s/b Storage Networking Industry Association (SNIA)

Intel 2) OSD not in figure 1 (Rejected) [346]
PDF Page 20

Figure 1 shows the relationship of this standard to the other standards

"This standard" is not shown in the figure.

Reason for Rejection: This standard is a device-type specific command set as indicated by entry a) in the a,b,c list on the same page above the cited text. Therefore, this standard is represented in figure 1.

Intel 3) OSD is in figure 1 (No Action Taken) [347]
PDF Page 20, Figure 1

"Device-Type Specific Command Sets", OSD sits here?

Editor's Note: Yes.

Intel 4) 'Inherited' Is The Wrong Word (Accepted, Editorial) [348]
PDF Page 36, 4.6.1 Stored data object types / b) Partition:

"(i.e. inherited)"

This is inconsistent with the normal usage of inheritance in OOA. Inheritance says a Partition is a "type of" Root object, and therefore has, at least as a subset, the same structure of attributes, though not necessarily the same values for each attribute. The default values of attributes are specified at object instantiation, and normally does not depend on the values of the object "type" it inherits from. What you're calling 'inheritance' here is more like functional decomposition, not inheritance as used in OOA.

Editor's Note: Three instances of the cited parenthetical expression will be removed.

Intel 5) contain [s/b] specify (Rejected) [349]
PDF Page 36, 4.6.1 Stored data object types / a) Root object

"contain" s/b specify

Reason for Rejection: The cited text more accurately represents the nature of attributes as variables as it is currently written, to wit:

Its attributes (see 4.7) [contain](#) global characteristics for the OSD logical unit (e.g., ...

Intel 6) Abstracted? (Accepted, Editorial) [350]

PDF Page 36, 4.5 Description of the OSD Architecture

"Abstracted" s/b "The term, 'abstracted',"

Editor's Note: This comment will be resolved as described in the response to comment Veritas 16).

Intel 7) between [s/b] among (Accepted, Editorial) [351]

PDF Page 37, 4.6.4 Partitions

"between" s/b among

Intel 8) Wordy partition capacity (Rejected) [352]

PDF Page 37, 4.6.4 Partitions

"There may be many partitions, up to the capacity of the OSD logical unit."

alt wording:

There may be as many partitions as can be accommodated within the capacity of the OSD logical unit.

Reason for Rejection: The wording changes in comment IBM 18) seem better.

Intel 9) Insert missing comma (Accepted, Editorial) [353]

PDF Page 37

root object s/b root object,

Intel 10) How does REMOVE COLLECTION work? (No Action Taken) [354]

PDF Page 38, 4.6.6 Collections

"deleted using the REMOVE COLLECTION command"

Does this also delete the user objects in the collection? If so, what happens to the other collections of which the user object was a member? Does removing all the individual user objects in a collection also delete the collection containing them?

Editor's Note: Removing a collection does not remove the user object that are members of the collection. Only the user objects' membership in the collection is removed and even that can be prevented by setting the FCR (force collection removal) bit to zero (see 6.16).

Removing all the user objects from a collection does not automatically cause the collection to be removed. That is how a collection can come to have zero user objects as members (see the response to comment Intel 11).

Intel 11) Questions about collections (No Action Taken) [355]

PDF Page 38, 4.6.6 Collections

"collections each of which may contain zero or more user objects."

Is it really a collection if it has less than two elements? Can collections collect collections?

Editor's Note: A collection might once have contained hundreds of user objects, but one by one they were removed until finally there are none. The collection does not disappear until it is explicitly removed, and automatic removal of a collection that contains too few user objects would produce undesirable confusion for initiators. Therefore, the cited statement is correct.

Collections cannot contain collections. Addition to a collection is accomplished by setting an attribute in the Collections attributes page. Only user objects have Collections attributes pages. Therefore, only user object can be added to a collection.

Intel 12) Get attributes suggestions (Rejected) [356]

PDF Page 39, 4.7.2 Command function ordering for commands that get and/or set attributes

- "3) Process any set attributes command functions specified in the CDB;
- "4) Process any get attributes command functions specified in the CDB."

Why in this order? Generally the application client already knows the value of set attributes, so wouldn't it be more useful to return the value of the get attributes as they were prior to processing the set attributes? A useful extension would be to check the get attributes against some criteria before processing either the write command or set attributes functions.

Reason for Rejection: The case motivating the currently specified behavior is truncating a user object after writing it. In such a case, it is generally preferable to return the length of the user object after the command rather than before. The GET ATTRIBUTES is available for the cases where the reverse is needed.

Intel 13) store [s/b] stored (Accepted, Editorial) [357]

PDF Page 39, 4.7.1 Overview

"store" s/b stored

Intel 14) Why GET ATTRIBUTES and SET ATTRIBUTES? (No Action Taken) [358]

PDF Page 40, 4.7.2 Command function ordering for commands that get and/or set attributes

- "3) Process any set attributes command functions specified in the CDB."

This doesn't sound like an expected function of a GET ATTRIBUTES command. If both GET ATTRIBUTES and SET ATTRIBUTES are capable of setting and getting attributes, why have two commands? Why not simply have an ATTRIBUTES command that does both get and set functions?

Editor's Note: It is considered desirable to have one command that gets attribute values before setting them (i.e., GET ATTRIBUTES) and a different command that sets attributes before getting them (i.e., SET ATTRIBUTES). See the response to comment Intel 12).

Intel 15) set [s/b] collection (Rejected) [359]

PDF Page 43, 4.9.1 Basic security model

"...set of OSD objects." s/b "...collection of OSD objects."?

Reason for Rejection: Sorry. 'Collection' is a glossary defined term and the proposed usage does not match the glossary definition.

Intel 16) Example sounds like a requirement (Rejected) [360]

PDF Page 45, 4.9.3.1 Introduction

"...the preparations for the ALLDATA security method include the preparations that are necessary for the CMDRSP security method)."

This sounds like a "shall" requirement, and if so, should not be written as a parenthetical "e.g.".

Reason for Rejection: The cited text is a statement of fact. The design of the ALLDATA security method just happens to include everything in the CMDRSP security method as a functional subset.

Intel 17) Is 'propriety' the right word? (No Action Taken) [361]

PDF Page 45, 4.9.1 Basic security model

"propriety of the application client's actions." Is the common English definition of "propriety" sufficient for this context?

Editor's Note: 'Propriety' passed muster with the security guru reviewer. In the absence of a better proposal, no changes will be made.

Intel 18) 'considered' is not very precise (Rejected) [362]

PDF Page 48, 4.9.3.4 The CMDRSP security method (second to last paragraph)

"the presence of malicious entities perpetrating a denial of service attack should be considered."

The term, "considered", isn't very precise. Perhaps, "the application client should generate a user alert that a possible denial of service attack has been detected."

Reason for Rejection: The imprecise term 'considered' is used because what happens is outside the scope of the standard. For example, the propose change to generating a user alert calls upon a feature (i.e., user alerts) that is not defined in this standard or any other SCSI standard.

It is always difficult for SCSI standards to require specific behavior on the part of application clients. The preferred approach is to specify the device server's response to application client misbehavior. However, the device server cannot be part of the solution in the cited instance.

Intel 19) reduce [s/b] limit (Rejected) [363]

PDF Page 68, 4.9.6.3.1 Introduction

"reduce" Perhaps, "limit" is a better word.

Reason for Rejection: In the cited sentence, reduce (to diminish) seems better than limit (to restrict), to whit:

Device servers may process commands containing request nonces that exceed the sum of the clock attribute in the Root Information attributes page (see 7.1.2.8) plus newest valid nonce attribute in the Partition Security attributes page (see 7.1.2.21) and use one or both of the following methods to [reduce](#) the amount of resources required to remember every request nonce ever received: ...

Intel 20) reduce [s/b] limit (Rejected) [364]

PDF Page 68, 4.9.6.3.2 Capability restrictions with far in the future nonces

"reduce" Perhaps, "limit" is a better word.

Reason for Rejection: See response to comment Intel 19).

Intel 21) 'drive key' is un-SCSI (Accepted, Editorial) [365]

PDF Page 69, Table 22 - OSD secret key hierarchy

See comment HP 51)

"Drive" Is this the right term in SCSI context? Does this mean a "disk drive"? Perhaps "Target", or "Device", or "Device Server" is better suited?

Editor's Note: 'drive key' and 'device key' will be changed globally to 'root key'. Details of how this comment is being resolved can be found in 04-193.

9. Lingua Data

Joe Breher from Lingua Data submitted the following comments as part of the LSI Logic Corp. (see page 120) ballot.

Lingua 1) Make ISO part number correct (Accepted, Editorial) [366]

Title page (i), lower right corner

"ISO/IEC 14776-131: 200x" conflicts with ISO/IEC # on pg 3, which lists doc as "ISO/IEC 14776-381"

Which is proper?

Editor's Note: Actually, none of the cited ISO part numbers are correct. Per 03-146r4, the ISO part number for OSD is 391. The necessary changes will be made.

Lingua 2) NCITS [s/b] INCITS (Accepted, Editorial) [367]

pg vii

"Secretariat National Committe for Information Technology Standards" s/b "Secretariat InterNational Committe for Information Technology Standards"?

Lingua 3) NCITS [s/b] INCITS (Accepted, Editorial) [368]

pg xvii, penultimate paragraph

"Secretariat, National Committe for Information Technology Standards" s/b "Secretariat, InterNational Committe for Information Technology Standards"?

Lingua 4) SNIA, not SINA (Accepted, Editorial) [369]

pg xviii, last sentence

"Storage Industry Network Association" s/b "Storage Networking Industry Association"

Lingua 5) an SCSI [s/b] a SCSI (Accepted, Editorial) [370]

pg 1, section 1, list item a

"...data over an SCSI..." s/b "...data over a SCSI..." to match all other appearances of 'a SCSI'

Lingua 6) Devices [s/b] Device (Accepted, Editorial) [371]

pg 3, 2/3 down page

"Object-based Storage Devices Commands" s/b "Object-Based Storage Device Commands" to match title

Lingua 7) Should SAM-2 be referenced? (Rejected) [372]

pg 4, sec 2.2

should SAM-2 be listed as an approved reference?

Editor's Note: No. SAM-2 is the last parallel bus SCSI Architecture Model. SAM-2 also makes Autosense support optional. Since no parallel bus products support the bidirectional commands required by OSD and since OSD requires Autosense, SAM-3 is the only viable SCSI Architecture Model reference for OSD.

Lingua 8) lost [s/b] loss (Accepted, Editorial) [373]

pg 8, sec 3.1.44 stable storage

"...may result in the lost of data..." s/b "...may result in the loss of data..."

Lingua 9) Add User_Object_ID to user object definition (Rejected) [374]

pg 8, sec 3.1.48 user object

"...referenced by byte offset..." s/b "...referenced by User_Object_ID (see 3.1.49) identifying the OSD object, and the byte offset..."

Reason for Rejection: The current text says that the data is referenced to byte offset, to wit:

An OSD object (see 3.1.27) that contains **user data** (see 4.6.1) **that is referenced by byte offset** within the OSD object.

If one desires to describe the addressing for a user object as well as the data in that user object, then both the User_Object_ID and Partition_ID must be included. Doing that will make the definition so complex as to obscure the main point.

Lingua 10) Typo in Volatile Cache definition (Accepted, Editorial) [375]

pg 8, sec 3.1.51 volatile cache

See also comment Lingua 10)

"Storage is lost..." s/b "Storage that is lost..."

Lingua 11) Clarify character encoding description (Accepted, Editorial) [376]

pg 11, sec 3.6.1, 1st paragraph thereof

"When this standard requires one or more bytes to contain specific encoded **character**, the specific characters are enclosed in double quotation marks. The double quotation marks identify the start and end of the characters that are required to be encoded but are not themselves to be encoded. The characters that are to be encoded are shown in exactly the case that is to be encoded."

s/b

"When this standard requires one or more bytes to contain specific encoded **characters**, the specific characters are enclosed in double quotation marks. The double quotation marks identify the start and end of the characters that are required to be encoded but **these quotation marks** are not themselves to be encoded. The characters that are to be encoded are shown in exactly the case that is to be encoded."

plural in first sentence, clarify with 'these taction marks', add 'be' to 2nd sentence

Editor's Note: The cited text will be modified as follows:

When this standard requires one or more bytes to contain specific encoded ~~character~~ **characters**, the specific characters are enclosed in double quotation marks. The double quotation marks identify the start and end of the characters that are required to be encoded but ~~the quotation marks~~ are not ~~themselves~~ to **be** encoded. The characters that are to be encoded are shown in exactly the case that is to be encoded.

The differences from the proposed text are use of 'the' quotation marks, not 'these', and deletion of the now redundant 'themselves'.

Lingua 12) Add cross reference to user objects (Accepted, Editorial) [377]

pg 15, third para up from bottom

"The storage management component is focused on mapping logical constructs (e.g., files or database entries) to the physical organization of the storage media. In the OSD model, the logical constructs are called user objects. The root object (see 4.6.3), partitions (see 4.6.4), and collections (see 4.6.6) provide additional navigational aids for user objects."

s/b

"The storage management component is focused on mapping logical constructs (e.g., files or database entries) to the physical organization of the storage media. In the OSD model, the logical constructs are called user objects (see 4.6.5). The root object (see 4.6.3), partitions (see 4.6.4), and collections (see 4.6.6) provide additional navigational aids for user objects."

added reference to 4.6.5 for consistency with other references in para.

Lingua 13) Remove policy/storage manager from model (Rejected) [378]

pg 16 - 3 references on page

Consider eliminating references to 'policy/storage manager'. This page is the only place this term appears. It adds nothing of value to the spec.

Reason for Rejection: The resolution for comment IBM 141) gives the policy/storage manager responsibility for constructing capabilities and managing the object version tag.

Lingua 14) Non-Support for the FORMAT OSD Command (Accepted, Substantive) [379]

The author marked this comment as technical.

pg 18, list item a)

See also comments LSI 5), LSI 6), LSI 7), and Seagate 2)

I would like to see the ability to have a mechanical 'lockout' (e.g. by means of a jumper) to prevent FORMAT OSD from being executed. This would necessitate this command to be removed from this list.

Editor's Note: During the 20 May discussion of this text, it was agreed that the text derives from some historical concerns that are no longer valid. The entirety paragraph containing the a, b, c list will be removed.

To facilitate preventing the FORMAT OSD command from being executed, the FORMAT OSD entry in the table listing OSD commands will be changed from **M (Mandatory)** to **O (Optional)**.

With support for FORMAT OSD being optional, the second to the last paragraph in 4.6.1 (Stored data object types, page 18, immediately following the cited text) is no longer valid and will be removed:

~~The required corrective actions for an error shall never include use of the FORMAT OSD command.~~

Note that the response to comment LSI 7) will result in the removal of the last paragraph in 4.6.1, in addition to all the other deletions described in the response to this comment.

Lingua 15) Add Current Command attributes page to table 3 (Rejected) [380]

pg 21, table 3

add Current Command attribute page to list

Reason for Rejection: The Current Command attributes page already appears in table 3. The Current Command attributes page (page number FFFF FFEh) is one of the attributes pages that is associated with any OSD object type. Since future versions of OSD might define other attributes pages that are associated with any OSD object type in addition to the Current Command attributes page, it is not appropriate to specifically name the Current Command attributes page in table 3.

Lingua 16) FrameMaker cross reference marker not on the correct paragraph (Accepted, Editorial) [381]

pg 23, sec 4.8.1, list item c)

"(see)" s/b "(see 7.1.2.13)"

Editor's Note: List entry c) is supposed to be the User Object Quotas attributes page. The problem is a technical gaff in the FrameMaker source file. The cross reference marker used to generate the text in entry c of the a, b, c list is located on the paragraph before the header for the User Object Quotas attributes page subclause. The correct location for the cross reference marker is on the header. The cross reference marker will be moved and the cross references will be regenerated to correct this problem.

Lingua 17) FrameMaker cross reference marker not on the correct paragraph (Accepted, Editorial) [382]

pg 23, sec 4.8.1, list

add item "d) User Object Quotas attributes page (see 7.1.2.14)"

Editor's Note: List entry c) is supposed to be the User Object Quotas attributes page. The problem is a technical gaff in the FrameMaker source file. The cross reference marker used to generate the text in entry c of the a, b, c list is located on the paragraph before the header for the User Object Quotas attributes page subclause. The correct location for the cross reference marker is on the header. The cross reference marker will be moved and the cross references will be regenerated to correct this problem.

Lingua 18) Put Security Method values in the Model clause (Rejected) [383]

pg 26, table 6

Add column for Value

Values corresponding to each of the Security methods currently undefined

Reason for Rejection: The cited values are properties of the fields that describe the Security Methods, not part of the Security Model. Placing the values in the model clause might suggest to some readers that there is a numeric relationship between the Security Methods. There is no such relationship and the standard must reflect this.

Lingua 19) Channel ID issue (Unresolved) [384]

The author marked this comment as technical.
 pg 27, sec 4.9.3.3, enumerated list item 2)
 See comment EMC 5)

"Target identifier" s/b "Logical Unit Name"

OK, I see that the target identifier at least *looks* symmetric to the initiator identifier. However, the relationship which has security attributes is really the relationship between the app client and the LU. The fact that we have no app client name seems poor justification for similarly crippling the other side of the connection. I would advocate that the best we can do, given the lack of app client names, is create a security token from the Initiator Identifier and the Logical Unit Name.

Editor's Note: Details of how this comment is being resolved can be found in 04-193.

Lingua 20) the same [s/b] the same protections (Accepted, Editorial) [385]

pg 28, 2nd para
 See comment EMC 5)

"provides the same as the ALLDATA security method." s/b "provides the same protections as the ALLDATA security method."

Lingua 21) Correct security methods in SPC-3 (No Action Taken) [386]

for SPC-3

The latest version of SPC-3 on the T10 website is spc3r17. Therein, clause 4.5.2.9, paragraph 1 still refers to security methods "2 or 3". These references should be changed to "CMDRSP or ALLDATA". If I just missed a proposal, I apologize.

Editor's Note: The problem no longer exists in SPC-3 r19. See 04-095 for details regarding removing the OSD-specific sense data descriptors from SPC-3 and inserting them in this standard.

Lingua 22) CMDRSP protects more than just command and CDB parameters (Accepted, Editorial) [387]

pg 29, last para of 4.9.3.4

"The CMDRSP security method protects against corruption of the command and CDB parameters while avoiding the overhead that may be required to protect all transferred data."

s/b

"The CMDRSP security method protects against corruption of the command and CDB parameters, as well as the response STATUS and Sense parameters, while avoiding the overhead that may be required to protect all transferred data."

Editor's Note: The cited sentence will be modified as follows:

The CMDRSP security method protects against corruption of the command, **command parameter data, status, and sense data and CDB parameters** while avoiding the overhead that may be required to protect all transferred data.

Lingua 23) Delete redundant paragraph (Accepted, Editorial) [388]

bottom of pg 34

The last paragraph on this page, along with the associated table 13, are semantically identical to the preceding paragraph. Suggest keeping the version with the table.

Lingua 24) field says 'creation' but attribute says 'created' (Accepted, Editorial) [389]

bottom of page 34 - same material as previous comment

references to "creation time" s/b "created time" In order to match terminology from attributes pages.

Editor's Note: To fully match the capability field name to the attribute name, the OBJECT CREATION TIME field will be changed globally to the OBJECT CREATED TIME field. Also, any other uses of 'creation time' will be changed to 'created time' as noted.

Lingua 25) field says 'creation' but attribute says 'created' (Accepted, Editorial) [390]

bottom of page 34 - same material as previous comment

Eliminate references to 'creation time from Root Timestamps attributes page'. This page does not have a creation time attribute.

Editor's Note: This comment will be resolved as described in the response to comment Lingua 24).

Lingua 26) OSO [s/b] OSD (Accepted, Editorial) [391]

pg 35, table 14

"OSO object type" s/b "OSD object type"

Lingua 27) Clarify capabilities tables 19 and 20 introduction (Accepted, Editorial) [392]

pg 37, 1st para of 4.9.4.4

See also comment Veritas 62)

"The validity of a specific command and some of the function-related fields in that command is determined by the presence of specific combinations of values in capability fields as shown in table 19. Any command may retrieve or set attributes and combinations of capability fields that allow those functions are shown in table 20."

s/b

"The validity of a specific command and some of the command function-related fields in that command is determined by the presence of specific combinations of values in capability fields as shown in table 19. Any command may retrieve or set attributes, and combinations of capability fields that allow those command functions are shown in table 20."

Changed for clarity: 'command function' is a defined term. More important for comprehension was adding the comma.

Editor's Note: The added comma does not work for the editor or for the author of comment Veritas 62). The hyphenation of command function-related seems dubious. The cited sentences will be modified as follows:

The validity of a specific command and some of the ~~function-related command function~~ (see 3.1.10) related fields in that command is determined by the presence of specific combinations of values in capability fields as shown in table 19. Any command may retrieve or set attributes. ~~and combinations~~ The combinations of capability fields that allow those functions are shown in table 20.

Lingua 28) At least one row per command in table 19 (Accepted, Editorial) [393]

Table 19

Table is hard to navigate. I tend to use the command as the primary index. If this navigation method is shared by others, I would suggest creating separate rows for "FORMAT OSD" and "PERFORM SCSI COMMAND".

Editor's Note: The FORMAT OSD command will be placed in a separate row (see 04-189). The text in the row containing PERFORM SCSI COMMAND will be line wrapped so that the command name starts on the left margin.

Lingua 29) Increase Granularity of Permissions for Commands Accessed via PERFORM SCSI COMMAND (Rejected) [394]

The author marked this comment as technical.
pg 40

Should not require Root, DEV_MGMT, and GLOBAL for (e.g.) LOG SENSE under PERFORM SCSI COMMAND.

Reason for Rejection: Yes, use of the LOG SENSE command should require ROOT, DEV_MGMT, and GLOBAL permissions. It is a command function outside the scope of the OSD command set defined in this standard. Unless or until, this standard defines an OSD-specific log page, accessing log pages is beyond the intended command usage defined by this standard and requiring substantial permissions to send a LOG SENSE command is completely appropriate.

Lingua 30) CREATE and CREATE AND WRITE can retrieve attributes from any created user object (Accepted, Substantive) [395]

pg 42, table 20, entry for USER, GET_ATTR, NONE

"As part of a CREATE command or CREATE AND WRITE command, the retrieval of attributes from any attributes page associated with the user object with a Partition_ID matching the value in the credential PARTITION_ID field and the largest valued User_Object_ID created."

s/b

"As part of a CREATE command or CREATE AND WRITE command, the retrieval of attributes from any attributes page associated with the user object with a Partition_ID matching the value in the credential PARTITION_ID field and the largest valued User_Object_ID created by that command."

Editor's Note: Actually, a CREATE command or CREATE AND WRITE command is allowed to retrieve attributes from all created user objects using the list format for attributes retrieval (see 7.1.3.4). Therefore, the cited sentence will be modified as follows:

As part of a CREATE command or CREATE AND WRITE command, the retrieval of attributes from any attributes page associated with **any user object created by the command. ~~the user object with a Partition_ID matching the value in the credential PARTITION_ID field and the largest valued User_Object_ID created.~~**

For the same reason, the USER/SET_ATTR/NONE text will be modified as follows (including the addition of commas shown in the response to comment Lingua 31):

As part of a CREATE command or CREATE AND WRITE command, the setting of attributes in any attributes page, other than User Object Security attributes page, associated with **any user object created by the command. ~~the user object with a Partition_ID matching the value in the credential PARTITION_ID field and the largest valued User_Object_ID created.~~**

For the same reason, the USER/SET_ATTR and SECURITY/NONE text will be modified as follows:

As part of a CREATE command or CREATE AND WRITE command, the setting of attributes in any attributes page associated with ~~any user object created by the command. the user object with a Partition_ID matching the value in the credential PARTITION_ID field and the largest valued User_Object_ID created.~~

Lingua 31) Add commas to clarify security attributes page exceptions (Accepted, Editorial) [396]

The author marked this comment as technical.

pg 42, 43, table 19, *each* entry with SET_ATTR, but without SECURITY (7 items)

Attribute-related functions are misconstruction due to ordering of terms. Terms need order reversal. Example:

"The setting of attributes in any attributes page other than User Object Security attributes page associated with the user object with a Partition_ID matching the value in the credential PARTITION_ID field and a User_Object_ID matching the value in the object descriptor SINGLE OBJECT_ID field."

s/b

"The setting of attributes in any attributes page associated with the user object with a Partition_ID matching the value in the credential PARTITION_ID field and a User_Object_ID matching the value in the object descriptor SINGLE OBJECT_ID field, with the exception of than the User Object Security attributes page."

Editor's Note: If there were more than one exception to list, the sentences as written might be viewed as to complex, but that is not the case. As written, the exception is stated as close as possible to the modified text. This is good. The addition of commas around the exception text is all that is needed to clarify matters. For example:

The setting of attributes in any attributes page, other than User Object Security attributes page, associated with the user object with a Partition_ID matching the value in the credential PARTITION_ID field and a User_Object_ID matching the value in the object descriptor SINGLE OBJECT_ID field.

Lingua 32) Incorrect cross reference (Accepted, Editorial) [397]

pg 45, list item 3) B) a)

"(see 4.9.3.2)" s/b "(see 4.9.3.3)"

Lingua 33) Use small caps for field contents (Rejected) [398]

pg 46, para 2

"If the capability object type, permissions bit mask, and object descriptor do not allow a command..."

s/b

"If the capability OBJECT TYPE, PERMISSIONS BIT MASK, and OBJECT DESCRIPTOR are set to values that do not allow a command..."

Reason for Rejection: The nomenclature in SCSI standards is either "the contents of the FIELD_NAME field" or "the field_name", with the use of small caps exactly as shown. The latter form is shorter, appropriate in the cited context, and easier to read.

Lingua 34) Add Partition_ID Field in FORMAT OSD Command (Rejected) [399]

The author marked this comment as technical.
pg 46, 4.9.5.3, list item 1) B) a)

Why does the Credential's partition id not match that of the command? Intuitively, it would seem to make sense that it should. Indeed, this mechanism has no possibility of working for commands to LIST all User objects in a partition, as that would require a nonzero PARTITION_ID.

Additionally, FORMAT OSD is the only command listed herein that does not have a PARTITION_ID field.

I am herein advocating that a PARTITION ID field be added to FORMAT OSD (which must be set to 0), and eliminating this entire special case (all of item 1 B a).

Reason for Rejection: The requested change would break the security model.

Lingua 35) Task Management Functions Are Not Directed To User Objects (Rejected) [400]

The author marked this comment as technical.
pg 46, 4.9.5.3, list item 1) B) b)
see also: comment Lingua 44)

While my particular view is that Task Management Functions are directed to LUs, (not (e.g) User Objects), the definition in 6.13 of this command allows the direction of certain TMFs to User Objects. If this (ill advised) construct is to survive, then the PARTITION_ID field cannot be herein set to 0. (I have more to say about TMFs being sent to User objects later on...)

Reason for Rejection: An ABORT TASK task management function is directed to a single task. In this standard that task is restricted to acting on a single object. Therefore at least one task management function is directed to an object (albeit with one level of indirection).

Lingua 36) Credential integrity check value is not used in a reconstructed credential (Accepted, Editorial) [401]

pg 46, 4.9.5.3

add list item "3) Concatenating the computed credential integrity check value (see 4.9.5.4)."

Editor's Note: Actually, the CREDENTIAL INTEGRITY CHECK VALUE field is not used in a reconstructed credential. The credential integrity check value is used as the capability key for computing other integrity check values, but the use of the CREDENTIAL INTEGRITY CHECK VALUE field in the defined credential format is not important to the algorithms.

To clarify this, the following sentence will be added to the end of 4.9.5.3:

The CREDENTIAL INTEGRITY CHECK VALUE field is not used in a reconstructed credential.

Lingua 37) An integrity check values [s/b] Integrity check values (Accepted, Editorial) [402]

pg 49, 1st para of 4.9.7

"An integrity check values..." s/b "Integrity check values..."

Lingua 38) rest [s/b] reset (Accepted, Editorial) [403]

pg 53, item a) A)

"...rest event..." s/b "...reset event..."

Lingua 39) lost [s/b] loss (Accepted, Editorial) [404]
pg 53, item b)

"...lost of storage..." s/b "...loss of storage..."

Lingua 40) is [s/b] of (Accepted, Editorial) [405]
pg 53

"Individual OBSD (see 3.1.26) implementations may use whatever technologies they choose to implement stable storage (e.g., an OBSD may implement stable storage as a combination **is** non-volatile random access memory and disk devices)."

s/b

"Individual OBSD (see 3.1.26) implementations may use whatever technologies they choose to implement stable storage (e.g., an OBSD may implement stable storage as a combination **of** non-volatile random access memory and disk devices)."

(replace 'is' with 'of')

Lingua 41) zero [s/b] NOSEC (Accepted, Editorial) [406]
pg 59, 1st sentence fragment

"OSD logical unit is not zero,..." s/b "OSD logical unit is not NOSEC..."

Lingua 42) Eliminate Page Formats for Attributes (Rejected) [407]
The author marked this comment as technical.
pg 62, 63

Eliminate entire mechanism of page oriented get & set attributes.

There is absolutely no reason to have both these formats. No device can assume that all initiators will only use the page oriented mechanism. Accordingly, they will have to implement the list oriented mechanism, which is a functional superset of the page oriented, as well as the page oriented. This illusory flexibility of mechanisms brings no benefit whatsoever. It does, however carry a requirement to develop and test both mechanisms where only one would do.

Reason for Rejection: The page format allows an organization and simplification of attribute setting and retrieval.

Lingua 43) taSk [s/b] Task (Accepted, Editorial) [408]
page 92, table 57

"SAM-3 taSk..." in header s/b "SAM-3 task..."

Lingua 44) Require Task Management Functions be Sent To the Root Object (No Action Taken) [409]

The author marked this comment as technical.

sec 6.13

see also: comment Lingua 35)

Require that all TMFs be sent to root object.

TMFs are NOT specific to objects -

They operate on Tasks in the Task Set

Operation is performed by the Task Manager

- this all happens in the LU, but OUTSIDE the Device Server

Also, the Task Router isn't even in the LU - it is in the Port

Device Server knows only the Current Task - relies on the Task Manager to feed it in an efficient manner.

Editor's Note: This comment was withdrawn by its author.

10. LSI Logic Corp.

John Lohmeyer from LSI Logic Corp. submitted the following comments on a No vote.

LSI 1) Remove revision history (Accepted, Editorial) [410] Page iii, Revision Information

The Revision History should be deleted before forwarding (but you already knew that).

LSI 2) Insert SNIA OSD TWG membership (Unresolved) [411] Page xviii, Foreword

Can we get a list of the SNIA OSD TWG members to insert in the foreword?

Editor's Note: Awaiting a decision from SNIA legal counsel.

LSI 3) rest [s/b] reset (Accepted, Editorial) [412] Page 8, 3.1.51 volatile cache

'rest event' should be 'reset event'.

LSI 4) e.g. [s/b] i.e. (Accepted, Editorial) [413] Page 16, 4.4 Elements of the example configuration

The "e.g." in the paragraph below Figure 3 should be "i.e." because it includes the phrase "or other storage devices".

LSI 5) Remove requirement for GOOD status (Accepted, Substantive) [414] Page 18, 4.6.1 Stored data object types See also comments Lingua 14), and Seagate 2)

The third-from-last paragraph requires an OSD device to report GOOD status even with a head crash for FORMAT OSD and several other OSD commands. I suspect this was not intended. Either remove this paragraph or re-word it to deal with all error conditions.

Editor's Note: This comment will be resolved as described in the response to comment Lingua 14), specifically the cited paragraph and its associated a, b, c list will be removed for the reasons described in the response to comment Lingua 14).

LSI 6) Confusing FORMAT OSD requirement to be removed (Accepted, Substantive) [415] Page 18, 4.6.1 Stored data object types, Penultimate paragraph See also comments Lingua 14), and Seagate 2)

The phrase, shall never, seems inappropriate for a standard. I think this paragraph should be replaced with, "The FORMAT OSD command is not intended to be used for error recovery."

Editor's Note: This comment will be resolved as described in the response to comment Lingua 14), specifically the cited sentence will be removed for the reasons described in the response to comment Lingua 14).

LSI 7) Remove potentially troublesome error recovery statement (Accepted, Substantive) [416]

Page 18, 4.6.1 Stored data object types, Last paragraph
See also comments Lingua 14), and Seagate 2)

This paragraph adds no value and should be deleted.

Editor's Note: Since the last paragraph in 4.6.1 is states no requirements and may conflict with future efforts to standardize error recovery, it will be removed, to whit:

~~A vendor specific process may be used to verify the integrity and perform the recovery of an OSD logical unit.~~

LSI 8) 'and' is on the wrong 1,2,3 list entry (Accepted, Editorial) [417]

Page 20, 4.7.2 Command function ordering for commands that get and/or set attributes, second paragraph

The "and" at the end of item 2) seems misplaced. I think it belongs at the end of item 3).

Editor's Note: The 'and' is misplaced. It will be moved.

LSI 9) All OSD object types is confusing (Accepted, Editorial) [418]

Page 20, 4.7.1 Attributes parameters, fifth paragraph
See also comment Veritas 36)

The "All OSD object types" item is confusing. Perhaps an example might help.

Editor's Note: The following modifications will be made:

- In table 3 (pg 21), the F000 0000h row will be modified as follows: "Any OSD object type (i.e., root, partition, collection, or user)"
- In 4.7.1 (pg 20), the cited text will be modified as follows: "e) ~~All~~ Any OSD object ~~types~~ type (see table 3 in 4.7.3)"
- The first sentence following the a,b,c list will be modified as follows: "With the exception of attributes pages in the attributes page number range assigned to ~~all~~ any OSD object ~~type types (see table 3 in 4.7.3)~~, the same attributes page shall not be associated with more than one OSD object."
- The last sentence before table 4 (pg 22) will be modified as follows: "Except for the attributes page numbers that apply to ~~all~~ any OSD object ~~type types~~ (i.e., F000 0000h through FFFF FFFFh), the ranges of attributes page numbers shown in table 3 are subdivided as shown in table 4."

LSI 10) FrameMaker cross reference marker not on the correct paragraph (Accepted, Editorial) [419]

Page 23, 4.8.1 Introduction

Item c) in the first list references a blank destination.

Editor's Note: The problem is a technical gaff in the FrameMaker source file. The cross reference marker used to generate the text in entry c of the a, b, c list is located on the paragraph before the header for the User Object Quotas attributes page subclause. The correct location for the cross reference marker is on the header. The cross reference marker will be moved and the cross references will be regenerated to correct this problem.

LSI 11) by command [s/b] by a command (Accepted, Editorial) [420]

Page 24, 4.8.4 Changing quotas, first paragraph

The phrase "by command" should be "by a command".

LSI 12) far in the future nonce(s) [s/b] far-in-the-future nonce(s) (Accepted, Editorial) [421]

Page 48, 4.9.6.2 Device server validation of request nonces (global)

The phrase "far in the future" should be "far-in-the-future".

Editor's Note: The requested change will be made in the following locations: the last words on pg 48, in the last sentence of 4.9.6.2, the heading and first sentence in 4.9.6.3.2, the heading and first sentence 4.9.6.3.3, and in 7.1.2.21 (Partition Security attributes page) in the second paragraph before table 100 (pg 136).

LSI 13) rest [s/b] reset (Accepted, Editorial) [422]

Page 53, 4.10 Data persistence model

The phrase "rest event" should be "reset event".

LSI 14) Admit that no attribute page numbers are assigned by other standards (Accepted, Editorial) [423]

Page 152, Annex A

See also comment IBM 140)

Since there are none at this point in time, a note saying so would make this annex less confusing.

Editor's Note: The cited sentence will be modified as follows:

At the time of publication, no attribute page numbers are assigned by other standards. The attributes page numbers ~~assigned available for assignment~~ by other standards ~~at the time of publication is~~ are shown in table A.1.

11. Panasas

David Nagel from Panasas submitted the following comments as part of the ENDL Texas (see page 42) ballot.

Panasas 1) user object [s/b] partition (Accepted, Editorial) [424]

PDF pg 37, pg 18, 4.6.4, p 2, s 2

user object [s/b] partition [The ID under discussion is a Partition_ID not a User_Object_ID.]

Panasas 2) Reporting credential and capability errors is critical (Accepted, Substantive) [425]

The author marked this comment as technical.

PDF pg 39, pg 20, 4.7.2

There is no description of the order in which commands are processed. In the Error Handling document I posted in January, there were two ordering that the error group thought essential. They are:

The reporting of errors during initial command processing is important to avoid leaking any information when security may be compromised. Therefore, the REPORTING of errors should proceed in the following order.

1. Any security errors against the integrity of the credential (and cmd/data if requested) should be reported first. The including verifying any INTEGRITY CHECK values, key version, credential format, and checksum algorithm requested.

2. Any errors due to NONCE processing must be reported.

3. Any errors in the capability permissions, expired capability, unknown object, invalid create time or version number

Only the first error in this ordering should be reported. This ordering does not mandate that a device process a command in this order. Only that the device report errors in the specified order.

Editor's Note: This issue will be resolved as described in the responses to comment EMC 7) and comment IBM 141).

Panasas 3) partition [s/b] collection (Accepted, Editorial) [426]

PDF pg 42, pg 23, 4.7.5, Table 5, r 3, c 3

partition [s/b] collection [This row is describing the Collection Directory, not the Partition Directory.]

Panasas 4) When to Test Quotas (Accepted, Substantive) [427]

The author marked this comment as technical.
PDF pg 43, pg 24, 4.8.3

The spec says that 'tests for quota errors may be made at any time during the processing of a command.' This presents several problems. First, writes may be partially complete when quota errors occur, terminating in a partial write, which might have partially overwritten data. Second, if (quota < already-used-space), any space-consuming command should immediately terminate. However, because quota tests can be performed at any time during the processing of the command, it is possible for the OSD to allow a write-like command to consume more space before terminating with an error.

We recommend that tests for quota should occur before any data or attributes have been modified or newly written. In other words, a command that would result in exceeding quota is terminated before any permanent OSD object-state has been modified.

Editor's Note: The following new paragraph will be added to 4.8.2 (Quota errors):

The device server shall not terminate a command for quota errors after any user data or attributes have been modified.

The following paragraph will be removed from 4.8.3 (Quota testing):

~~Tests for quota errors may be made at any time during the processing of a command. The processing of a command may be partially completed at the time a quota error is detected. The device server is not required to restore the state of the OSD logical unit to the state that was present before processing was begun for the command in which the quota error has been detected.~~

Beyond that, everything else is just an implementation detail.

Panasas 5) Per-Partition Security Is Wrong (Unresolved) [428]

The author marked this comment as technical.
PDF pg 45, pg 26, 4.9.3.1, 1st s after Table 6

Why is security defined on a per-partition basis? I thought that the partition (and/or root object) encoded the MINIMUM level of security and that every capability encoded the actual level of security used to generate the capability?

This is important because if every object w/in a partition (or root device) must use the same level of security, then providing different clients with different security needs becomes impossible. For example, a client that sits behind a firewall may need a much lower level of security than a client outside the firewall.

Panasas 6) FrameMaker cross reference marker not on the correct paragraph (Accepted, Editorial) [429]

PDF pg 90, pg 71, 6.2, 2nd p on pg

This paragraph is a total nonsense. A sentence ends without a period. There is a missing cross reference.

Editor's Note: The problem is a technical gaff in the FrameMaker source file. The cross reference marker used to generate the text the User Object Quotas attributes page subclause cross reference for the maximum user object length attribute description is located on the wrong paragraph. The correct location for the cross reference marker is on the header. The cross reference marker will be moved and the cross references will be regenerated to correct this problem.

Panasas 7) Collection_Object_IDs in same name space as User_Object_IDs (No Action Taken) [430]

PDF pg 92, pg 73, 6.3, 1st a,b,c list on pg, item a

Actually, the device server shall not allow the same ID to be associated with more than one user-object or collection-object (remember that user and collection objects share the same namespace)

Editor's Note: True, and the condition is covered by item b in the cited list, to wit:

Within a partition, the device server shall not allow:

- a) The same User_Object_ID to be associated with more than one user object at any point in time; or
- b) A User_Object_ID to have the same value as any assigned Collection_Object_ID.

12. Seagate Technology

Gerald Houlder from Seagate Technology submitted the following comments on a No vote.

Seagate 1) partition [s/b] collection (Accepted, Editorial) [431]

Seagate (RiedelE) #1

Page 17

"The data component of a [partition](#) is the list of User_Object_IDs." should be "The data component of a [collection](#) is the list of User_Object_IDs."

Seagate 2) Confusing FORMAT OSD requirement to be removed (Accepted, Substantive) [432]

Seagate (RiedelE) #2

Page 18

See also comments Lingua 14), LSI 5), LSI 6), and LSI 7)

This sentence is confusing in what it is trying to limit. Perhaps it could just read "the above commands shall never..." rather than seeming to apply to any error anywhere in the document.

Editor's Note: This comment will be resolved as described in the response to comment Lingua 14), specifically the cited sentence will be removed for the reasons described in the response to comment Lingua 14).

Seagate 3) Insert missing comma (Accepted, Editorial) [433]

Seagate (RiedelE) #3

Page 18

missing ","

Editor's Note: The cited text will be modified as follows:

The combination of Partition_ID and User_Object_ID uniquely identifies the root object, each partition, each collection and each user object. Partition_ID and User_Object_ID values are assigned as shown in table 2.

Seagate 4) Enhance list of commands that cannot be directed to a collection (Rejected) [434]

Seagate (RiedelE) #4

Page 19

See also comment IBM 19), comment IBM 21), and comment Other 1)

should also include (i.e. exclude) CREATE, REMOVE, CREATE&WRITE

Reason for Rejection: The propose list is not complete and cannot be maintained completely over time (e.g., what about REMOVE PARTITION, SET KEY, soon to be FLUSH PARTITION, and ...?).

The intent of the cited paragraph is to state that a collection does not include data. With the exception of the CREATE AND WRITE command the proposed changes only confuse that intent.

Comment Other 1) describes the only changes that will be made in the cited paragraph.

Seagate 5) store [s/b] stored (Accepted, Editorial) [435]

Seagate (RiedelE) #5
Page 20

should be "stored" with "d"

Editor's Note: The cited text will be modified as follows:

The GET ATTRIBUTES command (see 6.9) and SET ATTRIBUTES command (see 6.18) allow attributes to be retrieved and ~~store~~ stored without performing other command functions (see 3.1.10).

Seagate 6) FrameMaker cross reference marker not on the correct paragraph (Accepted, Editorial) [436]

Seagate (RiedelE) #6
Page 23

dangling reference to nowhere

Editor's Note: The problem is a technical gaff in the FrameMaker source file. The cross reference marker used to generate the text in entry c of the a, b, c list is located on the paragraph before the header for the User Object Quotas attributes page subclause. The correct location for the cross reference marker is on the header. The cross reference marker will be moved and the cross references will be regenerated to correct this problem.

Seagate 7) Add Security Method Field to CDB (Accepted, Substantive) [437]

Seagate (RiedelE) #7
Page 26

the information about which security method a request was prepared for needs to be in the CDB. The device must be able to report an error (via the "may complete" below) if the request was prepared with ALLDATA and CMDRSP is in use. This means that the client is not getting the security level they think they are, and that should be reportable as an error. Irrespective of this, the security level needs to be known in the CDB for reasonable performance at the device (no wasted checksum work).

Editor's Note: The cited paragraph will be removed in its entirety because the text describing the new Capability field proposed in comment Seagate 9) will define the error to be reported for the wrong security method. The specific change to be made is:

~~A command prepared for a security mode other than the one the device server uses for processing may complete without errors (e.g., a command prepared for the ALLDATA security method may complete without errors reported by the device server if the CMDRSP security method is in use because the preparations for the ALLDATA security method include the preparations that are necessary for the CMDRSP security method).~~

Seagate 8) Capability Validation is NOT a Security Function (Accepted, Substantive) [438]

Seagate (RiedelE) #8
Page 27

this should be reworded to be specific about which security features are no used. E.g. permission bits are still checked, version numbers are still checked - just no hashes checked or keys used. The current wording of "no OSD security feature" is imprecise and too broad. For example, are SET KEY and SET MASTER KEY "security features"?

Editor's Note: This comment will be resolved as described in the response to comment IBM 141).

Seagate 9) Add Security Method to Capability (Accepted, Substantive) [439]

Seagate (RiedelE) #9
Page 33

capability needs to contain an indication of the security method used to create it - as already noted on page 26 (see comment Seagate 7).

Editor's Note: The requested change has been included in the response to comment IBM 141). Details of how this comment is being resolved can be found in 04-193.

Seagate 10) Delete redundant paragraph (Accepted, Editorial) [440]

Seagate (RiedelE) #10
Page 34

this paragraph is redundant with Table 13 and should be removed.

Editor's Note: Actually, the cited paragraph is redundant with the paragraph that introduces table 13 and table 13. The statement about the meaning of zero appears only in the introductory paragraph. Still, the cited paragraph will be deleted, to wit:

~~The OBJECT CREATION TIME field specifies the contents of the creation time attribute in the User Object Timestamps attributes page (see 7.1.2.18), Collection Timestamps attributes page (see 7.1.2.17), Partition Timestamps attributes page (see 7.1.2.16), or Root Timestamps attributes page (see 7.1.2.15), for the OSD object to which the credential applies. A value of zero specifies that any object creation time is allowed.~~

The OBJECT CREATION TIME field specifies the contents of the creation time attribute for the OSD object (see table 13) to which the credential applies. A value of zero specifies that any object creation time is allowed.

Seagate 11) Add Permission Bit for APPEND (Accepted, Substantive) [441]

Seagate (RiedelE) #11
Page 35

request addition of a separate permissions bit for APPEND - want to allow clients to append at end of file without being able to scribble in the middle

Editor's Note: Bit 0 in byte 31 of table 15 will be changed from **Reserved** to **APPEND**.

The following changes will be made in the text describing the bits in table 15:

A WRITE bit set to one allows ~~processing of the WRITE command (see 6.21), write access to the data in a user object,~~ but not ~~to the access to user object~~ attributes. A WRITE bit set to zero prohibits ~~processing of the WRITE command. write access to the data in a user object.~~

...

An OBJ_MGMT (object management) bit set to one allows command functions that may change how the OSD logical unit handles an OSD object without affecting the stored data, stored attributes, commands in the task set, or security for the OSD object. A OBJ_MGMT bit set to zero prohibits such command functions.

~~An APPEND bit set to one allows processing of the APPEND command (see 6.2), but not access to user object attributes. A APPEND bit set to zero prohibits processing of the APPEND command.~~

Seagate 12) Add Permission Bit for APPEND (Accepted, Substantive) [442]

Seagate (RiedelE) #12
Page 37

APPEND instead of WRITE in row 1 of table 19

Seagate 13) Incorrect cross reference (Accepted, Editorial) [443]

Seagate (RiedelE) #13
Page 45, 4.9.5.2, 3) A) a)

4.9.3.2 should be 4.9.3.3

Seagate 14) is [s/b] in (Accepted, Editorial) [444]

Seagate (RiedelE) #14
Page 45, 4.9.5.2, p 1, s 1

"is" should be "in"

Seagate 15) Working key selection wording is confusing (Accepted, Editorial) [445]

Seagate (RiedelE) #15
Page 47

order of the phrases makes this confusing - the key is the partition, not the KEY VERSION field - suggest moving "identified by the KV field in the cap" to the end of the sentence, i.e.:

If the OBJECT TYPE field in the capability (see 4.9.4.3) contains COLLECTION or USER, the authentication working key ~~identified by the KEY VERSION field in the capability~~ for the partition identified by the PARTITION_ID field in the CDB ~~identified by the KEY VERSION field in the capability~~;

Editor's Note: The proposed change does not diminish the confusion. It only moves the confusion around, making the CDB of concern be the CDB identified by the KEY VERSION field. The real problem is that word 'for' does not separate what precedes it from what follows sufficiently. The sentence will be modified as follows:

If the OBJECT TYPE field in the capability (see 4.9.4.3) contains COLLECTION or USER, the authentication working key:

- a) Identified ~~identified~~ by the KEY VERSION field in the capability; and
- b) Associated with ~~for~~ the partition identified by the PARTITION_ID field in the CDB;

Seagate 16) Add example regarding changing security version key values (Rejected) [446]

Seagate (RiedelE) #16
Page 47

this is confusing to the non-cognoscenti. Please include "(e.g. if the version key is changed to a value that has been used before, then capabilities that had been invalid may become valid again)" or similar explanatory words

Reason for Rejection: The request change represents a degree of tutoring that is not appropriate for a standard. The existing text states the recommended behavior. That is all the standard should say.

Seagate 17) Require Command Terminations if working keys restrict far-in-the-future nonces (Rejected)

[447]

Seagate (RiedelE) #17

Page 49

This "Should" should be "Shall", there really is no way to get around not doing this.

Reason for Rejection: Sure there is a way to get around doing this. The device server processes the command. Circumstances may oblige the device server to terminate commands, but there is no obvious way in which the interface between the application client and the device server breaks down if some commands are terminated and others are processed.

Seagate 18) SHA1 shall not be truncated (Accepted, Substantive) [448]

Seagate (RiedelE) #18

Page 50

See also the following comments IBM 56), IBM 141), Seagate 19), and Veritas 69)

Why is this truncation suggested? Perhaps the concern is over space in the CDB. Using the full 20 byte SHA1 would more closely follow usage of SHA1 in other areas and enable interoperability (as well as higher security). A truncation to 16 bytes is also common usage. 12 bytes is very much non-standard usage.

Editor's Note: In the cited subclause, the following paragraph will be removed:

~~If the field specified to contain the integrity check value is has fewer bytes than the output of the cryptographic function, then the cryptographic function output shall be truncated to fit in the field. The output of an HMAC-SHA1 function is 20 bytes. It is common in this standard to specify field sizes of only 12 bytes for integrity check values, meaning that the HMAC-SHA1 function output of 160 bits is truncated into 96 bits.~~

All 12 byte integrity check value fields will be increased to 20 bytes. Details of field size changes can be found in 04-193.

Seagate 19) is has [s/b] has (Accepted, Editorial) [449]

Seagate (RiedelE) #19

Page 50

remove "is"

Editor's Note: This comment would have been accepted, except that comment Seagate 18) specifies that the entire cited paragraph be removed.

Seagate 20) Integrity check values for SET KEY and SET MASTER KEY not clear (Accepted, Editorial)

[450]

Seagate (RiedelE) #20

Page 51

is this [SET KEY and SET MASTER KEY integrity check value] computation/authentication required regardless of the security level? If so, it should state that.

Editor's Note: No, the computation and verification is required only if specified by the security method. To clarify this the first sentence in the subclause will be modified as follows:

Like all other commands, the SET KEY command and SET MASTER KEY command **may** use a credential and the capability contained in that credential to authorize processing of the command **as described in 4.9**.

Seagate 21) Working Key Generation Key Values Are Not Saved (Accepted, Substantive) [451]

Seagate (RiedelE) #21
Page 51

OSD saves only one key (auth key) for working keys, gen key is not stored (or "saved")

Editor's Note: Including the changes described in the resolutions for comments HP 51), HP 120), and Intel 21), the cited text will be modified as follows:

~~The OSBD shall save two~~ Two secret key values shall be saved for each master, drive root, and partition, ~~and working~~ secret key as follows:

Seagate 22) Prohibiting changes in unused Data-In Buffer bytes (Accepted, Substantive) [452]

Seagate (RiedelE) #22
Page 55
See also comment Veritas 76)

this is very confusing - the device is creating these bytes, how can it not "alter" them? does this mean fill with zeros? Suggest that this phrase be stricken.

Editor's Note: The device server is transmitting the bytes to an application client's buffer. The requirement is that the device server not act in such a way as to cause the unused bytes in the application client's buffer not be altered. This is a critical requirement. Its absence will require initiators to construct special areas for all interactions with OSD devices and to copy data to/from those areas with all the added overhead that entails.

To clarify the requirement, the cited text will be modified as follows:

The device server shall not ~~alter send data to the initiator device that causes~~ unused bytes in the Data-In Buffer to be overwritten.

Seagate 23) Inconsistencies in Data-In Buffer integrity check value segment (Accepted, Editorial) [453]

Seagate (RiedelE) #23
Page 55

the size of this [table 25] segment and the wording needs to be consistent with Table 9 on page 31. Likely also should provide a ref to that table/section.

Editor's Note: Maintaining consistency between the two tables could be very difficult. For example, changing the SHA1 truncation (see comment Seagate 18) will affect both tables. Therefore, the sizes will be made variable. Also, a cross reference to 4.9.3.5 will be added to the Data-In Buffer integrity check value segment description.

Seagate 24) Inconsistencies in Data-Out Buffer integrity check value segment (Accepted, Editorial) [454]

Seagate (RiedelE) #24
Page 56

[table 27] needs to be consistent with Table 8 on page 30

Editor's Note: Maintaining consistency between the two tables could be very difficult. For example, changing the SHA1 truncation (see comment Seagate 18) will affect both tables. Therefore, the sizes will be made variable. Also, a cross reference to 4.9.3.5 will be added to the Data-Out Buffer integrity check value segment description.

Seagate 25) Missing Collection_Object_ID parameter description (Rejected) [455]

Seagate (RiedelE) #25
Page 67

a section 5.2.9 on Collection_Object_ID is missing here

Reason for Rejection: The COLLECTION_OBJECT_ID field appears in only two commands, LIST COLLECTION and REMOVE COLLECTION. With the changes resulting from these letter ballot comments resolutions, a third command will be added, FLUSH COLLECTION. This is not a sufficient number of commands for the field to be considered common to many CDBs. So the COLLECTION_OBJECT_ID field in situ in the commands were it is needed. The situation is similar to that of the ALLOCATION LENGTH field in the LIST and LIST COLLECTION commands or the SEED field in the SET KEY and SET MASTER KEY commands.

Seagate 26) Update Timestamp Bypass Mechanism (Accepted, Substantive) [456]

Seagate (RiedelE) #26
Page 67

New Timestamp Bypass Mechanism for OSDv1 Rev 10 (as agreed upon)

The 1-byte timestamp bypass attribute remains the same except a new value is defined as follows:

00h Timestamps are updated
01h - 7Eh Reserved
7Fh Timestamps are not updated
80h - DFh Reserved
E0h - FEh Vendor specific
FFh Use CDB timestamps control field

A timestamps control field remains the same except that the 1-bit ATTR-BYP field is eliminated (merged into the other field). The new values of the timestamps control field are defined as:

0h Timestamps are updated
01h - 7Eh Reserved
7Fh Timestamps are not updated
80h - DFh Reserved
E0h - FFh Vendor specific

Editor's Note: Including the changes described in comment Seagate 48), table 84 will be modified as follows:

Table 84 — Timestamp bypass attribute values

Value	Description
00h	Timestamps shall be updated as described in the subclause that defines them
01h to 7Eh	Reserved
7Fh	Timestamps shall not be updated ^a
80h to DFh	Reserved
E0 to FFh FEh	Vendor specific
FFh	Timestamps shall be updated as specified by the TIMESTAMPS CONTROL field in the CDB (see 5.2.7)

^a A timestamp attribute that has never been updated shall have a length of six and a value of zero. Bypassing a timestamp update shall not affect any previously established timestamp attribute values.

Including the changes described by comment Veritas 82), 5.2.7 (Timestamps control) will be modified as follows:

The If the timestamp bypass attribute in the applicable Root Timestamps attributes page (see 7.1.2.15) or Partition Timestamps attributes page (see 7.1.2.16) contains FFh, the `TIMESTAMPS CONTROL` field specifies the timestamp update policy (see table 38 37) for the command.

Table 37 — Timestamps control format

Bit	7	6	5	4	3	2	1	0
Byte								
12	ATTR_BYP	TIMESTAMP-BYPASS						

If the `ATTR_BYP` (attributes bypass) bit is set to zero, the contents of the timestamp bypass field shall be ignored and the timestamps shall be updated as specified by the timestamp bypass attribute in the applicable Root Timestamps attributes page (see 7.1.2.15) or Partition Timestamps attributes page (see 7.1.2.16).

If the `ATTR_BYP` bit is set to one, the timestamps shall be updated as specified by the `TIMESTAMP-BYPASS` field (see table 38) in the timestamps control byte.

Table 38 — Timestamp bypass Timestamps control values

Value	Description
0h	Timestamps shall updated as described in the subclause that defines them
01h to 7Eh	Reserved
7Fh	Timestamps shall not be updated
80h to DFh	Reserved
E0 to FFh	Vendor specific

A timestamp attribute (see 7.1) that has never been updated shall have a length of six and a value of zero.

Bypassing a timestamp update shall not affect any previously ~~updated or not updated~~ established timestamp attribute values.

Seagate 27) FrameMaker cross reference marker not on the correct paragraph (Accepted, Editorial) [457]

Seagate (RiedelE) #27
Page 71

some type of cut & paste error starting at "The COLLECTIONS..." - also happens on page 75 and page 103 in a similar paragraph.

Editor’s Note: The problem is a technical gaff in the FrameMaker source file. The cross reference marker used to generate the text the User Object Quotas attributes page subclause cross reference for the maximum user object length attribute description is located on the wrong paragraph. The correct location for the cross reference marker is on the header. The cross reference marker will be moved and the cross references will be regenerated to correct this problem.

Seagate 28) FLUSH OBJECT command should include collections (Rejected) [458]

Seagate (RiedelE) #28
Page 79, FLUSH OBJECT command

include "or collection object"

Reason for Rejection: As agreed in the response to comment HP 76), a separate FLUSH COLLECTION command will be defined, making the suggested change inappropriate for the FLUSH OBJECT command.

Seagate 29) Reference subclause for COLLECTION_OBJECT_ID description (Rejected) [459]

Seagate (RiedelE) #29
Page 79

reference 5.2.9 about COLLECTION_OBJECT_IDs

Reason for Rejection: As noted in the response to comment Seagate 25), there are not enough repeated uses of the COLLECTION_OBJECT_ID field to motivate writing a separate subclause to define the field.

Seagate 30) FORMAT OSD should not change security attributes (Accepted, Editorial) [460]

Seagate (RiedelE) #30
Page 80

the security level is an attribute and should not be set to defaults (should be left as-is)

Editor's Note: The cited text will be modified as follows:

... and set the attributes for the root object and partition zero ~~to defaults as defined by this standard.~~

Similar wording was reviewed without comment on the page 81.

Seagate 31) Describe intended use of LIST IDENTIFIER field (Rejected) [461]

Seagate (RiedelE) #31
Page 85

could there be some additional explanatory text here? something like "The LIST IDENTIFIER is intended only to allow a device to know if the LSTCHG bit should be set - i.e. for the device to determine if the list of OSD objects has changed since a given LIST IDENTIFIER was returned. The position in the list and the location of restart is completely determined by the CONTINUATION OBJECT_ID and does not use the LIST IDENTIFIER." This makes it more clear what each piece is used for and does not seem to exceed the amount of "for example" text allowed in other parts of the standard.

Reason for Rejection: The purpose of the field is no relevant. From the point of view of the standard, the device server may use the contents of the LIST IDENTIFIER field for any purpose the device server needs. Attempting to suggest a restricted use of the LIST IDENTIFIER field is nothing less than attempting to force one device server implementation philosophy on all implementations. That is not the purpose of standards.

If there were interoperability issues that adding requirements would address, the case would be different. However, the LIST IDENTIFIER field delivered by the device server to the application client in one command and returned unmodified by the application client to the same device server that provided the original value in another command. In such a usage scenario, there can be no interoperability concerns. The contents of the LIST IDENTIFIER field belong to the device server and no standardization value is served by constraining the usage of the LIST IDENTIFIER field.

Seagate 32) Explain native command usage in PERFORM SCSI COMMAND definition (Rejected) [462]

Seagate (RiedelE) #32
Page 90

To increase clarity, suggest that there be a footnote in this table that marks INQUIRY, REPORT LUNS, REQUEST SENSE, and TEST UNIT READY and reads "these commands may be issued as normal SCSI commands at all security levels. If further security protection is desired, then PERFORM SCSI COMMAND can be used to issue them as higher security. See Table 39. All other commands in this table may be issued only via PERFORM SCSI COMMAND when operating at higher security levels."

Reason for Rejection: The requested information appears in table 39, subclause 6.1. Duplicating requirements in more than one place leads to conflicting requirements and disparate implementations.

Seagate 33) Removing a non-existent collection should not be an error (Accepted, Substantive) [463]

Seagate (RiedelE) #33
Page 96

Why doesn't this text about non-existent object errors occur on page 95 for plain REMOVE and plain objects?

Editor's Note: An excellent point. Why should the device server be required to debug application client software by reporting an error when a non-existent collection is removed when no such requirement exists for user objects? Successful completion of a remove command indicates that the identified object (user or collection) is no longer present in the OBSD. What is the difference if the identified object was never there in the first place? What is an application client going to do differently if it is told that the object it tried to remove was not there?

The cited text will be modified as follows:

The contents of the COLLECTION_OBJECT_ID field specify the Collection_Object_ID (see 4.6.6) the collection to be removed. ~~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.~~

N.B. This change does not prohibit device servers from returning errors for commands that attempt to remove non-existent objects. It simply withdraws the requirement that they do so for collections.

Seagate 34) CREATE AND WRITE [s/b] WRITE & FrameMaker cross-ref bug (Accepted, Editorial) [464]

Seagate (RiedelE) #34
Page 103

this should be just WRITE not CREATE AND WRITE, ditto next paragraph. Also, the funny cut & paste from The COLLECTIONS... happens here as well - mentioned in a previous comment.

Editor's Note: CREATE AND WRITE will be changed WRITE.

The other problem is a technical gaff in the FrameMaker source file. The cross reference marker used to generate the text the User Object Quotas attributes page subclause cross reference for the maximum user object length attribute description is located on the wrong paragraph. The correct location for the cross reference marker is on the header. The cross reference marker will be moved and the cross references will be regenerated to correct the problem.

Seagate 35) Attribute number 0h exceptions list should include Collection Directory (Accepted, Editorial)
[465]

Seagate (RiedelE) #35
Page 106

should include "and Collection Directory"

Editor's Note: The cited text will be modified as follows:

With the exception of the Root Directory, ~~and~~ Partition Directory, ~~and~~ Collection Directory attributes pages, all attributes pages defined by this standard shall contain an identification of the page in attribute number 0h.

Seagate 36) Add "lists format only" note (Rejected) [466]

Seagate (RiedelE) #36
Page 107
See also: comment HP 85) and comment Veritas 108)

Perhaps add a note here that this information can only be retrieved with a list format get attributes.

Reason for Rejection: Actually, page format can be attempted with any page and those with no page format defined will return the null attributes page (see the last sentence in 7.1.1). A better way to provide the requested information is to add a column to table 67 (the list of attributes pages) as proposed and accepted for comment HP 85).

Seagate 37) Partition Directory lists only partition attributes pages (Accepted, Substantive) [467]

Seagate (RiedelE) #37
Page 108

The first sentence of 7.1.2.5 (Partition Directory attributes page) should be modified as follows:

The Partition Directory attributes page (P+0h) shall contain one attribute for every partition attributes page number accessible to the partition ~~or any user object within the partition.~~

Seagate 38) Logical unit device identifier attribute disappeared from Root Information attributes page (No Action Taken) [468]

Seagate (RiedelE) #38
Page 111

"logical unit id" is missing from this list - was present in Rev 8

Editor's Note: Because of the changes in the definition for the OSD System ID attribute, the r09 OSD System ID attribute contains exactly the same information as the r08 logical unit device identifier attribute. The only difference is the format in which the information is presented and one additional identifier format option allowed by the r09 OSD System ID attribute definition. Having two attributes that contain basically the same information seems unnecessarily redundant.

Seagate 39) Copy new partition username from partition 0, not from the root (Accepted, Substantive) [469]

Seagate (RiedelE) #39
Page 113

This should be the "username" attribute from partition 0, rather than the OSD name.

Editor's Note: The cited text will be modified as follows:

A CREATE PARTITION command (see 6.5) shall copy the ~~OSD name attribute from the Root Information attributes page (see 7.1.2.8)~~ username attribute from the Partition Information attributes page for partition zero to the new Partition Information attributes page.

Seagate 40) "Default" not used consistently (Rejected) [470]

Seagate (RiedelE) #40
Page 116

The term "Default" in these type of attributes is used inconsistently. We believe that all attributes that are "copied into" other attributes on creation/whatnot (e.g. Partition object count in this list) are prefaced with the term "Default"

Reason for Rejection: By the proposed reasoning, the change accepted for comment Seagate 39) would result in the 'username' attribute having to be named the 'default username' attribute for the Partition Information page for partition zero but just the 'username' attribute in all other pages. In short, the use of "default" cannot be applied consistently.

Seagate 41) "Quota" not used consistently (Rejected) [471]

Seagate (RiedelE) #41
Page 116

The term "quota" should be used consistently - e.g. in this list "Partition count" should be "Partition count quota" just like "Partition capacity quota"

Reason for Rejection: If the attribute appears in a quotas attributes page, then the attribute is a quota. Since all references to attributes include both the attribute and the attribute page, there will never be any lack of clarity.

The partition capacity quota is the exception was made because of potential confusion with non-quota attributes such as used capacity.

It is unacceptable to use the exception to justify an unnecessary rule.

Seagate 42) "Quota" not used consistently (Rejected) [472]

Seagate (RiedelE) #42
Page 118

consistent use of "quota" - suggest "Capacity quota" and "Object count quota"

Reason for Rejection: See response to comment Seagate 42).

Seagate 43) FrameMaker cross reference marker not on the correct paragraph (Accepted, Editorial) [473]

Seagate (RiedelE) #43
Page 118

the proverbial cut & paste problem starting at "The COLLECTION..."

Editor's Note: The problem is a technical gaff in the FrameMaker source file. The cross reference marker used to generate the text the User Object Quotas attributes page subclause cross reference for the maximum user object length attribute description is located on the wrong paragraph. The correct location for the cross reference marker is on the header. The cross reference marker will be moved and the cross references will be regenerated to correct this problem.

Seagate 44) Require Zero User Object Length in Partition Zero (Accepted, Substantive) [474]

Seagate (IrenS) #44
Page 118

Add the following statement: "For partition zero, the default maximum user object length attribute will be zero".

Editor's Note: The following will be inserted at the end of the cited paragraph:

For partition zero, the value for the default maximum user object length attribute shall be zero.

Seagate 45) Require Zero Collections Per User Object in Partition Zero (Accepted, Substantive) [475]

Seagate (IrenS) #45
Page 118

Add the following statement: "For partition zero, the collections per user object attribute will be set to zero."

Editor's Note: The following will be inserted at the end of the cited paragraph:

For partition zero, the value for the collections per user object attribute shall be zero.

Seagate 46) Require Zero Object Count in Partition Zero (Accepted, Substantive) [476]

Seagate (IrenS) #46
Page 118

Add the following statement: "For partition zero, the object count attribute will be set to zero."

Editor's Note: The following will be inserted at the end of the cited paragraph:

For partition zero, the value for the object count attribute shall be zero.

Seagate 47) the allow [s/b] allowed (Accepted, Editorial) [477]

Seagate (IrenS) #47
Page 120
See also comment IBM 114)

change "the allow" to "allowed"

Editor's Note: The cited text will be modified as follows:

The maximum user object length attribute (number 1h) specifies the maximum value ~~the allow~~ allowed in the user object logical length attribute of the User Object Information attributes page (see 7.1.2.11).

Seagate 48) shall [s/b] shall be (Accepted, Editorial) [478]

Seagate (IrenS) #48
Page 121, table 84

change "shall update" to "shall be updated"

Editor's Note: The cited text will be modified as follows:

Timestamps shall **be** updated as described in the subclause that defines them

Seagate 49) the each [s/b] each (Accepted, Editorial) [479]

Seagate (IrenS) #49
Page 147

change "the each" to "each"

Editor's Note: The cited text will be modified as follows:

The attributes list entry format shown in table 113 is used for returning ~~the~~ each attribute value for each user object requested by a CREATE command (see 6.3) that creates more than one user object.

Seagate 50) SET ATTRIBUTES service action is 880Fh (Accepted, Editorial) [480]

Page 98, table 62

The diagrams of the CDBs for GET ATTRIBUTES and SET ATTRIBUTES both show the service action as 880Eh. The SET ATTRIBUTES should be 880Fh.

Seagate 51) LIST COLLECTION listing partition contents is a partition function (Accepted, Substantive)
 [481]

Page 39, 4.9.4.4 (Credentials and commands allowed), table 19

Modify table 19 as follows:

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 LIST COLLECTION command with the CDB PARTITION_ID field containing a value that matches the contents of the credential PARTITION_ID field and the CDB COLLECTION_OBJECT_ID field containing a value other than zero that matches the contents of the object descriptor SINGLE OBJECT_ID field. If the object descriptor SINGLE OBJECT_ID field contains zero, a LIST COLLECTION command is allowed to list the collections defined in the partition.	COLLECTION	READ	1OBJECT
A LIST COLLECTION command with the CDB PARTITION_ID field containing a value that matches the object descriptor SINGLE OBJECT_ID field and the CDB COLLECTION_OBJECT_ID field equal to zero.	PARTITION	READ	1OBJECT
...			

13. Veritas Software

Roger Cummings from Veritas Software submitted the following comments on a No vote.

Veritas 1) Peer-to-peer does not align with common usage (Accepted, Editorial) [482]

PDF pg 3, pg iii, Abstract

"peer-to-peer" is not used throughout the document, and is only used in a limited sense in SAM-3. It adds no value, and does not align with common usage.

Proposed Resolution:

Delete "peer-to-peer"

Veritas 2) Peer-to-peer does not align with common usage (Accepted, Editorial) [483]

PDF pg 17, pg xvii, Foreword

"peer-to-peer" is not used throughout the document, and is only used in a limited sense in SAM-3. It adds no value, and does not align with common usage.

Proposed Resolution:

Delete "peer-to-peer"

Veritas 3) Align attributes definition with OSD command names (Rejected) [484]

PDF pg 25, pg 6, 3.1.3 Attributes

Section 3.1.3 defines attributes as data "not accessible via read or write command functions". Since any command (including READ and WRITE) can potentially access and change attributes, is this true?

Proposed Resolution:

Define attributes as data associated with an object that 1) requires special capabilities to access or change, and 2) is addressed by attribute number instead of offset (like user data).

Reason for Rejection: Examine the cited glossary entry carefully:

3.1.3 attributes: Data (sometimes called meta data) that is associated with an OSD object (see 3.1.27) that is not accessible via [read or write command functions \(see 3.1.10\)](#). See 4.7.

There is no discussion of commands (e.g., READ and WRITE). There is only discussion of command functions (i.e., read and write) and there is a cross reference to the glossary entry that explains the distinction, to wit:

3.1.10 command function: One unit of work within a single command (see 3.1.8). [This standard extends the SAM-3 definition of command to allow multiple command functions to be requested by a single command.](#)

Every OSD command function requested by an application client requires a special capability. Attributes are not distinguished by that property. The addressing of both attributes and user data requires too explanation to be useful in a glossary entry.

The reference to read and write command functions most appropriately describes what differentiates attributes from other data. Every effort has been made to use and extend normal SCSI terminology in a clear and consistent way to describe attributes.

Veritas 4) from single [s/b] from a single (Accepted, Editorial) [485]

PDF pg 25, pg 6, 3.1.6 collection

from single

Proposed Resolution:

from a single

Veritas 5) object-based storage device definition is too redundant (Rejected) [486]

PDF pg 26, pg 7, 3.1.26 object-based storage device (OBSD)

"A SCSI device that implements this standard in which data is organized and accessed as objects." has much redundancy.

Proposed Resolution:

"A SCSI device that implements this standard."

Reason for Rejection: Hammering the point home seems worthwhile and no other reviewers complained about the cited text.

Veritas 6) lost [s/b] loss (Accepted, Editorial) [487]

PDF pg 27, pg 8, 3.1.44 stable storage

lost

Proposed Resolution:

loss

Veritas 7) greenwich mean time [s/b] Greenwich Mean Time (Accepted, Editorial) [488]

PDF pg 27, pg 8, 3.1.47 universal time (UT)

greenwich mean time

Proposed Resolution:

Greenwich Mean Time

Veritas 8) Two types in Volatile Cache definition (Accepted, Editorial) [489]

PDF pg 27, pg 8, 3.1.51 Volatile Cache

Storage Is lost after a power on or rest event

Proposed Resolution:

Storage whose contents are lost after a power on or reset event

Editor's Note: 'Storage is lost' will be changed to 'Storage that is lost' (see comment Lingua 10). 'rest' will be changed to 'reset'.

Veritas 9) Devices [s/b] Device (Accepted, Editorial) [490]

PDF pg 28, pg 9, 3.2 Acronyms

OSD Object-based Storage Devices Commands (this standard, see clause 1)

Proposed Resolution:

OSD Object-based Storage Device Commands (this standard, see clause 1)

Veritas 10) to encoded [s/b] to be encoded (Accepted, Editorial) [491]

PDF pg 30, pg 11, 3.6.1 Notation for byte encoded character strings, 1st para, 3rd line

See also Lingua 11)

to encoded

Proposed Resolution:

to be encoded

Editor's Note: A complete description of all the changes in the cited paragraph can be found in the response to comment Lingua 11).

Veritas 11) same writing [s/b] same as writing (Accepted, Editorial) [492]

PDF pg 30, pg 11, 3.6.1 Notation for byte encoded character strings, 3st para, 2nd line

same writing

Proposed Resolution:

same as writing

Veritas 12) Use UTF-8 instead ASCII (Rejected) [493]

PDF pg 30, pg 11, 3.6.1 Notation for byte encoded character strings

ASCII is not listed in acronyms or references. Also, is it not current policy to reference UTF-8 instead?

Proposed Resolution:

Correct reference

Reason for Rejection: The editor is not aware of any T10 policy requiring UTF-8 instead of ASCII. No request to use UTF-8 has been received from the SNIA OSD Technical Working Group.

Veritas 13) Encryption (Privacy) Of Credential Transmission To Application Clients (Rejected) [494]

PDF pg 35, pg 16, 4.4 Elements of the example configuration
see comment EMC 1)

Section 4.4 states "When sending credentials to an application client, the security manager shall use a private, authenticated communications mechanism. The security manager may reside in the OBSD, in applications clients, or as a separate entity, but the security requirements on the communications mechanism shall not change based on the location of the security manager."

Communications between the security manager and clients are outside the scope of this standard, so it doesn't make sense to place mandatory requirements on it.

Proposed Resolution:

Change "shall" to "should".

Reason for Rejection: The security of Credential usage requires secret transmission of Credentials between the Security Manager and the Application Client (see comment EMC 1).

Veritas 14) Define what is outside the scope of this standard (Rejected) [495]

PDF pg 35, pg 16, 4.4 Elements of the example configuration

This section should indicate the parts of the example configuration that are outside of the scope of the document.

Proposed Resolution:

Add two sentences in appropriate places. "The detailed operation of the Policy/Storage Manager and its interface to the service delivery subsystem is outside of the scope of this standard." "Definition of the interfaces between the Security manager and an Initiator Device, and between the Security Manager and an OBSD, are outside of the scope of this standard."

Reason for Rejection: What is outside the scope of this standard is subject to incorrectly narrow misinterpretation (see comment Veritas 13). What is outside the scope of this standard is constantly being enlarged (e.g., the changes made in the resolution for comment IBM 141) give the policy/storage manager responsibility for constructing capabilities and managing the object version tag). Attempting to define what is outside the scope of the standard is more likely to produce conflicts between definitions in the standard than it is likely to produce enlightenment for readers.

Veritas 15) or as [s/b] or exist as (Rejected) [496]

PDF pg 35, pg 16, 4.4 Elements of the example configuration, 3rd para under figure 3

or as

Proposed Resolution:

or exist as

Reason for Rejection: The current wording is acceptable English, reading as 'The security manager may reside ... as a separate entity', to whit:

The security manager may reside in the OBSD, in applications clients, or **as a separate entity**, but the security requirements on the communications mechanism shall not change based on the location of the security manager.

Veritas 16) "Abstracted subsets" not clear (Accepted, Editorial) [497]

PDF pg 36, pg 17, 4.5 Description of the OSD Architecture

"Abstracted subsets" is unclear and does not reflect the terminology used elsewhere in the document.

Proposed Resolution:

"Data is stored in abstract containers called objects that contain operating system and application constructs."

Editor's Note: The cited text will be modified as follows:

Data is stored in ~~abstracted subsets~~ abstract containers by the OBSD (see 3.1.26). ~~Abstracted indicates that the data~~ Data in the abstract containers is not addressable using LBAs (Logical Block Addresses).

The remainder of the proposed changes exceed what is needed in this context.

Veritas 17) OSD can assign user object IDs (Rejected) [498]

PDF pg 36, pg 17, 4.5 Description of the OSD Architecture

The sentence "The OSD logical unit allocates space for data and delivers to the application client a unique identifier." is incorrect. The unique identifier may be given to the OSD LU at part of the CREATE command.

Proposed Resolution:

Change to something like: "The data is addressed using a unique identifier which can either be assigned by the OSD logical unit or be given to the OSD logical unit by the client. The OSD logical unit allocates space for the data."

Reason for Rejection: Application Clients cannot assign User_Object_IDs. They can only request that the logical unit assign a specified User_Object_ID.

The table footnote wording in Table 2 clearly indicates this, to wit:

"b — User_Object_ID values assigned by the OSD logical unit **in response to application client requests.**"

The logical unit is **always** the ID assignment authority. Any other architecture will require the definition of a locking mechanism to coordinate Application Client assignments of User_Object_IDs.

Veritas 18) An OSD contains [s/b] An OBSD contains (Accepted, Editorial) [499]

PDF pg 36, pg 17, 4.6.1 Stored data object types

"An OSD contains" should reference an OBSD?

Proposed Resolution:

"An OBSD contains"

Veritas 19) Clarify OSD logical unit (Accepted, Editorial) [500]

PDF pg 36, pg 17, 4.6.1 Stored data object types

Is "OSD Logical Unit" a synonym for OBSD? Can an OBSD have multiple LUs?

OSD Logical Unit is not defined in the glossary but is used throughout the document.

Proposed Resolution:

Either replace OSD Logical Unit with another term or add it to the glossary and explain the relationship to OBSD.

Editor's Note: The following changes will be made:

- The following glossary entry will be added: "**3.1.x OSD logical unit: A logical unit within an OBSD (see 3.1.26).**"
- Including the changes described in comment Veritas 18), the first sentence of 4.6.1 will be modified as follows: "An ~~OSD~~ **OBSD** contains **one or more logical units with** the following types of stored data objects:"

Veritas 20) Eliminate Root Object Data (Accepted, Editorial) [501]

PDF pg 36, pg 17, 4.6.1 Stored Data Object Types

In discussion of the Root Object, the sentence "Its data contains the list of Partition_IDs." appears. This is misleading since it implies that a client may issue a READ command to the Root Object. In fact, a READ command issued to the Root Object is an error.

Proposed Resolution:

Change to "The Root Object maintains the list of PARTITION_IDs contained on the logical unit. This list is accessed with the LIST command"

Editor's Note: The cited text will be modified as follows:

~~Its data contains the list of Partition_IDs.~~ The root object contains a list of Partition_IDs for the partitions in the logical unit that may be retrieved using the LIST command (see 6.10).

Veritas 21) Eliminate Partition Data (Accepted, Editorial) [502]

PDF pg 36, pg 17, 4.6.1 Stored Data Object Types

In discussion of Partition Objects, the sentence "The data component of a partition is the list of User_Object_IDs." appears. This is misleading since it implies that a client may issue a READ command to a Partition Object. In fact, a READ command issued to a Partition Object is an error.

Proposed Resolution:

Change to "Partition Objects maintain the list of USER_OBJECT_IDs contained in the partition. This list is accessed with the LIST command"

Editor's Note: The cited text will be modified as follows:

~~Its data contains the list of Partition_IDs.~~ Each partition contains a list of User_Object_IDs and Collection_Object_IDs contained in the partition that may be retrieved using the LIST command (see 6.10).

Veritas 22) Is the root object a member of partition 0? (No Action Taken) [503]

PDF pg 36, pg 17, 4.6.1 Stored data object Types

An OSD logical unit shall always contain a root object and an OSD object for partition zero with at least the attributes (see 4.7) defined by this standard.

Is the root object a member of partition 0? 3.1.31 suggests yes.

Editor's Note: As described in the response to comment IBM 20), partition zero represents the root object in those matters where a partition is needed.

Veritas 23) Eliminate Collection Data (Accepted, Editorial) [504]

PDF pg 36, pg 17, 4.6.1 Stored Data Object Types

In discussion of Collection Objects, the sentence "The data component of a partition is the list of User_Object_IDs." appears. This is misleading since it implies that a client may issue a READ command to a Collection Object.

In fact, a READ command issued to a Collection Object is an error. It is also wrong since it is supposed to be about Collection Object, not Partition Objects.

Proposed Resolution:

Change to "Collection Object maintains the list of USER_OBJECT_IDs contained in the collection. This list is accessed with the LIST command."

Editor's Note: The cited text will be modified as follows:

~~Its data contains the list of Partition_IDs.~~ Each collection contains a list of User_Object_IDs contained in the collection that may be retrieved using the LIST COLLECTION command (see 6.11).

Veritas 24) Multiple Object Operations in this standard (Accepted, Editorial) [505]

PDF pg 36, pg 17, 4.6.1 Stored Data Object Types

In the discussion of Collection Objects, the sentence "It is used for fast indexing of user objects and operations involving multiple user objects." appears. There currently are no operations that involve multiple objects.

Proposed Resolution:

Strike "and operations involving multiple user objects" from the sentence.

Veritas 25) Clients Can Assign Partition/User Object IDs (Rejected) [506]

PDF pg 37, pg 18 Identifying OSD Objects

see also: comment Veritas 27), comment Veritas 33), and comment Veritas 34)

The notes for Table 2 suggest that PARTITION_IDs and USER_OBJECT_ID's are always assigned by the OSD logical unit. This is not true. Clients can assign both PARTITION_IDs and USER_OBJECT_IDs.

Proposed Resolution

Change the notes on Table 2 to reflect the fact that PARTITION_IDs and USER_OBJECT_IDs may be assigned by the client.

Reason for Rejection: Application Clients cannot assign Partition_IDs or User_Object_IDs. They can only request that the logical unit assign a specified Partition_ID or User_Object_ID.

The table footnote wording in Table 2 clearly indicates this, to wit:

"a — Partition_ID values assigned by the OSD logical unit **in response to application client requests.**

"b — User_Object_ID values assigned by the OSD logical unit **in response to application client requests.**"

The logical unit is **always** the ID assignment authority. Any other architecture will require the definition of a locking mechanism to coordinate Application Client assignments of Partition_IDs or User_Object_IDs.

Veritas 26) Insert missing comma (Accepted, Editorial) [507]

PDF pg 37, pg 18, 4.6.2 Identifying OSD objects

object each

Proposed Resolution:

object, each

Veritas 27) Clients Can Assign Partition IDs (Rejected) [508]

PDF pg 37, pg 18, 4.6.4 Partitions

see also: comment Veritas 25)

The sentence "Partitions have a User_Object_ID of zero and a Partition_ID (see 4.6.2) that is assigned by the OSD logical unit when the user object is created." is not accurate. Partition_ID can be assigned by the client.

Proposed Resolution

Change sentence to reflect Client assigned Partition_IDs.

Reason for Rejection: See response to comment Veritas 25).

Veritas 28) partition OSD object [s/b] OSD partition object (Rejected) [509]

PDF pg 37, pg 18, 4.6.4 Partitions

partition OSD

Proposed Resolution:

OSD partition

Reason for Rejection: The OSD editing review meetings specifically requested the definition of the term 'OSD object' (see 3.1.27) so that in text such as the cited sentence OSD objects could be clearly differentiated from SAM-3 objects.

Veritas 29) user object [s/b] partition (Accepted, Editorial) [510]

PDF pg 37, pg 18, 4.6.4 Partitions, 2nd para

user object

Proposed Resolution:

partition

Veritas 30) partition OSD object [s/b] OSD partition object (Rejected) [511]

PDF pg 37, pg 18, 4.6.4 Partitions, 3rd para

partition OSD

Proposed Resolution:

an OSD partition

Reason for Rejection: The OSD editing review meetings specifically requested the definition of the term 'OSD object' (see 3.1.27) so that in text such as the cited sentence OSD objects could be clearly differentiated from SAM-3 objects.

Veritas 31) User objects contain Partition_IDs (Rejected) [512]

PDF pg 38, pg 19, 4.6.5 User objects

User objects have the Partition_ID of

Proposed Resolution:

User objects contain the Partition_ID of

Reason for Rejection: User objects contain user data, not a Partition_ID.

Veritas 32) Make user object membership in one partition an enforceable requirement (Rejected) [513]
PDF pg 38, pg 19, 4.6.5 User objects

A user object is a member of only one partition.

Proposed Resolution:

A user object shall be a member of only one partition.

Reason for Rejection: A user object's membership in just one partition is a consequence of the OSD model. There simply is no way defined to cause a user object to be a member of more than one partition. There is no requirement placed on OSDs regarding user object membership in a single partition. None is needed. It is just a fact of life, in keeping with the phrasing of the cited text.

Veritas 33) Clients Can Assign User Object IDs (Rejected) [514]
PDF pg 38, pg 19, 4.6.5 User Objects
see also: comment Veritas 25)

The sentence "User objects have the Partition_ID of the partition to which they belong and a User_Object_ID (see 4.6.2) that is assigned by the OSD logical unit when the user object is created." is not accurate. User_Object_IDs can be assigned by the client.

Proposed Resolution

Change sentence to reflect Client assigned User_Object_IDs.

Reason for Rejection: See response to comment Veritas 25).

Veritas 34) Clients Can Assign Collection_Object_IDs (Rejected) [515]
PDF pg 38, pg 19, 4.6.6 Collections
see also: comment Veritas 25)

The sentence "Collections have the Partition_ID of the partition to which they belong and a Collection_Object_ID (see 4.6.2) that is assigned by the OSD logical unit when the collection is created." is inaccurate. Collection_Object_IDs can be assigned by the client.

Proposed Resolution

Change sentence to reflect Client assigned Collection_Object_IDs

Reason for Rejection: See response to comment Veritas 25).

Veritas 35) SBC-based systems is unclear (Accepted, Editorial) [516]

PDF pg 39, pg 20, 4.7.1 Overview

"SBC-based systems" is unclear - does this mean to exclude SBC-2 based systems?

Proposed Resolution:

"systems that use random access to blocks"

Editor's Note: The cited sentence will be modified as follows:

File systems and other **systems based on the traditional storage model (see 4.3) ~~SBC-based systems~~** store both user data and meta data.

Veritas 36) All OSD object types is confusing (Accepted, Editorial) [517]

PDF pg 39, pg 20 4.7.1 Overview

The sentence "With the exception of attributes pages in the attributes page number range assigned to all OSD object types (see table 3 in 4.7.3), the same attributes page shall not be associated with more than one OSD object." doesn't seem right. Is this a typo? Should it be "more than one OSD object type"?

Proposed resolution

Fix as needed.

Editor's Note: This comment will be resolved as described in the response to comment LSI 9) and in addition "more than one OSD object" will be changed to "more than one OSD object **type**".

Veritas 37) store [s/b] stored (Accepted, Editorial) [518]

PDF pg 39, pg 20, 4.7.1 Overview, 3rd para

store

Proposed Resolution:

stored

Veritas 38) Add Current Command attributes page to overview (Rejected) [519]

PDF pg 39, pg 20 4.7.1 Overview

The special nature of the current command page (i.e. it can always be retrieved) should be noted here.

Proposed Resolution

Call out the special nature of the current command page.

Reason for Rejection: The special nature of the Current Command attributes page is covered in the overview subclause because the Current Command attributes page is one of the pages that are associated with any OSD object type. Since future versions of OSD might define other attributes pages that are associated with any OSD object type in addition to the Current Command attributes page, it is not appropriate to specifically name the Current Command attributes page in the overview.

Veritas 39) an command [s/b] a command (No Action Taken) [520]

PDF pg 39, pg 20, 4.7.2 Command function ordering for commands that get and/or set attributes

an command

Proposed Resolution:

a command

Reason for Rejection: A search of the OSD r09 PDF found no instances of 'an command' to change.

Veritas 40) command functions [s/b] command function (Rejected) [521]

PDF pg 39, pg 20, 4.7.2 Command function ordering for commands that get and/or set attributes

the command functions

Proposed Resolution:

command function

Reason for Rejection: It is not in the long-term interests of OSD to limit a command to a single command function.

Veritas 41) GET ATTRIBUTES Processing Order (Rejected) [522]

PDF pg 40, pg 21, 4.7.2 processing of a GET ATTRIBUTES command...

2) Process any get attributes command functions specified in the CDB; and

Proposed Resolution:

(Question) I do not understand why getting attributes happens prior to the set attributes in this case alone. I think it should be consistent with the other two cases (below) and the general case (above).

Reason for Rejection: The processing order for GET ATTRIBUTES is defined to give a distinct purpose to having both a GET ATTRIBUTES command and a SET ATTRIBUTES command. Providing a test-and-set functionality may prove valuable to specific application client implementations.

Veritas 42) Publicly available vendor specific attribute page definitions enhance this standard (Accepted, Editorial) [523]

PDF pg 41, pg 22, 4.7.3 Attributes pages

The distinction between "Defined by OBSD (see 3.1.26) manufacturer product specifications" and "Vendor specific" in Table 4 is unclear. The former sure sounds vendor specific to me - is the difference that those pages are read-only?

Editor's Note: Any publicly available documentation of attributes pages enhances the value of this standard. Encouraging such publication of attributes pages definitions is the purpose of the cited attribute page number value range.

The cited text will be changed to:

Defined by **publicly available** ~~OBSD (see 3.1.26)~~ manufacturer ~~product~~ specifications

Veritas 43) insert 'in order' (Accepted, Editorial) [524]

PDF pg 42, pg 23, 4.7.5 Attributes directories, 2nd para under table

in attribute number 0h to make

Proposed Resolution:

to attribute number 0h in order to make

Veritas 44) FrameMaker cross reference marker not on the correct paragraph (Accepted, Editorial) [525]

PDF pg 42, pg 23,

Reference on item c) is incomplete

Editor's Note: The problem is a technical gaff in the FrameMaker source file. The cross reference marker used to generate the text in entry c of the a, b, c list is located on the paragraph before the header for the User Object Quotas attributes page subclause. The correct location for the cross reference marker is on the header. The cross reference marker will be moved and the cross references will be regenerated to correct this problem.

Veritas 45) OSD Device [s/b] OBSD (Accepted, Editorial) [526]

PDF pg 44, pg 25, 4.9.1 Basic security model

In Figure 4 "OSD Device" should be OBSD.

Veritas 46) meaning that [s/b] consequently (Rejected) [527]

PDF pg 44, pg 25, 4.9.1 Basic security model, next to last paragraph on page

value, meaning that application

Proposed Resolution:

value. Consequently, the application

Reason for Rejection: There are Letter Ballot reviewers whose distaste for 'ly' words is such that they would request the opposite.

Veritas 47) does allow [s/b] allows (Accepted, Editorial) [528]

PDF pg 44, pg 25, 4.9.1 Basic security model, next to last paragraph on page

does allow delegation of a credential if a application client

Proposed Resolution:

allows delegation of a credential in an application client

Veritas 48) Add a sentence break (Rejected) [529]

PDF pg 44, pg 25, 4.9.1 Basic security model, last paragraph on page

perform and sends those capabilities in those credentials

Proposed Resolution:

perform. It sends those capabilities and those credentials

Reason for Rejection: The proposed change introduces an 'it' that might reference either the application client or the security manager. The grammatical structure of the cited single sentence does not introduce such confusion.

Veritas 49) Application Client Trust Assumption Rewording (Accepted, Editorial) [530]

PDF pg 44, pg 25, 4.9.1 Basic security model, last paragraph on page

While the application client is not trusted to follow this protocol, the protocol is structured in a way that makes it in the application client's self-interest to follow the protocol. An application client that does not follow the protocol is unlikely to receive service from the OSD device server.

Proposed Resolution:

While the application client is not trusted to follow this protocol, an application client that does not follow the protocol is unlikely to receive service from the OSD device server.

Veritas 50) send [s/b] sent (Accepted, Editorial) [531]

PDF pg 45, pg 26, 4.9.1 Basic security model, first paragraph on page

send by

Proposed Resolution:

sent by

Veritas 51) and the [s/b] and by the (Accepted, Editorial) [532]

PDF pg 45, pg 26, 4.9.2 Trust assumptions, 4th paragraph

and the

Proposed Resolution:

and by the

Veritas 52) security mode [s/b] security method (Accepted, Editorial) [533]

PDF pg 45, pg 26, 4.9.2 Trust assumptions, 2nd paragraph under table

mode other

Proposed Resolution:

method other

Editor's Note: This comment would have been accepted, except that the response to comment Seagate 7) specifies that the entire cited paragraph be removed.

Veritas 53) Incorrect cross reference (Accepted, Editorial) [534]

PDF pg 46, pg 27, 4.9.3.2 The NOSEC security method

PERFORM SCSI COMMAND command (see 6.11)

Proposed Resolution:

PERFORM SCSI COMMAND command (see 6.12)

Veritas 54) should [s/b] could & add cross reference (Rejected) [535]

PDF pg 47, pg 28, 4.9.3.3 The CAPKEY security method

The credential should appear to have been tampered with when this is not truly the case. The problems caused by multiple SCSI domains may be solved by using a different security method (e.g., the CMDRSP security method).

Proposed Resolution:

The credential could appear to have been tampered with when this is not actually the case. The problems caused by multiple SCSI domains can be solved by using a different security method (e.g., the CMDRSP security method, see 4.9.3.4).

Reason for Rejection: The word 'could' is not allowed in T10 standards (see 7.1.1 in T10/01-313r2). The alternative 'may' was debated in the SNIA OSD TWG and dismissed as too weak. The word 'shall' is too strong because there is no way to establish the certainty necessary to use 'shall'. The word 'should' was/is the only remaining option.

The requested cross reference is to a subclause that begins one paragraph subsequent to the location where the reference has been requested.

Veritas 55) Why denial of service? (No Action Taken) [536]

PDF pg 48, pg 29, 4.9.3.4 The CMDRSP security method, 1st paragraph under list.

presence of malicious entities perpetrating a denial of service attack should be considered.

(question) why this set of circumstances suggest a denial of service attack?

Editor's Note: The malicious entities may be altering the response information in ways that guarantee failures in the response integrity check value validation, thus denying service.

Veritas 56) Why denial of service? (No Action Taken) [537]

PDF pg 51, pg 32, 4.9.3.5 The ALLDATA security method, first paragraph at top of page

presence of malicious entities perpetrating a denial of service attack should be considered.

(question) why this set of circumstances suggest a denial of service attack?

Editor's Note: The malicious entities may be altering the response information in ways that guarantee failures in the response integrity check value validation, thus denying service.

Veritas 57) to [s/b] into (Accepted, Editorial) [538]

PDF pg 52, pg 33, 4.9.4.2 Capability key, first paragraph at top of page

copied to the CDB.

Proposed Resolution:

copied into the CDB.

Veritas 58) as follows [s/b] by (Accepted, Editorial) [539]

PDF pg 52, pg 33, 4.9.4.2 Capability key, first paragraph at top of page

1) Reconstructing the credential containing

Proposed Resolution:

1) Reconstruct the credential containing

Editor's Note: Since there is already a verb in the sentence, this comment will be resolved by modifying the cited text as follows:

Since the capability does not include the CREDENTIAL INTEGRITY CHECK VALUE field, the device server needs to compute the capability key for each processed command **by as follows:**

- 1) Reconstructing the credential containing the CDB capability as described in 4.9.5.3; and
- 2) Computing the credential integrity check value for the reconstructed credential using the algorithm, inputs, and secret key specified in 4.9.5.4.

Veritas 59) as follows [s/b] by (Accepted, Editorial) [540]

PDF pg 52, pg 33, 4.9.4.2 Capability key, first paragraph at top of page

2) Computing the credential integrity check

Proposed Resolution:

2) Compute the credential integrity check

Editor's Note: This comment will be resolved as shown in the response to comment Veritas 58)

Veritas 60) attributes [s/b] attribute (Rejected) [541]

PDF pg 53, pg 34, 4.9.4.3 Capability format, 2nd paragraph under table

supported integrity check value algorithm attributes in the Root Security

Proposed Resolution:

supported integrity check value algorithm attribute in the Root Security

Reason for Rejection: There are 16 supported integrity check value algorithm attributes (numbers 8000 0000h to 8000 000Fh) in the Root Security attributes page (see table 94 in 7.1.2.20).

Veritas 61) of OSD object [s/b] of the OSD object (Accepted, Editorial) [542]

PDF pg 55, pg 36, 4.9.4.3 Capability format, 3rd paragraph from bottom of page

of OSD object

Proposed Resolution:

of the OSD object

Veritas 62) Add a sentence break in place of 'and' (Accepted, Editorial) [543]

PDF pg 56, pg 37, 4.9.4.4 Credentials and commands allowed

attributes and combinations of capability fields

Proposed Resolution:

attributes. Combinations of capability fields

Editor's Note: The change proposed by this comment is included in the response to comment Lingua 27).

Veritas 63) three editorial changes in one sentence (Accepted, Editorial) [544]

PDF pg 56, pg 37, 4.9.4.4 Credentials and commands allowed

the retrieval and setting of attributes by combining the permission bits values described

Proposed Resolution:

retrieving and setting attributes by combining permission bit values described

Editor's Note: The cited sentence will be modified as follows:

A single credential for a single object type may allow processing of multiple command functions (e.g., read and write) as well as ~~the retrieval~~ retrieving and setting ~~of~~ attributes by combining the permission bits values described in multiple rows of table 19 and table 20.

All requested changes are included, except the removal of the word 'the'.

Veritas 64) FORMAT OSD should require GLOBAL permission? (Accepted, Substantive) [545]

PDF pg 58, pg 39, Table 19, 3rd row

both equal to zero or a FORMAT OSD command.

Proposed Resolution:

both equal to zero. Also a FORMAT OSD command. (question) Why doesn't FORMAT OSD require global permission?

Editor's Note: The FORMAT OSD command will be changed to required GLOBAL permission. This requires a totally new row for the FORMAT OSD command, overtaking the need to add the suggested sentence break. See 04-189.

Veritas 65) Capability Validation is NOT a Security Function (Accepted, Substantive) [546]

PDF pg 64, pg 45, 4.9.5.2 Credential and capability validation

A NOSEC target should do everything listed in this section except for validating the integrity check values.

Proposed resolution

"This subclause describes the process for validating a capability received in a CDB. If the security method is NOSEC, the integrity check defined herein shall not be performed.

Editor's Note: Oh, would that the changes were this easy. This comment will be resolved as described in the response to comment IBM 141).

Veritas 66) is [s/b] in (Accepted, Editorial) [547]

PDF pg 64, pg 45, 4.9.5.2 Credential and capability validation, 1st para

method is use

Proposed Resolution:

method in use

Veritas 67) is [s/b] in (Accepted, Editorial) [548]

PDF pg 64, pg 45, 4.9.5.2 Credential and capability validation, 3rd para

method is use

Proposed Resolution:

method in use

Veritas 68) An integrity [s/b] Integrity (Accepted, Editorial) [549]

PDF pg 68, pg 49, 4.9.7 Integrity check values

key. An integrity

Proposed Resolution:

key. Integrity

Veritas 69) is has [s/b] has (Accepted, Editorial) [550]

PDF pg 69, pg 50, 4.9.7 Integrity check values

check value is has

Proposed Resolution:

check value has

Editor's Note: This comment would have been accepted, except that comment Seagate 18) specifies that the entire cited paragraph be removed.

Veritas 70) Remove redundant 'defining the' (Accepted, Editorial) [551]

PDF pg 69, pg 50, 4.9.8.1 Introduction

c) The definition of the defining the SET KEY

Proposed Resolution:

c) The definition of the SET KEY

Veritas 71) As [s/b] As a (Accepted, Editorial) [552]

PDF pg 69, pg 50, Table 22 OSD secret key hierarchy

d As dual purpose number,

Proposed Resolution:

d As a dual purpose number,

Veritas 72) insert 'a' (Accepted, Editorial) [553]

PDF pg 70, pg 51, 4.9.8.1 Introduction

may be protected by tamper resistant software shield.

Proposed Resolution:

may be protected by a tamper resistant software shield.

Veritas 73) rest [s/b] reset (Accepted, Editorial) [554]

PDF pg 72, pg 53, 4.10 Data persistence model

rest event (see SAM-3)

Proposed Resolution

rest event (see SAM-3)

Veritas 74) lost [s/b] loss (Accepted, Editorial) [555]

PDF pg 72, pg 53, 4.10 Data persistence model

the lost of data

Proposed Resolution:

the loss of data

Veritas 75) is [s/b] of (Accepted, Editorial) [556]

PDF pg 72, pg 53, 4.10 Data persistence model

combination is non-volatile random

Proposed Resolution:

combination of non-volatile random

Veritas 76) Prohibiting changes in unused Data-In Buffer bytes (Accepted, Substantive) [557]

PDF pg 74, pg 74, 4.11.2 OSD Data-In Buffers Format
See also comment Seagate 22)

The sentence "The device server shall not alter unused bytes in the Data-In Buffer" doesn't make sense. The device server has to send these bytes to the initiator and the bytes therefore will be modified in the initiator buffers.

Proposed Resolution

Replace offending sentence with "The contents of the unused bytes in the Data-In Buffer is undefined".

Editor's Note: This comment will be resolved as described in the response to comment Seagate 22).

Veritas 77) zero [s/b] NOSEC (Accepted, Editorial) [558]

PDF pg 78, pg 59, 4.15 Reservations

OSD logical unit is not zero, the PERSISTENT RESERVE IN

Proposed Resolution:

OSD logical unit is not NOSEC, the PERSISTENT RESERVE IN

Veritas 78) containing service action [s/b] containing a service action (Accepted, Editorial) [559]

PDF pg 80, pg 61, 5.1 OSD CDB format

a CDB containing service action

Proposed Resolution:

a CDB containing a service action

Veritas 79) Remove redundant 'there is' (Accepted, Editorial) [560]

PDF pg 83, pg 64, 5.2.1.3 Get and set attributes lists

zero specifies that there is no get attributes list is

Proposed Resolution:

zero specifies that no get attributes list is

Veritas 80) When both DPO and FUA are set (Rejected) [561]

PDF pg 84, pg 65, 5.2.3 Options byte

Is it necessary to add a paragraph explaining behavior if both DPO and FUA are set?

Proposed Resolution:

If both DPO and FUA each have a value of one, then the device server should place no data in the volatile cache and the device server shall not return status until the data transferred by this command has been written to stable storage.

Reason for Rejection: The proposed text additions does not say anything that is not already stated in the cited subclause. It simply restates and combines the requirements already stated for behavior when the DPO bit is set to one and when the FUA bit is set to one.

Veritas 81) PARTITION_ID field and CREATE PARTITION (Rejected) [562]

PDF pg 85, pg 66, 5.2.4 Partition_ID

"If the partition identified by the PARTITION_ID field does not exist, the command shall be terminated ...". What about CREATE_PARTITION with PARTITION_ID != 0?

Proposed Resolution

Call out exception for CREATE_PARTITION.

Reason for Rejection: The field name in the CREATE PARTITION command is REQUESTED_PARTITION_ID, not PARTITION_ID. The cited subclause is not referenced by the CREATE PARTITION command.

Veritas 82) updated or not updated [s/b] established (Accepted, Editorial) [563]

PDF pg 86, pg 67, 5.2.7 Timestamps control

previously updated or not updated timestamp attribute

Proposed Resolution:

previously established timestamp attribute

Veritas 83) USER_OBJECT_ID field and CREATE and CREATE AND WRITE (Rejected) [564]

PDF pg 87, pg 68, 5.2.8 User_Object_ID

"If the user object identified by the USER_OBJECT_ID field does not exist, the command shall be terminated ...". What about CREATE & CREATE_AND_WRITE with User_Object_ID != 0?

Proposed Resolution

Call out exception for CREATE & CREATE_AND_WRITE.

Reason for Rejection: The field name in the CREATE command and the CREATE AND WRITE command is REQUESTED_USER_OBJECT_ID, not USER_OBJECT_ID. The cited subclause is not referenced by the CREATE command or the CREATE AND WRITE command.

Veritas 84) Persistent Reservations Should Be Optional (Accepted, Substantive) [565]

PDF pg 87, pg 68, Table 39 Commands for OSD type devices (part 1 of 2)
see also: comment ENDL 3)

PRESISTENT RESERVE IN and PRESISTENT RESERVE OUT should be optional commands, because capability checking other than the integrity check computation should happen for NOSEC

Proposed resolution

Modify table to declare PERSISTENT RESERVE IN and PRESISTENT RESERVE OUT are optional.

Editor's Note: This comment will be resolved as described in the response to comment ENDL 3).

Veritas 85) Give that sentence a verb (Accepted, Editorial) [566]

PDF pg 87, pg 68, Table 39 Commands for OSD type devices (part 1 of 2)

8800h and 8F7Fh that not listed

Proposed Resolution:

8800h and 8F7Fh that are not listed

Veritas 86) Give that sentence a verb (Accepted, Editorial) [567]

PDF pg 88, pg 69, Table 39 Commands for OSD type devices (part 2 of 2)

8800h and 8F7Fh that not listed

Proposed Resolution:

8800h and 8F7Fh that are not listed

Veritas 87) FrameMaker cross reference marker not on the correct paragraph (Accepted, Editorial) [568]

PDF pg 90, pg 71, 6.2 APPEND

maximum user object length attribute in the The COLLECTIONS PER USER OBJECT field contains the value of the collections per user object attribute. (see), then a

Proposed Resolution:

maximum user object length attribute (see 7.1.2.14), then a

Editor's Note: The problem is a technical gaff in the FrameMaker source file. The cross reference marker used to generate the text the User Object Quotas attributes page subclause cross reference for the maximum user object length attribute description is located on the wrong paragraph. The correct location for the cross reference marker is on the header. The cross reference marker will be moved and the cross references will be regenerated to correct this problem.

Veritas 88) Insert a comma (Accepted, Editorial) [569]

PDF pg 93, pg 74, 6.4 CREATE AND WRITE

If the REQUESTED USER_OBJECT_ID field contains zero any User_Object_ID may be assigned.

Proposed Resolution:

If the REQUESTED USER_OBJECT_ID field contains zero, any User_Object_ID may be assigned.

Editor's Note: The proposed change will be made in both the CREATE command and the CREATE AND WRITE command definitions.

Veritas 89) to the [s/b] when (Rejected) [570]

PDF pg 94, pg 75, 6.4 CREATE AND WRITE (3 instances)

in 4.8.3 apply to the testing of the

Proposed Resolution:

in 4.8.3 apply when testing the

Editor's Note: Recent T10 lore has attributed a degree of imprecision to the word 'when'. Since its use is not necessary in the cited cases, no changes will be made.

Veritas 90) FrameMaker cross reference marker not on the correct paragraph (Accepted, Editorial) [571]

PDF pg 94, pg 75, 6.4 CREATE AND WRITE

maximum user object length attribute in the The COLLECTIONS PER USER OBJECT field contains the value of the collections per user object attribute. (see), then a

Proposed Resolution:

maximum user object length attribute (see 7.1.2.14), then a

Editor's Note: The problem is a technical gaff in the FrameMaker source file. The cross reference marker used to generate the text the User Object Quotas attributes page subclause cross reference for the maximum user object length attribute description is located on the wrong paragraph. The correct location for the cross reference marker is on the header. The cross reference marker will be moved and the cross references will be regenerated to correct this problem.

Veritas 91) of partition [s/b] of the partition (Accepted, Editorial) [572]

PDF pg 95, pg 76, 6.5 CREATE COLLECTION

Partition_ID of partition in which

Proposed Resolution:

Partition_ID of the partition in which

Veritas 92) to the [s/b] when (Rejected) [573]

PDF pg 96, pg 77, 6.5 CREATE COLLECTION

in 4.8.3 apply to the testing of the

Proposed Resolution:

in 4.8.3 apply when testing the

Editor's Note: Recent T10 lore has attributed a degree of imprecision to the word 'when'. Since its use is not necessary in the cited case, no changes will be made.

Veritas 93) shall terminated [s/b] shall be terminated (Accepted, Editorial) [574]
PDF pg 96, pg 77, 6.6 CREATE PARTITION

the command shall terminated

Proposed Resolution:

the command shall be terminated

Editor's Note: The proposed change will be made in both the CREATE COLLECTION command and the CREATE PARTITION command definitions.

Veritas 94) to the [s/b] when (Rejected) [575]
PDF pg 96, pg 77, 6.6 CREATE PARTITION

in 4.8.3 apply to the testing of the

Proposed Resolution:

in 4.8.3 apply when testing then

Editor's Note: Recent T10 lore has attributed a degree of imprecision to the word 'when'. Since its use is not necessary in the cited case, no changes will be made.

Veritas 95) for returned list [s/b] for the returned list (Accepted, Editorial) [576]
PDF pg 103, pg 84, 6.10 LIST

has allocated for returned list.

Proposed Resolution:

has allocated for the returned list.

Editor's Note: The proposed change will be made in both the LIST command and the LIST COLLECTION command definitions.

Veritas 96) that command [s/b] that new command (Accepted, Editorial) [577]
PDF pg 104, pg 85, 6.10 LIST

INITIAL OBJECT_ID field of that command.

Proposed Resolution:

INITIAL OBJECT_ID field of that new command.

Editor's Note: The proposed change will be made in both the LIST command and the LIST COLLECTION command definitions.

Veritas 97) that command [s/b] that new command (Accepted, Editorial) [578]
 PDF pg 107, pg 88, 6.11 LIST COLLECTION

LIST IDENTIFIER field of that command.

Proposed Resolution:

LIST IDENTIFIER field of that new command.

Editor's Note: The proposed change will be made in both the LIST command and the LIST COLLECTION command definitions.

Veritas 98) Make support for PERFORM SCSI COMMAND mandatory (No Action Taken) [579]
 PDF pg 108, pg 89, 6.12 PERFORM SCSI COMMAND

An OSBD should support this command even when the security level is NOSEC.

Proposed resolution

Clarify that NOSEC devices SHALL support PERFORM SCSI COMMAND

Editor's Note: Table 39 clearly indicates that support for the PERFORM SCSI COMMAND command is mandatory. Nothing in the current standard allows rejecting a PERFORM SCSI COMMAND command based on the security method currently in effect. There is only a statement that the PERFORM SCSI COMMAND command is useful when the security method is something other than NOSEC, but that is followed immediately by a statement showing more general usefulness for the PERFORM SCSI COMMAND command. Surely, the market can ferret out OSBDs that inappropriately reject the PERFORM SCSI COMMAND command.

Veritas 99) Make support for PERFORM TASK MANAGEMENT FUNCTION mandatory (No Action Taken) [580]

PDF pg 110, pg 91, 6.13 PERFORM TASK MANAGEMENT FUNCTION

An OSD should support this command even when the security level is NOSEC.

Proposed resolution

Clarify that NOSEC devices SHALL support PERFORM TASK MANAGEMENT FUNCTION

Editor's Note: Table 39 clearly indicates that support for the PERFORM TASK MANAGEMENT FUNCTION command is mandatory. Nothing in the current standard allows rejecting a PERFORM SCSI TASK MANAGEMENT FUNCTION command based on the security method currently in effect. There is only a statement that the PERFORM SCSI TASK MANAGEMENT FUNCTION command is useful when the security method is something other than NOSEC, but that is followed immediately by a statement showing more general usefulness for the PERFORM SCSI TASK MANAGEMENT FUNCTION command. Surely, the market can ferret out OSBDs that inappropriately reject the PERFORM SCSI TASK MANAGEMENT FUNCTION command.

Veritas 100) object [s/b] objects (Accepted, Editorial) [581]
 PDF pg 115, pg 96, 6.16 REMOVE COLLECTION

the collection contains user object,

Proposed Resolution:

the collection contains user objects,

Veritas 101) to the [s/b] when (Rejected) [582]
 PDF pg 122, pg 103, 6.21 WRITE (2 instances)

in 4.8.3 apply to the testing of the

Proposed Resolution:

described in 4.8.3 apply when testing the

Editor's Note: Recent T10 lore has attributed a degree of imprecision to the word 'when'. Since its use is not necessary in the cited cases, no changes will be made.

Veritas 102) CREATE AND WRITE [s/b] WRITE (Accepted, Editorial) [583]
 PDF pg 122, pg 103, 6.21 WRITE

References to CREATE_AND_WRITE in last 2 paragraphs don't belong. This section is on WRITE.

Proposed Resolution

Change CREATE_AND_WRITE to WRITE.

Veritas 103) Who defines page format? (No Action Taken) [584]
 PDF pg 123, pg 104, 7.1.1 Attributes parameter formats

The sentence "Those attributes pages that do not have a defined page format are not accessible via page format parameter data (e.g., the Root Directory attributes page defined in 7.1.2.4)." What does this mean? Who defines the page formats? This spec? Can vendors define pages that can be accessed with page format?

What is the Root Directory page an example of?

Editor's Note: The entity that defines an attributes page may or may not choose to define a page format. Since vendors can define attributes pages, they can also define page formats for those attributes pages.

This standard defines page formats for some of the attributes pages it defines but not for others. The Root Directory attributes page is an example of an attributes page defined by this standard for which no attributes page format is defined. The Root Quotas attributes page is an example of an attributes page defined by this standard for which a page format is defined.

Veritas 104) Clarify attribute access limitations (Rejected) [585]
 PDF pg 123, pg 104, 7.1.1 Attributes parameter format

The sentence "Attribute access is limited to the". Is this for both page and list format?

Proposed resolution

Clarify if that this applies to both page and list formats.

Reason for Rejection: The cited statement is an overview statement. The precise details are sufficiently complex as to require reading all of clause 7. There is no practical way to provide any more detail in a manner that is consistent with the level of coverage in the cited subclause.

Veritas 105) Use Length FFFF FFFFh for Undefined Attributes (Accepted, Substantive) [586]

PDF pg 123, pg 104, 7.1.1 Attributes parameter format

Then sentence "If an attribute value that has not been previously established is requested, a list entry format value (see 7.1.3.3) having zero in the ATTRIBUTE LENGTH field shall be returned."

After considering the issue, we think it would be nice to distinguish a zero length attribute from an empty attribute.

Proposed resolution

Reserve the attribute length FFFF FFFF for empty attributes. Change wording here and in 7.1.3.3 to reflect this change.

Editor's Note: In 7.1.1 (Attributes parameter formats) the cited text will be modified as follows:

If an attribute value that has not been previously established is requested, a list entry format value (see 7.1.3.3) having ~~zero~~ FFFF FFFFh in the ATTRIBUTE LENGTH field shall be returned.

No text can be found to change in 7.1.3.3 (List entry format for retrieved attributes and for setting attributes for this OSD object).

Veritas 106) Replace 'them' with specific wording (Accepted, Substantive) [587]

PDF pg 123, pg 104, 7.1.1 Attributes parameter formats

allows retrieval of them in formatted pages

Proposed Resolution:

allows retrieval of attribute values in formatted pages

Editor's Note: Since each attribute includes both a length and a value, 'them' will be changed to 'attributes'

Veritas 107) Clarify title of table 67 (Accepted, Editorial) [588]

PDF pg 124, pg 105, Table 67 - Attributes pages

The following attribute page ranges are not described in this table:

80 to P - 1 (2FFF FFFFh)
 P + 80 (3000 0080h) to C - 1 (5FFF FFFFh)
 C + 80 (6000 0080h) to 8FFF FFFFh)
 R + 80 (9000 0080h) to EFFF FFFFh)

Proposed resolution

Describe the missing ranges in the table, and who each range is defined by.

Editor's Note: The cited ranges of attributes pages are described in table 3 and table 4. That information should not be duplicated here.

The title for table 67 is misleading. It will be modified as follows: "Attributes pages **defined by this standard**"

Veritas 108) Indicate which attributes pages have a page format defined (Accepted, Editorial) [589]

PDF pg 124, pg 105, Table 67 - Attributes pages
See also: comment Seagate 36) and comment HP 85)

Not all attribute pages on this table can be accessed in page format (For example, the information pages)

This information should appear in this table.

Proposed resolution

Add a "Page Format Available" column to this table.

Editor's Note: This comment will be resolved as described in the response to comment HP 85).

Veritas 109) number partitions [s/b] number of partitions (Accepted, Editorial) [590]

PDF pg 131, pg 112, 7.1.2.8 Root Information attributes page

(number C0h) shall contain the number partitions present

Proposed Resolution:

(number C0h) shall contain the number of partitions present

Veritas 110) OSD Initialized Attributes Should be 'Yes' In OSD Provided (No Action Taken) [591]

PDF pg 135, pg 116, Table 77, Root Quotas attribute page contents

Fields that are set by the OSD when FORMAT'ed should be marked as 'YES' in the "OSD Logical Unit Provided" column.

Proposed resolution

All "OSD Logical Unit Provided" entries currently marked 'NO' should be changed to 'YES'.

Editor's Note: This comment was withdrawn by its author.

Veritas 111) OSD Initialized Attributes Should be 'Yes' In OSD Provided (No Action Taken) [592]

PDF pg 137, pg 118, Table 79, Partition Quotas attribute page contents

Fields that are set by the OSD at CREATE_PARTITION time should be marked as 'YES' in the "OSD Logical Unit Provided" column.

Proposed resolution

All "OSD Logical Unit Provided" entries currently marked 'NO' should be changed to 'YES'.

Editor's Note: This comment was withdrawn by its author.

Veritas 112) FrameMaker cross reference marker not on the correct paragraph (Accepted, Editorial) [593]
 PDF pg 137, pg 118, 7.1.2.13 Partition Quotas attributes page

maximum user object length attribute in the The COLLECTIONS PER USER OBJECT field contains the value of the collections per user object attribute. (see)

Proposed Resolution:

maximum user object length attribute in the User Object Quotas attributes page (see 7.1.2.14)

Editor's Note: The problem is a technical gaff in the FrameMaker source file. The cross reference marker used to generate the text the User Object Quotas attributes page subclause cross reference for the maximum user object length attribute description is located on the wrong paragraph. The correct location for the cross reference marker is on the header. The cross reference marker will be moved and the cross references will be regenerated to correct this problem.

Veritas 113) OSD Initialized Attributes Should be 'Yes' In OSD Provided (No Action Taken) [594]
 PDF pg 138, pg 119, Table 81, User Object Quotas attribute page contents

Fields that are set by the OSD at CREATE or CREATE_AND_WRITE time should be marked as 'YES' in the "OSD Logical Unit Provided" column.

Proposed resolution

All "OSD Logical Unit Provided" entries currently marked 'NO' should be changed to 'YES'.

Editor's Note: This comment was withdrawn by its author.

Veritas 114) the allow [s/b] allowed (Accepted, Editorial) [595]
 PDF pg 139, pg 120, 7.1.2.14 User Object Quotas attributes page

(number 1h) specifies the maximum value the allow in the user

Proposed Resolution:

(number 1h) specifies the maximum value allowed in the user

Veritas 115) get and set [s/b] get (Rejected) [596]
 PDF pg 140, pg 121, 7.1.2.15 Root Timestamps attributes page (2 instances)

whose CDB get and set attributes parameters (see 5.2.1)

Proposed Resolution:

whose CDB get attributes parameters (see 5.2.1)

Reason for Rejection: The name of the CDB section (i.e., set of fields) is 'get and set attributes parameters' as is shown by the title of 5.2.1 and in every CDB format shown in clause 6. The current wording is correct.

Veritas 116) Does Changing the User Object Logical Length Update the Data Modified Timestamp? (Accepted, Substantive) [597]

PDF pg 146, pg 127, 7.1.2.18 User Object Timestamp attributes page

Does modifying the logical size of an object by writing to "User Object Logical Length" attribute of the User Object Information page update the Data Modified Time? We think that modifying the logical length should update both the data modified time and the attribute modified time.

Proposed Resolution

Update description to include changing the logical length as something that causes the data modified time to be updated.

Editor's Note: The cited paragraph will be modified as follows:

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 **command that changed the value of the user object logical length attribute in the User Object Information attributes page (see 7.1.2.11) or that stored data in the user object (i.e., WRITE command (see 6.21), APPEND command (see 6.2), or CREATE AND WRITE command (see 6.4)) that stored data in the user object.**

Veritas 117) Clarify adjustable clock attribute definition (Accepted, Editorial) [598]

PDF pg 151, pg 132, 7.1.2.20 Root Security attributes page

to maintain value in the adjustable clock attribute value outside

Proposed Resolution:

to maintain a value in the adjustable clock attribute is outside

Editor's Note: The cited sentence will be modified as follows:

The mechanism used to maintain **value in** the adjustable clock attribute value **is** outside the scope of the standard.

Veritas 118) an request [s/b] a request (Accepted, Editorial) [599]

PDF pg 156, pg 137, 7.1.2.21 Partition Security attributes page (two instances)

(see 5.2.1.3) contains an request

Proposed Resolution:

(see 5.2.1.3) contains a request

Veritas 119) an request [s/b] a request (Accepted, Editorial) [600]

PDF pg 159, pg 140, 7.1.2.22 Collection Security attributes page

(see 5.2.1.3) contains an request

Proposed Resolution:

(see 5.2.1.3) contains a request

Veritas 120) an request [s/b] a request (Accepted, Editorial) [601]
PDF pg 160, pg 141, 7.1.2.23 User Object Security attributes page

(see 5.2.1.3) contains an request

Proposed Resolution:

(see 5.2.1.3) contains a request

Veritas 121) in the in the [s/b] in the (Accepted, Editorial) [602]
PDF pg 165, pg 146, 7.1.3.3 List entry format....

by the ATTRIBUTE NUMBER field in the in the attributes page

Proposed Resolution:

by the ATTRIBUTE NUMBER field in the attributes page

14. Late Comments

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

Other 1) Clarify lack of data area in the root object, partitions, and collections (Accepted, Editorial) [603]
 pgs 18&19, 4.6.3 p 2, 4.6.4 p 5, 4.6.6 p 6
 See also comment IBM 19), comment IBM 21), and comment Seagate 4)

The subclauses describing the root object, partitions, and collections all conclude with a paragraph intended to convey that the object does not have a read/write data area. There has been some confusion over this intent, resulting in requests to add a complete list of the commands prohibited. Such a list is impractical to maintain correctly over time.

To clarify the intent of the three cited paragraphs, add the following to the beginning of each one:

{The root object | A partition | A collection} does not contain a read/write data area.

Other 2) partition for [s/b] partition that represents (Accepted, Editorial) [604]
 pg 7, 3.1.31 partition zero

Modify the second sentence of the partition zero definition as follows:

The partition ~~for~~ that represents the root object (see 3.1.33).

Other 3) Correct object count testing wording (Accepted, Editorial) [605]
 pg 73, 6.3 (CREATE), last p
 pg 75, 6.4 (CREATE AND WRITE), 3rd from last p, and
 pg 77, 6.3 (CREATE COLLECTION), last p in subclause

Modify the CREATE command discussion of object count testing as follows:

If a CREATE command causes the value in the number of collections and user objects attribute in the Partition Information attributes page (see 7.1.2.9) ~~sum of the number of collections and the number of user objects in the partition~~ to exceed the value in the object count attribute in the Partition Quotas attributes page (see 7.1.2.13), then a quota error shall be generated (see 4.8.2). The quota testing principles described in 4.8.3 apply to the testing of the object count quota.

Modify the CREATE AND WRITE command discussion of object count testing as follows:

If a CREATE AND WRITE command causes the value in the number of collections and user objects attribute in the Partition Information attributes page (see 7.1.2.9) ~~sum of the number of collections and the number of user objects in the partition~~ to exceed the value in the object count attribute in the Partition Quotas attributes page (see 7.1.2.13), then a quota error shall be generated (see 4.8.2). The quota testing principles described in 4.8.3 apply to the testing of the object count quota.

Modify the CREATE COLLECTION command discussion of object count testing as follows:

If a CREATE COLLECTION command causes the value in the number of collections and user objects attribute in the Partition Information attributes page (see 7.1.2.9) ~~sum of the number of collections and the number of user objects in the partition~~ to exceed the value in the object count attribute in the Partition Quotas attributes page (see 7.1.2.13), then a quota error shall be generated (see 4.8.2). The quota testing principles described in 4.8.3 apply to the testing of the object count quota.

Other 4) Correct partition count testing wording (Accepted, Editorial) [606]
pg 78, 6.6 (CREATE PARTITION), last p in subclause

Modify the cited paragraph as follows:

If a CREATE PARTITION command causes the **value in the number of partitions attribute in the Root Information attributes page (see 7.1.2.8) number of partitions** to exceed the value in the partition count attribute in the Root Quotas attributes page (see 7.1.2.12), then a quota error shall be generated (see 4.8.2). The quota testing principles described in 4.8.3 apply to the testing of the object count quota.

Other 5) Use consistent notation in table B.1 (Accepted, Editorial) [607]
pg 153, Table B.1 (Numerical order OSD service action codes)

Change:

~~8810h—Reserved~~
~~8811h—Reserved~~

to:

8810h - 8811h Reserved

for consistency with:

8813h - 8814h Reserved

Other 6) Define values of the x+0h attributes in the directory attributes pages (Accepted, Substantive) [608]

pg 107, 7.1.2.4 (Root Directory attributes page)
pg 108, 7.1.2.5 (Partition Directory attributes page)
pg 109, 7.1.2.6 (Collection Directory attributes page)
pg 110, 7.1.2.7 (User Object Directory attributes page)
See also comment HP 90)

In 7.1.2.4 (Root Directory attributes page) insert the following immediately after the a,b,c list:

The Root Directory page identification attribute (number R+0h) shall have the format described in 7.1.2.2 with the **VENDOR IDENTIFICATION** field containing the ASCII characters "INCITS" and the **ATTRIBUTES PAGE IDENTIFICATION** field containing the ASCII characters "T10 Root Directory".

In 7.1.2.5 (Partition Directory attributes page) insert the following immediately after the a,b,c list:

The Partition Directory page identification attribute (number P+0h) shall have the format described in 7.1.2.2 with the **VENDOR IDENTIFICATION** field containing the ASCII characters "INCITS" and the **ATTRIBUTES PAGE IDENTIFICATION** field containing the ASCII characters "T10 Partition Directory".

In 7.1.2.4 (Root Directory attributes page) insert the following immediately after the a,b,c list:

The Root Directory page identification attribute (number R+0h) shall have the format described in 7.1.2.2 with the **VENDOR IDENTIFICATION** field containing the ASCII characters "INCITS" and the **ATTRIBUTES PAGE IDENTIFICATION** field containing the ASCII characters "T10 Root Directory".

In 7.1.2.7 (User Object Directory attributes page) insert the following immediately after the a,b,c list:

The User Object Directory page identification attribute (number 0h) shall have the format described in 7.1.2.2 with the `VENDOR IDENTIFICATION` field containing the ASCII characters "INCITS" and the `ATTRIBUTES PAGE IDENTIFICATION` field containing the ASCII characters "T10 User Object Directory".

Other 7) Remove 'successfully' globally (Accepted, Editorial) [609]
Global

Remove all occurrences of 'successfully'. If it happened, it was successful.

Other 8) Increasing a user object's logical length can produce a quota error (Accepted, Substantive) [610]
pg 115, 7.1.2.11 (User Object Information attributes page)

Modify the definition of the user object logical length attribute as follows:

The user object logical length attribute (number 82h) specifies the largest valued byte number written in the associated user object. Setting the user object logical length attribute to a value that is smaller than the user object's logical length known to the OSD device server shall cause the user object to be truncated to the specified length. Setting the user object logical length attribute to a value that is larger than the user object's logical length known to the OSD device server shall cause unwritten bytes to be added at the end of the user object.

At attempt to set the user object logical length attribute a value that is larger than the value in the maximum user object length attribute in the User Object Quotas attributes page (see 7.1.2.14) shall generate a quota error (see 4.8.2). The quota testing principles described in 4.8.3 apply to the testing of the maximum user object length quota.

If setting the user object logical length attribute to a value that is larger than the user object's logical length known to the OSD device server causes the value in the used capacity attribute in the Partition Information attributes page (see 7.1.2.9) to exceed the value in the capacity quota attribute in the Partition Quotas attributes page (see 7.1.2.13), then a quota error shall be generated (see 4.8.2). The quota testing principles described in 4.8.3 apply to the testing of the capacity quota.

Other 9) Remove comma in table 39 column heading (Accepted, Editorial) [611]
pg 68, 6.1 (Summary of commands for OSD type devices), table 39

Service action: ^a [s/b] Service action ^a

Other 10) ACA requirements (Accepted, Substantive) [612]
pg 58, new subclause before/after 4.14 (Linked commands)

Insert the following new subclause either before or after 4.14 (Linked commands):

4.xx Auto contingent allegiance

OSD logical units that are not capable of accepting a command with the ACA task attribute (see SAM-3) at any time and performing all data transfer operations that the command requests shall set the NORMACA bit to zero in the Standard INQUIRY data (see SPC-3).