

Open: CPQ #28 Page 35 Section 6.8 14

Open: CPQ #29 Page 35 Section 6.8 14

Open: CPQ #30 Page 35 Section 6.8 14

Open: CPQ #31 Page 38 Section 6.9 14

Open: CPQ #32 Page 39 Section 6.9 14

Open: CPQ #33 Page 39 Section 6.9 14

Open: CPQ #34 Page 41 Section 6.12 14

Open: CPQ #35 Page 41 Section 6.12 14

Open: CPQ #36 Page 42 Section 6.13 14

Open: CPQ #37 Page 43 Section 7.1 14

Open: CPQ #38 Page 52 Annex B 14

Open: CPQ #39 Page 52 Annex B 15

Open: CPQ #40 Page 62 Annex C 15

Open: HP #1 15

Open: HP #2 15

Open: HP Page 1 Line 11 15

Open: HP Page 1 Line 19 16

Open: HP Page 2 Line 28 16

Open: HP Page 8 Line 4 16

Open: HP Page 8 Line 17 16

Open: HP Page 8 Line 20 16

Open: HP Page 10 Line 12 16

Open: HP Page 11 Line 36 16

Open: HP Page 12 Line 40 16

Open: HP Page 14 Line 24 17

Open: HP Page 15 Line 24 17

Open: HP Page 16 Line 28 17

Open: HP Page 20 Line 4 17

Open: HP Page 25 Line 1 17

Open: HP Page 25 Line 1 18

Open: HP Page 25 Line 32 18

Open: HP Page 25 Line 32 18

Open: HP Page 54 Line 1 18

Open: HP Page 54 Line 23 18

Open: HP Page 54 Line 28 18

Open: HP Page 54 Line 28 18

Open: HP Page 54 Line 28 18

Open: HP Page 55 Line 9 18

Open: HP Page 55 Line 25 19

Open: HP Page 55 Line 25 19

Open: IBM / Tivoli 0 19

Open: IBM / Tivoli #1 PDF Page 3 Tivoli comment from George Penokie 19

Open: IBM / Tivoli #2 PDF Page 4 Tivoli comment from George Penokie 19

Open: IBM / Tivoli #3 PDF Page 4 Tivoli comment from George Penokie 19

Open: IBM / Tivoli #4 PDF Page 6 Tivoli comment from George Penokie 20

Open: IBM / Tivoli #5 PDF Page 11 Tivoli comment from George Penokie 20

Open: IBM / Tivoli #6 PDF Page 12 Tivoli comment from George Penokie 20

Open: IBM / Tivoli #7 PDF Page 12 Tivoli comment from George Penokie 20

Open: IBM / Tivoli #8 PDF Page 13 Tivoli comment from George Penokie 20

Open: IBM / Tivoli #9 PDF Page 16 Tivoli comment from George Penokie 20

Open: IBM / Tivoli #10 PDF Page 16 Tivoli comment from George Penokie 20

Open: IBM / Tivoli #11 PDF Page 16 Tivoli comment from George Penokie 20

Open: IBM / Tivoli #12 PDF Page 16 Tivoli comment from George Penokie 20

Open: IBM / Tivoli #13 PDF Page 16 Tivoli comment from George Penokie 20

Open: IBM / Tivoli #122 PDF Page 50 Tivoli comment from George Penokie 30

Open: IBM / Tivoli #123 PDF Page 51 Tivoli comment from George Penokie 30

Open: IBM / Tivoli #124 PDF Page 51 Tivoli comment from George Penokie 30

Open: IBM / Tivoli #125 PDF Page 53 Tivoli comment from George Penokie 30

Open: IBM / Tivoli #126 PDF Page 54 Tivoli comment from George Penokie 30

Open: IBM / Tivoli #127 PDF Page 56 Tivoli comment from George Penokie 31

Open: IBM / Tivoli #128 PDF Page 56 Tivoli comment from George Penokie 31

Open: IBM / Tivoli #129 PDF Page 58 Tivoli comment from George Penokie 31

Open: IBM / Tivoli #130 PDF Page 58 Tivoli comment from George Penokie 31

Open: IBM / Tivoli #131 PDF Page 64 Tivoli comment from George Penokie 31

Open: IBM / Tivoli #132 PDF Page 64 Tivoli comment from George Penokie 31

Open: IBM / Tivoli #133 PDF Page 65 Tivoli comment from George Penokie 31

Open: IBM / Tivoli #134 PDF Page 65 Tivoli comment from George Penokie 31

Open: IBM / Tivoli #135 PDF Page 67 Tivoli comment from George Penokie 31

Open: IBM / Tivoli #136 PDF Page 68 Tivoli comment from George Penokie 31

Open: IBM / Tivoli #137 PDF Page 68 Tivoli comment from George Penokie 31

Open: IBM / Tivoli #138 PDF Page 68 Tivoli comment from George Penokie 31

Open: IBM / Tivoli #139 PDF Page 68 Tivoli comment from George Penokie 31

Open: IBM / Tivoli #140 PDF Page 68 Tivoli comment from George Penokie 32

Open: IBM / Tivoli #141 PDF Page 68 Tivoli comment from George Penokie 32

Open: IBM / Tivoli #142 PDF Page 68 Tivoli comment from George Penokie 32

Open: IBM / Tivoli #143 PDF Page 68 Tivoli comment from George Penokie 32

Open: IBM / Tivoli #144 PDF Page 68 Tivoli comment from George Penokie 32

Open: IBM / Tivoli #145 PDF Page 69 Tivoli comment from George Penokie 32

Open: IBM / Tivoli #146 PDF Page 69 Tivoli comment from George Penokie 32

Open: IBM / Tivoli #147 PDF Page 73 Tivoli comment from George Penokie 32

Open: IBM / Tivoli #148 PDF Page 73 Tivoli comment from George Penokie 32

Open: IBM / Tivoli #149 32

Open: IBM / Tivoli #150 32

Open: IBM / Tivoli #151 32

Open: IBM / Tivoli #152 32

Open: intel0001 Sect:1 Pg:1 Ln:32 33

Open: intel0002 Sect:1 Pg:1 Ln:37 33

Open: intel0003 Sect:3.1.11 Pg:4 Ln:27 33

Open: intel0004 Sect:3.1.14 Pg:4 Ln:34 33

Open: intel0005 Sect:3.1.15 Pg:4 Ln:35 33

Open: intel0006 Sect:3.1.16 Pg:4 Ln:39 33

Open: intel0007 Sect:3.1.17 Pg:4 Ln:41 33

Open: intel0008 Sect:3.1.17 Pg:4 Ln:43 C 33

Open: intel0009 Sect:3.1.23 Pg:5 Ln:6 33

Open: intel0010 Sect:3.1.27 Pg:5 Ln:12 C 33

Open: intel0011 Sect:3.1.28 Pg:5 Ln:15 33

Open: intel0012 Sect:3.3.9 Pg:6 Ln:12 33

Open: intel0013 Sect:4 Pg:8 Ln:1 C 33

Open: intel0014 Sect:4.2 Pg:8 Ln:46 33

Open: intel0015 Sect:4.2 Pg:9 Ln:2 C 33

Open: intel0016 Sect:4.2 Pg:9 Ln:2 34

Open: intel0017 Sect:4.2 Pg:9 Ln:29 34

Open: intel0018 Sect:4.2 Pg:9 Ln:50 C 34

Open: intel0019 Sect:4.2 Pg:10 Ln:17 34

Open: intel0020 Sect:4.2 Pg:10 Ln:22 34

Open: intel0021 Sect:4.2 Pg:10 Ln:26 34

Open: intel0022 Sect:4.2 Pg:10 Ln:28 34

Open: intel0023 Sect:4.2 Pg:10 Ln:30 34

Open: intel0024 Sect:4.2 Pg:10 Ln:42 34

Open: intel0025 Sect:4.2 Pg:11 Ln:2 34

Open: intel0026 Sect:4.2 Pg:11 Ln:5 34

Open: intel0027 Sect:4.2 Pg:11 Ln:9 34

Open: intel0028 Sect:4.2 Pg:11 Ln:9 34

Open: intel0029 Sect:4.2 Pg:11 Ln:13 35

Open: intel0030 Sect:4.2 Pg:11 Ln:16 35

Open: intel0031 Sect:4.2 Pg:11 Ln:16 35

Open: intel0032 Sect:4.2 Pg:11 Ln:13 35

Open: intel0033 Sect:4.3 Pg:11 Ln:18 35

Open: intel0034 Sect:4.3 Pg:11 Ln:21 35

Open: intel0035 Sect:4.3 Pg:11 Ln:24 35

Open: intel0036 Sect:4.4 Pg:12 Ln:23 C 35

Open: intel0037 Sect:4.5 Pg:12 Ln:45 35

Open: intel0038 Sect:4.5 Pg:13 Ln:15 35

Open: intel0039 Sect:5.1.1 Pg:14 Ln:20 35

Open: intel0040 Sect:5.1.1 Pg:14 Ln:26 35

Open: intel0041 Sect:5.1.1 Pg:14 Ln:31 35

Open: intel0042 Sect:5.1.1 Pg:14 Ln:44 36

Open: intel0043 Sect:5.1.3 Pg:15 Ln:16 36

Open: intel0044 Sect:5.1.3 Pg:15 Ln:22 36

Open: intel0045 Sect:5.1.3 Pg:15 Ln:27 36

Open: intel0046 Sect:5.1.3 Pg:15 Ln:36 36

Open: intel0047 Sect:5.3 Pg:16 Ln:24 36

Open: intel0048 Sect:5.3 Pg:16 Ln:27 36

Open: intel0049 Sect:5.3 Pg:16 Ln:40 C 36

Open: intel0050 Sect:5.3 Pg:16 Ln:42 C 36

Open: intel0051 Sect:5.3 Pg:16 Ln:46 C 36

Open: intel0052 Sect:5.3 Pg:16 Ln:46 C 36

Open: intel0053 Sect:5.3 Pg:16 Ln:46 C 37

Open: intel0054 Sect:5.3 Pg:17 Ln:1 37

Open: intel0055 Sect:5.3 Pg:17 Ln:1 C 37

Open: intel0056 Sect:5.4.1 Pg:17 Ln:6 37

Open: intel0057 Sect:5.4.1 Pg:17 Ln:6 37

Open: intel0058 Sect:5.4.1 Pg:17 Ln:6 37

Open: intel0059 Sect:5.4.1 Pg:17 Ln:36 37

Open: intel0060 Sect:5.4.1 Pg:18 Ln:2 37

Open: intel0061 Sect:5.4.1 Pg:18 Ln:4 37

Open: intel0062 Sect:5.4.1 Pg:18 Ln:4 37

Open: intel0063 Sect:5.4.1 Pg:18 Ln:17 37

Open: intel0064 Sect:5.4.1 Pg:18 Ln:17 C 37

Open: intel0065 Sect:5.4.1 Pg:18 Ln:24 38

Open: intel0066 Sect:5.4.1 Pg:18 Ln:32 38

Open: intel0067 Sect:5.4.1 Pg:18 Ln:43 38

Open: intel0068 Sect:5.4.1 Pg:18 Ln:47 38

Open: intel0069 Sect:5.4.2.2 Pg:18 Ln:49 38

Open: intel0070 Sect:5.4.2.2 Pg:19 Ln:16 38

Open: intel0071 Sect:5.4.2.2 Pg:19 Ln:16 38

Open: intel0072 Sect:5.4.2.2 Pg:19 Ln:18 38

Open: intel0073 Sect:5.4.2.4 Pg:19 Ln:44 38

Open: intel0074 Sect:5.4.2.4 Pg:19 Ln:48 38

Open: intel0075 Sect:5.4.2.4 Pg:20 Ln:1 38

Open: intel0076 Sect:5.4.2.5 Pg:20 Ln:6 38

Open: intel0077 Sect:5.4.2.5 Pg:20 Ln:8 38

Open: intel0078 Sect:5.4.2.5 Pg:20 Ln:10 39

Open: intel0079 Sect:5.4.2.5 Pg:20 Ln:11 39

Open: intel0080 Sect:5.4.2.5 Pg:20 Ln:11 39

Open: intel0081 Sect:5.4.2.5 Pg:20 Ln:12 39

Open: intel0082 Sect:5.4.2.5 Pg:20 Ln:13 39

Open: intel0083 Sect:5.4.2.5 Pg:20 Ln:29 39

Open: intel0084 Sect:5.4.2.5 Pg:20 Ln:29 39

Open: intel0085 Sect:5.4.2.5 Pg:20 Ln:31 39

Open: intel0086 Sect:5.4.2.5 Pg:20 Ln:33 39

Open: intel0087 Sect:5.4.2.5 Pg:20 Ln:35 39

Open: intel0088 Sect:5.4.2.5 Pg:20 Ln:36 39

Open: intel0089 Sect:5.4.2.5 Pg:20 Ln:43 39

Open: intel0090 Sect:5.4.2.5 Pg:20 Ln:47 39

Open: intel0091 Sect:5.4.2.5 Pg:21 Ln:1 39

Open: intel0092 Sect:5.4.2.5 Pg:21 Ln:12 39

Open: intel0093 Sect:5.4.2.5 Pg:21 Ln:44 39

Open: intel0094 Sect:6.1 Pg:23 Ln:7 40

Open: intel0095 Sect:6.1 Pg:23 Ln:24 40

Open: intel0096 Sect:6.1 Pg:23 Ln:46 40

Open: intel0097 Sect:6.1 Pg:23 Ln:48 40

Open: intel0098 Sect:6.1 Pg:24 Ln:2 C 40

Open: intel0099 Sect:6.2 Pg:24 Ln:2 40

Open: intel0100 Sect:6.2 Pg:24 Ln:41 40

Open: intel0101 Sect:6.3 Pg:27 Ln:4 40

Open: intel0102 Sect:6.3 Pg:27 Ln:40 40

Open: intel0103 Sect:6.3 Pg:27 Ln:45 C 40

Open: intel0104 Sect:6.4 Pg:29 Ln:3 40

Open: intel0105 Sect:6.4 Pg:29 Ln:40 40

Open: intel0106 Sect:6.5 Pg:30 Ln:20 40

Open: intel0107 Sect:6.6 Pg:31 Ln:3 C 40

Open: intel0108 Sect:6.6 Pg:31 Ln:30 40

Open: intel0109 Sect:6.6 Pg:31 Ln:45 41

Open: intel0110 Sect:6.8 Pg:34 Ln:14 41

Open: intel0111 Sect:6.8 Pg:34 Ln:40 41

Open: intel0112 Sect:6.9 Pg:36 Ln:36 41

Open: intel0113 Sect:6.9 Pg:37 Ln:9 41

Open: intel0114 Sect:6.9 Pg:37 Ln:26 41

Open: intel0115 Sect:6.9 Pg:37 Ln:26 41

Open: intel0116 Sect:6.9 Pg:37 Ln:26 C 41

Open: intel0117 Sect:6.9 Pg:37 Ln:47 41

Open: intel0118 Sect:6.9 Pg:38 Ln:17 41

Open: intel0119 Sect:6.9 Pg:38 Ln:24 41

Open: intel0120 Sect:6.9 Pg:38 Ln:31 41

Open: intel0121 Sect:6.9 Pg:38 Ln:45 41

Open: intel0122 Sect:6.9 Pg:39 Ln:1 41

Open: intel0123 Sect:6.9 Pg:39 Ln:18 41

Open: intel0124 Sect:6.9 Pg:39 Ln:31 41

Open: intel0125 Sect:6.9 Pg:39 Ln:33 42

Open: intel0126 Sect:6.9 Pg:39 Ln:30 42

Open: intel0127 Sect:6.11 Pg:40 Ln:43 C 42

Open: intel0128 Sect:6.12 Pg:41 Ln:31 42

Open: intel0129 Sect:6.12 Pg:41 Ln:43 42

Open: intel0130 Sect:6.12 Pg:41 Ln:44 C 42

Open: intel0131 Sect:6.12 Pg:42 Ln:1 42

Open: intel0132 Sect:6.12 Pg:42 Ln:7 42

Open: intel0133 Sect:7.2 Pg:43 Ln:21 42

Open: intel0134 Sect:7.2 Pg:43 Ln:25 42

Open: intel0135 Sect:7.2 Pg:44 Ln:1 42

Open: intel0136 Sect:7.2 Pg:44 Ln:1 42

Open: intel0137 Sect:7.2 Pg:44 Ln:7 42

Open: intel0138 Sect:7.2 Pg:44 Ln:17 42

Open: intel0139 Sect:7.2 Pg:44 Ln:19 42

Open: intel0140 Sect:7.2 Pg:44 Ln:24 43

Open: intel0141 Sect:7.2 Pg:44 Ln:24 43

Open: intel0142 Sect:7.2 Pg:44 Ln:28 43

Open: intel0143 Sect:A.1 Pg:45 Ln:11 43

Open: intel0144 Sect:A.1 Pg:45 Ln:29 43

Open: intel0145 Sect:A.1 Pg:46 Ln:16 43

Open: intel0146 Sect:A.1 Pg:46 Ln:43 43

Open: intel0147 Sect:A.3 Pg:47 Ln:11 43

Open: intel0148 Sect:A.4.1 Pg:48 Ln:44 43

Open: intel0149 Sect:B.3.1.7 Pg:52 Ln:35 43

Open: intel0150 Sect:B.3.1.2 Pg:52 Ln:23 43

Open: intel0151 Sect:B.3.1.14 Pg:53 Ln:1 43

Open: intel0152 Sect:B.3.1.16 Pg:53 Ln:5 43

Open: intel0153 Sect:B.3.2 Pg:53 Ln:20 43

Open: intel0154 Sect:B.4 Pg:54 Ln:50 C 43

Open: intel0155 Sect:B.4 Pg:55 Ln:17 43

Open: intel0156 Sect:B.5 Pg:56 Ln:2 43

Open: intel0157 Sect:B.5 Pg:56 Ln:15 44

Open: intel0158 Sect:B.5 Pg:56 Ln:17 44

Open: intel0159 Sect:B.5 Pg:56 Ln:36 44

Open: intel0160 Sect:B.5 Pg:56 Ln:41 44

Open: intel0161 Sect:B.5 Pg:56 Ln:47 44

Open: intel0162 Sect:B.5 Pg:56 Ln:48 44

Open: intel0163 Sect:B.6.2 Pg:57 Ln:13 44

Open: intel0164 Sect:B.6.3 Pg:57 Ln:25 44

Open: intel0165 Sect:B.6.4 Pg:57 Ln:38 44

Open: intel0166 Sect:B.6.4 Pg:57 Ln:38 44

Open: intel0167 Sect:B.6.4 Pg:57 Ln:38 44

Open: intel0168 Sect:B.6.4 Pg:57 Ln:42 44

Open: intel0169 Sect:B.6.5 Pg:57 Ln:46 44

Open: intel0170 Sect:B.6.5 Pg:58 Ln:1 44

Open: intel0171 Sect:B.7 Pg:58 Ln:37 44

Open: intel0172 Sect:B.7 Pg:59 Ln:7 45

Open: intel0173 Sect:B.7 Pg:60 Ln:23 C 45

Open: intel0174 Sect:B.7 Pg:60 Ln:24 C 45

Open: intel0175 Sect:B.7 Pg:60 Ln:26 C 45

Open: intel0176 Sect:B.7 Pg:60 Ln:46 C 45

Open: intel0177 Sect:B.7 Pg:61 Ln:13 C 45

Open: intel0178 Sect:B.7 Pg:61 Ln:16 45

Open: intel0179 Sect:B.7 Pg:61 Ln:16 45

Open: Comment from IBTA by William Futral (Intel): 45

Open: OD 1 Page 13, lines 5-7, multiple RDMA writes on the same channel store data in order. .
45

Open: OD 2 Page 13, line 13, RDMA read operations may complete in any order. 46

Open: OD 3 46

Open: OD 4 Page 56, tables B.2 and B.3 47

Open: OD 5 Pages 4 and 5, glossary terms, and their use throughout the document, Clause 4 47
Open: OD 6 Page 11 lines 20-22, normal and solicited message reception. 47
Open: OD 7 RDMA communication service specific opcode 48
Open: OD 8 Page 18, lines 13-37 and elsewhere, data buffer format code and count values. . . 48
Open: Troika Networks, Inc.: 49
Open: Woven Electronics: 49

Open: Brocade1 001 (E) Page: Many Locator: Many

The word 'which' is used inappropriately in many places. Suggested Solution: Do a global search for the word which and replace it with one of the following corrections: A) the word 'that'. B) a new sentence construction that does not require the word. C) nothing. (Which can simply be removed in many cases.)

Open: Brocade1 002 (E) Page: 59 Locator: B.7, figure B.6

The word 'must' is used inappropriately. Suggested Solution: The line 'At least one IB I/O controller must be present' should be replaced. I am not sure if this is a requirement that at one or more controllers shall be present. If so, wording like 'At least one IB I/O controller shall be present' is appropriate.

Open: Brocade1 003 (E) Page: vii Locator: Foreword, line 3

X3.269 is not the proper name Suggested Solution: This value is not correct and should be marked as TBD or XXX or something like that. In any case, it is an NCITS document, not an X3 document.

Open: Brocade1 004 (E) Page: vii Locator: Foreword, line 8

"by National' s/b 'by the National'" Suggested Solution: Correct as requested.

Open: Brocade1 005 (E) Page: viii Locator: Introduction, line 7

"The working draft SCSI' s/b 'The SCSI'" Suggested Solution: This correction should be made now, even though the document is still a working draft, because it is clearly labeled in lots of places that it is a draft, but the text in it is intended to be the content of the standard.

Open: Brocade1 006 (E) Page: 1 Locator: Title, line 6

"The working draft SCSI' s/b 'The SCSI'" Suggested Solution: This correction should be made now, even though the document is still a working draft, because it is clearly labeled in lots of places that it is a draft, but the text in it is intended to be the content of the standard.

Open: Brocade1 007 (E) Page: 3 Locator: 2.1, lines 32-35

Global Engineering should be included here as well, since the drafts are not available from ANSI or NCITS. Suggested Solution: Include Global Engineering as a document source. Include www.t10.org as a document source for standards in development.

Open: Brocade0 001 (E) Page: Locator:

The draft now seems to equate 'SRP target port' and 'IB service', so an SRP target port is designated by a ServiceID. This implies there can be many ports per IOC. This is a significant change from prior drafts where the target port was equated with an IOC, and there was just a single ServiceID per port. It requires a different model for software (OSs or whatever) to manage which hosts have access to which devices in a multi-host environment. Previously, access control was needed only to the level of IOCs, the draft now implies a need to manage not only who can use which IOCs, but which devices within an IOC. Suggested Solution: No solution required if interpretation is correct and implications are understood

Open: Brocade0 002 (E) Page: 52 Locator: B.3.1.2, lines 22-23

The definition of 'IB channel adapter GUID' implies it is the Node GUID but doesn't say so; might as well be explicit Suggested Solution: 'An IB Node GUID that uniquely identifies an IB channel adapter'

Open: Brocade0 003 (E) Page: 52 Locator: B.3.1.9, lines 39-40

The definition of 'IB I/O controller GUID' implies it is the IOControllerProfile GUID but doesn't say so; might as well be explicit Suggested Solution: 'An IB IOControllerProfile GUID that uniquely identifies an IB channel adapter'

Open: Brocade0 004 (E) Page: 54 Locator: Line 50

IB GUIDs can have link-local scope and thus may not be 'globally' unique Suggested Solution: Change to 'unique within a subnet', or 'either unique within a subnet or globally unique'

Open: Brocade0 005 (E) Page: 55 Locator: Table B.1, lines 16-17

IB GUIDs can have link-local scope and thus may not be unique 'worldwide' Suggested Solution: Change 'worldwide' to 'IB subnet or worldwide'

Open: Brocade0 006 (E) Page: 55 Locator: Figure B.3

Figure B.3's equating of 'SRP Target Ports' with 'IB consumers' is problematic. A 'target port' is a sort of service access point--somewhere where interested parties initially go to obtain service, but without any implication that that's where the service is actually provided. (In IB, it's the Connection Manager that receives the initial connection request, interprets the ServiceID contained therein, and performs some magic that results in the instantiation of a QP bound to some entity that actually provides the target services). This target-services-providing entity fits the definition of 'IB consumer'. But the mapping of ServiceIDs-cum-SRP target ports onto such entities is clearly a matter of implementation, and could be one-to-many, many-to-one, or many-to-many Suggested Solution: One possibility: to the left of the IB Consumers show a table/list of service IDs within each IB I/O unit and label these entries as SRP Target Ports; use arrows to show a mapping from the entries to the IB Consumers, with e.g. one Consumer mapped to two IDs and another mapped to one ID to show that the mappings are not always 1 to 1. A further refinement might be to use another set of arrow between the Consumers and the QPs to show that this mapping is also not 1 to 1

Open: Brocade0 007 (E) Page: 56 Locator: Line 16

Missing word Suggested Solution: 'used by the SRP initiator port'?

Open: Brocade0 008 (E) Page: 56 Locator: Line 21

Names of IB attributes are incomplete Suggested Solution: 'IOUnitInfo, IOControllerProfile, and ServiceEntries'

Open: Brocade0 009 (E) Page: 56 Locator: Lines 48-49

'I/O' is broken across lines (and pages) Suggested Solution: Make sure the slash in 'I/O' is non-breaking

Open: Brocade0 010 (E) Page: 56 Locator: Lines 48-49

The phrase 'processor unit or IB I/O controller' makes an incorrect distinction; target ports can only be found on IB I/O controllers by definition, whether or not the I/O controller embodies a processor unit Suggested Solution: Omit 'processor unit or'

Open: Brocade0 011 (E) Page: 57 Locator: Lines 13-14

'IB I/O controllers acting as SRP target ports' could be construed as a 1-to-1 correspondence between controllers and target ports Suggested Solution: 'IB I/O controllers making SRP target ports available' or 'IB I/O controllers hosting SRP target ports'?

Open: Brocade0 012 (E) Page: 61 Locator: Line 1

'An IB I/O controller acting as an SRP target ports' could be construed as a 1-to-1 correspondence between controllers and target ports Suggested Solution: 'And IB I/O controller making SRP target ports available' or 'An IB I/O controller hosting SRP target ports'?

Open: Brocade0 013 (E) Page: 61 Locator: Lines 4-5

'IB I/O controllers acting as SRP target ports' could be construed as a 1-to- 1 correspondence between controllers and target ports Suggested Solution: 'IB I/O controllers making SRP target ports available' or 'IB I/O controllers hosting SRP target ports'?

Open: CPQ #1 Page a Title page

Remove: American National Standard for Information Systems - and change "working draft SCSI RDMA Protocol" to "SCSI RDMA Protocol"

Open: CPQ #2 Page a General

Update the PDF properties title and author

Open: CPQ #3 Page ii General

Remove revision history, line numbers, change bars, etc. from final version

Open: CPQ #4 Page 2 Section 1

Delete CAM from figure 1 Delete these SCSI-2 standards from the example standards list: Serial Storage Architecture SCSI-2 Protocol SSA-S2P [ANSI X3.294:1996] Common Access Method: SCSI Common Access Method CAM [ISO/IEC 9316-421] [ANSI X3.232:1996]

Open: CPQ #5 Page 2 Section 1

Change Fiber to Fibre

Open: CPQ #6 Page 5 Add:

3.1.8 autosense data: Sense data (see 3.1.49) that is returned in the SRP_RSP IU payload. See SAM-2. 3.1.49 sense data: Data returned to an application client as a result of an autosense operation, asynchronous event report, or REQUEST SENSE command. See SPC-2.

Open: CPQ #7 Page 16 Section 5.3

This section should mention the SRP_CRED_REQ and SRP_CRED_RSP IUs, which are dedicated to flow control service.

Open: CPQ #8 Page 18 Section 5.4.2.1

Table 2 Remove period from "NO DATA BUFFER DESCRIPTOR PRESENT."

Open: CPQ #9 Page 18 Section 5.4.2.1

Table 2 There is no reference to note b. It probably needs to be in the 2h row buffer descriptor length cell, where "count" is used

Open: CPQ #10 Page 18 Section 5.4.2.1

Table 2 Add a period at the end of note c.

Open: CPQ #11 Page 19 Section 5.4.2.4

Add a fairly content-free table showing a direct data buffer containing a memory descriptor so this section has a visual reference like the indirect section.

Open: CPQ #12 Page 20 Section 5.4.2.5

Table 5 note a count should be defined with a note b similar to that in table 2

Open: CPQ #13 Page 20 Section 5.4.2.5

Table 4 If n is zero in $16*n+19$, then the table shows byte 20 followed by byte 19. Remove the 20 and that numbering problem is eluded.

Open: CPQ #14 Page 25 Section 6.2

Change: "maximum length" to "maximum length in bytes"

Open: CPQ #15 Page 25 Section 6.2 and elsewhere

I thought we decided that TAG fields don't have bits labeled (MSB)/(LSB).

Open: CPQ #16 Page 25 Section 6.2

Table 9 The REQUIRED BUFFER FORMATS cell is missing the horizontal lines present in other multibyte cells

Open: CPQ #17 Page 26 Section 6.2

Table 10 Remove period from first Reserved. row

Open: CPQ #18 Page 27 Section 6.3

Change (two places): maximum length to "maximum length in bytes"

Open: CPQ #19 Page 29 Section 6.4

Table 14 Capitalize Reserved

Open: CPQ #20 Page 29 Section 6.4

Table 13 The SUPPORTED BUFFER FORMATS cell is missing the horizontal lines present in other multibyte cells

Open: CPQ #21 Page 31 Section 6.6

Table 17 Add period after Reserved or remove from other rows

Open: CPQ #22 Page 33 Section 6.7

Rename TASK MANAGEMENT FLAGS to TASK MANAGEMENT FUNCTION. It doesn't really contain flags.

Open: CPQ #23 Page 33 Section 6.7

Table 19 end each row with a period (or don't)

Open: CPQ #24 Page 33 Section 6.7

Table 19 Change Codes to Code.

Open: CPQ #25 Page 33 Section 6.7

Table 19 Remove small caps from TABLE.

Open: CPQ #26 Page 34 Section 6.8

Table 20 Per Patrick Fitzgerald at JNI, please require that DATA-OUT BUFFER DESCRIPTOR and DATA-IN BUFFER DESCRIPTOR start on 8-byte aligned boundaries. The ADDITIONAL CDB field is only 4 byte aligned.

Open: CPQ #27 Page 34 Section 6.8

Table 20 footnotes Change: length to: length in bytes

Open: CPQ #28 Page 35 Section 6.8

Table 21 SAM-2 rev 20 still requires that untagged tasks be supported by all protocols. 01-318 will remove this requirement and make SRP legal.

Open: CPQ #29 Page 35 Section 6.8

Table 21 Change a to an in the ACA row

Open: CPQ #30 Page 35 Section 6.8

Table 21 Remove small caps from TABLE

Open: CPQ #31 Page 38 Section 6.9

After: The STATUS field contains the status of a task that completes. See the SAM-2 standard for a list of status codes. Add this sentence and a table: Some of the status codes defined in SAM-2 are listed in table xx. Table xx - Some STATUS codes 00h GOOD 02h CHECK CONDITION 08h BUSY 18h RESERVATION CONFLICT 28h TASK SET FULL 30h ACA ACTIVE 40h TASK ABORTED This helps save the reader a reference to SAM-2 for the most popular fields.

Open: CPQ #32 Page 39 Section 6.9

Remove from 2nd sentence of SENSE DATA paragraph: as specified by the SCSI Primary Commands-2 standard.

Open: CPQ #33 Page 39 Section 6.9

Reword the SENSE DATA paragraph to focus on the term autosense which is defined in SAM-2 rather than the REQUEST SENSE command in SPC-2. Change: The SENSE DATA field contains the information specified by the SCSI Primary Commands-2 standard for presentation by the REQUEST SENSE command. The proper sense data shall be presented when a SCSI status byte of CHECK CONDITION is presented by the SCSI Primary Commands -2 standard. to: The SENSE DATA field contains the autosense data (see SCSI Architecture Model - 2) when a SCSI STATUS byte of CHECK CONDITION is presented.

Open: CPQ #34 Page 41 Section 6.12

Change report an asynchronous event. to: report an asynchronous event (see SAM-2).

Open: CPQ #35 Page 41 Section 6.12

Add sentence to first paragraph: Parameters managing the use of asynchronous event reporting are contained in the Control mode page (see SPC-2). This sentence is in SAM-2, but a direct reference from SRP seems helpful.

Open: CPQ #36 Page 42 Section 6.13

Reword the SENSE DATA paragraph like in 6.9, but don't call it autosense here, call it "sense data for the event".

Open: CPQ #37 Page 43 Section 7.1

Table 29 Section 7.3 LUN should be LU (this is broken in SPC too) - the logical unit number is irrelevant here.

Open: CPQ #38 Page 52 Annex B

Change (many places): Infiniband to: InfiniBand

Open: CPQ #39 Page 52 Annex B

There are too many TMs. There only needs to be one per page or one per the whole section.

Open: CPQ #40 Page 62 Annex C

Ralph Weber agreed to put alias formats for each protocol in SPC-3, so this annex can be removed.

Open: HP #1

Need a mandatory requirement to persistently report service names (DevMgtGet-Resp(ServiceEntries)) across IOU/IOC power cycles in order to persistently identify an SRP target port.

Description: Table B.8 describes the format of service name as SRP.T10:xxxxxxxxxxxxxxxx. Since the string xxxxxxxxxxxxxxxx in the service name identifies the 64 bit extension identifier value used to construct the SRP target port identifier, it is required that the service name reported by an IOU for a given SRP target port to be persistent across IOU/IOC power cycles. IB boot records contain SRP initiator port identifier, SRP target port identifier and logical unit name to locate an SRP boot LUN and the assumption is that the target port ID is persistent.

Open: Hp #2

These informal comments are the result of a newcomer's first in-depth reading of the SRP specification. I hope they will suggest avenues for further improvement, but they are not formulated at this time as specific requests for changes.

These comments derive from my work on iSCSI, and are in anticipation of development of iWARP, which will be an RDMA protocol for IP networks. IWARP is intended to provide a standard protocol-independent means of doing direct data placement into host memory, without the need for anonymous reassembly buffers. We anticipate that iSCSI and other Internet storage protocols such as CIFS and NFS will be adapted to iWARP. Inclusion of a formalized RDMA transport layer in the IP storage protocol stack places iSCSI on a path to converge with SRP.

Each protocol can learn from the other. Today, SRP, while meant to be generally applicable, is demonstrably applicable only to InfiniBand. iSCSI's applicability is similarly limited to IP networks. In the future, we may be able to engineer a single SCSI transport that works both with InfiniBand's RDMA service and with iWARP.

These are my personal comments, and are not meant to reflect an HP consensus. We at HP have not yet taken the time to form an internal consensus on SRP.

Open: HP Page 1 Line 11

It is not clear at the outset just what kind of standard SRP is. The text says that "the SCSI family of standards provides for many different transport protocols?" Is SRP a transport protocol? The text continues, "This standard defines the rules for exchanging information between SCSI devices using an RDMA communication service." So SRP is a mapping from SCSI to an abstract RDMA communication service? What then is the SCSI transport? Is it the combination of SRP and the underlying real RDMA communication service? The standard continues, "Other SCSI transport protocol standards?" So, perhaps SRP is a SCSI transport. A statement along these lines would help a lot: "SRP, in combination with a compatible underlying RDMA communication service, is a SCSI transport. This document defines SRP and the requirements that SRP has for the underlying RDMA communication service."

Open: HP Page 1 Line 19

"Figure 1 shows the relationship of this standard to the other standards?" But it doesn't. The SRP standard is not identified in the figure. Despite the disclaimer, layering of the blocks does suggest a hierarchy, protocol stack and system architecture. But the figure does not indicate the applicability of SRP to the implementation of a SCSI transport, as far as I can tell.

Open: HP Page 2 Line 28

SRP is included in a list of transport protocols. So it is a transport protocol. But certainly it is not a complete transport protocol. A discussion of how SRP is used in combination with an underlying RDMA service and its transport protocol to form a SCSI transport protocol would be very instructive to the reader. This would involve a layering diagram-why not?

Open: HP Page 8 Line 4

It would be useful to say at the beginning of clause 4 that the purpose of clause 4 is to describe an abstract RDMA service that is suitable for supporting SRP. That is, to define SRP's requirements of an underlying RDMA service.

Open: HP Page 8 Line 17

"This clause describes various functions that may be provided?" Don't you mean to say that this clause describes various functions that must be provided by an RDMA service, in support of SRP? How the function is provided is immaterial, and of course it can be provided through further functional decomposition. Why mention it? Generally, this whole clause 4 seems to be descriptive of RDMA services in general, but not prescriptive in terms of SRP's requirements. It is difficult to separate descriptive information from requirements.

Open: HP Page 8 Line 20

"Annex B describes the mapping of these functions?" Is it the intention of SRP to work with other RDMA services besides InfiniBand? If so, it might be useful to mention that future revisions of the standard may include other Annexes that define the mapping of SRP to other RDMA services.

Open: HP Page 10 Line 12

SRP is deficient in not providing a security protocol for client (initiator) authentication. Is the notion of "other parameters required by the RDMA communication service" to be interpreted as suggesting that the RDMA service itself should provide authentication? Given that SCSI port names are conveyed by SRP, this doesn't seem possible. (The RDMA service will have its own names for its end nodes, but they're not related to SCSI/SRP port names.)

Open: HP Page 11 Line 36

"An RDMA communication service may require?" This sounds to vague and inclusive. What does SRP require of the RDMA service? That's all that should be defined in clause 4. It seems like SRP either will depend on the RDMA service's providing flow control for messages, or it will provide its own flow control. If SRP provides its own flow control, and doesn't depend on flow control from the RDMA service, then there is no reason to discuss flow control except maybe to mention that it is not required.

Open: HP Page 12 Line 40

4.5 Ordering and Reliability. Very glad to see this here. Wish it were in SAM-2.

Open: HP Page 14 Line 24

"Server address" probably should be "server identifier".

Open: HP Page 15 Line 24

Establishing multiple connections between an I,T port pair is an interesting concept, but may not be very useful, ultimately. The paragraph states that all such RDMA channels are associated with the single I_T nexus. While there is no ordering assumed between different RDMA channels (15-41), this channel independence cannot be maintained once the tasks are forwarded to the SCSI layer, where the RDMA channel allegiance of the task is forgotten, and only the I_T information is retained. Effectively, the tasks will merge from multiple transmission channels into a single queue as they transition from SRP to SCSI, and the original partial order will be replaced by a total order. Correct operation will result, but performance will suffer. Perhaps the only practical use of this construct is for the asynchronous transmission of task management requests, as in the given example.

Open: HP Page 16 Line 28

A request windowing scheme would be easier to describe than this request limit mechanism. Race conditions would not be an issue.

Open: HP Page 20 Line 4

Indirect data buffer descriptor. I don't see a good use for this facility in an IO application such as SRP, and I question its inclusion here. The channel adapter local to the memory that is to be read or written (typically the channel adapter of the Initiator) can use a scatter/gather list (SGL) to define an arbitrary virtual memory segment for an I/O buffer, and assign it a unique memory handle. This segment can then be read or written, starting at any offset, and in any order, by the target's RDMA mechanism's simply generating a series of RDMA reads or writes, always referring to the same memory handle, but using different offsets and lengths for each operation. (For example, a series of RDMA writes to increasing offsets, eventually filling the memory segment.) The direct data buffer descriptor format is sufficient for this operation, because the SGL provides for scatter/gather to buffets that start and end at arbitrary addresses in physical memory (not just page-aligned addresses), just as a traditional DMA controller does.

The only motivation I can find for the indirect model is to reduce the number of SGLs (or mapped memory regions) that the initiator's channel adapter must deal with. Unfortunately, the use of the indirect mechanism means that we must trust the target devices that share a memory region not to step on each other through misoperation or by deliberately generating invalid memory descriptors. While this is the truest form of remote DMA, because it leave the matter of address generation to the target device, it also leave the initiator exposed to target device misoperation, or worse.

I am not sufficiently familiar with IB HCA architecture to know whether such HCAs are limited to mapping only regions of contiguous pages, which would necessitate including the indirect data buffer descriptor method to support non-page-oriented IO.

Open: HP Page 25 Line 1

Login request. The statement that the login request "shall only be sent during RDMA channel establishment" seems to me overly restrictive on the RDMA model. Furthermore, I'm not sure I discern in clause 4 that the RMDA service must transport SRP login information during its own connection establishment, although this requirement is made clear in clause 5, line 14-13. It would seem quite natural to establish an RDMA connection first, and then log in SRP using the RDMA connection. (As an example, iSCSI establishes a TCP connection, and then logs it into a new or existing iSCSI session.)

Open: HP Page 25 Line 1

Login request. Need to resolve how security protocols are handled in the SRP world. The login request does not contain any provision for initiator port authentication to the target.

Open: HP Page 25 Line 32

So port identifiers are 16 bytes. But SAM-2 rev. 17 allows 8 bytes only, and iSCSI allows 260 bytes or more (still in discussion). These differences need to be rationalized. It would be best if SCSI itself would adopt a naming convention for its ports, rather than delegating this crucial task to its many transports. If SCSI were to name its ports, then SRP would only have to convey the SCSI port identifier passed down the stack by SCSI, and not make provision for conveying an identifier defined by a lower-level transport.

Open: HP Page 25 Line 32

The port identifier fields, at 16B, are too small to carry identifiers as used by iSCSI. This may prove problematical as we attempt to merge iSCSI and SRP for use with iWARP.

Open: HP Page 54 Line 1

SRP annex. Are Queue Pairs (QP) in one-to-one correspondence with IB consumers?

Open: HP Page 54 Line 23

"An IB I/O unit?contains an IB channel adapter." Why restrict it to a single channel adapter? In Figure B.3 the analogous (but nameless) initiator unit-defined by the dashed lines-is shown with multiple channel adapters. An iSCSI device is conceived as having multiple channel adapters (known informally as channel groups and in the specification as portal groups). OTOH, since an IB I/O unit is not named (it has no GUID associated with it), is there any purpose to the architecture's defining it?

Open: HP Page 54 Line 28

Figure B.2. Can I/O controllers be virtual objects?

Open: HP Page 54 Line 28

Figure B.2. What is the purpose of allowing multiple IB consumers per IB I/O controller? Is it so that multiple IB connections can be terminated within an IB I/O controller? (This relates to the question above about correspondence between QPs and IB consumers.)

Open: HP Page 54 Line 28

Figure B.2. How are shared LUs modeled? Do SRP target ports contain the "task router" function described recently by Penokie? Can two IB I/O controllers have an underlying LU in common, or is this functionality restricted to two IB consumers within the same IB I/O controller?

Open: HP Page 55 Line 9

Table B.1. IB port GUID is described as "Identifies an IB port within an IB channel adapter". This can be taken to mean that the naming scope for IB port is within a single channel adapter. I doubt that is the intention, since IB port GUIDs are globally unique. Similar comment for IB I/O controller GUID?with the further observation that IB I/O units themselves are not named, and so cannot form a naming scope. It seems to me that the first three lines of this table should read, "Identifies a _____", without qualification. It is incidental, isn't it, that an IB port is contained in an IB channel adapter (and an IB I/O controller is contained in an IB I/O unit)? The fact that the discovery process finds IB channel adapters, and then IB I/O controllers, and then IB consumers, utilizing the containment properties, seems irrelevant to describing the naming architecture, when globally unique names are used.

Open: HP Page 55 Line 25

Figure B.3. What is the object indicated by the dashed lines in the initiator model, analogous to the IB I/O unit in the target model?

Open: HP Page 55 Line 25

Figure B.3. and 56-1. Table B.2. The rules for constructing initiator ports seem entirely too lax. The text says, "Initiator port identifier should be constructed?" And then the Table indicates that GUID, for example, is the channel adapter GUID. Is there no meaning associated with the initiator port ID? Is the only design goal that the 16B port ID be globally unique? Will any GUID do at all? If so, let's be explicit about this, and let's not make any suggestions about the origin (and possible meaning) of the port name.

But it would be a better model, I think, for the "GUID" used in the initiator port ID to be associated not with the IB channel adapter, but instead with the (unnamed) SRP initiator device. It is the SRP initiator device that is associated with a naming domain such as an operating system image. IB channel adapters will be shared among operating system images, and using them as a naming domain would require that the operating system images cooperate, or that the selection of port identifier extension be delegated to the virtual machine layer, both of which are undesirable.

While we're at it, let's decouple the naming of SRP ports entirely from IB. Although SCSI really should be the layer that names its ports, let's for the moment assume that SCSI continues to delegate port naming to its transport. But let's assume further that SRP accepts the responsibility to name its ports, and doesn't delegate it further to IB. SRP can then generate its own name for SRP initiator device, with an identifier extension to make a unique port name. Analogously, SRP can name the entities identified in the figure as SRP target devices. SRP could adopt a naming scheme that uses 16B "GUIDs" analogously to IB's, and it could draw from the same naming assignment authority that IB uses. But this is not the same as saying that IB defines SRP's port names, and in fact, the description of SRP port naming would be moved from the IB annex to the main SRP text.

This change would require that during the discovery process, the IB I/O unit return the full name of the SRP port from its Service Entries table, in step 3.

This approach to naming ports brings SRP much closer to iSCSI. What is unresolved is iSCSI's use of long text strings to name iSCSI devices vs. the use of more compact GUID numbers. The two mechanisms could be combined with the introduction of a name service that dereferences string IDs to GUIDs.

Open: IBM / Tivoli 0

In my comments the notation 'Page xx' refers to all pages in the standard not roman numeral xx. All comments are editorial unless indicated with a '(T)' at the start of the comment.

Open: IBM / Tivoli #1 PDF Page 3 Tivoli comment from George Penokie

Page c - The page numbering in first part of the front matter is a,b,c, and d instead of roman numerals. This needs to be corrected.

Open: IBM / Tivoli #2 PDF Page 4 Tivoli comment from George Penokie

Page c - d - The Revision list needs to be removed before public review.

Open: IBM / Tivoli #3 PDF Page 4 Tivoli comment from George Penokie

All - All the line numbers need to be removed throughout the document.

Open: IBM / Tivoli #4 PDF Page 6 Tivoli comment from George Penokie

All - The printing date information at the bottom of every page needs to be removed.

Open: IBM / Tivoli #5 PDF Page 11 Tivoli comment from George Penokie

Page vii - Forward - the BSR number x3.269-199x is not correct for this standard. It should be 'NCITS.xxx-200x' until the actual number is assigned.

Open: IBM / Tivoli #6 PDF Page 12 Tivoli comment from George Penokie

Page viii - Line 7 - The statement 'The working draft SCSI RDMA Protocol (SRP) standard is divided into the following clauses:' should be 'The SCSI RDMA Protocol standard is divided into the following clauses:

Open: IBM / Tivoli #7 PDF Page 12 Tivoli comment from George Penokie

All - The acronym SRP should be replaced with 'SCSI RDMA Protocol' in all cases in this document.

Open: IBM / Tivoli #8 PDF Page 13 Tivoli comment from George Penokie

Page 1-2 - The following standards should be removed from the list: FC-AL, FC-PH, FC-PH-2, SPI-3, FCP, SPC, and RMC.

Open: IBM / Tivoli #9 PDF Page 16 Tivoli comment from George Penokie

Page 4 - section 3.1.1 - The last sentence implies that SRP_LOGIN_RSP is the only use for accept data. I believe this is not correct. This should be stated to be an example of accept data.

Open: IBM / Tivoli #10 PDF Page 16 Tivoli comment from George Penokie

All - The full name of a standard should always be used instead of the acronym. This should be change throughout the document.

Open: IBM / Tivoli #11 PDF Page 16 Tivoli comment from George Penokie

Page 4 - line 19 and others - when SRP is used and it is referring to this document then it should be changed to 'this standard'. Line 19 is one case where this appears to be true.

Open: IBM / Tivoli #12 PDF Page 16 Tivoli comment from George Penokie

Page 4 - section 3.1.13 - The statement 'An externally addressable object...' should be 'An.addressable object...'. The term externally implies that the addressing is outside the standard.

Open: IBM / Tivoli #13 PDF Page 16 Tivoli comment from George Penokie

Page 4 - section 3.1.15 - The last sentence implies that SRP_LOGIN_REQ is the only use for login data. If this is not correct. Then this should be stated to be an example of login data.

Open: IBM / Tivoli #14 PDF Page 16 Tivoli comment from George Penokie

Page 4 - section 3.1.15 - The statement '...server agent or consumer...' should be '....server agent or server consumer...'

Open: IBM / Tivoli #15 PDF Page 16 Tivoli comment from George Penokie

Page 4 - section 3.1 - The terms client consumer, server agent, and server consumer should be definitions is the glossary.

Open: IBM / Tivoli #16 PDF Page 17 Tivoli comment from George Penokie

Page 5 - section 3.1.22 - The statement '...server agent or consumer...' should be '...server agent or server consumer...'

Open: IBM / Tivoli #17 PDF Page 17 Tivoli comment from George Penokie

Page 5 - section 3.2 - line 34 - The acronym for SRP implies that in almost all cases SRP should be changed to 'this standard'.

Open: IBM / Tivoli #18 PDF Page 20 Tivoli comment from George Penokie

Page 8 - line 5 - The statement 'by means of' should be change to 'using'.

Open: IBM / Tivoli #19 PDF Page 20 Tivoli comment from George Penokie

Page 8 - line 44 - The statement 'established and disconnected' should be either 'established and removed' or 'connected and disconnected'. In this case I think the first option is better. The wording in the remaining document must then be made to match this change.

Open: IBM / Tivoli #20 PDF Page 20 Tivoli comment from George Penokie

Pages 8 - 11 - section 4.2 - This clause should be broken in subclauses and there should be references added between the steps in the figure and the text descriptions of those steps. This will help the reader relate the figures flow to the text.

Open: IBM / Tivoli #21 PDF Page 21 Tivoli comment from George Penokie

Page 9 - lines 7-9 - The for example text should be change to (e.g., ...).

Open: IBM / Tivoli #22 PDF Page 21 Tivoli comment from George Penokie

Page 9 - line 2 - The statement '...directed to a server and, if...' is not clear because there is a server agent and a server consumer. Which is this server supposed to be?

Open: IBM / Tivoli #23 PDF Page 21 Tivoli comment from George Penokie

Page 9 - line 5 - The statement '...identify the server with which...' is not clear because there is a server agent and a server consumer. Which is this server supposed to be?

Open: IBM / Tivoli #24 PDF Page 21 Tivoli comment from George Penokie

Page 9 - Figure 3 - line 40 - The arrow exiting to the right seems to be a dead end. Where does the flow go from there. All the other exit points are clear that one is not.

Open: IBM / Tivoli #25 PDF Page 22 Tivoli comment from George Penokie

(T) Page 10 - line 12 - This states '...the server identifier shall identify one or more SRP target ports, and the login data...'. How is it possible for a single server identifier to identify more than one SRP port? SCSI requires all target port identifiers be unique within a domain.

Open: IBM / Tivoli #26 PDF Page 22 Tivoli comment from George Penokie

Page 10 - at least lines 2-15 - The term 'server' is used by itself several times. There needs to be a qualifier on server so the reader does not assume that server equates to server agent and server consumer.

Open: IBM / Tivoli #27 PDF Page 22 Tivoli comment from George Penokie

Page 10 - lines 28 - 29 - The statement 'With SRP the reject data includes an SRP_LOGIN_REJ response (see 6.4).' Is confusing in that it implies the SRP (which is this standard) has additional requirements than what was just specific in the sentence before. That does not compute and needs to be fixed.

Open: IBM / Tivoli #28 PDF Page 22 Tivoli comment from George Penokie

Page 10 - lines 31 -32 - Is it possible for an RDMA channel to be successfully established and not operational? If not then the statement 'and is operational' should be deleted. If so then it needs to be explained how it is possible.

Open: IBM / Tivoli #29 PDF Page 22 Tivoli comment from George Penokie

Page10 - line 34 - The statement '...server agent or consumer...' should be '....server agent or server consumer...'. This needs to be looked for throughout the document and corrected.

Open: IBM / Tivoli #30 PDF Page 22 Tivoli comment from George Penokie

Page 10 - line 35 - The statement 'With SRP the accept data includes an SRP_LOGIN_RSP response (see 6.3).' Is confusing in that it implies the SRP (which is this standard) has additional requirements than what was just specific in the sentence before. That does not compute and needs to be fixed.

Open: IBM / Tivoli #31 PDF Page 22 Tivoli comment from George Penokie

Page 10 - lines 44-45 - The statement 'With SRP the login data includes an SRP_LOGIN_REQ request (see 6.2)...' Is confusing in that it implies the SRP (which is this standard) has additional requirements than what was just specific in the sentence before. That does not compute and needs to be fixed.

Open: IBM / Tivoli #32 PDF Page 22 Tivoli comment from George Penokie

Page 10 - lines 43 - 44 - The sentence 'The server agent is provided the login data from the client consumer's request in addition to RDMA communication service specific data.' is awkward. It would be better stated as 'The server agent receives the login data and RDMA communication service specific data from the client consumer's request.'

Open: IBM / Tivoli #33 PDF Page 23 Tivoli comment from George Penokie

Page 11 - line 2 - The statement 'With SRP the reject data shall contain an SRP_LOGIN_REJ response (see 6.4).' Is confusing in that it implies the SRP (which is this standard) has additional requirements than what was just specific in the sentence before. That does not compute and needs to be fixed.

Open: IBM / Tivoli #34 PDF Page 23 Tivoli comment from George Penokie

Page 11 - lines 5 - 6 - The statement 'With SRP the accept data shall contain an SRP_LOGIN_RSP response (see 6.3)...' Is confusing in that it implies the SRP.(which is this standard) has additional requirements than what was just specific in the sentence before. That does not compute and needs to be fixed.

Open: IBM / Tivoli #35 PDF Page 23 Tivoli comment from George Penokie

Page 11 - line 11 - The term 'such' should be deleted.

Open: IBM / Tivoli #36 PDF Page 23 Tivoli comment from George Penokie

Page 11 - lines 30-31 - The statement '...to deliver the message to the other consumer associated with the specified RDMA channel (the receiving consumer).' should be changed to '...to deliver the message to the receiving consumer.' There is no need to redefine what a receiving consumer is as that is done in the first paragraph of this section.

Open: IBM / Tivoli #37 PDF Page 23 Tivoli comment from George Penokie

Pages 11 - 12 - section 4.4 - This clause should be broken in subclauses. For example at least an overview, one for read RDMA, and one for write RDMA. PDF Page 24

Open: IBM / Tivoli #38 PDF Page 23 Tivoli comment from George Penokie

Page 12 - line 5 - The statement 'as well' should be deleted.

Open: IBM / Tivoli #39 PDF Page 23 Tivoli comment from George Penokie

Page 12 - line 14 - The following statement 'Such information may be communicated by an application protocol.' Does not seem relevant to this standard and should be deleted.

Open: IBM / Tivoli #40 PDF Page 23 Tivoli comment from George Penokie

Page 12 - lines 41-43 - This paragraph contains information that is not useful and should be deleted. It essentially states that RDMA communication has characteristics defined here and those not defined here are out side the scope of this standard. That is true but it is also true for every clause in this standard.

Open: IBM / Tivoli #41 PDF Page 23 Tivoli comment from George Penokie

Page 12 - line 45 - The statement 'or else' should be just 'or'.

Open: IBM / Tivoli #42 PDF Page 23 Tivoli comment from George Penokie

Page 12 - line 46 - The term 'exactly' should be deleted. There is no difference between 'exactly once' and 'once'.

Open: IBM / Tivoli #43 PDF Page 23 Tivoli comment from George Penokie

Page many - The terms Write and Read in RDMA Write and RDMA Read should not be capitalized.

Open: IBM / Tivoli #44 PDF Page 25 Tivoli comment from George Penokie

Page 13 - line 14 - The term 'satisfy' should be changed to 'meet'.

Open: IBM / Tivoli #45 PDF Page 26 Tivoli comment from George Penokie

Page 14 - line 8 - The statement 'I_T nexus' is correct but there is no reference to where one would find out more about what it is. This needs to be added.

Open: IBM / Tivoli #46 PDF Page 26 Tivoli comment from George Penokie

Page 14 - line 7 - The statement 'for its lifetime' is not clear. It should be stated as 'as long as it is established'. This ties it to the previous section. Note this assumes that the term established in 4.2 is not changed.

Open: IBM / Tivoli #47 PDF Page 26 Tivoli comment from George Penokie

Page 14 - lines 24-28 - This whole paragraph does not look like it belongs here or anywhere and it should be deleted. It appears to be attempting to defines things that are either already defined in section 4 or don't need to be defined.

Open: IBM / Tivoli #48 PDF Page 26 Tivoli comment from George Penokie

Page 14 - line 40 - The statement 'that were contained in SRP_CMD requests (see 6.8)' should be deleted as it is redundant with the statement 'outstanding SCSI tasks'.

Open: IBM / Tivoli #49 PDF Page 26 Tivoli comment from George Penokie

(T) Page 14 - line 43 - The statement '...an SRP target port should send an...' gives inadequate guidance to a target implementor. This should be required to send the SRP_T_LOGOUT or not send it. Or it should be specified when it is required to be sent and when it is not required to be sent.

Open: IBM / Tivoli #50 PDF Page 26 Tivoli comment from George Penokie

(T) Page 15 - line 4 - I recommend adding into this list a statement that other SCSI related parameters (e.g., mode pages, logs) not be effected by the disconnect. This should avoid the hole the FC has doug for itself in this area.

Open: IBM / Tivoli #51 PDF Page 27 Tivoli comment from George Penokie

Page 15 - line 18 - The statement '...operation, if accepted, may allow...' should be '...operation may allow...'. The if accepted is redundant with may.

Open: IBM / Tivoli #52 PDF Page 27 Tivoli comment from George Penokie

Page 15 - line 36 - The term 'may' should be deleted.

Open: IBM / Tivoli #53 PDF Page 27 Tivoli comment from George Penokie

Page 15 - lines 36-40 - the format of the e.g is incorrect. It should be...'standards (e.g., ...)'.

Open: IBM / Tivoli #54 PDF Page 27 Tivoli comment from George Penokie

Page 15 - line 49 - The statement 'as well as' should be 'or'.

Open: IBM / Tivoli #55 PDF Page 28 Tivoli comment from George Penokie

Page 16 - line 3 - The term 'initiation' should be 'start' or 'beginning'.

Open: IBM / Tivoli #56 PDF Page 28 Tivoli comment from George Penokie

Page 16 - line 5 - The term 'all' should be 'the'.

Open: IBM / Tivoli #57 PDF Page 28 Tivoli comment from George Penokie

Page 16 - lines 7-8 - I am not aware of a SCSI command that specifies that status not be returned. If there is such a thing then an e.g., would be helpful. If there is no such thing then this item should be deleted.

Open: IBM / Tivoli #58 PDF Page 28 Tivoli comment from George Penokie

Page 16 - line 18 - What is the 'it' referring to? The 'it' needs to be replaced with whatever 'it' is.

Open: IBM / Tivoli #59 PDF Page 28 Tivoli comment from George Penokie

Page 16 - line 23 - The term 'might' should be 'may'.

Open: IBM / Tivoli #60 PDF Page 28 Tivoli comment from George Penokie

Page 16 - line 23 - What is the 'it' referring to? The 'it' needs to be replaced with whatever 'it' is.

Open: IBM / Tivoli #61 PDF Page 28 Tivoli comment from George Penokie

Page 16 - line 24 - The statement '...to at most one...' seems redundant. It should be '...to one...'.

Open: IBM / Tivoli #62 PDF Page 28 Tivoli comment from George Penokie

Page 16 - lines 28-29 - The statement '...present in most information units...' is troublesome. There either needs to be a list of the IUs that have the field or a reference to a place that would tell my which IUs have or do not have the field.

Open: IBM / Tivoli #63 PDF Page 28 Tivoli comment from George Penokie

Page 16 and others? - The when to use small caps rule is not being followed here. The rule is that small caps are only used when the field is being named (e.g., xxx field would

have the xxx in small caps). When contents of the field is being called out it is not in small caps (e.g. request limit and request limit delta are both signed...').

Open: IBM / Tivoli #64 PDF Page 28 Tivoli comment from George Penokie

Page 16 - line 49 - The sentence starting with 'An SRP port shall not specify a negative...' should be a separate item in the list.

Open: IBM / Tivoli #65 PDF Page 28 Tivoli comment from George Penokie

(T) Page 16 - section 5.3 - This section on flow control seems overly complex for what appears to be actually needed. The only SRP request that even needs to have multiple outstanding requests in the command. All others should not be streamed but should be interlocked and some should be allowed to occur at any time. This all needs to be looked at to make sure the design point is what we really want.

Open: IBM / Tivoli #66 PDF Page 29 Tivoli comment from George Penokie

Page 17 - Figure 4 - The way the arrows are pointing for the virtual address implies that it is not the address of the first byte of the memory segment. It currently implies that it is the space from the memory handle to the beginning of the memory segment which is the memory region. It is also not clear as to what the boundaries are of the memory region. The current drawing implies it is only the area above the memory segment. I do not believe that is correct so it needs to be fixed.

Open: IBM / Tivoli #67 PDF Page 29 Tivoli comment from George Penokie

Page 17 - line 26 - There is no indication as to what kind of value the memory handle is. This would normally not be a problem except that the other two fields to explicitly indicate that they are unsigned integer values. I generally consider all fields to be unsigned integers but in this case there is doubt cast about that assumption.

Open: IBM / Tivoli #68 PDF Page 30 Tivoli comment from George Penokie

Page 18 - line 1 - The statement 'A SRP...' should be 'An SRP...' This needs to be checked for throughout the document and corrected.

Open: IBM / Tivoli #69 PDF Page 30 Tivoli comment from George Penokie

Page 18 - line 3 - The statement '...within its memory segment.' should be '...within the memory segment.'.

Open: IBM / Tivoli #70 PDF Page 30 Tivoli comment from George Penokie

Page 18 - line 2 - The statement 'SRP target ports shall only issue the appropriate type of RDMA operation for a memory descriptor' appears to be restating what was stated in the previous sentence and therefore should be deleted. The sentence would then read 'SRP target ports shall ensure that each RDMA operation...'.

Open: IBM / Tivoli #71 PDF Page 30 Tivoli comment from George Penokie

Page 18 - line 3 - There needs to be a connection between the text above the a.b.c list and the list. Something like 'segment by using the following rules:'.

Open: IBM / Tivoli #72 PDF Page 30 Tivoli comment from George Penokie

Page 18 - lines 15-17 - The sentences 'The format of each data buffer descriptor is specified by a format code value. Some data buffer descriptor format code values use the contents of a count field to further specify the data buffer descriptor format.' should be deleted as the information is a duplicate of what is in table 2.

Open: IBM / Tivoli #73 PDF Page 30 Tivoli comment from George Penokie

Page 18 - table 2 - line 35 - footnote c - There statement 'and and' should be just 'and' and there is not period at the end of the sentence.

Open: IBM / Tivoli #74 PDF Page 30 Tivoli comment from George Penokie

Page 18 - table 2 - line 27 - The equation $20+16*\text{count}$ should be change to $20 + 16 \times \text{count}$. This change from * to x should be make throughout the document.

Open: IBM / Tivoli #75 PDF Page 30 Tivoli comment from George Penokie

Page 18 - table 2 - footnote b - This should have a reference from the cell with 'count' in it.

Open: IBM / Tivoli #76 PDF Page 30 Tivoli comment from George Penokie

Page 18 - lines 43-45 - The sentence 'An SRP initiator port shall not specify a data buffer descriptor format that was not indicated in the REQUIRED BUFFER FORMATS field value for that RDMA channel.' does not make sense. How can the initiator port be indicating the buffer formats in the REQUIRED BUFFER FORMATS field and at the same time not specifying the buffer formats in the REQUIRED BUFFER FORMATS field that were not indicated in the in the REQUIRED BUFFER FORMATS field. This is circular and needs to be fixed.

Open: IBM / Tivoli #77 PDF Page 30 Tivoli comment from George Penokie

Page 18 - line 41 - There should be a reference to table 2 as follows 'data buffer descriptor formats (see table 2)'.

Open: IBM / Tivoli #78 PDF Page 30 Tivoli comment from George Penokie

Page 18 - line 47 - The statement '...RDMA channel and...' should be '...RDMA channel request and...'.

Open: IBM / Tivoli #79 PDF Page 30 Tivoli comment from George Penokie

Page 18 - line 49 - The statement '...RDMA channel and...' should be '...RDMA channel request and...'.

Open: IBM / Tivoli #80 PDF Page 30 Tivoli comment from George Penokie

Page 18 - line 40 - There should be a reference to table 3 as follows 'The REQUIRED BUFFER FORMATS field (see table 3)...'.

Open: IBM / Tivoli #81 PDF Page 31 Tivoli comment from George Penokie

Page 19 - line 4 - I believe the 'and' should be an 'or'. I don't believe a target port would do both IU at the same time.

Open: IBM / Tivoli #82 PDF Page 31 Tivoli comment from George Penokie

Page 19 - lines 3-4 - There should be a reference to table 3 as follows 'The SUPPORTED BUFFER FORMATS field (see table 3)...'.

Open: IBM / Tivoli #83 PDF Page 31 Tivoli comment from George Penokie

Page 19 - line 8 - The statement '...contents of the REQUIRED BUFFER...' should be '...contents of both the REQUIRED BUFFER...'.

Open: IBM / Tivoli #84 PDF Page 31 Tivoli comment from George Penokie

(T) Page 19 - line 18 and line 28 - Why is that when the IDBD bit and the DDBD bit is set to zero it is a should instead of a shall? This should be changed to a shall unless there is some good reason.

Open: IBM / Tivoli #85 PDF Page 31 Tivoli comment from George Penokie

Page 19 - note 2 - This note should be a note. It should be part of the main text. It should also be restated as: 'The length of requests sent by an SRP initiator port, as determined by the data buffer descriptor formats, shall be limited to the MAXIMUM INITIATOR TO TARGET IU LENGTH field (see xxx) returned in the SRP_LOGIN_RSP response.'

Open: IBM / Tivoli #86 PDF Page 31 Tivoli comment from George Penokie

(T) Page 19 - lines 39 - 40 - The sentence 'SRP target ports are not required to check the contents of the count field.' should be changed to 'SRP target ports shall ignore the contents of the count field.'

Open: IBM / Tivoli #87 PDF Page 31 Tivoli comment from George Penokie

(T) Page 19 - lines 44 - 45 - The sentence 'SRP target ports are not required to check the contents of the count field.' should be changed to 'SRP target ports shall ignore the contents of the count field.'

Open: IBM / Tivoli #88 PDF Page 31 Tivoli comment from George Penokie

Page 19 and others - line 39 and others - The term 'count field' is used in many places. First there are two of them so it should be 'count fields'. Second is not clear that these are the count fields in the SRP_CMD request. I recommend changing 'count field' to 'count fields in the SRP_CMD request' in all places in the main body text.

Open: IBM / Tivoli #89 PDF Page 32 Tivoli comment from George Penokie

Page 20 - line 8 - The statement 'count field' should be 'DATA-OUT BUFFER DESCRIPTOR COUNT field (or DATA-IN BUFFER DESCRIPTOR COUNT field)'

Open: IBM / Tivoli #90 PDF Page 32 Tivoli comment from George Penokie

Page 20 - line 12 - A reference to table 5 should be added to the end of the paragraph.

Open: IBM / Tivoli #91 PDF Page 32 Tivoli comment from George Penokie

Page 20 - table 4 - line 29 - Footnote a - It's not clear which count field is being referred to. Is it the one in table 2 or the ones in the SRP_CMD_ request. This needs to be fixed with the proper terminology and a reference to the correct place.

Open: IBM / Tivoli #92 PDF Page 32 Tivoli comment from George Penokie

Page 20 - line 34 - The statement 'The DATA LENGTH field of the INDIRECT TABLE MEMORY DESCRIPTOR field value contains...' is not correct. It should be 'The DATA LENGTH field of the memory descriptors in the indirect table contains...'

Open: IBM / Tivoli #93 PDF Page 32 Tivoli comment from George Penokie

Page 20 - line 39 - The sentence 'SRP target port behavior when the TOTAL LENGTH field contains any other value is vendor specific.' should be moved to the end of the paragraph and restated as 'If the TOTAL LENGTH field value is not equal to the sum of the DATA LENGTH field values the SRP target port's behavior shall be vendor specific.'

Open: IBM / Tivoli #94 PDF Page 32 Tivoli comment from George Penokie

Page 20 - line 42 - It's not clear which count field is being referred to. Is it the one in table 2 or the ones in the SRP_CMD_ request. This needs to be fixed with the proper terminology and a reference to the correct place.

Open: IBM / Tivoli #95 PDF Page 32 Tivoli comment from George Penokie

Page 20 - line 47 - This should be the start of a new subclause. Something like 'SRP target port indirect data restrictions'. PDF Page 33

Open: IBM / Tivoli #96 PDF Page 32 Tivoli comment from George Penokie

Page 21 - line 7 - This paragraph should be the start of a new subclause titled something like 'Examples of Indirect data buffers'.

Open: IBM / Tivoli #97 PDF Page 32 Tivoli comment from George Penokie

(T) Page 20 and 21 - The possibility of having both a data-in and a data-out buffer is not described here. Why not? This needs to be fixed.

Open: IBM / Tivoli #98 PDF Page 32 Tivoli comment from George Penokie

Page 21 - lines 12 and 13 - The term 'might' should be changed to 'may'. This should be done throughout this document.

Open: IBM / Tivoli #99 PDF Page 35 Tivoli comment from George Penokie

Page 23 - line 48 - The statement 'A requestor shall provide a TAG value in each SRP request that is unique among all of the requestor's outstanding SRP requests with a particular responder. A responder shall copy the TAG value from each SRP request to the SRP request's SRP response. Responders are not required to check whether the TAG values of outstanding SRP requests are unique.' should be 'Each SRP request shall contain a TAG value that is unique among all of the outstanding SRP requests from a particular SRP initiator port. Each SRP response shall contain a copy of the TAG value from the corresponding SRP request. Responders are not required to check whether the TAG values are unique.'

Open: IBM / Tivoli #100 PDF Page 37 Tivoli comment from George Penokie

Page 25 - line 2 - The term 'conveys' should be changed to 'sends'.

Open: IBM / Tivoli #101 PDF Page 37 Tivoli comment from George Penokie

Page 25 - line 42 - The statement '...wishes to send...' should be changed to '...sends...'

Open: IBM / Tivoli #102 PDF Page 37 Tivoli comment from George Penokie

Page 25 - line 42 - The statement '....be 64 or larger.' should be '....be greater than or equal to 64.' or '...be greater than 63.'

Open: IBM / Tivoli #103 PDF Page 38 Tivoli comment from George Penokie

Page 26 - lines 1-2 - The statement 'The MULTI-CHANNEL ACTION field identifies how an SRP target port treats any existing RDMA channel associated with the same I_T nexus. The MULTI-CHANNEL ACTION field is defined in table 10.' should be changed to 'The MULTI-CHANNEL ACTION field (see table 10) indicates how an SRP target port handles existing RDMA channels.associated with the same I_T nexus.'

Open: IBM / Tivoli #104 PDF Page 38 Tivoli comment from George Penokie

Page 26 - table 10 - All the codes except for the 2 that are defined need to be listed as reserved. The row should have '02h - FFh' in the action column and 'reserved' in the description column.

Open: IBM / Tivoli #105 PDF Page 39 Tivoli comment from George Penokie

Page 27 - line 2 - The term 'conveys' should be changed to 'sends'.

Open: IBM / Tivoli #106 PDF Page 40 Tivoli comment from George Penokie

Page 28 - lines 1-2 - The statement 'MULTI-CHANNEL RESULT identifies how the SRP target port treated existing RDMA channels associated with the same I_T nexus. Table 12 defines this field.' should be changed to 'The MULTI-CHANNELRESULT field (see table

12) indicates how an SRP target port handles existing RDMA channels associated with the same I_T nexus.'

Open: IBM / Tivoli #107 PDF Page 40 Tivoli comment from George Penokie

Page 28 - table 12 - All the codes except for the 3 that are defined need to be listed as reserved. The row should have '03h - FFh' in the action column and 'reserved' in the description column.

Open: IBM / Tivoli #108 PDF Page 42 Tivoli comment from George Penokie

Page 30 - line 4 - The statement '...failed, rendering it non-operational.' should be changed to '...failed.'

Open: IBM / Tivoli #109 PDF Page 43 Tivoli comment from George Penokie

Page 31 - line 4 - The statement '...failed, rendering it non-operational.' should be changed to '...failed.'

Page 32 - line 2 - The term 'conveys' should be changed to 'sends'.

Open: IBM / Tivoli #110 PDF Page 44 Tivoli comment from George Penokie

Page 32 - lines 37-38 - The statement '...logical unit component of the nexus for the task management request.' should be changed to '...logical unit to which to send task management request.'

Open: IBM / Tivoli #111 PDF Page 46 Tivoli comment from George Penokie

Page 34 - line 2 - The term 'conveys' should be changed to 'sends'.

Open: IBM / Tivoli #112 PDF Page 46 Tivoli comment from George Penokie

Page 34 - 35 - Table 20 - This table splits up a paragon and worse a sentence. This needs to be fixed.

Open: IBM / Tivoli #113 PDF Page 46 Tivoli comment from George Penokie

Page 34 - table 20 - The notation 'do' and 'di' are confusing when placed into a sentence (as in the footnotes). They should be changed to 'x' and 'y'.

Open: IBM / Tivoli #114 PDF Page 48 Tivoli comment from George Penokie

Page 36 - line 2 - The term 'conveys' should be changed to 'sends'.

Open: IBM / Tivoli #115 PDF Page 48 Tivoli comment from George Penokie

Page 36 - line 6 - The statement '...message capable of containing...' should be changed to '...message containing...'.

Open: IBM / Tivoli #116 PDF Page 49 Tivoli comment from George Penokie

Page 37 - The statement 'set to 1' should be 'set to one' and the statement 'set to 0' should be 'set to zero' in all cases throughout this document.

Open: IBM / Tivoli #117 PDF Page 49 Tivoli comment from George Penokie

Page 37 - line 44 - The statement 'are not reliable and' should be deleted as it contains no useful information.

Open: IBM / Tivoli #118 PDF Page 50 Tivoli comment from George Penokie

Page 38 - line 3 - Add a reference to the RSP_CODE values table (table 24) at the end of this paragraph.

Open: IBM / Tivoli #119 PDF Page 50 Tivoli comment from George Penokie

(T) Page 38 - lines 15-17 - The statement 'If DOUNDER is set to 1, a transfer that did not use the entire data-out buffer was performed and the value of DATA-OUT RESIDUAL COUNT shall be equal to: data-out buffer length - highest offset of any data-out byte transferred - 1' needs to be changed to 'If DOUNDER is set to one and a transfer that did not fill the entire data-out buffer was performed the value of DATA-OUT RESIDUAL COUNT is defined as follows: DATA-OUT RESIDUAL COUNT = (data-out buffer length) - (highest offset of any data-out byte transmitted + 1)'

Open: IBM / Tivoli #120 PDF Page 50 Tivoli comment from George Penokie

(T) Page 38 - lines 22-23 - The statement 'DATA-OUT RESIDUAL COUNT shall be equal to: data-out transfer length required by command - data-out buffer length' needs to be changed to 'The DATA-OUT RESIDUAL COUNT is defined as follows: DATA-OUT RESIDUAL COUNT = (Transfer length required by command) - (data-out buffer length)'

Open: IBM / Tivoli #121 PDF Page 50 Tivoli comment from George Penokie

(T) Page 38 - lines 34-36 - The statement 'If DIUNDER is set to 1, a transfer that did not fill the entire data-in buffer was performed and the value of DATA-IN RESIDUAL COUNT shall be equal to: data-in buffer length - highest offset of any data-in byte transferred - 1' needs to be changed to 'If DIUNDER is set to one and a transfer that did not fill the entire data-in buffer was performed the value of DATA-IN RESIDUAL COUNT is defined as follows: DATA-IN RESIDUAL COUNT = (data-in buffer length) - (highest offset of any data-in byte transmitted + 1)'

Open: IBM / Tivoli #122 PDF Page 50 Tivoli comment from George Penokie

(T) Page 38 - lines 41-43 - The statement 'DATA-IN RESIDUAL COUNT shall be equal to: data-in transfer length required by command - data-in buffer length' needs to be changed to 'The DATA-IN RESIDUAL COUNT is defined as follows: DATA-IN RESIDUAL COUNT = (Transfer length required by command) - (data-in buffer length)'

Open: IBM / Tivoli #123 PDF Page 51 Tivoli comment from George Penokie

Page 39 - line 1 - The term 'certain' should be deleted.

Open: IBM / Tivoli #124 PDF Page 51 Tivoli comment from George Penokie

(T) Page 39 - lines 30 - 41 - All this should be deleted and replaced with 'The SENSE DATA field contains the autosense data specified by the SCSI Primary Commands-2 standard. The proper SENSE DATA shall be presented when the SCSI status byte of CHECK CONDITION is presented as specified by the SCSI Primary Commands-2 standard. If no conditions requiring the presentation of SCSI sense data have occurred, the SENSE DATA field shall not be included in the SRP_RSP response and the RSPVALID bit shall be zero. SRP devices shall perform autosense.'

Open: IBM / Tivoli #125 PDF Page 53 Tivoli comment from George Penokie

Page 41 - line 2 - The term 'conveys' should be changed to 'sends'.

Open: IBM / Tivoli #126 PDF Page 54 Tivoli comment from George Penokie

(T) Page 42 - lines 3-13 - All this should be deleted and replaced with the following 'The SENSE DATA field contains sense data specified by the SCSI Primary Commands-2 standard.'. This is AER not a check condition they are different things. The only thing that should be stated here is that sense data is returned.

Open: IBM / Tivoli #127 PDF Page 56 Tivoli comment from George Penokie

Page 44 and others - line 16 and others - The term 'set to 0' and 'set to 1' should be 'set to one' and 'set to zero'.

Open: IBM / Tivoli #128 PDF Page 56 Tivoli comment from George Penokie

Page 44 - line 19 - The term 'all' should be deleted as it is redundant.

Open: IBM / Tivoli #129 PDF Page 58 Tivoli comment from George Penokie

Page 46 - figure A.2 and A.3 - line 15 and 43 - The statement '(SRP initiator' should be '(SRP initiator port)'.

Open: IBM / Tivoli #130 PDF Page 58 Tivoli comment from George Penokie

Page 46 and others- lines 22-26 and others - The 1,2,3 list should not have line spaces between numbered items. This should be fixed in all cases

Open: IBM / Tivoli #131 PDF Page 64 Tivoli comment from George Penokie

Page 52 - line 20 - The term 'executes' should be changed to 'processes'.

Open: IBM / Tivoli #132 PDF Page 64 Tivoli comment from George Penokie

Page 52 - line 32 - The statement '...a device or component...' should be 'an IB device or component...'.

Open: IBM / Tivoli #133 PDF Page 65 Tivoli comment from George Penokie

Page 53 - line 20 - There seems to be no definition of what a 'connection manager' is. This should, at least, be added to the glossary.

Open: IBM / Tivoli #134 PDF Page 65 Tivoli comment from George Penokie

Page 53 - section B.3.2 - The abbreviation IOC needs to be added to the list.

Open: IBM / Tivoli #135 PDF Page 67 Tivoli comment from George Penokie

Page 55 - lines 1-2 - The sentence 'The IB more IB LIDs and IB GUIDs corresponding to an IB port GUID or IB channel adapter GUID.' does not seem to be a complete sentence and is not clear as to what it is trying to state. This needs to be fixed.

Open: IBM / Tivoli #136 PDF Page 68 Tivoli comment from George Penokie

(T) Page 56 - line 2 - Why is the should not a shall. I believe it should be changed to a shall.

Open: IBM / Tivoli #137 PDF Page 68 Tivoli comment from George Penokie

Page 56 - line 15 - The statement '...field should an IB GUID...' should be '...field should be an IB GUID...'.

Open: IBM / Tivoli #138 PDF Page 68 Tivoli comment from George Penokie

Page 56 - line15 - The statement '...port, e.g. the...SRP initiator port.' should be '...port (e.g., the...SRP initiator port)'.

Open: IBM / Tivoli #139 PDF Page 68 Tivoli comment from George Penokie

Page 56 - lines 15-16 - The statement 'the IB channel adapter GUID for an IB channel adapter used the SRP initiator port.' is not very clear as to what it is. This needs to be fixed.

Open: IBM / Tivoli #140 PDF Page 68 Tivoli comment from George Penokie

Page 56 - line 20 - There is not clue as what a 'device management agent' is. This could be fixed by replacing 'device management agent' with the more generic term 'entity'.

Open: IBM / Tivoli #141 PDF Page 68 Tivoli comment from George Penokie

Page 56 - line 22 - The term 'indicated' is confusing in this sentence. A better term would be 'identified'.

Open: IBM / Tivoli #142 PDF Page 68 Tivoli comment from George Penokie

Page 56 - line 39 - The term 'indicated' is confusing in this sentence. A better term would be 'identified'.

Open: IBM / Tivoli #143 PDF Page 68 Tivoli comment from George Penokie

Page 56 - lines 41-42 - This sentence seems out of place here. I should be moved to right after figure B.3.

Open: IBM / Tivoli #144 PDF Page 68 Tivoli comment from George Penokie

Page 56 - line 49 and page 57 - line 1 - The term ' IB I/O ' has been split across lines (and in this case across pages) at the /. This needs to be fixed so it will not happen. There is an option in frame that if selected will prevent this. It should be enabled for this document.

Open: IBM / Tivoli #145 PDF Page 69 Tivoli comment from George Penokie

Page 57 - line 34 - The 'it' at the beginning of the sentence should be replaced with whatever the 'it' is.

Open: IBM / Tivoli #146 PDF Page 69 Tivoli comment from George Penokie

Page 57 - line 46 and page 58 - line 1 - Why is the receive data-out mapped to RDMA requests and send data-in mapped to RDMA WRITE packets? One is a 'request' the other a 'packet' this seems strange shouldn't they be the same?

Open: IBM / Tivoli #147 PDF Page 73 Tivoli comment from George Penokie

Page 61 - table B.8 - line 31 - The statement '(binary zeros)' should be '(i.e., binary zeros)'.

Open: IBM / Tivoli #148 PDF Page 73 Tivoli comment from George Penokie

SRP does not define any format for the 3rd party device identifier for third party reservations. This needs to be added to comply with requirements in SPC-3.

Open: IBM / Tivoli #149

p50 line 11. "See 4x1" is a typo. I think this should be "See 4.4".

Open: IBM / Tivoli #150

p50 line 14. "Sever" should be "server".

Open: IBM / Tivoli #151

p50 line 35. "Sever" should be "server".

Open: IBM / Tivoli #152

p57 section B.6.5. The descriptions for data-in and data-out are not symmetrical. One is described in terms of an "RDMA READ Request" and the other in terms of "one or more RDMA WRITE packets". I think the rules are the same for both data-in and data-out (please let me know if I'm incorrect in that assumption). Describing them differently implies that they are somehow different, and generates unnecessary confusion. (This is the same as Tivoli comment number 146).

Open: intel0001 Sect:1 Pg:1 Ln:32

Transport protocol s/b 'SCSI Protocol' Suggest shading box to clarify what we're doing in this spec

Open: intel0002 Sect:1 Pg:1 Ln:37

Remove 'Physical'

Open: intel0003 Sect:3.1.11 Pg:4 Ln:27

inconsistent use of 'the' before SRP - suggest no 'the'

Open: intel0004 Sect:3.1.14 Pg:4 Ln:34

Is it necessary to specify field size in definition?

Open: intel0005 Sect:3.1.15 Pg:4 Ln:35

'Application protocol' is not defined, thus what constitutes app proto data is unclear

Open: intel0006 Sect:3.1.16 Pg:4 Ln:39

Key feature is that data placement is under control of receiver

Open: intel0007 Sect:3.1.17 Pg:4 Ln:41

'path' is a poor term, implies routing

Open: intel0008 Sect:3.1.17 Pg:4 Ln:43 C

'a transport protocol or service' - which is it? There appears to be an abstraction layering problem Using 'service' to define a service suggests we don't have a clean definition - we don't

Open: intel0009 Sect:3.1.23 Pg:5 Ln:6

rewrite as 'specific to an RDMA comm service'

Open: intel0010 Sect:3.1.27 Pg:5 Ln:12 C

TP ID ' within an RDMA comm service' - another abstraction issue - what is a service?

Open: intel0011 Sect:3.1.28 Pg:5 Ln:15

Any reason to spec field size?

Open: intel0012 Sect:3.3.9 Pg:6 Ln:12

reported as AN error

Open: intel0013 Sect:4 Pg:8 Ln:1 C

Clause 4 alternates between being a generic overview of RDMA, including discussion of features not used by SRP (e.g., solicited events in 4.3), and being normative (numerous SHALLs), which seems out of place in a clause entitled '...model'

Suggest separating the architectural model from the normative.

Open: intel0014 Sect:4.2 Pg:8 Ln:46

Seems redundant to Line 10 above.

Open: intel0015 Sect:4.2 Pg:9 Ln:2 C

Model is unclear: "A client consumer requests that the RDMA communication service establish an RDMA channel."

But RDMA_CS is defined as a protocol. The sense should be that the client requests a SERVICE PROVIDER establish a channel.

Open: intel0016 Sect:4.2 Pg:9 Ln:2

"The request is directed to a server" - Ambiguous

There are several standard meanings for 'server' - a piece of HW, a process, etc.

Open: intel0017 Sect:4.2 Pg:9 Ln:29

Should we add 'and validate' to 'Determine'?

Open: intel0018 Sect:4.2 Pg:9 Ln:50 C

We need a similar diagram for channel teardown.

Open: intel0019 Sect:4.2 Pg:10 Ln:17

(Many places in this clause) Some formatting is needed to set off model-specific terms such as "channel establishment failure response" - suggest bold or small caps. This would making parsing and understanding much easier.

Open: intel0020 Sect:4.2 Pg:10 Ln:22

Given the vague definition of RDMA CS, it's hard to tell what 'internal to the RDMA communication service' does or does not mean.

Open: intel0021 Sect:4.2 Pg:10 Ln:26

"An RDMA channel rejected response returns reject data" s/b "Rejection" data

Open: intel0022 Sect:4.2 Pg:10 Ln:28

'With SRP the reject data includes' - near duplicate of page 11, Line 2

Open: intel0023 Sect:4.2 Pg:10 Ln:30

'service specific data' s/b 'service-specific data' (global replace)

Open: intel0024 Sect:4.2 Pg:10 Ln:42

'requests that are acceptable to the RDMA communication service shall be passed to the server agent.' (SHALL in model clause.) What does it mean to be acceptable to the service? As there is no mapping of 'Server Agent' to any entity, on what is this requirement placed? Can this requirement be stated in SRP or Annex B -specific terms?

Open: intel0025 Sect:4.2 Pg:11 Ln:2

'reject(ion) data shall contain an SRP_LOGIN_REJ...' (SHALL) Do we need a subclause similar to '4.5 Ordering and Reliability' to capture size issues, so we can specify requirements on underlying interconnects? (e.g., Must be able to return _REJ as part of connection establishment protocol.)

Open: intel0026 Sect:4.2 Pg:11 Ln:5

'accept data' s/b 'acceptance data'

Open: intel0027 Sect:4.2 Pg:11 Ln:9

It is unclear how an RDMA comm svc requests that a channel be disconnected.

Open: intel0028 Sect:4.2 Pg:11 Ln:9

Need to discuss the case of a channel being destroyed due to an error.

Open: intel0029 Sect:4.2 Pg:11 Ln:13

'A disconnect request causes an RDMA channel to become non-operational.' Is this a request by a consumer to the local CS provider, or to the remote client, server agent,...?

Open: intel0030 Sect:4.2 Pg:11 Ln:16

'may or may not' Since 'May' and 'May Not' are both defined to be equivalent to 'May or May Not', there appears to be no reason to include both. (global)

Open: intel0031 Sect:4.2 Pg:11 Ln:16

Suggest: 'The completion status of operations... is indeterminate.'

Open: intel0032 Sect:4.2 Pg:11 Ln:13

'disconnect request' s/b 'disconnection request' (global)

Open: intel0033 Sect:4.3 Pg:11 Ln:18

'An RDMA channel may allow its consumers to exchange messages.' One that did not would be useless for the present case, wouldn't it?

Open: intel0034 Sect:4.3 Pg:11 Ln:21

'may provide normal and solicited message reception notification,' Since not used by SRP, why included?

Open: intel0035 Sect:4.3 Pg:11 Ln:24

'providing the following to an RDMA communication service' Again, CS model issue - how do you provide this to a protocol?

Open: intel0036 Sect:4.4 Pg:12 Ln:23 C

'An RDMA communication service is not required to provide a way for a requesting consumer to determine whether the data has been written into the specified range of addresses in registered memory.' If the target does not know whether a write has completed, how does it know when to send status, and whether status is good or not?

Open: intel0037 Sect:4.5 Pg:12 Ln:45

'or else disconnect the RDMA channel.' 'destroy' is a better term to reflect the error case.

Open: intel0038 Sect:4.5 Pg:13 Ln:15

disconnect s/b destroy

Open: intel0039 Sect:5.1.1 Pg:14 Ln:20

'An SRP target port shall not accept a new RDMA channel unless its SRP target port identifier matches the value in the SRP_LOGIN_REQ request.' As we have not defined 'match', do we need to explicitly allow wildcards?

Open: intel0040 Sect:5.1.1 Pg:14 Ln:26

Additional - spelling

Open: intel0041 Sect:5.1.1 Pg:14 Ln:31

'Prior to requesting that an RDMA channel be disconnected, an SRP initiator port may send an SRP_I_LOGOUT' s/b SHALL send

Open: intel0042 Sect:5.1.1 Pg:14 Ln:44

'Prior to requesting that an RDMA channel be disconnected, an SRP target port should send an SRP_T_LOGOUT request' s/b 'SHALL send'

Open: intel0043 Sect:5.1.3 Pg:15 Ln:16

'Following acceptance of a login specifying single RDMA channel operation that single RDMA channel' Add comma after 'operation'

Open: intel0044 Sect:5.1.3 Pg:15 Ln:22

'shall not accept such a login' What _REJ reason code is returned?

Open: intel0045 Sect:5.1.3 Pg:15 Ln:27

identifoer

Open: intel0046 Sect:5.1.3 Pg:15 Ln:36

Break E.g. sentence into two or more sentences, or write as a note.

Open: intel0047 Sect:5.3 Pg:16 Ln:24

Suggest creating 5.3.1 Initiator Requests, and 5.3.2 Target Requests, to discuss separately. _Many_ reviewers have become confused with 'SRP target ports shall limit...' Add pointer to Table 7 and emphasis that these are target-initiated SRP requests, _not_ RDMA requests.

Open: intel0048 Sect:5.3 Pg:16 Ln:27

'credit based' s/b 'credit-based'

Open: intel0049 Sect:5.3 Pg:16 Ln:40 C

'An SRP initiator port shall not send an SRP request on any RDMA channel whose REQUEST LIMIT has a value less than or equal to zero.' What is Target Port response to this?

Open: intel0050 Sect:5.3 Pg:16 Ln:42 C

'To ensure that task management requests may be sent, an SRP initiator port may choose to send commands only when REQUEST LIMIT is greater than one' Since Targ-Port can remove an arbitrary number of credits at any time, Init Port can be prohibited from performing Task Mgmt or sending SRP_I_LOGOUT.

Open: intel0051 Sect:5.3 Pg:16 Ln:46 C

'An SRP initiator port shall add...whenever it receives an information unit on that RDMA channel' What does 'receive' mean? Received at what layer? There may be a significant delay between receiving and reading.

Open: intel0052 Sect:5.3 Pg:16 Ln:46 C

Target Port maintains, implicitly or explicitly, a value representing its view of the number of free request contexts (Call this Target Request Limit TRL) When there are no requests outstanding, TRL will be equal to the initiator's REQUEST LIMIT (IRL).

The description in 5.3 only describes IRL, but TRL may differ from IRL, and there is no definition of when IRL is changed. Specifically, when TargPort sends SRP_CRED_REQ with a negative value, when does TP update TRL? It only makes sense to update upon receipt of SRP_CRED_RSP, but that is not stated.

Rewrite to describe with state variable at IP and at TP, and rules for updating.

Open: intel0053 Sect:5.3 Pg:16 Ln:46 C

When can TPort be sure that IPort has seen the REQ_LIMIT_DELTA in an SRP_RSP?
(Receipt of transport ACK is not enough)

Open: intel0054 Sect:5.3 Pg:17 Ln:1

'An SRP target port shall not specify a negative value of REQUEST LIMIT DELTA that might cause REQUEST LIMIT to drop below -2^{31} ' Given wrapping, it's impossible to drop below -2^{31} in 32-bit 2's comp. Would -2^{16} be negative enough?

Open: intel0055 Sect:5.3 Pg:17 Ln:1 C

'An SRP target port shall account for all possible race conditions to meet these requirements.' Remove this sentence.

Open: intel0056 Sect:5.4.1 Pg:17 Ln:6

'memory segment' and 'memory region' need to be defined before use.

Open: intel0057 Sect:5.4.1 Pg:17 Ln:6

'identifies the byte address' Isn't the interpretation of a VA up to the particular interconnect/transport?

Open: intel0058 Sect:5.4.1 Pg:17 Ln:6

(Memory Handle) 'The SRP initiator port shall use this value to locate the region.' It doesn't appear to be within our scope to define initiator memory controller implementations. Remove this sentence.

Open: intel0059 Sect:5.4.1 Pg:17 Ln:36

Drawing seems to indicate that memory addresses increase moving downward. Should be explicit.

Open: intel0060 Sect:5.4.1 Pg:18 Ln:2

'SRP target ports shall only issue the appropriate type of RDMA operation for a memory descriptor,' Add: 'depending on whether the descriptor was a DATA-IN or DATA-OUT descriptor'

Open: intel0061 Sect:5.4.1 Pg:18 Ln:4

'a) The RDMA operations VIRTUAL ADDRESS shall be greater' Should specify STARTING address.

Although VIRTUAL ADDRESS is a field name in Table 1, the field may have a different name in a particular interconnect's request format. Should not be in CAPS.

Open: intel0062 Sect:5.4.1 Pg:18 Ln:4

'Some data buffer descriptor format code values' s/b 'descriptor formats'

Open: intel0063 Sect:5.4.1 Pg:18 Ln:17

'use the contents of a count field to further specify the data buffer descriptor format.' specify -> describe

Open: intel0064 Sect:5.4.1 Pg:18 Ln:17 C

'use the contents of a count field to further specify the data buffer descriptor format.' 'count' is essentially a pointer to another field someplace, but this is far from obvious when reading. Suggest we define a format for 'virtual fields', e.g, '*COUNT', or 'vCOUNT', which the reader could easily recognize. Clause 3 would contain a table allowing

*COUNT to be looked up as 'SRP_CMD DATA_OUT BUFFER DESCRIPTOR COUNT
or SRP_CMD DATA_IN BUFFER DESCRIPTOR COUNT, as appropriate'

Open: intel0065 Sect:5.4.1 Pg:18 Ln:24

Remove period after PRESENT

Open: intel0066 Sect:5.4.1 Pg:18 Ln:32

Note 'b' is not referenced above, probably s/b on 'count'

Open: intel0067 Sect:5.4.1 Pg:18 Ln:43

'initiator port may specify in SRP_CMD requests (see 6.8) sent on that RDMA channel. An SRP initiator port shall not specify a data buffer descriptor format that was not indicated in the REQUIRED BUFFER FORMATS field value for that RDMA channel. 'What is target response if it does?'

'SRP target ports are not required to check SRP_CMD requests for data buffer descriptor formats that were not indicated in the REQUIRED BUFFER FORMATS field value.' Not clear - are they required to validate that they did a valid format?

Open: intel0068 Sect:5.4.1 Pg:18 Ln:47

'An SRP target port may accept an RDMA channel and' s/b 'channel establishment request'

Open: intel0069 Sect:5.4.2.2 Pg:18 Ln:49

shall reject the RDMA channel and return after channel, add 'establishment request'

Open: intel0070 Sect:5.4.2.2 Pg:19 Ln:16

indirect data buffer descriptor (IDBD) Use caps or formatting to set off these field names

Open: intel0071 Sect:5.4.2.2 Pg:19 Ln:16

if the SRP initiator port may specify the INDIRECT s/b 'if the TP will accept...'

Open: intel0072 Sect:5.4.2.2 Pg:19 Ln:18

does not use (Sense is that IP forebears use of indirect) shall not use?

Open: intel0073 Sect:5.4.2.4 Pg:19 Ln:44

'sixteen bytes' Previously defined in Table 2 - eschew multiple definitions

Open: intel0074 Sect:5.4.2.4 Pg:19 Ln:48

target port shall only issue RDMA Read operations using the memory descriptor contained in the direct data buffer descriptor. Statement does not have desired effect - limits what you can read, but does not limit accesses to READs. s/b 'shall issue only RDMA Reads when using'

Open: intel0075 Sect:5.4.2.4 Pg:20 Ln:1

shall issue only RDMA Writes...

Open: intel0076 Sect:5.4.2.5 Pg:20 Ln:6

format code value 'value' appears to be superfluous

Open: intel0077 Sect:5.4.2.5 Pg:20 Ln:8

'The length....sixteen bytes.' Drop sentence - redundant to Table 2

Open: intel0078 Sect:5.4.2.5 Pg:20 Ln:10

'An indirect data buffer is comprised of one or more memory segments' Need a real definition.

Open: intel0079 Sect:5.4.2.5 Pg:20 Ln:11

segments may or may not be contiguous. s/b 'may be discontinuous'

Open: intel0080 Sect:5.4.2.5 Pg:20 Ln:11

remove 'may be in a single memory region'

Open: intel0081 Sect:5.4.2.5 Pg:20 Ln:12

of the memory segments (ADD: listed in an IBDB)

Open: intel0082 Sect:5.4.2.5 Pg:20 Ln:13

may have any length As the length field is finite, so is the segment length

Open: intel0083 Sect:5.4.2.5 Pg:20 Ln:29

value contained in the data buffer descriptor's count field. Implies that the field is contained within the DBD

Open: intel0084 Sect:5.4.2.5 Pg:20 Ln:29

'count' Suggest replacing with 'PMDL Length'

Open: intel0085 Sect:5.4.2.5 Pg:20 Ln:31

DESCRIPTOR field value is a memory descriptor Suggest: DESCRIPTOR field contains a memory descriptor

Open: intel0086 Sect:5.4.2.5 Pg:20 Ln:33

concatenated together 'together' is redundant Stamp Out and Abolish Redundancy!

Open: intel0087 Sect:5.4.2.5 Pg:20 Ln:35

contains the number of memory descriptors in the indirect table times sixteen. Suggest: contains the length, in bytes, of the indirect table (16 bytes * number of descriptors in table)

Open: intel0088 Sect:5.4.2.5 Pg:20 Ln:36

MEMORY DESCRIPTOR field value contains any other drop 'value'

Open: intel0089 Sect:5.4.2.5 Pg:20 Ln:43

list of n memory descriptors Use bold or something to set off n

Open: intel0090 Sect:5.4.2.5 Pg:20 Ln:47

shall only issue s/b shall issue only

Open: intel0091 Sect:5.4.2.5 Pg:21 Ln:1

shall only issue s/b shall issue only (also Ln 4)

Open: intel0092 Sect:5.4.2.5 Pg:21 Ln:12

All four..., each might..., or several might be...

Awkward - generalize to: segments may be in different memory regions

Open: intel0093 Sect:5.4.2.5 Pg:21 Ln:44

value contains Drop: value (i.e.,) Add: in bytes

Open: intel0094 Sect:6.1 Pg:23 Ln:7

Add M/O column, or statement that all are mandatory.

Open: intel0095 Sect:6.1 Pg:23 Ln:24

Add space between Tables 6 and 7 to clarify distinction between I>T and T>I requests.

Open: intel0096 Sect:6.1 Pg:23 Ln:46

shall send SRP_T_LOGOUT What reason code?

Open: intel0097 Sect:6.1 Pg:23 Ln:48

Need to define requestor, responder. Much reviewer confusion wrt Targ as requestor.

Open: intel0098 Sect:6.1 Pg:24 Ln:2 C

Responders are not required to check whether the TAG values of outstanding SRP requests are unique. Since duplicate tags would likely cause a credit leak (one response for two requests), this could lead to deadlock, as InitReqLimit and TRL would be out of sync. We either need to require verification of uniqueness, or provide a ReqLimit re-sync mechanism.

Open: intel0099 Sect:6.2 Pg:24 Ln:2

shall only be sent during RDMA s/b: shall be sent only during RDMA

Open: intel0100 Sect:6.2 Pg:24 Ln:41

maximum length Add: in bytes

Open: intel0101 Sect:6.3 Pg:27 Ln:4

shall only be sent s/b: shall be sent only

Open: intel0102 Sect:6.3 Pg:27 Ln:40

maximum length Add: in bytes

Open: intel0103 Sect:6.3 Pg:27 Ln:45 C

52 or larger AER_REQ requires 56

Open: intel0104 Sect:6.4 Pg:29 Ln:3

a(n) SRP target

Open: intel0105 Sect:6.4 Pg:29 Ln:40

too large Need a way to specify, so that Init does not have to guess

Open: intel0106 Sect:6.5 Pg:30 Ln:20

delay a vendor specific time Wait for transport ACK or timeout error at least

Open: intel0107 Sect:6.6 Pg:31 Ln:3 C

An SRP_T_LOGOUT request may also be used to notify the SRP initiator port that an RDMA channel has failed, rendering it non-operational. If the channel has failed, it won't be able to carry this IU. We DO need a way to report failures.

Open: intel0108 Sect:6.6 Pg:31 Ln:30

There are no references in spec to reason codes 2,3, 6-9. Do we need some SHALLS pointing to them?

Open: intel0109 Sect:6.6 Pg:31 Ln:45

delay a vendor... Reference: xport ack or timeout

Open: intel0110 Sect:6.8 Pg:34 Ln:14

COUNT Change to PMDL Length

Open: intel0111 Sect:6.8 Pg:34 Ln:40

Add ref a,b to notes below

Open: intel0112 Sect:6.9 Pg:36 Ln:36

Since SENSE DATA length is 7 bytes + a one-byte length field, at least the top two bytes s/b reserved. We may want to have this field be that one-byte length field, with 7 assumed, as in SPC.

Open: intel0113 Sect:6.9 Pg:37 Ln:9

length of the...buffer Ref 5.4 for length determination

Open: intel0114 Sect:6.9 Pg:37 Ln:26

indicates (that) the contentsshall be ignored and (that) the

Open: intel0115 Sect:6.9 Pg:37 Ln:26

The(value of the) SENSE DATA LIST LENGTH field (be a multiple of four).

Open: intel0116 Sect:6.9 Pg:37 Ln:26 C

SENSE DATA LIST LENGTH shall contain the length of the truncated SENSE DATA field. This is at odds with SPC-2, which returns the total length. How would you know that you had missed some Sense Data?

Open: intel0117 Sect:6.9 Pg:37 Ln:47

shall contain a length of 4 Also defined in Table 23 - refer to table instead

Open: intel0118 Sect:6.9 Pg:38 Ln:17

structure eqn as DOBL - (offset + 1) Much easier to understand (global change to all similar eqns) Formatting - more white space above and below, use bold font

Open: intel0119 Sect:6.9 Pg:38 Ln:24

may or may not not is the more worrisome case (more so for Ln 25)

Open: intel0120 Sect:6.9 Pg:38 Ln:31

Some commands may have a non-zero residual Add: e.g., INQUIRY

Open: intel0121 Sect:6.9 Pg:38 Ln:45

may not

Open: intel0122 Sect:6.9 Pg:39 Ln:1

certian (SRP) protocol errors

Open: intel0123 Sect:6.9 Pg:39 Ln:18

Would there ever be a case where a RSP of NO FAILURE was returned?

Open: intel0124 Sect:6.9 Pg:39 Ln:31

sense data shall be presented presented s/b returned Also Ln 32,33

Open: intel0125 Sect:6.9 Pg:39 Ln:33

whose Use whose wrt people only

Open: intel0126 Sect:6.9 Pg:39 Ln:30

SPC-2 Annex C references SPC-3 - which?

Open: intel0127 Sect:6.11 Pg:40 Ln:43 C

See comments on 5.3 for CRED_RSP issues

Open: intel0128 Sect:6.12 Pg:41 Ln:31

Don't need four bytes for SENSE data length (7 + 1 byte)

Open: intel0129 Sect:6.12 Pg:41 Ln:43

The (value of) the SENSE DATA Len field (shall be a multiple of four.)

Open: intel0130 Sect:6.12 Pg:41 Ln:44 C

If no sense data is provided, What would the point be - to force Init to issue Req Sense Request? Should this be allowed?

Open: intel0131 Sect:6.12 Pg:42 Ln:1

SENSE DATA LIST LENGTH shall contain the length of the truncated SENSE DATA field. Appears to violate SPC-2.

Open: intel0132 Sect:6.12 Pg:42 Ln:7

presented s/b (returned in response to)

Open: intel0133 Sect:7.2 Pg:43 Ln:21

The following subclause defines the fields in the disconnect-reconnect mode Nope - same subclause

Open: intel0134 Sect:7.2 Pg:43 Ln:25

Gray-out or mark as Reserved the fields that are reserved for SRP. There's a lot of noise for the two fields that are used...

Open: intel0135 Sect:7.2 Pg:44 Ln:1

SRP devices shall only use (the) disconnect-reconnect page parameter fields Use formatting for disconnect-reconnect

Open: intel0136 Sect:7.2 Pg:44 Ln:1

SRP devices shall only use ...fields defined below. What about the standard mode page header fields?

Open: intel0137 Sect:7.2 Pg:44 Ln:7

field shall not be implemented by SRP target ports Define in terms of behavior, not implementation. Appears to have been covered by para above.

Open: intel0138 Sect:7.2 Pg:44 Ln:17

If the EMDP bit is set to 0, the SRP target port shall generate (RDMA requests with) continuously increasing () addresses for a single SCSI command.

Open: intel0139 Sect:7.2 Pg:44 Ln:19

affect the order of frames within an RDMA. What's a frame? Within an RDMA what?

Open: intel0140 Sect:7.2 Pg:44 Ln:24

Open: intel0141 Sect:7.2 Pg:44 Ln:24

protocol specific s/b protocol-specific (also Ln 27)

Open: intel0142 Sect:7.2 Pg:44 Ln:28

LUN -> PORT

Open: intel0143 Sect:A.1 Pg:45 Ln:11

Top right box s/b Device Server?

Open: intel0144 Sect:A.1 Pg:45 Ln:29

four step, two step s/b four-step, two-step (global)

Open: intel0145 Sect:A.1 Pg:46 Ln:16

Need close paren after initiator

Open: intel0146 Sect:A.1 Pg:46 Ln:43

Close paren

Open: intel0147 Sect:A.3 Pg:47 Ln:11

See table A.1 for the definitions of the names used within Don't see names there - objects?

Open: intel0148 Sect:A.4.1 Pg:48 Ln:44

Use bold for EXECUTE COMMAND

Open: intel0149 Sect:B.3.1.7 Pg:52 Ln:35

IBTA uses caps for G S I

Open: intel0150 Sect:B.3.1.2 Pg:52 Ln:23

Do we need to define, spell out GUID?

Open: intel0151 Sect:B.3.1.14 Pg:53 Ln:1

Ports also present on switches.

Open: intel0152 Sect:B.3.1.16 Pg:53 Ln:5

Speel out QP, use IBTA definition.

Open: intel0153 Sect:B.3.2 Pg:53 Ln:20

IBTA uses caps for R T U

Open: intel0154 Sect:B.4 Pg:54 Ln:50 C

Each IB GID is globally unique, Not true - see IBA Vol 1 4.1.1

Open: intel0155 Sect:B.4 Pg:55 Ln:17

worldwide Varies - see IBA Vol 1, 4.1.1

Open: intel0156 Sect:B.5 Pg:56 Ln:2

An SRP initiator device is one or more IB consumers may consist of

Open: intel0157 Sect:B.5 Pg:56 Ln:15

The GUID field should (be) an IB GUID available to the SRP initiator port, Must it be a GUID, an IB GUID,?

Open: intel0158 Sect:B.5 Pg:56 Ln:17

The IDENTIFIER EXTENSION field shall be chosen by the SRP initiator port to ensure that all SRP initiator port identifiers are unique. Over what domain?

Open: intel0159 Sect:B.5 Pg:56 Ln:36

[containing] the SRP target port. providing?

Open: intel0160 Sect:B.5 Pg:56 Ln:41

The service delivery subsystem contains queue pairs, IB channel adapters, IB ports, and the InfiniBand TM Architecture fabric. Contains exclusively? How does this map to Clause 4 RDMA Comm Service?

Open: intel0161 Sect:B.5 Pg:56 Ln:47

general service interface IBTA uses caps

Open: intel0162 Sect:B.5 Pg:56 Ln:48

/ (breaks across page) O Remove slash from FRAME list of characters for line breaks.

Open: intel0163 Sect:B.6.2 Pg:57 Ln:13

open IBA connections use establish instead

Open: intel0164 Sect:B.6.3 Pg:57 Ln:25

Port and CM Redirection or Port Redirection. Very hard to parse - use bold or underscores inside the names

Open: intel0165 Sect:B.6.4 Pg:57 Ln:38

SRP_LOGOUT IU list as T_LOGOUT, I_LOGOUT or define as a virtual field

Open: intel0166 Sect:B.6.4 Pg:57 Ln:38

CM disconnect request use caps -it's not generic

Open: intel0167 Sect:B.6.4 Pg:57 Ln:38

The sender may disconnect if its send queue has transitioned to (THE) error state. What do you mean by disconnect here - local action?

Open: intel0168 Sect:B.6.4 Pg:57 Ln:42

The receiver of an SRP_LOGOUT IU shall respond with an InfiniBand TM Architecture transport acknowledgement and disconnect. Destroy QP, send DREQ, ...?

Open: intel0169 Sect:B.6.5 Pg:57 Ln:46

to an ... RDMA READ Request. One or more requests.

Open: intel0170 Sect:B.6.5 Pg:58 Ln:1

WRITE packets WRITE requests

Open: intel0171 Sect:B.7 Pg:58 Ln:37

outcommands

Open: intel0172 Sect:B.7 Pg:59 Ln:7

Why list ChangeID and OptionROM to say we don't care about them?

Open: intel0173 Sect:B.7 Pg:60 Ln:23 C

Send Message Depth Reserved -> Maximum Initiator Request Limit This allows initiators to efficiently allocate buffers

Open: intel0174 Sect:B.7 Pg:60 Ln:24 C

RDMA Read Depth reserved -> Maximum IOC-issued RDMA depth Allows inits to efficiently allocate RDMA resources

Open: intel0175 Sect:B.7 Pg:60 Ln:26 C

Send Message Size rsvd -> MAXIMUM INITIATOR TO TARGET IU SIZE Eliminates need to guess this value

Open: intel0176 Sect:B.7 Pg:60 Ln:46 C

This field is expected to be marked obsolete in future versions of the InfiniBand TM Architecture Not for T10/ANSI to say

Open: intel0177 Sect:B.7 Pg:61 Ln:13 C

Is :reserved a literal? If not, express as :zzzz, explain below that it is reserved.

Open: intel0178 Sect:B.7 Pg:61 Ln:16

No references to Table notes.

Open: intel0179 Sect:B.7 Pg:61 Ln:16

padded s/b extended

Open: Comment from IBTA by William Futral (Intel):

The IBTA Application Working Group understands that the SRP document is out for review and would like to offer the following comment.

The value assigned to I/O Class field in Table B.7 of the SRP document needs to be changed as a result of a change made to the format of this component in the latest InfiniBand(TM) Identifiers Annex, which is a supplement to InfiniBand(TM) Architecture Specification Volume 1.

Attached is a PDF document that contains the new wording in the IBTA Annex (see T10/01-319).

A Class Category needs to be selected for the SRP protocol and inserted in the I/O class field in place of the 0xFF value currently stated. For example, if the Storage Class was selected, the value for I/O class in your Table b.7 would become 0x0100.

Bill Futral
Application Working Group Co Chair
InfiniBand Trade Association

Open: OD 1 Page 13, lines 5-7, multiple RDMA writes on the same channel store data in order.

Some RDMA communication services (e.g. iWARP) are unable to ensure strict ordering of overlapping RDMA Write operations during normal operation. While methods are available to ensure strict ordering, invoking them for all RDMA Writes would severely affect performance.

SAM-2 does not specify the result of multiple commands to overlapping buffers in most cases. It is unclear whether it specifies the result in any situation (see T10/01-309). Overlapping transfers, also called data overlay, within a single command is unusual enough that some SCSI protocols routinely prohibit it.

This requirement should be removed from SRP. It should be replaced with a statement that overlapping transfers may yield unpredictable results unless the RDMA client (SRP) takes special precautions. The nature of said special precautions, if any, are RDMA communication service specific. A section should be added to clause 5 discussing data overlay to specify that SRP target ports shall take said special precautions whenever data overlay occurs within a command.

Open: OD 2 Page 13, line 13, RDMA read operations may complete in any order.

While this states that RDMA Read operations may complete in any order, it is not clear what data they are required to return. See the first example in T10/01-309r0.

If T10/01-309r0 is accepted, this should be clarified to indicate that the data returned by RDMA Read operations need not reflect concurrent RDMA Writes that precede the RDMA Read.

If T10/01-309r0 is not accepted, this should be changed to require that RDMA Reads and RDMA Writes to overlapping locations are strictly ordered for memory access.

Open: OD 3

Page 14, RDMA channel disconnection
Page 15, Multiple independent RDMA channel operation
Page 16, lines 9 and 10 (list items b and c)
Page 27, SRP_LOGIN_RSP response
Page 30, SRP_I_LOGOUT request
Page 31, SRP_T_LOGOUT request

One of the characteristics of a network or fabric communication service is that errors affecting a channel can rarely be reported using that channel. In the context of SRP, many errors that disconnect an RDMA channel will be reported to one consumer but not the other. The consumer receiving the report cannot use the same RDMA channel to notify the other consumer, as the channel is no longer operational.

It is nonetheless useful for both consumers to know that an RDMA channel has failed. When using multiple independent RDMA channels, the consumers could use one of the other channels to report a channel failure. SRP should be extended to support this. This should be mandatory behavior whenever multiple channels are used between the same SRP initiator port and the same SRP target port. The following paragraphs summarize the changes to SRP to accomplish this.

The SRP_LOGIN_RSP response should return a channel handle. The channel handle shall be non-zero and unique among all channels in use on the same I_T nexus. Zero is valid if and only if the SRP target port only supports one channel per nexus. The channel handle should be a 16-bit field in bytes 28 and 29.

The SRP_I_LOGOUT and SRP_T_LOGOUT requests should specify an optional channel handle. The channel handle should be a 16-bit field in bytes 2 and 3. If the channel handle is zero, it specifies that the channel on which the request was sent is being logged out; no response is generated. This is identical to the behavior currently specified by SRP. If the channel handle is non-zero then the specified channel is being logged out. A response is generated to confirm the logout and to indicate that all outstanding requests on that channel have been discarded. Targets shall not use of a non-zero channel handle that specifies the channel on which the SRP_T_LOGOUT request is sent. Use of a non-zero channel

handle that specifies the channel on which the SRP_I_LOGOUT request is sent results in target specific behavior.

Extend the discussion of RDMA channel disconnection (page 14) and multiple independent RDMA channel operation (page 15) to require that targets report disconnection using an alternate channel if one is available.

Amend the list of requests that do not have responses on page 16 to say that SRP_I_LOGOUT and SRP_T_LOGOUT do not have responses when the channel handle is zero, but do have responses when the channel handle is non-zero.

Note that this change cannot be straightforwardly added in an SRP-2. An initiator or target that ignores the channel handle field (because it was reserved in SRP) would logout the wrong channel.

Open: OD 4 Page 56, tables B.2 and B.3

SRP port identifiers for Infiniband are 128-bit identifiers with an embedded GUID (EUI-64). Infiniband GIDs are 128-bit identifiers with an embedded GUID (EUI-64). Unfortunately they are formatted incompatibly. Annex B specifies that the EUI-64 occupies the most significant bytes of an SRP port identifier while the EUI-64 occupies the least significant bytes of an InfiniBand GID or IPv6 formatted address. The bytes not occupied by the EUI-64 are also different.

Having conflicting formats of otherwise equivalent identifiers is guaranteed to lead to interoperability problems. Various people have stated (in SRP working groups) that they expect to identify SRP targets using IPv6 formatted identifiers. SRP should be changed to satisfy this.

A new informative annex should be added recommending that SRP port identifiers adhere to IPv6 address formatting conventions and use one of the three forms listed below. Annex B should require that InfiniBand SRP port identifiers be one of the three forms listed below.

1. The Link-Local prefix (FE80h:0:0:0::/64) concatenated with an EUI-64.
2. The Site-Local prefix (FEC0h:0:0:0::/48) concatenated with 16-bit locally administered value concatenated with an EUI-64.
3. Any value configured manually or by a system management agent.

Open: OD 5 Pages 4 and 5, glossary terms, and their use throughout the document, Clause 4

When SRP was proposed and for much of its development no satisfactory glossary of RDMA terms was available. Available external documents used definitions specific to particular implementations. That has recently changed. See the message titled "iWARP Glossary" posted to the yahoo RDMA reflector on September 27, 2001 by Jim Wendt. It would be beneficial if SRP were changed to use the same terms and definitions.

Open: OD 6 Page 11 lines 20-22, normal and solicited message reception.

This feature is described in the RDMA communication service model, yet not used by SRP. Interrupt mitigation is important in high end systems. Therefore this should be supported by SRP information units. A description of how to do so follows.

Define a bit to be included in all SRP information units. Recommend this be bit 0 of byte 1 and called noturg (notification urgency or not urgent, take your pick).

In initiator to target requests, noturg specifies the notification urgency for the response. The initiator may set it to any value.

In target to initiator responses, noturg specifies the notification urgency. The target shall copy it from the request.

In target to initiator requests, noturg shall be zero. Specify this individually in each request, not as a general rule, so that it may be changed for future requests.

In initiator to target responses, the target shall ignore noturg.

In Annex B, specify that the target shall send information units with solicited event notification enabled if noturg is zero. The target shall send information units with solicited event notification disabled if noturg is one. The initiator shall ignore noturg and send all information units with solicited event notification enabled.

Open: OD 7 RDMA communication service specific opcode

SRP currently requires RDMA Read support for practical operation. However there are RDMA communication services that do not support an RDMA Read. So-called Unreliable Connections on InfiniBand are an example. Note that these have the same reliability characteristics as most existing SCSI protocols (e.g. FCP). Various people have suggested that they would be the most natural service for storage access, except for the lack of RDMA Read. Some VI Architecture implementations also lack RDMA Read.

It is straightforward to emulate an RDMA Read. The target sends a request to the initiator identifying the data to be read. The initiator responds with an RDMA Write supplying the required data, then a response to indicate completion. All that is missing is SRP opcodes that could be used for the request and response.

This is one example of a need for an RDMA communication service specific operation. Others might be required in the future for as yet unanticipated reasons. The purpose of defining this now is to describe proper behavior for an initiator that does not recognize the request.

The following could be defined using a new pair of opcodes or as an extension to the existing SRP_CRED_REQ and SRP_CRED_RSP. I don't particularly care which is used.

Define a target to initiator request. It is formatted identically to SRP_CRED_REQ with the addition of an action code field and action code specific parameters. I recommend a 16-bit action code field. The action code specific parameters may be any length (including zero) provided the total request length is within the limit agreed to during login.

Define the corresponding initiator to target response. It is formatted identically to SRP_CRED_RSP with the addition of an action code, an action response code and action code specific parameters. The action code is an echo of the value in the request (could be omitted). The action response code indicates the outcome of the action. Define value zero to designate the action is not supported, all other values reserved. The action code specific parameters may be any length (including zero) provided the total request length is within the limit agreed to during login. If the response code indicates the action was not supported, the action code specific parameters shall be zero length.

Open: OD 8 Page 18, lines 13-37 and elsewhere, data buffer format code and count values.

The combination of a data buffer format code and a data buffer format count is awkward. Their interpretation is interdependent. We really have a single 12-bit field. It would simplify the description (and probably the implementation) if we had a single encoded data buffer format field. The following is a suggested way to encode an 8-bit data buffer format code:

00h NO DATA BUFFER DESCRIPTOR PRESENT
 01h DIRECT DATA BUFFER DESCRIPTOR
 02h – 0Fh Reserved

1xh INDIRECT DATA BUFFER DESCRIPTOR
10h INDIRECT DATA BUFFER DESCRIPTOR WITH NO PARTIAL MEMORY
DESCRIPTOR LIST
11h INDIRECT DATA BUFFER DESCRIPTOR WITH 1 ENTRY PARTIAL MEMORY
DESCRIPTOR LIST
12h INDIRECT DATA BUFFER DESCRIPTOR WITH 2 ENTRY PARTIAL MEMORY
DESCRIPTOR LIST
etc.
1Fh INDIRECT DATA BUFFER DESCRIPTOR WITH 15 ENTRY PARTIAL MEMORY
DESCRIPTOR LIST
20h to FFh Reserved

These values would occupy bytes 6 and 7 of SRP_CMD, byte 5 would be reserved.

Texas Instruments

This has the appearance of a draft copy, not a final review copy. Change bars and line numbers should not be on a letter ballot document.

Open: Troika Networks, Inc.:

The TYPE code value of 80h in tble 13 is incorrect according to table 6 and should be value C2h.

Open: Woven Electronics:

Can not Contribute