

To: T10 Technical Committee
From: Rob Elliott (elliott@hp.com)
Date: 24 September 2007
Subject: 07-263r1 SAM-4 SCSI Initiator Port and Target Port capabilities annex

Revision history

Revision 0 (6 June 2007) First revision

Revision 1 (24 September 2007) Incorporated comments from July 2007 and September 2007 CAP WGs. For SAM-4, just keep the informative annex; defer defining the attributes and any reference to the fence and ordering topics until SAM-5 (proposal 07-429).

Related documents

T10/sam4r12 - SCSI Architecture Model - 4 (SAM-4) revision 12

T10/sas2r11 - Serial Attached SCSI - 2 (SAS-2) revision 11

T10/adt2r04a - Automation Device Transport - 2 (ADT-2) revision 4a

T10/fcp4r00 - Fibre Channel Protocol - 4 (FCP-4) revision 0

T10/srp-r16a - SCSI RDMA Protocol (SRP) revision 16a

T10/06-341 - SAM-4 Request and Response Fence for protocol services (Rob Elliott, HP)

T10/07-072 - FCP-4 QUERY TASK task management function (Rob Elliott, HP)

T10/07-143 - FCP-4 QUERY TASK SET task management function (Rob Elliott, HP)

T10/07-144 - FCP-4 QUERY UNIT ATTENTION task management function (Rob Elliott, HP)

T10/07-157 - SAM-4 Changes requested from 3/2007 editing session (George Penokie, IBM)

T10/07-429 - SAM-5 SCSI Initiator and Target Port capability attributes (Rob Elliott, HP)

IETF/RFC 3270 - iSCSI

IETF/draft-ietf-ips-iscsi-impl-guide-08.txt - iSCSI Implementor's Guide revision 8 (by Mallikarjun Chadalapaka, HP). Available as <http://www.ietf.org/internet-drafts/draft-ietf-ips-iscsi-impl-guide-08.txt> through November 2007.

Overview

Each SCSI transport protocol supports a different subset of the features in SAM-4. For example, ADT only supports 16 byte CDBs, while SAS, FCP, ADT, SRP, and iSCSI all support at least 260 bytes (but their maximum sizes are not all identical). The Retry Delay Timer feature is only supported in SAS-2 and FCP-4; it has not yet been added to SRP, iSCSI or ADT.

All protocols currently have a maximum Data Offset field in their DATA frames of 32 bits (meaning 0 through $2^{32}-1$ bytes of data can be transferred for a single command), yet some CDB structures (particularly in the OSD command set) allow commands to be specified that transfer more than 2^{32} bytes of data.

A SCSI port may not always implement all the features defined by its protocol, particularly as the transport protocols are updated (e.g., FCP-3 to FCP-4). Sometimes the protocol explicitly defines a feature as optional. For example, if Sequence Level Recovery is being used, FCP-4 prohibits using bidirectional commands. The hardware, device driver or operating system may impose constraints as well. For example, most operating systems do not support bidirectional commands even though all the SCSI transport protocols nominally support them.

Conceptually, an application client needs to understand the limits of the SCSI initiator port, and a device server needs to understand the limits of the SCSI target port - both those limits imposed by the SCSI transport protocol it uses and any extra limits imposed by the implementation (e.g., just because the protocol supports bidirectional commands doesn't mean the SCSI port always supports them). In 07-429, these are proposed to be reported in new UML attributes for initiator port and target port objects (defined in the SCSI Initiator Port and SCSI Target Port classes). For the Execute Command () and task management procedure call input and output arguments, attributes are defined to indicate whether certain arguments are supported (e.g., Retry Delay Timer) and indicate the maximum values for other arguments (e.g., CDB size). Attributes indicating other SCSI initiator port information like the SCSI transport protocol are also defined (the port name and the port identifier were already defined as attributes).

This does not mean an HBA (or the controller on a target side) is required to literally implement a data structure containing these values; it just means the HBA and its device driver must understand the limits and

reject any attempts by an application client to exceed those limits. For example, if an HBA doesn't support a full 260 byte CDB, it should report an error to the application client rather than silently truncating a CDB, or have a interface that simply does not allow such a CDB to be specified. The T11.5 SM-HBA-2 standard could be enhanced to report some or all of these attributes.

In 07-429, a new Service Response is proposed to represent returning an error to the application client. SAM-4 doesn't really cover target ports returning errors to device servers; that is not changed by this proposal.

In 07-263r1, an informative annex is included listing the values supported by each of the SCSI transport protocols.

Suggested changes to SAM-4

Annex B

(informational) [\[all new\]](#)

SCSI Initiator Port and SCSI Target Port attributes supported by SCSI transport protocols

Table B.1 lists the values of the SCSI Initiator Port attributes that a SCSI initiator port using each different SCSI transport protocol is able to return, and the values of the SCSI Target Port attributes that a SCSI target port using that SCSI transport protocol is able to return.

Table B.1 — SCSI Initiator Port and SCSI Target Port attributes supported by SCSI transport protocols
(part 1 of 2)

Attribute	ADT-2	FCP-4	iSCSI	SAS-2	SRP
Maximum CDB length (in bytes) ^a	16	268	65 550	268	268
Task identifier size (in bits)	3	16	32	16	64
Task Attributes supported	SIMPLE, HEAD OF QUEUE, ORDERED, and ACA				
Maximum Data-In Buffer Size (in bytes)	FFFFFFFFh				
Maximum Data-Out Buffer Size (in bytes)	FFFFFFFFh				
Maximum CRN ^b	zero	FFh	zero	zero	zero
Task Priority supported	no	yes	no	yes	no
Maximum Sense Data Length (in bytes) ^c	FFFFh	FFFFFFFFh	FFFFh	3E8h ^e	FFFFFFFFh
Retry Delay Timer supported	no	yes	no	yes	no
^a SPC-4 defines the maximum length of a CDB as being 260 bytes. ^b A Maximum CRN of zero indicates that CRN is not supported. ^c SPC-4 defines the maximum length of sense data as being 252 bytes. ^d The task management function name is not an argument to a procedure call; it is the name of the procedure call itself. ^e 3E8h represents 1 000 bytes, which keeps the sense data in one RESPONSE frame.					

Table B.1 — SCSI Initiator Port and SCSI Target Port attributes supported by SCSI transport protocols
(part 2 of 2)

Attribute	ADT-2	FCP-4	iSCSI	SAS-2	SRP
Additional Response Information supported	no	yes	no	yes	no
Bidirectional Commands supported	yes				
Task Management Functions supported ^d	ABORT TASK, ABORT TASK SET, CLEAR TASK SET, LOGICAL UNIT RESET, CLEAR ACA, QUERY TASK	All	ABORT TASK, ABORT TASK SET, CLEAR TASK SET, LOGICAL UNIT RESET, CLEAR ACA	All	ABORT TASK, ABORT TASK SET, CLEAR TASK SET, LOGICAL UNIT RESET, CLEAR ACA
<p>^a SPC-4 defines the maximum length of a CDB as being 260 bytes.</p> <p>^b A Maximum CRN of zero indicates that CRN is not supported.</p> <p>^c SPC-4 defines the maximum length of sense data as being 252 bytes.</p> <p>^d The task management function name is not an argument to a procedure call; it is the name of the procedure call itself.</p> <p>^e 3E8h represents 1 000 bytes, which keeps the sense data in one RESPONSE frame.</p>					

Editor’s Note 1: ADT-2 is expected to add QUERY TASK SET and QUERY UNIT ATTENTION before it is published. As of revision 4a, it does not include them.

Editor’s Note 2: iSCSI is likely to add Task Priority, Retry Delay Timer, QUERY TASK SET, and QUERY UNIT ATTENTION via a “SAM-4 enhancements” RFC in the future.

Editor’s Note 3: An SBP-2 (1394) column is not included. It is considered a SAM-2 compliant protocol like SPI-5 (parallel SCSI) so isn’t really covered by later versions of SAM. Formatting will become an issue when another column is added.

[\[end of all new section\]](#)