

This document mimics the structure of the SCSI on Fibre Channel (FCP) document of the SCSI-3 specification.

References:

- 1) ST standard
- 2) SCSI-3 standards

The ST Protocol for SCSI standard is part of the SCSI-3 family of standards developed by X3T10 to facilitate the use of the SCSI command sets for many different types of devices across many different types of physical interconnects [not really, but it probably should be...]. The master architectural document of the family of standards is ANSI X3.270-199X, Information Technology - SCSI-3 Architecture Model (SAM). The SAM document contains a guide to other documents in the SCSI-3

Clause 1 describes how SCSI operations are structurally mapped on to ST operations in the SCSI on ST protocol.

Clause 2 defines specific message formats for the SCSI on ST protocol.

Clause 3 defines error recovery procedures.

1 General

=====

1.1 Structure and Concepts

=====

Scheduled Transfer (ST) is a data transfer protocol which may be implemented on a wide variety of Lower Layer Protocols (LLP), which supports a connection-oriented data transfer model with multiple outstanding data transfers per connection.

Three kinds of functional management are defined by this SCSI on ST specification:

- 1) Connection management
- 2) Device management
- 3) Task management

For the purposes of describing this protocol, we define the term ``transaction'' which consists of one or more distinct ST sequences (e.g. read sequence, write sequence, PMR Put sequence), which are logically associated to implement a complete function from the list above.

The correspondence between SCSI functions and elements of the SCSI on ST protocol are:

SCSI	SCSI on ST
I/O Operation	Transaction
Command Service Request	Request_To_Send or Request_To_Receive operation
Data delivery request	Clear_To_Send operation
Data delivery action	Data operation (STU)
Command service response	Put sequence

1.2 Common ST Operation Characteristics

=====

All ST control operations used in the SCSI on ST protocol shall have the ST Interrupt flag bit set.

The non-final ST Data operations of all data sequences (regular data, or Status Put), shall have the ST Interrupt and Last flag bits cleared, and the ST Silent flag bit set.

The final ST Data operation of all data sequences shall have the ST Interrupt and Last flag bits set and the ST Silent flag bit cleared.

All ST Data operations for a SCSI on ST connection shall be sent on the ST Data Channel specified in the connection setup protocol. ST Data Channel flag bits shall be appropriately set to reflect this.

The SCSI on ST protocol does not use the operation pairs for reliable data movement as described in section 10.2 of the ST protocol specification. Instead, entire SCSI on ST transactions are retried as described in section 3.2, Operation Timeouts.

1.3 Connection management

=====

A SCSI on ST connection is an extension of the basic ST Virtual Connection.

1.3.1 Connection Setup

=====

A SCSI on ST connection is established by sending an ST Request_Connect operation to the well known port for SCSI on ST (TBD), with SCSI on ST connection parameters in the optional payload.

If the connection request is acceptable to the SCSI target, it responds with an ST Connection_Answer operation, with its SCSI on ST connection parameters in the optional payload.

If the SCSI on ST connection parameters are unacceptable to the connection responder, the responder shall reject the Virtual Connection attempt by setting the ST Reject Flag bit, with an appropriate reason code, as specified in section 2.1.6

If the SCSI on ST connection parameters are unacceptable to the connection initiator, the initiator shall tear down the Virtual Connection with the standard ST Virtual Connection teardown sequence.

If the SCSI on ST connection parameters indicate that the connection initiator can function as a SCSI initiator, and the connection responder intends to function as a SCSI target, the connection responder shall request the connection initiator to expose a Persistent Memory Region, called the Status Persistent Memory Region, of a size specified in the Request_Connect optional payload, for the ST Data Channel specified in the Request_Connect optional payload,

with an ST Request_Memory_Region operation. The connection initiator shall expose the Status Persistent Memory Region and respond with an appropriate Memory_Region_Available operation. The initial Offset of the Status Persistent Memory Region must be a multiple of 512 bytes, the Status Persistent Memory Region segment size. The connection responder will record the parameters of the Status Persistent Memory Region with the Virtual Connection state.

When the Connection_Answer operation is received by the connection initiator, if the connection parameters specify that the connection responder can function as a SCSI initiator and the connection initiator intends to function as a SCSI target, the connection initiator shall request that the connection responder expose a Status Persistent Memory Region, as described above.

A SCSI on ST Virtual Connection endpoint shall not initiate any SCSI operations until the Status Persistent Memory Region has been requested, exposed and the ST Memory_Region_Available operation sent. A SCSI on ST initiator shall keep a timer awaiting the ST Request_Memory_Region operation, and if none arrives within the duration of the timer, tear down the ST Virtual Connection with the standard ST Virtual Connection teardown sequence and retry the SCSI on ST connection sequence.

A SCSI on ST target shall respond to any data transfer request (Request_To_Send or Request_To_Receive) with a Request_Answer with a reason code of Status Persistent Memory Region Not Established (TBD) until the ST Memory_Region_Available operation is received.

If the initial offset returned in the ST Memory_Region_Available operation for the SCSI on ST Status Persistent Memory Region is not a multiple of the the status segment size, 512 bytes, a SCSI on ST target shall respond to any data transfer request with a Request_Answer, with a reason code of Status Persistent Memory Region Unaligned (TBD).

If a SCSI on ST Virtual Connection endpoint intends to function as a SCSI initiator, and the remote endpoint indicates that an authentication handshake is required, the SCSI initiator endpoint shall begin the authentication handshake as described in TBD (TODO).

If authentication is required, a SCSI on ST Virtual Connection endpoint shall not initiate any SCSI operations until the authentication handshake has been successfully performed. A SCSI on ST target shall respond to any data transfer request (Request_To_Send or Request_To_Receive) with a Request_Answer with a reason code of Authentication Required (TBD) until the SCSI on ST authentication handshake is performed. A SCSI on ST target may adopt any suitable policy for reclaiming Virtual Connection resources if authentication is required and a successful SCSI on ST authentication sequence is not performed by a SCSI on ST initiator. Specifically, such a policy should protect against denial of service attacks by unauthorized SCSI on ST connection initiators.

1.3.2 Connection Teardown

=====

SCSI on ST connections shall be torn down using the standard ST Virtual Connection teardown sequence.

1.3.3 Virtual Circuit Keep Alive

=====

A SCSI on ST endpoint which behaves as a SCSI target shall implement ST Virtual Connection keep-alive sequence as described in the ST specification. This keep-alive sequence is the primary mechanism for SCSI on ST targets to reclaim resources associated with a broken Virtual Connection.

A SCSI on ST endpoint which behaves as a SCSI target may implement the ST Virtual Connection keep-alive sequence as described in the ST specification.

1.4 Device management

=====

An application client begins a SCSI on ST I/O transaction when it provides to the SCSI on ST layer a request for an Execute command service. A single request or a list of linked requests may be presented to the software interface. The SCSI on ST layer then performs the following actions using ST services to perform the SCSI command.

The ST port that is the SCSI initiator for the command starts a transaction by sending an ST Request_To_Send, ST Request_To_Receive, or SCSI on ST Zero_Length_Command (see section 2.2, ST NOP Operations in SCSI On ST) operation to begin an ST device management sequence. A Request_To_Send is sent if the SCSI command involves data transfer from the SCSI initiator to the SCSI target, or no data transfer. A Request_To_Receive is sent if the SCSI command involves data transfer from the SCSI target to the SCSI initiator. A SCSI on ST Zero_Length_Command is sent if the command involves no data transfer.

The ST Request_To_Send, Request_To_Receive or SCSI on ST Zero_Length_Command operation has an optional payload including some command control flags, addressing information, and the SCSI command descriptor block (CDB).

This ST Request_To_Send, Request_To_Receive or SCSI on ST Zero_Length_Command operation is the Execute Command service request and starts the I/O Operation. The transaction that is started is identified by the ST Initiator ID, qualified by the ST Virtual Connection. This Initiator ID is set appropriately on all ST operations associated with the transaction.

When the device server for the command has completed the interpretation of the command, has determined that data transfer is required, and is prepared to request the data delivery service, it indicates this to the initiator, and data is transferred, as described below. Note that the amount of data transferred is the minimum of:

- 1) the transfer length specified in the ST Request_To_Send or Request_To_Receive operation which initiated the transaction
- 2) the amount of data the device server requires to be transferred

as specified in the SCSI CDB.

If the transaction was initiated with an ST Request_To_Send operation, the device server sends one or more ST Clear_To_Sends for the amount of data which it is prepared to receive.

The initiator then responds with sequences of one or more ST Data operations to satisfy the outstanding Clear_To_Sends received from the device server. The last ST Data operation in each sequence shall have the ST Interrupt and Last flag bits set, and the ST Silent flag bit cleared.

When the device server is prepared to receive additional data, it may send one or more additional ST Clear_To_Sends, to which the initiator will respond with additional data sequences, until the entire, required amount of data has been transferred.

If the transaction was initiated with an ST Request_To_Receive operation, the device server ``echos'' the ST Request_To_Receive operation back to the initiator as an ST Request_To_Send operation. Data is then transferred from the device server to the initiator as above, except with the initiator sending ST Clear_To_Send operations, and the device server sending ST Data operation sequences.

After all the data has been transferred, the device server transmits the Execute Command service response by sending an ST Put sequence to the Status Persistent Memory Region defined by the connection setup protocol, at an address (bufx/offset) computed from the Initiator ID for the transaction. This ST Put sequence contains the SCSI status and, if an unusual condition has been detected, the SCSI REQUEST SENSE information and the FCP response information describing the condition. The command Status Put terminates the command. The SCSI logical unit determines whether additional commands will be performed in the SCSI on ST I/O Operation. If this is the last or only command executed in the SCSI on ST I/O Operation, the SCSI on ST I/O Operation and the transaction are terminated.

When the command is completed, returned information is used to prepare and return the Execute Command service confirmation information to the software that requested the operation. The returned status indicates whether or not the command was successful. The successful completion of the command indicates that the SCSI device performed the desired operations with the transferred data and that the information was successfully transferred to or from the SCSI initiator. A protocol is provided to present error information if a command is not successful.

If the command is linked to another command, the ST Put payload shall contain the proper status indicating that another command will be executed. The initiator shall continue the same transaction with a SCSI on ST Request_To_Send or Request_To_Receive operation, beginning the next SCSI command. All SCSI commands linked in the SCSI on ST I/O Operation except the last are executed in the manner described above.

The number of FCP I/O Operations that may be active at one time depends on the queuing capabilities of the particular SCSI devices and the number of concurrent exchanges supported by the ST ports.

A brief way of summarizing this is that a SCSI on ST transaction consists of a pair of:

- 1) an ST data transfer operation
- 2) an ST Put sequence with status

for each linked command in a single SCSI Execute service request.

1.4.1 Underrun Residual Handling

=====

If the amount of data the device server requires to be transferred as specified in the SCSI CDB is less than the length of the ST data transfer specified in the ST Request_To_Send or Request_To_Receive operation, either as a result of a successful or unsuccessful SCSI command completion, this is called an underrun residual condition.

If an ST Write sequence in a SCSI on ST transaction results in an underrun residual, the SCSI on ST target shall receive all data it has requested with outstanding ST Clear_To_Send operations before performing the SCSI on ST status Put operation which signifies the completion of the SCSI command.

If an ST Read sequence in a SCSI on ST transaction results in an underrun residual, and the SCSI on ST initiator has sent ST Clear_To_Send operations for this ST Read sequence for data beyond the final block in which the SCSI on ST target has transferred data, the SCSI on ST initiator shall ensure that all Clear_To_Send operations have been received by the SCSI on ST target by sending an ST End operation, with an optional payload that indicates the highest block number of Clear_To_Send operations which have been sent. The SCSI on ST target shall respond to the ST End operation with an ST End_Ack operation as described by the ST protocol, after all outstanding Clear_To_Send operations have been received.

1.4.2 Third Party SCSI Commands

=====

Certain third-party SCSI commands and parameters specify a 64-bit field that is defined to access other SCSI devices addressable from that port. These commands include COPY, RESERVE, and several others.

The ST protocol does not specify any form of addressing. However, it is usually run on a Lower Level Protocol which does specify an address format. Therefore, addressing format for third-party SCSI commands in the SCSI on ST protocol is a function of the Lower Level Protocol addressing format.

These addressing formats are beyond the scope of this specification.

1.5 Task Management

=====

An application client requests a task management function when a task or some group of tasks must be aborted or terminated.

Some SCSI task management functions are mapped into ST control sequences and others are mapped into control bits in the SCSI on ST optional payload of an ST Request_To_Send operation, as part of a complete transaction.

Task management functions that use the ST Request_To_Send operation are sent as the first operation in a new transaction. Task management functions which are mapped to a SCSI on ST transaction end with a SCSI on ST Status Put that indicates whether it was correctly accepted. The CDB field in a SCSI on ST optional payload that perform task management functions is ignored. The mappings are:

SCSI	SCSI on ST
TARGET RESET	SCSI on ST transaction with target reset bit set
ABORT TASK	ST End sequence
ABORT TASK SET	SCSI on ST transaction with Abort Task Set bit set
CLEAR TASK SET	SCSI on ST transaction with Clear Task Set bit set
CLEAR ACA	SCSI on ST transaction with Clear Auto Contingent Allegiance bit set (only ACA is supported by target)

Note that this mapping expands the impact of the ST End sequence from terminating a single data transfer to terminating an entire SCSI on ST transaction.

2 SCSI On ST Protocol Operation Formats

2.1 Optional Payload for Connection_Request and Connection_Answer

Field Name	Description	Size
TYPE	ST optional payload type code	1 byte
LENGTH	ST optional payload size field	1 byte
MAX_TARGET	Maximum supported target address	2 bytes
FLAGS	SCSI on ST feature flags	4 bytes
SPMR_SIZE	Status Persistent Memory Region Size	4 bytes

2.1.1 TYPE

Always 3, indicating that the ST optional payload field is valid and contains a ULP parameter

2.1.2 LENGTH

Always 32, indicating that the ST optional payload field is 32 bytes long.

2.1.3 MAX_TARGET

=====

Maximum supported value of the TARGET field (see section 2.3.3) in the optional payload for ST Request_To_Send and Request_To_Receive operations.

2.1.4 FLAGS

=====

The FLAGS field contains a number of supported and required feature flags:

Bit Definition	Bit

Byte 0 (most significant)	
RECEIVE_ALIGN	7-5
SEND_ALIGN	4-0
Byte 1	
reserved	7-0
Byte 2	
reserved	7-2
ONE_CTS_READ	0
Byte 3	
reserved	7-5
DATA_CHANNEL	4-3
AUTHENTICATE	2
TARGET	1
INITIATOR	0

INITIATOR indicates that the SCSI on ST endpoint is capable of performing SCSI initiator functions.

TARGET indicates that the SCSI on ST endpoint is capable of performing SCSI target functions.

AUTHENTICATE indicates that an authentication handshake must be performed during a SCSI on ST connection setup before data may be transferred.

DATA_CHANNEL indicates the ST Data Channel for which to request the SCSI on ST Status Persistent Memory Region, and on which all SCSI on ST Data operations, including SCSI on ST Status Put sequence shall be performed.

ONE_CTS_READ indicates that a SCSI on ST target requires an initiator to send only single ST Clear_To_Send operation to cover the entire data transfer of a SCSI read operation.

SEND_ALIGN is the log, base 2, of the minimum required alignment of the offset parameter of any ST Clear_To_Send operations sent to the SCSI on ST Virtual Connection endpoint specifying the SEND_ALIGN value.

RECEIVE_ALIGN is the log, base 2, of the minimum required alignment of the offset parameter of any ST Data operations sent to the SCSI on ST Virtual Connection endpoint specifying the RECEIVE_ALIGN value. This, necessarily implies that a SCSI on ST endpoint will never perform a Clear_To_Send operation with an initial offset which does not conform

to the RECEIVE_ALIGN value it specified during the SCSI on ST connection setup.

2.1.5 SPMR_SIZE

=====

This field indicates the number of 512 byte Status Persistent Memory Region segments a SCSI on ST target should request that a SCSI on ST initiator should expose.

In other words, the size of the Status Persistent Memory Region requested shall be:

SPMR_SIZE * 512 bytes

2.1.6 SCSI on ST Virtual Connection Reject Reason Codes

=====

If the parameters of a SCSI on ST Virtual Circuit are unacceptable, the Virtual Circuit request shall be rejected, or closed with an appropriate reason code:

Definition	Value (hex)
-----	-----
Single CTS Read unsupported	TBD
Target function unsupported	TBD
Busy (no resources)	TBD

2.2 ST NOP Operations In SCSI On ST

=====

SCSI on ST distinguishes different uses of the ST NOP operation using the 16 bit opaque Param field as an operation code:

Definition	Value (hex)
-----	-----
Zero_Length_Command	0

The remainder of the opaque fields SCSI on ST Zero_Length_Command operation are unused. The optional payload of the Zero_Length_Command operation contains the optional payload for SCSI on ST command operations as described in 2.3.

2.3 Optional Payload For SCSI on ST Command Operations

=====

The optional payload for Request_To_Send, Request_To_Receive, and Zero_Length_Command operations is organized as:

Field Name	Description	Size
-----	-----	-----
TYPE	ST optional payload type code	1 byte
LENGTH	ST optional payload size field	1 byte
TARGET	Attached target address	2 bytes
LUN	Logical Unit Number	8 bytes

CNTL	Control Field	4 bytes
CDB	SCSI command descriptor block	16 bytes

2.3.1 TYPE

=====

Always 3, indicating that the ST optional payload field is valid and contains a ULP parameter

2.3.2 LENGTH

=====

Always 32, indicating that the ST optional payload field is 32 bytes long.

2.3.3 TARGET

=====

Selects one of a set of individual SCSI targets addressable through a single SCSI on ST Virtual Connection. The maximum allowable value of the TARGET field is established in the SCSI on ST connection setup protocol.

2.3.4 LUN

=====

The FCP logical unit number is the address of the desired logical unit in the attached subsystem. The FCP_LUN field is specified by X3.230-199X for all IUs of Category 6.

Each SCSI FCP target shall accept an INQUIRY command to the logical unit addressed by the FCP_LUN zero (0000 0000 0000 0000 hexadecimal). Using the SCSI INQUIRY information, the initiator can determine the SCSI device type, manufacturer, and model of the logical unit. If the logical unit at LUN 0 specifies a device model that has a defined addressing structure, the initiator can use that information to determine what other logical units are implemented in the target. The structure of the FCP_LUN field is not specified by the FCP. The structure is specified by the SCSI device model and may be vendor unique for device models that do not have a defined address structure. An example of a four-layer hierarchical address is given in annex C.

If the FCP_LUN address locates a valid logical unit, the command shall be executed according to standard SCSI behavior. Behavior may include successful execution of the command, presentation of errors associated with the command, or rejection of the command. If the addressed logical unit does not exist, the responses shall follow the SCSI-3 rules for selection of invalid logical units.

2.3.4 CNTL

=====

The CNTL field contains a number of control flags and control bits:

Bit Definition	Bit

Byte 0 (most significant), reserved	
reserved	7-0
Byte 1, Task codes	
reserved	7-3
TASK ATTRIBUTE	2-0
Byte 2, Task management flags	
reserved	7
CLEAR ACA	6
TARGET RESET	5
reserved	4-3
CLEAR TASK SET	2
ABORT TASK SET	1
reserved	0
Byte 3, Execution management codes	
reserved	7-2
READ DATA	1
WRITE DATA	0

2.3.4.1 Task Codes, Byte 1

=====

TASK ATTRIBUTE can be:

Value, bits 2-0	TASK ATTRIBUTE
-----	-----
000 b	SIMPLE_Q
001 b	HEAD_OF_Q
010 b	ORDERED_Q
100 b	ACA_Q
101 b	UNTAGGED
others	reserved

SIMPLE_Q requests that the task be managed according to the rules for a SIMPLE task attribute.

HEAD_OF_Q requests that the task be managed according to the rules for a HEAD OF QUEUE task attribute.

ORDERED_Q requests that the task be managed according to the rules for an ORDERED task attribute.

ACA_Q requests that the task be managed according to the rules for an automatic contingent allegiance (ACA) task attribute.

UNTAGGED requests that the task be managed according to the rules for an untagged task. Only one untagged task can exist for each logical unit identifier / initiator pair. Requesting a second untagged command for the same logical unit identifier / initiator pair shall be treated as an overlapped command. [X3.270-199X] An untagged task is scheduled according to the rules for a SIMPLE task attribute.

2.3.4.2 Task Management Flags, Byte 2

=====

If any task management flag is set to 1, the CDB, and CDB related FCP_CNTL flags (task codes and execution management codes) are not valid and are ignored. No more than one task management flag shall be

set to 1 in any ST Request_To_Send or Request_To_Receive operation.

CLEAR ACA, when set to 1, causes the ACA condition to be cleared. When the task manager clears the auto contingent allegiance condition, any task within that task set may be completed subject to the rules for task management specified by X3.270-199X.

The use of the ACA bit in the CDB control field and the implementation of ACA is described in X3.270-199X.

If the ACA bit in the CDB control field is set to 0, the automatic sense operation performed by the presentation of the FCP_RSP IU shall clear the ACA condition. Depending on the MODE SELECT parameters that have been established, additional SCSI on ST I/O operations may have to be aborted by the recovery abort

TARGET RESET, when set to 1, performs a reset to the SCSI device as defined in X3.270-199X. TARGET RESET resets all tasks for all initiators and resets all internal states of the target to their initial power on and default values. A unit attention condition is created for all initiators. The initiator and target clear all resources.

CLEAR TASK SET causes all tasks from all initiators in the specified task set to be aborted as defined in X3.270-199X. A unit attention condition is created for all initiators other than the initiator that sent the CLEAR TASK SET that had tasks in the task set.

ABORT TASK SET causes all tasks in the task set from the initiator requesting the ABORT TASK SET to be aborted as defined in X3.270-199X. The bit is mandatory in SCSI on ST.

2.3.4.3 Execution management codes, Byte 3 =====

READ DATA may only be set on an ST Request_To_Receive operation. When set to 1, specifies that the initiator expects data for the task to be in the direction opposite to the direction of the ST Request_To_Receive operation. This is a SCSI read-type operation.

WRITE DATA may only be set on an ST Request_To_Send operation. When set to 1, specifies that the initiator expects data for the task to be in the same direction as the ST Request_To_Send operation. This is a SCSI write operation.

If both READ DATA and WRITE DATA are set to 0, there shall be no data and ST Request_To_Send or Request_To_Receive transfer length shall be 0.

2.3.5 CDB =====

The CDB field contains the actual CDB to be interpreted by the addressed logical unit. The maximum CDB length is 16 bytes. The CDB is not valid and is ignored if any task management flag is set to 1. The command byte of the CDB shall be located at the first byte

transmitted of the CDB field. The CDB's first byte defines the length and basic format of the remainder of the CDB as described in X3.270-199X. Bytes beyond the last byte of the CDB are not defined by SCSI on ST, shall be ignored by the target, and may have any value.

2.4 SCSI on ST Status Put Sequence

=====

The SCSI on ST Status Put sequence shall be directed to an address (bufx/offset) which is computed as:

```

address = (Status Persistent Memory Region initial bufx * initiator bufsize
          + Status Persistent Memory Region initial offset)
          + 512 * transaction Initiator ID
Status Put bufx = address / initiator bufsize
Status Put offset = address % initiator bufsize

```

In addition, the transaction Initiator ID is sent in the ST S_id field, defined as opaque in the ST specification, of the ST Data operations(s) of the SCSI on ST Status Put. This is a departure from all other ST operations for a transaction, where the target puts the Initiator ID in the ST D_id field. The D_id field of SCSI on ST Status Put operations contain the responder ID for the Persistent Memory Region identified in the ST Memory_Region_Available operation received as part of the SCSI on ST connection setup protocol.

The SCSI on ST Status Put sequence shall consist of a single ST Data STU, with the entire SCSI on ST Status payload. The SCSI on ST Status Put Data STU shall have the ST Interrupt and Last flags bits set, and the ST Silent flag bit cleared.

The payload of the SCSI on ST Status Put sequence is:

Field Name	Description	Size
reserved		8 bytes
STATUS	Field validity and SCSI status	4 bytes
RESID	Residual count	4 bytes
SNS_LEN	Length of SNS_INFO field	4 bytes
RSP_LEN	Length of RSP_INFO field	4 bytes
RSP_INFO	SCSI on ST response information	m bytes (0, 4 or 8)
SNS_INFO	SCSI sense information	n bytes

2.4.1 STATUS

=====

The STATUS field is normally 0 upon successful completion of a SCSI on ST I/O Operation. If command linking is being performed, an SCSI on ST Status Put is performed for each command executed. A value of 0 means no error, and in addition indicates that no other information is present in the status put. For linked commands, INTERMEDIATE status or INTERMEDIATE CONDITION MET status indicates successful completion of a command with no other information valid if all other fields are 0.

The format of the FCP_STATUS field is:

Bit Definition	Bit
----------------	-----

```

-----
Byte 0 (most significant)
    reserved                7-0
Byte 1
    reserved                7-0
Byte 2
    reserved                7-4
    RESID_UNDER             3
    RESID_OVER              2
    SNS_LEN_VALID           1
    RSP_LEN_VALID           0
Byte 3
    SCSI status byte        7-0

```

RESID_UNDER indicates that the RESID field is valid and contains the count of bytes that were expected to be transferred, but were not transferred.

RESID_OVER indicates that the RESID field is valid and contains the count of bytes that could not be transferred because the ST transfer length was not sufficient.

SNS_LEN_VALID indicates that the SNS_LEN field is valid and contains the count of bytes in the SNS_INFO field.

RSP_LEN_VALID indicates that the RSP_LEN field is valid and contains the count of bytes in the RSP_INFO field.

Byte 3 (LSB) contains the status byte from the SCSI logical unit. The status byte codes are defined by X3.270-199X.

2.4.2 RESID

=====

If the RESID_UNDER or the RESID_OVER bit is 1, the RESID field contains a count of the number of residual data bytes which were not transferred in the DATA IUs for this SCSI command. Upon successful completion of a SCSI on ST I/O Operation, the residual value is normally 0 and the RESID value is not valid. Devices having indeterminate data lengths may have a nonzero residual byte count after completing valid operations. Targets are not required to verify that the data length implied by the contents of the CDB will create an overrun or underrun before beginning execution of an SCSI command.

If the RESID_UNDER bit is set, a transfer that did not fill the buffer to the expected displacement, specified by the ST transfer length, was performed and the value of RESID is a number equal to:

ST transfer length - highest offset of any byte transmitted

A condition of RESID_UNDER may not be an error for some devices and some commands.

If the RESID_OVER bit is set, the transfer was truncated because the data transfer required by the SCSI command extended beyond the displacement value of DL. Those bytes that could be transferred without violating the ST transfer length value may be transferred. The

RESID is a number equal to:

(Transfer length required by command) - ST transfer length

If a condition of RESID_OVER is detected, the termination state of the SCSI on ST I/O operation is not certain. Data may or may not have been transferred and the SCSI status byte may or may not provide correct command completion information.

If the RESID_UNDER and the RESID_OVER bits are 0, the RESID field is not meaningful and may have any value.

2.4.3 SNS_LEN

=====

If the SNS_LEN_VALID bit is 1, the SNS_LEN field specifies the number of valid bytes of SNS_INFO.

If the SNS_LEN_VALID bit is 0, the SNS_LEN field is not valid and no SNS_INFO is provided.

The SNS_LEN field is always included in the SCSI on ST Status Put.

2.4.4 RSP_LEN

=====

If the RSP_LEN_VALID bit is 1, the RSP_LEN field specifies the number of valid bytes of RSP_INFO. The number shall be 0, 4, or 8. Other values of length are reserved for future standardization.

An RSP_LEN of 0 specifies that no bytes of response information are being provided.

If the RSP_LEN_VALID bit is 0, the RSP_LEN field is not valid and no RSP_INFO is provided.

The RSP_LEN field is always included in the SCSI on ST Status Put.

2.4.5 RSP_INFO

=====

The RSP_INFO field contains information describing only the protocol failures detected during the execution of a SCSI on ST I/O operation. The RSP_INFO does not contain SCSI logical unit error information, since that is contained in the SNS_INFO field as described in 2.4.6. The RSP_INFO field shall contain valid information if the target detects any of the conditions indicated by an FCP RSP_CODE. The format of the RSP_INFO field is:

Bit Definition	Bit

Byte 0 (most significant)	
reserved	7-0
Byte 1	
reserved	7-0
Byte 2	
reserved	7-0

The valid RSP_CODE values are:

RSP_CODE definition	Value (hexadecimal)
No failure or Task Management function complete	0
reserved	1
SCSI on ST CTS or RTS payload fields invalid	2
reserved	3
Task Management Function Not Supported	4
Task Management Function Failed	5
Nonexistent target	6
Busy (no resources)	7

2.4.6 SNS_INFO

=====

The SNS_INFO field contains the information specified by X3.131-1994 for presentation by the REQUEST SENSE command. The proper SNS_INFO shall be presented when the SCSI status byte of CHECK CONDITION or COMMAND TERMINATED is presented as specified by X3.270-199X. SCSI on ST devices should always use autosense.

3 ERROR RECOVERY PROCEDURES

=====

3.1 Virtual Circuit Keep Alive

=====

A SCSI on ST endpoint may implement ST Virtual Connection keep-alive sequences as described in the ST specification.

If an ST Virtual Connection is discovered to be dead by a failure of the keep-alive test, all SCSI on ST transactions on the Virtual Circuit shall be declared dead and their resources reclaimed, as well as the resources for the Virtual Connection itself.

3.2 Operation Timeouts

=====

A SCSI on ST initiator shall maintain timeouts on all outstanding device and task management operations from the time a Request_To_Send, Request_To_Receive or Zero_Length_Command operation is sent until the final Data and SCSI on ST status Put operations are received.

The operation timeout shall be at least the sum of:

- 1) the maximum time required to perform the SCSI command and transfer the associated data
- 2) the maximum time an ST operation can remain in the network

This timeout ensures that no ST operations remain in the network or will be subsequently generated by a SCSI on ST target for a SCSI on ST transaction, which guarantees that a SCSI on ST initiator may reuse

the I-id associated with the SCSI on ST transaction immediately.

If an operation timeout expires, the initiator shall first determine the viability of the SCSI on ST Virtual Connection, with an ST Request_State/Request_State_Response sequence. If a Request_State_Response operation is not received within an appropriate period of time, the Request_State/Request_State_Response sequence shall be retried an appropriate number of times, after which the Virtual Connection shall be declared dead, and its resources reclaimed.

If a Request_State_Response is received, the operation which timed out shall be aborted with an ST End sequence. If an ST End_Ack operation is not received within the maximum an ST operation can remain in the network, the ST End sequence shall be retried an appropriate number of times, after which the SCSI on ST transaction shall be declared dead and its resources reclaimed.

A SCSI on ST target shall not maintain timeouts on outstanding transactions. A SCSI on ST target's resources for a transaction shall only be reclaimed when:

- 1) the target completes the transaction
- 2) the initiator ends the transaction with an ST End operation
- 3) the SCSI on ST Virtual Connection is ended

The requirement that the SCSI on ST initiator perform an ST End sequence for timed out operations ensures that a SCSI on ST target can deterministically reclaim resources associated with a failed transaction.

\$Id: scsiOnSt.txt,v 1.9 1999/03/03 21:53:10 steph Exp steph \$