Accredited Standards Committee
X3, Information Processing Systems

Doc: X3T10.1/96a134r1
Date: June 24, 1996
Project: X3T10.1 / 1147
Ref Doc.: SSA-TL2
Reply to: Adge Hawes

To: X3T10.1 Membership
From: Adge Hawes

Subject: Third-party Data Transfers

BACKGROUND

Conventional data transfer in an SSA web is between an Initiator and a Target. An Initiator usually sends a read or write command to a Target, and the Target responds by accepting data from (in the case of a Write), or sending data to (in the case of a Read), the Initiator. Flow of the data is governed by the three Application SMSs (the Data SMSs) that allow the Target, with its limited buffer resources, rather than the Initiator, to dictate when the transfer takes place.

When a Read operation occurs, the Target knows where to send the data because it was told the return path necessary to reach the Initiator when the Initiator 'registered' itself with the Target during the configuration process. For a Write operation, the Target does not care where the data came from, it simply arrives at its port (with the appropriate Tag and Return Path ID) when its address reaches zero.

There is, however, a desire for Third-Party transfer of data to take place. Third party data transfers occur when data is transferred between two SSA agents, neither of which are Initiators. Examples of 3rd party data transfer are:

- Third Party Copy (SCSI command)
- Tape Backup, where the tape is part of the web (no Initiator involvement in the data transfer)
- XORing drives, where the Targets themselves reduce the amount of link traffic by performing some of the fundamental RAID operations.

Currently there is no mechanism within SSA to allow this to happen, because there is no way to instruct a target to transfer to a different address than the Initiator that originated the command. This document proposes a new SMS to allow 3rd Party transfers to take place.

Previous suggestions for achieving 3rd-party transfer have involved instructing the data source node to become a (temporary) Initiator, and executing a conventional WRITE transfer to the third party. While this has the benefit of simplicity, it does have the drawback of burdening the data source node (which spends most of its life being a mere Target) with the added duties of Initiatorship, i.e. extra SMS support, data flow control, and error handling. It is the intent of this proposal to keep the Initiator duties firmly within the Initiator that issued the third party transfer, and require nothing more than Data redirection from the Data source node.

PROPOSAL

The proposal is to introduce a new application SMS for TL2 called DATA REDIRECT. This new SMS is based on DATA REPLY but does not replace it. The reasons for not replacing DATA REPLY are:

1. TL1 level devices may still be sent the old DATA REPLY messages,
2. Even third-party-capable devices will probably execute non third-party transfers most of the time, and may therefore use DATA REPLY.

DATA REDIRECT has a similar field structure to DATA REPLY, but adds information that allows the Target to return its data on another path, not back to the Initiator. This extra information consists of:

1. The port that the Target should use to send the data out, and

2. The path component that the Target should place in the address field.

Therefore, the following description should be added after 11.3.3, DATA REQUEST SMS, and Table 22 and section 11.3.1 should be updated as shown.

DATA REDIRECT SMS

The DATA REDIRECT SMS defined in Table 53 is sent in response to a DATA READY SMS, when the data is to be sent to a node other than the one that issued the command. The node receiving the DATA REDIRECT SMS sends the requested Data frames from the port specified and along the path indicated in the SMS. More than one DATA REDIRECT SMS may be sent in response to a single DATA READY SMS. A DATA REDIRECT SMS is considered outstanding from the time it is sent until the last Data frame associated with it is received by the third party.

Table 53 - DATA REDIRECT SMS

Byte	Bit 7	6	5	4	3	2	1	Bit 0
0	SMS CODE (23h)							
1	reserved							
2	TAG							
3	TAG							
4	RETURN PATH ID							
5	RETURN PATH ID							
6	RETURN PATH ID							
7	RETURN PATH ID							
8	BYTE COUNT							
9	BYTE COUNT							
10	BYTE COUNT							
11	BYTE COUNT							
12	CHANNEL							
13	CHANNEL							
14	TRANSFER ID							
15	TRANSFER ID							
16	OUT PATH							
17	OUT PATH							
18	OUT PATH							
19	OUT PATH							
20	OUT PORT							

The TAG field is copied from the associated DATA READY SMS.

The RETURN PATH ID field identifies the node that issued the DATA REDIRECT SMS. The RETURN PATH ID field and TAG field serve to identify the I/O process. If the RETURN PATH ID field does not match the corresponding field of the ULP command SMS, an Asynchronous Alert is generated with an ALERT CODE value of INVALID RETURN PATH OR RETURN PATH ID. The node then waits for another DATA REDIRECT SMS or until the I/O process is terminated. The RETURN PATH ID field specifies the node issuing the DATA REDIRECT SMS, since tags may be redundant across Configutor nodes.

If the modulo eight rules regarding the BYTE COUNT field are violated, the SMS is rejected by invoking the Asynchronous Alert process with an ALERT CODE value of INVALID FIELD.

The BYTE COUNT field specifies the number of bytes that the node is willing to accept during the transfer associated with this DATA REDIRECT SMS. If it is less than the BYTE COUNT field of the corresponding DATA READY SMS, then the node shall send another DATA REDIRECT SMS when it is ready to accept more data. The value of the BYTE COUNT field shall be a multiple of eight unless this SMS transfers the last byte requested by the command. If the modulo eight rules regarding the BYTE COUNT field are violated, the SMS is rejected by invoking the Asynchronous Alert process with an ALERT CODE value of INVALID FIELD.

The CHANNEL field specifies the channel component of the ADDRESS field for the Data frames. The path component is obtained from the OUT PATH field.

The TRANSFER ID field is a unique identifier that was generated by the node that sent the DATA READY SMS. The value is unique among all outstanding DATA READY SMSs between a given pair of nodes. The node receiving the DATA READY SMS shall copy the TRANSFER ID into all associated DATA REDIRECT SMSs to associate the DATA REDIRECT SMS to a DATA READY SMS.

The OUT PATH field specifies the path component of the ADDRESS field for the Data frames. The channel component is obtained from the CHANNEL field.

The OUT PORT field specifies the port that should be used to send the Data frames.

The DATA REDIRECT SMSs for a single DATA READY SMS may contain different CHANNEL values. If two consecutive DATA REDIRECT SMSs for the same DATA READY SMS contain the same CHANNEL value, then the node shall complete the first DATA REDIRECT SMS's transfer before beginning the second DATA REDIRECT SMS's transfer. If two consecutive DATA REDIRECT SMSs for the same DATA READY SMS contain different CHANNEL values, the node may overlap the data transfers.

In the case of an invalid field other than RETURN PATH ID, then the DATA REDIRECT SMS is acknowledged with a RESPONSE SMS with a RETURN CODE of INVALID FIELD. In the case of an invalid RETURN PATH ID field value, an SMS UNEXPECTED, or TOO MANY SMS OUTSTANDING, an ASYNC ALERT SMS is generated with the appropriate meaning codes. The associated I/O processes identified by the tag value (if any) is not terminated, instead the DATA REDIRECT SMS is rejected with a RESPONSE SMS or an ASYNC ALERT SMS.

Update table 22 as shown below:

The SMS codes that are valid for SSA-TL2 and associated frame type and node type support is defined in Table 22.

Table 22 - SSA-TL2 messages supported

SMS	SMS CODE	SMS FRAME TYPE	Node type support	
			Sent from	Received by
QUERY NODE	00h	Privileged	Configutor	all
QUERY NODE REPLY	01h	Privileged	all	Configutor
CONFIGURE PORT	02h	Privileged	Master	all
RESPONSE	03h	Privileged	all	Configutor
ASYNC ALERT	04h	Privileged	all	Master
MASTER ALERT	05h	Privileged	Master	Configutor
QUIESCE	06h	Privileged	Configutor	all
ASYNC REPLY	07h	Privileged	Master	all
QUERY PROTOCOL	08h	Privileged	Configutor	all
QUERY PROTOCOL REPLY	09h	Privileged	all	Configutor
QUERY PORT	0Ah	Privileged	Configutor	all
QUERY PORT REPLY	0Bh	Privileged	all	Configutor
QUERY SWITCH	0Ch	Privileged	Configutor	all
QUERY SWITCH REPLY	0Dh	Privileged	all	all
REQUEST SAT REGION	0Eh	Privileged	Configutor	Master
REQUEST SAT REGION REPLY	0Fh	Privileged	Master	Configutor
reserved for Privileged frames	10h-1Fh	Privileged	reserved	reserved
DATA READY	20h	Application	all	all
DATA REPLY	21h	Application	all	all
DATA REQUEST	22h	Application	all	all
DATA REDIRECT	23h	Application	all	all
reserved for Application frames	24~~3~~h-7Fh	Application	reserved	reserved
Defined by ULP	see 11.2.6	defined in ULP		

Also update section 11.3.1, DATA READY SMS as shown below:

DATA READY SMS

The DATA READY SMS defined in Table 50 is sent before the outbound transfer of data. The receiving node returns one or more DATA REPLY SMSs (or DATA REDIRECT SMSs if third party transfers are required) for each DATA READY SMS. The node optionally uses more than one DATA READY SMS to transfer all of the data for a I/O process. A DATA READY SMS is considered outstanding from the time it is sent until the last Data frame associated with it is sent.

Table 50 - DATA READY SMS

Byte	Bit 7	6	5	4	3	2	1	Bit 0
0	SMS CODE (20h)							
1	reserved							
2	TAG							
3	TAG							
4	BYTE OFFSET							
5	BYTE OFFSET							
6	BYTE OFFSET							
7	BYTE OFFSET							
8	BYTE COUNT							
9	BYTE COUNT							
10	BYTE COUNT							
11	BYTE COUNT							
12	reserved							
13	reserved							
14	TRANSFER ID							
15	TRANSFER ID							

The TAG field contains a copy of the TAG field in the ULP command SMS and allows the node to associate this SMS with the correct I/O process.

The BYTE OFFSET field indicates the byte offset of the first byte to be transferred, relative to first byte requested by the I/O process. The value of the BYTE OFFSET field shall be a multiple of eight.

The BYTE COUNT field specifies the number of bytes to be transferred. The value of the BYTE COUNT field shall be a multiple of eight unless this SMS transfers the last byte requested by the command.

The sum of the BYTE COUNT and BYTE OFFSET fields in a DATA READY SMS shall not exceed the transfer length of the I/O process. Checking for violation of this rule is optional.

> Implementer's note 1: Since data maybe retransmitted, the Configutor node cannot compare the sum of the BYTE COUNT fields in all the associated DATA READY SMSs with the expected total byte count of the I/O process.

The TRANSFER ID field is a unique identifier generated by the node to identify the DATA READY SMS. The value shall be unique among all outstanding DATA READY SMSs between a given pair of nodes. The node receiving the DATA READY SMS shall copy the TRANSFER ID into all associated DATA REPLY or DATA REDIRECT SMSs to associate the DATA REPLY or DATA REDIRECT SMS to a DATA READY SMS.

The protocol for the DATA READY and DATA REPLY (or DATA REDIRECT) SMSs is as follows if the DDRM bit is cleared:

1) When the node is ready to transfer data it sends a DATA READY SMS. This specifies how many bytes it currently has available and the starting offset within the data requested by the command along with a transfer ID.
2) If the BYTE COUNT, BYTE OFFSET, or TRANSFER ID do not conform to the rules above, then the receiving node aborts the I/O process as allowed by the upper level protocol. Otherwise the receiving node allocates a channel to receive the data.

3) If the data is to be sent to a third party, the node that received the DATA READY SMS waits to receive a DATA REQUEST SMS from the third party node. This DATA REQUEST SMS indicates how many bytes the third party node is willing to accept, and the channel it wishes to accept them on.
4) The node receiving the DATA READY SMS responds with a DATA REPLY or DATA REDIRECT SMS indicating the CHANNEL and BYTE COUNT that it or the third party is willing to currently accept (and OUT PATH and OUT PORT in the case of third party transfers) and copies the TRANSFER ID from the associated DATA READY SMS.
5) The node sends Data frames addressed to the specified channel, and (in the case of a DATA REDIRECT) to the correct path.
6) If all of the data offered by the previous DATA READY SMS has not been transferred then the node returns to step 3 using the same or a different channel. The node may send a DATA REPLY or DATA REDIRECT SMS for an outstanding DATA READY SMS whenever there is at most one outstanding DATA REPLY or DATA REDIRECT SMS for that DATA READY SMS. A DATA REPLY or DATA REDIRECT SMS shall not span multiple DATA READY SMSs.
7) If all of the data requested by the I/O process has not been transferred then the node returns to step 1.
8) The node may send a DATA READY SMS when either there are no outstanding DATA READY SMSs for the command, or, at least one DATA REPLY or DATA REDIRECT has been received for each DATA READY that is outstanding for the command.

A Data frame shall not span multiple DATA REPLY SMSs, DATA REDIRECT SMSs or DATA REQUEST SMSs.