Accredited Standards Committee
X3, Information Processing Systems

Doc: X3T10.1/96a145r3
Date: September 16, 1996
Project: X3T10.1 / 1147D
Ref Doc.: SSA-TL2 rev 2
Reply to: Mark DeWilde

To: X3T10.1 Membership
From: Mark DeWilde

Subject: TL2 Configutor Table Management Support

BACKGROUND

This proposal updates R1, by adding changes suggested by John Scheible in his reflector comments. The original proposal was the result of an effort to simplify 95A210R2, entitled “Interlocked Election of Master”. This proposal is being made as a separate document, since it is not a part of the election process itself, but will facilitate the simplification of the process and can stand alone on it’s own right. The interlocked process described in 95a210r2 is complicated excessively by the necessity to track all possible web changes in order not to lose knowledge of missing Configutors. This necessity is caused by the fact that the only current methods of releasing initiator table space are through the use of resets and the third party quiesce. Only master nodes that have received the Asynchronous Alerts indicating the loss of a Configutor have the necessary information to perform this function. These Master nodes may, however, lose mastership prior to being able to perform the necessary third party quiesces. The result is loss of available Configutor table space on targets, which may only be recovered by a reset. This proposal seeks to add the additional SMSs and privileged data frame support necessary for a “new” master lacking historical knowledge of the web to query and repair Configutor tables on target nodes without the more obtrusive node reset method being employed. This revision captures the results of discussions on the previous revision and simplifies the protocol changes required.

John Scheible took revision 2 of this proposal and formatted it in the proper format for SSA-TL2. The resulting revision 3 is shown below. The main changes are that the DELETE RETURN PATH ID SMS has a RETURN PATH field in bytes 4-7 to be used to return the resulting RESPONSE SMS, and a description of the function of this SMS (along with all the field descriptions). This proposal uses Configutor table STATUS field values defined in 96a163r1.

PROPOSAL

Add SMS support for the following new functions:

a) Query Registration: Requests that a target prepare to send it’s Configutor table and provide its size
b) Delete Return Path ID: Deletes specific table entry from Configutor table of target

Define a Configutor table cleanup process which becomes a duty of the master assuming mastership of the web. This process is not a time-critical process, but may proceed as a background process since it serves to recover lost initiator table space.

Add a flag to Configutor table entries that is set to indicate that the table entry has been reported to the master in response to the query registration SMS. Clear this flag if the table entry has been changed in any way since the last reporting.

The following section is a set of suggested changes to the TL2 document to implement this proposal.

a) Add the following two new SMSs to clause 11.
b) Add the following two new SMSs to Table 21.
c) Add the following Configutor Table Maintenance process to clause 12.
d) Add the Configutor Table Maintenance process to the list under clause 12.

11.2.x QUERY REGISTRATION SMS

The QUERY REGISTRATION SMS defined in Table 1 is used by the Master to set up a target node to transfer its current Configutor table information back to the Master, and to request the target to provide the Master with the number of bytes that will be transferred.

Table 1 - QUERY REGISTRATION SMS

Byte	Bit 7	6	5	4	3	2	1	Bit 0
0	SMS CODE (10h)							
1	reserved							
2	TAG							
3	TAG							
4	RETURN PATH							
5	RETURN PATH							
6	RETURN PATH							
7	RETURN PATH							

The TAG field is returned in the QUERY REGISTRATION REPLY SMS. The TAG is assigned by the Configutor node and it shall be unique among the TAG values that are currently active from that Configutor node.

The RETURN PATH field specifies the path component that shall be placed in the path component of the ADDRESS field of the resulting QUERY REGISTRATION REPLY SMS and DATA READY SMS.

When a node receives a QUERY REGISTRATION SMS, it prepares to send its Configutor table data to the Configutor requesting it. The node sends a DATA READY SMS indicating the total table data size in the BYTE COUNT field. The node shall set the BYTE OFFSET field to zero. In response to the DATA READY SMS, the node requesting the table data sends a DATA REPLY SMS to begin the data transfer. The Configutor Table Data is sent as a set of records, one for each table entry. Each record contains 18 bytes, so that a data frame may contain up to 8 table entries. The format for a record and a data frame is shown in Table 2.

Table 2 - Configutor table returned data entry format

Byte	Bit 7	6	5	4	3	2	1	Bit 0
0	RETURN PATH ID							
1	RETURN PATH ID							
2	RETURN PATH ID							
3	RETURN PATH ID							
4	RETURN PATH							
5	RETURN PATH							
6	RETURN PATH							
7	RETURN PATH							
8	UNIQUE ID							
9	UNIQUE ID							
10	UNIQUE ID							
11	UNIQUE ID							
12	UNIQUE ID							
13	UNIQUE ID							
14	UNIQUE ID							
15	UNIQUE ID							
16	PSTATUS	PORT-1						
17	reserved							

The PSTATUS bit is cleared if the table entry was valid, but has since been quiesced and is reserved for the Configutor identified by the UNIQUE ID field value. The PSTATUS bit is set if the table entry is currently valid, and the RETURN PATH ID field value is assigned to the Configutor node identified by the UNIQUE ID FIELD value and can be addressed by the RETURN PATH field value on the port identified by one more than the PORT-1 field value.

11.2.y DELETE RETURN PATH ID SMS

The DELETE RETURN PATH ID SMS defined in Table 3 is used by the Master to request the target to delete the Configutor table entry identified by the port number and Return Path ID indicated in the SMS.

Table 3 - DELETE RETURN PATH ID SMS

Byte	Bit 7	6	5	4	3	2	1	Bit 0
0	SMS CODE (11h)							
1	reserved							
2	TAG							
3	TAG							
4	RETURN PATH							
5	RETURN PATH							
6	RETURN PATH							
7	RETURN PATH							
4	RETURN PATH ID							
5	RETURN PATH ID							
6	RETURN PATH ID							
7	RETURN PATH ID							
16	reserved	PORT-1						

The TAG field is returned in the RESPONSE SMS. The TAG value is assigned by the Configutor node and it shall be unique among the TAG values that are currently active from that Configutor node.

The RETURN PATH field specified the Path component that shall be placed in the Path component of the ADDRESS field of the resulting RESPONSE SMS.

The RETURN PATH ID field is the RETURN PATH ID field value sent by the target node to the Master as a result of the SEND REGISTRATION SMS that was identified by the Master as Invalid.

The PORT-1 field is the PORT-1 field value sent by the target node to the Master as a result of the SEND REGISTRATION SMS that was identified by the Master as Invalid.

The destination node shall find the Configutor table entry that meets the following conditions:

a) matches the Return Path ID identified by the RETURN PATH ID field value;

b) it is assigned to the port identified by the PORT-1 field value,

c) it has a REPORTED status (see 9.2.4).

If a matching entry is found, then the STATUS field of the entry is changed to FREE and a RESPONSE SMS is sent with a RETURN CODE field value of REQUESTED FUNCTION WAS COMPLETED SUCCESSFULLY. If a matching entry is not found, then a RESPONSE SMS is generated with a RETURN CODE field value of REQUESTED FUNCTION FAILED. In either case, the TAG field value is used in the RESPONSE SMS and the RESPONSE SMS is sent to the address indicated by the RETURN PATH field value and the port the DELETE RETURN PATH ID SMS was received on.

12.4 Configutor Table Maintenance process

The Master shall begin the Configutor Table Maintenance Process when either of the following events has occurred, and will perform the maintenance on the set of nodes indicated:

a) A new Master takes over the Web, and process all responders in the web.

b) A new Sub-Web has been added to the web, and process the responders in the new sub-web

This process is a low priority process that is run after the Master has built it's Topology table of the web and after all ports have been placed in normal mode. The Master performs the following operations on each node in the system that has a Configutor table, and has come under the Master's control since the completion of its last Configutor Table Maintenance process:

a) The Master sends a QUERY REGISTRATION SMS to the target node and waits for the reply.

b) When the Master receives the associated Data Ready SMS, it reserves buffer space for as much of the Configutor table space as possible, and sets up a channel to receive the data. The Master reserves space for a clean-up list which will contain the Return Path IDs and port numbers for entries to be deleted.

c) The Master sends a Data Reply SMS to the target providing it with the channel number and byte count for the transfer.

d) When the Master receives the table data, it verifies that each Return Path ID points to the correct Unique ID node relative to the indicated port number. If it does not, then the Master adds the Return Path ID and port number to the clean-up list. The responder sets a flag in each table entry indicating it has been sent to the Master in response to a QUERY REGISTRATION SMS. This flag is cleared if the table entry is subsequently changed via a quiesce, Total Reset or Absolute Reset frame.

e) The Master sends additional Data Reply SMSs with offsets and byte counts necessary to complete the total table contents as indicated by the Data Ready SMS. On each portion of the table retrieved, step d is repeated.

f) The Master sends DELETE RETURN PATH ID SMSs to the target node for each entry in the cleanup list.

g) The Master waits for a RESPONSE SMS for each DELETE RETURN PATH ID SMS that it sent.

h) The Master sets an internal flag to indicate that the target node has been maintained since the Master took control of the node.

Sincerely,

Mark A. DeWilde
Principal System Architect
Pathlight Technologies
Voice: (607)266-4000 X-403
FAX: (607)266-0352
Email: mark@pathlight.com

and

John Scheible
Senior Engineer/Scientist
IBM Corporation
Voice: (512) 823-8208
FAX: (512) 838-3822
Email: Scheible@vnet.ibm.com