

Accredited Standards Committee
X3, Information Processing Systems

Doc: X3T10.1/96a145r1
Date: 10 May, 1996
Project: 1147D
Ref Doc.: X3T10.1/1147D
Reply to: Mark DeWilde

To: X3T10.1 Membership
From: Mark DeWilde
Subject: TL2 Configurator Table Management Support

BACKGROUND

This proposal is 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 configurators. 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 async alerts indicating the loss of a configurator 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 configurator 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 configurator 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.

PROPOSAL

Add SMS support for the following new functions:

Query Registration: Requests that a target node prepare to send it's configurator table and provide its size

Delete Return Path ID: Deletes specific table entry from configurator table of target

Define a configurator 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.

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

SUGGESTED CHANGES

Add the following sections:

11.2.15 QUERY REGISTRATION SMS

The QUERY REGISTRATION SMS defined in table XX is used by the Master to set up a target node to transfer its current configurator 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 XX- QUERY REGISTRATION SMS

Byte	Bit 7	6	5	4	3	2	1	0
0	SMS CODE(0Eh)							

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 Configurator node and it shall be unique among the TAG values that are currently active from that Configurator 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.

When a node receives a Query Registration SMS, it prepares to send its configurator table data to the Configurator 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 offset field to zero. In response to the Data_Ready, the node requesting the table data sends a Data Reply to begin the data transfer. The Configurator 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 8 table entries. The format for a record and a data frame is shown in table YY.

TABLE YY. CONFIGUTOR TABLE RETURNED DATA FORMAT

Byte	table entry number	Field Contents	
0-3	1	Return Path ID	
4-7	1	Return Path	
8-15	1	Unique ID	
16	1	Port Status	Port Number
17-127	2-8	repeats of table records for other table entries	

The port status field is set to 0 if the Unique ID has been quiesced without the ITF bit set, or 1 if the entry represents an a valid return path ID from the responder’s point of view.

11.2.18 DELETE RETURN PATH ID SMS

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

TABLE YY- DELETE RETURN PATH ID SMS

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

5	RETURN PATH ID
6	RETURN PATH ID
7	RETURN PATH ID
8	PORT
9	reserved

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

The RETURN PATH ID field is a RETURN PATH ID 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 field identifies the port in the target node that is related to the invalid return path id.

Add the following section to section 12:

12.3 Configurator Table Maintenance Process

Whenever there has been a change in the web configuration or a change of Master, the Master shall begin the Configurator Table Maintenance Process. This process is a low priority process that is run after the master has built its 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 configurator table, and has come under the Master's control since the completion of its last Configurator table Maintenance Process:

- 1) The Master sends a QUERY REGISTRATION SMS to the target node and waits for the reply.
- 2) When the Master receives the associated Data Ready SMS, it reserves buffer space for as much of the configurator 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.
- 3) The Master sends a Data Reply SMS to the target providing it with the channel number and bytecount for the transfer.
- 4) 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.
- 5) 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 4 is repeated.
- 6) The Master sends DELETE RETURN PATH ID SMSs to the target node for each entry in the cleanup list.
- 7) The master waits for a RESPONSE SMS for each DELETE RETURN PATH ID SMS that it sent.
- 8) 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