Linear 1:1 Trunk Protection Switching

Zehavit Alon / Nokia Siemense Networks

January 2008

Introduction

PBB-TE supports 1:1 bi-directional linear protection. To guarantee protection, two PBB-TE trunks[1] are provisioned: one trunk is configured as the "working" trunk and the other as the "protection" trunk. In normal conditions, traffic is transmitted over the "working" trunk. Normal conditions imply that no failure has occurred on the "working" trunk and that there is no administrative configuration or request which would cause traffic to be transmitted over the "protection" trunk. In the event of either a failure or a specific administrative request, traffic is switched to the "protection" trunk.

Protection switching makes use of CCMs, as defined in 802.1ag. CCM is the standard Ethernet mechanism that detects and signals connectivity failures.

The protection switching mechanism requires that both the "working" and "protection" trunks are monitored; this is realized by CCMs. Each of the PBB-TE trunks is monitored by an independent Maintenance Association. One MA is set to monitor the "working" PBB-TE trunk, while another MA is set to monitor the protection PBB-TE trunk. Each MA contains two MEPs: one is located in the CBP on the near trunk edge while the other is located on the CBP on the far trunk edge. The absence of CCM from the far edge or specific information received in one of the CCM's TLVs indicates to the near edge that trunk connectivity has been disrupted. While the near edge can determine a failure in connectivity, the far edge must be notified about this condition. The near edge signals the failure condition to the far edge by sending a CCM message with an RDI flag (as defined in 802.1ag).

When a failure is detected in the PBB-TE trunk, trunk protection switching is automatically executed. Since both trunk edges are notified of the failure (either because of missing CCMs or specific information in the CCMs, or owing to receipt of an RDI) both edges switch traffic over the same PBB-TE trunk (even if the failure occurred in one direction only).

When the failure is eliminated, traffic may be switched back to the "working" trunk according to the type of configuration (revertive/non-revertive mode).

The main purpose of automatic protection switching is to guarantee the availability of resources in the event of failure, and to ensure that switchover is achieved in less than 50 ms so that the network will not be affected. Nevertheless, certain maintenance operations can benefit from network flexibility and from the ability to transmit traffic over either trunks (either over the "working" trunk or the "protection" trunk). For these reasons, the mechanism must enable traffic to be manually switched over from one trunk to the other.

[1] See PBB-TE Trunk definition in clause 3 of 802.1Qay /D1.1

Management Requests

The protection switching mechanism must allow for manual operation by the network operator, regardless of the network state (i.e. even when the network is performing without problems). Examples and requirements for such maintenance operations can be found in: http://www.ieee802.org/1/files/public/docs2007/ay-roese-APS-protocol-1107-v01.pdf.

Different management tasks require different switching capabilities. Deterministic network behavior is needed to handle each task and a relevant management command should be supported.

The required management switching capabilities are defined in:
http://www.ieee802.org/1/files/public/docs2007/ay-mcguire-linear-121-protsw-0709-v1.pdf
These support the following operations:

- lockout protection
- manual switch

The corresponding administrative management commands should be:

1. Lockout protection / Clear Lockout Protection (disable/enable protection switching functionality)
2. Manual Switch to Working / Protection (switched only if the trunk's operational state is up)
3. Force Switch to Working / Protection (switched without checking the trunk's operational state)

The trigger for protection switching should be unambiguous for both ends, thus there should be a mechanism to indicate to the far edge the exact trigger to switch over in order to avoid undesirable behavior.

CCM

The existing CCM mechanism should be enhanced to signal a switchover operation via an administrative management request. This can be done by defining an optional TLV in the CCM message and by providing an explicit indication that this TLV appears in the CCM message.

The existence of the new, optional TLV is indicated by a bit in the Flags field which is defined in the common header of the CFM PDUs. According to the 802.1ag, the use of the Flags is defined differently for each CFM PDU type.

For the CCM PDU, the Flags field of the Common CFM Header is currently split into three parts:

1. The RDI field (one bit, the most significant bit)
2. The Reserved field (4 bits)
3. The CCM Interval field (The least significant three bits of the Flags field constitute the CCM Interval field.).

It is proposed that one of the reserved bits be used to indicate that an administrative management request is included in the CCM. In this document, this bit is referred to as the "ADMIN" bit.

.

A new optional "Admin" TLV (as defined below) should be used to indicate a specific administrative management request. This TLV should be added to the CCM only when an administrative management request is received and carried out by the node.

The format of the "Admin" TLV structure is defined as follows:

Type = 5
Length
Value - See table below

The following lists the valid values for the "Value" field:

mnemonic	**Administrative Management Request**	**Value**
LoP	Lockout Protection	1
CLoP	Clear lockout protection	2
MStP	Manual Switch to protection	3
MStW	Manual Switch to working	4
FStP	Force Switch to protection	5
FStW	Force Switch to working	6
CS	Clear the manual switch command	7

Whenever an administrative management request is received by the MEP at the near edge, the following operations should be performed:

1. Execute the appropriate operation (e.g. force switch, etc.). If the command was executed successfully, perform the following:
    a. Set the protocol version in the CCM to version 2. (In version 1, the reserved bits in the Flag field are ignored.)
    b. Set the "ADMIN" flag in the CCM to "true".
    c. Add the "ADMIN" TLV with the appropriate value (as defined in the table above) to the CCM.

When the MEP on the far edge receives a CCM, it should perform the following:

1. Inspect the CCM flags. If the "ADMIN" flag is set to "true", it should look for the "ADMIN" TLV.
2. Perform the administrative management request, as indicated by the management TLV.

Conclusion

The mechanism defined in this paper addresses the requirements for protection switching and is based on the functions and messages for as defined in 802.1ag, with a minor extension to the CCM. There is no need to signal any further information between the two edges of the trunk.

At the IEEE 802.1 interim meeting on September 2007, it was proposed that the APS protocol be used to signal management requests (as defined in ITU-T G.8031). As clearly indicated in http://www.ieee802.org/1/files/public/docs2007/ay-martin-protection-0907-v03.pdf, most of the APS information in the APS PDU is redundant for PBB-TE. Nevertheless, nodes will need to create a new PDU and to parse it, even though most of the data in the PDU will be ignored by the receiving nodes.
Out of the four octets of APS-specific information, only three values are defined for PBB-TE (three values out of the four bits from the first octet (Request state) are valid - 1111 - Lockout of protection, 1101 – forced switch, 0111 - manual switch.). The other values of this request state are ignored as well as the next four bits of the first octet and the other three octets, as the protection type is fixed to 1:1 bi-directional.

The APS protocol is asymmetrical, since all the messages are carried by the "protection" trunk only. There is no such requirement for PBB-TE, and the ability to send the messages over both trunks is an advantage. This limitation is reflected in the switching capabilities of APS where a force and a manual switch can only be triggered from the "working" trunk to the "protection" trunk, and not from the "protection" trunk to the "working" trunk.

In the proposed mechanism, switchover is performed by the CCM state machine immediately after a failure is detected or signaled. However, in APS, another state machine controls the actual switching and the CCM is only used for failure detection. Involving two state machines may increase the delay between the detection of a failure and the execution of a switchover.
Using the existing 802.1ag messages with a new TLV to signal the three required indications is a much simpler and risk-free method compared ti the APS state machine which is complicated and error-prone

The entire operation of PBB-TE depends on the correct and logical configuration of the network (PBB-TE VID range, the trunks ("working" and "protection"), etc.). The correct and consistent configuration of the network is the responsibility of the network administrator. When all configurations are performed correctly, there is no need to signal configuration information (like revertive/non-revertive mode) to verify consistency.

It may be misleading to use the APS PDU, since most of the information in the PDU is not relevant for PBB-TE, and one might infer that the entire protocol and state machine are implemented. Please note, the ITU-T is working on protection switching and can use the full-blown APS. What should a node do when it receives an APS message with a request that it does not recognize?