Proposed Specification of PBB-TE Protection with Load Sharing

26.10.4 Protection switching with load sharing for point-to-point TE service instances

Protection without load sharing carries all backbone service instances on one TE service instance at a time, either the working TE service instance or the protection TE service instance. Thus all of the traffic transported by the protection group follows one path through the network and two paths must be provided, each with full traffic capacity reserved.

Protection with load sharing allows the backbone service instances carried by the protection group to be distributed over a number of TE service instances. The working transport entity is thus a set of TE service instances, as shown in Figure 26-xx, which may take diverse paths across the network. The CBP distributes the backbone service instances carried by the protection group across these working TE service instances. Each TE service instance must provide only a fraction of the total traffic capacity for the group to carry its share of the working traffic.

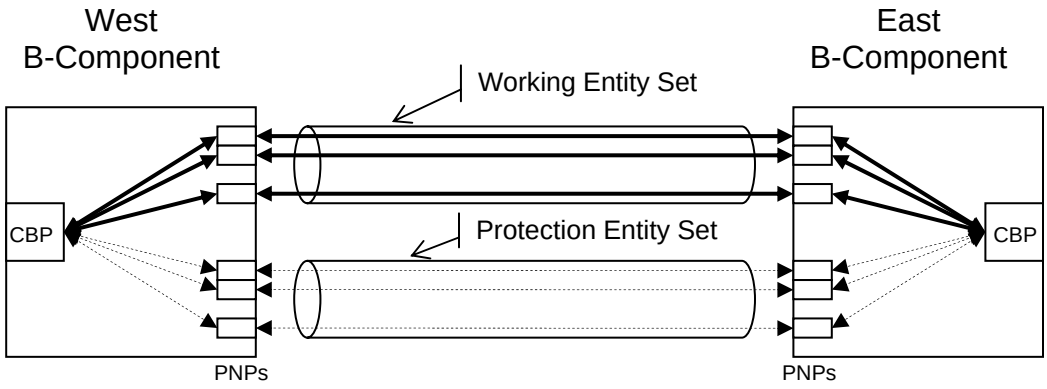

Figure 26-xx — Working and Protection Sets

In the presence of a failure, or an operator switch request, all traffic from the affected TE service instance is moved to the protection transport entity which, as with the working entity, is a set of TE service instances. The protection traffic (the set of backbone service instances normally carried by the affected TE service instance) is distributed across the set of protection TE service instances, as shown in Figure 26-yy. Each of the TE service instances in the protection set must provide only a fraction of the total traffic capacity of the affected TE service instance.

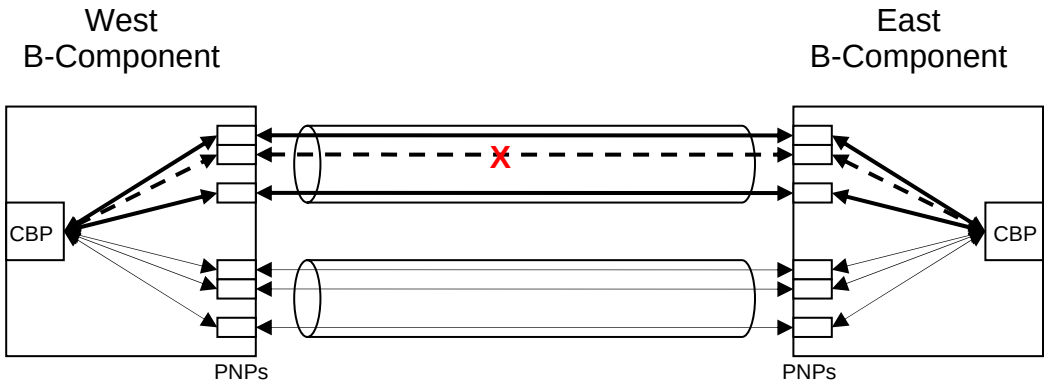

Figure 26-yy — Mapping Traffic to the Protection Entity

By distributing the traffic carried by the protection group in this way bandwidth requirements for the protection group can be reduced, providing more efficient service protection (see Annex M for some example use cases). Using the 1:1 protection model, only a single failure among the working set of TESI is protected. If multiple failures occur, some backbone service instances will be affected.

The working and protection TE service instance sets may be combined so that each TE service instance carries some normal traffic and also provides protection bandwidth for other members of the protection group, as shown in Figure 26-zz. This approach reduces the number of TE service instances required in the protection group without sacrificing the bandwidth advantage provided by load sharing.

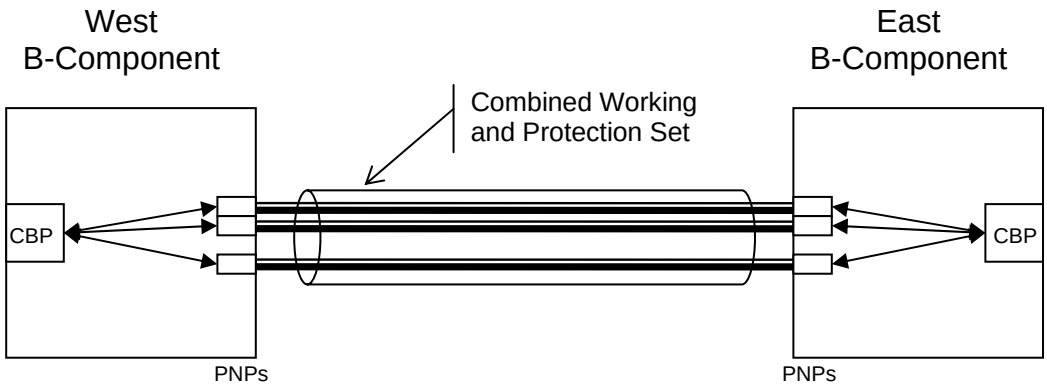

Figure 26-zz — A Combined Protection Group

Backbone service instances (BSIs), identified by I-SID, at a CBP may be assigned to a load sharing protection group. Each BSI assigned to a load sharing protection group has the following management information:

a) Working TE service instance – A configuration option that specifies a member of the working entity set for this BSI.
b) Protection TE service instance – A configuration option that specifies a member of the protection entity set for this BSI's working TE service instance to carry the BSI in the event its working TE service instance fails or is subject to an operator switch request.
c) Active TE service instance – Status option that specifies the current TE service instance carrying this BSI.

At most one TE service instance may have its traffic switched to its protection set at any given time. The highest priority request governs the protection switching action as described below.

26.10.4.1 Load sharing protection state machines

In load sharing mode, protection and traffic allocation can be handled by a single state machine for the protection group. However, to simplify the state diagram it is convenient to specify a state machine per protection group member. The protection group state machine is the set of member state machines operating together.

26.10.4.2 State machine timers

The load sharing mode uses the same state machine timers as defined in clause 26.10.3.2.

26.10.4.3 Protection switching variables

In the load sharing mode there may be more than two TE service instances in the protection group. Therefore some modifications are required to the original 1:1 control model to adapt requests that refer to transport entities by role (working or protection) to refer to a specific transport entity in the protection group (since role identification is not effective). These modifications are:

1) Force Switch (FS) and Lock Out (LO) are defined in terms of the working and protecting entities in the original model. However they have the same purpose – to move traffic off a transport entity, for example, to perform maintenance. In the load sharing just one request is needed for this purpose, targeted at a specific member of the protection group. So in the load sharing case LO of a transport entity corresponds to LO-P in its protecting role(s) and FS in its working role. Since only one request is needed to cover all the roles played by the transport entity, only LO is defined. The LOi request will unconditionally remove all traffic from TE service instance i in the protection group.

2) Similarly one MS command can serve to remove traffic from a transport entity conditionally, that is, as long as the entity is not needed in its protecting role(s). The MSi request will conditionally remove all traffic from TE service instance i in the protection group.

3) Finally, each Signal Fail (SF) condition is identified by the related TE service instance. So in the load sharing case there are n SF conditions SF1, SF2 ... SFn – one for each TE service instance in the protection group.

The protection requests have the same priority relationship as in the original 1:1 model, as shown in the following table:

Table 26-9—Protection Request Hierarchy

Request	Priority
Forced Switch/Lock Out (LOi)	highest
Signal Fail (SFi)	
Manual Switch (MSi)	
Wait to Restore (WTRi)	
Do Not Revert (DNRi)	
No Request (NR)	lowest

A request applies to one member of the protection group (TE service instance). For example, SFi indicates that protection group member i has a signal fail condition and a switch to that member's protection set is requested. NR indicates that there is no active request for the protection group.

At each end of the protection group the following variables are present:

a) one Local Request – whose value is one of LOx, MSx, NR.
b) one Remote Request/State – whose value is one of LOx, SFx, MSx, WTRx, DNRx, NR.
c) n local SF variables (SF1..SFn) – whose values are set or cleared by TE service instance connectivity monitoring.

Within the context of the priorities shown in the table above, a request targeted to lower numbered protection group member takes priority over the same kind of request targeted to a higher numbered member, and a local request takes priority over an identical remote request.

Based on the variables and priorities described above, one input request (Rin) is asserted to the set of state machines for the protection group which is the highest priority request that is currently set. One output request variable is set by the state machine (Rout) and this variable is communicated to the far end protection group state machine, appearing there as Remote Request/State. Whenever a state transition occurs Rout is set to Rin unless Rin is R-LOx, R-SFx, or R-MSx in which case Rout is set to NR.

If HoldOffTime is provisioned, SF is not asserted until after the Hold-Off timer expires.

26.10.4.4 Protection switching procedures

The following procedures are defined for the load sharing protection state machine:
a) mapDataToWorking(i) (26.10.4.4.1);

b) mapDataToProtection(i) (26.10.4.4.2).

26.10.4.4.1 mapDataToWorking(i)

Maps the customer service(s) that are to be transported by a TE protection group and whose working TE service instance is protection group member i to member i by:

a) Setting the VID value of the corresponding I-SID entry(ies) in the backbone service instance table to the ESP-VID for member i.

26.10.4.4.2 mapDataToProtection(i)

Maps the customer service(s) that are to be transported by a TE protection group and whose working TE service instance is protection group member i, to member i's protection set by:

a) Setting the VID value of the corresponding I-SID entry(ies) in the backbone service instance table to the ESP-VID of the protection TE service instance for that BSI.

26.10.4.5 Protection group member state machine

There is one state machine for each TE service instance in a protection group. The state machines use the variables that are presented in clause 26.10.4.3 and procedures described in clause 26.10.4.4. The state machines for revertive and non-revertive mode are shown separately for clarity.

In these state machine diagrams the following conventions are used:

- The predicate Rin == XX where XX is the highest priority input request is shown simply as XX.
- A Remote Request is shown as R-XX (e.g., R-LOi) and a Local Request is shown as L-XX (e.g., L-MSj).
- If a request predicate is shown as XX it is the logical "or" of the local and remote requests, for example, SFj == L-SFj || R-SFj.
- The subscript "i" refers to the TE service instance whose working traffic is controlled by the state machine and the subscript "j" refers to any other TE service in the protection group.

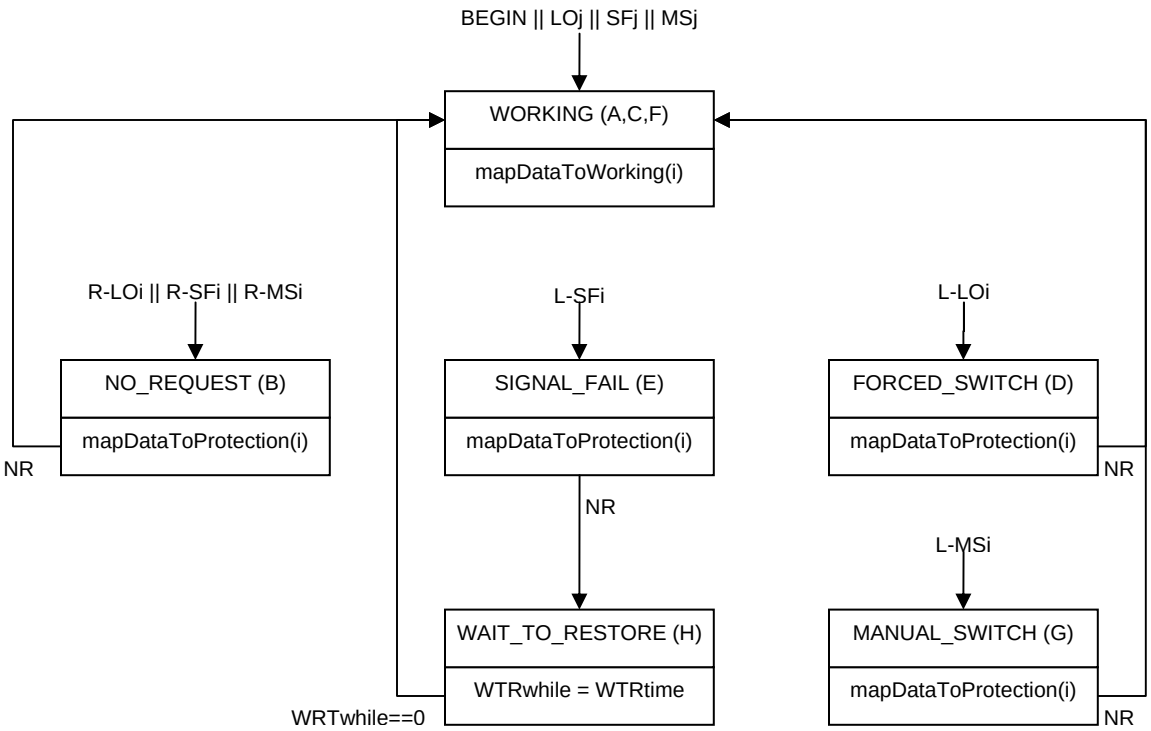

Figure 26-13 — 1:1 load sharing protection member state machine (revertive)

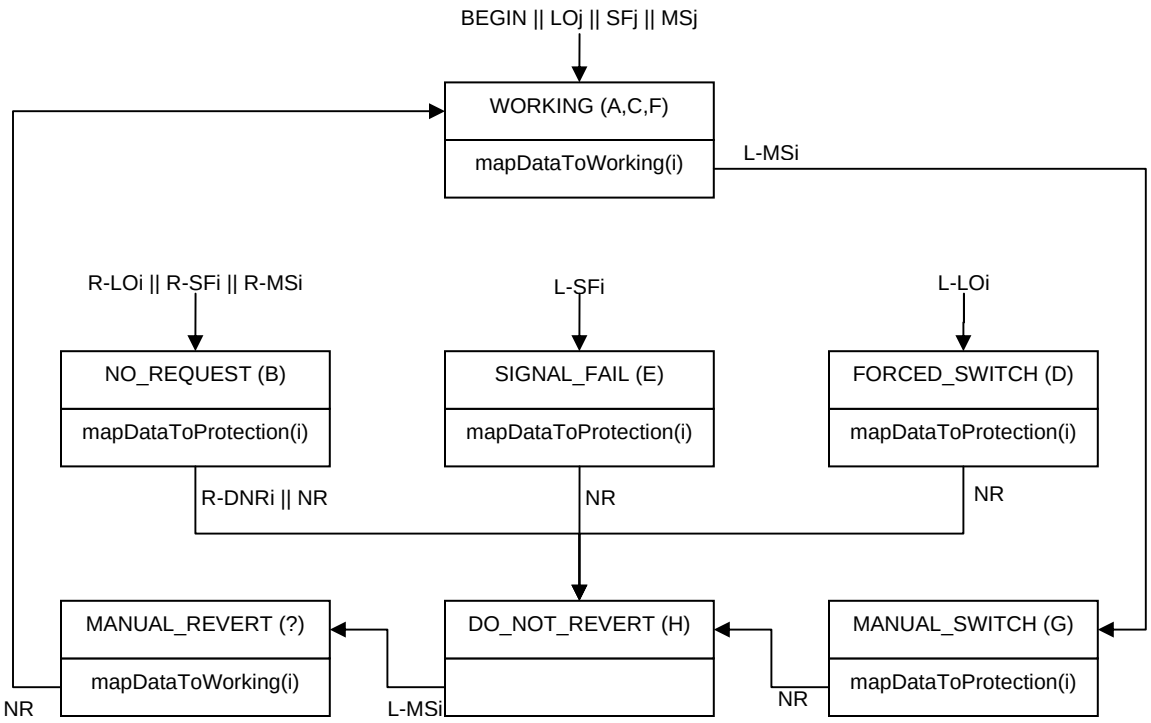

Figure 26-14 — 1:1 load sharing protection member state machine (non-revertive)