Partial Fault Notification and Shared Protection within PBB-TE

Jan Bialkowski, Infinera, jbialkowski@infinera.com

Linda Dunbar, Huawei, ldunbar@huawei.com

Raj Jain, Washington University in Saint Louis, jain@wustl.edu

Alan McGuire, BT, alan.mcguire@bt.com

John Messenger, ADVA Optical, jmessenger@advaoptical.com

Stephen Oliva, Sprint, steve.oliva@sprint.com

Hayim Porat, Ethos-Networks, hayim.porat@ethos-networks.com

Himanshu Shah, Ciena Corporation, hshah@ciena.com

Muneyoshi Suzuki, NTT, suzuki.muneyoshi@lab.ntt.co.jp

Overview

❑ Shared Protection: TDM vs. Ethernet

❑ Partial Faults

❑ Managing Partial Faults

❑ Goals for PBB-TE

Shared Protection: TDM vs. Ethernet

❑ Protection in TDM Networks:

   ➢ Protect against a single failure (either an intermediate node
      or a link),
      but does not protect against double failures

   ➢ Link protection and end-to-end path protection

   ➢ Link is either up or down. (No notion of 90% up)

❑ Protection in Ethernet Networks:

   ➢ Not limited to single failure protection.
      Link aggregation can easily protect against multiple link
      failures.

   ➢ Protection path bandwidth can be less than the primary path
      bandwidth.
      Subset of service instances can be switched over to
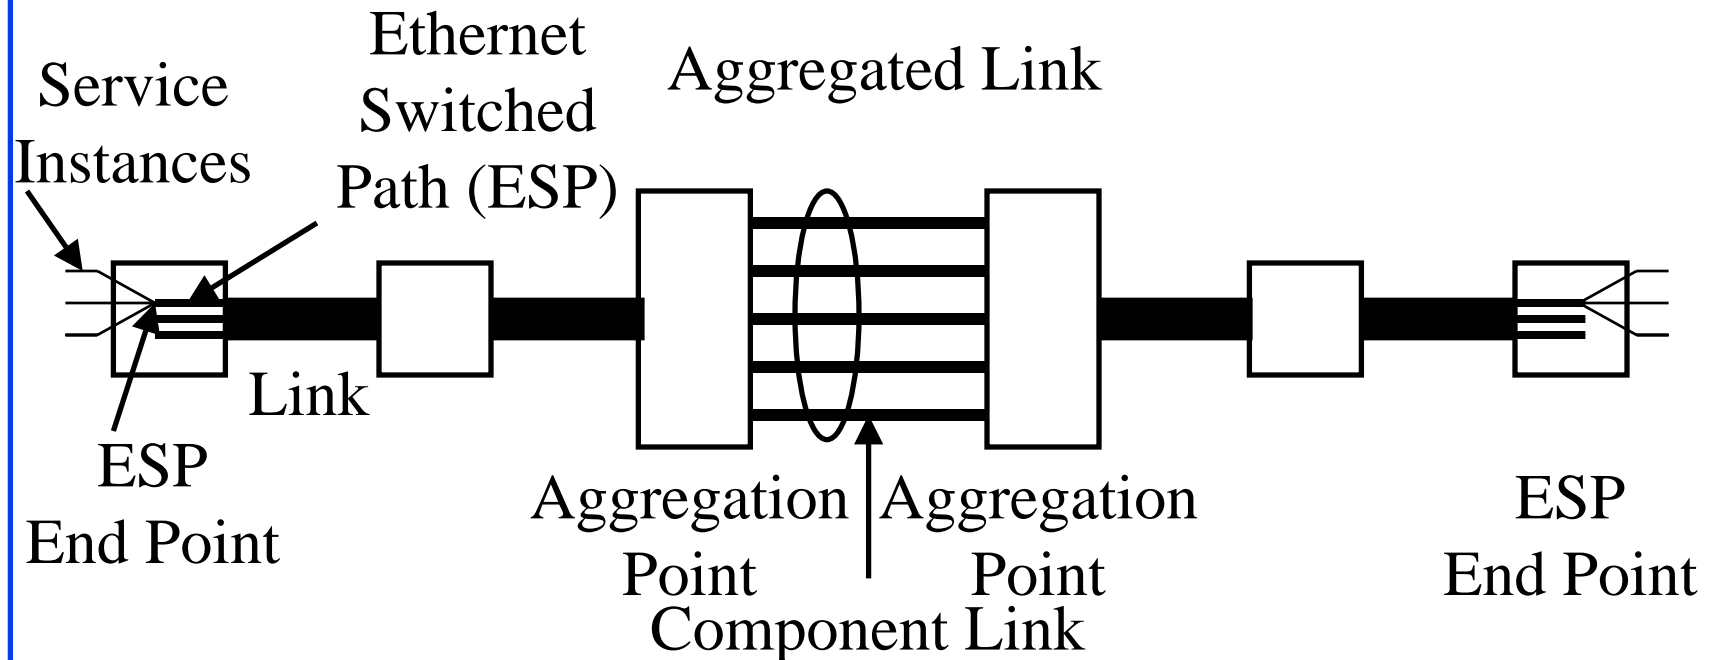   protection path.

Partial Faults
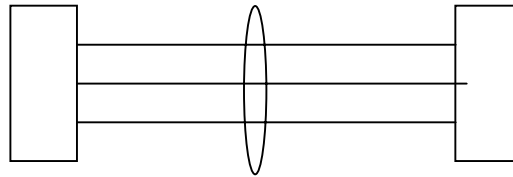
802.3ad Link Aggregation

❑ Physical link bundles are common in packet networks

➢ Several GE links form a large logical link based on 802.3ad.
For example, five 1GE form a 5G link

➢ High speed Ethernet links (>10G) will generally use link aggregation.

❑ Partial faults can exist in a packet transport network

➢ When one of physical links fails in a link bundle, the logical link still
exists and is able to transport some packets but not at its maximum
throughput

➢ Carriers want to continue using a partially failed link for packet
transportation and minimize the service impact

➢ Particularly important for high-speed (high-cost) links

➢ For example, partial failure in 3ad link is notified to OSPF-TE, which
in turn will advertise reduction in available bandwidth.

Terminology

Service
Instances

Ethernet
Switched
Path (ESP)

Aggregated Link

Link

ESP
End Point

Aggregation
Point

Aggregation
Point

Component Link

ESP
End Point

- ❑ Each Ethernet Switch Path (ESP) carries multiple service instances
- ❑ Each link carries multiple ESPs
- ❑ Each ESP can go through only one component link.
  Each component link can have multiple ESPs going though it.

Composite Link

❑ ITU SG15 Q12 has recently introduced a concept of "Composite Link": Multiple parallel links between the same subnetworks bundled together into a single composite link.

➢ Each component link is independent. A composite link may not guarantee of the symbol sequence if symbols are transported through different component links.
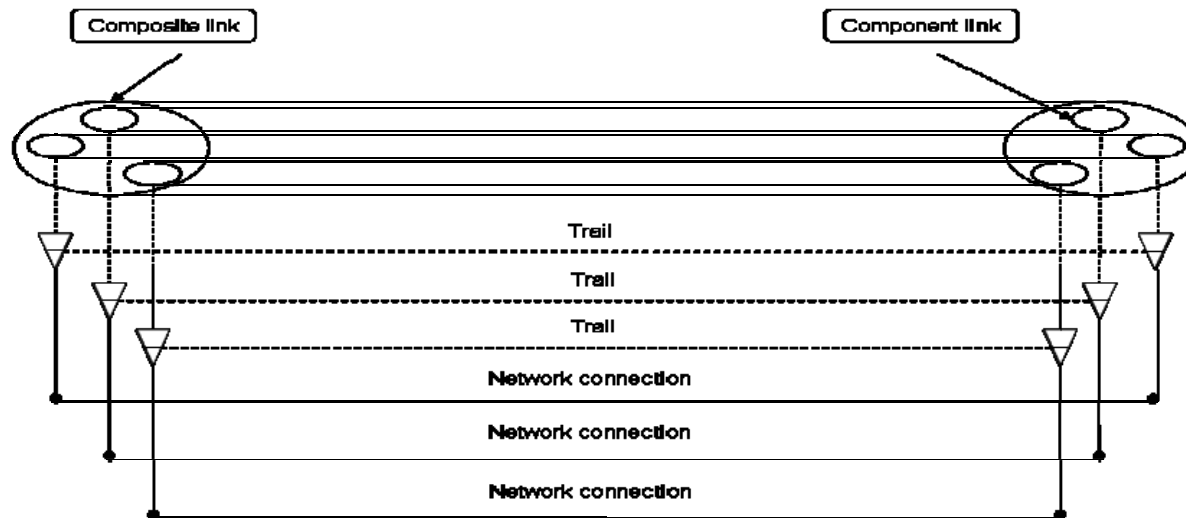
Possible Approaches

A. Not to use link aggregation or composite links but to let the PBB-TE network see each link as separate and to "load share" by TE. Some operators may do this.

B. Recognize that link aggregation or composite links are important and handle the problem of partial faults.

Managing Partial Faults

❑ Managing partial faults consists of **detecting** partial faults and **responding** to partial faults

❑ When aggregated link capacity is reduced to less than the total bandwidth of Ethernet Switched Paths (ESPs) going through the aggregated link group:

1. Aggregation Point Processing:

   ❑ Either shutting down some ESPs traversing through this aggregated group which will trigger the selected ESPs to detect connectivity loss immediately to trigger needed protection, or

   ❑ Notifying all the ESPs traversing through the aggregated group about reduced capacity. Letting ESP end points to process needed action.

Managing Partial Faults (Cont)

2.   ESP end point processing:

   ❑   If connectivity loss is detected, perform protection action if there is one.

   ❑   If notification of reduced bandwidth is received, the ESP end points can either switch some of the service instances or entire ESP to a different path (protection ESP), and/or shut down some of the service instances in the ESP or the entire ESP.

Goals for PBB-TE

❑ Define a **mechanism for notification** of partial faults to be propagated to ESP end nodes (Source or Destination).

Based on service provider's policies, this will:

1. Allow ESP end nodes to switch a portion of service instances bundled within the affected ESPs to protection path if protection path is available

2. Allow ESP end nodes to shutdown a portion of service instances to guarantee that the remaining services are transported without interference or notify remaining services about reduced bandwidth

   The shutdown option is useful when the selected service instances are segments to other services. By disabling these service instances, CC messages for the service instances, if there are any, can't get through, which will trigger protection at an outer layer.

Benefits of Notification

❑ Disabling selected service instances instead of letting intermediate nodes silently drop packets with lower priority has the following advantages:

➢ Dropping lower priority packets is not predictable. Some CC can still go through even though some data packets might be dropped when partial fault occurs.

➢ For TE traffic, it is important to have predicable behavior. It is likely better for an outer service layer to be aware of a shutdown than to have unpredictable behavior.

➢ Dropping packets may work for voice and video. For data traffic, all dropped packets have to be retransmitted.

Summary

❑ PBB-TE should provide **mechanisms for indicating** reduced capacity to originating end-points, which allows them to redistribute or selectively shutdown the service instances.