

Fast ReRoute for Traffic Engineered Ethernet

Bob Sultan

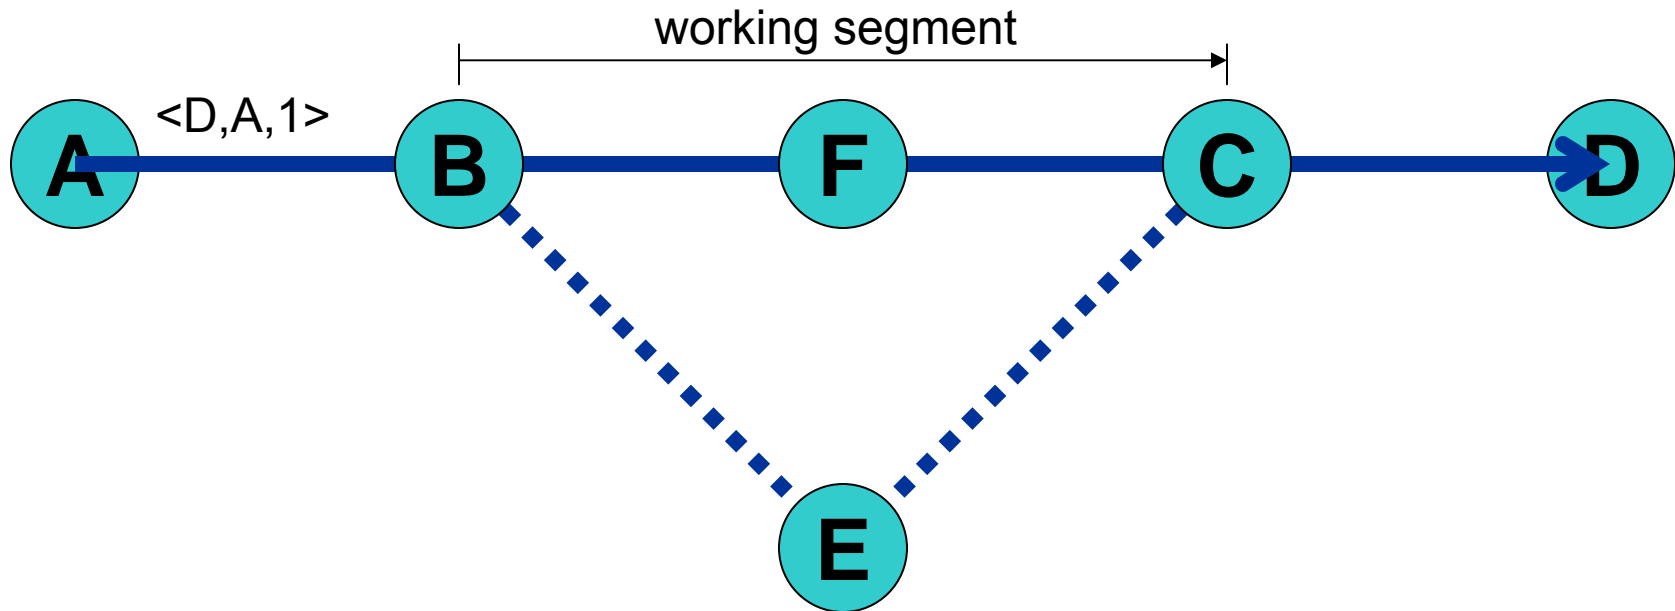
John Lemon

Deng Zhusheng

Abhay Karandikar


M Vinod Kumar

Working Segment




- A 'working segment' is a portion of an ESP protected by the existence of an alternate path between its endpoints.
- In the figure above, B-F-C is a working segment of the ESP <D,A,1>

Protection Segment



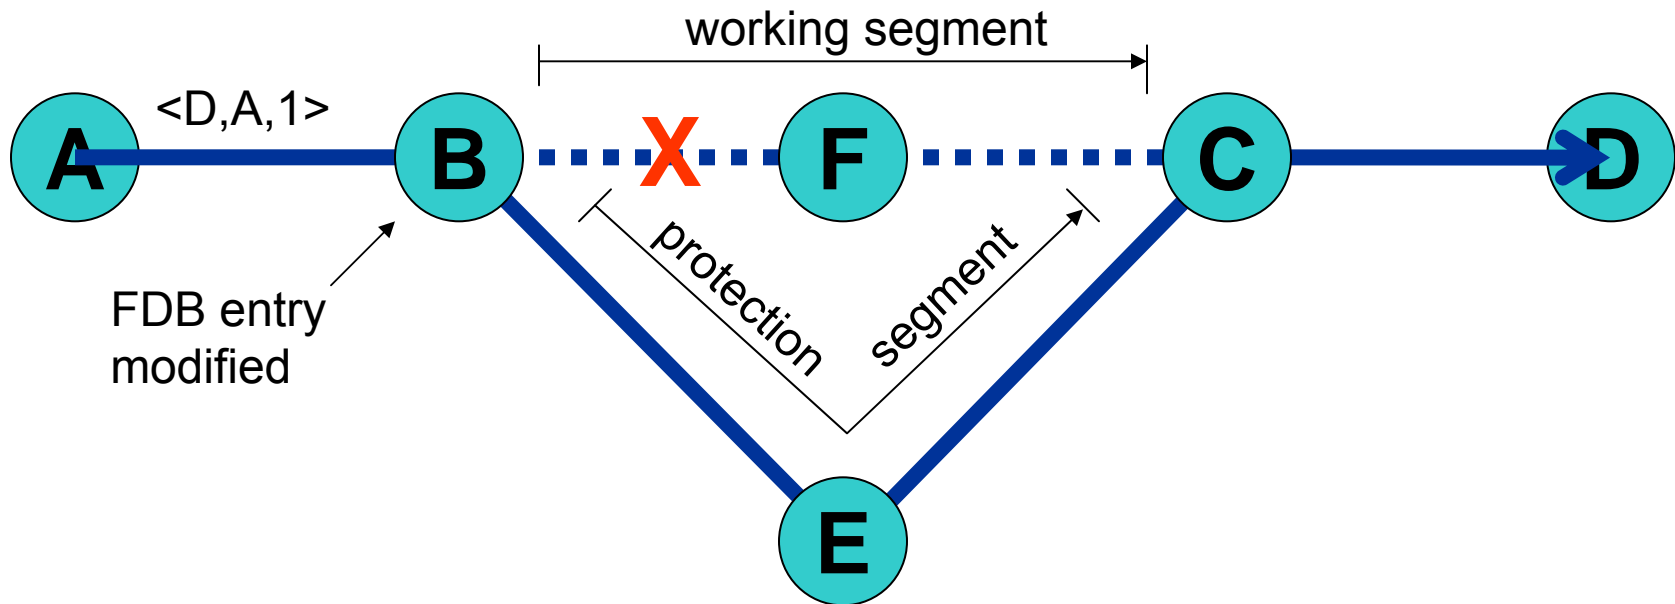
- The alternate path associated with a working segment is called a 'protection segment'
- Traffic traverses a working segment or the associated protection segment, but not both.
- Traffic is carried on the *working segment* in the absence of faults on the working segment and on the *protection segment* in the presence of faults on the working segment.

Segment




- **As a notational convenience, we will use the term ‘segment’ to refer to a working segment and its associated protection segment.**

Re-Route




- A fault occurring on the working segment B-F-C of ESP $\langle D,A,1 \rangle$ results in modification of the FDB at segment ingress B to direct traffic on the protection segment B-E-C
- The FDB is restored when the fault is cleared and traffic is again directed on the working segment

Multiple segments per ESP



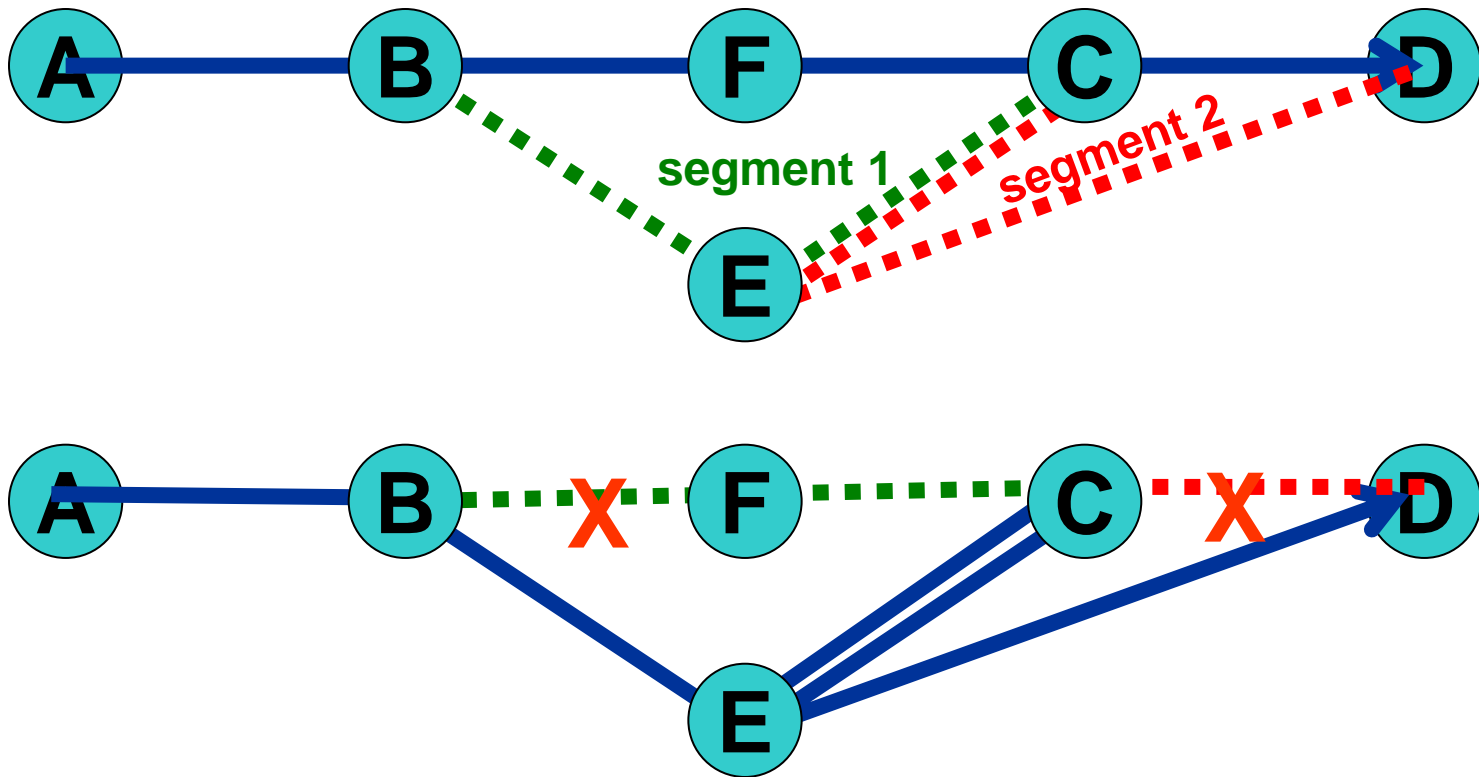
- **Multiple segments may be defined along the path of an ESP**
- **Segments may overlap (as above)**

Non-overlapping segments




- **Segments may be non-overlapping**
- **Non-overlapping, but contiguous segments are shown above**

'No backtrack' rule




- Failures of contiguous segments having a link common to the protection segments can result in the 'backtracking' of traffic along the path of the ESP.
- Backtracking should be prohibited.

Avoiding backtracking



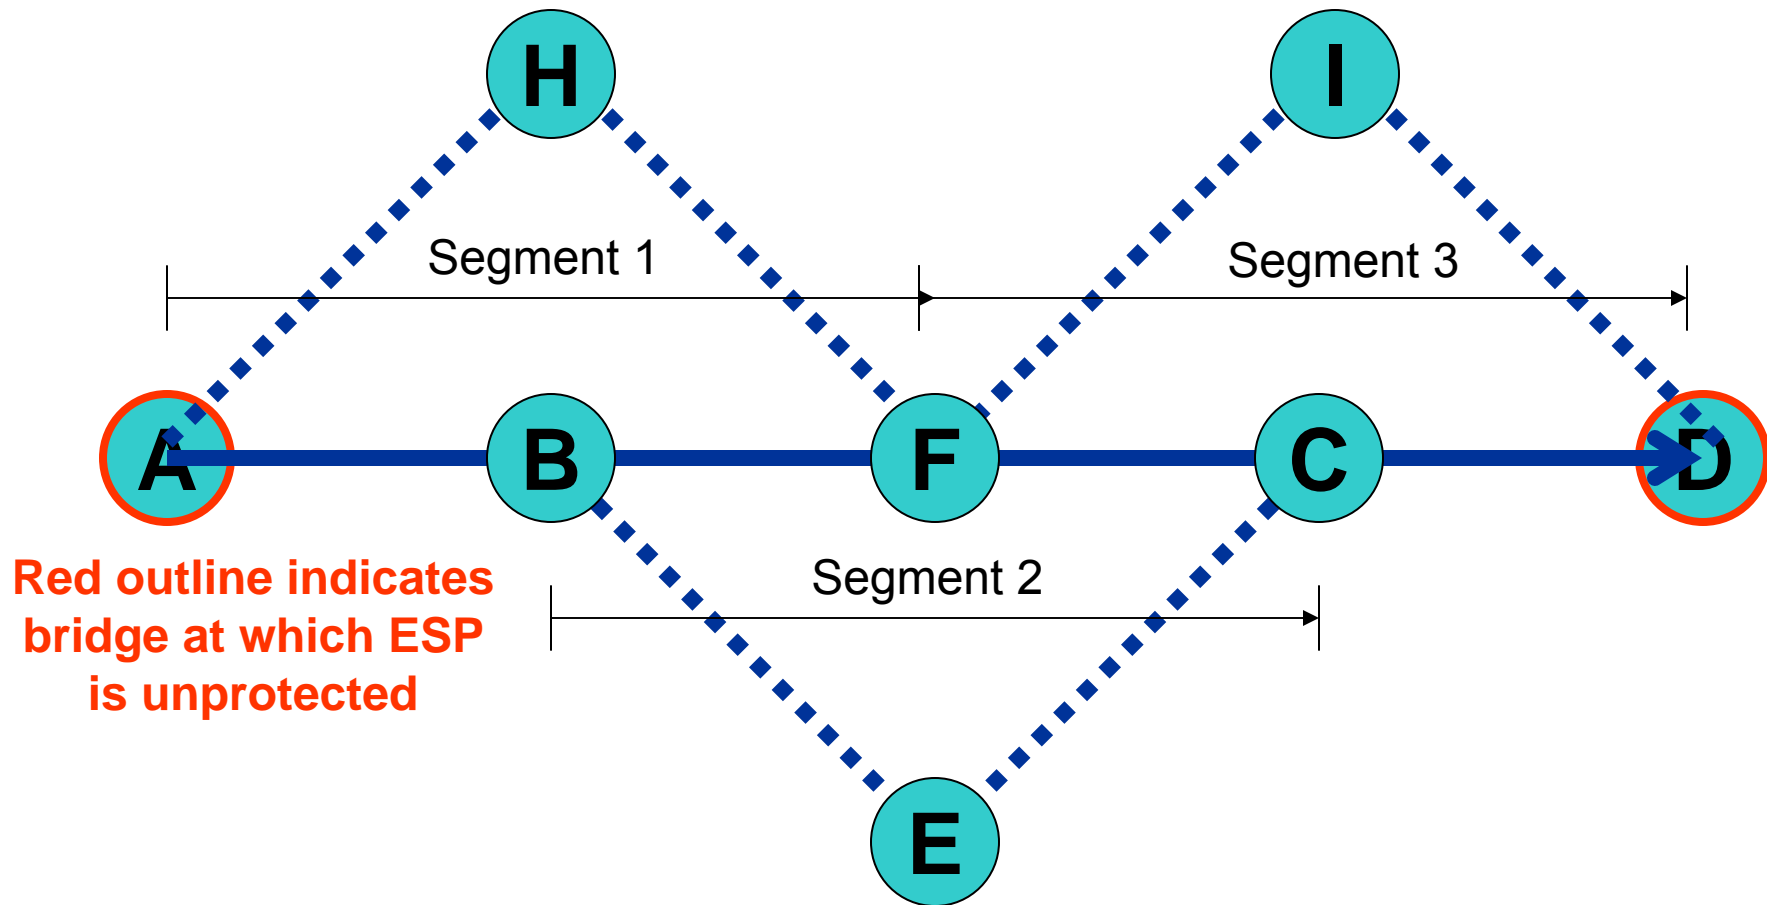
- The figure above shows the desired result, that the path of ESP traffic should bypass the portion of the path in which backtracking would occur.
- The operator can also avoid backtracking by defining a working segment A-F-C-D and protection segment B-E-D

Non-contiguous segments



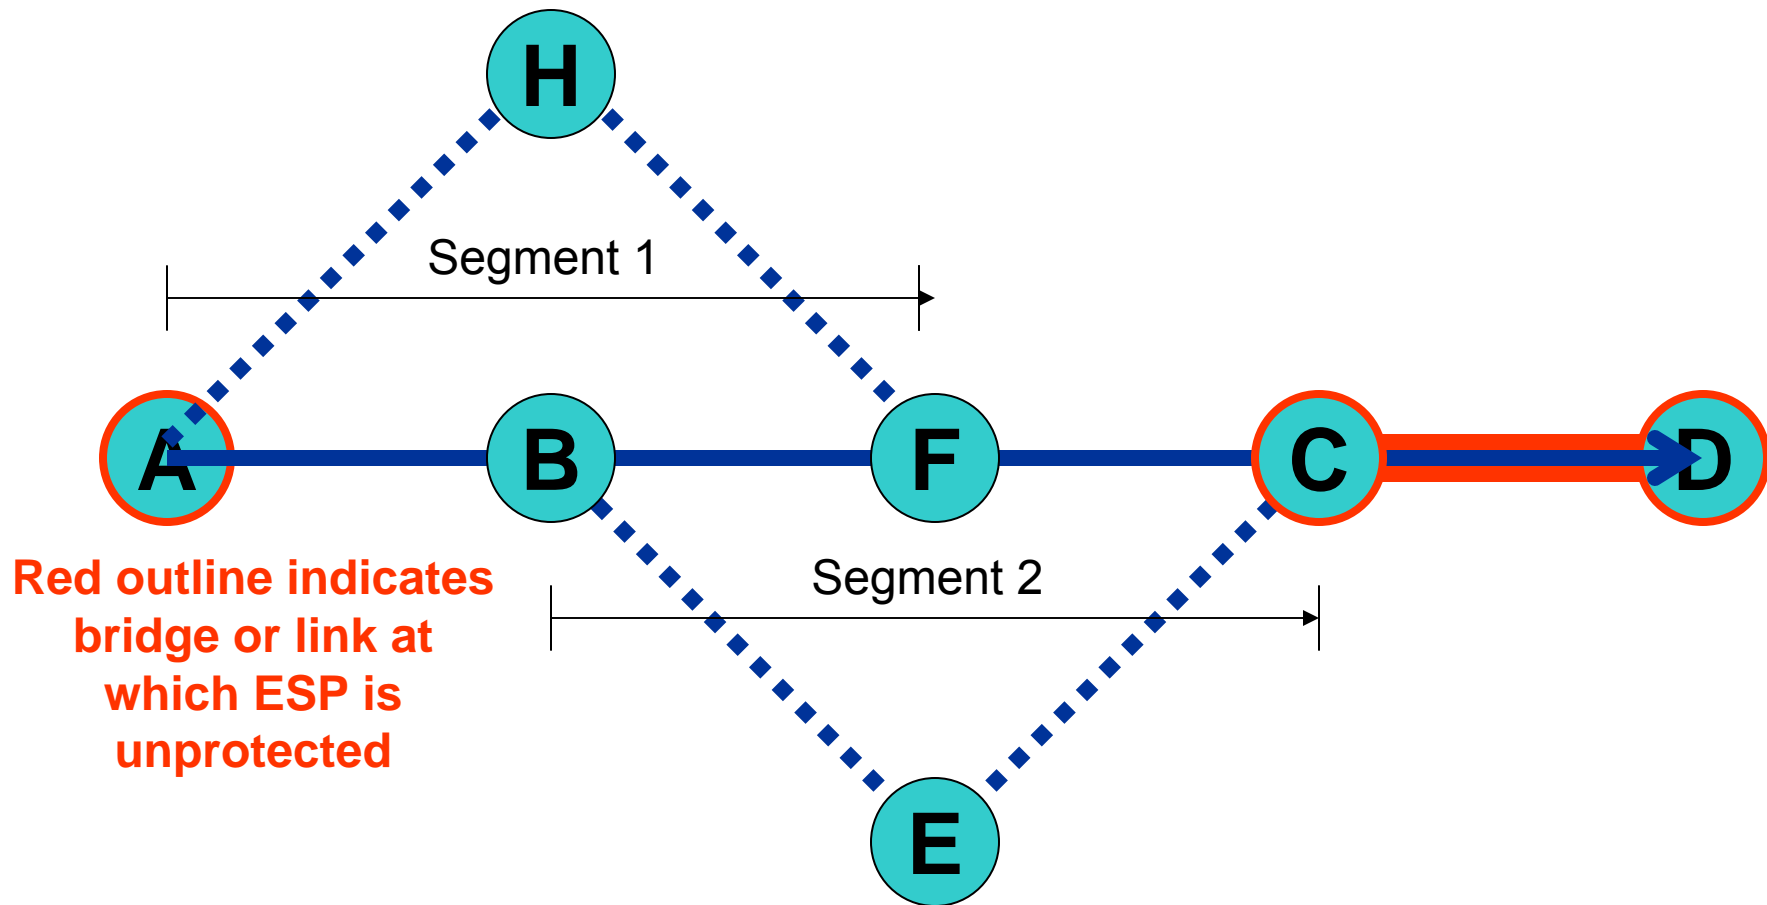
- **Segments may be non-contiguous**

Protection Scope can be end-to-end



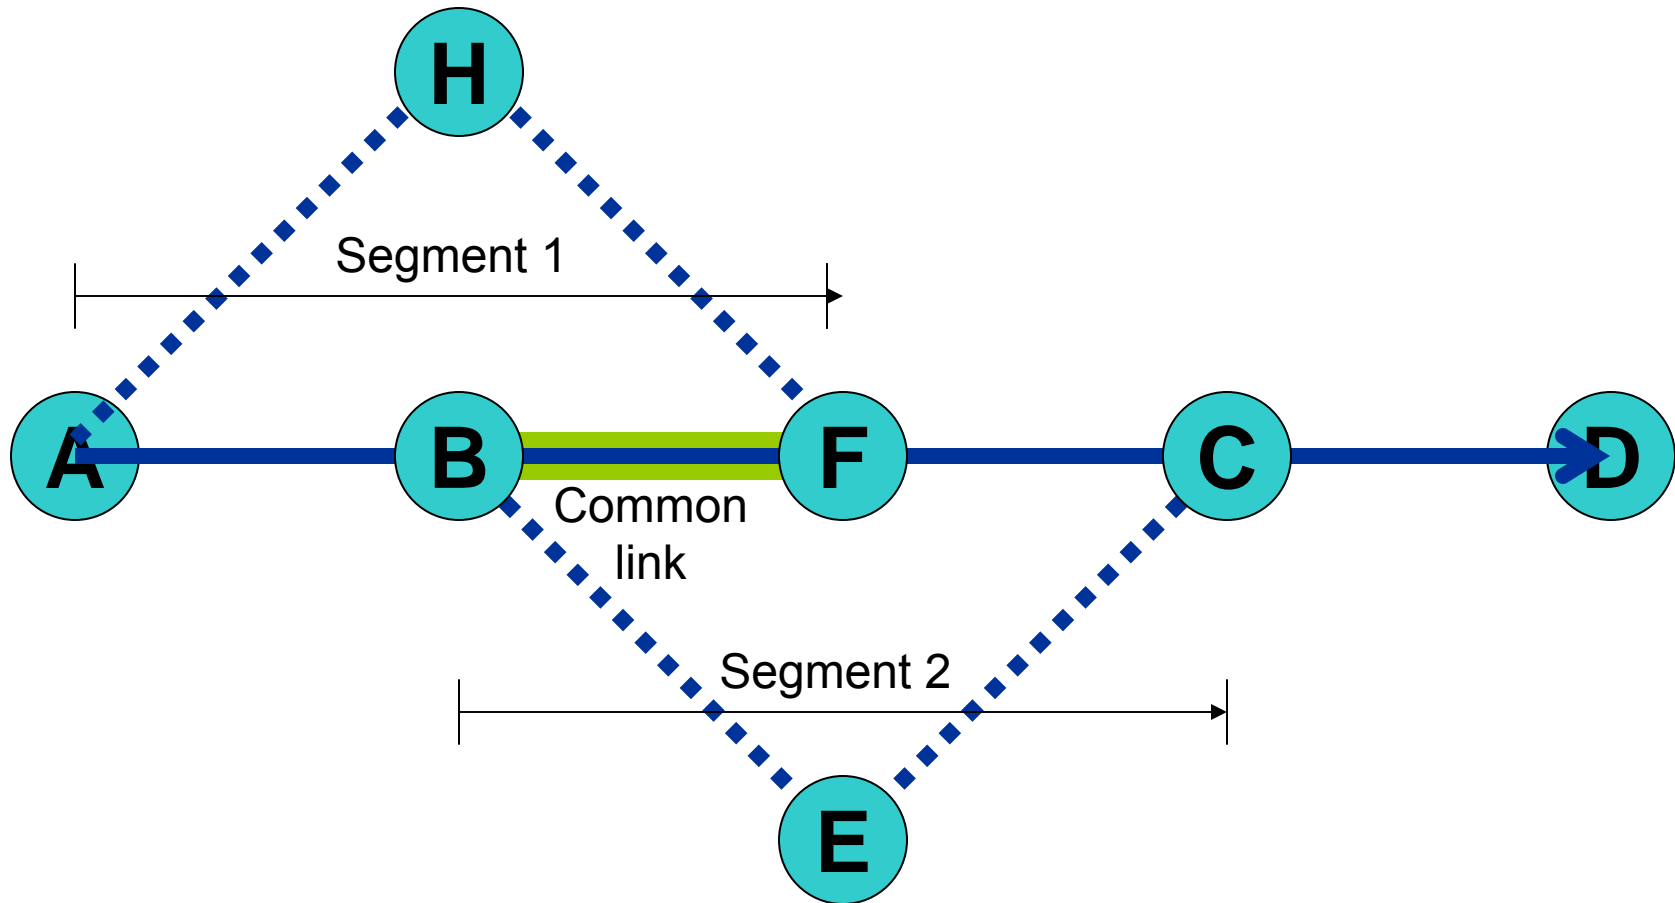
- **Segments can be defined to protect *all* links and bridges along the path of the ESP (excluding endpoint bridges)**

Scope may exclude some links or bridges



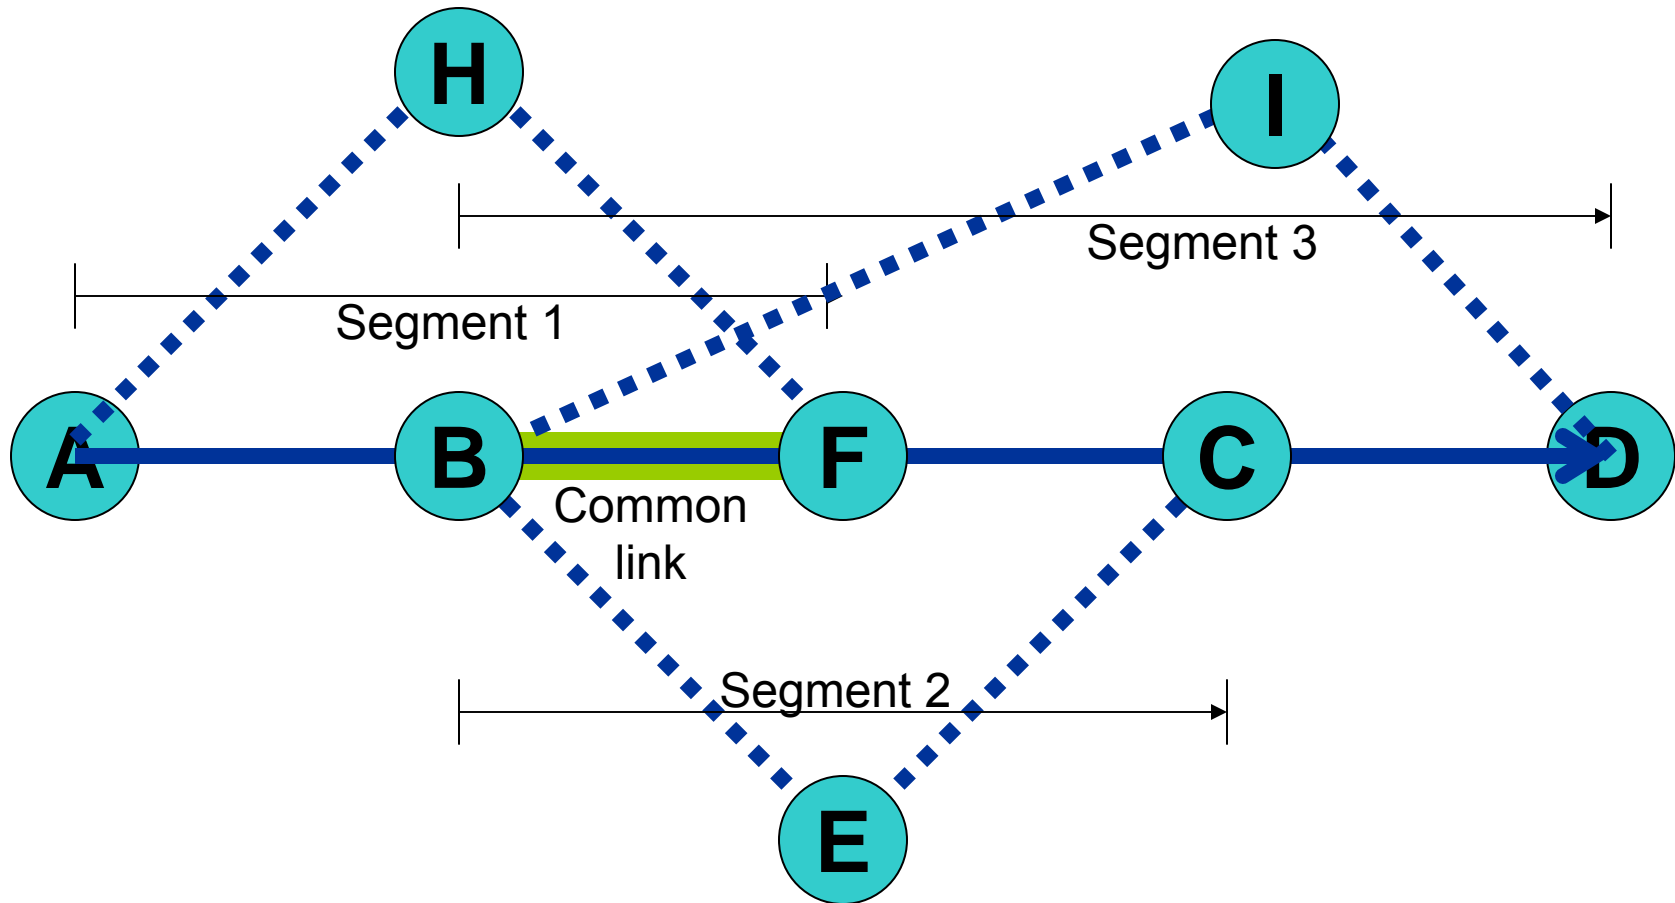
- **Segments can be defined such that some links and/or bridges in the interior of the ESP are not protected.**

Common link



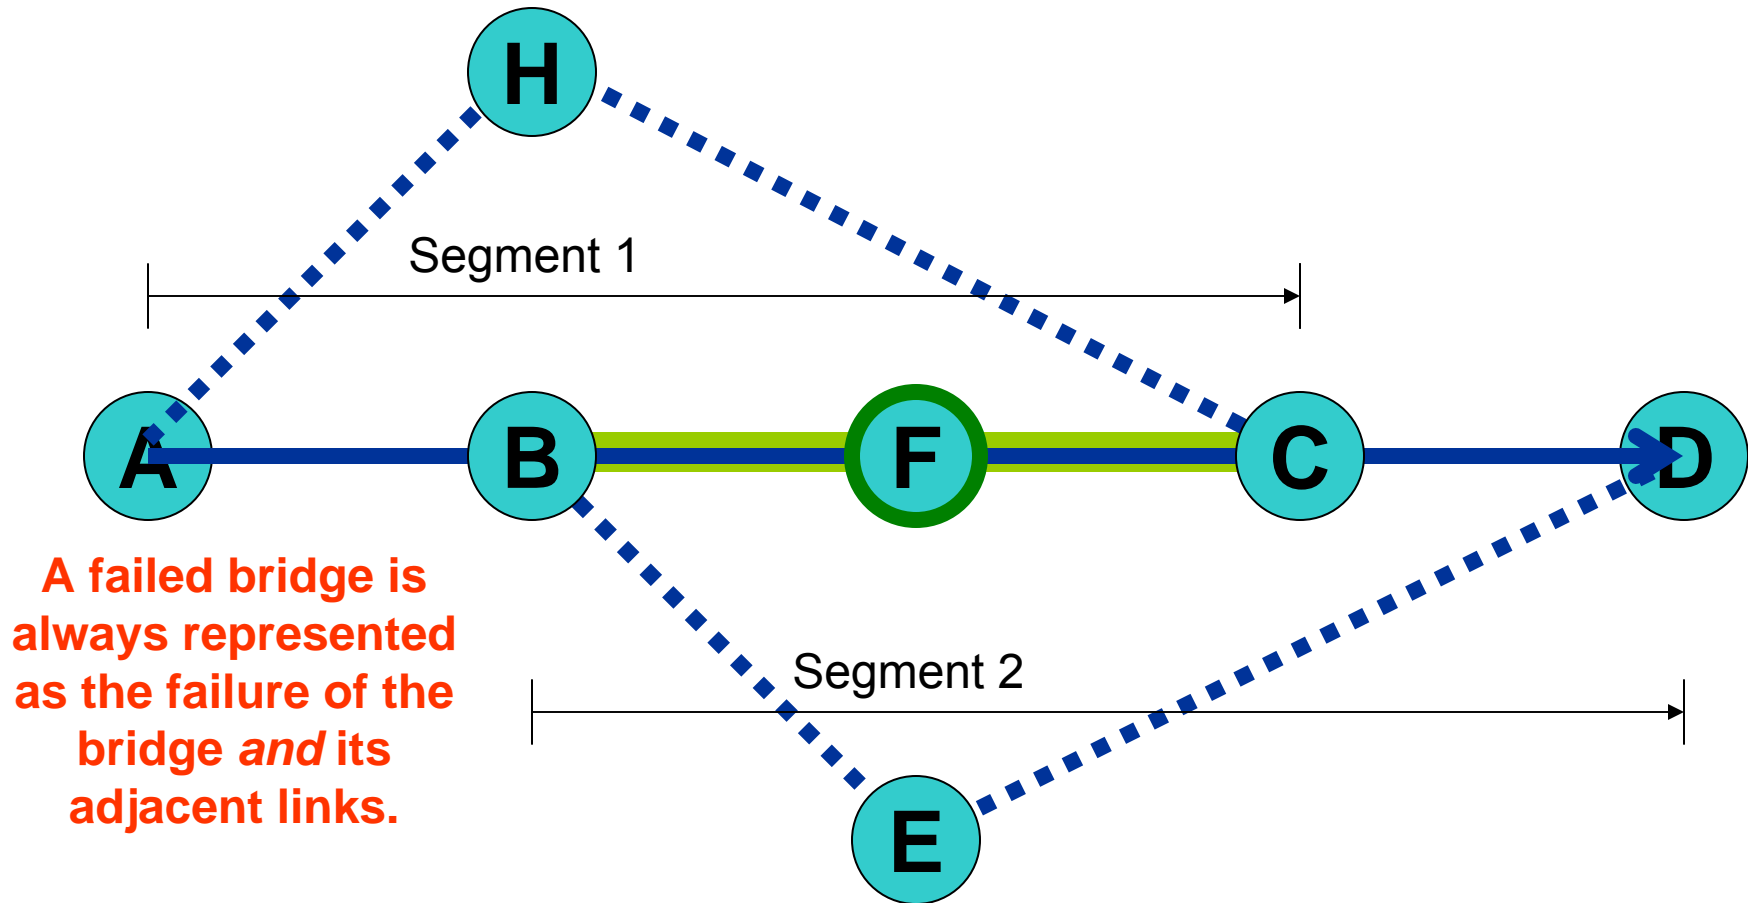
- **Overlapping segments will have at least one common link.**

Multiple segments sharing a link



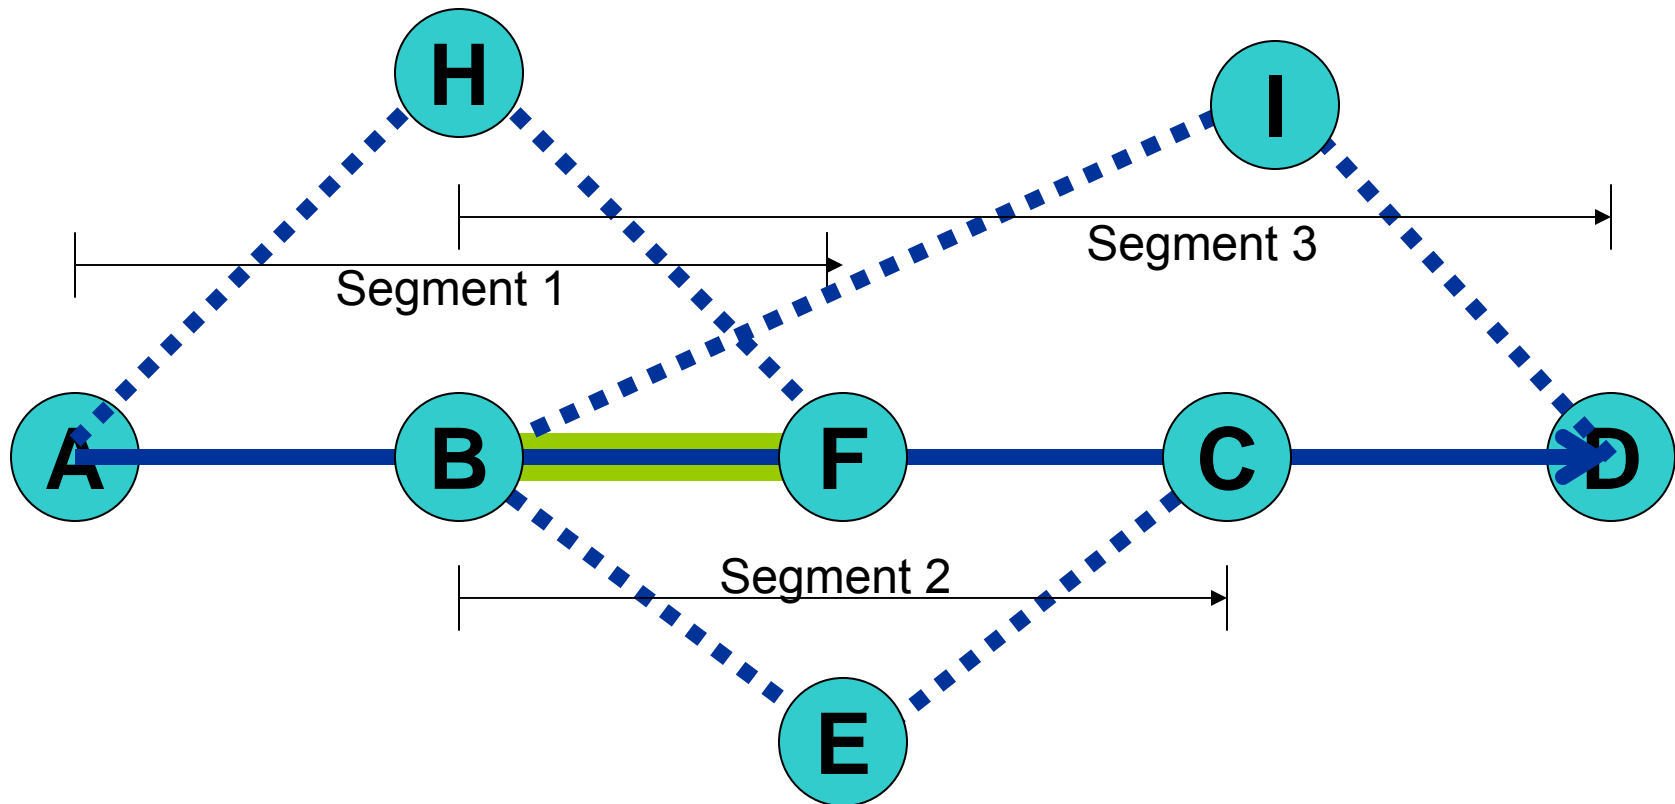
- **Multiple segments can share a link**

Common bridge




- **Overlapping segments can share a bridge**
- **Like a shared link, a bridge can be shared by an arbitrary number of segments.**

Only one segment takes protection action

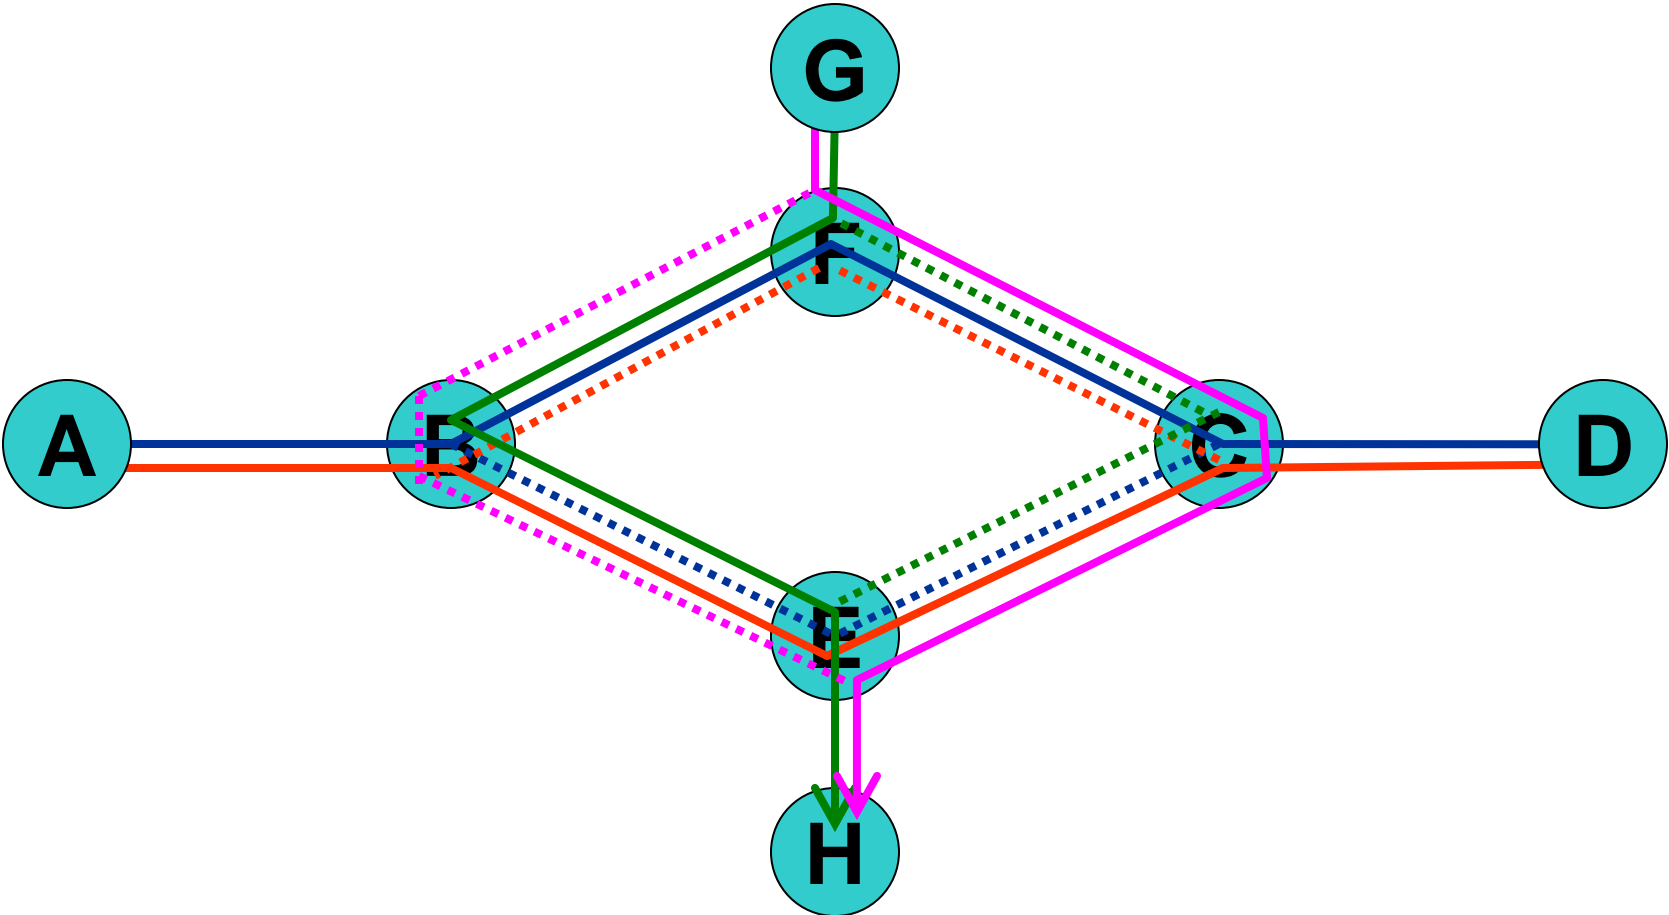






- **Only one of the segments sharing a link is permitted to take a protection action on the failure of that link.**
- **Any methods of supporting FRR-TE should specify how this rule is enforced.**

Let's redraw the picture like this...




Draw four ESPs following different paths



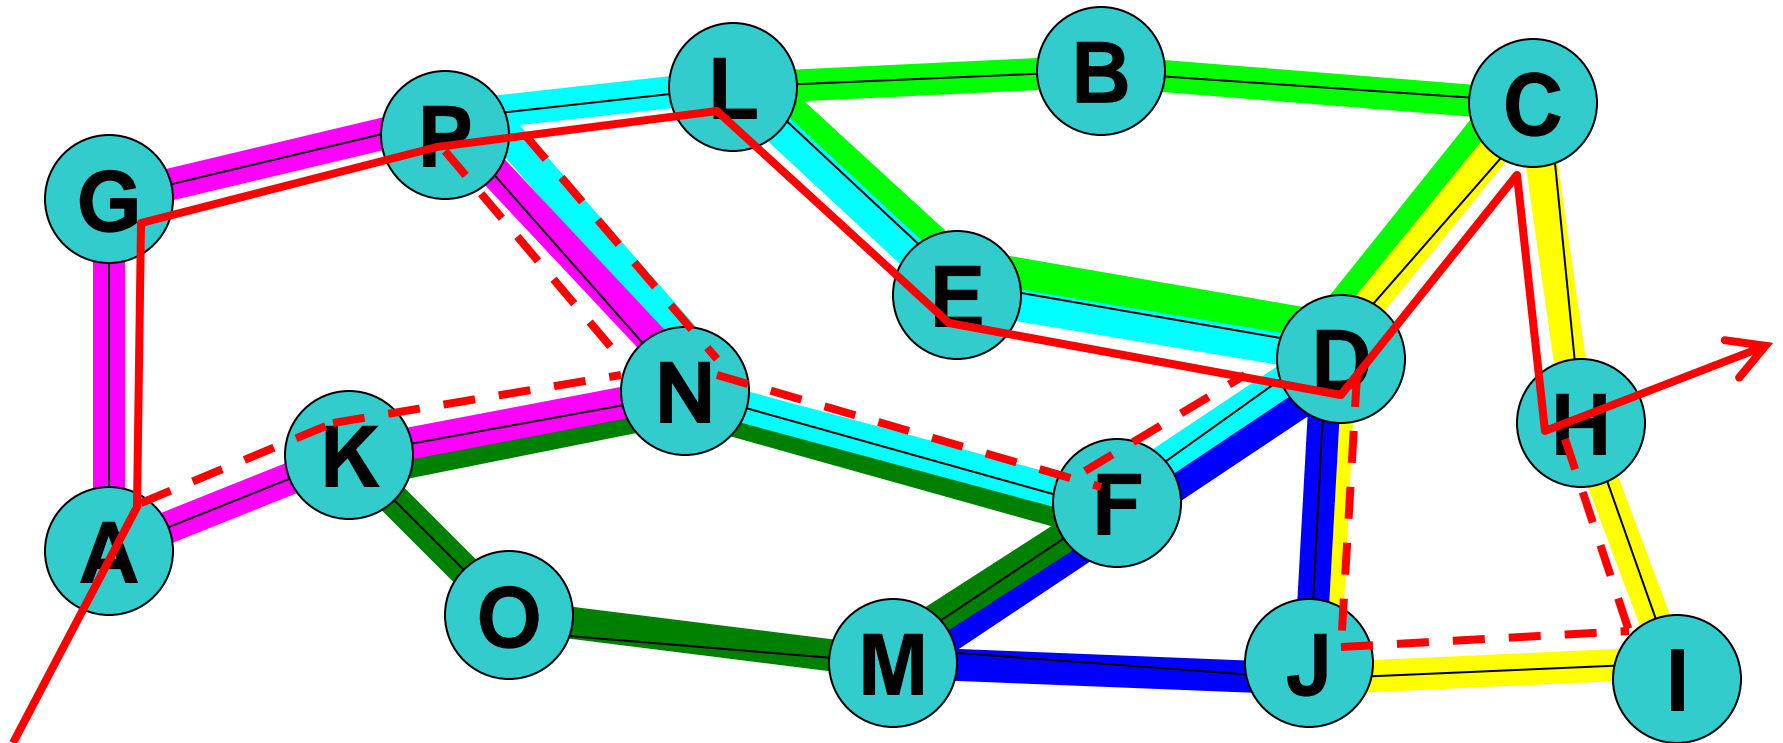
	ESP	path	W/SEG	P/SEG
	<A,D,1>	A-B-E-C-D	B-E-C	B-F-C
	<A,D,2>	A-B-F-C-D	B-F-C	B-E-C
	<G,H,1>	G-F-B-E-H	F-B-E	F-C-E
	<G,H,2>	G-F-C-E-H	F-C-E	F-B-E

Observation



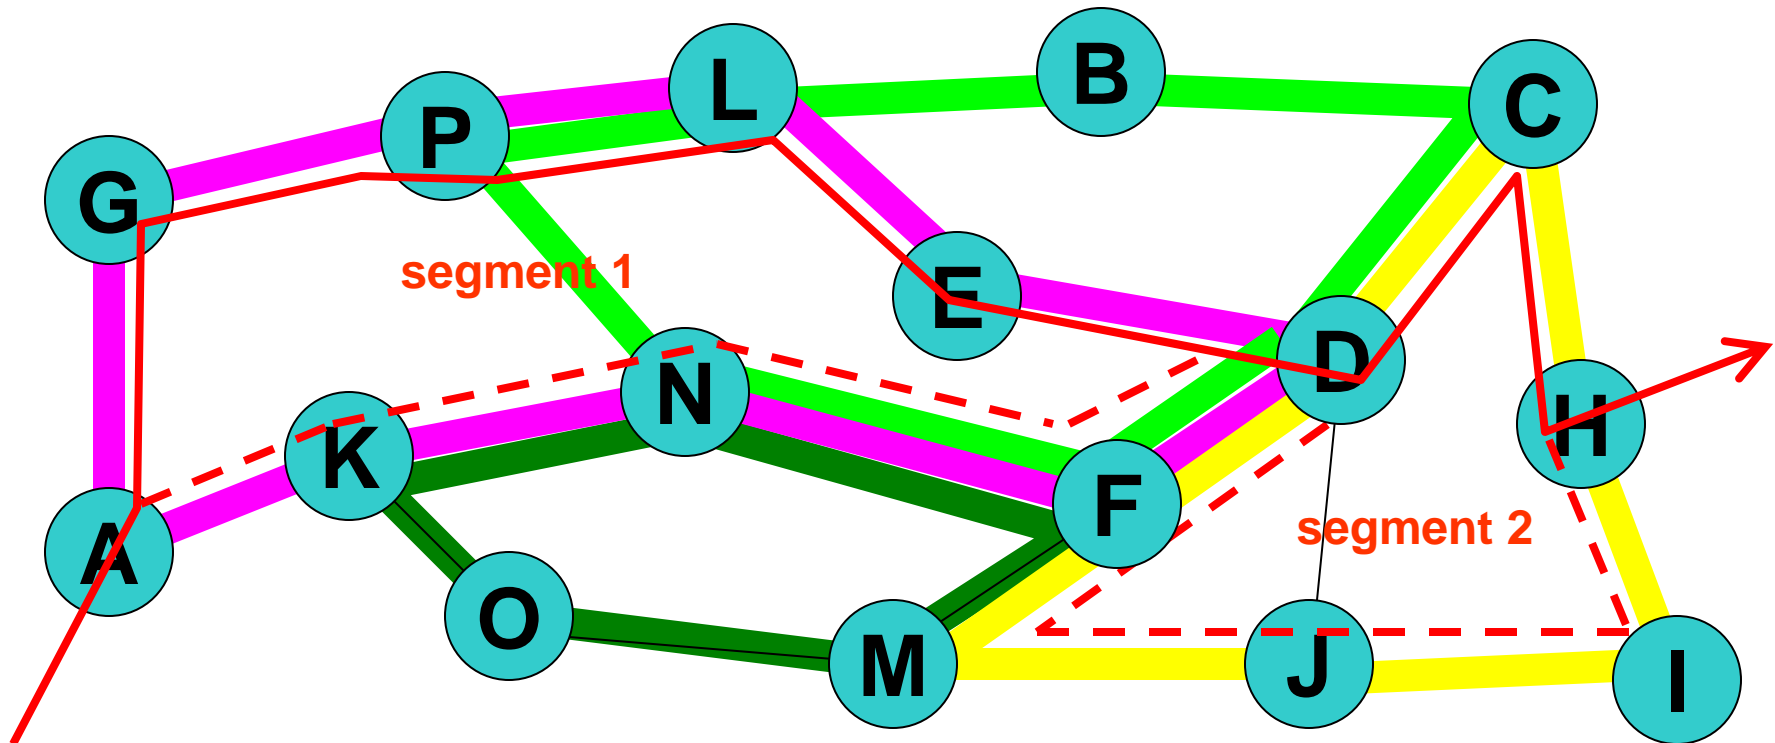
- Many ESPs may map their working and protection segments to the same *closed* physical path
- We call such a path a *cycle*

Mapping segments to cycles



- The network designer or operator identifies cycles in the network topology in order to create the segments associated with an ESP.


Cycles may overlap



In the above figure, red ESP segment 1 is mapped to the violet cycle;
red ESP segment 2 is mapped to the yellow cycle.

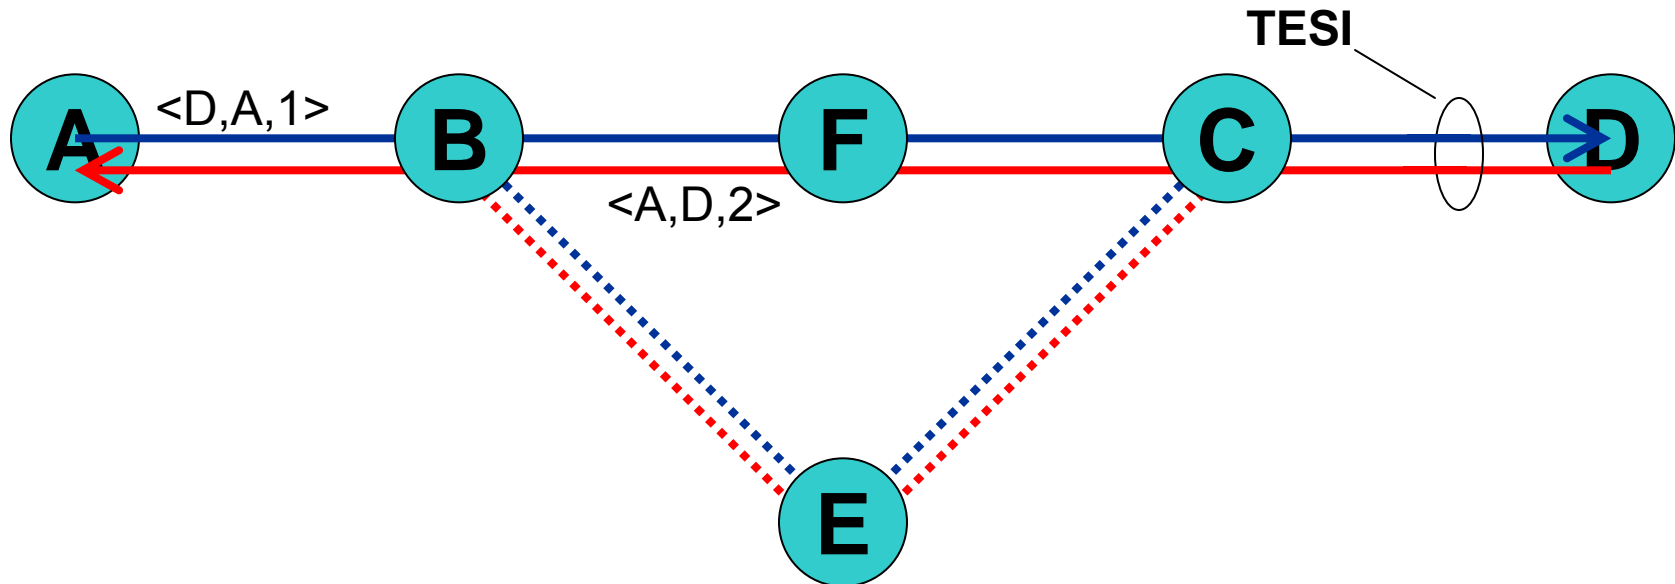
- **Cycles may overlap (violet and green cycles)**
- **Not every closed path need be identified as a cycle (only those to which segments are mapped)**

Another Observation



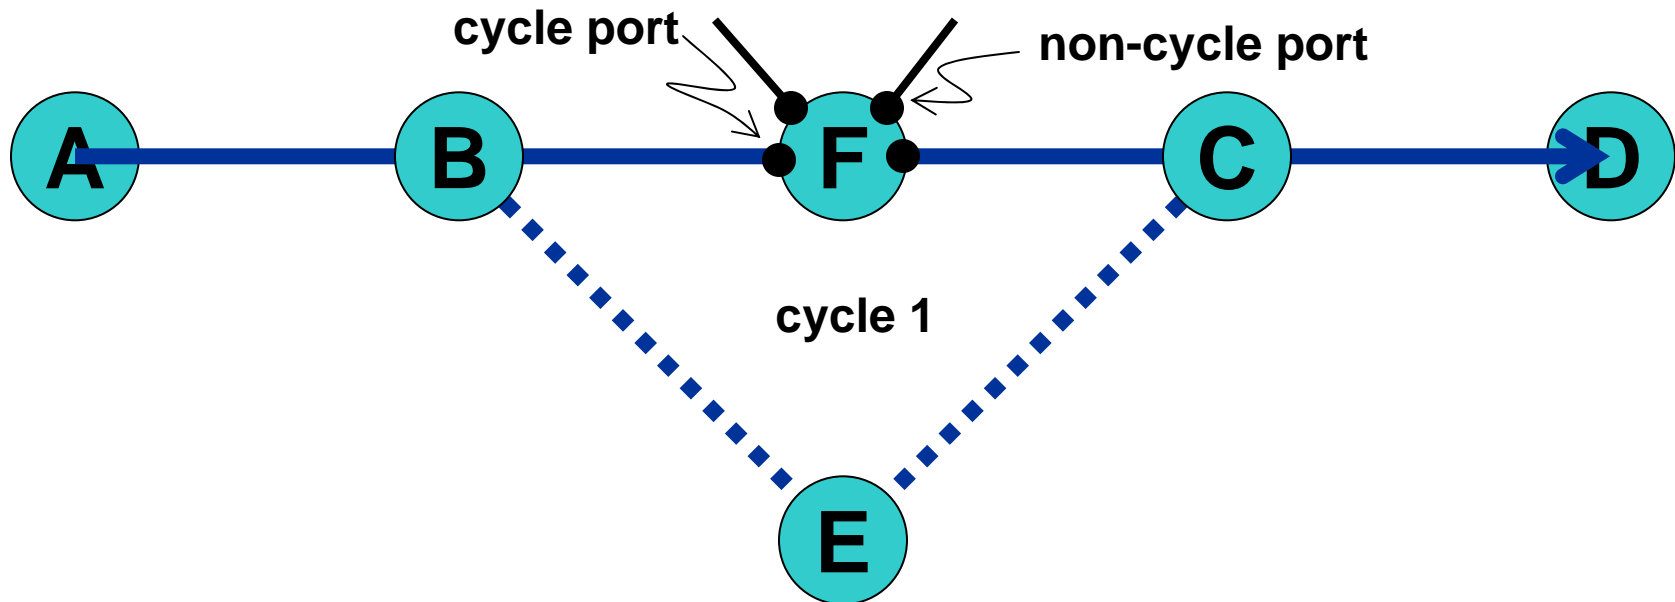
- **Methods of deploying FRR-TE may require the distribution of information to all bridges within a segment.**
- **In this case, it is efficient to distribute the information one time for the cycle rather than once per segment associated with that cycle**

One more thing....



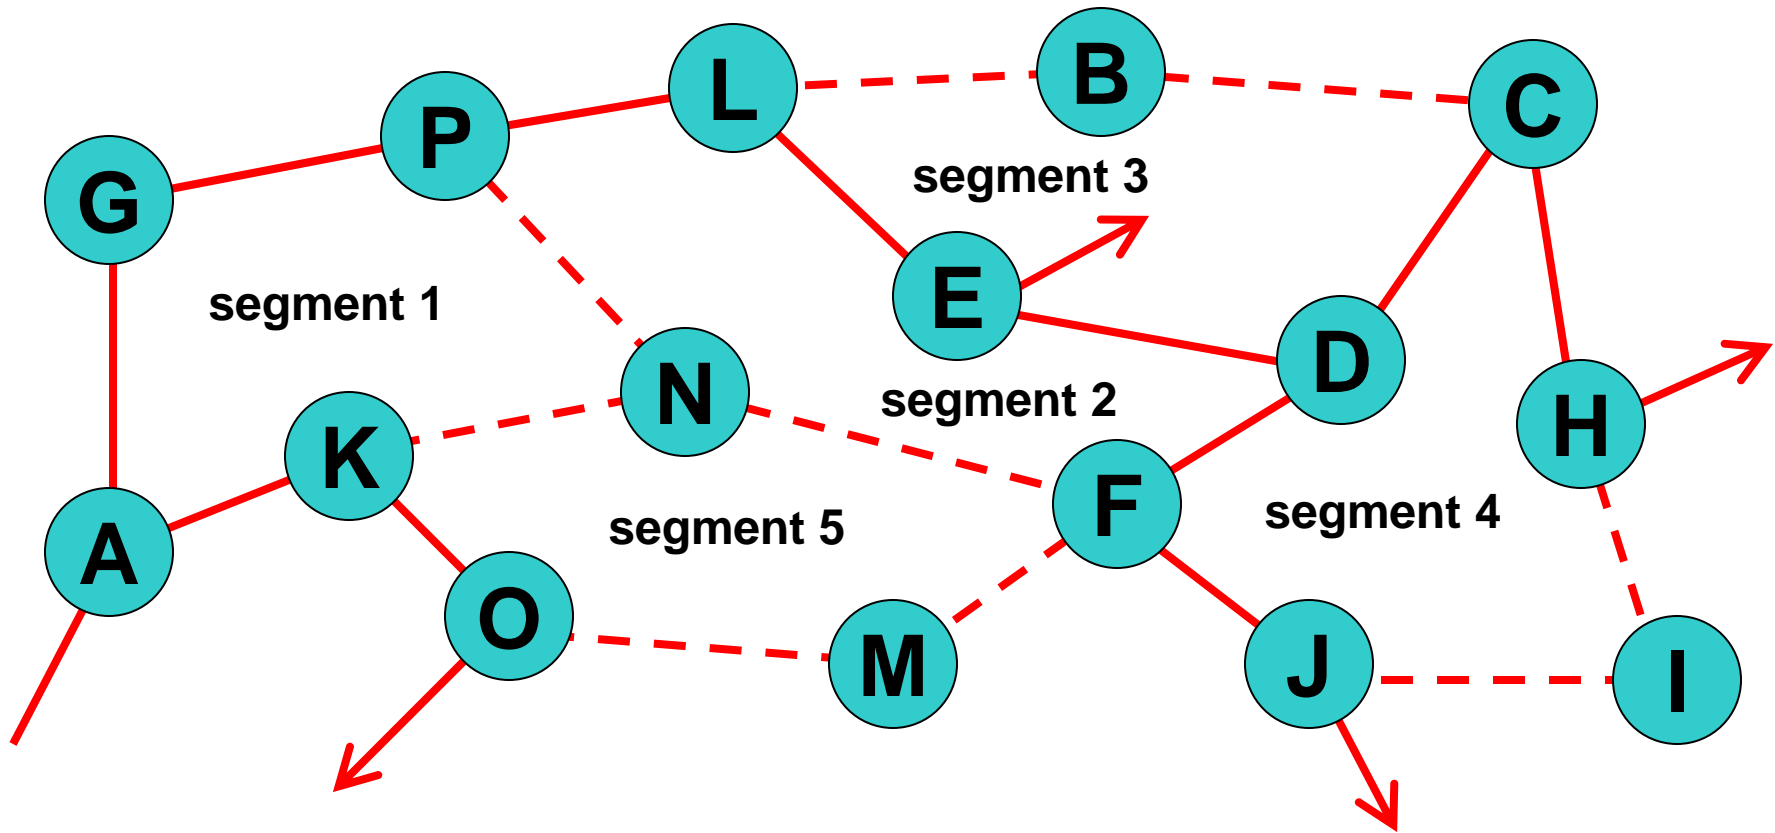
- ***Everything* we have said so far about a PtP ESP applies to a PtP TESI comprised of a pair of co-routed counter-directional PtP ESPs**
- **The actions described are performed for both ESPs in the pair**
- **We don't show this explicitly because it results in very complex pictures**

Efficient Provisioning



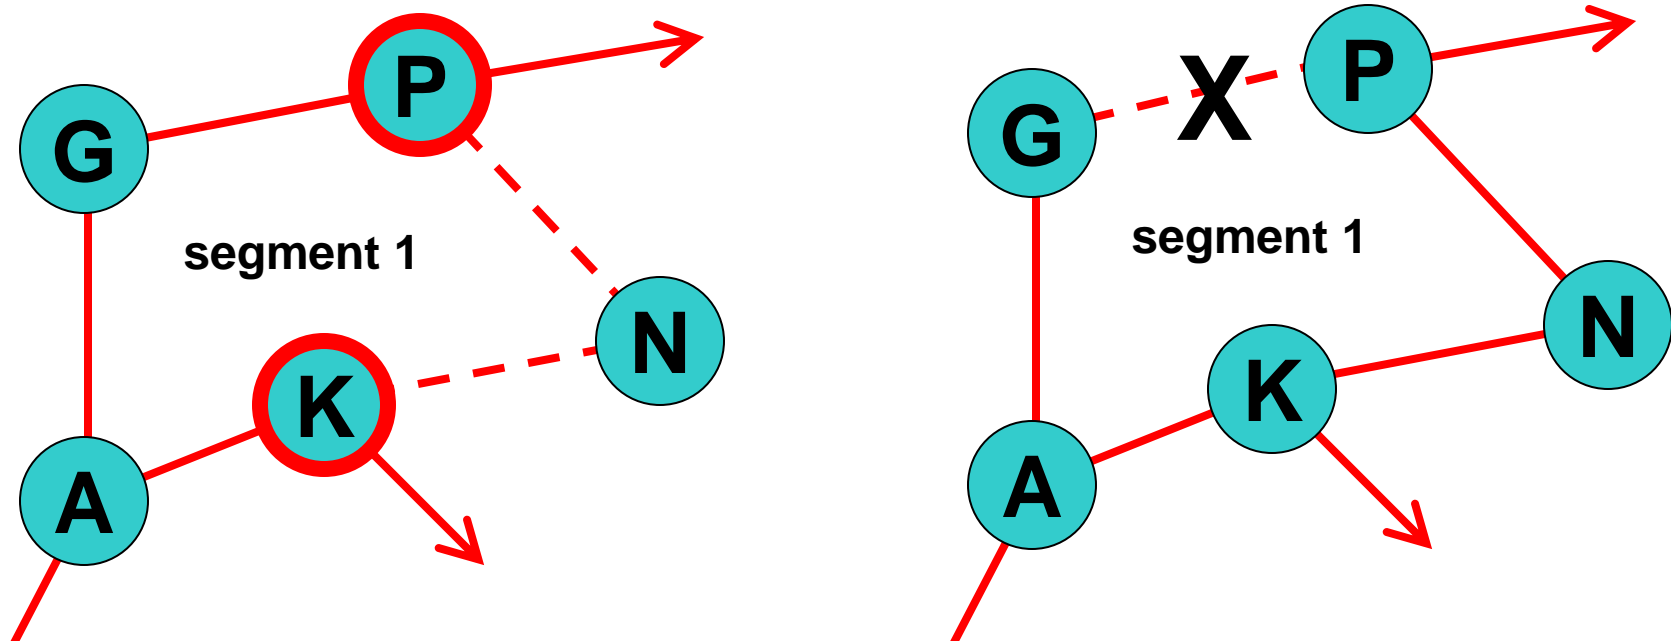
- A bridge that is associated with only one cycle (ie., not shared among cycles) need not have an explicit FDB entry for a segment transiting that bridge
- The FDB of each bridge contains a 'wildcard' entry that matches any TE frame for which a more specific FDB match is not found

FRR-TE with P2MP ESP



- **In the absence of faults on segment 1:**
 - P is provisioned so as not to forward on P-N
 - K is provisioned so as not to forward on K-N

FRR-TE with P2MP ESP



- We can say that P and K are the *stop-points* for segment 1.
- A stop-point receiving failure notification on its working-side no longer stops frames from continuing on the segment.
- This method does send unnecessary traffic on link A-G following protection switch.

Wildcard FDB Entry

The wildcard entry specifies the two 'cycle' ports of the bridge as the 'outbound ports' of the FDB entry.

A frame received on one cycle port is forwarded on the other cycle port. Thus the wildcard entry allows a frame to transit the bridge, following the cycle.

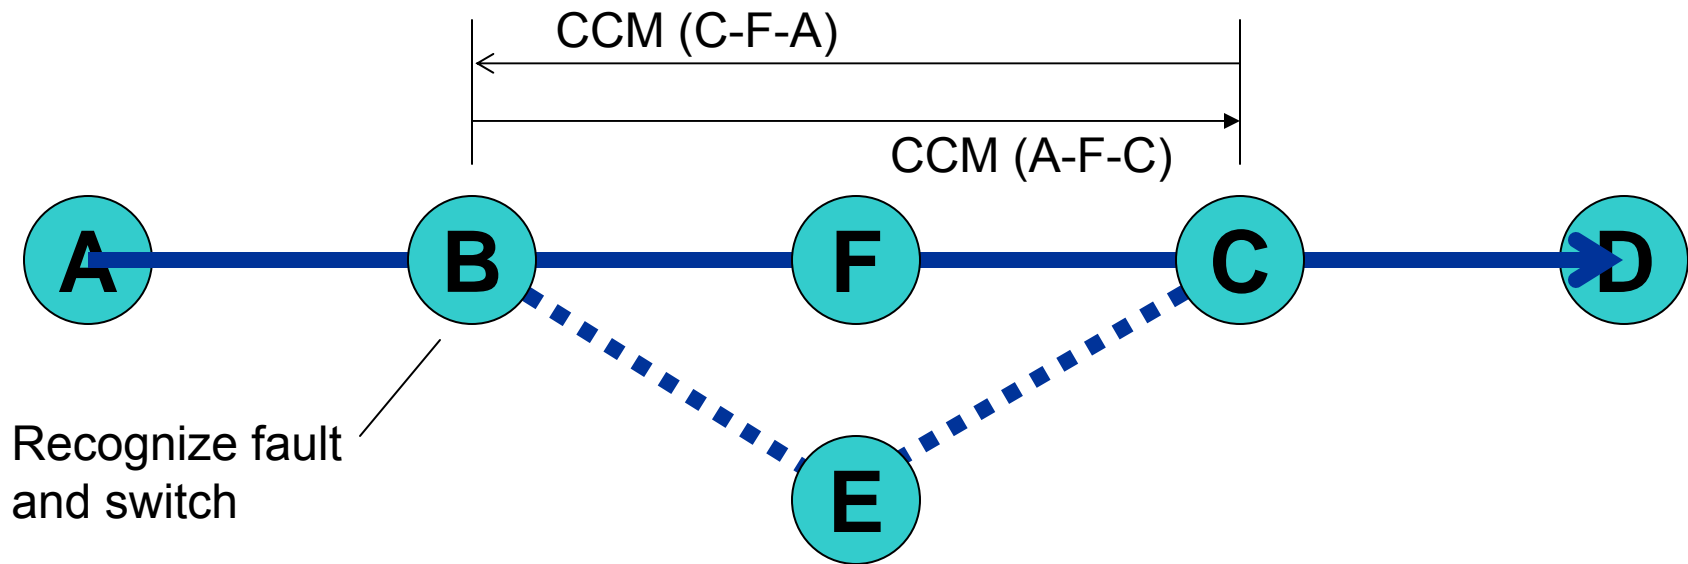
A bridge at which a TESI enters or exits a cycle must be configured with explicit FDB entries.

Wildcard FDB Entry Specification

The wildcard entry meets the conditions described in 802.1Q subclause clause 8.8.1 Static Filtering Entries (amended as 802.1Qay):

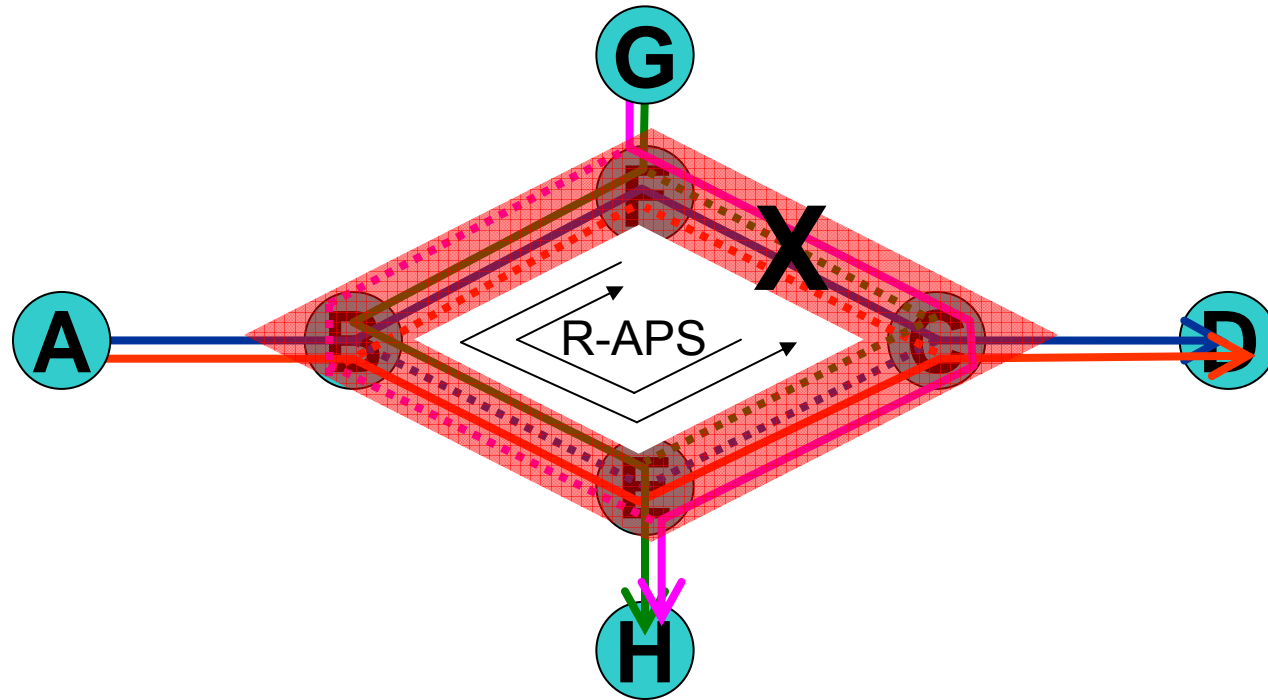
- a) MAC address for which no more specific static filtering entry exists (items 3-5).**
- b) VID(s) associated with the TE-MSTID**
- c) Port map with entries associated with the two cycle ports of the bridge**

Detecting faults using CFM



- **Segment ingress B detects connectivity failure**
- **Switch to protection segment**
- **Restore working segment when connectivity restored**
- **This is a simplified explanation. For complete proposal see: [ay-Abhay-Protection-Switching-for-P2MP-0508.ppt](#)**

Detecting faults using G.8032 R-APS



- Propagate Ring-APS on the cycle shared by a set of segments
- Each segment ingress determines whether the fault lies on the working segment and, if so performs protection switch.

CFM and G.8032 both standard solutions

- **CFM solution implies CCM traffic per segment but is a straightforward solution.**
- **R-APS solution uses information about the mapping of segments to cycles, but is a highly scalable solution.**
- **Since both are already standards, there is no reason to disallow either.**

Requirements

- **Preserves the connectivity of a TESI segment in the presence of the failure of a single bridge or link within that segment.**
- **Preserves connectivity across bi-connected TESI segments in the presence of a failure at the point of interconnection or a failure of bridges or links within both of the interconnected segments.**
- **Allows a bridge or link to be shared among multiple (overlapping) segments of an ESP.**
- **Prohibits ESP traffic from crossing a link in both directions (backtracking).**

Requirements

- **Operates independently of any 1:1 end-to-end TESI protection that may be deployed. A 1:1 TESI protection action is not performed until sufficient time elapses to allow the associated fault to be corrected by Fast Re-Route.**
- **Supports Fast Re-Route for both PtP and PtMP TESI.**
- **Segments are provisioned.**
- **Cycles *may* be provisioned to provide efficient fault notification.**
- **A ‘wildcard’ FDB entry can be used at a bridge that is not associated with multiple cycles and that lies within the segment.**

Title

- PAR for an amendment to an existing Standard 802.1Q-2005
- P802.1Qbb (or Qbc, etc., as appropriate)
- IEEE Standard for Local and Metropolitan Area Networks---Virtual Bridged Local Area Networks - Amendment: Fast ReRoute for Traffic Engineered Ethernet (FRR-TE)

Scope

- This standard specifies protocols, procedures, and managed objects to support the rapid restoration of a Traffic Engineered Service Instance (TESI) connectivity when a single failure occurs on a provisioned segment of the TESI.
- The scope also includes all the requirements listed in slides 28 and 29.

Purpose

- Provide rapid restoration of TESI connectivity for a single failure within each provisioned segment of a TESI. Limit re-routing of traffic to the TESI segment on which the failure occurs.

Need

- It is anticipated that Traffic Engineered bridged networks will be widely deployed when the PBB-TE (IEEE 802.1Qay) standard becomes available. Currently, only end-to-end 1:1 TESI protection is specified. Localized protection is supported by synchronous transport networks traditionally used to provide Traffic Engineered service and by Fast ReRoute (FRR) in MPLS-based packet transport networks. The absence of FRR capability in Bridged networks puts Bridging technology at a competitive disadvantage. Further, 802.1Qay does not specify a method of protection for a PtMP TESI. This requirement can be met using localized protection.

Stakeholders

- Vendors, users, administrators, designers, customers, and owners of traffic-engineered bridged networks.

Other standards with similar scope

- There are no standards solving this problem for IEEE 802.1Q Traffic Engineered Service Instances.

Five Criteria

Broad Market Potential

A standards project authorized by IEEE 802 shall have a broad market potential. Specifically, it shall have the potential for:

- **Broad sets of applicability.**
 - The commercial provision of Traffic Engineered services is a large and growing business.
- **Multiple vendors and numerous users.**
 - The same large body of vendors and users having a requirement for IEEE 802.1Qay.
- **Balanced costs (LAN versus attached stations).**
 - This project does not materially alter the existing cost structure of bridged networks.

Compatibility

- IEEE 802 defines a family of standards. All standards shall be in conformance with the IEEE 802.1 Architecture, Management, and Interworking documents as follows: 802. Overview and Architecture, 802.1D, 802.1Q, and parts of 802.1f. If any variances in conformance emerge, they shall be thoroughly disclosed and reviewed with 802.
 - This PAR is for an amendment to 802.1Q, thus ensuring compatibility.
- Each standard in the IEEE 802 family of standards shall include a definition of managed objects that are compatible with systems management standards.
 - Such a definition will be included.

Distinct Identity

Each IEEE 802 standard shall have a distinct identity. To achieve this, each authorized project shall be:

- Substantially different from other IEEE 802 standards.
 - This project will amend the only IEEE 802 standard defining Traffic Engineered bridged networks.
- One unique solution per problem (not two solutions to a problem).
 - There are no other standard solutions to localized recovery in a Traffic Engineered bridged network.
- Easy for the document reader to select the relevant specification.
 - This project will amend the only IEEE 802 standard defining Traffic Engineered bridged networks.

Technical Feasibility

For a project to be authorized, it shall be able to show its technical feasibility.

At a minimum, the proposed project shall show:

- Demonstrated system feasibility.
 - Several bridge vendors offer products that offer capabilities substantially the same as those defined by this project.
- Proven technology, reasonable testing.
 - Several bridge vendors offer products that offer capabilities substantially the same as those defined by this project. Compliance with the project can be tested using straightforward extensions of existing test tools for bridged networks.
- Confidence in reliability.
 - The reliability of the modified protocols will be not be measurably worse than that of the existing Traffic Engineered Bridged networks.

Economic Feasibility

For a project to be authorized, it shall be able to show economic feasibility (so far as can reasonably be estimated) for its intended applications. At a minimum, the proposed project shall show:

- Known cost factors, reliable data.
 - This project introduces no hardware costs beyond the minimal and well-known resources consumed by extending an existing software protocol.
- Reasonable cost for performance.
 - The cost of upgrading software and configuring a priori knowledge of the overall system topology is reasonable for the significant reduction in the time required to recover from a network failure.
- Consideration of installation costs.
 - The cost of installing enhanced software, in exchange for improved network performance, is familiar to vendors and users of bridged networks.