Mutually Exclusive Virtual TE Links

[DRAFT-MELG] draft-beeram-ccamp-melg

[DRAFT-SRCLG] draft-beeram-ccamp-srclg

Vishnu Pavan Beeram, Igor Bryskin, et al.,

Virtual TE-Link

Virtual TE-Link - Definition (RFC 6001):

- Advertised into a Client Network Domain
- Represents a potentiality to setup an LSP in the Server Network Domain to support the advertised TE-Link.
- Follow the same rules that are defined for the advertising, processing and use of regular TE-Links

Limitations with existing definition:

- No strict guidelines on how the underlying server LSP (what path) needs to get set up.
- Characteristics of the underlying server path not determined until the Virtual TE-Link gets committed.
  - Some important characteristics of the Virtual TE-Link (e.g. shared-risk and delay) not known to the client until the corresponding server LSP is set up
    - Mutual Exclusivity IS a key characteristic (more on this later..)

Virtual TE-Link: An enhanced view

Virtual TE-Link – An enhanced view (used by [DRAFT-MELG] and [DRAFT-SRCLG]):

- Aware of the key characteristics of the underlying server-path (while still uncommitted).
- Creation/Maintenance is driven by policy.
  - Policy determines which Virtual TE-Link to create (which end-points) and how the underlying server LSP (what path) needs to get set up.
- A Virtual TE-Link remains a Virtual TE-Link through-out its life-time.
  - It may get committed and uncommitted from time to time – but never loses its “Virtual” property.

The basic idea behind this “enhanced view” is that it makes the “Virtual TE-Link” get as close as it can to representing a “Real TE-Link”.

Mutually Exclusive Virtual TE Links

Mutual Exclusivity is a property that is specific to “Virtual” TE-Links.

Mutual Exclusivity comes into play when multiple Virtual TE-Links are dependent on the usage of the same underlying server resource.

- Since not all of these Virtual TE-Links can get committed at the same time, they are deemed to be mutually exclusive.

The existence of this “mutual exclusivity” property would need to be advertised into the Client TE Domain.

- This is of relevance to client path computation engines; especially those that are capable of doing concurrent computations.
- If this information is absent, there exists the risk of yielding erroneous concurrent path computation results where only a subset of the computed paths get successfully provisioned.

Static vs Dynamic

The “Mutual Exclusivity” property of a Virtual TE-Link can be either static or dynamic in nature.

Static Mutual Exclusivity

- Permanent within a given network configuration
- Comes into play when multiple Virtual TE-Links are dependent on the same non-shareable underlying server resource.
- Resource used in its entirety by a single Virtual TE-Link when committed; Such resources exist only in the WDM layer.

Dynamic Mutual Exclusivity

- Temporary within a given network configuration
- Comes into play when multiple Virtual TE-Links are dependent on the same shareable resource
- Mutual exclusivity exists when the amount of the resource that is available for sharing is limited; It ceases to exist when there is sufficient amount of the resource to accommodate all corresponding Virtual TE-Links; Such resources exist in all layers.

Construct Requirements - Static vs. Dynamic Mutual Exclusivity

- The advertisement paradigm of the TE construct required to carry static mutual exclusivity information is quite different from that of the TE construct required to carry dynamic mutual exclusivity information.
  - Static mutual exclusivity information can get advertised per TE-Link using a simple sub-TLV construct.
    - No scaling issues with this approach.
  - Advertising dynamic mutual exclusivity information per TE-link poses serious scaling concerns and hence requires a different type of construct/paradigm.
  - The TE construct for carrying static mutual exclusivity information is introduced in [DRAFT-MELG]; The construct for carrying the dynamic mutual exclusivity information is discussed in [DRAFT-SRcLG].

Mutually Exclusive Link Group [MELG]

draft-beeram-ccamp-melg

Vishnu Pavan Beeram (Ed), John Drake, Gert Grammel, Cyril Margaria (Juniper Networks)

Igor Bryskin (Ed) (ADVA Optical Networking)

Manuel Paul, Ruediger Kunze (Deutsche Telekom)

Friedrich Armbruster (Coriant GmbH)

Oscar Gonzalez de Dios (Telefonica)

Fatai Zhang (Huawei Technologies)

Daniele Ceccarelli (Ericsson)

Wes Doonan

Static Mutual Exclusivity

Client
TE Info

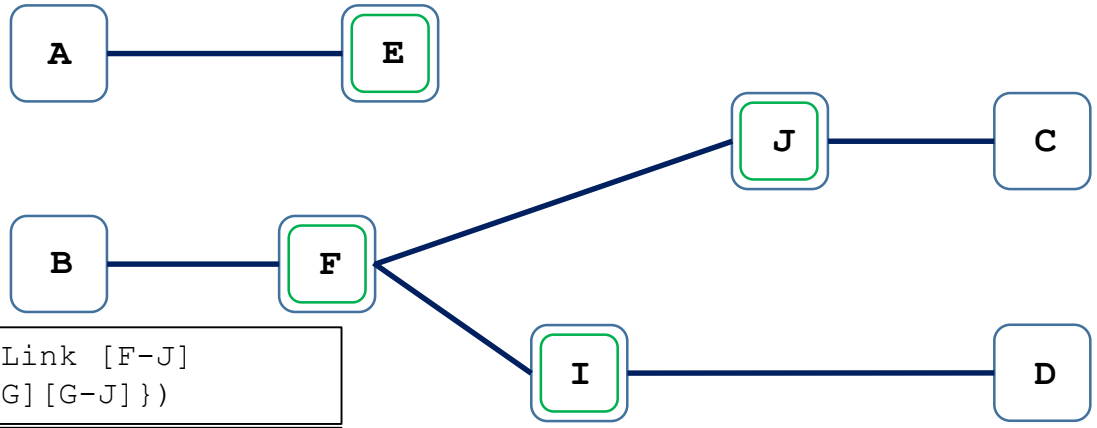


```
Create Virtual TE Link [F-J]
(Server-Path: {[F-G][G-J]})
```

```
Create Virtual TE Link [F-I]
(Server-Path: {[F-H][H-I]})
```

[F-J] and [F-I] depend on the usage of the **same non-shareable resource** – source transponder

Only one of the two Virtual TE Links can be committed at any given point of time.

Physical
Topology

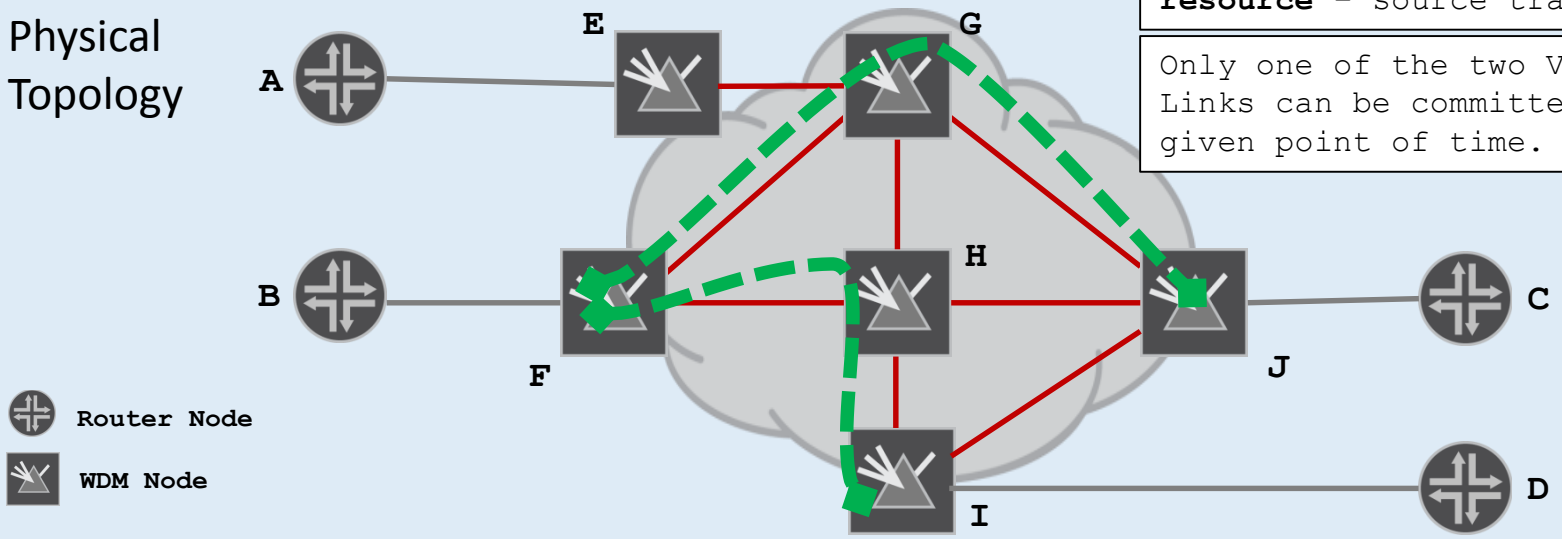

MELG – Mutually Exclusive Link Group

MELG - Construct for carrying static mutual exclusivity information in the Virtual TE-Link advertisement.

Construct Usage:

- Primary Purpose: Indicates via a network unique identifier that the advertised TE-Link belongs to one or more MELGs.
  - Client Path Computation function can decide on whether two or more Virtual TE-Links are mutually exclusive or not by finding an overlap of advertised MELGs
- Additionally, it indicates whether the advertised Virtual TE-Link is committed or not at the time of advertising
  - This allows the Computation function to show preference to already committed links.

MELG – Construct Format

```
Name: MELG
Type: TBD
Length: Variable
```

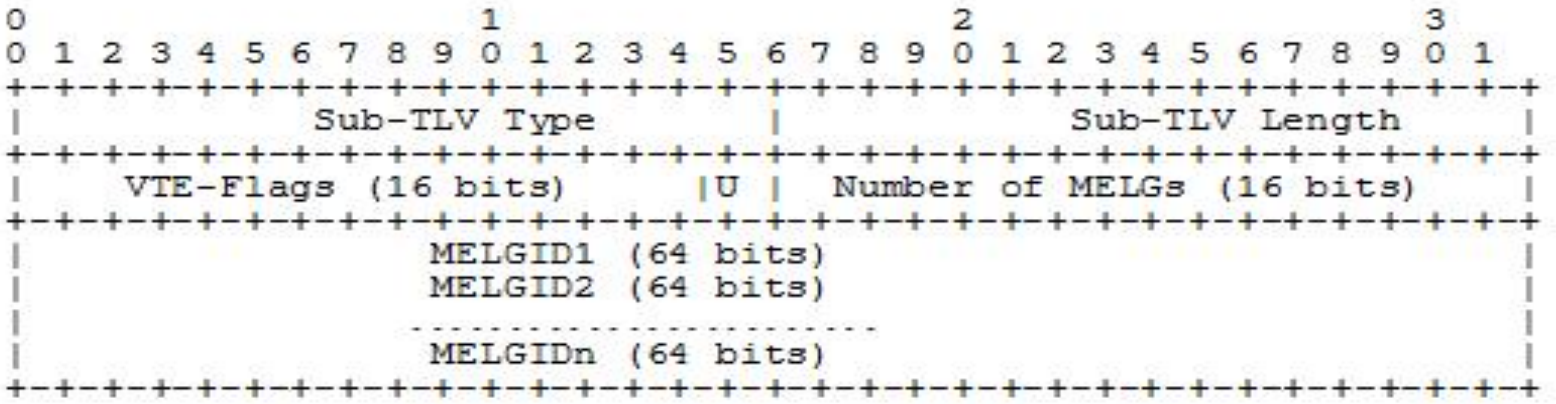


```
Number of MELGs:                 number of MELGS advertised for the
                                 Virtual TE Link;
VTE-Flags:                       Virtual TE Link specific flags;
MELGID1,MELGID2,...,MELGIDn:     64-bit network domain unique numbers
                                 associated with each of the advertised
                                 MELGs

Currently defined Virtual TE Link specific flags are:
   U bit (bit 1): Uncommitted, if set, the Virtual TE Link is
   uncommitted at the time of the advertising (i.e. the server layer
   network LSP is not set up); if cleared, the Virtual TE Link is
   committed (i.e. the server layer LSP is fully provisioned and
   functioning). All other bits of the "VTE-Flags" field are
   reserved for future use and MUST be cleared.
```

Shared Resource Link Group [SRcLG]

draft-beeram-ccamp-srclg

Vishnu Pavan Beeram (Ed), Cyril Margaria (Juniper Networks)

Igor Bryskin (Ed) (ADVA Optical Networking)

Fatai Zhang (Huawei Technologies)

Dynamic Mutual Exclusivity

Client
TE Info

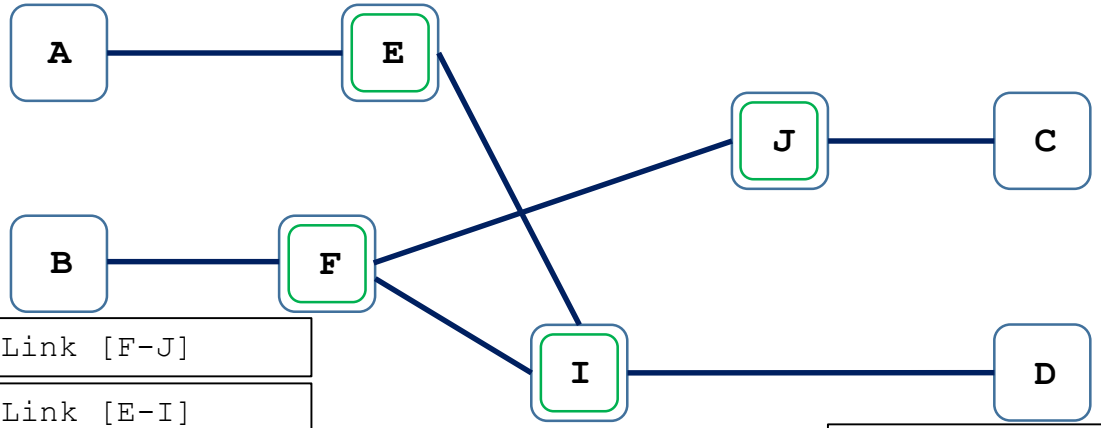

Create Virtual TE Link [F-J]

Create Virtual TE Link [E-I]

Create Virtual TE Link [F-I]

All 3 Server-Paths traverse {H-I}

Assumption: WDM Link {H-I} has only 2 lambda channels available.

Only two of the three Virtual TE Links – [E-I], [F-I] and [F-J] can be concurrently committed.

Mutual Exclusivity ceases to exist when third lambda channel becomes available on {H-I})

Physical
Topology

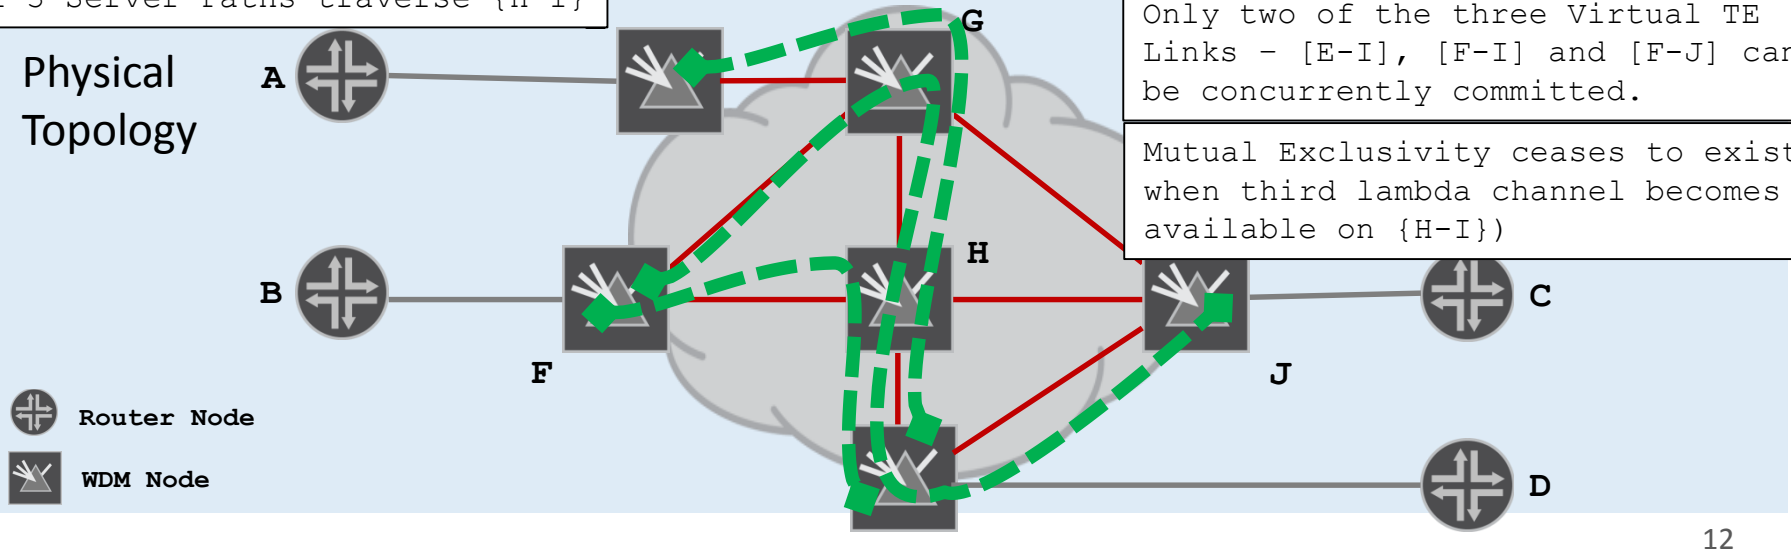

Shared Resource Link Group [SRcLG]

- Shared Resource Link Group (SRcLG) is meaningful only in the context of Virtual TE-Links.
- SRcLG represents a set of Virtual TE-Links that depend on the usage of a shared server-layer resource that has a variable bandwidth capacity and as a result may sometimes not be able to simultaneously accommodate all corresponding Virtual TE-Links in the set.
  - Since dynamic mutual exclusivity comes into play only when the underlying server resource is shareable, all Virtual TE-Links in an SRcLG would also belong to the same SRLG.
  - As is the case with SRLGs, a given Virtual TE-Link may belong to multiple SRcLGs.

SRcLG - Construct

- In terms of the TE construct, an SRcLG is nothing but an SRLG with some additional information to help determine which and how many of the corresponding Virtual TE Links can get committed simultaneously.
  - This additional information is the per-priority available shared resource bandwidth associated with a given SRcLG.
  - Since an SRcLG cannot exist without the presence of a corresponding SRLG, the SRcLG is identified by the corresponding 32-bit SRLG-ID.
  - Unlike the SRLG construct or the MELG construct, the SRcLG construct does not get advertised per TE-Link (to avoid scaling concerns).

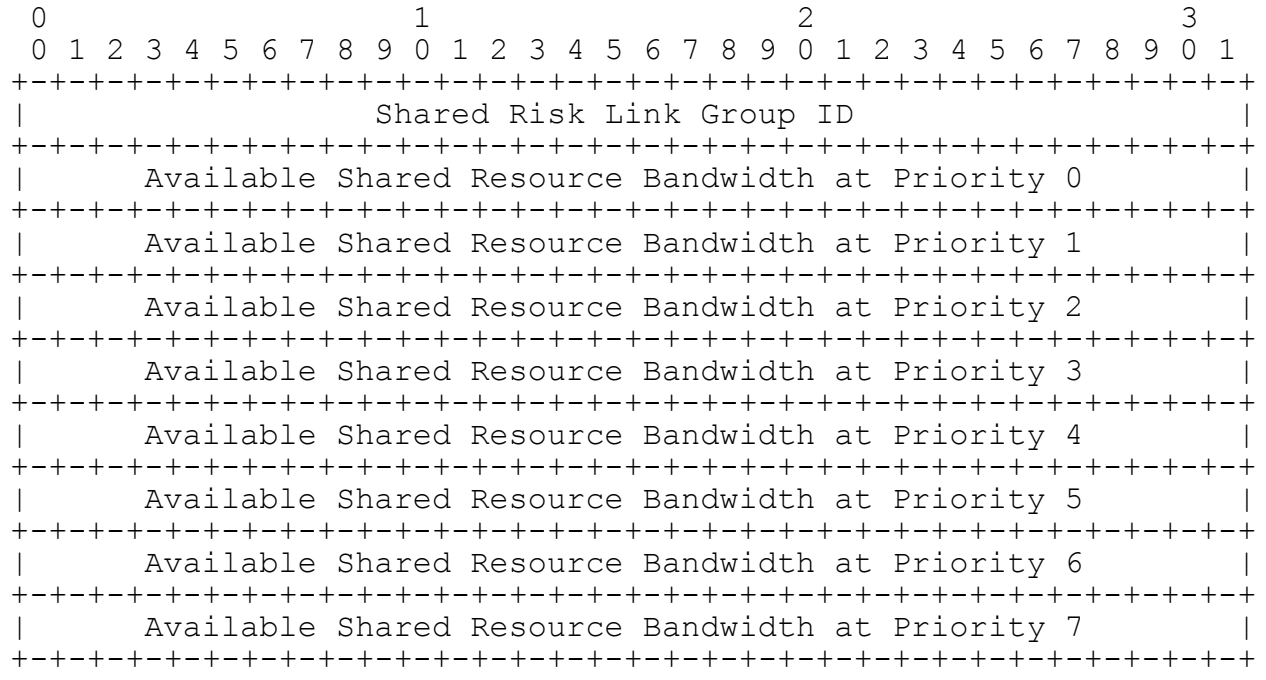

Advertising Rules

- As far as the advertisement of a Virtual TE-Link is concerned, there is no perceived difference between SRLG and SRcLG.
  - The 32-bit IDs of all SRcLGs that a Virtual TE-Link belongs to are advertised via the SRLG construct.
- Additionally, all SRcLG information associated with a given Virtual Topology is advertised into the Client TE Domain by the provider of the Virtual Topology.
  - It is the responsibility of this provider to keep the bandwidth availability information for each SRcLG current with timely updates.
  - The draft envisions that one or more server domain OSPF/ISIS TE speakers will be tasked to provide these timely updates. This TE speaker may advertise all SRcLG information (that it is responsible for) in the same OSPF-LSA/ISIS-LSP or advertise each SRcLG TLV separately – one in each OSPF-LSA/ISIS-LSP.

Processing Rules

- Intended consumer of the SRcLG information is the PCE in the Client TE Domain.
  - Client PCE should take this advertised information into account when performing path selection for services over the Virtual Topology provided by the network domain.
    - In particular, this information should be used when deciding how many Virtual TE-Links could be accommodated simultaneously on a given SRcLG at a given priority level.

Thank You