CCAMP Working Group
Internet Draft
Intended status: Standards Track

Vishnu Pavan Beeram (Ed)
Juniper Networks
Igor Bryskin (Ed)
ADVA Optical Networking

Expires: January 15, 2013

July 15, 2013

# Mutually Exclusive Link Group (MELG)
## draft-beeram-ccamp-melg-01.txt

## Status of this Memo



## Copyright Notice

Abstract

This document introduces the concept of MELG ("Mutually Exclusive Link Group") and discusses its usage in the context of mutually exclusive Virtual TE Links.

Conventions used in this document

The key words "MUST", "MUST NOT", "REQUIRED", "SHALL", "SHALL NOT", "SHOULD", "SHOULD NOT", "RECOMMENDED", "MAY", and "OPTIONAL" in this document are to be interpreted as described in RFC-2119 [RFC2119].

Table of Contents



## 1. Introduction

A Virtual TE Link (as defined in [RFC6001]) advertised into a Client Network Domain represents a potentiality to setup an LSP in the Server Network Domain to support the advertised TE link. The Virtual TE Link gets advertised like any other TE link and follows exactly the same rules that are defined for the advertising, processing and use of regular TE links [RFC4202]. However, "mutual exclusivity" is one attribute that is specific to Virtual TE links. This document discusses the need to advertise this information and the means to do so.

## 2. Mutually Exclusive Virtual TE Links

Consider the network topology depicted in Figure 1a. This is a typical packet optical transport deployment scenario where the WDM layer network domain serves as a Server Network Domain providing transport connectivity to the packet layer network Domain (Client Network Domain).

```
                                  |
                                  | +---+            /-\
                                  | |   | Router    (   ) WDM
                                  | +---+ Node       \-/  node
                                  |________________________________

     +---+        /-\          /-\           /-\          +---+
     | B |-------( E )--------( G )---------( J )---------| C |
     +---+        \-/          \-/           \-/          +---+
                              /   \         /   \
                             /     \       /     \
                            /       \     /       \
                           /         \   /         \
                          /           \ /           \
         +---+          /-\           /-\           /-\          +---+
         | A |---------( F )---------( H )---------( I )---------| D |
         +---+          \-/           \-/           \-/          +---+
```

Figure 1a: Sample topology

```
         -------------                        |  [ ] Client TE Node
         | Client TE |                        |  +++ Client TE Link
         | DataBase  |                        |_____________________
         -------------
            [B] ++++++++ [E]                  [J] +++++++++ [C]


            [A] ++++++++ [F]                  [I] +++++++++ [D]
```

Figure 1b: Client TE Database

Nodes A, B, C and D are IP routers that are connected to an Optical WDM transport network. E, F, G, H, I and J are WDM nodes that constitute the Server Network Domain. The border nodes (E, F, I and J) operate in both the server and client domains. Figure 1b depicts how the Client Network Domain TE topology looks like when there are no Client TE Links provisioned across the optical domain.

```
                                | *****  F-J WDM Path (lambda 192000)
                                | @@@@@  E-I WDM Path (lambda 192000)
                                |________________________________

 +---+        /-\ @@@@@@@@ /-\           /-\          +---+
 | B |-------( E )--------( G )---------( J )---------| C |
 +---+        \-/         *\-/*@       @*\-/@         +---+
                         */   \*@     @*/   \@
                        */     \*@   @*/     \@
                       */       \*@ @*/       \@
                      */         \*@*/         \@
                     */           \*/           \@
     +---+          /-\           /-\           /-\          +---+
     | A |---------( F )---------( H )---------( I )---------| D |
     +---+          \-/           \-/           \-/          +---+
```

Figure 2a: Mutually Exclusive potential WDM paths

```
      ------------   |  TE-Links E-I and F-J are mutually exclusive
      | Client-TE|   |  Advertised with MELG-ID - 25/192000
      | Database |   |  [SRLG-ID 25; Shared Resource ID 192000]
      ------------   |_____________________________________________

        [B] ++++++++ [E]                      [J] +++++++++ [C]
                         ++++            +++++
                             +++      +++
                                ++++++
                             +++      +++
                         ++++            +++++
        [A] ++++++++ [F]                      [I] +++++++++ [D]
```

Figure 2b: Client TE Database - Mutually Exclusive Virtual TE Links

Now consider augmenting the Client TE topology by creating a couple of Virtual TE Links across the optical domain. The potential paths in the WDM network catering to these two virtual TE links are as shown in Fig 2a and the corresponding augmented Client TE topology is as illustrated in Fig 2b.

In this particular example, the potential paths in the WDM layer network supporting the Virtual TE Links not only intersect, but also require the usage of the same resource (lambda channel 192000) on the intersection. Because the Virtual TE Links depend on the same uncommitted network resource, only one of them could get activated at any given time. In other words they are mutually exclusive. This scenario is encountered when the potential paths depend on any common physical resource (e.g. transponder, regenerator, wavelength converter, etc.) that could be used by only one Server Network Domain LSP at a time.

For a Client Network Domain path computation function (especially a centralized one) that is capable of concurrent computation of multiple paths, it is important to know the existence of mutually exclusive relationship between Virtual TE Links. Absent this information, there exists the risk of yielding erroneous concurrent path computation results where only a subset of the computed paths can get successfully provisioned. This document introduces the concept of Mutually Exclusive Link Group to address this problem.

The Virtual TE Link mutual exclusivity attribute can be either (a) Static or (b) Dynamic.

In case (a), the Virtual TE Link mutual exclusivity exists permanently within a given network configuration. This type of mutual exclusivity comes into play when two or more Virtual TE Links depend on a Server Network Domain resource that could be used in its entirety by only one Virtual TE Link (when committed).

In case (b), the Virtual TE Link mutual exclusivity exists temporarily within a given network configuration. This type of mutual exclusivity comes into play when two or more Virtual TE Links depend on a Server Network Domain resource that could be shared simultaneously in some limited capacity by several Virtual TE Links (when committed). Consider, for example, a situation when three Virtual TE Links depend on a Server Domain WDM link that currently has two lambda channels available. Consider further that each Virtual TE Link (in order to be committed) requires one lambda channel to be allocated on said WDM link. Obviously, under these conditions only two out of three Virtual TE Links could be concurrently committed. Such Virtual TE Link mutual exclusivity is dynamic and temporary because as soon as additional lambda channels become available on the WDM link, the Virtual TE Link mutual exclusivity will cease to exist - it will become possible to commit all three Virtual TE Links concurrently.

This revision of the draft discusses only the Static Virtual TE Link mutual exclusivity. Dynamic Virtual TE Link mutual exclusivity will be addressed in later revisions.

## 3. Mutually Exclusive Link Group

The Mutually Exclusive Link Group (MELG) construct defined in this document has 2 purposes

- To indicate via a separate network unique number (MELG ID) an element or a situation that makes the advertised Virtual TE Link belong to one or more Mutually Exclusive Link Groups. Path computing element will be able to decide on whether two or more Virtual TE Links are mutually exclusive or not by finding an overlap of advertised MELGs (similar to deciding on whether two or more TE links share fate or not by finding common SRLGs)

- To indicate whether the advertised Virtual TE Link is committed or not at the moment of the advertising. Such information is important for a path computation element: Committing new Virtual TE links (vs. re-using already committed ones) has a consequence of allocating more server layer resources and disabling other Virtual TE Links that have common MELGs with newly committed Virtual TE Links; Committing a new Virtual TE Link also means a longer setup time for the Client LSP and higher risk of setup-failure.

## 4. Protocol Extensions

### 4.1. OSPF

The MELG is a sub-TLV of the top level TE Link TLV. It may occur at most once within the Link TLV. The format of the MELGs sub-TLV is defined as follows:

Name: MELG
Type: TBD
Length: Variable

```
 0                   1                   2                   3
 0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
|             Sub-TLV Type      |          Sub-TLV Length       |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
|    VTE-Flags (16 bits)    |U |  Number of MELGs (16 bits)     |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
|                 MELGID1 (64 bits)                             |
|                 MELGID2 (64 bits)                             |
|                ........................                       |
|                 MELGIDn (64 bits)                             |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
```

```
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.

Note: A Virtual TE Link advertisement MAY include MELGs sub-TLV with zero MELGs for the purpose of communicating to the TE domain whether the Virtual TE Link is currently committed or not.

### 4.2. ISIS

The MELG TLV (of type TBD) contains a data structure consisting of:

```
6         octets of System ID
1         octet of Pseudonode Number
1         octet Flag
4         octets of IPv4 interface address or 4 octets of a Link
          Local Identifier
4         octets of IPv4 neighbor address or 4 octets of a Link
          Remote Identifier
2         octets MELG-Flags
```

```
      2        octets - Number of MELGs
      variable List of MELG values, where each element in the list
                   has 8 octets
```

The following illustrates encoding of the value field of the MELG TLV.

```
    0                   1                   2                   3
    0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1
   +-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
   |                           System ID                           |
   +-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
   |        System ID (cont.)      |Pseudonode num |    Flags      |
   +-+-+-+-+-+-+-+-+-+-+-+-+-+-+-++-+-+-+-+-+-+-+-+-+-++-+-+-+-+-+-+-+
   |          Ipv4 interface address/Link Local Identifier         |
   +-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
   |          Ipv4 neighbor address/Link Remote Identifier         |
   +-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
   |    VTE-Flags (16 bits)     |U |  Number of MELGs (16 bits)    |
   +-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
   |                 MELGID1 (64 bits)                             |
   |                 MELGID2 (64 bits)                             |
   |                ........................                       |
   |                 MELGIDn (64 bits)                             |
   +-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
```

The neighbor is identified by its System ID (6 octets), plus one octet to indicate the pseudonode number if the neighbor is on a LAN interface.

The least significant bit of the Flag octet indicates whether the interface is numbered (set to 1) or unnumbered (set to 0). All other bits are reserved and should be set to 0.

The length of the TLV is 20 + 8 * (number of MELG values).

The semantics of "VTE-Flags", "Number of MELGs" and "MELGID Values" are the same as the ones defined under OSPF extensions.

The MELG TLV MAY occur more than once within the IS-IS Link State Protocol Data Units.

## 5. Security Considerations

TBD

## 6. IANA Considerations

### 6.1. OSPF

IANA is requested to allocate a new sub-TLV type for MELG (as defined in Section 4.1) under the top-level TE Link TLV.

### 6.2. ISIS

IANA is requested to allocate a new IS-IS TLV type for MELG (as defined in Section 4.2).

## 7. Normative References

[RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, March 1997.

[RFC4202] K.Kompella, Y.Rekhter, "Routing Extensions in Support of Generalized Multi-Protocol Label Switching (GMPLS)", RFC4202, October 2005.

[RFC6001] D.Papadimitriou, M.Vigoureax, K.Shiomoto, D.Brungard and JL. Le Roux, "GMPLS Protocol Extensions for Multi-Layer and Multi-Region Networks", RFC 6001, October 2010.

## 8. Acknowledgments

Chris Bowers [cbowers@juniper.net]

Authors' Addresses

Vishnu Pavan Beeram
Juniper Networks
Email: vbeeram@juniper.net

Igor Bryskin
ADVA Optical Networking
Email: ibryskin@advaoptical.com

John Drake
Juniper Networks
Email: jdrake@juniper.net

Gert Grammel
Juniper Networks
Email: ggrammel@juniper.net

Wes Doonan
ADVA Optical Networking
Email: wdoonan@advaoptical.com

Manuel Paul
Deutsche Telekom
Email: Manuel.Paul@telekom.de

Ruediger Kunze
Deutsche Telekom
Email: Ruediger.Kunze@telekom.de

Oscar Gonzalez de Dios
Telefonica
Email: ogondio@tid.es

Cyril Margaria
Coriant GmbH
Email: cyril.margaria@coriant.com

Friedrich Armbruster
Coriant GmbH
Email: friedrich.armbruster@coriant.com

Daniele Ceccarelli
Ericsson
Email: daniele.ceccarelli@ericsson.com