

Internet Engineering Task Force (IETF)
Request for Comments: 6003
Updates: 3471, 3473
Category: Standards Track
ISSN: 2070-1721

D. Papadimitriou
Alcatel-Lucent
October 2010

Ethernet Traffic Parameters

Abstract

This document describes the support of Metro Ethernet Forum (MEF) Ethernet traffic parameters as described in MEF10.1 when using Generalized Multi-Protocol Label Switching (GMPLS) Resource ReSerVation Protocol - Traffic Engineering (RSVP-TE) signaling.

Status of This Memo

This is an Internet Standards Track document.

This document is a product of the Internet Engineering Task Force (IETF). It represents the consensus of the IETF community. It has received public review and has been approved for publication by the Internet Engineering Steering Group (IESG). Further information on Internet Standards is available in Section 2 of RFC 5741.

Information about the current status of this document, any errata, and how to provide feedback on it may be obtained at <http://www.rfc-editor.org/info/rfc6003>.

Copyright Notice

Copyright (c) 2010 IETF Trust and the persons identified as the document authors. All rights reserved.

This document is subject to BCP 78 and the IETF Trust's Legal Provisions Relating to IETF Documents (<http://trustee.ietf.org/license-info>) in effect on the date of publication of this document. Please review these documents carefully, as they describe your rights and restrictions with respect to this document. Code Components extracted from this document must include Simplified BSD License text as described in Section 4.e of the Trust Legal Provisions and are provided without warranty as described in the Simplified BSD License.

1. Introduction

Per [RFC3471], Generalized Multi-Protocol Label Switching (GMPLS) allows the inclusion of technology-specific parameters in signaling. This document introduces Ethernet SENDER_TSPEC and FLOWSPEC-specific objects in support of Metro Ethernet Forum (MEF) Ethernet traffic parameters as specified in [MEF10.1] and ITU-T Ethernet Service Switching as discussed in [RFC6004]. For example:

- o For Ethernet Private Line (EPL) services [MEF6], these traffic parameters are applicable to each Ethernet Virtual Connection (EVC) crossing a given port.
- o For Ethernet Virtual Private Line (EVPL) services [MEF6], these traffic parameters are applicable per Ethernet Virtual Connection (EVC) with a single or multiple Class of Service (CoS), independent of its associated Virtual LAN ID (VID) or set of VIDs.

Association between EVC and VIDs is detailed in [MEF10.1]. The format and encoding of the VID (or set of VIDs) is documented in a companion document [RFC6004].

This does not preclude broader usage of the Ethernet SENDER_TSPEC and FLOWSPEC-specific objects specified this document. For instance, they may also be used for signaling Ethernet Label Switched Paths (LSPs), in the Generalized Label Request (see [RFC3471]), the Switching Type field is set to Layer 2 Switching Capability (L2SC) and the LSP Encoding Type field to Ethernet.

2. 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].

Moreover, the reader is assumed to be familiar with the terminology in [MEF10.1] as well as in [RFC3471] and [RFC3473].

3. Overview

In GMPLS RSVP-TE [RFC3473], the SENDER_TSPEC object is used on a Path message to indicate the bandwidth that is requested for the LSP being established, and the FLOWSPEC object is used on a Resv message to indicate the bandwidth actually reserved for the LSP. The Ethernet SENDER_TSPEC/FLOWSPEC object includes the Ethernet link type (switching granularity) of the requested LSP and the MTU value for

the LSP. Other information about the requested bandwidth characteristics of the LSP are carried in the Bandwidth Profile as a TLV within the Ethernet SENDER_TSPEC/FLOWSPEC object.

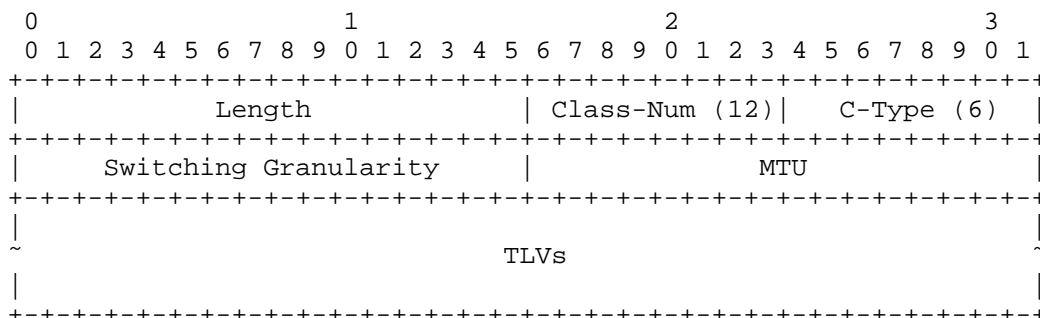
The Ethernet SENDER_TSPEC/FLOWSPEC object includes the Ethernet link type (switching granularity) of the requested LSP and the MTU value for the LSP.

The Bandwidth Profile defines the set of traffic parameters applicable to a sequence of Service Frames, referred to as bandwidth profile parameters (as specified in [MEF10.1]):

- o Committed Rate: indicates the rate at which traffic commits to be sent to the Ethernet LSP. The committed rate is described in terms of the CIR (Committed Information Rate) and CBS (Committed Burst Size) traffic parameters.
 - o CIR is defined as the average rate (in bytes per unit of time) up to which the network is committed to transfer frames and meets its performance objectives.
 - o CBS defines a limit on the maximum number of information units (e.g., bytes) available for a burst of frames sent at the interface speed to remain CIR-conformant.
- o Excess Rate: indicates the extent by which the traffic sent on an Ethernet LSP exceeds the committed rate. The Excess Rate is described in terms of the EIR (Excess Information Rate) and EBS (Excess Burst Size) traffic parameters.
 - o EIR is defined as the average rate (in bytes per unit of time), in excess of the CIR, up to which the network may transfer frames without any performance objectives.
 - o EBS defines a limit on the maximum number of information units (e.g., bytes) available for a burst of frames sent at the interface speed to remain EIR-conformant.
- o Color mode (CM): indicates whether the "color-aware" or "color-blind" property is employed by the bandwidth profile.
- o Coupling flag (CF): allows the choice between two modes of operation of the rate enforcement algorithm.

4. Ethernet SENDER_TSPEC Object

The Ethernet SENDER_TSPEC object (Class-Num = 12, Class-Type = 6) has the following format:



Switching Granularity (SG): 16 bits

This field indicates the type of link that comprises the requested Ethernet LSP.

The permitted Ethernet Link Type values are:

| Value | Switching Granularity |
|-------|--|
| 0 | Provided in signaling. See [RFC6004]. |
| 1 | Ethernet Port (for port-based service) |
| 2 | Ethernet Frame (for EVC-based service) |
| 255 | Reserved |

Values 0 to 2 are specified by the present document. Values 3 through 239 are to be assigned by IANA via Standards Action [RFC5226]. Value 255 is reserved by the present document (its Length is to be determined by the RFC that will specify it).

Values 240 through 254 are reserved for vendor-specific use.

Values 256 through 65535 are not assigned at this time.

MTU: 16 bits


This is a two-octet value indicating the MTU in octets.

The MTU field MUST NOT take a value smaller than 46 bytes for Ethernet v2 [ETHv2] and 38 bytes for IEEE 802.3 [IEEE802.3].

TLV (Type-Length-Value):

The Ethernet SENDER_TSPEC object MUST include at least one TLV and MAY include more than one TLV.

Each TLV MUST have the following format:



Type: 16 bits

Defined values are:

| Type | Length | Format | Description |
|------|--------|-----------------|--------------------------------------|
| 0 | - | Reserved | Reserved value |
| 1 | - | Reserved | Reserved value |
| 2 | 24 | see Section 3.1 | Ethernet Bandwidth Profile [MEF10.1] |
| 3 | 8 | [RFC6004] | Layer 2 Control Protocol (L2CP) |
| 255 | - | Reserved | Reserved value |

Values 0, 1, and 255 are reserved by the present document. Values 2 and 3 are specified by the present document.

Values 4 through 239 are to be assigned by IANA via Standards Action [RFC5226].

Values 240 through 254 are reserved for vendor-specific use.

Values 256 through 65535 are not assigned at this time.


Length: 16 bits

Indicates the length in bytes of the whole TLV including the Type and Length fields. A value field whose length is not a multiple of four MUST be zero-padded (with trailing zeros) so that the TLV is four-octet aligned.

4.1. Ethernet Bandwidth Profile TLV

The Type 2 TLV specifies the Ethernet Bandwidth Profile (BW profile). It defines an upper bound on the volume of the expected service frames belonging to a particular Ethernet service instance. The Ethernet SENDER_TSPEC object MAY include more than one Ethernet Bandwidth Profile TLV.

The Type 2 TLV has the following format:



Profile: 8 bits

This field is defined as a bit vector of binary flags. The following flags are defined:

- Flag 1 (bit 0): Coupling Flag (CF)
- Flag 2 (bit 1): Color Mode (CM)

Where bit 0 is the low order bit. Other flags are reserved, they SHOULD be set to zero when sent, and SHOULD be ignored when received.

A flag is set to value 1 to indicate that the corresponding metering profile is requested.

The Flag 1 (CF) allows the choice between two modes of operation of the rate enforcement algorithm.

The Flag 2 (CM) indicates whether the color-aware or color-blind property [MEF10.2] is employed by the bandwidth profile. When Flag 2 is set to value 0 (1), the bandwidth profile algorithm is said to be in color-blind (color-aware) mode.

Index: 8 bits

The Index field is used to reference bandwidth allocated for a given traffic class in case a multiple-class LSP is being requested. The Index field value MUST correspond to at least one of the Class-Type values included either in the CLASSTYPE object [RFC4124] or in the EXTENDED_CLASSTYPE object [MCOS].

A given index value j can be associated to at most N Class-Type values CT_i ($i \leq N$) of the EXTENDED_CLASSTYPE object. This association applies when a set of one or more CT_i s maps to a single (shared) BW profile. An example of value setting consists in assigning an arbitrary value comprised within the range [0x08,0xF8] associated to a set of CT_i , the values in the range [0xF8,0xFF] being selected for reserved sets. This allows mapping to one of 248 predefined CT_i sets.

A given index value j can be associated to a single CT_i (1:1 correspondence). In this case, the index value setting consists in assigning the 3 least significant bits of the Index field itself to the CT_i value itself (comprised in the range [0x00,0x07]). This applies in case a single CT_i maps a single (dedicated) BW profile or multiple (dedicated) BW profiles. In the former case, the Ethernet SENDER_TSPEC object includes a single Ethernet Bandwidth Profile TLV. In the latter case, the Ethernet SENDER_TSPEC includes a set of more than one Ethernet Bandwidth Profile TLVs (whose respective index value is associated to a single CT_i value).

Note that the current specification allows for combining shared and dedicated BW profiles to the same LSP. That is, an Ethernet SENDER_TSPEC object MAY include multiple Ethernet Bandwidth Profile TLVs whose respective index can be associated on a 1:1 basis to a single CT_i or to a set of multiple CT_i s.

For each subobject of the EXTENDED_CLASSTYPE object [MCOS]:

- o Each CT_i value SHOULD correspond 1:1 to the MEF Customer Edge VLAN CoS (CE-VLAN CoS).
- o The BW requested per CT_i field MAY be used for bandwidth accounting purposes.

By default, the value of the Index field MUST be set to 0.

Reserved: 16 bits

These bits SHOULD be set to zero when sent and MUST be ignored when received.

CIR (Committed Information Rate): 32 bits

The value of the CIR is in units of bytes per second. The CIR is encoded as a 32-bit IEEE single-precision floating-point number (see [RFC4506]).

The CIR value MUST be greater than or equal to 0.

CBS (Committed Burst Size): 32 bits

The value of the CBS is in units of bytes. The CBS is encoded as a 32-bit IEEE single-precision floating-point number (see [RFC4506]).

When CIR is strictly greater than 0 ($CIR > 0$), the CBS MUST be greater than or equal to the maximum frame size.

EIR (Excess Information Rate): 32 bits

The value of the EIR is in units of bytes per second. The EIR is encoded as a 32-bit IEEE single-precision floating-point number (see [RFC4506]).

The EIR value MUST be greater than or equal to 0.

EBS (Excess Burst Size): 32 bits

The value of the EBS is in units of bytes. The EBS is encoded as a 32-bit IEEE single-precision floating-point number (see [RFC4506]).

When EIR is strictly greater than 0 ($EIR > 0$), the EBS MUST be greater than or equal to the maximum frame size.

5. Ethernet FLOWSPEC Object

The Ethernet FLOWSPEC object (Class-Num = 9, Class-Type = 6) has the same format as the Ethernet SENDER_TSPEC object.

6. Ethernet ADSPEC Object

There is no ADSPEC object associated with the Ethernet SENDER_TSPEC object.

Either the ADSPEC object is omitted or an IntServ ADSPEC with the Default General Characterization Parameters and Guaranteed Service fragment is used, see [RFC2210].

7. Processing

The Ethernet SENDER_TSPEC and FLOWSPEC objects specified in this document MAY be used for signaling Ethernet LSP. For signaling such an LSP, in the Generalized LABEL_REQUEST object (see [RFC3471]), the Switching Type field MUST be set to the value 51 (L2SC) and the LSP Encoding Type field MUST be set to the value 2 (Ethernet).

The Ethernet SENDER_TSPEC object carries the traffic specification generated by the RSVP session sender. The Ethernet SENDER_TSPEC object SHOULD be forwarded and delivered unchanged to both intermediate and egress nodes.

The Ethernet FLOWSPEC object carries reservation request information generated by receivers. As with any FLOWSPEC object, the Ethernet FLOWSPEC object flows upstream toward the ingress node.

Intermediate and egress nodes MUST verify that the node itself and the interfaces on which the LSP will be established can support the requested Switching Granularity, MTU and values included in subobject TLVs. These nodes MUST be configured with the same predefined CT sets as the index value signaled as part of the Index field of the Ethernet Bandwidth Profile TLV (see Section 4.1). If the requested value(s) cannot be supported, the receiver node MUST generate a PathErr message with the error code "Traffic Control Error" and the error value "Service unsupported" (see [RFC2205]).

In addition, if the MTU field is received with a value smaller than the minimum transfer unit size of the Ethernet frame (e.g., 46 bytes for Ethernet v2, 38 bytes for IEEE 802.3), the node MUST generate a PathErr message with the error code "Traffic Control Error" and the error value "Bad Tspec value" (see [RFC2205]).

Error processing of the CLASSTYPE object follows rules defined in [RFC4124]. Error processing of the EXTENDED_CLASSTYPE object follows rules defined in [MCOS]. Moreover, a Label Switching Router (LSR) receiving a Path message with the EXTENDED_CLASSTYPE object, which recognizes the object and the particular Class-Type but does detect a mismatch in the index values, MUST send a PathErr message towards the sender with the error code "Extended Class-Type Error" and the error value "Class-Type mismatch" (see [RFC2205]).

8. Security Considerations

This document introduces no new security considerations to [RFC3473].

GMPLS security is described in Section 11 of [RFC3471] and refers to [RFC3209] for RSVP-TE. Further details of MPLS-TE and GMPLS security can be found in [RFC5920].

9. IANA Considerations

IANA maintains registries and sub-registries for RSVP-TE as used by GMPLS. IANA has made allocations from these registries as set out in the following sections.

9.1. RSVP Objects Class Types

This document introduces two new Class Types for existing RSVP objects. IANA has made allocations from the "Resource ReSerVation Protocol (RSVP) Parameters" registry using the "Class Names, Class Numbers, and Class Types" sub-registry.

| Class Number | Class Name | Reference |
|--------------|------------|-----------|
| ----- | ----- | ----- |
| 9 | FLOWSPEC | [RFC2205] |

Class Type (C-Type):

| | | |
|---|-----------------------|-----------|
| 6 | Ethernet SENDER_TSPEC | [RFC6003] |
|---|-----------------------|-----------|

| Class Number | Class Name | Reference |
|--------------|--------------|-----------|
| ----- | ----- | ----- |
| 12 | SENDER_TSPEC | [RFC2205] |

Class Type (C-Type):

| | | |
|---|-----------------------|-----------|
| 6 | Ethernet SENDER_TSPEC | [RFC6003] |
|---|-----------------------|-----------|

9.2. Ethernet Switching Granularities

IANA maintains a registry of GMPLS parameters called "Generalized Multi-Protocol Label Switching (GMPLS) Signaling Parameters".

IANA has created a new sub-registry called "Ethernet Switching Granularities" to contain the values that may be carried in the Switching Granularity field of the Ethernet SENDER_TSPEC object.

Values are as follows:

| | |
|-----------|----------------------------------|
| 0-2 | See below. |
| 3-239 | Unassigned |
| 240-254 | Reserved for Vendor-Specific Use |
| 255 | Reserved |
| 256-65535 | Not assigned at this time |

The registration procedure is Standards Action.

Initial entries in this sub-registry are as follows:

| Value | Switching Granularity | Reference |
|-------|--|--------------------|
| 0 | Provided in signaling. | [RFC6003][RFC6004] |
| 1 | Ethernet Port (for port-based service) | [RFC6003] |
| 2 | Ethernet Frame (for EVC-based service) | [RFC6003] |
| 255 | Reserved | [RFC6003] |

9.3. Ethernet Sender TSpec TLVs

IANA maintains a registry of GMPLS parameters called "Generalized Multi-Protocol Label Switching (GMPLS) Signaling Parameters".

IANA has created a new sub-registry called "Ethernet Sender TSpec TLVs / Ethernet Flowspec TLVs" to contain the TLV type values for TLVs carried in the Ethernet SENDER_TSPEC object.

Values are as follows:

| | |
|-----------|----------------------------------|
| 0-3 | See below. |
| 4-239 | Unassigned |
| 240-254 | Reserved for Vendor-Specific Use |
| 255 | Reserved |
| 256-65535 | Not assigned at this time |

The registration procedure is Standards Action.

Initial entries in this sub-registry are as follows:

| Type | Description | Reference |
|------|---------------------------------|-----------|
| 0 | Reserved | [RFC6003] |
| 1 | Reserved | [RFC6003] |
| 2 | Ethernet Bandwidth Profile | [RFC6003] |
| 3 | Layer 2 Control Protocol (L2CP) | [RFC6003] |
| 255 | Reserved | [RFC6003] |

9.4. Ethernet Bandwidth Profiles

IANA maintains a registry of GMPLS parameters called "Generalized Multi-Protocol Label Switching (GMPLS) Signaling Parameters".

IANA has created a new sub-registry called "Ethernet Bandwidth Profiles" to contain bit flags carried in the Ethernet Bandwidth Profile TLV of the Ethernet SENDER_TSPEC object.

Bits are to be allocated by IETF Standards Action. Bits are numbered from bit 0 as the low order bit. Initial entries are as follows:

| Bit | Hex | Description | Reference |
|-----|------|--------------------|-----------|
| 0 | 0x01 | Coupling Flag (CF) | [RFC6003] |
| 1 | 0x02 | Color Mode (CM) | [RFC6003] |

10. Acknowledgments

Many thanks to Adrian Farrel for his comments. Lou Berger provided the input on control traffic processing.

11. References

11.1. Normative References

- [MEF10.1] The MEF Technical Specification, "Ethernet Services Attributes Phase 2", MEF 10.1, November 2006.
- [RFC2205] Braden, R., Ed., Zhang, L., Berson, S., Herzog, S., and S. Jamin, "Resource ReSerVation Protocol (RSVP) -- Version 1 Functional Specification", RFC 2205, September 1997.
- [RFC2210] Wroclawski, J., "The Use of RSVP with IETF Integrated Services", RFC 2210, September 1997.
- [RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, March 1997.
- [RFC3209] Awduche, D., Berger, L., Gan, D., Li, T., Srinivasan, V., and G. Swallow, "RSVP-TE: Extensions to RSVP for LSP Tunnels", RFC 3209, December 2001.
- [RFC3471] Berger, L., Ed., "Generalized Multi-Protocol Label Switching (GMPLS) Signaling Functional Description", RFC 3471, January 2003.

- [RFC3473] Berger, L., Ed., "Generalized Multi-Protocol Label Switching (GMPLS) Signaling Resource ReserVation Protocol-Traffic Engineering (RSVP-TE) Extensions", RFC 3473, January 2003.
- [RFC4124] Le Faucheur, F., Ed., "Protocol Extensions for Support of Diffserv-aware MPLS Traffic Engineering", RFC 4124, June 2005.
- [RFC4506] Eisler, M., Ed., "XDR: External Data Representation Standard", STD 67, RFC 4506, May 2006.
- [RFC5226] Narten, T. and H. Alvestrand, "Guidelines for Writing an IANA Considerations Section in RFCs", BCP 26, RFC 5226, May 2008.
- [RFC6004] Berger, L. and D. Fedyk, "Generalized MPLS (GMPLS) Support for Metro Ethernet Forum and G.8011 Ethernet Services", RFC 6004, October 2010.

11.2. Informative References

- [ETHv2] Digital, Intel, and Xerox, "The Ethernet -- A Local Area Network: Data Link Layer and Physical Layer Specifications", Version 2.0, November 1982.
- [IEEE802.3] IEEE 802.3 LAN/MAN CSMA/CD (Ethernet) Access Method, IEEE Standard for Information technology- Specific requirements - Part 3: Carrier Sense Multiple Access with Collision Detection (CMSA/CD) Access Method and Physical Layer Specifications, IEEE 802.3-2008.
- [MCOS] Minei, I., Gan, D., Kompella, K., and X. Li, "Extensions for Differentiated Services-aware Traffic Engineered LSPs", Work in Progress, June 2006.
- [MEF6] The Metro Ethernet Forum, "Ethernet Services Definitions - Phase I", MEF 6, June 2004.
- [MEF10.2] The MEF Technical Specification, "Ethernet Services Attributes Phase 2", MEF 10.2, October 2009.
- [RFC5920] Fang, L., Ed., "Security Framework for MPLS and GMPLS Networks", RFC 5920, July 2010.

Author's Address

Dimitri Papadimitriou
Alcatel-Lucent Bell
Copernicuslaan 50
B-2018 Antwerpen, Belgium
Phone: +32 3 2408491
EMail: dimitri.papadimitriou@alcatel-lucent.be