

Network Working Group
Request for Comments: 5467
Category: Experimental

L. Berger
LabN
A. Takacs
Ericsson
D. Caviglia
Ericsson
D. Fedyk
Nortel
J. Meuric
France Telecom
March 2009

GMPLS Asymmetric Bandwidth Bidirectional Label Switched Paths (LSPs)

Status of This Memo

This memo defines an Experimental Protocol for the Internet community. It does not specify an Internet standard of any kind. Discussion and suggestions for improvement are requested. Distribution of this memo is unlimited.

Copyright Notice

Copyright (c) 2009 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 in effect on the date of publication of this document (<http://trustee.ietf.org/license-info>). Please review these documents carefully, as they describe your rights and restrictions with respect to this document.

This document may contain material from IETF Documents or IETF Contributions published or made publicly available before November 10, 2008. The person(s) controlling the copyright in some of this material may not have granted the IETF Trust the right to allow modifications of such material outside the IETF Standards Process. Without obtaining an adequate license from the person(s) controlling the copyright in such materials, this document may not be modified outside the IETF Standards Process, and derivative works of it may not be created outside the IETF Standards Process, except to format it for publication as an RFC or to translate it into languages other than English.

Abstract

This document defines a method for the support of GMPLS asymmetric bandwidth bidirectional Label Switched Paths (LSPs). The presented approach is applicable to any switching technology and builds on the original Resource Reservation Protocol (RSVP) model for the transport of traffic-related parameters. The procedures described in this document are experimental.

Table of Contents

1. Introduction	2
1.1. Background	3
1.2. Approach Overview	3
1.3. Conventions Used in This Document	4
2. Generalized Asymmetric Bandwidth Bidirectional LSPs	4
2.1. UPSTREAM_FLOWSPEC Object	5
2.1.1. Procedures	5
2.2. UPSTREAM_TSPEC Object	5
2.2.1. Procedures	5
2.3. UPSTREAM_ADSPEC Object	6
2.3.1. Procedures	6
3. Packet Formats	6
4. Compatibility	7
5. IANA Considerations	8
5.1. UPSTREAM_FLOWSPEC Object	8
5.2. UPSTREAM_TSPEC Object	8
5.3. UPSTREAM_ADSPEC Object	8
6. Security Considerations	8
7. References	9
7.1. Normative References	9
7.2. Informative References	9
Appendix A. Alternate Approach Using ADSPEC Object.....	11
A.1. Applicability	11
A.2. Overview	11
A.3. Procedures	12
A.4. Compatibility	13

1. Introduction

GMPLS [RFC3473] introduced explicit support for bidirectional Label Switched Paths (LSPs). The defined support matched the switching technologies covered by GMPLS, notably Time Division Multiplexing (TDM) and lambdas; specifically, it only supported bidirectional LSPs with symmetric bandwidth allocation. Symmetric bandwidth requirements are conveyed using the semantics objects defined in [RFC2205] and [RFC2210].

Recent work ([GMPLS-PBBTE] and [MEF-TRAFFIC]) has looked at extending GMPLS to control Ethernet switching. In this context, there has been discussion of the support of bidirectional LSPs with asymmetric bandwidth. (That is, bidirectional LSPs that have different bandwidth reservations in each direction.) This discussion motivated the extensions defined in this document, which may be used with any switching technology to signal asymmetric bandwidth bidirectional LSPs. The procedures described in this document are experimental.

1.1. Background

Bandwidth parameters are transported within RSVP ([RFC2210], [RFC3209], and [RFC3473]) via several objects that are opaque to RSVP. While opaque to RSVP, these objects support a particular model for the communication of bandwidth information between an RSVP session sender (ingress) and receiver (egress). The original model of communication, defined in [RFC2205] and maintained in [RFC3209], used the SENDER_TSPEC and ADSPEC objects in Path messages and the FLOWSPEC object in Resv messages. The SENDER_TSPEC object was used to indicate a sender's data generation capabilities. The FLOWSPEC object was issued by the receiver and indicated the resources that should be allocated to the associated data traffic. The ADSPEC object was used to inform the receiver and intermediate hops of the actual resources allocated for the associated data traffic.

With the introduction of bidirectional LSPs in [RFC3473], the model of communication of bandwidth parameters was implicitly changed. In the context of [RFC3473] bidirectional LSPs, the SENDER_TSPEC object indicates the desired resources for both upstream and downstream directions. The FLOWSPEC object is simply confirmation of the allocated resources. The definition of the ADSPEC object is either unmodified and only has meaning for downstream traffic, or is implicitly or explicitly ([RFC4606] and [MEF-TRAFFIC]) irrelevant.

1.2. Approach Overview

The approach for supporting asymmetric bandwidth bidirectional LSPs defined in this document builds on the original RSVP model for the transport of traffic-related parameters and GMPLS's support for bidirectional LSPs. An alternative approach was considered and rejected in favor of the more generic approach presented below. For reference purposes only, the rejected approach is summarized in Appendix A.

The defined approach is generic and can be applied to any switching technology supported by GMPLS. With this approach, the existing SENDER_TSPEC, ADSPEC, and FLOWSPEC objects are complemented with the addition of new UPSTREAM_TSPEC, UPSTREAM_ADSPEC, and

UPSTREAM_FLOWSPEC objects. The existing objects are used in the original fashion defined in [RFC2205] and [RFC2210], and refer only to traffic associated with the LSP flowing in the downstream direction. The new objects are used in exactly the same fashion as the old objects, but refer to the upstream traffic flow. Figure 1 shows the bandwidth-related objects used for asymmetric bandwidth bidirectional LSPs.




Figure 1: Generic Asymmetric Bandwidth Bidirectional LSPs

The extensions defined in this document are limited to Point-to-Point (P2P) LSPs. Support for Point-to-Multipoint (P2MP) bidirectional LSPs is not currently defined and, as such, not covered in this document.

1.3. 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 [RFC2119].

2. Generalized Asymmetric Bandwidth Bidirectional LSPs

The setup of an asymmetric bandwidth bidirectional LSP is signaled using the bidirectional procedures defined in [RFC3473] together with the inclusion of the new UPSTREAM_FLOWSPEC, UPSTREAM_TSPEC, and UPSTREAM_ADSPEC objects.

The new upstream objects carry the same information and are used in the same fashion as the existing downstream objects; they differ in that they relate to traffic flowing in the upstream direction while the existing objects relate to traffic flowing in the downstream direction. The new objects also differ in that they are used on messages in the opposite directions.

2.1. UPSTREAM_FLOWSPEC Object

The format of an UPSTREAM_FLOWSPEC object is the same as a FLOWSPEC object. This includes the definition of class types and their formats. The class number of the UPSTREAM_FLOWSPEC object is 120 (of the form 0bbbbbbb).

2.1.1. Procedures

The Path message of an asymmetric bandwidth bidirectional LSP MUST contain an UPSTREAM_FLOWSPEC object and MUST use the bidirectional LSP formats and procedures defined in [RFC3473]. The C-Type of the UPSTREAM_FLOWSPEC object MUST match the C-Type of the SENDER_TSPEC object used in the Path message. The contents of the UPSTREAM_FLOWSPEC object MUST be constructed using a format and procedures consistent with those used to construct the FLOWSPEC object that will be used for the LSP, e.g., [RFC2210] or [RFC4328].

Nodes processing a Path message containing an UPSTREAM_FLOWSPEC object MUST use the contents of the UPSTREAM_FLOWSPEC object in the upstream label and the resource allocation procedure defined in Section 3.1 of [RFC3473]. Consistent with [RFC3473], a node that is unable to allocate a label or internal resources based on the contents of the UPSTREAM_FLOWSPEC object MUST issue a PathErr message with a "Routing problem/MPLS label allocation failure" indication.

2.2. UPSTREAM_TSPEC Object

The format of an UPSTREAM_TSPEC object is the same as a SENDER_TSPEC object. This includes the definition of class types and their formats. The class number of the UPSTREAM_TSPEC object is 121 (of the form 0bbbbbbb).

2.2.1. Procedures

The UPSTREAM_TSPEC object describes the traffic flow that originates at the egress. The UPSTREAM_TSPEC object MUST be included in any Resv message that corresponds to a Path message containing an UPSTREAM_FLOWSPEC object. The C-Type of the UPSTREAM_TSPEC object MUST match the C-Type of the corresponding UPSTREAM_FLOWSPEC object. The contents of the UPSTREAM_TSPEC object MUST be constructed using a format and procedures consistent with those used to construct the FLOWSPEC object that will be used for the LSP, e.g., [RFC2210] or [RFC4328]. The contents of the UPSTREAM_TSPEC object MAY differ from contents of the UPSTREAM_FLOWSPEC object based on application data transmission requirements.

When an UPSTREAM_TSPEC object is received by an ingress, the ingress MAY determine that the original reservation is insufficient to satisfy the traffic flow. In this case, the ingress MAY issue a Path message with an updated UPSTREAM_FLOWSPEC object to modify the resources requested for the upstream traffic flow. This modification might require the LSP to be re-routed, and in extreme cases might result in the LSP being torn down when sufficient resources are not available.

2.3. UPSTREAM_ADSPEC Object

The format of an UPSTREAM_ADSPEC object is the same as an ADSPEC object. This includes the definition of class types and their formats. The class number of the UPSTREAM_ADSPEC object is 122 (of the form 0bbbbbbb).

2.3.1. Procedures

The UPSTREAM_ADSPEC object MAY be included in any Resv message that corresponds to a Path message containing an UPSTREAM_FLOWSPEC object. The C-Type of the UPSTREAM_TSPEC object MUST be consistent with the C-Type of the corresponding UPSTREAM_FLOWSPEC object. The contents of the UPSTREAM_ADSPEC object MUST be constructed using a format and procedures consistent with those used to construct the ADSPEC object that will be used for the LSP, e.g., [RFC2210] or [MEF-TRAFFIC]. The UPSTREAM_ADSPEC object is processed using the same procedures as the ADSPEC object and, as such, MAY be updated or added at transit nodes.

3. Packet Formats

This section presents the RSVP message-related formats as modified by this section. This document modifies formats defined in [RFC2205], [RFC3209], and [RFC3473]. See [RSVP-BNF] for the syntax used by RSVP. Unmodified formats are not listed. Three new objects are defined in this section:

Object name	Applicable RSVP messages
-----	-----
UPSTREAM_FLOWSPEC	Path, PathTear, PathErr, and Notify (via sender descriptor)
UPSTREAM_TSPEC	Resv, ResvConf, ResvTear, ResvErr, and Notify (via flow descriptor list)
UPSTREAM_ADSPEC	Resv, ResvConf, ResvTear, ResvErr, and Notify (via flow descriptor list)

The format of the sender description for bidirectional asymmetric LSPs is:

```
<sender descriptor> ::= <SENDER_TEMPLATE> <SENDER_TSPEC>
                        [ <ADSPEC> ]
                        [ <RECORD_ROUTE> ]
                        [ <SUGGESTED_LABEL> ]
                        [ <RECOVERY_LABEL> ]
                        <UPSTREAM_LABEL>
                        <UPSTREAM_FLOWSPEC>
```

The format of the flow descriptor list for bidirectional asymmetric LSPs is:

```
<flow descriptor list> ::= <FF flow descriptor list>
                        | <SE flow descriptor>

<FF flow descriptor list> ::= <FLOWSPEC>
                        <UPSTREAM_TSPEC> [ <UPSTREAM_ADSPEC> ]
                        <FILTER_SPEC>
                        <LABEL> [ <RECORD_ROUTE> ]
                        | <FF flow descriptor list>
                        <FF flow descriptor>

<FF flow descriptor> ::= [ <FLOWSPEC> ]
                        [ <UPSTREAM_TSPEC> ] [ <UPSTREAM_ADSPEC> ]
                        <FILTER_SPEC> <LABEL>
                        [ <RECORD_ROUTE> ]

<SE flow descriptor> ::= <FLOWSPEC>
                        <UPSTREAM_TSPEC> [ <UPSTREAM_ADSPEC> ]
                        <SE filter spec list>
```

<SE filter spec list> is unmodified by this document.

4. Compatibility

This extension reuses and extends semantics and procedures defined in [RFC2205], [RFC3209], and [RFC3473] to support bidirectional LSPs with asymmetric bandwidth. To indicate the use of asymmetric bandwidth, three new objects are defined. Each of these objects is defined with class numbers in the form 0bbbbbbb. Per [RFC2205], nodes not supporting this extension will not recognize the new class numbers and should respond with an "Unknown Object Class" error. The error message will propagate to the ingress, which can then take action to avoid the path with the incompatible node or may simply terminate the session.

5. IANA Considerations

IANA has assigned new values for namespaces defined in this section and reviewed in this subsection.

The IANA has made the assignments described below in the "Class Names, Class Numbers, and Class Types" section of the "RSVP PARAMETERS" registry.

5.1. UPSTREAM_FLOWSPEC Object

A new class named UPSTREAM_FLOWSPEC has been created in the 0bbbbbbb range (120) with the following definition:

Class Types or C-types:

Same values as FLOWSPEC object (C-Num 9)

5.2. UPSTREAM_TSPEC Object

A new class named UPSTREAM_TSPEC has been created in the 0bbbbbbb range (121) with the following definition:

Class Types or C-types:

Same values as SENDER_TSPEC object (C-Num 12)

5.3. UPSTREAM_ADSPEC Object

A new class named UPSTREAM_ADSPEC has been created in the 0bbbbbbb range (122) with the following definition:

Class Types or C-types:

Same values as ADSPEC object (C-Num 13)

6. Security Considerations

This document introduces new message objects for use in GMPLS signaling [RFC3473] -- specifically the UPSTREAM_TSPEC, UPSTREAM_ADSPEC, and UPSTREAM_FLOWSPEC objects. These objects parallel the exiting SENDER_TSPEC, ADSPEC, and FLOWSPEC objects but are used in the opposite direction. As such, any vulnerabilities that are due to the use of the old objects now apply to messages flowing in the reverse direction.

From a message standpoint, this document does not introduce any new signaling messages or change the relationship between LSRs that are adjacent in the control plane. As such, this document introduces no additional message- or neighbor-related security considerations.

See [RFC3473] for relevant security considerations, and [SEC-FRAMEWORK] for a more general discussion on RSVP-TE security discussions.

7. References

7.1. Normative References

- [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.
- [RFC3473] Berger, L., Ed., "Generalized Multi-Protocol Label Switching (GMPLS) Signaling Resource Reservation Protocol-Traffic Engineering (RSVP-TE) Extensions", RFC 3473, January 2003.

7.2. Informative References

- [GMPLS-PBBTE] Fedyk, D., et al "GMPLS Control of Ethernet", Work in Progress, July 2008.
- [MEF-TRAFFIC] Papadimitriou, D., "MEF Ethernet Traffic Parameters," Work in Progress, October 2008.
- [RFC4606] Mannie, E. and D. Papadimitriou, "Generalized Multi-Protocol Label Switching (GMPLS) Extensions for Synchronous Optical Network (SONET) and Synchronous Digital Hierarchy (SDH) Control", RFC 4606, August 2006.

- [RFC4328] Papadimitriou, D., Ed., "Generalized Multi-Protocol Label Switching (GMPLS) Signaling Extensions for G.709 Optical Transport Networks Control", RFC 4328, January 2006.
- [RSVP-BNF] Farrel, A. "Reduced Backus-Naur Form (RBNF) A Syntax Used in Various Protocol Specifications", Work in Progress, November 2008.
- [SEC-FRAMEWORK] Fang, L., Ed., "Security Framework for MPLS and GMPLS Networks", Work in Progress, November 2008.

A. Appendix A: Alternate Approach Using ADSPEC Object

This section is included for historic purposes and its implementation is NOT RECOMMENDED.

A.1. Applicability

This section presents an alternate method for the support of asymmetric bandwidth bidirectional LSP establishment with a single RSVP-TE signaling session. This approach differs in applicability and generality from the approach presented in the main body of this document. In particular, this approach is technology-specific; it uses the ADSPEC object to carry traffic parameters for upstream data and requires the Metro Ethernet Forum (MEF) Ethernet Traffic Parameter, while the approach presented above is suitable for use with any technology.

The generalized asymmetric bandwidth bidirectional LSP presented in the main body of this document has the benefit of being applicable to any switching technology, but requires support for three new types of object classes, i.e., the UPSTREAM_TSPEC, UPSTREAM_ADSPEC, and UPSTREAM_FLOWSPEC objects.

The solution presented in this section is based on the Ethernet-specific ADSPEC object, and is referred to as the "ADSPEC Object" approach. This approach limits applicability to cases where the [MEF-TRAFFIC] traffic parameters are appropriate, and to switching technologies that define no use for the ADSPEC object. While ultimately it is this limited scope that has resulted in this approach being relegated to an Appendix, the semantics of this approach are quite simple in that they only require the definition of a new ADSPEC object C-Type.

In summary, the "ADSPEC Object" approach presented in this section SHOULD NOT be implemented.

A.2. Overview

The "ADSPEC Object" approach is specific to Ethernet and uses [MEF-TRAFFIC] traffic parameters. This approach is not generic and is aimed at providing asymmetric bandwidth bidirectional LSPs for just Ethernet transport. With this approach, the ADSPEC object carries the traffic parameters for the upstream data flow. SENDER_TSPEC object is used to indicate the traffic parameters for the downstream data flow. The FLOWSPEC object provides confirmation of the allocated downstream resources. Confirmation of the upstream resource allocation is a Resv message, as any resource allocation

failure for the upstream direction will always result in a PathErr message. Figure 2 shows the bandwidth-related objects used in the first approach.




Figure 2: Asymmetric Bandwidth Bidirectional LSPs Using ADSPEC Object

In the "ADSPEC Object" approach, the setup of an asymmetric bandwidth bidirectional LSP would be signaled using the bidirectional procedures defined in [RFC3473] together with the inclusion of a new ADSPEC object. The new ADSPEC object would be specific to Ethernet and could be called the Ethernet Upstream Traffic Parameter ADSPEC object. The Ethernet Upstream Traffic Parameter ADSPEC object would use the Class-Number 13 and C-Type UNASSIGNED (this approach should not be implemented). The format of the object would be the same as the Ethernet SENDER_TSPEC object defined in [MEF-TRAFFIC].

This approach would not modify behavior of symmetric bandwidth LSPs. Per [MEF-TRAFFIC], such LSPs are signaled either without an ADSPEC or with an INTSERV ADSPEC.

The defined approach could be reused to support asymmetric bandwidth bidirectional LSPs for other types of switching technologies. All that would be needed would be to define the proper ADSPEC object.

A.3. Procedures

Using the approach presented in this section, the process of establishing an asymmetric bandwidth bidirectional LSP would follow the process of establishing a symmetric bandwidth bidirectional LSP, as defined in Section 3 of [RFC3473], with two modifications. These modifications would be followed when an incoming Path message is received containing an Upstream_Label object and the Ethernet Upstream Traffic Parameter ADSPEC object.

The first modification to the symmetric bandwidth process would be that when allocating the upstream label, the bandwidth associated with the upstream label would be taken from the Ethernet Upstream Traffic Parameter ADSPEC object, see Section 3.1 of [RFC3473].

Consistent with [RFC3473], a node that is unable to allocate a label or internal resources based on the contents of the ADSPEC object, would issue a PathErr message with a "Routing problem/MPLS label allocation failure" indication.

The second modification would be that the ADSPEC object would not be modified by transit nodes.

A.4. Compatibility

The approach presented in this section reuses semantics and procedures defined in [RFC3473]. To indicate the use of asymmetric bandwidth, a new ADSPEC object C-type would be defined. Per [RFC2205], nodes not supporting the approach should not recognize this new C-type and respond with an "Unknown object C-Type" error.

Authors' Addresses

Lou Berger
LabN Consulting, L.L.C.

EMail: lberger@labn.net

Attila Takacs
Ericsson
1. Laborc u.
1037 Budapest, Hungary

Phone: +36-1-4377044
EMail: attila.takacs@ericsson.com

Diego Caviglia
Ericsson
Via A. Negrone 1/A
Genova-Sestri Ponente, Italy

Phone: +390106003738
EMail: diego.caviglia@ericsson.com

Don Fedyk
Nortel Networks
600 Technology Park Drive
Billerica, MA, USA 01821

Phone: +1-978-288-3041
EMail: dwfedyk@nortel.com

Julien Meuric
France Telecom
Research & Development
2, avenue Pierre Marzin
22307 Lannion Cedex - France

Phone: +33 2 96 05 28 28
EMail: julien.meuric@orange-ftgroup.com

