Network Working Group Peter Ashwood-Smith (Nortel Networks Corp.)
Internet Draft Ayan Banerjee (Calient Networks)
Expiration Date: January 2002 Lou Berger (Movaz Networks)
Greg Bernstein (Ciena Corporation)
John Drake (Calient Networks)
Yanhe Fan (Axiowave Networks)
Don Fedyk (Nortel Networks Corp.)
Kireeti Kompella (Juniper Networks, Inc.)
Eric Mannie (EBONE)
Jonathan P. Lang (Calient Networks)
Bala Rajagopalan (Tellium, Inc.)
Yakov Rekhter (Juniper Networks, Inc.)
Debanjan Saha (Tellium, Inc.)
Vishal Sharma (Jasmine Networks)
George Swallow (Cisco Systems)
Z. Bo Tang (Tellium, Inc.)

July 2001

Generalized MPLS Signaling - CR-LDP Extensions

draft-ietf-mpls-generalized-cr-ldp-04.txt

Status of this Memo

Abstract

This document describes extensions to CR-LDP signaling required to support Generalized MPLS. Generalized MPLS extends MPLS to encompass time-division (e.g. SONET ADMs), wavelength (optical lambdas) and spatial switching (e.g. incoming port or fiber to outgoing port or fiber). This document presents a CR-LDP specific description of the extensions. An RSVP-TE specific description can be found in [GMPLS-RSVP]. A generic functional description is presented in [GMPLS-SIG].

Contents

Changes from previous version:

- o Fixed Label Set format (for LDP)
- o Removed unassigned values
- o Added Switching type of LSP being requested
- o Added Administrative Status Information (based on last call comments)
- o Added section on Control Channel Separation
  (based on last call comments)
  Covers:
  - Separation of control and data channels

1. Introduction

Generalized MPLS extends MPLS from supporting packet (PSC) interfaces and switching to include support of three new classes of interfaces and switching: Time-Division Multiplex (TDM), Lambda Switch (LSC) and Fiber-Switch (FSC). A functional description of the extensions to MPLS signaling needed to support the new classes of interfaces and switching is provided in [GMPLS-SIG]. This document presents CR-LDP specific formats and mechanisms needed to support all four classes of interfaces. RSVP-TE extensions can be found in [GMPLS-RSVP].

[GMPLS-SIG] should be viewed as a companion document to this document. The format of this document parallels [GMPLS-SIG]. It should be noted that the RSVP-TE specific version of Generalized MPLS includes RSVP specific support for rapid failure notification, see Section 4 [GMPLS-RSVP]. For CR-LDP there is not currently a similar mechanism. When a failure is detected it will be propagated with RELEASE/WITHDRAW messages radially outward from the point of failure. Resources are to be released in this phase and actual resource information may be fed back to the source using a feedback mechanisms.

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. Label Related Formats

This section defines formats for a generalized label request, a generalized label, support for waveband switching, suggested label and label sets.

2.1. Generalized Label Request

A REQUEST message SHOULD contain as specific an LSP Encoding Type as possible to allow the maximum flexibility in switching by transit LSRs. A Generalized Label Request TLV is set by the ingress node, transparently passed by transit nodes, and used by the egress node.

The format of a Generalized Label Request is:

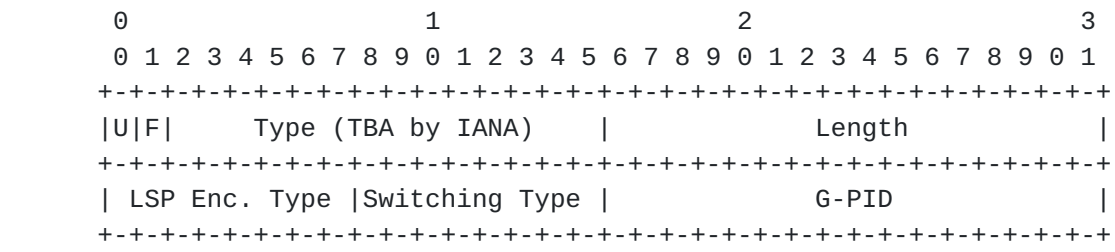

See [GMPLS-SIG] for a description of parameters.

2.1.1. Procedures

A node processing a REQUEST message containing a Generalized Label Request must verify that the requested parameters can be satisfied by the incoming interface, the node and by the outgoing interface. The node may either directly support the LSP or it may use a tunnel (FA), i.e., another class of switching. In either case, each parameter must be checked.

Note that local node policy dictates when tunnels may be used and when they may be created. Local policy may allow for tunnels to be dynamically established or may be solely administratively controlled. For more information on tunnels and processing of ER hops when using tunnels see [MPLS-HIERARCHY].

Transit and egress nodes MUST verify that the node itself and, where appropriate, that the outgoing interface or tunnel can support the requested LSP Encoding Type. If encoding cannot be supported, the node MUST generate a NOTIFICATION message, with a "Routing problem/Unsupported Encoding" indication.

The G-PID parameter is normally only examined at the egress. If the indicated G-PID cannot be supported then the egress MUST generate a NOTIFICATION message, with a "Routing problem/Unsupported GPID" indication. In the case of PSC and when penultimate hop popping (PHP) is requested, the penultimate hop also examines the (stored) G-PID during the processing of the MAPPING message. In this case if the G-PID is not supported, then the penultimate hop MUST generate a

NOTIFICATION message with a "Routing problem/Unacceptable label value" indication. The generated NOTIFICATION message MAY include an Acceptable Label Set, see Section 4.

When an error message is not generated, normal processing occurs. In the transit case this will typically result in a REQUEST message being propagated. In the egress case and PHP special case this will typically result in a MAPPING message being generated.

2.1.2. Bandwidth Encoding

Bandwidth encodings are carried in the CR-LDP Traffic Parameters TLV. See [GMPLS-SIG] for a definition of values to be used for specific signal types. These values are set in the Peak and Committed Data Rate fields of the Traffic Parameters TLV. Other bandwidth/service related parameters in the TLV are ignored and carried transparently.

2.2. Generalized Label

The format of a Generalized Label is:

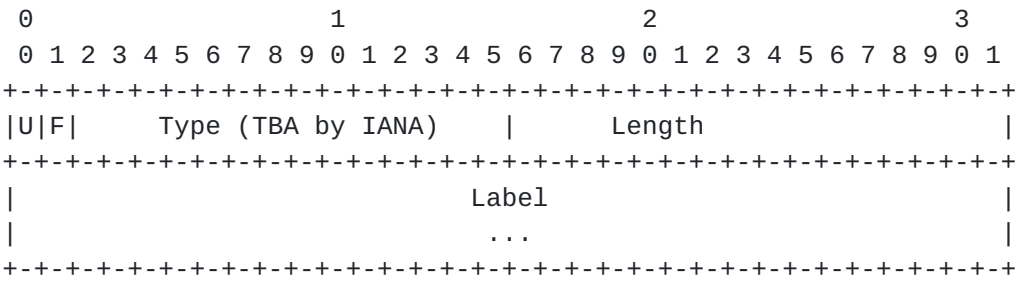

See [GMPLS-SIG] for a description of parameters and encoding of labels.

2.2.1. Procedures

The Generalized Label travels in the upstream direction in MAPPING messages.

The presence of both a generalized and normal label TLV in a MAPPING message is a protocol error and should treated as a malformed message by the recipient.

The recipient of a MAPPING message containing a Generalized Label

verifies that the values passed are acceptable. If the label is unacceptable then the recipient MUST generate a NOTIFICATION message with a "Routing problem/MPLS label allocation failure" indication. The generated NOTIFICATION message MAY include an Acceptable Label Set, see Section 4.

2.3. Waveband Switching

Waveband switching uses the same format as the generalized label, see section 2.2. The type (0x0903) is assigned for the Waveband Label.

In the context of waveband switching, the generalized label has the following format:

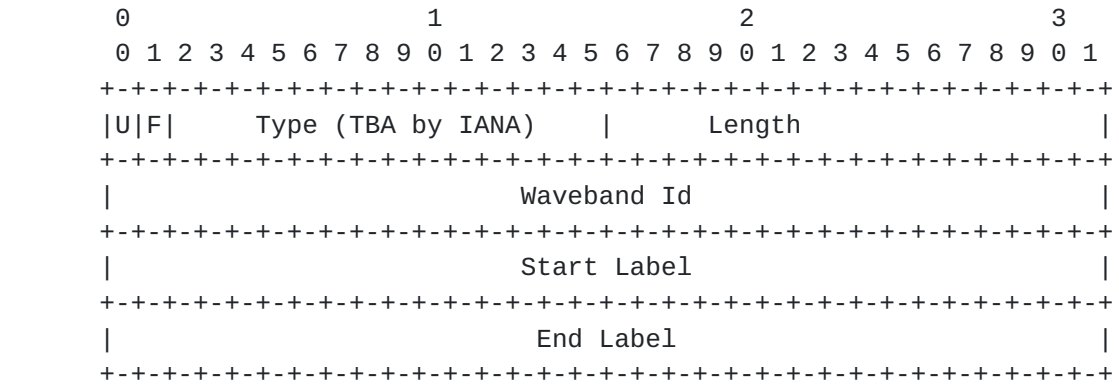

See [GMPLS-SIG] for a description of parameters.

2.3.1. Procedures

The procedures defined in Section 2.2.1 apply to waveband switching. This includes generating a NOTIFICATION message with a "Routing problem/MPLS label allocation failure" indication if any of the label fields are unrecognized or unacceptable.

Additionally, when a waveband is switched to another waveband, it is possible that the wavelengths within the waveband will be mirrored about a center frequency. When this type of switching is employed, the start and end label in the waveband label TLV MUST be flipped before forwarding the label TLV with the new waveband Id. In this manner an egress/ingress LSR that receives a waveband label which has these values inverted, knows that it must also invert its egress association to pick up the proper wavelengths. Without this mechanism and with an odd number of mirrored switching operations, the egress LSRs will not know that an input wavelength of say L1 will emerge from the waveband tunnel as L100.

This operation MUST be performed in both directions when a bidirectional waveband tunnel is being established.

2.4. Suggested Label

The format of a suggested label is identical to a generalized label. It is used in REQUEST messages. Suggested Label uses type = 0x904.

Errors in received Suggested Labels MUST be ignored. This includes any received inconsistent or unacceptable values.

2.5. Label Set

The format of a Label_Set is:

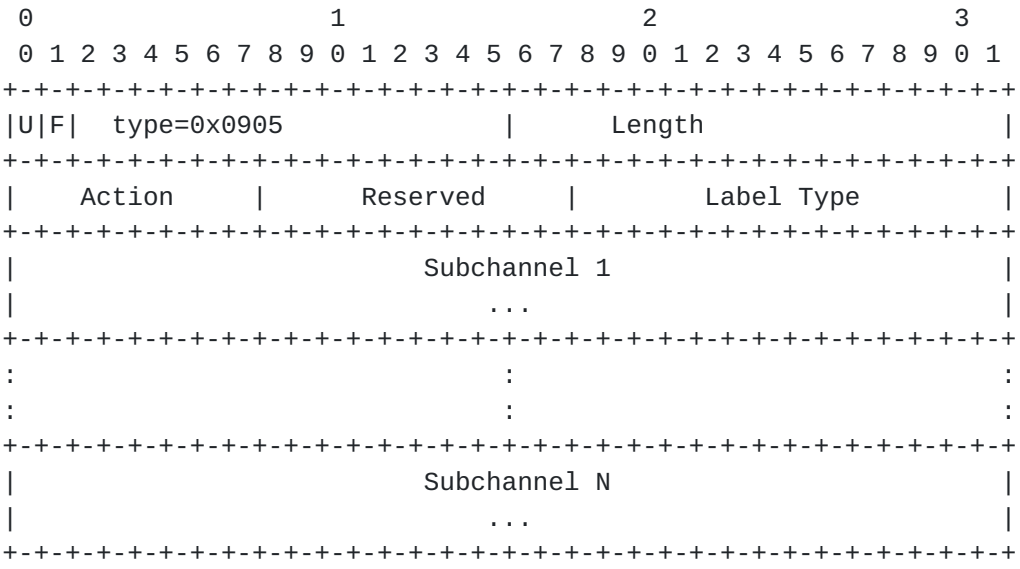

Label Type: 14 bits

Indicates the type and format of the labels carried in the TLV. Values match the TLV type of the appropriate Label TLV.

See [GMPLS-SIG] for a description of other parameters.

2.5.1. Procedures

A Label Set is defined via one or more Label_Set TLVs. Specific labels/subchannels can be added to or excluded from a Label Set via Action zero (0) and one (1) TLVs respectively. Ranges of labels/subchannels can be added to or excluded from a Label Set via Action two (2) and three (3) TLVs respectively. When the Label_Set TLVs only list labels/subchannels to exclude, this implies that all other labels are acceptable.

The absence of any Label_Set TLVs implies that all labels are acceptable. A Label Set is included when a node wishes to restrict the label(s) that may be used downstream.

On reception of a REQUEST message, the receiving node will restrict its choice of labels to one, which is in the Label Set. Nodes capable of performing label conversion may also remove the Label Set prior to forwarding the REQUEST message. If the node is unable to pick a label from the Label Set or if there is a problem parsing the Label_Set TLVs, then the request is terminated and a NOTIFICATION message with a "Routing problem/Label Set" indication MUST be generated. It is a local matter if the Label Set is stored for later selection on the MAPPING or if the selection is made immediately for propagation in the MAPPING.

On reception of a REQUEST message, the Label Set represented in the message is compared against the set of available labels at the downstream interface and the resulting intersecting Label Set is forwarded in a REQUEST message. When the resulting Label Set is empty, the REQUEST must be terminated, and a NOTIFICATION message, and a "Routing problem/Label Set" indication MUST be generated. Note that intersection is based on the physical labels (actual wavelength/band values) which may have different logical values on different links, as a result it is the responsibility of the node to map these values so that they have a consistent physical meaning, or to drop the particular values from the set if no suitable logical label value exists.

When processing a MAPPING message at an intermediate node, the label propagated upstream MUST fall within the Label Set.

Note, on reception of a MAPPING message a node that is incapable of performing label conversion has no other choice than to use the same physical label (wavelength/band) as received in the MAPPING message. In this case, the use and propagation of a Label Set will significantly reduce the chances that this allocation will fail.

3. Bidirectional LSPs

Bidirectional LSP setup is indicated by the presence of an Upstream Label in the REQUEST message. An Upstream Label has the same format as the generalized label, see Section 2.2. Upstream Label uses type=0x0906

3.1. Procedures

The process of establishing a bidirectional LSP follows the establishment of a unidirectional LSP with some additions. To support bidirectional LSPs an Upstream Label is added to the REQUEST message. The Upstream Label MUST indicate a label that is valid for forwarding at the time the REQUEST message is sent.

When a REQUEST message containing an Upstream Label is received, the receiver first verifies that the upstream label is acceptable. If the label is not acceptable, the receiver MUST issue a NOTIFICATION message with a "Routing problem/Unacceptable label value" indication. The generated NOTIFICATION message MAY include an Acceptable Label Set, see Section 4.

An intermediate node must also allocate a label on the outgoing interface and establish internal data paths before filling in an outgoing Upstream Label and propagating the REQUEST message. If an intermediate node is unable to allocate a label or internal resources, then it MUST issue a NOTIFICATION message with a "Routing problem/Label allocation failure" indication.

Terminator nodes process REQUEST messages as usual, with the exception that the upstream label can immediately be used to transport data traffic associated with the LSP upstream towards the initiator.

When a bidirectional LSP is removed, both upstream and downstream labels are invalidated and it is no longer valid to send data using the associated labels.

4. Notification on Label Error

This section defines the Acceptable_Label_Set TLV to support Notification on Label Error per [GMPLS-SIG]. An Acceptable_Label_Set TLV uses a type value of 0x0907. The remaining contents of the TLV have the identical format as the Label_Set TLV, see Section 2.5.

Acceptable_Label_Set TLVs may be carried in NOTIFICATION messages. The procedures for defining an Acceptable Label Set follow the procedures for defining a Label Set, see Section 2.5.1. Specifically, an Acceptable Label Set is defined via one or more Acceptable_Label_Set TLVs. Specific labels/subchannels can be added to or excluded from an Acceptable Label Set via Action zero (0) and one (1) TLVs respectively. Ranges of labels/subchannels can be added to or excluded from an Acceptable Label Set via Action two (2) and three (3) TLVs respectively. When the Acceptable_Label_Set TLVs only list labels/subchannels to exclude, this implies that all other labels are acceptable.

The inclusion of Acceptable_Label_Set TLVs is optional. If included, the NOTIFICATION message SHOULD contain a "Routing problem/Unacceptable label value" indication. The absence of Acceptable_Label_Set TLVs does not have any specific meaning.

5. Explicit Label Control

The Label ER-Hop is defined as follows:

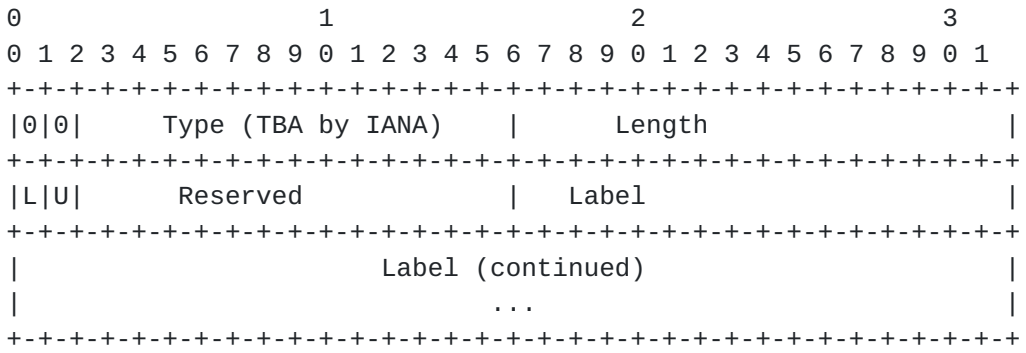

See [GMPLS-SIG] for a description of L, U and Label parameters.

5.1. Procedures

The Label ER-Hop follows a ER-Hop containing the IP address, or the interface identifier [MPLS-UNNUM], associated with the link on which it is to be used. The preceding ER-Hop must be strict. Up to two label ER-Hops may be present, one for the downstream label and one for the upstream label. The following SHOULD result in "Bad EXPLICIT_ROUTE" errors:

- If the first label ER-Hop is not preceded by a ER-Hop containing an IP address, or a interface identifier [MPLS-UNNUM], associated with an output link.
- For a label ER-Hop to follow a ER-Hop that has the L-bit set
- On unidirectional LSP setup, for there to be a label ER-Hop with the U-bit set
- For there to be two label ER-Hops with the same U-bit values

To support the label ER-Hop, a node must check to see if the ER-Hop following its associate address/interface is a label ER-Hop. If it is, one ER-Hop is examined for unidirectional LSPs and two ER-Hops for bidirectional LSPs. If the U-bit of the ER-Hop being examined is clear (0), then value of the label is copied into a new Label_Set TLV. This Label_Set TLV MUST be included on the corresponding outgoing MAPPING message.

If the U-bit of the ER-Hop being examined is set (1), then value of the label is label to be used for upstream traffic associated with the bidirectional LSP. If this label is not acceptable, a "Bad EXPLICIT_ROUTE" error SHOULD be generated. If the label is acceptable, the label is copied into a new Upstream Label TLV. This Upstream Label TLV MUST be included on the corresponding outgoing MAPPING message.

After processing, the label ER-Hops are removed from the ER.

Note an implication of the above procedures is that the label ER-Hop should never be the first ER-Hop in a newly received message. If the label ER-Hop is the first ER-Hop an a received ER, then it SHOULD be treated as a "Bad strict node" error.

Procedures by which an LSR at the head-end of an LSP obtains the information needed to construct the Label ER-Hop are outside the scope of this document.

6. Protection TLV

The use of the Protection TLV is optional. The TLV is included to indicate specific protection attributes of an LSP.

The format of Protection Information TLV is:

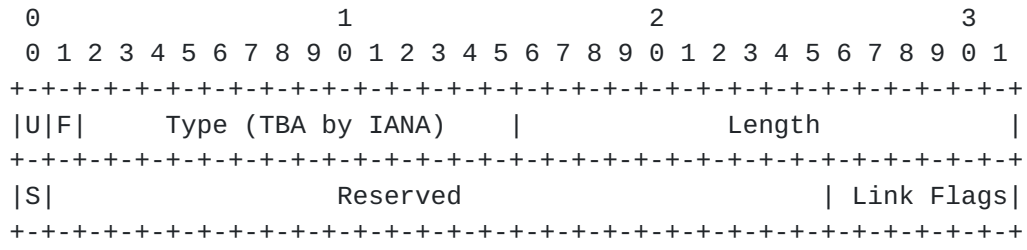

See [GMPLS-SIG] for a description of parameters.

6.1. Procedures

Transit nodes processing a REQUEST message containing a Protection TLV MUST verify that the requested protection can be satisfied by the outgoing interface or tunnel (FA). If it cannot, the node MUST generate a NOTIFICATION message, with a "Routing problem/Unsupported Link Protection" indication.

7. Administrative Status Information

Administrative Status Information is carried in the Admin Status TLV. The TLV provides information related to the administrative state of a particular LSP. The information is used in two ways. In the first, the object is carried in REQUEST and MAPPING messages to indicate the administrative state of an LSP. In the second, the TLV is carried in a REQUEST and MAPPING message with the modification bit set to request a change to the administrative state of an LSP.

7.1. Admin Status Object

The use of the Admin Status TLV is optional.

The format of the TLV is:

The format of Admin Status TLV in REQUEST, MAPPING Messages is:

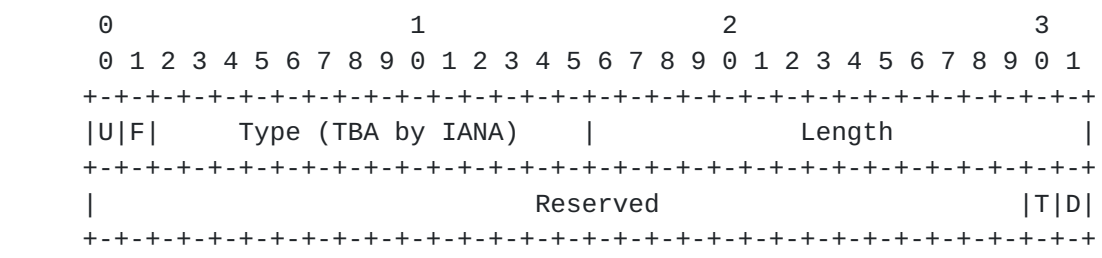

The U Bit Should be set to 1. The F bit should be set to 1.

See [GMPLS-SIG] for a description of parameters.

7.2. REQUEST and MAPPING Message Procedures

The Admin Status TLV is used to notify each node along the path of the status of the LSP. Status information is processed by each node based on local policy and then propagated in the corresponding outgoing messages. The TLV is inserted in REQUEST messages at the discretion of the ingress node.

Transit nodes receiving a REQUEST message containing an Admin Status TLV, update their local state, take any appropriate local action based on the indicated status and then propagate the received Admin Status TLV in the outgoing REQUEST message.

Egress nodes receiving a REQUEST message containing an Admin Status TLV, also update their local state and take any appropriate local action based on the indicated status.

The subsequent MAPPING message MUST carry back the Admin Status TLV if the corresponding request message had the Admin Status TLV.

7.3. Modification Message Procedures

Subsequent messaging Admin Status messaging is all performed by REQUEST and MAPPING Messages using the modification action indicator flag. The ingress may begin the propagation of a Message with an Admin Status TLV. Each subsequent node propagates the REQUEST with the Admin Status TLV from the ingress to the egress and then the egress node returns the MAPPING messages back Upstream carrying the Admin Status TLV. A modification message of this type would typically only carry the mandatory CR-LDP TLVs and the Admin Status

TLV.

7.3.1. Deletion procedure

In some circumstances, particularly optical networks, it is useful to set the administrative status of an LSP before tearing it down. In such circumstances the procedure SHOULD be followed when deleting an LSP:

The ingress node precedes an LSP deletion by inserting an Admin Status TLV in a REQUEST Message with the modification action indicator set to modify message and setting the Down (D) bit.

Transit and egress nodes process the Admin Status TLV as described above.

Upon receiving the Admin Status TLV with the Down (D) bit set in the MAPPING message with the modification action indicator set to modify the ingress node sends a RELEASE message downstream to remove the LSP and normal CR-LDP processing takes place.

7.3.2. Compatibility

It is possible that some nodes along an LSP will not support the Admin Status TLV. In the case of a non-supporting transit node, the TLV will pass through the node unmodified and normal processing can continue. In the case of a non-supporting egress node, the Admin Status TLV may not be reflected back in the MAPPING Message. In this case, the ingress SHOULD continue to set the contents of the object normally but, when processing an LSP deletion, it MUST NOT wait for an updated Admin Status TLV in a MAPPING message before issuing a RELEASE/WITHDRAW message.

7.3.3. Notify Message Procedures

Intermediate and egress nodes may trigger the setting of administrative status before a deletion via the use of REQUEST messages. To accomplish this, an intermediate or egress node generates a REQUEST message with the modification action indicator set to modify and with the corresponding upstream. The Admin Status TLV MUST be included in the message, with the Down (D) bit set.

An ingress node receiving a MAPPING message containing an Admin Status TLV with the Down (D) bit set, SHOULD initiate the deletion procedure described in the previous section.

8. Control Channel Separation

This section provides the protocol specific formats and procedures to required support a control channel not being in-band with a data channel.

8.1. Interface Identification

The choice of the data interface to use is always made by the sender of the REQUEST message. The choice of the data interface is indicated by the sender of the REQUEST message by including the data channel's interface identifier in the message using a new Interface TLV. type. For bidirectional LSPs, the sender chooses the data interface in each direction. In all cases but bundling [MPLS-BUNDLE] the upstream interface is implied by the downstream interface. For bundling, the path sender explicitly identifies the component interface used in each direction.
The format of Interface ID in REQUEST, MAPPING Messages is:

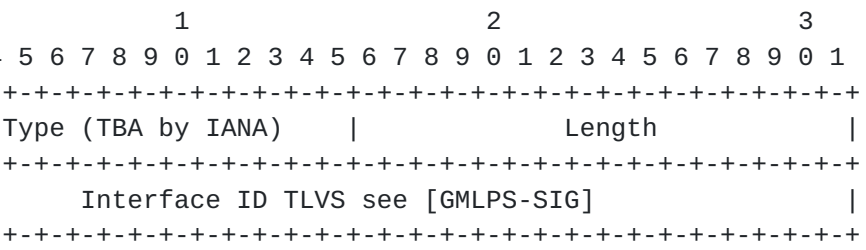

The U Bit Should be set to 0. The F bit should be set to 0.

See [GMPLS-SIG] for a description of parameters.

See [CR-LDP] for a description of signaling address. See [GMPLS-SIG] for a description of parameters and encoding of TLVs.

8.2. Procedures

An IF_ID TLV is used on links where there is not a one-to- one association of a control channel to a data channel, see [GMPLS- SIG].

The LDP session uses the IF_ID TLV to identify the data channel(s) associated with the LSP. For a unidirectional LSP, a forward channel MUST be indicated. For a bidirectional LSP that use bundled links, a reverse channel MUST be indicated. Data channels are specified from the viewpoint of the sender of the REQUEST message.

9. Fault Handling

In optical transport networks, failures in the out-of-fiber signaling communication or optical control plane should not have service impact on the existing optical connections. Under such circumstances, a mechanism MUST exist to detect a signaling communication failure and a recovery procedure SHALL guarantee connection integrity at both ends of the signaling channel.

The LDP Fault tolerant draft [LDP-FT] specifies the procedures for recovering LDP and CR-LDP sessions under failure. Please refer to this draft for procedures on recovering optical connections. Currently the Fault tolerant draft covers many of the common failure modes for a separated control and data plane.

10. Acknowledgments

This draft is the work of numerous authors and consists of a composition of a number of previous drafts in this area. A list of the drafts from which material and ideas were incorporated follows:

draft-saha-rsvp-optical-signaling-00.txt
draft-lang-mpls-rsvp-oxc-00.txt
draft-kompella-mpls-optical-00.txt
draft-fan-mpls-lambda-signaling-00.txt

Valuable comments and input were received from a number of people, notably Adrian Farrel.

11. Security Considerations

This draft introduce no new security considerations to [CR-LDP].

12. References

[CR-LDP] Jamoussi et al., "Constraint-Based LSP Setup using LDP", draft-ietf-mpls-cr-ldp-05.txt, Feb., 2001.

[MPLS-HIERARCHY] Kompella, K., and Rekhter, Y., "LSP Hierarchy with MPLS TE", Internet Draft, draft-ietf-mpls-lsp-hierarchy-02.txt, Feb., 2001.

[MPLS-UNNUM] Kompella, K., Rekhter, Y., "Signalling Unnumbered Links in CR-LDP", Internet Draft, draft-ietf-mpls-crldp-unnum-01.txt, February 2001

[GMPLS-RSVP] Ashwood-Smith, P. et al, "Generalized MPLS Signaling - RSVP-TE Extensions", Internet Draft, draft-ietf-mpls-generalized-rsvp-te-01.txt, February 2001.

[GMPLS-SIG] Ashwood-Smith, P. et al, "Generalized MPLS - Signaling Functional Description", Internet Draft, draft-ietf-mpls-generalized-signaling-02.txt, February 2001.

[LDP-FT] Farrel, A. et al, "Fault Tolerance for LDP and CR-LDP", Internet Draft, draft-ietf-mpls-ldp-ft-02.txt, February 2001.

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

13. Authors' Addresses

Peter Ashwood-Smith
Nortel Networks Corp.
P.O. Box 3511 Station C,
Ottawa, ON K1Y 4H7
Canada
Phone: +1 613 763 4534
Email: petera@nortelnetworks.com

Ayan Banerjee
Calient Networks
5853 Rue Ferrari
San Jose, CA 95138
Phone: +1 408 972-3645
Email: abanerjee@calient.net

Lou Berger
Movaz Networks, Inc.
7926 Jones Branch Drive
Suite 615
McLean VA, 22102
Phone: +1 703 847-1801
Email: lberger@movaz.com

Greg Bernstein
Ciena Corporation
10480 Ridgeview Court
Cupertino, CA 94014
Phone: +1 408 366 4713
Email: greg@ciena.com

John Drake
Calient Networks
5853 Rue Ferrari
San Jose, CA 95138
Phone: +1 408 972 3720
Email: jdrake@calient.net

Yanhe Fan
Axiowave Networks, Inc.
100 Nickerson Road
Marlborough, MA 01752
Phone: +1 508 460 6969 Ext. 627
Email: yfan@axiowave.com

Don Fedyk
Nortel Networks Corp.
600 Technology Park
Billerica MA 01821
Phone: +1 978 288 3041
Fax: +1 978 288 0620
Email: dwfedyk@nortelnetworks.com

Kireeti Kompella
Juniper Networks, Inc.
1194 N. Mathilda Ave.
Sunnyvale, CA 94089
Email: kireeti@juniper.net

Jonathan P. Lang
Calient Networks
25 Castilian
Goleta, CA 93117
Email: jplang@calient.net

Eric Mannie
EBONE
Terhulpsesteenweg 6A
1560 Hoeilaart - Belgium
Phone: +32 2 658 56 52
Mobile: +32 496 58 56 52
Fax: +32 2 658 51 18
Email: eric.mannie@ebone.com

Bala Rajagopalan
Tellium, Inc.
2 Crescent Place
P.O. Box 901
Oceanport, NJ 07757-0901
Phone: +1 732 923 4237
Fax: +1 732 923 9804
Email: braja@tellium.com

Yakov Rekhter
Juniper Networks, Inc.
Email: yakov@juniper.net

Debanjan Saha
Tellium Optical Systems
2 Crescent Place
Oceanport, NJ 07757-0901
Phone: +1 732 923 4264
Fax: +1 732 923 9804
Email: dsaha@tellium.com

Vishal Sharma
Jasmine Networks, Inc.
3061 Zanker Road, Suite B
San Jose, CA 95134
Phone: +1 408 895 5030
Fax: +1 408 895 5050
Email: vsharma@jasminenetworks.com

George Swallow
Cisco Systems, Inc.
250 Apollo Drive
Chelmsford, MA 01824
Voice: +1 978 244 8143
Email: swallow@cisco.com

Z. Bo Tang
Tellium, Inc.
2 Crescent Place
P.O. Box 901
Oceanport, NJ 07757-0901
Phone: +1 732 923 4231
Fax: +1 732 923 9804
Email: btang@tellium.com

Generated on: Fri Jul 20 16:49:49 EDT 2001