

Network Working Group
Request for Comments: 3472
Category: Standards Track

P. Ashwood-Smith, Editor
Nortel Networks
L. Berger, Editor
Movaz Networks
January 2003

Generalized Multi-Protocol Label Switching (GMPLS) Signaling Constraint-based Routed Label Distribution Protocol (CR-LDP) Extensions

Status of this Memo

This document specifies an Internet standards track protocol for the Internet community, and requests discussion and suggestions for improvements. Please refer to the current edition of the "Internet Official Protocol Standards" (STD 1) for the standardization state and status of this protocol. Distribution of this memo is unlimited.

Copyright Notice

Copyright (C) The Internet Society (2003). All Rights Reserved.

Abstract

This document describes extensions to Multi-Protocol Label Switching (MPLS) Constraint-based Routed Label Distribution Protocol (CR-LDP) signaling required to support Generalized MPLS. Generalized MPLS extends the MPLS control plane to encompass time-division (e.g., Synchronous Optical Network and Synchronous Digital Hierarchy, SONET/SDH), 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. A generic functional description can be found in separate documents.

Table of Contents

1. Introduction	2
2. Label Related Formats	3
2.1 Generalized Label Request	3
2.2 Generalized Label	4
2.3 Waveband Switching	5
2.4 Suggested Label	6
2.5 Label Set	6
3. Bidirectional LSPs	8
3.1 Procedures	8
4. Notification on Label Error	9
5. Explicit Label Control	9
5.1 Procedures	9

6.	Protection TLV	10
6.1	Procedures	11
7.	Administrative Status Information	11
7.1	Admin Status TLV	11
7.2	REQUEST and MAPPING Message Procedures	12
7.3	Notification Message Procedures	13
8.	Control Channel Separation	14
8.1	Interface Identification	14
8.2	Errored Interface Identification	15
9.	Fault Handling	17
10	Acknowledgments	17
11.	Security Considerations	17
12.	IANA Considerations	17
13.	Intellectual Property Considerations	18
14.	References	18
14.1	Normative References	18
14.2	Informative References	19
15.	Contributors	19
16.	Editors' Addresses	22
17.	Full Copyright Statement	23

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 [RFC3471]. This document presents CR-LDP specific formats and mechanisms needed to support all four classes of interfaces. RSVP-TE extensions can be found in [RFC3473].

[RFC3471] should be viewed as a companion document to this document. The format of this document parallels [RFC3471]. It should be noted that the RSVP-TE specific version of Generalized MPLS includes RSVP specific support for rapid failure notification, see Section 4 [RFC3473]. 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 (Label Switched Path) Encoding Type as possible to allow the maximum flexibility in switching by transit LSRs. A Generalized Label Request Type, Length, and Value (TLV) is set by the ingress node, transparently passed by transit nodes, and used by the egress node. The Switching Type field may also be updated hop-by-hop.

The format of a Generalized Label Request is:

```

      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
+-----+-----+-----+-----+-----+-----+-----+-----+-----+
|U|F|          Type (0x0824)          |          Length          |
+-----+-----+-----+-----+-----+-----+-----+-----+-----+
| LSP Enc. Type | Switching Type |          G-PID          |
+-----+-----+-----+-----+-----+-----+-----+-----+-----+

```

See [RFC3471] 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 (Explicit Route) 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.

Nodes MUST verify that the type indicated in the Switching Type parameter is supported on the corresponding incoming interface. If the type cannot be supported, the node MUST generate a NOTIFICATION message with a "Routing problem/Switching Type" 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 G-PID" 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 [RFC3471] 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 [RFC3471] 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 be 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 0x0828 is assigned for the Waveband Label.

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

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	2	3	4	5	6	7	8	9																								
U F																Type (0x0828)																Length																															
																Waveband Id																																															
																Start Label																																															
																End Label																																															

See [RFC3471] 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 swapped 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

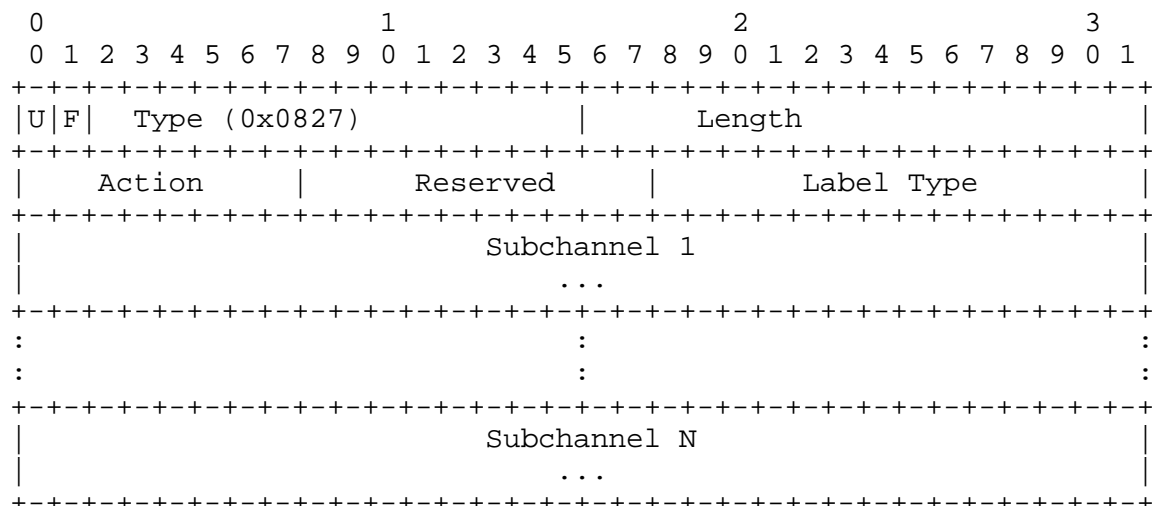
The format of a suggested label is identical to a generalized label. It is used in REQUEST messages. Suggested Label uses type = 0x904.

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

Per [RFC3471], if a downstream node passes a label value that differs from the suggested label upstream, the upstream LSR MUST either reconfigure itself so that it uses the label specified by the downstream node or generate a NOTIFICATION message with a "Routing problem/Unacceptable label value" indication. Furthermore, an ingress node SHOULD NOT transmit data traffic using a suggested label until the downstream node passes corresponding a label upstream.

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 [RFC3471] 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 message or if the selection is made immediately for propagation in the MAPPING message.

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 = 0x0826.

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 [RFC3471]. An Acceptable Label Set TLV uses a type value of 0x082a. 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 TLV is defined as follows:

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	2	3	4	5	6	7	8	9										
0 0										Type (0x0829)																				Length																			
L U										Reserved																				Label																			
										Label (continued)																																							
										...																																							

See [RFC3471] 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. 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:

- o 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.
- o For a label ER-Hop to follow a ER-Hop that has the L-bit set.
- o On unidirectional LSP setup, for there to be a label ER-Hop with the U-bit set.
- o 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 REQUEST 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 REQUEST 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 on 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:

```

      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
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
|U|F|           Type (0x0835)           |           Length           |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
|S|                               Reserved                               | Link Flags|
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+

```

See [RFC3471] 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 TLV is carried in REQUEST and MAPPING messages to indicate the administrative state of an LSP. In the second, the TLV is carried in Notification message to request a change to the administrative state of an LSP.

7.1. Admin Status TLV

The use of the Admin Status TLV is optional. It uses Type = 0x082b. The format of the TLV is:

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

```

      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
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
|U|F|           Type (0x082b)           |           Length           |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
|R|                               Reserved                               |T|A|D|
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+

```

See [RFC3471] 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. Each node processes status information 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. The absence of the TLV is the equivalent to receiving a TLV containing values all set to zero.

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.

Edge 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. When the ADMIN Status TLV is received with the R bit set, the receiving edge node should reflect the received values in a corresponding MAPPING message. Specifically, if an egress node receives a Request message with the R bit of the Admin_Status TLV set and the node the node SHOULD send a Mapping message containing an Admin_Status TLV with the same values set, with the exception of the R bit, as received in the corresponding Request message.

7.2.1. 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 from the ingress:

- o The ingress node precedes an LSP deletion by inserting an Admin Status TLV in a Notification Message setting the Reflect (R) and Delete (D) bits.
- o Transit nodes process the Admin Status TLV by passing the Notification message. The egress node May respond with a Notification message with the Admin Status TLV.
- o Upon receiving the Admin Status TLV with the Delete (D) bit set in the Notification message, the egress SHOULD respond with a LABEL WITHDRAW message and normal CR-LDP processing takes place.

In such circumstances the procedure SHOULD be followed when deleting an LSP from the egress:

- o The egress node indicates its desire for deletion by inserting an Admin Status TLV in a Notification message and setting Delete (D) bit.
- o Transit nodes process the Admin Status TLV as described above.
- o Upon receiving the Admin Status TLV with the Delete (D) bit set in the Notification message, the ingress node sends a LABEL RELEASE message downstream to remove the LSP and normal CR-LDP processing takes place.

7.3. Notification Message Procedures

Subsequent messaging Admin Status messaging may be performed by Notification Messages. The ingress may begin the propagation of a Notification Message with an Admin Status TLV. Each subsequent node propagates the Notification with the Admin Status TLV from the ingress to the egress and then the egress node returns the Notification messages back Upstream carrying the Admin Status TLV.

Intermediate and egress nodes may trigger the setting of administrative status via the use of Notification messages. To accomplish this, an intermediate or egress node generates a Notification message with the corresponding upstream notify session information. The Admin Status TLV MUST be included in the session information, with the appropriate bit or bits set. The Reflect (R) bit MUST NOT be set.

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

7.3.1. Compatibility and Error Procedures

Some special processing is required in order to cover the case of nodes that do not support the Admin Status TLV and other error conditions. Specifically, a node that sends a Notification message containing an Admin Status TLV with the Down (D) bit set MUST verify that it receives a corresponding LABEL RELEASE message within a configurable period of time. By default this period of time SHOULD be 30 seconds. If the node does not receive such a LABEL RELEASE message, it SHOULD send a Label Release message downstream and a LABEL WITHDRAW message upstream.

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, the upstream interface is implied by the downstream interface. For bundling, the REQUEST sender explicitly identifies the component interface used in each direction.

8.1.1. Interface ID TLV

The format of IPV4 Interface ID in REQUEST, MAPPING Messages is:

```

      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
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
|U|F|      Type (0x082d)      |      Length      |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
|                               IPv4 Next/Previous Hop Address      |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
|                               Logical Interface ID                |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
|                               Interface ID TLVS see [RFC3471]      |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+

```

The format of IPV6 Interface ID TLV in REQUEST, MAPPING Messages is:

```

      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
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
|U|F|      Type (0x082e)      |      Length      |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
|                               IPv6 Next/Previous Hop Address      |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
|                               Logical Interface ID                |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
|                               Interface ID TLVS see [RFC3471]      |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+

```

See [RFC3471] for a description of parameters.

See [RFC3212] for a description of signaling address. See [RFC3471] for a description of parameters and encoding of TLVs.

8.1.2. 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 [RFC3471].

The LDP session uses the IF_ID TLV to identify the data channel(s) associated with the LSP. For a unidirectional LSP, a downstream data channel MUST be indicated. For bidirectional LSPs, a common downstream and upstream data channel is normally indicated. In the special case where a bidirectional LSP that traverses a bundled link, it is possible to specify a downstream data channel that differs from the upstream data channel. Data channels are specified from the viewpoint of the sender of a REQUEST message. The IF_ID TLV SHOULD NOT be used when no TLVs are needed.

A node receiving one or more IF_ID TLVs in a REQUEST message saves their values and returns them in the subsequent MAPPING message sent to the node that originated the TLVs.

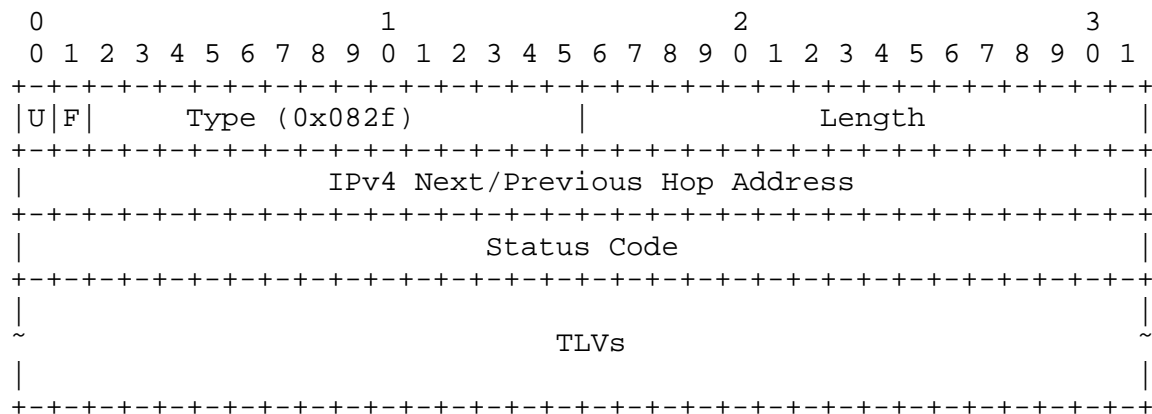
Note, the node originating an IF_ID TLV MUST ensure that the selected outgoing interface, as specified in the IF_ID TLV, is consistent with an ERO. A node that receives an IF_ID TLV SHOULD check whether the information carried in this TLV is consistent with the information carried in a received ERO, and if not it MUST send a LABEL ABORT Message with the error code "Routing Error" and error value of "Bad Explicit Routing TLV Error" toward the sender. This check CANNOT be performed when the initial ERO subobject is not the incoming interface.

8.2. Errored Interface Identification

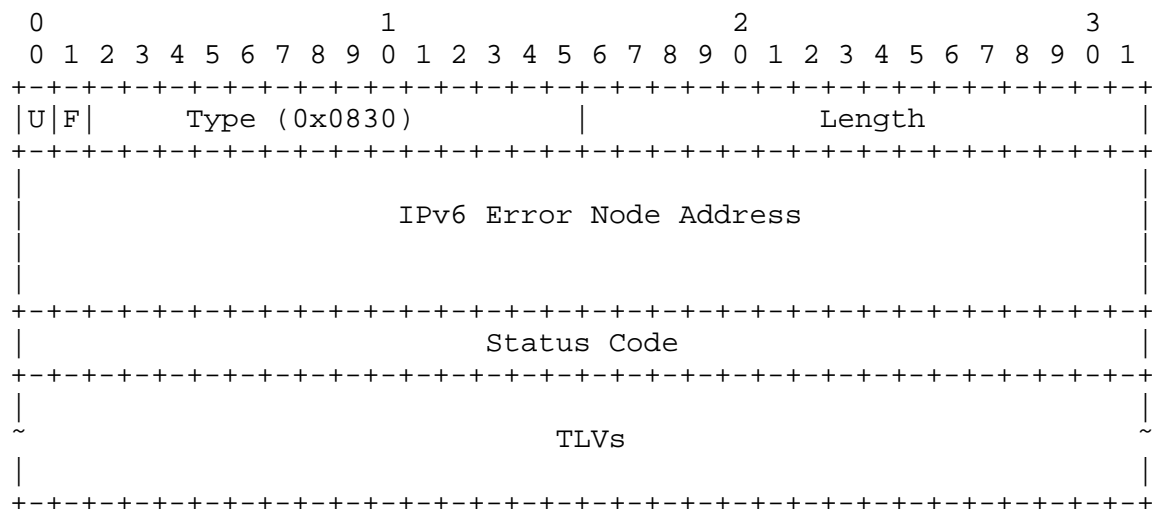
There are cases where it is useful to indicate a specific interface associated with an error. To support these cases the IF_ID Status TLV are defined.

8.2.1. IF_ID Status TLVs

The format of the IPv4 IF_ID Status TLV is:



The format of the IPv6 IF_ID Status TLV is:



See [RFC3036] for a description of status code value fields. See [RFC3471] for a description of parameters and encoding of TLVs.

8.2.2. Procedures

Nodes wishing to indicate that an error is related to a specific interface SHOULD use the appropriate IF_ID Status TLV in the corresponding LABEL WITHDRAW or LABEL RELEASE message. IF_ID Status TLV SHOULD be generated and processed as any other Status TLV, see [RFC3036].

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

10. Acknowledgments

This document is the work of numerous authors and consists of a composition of a number of previous documents in this area.

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

11. Security Considerations

This document introduces no new security considerations to [RFC3212].

12. IANA Considerations

This document uses the LDP [RFC3036] name spaces, see <http://www.iana.org/assignments/ldp-namespaces>, which lists the assignments for the following TLVs:

- o Generalized Label Request (TLV 0x0824)
- o Generalized Label (TLV 0x0825)
- o Upstream Label (TLV 0x0826)
- o Label Set (TLV 0x0827)
- o Waveband Label (TLV 0x0828)
- o ER-Hop (TLV 0x0829)
- o Acceptable Label Set (TLV 0x082a)
- o Admin Status (TLV 0x082b)
- o Interface ID (TLV 0x082c)
- o IPV4 Interface ID (TLV 0x082d)
- o IPV6 Interface ID (TLV 0x082e)
- o IPV4 IF_ID Status (TLV 0x082f)
- o IPV6 IF_ID Status (TLV 0x0830)
- o Protection (TLV 0x0835)

13. Intellectual Property Considerations

This section is taken from Section 10.4 of [RFC2026].

The IETF takes no position regarding the validity or scope of any intellectual property or other rights that might be claimed to pertain to the implementation or use of the technology described in this document or the extent to which any license under such rights might or might not be available; neither does it represent that it has made any effort to identify any such rights. Information on the IETF's procedures with respect to rights in standards-track and standards-related documentation can be found in BCP-11. Copies of claims of rights made available for publication and any assurances of licenses to be made available, or the result of an attempt made to obtain a general license or permission for the use of such proprietary rights by implementors or users of this specification can be obtained from the IETF Secretariat.

The IETF invites any interested party to bring to its attention any copyrights, patents or patent applications, or other proprietary rights which may cover technology that may be required to practice this standard. Please address the information to the IETF Executive Director.

14. References

14.1. Normative References

- [RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels," BCP 14, RFC 2119. March 1997.
- [RFC3036] Andersson, L., Doolan, P., Feldman, N., Fredette, A. and B. Thomas, "LDP Specification", RFC 3036, January 2001.
- [RFC3212] Jamoussi, B., Andersson, L., Callon, R., Dantu, R., Wu, L., Doolan, P., Worster, T., Feldman, N., Fredette, A., Girish, M., Gray, E., Heinanen, J., Kilty, T. and A. Malis, "Constraint-Based LSP Setup using LDP", RFC 3212, January 2002.
- [RFC3471] Berger, L., Editor, "Generalized Multi-Protocol Label Switching (GMPLS) Signaling Functional Description", RFC 3471, January 2003.

14.2. Informative References

- [LDP-FT] Farrel, A., et al, "Fault Tolerance for LDP and CR-LDP", Work in Progress.
- [MPLS-HIERARCHY] Kompella, K. and Y. Rekhter, "LSP Hierarchy with MPLS TE", Work in Progress.
- [MPLS-UNNUM] Kompella, K., Rekhter, Y. and A. Kullberg, "Signalling Unnumbered Links in CR-LDP", Work in Progress.
- [RFC2026] Bradner, S., "The Internet Standards Process -- Revision 3," BCP 9, RFC 2026, October 1996.
- [RFC3473] Berger, L., Editor, "Generalized Multi-Protocol Label Switching (GMPLS) Signaling - Resource ReserVation Protocol-Traffic Engineering (RSVP-TE) Extensions", RFC 3473, January 2003.

15. Contributors

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
EMail: gregb@grotto-networking.com

Yanhe Fan
Axiowave Networks, Inc.
200 Nickerson Road
Marlborough, MA 01752

Phone: + 1 774 348 4627
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

Jonathan P. Lang
EMail: jplang@ieee.org

Eric Mannie
Independent Consultant
2 Avenue de la Folle Chanson
1050 Brussels
Belgium

EMail: eric_mannie@hotmail.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

Debanjan Saha
EMail: debanjan@acm.org

Vishal Sharma
Metanoia, Inc.
1600 Villa Street, Unit 352
Mountain View, CA 94041-1174

Phone: +1 650-386-6723
EMail: v.sharma@ieee.org

George Swallow
Cisco Systems, Inc.
250 Apollo Drive
Chelmsford, MA 01824

Phone: +1 978 244 8143
EMail: swallow@cisco.com

Z. Bo Tang
EMail: botang01@yahoo.com

16. Editors' 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

Lou Berger
Movaz Networks, Inc.
7926 Jones Branch Drive
Suite 615
McLean VA, 22102

Phone: +1 703 847-1801
EMail: lberger@movaz.com

17. Full Copyright Statement

Copyright (C) The Internet Society (2003). All Rights Reserved.

This document and translations of it may be copied and furnished to others, and derivative works that comment on or otherwise explain it or assist in its implementation may be prepared, copied, published and distributed, in whole or in part, without restriction of any kind, provided that the above copyright notice and this paragraph are included on all such copies and derivative works. However, this document itself may not be modified in any way, such as by removing the copyright notice or references to the Internet Society or other Internet organizations, except as needed for the purpose of developing Internet standards in which case the procedures for copyrights defined in the Internet Standards process must be followed, or as required to translate it into languages other than English.

The limited permissions granted above are perpetual and will not be revoked by the Internet Society or its successors or assigns.

This document and the information contained herein is provided on an "AS IS" basis and THE INTERNET SOCIETY AND THE INTERNET ENGINEERING TASK FORCE DISCLAIMS ALL WARRANTIES, EXPRESS OR IMPLIED, INCLUDING BUT NOT LIMITED TO ANY WARRANTY THAT THE USE OF THE INFORMATION HEREIN WILL NOT INFRINGE ANY RIGHTS OR ANY IMPLIED WARRANTIES OF MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE.

Acknowledgement

Funding for the RFC Editor function is currently provided by the Internet Society.

