

February 1999

Constraint-Based LSP Setup using LDP

[draft-ietf-mpls-cr-ldp-01.txt](#)

Status of this Memo

This document is an Internet-Draft and is in full conformance with all provisions of [Section 10 of RFC2026](#).

Internet-Drafts are working documents of the Internet Engineering Task Force (IETF), its areas, and its working groups. Note that other groups may also distribute working documents as Internet-Drafts.

Internet-Drafts are draft documents valid for a maximum of six months and may be updated, replaced, or obsoleted by other documents at any time. It is inappropriate to use Internet-Drafts as reference material or to cite them other than as "work in progress."

The list of current Internet-Drafts can be accessed at <http://www.ietf.org/ietf/1id-abstracts.txt>

The list of Internet-Draft Shadow Directories can be accessed at <http://www.ietf.org/shadow.html>.

Distribution of this memo is unlimited.

Copyright Notice

Copyright (C) The Internet Society (1998). All Rights Reserved.

Abstract

Label Distribution Protocol (LDP) is defined in [[LDP](#)] for distribution of labels inside one MPLS domain. One of the most important services that may be offered using MPLS in general and LDP in particular is support for constraint-based routing of traffic across the routed network. Constraint-based routing offers the opportunity to extend the information used to setup paths beyond what is available for the routing protocol. For instance, an LSP can be setup based on explicit route constraints, QoS constraints, and others. Constraint-based routing (CR) is a mechanism used to meet Traffic Engineering requirements that have been proposed by [[FRAME](#)], [[ARCH](#)] and [[TER](#)]. These requirements may be met by extending LDP for support of constraint-based routed label switched paths (CRLSPs).

Other uses exist for CRLSPs as well ([[VPN1](#)], [[VPN2](#)] and [[VPN3](#)]).

This draft specifies mechanisms and TLVs for support of CRLSPs using LDP. The Explicit Route object and procedures are extracted from [[ER](#)].

Table of Contents

1.	Introduction	3
2.	Constraint-based Routing Overview	3
2.1	Strict and Loose Explicit Routes	4
2.2	Traffic Characteristics	4
2.3	Pre-emption	5
2.4	Route Pinning	5
2.5	Resource Class	5
3.	Solution Overview	5
3.1	Required Messages and TLVs	7
3.2	Label Request Message	7
3.3	Label Mapping Message	8
3.4	Notification Message	9
3.5	Release & Withdraw Messages	9
4.	Protocol Specification	9
4.1	Explicit Route TLV (ER-TLV)	10
4.2	Explicit Route Hop TLV	10
4.3	Traffic Parameters TLV	12
4.3.1	Semantics	13
4.3.1.1	Frequency	13
4.3.1.2	Peak Rate	14
4.3.1.3	Committed Rate	14
4.3.1.4	Excess Burst Size	14
4.3.1.5	Peak Rate Token Bucket.....	14
4.3.1.6	Committed Data Rate Token Bucket	15
4.3.1.7	Weight	16
4.3.2	Procedures	16
4.3.2.1	Label Request Message	16
4.3.2.2	Label Mapping Message	16
4.3.2.3	Notification Message	17
4.4	Preemption TLV	18
4.5	LSPID TLV	18
4.6	Resource Class TLV	19
4.7	ER-Hop Semantics	19
4.7.1	ER-Hop 1 TLV IPv4 Prefix	20
4.7.2	ER-Hop 2 TLV IPv6 Prefix	20
4.7.3	ER-Hop 3 TLV AS Number	21
4.7.4	ER-Hop 4 TLV LSPID	21
4.8	Processing of the ER-TLV	22
4.8.1	Selection of the next hop	22
4.8.2	Adding the Label Request Message to the next hop	24
4.9	Route Pinning TLV	24
4.10	CR-LSP FEC Element	24
4.11	Error Subcodes	25

5.	Security Considerations	26
6.	Acknowledgement	26
7.	References	26
8.	Author Information	28
Appendix A	CRLSP Establishment Examples	30
A.1	Strict Explicit Route Example	30
A.2	Node Groups and Specific Nodes Example	31
Appendix B	QoS Service Examples	34
B.1	Service Examples	34
B.2	Establishing CR-LSP Supporting Real-Time Applications.	35
B.3	Establishing CR-LSP Delay Insensitive Applications ...	36

[1.](#) Introduction

The need for constraint-based routing (CR) in MPLS has been explored elsewhere [[ARCH](#)], [[FRAME](#)], and [[TER](#)]. Explicit routing is a subset of the more general constraint-based routing function. At the MPLS WG meeting held during the Washington IETF there was consensus that LDP should support explicit routing of LSPs with provision for indication of associated (forwarding) priority. In the Chicago meeting, a decision was made that support for explicit path setup in LDP will be moved to a separate document. This document provides that support and it has been accepted as a working document in the Orlando meeting. This specification proposes an end-to-end setup mechanism of a constraint-based routed LSP (CRLSP) initiated by the ingress LSR. We also specify mechanisms to provide means for reservation of resources using LDP.

This document introduce TLVs and procedures that provide support for:

- Strict and Loose Explicit Routing
- Specification of Traffic Parameters
- Route Pinning
- CRLSP Pre-emption though setup/holding priorities
- Handling Failures
- LSPID
- Resource Class

[Section 2](#) introduces the various constraints defined in this specification. [Section 3](#) outlines the CR-LDP solution. [Section 4](#) defines the TLVs and procedures used to setup constraint-based routed label switched paths. [Appendix A](#) provides several examples of CR-LSP path setup. [Appendix B](#) provides Service Definition Examples.

[2.](#) Constraint-based Routing Overview

Constraint-based routing is a mechanism that supports the Traffic Engineering requirements defined in [[TER](#)]. Explicit Routing is a subset of the more general constraint-based routing where the

constraint is the explicit route (ER). Other constraints are defined to provide a network operator with control over the path taken by an LSP. This section is an overview of the various constraints supported by this specification.

2.1 Strict and Loose Explicit Routes

Like any other LSP an CRLSP is a path through an MPLS network. The difference is that while other paths are setup solely based on information in routing tables or from a management system, the constraint-based route is calculated at one point at the edge of network based on criteria, including but not limited to routing information. The intention is that this functionality shall give desired special characteristics to the LSP in order to better support the traffic sent over the LSP. The reason for setting up CRLSPs, might be that one wants to assign certain bandwidth or other Service Class characteristics to the LSP, or that one wants to make sure that alternative routes use physically separate paths through the network.

An explicit route is represented in a Label Request Message as a list of nodes or groups of nodes along the constraint-based route. When the CRLSP is established, all or a subset of the nodes in a group may be traversed by the LSP. Certain operations to be performed along the path can also be encoded in the constraint-based route.

The capability to specify, in addition to specified nodes, groups of nodes, of which a subset will be traversed by the CRLSP, allows the system a significant amount of local flexibility in fulfilling a request for a constraint-based route. This allows the generator of the constraint-based route to have some degree of imperfect information about the details of the path.

The constraint-based route is encoded as a series of ER-Hops contained in a constraint-based route TLV. Each ER-Hop may identify a group of nodes in the constraint-based route. A constraint-based route is then a path including all of the identified groups of nodes.

To simplify the discussion, we call each group of nodes an abstract node. Thus, we can also say that a constraint-based route is a path including all of the abstract nodes, with the specified operations occurring along that path.

2.2 Traffic Characteristics

The traffic characteristics of a path are described in the Traffic Parameters TLV in terms of a peak rate, committed rate, and service granularity. The peak and committed rates describe the bandwidth constraints of a path while the service granularity can be used to specify a constraint on the delay variation that the CRLDP MPLS domain may introduce to a path's traffic.

2.3 Pre-emption

CR-LDP signals the resources required by a path on each hop of the route. If a route with sufficient resources can not be found, existing paths may be rerouted to reallocate resources to the new path. This is the process of path pre-emption. Setup and holding priorities are used to rank existing paths (holding priority) and the new path (setup priority) to determine if the new path can pre-empt an existing path.

The setupPriority of a new CRLSP and the holdingPriority attributes of the existing CRLSP are used to specify priorities. Signaling a higher holding priority expresses that the path, once it has been established, should have a lower chance of being pre-empted. Signaling a higher setup priority expresses the expectation that, in the case that resource are unavailable, the path is more likely to pre-empt other paths. The exact rules determining bumping are an aspect of network policy.

The allocation of setup and holding priority values to paths is an aspect of network policy.

The setup and holding priority values range from zero (0) to seven (7). The value zero (0) is the priority assigned to the most important path. It is referred to as the highest priority. Seven (7) is the priority for the least important path. The use of default priority values is an aspect of network policy.

The setupPriority of a CRLSP should not be higher (numerically less) than its holdingPriority since it might bump an LSP and be bumped by next "equivalent" request.

2.4 Route Pinning

Route pinning is applicable to segments of an LSP that are loosely routed - i.e. those segments which are specified with a next hop with the 'L' bit set or where the next hop is an "abstract node". A CRLSP may be setup using route pinning if it is undesirable to change the path used by an LSP because a better next hop becomes available at some LSR along the loosely routed portion of the LSP.

2.5 Resource Class

Network resources may be classified in various ways by the network operator. These classes are also known as "colors" or "administrative groups". When an CR-LSP is being established, it's necessary to indicate which resource classes the CR-LSP can draw from.

3. Solution Overview

CRLSP over LDP Specification is designed with the following goals:

1. Meet the requirements outlined in [TER] for performing traffic engineering and provide a solid foundation for performing more general constraint-based routing.
2. Build on already specified functionality that meets the requirements whenever possible. Hence, this specifications is based on [LDP] and the Explicit Route object and procedures defined in [ER].
3. Keep the solution simple.

In this document, support for unidirectional point-to-point CRLSPs is specified. Support for point-to-multipoint, multipoint-to-point, is for further study (FFS).

Support for constraint-based routed LSPs in this specification depends on the following minimal LDP behaviors as specified in [LDP]:

- Basic and/or Extended Discovery Mechanisms.
- Use the Label Request Message defined in [LDP] in downstream on demand label advertisement mode with ordered control.
- Use the Label Mapping Message defined in [LDP] in downstream on demand mode with ordered control.
- Use the Notification Message defined in [LDP].
- Use the Withdraw and Release Messages defined in [LDP].
- Use the Loop Detection (in the case of loosely routed segments of a CRLSP) mechanisms defined in [LDP].

In addition, the following functionality is added to what's defined in [LDP]:

- The Label Request Message used to setup a CRLSP includes one or more CR-TLVs defined in [Section 4](#). For instance, the Label Request Message may include the ER-TLV.
- An LSR implicitly infers ordered control from the existence of one or more CR-TLVs in the Label Request Message. This means that the LSR can still be configured for independent control for LSPs established as a result of dynamic routing. However, when a Label Request Message includes one or more of the CR-TLVs, then ordered control is used to setup the CRLSP. Note that this is also true for the loosely routed parts of a CRLSP.
- New status codes are defined to handle error notification for failure of established paths specified in the CR-TLV.

Examples of CRLSP establishment are given in [Appendix A](#) to illustrate how the mechanisms described in this draft work.

[3.1](#) Required Messages and TLVs

Any Messages, TLVs, and procedures not defined explicitly in this document are defined in the [\[LDP\]](#) Specification. The state transitions which relate to CR-LDP messages can be found in [\[LDP-STATE\]](#).

The following subsections are meant as a cross reference to the [\[LDP\]](#) document and indication of additional functionality beyond what's defined in [\[LDP\]](#) where necessary.

[3.2](#) Label Request Message

The Label Request Message is as defined in 3.5.8 of [\[LDP\]](#) with the following modifications (required only if any of the CR-TLVs is included in the Label Request Message):

- Only a single FEC-TLV may be included in the Label Request Message. The CR-LSP FEC TLV should be used.
- The Return Message ID TLV is MANDATORY.
- The Optional Parameters TLV includes the definition of any of the Constraint-based TLVs specified in [Section 4](#).
- The Procedures to handle the Label Request Message are augmented by the procedures for processing of the CR-TLVs as defined in [Section 4](#).

The encoding for the CR-LDP Label Request Message is 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																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																								
+	-	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+

3.3 Label Mapping Message

The Label Mapping Message is as defined in 3.5.7 of [[LDP](#)] with the following modifications:

- Only a single Label-TLV may be included in the Label Mapping Message.
- The Label Mapping Message MUST include Label Request Message ID TLV.
- The Label Mapping Message MUST include LSPID TLV.
- The Label Mapping Message Procedures are limited to downstream on demand ordered control mode.

A Mapping message is transmitted by a downstream LSR to an upstream LSR under one of the following conditions:

1. The LSR is the egress end of the CRLSP and an upstream mapping has been requested.
2. The LSR received a mapping from its downstream next hop LSR for an CRLSP for which an upstream request is still pending.

The encoding for the CR-LDP Label Mapping Message is 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								
U Label Mapping (0x0400)										Message Length																													
Message ID																																							
FEC TLV																																							
Label TLV																																							
Label Request Message ID TLV (mandatory)																																							
LSPID TLV (CR-LDP, mandatory)																																							
Traffic TLV (CR-LDP, optional)																																							

3.4 Notification Message

The Notification Message is as defined in Section 3.5.1 of [LDP] and the Status TLV encoding is as defined in Section 3.4.7 of [LDP].

Establishment of an Explicitly Routed LSP may fail for a variety of reasons. All such failures are considered advisory conditions and they are signaled by the Notification Message.

Notification Messages carry Status TLVs to specify events being signaled. New status codes are defined in [Section 4.11](#) to signal error notifications associated with the establishment of a CRLSP and the processing of the CR-TLV.

The Notification Message must carry the LSPID TLV of the corresponding CRLSP.

3.5 Release and Withdraw Messages

The Label Release and Label Withdraw Messages are used as specified in [\[LDP\]](#) to clear CR-LSPs. These message may also carry the LSPID TLV.

4. Protocol Specification

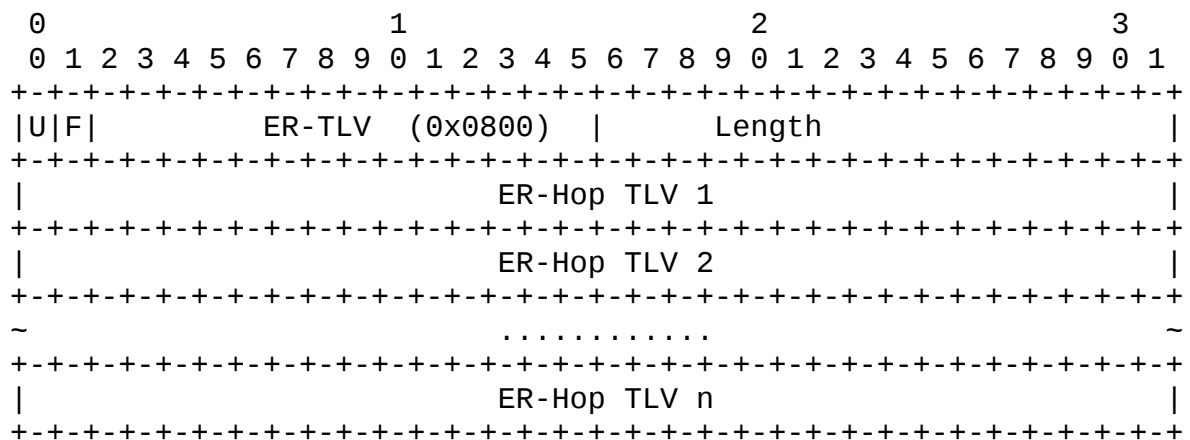
The Label Request Messages defined in [LDP] optionally carries one or more of the optional Constraint-based Routing TLVs (CR-TLVs) defined in this section. If needed, other constraints can be supported later through the definition of new TLVs. In this specification, the following TLVs are defined:

- Explicit Route TLV

- Explicit Route Hop TLV
- Traffic Parameters TLV
- Preemption TLV
- LSPID TLV
- Route Pinning TLV
- Resource Class TLV
- CRLSP FEC TLV

4.1 Explicit Route TLV (ER-TLV)

The ER-TLV is an object that specifies the path to be taken by the LSP being established. It is composed of one or more Explicit Route Hop TLVs (ER-Hop TLVs) defined in [Section 4.2](#).



U bit
Unknown TLV bit. As defined in [LDP].

F bit
Forward unknown TLV bit. As defined in [LDP].

Type


A two byte field carrying the value of the ER-TLV type which is 0x800.

Length
Specifies the length of the value field in bytes.

ER-Hop TLVs
One or more ER-Hop TLVs defined in [Section 4.2](#).

4.2 Explicit Route Hop TLV (ER-Hop TLV)

The contents of an ER-TLV are a series of variable length ER-Hop TLVs. Each ER-Hop TLV has the form:



U bit

Unknown TLV bit. As defined in [LDP].

F bit

Forward unknown TLV bit. As defined in [LDP].

ER-Hop Type

A fourteen-bit field indicating the type of contents of the ER-Hop. Currently defined values are:

Value	Type
0x801	IPv4 prefix
0x802	IPv6 prefix
0x803	Autonomous system number
0x804	LSPID

Length

Specifies the length of the value field in bytes.

L bit

The L bit is an attribute of the ER-Hop. The L bit is set if the ER-Hop represents a loose hop in the explicit route. If the bit is not set, the ER-Hop represents a strict hop in the explicit route.

The L bit in the ER-Hop is a one-bit attribute. If the L bit is set, then the value of the attribute is "loose." Otherwise, the value of the attribute is "strict." For brevity, we say that if the value of the ER-Hop attribute is loose then it is a "loose ER-Hop." Otherwise, it's a "strict ER-Hop." Further, we say that the abstract node of a strict or loose ER-Hop is a strict or a loose node, respectively. Loose and strict nodes are always interpreted relative to their prior abstract nodes.

The path between a strict node and its prior node MUST include only network nodes from the strict node and its prior abstract node.

The path between a loose node and its prior node MAY include other network nodes which are not part of the strict node or its prior abstract node.

Contents

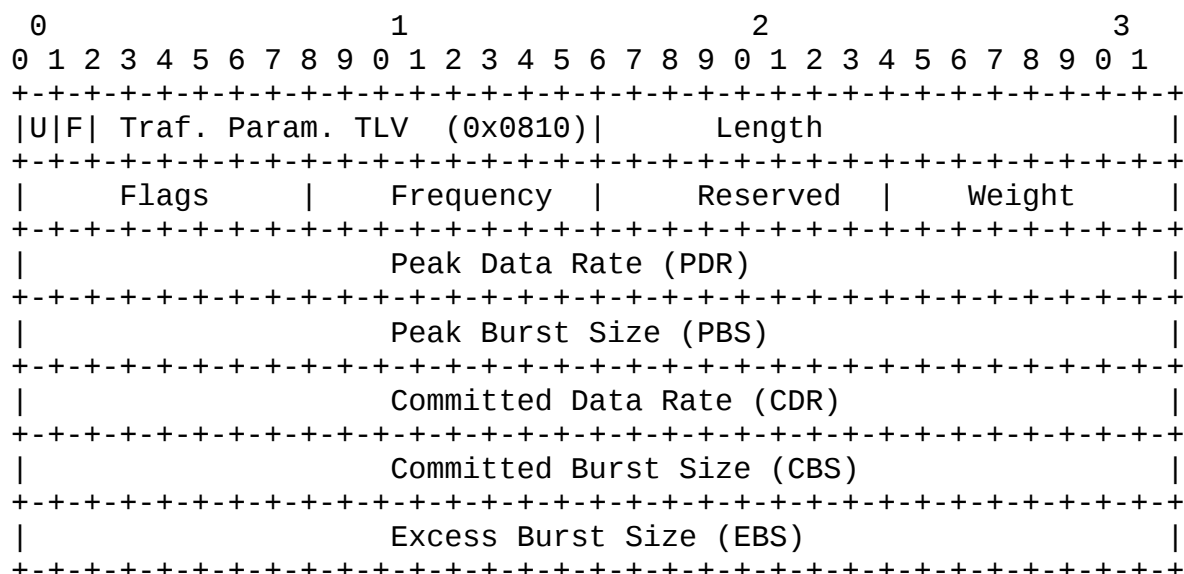
A variable length field containing the node or abstract node that is the consecutive nodes that make up the explicit routed LSP.

4.3 Traffic Parameters TLV

The following sections describe the CRLSP Traffic Parameters. The required characteristics of a CRLSP are expressed by the Traffic Parameter values.

A Traffic Parameters TLV, is used to signal the Traffic Parameter values. The Traffic Parameters are defined in the subsequent sections.

The Traffic Parameters TLV contains a Flags field, a Frequency, a Weight, and the five Traffic Parameters PDR, PBS, CDR, CBS, EBS. The Traffic Parameters TLV is shown below:



U bit

Unknown TLV bit. As defined in [LDP].

F bit

Forward unknown TLV bit. As defined in [LDP].

Type

A fourteen-bit field carrying the value of the ER-TLV type which is 0x810.

Length

Specifies the length of the value field in bytes.

Flags

The Flags field is shown below:

```

+---+---+---+---+---+---+---+---+
| Res |F6|F5|F4|F3|F2|F1|
+---+---+---+---+---+---+---+---+

```

Res - These bits are reserved.
Zero on transmission.
Ignored on receipt.

F1 - Corresponds to the PDR.
F2 - Corresponds to the PBS.
F3 - Corresponds to the CDR.
F4 - Corresponds to the CBS.
F5 - Corresponds to the EBS.
F6 - Corresponds to the Weight.

Each flag F_i is a Negotiable Flag corresponding to a Traffic Parameter. The Negotiable Flag value zero denotes NotNegotiable and value one denotes Negotiable.

Frequency

The Frequency field is coded as an 8 bit unsigned integer with the following code points defined:

0	- Unspecified
1	- Frequent
2	- VeryFrequent
3-255	- Reserved

Reserved

Zero on transmission. Ignored on receipt.

Weight

An 8 bit unsigned integer indicating the weight of the CRLSP. Valid weight values are from 1 to 255. The value 0 means that weight is not applicable for the CRLSP.

Traffic Parameters

Each Traffic Parameter is encoded as a 32 bit IEEE single-precision floating point number. A value of positive infinity is represented as an IEEE single-precision floating-point number with an exponent of all ones (255) and a sign and mantissa of all zeros. The values PDR and CDR are in units of bytes per second. The values PBS, CBS and EBS are in units of bytes.

The value of PDR MUST be greater than or equal to the value of CDR in a correctly encoded Traffic Parameters TLV.

[4.3.1 Semantics](#)

[4.3.1.1 Frequency](#)

The Frequency specifies at what granularity the CDR allocated to the CRLSP is made available. The value VeryFrequently means that the available rate should average at least the CDR when measured over any time interval equal to or longer than the shortest packet time at the CDR. The value Frequently means that the available rate should average at least the CDR when measured over any time interval equal to or longer than a small number of shortest packet times at the CDR. The value Unspecified means that the CDR MAY be provided at any granularity.

[4.3.1.2](#) Peak Rate

The Peak Rate defines the maximum rate at which traffic SHOULD be sent to the CRLSP. The Peak Rate is useful for the purpose of resource allocation. If resource allocation within the MPLS domain depends on the Peak Rate value then it should be enforced at the ingress to the MPLS domain.

The Peak Rate is defined in terms of the two Traffic Parameters PDR and PBS, see [section 4.3.1.5](#) below.

[4.3.1.3](#) Committed Rate

The Committed Rate defines the rate that the MPLS domain commits to be available to the CRLSP.

The Committed Rate is defined in terms of the two Traffic Parameters CDR and CBS, see [section 4.3.1.6](#) below.

[4.3.1.4](#) Excess Burst Size

The Excess Burst Size may be used at the edge of an MPLS domain for the purpose of traffic conditioning. The EBS MAY be used to measure the extent by which the traffic sent on a CRLSP exceeds the committed rate.

The possible traffic conditioning actions, such as passing, marking or dropping, are specific to the MPLS domain.

The Excess Burst Size is defined together with the Committed Rate, see [section 4.3.1.6](#) below.

[4.3.1.5](#) Peak Rate Token Bucket

The Peak Rate of a CRLSP is specified in terms of a token bucket P with token rate PDR and maximum token bucket size PBS.

The token bucket P is initially (at time 0) full, i.e., the token count $T_p(0) = PBS$. Thereafter, the token count T_p , if less than PBS, is incremented by one PDR times per second. When a packet of size B bytes arrives at time t, the following happens:

- o If $T_p(t) - B \geq 0$, the packet is not in excess of the peak rate and T_p is decremented by B down to the minimum value of 0, else
- o the packet is in excess of the peak rate and T_p is not decremented.

Note that according to the above definition, a positive infinite value of either PDR or PBS implies that arriving packets are never in excess of the peak rate.

The actual implementation of a LSR doesn't need to be modeled according to the above formal token bucket specification.

4.3.1.6 Committed Data Rate Token Bucket

The committed rate of a CRLSP is specified in terms of a token bucket C with rate CDR. The extent by which the offered rate exceeds the committed rate MAY be measured in terms of another token bucket E , which also operates at rate CDR. The maximum size of the token bucket C is CBS and the maximum size of the token bucket E is EBS.

The token buckets C and E are initially (at time 0) full, i.e., the token count $T_c(0) = CBS$ and the token count $T_e(0) = EBS$. Thereafter, the token counts T_c and T_e are updated CDR times per second as follows:

- o If T_c is less than CBS, T_c is incremented by one, else
- o if T_e is less than EBS, T_e is incremented by one, else
- o neither T_c nor T_e is incremented.

When a packet of size B bytes arrives at time t , the following happens:

- o If $T_c(t) - B \geq 0$, the packet is not in excess of the Committed Rate and T_c is decremented by B down to the minimum value of 0, else
- o if $T_e(t) - B \geq 0$, the packet is in excess of the Committed Rate but is not in excess of the EBS and T_e is decremented by B down to the minimum value of 0, else
- o the packet is in excess of both the Committed Rate and the EBS and neither T_c nor T_e is decremented.

Note that according to the above specification, a CDR value of positive infinity implies that arriving packets are never in excess of either the Committed Rate or EBS. A positive infinite value of either CBS or EBS implies that the respective limit cannot be

exceeded.

The actual implementation of a LSR doesn't need to be modeled according to the above formal specification.

4.3.1.7 Weight

The weight determines the CRLSP's relative share of the possible excess bandwidth above its committed rate. The definition of "relative share" is MPLS domain specific.

4.3.2 Procedures

4.3.2.1 Label Request Message

If an LSR receives an incorrectly encoded Traffic Parameters TLV in which the value of PDR is less than the value of CDR then it MUST send a Notification Message including the Status code Traffic Parameters Unavailable to the upstream LSR from which it received the erroneous message.

If a Traffic Parameter is indicated as Negotiable in the Label Request Message by the corresponding Negotiable Flag then an LSR MAY replace the Traffic Parameter value with a smaller value.

If the Weight is indicated as Negotiable in the Label Request Message by the corresponding Negotiable Flag then an LSR may adjust replace the Weight value with a lower value (down to 1).

If, after possible Traffic Parameter negotiation, an LSR can support the CRLSP Traffic Parameters then the LSR MUST reserve the corresponding resources for the CRLSP.

If, after possible Traffic Parameter negotiation, an LSR cannot support the CRLSP Traffic Parameters then the LSR MUST send a notification message that contains the Resource Unavailable status code.

4.3.2.2 Label Mapping Message

If an LSR receives an incorrectly encoded Traffic Parameters TLV in which the value of PDR is less than the value of CDR then it MUST send a Label Release message containing the Status code Traffic Parameters Unavailable to the LSR from which it received the erroneous message.

The egress LSR MUST include the (possibly negotiated) Traffic Parameters and Weight in the Label Mapping message.

The Traffic Parameters and the Weight in a Label Mapping message MUST be forwarded unchanged.


An LSR SHOULD adjust the resources that it reserved for a CRLSP when it receives a Label Mapping Message if the Traffic Parameters differ from those in the corresponding Label Request Message.

4.3.2.3 Notification Message

If an LSR receives a Notification Message for a CRLSP, it SHOULD release any resources that it possibly had reserved for the CRLSP.

In addition, on receiving a Notification Message from a Downstream LSR that is associated with a Label Request from an upstream LSR, the local LSR MUST propagate the Notification message using the procedures in [LDP].

4.4 Preemption TLV



U bit
Unknown TLV bit. As defined in [LDP].

F bit
Forward unknown TLV bit. As defined in [LDP].

Type
A fourteen-bit field carrying the value of the Preemption-TLV type which is 0x810.

Length
Specifies the length of the value field in bytes.

Reserved
Zero on transmission. Ignored on receipt.

SetPrio
A SetupPriority of value zero (0) is the priority assigned to the most important path. It is referred to as the highest priority. Seven (7) is the priority for the least important path. The higher the setup priority, the more paths CR-LDP can bump to set up the path.

HoldPrio
A HoldingPriority of value zero (0) is the priority assigned to the most important path. It is referred to as the highest priority. Seven (7) is the priority for the least important path.

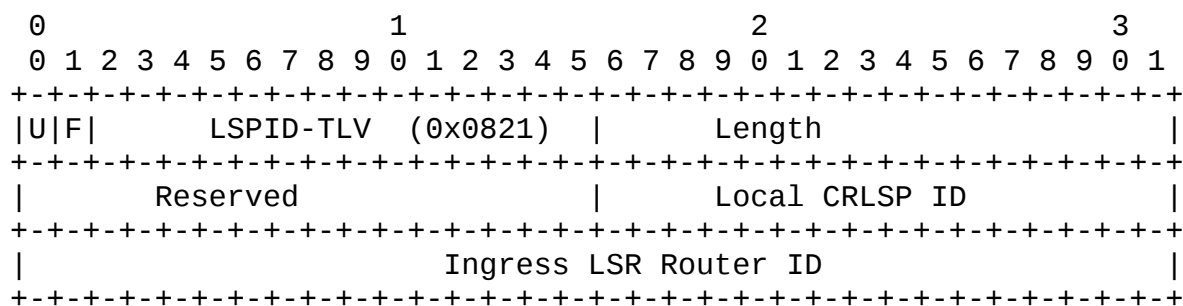
The higher the holding priority, the less likely it is for CR-LDP to reallocate its bandwidth to a new path.

4.5 LSPID TLV

LSPID is a unique identifier of a CRLSP within an MPLS network.

The LSPID is composed of the ingress LSR Router ID and a Locally unique CRLSP ID to that LSR.

The LSPID is useful in network management, in CR-LSP repair, and in using an already established CR-LSP as a hop in an ER-TLV.



U bit

Unknown TLV bit. As defined in [LDP].

F bit

Forward unknown TLV bit. As defined in [LDP].

Type

A fourteen-bit field carrying the value of the LSPID-TLV type which is 0x821.

Length

Specifies the length of the value field in bytes.

Reserved

Zero on transmission. Ignored on receipt.

Local CRLSP ID

The Local LSP ID is an identifier of the CRLSP locally unique within the Ingress LSR originating the CRLDP.


Ingress LSR Router ID

A 4 byte field indicating the Ingress LSR ID.

4.6 Resource Class (Color) TLV

The Resource Class as defined in [TER] is used to specify which links are acceptable by this CRLSP. This information allows for the

networks topology to be pruned.



U bit

Unknown TLV bit. As defined in [LDP].

F bit

Forward unknown TLV bit. As defined in [LDP].

Type

A fourteen-bit field carrying the value of the ResCls-TLV type which is 0x822.

Length

Specifies the length of the value field in bytes.

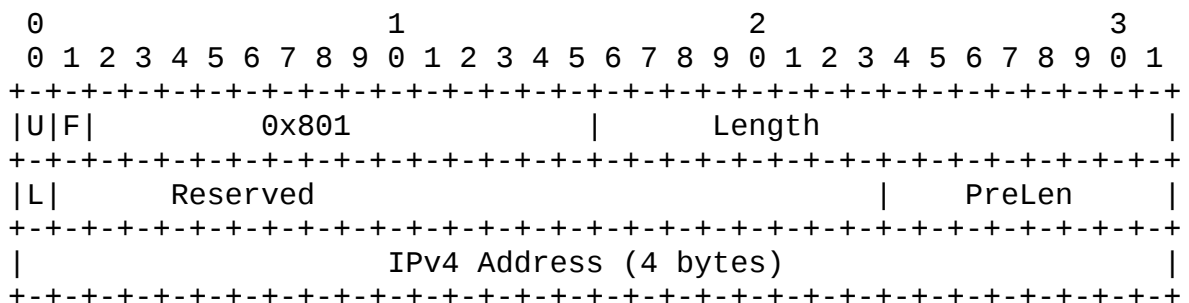
RsC1s

The Resource Class bit mask indicating which of the 32 "administrative groups" or "colors" of links the CRLSP can traverse.

4.7 ER-Hop semantics

4.7.1. ER-Hop 1: The IPv4 prefix

The abstract node represented by this ER-Hop is the set of nodes which have an IP address which lies within this prefix. Note that a prefix length of 32 indicates a single IPv4 node.



U bit

Unknown TLV bit. As defined in [LDP].

F bit

Forward unknown TLV bit. As defined in [LDP].

Type

IPv4 Address 0x801

Length

Specifies the length of the value field in bytes.

L Bit

Set to indicate Loose hop.

Cleared to indicate a strict hop.

Reserved

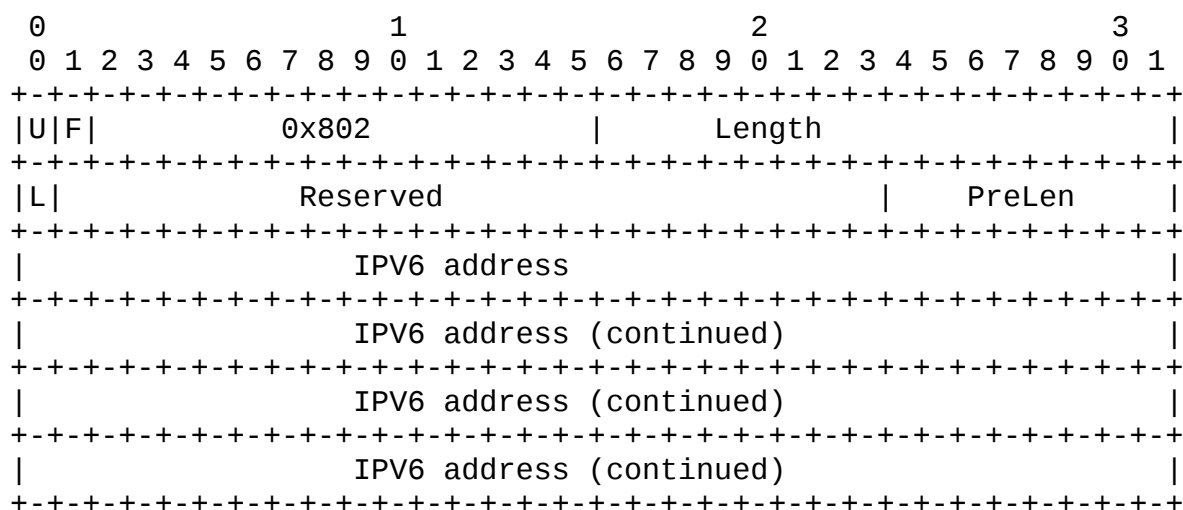
Zero on transmission. Ignored on receipt.

PreLen

Prefix Length 1-32

IP Address

A four byte field indicating the IP Address.

4.7.2. ER-Hop 2: The IPv6 address

U bit

Unknown TLV bit. As defined in [[LDP](#)].

F bit

Forward unknown TLV bit. As defined in [[LDP](#)].

Type

0x802 IPv6 address

Length

Specifies the length of the value field in bytes.

L Bit

Set to indicate Loose hop.

Cleared to indicate a strict hop.

Reserved

Zero on transmission. Ignored on receipt.

PreLen


Prefix Length 1-128

IPv6 address

A 128-bit unicast host address.

4.7.3. ER-Hop 32: The autonomous system number

The abstract node represented by this ER-Hop is the set of nodes belonging to the autonomous system.



U bit

Unknown TLV bit. As defined in [[LDP](#)].

F bit

Forward unknown TLV bit. As defined in [[LDP](#)].

Type

AS Number 0x803

Length

Specifies the length of the value field in bytes.

L Bit

Set to indicate Loose hop.

Cleared to indicate a strict hop.

Reserved

Zero on transmission. Ignored on receipt.

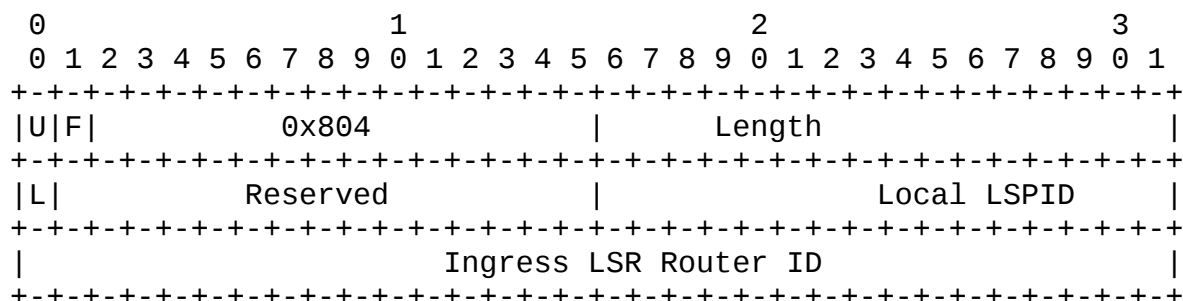
AS Number

Autonomous System number

4.7.4. ER-Hop 4: LSPID

The LSPID is used to identify the tunnel ingress point as the next hop in the ER. This ER-Hop allows for stacking new CR-LSPs within an already established CR-LSP. It also allows for splicing the CR-LSP

being established with an existing CR-LSP.



U bit

Unknown TLV bit. As defined in [[LDP](#)].

F bit

Forward unknown TLV bit. As defined in [[LDP](#)].

Type

LSPID 0x804

Length

Specifies the length of the value field in bytes.

L Bit

Set to indicate Loose hop.

Cleared to indicate a strict hop.

Reserved

Zero on transmission. Ignored on receipt.

Local LSPID

A 2 byte field indicating the LSPID which is unique with reference to the its Ingress LSR.

Ingress LSR Router ID

A 4 byte field indicating the Ingress LSR ID.

[4.8](#). Processing of the Explicit Route TLV

[4.8.1](#). Selection of the next hop

A Label Request Message containing a explicit route TLV must determine the next hop for this path. Selection of this next hop may involve a selection from a set of possible alternatives. The mechanism for making a selection from this set is implementation dependent and is outside of the scope of this specification. Selection of particular paths is also outside of the scope of this specification, but it is assumed that each node will make a best effort attempt to determine a loop-free path. Note that such best

efforts may be overridden by local policy.

To determine the next hop for the path, a node performs the following steps:

- 1) The node receiving the Label Request Message must first evaluate the first ER-Hop. If the L bit is not set in the first ER-Hop and if the node is not part of the abstract node described by the first ER-Hop, it has received the message in error, and should return a "Bad initial ER-Hop" error. If the L bit is set and the local node is not part of the abstract node described by the first ER-Hop, the node selects a next hop that is along the path to the abstract node described by the first ER-Hop. If there is no first ER-Hop, the message is also in error and the system should return a "Bad Explicit Routing TLV" error.
- 2) If there is no second ER-Hop, this indicates the end of the explicit route. The explicit route TLV should be removed from the Label Request Message. This node may or may not be the end of the LSP. Processing continues with [section 4.8.2](#), where a new explicit route TLV may be added to the Label Request Message.
- 3) If the node is also a part of the abstract node described by the second ER-Hop, then the node deletes the first ER-Hop and continues processing with step 2, above. Note that this makes the second ER-Hop into the first ER-Hop of the next iteration.
- 4) The node determines if it is topologically adjacent to the abstract node described by the second ER-Hop. If so, the node selects a particular next hop which is a member of the abstract node. The node then deletes the first ER-Hop and continues processing with [section 4.8.2](#).
- 5) Next, the node selects a next hop within the abstract node of the first ER-Hop that is along the path to the abstract node of the second ER-Hop. If no such path exists then there are two cases:
 - 5a) If the second ER-Hop is a strict ER-Hop, then there is an error and the node should return a "Bad strict node" error.
 - 5b) Otherwise, if the second ER-Hop is a loose ER-Hop, then the node selects any next hop that is along the path to the next abstract node. If no path exists within the MPLS domain, then there is an error, and the node should return a "Bad loose node" error.
- 6) Finally, the node replaces the first ER-Hop with any ER-Hop that denotes an abstract node containing the next hop. This is necessary so that when the explicit route is received by the next hop, it will be accepted.

7) Progress the Label Request Message to the next hop.

4.8.2. Adding ER-Hops to the explicit route TLV


After selecting a next hop, the node may alter the explicit route in the following ways.

If, as part of executing the algorithm in [section 4.8.1](#), the explicit route TLV is removed, the node may add a new explicit route TLV.

Otherwise, if the node is a member of the abstract node for the first ER-Hop, then a series of ER-Hops may be inserted before the first ER-Hop or may replace the first ER-Hop. Each ER-Hop in this series must denote an abstract node that is a subset of the current abstract node.

Alternately, if the first ER-Hop is a loose ER-Hop, an arbitrary series of ER-Hops may be inserted prior to the first ER-Hop.

4.9 Route Pinning TLV



U bit

Unknown TLV bit. As defined in [[LDP](#)].

F bit

Forward unknown TLV bit. As defined in [[LDP](#)].

Type

Pinning-TLV type 0x823

Length

Specifies the length of the value field in bytes.

P Bit

The P bit is set to 1 to indicate that route pinning is requested.
The P bit is set to 0 to indicate that route pinning is not requested

Reserved

Zero on transmission. Ignored on receipt.

4.10 CRLSP FEC Element

Status Code	Type
-----	-----
Bad Explicit Routing TLV Error	0x04000001
Bad Strict Node Error	0x04000002
Bad Loose Node Error	0x04000003
Bad Initial ER-Hop Error	0x04000004
Resource Unavailable	0x04000005
Traffic Parameters Unavailable	0x04000006
Setup abort	0x04000007

5. Security

Pre-emption has to be controlled by the MPLS domain.

Resource reservation requires the LSRs to have an LSP admission control function.

Normal routing can be bypassed by Traffic Engineered LSPs.

6. Acknowledgments

The messages used to signal the CRLSP setup are based on the work done by the [LDP] team. The Explicit Route object and procedures used in this specification are based on [ER].

The authors would also like to acknowledge the careful review and comments of Ken Hayward, Greg Wright, Geetha Brown, Brian Williams, Paul Beaubien, Matthew Yuen, Liam Casey, and Ankur Anand.

7. References

- [LDP] Andersson et al, "Label Distribution Protocol Specification" work in progress ([draft-ietf-mpls-ldp-03](#)), Feb. 1999.
- [ARCH] Rosen et al, "Multiprotocol Label Switching Architecture", work in progress ([draft-ietf-mpls-arch-04](#)), Feb. 1999.
- [FRAME] Callon et al, "Framework for Multiprotocol Label Switching", work in progress ([draft-ietf-mpls-framework-02](#)), November 1997.
- [TER] Awduche et al, "Requirements for Traffic Engineering Over MPLS", work in progress ([draft-ietf-mpls-traffic-eng-00](#)), August 1998.
- [ER] Guerin et al, "Setting up Reservations on Explicit Paths using RSVP", work in progress ([draft-guerin-expl-path-rsvp-01](#))
November 1997.

- [VPN1] Heinanen et al, "MPLS Mappings of Generic VPN Mechanisms", work in progress ([draft-heinanen-generic-vpn-mpls-00](#)), August 1998.
- [VPN2] Jamieson et al, "MPLS VPN Architecture" work in progress ([draft-jamieson-mpls-vpn-00](#)), August 1998.
- [VPN3] T. Li, "CPE based VPNs using MPLS", work in progress ([draft-li-mpls-vpn-00.txt](#)), October 1998.
- [LDP-STATE] L. Wu, et. al., "LDP State Machine" work in progress ([draft-ietf-mpls-ldp-state-00](#)), Feb 1999.

8. Author Information

Osama S. Aboul-Magd
Nortel Networks
P O Box 3511 Station C
Ottawa, ON K1Y 4H7
Canada
phone: +1 613 763-5827
osama@NortelNetworks.com

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

Ram Dantu
Alcatel USA Inc.
IP Competence Center
[1201](#) E. Campbell Road., 446-315
Richardson, TX USA., 75081-2206
Phone: 972 996 2938
Fax: 972 996 5902
ram.dantu@aud.alcatel.com

Nancy Feldman
IBM Corp.
[17](#) Skyline Drive
Hawthorne NY 10532
Phone: 914-784-3254
nkf@us.ibm.com

Eric Gray
Lucent Technologies, Inc
[1600](#) Osgood St.
North Andover, MA 01847
Phone: 603-659-3386
ewgray@lucent.com

Juha Heinanen
Telia Finland, Inc.
Myyrmaentie 2
[01600](#) VANTAA
Finland
Tel: +358 41 500 4808
jh@telia.fi

Loa Andersson
Director Bay Architecture Lab, EMEA
Kungsgatan 34, PO Box 1788
111 97 Stockholm, Sweden
phone: +46 8 441 78 34
mobile +46 70 522 78 34
loa_andersson@baynetworks.com

Ross Callon
IronBridge Networks
55 Hayden Avenue,
Lexington, MA 02173
Phone: +1-781-402-8017
rcallon@ironbridgenetworks.com

Paul Doolan
Ennovate Networks
330 Codman Hill Rd
Marlborough MA 01719
Phone: 978-263-2002
pdoolan@ennovatenetworks.com

Andre Fredette
Nortel Networks
3 Federal Street
Billerica, MA 01821
fredette@baynetworks.com

Joel M. Halpern
Newbridge Networks Inc.
593 Herndon Parkway
Herndon, VA 20170
phone: 1-703-736-5954
jhalpern@newbridge.com

Fiffi Hellstrand
Ericsson Telecom AB
S-126 25 STOCKHOLM
Sweden
Tel: +46 8 719 4933
etxfiff@etxb.ericsson.se

Bilel Jamoussi
Nortel Networks
P O Box 3511 Station C
Ottawa, ON K1Y 4H7
Canada
phone: +1 613 765-4814
jamoussi@NortelNetworks.com

Andrew G. Malis
Ascend Communications, Inc.
[1](#) Robbins Road
Westford, MA 01886
phone: 978 952-7414
fax: 978 392-2074
malis@ascend.com

Kenneth Sundell
Ericsson
SE-126 25 Stockholm
Sweden
kenneth.sundell@etx.ericsson.se

Tom Worster
General DataComm, Inc.
[5](#) Mount Royal Ave.
Marlboro MA 01752
tom.worster@gdc.com

Timothy E. Kilty
Northchurch Communications
5 Corporate Drive,
Andover, MA 018110
phone: 978 691-4656
tkilty@northhc.com

Muckai K Girish
SBC Technology Resources, Inc.
4698 Willow Road
Pleasanton, CA 94588
Phone: (925) 598-1263
Fax: (925) 598-1321
mgirish@tri.sbc.com

Pasi Vaananen
Nokia Telecommunications
3 Burlington Woods Drive, Suite 250
Burlington, MA 01803
Phone: +1-781-238-4981
pasi.vaananen@ntc.nokia.com

Liwen Wu
Alcatel U.S.A
44983 Knoll Square
Ashburn, Va. 20147
USA
Phone: (703) 724-2619
FAX: (703) 724-2005
liwen.wu@adn.alcatel.com

Appendix A: CRLSP Establishment Examples

A.1 Strict Explicit Route Example

This appendix provides an example for the setup of a strictly routed CRLSP. In this example, each abstract node is represented by a specific node.

The sample network used here is a four node network with two edge LSRs and two core LSRs as follows:

```
          a          b          c
LSR1-----LSR2-----LSR3-----LSR4
```

LSR1 generates a Label Request Message as described in [Section 3.1](#) of this draft and sends it to LSR2. This message includes the CR-TLV.

The ER-TLV is composed by a vector of three ER-Hop TLVs <a, b, c>. The ER-Hop TLVs used in this example are of type 0x0801 (IPv4 prefix) with a prefix length of 32. Hence, each ER-Hop TLV identifies a specific node as opposed to a group of nodes.

At LSR2, the following processing of the ER-TLV per [Section 4.8.1](#) of this draft takes place:

- 1) The first hop <a> is part of the abstract node LSR2. Therefore, the first step passes the test. Go to step 2.
- 2) There is a second ER-Hop, . Go to step 3.
- 3) LSR2 is not part of the abstract node described by the second ER-Hop . Go to Step 4.
- 4) LSR2 determines that it is topologically adjacent to the abstract node described by the second ER-Hop . LSR2 selects a next hop (LSR3) which is the abstract node. LSR2 deletes the first ER-Hop <a> from the ER-TLV which now becomes <b , c>. Go to [Section 4.8.2](#).

At LSR2, the following processing of [Section 4.8.2](#) takes place:

Executing algorithm 4.8.1 did not result in the removal of the ER-TLV.

Also, LSR2 is not a member of the abstract node described by the first ER-Hop .

Finally, the first ER-Hop is a strict hop.

Therefore, processing [section 4.8.2](#) does not result in the insertion of new ER-Hops. The selection of the next hop has been

already done is step 4 of [Section 4.8.1](#) and the processing of the ER-TLV is completed at LSR2. In this case, the Label Request Message including the ER-TLV <b, c> is progressed by LSR2 to LSR3.

At LSR3, a similar processing to the ER-TLV takes place except that the incoming ER-TLV = <b, c> and the outgoing ER-TLV is <c>.

At LSR4, the following processing of [section 4.8.1](#) takes place:

- 1) The first hop <c> is part of the abstract node LSR4. Therefore, the first step passes the test. Go to step 2.
- 2) There is no second ER-Hop, this indicates the end of the CRLSP. The ER-TLV is removed from the Label Request Message. Processing continues with [Section 4.8.2](#).

At LSR4, the following processing of [Section 4.8.2](#) takes place:

Executing algorithm 4.8.1 resulted in the removal of the ER-TLV. LSR4 does not add a new ER-TLV.

Therefore, processing [section 4.8.2](#) does not result in the insertion of new ER-Hops. This indicates the end of the CRLSP and the processing of the ER-TLV is completed at LSR4.

At LSR4, processing of [Section 3.2](#) is invoked. The first condition is satisfied (LSR4 is the egress end of the CRLSP and upstream mapping has been requested). Therefore, a Label Mapping Message is generated by LSR4 and sent to LSR3.

At LSR3, the processing of [Section 3.2](#) is invoked. The second condition is satisfied (LSR3 received a mapping from its downstream next hop LSR4 for a CRLSP for which an upstream request is still pending). Therefore, a Label Mapping Message is generated by LSR3 and sent to LSR2.

At LSR2, a similar processing to LSR 3 takes place and a Label Mapping Message is sent back to LSR1 which completes the end-to-end CRLSP setup.

[A.2](#). Node Groups and Specific Nodes Example

A request at an ingress LSR to setup a CRLSP might originate from a management system or an application, the details are implementation specific.

The ingress LSR uses information provided by the management system or the application and possibly also information from the routing database to calculate the explicit route and to create the Label Request Message.

The Label request message carries together with other necessary information a ER-TLV defining the explicitly routed path. In our example the list of hops in the ER-Hop TLV is supposed to contain an abstract node representing a group of nodes, an abstract node representing a specific node, another abstract node representing a group of nodes, and an abstract node representing a specific egress point.

In--{Group 1}--{Specific A}--{Group 2}--{Specific Out: B}

The ER-TLV contains four ER-Hop TLVs:

1. An ER-Hop TLV that specifies a group of LSR valid for the first abstract node representing a group of nodes (Group 1).
2. An ER-Hop TLV that indicates the specific node (Node A).
3. An ER-Hop TLV that specifies a group of LSRs valid for the second abstract node representing a group of nodes (Group 2).
4. An ER-Hop TLV that indicates the specific egress point for the CRLSP (Node B).

All the ER-Hop TLVs are strictly routed nodes.

The setup procedure for this CRLSP works as follows:

1. The ingress node sends the Label Request Message to a node that is a member the group of nodes indicated in the first ER-Hop TLV, following normal routing for the specific node (A).
2. The node that receives the message identifies itself as part of the group indicated in the first ER-Hop TLV, and that it is not the specific node (A) in the second. Further it realizes that the specific node (A) is not one of its next hops.
3. It keeps the ER-Hop TLVs intact and sends a Label Request Message to a node that is part of the group indicated in the first ER-Hop TLV (Group 1), following normal routing for the specific node (A).
4. The node that receives the message identifies itself as part of the group indicated in the first ER-Hop TLV, and that it is not the specific node (A) in the second ER-Hop TLV. Further it realizes that the specific node (A) is one of its next hops.
5. It removes the first ER-Hop TLVs and sends a Label Request Message to the specific node (A).
6. The specific node (A) recognizes itself in the first ER-Hop TLV. Removes the specific ER-Hop TLV.

7. It sends a Label Request Message to a node that is a member of the group (Group 2) indicated in the ER-Hop TLV.
8. The node that receives the message identifies itself as part of the group indicated in the first ER-Hop TLV, further it realizes that the specific egress node (B) is one of its next hops.
9. It sends a Label Request Message to the specific egress node (B).
10. The specific egress node (B) recognizes itself as the egress for the CRLSP, it returns a Label Mapping Message, that will traverse the same path as the Label Request Message in the opposite direction.

[Appendix B](#). QoS Service Examples

[B.1](#) Service Examples

Construction of an end-to-end service is the result of the rules enforced at the edge and the treatment that packets receive at the network nodes. The rules define the traffic conditioning actions that are implemented at the edge and they include policing with pass, mark, and drop capabilities. The edge rules are expected to be defined by the mutual agreements between the service providers and their customers and they will constitute an essential part of the SLA. Therefore edge rules are not included in the signaling protocol.

Packets treatment at a network node is usually referred to as the local behavior. Local behavior could be specified in many ways. One example for local behavior specification is the service frequency introduced in section 4.3.2.1., together with the resource reservation rules implemented at the nodes.

Edge rules and local behaviors can be viewed as the main building blocks for the end-to-end service construction. The following table illustrates the applicability of the building block approach for constructing different services including those defined for ATM.

Service Examples	PDR	PBS	CDR	CBS	EBS	Service Frequency	Conditioning Action
DS	S	S	=PDR	=PBS	0	Frequent	drop>PDR
TS	S	S	S	S	0	Unspecified	drop>PDR, PBS mark>CDR, CBS
BE	inf	inf	inf	inf	0	Unspecified	-
FRS	S	S	CIR	~B_C	~B_E	Unspecified	drop>PDR, PBS mark>CDR, CBS, EBS
ATM-CBR	PCR	CDVT	=PCR	=CDVT	0	VeryFrequent	drop>PCR
ATM-VBR.3(rt)	PCR	CDVT	SCR	MBS	0	Frequent	drop>PCR mark>SCR, MBS
ATM-VBR.3(nrt)	PCR	CDVT	SCR	MBS	0	Unspecified	drop>PCR mark>SCR, MBS
ATM-UBR	PCR	CDVT	-	-	0	Unspecified	drop>PCR
ATM-GFR.1	PCR	CDVT	MCR	MBS	0	Unspecified	drop>PCR

ATM-GFR.2	PCR	CDVT	MCR	MBS	0	Unspecified	drop>PCR mark>MCR, MFS
int-serv-CL	p	m	r	b	0	Frequent	drop>p drop>r, b

S= User specified

In the above table, the DS refers to a delay sensitive service where the network commits to deliver with high probability user datagrams at a rate of PDR with minimum delay and delay requirements. Datagrams in excess of PDR will be discarded.

The TS refers to a generic throughput sensitive service where the network commit to deliver with high probability user datagrams at a rate of at least CDR. The user may transmit at a rate higher than CDR but datagrams in excess of CDR would have a lower probability of being delivered.

The BE is the best effort service and it implies that there are no expected service guarantees from the network.

B.2. Establishing CR-LSP Supporting Real-Time Applications

In this scenario the customer needs to establish an LSP for supporting real-time applications such voice and video. The Delay-sensitive (DS) service is requested in this case.

The first step is the specification of the traffic parameters in the signaling message. The two parameters of interest to the DS service are the PDR and the PBS and their values are specified by the user based on his requirements. Since all the traffic parameters are included in the signaling message, appropriate values must be assigned to all of them. For DS service, the CDR and the CBS values are set equal to the PDR and the PBS respectively. An indication of whether the parameter values are subject to negotiation is flagged.

The transport characteristics of the DS service requires that Frequent frequency to be requested to reflect the real-time delay requirements of the service.

In addition to the transport characteristics, both the network provider and the customer need to agree on the actions enforced at the edge. The specification of those actions is expected to be a part of the service level agreement (SLA) negotiation and is not included in the signaling protocol. For DS service, the edge action is to drop packets that exceed the PDR and the PBS specifications.

The signaling message will be sent in the direction of the ER path and the LSP is established following the normal LDP procedures. Each

LSR applies its admission control rules. If sufficient resources are not available and the parameter values are subject to negotiation, then the LSR could negotiate down either the PDR, the PBS, or both. The new parameters values are echoed back in the Label Mapping Message. LSRs might need to re-adjust their resource reservations based on the new traffic parameter values.

B.3. Establishing CR-LSP Supporting Delay Insensitive Applications

In this example we assume that a throughput sensitive (TS) service is requested. For resource allocation the user assigns values for PDR, PBS, CDR, and CBS. The negotiation flag is set if the traffic parameters are subject to negotiation.

Since the service is delay insensitive by definition, the Unspecified frequency is signaled to indicate that the service frequency is not an issue.

Similar to the previous example, the edge actions are not subject for signaling and are specified in the service level agreement between the user and the network provider.

For TS service, the edge rules might include marking to indicate high discard precedence values for all packets that exceed CDR and the CBS. The edge rules will also include dropping of packets that are do not conform to either PDR and PBS.

Each LSR of the LSP is expected to run its admission control rules and negotiate traffic parameters down if sufficient resources do not exist. The new parameters values are echoed back in the Label Mapping Message. LSRs might need to re-adjust their resources based on the new traffic parameter values.