

Network Working Group  
Internet Draft  
Expiration Date: December 2000

P. Ashwood-Smith  
A. Paraschiv  
Nortel Networks

June 2000

## **MPLS LDP Query Message Description**

[draft-paraschiv-mpls-lsp-query-00.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>.

To view the current status of any Internet-Draft, please check the "1id-abstracts.txt" listing contained in an Internet-Drafts Shadow Directory, see <http://www.ietf.org/shadow.html>.

### Abstract

This document describes the encoding and procedures for two new LDP messages, the Query Message and Query-Reply Message. An LER sends a Query message when it needs to find out information about an LSP. The Query message is sent for an established LSP. The Query message can be used for LDP LSPs as well as for CR-LSPs. The queried data is encoded into the Query-Reply message and partially into the Query Message.



## Contents

<a href="#">1</a>	Introduction .....	<a href="#">3</a>
<a href="#">2</a>	Overview .....	<a href="#">3</a>
<a href="#">2.1</a>	LDP Overview .....	<a href="#">3</a>
<a href="#">2.2</a>	CR-LDP Overview .....	<a href="#">4</a>
<a href="#">3</a>	LDP Message Structure Overview .....	<a href="#">4</a>
<a href="#">4</a>	Query Message .....	<a href="#">5</a>
<a href="#">4.1</a>	Query Message encoding .....	<a href="#">6</a>
<a href="#">4.2</a>	Query Message Procedures .....	<a href="#">7</a>
<a href="#">4.3</a>	Query-Reply Message encoding .....	<a href="#">8</a>
<a href="#">4.4</a>	Query-Reply Message Procedures .....	<a href="#">10</a>
<a href="#">4.5</a>	Query TLV .....	<a href="#">10</a>
<a href="#">4.6</a>	Query Label TLV .....	<a href="#">11</a>
<a href="#">4.7</a>	Query Merge Flags TLV .....	<a href="#">12</a>
<a href="#">4.8</a>	Label TLV .....	<a href="#">12</a>
<a href="#">5</a>	Acknowledgments .....	<a href="#">13</a>
<a href="#">6</a>	References .....	<a href="#">13</a>
<a href="#">7</a>	Author's Addresses .....	<a href="#">14</a>



Changes from previous version:

- o First revision.

## **1. Introduction**

The original Multiprotocol Label Switching (MPLS) architecture [MPLS-ARCH] was been defined to support the forwarding of data based on a label. The MPLS architecture does not assume a single label distribution protocol. A number of different label distribution protocols are being standardized. This draft describes the query mechanism for an LSP or CR-LSP. It specifies procedures and encodings for the two new messages added for the query mechanism. Extensions to the RSVP-TE to provide the same functionality are subject for future study and will be covered in future draft versions.

The two new LDP messages are: Query Message and Query-Reply Message. The following new TLVs are added to accommodate the encodings for the new query messages:

- Query TLV
- Query Label TLV
- Query Merge Flags TLV
- Label TLV.

## **2. Overview**

### **2.1. LDP Overview**

Label Distribution Protocol (LDP) defined in [LDP Specification] contains a set of procedures and messages by which Label Switched Routers (LSR) establish Label Switch Paths (LSP) through a network by mapping network layer routing information directly to data-link layer switched paths. LDP associates a Forwarding Equivalence Class (FEC) with each LSP it creates. The FEC associated with an LSP specifies which packets are mapped to that LSP.

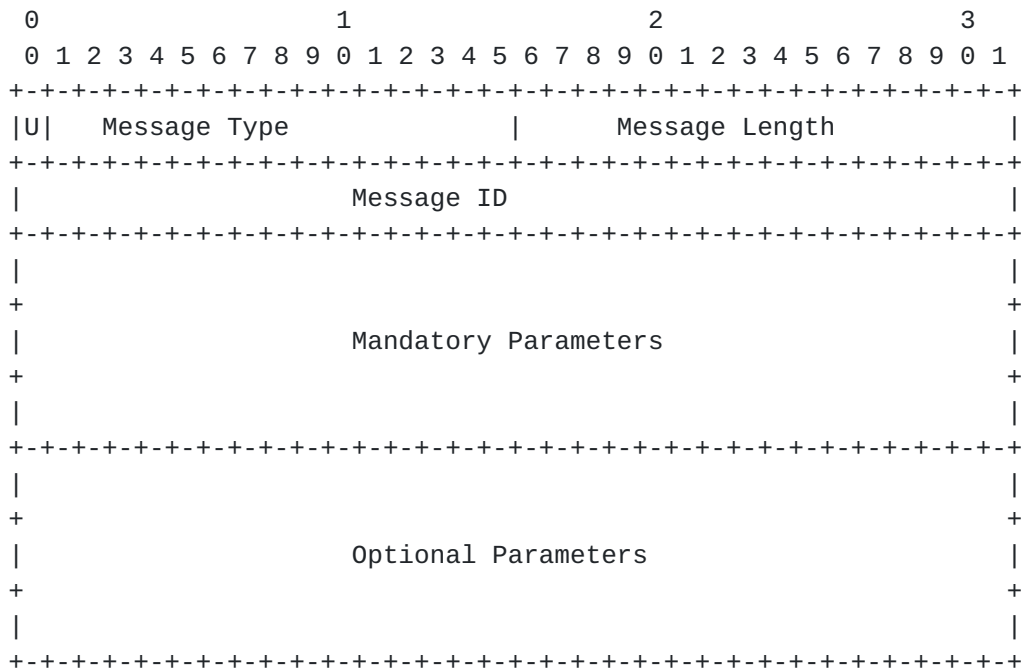


## 2.2. CR-LDP Overview

As described in [Constraint-Base LSP Setup using LDP], Constraint Base Routing (CR-LDP) offers the opportunity to extend the information used to setup paths beyond what is available for the routing protocol.

## 3. LDP Message Structure Overview

As described in LDP Specification draft, all LDP messages have the following format:



### U bit

Unknown message bit. Upon receipt of an unknown message, if U is clear (=0), a notification is returned to the message originator; if U is set (=1), the unknown message is silently ignored. The sections following that define messages specify a value for the U-bit.

### Message Type

Identifies the type of message

### Message Length

Specifies the cumulative length in octets of the Message ID, Mandatory Parameters, and Optional Parameters.





#### Message ID

32-bit value used to identify this message. Used by the sending LSR to facilitate identifying notification messages that may apply to this message. An LSR sending a notification message in response to this message should include this Message Id in the Status TLV carried by the notification message; see Section "Notification Message".

#### Mandatory Parameters

Variable length set of required message parameters. Some messages have no required parameters.

For messages that have required parameters, the required parameters MUST appear in the order specified by the individual message specifications in the sections that follow.

#### Optional Parameters

Variable length set of optional message parameters. Many messages have no optional parameters.

For messages that have optional parameters, the optional parameters may appear in any order.

## 4. Query Message

This sections describes the Query message and its encodings and procedures. This message is meant to be used to gather information about an LSP. It can be sent at any time for an established LSP. The draft currently describes the procedures for the cases when the Query Message is initiated by the ingress LER. Future versions of the draft may add the procedures for the query message when issued from a core LSR or from egress.

The Query Message can be used to gather information about:

- LSRs which form the LSP
- labels along the LSP
- information on what LSRs are merging points along the path
- unused bandwidth (as described in "Improving Topology Data Base Accuracy with LSP Feedback")
- congestion status
- round trip delay
- anything that is needed in the future and can be computed and encoded in a TLV.

The queried information is going to be encoded in the Query-Reply message which is sent back by the egress node, as a response to the



Query message.

#### 4.1. Query Message encoding

The encoding for the Query message 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	2	3	4	5	6	7	8	9																																								
0																Query (0x0409)																																Message Length																															
																Message ID																																																															
																Label TLV																																																															
																Query TLV																																																															
																Hop Count TLV																																																															
																Optional Parameters																																																															

##### Message ID

32-bit value used to identify this message.

##### Label TLV

The label associated to the LSP which is queried. This TLV is a list of Generalized Label TLVs [OPTICAL reference]. The Generalized Label TLV provides a more generic encoding for different types of labels. Most of the time the list has one element; this is the case when the LSP is not tunneled. For tunneled LSPs, the Label TLV has more than one element; it has to behave like a label stack (it contains the previous label and the tunnel's label). See Section Label TLV for more information on Label TLV encoding.

##### Query TLV

What to query. See Section Query TLV for encoding.

##### Hop Count TLV

Specifies the number of LSR hops that can still be traversed before the message is dropped due to loop detection. It is initialized to the max value of 255 (or the configured value, if any). Every LSR that receives the Query Message has to subtract 1 from the Hop Count value. The Query message should be dropped if



the hop count value becomes zero; a Notification signaling Loop Detection should be send in reply to the sender of the message. See [LDP Specification] for Hop Count TLV encoding.

#### Optional Parameters

This variable length field contains 0 or more parameters, each encoded as a TLV.

Optional Parameter		Length	Value
ER	TLV	var	See CR-LDP

The ER TLV is a list of hops. It is used when the Query flag Q3 is set. Every LSR should add its IP address. The address to be added should be the outgoing interface address. Addresses are organized as a last-in-first-out stack (the first address in TLV is considered the top). By carrying this TLV in the Query Message and preserving this order for the hops, we allow the possibility to interwork the Query Message with the RSVP Path message.

### **4.2. Query Message Procedures**

The LER ingress initiates the Query message. It populates the Query TLV Parameters according to what kind of information it wants to gather. The query for an LSP is done by its label. The only data that the Query Message carries is the list of hops. This way, each node along the path will have a complete route from source to destination. This is useful for network management.

Upon receiving a Query Message, an LSR decodes the label to identify which LSP is queried. If it cannot find the LSP which is using the label, it sends back a Notification message. Otherwise, it checks which is the out label which is bound to the queried in-label and which is the downstream LSR peer. It replaces the in-label from Label TLV with the out-label used by the LSP. It then passes the Query message to the downstream peer. When the Query message gets to a tunnel, it has to be able to handle both the previous label and the tunnel's label. The Label TLV behaves like a label stack. The previous label is pushed and the tunnel label is used. At the end of the tunnel, we need to pop the stack and start substituting the lower level labels.

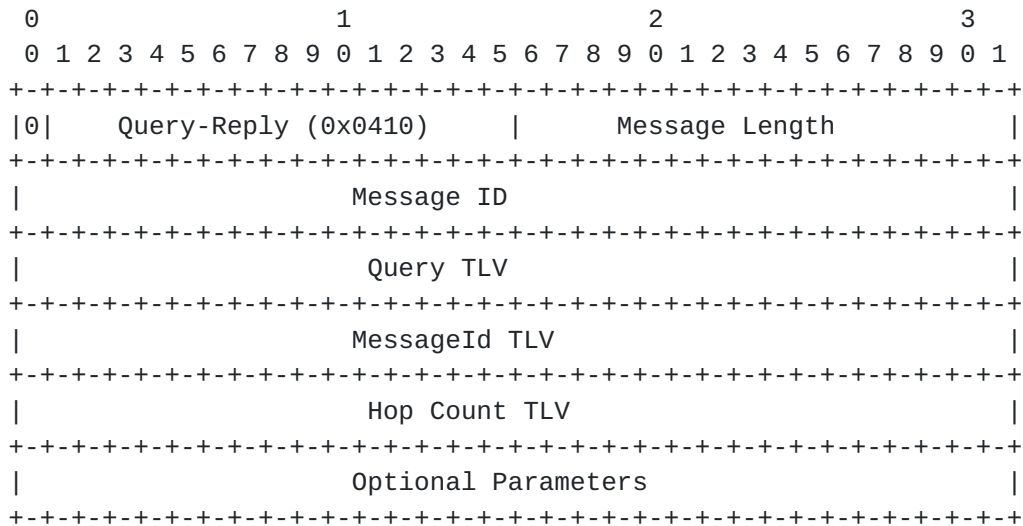
Upon receiving the Query message, the egress node has to reply with a Query Reply Message. This Query-Reply Message contains the Query TLV which was received in the Query Message. The Query TLV tells the egress LER which information is being queried and allows intermediate



LSRs to piggy back their own queried information on the Query reply message.

#### 4.3. Query-Reply Message encoding

The encoding for the Query-Reply message is:



##### Message ID

32-bit value used to identify this message.

##### Query TLV

What is to be queried. See Section Query TLV for encoding.

##### Message Id TLV

The value of this parameter is the message id of the corresponding Query message.

##### Hop Count TLV

Specifies the number of LSR hops that can still be traversed before the message is dropped due to loop detection. It is initialized to the max value of 255 (or the configured value, if any). Every LSR that receives the Query Message has to subtract 1 from the Hop Count value. The Query message should be dropped if the hop count value becomes zero. See [LDP Specification] for Hop Count TLV encoding.

##### Optional Parameters

This variable length field contains 0 or more parameters, each encoded as a TLV. The optional parameters are:





Optional Parameter	Length	Value
ER TLV	var	See CR-LDP
Query Label TLV	var	See Query Label TLV section
IPv4 specified link feedback TLV		See Improving Topology ...
Query Merge Flags TLV	var	See Query Merge Flags TLV section

For simplicity we reuse here few TLV types defined for CR-LDP and LDP. They are:

- IPv4 specified link feedback TLV
- ER TLV
- Generalized Label TLV (used in the Query Label TLV encoding)
- Hop Count TLV.

The IPv4 specified link feedback TLV is used when the Q1 flag from the Query TLV is set. It is used by the egress LER to encode the bandwidth. For more information on query flags, Q1, ... Q6, refer to Query TLV section.

The ER TLV is a list of hops. It is used when the Query flag Q3 is set. Every LSR should add its IP address. The address to be added should be the outgoing interface address. Addresses are organized as a last-in-first-out stack (the first address in TLV is considered the top). By carrying this TLV in the Query-Reply Message and preserving this order for the hops, we allow the possibility to interwork the Query-Reply Message with the RSVP Resv message.

The Query Label TLV is a list of labels. It is used when the Query flag Q2 is set. It is populated with the labels used for the path which is queried. For tunneled LSPs, the Query Label TLV represents a list of labels associated to the lowest level tunnel.

If Q3 and Q2 flags are set, the labels should be encoded in the same order as the hops.

Query Merge Flags TLV is a bit mask. It has variable length and every bit in the mask will correspond to an LSR along the path. Its length is rounded up to the next byte. If Q6 is set, every LSP along the path will have to set its corresponding bit in the mask. The bits have to be set in the same order as the labels and hops. Usually, Q6 is set when Q2 set and/or Q3 set.

For more information for the TLV encodings of the TLVs which are reused, please see CR-LDP draft, LDP draft and IMPROVING TOPOLOGY DATA BASE ACCURACY WITH LSP FEEDBACK VIA CR\_LDP draft.



#### 4.4. Query-Reply Message Procedures

A Query-Reply message is initiated by an egress node which receives a Query message, if the egress is able to identify the queried LSP. If not (maybe the LSP was just torn down), the egress replies with a Notification message.

Upon receiving the Query message, the egress node has to reply with a Query Reply message. The egress node has to encode into the Query-Reply message a MessageId Tlv. The mapping between a Query and a Query-Reply Message is done based on the message id. Besides the MessageId Tlv, the egress has to encode the information that was queried (bandwidth, path, etc).

After the encoding is done, the query reply message is sent back, on the reversed path, toward the ingress. Every LSR across the LSP has to encode its information according to what query flags are set.

#### 4.5. Query TLV

The Query TLV is used to specify the information being queried. The Query TLV travels in the Query message to the egress node, where it is copied into a reverse flowing Query Reply message and used by the egress and intermediate LSRs to know what information is being queried.

The format for the Query 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
+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+
|0|0|  Query TLV  (0x0840)      |      Length      |
+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+
| Query  Flags  |      Reserved      |
+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+

```

Query Flags can be set according to what the Query is used for.

```

+--+--+--+--+--+--+--+--+--+--+
|Re|Q6|Q5|Q4|Q3|Q2|Q1|
+--+--+--+--+--+--+--+--+--+--+

```

They can be:



- Q1 : query the bandwidth; if set, the LSR that receives the Query message has to encode the bandwidth that is available on the link (unused bandwidth);
- Q2 : query the labels which are associated to each hop in the path;
- Q3 : query the LSRs which form the LSP which is queried; if set, the LSR that received the Query-Reply message has to encode the current hop in the ER-TLV
- Q4 : reserved for congestion status; < format - TBD >
- Q5 : query the round trip delay; if set, the edge egress LSR should fill in the delay; < format - TBD >
- Q6 : query which LSPs along the path are merging points; if set, the LSR that receives the Query message has to encode if it is a merging point; the encoding is done in the Query Merge flags TLV.

#### 4.6. Query Label TLV

The Query Label TLV is used to encode the labels used along the path which is queried.

The format for the Query Label 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
+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+
|0|0| Query Label TLV (0x0841)|          Length          |
+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+
|          Generalized Label TLV 1          |
+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+
~
+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+
|          Generalized Label TLV n          |
+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+

```

Generalized Label TLV is used to encode labels along the path. Please refer to [OPTICAL reference] for more information on the Generalized Label TLV encoding. If the Q2 flag is set, every LSR has to encode the out-label corresponding to the queried LSP. In the Query Label TLV, labels are organized as a last-in-first-out stack (the first label in TLV is considered the top). They should be encoded in the same order as the hops and the merge flags.

Note: The Query Label TLV can use the Label TLV which is defined in LDP draft. The same note applies to the Label TLV that encodes the label which is queried.



#### 4.7. Query Merge Flags TLV

The Query Merge Flags TLV is used to encode the information about which LSRs along the path the queried LSP is being merged into.

The format for the Query Label 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
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
|0|0| Merge Flags TLV (0x0842) |      Length      |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
|                                     Number of merge flags                                     |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
| Merge flags   |           |           |           |           | ~
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
~                                                         ~
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
|                                                         |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+

```

The Query Merge Flags TLV has 4 bytes field to store the number of merge flags. This number is equal to the number of LSRs which are traversed by the Query-Reply Message. Each bit in the Merge flags field represents the merge info for an LSR. The bit is set to 0 if the LSR does not do merge for the queried LSP and is set to 1 otherwise. The length is going to be rounded up to the next byte. Every LSR which is asked to encode the merging info into this TLV has to update the Number of merge flags and to set its corresponding bit accordingly. If the Number of merge flags is already multiple of 8, then a new byte needs to be added to this TLV. The length of the TLV has to be updated as well.

#### 4.8. Label TLV

The Label TLV is used to encode the label stack of the labels which are queried (along the queried LSP).





The format for the Query Label 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
+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+
|0|0|          Label TLV (0x0843)|          Length                    |
+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+
|          Generalized Label TLV 1                                   |
+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+
~                                                                    ~
+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+
|          Generalized Label TLV n                                   |
+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+

```

The Label TLV is a stack of Generalized Labels. The top of the stack is in the right octets of the TLV encoding. The Label TLV has to be updated every time the Query message is processed by an LSR. The encoding of the labels should follow the last-in-first-out stack model.

Generalized Label TLV is used to encode labels. Please refer to [OPTICAL reference] for more information on the Generalized Label TLV encoding.

Note: The Label TLV can use the Label TLV which is defined in LDP draft.

## 5. Acknowledgments

The authors would like to acknowledge the careful review and comments of Jean-Pierre Coupal and Steve Hamilton.

## 6. References

- [CR-LDP] Jamoussi et al., "Constraint-Based LSP Setup using LDP", [draft-ietf-mpls-cr-ldp-03.txt](#), October 1999
- [LDP] Andersson et al., "LDP Specification", [draft-ietf-mpls-ldp-07.txt](#), May 2000.
- [RSVP-RR] Berger L., Gan D., Swallow G., Pan P., Tommasi F., Molendini S., "RSVP Refresh Overhead Reduction Extensions", [draft-ietf-rsvp-refresh-reduct-04.txt](#).



[CR-LDP] Peter Ashwood-Smith et al., "Improving Topology Data Base Accuracy with LSP Feedback", [draft-ietf-mpls-te-feed-01.txt](#)

Ashwood-Smith P., Berger L., "Generalized MPLS-Signaling Functional Description" [draft-ashwood-generalized-mpls-signaling-00.txt](#)

## **7. Author's Addresses**

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

Antonela Paraschiv  
Nortel Networks Corp.  
600 Technology Park Drive  
Billerica, MA 01821  
USA  
phone: +1 978-288-6136  
antonela@nortelnetworks.com

## **Full Copyright Statement**

Copyright (C) The Internet Society (date). 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.

