

Network Working Group
Internet Draft
Expiration Date: January 2004

P. Ashwood-Smith
A. Paraschiv
D. Allan
Nortel Networks

June 2003

Multi Protocol Label Switching Label Distribution Protocol Query Message Description

[draft-ietf-mpls-lsp-query-09.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>.

Abstract

This document describes the encoding and procedures for three new Label Distribution Protocol (LDP) messages: Query Message, Query-Reply Message and Partial Query-Reply Message. A Label Edge Router (LER) sends a Query message when it needs to find out information about an established Label Switched Path (LSP). The Query message can be used for LDP LSPs as well as for Constraint-Based Label Switched Paths (CR-LSPs). The queried data is encoded into the Query-Reply messages.

Contents

1.	Introduction.....	3
2.	Specification.....	3
3.	Overview.....	3
3.1	LDP Overview.....	3
3.2	CR-LDP Overview.....	4
4.	LDP Message Structure Overview.....	4
5.	LSRs with constraints in handling the query messages.....	5
5.1	LSR does not support the query messages.....	6
5.2	LSR cannot share any information.....	6
5.3	LSR cannot share some of the queried information.....	6
5.4	LSR can share the queried information.....	6
6.	Query Message.....	7
6.1	Query Message encoding.....	7
6.2	Query Message Procedures.....	9
7.	Reply Messages.....	10
7.1	Query-Reply Message encoding.....	11
7.2	Query-Reply Message Procedures.....	12
7.3	Partial Query-Reply Message encoding.....	13
7.4	Partial Query-Reply Message Procedures.....	13
8.	Query TLVs.....	13
8.1	Query Label TLV.....	14
8.2	Query Merge Flags TLV.....	15
8.3	Query Payload TLV.....	16
8.4	Status code summary.....	16
9.	Security Considerations.....	17
10.	IANA Considerations.....	17
10.1	Message Type Space Extension.....	17
10.2	TLV Type Name Space Extension.....	17
10.3	Status Code Space Extension.....	18
11.	Acknowledgments.....	18
12.	Normative References.....	18
13.	Author's Addresses.....	18

Changes from previous version:

- o Additions to support ECMP
- o Editorial changes
- o Elaboration on security considerations
- o Clarification to message fragmenting issues.

[Editor's note: This section has to be removed prior to publication]

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 memo describes the query mechanism for a Label Switched Path (LSP) or Constraint Based LSP (CR-LSP). It specifies procedures and encodings for the new messages added for the query mechanism.

The new LDP messages are: Query Message, Query-Reply Message and Partial Query-Reply Message. The Partial Query Reply is almost Identical to the Query Reply message; therefore all references to the Query Reply message imply the Partial Query Reply Message as well unless explicitly noted.

The following new TLVs are added to accommodate the encodings for the new query messages:

- Query TLV
- Query Label TLV
- Query Merge Flags TLV
- Query Payload TLV

LDP uses the TCP transport for session, advertisement and notification messages; i.e., for everything but the UDP-based discovery mechanism. The messages which are added to support the query mechanism are sent over TCP as well.

2. Specification

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](#).

3. Overview

3.1 LDP Overview

Paraschiv, et. al.

Expires December 2003

Page 3

Internet Draft

draft-ietf-mppls-lsp-query-09.txt

June 2003

Label Distribution Protocol (LDP) defined in [4] 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.

3.2 CR-LDP Overview

As described in CR-LDP [1], Constraint Base Routing offers the opportunity to extend the information used to setup paths beyond what is available from the routing protocol. For instance, an LSP can be setup based on explicit route constraints, QoS constraints, and other constraints.

4. LDP Message Structure Overview

The LDP message format is specified in the LDP Specification [4]. All LDP messages have 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										
U										Message Type																				Message Length											
										Message ID																															
										Mandatory Parameters																															

```

+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
|
+
|
+
|
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
Optional Parameters
|
+
|
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+

```

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.

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. It is 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 MessageId in the Status TLV carried by the 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.

5. LSRs with constraints in handling the query messages

Upon receiving a Query message, an LSR has to behave according to its configuration constraints in handling the query messages and returning the queried information. The following cases were identified:

- The LSR does not support the code to handle the messages for the query mechanism
- The LSR supports the code to handle the messages for the query mechanism, but it is configured not to return any data
- The LSR supports the code to handle the messages for the query mechanism, but it is configured not to return part of the queried data
- The LSR supports the code to handle the messages for the query mechanism, and it is configured to return all the data which is queried.

This memo provides flexibility to handle each of the above cases.

5.1 LSR does not support the query messages

In this case, the LSR has to behave as if it received an unknown message type. It therefore honors the U bit. Where this is in response to query with Query flag Q8 clear, the originating LSR may then repeat the request with Q8 set (partial repliers requested) to obtain LSP information up to the non-implementing LSR along the LSP path.

5.2 LSR cannot share any information

In this case, the LSR is able to decode and process the query messages. However, it is configured to hide all the data. It should propagate the message after it encodes a zero-length TLV for its hop in the hop list in the Query message. When Query-Reply message is received from downstream, the LSR is requested to propagate the reply message upstream after it encodes the zero-length TLVs for the queried data. When the ingress receives back the reply, it can identify which TLVs are empty; it can therefore ignore the zero-length TLVs and process the rest of the data.

Note: zero-length TLV encoding can be used for all types of queried

information except the merge information. The LSR is requested to signal the fact that the merging information is private by encoding a special value in the corresponding merge bits (for more information on the merge flags values please refer to [Section 7.2](#) of this memo - Query Merge Flags TLV).

[5.3](#) LSR cannot share some of the queried information

In this case, the LSR is able to decode and process the query messages. It has to propagate the query messages. It has to encode values for the data types that it is willing to return and zero-length TLVs for values for the data that is hidden.

Note: zero-length TLV encoding can be used for all types of queried information except the merge information. The LSR is requested to signal the fact that the merging information is private by encoding a special value in the corresponding merge bits (for more information on the merge flags values please refer to [Section 7.2](#) of this memo - Query Merge Flags TLV).

[5.4](#) LSR can share the queried information

In this case, the LSR's behavior has to follow the query and replies procedures described in the following sections of this memo.

In order to have consistency among data encoded in the query and reply messages, each LSR which can propagate the messages has to encode its information in the query and in the reply messages.

Paraschiv, et. al.

Expires December 2003

Page 6

Internet Draft

[draft-ietf-mpls-lsp-query-09.txt](#)

June 2003

The decision that an LSR can share the queried information has to be controlled through configuration flags. This way each node along the path can protect its data if they consider it private.

Note: It would be more efficient to control/restrict the private data per MPLS cloud (inter MPLS domain) and not per LSR node (inter and intra MPLS domain). When there are different MPLS clouds which have nodes belonging to different vendors, the control of the private information could be restricted to the boundary nodes. Within an MPLS domain, there should be no restrictions on the queried information. It would be useful to have some knowledge on which are the nodes on the boundaries and have only those hiding the queried data. Because there is no mechanism to identify which are the boundary nodes, this is subject for future study.

6. Query Message

This section 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. This memo currently describes the procedures for the cases when the Query Message is initiated by the ingress LER. Additional procedures may be added in the future for the query message when issued from an 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" [\[2\]](#))
- What path any specific payload may follow in the presence of load spreading mechanisms such as ECMP
- Anything that is needed in the future and can be computed and encoded in a TLV.

It should be noted that useful information can still be extracted from partial responses due constraints on handling messages placed on some LSRs along the LSP. In the fault diagnosis scenario (particularly in the presence of ECMP), those LSRs that at a minimum return their own LSR ID facilitate diagnosis when combined with data plane tracing tools as those LSRs can act as segmentation boundaries for fault isolation.

The queried information is encoded in the Query-Reply message which is sent back upstream, as a response to the Query message.

6.1 Query Message encoding

Paraschiv, et. al.

Expires December 2003

Page 7

Internet Draft

draft-ietf-mppls-lsp-query-09.txt

June 2003

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 type TBD IANA																Message Length																			
+-+-+-----+-+-+-----+-+-+-----+-+-+-----+-+-+-----+-+-+-----+-+-+-----+																																							
				Message ID																																			
+-+-+-----+-+-+-----+-+-+-----+-+-+-----+-+-+-----+-+-+-----+-+-+-----+																																							
				Query Label TLV																																			
+-+-+-----+-+-+-----+-+-+-----+-+-+-----+-+-+-----+-+-+-----+-+-+-----+																																							


```

|                                     Query TLV                                     |
+-+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+
|                                     Hop Count TLV                               |
+-+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+
|                                     Optional Parameters                           |
+-+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+

```

Message ID
 32-bit value used to identify this message.

Query Label TLV
 The label associated to the LSP which is queried. This TLV is a list of Generalized Label TLVs [3]. 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 8.1](#) of this memo - Query Label TLV - for more information on Label TLV encoding.

Query TLV
 What to query. See [Section 8](#) of this memo - 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 default 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 sent in reply to the sender of the message. See [4] 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
Paraschiv, et. al.	Expires December 2003		Page 8
Internet Draft	draft-ietf-mpls-lsp-query-09.txt		June 2003
ER TLV	var	See [1]	
LSPID TLV	var	See [1]	
FEC TLV	var	See [4]	
Query Payload TLV	var	See below	

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.

FEC TLV is used when penultimate hop popping (PHP) is in use, for LDP.

LSPID TLV is used when PHP is in use, for CR-LDP.

Query Payload TLV is used where ECMP is deployed in order to qualify the specific LSP hops chosen by intermediate LSRs.

For more details on the FEC, LSPID, or Query Payload TLVs usage, please refer to the Query Message Procedures below.

6.2 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 could carry is the list of hops. This way, each node along the path can have a complete route from source to destination. This is useful for network management. Please note that this parameter is optional. If the Query message does not contain the ER TLV, it should be propagated by LSRs along the path without the ER TLV.

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 with "No LSP to query" status. 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 Query Label TLV with the out-label used by the LSP. It then passes the Query message to the downstream peer.

Where there is more than one valid downstream peer, AND the Payload Query TLV is present, the LSR may use the Payload Query TLV to select the appropriate downstream LSP peer using the same algorithms that it would use if the payload was carried in the LSP data plane. If the Payload Query TLV is not present but the ECMP implementation requires the payload information for downstream path selection, it

sends back a Notification message with "Payload Query TLV Required" status.

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 Query 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. The Query-Reply Message contains the Query TLV which was received in the Query Message. The Query TLV tells the LSRs along the path which information is being queried and allows intermediate LSRs to piggy back their own queried information on the Query reply message.

The initiator of a Query reply might not receive a response back. In this case, it is the initiator's responsibility to decide if and when to retry.

If PHP is in use, the Query Message sent from the PHP to the destination node needs to carry the FEC (for LDP) or LSPID (for CR-LDP). This information is required because the label between the PHP and the destination is a special label and the destination cannot uniquely identify the queried LSP just by using the label value. If the PHP does not encode the FEC or the LSPID, the destination node should reply with a notification message with "Ambiguous label value".

7. Reply Messages

These messages are propagated upstream. There are two types of reply messages:

- Query-Reply message (final reply)
- Partial Query-Reply message.

The Reply messages carry the queried information upstream. A Query-Reply message is sent in response to a Query message. The ingress which initiated the Query message is interested to gather the information from all the nodes along the queried LSP. However, there are situations in which the Query message does not reach the end point of the queried LSP. In these scenarios it would be useful if the ingress LSR gathered at least some information about the LSRs which are along the path, up to the one that failed. The Partial Query-Reply message provides this mechanism. It is recommended to use the Partial Query-Reply messages when a Query message fails.

Both reply messages are described in the following sections.

7.1 Query-Reply Message encoding

This message is generated by the end point of the LSP. It is propagated upstream, by each LSR along the path.

The encoding for the Query-Reply 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
+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+
|0| Query-Reply Type TBD IANA |      Message Length      |
+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+
|
|      Message ID
|
+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+
|
|      Query TLV
|
+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+
|
|      MessageId TLV
|
+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+
|
|      Optional Parameters
|
+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+

```

Message ID

32-bit value used to identify this message.

Query TLV

What is to be queried. See [Section 7](#) of this memo - Query TLV - for encoding.

MessageId TLV

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

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 [1]
Query Label TLV	var	See Query Label TLV section
IPv4/6 specified link feedback TLV	var	See [2]

Query Merge Flags TLV	var	See Query Merge Flags TLV section
-----------------------	-----	-----------------------------------

The TLV types are defined in CR-LDP [1] and LDP [4]. They are:

- IPV4/6 specified link feedback TLV
- ER TLV

Paraschiv, et. al.	Expires December 2003	Page 11
Internet Draft	draft-ietf-mpls-lsp-query-09.txt	June 2003

- Generalized Label TLV (used in the Query Label TLV encoding)
- Hop Count TLV.

The IPV4/6 specified link feedback TLV is used when the Q1 flag from the Query TLV is set. It is used to encode the bandwidth information. For more information on query flags, Q1, Q2, Q3 and Q4, 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).

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 list of pairs of bits. It has variable length and every two bits in the mask will correspond to an LSR along the path. Its length is rounded up to the next byte. If Q4 is set, every LSP along the path will have to set its corresponding bits in the mask. The bits have to be set in the same order as the labels and hops. Usually, Q4 is set when Q2 set and/or Q3 set.

For more information for the TLV encodings of the TLVs which are used, please see [1], [4] and [2].

[7.2 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, the egress replies with a Notification message with "No LSP

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).

Paraschiv, et. al.	Expires December 2003	Page 12
Internet Draft	draft-ietf-mpls-lsp-query-09.txt	June 2003

The Partial Query-Reply message is initiated by LSRs along the queried path. The message is generated only if the following rules apply:

- The encoding for the Partial Query-Reply message is identical to the Query-Reply, except the message type. For more details on the encoding please refer to the Query-Reply encoding.

The procedures are similar to the Query-Reply's procedures. Upon

receiving a Query message, an LSR will check the flag from the Query message (Q8) which signals if the partial replies are requested by the ingress node. If the flag is set, the LSR has to check next if it is configured to fulfill this request. If the LSR supports partial replies, it has to create a Partial Query-Reply and encode the queried data and send it upstream like any Query-Reply messages. It has then to process the Query message according to the Query message procedures. When an LSR receives a Partial Query-Reply from upstream, it should encode its information according to what is queried and propagate the message. It is recommended to use the Partial Query mechanism when the Query message fails (when the ingress LER does not receive a Query-Reply message in response to a Query message).

8. Query TLVs

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 messages and used by

Paraschiv, et. al.

Expires December 2003

Page 13

Internet Draft

[draft-ietf-mpls-lsp-query-09.txt](#)

June 2003

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  Type TBD IANA |          Length          |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
| Query  Flags  |          Reserved          |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+

```

Query Flags can be set according to what the Query is used for. A query flag is set when it is 1.

```

+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
|Q8|Reserved|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 : 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.
- Q8 : if set, the ingress requests partial query-replies;
each LSR along the path is signaled to send a Partial Query Reply.

The reserved bits need to be set to zero on transmission and must be ignored on receipt. They might be used in the future to signal other types of queried information. The Query Flags can only be defined by updating this memo.

8.1 Query Label TLV

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

Note: Query Label TLV is used in both Query and Query Reply. It is a required parameter in the Query message and it is an optional parameter in Query Reply message. When being used in the Query message, it carries the label or stack of labels which are being

Paraschiv, et. al.

Expires December 2003

Page 14

Internet Draft

[draft-ietf-mpls-lsp-query-09.txt](#)

June 2003

followed and queried. When being used in the Query Reply message, it carries the list of labels which make up the queried path.

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 TBD IANA |          Length          |
+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+
|          Generalized Label TLV 1          |
+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+
~
+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+
|          Generalized Label TLV n          |
+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+

```

Generalized Label TLV is used to encode labels along the path. Please refer to [3] 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.

8.2 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 Merge Flags 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 TBD IANA |          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. The Merge flags field contains the merge information. It is a variable length field which is rounded up to the next byte. Each pair of bits in the Merge flags

field carries the merge information for one LSR. The valid values for the merge bits for an LSR are:

- 01 - LSR does not do merge for the queried LSP
- 10 - LSR does merge for the queried LSP
- 00 - LSR cannot share the merge information.

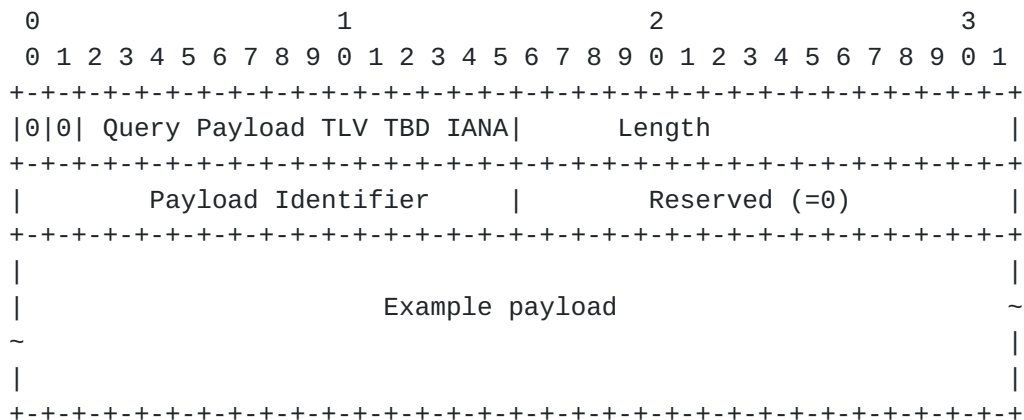
Every LSR which is asked to encode the merging info has to update the number of merge flags and to set its corresponding bits accordingly.

8.3 Query Payload TLV

The Query Payload TLV is used to encode the payload that LSRs

implementing ECMP or similar load spreading should use when selecting a downstream peer.

The format for the Query Payload TLV is:



The Query Payload TLV carries an example of the payload of the LSP that is of interest for tracing purposes. The payload example may be:

- an MPLS label stack,
- an MPLS label stack followed by an IPv4 or IPv6 packet header or
- an IPv4 or IPv6 packet header.

The contents of the Query Payload TLV are identified via use of the payload identifier (currently used as a bitmap)

- 0x01 - label stack present
- 0x02 - IP packet present

When the IP packet present flag is set, the IP version is determined by examining the version number field of the packet header in the example payload.

8.4 Status code summary

Status Code	E	Status Data	Section Title
Paraschiv, et. al.		Expires December 2003	Page 16
Internet Draft		draft-ietf-mpls-lsp-query-09.txt	June 2003
No LSP to query	1	TBD IANA	"Query Message Procedures..."
Ambiguous label value	1	TBD IANA	"Query Message Procedures..."
Payload Query TLV Required	1	TBD IANA	"Query Message Procedures..."
Message space exhausted	1	TBD IANA	"Query Message Procedures..."

9. Security Considerations

The Query mechanism inherits the same security mechanism described in Section 5.0 of [4]. The Query mechanism provides an additional security measure for cases when a node cannot share the queried information. Such nodes have the option of hiding their information, if their configuration requires it. Please refer to [Section 4](#) of this memo - "Behavior of LSRs with constraints in handling the query messages" - for more details.

The Query mechanism focuses on the ability to trace the control plane, the outcome of which may then be compared against the output of data plane tracing tools to ensure system consistency. If the control plane maintains a chain of trust from LSP ingress to egress using mechanisms to secure individual adjacencies against known or yet to be developed attacks, then LSP Query will produce useful information. Where the chain of trust is incomplete, the partial-query reply mechanism may be used to trace the control plane up to the point of the boundary between trusted and untrusted network elements on the LSP path.

10. IANA Considerations

[RFC 3036](#) [4] defines several name spaces including the Message Type Name Space, the TLV Type Name Space, and the Status Code Name Space. This document makes the following assignments within those spaces.

10.1 Message Type Space Extension

The message types for Query Message, Query-Reply Message and Partial Query-Reply Message are as follows:

Message	Type
-----	-----
Query Message	TBD IANA
Query Reply Message	TBD IANA
Partial Query Reply Message	TBD IANA

10.2 TLV Type Name Space Extension

The TLV types for Query TLV, Query Label TLV and Query Merge Flags

Paraschiv, et. al. Expires December 2003 Page 17

Internet Draft [draft-ietf-mpls-lsp-query-09.txt](#) June 2003

TLV are as follows:

TLV	Type
-----	------

-----	-----
Query TLV	TBD IANA
Query Label TLV	TBD IANA
Query Merge Flags TLV	TBD IANA
Query Payload TLV	TBD IANA

10.3 Status Code Space Extension

The Status codes are as follows:

Status Code	Type
-----	-----
No LSP to query	TBD IANA
Ambiguous label value	TBD IANA
Payload Query TLV Required	TBD IANA
Message space exhausted	TBD IANA

11. Acknowledgments

The authors would like to acknowledge the careful review and comments of Jean-Pierre Coupal, Steve Hamilton, Don Fedyk, Gregory Wright and Adrian Farrel.

12. Normative References

- [1] B. Jamoussi et al., "Constraint-Based LSP Setup using LDP", [RFC 3212](#), January 2002
- [2] Peter Ashwood-Smith et al., "Improving Topology Data Base Accuracy with LSP Feedback", Work in Progress, [draft-ietf-mpls-te-feed-03.txt](#).
- [3] Ashwood-Smith P., Berger L., "Generalized MPLS - Signaling Functional Description", Work in Progress, [draft-ietf-mpls-generalized-signaling-07.txt](#).
- [4] Andersson et al., "LDP Specification", [RFC 3036](#), January 2001.

13. 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

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

petera@nortelnetworks.com

antonela@nortelnetworks.com

Dave Allan
Nortel Networks Corp.
P.O. Box 3511 Station C,
Ottawa, ON K1Y 4H7
Canada
Phone: +1 613-763-6362
dallan@nortelnetworks.com

Full Copyright Statement

Copyright (C) The Internet Society (2002). 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.

