

Network Working Group
Request for Comments: 5561
Updates: 5036
Category: Standards Track

B. Thomas
K. Raza
Cisco Systems, Inc.
S. Aggarwal
R. Aggarwal
Juniper Networks
JL. Le Roux
France Telecom
July 2009

LDP Capabilities

Abstract

A number of enhancements to the Label Distribution Protocol (LDP) have been proposed. Some have been implemented, and some are advancing toward standardization. It is likely that additional enhancements will be proposed in the future. This document defines a mechanism for advertising LDP enhancements at session initialization time, as well as a mechanism to enable and disable enhancements after LDP session establishment.

Status of This Memo

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

Copyright Notice

Copyright (c) 2009 IETF Trust and the persons identified as the document authors. All rights reserved.

This document is subject to BCP 78 and the IETF Trust's Legal Provisions Relating to IETF Documents in effect on the date of publication of this document (<http://trustee.ietf.org/license-info>). Please review these documents carefully, as they describe your rights and restrictions with respect to this document.

This document may contain material from IETF Documents or IETF Contributions published or made publicly available before November 10, 2008. The person(s) controlling the copyright in some of this material may not have granted the IETF Trust the right to allow modifications of such material outside the IETF Standards Process. Without obtaining an adequate license from the person(s) controlling

the copyright in such materials, this document may not be modified outside the IETF Standards Process, and derivative works of it may not be created outside the IETF Standards Process, except to format it for publication as an RFC or to translate it into languages other than English.

Table of Contents

1. Introduction	2
1.1. Conventions Used in This Document	3
2. The LDP Capability Mechanism	3
2.1. Capability Document	4
3. Specifying Capabilities in LDP Messages	4
3.1. Backward Compatibility TLVs	6
4. Capability Message	6
5. Note on Terminology	7
6. Procedures for Capability Parameters in Initialization Messages	7
7. Procedures for Capability Parameters in Capability Messages	8
8. Extensions to Error Handling	9
9. Dynamic Capability Announcement TLV	9
10. Backward Compatibility	10
11. Security Considerations	10
12. IANA Considerations	11
13. Acknowledgments	11
14. References	11
14.1. Normative References	11
14.2. Informative References	11

1. Introduction

A number of enhancements to LDP as specified in [RFC5036] have been proposed. These include LDP Graceful Restart [RFC3478], Fault Tolerant LDP [RFC3479], multicast extensions [MLDP], signaling for Layer 2 circuits [RFC4447], a method for learning labels advertised by next-next-hop routers in support of fast reroute node protection [NNHOP], upstream label allocation [UPSTREAM_LDP], and extensions for signaling inter-area Label Switched Paths (LSPs) [RFC5283]. Some have been implemented, and some are advancing toward standardization. It is also likely that additional enhancements will be implemented and deployed in the future.

This document proposes and defines a mechanism for advertising LDP enhancements at session initialization time. It also defines a mechanism to enable and disable these enhancements after LDP session establishment.

LDP capability advertisement provides means for an LDP speaker to announce what it can receive and process. It also provides means for a speaker to inform peers of deviations from behavior specified by [RFC5036]. An example of such a deviation is LDP Graceful Restart, where a speaker retains MPLS forwarding state for LDP-signaled LSPs when its LDP control plane goes down. It is important to point out that not all LDP enhancements require capability advertisement. For example, upstream label allocation requires capability advertisement, but inbound label filtering, where a speaker installs forwarding state for only certain Forwarding Equivalence Classes (FECs), does not.

1.1. Conventions Used in This Document

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].

This document uses the terms "LDP speaker" and "speaker" interchangeably.

2. The LDP Capability Mechanism

Enhancements are likely to be announced during LDP session establishment as each LDP speaker advertises capabilities corresponding to the enhancements it desires.

Beyond that, capability advertisements may be used to dynamically modify the characteristics of the session to suit the changing conditions. For example, an LSR capable of a particular enhancement in support of some "feature" may not have advertised the corresponding capability to its peers at session establishment time because the feature was disabled at that time. Later, an operator may enable the feature, at which time the LSR would react by advertising the corresponding capability to its peers. Similarly, when an operator disables a feature associated with a capability, the LSR reacts by withdrawing the capability advertisement from its peers.

The LDP capability advertisement mechanism operates as follows:

- Each LDP speaker is assumed to implement a set of enhancements, each of which has an associated capability. At any time, a speaker may have none, one, or more of those enhancements "enabled". When an enhancement is enabled, the speaker advertises the associated capability to its peers. By advertising the capability to a peer, the speaker asserts that it shall perform the protocol actions specified for the associated enhancement. For example, the actions

may require the LDP speaker to receive and process enhancement-specific messages from its peer. Unless the capability has been advertised, the speaker will not perform protocol actions specified for the corresponding enhancement.

- At session establishment time, an LDP speaker MAY advertise a particular capability by including an optional parameter associated with the capability in its Initialization message.
- There is a well-known capability called Dynamic Capability Announcement that an LDP speaker MAY advertise in its Initialization message to indicate that it is capable of processing capability announcements following a session establishment.

If a peer had advertised the Dynamic Capability Announcement capability in its Initialization message, then at any time following session establishment, an LDP speaker MAY announce changes in its advertised capabilities to that peer. To do this, the LDP speaker sends the peer a Capability message that specifies the capabilities being advertised or withdrawn.

2.1. Capability Document

When the capability advertisement mechanism is in place, an LDP enhancement requiring LDP capability advertisement will be specified by a document that:

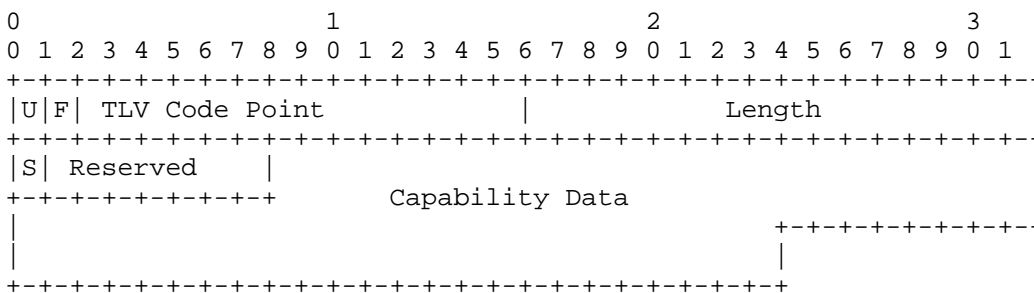
- Describes the motivation for the enhancement;
- Specifies the behavior of LDP when the enhancement is enabled. This includes the procedures, parameters, messages, and TLVs required by the enhancement;
- Includes an IANA considerations section that requests IANA assignment of a code point (from TLV Type namespace) for the optional capability parameter corresponding to the enhancement.

The capability document MUST also describe the interpretation and processing of associated capability data, if present.

3. Specifying Capabilities in LDP Messages

This document uses the term "Capability Parameter" to refer to an optional parameter that may be included in Initialization and Capability messages to advertise a capability.

The format of a "Capability Parameter" TLV is as follows:



where:

U-bit:

Unknown TLV bit, as described in [RFC5036]. The value could be either 0 or 1 as specified in the Capability document associated with the given capability.

F-bit:

Forward unknown TLV bit, as described in [RFC5036]. The value of this bit MUST be 0 since a Capability Parameter TLV is sent only in Initialization and Capability messages, which are not forwarded.

TLV Code Point:

The TLV type that identifies a specific capability. This is an IANA-assigned code point (from TLV Type namespace) for a given capability as requested in the associated capability document.

S-bit:

The State Bit. It indicates whether the sender is advertising or withdrawing the capability corresponding to the TLV code point. The State Bit value is used as follows:

- 1 - The TLV is advertising the capability specified by the TLV code point.
- 0 - The TLV is withdrawing the capability specified by the TLV code point.

Capability Data:

Information, if any, about the capability in addition to the TLV code point required to fully specify the capability.

The method for interpreting and processing this data is specific to the TLV code point and MUST be described in the document specifying the capability.

An LDP speaker MUST NOT include more than one instance of a Capability Parameter (as identified by the same TLV code point) in an Initialization or Capability message. If an LDP speaker receives more than one instance of the same Capability Parameter type in a message, it SHOULD send a Notification message to the peer before terminating the session with the peer. The Status Code in the Status TLV of the Notification message MUST be Malformed TLV value, and the message SHOULD contain the second Capability Parameter TLV of the same type (code point) that is received in the message.

3.1. Backward Compatibility TLVs

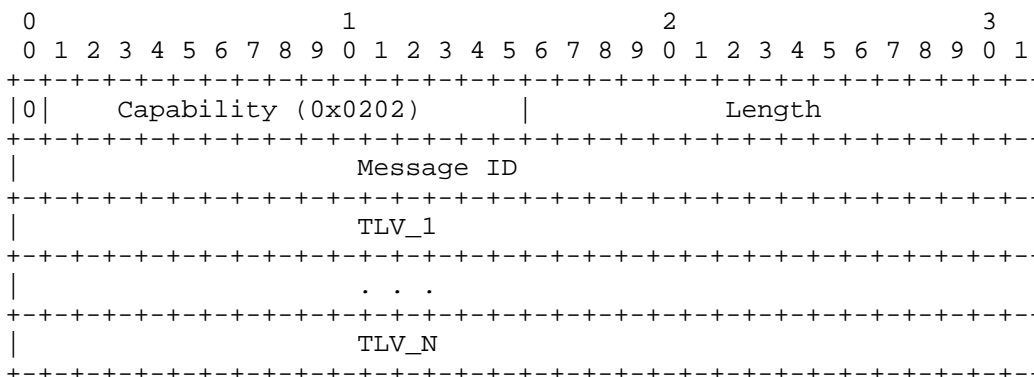
LDP extensions that require advertisement or negotiation of some capability at session establishment time typically use TLVs that are included in an Initialization message. To ensure backward compatibility with existing implementations, such TLVs continue to be supported in an Initialization message and are known in this document as "Backward Compatibility TLVs". A Backward Compatibility TLV plays the role of a "Capability Parameter" TLV; that is, the presence of a Backward Compatibility TLV has the same meaning as a Capability Parameter TLV with the S-bit set for the same capability.

One example of a Backward Capability TLV is the "FT Session TLV" that is exchanged in an Initialization message between peers to announce LDP Fault Tolerance [RFC3479] capability.

4. Capability Message

The LDP Capability message is used by an LDP speaker to announce changes in the state of one or more of its capabilities subsequent to session establishment.

The format of the Capability message is as follows:



where TLV_1 through TLV_N are Capability Parameter TLVs. The S-bit of each of the TLVs specifies the new state for the corresponding capability.

Note that Backward Compatibility TLVs (see Section 3.1) MUST NOT be included in Capability messages. An LDP speaker that receives a Capability message from a peer that includes Backward Compatibility TLVs SHOULD silently ignore these Backward Compatibility TLVs and continue processing the rest of the message.

5. Note on Terminology

The following sections in this document talk about enabling and disabling capabilities. The terminology "enabling (or disabling) a capability" is short hand for "advertising (or withdrawing) a capability associated with an enhancement". Bear in mind that it is an LDP enhancement that is being enabled or disabled, and that it is the corresponding capability that is being advertised or withdrawn.

6. Procedures for Capability Parameters in Initialization Messages

The S-bit of a Capability Parameter in an Initialization message MUST be 1 and SHOULD be ignored on receipt. This ensures that any Capability Parameter in an Initialization message enables the corresponding capability.

An LDP speaker determines the capabilities of a peer by examining the set of Capability Parameters present in the Initialization message received from the peer.

An LDP speaker MAY use a particular capability with its peer after the speaker determines that the peer has enabled that capability.

These procedures enable an LDP speaker S1, that advertises a specific LDP capability C, to establish an LDP session with speaker S2 that does not advertise C. In this situation, whether or not capability C may be used for the session depends on the semantics of the enhancement associated with C. If the semantics do not require both S1 and S2, advertise C to one another, then S2 could use it; i.e., S1's advertisement of C permits S2 to send messages to S1 used by the enhancement.

It is the responsibility of the capability designer to specify the behavior of an LDP speaker that has enabled a certain enhancement, advertised its capability and determines that its peer has not advertised the corresponding capability. The document specifying procedures for the capability MUST describe the behavior in this situation. If the specified procedure is to terminate the session, then the LDP speaker SHOULD send a Notification message to the peer before terminating the session. The Status Code in the Status TLV of the Notification message MUST be Unsupported Capability, and the message SHOULD contain the unsupported capability (see Section 8 for more details).

An LDP speaker that supports capability advertisement and includes a Capability Parameter in its Initialization message MUST set the TLV U-bit to 0 or 1, as specified by Capability document. The LDP speaker should set the U-bit to 1 if the capability document allows it to continue with a peer that does not understand the enhancement, and set the U-bit to 0 otherwise. If a speaker receives a message containing unsupported capability, it responds according to the U-bit setting in the TLV. If the U-bit is 1, then the speaker MUST silently ignore the Capability Parameter and allow the session to be established. However, if the U-bit is 0, then speaker SHOULD send a Notification message to the peer before terminating the session. The Status Code in the Status TLV of the Notification message MUST be Unsupported Capability, and the message SHOULD contain the unsupported capability (see Section 8 for more details).

7. Procedures for Capability Parameters in Capability Messages

An LDP speaker MUST NOT send a Capability message to a peer unless its peer advertised the Dynamic Capability Announcement capability in its session Initialization message. An LDP speaker MAY send a Capability message to a peer if its peer advertised the Dynamic Capability Announcement capability in its session Initialization message (see Section 9).

An LDP speaker determines the capabilities enabled by a peer by determining the set of capabilities enabled at session initialization (as specified in Section 6) and tracking changes to that set made by Capability messages from the peer.

An LDP speaker that has enabled a particular capability MAY use the enhancement corresponding to the capability with a peer after the speaker determines that the peer has enabled the capability.

8. Extensions to Error Handling

This document defines a new LDP status code named Unsupported Capability. The E-bit of the Status TLV carried in a Notification message that includes this status code MUST be set to 0.

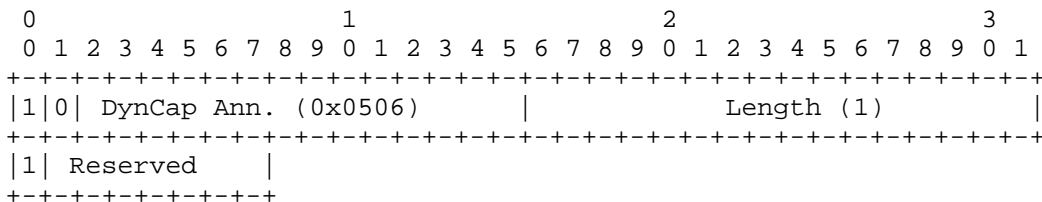
In addition, this document defines a new LDP TLV, named Returned TLVs, that MAY be carried in a Notification message as an Optional Parameter. The U-bit setting for a Returned TLVs TLV in a Notification message SHOULD be 1, and the F-bit setting SHOULD be 0.

When the Status Code in a Notification message is Unsupported Capability, the message SHOULD specify the capabilities that are unsupported. When the Notification message specifies the unsupported capabilities, it MUST include a Returned TLVs TLV. The Returned TLVs TLV MUST include only the Capability Parameters for unsupported capabilities, and the Capability Parameter for each such capability SHOULD be encoded as received from the peer.

When the Status Code in a Notification Message is Unknown TLV, the message SHOULD specify the TLV that was unknown. When the Notification message specifies the TLV that was unknown, it MUST include the unknown TLV in a Returned TLVs TLV.

9. Dynamic Capability Announcement TLV

The Dynamic Capability Announcement TLV is a Capability Parameter defined by this document with following format:



The value of the U-bit for the Dynamic Capability Announcement Parameter TLV MUST be set to 1 so that a receiver MUST silently ignore this TLV if unknown to it, and continue processing the rest of the message. There is no "Capability Data" associated with this TLV and hence the TLV length MUST be set to 1.

The Dynamic Capability Announcement Parameter MAY be included by an LDP speaker in an Initialization message to signal its peer that the speaker is capable of processing Capability messages.

An LDP speaker MUST NOT include the Dynamic Capability Announcement Parameter in Capability messages sent to its peers. Once enabled during session initialization, the Dynamic Capability Announcement capability cannot be disabled. This implies that the S-bit is always 1 for the Dynamic Capability Announcement.

An LDP speaker that receives a Capability message from a peer that includes the Dynamic Capability Announcement Parameter SHOULD silently ignore the parameter and process any other Capability Parameters in the message.

10. Backward Compatibility

From the point of view of the LDP capability advertisement mechanism, an [RFC5036]-compliant peer has label distribution for IPv4 enabled by default. To ensure compatibility with an [RFC5036]-compliant peer, LDP implementations that support capability advertisement have label distribution for IPv4 enabled until it is explicitly disabled and MUST assume that their peers do as well.

Section 3.1 introduces the concept of Backward Compatibility TLVs that may appear in an Initialization message in the role of a Capability Parameter. This permits existing LDP enhancements that use an ad hoc mechanism for enabling capabilities at session initialization time to continue to do so.

11. Security Considerations

[MPLS_SEC] describes the security framework for MPLS networks, whereas [RFC5036] describes the security considerations that apply to the base LDP specification. The same security framework and considerations apply to the capability mechanism described in this document.

12. IANA Considerations

This document specifies the following code points assigned by IANA:

- LDP message code point for the Capability message (0x0202).
- LDP TLV code point for the Dynamic Capability Announcement TLV (0x0506).
- LDP TLV code point for the Returned TLVs TLV (0x0304).
- LDP Status Code code point for the Unsupported Capability Status Code (0x0000002E).

13. Acknowledgments

The authors wish to thank Enke Chen, Vanson Lim, Ina Minei, Bin Mo, Yakov Rekhter, and Eric Rosen for their comments.

14. References

14.1. Normative References

- [RFC5036] Andersson, L., Ed., Minei, I., Ed., and B. Thomas, Ed., "LDP Specification", RFC 5036, October 2007.
- [RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, March 1997.
- [RFC3479] Farrel, A., Ed., "Fault Tolerance for the Label Distribution Protocol (LDP)", RFC 3479, February 2003.

14.2. Informative References

- [RFC5283] Decraene, B., Le Roux, JL., and I. Minei, "LDP Extension for Inter-Area Label Switched Paths (LSPs)", RFC 5283, July 2008.
- [MLDP] Minei, I., Ed., Kompella, K., Wijnands, I., Ed., and B. Thomas, "Label Distribution Protocol Extensions for Point-to-Multipoint and Multipoint-to-Multipoint Label Switched Paths", Work in Progress, April 2009.
- [NNHOP] Shen, N., Chen, E., and A. Tian, "Discovery LDP Next-Next-hop Labels", Work in Progress, May 2005.

- [RFC4447] Martini, L., Ed., Rosen, E., El-Aawar, N., Smith, T., and G. Heron, "Pseudowire Setup and Maintenance Using the Label Distribution Protocol (LDP)", RFC 4447, April 2006.
- [RFC3478] Leelanivas, M., Rekhter, Y., and R. Aggarwal, "Graceful Restart Mechanism for Label Distribution Protocol", RFC 3478, February 2003.
- [UPSTREAM_LDP] Aggarwal R., and J.L. Le Roux, "MPLS Upstream Label Assignment for LDP" Work in Progress, July 2008.
- [MPLS_SEC] Fang, L., Ed., "Security Framework for MPLS and GMPLS Networks", Work in Progress, March 2009.

Authors' Addresses

Bob Thomas
Cisco Systems, Inc.
1414 Massachusetts Ave.
Boxborough, MA 01719
EMail: bobthomas@alum.mit.edu

Shivani Aggarwal
Juniper Networks
1194 North Mathilda Ave.
Sunnyvale, CA 94089
EMail: shivani@juniper.net

Rahul Aggarwal
Juniper Networks
1194 North Mathilda Ave.
Sunnyvale, CA 94089
EMail: rahul@juniper.net

Jean-Louis Le Roux
France Telecom
2, Avenue Pierre-Marzin
22307 Lannion Cedex, France
EMail: jeanlouis.leroux@orange-ftgroup.com

Kamran Raza
Cisco Systems, Inc.
2000 Innovation Dr.
Kanata, ON K2K 3E8, Canada
EMail: skraza@cisco.com