

Internet Engineering Task Force (IETF)
Request for Comments: 5919
Category: Standards Track
ISSN: 2070-1721

R. Asati
P. Mohapatra
Cisco Systems
E. Chen
Huawei Technologies
B. Thomas
August 2010

Signaling LDP Label Advertisement Completion

Abstract

There are situations following Label Distribution Protocol (LDP) session establishment where it would be useful for an LDP speaker to know when its peer has advertised all of its labels. The LDP specification provides no mechanism for an LDP speaker to notify a peer when it has completed its initial label advertisements to that peer. This document specifies means for an LDP speaker to signal completion of its initial label advertisements following session establishment.

Status of This Memo

This is an Internet Standards Track document.

This document is a product of the Internet Engineering Task Force (IETF). It represents the consensus of the IETF community. It has received public review and has been approved for publication by the Internet Engineering Steering Group (IESG). Further information on Internet Standards is available in Section 2 of RFC 5741.

Information about the current status of this document, any errata, and how to provide feedback on it may be obtained at <http://www.rfc-editor.org/info/rfc5919>.

Copyright Notice

Copyright (c) 2010 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 (<http://trustee.ietf.org/license-info>) in effect on the date of publication of this document. Please review these documents carefully, as they describe your rights and restrictions with respect to this document. Code Components extracted from this document must include Simplified BSD License text as described in Section 4.e of the Trust Legal Provisions and are provided without warranty as described in the Simplified BSD License.

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	3
1.1. Applicability - Label Advertisement Mode	3
2. Specification Language	3
3. Unrecognized Notification Capability	4
4. Signaling Completion of Label Advertisement	4
4.1. Missing Expected End-of-LIB Notifications	5
5. Usage Guidelines	6
5.1. LDP-IGP Sync	6
5.2. LDP Graceful Restart	7
5.3. Wildcard Label Request	7
6. Security Considerations	8
7. IANA Considerations	8
8. Acknowledgments	8
9. References	8
9.1. Normative References	8
9.2. Informative References	9

1. Introduction

There are situations following LDP session establishment where it would be useful for an LDP speaker to know when its peer has advertised all of the labels from its Label Information Base (LIB). For example, when an LDP speaker is using LDP-IGP synchronization procedures [RFC5443], it would be useful for the speaker to know when its peer has completed advertisement of its IP label bindings. Similarly, after an LDP session is re-established when LDP Graceful Restart [RFC3478] is in effect, it would be helpful for each peer to signal the other after it has advertised all its label bindings.

The LDP specification [RFC5036] provides no mechanism for an LDP speaker to notify a peer when it has completed its initial label advertisements to that peer.

This document specifies use of a Notification message with the End-of-LIB Status Code for an LDP speaker to signal completion of its label advertisements following session establishment.

RFC 5036 implicitly assumes that new Status Codes will be defined over the course of time. However, it does not explicitly define the behavior of an LDP speaker that does not understand the Status Code in a Notification message. To avoid backward compatibility issues, this document specifies use of the LDP capability mechanism [RFC5561] at session establishment time for informing a peer that an LDP speaker is capable of handling a Notification message that carries an unrecognized Status Code.

1.1. Applicability - Label Advertisement Mode

The mechanisms specified in this document are deemed useful to LDP peering using the 'Downstream Unsolicited' label advertisement mode [RFC5036]. They are not deemed useful to any LDP peering using the 'Downstream on Demand' label advertisement mode since the LDP speaker would request particular label binding(s) from the peer anyway and know when it has received them.

2. Specification Language

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. Unrecognized Notification Capability

An LDP speaker MAY include a Capability Parameter [RFC5561] in the Initialization message to inform a peer that it ignores Notification Messages that carry a Status Type-Length-Value (TLV) with a non-fatal Status Code unknown to it.

The Capability Parameter for the Unrecognized Notification capability is a TLV with the following format:




Figure 1: Unrecognized Notification Capability Format

Where:

U- and F-bits: MUST be 1 and 0, respectively, as per Section 3 of LDP Capabilities [RFC5561].

Unrecognized Notif: 0x0603

S-bit: MUST be 1 (indicates that capability is being advertised).

Upon receiving a Notification with an unrecognized Status Code, an LDP speaker MAY generate a console or system log message for trouble shooting purposes.

4. Signaling Completion of Label Advertisement

An LDP speaker that conforms to this specification SHOULD signal completion of its label advertisements to a peer by means of a Notification message, if its peer has advertised the Unrecognized Notification capability during session establishment. The LDP speaker SHOULD send the Notification message (per Forwarding Equivalence Class (FEC) Type) to a peer even if the LDP speaker has zero Label bindings to advertise to that peer.

Such a Notification message MUST carry:

- A status TLV (with TLV E- and F-bits set to zero) that carries an End-of-LIB Status Code (0x0000002F).

- A FEC TLV with the Typed Wildcard FEC Element [RFC5918] that identifies the FEC type for which initial label advertisements have been completed. In terms of Section 3.5.1 of RFC 5036, this TLV is an "Optional Parameter" of the Notification message.

An LDP speaker MUST NOT send a Notification that carries a Status TLV with the End-of-LIB Status Code to a peer unless the peer has advertised the Unrecognized Notification capability during session establishment.

This applies to any LDP peers discovered via either basic discovery or extended discovery mechanisms (per Section 2.4 of [RFC5036]).

4.1. Missing Expected End-of-LIB Notifications

There is no guarantee that an LDP speaker will receive (or send) an End-of-LIB Notification from (or to) a peer even if the LDP speaker has signaled the Unrecognized Notification capability (Section 3).

Although it is expected that an LDP speaker supporting the Unrecognized Notification capability would support sending and receiving an End-of-LIB Notification, it is not mandatory by definition.

Please note that this is not a concern since the LDP speaker would simply ignore the received Notification with an End-of-LIB status code (or any status code) that is not recognized or supported, by definition.

To deal with the possibility of missing End-of-LIB Notifications after the LDP session establishment, an LDP speaker MAY time out receipt of an expected End-of-LIB Notification. An LDP speaker SHOULD start a per-peer internal timer, called 'EOL Notification' timer (the default value of 60 seconds is RECOMMENDED, though the value of this timer SHOULD be configurable) immediately following the LDP session establishment.

This timer is reset by the subsequent label advertisement, and stopped by the End-of-LIB Notification message. Lacking any label advertisement from the peer, the timer would expire, causing the LDP speaker to behave as if it had received the End-of-LIB notification from the peer.

If the End-of-LIB Notification message is received after the timer expires, then the message SHOULD be ignored.

5. Usage Guidelines

The FECs known to an LDP speaker and the labels the speaker has bound to those FECs may change over the course of time. This makes it difficult to determine when an LDP speaker has advertised "all" of its label bindings for a given FEC type. Ultimately, this determination is a judgment call the LDP speaker makes. The following guidelines may be useful.

An LDP speaker is assumed to "know" a set of FECs. Depending on a variety of criteria, such as:

- the label distribution control mode in use (Independent or Ordered);
- the set of FECs to which the speaker has bound local labels;
- configuration settings that may constrain which label bindings the speaker may advertise to peers.

The speaker can determine the set of bindings for a given FEC type that it is permitted to advertise to a given peer.

LDP-IGP Sync, LDP Graceful Restart, and the response to a Wildcard Label Request [RFC5918] are situations that would benefit from End-of-LIB Notification. In these situations, after an LDP speaker completes its label binding advertisements to a peer, sending an End-of-LIB Notification to the peer makes their outcome deterministic. The following subsections further explain each of these situations one by one.

5.1. LDP-IGP Sync

The LDP-IGP Synchronization [RFC5443] specifies a mechanism by which directly connected LDP speakers may delay the use of the link (between them) for transit IP traffic forwarding until the labels required to support IP-over-MPLS traffic forwarding have been distributed and installed.

Without an End-of-LIB Notification, the speaker must rely on some heuristic to determine when it has received all of its peer's label bindings. The heuristic chosen could cause LDP to signal the IGP too soon (in which case, the likelihood that traffic will be dropped increases) or too late (in which case, traffic is kept on sub-optimal paths longer than necessary).

Following session establishment, with a directly connected peer that has advertised the Unrecognized Notification capability, an LDP speaker using LDP-IGP Sync may send the peer an End-of-LIB Notification after it completes advertisement of its IP label bindings to the peer. Similarly, the LDP speaker may use the End-of-LIB Notification received from a directly connected peer to determine when the peer has completed advertisement of its label bindings for IP prefixes. After receiving the notification, the LDP speaker should consider LDP to be fully operational for the link and should signal the IGP to start advertising the link with normal cost.

5.2. LDP Graceful Restart

LDP Graceful Restart [RFC3478] helps to reduce the loss of MPLS traffic caused by the restart of a router's LDP component. It defines procedures that allow routers capable of preserving MPLS forwarding state across the restart to continue forwarding MPLS traffic using forwarding state installed prior to the restart for a configured time period.

The current behavior without End-of-LIB Notification is as follows: the restarting router and its peers consider the preserved forwarding state to be usable but stale until it is refreshed by receipt of new label advertisements following re-establishment of new LDP sessions or until the time period expires. When the time period expires, any remaining stale forwarding state is removed by the router.

Receiving End-of-LIB Notification from a peer in an LDP Graceful Restart scenario enables an LDP speaker to stop using stale forwarding information learned from that peer and to recover the resources it requires without having to wait until the time period expiry. The time period expiry can still be used if the End-of-LIB Notification message is not received.

5.3. Wildcard Label Request

When an LDP speaker receives a Label Request message for a Typed Wildcard FEC (e.g., a particular FEC Element Type) from a peer, the LDP speaker determines the set of bindings (as per any local filtering policy) to advertise to the peer for the FEC type specified by the request. Assuming the peer had advertised the Unrecognized Notification capability at session initialization time, the speaker should send the peer an End-of-LIB Notification for the FEC type when it completes advertisement of the permitted bindings.

As in the previous applications, receipt of the Notification eliminates uncertainty as to when the peer has completed its advertisements of label bindings for the requested Wildcard FEC Element Type.

6. Security Considerations

No security considerations beyond those that apply to the base LDP specification [RFC5036] and that are further described in [RFC5920] apply to signaling the End-of-LIB condition as described in this document.

7. IANA Considerations

This document introduces a new LDP Status Code and a new LDP Capability.

IANA has assigned the 'End-of-LIB' status code (0x0000002F) from the Status Code Name Space. [RFC5036] partitions the Status Code Name Space into 3 regions: IETF Consensus region, First Come First Served region, and Private Use region. The code point 0x0000002F is from the IETF Consensus range.

IANA has assigned the 'Unrecognized Notification' capability (0x0603) from the TLV Type name space. [RFC5036] partitions the TLV Type name space into 3 regions: IETF Consensus region, Vendor Private Use region, and Experimental Use region. The code point 0x0603 is from the IETF Consensus range.

8. Acknowledgments

The authors would like to recognize Kamran Raza, who helped to formulate this draft.

The authors would like to thank Ina Minei, Alia Atlas, Yakov Rekhter, Loa Andersson, and Luyuan Fang for their valuable feedback and contributions.

9. References

9.1. Normative References

- [RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, March 1997.
- [RFC5036] Andersson, L., Ed., Minei, I., Ed., and B. Thomas, Ed., "LDP Specification", RFC 5036, October 2007.

- [RFC5561] Thomas, B., Raza, K., Aggarwal, S., Aggarwal, R., and JL. Le Roux, "LDP Capabilities", RFC 5561, July 2009.
- [RFC5918] Asati, R., Minei, I., and B. Thomas, "Label Distribution Protocol (LDP) 'Typed Wildcard' Forward Equivalence Class (FEC)", RFC 5918, August 2010.

9.2. Informative References

- [RFC3478] Leelanivas, M., Rekhter, Y., and R. Aggarwal, "Graceful Restart Mechanism for Label Distribution Protocol", RFC 3478, February 2003.
- [RFC5443] Jork, M., Atlas, A., and L. Fang, "LDP IGP Synchronization", RFC 5443, March 2009.
- [RFC5920] Fang, L., Ed., "Security Framework for MPLS and GMPLS Networks", RFC 5920, July 2010.

Authors' Addresses

Rajiv Asati
Cisco Systems
7025-6 Kit Creek Rd.
Research Triangle Park, NC 27709-4987
EMail: rajiva@cisco.com

Pradosh Mohapatra
Cisco Systems
3750 Cisco Way
San Jose, CA 95134
EMail: pmohapat@cisco.com

Emily Chen
Huawei Technologies
No. 5 Street, Shangdi Information, Haidian
Beijing, China
EMail: chenying220@huawei.com

Bob Thomas
EMail: bobthomas@alum.mit.edu

