

MPLS Working Group
Internet Draft
Intended status: Standards Track
Expires: August 15, 2025

C. Lin
New H3C Technologies
Z. Huang
CRC Corporation
February 15, 2025

Label Distribution Protocol (LDP) Send Hold Timer
draft-lin-mpls-ldp-holdtimer-03

Abstract

This document defines the SendHoldTimer and SendHoldTimer_Expires event in the LDP protocol. The implementation of SendHoldTimer helps address situations where the local system detects that the remote system has not processed LDP messages but has not terminated the LDP session. The document specifies that when the SendHoldTimer expires, the local system should close the LDP session connection, rather than solely relying on the remote system for session termination. This document updates RFC5036.

Status of this Memo

This Internet-Draft is submitted in full conformance with the provisions of BCP 78 and BCP 79.

Internet-Drafts are working documents of the Internet Engineering Task Force (IETF). Note that other groups may also distribute working documents as Internet-Drafts. The list of current Internet-Drafts is at <https://datatracker.ietf.org/drafts/current/>.

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."

This Internet-Draft will expire on August 15, 2025.

Copyright Notice

Copyright (c) 2025 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.

Table of Contents

1. Introduction.....	3
1.1. Conventions and Terminology.....	3
2. Problem Statement.....	3
3. SendHoldTimer - Changes to RFC 5036	4
3.1. State Machine.....	4
3.2. Status Code.....	6
4. Security Considerations.....	7
5. IANA Considerations.....	7
6. Acknowledgements.....	7
7. References.....	7
7.1. Normative References.....	7
Authors' Addresses.....	9

1. Introduction

As described in [I-D.ietf-idr-bgp-sendholdtimer], any upper-layer protocol that uses TCP for transport can encounter similar situations where the remote system is unable to read TCP data due to a failure, leading to the inability to terminate the BGP connection.

LDP also requires a BGP-like send timeout mechanism [I-D.ietf-idr-bgp-sendholdtimer] to resolve this issue.

This document defines the `SendHoldTimer` and `SendHoldTimer_Expires` mechanisms in the LDP (RFC5036) to detect and handle situations where the peer is unable to receive LDP messages.

Failure to terminate a blocked LDP session may result in denial-of-service (DoS) attacks, leading to an inability to generate and disseminate LDP messages and causing the use of outdated LDP LSPs during forwarding. This specification aims to address this by requiring the termination of a session when the local system detects that the remote system is unable to process any LDP messages during the `SendHoldTime`.

With this specification, blocked connections can be terminated by the remote system as well as by the local system between LDP devices.

1.1. Conventions and Terminology

The key words "MUST", "MUST NOT", "REQUIRED", "SHALL", "SHALL NOT", "SHOULD", "SHOULD NOT", "RECOMMENDED", "NOT RECOMMENDED", "MAY", and "OPTIONAL" in this document are to be interpreted as described in BCP 14 [RFC2119] [RFC8174] when, and only when, they appear in all capitals, as shown here.

2. Problem Statement

An example of a network fault condition: Since LDP sessions operate over TCP (RFC9293), they are established between Label Switch Routers (LSRs), and FEC-label mappings are exchanged between LSRs via LDP messages.

In an established session between LSRs, one LSR may declare a zero-size TCP receive window (RCV.WND) to the other. This condition prevents the local system from sending essential LDP messages such as `KEEPALIVE`, `Label Request`, and `Label Mapping` to the remote peer through the network socket.

In the absence of an implementation of the SendHoldTimer concept, failures or overloaded remote peers may cause data to accumulate continuously on the local system's LDP socket, resulting in the use of outdated FEC-label mappings for forwarding and potentially leading to forwarding failures and traffic loss.

In typical scenarios, LDP implementations cannot observe the current receive window size of the underlying subsystem (e.g., TCP), and there is no LDP mechanism to terminate such blocked sessions. Therefore, LDP implementations cannot uniformly handle this situation.

This document provides a mechanism that enables an LDP implementation to detect whether the TCP socket between LSRs is actively transmitting data or is stalled. In the event of a stalled state, the LDP session can be restarted to restore normal operation of the LDP protocol.

3. SendHoldTimer - Changes to RFC 5036

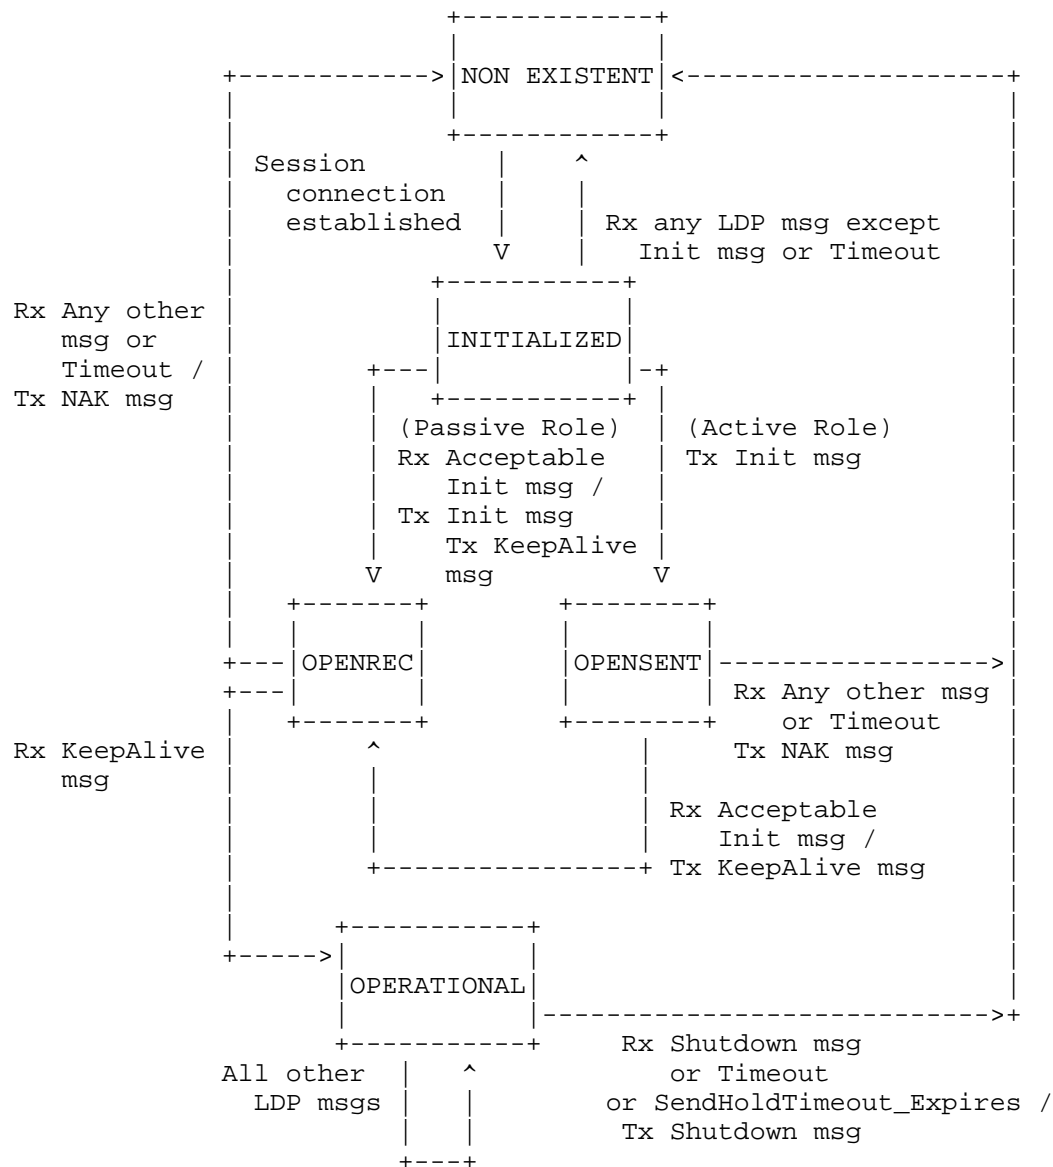
3.1. State Machine

In RFC5036, Section 2 describes the LDP state machine. To modify it and add handling for SendHoldTimer and SendHoldTimer_Expires.

In Section 2.5.4, the Initialization State Machine is described. When receiving a Shutdown message or modifying the LDP state machine, the handling for the SendHoldTimer Expires event needs to be added in the OPERATIONAL state.

This modification ensures that the LDP state machine includes specific handling for the SendHoldTimer Expires event in the OPERATIONAL state, thereby addressing the timing-related aspects of LDP protocol operation and session management.

OPERATIONAL	Receive Shutdown msg	NON EXISTENT
	Action: Transmit Shutdown msg and close transport connection	
	Receive other LDP msgs	OPERATIONAL
	Timeout	NON EXISTENT
	Action: Transmit Shutdown msg and close transport connection	
	SendHoldTimer Expires	NON EXISTENT
	Action: Transmit Shutdown msg and close transport connection	



In Section 2.5.6, which covers the maintenance process of LDP sessions, it is necessary to add the handling of the SendHoldTimer.

Old Text:

2.5.6. Maintaining LDP Sessions

...

After an LDP session has been established, an LSR must arrange that its peer receive an LDP PDU from it at least every KeepAlive time period to ensure the peer restarts the session KeepAlive Timer. The LSR may send any protocol message to meet this requirement. In circumstances where an LSR has no other information to communicate to its peer, it sends a KeepAlive message.

Next Text:

...

After an LDP session has been established, an LSR must arrange that its peer receive an LDP PDU from it at least every KeepAlive time period to ensure the peer restarts the session KeepAlive Timer. The LSR may send any protocol message to meet this requirement. In circumstances where an LSR has no other information to communicate to its peer, it sends a KeepAlive message.

After successfully sending an LDP message, the SendHoldTimer should be reset. This addition ensures the appropriate resetting of the SendHoldTimer during the maintenance of LDP sessions in accordance with the protocol's requirement for timely and responsive session management.

3.2. Status Code

A new Status Code "Hold Timer Expired" should be added for the event of SendHoldTimer_Expires. This code is used to notify the remote end of the reason for actively closing the session.

Status Code	E	Status Data	Section Title
Hold Timer Expired	1	TBD	"Events Signaled by ..."

This new Status Code "Hold Timer Expired" should be included in Notification Messages. In section 3.5.1.2 "Events Signaled by Notification Messages," add a description for Send Hold Timer Expiration:

Send Hold Timer Expiration

This is a fatal error signaled by the SendHold Timer Expired Status Code.

4. Security Considerations

This specification addresses the vulnerability of LSRs to potential attacks, where LDP peers pretend to be unable to process LDP messages, leading to the disruption of the LDP protocol.

In other aspects, this specification does not alter the security characteristics of LDP.

5. IANA Considerations

A new Status Code "Hold Timer Expired" should be added for the event of `SendHoldTimer_Expires`.

Status Code	E	Status Data	Section Title
Hold Timer Expired	1	TBD	"Events Signaled by ..."

6. Acknowledgements

The authors wish to thank Job Snijders, Ben Cartwright-Cox, and Yingzhen Qu for their work on the BGP Send Hold Timer concept.

7. References

7.1. Normative References

- [I-D.ietf-idr-bgp-sendholdtimer] Snijders, J., Cartwright-Cox, B., and Y. Qu, "Border Gateway Protocol 4 (BGP-4) Send Hold Timer", Work in Progress, Internet-Draft, draft-ietf-idr-bgp-sendholdtimer, <<https://datatracker.ietf.org/doc/html/draft-ietf-idr-bgp-sendholdtimer>>
- [RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, DOI 10.17487/RFC2119, March 1997, <<https://www.rfc-editor.org/info/rfc2119>>.
- [RFC8174] Leiba, B., "Ambiguity of Uppercase vs Lowercase in RFC 2119 Key Words", BCP 14, RFC 8174, DOI 10.17487/RFC8174, May 2017, <<https://www.rfc-editor.org/info/rfc8174>>.
- [RFC5036] Acreo AB, "LDP Specification", RFC 5036, October 2007, <<https://www.rfc-editor.org/info/rfc5036>>.

Authors' Addresses

Changwang Lin
New H3C Technologies
Beijing
China
Email: linchangwang.04414@h3c.com

Zhibo Huang
China Railway Construction Corporation Limited
Email: huangzhibo.gj@crcc.cn

