

Network Working Group
Request for Comments: 3037
Category: Informational

B. Thomas
Cisco Systems, Inc.
E. Gray
Zaffire, Inc.
January 2001

LDP Applicability

Status of this Memo

This memo provides information for the Internet community. It does not specify an Internet standard of any kind. Distribution of this memo is unlimited.

Copyright Notice

Copyright (C) The Internet Society (2001). All Rights Reserved.

Abstract

Multiprotocol Label Switching (MPLS) is a method for forwarding packets that uses short, fixed-length values carried by packets, called labels, to determine packet nexthops. A fundamental concept in MPLS is that two Label Switching Routers (LSRs) must agree on the meaning of the labels used to forward traffic between and through them. This common understanding is achieved by using a set of procedures, called a label distribution protocol, by which one LSR informs another of label bindings it has made. This document describes the applicability of a set of such procedures called LDP (for Label Distribution Protocol) by which LSRs distribute labels to support MPLS forwarding along normally routed paths.

1. LDP Applicability

A label distribution protocol is a set of procedures by which one Label Switching Router (LSR) informs another of the meaning of labels used to forward traffic between and through them.

The MPLS architecture allows for the possibility of more than a single method for distributing labels, and a number of different label distribution protocols are being standardized. Existing protocols have been extended so that label distribution can be piggybacked on them, and new protocols have been defined for the explicit purpose of distributing labels.

This document describes the applicability of the Label Distribution Protocol (LDP), a new protocol for label distribution designed to support label distribution for MPLS forwarding along normally routed paths as determined by destination-based routing protocols. This is sometimes called MPLS hop-by-hop forwarding.

LDP, together with an IP routing plane and software to program ATM switch or Frame Relay switch cross-connect tables, can implement IP in a network of ATM and/or Frame Relay switches without requiring an overlay or the use of ATM-specific or Frame Relay-specific addressing or routing.

LDP is also useful in situations that require efficient hop-by-hop routed tunnels, such as MPLS-based VPN architectures [RFC2574] and tunneling between BGP border routers.

In addition, LDP includes a mechanism that makes it possible to extend it to support MPLS features that go beyond best effort hop-by-hop forwarding.

As a stand-alone protocol for distributing labels LDP does not rely on the presence of specific routing protocols at every hop along an LSP path in order to establish an LSP. Hence LDP is useful in situations in which an LSP must traverse nodes which may not all support a common piggybacked approach to distributing labels.

Traffic Engineering [TE] is expected to be an important MPLS application. MPLS support for Traffic Engineering uses explicitly routed LSPs, which need not follow normally-routed (hop-by-hop) paths.

Explicitly routed LSPs may be setup by CR-LDP [CRLDP-AS], a set of extensions to LDP, or by RSVP-TE [RSVP-TE-AS], a set of extensions to RSVP. There is currently no consensus on which of these protocols is technically superior. Therefore, network administrators should make a choice between the two based upon their needs and particular situation.

2. Requirement Level

The "requirement level" [RFC2026] for LDP is:

Implementation of LDP is recommended for devices that perform MPLS forwarding along normally routed paths as determined by destination-based routing protocols.

3. Feature Overview

LDP associates a Forwarding Equivalence Class (FEC) [RFC3031] with each label it distributes. Two LSRs which use LDP to exchange FEC-label binding information are known as "LDP Peers", and we speak of there being an "LDP Session" between them.

LDP uses TCP for session communication. Use of TCP ensures that session messages are reliably delivered, and that distributed labels and state information associated with LSPs need not be refreshed periodically.

LDP includes a mechanism by which an LSR can discover potential LDP peers. The discovery mechanism makes it unnecessary for operators to explicitly configure each LSR with its LDP peers.

When an LSR discovers another LSR it follows the LDP session setup procedure to establish an LDP session. By means of this procedure the LSRs establish a session TCP connection and use it to negotiate parameters for the session, such as the label distribution method to be used (see below). After the LSRs agree on the parameters, the session is operational and the LSRs use the TCP connection for label distribution.

LDP supports two different methods for label distribution. An LSR using Downstream Unsolicited distribution advertises FEC-label bindings to its peers when it is ready to forward packets in the FEC by means of MPLS. An LSR using Downstream on Demand distribution provides FEC-label bindings to a peer in response to specific requests from the peer for a label for the FEC.

LDP allows LSRs flexibility in strategies for retaining learned labels. An LSR using liberal label retention stores all labels learned from peers regardless of whether it currently needs them for forwarding, whereas one using conservative label retention stores only labels for which it has immediate use and releases unneeded labels to the peer that advertised them.

In addition, LDP allows flexibility in strategies for when to advertise FEC-label bindings. An LSR using independent control mode advertises FEC-label bindings to peers whenever it sees fit, whereas one using ordered control advertises bindings only when it has previously received a label for the FEC from the FEC next hop or it is an MPLS egress for the FEC.

Downstream on Demand distribution with conservative label retention and ordered control is appropriate in situations where labels are a relatively scarce resource that must be conserved, and Downstream

Unsolicited distribution with liberal label retention and independent control is appropriate where labels are plentiful and need not be carefully conserved. However, the protocol permits other combinations of distribution method, label retention mode and control mode, including hybrid variants of them.

LDP defines a mechanism for loop detection to protect against forwarding loops in LSPs that traverse non-TTL MPLS clouds; see [RFC3031] for discussion of situations which may benefit from this mechanism. The loop detection mechanism is optional in the sense that it may be disabled by LSR configuration. However, an LDP-compliant LSR must implement it.

LDP includes an extension mechanism which supports the development of vendor-private and experimental features. This mechanism defines procedures for introducing new types of messages and TLVs, methods an LSR can use for detecting such messages and TLVs, and procedures an LSR must follow when it receives a message or TLV it does not implement. While it is not possible to make every future enhancement backwards compatible, these procedures facilitate the introduction of new capabilities in MPLS networks that include older implementations that do not recognize them.

4. Scalability Considerations

The following factors may influence the scalability of LDP implementations:

- LDP label distribution is incremental, requiring no periodic refresh of FEC-label bindings.
- In situations where label resources may be scarce (ATM and Frame Relay links) the use of the Downstream on Demand distribution method with conservative label retention ensures that only those labels required to support normally-routed paths are allocated and distributed.
- In situations where label resources are not scarce, the use of the Downstream Unsolicited method with liberal label retention ensures that changes in FEC next-hop from one LDP peer to another require no distribution action to update previously distributed labels.
- Limitations on the number of TCP connections an LSR supports limit the number of LDP peers the LSR can support.

- Use of the optional path vector based loop detection mechanism imposes additional memory and processing requirements on an LSR, as well as additional LDP traffic. Both impact scalability.

5. Security Considerations

LDP defines the optional use of the TCP MD5 Signature Option to protect against the introduction of spoofed TCP segments into LDP session connection streams. LDP use of the TCP MD5 Signature Option is similar to BGP [RFC1771] use of the option specified in [RFC2385].

6. References

- [CRLDP-AS] J. Ash, M. Girish, E. Gray, B. Jamoussi, G. Wright, "Applicability Statement for CR-LDP", Work in Progress, September 1999.
- [RFC1771] Rekhter, Y. and T. Li, "A Border Gateway Protocol 4 (BGP-4)", RFC 1771, March 1995.
- [RFC2026] Bradner, S., "The Internet Standards Process -- Revision 3", BCP 9, RFC 2026, October 1996.
- [RFC2385] Heffernan, A., "Protection of BGP Sessions via the TCP MD5 Signature Option", RFC 2385, August 1998.
- [RFC2547] Rosen, E. and Y. Rekhter, "BGP/MPLS VPNs", RFC 2547, March 1999.
- [RFC3036] Andersson, L., Doolan, P., Feldman, N., Fredette, A. and B. Thomas, "LDP Specification", RFC 3036, January 2001.
- [RFC3031] Rosen, E., Viswanathan, A. and R. Callon, "Multiprotocol Label Switching Architecture", RFC 3031, January 2001.
- [RSVP-TE-AS] Awduche, D., Hannan, A. and X. Xiao, "Applicability State for Extensions to RSVP for LSP-Tunnels", Work in Progress, April 2000.

7. Authors' Addresses

Eric Gray
Zaffire, Inc
2630 Orchard Parkway,
San Jose, CA 95134-2020

Phone: 408-894-7362
EMail: ewgray@mindspring.com

Bob Thomas
Cisco Systems, Inc.
250 Apollo Dr.
Chelmsford, MA 01824

Phone: 978-244-8078
EMail: rhthomas@cisco.com

8. Full Copyright Statement

Copyright (C) The Internet Society (2001). 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.

This document and the information contained herein is provided on an "AS IS" basis and THE INTERNET SOCIETY AND THE INTERNET ENGINEERING TASK FORCE DISCLAIMS ALL WARRANTIES, EXPRESS OR IMPLIED, INCLUDING BUT NOT LIMITED TO ANY WARRANTY THAT THE USE OF THE INFORMATION HEREIN WILL NOT INFRINGE ANY RIGHTS OR ANY IMPLIED WARRANTIES OF MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE.

Acknowledgement

Funding for the RFC Editor function is currently provided by the Internet Society.