MPLS Working Group
INTERNET-DRAFT
Intended Status: Proposed Standard
Expires: March 23, 2015

Santosh Esale
Raveendra Torvi
Chris Bowers
Juniper Networks

Luay Jalil
Verizon

U. Chunduri
Ericsson Inc.

Zhenbin Li
Huawei

September 19, 2014

Application-aware Targeted LDP
draft-esale-mpls-app-aware-tldp-01

Abstract

Recent targeted LDP applications such as remote loop-free alternates and BGP auto discovered pseudowire may automatically establish a tLDP session to any LSR in a network. The initiating LSR has information about the targeted applications to administratively control initiation of the session. However the responding LSR has no such information to control acceptance of this session. This document defines a mechanism to advertise and negotiate Targeted Applications Capability during LDP session initialization. As the responding LSR becomes aware of targeted applications, it may establish a limited number of tLDP sessions for certain applications. In addition, each targeted application is mapped to LDP Forwarding Equivalence Class (FEC) Elements to advertise only necessary LDP FEC-label bindings over the session.

Status of this Memo

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

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/1id-abstracts.html

The list of Internet-Draft Shadow Directories can be accessed at http://www.ietf.org/shadow.html

Table of Contents

1 Introduction 4
1.1 Terminology 4
2. Targeted Application Capability 5
3. Targeted Application Capability Procedures 6
4. Interaction of Targeted Application Capabilities and State Advertisement Control Capabilities 8
5. Targeted Application capability in LDP messages 9
5.1 TAC in LDP Initialization message 9
5.2 TAC in LDP Capability message 10
6. Use cases 10
6.1 Remote LFA Automatic Targeted session 10
6.2 FEC 129 Auto Discovery Targeted session 11
6.3 LDP over RSVP and Remote LFA targeted session 11
6.4 mLDP node protection targeted session 12
7 Security Considerations 12
8 IANA Considerations 12
9. Acknowledgments 13
10 References 13

10.1 Normative References 13
10.2 Informative References 14
Authors' Addresses 14

1 Introduction

LDP can use the extended discovery mechanism to establish a tLDP adjacency and subsequent session as described in [RFC5036]. An LSR initiates extended discovery by sending a tLDP Hello to a specific address. The remote LSR decides either to accept or ignore a tLDP Hello based on local configuration only. For an application such as FEC 128 pseudowire, the remote LSR is configured with the source LSR address, so the remote LSR can use that information to accept or ignore a given tLDP Hello.

Applications such as Remote LFA and BGP auto discovered pseudowire automatically initiate asymmetric extended discovery to any LSR in a network based on local state only. With these applications, the remote LSR is not explicitly configured with the source LSR address. so the remote LSR either responds to all LDP requests or ignores all LDP requests.

In addition, since the session is initiated and established after adjacency formation, the responding LSR has no targeted applications information to choose the targeted application it is configured to support. Also, the initiating LSR may employ a limit per application on locally initiated automatic tLDP sessions, however the responding LSR has no such information to employ a similar limit on the incoming tLDP sessions. Further, the responding LSR does not know whether the source LSR is establishing a tLDP session for a configured or an automatic application or both.

This document proposes and describes a solution to advertise Targeted Application Capability, consisting of a targeted application list, during initialization of a tLDP session. It also defines a mechanism to enable a new application and disable an old application after session establishment. This capability advertisement provides the responding LSR with the necessary information to control the acceptance of tLDP sessions per application. For instance, an LSR may accept all BGP auto discovered tLDP sessions but may only accept limited number of Remote LFA tLDP sessions.

Also, targeted LDP application is mapped to LDP FEC element type to advertise specific application FECs only, avoiding the advertisement of other unnecessary FECs over a tLDP session.

1.1 Terminology

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 RFC 2119 [RFC2119].

2. Targeted Application Capability

An LSR MAY advertise that it is capable to negotiate a targeted LDP application list over a tLDP session by using the Capability Advertisement as defined in [RFC5561].

A new optional capability TLV is defined, 'Targeted Application Capability (TAC)'. Its encoding is as follows:

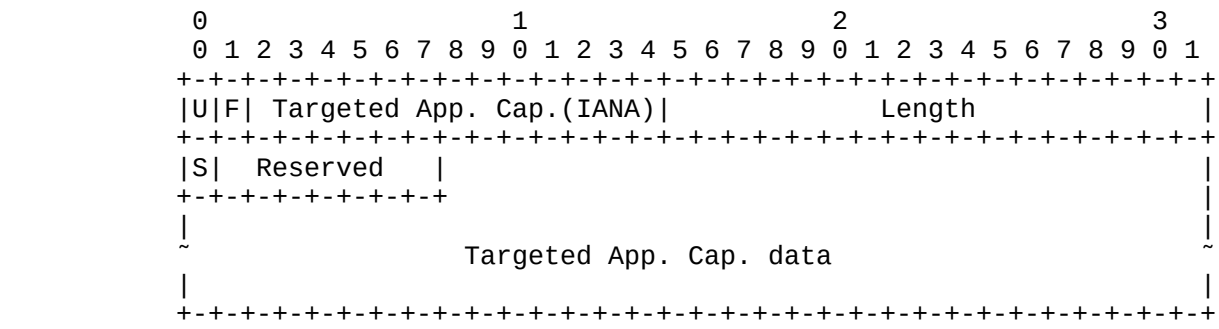

As described in [RFC5561]
U: set to 1. Ignore, if not known.
F: Set to 0. Do not forward.
S: MUST be set to 1 or 0 to advertise or withdraw the Targeted Application Capability TLV respectively.

Targeted Application Capability data:
A Targeted Applications Capability data consists of none, one or more 16 bit Targeted Application Elements. Its encoding is as follows:

Targeted Application Element(TAE)

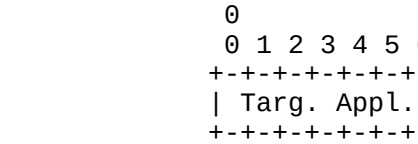

Targeted Application Identifier (TA-Id):

0x01: LDPv4 Tunneling
0x02: LDPv6 Tunneling
0x03: mLDP Tunneling
0x04: LDPv4 Remote LFA
0x05: LDPv6 Remote LFA
0x06: LDP FEC 128 Pseudowire

0x07: LDP FEC 129 Pseudowire
0x08: LDPv4 Session Protection
0x09: LDPv6 Session Protection
0x0A: LDP ICCP
0x0B: LDP P2MP PW
0x0C: mLDP node protection

E-bit: The enable bit indicates whether the sender is advertising or withdrawing the Targeted Application. The E-bit value is used as follows:

1 - The TAE is advertising the targeted application.
0 - The TAE is withdrawing the targeted application.

The length of TAC depends on the number of TAEs. For instance, if two TAEs are added, the length is set to 5. If both the peers advertise TAC, an LSR decides to establish or close a tLDP session based on the negotiated targeted application list.

For instance, suppose a initiating LSR advertises A, B and C as TA-Ids. Further, suppose the responding LSR advertises C, D and E as TA-Ids. Than the negotiated TA-Id, as per both the LSRs is C. In the second instance, suppose a initiating LSR advertises A, B and C as TA-Ids and the responding LSR, which acts as a passive LSR, advertises all the applications - A, B, C, D and E that it supports over this session. Than the negotiated targeted application as per both the LSRs are A, B and C. In the last instance, suppose the initiating LSR advertises A, B and C as a TA-Ids and the responding LSR advertises D and E as TA-Ids, than the negotiated targeted applciations as per both the LSRs is none. The Responding LSR sends 'Session Rejected/Targeted Application Capability Mis-Match' Notification message to the initiating LSR and may close the session.

3. Targeted Application Capability Procedures

At tLDP session establishment time, a LSR MAY include a new capability TLV, Targeted Application Capability (TAC) TLV, as an optional TLV in the LDP Initialization message. The TAC TLV's Capability data MUST consists of none, one or more Targeted Application Element(TAE) each pertaining to a unique Targeted Application Identifier(TA-Id) that a LSR supports over the session. If the receiver LSR receives the same TA-Id in more than one TAE, it MUST process the first element and ignore the duplicate elements. If the receiver LSR receives an unknown TA-Id in a TAE, it MUST silently ignore such a TAE and continue processing the rest of the TLV.

If the receiver LSR does not receive the TAC in the Initialization message or it does not understand the TAC TLV, the TAC negotiation MUST be considered unsuccessful and the session establishment MUST proceed as per [RFC5036]. On the receipt of a valid TAC TLV, an LSR MUST generate its own TAC TLV with TAEs consisting of unique TA-Ids that it supports over the tLDP session. If there is at least one TAE common between the TAC TLV it has received and its own, the session MUST proceed to establishment as per [RFC5036]. If not, A LSR MUST send a 'Session Rejected/Targeted Application Capability Mis-Match' Notification message to the peer and close the session. The initiating LSR playing the passive role in LDP session establishment MAY tear down the corresponding tLDP adjacency.

When the responding LSR playing the active role in LDP session establishment receives a 'Session Rejected/Targeted Application Capability Mis-Match' Notification message, it MUST set its session setup retry interval to a maximum value, as 0xffff. The session MAY stay in non-operational state. When it detects a change in the initiating LSR configuration or local LSR configuration pertaining to TAC TLV, it MUST clear the session setup back off delay associated with the session to re-attempt the session establishment. A LSR detects configuration change on the other LSR with the receipt of tLDP Hello message that has a higher configuration sequence number than the earlier tLDP Hello message.

When the initiating LSR playing the active role in LDP session establishment receives a 'Session Rejected/Targeted Application Capability Mis-Match' Notification message, either it MUST set its session setup retry interval to a maximum value, as 0xffff or it MUST teard down the corresponding tLDP adjacency with the session. This also leads to destruction of the session.

If it sets the session setup retry interval to maximum, the session MAY stay in a non-operational state. When this LSR detects a change in the responding LSR configuration or its own configuration pertaining to TAC TLV, it MUST clear the session setup back off delay associated with the session to re-attempt the session establishment.

If it decides to tear down the associated tLDP adjacency, the session is destroyed on the initiating as well as the responding LSR. The initiating LSR MAY take appropriate actions if it is unable to bring up the tLDP session. For instance, if an automatic session intended to support the Remote LFA application is rejected by the responding LSR, the initiating LSR may inform the IGP to calculate another PQ node [I-D.draft-ietf-rtgwg-remote-lfa] for the route or set of routes. More specific actions are a local matter and outside the scope of this document.

After a tLDP session has been established with TAC capability, the initiating and responding LSR MUST distribute FEC-label bindings for the negotiated applications only. For instance, if the tLDP session is established for BGP auto discovered pseudowire, only FEC 129 label bindings MUST be distributed over the session. Similarly, a LSR operating in downstream on demand mode MUST request FEC-label bindings for the negotiated applications only.

If the Targeted Application Capability and Dynamic Capability, as described in RFC 5561, are negotiated during session initialization, TAC MAY be re-negotiated after session establishment by sending an updated TAC TLV in LDP Capability message. The updated TLV MUST consist of one or more TAEs with E-bit set or E-bit off to advertise or withdraw the new and old application respectively. This may lead to advertisements or withdrawals of certain types of FEC-Label bindings over the session or tear down of the tLDP adjacency and subsequently the session.

The Targeted Application Capability is advertised on tLDP session only. If the tLDP session changes to link session, a LSR should withdraw it with S bit set to 0, which indicates wildcard withdral of all TAE elements. Similarly, if the link session changes to tLDP, a LSR should advertise it via the Capability message. If the capability negotiation fails, this may lead to destruction of the tLDP session.

Also, currently the remote LSR accepts asymmetric extended Hellos by default or by appropriate configuration. With this document, it should accept by default in order to then accept or reject the tLDP session based on the application information.

4. Interaction of Targeted Application Capabilities and State Advertisement Control Capabilities

As described in this document, the set of Targeted Application Elements negotiated between two LDP peers advertising TAC represents the willingness of both peers to advertise state information for a set of applications. The set of applications negotiated by the TAC mechanism is symmetric between the two LDP peers. In the absence of further mechanisms, two LDP peers will both advertise state information for the same set of applications.

As described in [I-D.draft-ietf-mpls-ldp-ip-pw-capability], State Advertisement Control(SAC) TLV can be used by an LDP speaker to communicate its interest or disinterest in receiving state information from a given peer for a particular application. Two LDP peers can use the SAC mechanism to create asymmetric advertisement of state information between the two peers for any particular application.

For a given tLDP session, the TAC mechanism can be used without the SAC mechanism, and the SAC mechanism can be used without the TAC mechanism. It is useful to discuss the behavior when TAC and SAC mechanisms are used on the same tLDP session. The TAC mechanism takes precedence over the SAC mechanism with respect to enabling applications for which state information will be advertised. For an tLDP session using the TAC mechanism, the LDP peers MUST NOT advertise state information for an application that has not been negotiated in the most recent Targeted Application Elements list (referred to as an un-negotiated application). This is true even if one of the peers announces its interest in receiving state information that corresponds to the un-negotiated application by sending a SAC TLV. In other words, when TAC is being used, SAC cannot enable state information advertisement for applications that have not been enabled by TAC.

On the other hand, the SAC mechanism takes precedence over the TAC mechanism with respect to disabling state information advertisements. If an LDP speaker has announced its disinterest in receiving state information for a given application to a given peer using the SAC mechanism, its peer MUST NOT send state information for that application, even if the two peers have negotiated that the corresponding application via the TAC mechanism.

For the purposes of determining the correspondence between targeted applications defined in this document and application state as defined in [I-D.draft-ietf-mpls-ldp-ip-pw-capability] an LSR MAY using the following mappings:

LDPv4 Tunneling - IPv4 Prefix-LSPs
LDPv6 Tunneling - IPv6 Prefix-LSPs
LDPv4 Remote LFA - IPv4 Prefix-LSPs
LDPv6 Remote LFA - IPv6 Prefix-LSPs
LDP FEC 128 Pseudowire - FEC128 P2P-PW
LDP FEC 129 Pseudowire - FEC129 P2P-PW
LDPv4 Session Protection - IPv4 Prefix-LSPs
LDPv6 Session Protection - IPv6 Prefix-LSPs

An LSR MAY map Targeted Application to LDP capability as follows:

mLDP Tunneling - P2MP Capability, MP2MP Capability

5. Targeted Application capability in LDP messages

5.1 TAC in LDP Initialization message

1. The S-bit of the Targeted Application Capability TLV MUST be set to 1 to advertise Targeted Application Capability and

SHOULD be ignored on the receipt.

2. The E-bit of the Targeted Application Element MUST be set to 1 to enable Targeted application.

3. An LSR MAY add State Control Capability by mapping Targeted Application Element to State Advertisement Control (SAC) Elements as defined in Section 4.

4. The LSR MAY add a different KeepAlive Time [RFC5036] value for an automatic tLDP session.

5.2 TAC in LDP Capability message

The initiating or responding LSR may re-negotiate the TAC after local configuration change with the Capability message.

1. The S-bit of Targeted Application Capability is set to 1 or 0 to advertise or withdraw it.

2. After configuration change, If there is no common TAE between its new TAE list and peers TAE list, the LSR MUST send a 'Session Rejected/Targeted Application Capability Mis-Match' Notification message and close the session.

3. If there is a common TAE, a LSR MAY also update SAC Capability based on updated TAC as described in section 4 and sends the updated TAC and SAC capabilities in a Capability message to the peer.

4. A receiving LSR processes the Capability message with TAC TLV. If the S-bit is set to 0, the TAC is disbaled for the session. After that, the session may remain in established state or torn down based on [RFC5036] rules.

5. If the S-bit is set to 1, a LSR process a list of TAEs from TACs capability data with E-bit set to 1 or 0 to update the peers TAE. Also, it udpates the negotiated TAE list over the tLDP session.

6. Use cases

6.1 Remote LFA Automatic Targeted session

An LSR determines that it needs to form an automatic tLDP session to remote LSR based on IGP calculation as described in [I-D.draft-ietf-rtgwg-remote-lfa] or some other mechanism, which is outside the scope

of this document. The LSR forms the tLDP adjacency and during session setup, constructs an Initialization message with Targeted Applications Capability (TAC) with Targeted Application Element (TAE) as Remote LFA. The receiver LSR processes the LDP Initialization message and verifies whether it is configured to accept a Remote LFA tLDP session. If it is, it may further verify that establishing such a session does not exceed the configured limit for Remote LFA sessions. If all these conditions are met, the receiver LSR may respond back with an Initialization message with TAC corresponding to Remote LFA, and subsequently the session may be established.

After the session has been established with TAC capability, the sender and receiver LSR distribute IPv4 or IPv6 FEC label bindings over the session. Further, the receiver LSR may determine that it does not need these FEC label bindings. So it may disable the receipt of these FEC label bindings by mapping targeted application element to state control capability as described in section 4.

6.2 FEC 129 Auto Discovery Targeted session

BGP auto discovery or other mechanisms outside the scope of this document MAY determine whether an LSR needs to initiate an auto-discovery tLDP session with a border LSR. Multiple LSRs MAY try to form an auto discovered tLDP session with a border LSR. So, a service provider may want to limit the number of auto discovered tLDP sessions a border LSR may accept. As described in Section 3, LDP may convey targeted applications with TAC TLV to border LSR. A border LSR may establish or reject the tLDP session based on local administrative policy. Also, as the receiver LSR becomes aware of targeted applications, it can also employ an administrative policy for security. For instance, it can employ a policy ’accept all auto-discovered session from source-list’.

Moreover, the sender and receiver LSR MUST exchange FEC 129 label bindings only over the tLDP session.

6.3 LDP over RSVP and Remote LFA targeted session

A LSR may want to establish a tLDP session to a remote LSR for LDP over RSVP tunneling and Remote LFA applications. The sender LSR may add both these applications as a unique Targeted Application Element in the Targeted Application Capability data of a TAC TLV. The receiver LSR MAY have reached a configured limit for accepting Remote LFA automatic tLDP sessions, but it may also be configured to accept LDP over RSVP tunneling. In such a case, the tLDP session is formed for both LDP over RSVP and Remote LFA applications as both needs same FECs - IPv4 and/or IPv6.

Also, the sender and the receiver LSR MUST distributes IPv4 and or IPv6 FEC label bindings only over the tLDP ession.

6.4 mLDP node protection targeted session

A merge point LSR may determines that it needs to form automatic tLDP session to the upstream point of local repair (PLR) LSR for MP2P and MP2MP LSP node protection as described in the [I-D.draft-ietf-mpls-mldp-node-protection] or other documents, which is outside the scope of this document. The MPT LSR may add a new targeted LDP application - mLDP node protection, as a unique TAE in the Targeted Application Capability Data of a TAC TLV and send it in the Initialization message to the PLR. If the PLR is configured for mLDP node protection and establishing this session does not exceed the limit of either mLDP node protection sessions or automatic tLDP sessions, the PLR may decide to accept this session. Further, the PLR responds back with the initialazation message with a TAC TLV that has one of the TAEs as - mLDP node protection and the session procceds to establishemt as per RFC 5036.

7 Security Considerations

The Capability procedure described in this document will apply and does not introduce any change to LDP Security Considerations section described in [RFC5036].

8 IANA Considerations

This document requires the assignment of a new code point for a Capability Parameter TLVs from the IANA managed LDP registry "TLV Type Name Space", corresponding to the advertisement of the Targeted Applications capability. IANA is requested to assign the lowest available value after 0x050B.

Value	Description	Reference
TBD1	Targeted Applications capability	[This draft]

This document requires the assignment of a new code point for a status code from the IANA managed registry "STATUS CODE NAME SPACE" on the Label Distribution Protocol (LDP) Parameters page, corresponding to the notification of session Rejected/Targeted Application Capability Mis-Match. IANA is requested to assign the lowest available value after 0x0000004B.

Value	Description	Reference

TBD2 Session Rejected/Targeted
Application Capability Mis-Match [This draft]

This document also creates a new name space 'the LDP Targeted Application Element type' on the Label Distribution Protocol (LDP) Parameters page, that is to be managed by IANA. The range is 0-255, with the following values requested in this document.

0x00: Reserved
0x01: LDPv4 Tunneling
0x02: LDPv6 Tunneling
0x03: mLDP Tunneling
0x04: LDPv4 Remote LFA
0x05: LDPv6 Remote LFA
0x06: LDP FEC 128 Pseudowire
0x07: LDP FEC 129 Pseudowire
0x08: LDPv4 Session Protection
0x09: LDPv6 Session Protection
0x0A: LDP ICCP
0x0B: LDP P2MP PW
0x0C: mLDP node protection
0xFF: Reserved

The allocation policy for this space is 'Standards Action'.

9. Acknowledgments

The authors wish to thank Nischal Sheth, Hassan Hosseini, Kishore Tiruveedhula, Kamran Raza and Loa Andersson for doing the detailed review. Thanks to Manish Gupta and Martin Ehlers for their input to this work and for many helpful suggestions.

10 References

10.1 Normative References

[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.

[I-D.draft-ietf-mpls-ldp-ip-pw-capability] Kamran Raza, Sami Boutros, "Disabling IPoMPLS and P2P PW LDP Application's State Advertisement", draft-ietf-mpls-ldp-ip-pw-capability-07

(work in progress), April 27, 2014.

[I-D.draft-ietf-mpls-mldp-node-protection] IJ. Wijnands, E. Rosen, K. Raza, J. Tantsura, A. Atlas, Q. Zhao, "mLDP Node Protection", draft-ietf-mpls-mldp-node-protection-01 (work in progress), February 13, 2014.

[RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, March 1997.

10.2 Informative References

[I-D.draft-ietf-rtgwg-remote-lfa] S. Bryant, C. Filsfils, S. Previdi,M. Shand, "Remote LFA FRR", draft-ietf-rtgwg-remote-lfa-06 (work in progress), May 23, 2014.

[RFC6074] E. Rosen, B. Davie, V. Radoaca, and W. Luo, "Provisioning, Auto-Discovery, and Signaling in Layer 2 Virtual Private Networks (L2VPNs)"

[RFC4762] M. Lasserre, and V. Kompella, "Virtual Private LAN Service VPLS) Using Label Distribution Protocol (LDP) Signaling", RFC 4762, January 2007.

[RFC4447] L. Martini, E. Rosen, El-Aawar, T. Smith, and G. Heron, "Pseudowire Setup and Maintenance using the Label Distribution Protocol", RFC 4447, April 2006.

[RFC5331] Aggarwal, R., Rekhter, Y., and E. Rosen, "MPLS Upstream Label Assignment and Context-Specific Label Space", RFC 5331, August 2008.

Authors' Addresses

Santosh Esale
Juniper Networks
1194 N. Mathilda Ave.
Sunnyvale, CA 94089
US
EMail: sesale@juniper.net

Raveendra Torvi
Juniper Networks
10 Technology Park Drive.
Westford, MA 01886
US

EMail: rtorvi@juniper.net

Chris Bowers
Juniper Networks
1194 N. Mathilda Ave.
Sunnyvale, CA 94089
US
EMail: cbowers@juniper.net

Luay Jalil
Verizon
1201 E Arapaho Rd.
Richardson, TX 75081
US
Email: luay.jalil@verizon.com

Uma Chunduri
Ericsson Inc.
300 Holger Way
San Jose, California 95134
USA
Email: uma.chunduri@ericsson.com

Zhenbin Li
Huawei Bld No.156 Beiqing Rd.
Beijing 100095
China
Email: lizhenbin@huawei.com