MPLS Working Group
Internet Draft
Intended status: Standards Track
Expires: November 9, 2013

Kamran Raza
Sami Boutros

Cisco Systems

May 10, 2013

Disabling IPoMPLS and P2P PW LDP Application's State Advertisement

draft-ietf-mpls-ldp-ip-pw-capability-05.txt

Status of this Memo

Copyright Notice

Abstract

Currently, no LDP capability is exchanged for LDP applications like IP Label Switching and L2VPN P2P PW signaling. When an LDP session comes up, an LDP speaker may unnecessarily advertise its local state for such LDP applications even when the peer session is established for some other applications like mLDP or ICCP. This document defines a solution by which an LDP speaker announces to its peer its disinterest in such non-negotiated applications. This, in turn, disables the advertisement of corresponding application state, which would have otherwise be advertised by default, over the established LDP session.

Table of Contents

1. Introduction

LDP Capabilities [RFC5561] introduced a mechanism to negotiate LDP capabilities for a given feature between peer LSRs. The capability mechanism insures that no unnecessary state is exchanged between peer LSRs unless the corresponding feature capability is successfully negotiated between the peers.

While new LDP features and applications, such as Typed Wildcard FEC [RFC5918], Inter-Chassis Communication Protocol [ICCP], mLDP [RFC6388], and L2VPN P2MP PW [P2MP-PW] make use of LDP capabilities framework for their feature negotiation, the earlier LDP features and applications like IP Label Switching and L2VPN P2P PW signaling [RFC4447] [RFC4762] may cause LDP speakers to exchange application state unnecessarily even when the given application is not enabled on one of the LDP speakers participating in a given session. For example, when bringing up and using an LDP peer session with a remote PE LSR for purely ICCP signaling reasons, an LDP speaker may unnecessarily advertise labels for IP (unicast) prefixes to this ICCP related LDP peer.

Another example of unnecessary state advertisement can be cited when LDP is to be deployed in an IP dual-stack environment. For instance, an LSR that is locally enabled for both IPv4 and IPv6 label switching may advertise label bindings for both IPv4 and IPv6 address families towards an LDP peer that is interested in IPv4 prefix labels only. In this case, the advertisement of IPv6 prefix labels to the peer is unnecessary, as well as wasteful, from the point of view of LSR memory/CPU and network resource consumption.

To avoid this unnecessary state advertisement and exchange, currently an operator is typically required to configure and define some sort of filtering policies on the LSR, which introduces operational overhead and complexity for such deployments.

This document defines an LDP Capabilities [RFC5561] based solution by which an LDP speaker may announce to its peer(s) its disinterest (or non-support) for state related to IP Label Switching and/or L2VPN P2P PW Signaling application at the time of session establishment. This helps in avoiding unnecessary state advertisement for such feature applications. The document also states the mechanics to dynamically

disable or enable the state advertisement for such applications during the session lifetime. The non-interesting state of an application depends on the type of application and is described later in section 3.1.

2. 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 RFC-2119 [RFC2119].

The term "IP" in this document refers to both IPv4 and IPv6 unicast address families.

This document uses shorthand terms "IPoMPLS" to refer to IP Label Switching application, and "P2P PW" to refer to L2VPN PW signaling for FEC 128 and FEC 129 P2P Pseudowires.

3. Non-negotiated LDP applications

For the applications that existed prior to the definition of LDP Capabilities framework [RFC5561], an LDP speaker typically advertises, without waiting for any capabilities exchange and negotiation, its corresponding application state to its peers right after the session establishment. These early LDP applications include:

- IPv4/IPv6 Label Switching ("IPoMPLS")

- L2VPN P2P PW signaling ("P2P PW")

To disable unnecessary state advertisement for such LDP applications over an established LDP session, a new capability is introduced in this document. This new capability controls the advertisement of application state and enables an LDP speaker to notify its peer its disinterest in the state of one or more of these "Non-negotiated" LDP applications at the time of session establishment. Upon receipt of such capability, the receiving LDP speaker, if supporting the capability, disables the advertisement of the state related to the application towards the sender. This new capability can also be sent later in a Capability message to either disable a previously enabled application's state advertisement or to enable a previously disabled application's state advertisement.

3.1. Non-interesting State

So far, this document has used the term application "state" to generically refer to some non-interesting state. Now, let us further specify and clarify this term:

- A non-interesting state of a non-negotiated application refers to the application state which is of a no interest to an LSR and need not be advertised to the LSR;
- This state MUST NOT be advertised in any of the LDP protocol messages;
- This state is dependent on application type and specified accordingly.

For IPoMPLS application type, the non-interesting state refers to any state related to IP Prefix FEC (such as label bindings, LDP Status). This document, however, does not classify IP address bindings as a non-interesting state and allows the advertisement of IP Address bindings to facilitate other LDP applications (such as mLDP) that depend on learning of peer addresses over an LDP session for their correct operation.

For P2P PW application type, the non-interesting state refers to any state related to P2P PW FEC (such as label bindings, MAC [address] withdrawal, and LDP PW Status).

From now onward in this document, the term "state" will mean to refer to the "non-interesting state" for an application, as defined in this section.

4. Controlling State Advertisement for Non-negotiated LDP Applications

To control advertisement of non-interesting state related to non-negotiated LDP applications, namely IPoMPLS and P2P PW signaling, a new capability TLV is defined as follows.

4.1. State Advertisement Control Capability

The "State Advertisement Control Capability" is a new Capability Parameter TLV defined in accordance with section 3 of LDP Capabilities specification [RFC5561]. The format of this new TLV is as follows:

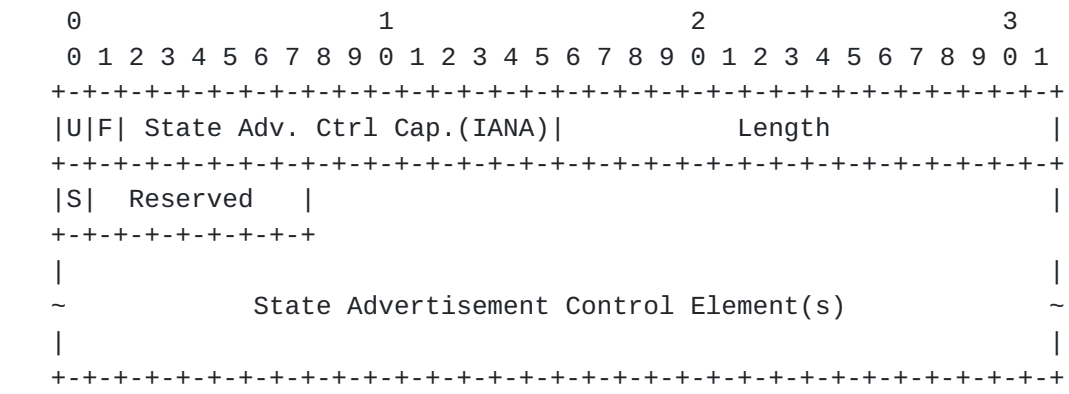

Figure 1: Format of an "State Advertisement Control Capability" TLV

The value of the U-bit for the 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. Whereas, The value of F-bit MUST be set to 0. Once advertised, this capability cannot be withdrawn; thus S-bit MUST be set to 1 both in an Initialization and Capability message.

The capability data associated with this State Advertisement Control (SAC) Capability TLV is one or more State Advertisement Control (SAC) Elements, where each element indicates enabling/disabling of advertisement of non-interesting state for a given application. The format of a SAC Element is defined as follows:

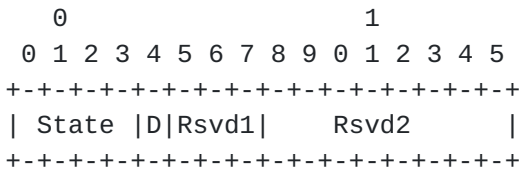

Figure 2: Format of an "State Advertisement Control Element"

Where:

State: Defines the type of application state (to be controlled).
The value of this field is defined as follows:
1: IPv4 Label switching
2: IPv6 Label switching
3: P2P PW FEC128 signaling
4: P2P PW FEC129 signaling
0, 5-15: Reserved.

D bit: Controls the advertisement of the state:
1: Disable state advertisement
0: Enable state advertisement
When sent in an Initialization message, D bit MUST be set to 1.

Rsvd1, Rsvd2: Reserved for future use. MBZ on transmit and ignored on receipt.

The "Length" field of SAC Capability TLV depends on the number of SAC Elements present in the TLV. For example, if there are two elements present, then the Length field is set to 5 octets. A receiver of this capability TLV can deduce number of elements present in the TLV by using the Length field.

From now onward, this document uses the term "element" to refer to a SAC Element.

As described earlier, SAC Capability TLV MAY be included by an LDP speaker in an Initialization message to signal to its peer LSR that state advertisement for one or more application(s) need to be disabled on the given peer session. This TLV can also be sent later in a Capability message to selectively enable or disable these applications. A SAC Capability TLV MUST contain elements with distinct state types and the TLV MUST NOT contain the same state type more than once. If a receiver receives such a malformed TLV, it SHOULD discard this TLV and continue processing rest of the message.

To control more than one application state, a sender LSR can either send a single capability TLV in a message with multiple elements present, or can send separate messages with capability TLV specifying one or more elements. A receiving LSR, however, MUST treat each incoming capability TLV for a given application state type as an update to its existing policy for the given type.

To understand capability updates from an example, let us consider 2 LSRs, S (LDP speaker) and P (LDP peer), both of which support all the non-negotiated applications listed earlier. By default, these LSR will advertise state for these applications, as configured, to their peer as soon as an LDP session is established. Now assume that P receives from S a SAC capability in the Initialization message with "IPv6 Label switching" and "P2P PW FEC129" states disabled. This updates P's outbound policy towards S to advertise state related to only "IPv4 Label switching" and "P2P PW FEC 128" applications. Later, P receives another capability update from S via a Capability message with "IPv6 Label switching" enabled and "P2P PW FEC128" disabled. This results in P's outbound policy towards S to advertise both IPv4 and IPv6 Label switching state, and

disable both P2P PW FEC128 and FEC 129 signaling. Finally, P receives another update from S via a Capability message that specifies to disable all four non-negotiated applications state, resulting in P outbound policy towards S to block/disable state for all these applications, and only advertise state for any other application, as applicable.

5. Capabilities Procedures

The SAC capability conveys the desire of an LSR to disable the receipt of unwanted/unnecessary state from its LDP peer. This capability is uni-lateral and uni-directional in nature, and a receiving LSR is not required to send a similar capability TLV in an Initialization or Capability message towards the sender of this capability. This unilateral behavior conforms to the procedures defined in the Section 6 of LDP Capabilities [RFC5561].

After this capability is successfully negotiated (i.e. sent by an LSR and received/understood by its peer), then the receiving LSR MUST NOT advertise any state related to the disabled applications towards the capability sending LSR until and unless these application states are explicitly enabled again via a capability update. Upon receipt of a capability update to disable an enabled application [state] during the lifetime of a session, the receiving LSR MUST also withdraw from the peer any previously advertised state (corresponding to the disabled application).

If a receiving LDP speaker does not understand the SAC capability TLV, then it MUST respond to the sender with "Unsupported TLV" notification as described in LDP Capabilities [RFC5561]. If a receiving LDP speaker does not understand or does not support an application specified in an application control element, it SHOULD silently ignore/skip such an element and continue processing rest of the TLV.

5.1. State Control Capability in an Initialization message

LDP Capabilities [RFC5561] framework dictates that the S-bit of capability parameter in an Initialization message MUST be set to 1 and SHOULD be ignored on receipt.

An LDP speaker determines (e.g. via some local configuration or default policy) if it needs to disable IPoMPLS and/or P2P PW applications with a peer LSR. If there is a need to disable, then the SAC TLV needs to be included in the Initialization message with respective SAC elements included with their D bit set to 1.

An LDP speaker that supports the SAC capability MUST interpret the capability TLV in a received Initialization message such that it disables the advertisement of the application state towards the capability sending LSR for IPoMPLS and/or P2P PW applications if their SAC element's D bit is set to 1.

5.2. State Control capability in a Capability message

If the LDP peer supports "Dynamic Announcement Capability" [RFC5561], then an LDP speaker may send SAC capability in a Capability message towards the peer. Once advertised, these capabilities cannot be withdrawn and hence the S-bit of the TLV MUST be set to 1 when sent in a Capability message.

An LDP speaker may decide to send this TLV towards an LDP peer if one or more of its IPoMPLS and/or P2P PW signaling applications get disabled, or if previously disabled application gets enabled again. In this case, the LDP speaker constructs the TLV with appropriate SAC element(s) and sends the corresponding capability TLV in a Capability message.

Upon receipt of this TLV in a Capability message, the receiving LDP speaker reacts in the same manner as it reacts upon the receipt of this TLV in an Initialization message. Additionally, the peer withdraws/advertises the application state from/to the capability sending LDP speaker according to the capability update.

6. Operational Examples

6.1. Disabling IPoMPLS and P2P PW applications on an ICCP session

Consider two PE routers, LSR1 and LSR2, which understand/support SAC capability TLV, and have an established LDP session to exchange ICCP state related to dual-homed devices connected to these LSRs. Let us assume that both LSRs are provisioned not to exchange any state for IPoMPLS (IPv4/IPv6) and P2P PW (FEC128/129) application.

To indicate their disinterest in these applications, the LSRs will include a SAC capability TLV (with 4 SAC elements corresponding to these 4 applications with D bit set to 1 for each one) in the Initialization message. Upon receipt of this TLV in Initialization message, the receiving LSR will disable the advertisement of IPv4/IPv6 label bindings, as well as P2P PW FEC128/129 signaling, towards its peer after session establishment.

6.2. Disabling IPoMPLS application on a L2VPN/PW T-LDP session

Now, consider LSR1 and LSR2 have an established T-LDP session for P2P PW application to exchange label bindings for FEC 128/129. Given that there is no need to exchange IP (v4/v6) label bindings between the PE LSRs over a PW T-LDP session in most typical deployments, let us assume that LSRs are provisioned to disable IPoMPLS (IPv4/IPv6) application state on given PW session.

To indicate their disinterest in IPoMPLS application over a PW T-LDP session, the LSRs will follow/apply the same procedures to disable IPv4 and IPv6 label switching as described in previous section. As a result, only P2P PW related state will be exchanged between these LSRs over this T-LDP session.

6.3. Disabling IPoMPLS application dynamically on an established IP/PW LDP session

Assume that LSRs from previous sections were initially provisioned to exchange both IPoMPLS and P2P PW state over the session between them, and also support "Dynamic Announcement" Capability [RFC5561]. Now, assume that LSR1 is dynamically provisioned to disable IPoMPLS (IPv4/IPv6) over T-LDP session with LSR2. In this case, LSR1 will send SAC capability TLV in a Capability message towards LSR2 with application control elements defined for IPv4 and IPv6 label switching with D bit set to 1. Upon receipt of this TLV, LSR2 will disable IPoMPLS application state(s) towards LSR1 and withdraw all previous application state from LSR1. To withdraw label bindings from its peer, LSR2 MAY use a single Prefix FEC Typed Wildcard Label Withdraw message [RFC5918].

This dynamic disability of IPoMPLS application does not impact L2VPN P2P PW application on the given session, and both LSRs should continue to exchange PW Signaling application related state.

6.4. Disabling IPoMPLS application on an mLDP-only session

Now assume that LSR1 and LSR2 have formed an LDP session to exchange mLDP application state only. In typical deployments, LSR1 and LSR2 also exchange label bindings for IP prefixes over an mLDP session, which is unnecessary and wasteful for an mLDP-only LSR.

Using the procedures defined earlier, an LSR can indicate its disinterest in IPoMPLS application state to its peer upon session establishment time or dynamically later via LDP capabilities update.

Reference to section 3.1, the peer disables the advertisement of any state related to IP Prefix FECs, but still advertises IP address bindings that are required for the correct operation of mLDP.

6.5. Disabling unwanted IP state advertisement by an IP dual-stack LSR

In IP dual-stack scenarios, an LSR2 may advertise unnecessary state (e.g. IPv6 prefix label bindings) towards peer LSR1 corresponding to IPv6 label switching application once a session is established mainly for exchanging state for IPv4. The similar scenario also applies when advertising IPv4 label switching state on a session meant for IPv6. The SAC capability and its procedures defined in this document can help to avoid such unnecessary state advertisement.

Consider IP dual-stack environment where LSR2 is enabled for IPoMPLS application for both IPv4 and IPv6, but LSR1 is enabled for (or interested in) only IPv4oMPLS Label switching. To avoid receiving unwanted state advertisement for IPv6oMPLS Label switching application from LSR2, LSR1 can send SAC capability with element for IPv6 label switching with D bit set to 1 in the Initialization message towards LSR2 at the time of session establishment. Upon receipt of this capability, LSR2 will disable all IPv6 label binding advertisement towards LSR1. If IPv6oMPLS Label switching application is later enabled on LSR1, LSR1 can update the capability by sending SAC capability in a Capability message towards LSR2 to enable IPv6oMPLS Label switching application dynamically.

7. Security Considerations

The proposal introduced in this document does not introduce any new security considerations beyond that already apply to the base LDP specification [RFC5036] and [RFC5920].

8. IANA Considerations

This document defines a new LDP cpability parameter TLV. IANA is requested to assign the lowest available value after 0x0500 from "TLV Type Name Space" in the "Label Distribution Protocol (LDP) Parameters" registry within "Label Distribution Protocol (LDP) Name Spaces" as the new code point for the LDP TLV code point.

Range	Description	Reference	Notes/Registration Date
TBA	State Advertisement Control Capability	This document	

9. References

9.1. Normative References

[RFC5036] L. Andersson, I. Minei, and B. Thomas, "LDP Specification", RFC 5036, September 2007.

[RFC5561] B. Thomas, K. Raza, S. Aggarwal, R. Aggarwal, and JL. Le Roux, "LDP Capabilities", RFC 5561, July 2009.

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

9.2. Informative References

[RFC5918] R. Asati, I. Minei, and B. Thomas, "Label Distribution Protocol Typed Wildcard FEC", RFC 5918, August 2010.

[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.

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

[P2MP-PW] Martini, L. et. al, "Signaling Root-Initiated Point-to-Multipoint Pseudowires using LDP", draft-ietf-pwe3-p2mp-pw-04.txt, Work in Progress, March 2012.

[ICCP] L. Martini, S. Salam, A. Sajassi, and S. Matsushima, "Inter-Chassis Communication Protocol for L2VPN PE Redundancy", draft-ietf-pwe3-iccp-11.txt, Work in Progress, February 2013.

[RFC6388] I. Minei, I. Wijnands, K. Kompella, and B. Thomas, "LDP Extensions for P2MP and MP2MP LSPs", RFC 6388, November 2011.

[RFC5920] L. Fang, et al., "Security Framework for MPLS and GMPLS Networks", RFC 5920, July 2010.

10. Acknowledgments

The authors would like to thank Eric Rosen for his valuable input and comments. We also acknowledge Karthik Subramanian and IJsbrand Wijnands for bringing up mLDP use case.

This document was prepared using 2-Word-v2.0.template.dot.

Authors' Addresses

Kamran Raza
Cisco Systems, Inc.,
2000 Innovation Drive,
Ottawa, ON K2K-3E8, Canada.
E-mail: skraza@cisco.com

Sami Boutros
Cisco Systems, Inc.
3750 Cisco Way,
San Jose, CA 95134, USA.
E-mail: sboutros@cisco.com