Network Working Group
Internet Draft
Intended status: Standards Track
Expires: October 19, 2011

Kamran Raza
Cisco Systems

Sami Boutros
Cisco Systems

April 20, 2011

LDP IP and PW Capability

draft-raza-mpls-ldp-ip-pw-capability-01.txt

Status of this Memo

Copyright Notice

Abstract

Currently, no LDP capability is exchanged for LDP applications like IP label switching and L2VPN/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 may be established for some other applications like ICCP. This document proposes a solution by which an LDP speaker announces its "incapability" or disability or non-support for IP label switching or L2VPN/PW application, hence disabling corresponding application state exchange over the established LDP session.

Table of Contents

1. Introduction

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

While new features and applications, such as Typed Wildcard FEC [RFC5918], Inter-Chassis Communication Protocol [ICCP], mLDP [MLDP], make use of LDP capabilities framework for their feature negotiation, the earlier LDP features and applications like IP label switching and L2VPN/PW signaling [RFC4447] may cause unnecessary state exchange between LDP peers 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 purposes, the LDP speaker may unnecessarily advertise labels for IP (unicast) prefixes to this ICCP related LDP peer as per its default behavior. To avoid this unnecessary state advertisement and exchange, currently customers are typically required to configure/define some sort of LDP state (label) filtering policies on the box, which introduces operational overhead and complexity.

This document proposes a solution by which an LDP speaker may announce its "incapability" (or disability) to its peer for IP Label Switching and/or L2VPN/PW Signaling application at session establishment time. This helps avoiding unnecessary state exchange for such feature applications. The proposal also state the mechanics to enable previously disabled application later during the session lifetime. The document introduces two new LDP Capabilities for IP label switching and L2VPN/PW applications to implement this proposal.

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 "IP unicast", unless otherwise explicitly stated.

3. Non-negotiated LDP applications

For the applications that existed before LDP Capabilities [RFC5561] mechanics were defined, LDP speaker may advertise relevant application state to its peers after session establishment without waiting for any capabilities exchange and negotiation.

Currently, the most important non-negotiated applications include:

- IP [v4 and v6] label switching
- L2VPN/PW signaling

To disable unnecessary state exchange for such LDP applications, two new capabilities are being introduced in this document. These new capabilities allow an LDP speaker to notify its LDP peer at the session establishment time when one or more LDP "Non-negotiated applications" are not required/configured on the sender side. Upon receipt of such capability, if supported, the receiving LDP speaker MUST disable the advertisement of application state towards the sender. These capabilities can also be sent later in a Capability message to either disable these applications, or to enable previously disabled applications.

3.1. Application Control Capabilities

To control advertisement of state related to non-negotiated LDP applications, namely IP Label switching and L2VPN/PW signaling, two new capability TLVs are defined as described in the following subsections.

3.1.1. IP Label Switching Capability TLV

The IP Label Switching capability is a new Capability Parameter defined with the following format:

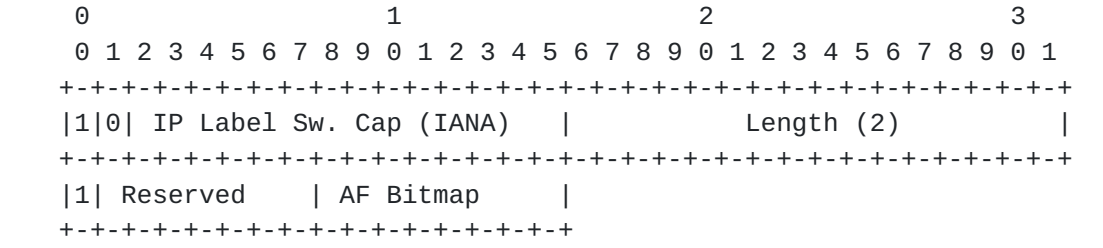

The value of the U-bit for the IP capability parameter TLV MUST be set to 1 so that a receiver MUST silently ignore this TLV if unknown

to it, and continue processing the rest of the message. Once advertised, this capability cannot be withdrawn and hence the S-bit must always be set to 1 both in Initialization message and Capability message. The capability data associated with this TLV is 1 octet long "Address Family Bitmap", and hence the TLV length MUST be set to 2.

The Capability data "Address Family Bitmap" is defined as:

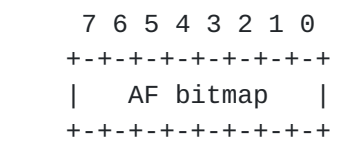

Where:

bit0: IPv4 label switching application

bit1: IPv6 label switching application

bit2-7: Reserved.

A bit in the bitmap is set to 0 or 1 to disable or enable respectively a corresponding IP application.

As described earlier, "IP Label Switching" Capability Parameter TLV MAY be included by an LDP speaker in an Initialization message to signal to its peer LSR that state exchange for IPv4 and/or IPv6 application(s) need to be disabled on a given peer session. This TLV can also be sent later in a Capability message to selectively enable or disable IPv4/v6 label switching application(s).

3.1.2. PW Signaling Capability TLV

The "PW Signaling" capability is a new Capability Parameter defined with the following format:

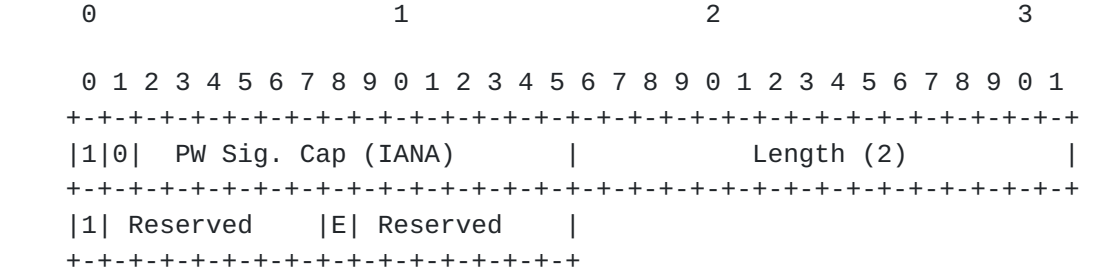

The value of the U-bit for the PW capability parameter TLV MUST be set to 1 so that a receiver MUST silently ignore this TLV if unknown to it, and continue processing the rest of the message. Once advertised, this capability cannot be withdrawn and hence the S-bit MUST always be set to 1 in Initialization message or Capability message. The capability data associated with this TLV is 1 octet long and hence the TLV length MUST be set to 2.

The capability data is defined as following byte:

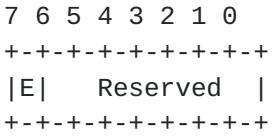

Where E-bit (Enable bit) is used to control PW signaling application by setting it to 0 and 1 to disable and enable the application respectively.

As described earlier, PW Signaling Capability Parameter TLV MAY be included by an LDP speaker in an Initialization message to signal to its peer LSR that state exchange for PW application need to be disabled on given peer session. This TLV can also be sent later in a Capability message to enable/disable the PW Signaling application.

3.2. Procedures for Application Control Capabilities in an Initialization message

LDP Capabilities [RFC5561] dictate 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 they need to disable IP and/or L2VPN/PW applications with a peer LSR. If there is a need to disable, then the IP and/or PW application capability TLVs need to be included in the Initialization message with respective application bits set to 0 to indicate application disable, where the application bit refers to a bit in "Address Family Bitmap" of the "IP Label Switching" Capability or E-bit in "PW Signaling" Capability.

An LDP speaker that supports the "IP Label Switching" and/or "PW Signaling" capability MUST interpret those TLVs in a received Initialization message such that it disables the advertisement of the

application state towards the sender LSR for IP (v4 and/or v6) and/or L2VPN/PW applications if their application control bits are set to 0. If a receiving LDP speaker does not understand the capability TLVs, then it MUST respond to the sender with "Unsupported TLV" Notification as described in LDP Capabilities [RFC5561]. Upon receipt of such Notification, the sender MAY still continue to block/disable its outbound state advertisement towards the peer for the requested disabled applications.

Once this capability has been sent by sender LSR and received and understood by the receiver LSR, then both these LSRs MUST NOT exchange any state related to the disabled applications until and unless these applications are explicitly enabled again (e.g. via the same Capability TLV sent in a Capability message with corresponding application control bit set to 1).

"IP Label Switching" and "PW Signaling" capability TLVs are unilateral/uni-directional in nature. This means that the receiving LSR may not need to send a similar capability TLV in an Initialization or Capability message towards the sender. This unilateral behavior also conforms to the procedures defined in the Section 6 of LDP Capabilities [RFC5561].

3.3. Procedures for Application Control capabilities in a Capability message

If the LDP peer supports "Dynamic Announcement Capability" [RFC5561], then an LDP speaker can send IP Label Switching and/or PW Signaling capability in a Capability message. 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 any of its IP and/or L2VPN/PW signaling applications gets disabled, or if previously disabled IP and/or L2VPN/PW applications gets enabled again. In this case, LDP speaker constructs the TLVs with appropriate application control bitmap and sends the corresponding capability TLVs in a Capability message. Furthermore, the LDP speaker also withdraws application(s) related advertised state (such as label bindings) from its peer.

Upon receipt of those TLVs in a Capability message, the receiving LDP speaker reacts in the same manner as it reacts upon the receipt of those TLVs in an Initialization message. Additionally, the receiving LDP speaker withdraws the application(s) related advertised state (such as label bindings) from the sending LDP speaker. If the receiving LDP speaker does not understand or support either Dynamic

Announcement capability or received Application Control capability TLV ("IP Label Switching" or "PW Signaling"), it MUST respond with "Unsupported Capability" notification to the sender of the Capability message.

4. Operational Examples

4.1. Disabling IP/PW label applications on an ICCP session

Consider two PE routers, LSR1 and LSR2, which understand/support "IP Label Switching" and "PW Signaling" capability TLVs. These LSR have an established LDP session due to ICCP application in order to exchange ICCP state related to dual-homed devices connected to these LSRs. Let us assume that LSR1 is provisioned not to exchange any label bindings related to IP (v4/v6) prefixes and PW layer2 FEC (FEC128/129) with LSR2.

To indicate its "disability" for the IP/PW applications, the LSR1 will include both the "IP Label Switching" capability TLV (with bit0-1 of "Address Family Bitmap" set to 0) and "PW Signaling" capability TLV (with E-bit set to 0) in the Initialization message. Upon receipt of those TLVs in Initialization message, the LSR2 will disable any IP/PW address/label binding state advertisement towards LSR1 after session establishment.

The LSR1 will also disable any IP/PW address/label binding state towards LSR2, irrespective of the fact whether or not LSR2 could disable the corresponding application state advertisement towards LSR1.

4.2. Disabling IP Label Switching application on a L2VPN/PW session

Now, consider LSR1 and LSR2 have an established session due to L2VPN/PW application just to exchange PW (FEC128/129) label bindings for VPWS/VPLS services amongst them. Since in most typical deployments, there is no need to exchange IP (v4/v6) address/label bindings amongst the PE LSRs, let us assume that LSR1 is provisioned to disable IP (v4/v6) application on given PW session towards LSR2.

To indicate its disability for IP application, the LSR1 will include the "IP Label Switching" capability TLV in the Initialization message with bit0-1 (IPv4, IPv6) in "Address Family Bitmap" set to zero. Upon receipt of this TLV in Initialization message, the LSR2 will disable any IP address/label binding state advertisement towards LSR1.

The LSR1 will also disable any IP address/label binding state towards LSR2, irrespective of the fact whether or not LSR2 could disable the corresponding IP application state advertisement towards LSR1.

4.3. Disabling IP application dynamically on an established IP/PW session

Assume that LSRs from previous sections were initially provisioned to exchange both IP and PW state over the session between them, and also support "Dynamic Announcement" capability [RFC5561]. Now, assume that LSR1 is dynamically provisioned to disable IP label switching with LSR2. In this case, LSR1 will first withdraw all its IP label state by sending a single Label Withdraw message with IP Prefix Typed Wildcard FEC using the mechanics described in [RFC5918], and Address Withdraw message to withdraw its addresses. LSR1 will also send IP Label Switching capability TLV in Capability message towards LSR2 with bit0-1 (IPv4, IPv6) in "Address Family Bitmap" set to zero. Upon receipt of this TLV, LSR2 will also disable IP label switching towards LSR1 and withdraw all previous IP application label/address state using the same mechanics as described earlier for LSR1. The disability of IP label switching dynamically should not impact L2VPN/PW application on given session, and both LSRs should continue to exchange PW Signaling application related state.

5. 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].

6. IANA Considerations

The document introduces following two new capability parameter TLVs and requests following LDP TLV code point assignment by IANA:

- "IP Label Switching" Capability TLV (requested codepoint: 0x50C)
- "PW Signaling" Capability TLV (requested codepoint: 0x50D)

7. Conclusions

The document proposed a solution using LDP Capabilities [RFC5561] mechanics to disable unnecessary state exchange, if/as desired, between LDP peers for currently non-negotiated IP/PW applications.

8. References

8.1. Normative References

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

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

[ICCP] Martini, L., Salam, S., and Matsushima, S., "Inter-Chassis Communication Protocol for L2VPN PE Redundancy", draft-ietf-pwe3-iccp-04.txt, Work in Progress, October 2010.

[MLDP] Minei, I., Kompella, K., Wijnands, I., and Thomas, B., "LDP Extensions for Point-to-Multipoint and Multipoint-to-Multipoint Label Switched Paths", draft-ietf-mpls-ldp-p2mp-10.txt, Work in Progress, July 2010.

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

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

8.2. Informative References

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

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

9. Acknowledgments

The authors would like to thank Eric Rosen for his valuable input and comments.

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

Authors' Addresses

Kamran Raza
Cisco Systems, Inc.,
2000 Innovation Drive,
Kanata, 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