MPLS Working Group
Internet-Draft
Intended status: Standards Track
Expires: January 9, 2014

IJ. Wijnands, Ed.
Cisco
A. Gulko
Thomson Reuters
U. Joorde
Deutsche Telekom
J. Tantsura
Ericsson

July 8, 2013

mLDP in-band signalling Wildcard encoding
draft-wijnands-mpls-mldp-in-band-wildcard-encoding-00

Abstract

Documents [RFC6826] and [I-D.ietf-l3vpn-mldp-vrf-in-band-signaling] define a solution to splice an IP multicast tree together with a multipoint LSP in the global or VRF context. In these drafts the Multipoint Label Distribution Protocol (mLDP) Opaque TLV encodings have been documented for Source specific and Bidir IP multicast trees. For each IP multicast tree a multipoint LSP is created. There are scenarios where it is beneficial to support shared trees and allow aggregation such that fewer multipoint LSPs are created in the network. This document defines wildcard encodings to be used for the Source or Group fields of the existing opaque encodings. With the wildcard encoding it is possible to create a single multipoint LSP that is used to represent *all* sources for a given multicast group or *all* groups for a given source.

Status of this Memo

Copyright Notice

Table of Contents

1. Terminology and Definitions

PIM: Protocol Independent Multicast.

IGMP: Internet Group Management Protocol.

MLD: Multicast Listener Discovery.

IP multicast tree: An IP multicast distribution tree identified by a IP multicast group address and optionally a Source IP address, also referred to as (S,G) and (*,G).

MP-LSP: A P2MP or MP2MP LSP.

PIM-ASM: PIM Any Source Multicast.

PIM-SSM: PIM Source Specific Multicast.

PIM-SM: PIM Sparse-mode Multicast.

RP: The PIM Rendezvous Point.

mLDP: Multipoint LDP.

In-band signaling: Using the opaque value of a mLDP FEC element to carry the (S,G) or (*,G) identifying a particular IP multicast tree.

Ingress LSR: Source of the P2MP LSP, also referred to as root node.

Egress LSR: A LSR that has receivers attached, also referred to as leaf node.

Threshold Infinity: A PIM-SM procedure where no source specific multicast (S,G) trees are created for multicast packets that are forwarded down the shared tree (*,G).

2. Introduction

Documents [RFC6826] and [I-D.ietf-l3vpn-mldp-vrf-in-band-signaling] define a solution to splice an IP multicast tree together with a multipoint LSP in the global or VRF context. In these drafts the Multipoint Label Distribution Protocol (mLDP) Opaque TLV encodings have been documented for Source specific and Bidir IP multicast trees. For each IP multicast tree a mLDP MP-LSP is created. There are scenarios where it is beneficial to support shared trees and allow aggregation such that fewer multipoint LSPs are created in the

network. This document defines wildcard encodings that can be used in Source or Group fields of the existing Opaque TLV encodings. With the wildcard encoding it is possible to create a single multipoint LSP used to represent *all* sources for a given multicast group or *all* groups for a given source.

The behaviour of an mLDP in-band signalled multipoint LSPs containing a wildcard entry follows the procedures defined in [RFC6826] and [I-D.ietf-l3vpn-mldp-vrf-in-band-signaling]. This draft does not talk about already defined procedures but only documents the differences.

There are a few scenarios (not limited to) where wildcard encoding is useful, for example;

o PIM Shared tree forwarding with threshold infinity.

o IGMP/MLD proxying.

o Selective Source mapping.

These scenarios are discussed in this draft below.

3. PIM shared tree forwarding

PIM [RFC4601] has the concept of a shared tree, known as (*,G). This means, *all* Sources for a given Group in the ASM range. The (*,G) is built towards the Rendezvous Point (RP) that typically joins *all* multicast sources for this group. The RP will then forward all the IP multicast packets for this group down the (*,G) tree towards the receivers. There are several procedures how the RP learns about these sources, for example; PIM Registers [RFC4601], MSDP [RFC3618] or a Source that is directly connected to the RP. In some cases, the last hop routers does not wish to join the source trees, and expect to receive all the traffic for group G from the (*,G) tree; in this case, we say that the last hop routers have 'threshold infinity' for group G. This is optional behaviour documented in the [RFC4601]. This is often used in deployments where the RP is between the multicast sources and the multicast receivers for group G, i.e., the shortest path from any source to any receiver of the group goes through the RP. In this scenario, there is no advantage for a last hop router to join a source tree for the group, since joining a source tree would not change the path of the multicast data from the source. The only effect of executing the complicated procedures for joining a source tree and pruning the source off the shared tree would be to an increase of the amount of multicast routing state.

This deployment model can be implemented using wildcards. The egress router will splice the (*,G) IP multicast tree to a mLDP Multipoint LSP where the Source address is encoded as wildcard entry. In scenarios where it does not make sense to apply "threshold infinity" to a given ASM group, a more complex set of procedures are needed, as per [I-D.rekhter-pim-sm-over-mldp].

4. IGMP/MLD proxying

There are scenarios where the multicast senders and receivers are directly connected to a MPLS routing domain and mLDP is available. In these cases we can apply "IGMP/MLD proxying" and avoid using PIM as a multicast routing protocol to transport multicast packets from the senders to the receivers. The senders and receivers consider the

MPLS domain to be single hop between each other. [RFC4605] documents procedures where a multicast routing protocol is not nessesary to build a 'simple tree'. Within the MPLS domain mLDP will be used to build a 'spanning tree' to avoid looping and duplication of packets, but for the point of view of the senders and receivers this is hidden. The procedures as defined [RFC4605] are applicable since the senders and receivers are considered to be one hop away from each other.

For mLDP to build a tree, it needs to know the root of the tree. Following the procedures as defined in [RFC4605] we depend on manual configuration of the mLDP root for the ASM multicast group. The Source will be encoded as a wildcard entry.

5. Selective Source mapping

In IPTV deployments, rather often, the content servers are co-located in a few sites. Popular channels are often statically configured and always forwarded over the core MPLS network to the egress routers. Since these channels are statically defined, they MAY also be forwarded over a multipoint LSP with wildcard encoding. The sort of wildcard encoding that needs to be used (Source and/or Group) depends on the Source/Group allocation policy of the IPTV provider. Other options are to use MSDP [RFC3618] or BGP AD [RFC6513] for source discovery by the ingress LSR. Based on the received wildcard, the ingress LSR can make a selection out of the IP multicast streams it has state for.

6. Wildcard Source

When the IP multicast component on the ingress LSR has received a

wildcard source from mLDP it may have been initiated by one of the scenarios described in this draft. How the wildcard source is to be interpretted is a local matter and follows the rules below;

1. If PIM is enabled and the group is a non-bidirectional ASM group, the wildcard source is treated as having received a (*,G) IGMP/MLD report from a downstream node and the procedures as defined in [RFC4601] are followed.

2. If PIM is enabled and the group mode is PIM-SSM, all multicast sources known for the group on the root node must be forwarded down the multipoint LSP.

3. If PIM is not enabled for this group, the wildcard source is treated as having received a (*,G) IGMP/MLD report from a downstream node and the procedures as defined in [RFC4605] are followed.

The IP multicast component on the egress LSR determins when a Wildcard Source is to be used in a mLDP Opaque TLV encoding. How the IP multicast component determins this is a local matter and potentially subjected to explicit user configuration. It MAY however use the following rules (with or without explicit user configuration);

1. If PIM is enabled, the group is a non-bidirectional ASM group and the RP is reachable via a BGP route, a Wildcard Source encoding MAY be used to signal group membership (*,G) to the ingress LSR, using the BGP next hop as the ingress LSR (root of the LSP). Also see Section 8.

2. If PIM is not enabled for this group and an IGMP/MLD group membership report has been received, the IP multicast component may use a Wildcard Source encoding to signal the group membership (*,G) to a Proxy device (root of the LSP). The procedures how to determin the Proxy device for a given group are defined in [RFC4605].

The wildcard source encoding MUST NOT appear in the "Bidir TLVs" that are defined in [RFC6826] sections 3.3 and 3.4."

A wildcard group in combination with a wildcard source encoding is under investigation.

7. Wildcard Group

When the IP multicast component on the root node receives a wildcard

group encoding, the root node SHOULD apply the wildcard encoding to

the existing IP multicast routing table and forward all the IP multicast stream(s) that match the given Source. Note, this behaviour is independent of the PIM group mode (ie. ASM or SSM).

The IP multicast component on the egress LSR determins when a Wildcard Group is to be used in a mLDP Opaque TLV encoding. How the IP multicast component determins this is a local matter and subjected to explicit user configuration.

The wildcard group encoding for PIM bidir is under investigation.

A wildcard source in combination with a wildcard group encoding is under investigation.

8. Root node address discovery

Documents [RFC6826] and [I-D.ietf-l3vpn-mldp-vrf-in-band-signaling] describe procedures to discover the mLDP root node address by using the Source of the IP multicast stream. When a wildcard source encoding is used, PIM is enabled and the group is a non-bidirectional ASM group, a similar procedure s applied. The only different with the above mentioned procedures is that the Proxy device or RP address is used instead of the Source to discover the mLDP root node address.

In all other cases some sort of manual configuration is applied in order to find the root node. Note, finding the root node is a local implementation matter and not limited to the solutions mentioned in this draft.

9. Anycast RP

With in-band signalling there is likely no RP to Group mappings distribution taking place over the MPLS core to the different IP multicast sites. The RP address is likely statically configured on each multicast site. In these cases it makes sense to configure an Anycast RP Address to provide redundancy. See [RFC3446] for more details.

10. Wildcard encoding

The source and group fields in the Transit IPv4, IPv6, VPNv4 and VPNv6 Source TLVs, as documented in [RFC6826] and [I-D.ietf-l3vpn-mldp-vrf-in-band-signaling] only allow valid IP addresses to be encoded. This document proposes to use a source/

group field of *all* zero's to be used as wildcard encoding.

11. Acknowledgements

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

12. IANA Considerations

There are no new allocations required from IANA.

13. Security Considerations

There are no security considerations other then ones already mentioned in [RFC6826] and [I-D.ietf-l3vpn-mldp-vrf-in-band-signaling].

14. References

14.1. Normative References

[I-D.ietf-l3vpn-mldp-vrf-in-band-signaling]
Wijnands, I., Hitchen, P., Leymann, N., Henderickx, W., and a. arkadiy.gulko@thomsonreuters.com, "Multipoint Label Distribution Protocol In-Band Signaling in a VRF Context", draft-ietf-l3vpn-mldp-vrf-in-band-signaling-01 (work in progress), June 2013.

[RFC4605] Fenner, B., He, H., Haberman, B., and H. Sandick, "Internet Group Management Protocol (IGMP) / Multicast Listener Discovery (MLD)-Based Multicast Forwarding ("IGMP/MLD Proxying")", RFC 4605, August 2006.

[RFC6826] Wijnands, IJ., Eckert, T., Leymann, N., and M. Napierala, "Multipoint LDP In-Band Signaling for Point-to-Multipoint and Multipoint-to-Multipoint Label Switched Paths", RFC 6826, January 2013.

14.2. Informative References

[I-D.rekhter-pim-sm-over-mldp]
Rekhter, Y. and R. Aggarwal, "Carrying PIM-SM in ASM mode Trees over P2MP mLDP LSPs",

draft-rekhter-pim-sm-over-mldp-04 (work in progress), August 2011.

[RFC3446] Kim, D., Meyer, D., Kilmer, H., and D. Farinacci, "Anycast Rendevous Point (RP) mechanism using Protocol Independent Multicast (PIM) and Multicast Source Discovery Protocol (MSDP)", RFC 3446, January 2003.

[RFC3618] Fenner, B. and D. Meyer, "Multicast Source Discovery Protocol (MSDP)", RFC 3618, October 2003.

[RFC4601] Fenner, B., Handley, M., Holbrook, H., and I. Kouvelas, "Protocol Independent Multicast - Sparse Mode (PIM-SM): Protocol Specification (Revised)", RFC 4601, August 2006.

[RFC6513] Rosen, E. and R. Aggarwal, "Multicast in MPLS/BGP IP VPNs", RFC 6513, February 2012.

Authors' Addresses

IJsbrand Wijnands (editor)
Cisco
De kleetlaan 6a
Diegem, 1831
Belgium

Phone:
Email: ice@cisco.com

Arkadiy Gulko
Thomson Reuters
195 Broadway
New York NY 10007
USA

Email: arkadiy.gulko@thomsonreuters.com

Uwe Joorde
Deutsche Telekom

Hammer Str. 216-226
Muenster D-48153
DE

Email: Uwe.Joorde@telekom.de

Jeff Tantsura
Ericsson
300 Holger Way
San Jose, California 95134
USA

Email: Jeff.Tantsura@ericsson.com