Network Working Group
Internet-Draft
Intended status: Standards Track
Expires: May 22, 2008

Y. Brehon
M. Vigoureux
L. Ciavaglia
Alcatel-Lucent
M. Chaitou
France Telecom
Z. Ali
Cisco Systems, Inc.
November 19, 2007

IGP Routing Protocol Extensions for Discovery of Upstream Label Assignment Node Capability
draft-brevigcia-ccamp-mpls-upstream-node-cap-01.txt

Status of this Memo

Copyright Notice

Abstract

This document proposes an extension to [TE-NODE-CAP] in order to support additional TE node capabilities in the context of Point-to-MultiPoint (P2MP) LSP routing and fast reroute protection. As for point-to-point LSP, nesting P2MP LSPs, i.e., MPLS P2MP hierarchy, is a desirable traffic-engineering feature. However, nesting P2MP LSPs requires a mechanism to coordinate the label allocation of the inner P2MP LSP between the egress nodes of the P2MP Tunnel as described in [MPLS-UPSTREAM]. To solve this issue, a new label allocation scheme called Upstream Label Assignment (ULA) is defined. Network elements responsible for the route computation of P2MP LSP should be aware of the nodes that support this functionality. For that purpose, we define a new bit flag to the TE Node Capability Descriptor as defined in [TE-NODE-CAP]. The bit flag (U) enables a node to advertise its capability to accept Upstream Label Assignment.

Table of Contents

1. Terminology

This document uses terminologies defined in [RFC3031], [RFC3209] and [RFC4461].

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. Introduction

This document proposes an extension to [TE-NODE-CAP] in order to support additional TE node capabilities in the context of Point-to-MultiPoint (P2MP) LSP routing and Fast ReRoute protection. As for point-to-point LSP, nesting P2MP LSPs, i.e., MPLS P2MP hierarchy, is a desirable traffic-engineering feature. P2MP Fast ReRoute (FRR) [P2MP-FRR] is a typical application of P2MP LSPs nesting that is likely to be deployed in MPLS-TE networks transporting multicast traffic.

However, nesting P2MP LSPs requires a mechanism to coordinate the label allocation of the inner P2MP LSP between the egress nodes of the P2MP Tunnel as described in [MPLS-UPSTREAM]. To solve this issue, a new label allocation scheme called Upstream Label Assignment (ULA) is defined where the ingress node of the P2MP Tunnel allocates a single label to the egress nodes of the P2MP Tunnel for the nested P2MP LSP. Use of this technique raises an additional issue: as the upstream neighbor now assigns the label, two different upstream nodes may allocate the same label value to the same receiver(s) for two different P2MP LSPs nested in different P2MP Tunnels. The egress nodes cannot anymore distinguish the LSPs based on the incoming label value. To overcome this issue, [MPLS-UPSTREAM] defines a Context-specific Label Space (CLS). The egress node must now disambiguate the label of the inner LSP by defining a per-upstream-neighbour label space. As defined in [MPLS-UPSTREAM], downstream LSRs maintain separate label space for each unique root (a P2MP Tunnel head-end) and MUST be able to determine the root of the P2MP Tunnels. The root is identified by the head-end IP address of the Tunnel. If the same upstream node uses different head-end IP addresses for different tunnels then the downstream nodes MUST maintain a different Upstream Neighbor Label Space for each such head-end IP address.

[RSVP-UPSTREAM] defines extensions to [RFC4875] to support the advertisement of the ULA capability between adjacent nodes - i.e., between nodes which have a signaling adjacency. Unfortunately, when nesting P2MP LSPs in P2MP Tunnels, the ingress nodes and the egress nodes usually do not have such a signaling adjacency. Nevertheless, the knowledge of this capability is crucial when calculating the routes of nested P2MP LSPs over P2MP tunnels (either by the ingress node or by a Path Computation Element, PCE). If the ingress of the (nested) P2MP LSP or the PCE does not have a control adjacency with the egress nodes of the P2MP Tunnel, LSP setup will be tried and will fail if at least one egress node does not support the ULA capability. This is a trial-and-error approach, which can reveal inefficient and time and resource consuming.

The idea is thus to extend the MPLS/GMPLS routing protocols (OSPF-TE

and IS-IS-TE) to allow LSRs to inform all nodes within a network domain of a node's capability to receive upstream-assigned labels and process them accordingly. Using the routing protocol guarantees this information will be distributed to all nodes, which should perform route calculations, independently of the signaling protocol used for establishing the LSPs (e.g. RSVP-TE).

For that purpose, this document defines a new bit flag to the TE Node Capability Descriptor as defined in [TE-NODE-CAP]. The bit flag (U) enables a node to advertise its capability to accept Upstream Label Assignment.

3. Advantages of the solution

The advertisement of the ULA capability across the network brings additional Traffic Engineering possibilities to better manage P2MP TE LSPs.

A first advantage of the proposed solution concerns P2MP TE path computation. When transporting multicast traffic over their MPLS networks, service providers and operators will often establish P2MP TE Tunnels and nest the client P2MP LSPs into them in order to keep control of the planning and resource allocation in their networks.

As described briefly above, remote nodes at the endpoints of tunnels do not usually establish signaling adjacency because this would result in a fully connected graph where each node would have a control adjacency with all other nodes in the network. Similarly, if the network operator uses a PCE to calculate P2MP TE paths, the knowledge of the ULA capability cannot be advertised by the signaling protocols. Therefore, in order to avoid these blocking situations and to allow remote nodes to efficiently calculate TE P2MP paths with all the relevant information, disseminating the node capability to accept upstream-assigned labels through IGP routing protocols appears as a desirable feature and seems a scalable and efficient approach.

Moreover, if an operator wishes to setup P2MP tunnels to transport P2MP LSPs, the egress nodes of the P2MP tunnel will have to support ULA. Therefore, it is pointless to setup a P2MP tunnel to only afterwards fail in all nested P2MP LSP establishments; it is much more efficient to have the relevant information for the P2MP tunnel setup right from the start.

A second advantage of the proposed solution concerns P2MP fast reroute protection. As described in [P2MP-FRR], in the P2MP Facility Backup method, a P2MP Bypass Tunnel can be used to protect a set of P2MP TE LSPs. During failure, the same backup label MUST be used for all S2L sub-LSPs of a given backup P2MP LSP, tunneled within the same P2MP Bypass Tunnel to avoid data replication and traffic mis-routing in the Merge Points (MP). The Point of Local Repair (PLR) assigns the same label to all the Merge Points (MP) using the Upstream Label Assignment procedure.

The backup P2MP LSPs and the P2MP Bypass tunnel have to be established prior to the failure and to work properly, the mechanism needs to know the capability of the remote nodes to accept upstream-assigned labels. If some egress nodes do not support ULA, then the PLR will establish dedicated P2P Tunnels towards them.

In P2MP FRR protection, the knowledge of the ULA capability is vital

and particularly important in order to limit the number of trials before the appropriate backup LSP(s) are found and established.

Globally, the proposed solution to transport the ULA capability bit in IGP routing protocols enables:

- o a scalable dissemination of the P2MP node capabilities,

- o a workable fast reroute protection mechanism,

- o a higher reliability/robustness of the signaling phase.

4. New value to the TE Node Capability Descriptor

The TE Node Capability Descriptor contains a variable length set of bit flags, where each bit corresponds to a given TE node capability.

Currently, five TE Node Capabilities are defined in [TE-NODE-CAP]. This document defines a new TE Node Capability:

- U bit: when set, this flag indicates that the LSR accepts Upstream Label Assignment ([RSVP-UPSTREAM]);

The following bit is added to the OSPF TE Node Capability Descriptor TLV:

Bit	Capabilities
5	U bit: If set this indicates that the LSR accepts Upstream Label Assignment ([RSVP-UPSTREAM]).

The following bit is added to IS-IS TE Node Capability Descriptor sub-TLV:

Bit	Capabilities
5	U bit: If set this indicates that the LSR accepts Upstream Label Assignment ([RSVP-UPSTREAM]).

5. Elements of procedure

*** no new element introduced by this draft ***

6. Backward Compatibility

*** no new element introduced by this draft ***

7. Security Considerations

*** no new element introduced by this draft ***

8. IANA considerations

8.1. Capability Registry

[OSPF-CAP] defines a new code point registry for TLVs carried in the Router Information LSA defined in [OSPF-CAP].

IANA is requested to make assignments for the TE node capability defined in this document (see Section 4) using the following suggested values, in the Link State Routing TE Capabilities registry:

Bit No.	Name	Reference
5	U bit: Upstream Label Assignment capability	[This.I-D]

9. References

[TE-NODE-CAP] J.P. Vasseur, J.L. Le Roux et al., "IGP Routing Protocol Extensions for Discovery of Traffic Engineering Node Capabilities", draft-ietf-ccamp-te-node-cap-05, work in progress.

[MPLS-UPSTREAM] R. Aggarwal, Y. Rekhter, E. Rosen, "MPLS Upstream Label Assignment and Context Specific Label Space", draft-ietf-mpls-upstream-label-03, work in progress.

[P2MP-FRR] J. L. Le Roux, R. Aggarwal, J.P. Vasseur, M. Vigoureux, "P2MP MPLS-TE Fast Reroute with P2MP Bypass Tunnels", draft-ietf-mpls-p2mp-te-bypass-01, work in progress.

[RSVP-UPSTREAM] R. Aggarwal, J. L. Le Roux, "MPLS Upstream Label Assignment for RSVP-TE", draft-ietf-mpls-rsvp-upstream-02, work in progress.

[RFC4875] R. Aggarwal, D. Papadimitriou, S. Yasukawa, et. al. "Extensions to RSVP-TE for point-to-multipoint TE LSPs", RFC 4875.

[OSPF-CAP] Lindem, A., Shen, N., Aggarwal, R., Shaffer, S., Vasseur, J.P., "Extensions to OSPF for advertising Optional Router Capabilities", draft-ietf-ospf-cap, work in progress.

Authors' Addresses

Yannick Brehon
Alcatel-Lucent
Route de Villejust
Nozay 91620
France

Email: Yannick.Brehon@alcatel-lucent.fr

Martin Vigoureux
Alcatel-Lucent
Route de Villejust
Nozay 91620
France

Email: Martin.Vigoureux@alcatel-lucent.fr

Laurent Ciavaglia
Alcatel-Lucent
Route de Villejust
Nozay 91620
France

Email: Laurent.ciavaglia@alcatel-lucent.fr

Mohamad Chaitou
France Telecom
2, avenue Pierre-Marzin
Lannion 22307
France

Email: Mohamad.Chaitou@orange-ftgroup.com

Zafar Ali
Cisco Systems, Inc.
100 South Main St. #200
Ann Arbor MI 48104
USA

Email: zali@cisco.com

Full Copyright Statement

Intellectual Property

Acknowledgment

Funding for the RFC Editor function is provided by the IETF Administrative Support Activity (IASA).