

Network Working Group
Internet-Draft
Intended status: Informational
Expires: November 24, 2009

P. Francis
MPI-SWS
X. Xu
Huawei
May 23, 2009

MPLS Tunnels for Virtual Aggregation
draft-ietf-grow-va-mpls-00.txt

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/ietf/lid-abstracts.txt>.

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

This Internet-Draft will expire on November 24, 2009.

Copyright Notice

Copyright (c) 2009 IETF Trust and the persons identified as the document authors. All rights reserved.

This document is subject to [BCP 78](#) and the IETF Trust's Legal Provisions Relating to IETF Documents in effect on the date of publication of this document (<http://trustee.ietf.org/license-info>). Please review these documents carefully, as they describe your rights and restrictions with respect to this document.

Abstract

The document "FIB Suppression with Virtual Aggregation" [[I-D.francis-intra-va](#)] describes how FIB size may be reduced. The

latest revision of that draft refers generically to tunnels, and leaves it to other documents to define the usage and signaling methods for specific tunnel types. This document provides those definitions for MPLS Label Switched Paths (LSP), without tag stacking.

Table of Contents

- [1. Introduction](#) [3](#)
- [1.1. Requirements notation](#) [3](#)
- [1.2. Changes from Previous Versions](#) [3](#)
- [2. Tunneling Requirements](#) [3](#)
- [3. Tunneling Specification for MPLS](#) [4](#)
- [4. IANA Considerations](#) [4](#)
- [5. Security Considerations](#) [4](#)
- [6. Normative References](#) [4](#)
- [Authors' Addresses](#) [5](#)

1. Introduction

This document specifies how to use and signal the tunnels required by [[I-D.francis-intra-va](#)], "FIB Suppression with Virtual Aggregation", for MPLS. This document is limited to MPLS without tag stacking. This document adopts the terminology of [[I-D.francis-intra-va](#)]. This document covers the behavior for both VA routers and legacy routers.

1.1. Requirements notation

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 [[RFC2119](#)].

1.2. Changes from Previous Versions

This document was previously published as [draft-francis-va-tunnels-mpls-00](#). No substantive changes were made from that revision.

2. Tunneling Requirements

According to [[I-D.francis-intra-va](#)], VA has the following tunnel-related requirements. The requirement numbers here (R1 - R5) are cited by [[I-D.francis-intra-va](#)].

- R1: Legacy routers and APRs must be able to detunnel packets addressed to themselves at their BGP NEXT_HOP address. They must be able to signal the tunnel information needed by other routers to initiate these tunneled packets.
- R2: Border VA routers must be able to detunnel packets targeted to neighboring remote ASBRs. They must be able to forward these packets to the targeted remote ASBR without doing a FIB lookup. They must be able to signal the tunnel information needed by other routers to initiate these tunneled packets.

- R3: VA routers must be able to initiate tunneled packets targeted to any BGP NEXT_HOP address (i.e. those for APRs, legacy routers, or remote ASBRs).
- R4: Legacy routers may optionally be able to initiate tunneled packets targeted to any BGP NEXT_HOP address (i.e. those for APRs, legacy routers, or remote ASBRs). The MPLS tunnels defined in this document allow this capability.
- R5: All routers must be able to forward all tunneled packets.

3. Tunneling Specification for MPLS

VA utilizes a straight-forward application of MPLS. The tunnels are MPLS Label Switched Paths (LSP), and are signaled using either the Label Distribution Protocol (LDP) [[RFC5036](#)]. (Note that usage of RSVP-TE [[RFC3209](#)] to signal these tunnels, in particular the scalability of configuring so many tunnels, is for further study.) All routers (VA and legacy alike) must run LDP, as required by R5. A legacy router that cannot run LDP and initiate LSPs terminating at itself cannot participate in a VA domain.

Requirements R1 and R2 require that routers initiate tunnels. This is done by importing the full BGP NEXT_HOP address (/32 if IPv4, /128 if IPv6) into the IGP (i.e. OSPF [[RFC2328](#)]), and initiating Downstream Unsolicited tunnels to all IGP neighbors with the full BGP NEXT_HOP address as the Forwarding Equivalence Class (FEC).

Note that in the case of requirement R2, the BGP NEXT_HOP address is that of the remote ASBR, not that of the router that is initiating the LSP (i.e. the local ASBR VA router). Strictly speaking, this is non-standard behavior---normally it is the router owning the FEC address that initiates signaling. Nevertheless routers can employ existing Penultimate Hop Popping (PHP) mechanisms in the data plane for forwarding packets to remote ASBRs.

Requirements R3 and R4 should naturally be satisfied through normal MPLS usage. In other words, the LSP to the BGP NEXT_HOP address should automatically be the preferred method to routing the packet towards the BGP NEXT_HOP address.

[4.](#) IANA Considerations

There are no IANA considerations.

[5.](#) Security Considerations

Because this document describes a (near) standard application of intra-domain MPLS, there are no new security considerations beyond those already described in [[I-D.francis-intra-va](#)].

[6.](#) Normative References

[[I-D.francis-intra-va](#)]

Francis, P., Xu, X., Ballani, H., Jen, D., Raszuk, R., and L. Zhang, "FIB Suppression with Virtual Aggregation",

Francis & Xu

Expires November 24, 2009

[Page 4]

Internet-Draft

VA MPLS Tunnels

May 2009

[draft-francis-intra-va-01](#) (work in progress), April 2009.

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

[RFC2328] Moy, J., "OSPF Version 2", STD 54, [RFC 2328](#), April 1998.

[RFC3209] Awduche, D., Berger, L., Gan, D., Li, T., Srinivasan, V., and G. Swallow, "RSVP-TE: Extensions to RSVP for LSP Tunnels", [RFC 3209](#), December 2001.

[RFC5036] Andersson, L., Minei, I., and B. Thomas, "LDP Specification", [RFC 5036](#), October 2007.

Authors' Addresses

Paul Francis
Max Planck Institute for Software Systems
Gottlieb-Daimler-Strasse
Kaiserslautern 67633
Germany

Email: francis@mpi-sws.org

Xiaohu Xu

Huawei Technologies

No.3 Xixi Rd., Shang-Di Information Industry Base, Hai-Dian District

Beijing, Beijing 100085

P.R.China

Phone: +86 10 82836073

Email: xuxh@huawei.com