

Network Working Group	P. Francis
Internet-Draft	MPI-SWS
Intended status: Informational	X. Xu
Expires: November 29, 2009	Huawei
	May 28, 2009

[TOC](#)

## 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/1id-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 29, 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](#)] ([Francis, P., Xu, X., Ballani, H., Jen, D., Raszuk, R., and L. Zhang, "FIB Suppression with Virtual Aggregation," April 2009.](#)) 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](#)
    - [1.1. Requirements notation](#)
    - [1.2. Changes from Previous Versions](#)
  - [2. Tunneling Requirements](#)
  - [3. Tunneling Specification for MPLS](#)
  - [4. IANA Considerations](#)
  - [5. Security Considerations](#)
  - [6. Normative References](#)
  - [§ Authors' Addresses](#)
- 

## 1. Introduction

[TOC](#)

This document specifies how to use and signal the tunnels required by [\[I-D.francis-intra-va\] \(Francis, P., Xu, X., Ballani, H., Jen, D., Raszuk, R., and L. Zhang, "FIB Suppression with Virtual Aggregation," April 2009.\)](#), "FIB Suppression with Virtual Aggregation", for MPLS. This document is limited to MPLS without tag stacking. This document adopts the terminology of [\[I-D.francis-intra-va\] \(Francis, P., Xu, X., Ballani, H., Jen, D., Raszuk, R., and L. Zhang, "FIB Suppression with Virtual Aggregation," April 2009.\)](#). This document covers the behavior for both VA routers and legacy routers.

---

### 1.1. Requirements notation

[TOC](#)

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

---

### 1.2. Changes from Previous Versions

[TOC](#)

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

---

## 2. Tunneling Requirements

[TOC](#)

According to [[I-D.francis-intra-va](#)] ([Francis, P., Xu, X., Ballani, H., Jen, D., Raszuk, R., and L. Zhang, "FIB Suppression with Virtual Aggregation," April 2009.](#)), VA has the following tunnel-related requirements. The requirement numbers here (R1 - R5) are cited by [[I-D.francis-intra-va](#)] ([Francis, P., Xu, X., Ballani, H., Jen, D., Raszuk, R., and L. Zhang, "FIB Suppression with Virtual Aggregation," April 2009.](#)).

- 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

[TOC](#)

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](#)] ([Andersson, L., Minei, I., and B. Thomas, "LDP Specification," October 2007.](#)). (Note that usage of RSVP-TE [[RFC3209](#)] ([Awduche, D., Berger, L., Gan, D., Li, T., Srinivasan, V., and G. Swallow, "RSVP-TE: Extensions to RSVP for LSP Tunnels," December 2001.](#)) 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](#)] ([Moy, J., "OSPF Version 2," April 1998.](#))), 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

[TOC](#)

There are no IANA considerations.

---

#### 5. Security Considerations

[TOC](#)

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](#)] ([Francis, P., Xu, X., Ballani, H., Jen, D., Raszuk, R., and L. Zhang, "FIB Suppression with Virtual Aggregation," April 2009.](#)).

---

#### 6. Normative References

[TOC](#)

[I-D.francis-intra-va]	Francis, P., Xu, X., Ballani, H., Jen, D., Raszuk, R., and L. Zhang, " <a href="#">FIB Suppression with Virtual Aggregation</a> ," draft-francis-intra-va-01 (work in progress), April 2009 ( <a href="#">TXT</a> ).
[RFC2119]	<a href="#">Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels," BCP 14, RFC 2119, March 1997</a> ( <a href="#">TXT</a> , <a href="#">HTML</a> , <a href="#">XML</a> ).

[RFC2328]	Moy, J., " <a href="#">OSPF Version 2</a> ," STD 54, RFC 2328, April 1998 ( <a href="#">TXT</a> , <a href="#">XML</a> ).
[RFC3209]	Awduch, D., Berger, L., Gan, D., Li, T., Srinivasan, V., and G. Swallow, " <a href="#">RSVP-TE: Extensions to RSVP for LSP Tunnels</a> ," RFC 3209, December 2001 ( <a href="#">TXT</a> ).
[RFC5036]	Andersson, L., Minei, I., and B. Thomas, " <a href="#">LDP Specification</a> ," RFC 5036, October 2007 ( <a href="#">TXT</a> ).

## Authors' Addresses

[TOC](#)

Paul Francis	
Max Planck Institute for Software Systems	
Gottlieb-Daimler-Strasse	
Kaiserslautern 67633	
Germany	
Email: <a href="mailto:francis@mpi-sws.org">francis@mpi-sws.org</a>	
Xiaohu Xu	
Huawei Technologies	
No.3 Xinxin Rd., Shang-Di Information Industry Base, Hai-Dian District	
Beijing, Beijing 100085	
P.R.China	
Phone: +86 10 82836073	
Email: <a href="mailto:xuxh@huawei.com">xuxh@huawei.com</a>	