

Network Working Group
Internet-Draft
Intended status: Informational
Expires: January 2, 2012

P. Francis
MPI-SWS
X. Xu
Huawei
H. Ballani
Cornell U.
D. Jen
UCLA
R. Raszuk
Cisco
L. Zhang
UCLA
July 1, 2011

Auto-Configuration in Virtual Aggregation
draft-ietf-grow-va-auto-04.txt

Abstract

Virtual Aggregation as specified in [I-D.ietf-grow-va] requires configuration of a static "VP-List" on all routers. The VP-List allows routers to know which prefixes may or may not be FIB-installed. This draft specified an optional method of determining this that requires far less configuration. Specifically, it requires the configuration of a "VP-Range" in ASBRs connected to transit and peer ISPs. A Non-transitive Extended Communities Attribute is used to convey to other routers that a given route can be FIB-suppressed.

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). Note that other groups may also distribute working documents as Internet-Drafts. The list of current Internet-Drafts is at <http://datatracker.ietf.org/drafts/current/>.

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."

This Internet-Draft will expire on January 2, 2012.

Copyright Notice

Copyright (c) 2011 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 (<http://trustee.ietf.org/license-info>) in effect on the date of publication of this document. Please review these documents carefully, as they describe your rights and restrictions with respect to this document. Code Components extracted from this document must include Simplified BSD License text as described in Section 4.e of the Trust Legal Provisions and are provided without warranty as described in the Simplified BSD License.

Table of Contents

- 1. Introduction 3
 - 1.1. Requirements notation 3
- 2. Specification 3
- 3. IANA Considerations 4
- 4. Security Considerations 5
- 5. Acknowledgements 5
- 6. References 5
 - 6.1. Normative References 5
 - 6.2. Informative References 5
- Authors' Addresses 5

1. Introduction

As the current VA specification stands ([I-D.ietf-grow-va]), routers have to know which prefixes they must FIB-install and which they need not FIB-install. The VP-List tells them this: they must FIB-install routes to Virtual Prefixes (VP), and they need not FIB-install routes to prefixes that fall within VPs for which they are not an Aggregation Point Router (APR). The same VP-List must be installed in every router.

This draft specifies an optional alternative to the VP-List that requires far less configuration. Specifically, a list of one or more "VP-Ranges" is configured in ASBRs --- typically ASBRs that do not connect to customer networks. These ASBRs then simply tag routes as to whether the route can be suppressed. This is simpler than the current configured VP-List approach in two regards. First, fewer routers need to be configured. Second, the VP-Range is simpler than the VP-List. In most cases, once an ISP is past its initial VA roll-out phase, the VP-Range consists of a single 0/0 entry.

This draft uses terms defined in [I-D.ietf-grow-va].

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].

2. Specification

With the "VP-Range" approach to determining suppressability, certain ASBRs are designated as "tagging routers". Tagging routers explicitly tag routes with a Non-transitive Extended Communities Attribute that indicates whether the route can be FIB-suppressed. All ASBRs that connect to one or more transit provider ISPs MUST be tagging routers. ASBRs that connect to one or more peer ISPs SHOULD be tagging routers. ASBRs that connect to customer networks SHOULD NOT be tagging routers.

Tagging routers are configured with a "VP-Range" list. This consists of the ranges of IP address that are collectively covered by all VPs in the AS. In a mature deployment of VA, the range would amount to all IP addresses, in which case the VP-Range is simply 0/0. Early in VA deployment, when an ISP is still in the testing or roll-out phase, the VP-Range may consist of multiple entries.

Tagging routers SHOULD tag any route whose prefix falls within the

VP-Range with a "can-suppress" tag, with the following exceptions:

1. Tagging routers MUST NOT tag VP routes with can-suppress (where a VP route is that route to the VP that the router originates in its role as an APR).
2. If the ISP has a policy of FIB-installing customer routes, then routes received from customers SHOULD NOT be tagged with can-suppress.

The can-suppress tag itself is an Extended Communities Attribute [RFC4360] to be assigned by IANA from the "well-known" pool define in [I-D.ietf-idr-reserved-extended-communities]. The Transitive Bit MUST be set to value 1 (the community is non-transitive across ASes).

Routers install or suppress FIB entries according to the following rules. Note that tagging routers conceptually follow these rules after tagging (or not tagging) the route. Note also that these rules apply only to the route used by the router as the best route. In other words, if a router receives two routes for the same prefix, and one route is tagged can-suppress and the other is not, the router follows these rules only with respect to the route that it selects as the best route.

1. Routes without the can-suppress tag MUST be FIB-installed.
2. APRs MUST FIB-install routes for sub-prefixes that fall within the APR's VPs, whether or not the route is tagged can-suppress.
3. Otherwise, routers MAY FIB-suppress routes tagged as can-suppress.

3. IANA Considerations

IANA is requested to assign, from the registry "BGP Assigned non-transitive extended communities", a value TBD for "VA can suppress":

Registry Name: BGP Assigned non-transitive extended communities

Name	Type Value
----	-----
VA can suppress	TBD

4. Security Considerations

As of this writing, there are no known new security threats introduced by this draft.

5. Acknowledgements

The authors would like to thank Wes George and Bruno Decraene for their reviews and suggestions.

6. References

6.1. Normative References

[I-D.ietf-grow-va]

Francis, P., Xu, X., Ballani, H., Jen, D., Raszuk, R., and L. Zhang, "FIB Suppression with Virtual Aggregation", draft-ietf-grow-va-04 (work in progress), Oct 2009.

[I-D.ietf-idr-reserved-extended-communities]

Decraene, B. and P. Francois, "Assigned BGP extended communities", draft-ietf-idr-reserved-extended-communities-01 (work in progress), May 2011.

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

[RFC4360] Sangli, S., Tappan, D., and Y. Rekhter, "BGP Extended Communities Attribute", RFC 4360, February 2006.

6.2. Informative References

Authors' Addresses

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

Phone: +49 631 930 39600
Email: francis@mpi-sws.org

Xiaohu Xu
Huawei Technologies
No.3 Xinxu Rd., Shang-Di Information Industry Base, Hai-Dian District
Beijing, Beijing 100085
P.R.China

Phone: +86 10 82836073
Email: xuxh@huawei.com

Hitesh Ballani
Cornell University
4130 Upson Hall
Ithaca, NY 14853
US

Phone: +1 607 279 6780
Email: hitesh@cs.cornell.edu

Dan Jen
UCLA
4805 Boelter Hall
Los Angeles, CA 90095
US

Phone:
Email: jenster@cs.ucla.edu

Robert Raszuk
Cisco Systems, Inc.
170 West Tasman Drive
San Jose, CA 95134
USA

Phone:
Email: raszuk@cisco.com

Lixia Zhang
UCLA
3713 Boelter Hall
Los Angeles, CA 90095
US

Phone:
Email: lixia@cs.ucla.edu

