

Network Working Group  
Internet-Draft  
Intended status: Informational  
Expires: March 3, 2012

L. Cheng  
M. Sun  
ZTE Corporation  
August 31, 2011

Auto-Configuration Extension in Virtual Aggregation  
draft-cheng-grow-va-auto-extensions-00

Abstract

Auto-Configuration in Virtual Aggregation as specified in [I-D.ietf-grow-va-auto] requires configuration of a "VP-range list" in ASBRs connected to transit and peer ISPs. These ASBRs simply tag some routes whose prefix falls within the VP-Range with a "can-suppress" tag to indicate whether these routes should be FIB installed. This draft specified an extended auto-configuration mechanism in Virtual Aggregation to support the configuration of both "VP-List" and "popular prefixes". Specifically, based on this mechanism, the ratio of lost packets when VP routes fail could be minimized.

Status of this Memo

This Internet-Draft is submitted 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 March 3, 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 Language . . . . .	4
2. Terminology . . . . .	4
3. Specification . . . . .	4
3.1. Routes Classification and Routes Tagging . . . . .	4
3.2. Routes Installation . . . . .	5
3.3. Implementation . . . . .	5
4. Operation under Special Scenario . . . . .	7
4.1. Non-tagging Routers Operation . . . . .	7
4.2. Tagging Routers Operation . . . . .	8
4.3. Implementation . . . . .	8
5. IANA Considerations . . . . .	9
6. Security Considerations . . . . .	9
7. References . . . . .	9
7.1. Normative References . . . . .	9
7.2. Informative References . . . . .	9
Authors' Addresses . . . . .	10

## 1. Introduction

Virtual Aggregation specified in [I-D.ietf-grow-va] requires configuration of a static "VP-List" on all routers. "VP-List" allows routers to know which prefixes may or may not be FIB installed. Auto-configuration mechanism [I-D.ietf-grow-va-auto] provides an optional method for routers to do routes decision with less configuration.

Auto-configuration is an optional alternative to the VP-list that requires far less configuration. However, further concentrates should be focused on some scenarios where packets transmission maybe seriously influenced based on this mechanism. Furthermore, this mechanism could also be extended to provide more excellent service.

This draft specifies Auto-Configuration Extension Operation, which includes the following two aspects:

- o VP routes to be specified particularly. Based on current auto-configuration, tagging routers must not tag VP routes with can-suppress tag. If the ISP has a policy of FIB-installing customer routes, then routes received from customers should also not be tagged. Consequently, there may be three kinds of routes are non-tagged in the AS: routes whose prefix out scope of VP-Range, VP routes and customer routes. According to these tagging rules, non-tagging routers will not be able to identify VP routes. As a result, in the case where all VP routes for a given VP are withdrawn, non-tagging routers would not be able to FIB-install sub-prefixes within the VP. This will influence the normal transmission of data packets seriously.
- o Extensions to realize popular prefixes auto-configuration. As specified in [I-D.ietf-grow-va], deployment of Visual Aggregation will cause path stretch. To minimize the latency and load associated with the longer path, ISP could measure traffic volume over time and install the high volume prefixes. These prefixes which are within a VP, but still be FIB installed are called popular prefixes. Furthermore, popular prefixes could also be consisted of policy-based prefixes and static list prefixes. Customer prefixes could be considered as one kind of policy-based prefixes. Consequently, tagging routers could also be configured with a popular prefixes list, and realize popular prefixes auto-configuration by not tag routes whose prefix falls within this list.

### 1.1. Requirements Language

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. Terminology

This draft uses terms defined in [I-D.ietf-grow-va]. This section defines some new terms used in this document.

**Tagging router:** ASBRs which are configured with "VP range" and "Popular-Prefix list". These routers tag routes with different tags based on route type. Typically, all ASBRs that connect to one or more transit provider ISPs must be configured as tagging routers. ASBRs that connect to one or more peer ISPs should be configured as tagging routers. ASBRs that connect to customer networks should not be configured as tagging routers.

**Non-tagging router:** The VA routers in AS which are not tagging routers.

**Popular-Prefix list:** List of popular prefixes.

**Suppress tag:** Tags used by tagging routers to tag routes. Routes with this tag may not be FIB installed by routers. This tag could be attached to a route as a Non-transitive Extended Communities Attribute.

**Install tag:** Tags used by tagging routers to tag routes. Router with this tag must be FIB installed by any router. This tag could be attached to a route as a Non-transitive Extended Communities Attribute.

## 3. Specification

### 3.1. Routes Classification and Routes Tagging

With this extended auto-configuration approach, every tagging router will be configured with the same "VP-range list" and "popular prefix list".

"VP-range list" consists of the ranges of IP address that are collectively covered by all VPs in the AS [I-D.ietf-grow-va-auto]. "Popular-Prefix list" is a list of popular prefixes. These popular prefixes are all regular prefixes, and could be selected by ISPs

individually based on their requirements.

With the extended auto-configuration approach, ASBRs which are tagging routers first classify all routes into three types based on the "VP range" and "Popular-Prefix list" configured in them:

Type1: VP routes which MUST be FIB installed by any router;  
Type2: Routes whose prefix falls within "Popular-Prefix list", and routes whose prefix is not fall within "VP range";  
Type3: Routes whose prefix falls within "VP range", and meantime are out scope of Type 1 and Type 2 routes.

Tagging routers tag routes explicitly according to route types.

1. All VP routes (type 1) MUST be tagged with a "install tag".
2. All routes falls within Type 2 SHOULD NOT be tagged.
3. All routes falls within Type 3 MAYBE tagged with a "suppress tag".

### 3.2. Routes Installation

Routers install or suppress FIB entries according to the following rules.

1. Routes with "install tag" MUST be FIB-installed.
2. Routes without any tag SHOULD be FIB-installed.
3. Routes with "suppress tag" MAY be FIB-suppressed.
4. APRs MUST FIB-install routes for sub-prefixes that fall within the APRs!\_ VPs, whether or not the route is tagged.

Note: tagging routers conceptually follow these rules after tagging (or not tagging) the route.

Note: these rules apply only to the route used by the routers as the best route.

### 3.3. Implementation

An instance of mechanism operation is depicted in this subsection. Consider the scenario depicted in Figure. 1.

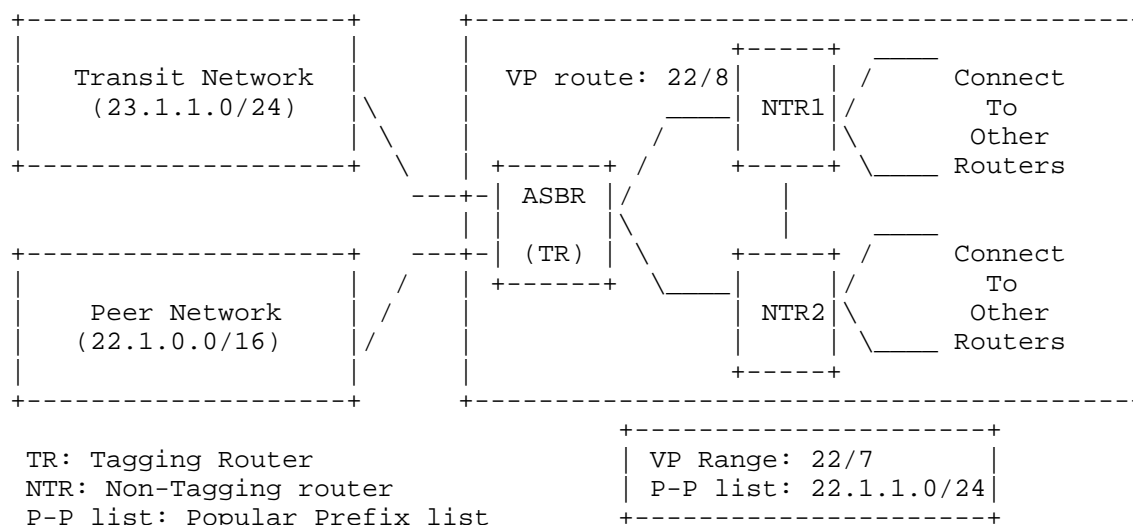


Figure. 1

In this situation, an ASBR connected to transit network (23.1.1.0/24) and peer network (22.1.0.0/16) is selected to be a TR (tagging router). TR is configured with a VP range 22/7 and a popular prefix list contains 23.1.1.0/24. The NTR1 (Non-Tagging Router 1) is an APR (Aggregate Point Router) who announces a VP route 22/8.

To describe operation of different elements, assume that following routes will be received by TR, NTR1 and NTR2.

Routes	Prefix
1	22/8
2	22.1.1.1/32
3	22.1.0.1/32
4	23.1.1.1/32

#### TR Operation:

- o For route with prefix 22/8, as this route is a VP route announced by NTR1, TR will tag it with a "install tag".
- o The prefix 22.1.1.1/32 falls within the popular prefix list, TR will not tag it, although the route's prefix falls within the VP range.

- o TR perceives that prefix 22.1.0.1/32 falls within the VP range and is not a popular prefix. This route will be tagged with a "suppress tag".
- o For route with prefix 23.1.1.1/32, as the prefix falls within VP range 22/7, this route will also be tagged with a "suppress tag".

#### NTR Operation:

- o For route with prefix 22/8, as this route is tagged with "install tag", all NTRs must FIB install it.
- o For route with prefix 22.1.1.1/32, as this route is not tagged, NTRs should FIB install it. Especially for NTR1, as the prefix falls within the VP (22/8) it announced, it must FIB install the route.

It should be noticed that in this scenario, NTR1 is an APR, and NTR2 is a non-APR. These two kinds of routers will implement different operation upon some suppress tagged routes.

- o According to NTR2, as the router is a non-APR, all routes with "suppress tag" should not be installed. As a result, NTR2 will not FIB install 22.1.0.1/32 and 23.1.1.1/32.
- o Now consider the operation of NTR1.
  - \* For the route 22.1.0.1/32 with "suppress tag", NTR1 perceives that prefix falls within 22/8, and will FIB install the route.
  - \* According to route 23.1.1.1/32 with "suppress tag", NTR1 doesn't have to FIB install it, as NTR1 is not the APR for this route.

## 4. Operation under Special Scenario

Based on analysis in section 1, when VP routes are not tagged specially, VP routes failing will influence the packets transmission seriously.

From perspective of non-tagging routers, VP routes could be identified through the "install tag" based on extended auto-configuration mechanism. This section assumes a special scenario that all VP routes for a given VP are withdrawn. Proper operation of tagging routers and non-tagging routers is described as following.

### 4.1. Non-tagging Routers Operation

When the non-tagging routers find that all VP routes for a given VP withdrawn, they will immediately look up the routing table in the RIB, select and FIB install the suppress tagged routes whose prefixes

fall within the withdrawn VP.

#### 4.2. Tagging Routers Operation

When the tagging routers find that all VP routes for a given VP are withdrawn, they will implement the following operation:

- o Look up the routing table in the RIB, select and FIB install the suppress tagged routes whose prefixes fall within the withdrawn VP;
- o Record VP information of the withdrawn VP routes;
- o When there is an invalid record for a given VP, all routes received by the tagging routers whose prefixes falls within this VP should not be tagged.

According to existence of invalid VP records, once receiving a VP route, tagging routers will compare the VP prefix received with the VP prefixes recorded. If the received prefix match an invalid prefix record, they will implement the following operation:

- o Delete the invalid VP prefix record;
- o Tag received VP route with "install tag";
- o Tag all Type 3 routes whose prefix falls within this VP with "suppress tag".

#### 4.3. Implementation

Consider the implementation usecase depicted in Figure. 1. Assume that all VP routes for the VP 22/8 are withdrawn.

According to this problem, NTRs (include APRs) check their routing table in RIB, and find out suppress tagged routes whose prefixes fall within the VP 22/8, such as routes with 22.1.0.1/32. NTRs will FIB install these routes, and forward packets based on them.

According to this problem, TRs will also FIB install the suppressed tagged routes whose prefixes fall within the VP 22/8. Furthermore, TRs will record the prefix information of VP 22/8. During the period that VP 22/8 is withdrawn, all routes whose prefix falls within the VP will not be tagged by TRs, indicates that routes such like 22.1.1.1/32 and 22.1.0.1/32 should be FIB installed.

Every TR maintains a record of the withdrawn VP 22/8. When the TR receives a new VP route with prefix 22/8, it will consider this situation as "VP Recovery". Withdrawn record of 22/8 will be deleted, and the new VP route will be tagged with "install tag". Furthermore, the Type 3 routes whose prefixes fall within the 22/8 such as 22.1.0.1/32 will be tagged with "suppress tag".



## 5. IANA Considerations

IANA is requested to assign, from the registry "BGP Assigned non-transitive extended communities", values TBD for "must install" and "can suppress".

Registry Name: BGP Assigned non-transitive extended communities

Name	Type Value
-----	-----
must install	TBD
can suppress	TBD

## 6. Security Considerations

TBD

## 7. References

### 7.1. Normative References

- [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.

### 7.2. Informative References

- [I-D.ietf-grow-val]  
Francis, P., Xu, X., Ballani, H., Jen, D., Raszuk, R., and L. Zhang, "FIB Suppression with Virtual Aggregation", draft-ietf-grow-val-05 (work in progress), June 2011.
- [I-D.ietf-grow-val-auto]  
Francis, P., Xu, X., Ballani, H., Jen, D., Raszuk, R., and L. Zhang, "Auto-Configuration in Virtual Aggregation", draft-ietf-grow-val-auto-04 (work in progress), June 2011.
- [I-D.ietf-idr-reserved-extended-communities]  
Decraene, B. and B. Decraene, "Assigned BGP extended communities", draft-ietf-idr-reserved-extended-communities-01 (work in progress), May 2011.

Authors' Addresses

Li Cheng  
ZTE Corporation  
Zijinghua Road No.68  
NanJing, Yuhuatai District 210012  
P.R.China

Email: cheng.li2@zte.com.cn

Mo Sun  
ZTE Corporation  
Zijinghua Road No.68  
NanJing, Yuhuatai District 210012  
P.R.China

Phone: +86-025-52871474

Email: sun.mo@zte.com.cn