

Network Working Group
Internet-Draft
Intended status: Standards Track
Expires: May 19, 2018

A. Azimov
E. Bogomazov
Qrator Labs
R. Bush
Internet Initiative Japan
K. Patel
Arrcus, Inc.
K. Sriram
US NIST
November 15, 2017

New definition of ISP internal eBGP border using BGP Roles
draft-ymbk-idr-isp-border-02

Abstract

This document proposes a new definition of ISP borders using BGP Roles. It may be used to improve the BGP best path selection algorithm for better support of hot-potato routing between different internal ASNs of an ISP. It may also be used to enable transmission of local attributes between different internal ASNs of an ISP.

Requirements Language

The key words "MUST", "MUST NOT", "REQUIRED", "SHALL", "SHALL NOT", "SHOULD", "SHOULD NOT", "RECOMMENDED", "MAY", and "OPTIONAL" are to be interpreted as described in [RFC 2119](#) [[RFC2119](#)] only when they appear in all upper case. They may also appear in lower or mixed case as English words, without normative meaning.

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 May 19, 2018.

Internet-Draft

Network border definition

November 2017

Copyright Notice

Copyright (c) 2017 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	2
2.	Changes in BGP decision process	3
3.	Local Attributes Transmission	3
4.	IANA Considerations	3
5.	Normative References	3
	Authors' Addresses	4

[1.](#) Introduction

The BGP best path selection algorithm ([Section 9.1.2.2 of \[RFC4271\]](#)) has a very clear definition of a network border: different ASNs – different networks. It differs from some real world situations when two networks become one business entity and want to operate as one network.

Today BGP does not provide any robust or automated support for such merging networks:

- o There is no support for carrying local attributes through this border,
- o Hot-potato routing, implemented by eBGP being preferred to iBGP, does not work, and
- o Route Leak prevention inside such a united network can not be easily automated.

In [[I-D.ietf-idr-bgp-open-policy](#)] BGP Roles were introduced - a configuration option that strongly enforces agreement on real-world peering relations between two BGP speakers. This configuration option can accept values of: Peering, Customer, Provider and

Internal. These values could be used in a new ISP border definition: Internal vs. External. With this definition of network borders, it becomes easy to allow robust propagation of local attributes between different ASNs of one ISP. It could be also used to improve the hot-potato routing mechanism: where routes learned from External BGP connections should be preferred over Internal, even those which cross the ISP's internal AS/AS boundary.

[2.](#) Changes in BGP decision process

To improve hot-potato routing for networks with multiple ASNs we propose to insert before d) in [Section 9.1.2.2 of \[RFC4271\]](#) next step:

If at least one of the candidate routes was received via a BGP session with External (Peer, Provider, Customer) role, remove from consideration all routes that were received via BGP sessions with an Internal role.

While this step will improve traffic control for ISPs with multiple ASNs it will have no affect on ISPs with single ASN.

[3.](#) Local Attributes Transmission

Propagation of local attributes through an ISP's internal AS/AS border could be enabled only if both sides set Internal roles in their BGP Open negotiation. Different attributes may still have different transmission policy:

- o iOTC attribute from [[I-D.ietf-idr-bgp-open-policy](#)] MUST be sent to enforce route leak prevention,
- o LOCAL_PREF attribute MAY be sent, and
- o MED attribute MAY be sent without changes.

[4.](#) IANA Considerations

This document has no IANA actions.

[5.](#) Normative References

[I-D.ietf-idr-bgp-open-policy]

Azimov, A., Bogomazov, E., Bush, R., Patel, K., and K. Sriram, "Route Leak Prevention using Roles in Update and Open messages", [draft-ietf-idr-bgp-open-policy-01](#) (work in progress), July 2017.

Azimov, et al.

Expires May 19, 2018

[Page 3]

Internet-Draft

Network border definition

November 2017

[RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", [BCP 14](#), [RFC 2119](#), DOI 10.17487/RFC2119, March 1997, <<https://www.rfc-editor.org/info/rfc2119>>.

[RFC4271] Rekhter, Y., Ed., Li, T., Ed., and S. Hares, Ed., "A Border Gateway Protocol 4 (BGP-4)", [RFC 4271](#), DOI 10.17487/RFC4271, January 2006, <<https://www.rfc-editor.org/info/rfc4271>>.

Authors' Addresses

Alexander Azimov
Qrator Labs

Email: aa@qrator.net

Eugene Bogomazov
Qrator Labs

Email: eb@qrator.net

Randy Bush
Internet Initiative Japan

Email: randy@psg.com

Keyur Patel
Arrcus, Inc.

Email: keyur@arrcus.com

Kotikalapudi Sriram
US NIST

Email: ksriram@nist.gov