

Network Working Group
Internet-Draft
Intended status: Informational
Expires: January 13, 2021

D. Li
J. Wu
Tsinghua
Y. Gu
Huawei
L. Qin
Tsinghua
T. Lin
H3C
July 12, 2020

Source Address Validation Architecture (SAVA): Intra-domain Use Cases
draft-li-sava-intra-domain-use-cases-00

Abstract

This document identifies scenarios where existing approaches for detection and mitigation of source address spoofing don't perform perfectly. Either Ingress ACL filtering [[RFC3704](#)], unicast Reverse Path Forwarding (uRPF) [[RFC3704](#)], feasible path uRPF [[RFC 3704](#)], or Enhanced Feasible-Path uRPF [[RFC8704](#)] has limitations regarding either automated implementation objective or detection accuracy objective (0% false positive and 0% false negative). This document identifies two such scenarios and provides solution discussions.

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 [RFC 2119](#) [[RFC2119](#)].

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 <https://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 13, 2021.

Internet-Draft

SAVA Intra-domain Use Cases

July 2020

Copyright Notice

Copyright (c) 2020 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 (<https://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
1.1.	Source Address Validation	2
1.2.	Existing SAV Techniques Overview	3
1.3.	SAV Requirements and Challenges	4
2.	Terminology	6
3.	Problem Statement	7
3.1.	SAVA Intra-domain Use Case 1: Intra-AS Multi-homing . . .	7
3.2.	SAVA Intra-domain Use Case 2: Inter-AS Multihoming . . .	8
4.	Solution Consideration	10
5.	Security Considerations	11
6.	Contributors	11
7.	Acknowledgments	11
8.	Normative References	11
	Authors' Addresses	14

[1.](#) Introduction[1.1.](#) Source Address Validation

The Internet is open to traffic, which means that a sender can generate traffic and send to any receiver in the Internet without permission of the receiver. Although this openness design improves the scalability of the Internet, it also leaves security risks, namely, a sender can forge his/her source address when sending the packets, which is also well known as source address spoofing.

Due to the lack of source address spoofing detection mechanism, Denial of Service (DoS) attacks seriously compromise network security. By employing source address spoofing, attackers can well hide themselves and pin the blame on the destination networks. Administrators often spend a lot of effort identifying attack packets

without being able to locate the attacker's true source address. In addition to DOS attacks, source address spoofing is also used in a multitude of ways. The threats of source address spoofing have been well documented in [[RFC6959](#)]. To briefly summarize, the possible attacks by source address spoofing includes single-packet attack, flood-based DoS, poisoning attack, spoof-based worm/malware propagation, reflective attack, accounting subversion, man-in-the-middle attack, third-party recon, etc.

[1.2.](#) Existing SAV Techniques Overview

Source address validation (SAV) verifies the authenticity of the packet's source address to detect and mitigate source address spoofing [[RFC2827](#)]. Source Address Validation Improvement (SAVI) method [[RFC7039](#)] implements SAV at a fine granularity of host-level IP address validation. The unicast Reverse Path Forwarding (uRPF) techniques (such as Strict uRPF, Feasible uRPF and Loose uRPF) [[RFC3704](#)] are particularly designed to perform SAV in the granularity of IP network. The Enhanced Feasible-Path Unicast Reverse Path Forwarding (EFP-uRPF) methods [[RFC8704](#)] further improve Feasible uRPF to reduce false positives in the case of inter-AS routing asymmetry and multihoming.

SAVI, typically performed at the access network, is enforced in switches, where the mapping relationship between an IP address and other "trust anchor" is maintained. A "trust anchor" can be link-layer information (such as MAC address), physical port of a switch to connect a host, etc. It enforces hosts to use legitimate IP source addresses. However, given numerous access networks managed by different operators, it is far away from practice for all the access networks to simultaneously deploy SAVI. Therefore, in order to mitigate the security risks raised by source address spoofing, SAV performed in network border routers is also necessary. Although it does not provide the same filtering granularity as SAVI does, it still helps the tracing of spoofing to a minimized network range.

Ingress ACLs [[RFC2827](#)], typically performed at the network border routers, is performed by manually maintaining a traffic filtering access list which contains acceptable source address for each interface. Only packets with a source address encompassed in the access list can be accepted. It strictly specifies the source address space of incoming packets. However, manual configuration brings scalability and reliability issues.

Strict uRPF, typically performed at the network (IGP areas or ASes) boarder routers, requires that a data packet can be only accepted when the FIB contains a prefix that encompasses the source address and the corresponding out-interface matches the data incoming

interface. It has the advantages of simple operation, easy deployment, and automatic update. However, in the case of multihoming, when the data incoming interface is different from the out-interface of the packet source IP address, using the longest prefix match, also referred to as asymmetric routing of data packets, Strict uRPF exhibits false positive.

Loose uRPF, sacrificing the directionality of Strict uRPF, only requires that the packet's source IP exists as a FIB entry. Intuitively, Loose uRPF cannot prevent the attacker from forging a source address that already exists in the FIB.

Feasible uRPF, typically performed at the network border routers, helps mitigate false positive of Strict uRPF in the multihoming scenarios. Instead of installing only the best route into FIB as Strict uRPF does, Feasible uRPF installs all alternative paths into the FIB. It helps reduce false positive filtering compared with the Strict uRPF, in the case when multiple paths are learnt from different interfaces. However, it should be noted that Feasible uRPF only works when multiple paths is learnt. There are cases when a device only learns one path but still has packets coming from other valid interfaces.

EFP-uRPF, specifically performed at the AS border routers, further improves Feasible uRPF in the inter-AS scenario. An AS performing EFP-uRPF maintains an individual RPF list on every customer/peer interface. It introduces two algorithms (i.e., Algorithm A and Algorithm B) regarding different application scenarios. In the case that a customer interface fails to learn any route from the directly

connected customer AS, enabling Algorithm A at this customer interface may exhibit false positive filtering. In this case, enabling Algorithm B may mitigate the false positive. However, in case of directly connected customer ASes spoofing each other, Algorithm B exhibits false negative.

[1.3.](#) SAV Requirements and Challenges

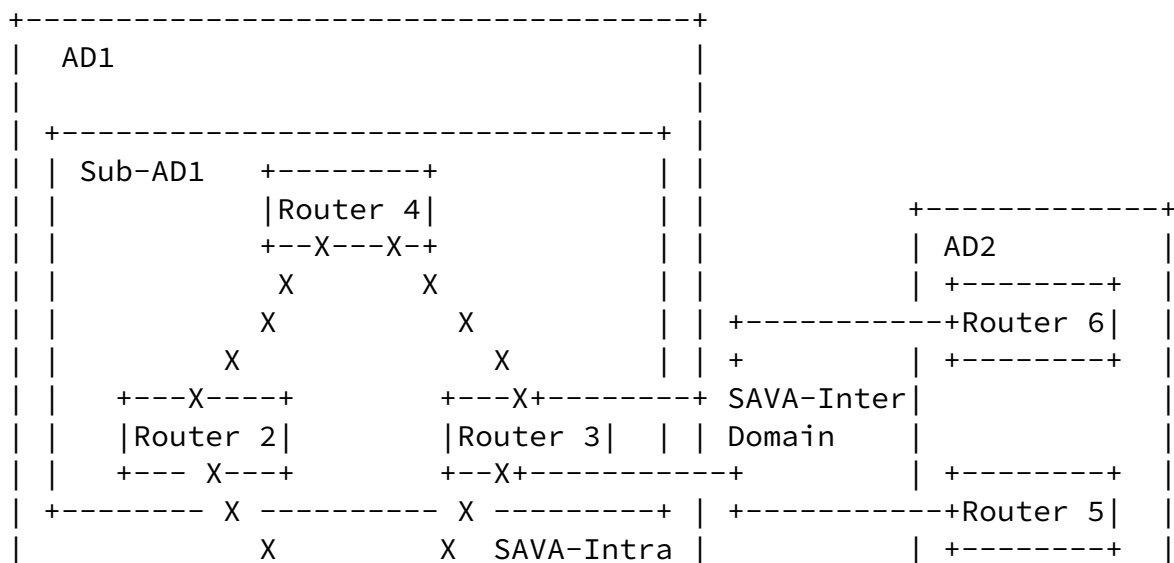
As the above overview indicates, to evaluate the quality of a specific SAV technique, one should balance between two general requirements: precise filtering and automatic implementation.

- o Precise filtering: Two important indicators for precise filtering. 1) 0% false positive. If legitimate packets may be dropped, it can seriously affect the user's internet experience. 2) 0% false negative. If some packets with a forged source address may pass through the SAV smoothly, it will pose potential security risks.

- o Automatic implementation: In reality, the address space of an administrative domain (AD) may grow or update, and the routing policy within the address domain may be dynamically adjusted. One solution that relies entirely on manual configuration is neither scalable nor easy to deploy.

Then to consider the whole network SAV solution, one should never rely on a single point but a systematic SAV technique combination deployed at different network levels. As shown in Figure 1, packet filtering at different levels from the access network to the AS boarder are all needed. In Figure 1, the administrative domain (AD) concept is used, which refers to a network domain managed by the same operator (OTT, ISP and so on). One AD is allowed to be divided into several sub-ADs and managed by different inner groups. There may exist different levels for sub-ADs. For example, sub-AD1 is the upper level compared to sub-AD2, meaning that sub-AD2 needs to connect through sub-AD1 for external reachability (i.e., networks outside AD1). So filtering at sub-AD boarders (between different levels and within the same level) is also necessary. Further, different sub-ADs can belong to one single AS or multiple ASes, which makes the filtering at the sub-AD boarders either intra-AS filtering or inter-AS filtering. In the rest of this document, we use the

July 2020



[3.](#) Problem Statement

As stated in [Section 1.3](#), existing methods, e.g., Loose/Strict mode uRPF, FP-uRPF, EFP-uRPF are not able to achieve 100% accurate filtering (i.e., 0% FN and 0% FP) in certain scenarios. This document specifically indicate two typical intra-domain cases that conventional approaches fail to cover: 1) all sub-ADs are under the same AS; 2) sub-ADs are under different ASes.

[3.1.](#) SAVA Intra-domain Use Case 1: Intra-AS Multi-homing

Figure 2 illustrates an intra-AS multihoming case, where sub-AD1, sub-AD2 and sub-AD3 are under the same AS.

Router 1 is multi-homed to Router 2 and Router 3. Router 1 doesn't announce any of its routes to Router 2 nor Router 3. Static routes are configured on Router 1, Router 2 and Router 3. Supposedly, both Router 2 and Router 3 should have static routes P1/P2 with Router 1 as next hop configured. However, due to configuration error, or traffic control purpose, on Router 3, no P1/P2 static routes are configured. Router 2 and Router 3 are connected with ISIS or OSPF. P1/P2 are flooded from Router 2 to Router 3.

Router 5 is single-homed to Router 3. Router 5 announces P3 to Router 3 using ISIS or OSPF. Router 3 floods P3 to Router 2 .

Now suppose two data flow coming from Router 1 to Router 3: Flow 1 with source IP as P1, and Flow 2 with source IP as P3 (IP spoofing). Using existing SAV methods at Router 3, Flow 1 is supposed to be passed, while Flow 2 is supposed to be dropped.

- o Loose uRPF: works for Flow 1, but fails for Flow 2.
- o Strict uRPF: works for Flow 2, but fails for Flow 1 (the incoming interface does not match P1/P2's out-interface).
- o FP-uRPF: works for Flow 2, but fails for Flow 1 (no feasible path for P1/P2 other than the best route exists).

- o EFP-uRPF: does not apply at the intra-AS case.

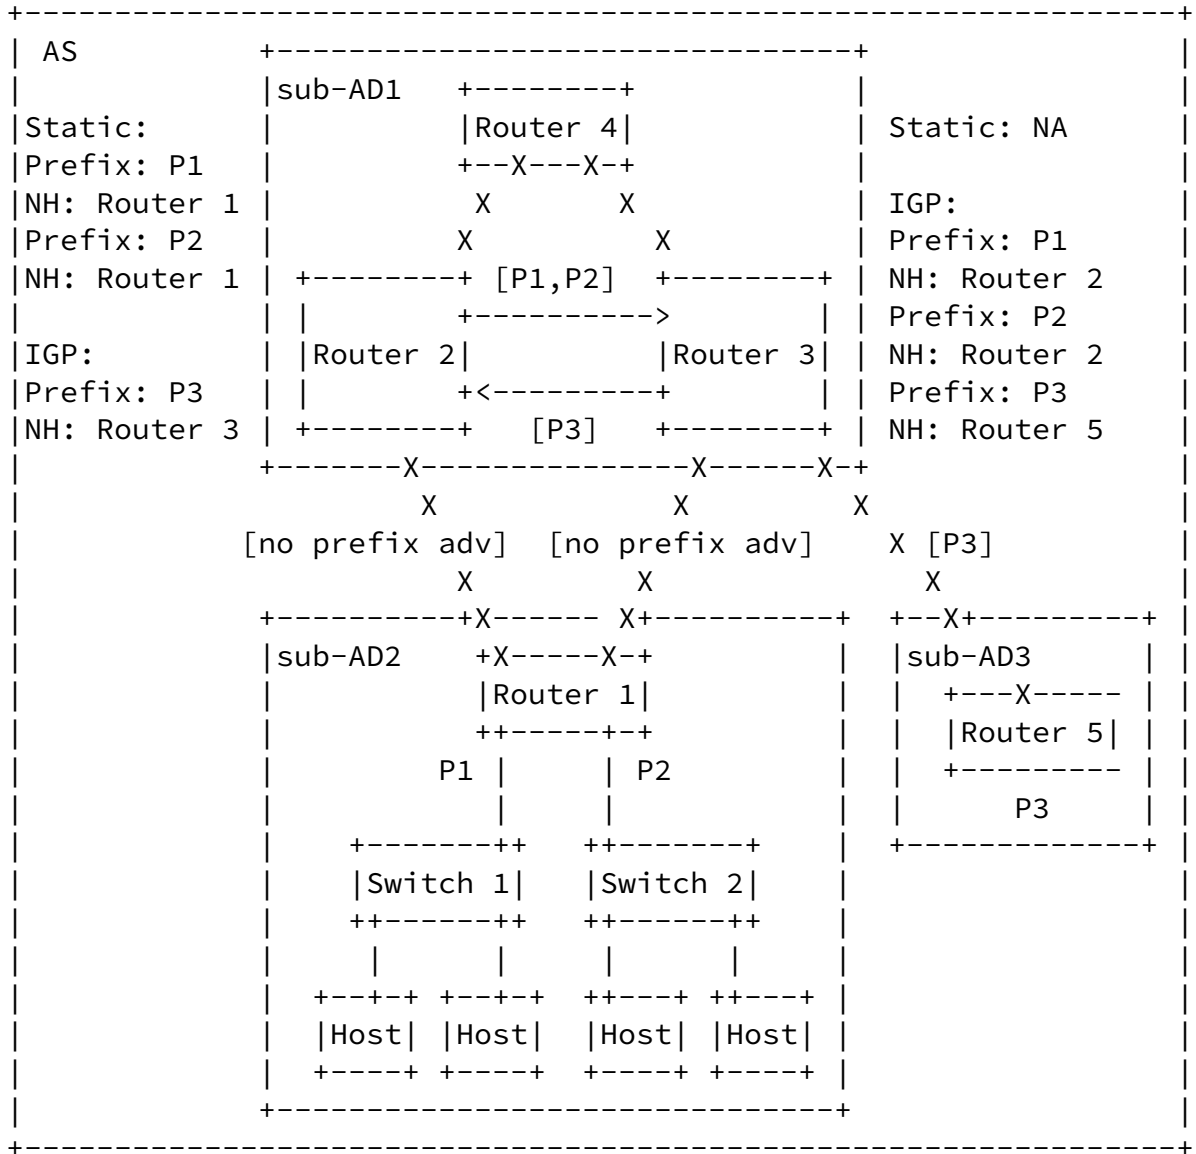


Figure 2: Asymmetric data flow in the Intra-AS scenario

3.2. SAVA Intra-domain Use Case 2: Inter-AS Multihoming

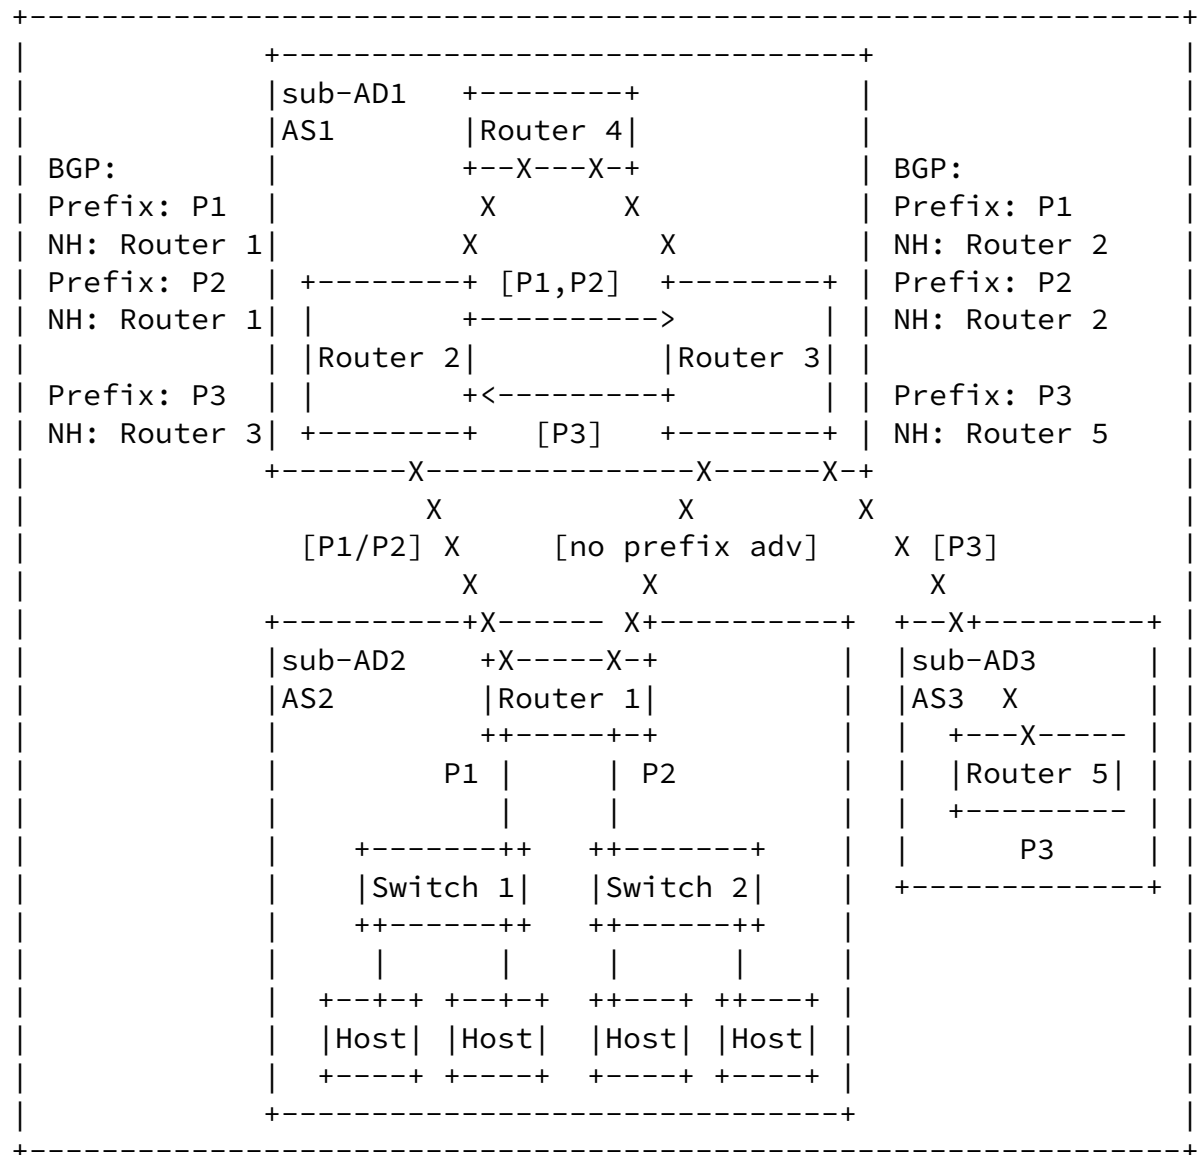
Figure 3 illustrates an inter-AS multihoming case, where sub-AD1, sub-AD2 and sub-AD3 are under three different ASes.

Router 1 (AS2) is multi-homed to Router 2 (AS1) and Router 3 (AS1). Router 1 announces P1/P2 to Router 2 through BGP. Router 1 doesn't announce any of its routes to Router 3 due to policy control. P1/P2 are propagated from Router 2 to Router 3 through BGP.

Router 5 (AS3) is single-homed to Router 3 (AS1). Router 5 announces P3 to Router 3 through BGP. Router 3 propagates P3 to Router 2 through BGP.

Now suppose two data flow coming from Router 1 to Router 3: Flow 1 with source IP as P1, and Flow 2 with source IP as P3 (IP spoofing). Using existing SAV methods at Router 3, Flow 1 is supposed to be passed, while Flow 2 is supposed to be dropped.

- o Loose uRPF: works for Flow 1, but fails for Flow 2.
- o Strict uRPF: works for Flow 2, but fails for Flow 1 (the incoming interface does not match P1/P2's out-interface).
- o FP-uRPF: works for Flow 2, but fails for Flow 1 (no feasible path for P1/P2 other than the best route exists).
- o EFP-uRPF: works for Flow 1, but fails for Flow 2 using Algorithm B. Works for Flow 2, but fails for Flow 1 when using Algorithm A.



4. Solution Consideration

In the inter-AS multi-homing scenario, EFP-uRPF further improves FR-uRPF's directionality, thanks to the availability of the route origin information. More specifically, the construction of RPF lists using EFP-uRPF Algorithm A or B is augmented with data from Route Origin Authorization (ROA) [[RFC6482](#)], as well as Internet Routing Registry (IRR) data, making EFP-uRPF performing better than FR-uRPF regarding directionality. In fact, the global availability of ROA and IRR databases provides a way for the multiple transit providers of the

same multihomed network to share such information without extra way of data synchronization.

In addition, although ERP-uRPF is striving for more accurate RPF list construction, there's still currently no way of constructing an 100%-accurate RPF list in the case shown in Figure 3. In order to to conquer such problem, it could help if devices in the upper level sub-AD(s) (i.e., Router 2 and Router 3) can share more information with each other through certain way.

What's worse, in case of the intra-AS multi-homing, as indicated in Figure2, there's no such prefix to sub-AD mapping (e.g., P3 originates from sub-AD3) database publicly available as ROA or IRR database, or automatically retrievable as RPKI ROA through RTR protocol [[RFC8210](#)]. Thus, enhancing such information sharing between devices of the upper level sub-AD(s) (i.e., Router 2 and Router 3) for the same multi-homed network, by extending certain routing protocols, could be a possible way.

[5.](#) Security Considerations

TBD

[6.](#) Contributors

TBD

[7.](#) Acknowledgments

TBD

[8.](#) Normative References

[I-D.brockners-inband-oam-requirements]

Brockners, F., Bhandari, S., Dara, S., Pignataro, C., Gredler, H., Leddy, J., Youell, S., Mozes, D., Mizrahi, T., Lapukhov, P., and r. remy@barefootnetworks.com, "Requirements for In-situ OAM", [draft-brockners-inband-oam-requirements-03](#) (work in progress), March 2017.

[I-D.ietf-grow-bmp-adj-rib-out]

Evens, T., Bayraktar, S., Lucente, P., Mi, K., and S. Zhuang, "Support for Adj-RIB-Out in BGP Monitoring Protocol (BMP)", [draft-ietf-grow-bmp-adj-rib-out-07](#) (work in progress), August 2019.

Li, et al.

Expires January 13, 2021

[Page 11]

Internet-Draft

SAVA Intra-domain Use Cases

July 2020

[I-D.ietf-grow-bmp-local-rib]

Evens, T., Bayraktar, S., Bhardwaj, M., and P. Lucente, "Support for Local RIB in BGP Monitoring Protocol (BMP)", [draft-ietf-grow-bmp-local-rib-07](#) (work in progress), May 2020.

[I-D.ietf-netconf-yang-push]

Clemm, A. and E. Voit, "Subscription to YANG Datastores", [draft-ietf-netconf-yang-push-25](#) (work in progress), May 2019.

[I-D.openconfig-rtgwg-gnmi-spec]

Shakir, R., Shaikh, A., Borman, P., Hines, M., Lebsack, C., and C. Morrow, "gRPC Network Management Interface (gNMI)", [draft-openconfig-rtgwg-gnmi-spec-01](#) (work in progress), March 2018.

[I-D.song-ntf]

Song, H., Zhou, T., Li, Z., Fioccola, G., Li, Z., Martinez-Julia, P., Ciavaglia, L., and A. Wang, "Toward a Network Telemetry Framework", [draft-song-ntf-02](#) (work in progress), July 2018.

[RFC1157]

Case, J., Fedor, M., Schoffstall, M., and J. Davin, "Simple Network Management Protocol (SNMP)", [RFC 1157](#),

DOI 10.17487/RFC1157, May 1990,
<<https://www.rfc-editor.org/info/rfc1157>>.

[RFC1191] Mogul, J. and S. Deering, "Path MTU discovery", [RFC 1191](#), DOI 10.17487/RFC1191, November 1990,
<<https://www.rfc-editor.org/info/rfc1191>>.

[RFC1195] Callon, R., "Use of OSI IS-IS for routing in TCP/IP and dual environments", [RFC 1195](#), DOI 10.17487/RFC1195, December 1990, <<https://www.rfc-editor.org/info/rfc1195>>.

[RFC1213] McCloghrie, K. and M. Rose, "Management Information Base for Network Management of TCP/IP-based internets: MIB-II", STD 17, [RFC 1213](#), DOI 10.17487/RFC1213, March 1991,
<<https://www.rfc-editor.org/info/rfc1213>>.

[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>>.

Li, et al.

Expires January 13, 2021

[Page 12]

Internet-Draft

SAVA Intra-domain Use Cases

July 2020

[RFC2827] Ferguson, P. and D. Senie, "Network Ingress Filtering: Defeating Denial of Service Attacks which employ IP Source Address Spoofing", [BCP 38](#), [RFC 2827](#), DOI 10.17487/RFC2827, May 2000, <<https://www.rfc-editor.org/info/rfc2827>>.

[RFC3209] Awduche, D., Berger, L., Gan, D., Li, T., Srinivasan, V., and G. Swallow, "RSVP-TE: Extensions to RSVP for LSP Tunnels", [RFC 3209](#), DOI 10.17487/RFC3209, December 2001,
<<https://www.rfc-editor.org/info/rfc3209>>.

[RFC3704] Baker, F. and P. Savola, "Ingress Filtering for Multihomed Networks", [BCP 84](#), [RFC 3704](#), DOI 10.17487/RFC3704, March 2004, <<https://www.rfc-editor.org/info/rfc3704>>.

[RFC3719] Parker, J., Ed., "Recommendations for Interoperable Networks using Intermediate System to Intermediate System (IS-IS)", [RFC 3719](#), DOI 10.17487/RFC3719, February 2004,
<<https://www.rfc-editor.org/info/rfc3719>>.

- [RFC3988] Black, B. and K. Kompella, "Maximum Transmission Unit Signalling Extensions for the Label Distribution Protocol", [RFC 3988](#), DOI 10.17487/RFC3988, January 2005, <<https://www.rfc-editor.org/info/rfc3988>>.
- [RFC6232] Wei, F., Qin, Y., Li, Z., Li, T., and J. Dong, "Purge Originator Identification TLV for IS-IS", [RFC 6232](#), DOI 10.17487/RFC6232, May 2011, <<https://www.rfc-editor.org/info/rfc6232>>.
- [RFC6241] Enns, R., Ed., Bjorklund, M., Ed., Schoenwaelder, J., Ed., and A. Bierman, Ed., "Network Configuration Protocol (NETCONF)", [RFC 6241](#), DOI 10.17487/RFC6241, June 2011, <<https://www.rfc-editor.org/info/rfc6241>>.
- [RFC6959] McPherson, D., Baker, F., and J. Halpern, "Source Address Validation Improvement (SAVI) Threat Scope", [RFC 6959](#), DOI 10.17487/RFC6959, May 2013, <<https://www.rfc-editor.org/info/rfc6959>>.
- [RFC7039] Wu, J., Bi, J., Bagnulo, M., Baker, F., and C. Vogt, Ed., "Source Address Validation Improvement (SAVI) Framework", [RFC 7039](#), DOI 10.17487/RFC7039, October 2013, <<https://www.rfc-editor.org/info/rfc7039>>.

- [RFC7752] Gredler, H., Ed., Medved, J., Previdi, S., Farrel, A., and S. Ray, "North-Bound Distribution of Link-State and Traffic Engineering (TE) Information Using BGP", [RFC 7752](#), DOI 10.17487/RFC7752, March 2016, <<https://www.rfc-editor.org/info/rfc7752>>.
- [RFC7854] Scudder, J., Ed., Fernando, R., and S. Stuart, "BGP Monitoring Protocol (BMP)", [RFC 7854](#), DOI 10.17487/RFC7854, June 2016, <<https://www.rfc-editor.org/info/rfc7854>>.

- [RFC8210] Bush, R. and R. Austein, "The Resource Public Key Infrastructure (RPKI) to Router Protocol, Version 1", [RFC 8210](#), DOI 10.17487/RFC8210, September 2017, <<https://www.rfc-editor.org/info/rfc8210>>.
- [RFC8231] Crabbe, E., Minei, I., Medved, J., and R. Varga, "Path Computation Element Communication Protocol (PCEP) Extensions for Stateful PCE", [RFC 8231](#), DOI 10.17487/RFC8231, September 2017, <<https://www.rfc-editor.org/info/rfc8231>>.
- [RFC8704] Sriram, K., Montgomery, D., and J. Haas, "Enhanced Feasible-Path Unicast Reverse Path Forwarding", [BCP 84](#), [RFC 8704](#), DOI 10.17487/RFC8704, February 2020, <<https://www.rfc-editor.org/info/rfc8704>>.

Authors' Addresses

Dan Li
Tsinghua
Beijing
China

Email: tolidan@tsinghua.edu.cn

Jianping Wu
Tsinghua
Beijing
China

Email: jianping@cernet.edu.cn

Yunan Gu
Huawei
Beijing
China

Email: guyunan@huawei.com

Lancheng Qin
Tsinghua
Beijing
China

Email: qlc19@mails.tsinghua.edu.cn

Tao Lin
H3C
Beijing
China

Email: lintao@h3c.com