Workgroup: Network Working Group
Internet-Draft: draft-qin-savnet-incentive-03
Published: 30 November 2022
Intended Status: Informational
Expires: 3 June 2023
Authors: L. Qin
Tsinghua University
D. Li
Tsinghua University
J. Wu
Tsinghua University
L. Chen
Zhongguancun Laboratory
F. Gao
Zhongguancun Laboratory

The Incentive Consideration for Defense Against Reflection Attacks

Abstract

Source address spoofing remains a significant challenge in today's Internet. Although source address validation (SAV) mechanisms, such as ingress filtering [RFC2827], unicast Reverse Path Forwarding (uRPF) [RFC3704], and the Enhanced Feasible-Path Unicast Reverse Path Forwarding (EFP-uRPF) [RFC8704], have been proposed for a long time, some ASes have not deployed SAV due to the problems of existing SAV mechanism, such as inaccurate validation, misaligned incentive, or other overhead concerns. This document specifically explains the misaligned incentive problem of existing SAV mechanisms and clarifies the direct incentive that a new SAV mechanism should achieve.

Requirements Language

The key words "MUST", "MUST NOT", "REQUIRED", "SHALL", "SHALL NOT", "SHOULD", "SHOULD NOT", "RECOMMENDED", "NOT RECOMMENDED", "MAY", and "OPTIONAL" in this document are to be interpreted as described in BCP 14 [RFC2119][RFC8174] when, and only when, they appear in all capitals, as shown here.

Status of This Memo

This Internet-Draft will expire on 3 June 2023.

Table of Contents

1. Introduction

Source address spoofing is one of the most important security threats in the Internet. By using forged source IP addresses, attackers can well hide their real identities and carry out various

malicious attacks [RFC6959], among which reflection attack is the most common and harmful. In the reflection attack, the attacker spoofs the victim's source IP address and sends requests to servers with reflection and amplification functions, such as DNS or NTP servers. Upon receiving the requests, these servers will reply a large number of responses to the victim, resulting in a large-scale Distributed Denial of Service (DDoS) attack to the victim.

To mitigate source address spoofing, several source address validation (SAV) mechanisms (e.g., ingress filtering [RFC2827], unicast Reverse Path Forwarding (uRPF) [RFC3704], and the Enhanced Feasible-Path Unicast Reverse Path Forwarding (EFP-uRPF) [RFC8704]) have been proposed to identify and reject traffic with forged source IP addresses. However, some ASes have not deployed SAV due to the problems of existing SAV mechanism. Source address spoofing remains a significant challenge in today's Internet.

To help narrow the gap of existing SAV mechanisms, [draft-li-savnet-intra-domain-problem-statement] and [draft-wu-savnet-inter-domain-problem-statement] summarize the fundamental problems of existing SAV mechanisms and define the requirements for new SAV mechanisms. This document further explains the misaligned incentive problem of existing SAV mechanisms and specifies the direct incentive that a new SAV mechanism should achieve. The direct incentive refers to a network deploying SAV can protect itself from being the victim of source address spoofing attacks, especially the most important reflection attacks.

2. Terminology

SAV: Source Address Validation, i.e. validating the authenticity of a packet's source IP address.

Three roles in a reflection attack:

* Attacker. A malicious host that spoofs the victim's source IP address when sending a request to the reflector.

* Reflector. A reflective server (e.g., DNS or NTP server) that receives the forged request and responds to the victim.

* Victim. An innocent host that receives a lot of responses from the reflector, resulting in a DoS attack.

Two results in the incentive comparison between EFP-uRPF and the new SAV mechanism:

* "FAIL" means the victim network cannot help itself prevent the reflection attack by deploying SAV (EFP-uRPF or the new SAV mechanism).

*"WORK" means the victim network can help itself prevent the reflection attack by deploying SAV (EFP-uRPF or the new SAV mechanism).

3. The Importance of Direct Incentive for SAV Deployment

Ingress filtering, or BCP38 [RFC2827] requires the network to implement SAV filtering on its outgoing traffic. If all networks deploy BCP38 and only allow outgoing traffic with legitimate source addresses, source address spoofing can be effectively prevented. However, although BCP38 has been proposed for more than 20 years and is highly recommended by the Mutually Agreed Norms for Routing Security (MANRS), some ASes still do not deploy BCP38. One main reason is that operators lack incentive to deploy BCP38 in their networks. Specifically, BCP38 only prevents the AS who deploys SAV from originating spoofed traffic but does not protect the AS from receiving spoofed traffic or being the victim of an attack. The benefits from deploying BCP38 do not flow to the deployed network, but to the rest of the Internet. As a result, some ASes are reluctant to deploy BCP38 and prefer to wait for others to deploy.

The deployment problem faced by BCP38 tells us that a good SAV mechanism must provide direct incentive/benefits to the deployed network. If a network deploys SAV but finds that it only helps other networks, the network will not be motivated to deploy SAV. If a network deploys SAV and finds that sometimes it can help itself (compared with not deploying), the network will be more motivated to deploy SAV.

4. The Demand for Defense Against Reflection Attack

Nowadays, reflection attack has become one of the most common attacks based on source address spoofing. However, the victim network in a reflection attack may not receive the spoofed request. If an intermediate network deploys SAV to protect itself from receiving spoofed-source traffic, it can help prevent the reflection attack when receiving the spoofed request. Therefore, to mitigate reflection attacks, customer or user networks are increasingly asking their upstreaming providers to deploy SAV as close to the source as possible and to protect their source addresses from being forged. Considering the market demand from customer or user networks, network operators would be willing to improve their competitiveness by providing defense against reflection attacks, so they would attract more users and gain more profits.

However, BCP38 is not aligned with the demand for defense against reflection attacks. The operator who deploys BCP38 neither protects itself from receiving spoofed traffic nor protects its customer or user networks from reflection attacks. More recently, RFC8704 or

BCP84 [RFC8704] proposes the Enhanced Feasible-Path Unicast Reverse Path Forwarding (EFP-uRPF) and recommends operators to apply EFP-uRPF at customer interfaces in most inter-domain scenarios. Different from BCP38, EFP-uRPF provides some direct incentive, as it aims to protect the AS who deploys SAV from receiving spoofed traffic from customer interfaces. Nonetheless, EFP-uRPF is essentially performing ingress filtering at a higher aggregation point (i.e., the top AS of a customer cone). It only validates traffic from customer interfaces but does not validate traffic from provider and peer interfaces. The operator who deploys EFP-uRPF only prevents its customer cone from originating spoofed traffic, but does not protect itself and its customer cone from receiving spoofed traffic or being the victim of a reflection attack from ASes outside the customer cone. Moreover, the victim network will not gain additional protection against reflection attack even if it also deploys EFP-uRPF. Therefore, EFP-uRPF cannot perfectly meet the demand for defense against reflection attacks. If the new SAV mechanism could be well-aligned with the demand for defense against reflection attacks, networks would be more willing to deploy the new SAV mechanism.

5. Incentive Comparison Between EFP-uRPF and the New SAV Mechanism

In the following, we use reflection attack as an example to measure the incentive that EFP-uRPF or the new SAV mechanism can provide to the victim network in the reflection attack. Since there is no mature new SAV mechanism yet, we assume the new SAV mechanism could meet the following requirements proposed in [draft-wu-savnet-inter-domain-problem-statement]:

- Validate traffic from all directions.

- Match the real data-plane forwarding path originated from each deployed AS.

We particularly focus on the partial deployment cases, since it is not practical to require all ASes in the Internet to deploy SAV simultaneously. We first simplify the participants in a reflection attack into three roles (attacker network, reflector network, and victim network) and enumerate different attack scenarios by changing the relative positions of the three roles. In each attack scenario, we suppose the victim network always deploys SAV mechanism (EFP-uRPF or new SAV), because only the victim can get benefit from the defense against reflection attacks. Then, for any deployment case of the other two networks (i.e., attacker network and reflector network), we make the theoretical analysis to check whether the reflection attack can be prevented. If so, the victim network would have strong motivation to deploy SAV; if not, the victim network would have weak motivation to deploy SAV.

5.1. Scenario 1

Figure 1 shows the first reflection attack scenario where the reflector network is located between the attacker network and the victim network. The attacker spoofs the source address of the victim and sends a forged request to the reflector. After receiving the request from attacker, the reflector responds to the victim.

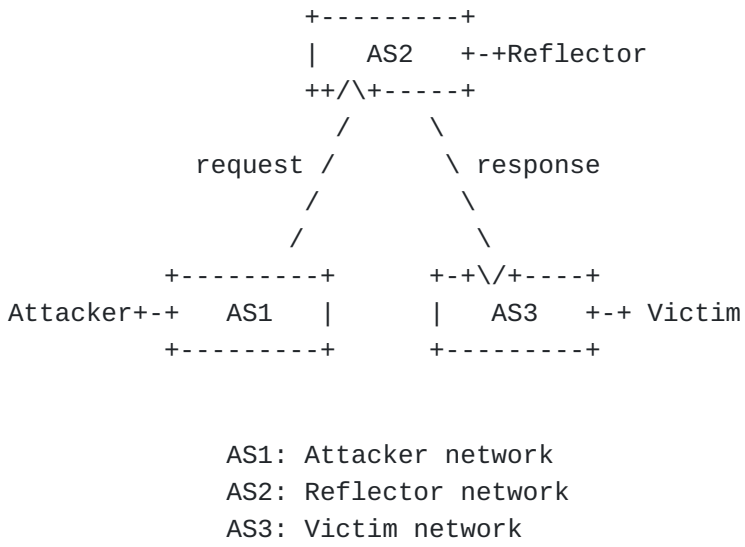

Figure 1: The first reflection attack scenario.

5.1.1. Case 1: only AS3 deploys SAV

Relationship between AS1 and AS2	Relationship between AS2 and AS3	EFP-uRPF algorithm A	EFP-uRPF algorithm B	New SAV
P2C	P2C	FAIL	FAIL	FAIL
P2P	P2C	FAIL	FAIL	FAIL
C2P	C2P	FAIL	FAIL	FAIL
C2P	P2P	FAIL	FAIL	FAIL
C2P	P2C	FAIL	FAIL	FAIL

Table 1: All SAV mechanisms would fail if only AS3 deploys SAV in scenario 1

Table 1 shows the effectiveness of EFP-uRPF and new SAV mechanism against the reflection attack under different relationships among AS1, AS2, and AS3. We omit combinations of AS commercial relationships that violate valley-free principle. If only the victim network deploys SAV, both EFP-uRPF and new SAV mechanism would fail to prevent the reflection attack in scenario 1, because the victim network does not receive the forged request at all.

5.1.2. Case 2: AS1 and AS3 deploy SAV

Relationship between AS1 and AS2	Relationship between AS2 and AS3	EFP-uRPF algorithm A	EFP-uRPF algorithm B	New SAV
P2C	P2C	FAIL	FAIL	WORK
P2P	P2C	FAIL	FAIL	WORK
C2P	C2P	FAIL	FAIL	WORK
C2P	P2P	FAIL	FAIL	WORK
C2P	P2C	FAIL	FAIL	WORK

Table 2: New SAV mechanism could work best if AS1 and AS3 deploy SAV in scenario 1

Table 2 shows that new SAV mechanism works best when victim network and attacker network deploy SAV. If AS1 and AS3 deploy new SAV mechanism, AS1 could learn that traffic with victim's source address must come from outside the AS, not inside the AS. Therefore, the new SAV mechanism in AS1 could successfully detect the forged request and prevent the reflection attack. However, since EFP-uRPF in AS1 does not verify outgoing traffic, EFP-uRPF would fail in this deployment case.

5.1.3. Case 3: AS2 and AS3 deploy SAV

Relationship between AS1 and AS2	Relationship between AS2 and AS3	EFP-uRPF algorithm A	EFP-uRPF algorithm B	New SAV
P2C	P2C	FAIL	FAIL	WORK
P2P	P2C	FAIL	FAIL	WORK
C2P	C2P	WORK	WORK	WORK
C2P	P2P	WORK	WORK	WORK
C2P	P2C	WORK	FAIL	WORK

Table 3: New SAV mechanism could work best if AS2 and AS3 deploy SAV in scenario 1

As shown in Table 3, new SAV mechanism works best when victim network and reflector network deploy SAV. If AS2 and AS3 deploy new SAV mechanism, AS2 could learn that traffic with victim's source address must come from AS3, so it would block the forged request from AS1. If AS2 and AS3 deploy EFP-uRPF, since EFP-uRPF only work for traffic from customer interfaces, EFP-uRPF algorithm A and algorithm B both fail when AS1 is the provider/peer of AS2. EFP-uRPF algorithm A works well when AS1 is the customer of AS2, but EFP-uRPF algorithm B still fails when AS1 and AS3 are both in the customer cone of AS2, because EFP-uRPF algorithm B cannot identify source address spoofing between ASes in the same customer cone.

5.1.4. Case 4: AS1, AS2, and AS3 deploy SAV

Relationship between AS1 and AS2	Relationship between AS2 and AS3	EFP-uRPF algorithm A	EFP-uRPF algorithm B	New SAV
P2C	P2C	FAIL	FAIL	WORK
P2P	P2C	FAIL	FAIL	WORK
C2P	C2P	WORK	WORK	WORK
C2P	P2P	WORK	WORK	WORK
C2P	P2C	WORK	FAIL	WORK

Table 4: New SAV mechanism could work best if AS1, AS2, and AS3 deploy SAV in scenario 1

In scenario 1, the new SAV mechanism would still work best when all three roles deploy SAV. When they deploy the new SAV mechanism, both AS1 and AS2 could effectively identify and block the forged request. When they deploy EFP-uRPF, only AS2 sometimes could prevent the reflection attack, with the same results as Section 4.1.3.

5.2. Scenario 2

Figure 2 shows the second reflection attack scenario. In scenario 2, the victim network is located between the attack network and the reflector network. When attacker sends a forged request to the reflector, the request first arrives at the victim network and then be forwarded to the reflector network. Subsequently, the reflector responds to the victim.

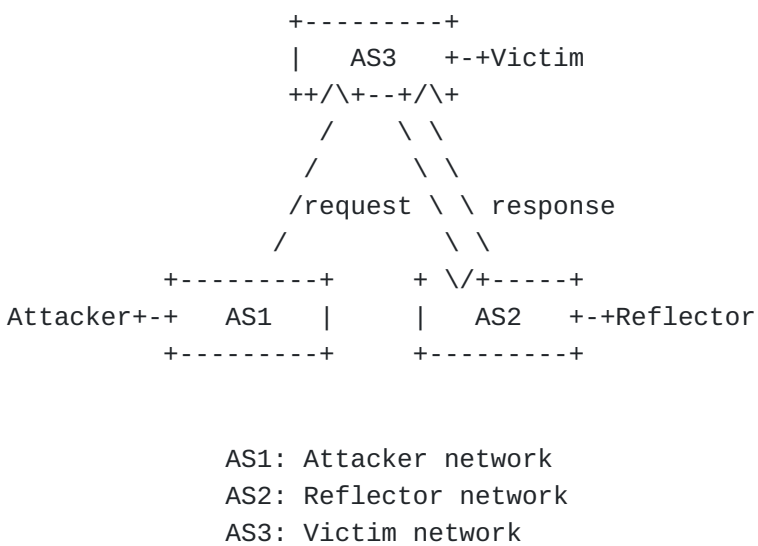

Figure 2: The second reflection attack scenario.

5.2.1. Case 1: only AS3 deploys SAV

Relationship between AS1 and AS3	Relationship between AS3 and AS2	EFP-uRPF algorithm A	EFP-uRPF algorithm B	New SAV
P2C	P2C	FAIL	FAIL	WORK
P2P	P2C	FAIL	FAIL	WORK
C2P	C2P	WORK	WORK	WORK
C2P	P2P	WORK	WORK	WORK
C2P	P2C	WORK	WORK	WORK

Table 5: New SAV mechanism could work best if only AS3 deploys SAV in scenario 2

Table 5 shows the effectiveness of EFP-uRPF and new SAV mechanism when only AS3 in scenario 2 deploys SAV. If AS3 deploys the new SAV mechanism, it could reject the forged request when it receives the forged request. If AS3 deploys EFP-uRPF, it only works when AS1 is the customer of AS3 because EFP-uRPF only implements SAV filtering at customer interfaces.

We also compare EFP-uRPF and the new SAV mechanism in the following three deployment cases. We find that if the SAV mechanism is EFP-uRPF algorithm A or EFP-uRPF algorithm B, only the victim network in scenario 2 would have the possibility to reject the forged request by implementing SAV. Even if attacker network or reflector network also deploys EFP-uRPF, it could not provide additional assistance to victim network. Therefore, on the basis that the victim network has deployed SAV, new SAV mechanism would always work best in different deployment cases.

5.2.2. Case 2: AS1 and AS3 deploy SAV

Relationship between AS1 and AS3	Relationship between AS3 and AS2	EFP-uRPF algorithm A	EFP-uRPF algorithm B	New SAV
P2C	P2C	FAIL	FAIL	WORK
P2P	P2C	FAIL	FAIL	WORK
C2P	C2P	WORK	WORK	WORK
C2P	P2P	WORK	WORK	WORK
C2P	P2C	WORK	WORK	WORK

Table 6: New SAV mechanism could work best if AS1 and AS3 deploy SAV in scenario 2

5.2.3. Case 3: AS2 and AS3 deploy SAV

Relationship between AS1 and AS3	Relationship between AS3 and AS2	EFP-uRPF algorithm A	EFP-uRPF algorithm B	New SAV
P2C	P2C	FAIL	FAIL	WORK

Relationship between AS1 and AS3	Relationship between AS3 and AS2	EFP-uRPF algorithm A	EFP-uRPF algorithm B	New SAV
P2P	P2C	FAIL	FAIL	WORK
C2P	C2P	WORK	WORK	WORK
C2P	P2P	WORK	WORK	WORK
C2P	P2C	WORK	WORK	WORK

Table 7: New SAV mechanism could work best if AS2 and AS3 deploy SAV in scenario 2

5.2.4. Case 4: AS1, AS2, and AS3 deploy SAV

Relationship between AS1 and AS3	Relationship between AS3 and AS2	EFP-uRPF algorithm A	EFP-uRPF algorithm B	New SAV
P2C	P2C	FAIL	FAIL	WORK
P2P	P2C	FAIL	FAIL	WORK
C2P	C2P	WORK	WORK	WORK
C2P	P2P	WORK	WORK	WORK
C2P	P2C	WORK	WORK	WORK

Table 8: New SAV mechanism could work best if AS1, AS2, and AS3 deploy SAV in scenario 2

5.3. Scenario 3

Figure 3 shows the third reflection attack scenario. The attacker network is located between the victim network and the reflector network. Attacker spoofs victim's source address in the request sent to reflector. Reflector receives the request from the attacker network and sends a response to the victim network via the attacker network.

Below we make the incentive comparison between EFP-uRPF and the new SAV mechanism in scenario 3. By varying SAV deployment status of attacker network and reflector network, we find all SAV mechanisms would fail in preventing the reflection attack in this scenario. For victim network, it does not receive the forged request. For attacker network and reflector network, SAV in their networks could not identify this spoofing because the forged source address (i.e., victim's source address) shares the same valid incoming interface with the actual one (i.e., attacker's source address) in the SAV rules.

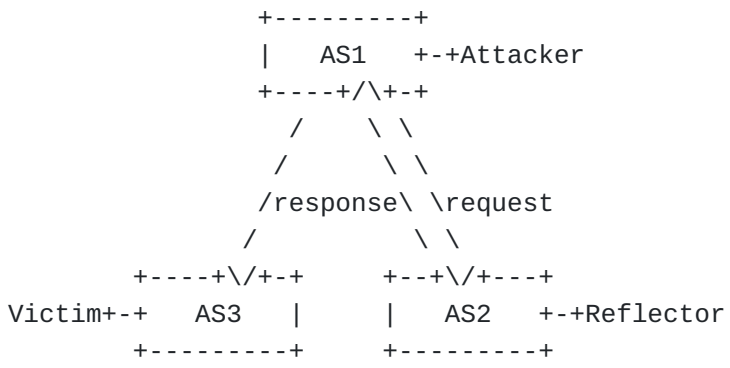

Figure 3: The third reflection attack scenario.

5.3.1. Case 1: only AS3 deploys SAV

Relationship between AS3 and AS1	Relationship between AS1 and AS2	EFP-uRPF algorithm A	EFP-uRPF algorithm B	New SAV
P2C	P2C	FAIL	FAIL	FAIL
P2P	P2C	FAIL	FAIL	FAIL
C2P	C2P	FAIL	FAIL	FAIL
C2P	P2P	FAIL	FAIL	FAIL
C2P	P2C	FAIL	FAIL	FAIL

Table 9: All SAV mechanisms would fail if only AS3 deploys SAV in scenario 3

5.3.2. Case 2: AS1 and AS3 deploy SAV

Relationship between AS3 and AS1	Relationship between AS1 and AS2	EFP-uRPF algorithm A	EFP-uRPF algorithm B	New SAV
P2C	P2C	FAIL	FAIL	FAIL
P2P	P2C	FAIL	FAIL	FAIL
C2P	C2P	FAIL	FAIL	FAIL
C2P	P2P	FAIL	FAIL	FAIL
C2P	P2C	FAIL	FAIL	FAIL

Table 10: All SAV mechanisms would fail if AS1 and AS3 deploy SAV in scenario 3

5.3.3. Case 3: AS2 and AS3 deploy SAV

Relationship between AS3 and AS1	Relationship between AS1 and AS2	EFP-uRPF algorithm A	EFP-uRPF algorithm B	New SAV
P2C	P2C	FAIL	FAIL	FAIL
P2P	P2C	FAIL	FAIL	FAIL
C2P	C2P	FAIL	FAIL	FAIL
C2P	P2P	FAIL	FAIL	FAIL
C2P	P2C	FAIL	FAIL	FAIL

Table 11: All SAV mechanisms would fail if AS2 and AS3 deploy SAV in scenario 3

5.3.4. Case 4: AS1, AS2, and AS3 deploy SAV

Relationship between AS3 and AS1	Relationship between AS1 and AS2	EFP-uRPF algorithm A	EFP-uRPF algorithm B	New SAV
P2C	P2C	FAIL	FAIL	FAIL
P2P	P2C	FAIL	FAIL	FAIL
C2P	C2P	FAIL	FAIL	FAIL
C2P	P2P	FAIL	FAIL	FAIL
C2P	P2C	FAIL	FAIL	FAIL

Table 12: All SAV mechanisms would fail if AS1, AS2, and AS3 deploy SAV in scenario 3

6. Summary

Overall, neither the new SAV mechanism nor EFP-uRPF could completely prevent the reflection attack. But for any attack scenario or deployment case, we find that the new SAV mechanism could work better or not worse than EFP-uRPF. It is worth noting that AS1 and AS2 in above scenarios can also be targets of reflection attacks from other networks. Therefore, a network could have more incentive to deploy the new SAV mechanism, because it would have high probability of defending against reflection attacks

7. Acknowledgments

TBD

8. Normative References

[draft-li-savnet-intra-domain-problem-statement]
Li, D., Wu, J., Qin, L., Huang, M., and N. Geng, "Source Address Validation in Intra-domain Networks (Intra-domain

SAVNET) Gap Analysis, Problem Statement and Requirements", 30 November 2022.

[draft-wu-savnet-inter-domain-problem-statement] Wu, J., Li, D., Qin, L., Huang, M., and N. Geng, "Source Address Validation in Inter-domain Networks (Inter-domain SAVNET) Gap Analysis, Problem Statement and Requirements", 30 November 2022.

[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>.

[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>.

[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>.

[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>.

[RFC8174] Leiba, B., "Ambiguity of Uppercase vs Lowercase in RFC 2119 Key Words", BCP 14, RFC 8174, DOI 10.17487/RFC8174, May 2017, <https://www.rfc-editor.org/info/rfc8174>.

[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

Lancheng Qin
Tsinghua University
Beijing
China

Email: qlc19@mails.tsinghua.edu.cn

Dan Li
Tsinghua University

Beijing
China

Email: tolidan@tsinghua.edu.cn

Jianping Wu
Tsinghua University
Beijing
China

Email: jianping@cernet.edu.cn

Li Chen
Zhongguancun Laboratory
Beijing
China

Email: Lichen@zgclab.edu.cn

Fang Gao
Zhongguancun Laboratory
Beijing
China

Email: gaofang@zgclab.edu.cn