Workgroup: Network Working Group
Internet-Draft:
draft-li-savnet-intra-domain-problem-
statement-07
Published: 11 March 2023
Intended Status: Informational
Expires: 12 September 2023
Authors: D. Li, Tsinghua University
J. Wu, Tsinghua University
L. Qin, Tsinghua University
M. Huang, Huawei
N. Geng, Huawei

Source Address Validation in Intra-domain Networks Gap Analysis, Problem Statement, and Requirements

Abstract

This document provides the gap analysis of existing intra-domain source address validation mechanisms, describes the fundamental problems, and defines the requirements for technical improvements.

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 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 12 September 2023.

Copyright Notice

Table of Contents

1. Introduction

Source Address Validation (SAV) is important for defending against source address spoofing attacks and allowing accurate traceback. A multi-fence architecture called Source Address Validation Architecture (SAVA) [RFC5210] was proposed to validate source addresses at three levels: access network SAV, intra-domain SAV, and inter-domain SAV. When SAV is not fully enabled at the edge of the Internet, the multi-fence architecture can help enhance the validation across the whole Internet and thus reduce the opportunities of launching source address spoofing attacks.

Particularly, access network SAV ensures that a host uses a valid address assigned to the host statically or dynamically. In this way, the host cannot use the source address of another host. There are many mechanisms for SAV in access networks. Static ACL rules can be manually configured for validation by specifying which source addre-

sses are acceptable or unacceptable. Dynamic ACL is another efficient mechanism which is associated with authentication servers (e.g., RADIUS and DIAMETER). The servers receive access requests and then install or enable ACL rules on the device to permit particular users' packets. SAVI [RFC7039] represents a kind of mechanism enforcing that the legitimate IP address of a host matches the link-layer property of the host's network attachment. For example, SAVI solution for DHCP [RFC7513] creates a binding between a DHCPv4/DHCPv6-assigned IP address and a link-layer property (like MAC address or switch port) on a SAVI device. IP Source Guard (IPSG) [IPSG] combined with DHCP snooping is an implementation of SAVI solution for DHCP. Cable Source-Verify [cable-verify] also shares some features of SAVI and is used in cable modem networks. Cable modem termination system (CMTS) devices with Cable Source-Verify maintain the bindings of the CPE's IP address, the CPE's MAC address, and the corresponding cable modem identifier. When receiving packets, the device will check the validity of the packets according to the bindings.

Given numerous access networks managed by different operators throughout the world, it is difficult to require all access networks to effectively deploy SAV. Therefore, intra-domain SAV and inter-domain SAV are needed to block spoofing traffic as close to the source as possible. Both intra-domain SAV and inter-domain SAV usually perform validation at the granularity of IP prefixes, which is coarser than the validation granularity of access network SAV, as an IP prefix covers a range of IP addresses.

This document focuses on the analysis of intra-domain SAV. In contrast to inter-domain SAV, intra-domain SAV does not require collaboration between different ASes. The SAV rules can be generated by the AS itself. Consider an AS X which provides its own subnets with the connectivity to other ASes. The intra-domain SAV for AS X has two goals: i) blocking the illegitimate packets originated from the local subnets of AS X with spoofed source addresses; and ii) blocking the illegitimate packets coming from other ASes which spoof the source addresses of AS X.

Figure 1 illustrates the function of intra-domain SAV with two cases. Case i shows that AS X forwards spoofed packets originated from its subnets to other ASes (e.g., AS Y). If AS X deploys intra-domain SAV, the spoofed packets from its own subnet can be blocked by AS X itself (i.e., Goal i). Case ii shows that AS X receives the packets which spoof AS X's source addresses from other ASes (e.g., AS Y). If AS X deploys intra-domain SAV, the spoofed packets from AS Y can be blocked by AS X (i.e., Goal ii).

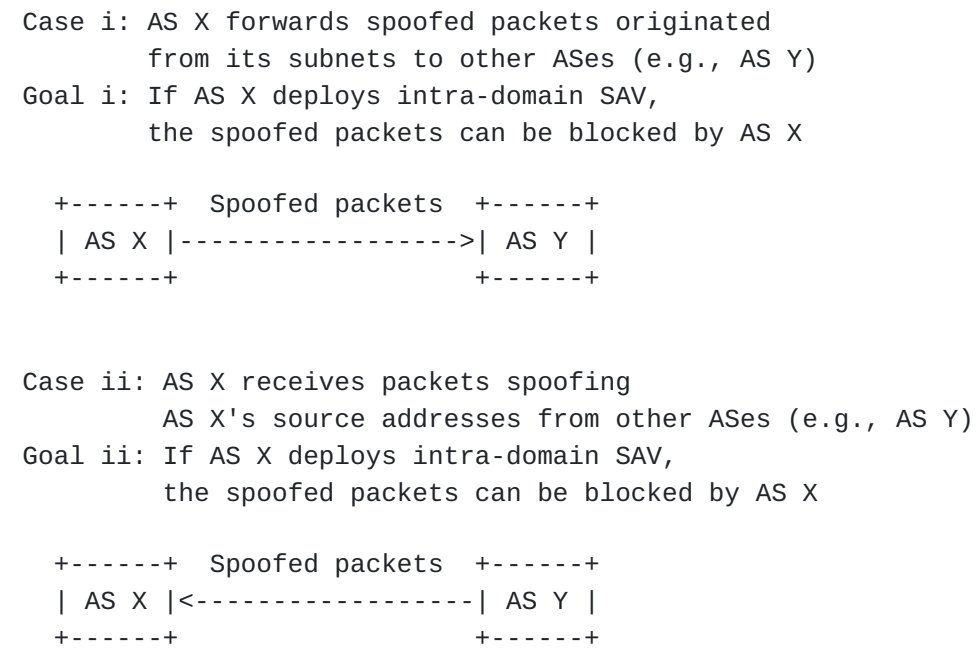

Figure 1: An example for illustrating intra-domain SAV

There are many mechanisms for intra-domain SAV. This document provides the gap analysis of existing intra-domain SAV mechanisms. According to the gap analysis, the document concludes the main problems of existing mechanisms and describes the requirements for future intra-domain SAV mechanisms.

2. Terminology

SAV Rule: The rule that indicates the validity of a specific source IP address or source IP prefix.

SAV Table: The table or data structure that implements the SAV rules and is used for source address validation in the data plane.

Improper Block: The validation results that the packets with legitimate source addresses are blocked improperly due to inaccurate SAV rules.

Improper Permit: The validation results that the packets with spoofed source addresses are permitted improperly due to inaccurate SAV rules.

3. Existing Mechanisms

Ingress filtering [RFC2827][RFC3704] is the current practice of intra-domain SAV. This section briefly introduces the existing intra-domain SAV mechanisms.

*ACL-based ingress filtering [RFC2827][RFC3704] is a typical mechanism for intra-domain SAV. ACL rules can be configured for blocking or permitting packets with specific source addresses. This mechanism can be applied at the downstream interfaces of edge routers connecting the subnets or at the downstream interfaces of aggregation routers [manrs-antispoofing]. The validation at downstream interfaces will prevent local subnets from spoofing source prefixes of other subnets. Besides, at the upstream interfaces of routers connecting other ASes, ACL can be enabled for blocking packets with disallowed source prefixes, such as the internal source prefixes owned by the subnets [nist-rec]. In any application scenario, ACL rules should be updated in time to be consistent with the latest filtering criteria.

*Strict uRPF [RFC3704] is another commonly used mechanism for SAV in intra-domain networks. Routers deploying strict uRPF accept a data packet only when i) the local FIB contains a prefix encompassing the packet's source address and ii) the corresponding outgoing interface for the prefix in the FIB matches the packet's incoming interface. Otherwise, the packet will be blocked. Strict uRPF is usually used at downstream interfaces of edge routers connecting local subnets.

*Loose uRPF [RFC3704] takes a looser validation mechanism than strict uRPF to avoid improper block. A packet will be accepted if the local FIB contains a prefix encompassing the packet's source address. The incoming interface of the packet is not checked. Upstream interfaces can enable loose uRPF for blocking non-global addresses [nist-rec].

*Carrier Grade NAT has some operations on the source addresses of packets, but is not an anti-spoofing tool, as described in [manrs-antispoofing]. If the source address of a packet is in the INSIDE access list, the NAT rule can translate the source address to an address in the pool OUTSIDE. The NAT rule cannot judge whether the source address is spoofed or not. In addition, the packet with a spoofed source address will be forwarded directly if the spoofed source address is not included in the INSIDE access list. Therefore, Carrier Grade NAT cannot help block or traceback spoofed packets, and other SAV mechanisms are still needed.

4. Gap Analysis

Existing intra-domain SAV mechanisms either require high operational overhead or have limitations in accuracy. They may improperly block the traffic with legitimate source addresses (i.e., improper block) or improperly permit the traffic with spoofed source addresses (i.e., improper permit).

4.1. Outbound Traffic Validation

Outbound traffic validation is implemented at downstream interfaces of routers to validate the packets from directly connected subnets. As described previously, ACL rules can be configured at downstream interfaces for ingress filtering. These rules need to be updated when prefixes or topologies of subnets change. If ACL rules are not updated in time, improper block or improper permit may occur. To ensure the accuracy of SAV in dynamic networks, high operational overhead will be induced to achieve timely updates for ACL configurations.

Strict uRPF can also be used for outbound traffic validation, but there may be improper block problem in multi-homing scenario. Figure 2 shows such a case. In the figure, Subnet 1 owns prefix 10.0.0.0/15 and is attached to two edge routers, i.e., Router 1 and Router 2. For the load balance purpose of inbound traffic, Subnet 1 expects the inbound traffic destined for 10.1.0.0/16 to come only from Router 1 and the inbound traffic destined for 10.0.0.0/16 to come only from Router 2. To this end, Router 1 only learns the route to sub prefix 10.1.0.0/16 from Subnet 1, while Router 2 only learns the route to the other sub prefix 10.0.0.0/16 from Subnet 1. Then, Router 1 and Router 2 advertise the learned sub prefix to the other routers in the AS through intra-domain routing protocols such as OSPF or IS-IS. Finally, Router 1 learns the route to 10.0.0.0/16 from Router 3, and Router 2 learns the route to 10.1.0.0/16 from Router 3. The FIBs on Router 1 and Router 2 are shown in the figure. Although Subnet 1 does not expect inbound traffic destined for 10.0.0.0/16 to come from Router 1, it may send outbound traffic with source addre-sses of prefix 10.0.0.0/16 to Router 1 for load balance of outbound traffic. As a result, there is asymmetric routing between Subnet 1 and Router 1. Similarly, Subnet 1 may also send outbound traffic with source addresses of prefix 10.1.0.0/16 to Router 2, resulting in asymmetric routing between Subnet1 and Router 2.

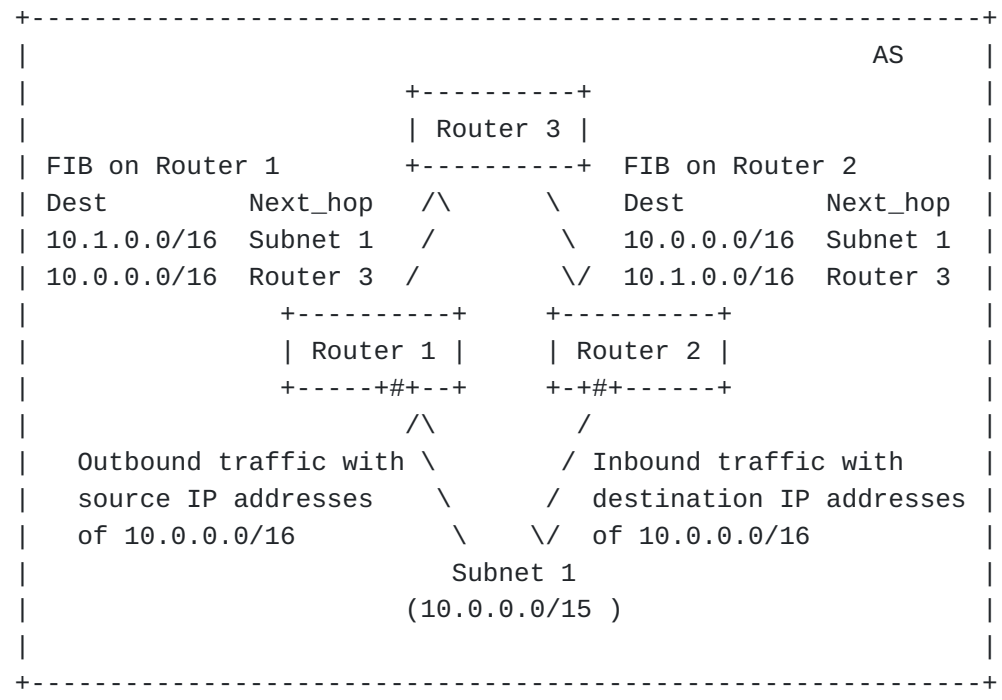

Figure 2: Asymmetric routing in the multi-homed subnet scenario

Strict uRPF takes the entries in FIB for SAV. It can improperly block the packets with legitimate source prefixes when asymmetric routing exists. In the figure, if Router 1 applies strict uRPF at interface '#', the SAV rule is that Router 1 only accepts packets with source addresses of 10.1.0.0/16 from Subnet 1. Therefore, when Subnet 1 sends packets with source addresses of 10.0.0.0/16 to Router 1, strict uRPF at Router 1 will improperly block these legitimate packets. Similarly, when Router 2 with strict uRPF deployed receives packets with source addresses of prefix 10.1.0.0/16 from Subnet 1, it will also improperly block these legitimate packets. Therefore, strict uRPF may cause improper block problem in the case of asymmetric routing.

4.2. Inbound Traffic Validation

Inbound traffic validation is performed at upstream interfaces of border routers to validate the packets from other ASes. Figure 3 shows an example of inbound SAV. In the figure, Router 3 and Router 4 deploy SAV mechanisms at interface '#' for validating external packets. Hence, there are multiple points for inbound traffic validation for the AS.

ACL-based ingress filtering is usually used for validating inbound traffic. By configuring specified ACL rules, inbound packets with disallowed source prefixes (e.g., non-global addresses or the internal source prefixes) can be blocked. As mentioned above, ACL-based ingress filtering requires timely updates when the routing

status changes dynamically. When the ACL rules are not updated in time, there may be improper block or improper permit problems. The operational overhead of maintaining updated ACL rules will be extremely high when there are multiple inbound validation points as shown in Figure 3.

Loose uRPF is another inbound SAV mechanism and is more adaptive than ACL-based rules. But it sacrifices the directionality of SAV and has limited blocking capability, because it allows packets with source addresses that exist in the FIB table at all router interfaces.

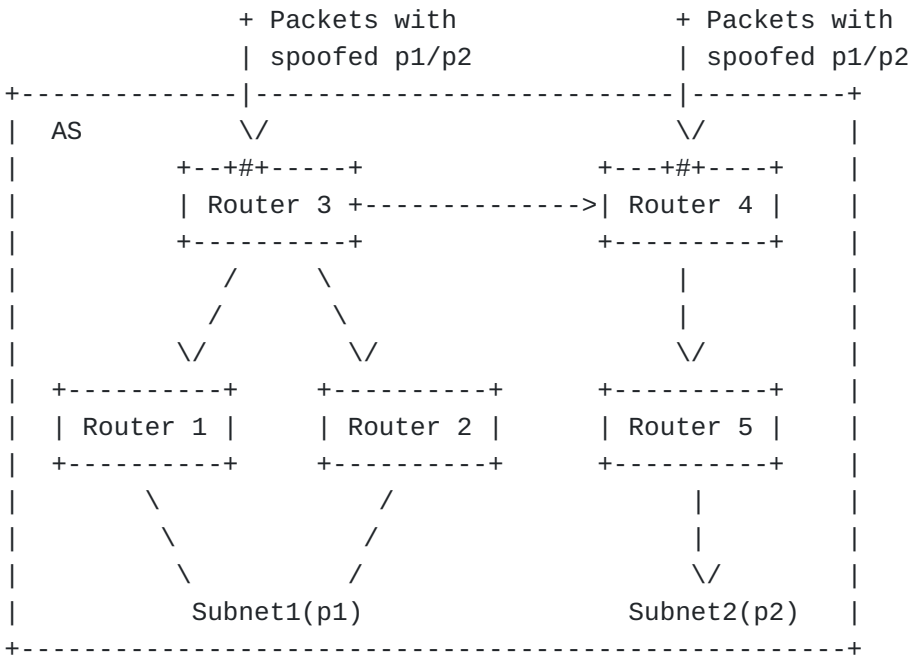

Figure 3: A scenario of inbound SAV

5. Problem Statement

Accurate validation and low operational overhead are two important design goals of intra-domain SAV mechanisms. As analyzed above, asymmetric routing and dynamic networks are two challenging scenarios for the two goals. In these scenarios, existing SAV mechanisms have problems of inaccurate validation or high operational overhead.

ACL-based SAV relies on manual configurations and thus requires high operational overhead in dynamic networks. Operators have to manually update the ACL-based filtering rules in time when the prefix or topology changes. Otherwise, improper block or improper permit problems may appear.

Strict uRPF-based SAV can automatically update SAV rules, but may improperly block legitimate traffic under asymmetric routing. The root cause is that strict uRPF leverages the local FIB table to determine the incoming interface for source addresses, which may not match the real data-plane forwarding path from the source, due to the existence of asymmetric routes. Hence, it may mistakenly consider a valid incoming interface as invalid, resulting in improper block problem; or it may consider an invalid incoming interface as valid, resulting in improper permit problem.

Loose uRPF is also an automated SAV mechanism but its SAV rules are overly loose. Most spoofed packets will be improperly permitted by adopting loose uRPF.

6. Requirements for New SAV Mechanisms

This section lists the requirements which can be a guidance for narrowing the gaps of existing intra-domain SAV mechanisms. The requirements can be fully or partially fulfilled when designing new intra-domain SAV mechanisms.

6.1. Automatic Update

The new intra-domain SAV mechanism MUST be able to automatically adapt to network dynamics such as routing change or prefix change, instead of relying on manual update.

6.2. Accurate Validation

The new intra-domain SAV mechanism needs to improve the validation accuracy upon existing intra-domain SAV mechanisms. Improper block must be avoided to guarantee that legitimate traffic will not be blocked. Improper permit must be reduced as much as possible. To avoid improper block in asymmetric routing scenario, it is better that the real forwarding path in the data plane can be learned and incoming interface for a certain prefix can be set accordingly. In case when the real forwarding path in the data plane cannot be learned, the learned paths must cover the real forwarding paths so as to avoid improper block. Further, by finding the least number of paths while covering all the real forwarding paths, improper permit can be minimized.

6.3. Working in Incremental/Partial Deployment

The new intra-domain SAV mechanism SHOULD NOT assume pervasive adoption. Some routers may not be able to be easily upgraded for supporting the new SAV mechanism due to their limitations of capabilities, versions, or vendors. The mechanism SHOULD be able to provide protection even when it is partially deployed. The effectiveness of protection for the new intra-domain SAV mechanism

under partial deployment SHOULD be no worse than existing mechanisms.

7. Intra-domain SAV Scope

The new intra-domain SAV mechanism should work in the same scenarios as existing intra-domain SAV mechanisms. Generally, it includes all IP-encapsulated scenarios:

- *Native IP forwarding: including both forwarding based on global routing table and CE site forwarding of VPN.
- *IP-encapsulated Tunnel (IPsec, GRE, SRv6, etc.): focusing on the validation of the outer layer IP address.
- *Validating both IPv4 and IPv6 addresses.

Scope does not include:

- *Non-IP packets: including MPLS label-based forwarding and other non-IP-based forwarding.

In addition, the new intra-domain SAV mechanism should avoid data-plane packet modification. Existing architectures or protocols or mechanisms can be used in the new SAV mechanism to achieve better SAV function.

8. Security Considerations

The new intra-domain SAV mechanism MUST NOT introduce additional security vulnerabilities or confusion to the existing intra-domain architectures or control or management plane protocols. Similar to the security scope of intra-domain routing protocols, intra-domain SAV mechanism should ensure integrity and authentication of protocol packets that deliver the required SAV information.

The new intra-domain SAV mechanism does not provide protection against compromised or misconfigured routers that poison existing control plane protocols. Such routers can not only disrupt the SAV function, but also affect the entire routing domain.

9. IANA Considerations

This document does not request any IANA allocations.

10. Acknowledgements

Many thanks to the valuable comments from: Jared Mauch, Barry Greene, Fang Gao, Anthony Somerset, Kotikalapudi Sriram, Yuanyuan Zhang, Igor Lubashev, Alvaro Retana, Joel Halpern, Aijun Wang,

Michael Richardson, Li Chen, Gert Doering, Mingxing Liu, Libin Liu, John O'Brien, Roland Dobbins, etc.

11. References

11.1. Normative References

[inter-domain] "Source Address Validation in Inter-domain Networks Gap Analysis, Problem Statement, and Requirements", 2023, <https://datatracker.ietf.org/doc/draft-wu-savnet-inter-domain-problem-statement/>.

[manrs-antispoofing] MANRS, "MANRS Implementation Guide", 2023, <https://www.manrs.org/netops/guide/antispoofing/>.

[nist-rec] Sriram, K. and D. Montgomery, "Resilient Interdomain Traffic Exchange: BGP Security and DDos Mitigation", 2019, <https://www.nist.gov/publications/resilient-interdomain-traffic-exchange-bgp-security-and-ddos-mitigation>.

[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>.

[RFC5210] Wu, J., Bi, J., Li, X., Ren, G., Xu, K., and M. Williams, "A Source Address Validation Architecture (SAVA) Testbed and Deployment Experience", RFC 5210, DOI 10.17487/RFC5210, June 2008, <https://www.rfc-editor.org/info/rfc5210>.

[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>.

11.2. Informative References

[cable-verify]

Cisco, "Cable Source-Verify and IP Address Security", 2021, <https://www.cisco.com/c/en/us/support/docs/broadband-cable/cable-security/20691-source-verify.html>.

[IPSG] Cisco, "Configuring DHCP Features and IP Source Guard", 2016, <https://www.cisco.com/c/en/us/td/docs/switches/lan/catalyst2960/software/release/12-2_53_se/configuration/guide/2960scg/swdhcp82.html>.

[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>.

[RFC7513] Bi, J., Wu, J., Yao, G., and F. Baker, "Source Address Validation Improvement (SAVI) Solution for DHCP", RFC 7513, DOI 10.17487/RFC7513, May 2015, <https://www.rfc-editor.org/info/rfc7513>.

Authors' Addresses

Dan Li
Tsinghua University
Beijing
China

Email: tolidan@tsinghua.edu.cn

Jianping Wu
Tsinghua University
Beijing
China

Email: jianping@cernet.edu.cn

Lancheng Qin
Tsinghua University
Beijing
China

Email: qlc19@mails.tsinghua.edu.cn

Mingqing Huang
Huawei
Beijing
China

Email: huangmingqing@huawei.com

Nan Geng

Huawei
Beijing
China

Email: gengnan@huawei.com