idr
Internet-Draft
Intended status: Standards Track
Expires: 6 January 2026

T. Tong
China Unicom
D. Li
Tsinghua University
N. Geng
Huawei
N. Wang
China Unicom
S. Zhuang
Huawei
J. Zhao
China Unicom

5 July 2025

Advertisement of Multi-Sourced SAV Rules using BGP Link-State draft-tong-idr-bgp-ls-sav-rule-02

Abstract

This document describes the protocol extensions of BGP Link-State to collect source address validation (SAV) rules generated by different protocols/mechanisms, to facilitate multi-sourced SAV rules monitoring and management.

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 6 January 2026.

Copyright Notice

Table of Contents

1. Introduction

Source Address Validation (SAV) can efficiently prevent source address spoofing-based attacks. SAV rules, which indicate the valid/invalid incoming interfaces of a specific source IP address or source IP prefix, are installed on routers for checking the source addresses of received packets.

SAV rules can be generated by static configuration, management tools, or based on different routing protocols such as OSPFv2, OSPFv3, IS-IS, BGP, or their extensions [I-D.ietf-savnet-intra-domain-architecture][I-D.ietf-savnet-inter-domain-architecture]. Due to the requirements of application scenarios, a router may use more than one tool at the same time to obtain SAV rules. Therefore, the rules on the router will be multi-sourced, which complicates management. What is more challenging is that there may exist conflicts among these multi-sourced rules and the rules can be dynamic.

To facilitate SAV rules monitoring and management, this document proposes to extend BGP-LS ([RFC9552]) for advertising SAV rules on routers to a centralized server. The centralized server can effectively collect multi-sourced SAV rules from routers. For the purpose of advertising SAV rules within BGP-LS advertisements, two new NLRIs called SAV Rule NLRIs are proposed for IPv4 and IPv6, respectively.

1.1. 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.

2. BGP-LS NLRI Advertisement for SAV Rules

The "Link-State NLRI" defined in [RFC9552] is extended to carry the SAV rule information. The format of "Link-State NLRI" is defined in [RFC9552] as follows:

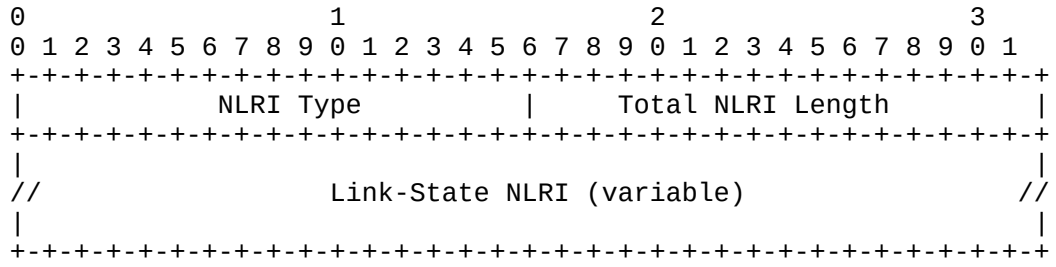

Figure 1: Link-State NLRI

This document defines two new NLRI Types known as SAV Rule NLRIs (values are TBD) for the advertisement of SAV rule Information.

2.1. SAV Rule NLRIs

This document defines SAV Rule NLRI Types with their common format as shown in the following figure:

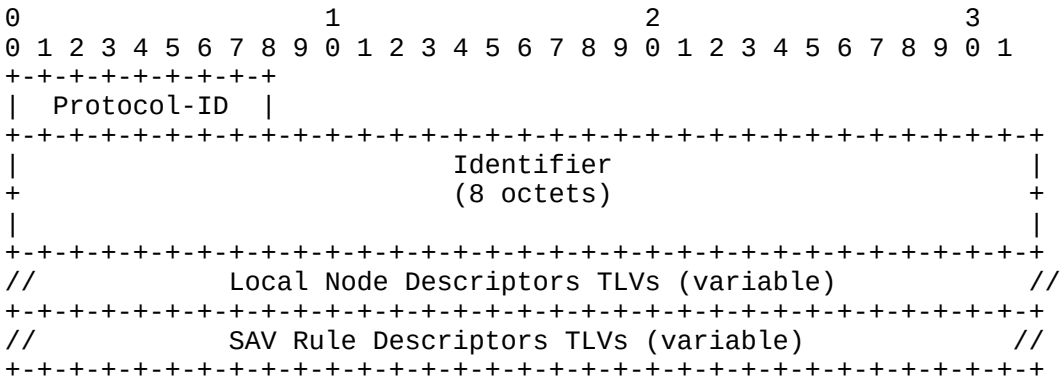

Figure 2: BGP-LS SAV Rule NLRI

The fields are defined as follows:

* Protocol-ID: Specifies the source of SAV rules in this NLRI. Protocol-ID values defined in RFC9552 [RFC9086] can be reused.

* Identifier: An 8 octet value as defined in [RFC9552].

* Local Node Descriptors TLV: Contains Node Descriptors for the nodes storing SAV rules. This is a mandatory TLV in SAV Rule NLRIs. The Type is 256. The length of this TLV is variable. The value contains one or more Node Descriptor sub-TLVs defined in [RFC9552].

* SAV Rule Descriptors TLVs: There can be one or more SAV Rule Descriptors TLVs for carrying SAV rules.

2.2. SAV Rule Descriptors TLVs

The SAV Rule Descriptor field is a set of TLV triplets. SAV Rule Descriptors TLVs identify a set of SAV rules having the same set of valid interfaces as defined in [I-D.ietf-savnet-general-sav-capabilities]. The following TLVs are valid as SAV Rule Descriptors in the SAV Rule NLRI:

TLV Code Point	Description	Length
TBD	Interface Name	variable
TBD	Interface Group	4
TBD	SAV Prefix	variable

Figure 3: SAV Rule Descriptor TLVs

2.2.1. Interface Name TLV

An Interface Name TLV is used to identify one valid interface of the source prefixes carried in SAV Prefix TLVs. The format of Interface Name TLV is as follows:

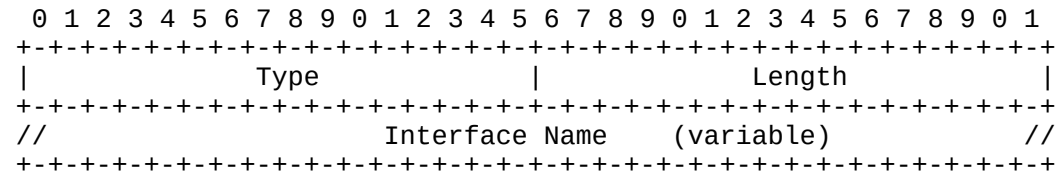

Figure 4: Interface Name TLV

There can be zero, one or more Interface Name TLVs in the SAV Rule Descriptor field.

2.2.2. Interface Group TLV

An Interface Group TLV is to identify a group of valid interfaces of the source prefixes carried in SAV Prefix TLVs. The format of Interface Group TLV is as follows:

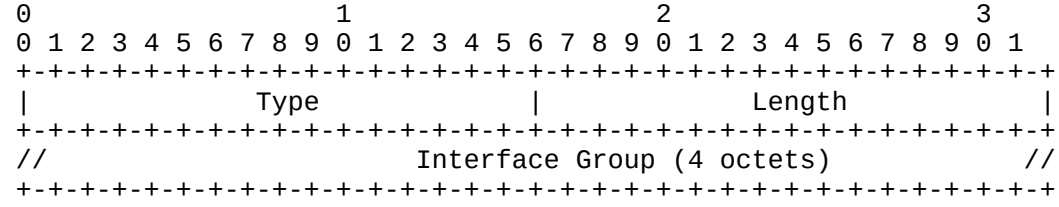

Figure 5: Interface Group TLV

The Interface Group value can have either a local meaning or a global meaning. On the one hand, it can be a local interface property on the target routers, and the meaning of it depends on the configurations of network administrator [I-D.ietf-idr-flowspec-interfaceset]. On the other hand, a global meaning Group Identifier field carries an AS number, which represents all the interfaces connected to the neighboring AS with the AS number. [I-D.geng-idr-flowspec-sav]

Interface Group value can also be an Interface ID for identifying a specific interface.

There can be zero, one or more Interface Group TLVs in the SAV Rule Descriptor field. Interface Group TLVs can be used together with Interface Name TLVs.

When there is neither an Interface Name TLV nor an Interface Group TLV, the source prefixes carried in SAV Prefix TLVs are considered valid for all the interfaces on the router.

2.2.3. SAV Prefix TLV

A SAV Prefix TLV carries one IP address prefix (IPv4 or IPv6). The format of SAV Prefix TLV is as follows:

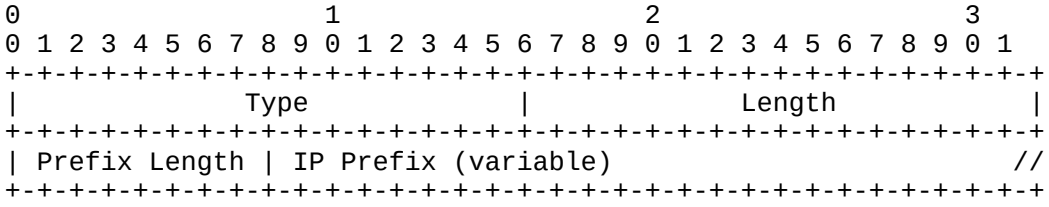

Figure 6: SAV Prefix TLV

There can be one or more SAV Prefix TLVs in the SAV Rule Descriptor field. The IPv4 SAV Prefix TLVs will only appear in the IPv4 SAV Rule NLRI, and The IPv6 SAV Prefix TLVs are only for the IPv6 SAV Rule NLRI

There can be more than one SAV mechanisms based on the same source (identified by Protocol-ID). In order to distinguish the different sources of rules in a more fine-grained manner, the Type field needs to be allocated for multiple values, and each value identifies a specific SAV mechanism based on the same source identified by Protocol-ID.

3. BGP-LS Attribute for SAV Mode

The BGP-LS Attribute, an optional and non-transitive BGP Attribute, is used to carry the validation mode information of SAV rules {I-D.ietf-savnet-general-sav-capabilities}. The following SAV Mode Attribute TLV is defined for the BGP-LS Attribute associated with a SAV Rule NLRI:

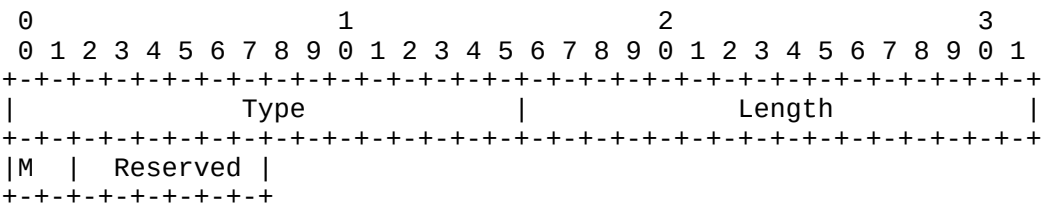

Figure 7: SAV Mode TLV

The SAV Mode TLV carries a Mode Flag (M flag is shown in the figure and occupies two bits) describing the validation mode attribute.

* When M flag is set to 00, the mode is Mode 1: interface-based source prefix allowlist. The NLRI carries the source prefixes and interfaces. Only the carried prefixes are valid on the carried interfaces, and any other prefixes are invalid on these interfaces.

* When M flag is set to 01, the mode is Mode 2: interface-based source prefix blocklist. The NLRI carries the source prefixes and interfaces. Only the carried prefixes are invalid on the carried interfaces, and any other prefixes are valid on these interfaces.

* When M flag is set to 10, the mode is Mode 3: prefix-based interface allowlist. The NLRI carries the source prefixes and interfaces. Only the carried interfaces are valid for the carried prefixes, and any other interfaces are invalid for those prefixes. Any other prefixes will not be validated.

* When M flag is set to 11, the mode is Mode 4: prefix-based interface blocklist. The NLRI carries the source prefixes and interfaces. Only the carried interfaces are invalid for the carried prefixes, and any other interfaces are valid for those prefixes. Any other prefixes will not be validated.

4. BGP-LS Attribute for SAV Actions

SAV actions in this document adopt the traffic filtering actions defined in [RFC8955] and [RFC8956].

Traffic filtering actions defined in [RFC 8955] include traffic-rate-bytes, traffic-rate-packets, traffic-action, rt-redirect, and traffic-marking, which are applicable to IPv4 and IPv6. Rt-redirect-ipv6 is a new traffic filtering action defined in [RFC 8956], which is applicable to IPv6. The encapsulation formats of SAV actions are consistent with the encapsulation formats defined in [RFC 8955] and [RFC 8956].

A SAV rule may match multiple SAV actions, and there may be conflicts among these SAV actions. Section 7.7 of [RFC 8955] describes the conflicts among Traffic filtering actions.

5. Example of Validation Modes and SAV rule NLRI Configuration

In this section, we provide examples of how to configure SAV rule NLRI for the four validation modes. The SAV rule NLRI can carry zero, one, or multiple interfaces/interface groups and one or more prefixes.

5.1. Mode 1: Interface-based prefix allowlist

In this mode, only the prefixes carried in the SAV rule NLRI are considered valid on the specified interface. All other prefixes arriving at this interface are considered invalid.

Example:

* SAV rule NLRI: prefix = 192.168.1.0/24, interfaces = [Interface A]

* Validation: Any packet with a source prefix of 192.168.1.0/24 arriving at Interface A is valid. All other prefixes on Interface A are invalid.

5.2. Mode 2: Interface-based prefix blocklist

In this mode, the prefixes carried in the SAV rule NLRI are considered invalid on the specified interface. All other prefixes arriving at this interface are considered valid.

Example:

* SAV rule NLRI: prefix = 10.0.0.0/8, interfaces = [Interface B]

* Validation: Any packet with a source prefix of 10.0.0.0/8 arriving at Interface B is invalid. All other prefixes on Interface B are valid.

5.3. Mode 3: Prefix-based interface allowlist

In this mode, for a specific source prefix, only the interfaces carried in the SAV rule NLRI are considered valid. Packets with this source prefix arriving at other interfaces are invalid. Other prefixes are not checked.

Example:

* SAV rule NLRI: prefix = 172.16.0.0/16, interfaces = [Interface C, Interface D]

* Validation: Packets with a source prefix of 172.16.0.0/16 are valid only if they arrive at Interface C or Interface D. This prefix arriving at other interfaces is invalid. Other prefixes are not checked.

5.4. Mode 4: Prefix-based interface blocklist

In this mode, for a specific source prefix, the interfaces carried in the SAV rule NLRI are considered invalid. Packets with this source prefix arriving at other interfaces are valid. Other prefixes are not checked.

Example:

* SAV rule NLRI: prefix = 192.168.2.0/24, interfaces = [Interface E]

* Validation: Packets with a source prefix of 192.168.2.0/24 arriving at Interface E are invalid. This prefix arriving at other interfaces is valid. Other prefixes are not checked.

6. Procedures

SAV rules only exist on the routers running SAV mechanisms/protocols and the controller, these routers are usually access routers or boundary routers. This document describes extensions to the BGP-LS NLRI. The Routers running SAV mechanisms/protocols establish BGP-LS sessions with the controller respectively to report multi-sourced SAV rules.

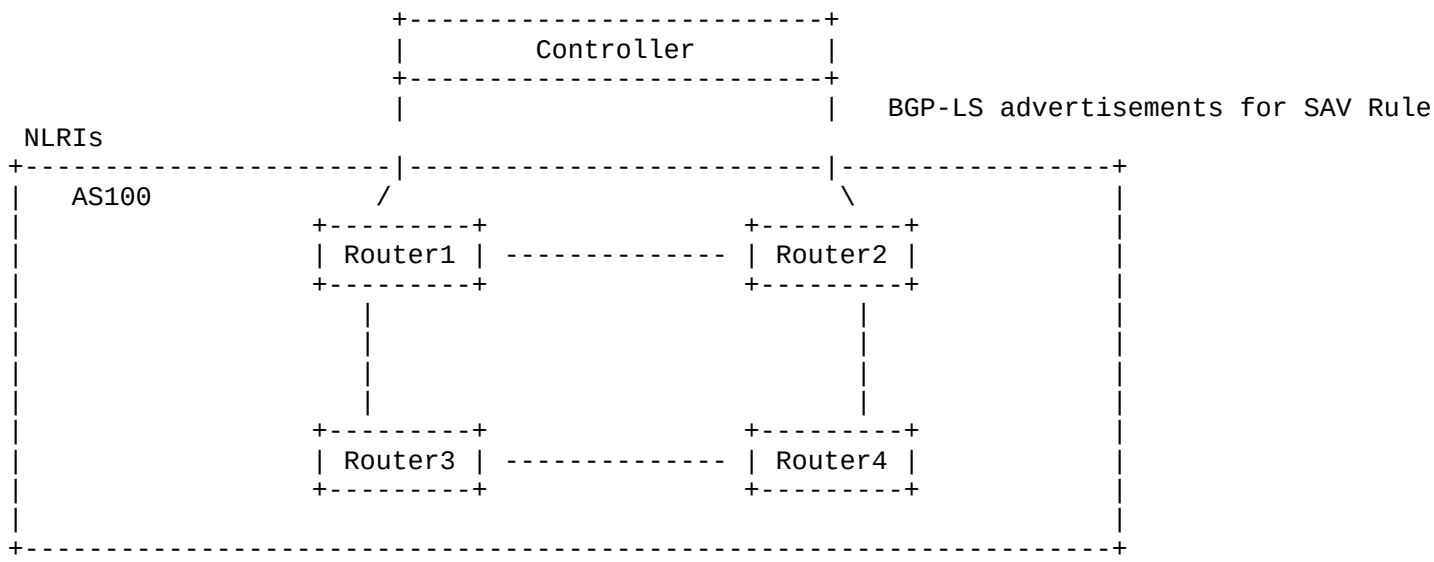

Figure 8: Advertisement of SAV Rules using BGP-LS

Based on Figure 8, the process of reporting SAV rules via BGP-LS is described as follows: Step 1: R1 and R2 run SAV mechanism/protocol, and generate multi-sourced SAV rules. Step 2: R1 and R2 respectively establish BGP-LS sessions with the controller. Step 3: R1 and R2 generate BGP-LS advertisements for the SAV Rule NLRIs. Step 4: R1 and R2 report multi-sourced SAV rules to the controller through the sav rule NLRIs (as defined in Section 2). This enables the controller to monitor and manage multi-sourced SAV rules.

7. Manageability Considerations

The Existing BGP operational and management procedures apply to this document. No new procedures are defined in this document. The considerations as specified in [RFC9552] apply to this document.

8. IANA Considerations

This section describes the code point allocation by IANA for this document.

8.1. "BGP-LS NLRI-Types" registry

This document requests assigning code-points from the registry for SAV Rule NLRIs:

Type	NLRI Type
TBD	IPv4 SAV Rule NLRI
TBD	IPv6 SAV Rule NLRI

8.2. "BGP-LS SAV Rule Descriptors TLVs" registry

This document requests assigning code-points from the registry for BGP-LS SAV Rule Descriptors TLVs based on Figure 3.

8.3. "BGP-LS SAV Mode Attribute TLV" registry

This document requests assigning a code-point from the registry for the BGP-LS SAV Mode attribute TLV.

9. Security Considerations

Procedures and protocol extensions defined in this document do not affect the base BGP security model. See [RFC6952] for details. The security considerations of the base BGP-LS specification as described in [RFC9552] also apply.

10. References

10.1. Normative References

[RFC9552] Talaulikar, K., Ed., "Distribution of Link-State and Traffic Engineering Information Using BGP", RFC 9552, DOI 10.17487/RFC9552, December 2023, <https://www.rfc-editor.org/info/rfc9552>.

[I-D.ietf-savnet-general-sav-capabilities] Huang, M., Cheng, W., Li, D., Geng, N., and L. Chen, "General Source Address Validation Capabilities", Work in Progress, Internet-Draft, draft-ietf-savnet-general-sav-capabilities-01, 24 June 2025, <https://datatracker.ietf.org/doc/html/draft-ietf-savnet-general-sav-capabilities-01>.

[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>.

[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>.

10.2. Informative References

[RFC6952] Jethanandani, M., Patel, K., and L. Zheng, "Analysis of BGP, LDP, PCEP, and MSDP Issues According to the Keying and Authentication for Routing Protocols (KARP) Design Guide", RFC 6952, DOI 10.17487/RFC6952, May 2013, <https://www.rfc-editor.org/info/rfc6952>.

[RFC8955] Loibl, C., Hares, S., Raszuk, R., McPherson, D., and M. Bacher, "Dissemination of Flow Specification Rules", RFC 8955, DOI 10.17487/RFC8955, December 2020, <https://www.rfc-editor.org/info/rfc8955>.

[RFC8956] Loibl, C., Ed., Raszuk, R., Ed., and S. Hares, Ed., "Dissemination of Flow Specification Rules for IPv6", RFC 8956, DOI 10.17487/RFC8956, December 2020, <https://www.rfc-editor.org/info/rfc8956>.

[RFC9086] Previdi, S., Talaulikar, K., Ed., Filsfils, C., Patel, K., Ray, S., and J. Dong, "Border Gateway Protocol - Link State (BGP-LS) Extensions for Segment Routing BGP Egress Peer Engineering", RFC 9086, DOI 10.17487/RFC9086, August 2021, <https://www.rfc-editor.org/info/rfc9086>.

[I-D.ietf-savnet-intra-domain-architecture] Li, D., Wu, J., Qin, L., Geng, N., and L. Chen, "Intra-domain Source Address Validation (SAVNET) Architecture", Work in Progress, Internet-Draft, draft-ietf-savnet-intra-domain-architecture-02, 13 April 2025, <https://datatracker.ietf.org/doc/html/draft-ietf-savnet-intra-domain-architecture-02>.

[I-D.ietf-savnet-inter-domain-architecture] Li, D., Chen, L., Geng, N., Liu, L., and L. Qin, "Inter-domain Source Address Validation (SAVNET) Architecture", Work in Progress, Internet-Draft, draft-ietf-savnet-inter-domain-architecture-01, 3 March 2025, <https://datatracker.ietf.org/doc/html/draft-ietf-savnet-inter-domain-architecture-01>.

[I-D.ietf-idr-flowspec-interfaceset] Litkowski, S., Simpson, A., Patel, K., Haas, J., and L. Yong, "Applying BGP flowspec rules on a specific interface set", Work in Progress, Internet-Draft, draft-ietf-idr-

flowspec-interfaceset-05, 18 November 2019, <https://datatracker.ietf.org/doc/html/draft-ietf-idr-flowspec-interfaceset-05>.

[I-D.geng-idr-flowspec-sav]
Geng, N., Li, D., tongtian124, and M. Huang, "BGP Flow Specification for Source Address Validation", Work in Progress, Internet-Draft, draft-geng-idr-flowspec-sav-05, 14 April 2025, <https://datatracker.ietf.org/doc/html/draft-geng-idr-flowspec-sav-05>.

Authors' Addresses

Tian Tong
China Unicom
Beijing
China
Email: tongt5@chinaunicom.cn

Dan Li
Tsinghua University
Beijing
China
Email: tolidan@tsinghua.edu.cn

Nan Geng
Huawei
Beijing
China
Email: gengnan@huawei.com

Nan Wang
China Unicom
Beijing
China
Email: wangn161@chinaunicom.cn

Shunwan Zhuang
Huawei
Beijing
China
Email: zhuangshunwan@huawei.com

Jing Zhao
China Unicom
Beijing
China
Email: zhaoj501@chinaunicom.cn