

BESS
Internet-Draft
Intended status: Standards Track
Expires: 28 May 2026

T. Murakami
K. Patel
Arrcus, Inc.
S. Matsushima
SoftBank
J. Zhang
Juniper Networks
S. Agrawal
D. Voyer
Cisco Systems
24 November 2025

BGP Extensions for the Mobile User Plane (MUP) SAFI
draft-ietf-bess-mup-safi-00

Abstract

This document defines a new SAFI known as a BGP Mobile User Plane (BGP-MUP) SAFI to support MUP Extensions and a extended community for BGP. This document also provides BGP signaling and procedures for the new SAFI to convert mobile session information into appropriate IP forwarding information. These extensions can be used by operators between a PE, and a Controller for integrating mobile user plane into BGP MUP network using the IP based routing.

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 28 May 2026.

Copyright Notice

Copyright (c) 2025 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 Revised BSD License text as described in Section 4.e of the Trust Legal Provisions and are provided without warranty as described in the Revised BSD License.

Table of Contents

1. Introduction	3
1.1. Requirements Notation	3
2. Terminology	3
3. BGP MUP Extensions	4
3.1. BGP MUP SAFI	4
3.1.1. BGP Interwork Segment Discovery route	5
3.1.2. BGP Direct Segment Discovery route	6
3.1.3. BGP Type 1 Session Transformed (ST) Route	6
3.1.4. BGP Type 2 Session Transformed (ST) Route	8
3.2. BGP MUP Extended Community	10
3.3. Operation	10
3.3.1. Generation of the Interwork Segment Discovery route	11
3.3.2. Withdrawal of the Interwork Segment Discovery route	11
3.3.3. Processing of the Interwork Segment Discovery routes	11
3.3.4. Generation of the Direct Segment Discovery route	12
3.3.5. Withdrawal of the Direct Segment Discovery route	13
3.3.6. Processing of the Direct Segment Discovery routes	13
3.3.7. Generation of the Type 1 ST Route	14
3.3.8. Withdrawing of the Type 1 ST Route	14
3.3.9. Processing of the Type 1 ST Route	15
3.3.10. Generation of the Type 2 ST Route	15
3.3.11. Withdrawing of the Type 2 ST Route	16
3.3.12. Processing of the Type 2 ST Route	17
4. Security Considerations	17
5. IANA Considerations	18
6. Contributors	18
7. References	19
7.1. Normative References	19
7.2. Informative References	20
Authors' Addresses	21

1. Introduction

[I-D.mhkk-dmm-srv6mup-architecture] defines the Segment Routing IPv6 Mobile User Plane (SRv6 MUP) architecture for Distributed Mobility Management. As part of the architecture, the document defines a new SRv6 segment type called as a MUP Segment, new routing information that can be carried within BGP, and advertised from a PE and a MUP Controller.

This document defines a new SAFI known as a BGP Mobile User Plane (BGP-MUP) SAFI to support MUP Extensions for BGP. This draft also provides BGP signaling and procedures for the new SAFI to convert mobile session information into appropriate IP routing information. These extensions can be used by operators between the PE and a MUP Controller for integrating mobile user plane into Segment Routing network using the IP based routing information. These extensions also work with routing instances accommodating two new well-known segment types known as Interwork and Direct [I-D.mhkk-dmm-srv6mup-architecture]. Finally, the BGP encoding and procedures defined in this document that use SRv6 as the forwarding fabric follow the SRv6 MUP architecture defined in [I-D.mhkk-dmm-srv6mup-architecture]. The BGP extensions to build networks that use forwarding mechanisms other than SRv6 (SRv6 MUP) are outside the scope of this document.

1.1. Requirements Notation

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. Terminology

MUP: Mobile User Plane

MUP Segment: Representation of mobile user plane segment

PE: Provider Edge node in an SR network

MUP Controller: Controller node for an SR network

UE: User Equipment, as per [TS.23501]

gNodeB: 3GPP-compliant implementation of 5G-NR base station, as per [TS.23501]

UPF: 3GPP-compliant implementation of User Plane Function, as per [TS.23501]

TEID: Tunnel Endpoint Identifier, as per [TS.29281]

3. BGP MUP Extensions

3.1. BGP MUP SAFI

This draft defines a new BGP SAFI known as a BGP-MUP SAFI. The value of this SAFI is to be assigned by IANA from the Subsequent Address Family Identifiers (SAFI) registry.

This document also defines a new BGP NLRI known as the BGP-MUP NLRI. The following is the format of the BGP-MUP NLRI:

```

+-----+
| Architecture Type (1 octet) |
+-----+
| Route Type (2 octets) |
+-----+
| Length (1 octet) |
+-----+
| Route Type specific (variable) |
+-----+

```

The Architecture Type field defines the encoding of the rest of BGP-MUP NLRI for a given Mobile User Plane architecture. This draft defines the following architecture type for BGP-MUP NLRI:

```
+ 1 - 3gpp-5g;
```

The Route Type field defines the encoding of the rest of BGP-MUP NLRI (Route Type specific BGP-MUP NLRI) for a given architecture type.

The Length field indicates the length in octets of the Route Type specific field of the BGP-MUP NLRI.

This document defines the following Route Types for BGP-MUP NLRI.

```

+ 1 - Interwork Segment Discovery route;
+ 2 - Direct Segment Discovery route;
+ 3 - Type 1 Session Transformed (ST) route;
+ 4 - Type 2 Session Transformed (ST) route;

```

These Route Types are applicable for the 3gpp-5G architecture type as per [I-D.mhkk-dmm-srv6mup-architecture]. Other new architectures can share them if it is applicable as well.

The BGP-MUP NLRI is carried in BGP [RFC4271] using BGP Multiprotocol Extensions [RFC4760] with an Address Family Identifier (AFI) of 1 or 2 and a Subsequent AFI (SAFI) of BGP-MUP. The NLRI field in the MP_REACH_NLRI/MP_UNREACH_NLRI attribute contains the BGP-MUP NLRI (encoded as specified above). The value of the AFI field in the MP_REACH_NLRI/MP_UNREACH_NLRI attribute that carries the BGP-MUP NLRI determines whether the addresses carried in the routes are IPv4 or IPv6 addresses (AFI 1 indicates IPv4 addresses, AFI 2 indicates IPv6 addresses).

In order for two BGP speakers to exchange BGP-MUP NLRI, they must use a BGP Capabilities Advertisement to ensure that they both are capable of properly processing such an NLRI. This is done as specified in [RFC4760], by using capability code 1 (multiprotocol BGP) with an AFI of 1 or 2 and an SAFI of BGP-MUP.

This document defines 4 Route Types for 3gpp-5G architecture type. Any other Route Types MUST be silently ignored upon a receipt if a BGP speaker supports only 3gpp-5G architecture type. An implementation MAY log an error when such Route Types are ignored. An implementation MAY consider retrieving any discarded Route Types by simply resetting the session or issuing a Route-REFRESH message [RFC2918] if the Route Refresh Capability is successfully negotiated.

The following sections describes the format of the Route Type specific BGP-MUP NLRI for various Route Types defined in this document.

3.1.1. BGP Interwork Segment Discovery route

A BGP Interwork Segment Discovery route Type specific BGP-MUP NLRI consist of the following:

```
+-----+
|          RD   (8 octets)          |
+-----+
|   Prefix Length (1 octet)   |
+-----+
|      Prefix (variable)      |
+-----+
```

The Interwork Segment Discovery route Type NLRI consist of RD which is encoded as described in [RFC4364]. It also has an prefix associated with Interwork segment connected locally. For the purpose of BGP route key processing, only the RD, Prefix Length and Prefix are considered to be part of the prefix in the NLRI.

In 3GPP 5G specific case, a prefix used for a N3 interface of the gNodeB connected locally MAY be used as this prefix. The prefix length of one octet indicating length of the prefix. If the AFI is IPv4, then the maximum value of the Prefix Length is 32 bits otherwise it is considered as a malformed NLRI. If the AFI is IPv6, then the maximum value of of the Prefix length is 128 bits otherwise it is considered as a malformed NLRI. A BGP speaker MUST handle such a malformed NLRI as a "Treat-as-withdraw" [RFC7606]. A BGP speaker MUST skip such NLRIs and continue processing of rest of the Update message.

3.1.2. BGP Direct Segment Discovery route

A BGP Direct Segment Discovery route Type specific BGP-MUP NLRI consist of the following:

```

+-----+
|          RD   (8 octets)          |
+-----+
|          Address (4 or 16 octets)  |
+-----+
```

The Direct Segment Discovery route Type NLRI consist of RD which is encoded as described in [RFC4364]. It also has an Address of originating BGP speaker. For the purpose of BGP route key processing, only the RD and Address are considered to be part of the prefix in the NLRI.

If the AFI is IPv4 then the address length is 4 octets otherwise it is considered as a malformed NLRI. If the AFI is IPv6 then the address length is 16 octets otherwise it is considered as a malformed NLRI. A BGP speaker MUST handle such a malformed NLRI as a "Treat-as-withdraw" [RFC7606]. A BGP speaker MUST skip such NLRIs and continue processing of rest of the Update message.

3.1.3. BGP Type 1 Session Transformed (ST) Route

A BGP Type 1 ST Route Type specific BGP-MUP NLRI consist of the following:

	RD (8 octets)	
	Prefix Length (1 octet)	
	Prefix (variable)	
	Architecture specific (variable)	

The Type 1 ST Route Type NLRI consist of RD which is encoded as described in [RFC4364]. It also has Prefix length of one octet indicating length of the Prefix. For the purpose of BGP route key processing, only the RD, Prefix Length and Prefix are considered to be part of the prefix in the NLRI.

In 3GPP 5G specific case, Prefix is the prefix allocated to a UE. If the AFI is IPv4, then the maximum value of the Prefix Length field is 32. If the AFI is IPv6, then the maximum value of the Prefix Length field is 128. Any other length field is considered a malformed NLRI. A BGP speaker MUST handle such a malformed NLRI as a "Treat-as-withdraw" [RFC7606]. A BGP speaker MUST skip such NLRIs and continue processing of rest of the Update message.

The architecture specific encoding values will follow after the variable length Prefix.

3.1.3.1. 3gpp-5g Specific BGP Type 1 ST Route

A BGP 3gpp-5g Type 1 ST Route Type specific BGP-MUP NLRI consist of the following:

	TEID (4 octets)	
	QFI (1 octet)	
	Endpoint Address Length (1 octet)	
	Endpoint Address (variable)	
	Source Address Length (1 octet)	
	Source Address (variable)	

The TEID has a fixed length of 4 octets. The TEID value of 0 is considered as an invalid and a malformed TEID. A BGP speaker MUST handle such a malformed NLRI as a "Treat-as-withdraw" [RFC7606]. A BGP speaker MUST skip such NLRIs and continue processing of rest of the Update message.

The QFI has a fixed length of 1 octet.

The Endpoint Address Length has a fixed length of 1 octet. Endpoint Address field contains of an IPv4 address, then the value of the Endpoint Address Length field is 32. If the Endpoint Address field contains of an IPv6 Address, then the value of the Endpoint Address Length field is 128. Any other value is considered as an invalid and a malformed Endpoint Address. A BGP speaker MUST handle such a malformed NLRI as a "Treat-as-withdraw" [RFC7606]. A BGP speaker MUST skip such NLRIs and continue processing of rest of the Update message.

The Source Address Length has a fixed length of 1 octet. Source Address Length is 0 bytes then the Source address is not carried within the NLRI. A BGP speaker MAY have a local configuration for using a Source address. The exact mechanism of local configuration is outside the scope of this document. Source Address field contains of an IPv4 address, then the value of the Source Address Length field is 32. If the Source Address field contains of an IPv6 Address, then the value of the Source Address Length field is 128. Any other value is considered as an invalid and a malformed Endpoint Address. A BGP speaker MUST handle such a malformed NLRI as a "Treat-as-withdraw" [RFC7606]. A BGP speaker MUST skip such NLRIs and continue processing of rest of the Update message.

The NLRI architecture field MUST be encoded as shown above if a BGP speaker receives 3gpp-5g specific BGP Type 1 ST route. Otherwise the NLRI is considered as a malformed. A BGP speaker MUST handle such a malformed NLRI as a "Treat-as-withdraw" [RFC7606]. A BGP speaker MUST skip such NLRIs and continue processing of rest of the Update message.

3.1.4. BGP Type 2 Session Transformed (ST) Route

A BGP Type 2 ST Route Type specific BGP-MUP NLRI consist of the following:

	RD (8 octets)	
	Endpoint Length (1 octet)	
	Endpoint Address (variable)	
	Architecture specific Endpoint Identifier (variable)	

The Type 2 ST Route Type NLRI consist of RD which is encoded as described in [RFC4364]. It also has Endpoint length of one octet indicating length of the Endpoint Address and the Architecture specific Endpoint Identifier as per

[I-D.mhkk-dmm-srv6mup-architecture] defines aggregation capability by the Type2 ST Route. If the AFI is IPv4 and the Endpoint Length is longer than 32 then the Architecture specific endpoint Identifier field exists with the IPv4 Endpoint Address. If the AFI is IPv6 and the Endpoint Length is longer than 128 then the Architecture specific endpoint Identifier field exists with then the IPv6 Endpoint Address. For the purpose of BGP route key processing, only the RD, Endpoint Address and Architecture specific Endpoint Identifier are considered to be part of the prefix in the NLRI.

In 3GPP 5G specific case, the Endpoint Address is a N3 interface address of the UPF. If the AFI is IPv4, then the maximum Endpoint length is 64 otherwise it is considered as a malformed NLRI. If the AFI is IPv6, then the maximum Endpoint length is 160 otherwise it is considered as a malformed NLRI. A BGP speaker MUST handle such a malformed NLRI as a "Treat-as-withdraw" [RFC7606]. A BGP speaker MUST skip such NLRIs and continue processing of rest of the Update message.

The architecture specific encoding values will follow after the variable length Prefix.

3.1.4.1. 3gpp-5g Specific BGP Type 2 ST Route

A BGP 3gpp-5g Type 2 ST Route Type specific BGP-MUP NLRI consist of the following:

	TEID (0-4 octets)	
--	-------------------	--

The maximum length of TEID is 4 octets. The TEID value of 0 is considered as an invalid and a malformed TEID. A BGP speaker MUST handle such a malformed NLRI as a "Treat-as-withdraw" [RFC7606]. A BGP speaker MUST skip such NLRIs and continue processing of rest of the Update message.

The NLRI architecture field MUST be encoded as shown above if a BGP speaker receives 3gpp-5g specific BGP Type 2 ST route. Otherwise the NLRI is considered as a malformed. A BGP speaker MUST handle such a malformed NLRI as a "Treat-as-withdraw" [RFC7606]. A BGP speaker MUST skip such NLRIs and continue processing of rest of the Update message.

3.2. BGP MUP Extended Community

This document defines a new BGP Extended community known as BGP MUP Extended community as per [I-D.mhkk-dmm-srv6mup-architecture].

This is a BGP MUP specific Extended Community, of an extended type, and is transitive across AS boundaries [RFC4360].

The Type value of this Extended community is set to MUP type, to be assigned by IANA from BGP Extended community transitive registry. The Sub-Type value is set to Direct-Type Segment Identifier type, to be assigned by IANA from BGP Extended community transitive registry. The value field of the community is set to the 6 bytes of configurable segment identifier value.

The usage of this Extended community is described in the Section 3.3.12 and Section 3.3.10

3.3. Operation

BGP speakers acting as a PE, and a MUP Controller MUST establish a BGP session to exchange BGP-MUP NLRIs for both, IPv4 and IPv6 AFIs. Once the session is established successfully, BGP speakers can exchange the Discovery routes as well as Session Transformed routes. This information is specific to a given routing instance. BGP-MUP SAFI is expected to work with routing instances accomodationg MUP segments. In 3GPP 5G specific case, the routing instances are depicted as N3RAN and N6DN instances defined in [I-D.mhkk-dmm-srv6mup-architecture]. The subsequent sections describes procedures of generating and processing of each route types.

3.3.1. Generation of the Interwork Segment Discovery route

The Interwork Segment Discovery route is generated by the PE when a routing instance accommodates an Interwork type MUP Segment, e.g., N3 interfaces or routes on RAN side in 3GPP 5G specific case. It generates route per each N3RAN IP prefix and stores the route in the routing instance of N3RAN. The IP prefix MAY include a gNodeB address which is connecting to the PE. The BGP AFI for BGP MP_REACH_NLRI attribute to carry the Discovery route is decided based on the AFI of the prefix.

When advertising the Interwork Segment Discovery route, a PE MUST attach the export BGP Route Target Extended Community of the associated routing instance.

When advertising the Interwork Segment Discovery route, a PE MUST use the IPv6 address of the PE as the nexthop address in the MP_REACH_NLRI attribute.

The Interwork Segment Discovery route update MUST have a prefix SID attribute which the SID consists of the PE locator followed by a function. In 3GPP 5G specific case, if the BGP AFI is IPv4, the function MUST be GTP4.E [I-D.ietf-dmm-srv6-mobile-uplane], or MUST be GTP6.E [I-D.ietf-dmm-srv6-mobile-uplane] if the BGP AFI is IPv6.

3.3.2. Withdrawal of the Interwork Segment Discovery route

The Interwork Segment Discovery route is withdrawn by the PE when it detects that the MUP Segment no longer present for the N3RAN. The BGP AFI for BGP MP_UNREACH_NLRI attribute to carry the Discovery route is decided based on the AFI of the prefix.

When withdrawing the Interwork Segment Discovery route, a PE MUST attach the export BGP Route Target Extended Community of the associated routing instance.

3.3.3. Processing of the Interwork Segment Discovery routes

A BGP speaker acting as a PE MAY keep the received MUP Interwork Segment Discovery routes advertised from other PEs. The receiving BGP speaker will perform the importing of the received MUP Interwork Segment Discovery routes in the configured routing instance based on the Route Target extended communities. The IP prefixes for the received segments are imported into the configured routing instance table. Thereby the receiving BGP speaker can provide network connectivity between the nodes that exist in the segments. A BGP speaker MAY discard the received Interwork Segment Discovery route if the speaker fails to import the route based on the Route Target

extended communities.

The BGP speaker receiving the Interwork Segment Discovery routes SHOULD ignore the nexthop in the MP_REACH_NLRI attribute. However, the receiving BGP speaker MUST ensure that the value of Address filed in the NLRI is an address of the originator of the locator value in the prefix SID attribute. The originator of the locator value can be resolved from the IPv6 IGP table. If the result of the match is not identical then the receiving BGP speaker MUST consider it as a malformed NLRI and the "Treat-as-withdraw procedure of [RFC7606] is applied. A BGP speaker MUST skip such NLRIs and continue processing of rest of the Update message.

When a BGP speaker receives the Interwork Segment Discovery routes with a MP_REACH_NLRI attribute without a prefix SID attribute, then it MUST be treated as if it contained a malformed prefix SID attribute and the "Treat-as-withdraw procedure of [RFC7606] is applied. A BGP speaker MUST skip such NLRIs and continue processing of rest of the Update message.

When a BGP speaker receives an MP_UNREACH_NLRI attribute update message it MUST delete the withdrawn Interwork Segment Discovery route from the routing instance table where it was created.

3.3.4. Generation of the Direct Segment Discovery route

The Direct Segment Discovery route is generated by the PE when a routing instance accommodates a Direct type MUP Segment, e.g., N6 interface or routes on DN side in 3GPP 5G specific case. It generates the Direct Segment Discovery route per each routing instance for the MUP Segment. The address in the BGP-MUP NLRI MUST be a unique PE identifier. The BGP AFI for BGP MP_REACH_NLRI attribute to carry the Direct Segment Discovery route is decided based on the AFI of the PE identifier [I-D.mhkk-dmm-srv6mup-architecture].

When announcing the Direct Segment Discovery route, a PE MUST attach a BGP MUP Extended community of the associated routing instance. The sub-type of the Extended community is Direct-Type Segment Identifier.

When advertising the Direct Segment Discovery route, a PE MUST use the IPv6 address of the PE as the nexthop address in the MP_REACH_NLRI attribute.

The Direct Segment Discovery route update MUST have a prefix SID attribute which the SID consists of the PE locator followed by a function. The function MAY be End.DT4/6 or End.DX4/6.

3.3.5. Withdrawal of the Direct Segment Discovery route

The Direct Segment Discovery route is withdrawn by the PE when it detects that the MUP Segment no longer present for the routing instance. The BGP AFI for BGP MP_UNREACH_NLRI attribute to carry the Discovery route is decided based on the AFI of the PE identifier.

When withdrawing the Direct Segment Discovery route, a BGP speaker MUST attach a BGP MUP Extended community of the associated routing instance.

3.3.6. Processing of the Direct Segment Discovery routes

A BGP speaker acting as a PE MAY keep the received MUP Direct Segment Discovery routes advertised from other PEs. The receiving BGP speaker will perform the importing of the received MUP Direct Segment Discovery routes in the configured routing instance based on the Route Target extended communities. The IP prefixes for the received segments are imported into the configured routing instance table. Thereby the receiving BGP speaker can provide network connectivity between the nodes that exist in the segments. A BGP speaker MAY discard the received Direct Segment Discovery route if the speaker fails to import the route based on the Route Target extended communities.

The BGP speaker receiving the Direct Segment Discovery routes SHOULD ignore the nexthop in the MP_REACH_NLRI attribute. However, the receiving BGP speaker MUST ensure that the received nexthop value in the MP_REACH_NLRI attribute is identical to the originator of the locator value in the prefix SID attribute. The originator of the locator value can be resolved from the IPv6 IGP table. If the result of the match is not identical then the receiving BGP speaker MUST consider it as a malformed NLRI and the "Treat-as-withdraw procedure of [RFC7606] is applied. A BGP speaker MUST skip such NLRIs and continue processing of rest of the Update message.

When a BGP speaker receives a MP_REACH_NLRI attribute update message with a Direct Segment Discovery route without a prefix SID attribute, than it MUST be treated as if it contained a malformed prefix SID attribute and the "Treat-as-withdraw procedure of [RFC7606] is applied. A BGP speaker MUST skip such NLRIs and continue processing of rest of the Update message.

When a BGP speaker receives an MP_UNREACH_NLRI attribute update message it MUST delete the withdrawn Direct Segment Discovery route from the routing instance table where it was created.

3.3.7. Generation of the Type 1 ST Route

A BGP speaker acting as a MUP Controller generates Type 1 ST route from corresponding session information through a northbound API as per [I-D.mhkk-dmm-srv6mup-architecture]. The northbound API definition for ST1 route creation is out of scope of this document.

In 3GPP 5G specific case, to compose a Type 1 ST route the controller acquires UE address or prefix, Tunnel Endpoint address of GTP [TS.29281] tunnel, TEID, and QFI for access side as input parameters from the northbound API. Controller also acquires optionally the associated Source Address. The controller decides the RD of the Type 1 ST route based on operator policy.

When advertising the Type 1 ST route, the controller SHOULD attach a Route Target Extended community which the PEs are importing into the routing instance for the corresponding Direct segment.

The MUP Controller MUST set the nexthop of the route to the address of the controller.

The controller MUST announce this route using a AFI of the route and the SAFI of BGP-MUP to all other BGP speakers within the SRv6 domain.

3.3.8. Withdrawing of the Type 1 ST Route

A BGP speaker acting as a MUP Controller withdraws Type 1 ST route based on deletion of corresponding session information through a northbound API as per [I-D.mhkk-dmm-srv6mup-architecture]. The northbound API definition for ST1 route withdraw is out of scope of this document.

In 3GPP 5G specific case, to withdraw a Type 1 ST route the controller acquires the UE address or prefix, Tunnel Endpoint address of GTP [TS.29281] tunnel, TEID, and QFI for access side as input parameters from the northbound API. Controller also acquires optionally Source Address. The controller MUST advertise the withdraws of the Type 1 ST route.

When withdrawing the Type 1 ST route, the controller SHOULD attach the Route Target Extended community which the PEs are importing into the routing instance accomodating the corresponding Direct segment to the Route Target Extended community.

3.3.9. Processing of the Type 1 ST Route

A BGP speaker acting as a PE MAY keep the received MUP Type 1 ST routes advertised from the MUP Controller. The receiving BGP speaker will perform the importing of the received MUP Type 1 ST routes in the configured N6DN routing instance based on the Route Target extended communities. A BGP speaker MAY discard the received Type 1 ST route if the speaker fails to import the route based on the Route Target extended communities.

In case of a BGP speaker receiving a Type 1 ST routes is a PE, the PE SHOULD use the received Tunnel Endpoint Address in this NLRI as a key to lookup the associated Interwork Segment Discovery route and extract the locator and the function in the prefix SID attribute of the Interwork route.

In 3GPP 5G specific case, the PE extracts TEID, QFI, Tunnel Endpoint address and the Source address if it is present from the NLRI. A PE MAY use a locally configured Source address in an event if the Source address is not present in Type 1 ST route. The Source address is used as the source address in the SRv6 Header. The PE SHOULD generate the forwarding SID for GTP4/6.E based on the procedures mentioned in the [I-D.ietf-dmm-srv6-mobile-uplane]. If the PE cannot generate the prefix SID, then it SHOULD mark the received Type 1 ST route as an invalid route. If the PE does not have a Source address for the given route, then the PE MAY hold such an invalid route until it is withdrawn. Similarly, the PE MAY hold such an invalid route until the route as a valid route upon successful generation of prefix SID.

The PE receiving Type 1 ST routes SHOULD ignore the received nexthop in the MP_REACH_NLRI attribute.

The PE receiving Type 1 ST routes in MP_UNREACH_NLRI attribute MUST delete all the routes from the associated routing instance. In an event where the received Type 1 ST route in MP_UNREACH_NLRI attribute does not have a Source address a receiving PE MUST delete all the matching Type 1 ST Routes with different Source addresses.

3.3.10. Generation of the Type 2 ST Route

A BGP speaker acting as a MUP Controller generates Type 2 ST route from corresponding session information through a northbound API, or pre-defined configuration as per [I-D.mhkk-dmm-srv6mup-architecture]. The northbound API definition for ST2 route creation is out of scope of this document.

In 3GPP 5G specific case, to compose a Type 2 ST route the controller acquires the Endpoint consists of Endpoint address of GTP [TS.29281] tunnel and TEID for core side with the effective length of the Endpoint as input parameters. The controller decides the RD of the Type 2 ST route based on operator policy.

When advertising the Type 2 ST route, the controller SHOULD attach a BGP MUP Extended community corresponding to the Direct segment. The sub-type of the Extended community is Direct-Type Segment Identifier. This Segment Identifier is generated from the information received through a northbound API, or a pre-defined configuration as per [I-D.mhkk-dmm-srv6mup-architecture]. The northbound API definition for receiving this information is out of scope of this document.

The controller MUST also attach a Route Target Extended community of the routing instances in the PE accomodating the corresponding Interwork segment.

The controller MUST set the nexthop of the route to the address of the MUP Controller.

3.3.11. Withdrawing of the Type 2 ST Route

A BGP speaker acting as a MUP Controller withdraws Type 2 ST route based on deletion of corresponding session information through a northbound API as per [I-D.mhkk-dmm-srv6mup-architecture]. The northbound API definition for ST2 route withdraw is out of scope of this document.

In 3GPP 5G specific case, to withdraw a Type 2 ST route the controller acquires the Endpoint consists of Endpoint address of GTP [TS.29281] tunnel and TEID for core side with the effective length of the Endpoint as input parameters. The controller MUST advertise the withdraws of the Type 2 ST route.

When withdrawing the Type 2 ST route, the controller SHOULD attach the BGP MUP Extended community corresponding to the Direct segment, and the Route Target Extended community which the PEs are importing into the routing instance accomodating the corresponding Interwork segment to the Route Target Extended community.

3.3.12. Processing of the Type 2 ST Route

A BGP speaker acting as a PE MAY keep the received MUP Type 2 ST routes advertised from the MUP Controller. The receiving BGP speaker will perform the importing of the received MUP Type 2 ST routes in the configured N3RAN routing instance based on the Route Target extended communities. A BGP speaker MAY discard the received Type 2 ST route if the speaker fails to import the route based on the Route Target extended communities.

The BGP speaker receiving the Type 2 ST routes SHOULD ignore the received nexthop in the MP_REACH_NLRI attribute.

A PE receiving the Type 2 ST routes in a MP_REACH_NLRI attribute without a BGP MUP Extended community SHOULD consider the route as a malformed route. The PE MUST handle such a malformed NLRI as a "Treat-as-withdraw" [RFC7606].

The PE receiving Type 2 ST route with a BGP MUP Extended Community should extract the received segment identifier from the community. The segment identifier is used to resolve an appropriate Direct segment routing instance.

4. Security Considerations

The mechanisms described in this document could reuse the existing BGP security mechanisms [RFC4271] [RFC4272]. The security model and threats specific to Provider Provisioned VPNs, including L3VPNs, are discussed in [RFC4111]. The method defined in [RFC5925] SHOULD be used where authentication of BGP control packets is needed.

The PEs and MUP Controller SHOULD NOT establish BGP sessions with other BGP speakers in the domains which are not trusted without any explicit configuration or an operator intervention. Usage of procedures defined in [RFC5925] SHOULD be enforced at such boundaries to ensure the proper authentication of BGP control packets.

Furthermore, [RFC5925] will not help to keep the BGP messages private. To protect the BGP messages exchanged between BGP speakers from eavesdrop, establishing BGP sessions over encrypted paths SHOULD be considered.

PEs SHOULD impose an upper bound on number of routes they should store to protect their control plane load. For example, BGP implementations MAY provide a configuration knob to impose an upper bound on Type 1 ST Routes.

5. IANA Considerations

This document defines a new BGP SAFI known as a BGP-MUP SAFI. IANA is requested to assign the value for the new SAFI from the Subsequent Address Family Identifiers (SAFI) registry.

This document defines a new Architecture Type for a BGP-MUP SAFI. IANA is requested to create a new Architecture Type NLRI registry for BGP-MUP SAFI. Furthermore, IANA is also requested to assign values for the following Architecture Types from the newly created BGP-MUP Architecture Type NLRI registry:

- + 1 - 3gpp-5g;

This document defines new NLRIs for a BGP-MUP SAFI. IANA is requested to create a new NLRI registry for BGP-MUP SAFI. Furthermore, IANA is also requested to assign values for the following NLRIs from the newly created BGP-MUP NLRI registry:

- + 1 - Discovery Internetwork route;
- + 2 - Direct Segment Discovery route;
- + 3 - Type 1 Session Transformed (ST) route;
- + 4 - Type 2 Session Transformed (ST) route;

This document defines a new BGP Extended Community called "SRv6 MUP Extended Community". This Community is of an extended type and is transitive. IANA is requested to assign the Type and the Sub-Type value for this Community from the Transitive Extended Community registry.

6. Contributors

In addition to the authors listed on the front page, the following individuals have also made significant contributions to the draft:

Katsuhiro Horiba
SoftBank
Email: katsuhiro.horiba@g.softbank.co.jp

Yuya Kawakami
SoftBank
Email: yuya.kawakami01@g.softbank.co.jp

Kalyani Rajaraman
Email: kalyanir@yahoo.com

7. References

7.1. Normative References

- [I-D.mhkk-dmm-srv6mup-architecture]
Matsushima, S., Horiba, K., Khan, A., Kawakami, Y., Murakami, T., Patel, K., Kohno, M., Kamata, T., Camarillo, P., Horn, J., Voyer, D., Zadok, S., Meilik, I., Agrawal, A., and K. Perumal, "Mobile User Plane Architecture using Segment Routing for Distributed Mobility Management", Work in Progress, Internet-Draft, draft-mhkk-dmm-srv6mup-architecture-06, 23 October 2023, <<https://datatracker.ietf.org/doc/html/draft-mhkk-dmm-srv6mup-architecture-06>>.
- [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>>.
- [RFC2918] Chen, E., "Route Refresh Capability for BGP-4", RFC 2918, DOI 10.17487/RFC2918, September 2000, <<https://www.rfc-editor.org/info/rfc2918>>.
- [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>>.
- [RFC4272] Murphy, S., "BGP Security Vulnerabilities Analysis", RFC 4272, DOI 10.17487/RFC4272, January 2006, <<https://www.rfc-editor.org/info/rfc4272>>.
- [RFC4360] Sangli, S., Tappan, D., and Y. Rekhter, "BGP Extended Communities Attribute", RFC 4360, DOI 10.17487/RFC4360, February 2006, <<https://www.rfc-editor.org/info/rfc4360>>.
- [RFC4760] Bates, T., Chandra, R., Katz, D., and Y. Rekhter, "Multiprotocol Extensions for BGP-4", RFC 4760, DOI 10.17487/RFC4760, January 2007, <<https://www.rfc-editor.org/info/rfc4760>>.

- [RFC7606] Chen, E., Ed., Scudder, J., Ed., Mohapatra, P., and K. Patel, "Revised Error Handling for BGP UPDATE Messages", RFC 7606, DOI 10.17487/RFC7606, August 2015, <<https://www.rfc-editor.org/info/rfc7606>>.
- [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>>.

7.2. Informative References

- [I-D.ietf-dmm-srv6-mobile-uplane] Matsushima, S., Filsfils, C., Kohno, M., Camarillo, P., and D. Voyer, "Segment Routing IPv6 for Mobile User Plane", Work in Progress, Internet-Draft, draft-ietf-dmm-srv6-mobile-uplane-24, 17 January 2023, <<https://datatracker.ietf.org/doc/html/draft-ietf-dmm-srv6-mobile-uplane-24>>.
- [RFC4111] Fang, L., Ed., "Security Framework for Provider-Provisioned Virtual Private Networks (PPVPNs)", RFC 4111, DOI 10.17487/RFC4111, July 2005, <<https://www.rfc-editor.org/info/rfc4111>>.
- [RFC4364] Rosen, E. and Y. Rekhter, "BGP/MPLS IP Virtual Private Networks (VPNs)", RFC 4364, DOI 10.17487/RFC4364, February 2006, <<https://www.rfc-editor.org/info/rfc4364>>.
- [RFC5925] Touch, J., Mankin, A., and R. Bonica, "The TCP Authentication Option", RFC 5925, DOI 10.17487/RFC5925, June 2010, <<https://www.rfc-editor.org/info/rfc5925>>.
- [RFC6811] Mohapatra, P., Scudder, J., Ward, D., Bush, R., and R. Austein, "BGP Prefix Origin Validation", RFC 6811, DOI 10.17487/RFC6811, January 2013, <<https://www.rfc-editor.org/info/rfc6811>>.
- [RFC8205] Lepinski, M., Ed. and K. Sriram, Ed., "BGPsec Protocol Specification", RFC 8205, DOI 10.17487/RFC8205, September 2017, <<https://www.rfc-editor.org/info/rfc8205>>.
- [RFC9012] Patel, K., Van de Velde, G., Sangli, S., and J. Scudder, "The BGP Tunnel Encapsulation Attribute", RFC 9012, DOI 10.17487/RFC9012, April 2021, <<https://www.rfc-editor.org/info/rfc9012>>.

[TS.23501] 3GPP, "System architecture for the 5G System (5GS)", 3GPP TS 23.501 17.2.0, 24 September 2021, <<http://www.3gpp.org/ftp/Specs/html-info/23501.htm>>.

[TS.29281] 3GPP, "General Packet Radio System (GPRS) Tunnelling Protocol User Plane (GTPv1-U)", 3GPP TS 29.281 16.1.0, September 2020.

Authors' Addresses

Tetsuya Murakami
Arrcus, Inc.
2077 Gateway Place, Suite 400
San Jose, CA 95110
United States of America
Email: tetsuya@arrcus.com

Keyur Patel
Arrcus, Inc.
2077 Gateway Place, Suite 400
San Jose, CA 95110
United States of America
Email: keyur@arrcus.com

Satoru Matsushima
SoftBank
Japan
Email: satoru.matsushima@g.softbank.co.jp

Jeffrey Zhang
Juniper Networks
United States of America
Email: zzhang@juniper.net

Swadesh Agrawal
Cisco Systems
United States of America
Email: swaagraw@cisco.com

Daniel Voyer
Cisco Systems
Montreal
Canada

Email: davoyer@cisco.com