

Internet Engineering Task Force
Internet-Draft
Intended status: Standards Track
Expires: 1 December 2025

S. Matsushima
SoftBank
K. Horiba

Y. Kawakami
SoftBank
T. Murakami
K. Patel
Arrcus, Inc
J. Horn
Cisco Systems, Inc.
30 May 2025

Mobile User Plane Architecture for Distributed Mobility Management
draft-ietf-dmm-mup-architecture-00

Abstract

This document defines the Mobile User Plane (MUP) architecture for Distributed Mobility Management. The requirements for Distributed Mobility Management described in [RFC7333] can be satisfied by routing fashion.

In MUP Architecture, session information between the entities of the mobile user plane is turned to routing information so that mobile user plane can be integrated into dataplane.

MUP architecture is designed to be pluggable user plane part of existing mobile service architectures, enabled by auto-discovery for the use plane. Segment Routing provides network programmability for a scalable option with it.

While MUP architecture itself is independent from a specific dataplane protocol, several dataplane options are available for the architecture. This document describes IPv6 dataplane in Segment Routing case (SRv6 MUP) due to the DMM requirement, and is suitable for mobile services which require a large IP address space.

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 1 December 2025.

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 Language	5
2. Terminology	5
3. Architecture Principles	6
4. Mobile User Plane Segment	7
4.1. Dataplane Considerations	8
5. Distribution of Mobile User Plane Segment Information	8
5.1. Direct Segment Discovery (DSD) Route	9
5.2. Interwork Segment Discovery (ISD) Route	10
6. Distribution of Session Transformed Route	10
6.1. Type 1 Session Transformed (ST1) Route	10
6.2. Type 2 Session Transformed (ST2) Route	11
6.3. MUP Controller	11
7. Illustration	11
7.1. SR Network Accommodating Existing Mobile Network Services	12
7.2. MUP PE Deployment at All SR Domain Edges	13
7.3. Adding Direct Segment with New MUP PE	15
7.4. Collapsed MUP PE Deployment	17
8. IANA Considerations	19
9. Security Considerations	19
10. References	19
10.1. Normative References	19
10.2. Informative References	20
Acknowledgements	21

Contributors	21
Additional Authors	21
Authors' Addresses	22

1. Introduction

Mobile service systems require IP connectivity for communication between the entities defined by mobile service architectures for the mobile service systems. [RFC5213][TS.23501].

In PMIPv6 [RFC5213], IP connectivity is required between LMA (Local Mobility Anchor) and MAG (Mobility Access Gateway), as well as LMA and Internet. In 3GPP 5G [TS.23501], IP connectivity for N3 interface between gNodeB(es) and UPFs (User Plane Function) is required, as well as for N6 interface between UPFs and DNS (Data Network).

These IP connectivities may be covered by multiple dataplane networks, such as IPv4, IPv6, MPLS, or bunch of dataplane protocols. When just one dataplane protocol network is adopted for simplicity, it is expected that the address space of the dataplane network should be large enough to cover a vast number of nodes, such as millions of base stations. For this reason, use of IPv6 dataplane looks sufficiently suitable.

IPv6 dataplane has been able to instantiate Segment Routing over IPv6 (SRv6) with network programming capability described in [RFC8986].

SRv6 network programmability enhances IPv6 dataplane to be integrated with mobile user plane [RFC9433]. It will make an entire IPv6 network support the user plane in a very efficient distributed routing fashion.

On the other hand, the requirements for Distributed Mobility Management (DMM) described in [RFC7333] can be satisfied by session management based solutions. [RFC8885] defines protocol extension to PMIPv6 for the DMM requirements. 3GPP 5G defines an architecture in which multiple session anchors can be added to one PDU session by the session management.

As a reminder, the user plane related requirements in [RFC7333] are reproduced here:

REQ1: Distributed mobility management

IP mobility, network access solutions, and forwarding solutions provided by DMM MUST enable traffic to avoid traversing a single mobility anchor far from the optimal route. It is noted that the requirement on distribution applies to the data plane only.

REQ3: IPv6 deployment

DMM solutions SHOULD target IPv6 as the primary deployment environment and SHOULD NOT be tailored specifically to support IPv4, particularly in situations where private IPv4 addresses and/or NATs are used.

REQ4: Existing mobility protocols

A DMM solution MUST first consider reusing and extending IETF standard protocols before specifying new protocols.

REQ5: Coexistence with deployed networks/hosts and operability across different networks

A DMM solution may require loose, tight, or no integration into existing mobility protocols and host IP stacks. Regardless of the integration level, DMM implementations MUST be able to coexist with existing network deployments, end hosts, and routers that may or may not implement existing mobility protocols. Furthermore, a DMM solution SHOULD work across different networks, possibly operated as separate administrative domains, when the needed mobility management signaling, forwarding, and network access are allowed by the trust relationship between them.

This document defines the Mobile User Plane (MUP) architecture for Distributed Mobility Management. MUP is not a mobility management system itself, but an architecture enables the dataplanes to integrate mobile user plane into it for the IP networks.

Although MUP architecture is independent from a specific dataplane protocol, this document describes IPv6 dataplane in Segment Routing case (SRv6 MUP) due to the DMM requirement, and as a suitable solution for scalable mobile service deployments. Other dataplane options is out of scope of this document, and may be described in the future.

In this routing paradigm, a session information from a mobile control plane of a mobile service system will be transformed to routing information. It means that any MUP dataplane nodes become functional to the session instead of the mobile user plane specific nodes for the role of anchor or intermediate points. The user plane anchor and intermediate functions can be supported by MUP enabled networks (REQ1), not to mention that MUP will naturally be deployed over IPv6 networks (REQ3).

MUP architecture is independent from the mobile service system. For the requirements (REQ4, 5), MUP architecture is designed to be pluggable user plane part of existing mobile service architectures. Those existing architectures are for example defined in [RFC5213], [TS.23501], or if any.

The level of MUP integration for mobile service systems running based on the existing mobile service architecture will be varied and depending on the level of MUP awareness of the control and user plane entities.

Specifying how to modify the existing architecture to integrate MUP is out of scope of this document. What this document provides for the existing architecture is an interface for MUP which the existing or future architectures can easily integrate.

1.1. Requirements Language

The key words "MUST", "MUST NOT", "REQUIRED", "SHALL", "SHALL NOT", "SHOULD", "SHOULD NOT", "RECOMMENDED", "MAY", and "OPTIONAL" in this document are to be interpreted as described in RFC 2119 [RFC2119].

2. Terminology

MUP: Mobile User Plane

MUP Segment: Representation of mobile user plane segment

MUP PE: MUP aware Provider Edge node

MUP Controller: Controller node for MUP networks

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

MN: Mobile Node, as per [RFC5213]

3. Architecture Principles

This section describes the principles for MUP architecture that guide its design and operation.

The first key principle is the abstraction of the mobile user plane. A network segment consisting of a mobile service is abstracted and represented as a MUP segment. It is noted that MUP segment described in this document is NOT Segment Routing[RFC8402] specific, as "segment" is widely common terminology through many networking technologies, and is defined in each technology context.

A MUP PE may accommodate MUP segment(s), such as an Interwork Segment and/or a Direct Segment described in Section 4. Figure 1 depicts the overview.

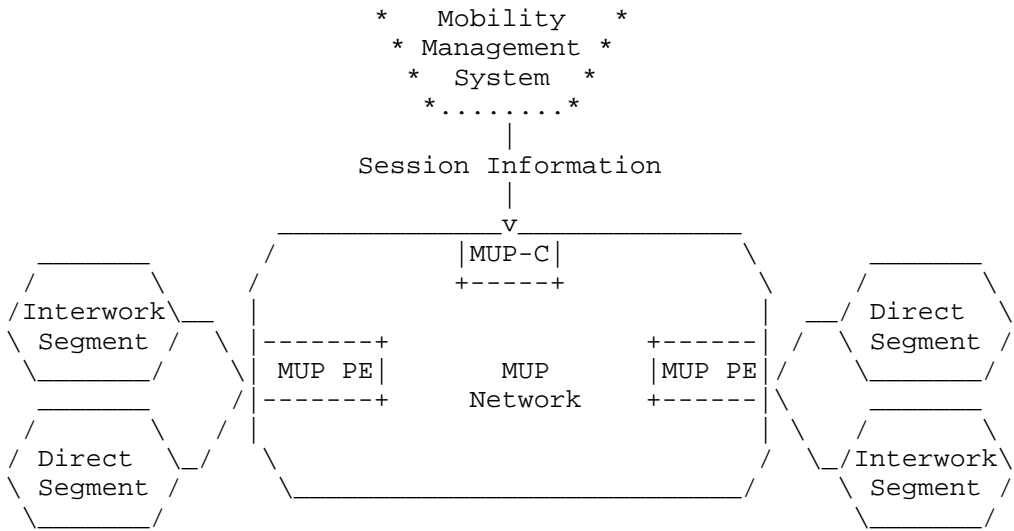


Figure 1: Overview of MUP Architecture

The second principle is auto-discovery for MUP segment. A MUP PE should be able to discover a MUP segment in a remote MUP PE. In MUP architecture, the remote MUP PE should advertise an auto-discovery route for a hosted MUP segment. The MUP PE can discover the MUP segment in the remote MUP PE when the MUP PE finds the MUP segment information in the received auto-discovery route from the remote MUP PE.

Section 5 in this document defines auto-discovery route for each type of MUP segment discovery.

It is noted that the auto-discovery route must be independent from a specific dataplane. But with an attribute for the specific dataplane, the auto-discovery route indicates specific dataplane behavior required for the MUP segment.

The third principle is that transforming session information to routing information. Assuming session information for a UE or a MN includes a pair of endpoints between the entities of the mobile user plane, a MUP Controller (MUP-C) advertises MUP PE(s) routing information for the UE or the MN, transformed from the input of corresponding session information in mobile service systems. In MUP architecture, it is called session-transformed route.

Section 6 in this document defines each type of session-transformed route. The session-transformed route must also be dataplane independent.

A MUP PE should resolve reachability for a received session-transformed route (ST route for short). When the MUP PE succeeds to resolve the ST route reachability with a MUP segment in local, or in remote via an auto-discovery route, and an appropriate dataplane behavior indicated in it for the ST route, the MUP PE can perform the dataplane action to the corresponding packet received from a local MUP segment.

The MUP PE sends the packet toward the egress MUP segment after the dataplane action applied to the packet. The egress MUP segment may exist in local, or in a remote MUP PE. In latter case, the remote MUP PE applies the dataplane action indicated by the received packet, and sends it out to the egress MUP segment.

The illustrations are described in Section 7.

To carry these new routing information, this architecture requires extending the existing routing protocols. Any routing protocol can be used to carry this information but this document recommends using BGP. Thus, this document describes extensions on BGP as an example.

4. Mobile User Plane Segment

This document defines two types of Mobile User Plane (MUP) segment. A MUP segment represents a network segment consisting of a mobile service. The MUP segment can be created by a MUP PE which provides connectivity for the mobile user plane.

Direct Segment is a type of MUP segment that provides connectivity between MUP segments through the MUP networks. Interwork Segment is another type of MUP segment. It provides connectivity between a user plane protocol of existing or future mobile service architecture and other MUP segments through the MUP networks.

A MUP PE may be instantiated as a physical node or a virtual node. The MUP PE may also be instantiated on a device which accommodates a mobile user plane node of a mobile service system.

4.1. Dataplane Considerations

As in Section 1, this document describes IPv6 dataplane in Segment Routing case due to the DMM requirement, and as a suitable solution for scalable mobile service deployments.

When SRv6 is adopted as the dataplane, an SRv6 SID (Segment Identifier) can represent a MUP segment. The SID can be any behavior defined in [RFC8986], [RFC9433], or any other extensions for further use cases. The behavior of the MUP segment will be chosen by the role of the representing MUP segment.

For example, in case of a MUP PE interfaces to 5G user plane on the access side defined as "N3" in [TS.23501], the MUP PE accommodates the N3 network as Interwork Segment in a routing instance and then the behavior of created segment SID by the MUP PE will be "End.M.GTP4.E", or "End.M.GTP6.E". In this case, the MUP PE may associate the SID to the routing instance for the N3 access network (N3RAN).

Another example here is that a MUP PE interfaces to 5G DN on the core side defined as "N6" in [TS.23501], the MUP PE accommodates the N6 network in a routing instance as Direct Segment and then the behavior of the created segment SID by the MUP PE will be "End.DT4", "End.DT6", or "End.DT2". In this case, the MUP PE may associate the SID to the routing instance for the N6 data network (N6DN).

5. Distribution of Mobile User Plane Segment Information

Distribution of MUP segment information can be done by advertising routing information with the MUP segment for mobile service. A MUP PE distributes MUP segment information when a MUP segment is connected to the MUP PE.

A MUP Segment Discovery route is routing information that associates the MUP segment with network reachability, and assistant metadata if applicable. This document defines the basic discovery route types, Direct Segment Discovery (DSD for short) route, and Interwork Segment

Discovery (ISD for short) route. Other types of segment discovery route may be mobile service architecture specific. Defining the architecture specific network reachability is out of scope of this document and it will be specified in another document.

To carry the assistant metadata for the MUP Segment Discovery route, this document defines MUP extended community. When a metadata applicable for a set of MUP Segment, a MUP extended community carries the metadata in the corresponding MUP Segment Discovery routes. The MUP extended community must be structured to indicate types of MUP segment specific metadata. Section 5.1 describes Direct Segment type MUP extended community, and is illustrated in Section 7.

MUP Segment Discovery route MUST be used only for resolving reachability for the ST routes. The connectivity among the routing instances for MUP Segments may be advertised as VPN routes. This is to avoid forwarding entries to the prefixes of the MUP Segment mingled in the other type of routing instance.

A MUP PE may discard the received MUP Segment Discovery route if the Route Target extended communities of the route does not meet the MUP PE's import policy.

5.1. Direct Segment Discovery (DSD) Route

When a MUP PE accommodates a network, a service, or an any type of resource through an interface or a routing instance as a Direct Segment, the MUP PE advertises the corresponding Direct Segment Discovery (DSD) route for the interface or the routing instance to the SR domain. The DSD route includes an address to indicate the Direct Segment in the MUP PE in the network layer reachability information (NLRI) with an extended community indicating the corresponding Direct Segment. Dataplane specific attribute should be attached to the DSD route, as a auto-discovery route, which itself must be dataplane independent defined in Section 3.

For example in 3GPP 5G specific case with SRv6 dataplane, an MUP PE may connect to N6 interface on a DN side, and a DSD route for the DN will be advertised with an address of the MUP PE in NLRI, an corresponding SRv6 SID in the SRv6 specific attribute, and a Direct Segment type MUP extended community to the routing instance for the DN from the MUP PE.

When a MUP PE receives a DSD route from other PEs, the MUP PE keeps the received DSD route in the RIB. The MUP PE uses the received DSD route to resolve Type 2 Session Transformed (ST2 for short) routes reachability, described in Section 6.2. If the DSD route resolves reachability for the endpoints, and match the Direct Segment extended community of the ST2 routes, the MUP PE updates the FIB entry for the ST2 route with the SID of the matched DSD route.

5.2. Interwork Segment Discovery (ISD) Route

When a PE accommodates a network through an interface or a routing instance for the user plane protocol of the mobile service architecture as an Interwork Segment, the PE advertises the corresponding Interwork Segment Discovery (ISD) route with the prefixes of the Interwork Segment. Dataplane specific attribute should be attached to the ISD route, as a auto-discovery route, which itself must be dataplane independent defined in Section 3.

For example in 3GPP 5G specific case with SRv6 dataplane, an MUP PE may connect to N3 network accommodating a RAN, and a ISD route will be advertised with IP prefix(es) of the N3 RAN in NLRI, and an corresponding SRv6 SID in the SRv6 specific attribute.

When a MUP PE receives a ISD route, the MUP PE keeps the received ISD routes in the RIB. The MUP PE uses the received ISD routes to resolve the reachability for remote endpoint of Type 1 Session Transformed (ST1 for short) routes, described in Section 6.1. If the ISD route resolves the reachability for ST1 routes, the MUP PE updates the FIB entry for the prefix of ST1 route with the SID of the matched ISD route.

The MUP PE may also use the received ISD routes to resolve the reachability for the endpoint address in the ST2 route prefix. If the ISD route resolves the reachability for ST2 routes, the MUP PE updates the FIB entry for the prefix of ST2 route with the SID of the matched ISD route.

6. Distribution of Session Transformed Route

MUP architecture defines two types of session transformed route.

6.1. Type 1 Session Transformed (ST1) Route

First type route, called Type 1 Session Transformed (ST1) route, encodes IP prefix(es) for a UE or MN in a BGP MP-NLRI attribute with associated session information of the tunnel endpoint identifier on the access side. The MUP-C advertises the ST1 route with the Route Target extended communities for the UE or MN to the MUP networks.

A MUP PE may receive the ST1 routes from the MUP-C in the MUP networks. The MUP PE may keep the received ST1 routes advertised from the MUP-C. The receiving MUP PE will perform the importing of the received ST1 routes in the configured routing instances based on the Route Target extended communities. A MUP PE may discard the received ST1 route if the MUP PE fails to import the route based on the Route Target extended communities.

6.2. Type 2 Session Transformed (ST2) Route

Second type route, called Type 2 Session Transformed (ST2) route, encodes the tunnel endpoint identifier of the session in a BGP MP-NLRI attribute with the nature of tunnel endpoint base packet handling. Longest match algorithm for the prefix in this type of session transformed route should be applicable to aggregate the routes for scale. The MUP-C advertises the ST2 route with a Direct Segment extended community to indicate corresponding Direct Segment and tunnel decapsulation in addition to the Route Target extended communities

A MUP PE may receive the ST2 routes from the MUP-C in the SR domain. The MUP PE may keep the received ST2 routes advertised from the MUP-C. The receiving MUP PE will perform the importing of the received ST2 routes in the configured routing instances based on the Route Target extended communities. A MUP PE may discard the received ST2 route if the MUP PE fails to import the route based on the Route Target extended communities.

6.3. MUP Controller

A MUP Controller (MUP-C) provides an API. A consumer of the API inputs session information for a UE or a MN from mobile service system. The MUP-C transforms the received session information to routing information and will advertise the session transformed routes with the corresponding extended communities to the MUP networks.

The received session information is expected to include the UE or MN IP prefix(es), tunnel endpoint identifiers for both ends, and any other attributes for the mobile networks. For example in a 3GPP 5G specific case, the tunnel endpoint identifier will be a pair of the F-TEIDs on both the N3 access side (RAN) and core side (UPF).

7. Illustration

This section illustrates possible MUP deployments with SRv6 dataplane. 3GPP 5G is an example mobile service for the deployment cases in this section.

7.1. SR Network Accommodating Existing Mobile Network Services

Figure 2 shows how SR networks can accommodate existing mobile network service before enabling MUP. The PEs S1, S2, and S3 compose an SR network. A routing instance is configured to each network of the mobile service. N6DN in S1 and S2 are providing connectivity to edge servers and the Internet respectively.

VRF (Virtual Routing Forwarding) is the routing instance to accommodate MUP segments in this section. All example cases in this section follow the typical routing policy control using the BGP extended community described in [RFC4360] and [RFC4684]

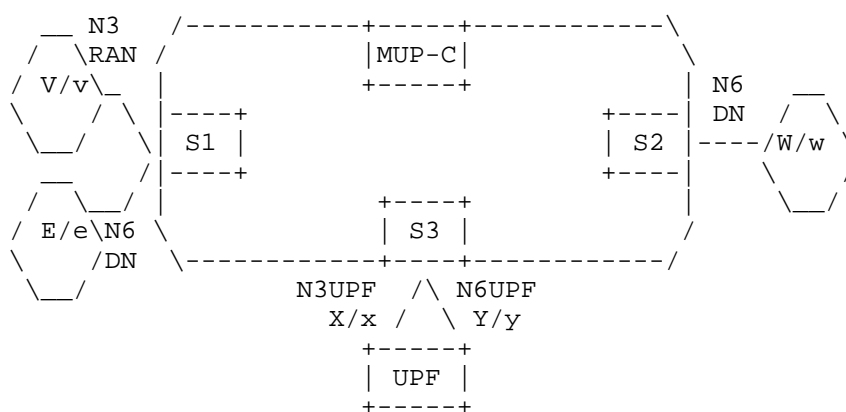


Figure 2

The following routing instances are configured:

- * N3RAN in S1
 - export route V/v with route-target (RT) community C1
 - import routes which have route-target (RT) community C1 and C2
- * N6DN in S1
 - export route E/e with RT C4
 - import routes which have RT C3 and C4
- * N6DN in S2
 - export route W/w with RT C4

- import routes which have RT C3 and C4
- * N3UPF in S3
 - export route X/x with RT C2
 - import routes which have RT C1
- * N6UPF in S3
 - export route Y/y with RT C3
 - import routes which have RT C4

Note: The above configurations are just to provide typical IP connectivity for 3GPP 5G. When the above configurations have been done, each endpoint in V/v and X/x can communicate through S1 and S3, but they can not communicate with nodes in E/e, W/w and Y/y.

7.2. MUP PE Deployment at All SR Domain Edges

Here, the PEs S1, S2 and S3 are configured to enable MUP as follows:

- * S1
 - advertises Interwork type discovery route: V/v with SID S1::
 - set S1:: behavior End.M.GTP4.E or End.M.GTP6.E
- * S1
 - advertise Direct type discovery route: MUP Direct Segment community D1 and SID S1:1::
 - set S1:1:: behavior End.DT4 or End.DT6 for the N6DN in S1
- * S2
 - advertise Direct type route: MUP Direct Segment community D1 and SID S2::
 - set S2:: behavior End.DT4 or End.DT6 for the N6DN in S2

S1 adopts the local N6DN to prioritize the closer segment for the same Direct Segment. Another PE may adopt D1 from S2, if the PE has no local N6DN for D1 and closer to S2 than S1.

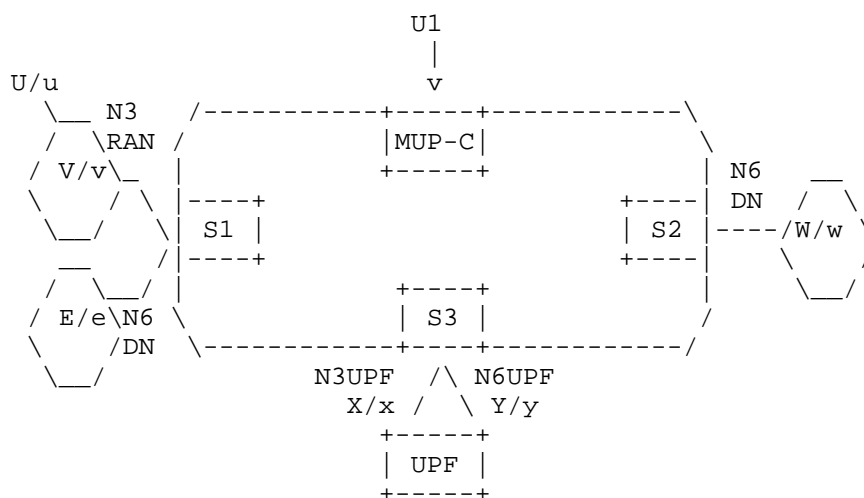


Figure 3

* MUP-C

- attach the MUP Direct Segment ID D1 and RT C3 to the DN in U1
- transforms UE's prefix U/u, the F-TEID on access side (gNB) and QFI in U1 to the ST1 route for the prefix U/u with the F-TEID, the QFI, and RT C3
- transforms F-TEID on core side (UPF) X in U1 to the ST2 route for X with MUP segment-ID D1 and RT C2

Then N3RAN and N6DN import route X and U/u respectively. S1 and S2 resolves U/u's remote endpoint with V/v and then install SID S1:: for U/u in FIB. S1:: will not appear in the packet from E/e to U/u over the wire.

As S1 adopts local N6DN for D1, N3RAN in S1 decapsulates GTP-U packets from V/v to X and then lookup the inner packets from U/u in N6DN after the decapsulation.

Note: When the above configurations have been done, MUP is applied only to the packets from/to U/u. Each endpoint in U/u, W/w and E/e can communicate through S1 and S2. The rest of traffic from/to other UEs go through the usual 3GPP 5G user plane path using UPF via S3.

- import routes which have RT C1
- * N6UPF in S3 (same with the previous case)
 - export route Y/y with RT C3
 - import routes which have RT C4
- * N6DN in S4
 - export route E/e with RT C4
 - import routes which have RT C3 and C4

Here, the PEs are configured to enable MUP as following:

- * S1 (same with the previous case)
 - advertises Interwork type route: V/v with SID S1::
 - set S1:: behavior End.M.GTP4.E or End.M.GTP6.E
- * S1
 - advertise Direct type route: MUP Direct Segment community D1 for the local N6DN
 - set S1:1:: behavior End.DT4 or End.DT6 for the N6DN in S1
- * S2 (same with the previous case)
 - advertise Direct type route: MUP Direct Segment community D1 and SID S2::
 - set S2:: behavior End.DT4 or End.DT6 for the N6DN in S2
- * S4
 - advertise Direct type route: MUP Direct Segment community D2 and SID S4::
 - set S4:: behavior End.DT4 or End.DT6 for the N6DN in S4

As in the previous case, S1 adopts the local N6DN for D1 as long as S1 prioritizes the closer segment for the same MUP Direct Segment. The Direct type route from S4 for D2 with SID S4:: will be kept in S1.

- * MUP-C (same with the previous case)
 - attach the MUP Direct Segment ID D1 and RT C3 to the DN in U1
 - transforms UE's prefix U/u, the F-TEID on access side (gNB) and QFI in U1 to the ST1 route for the prefix U/u with the F-TEID, the QFI, and RT C3
 - transforms F-TEID on core side (UPF) X in U1 to the ST2 route for X with MUP Direct Segment community D1 and RT C2

Then N3RAN and N6DN import route X and U/u respectively. S2 and S4 resolve U/u's remote endpoint with V/v and then install SID S1:: for U/u in FIB.

As in the previous case, S1 adopts local N6DN for D1, N3RAN in S1 decapsulates GTP-U packets from V/v to X and then lookup the inner packets from U/u in N6DN after the decapsulation.

For D2 on the other hand, no corresponding N6DN existed in S1. However, E/e with RT C4 from S4 is imported into N6DN in S1 as a VPN route, E/e is reachable from U/u via N6DN for D1 in S1.

If a session U1' includes the DN corresponding to D2, MUP-C advertises ST2 route X' with MUP Direct Segment community D2, and then N3RAN in S1 instantiates H.M.GTP4.D or End.M.GTP6.D for X with S4:: as the last SID in the received Direct type route from S4.

Note: When the above configurations have been done, MUP is applied only to the packets from/to U/u. Each endpoint in U/u, W/w and E/e can communicate through S1, S2 and S4. The rest of traffic from/to other UEs go through the usual 3GPP 5G user plane path using UPF via S3.

7.4. Collapsed MUP PE Deployment

In this case only S1 enables MUP in a collapsed fashion. S2 and S3 are L3VPN PEs without MUP capability. In this section, S2 and S3 are illustrated as SRv6 nodes. But they can be non-SR nodes if S1 provides SR independent connectivity to S2 and S3.

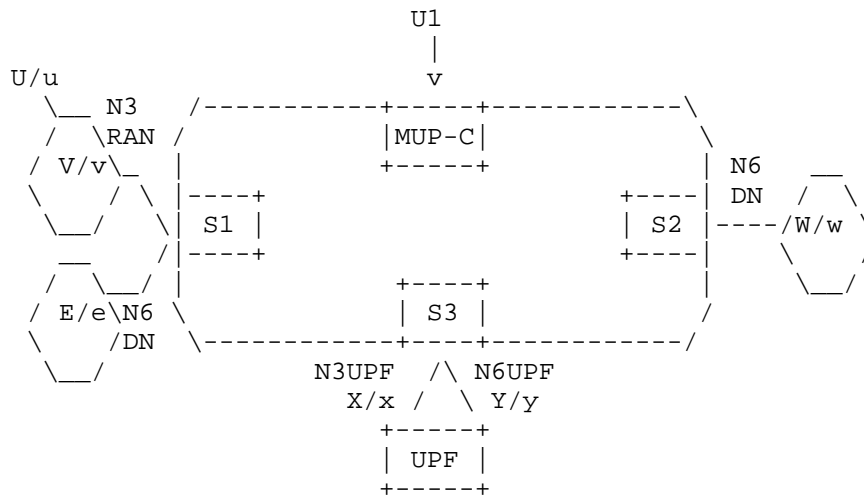


Figure 5

The difference between the previous case in Section 7.1 for the routing instance configuration is following:

- * N6DN in S1
 - export route E/e with RT C4
 - import routes which have RT C3, C4 and C5

Here, S1 is configured to enable MUP and S2 as an L3VPN PE is configured as follows:

- * S1
 - may not advertise Interwork type discovery route for V/v
 - may not advertise Direct type discovery route with MUP Direct Segment community D1 and S1:1::
 - set S1:1:: behavior End.DT4 or End.DT6 for the N6DN in S1
- * S2
 - set S2:: behavior End.DT4 or End.DT6 for the N6DN in S2

Now, session information U1 is added to the MUP Controller, MUP-C, and MUP-C and S1 is configured to transform U1 to the routes as follows:

* MUP-C

- attach the MUP Direct Segment ID D1 and RT C5 to the DN in U1
- transforms UE's prefix U/u, the F-TEID on access side (gNB) and QFI in U1 to the ST1 route for the prefix U/u with the F-TEID, the QFI, and RT C5
- transforms F-TEID on core side (UPF) X in U1 to the ST2 route for X with MUP Direct Segment community D1 and RT C2

* S1

- advertises U/u as an L3VPN route with RT C4 and SID S1:1::, when the ST1 route is imported into the N6DN

Then the N3RAN and N6DN import route X and U/u respectively. S1 resolves U/u's remote endpoint with V/v and then create the corresponding GTP encap entry for U/u into the N3RAN FIB. S2 will create a regular L3VPN routing entry for U/u with SID S1:1:: in the N6DN when S2 imports the L3VPN route with RT C4 for U/u advertised from S1.

As S1 adopts local N6DN for D1, N3RAN in S1 decapsulates GTP-U packets from V/v to X and then lookup the inner packets from U/u in N6DN after the decapsulation.

Note: When the above configurations have been done, MUP is applied only to the packets from/to U/u. Each endpoint in U/u, W/w and E/e can communicate through S1 and S2. The rest of traffic from/to other UEs go through the usual 3GPP 5G user plane path using UPF via S3.

8. IANA Considerations

This memo includes no request to IANA.

9. Security Considerations

TBD.

10. References

10.1. Normative References

- [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>>.
- [RFC7333] Chan, H., Ed., Liu, D., Seite, P., Yokota, H., and J. Korhonen, "Requirements for Distributed Mobility Management", RFC 7333, DOI 10.17487/RFC7333, August 2014, <<https://www.rfc-editor.org/info/rfc7333>>.
- [RFC8402] Filsfils, C., Ed., Previdi, S., Ed., Ginsberg, L., Decraene, B., Litkowski, S., and R. Shakir, "Segment Routing Architecture", RFC 8402, DOI 10.17487/RFC8402, July 2018, <<https://www.rfc-editor.org/info/rfc8402>>.
- [RFC8986] Filsfils, C., Ed., Camarillo, P., Ed., Leddy, J., Voyer, D., Matsushima, S., and Z. Li, "Segment Routing over IPv6 (SRv6) Network Programming", RFC 8986, DOI 10.17487/RFC8986, February 2021, <<https://www.rfc-editor.org/info/rfc8986>>.

10.2. Informative References

- [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>>.
- [RFC4684] Marques, P., Bonica, R., Fang, L., Martini, L., Raszuk, R., Patel, K., and J. Guichard, "Constrained Route Distribution for Border Gateway Protocol/MultiProtocol Label Switching (BGP/MPLS) Internet Protocol (IP) Virtual Private Networks (VPNs)", RFC 4684, DOI 10.17487/RFC4684, November 2006, <<https://www.rfc-editor.org/info/rfc4684>>.
- [RFC5213] Gundavelli, S., Ed., Leung, K., Devarapalli, V., Chowdhury, K., and B. Patil, "Proxy Mobile IPv6", RFC 5213, DOI 10.17487/RFC5213, August 2008, <<https://www.rfc-editor.org/info/rfc5213>>.
- [RFC8885] Bernardos, CJ., de la Oliva, A., Giust, F., Z炭単iga, JC., and A. Mourad, "Proxy Mobile IPv6 Extensions for Distributed Mobility Management", RFC 8885, DOI 10.17487/RFC8885, October 2020, <<https://www.rfc-editor.org/info/rfc8885>>.

[RFC9433] Matsushima, S., Ed., Filsfils, C., Kohno, M., Camarillo, P., Ed., and D. Voyer, "Segment Routing over IPv6 for the Mobile User Plane", RFC 9433, DOI 10.17487/RFC9433, July 2023, <<https://www.rfc-editor.org/info/rfc9433>>.

[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>>.

Acknowledgements

The authors would like to thank Jeffrey Zhang, Keiichi Makizono, Hideyuki Sasaki, Teruyuki Oya, Yoichi Funabiki, Daisuke Koshiro, Derek Yeung, Kalyani Rajaraman, Csaba Keszei, Daiki Kijima, Ryuma Seno, Keiji Hara, Mikio Nakajima, Masayuki Yamazaki, Takuto Shiozawa, Manabu Mikami, Takaya Watanabe, Hidenori Aita, Hiroyuki Fukumori, Leo Fujita, Toshihiro Yokomine, Toshikazu Morioka, Yadav Rituraji, Junya Yoshino, Takayuki Koshikawa, Kazunari Nagai, Tadayuki Okuda, Takahiro Hara, Yuya Kusakabe, Takeru Hayasaka, Hideki Takase, Yutaka Kikuchi, Kazuma Nishiuchi, Mitsuhiro Osaki, Suresh Krishnan, Sri Gundavelli, Zafar Ali, Luay Jalil, Markus Amend, Marco Liebsch, and Lionel Morand for their reviews and discussions with the authors.

Contributors

Clarence Filsfils
Cisco Systems, Inc.
Belgium
Email: cf@cisco.com

Ketan
Cisco Systems
India
Email: ketant.ietf@gmail.com

Additional Authors

Ashiq Khan
SoftBank
Japan
Email: ashiq.khan@g.softbank.co.jp

Miya Kohno
Cisco Systems, Inc.
Japan

Email: mkohno@cisco.com

Teppei Kamata
Cisco Systems, Inc.
Japan
Email: tkamata@cisco.com

Pablo Camarillo Garvia
Cisco Systems, Inc.
Spain
Email: pcamaril@cisco.com

Daniel Voyer
Bell Canada
Canada
Email: daniel.voyer@bell.ca

Shay Zadok
Broadcom
Israel
Email: shay.zadok@broadcom.com

Israel Meilik
Broadcom
Israel
Email: israel.meilik@broadcom.com

Ashutosh Agrawal
Intel
United States of America
Email: ashutosh.agrawal@intel.com

Kumaresh Perumal
Intel
United States of America
Email: kumaresh.perumal@intel.com

Authors' Addresses

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

Katsuhiro Horiba
Japan
Email: katsuhiro.horiba@gmail.com

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

Tetsuya Murakami
Arrcus, Inc.
United States of America
Email: tetsuya@arrcus.com

Keyur Patel
Arrcus, Inc.
United States of America
Email: keyur@arrcus.com

Jakub Horn
Cisco Systems, Inc.
Czech Republic
Email: jakuhorn@cisco.com