

Workgroup: Internet Engineering Task Force

Internet-Draft:

draft-mhkk-dmm-srv6mup-architecture-03

Published: 19 March 2022

Intended Status: Standards Track

Expires: 20 September 2022

Authors: S. Matsushima K. Horiba A. Khan Y. Kawakami
SoftBank SoftBank SoftBank SoftBank
T. Murakami K. Patel M. Kohno
Arrcus, Inc Arrcus, Inc Cisco Systems, Inc.
T. Kamata P. Camarillo
Cisco Systems, Inc. Cisco Systems, Inc.
D. Voyer S. Zadok I. Meilik A. Agrawal
Bell Canada Broadcom Broadcom Intel
K. Perumal J. Horn
Intel Cisco Systems, Inc.

Segment Routing IPv6 Mobile User Plane Architecture for Distributed Mobility Management

Abstract

This document defines the Segment Routing IPv6 Mobile User Plane (SRv6 MUP) architecture for Distributed Mobility Management. The requirements for Distributed Mobility Management described in [RFC7333] can be satisfied by routing fashion.

Mobile services are deployed over several parts of IP networks. A Segment Routing over IPv6 (SRv6) network can accommodate all, or part of those networks thanks to the large address space of IPv6 and the network programming capability described in [RFC8986].

Segment Routing IPv6 Mobile User Plane Architecture can incorporate existing session based mobile networks. By leveraging SRv6 network programmability, mobile user plane can be integrated into the SRv6 data plane. In that routing paradigm, session information between the entities of the mobile user plane is turned to routing information.

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 20 September 2022.

Copyright Notice

Copyright (c) 2022 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
1.1.	Requirements Language
2.	Terminology
3.	Architecture Overview
4.	Mobile User Plane Segment
5.	Distribution of Mobile User Plane Segment Information
5.1.	Direct Segment Discovery Route
5.2.	Interwork Segment Discovery Route
6.	Distribution of Session Transformed Route
6.1.	Type 1 Session Transformed Route
6.2.	Type 2 Session Transformed Route
6.3.	MUP Controller
7.	Illustration
8.	IANA Considerations
9.	Security Considerations
10.	References
10.1.	Normative References
10.2.	Informative References
	Authors' Addresses

1. Introduction

Mobile services require IP connectivity for communication between the entities of mobile service architecture [[RFC5213](#)][[TS.23501](#)]. To provide the IP connectivity, Segment Routing (SR) [[RFC8402](#)] can be a candidate solution.

In PMIPv6 [[RFC5213](#)], IP connectivity between LMA and MAG can be provided over SR networks, as well as LMA and Internet. In 3GPP 5G [[TS.23501](#)], IP connectivity for N3 interface between gNodeB(es) and UPFs can also be provided by SR, as well as for N6 interface between UPFs and DNS (Data Network).

These IP connectivities may be covered by multiple SR networks, or just one SR network, depending on the size of the deployment. In the latter case, it is expected that the address space of the SR network should be large enough to cover a vast number of nodes, such as millions of base stations. For this reason, use of IPv6 for the SR dataplane looks sufficiently suitable.

SRv6 is an instantiation of SR over IPv6 dataplane in which a single network can accommodate all entities of mobile services thanks to the huge available address space and network programming capability described in [[RFC8986](#)].

Meanwhile, SRv6 network programmability enhances SRv6 dataplane to be integrated with mobile user plane [[I-D.ietf-dmm-srv6-mobile-uplane](#)]. It will make an entire SRv6 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 mobility 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

across different networks

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 Segment Routing IPv6 Mobile User Plane (SRv6 MUP) architecture for Distributed Mobility Management. SRv6 MUP is not a mobility management system itself, but an architecture to integrate mobile user plane into the SRv6 data plane.

In this routing paradigm, session information from a mobility management system will be transformed to routing information. It means that mobile user plane specific nodes for the anchor or intermediate points are no longer required. The user plane anchor and intermediate functions can be supported by SR throughout an SR domain (REQ1), not to mention that SRv6 MUP will naturally be deployed over IPv6 networks (REQ3).

SRv6 MUP architecture is independent from the mobility management system. For the requirements (REQ4, 5), SRv6 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 SRv6 MUP integration for mobile networks running based on the existing architecture will be varied and depending on the level of SRv6 awareness of the control and user plane entities.

Specifying how to modify the existing architecture to integrate SRv6 MUP is out of scope of this document. What this document provides for the existing architecture is an interface for SRv6 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

PE: Provider Edge node in an SR network

MUP Controller: Controller node for an SR network

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

MN: Mobile Node, as per [[RFC5213](#)]

3. Architecture Overview

SRv6 MUP architecture defined in this document introduces new segment types of MUP segment called "Direct segment", and "Interwork Segment". An SR node of PE accommodates an Interwork Segment and/or a Direct Segment. [Figure 1](#) depicts the overview.




Figure 1: Overview of SRv6 MUP Architecture

This document also defines new routing information called "Segment Discovery route" and "Session Transformed route". 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 new 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 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 SRv6 network. 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 SRv6 networks.

An SRv6 SID (Segment Identifier) can represent a MUP segment. The SID can be any behavior defined in [[RFC8986](#)], [[I-D.ietf-dmm-srv6-mobile-uplane](#)], 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 PE interfaces to 5G user plane on the access side defined as "N3" in [[TS.23501](#)], the PE accommodates the N3 network as Interwork Segment in a routing instance and then the behavior of created segment SID by the PE will be "End.M.GTP4.E", or "End.M.GTP6.E". In this case, the PE may associate the SID to the routing instance for the N3 access network (N3RAN).

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

A MUP Segment Discovery route is routing information that associates the MUP segment with network reachability. This document defines the basic discovery route types, Direct Segment Discovery route, and Interwork Segment Discovery 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.

5.1. Direct Segment Discovery Route

When a PE accommodates a network through an interface or a routing instance as a Direct Segment, the PE advertises the corresponding Direct Segment Discovery route for the interface or the routing instance. The Direct Segment Discovery route includes an address of the PE in the network reachability information with an extended community indicating the corresponding Direct Segment, and SID of the routing instance to the SR domain.

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

When a PE receives a Direct Segment Discovery route from other PEs, the PE keeps the received Direct Segment Discovery route in the RIB. The PE uses the received Direct Segment Discovery route to resolve Type 2 session transformed routes reachability, described in [Section 6.2](#). If the Direct Segment Discovery route resolves reachability for the endpoints, and match the Direct Segment extended community of the Type 2 session transformed routes, the PE updates the FIB entry for the Type 2 session transformed route with the SID of the matched Direct Segment Discovery route.

5.2. Interwork Segment Discovery 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 route with the prefixes of the Interwork Segment and the corresponding SID of the prefixes to the SR domain.

For example in 3GPP 5G specific case, an Interwork Segment Discovery route for N3 network accommodating RAN will be incorporated in an N3RAN segment discovery route associated with a RAN segment SID.

When a PE receives a Interwork Segment Discovery route, the PE keeps the received Interwork Segment Discovery routes in the RIB. The PE uses the received Interwork Segment Discovery routes to resolve the

reachability for remote endpoint of Type 1 session transformed routes, described in [Section 6.1](#). If the Interwork Segment Discovery route resolves the reachability for Type 1 session transformed routes, the PE updates the FIB entry for the prefix of Type 1 session transformed route with the SID of the matched MUP segment discovery route.

The received Interwork Segment Discovery routes MUST be used only to resolve reachability for the remote endpoints of Type 1 session transformed routes. The connectivity among the routing instances for Interwork Segments may be advertised as VPN routes. This is to avoid forwarding entries to the prefixes of Interwork Segment mingled in the other type of routing instance. A PE may discard the received Interwork segment discovery route if the Route Target extended communities of the route does not meet the PE's import policy.

6. Distribution of Session Transformed Route

SRv6 MUP architecture defines two types of session transformed route.

6.1. Type 1 Session Transformed Route

First type route, called Type 1 Session Transformed 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 controller advertises the Type 1 Session Transformed route with the Route Target extended communities for the UE or MN to the SR domain.

A PE may receive the Type 1 Session Transformed routes from the MUP Controller in the SR domain. The PE may keep the received Type 1 Session Transformed routes advertised from the MUP Controller. The receiving PE will perform the importing of the received Type 1 Session Transformed routes in the configured routing instances based on the Route Target extended communities. A PE may discard the received Type 1 Session Transformed route if the PE fails to import the route based on the Route Target extended communities.

6.2. Type 2 Session Transformed Route

Second type route, called Type 2 Session Transformed route, encodes the tunnel endpoint identifier of the session on the core side in a BGP MP-NLRI attribute with the nature of tunnel decapsulation. Longest match algorithm for the prefix in this type of session transformed route should be applicable to aggregate the routes for scale. The MUP controller advertises the Type 2 Session Transformed route with the Route Target and Direct Segment extended communities for the endpoint to the SR domain.

A PE may receive the Type 2 Session Transformed routes from the MUP Controller in the SR domain. The PE may keep the received Type 2 Session Transformed routes advertised from the MUP Controller. The receiving PE will perform the importing of the received Type 2 Session Transformed routes in the configured routing instances based on the Route Target extended communities. A PE may discard the received Type 2 Session Transformed route if the PE fails to import the route based on the Route Target extended communities.

6.3. MUP Controller

A MUP controller provides a northbound API. A consumer of the API inputs session information for a UE or a MN from mobility management system. The MUP controller transforms the received session information to routing information and will advertise the session transformed routes with the corresponding extended communities to the SR domain.

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 shows an illustration of SRv6 MUP deployment. The example deployment cases here is 3GPP 5G.

Before enabling SRv6 MUP, how SRv6 networks can accommodate existing mobile network service shown in [Figure 2](#). The PEs of S1, S2, and S3 join an SR network. A routing instance is configured to each network of the mobile service. N6DN in S1 and S2 are supposed to provide 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

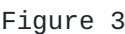
*N6UPF in S3

- * -export route Y/y with RT C3
- import routes which have RT C4

Note:

Here, the PEs are configured to enable SRv6 MUP as following:

S1 here adopts the local N6DN to prioritize 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.



Now, session information U1 is put to a MUP Controller, MUP-C, and MUP-C is configured to transforms U1 to the routes as follows:

*MUP-C

* -attach the 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 Type 1 session transformed 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 Type 2 session transformed 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 be appeared 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, SRv6 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.

Another case shown in [Figure 4](#) is that S4 joins the SR network and accommodates edge servers in the N6DN in S4.

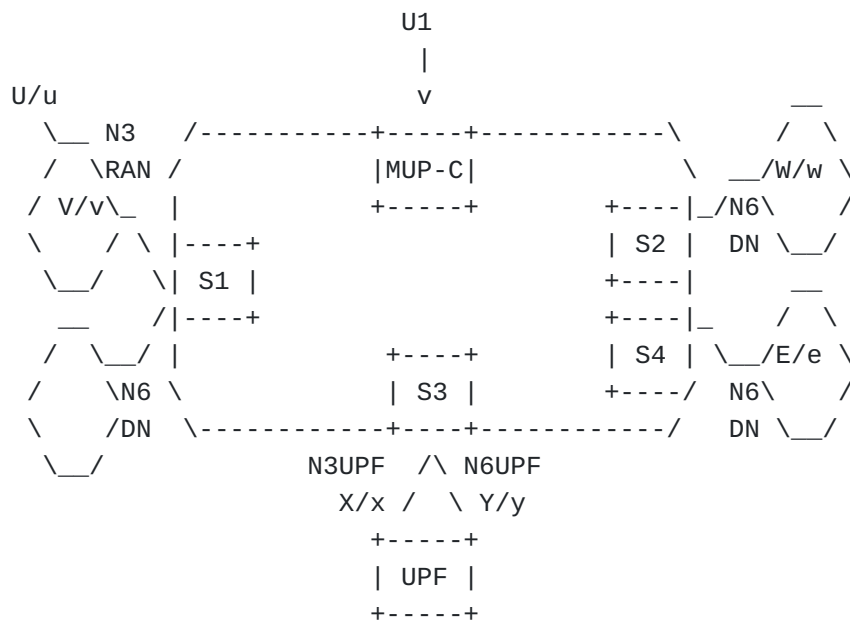


Figure 4

The following routing instances are configured:

*N3RAN in S1 (same with the previous case)

- * -export route V/v with RT C1
- import routes which have RT C1 and C2

*N6DN in S1

- * -export no route
- import routes which have RT C4

*N6DN in S2 (same with the previous case)

- * -export route W/w with RT C4
- import routes which have RT C3 and C4

*N3UPF in S3 (same with the previous case)

- * -export route X/x with RT C2
- 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 SRv6 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 same as the previous case, S1 adopts the local N6DN for D1 as long as S1 prioritizes 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 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 Type 1 session transformed 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 Type 2 session transformed 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 same as 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 DN corresponding to D2, MUP-C advertises Type 2 session transformed 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, SRv6 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.

8. IANA Considerations

This memo includes no request to IANA.

9. Security Considerations

TBD.

10. References

10.1. Normative References

[I-D.ietf-dmm-srv6-mobile-uplane]

Matsushima, S., Filsfils, C., Kohno, M., Garvia, P. C., Voyer, D., and C. E. Perkins, "Segment Routing IPv6 for Mobile User Plane", Work in Progress, Internet-Draft, draft-ietf-dmm-srv6-mobile-uplane-18, 18 February 2022, <<https://datatracker.ietf.org/doc/html/draft-ietf-dmm-srv6-mobile-uplane-18>>.

[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, C.J., de la Oliva, A., Giust, F., Zúñiga, J.C., 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>>.
- [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>>.

Authors' Addresses

Satoru Matsushima
SoftBank
Japan

Email: satoru.matsushima@g.softbank.co.jp

Katsuhiro Horiba
SoftBank
Japan

Email: katsuhiro.horiba@g.softbank.co.jp

Ashiq Khan
SoftBank
Japan

Email: ashiq.khan@g.softbank.co.jp

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

Miya Kohno
Cisco Systems, Inc.
Japan

Email: mkohno@cisco.com

Teppey 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

Jakub Horn
Cisco Systems, Inc.
Czech Republic

Email: jakuhorn@cisco.com