

DMM WG
Internet-Draft
Intended status: Standards Track
Expires: March 23, 2019

S. Gundavelli
Cisco
M. Liebsch
NEC
S. Matsushima
SoftBank
September 19, 2018

Mobility-aware Floating Anchor (MFA)
draft-gundavelli-dmm-mfa-01.txt

Abstract

IP mobility protocols are designed to allow a mobile node to remain reachable while moving around in the network. The currently deployed mobility management protocols are anchor-based approaches, where a mobile node's IP sessions are anchored on a central node. The mobile node's IP traffic enters and exits from this anchor node and it remains as the control point for all subscriber services. This architecture based on fixed IP anchors comes with some complexity and there is some interest from the mobile operators to eliminate the use of fixed anchors, and other residual elements such as the overlay tunneling that come with it.

This document describes a new approach for realizing a mobile user-plane that does not require fixed IP anchors. The architectural-basis for this approach is the separation of control and user plane, and the use of programmability constructs of the user-plane for traffic steering. This approach is referred to as, Mobility-aware Floating Anchor (MFA).

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 <http://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 March 23, 2019.

Copyright Notice

Copyright (c) 2018 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 (<http://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 Simplified BSD License text as described in Section 4.e of the Trust Legal Provisions and are provided without warranty as described in the Simplified BSD License.

Table of Contents

1. Introduction	3
2. Conventions and Terminology	4
2.1. Conventions	4
2.2. Terminology	4
3. Overview	6
3.1. The Network Topology Database	9
3.2. The Node Location Database	9
3.3. Determination of the Correspondent Node Anchor	10
3.4. Traffic Steering Approaches	10
3.5. Mobile Node Attachment Triggers	12
3.6. Programming the User-plane	12
4. Life of a Mobile Node in a MFA Domain	14
4.1. MN's Initial Attachment to a MFA Domain	15
4.2. MN's Roaming within the MFA Domain	17
4.3. Traffic Steering State Removal	20
4.4. Mobile Node's new IP flows	21
5. MFA in 5G System Architecture	21
6. IANA Considerations	23
7. Security Considerations	23
8. Acknowledgements	23
9. References	23
9.1. Normative References	23
9.2. Informative References	23
Authors' Addresses	24

1. Introduction

IP mobility protocols are designed to allow a mobile node to remain reachable while moving around in the network. The currently deployed mobility management protocols are anchor-based approaches, where a mobile node's IP sessions are anchored on a central node. The mobile node's IP traffic enters and exits from this anchor node and it remains as the control point for all subscriber services. This architecture based on fixed IP anchors comes with some complexity and there is some interest from the mobile operators to eliminate the use of fixed anchors, and other residual elements such as the overlay tunneling that come with it. Some of the key objectives for this effort are listed below.

- o Access-agnostic, shared user-plane that can be used for multiple access technologies
- o Optimized Routing for the mobile node's IP flows with topology awareness and leveraging the transport QoS
- o Elimination of overlay tunnels from the user-plane network for avoiding packet fragmentation, and reducing encapsulation related packet-size overhead
- o Elimination of centralized mobility anchors and shift towards a distributed mobility architecture, leveraging the edge compute at radio-access network for offloading some of the subscriber management services
- o Co-existence with control-plane and user-plane separated architecture; a stateless user-plane with no tunnels, and a control plane with the business/service logic
- o Support for services including accounting, charging, lawful-interception and other user plane services

Currently, there is a study item in 3GPP to explore options for simplifying the mobile user-plane. There are few proposals in IETF, which are presented as candidate solutions for user-plane simplification. However, each of these proposals come with certain complexity and do not leverage the 3GPP control plane, or the programmability aspects of the user-plane. For example, ILA defines a translation scheme without the need for overlay tunnels, but it also introduces significant amount of translation related state in the user-plane, and additionally introduces a new control-plane protocol for managing the mapping tables and the cache states. Therefore, we believe that none of the currently known approaches can adequately meet the stated goals for user-plane simplification.

This document describes a new approach for realizing a mobile user-plane that does not require any fixed IP anchors. The first-hop router on the link where the mobile node is attached remains as the IP anchor and thereby eliminating the need for IP tunneling to some central anchor node. Even when the mobile node moves in the network and changes its point of attachment, the IP anchor is always the first-hop router on that new link. The MFA entities will track the mobile node's movements in the network and will ensure the mobile node's IP flows always take the most optimal routing path. This is achieved by MFA entities programming the needed traffic steering rules for moving mobile node's IP packets directly between the correspondent node and the mobile node's edge anchor, which can be relocated to a new edge, e.g. in case of mobility. Furthermore, this approach does not require a new control-plane protocol, but instead leverages the SDN interfaces of the user-plane, and the mobility events in the control-plane for managing IP mobility. The architectural basis for this approach is the separation of control and user plane, and the use of programmability constructs of the user-plane for traffic steering. This approach is referred to as, Mobility-aware Floating Anchor (MFA). The rest of the document explains the operational details of the MFA approach.

2. Conventions and Terminology

2.1. Conventions

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].

2.2. Terminology

All mobility related terms used in this document are to be interpreted as defined in the IETF mobility specifications, including [RFC5213] and [RFC6275]. Additionally, this document uses the following terms:

MFA Domain

MFA domain refers to the network where the mobility management of a mobile node is handled by the MFA entities. The MFA domain includes MFA mobile node anchors, MFA corresponding node anchors, and MFA node controller, between which security associations can be set up for authorizing the configuration of traffic steering policies and other mobility management functions.

MFA Mobile Node Anchor (MFA-MNA)

Its an MFA function located in the user-plane network very close to the layer-2 access-point to where the mobile node is attached. It is typically on the first-hop router for the mobile node's IP traffic. The node hosting this function is required to support the standard IPv6 packet forwarding function, FPC or a similar interface for policy configuration, and packet steering functions such as based on SRv6 or alternative means that can support per-flow or per-flow-aggregate traffic steering. Typically, the MFA-MNA function will be collocated with the User Plane Function (UPF) in the 3GPP 5G system architecture.

MFA Corresponding Node Anchor (MFA-CNA)

Its an MFA function located in the user-plane node in the path between the mobile node and the correspondent node. If the correspondent node is another mobile node in the MFA domain, then the MFA-CNA is on the first hop router on the link shared with the correspondent node. The node hosting this function is required to support the standard IPv6 packet forwarding function, FPC or a similar interface for policy configuration, and packet steering functions such as based on SRv6 or alternative means that can support per-flow or per-flow-aggregate traffic steering. Typically, the MFA-CNA function will be collocated with the IP forwarding nodes on the N6 interface of the 3GPP 5G system architecture.

MFA Node

A generic term used for referring to MFA-MNA, or the MFA-CNA.

MFA Node-Controller (MFA-NC)

The is the function that controls the forwarding policies on the MFA-MNA and MFA-CNA nodes. This entity interfaces with the MFA node using the FPC interface [I-D.ietf-dmm-fpc-cdp], or a similar interface that support user-plane policy configuration. This is typically co-located with the SMF, or the AMF functions in the 3GPP 5G system architecture, and on WLAN controller in the case of Wi-Fi access architectures.

Node Location Database (NLDB)

A database that contains the location information of every mobile node that is part of the MFA domain and is currently attached to the network.

Network Topology Database (NTDB)

A database that contains the MFA node information along with the link state and directly connected neighbor information.

Home Network Prefix (HNP)

An IPv6 prefix assigned to the mobile node. This prefix is hosted by the MFA-MNA on the access link shared with the mobile node. The network will provide mobility support for the HNP prefixes. A meta-data tag indicating the mobility property [I-D.ietf-dmm-ondemand-mobility] is included in router advertisements and in address assignment related protocol messages.

Local Network Prefix (LNP)

An IPv6 prefix assigned to the mobile node. This prefix is hosted by the MFA-MNA on the access link shared with the mobile node. The network will not provide mobility support for the LNP prefixes. A meta-data tag indicating that there is no mobility support [I-D.ietf-dmm-ondemand-mobility] is included in router advertisements and in address assignment related protocol messages.

3. Overview

This specification describes the MFA protocol. The MFA protocol is designed for providing mobility management support to a mobile node without the need for a fixed IP anchor. In this approach the mobile node's IP session is always anchored on the first-hop router sharing the link with the mobile node. The entities in the MFA domain track the mobile node's movements in the MFA domain and will provision the forwarding states in the user-plane nodes for optimal routing and for ensuring the anchor is always the first-hop router. Any time the mobile node moves within the MFA domain and resulting in the mobile node's IP flows going through the previous anchor, the mobility entities detect this event and a corrective action is taken by provisioning the forwarding nodes with the path stitching rules. The result of this approach is an user-plane with no fixed anchors, and dynamically programmed user-plane for mobility and optimal packet routing.

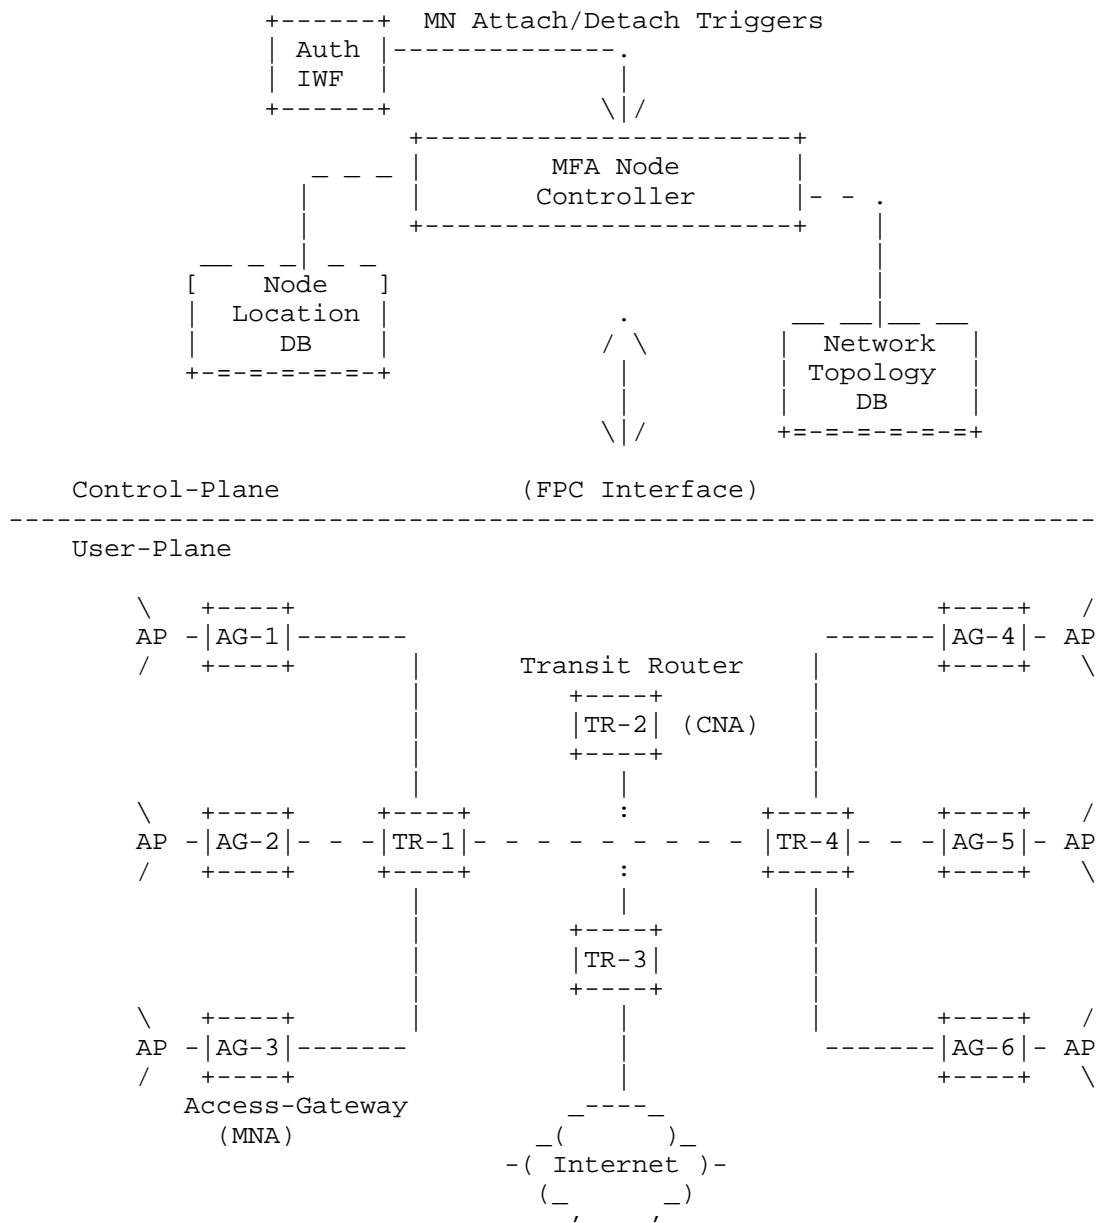
The following are the key functional entities in the MFA domain:

- o MFA Node Controller (MFA-NC)
- o MFA Mobile Node Anchor (MFA-MNA)

- o MFA Correspondent Node anchor (MFA-CNA)

The MFA-NC is typically collocated with the access network specific control-plane functions. It interfaces with the radio network/ authentication functions for detecting the mobile node's movements in the MFA domain for managing the forwarding states in the user-plane entities, MFA-MNA and MFA-CNA. The MFA node controller requires access to node location database and network topology database. These are the conceptual entities that can be realized using existing elements that are already present in different access architectures.

The MFA-MNA and the MFA-CNA are the functions in the user-plane network and they are collocated with the elements in the network that perform IP packet forwarding functions. The MFA-MNA is typically located on the first-hop router and whereas the MFA-CNA can be collocated with the access-gateways and transit routers. These entities interface with the MNA-NC using FPC ([I-D.ietf-dmm-fpc-cpdp]), or an alternative interface), for managing the IP forwarding policies.



* MFA-MNA is collocated with the access gateways

** MFA-CNA is collocated with the access gateways and transit routers

Figure 1: Example of a MFA Domain

3.1. The Network Topology Database

The network topology database contains the complete and the current information about all the MFA nodes in the network. The information includes the capabilities of each node, supported functions, supported interfaces with the interface-type, connected neighbors, hosted prefixes on each link, security configuration and other related configuration elements. The topology database can be used to determine the route between two nodes within the MFA domain, or the best exit gateway for reaching a correspondent node outside the MFA domain.

3.2. The Node Location Database

The node location database consists of location information of each mobile node that is currently attached to the MFA domain. It also includes the type of attachment, previous anchor, and other information elements, such as the mobile node's connection status and detailed or approximate location (e.g. tracking area) in case of device dormancy. Typically, the MFA entities obtain this information from the control-plane functions in the access network. For example, a WLAN controller and the authentication functions will be able to provide this information in IEEE 802.11 based networks. In 5G system architecture this information can be obtained from AMF/SMF functions.

Below diagram is an example NLDB database.

MN Identifier	Current Anchor	Previous Anchor	Handover Type
MN1@ietf.org	AG1	-	NEW_ATTACH
MN2@ietf.org	AG6	AG2	HANDOVER
MN3@ietf.org	-	AG4	UNKNOWN

Figure 2: Example NLDB Table

3.3. Determination of the Correspondent Node Anchor

The anchor for a correspondent node is a MFA node that is closest to the correspondent node and is in path for all the MN-CN IP traffic flows. The MFA node controller leverages the topology database for the CN-anchor determination.

If the correspondent node is another mobile node in the MFA domain, then the CN-Anchor for that correspondent node is the access gateway to which it is currently attached.

If the correspondent node is outside the MFA domain, then the CN-anchor is typically the exit gateway, or any MFA node that is always in path for reaching the CN's network. This is typically the PE router of the data center that hosts the correspondent node service, or a programmable data plane node inside the data center.

The below illustration is an example topology of a MFA domain. The domain consists of MFA nodes, mobile and correspondent nodes. A query for CN2's anchor should result in finding AG4, as that is the MFA node in the traffic path and closest to CN2. Similarly, the query for CN3's anchor which is outside the MFA domain should result in finding TR3 as that is the last exit gateway in the MFA domain and closest to the CN3.




Figure 3: CN Anchor Determination - Example Topology

3.4. Traffic Steering Approaches

The MFA nodes support traffic steering approaches for moving the mobile node's IP traffic between the MFA nodes over the most optimal routing path. Segment Routing for IPv6 (SRv6) is one approach that this specification focuses on for steering the traffic between two points in the network, whereas the MFA-NC can utilize the available information from Network Topology- and Node Location Database to enforce policies in the MFA nodes in support of alternative data

plane protocols to enable traffic steering. Future versions of the document may include information about additional mechanisms.

When using SRv6 for traffic steering, the approaches specified in [I-D.ietf-dmm-srv6-mobile-uplane] and [I-D.filsfils-spring-srv6-network-programming] will be leveraged for moving the mobile node's IP traffic between the MFA-MNA and the MFA-CNA nodes. The SRv6 policy including the SID information and the associated functions are pushed from the MFA Node controller to the MFA nodes. This document mostly leverages the functions specified in those documents, but may require some changes to the SRv6 functions for reporting the flow meta-data of the non-optimal traffic flows to the MFA node controller. The definitions of those SRv6 functions will be specified in either in the future revisions of this document, or in other IETF documents.

The following table captures the possible SRv6 function activation when IP traffic steering approach is in use. This is only an example.

FLOW DIRECTION	MN-Anchor	CN-Anchor
MN to CN	Variant of T.Insert (Transit with insertion of SRv6 policy and may require trigger to MFA-NC such as activation of Flow.Report)	Variant of End.X (Or, End.B6, instantiation of a binding SID); Or, End.T for internet traffic
CN to MN	Variant of End.X (Layer-3 cross connect (Or, End.B6, instantiation of a binding SID	Variant of T.Insert (Transit with insertion of SRv6 policy and may require trigger to MFA-NC such as activation of Flow.Report.

Figure 4: Using SRv6 for Traffic Steering - Example

3.5. Mobile Node Attachment Triggers

The MFA domain relies on the access network for certain key events related to the mobile node's movements in the network. These events include:

- o INITIAL_ATTACH - Initial Attachment of the mobile node to the MFA domain
- o HANDOVER - Layer-2/Layer-3 Handover of the mobile node within the MFA Domain
- o DETACH - Detachment of the mobile node from the MFA domain
- o UNKNOWN - State of the mobile node is Unknown; TBD

The MFA node controller interfaces with the radio network and the authentication infrastructure for these events. These events drive the policy configuration on the MFA nodes.

3.6. Programming the User-plane

The MFA-NC leverages suitable southbound semantics and operation to enforce traffic steering rules in the selected access gateways (AG) and/or transient routers (TR). One suitable data model and operation is being specified in [I-D.ietf-dmm-fpc-cpdp] for Forwarding Policy Configuration (FPC). The model and operation applies in between a FPC Client function and an FPC Agent function.

A deployment of FPC with the specification per this document about MFA, the FPC Client is co-located with the MFA-NC, whereas the FPC Agent function is co-located with functions that enforce user plane configuration per the rules received from the FPC Client. The FPC Agent can either reside on an transport network- or SDN controller and be in charge of the configuration of multiple user plane nodes (MFA-TR, MFA-MA, MFA-CA), or an FPC Agent resides on each MFA node.

The following figure schematically draws an example how FPC can integrate with the functional MFA architecture per this specification. The example assumes that MFA nodes can be programmatically configured by an SDN Controller. Details about whether a single or multiple distributed SDN Controllers are deployed are left out.

The FPC data model includes the following components:

Data Plane Nodes (DPN) Model:

Representation of nodes in the data plane which can be selected and enforce rules per the control plane's directives. DPNs take a particular role, which is identified in the model. In the context of this document, the role of a DPN can be, for example, an anchor node or a transit router.

Topology Model:

Representation of DPNs in the network and associate in between DPNs. The FPC Client and Agent use the Topology to select most appropriate data plane node resources for a communication. In the context of this document, Topology has can be leveraged to implement the NTDB for the selection of steering paths and associated DPNs which function as MFA-MNA, MFA-CNA, or MFA-TR.

Policy Model:

Defines and identifies rules for enforcement at DPNs.

Mobility-Context:

Holds information associated with a mobile node and its mobility sessions. In the context of this document, Mobility-Context can be enriched with traffic steering related rules.

Monitor:

Provides mechanisms to register monitors (traffic, events) in the data plane and define status reporting schedules, which can be periodic or event-based. In the context of this document, Monitors may be used to detect traffic from a CN to an MN on an MFA node, which could result in a notification to the MFA-NC for path optimization and associated steering of traffic to the MN's current MFA-MNA.

Please refer to [I-D.ietf-dmm-fpc-cpd] for model and operational details.

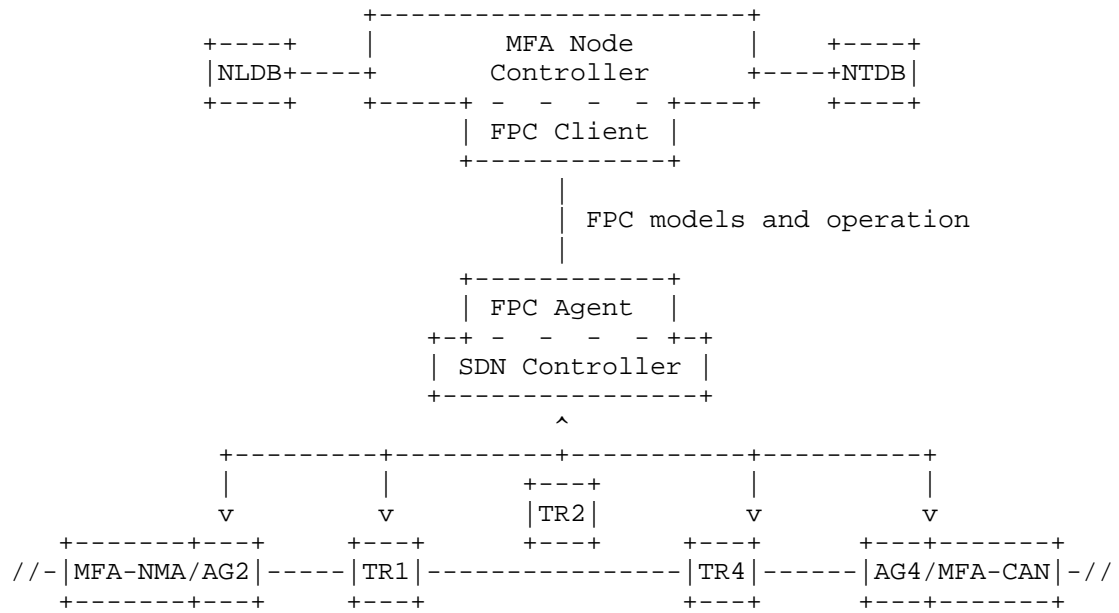


Figure 5: Deployment of the FPC models and operation in between the MFA-NC and MFA nodes on the user plane

4. Life of a Mobile Node in a MFA Domain

Reference Topology

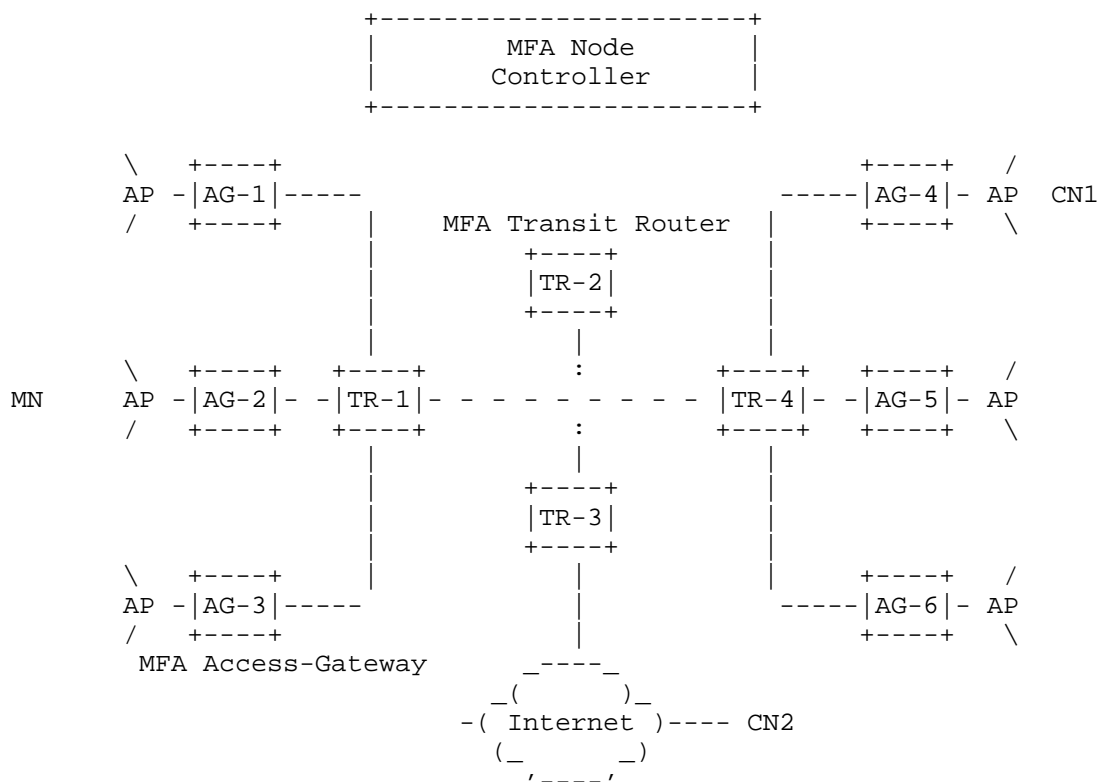


Figure 6: Reference Topology

4.1. MN's Initial Attachment to a MFA Domain

A mobile node, MN enters the MFA domain and attaches to the access point on the gateway AG-2.

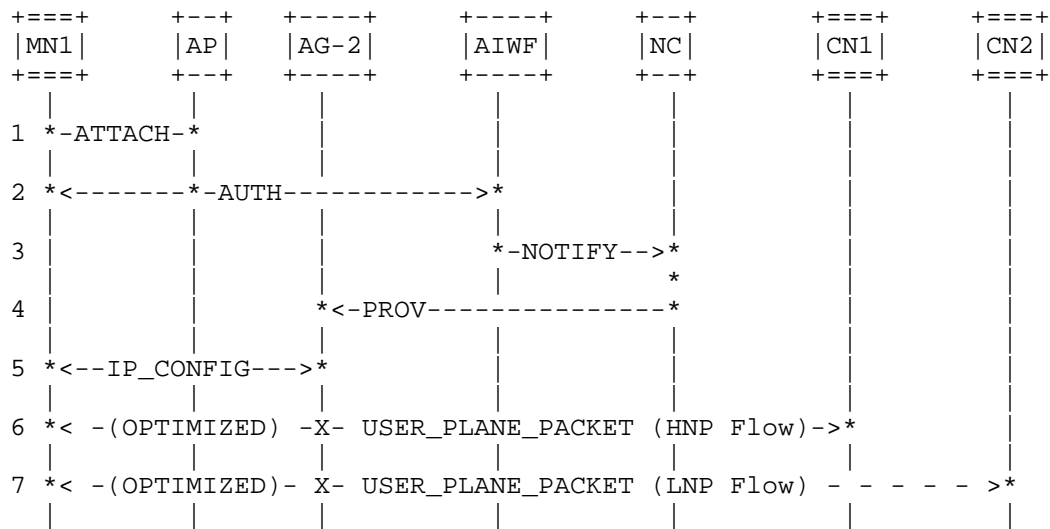


Figure 7: Mobile Node's Initial Attachment to a MFA Domain

- o 1-ATTACH: The mobile node with NAI (MN1@ietf.org) performs a layer-2 attach to the access point. This access point is connected to the access-gateway, AG-2, over a layer-2 link. The mobile node anchor function is supported on AG-2 and is active.
- o 2-AUTH: The mobile node completes the access authentication access technology specific access mechanisms. The mobile node's identity is established and is authorized for MFA domain access. The Authentication interworking (AUTH-IWK) function records the mobile node's identity, type of attach as `INITIAL_ATTACH`, and the current location of the mobile node in the access-network, to the node location database.
- o 3-NOTIFY: The Auth-IWK function delivers the attach event to the MFA node controller. The information elements that are delivered include the mobile node identifier (MN-1@ietf.org), type of attach as `INITIAL_ATTACH`, and the identity of the access gateway, which is AG-2.
- o 4-PROV: The NC provisions AG-2 for hosting the MN's home-network prefix(es). The assigned prefixes are HNP, H1::/64 and LNP, L1::/64. These prefixes are from a larger aggregate block (Ex: H1::/48; L1::/48) which are topologically anchored on AG-2. The policies for hosting the HNP prefixes on the link are provisioned using FPC interface. The AG-2 will include meta-data in the IPv6 RA messages for indicating the properties of the prefixes; H1::/64

as the prefix with mobility support and L1 as the prefix with no mobility support.

- o 5-IP_CONFIG: The mobile node generates one or more IPv6 addresses using the prefixes H1 and L1. The generated addresses are tagged with the property meta-data in the host's source address policy table. This allows the applications on the mobile node to pick the addresses based on the application's mobility requirements.
- o 6-USER_PLANE_PACKET: The mobile node establishes IP flow with CN1. The source address is based on the prefix H1. This IP address will have mobility support. The packets associated with this flow will take the optimized routing path. There are no tunnels, or special traffic steering rules in the network.
- o 7-USER_PLANE_PACKET: The mobile node establishes IP flow with CN2. The source address is based on the prefix L1. This IP address will not have mobility support. There are no tunnels, or special traffic steering rules in the network.

4.2. MN's Roaming within the MFA Domain

The mobile node roams and changes its point of attachment. It was initially attached to the access network on AG-2 and now it attaches to access network on AG-6. At the time of roaming, the mobile node had two active IPv6 prefixes HNP, H1::/64 and LNP, L1::/64 and there were two active IP flows, one to CN1 using an IPv6 address from the prefix H1::/64 and another flow to CN2 using an IPv6 address from the prefix L1::/64. The MFA network will ensure the prefix H1::/64 will be routable on the new network and the active flow to CN1 will survive, however the prefix L1::/64 will not be routable in the new access network and therefore the flow to CN2 will not survive.




Figure 8: Mobile Node's Roaming within the MFA Domain

- o 1-ATTACH: The mobile node with NAI (MN1@ietf.org) roams in the network from AG-2 to AG-6.
- o 2-AUTH: The mobile node completes the handover to the new access network using access network specific security mechanisms. The Auth-IWK function updates the mobile node's location in the node-location database. The updated entry in the node location database will include the mobile node's NAI, attach type as HANDOVER, and the current access-network location as AG-6.
- o 3-NOTIFY: The Auth-IWK function delivers the handover event to the MFA node controller. The information elements that are delivered include the mobile node identifier (MN-1@ietf.org), type of attach as HANDOVER, and the identity of the access gateway as AG-6.
- o 4-IP_PROV: The NC provisions AG-6 for hosting the MN's home-network prefix and local network prefix. The home network prefix, H1::/64 is from the previous anchor, AG-2 and is not topologically anchored on AG-6. However, for supporting mobility the prefix is hosted on the access link while the mobile node is attached to that access network and till there are active flows. The NC also

provisions AG-6 for hosting a new local network prefix, L2::/64. This prefix, L2::/64 is from a larger aggregate block that is topologically anchored on AG-6. The AG-6 will include meta-data in the IPv6 RA messages for indicating the properties of the prefixes; H1::/64 as the prefix with mobility support and L2::/64 as the prefix with no mobility support. The NC also provisions a traffic steering rule to steer all uplink IP traffic with source address H1::/64 through the previous anchor AG-2.

- o 5-IP_PROV: The NC provisions AG-2 to steer all IP traffic to destination addresses matching the prefix, H1::/64 to AG-6, and it also provisions a rule to report flow meta-data of those flows taking the non-optimal traffic path through AG-2. This essentially allows the NC to learn about any mobile node's IP flows still going through AG-2, so it can stitch the optimized path for those flows and remove AG-2 from the path for those flows.
- o 6-IP_CONFIG: The prefix H1::/64, obtained at the new location, will continue to be available on the new access link. The new local network prefix L2::/64 will also be available on the new access link and will be marked as a prefix with no mobility property. The mobile node may generate one, or more IPv6 addresses using the prefix L2::/64. The prefix L1::/64 is no longer hosted on the new link and the mobile node will remove it from interface configuration.
- o 7-USER_PLANE_PACKET: Any uplink IP link from CN1 will come to AG-2, as it's the topological anchor for that address/prefix and AG-2 will steer the traffic directly to AG-6. On detecting an IP flow with the IP address belonging to prefix H1::/64, AG-2 will report the CN1-MN1 flow meta-data to NC.
- o 8-Report: The NC on receiving this event will lookup the CN anchor for the flow in its node location database. If the CN is another MN within the MFA domain, its current anchor information is retrieved from the node location database. However, if the CN is a node outside the MFA domain, the anchor for this node can be any transit router in the MFA domain which is always in path for that destination. The CN-anchor determination for nodes outside the MFA domain will be based on the network topology database.
- o 9-FLOW_STEERING: The NC inserts a IP traffic steering rule on AG-6 to steer the MN1-CN1's IP flows using H1::/64 directly to CN1's anchor which is CN1-A, and bypassing AG-2.
- o 10-FLOW_STEERING: The NC inserts a IP traffic steering rule on CN1-A to steer the MN1-CN1 IP flows using H1::/64 directly to

MN1's current anchor which is AG-6, and bypassing AG-2.

- o 11-USER_PLANE_PACKET: The MN1-CN1's IP flows using H1::/64 will be steered directly from CN1-A to AG-6; AG-2 will not be in the path.

4.3. Traffic Steering State Removal

The mobile node's IP flows that were established at the previous location are no longer active. The steering state that was introduced at AG-6 and CN1-A will be removed on detecting the inactive flows. The network may also optionally choose to withdraw the prefix H1::/64 and may assign a new HNP prefix which are topologically anchored in the new location.

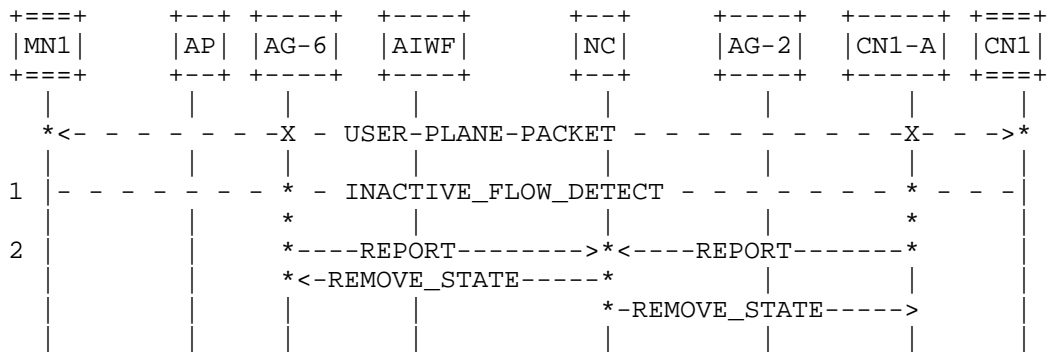


Figure 9: State Removal

- o 1-IN_ACTIVE_FLOW_DETECT: At some point the MN1-CN1 flow using the prefix H1::/64 is no longer active.
- o 2-REPORT: Both AG-6 and AG-2 will detect the inactive flows and may report this event to the NC. The steering state associated with MN1-CN1 flow using the prefix H1::/64 may be removed prior to reporting to the NC. Optionally, the NC on receiving the INACTIVE_FLOW_DETECT event may provision AG-6 and CN1-A to remove the steering state.
- o 4-REMOVE_STATE:

4.4. Mobile Node's new IP flows

The mobile node's IP flows that were established at the previous location are no longer active and any created steering state was removed. The network may optionally choose to withdraw the prefix H1::/64 and may assign a new HNP prefix which is topologically anchored in the new location. All new IP flows will use the new prefix and the flows will take optimal routing path.




Figure 10: New Flows

- o 1-USER_PLANE_PACKET: The mobile node's has established some IP flows using the IP address from the new HNP and LNP assigned at the new location. These IP flows will take optimal routing path and there is no need for any steering state, or the use of tunnels in the network for the mobile node's traffic.

5. MFA in 5G System Architecture

3GPP is specifying the 5G System Architecture, which follows a split between control- and data plane. Key control plane functions, which have interfaces to the data plane, are the Access Network and Mobility Management Function (AMF), and the Session Management Function (SMF). AMF and SMF cooperate to set up data plane nodes in the (radio) access network ((R)AN) and the core network, which comprises one or multiple User Plane Functions (UPF). As soon as a mobile node (UE) attaches to the network, as Packet Data Unit (PDU) Session is established and the SMF in the control plane selects one UPF as PDU Session Anchor, which serves also as IP address anchor. The SMF may select one more UPF on the path in between the PDU Session Anchor and the (R)AN, which enables routing traffic in between the UE and a local packet data network (PDN) with a correspondent node or service without the need to traverse the PDU Session Anchor.

In the view of MFA, each UPF can represent a locator for the UE's downlink traffic on the N9 as well as on the N6 reference point in

the 5G System Architecture. Since the SMF is in charge of UPF selection and configuration, the MFA-NC can leverage the SMF to retrieve node location information per this specification's procedure to access the NLDB from the MFA-NC. For MFA node selection and traffic steering, the MFA-NC may need more information about the data plane in terms of the transport network nodes and topology. Details about the NTDB are left out of this version of the document, but a realization may exploit available Topology information per [I-D.ietf-dmm-fpc-cpdp].

In the figure below, a UE's UPFs can function as MFA nodes, either as MFA-MNA or as MFA-CNA in case of mobile to mobile communication. Other transport network nodes, which may function as MFA-CNA for the UE's communication with a (non-mobile) correspondent node or service, are not explicitly depicted in the below figure. The MFA function can be tightly coupled with a UFP (co-located) or loosely coupled (separated). The MFA-NC utilized the FPC models and operation to enforce traffic steering policies in the MFA nodes. In case of loose coupling, the SMF utilizes the N4 protocol per the 3GPP standard to configure the selected UPF, whereas the MFA-NC uses FPC to enforce policies in the associated (loosely coupled) MFA node. In case of tight coupling, the MFA-NC may be co-located with the SMF and a single reference point and associated protocol may be used in between the SMF/MFA-NC and a UPF/MFA node.




Figure 11: New Flows

6. IANA Considerations

TBD

7. Security Considerations

This specification allows a mobility node controller to provision IP traffic steering policies on the user plane nodes. It essentially leverages the FPC interface [I-D.ietf-dmm-fpc-cdpd] for interfacing with the user-plane anchor nodes. The security considerations specified in the FPC specification are sufficient for securing the messages carried on this interface.

The traffic steering rules that are provisioned on the MFA nodes by the MFA node controller are the standard policy rules that the FPC interface defines and does not require any new security considerations.

8. Acknowledgements

TBD

9. References

9.1. 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>>.

9.2. Informative References

[I-D.filsfils-spring-srv6-network-programming]
Filsfils, C., Leddy, J., daniel.voyer@bell.ca, d., daniel.bernier@bell.ca, d., Steinberg, D., Raszuk, R., Matsushima, S., Lebrun, D., Decraene, B., Peirens, B., Salsano, S., Naik, G., Elmalky, H., Jonnalagadda, P., Sharif, M., Ayyangar, A., Mynam, S., Henderickx, W., Bashandy, A., Raza, K., Dukes, D., Clad, F., and P. Camarillo, "SRv6 Network Programming", draft-filsfils-spring-srv6-network-programming-03 (work in

progress), December 2017.

[I-D.ietf-dmm-fpc-cpdp]

Matsushima, S., Bertz, L., Liebsch, M., Gundavelli, S., Moses, D., and C. Perkins, "Protocol for Forwarding Policy Configuration (FPC) in DMM", draft-ietf-dmm-fpc-cpdp-09 (work in progress), October 2017.

[I-D.ietf-dmm-ondemand-mobility]

Yegin, A., Moses, D., Kweon, K., Lee, J., Park, J., and S. Jeon, "On Demand Mobility Management", draft-ietf-dmm-ondemand-mobility-13 (work in progress), January 2018.

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

Matsushima, S., Filsfils, C., Kohno, M., daniel.voyer@bell.ca, d., and C. Perkins, "Segment Routing IPv6 for Mobile User-Plane", draft-ietf-dmm-srv6-mobile-uplane-00 (work in progress), November 2017.

[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>>.

[RFC6275] Perkins, C., Ed., Johnson, D., and J. Arkko, "Mobility Support in IPv6", RFC 6275, DOI 10.17487/RFC6275, July 2011, <<https://www.rfc-editor.org/info/rfc6275>>.

Authors' Addresses

Sri Gundavelli
Cisco
170 West Tasman Drive
San Jose, CA 95134
USA

Email: sgundave@cisco.com

Marco Liebsch
NEC
Kurfuersten-Anlage 36
D-69115 Heidelberg,
Germany

Email: liebsch@neclab.eu

Satoru Matsushima
SoftBank
Tokyo,
Japan

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

