DMM Working Group
Internet-Draft
Intended status: Standards Track
Expires: April 25, 2014

M. Liebsch
NEC
P. Seite
Orange-France Telecom
G. Karagiannis
University of Twente
October 22, 2013

Distributed Mobility Management - Framework & Analysis
draft-liebsch-dmm-framework-analysis-02.txt

Abstract

Mobile operators consider the distribution of mobility anchors to enable offloading some traffic from their core network. The Distributed Mobility Management (DMM) Working Group is investigating the impact of decentralized mobility management to existing protocol solutions, while taking into account well defined requirements, which are to be met by a future solution. This document discusses DMM using a functional framework. Functional Entities to support DMM as well as reference points between these Functional Entities are introduced and described. The described functional framework allows distribution and co-location of Functional Entities and build a DMM architecture that matches the architecture of available protocols. Such methodology eases the analysis of best current practices with regard to functional and protocol gaps.

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 April 25, 2014.

Copyright Notice

Table of Contents

1. Introduction

The concept of Distributed Mobility Management (DMM) is based on the distribution of mobility anchors towards the access networks to provide mobile nodes with local anchors and enable optimized routing of traffic above anchor level to any kind of serving point, e.g. distributed content caches. The closer mobility anchors are located to mobile nodes, the more a mobile node's handover may necessitate the assignment of a new mobility anchor. Continuity of a mobile node's IP address or IP address prefix enables IP session continuity, but creates the problem of routing downlink packets to the mobile node's current mobility anchor. Different solutions and associated extensions to IP mobility management protocols are being discussed to maintain a mobile node's IP session after mobility anchor relocation, including solutions that are based on existing protocols.

This document defines a functional framework for DMM and describes an initial set of well defined functional entities (FE), which are required to support IP address continuity in a network with distributed mobility anchors. Having identified the function of each FE as well as required interfaces between FEs allows different constellations of FEs, either by co-locating or distributing them. Functional frameworks have been successfully used within and outside of the IETF, such as the ITU-T [ITU-TY2018][ITU-TY2804], to support the thorough analysis of protocols gaps with existing protocols and to enable the design of suitable solutions. Due to the complexity of the DMM problem and solutions space, we consider such framework of particular importance for performing a Gap Analysis while assigning the defined FEs to architecture components of existing protocols and to build suitable solutions for DMM based on extensions to a single or multiple existing protocols and architecture components.

This version of the draft introduces a basic set of FEs and interfaces between these FEs to support IP address continuity in DMM, without being specific to the used mobility management protocol, which operates below the mobility anchor. The functional framework as per this draft is protocol agnostic, such that it can apply to (1) solutions that are solely based on existing IP mobility protocols and to (2) solutions which get support from non-mobility protocols.

The framework enables the analysis of existing protocols' suitability to support DMM and allows building optimized solutions for DMM without being limited to the mobility protocol suites. Some examples how the framework can support the identification of required protocol extensions to existing mobility management protocol or alternatively the support from non-mobility protocols to design suitable DMM solutions on system level are described in Appendix A.

Appendix B defines an additional set of functional entities, which enables multicast support in DMM and can complement the framework for DMM unicast support.

2. Conventions and Terminology

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 [RFC2119].

3. Functional Architecture for DMM Support

The framework introduces four additional functional entities (FE) which are relevant complement existing mobility- and transport networks to enable DMM support for unicast traffic and to meet essential DMM requirements as per [I-D.ietf-dmm-requirements], such as enabling temporary IP address continuity after a mobile node got assigned a new mobility anchor. Further FEs may be needed to enable advanced features, such as simultaneous use of an imported mobile node HoA or HNP to maintain ongoing data sessions and a new HoA or HNP, which is allocated by the mobile node's new mobility anchor after handover. Additional FEs are not considered in this revision of the draft, but can be introduced easily in future versions of the draft and considered for the BCP discussion and gap analysis.

The following FEs are currently considered as existing functional entities to build the mobility- and transport network:

- FE_R: Functional Entity of a standard IP Router / Switch
- FE_MA_C: Functional Entity Mobility Anchor, Control Plane
- FE_MA_U: Functional Entity Mobility Anchor, User Plane
- FE_MU_C: Functional Entity Mobile User Client, Control Plane
- FE_MU_U: Functional Entity Mobile User Client, User Plane

The list comprises a generic router/switch function FE_R that's supposed to build the transport network. It has no particular function that's specific to DMM, but performs routing according to a longest prefix match. Deployment specific aspects, such as the use of IP/MPLS, are not (yet) considered in this draft.

The entities FE_MA_C and FE_MA_U represent the unmodified functions of the mobility architecture's mobility anchor. In Mobile IPv6, these function would be co-located with the Home Agent, in Proxy Mobile IPv6, these functions would be co-located with the Local Mobility Anchor (LMA). In a cellular IP (CIP) enabled domain, these functions would be co-located with the domain's CIP Gateway.

The entities FE_MU_C and FE_MU_U represent the existing user client functions, that send location updates to the mobility anchor. In Mobile IPv6, these functions are co-located with the Mobile Node, whereas in Proxy Mobile IPv6, these functions are co-located with the Mobile Access Gateway.

So far, this draft defines four DMM-specific FEs, which can be either

distributed or co-located with existing FEs of the mobility- or routing plane. One or more of the following FEs are currently assumed to add to an existing mobility- and transport network to enable DMM support for IP address continuity:

- FE_MCTX: Functional Entity Mobility Context Transfer
- FE_I: Functional Entity Ingress to DMM plane
- FE_E: Functional Entity Egress of DMM plane
- FE_IEC: Functional Entity for Ingress/Egress Control

Note: No all FEs or reference points between FEs may be relevant for a DMM-enabled solution that is based on existing protocols and the associated architecture. Which functions are relevant to complement an existing protocol and architecture depends on the identified gaps.

The task of the FE_MCTX is to export relevant binding cache information, such as the mobile node's HoA or HNP, from the mobile node's previous mobility anchor (pMA) during mobility anchor relocation to enable IP address continuity after mobility anchor relocation. Furthermore, the function allows importing mobility context on the mobile node's new mobility anchor. Imported HoA/HNP of a mobile node will be treated as identifier and non-routable IP address (prefix), as it probably does not match the new mobility anchor's location in the topology. Furthermore, the FE_MCTX can provide mobility context to the FE_IEC to allow keeping these policies updated, which allow forwarding of packets to the MN's currently used mobility anchor.

The function FE_I enables the ingress level of indirection by means of deviating from the standard routing path of the mobile node's downlink packets, which carry the mobile node's HoA/HNP in the destination IP address field of their IP header. The FE_I can retrieve information from a control function (FE_IEC) to establish forwarding of the mobile node's packets to the appropriate DMM egress function (FE_E). Forwarding can be for example accomplished by an IP tunnel to the egress function, address translation to a routable IP address or other means.

The function FE_E receives downlink packets being forwarded by the DMM ingress function FE_I, e.g. by terminating a forwarding tunnel. The state on the FE_I can be established through the DMM ingress/ egress control function (FE_IEC) and is used to identify an MN's received packets and deliver them to the MN's current mobility anchor (FE_MA). If the FE_E is co-located with the FE_MA, the delivery is a local operation. If the FE_E is not co-located with the FE_MA, other

techniques, such as host-routes or technology such as OpenFlow may be used to deliver the packets to the mobile node's current mobility anchor. If not co-located with the FE_MA, the FE_E is supposed to be located close to the mobile node's current FE_MA.

The function FE_IEC represents a control function, that establishes, updates and removes policies (per-host or grouped) in the FE_I and the FE_E to allow forwarding of a mobile node's downlink packets towards the mobile node's current mobility anchor.

The mobile node's IP address (prefix) is carried in the source address field of the uplink packet. This source address is thus topologically incorrect after mobile node's handover. When IP routers of the mobility domain do not apply filtering according to the source addresses, uplink packets can be assumed to be routable and no specific operation is required.

If source address filtering is used, relevant routers need to be reconfigured to exclude the mobile node's IP address from filtering rules. If such filtering is performed by a mobility anchor or a Proxy Mobile IPv6 Mobile Access Gateway (MAG), local mobility functions on these routers should perform the task to reconfigure the local filter rules for uplink traffic.

When traffic indirection also applies to the uplink, e.g. to enable bidirectional tunneling to ensure that downlink and uplink data packets always traverse the same ingress/egress functions, FE_E and FE_I functions come into play on the uplink path. Downlink FE_I and FE_E become respectively FE_E and FE_I on the uplink. The uplink FE_I forwards a mobile node's packets to the FE_E corresponding to the downlink FE_I that has sent packets with the mobile node's address in the destination address field. The FE_I can also retrieve the information from the FE_IEC.

Figure 1 illustrates how the four DMM-specific FEs complement existing FEs of the mobility architecture. These DMM-specific FEs and associated operation on the interfaces between them can be realized by existing protocols, extensions to them or new protocols. Figure 1 separates the data plane from the control plane.

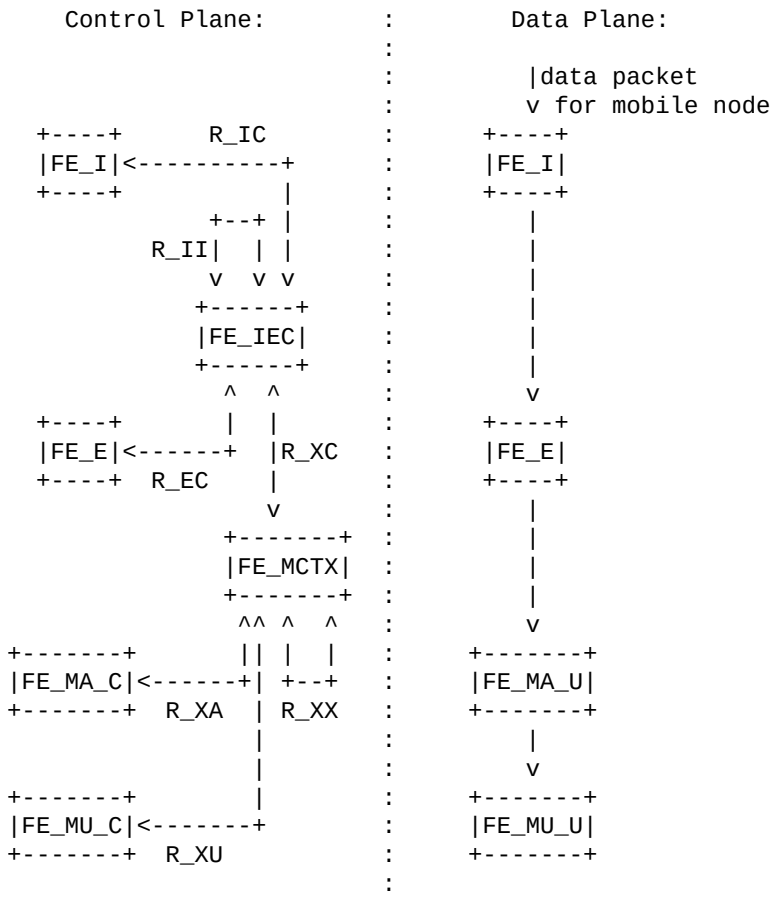

Figure 1: Basic set of functional entities (FE) and interfaces to enable IP-address continuity in DMM

The reference points between FEs comprise the following features:

- R_XA: Enables the FE_MCTX to retrieve mobility context information from the FE_MA of the MN's mobility anchor. Such information includes for example the MN's Home Address (HoA) or Home Network Prefix (HNP). In the network of the MN's new mobility anchor, the reference point enables the FE_MCTX to provide the MN's mobility context to the associated FE_MA, that imports the MN's mobility context to enable IP address continuity.

- R_XU: Enables the FE_MCTX to retrieve mobility context information from the mobile user client control function, the FE_MU_C. In host

mobility management, this function is located on the Mobile Node, who could support DMM operation by notifying the FE_MCTX through this reference point.

- R_XX: Enables the direct transfer of an MN's mobility context between two functions FE_MCTX, which are typically located in the network of the MN's previous and new mobility anchor respectively.

- R_IC: Enables the FE_IEC to provide policies to the FE_I, which are used to forward the MN's downlink packets towards the MN's new mobility anchor and the associated FE_E. These policies can be provided to the FE_I in an unsolicited manner or on request by the FE_I.

- R_EC: Enables the FE_IEC to provide policies to the FE_E, which are used at the FE_E to identify received packets that belong to a particular MN and deliver these packets to the MN's new mobility anchor. Such policies could include, for example, tunnel endpoint information, flow identification rules or other identification and addressing rules.

- R_XC: Enables initialization and update of the FE_IEC about the MN's mobility context as well as about its current location as represented by the FE_E in the network of the MN's current mobility anchor.

- R_II: Multiple instances of an FE_IEC can be deployed to build a DMM architecture, e.g. to distribute load and scale better, or distribute tasks associated with the FE_IEC to enable cooperative solutions.

4. Different Constellations of Functional Entities

The defined FEs can be grouped or distributed to build a DMM architecture that considers new architecture components or that is based on components of existing protocols. As a starting point, this section depicts and describes two deployment variants, which reflect the current understanding of the WG how DMM could be accomplished using existing protocol specifications as base. Variants of these two deployment models or entirely new models are possible and can be added to future versions of this document.

Note: This section is incomplete and needs further input on different deployment models and variants.

4.1. Condensed Deployment: Mobility Protocol Centric Solutions

Mobility protocol centric solutions aim at extensions to available mobility protocols to enable DMM, without being dependent on any external, non-mobility component and protocol. IP address continuity is typically established on the control plane by extensions to the mobility protocol to convey an MN's mobility context to a new mobility anchor, and on the data plane by the establishment of a forwarding tunnel between mobility anchors to deliver downlink packets from the originally assigned mobility anchor to the MN's currently used mobility anchor after anchor relocation. Alternatively, IP address continuity is enabled by using multiple mobility anchors simultaneously, whereas the mobile node's IP address(es) remain anchored at the topologically correct anchor point. These approaches differ in the level of extensions to the mobility protocols and in the support of certain features on the mobile node, such as the simultaneous use of multiple mobility anchors and associated Home Addresses. They have in common the sub-optimal routing path, as the mobile node's downlink traffic needs to traverse the location of the IP addresses topologically correct mobility anchor.

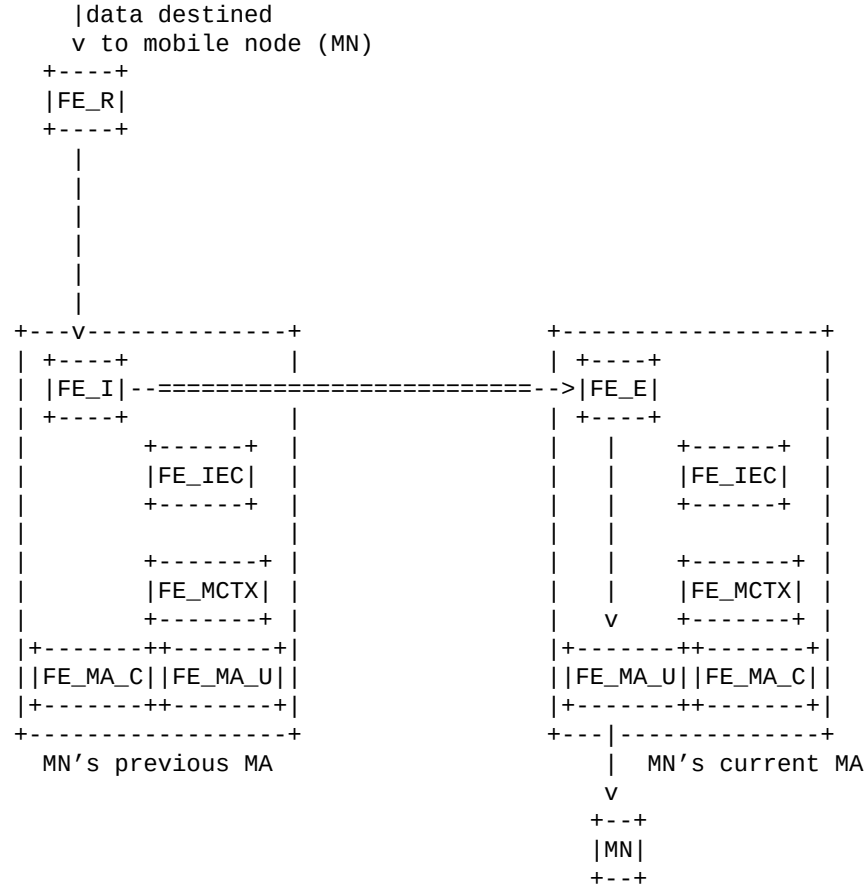

Figure 2: Condensed Deployment: Mobility Protocol Centric Solutions

4.2. Cooperative Deployment: Distributed Architecture

A distributed architecture considers protocol operation between distributed FEs, aiming at a DMM solution that's to a large extent independent of the mobility architecture and protocol. A further goal is to establish optimal routing paths for the MN's traffic after the MN's mobility anchor has been relocated and IP address continuity must be provided.

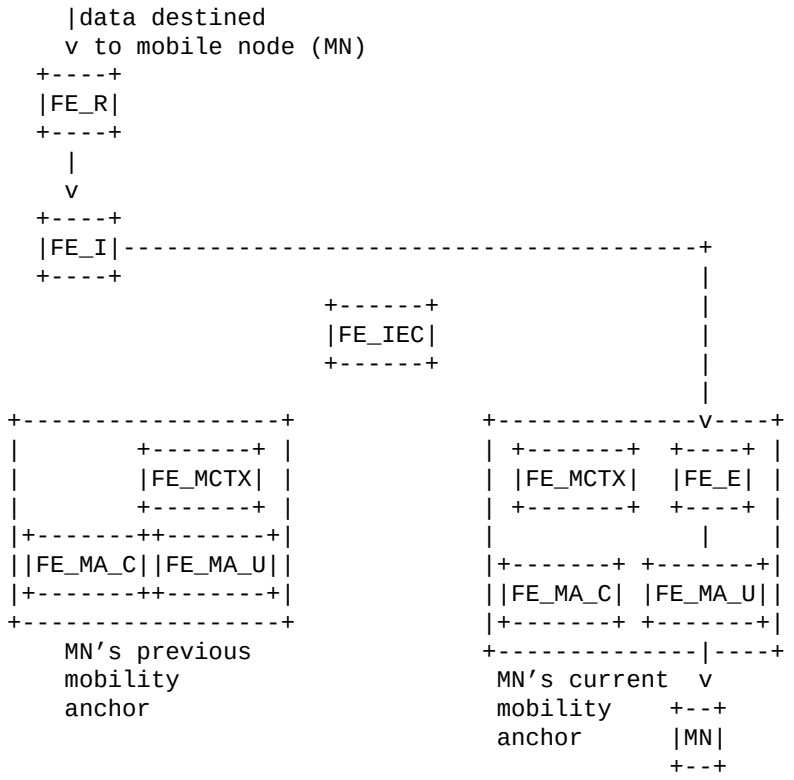

Figure 3: Cooperative Deployment: Distributed Architecture

5. Security Considerations

Different constellations of Functional Entities may allow re-use of existing protocols' security mechanisms to protect DMM protocol operation. In particular in a distributed model, new interfaces must be protected, e.g. to counteract unauthorized packet redirection to a different, possibly malicious mobility anchor. Details about security threats will be studied when the placement of Functional Entities for a selected set of preferred deployment models becomes mature.

6. IANA Considerations

As this document represents a framework and no protocol specification, there is no need for IANA actions.

7. References

7.1. Normative References

[I-D.ietf-dmm-requirements]
Chan, A., Liu, D., Seite, P., Yokota, H., and J. Korhonen, "Requirements for Distributed Mobility Management", draft-ietf-dmm-requirements-09 (work in progress), September 2013.

[RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, March 1997.

[RFC4601] Fenner, B., Handley, M., Holbrook, H., and I. Kouvelas, "Protocol Independent Multicast - Sparse Mode (PIM-SM): Protocol Specification (Revised)", RFC 4601, August 2006.

[RFC4605] Fenner, B., He, H., Haberman, B., and H. Sandick, "Internet Group Management Protocol (IGMP) / Multicast Listener Discovery (MLD)-Based Multicast Forwarding ("IGMP/MLD Proxying")", RFC 4605, August 2006.

[RFC6224] Schmidt, T., Waehlisch, M., and S. Krishnan, "Base Deployment for Multicast Listener Support in Proxy Mobile IPv6 (PMIPv6) Domains", RFC 6224, April 2011.

7.2. Informative References

[ITU-TY2018]
"ITU-T Y.2018, Mobility management and control framework and architecture within the NGN transport stratum".

[ITU-TY2804]
"ITU-T Q.1707/Y.2804, Generic framework of mobility management for next generation networks".

Appendix A. How the framework can support a gap analysis! Some examples..

A Gap analysis can be performed according to different deployment models and variants as summarized in Section 4. A suitable set of DMM FEs can be mapped to the architecture of existing protocols from within or beyond the IP mobility protocol solution space to analyze and identify gaps in the chosen protocols to support and optimize DMM operations. This section provides a few examples about the mapping of DMM FEs to mobility protocol FEs and non-mobility protocol FEs. Common goal is to enable DMM support, either in a mobility protocol centric manner or by means of a distributed architecture, relying on the support and associated collaboration with non-mobility protocol functions, such as routing. As examples for the distributed architecture, the Locator-Identifier Split Protocol (LISP) and the iBGP have been used to enable traffic indirection in the routing plane above the topological level of distributed mobility anchors.

A.1. Condensed Deployment using Mobile IPv6

Note: A detailed example needs to be added in a next revision.

Description: Framework mapping to existing Mobile IPv6 architecture. Technical approach is the establishment of a forwarding tunnel between previous HA and new HA to enable IP address continuity after anchor relocation. Approach is the identification of missing protocol functions in Mobile IPv6 as expected from DMM functional entities as per this specification to enable full DMM support.

A.2. Condensed Deployment using Proxy Mobile IPv6

Note: A detailed example needs to be added in a next revision.

Description: Framework mapping to existing Proxy Mobile IPv6 architecture. Technical approach is the establishment of a forwarding tunnel between previous LMA and new LMA to enable IP address continuity after anchor relocation. Approach is the identification of missing protocol functions in Proxy Mobile IPv6 as expected from DMM functional entities as per this specification to enable full DMM support.

A.3. Cooperative Deployment using LISP

This example utilizes LISP Tunnel Ingress Routers (TIR) to perform the LISP map and encap procedure and tunnel packets to the mobile node’s current mobility anchor (Figure 4). The mobile node’s IP address is assumed routable above TIR level. TIRs can be for example deployed close to a mobile operator’s IXP or close to operator-owned

traffic sources, such as a mobile Content Delivery Network (CDN). A TIR, which receives data packets destined to the mobile node, can consult the LISP Mapping Database (DB) to resolve the mobile node's IP address into its current locator, which is the mobile node's currently used mobility anchor. The mobility anchor has to terminate the LISP tunnel at the Tunnel Egress Router (TER) function and forward the data packets to the mobile node's current location according to the utilized mobility management protocol. An identified gap in a setup with LISP is the dynamic update of the Mapping Database and the update of already established states in TIRs in case the mobile node's location has changed from one mobility anchor to another mobility anchor.

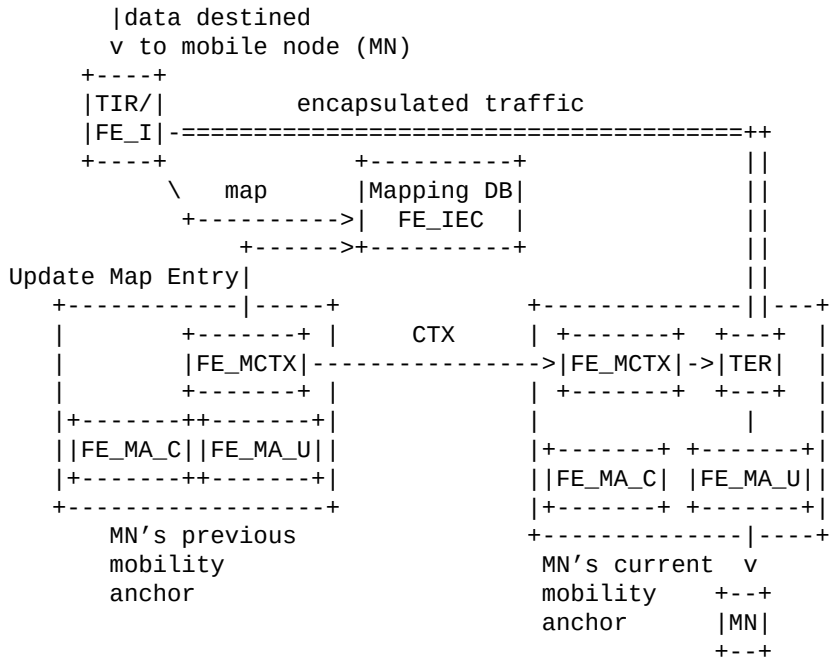

Figure 4: Example: DMM indirection at LISP TIRs

A.4. Cooperative Deployment using iBGP

This example utilizes the iBGP to establish per-host or group states in iBGP routers and forward a mobile node's packets hop-by-hop to its currently used mobility anchor. Figure 5 depicts an iBGP router with co-located FE_E and FE_I to receive data packets and to forward these packets to the next hop according to the routing state as per iBGP

update. The FE_IEC can be represented by the iBGP component to enable the setup of distributed routing states in distributed iBGP routers to direct the mobile node's data packets to its current mobility anchor. Hence, the FE_IEC is distributed in all iBGP routers to collaborate in the setup of host routes. The mobility anchor itself must implement iBGP to contribute to the distribution and update of host routes, e.g. after the mobile node changed its mobility anchor while IP address continuity must be supported. Since iBGP has been designed to propagate routing states to distributed routers, only minor protocol gaps may be identified in a detailed analysis. Beyond protocol gaps, further aspects need to be analyzed in such setup, which include limitations in scalability and route update latency.

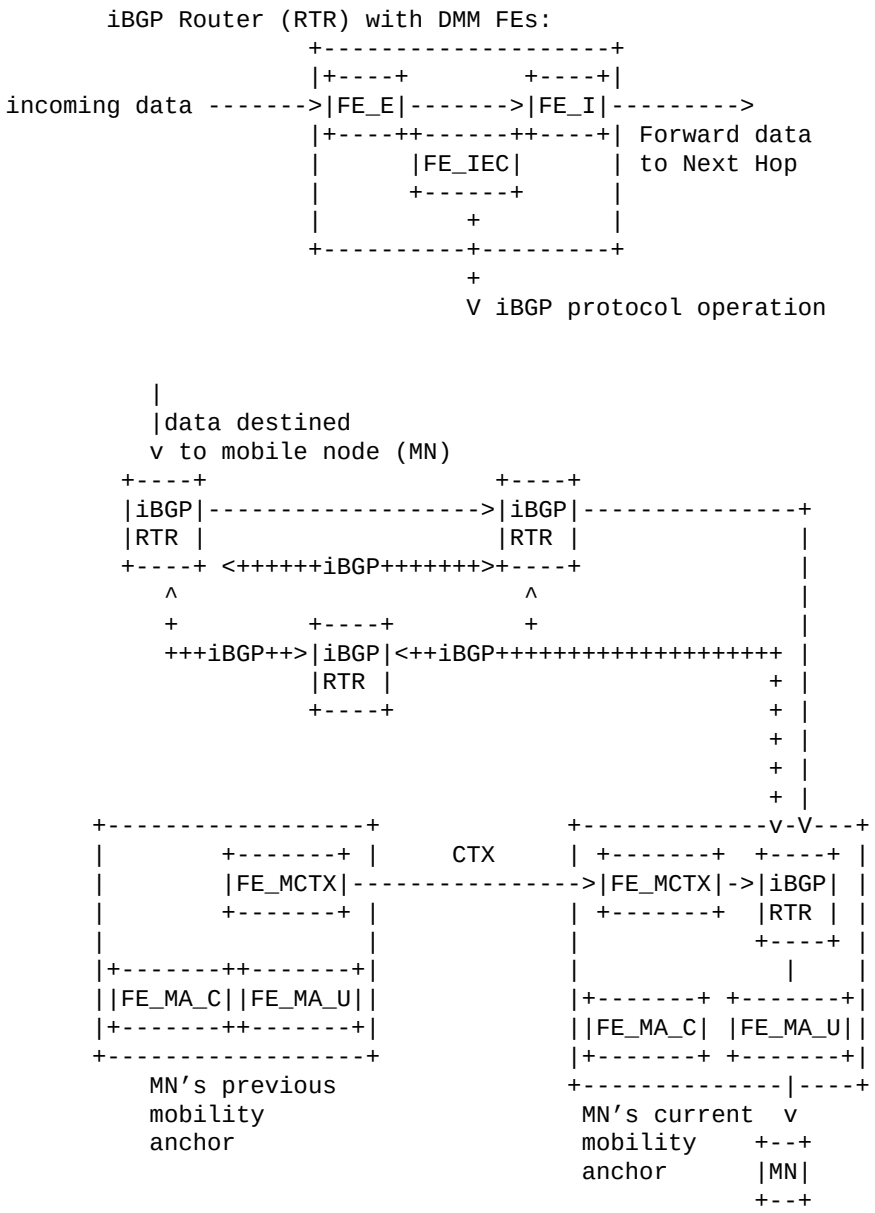

Figure 5: Example: DMM indirection at iBGP routers

Appendix B. Functional Architecture for Multicast DMM Support

The framework for multicast DMM support is similar to the framework for unicast DMM support introduced in Section 3 with the main difference that the additional introduced features are needed to support the multicast control and user plane. This framework, similar to the one introduced in Section 3, introduces four DMM-specific, with the main difference that these FEs are able to support multicast traffic, instead of unicast traffic. Additional FEs might be needed but are not considered in this revision of the draft, but can be introduced easily in future versions of the draft and considered for the BCP discussion and gap analysis.

The following FEs are currently considered as existing multicast based Functional entities to build the mobility- and transport network:

- FE_MR: Functional Entity of a standard Multicast IP Router / Switch. This FE can be incorporated to support the functionality of a Rendezvous Point (RP) and of a Designated Router (DR), see e.g., [RFC4601].

- FE_MLD-P: Functional Entity of a standard Multicast Listener Discovery Proxy (MLSD-P) used to provide MLD based forwarding, following the operation defined in e.g., [RFC4605] and [RFC6224].

- FE_MA_C_M: Functional Entity Mobility Anchor, Control Plane, for the support of multicast traffic

- FE_MA_U_M: Functional Entity Mobility Anchor, User Plane, for the support of multicast traffic

- FE_MU_C: Functional Entity Mobile User Client, Control Plane, for the support of unicast and multicast traffic. In case of multicast traffic the FE_MU_C can operate as multicast sender and multicast listener.

- FE_MU_U: Functional Entity Mobile User Client, User Plane, for the support of unicast and multicast traffic. In case of multicast traffic the FE_MU_U can operate as multicast sender and multicast listener.

The four DMM-specific FEs used to support multicast traffic are listed below.

- FE_MCTX_M: Functional Entity Mobility Context Transfer, used for the support of multicast traffic.

- FE_I_M: Functional Entity Ingress to DMM plane, used for the support of multicast traffic.

- FE_E_M: Functional Entity Egress of DMM plane, used for the support of multicast traffic.

- FE_IEC_M: Functional Entity for Ingress/Egress Control, used for the support of multicast traffic.

These FEs support similar features as the ones supported by the FE_MCTX, FE_I, FE_E, FE_IEC FEs, respectively, described in Section 3, with the main difference that they are used for the support of the multicast control and user planes, instead of the unicast control and user planes.

Figure 6 depicts the basic set of functional entities (FE) and interfaces to enable IP-address continuity in multicast based DMM. The four DMM-specific FEs and their associated operation on the interfaces between them can be realized by existing protocols, extensions to them or new protocols.

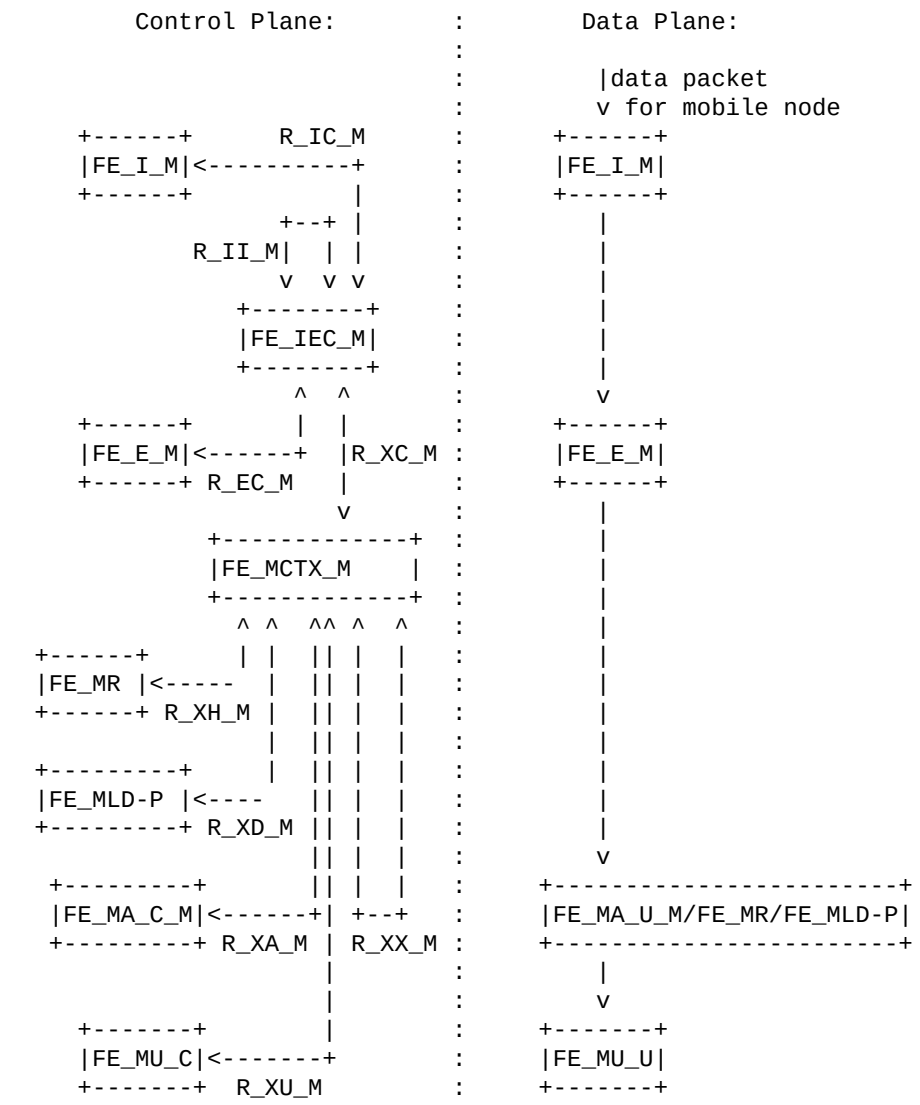

Figure 6: Basic set of functional entities (FE) and interfaces to enable IP-address continuity in multicast based DMM

The reference points between FEs are shown in Figure 6. In particular the features comprised by the reference points R_XA_M, R_XU_M, R_XX_M, R_IC_M, R_EC_M, R_XC_M, R_II_M, are similar to the ones supported by the reference points R_XA, R_XU, R_XX, R_IC, R_EC, R_XC, R_II, respectively, described in Section 3, with the difference that they are used to support the multicast based control plane,

instead of supporting the unicast based control plane.

Two additional reference points are added that are comprising the following features:

- R_XH_M: Enables the FE_MCTX_M to retrieve MR routing based information from FE_MR following the operation defined in e.g., [RFC4601].

- R_XD_M: Enables the FE_MCTX_M to retrieve Multicast Listener Discovery forwarding information from FE_MLD-P following the operation defined in e.g., [RFC4605] and [RFC6224].

Appendix C. Change Notes

Changes in version 01:

- Introduced functional split between existing Mobility Anchor Control- and User-Plane
- Introduced functional split of existing mobile user client Control- and User-Plane
- Added uplink routing considerations in DMM architecture
- Description of a first DMM Multicast framework in the Appendix
- Added examples to the appendix about how to use the framework for a gap analysis and for the design of optimized DMM solutions

Authors' Addresses

Marco Liebsch
NEC Laboratories Europe
NEC Europe Ltd.
Kurfuersten-Anlage 36
D-69115 Heidelberg,
Germany

Phone: +49 6221 4342146
Email: liebsch@neclab.eu

Pierrick Seite
Orange-France Telecom
4, rue du Clos Courtel, BP 91226
Cesson-Sevigne, 35512
France

Phone:
Email: pierrick.seite@orange-ftgroup.com

Georgios Karagiannis
University of Twente
AE Enschede, 7500
Netherlands

Phone: +31 53 4894099
Email: karagian@cs.utwente.nl