DMM
Internet-Draft
Intended status: Informational
Expires: January 9, 2017

H. Chan
X. Wei
Huawei Technologies
J. Lee
Sangmyung University
S. Jeon
Instituto de Telecomunicacoes
F. Templin
Boeing Research and Technology

July 8, 2016

Distributed Mobility Anchoring
draft-chan-dmm-distributed-mobility-anchoring-08

Abstract

This document defines distributed mobility anchoring. Multiple anchors and nodes are configured with appropriate mobility functions and work together to enable mobility solutions. Example solution is mid-session switching of the IP prefix anchor. Without ongoing session requiring session continuity, a flow can be started or re-started using the new IP prefix which is allocated from the new network and is therefore anchored to the new network. With ongoing session, the anchoring of the prior IP prefix may be relocated to the new network to enable session continuity.

Table of Contents

1. Introduction

A key requirement in distributed mobility management [RFC7333] is to enable traffic to avoid traversing single mobility anchor far from the optimal route. Distributed mobility management solutions do not make use of centrally deployed mobility anchor [Paper-Distributed.Mobility]. As such, the traffic of a flow SHOULD be able to change from traversing one mobility anchor to traversing another mobility anchor as the mobile node moves, or when changing operation and management requirements call for mobility anchor switching, thus avoiding non-optimal routes. This draft proposes distributed mobility anchoring to enable making such route changes.

Distributed mobility anchoring employs multiple anchors in the data plane. In general, the control plane function may be separate from the data plane functions and be centralized but may also co-located with the data plane function at these distributed anchors. Different configurations (Section 3.1) of distributed anchoring are then possible. Yet the distributed anchors need to have expected behaviors (Section 3.2).

A mobile node (MN) attached to an access router of a network may be allocated an IP prefix which is anchored to that router. It may then use the IP address configured from this prefix as the source IP address to run a flow with its correspondent node (CN). When there are multiple anchors, the flow may need to select the anchor when it is initiated (Section 4). Using an anchor in MN’s network of attachment has the advantage that the packets can simply be forwarded according to the forwarding table. Although the anchor is in the MN’s network of attachment when the flow was initiated, the MN may later move to another network, so that the IP address no longer belongs to the new network of attachment of the MN. Whether the flow needs session continuity will determine how to ensure that the IP address of the flow will be anchored to the new network of attachment. If the ongoing IP flow can cope with an IP prefix/ address change, the flow can be reiniated with a new IP address anchored in the new network (Section 4.1.1). On the other hand, if the ongoing IP flow cannot cope with such change, the IP address anchoring can be moved from the original network to the new network (Section 4.2).

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].

All general mobility-related terms and their acronyms used in this document are to be interpreted as defined in the Mobile IPv6 base specification [RFC6275], the Proxy Mobile IPv6 specification [RFC5213], and the DMM current practices and gap analysis [RFC7429]. This includes terms such as mobile node (MN), correspondent node

(CN), home agent (HA), home address (HoA), care-of-address (CoA), local mobility anchor (LMA), and mobile access gateway (MAG).

In addition, this document uses the following term:

Home network of an application session (or of an HoA): the network that has allocated the IP address (HoA) used for the session identifier by the application running in an MN. An MN may be running multiple application sessions, and each of these sessions can have a different home network.

IP prefix/address anchoring: An IP prefix, i.e., Home Network Prefix (HNP), or address, i.e., Home Address (HoA), allocated to a mobile node is topologically anchored to a node when the anchor node is able to advertise a connected route into the routing infrastructure for the allocated IP prefix.

Internetwork Location Management (LM) function: managing and keeping track of the internetwork location of an MN. The location information may be a binding of the IP advertised address/prefix, e.g., HoA or HNP, to the IP routing address of the MN or of a node that can forward packets destined to the MN. It is a control plane function.

In a client-server protocol model, location query and update messages may be exchanged between a Location Management client (LMc) and a Location Management server (LMs).

With separation of control plane and data plane, the LM function is in the control plane. It may be a logical function at the control plane node, control plane anchor, or mobility controller.

It may be distributed or centralized.

Forwarding Management (FM) function: packet interception and forwarding to/from the IP address/prefix assigned to the MN, based on the internetwork location information, either to the destination or to some other network element that knows how to forward the packets to their destination.

This function may be used to achieve indirection. With separation of control plane and data plane, FM may split into a FM function in the data plane (FM-DP) and a FM function in the control plane (FM-CP).

FM-DP may be distributed with distributed mobility management. It may be a function in a data plane anchor or data plane node.

FM-CP may be distributed or centralized. It may be a function in a control plane node, control plane anchor or mobility controller.

Security Management (SM) function: The security management function controls security mechanisms/protocols providing access control, integrity, authentication, authorization, confidentiality, etc. for the control plane and data plane.

This function resides in all nodes such as control plane anchor, data plane anchor, mobile node, and correspondent node.

3. Distributed anchoring

3.1. Distributed anchoring configurations

The mobility functions may be implemented in different configurations of distributed anchoring in architectures separating the control and data planes. The separation as described in [I-D.wt-dmm-deployment-models] has defined home control plane anchor (Home-CPA), home data plane anchor (Home-DPA), access control plane node (Access-CPN), and access data plane node (Access-DPN), which are respectively abbreviated as CPA, DPA, CPN, and DPN here. Some configurations are described in [I-D.sijeon-dmm-deployment-models].

Figure 1 shows 4 configurations of network-based mobility management. In each configuration, an MN is allocated an IP prefix/address IP1 and is using IP1 to communicate with a correspondent node (CN) not shown in the figure. The flow of this communication session is shown as flow(IP1, ...) which uses IP1 and other parameters.

```
       (a)                (b)                (c)                (d)
                           +-----+                               +-----+
                           |LMs  |                               |LMs  |
                           +-----+                               +-----+

 +------------+     +------------+     +------------+     +------------+
 |CPA:        |     |CPA:        |     |CPA:        |     |CPA:        |
 |FM-CP, LM   |     |FM-CP, LMc  |     |FM-CP, LMs  |     |FM-CP, LMp  |
 +------------+     +------------+     +------------+     +------------+
 +------------+     +------------+     +------------+     +------------+
 |DPA(IPa1):  |     |DPA(IPa1):  |     |DPA(IPa1):  |     |DPA(IPa1):  |
 |anchors IP1 |     |anchors IP1 |     |anchors IP1 |     |anchors IP1 |
 |FM-DP       |     |FM-DP       |     |FM-DP       |     |FM-DP       |
 +------------+     +------------+     +------------+     +------------+


                                       +------------+     +------------+
                                       |CPN:        |     |CPN:        |
                                       |FM-CP, LMc  |     |FM-CP, LMc  |
                                       +------------+     +------------+
                                       +------------+     +------------+
                                       |DPN(IPn1):  |     |DPN(IPn1):  |
                                       |FM-DP       |     |FM-DP       |
                                       +------------+     +------------+


 +------------+     +------------+     +------------+     +------------+
 |MN(IP1)     |     |MN(IP1)     |     |MN(IP1)     |     |MN(IP1)     |
 |flow(IP1,..)|     |flow(IP1,..)|     |flow(IP1,..)|     |flow(IP1,..)|
 +------------+     +------------+     +------------+     +------------+
```

Figure 1. (a) FM-CP and LM at CPA, FM-DP at DPA; (b) Separate LMs, FM-CP and LMc at CPA, FM-DP at DPA; (c) FM-CP and LMs at CPA, FM-DP at DPA, FM-CP and LMc at CPN, FM-DP at DPN; (d) Separate LMs, FM-CP and LMp at CPA, FM-DP at DPA, FM-CP and LMc at CPN, FM-DP at DPN.

In Figures 1(a), 1(b), 1(c), and 1(d), the IP address of the MN, IP1, is anchored to the DPA which has the IP prefix/address IPa1. The data plane is distributed so that there may be multiple instances of the DPA (not shown). The control plane may either be distributed or centralized. When the CPA co-locates with the distributed DPA there will be multiple instances of the co-located CPA and DPA (not shown).

In Figure 1(a) and Figure 1(b), the network is flat with FM-DP at the distributed DPA.

In Figure 1(a), LM and FM-CP co-locate at CPA. Then LM may be distributed or centralized according to whether the CPA is distributed or centralized.

Figure 1(b) differs from Figure 1(a) in that the LM function is split into a server LMs and a client LMc. LMc and FM-CP are at the CPA. The LMs may be centralized whereas the LMc may be distributed or centralized according to whether the CPA is distributed or centralized.

In Figure 1(c) and Figure 1(d), the network is hierarchical where there may be multiple DPN’s for each DPA. There is FM-DP at each of the distributed DPA and at each of the distributed DPN.

In Figure 1(c), LMs and FM-CP are at the CPA. In addition, there are FM-CP and LMc at the CPN. Again, LMs may be distributed or centralized according to whether the CPA is distributed or centralized. The CPA may co-locate with DPA or may separate.

Figure 1(d) differs from Figure 1(c) in that the LMs is separated out, and a proxy LMp is added between the LMs and LMc. LMp and FM-CP are at the CPA. Again, there are FM-CP and LMc at the CPN. The LMs may be centralized whereas the LMp may be distributed or centralized according to whether the CPA is distributed or centralized.

Host-based variants of the mobility function configurations from Figures 1(c) and 1(d) are shown in Figures 2(a) and 2(b) where the role to perform mobility functions by CPN and DPN are now taken by the MN. The MN then need to possess the mobility functions FM and LMc.

```
         (a)                 (b)
                              +-----+
                              |LMs  |
                              +-----+

   +------------+      +------------+
   |CPA:        |      |CPA:        |
   |FM-CP, LMs  |      |FM-CP, LMp  |
   +------------+      +------------+
   +------------+      +------------+
   |DPA(IPa1):  |      |DPA(IPa1):  |
   |anchors IP1 |      |anchors IP1 |
   |FM-DP       |      |FM-DP       |
   +------------+      +------------+


   +------------+      +------------+
   |MN(IP1)     |      |MN(IP1)     |
   |flow(IP1,..)|      |flow(IP1,..)|
   |FM,    LMc  |      |FM,    LMc  |
   +------------+      +------------+
```

Figure 2. (a) FM-CP and LMs at CPA, FM-DP at DPA, FM and LMc at MN; (b) Separate LMs, FM-CP and LMp at CPA, FM-DP at DPA, FM and LMc at MN.

In Figure 2(a) and Figure 2(b), FM-DP is at the distributed DPA as before.

In Figure 2(a), LMs and FM-CP are at the CPA. The LMs may be distributed or centralized according to whether the CPA is distributed or centralized.

Figure 2(b) differs from Figure 2(a) in that the LMs is separated out and the proxy LMp is added between the LMs and LMc. LMp and FM-CP are at the CPA. The FMs may be centralized whereas the LMp may be distributed or centralized according to whether the CPA is distributed or centralized.

3.2. Distributed anchoring behaviors and message information elements

The behaviors of distributed anchoring are defined in this section in order that they may work together in expected manners to produce a distributed mobility solution. The needed information elements are passed as message parameters.

3.2.1. Location management behaviors and message information elements

It is seen in (Section 3.1) that

(1) LMs may be a separate server or may co-locate with LMc at CPA;
(2) LMc may be at CPA, CPN, or MN.

Example LM design may consists of a distributed database of LMs servers in a pool of distributed servers. The location information about the prefix/address of a MN is primarily at a given LMs. Peer LMs may exchange the location information with each other. LMc may retrieve a given record or send a given record update to LMs.

Location information bebaviors:

(LM:1) LM may manage the location information in a client-server database system. The example LM database functions are:

(LM:1-1) LMc may query LMs about location information for a prefix of MN (pull).
Parameters:
IP prefix of MN.

(LM:1-2) LMs may reply to LMc query about location information for a prefix of MN (pull).
Parameters:
IP prefix of MN,
IP address of FM-DP/DPA/DPN to forward the packets of the flow.

(LM:1-3) LMs may inform LMc about location information for a prefix of MN (push).
Parameters:
IP prefix of MN,
IP address of FM-DP/DPA/DPN to forward the packets of the flow.

(LM:1-4) LMc may inform LMs about update location information for a prefix of MN.
Parameters:
IP prefix of MN,
IP address of FM-DP/DPA/DPN to forward the packets of the flow.

(LM:2) The LM may be a distributed database with multiple LMs servers. For example:

(LM:2-1) A LMs may join a pool of LMs servers.

Parameters:
IP address of the LMs,
IP prefixes for which the LMs will host the primary location information.

(LM:2-2) LMs may query a peer LMs about location information for a prefix of MN.
Parameters:
IP prefix.

(LM:2-3) LMs may reply to a peer LMs about location information for a prefix of MN.
Parameters:
IP prefix of MN,
IP address of FM-DP/DPA/DPN to forward the packets of the flow.

3.2.2. Forwarding management behaviors and message information elements

It is seen in (Section 3.1) that

(1) FM-CP may be at CPA, CPN, MN;
(2) FM-DP may be at DPA, DPN, MN.

The FM behaviors and message information elements are:

(FM:1) With distributed FM functions, the role of FM for a flow may pass to another FM as the DPA or DPN changes.

(FM:2) In addition to above, a flow/session may be stateful for the required information for QoS, charging, etc. are needed. These states need to be transferred from the old anchor to the new anchor.

(FM:3) An anchor may act on packets on a per flow basis and perform the changes to the forwarding path upon a change of point of attachment of a MN:

(FM:3-1) FM filters the packets up to the granularity of a flow.
Example matching parameters are the 5-tuple of a flow.

(FM:3-2) FM makes the necessary changes to the forwarding path of a flow.
Example mechanism is through forwarding table update activated by DHCPv6-PD.

(FM:3-3) FM reverts the previously made changes to the forwarding path of a flow when such changes are no longer needed, e.g., when ongoing flows using an IP prefix/address requiring session continuity have closed.
Example mechanism is through expiration of DHCPv6-PD.

(FM:4) An anchor may discover and be discovered such as through an anchor registration system:

(FM:4-1) FM registers and authenticates itself with a centralized mobility controller.
Parameters:
IP address of DPA and its CPA;
IP prefix anchored to the DPA.

(FM:4-2) registration reply: acknowledge of registration and echo the input parameters.

(FM:4-3) FM discovers the FM of another IP prefix by querying the mobility controller based on the IP prefix.
Parameters:
IP prefix of MN.

(FM:4-4) when making anchor discovery FM expects the answer parameters as: IP address of DPA to which IP prefix of MN is anchored; IP prefix of the corresponding CPA.

(FM:5) With separation of control plane function and data plane function, these function must work together.

(FM:5-1) CPA/FM-CP sends forwarding table updates to DPA/FM-DP.
Parameters:
new forwarding table entries to add;
expired forwarding table entries to delete.

(FM:5-2) DPA/FM-DP sends to CPA/FM-CP about its status and load.
Parameters:
state of forwarding function being active or not;
loading percentage.

(FM:6) An anchor can buffer packets of a flow in a mobility event:

(FM:6-1) CPA/FM-CP informs DPA/FM-DP to buffer packets of a flow.
Trigger:
MN leaves DPA in a mobility event.
Parameters:
IP prefix of the flow for which packets need to be buffered.

(FM:6-2) CPA/FM-CP on behalf of a new DPA/FM-DP informs the CPA/FM-CP of the prior DPA/FM-DP that it is ready to receive any buffered packets of a flow.
Parameters:
destination IP prefix of the flow’s packets;
IP address of the new DPA.

4. Example mobility solutions with distributed anchoring

The IP prefix/address at the MN’s side of a flow may be anchored at the access router to which the MN is attached. For example, when an MN attaches to a network (Net1) or moves to a new network (Net2), it is allocated an IP prefix from that network. It configures from this prefix an IP address which is typically a dynamic IP address. It then uses this IP address when a flow is initiated. Packets to the MN in this flow are simply forwarded according to the forwarding table.

There may be multiple IP prefixes/addresses to choose from. They may be from the same access network or different access networks. The network may advertise these prefixes with cost options [I-D.mccann-dmm-prefixcost] so that the mobile node may choose the one with the least cost. In addition, these IP prefixes/addresses may be of different types regarding whether mobility support is needed [I-D.ietf-dmm-ondemand-mobility]. A flow will need to choose the appropriate one according to whether it needs IP mobility support.

4.1. IP mobility support only when needed

IP mobility support may be provided only when needed instead of being provided by default. The simplest configuration in this case is shown in Figures 1(a) and 1(b) in Section 3.1 for which the LM and FM functions are utilized only when needed.

A straightforward choice of mobility anchoring is for a flow to use the IP prefix of the network to which the MN is attached when the flow is initiated [I-D.seite-dmm-dma].

4.1.1. Not needed: Changing to the new IP prefix/address

When IP mobility support is not needed for a flow, the LM and FM functions are not utilized so that the configuration from Figures 1(a) and 1(b) in Section 3.1 simplifies to that shown in Figure 3.

```
Net1                                                       Net2

+---------------+                                          +---------------+
|AR1            |                                          |AR2            |
+---------------+                                          +---------------+
|CPA:           |                                          |CPA:           |
|---------------|                                          |---------------|
|DPA(IPa1):     |                                          |DPA(IPa2):     |
|anchors IP1    |                                          |anchors IP2    |
+---------------+                                          +---------------+


+...............+                                          +---------------+
.MN(IP1)        .                 move                     |MN(IP2)        |
.flow(IP1,...)  .               =======>                   |flow(IP2,...)  |
+...............+                                          +---------------+
```

Figure 3. Changing to the new IP prefix/address. MN running a flow using IP1 in Net1 changes to running a flow using IP2 in Net2.

When there is no need to provide IP mobility to a flow, the flow may use a new IP address acquired from a new network as the MN moves to the new network.

Regardless of whether IP mobility is needed, if the flow has terminated before the MN moves to a new network, the flow may subsequently restart using the new IP address allocated from the new network.

When session continuity is needed, even if a flow is ongoing as the MN moves, it may still be desirable for the flow to change to using the new IP prefix configured in the new network. The flow may then close and then restart using a new IP address configured in the new network. Such a change in the IP address of the flow may be enabled using a higher layer mobility support which is not in the scope of this document.

In Figure 3, a flow initiated while the MN was in Net1 has terminated before the MN moves to a new network Net2. After moving to Net2, the MN uses the new IP prefix anchored in Net2 to start a new flow. The packets may then be forwarded without requiring IP layer mobility support.

The call flow is outlined in Figure 4.

```
 MN                    p-AR          n-AR                             CN
  |MN attaches to p-AR: |             |                                |
  |acquire MN-ID and profile          |                                |
  |--RS---------------->|             |                                |
  |                     |             |                                |
  |<----------RA(HNP1)--|             |                                |
  |                     |             |                                |
Allocated prefix HNP1
IP1 address configuration
  |                     |             |                                |
  |<-Flow(IP1,IPcn,...)-+--------------------------------------------->|
  |                     |             |                                |
  |MN detaches from p-AR|             |                                |
  |MN attaches to n-AR  |             |                                |
  |                     |             |                                |
  |--RS------------------------------>|                                |
  |                     |             |                                |
  |<--------------RA(HNP2)------------|                                |
  |                     |             |                                |
Allocated prefix HNP2
IP2 address configuration
  |                     |             |                                |
  |<-new Flow(IP2,IPcn,...)-----------+------------------------------->|
  |                     |             |                                |
```

Figure 4. A flow uses the IP allocated from the network at which the MN is attached when the flow is initiated.

The security management function in the anchor node at a new network must allow to assign a valid IP prefix/address to a mobile node.

4.1.2. Needed: Providing IP mobility support

When IP mobility is needed for a flow, the LM and FM functions in Figures 1(a) and 1(b) in Section 3.1 are utilized. The mobility support may be provided by IP prefix anchor switching to the new network to be described in Section 4.2 or by using other mobility management methods ([Paper-Distributed.Mobility.PMIP] and [Paper-Distributed.Mobility.Review]). Then the flow may continue to use the IP prefix from the prior network. Yet some time later, the user application for the flow may be closed. If the application is started again, the new flow may not need to use the prior network's IP address to avoid having to invoke IP mobility support. This may be the case where a permanent IP prefix/address is not used. The flow may then use the new IP prefix in the network where the flow is being initiated. Routing is again kept simpler without employing IP mobility and will remain so as long as the MN has not moved away from that network.

The call flow in this case is outlined in Figure 5.

```
 MN                    p-AR          n-AR                             CN
  |MN attaches to p-AR: |             |                                |
  |acquire MN-ID and profile          |                                |
  |--RS---------------->|             |                                |
  |                     |             |                                |
  |<----------RA(HNP1)--|             |                                |
  |                     |             |                                |
Allocated prefix HNP1
IP1 address configuration
  |                     |             |                                |
  |<-Flow(IP1,IPcn,...)-+--------------------------------------------->|
  |                     |             |                                |
  |MN detach from p-AR  |             |                                |
  |MN attach to n-AR    |             |                                |
  |                     |             |                                |
  |--RS------------------------------>|                                |

IP mobility support such as that described in next sub-section
  |<--------------RA(HNP2,HNP1)-------|                                |
  |                     |             |                                |
  |<-Flow(IP1,IPcn,...)---------------+------------------------------->|
  |                     |             |                                |
Allocated prefix HNP2
IP2 address configuration
  |                     |             |                                |
Flow(IP1,IPcn) teminates
  |                     |             |                                |
  |<-new Flow(IP2,IPcn,...)-----------+------------------------------->|
  |                     |             |                                |
```

Figure 5. A flow uses the IP allocated from the network at which the MN is attached when the flow is initiated.

To provide IP mobility support with distributed anchoring, the distributed anchors may need to message with each other. When such messaging is needed, the anchors may need to discover each other as described in the FM behaviors and information elements (FM:2) in Section 3.2.2.

Then the anchors need to properly forward the packets of the flows as described in the FM behaviors and information elements (FM:1) in Section 3.2.2.

If there are in-flight packets toward the old anchor while the MN is moving to the new anchor, it may be necessary to buffer these packets and then forward to the new anchor after the old anchor knows that the new anchor is ready. Such are described in the FM behaviors and information elements (FM:4) in Section 3.2.2.

4.2. IP prefix/address anchor switching to the new network

The IP prefix/address anchoring may move without changing the IP prefix/address of the flow. Here the LM and FM functions in Figures 1(a) and 1(b) in Section 3.1 are implemented as shown in Figure 6.

```
Net1                                                  Net2

+---------------+                                     +---------------+
|AR1            |                                     |AR2            |
+---------------+                                     +---------------+
|CPA:           |                                     |CPA:           |
|LM:IP1<-->IPa2 |                                     |LM:IP1<-->IPa2 |
|---------------|                                     |---------------|
|DPA(IPa1):     |                                     |DPA(IPa2):     |
|anchors IP1    |               move                  |anchors IP2,IP1|
|FM:DHCPv6-PD   |              =======>               |FM:DHCPv6-PD   |
+---------------+                                     +---------------+


+...............+                                     +---------------+
.MN(IP1)        .               move                  |MN(IP2,IP1)    |
.flow(IP1,...)  .              =======>               |flow(IP1,...)  |
+...............+                                     +---------------+
```

Figure 6. IP prefix/address anchor switching to the new network. MN with flow using IP1 in Net1 continues to run the flow using IP1 as it moves to Net2.

As an MN with an ongoing session moves to a new network, the flow may preserve session continuity by moving the anchoring of the original IP prefix/address of the flow to the new network. An example is in the use of BGP UPDATE messages to change the forwarding table entries as described in [I-D.mccann-dmm-flatarch] and also for 3GPP Evolved Packet Core (EPC) network in [I-D.matsushima-stateless-uplane-vepc]. However, the response time and scalability of using a distributed routing protocol to update forwarding tables may be controversial.

Use of a centralized routing protocol with a centralized control plane as described in Section 4.2.1 will be more scalable.

The location management provides information about which IP prefix from an AR in the original network is being used by a flow in which AR in a new network. Such information needs to be deleted or updated when such flows have closed so that the IP prefix is no longer used in a different network. The LM behaviors are described in Section 3.2.1.

The FM functions are implemented through the DHCPv6-PD protocol. Here the anchor behavior to properly forward the packets for a flow as described in the FM behaviors and information elements FM:1 in Section 3.2.2 is realized by changing the anchor with DHCPv6-PD and also by reverting such changes later after the application has already closed and when the DHCPv6-PD timer expires. If there are in-flight packets toward the old anchor while the MN is moving to the new anchor, it may be necessary to buffer these packets and then forward to the new anchor after the old anchor knows that the new anchor is ready. Such are described in the FM behaviors and information elements FM:4 in Section 3.2.2. The anchors may also need to discover each other as described in the FM behaviors and information elements FM:2.

The security management function in the anchor node at a new network must allow to assign the original IP prefix/address used by the mobile node at the previous (original) network. As the assigned original IP prefix/address is to be used in the new network, the security management function in the anchor node must allow to advertise the prefix of the original IP address and also allow the mobile node to send and receive data packets with the original IP address.

The security management function in the mobile node must allow to configure the original IP prefix/address used at the previous (original) network when the original IP prefix/address is assigned by the anchor node in the new network. The security management function in the mobile node also allows to use the original IP address for the previous flow in the new network.

4.2.1. Centralized control plane

An example of IP prefix anchor switching is in the case where Net1 and Net2 both belong to the same operator network with separation of control and data planes ([I-D.liu-dmm-deployment-scenario] and [I-D.matsushima-stateless-uplane-vepc]), where the controller may send to the switches/routers the updated information of the forwarding tables with the IP address anchoring of the original IP prefix/address at AR1 moved to AR2 in the new network. That is, the IP address anchoring in the original network which was advertising the prefix will need to move to the new network. As the anchoring in the new network advertises the prefix of the original IP address in the new network, the forwarding tables will be updated so that packets of the flow will be forwarded according to the updated forwarding tables. The configuration in Figures 1(a) and 1(b) in Section 3.1 for which FM-CP and LM are centralized and FM-DP's are distributed. applies here. Figure 7 shows its implementation where LM is a binding between the original IP prefix/address of the flow and the IP address of the new DPA, whereas FM uses the DHCPv6-PD protocol.

```
Net1                                                         Net2
+-----------------------------------------------------------------+
|                            CPA:                                 |
|                            LM:IP1<-->IPa2                       |
|                            FM-CP                                |
+-----------------------------------------------------------------+


+---------------+                                      +---------------+
|AR1            |                                      |AR2            |
+---------------+                                      +---------------+
|DPA(IPa1):     |                                      |DPA(IPa2):     |
|anchors IP1    |                  move                |anchors IP2,IP1|
|FM:DHCPv6-PD   |                =======>              |FM:DHCPv6-PD   |
+---------------+                                      +---------------+


+...............+                                      +---------------+
.MN(IP1)        .                  move                |MN(IP2,IP1)    |
.flow(IP1,...)  .                =======>              |flow(IP1,...)  |
+...............+                                      +---------------+
```

Figure 7. IP prefix/address anchor switching to the new network with with LM and FM-CP in a centralized control plane whereas the FM-DP's are distributed.

The call flow in Figure 8 shows that MN is allocated HNP1 when it attaches to the p-AR. A flow running in MN may or may not need IP mobility. If it does, it may continue to use the previous IP prefix. If it does not, it may use a new IP prefix allocated from the new network.

```
 MN                    p-AR          n-AR             DHCP Servers   CN
  |MN attaches to p-AR: |             |                     |          |
  |acquire MN-ID and profile          |                     |          |
  |--RS---------------->|             |                     |          |
  |<----------RA(HNP1)--|             |                     |          |
  |                     |             |             Allocate MN-HNP1   |
IP addr config          |             |                     |          |
  |                     |             |                     |          |
  |<-Flow(IP1,IPcn,...)-+--------------------------------------------->|
  |                     |             |                     |          |
  |MN detach from p-AR  |             |                     |          |
  |MN attach to n-AR    |             |                     |          |
  |                     |             |                     |          |
  |--RS------------------------------>|                     |          |
  |                     |             |                     |          |
  |                     |------DHCPv6 release-------------->|          |
  |                     |             |                     |          |
  |                     |             |--DHCPv6 PD request->|          |
  |                     |             |<-DHCPv6 PD reply--->|          |
  |                     |             |                     |          |
  |                 forwarding table updates                |          |
  |                     |             |                     |          |
  |<--------------RA(HNP2,HNP1)-------|                     |          |
  |                     |             |             Allocate MN-HNP2   |
IP addr config          |             |                     |          |
  |                     |             |                     |          |
  |<-Flow(IP1,IPcn,...)---------------+------------------------------->|
  |                     |             |                     |          |
  |  Flow(IP1,IPcn,...) terminates    |                     |          |
  |                     |             |                     |          |
  |                     | DHCPv6-PD timeout                 |          |
  |                     |             |                     |          |
  |                 forwarding table updates                |          |
  |                     |             |                     |          |
  |                     |             |                     |          |
  |<-new Flow(IP2,IPcn,...)-----------+------------------------------->|
  |                     |             |                     |          |
```

Figure 8. DMM solution. MN with flow using IP1 in Net1 continues to run the flow using IP1 as it moves to Net2.

As the MN moves from p-AR to n-AR, the p-AR as a DHCP client may send a DHCP release message to release the HNP1. It is now necessary for n-AR to learn the IP prefix of the MN from the previous network so that it will be possible for Net2 to allocate both the previous network prefix and the new network prefix. The network may learn the previous prefix in different methods. For example, the MN may provide its previous network prefix information by including it to the RS message [I-D.jhlee-dmm-dnpp].

Knowing that MN is using HNP1, the n-AR sends to a DHCP server a DHCPv6-PD request to move the HNP1 to n-AR. The server sends to n-AR a DHCPv6-PD reply to move the HNP1. Then BGP route updates will take place here.

In addition, the MN also needs a new HNP in the new network. The n-AR may now send RA to n-AR, with prefix information that includes HNP1 and HNP2. The MN may then continue to use IP1. In addition, the MN is allocated the prefix HNP2 with which it may configure its IP addresses. Now for flows using IP1, packets destined to IP1 will be forwarded to the MN via n-AR.

As such flows have terminated and DHCP-PD has timed out, HNP1 goes back to Net1. MN will then be left with HNP2 only, which it will use when it now starts a new flow.

The anchor behavior to properly forward the packets for a flow as described in the FM behaviors and information elements (FM:1) in Section 3.2.2 is realized by changing the anchor with DHCPv6-PD and undoing such changes later when its timer expires and the application has already closed. With the anchors being separated in control and data planes with LMs and FM-CP centralized in the same control plane, messaging between anchors and the discovery of anchors become internal to the control plane. However, the centralized FM-CP needs to communicate with the distributed FM-DP as described as described in the FM behaviors and information elements (FM:3). Such may be realized by the appropriate messages in [I-D.ietf-dmm-fpc-cpdp]. Again, if there are in-flight packets toward the old anchor while the MN is moving to the new anchor, it may be necessary to buffer these packets and then forward to the new anchor after the old anchor knows that the new anchor is ready. The corresponding FM behaviors and information elements (FM:4) are however realized by the internal behavior in the control plane together with signaling between the control plane and distributed data plane.

4.2.2. Hierarchical network

The configuration for a hierarchical network is shown in Figures 1(c) and 1(d) in Section 3.1. With centralized control and with a centralized anchor, LM, CPA, CPN are co-located at the centralized control, and there is an AR with the DPA function supporting multiple forwarding switches (FW’s) each with a DPN function. A mobility event in this configuration involving change of FW but not of AR is shown in Figure 9.

Here the IP prefix allocated to the MN is anchored at the access router (AR) supporting the old FW to which the MN was originally attached as well as the new FW to which the MN has moved.

The realization of LM may bet the binding between the IP prefix/ address of the flow used by the MN and the IP address of the DPN to which MN has moved. The implementation of FM to enable change of FW without changing AR may be accomplished using tunneling between the AR and the FW as described in [I-D.korhonen-dmm-local-prefix] and in [I-D.templin-aerolink] or using some other L2 mobility mechanism.

```
Net1                                                        Net2
+-----------------------------------------------------------------------+
|                             CPA,CPN:                                  |
|                             LM:IP1<-->IPn2                            |
|                             FM-CP                                     |
+-----------------------------------------------------------------------+


                           +----------------+
                           |AR1             |
                           +----------------+
                           |DPA(IPa1):      |
                           |anchors IP1     |
                           |FM:DHCPv6-PD    |
                           +----------------+


+----------------+                                      +----------------+
|FW1             |                                      |FW2             |
+----------------+                move                  +----------------+
|DPN(IPn1):      |              =======>                |DPN(IPn2):      |
+----------------+                                      +----------------+


+................+                                      +----------------+
.MN(IP1)         .                move                  |MN(IP2)         |
.flow(IP1,...)   .              =======>                |flow(IP1,...)   |
+................+                                      +----------------+
```

Figure 9. Mobility without involving change of IP anchoring in a network with hierarchy in which the IP prefix allocated to the MN is anchored at an Edge Router supporting multiple access routers to which the MN may connect.

Here, the LM behaviors and information elements described in Section 3.2.1 provides information of which IP prefix from its FW needs to be used by a flow using which new FW. The anchor behaviors to properly forward the packets of a flow described in the FM behaviors and information elements (FM:1) may be realized with PMIPv6 protocol ([I-D.korhonen-dmm-local-prefix]) or with AERO protocol ([I-D.templin-aerolink]) to tunnel between the AR and the FW.

4.2.3. Hierarchical network with anchoring change

The configuration for a hierarchical network is still shown in Figures 1(c) and 1(d) in Section 3.1. Again, with centralized control and with a centralized anchor, LM, CPA, CPN are co-located at the centralized control, and there is an AR with the DPA function supporting multiple forwarding switches (FW’s) each with a DPN function. However, the mobility event involving change of FW may also involve a change of AR. Such configuration is shown in Figure 10.

This deployment case involves both a change of anchor from AR1 to AR2 and a network hierarchy AR-FW. It can be realized by a combination of changing the IP prefix/address anchoring from AR1 to AR2 with the mechanism as described in Section 4.2.1 and then forwarding the packets with network hierarchy AR-FW as described in Section 4.2.2.

To change AR, AR1 acting as a DHCP-PD client may exchange message with the DHCP server to release the prefix IP1. Meanwhile, AR2 acting as a DHCP-PD client may exchange message with the DHCP server to delegate the prefix IP1 to AR2.

```
Net1                                                    Net2
+-------------------------------------------------------------------------+
|                            CPA,CPN:                                     |
|                            LM:IP1<-->IPa2,IPn2                          |
|                            FM-CP                                        |
+-------------------------------------------------------------------------+

                            +---------------+
                            |Aggregate Point|
                            |---------------|
                            |FM,    LM      |
                            +---------------+


+---------------+                                       +---------------+
|AR1            |                                       |AR2            |
+---------------+                                       +---------------+
|DPA(IPa1):     |                                       |DPA(IPa2):     |
|anchors IP1    |                 move                  |anchors IP2,IP1|
|FM:DHCPv6-PD   |               =======>                |FM:DHCPv6-PD   |
+---------------+                                       +---------------+


+---------------+                                       +---------------+
|FW1            |                                       |FW2            |
+---------------+                 move                  +---------------+
|DPN(IPn1):     |               =======>                |DPN(IPn2):     |
+---------------+                                       +---------------+


+...............+                                       +---------------+
.MN(IP1)        .                 move                  |MN(IP2,IP1)    |
.flow(IP1,...)  .               =======>                |flow(IP1,...)  |
+...............+                                       +---------------+
```

Figure 10. Mobility involving change of IP anchoring in a network with hierarchy in which the IP prefix allocated to the MN is anchored at an Edge Router supporting multiple access routers to which the MN may connect.

5. Security Considerations

TBD

6. IANA Considerations

This document presents no IANA considerations.

7. Contributors

This document has benefited from other work on mobility solutions using BGP update, on mobility support in SDN network, on providing mobility support only when needed, and on mobility support in enterprise network. These work have been referenced. While some of these authors have taken the work to jointly write this document, others have contributed at least indirectly by writing these drafts. The latter include Philippe Bertin, Dapeng Liu, Satoru Matushima, Peter McCann, Pierrick Seite, Jouni Korhonen, and Sri Gundavelli.

Valuable comments have also been received from John Kaippallimil, ChunShan Xiong, and Dapeng Liu.

8. References

8.1. Normative References

[I-D.ietf-dmm-fpc-cpdp]
Liebsch, M., Matsushima, S., Gundavelli, S., Moses, D., and L. Bertz, "Protocol for Forwarding Policy Configuration (FPC) in DMM", draft-ietf-dmm-fpc-cpdp-03 (work in progress), March 2016.

[I-D.ietf-dmm-ondemand-mobility]
Yegin, A., Moses, D., Kweon, K., Lee, J., and J. Park, "On Demand Mobility Management", draft-ietf-dmm-ondemand-mobility-07 (work in progress), July 2016.

[I-D.jhlee-dmm-dnpp]
Lee, J. and Z. Yan, "Deprecated Network Prefix Provision", draft-jhlee-dmm-dnpp-01 (work in progress), April 2016.

[I-D.korhonen-dmm-local-prefix]
Korhonen, J., Savolainen, T., and S. Gundavelli, "Local Prefix Lifetime Management for Proxy Mobile IPv6", draft-korhonen-dmm-local-prefix-01 (work in progress), July 2013.

[I-D.liu-dmm-deployment-scenario]
Liu, V., Liu, D., Chan, A., Lingli, D., and X. Wei, "Distributed mobility management deployment scenario and architecture", draft-liu-dmm-deployment-scenario-05 (work in progress), October 2015.

[I-D.matsushima-stateless-uplane-vepc]
Matsushima, S. and R. Wakikawa, "Stateless user-plane architecture for virtualized EPC (vEPC)", draft-matsushima-stateless-uplane-vepc-06 (work in progress), March 2016.

[I-D.mccann-dmm-flatarch]
McCann, P., "Authentication and Mobility Management in a Flat Architecture", draft-mccann-dmm-flatarch-00 (work in progress), March 2012.

[I-D.mccann-dmm-prefixcost]
McCann, P. and J. Kaippallimalil, "Communicating Prefix Cost to Mobile Nodes", draft-mccann-dmm-prefixcost-03 (work in progress), April 2016.

[I-D.seite-dmm-dma]
Seite, P., Bertin, P., and J. Lee, "Distributed Mobility Anchoring", draft-seite-dmm-dma-07 (work in progress), February 2014.

[I-D.sijeon-dmm-deployment-models]
Jeon, S. and Y. Kim, "Deployment Models for Distributed Mobility Management", draft-sijeon-dmm-deployment-models-03 (work in progress), July 2016.

[I-D.templin-aerolink]
Templin, F., "Asymmetric Extended Route Optimization (AERO)", draft-templin-aerolink-67 (work in progress), June 2016.

[I-D.wt-dmm-deployment-models]
Gundavelli, S., "DMM Deployment Models and Architectural Considerations", draft-wt-dmm-deployment-models-00 (work in progress), April 2016.

[RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, DOI 10.17487/RFC2119, March 1997, <http://www.rfc-editor.org/info/rfc2119>.

[RFC5213] Gundavelli, S., Ed., Leung, K., Devarapalli, V., Chowdhury, K., and B. Patil, "Proxy Mobile IPv6", RFC 5213, DOI 10.17487/RFC5213, August 2008, <http://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, <http://www.rfc-editor.org/info/rfc6275>.

[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, <http://www.rfc-editor.org/info/rfc7333>.

[RFC7429] Liu, D., Ed., Zuniga, JC., Ed., Seite, P., Chan, H., and CJ. Bernardos, "Distributed Mobility Management: Current Practices and Gap Analysis", RFC 7429, DOI 10.17487/RFC7429, January 2015, <http://www.rfc-editor.org/info/rfc7429>.

8.2. Informative References

[Paper-Distributed.Mobility]
Lee, J., Bonnin, J., Seite, P., and H. Chan, "Distributed IP Mobility Management from the Perspective of the IETF: Motivations, Requirements, Approaches, Comparison, and Challenges", IEEE Wireless Communications, October 2013.

[Paper-Distributed.Mobility.PMIP]
Chan, H., "Proxy Mobile IP with Distributed Mobility Anchors", Proceedings of GlobeCom Workshop on Seamless Wireless Mobility, December 2010.

[Paper-Distributed.Mobility.Review]
Chan, H., Yokota, H., Xie, J., Seite, P., and D. Liu, "Distributed and Dynamic Mobility Management in Mobile Internet: Current Approaches and Issues", February 2011.

Authors' Addresses

H Anthony Chan
Huawei Technologies
5340 Legacy Dr. Building 3
Plano, TX 75024
USA

Email: h.a.chan@ieee.org

Xinpeng Wei
Huawei Technologies
Xin-Xi Rd. No. 3, Haidian District
Beijing, 100095
P. R. China

Email: weixinpeng@huawei.com

Jong-Hyouk Lee
Sangmyung University
708 Hannuri Building
Cheonan 330-720
Korea

Email: jonghyouk@smu.ac.kr

Seil Jeon
Instituto de Telecomunicacoes
Campus Universitario de Santiago
Aveiro 3810-193
Portugal

Email: seiljeon@av.it.pt

Fred L. Templin
Boeing Research and Technology
P.O. Box 3707
Seattle, WA 98124
USA

Email: fltemplin@acm.org

Distributed Mobility Management [dmm] C. Perkins
Internet-Draft Futurewei
Intended status: Standards Track V. Devarapalli
Expires: September 19, 2018 Vasona Networks
March 18, 2018

MN Identifier Types for RFC 4283 Mobile Node Identifier Option
draft-ietf-dmm-4283mnids-08.txt

Abstract

Additional Identifier Type Numbers are defined for use with the Mobile Node Identifier Option for MIPv6 (RFC 4283).

Status of This Memo

Copyright Notice

Table of Contents

1. Introduction

The Mobile Node Identifier Option for MIPv6 [RFC4283] has proved to be a popular design tool for providing identifiers for mobile nodes during authentication procedures with AAA protocols such as Diameter [RFC3588]. To date, only a single type of identifier has been specified, namely the MN NAI. Other types of identifiers are in common use, and even referenced in RFC 4283. In this document, we propose adding some basic types that are defined in various telecommunications standards, including types for IMSI [ThreeGPP-IDS], P-TMSI [ThreeGPP-IDS], IMEI [ThreeGPP-IDS], and GUTI [ThreeGPP-IDS]. In addition, we specify the IPv6 address itself and IEEE MAC-layer addresses as mobile node identifiers. Defining identifiers that are tied to the physical elements of the device (

MAC address etc.) help in deployment of Mobile IP because in many cases such identifiers are the most natural means for uniquely identifying the device, and will avoid additional look-up steps that might be needed if other identifiers were used.

2. Terminology

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

3. New Mobile Node Identifier Types

The following types of identifiers are commonly used to identify mobile nodes. For each type, references are provided with full details on the format of the type of identifer.

Mobile Node Identifier Description

Identifier Type	Description	Reference
IPv6 Address		[RFC4291]
IMSI	International Mobile Subscriber Identity	[ThreeGPP-IDS]
P-TMSI	Packet-Temporary Mobile Subscriber Identity	[ThreeGPP-IDS]
GUTI	Globally Unique Temporary ID	[ThreeGPP-IDS]
EUI-48 address	48-bit Extended Unique Identifier	[IEEE802]
EUI-64 address	64-bit Extended Unique Identifier-64 bit	[IEEE802]
DUID	DHCPv6 Unique Identifier	[RFC3315]

Table 1

4. Descriptions of MNID types

In this section descriptions for the various MNID types are provided.

4.1. Description of the IPv6 address type

The IPv6 address [RFC4291] is encoded as a 16 octet string containing a full IPv6 address which has been assigned to the mobile node. The IPv6 address MUST be a unicast routable IPv6 address. Multicast addresses, link-local addresses, and the unspecified IPv6 address MUST NOT be used. IPv6 Unique Local Addresses (ULAs) MAY be used, as long as any security operations making use of the ULA also take into account the domain in which the ULA is guaranteed to be unique.

4.2. Description of the IMSI MNID type

The International Mobile Subscriber Identity (IMSI) [ThreeGPP-IDS] is at most 15 decimal digits (i.e., digits from 0 through 9). The IMSI MUST be encoded as a string of octets in network order (i.e., high-to-low for all digits), where each digit occupies 4 bits. If needed for full octet size, the last digit MUST be padded with 0xf. For example an example IMSI 123456123456789 would be encoded as follows:

```
0x12, 0x34, 0x56, 0x12, 0x34, 0x56, 0x78, 0x9f
```

4.3. Description of the EUI-48 address type

The IEEE EUI-48 address [IEEE802-eui48] is encoded as 6 octets containing the IEEE EUI-48 address.

4.4. Description of the EUI-64 address type

The IEEE EUI-64 address [IEEE802-eui64] is encoded as 8 octets containing the full IEEE EUI-64 address.

4.5. Description of the DUID type

The DUID is the DHCPv6 Unique Identifier (DUID) [RFC3315]. There are various types of DUID, which are distinguished by an initial two-octet type field. Clients and servers MUST treat DUIDs as opaque values and MUST only compare DUIDs for equality.

5. Security Considerations

This document does not introduce any security mechanisms, and does not have any impact on existing security mechanisms.

Mobile Node Identifiers such as those described in this document are considered to be private information. If used in the MNID extension as defined in [RFC4283], the packet including the MNID extension MUST be encrypted so that no personal information or trackable identifiers is inadvertently disclosed to passive observers. Operators can potentially apply IPsec Encapsulating Security Payload (ESP) [RFC4303], in transport mode, with confidentiality and integrity protection for protecting the identity and location information in Mobile IPv6 signaling messages.

Some MNIDs contain sensitive identifiers which, as used in protocols specified by other SDOs, are only used for signaling during initial network entry. In such protocols, subsequent exchanges then rely on a temporary identifier allocated during the initial network entry. Managing the association between long-lived and temporary identifiers is outside the scope of this document.

6. IANA Considerations

The new mobile node identifier types defined in the document should be assigned values from the "Mobile Node Identifier Option Subtypes" registry. The following values should be assigned.

New Mobile Node Identifier Types

Identifier Type	Identifier Type Number
IPv6 Address	2
IMSI	3
P-TMSI	4
EUI-48 address	5
EUI-64 address	6
GUTI	7
DUID-LLT	8
DUID-EN	9
DUID-LL	10
DUID-UUID	11
	12-15 reserved
	16-255 unassigned

Table 2

See Section 4 for additional information about the identifier types. Future new assignments are to be made only after Expert Review [RFC8126]. The expert must ascertain that the identifier type allows unique identification of the mobile device; since all MNIDs require encryption there is no additional privacy exposure attendent to the use of new types.

7. Acknowledgements

The authors wish to acknowledge Hakima Chaouchi, Tatuya Jinmei, Jouni Korhonen, Sri Gundavelli, Suresh Krishnan, Dapeng Liu, Dale Worley, Joseph Salowey, Linda Dunbar, and Mirja Kuehlewind for their helpful comments.

8. References

8.1. 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>.

[RFC3315] Droms, R., Ed., Bound, J., Volz, B., Lemon, T., Perkins, C., and M. Carney, "Dynamic Host Configuration Protocol for IPv6 (DHCPv6)", RFC 3315, DOI 10.17487/RFC3315, July 2003, <https://www.rfc-editor.org/info/rfc3315>.

[RFC4283] Patel, A., Leung, K., Khalil, M., Akhtar, H., and K. Chowdhury, "Mobile Node Identifier Option for Mobile IPv6 (MIPv6)", RFC 4283, DOI 10.17487/RFC4283, November 2005, <https://www.rfc-editor.org/info/rfc4283>.

[RFC4291] Hinden, R. and S. Deering, "IP Version 6 Addressing Architecture", RFC 4291, DOI 10.17487/RFC4291, February 2006, <https://www.rfc-editor.org/info/rfc4291>.

[RFC4303] Kent, S., "IP Encapsulating Security Payload (ESP)", RFC 4303, DOI 10.17487/RFC4303, December 2005, <https://www.rfc-editor.org/info/rfc4303>.

[RFC8126] Cotton, M., Leiba, B., and T. Narten, "Guidelines for Writing an IANA Considerations Section in RFCs", BCP 26, RFC 8126, DOI 10.17487/RFC8126, June 2017, <https://www.rfc-editor.org/info/rfc8126>.

8.2. Informative References

[EANUCCGS] EAN International and the Uniform Code Council, "General EAN.UCC Specifications Version 5.0", Jan 2004.

[EPC-Tag-Data] EPCglobal Inc., "EPC(TM) Generation 1 Tag Data Standards Version 1.1 Rev.1.27 http://www.gs1.org/gsmp/kc/epcglobal/tds/tds_1_1_rev_1_27-standard-20050510.pdf", January 2005.

[IEEE802] IEEE, "IEEE Std 802: IEEE Standards for Local and Metropolitan Networks: Overview and Architecture", 2001.

[IEEE802-eui48]
IEEE, "Guidelines for 48-Bit Global Identifier (EUI-48) https://standards.ieee.org/develop/regauth/tut/eui48.pdf", 2001.

[IEEE802-eui64]
IEEE, "Guidelines for 64-Bit Global Identifier (EUI-64) https://standards.ieee.org/develop/regauth/tut/eui.pdf64", 2001.

[RFC3588] Calhoun, P., Loughney, J., Guttman, E., Zorn, G., and J. Arkko, "Diameter Base Protocol", RFC 3588, DOI 10.17487/RFC3588, September 2003, <https://www.rfc-editor.org/info/rfc3588>.

[RFID-DoD-spec]
Department of Defense, "United States Department of Defense Suppliers Passive RFID Information Guide (Version 15.0)", January 2010.

[RFID-framework]
Institut National des Telecommunication, ""Heterogeneous RFID framework design, analysis and evaluation"", July 2012.

[ThreeGPP-IDS]
3rd Generation Partnership Project, "3GPP Technical Specification 23.003 V8.4.0: Technical Specification Group Core Network and Terminals; Numbering, addressing and identification (Release 8)", March 2009.

[TRACK-IoT]
IPv6.com, ""Heterogeneous IoT Network : TRACK-IoT"", March 2012.

[Using-RFID-IPv6]
IPv6.com, ""Using RFID & IPv6"", September 2006.

Appendix A. RFID types

The material in this non-normative appendix was originally composed for inclusion in the main body of the specification, but was moved into an appendix because there was insufficient support for allocating RFID types at this time. It was observed that RFID-based mobile devices may create privacy exposures unless confidentiality is assured for signaling. A specification for eliminating unauthorized RFID tracking based on layer-2 addresses would be helpful.

Much of the following text is due to contributions from Hakima Chaouchi. For an overview and some initial suggestions about using RFID with IPv6 on mobile devices, see [Using-RFID-IPv6].

In the context of IoT and industry 4.0 vertical domain, efficient inventory and tracking items is of major interest, and RFID technology is the identification technology in the hardware design of many such items.

The "TRACKIOT: Heterogeneous IoT control" project ([TRACK-IoT], [RFID-framework]) explored Mobile IPv6 as a mobility management protocol for RFID-based mobile devices.

1. Passive RFID tags (that have no processing resources) need to be handled by the gateway (likely also the RFID Reader), which is then the end point of the mobility protocol. It is also the point where the CoA will be created based on some combination such as the RFID tag and the prefix of that gateway. The point here is to offer the possibility to passive RFID items to get an IPv6 address and take advantage of the mobility framework to follow the mobile device (passive tag on the item). One example scenario that has been proposed, showing the need for mobility management of passive RFID items, would be pieces of art tagged with passive tags that need to be monitored while transported.
2. Using active RFID tags (where processing resource is available on the tag), the end point of the mobility protocol can be pushed up to the RFID Active tag. We name it also an identification sensor. Use cases include active RFID tags for traceability of cold food respect during mobility (transport) of food. Mobility of cars equiped with active RFID tags that we already use for toll payement can be added with mobility management.

One major effort of connecting IETF efforts to the EPCGlobal (RFID standardisation) led to the ONS (DNS version applied for RFID logical names and page information retrieval). Attempts have tried to connect IPv6 on the address space to RFID identifier format. Other initiatives started working on gateways to map tag identifiers with IPv6 addresses and build signaling protocols for the application level. For instance tracking of mobile items equipped with a tag can be triggered remotely by a remote correspondent node until a visiting area where a mobile item equipped with an RFID tag is located. An RFID reader will be added with an IPv6 to RFID tag translation. One option is to build a Home IPv6 address of that tagged item by using the prefix of the Home agent combined with the tag RFID identifier of the mobile item; as the tag ID is unique, the home IPv6 address of that item will be also unique. Then the visiting RFID reader will compose the IPV6 care of address of the tagged mobile item by combining the prefix of the RFID reader with the tag ID of the item).

MIPv6 can then provide normally the mobility management of that RFID tagged item. A different useful example of tagged items involves items of a factory that can be tracked while they are transported, especially for real time localisation and tracking of precious items transported without GPS. An automotive car manufacturer can assign IPv6 addresses corresponding to RFID tagged cars or mechanical car parts, and build a tracking dataset of the mobility not only of the cars, but also of the mechanical pieces.

The Tag Data standard promoted by Electronic Product Code(TM) (abbreviated EPC) [EPC-Tag-Data] supports several encoding systems or schemes, which are commonly used in RFID (radio-frequency identification) applications, including

o RFID-GID (Global Identifier),
o RFID-SGTIN (Serialized Global Trade Item Number),
o RFID-SSCC (Serial Shipping Container),
o RFID-SGLN (Global Location Number),
o RFID-GRAI (Global Returnable Asset Identifier),
o RFID-DOD (Department of Defense ID), and
o RFID-GIAI (Global Individual Asset Identifier).

For each RFID scheme except GID, there are three representations:

o a 64-bit binary representation (for example, SGLN-64) (except for GID)
o a 96-bit binary representation (SGLN-96)
o a representation as a URI

The URI representation for the RFID is actually a URN. The EPC document has the following language:

> All categories of URIs are represented as Uniform Reference Names (URNs) as defined by [RFC2141], where the URN Namespace is epc.

The following list includes the above RFID types.

Mobile Node RFID Identifier Description

Identifier Type	Description	Reference
RFID-SGTIN-64	64-bit Serialized Global Trade Item Number	[EPC-Tag-Data]
RFID-SSCC-64	64-bit Serial Shipping Container	[EPC-Tag-Data]
RFID-SGLN-64	64-bit Serialized Global Location Number	[EPC-Tag-Data]
RFID-GRAI-64	64-bit Global Returnable Asset Identifier	[EPC-Tag-Data]
RFID-DOD-64	64-bit Department of Defense ID	[RFID-DoD-spec]
RFID-GIAI-64	64-bit Global Individual Asset Identifier	[EPC-Tag-Data]
RFID-GID-96	96-bit Global Identifier	[EPC-Tag-Data]
RFID-SGTIN-96	96-bit Serialized Global Trade Item Number	[EPC-Tag-Data]
RFID-SSCC-96	96-bit Serial Shipping Container	[EPC-Tag-Data]
RFID-SGLN-96	96-bit Serialized Global Location Number	[EPC-Tag-Data]
RFID-GRAI-96	96-bit Global Returnable Asset Identifier	[EPC-Tag-Data]
RFID-DOD-96	96-bit Department of Defense ID	[RFID-DoD-spec]
RFID-GIAI-96	96-bit Global Individual Asset Identifier	[EPC-Tag-Data]
RFID-GID-URI	Global Identifier represented as URI	[EPC-Tag-Data]
RFID-SGTIN-URI	Serialized Global Trade Item Number represented as URI	[EPC-Tag-Data]
RFID-SSCC-URI	Serial Shipping Container represented as URI	[EPC-Tag-Data]
RFID-SGLN-URI	Global Location Number represented as URI	[EPC-Tag-Data]
RFID-GRAI-URI	Global Returnable Asset Identifier represented as URI	[EPC-Tag-Data]
RFID-DOD-URI	Department of Defense ID represented as URI	[RFID-DoD-spec]
RFID-GIAI-URI	Global Individual Asset Identifier represented as URI	[EPC-Tag-Data]

Table 3

A.1. Description of the RFID types

The General Identifier (GID) that is used with RFID is composed of three fields - the General Manager Number, Object Class and Serial Number. The General Manager Number identifies an organizational entity that is responsible for maintaining the numbers in subsequent fields. GID encodings include a fourth field, the header, to guarantee uniqueness in the namespace defined by EPC.

Some of the RFID types depend on the Global Trade Item Number (GTIN) code defined in the General EAN.UCC Specifications [EANUCCGS]. A GTIN identifies a particular class of object, such as a particular kind of product or SKU.

The EPC encoding scheme for SGTIN permits the direct embedding of EAN.UCC System standard GTIN and Serial Number codes on EPC tags. In all cases, the check digit is not encoded. Two encoding schemes are specified, SGTIN-64 (64 bits) and SGTIN-96 (96 bits).

The Serial Shipping Container Code (SSCC) is defined by the EAN.UCC Specifications. Unlike the GTIN, the SSCC is already intended for assignment to individual objects and therefore does not require additional fields to serve as an EPC pure identity. Two encoding schemes are specified, SSCC-64 (64 bits) and SSCC-96 (96 bits).

The Global Location Number (GLN) is defined by the EAN.UCC Specifications. A GLN can represent either a discrete, unique physical location such as a warehouse slot, or an aggregate physical location such as an entire warehouse. In addition, a GLN can represent a logical entity that performs a business function such as placing an order. The Serialized Global Location Number (SGLN) includes the Company Prefix, Location Reference, and Serial Number.

The Global Returnable Asset Identifier (GRAI) is defined by the General EAN.UCC Specifications. Unlike the GTIN, the GRAI is already intended for assignment to individual objects and therefore does not require any additional fields to serve as an EPC pure identity. The GRAI includes the Company Prefix, Asset Type, and Serial Number.

The Global Individual Asset Identifier (GIAI) is defined by the General EAN.UCC Specifications. Unlike the GTIN, the GIAI is already intended for assignment to individual objects and therefore does not require any additional fields to serve as an EPC pure identity. The GRAI includes the Company Prefix, and Individual Asset Reference.

The DoD Construct identifier is defined by the United States Department of Defense (DoD). This tag data construct may be used to encode tags for shipping goods to the DoD by a supplier who has already been assigned a CAGE (Commercial and Government Entity) code.

A.1.1. Description of the RFID-SGTIN-64 type

The RFID-SGTIN-64 is encoded as specified in [EPC-Tag-Data]. The SGTIN-64 includes five fields: Header, Filter Value (additional data that is used for fast filtering and pre-selection), Company Prefix Index, Item Reference, and Serial Number. Only a limited number of Company Prefixes can be represented in the 64-bit tag.

A.1.2. Description of the RFID-SGTIN-96 type

The RFID-SGTIN-96 is encoded as specified in [EPC-Tag-Data]. The SGTIN-96 includes six fields: Header, Filter Value, Partition (an indication of where the subsequent Company Prefix and Item Reference numbers are divided), Company Prefix Index, Item Reference, and Serial Number.

A.1.3. Description of the RFID-SSCC-64 type

The RFID-SSCC-64 is encoded as specified in [EPC-Tag-Data]. The SSCC-64 includes four fields: Header, Filter Value, Company Prefix Index, and Serial Reference. Only a limited number of Company Prefixes can be represented in the 64-bit tag.

A.1.4. Description of the RFID-SSCC-96 type

The RFID-SSCC-96 is encoded as specified in [EPC-Tag-Data]. The SSCC-96 includes six fields: Header, Filter Value, Partition, Company Prefix, and Serial Reference, as well as 24 bits that remain Unallocated and must be zero.

A.1.5. Description of the RFID-SGLN-64 type

The RFID-SGLN-64 type is encoded as specified in [EPC-Tag-Data]. The SGLN-64 includes five fields: Header, Filter Value, Company Prefix Index, Location Reference, and Serial Number.

A.1.6. Description of the RFID-SGLN-96 type

The RFID-SGLN-96 type is encoded as specified in [EPC-Tag-Data]. The SGLN-96 includes six fields: Header, Filter Value, Partition, Company Prefix, Location Reference, and Serial Number.

A.1.7. Description of the RFID-GRAI-64 type

The RFID-GRAI-64 type is encoded as specified in [EPC-Tag-Data]. The GRAI-64 includes five fields: Header, Filter Value, Company Prefix Index, Asset Type, and Serial Number.

A.1.8. Description of the RFID-GRAI-96 type

The RFID-GRAI-96 type is encoded as specified in [EPC-Tag-Data]. The GRAI-96 includes six fields: Header, Filter Value, Partition, Company Prefix, Asset Type, and Serial Number.

A.1.9. Description of the RFID-GIAI-64 type

The RFID-GIAI-64 type is encoded as specified in [EPC-Tag-Data]. The GIAI-64 includes four fields: Header, Filter Value, Company Prefix Index, and Individual Asset Reference.

A.1.10. Description of the RFID-GIAI-96 type

The RFID-GIAI-96 type is encoded as specified in [EPC-Tag-Data]. The GIAI-96 includes five fields: Header, Filter Value, Partition, Company Prefix, and Individual Asset Reference.

A.1.11. Description of the RFID-DoD-64 type

The RFID-DoD-64 type is encoded as specified in [RFID-DoD-spec]. The DoD-64 type includes four fields: Header, Filter Value, Government Managed Identifier, and Serial Number.

A.1.12. Description of the RFID-DoD-96 type

The RFID-DoD-96 type is encoded as specified in [RFID-DoD-spec]. The DoD-96 type includes four fields: Header, Filter Value, Government Managed Identifier, and Serial Number.

A.1.13. Description of the RFID URI types

In some cases, it is desirable to encode in URI form a specific encoding of an RFID tag. For example, an application may prefer a URI representation for report preparation. Applications that wish to manipulate any additional data fields on tags may need some representation other than the pure identity forms.

For this purpose, the fields as represented the previous sections are associated with specified fields in the various URI types. For instance, the URI may have fields such as CompanyPrefix, ItemReference, or SerialNumber. For details and encoding specifics, consult [EPC-Tag-Data].

Authors' Addresses

Charles E. Perkins
Futurewei Inc.
2330 Central Expressway
Santa Clara, CA 95050
USA

Phone: +1-408-330-4586
Email: charliep@computer.org

Vijay Devarapalli
Vasona Networks
2900 Lakeside Drive, Suite 180
Santa Clara, CA 95054
USA

Email: dvijay@gmail.com

DMM Working Group
Internet-Draft
Intended status: Standards Track
Expires: 27 March 2021

S. Matsushima
SoftBank
L. Bertz
Sprint
M. Liebsch
NEC
S. Gundavelli
Cisco
D. Moses
Intel Corporation
C.E. Perkins
Futurewei

23 September 2020

Protocol for Forwarding Policy Configuration (FPC) in DMM
draft-ietf-dmm-fpc-cpdp-14

Abstract

This document describes a way, called Forwarding Policy Configuration (FPC) to manage the separation of data-plane and control-plane. FPC defines a flexible mobility management system using FPC agent and FPC client functions. A FPC agent provides an abstract interface to the data-plane. The FPC client configures data-plane nodes by using the functions and abstractions provided by the FPC agent for the data-plane nodes. The data-plane abstractions presented in this document are extensible in order to support many different types of mobility management systems and data-plane functions.

Status of This Memo

Table of Contents

1. Introduction

This document describes Forwarding Policy Configuration (FPC), a system for managing the separation of control-plane and data-plane. FPC enables flexible mobility management using FPC client and FPC agent functions. A FPC agent exports an abstract interface representing the data-plane. To configure data-plane nodes and functions, the FPC client uses the interface to the data-plane offered by the FPC agent.

Control planes of mobility management systems, or related applications which require data-plane control, can utilize the FPC client at various levels of abstraction. FPC operations are capable of directly configuring a single Data-Plane Node (DPN), as well as multiple DPNs, as determined by the data-plane models exported by the FPC agent.

A FPC agent represents the data-plane operation according to several basic information models. A FPC agent also provides access to Monitors, which produce reports when triggered by events or FPC Client requests regarding Mobility Contexts, DPNs or the Agent.

To manage mobility sessions, the FPC client assembles applicable sets of forwarding policies from the data model, and configures them on the appropriate FPC Agent. The Agent then renders those policies into specific configurations for each DPN at which mobile nodes are attached. The specific protocols and configurations to configure a DPN from a FPC Agent are outside the scope of this document.

A DPN is a logical entity that performs data-plane operations (packet movement and management). It may represent a physical DPN unit, a sub-function of a physical DPN or a collection of physical DPNs (i.e., a "virtual DPN"). A DPN may be virtual -- it may export the FPC DPN Agent interface, but be implemented as software that controls other data-plane hardware or modules that may or may not be FPC-compliant. In this document, DPNs are specified without regard for whether the implementation is virtual or physical. DPNs are connected to provide mobility management systems such as access networks, anchors and domains. The FPC agent interface enables establishment of a topology for the forwarding plane.

When a DPN is mapped to physical data-plane equipment, the FPC client can have complete knowledge of the DPN architecture, and use that information to perform DPN selection for specific sessions. On the other hand, when a virtual DPN is mapped to a collection of physical DPNs, the FPC client cannot select a specific physical DPN because it is hidden by the abstraction; only the FPC Agent can address the specific associated physical DPNs. Network architects have the flexibility to determine which DPN-selection capabilities are performed by the FPC Agent (distributed) and which by the FPC client (centralized). In this way, overlay networks can be configured without disclosing detailed knowledge of the underlying hardware to the FPC client and applications.

The abstractions in this document are designed to support many different mobility management systems and data-plane functions. The architecture and protocol design of FPC is not tied to specific types of access technologies and mobility protocols.

2. 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].

Attribute Expression: The definition of a template Property. This includes setting the type, current value, default value and if the attribute is static, i.e. can no longer be changed.

Domain: One or more DPNs that form a logical partition of network resources (e.g., a data-plane network under common network administration). A FPC client (e.g., a mobility management system) may utilize a single or multiple domains.

DPN: A data-plane node (DPN) is capable of performing data-plane features. For example, DPNs may be switches or routers, regardless of whether they are realized as hardware or purely in software.

FPC Client: A FPC Client is integrated with a mobility management system or related application, enabling control over forwarding policy, mobility sessions and DPNs via a FPC Agent.

Mobility Context: A Mobility Context contains the data-plane information necessary to efficiently send and receive traffic from a mobile node. This includes policies that are created or modified during the network’s operation - in most cases, on a per-flow or per session basis. A Mobility-Context represents the mobility sessions (or flows) which are active on a mobile node. This includes associated runtime attributes, such as tunnel endpoints, tunnel identifiers, delegated prefix(es), routing information, etc. Mobility-Contexts are associated to specific DPNs. Some pre-defined Policies may apply during mobility signaling requests. The Mobility Context supplies information about the policy settings specific to a mobile node and its flows; this information is often quite dynamic.

Mobility Session: Traffic to/from a mobile node that is expected to survive reconnection events.

Monitor: A reporting mechanism for a list of events that trigger notification messages from a FPC Agent to a FPC Client.

Policy: A Policy determines the mechanisms for managing specific traffic flows or packets. Policies specify QoS, rewriting rules for packet processing, etc. A Policy consists of one or more rules. Each rule is composed of a Descriptor and Actions. The Descriptor in a rule identifies packets (e.g., traffic flows), and the Actions apply treatments to packets that match the Descriptor in the rule. Policies can apply to Domains, DPNs, Mobile Nodes, Service-Groups, or particular Flows on a Mobile Node.

Property: An attribute-value pair for an instance of a FPC entity.

Service-Group: A set of DPN interfaces that support a specific data-plane purpose, e.g. inbound/ outbound, roaming, subnetwork with common specific configuration, etc.

Template: A recipe for instantiating FPC entities. Template definitions are accessible (by name or by a key) in an indexed set. A Template is used to create specific instances (e.g., specific policies) by assigning appropriate values into the Template definition via Attribute Expression.

Template Configuration The process by which a Template is referenced (by name or by key) and Attribute Expressions are created that change the value, default value or static nature of the Attribute, if permitted. If the Template is Extensible, new attributes MAY be added.

Tenant: An operational entity that manages mobility management systems or applications which require data-plane functions. A Tenant defines a global namespace for all entities owned by the Tenant enabling its entities to be used by multiple FPC Clients across multiple FPC Agents.

Topology: The DPNs and the links between them. For example, access nodes may be assigned to a Service-Group which peers to a Service-Group of anchor nodes.

3. FPC Design Objectives and Deployment

Using FPC, mobility control-planes and applications can configure DPNs to perform various mobility management roles as described in [I-D.ietf-dmm-deployment-models]. This fulfills the requirements described in [RFC7333].

This document defines FPC Agent and FPC Client, as well as the information models that they use. The attributes defining those models serve as the protocol elements for the interface between the FPC Agent and the FPC Client.

Mobility control-plane applications integrate features offered by the FPC Client. The FPC Client connects to FPC Agent functions. The Client and the Agent communicate based on information models described in Section 4. The models allow the control-plane to configure forwarding policies on the Agent for data-plane communications with mobile nodes.

Once the Topology of DPN(s) and domains are defined on an Agent for a data plane, the DPNs in the topology are available for further configuration. The FPC Agent connects those DPNs to manage their configurations.

A FPC Agent configures and manages its DPN(s) according to forwarding policies requested and Attributes provided by the FPC Client. Configuration commands used by the FPC agent to configure its DPN node(s) may be specific to the DPN implementation; consequently the method by which the FPC Agent carries out the specific configuration for its DPN(s) is out of scope for this document. Along with the data models, the FPC Client (on behalf of control-plane and applications) requests that the Agent configures Policies prior to the time when the DPNs start forwarding data for their mobility sessions.

This architecture is illustrated in Figure 1. A FPC Agent may be implemented in a network controller that handles multiple DPNs, or (more simply) an FPC Agent may itself be integrated into a DPN.

This document does not specify a protocol for the FPC interface; it is out of scope.

```
                           +-------------------------+
                           | Mobility Control-Plane  |
                           |          and            |
                           |      Applications       |
                           |+-----------------------+|
                           ||      FPC Client       ||
                           |+----------^------------+|
                           +-----------|-------------+
               FPC interface protocol  |
                       +---------------+-----------------+
                       |                                 |
         Network       |                                 |
         Controller    |                      DPN        |
           +-----------|-------------+        +----------|---------+
           |+----------v------------+|        |+---------v--------+|
           ||   [Data-plane model]  ||        ||[Data-plane model]||
           ||       FPC Agent       ||        ||    FPC Agent     ||
           |+-----------------------+|        |+------------------+|
           |+------------+----------+|        |                    |
           ||SB Protocol |FPC Client||        |  DPN Configuration |
           ||   Modules  |  Module  ||        +--------------------+
           |+------^-----+----^-----+|
           +-------|----------|------+
                   |          |
         Other     |          | FPC interface
        southbound |          | protocol
         protocols |          |
                   |          +-----------------+
                   |                            |
       DPN         |                 DPN        |
        +----------|---------+       +----------|---------+
        |+---------v--------+|       |+---------v--------+|
        ||  Configuration   ||       ||[Data-plane model]||
        || Protocol module  ||       ||    FPC Agent     ||
        |+------------------+|       |+------------------+|
        |                    |       |                    |
        | DPN Configuration  |       |  DPN Configuration |
        +--------------------+       +--------------------+
```

Figure 1: Reference Forwarding Policy Configuration (FPC) Architecture

The FPC architecture supports multi-tenancy; a FPC enabled data-plane supports tenants of multiple mobile operator networks and/or applications. It means that the FPC Client of each tenant connects to the FPC Agent and it MUST partition namespace and data for their data-planes. DPNs on the data-plane may fulfill multiple data-plane roles which are defined per session, domain and tenant.

Multi-tenancy permits the paritioning of data-plane entities as well as a common namespace requirement upon FPC Agents and Clients when they use the same Tenant for a common data-plane entity.

FPC information models often configuration to fit the specific needs for DPN management of a mobile node's traffic. The FPC interfaces in Figure 1 are the only interfaces required to handle runtime data in a Mobility Context. The Topology and some Policy FPC models MAY be pre-configured; in that case real-time protocol exchanges are not required for them.

The information model provides an extensibility mechanism through Templates that permits specialization for the needs of a particular vendor's equipment or future extension of the model presented in this specification.

4. FPC Mobility Information Model

The FPC information model includes the following components:

DPN Information Model,
Topology Information Model,
Policy Information Model,
Mobility-Context, and
Monitor, as illustrated in Figure 2.

```
                   :
                   |
                   +-[FPC Mobility Information Model]
                   |          |
                   |          +-[Topology Information Model]
                   |          |
                   |          +-[Policy Information Model]
                   |          |
                   |          +-[Mobility-Context]
                   |          |
                   |          +-[Monitor]
                   |
```

Figure 2: FPC Information Model structure

4.1. Model Notation and Conventions

The following conventions are used to describe the FPC information models.

Information model entities (e.g. DPNs, Rules, etc.) are defined in a hierarchical notation where all entities at the same hierarchical level are located on the same left-justified vertical position sequentially. When entities are composed of sub-entities, the sub-entities appear shifted to the right, as shown in Figure 3.

```
                              |
                              +-[entity2]
                              |          +-[entity2.1]
                              |          +-[entity2.2]
```

Figure 3: Model Notation - An Example

Some entities have one or more qualifiers placed on the right hand side of the element definition in angle-brackets. Common types include:

List: A collection of entities (some could be duplicated)

Set: A nonempty collection of entities without duplications

Name: A human-readable string

Key: A unique value. We distinguish 3 types of keys:

U-Key: A key unique across all Tenants. U-Key spaces typically involve the use of registries or language specific mechanisms that guarantee universal uniqueness of values.

G-Key: A key unique within a Tenant

L-Key: A key unique within a local namespace. For example, there may exist interfaces with the same name, e.g. "if0", in two different DPNs but there can only be one "if0" within each DPN (i.e. its local Interface-Key L-Key space).

Each entity or attribute may be optional (O) or mandatory (M). Entities that are not marked as optional are mandatory.

```
 The following example shows 3 entities:
     -- Entity1 is a globally unique key, and optionally can have
                an associated Name
     -- Entity2 is a list
     -- Entity3 is a set and is optional
                +
                |
                +-[entity1] <G-Key> (M), <Name> (O)
                +-[entity2] <List>
                +-[entity3] <Set> (O)
                |
                +
```

Figure 4

When expanding entity1 into a modeling language such as YANG it would result in two values: entity1-Key and entity1-Name.

To encourage re-use, FPC defines indexed sets of various entity Templates. Other model elements that need access to an indexed model entity contain an attribute which is always denoted as "entity-Key". When a Key attribute is encountered, the referencing model element may supply attribute values for use when the referenced entity model is instantiated. For example: Figure 5 shows 2 entities:

EntityA definition references an entityB model element.

EntityB model elements are indexed by entityB-Key.

Each EntityB model element has an entityB-Key which allows it to be uniquely identified, and a list of Attributes (or, alternatively, a Type) which specifies its form. This allows a referencing entity to create an instance by supplying entityB-Values to be inserted, in a Settings container.

```
                         .
                         .
                         |
                         +-[entityA]
                         |      +-[entityB-Key]
                         |      +-[entityB-Values]
                         .
                         .
                         |
                         +-[entityB] <L-Key> (M) <Set>
                         |      +-[entityB-Type]
                         .
                         .
```

Figure 5: Indexed sets of entities

Indexed sets are specified for each of the following kinds of entities:

Domain (See Section 4.9.3)
DPN (See Section 4.9.4)
Policy (See Section 4.9.5)
Rule (See Section 4.9.5)
Descriptor (See Figure 12)
Action (See Figure 12)
Service-Group (See Section 4.9.2, and
Mobility-Context (See Section 4.9.6)

As an example, for a Domain entity, there is a corresponding attribute denoted as "Domain-Key" whose value can be used to determine a reference to the Domain.

4.2. Templates and Attributes

In order to simplify development and maintenance of the needed policies and other objects used by FPC, the Information Models which are presented often have attributes that are not initialized with their final values. When an FPC entity is instantiated according to a template definition, specific values need to be configured for each such attribute. For instance, suppose an entity Template has an Attribute named "IPv4-Address", and also suppose that a FPC Client instantiates the entity and requests that it be installed on a DPN. An IPv4 address will be needed for the value of that Attribute before the entity can be used.

```
+-[Template] <U-Key, Name> (M) <Set>
|       +-[Attributes] <Set> (M)
|       +-[Extensible ~ FALSE]
|       +-[Entity-State ~ Initial]
|       +-[Version]
```

Figure 6: Template entities

Attributes: A set of Attribute names MAY be included when defining a Template for instantiating FPC entities.

Extensible: Determines whether or not entities instantiated from the Template can be extended with new non-mandatory Attributes not originally defined for the Template. Default value is FALSE. If a Template does not explicitly specify this attribute, the default value is considered to be in effect.

Entity-State: Either Initial, PartiallyConfigured, Configured, or Active. Default value is Initial. See Section 4.6 for more information about how the Entity-Status changes during the configuration steps of the Entity.

Version: Provides a version tag for the Template.

The Attributes in an Entity Template may be either mandatory or non-mandatory. Attribute values may also be associated with the attributes in the Entity Template. If supplied, the value may be either assigned with a default value that can be reconfigured later, or the value can be assigned with a static value that cannot be reconfigured later (see Section 4.3).

It is possible for a Template to provide values for all of its Attributes, so that no additional values are needed before the entity can made Active. Any instantiation from a Template MUST have at least one Attribute in order to be a useful entity unless the Template has none.

4.3. Attribute-Expressions

The syntax of the Attribute definition is formatted to make it clear. For every Attribute in the Entity Template, six possibilities are specified as follows:

'[Att-Name:]' Mandatory Attribute is defined, but template does not provide any configured value.

'[Att-Name: Att-Value]' Mandatory Attribute is defined, and has a statically configured value.

'[Att-Name: ~ Att-Value]' Mandatory Attribute is defined, and has a default value.

'[Att-Name]' Non-mandatory Attribute may be included but template does not provide any configured value.

'[Att-Name = Att-Value]' Non-mandatory Attribute may be included and has a statically configured value.

'[Att-Name ~ Att-Value]' Non-mandatory Attribute may be included and has a default value.

So, for example, a default value for a non-mandatory IPv4-Address attribute would be denoted by [IPv4-Address ~ 127.0.0.1].

After a FPC Client identifies which additional Attributes have been configured to be included in an instantiated entity, those configured Attributes MUST NOT be deleted by the FPC Agent. Similarly, any statically configured value for an entity Attribute MUST NOT be changed by the FPC Agent.

Whenever there is danger of confusion, the fully qualified Attribute name MUST be used when supplying needed Attribute Values for a structured Attribute.

4.4. Attribute Value Types

For situations in which the type of an attribute value is required, the following syntax is recommended. To declare than an attribute has data type "foo", typecast the attribute name by using the parenthesized data type (foo). So, for instance, [(float) Max-Latency-in-ms:] would indicate that the mandatory Attribute "Max-Latency-in-ms" requires to be configured with a floating point value before the instantiated entity could be used. Similarly, [(float) Max-Latency-in-ms: 9.5] would statically configure a floating point value of 9.5 to the mandatory Attribute "Max-Latency-in-ms".

4.5. Namespace and Format

The identifiers and names in FPC models which reside in the same Tenant must be unique. That uniqueness must be maintained by all Clients, Agents and DPNs that support the Tenant. The Tenant namespace uniqueness MUST be applied to all elements of the tenant model, i.e. Topology, Policy and Mobility models.

When a Policy needs to be applied to Mobility-Contexts in all Tenants on an Agent, the Agent SHOULD define that policy to be visible by all Tenants. In this case, the Agent assigns a unique identifier in the Agent namespace and copies the values to each Tenant. This effectively creates a U-Key although only a G-Key is required within the Tenant.

The notation for identifiers can utilize any format with agreement between data-plane agent and client operators. The formats include but are not limited to Globally Unique IDentifiers (GUIDs), Universally Unique IDentifiers (UUIDs), Fully Qualified Domain Names (FQDNs), Fully Qualified Path Names (FQPNs) and Uniform Resource Identifiers (URIs). The FPC model does not limit the format, which could dictate the choice of FPC protocol. Nevertheless, the identifiers which are used in a Mobility model should be considered to efficiently handle runtime parameters.

4.6. Configuring Attribute Values

Attributes of Information Model components such as policy templates are configured with values as part of FPC configuration operations. There may be several such configuration operations before the template instantiation is fully configured.

Entity-Status indicates when an Entity is usable within a DPN. This permits DPN design tradeoffs amongst local storage (or other resources), over the wire request size and the speed of request processing. For example, DPN designers with constrained systems MAY only house entities whose status is Active which may result in sending over all policy information with a Mobility-Context request. Storing information elements with an entity status of "PartiallyConfigured" on the DPN requires more resources but can result in smaller over the wire FPC communication and request processing efficiency.

When the FPC Client instantiates a Policy from a Template, the Policy-Status is "Initial". When the FPC Client sends the policy to a FPC Agent for installation on a DPN, the Client often will configure appropriate attribute values for the installation, and accordingly changes the Policy-Status to "PartiallyConfigured" or "Configured". The FPC Agent will also configure Domain-specific policies and DPN-specific policies on the DPN. When configured to provide particular services for mobile nodes, the FPC Agent will apply whatever service-specific policies are needed on the DPN. When a mobile node attaches to the network data-plane within the topology under the jurisdiction of a FPC Agent, the Agent may apply policies and settings as appropriate for that mobile node. Finally, when the mobile node launches new flows, or quenches existing flows, the FPC

Agent, on behalf of the FPC Client, applies or deactivates whatever policies and attribute values are appropriate for managing the flows of the mobile node. When a "Configured" policy is de-activated, Policy-Status is changed to be "Active". When an "Active" policy is activated, Policy-Status is changed to be "Configured".

Attribute values in DPN resident Policies may be configured by the FPC Agent as follows:

Domain-Policy-Configuration: Values for Policy attributes that are required for every DPN in the domain.

DPN-Policy-Configuration: Values for Policy attributes that are required for every policy configured on this DPN.

Service-Group-Policy-Configuration: Values for Policy attributes that are required to carry out the intended Service of the Service Group.

MN-Policy-Configuration: Values for Policy attributes that are required for all traffic to/from a particular mobile node.

Service-Data-Flow-Policy-Configuration: Values for Policy attributes that are required for traffic belonging to a particular set of flows on the mobile node.

Any configuration changes MAY also supply updated values for existing default attribute values that may have been previously configured on the DPN resident policy.

Entity blocks describe the format of the policy configurations.

4.7. Entity Configuration Blocks

As described in Section 4.6, a Policy Template may be configured in several stages by configuring default or missing values for Attributes that do not already have statically configured values. A Policy-Configuration is the combination of a Policy-Key (to identify the Policy Template defining the Attributes) and the currently configured Attribute Values to be applied to the Policy Template. Policy-Configurations MAY add attributes to a Template if Extensible is True. They MAY also refine existing attributes by:

assign new values if the Attribute is not static

make attributes static if they were not

make an attribute mandatory

A Policy-Configuration MUST NOT define or refine an attribute twice. More generally, an Entity-Configuration can be defined for any configurable Indexed Set to be the combination of the Entity-Key along with a set of Attribute-Expressions that supply configuration information for the entity's Attributes. Figure 7 shows a schematic representation for such Entity Configuration Blocks.

```
[Entity Configuration Block]
|       +-[Entity-Key] (M)
|       +-[Attribute-Expression] <Set> (M)
```

Figure 7: Entity Configuration Block

This document makes use of the following kinds of Entity Configuration Blocks:

Descriptor-Configuration

Action-Configuration

Rule-Configuration

Interface-Configuration

Service-Group-Configuration

Domain-Policy-Configuration

DPN-Policy-Configuration

Policy-Configuration

MN-Policy-Configuration

Service-Data-Flow-Policy-Configuration

4.8. Information Model Checkpoint

The Information Model Checkpoint permits Clients and Tenants with common scopes, referred to in this specification as Checkpoint BaseNames, to track the state of provisioned information on an Agent. The Agent records the Checkpoint BaseName and Checkpoint value set by a Client. When a Client attaches to the Agent it can query to determine the amount of work that must be executed to configure the Agent to a specific BaseName / checkpoint revision.

Checkpoints are defined for the following information model components:

Service-Group

DPN Information Model

Domain Information Model

Policy Information Model

4.9. Information Model Components

4.9.1. Topology Information Model

The Topology structure specifies DPNs and the communication paths between them. A network management system can use the Topology to select the most appropriate DPN resources for handling specific session flows.

The Topology structure is illustrated in Figure 8 (for definitions see Section 2):

```
|
+-[Topology Information Model]
|           +-[Extensible: FALSE]
|           +-[Service-Group]
|           +-[DPN] <Set>
|           +-[Domain] <Set>
```

Figure 8: Topology Structure

4.9.2. Service-Group

Service-Group-Set is collection of DPN interfaces serving some data-plane purpose including but not limited to DPN Interface selection to fulfill a Mobility-Context. Each Group contains a list of DPNs (referenced by DPN-Key) and selected interfaces (referenced by Interface-Key). The Interfaces are listed explicitly (rather than referred implicitly by its specific DPN) so that every Interface of a DPN is not required to be part of a Group. The information provided is sufficient to ensure that the Protocol, Settings (stored in the Service-Group-Configuration) and Features relevant to successful interface selection is present in the model.

```
|
+-[Service-Group] <G-Key>, <Name> (O) <Set>
|           +-[Extensible: FALSE]
|           +-[Role] <U-Key>
|           +-[Protocol] <Set>
|           +-[Feature] <Set> (O)
|           +-[Service-Group-Configuration] <Set> (O)
|           +-[DPN-Key] <Set>
|           |           +-[Referenced-Interface] <Set>
|           |           |       +-[Interface-Key] <L-Key>
|           |           |       +-[Peer-Service-Group-Key] <Set> (O)
```

Figure 9: Service Group

Each Service-Group element contains the following information:

Service-Group-Key: A unique ID of the Service-Group.

Service-Group-Name: A human-readable display string.

Role: The role (MAG, LMA, etc.) of the device hosting the interfaces of the DPN Group.

Protocol-Set: The set of protocols supported by this interface (e.g., PMIP, S5-GTP, S5-PMIP etc.). The protocol MAY be only its name, e.g. 'gtp', but many protocols implement specific message sets, e.g. s5-pmip, s8-pmip. When the Service-Group supports specific protocol message sub-subsets the Protocol value MUST include this information.

Feature-Set: An optional set of static features which further determine the suitability of the interface to the desired operation.

Service-Group-Configuration-Set: An optional set of configurations that further determine the suitability of an interface for the specific request. For example: SequenceNumber=ON/OFF.

DPN-Key-Set: A key used to identify the DPN.

Referenced-Interface-Set: The DPN Interfaces and peer Service-Groups associated with them. Each entry contains

Interface-Key: A key that is used together with the DPN-Key, to create a key that is refers to a specific DPN interface definition.

Peer-Service-Group-Key: Enables location of the peer Service-Group for this Interface.

4.9.3. Domain Information Model

A Domain-Set represents a group of heterogeneous Topology resources typically sharing a common administrative authority. Other models, outside of the scope of this specification, provide the details for the Domain.

```
   |
   +-[Domain] <G-Key>, <Name> (O) <Set>
   |       +-[Domain-Policy-Configuration] (O) <Set>
   |
```

Figure 10: Domain Information Model

Each Domain entry contains the following information:

Domain-Key: Identifies and enables reference to the Domain.

Domain-Name: A human-readable display string naming the Domain.

4.9.4. DPN Information Model

A DPN-Set contains some or all of the DPNs in the Tenant's network. Some of the DPNs in the Set may be identical in functionality and only differ by their Key.

```
   |
   +-[DPN] <G-Key>, <Name> (O) <Set>
   |     +-[Extensible: FALSE]
   |     +-[Interface] <L-Key> <Set>
   |     |      +-[Role] <U-Key>
   |     |      +-[Protocol] <Set>
   |     |      +-[Interface-Configuration] <Set> (O)
   |     +-[Domain-Key]
   |     +-[Service-Group-Key] <Set> (O)
   |     +-[DPN-Policy-Configuration] <List> (M)
   |     +-[DPN-Resource-Mapping-Reference] (O)
```

Figure 11: DPN Information Model

Each DPN entry contains the following information:

DPN-Key: A unique Identifier of the DPN.

DPN-Name: A human-readable display string.

Domain-Key: A Key providing access to the Domain information about the Domain in which the DPN resides.

Interface-Set: The Interface-Set references all interfaces (through which data packets are received and transmitted) available on the DPN. Each Interface makes use of attribute values that are specific to that interface, for example, the MTU size. These do not affect the DPN selection of active or enabled interfaces. Interfaces contain the following information:

Role: The role (MAG, LMA, PGW, AMF, etc.) of the DPN.

Protocol (Set): The set of protocols supported by this interface (e.g., PMIP, S5-GTP, S5-PMIP etc.). The protocol MAY implement specific message sets, e.g. s5-pmip, s8-pmip. When a protocol implements such message sub-subsets the Protocol value MUST include this information.

Interface-Configuration-Set: Configurable settings that further determine the suitability of an interface for the specific request. For example: SequenceNumber=ON/OFF.

Service-Group-Set: The Service-Group-Set references all of the Service-Groups which have been configured using Interfaces hosted on this DPN. The purpose of a Service-Group is not to describe each interface of each DPN, but rather to indicate interface types for use during the DPN selection process, when a DPN with specific interface capabilities is required.

DPN-Policy-Configuration: A list of Policies that have been configured on this DPN. Some may have values for all attributes, and some may require further configuration. Each Policy-Configuration has a key to enable reference to its Policy-Template. Each Policy-Configuration also has been configured to supply missing and non-default values to the desired Attributes defined within the Policy-Template.

DPN-Resource-Mapping-Reference (O): A reference to the underlying implementation, e.g. physical node, software module, etc. that supports this DPN. Further specification of this attribute is out of scope for this document.

4.9.5. Policy Information Model

The Policy Information Model defines and identifies Rules for enforcement at DPNs. A Policy is basically a set of Rules that are to be applied to each incoming or outgoing packet at a DPN interface. Rules comprise Descriptors and a set of Actions. The Descriptors, when evaluated, determine whether or not a set of Actions will be performed on the packet. The Policy structure is independent of a policy context.

In addition to the Policy structure, the Information Model (per Section 4.9.6) defines Mobility-Context. Each Mobility-Context may be configured with appropriate Attribute values, for example depending on the identity of a mobile node.

Traffic descriptions are defined in Descriptors, and treatments are defined separately in Actions. A Rule-Set binds Descriptors and associated Actions by reference, using Descriptor-Key and Action-Key. A Rule-Set is bound to a policy in the Policy-Set (using Policy-Key), and the Policy references the Rule definitions (using Rule-Key).

```
              |
              +-[Policy Information Model]
              |      +-[Extensible:]
              |      +-[Policy-Template] <G-Key> (M) <Set>
              |      |       +-[Policy-Configuration] <Set> (O)
              |      |       +-[Rule-Template-Key] <List> (M)
              |      |       |       +-[Precedence] (M)
              |      +-[Rule-Template] <L-Key> (M) <Set>
              |      |       +-[Descriptor-Match-Type] (M)
              |      |       +-[Descriptor-Configuration] <Set> (M)
              |      |       |       +-[Direction] (O)
              |      |       +-[Action-Configuration] <Set> (M)
              |      |       |       +-[Action-Order] (M)
              |      |       +-[Rule-Configuration] (O)
              |      +-[Descriptor-Template] <L-Key> (M) <Set>
              |      |       +-[Descriptor-Type] (O)
              |      |       +-[Attribute-Expression] <Set> (M)
              |      +-[Action-Template] <L-Key> (M) <Set>
              |              +-[Action-Type] (O)
              |      |       +-[Attribute-Expression] <Set> (M)
```

Figure 12: Policy Information Model

The Policy structure defines Policy-Set, Rule-Set, Descriptor-Set, and Action-Set, as follows:

Policy-Template: <Set> A set of Policy structures, indexed by Policy-Key, each of which is determined by a list of Rules referenced by their Rule-Key. Each Policy structure contains the following:

Policy-Key: Identifies and enables reference to this Policy definition.

Rule-Template-Key: Enables reference to a Rule template definition.

Rule-Precedence: For each Rule identified by a Rule-Template-Key in the Policy, specifies the order in which that Rule must be applied. The lower the numerical value of Precedence, the higher the rule precedence. Rules with equal precedence MAY be executed in parallel if supported by the DPN. If this value is absent, the rules SHOULD be applied in the order in which they appear in the Policy.

Rule-Template-Set: A set of Rule Template definitions indexed by Rule-Key. Each Rule is defined by a list of Descriptors (located by Descriptor-Key) and a list of Actions (located by Action-Key) as follows:

Rule-Template-Key: Identifies and enables reference to this Rule definition.

Descriptor-Match-Type Indicates whether the evaluation of the Rule proceeds by using conditional-AND, or conditional-OR, on the list of Descriptors.

Descriptor-Configuration: References a Descriptor template definition, along with an expression which names the Attributes for this instantiation from the Descriptor-Template and also specifies whether each Attribute of the Descriptor has a default value or a statically configured value, according to the syntax specified in Section 4.2.

Direction: Indicates if a rule applies to uplink traffic, to downlink traffic, or to both uplink and downlink traffic. Applying a rule to both uplink and downlink traffic, in case of symmetric rules, eliminates the requirement for a separate entry for each direction. When not present, the direction is implied by the Descriptor's values.

Action-Configuration: References an Action Template definition, along with an expression which names the Attributes for this instantiation from the Action-Template and also specifies whether each Attribute of the Action has a default value or a statically configured value, according to the syntax specified in Section 4.2.

Action-Order: Defines the order in which actions are executed when the associated traffic descriptor selects the packet.

Descriptor-Template-Set: A set of traffic Descriptor Templates, each of which can be evaluated on the incoming or outgoing packet, returning a TRUE or FALSE value, defined as follows:

Descriptor-Template-Key: Identifies and enables reference to this descriptor template definition.

Attribute-Expression: An expression which defines an Attribute in the Descriptor-Template and also specifies whether the Template also defines a default value or a statically configured value for the Attribute of the Descriptor has, according to the syntax specified in Section 4.2.

Descriptor-Type: Identifies the type of descriptor, e.g. an IPv6 traffic selector per [RFC6088].

Action-Template-Set: A set of Action Templates defined as follows:

Action-Template-Key: Identifies and enables reference to this action template definition.

Attribute-Expression: An expression which defines an Attribute in the Action-Template and also specifies whether the Template also defines a default value or a statically configured value for the Attribute of the Action has, according to the syntax specified in Section 4.2.

Action-Type: Identifies the type of an action for unambiguous interpretation of an Action-Value entry.

4.9.6. Mobility-Context Information Model

The Mobility-Context structure holds entries associated with a mobile node and its mobility sessions (flows). It is created on a DPN during the mobile node's registration to manage the mobile node's flows. Flow information is added or deleted from the Mobility-Context as needed to support new flows or to deallocate resources for flows that are deactivated. Descriptors are used to characterize the nature and resource requirement for each flow.

Termination of a Mobility-Context implies termination of all flows represented in the Mobility-Context, e.g. after deregistration of a mobile node. If any Child-Contexts are defined, they are also terminated.

```
     +-[Mobility-Context] <G-Key> <Set>
     |            +-[Extensible:~ FALSE]
     |            +-[Delegating-IP-Prefix:] <Set> (O)
     |            +-[Parent-Context] (O)
     |            +-[Child-Context] <Set> (O)
     |            +-[Service-Group-Key] <Set> (O)
     |            +-[Mobile-Node]
     |            |       +-[IP-Address] <Set> (O))
     |            |       +-[MN-Policy-Configuration] <Set>
     |            +-[Domain-Key]
     |            |       +-[Domain-Policy-Configuration] <Set>
     |            +-[DPN-Key] <Set>
     |            |       +-[Role]
     |            |       +-[DPN-Policy-Configuration] <Set>
     |            |       +-[ServiceDataFlow] <L-Key> <Set> (O)
     |            |       |       +-[Service-Group-Key] (O)
     |            |       |       +-[Interface-Key] <Set>
     |            |       |       +-[ServiceDataFlow-Policy-
     |                                  Configuration] <Set> (O)
     |            |       |       |       +-[Direction]
```

Figure 13: Mobility-Context Information Model

The Mobility-Context Substructure holds the following entries:

Mobility-Context-Key: Identifies a Mobility-Context

Delegating-IP-Prefix-Set: Delegated IP Prefixes assigned to the Mobility-Context

Parent-Context: If present, a Mobility Context from which the Attributes and Attribute Values of this Mobility Context are inherited.

Child-Context-Set: A set of Mobility Contexts which inherit the Attributes and Attribute Values of this Mobility Context.

Service-Group-Key: Service-Group(s) used during DPN assignment and re-assignment.

Mobile-Node: Attributes specific to the Mobile Node. It contains the following

IP-Address-Set IP addresses assigned to the Mobile Node.

MN-Policy-Configuration-Set For each MN-Policy in the set, a key and relevant information for the Policy Attributes.

Domain-Key: Enables access to a Domain instance.

Domain-Policy-Configuration-Set: For each Domain-Policy in the set, a key and relevant information for the Policy Attributes.

DPN-Key-Set: Enables access to a DPN instance assigned to a specific role, i.e. this is a Set that uses DPN-Key and Role as a compound key to access specific set instances.

Role: Role this DPN fulfills in the Mobility-Context.

DPN-Policy-Configuration-Set: For each DPN-Policy in the set, a key and relevant information for the Policy Attributes.

ServiceDataFlow-Key-Set: Characterizes a traffic flow that has been configured (and provided resources) on the DPN to support data-plane traffic to and from the mobile device.

Service-Group-Key: Enables access to a Service-Group instance.

Interface-Key-Set: Assigns the selected interface of the DPN.

ServiceDataFlow-Policy-Configuration-Set: For each Policy in the set, a key and relevant information for the Policy Attributes.

Direction: Indicates if the reference Policy applies to uplink or downlink traffic, or to both, uplink- and downlink traffic. Applying a rule to both, uplink- and downlink traffic, in case of symmetric rules, allows omitting a separate entry for each direction. When not present the value is assumed to apply to both directions.

4.9.7. Monitor Information Model

Monitors provide a mechanism to produce reports when events occur. A Monitor will have a target that specifies what is to be watched.

The attribute/entity to be monitored places certain constraints on the configuration that can be specified. For example, a Monitor using a Threshold configuration cannot be applied to a Mobility-Context, because it does not have a threshold. Such a monitor configuration could be applied to a numeric threshold property of a Context.

```
                              |
                              +-[Monitor] <G-Key> <List>
                              |          +-[Extensible:]
                              |          +-[Target:]
                              |          +-[Deferrable]
                              |          +-[Configuration]
```

Figure 14: Monitor Substructure

Monitor-Key: Identifies the Monitor.

Target: Description of what is to be monitored. This can be a Service Data Flow, a Policy installed upon a DPN, values of a Mobility-Context, etc. The target name is the absolute information model path (separated by '/') to the attribute / entity to be monitored.

Deferrable: Indicates that a monitoring report can be delayed up to a defined maximum delay, set in the Agent, for possible bundling with other reports.

Configuration: Determined by the Monitor subtype. The monitor report is specified by the Configuration. Four report types are defined:

* "Periodic" reporting specifies an interval by which a notification is sent.

* "Event-List" reporting specifies a list of event types that, if they occur and are related to the monitored attribute, will result in sending a notification.

* "Scheduled" reporting specifies the time (in seconds since Jan 1, 1970) when a notification for the monitor should be sent. Once this Monitor's notification is completed the Monitor is automatically de-registered.

* "Threshold" reporting specifies one or both of a low and high threshold. When these values are crossed a corresponding notification is sent.

5. Security Considerations

Detailed protocol implementations for DMM Forwarding Policy Configuration must ensure integrity of the information exchanged between a FPC Client and a FPC Agent. Required Security Associations may be derived from co-located functions, which utilize the FPC Client and FPC Agent respectively.

General usage of FPC MUST consider the following:

FPC Naming Section 4.5 permits arbitrary string values but a user MUST avoid placing sensitive or vulnerable information in those values.

Policies that are very narrow and permit the identification of specific traffic, e.g. that of a single user, SHOULD be avoided.

6. IANA Considerations

TBD

7. Work Team Participants

Participants in the FPSM work team discussion include Satoru Matsushima, Danny Moses, Sri Gundavelli, Marco Liebsch, Pierrick Seite, Alper Yegin, Carlos Bernardos, Charles Perkins and Fred Templin.

8. References

8.1. 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>.

[RFC6088] Tsirtsis, G., Giarreta, G., Soliman, H., and N. Montavont, "Traffic Selectors for Flow Bindings", RFC 6088, DOI 10.17487/RFC6088, January 2011, <https://www.rfc-editor.org/info/rfc6088>.

8.2. Informative References

[I-D.bertz-dime-policygroups]
Bertz, L. and M. Bales, "Diameter Policy Groups and Sets", Work in Progress, Internet-Draft, draft-bertz-dime-policygroups-06, 18 June 2018, <http://www.ietf.org/internet-drafts/draft-bertz-dime-policygroups-06.txt>.

[I-D.ietf-dmm-deployment-models]
Gundavelli, S. and S. Jeon, "DMM Deployment Models and Architectural Considerations", Work in Progress, Internet-Draft, draft-ietf-dmm-deployment-models-04, 15 May 2018, <http://www.ietf.org/internet-drafts/draft-ietf-dmm-deployment-models-04.txt>.

[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>.

Appendix A. Implementation Status

Three FPC Agent implementations have been made to date. The first was based upon Version 03 of the draft and followed Model 1. The second follows Version 04 of the document. Both implementations were OpenDaylight plug-ins developed in Java by Sprint. Version 04 is now primarily enhanced by GS Labs. Version 03 was known as fpcagent and version 04's implementation is simply referred to as 'fpc'. A third has been developed on an ONOS Controller for use in MCORD projects.

fpcagent's intent was to provide a proof of concept for FPC Version 03 Model 1 in January 2016 and research various errors, corrections and optimizations that the Agent could make when supporting multiple DPNs.

As the code developed to support OpenFlow and a proprietary DPN from a 3rd party, several of the advantages of a multi-DPN Agent became obvious including the use of machine learning to reduce the number of Flows and Policy entities placed on the DPN. This work has driven new efforts in the DIME WG, namely Diameter Policy Groups [I-D.bertz-dime-policygroups].

A throughput performance of tens per second using various NetConf based solutions in OpenDaylight made fpcagent, based on version 03, undesirable for call processing. The RPC implementation improved throughput by an order of magnitude but was not useful based upon FPC's Version 03 design using two information models. During this time the features of version 04 and its converged model became attractive and the fpcagent project was closed in August 2016. fpcagent will no longer be developed and will remain a proprietary implementation.

The learnings of fpcagent has influenced the second project, fpc. Fpc is also an OpenDaylight project but is an open source release as the Opendaylight FpcAgent plugin (https://wiki.opendaylight.org/view/Project_Proposals:FpcAgent). This project is scoped to be a fully compliant FPC Agent that supports multiple DPNs including those that communicate via OpenFlow. The following features present in this draft and others developed by the FPC development team have already led to an order of magnitude improvement.

Migration of non-realtime provisioning of entities such as topology and policy allowed the implementation to focus only on the rpc.

Using only 5 messages and 2 notifications has also reduced implementation time.

Command Sets, an optional feature in this specification, have eliminated 80% of the time spent determining what needs to be done with a Context during a Create or Update operation.

Op Reference is an optional feature modeled after video delivery. It has reduced unnecessary cache lookups. It also has the additional benefit of allowing an Agent to become cacheless and effectively act as a FPC protocol adapter remotely with multi-DPN support or co-located on the DPN in a single-DPN support model.

Multi-tenant support allows for Cache searches to be partitioned for clustering and performance improvements. This has not been capitalized upon by the current implementation but is part of the development roadmap.

Use of Contexts to pre-provision policy has also eliminated any processing of Ports for DPNs which permitted the code for CONFIGURE and CONF_BUNDLE to be implemented as a simple nested FOR loops (see below).

Initial v04 performance results without code optimizations or tuning allow reliable provisioning of 1K FPC Mobility-Contexts processed per second on a 12 core server. This results in 2x the number of transactions on the southbound interface to a proprietary DPN API on the same machine.

fpc currently supports the following:

1 proprietary DPN API

Policy and Topology as defined in this specification using OpenDaylight North Bound Interfaces such as NetConf and RestConf

CONFIG and CONF_BUNDLE (all operations)

DPN assignment, Tunnel allocations and IPv4 address assignment by the Agent or Client.

Immediate Response is always an OK_NOTIFY_FOLLOWS.

```
assignment system (receives rpc call):
  perform basic operation integrity check
  if CONFIG then
    goto assignments
    if assignments was ok then
      send request to activation system
      respond back to client with assignment data
    else
      send back error
    end if
  else if CONF_BUNDLE then
    for each operation in bundles
    goto assignments
    if assignments was ok then
      hold onto data
    else
      return error with the assignments that occurred in
        prior operations (best effort)
    end if
    end for
    send bundles to activation systems
  end if

assignments:
  assign DPN, IPv4 Address and/or tunnel info as required
  if an error occurs undo all assignments in this operation
  return result

activation system:
  build cache according to op-ref and operation type
  for each operation
    for each Context
      for each DPN / direction in Context
        perform actions on DPN according to Command Set
      end for
    end for
  end for
  commit changes to in memory cache
  log transaction for tracking and notification
                                (CONFIG_RESULT_NOTIFY)
```

Figure 15: fpc pseudo code

For further information please contact Lyle Bertz who is also a co-author of this document.

NOTE: Tenant support requires binding a Client ID to a Tenant ID (it is a one to many relation) but that is outside of the scope of this specification. Otherwise, the specification is complete in terms of providing sufficient information to implement an Agent.

Authors' Addresses

Satoru Matsushima
SoftBank
1-9-1,Higashi-Shimbashi,Minato-Ku,
Japan

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

Lyle Bertz
6220 Sprint Parkway
Overland Park KS, 66251,
United States of America

Email: lylebe551144@gmail.com

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

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

Sri Gundavelli
Cisco
170 West Tasman Drive
San Jose, CA 95134
United States of America

Email: sgundave@cisco.com

Danny Moses

Email: danny.moses@intel.com

Charles E. Perkins
Futurewei Inc.
2330 Central Expressway
Santa Clara, CA 95050
United States of America

Phone: +1-408-330-4586
Email: charliep@computer.org

DMM Working Group A. Yegin
Internet-Draft Actility
Intended status: Informational D. Moses
Expires: January 31, 2020 Intel
S. Jeon
Sungkyunkwan University
July 30, 2019

On Demand Mobility Management
draft-ietf-dmm-ondemand-mobility-18

Abstract

Applications differ with respect to whether they need session continuity and/or IP address reachability. The network providing the same type of service to any mobile host and any application running on the host yields inefficiencies, as described in [RFC7333]. This document defines a new concep of enabling applications to influence the network's mobility services (session continuity and/or IP address reachability) on a per-Socket basis, and suggests extensions to the networking stack's API to accomodate this concept.

Status of This Memo

Copyright Notice

Table of Contents

1. Introduction

In the context of Mobile IP [RFC5563][RFC6275][RFC5213][RFC5944], the following two attributes are defined for IP service provided to mobile hosts:

- Session Continuity

The ability to maintain an ongoing transport interaction by keeping the same local end-point IP address throughout the life-time of the IP socket despite the mobile host changing its point of attachment within the IP network topology. The IP address of the host may change after closing the IP socket and before opening a new one, but that does not jeopardize the ability of applications using these IP sockets to work flawlessly. Session continuity is essential for mobile hosts to maintain ongoing flows without any interruption.

- IP Address Reachability

The ability to maintain the same IP address for an extended period of time. The IP address stays the same across independent sessions, and even in the absence of any session. The IP address may be published in a long-term registry (e.g., DNS), and is made available for serving incoming (e.g., TCP) connections. IP address reachability is essential for mobile hosts to use specific/published IP addresses.

Mobile IP is designed to provide both session continuity and IP address reachability to mobile hosts. Architectures utilizing these protocols (e.g., 3GPP, 3GPP2, WIMAX) ensure that any mobile host attached to the compliant networks can enjoy these benefits. Any application running on these mobile hosts is subjected to the same treatment with respect to session continuity and IP address reachability.

Achieving session continuity and IP address reachability with Mobile IP incurs some cost. Mobile IP protocol forces the mobile host's IP traffic to traverse a centrally-located router (Home Agent, HA), which incurs additional transmission latency and use of additional network resources, adds to the network CAPEX and OPEX, and decreases the reliability of the network due to the introduction of a single point of failure [RFC7333]. Therefore, session continuity and IP address reachability SHOULD be provided only when necessary.

In reality not every application may need these benefits. IP address reachability is required for applications running as servers (e.g., a web server running on the mobile host). But, a typical client application (e.g., web browser) does not necessarily require IP address reachability. Similarly, session continuity is not required for all types of applications either. Applications performing brief communication (e.g., text messaging) can survive without having session continuity support.

Furthermore, when an application needs session continuity, it may be able to satisfy that need by using a solution above the IP layer, such as MPTCP [RFC6824], SIP mobility [RFC3261], or an application-layer mobility solution. These higher-layer solutions are not subject to the same issues that arise with the use of Mobile IP since they can utilize the most direct data path between the end-points. But, if Mobile IP is being applied to the mobile host, the higher-layer protocols are rendered useless because their operation is inhibited by Mobile IP. Since Mobile IP ensures that the IP address of the mobile host remains fixed (despite the location and movement of the mobile host), the higher-layer protocols never detect the IP-layer change and never engage in mobility management.

This document proposes a solution for applications running on mobile hosts to indicate when establishing the network connection ('on demand') whether they need session continuity or IP address reachability. The network protocol stack on the mobile host, in conjunction with the network infrastructure, provides the required type of service. It is for the benefit of both the users and the network operators not to engage an extra level of service unless it is absolutely necessary. It is expected that applications and networks compliant with this specification will utilize this solution to use network resources more efficiently.

2. Notational Conventions

The key words "MUST", "MUST NOT", "REQUIRED", "SHALL", "SHALL NOT", "SHOULD", "SHOULD NOT", "RECOMMENDED", "NOT RECOMMENDED", "MAY", and "OPTIONAL" in this document are to be interpreted as described in BCP 14 , [RFC2119] [RFC8174] when, they appear in all capitals, as shown here.

3. Solution

3.1. High-level Description

Enabling applications to indicate their mobility service requirements e.g. session continuity and/or IP address reachability, comprises the following steps:

- The application indicates to the network stack (local to the mobile host) the desired mobility service.

- The network stack assigns a source IP address based on an IP prefix with the desired services that was previously provided by the network. If such an IP prefix is not available, the network stack performs the additional steps below.

- The network stack sends a request to the network for a new source IP prefix that is associated with the desired mobility service.

- The network responds with the suitable allocated source IP prefix (or responds with a failure indication).

- If the suitable source IP prefix was allocates, the network stack constructs a source IP address and provides it to the application.

This document specifies the new address types associated with mobility services and details the interaction between the applications and the network stack steps. It uses the Socket interface as an example for an API between applications and the network stack. Other steps are outside the scope of this document.

3.2. Types of IP Addresses

Four types of IP addresses are defined with respect to mobility management.

- Fixed IP Address

A Fixed IP address is an address with a guarantee to be valid for a very long time, regardless of whether it is being used in any packet to/from the mobile host, or whether or not the mobile host is connected to the network, or whether it moves from one point-of-attachment to another (with a different IP prefix) while it is connected.

Fixed IP addresses are required by applications that need both session continuity and IP address reachability.

- Session-lasting IP Address

A session-lasting IP address is an address with a guarantee to be valid throughout the life-time of the socket(s) for which it was requested. It is guaranteed to be valid even after the mobile host had moved from one point-of-attachment to another (with a different IP prefix).

Session-lasting IP addresses are required by applications that need session continuity but do not need IP address reachability.

- Non-persistent IP Address

This type of IP address has no guarantee to exist after a mobile host moves from one point-of-attachment to another, and therefore, no session continuity nor IP address reachability are provided. The IP address is created from an IP prefix that is obtained from the serving IP gateway and is not maintained across gateway changes. In other words, the IP prefix may be released and replaced by a new one when the IP gateway changes due to the movement of the mobile host forcing the creation of a new source IP address with the updated allocated IP prefix.

- Graceful Replacement IP Address

In some cases, the network cannot guarantee the validity of the provided IP prefix throughout the duration of the opened socket, but can provide a limited graceful period of time in which both the original IP prefix and a new one are valid. This enables the application some flexibility in the transition from the existing source IP address to the new one.

This gracefulness is still better than the non-persistence type of address for applications that can handle a change in their source IP address but require that extra flexibility.

Applications running as servers at a published IP address require a Fixed IP Address. Long-standing applications (e.g., an SSH session) may also require this type of address. Enterprise applications that connect to an enterprise network via virtual LAN require a Fixed IP Address.

Applications with short-lived transient sessions can use Session-lasting IP Addresses. For example: Web browsers.

Applications with very short sessions, such as DNS clients and instant messengers, can utilize Non-persistent IP Addresses. Even though they could very well use Fixed or Session-lasting IP Addresses, the transmission latency would be minimized when a Non-persistent IP Addresses are used.

Applications that can tolerate a short interruption in connectivity can use the Graceful-replacement IP addresses. For example, a streaming client that has buffering capabilities.

3.3. Granularity of Selection

IP address type selection is made on a per-socket granularity. Different parts of the same application may have different needs. For example, the control-plane of an application may require a Fixed IP Address in order to stay reachable, whereas the data-plane of the same application may be satisfied with a Session-lasting IP Address.

3.4. On Demand Nature

At any point in time, a mobile host may have a combination of IP addresses configured. Zero or more Fixed, zero or more Session-lasting, zero or more Non-persistent and zero or more Graceful-Replacement IP addresses may be configured by the IP stack of the host. The combination may be as a result of the host policy, application demand, or a mix of the two.

When an application requires a specific type of IP address and such an address is not already configured on the host, the IP stack SHALL attempt to configure one. For example, a host may not always have a Session-lasting IP address available. When an application requests one, the IP stack SHALL make an attempt to configure one by issuing a request to the network. If the operation fails, the IP stack SHALL fail the associated socket request and return an error. If successful, a Session-lasting IP Address gets configured on the mobile host. If another socket requests a Session-lasting IP address at a later time, the same IP address may be served to that socket as well. When the last socket using the same configured IP address is closed, the IP address may be released or kept for future applications that may be launched and require a Session-lasting IP address.

In some cases it might be preferable for the mobile host to request a new Session-lasting IP address for a new opening of an IP socket (even though one was already assigned to the mobile host by the network and might be in use in a different, already active IP sockets). It is outside the scope of this specification to define criteria for choosing to use available addresses or choosing to request new ones. It supports both alternatives (and any combination).

It is outside the scope of this specification to define how the host requests a specific type of prefix and how the network indicates the type of prefix in its advertisement or in its reply to a request.

The following are matters of policy, which may be dictated by the host itself, the network operator, or the system architecture standard:

- The initial set of IP addresses configured on the host at boot time.

- Permission to grant various types of IP addresses to a requesting application.

- Determination of a default address type when an application does not make any explicit indication, whether it already supports the required API or it is just a legacy application.

4. Backwards Compatibility Considerations

Backwards compatibility support is REQUIRED by the following 3 types of entities:

- The Applications on the mobile host

- The IP stack in the mobile host

- The network infrastructure

4.1. Applications

Legacy applications that do not support the On-Demand functionality will use the legacy API and will not be able to take advantage of the On-Demand Mobility feature.

Applications using the new On-Demand functionality should be aware that they may be executed in legacy environments that do not support it. Such environments may include a legacy IP stack on the mobile host, legacy network infrastructure, or both. In either case, the API will return an error code and the invoking applications may just give up and use legacy calls.

4.2. IP Stack in the Mobile Host

New IP stacks (that implement On Demand functionality) MUST continue to support all legacy operations. If an application does not use On-Demand functionality, the IP stack MUST respond in a legacy manner.

If the network infrastructure supports On-Demand functionality, the IP stack SHOULD follow the application request: If the application requests a specific address type, the stack SHOULD forward this request to the network. If the application does not request an address type, the IP stack MUST NOT request an address type and leave it to the network's default behavior to choose the type of the allocated IP prefix. If an IP prefix was already allocated to the host, the IP stack uses it and may not request a new one from the network.

4.3. Network Infrastructure

The network infrastructure may or may not support the On-Demand functionality. How the IP stack on the host and the network infrastructure behave in case of a compatibility issue is outside the scope of this API specification.

4.4. Merging this work with RFC5014

[RFC5014] defines new flags that may be used with setsockopt() to influence source IP address selection for a socket. The list of flags include: source home address, care-of address, temporary address, public address CGA (Cryptographically Created Address) and non-CGA. When applications require session continuity service, they SHOULD NOT set the flags specified in [RFC5014].

However, if an application erroneously performs a combination of (1) Use setsockopt() to set a specific option (using one of the flags specified in [RFC5014]) and (2) Selects a source IP address type, the

IP stack will fulfill the request specified by (2) and ignore the flags set by (1).

5. Security Considerations

The different service types (session continuity types and address reachability) associated with the allocated IP address types, may be associated with different costs. The cost to the operator for enabling a type of service, and the cost to applications using a selected service. A malicious application may use these to generate extra billing of a mobile subscriber, and/or impose costly services on the mobile operator. When costly services are limited, malicious applications may exhaust them, preventing other applications on the same mobile host from being able to use them.

Mobile hosts that enables such service options, should provide capabilities for ensuring that only authorized applications can use the costly (or limited) service types.

The ability to select service types requires the exchange of the association of source IP prefixes and their corresponding service types, between the mobile host and mobile network. Exposing these associations may provide information to passive attackers even if the traffic that is used with these addressed is encrypted.

To avoid profiling an application according to the type of IP addresses, it is expected that prefixes provided by the mobile operator are associated to various type of addresses over time. As a result, the type of address could not be associated to the prefix, making application profiling based on the type of address harder.

The application or the OS should ensure that IP addresses regularly change to limit IP tracking by a passive observer. The application should regularly set the On Demand flag. The application should be able to ensure that session lasting IP addresses are regularly changed by setting a lifetime for example handled by the application. In addition, the application should consider the use of graceful replacement IP addresses.

Similarly, the OS may also associated IP addresses with a lifetime. Upon receiving a request for a given type of IP address, after some time, the OS should request a new address to the network even if it already has one IP address available with the requested type. This includes any type of IP address. IP addresses of type graceful replacement or non persistent should be regularly renewed by the OS.

The lifetime of an IP address may be expressed in number of seconds or in number of bytes sent through this IP address.

6. IANA Considerations

This document has no IANA considerations.

7. Contributors

This document was merged with [I-D.sijeon-dmm-use-cases-api-source]. We would like to acknowledge the contribution of the following people to that document as well:

Sergio Figueiredo
Altran Research, France
Email: sergio.figueiredo@altran.com

Younghan Kim
Soongsil University, Korea
Email: younghak@ssu.ac.kr

John Kaippallimalil
Huawei, USA
Email: john.kaippallimalil@huawei.com

8. Acknowledgements

We would like to thank Wu-chi Feng, Alexandru Petrescu, Jouni Korhonen, Sri Gundavelli, Dave Dolson Lorenzo Colitti and Daniel Migault for their valuable comments and suggestions on this work.

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>.

[RFC5014] Nordmark, E., Chakrabarti, S., and J. Laganier, "IPv6 Socket API for Source Address Selection", RFC 5014, DOI 10.17487/RFC5014, September 2007, <https://www.rfc-editor.org/info/rfc5014>.

[RFC8174] Leiba, B., "Ambiguity of Uppercase vs Lowercase in RFC 2119 Key Words", BCP 14, RFC 8174, DOI 10.17487/RFC8174, May 2017, <https://www.rfc-editor.org/info/rfc8174>.

9.2. Informative References

[I-D.sijeon-dmm-use-cases-api-source]
Jeon, S., Figueiredo, S., Kim, Y., and J. Kaippallimalil, "Use Cases and API Extension for Source IP Address Selection", draft-sijeon-dmm-use-cases-api-source-07 (work in progress), September 2017.

[RFC3261] Rosenberg, J., Schulzrinne, H., Camarillo, G., Johnston, A., Peterson, J., Sparks, R., Handley, M., and E. Schooler, "SIP: Session Initiation Protocol", RFC 3261, DOI 10.17487/RFC3261, June 2002, <https://www.rfc-editor.org/info/rfc3261>.

[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>.

[RFC5563] Leung, K., Dommety, G., Yegani, P., and K. Chowdhury, "WiMAX Forum / 3GPP2 Proxy Mobile IPv4", RFC 5563, DOI 10.17487/RFC5563, February 2010, <https://www.rfc-editor.org/info/rfc5563>.

[RFC5944] Perkins, C., Ed., "IP Mobility Support for IPv4, Revised", RFC 5944, DOI 10.17487/RFC5944, November 2010, <https://www.rfc-editor.org/info/rfc5944>.

[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>.

[RFC6824] Ford, A., Raiciu, C., Handley, M., and O. Bonaventure, "TCP Extensions for Multipath Operation with Multiple Addresses", RFC 6824, DOI 10.17487/RFC6824, January 2013, <https://www.rfc-editor.org/info/rfc6824>.

[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>.

Appendix A. Conveying the Desired Address Type

Following are some suggestions of possible extensions to the Socket API for enabling applications to convey their session continuity and address reachability requirements.

[RFC5014] introduced the ability of applications to influence the source address selection with the IPV6_ADDR_PREFERENCE option at the IPPROTO_IPV6 level. This option is used with setsockopt() and getsockopt() calls to set/get address selection preferences.

One alternative is to extend the defintion of the IPV6_ADDR_REFERENCE opion with flags that express the invoker's desire. An "OnDeman" field could contains one of the values: FIXED_IP_ADDRESS, SESSION_LASTING_IP_ADDRESS, NON_PERSISTENT_IP_ADDRESS or GRACEFUL_REPLACEMENT_IP_ADDRESS.

Another alternative is to define a new Socket function used by the invoker to convey its desire. This enables the implementation of two behaviors of Socket functions: The existing "setsockotp()" is a function that returns after executing, and the new "setsc()" (Set Service Contionuity) function that may initaite a request for the desired service, and wait until the network responds with the allocated resources, before returning to the invoker.

After obtaining an IP address with the desired behavior the application can call the bind() Socket function to associate that received IP address with the socket.

Authors' Addresses

Alper Yegin
Actility
Istanbul
Turkey

Email: alper.yegin@actility.com

Danny Moses
Intel Corporation
Petah Tikva
Israel

Email: danny.moses@intel.com

Seil Jeon
Sungkyunkwan University
Suwon
South Korea

Email: seiljeon@skku.edu

DMM WG
Internet-Draft
Intended status: Informational
Expires: October 5, 2016

S. Gundavelli
Cisco

April 3, 2016

DMM Deployment Models and Architectural Considerations
draft-wt-dmm-deployment-models-00.txt

Abstract

This document identifies the deployment models for Distributed Mobility Management architecture.

Status of this Memo

This Internet-Draft will expire on October 5, 2016.

Copyright Notice

Table of Contents

1. Overview

One of the key aspects of the Distributed Mobility Management (DMM) architecture is the separation of control plane (CP) and data plane (DP) functions of a network element. While data plane elements continue to reside on customized networking hardware, the control plane resides as a software element in the cloud. This is usually referred to as CP-DP separation and is the basis for the IETF’s DMM Architecture. This approach of centralized control plane and distributed data plane allows elastic scaling of control plane and efficient use of common data plane that is agnostic to access architectures.

This document identifies the functions in the DMM architecture and the supported deployment models.

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 the mobility related terms are to interpreted as defined in [RFC6275], [RFC5213], [RFC5844], [RFC7333], [RFC7429], [I-D.ietf-sfc-nsh] and [I-D.ietf-dmm-fpc-cpdp]. Additionally, this document uses the following terms:

Home Control-Plane Anchor (H-CPA)

The Home-CPA function hosts the mobile node’s mobility session. There can be more than one mobility session for a mobile node [MN] and those sessions may be anchored on the same or different Home-CPA’s. The home-CPA will interface with the home-dpa for managing the forwarding state.

Home Data Plane Anchor (Home-DPA)

The Home-DPA is the topological anchor for the mobile node’s IP address/prefix(es). The Home-DPA is chosen by the Home-CPA on a session-basis. The Home-DPA is in the forwarding path for all the mobile node’s IP traffic.

Access Control Plane Node (Access-CPN)

The Access-CPN is responsible for interfacing with the mobile node's Home-CPA and with the Access-DPN. The Access-CPN has a protocol interface to the Home-CPA.

Access Data Plane Node (Access-DPN)

The Access-DPN function is hosted on the first-hop router where the mobile node is attached. This function is not hosted on a layer-2 bridging device such as a eNode(B) or Access Point.

3. DMM Architectural Overview

Following are the key goals of the Distributed Mobility Management architecture.

1. Separation of control and data Plane

2. Aggregation of control plane for elastic scaling

3. Distribution of the data plane for efficient network usage

4. Elimination of mobility state from the data plane

5. Dynamic selection of control and data plane nodes

6. Enabling the mobile node with network properties

7. Relocation of anchor functions for efficient network usage

3.1. DMM Service Primitives

The functions in the DMM architecture support a set of service primitives. Each of these service primitives identifies a specific service capability with the exact service definition. The functions in the DMM architecture are required to support a specific set of service primitives that are mandatory for that service function. Not all service primitives are applicable to all DMM functions. The below table identifies the service primitives that each of the DMM function SHOULD support. The marking "X" indicates the service primitive on that row needs to be supported by the identified DMM function on the corresponding column; for example, the IP address management must be supported by Home-CPA function.

Service Primitive	H-CPA	H-DPA	A-CPN	A-DPN	MC	RC
IP Management	X				X	
IP Anchoring		X				
MN Detect			X	X		
Routing		X		X		
Tunneling		X		X		
QoS Enforcement		X		X		
FPC Client	X		X		X	
FPC Agent		X		X		X
NSH Classifier		X		X		

Figure 1: Mapping of DMM functions

3.2. DMM Functions and Interfaces

3.2.1. Home Control-Plane Anchor (H-CPA):

The Home-CPA function hosts the mobile node's mobility session. There can be more than one mobility session for a mobile node and those sessions may be anchored on the same or different Home-CPA's. The home-CPA will interface with the homd-dpa for managing the forwarding state.

There can be more than one Home-CPA serving the same mobile node at a given point of time, each hosting a different control plane session.

The Home-CPA is responsible for life cycle management of the session, interfacing with the policy infrastructure, policy control and interfacing with the Home-DPA functions.

The Home-CPA function typically stays on the same node. In some special use-cases (Ex: Geo-Redundancy), the session may be migrated to a different node and with the new node assuming the Home-CPA role for that session.

3.2.2. Home Data-Plane Anchor (H-DPA):

The Home-DPA is the topological anchor for the mobile node’s IP address/prefix(es). The Home-DPA is chosen by the Home-CPA/MC on a session-basis. The Home-DPA is in the forwarding path for all the mobile node’s IP traffic.

As the mobile node roams in the mobile network, the mobile node’s access-DPN may change, however, the Home-DPA does not change, unless the session is migrated to a new node.

The Home-DPA interfaces with the Home-CPA/MC for all IP forwarding and QoS rules enforcement.

The Home-DPA and the Access-DPN functions may be collocated on the same node.

3.2.3. Access Control Plane Node (Access-CPN)

The Access-CPN is responsible for interfacing with the mobile node’s Home-CPA and with the Access-DPN. The Access-CPN has a protocol interface to the Home-CPA.

The Access-CPN is responsible for the mobile node’s Home-CPA selection based on: Mobile Node’s Attach Preferences, Access and Subscription Policy, Topological Proximity and Other Considerations.

The Access-CPN function is responsible for MN’s service authorization. It will interface with the access network authorization functions.

3.2.4. Access Data Plane Node (Access-DPN)

The Access-DPN function is hosted on the first-hop router where the mobile node is attached. This function is not hosted on a layer-2 bridging device such as a eNode(B) or Access Point.

The Access-DPA will have a protocol interface to the Access-CPA.

The Access-DPN and the Home-DPA functions may be collocated on the same node.

3.2.5. DMM Function Mapping to other Architectures

Following table identifies the potential mapping of DMM functions to protocol functions in other system architectures.

FUNCTION	PMIPv6	MIPv6	IPsec	3GPP	Broadband
Home-CPA	LMA-CPA	HA-CPA	IKE-CPA	PGW-CPA	BNG-CPA
Home-DPA	LMA-DPA	HA-DPA	IKE-DPA	PGW-DPA	BNG-DPA
Access-CPN	MAG-CPN	-	-	SGW-CPN	RG-CPN
Access-DPN	MAG-DPN	-	-	SGW-DPN	RG-DPN

Figure 2: Mapping of DMM functions

4. Deployment Models

This section identifies the key deployment models for the DMM architecture.

4.1. Model-1: Split Home Anchor Mode

In this model, the control and the data plane functions of the home anchor are separated and deployed on different nodes. The control plane function of the Home anchor is handled by the Home-CPA and where as the data plane function is handled by the Home-DPA. In this model, the access node operates in the legacy mode with the integrated control and user plane functions.

The FPC interface defined in [I-D.ietf-dmm-fpc-cpdp] allows the control plane functions to interact with the data plane for the subscriber’s forwarding state management.

```
                          +============+
                          |   Policy   |
             . . . . . . .|  Function  |. . . . . . .
             .            +============+            .
             .                                      .
             .                                      .
       +============+       {PMIPv6/GTP}      +============+
       |            |- - - - - - - - - - - - -|  Home-CPA  |
       |            |                         +============+
       |            |                               .
       |            |                               .  FPC
       | Access Node|                               .
       |            |                               .
       | (CPN + DPN)|                               .
       |            |                         +============+
       |   Legacy   |. . . . . . . . . . . . .|  Home-DPA  |
       +============+   UP {Tunnel/Route}     +============+
                .
                .
               [MN]
```

Figure 3: Split Home Anchor Mode

4.2. Model-2: Seperated Control and User Plane Mode

In this model, the control and the data plane functions on both the home anchor and the access node are seperated and deployed on different nodes. The control plane function of the Home anchor is handled by the Home-CPA and where as the data plane function is handled by the Home-DPA. The control plane function of the access node is handled by the Access-CPN and where as the data plane function is handled by the Access-DPN.

The FPC interface defined in [I-D.ietf-dmm-fpc-cpdp] allows the control plane functions of the home and access nodes to interact with the respective data plane functions for the subscriber’s forwarding state management.

```
                          +============+
                          |   Policy   |
             . . . . . . .|  Function  |. . . . . . .
             .            +============+            .
             .                                      .
             .                                      .
             .                                      .
             .                                      .
       +============+    {PMIPv6/GTP}       +============+
       | Access-CPN |- - - - - - - - - - - -|  Home-CPA  |
       +============+                       +============+
             .                                      .
             .  FPC                                 .  FPC
             .                                      .
             .                                      .
             .                                      .
       +============+                       +============+
       | Access-DPN |. . . . . . . . . . .  |  Home-DPA  |
       +============+   UP {Tunnel/Route}   +============+
              .
              .
             [MN]
```

Figure 4: Seperated Control and User Plane Mode

4.3. Model-3: Centralized Control Plane Mode

In this model, the control-plane functions of the home and the access nodes are collapsed. This is a flat architecture with no signaling protocol between the access node and home anchors. The interface between the Home-CPA and the Access-DPN is internal to the system.

The FPC interface defined in [I-D.ietf-dmm-fpc-cpdp] allows the mobility controller to interact with the respective data plane functions for the subscriber’s forwarding state management.

```
                   +=======================+     +============+
                   | Home-CPA + Access-CPN |     |  Policy    |
                   |                       |-----| Function   |
                   +=======================+     +============+
                              .
                           .     .
                        .          .
                  FPC .              .  FPC
                       .               .
                     .                   .
         +============+                   +============+
         | Access-DPN |. . . . . . .  . . . .|  Home-DPA  |
         +============+ UP {Tunnel/Route}  +============+
               .
               .
              [MN]
```

Figure 5: Centralized Control Plane Mode

4.4. Model-4: Data Plane Abstraction Mode

In this model, the data plane network is completely abstracted from the control plane. There is a new network element, Routing Controller which abstracts the entire data plane network and offers data plane services to the control plane functions. The control plane functions, Home-CPA and the Access-CPN interface with the Routing Controller for the forwarding state management.

The FPC interface defined in [I-D.ietf-dmm-fpc-cpdp] allows the Home-CPA and Access-CPN functions to interface with the Routing Controller for subscriber's forwarding state management.

```
                          +============+
                          |   Policy   |
             . . . . . . .|  Function  |. . . . . . .
             .            +============+            .
             .                                      .
             .                                      .
             .                                      .
       +============+    {PMIPv6/GTP}        +============+
       | Access-CPN |- - - - - - - - - - - - |  Home-CPA  |
       +============+                        +============+
             .                                      .
             .                                      .
             .                                      .
             .           +============+             .
             . . . . . . |  Routing   | . . . . . . .
                         | Controller |
                         +============+
                                .
                             .     .
                          .          .  BGP/Others
                        .              .
                      .                  .
                    .                      .
          +============+                 +============+
          | Access-DPN |. . . . . .  . . . .|  Home-DPA  |
          +============+ UP {Tunnel/Route}  +============+
                .
                .
               [MN]
```

Figure 6: Data Plane Abstraction Mode

4.5. On-Demand Control Plane Orchestration Mode

In this model, there is a new function Mobility Controller which manages the orchestration of Access-CPN and Home-CPA functions. The Mobility Controller allocates the Home-CPA and Access-DPN

```
   + - - - - - - - - - - - - - - - - - - - - - - - - - - - -+
   |    +----------+     +----------+     +----------+    |
        |Access-CPN|     |Access-CPN|     |Access-CPN|
   |    +----------+     +----------+     +----------+    |

   |    +----------+     +----------+     +----------+    |
        | Home-CPA |     | Home-CPA |     | Home-CPA |
   |    +----------+     +----------+     +----------+    |
   + - - - - - - - - - - - - - - - - - - - - - - - - - - - -+
              .                 .
              .                 .
              .                 .
              .          +============+     +============+
              .          |   Mobility |     |  Policy    |
              .          | Controller |-----| Function   |
              .          +============+     +============+
              .
              .
              .
              .          +============+
              . . . . . .|  Routing   |
                         | Controller |
                         +============+
                                .
                                .
                                .
   + - - - - - - - - - - - - - - - - - - - - - - - - - - - -+
   |    +----------+     +----------+     +----------+    |
        |Access-DPN|     |Access-DPN|     |Access-DPN|
   |    +----------+     +----------+     +----------+    |

   |    +----------+     +----------+     +----------+    |
        | Home-DPA |     | Home-DPA |     | Home-DPA |
   |    +----------+     +----------+     +----------+    |
   + - - - - - - - - - - - - - - - - - - - - - - - - - - - -+
```

Figure 7: On-Demand CP Orchestration Mode

5. IANA Considerations

This document does not require any IANA actions.

6. Security Considerations

The control-plane messages exchanged between a Home-CPA and the Home-DPA must be protected using end-to-end security associations with data-integrity and data-origination capabilities.

IPsec ESP in transport mode with mandatory integrity protection should be used for protecting the signaling messages. IKEv2 should be used to set up security associations between the Home-CPA and Home-DPA.

There are no additional security considerations other than what is presented in the document.

7. Work Team

This document reflects contributions from the following work team members:

Seil Jeon

seiljeon@av.it.pt

Younghan Kim

younghak@ssu.ac.kr

Vic Liu

liuzhiheng@chinamobile.com

Danny S Moses

danny.moses@intel.com

Marco Liebsch

liebsch@neclab.eu

Carlos Jesus Bernardos Cano

cjbc@it.uc3m.es

8. Acknowledgements

This document is a result of DMM WT#4 team discussions and ideas taken from several DMM WG presentations and documents including, draft-sijeon-dmm-deployment-models, draft-liu-dmm-deployment-scenario and others. The work teams would like to thank the authors of these documents and additionally the discussions in DMM Working group that helped shape this document.

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, <http://www.rfc-editor.org/info/rfc2119>.

9.2. Informative References

[I-D.ietf-dmm-fpc-cpdp] Liebsch, M., Matsushima, S., Gundavelli, S., Moses, D., and L. Bertz, "Protocol for Forwarding Policy Configuration (FPC) in DMM", draft-ietf-dmm-fpc-cpdp-03 (work in progress), March 2016.

[I-D.ietf-sfc-nsh] Quinn, P. and U. Elzur, "Network Service Header", draft-ietf-sfc-nsh-04 (work in progress), March 2016.

[RFC5213] Gundavelli, S., Ed., Leung, K., Devarapalli, V., Chowdhury, K., and B. Patil, "Proxy Mobile IPv6", RFC 5213, DOI 10.17487/RFC5213, August 2008, <http://www.rfc-editor.org/info/rfc5213>.

[RFC5844] Wakikawa, R. and S. Gundavelli, "IPv4 Support for Proxy Mobile IPv6", RFC 5844, DOI 10.17487/RFC5844, May 2010, <http://www.rfc-editor.org/info/rfc5844>.

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

[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, <http://www.rfc-editor.org/info/rfc7333>.

[RFC7429] Liu, D., Ed., Zuniga, JC., Ed., Seite, P., Chan, H., and CJ. Bernardos, "Distributed Mobility Management: Current Practices and Gap Analysis", RFC 7429, DOI 10.17487/RFC7429, January 2015, <http://www.rfc-editor.org/info/rfc7429>.

Author's Address

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

Email: sgundave@cisco.com