

Network Working Group
Internet-Draft
Intended status: Standards Track
Expires: April 5, 2013

Y. Cui
Tsinghua University
X. Xu
WD. Wang
XM. Li
Beijing University of Posts and
Telecommunications
YZ. Huo
W. Luo
ZTE Corporation
October 2, 2012

Seamless Handover for Multiple-Access Mobile Node in PMIPv6
draft-cui-netext-pmipv6-shpmipv6-00

Abstract

Proxy Mobile IPv6 (PMIPv6), specified in [RFC5213], provide a mobile node(MN) which requires no additional modification to MN with IP mobility. Fast Handover for Proxy Mobile IPv6 (FHPMIPv6), specified in[RFC5949], proposed two modes of fast handover, both of them use single interface to transmate packets during handover, which requires it to buffer packets in MAGs when interface performs handover. Buffer packets in MAGs result in additional overhead, and increase packets transmission delay. Unlike FHPMIPv6, this document proposed a seamless handover scheme for multi-access mobile node with IP mobility when one of MN's network interface performs handover from one MAG to another. This scheme uses some other interface of the multi-access mobile node to help process packets while handovering.

Status of This Memo

This Internet-Draft is submitted in full conformance with the provisions of BCP 78 and BCP 79.

Internet-Drafts are working documents of the Internet Engineering Task Force (IETF). Note that other groups may also distribute working documents as Internet-Drafts. The list of current Internet-Drafts is at <http://datatracker.ietf.org/drafts/current/>.

Internet-Drafts are draft documents valid for a maximum of six months and may be updated, replaced, or obsoleted by other documents at any time. It is inappropriate to use Internet-Drafts as reference material or to cite them other than as "work in progress."

This Internet-Draft will expire on April 5, 2013.

Copyright Notice

Copyright (c) 2012 IETF Trust and the persons identified as the document authors. All rights reserved.

This document is subject to BCP 78 and the IETF Trust's Legal Provisions Relating to IETF Documents (<http://trustee.ietf.org/license-info>) in effect on the date of publication of this document. Please review these documents carefully, as they describe your rights and restrictions with respect to this document. Code Components extracted from this document must include Simplified BSD License text as described in Section 4.e of the Trust Legal Provisions and are provided without warranty as described in the Simplified BSD License.

Table of Contents

1. Introduction	3
2. Requirements Language	3
3. Terminology	3
4. Protocol overview	4
4.1. Protocol Operation	5
4.2. Mobile Node considerations	7
4.3. Mobile Access Gateway considerations	7
4.4. Local Mobility Anchor considerations	8
5. Message Formats	8
5.1. Streamless Handover Initiate (SHI) message	8
5.2. Streamless Handover Acknowledge (SHAck) message	9
6. IANA Considerations	10
7. Security Considerations	10
8. Normative References	10

1. Introduction

With the development of internet access technologies and mobile terminal equipment, more and more hosts are operating in multiple-interfaces, thus a terminal having access to multiple heterogeneous network domain simultaneously has become possible. Proxy Mobile IPv6 is a network-based mobility protocol, it provides mobility support for mobile node and requires no additional modification.

RFC 5949 FHPMIPv6 is a fast handover extension for PMIPv6, the document proposed two modes of fast handover: reactive mode and predictive mode. The main idea of the two modes of operations is to establish a bi-directional tunnel between the Previous Mobile Access Gateway (PMAG) and the New Mobile Access Gateway (NMAG). So, packets destined for the Mobile Node are forward from the PMAG to the NMAG over this tunnel. Both of the two modes of fast handover improve the handover performance in terms of packet loss and latency, while none of them takes full advantage of multi-access features of the mobile node, as in both of the two handover modes, packets transmission on the handover interface should be buffered at the PMAG or NMAG which increases the requirement of storage volume for the MAG. When there are many MNs are handovering within the coverage area of the same MAG, some packets may be lost due to cache insufficiency. The two modes adopt cached and forwarded to deal with the packet while handover will greatly increase the transmission delay, that may be deadly to delay-sensitive applications.

This document propose a seamless handover scheme for multiple-access mobile node in PMIPv6, compared with the two kinds of handover modes mentioned above. This seamless handover scheme doesn't need to buffer the packets in MAG, which reduces the requirements on the MAG cache, while reducing the transmission delay at the same time.

2. Requirements Language

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

3. Terminology

The following terminologies used in this document are define in RFC5213:

Local Mobility Anchor (LMA).

Mobile Access Gateway (MAG).

Proxy Mobile IPv6 Domain (PMIPv6-Domain).

The following terminologies used in this document are define in RFC5949:

Previous Mobile Access Gateway (PMAG).

New Mobile Access Gateway (NMAG).

The following terminologies are define and used in this document:

Stable Mobile Access Gateway (SMAG)

while one of MN's interface is handovering, The MAGs that connect with some other interface of MN are called SMAG.

4. Protocol overview




Figure 1 reference network for Multiple-Access Mobile Node handover

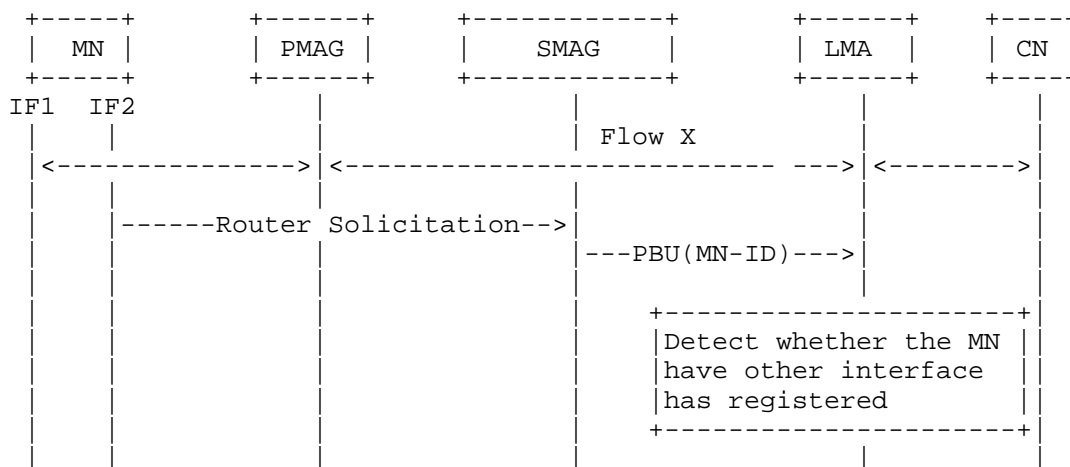
In order to alleviate the packet loss during handover, RFC5949 proposed two kinds of fast handover schemes. In both of the two handover scheme, the downlink packets need to be buffered either at

the PMAG or NMAG, depending on when the packet forwarding is performed. This buffer and forwarding mechanism increase the cache overhead in MAG and increase the data transmission delay. In this document, we assume that mobile node have multiple network interface access different MAG in the same PMIPv6-Domain and support weak host model, that means MN can receive any locally destined packet regardless of the network interface on which the packet was received. The deployment scenario is illustrated in Figure 1.

In order to improve the performance during handover and reduce the demand of the MAG buffer capacity, this document specifies a bi-directional tunnel between the PMAG and SMAG to forward packets for mobile node. If an interface is handovering, the packets transmission on this interface was forwarded to SMAG then forwarded to some other interface of MN. In order to build a bi-directional tunnel between the PMAG and the SMAG, a new message called Streamless Handover Initiate(SHI) and Streamless Handover Acknowledge (SHIA) was define in Section 5. When multi-interface MN attach to MAG, MAG will send PBU register message to LMA, then receive a PBA message if register succeeded, MAG will send SHI message to MAGs that connect with MN's interface. Necessary extensions to LMA and MAG need to support this handover scheme and the extensions are define in section 4.3 and section 4.4.

4.1. Protocol Operation

Unlike Predictive Fast Handover and Reactive Fast Handover, this protocol build a bi-directional tunnel between MAGs that different interfaces of the mobile node connects to. The sequence of event for the seamless handover scheme for Multiple-Access Mobile Node is illustrated in Figure 2.



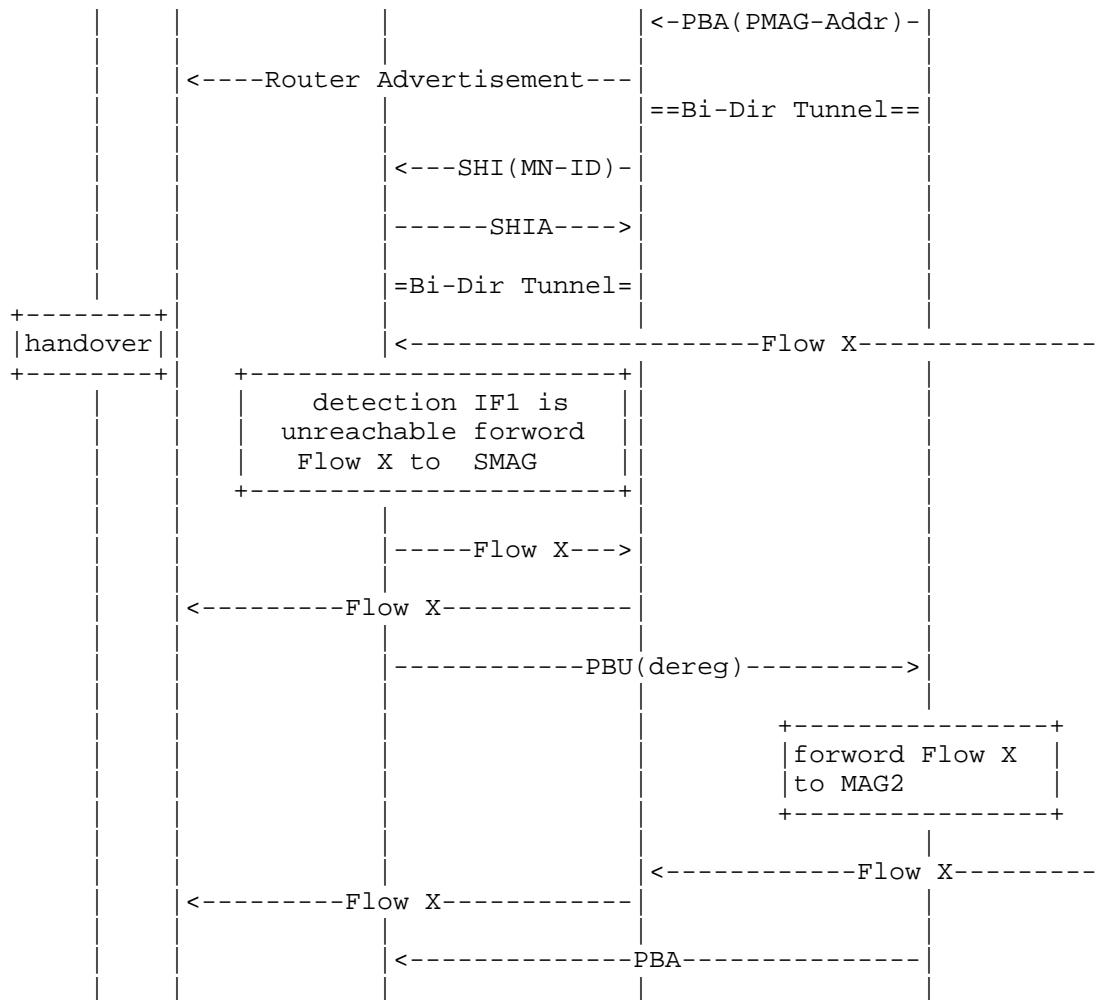


Figure 2 the signaling process of streamless handover scheme for Multiple-Access Mobile Node

The detailed descriptions are as follows:

- o In the proxy mobile ipv6 network domain, MN has multiple interface as illustrated in Figure 1, assumes that interface 1(IF1) has already accessed PMIPv6 domain and data flow X transmitted through the interface.
- o IF2 accesses SMAG, and SMAG sends PBU message to LMA. If register success, LMA send back PBA message to SMAG.

- o When SMAG receives PBA message, it sends SHI message to PMAG, noticing that the SHI message must include the MN-ID option. when PMAG receives SHI message and finds that the MN connects to it, PMAG sends a SHIA message to SMAG, otherwise PMAG send back MN not attached SHIA message. When all this done, PMAG and SMAG detect whether there exists any tunnel between them, if not, it will build a bi-directional tunnel between them. notice that the tunnel between MAGs are per-MAG-MAG.
- o When IF1 performs a handover, first, if PMAG detects IF1 is unreachable, it change the router and forwards the packet that destination address is IF1 to SMAG. In this case , the transmission path of flow X is LMA->PMAG->SMAG->IF2. Then PMAG sends the DeReg PBU message to LMA.
- o LMA receives the DeReg PBU message, first it changes the router and forwards the packet of destination address IF1 to SMAG,.In this case, the transmission path of flow X is LMA->SMAG->IF2. Then LMA sends back DeReg PBA message to PMAG.

4.2. Mobile Node considerations

In this document, we assume that mobile node has multiple network interfaces, and those interfaces access to the same PMIPv6-domain. and all of the MN's network interfaces configuration the same home network prefix. In order to support MNs that receive any locally destined packet regardless of the network interface on which the packet is received, the mobile node must support the weak host model. While interface is handovering, it may re-config its IP address and MN may not accept the packet that the destination address is the handover interface, in this document, we assume MN can accept the packet that the destination address is the handover interface's IP address temporarily while the interface is handovering(details are out of the scope of this document).

4.3. Mobile Access Gateway considerations

In the seamless handover scheme, when MAG receive a PBA message, it need to send SHI message to some other MAGs that connect to MN, in this document we assume that MAG knows the ip address of those MAG. Notice that the SHI message at least includes MN-Id option. When MAG receives SHI message, it detects whether the MN has a interface connected with it, if so, MAG sends SHIA response message, and bulids a bi-directional tunnel, otherwise, sends the response message of no such node.

When MAG detects the departure of the MN's network interface, it configures routing manner, the packets that sent to the interface are

forwarded through to SMAG through tunnels. As all the network interfaces of MN's configured the same home network prefix, MAG can forward packets to MN by prefix match.

4.4. Local Mobility Anchor considerations

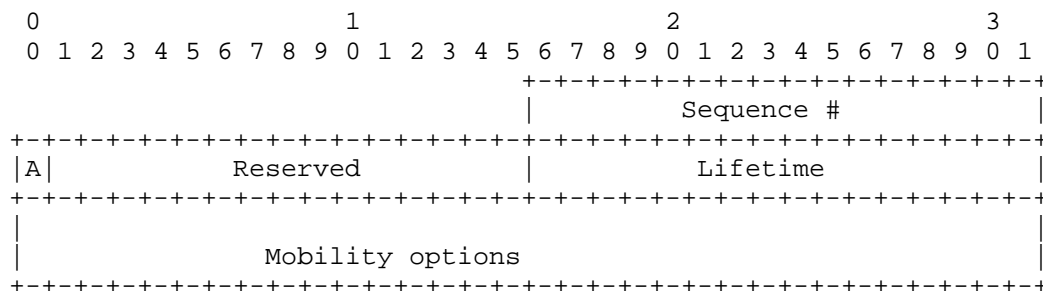
When LMA receives a PBU message, it needs to detect whether the MN has another interface accessed the PMIPv6-domain, and associate all MN's interface, in this document, we assume that LMA support flow mobility, as [I-D.ietf-netext-pmipv6-flowmob] described

5. Message Formats

This section defines new mobility header messages for seamless handover .

5.1. Streamless Handover Initiate (SHI) message

This message is created to build associate between MAGs that different interfaces of MN connect to. The format of the Message Data field in the Mobility Header is as follows:



Sequence #

Must be set by the sender so replies can be matched to this message.

'A' flag

The Acknowledge (A) bit is set to request a Streamless Handover Acknowledge be returned upon receipt of the SHI message.

Reserved

These fields are unused. They MUST be initialized to zero by the sender and MUST be ignored by the receiver

Liftime

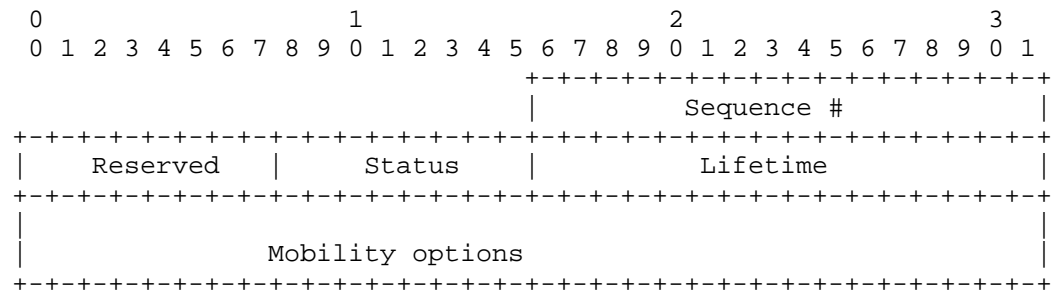
16-bit unsigned integer. It represents the tunnel survival time.

Mobility Option

Same as [RFC5213]

5.2. Streamless Handover Acknowledge (SHAck) message

The Streamless Handover Acknowledge is used to acknowledge receipt of a SHI message. The format of the Message Data field in the Mobility Header is as follows:



Sequence#

The Sequence Number in the Streamless Handover Acknowledge is copied from the Sequence Number field in the SHI message.

Reserved

These fields are unused. They MUST be initialized to zero by the sender and MUST be ignored by the receiver.

Lifetime

16-bit unsigned integer. It represents the tunnel survival time.

Status

0: succeeded

128: Reason unspecified

129: MN not attached

6. IANA Considerations

TBD

7. Security Considerations

TBD

8. Normative References

- [I-D.ietf-netext-pmipv6-flowmob] Bernardos, C., "Proxy Mobile IPv6 Extensions to Support Flow Mobility", draft-ietf-netext-pmipv6-flowmob-04 (work in progress), July 2012.
- [RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, March 1997.
- [RFC5213] Gundavelli, S., Leung, K., Devarapalli, V., Chowdhury, K., and B. Patil, "Proxy Mobile IPv6", RFC 5213, August 2008.
- [RFC5949] Yokota, H., Chowdhury, K., Koodli, R., Patil, B., and F. Xia, "Fast Handovers for Proxy Mobile IPv6", RFC 5949, September 2010.

Authors' Addresses

Yong Cui
Tsinghua University
Department of Computer Science, Tsinghua University
Beijing 100084
P.R.China

Phone: +86-10-6260-3059
EMail: cuiyong@tsinghua.edu.cn

Xin Xu
Beijing University of Posts and Telecommunications
Tsinghua University FIT Building 4-103
Beijing 100084
P.R.China

EMail: xuxin1988@gmail.com

Wendong Wang
Beijing University of Posts and Telecommunications
Room 609, teaching building 3,BUPT
Beijing 100876
P.R.China

EMail: wdwang@bupt.edu.cn

XiMing Li
Beijing University of Posts and Telecommunications
Tsinghua University FIT Building 4-103
Beijing 100084
P.R.China

EMail: xml@bupt.edu.cn

Yuzhen Huo
ZTE Corporation
No.68 Zijinghua Rd.,Yuhuatai District
Nanjing 210012
P.R.China

EMail: huo.yuzhen@zte.com.cn

Wen Luo
ZTE Corporation
No.68 Zijinghua Rd.,Yuhuatai District
Room 609, teaching building 3,BUPT 210012
P.R.China

EMail: EMail:luo.wen@zte.com.cn

NETEXT WG
Internet-Draft
Intended status: Informational
Expires: April 25, 2013

T. Melia, Ed.
Alcatel-Lucent
S. Gundavelli, Ed.
Cisco
October 22, 2012

Logical Interface Support for multi-mode IP Hosts
draft-ietf-netext-logical-interface-support-06.txt

Abstract

A Logical Interface is a software semantic internal to the host operating system. This semantic is available in all popular operating systems and is used in various protocol implementations. The Logical Interface support is required on the mobile node operating in a Proxy Mobile IPv6 domain, for leveraging various network-based mobility management features such as inter-technology handoffs, multihoming and flow mobility support. This document explains the operational details of Logical Interface construct and the specifics on how the link-layer implementations hide the physical interfaces from the IP stack and from the network nodes on the attached access networks. Furthermore, this document identifies the applicability of this approach to various link-layer technologies and analyzes the issues around it when used in context with various mobility management features.

Status of this Memo

This Internet-Draft is submitted in full conformance with the provisions of BCP 78 and BCP 79.

Internet-Drafts are working documents of the Internet Engineering Task Force (IETF). Note that other groups may also distribute working documents as Internet-Drafts. The list of current Internet-Drafts is at <http://datatracker.ietf.org/drafts/current/>.

Internet-Drafts are draft documents valid for a maximum of six months and may be updated, replaced, or obsoleted by other documents at any time. It is inappropriate to use Internet-Drafts as reference material or to cite them other than as "work in progress."

This Internet-Draft will expire on April 25, 2013.

Copyright Notice

Copyright (c) 2012 IETF Trust and the persons identified as the document authors. All rights reserved.

This document is subject to BCP 78 and the IETF Trust's Legal Provisions Relating to IETF Documents (<http://trustee.ietf.org/license-info>) in effect on the date of publication of this document. Please review these documents carefully, as they describe your rights and restrictions with respect to this document. Code Components extracted from this document must include Simplified BSD License text as described in Section 4.e of the Trust Legal Provisions and are provided without warranty as described in the Simplified BSD License.

Table of Contents

1. Introduction	4
2. Terminology	5
3. Hiding Link-layer Technologies - Approaches and Applicability	6
3.1. Link-layer Abstraction - Approaches	6
3.2. Applicability Statement	7
3.2.1. Link layer support	8
3.2.2. Logical Interface	8
4. Technology Use cases	10
5. Logical Interface Functional Details	11
5.1. Configuration of a Logical Interface	12
5.2. MTU considerations for a Logical Interface	13
5.3. Supported Link models for a logical interface	13
5.4. Link-layer Identifier Selection for a Logical Interface	13
5.5. ND Considerations for Logical Interface	14
5.6. Provisioning Domain Considerations	15
5.7. Logical Interface Forwarding Conceptual Data Structures	15
6. Logical Interface Use-cases in Proxy Mobile IPv6	17
6.1. Multihoming Support	17
6.2. Inter-Technology Handoff Support	18
6.3. Flow Mobility Support	20
7. IANA Considerations	21
8. Security Considerations	22
9. Authors	23
10. Acknowledgements	23

11. References	24
11.1. Normative References	24
11.2. Informative References	24
Authors' Addresses	25

1. Introduction

Proxy Mobile IPv6 [RFC5213] is a network-based mobility protocol. Some of the key goals of the protocol include support for multihoming, inter-technology handoffs and flow mobility support. The base protocol features specified in [RFC5213] and [RFC5844] allow the mobile node to attach to the network using multiple interfaces (simultaneously or sequentially), or to perform handoff between different interfaces of the mobile node. However, for supporting these features, the mobile node is required to be activated with specific software configuration that allows the mobile node to either perform inter-technology handoffs between different interfaces, attach to the network using multiple interfaces, or perform flow movement from one access technology to another. This document analyzes from the mobile node's perspective a specific approach that allows the mobile node to leverage these mobility features. Specifically, it explores the use of the Logical Interface support, a semantic available on most operating systems.

A Logical Interface is a construct internal to the operating system. It is an approach where the link-layer implementations hide the physical interfaces from the IP stack and from the network nodes on the attached access networks. This semantic is widely available in all popular operating systems. Many applications such as Mobile IP client [RFC6275] and IPsec VPN client [RFC4301] rely on this semantic for their protocol implementation and the same semantic can also be useful in this context. Specifically, the mobile node can use the logical interface configuration for leveraging various network-based mobility management features provided by the Proxy Mobile IPv6 domain [RFC5213].

The rest of the document provides the operational details of a Logical Interface on the mobile node and the inter-working between a mobile node using logical interface and network elements in the Proxy Mobile IPv6 domain when supporting some of the mobility management features. It also analyzes the issues involved with this approach and characterizes the contexts in which such usage is appropriate.

2. Terminology

All the mobility related terms used in this document are to be interpreted as defined in Proxy Mobile IPv6 specifications, [RFC5213] and [RFC5844]. In addition, this document introduces the following terms:

PIF (Physical Interface) - a network interface card attached to an host providing network connectivity (e.g. an Ethernet card, a WLAN card, an LTE interface).

LIF (Logical Interface) - It is a virtual interface in the IP stack. It appears just as any other physical interface, provides similar semantics with respect to packet transmit and receive functions to the upper layers in the IP stack. However, it is only logical construct and is not a representation of an instance of any physical hardware.

VLL-ID (Virtual Link-layer ID) - a virtual link-layer address configured on the logical interface. This identifier can be randomly generated, or configured based on the link-layer address of one of the physical interface.

Sub-If (Sub Interface) - a physical interface that is part of a logical interface construct. For example, a logical interface may have been created abstracting two physical interfaces, LTE and WLAN. These physical interfaces, LTE and WLAN are referred to as sub-interfaces of that logical interface. In some cases, a sub-interface can also be another logical interface, such as an IPsec tunnel interface.

3. Hiding Link-layer Technologies - Approaches and Applicability

There are several techniques/mechanisms that allow hiding access technology changes or movement from host IP layer. This section classifies these existing techniques into a set of generic approaches, according to their most representative characteristics. Later sections of this document analyze the applicability of these solution approaches for supporting features such as, inter-technology handovers and IP flow mobility support for a mobile node in a Proxy Mobile IPv6 domain [RFC5213].

3.1. Link-layer Abstraction - Approaches

The following generic mechanisms can hide access technology changes from host IP layer:

- o Link-layer Support - Certain link-layer technologies are able to hide physical media changes from the upper layers (see Figure 1). For example, IEEE 802.11 is able to seamlessly change between IEEE 802.11a/b/g physical layers. Also, an 802.11 STA can move between different Access Points within the same domain without the IP stack being aware of the movement. In this case, the IEEE 802.11 MAC layer takes care of the mobility, making the media change invisible to the upper layers. Another example is IEEE 802.3, that supports changing the rate from 10Mbps to 100Mbps and to 1000Mbps.

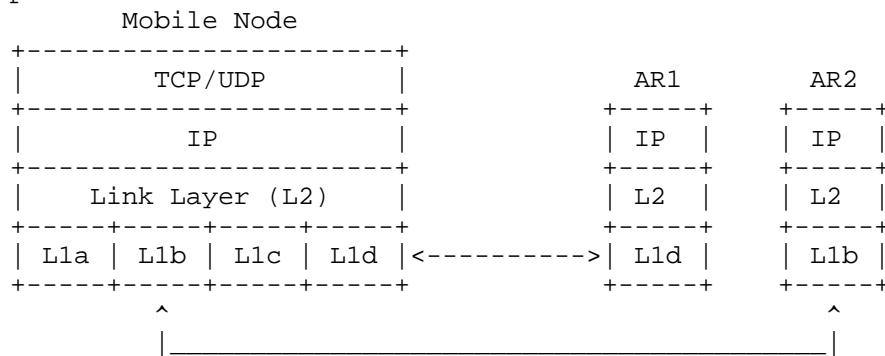


Figure 1: Link layer support solution architecture

There are also other examples with more complicated architectures, like for instance, 3GPP EPC [TS23401]. In this case, a UE can move (inter-RA handover) between GERAN/UTRAN/E-UTRAN, being this movement invisible to the IP layer at the UE, and also to the LMA logical component at the PGW. The link layer stack at the UE (i.e. PDCP and RLC layers), and the GTP between the RAN and the SGW (which plays the role of inter-3GPP AN mobility anchor) hide

this kind of mobility, which is not visible to the IP layer of the UE (see Figure 2).




Figure 2: 3GPP LTE/EPC data plane architecture (GTP option)

- o Logical interface: this refers to solutions (see Figure 3) that logically group/bond several physical interfaces so they appear to the upper layers (i.e. IP) as one single interface (where application sockets bind). Depending on the OS support, it might be possible to use more than one physical interface at a time -- so the node is simultaneously attached to different media -- or just to provide a fail-over mode. Controlling the way the different media is used (simultaneous, sequential attachment, etc) is not trivial and requires additional intelligence and/or configuration at the logical interface device driver. An example of this type of solution is the Logical interface, which is defined in this document, or the bonding driver (a Linux implementation).

3.2. Applicability Statement

We now focus on the applicability of the above solutions against the following requirements:

- o multi technology support
- o sequential vs. simultaneous access

3.2.1. Link layer support

Link layer mobility support applies to cases when the same link layer technology is used and mobility can be fully handled at these layers. One example is the case where several 802.11 access points are deployed in the same subnet and all of them share higher layer resources such as DHCP server, IP gateway, etc. In this case the access points can autonomously (or with the help of a central box) communicate and control the STA association changes from one AP to another, without the STA being aware of the movement. This type of scenario is applicable to cases when the different points of attachment (i.e. access points) belong to the same network domain, e.g. Enterprise, hotspots from same operator, etc.

This type of solution does not typically allow for simultaneous attachment to different access networks, and therefore can only be considered for inter-access technology handovers, but not for flow mobility. Existing RFC 5213 handover hint mechanisms could benefit from link layer information (e.g. triggers) to detect and identify MN handovers.

Link layer support is not applicable when two different access technologies are involved (e.g. 802.11 WLAN and 802.16 WiMAX) and the same is true when the same access technology expands over multiple network domains. This solution does not impose any change at the IP layer since changes in the access technology occur at layer two.

3.2.2. Logical Interface

The use of a logical interface allows the mobile node to provide a single interface view to the layers above IP (thus not changing the IP layer itself). Upper layers can bind to this interface, which hides inner inter-access technology handovers or data flow transfers among different physical interfaces.

This type of solution may support simultaneous attachment, in addition to sequential attachment. It requires additional support at the node and the network in order to benefit from simultaneous attachment. For example special mechanisms are required to enable addressing a particular interface from the network (e.g. for flow mobility). In particular extensions to PMIPv6 are required in order to enable the network (i.e., the MAG and LMA) to deal with logical interface, instead to IP interfaces as current RFC5213 does. RFC5213 assumes that each physical interface capable of attaching to a MAG is an IP interface, while the logical interface solution groups several physical interfaces under the same IP logical interface.

It is therefore clear that the Logical Interface approach satisfies

the multi technology and the sequential vs: simultaneous access support.

4. Technology Use cases

The 3GPP has defined the Evolved Packet Core (EPC) for heterogeneous wireless access. A mobile device equipped with 3GPP and non-3GPP wireless technologies can simultaneously or sequentially connect any of the available devices and receive IP services through any of them. This document focuses on the simultaneous/sequential use of these technologies and on the use cases that derive.

As mentioned in the previous sections the Logical Interface construct is required to hide the specifics of each technology in the context of network based mobility (e.g. in PMIPv6 deployments). The LIF concept can be used with at least the following technologies: 3GPP access technologies (3G, LTE), WIMAX access technology and IEEE 802.11 access technology.

3GPP In most OS implementations the connection setup establishes a PPP interface through the IPCP and IPv6CP protocol [RFC5072]. In this case the PPP interface does not have any L2 address assigned and does not generate any ARP or ND message for layer two address resolution. Conversely recent implementations configure an ethernet alike interface at OS level hiding to the upper layers the PPP nature of the connection. It has been verified (Android platform) that in these cases the ethernet alike interface configures a random L2 MAC address and uses this address as source link layer address option carried in the ND messages. ARP is also run between the mobile device and the remote peer (the network is a /30 address space).

WIMAX In WiMAX system also, the connection between the mobile station (MS) and the access router (AR) is a point-to-point link. The MS auto configures an address based on the prefix advertised by the AR or is assigned an address via DHCPv6. The stateless address auto-configuration is performed as per [RFC4861] and the IPv6 address is formed by adding an IID to the prefix learnt from Router Advertisement. IPv6 packets sent or received by the MS are identified by specific IDs, by which the AR can map them to the corresponding tunnel in the network.

5. Logical Interface Functional Details

This section identifies the functional details of a logical interface and provides some implementation considerations.

On most operating systems, a network interface is associated with a physical device that offers the services for transmitting and receiving IP packets to the applications on the host. In some configurations, a network interface can also be implemented as a logical interface which does not have the inherent capability to transmit, or receive packets on a physical medium, but relies on other physical interfaces for such services. Example of such configuration is an IP tunnel interface.

General overview of a logical interface is shown in Figure 3. The logical interface allows heterogeneous attachment while leaving the change in the media transparent to the IP stack. Simultaneous and sequential network attachment procedures are possible enabling inter-technology and flow mobility scenarios.

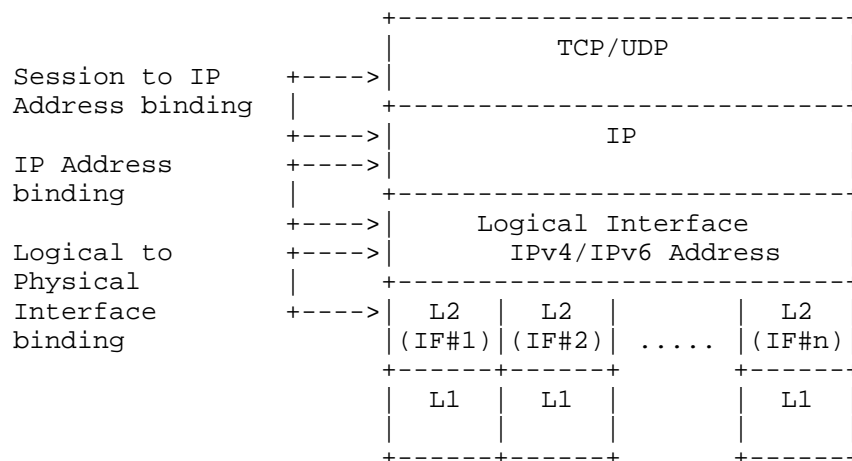


Figure 3: General overview of logical interface

From the perspective of the IP stack and the applications, a Logical interface is just another interface. In fact, the logical interface is only visible to the IP and upper layers when enabled. A host does not see any operational difference between a Logical and a physical interface. As with physical interfaces, a Logical interface is represented as a software object to which IP address configuration is bound. However, the Logical interface has some special properties which are essential for enabling inter-technology handover and flow-mobility features. Following are those properties:

1. The logical interface has a relation to a set of physical interfaces (sub-interfaces) on the host that it is abstracting. These sub-interfaces can be attached or detached from the Logical Interface at any time. The sub-interfaces attached to a Logical interface are not visible to the IP and upper layers.
 2. The logical Interface may either use a virtual interface identifier independent of the interface identifiers of its sub-interfaces, or it may use the link-layer identifier from one of its sub-interfaces.
 3. The logical interface has the path awareness with respect to the attached IP networks. For example, the logical interface may be bound to two IP networks, CAFE::/64 and BABA::/64, each of these prefixes may have been hosted on access networks attached through different sub-interfaces, WLAN and LTE. The logical interface has the path awareness with respect to IP network to sub-interface mapping.
 4. The logical interface may be attached to multiple access technologies with different link MTU values. The adopted MTU value for the logical interface must be lowest MTU value across those access technologies.
 5. The Transmit/Receive functions of the logical interface are mapped to the Transmit/Receive services exposed by the sub-interfaces. This mapping is dynamic and any change is not visible to the upper layers of the IP stack.
 6. The logical interface adapts to the point-to-point link model.
 7. The logical interface maintains IP flow information for each of its sub-interfaces. A conceptual data structure is maintained for this purpose. The host may populate this information based on tracking each of the sub-interface for the active flows.
- 5.1. Configuration of a Logical Interface

A host may be statically configured with the logical interface configuration, or an application such as a connection manager on the host may dynamically create it. Furthermore, the set of sub-interfaces that are part of a logical interface construct may be a fixed set, or may be kept dynamic, with the sub-interfaces getting added or deleted as needed. The specific details related to these configuration aspects are implementation specific and is outside the scope of this document.

5.2. MTU considerations for a Logical Interface

The link MTU (maximum transmission unit) value configured on a logical interface should be the lowest of the MTU values supported across any of the physical interfaces that are part of that logical interface construct. The MTU value should be configured as part of the logical interface creation on the host.

Furthermore, this value must be updated any time there is a change to the logical interface construct, such as when interfaces are added or deleted from the logical interface setup. Any time there is an inter-technology handover between two access technologies, the applications on the host bound to the IP address configuration on the logical interface will not detect the change and will continue to use the MTU value of the logical interface for the outbound packets, which is never greater than the MTU value on that supported access network. However, the access network may continue to deliver the packets conforming to the MTU value supported on that access technology and the logical interface should be able to receive those packets from the physical interface attached to that network. This approach of MTU configuration will ensure there is no IP packet fragmentation after inter-technology handovers.

5.3. Supported Link models for a logical interface

As per the base Proxy Mobile IPv6 specification [RFC5213] the media underneath the physical interface has to be bound to a point-to-point link [RFC5213]. Access technologies that provides a shared media (e.g., IEEE 802.11) can be supported as long as they provide a point-to-point link [RFC4861]. The details of how a shared media provides a point to point link are link layer specific and/or operational matters that are out of scope of this document. For example IEEE 802.11 media can provide a point-to-point link via the appropriate use of IEEE 802.1Q VLAN header where a distinct VLAN is configured between the MAG and each of the mobile node, or by the approach of MAG transmitting multicast packets as layer-2 unicast packets [RFC6085] and thereby preserving the point-to-point link properties on a shared link.

5.4. Link-layer Identifier Selection for a Logical Interface

The logical Interface may be configured to use the link-layer identifier from one of its sub-interfaces, or an identifier independent of the link-layer identifiers of the sub-interfaces. Following are the considerations.

- o In access architectures where it is possible to adopt a virtual link-layer identifier and use it for layer-2 communications in any of the access networks, a virtual identifier (VLL-Id) may be used. The specifics on how that identifier is chosen is out side the scope of this document. This identifier may be used for all link-layer communications. This identifier may also be used as the interface identifier when generating IPv6 global or link-local addresses, based on Stateless Autoconfiguration [RFC4862]
- o In access architectures, where the link-layer identifier is associated with a specific access technology, it will not be possible for the logical interface to adopt a virtual identifier and it use it across different access networks. In such networks, the logical interface must use the identifier of the respective sub-interface through which a packet is being transmitted. However, if more than one access technology domains that are part of the logical interface have such requirement, then the logical interface will not be able to support such configuration.

5.5. ND Considerations for Logical Interface

The following are the considerations related to supporting Neighbor Discovery [RFC4861] on a logical interface.

- o Any Neighbor Discovery messages, such as Router Solicitation, Neighbor Solicitation Neighbor Advertisement messages that the host sends to a multicast destination address of link-local scope such as, all-nodes, all-routers, solicited-node multicast group addresses, using either an unspecified (::) source address, or a link-local address configured on the logical interface will be replicated and forwarded on each of the sub-interfaces under that logical interface. However, if the destination address is a unicast address and if that target is known to exist on a specific sub-interface, the packet will be forwarded only on that specific sub-interface and will not be replicated on all sub-interfaces.
- o Any Neighbor Discovery messages, such as Router Advertisement, that the host receives from any of its sub-interfaces part of the logical interface, will be associated with the logical interface, i.e., in some implementations the packet will appear on the input interface of the logical interface.
- o When using Stateless Address Autoconfiguraion [RFC4862] for generating IPv6 address configuration on the logical interface, the host may use any of the IPv6 prefixes received from the Router Advertisement messages that it received from any of its sub-interfaces.

- o The response to a Neighbor Discovery message received for a unicast, link-specific multicast group address, will be sent on the same sub-interface path where the packet was received.
- o When using DHCPv4 [RFC2131] for obtaining address configuration for the logical interface, the value in the chaddr field in the DHCP messages will be based on the link-layer identifier scheme chosen by the logical interface.

5.6. Provisioning Domain Considerations

The considerations related to the support of multiple provisioning domains in a multi-interface host is documented in [RFC6418]. These considerations specifically focus on the aspects related to DNS configuration. However, from the perspective of logical interface support, these considerations are not applicable, as the logical interface support is relevant only for a single provisioning domain. The key motivation for logical interface support is inter-technology handovers and the handovers are always in the context of a single provisioning domain.

5.7. Logical Interface Forwarding Conceptual Data Structures

The logical interface maintains the list of sub-interfaces that are part of the logical interface. This conceptual data structure is called as the LIF Table. The logical interface also maintains the list of flows associated with a given sub-interface and this conceptual data structure is called as the PIF Table. Both of these data structures have to be associated with a logical interface, and are depicted in Figure 4

LIF TABLE		FLOW table	
+=====+		+=====+	
PIF_ID	FLOW RoutingPolicies	FLOW ID	Physical_Intf_Id
	Home Network Prefix		-----+
	Link Layer Address	FLOW_ID	Physical_Intf_Id
	Status	+=====+	
+-----+			
PIF_ID	FLOW RoutingPolicies		
	Home Network Prefix		
	Link Layer Address		
	Status		
+-----+			
....		
+=====+			

Figure 4

The LIF table maintains the mapping between the LIF and each PIF associated to the LIF (refer to property #3, Figure 3). For each PIF entry the table should store the associated Routing Policies, the Home Network Prefix received in Router Advertisement, the configured link layer Address (as described above) and the Status of the PIF (e.g. active, not active). The method by which the Routing Policies are configured on the host is out of scope for this document.

The FLOW table allows the logical interface to properly route each IP flow over the right interface. The logical interface can identify the flows arriving on its sub-interfaces and associate them to those sub-interfaces. This approach is similar to reflective QoS performed by the IP routers. For locally generated traffic (e.g. unicast flows), the logical interface should perform interface selection based on the Flow Routing Policies. In case traffic of an existing flow is suddenly received from the network on a different sub-interface than the one locally stored, the logical interface should interpret the event as an explicit flow mobility trigger from the network and it should update the PIF_ID parameter in the FLOW table. Similarly, locally generated events from the sub-interfaces, or configuration updates to the local policy rules can cause updates to the table and hence trigger flow mobility.

6. Logical Interface Use-cases in Proxy Mobile IPv6

This section explains how the Logical interface support on the mobile node can be used for enabling some of the Proxy Mobile IPv6 protocol features.

6.1. Multihoming Support

A mobile node with multiple interfaces can attach simultaneously to the Proxy Mobile IPv6 domain. Each of the attachment links are assigned a unique set of IPv6 prefixes. If the host is configured to use Logical interface over the physical interface through which it is attached, following are the related considerations.

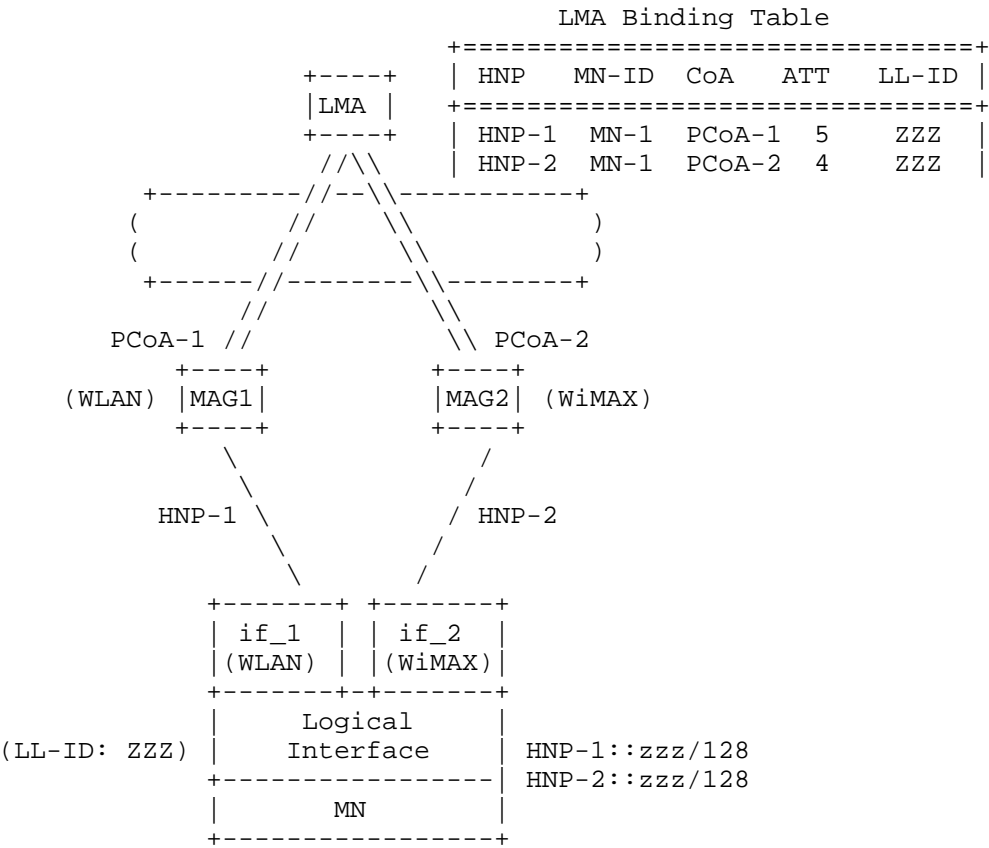


Figure 5: Multihoming Support

- o The mobile node detects the advertised prefixes from the MAG1 and MAG2 as the on link prefixes on the link to which the Logical interface is attached.
- o The mobile node can generate address configuration using stateless auto configuration mode from any of those prefixes.
- o The applications can be bound to any of the addresses bound to the Logical interface and that is determined based on the source address selection rules.
- o The host has path awareness for the hosted prefixes based on the received Router Advertisement messages. Any packets with source address generated using HNP_1 will be routed through the interface if_1 and for packets using source address from HNP_2 will be routed through the interface if_2.

6.2. Inter-Technology Handoff Support

The Proxy Mobile IPv6 protocol enables a mobile node with multiple network interfaces to move between access technologies, but still retaining the same address configuration on its attached interface. The protocol enables a mobile node to achieve address continuity during handoffs. If the host is configured to use Logical interface over the physical interface through which it is attached, following are the related considerations.

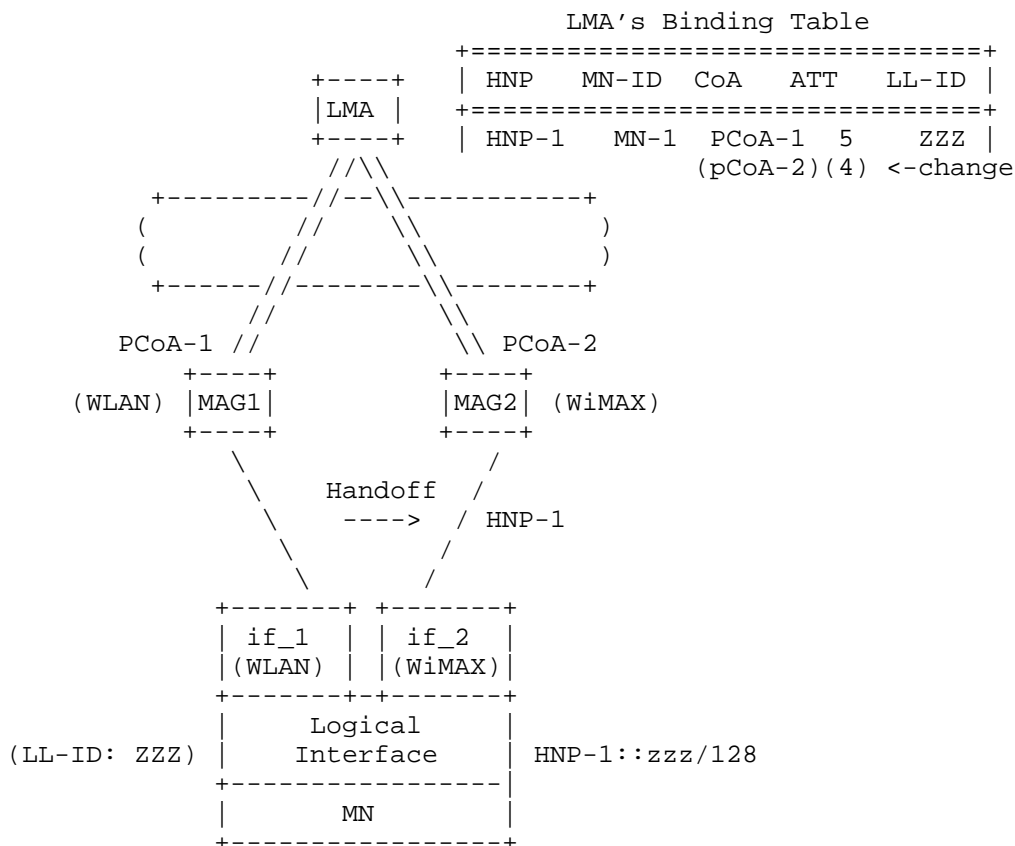


Figure 6: Inter-Technology Handoff Support

- o When the mobile node performs an handoff between if_1 and if_2, the change will not be visible to the applications of the mobile node. It will continue to receive Router Advertisements from the network, but from a different sub-interface path.
- o The protocol signaling between the network elements will ensure the local mobility anchor will switch the forwarding for the advertised prefix set from MAG1 to MAG2.
- o The MAG2 will host the prefix on the attached link and will include the home network prefixes in the Router Advertisements that it sends on the link.

6.3. Flow Mobility Support

For supporting flow mobility support, there is a need to support vertical handoff scenarios such as transferring a subset of prefix(es) (hence the flows associated to it/them) from one interface to another. The mobile node can support this scenario by using the Logical interface support. This scenario is similar to the Inter-technology handoff scenario defined in Section 6.2, only a subset of the prefixes are moved between interfaces.

Additionally, IP flow mobility in general initiates when the LMA decides to move a particular flow from its default path to a different one. The LMA can decide on which is the best MAG that should be used to forward a particular flow when the flow is initiated e.g. based on application policy profiles) and/or during the lifetime of the flow upon receiving a network-based or a mobile-based trigger.

As an example of mobile-based triggers, the LMA could receive input (e.g. by means of a layer 2.5 function via L3 signaling [RFC5677]) from the MN detecting changes in the mobile wireless environment (e.g. weak radio signal, new network detected, etc.). Upon receiving these triggers, the LMA can initiate the flow mobility procedures. For instance, when the mobile node only supports single-radio operation (i.e. one radio transmitting at a time), only sequential (i.e. not simultaneous) attachment to different MAGs over different media is possible. In this case layer 2.5 signaling can be used to perform the inter-access technology handover and communicate to the LMA the desired target access technology, MN-ID, Flow-ID and prefix.

7. IANA Considerations

This specification does not require any IANA Actions.

8. Security Considerations

This specification explains the operational details of Logical interface on an IP host. The Logical Interface implementation on the host is not visible to the network and does not require any special security considerations.

9. Authors

This document reflects contributions from the following authors (listed in alphabetical order):

Carlos Jesus Bernardos Cano

cjbc@it.uc3m.es

Antonio De la Oliva

aoliva@it.uc3m.es

Yong-Geun Hong

yonggeun.hong@gmail.com

Kent Leung

kleung@cisco.com

Tran Minh Trung

trungtm2909@gmail.com

Hidetoshi Yokota

yokota@kddilabs.jp

Juan Carlos Zuniga

JuanCarlos.Zuniga@InterDigital.com

10. Acknowledgements

The authors would like to acknowledge prior discussions on this topic in NETLMM and NETEXT working groups. The authors would also like to thank Joo-Sang Youn, Pierrick Seite, Rajeev Koodli, Basavaraj Patil, Peter McCann, and Julien Laganier for all the discussions on this topic.

11. References

11.1. Normative References

- [RFC4861] Narten, T., Nordmark, E., Simpson, W., and H. Soliman, "Neighbor Discovery for IP version 6 (IPv6)", RFC 4861, September 2007.
- [RFC4862] Thomson, S., Narten, T., and T. Jinmei, "IPv6 Stateless Address Autoconfiguration", RFC 4862, September 2007.
- [RFC5213] Gundavelli, S., Leung, K., Devarapalli, V., Chowdhury, K., and B. Patil, "Proxy Mobile IPv6", RFC 5213, August 2008.
- [RFC5844] Wakikawa, R. and S. Gundavelli, "IPv4 Support for Proxy Mobile IPv6", RFC 5844, May 2010.

11.2. Informative References

- [RFC2131] Droms, R., "Dynamic Host Configuration Protocol", RFC 2131, March 1997.
- [RFC4301] Kent, S. and K. Seo, "Security Architecture for the Internet Protocol", RFC 4301, December 2005.
- [RFC5072] S.Varada, Haskins, D., and E. Allen, "IP Version 6 over PPP", RFC 5072, September 2007.
- [RFC5677] Melia, T., Bajko, G., Das, S., Golmie, N., and JC. Zuniga, "IEEE 802.21 Mobility Services Framework Design (MSFD)", RFC 5677, December 2009.
- [RFC6085] Gundavelli, S., Townsley, M., Troan, O., and W. Dec, "Address Mapping of IPv6 Multicast Packets on Ethernet", RFC 6085, January 2011.
- [RFC6275] Perkins, C., Johnson, D., and J. Arkko, "Mobility Support in IPv6", RFC 6275, July 2011.
- [RFC6418] Blanchet, M. and P. Seite, "Multiple Interfaces and Provisioning Domains Problem Statement", RFC 6418, November 2011.
- [TS23401] "3rd Generation Partnership Project; Technical Specification Group Services and System Aspects; General Packet Radio Service (GPRS) enhancements for Evolved Universal Terrestrial Radio Access Network (E-UTRAN) access.", 2009.

Authors' Addresses

Telemaco Melia (editor)
Alcatel-Lucent
Route de Villejust
Nozay 91620
France

Email: telemaco.melia@alcatel-lucent.com

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

Email: sgundave@cisco.com

NETEXT WG
Internet-Draft
Intended status: Standards Track
Expires: June 21, 2014

X. Zhou
ZTE Corporation
J. Korhonen
Broadcom
C. Williams
Consultant
S. Gundavelli
Cisco
CJ. Bernardos
UC3M
December 18, 2013

Prefix Delegation Support for Proxy Mobile IPv6
draft-ietf-netext-pd-pmip-14

Abstract

This specification defines extensions to the Proxy Mobile IPv6 protocol for allowing a mobile router in a Proxy Mobile IPv6 domain to obtain IP prefixes for its attached mobile networks using DHCPv6 prefix delegation. Network-based mobility management support is provided for those delegated IP prefixes just as it is provided for the mobile node's home address. Even if the mobile router performs a handoff and changes its network point of attachment, mobility support is ensured for all the delegated IP prefixes and for all the IP nodes in the mobile network that use IP address configuration from those delegated IP prefixes.

Status of this Memo

This Internet-Draft is submitted in full conformance with the provisions of BCP 78 and BCP 79.

Internet-Drafts are working documents of the Internet Engineering Task Force (IETF). Note that other groups may also distribute working documents as Internet-Drafts. The list of current Internet-Drafts is at <http://datatracker.ietf.org/drafts/current/>.

Internet-Drafts are draft documents valid for a maximum of six months and may be updated, replaced, or obsoleted by other documents at any time. It is inappropriate to use Internet-Drafts as reference material or to cite them other than as "work in progress."

This Internet-Draft will expire on June 21, 2014.

Copyright Notice

Copyright (c) 2013 IETF Trust and the persons identified as the document authors. All rights reserved.

This document is subject to BCP 78 and the IETF Trust's Legal Provisions Relating to IETF Documents (<http://trustee.ietf.org/license-info>) in effect on the date of publication of this document. Please review these documents carefully, as they describe your rights and restrictions with respect to this document. Code Components extracted from this document must include Simplified BSD License text as described in Section 4.e of the Trust Legal Provisions and are provided without warranty as described in the Simplified BSD License.

Table of Contents

1. Introduction	4
2. Terminology	6
3. Solution Overview	7
3.1. Stated Assumptions	7
3.2. Deployment Models	8
3.2.1. Delegating Router co-located with Mobile Access Gateway	8
3.2.2. Delegating Router co-located with Local Mobility Anchor	9
3.2.3. Static Configuration of Delegated Mobile Network Prefixes	11
4. Message formats	12
4.1. Delegated Mobile Network Prefix Option	12
4.2. Status Codes	14
5. Operational Details	14
5.1. MAG Considerations	14
5.1.1. Extension to Binding Update List Entry Data Structure	14
5.1.2. Signaling Considerations	14
5.1.3. DHCP - MAG Interactions	16
5.1.3.1. Delegating Router co-located with Mobile Access Gateway	16
5.1.3.2. Delegating Router co-located with Local Mobility Anchor	18
5.1.4. Packet Forwarding	19
5.2. LMA Considerations	20
5.2.1. Extensions to Binding Cache Entry Data Structure	20
5.2.2. Signaling Considerations	20
5.2.3. Packet Forwarding	22
5.3. Security Policy Database (SPD) Example Entries	22
6. Security Considerations	23
7. IANA Considerations	24

8. Acknowledgments	24
9. References	25
9.1. Normative References	25
9.2. Informative References	25
Authors' Addresses	26

1. Introduction

Proxy Mobile IPv6 [RFC5213] enables network-based mobility management support for an IP host without requiring its participation in any IP mobility signaling. In Proxy Mobile IPv6 (PMIPv6), the mobile access gateway (MAG) performs the mobility management function on behalf of the mobile node (MN). The local mobility anchor (LMA) is the home agent for the MN and the topological anchor point. The mobility elements (LMA and MAGs) in the network allow an IP host to obtain an IPv4 address and/or a set of IPv6 addresses and be able to obtain IP mobility support for those IP address(es) within the Proxy Mobile IPv6 domain. In this context, the mobility management support is enabled for an individual IP host, which is the mobile node. The IPv4 home address, or the IPv6 home network prefixes are logically bound to the link shared between the mobile access gateway and the mobile node and only the mobile node can use those IP address(es) by configuring them on the interface attached to that link. Currently, there is no mobility support for the mobile networks attached to a mobile router in a Proxy Mobile IPv6 domain.

This specification defines extensions to the Proxy Mobile IPv6 protocol (a new mobility option for carrying delegated prefix information in proxy binding update and proxy binding acknowledgement messages) for allowing mobility support to the mobile networks attached to a mobile router. The mobile router can request the mobility entities in the Proxy Mobile IPv6 domain for one or more delegated IP prefixes using DHCP Prefix Delegation extensions [RFC3633], or through other means such as static configuration, or access technology specific mechanisms. The mobility entities in the PMIPv6 network provide network-based mobility management support for those delegated prefixes just as it is supported for a home address. The delegated prefixes are hosted in the mobile network attached to the mobile router. IP mobility is ensured for all the IP nodes in the mobile network, even as the mobile router performs a handoff by changing its point of network attachment within the Proxy Mobile IPv6 domain. The local mobility anchor in the Proxy Mobile IPv6 domain will not track the individual IP sessions for all the IP nodes in the mobile network, it only tracks a single mobile router session that is hosting the mobile network and associates the delegated IP prefixes with that session. Although the protocol solution defined in this specification also allows signaling IPv4 subnets between the mobile access gateway and the local mobility anchor, the delegation of IPv4 subnets to the mobile router is out of scope of this specification.




Figure 1: Mobile Router in Proxy Mobile IPv6 Domain

Within the context of this document, the definition of a mobile router extends that of a mobile node definition from [RFC5213], by adding routing capability between the mobile network and the point of attachment of the mobile router. The network of nodes part of the mobile network are referred to as locally fixed nodes (LFN) and they all move with the mobile router as a single cluster. As the mobile router moves, the LFNs are not aware of the mobility of the MR to a new point of attachment. Figure 1 illustrates a mobile router in a Proxy Mobile IPv6 domain.

The rest of the document identifies the protocol extensions and the operational details of the local mobility anchor and mobile access gateway for supporting this specification.

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].

All the mobility related terms used in this document are to be interpreted as defined in Proxy Mobile IPv6 specifications [RFC5213] and [RFC5844]. All the DHCP related terms are to be interpreted as defined in DHCPv6-PD for NEMO [RFC6276], DHCPv6-PD [RFC3633] and Subnet Allocation Option for DHCPv4 [RFC6656]. This document also provides a context-specific explanation to the following terms used in this document, and originally defined in the Mobile Network terminology document [RFC4885].

Mobile Router (MR)

The term mobile router is used to refer to an IP router whose mobility is managed by the network while being attached to a Proxy Mobile IPv6 domain. The mobile router is a mobile node as defined in [RFC5213], but with additional capabilities for supporting an attached mobile network. The MR's interface used for attachment to the mobile access gateway is referred to as the egress interface. Any MR's interface used for attachment to the mobile network is referred to as ingress interface. The mobility entities in the Proxy Mobile IPv6 domain provide mobility for the IPv4/IPv6 address(es) assigned to the mobile node's egress link and also mobility support to the network prefixes hosted in the network attached to the mobile router.

Mobile Network

It is an IP network attached to a mobile router. There can be many IP nodes in this IP network. The mobile router is a gateway for these IP nodes for reaching other IP networks or the Internet. The mobile router and the attached IP networks move as a single cluster.

Delegated Mobile Network Prefix (DMNP)

The Delegated Mobile Network Prefix is an IPv4/IPv6 prefix delegated to a mobile router and is hosted in the mobile network. The IP nodes in the mobile network will be able to obtain IP address configuration from the delegated mobile network prefix and will have IP mobility support for that address configuration. The DMNP is topologically anchored on the local mobility anchor and the mobility elements in the Proxy Mobile IPv6 domain provide IP mobility support for the prefix, by forwarding the mobile network

traffic to the mobile router.

Locally Fixed Node (LFN)

A Locally Fixed Node is an IP node in the mobile network. As the mobile router performs a handoff and changes its network point of attachment, the locally fixed node moves along with the mobile router.

3. Solution Overview

This section provides an overview of the operation of this specification, as well as lists the stated assumptions. This specification references three different deployment scenarios and explains the protocol operation.

3.1. Stated Assumptions

- o The mobile router is a mobile node as defined in [RFC5213], but with additional capabilities for routing IP packets between its egress interface (interface used for attachment to the mobile access gateway) and any of its ingress interfaces (interface used for attachment to the mobile network).
- o The specification assumes that a mobile router is an IPv4 and/or IPv6 router without any capability for mobility management.
- o The mobile router can obtain the delegated IP prefix(es) for its attached mobile networks using DHCPv6 Prefix Delegation, Static configuration, or through mechanisms specific to the access technology. This document assumes DHCPv6 Prefix Delegation [RFC3633] and in conjunction with the Prefix Exclude Option [RFC6603] as the default mechanism for prefix assignment to the mobile node. It defines an interworking between the mobility entities and the DHCPv6 functional elements in a non-normative way. The mechanism how to delegate IPv4 subnets to a mobile router is out of scope of this specification.
- o The mobile router obtains the IP address configuration for its egress roaming interface as specified in [RFC5213] and [RFC5844]. The mobile router along with its mobile networks will be able to perform handoff and change its point of attachment in the network and will be able to retain IP mobility support.
- o When using DHCPv6 Prefix Delegation, this document assumes that the mobile router uses its egress interface when making DHCPv6 requests.

3.2. Deployment Models

This section explains the protocol operation for supporting prefix delegation support in Proxy Mobile IPv6 for the following three deployment models: i) Delegating router co-located with mobile access gateway, ii) Delegating router co-located with local mobility anchor, and iii) Static configuration of delegated prefixes. High-level message call flows between the mobile router, mobile access gateway and the local mobility anchor are presented while explaining the protocol operation.

3.2.1. Delegating Router co-located with Mobile Access Gateway

In this deployment scenario, the delegating router (DR) function, as specified in [RFC3633], is co-located with the mobile access gateway, and a requesting router (RR) function is enabled on the mobile router.

Figure 2 shows the high-level message call flow for this case. The mobile router attaches to the mobile access gateway, which triggers the Proxy Mobile IPv6 signaling between the mobile access gateway and the local mobility anchor, setting up the bi-directional tunnel between them (regular Proxy Mobile IPv6 registration). After that, the DHCPv6 requesting router function running on the mobile router sends a Solicit message requesting a prefix. This message is received by the DHCPv6 delegating router function running on the mobile access gateway. The mobile access gateway then sends a proxy binding update message including a delegated mobile network prefix (DMNP) option carrying the ALL_ZERO value [RFC5213]. This serves as a request for the local mobility anchor to allocate a set of delegated prefixes, conveyed back in one or more DMNP options in a proxy binding acknowledgment message. The DHCPv6-PD signaling is then completed as described in [RFC3633], finalizing with the delegating router sending a Reply message conveying the delegated prefixes. If the requesting router includes a Rapid Commit option in its Solicit message, it is preferable that the MAG respond directly with a Reply rather than with an Advertise message, as described in [RFC3315], Section 17.2.3.

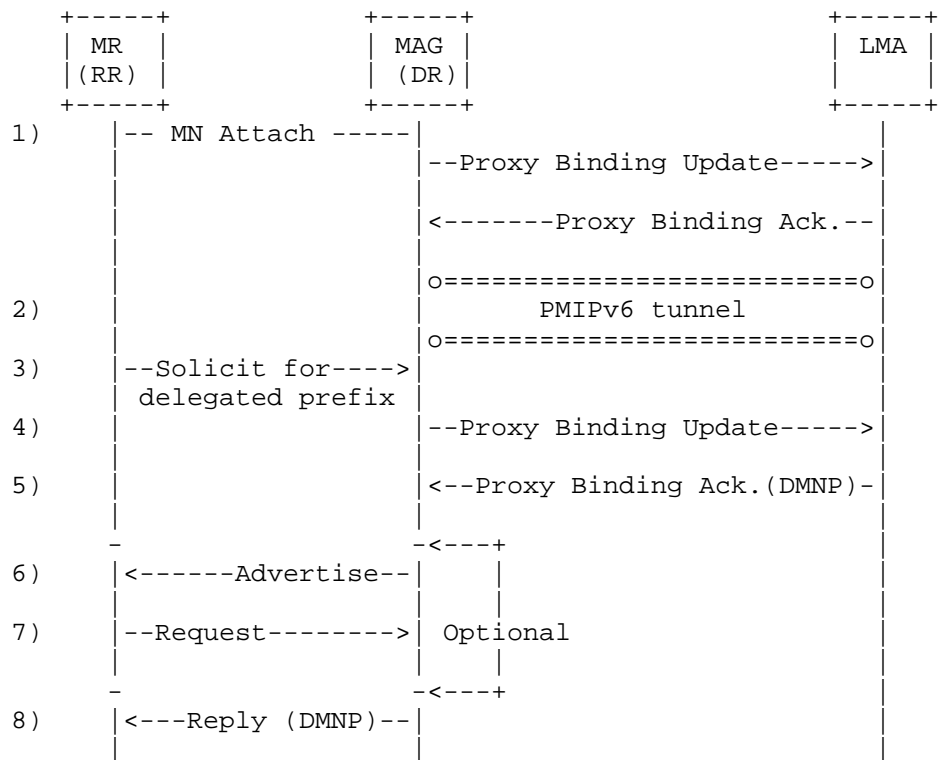


Figure 2: Delegating Router co-located with Mobile Access Gateway

From an operational point of view, this is the simplest deployment option, as it keeps a single protocol interface between the mobile access gateway and the local mobility anchor.

3.2.2. Delegating Router co-located with Local Mobility Anchor

In this deployment scenario, the delegating router (DR) function, as specified in [RFC3633], is co-located with the local mobility anchor, the requesting router (RR) function is enabled on the mobile router and a DHCPv6 Relay Agent (DRA) function, is co-located on the mobile access gateway.

Figure 3 shows the high-level message call flow for this case. The mobile router attaches to the mobile access gateway, which triggers the Proxy Mobile IPv6 signaling between the mobile access gateway and the local mobility anchor, setting up the bi-directional tunnel between them (regular Proxy Mobile IPv6 registration). After that, the DHCPv6 requesting router function running on the mobile router requests a prefix by sending a Solicit message. This message is

received by the DHCPv6 relay agent function running on the mobile access gateway, which then completes the DHCPv6 signaling, according to [RFC3315]. The relay agent function SHOULD include the relay agent remote-id option [RFC4649] into Relay-forward messages with appropriate identity information to enable correlation of mobile router identities used over DHCPv6 and PMIPv6.

Once the mobile access gateway gets the set of delegated prefixes from the delegating router function running on the local mobility anchor, the MAG conveys the delegated prefixes in a proxy binding update. This ensures that the local mobility anchor properly routes the traffic addressed to the delegated prefixes via the PMIPv6 tunnel established with the mobile access gateway, and that mobility is provided to these prefixes while the mobile router roams within the PMIPv6 domain. Note that the relay agent function in the mobile access gateway has to queue the Reply message for the duration of the PMIPv6 signaling (steps 10 and 11) before forwarding the Reply message to the requesting router. While this does not change anything from the DHCPv6-PD protocol point of view, implementations will need to account for interactions between the timing of PMIPv6 signaling and the DHCPv6 timeout/retry logic.

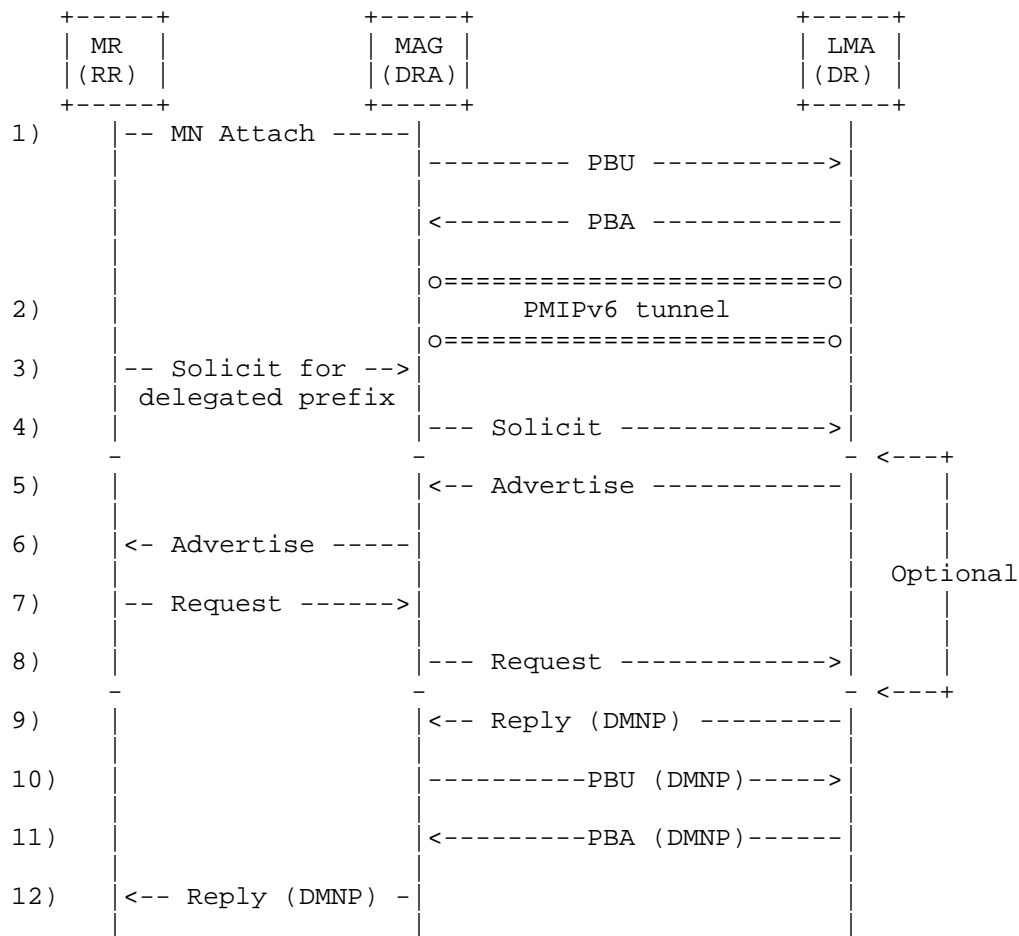


Figure 3: Delegating Router co-located with Local Mobility Anchor

The DR function can also be on the located in other entities of the home network different from the LMA. This deployment model requires some interworking between the DR and the LMA and is out of scope for this specification. Note that this additional interworking would have no impact on the protocol between the LMA and MAG defined in this document.

3.2.3. Static Configuration of Delegated Mobile Network Prefixes

In this deployment scenario, the delegated mobile network prefixes of the mobile router are statically configured in the mobile node's policy profile [RFC5213]. The delegated mobile network prefixes are statically configured in the mobile network attached to the mobile

router. The mobile router is the default-router for the mobile networks.

Figure 4 shows a high-level message call flow for this example. The mobile access gateway obtains statically configured mobile network prefixes from the policy profile and registers them with the local mobility anchor using the extensions specified in this document, that is, the use of the delegated mobile network prefix (DMNP) option in the Proxy Mobile IPv6 signaling. There is no explicit trigger from the mobile router for registering, or de-registering those prefixes. As long as there is a mobility session for the mobile router's home address, the local mobility anchor enables mobility support for the mobile network prefixes.




Figure 4: Static Configuration of Delegated Mobile Network Prefixes


4. Message formats

This section defines extensions to Proxy Mobile IPv6 [RFC5213] protocol messages.

4.1. Delegated Mobile Network Prefix Option

A new mobility header option, Delegated Mobile Network Prefix option is defined for use with Proxy Binding Update and Proxy Binding Acknowledgment messages exchanged between a local mobility anchor and a mobile access gateway. This option is used for exchanging the mobile router's IPv4/IPv6 delegated mobile network prefix. There can be multiple instances of the Delegated Mobile Network Prefix option present in a message.

The Delegated Mobile Network Prefix option has an alignment requirement of $8n+2$. Its format is as follows:



Type

<IANA-1>: To be assigned by IANA.

Length

8-bit unsigned integer indicating the length of the option in octets, excluding the type and length fields.

IPv4 Prefix (V)

If the IPv4 Prefix (V) flag is set to a value of (1), then it indicates that the prefix that is included in the DMNP field is an IPv4 prefix. If the IPv4 Prefix (V) flag is set to a value of (0), then it indicates that the prefix that is included in the DMNP field is an IPv6 prefix.

Reserved

This field is unused for now. The value MUST be initialized to 0 by the sender and MUST be ignored by the receiver.

Prefix Length

8-bit unsigned integer indicating the prefix length of the prefix contained in the option.

Delegated Mobile Network Prefix

Contains a mobile router's 4-byte IPv4 or a 16-byte IPv6 Delegated Mobile Network Prefix.

4.2. Status Codes

This document defines the following new status code values for use in the Proxy Binding Acknowledgement message. These values have been allocated from the same number space as defined in Section 6.1.8 of [RFC6275].

NOT_AUTHORIZED_FOR_DELEGATED_MNP: <IANA-2>

Not Authorized for delegated mobile network prefix

REQUESTED_DMNP_IN_USE: <IANA-3>

Requested delegated mobile network prefix is in use

5. Operational Details

5.1. MAG Considerations

5.1.1. Extension to Binding Update List Entry Data Structure

In order to support this specification, the conceptual Binding Update List Entry (BULE) data structure [RFC5213] needs to be extended to include a delegated mobile network prefix (DMNP) list. Each entry in the list is used for storing an IPv4/IPv6 mobile network prefix delegated to the mobile router.

5.1.2. Signaling Considerations

During the mobile router's initial attachment procedure, the mobile access gateway obtains the mobile router's policy profile, as per the procedures defined in [RFC5213]. The mobile node's policy profile defined in [RFC5213] is extended to include a parameter which indicates Delegated Prefix support. If the policy profile indicates that the mobile router is authorized for Delegated Prefix support, then the considerations described next apply.

The mobile access gateway MUST include one or more Delegated Mobile Network Prefix (DMNP) options in the Proxy Binding Update message in order to request the local mobility anchor to allocate delegated mobile network prefix(es) for the mobile router.

If the mobile access gateway requests the local mobility anchor to perform the prefix assignment, then:

- o There MUST be exactly one instance of the Delegated Mobile Network Prefix option with ALL_ZERO value and with the (V) flag set to a value of (0). This serves as a request to the local mobility anchor to allocate a set of delegated IPv6 mobile network prefixes.
- o There MUST be exactly one instance of the Delegated Mobile Network Prefix option with ALL_ZERO value and with the (V) flag set to a value of (1). This serves as a request to the local mobility anchor to allocate a set of delegated IPv4 mobile network prefixes.
- o If the received Proxy Binding Acknowledgement message has the status field value set to NOT_AUTHORIZED_FOR_DELEGATED_MNP (Not Authorized for delegated mobile network prefix), the mobile access gateway MUST NOT enable mobility support for any of the prefixes in the mobile network and prefix delegation support has to be disabled.
- o If the received Proxy Binding Acknowledgement message has the status field value set to REQUESTED_DMNP_IN_USE (Requested delegated mobile network prefix is in use), the mobile access gateway MUST NOT enable mobility support for the requested prefixes. The mobile access gateway MAY choose to send Proxy Binding Update message requesting the local mobility anchor to perform the prefix assignment.

If the mobile access gateway provides the local mobility anchor with the prefix(es) that wants to get allocated, then:

- o There MUST be exactly one instance of the Delegated Mobile Network Prefix option with NON_ZERO prefix value [RFC5213] for each of the mobile network prefixes that the mobile access gateway is requesting the local mobility anchor to allocate. The prefix value in the option is the prefix that is either statically configured for that mobile router in the mobile node's policy profile, or obtained via interactions with the DHCP PD functions. This serves as a request to the local mobility anchor to allocate the requested IPv4/IPv6 prefix.

If the received Proxy Binding Acknowledgement message has the status field value set to 0 (Proxy Binding Update accepted), the mobile access gateway has to apply the following considerations.

- o The delegated mobile network prefix (DMNP) list in the mobile router's Binding Update List entry has to be updated with the allocated prefix(es). However, if the received message was in response to a de-registration request with a lifetime value of

(0), then the delegated mobile network prefix list has to be removed along with the Binding Update List entry.

- o The mobile access gateway has to set up a policy-based route for forwarding the IP packets received from the mobile network (with the source IP address from any of the delegated IPv4/IPv6 mobile network prefixes) through the bidirectional tunnel set up for that mobile router. However, if the received message was in response to a de-registration request with a lifetime value of (0), then the created forwarding state has to be removed.

This specification assumes that all the mobile access gateways of a PMIPv6 Domain support the same prefix delegation mechanism. If there is any difference, it will result in delegated mobile network prefix(es) getting de-registered and the mobile network losing the prefix(es). This would result in the attached local fixed nodes losing the assigned IP addresses. The mobile router MAY explicitly deprecate these prefixes. Alternatively the lifetime of the addresses may expire.

5.1.3. DHCP - MAG Interactions

This section describes the interactions between the DHCP and PMIPv6 logical entities running on the mobile access gateway. This section is applicable only for deployments that use DHCPv6-based prefix delegation (i.e., it does not apply if static configuration is used). As described next, these interactions vary slightly depending on the considered deployment model at the mobile access gateway (described in Section 3.2).

The mobile router, acting as a "Requesting Router" as described in [RFC3633], sends a Solicit message including one or more IA_PD option(s) to the Delegating Router/DHCPv6 Relay Agent collocated on the mobile access gateway. This message provides the needed trigger for the mobile access gateway for requesting the local mobility anchor to enable delegated mobile network prefix support for that mobility session. We next describe the subsequent interactions depending on the deployment model.

5.1.3.1. Delegating Router co-located with Mobile Access Gateway

The mobile access gateway applies the considerations in Section 5.1.2 for requesting the local mobility anchor to enable delegated prefix support. For example, if the mobile router is soliciting an IPv4 prefix, the mobile access gateway includes in the Proxy Binding Update signaling a Delegated Mobile Network Prefix option with ALL_ZERO value and with the (V) flag set to a value of (1).

The mobile access gateway, upon successfully completing the Proxy Binding Update signaling with the local mobility anchor (following the considerations described in Section 5.1.2), adds the delegated mobile network prefixes to the binding update list. Then, the mobile access gateway provides the obtained prefixes to the DHCPv6 Delegating Router for prefix assignment. The way in which these prefixes are passed to the DHCPv6 delegating router function is beyond the scope of this document.

- o In case the Proxy Binding Update signaling with the local mobility anchor is not completed successfully, for example because the local mobility anchor is not authorized for delegated mobile network prefix or the requested prefix is in use, the DHCPv6 Delegating Router will send a Reply message to the Requesting Router with no IA_PREFIX suboptions and with a Status Code option as described in [RFC3633], section 11.2.

The standard DHCPv6 considerations will be applied with respect to the interactions between the Delegating Router and the Requesting Router. The Requesting Router is provided with the delegated prefix(es), which can then be then advertised in the mobile network, and therefore used by the locally fixed nodes to auto configure IP addresses allowing to gain access to the Internet.

Any time, the Requesting Router releases the delegated prefixes, the Delegating Router removes the assigned prefixes. To do so, the mobile access gateway will send an Updated Proxy Binding Update following the considerations described in Section 5.1.2 for de-registering those prefixes. The way in which the DHCPv6 Delegating Router triggers the mobile access gateway in order to de-register the prefixes is beyond the scope of this document.

In case the mobile router performs a handover and attaches to a different mobile access gateway, the following cases are possible:

- o The new mobile access gateway does not support the delegation of mobile network prefixes described in this specification. In this case, forwarding of the previously delegated mobile network prefixes is no longer performed.
- o The new mobile access gateway supports the delegation of mobile network prefixes described in this specification. There are two possible cases upon the reception of the SOLICIT message by the Delegating Router. If the MAG already knows the delegated mobile network prefixes, it conveys them in a DMNP option included in the Proxy Binding Update sent to the local mobility anchor, which then authorizes them based on: a) the content of the associated binding cache entry (if exists), b) the user profile (if the allocation is

static), or, c) checking that the delegated mobile network prefixes are not already allocated. On the other hand, if the mobile access gateway is not aware of the delegated mobile network prefixes, it will include 0.0.0.0 / ::0 in a DMNP option included in the Proxy Binding Update sent to the LMA, which will provide the right prefixes back in the Proxy Binding Acknowledgement based on a) the content of the associated binding cache entry (if exists), b) the profile (if static allocation is used), or c) dynamic assignment.

5.1.3.2. Delegating Router co-located with Local Mobility Anchor

A DHCPv6 Relay Agent function running on the mobile access gateway will forward the DHCP messages to the local mobility anchor which has the co-located Delegating Router function. The Requesting Router and the Delegating Router complete the DHCP messages related to prefix delegation.

During the DHCPv6 exchange, the standard DHCPv6 considerations apply with respect to the interactions between the Delegating Router, DHCPv6 Relay Agent and the Requesting Router.

The mobile access gateway learns from the co-located DHCPv6 Relay Agent the prefixes allocated by the Delegating Router. The way in which the mobile access gateway learns obtains this information from the DHCPv6 Relay Agent function is beyond the scope of this document.

The mobile access gateway will apply the considerations in Section 5.1.2 for requesting the local mobility anchor to enable delegated prefix support. The mobile access gateway will include exactly one instance of the Delegated Mobile Network Prefix option with NON_ZERO prefix value for each of the mobile network prefixes that the mobile access gateway is requesting the local mobility anchor to allocate. The prefix value(s) in the option will be the prefix(es) obtained via DHCP prefix delegation.

The mobile access gateway, upon successfully completing the Proxy Binding Update signaling with the local mobility anchor, will provide the obtained prefixes to the DHCPv6 Relay Agent for prefix assignment. The Delegating Router is provided with the delegated prefix(es) completing the standard DHCPv6 signaling. These prefixes can then be then advertised in the mobile network, and therefore used by the locally fixed nodes to auto configure IP addresses allowing to gain access to the Internet.

- o In case the Proxy Binding Update signaling with the local mobility anchor is not completed successfully, for example because the local mobility anchor is not authorized for delegated mobile

network prefix, the requested prefix is in use, or the delegated prefix(es) do not match the ones allocated by DHCP prefix delegation, the DHCPv6 Relay Agent MAY send a Reply message to the Requesting Router with no IA_PREFIX suboptions and with a Status Code option as described in [RFC3633], section 11.2.

In case the mobile router performs a handover and attaches to a different mobile access gateway, the following cases are possible:

- o The new mobile access gateway does not support the delegation of mobile network prefixes described in this specification. In this case, forwarding of the previously delegated mobile network prefixes is no longer performed.
- o The new mobile access gateway supports the delegation of mobile network prefixes described in this specification. There are two possible cases upon the reception of the SOLICIT message by the DHCPv6 Relay Agent. If the MAG already knows the delegated mobile network prefixes, it conveys them in a DMNP option included in the Proxy Binding Update sent to the local mobility anchor, which then authorizes them based on: a) the content of the associated binding cache entry (if exists), b) the user profile (if the allocation is static), or, c) checking that the delegated mobile network prefixes are not already allocated. On the other hand, if the mobile access gateway is not aware of the delegated mobile network prefixes, it will include 0.0.0.0 / ::0 in a DMNP option included in the Proxy Binding Update sent to the LMA, which will provide the right prefixes back in the Proxy Binding Acknowledgement based on a) the content of the associated binding cache entry (if exists), b) the profile (if static allocation is used), or c) dynamic assignment.

5.1.4. Packet Forwarding

On receiving an IP packet from a mobile router, the mobile access gateway before tunneling the packet to the local mobility anchor MUST ensure that there is an established binding for the mobile router and the source IP address of the packet is a prefix delegated to that mobile router. If the source address of the received IP packet is not part of the delegated mobile network prefix, then the mobile access gateway MUST NOT tunnel the packet to the local mobility anchor.

On receiving an IP packet from the bi-directional tunnel established with the local mobility anchor, the mobile access gateway MUST first decapsulate the packet (removing the outer header) and then use the destination address of the (inner) packet to forward it on the interface through which the mobile router is reachable.

The above forwarding considerations are not applicable to the IP traffic sent/received to/from the mobile router's home address (IPv4 HOA/HNP). For the mobile router's home address traffic, forwarding considerations from [RFC5213] and [RFC5844] continue to apply.

5.2. LMA Considerations

5.2.1. Extensions to Binding Cache Entry Data Structure

In order to support this specification, the conceptual Binding Cache Entry (BCE) data structure [RFC5213] needs to be extended to include the delegated mobile network prefix (DMNP) list. Each entry in the list represents a delegated mobile network prefix.

5.2.2. Signaling Considerations

If the Proxy Binding Update message does not include any Delegated Mobile Network Prefix option(s) (Section 4.1), then the local mobility anchor MUST NOT enable Delegated Prefix support for the mobility session, and the Proxy Binding Acknowledgment message that is sent in response MUST NOT contain any Delegated Mobile Network Prefix option(s).

If the Proxy Binding Update message includes one or more Delegated Mobile Network Prefix options, but the local mobility anchor is not configured with Delegated Prefix support, then the local mobility anchor will ignore the option(s) and process the rest of the option as specified in [RFC5213]. This would have no effect on the operation of the rest of the protocol. The Proxy Binding Acknowledgment message that is sent in response will not include any Delegated Mobile Network Prefix option(s).

If the Proxy Binding Update message has the Delegated Mobile Network Prefix option(s) and if the local mobility anchor is configured for Delegated Prefix support, then the local mobility anchor MUST enable Delegated Mobile Network Prefix option for that mobility session. The Proxy Binding Acknowledgment message that is sent in response MUST include the Delegated Mobile Network Prefix option(s). The following considerations apply.

- o If there is at least one instance of the Delegated Mobile Network Prefix option with a ALL_ZERO [RFC5213] prefix value, then this serves as a request for the local mobility anchor to perform the assignment of one or more delegated mobile network prefixes.
- * A Delegated Mobile Network option with ALL_ZERO value and with the (V) flag set to a value of (0), is a request for the local mobility anchor to allocate one or more IPv6 prefixes.

- * A Delegated Mobile Network option with ALL_ZERO value and with the (V) flag set to a value of (1), is a request for the local mobility anchor to allocate one or more IPv4 prefixes.
- * Inclusion of multiple instances of Delegated Mobile Network options with ALL_ZERO value, one with the (V) flag set to a value of (1), and another instance with the (V) flag set to a value of (0) is a request to allocate both IPv4 and IPv6 prefixes.
- o If there are no instances of the Delegated Mobile Network Prefix option present in the request with ALL_ZERO value, but has a specific prefix value, then this serves as a request for the local mobility anchor to perform the allocation of the requested prefix(es).
- * If any one of the requested prefixes are assigned to some other mobility node, or not from an authorized pool that the local mobility can allocate for that mobility session, then the Proxy Binding Update MUST be rejected by sending a Proxy Binding Acknowledgement message with Status field set to REQUESTED_DMNP_IN_USE (Requested delegated mobile network prefix is in use).

Upon accepting the Proxy Binding Update, the local mobility anchor MUST send a Proxy Binding Acknowledgement message with the Status field set to 0 (Proxy Binding Update accepted).

- o The message MUST include one instance of the Delegated Mobile Network Prefix option for each of the allocated IPv4/IPv6 delegated mobile network prefixes.
- o The delegated mobile network prefix (DMNP) list in the mobile router's Binding Cache entry has to be updated with the allocated prefix(es). However, if the request is a de-registration request with a lifetime value of (0), the delegated mobile network prefix list has to be removed along with the Binding Cache entry.
- o A route (or a platform-specific equivalent function that sets up the forwarding) for each of the allocated prefixes over the tunnel has to be added. However, if the request is a de-registration request, with a lifetime value of (0), all the IPv4/IPv6 delegated prefix routes created for that session have to be removed.

5.2.3. Packet Forwarding

The local mobility anchor MUST advertise a connected route into the routing infrastructure for the IP prefixes delegated to all of the mobile routers that it is serving. This step essentially enables the local mobility anchor to be a routing anchor for those IP prefixes and be able to intercept IP packets sent to those mobile networks.

On receiving a packet from a correspondent node with the destination address matching any of the mobile router's delegated mobile network prefixes, the local mobility anchor MUST forward the packet through the bi-directional tunnel set up with the mobile access gateway where the mobile router is attached.

On receiving an IP packet from the bi-directional tunnel established with the mobile access gateway, the local mobility anchor MUST first decapsulate the packet (removing the outer header) and then use the destination address of the (inner) packet for forwarding decision. The local mobility anchor MUST ensure that there is an established binding for the mobile router and the source IP address of the packet is a prefix delegated to a mobile router reachable over that bi-directional tunnel.

The above forwarding considerations are not applicable to the IP traffic sent/received to/from the mobile router's home address (IPv4 HOA/HNP). For the mobile router's home address traffic, forwarding considerations from [RFC5213] and [RFC5844] continue to apply.

5.3. Security Policy Database (SPD) Example Entries

The use of DHCPv6, as described in this document, requires message integrity protection and source authentication. The IPsec security mechanism used by Proxy Mobile IPv6 [RFC5213] for securing the signaling messages between the mobile access gateway and the local mobility anchor can be used for securing the DHCP signaling between the mobile access gateway and the local mobility anchor.

The Security Policy Database (SPD) and Security Association Database (SAD) entries necessary to protect the DHCP signaling is specified below. The format of these entries is based on [RFC4877] conventions. The SPD and SAD entries are only example configurations. A particular implementation of mobile access gateway and local mobility anchor implementation can configure different SPD and SAD entries as long as they provide the required security for protecting DHCP signaling messages.

For the examples described in this document, a mobile access gateway with address "mag_address_1", and a local mobility anchor with

address "lma_address_1" are assumed.

mobile access gateway SPD-S:

- IF local_address = mag_address_1 &
remote_address = lma_address_1 & proto = UDP &
local_port = any & remote_port = DHCP
Then use SA1 (OUT) and SA2 (IN)

mobile access gateway SAD:

- SA1(OUT, spi_a, lma_address_1, ESP, TRANSPORT):
local_address = mag_address_1 &
remote_address = lma_address_1 &
proto = UDP & remote_port = DHCP
- SA2(IN, spi_b, mag_address_1, ESP, TRANSPORT):
local_address = lma_address_1 &
remote_address = mag_address_1 &
proto = UDP & local_port = DHCP

local mobility anchor SPD-S:

- IF local_address = lma_address_1 &
remote_address = mag_address_1 & proto = UDP &
local_port = DHCP & remote_port = any
Then use SA2 (OUT) and SA1 (IN)

local mobility anchor SAD:

- SA2(OUT, spi_b, mag_address_1, ESP, TRANSPORT):
local_address = lma_address_1 &
remote_address = mag_address_1 &
proto = UDP & local_port = DHCP
- SA1(IN, spi_a, lma_address_1, ESP, TRANSPORT):
local_address = mag_address_1 &
remote_address = lma_address_1 &
proto = UDP & remote_port = DHCP

6. Security Considerations

The Delegated Mobile Network Prefix Option defined in this specification is for use in Proxy Binding Update and Proxy Binding Acknowledgement messages. This option is carried like any other mobility header option as specified in [RFC5213]. Therefore, it inherits from [RFC5213] its security guidelines and does not require any additional security considerations.

The use of DHCPv6 in this specification is as defined in DHCPv6 base specification [RFC3315] and DHCPv6 Prefix Delegation specifications [RFC3633]. The security considerations specified in those specifications apply to this document.

If IPsec is used, the IPsec security association that is used for protecting the Proxy Binding Update and Proxy Binding Acknowledgement, also needs to be used for protecting the DHCPv6 signaling between the mobile access gateway and the local mobility anchor. Considerations specified in Section 5.3 identify the extensions to security policy entries [RFC4301]

7. IANA Considerations

This document requires the following IANA actions.

- o Action-1: This specification defines a new Mobility Header option, Delegated Mobile Network Prefix option. This mobility option is described in Section 4.1. The type value <IANA-1> for this message needs to be allocated from the Mobility Options registry at <http://www.iana.org/assignments/mobility-parameters>. RFC Editor: Please replace <IANA-1> in Section 4.1 with the assigned value, and update this section accordingly.
- o Action-2: This document also defines two new status code values for use in the Proxy Binding Acknowledgement message, as described in Section 4.2. These status codes are, NOT_AUTHORIZED_FOR_DELEGATED_MNP (Not Authorized for delegated mobile network prefix) with a status code value of <IANA-2>, and REQUESTED_DMNP_IN_USE (Requested delegated mobile network prefix is in use) with a status code value of <IANA-3>. These values have to be assigned from the same number space as allocated for other status codes [RFC6275] and update this section accordingly.

8. Acknowledgments

The authors would like to acknowledge Ryuji Wakikawa, Alexandru Petrescu, Behcet Sarikaya, Seil Jeon, Basavaraj Patil, Brian Haberman and Michal Hoft for all the discussions and reviews of this draft.

The work of Carlos J. Bernardos has also been partially supported by the European Community's Seventh Framework Programme (FP7-ICT-2009-5) under grant agreement n. 258053 (MEDIEVAL project) and by the Ministry of Science and Innovation of Spain under the QUARTET project (TIN2009-13992-C02-01).

9. References

9.1. Normative References

- [RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, March 1997.
- [RFC3315] Droms, R., Bound, J., Volz, B., Lemon, T., Perkins, C., and M. Carney, "Dynamic Host Configuration Protocol for IPv6 (DHCPv6)", RFC 3315, July 2003.
- [RFC3633] Troan, O. and R. Droms, "IPv6 Prefix Options for Dynamic Host Configuration Protocol (DHCP) version 6", RFC 3633, December 2003.
- [RFC4301] Kent, S. and K. Seo, "Security Architecture for the Internet Protocol", RFC 4301, December 2005.
- [RFC4649] Volz, B., "Dynamic Host Configuration Protocol for IPv6 (DHCPv6) Relay Agent Remote-ID Option", RFC 4649, August 2006.
- [RFC4877] Devarapalli, V. and F. Dupont, "Mobile IPv6 Operation with IKEv2 and the Revised IPsec Architecture", RFC 4877, April 2007.
- [RFC5213] Gundavelli, S., Leung, K., Devarapalli, V., Chowdhury, K., and B. Patil, "Proxy Mobile IPv6", RFC 5213, August 2008.
- [RFC5844] Wakikawa, R. and S. Gundavelli, "IPv4 Support for Proxy Mobile IPv6", RFC 5844, May 2010.
- [RFC6275] Perkins, C., Johnson, D., and J. Arkko, "Mobility Support in IPv6", RFC 6275, July 2011.
- [RFC6276] Droms, R., Thubert, P., Dupont, F., Haddad, W., and C. Bernardos, "DHCPv6 Prefix Delegation for Network Mobility (NEMO)", RFC 6276, July 2011.
- [RFC6603] Korhonen, J., Savolainen, T., Krishnan, S., and O. Troan, "Prefix Exclude Option for DHCPv6-based Prefix Delegation", RFC 6603, May 2012.

9.2. Informative References

- [RFC4885] Ernst, T. and H-Y. Lach, "Network Mobility Support Terminology", RFC 4885, July 2007.
- [RFC6656] Johnson, R., Kinnear, K., and M. Stapp, "Description of Cisco Systems' Subnet Allocation Option for DHCPv4",

RFC 6656, July 2012.

Authors' Addresses

Xingyue Zhou
ZTE Corporation
No.50 Software Avenue, Yuhuatai District
Nanjing
China

Phone: +86-25-8801-4634
Email: zhou.xingyue@zte.com.cn

Jouni Korhonen
Broadcom
Porkkalankatu 24
Helsinki FIN-00180
Finland

Email: jouni.nospam@gmail.com

Carl Williams
Consultant
San Jose, CA
USA

Email: carlw@mcsr-labs.org

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

Email: sgundave@cisco.com

Carlos J. Bernardos
Universidad Carlos III de Madrid
Av. Universidad, 30
Leganes, Madrid 28911
Spain

Phone: +34 91624 6236
Email: cjbc@it.uc3m.es
URI: <http://www.it.uc3m.es/cjbc/>

NETEXT WG
Internet-Draft
Intended status: Standards Track
Expires: September 29, 2014

M. Liebsch
NEC
P. Seite
Orange
H. Yokota
KDDI Lab
J. Korhonen
Broadcom Communications
S. Gundavelli
Cisco
March 28, 2014

Quality of Service Option for Proxy Mobile IPv6
draft-ietf-netext-pmip6-qos-12.txt

Abstract

This specification defines a new mobility option, the Quality of Service (QoS) option, for Proxy Mobile IPv6. This option can be used by the local mobility anchor and the mobile access gateway for negotiating Quality of Service parameters for a mobile node's IP flows. The negotiated QoS parameters can be used for QoS policing and marking of packets to enforce QoS differentiation on the path between the local mobility anchor and the mobile access gateway. Furthermore, making QoS parameters available on the mobile access gateway enables mapping of these parameters to QoS rules that are specific to the access technology and allows those rules to be enforced on the access network using access technology specific approaches.

Status of this Memo

This Internet-Draft is submitted in full conformance with the provisions of BCP 78 and BCP 79.

Internet-Drafts are working documents of the Internet Engineering Task Force (IETF). Note that other groups may also distribute working documents as Internet-Drafts. The list of current Internet-Drafts is at <http://datatracker.ietf.org/drafts/current/>.

Internet-Drafts are draft documents valid for a maximum of six months and may be updated, replaced, or obsoleted by other documents at any time. It is inappropriate to use Internet-Drafts as reference material or to cite them other than as "work in progress."

This Internet-Draft will expire on September 29, 2014.

Copyright Notice

Copyright (c) 2014 IETF Trust and the persons identified as the document authors. All rights reserved.

This document is subject to BCP 78 and the IETF Trust's Legal Provisions Relating to IETF Documents (<http://trustee.ietf.org/license-info>) in effect on the date of publication of this document. Please review these documents carefully, as they describe your rights and restrictions with respect to this document. Code Components extracted from this document must include Simplified BSD License text as described in Section 4.e of the Trust Legal Provisions and are provided without warranty as described in the Simplified BSD License.

Table of Contents

1. Introduction	4
2. Conventions and Terminology	6
2.1. Conventions	6
2.2. Terminology	6
3. Overview of QoS Support in Proxy Mobile IPv6	9
3.1. Quality of Service Option - Usage Examples	11
3.2. Quality of Service Attributes - Usage Examples	13
4. Protocol Messaging Extensions	15
4.1. Quality of Service Option	15
4.2. Quality of Service Attribute	17
4.2.1. Per Mobile Node Aggregate Maximum Downlink Bit Rate	19
4.2.2. Per Mobile Node Aggregate Maximum Uplink Bit Rate	20
4.2.3. Per Mobility Session Aggregate Maximum Downlink Bit Rate	21
4.2.4. Per Mobility Session Aggregate Maximum Uplink Bit Rate	23
4.2.5. Allocation and Retention Priority	25
4.2.6. Aggregate Maximum Downlink Bit Rate	27
4.2.7. Aggregate Maximum Uplink Bit Rate	28
4.2.8. Guaranteed Downlink Bit Rate	29
4.2.9. Guaranteed Uplink Bit Rate	30
4.2.10. QoS Traffic Selector	31
4.2.11. QoS Vendor Specific Attribute	32
4.3. New Status Code for Proxy Binding Acknowledgement	33
4.4. New Notification Reason for Update Notification Message	33
4.5. New Status Code for Update Notification Acknowledgement Message	33

5.	Protocol Considerations	35
5.1.	Local Mobility Anchor Considerations	35
5.2.	Mobile Access Gateway Considerations	38
6.	QoS Services in Integrated WLAN-3GPP Networks	43
6.1.	Technical Scope and Procedure	43
6.2.	Relevant QoS Attributes	45
7.	IANA Considerations	47
8.	Implementation Status	50
9.	Security Considerations	52
10.	Acknowledgements	53
11.	References	54
11.1.	Normative References	54
11.2.	Informative References	54
Appendix A.	Information when implementing 3GPP QoS in IP transport network	56
A.1.	Mapping tables	56
A.2.	Use cases and protocol operations	57
A.2.1.	Handover of existing QoS rules	57
A.2.2.	Establishment of QoS rules	59
A.2.3.	Dynamic Update to QoS Policy	61
Appendix B.	Information when implementing PMIP based QoS support with IEEE 802.11e	63
Appendix C.	Information when implementing with a Broadband Network Gateway	67
Authors' Addresses	68

1. Introduction

Mobile operators deploy Proxy Mobile IPv6 (PMIPv6) [RFC5213] to enable network-based mobility management for mobile nodes (MN). Users can access Internet Protocol (IP) based services from their mobile device by using various radio access technologies. The currently supported mobile standards have adequate support for QoS-based service differentiation for subscriber traffic in cellular radio access networks. QoS policies are typically controlled by a policy control function, whereas the policies are enforced by one or more gateways in the infrastructure, such as the local mobility anchor and the mobile access gateway, as well as by access network elements. Policy control and in-band QoS differentiation for access to the mobile operator network through alternative non-cellular access technologies is not supported in the currently specified standards. All though support for IP session handovers and IP flow mobility across access technologies already exists in cellular standards [TS23.402], however, QoS policy handovers across access technologies has not received much attention so far.

Based on the deployment trends, Wireless LAN (WLAN) can be considered as the dominant alternative access technology to complement cellular radio access. Since the 802.11e extension provides QoS extensions to WLAN, it is beneficial to apply QoS policies to WLAN access, which enables QoS classification of downlink as well as uplink traffic between a mobile node and its local mobility anchor. For realizing this capability this specification identifies three functional operations:

- (a) Maintaining QoS classification during a handover between cellular radio access and WLAN access by means of establishing QoS policies in the handover target access network,
- (b) mapping of QoS classes and associated policies between different access systems and
- (c) establishment of QoS policies for new data sessions/flows, which are initiated while using WLAN access.

This document specifies an extension to the PMIPv6 protocol [RFC5213] to establish QoS policies for a mobile node's data traffic on the local mobility anchor and the mobile access gateway. QoS policies are conveyed in-band with PMIPv6 signaling using the specified QoS option and are enforced on the local mobility anchor for downlink traffic and on the mobile access gateway and its access network for the uplink traffic. The specified option allows association between IP session classification characteristics, such as a Differentiated Services Code Point (DSCP) [RFC2474], and the expected QoS class for

the IP session. This document specifies fundamental QoS attributes which apply on a per mobile node, mobility session or on a per-flow basis. The specified attributes are not specific to any access technology, but are compatible with the Third Generation Partnership Project (3GPP) and IEEE 802.11 Wireless LAN QoS specifications.

Additional QoS attributes can be specified and used with the QoS option, e.g. to represent more specific descriptions of latency constraints or jitter bounds. The specification of such additional QoS attributes as well as the handling of QoS policies between the mobile access gateway and the access network are out of scope for this specification.

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 used in this document are to be interpreted as defined in the Proxy Mobile IPv6 specifications [RFC5213], [RFC5844], and [RFC7077]. Additionally, this document uses the following abbreviations:

Aggregate Maximum Bit Rate (AMBR)

AMBR defines the upper limit on the bit-rate that can be provided by the network for a set of IP flows. IP packets exceeding the AMBR limit will be discarded by the rate-shaping function where the AMBR parameter is enforced. Variants of AMBR term can be defined by restricting the target set of IP flows on which the AMBR is applied to a mobile node, mobility session or flow direction. For example, Per Mobile Node Aggregate Maximum Downlink Bit Rate, Per Mobile Node Aggregate Maximum Uplink Bit Rate, Per Mobility Session Aggregate Maximum Downlink Bit Rate and Per Mobility Session Aggregate Maximum Uplink Bit Rate are used in this document.

Allocation and Retention Priority (ARP)

ARP is used in congestion situations when there are insufficient resources for meeting all services requests. It is used primarily by the Admission Control function to determine whether a particular service request must be rejected due to lack of resources, or if it must be honored by preempting an existing low-priority service.

Differentiated Services Code Point (DSCP)

In Differentiated Services Architecture [RFC2474], packets are classified and marked to receive a particular per-hop forwarding behavior on nodes along their path based on the marking present on the packet. This marking on IPv4 and IPv6 packets that defines a specific Per-hop behavior is known as DSCP. Refer to [RFC2474], [RFC2475], [RFC4594] and [RFC2983] for a complete explanation. Please also refer to

Downlink (DL) Traffic

The mobile node's IP packets that the mobile access gateway receives from the local mobility anchor is referred to as the Downlink traffic. The "Downlink" term used in the QoS attribute definition is always from the reference point of the mobile node and it implies traffic heading towards the mobile node.

Guaranteed Bit Rate (GBR)

GBR denotes the assured bit-rate that will be provided by the network for a set of IP flows. It is assumed that the network reserves the resources for supporting the GBR parameter. Variants of the GBR term can be defined by limiting the scope of the target IP flows on which the GBR is applied to a mobile node, mobility session or flow direction. For example, Guaranteed Downlink Bit Rate and Guaranteed Uplink Bit Rate are used in this document.

Mobility Session

The term mobility session, is defined in [RFC5213]. It refers to the creation or existence of state associated with the mobile node's mobility binding on the local mobility anchor and on the mobile access gateway.

QoS Service Request

A set of QoS parameters that are defined to be enforced on one or more mobile node's IP flows. The parameters at the minimum include a DSCP marking and additionally may include Guaranteed Bit Rate or Aggregate Maximum Bit Rate. The Quality of Service option defined in this document represents a QoS Service Request.

Service Identifier

In some mobility architectures, multiple services within the same mobility service subscription are offered to a mobile node. Each of those services provide a specific service (examples: Internet Service, Voice Over IP Service) and has an identifier called Service Identifier. 3GPP APN (Access Point Name) is an example of a Service Identifier. Refer to [RFC5149] for the definition of the Service Identifier and the mobility option used for carrying the Service Identifier.

Uplink (UL) Traffic

The mobile node's IP packets that the mobile access gateway forwards to the local mobility anchor is referred to as the Uplink traffic. The "Uplink" term used in the QoS attribute definitions is based on the reference point of the mobile node and "Uplink" implies traffic originating from the mobile node.

3. Overview of QoS Support in Proxy Mobile IPv6

The Quality of Service support in Proxy Mobile IPv6 specified in this document is based on the Differentiated-Services architecture ([RFC2474] and [RFC2475]). The access and the home network in the Proxy Mobile IPv6 domain are assumed to be DiffServ enabled, with every network node in the forwarding path for the mobile node's IP traffic being Diffserv compliant. The per-hop behavior for providing differential treatment based on the DiffServ marking in the packet is assumed to be supported in the Proxy Mobile IPv6 domain.

The local mobility anchor in the home network and the mobile access gateway in the access network define the network boundary between the access and the home network. These entities being the entry and exit points for the mobile node's IP traffic, are the logical choice for being chosen as the QoS enforcement points. The basic QoS functions such as marking, metering, policing and rate-shaping on the mobile node's IP flows can be enforced at these nodes.

The local mobility anchor and the mobile access gateway can negotiate the Quality of Service parameters for a mobile node's IP flows based on the signaling extensions defined in this document. The QoS services that can be enabled for a mobile node are for meeting both the quantitative performance requirements (such as Guaranteed Bit-Rate) and as well for realizing relative performance treatment by the ways of class-based differentiation. The subscriber's policy and the charging profile [TS22.115] is a key consideration for the mobility entities in the QoS service negotiation. The decision on the type of QoS services that are to be enabled for a mobile node is based on the subscriber profile and based on available network resources. The negotiated QoS parameters are used for providing QoS service differentiation on the path between the local mobility anchor and the mobile access gateway. The signaling related to QoS services is strictly between the mobility entities and does not result in per-flow state, or signaling to any other node in the network.

Figure 1: QoS Support

Figure 1 illustrates the support of QoS services in a Proxy Mobile IPv6 domain. The local mobility anchor and the mobile access gateway have negotiated QoS parameters for the mobility sessions belonging to MN-1 and MN-2. A Per-Session-AMBR of 1Mbps and 2 Mbps for MN-1 and MN-2 respectively. Furthermore, different IP flows from MN-1 and MN-2 are given different QoS service treatment. For example, a GBR of 64Kbps for Flow-1 and Flow-5 is assured, a DSCP marking

enforcement of "Z" on Flow-6, and MBR of 100 Kbps on Flow-5;

3.1. Quality of Service Option - Usage Examples

Use Case 1: Figure 2 illustrates a scenario where a local mobility anchor initiates a QoS service request to a mobile access gateway.

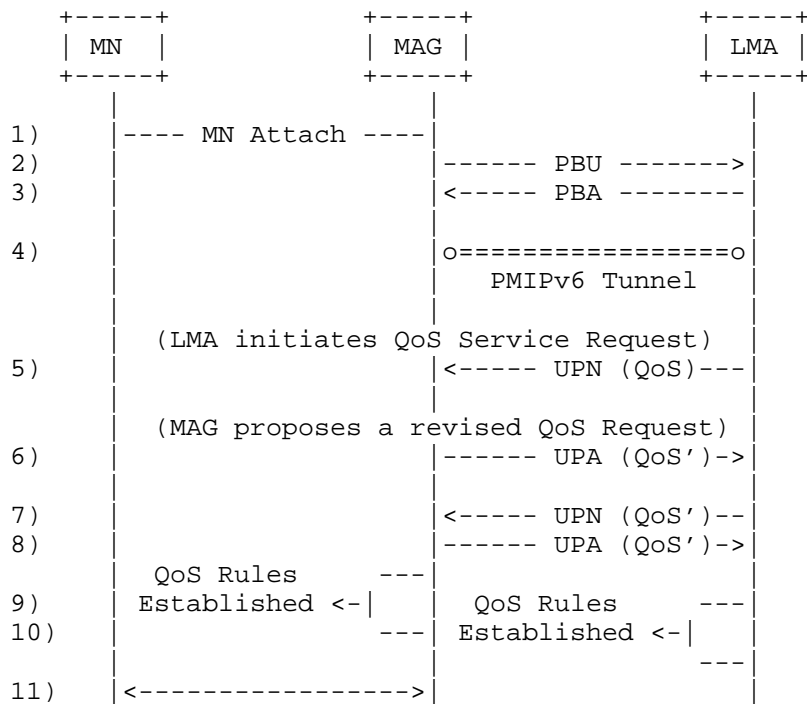


Figure 2: LMA Initiated QoS Service Request

- o (1) to (4): MAG detects the mobile node's attachment to the access link and initiates the signaling with the local mobility anchor. The LMA and MAG upon completing the signaling establish the mobility session and the forwarding state.
- o (5) to (8): The LMA initiates a QoS Service request to the mobile access gateway. The trigger for this service can be based on a trigger from a policy function and the specific details of that trigger are outside the scope of this document. The LMA sends an Update Notification message [RFC7077] to the MAG. The message includes the QoS option Section 4.1 which includes a set of QoS parameters. The MAG on determining that it cannot support the requested QoS service request for that mobile sends an Update Notification Acknowledgement message. The message contains a

revised QoS option with updated set of QoS attributes. The LMA accepts the revised QoS service request by sending a new Update Notification message including the updated QoS option.

- o (9) to (11): Upon successfully negotiating a QoS service request the MAG and the LMA install the QoS rules for that service request. Furthermore, the MAG (using access technology specific mechanisms) install the QoS rules on the access network.

Use Case 2: Figure 3 illustrates a scenario where a mobile access gateway initiates a QoS service request to a local mobility anchor.

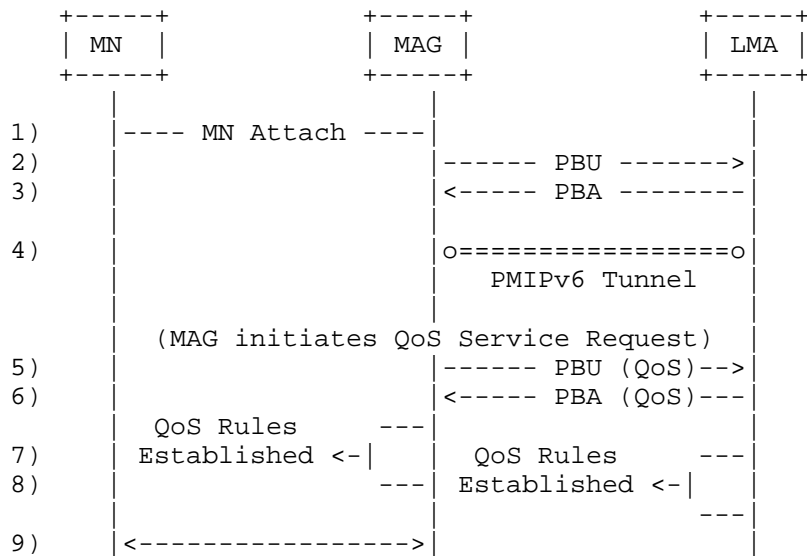


Figure 3: MAG Initiated QoS Service Request

- o (1) to (4): MAG detects the mobile node's attachment to the access link and initiates the signaling with the local mobility anchor. The LMA and MAG upon completing the signaling establish the mobility session and the forwarding state.
- o (5) to (6): The MAG initiates a QoS Service request to the local mobility anchor. The trigger for this service can be based on a trigger from the mobile node using access technology specific mechanisms. The specific details of that trigger are outside the scope of this document. The MAG sends a Proxy Binding Update message [RFC5213] to the LMA. The message includes the QoS option Section 4.1 which includes a set of QoS parameters. The LMA agrees to the proposed QoS service request by sending Proxy

Binding Acknowledgement message.

- o (7) to (9): Upon successfully negotiating a QoS service request the MAG and the LMA install the QoS rules for that service request. Furthermore, the MAG using access technology specific mechanisms install the QoS rules on the access network.

3.2. Quality of Service Attributes - Usage Examples

This section identifies the use-cases where the Quality of Service Option (Section 4.1) and its attributes (Section 4.2) defined in this document are relevant.

- o The subscription policy offered to a mobile subscriber requires the service provider to enforce Aggregate Maximum Bit Rate (AMBR) limits on the subscriber's IP traffic. The local mobility anchor and the mobile access gateway negotiate the uplink and the downlink AMBR values for the mobility session and enforce them in the access and the home network. The QoS option (Section 4.1) with the QoS Attributes, Per-Session-Agg-Max-DL-Bit-Rate (Section 4.2.3) and Per-Session-Agg-Max-UL-Bit-Rate (Section 4.2.4) are used for this purpose.
- o In Community Wi-Fi deployments, the residential gateway participating in the Wi-Fi service is shared between the home user and the community Wi-Fi users. In order to ensure the home user's Wi-Fi service is not impacted because of the community Wi-Fi service, the service provider enables Guaranteed Bit Rate (GBR) for the home user's traffic. The QoS option (Section 4.1) with the QoS Attributes, Guaranteed-DL-Bit-Rate (Section 4.2.8), Guaranteed-UL-Bit-Rate (Section 4.2.9) are used for this purpose.
- o A mobile user using the service provider's Voice over IP infrastructure establishes a VoIP call with some other user in the network. The negotiated call parameters for the VoIP call require a dedicated bandwidth of certain fixed value for the media flows associated with that VoIP session. The Application function in the VoIP infrastructure notifies the local mobility anchor to enforce the GBR limits on that IP flow identified by the flow definition. The QoS option (Section 4.1) with the QoS Attributes, Guaranteed-DL-Bit-Rate (Section 4.2.8), Guaranteed-UL-Bit-Rate (Section 4.2.9), QoS-Traffic-Selector (Section 4.2.10) are used for this purpose.
- o An emergency service may require network resources in conditions when the network resources have been fully allocated to other users and the network may be experiencing severe congestion and in such cases the service provider may want to revoke resources that

have been allocated and reassign them to emergency services. The local mobility anchor and the mobile access gateway negotiate Allocation and Retention Priority (ARP) values for the IP sessions associated with the emergency applications. The QoS option (Section 4.1) with the QoS Attribute, Allocation-Retention-Priority (Section 4.2.5) are used for this purpose.

4. Protocol Messaging Extensions

4.1. Quality of Service Option

The Quality of Service option is a mobility header option used by local mobility anchor and mobile access gateway for negotiating QoS parameters associated with a mobility session. This option can be carried in Proxy Binding Update (PBU) [RFC5213], Proxy Binding Acknowledgement (PBA) [RFC5213], Update Notification (UPN) [RFC7077] and Update Notification Acknowledgement (UPA) [RFC7077] messages. There can be more than one instance of the Quality of Service option in a single message. Each instance of the Quality of Service option represents a specific QoS service request.

The alignment requirement for this option is 4n.

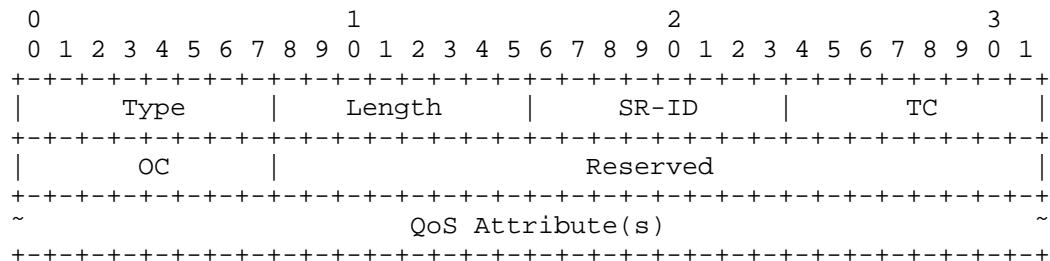


Figure 4: QoS Option

Type

<IANA-1>

Length

8-bit unsigned integer indicating the length of the option in octets, excluding the Type and Length fields.

Service Request Identifier (SR-ID)

A 8-bit unsigned integer used for identifying the QoS service request. Its uniqueness is within the scope of a mobility session. The local mobility anchor always allocates the identifier value. When the QoS Service request is initiated by a mobile access gateway, it sets the value to (0) and the local mobility anchor allocates and includes the value in the

response. For any QoS service requests initiated by a local mobility anchor, the Service Request Identifier is set to the allocated value.

Traffic Class (TC)

Traffic Class consists of a 6-bit DSCP field followed by a 2-bit reserved field.

Differentiated Services Code Point (DSCP)

A 6-bit unsigned integer indicating the code point value, as defined in [RFC2475] to be used for the mobile node's IP flows. When this DSCP marking needs to be applied only for a subset of mobile node's IP flows, there will be a Traffic Selector attribute (Section 4.2.10) in the option which provides the flow selectors. In the absence of any such traffic selector attribute, the DSCP marking applies to all the IP flows associated with the mobility session.

Two-bit Reserved Field

The last two-bits in the Traffic Class field are currently unused. These bits MUST be initialized by the sender to (0) and MUST be ignored by the receiver.

Operational Code (OC)

One-Octet Operational code indicates the type of QoS request.

RESPONSE: (0)

Response to a QoS request

ALLOCATE: (1)

Request to allocate QoS resources

DE-ALLOCATE: (2)

Request to de-Allocate QoS resources

MODIFY: (3)

Request to modify QoS parameters for a previously negotiated QoS service request

QUERY: (4)

Query to list the previously negotiated QoS service requests and that are still active

NEGOTIATE: (5)

Response to a QoS service request with a counter QoS proposal

Reserved: (6) to (255)

Currently not used. Receiver MUST ignore the option received with any value in this range.

Reserved

This field is unused for now. The value MUST be initialized to a value of (0) by the sender and MUST be ignored by the receiver.

QoS Attribute(s)

Zero or more Type-Length-Value (TLV) encoded QoS Attributes. The format of the QoS attribute is defined in Section 4.2. The interpretation and usage of the QoS attribute is based on the value in the "Type" field.

4.2. Quality of Service Attribute

This section identifies the format of a Quality of Service attribute. QoS attribute can be included in the Quality of Service option defined in Section 4.1. The latter part of this section identifies the QoS attributes defined by this specification.




Figure 5: Format of a Quality of Service Attribute

Type: 8-bit unsigned integer indicating the type of the QoS attribute. This specification reserves the following values.

(0) - Reserved

This value is reserved and cannot be used

- (1) - Per-MN-Agg-Max-DL-Bit-Rate
This QoS attribute, Per Mobile Node Aggregate Maximum Downlink Bit Rate, is defined in Section 4.2.1.
- (2) - Per-MN-Agg-Max-UL-Bit-Rate
This QoS attribute, Per Mobile Node Aggregate Maximum Uplink Bit Rate, is defined in Section 4.2.2.
- (3) - Per-Session-Agg-Max-DL-Bit-Rate
This QoS attribute, Per Mobility Session Aggregate Maximum Downlink Bit Rate, is defined in Section 4.2.3.
- (4) - Per-Session-Agg-Max-UL-Bit-Rate
This QoS attribute, Per Mobility Session Aggregate Maximum Uplink Bit Rate, is defined in Section 4.2.4.
- (5) - Allocation-Retention-Priority
This QoS attribute, Allocation and Retention Priority, is defined in Section 4.2.5.
- (6) - Aggregate-Max-DL-Bit-Rate
This QoS attribute, Aggregate Maximum Downlink Bit Rate, is defined in Section 4.2.6.
- (7) - Aggregate-Max-UL-Bit-Rate
This QoS attribute, Aggregate Maximum Uplink Bit Rate, is defined in Section 4.2.7.
- (8) - Guaranteed-DL-Bit-Rate
This QoS attribute, Guaranteed Downlink Bit Rate, is defined in Section 4.2.8.
- (9) - Guaranteed-UL-Bit-Rate
This QoS attribute, Guaranteed Uplink Bit Rate, is defined in Section 4.2.9.

(10) - QoS-Traffic-Selector

This QoS attribute, QoS Traffic Selector, is defined in Section 4.2.10.

(11) - QoS-Vendor-Specific-Attribute

This QoS attribute, QoS Vendor Specific Attribute, is defined in Section 4.2.11.

(12) to (254) - Reserved

These values are reserved for future allocation.

(255) - Reserved

This value is reserved and cannot be used

Length: 8-bit unsigned integer indicating the number of octets needed to encode the Value, excluding the Type and Length fields.

Value: The format of this field is based on the Type value.

4.2.1. Per Mobile Node Aggregate Maximum Downlink Bit Rate

This attribute, Per-MN-Agg-Max-DL-Bit-Rate, represents the maximum downlink bit-rate for a mobile node. It is a variant of the AMBR term defined in Section 2.2. This value is an aggregate across all mobility sessions associated with that mobile node.

This attribute can be included in the Quality of Service option defined in Section 4.1 and it is an optional attribute. There can only be a single instance of this attribute present in a QoS option.


When this attribute is present in a Proxy Binding Update sent by a mobile access gateway, or in a Update Notification message sent by a local mobility anchor, it indicates the maximum aggregate downlink bit-rate that is being requested for the mobile node at the peer.

When this attribute is present in a Proxy Binding Acknowledgement message, or in an Update Notification Acknowledgement message, it indicates the maximum aggregate downlink bit-rate that the peer agrees to offer.

If multiple mobility sessions are established for a mobile node, through multiple mobile access gateways and with sessions anchored either on a single local mobility anchor, or when spread out across multiple local mobility anchors, then it depends on the operator's

policy and the specific deployment as how the total bandwidth for the mobile node on each MAG-LMA pair is computed.

When a QoS option includes both the Per-MN-Agg-Max-DL-Bit-Rate attribute and the QoS Traffic Selector attribute (Section 4.2.10), then the QoS Traffic Selector attribute does not apply to this attribute.



- o Type: 1
- o Length: The length in octets of the attribute, excluding the Type and Length fields. This value is set to (6).
- o Reserved: This field is unused for now. The value MUST be initialized by the sender to 0 and MUST be ignored by the receiver.
- o Per-MN-Agg-Max-DL-Bit-Rate: is a 32-bit unsigned integer, and it indicates the aggregate maximum downlink bit-rate that is requested/allocated for all the mobile node's IP flows. The measurement units for Per-MN-Agg-Max-DL-Bit-Rate are bits-per-second.

4.2.2. Per Mobile Node Aggregate Maximum Uplink Bit Rate

This attribute, Per-MN-Agg-Max-UL-Bit-Rate, represents the maximum uplink bit-rate for the mobile node. It is a variant of the AMBR term defined in Section 2.2. This value is an aggregate across all mobility sessions associated with that mobile node.

This attribute can be included in the Quality of Service option defined in Section 4.1 and it is an optional attribute. There can only be a single instance of this attribute present in a QoS option.

When this attribute is present in a Proxy Binding Update sent by a mobile access gateway, or in an Update Notification message sent by the local mobility anchor, it indicates the maximum aggregate uplink bit-rate that is being requested for the mobile node at the peer.

When this attribute is present in a Proxy Binding Acknowledgement

message, or in an Update Notification Acknowledgement message, it indicates the maximum aggregate uplink bit-rate that the peer agrees to offer for that mobile node.

If multiple mobility sessions are established for a mobile node, through multiple mobile access gateways and with sessions anchored either on a single local mobility anchor, or when spread out across multiple local mobility anchors, then it depends on the operator's policy and the specific deployment as how the total bandwidth for the mobile node on each MAG-LMA pair is computed.

When a QoS option includes both the Per-MN-Agg-Max-UL-Bit-Rate attribute and the QoS Traffic Selector attribute (Section 4.2.10), then the QoS Traffic Selector attribute does not apply to this attribute.

0										1										2										3									
0	1	2	3	4	5	6	7	8	9	0	1	2	3	4	5	6	7	8	9	0	1	2	3	4	5	6	7	8	9	0	1	2	3	4	5	6	7	8	9
Type										Length										Reserved																			
Per-MN-Agg-Max-UL-Bit-Rate																																							

- o Type: 2
- o Length: The length in octets of the attribute, excluding the Type and Length fields. This value is set to (6).
- o Reserved: This field is unused for now. The value MUST be initialized by the sender to 0 and MUST be ignored by the receiver.
- o Per-MN-Agg-Max-UL-Bit-Rate: is of type unsigned 32-bit integer, and it indicates the aggregate maximum uplink bit-rate that is requested/allocated for the mobile node's IP flows. The measurement units for Per-MN-Agg-Max-UL-Bit-Rate are bits-per-second.

4.2.3. Per Mobility Session Aggregate Maximum Downlink Bit Rate

This attribute, Per-Session-Agg-Max-DL-Bit-Rate, represents the maximum downlink bit-rate for the mobility session. It is a variant of the AMBR term defined in Section 2.2.

This attribute can be included in the Quality of Service option defined in Section 4.1 and it is an optional attribute. There can only be a single instance of this attribute present in a QoS option.

When this attribute is present in a Proxy Binding Update sent by a mobile access gateway, or in an Update Notification message sent by the local mobility anchor, it indicates the maximum aggregate downlink bit-rate that is being requested for that mobility session.

When this attribute is present in a Proxy Binding Acknowledgement message, or in an Update Notification Acknowledgement message, it indicates the maximum aggregate downlink bit-rate that the peer agrees to offer for that mobility session.

When a QoS option includes both the Per-Session-Agg-Max-DL-Bit-Rate attribute and the QoS Traffic Selector attribute (Section 4.2.10), then the QoS Traffic Selector attribute does not apply to this attribute.

```

      0               1               2               3
      0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1
+-----+-----+-----+-----+-----+-----+-----+-----+
|      Type      |      Length      |S|E|      Reserved      |
+-----+-----+-----+-----+-----+-----+-----+-----+
|                                     Per-Session-Agg-Max-DL-Bit-Rate
+-----+-----+-----+-----+-----+-----+-----+-----+

```

- o Type: 3
- o Length: The length of the attribute in octets, excluding the Type and Length fields. This value is set to (6).
- o Service (S) flag: This flag is used for extending the scope of the target flows for Per-Session-Agg-Max-DL-Bit-Rate to mobile node's other mobility sessions sharing the same service identifier. 3GPP Access Point Name (APN) is an example of service identifier and that identifier is carried using the Service Selection mobility option [RFC5149].
 - * When the (S) flag is set to a value of (1), then the Per-Session-Agg-Max-DL-Bit-Rate is measured as an aggregate across all the mobile node's other mobility sessions sharing the same service identifier associated with this mobility session.
 - * When the (S) flag is set to a value of (0), then the target flows are limited to the current mobility session.
 - * The (S) flag MUST NOT be set to a value of (1), when there is no service identifier associated with the mobility session.

- o Exclude (E) flag: This flag is used to request that some flows be excluded from the target IP flows for which Per-Session-Agg-Max-DL-Bit-Rate is measured.
 - * When the (E) flag is set to a value of (1), then the request is for excluding the IP flows for which Guaranteed-DL-Bit-Rate (Section 4.2.8) is negotiated, from the flows for which Per-Session-Agg-Max-DL-Bit-Rate applies is measured.
 - * When the (E) flag is set to a value of (0), then the request is not to excluded any IP flows from the target IP flows for which Per-Session-Agg-Max-DL-Bit-Rate is measured.
 - * When the (S) flag and (E) flag are both set to a value of (1), then the request is for excluding all the IP flows sharing the service identifier associated with this mobility session, from the target flows for which Per-Session-Agg-Max-DL-Bit-Rate is measured.
- o Reserved: This field is unused for now. The value MUST be initialized by the sender to 0 and MUST be ignored by the receiver.
- o Per-Session-Agg-Max-DL-Bit-Rate: is a 32-bit unsigned integer, and it indicates the aggregate maximum downlink bit-rate that is requested/allocated for all the IP flows associated with that mobility session. The measurement units for Per-Session-Agg-Max-DL-Bit-Rate are bits-per-second.

4.2.4. Per Mobility Session Aggregate Maximum Uplink Bit Rate


This attribute, Per-Session-Agg-Max-UL-Bit-Rate, represents the maximum uplink bit-rate for the mobility session. It is a variant of the AMBR term defined in Section 2.2.

This attribute can be included in the Quality of Service option defined in Section 4.1 and it is an optional attribute. There can only be a single instance of this attribute present in a QoS option.

When this attribute is present in a Proxy Binding Update sent by a mobile access gateway, or in an Update Notification message [RFC7077] sent by the local mobility anchor, it indicates the maximum aggregate uplink bit-rate that is being requested for that mobility session.

When this attribute is present in a Proxy Binding Acknowledgement message, or in an Update Notification Acknowledgement [RFC7077] message, it indicates the maximum aggregate uplink bit-rate that the peer agrees to offer for that mobility session.

When a QoS option includes both the Per-Session-Agg-Max-UL-Bit-Rate attribute and the QoS Traffic Selector attribute (Section 4.2.10), then the QoS Traffic Selector attribute does not apply to this attribute.



- o Type: 4
- o Length: The length of the attribute in octets, excluding the Type and Length fields. This value is set to (6).
- o Service (S) flag: This flag is used for extending the scope of the target flows for Per-Session-Agg-Max-UL-Bit-Rate to mobile node's other mobility sessions sharing the same service identifier. 3GPP Access Point Name (APN) is an example of service identifier and that identifier is carried using the Service Selection mobility option [RFC5149].
 - * When the (S) flag is set to a value of (1), then the Per-Session-Agg-Max-UL-Bit-Rate is measured as an aggregate across all the mobile node's other mobility sessions sharing the same service identifier associated with this mobility session.
 - * When the (S) flag is set to a value of (0), then the target flows are limited to the current mobility session.
 - * The (S) flag MUST NOT be set to a value of (1), when there is no service identifier associated with the mobility session.
- o Exclude (E) flag: This flag is used to request that some flows be excluded from the target IP flows for which Per-Session-Agg-Max-UL-Bit-Rate is measured.
 - * SGS When the (E) flag is set to a value of (1), then the request is for excluding the IP flows for which Guaranteed-UL-Bit-Rate (Section 4.2.9) is negotiated, from the flows for which Per-Session-Agg-Max-UL-Bit-Rate is measured.
 - * When the (E) flag is set to a value of (0), then the request is not to exclude any IP flows from the target IP flows for which Per-Session-Agg-Max-UL-Bit-Rate is measured.

- * When the (S) flag and (E) flag are both set to a value of (1), then the request is for excluding all the IP flows sharing the service identifier associated with this mobility session, from the target flows for which Per-Session-Agg-Max-UL-Bit-Rate is measured.
- o Reserved: This field is unused for now. The value MUST be initialized by the sender to 0 and MUST be ignored by the receiver.
- o Per-Session-Agg-Max-UL-Bit-Rate: is a 32-bit unsigned integer, and it indicates the aggregate maximum uplink bit-rate that is requested/allocated for all the IP flows associated with that mobility session. The measurement units for Per-Session-Agg-Max-UL-Bit-Rate are bits-per-second.

4.2.5. Allocation and Retention Priority

This attribute, Allocation-Retention-Priority, represents allocation and retention priority for the mobility session or a set of IP flows. It is defined in Section 2.2.

This attribute can be included in the Quality of Service option defined in Section 4.1 and it is an optional attribute. There can only be a single instance of this attribute present in a QoS option.

When the QoS option includes both the Allocation and Retention Priority attribute and the QOS Traffic Selector attribute (Section 4.2.10), then the Allocation and Retention Priority attribute is to be applied at a flow level. The traffic selector in the QOS Traffic Selector attribute identifies the target flows.

When the QoS option including the Allocation and Retention Priority attribute does not include the QOS Traffic Selector attribute (Section 4.2.10), then the Allocation and Retention Priority attribute is to be applied to all the IP flows associated with that mobility session.

```

      0               1               2               3
      0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1
+-----+-----+-----+-----+-----+-----+-----+-----+
|      Type      |      Length      |      Reserved      |  PL  | PC  | PV  |
+-----+-----+-----+-----+-----+-----+-----+

```

- o Type: 5
- o Length: The length of the attribute in octets, excluding the Type and Length fields. This value is set to (10).

- o Reserved: This field is unused for now. The value MUST be initialized by the sender to 0 and MUST be ignored by the receiver.
- o Priority-Level (PL): is a 4-bit unsigned integer value. It is used to decide whether a mobility session establishment or modification request can be accepted; this is typically used for admission control of Guaranteed Bit Rate traffic in case of resource limitations. The priority level can also be used to decide which existing mobility session to pre-empt during resource limitations. The priority level defines the relative timeliness of a resource request.

Values 1 to 15 are defined, with value 1 as the highest level of priority.

Values 1 to 8 should only be assigned for services that are authorized to receive prioritized treatment within an operator domain. Values 9 to 15 may be assigned to resources that are authorized by the home network and thus applicable when a mobile node is roaming.

- o Preemption-Capability (PC): is a 2-bit unsigned integer value. It defines whether a service data flow can get resources that were already assigned to another service data flow with a lower priority level. The following values are defined:

Enabled (0): This value indicates that the service data flow is allowed to get resources that were already assigned to another IP data flow with a lower priority level.

Disabled (1): This value indicates that the service data flow is not allowed to get resources that were already assigned to another IP data flow with a lower priority level. The values (2) and (3) are reserved.

- o Preemption-Vulnerability (PV): is a 2-bit unsigned integer value. It defines whether a service data flow can lose the resources assigned to it in order to admit a service data flow with higher priority level. The following values are defined:

Enabled (0): This value indicates that the resources assigned to the IP data flow can be pre-empted and allocated to a service data flow with a higher priority level.

Disabled (1): This value indicates that the resources assigned to the IP data flow shall not be pre-empted and allocated to a service data flow with a higher priority level. The values (2)

and (3) are reserved.

4.2.6. Aggregate Maximum Downlink Bit Rate

This attribute, `Aggregate-Max-DL-Bit-Rate`, represents the maximum downlink bit-rate for the mobility session. It is a variant of the AMBR term defined in Section 2.2.

This attribute can be included in the Quality of Service option defined in Section 4.1 and it is an optional attribute. There can only be a single instance of this attribute present in a QoS option.

When this attribute is present in a Proxy Binding Update sent by a mobile access gateway, or in an Update Notification message sent by the local mobility anchor, it indicates the maximum aggregate bit-rate for downlink IP flows that is being requested.

When this attribute is present in a Proxy Binding Acknowledgement message, or in an Update Notification Acknowledgement message, it indicates the maximum aggregate downlink bit-rate that the peer agrees to offer.

When a QoS option includes both the `Aggregate-Max-DL-Bit-Rate` attribute and the `QOS-Traffic-Selector` attribute (Section 4.2.10), then the `Aggregate-Max-DL-Bit-Rate` attribute is to be enforced at a flow level and the traffic selectors present in the `QOS-Traffic-Selector` attribute identifies those target flows.

When the QoS option that includes the `Aggregate-Max-DL-Bit-Rate` attribute does not include the `QOS-Traffic-Selector` attribute (Section 4.2.10), then the `Aggregate-Max-DL-Bit-Rate` attribute is to be applied to all the IP flows associated with the mobility session.

```

      0               1               2               3
      0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
|   Type   |   Length   |   Reserved   |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
|                                     Aggregate-Max-DL-Bit-Rate                                     |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+

```

- o Type: 6
- o Length: The length of the attribute in octets, excluding the Type and Length fields. This value is set to (6).
- o Reserved: This field is unused for now. The value MUST be initialized by the sender to 0 and MUST be ignored by the

receiver.

- o Aggregate-Max-DL-Bit-Rate: is a 32-bit unsigned integer, and it indicates the aggregate maximum downlink bit-rate that is requested/allocated for downlink IP flows. The measurement units for Aggregate-Max-DL-Bit-Rate are bits-per-second.

4.2.7. Aggregate Maximum Uplink Bit Rate

This attribute, Aggregate-Max-UL-Bit-Rate, represents the maximum uplink bit-rate for the mobility session. It is a variant of the AMBR term defined in Section 2.2.

This attribute can be included in the Quality of Service option defined in Section 4.1 and it is an optional attribute. There can only be a single instance of this attribute present in a QoS option.

When this attribute is present in a Proxy Binding Update sent by a mobile access gateway, or in an Update Notification message sent by the local mobility anchor, it indicates the maximum aggregate uplink bit-rate that is being requested.

When this attribute is present in a Proxy Binding Acknowledgement message, or in an Update Notification Acknowledgement message, it indicates the maximum aggregate uplink bit-rate that the peer agrees to offer.

When a QoS option includes both the Aggregate-Max-UL-Bit-Rate attribute and the QOS-Traffic-Selector attribute (Section 4.2.10), then the Aggregate-Max-UL-Bit-Rate attribute is to be enforced at a flow level and the traffic selectors present in the QOS-Traffic-Selector attribute identifies those target flows.

When the QoS option that includes the Aggregate-Max-UL-Bit-Rate attribute does not include the QOS-Traffic-Selector attribute (Section 4.2.10), then the Aggregate-Max-UL-Bit-Rate attribute is to be applied to all the IP flows associated with the mobility session.

0									1									2									3												
0	1	2	3	4	5	6	7	8	9	0	1	2	3	4	5	6	7	8	9	0	1	2	3	4	5	6	7	8	9	0	1	2	3	4	5	6	7	8	9
Type									Length									Reserved																					
Aggregate-Max-UL-Bit-Rate																																							

- o Type: 7
- o Length: The length of the attribute in octets, excluding the Type and Length fields. This value is set to (6).
- o Reserved: This field is unused for now. The value MUST be initialized by the sender to 0 and MUST be ignored by the receiver.
- o Per-Session-Agg-Max-UL-Bit-Rate: is a 32-bit unsigned integer, and it indicates the aggregate maximum uplink bit-rate that is requested/allocated for all the IP flows associated with that mobility session. The measurement units for Aggregate-Max-UL-Bit-Rate are bits-per-second.

4.2.8. Guaranteed Downlink Bit Rate

This attribute, Guaranteed-DL-Bit-Rate, represents the assured bit-rate on the downlink path that will be provided for a set of IP flows associated with a mobility session. It is a variant of the GBR term defined in Section 2.2.


This attribute can be included in the Quality of Service option defined in Section 4.1 and it is an optional attribute. There can only be a single instance of this attribute present in a QoS option.

When this attribute is present in a Proxy Binding Update sent by a mobile access gateway, or in a Update Notification message sent by the local mobility anchor, it indicates the guaranteed downlink bit-rate that is being requested.

When this attribute is present in a Proxy Binding Acknowledgement message, or in an Update Notification Acknowledgement message, it indicates the guaranteed downlink bit-rate that the peer agrees to offer.

When a QoS option includes both the Guaranteed-DL-Bit-Rate attribute and the QOS-Traffic-Selector attribute (Section 4.2.10), then the Guaranteed-DL-Bit-Rate attribute is to be enforced at a flow level and the traffic selectors present in the QOS-Traffic-Selector attribute identifies those target flows.

When the QoS option that includes the Guaranteed-DL-Bit-Rate attribute does not include the QOS-Traffic-Selector attribute (Section 4.2.10), then the Guaranteed-DL-Bit-Rate attribute is to be applied to all the IP flows associated with the mobility session.



- o Type: 8
- o Length: The length of the attribute in octets, excluding the Type and Length fields. This value is set to (6).
- o Reserved: This field is unused for now. The value MUST be initialized by the sender to 0 and MUST be ignored by the receiver.
- o Guaranteed-DL-Bit-Rate: is of type unsigned 32-bit integer, and it indicates the guaranteed bandwidth in bits-per-second for downlink IP flows. The measurement units for Guaranteed-DL-Bit-Rate are bits-per-second.

4.2.9. Guaranteed Uplink Bit Rate

This attribute, Guaranteed-UL-Bit-Rate, represents the assured bit-rate on the uplink path that will be provided for a set of IP flows associated with a mobility session. It is a variant of the GBR term defined in Section 2.2.


This attribute can be included in the Quality of Service option defined in Section 4.1 and it is an optional attribute. There can only be a single instance of this attribute present in a QoS option.

When this attribute is present in a Proxy Binding Update sent by a mobile access gateway, or in a Update Notification message sent by the local mobility anchor, it indicates the guaranteed uplink bit-rate that is being requested.

When this attribute is present in a Proxy Binding Acknowledgement message, or in an Update Notification Acknowledgement message, it indicates the guaranteed uplink bit-rate that the peer agrees to offer.

When a QoS option includes both the Guaranteed-UL-Bit-Rate attribute and the QOS-Traffic-Selector attribute (Section 4.2.10), then the Guaranteed-UL-Bit-Rate attribute is to be enforced at a flow level and the traffic selectors present in the QOS-Traffic-Selector attribute identifies those target flows.

When the QoS option that includes the Guaranteed-UL-Bit-Rate attribute does not include the QoS-Traffic-Selector attribute (Section 4.2.10), then the Guaranteed-UL-Bit-Rate attribute is to be applied to all the IP flows associated with the mobility session.




- o Type: 9
- o Length: The length of the attribute in octets, excluding the Type and Length fields. This value is set to (6).
- o Reserved: This field is unused for now. The value MUST be initialized by the sender to 0 and MUST be ignored by the receiver.
- o Guaranteed-UL-Bit-Rate: is of type unsigned 32-bit integer, and it indicates the guaranteed bandwidth in bits-per-second for uplink IP flows. The measurement units for Guaranteed-UL-Bit-Rate are bits-per-second.

4.2.10. QoS Traffic Selector

This attribute, QoS-Traffic-Selector, includes the parameters used to match packets for a set of IP flows.

This attribute can be included in the Quality of Service option defined in Section 4.1 and it is an optional attribute.

When a QoS option that includes the QoS-Traffic-Selector also includes any one or more of the attributes, Allocation-Retention-Priority (Section 4.2.5), Aggregate-Max-DL-Bit-Rate (Section 4.2.6), Aggregate-Max-UL-Bit-Rate (Section 4.2.7), Guaranteed-DL-Bit-Rate (Section 4.2.8), and Guaranteed-UL-Bit-Rate (Section 4.2.9), then those included attributes are to be enforced at a flow level and the traffic selectors present in the QoS-Traffic-Selector attribute identifies those target flows. Furthermore, the DSCP marking in the QoS option is to be applied only to partial set of mobile node's IP flows and the traffic selectors present in the QoS-Traffic-Selector attribute identifies those target flows.




- o Type: 10
- o Length: The length of the attribute in octets, excluding the Type and Length fields.
- o Reserved: This field is unused for now. The value MUST be initialized by the sender to 0 and MUST be ignored by the receiver.
- o TS Format: An 8-bit unsigned integer indicating the Traffic Selector Format. The values are allocated from the "Traffic Selector Format" namespace for the traffic selector sub-option defined in [RFC6089]; those defined in [RFC6089] are repeated here for clarity. Value (0) is reserved and MUST NOT be used. When the value of TS Format field is set to (1), the format that follows is the IPv4 Binary Traffic Selector specified in section 3.1 of [RFC6088], and when the value of TS Format field is set to (2), the format that follows is the IPv6 Binary Traffic Selector specified in section 3.2 of [RFC6088].
- o Traffic Selector: variable-length field for including the traffic specification identified by the TS format field.

4.2.11. QoS Vendor Specific Attribute

This attribute is used for carrying vendor specific QoS attributes. The interpretation and the handling of this option is specific to the vendor implementation.

This attribute can be included in the Quality of Service option defined in Section 4.1 and it is an optional attribute. There can be multiple instances of this attribute with different sub-type values present in a single QoS option.



- o Type: 11
- o Length: The length of the attribute in octets, excluding the Type and Length fields.
- o Reserved: This field is unused for now. The value MUST be initialized by the sender to 0 and MUST be ignored by the receiver.
- o Vendor ID: The Vendor ID is the SMI (Structure of Management Information) Network Management Private Enterprise Code of the IANA-maintained Private Enterprise Numbers registry [SMI].
- o Sub-Type: An 8-bit field indicating the type of vendor-specific information carried in the option. The name space for this Sub-type is managed by the Vendor identified by the Vendor ID field.

4.3. New Status Code for Proxy Binding Acknowledgement

This document defines the following new Status Code value for use in Proxy Binding Acknowledgement message.

CANNOT_MEET_QOS_SERVICE_REQUEST (Cannot meet QoS Service Request):
<IANA-2>

4.4. New Notification Reason for Update Notification Message

This document defines the following new Notification Reason value for use in Update Notification message.

QOS_SERVICE_REQUEST (QoS Service Requested): <IANA-3>

4.5. New Status Code for Update Notification Acknowledgement Message

This document defines the following new Status code value for use in Update Notification Acknowledgement message.

CANNOT_MEET_QOS_SERVICE_REQUEST (Cannot meet QoS Service Request):

<IANA-4>

5. Protocol Considerations

5.1. Local Mobility Anchor Considerations

- o The conceptual Binding Cache entry data structure maintained by the local mobility anchor, described in Section 5.1 of [RFC5213], can be extended to store a list of negotiated Quality of Service requests to be enforced. There can be multiple such entries and each entry must include the Service Request Identifier, DSCP value and the attributes defined in Section Section 4.2.

LMA Receiving a QoS Service Request:

- o On receiving a Proxy Binding Update message with one or more instances of Quality of Service option included in the message, the local mobility anchor processes the option(s) and determines if the QoS service request for the proposed QoS service request(s) can be met. Each instance of the Quality of Service option represents a specific QoS service request. This determination to accept the request(s) can be based on policy configured on the local mobility anchor, available network resources, or based on other considerations.
- o If the local mobility anchor can support the proposed QoS service requests in entirety, then it sends a Proxy Binding Acknowledgement message with a status code value of (0).
 - * The message includes all the Quality of Service option instances copied (including all the option content) from the received Proxy Binding Update message. However, if the Operational Code field in the request is a QUERY, then the message includes all the Quality of Service option(s) reflecting the currently negotiated QoS service requests for that mobility session.
 - * The Operational Code field in each of the Quality of Service option(s) is set to RESPONSE.
 - * The local mobility anchor should enforce the Quality of Service rules for all the negotiated QoS service requests on the mobile node's uplink and downlink traffic.
- o If the local mobility anchor cannot support any of the requested QoS service requests in entirety, it rejects the request and sends a Proxy Binding Acknowledgement message with the status code value set to CANNOT_MEET_QOS_SERVICE_REQUEST (Cannot meet QoS Service Request).

- * The denial for QoS service request MUST NOT result in removal of the mobility session for that mobile node.
- * The Operational Code field in each of the Quality of Service option(s) is set to RESPONSE.
- * The Proxy Binding Acknowledgement message may include the Quality of Service option based on the following considerations.
 - + If the local mobility anchor cannot support QoS services for that mobile node, then Quality of Service option is not included in the Proxy Binding Acknowledgement message. This serves as an indication to the mobile access gateway that QoS services are not supported for that mobile node.
 - + If the local mobility anchor can support QoS services for that mobile node, but for a downgraded/revised QoS service request, or for a partial set of QoS service requests, the updated Quality of Service option(s) is included in the Proxy Binding Acknowledgement message. This includes the case, where the Attributes in a QoS option have conflicting requirements, Ex: Per-Session-Agg-Max-UL-Bit-Rate is lower than the Guaranteed-UL-Bit-Rate. The contents of each of the option (including the QoS attributes) reflect the QoS service parameters that the local mobility anchor can support for that mobile node. The Operational Code field in each of the Quality of Service option(s) is set to NEGOTIATE. This serves as an indication for the mobile access gateway to resend the Proxy Binding Update message with the revised QoS parameters.

LMA Sending a QoS Service Request:

- o The local mobility anchor, at any time, can initiate a QoS service request for mobile node, by sending an Update Notification message [RFC7077]. The Notification Reason in the Update Notification message is set to a value of QOS_SERVICE_REQUEST and the Acknowledgement Requested (A) flag set to a value of (1).
- * New QoS service request:
 - + The message includes a Quality of Service option with one or more QoS attributes included in the option.
 - + The Operational Code field in the Quality of Service option is set to ALLOCATE.

- + The Service Request Identifier is set to a value of (0).
- + The DSCP field in the Traffic Class (TC) field reflects the requested DSCP value.
- * Modification of an existing QoS Service Request:
 - + The message includes a Quality of Service option with the QoS attributes reflecting the updated values in the Attributes, and the updated list of Attributes.
 - + The Operational Code field in the Quality of Service option is set to MODIFY.
 - + The Service Request Identifier is set to a value that was allocated for that QoS service request.
 - + There can be more than one QoS service request in a single message. If so, the message includes an instance of a Quality of Service option for each of those service requests.
- * Deletion of an existing QoS Service Request:
 - + The Operational Code field in the Quality of Service option is set to DE-ALLOCATE.
 - + The Service Request Identifier is set to a value that was allocated for that QoS service request.
 - + The message includes a Quality of Service option with the QoS attributes reflecting the updated values for the attributes.
- * Query for the previously negotiated QoS Service Requests:
 - + The Operational Code field in the Quality of Service option is set to QUERY.
 - + The Service Request Identifier is set to a value of (0).
 - + The message includes a single instance of the Quality of Service option without including any QoS Attributes.
- o Handling a Response to the QoS Service Request:
 - * If the received Update Notification Acknowledgement [RFC7077] message has the status code field set to value of (0), the

local mobility anchor should enforce the Quality of Service rules for the negotiated QoS parameters on the mobile node's uplink and downlink traffic.

- * If the received Update Notification Acknowledgement message is with the status code field set to value of (CANNOT_MEET_QOS_SERVICE_REQUEST), the local mobility anchor applies the following considerations.
 - + The denial of QoS service request results in removal of any of the mobile node's Binding Cache entries.
 - + If the message did not include any Quality of Service option(s), then it is an indication from the mobile access gateway that QoS services are not enabled for the mobile node.
 - + If the Operational Code field in the Quality of Service option is set to a value of NEGOTIATE and the message includes one or more instances of the Quality of Service option, but the option contents reflect a downgraded/revised set of QoS parameters, then the local mobility anchor MAY choose to agree to proposed QoS service request by resending a new Proxy Binding Update message with the updated Quality of Service option.

General Considerations:

- o Any time the local mobility anchor removes a mobile node's mobility session by removing a Binding Cache entry [RFC5213], for which QoS resources have been previously allocated, those allocated resources are released.
- o Any time the local mobility anchor receives a Proxy Binding Update with HI hint = 3 (inter-MAG handover), the local mobility anchor when sending a Proxy Binding Acknowledgement message includes the QoS option(s) for each of the QoS service requests that are active for that mobile node. This allows the mobile access gateway to allocate QoS resources on the current path. This is relevant for the scenario where a mobile node performs an handover to a new mobile access gateway which is unaware of the previously negotiated QoS services.

5.2. Mobile Access Gateway Considerations

- o The conceptual Binding Update List entry data structure maintained by the mobile access gateway, described in Section 6.1 of [RFC5213], can be extended to store a list of negotiated Quality

of Service requests to be enforced. There can be multiple such entries and entry including the Service Request Identifier, DSCP value and the attributes defined in Section Section 4.2.

MAG Receiving a QoS Service Request:

- o On receiving a Update Notification message with one or more instances of Quality of Service option included in the message, the mobile access gateway processes the option(s) and determine if the QoS service request for the proposed QoS service request(s) can be met. Each instance of the Quality of Service option represents a specific QoS service request. This determination to accept the request(s) can be based on policy configured on the mobile access gateway, available network resources, or based on other considerations.
- o If the mobile access gateway can support the proposed QoS service requests in entirety, then it sends a an Update Notification Acknowledgement message with status code value of (0).
 - * The message includes all the Quality of Service option instances copied (including all the option content) from the received Update Notification message. However, if the Operational Code field in the request is a QUERY, then the message includes all the Quality of Service option(s) reflecting the currently negotiated QoS service requests for that mobility session.
 - * The Operational Code field in each of the Quality of Service option(s) is set to RESPONSE.
 - * The mobile access gateway should enforce the Quality of Service rules for all the negotiated QoS service requests on the mobile node's uplink and downlink traffic.
- o If the mobile access gateway cannot support any of the requested QoS service requests in entirety, then it rejects the request and send an Update Notification Acknowledgement message with the status code set to CANNOT_MEET_QOS_SERVICE_REQUEST (Cannot meet QoS Service Request).
 - * The denial for QoS service request MUST NOT result in removal of the mobility session for that mobile node.
 - * The Operational Code field in each of the Quality of Service option(s) is set to RESPONSE.

- * The Update Notification Acknowledgement message may include the Quality of Service option(s) based on the following considerations.
 - + If the mobile access gateway cannot support QoS services for that mobile node, then Quality of Service option is not included in the Update Notification Acknowledgement message. This serves as an indication to the local mobility anchor that QoS services are not supported for that mobile node.
 - + If the mobile access gateway can support QoS services for that mobile node, but for a downgraded/revise QoS service request, or for a partial set of QoS service requests, then the updated Quality of Service option(s) is included in the Update Notification Acknowledgement message. This includes the case, where the Attributes in a QoS option have conflicting requirements, Ex: Per-Session-Agg-Max-UL-Bit-Rate is lower than the Guaranteed-UL-Bit-Rate. The contents of each of the option (including the QoS attributes) reflect the QoS service parameters that the mobile access gateway can support for that mobile node. The Operational Code field in each of the Quality of Service option(s) is set to NEGOTIATE. This serves as an indication to the local mobility anchor to resend the Update Notification message with the revised QoS parameters.

MAG Sending a QoS Service Request:

- o The mobile access gateway, at any time, can initiate a QoS service request for a mobile node, by sending a Proxy Binding Update message. The QoS service request can be initiated as part of the initial Binding registration, or during binding re-registrations.
 - * New QoS service request:
 - + The message includes a Quality of Service option with one or more QoS attributes included in the option.
 - + The Operational Code field in the Quality of Service option is set to ALLOCATE.
 - + The Service Request Identifier is set to a value of (0).
 - + The DSCP value in the Traffic Class field reflects the requested DSCP value.

- * Modification of an existing QoS Service Request:
 - + The message includes a Quality of Service option with the QoS attributes reflecting the updated values in the Attributes, and the updated list of Attributes.
 - + The Operational Code field in the Quality of Service option is set to MODIFY.
 - + The Service Request Identifier is set to a value that was allocated for that QoS service request.
 - + There can be more than one QoS service request in a single message. If so, the message includes an instance of a Quality of Service option for each of those service requests.
- * Deletion of an existing QoS Service Request:
 - + The Operational Code field in the Quality of Service option is set to DE-ALLOCATE.
 - + The Service Request Identifier is set to a value that was allocated for that QoS service request.
 - + The message includes a Quality of Service option with the QoS attributes reflecting the updated values for the attributes.
- * Query for the previously negotiated QoS Service Requests:
 - + The Operational Code field in the Quality of Service option is set to QUERY.
 - + The Service Request Identifier is set to a value of (0).
 - + The message includes a single instance of the Quality of Service option without including any QoS Attributes.
- o Handling a Response to the QoS Service Request:
 - * If the received Proxy Binding Acknowledgement message has the status code field set to a value of (0), the mobile access gateway should enforce the Quality of Service rules for the negotiated QoS parameters on the mobile node's uplink and downlink traffic.

- * If the received Proxy Binding Acknowledgement message has the status code field set to a value of (CANNOT_MEET_QOS_SERVICE_REQUEST), the mobile access gateway applies the following considerations.
 - + The denial of QoS service request MUST NOT result in removal of any of the mobile node's Binding Update list entries.
 - + If the message did not include any Quality of Service option(s), then it is an indication from the local mobility anchor that QoS services are not enabled for the mobile node.
 - + If the Operational Code field in the Quality of Service option is set to a value of NEGOTIATE and the message includes one or more instances of the Quality of Service option, but the option contents reflect a downgraded/revised set of QoS parameters, then the mobile access gateway MAY choose to agree to proposed QoS service request by resending a new Proxy Binding Update message with the updated Quality of Service option.
- * General Considerations:
 - + There can be more than one QoS service request in a single message. If so, the message includes an instance of a Quality of Service option for each of those service requests. Furthermore, the DSCP value is different in each of those requests.
 - + Any time the mobile access gateway removes a mobile node's mobility session by removing a Binding Update List entry [RFC5213], for which QoS resources have been previously allocated, those allocated resources are released.

6. QoS Services in Integrated WLAN-3GPP Networks

6.1. Technical Scope and Procedure

The QoS option specified in this document can provide the equivalent level of QoS information defined in 3GPP, which is used to enforce QoS policies for IP flows, which have been established while the mobile node is attached to WLAN access, or moved from 3GPP to WLAN access. The QoS classification defined by the 3GPP specification is provided by Differentiated Services techniques in the IP transport network and translated as appropriate into WLAN QoS specification in WLAN access, the details of which are described in Appendix A and Appendix B.

Figure 6 illustrates a generalized architecture where the QoS option can be used. The QoS policies could be retrieved from a Policy Control Function (PCF), such as defined in current cellular mobile communication standards, which aims to assign an appropriate QoS class to a mobile node's individual flows. Alternatively, more static and default QoS rules could be made locally available, e.g. on a local mobility anchor, through administration.

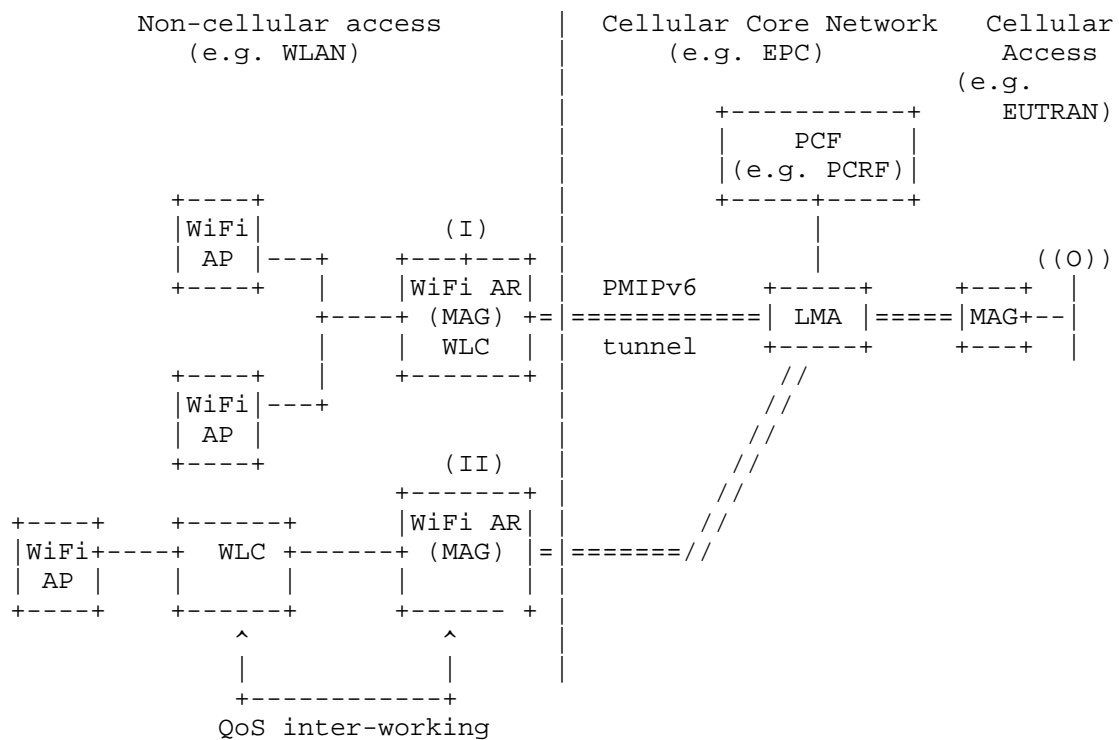


Figure 6: Architecture for QoS inter-working between cellular access and non-cellular access

During a mobile node's handover from cellular access to non-cellular access, e.g. a wireless LAN (WLAN) radio access network, the mobile node's QoS policy rules, as previously established on the local mobility anchor for the mobile node's communication through the cellular access network, are moved to the handover target mobile access gateway serving the non-cellular access network. Such non-cellular mobile access gateway can have an access technology specific controller or function co-located, e.g. a Wireless LAN Controller (WLC), as depicted in option (I) of Figure 6. Alternatively, the access specific architecture can be distributed and the access technology specific control function is located external to the mobile access gateway, as depicted in option (II). In this case, the mobile access gateway and the access technology specific control function (e.g. the WLC) must provide some protocol for QoS inter-working. Details of such inter-working are out of scope of this specification.

6.2. Relevant QoS Attributes

The QoS Option shall at least contain a DSCP value being associated with IP flows of a mobility session. The DSCP value should correspond to the 3GPP QoS Class Index (QCI), which identifies the type of service in term of QoS characteristics (e.g. conversational voice, streaming video, signalling, best effort,...); more details on DSCP and QCI mapping are given on section Appendix A. Optional QoS information could also be added. For instance, in order to comply with the bearer model defined in 3GPP [TS23.203], the following QoS parameters are conveyed for each PMIPv6 mobility session:

- o Default, non-GBR bearer (QCI=5-9)

- * DSCP=(BE, AF11, AF21, AF31, AF32)
- * Per-MN AMBR-UL/DL
- * Per-Session AMBR-UL/DL {S=1,E=1}
- * AARP

APN (Access Point Name) is provided via the Service Selection ID defined in [RFC5149]. If APN is not interpreted by Wi-Fi AP, the latter will police only based on Per-MN AMBR-UL/DL (without Per-Session AMBR-UL/DL) on the Wi-Fi link.

- o Dedicated, GBR bearer (QCI=1-4)

- * DSCP=(EF, AF41)
- * GBR-UL/DL
- * MBR-UL/DL
- * AARP
- * TS

Wi-Fi AP will perform the policy enforcement with the minimum bit-rate=GBR and the maximum bit-rate=MBR.

- o Dedicated, non-GBR bearer (QCI=5-9)

- * DSCP=(BE, AF11, AF21, AF31, AF32)
- * Per-MN AMBR-UL/DL

- * Per-Session AMBR-UL/DL {S=1,E=1}
- * AARP
- * TS

If APN is not interpreted by Wi-Fi AP, it will police based only on Per-MN AMBR-UL/DL (without Per-Session AMBR-UL/DL) on the Wi-Fi link.

If DSCP values follow the 3GPP specification and deployment, the code point can carry intrinsically additional attributes according to Figure 7.

For some optional QoS attributes the signalling can differentiate enforcement per mobility session and per IP flow. For the latter, as long as the AMBR constraints are met, the rule associated with the identified flow(s) overrules the aggregated rules which apply per Mobile Node or per Mobility Session. Additional attributes can be appended to the QoS option, but their definition and specification is out of scope of this document and left to their actual deployment.

7. IANA Considerations

This document requires the following IANA actions.

- o Action-1: This specification defines a new mobility option, the Quality of Service (QoS) option. The format of this option is described in Section 4.1. The type value <IANA-1> for this mobility option needs to be allocated from the Mobility Options registry at <<http://www.iana.org/assignments/mobility-parameters>>. RFC Editor: Please replace <IANA-1> in Section 4.1 with the assigned value and update this section accordingly.
- o Action-2: This specification defines a new mobility attribute format, Quality of Service attribute. The format of this attribute is described in Section Section 4.2. This attribute can be carried in the Quality of Service mobility option. The type values for this attribute need to be managed by IANA in a new Registry, the "Quality of Service Attribute Registry". This registry is maintained under "Mobile IPv6 Parameters" registry at <<http://www.iana.org/assignments/mobility-parameters>>. This specification reserves the following type values. All other values (12 - 254) are unassigned and may be assigned by IANA using the Specification Required policy [RFC5226]. Designated Expert reviewing the value assignment is expected to verify that the protocol extension follows the Proxy Mobile IPv6 architecture and does not raise backward compatibility issues with existing deployments.

Value	Description	Reference
0	Reserved	<this draft>
1	Per-MN-Agg-Max-DL-Bit-Rate	<this draft>
2	Per-MN-Agg-Max-UL-Bit-Rate	<this draft>
3	Per-Session-Agg-Max-DL-Bit-Rate	<this draft>
4	Per-Session-Agg-Max-UL-Bit-Rate	<this draft>
5	Allocation-Retention-Priority	<this draft>
6	Aggregate-Max-DL-Bit-Rate	<this draft>
7	Aggregate-Max-UL-Bit-Rate	<this draft>
8	Guaranteed-DL-Bit-Rate	<this draft>
9	Guaranteed-UL-Bit-Rate	<this draft>
10	QoS-Traffic-Selector	<this draft>
11	QoS-Vendor-Specific-Attribtute	<this draft>
255	Reserved	<this draft>

- o Action-3: This document defines a new status value, CANNOT_MEET_QOS_SERVICE_REQUEST (<IANA-2>) for use in Proxy Binding Acknowledgement message, as described in Section 4.3. This value is to be assigned from the "Status Codes" registry at <<http://www.iana.org/assignments/mobility-parameters>>. The allocated value has to be greater than 127. RFC Editor: Please replace <IANA-2> in Section 4.3 with the assigned value and update this section accordingly.
- o Action-4: This document defines a new Notification Reason, QOS_SERVICE_REQUEST (<IANA-3>) for use in Update Notification message [RFC7077] as described in Section 4.4. This value is to be assigned from the "Update Notification Reasons Registry" at <<http://www.iana.org/assignments/mobility-parameters>>. RFC Editor: Please replace <IANA-3> in Section 4.4 with the assigned value and update this section accordingly.

- o Action-5: This document defines a new Notification Reason, CANNOT_MEET_QOS_SERVICE_REQUEST (<IANA-4>) for use in Update Notification Acknowledgement message [RFC7077] as described in Section 4.5. This value is to be assigned from the "Update Notification Acknowledgement Status Registry" at <<http://www.iana.org/assignments/mobility-parameters>>. RFC Editor: Please replace <IANA-4> in Section 4.5 with the assigned value and update this section accordingly.

8. Implementation Status

Note to RFC Editor: Please remove this section and the reference to [RFC6982] before publication.

This section records the status of known implementations of the protocol defined by this specification at the time of posting of this Internet-Draft, and is based on a proposal described in [RFC6982]. The description of implementations in this section is intended to assist the IETF in its decision processes in progressing drafts to RFCs. Please note that the listing of any individual implementation here does not imply endorsement by the IETF. Furthermore, no effort has been spent to verify the information presented here that was supplied by IETF contributors. This is not intended as, and must not be construed to be, a catalog of available implementations or their features. Readers are advised to note that other implementations may exist.

According to [RFC6982], "this will allow reviewers and working groups to assign due consideration to documents that have the benefit of running code, which may serve as evidence of valuable experimentation and feedback that have made the implemented protocols more mature. It is up to the individual working groups to use this information as they see fit".

Cisco Implementation

Organization: Cisco

Description: QoS Extensions to Cisco IOS-based MAG and LMA Implementations. Engineering prototype code under development.

Coverage: Support includes QoS signaling from MAG to LMA based on PBU/PBA and LMA to MAG based on the recently standardized UPN/UPA messages. Implementation includes only a partial set of QoS attributes and support for other Attributes is under development. The QoS option is based on the Vendor-specific mobility option, but it has all the parameters defined in -07 version of the document. We have plans to show a demo in the next IETF.

Licensing: Closed. However, cisco has plans to release the MAG portion of the code for Linux as open source.

Implementation Experience: The feedback from the developer suggests that the protocol extensions needed for this specification proved to be reasonably straightforward. Numerous draft revisions were made based on the questions and comments from the developer. The effort to most part appears to be around

interfacing with the platform specific QoS features for enforcing the negotiated QoS parameters for a subscriber's IP session/flows. On Cisco IOS, there is a programmatic interface with rich semantics for interfacing with IOS MQC. It needs to be seen as how this can be realized on a Linux OS.

Contact: Sri Gundavelli (sgundave@cisco.com)

9. Security Considerations

The quality of service option defined in this specification is for use in Proxy Binding Update, Proxy Binding Acknowledgement, Update Notification, and Update Notification Acknowledgement messages. This option is carried in these message like any other mobility header option. [RFC5213] and [RFC7077] identify the security considerations for these signalling messages. The quality of service option when included in these signalling messages does not require additional security considerations.

10. Acknowledgements

The authors of this document thank the members of NetExt Working Group for the valuable feedback to different versions of this specification. In particular the authors want to thank Basavaraj Patil, Behcet Sarikaya, Charles Perkins, Dirk von Hugo, Mark Grayson, Tricci So, Ahmad Muhanna, Pete McCann, Byju Pularikkal, John Kaippallimalil, Rajesh Pazhyannur, Carlos J. Bernardos Cano, Michal Hoeft, Ryuji Wakikawa, Liu Dapeng, Seil Jeon, Georgios Karagiannis.

The authors would like to thank all the IESG reviewers and specially, Ben Campbell, Barry Leiba, Jari Arkko, Alissa Cooper, Stephen Farrell, Ted Lemon and Alia Atlas for their valuable comments and suggestions to improve this specification.

Finally, the authors would like to express sincere and profound appreciation to our Internet Area Director, Brian Haberman for his guidance and great support in allowing us to complete this work.

11. References

11.1. Normative References

- [RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, March 1997.
- [RFC5213] Gundavelli, S., Leung, K., Devarapalli, V., Chowdhury, K., and B. Patil, "Proxy Mobile IPv6", RFC 5213, August 2008.
- [RFC5226] Narten, T. and H. Alvestrand, "Guidelines for Writing an IANA Considerations Section in RFCs", BCP 26, RFC 5226, May 2008.
- [RFC5844] Wakikawa, R. and S. Gundavelli, "IPv4 Support for Proxy Mobile IPv6", RFC 5844, May 2010.
- [RFC6088] Tsirtsis, G., Giarreta, G., Soliman, H., and N. Montavont, "Traffic Selectors for Flow Bindings", RFC 6088, January 2011.
- [RFC7077] Krishnan, S., Gundavelli, S., Liebsch, M., Yokota, H., and J. Korhonen, "Update Notifications for Proxy Mobile IPv6", RFC 7077, November 2013.

11.2. Informative References

- [GSMA.IR.34] GSMA, "Inter-Service Provider IP Backbone Guidelines 5.0", May 2013.
- [IEEE802.11-2012] IEEE, "Part 11: Wireless LAN Medium Access Control(MAC) and Physical Layer (PHY) specifications", 2012.
- [RFC2474] Nichols, K., Blake, S., Baker, F., and D. Black, "Definition of the Differentiated Services Field (DS Field) in the IPv4 and IPv6 Headers", RFC 2474, December 1998.
- [RFC2475] Blake, S., Black, D., Carlson, M., Davies, E., Wang, Z., and W. Weiss, "An Architecture for Differentiated Services", RFC 2475, December 1998.
- [RFC2983] Black, D., "Differentiated Services and Tunnels", RFC 2983, October 2000.
- [RFC4594] Babiarz, J., Chan, K., and F. Baker, "Configuration

Guidelines for DiffServ Service Classes", RFC 4594, August 2006.

- [RFC5149] Korhonen, J., Nilsson, U., and V. Devarapalli, "Service Selection for Mobile IPv6", RFC 5149, February 2008.
- [RFC6089] Tsirtsis, G., Soliman, H., Montavont, N., Giaretta, G., and K. Kuladinithi, "Flow Bindings in Mobile IPv6 and Network Mobility (NEMO) Basic Support", RFC 6089, January 2011.
- [RFC6982] Sheffer, Y. and A. Farrel, "Improving Awareness of Running Code: The Implementation Status Section", RFC 6982, July 2013.
- [SMI] IANA, "PRIVATE ENTERPRISE NUMBERS", SMI Network Management Private Enterprise Codes, February 2011.
- [TS22.115] 3GPP, "Technical Specification Group Services and System Aspects, Service aspects; Charging and Billing", 2002.
- [TS23.203] 3GPP, "Policy and charging control architecture", 2013.
- [TS23.402] 3GPP, "Architecture enhancements for non-3GPP accesses", 2010.

Appendix A. Information when implementing 3GPP QoS in IP transport network

A.1. Mapping tables

Mapping between 3GPP QCI values and DSCP is defined in [GSMA.IR.34] as follows.

QCI	Traffic Class	DiffServ Per-Hop-Behavior	DSCP
1	Conversational	EF	101110
2	Conversational	EF	101110
3	Conversational	EF	101110
4	Streaming	AF41	100010
5	Interactive	AF31	011010
6	Interactive	AF32	011100
7	Interactive	AF21	010010
8	Interactive	AF11	001010
9	Background	BE	000000

Figure 7: QCI/DSCP Mapping Table

Mapping between QoS attributes defined in this document and 3GPP QoS parameters is as follows.

Section	PMIPv6 QoS Attribute	3GPP QoS Parameter
4.2.1	Per-MN-Agg-Max-DL-Bit-Rate	UE AMBR-DL
4.2.2	Per-MN-Agg-Max-UL-Bit-Rate	UE AMBR-UL
4.2.3	Per-Session-Agg-Max-DL-Bit-Rate Flags: (S=1, E=1)	APN AMBR-DL
4.2.4	Per-Session-Agg-Max-UL-Bit-Rate Flags: (S=1, E=1)	APN AMBR-UL
4.2.5	Allocation-Retention-Priority	ARP
4.2.6	Aggregate-Max-DL-Bit-Rate	MBR-DL
4.2.7	Aggregate-Max-UL-Bit-Rate	MBR-UL
4.2.8	Guaranteed-DL-Bit-Rate	GBR-DL
4.2.9	Guaranteed-UL-Bit-Rate	GBR-UL
4.2.10	QoS-Traffic-Selector	TFT

Figure 8: QoS attributes and 3GPP QoS parameters Mapping Table

A.2. Use cases and protocol operations

This subsections provide example message flow charts for scenarios where the QoS option extensions will apply as described in (Section 6.1), to the protocol operation for QoS rules establishment as shown in Appendix A.2.1 and Appendix A.2.2, and modification as show in Appendix A.2.3.

A.2.1. Handover of existing QoS rules

In Figure 9, the MN is first connected to the LTE network, and having a multimedia session such as a video call with appropriate QoS parameters set by the Policy Control Function. Then, the MN discovers a Wi-Fi AP (e.g., at home or in a cafe) and switches to it provided that Wi-Fi access has a higher priority when available. Not only is the session continued, but also the QoS is maintained after moving to the Wi-Fi access. In order for that to happen, the LMA delivers the QoS parameters according to the bearer type on the 3GPP

access to the MAG via the PMIPv6 signaling with the QoS option (OC=ALLOCATE, SR-ID, QoS attributes, etc.). The equivalent QoS treatment is provided by the Wi-Fi AP toward the MN on the Wi-Fi link.

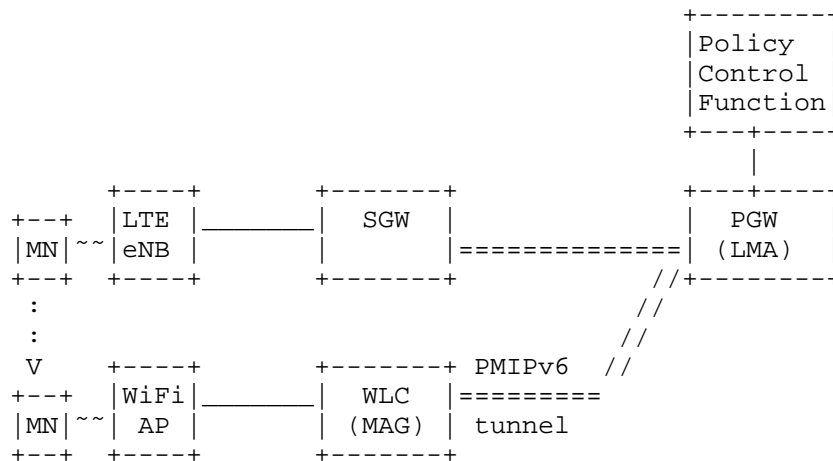
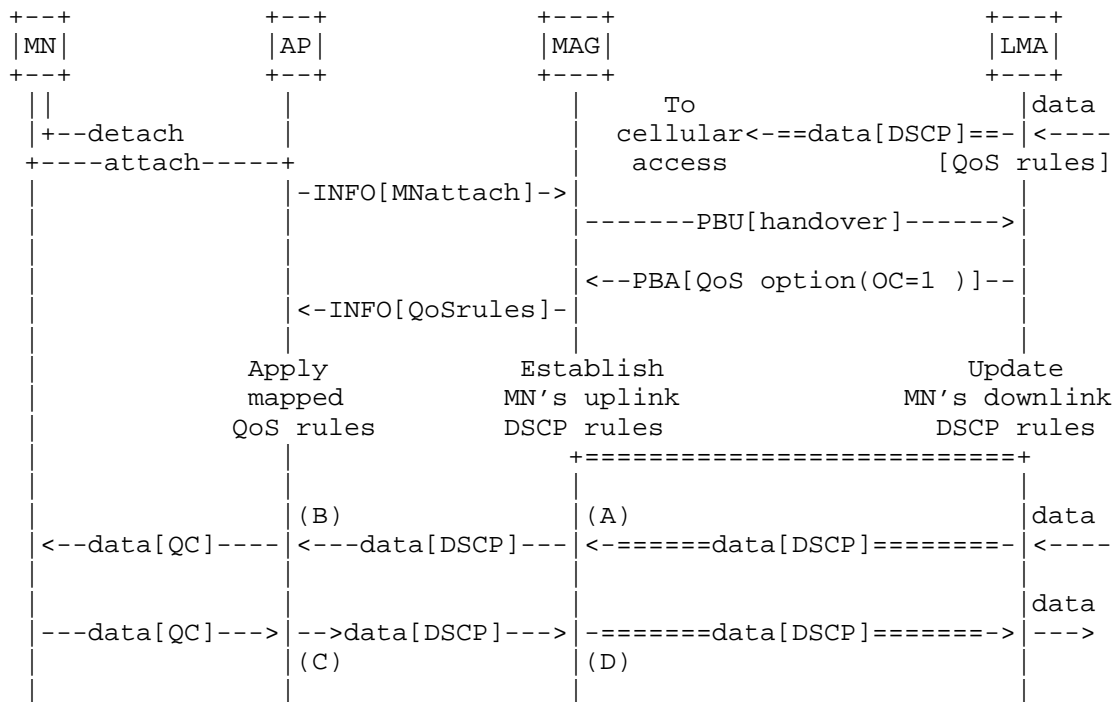


Figure 9: Handover Scenario (from LTE to WLAN)

Figure 10 shows an example of how the QoS rules can be conveyed and enforced between the LMA and MN in the case of handover from 3GPP access to WLAN access.



- (A): Apply DSCP at link to AP
 (B): Enforce mapped QoS rules to access technology
 (C): Map MN-indicated QoS Class (QC) to DSCP on the AP-MAG link, or validate MN-indicated QC and apply DSCP on the AP-MAG link according to QoS rules
 (D): Validate received DSCP and apply DSCP according to QoS rules

Figure 10: Handover of QoS rules

A.2.2. Establishment of QoS rules

A single operator has deployed both a fixed access network and a mobile access network. In this scenario, the operator may wish a harmonized QoS management on both accesses, but the fixed access network does not implement a QoS control framework. So, the operator chooses to rely on the 3GPP policy control function, which is a standard framework to provide a QoS control, and to enforce the 3GPP QoS policy on the Wi-Fi Access network. The PMIP interface is used to realize this QoS policy provisioning.

The use-case is depicted on Figure 11. The MN first attaches to the Wi-Fi network. During the attachment process, the LMA, which may

communicate with Policy Control Function (using procedures outside the scope of this document), provides the QoS parameters to the MAG via the QoS option (OC=ALLOCATE) in the PMIP signaling (i.e. PBA). Subsequently, an application on the MN may trigger the request for alternative QoS resources, e.g., by use of the WMM-API. The MN may request traffic resources be reserved using L2 signaling, e.g., sending an ADDTS message [IEEE802.11-2012]. The request is relayed to the MAG which includes the QoS parameters in the QoS option (OC=ALLOCATE) on the PMIP signaling (i.e. the PBU initiated upon flow creation). The LMA, in co-ordination with the PCF, can then authorize the enforcement of such QoS policy. Then, the QoS parameters are provided to the MAG via the QoS option (OC=ALLOCATE, SR-ID, QoS attributes, etc.) in the PMIP signaling and the equivalent QoS treatment is provided towards the MN on the Wi-Fi link.

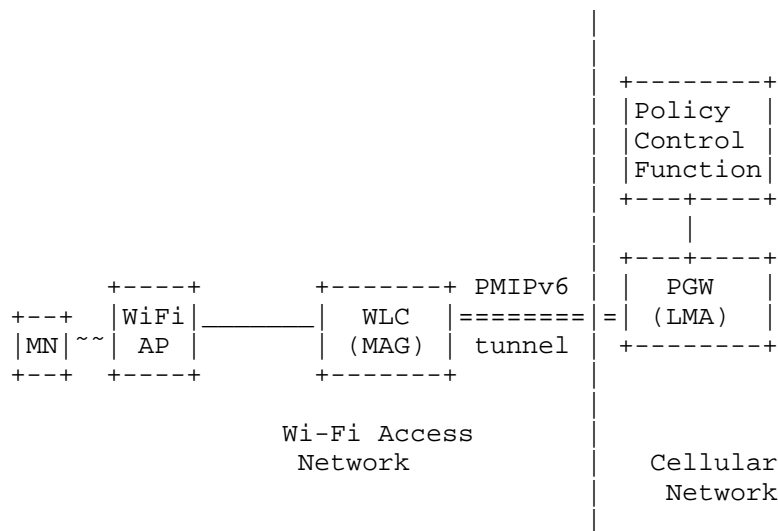



Figure 11: QoS policy provisioning

Figure 12 shows an example of how the QoS rules can be conveyed and enforced between the LMA and MN in the case of initial attachment to WLAN access.



(E): AP may enforce uplink QoS rules according to priority class set by the MN

(F): MAG can enforce a default QoS class until local mobility anchor has classified the new flow (notified with PBA) or mobile access gateway classifies new flow and proposes the associated QoS class to the local mobility anchor for validation (proposed with PBU, notification of validation result with PBA)

Figure 12: Adding new QoS Service Request for MN initiated flow

A.2.3. Dynamic Update to QoS Policy

A mobile node is attached to the WLAN access and has obtained QoS parameters from the LMA for that mobility session. Having obtained the QoS parameters, a new application, e.g. IMS application, gets launched on the mobile node that requires certain QoS support.

The application on the mobile node initiates the communications via a dedicated network function (e.g. IMS Call Session Control Function).

Once the communication is established, the application network function notifies the PCF about the new IP flow. The PCF function in turn notifies the LMA about the needed QoS parameters identifying the IP flow and QoS parameters. LMA sends an Update Notification message [RFC7077] to the MAG with the Notification Reason value set to "QOS_SERVICE_REQUEST". The MAG, on receiving the Update Notification message, completes the PBU/PBA signaling for obtaining the new QoS parameters via the QoS options (OC=MODIFY, SR-ID, QoS attributes, etc.). The MAG provisions the newly obtained QoS parameters on the access network to ensure the newly established IP flow gets its requested network resources.

Upon termination of the established IP flow, the application network function again notifies the PCF function for removing the established QoS parameters. The PCF notifies the LMA for withdrawing the QoS resources established for that voice flow. The LMA sends an Update Notification message to the MAG with the "Notification Reason" value set to "FORCE-REREGISTRATION". The MAG on receiving this message sends an Update Notification Acknowledgement and completes the PBU/PBA signaling for removing the existing QoS rules (OC=DE-ALLOCATE, SR-ID). The MAG then removes the QoS parameters from the corresponding IP flow and releases the dedicated network resources on the access network.

Appendix B. Information when implementing PMIP based QoS support with IEEE 802.11e

This section shows, as an example, the end-to-end QoS management with a 802.11e capable WLAN access link and a PMIP based QoS support.

The 802.11e, or Wi-Fi Multimedia (WMM), specification provides prioritization of packets for four types of traffic, or access categories (AC):

Voice (AC_VO): Very high priority queue with minimum delay. Time-sensitive data such as VoIP and streaming mode are automatically sent to this queue.

Video (AC_VI): High priority queue with low delay. Time-sensitive video data is automatically sent to this queue.

Best effort (AC_BE): Medium priority queue with medium throughput and delay. Most traditional IP data is sent to this queue.

Background (AC_BK): Lowest priority queue with high throughput. Bulk data that requires maximum throughput but is not time-sensitive (for example, FTP data) is sent to the queue.

The access point uses the 802.11e indicator to prioritize traffic on the WLAN interface. On the wired side, the access point uses the 802.1p priority tag and DiffServ code point (DSCP). To allow consistent QoS management on both wireless and wired interfaces, the access point relies on the 802.11e specification which define mapping between the 802.11e access categories and the IEEE 802.1D priority (802.1p tag). The end-to-end QoS architecture is depicted on Figure 13 and the 802.11e/802.1D priority mapping is reminded in the following table:

802.1e AC	802.1D priority
AC_VO	7,6
AC_VI	5,4
AC_BE	0,3
AC_BK	2,1

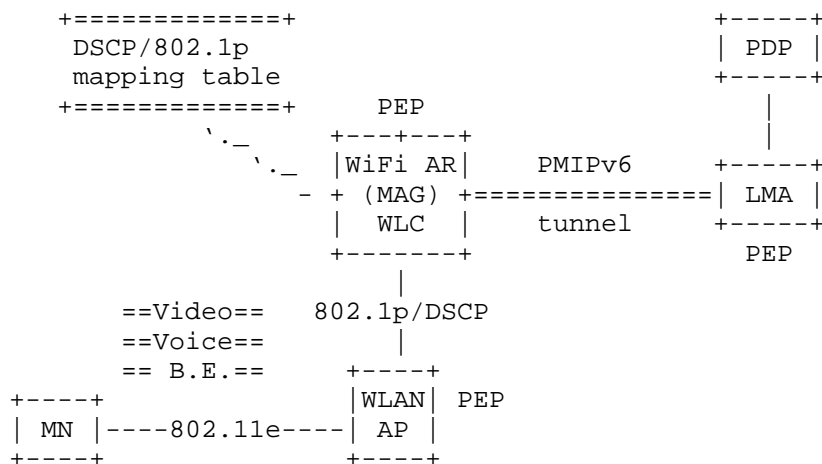


Figure 13: End-to-end QoS management with 802.11e

When receiving a packet from the MN, the AP checks whether the frame contains 802.11e markings in the L2 header. If not, the AP checks the DSCP field. If the uplink packet contains the 802.11e marking, the access point maps the access categories to the corresponding 802.1D priority as per the table above. If the frame does not contain 802.11e marking, the access point examines the DSCP field. If DSCP is present, the AP maps DSCP values to a 802.1p value (i.e 802.1D priority). This mapping is not standardized and may differ between operator; a mapping example given in the following table.

Type of traffic	802.1p	DSCP value
Network Control	7	56
Voice	6	46 (EF)
Video	5	34 (AF 41)
voice control	4	26 (AF 31)
Background Gold	2	18 (AF 21)
Background Silver	1	10 (AF 11)
Best effort	0,3	0 (BE)

The access point prioritizes ingress traffic on the Ethernet port

based on the 802.1p tag or the DSCP value. If 802.1p priority tag is not present, the access point checks the DSCP/802.1p mapping table. The next step is to map the 802.1p priority to the appropriate egress queue. When 802.11e support is enabled on the wireless link, the access point uses the IEEE standardized 802.1p/802.11e correspondence table to map the traffic to the appropriate hardware queues.

When the 802.11e capable client sends traffic to the AP, it usually marks packets with a DSCP value. In that case, the MAG/LMA can come into play for QoS renegotiation and call flows depicted in Appendix A apply. Sometimes, when communication is initiated on the WLAN access, the application does not mark upstream packets. If the uplink packet does not contain any QoS marking, the AP/MAG could determine the DSCP field according to traffic selectors received from the LMA. Figure 14 gives the call flow corresponding to that use-case and shows where QoS tags mapping does come into play. The main steps are as follows:

(A): during MN attachment process, the MAG fetches QoS policies from the LMA. After this step, both MAG and LMA are provisioned with QoS policies.

(B): the MN starts a new IP communication without making IP packets with DSCP tags. The MAG uses the traffic selector to determine the DSCP value, then it marks the IP packet and forwards within the PMIP tunnel.

(C): the LMA checks the DSCP value with respect to the traffic selector. If the QoS policies is valid, the LMA forwards the packet without renegotiating the QoS rules.

(D): when receiving a marked packet, the MAG, the AP and the MN use 802.11e (or WMM), 802.1p tags and DSCP values to prioritize the traffic.




Figure 14: Prioritization of a flow created on the WLAN access

Authors' Addresses

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

Email: liebsch@neclab.eu

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

Email: pierrick.seite@orange.com

Hidetoshi Yokota
KDDI Lab
2-1-15 Ohara
Saitama, Fujimino 356-8502
Japan

Email: yokota@kddilabs.jp

Jouni Korhonen
Broadcom Communications
Porkkalankatu 24
Helsinki FIN-00180
Finland

Email: jouni.nospam@gmail.com

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

Email: sgundave@cisco.com

NETEXT Working Group
Internet-Draft
Updates: 5213 (if approved)
Intended status: Standards Track
Expires: September 19, 2016

CJ. Bernardos, Ed.
UC3M
March 18, 2016

Proxy Mobile IPv6 Extensions to Support Flow Mobility
draft-ietf-netext-pmipv6-flowmob-18

Abstract

Proxy Mobile IPv6 allows a mobile node to connect to the same Proxy Mobile IPv6 domain through different interfaces. This document describes extensions to the Proxy Mobile IPv6 protocol that are required to support network based flow mobility over multiple physical interfaces.

This document updates RFC 5213. The extensions described in this document consist of the operations performed by the local mobility anchor and the mobile access gateway to manage the prefixes assigned to the different interfaces of the mobile node, as well as how the forwarding policies are handled by the network to ensure consistent flow mobility management.

Requirements Language

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

Status of This Memo

This Internet-Draft is submitted in full conformance with the provisions of BCP 78 and BCP 79.

Internet-Drafts are working documents of the Internet Engineering Task Force (IETF). Note that other groups may also distribute working documents as Internet-Drafts. The list of current Internet-Drafts is at <http://datatracker.ietf.org/drafts/current/>.

Internet-Drafts are draft documents valid for a maximum of six months and may be updated, replaced, or obsoleted by other documents at any time. It is inappropriate to use Internet-Drafts as reference material or to cite them other than as "work in progress."

This Internet-Draft will expire on September 19, 2016.

Copyright Notice

Copyright (c) 2016 IETF Trust and the persons identified as the document authors. All rights reserved.

This document is subject to BCP 78 and the IETF Trust's Legal Provisions Relating to IETF Documents (<http://trustee.ietf.org/license-info>) in effect on the date of publication of this document. Please review these documents carefully, as they describe your rights and restrictions with respect to this document. Code Components extracted from this document must include Simplified BSD License text as described in Section 4.e of the Trust Legal Provisions and are provided without warranty as described in the Simplified BSD License.

Table of Contents

1. Introduction	3
2. Terminology	3
3. Overview of the PMIPv6 flow mobility extensions	4
3.1. Use case scenarios	4
3.2. Basic Operation	5
3.2.1. MN sharing a common set of prefixes on all MAGs	5
3.2.2. MN with different sets of prefixes on each MAG	9
3.3. Use of PBU/PBA signaling	11
3.4. Use of flow-level information	12
4. Message Formats	12
4.1. Home Network Prefix	12
4.2. Flow Mobility Initiate (FMI)	13
4.3. Flow Mobility Acknowledgement (FMA)	14
5. Conceptual Data Structures	14
5.1. Multiple Proxy Care-of Address Registration	14
5.2. Flow Mobility Cache	15
6. Mobile Node considerations	16
7. IANA Considerations	16
8. Security Considerations	17
9. Authors	17
10. Acknowledgments	18
11. References	18
11.1. Normative References	18
11.2. Informative References	19
Author's Address	19

1. Introduction

Proxy Mobile IPv6 (PMIPv6), specified in [RFC5213], provides network based mobility management to hosts connecting to a PMIPv6 domain. PMIPv6 introduces two new functional entities, the Local Mobility Anchor (LMA) and the Mobile Access Gateway (MAG). The MAG is the entity detecting the Mobile Node's (MN) attachment and providing IP connectivity. The LMA is the entity assigning one or more Home Network Prefixes (HNP) to the MN and is the topological anchor for all traffic belonging to the MN.

PMIPv6 allows a mobile node to connect to the same PMIPv6 domain through different interfaces. This document specifies protocol extensions to Proxy Mobile IPv6 between the local mobility anchor and mobile access gateways to enable "flow mobility" and hence distribute specific traffic flows on different physical interfaces. It is assumed that the mobile node IP layer interface can simultaneously and/or sequentially attach to multiple MAGs, possibly over multiple media. One form to achieve this multiple attachment is described in [I-D.ietf-netext-logical-interface-support], which allows the mobile node supporting traffic flows on different physical interfaces regardless of the assigned prefixes on those physical interfaces. Another alternative is to configure the IP stack of the mobile node to behave according to the weak host model [RFC1122].

In particular, this document specifies how to enable "flow mobility" in the PMIPv6 network (i.e., local mobility anchors and mobile access gateways). In order to do so, two main operations are required: i) proper prefix management by the PMIPv6 network, and, ii) consistent flow forwarding policies. This memo analyzes different potential use case scenarios, involving different prefix assignment requirements, and therefore different PMIPv6 network extensions to enable "flow mobility".

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 [RFC2119].

The following terms used in this document are defined in the Proxy Mobile IPv6 [RFC5213]:

Local Mobility Agent (LMA).

Mobile Access Gateway (MAG).

Proxy Mobile IPv6 Domain (PMIPv6-Domain).

LMA Address (LMAA).

Proxy Care-of Address (Proxy-CoA).

Home Network Prefix (HNP).

The following terms used in this document are defined in the Multiple Care-of Addresses Registration [RFC5648] and Flow Bindings in Mobile IPv6 and Network Mobility (NEMO) Basic Support [RFC6089]:

Binding Identification Number (BID).

Flow Identifier (FID).

Traffic Selector (TS).

The following terms are defined and used in this document:

FMI (Flow Mobility Initiate). Message sent by the LMA to the MAG conveying the information required to enable flow mobility in a PMIPv6-Domain.

FMA (Flow Mobility Acknowledgement). Message sent by the MAG in reply to an FMI message.

FMC (Flow Mobility Cache). Conceptual data structure to support the flow mobility management operations described in this document.

3. Overview of the PMIPv6 flow mobility extensions

3.1. Use case scenarios

In contrast to a typical handover where connectivity to a physical medium is relinquished and then re-established, flow mobility assumes a mobile node can have simultaneous access to more than one network. In this specification, it is assumed that the local mobility anchor is aware of the mobile node's capabilities to have simultaneous access to both access networks and it can handle the same or a different set of prefixes on each access. How this is done is outside the scope of this specification.

There are different flow mobility scenarios. In some of them the mobile node might share a common set of prefixes among all its physical interfaces, whereas in others the mobile node might have a different subset of prefixes configured on each of the physical interfaces. The different scenarios are the following:

1. At the time of a new network attachment, the MN obtains the same prefix or the same set of prefixes as already assigned to an existing session. This is not the default behavior with basic PMIPv6 [RFC5213], and the LMA needs to be able to provide the same assignment even for the simultaneous attachment (as opposed to the handover scenario only).
2. At the time of a new network attachment, the MN obtains a new prefix or a new set of prefixes for the new session. This is the default behavior with basic PMIPv6 [RFC5213].

A combination of the two above-mentioned scenarios is also possible. At the time of a new network attachment, the MN obtains a combination of prefix(es) in use and new prefix(es). This is a hybrid of the two scenarios described before. The local policy determines whether the new prefix is exclusive to the new attachment or it can be assigned to an existing attachment as well.

The operational description of how to enable flow mobility in each of these scenarios is provided in Section 3.2.1 and Section 3.2.2.

The extensions described in this document support all the aforementioned scenarios.

3.2. Basic Operation

This section describes how the PMIPv6 extensions described in this document enable flow mobility support.

Both the mobile node and the local mobility anchor MUST have local policies in place to ensure that packets are forwarded coherently for unidirectional and bidirectional communications. The details about how this consistency is ensured are out of the scope of this document. Either the MN or the LMA can initiate IP flow mobility. If the MN makes the flow mobility decision, then the LMA follows that decision and updates its forwarding state accordingly. The network can also trigger mobility on the MN side via out-of-band mechanisms (e.g., 3GPP/ANDSF sends updated routing policies to the MN). In a given scenario and mobile node, the decision on IP flow mobility MUST be taken either by the MN or the LMA, but MUST NOT be taken by both.

3.2.1. MN sharing a common set of prefixes on all MAGs

This scenario corresponds to the first use case scenario described in Section 3.1. Extensions to basic PMIPv6 [RFC5213] signaling at the time of a new attachment are needed to ensure that the same prefix (or set of prefixes) is assigned to all the interfaces of the same mobile node that are simultaneously attached. Subsequently, no

further signaling is necessary between the local mobility anchor and the mobile access gateway and flows are forwarded according to policy rules on the local mobility anchor and the mobile node.

If the local mobility anchor assigns a common prefix (or set of prefixes) to the different physical interfaces attached to the domain, then every MAG already has all the routing knowledge required to forward uplink or downlink packets after the PBU/PBA registration for each MAG, and the local mobility anchor does not need to send any kind of signaling in order to move flows across the different physical interfaces (because moving flows is a local decision of the LMA). Optionally, signaling MAY be exchanged in case the MAG needs to know about flow level information (e.g., to link flows with proper QoS paths and/or inform the mobile node) [RFC7222].

The local mobility anchor needs to know when to assign the same set of prefixes to all the different physical interfaces of the mobile node. This can be achieved by different means, such as policy configuration, default policies, etc. In this document a new Handoff Indicator (HI) value ("Attachment over a new interface sharing prefixes", value {IANA-0}) is defined, to allow the mobile access gateway to indicate to the local mobility anchor that the same set of prefixes MUST be assigned to the mobile node. The considerations of Section 5.4.1 of [RFC5213] are updated by this specification as follows:

- o If there is at least one Home Network Prefix option present in the request with a NON_ZERO prefix value, there exists a Binding Cache entry (with all home network prefixes in the Binding Cache entry matching the prefix values of all Home Network Prefix options of the received Proxy Binding Update message), and the entry matches the mobile node identifier in the Mobile Node Identifier option of the received Proxy Binding Update message, and the value of the Handoff Indicator of the received Proxy Binding Update is equal to "Attachment over a new interface sharing prefixes".
 1. If there is an MN-LL-Identifier Option present in the request and the Binding Cache entry matches the Access Technology Type (ATT), and MN-LL-Identifier, the request MUST be considered as a request for updating that Binding Cache entry.
 2. If there is an MN-LL-Identifier Option present in the request and the Binding Cache entry does not match the Access Technology Type (ATT), and MN-LL-Identifier, the request MUST be considered as a request for creating a new mobility session sharing the same set of home network prefixes assigned to the existing Binding Cache entry found.

3. If there is not an MN-LL-Identifier Option present in the request, the request MUST be considered as a request for creating a new mobility session sharing the same set of home network prefixes assigned to the existing Binding Cache entry found.




Figure 1: Shared prefix across physical interfaces scenario

Next, an example of how flow mobility works in this case is shown. In Figure 1, a mobile node (MN1) has two different physical interfaces (if1 of access technology type ATT1, and if2 of access technology type ATT2). Each physical interface is attached to a different mobile access gateway, both of them controlled by the same local mobility anchor. Both physical interfaces are assigned the same prefix (pref1) upon attachment to the MAGs. If the IP layer at the mobile node shows one single logical interface (e.g., as described in [I-D.ietf-netext-logical-interface-support]), then the mobile node has one single IPv6 address configured at the IP layer: pref1::mn1. Otherwise, per interface IPv6 addresses (e.g., pref1::if1 and pref1::if2) would be configured; each address MUST be valid on every interface. We assume the first case in the following example (and in the rest of this document). Initially, flow X goes through MAG1 and flow Y through MAG2. At a certain point, flow Y can be moved to also go through MAG1. Figure 2 shows the scenario in which no flow-level information needs to be exchanged, so there is no

signaling between the local mobility anchor and the mobile access gateways.

Note that if different IPv6 addresses are configured at the IP layer, IP session continuity is still possible (for each of the configured IP addresses). This is achieved by the network delivering packets destined to a particular IP address of the mobile node to the right MN's physical interface where the flow is selected to be moved, and the MN also selecting the same interface when sending traffic back up link.

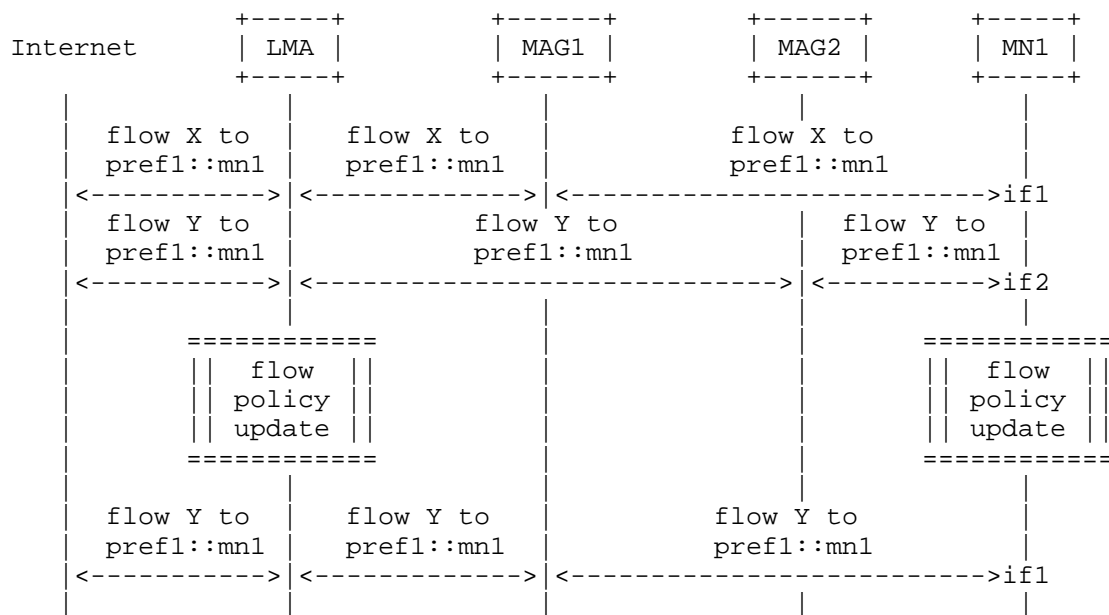


Figure 2: Flow mobility message sequence with common set of prefixes

Figure 3 shows the state of the different network entities after moving flow Y in the previous example. This document re-uses some of the terminology and mechanisms of the flow bindings and multiple care-of address registration specifications. Note that, in this case the BIDs shown in the figure are assigned locally by the LMA, since there is no signaling required in this scenario. In any case, alternative implementations of flow routing at the LMA MAY be used, as it does not impact on the operation of the solution in this case.




Figure 3: Data structures with common set of prefixes

3.2.2. MN with different sets of prefixes on each MAG

A different flow mobility scenario happens when the local mobility anchor assigns different sets of prefixes to physical interfaces of the same mobile node. This covers the second case, or a combination of scenarios, described in Section 3.1. In this case, additional signaling is required between the local mobility anchor and the mobile access gateway to enable relocating flows between the different attachments, so the MAGs are aware of the prefixes for which the MN is going to receive traffic, and local routing entries are configured accordingly.

In this case, signaling is required when a flow is to be moved from its original interface to a new one. Since the local mobility anchor cannot send a PBA message which has not been triggered in response to a received PBU message, the solution defined in this specification makes use of two mobility messages: Flow Mobility Indication and Flow Mobility Acknowledgement, which actually use the format of the Update Notifications for Proxy Mobile IPv6 defined in [RFC7077]. The trigger for the flow movement can be on the mobile node (e.g., by using layer-2 signaling with the MAG) or on the network (e.g., based

on congestion and measurements) which then notifies the MN for the final IP flow mobility decision (as stated in section 3.1). Policy management functions (e.g., 3GPP/ANDSF) can be used for that purpose, however, how the network notifies the MN is out of the scope of this document.

If the flow is being moved from its default path (which is determined by the destination prefix) to a different one, the local mobility anchor constructs a Flow Mobility Indication (FMI) message. This message includes a Home Network Prefix option for each of the prefixes that are requested to be provided with flow mobility support on the new MAG (note that these prefixes are not anchored by the target MAG, and therefore the MAG MUST NOT advertise them on the MAG-MN link), with the off-link bit (L) set to one. This message MUST be sent to the new target mobile access gateway, i.e. the one selected to be used in the forwarding of the flow. The MAG replies with a Flow Mobility Acknowledgement (FMA). The message sequence is shown in Figure 4.




Figure 4: Flow mobility message sequence when the LMA assigns different sets of prefixes per physical interface

The state in the network after moving a flow, for the case the LMA assigns a different set of prefixes is shown in Figure 5.

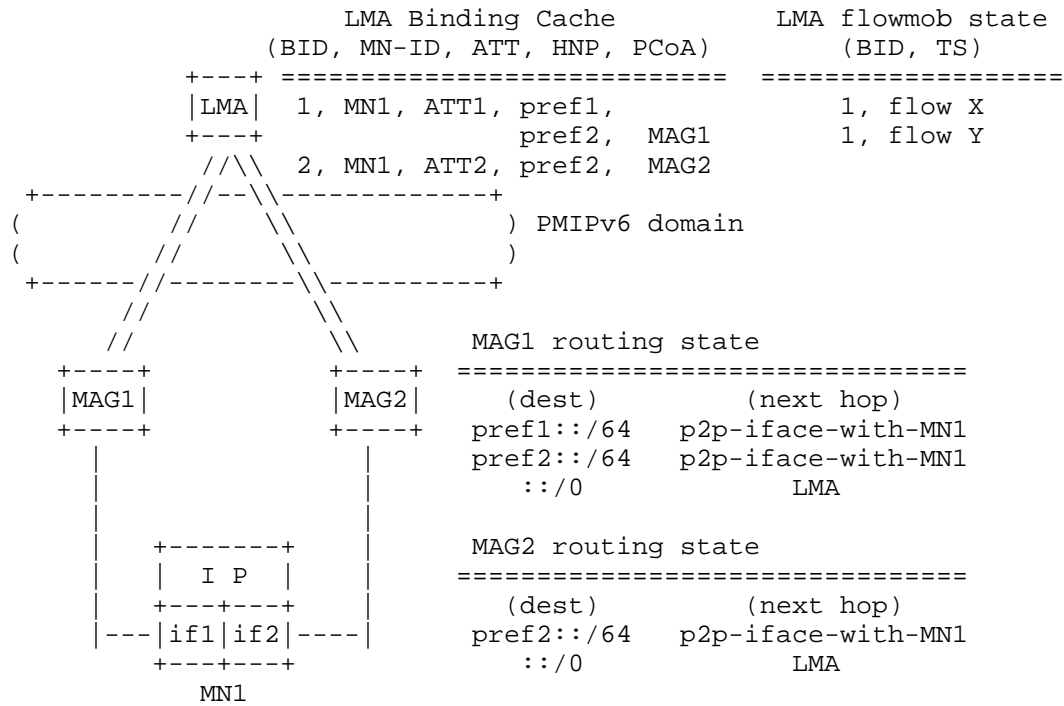


Figure 5: Data structures when the LMA assigns a different set of prefixes

3.3. Use of PBU/PBA signaling

This specification introduces the FMI/FMA signaling so the LMA can exchange with the MAG information required to enable flow mobility without waiting for receiving a PBU. There are however scenarios in which the trigger for flow mobility might be related to a new MN's interface attachment. In this case, the PBA sent in response to the PBU received from the new MAG can convey the same signaling that the FMI does. In this case the LMA MUST include in the PBA a Home Network Prefix option for each of the prefixes that are requested to be provided with flow mobility support on the new MAG with the off-link bit (L) set to one.

3.4. Use of flow-level information

This specification does not mandate flow-level information to be exchanged between the LMA and the MAG to provide flow mobility support. It only requires the LMA to keep flow-level state (Section 5.2). However, there are scenarios in which the MAG might need to know which flow(s) is/are coming within a prefix that has been moved, to link it/them to proper QoS path(s) and optionally inform the MN about it. This section describes the extensions used to include flow-level information in the signaling defined between the LMA and the MAG.

This specification re-uses some of the mobility extensions and message formats defined in [RFC5648] and [RFC6089], namely the Flow Identification Mobility Option and the Flow Mobility Sub-Options.

In case the LMA wants to convey flow-level information to the MAG, it MUST include in the FMI (or the PBA) a Flow Identification Mobility Option for all the flows that the MAG needs to be aware with flow granularity. Each Flow Identification Option MUST include a Traffic Selector Sub-Option including such flow-level information.

To remove a flow binding state at the MAG, the LMA simply sends a FMI (or PBA if it is in response to a PBU) message that includes flow identification options for all the flows that need to be refreshed, modified, or added, and simply omits those that need to be removed.

Note that even if a common set of prefixes is used, providing the MAG with flow-level information requires signaling to be exchanged in this case between the LMA and the MAG. This is done sending a FMI message (or a PBA if it is sent in response to a PBU).

4. Message Formats


This section defines modifications to the Proxy Mobile IPv6 [RFC5213] protocol messages.

This specification requires implementation of UPN [RFC7077] and UPA [RFC7077] messages with the specific Notification Reason and Status Code values as defined by this document. This document does not require implementation of any other aspects of [RFC7077].

4.1. Home Network Prefix

A new flag (L) is included in the Home Network Prefix option to indicate to the Mobile Access Gateway whether the conveyed prefix has to be hosted on-link or not on the point-to-point interface with the mobile node. A prefix is hosted off-link for the flow mobility

purposes defined in this document. The rest of the Home Network Prefix option format remains the same as defined in [RFC5213].



Off-link Home Network Prefix Flag (L):

The Off-link Home Network Prefix Flag is set to indicate to the Mobile Access Gateway that the home network prefix conveyed in the option is not to be hosted on-link, but has to be considered for flow mobility purposes and therefore added to the Mobile Access Gateway routing table. If the flag is set to 0, the Mobile Access Gateway assumes that the home network prefix has to be hosted on-link.

4.2. Flow Mobility Initiate (FMI)

The FMI message used in this specification is the Update Notification (UPN) message specified in [RFC7077]. The message format, transport and security consideration are as specified in [RFC7077]. The format of the message is specified in Section 4.1 of [RFC7077]. This specification does not modify the UPN message, however, it defines the following new notification reason value for use in this specification:

Notification Reason:

{IANA-1} - FLOW-MOBILITY. Request to add/refresh the prefix(es) conveyed in the Home Network Prefix options included in the message to the set of prefixes for which flow mobility is provided.

The Mobility Options field of an FMI MUST contain the MN-ID, followed by one or more Home Network Prefixes options. Prefixes for which flow mobility was provided that are not present in the message MUST be removed from the set of flow mobility enabled prefixes.

4.3. Flow Mobility Acknowledgement (FMA)

The FMA message used in this specification is the Update Notification Ack (UPA) message specified in Section 4.2 of [RFC7077]. The message format, transport and security consideration are as specified in [RFC7077]. The format of the message is specified in Section 4.2 of [RFC7077]. This specification does not modify the UPA message, however, it defines the following new status code values for use in this specification:

Status Code:

0: Success.

{IANA-2}: Reason unspecified.

{IANA-3}: MN not attached.

When Status code is 0, the Mobility Options field of an FMA MUST contain the MN-ID, followed by one or more Home Network Prefixes options.

5. Conceptual Data Structures

This section summarizes the extensions to Proxy Mobile IPv6 that are necessary to manage flow mobility.

5.1. Multiple Proxy Care-of Address Registration

The binding cache structure of the local mobility anchor is extended to allow multiple proxy care-of address (Proxy-CoA) registrations, and support the mobile node use the same address (prefix) beyond a single interface and mobile access gateway. The LMA maintains multiple binding cache entries for an MN. The number of binding cache entries for a mobile node is equal to the number of the MN's interfaces attached to any MAGs.

This specification re-uses the extensions defined in [RFC5648] to manage multiple registrations, but in the context of Proxy Mobile IPv6. The binding cache is therefore extended to include more than one proxy care-of address and to associate each of them with a binding identifier (BID). Note that the BID is a local identifier, assigned and used by the local mobility anchor to identify which entry of the flow mobility cache is used to decide how to route a given flow.

BID-PRI	BID	MN-ID	ATT	HNP(s)	Proxy-CoA
20	1	MN1	WiFi	HNP1,HNP2	IP1 (MAG1)
30	2	MN1	3GPP	HNP1,HNP3	IP2 (MAG2)

Figure 6: Extended Binding Cache

Figure 6 shows an example of extended binding cache, containing two binding cache entries (BCEs) of a mobile node MN1 attached to the network using two different access technologies. Both of the two attachments share the same prefix (HNP1) and are bound to two different Proxy-CoAs (two MAGs).

5.2. Flow Mobility Cache

Each local mobility anchor MUST maintain a flow mobility cache (FMC) as shown in Figure 7. The flow mobility cache is a conceptual list of entries that is separate from the binding cache. This conceptual list contains an entry for each of the registered flows. This specification re-uses the format of the flow binding list defined in [RFC6089]. Each entry includes the following fields:

- o Flow Identifier Priority (FID-PRI).
- o Flow Identifier (FID).
- o Traffic Selector (TS).
- o Binding Identifier (BID).
- o Action.
- o Active/Inactive.

FID-PRI	FID	TS	BIDs	Action	A/I
10	2	TCP	1	Forward	Active
20	4	UDP	1,2	Forward	Inactive

Figure 7: Flow Mobility Cache

The BID field contains the identifier of the binding cache entry which packets matching the flow information described in the TS field

will be forwarded to. When a flow is decided to be moved, the affected BID(s) of the table are updated.

Similar to flow binding described in [RFC6089], each entry of the flow mobility cache points to a specific binding cache entry identifier (BID). When a flow is moved, the local mobility anchor simply updates the pointer of the flow binding entry with the BID of the interface to which the flow will be moved. The traffic selector (TS) in flow binding table is defined as in [RFC6088]. TS is used to classify the packets of flows based on specific parameters such as service type, source and destination address, etc. The packets matching with the same TS will be applied the same forwarding policy. FID-PRI is the order of precedence to take action on the traffic. Action may be forward or drop. If a binding entry becomes 'Inactive' it does not affect data traffic. An entry becomes 'Inactive' only if all of the BIDs are de-registered.

The mobile access gateway MAY also maintain a similar data structure. In case no full flow mobility state is required at the MAG, the Binding Update List (BUL) data structure is enough and no extra conceptual data entries are needed. In case full per-flow state is required at the mobile access gateway, it SHOULD also maintain a flow mobility cache structure.

6. Mobile Node considerations

This specification assumes that the mobile node IP layer interface can simultaneously and/or sequentially attach to multiple MAGs, possibly over multiple media. The mobile node MUST be able to enforce uplink policies to select the right outgoing interface. One alternative to achieve this multiple attachment is described in [I-D.ietf-netext-logical-interface-support], which allows the mobile node supporting traffic flows on different physical interfaces regardless of the assigned prefixes on those physical interfaces. Another alternative is configuring the IP stack of the mobile node to behave according to the weak host model [RFC1122].

7. IANA Considerations

This specification establishes new assignments to the IANA mobility parameters registry:

- o Handoff Indicator Option type: the value {IANA-0} has to be assigned from the "Handoff Indicator Option type values" registry defined in <http://www.iana.org/assignments/mobility-parameters/mobility-parameters.xhtml#mobility-parameters-9>.

- o Update Notification Reason: the value ({IANA-1}) has to be assigned from the "Update Notification Reasons Registry" defined in <http://www.iana.org/assignments/mobility-parameters/mobility-parameters.xhtml#upn-reasons>.
- o Update Notification Acknowledgement Status: values ({IANA-2} and {IANA-3}) have to be assigned from the "Update Notification Acknowledgement Status Registry". Since {IANA-2} and {IANA-3} are used in error messages, their values have to be greater than 128 from the range defined in <http://www.iana.org/assignments/mobility-parameters/mobility-parameters.xhtml#upa-status>.

8. Security Considerations

The protocol signaling extensions defined in this document share the same security concerns of Proxy Mobile IPv6 [RFC5213] and do not pose any additional security threats to those already identified in [RFC5213] and [RFC7077].

The mobile access gateway and the local mobility anchor MUST use the IPsec security mechanism mandated by Proxy Mobile IPv6 [RFC5213] to secure the signaling described in this document.

9. Authors

This document reflects contributions from the following authors (in alphabetical order).

Kuntal Chowdhury

E-mail: kc@altiostar.com

Sri Gundavelli

E-mail: sgundave@cisco.com

Youn-Hee Han

E-mail: yhhan@kut.ac.kr

Yong-Geun Hong

E-mail: yonggeun.hong@gmail.com

Rajeev Koodli

E-mail: rajeevkoodli@google.com

Telemaco Melia

E-mail: telemaco.melia@googlemail.com

Frank Xia

E-mail: xiayangsong@huawei.com

10. Acknowledgments

The authors would like to thank Vijay Devarapalli, Mohana Dahamayanthi Jeyatharan, Kent Leung, Bruno Mongazon-Cazavet, Chan-Wah Ng, Behcet Sarikaya and Tran Minh Trung for their valuable contributions which helped generating this document.

The authors would also like to thank Juan-Carlos Zuniga, Pierrick Seite, Julien Laganier for all the useful discussions on this topic.

Finally, the authors would also like to thank Marco Liebsch, Juan-Carlos Zuniga, Dirk von Hugo, Fabio Giust and Daniel Corujo for their reviews of this document.

The work of Carlos J. Bernardos has been partially performed in the framework of the H2020-ICT-2014-2 project 5G NORMA.

11. References

11.1. 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>>.
- [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>>.
- [RFC5648] Wakikawa, R., Ed., Devarapalli, V., Tsirtsis, G., Ernst, T., and K. Nagami, "Multiple Care-of Addresses Registration", RFC 5648, DOI 10.17487/RFC5648, October 2009, <<http://www.rfc-editor.org/info/rfc5648>>.
- [RFC6088] Tsirtsis, G., Giarreta, G., Soliman, H., and N. Montavont, "Traffic Selectors for Flow Bindings", RFC 6088, DOI 10.17487/RFC6088, January 2011, <<http://www.rfc-editor.org/info/rfc6088>>.

- [RFC6089] Tsirtsis, G., Soliman, H., Montavont, N., Giaretta, G., and K. Kuladinithi, "Flow Bindings in Mobile IPv6 and Network Mobility (NEMO) Basic Support", RFC 6089, DOI 10.17487/RFC6089, January 2011, <<http://www.rfc-editor.org/info/rfc6089>>.
- [RFC7077] Krishnan, S., Gundavelli, S., Liebsch, M., Yokota, H., and J. Korhonen, "Update Notifications for Proxy Mobile IPv6", RFC 7077, DOI 10.17487/RFC7077, November 2013, <<http://www.rfc-editor.org/info/rfc7077>>.

11.2. Informative References

- [I-D.ietf-netext-logical-interface-support]
Melia, T. and S. Gundavelli, "Logical-interface Support for Multi-access enabled IP Hosts", draft-ietf-netext-logical-interface-support-13 (work in progress), February 2016.
- [RFC1122] Braden, R., Ed., "Requirements for Internet Hosts - Communication Layers", STD 3, RFC 1122, DOI 10.17487/RFC1122, October 1989, <<http://www.rfc-editor.org/info/rfc1122>>.
- [RFC7222] Liebsch, M., Seite, P., Yokota, H., Korhonen, J., and S. Gundavelli, "Quality-of-Service Option for Proxy Mobile IPv6", RFC 7222, DOI 10.17487/RFC7222, May 2014, <<http://www.rfc-editor.org/info/rfc7222>>.

Author's Address

Carlos J. Bernardos (editor)
Universidad Carlos III de Madrid
Av. Universidad, 30
Leganes, Madrid 28911
Spain

Phone: +34 91624 6236
Email: cjbc@it.uc3m.es
URI: <http://www.it.uc3m.es/cjbc/>

INTERNET-DRAFT
Intended Status: Informational
Expires: April 23, 2013

John Kaippallimalil
Huawei
October 20, 2012

Mapping PMIP Quality of Service in WiFi Network
draft-kaippallimalil-netext-pmip-qos-wifi-01

Abstract

This document proposes a model for mapping PMIP QoS parameters of a mobile network session to the corresponding connection at a WiFi Access Point. In congested network conditions, it is possible that a user's flows from the WiFi AP are metered and shaped at the WLC to match the bandwidth constraints or service priority of the user's subscription and PMIP QoS parameters. Applying similar QoS policing at the WiFi AP allows optimal use of scarce radio network resources. Currently, the WiFi AP does not have information on the MNs subscribed bandwidth, or relative priority of its flows or services for per user QoS handling at the WiFi AP. This document provides a model for mapping PMIP QoS to corresponding 802.11e QoS parameters.

Status of this Memo

This Internet-Draft is submitted to IETF in full conformance with the provisions of BCP 78 and BCP 79.

Internet-Drafts are working documents of the Internet Engineering Task Force (IETF), its areas, and its working groups. Note that other groups may also distribute working documents as Internet-Drafts.

Internet-Drafts are draft documents valid for a maximum of six months and may be updated, replaced, or obsoleted by other documents at any time. It is inappropriate to use Internet-Drafts as reference material or to cite them other than as "work in progress."

The list of current Internet-Drafts can be accessed at
<http://www.ietf.org/lid-abstracts.html>

The list of Internet-Draft Shadow Directories can be accessed at
<http://www.ietf.org/shadow.html>

Copyright and License Notice

Copyright (c) 2012 IETF Trust and the persons identified as the document authors. All rights reserved.

This document is subject to BCP 78 and the IETF Trust's Legal Provisions Relating to IETF Documents (<http://trustee.ietf.org/license-info>) in effect on the date of publication of this document. Please review these documents carefully, as they describe your rights and restrictions with respect to this document. Code Components extracted from this document must include Simplified BSD License text as described in Section 4.e of the Trust Legal Provisions and are provided without warranty as described in the Simplified BSD License.

Table of Contents

1. Introduction	3
1.1 Terminology	4
1.2 Definitions	4
1.3 Abbreviations	4
2. QoS Mechanisms	4
2.1 QoS in Mobile Networks	4
2.2 QoS in WiFi Networks	5
3. Policy Provisioning Architecture	5
4. Connections and QoS Mapping	7
4.1 Connection Model	7
4.2 PMIP - 802.11e QoS Configuration	8
5. Mapping Recommendations and Default Values	9
6. Next Steps	10
7. Security Considerations	11
8. IANA Considerations	11
9. References	11
9.1 Normative References	11
9.2 Informative References	11
Authors' Addresses	12

1 Introduction

This document provides a description of how the QoS profile for a PMIP session maps to QoS for the corresponding 802.11 connection segment of the MN (Mobile Node). When a mobile network user attaches via a WiFi access, the WLC (MAG) obtains a QoS profile for each PMIP session. [PMIP-QoS] proposes a mechanism by which QoS policy parameters in the mobile network are delivered from the LMA to the WLC (MAG) using PMIP QoS extensions. [PMIP-QoS] further describes how the DSCP value for the PMIP session is mapped to corresponding 802.1p value that may be used by IP backhaul network or WiFi APs to prioritize IP flows of a host (MN).

[PMIP-QoS] outlines a model in which the QoS in PMIP flows can be reflexively mapped to IP flows over 802.11 or backhaul network. The WiFi AP can infer the QoS priority associated with an IP flow based on the the DSCP value in the downstream packets of the PMIP flow, and apply the same priority to upstream packets of the flow. It should be noted that the WLC (MAG) uses DSCP priority as well as other parameters of the MN such as subscribed bandwidth and service priority in [PMIP-QoS] to police IP flows of the MN. In congested network conditions, it is possible that upstream flows from the WiFi AP are throttled by the WLC to match the bandwidth constraints or service priority. This will result in sub-optimal use of network resources. Currently, the WiFi AP does not have information on the MNs subscribed bandwidth, or relative priority of its flows or services. In addition to uneven policing between WiFi AP and WLC, when the WiFi network itself is congested, the MN subscribed bandwidth and service priority can be useful to schedule and use the radio network resources more effectively.

This proposal aims to provide the WiFi AP with per MN QoS profile to allow more effective overall use of network resources - both WiFi radio and IP backhaul. The QoS parameters needed are available to the WLC during MN authorization and establishment of the PMIP session with QoS extensions. Since an MN may establish tunneled IP flows, direct IP connections or offloaded connections, the relationship of PMIP QoS to 802.11e QoS is explained. It is possible that an MN (with a single 802.11 interface) has more than one PMIP session. The QoS policy for the MN may be applied by the AP to schedule and control WiFi radio network resources and upstream user flows in the IP backhaul network. If per session QoS policy is not available, the AP may be provisioned to apply QoS based on the subscribed QoS values obtained during 3GPP user authorization.

In order to provision QoS in the WiFi network, a consistent mapping of QoS parameters and values between 3GPP and 802.11e is needed. Recommendations to map 3GPP QCI to DSCP for mobility sessions are

available in [PMIP-QoS]. This document adds the configuration of QoS per PMIP mobility session to a WiFi radio access.

The rest of the document is organized as follows. Chapter 2 outlines the QoS mechanisms in 3GPP mobile networks and 802.11 networks. Chapter 3 provides an overview of the architecture in which QoS is provisioned on the WiFi AP. Chapters 4 and 5 describe the connection model in the access network and the QoS mapping itself.

1.1 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 RFC 2119 [RFC2119].

1.2 Definitions

1.3 Abbreviations

3GPP	Third Generation Partnership Project
AAA	Authentication Authorization Accounting
ARP	Allocation and Retention Priority
AP	Access Point
DSCP	Differentiated Services Code Point
EPC	Enhanced Packet Core
GBR	Guaranteed Bit Rate
MAG	Mobility Access Gateway
MBR	Maximum Bit Rate
MN	Mobile Node
PDN-GW	Packet Data Network Gateway
QCI	QoS Class Indicator
QoS	Quality of Service
Tspec	Traffic Conditioning Spec
WLC	Wireless Controller

2. QoS Mechanisms

2.1 QoS in Mobile Networks

3GPP has standardized QoS for EPC (Enhanced Packet Core) from Release 8 [TS 23.107]. 3GPP QoS policy configuration defines access agnostic QoS parameters that can be used to provide service differentiation in multi vendor and operator deployments. The concept of a bearer is

used as the basic construct for which the same QoS treatment is applied for uplink and downlink packet flows between the MN (host) and gateway [TS23.401]. A bearer may have more than one packet filter associated and this is called a Traffic Flow Template (TFT). The IP five tuple (IP source address, port, IP destination, port, protocol) identifies a flow.

The access agnostic QoS parameters associated with each bearer are QCI (QoS Class Identifier), ARP (Allocation and Retention Priority), MBR (Maximum Bit Rate) and optionally GBR (Guaranteed Bit Rate). QCI is a scalar that defines packet forwarding criteria in the network. Mapping of QCI values to DSCP is well understood and GSMA has defined standard means of mapping between these scalars [GSMA-IR34].

An MN may have more than one IP addresses associated with the same hardware (MAC) address corresponding to each of the networks than it is attached to. This corresponds to more than one PMIP mobility session for which QoS is provisioned in the WLC.

2.2 QoS in WiFi Networks

802.11e [802.11e] defined by IEEE provides an enhancement of the MAC layer in WiFi networks to support QoS. Basic 802.11 WiFi uses CSMA and collision avoidance to provide best effort access to the medium. 802.11e defines a Hybrid Coordination Function (HCF) that provides a priority based access and also admission control based access.

HCF contention based channel access provides prioritized access to the 802.11 medium. Four access categories (AC) are defined based on traffic type. Each arriving frame is mapped into one of four FIFO queues corresponding to different user priority (UP) values. The highest priority frame is transmitted when access is obtained in a contention window. Access categories and their mapping to 802.1D user priorities is provided [802.11e].

HCF controlled channel access uses a central coordinator to provide contention free access to the medium based on admission control. The HCCA (HCF Controlled Channel Access) based scheduling can use configured policies to grant exclusive access to a QSTA (user) for limited contention free slots.

3. Policy Provisioning Architecture

This section describes the architecture in which the PMIP QoS configuration of MN sessions is applied to the corresponding traffic

flows in the WiFi Access Point. Following MN attach to the WiFi network and authentication with the mobile network, the WLC gets QoS parameters and other policy for the authorized MN. When the PMIP connection is created, the PDN-GW returns QoS policy using [PMIP-QoS] extensions.

In [PMIP-QoS], the Access Point (AP) is not directly provisioned with QoS for an MN connection. As a result, the AP is only able to prioritize flows based on observed downlink DSCP values. Additionally, the AP does not know the maximum bandwidth of a subscriber or flow to be applied on the WiFi radio network. This can result in sub-optimal utilization of scarce WiFi network resources, and of the overall access network. This solution provides a description to provision the AP with QoS policy associated to an MN.

The paragraphs that follow outline the overall architecture and subsequent chapters provide details on QoS parameters provisioned in the AP.

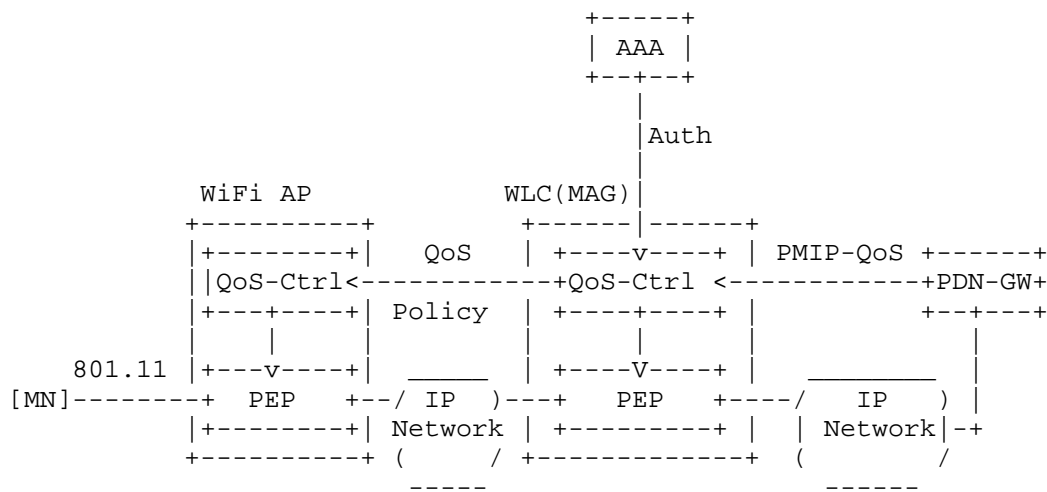


Figure 1: Architecture for provisioning QoS Policy on WiFi AP

Figure 1 provides an overview of the architecture in which QoS for an MN is provisioned on the AP. MN QoS policy from initial session authorization and PMIP connection establishment is provisioned in the WLC QoS-Ctrl (logical function). QoS-Ctrl in WLC installs QoS to the WLC PEP as described in [PMIP-QoS].

In this solution, the WLC translates the 3GPP QoS policy to equivalent parameters for 802.11e and IP flows and sends them to the WiFi AP. The protocols used to exchange QoS parameters between the

WLC and AP are not discussed in this document. The AP maps the received QoS policy configuration and applies them to upstream and downstream forwarding of data packets in the WiFi radio network. The AP also applies these QoS policies for upstream user IP flows to the WLC. The WLC should provide the AP with a policy that applies to each MN (MAC address in WiFi network) and parameters per IP flow. This model is described further in the following chapter.

4. Connections and QoS Mapping

4.1 Connection Model

MNs that attach to the mobile network have QoS policies associated to the corresponding PMIP connection. This section outlines the connection model for QoS mapping on the WiFi AP.

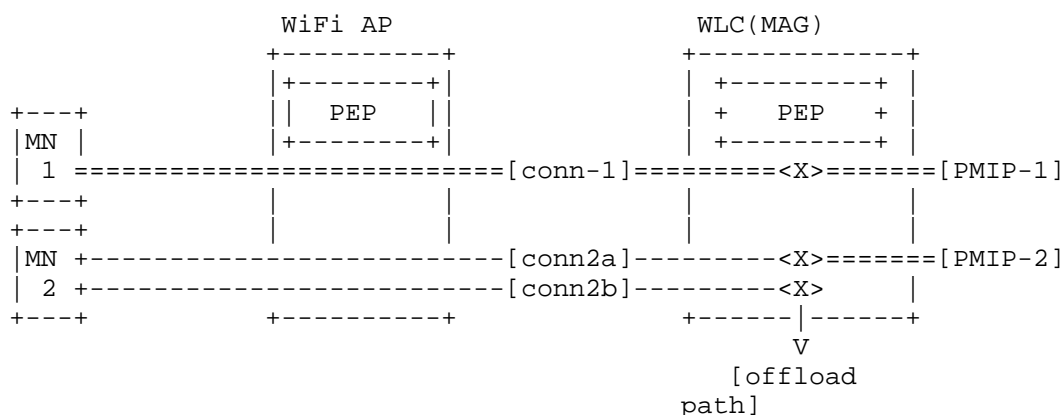


Figure 2: MN Connection and QoS Mapping

Figure 2 shows MN1 and MN2 attached to the WLC via a WiFi AP. An MN may have a tunneled connection to the mobile network (MN1, conn-1, PMIP-1), or an IP connection to the mobile network (MN2, conn2a, PMIP-2) and an IP connection that is offloaded at the WLC (MN2, conn2b, offload). The connection segment between MN and WLC may be IP connections, or tunneled connections such as IPSec.

For an MN - WLC connection segment with IP address configured via PMIP (e.g. MN2 conn2a), the corresponding PMIP QoS would be

applicable to flows with this IP address and MAC address.
For tunneled connections in MN - WLC segment (e.g. MN1 conn-1), MN1 first gets a local IP address from the WLC. MN1 then establishes an IPSec connection to the WLC and in the IKE signaling indicates parameters for PMIP connection setup. The corresponding PMIP QoS parameters are applicable to flows identified by local IP address, port (IPSec source port at MN1) and MAC address.
WiFi AP may use a default QoS profile for connection flows that are offloaded (e.g. MN2 conn2b in figure).

The WiFi AP would get QoS traffic filters corresponding to PMIP flows. These flows would be identified by the tuple {MAC, IP address, port}. The port field may be specified for tunneled or NATed connections and wildcarded otherwise.

4.2 PMIP - 802.11e QoS Configuration

The WiFi Access Point (AP) gets QoS configuration per IP session from the WLC. The QoS information per IP session provided to the AP includes:

- Hardware (MAC) address of host for which PMIP session is established.
- IP prefix or address of PMIP mobility session.
- IP port address of tunneled flow or NATed connection.
- DSCP. Diffserv PHB value of PMIP QoS for the mobility session.
- QCI. The WLC provides the 3GPP QCI value if available, for example, from authorization profile of APN (i.e. subscribed values per established PMIP mobility session).
- ARP (Allocation and Retention Priority). This value is obtained from the PMIP QoS for the mobility session. It determines the priority of a flow (1 has highest priority).
- MBR (Maximum Bit Rate) for mobility session uplink and downlink. This should not exceed the AMBR (Aggregate MBR) of the subscription.
- GBR (Guaranteed Bit Rate) for mobility session uplink and downlink, if required.

The WiFi AP uses the above QoS configuration to implement classification, admission control and forwarding of MN flows. The WiFi AP maps DSCP (or QCI) to 802.11e AC (Access Categories) for each IP session / hardware (MAC) address of the host (3GPP user). The mapping from DSCP or QCI to 802.11e AC is shown in table in chapter 4 below.

In the WiFi radio network, the AP uses 802.11e AC values for

contention (HCF) based forwarding based on priority. The AP schedules downstream flows in the WiFi radio network and for upstream IP backhaul to the WLC. For contention free scheduling (based on HCCA), the WiFi AP additionally uses the QoS configuration per user to admit flows based on 802.11e ADDTS (ADD TSpec) requests from the host (3GPP user). The WiFi AP may drop packets that fall outside the configured MBR and GBR. In case of severe radio congestion, the WiFi AP can use ARP in addition to DSCP drop precedence to determine the flows to be dropped.

5. Mapping Recommendations and Default Values

The table below outlines a recommended mapping between 3GPP QCI, and 802.11e Access Category (AC) priorities. QCI packet delay budget and packet error loss rate may be used by the WiFi access point in scheduling contention free access when HCCA scheduling is used.

QCI	DSCP	802.11e AC	Example 3GPP service
1	EF	3 AC_VO	conversational voice
2	EF	3 AC_VO	conversational video
3	EF	3 AC_VO	real-time gaming
4	AF41	2 AC_VI	buffered streaming
5	AF31	2 AC_VI	IMS signaling
6	AF31	2 AC_VI	buffered streaming
7	AF21	0 AC_BE	interactive gaming
8	AF11	0 AC_BE	web access
9	BE	1 AC_BK	e-mail

Table 1: QoS Mapping between QCI, WMM, 802.11e AC

6. Next Steps

This document has described a basic model for mapping PMIP QoS parameters to 802.11e parameters. However, there are a few questions that need to be explored further.

The draft that the protocol between WLC and AP is not considered further here. There needs to be some work to determine the protocol parameters and other details if it is desired that WLC and AP should interwork.

Another aspect is this draft does not describe multiple PDN connections per MN in much detail. This is work in progress in 3GPP and the results should be compatible with the model in this draft. RTC Web flows introduce cases where the same IP flow (5-tuple) can have multiple DSCP values - for example when the IP flow carries audio and video applications. This would impact the QCI and DSCP parameters of IP flows, but this is not only this description that is affected. A

Finally, the QoS values listed in the table in chapter 5 needs to be aligned with [PMIP-QoS] and GSMA.

7. Security Considerations

This document describes mapping of 3GPP QoS profile and parameters to IEEE 802.11e parameters. No security concerns are expected as a result of using this mapping.

8. IANA Considerations

No IANA assignment of parameters are required in this document.

9. References

9.1 Normative References

- [KEYWORDS] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, March 1997.
- [RFC1776] Crocker, S., "The Address is the Message", RFC 1776, April 1 1995.
- [TRUTHS] Callon, R., "The Twelve Networking Truths", RFC 1925, April 1 1996.

9.2 Informative References

- [EVILBIT] Bellovin, S., "The Security Flag in the IPv4 Header", RFC 3514, April 1 2003.
- [RFC5513] Farrel, A., "IANA Considerations for Three Letter Acronyms", RFC 5513, April 1 2009.
- [RFC5514] Vyncke, E., "IPv6 over Social Networks", RFC 5514, April 1 2009.
- [PMIP-QoS] Liebsch, et al., "Quality of Service Option for Proxy Mobile IPv6", draft-ietf-netext-pmip6-qos-00, June 2012.
- [RFC 2211] Wroclawski, J., "Specification of the Controlled Load Quality of Service", RFC 2211, September 1997.
- [RFC 2212] Shenker, S., Partridge, C., and R. Guerin, "Specification

of Guaranteed Quality of Service", RFC 2212, September 1997.

- [RFC 2216] Shenker, S., and J. Wroclawski, "Network Element QoS Control Service Specification Template", RFC 2216, September 1997.
- [TS23.107] Quality of Service (QoS) Concept and Architecture, Release 10, 3GPP TS 23.107, V10.2.0 (2011-12).
- [TS23.207] End-to-End Quality of Service (QoS) Concept and Architecture, Release 10, 3GPP TS 23.207, V10.0.0 (2011-03).
- [TS23.401] General Packet Radio Service (GPRS) enhancements for Evolved Universal Terrestrial Radio Access Network (E-UTRAN) access (Release 11), 3GPP TS 23.401, V11.2.0 (2012-06).
- [TS23.203] Policy and Charging Control Architecture, Release 11, 3GPP TS 23.203, V11.2.0 (2011-06).
- [TS29.212] Policy and Charging Control over Gx/Sd Reference Point, Release 11, 3GPP TS 29.212, V11.1.0 (2011-06).
- [802.11e] IEEE, "IEEE part 11: Wireless LAN Medium Access Control(MAC) and Physical Layer (PHY) specifications. Amendment 8: Medium Access Control (MAC) Quality of Service Enhancements" 802.11e-2005, 22 September 2005.
- [GSMA-IR34] Inter-Service Provider Backbone Guidelines 5.0, 22 December 2010

Authors' Addresses

John Kaippallimalil
5340 Legacy Drive, Suite 175
Plano Texas 75024

E-Mail: john.kaippallimalil@huawei.com

Netext WG
Internet-Draft
Intended status: Standards Track
Expires: April 25, 2013

S. Krishnan
Ericsson
S. Gundavelli
Cisco
M. Liebsch
NEC
H. Yokota
KDDI
J. Korhonen
Nokia Siemens Networks
October 22, 2012

Update Notifications for Proxy Mobile IPv6
draft-krishnan-netext-update-notifications-01

Abstract

Proxy Mobile IPv6 (PMIPv6) is a network based mobility management protocol that enables IP mobility for a host without requiring its participation in any mobility-related signaling. This document proposes a mechanism for the Local Mobility Anchor to asynchronously notify the Mobile Access Gateway about changes related to the mobility session.

Status of this Memo

This Internet-Draft is submitted in full conformance with the provisions of BCP 78 and BCP 79.

Internet-Drafts are working documents of the Internet Engineering Task Force (IETF). Note that other groups may also distribute working documents as Internet-Drafts. The list of current Internet-Drafts is at <http://datatracker.ietf.org/drafts/current/>.

Internet-Drafts are draft documents valid for a maximum of six months and may be updated, replaced, or obsoleted by other documents at any time. It is inappropriate to use Internet-Drafts as reference material or to cite them other than as "work in progress."

This Internet-Draft will expire on April 25, 2013.

Copyright Notice

Copyright (c) 2012 IETF Trust and the persons identified as the document authors. All rights reserved.

This document is subject to BCP 78 and the IETF Trust's Legal

Provisions Relating to IETF Documents
(<http://trustee.ietf.org/license-info>) in effect on the date of publication of this document. Please review these documents carefully, as they describe your rights and restrictions with respect to this document. Code Components extracted from this document must include Simplified BSD License text as described in Section 4.e of the Trust Legal Provisions and are provided without warranty as described in the Simplified BSD License.

Table of Contents

1. Introduction	3
2. Example use case	3
3. LMA Behavior	4
4. MAG Behavior	4
5. Message Formats	5
5.1. Update Notification(UPN)	5
5.2. Update Notification Acknowledgement(UPA)	5
6. Security Considerations	6
7. IANA Considerations	6
8. Acknowledgements	7
9. References	7
9.1. Normative References	7
9.2. Informative References	7
Authors' Addresses	7

1. Introduction

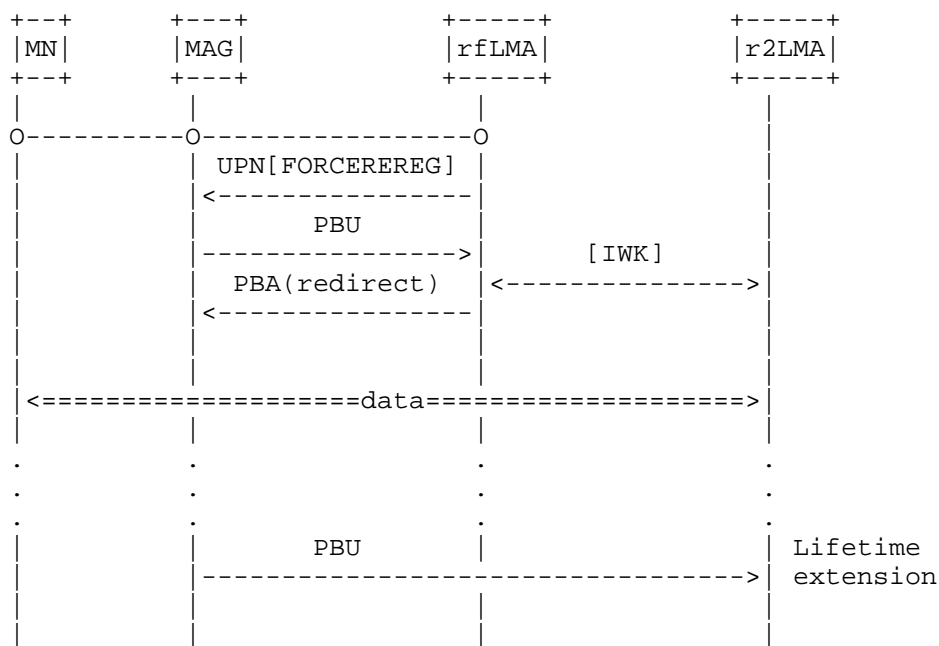
Proxy Mobile IPv6 [RFC5213] describes the protocol operations to maintain reachability and session persistence for a Mobile Node (MN) without the explicit participation from the MN in signaling operations at the Internet Protocol (IP) layer. In order to facilitate such network-based mobility, the PMIPv6 protocol defines a Mobile Access Gateway (MAG), which acts as a proxy for the Mobile IPv6 [RFC6275] signaling, and the Local Mobility Anchor (LMA) which acts similar to a Home Agent. The setup of the mobility session is initiated by the MAG by sending a PBU message and confirmed by the LMA in the PBA message. Once the mobility session is set up for a given lifetime, the LMA has no mechanism to inform the MAG about changes to the mobility session or any parameters related to the mobility session.

One such scenario where such a mechanism is needed is when the LMA wants to inform the MAG that it needs to reregister. It is possible to achieve a similar effect by using a much shorter lifetime for the mobility sessions but in several networks this results in an unacceptable, and mostly unnecessary, increase in the signaling load and overhead.

This document defines a new mobility header message for performing notifications and a corresponding mobility header message for the MAG to acknowledge the notification. While it is possible to use an existing mobility header type for this purpose, for instance the PMIPv6 Heartbeat message [RFC5847], the existing messages do not provide the required semantics. e.g. The Heartbeat message does not provide a reason why it was sent.

2. Example use case

Consider an use case where an LMA (r1LMA) wants to move over one or more mobility sessions from a given MAG to a different LMA (r2LMA) using [RFC6463]. e.g. In order to allow planned maintenance. The LMA could send an update notification to the MAG to force a re-registration for one or more MNs. The MAG tries to register and gets a redirect from the r1LMA towards the r2LMA.



3. LMA Behavior

The LMA sends the Update Notification message in response to a condition that is specified in the Notification Reason field. If the LMA requires an acknowledgement from the MAG concerning the UPN message, it MUST set the A bit to 1. If not it MUST set the A bit to 0. The LMA MAY retransmit the UPN messages if reliability is required for the specific Notification reason. If the UPN message is retransmitted, the LMA MUST reuse the same sequence number as the original message. If the LMA receives an UPA message with a failure Status (Status value >127) it SHOULD log an error.

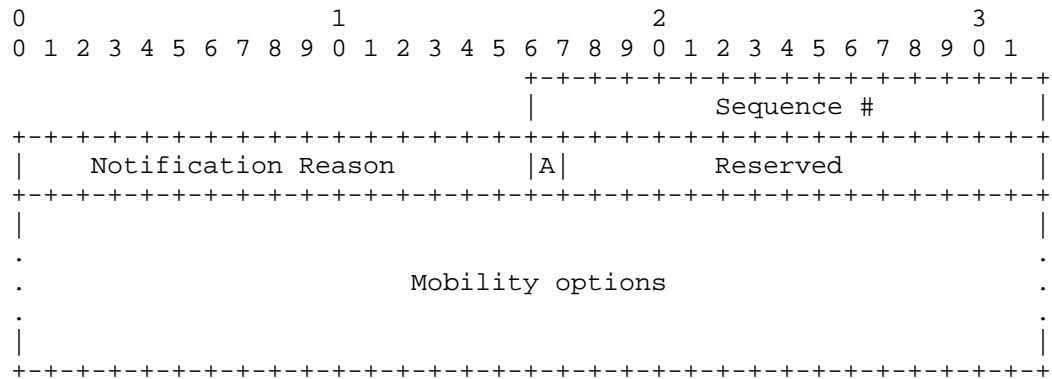
4. MAG Behavior

If a received Update Notification message has the A bit set to 1, the MAG MUST create and transmit an Update Notification Acknowledgement message in response to the UPN message. The sequence number of the UPA message MUST be copied from the UPN message that is being responded to. Depending on whether the message was processed successfully or not, the MAG MUST set the Status value in the UPA message to an appropriate value. The actual processing required on the MAG is out of the scope of this document and will be specified for each Notification reason.

5. Message Formats

5.1. Update Notification(UPN)

The LMA sends an UPN message to a MAG to notify the MAG that some information regarding the mobility session or parameters related to the mobility session has changed.



Sequence Number: A monotonically increasing integer. Set by the LMA and retained for retransmissions.

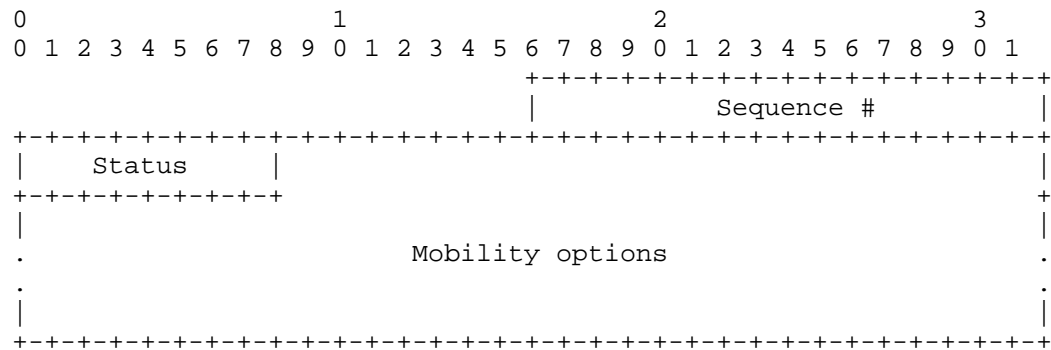
Acknowledgement Requested (A): If this bit is set, the MAG MUST send an UPA message in response to the received UPN message.

Notification Reason: Contains the code corresponding to the reason that caused the LMA to send the Update Notification to the MAG. This field does not contain any structure and MUST be treated as an enumeration.

Mobility Options: Contains a set of mobility options for the MAG to act upon. The set of mobility options that can be present in the message is related to the Notification Reason field in the message.

5.2. Update Notification Acknowledgement(UPA)

The MAG sends an UPA message to a LMA in order to acknowledge that it has received an UPN message with the A bit set.



Sequence Number: Copied from the UPN message being acknowledged.

Status: Specifies the result of the MAG's processing of the UPN message. The status codes between 0 and 127 signify successful processing of the UPN message and codes between 128 and 255 signify that an error occurred during processing of the UPN message.

Mobility Options: Contains a set of mobility options used to provide context to the LMA. The set of mobility options that can be present in the message is related to the Status field in the message.

6. Security Considerations

The protocol specified in this document uses the same security association as defined in [RFC5213] for use between the LMA and the MAG to protect the UPN messages. Support for integrity protection using IPsec is REQUIRED, but support for confidentiality is NOT REQUIRED.

7. IANA Considerations

The Update Notification message require a single Mobility Header Type (TBA1) from the Mobility Header Types registry at <http://www.iana.org/assignments/mobility-parameters>

The Update Notification Acknowledgement message require a single Mobility Header Type (TBA2) from the Mobility Header Types registry at <http://www.iana.org/assignments/mobility-parameters>

This document creates a new registry for Notification Reasons. The

allocation policy for this field is First Come, First Served.

This document creates a new registry for Status codes in the UPA message. The allocation policy for this field is First Come, First Served.

8. Acknowledgements

The authors would like to thank Basavaraj Patil, Rajeev Koodli and other members of netext working group for their valuable comments to improve 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, March 1997.
- [RFC5213] Gundavelli, S., Leung, K., Devarapalli, V., Chowdhury, K., and B. Patil, "Proxy Mobile IPv6", RFC 5213, August 2008.

9.2. Informative References

- [RFC5847] Devarapalli, V., Koodli, R., Lim, H., Kant, N., Krishnan, S., and J. Laganier, "Heartbeat Mechanism for Proxy Mobile IPv6", RFC 5847, June 2010.
- [RFC6275] Perkins, C., Johnson, D., and J. Arkko, "Mobility Support in IPv6", RFC 6275, July 2011.
- [RFC6463] Korhonen, J., Gundavelli, S., Yokota, H., and X. Cui, "Runtime Local Mobility Anchor (LMA) Assignment Support for Proxy Mobile IPv6", RFC 6463, February 2012.

Authors' Addresses

Suresh Krishnan
Ericsson
8400 Blvd Decarie
Town of Mount Royal, Quebec
Canada

Phone: +1 514 345 7900 x42871
Email: suresh.krishnan@ericsson.com

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

Email: sgundave@cisco.com

Marco Liebsch
NEC

Email: marco.liebsch@nw.neclab.eu

Hidetoshi Yokota
KDDI

Email: yokota@kddilabs.jp

Jouni Korhonen
Nokia Siemens Networks
Linnoitustie 6
FI-02600 Espoo
Finland

Email: jouni.nospam@gmail.com

NETEXT WG
Internet Draft
Intended status: Standard Track
Expires: April 16, 2013

S. Jeon
Instituto de Telecomunicacoes
B. Sarikaya
Huawei
R. L. Aguiar
Universidade de Aveiro
October 15, 2012

Network Mobility Support using Mobile MAG in Proxy Mobile IPv6 Domain

draft-sijeon-netext-mmag-pmip-00.txt

Abstract

This draft specifies IP mobility support protocol for moving network including mobile nodes (MNs) over Proxy Mobile IPv6 network by introducing a new functional entity, mobile MAG (mMAG) on the moving network. The mMAG takes charge of MN's movement detection, binding update on behalf of MNs as a mobile access gateway (MAG) does in PMIPv6 infrastructure. This protocol also supports IP session continuity for a mobile node to move between mobile network and fixed MAG. This protocol does not require any modification or extension on the MN.

Status of this Memo

This Internet-Draft is submitted to IETF in full conformance with the provisions of BCP 78 and BCP 79.

Internet-Drafts are working documents of the Internet Engineering Task Force (IETF), its areas, and its working groups. Note that other groups may also distribute working documents as Internet-Drafts.

Internet-Drafts are draft documents valid for a maximum of six months and may be updated, replaced, or obsoleted by other documents at any time. It is inappropriate to use Internet-Drafts as reference material or to cite them other than as "work in progress."

The list of current Internet-Drafts can be accessed at
<http://www.ietf.org/ietf/lid-abstracts.txt>

The list of Internet-Draft Shadow Directories can be accessed at
<http://www.ietf.org/shadow.html>

This Internet-Draft will expire on April 16, 2013.

Copyright Notice

Copyright (c) 2012 IETF Trust and the persons identified as the document authors. All rights reserved.

This document is subject to BCP 78 and the IETF Trust's Legal Provisions Relating to IETF Documents (<http://trustee.ietf.org/license-info>) in effect on the date of publication of this document. Please review these documents carefully, as they describe your rights and restrictions with respect to this document. Code Components extracted from this document must include Simplified BSD License text as described in Section 4.e of the Trust Legal Provisions and are provided without warranty as described in the Simplified BSD License.

Table of Contents

1. Introduction	3
2. Conventions and Terminology.....	3
3. Overview	4
3.1. Initial Attach	4
3.2. mMAG Handoff	6
3.3. Mobile Node Handoff.....	6
4. mMAG Operation	7
5. LMA Operation	7
6. MAG Operation	7
7. MN Operation	7
8. IANA Considerations	8
9. Security Considerations.....	8
10. References	8
10.1. Normative References.....	8
10.2. Informative References.....	8
Authors' Addresses	9

1. Introduction

Network mobility is a novel concept for handling a group of nodes within a moving vehicular area. It provides an effective way for wireless hosts to access the Internet through an intermediate router connecting to an external wireless wide access network.

Proxy Mobile IPv6 (PMIPv6) is a network-based IP mobility protocol, taking charge of host movement detection, binding update on behalf of mobile nodes (MNs). It does not require any modification on mobile node and thus provides better mobility performance compared to host-based mobility protocol, e.g. Mobile IPv6 [RFC6275]. However, it does not support network mobility in the specification [RFC5213].

NEMO Basic Support protocol (NEMO-BSP) [RFC3963] addressed this issue for allowing a host within a moving vehicle to continue their IP sessions when the vehicle is moving between access routers. However, NEMO-BSP employs a mobile router (MR), which requires MIPv6 client function having host-based mobility protocol feature. According to this fact, a MR introduced in NEMO-BSP is not aligned with network-based PMIPv6 approach, thus it is not suited to be used in PMIPv6 domain.

This draft describes network mobility support over PMIPv6 domain by introducing a new entity, called mobile MAG (mMAG) [N-PMIPv6], which is responsible for detecting MN's movement, performing mobility management operation on behalf of MNs, and managing binding update list as a MAG does.

This draft is based on stateless IPv6 address configuration for mMAG and MNs to configure their IP addresses not by using DHCP. This idea does not require any modification or extension on the MN.

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 RFC 2119 [RFC2119].

This document uses the terminology defined in [RFC5213]. In addition to, we defined Mobile MAG (mMAG) as follow.

- Mobile MAG (mMAG): A mobile router, which has a similar function to MAG defined in PMIPv6 specification.

3. Overview

3.1. Initial Attach

This sub-section describes initial attach of mMAG and MN. The mMAG is not a mobile router (MR) having Mobile IPv6 client functionality presented in [RFC3963]. The mMAG is the entity that performs the mobility management on behalf of MNs within mobile network. It has upstream and downstream interfaces; upstream interface is seen as normal MN to a MAG and it is connected as a tunnel with a LMA over PMIPv6 tunnel between a MAG and a LMA. Downstream interface is seen as fixed MAG in PMIPv6 infrastructure to attached MNs. LMA sees mMAG as a normal MN and then process initial attach process specified in [RFC5213] for normal MN.

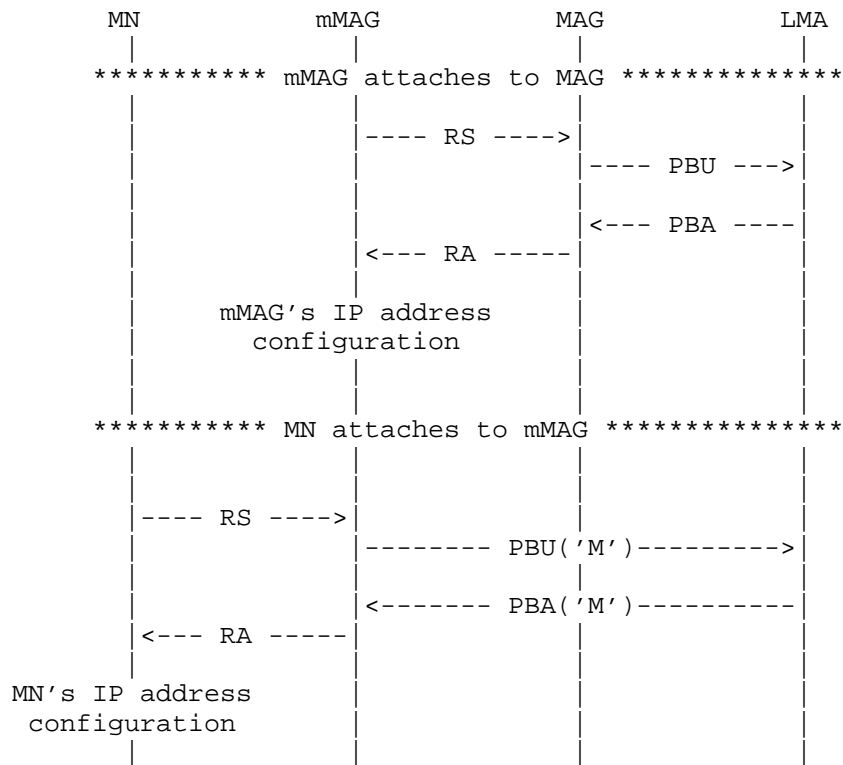


Figure 1 mMAG and MN Attachment - Signaling Call Flow

Figure 1 shows the signaling call flow when the mMAG enters the Proxy Mobile IPv6 domain. In order to enable network mobility service support, the mMAG should be attached first to PMIPv6 domain. The MAG on detecting the mMAG performs authentication and authorization process as it does normal MN in [RFC5213] and then sends Proxy Binding Update message to the LMA.

The LMA on receiving Proxy Binding Update message creates new Binding Cache entry and assigns new prefix and associates mMAG's ID and assigned prefix. The LMA sends Proxy Binding Acknowledgement message including assigned prefix to the MAG. The mMAG configures its IPv6 address based on the prefix received from Router Advertisement.

Subsequently, when a MN enters into wireless range of mobile network, the mMAG detects MN's attachment and performs NEMO service availability of attached MN through authentication and authorization processes. If is verified, the mMAG will be aware of the address of the LMA to which it belongs for attached MN. The mMAG then sends Proxy Binding Update message with setting 'M' flag to the associated LMA. The MAG as intermediate entity between mMAG and LMA will treat this message as normal packets originated from the mMAG and send it after encapsulating the message with destination IP address of LMA. On receiving encapsulated Proxy Binding Update message, the LMA will decapsulate and processes the message by adding the MN's ID to Binding Cache with setting 'M' flag indicating that this node belongs to a mobile network and store mMAG's source IP address as Proxy Care-of Address in Proxy Binding Update message.

The LMA then assigns and delivers new prefix to the mMAG by sending Proxy Binding Acknowledgement message with setting 'M' flag. The mMAG sends Router Advertisement message to the MN. The MN configures its stateless IPv6 address based on received prefix in Router Advertisement message.

The MAG is transparent for exchanging signaling messages and data packets between mMAG and LMA. 'M' flag is used to let the LMA know that additional Binding Cache entry lookup should be allowed when it receives packets destined MN's prefix. On receiving Proxy Binding Update message, the LMA sets 'M' flag in Binding Cache entry of the MN. As a result, when a LMA receives a packet destined MN's prefix within mobile network, it performs recursive look up processing. In a first look up, the LMA obtains the mMAG to which the MN is attached. And in a second look up, the LMA obtains a MAG to which the mMAG of the MN is attached. The packets will be tunneled with mMAG's IP address and MAG's address to which mMAG belongs, for destination IP address in inner/outer tunnel header, respectively.

3.2. mMAG Handoff

The mMAG's handoff is assumed that the mMAG moves to the newly attached mobile access gateway (n-MAG) from the previously attached mobile access gateway (p-MAG).

The mMAG's handoff process follows normal PMIPv6 handoff specified in [RFC5213]. The n-MAG detects mMAG attach and sends Proxy Binding Update message to mMAG's associated LMA by following standard PMIPv6 operation.

On receiving Proxy Binding Update message from n-MAG, the LMA will change the transport endpoint of the tunnel from p-MAG to n-MAG in LMA Binding Cache. The LMA is not required to perform additional operations for MNs within mMAG in Binding Cache due to mMAG's handoff because each binding of MN and mMAG, mMAG and MAG is managed separately. After updating Binding Cache entry of mMAG, the LMA sends Proxy Binding Acknowledgment message to the n-MAG. The n-MAG will send Router Advertisements containing the mMAG's home network prefix, and this will ensure the mMAG will not detect any change with respect to layer-3 attachment of its interface.

3.3. Mobile Node Handoff

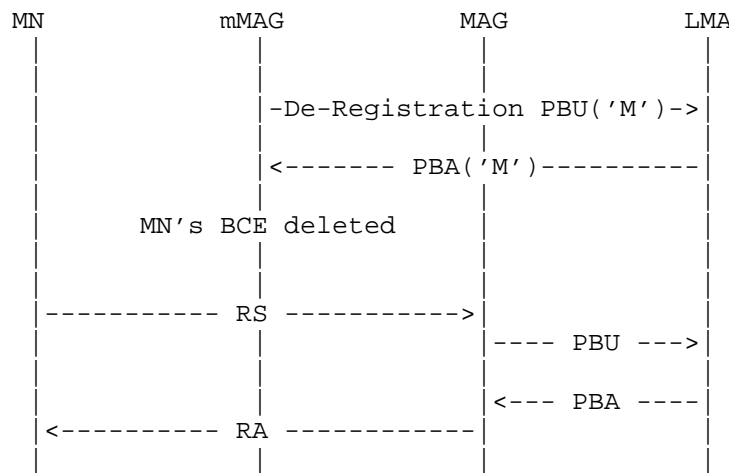


Figure 2 Mobile Node Handoff - Signaling Call Flow

Figure 2 shows the signaling call flow for the MN's handoff from mMAG to fixed MAG in same PMIPv6 domain. When the mMAG detects the MN detach, it will send de-registration Proxy Binding Update with the lifetime value of zero and setting 'M' flag to the LMA. Upon receiving the Proxy Binding Update message, the LMA waits for amount of time specified in [RFC5213], before it deletes the Binding Cache entry. On accepting Proxy Binding Acknowledgment message from the LMA the mMAG deletes the MN in the Binding Update List. On detecting new MN, the MAG performs initial attach operation following the specification in [RFC5213]. The LMA updates MN's Binding Cache entry by changing Proxy CoA with the MAG's address and setting 'M' flag to '0'.

4. mMAG Operation

A mMAG, a new functional entity, is responsible for taking charge of MNs within mobile network to detect the MN's movements to and from the access link and to send the Proxy Binding Update message on behalf of the MN to the LMA as a MAG does in this document. The mMAG has same data structure of Binding Update List a MAG has and it emulates attached MNs by sending Router Advertisements based on each MN's home network prefix in the Proxy Binding Update List. But when the mMAG sends the Proxy Binding Update message to the LMA, it is required to add 'M' flag in Proxy Binding Update message.

5. LMA Operation

When the LMA receives Proxy Binding Update message, it is not required to recognize where the message comes from fixed MAG or mMAG and to have knowledge of mMAG list not to extend PMIPv6 specification possibly. However, the LMA needs to have additional element called 'M' flag in Binding Cache to distinguish which kinds of MAG the node is attached. This is used to provide efficient mMAG handoff management, not requiring the changes of Binding Cache of MNs within mobile network due to mMAG's handoff and to forward the packets destined the MN that belongs to mMAG.

6. MAG Operation

A MAG is transparent for providing network mobility support. The mMAG attached to the MAG is treated as normal MN. No extension or modification is required to the MAG.

7. MN Operation

No extension is required.

8. IANA Considerations

This document makes no request of IANA.

9. Security Considerations

TBD

10. References

10.1. Normative References

- [RFC2119] S. Bradner, "Key words for use in RFCs to Indicate Requirement Levels", RFC 2119, March 1997.
- [RFC6275] C. Perkins, D. Johnson, and J. Arkko, "Mobility Support in IPv6", IETF RFC 6275, July 2011.
- [RFC5213] S. Gundavelli, K. Leung, V. Devarapalli, K. Chowdhury, and B. Patil, "Proxy Mobile IPv6", IETF RFC 5213, August 2008.
- [RFC3963] V. Devarapalli, R. Wakikawa, A. Petrescu, and P. Thubert, "Network Mobility (NEMO) Basic Support Protocol", RFC3963, January 2005.

10.2. Informative References

- [N-PMIPv6] I. Sogo, C. J. Bernardos, M. Calderon, A. Banchs, and A. Azcorra, "NEMO-Enabled Localized Mobility Support for Internet Access in Automotive Scenarios", IEEE Coms. Mag., vol.47, no.5, pp.152-159, May 2009.

Authors' Addresses

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

E-mail: seiljeon@av.it.pt

Behcet Sarikaya
Huawei
5340 Legacy Dr.
Plano, TX 75024, USA

E-mail: sarikaya@ieee.org

Rui L. Aguiar
Universidade de Aveiro
3810-193 Aveiro, Portugal

E-mail: ruilaa@ua.pt

NETEXT Working Group
Internet-Draft
Intended status: Standards Track
Expires: January 10, 2013

Y. Tu
C. Zhu
ZTE
CJ. Bernardos
UC3M
C. Williams
MCSR Labs
July 9, 2012

MN Status Option for Proxy Mobile IPv6
draft-tu-netext-mn-status-option-02

Abstract

IP flow mobility enables the ability of movement of selected flows from one access technology to another. This document extends the Proxy Mobile IPv6 signaling to convey mobile node's status information that can be used by the network to decide when and how perform flow mobility. It also defines options allowing the network getting information about different mobile node capabilities, which might be considered to decide how to tackle the node's mobility.

Status of this Memo

This Internet-Draft is submitted in full conformance with the provisions of BCP 78 and BCP 79.

Internet-Drafts are working documents of the Internet Engineering Task Force (IETF). Note that other groups may also distribute working documents as Internet-Drafts. The list of current Internet-Drafts is at <http://datatracker.ietf.org/drafts/current/>.

Internet-Drafts are draft documents valid for a maximum of six months and may be updated, replaced, or obsoleted by other documents at any time. It is inappropriate to use Internet-Drafts as reference material or to cite them other than as "work in progress."

This Internet-Draft will expire on January 10, 2013.

Copyright Notice

Copyright (c) 2012 IETF Trust and the persons identified as the document authors. All rights reserved.

This document is subject to BCP 78 and the IETF Trust's Legal Provisions Relating to IETF Documents (<http://trustee.ietf.org/license-info>) in effect on the date of

publication of this document. Please review these documents carefully, as they describe your rights and restrictions with respect to this document. Code Components extracted from this document must include Simplified BSD License text as described in Section 4.e of the Trust Legal Provisions and are provided without warranty as described in the Simplified BSD License.

Table of Contents

1. Introduction	3
2. Conventions used in this document	3
3. New MN status and capabilities options for PMIPv6	3
3.1. Overview	3
3.2. Use case scenarios	4
3.3. Solution	5
3.3.1. MAG considerations	5
3.3.2. LMA considerations	5
3.4. Mobile Node Status and Capabilities Options	6
3.4.1. Mobile Node Capability Status Option	6
3.4.2. Mobile Node Connectivity Status Option	7
4. Security Considerations	8
5. IANA Considerations	8
6. Contributors	9
7. Normative References	9
Authors' Addresses	9

1. Introduction

There are several use cases where it would be useful that the local mobility anchor (LMA) can decide to perform flow mobility from one access network to another, e.g., from 3GPP to WLAN or from WLAN to WiMAX. With current Proxy Mobile IPv6 specification [RFC5213], the LMA can only know the different access technologies the mobile node (MN) is attached to (this information is conveyed from the mobile access gateway to the local mobility anchor in the Access Technology Type option). No accurate information about the mobile node status (e.g., if it is in idle/power saving mode or experiencing low radio quality) is available at the LMA to aid it in the decision of when and how to perform flow mobility. It is therefore helpful to provide the LMA with additional information from the MN, so the LMA can trigger flow mobility actions with a lower risk of failure/data loss. This can be done by including mobile node status information in the signaling between the mobile access gateway and the local mobility anchor, and by enabling the mobile access gateway to update that information as needed.

It is also useful to support the mobile access gateway to convey information to the local mobility anchor about specific capabilities of attached mobile nodes. These capabilities may include, for example, the support of a logical interface (to hide from the IP stack the existence of multiple physical interfaces, which may be simultaneously attached to different MAGs), the availability of a dual IPv4/IPv6 stack at the mobile node, etc.

This document defines two new mobility options, the MN Capability Status option and the MN Connectivity Status option for Proxy Mobile IPv6 (PMIPv6), that can be used by the mobile access gateway (MAG) for carrying information to the local mobility anchor about the MN status with the correspondent access network and its capabilities.

2. Conventions used in this document

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

3. New MN status and capabilities options for PMIPv6

3.1. Overview

In some Proxy Mobile IPv6 deployments, a mobile network (e.g., the one defined by the 3GPP) needs to support multiple access

technologies, and the local mobility anchor can be triggered to decide which access technology will be used to move a particular IP flow according to the operator preferences and local policy. To guarantee the success of the flow mobility procedure from one access technology to another, a critical piece of information to help the LMA is the current mobile node status at the different access networks it might be attached to.

The mobile access gateway is the right PMIPv6 network entity to detect the mobile node status using, in addition to the mechanisms defined in RFC5213 [RFC5213], any access network specific mechanism that is available to detect the connectivity status of the attached mobile node.

The MAG can provide the mobile node status information to the LMA as part of the signaling exchange between MAG and LMA. Namely, the MAG can periodically, or triggered by a particular event, update the MN status to the LMA. How the LMA use this information is outside the scope of this document.

3.2. Use case scenarios

The approach specified in this document provides additional benefits to some use cases involving flow mobility among multiple access technologies. These use case scenarios are illustrated next:

- (a) The user is simultaneously attached and accessing some services from both WLAN and 3GPP networks, and for some time the network link connecting to the WLAN access network is going to be released for some purposes, such as scheduled maintenance. Triggered by that event, there are two choices the LMA can take to reallocate these IP flows, either to switch the affected flows to the 3GPP access, or to switch the flows to another WLAN access (if available). Without updated information about the status of the MN, the LMA can trigger an erroneous flow mobility decision (leading to a long delay and/or data losses), for example if a flow is moved to an network interface that is currently in idle state or perceiving a low signal quality.
- (b) At residential areas, during night there are more people using WLAN, and less people using a cellular access, hence for the VoIP service it might be better to switch some users to the cellular access. On the other hand, during the day, there are less people on WLAN and more people using cellular, so it might be better to use the WLAN to offload the cellular network. The LMA can move flows according to policies, but without accurate knowledge of the MN status the flow handoff may suffer from delays, data loss or other performance problems.

- (c) The user is accessing some services from both WLAN and cellular, and an FTP IP flow is initiated which may cause the bandwidth resources to be insufficient. The LMA may consider changing the flows for VoIP service from the WLAN to the 3GPP access according to the operator policies and other factors (e.g., user preferences). If the LMA only has information about which networks the MN is connected, but not the real status/quality of each of them, it might be that the LMA incorrectly decides to move a flow (e.g., if the 3GPP radio link connecting between MN and MAG is poor) resulting in then long delay or data loss.

3.3. Solution

3.3.1. MAG considerations

The MAG can retrieve the Mobile Node status from some other network elements, such as the Mobility Management Entity (MME) in 3GPP, the Paging controller in WiMAX, or the Access Point in WLAN. The MAG may periodically, or triggered by a specific event, update this information to the LMA, so that this information can be used as one of the factors to make the decision of flow mobility by the LMA. In particular, the MAG can also retrieve the radio quality of the MAG-MN link from other network elements (e.g., the eNB in 3GPP), and a threshold can be configured locally or downloaded (e.g., from the network management system) as the operator policies. As soon as the radio quality of the MAG-MN link drops below this threshold, the MAG updates this event to the LMA as a part of Mobile Node status information, which helps the LMA making the best decision of performing flow mobility.

In addition, some Mobile Node capabilities information, such as logical interface support or dual-stack availability, can also be carried from the MAG to the LMA, helping to make a flow mobility decision. How the MAG obtains this capabilities information is out of scope of this document.

3.3.2. LMA considerations

The LMA receives the mobile node status information from the MAG, and makes the decision of flow mobility for a specific IP flow according to the operator policies and other factors (e.g., user preferences and MN status). How the LMA uses this information is outside the scope of this document.

We consider next the use case (a) of Section 3.2 as an example. If the WLAN infrastructure is scheduled for maintenance, the LMA can check the operator policies with other factors to decide which is the best candidate access network to move the flows that were using the

WLAN. One example of possible prioritized list could be the following:

1. 3GPP access with MN status set to "connected".
2. Another available WLAN access infrastructure.
3. 3GPP access with MN status set to "idle".

In this case, if the MN status is "idle" in the 3GPP access network, the LMA will hand off the mobile node to another WLAN access without trying to wake up the mobile node and re-establish the link connecting the MN and the MAG.

Another example is the following, in which the priority of each access network is set in a different way:

1. 3GPP access with MN status set to "connected".
2. 3GPP access with MN status set to "idle".
3. Another available WLAN access infrastructure.

In this case, the LMA will first try to wake up the mobile node in the 3GPP access and re-establish the link connecting the MN and the MAG. If this procedure can be done successfully, the LMA will not attempt to hand off the mobile node to another WLAN access, but it will initiate the flow mobility to the 3GPP access.

3.4. Mobile Node Status and Capabilities Options

This section defines extensions to the Proxy Mobile IPv6 [RFC5213] Protocol messages.

3.4.1. Mobile Node Capability Status Option

A new option, called Mobile Node Capability Status Option, is defined to be included in the PMIPv6 signaling (e.g., PBU and PBA messages) exchanged between a local mobility anchor and a mobile access gateway. This option is used for conveying the mobile node's features support capability information. Its format is the following:




Figure 1: MN Capability Status Option

MN-Co-St Type

To be assigned by IANA.

MN-Co-St Lngth

8-bit unsigned integer indicating the length in octets of the option, excluding the type and length fields.

Reserved

This field is unused for now. The value MUST be initialized to 0 by the sender and MUST be ignored by the receiver.

L

1-bit unsigned integer indicating the capability of supporting the logical interface feature by the mobile node. The value is set to 1 when logical interface is supported by the mobile node, otherwise it is set to 0.

D

1-bit unsigned integer indicating the IPv6/IPv4 Dual Stack support of the mobile node. The value is set to 1 when IPv6/IPv4 Dual Stack is supported by the mobile node, otherwise it is set to 0.

M

1-bit unsigned integer indicating the Mobile IPv6 support of the mobile node. The value is set to 1 when Mobile IPv6 stack is supported by the mobile node, otherwise it is set to 0.

3.4.2. Mobile Node Connectivity Status Option

A new option, called Mobile Node Connectivity Status Option, is defined to be included in the PMIPv6 signaling (e.g., PBU and PBA messages) exchanged between a local mobility anchor and a mobile access gateway. This option is used for conveying to the network the mobile node's air link connectivity status information. Its format is the following:




Figure 2: MN Connectivity Status Option

MN-Co-St Type

To be assigned by IANA.

MN-Co-St Lngth

8-bit unsigned integer indicating the length in octets of the option, excluding the type and length fields.

Reserved

This field is unused for now. The value MUST be initialized to 0 by the sender and MUST be ignored by the receiver.

MN status

The status of the mobile node attached from a specific access network, such as WiFi, WiMAX and 3GPP. Currently the value of the MN status can be as follows:

- 1: connected,
- 2: disconnected,
- 3: idle/power saving mode,
- 4: reserved.

RQ

This field is used to indicate when the radio quality of the MAG-MN link dropped below a certain threshold configured by the operators. The value can be set as follow:

- 00: the radio quality of the MAG-MN link does not drop below the threshold,
- 01: the radio quality of the MAG-MN link drops below the threshold.

All other values are reserved.

4. Security Considerations

TBD

5. IANA Considerations

TBD

6. Contributors

The following people contributed to this document (in no specific order):

Yifeng Bi
ZTE
bi.yifeng@zte.com.cn

Tricci So
ZTE USA
tso@zteusa.com

7. Normative References

- [RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, March 1997.
- [RFC5213] Gundavelli, S., Leung, K., Devarapalli, V., Chowdhury, K., and B. Patil, "Proxy Mobile IPv6", RFC 5213, August 2008.

Authors' Addresses

Yangwei Tu
ZTE
Nanjing
Nanjing
China

Email: tu.yangwei@zte.com.cn

Chunhui Zhu
ZTE
Nanjing
Nanjing
China

Email: zhu.chunhui@zte.com.cn

Carlos J. Bernardos
Universidad Carlos III de Madrid
Av. Universidad, 30
Leganes, Madrid 28911
Spain

Phone: +34 91624 6236
Email: cjbc@it.uc3m.es
URI: <http://www.it.uc3m.es/cjbc/>

Carl Williams
MCSR Labs
USA

Phone:
Email: carlw@mcsr-labs.org
URI:

