

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.

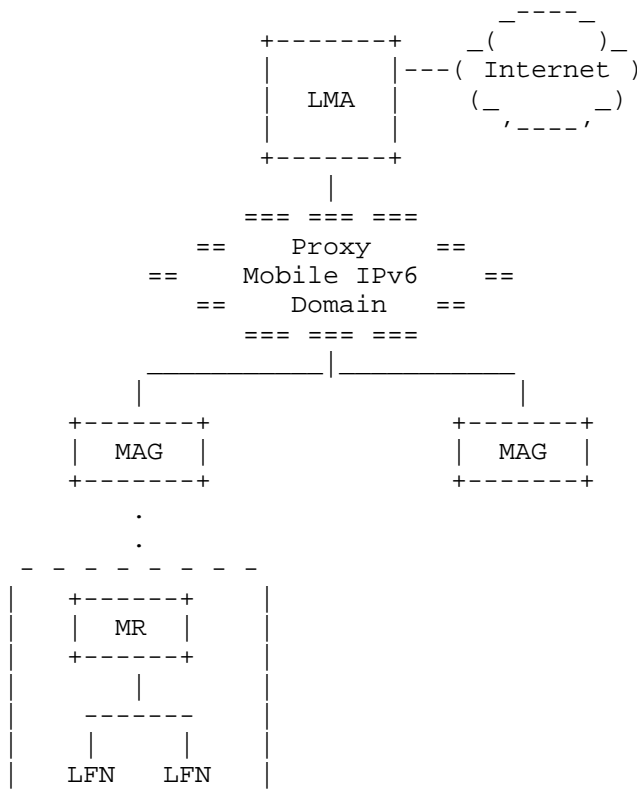


Figure 1: Mobile Router in Proxy Mobile IPv6 Domain

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.

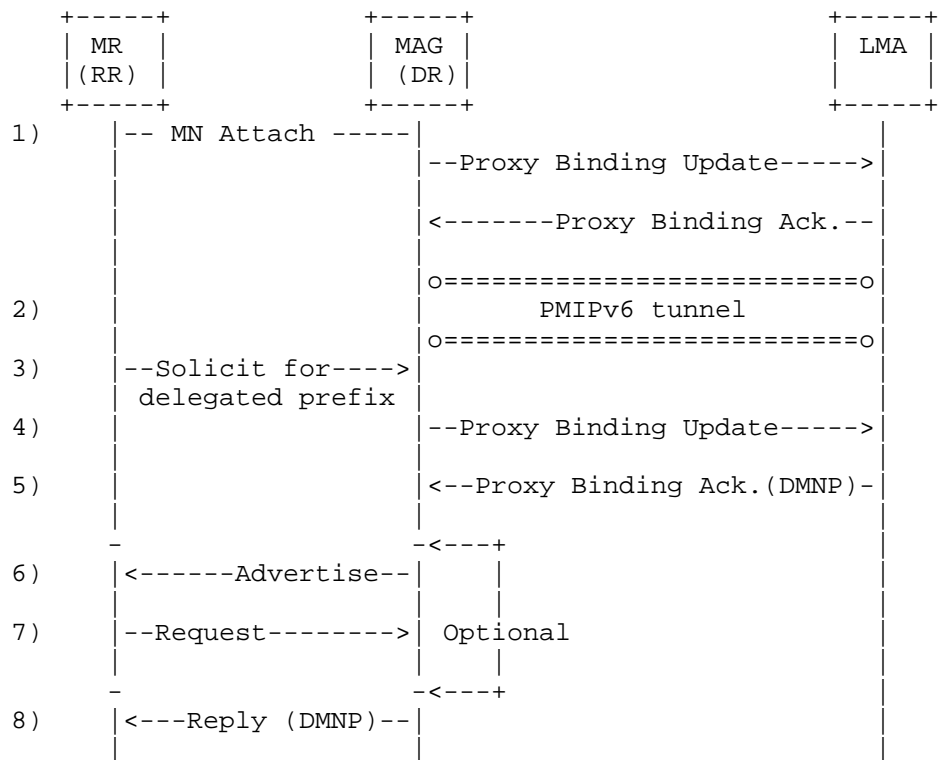


Figure 2: Delegating Router co-located with Mobile Access Gateway

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.

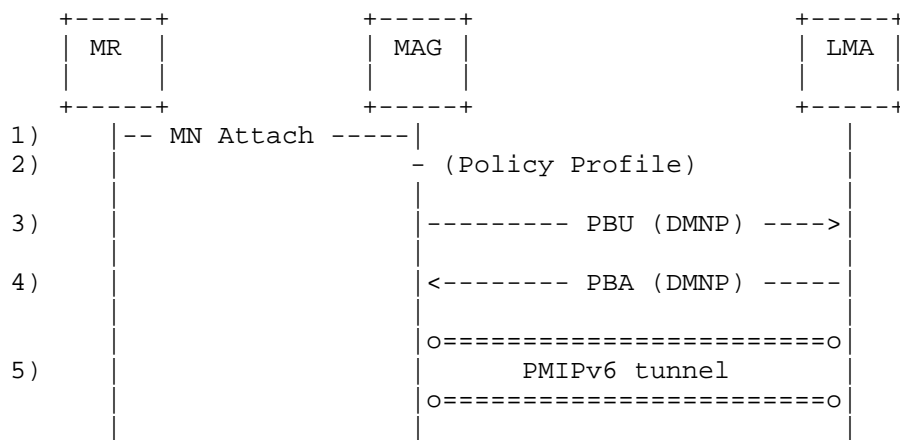


Figure 4: Static Configuration of Delegated Mobile Network Prefixes


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 Hoefft 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.

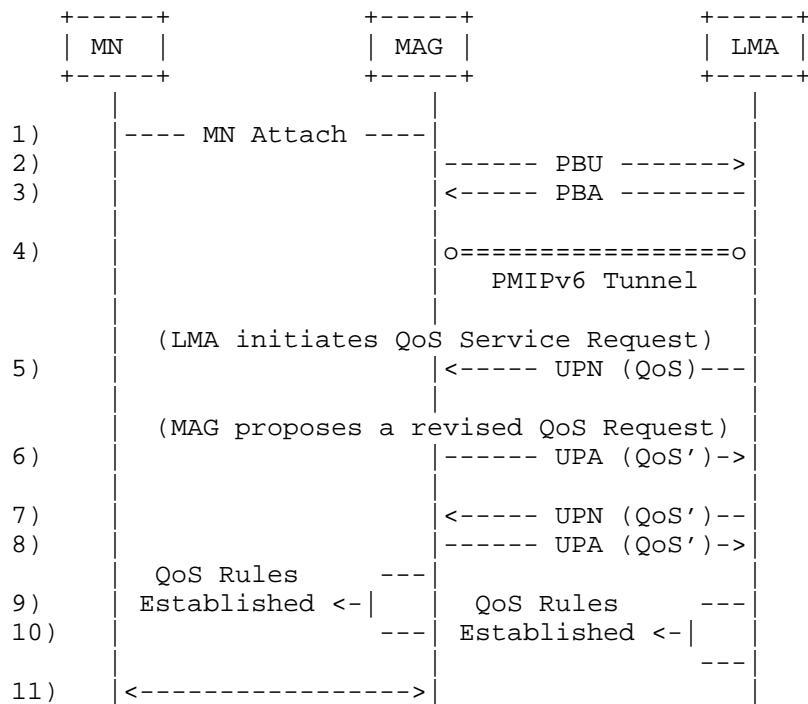


Figure 2: LMA Initiated QoS Service Request

- 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.



- 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.



- 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.

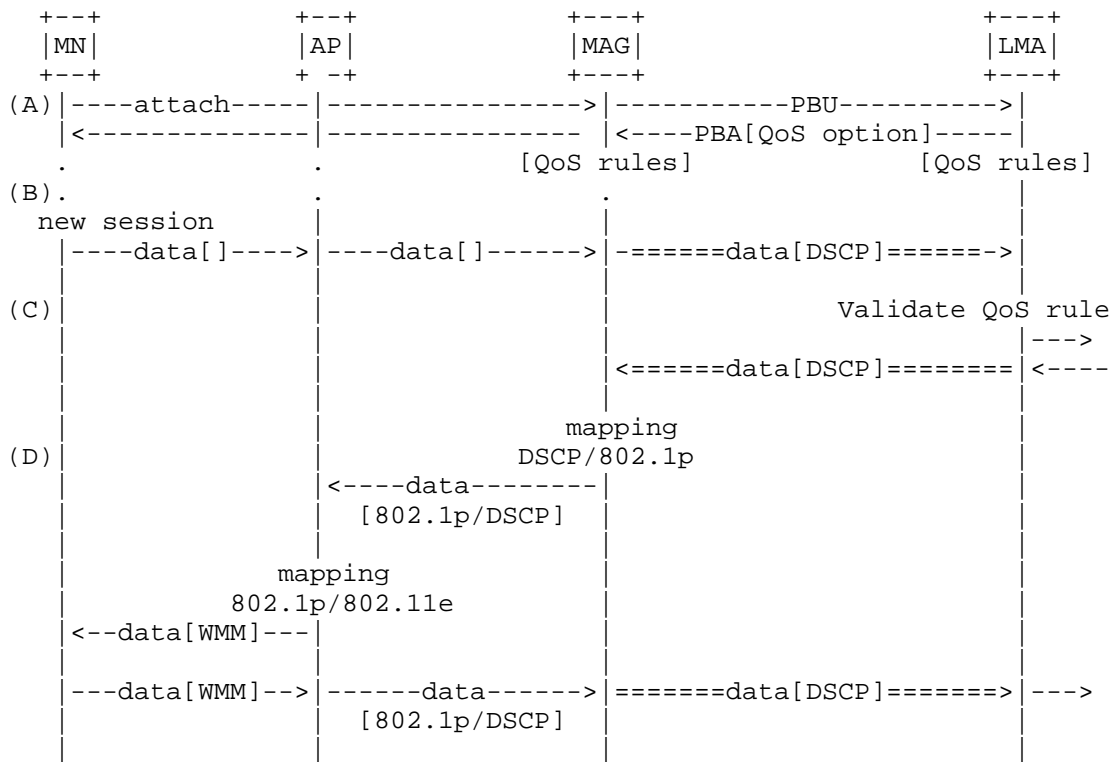


Figure 14: Prioritization of a flow created on the WLAN access

Appendix C. Information when implementing with a Broadband Network Gateway

This section shows an example of QoS interworking between the PMIPv6 domain and the broadband access. The Broadband Network Gateway (BNG) or Broadband Remote Access Server (BRAS) has the MAG function and the CPE (Customer Premise Equipment) or Residential Gateway (RG) is connected via the broadband access network. The MN is attached to the RG via e.g., Wi-Fi AP in the broadband home network. In the segment of the broadband access network, the BNG and RG are the Policy Enforcement Point (PEP) for the downlink and uplink traffic, respectively. The QoS information is downloaded from the LMA to the BNG via the PMIPv6 with the QoS option defined in this document. Based on the received QoS parameters (e.g., DSCP values), the broadband access network and the RG provide appropriate QoS treatment to the downlink and uplink traffic to/from the MN.




Figure 15: End-to-end QoS management with the broadband access network

In the segment of the broadband access network, QoS mapping between 3GPP QCI values and DSCP described in Section 6.2 is applied. In the segment of the broadband home network, if the MN is attached to the RG via Wi-Fi, the same QoS mapping as described in Appendix B can be applied.

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.

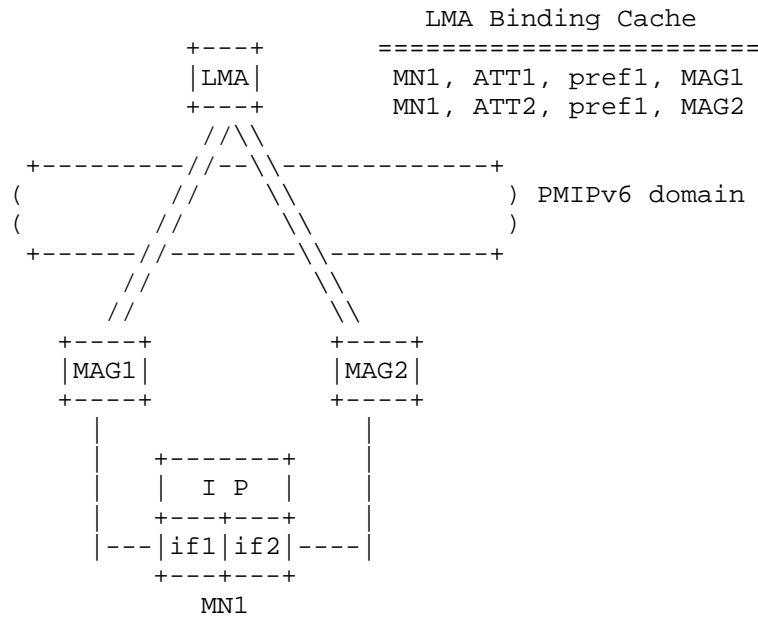


Figure 1: Shared prefix across physical interfaces scenario

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.

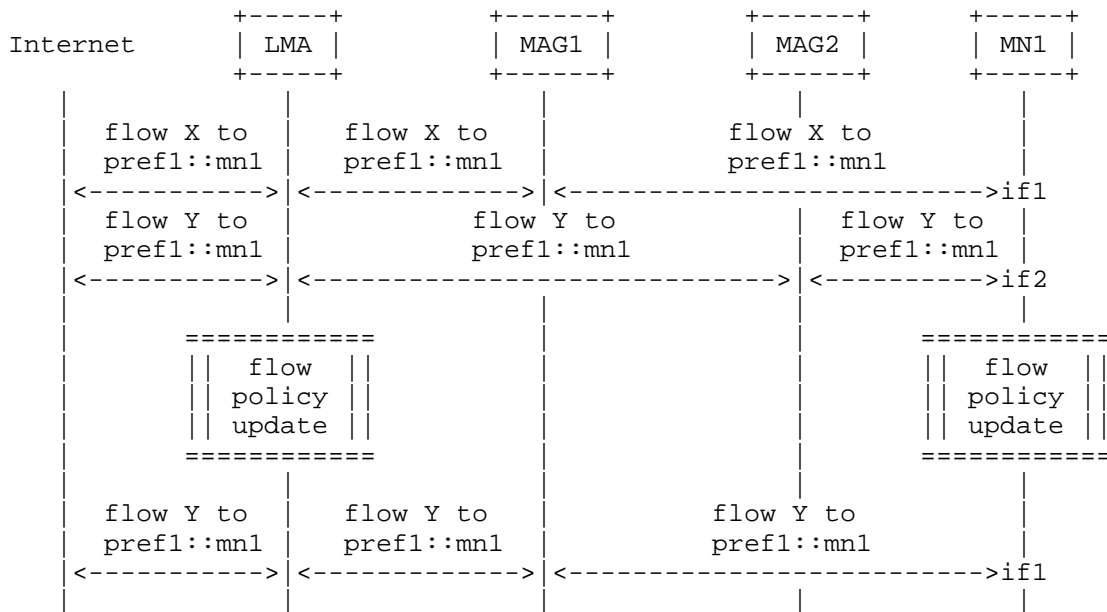


Figure 2: Flow mobility message sequence with common set of prefixes

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

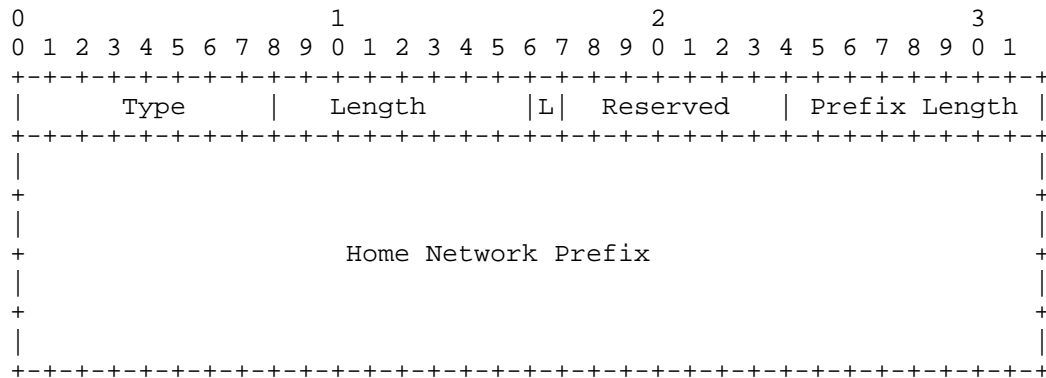
This section defines modifications to the Proxy Mobile IPv6 [RFC5213] protocol messages.

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/>

Netext
Internet-Draft
Intended status: Informational
Expires: July 9, 2015

Ravi. Valmikum
Unaffiliated
Rajeev. Koodli
Intel
January 5, 2015

EAP Attributes for Wi-Fi - EPC Integration
draft-ietf-netext-wifi-epc-eap-attributes-16

Abstract

With Wi-Fi emerging as a crucial access network for mobile service providers, it has become important to provide functions commonly available in 3G and 4G networks in Wi-Fi access networks as well. Such functions include Access Point Name (APN) Selection, multiple Packet Data Network (PDN) connections, and seamless mobility between Wi-Fi and 3G/4G networks.

The EAP-AKA (and EAP-AKA') protocol is required for mobile devices to access the mobile Evolved Packet Core (EPC) via Wi-Fi networks. This document defines a few new EAP attributes to enable the above-mentioned functions in such networks. The attributes are exchanged between a client (such as a Mobile Node) and its network counterpart (such as a AAA server) in the service provider's infrastructure.

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 July 9, 2015.

Copyright Notice

Copyright (c) 2015 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	2
1.1. APN Selection	3
1.2. Multiple APN Connectivity	4
1.3. Wi-Fi to E-UTRAN mobility	4
2. Terminology	4
3. Protocol Overview	4
3.1. Brief Introduction to EAP	4
3.2. IEEE 802.11 Authentication using EAP over 802.1X	5
4. New EAP Attributes	7
4.1. APN Selection	7
4.2. Connectivity Type	7
4.3. Wi-Fi to UTRAN/E-UTRAN Mobility	7
4.4. MN Serial ID	8
5. Attribute Extensions	8
5.1. AT_VIRTUAL_NETWORK_ID	8
5.2. AT_VIRTUAL_NETWORK_REQ	9
5.3. AT_CONNECTIVITY_TYPE	10
5.4. AT_HANDOVER_INDICATION	11
5.5. AT_HANDOVER_SESSION_ID	11
5.6. AT_MN_SERIAL_ID	12
6. Security Considerations	13
7. IANA Considerations	14
8. Acknowledgment	15
9. References	15
9.1. Normative References	15
9.2. Informative References	16
Appendix A. Change Log	17
Authors' Addresses	18

1. Introduction

Wi-Fi has emerged as a "trusted" access technology for mobile service providers; see [EPC2] for reference to the 3GPP description of "trusted" access. Advances in IEEE 802.11u [IEEE802.11u] and "HotSpot 2.0" [hs20] have enabled seamless roaming, in which a Mobile Node can select and connect to a Wi-Fi access network just as it

would roam into a cellular network. It has thus become important to provide certain functions in Wi-Fi which are commonly supported in licensed-spectrum networks such as 3G and 4G networks. This draft specifies a few new EAP attributes for a Mobile Node (MN) to interact with the network to support some of these functions (see below). These new attributes serve as a trigger for Proxy Mobile IPv6 network nodes to undertake the relevant mobility operations. For instance, when the Mobile Node requests and the network agrees for a new IP session (i.e., a new Access Point Name or APN in 3GPP), the corresponding attribute (defined below) acts as a trigger for the Mobile Anchor Gateway (MAG) to initiate a new mobility session with the Local Mobility Anchor (LMA). This document refers to [RFC6459] for the basic definitions of mobile network terminology (such as APN) used here.

The 3rd Generation Partnership Project (3GPP) networks support many functions that are not commonly implemented in a Wi-Fi network. This document defines EAP attributes that enable the following functions in Wi-Fi access networks using EAP-AKA' [RFC5448] and EAP-AKA [RFC4187]:

- o APN Selection
- o Multiple APN Connectivity
- o Wi-Fi to 3G/4G (UTRAN/EUTRAN) mobility

The attributes defined here are exchanged between the Mobile Node and the EAP server, typically realized as part of the AAA server infrastructure in a service provider's infrastructure. In particular, the Wi-Fi access network simply conveys the attributes to the service provider's core network where the EAP processing takes place [EPC]. Since these attributes share the same IANA registry, the methods are applicable to EAP-AKA', EAP-AKA, EAP-SIM [RFC4186] and, with appropriate extensions, are possibly applicable for other EAP methods as well. In addition to the trusted Wi-Fi access networks, the attributes are applicable to any trusted "non-3GPP" access network that uses the EAP methods and provides connectivity to the mobile EPC, which provides connectivity for 3G, 4G, and other non-3GPP access networks [EPC2].

1.1. APN Selection

The 3GPP networks support the concept of an APN (Access Point Name). This is defined in [GPRS]. Each APN is an independent IP network with its own set of IP services. When the MN attaches to the network, it may select a specific APN to receive desired services. For example, to receive generic Internet services, a user device may

select APN "Internet" and to receive IMS voice services, it may select APN "IMSvoice".

In a Wi-Fi access scenario, an MN needs a way of sending the desired APN name to the network. This draft specifies a new attribute to propagate the APN information via EAP. The agreed APN is necessary for the Proxy Mobile IPv6 MAG to initiate a new session with the LMA.

1.2. Multiple APN Connectivity

As an extension of APN Selection, an MN may choose to connect to multiple IP networks simultaneously. 3GPP provides this feature via additional Packet Data Protocol (PDP) contexts or additional Packet Data Network (PDN) connections, and defines the corresponding set of signaling procedures. In a trusted Wi-Fi network, an MN connects to the first APN via DHCPv4 or IPv6 Router Solicitation. This document specifies an attribute that indicates the MN's capability to support multiple APN connectivity. The specific connectivity types are also necessary for the Proxy Mobile IPv6 signaling.

1.3. Wi-Fi to E-UTRAN mobility

When operating in a multi-access network, an MN may want to gracefully handover its IP attachment from one access network to another. For instance, an MN connected to a 3GPP E-UTRAN network may choose to move its connectivity to a trusted Wi-Fi network. Alternatively, the MN may choose to connect using both access technologies simultaneously, and maintain two independent IP attachments. To implement these scenarios, the MN needs a way to correlate the UTRAN/E-UTRAN session with the new Wi-Fi session. This draft specifies an attribute to propagate E-UTRAN session identification to the network via EAP. This helps the network to correlate the sessions between the two Radio Access Network technologies and thus helps the overall handover process.

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].

3. Protocol Overview

3.1. Brief Introduction to EAP

EAP is defined as a generic protocol in [RFC3748]. EAP, combined with one of the payload protocols such as EAP-AKA' [RFC5448] can accomplish several things in a network:

- o Establish identity of the user (MN) to the network.
- o Authenticate the user during the first attach with the help of an authentication center that securely maintains the user credentials. This process is called EAP Authentication.
- o Re-authenticate the user periodically, but without the overhead of a round-trip to the authentication center. This process is called EAP Fast Re-Authentication.

This draft makes use of the EAP Authentication procedure. The use of EAP Fast Re-Authentication procedure is for further study. Both the EAP Authentication and EAP Fast Re-Authentication procedures are specified for trusted access network use in 3GPP. [TS-33.402]

3.2. IEEE 802.11 Authentication using EAP over 802.1X

In a Wi-Fi network, EAP is carried over the IEEE 802.1X Authentication protocol. The IEEE 802.1X Authentication is a transparent, payload-unaware mechanism to carry the authentication messages between the MN and the Wi-Fi network elements.

EAP, on the other hand, has multiple purposes. Apart from its core functions of communicating an MN's credentials to the network and proving the MN's identity, it also allows the MN to send arbitrary information elements to help establish the MN's IP session in the network. The following figure shows an example end-to-end EAP flow in the context of an IEEE 802.11 Wi-Fi network. We first define the terminology:

- o MN: Mobile Node
- o WAN: Wi-Fi Access Node, typically consisting of Wi-Fi Access Point and Wi-Fi Controller. In a PMIPv6 [RFC5213] network, the MAG functionality is located in the WAN, either in the Wi-Fi Access Point or in the Wi-Fi Controller.
- o AAA: The infrastructure node supporting the AAA server with the EAP methods (AKA, AKA', EAP-SIM). The end-points of the EAP method are the MN and the AAA server.
- o IPCN: IP Core Network. This includes the PMIPv6 LMA function.

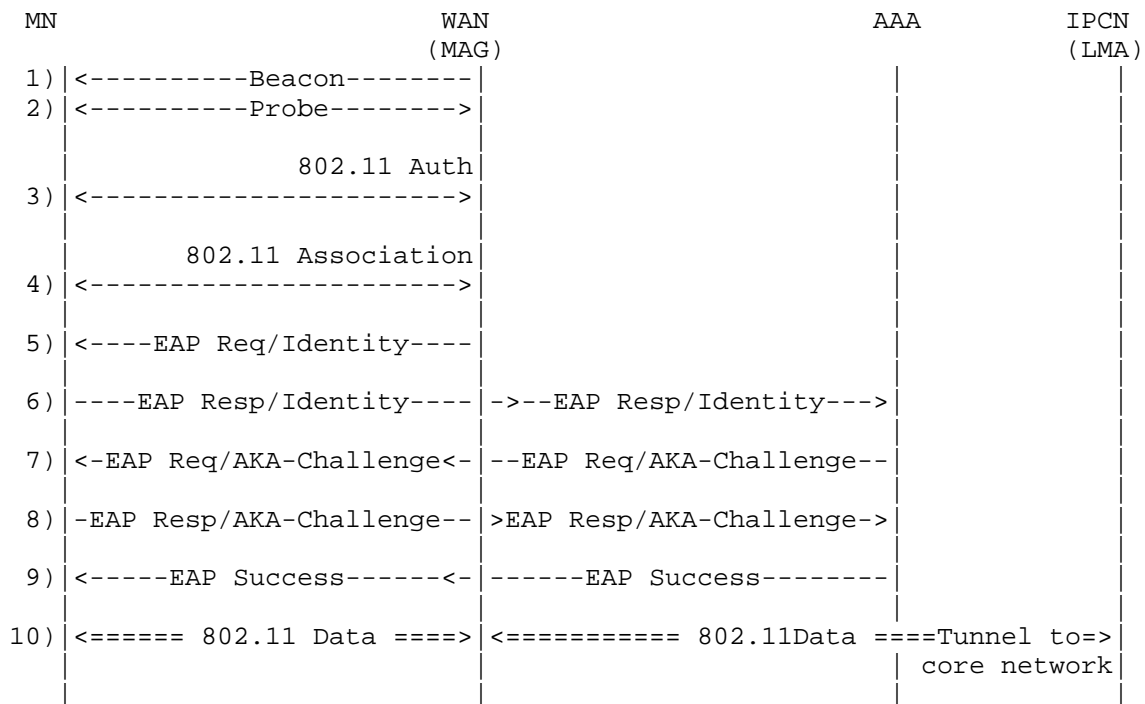


Figure 1: Example EAP Deployment

The figure shows separate Wi-Fi Access Point and Wi-Fi Access Controller, following the split-MAC model of CAPWAP [RFC5415]. A particular deployment may have the two functions within a single node.

1. An MN detects a beacon from a WAP in the vicinity.
2. The MN probes the WAP to determine suitability to attach (Verify SSID list, authentication type and so on).
3. The MN initiates the IEEE 802.11 Authentication with the Wi-Fi network. In WPA/WPA2 mode, this is an open authentication without any security credential verification.
4. The MN initiates 802.11 Association with the Wi-Fi network.
5. The Wi-Fi network initiates 802.1X/EAP Authentication procedures by sending EAP Request/Identity.
6. The MN responds with its permanent or temporary identity.

7. The Wi-Fi network challenges the MN to prove its identity by sending EAP Request/AKA-Challenge.
8. The MN calculates the security digest and responds with EAP Response/AKA-Challenge.
9. If the authentication is successful, the Wi-Fi network responds to the MN with EAP Success.
10. An end-to-End data path is available for the MN to start IP layer communication (DHCPv4, IPv6 Router Solicitation and so on).

4. New EAP Attributes

The following sections define the new EAP attributes and their usage.

4.1. APN Selection

In a Wi-Fi network, an MN includes the `AT_VIRTUAL_NETWORK_ID` attribute in the EAP-Response/AKA-Challenge to indicate the desired APN identity for the first PDN connection.

If the MN does not include the `AT_VIRTUAL_NETWORK_ID` attribute in the EAP-Response/AKA-Challenge, the network may select an APN by other means. This selection mechanism is outside the scope of this document.

An MN includes the `AT_VIRTUAL_NETWORK_REQ` attribute to indicate single or multiple PDN capability. In addition, a sub-type in the attribute indicates IPv4, IPv6, or dual IPv4v6 PDN connectivity.

4.2. Connectivity Type

An MN indicates its preference for connectivity using the `AT_CONNECTIVITY_TYPE` attribute in the EAP-Response/AKA-Challenge message. The preference indicates whether the MN wishes connectivity to the Evolved Packet Core (the so-called "EPC PDN connectivity") or Internet Offload (termed as "Non-Seamless Wireless Offload").

The network makes its decision and replies with the same attribute in the EAP Success message.

4.3. Wi-Fi to UTRAN/E-UTRAN Mobility

When a multi-access MN enters a Wi-Fi network, the following parameters are applicable in the EAP-Response/AKA-Challenge for IP session continuity from UTRAN/E-UTRAN.

- o AT_HANOVER_INDICATION: This attribute indicates to the network that the MN intends to continue the IP session from UTRAN/E-UTRAN. If a previous session can be located, network will honor this request by connecting the Wi-Fi access to the existing IP session.
- o AT_HANOVER_SESSION_ID: An MN MAY use this attribute to identify the session on UTRAN/E-UTRAN. If used, this attribute contains P-TMSI (Packet Temporary Mobile Subscriber Identity) if the previous session was on UTRAN or M-TMSI (Mobile Temporary Mobile Subscriber Identity) if the previous session was on E-UTRAN. This attribute helps the network correlate the Wi-Fi session to an existing UTRAN/E-UTRAN session.

4.4. MN Serial ID

The MN_SERIAL_ID attribute defines an MN's serial number, including International Mobile Equipment Identity (IMEI) and International Mobile Equipment Identity Software Version (IMEISV). The IMEI (or IMEISV) is used for ensuring a legitimate (and not a stolen) device is in use. As with the others, this attribute is exchanged with the service provider's AAA server. The MN_SERIAL_ID MUST NOT be propagated further by the AAA server to any other node.

5. Attribute Extensions

The format for the new attributes follows that in [RFC4187]. Note that the Length field value is inclusive of the first two bytes.

5.1. AT_VIRTUAL_NETWORK_ID

The AT_VIRTUAL_NETWORK_ID attribute identifies the virtual IP network that the MN intends to attach to. The implementation of the virtual network on the core network side is technology specific. For instance, in a 3GPP network, the virtual network is implemented based on the 3GPP APN primitive.

This attribute SHOULD be included in the EAP-Response/AKA-Challenge message.

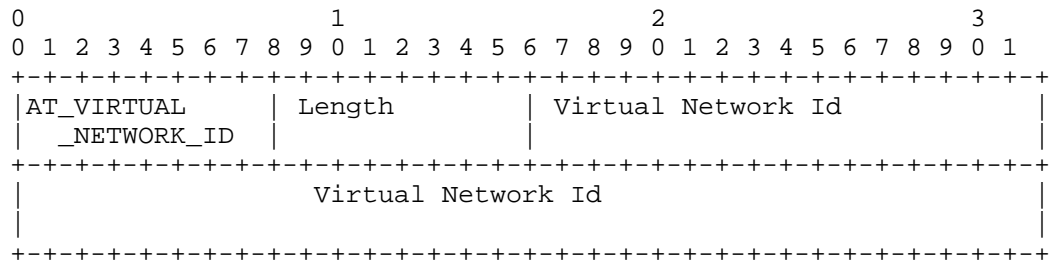


Figure 2: AT_VIRTUAL_NETWORK_ID EAP Attribute

Virtual Network Id:

An arbitrary octet string that identifies a virtual network in the access technology the MN is attaching to. For instance, in 3GPP E-UTRAN, this could be an APN. See [TS-23.003] for encoding of the field.

5.2. AT_VIRTUAL_NETWORK_REQ

When an MN intends to connect an APN, it SHOULD use this attribute to indicate different capabilities to the network. In turn, the network provides what is supported.

From the MN, this attribute can be included only in EAP-Response/Identity. From the network, it SHOULD be included in the EAP Request/AKA-Challenge message. In the MN-to-network direction, the Type field (below) indicates the MN's request. In the network-to-MN direction, the Type field indicates network's willingness to support the request; a present Type field value indicates the network support for that Type.




Figure 3: AT_VIRTUAL_NETWORK_REQ EAP Attribute

Virt-Net-Req Type:

Type can have one of the following values:

- o TBA IANA: Reserved

- o TBA IANA: Single PDN connection
- o TBA IANA : Multiple PDN connection. Can request Non-Seamless Wi-Fi Offload or EPC connectivity (see the Connectivity Type attribute below)

Virt-Net-Req Sub-type:

Sub-type can have one of the following values:

- o TBA IANA : Reserved
- o TBA IANA : PDN Type: IPv4
- o TBA IANA : PDN Type: IPv6
- o TBA IANA : PDN Type: IPv4v6

5.3. AT_CONNECTIVITY_TYPE

An MN uses this attribute to indicate whether it wishes the connectivity type to be Non-Seamless WLAN Offload or EPC. This attribute is applicable for multiple PDN connections only.

From the MN, this attribute can be included only in EAP-Response/Identity. From the network, it SHOULD be included in the EAP Request/AKA-Challenge message.




Figure 4: AT_CONNECTIVITY_TYPE EAP Attribute

Connectivity Type:

Connectivity Type can have one of the following values:

- o TBA IANA : Reserved
- o TBA IANA : Non-Seamless WLAN Offload (NSWO)
- o TBA IANA : EPC PDN connectivity

5.4. AT_HANOVER_INDICATION

This attribute indicates an MN's handover intention of an existing IP attachment.

This attribute SHOULD be included in the EAP-Response/AKA-Challenge message.




Figure 5: AT_HANOVER_INDICATION EAP Attribute

Handover Type:

- o 0 - the MN has no intention of handing over an existing IP session, i.e., the MN is requesting an independent IP session with the Wi-Fi network without disrupting the IP session with the UTRAN/E-UTRAN. In this case, no Session Id (Section 5.5) is included.
- o 1 - the MN intends to handover an existing IP session. In this case, MN MAY include a Session Id (Section 5.5) to correlate this Wi-Fi session with a UTRAN/E-UTRAN session.

5.5. AT_HANOVER_SESSION_ID

When an MN intends to handover an earlier IP session to the current access network, it may propagate a session identity that can help identify the previous session from UTRAN/E-UTRAN that the MN intends to handover. This attribute is defined as a generic octet string. The MN MAY include an E-UTRAN GUTI if the previous session was an E-UTRAN session. If the previous session was a UTRAN session, the MN MAY include UTRAN Global RNC ID (MCC, MNC, RNC Id) and P-TMSI concatenated as an octet string.

This attribute SHOULD be included in the EAP-Response/AKA-Challenge message.

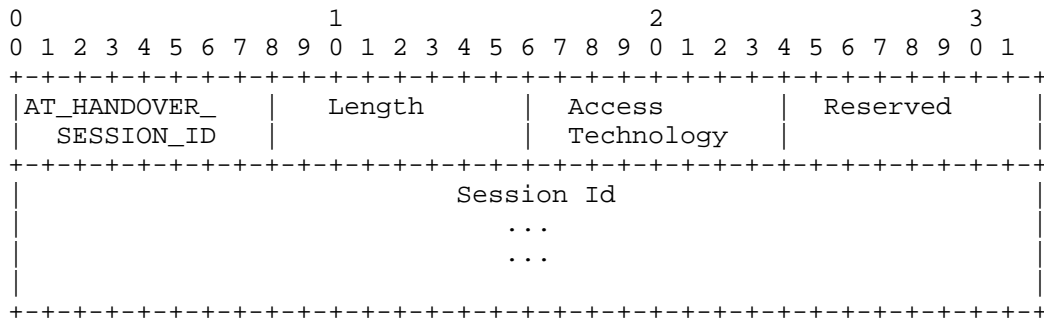


Figure 6: AT_HANOVER_SESSION_ID EAP Attribute

Access Technology:

This field represents the RAN technology from which the MN is undergoing a handover.

- o TBA IANA: Reserved
- o TBA IANA: UTRAN
- o TBA IANA: E-UTRAN

Session Id:

An octet string of variable length that identifies the session in the source access technology. As defined at the beginning of this section, the actual value is RAN technology dependent. For E-UTRAN, the value is GUTI. For UTRAN, the value is Global RNC Id (6 bytes) followed by P-TMSI (4 bytes). See [TS-23.003] for encoding of the field.

5.6. AT_MN_SERIAL_ID

This attribute defines the MN's machine serial number. Examples are International Mobile Equipment Identity (IMEI) and International Mobile Equipment Identity Software Version (IMEISV).

A network that requires the machine serial number for authorization purposes MUST send a request for the attribute in an EAP-Request/AKA-Challenge message. If the attribute is present, the MN SHOULD include the attribute in the EAP-Response/AKA-Challenge message. If the MN sends the attribute, it MUST be contained within an AT_ENCR_DATA attribute. An MN MUST NOT provide the attribute unless it receives the request from a network authenticated via EAP/AKA.

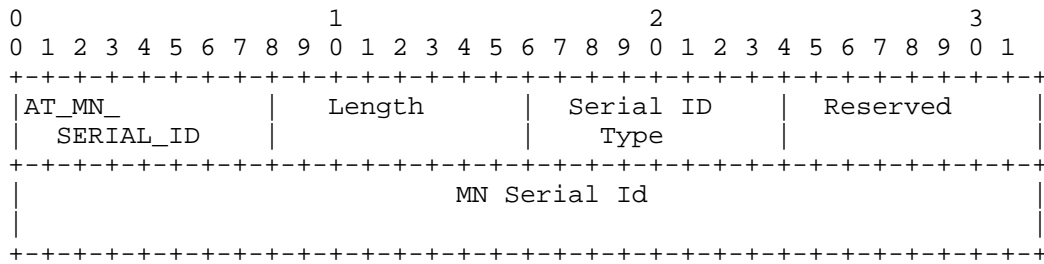


Figure 7: AT_MN_SERIAL_ID EAP Attribute

Serial ID Type:

This field identifies the type of the MN Identifier.

- o TBA IANA: Reserved
- o TBA IANA: IMEI
- o TBA IANA: IMEISV

MN Serial Id:

An arbitrary octet string that identifies the MN's machine serial number. The actual value is device-specific. See [TS-23.003] for encoding of the field. When sent by the network in the EAP-Request/ AKA-Challenge message, this field is not present, which serves as an indication for the MN to provide the attribute in the EAP-Response/ AKA-Challenge message.

AT_MN_SERIAL_ID attribute MUST only be used with methods which can provide mutual (network and device) authentication, such as AKA, AKA' and EAP-SIM

6. Security Considerations

This document defines new EAP attributes to extend the capability of the EAP-AKA protocol as specified in Section 8.2 of [RFC4187]. The attributes are passed between an MN and a AAA server in provider-controlled trusted Wi-Fi networks, where the Wi-Fi Access Network is a relay between the MN and the AAA server. The document does not specify any new messages or options to the EAP-AKA protocol.

The attributes defined here are fields which are used in existing 3G and 4G networks, where they are exchanged (in protocols specific to 3G and 4G networks) subsequent to the mobile network authentication (e.g., using the UMTS-AKA mechanism). For the operator-controlled Wi-

Fi access which is connected to the same core infrastructure as the 3G and 4G access, similar model is followed here with the EAP-AKA (or EAP-AKA', EAP-SIM) authentication. In doing so, these attribute processing, security-wise, is no worse than that in existing 3G and 4G mobile networks.

The attributes inherit the security protection (integrity, replay, and confidentiality) provided by the parameters in the AKA(') or SIM methods ; see Section 12.6 in [RFC4187]. Furthermore, RFC 4187 requires attributes exchanged in EAP-Request/AKA-Identity or EAP-Response/AKA-Identity to be integrity-protected with AT_CHECKCODE; see Section 8.2 in [RFC4187]. This requirement applies to the AT_CONNECTIVITY_TYPE and AT_VIRTUAL_NETWORK_REQ attributes defined in this document.

The AT_MN_SERIAL_ID attribute MUST have confidentiality protection provided by the AKA(') or EAP-SIM methods beyond the secure transport (such as private leased lines, VPN etc.) deployed by the provider of the trusted Wi-Fi service.

Use of identifiers such as IMEI could have privacy implications, wherein devices can be profiled and tracked. With additional information, this could also lead to profiling of user's network access patterns. Implementers should consult [hotos-2011] and references therein for a broader discussion and possible mitigation methods on the subject.

7. IANA Considerations

This document defines the following new skippable EAP-AKA attributes. These attributes need assignments from the "EAP-AKA and EAP-SIM Parameters" registry at <https://www.iana.org/assignments/eapsimakea-numbers>

- o AT_VIRTUAL_NETWORK_ID (Section 5.1) - TBA by IANA
- o AT_VIRTUAL_NETWORK_REQ (Section 5.2) - TBA by IANA
- o AT_CONNECTIVITY_TYPE (Section 5.3) - TBA IANA
- o AT_HANOVER_INDICATION (Section 5.4) - TBA by IANA
- o AT_HANOVER_SESSION_ID (Section 5.5) - TBA by IANA
- o AT_MN_SERIAL_ID (Section 5.6) - TBA by IANA

This document requests a new IANA registry "Trusted non-3GPP Access EAP Parameters". The range for both Types and Sub types in the

registry is 0 - 127, with 0 (zero) being a reserved value. The document requests IANA to make assignments in a monotonically increasing order in increments of 1, starting from 1. New assignments in this registry are made with the Specification Required policy [RFC5226].

The IANA Designated Expert should review the requirements for new assignments based on factors including, but not limited to, the source of request (e.g., standards bodies), deployment needs (e.g., industry consortium, operator community) and experimental needs (e.g., academia, industrial labs). A document outlining the purpose of new assignments should accompany the request. Such a document could be a standards document, or a research project description. The Designated Expert should consider that there is sufficient evidence of potential usage both on the end-points (e.g., Mobile Devices etc.) and the infrastructure (e.g., AAA servers, gateways etc.)

The document requests assignments from the new registry for the following fields defined in this document:

- o Virt-Net-Req Type (Section 5.2) - TBA by IANA
- o Virt-Net-Req Sub type (Section 5.2) - TBA by IANA
- o Connectivity Type (Section 5.3) - TBA IANA
- o Access Technology (Section 5.5) - TBA by IANA
- o Serial ID Type (Section 5.6) - TBA by IANA

8. Acknowledgment

Thanks to Sebastian Speicher for the review and suggesting improvements. Thanks to Mark Grayson for proposing the MN Serial ID attribute. And, thanks to Brian Haberman for suggesting a new registry.

9. References

9.1. Normative References

- [RFC4187] Arkko, J. and H. Haverinen, "Extensible Authentication Protocol Method for 3rd Generation Authentication and Key Agreement (EAP-AKA)", RFC4187, January 2006, <<http://tools.ietf.org/html/rfc4187>>.

- [RFC5226] Narten, T. and H. Alvestrand, "Guidelines for Writing an IANA Considerations Section in RFCs", BCP 26, RFC 5226, May 2008.
- [RFC6459] Korhonen, J., Soininen, J., Patil, B., Savolainen, T., Bajko, G., and K. Iisakkila, "IPv6 in 3rd Generation Partnership Project (3GPP) Evolved Packet System (EPS)", RFC 6459, January 2012.

9.2. Informative References

- [EPC] "General Packet Radio Service (GPRS); enhancements for Evolved Universal Terrestrial Radio Access Network (E-UTRAN) access, 3GPP TS 23.401 8.8.0, December 2009.", <<http://www.3gpp.org/ftp/Specs/html-info/23401.htm>>.
- [EPC2] "Architecture enhancements for non-3GPP accesses, 3GPP TS 23.402 8.8.0, December 2009.", <<http://www.3gpp.org/ftp/Specs/html-info/23402.htm>>.
- [GPRS] "General Packet Radio Service (GPRS); Service description, Stage 2, 3GPP TS 23.060, December 2006", <<http://www.3gpp.org/ftp/Specs/html-info/23060.htm>>.
- [IEEE802.11u] "802.11u-2011 - IEEE Standard for Information Technology- Telecommunications and information exchange between systems- Local and Metropolitan networks-specific requirements- Part II: Wireless LAN Medium Access Control (MAC) and Physical Layer (PHY) specifications: Amendment 9: Interworking with External Networks", , Feb 2011, <<http://standards.ieee.org/findstds/standard/802.11u-2011.html>>.
- [RFC3748] Aboba, B., Blunk, L., Vollbrecht, J., Carlson, J., and H. Levkowetz, Ed., "Extensible Authentication Protocol (EAP)", RFC3748, June 2004, <<http://www.ietf.org/rfc/rfc3748.txt>>.
- [RFC4186] Haverinen, H. and J. Salowey, "Extensible Authentication Protocol Method for Global System for Mobile Communications (GSM) Subscriber Identity Modules (EAP-SIM)", RFC 4186, January 2006.
- [RFC5213] Gundavelli, S., Leung, K., Devarapalli, V., Chowdhury, K., and B. Patil, "Proxy Mobile IPv6", RFC 5213, August 2008.

- [RFC5415] Calhoun, P., Montemurro, M., and D. Stanley, "Control And Provisioning of Wireless Access Points (CAPWAP) Protocol Specification", RFC5415, January 2009, <<http://www.ietf.org/rfc/rfc5415.txt>>.
- [RFC5448] Arkko, J., Lehtovirta, V., and P. Eronen, "Improved Extensible Authentication Protocol Method for 3rd Generation Authentication and Key Agreement (EAP-AKA')", RFC 5448, May 2009.
- [TS-23.003] "3rd Generation Partnership Project: Numbering, Addressing and Identification, 3GPP TS 23.003 12.2.0, March 2014.", , <<http://www.3gpp.org/ftp/Specs/html-info/23003.htm>>.
- [TS-33.402] "3GPP System Architecture Evolution (SAE); Security aspects of non-3GPP accesses, 3GPP TS 33.402 8.6.0, December 2009.", , <<http://www.3gpp.org/ftp/Specs/html-info/33402.htm>>.
- [hotos-2011] D. Wetherall et al., , "Privacy Revelations for Web and Mobile Apps.", Proceedings of the Hot Topics in Operating Systems (HotOS) , May 2011, <<https://www.usenix.org/legacy/events/hotos11/tech/>>.
- [hs20] "Hotspot 2.0 (Release 2) Technical Specification Package v1.0.0", , <<https://www.wi-fi.org/hotspot-20-release-2-technical-specification-package-v100>>.

Appendix A. Change Log

- o: Initial Draft
- o: v01: status to Informational, Updated References, Revised the Figure
- o: No changes from 01 to 02
- o: Per recent 3GPP updates, added the Connectivity Type attribute to allow indicating Non-Seamless WLAN Offload or EPC connectivity
- o: version-04: Revised AT_VIRTUAL_NETWORK_REQ to include 1) single PDN vs Multiple PDN connections, 2) PDN Types, and referred to NSWO Connectivity Type attribute

- o: version 05: Added AT_MN_SERIAL_ID. Revised the IANA Considerations section
- o: version 06, 07: various edits
- o: AD review revs
- o: version 09: IETF LC, Directorate review revs
- o: IANA Section revision, based on IANA interaction
- o: version 12 - clarified/revised: 1) IMEI purpose, 2) attributes requirement in PMIP6 signaling, 3) references to 802.11u, HotSpot 2.0 (seamless roaming) 4) References (normative/informative), 5) editorial corrections
- o: version 13 - revised AT_MN_SERIAL_ID processing per IESG DISCUSS
- o: version 14 -clarified usage of AT_MN_SERIAL_ID. Provided additional reference to "trusted" Wi-Fi access.
- o: version 15,16: Addressed IESG comments. Revised Figure.

Authors' Addresses

Ravi Valmikum
Unaffiliated
USA

Email: valmikum@gmail.com

Rajeev Koodli
Intel
USA

Email: rajeev.koodli@intel.com

INTERNET-DRAFT
Intended Status: Informational
Expires: August 14, 2014

John Kaippallimalil
Huawei
Rajesh S. Pazhyannur
Cisco
Parviz Yegani
Juniper
February 10, 2014

Mapping 802.11 QoS in a PMIPv6 Mobility Domain
draft-kaippallimalil-netext-pmip-qos-wifi-04

Abstract

This document provides recommendations on procedures and mapping of QoS parameters between 802.11 and PMIPv6. QoS parameters in 802.11 that reserve resources for 802.11 streams should be mapped to PMIPv6 QoS resources for IP sessions and flows. QoS reservation sequences in 802.11 should allow cases where MN initiate resource reservation, as well as cases where the network initiates resource reservation. Additionally, it should be possible for QoS parameters for PMIPv6 flows and mobility sessions to be mapped to 802.11 traffic stream reservations. The sequences and parameters to be mapped to provide a consistent behavior across 802.11 and PMIPv6 QoS are described here.

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) 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	3
1.1. Terminology	4
1.2. Definitions	5
1.3. Abbreviations	6
2. End-to-End QoS with no Admission Control	6
3. End-to-End QoS with Admission Control	8
3.1. Case A: MN Initiates QoS Request	9
3.2. Case B: Network Initiates QoS Signaling (802.11aa based)	11
3.3. Case C: Hybrid (Network Initiated for PMIP, MN initiated in 802.11)	12
3.4. Case D: Network Initiated Release	14
3.5. Case E: MN Initiated Release	16
3.6. Service Guarantees in 802.11	17
4. Mapping of QoS Parameters	17
4.1 Connection Mapping	18
4.2. QoS Class	18
4.3. Bandwidth	19
4.4. Preemption Priority	20
5. Security Considerations	20
6. IANA Considerations	21
7. References	21
7.1. Normative References	21
7.2. Informative References	21
Authors' Addresses	22
Appendix A: QoS in 802.11, PMIPv6 and 3GPP Networks	23
A.1. QoS in IEEE 802.11 Networks	23

A.2. QoS in PMIPv6 Mobility domain	23
A.3. QoS in 3GPP Networks	24

1. Introduction

802.11 networks can currently apply QoS policy by using ALG (Application Level Gateway) to detect an application (e.g. SIP signaling) and then install QoS for the corresponding IP flow on the Wireless LAN Controller (WLC)/ Access Point (AP). However, this is not a general mechanism and would require ALG or detection of application level semantics in the access to install the right QoS.

[PMIP-QoS] describes a application neutral procedure to obtain QoS for PMIPv6 flows and sessions. However, there are differences in parameters and procedures that need to be mapped between PMIPv6 QoS and 802.11. PMIPv6 has the notion of QoS for mobility sessions and flows while in 802.11 these should correspond to QoS for 802.11 data frames. Parameters in 802.11 QoS do not always have a one-to-one correspondence in PMIPv6 QoS. Further, 802.11 and PMIP QoS procedures need to be aligned based on whether QoS setup is triggered by the MN or pushed by the the network, as well as working with WMM or 802.11aa mechanisms.

This document provides information on using PMIPv6 QoS parameters for an MN connection over a 802.11 access network. The recommendations here allow for dynamic QoS policy information per Mobile Node (MN) and session to be configured by the 802.11 access network. PMIPv6 QoS signaling between MAG and LMA provisions the per MN QoS policies in the MAG. In the 802.11 access network modeled here, the MAG is located at the Access Point (AP)/ Wireless LAN Controller (WLC) . Figure 1 below provides an overview of the entities and protocols.

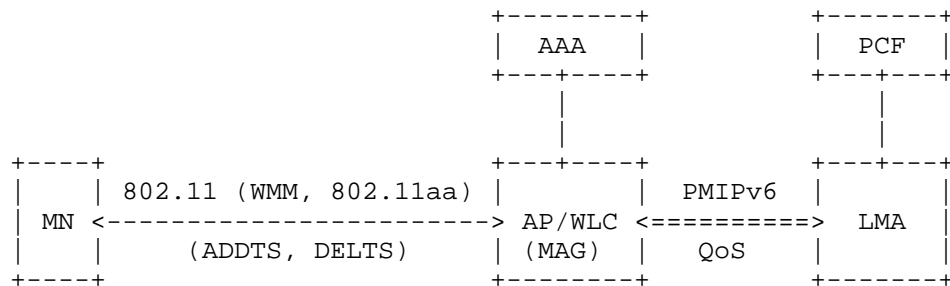


Figure 1: QoS Policy in 802.11 Access

MN and AP/WLC use 802.11 QoS mechanisms to setup admission controlled

flows. The AP/WLC is a MAG that requests for QoS policy from the LMA. The MN uses ADDTS (Add Traffic Stream) to setup QoS for a traffic stream between itself and the AP, and DELTS (Delete Traffic Stream) to delete that stream. In WMM [WMM 1.2.0], the AP advertises if admission control is mandatory for an access class. Admission control for best effort or background access classes is not recommended. In addition to WMM capability, 802.11aa allows for AP/WLC to support an ADDTS reservation request to the MN. This makes it simpler to support a PMIPv6 QoS request that is pushed to the AP/WLC.

The parameter mapping recommendations described here support the procedures by which the 3GPP network provisions QoS per application dynamically or during authorization of the Mobile Node (MN). However, the 802.11 procedures described here are not limited to work for just the 3GPP policy provisioning. If PMIPv6 QoS parameters can be provisioned on the MAG via mechanisms defined in [PMIP-QoS], the 802.11 procedures can be applied in general for provisioning QoS in a 802.11 network.

PMIPv6 QoS parameters need to be mapped to 802.11 QoS parameters. In some cases, there is no one-to-one mapping. And in other cases such as bandwidth, the values received in PMIP should be mapped to the right 802.11 parameters. This document provides recommendations to perform QoS mapping between PMIPv6 and 802.11 QoS.

[PMIP-QoS] does not explicitly describe how the QoS signaling and QoS sub-options map into corresponding signaling and parameters in the 802.11 access network. This mapping and the procedures in the 802.11 network to setup procedures are the focus of this document. The end-to-end flow spanning 802.11 access and PMIPv6 domain and the QoS parameters in both segments are described here. Thus, it provides a systematic way to map the various QoS parameters available in initial authorization, as well as setup of new sessions (such as a voice/video call). The mapping recommendations allow for proper provisioning and consistent interpretation between the various QoS parameters provided by PMIP QoS, and 802.11.

The rest of the document is organized as follows. Chapter 2 provides an overview of establishing mobility sessions with no admission control. These mechanisms are specified in [PMIP QoS] and outlined here since the mobility session established is the basis for subsequent admission controlled requests for flows. Chapter 3 describes how end to end QoS with 802.11 admission control is achieved. The mapping of parameters between 802.11 and PMIP QoS is described in Chapter 5.

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

Guaranteed Bit Rate (GBR)

GBR in a mobile network defines the guaranteed (reserved) bit rate resources of service data flow on a connection (bearer) [TS23.203].

Maximum Bit Rate (AMBR)

MBR represents the maximum bandwidth of a flow with reservation.

Aggregate Maximum Bit Rate (MBR)

AMBR represents the total bandwidth that all flows of a user is allowed. AMBR does not include flows with reservation.

Allocation Retention Priority (ARP)

ARP is used in the mobile network to determine the order in which resources for a flow may be preempted during severe congestion or other resource limitation. ARP of 1 is the highest priority while 15 is the lowest [TS23.203].

Peak Data Rate

In WMM, Peak Data Rate specifies the maximum data rate in bits per second. The Maximum Data Rate does not include the MAC and PHY overheads [WMM 1.2.0].

Mean Data Rate

This is the average data rate in bits per second. The Mean Data Rate does not include the MAC and PHY overheads [WMM1.2.0]

Minimum Data Rate

In WMM, Minimum Data Rate specifies the minimum data rate in bits per second. The Minimum Data Rate does not include the MAC and PHY overheads [WMM 1.2.0].

TSPEC

The TSPEC element in 802.11 contains the set of parameters that define the characteristics and QoS expectations of a traffic flow.

TCLAS

The TCLAS element specifies an element that contains a set of parameters necessary to identify incoming MSDU (MAC Service Data Unit) that belong to a particular TS (Traffic Stream) [802.11].

1.3. Abbreviations

3GPP	Third Generation Partnership Project
AAA	Authentication Authorization Accounting
AMBR	Aggregate Maximum Bit Rate
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
PCF	Policy Control Function
PDN-GW	Packet Data Network Gateway
QCI	QoS Class Indicator
QoS	Quality of Service
TCLAS	Type Classification
TSPEC	Traffic Conditioning Spec
WLC	Wireless Controller

2. End-to-End QoS with no Admission Control

PMIPv6 and 802.11 QoS with no admission control is specified in [PMIP QoS]. This section is provided as background here since prior to the establishment of an admission controlled flow, a mobility session as described here is established. IETF (RFC 4594) and GSMA have defined mapping between DSCP and IEEE 802.11 UP (User Priority). The AP/WLC (MAG) should be pre-configured to use the mapping from one of these specifications.

An MN that attempts to connect to a 802.11 network typically authenticates first and may have an authorization profile downloaded. The AP/WLC may use the QoS profile for the MN for policing flows. However, the network can obtain more dynamic policy that corresponds to current mobile network conditions and preferences using PMIP QoS.

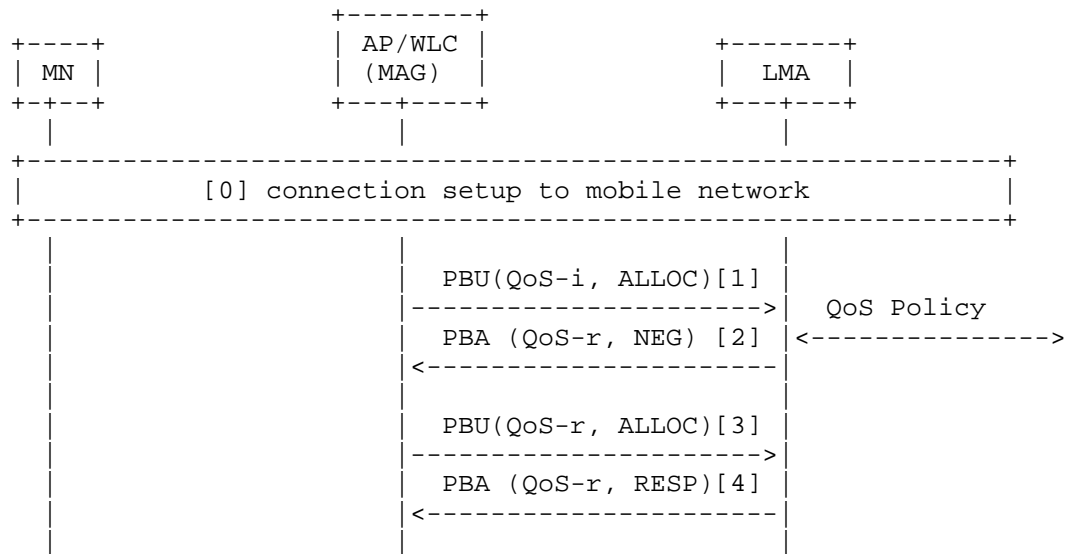


Figure 2: Default connection setup

- [0] MN signals to setup connection. The AP/WLC obtains an authorization profile that includes QoS information, or may have an administratively configured profile with QoS information.
- [1] The completion of 802.11 and IP setup serves as a trigger for the MAG (AP/WLC) to request for dynamic QoS parameters. The MAG sends a PBU containing QoS Option with operation code set to ALLOCATE, and DSCP, QoS Attributes set to initially authorized values for the MN's default connection (QoS-i).

This request is for QoS of all flows of a connectivity session of the MN and includes DSCP, Per-MN-Agg-Max-DL-Bit-Rate, Per-MN-Agg-Max-UL-Bit-Rate, Per-Session-Agg-Max-DL-Bit-Rate, Per-Session-Agg-Max-UL-Bit-Rate and Allocation-Retention-Priority fields derived from the MN initial authorization profile. The Traffic Selector field should not be present.

- [2] The LMA queries the policy server and obtains a response. The policy server may grant the QoS requested or may change the QoS levels based on network or other dynamic conditions (QoS-r in figure). This example assumes that the LMA cannot provide the QoS requested by the MAG.

The LMA sets the operational code to NEGOTIATE and responds with downgraded parameters for DSCP, Per-MN-Agg-Max-DL-Bit-Rate, Per-MN-Agg-Max-UL-Bit-Rate, Per-Session-Agg-Max-DL-Bit-Rate, Per-

Session-Agg-Max-UL-Bit-Rate and Allocation-Retention-Priority. The Traffic Selector field is not present since the provisioning applies to the entire PMIPv6 connectivity session.

[3] The MAG receives the downgraded QoS and sends a revised PBU with the QoS options that the LMA is prepared to offer. The operational code is set to ALLOCATE.

[4] The LMA can accept the requested QoS. The LMA sends a PBA message with the revised QoS options and operational code set to RESPONSE.

The new QoS values will be used by the MAG to police flows of the MN and will supercede earlier (or initially) provisioned QoS values. MAG polices session flows to not exceed Per-Session-Agg-Max-DL-Bit-Rate, Per-Session-Agg-Max-UL-Bit-Rate. If there are multiple sessions, the total bandwidth should not exceed Per-MN-Agg-Max-DL-Bit-Rate, Per-MN-Agg-Max-UL-Bit-Rate.

3. End-to-End QoS with Admission Control

This section outlines a few use cases to illustrate how parameters and mapping are applied for flows that require admission control. These cases illustrate the various provisioning sequences and mechanisms. It is not intended to be exhaustive.

The general procedure here is that a flow that requires admission control is part of a PMIPv6 connectivity session. QoS options for the overall session are provisioned as described in section 2. As a result of some application layer signaling, specific flows of the application may require admission controlled QoS which can be provisioned on a per flow basis.

There are two main types of interaction possible to provision QoS for flows that require admission control - one case is where the MN initiates the QoS request and the network provisions the resources. The second is where the network provisions resources as a result of some out of band signaling (like application signaling). In the second scenario, if the MN supports 802.11aa, the network can push the QoS configuration to the MN. If the MN only supports WMM QoS, then MN requests for QoS for the 802.11 segment and the MAG provisions based on QoS already provisioned for the MN. These three cases are described in sections 3.1 - 3.3.

In each of the sequences, QoS parameters need to be mapped between 802.11 and PMIPv6. The table below provides an overview of the mapping for establishing QoS for an admission controlled flow.

Further details of the parameters and mappings are provided in section 4.

MN <--> AP/WLC(802.11)	AP/WLC(MAG) <--> LMA PMIPv6
(TCLAS) TCP/UDP IP	Traffic Selector (IP flow)
(TCLAS) User Priority	DSCP
(TSPEC)Minimum Data Rate, DL	Guaranteed-DL-Bit-Rate
(TSPEC)Minimum Data Rate, UL	Guaranteed-UL-Bit-Rate
(TSPEC)Mean Data Rate UL/DL	-
(TSPEC)Peak Data Rate, DL	Aggregate-Max-DL-Bit-Rate
(TSPEC)Peak Data Rate, UL	Aggregate-Max-UL-Bit-Rate

Table 1: 802.11 - PMIPv6 QoS Parameter Mapping

3.1. Case A: MN Initiates QoS Request

During an MN flow setup that requires admission control in the 802.11 network, QoS parameters for the flow needs to be provisioned. This procedure outlines the case where the MN is configured (e.g. in SIM) to start the QoS signaling. In this case, the MN sends an ADDTS request indicating the QoS required for the flow. The AP/WLC (MAG) obtains the corresponding level of QoS to be granted to the flow by PMIPv6 PBU/PBA sequence with QoS options with the LMA. Details of the QoS provisioning for the flow are described below.




Figure 3: MN initiated QoS setup

[0] The MN has a best effort connectivity session as described in section 2. This allows the MN to perform application level signaling and setup.

[1] The trigger for MN to request QoS is an upper layer notification. This may be the result of end-to-end application signaling and setup procedures (e.g. SIP)

If the MN is configured to start QoS signaling, the MN sends an ADDTS request with TSPEC and TCLAS identifying the flow for which QoS is requested. The TSPECs for both uplink and downlink in this request should contain the Minimum Data Rate and Peak Data Rate .

[2] If there are sufficient resources at the AP/WLC to satisfy the request, the MAG (AP/WLC) sends a PBU with QoS options, operational code ALLOCATE and Traffic Selector identifying the flow. The Traffic selector is derived from the TCLAS to identify the flow requesting QoS. 802.11 QoS parameters in TSPEC are mapped to PMIPv6 parameters. The mapping of TCLAS and TSPEC parameters to PMIPv6 is shown in Table 1.

[3] The LMA obtains the authorized QoS for the flow and responds to the MAG with operational code set to RESPONSE. Mapping of PMIPv6

parameters to 802.11 TSPEC and TCLAS is shown in Table 1.

In networks like 3GPP, the reserved bandwidth for flows are accounted separately from the non-reserved session bandwidth. The Traffic Selector identifies the flow for which the QoS reservations are made.

- [4] The AP/WLC (MAG) provisions the corresponding QoS and replies with ADDTS Response containing authorized QoS in TSPEC and flow identification in TSPEC.

The AP/WLC polices these flows according to the QoS provisioning.

3.2. Case B: Network Initiates QoS Signaling (802.11aa based)

In some cases (e.g. LTE/SAE), the policy server in the network may be configured to initiate the policy reservation request for a flow. This use case illustrates how an MN and 802.11 network that support 802.11aa can provision QoS to flows of the MN that when the policy server pushes the reservation request.




Figure 4: Network initiated QoS setup with 802.11aa

- [0] The MN sets up best effort connectivity session as described in Case A. This allows the MN to perform application level

signaling and setup.

- [1] The policy server sends a QoS reservation request to the LMA. This is usually sent in response to an application that requests the policy server for higher QoS for some of its flows.

The LMA reserves resources for the flow requested.

- [2] LMA sends PMIP UPN (Update Notification) to the MAG with QoS parameters for the flow for which the LMA reserved resources in step [1]. In UPN, the operational code in QoS option is set to ALLOCATE and the Traffic Selector identifies the flow for QoS.

The LMA QoS parameters include Guaranteed-DL-Bit-Rate/Guaranteed-UL-Bit-Rate and Aggregate-Max-DL-Bit-Rate/Aggregate-Max-UL-Bit-Rate for the flow. In networks like 3GPP, the reserved bandwidth for flows are accounted separately from the non-reserved session bandwidth.

- [3] If there are sufficient resources to satisfy the request, the AP/WLC (MAG) sends an ADDTS Reserve Request (802.11aa) specifying the QoS reserved for the traffic stream including TSPEC and TCLAS element mapped from PMIP QoS Traffic Selector to identify the flow.

PMIPv6 parameters are mapped to TCLAS and TSPEC as shown in Table 1.

If there are insufficient resources at the AP/WLC, the MAG will not send an ADDTS message and will continue processing of step [5].

- [4] MN accepts the QoS reserved in the network and replies with ADDTS Reserve Response.
- [5] The MAG (AP/WLC) replies with UPA confirming the acceptance of QoS options and operational code set to RESPONSE. The AP/WLC police flows based on the new QoS.

If there are insufficient resources at the AP/WLC, the MAG sends a response with UPA status code set to CANNOT_MEET_QOS_SERVICE_REQUEST.

3.3. Case C: Hybrid (Network Initiated for PMIP, MN initiated in 802.11)

This use case outlines a scenario where an MN attaches to the 802.11

and then obtains services in the mobile network. When the MN attaches, PMIP signaling between the MAG and LMA establishes mobile connection and related QoS. Subsequently, the MN starts an application that requires dedicated bandwidth resources and signals that using TSPEC/ADDTS request. The details of this sequence are described below.




Figure 5: Network initiated QoS setup with WMM

- [0] The MN sets up best effort connectivity session as described in Case A. This allows the MN to perform application level signaling and setup.
- [1] The policy server sends a QoS reservation request to the LMA. This is usually sent in response to an application that requests the policy server for higher QoS for some of its flows.

The LMA reserves resources for the flow requested.

- [2] LMA sends PMIP UPN (Update Notification) to the MAG with QoS option operational code set to ALLOCATE and QoS parameters for which the LMA reserved resources in step [1]. In UPN, the Traffic selector field in QoS Option identifies the flow for QoS.

The LMA QoS parameters include Guaranteed-DL-Bit-

Rate/Guaranteed-UL-Bit-Rate and Aggregate-Max-DL-Bit-Rate/Aggregate-Max-UL-Bit-Rate for the flow. In networks like 3GPP, the reserved bandwidth for flows are accounted separately from the non-reserved session bandwidth. This is indicated by using the Traffic Selector in PMIPv6 QoS.

- [3] If there are sufficient resources to satisfy the request, the MAG (AP/WLC) replies with UPA confirming the acceptance of QoS options and operation code set to RESPONSE. If there are insufficient resources at the AP/WLC, the MAG may send a response with UPA status code set to CANNOT_MEET_QOS_SERVICE_REQUEST.

The AP/WLC can police flows based on the new QoS. However, the AP/WLC does not initiate QoS reservation signaling on 802.11 because either it or the MN does not support 802.11aa.

- [4] The trigger for the MN to request QoS is an upper layer notification. This may be the result of end-to-end application signaling and setup procedures (e.g. SIP)

The MN sends an ADDTS request with TSPEC and TCLAS identifying the flow for which QoS is requested. The TSPECs for both uplink and downlink in this request should contain the Minimum Data Rate and Peak Data Rate. The MAG maps PMIPv6 parameters obtained earlier as shown in Table 1.

If the MN supports only WMM QoS, TCLAS is not sent. The AP/WLC may identify the flow based on connection signaling (e.g. 3GPP 23.402, WCS), most recent updates from PMIP QoS (i.e. that in message [3] above), or some combination thereof.

- [5] The AP/WLC (MAG) provisions the corresponding QoS and replies with ADDTS Response containing authorized QoS in TSPEC.

The AP/WLC (MAG) may revise the offer to the MN based on PMIPv6 QoS reservation.

3.4. Case D: Network Initiated Release

QoS resources reserved for a session are released on completion of the session. When the application session completes, the policy server, or the MN may signal for the release of resources. In this use case, the network initiates the release of QoS resources.

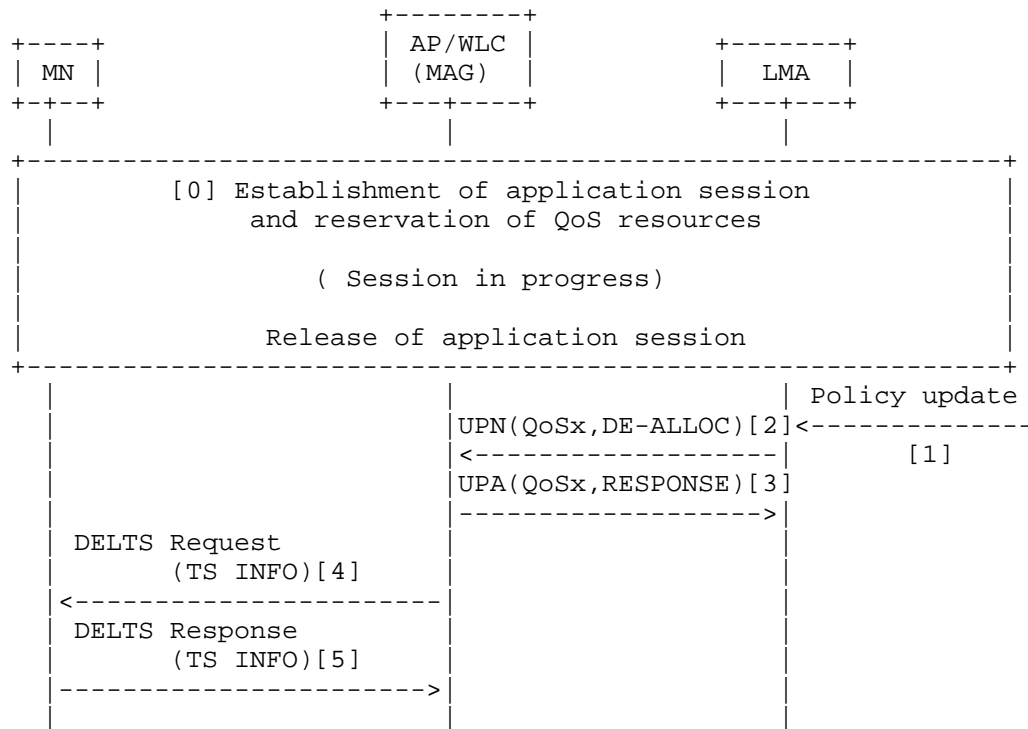


Figure 6: Network initiated QoS resource release

- [0] The MN establishes and reserves QoS resources as in use cases A, B or C.
When the application session terminates, the policy server receives notification that the session has terminated.
- [1] LMA receives a policy update indicating that QoS for flow (QoSx) should be released. The LMA releases local resources associated with the flow.
- [2] LMA sends a UPN with QoS options - Traffic Selector field identifying the flow for which QoS resources are to be released, and operation code set to DE-ALLOCATE. No additional LMA QoS parameters are sent.
- [3] MAG replies with UPA confirming the acceptance and operation code set to RESPONSE.
- [4] AP/WLC (MAG) releases local QoS resources associated with the flow. AP/WLC derives the corresponding 802.11 Traffic Stream from the PMIPv6 Traffic Selector. The AP sends a DELTS Request

with TS INFO identifying the reseravtion.

[5] MN sends DELTS Response confirming release.

Since the MN has completed the session, it may send a DELTS to explicitly request release QoS resources at AP. If the AP and MN are 802.11aa capable, the release of resources may also be signaled to the MN.

3.5. Case E: MN Initiated Release

QoS resources reserved for a session are released on completion of the session. When the application session completes, the policy server, or the MN may signal for the release of resources. In this use case, the network initiates the release of QoS resources.

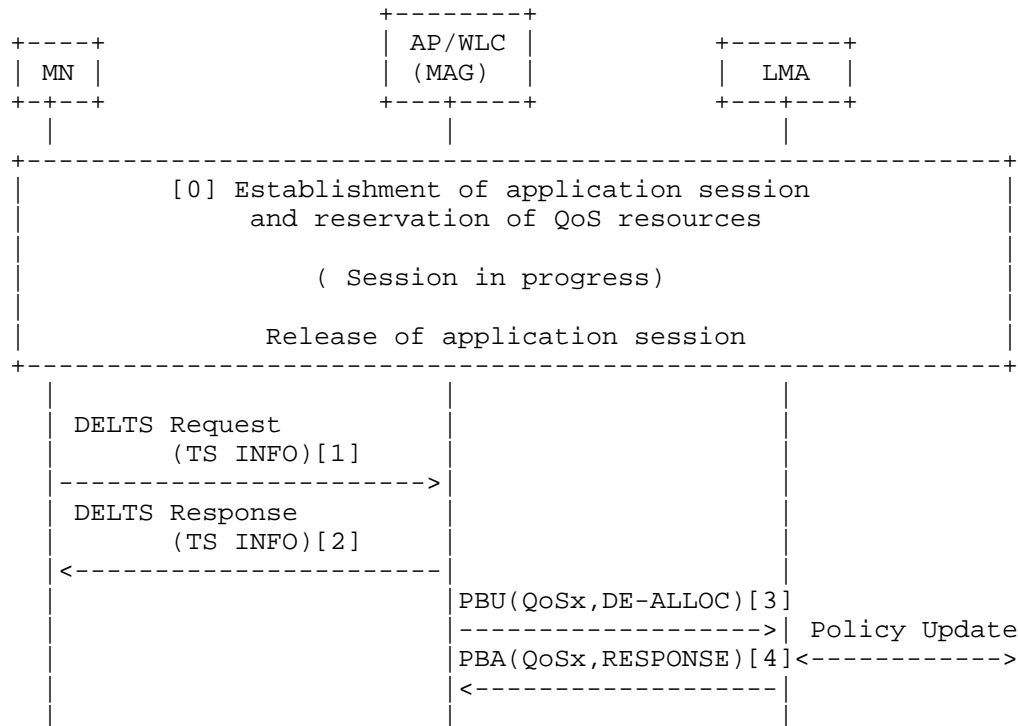


Figure 6: Network initiated QoS resource release

[0] The MN establishes and reserves QoS resources as in use cases A, B or C.

When the application session terminates, the MN prepares to release QoS resources.

- [1] MN releases its own internal resources and sends a DELTS Request to the AP/WLC with TS (Traffic Stream) INFO.
- [2] AP/WLC receives the DELTS request, releases local resources and responds to MN with a DELTS response.
- [3] AP/WLC (MAG) initiates a PBU with Traffic Selector constructed from TCLAS and PMIPv6 QoS parameters from TSPEC (QoSx) as shown in Table 1.
- [4] LMA receives the PBU, releases local resources and informs policy server. The LMA then responds with a PBA.

3.6. Service Guarantees in 802.11

The GBR - Guaranteed Bit Rate in mobile networks are used to request and commit resources in the network for providing the bandwidth requested. In 802.11 networks, a random backoff timer based on the access class only provides priority access to a shared medium. These mappings and recommendations allow the AP to schedule resources in a fair manner based on subscribed QoS and application request/policy server interaction.

However, there are no guaranteed or committed resources in the 802.11 network - only prioritization that gives better opportunity for frames to compete for a shared medium.

It should also be noted that unlike mobile networks which inform the MN about QoS for established or modified connections (bearers), there is no means for an MN in 802.11 networks to find out the QoS that a policy server requests to be granted. Thus, the application in MN should make its determination to downgrade a request based on SDP and media parameters to downgrade to a lower quality.

4. Mapping of QoS Parameters

This section outlines the handling of QoS parameters between 802.11 and PMIP QoS. 802.11 QoS reservations are made for an MN's data frames. PMIP QoS provisioning on the other hand is for IP sessions and flows. Parameters in PMIP QoS and 802.11 also need to be mapped according to the recommendations below.

4.1 Connection Mapping

TSPEC in 802.11 is used to reserve QoS for a traffic stream (MN MAC, TS(Traffic Stream) id). The QoS reservation is for 802.11 frames associated with an MN's MAC address. TCLAS element with Classifier 1 (TCP/UDP Parameters) should be used to identify a flow. The flow definition should use the specification in [PMIP-QoS] Traffic Selector. Thus, there is a one-to-one mapping between the TCLAS defined flow and that in Traffic Selector.

When an 802.11 QoS reservation is complete, it is identified by a Traffic Stream (TS) identifier. This corresponds to the flow in PMIPv6 Traffic Selector, and identified in TCLAS. For releasing QoS resources identified by a PMIPv6 Traffic selector, the AP/WLC uses the above relationship to determine the corresponding TS identifier to be sent in the DELTS request.

If the MN or AP/WLC is not able to convey TCLAS, the AP/WLC should use out of band methods to determine the IP flow for which QoS is requested. This includes correlation with connection signaling protocols (e.g. 3GPP 23.402 WCS) and Traffic Selector in most recent PMIP QoS updates.

4.2. QoS Class

Table 1 contains a mapping between Access Class (WMM AC) and 802.1D in 802.11 frames, and DSCP in IP data packets. The table also provides the mapping between Access Class (WMM AC) and DSCP for use in 802.11 TSPEC and PMIP QoS reservations.

QCI	DSCP	802.1D UP	WMM AC	Example Services
1	EF	6(VO)	3 AC_VO	conversational voice
2	EF	6(VO)	3 AC_VO	conversational video
3	EF	6(VO)	3 AC_VO	real-time gaming
4	AF41	5(VI)	2 AC_VI	buffered streaming
5	AF31	4(CL)	2 AC_VI	signaling
6	AF32	4(CL)	2 AC_VI	buffered streaming
7	AF21	3(EF)	0 AC_BE	interactive gaming
8	AF11	1(BE)	0 AC_BE	web access
9	BE	0(BK)	1 AC_BK	e-mail

Table 2: QoS Mapping between QCI/DSCP, 802.1D UP, WMM AC

The MN tags data packets with DSCP and 802.1D UP corresponding to the application and the subscribed policy or authorization. The AP/WLC polices sessions and flows based on these values and the QoS policy

for the MN.

For QoS reservations, TSPEC use WMM AC values and PMIP QoS uses corresponding DSCP values in Traffic Selector. 802.11 QoS Access Class AC_VO, AC_VI are used for QoS reservations. AC_BE, AC_BK should not be used in reservations.

4.3. Bandwidth

There are bandwidth parameters that need to be mapped for admission controlled flows and others for non-admission controlled flows.

Non-Admission Controlled Flows:

Flows and sessions that do not need QoS reservation have no need for equivalent mapping for 802.11. These sessions and flows are policed by the AP/WLC to ensure that QoS policy obtained initially (during MN authorization) or dynamically over PMIP QoS is not exceeded by the MN.

All connection sessions of the MN should not in total exceed Per-MN-Agg-Max-DL-Bit-Rate and Per-MN-Agg-Max-UL-Bit-Rate in the downlink and uplink directions respectively. The non-admission controlled flows of a single connectivity session of an MN should not exceed Per-Session-Agg-Max-DL-Bit-Rate and Per-Session-Agg-Max-UL-Bit-Rate in the downlink and uplink directions respectively.

Admission Controlled Flows:

For flows that require reservation, the 802.11 Minimum Data Rate should be equal to Guaranteed Bit Rate (GBR). If the MN requests Minimum Data Rate in ADDTS greater than GBR, then AP/WLC should reject the admission request in ADDTS Response.

MN <--> AP/WLC(802.11)	AP/WLC(MAG) <--> LMA PMIPv6
Minimum Data Rate, DL	Guaranteed-DL-Bit-Rate
Minimum Data Rate, UL	Guaranteed-UL-Bit-Rate
Mean Data Rate UL/DL	[a]
Peak Data Rate, DL	Aggregate-Max-DL-Bit-Rate
Peak Data Rate, UL	Aggregate-Max-UL-Bit-Rate

NOTE[a] AP/WLC may derive Mean Data Rate from Minimum and Maximum Data Rates. There is no equivalent parameter in PMIP QoS.

Table 3: Bandwidth Parameters for Admission Controlled Flows

During the QoS reservation procedure, if the MN requests Minimum Data Rate, or other parameters in excess of values authorized in PMIP QoS, the AP/WLC should deny the request in ADDTS Response. Bandwidth of admission controlled flows are policed according to the mappings in Table 2.

4.4. Preemption Priority

Mobile networks with resource reservation configure ARP (Allocation Retention Priority) during authorization and it is obtained in [PMIP QoS]. There is no corresponding configuration in 802.11 QoS. However, the AP/WLC may use ARP to determine priority during call setup and vulnerability to release of reserved QoS resources.

Parameter Allocation-Retention-Priority and sub fields of Priority, Preemption-Capability and Preemption-Vulnerability are used as defined in [PMIP-QoS].

When a new ADDTS request for reservation of QoS resources arrives, if there is sufficient free resources, the AP/WLC proceeds to allocate it. If there are insufficient resources, the AP/WLC may preempt existing calls based on the Preemption-Capability of the new call and Preemption-Vulnerability of established calls.

If the AP/WLC determines that an established flow with reserved resources should be released, the AP/WLC should inform the MN using ADDTS (802.11aa) and signal the LMA with a revised QoS reservation in PBU/PBA.

5. Security Considerations

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

6. IANA Considerations

No IANA assignment of parameters are required in this document.

7. References

7.1. 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.

7.2. 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-11, Feb 2014.
- [WMM 1.2.0] Wi-Fi Multimedia Technical Specification (with WMM-Power Save and WMM-Admission Control) Version 1.2.0
- [802.11aa] Wireless LAN Medium Access Control (MAC) and Physical Layer (PHY) Specification, Amendment 2: MAC Enhancements for Robust Audio Video Streaming, IEEE 802.11aa-2012.
- [802.11-2012] 802.11-2012 - IEEE Standard for Information technology-Telecommunications and information exchange between

systems Local and metropolitan area networks--Specific
requirements Part 11: Wireless LAN Medium Access Control
(MAC) and Physical Layer (PHY) Specifications

- [GSMA-IR34] Inter-Service Provider Backbone Guidelines 5.0, 22
December 2010
- [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).
- [TS23.402] Architecture Enhancements for non-3GPP accesses (Release
12), 3GPP TS 23.402, V12.2.0 (2013-09).
- [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).
- [TS29.273] 3GPP EPS AAA interfaces (Release 12), 3GPP TS 29.273
v12.1.0 (2013-09)

Authors' Addresses

John Kaippallimalil
5340 Legacy Drive, Suite 175
Plano, Texas 75024

E-Mail: john.kaippallimalil@huawei.com

Rajesh Pazhyannur
170 West Tasman Drive
San Jose, CA 95134

E-Mail: rpazhyan@cisco.com

Parviz Yegani
1194 North Mathilda Ave.
Sunnyvale, CA 94089-1206

E-Mail: pyegani@juniper.net

Appendix A: QoS in 802.11, PMIPv6 and 3GPP Networks

A.1. QoS in IEEE 802.11 Networks

IEEE 802.11-2012 [802.11-2012] provides an enhancement of the MAC layer in 802.11 networks to support QoS--EDCA (Enhanced Distributed Channel Access). EDCA uses a contention based channel access method to provide differentiated, distributed access using eight different UPs (User Priorities). EDCA also defines four access categories (AC) that provide support for the delivery of traffic. In EDCA, the random back-off timer and arbitration inter-frame space is adjusted according to the QoS priority. Frames with higher priority AC have shorter random back-off timers and arbitration inter-frame spaces. Thus, there is a better chance for higher priority frames to be transmitted. The Wi-Fi Alliance has created a specification referred to as WMM (Wi-Fi Multimedia) based on above.

The MN uses ADDTS (Add Traffic Specs) to setup QoS for a traffic stream between itself and the AP, and DELTS to delete that stream. In WMM [WMM 1.2.0], the AP advertises if admission control is mandatory for an access class. Admission control for best effort or background access classes is not recommended. The Wi-Fi Alliance has created a specification referred to as WMM-AC (Wi-Fi Multimedia Admission Control) based on the above.

A.2. QoS in PMIPv6 Mobility domain

[PMIP-QoS] defines a mobility option that can be used by the mobility entities in the Proxy Mobile IPv6 domain to exchange Quality of Service parameters associated with an MN's IP flows. Using the QoS option, the local mobility anchor and the mobile access gateway can

exchange available QoS attributes and associated values. QoS attributes include node and mobile session Aggregate Maximum Bit Rate (AMBR) for upstream and downstream, Guaranteed Bit Rate (GBR) for upstream and downstream, Maximum Bit Rate (MBR) for upstream and downstream and the Allocation Retention Priority (ARP).

[PMIP-QoS] does not explicitly describe how the QoS signaling and QoS sub-options map into corresponding signaling and parameters in the 802.11 access network. This mapping and the procedures in the 802.11 network to setup procedures are the focus of this document. The end-to-end flow spanning 802.11 access and PMIPv6 domain and the QoS parameters in both segments are described in subsequent sections.

A.3. QoS in 3GPP 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 [TS23.402]. 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].

The use cases in subsequent sections use 3GPP policy along with PMIP QoS for provisioning of QoS in the 802.11 network. However, this is exemplary and alternative policy architectures may be used in practice.

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

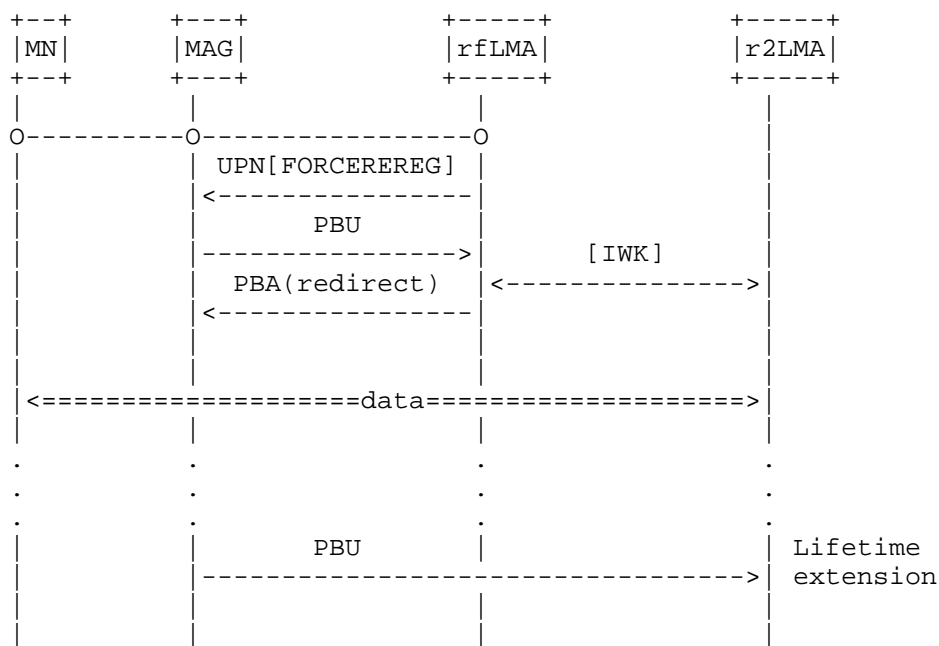
Proxy Mobile IPv6 [RFC5213] describes the protocol operations to maintain reachability and session persistence for a Mobile Node (MN) without the explicit participation from the MN in signaling operations at the Internet Protocol (IP) layer. In order to facilitate such network-based mobility, the PMIPv6 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.

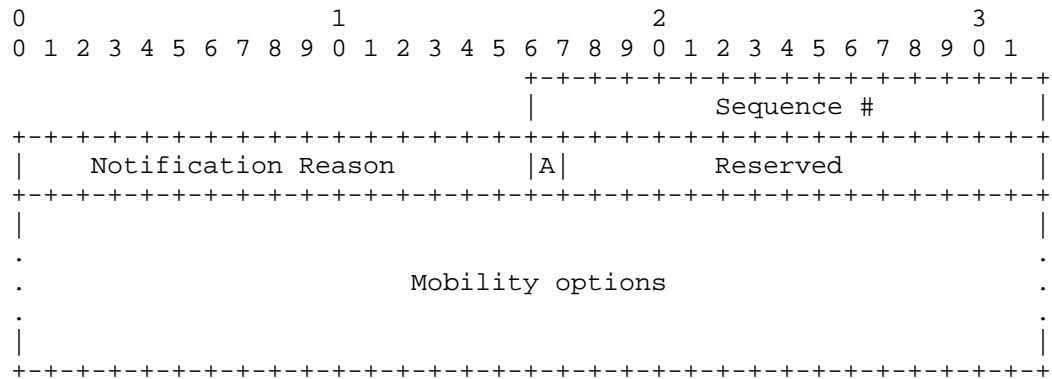
4. MAG Behavior

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.

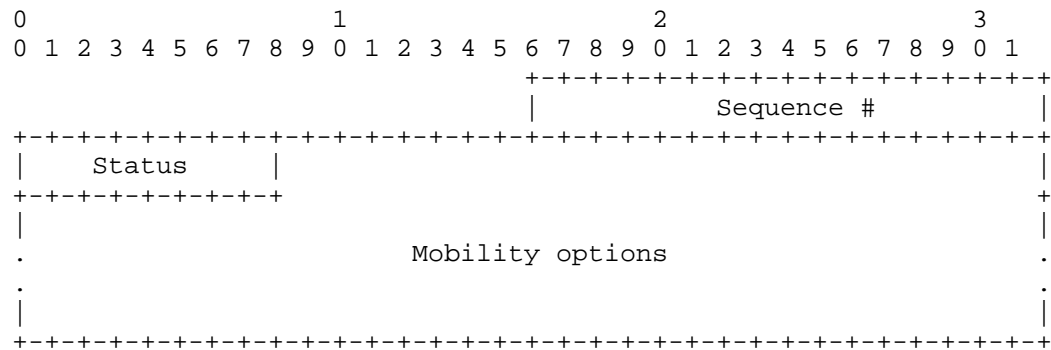
Acknowledgement Requested (A): If this bit is set, the MAG MUST send an UPA message in response to the received UPN message.

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
Internet-Draft
Intended status: Informational
Expires: January 12, 2013

A. Petrescu
M. Boc
C. Janneteau
CEA
July 11, 2012

Network Mobility with Proxy Mobile IPv6
draft-petrescu-netext-pmip-nemo-01.txt

Abstract

The Proxy Mobile IPv6 protocol supports Mobile Hosts moving independently, but not Mobile Routers in charge of moving networks.

This draft addresses this problem. The goal is to allow bidirectional communication between a Local Fixed Node (in the moving network) and a Correspondent Node (situated arbitrarily somewhere in the Internet). First, a mechanism of "prefix division" is presented, whereby the Home Network Prefix typically assigned by PMIPv6 to a MH is used by MR to form Mobile Network sub-Prefix(es); they are used by LFNs within the moving network to form addresses; this avoids changes in the PMIPv6 protocol specification. A second mechanism proposes enhancements to the use of the DHCPv6 Prefix Delegation protocol entities informing the PMIPv6 entities about the allocated MNP; this is achieved by equaling MNID and DUID.

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 12, 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.

Table of Contents

1. Requirements notation	3
2. Concentrated Description	4
2.1. HNP Division	4
2.2. DHCPv6-PD and PMIPv6 Enhancements	7
3. Security Considerations	11
4. Acknowledgements	12
5. Normative References	13
Appendix A. ChangeLog	14
Authors' Addresses	15

1. Requirements notation

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].

2. Concentrated Description

The term Mobile Router has several meanings. One of the agreed meanings at IETF, documented in terminology RFCs, is that of an entity implementing the Mobile IPv6 protocol with NEMOv6 extensions, and accomodating changes in its Care-of Address, maintaining a stable Home Address with the help of a Home Agent, and in charge of LFNs in a moving network whose addresses do not change. Another meaning is that of a router which moves around and does not necessarily change its IP address. In the context of this draft we consider this latter meaning. We ignore whether or not the MR runs Mobile IPv6.

The work presented in this draft is developped in the context of Proxy Mobile IPv6 [RFC5213]. With respect to prefix division, similar methods have been alluded to in the context of DHCPv6 Prefix Delegation by [I-D.krishnan-intarea-pd-epc] (with a slide presentation in the DHC WG at IETF77) and of OSPFv3 by draft-arkko-homenet-prefix-assignment-01.

Mechanisms for supporting Mobile Routers with PMIPv6 and DHCPv6 are presented in [I-D.ietf-netext-pd-pmip] and preceding individual drafts.

The methods presented in this draft are different from most if not all existing documented methods to accomodate moving networks with PMIPv6. In particular, the HNP Division offers several MNPs for use by LFNs, does not modify PMIPv6, does not require the use of DHCPv6-PD but has an inconveninent in that it may not accommodate Ethernet LFNs with SLAAC. The DHCPv6-PD and PMIPv6 enhancements offer MNPs potentially completely different than HNP, may use Ethernet LFNs with SLAAC, modify MAG, LMA, DHCP Relay and potentially DHCP Server.

Moreover, the PMIPv6 and DHCPv6 enhancements presented in this draft rely on the use of MNID being equal to the DUID, a feature absent from existing proposals. Also, with this mechanism the entity performing the allocation of an MNP is the DHCPv6 Server (and not the LMA).

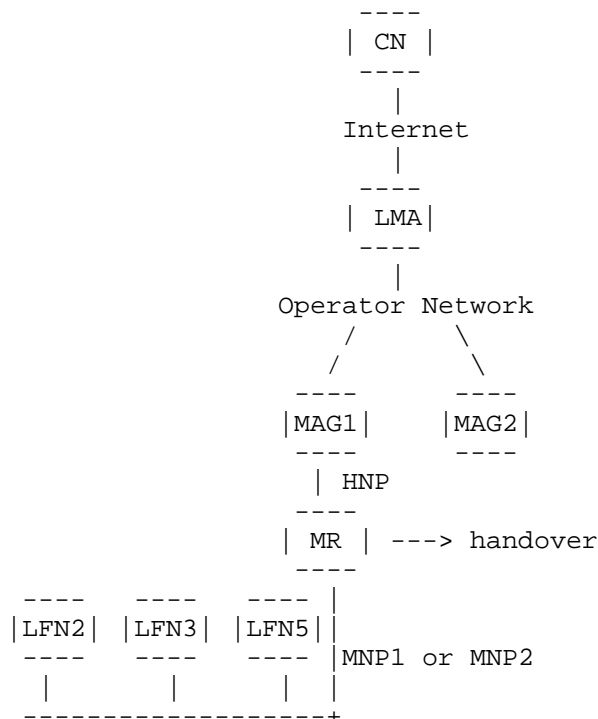
2.1. HNP Division

The mechanism "HNP Division" divides the Home Network Prefix into two or more Mobile Network Prefixes (MNPs).

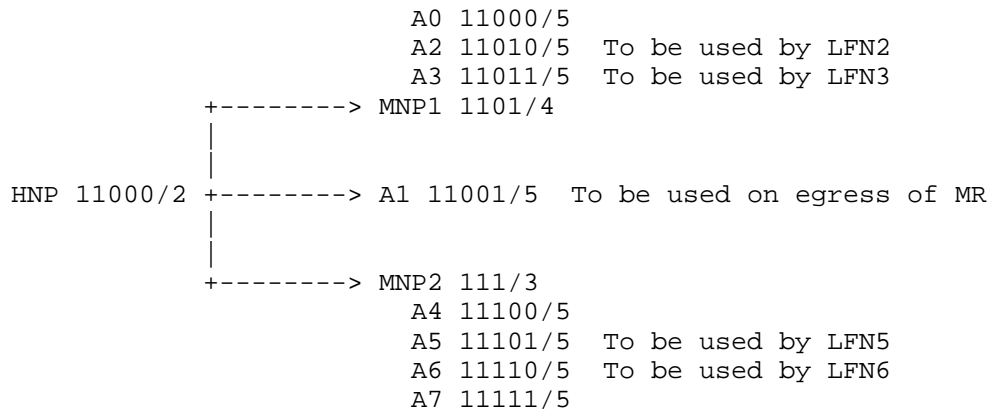
It is assumed that in a domain running PMIPv6 the LMA assigns a Home Network Prefix (HNP) to the Mobile Host. If we consider this Mobile Host to be a Mobile Router, in charge of a set of Local Fixed Nodes (LFNs) in a moving network, it is necessary to use a Mobile Network

Prefix (MNP) within the moving network. Simply using HNP to form addresses for LFNs, without modifying MR behaviour with respect to its routing table, is not sufficient.

The topology illustrated in the next figure depicts a domain where PMIPv6 is run, and a Mobile Router in charge of a set of LFNs forming a moving network.



For a HNP with prefix length 64, two or more MNPs are generated, each having a prefix length longer than 64. For brevity and without losing generality, we present a detailed division example for a fictitious addressing system whose "IP" addresses are of a maximum length of 5 bits (instead of 128 bits of IPv6).



In this example, the HNP/2 11000 is assigned by LMA to MR. The MR divides this into MNP1 1101/4 and MNP2 111/3, and an address A1 11001/5. The MNP1 and MNP2 are used to help LFNs within the moving network to configure full /5 addresses. This may be achieved either with DHCPv6 (MR or a DHCPv6 Server send these addresses) or with stateless address auto-configuration (MR or a Router send Router Advertisements containing MNP1 and/or MNP2).

In most PMIPv6 implementations for MHs, the MAG contains a routing table entry with respect to the allocated HNP. Depending on the nature of the link between MAG and MR, this entry has two different forms: [HNP, vif, *] in case of point-to-point links (typically used in cellular systems) and [HNP, eth, *] (typically used in WiFi hotspot shared links). The vif is a virtual interface, e.g. "ppp0", whereas eth is a real interface, e.g. "eth0".

In the case of point-to-point links, it is not necessary to add any additional behaviour for MR to work (LFN to be reachable from CN). It is sufficient for MR to perform HNP division as described above.

On the contrary, in the case of shared links, it is necessary to perform an operation of Neighbor Discovery proxying on the Mobile Router. When MAG receives a packet from CN addressed to LFN, it would solicit the MAC address of LFN on the MAG-MR link (even though LFN is not present on that link). For this reason, the MR must pretend it owns the IP address of LFN and respond to that solicitation with its own MAC address.


The HNP division mechanism requires that the MNP be part of the HNP (e.g. MNP must have the leftmost n bits the same as the prefix length of HNP), and its length be longer. In case of an HNP/64 and the use of Ethernet for LFNs, only the DHCPv6 protocol can be used by

LFNs, and not SLAAC, because stateless address auto-configuration is not possible for MNPs whose prefix length is longer than 64, the Interface ID being of length precisely 64 for Ethernet.

2.2. DHCPv6-PD and PMIPv6 Enhancements

A second mechanism considers the use of MNP completely different than HNP (may differ on the leftmost bit), hence the use of SLAAC with Ethernet LFNs and HNP/64 is possible, but whereby the PMIPv6 protocol implementation must be modified; this mechanism involves also the use of the DHCPv6 Prefix Delegation protocol.

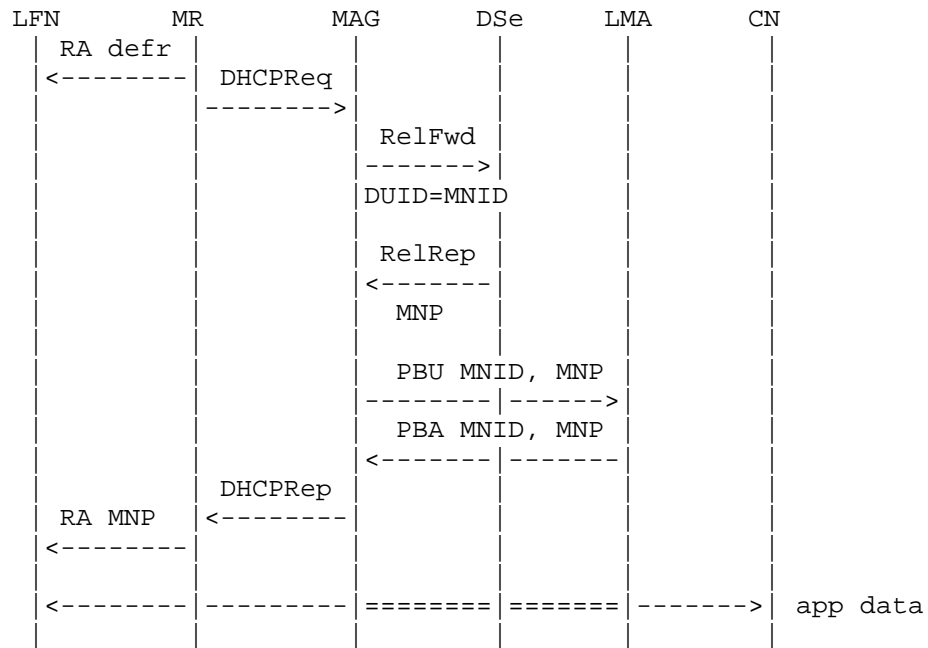
For this mechanism, we consider the following PMIP topology augmented with DHCP entities:



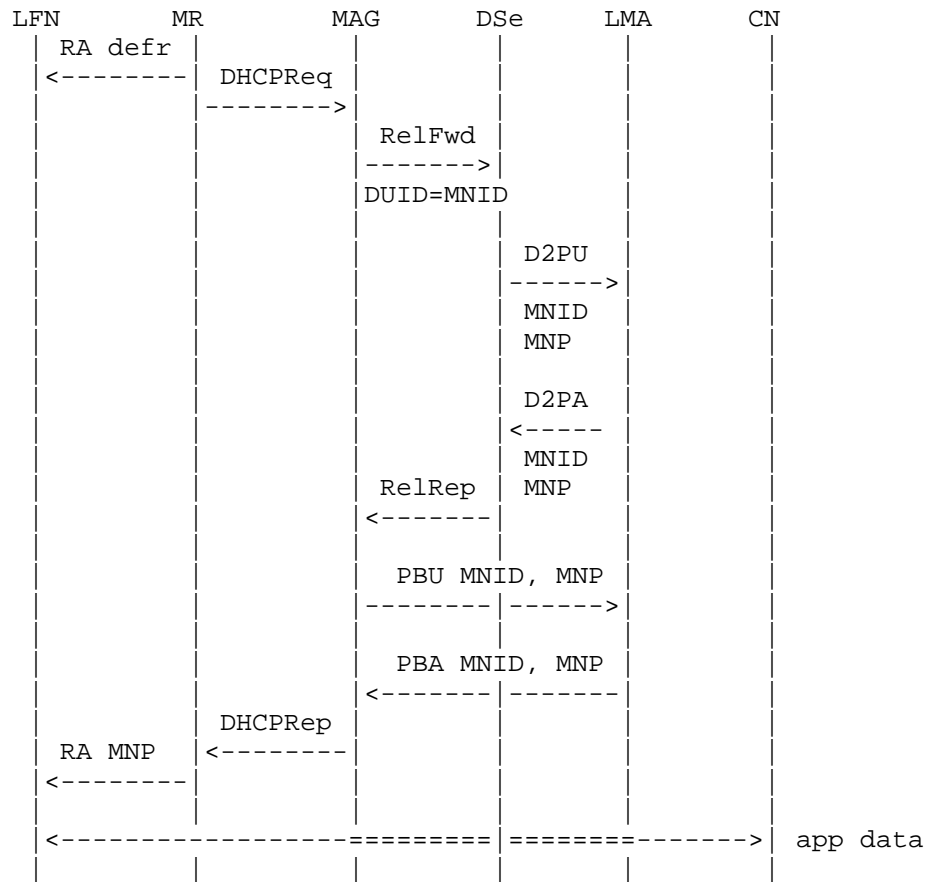
The DSe entity is a DHCPv6 Server. Each MAG also runs a DRe which is a DHCPv6 Relay.

It is necessary to modify the DRe, LMA and MAG behaviour. Depending

on deployment, it may be preferable to modify or to not modify the DHCPv6 Server as well. In case it is not acceptable to modify the DSe the following protocol is proposed:



In case it is not acceptable to modify the DSe the following protocol is proposed:



D2PU and D2PA are new message formats, to be further defined.

The structure of the PBU message is enhanced with respect to the original. Its structure is presented in the following figure:

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	0	1	2	3	4	5	6	7	8	9														
																+++++																+++++																+++++															
																																Sequence #																															
+++++																+++++																+++++																+++++															
A H L K M R P Q								Reserved																								Lifetime																															
+++++																+++++																+++++																+++++															

```

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
                                     +-----+-----+-----+
                                     | Option Type | Option Length |
+-----+-----+-----+-----+-----+-----+-----+-----+
| Subtype | Identifier ...
+-----+-----+-----+-----+-----+-----+-----+-----+

```

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										Prefix Length									
Mobile Network Prefix (MNP)																																							

'Q' flag in PBU: it must be set. It signifies this PBU is sent for an MNP (and not for an HNP).

Length field in the MNP Option: the length of the MNP as was assigned by DHCP.

3. Security Considerations

DHCPv6 and PMIPv6 have security options that should be used in this context as well.

Security risks exist in the process of MR performing proxy Neighbor Discovery on behalf of LFN, if done without explicit authorization provided by LFN.

Security risks exist when performing D2PU and D2PA.

4. Acknowledgements

The mechanisms described in this draft were inspired by several discussions on the NETEXT and intarea email lists. Contributors of these discussions are acknowledged here.

In the process of filing for patent applications the lawyers provided comments which led to better descriptions.

Administratively, this work has been performed in the framework of CELTIC project CP7-011 MEVICO. The authors would like to acknowledge the contributions of their colleagues, although the views expressed are those of the authors and do not necessarily represent the project.

5. Normative References

- [I-D.ietf-netext-pd-pmip]
Zhou, X., Korhonen, J., Williams, C., and S. Gundavelli,
"Prefix Delegation for Proxy Mobile IPv6",
draft-ietf-netext-pd-pmip-01 (work in progress),
October 2011.
- [I-D.krishnan-intarea-pd-epc]
Krishnan, S., Garneij, F., Korhonen, J., and T.
Savolainen, "Prefix Delegation in Evolved Packet Core
networks", draft-krishnan-intarea-pd-epc-00 (work in
progress), February 2010.
- [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.

Appendix A. ChangeLog

The changes are listed in reverse chronological order, most recent changes appearing at the top of the list.

From nil to draft-petrescu-nextex-pmip-nemo-00.txt:

- o The -00 version is mostly a placeholder containing the essence of the mechanisms.

From draft-petrescu-netext-pmip-nemo-00 to -01:

- o Updated the address of authors.
- o Aspects described in the draft are now implemented.

Authors' Addresses

Alexandru Petrescu
CEA, LIST
Communicating Systems Laboratory, Point Courrier 173
Palaiseau, F-91120
France

Phone: +33 169089223
Email: alexandru.petrescu@cea.fr

Michael Mathias Boc
CEA, LIST
Communicating Systems Laboratory, Point Courrier 173
Palaiseau, F-91120
France

Phone: +33 (0) 169083976
Email: michael.boc@cea.fr

Christophe Janneteau
CEA, LIST
Communicating Systems Laboratory, Point Courrier 173
Palaiseau, F-91120
France

Phone: +33 (0) 169089182
Email: christophe.janneteau@cea.fr

Network Working Group
Internet-Draft
Expires: December 9, 2012

B. Sarikaya
F. Xia
Huawei
June 7, 2012

PMIPv6 Multihoming Support for Flow Mobility
draft-sarikaya-netext-fb-support-extensions-02

Abstract

This document specifies extensions to Proxy Mobile IPv6 (PMIPv6) for flow mobility support. Binding cache, binding update list and home network prefix option are slightly extended to allow indicating the home interface and other interfaces.

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 December 9, 2012.

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. Terminology	3
3. Problem Statement	4
4. Multihoming Case	5
5. LMA Operation	5
5.1. Extensions to Binding Cache Entry	5
6. MAG Operation	6
6.1. Extensions to Binding Update List Entry Data Structure . .	7
7. Message Formats	7
8. Security Considerations	8
9. IANA considerations	8
10. Acknowledgements	8
11. References	8
11.1. Normative References	8
11.2. Informative references	9
Authors' Addresses	10

1. Introduction

In Mobile IPv6 [RFC6275] multi-homing is supported efficiently due to the use of home address. Mobile node uses its home address as the source address and all incoming traffic is directed to the home address (HoA). When multiple interfaces are concurrently active the home agent (HA) has to decide how to route incoming packets to different active interfaces. HA does this based on the flow bindings. MN has to register its active flows with the HA and HA keeps flow binding entries for each HoA. HA then forwards packets to one of the care-of addresses of an active interface after matching it with an ordered list of flow bindings.

Proxy Mobile IPv6 [RFC5213] lacks a similar mechanism because each active interface is treated separately and a different binding cache entry is created. This document proposes changes necessary to the local mobility anchor (LMA) behaviour so that flow mobility can seamlessly be supported in PMIPv6. The changes to the mobile node considered in [I-D.ietf-netext-logical-interface-support] are also needed to complement our solution on the host side.

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].

The terminology in this document is based on the definitions in [RFC5213], [RFC6089] in addition to the ones specified in this section:

Single-Radio MN: Consider MN with two interfaces. These interfaces are implemented in such a way that MN can keep one radio module (interface) active at a given time.

Dual/multiple-Radio MN: Consider MN with two interfaces. These interfaces are implemented in such a way that both radio modules can receive and transmit simultaneously.

Inter-technology handover: Sometimes called vertical handover. A multi-homed MN communicates with one interface at any time to conserve power. Each interface can support different access technology. Inter-technology handover occurs when MN moves out of coverage of one technology and moves into the coverage area of another technology which will result in switching of the communicating interface on MN
[I-D.ietf-netext-logical-interface-support].

3. Problem Statement

In base Proxy Mobile IPv6 when MN connects simultaneously with multiple interfaces each interface is treated independently and MN uses different source addresses when sending packet over these interfaces [RFC5213]. However in case of flow mobility, MN itself or LMA might wish to move one flow from one interface to the other. When a flow is moved from interface A to interface B, MN has to stop sending packets on interface A, i.e. it should set the source address to an address based on HNPs assigned to interface B. Forcing an MN to do this after a flow is moved is difficult currently and is one of the problems PMIPv6 flow mobility is facing.

The solution for this is to let MN always use a source address from HNPs assigned to its home interface. When multiple interfaces are active, incoming packets can be directed to different active interfaces based on flow state established at LMA.

In based Proxy Mobile IPv6 LMA treats each interface independently of the other interface(s) MN may have and tries to provide mobility support for each interface. LMA does not manage bindings from different interfaces of the mobile node in an integrated fashion. So LMA can not be in control of moving the flows in between interfaces.

The solution to this is to modify the way the binding cache is managed. Instead of creating an independent mobility session for each interface, the bindings from each interface are kept together so that the flows can be moved among interfaces. The extensions to the base protocol needed for this should be minimal.

When MN does an inter-technology handover, the new MAG sends a Proxy Binding Update (PBU) message to the local mobility anchor (LMA) to register the new proxy care-of address. In the PBU, MAG sets the access technology type (ATT) and handoff indicator (HI) values. If ATT is different from the one stored in the existing binding cache entry for this MN and if HI is set to 2 (Handoff between two different interfaces of the mobile node), LMA concludes that an inter-technology handover happened and assigns the same home network prefix(es) to MN which enables IP session continuity.

Setting the handoff indicator correctly is also not so easy. Most MAGs would tend to set HI to 1 (Attachment over a new interface) which would result in LMA setting new prefix(es) to MN and creating a new binding cache entry and allocating a new mobility session for this new interface. This behaviour as described in Section 5.4 of [RFC5213] needs to be changed.

4. Multihoming Case

When there is attachment over a new interface (HI value received in the Binding Update from MAG is 1) LMA creates a new binding cache entry and assigns the flag "S" defined in Section 5.1 to all home network prefixes assigned to this interface. Also the corresponding value is set to the (H) flag of the home network interface option defined in Section 7 in the binding acknowledgement sent to MAG. LMA MUST also include the home network prefixes with "H" flag in the BA message. This should enable MN continue to send packets with source addresses selected from HNPs with "H" flag on.

The new binding cache entry does not create a new mobility session. The entry is considered as a pointer to another binding the same MN has with LMA. MN may have as many such binding entries as it has active interfaces. These secondary binding cache entries are refreshed regularly by MAGs sending BUs. MAG MUST include HNPs both with "H" and "S" flags in the BU message. LMA refreshes the binding cache entry for the interface with only "S" flag.

5. LMA Operation

When LMA receives a Binding Update message which contains Handoff Indication set to a value of 1 LMA MUST create a new binding cache entry and assign new home network prefixes for this interface. In the binding cache entry these HNPs MUST be flagged with a value of 0 representing "S". This binding cache entry becomes part of the binding cache entry that contains home network prefixes with "H" flag. "H" and "S" flags are as defined in Section 5.1.

LMA sends home network prefixes assigned to the new interface in the Binding Acknowledgement message. LMA MUST also set the (H) Flag in HNP option to 0. In the same BA message, LMA MAY also send home network prefixes whose (H) flag is set to 1 in the same BA.

The modifications specified in this document allow a mobile node to have a single interface connected at a given moment and that interface has prefixes assigned an "S" flag, i.e. the binding with the home interface may have expired. In this case LMA MUST also store the home network prefixes with "H" flag in the binding cache entry.

5.1. Extensions to Binding Cache Entry

One flag associated with the following binding cache entry: list of IPv6 home network prefixes assigned to the mobile node's connected interface and prefix length. The flag is set to 1 representing "H"

if the connected interface is the home interface and flag is set to 0 representing "S" if the connected interface is not the home interface but it is one of the secondary interfaces.

The prefixes assigned after the very first PBU is received for this MN are assigned the "H" flag. The handoff between two different interfaces does not require the prefixes to be changed in order to allow session continuity. Because of this the flag (of "H" or "S") associated with the prefixes stays the same.

This specification also brings the change that binding cache entries for the same MN-Identifier are considered together. The number of entries is equal to the number of active interfaces of MN. If there is a single entry it is assumed that the flag value is "H", otherwise the prefixes with "H" flag should also be stored in the binding cache entry.

For an incoming packet, the destination address MUST be selected from the set of prefixes with "H" flag, i.e. MN always sends non-local packets with source address assigned from HNPs of its home interface. LMA decides to which interface to route this packet by consulting the flow mobility cache [I-D.ietf-netext-pmipv6-flowmob], similar to the case in Mobile IPv6 [RFC6089]. The packet will be matched against the flow descriptions [RFC6088] in the flow mobility cache and Proxy-CoA of the matching entry will be determined. Next, binding cache entry for this MN will be searched and the packet will be directed to the MAG to which the matching interface is connected.

6. MAG Operation

When MAG detects an attachment over a new interface it sets Handoff Indicator field to 1 as described in [RFC5213] in the Binding Update message that it sends to LMA.

MAG MUST store home network prefixes it receives in Binding Acknowledgement message from LMA together with the flag in the binding update list entry. If the flag is "S" MAG MUST also store all home network prefixes in the BA message whose flag is "H" in the corresponding binding update list entry. There will be a maximum of two sets of HNPs for each MN if the MAG is not connected to the home interface.

MAG receives packets from LMA, decapsulates them and searches the binding update list to find the corresponding entry (with the "H" flag) and sends them to the MN with the corresponding "S" flag.

6.1. Extensions to Binding Update List Entry Data Structure

A flag associated with the following binding update list entry: list of IPv6 home network prefixes assigned to the mobile node's connected interface the corresponding prefix length. The flag is set to 0 representing "H" if the connected interface is the home interface and is set to 1 representing "S" if the connected interface is not.

MAG MUST also store the home network prefixes with flag "H" in addition to the prefixes associated with the connected interface if the flag of the home network prefix assigned to the connected interface is "S". MAG determines these flag values from the home network prefix option's (H) flag.

7. Message Formats

Home Network Prefix Option defined in [RFC5213] is modified to include a flag to indicate if home network prefixes are associated with "H" flag or "S" flag in the binding cache entries.

This specification extends the Home Network Prefix Option with a new flag. The flag is shown and described below. All other fields are as described in [RFC5213].

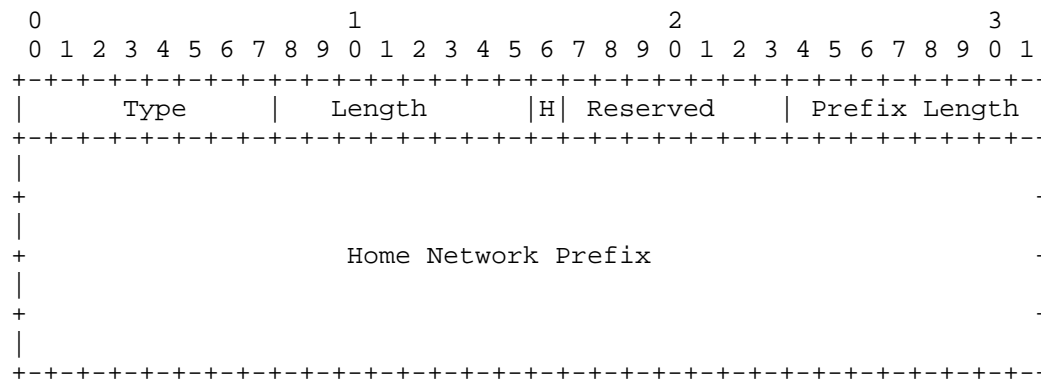


Figure 1: Home Network Prefix Option

H Flag

This flag is set to 1 when this prefix is assigned to a secondary interface of the mobile node, i.e. when the binding cache entry for this HNP has "S" flag set. This flag is set to 0 when this prefix is assigned to the firstly connecting or the only connected interface of the mobile node, i.e. when the binding cache entry

for this HNP has "H" flag set.

8. Security Considerations

This document does not define any new security issues. PMIPv6 security procedures apply.

9. IANA considerations

IANA is requested to add the H Flag into the reserved field in Home Network Prefix Option defined in Section 8.3 in [RFC5213] as the first bit, i.e. Bit number 16 and change the Reserved (R) field as follows:

This 7-bit field is unused for now. The value MUST be initialized to 0 by the sender and MUST be ignored by the receiver.

10. Acknowledgements

The authors thank Hidetoshi Yokota who provided valuable comments.

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.
- [RFC2629] Rose, M., "Writing I-Ds and RFCs using XML", RFC 2629, June 1999.
- [RFC6275] Perkins, C., Johnson, D., and J. Arkko, "Mobility Support in IPv6", RFC 6275, July 2011.
- [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.
- [RFC6088] Tsirtsis, G., Giarreta, G., Soliman, H., and N. Montavont, "Traffic Selectors for Flow Bindings", RFC 6088, January 2011.
- [RFC5213] Gundavelli, S., Leung, K., Devarapalli, V., Chowdhury, K.,

and B. Patil, "Proxy Mobile IPv6", RFC 5213, August 2008.

[RFC5648] Wakikawa, R., Devarapalli, V., Tsirtsis, G., Ernst, T.,
and K. Nagami, "Multiple Care-of Addresses Registration",
RFC 5648, October 2009.

11.2. Informative references

[I-D.ietf-netext-pmipv6-flowmob]
Cano, C., "Proxy Mobile IPv6 Extensions to Support Flow
Mobility", draft-ietf-netext-pmipv6-flowmob-03 (work in
progress), March 2012.

[I-D.ietf-netext-logical-interface-support]
Melia, T. and S. Gundavelli, "Logical Interface Support
for multi-mode IP Hosts",
draft-ietf-netext-logical-interface-support-05 (work in
progress), April 2012.

Authors' Addresses

Behcet Sarikaya
Huawei
5340 Legacy Dr.
Plano, TX 75074

Email: sarikaya@ieee.org

Frank Xia
Huawei
Nanjing, China

Phone:
Email: xiayangsong@huawei.com

