

NETEXT WG
Internet-Draft
Intended status: Standards Track
Expires: May 22, 2014

M. Liebsch
NEC
P. Seite
Orange
H. Yokota
KDDI Lab
J. Korhonen
Broadcom Communications
S. Gundavelli
Cisco
November 18, 2013

Quality of Service Option for Proxy Mobile IPv6
draft-ietf-netext-pmip6-qos-07.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 May 22, 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	5
2.	Conventions and Terminology	7
2.1.	Conventions	7
2.2.	Terminology	7
3.	Quality of Service Option - Usage Examples	10
4.	Protocol Messaging Extensions	13
4.1.	Quality of Service Option	13
4.2.	Quality of Service Attribute	15
4.2.1.	Per Mobile Node Aggregate Maximum Downlink Bit Rate	16
4.2.2.	Per Mobile Node Aggregate Maximum Uplink Bit Rate	17
4.2.3.	Per Mobility Session Aggregate Maximum Downlink Bit Rate	19
4.2.4.	Per Mobility Session Aggregate Maximum Uplink Bit Rate	20
4.2.5.	Allocation and Retention Priority	22
4.2.6.	Aggregate Maximum Downlink Bit Rate	24
4.2.7.	Aggregate Maximum Uplink Bit Rate	25
4.2.8.	Guaranteed Downlink Bit Rate	26
4.2.9.	Guaranteed Uplink Bit Rate	27
4.2.10.	QoS Traffic Selector	28
4.2.11.	QoS Vendor Specific Attribute	29
4.3.	New Status Code for Proxy Binding Acknowledgement	30
4.4.	New Notification Reason for Update Notification Message	30
4.5.	New Status Code for Update Notification Acknowledgement Message	31
5.	Protocol Considerations	32
5.1.	Local Mobility Anchor Considerations	32
5.2.	Mobile Access Gateway Considerations	34
6.	QoS Services in Integrated WLAN-3GPP Networks	38
6.1.	Technical Scope and Procedure	38
6.2.	Relevant QoS Attributes	40
6.3.	Protocol Operation	41
6.3.1.	Handover of existing QoS rules	42
6.3.2.	Establishment of QoS rules	43
7.	IANA Considerations	44
8.	Security Considerations	47
9.	Acknowledgements	48

10.	References	49
10.1.	Normative References	49
10.2.	Informative References	49
Appendix A.	General Use Cases	51
A.1.	Use Case A -- Handover of Available QoS Context	51
A.2.	Use Case B -- Establishment of new QoS Context in non-cellular Access	51
A.3.	Use Case C -- Dynamic Update to QoS Policy	52
Appendix B.	Information when implementing PMIP based QoS support with IEEE 802.11e	54
Appendix C.	Information when implementing with a Broadband Network Gateway	58
	Authors' Addresses	59

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. Current standardization effort considers strong QoS classification and enforcement for cellular radio access technologies. 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 QoS differentiation for access to the mobile operator network through alternative non-cellular access technologies is not yet considered, even though some of these access technologies are able to support QoS by appropriate traffic prioritization techniques. However, handover and IP Flow Mobility using alternative radio access technologies, such as IEEE802.16 and Wireless LAN according to the IEEE802.11 specification, are being considered by the standards [[TS23.402](#)], whereas inter-working with the cellular architecture to establish QoS policies in alternative access networks has not received much attention so far.

In particular Wireless LAN (WLAN) has been identified as an alternative 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. Three functional operations have been identified to accomplish this:

- (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 this IP session. This document specifies fundamental QoS attributes which apply per mobile node, and others that apply per Mobility Session. Additional attributes are specified, which can identify if they apply either per Mobility Session or per flow. The chosen attributes are compatible with the 3GPP 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 of 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](#)], [[RFC5845](#)] and [[RFC5846](#)]. Additionally, this document uses the following abbreviations:

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 that defines a specific Per-hop behavior is known as DSCP. Refer to [[RFC2474](#)], [[RFC2475](#)] and [[RFC4594](#)] for complete explanation.

QoS Profile

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. The Quality of Service option defined in this document represents a QoS Profile.

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. Example of such definitions that are used in this document are, Guaranteed Downlink Bit Rate and Guaranteed Uplink Bit Rate.

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. Example of such definitions that are used in this document are, 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.

Allocation and Retention Priority (AARP)

AARP is used in congestion situations when there are no sufficient 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 rejected by preempting an existing low-priority service.

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.

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.

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.

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 (Example: Internet Service, Voice Over IP Service) and have 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.

3. Quality of Service Option - Usage Examples

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

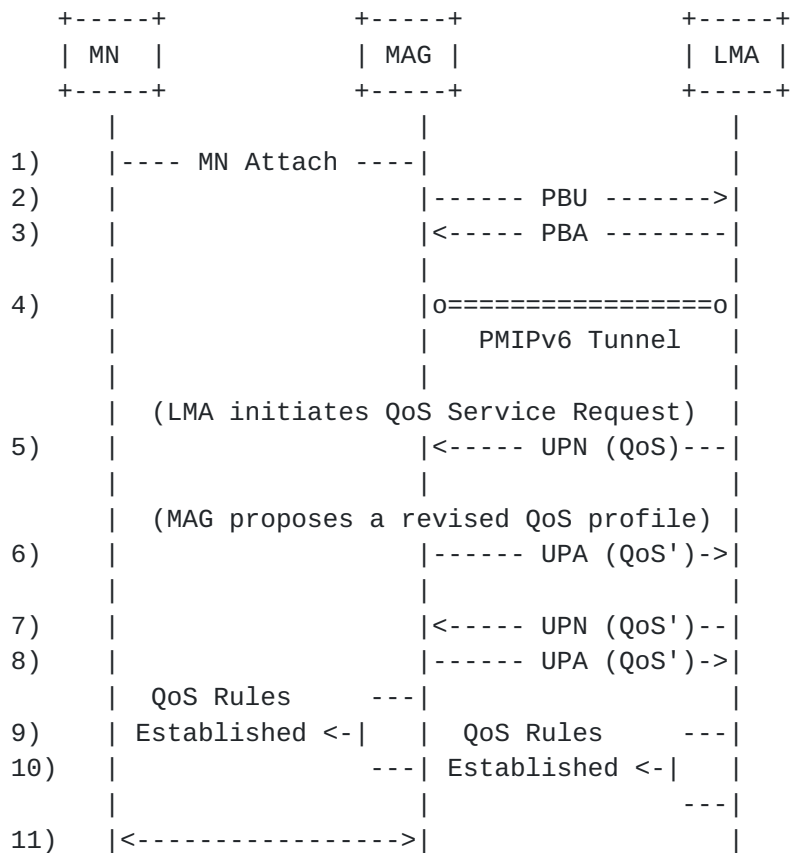


Figure 1: LMA Initiated QoS Service Request

- 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 [[I-D.ietf-netext-update-notifications](#)] to the MAG. The message includes the QoS option [Section 4.1](#) which includes a set of QoS parameters. The mobile access gateway on determining that it cannot support the requested QoS profile for that mobile sends an revised QoS profile in the Update Notification Acknowledgement message, which the LMA agrees to the proposed QoS profile by sending a new Update Notification message

with a modified QoS option.

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

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

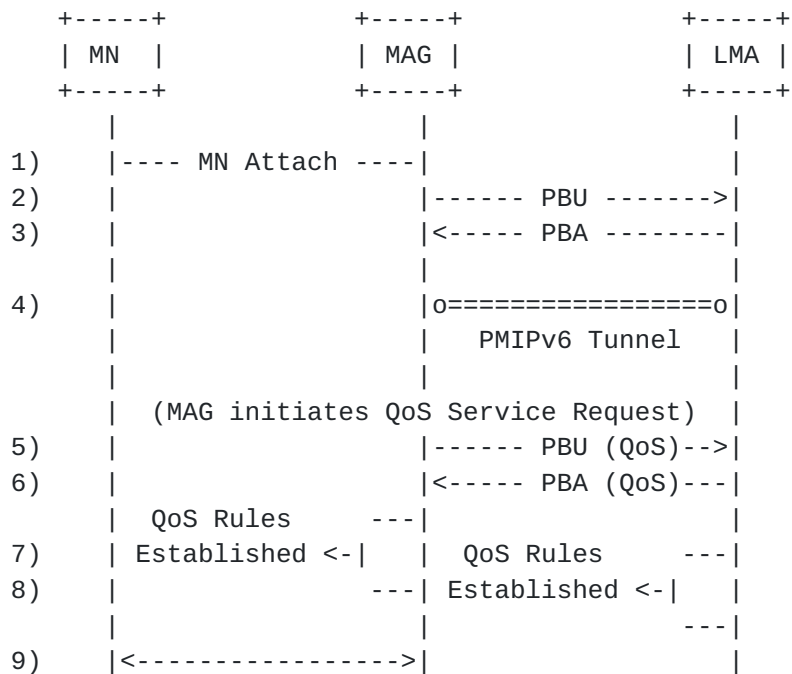


Figure 2: MAG Initiated QoS Service Request

- 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 profile by sending Proxy Binding Acknowledgement message.

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

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) [[I-D.ietf-netext-update-notifications](#)] and Update Notification Acknowledgement (UPA) [[I-D.ietf-netext-update-notifications](#)] 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 profile.

The alignment requirement for this option is 4n.




Figure 3: QoS Option

- o Type: <IANA-1>
- o Length: 8-bit unsigned integer indicating the length of the option in octets, excluding the Type and Length fields.
- o De-Allocate QoS Resources (D) Flag:
 - * When the (D) flag is set to a value of (1), it is an indication that the request is for removal of the QoS resources that have been previously allocated for this mobile node, identified using the profile identifier.
 - * When the (D) flag is set to a value of (0), then there is no special effect on the receiver.
 - * When the (M) flag is set to a value of (1), the (D) flag MUST be set to a value of (0).

- o Modify QoS Resources (M) Flag:
 - * When the (M) flag is set to a value of (1), it is an indication that the request is for modification of the QoS resources that have been previously allocated for this mobile node and identified using the profile identifier.
 - * When the (M) flag is set to a value of (0), then there is no special effect on the receiver.
 - * When the (M) and (D) flags are both set to a value of (0), then there is no special effect on the receiver.
 - * When the (D) flag is set to a value of (1), the (M) flag MUST be set to a value of (0).
- o Profile Identifier (PID): A 5-bit unsigned integer used for identifying the QoS profile. Its uniqueness is within the scope of mobility session. The local mobility anchor always allocates the profile identifier. When the QoS Service request is initiated by a mobile access gateway, it sets the PID value to (0) and the local mobility anchor allocates the PID value and includes that value in the response. For any QoS service requests initiated by a local mobility anchor, the PID value is set to the allocated value.
- 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 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.
- o QoS Attribute(s): Zero or more Type-Length-Value (TLV) encoded QoS Attributes, also referred to as sub-options. 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 is a sub-option that 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 4: 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](#).

(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.

[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 [[I-D.ietf-netext-update-notifications](#)] 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 [[I-D.ietf-netext-update-notifications](#)] 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.



- 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.

- 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.

- 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 there is no special effect on the receiver.
- * 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 there is no special effect on the receiver.
 - * 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 [[I-D.ietf-netext-update-notifications](#)] 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 [[I-D.ietf-netext-update-notifications](#)] 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.

```

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

```

- 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 there is no special effect on the receiver.
 - * 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 there is no special effect on the receiver.
 - * 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 4.2.8](#)), 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 4.2.8](#)), 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																			
+--+--+--+--+--+--+--+--+--+										+--+--+--+--+--+--+--+--+--+										+--+--+--+--+--+--+--+--+--+										+--+--+--+--+--+--+--+--+--+									
										Priority-Level																													
+--+--+--+--+--+--+--+--+--+										+--+--+--+--+--+--+--+--+--+										+--+--+--+--+--+--+--+--+--+										+--+--+--+--+--+--+--+--+--+									
Preemption-Capability										Preemption-Vulnerability																													
+--+--+--+--+--+--+--+--+--+										+--+--+--+--+--+--+--+--+--+										+--+--+--+--+--+--+--+--+--+										+--+--+--+--+--+--+--+--+--+									

- 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: is of type unsigned 32-bit integer, and 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: 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.

- o Preemption-Vulnerability: 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.

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 [[I-D.ietf-netext-update-notifications](#)] 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 [[I-D.ietf-netext-update-notifications](#)] 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	2	3	4	5	6	7	8	9	
Type										Length										Reserved																				
Aggregate-Max-DL-Bit-Rate																																								

- o Type: 3
- o Length: The length of the attribute in octets, excluding the Type and Length fields. This value is set to (6).
- o 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 [[I-D.ietf-netext-update-notifications](#)] 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 [[I-D.ietf-netext-update-notifications](#)] 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 this attribute is present in a Proxy Binding Acknowledgement message, or in an Update Notification Acknowledgement [[I-D.ietf-netext-update-notifications](#)] 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.

```

      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   |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
|                                     |
|               Guaranteed-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 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

[I-D.ietf-netext-update-notifications] 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 [I-D.ietf-netext-update-notifications] 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.

```

      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           |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
|                                     Guaranteed-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 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. There can only be a single instance of this attribute present in a QoS option.

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.

```

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   |   TS Format   |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
~                               Traffic Selector ...                               ~
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+

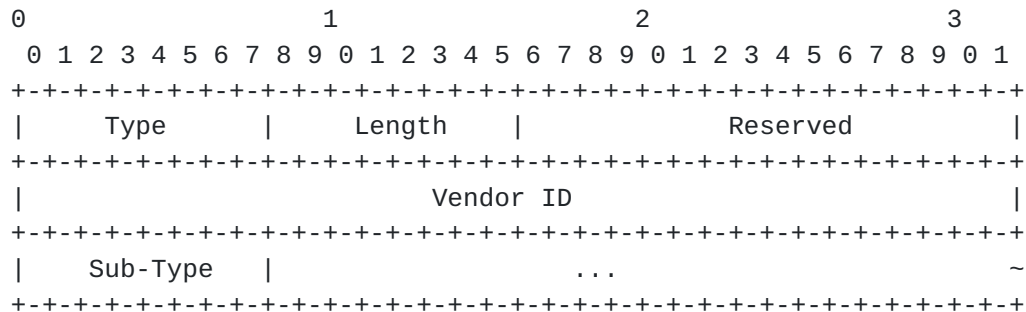
```

- o Type: 8
- 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. 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 opaque 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: 9
- 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 the negotiated Quality of Service profile(s) to be enforced. There can be multiple such profiles and each profile must include the profile identifier and the attributes defined in [Section 4.2](#).

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 MUST process the option(s) and determine if the QoS service request for the proposed QoS profile(s) can be met. Each instance of the Quality of Service option represents a specific QoS profile. This determination 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 profile(s) in entirety, then it MUST send a Proxy Binding Acknowledgement message with a status code value of (0). The message MUST include all the Quality of Service option instances copied (including all the option content) from the received Proxy Binding Update message. The local mobility anchor MUST enforce the Quality of Service rules for all the proposed QoS profile(s) on the mobile node's uplink and downlink traffic. However, if the De-Allocate QoS Resources (D) flag in the received Quality of Service option is set to a value of (1), then the QoS resources previously allocated have to be de-allocated.
- o If the local mobility anchor cannot support the requested QoS profile(s) in entirety then it MUST reject the request and send 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 any existing mobility session for that mobile node. The Proxy Binding Acknowledgement message may include the Quality of Service option based on the following considerations. Rest of the Proxy Binding Acknowledgement message MUST be as specified in [\[RFC5213\]](#) and [\[RFC5844\]](#).
 - * If the local mobility anchor cannot support QoS services for that mobile node and for any QoS profile, then the Quality of Service option MUST NOT be 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 profile(s) or for a partial set of QoS profiles, then the Quality of Service option(s) MUST be included in the Proxy Binding Acknowledgement message. The contents of each of the option (including the QoS attributes) MUST reflect the QoS profile that the local mobility anchor can support for that mobile node. This serves as an indication for the mobile access gateway to resend the Proxy Binding Update message with the proposed QoS profile(s).

Sending a QoS Service Request:

- o The local mobility anchor, at any time during a mobility session, can initiate QoS service request by sending an Update Notification message [[I-D.ietf-netext-update-notifications](#)] with the Notification Reason set to a value of QOS_SERVICE_REQUEST and with the Acknowledgement Requested (A) flag set to a value of (1).
 - * The message MUST include the Quality of Service option(s) with the QoS attributes reflecting the requested QoS profile. Each instance of the Quality of Service option represents a specific QoS profile. The profile identifier MUST be set to a unique identifier value that will be allocated for that profile upon a successful negotiation of the QoS service request.
 - * If the request is for updating any of the parameters of an existing, negotiated QoS service request, the local mobility anchor MUST set the profile identifier to the identifier value allocated for that QoS service request. The Quality of Service option SHOULD have the updated values for the attributes. The (M) flag in the request MUST be set to a value of (1).
 - * If the request is for withdrawal of a currently negotiated QoS service request, the Quality of Service option MUST include the QoS parameters, DSCP value and the profile identifier matching that service request. The (D) flag in the request MUST be set to a value of (1).
- o The response to the Update Notification message for QoS service request must be handled as follows.
 - * If the received Update Notification Acknowledgement [[I-D.ietf-netext-update-notifications](#)] message is with the status code field set to value of (0), the local mobility

anchor MUST enforce the Quality of Service rules for the negotiated QoS profile(s) 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 MUST apply the following considerations.
 - + If the message did not include any Quality of Service option(s), then it is indication from the mobile access gateway that QoS services are not enabled for the mobile node.
 - + If the message includes one more instances of the Quality of Service option, but the option contents reflect a downgraded/revised QoS profile, then the local mobility anchor MAY choose to agree to the proposed QoS profile(s) by resending a new Update Notification message with the revised QoS profile(s). If the proposed QoS profile(s) are not acceptable to the local mobility anchor, then there is no further action needed.

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 MUST be 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 MUST include the QoS option(s) for each of the QoS profiles 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 the negotiated Quality of Service profile(s) to be enforced. There can be multiple such profiles and each profile must include the profile identifier, DSCP value and the attributes defined in [Section 4.2](#).

Receiving a QoS Service Request:

- o On receiving an Update Notification message with one or more instances of Quality of Service option included in the message, the mobile access gateway MUST process the option(s) and determine if the QoS service request for the proposed QoS profile(s) can be met. Each instance of the Quality of Service option represents a specific QoS profile. This determination can be based on policy configured on the mobile access gateway, available network resources in the access network, or based on other considerations.
- o If the mobile access gateway can support all the proposed QoS profile(s) in entirety, then it MUST send an Update Notification Acknowledgement message to the local mobility anchor with the status code value of (0). The message MUST include all the Quality of Service option instances copied (including all the option content) from the received Update Notification message. The mobile access gateway MUST enforce the Quality of Service rules for all the proposed QoS profile(s) on the mobile node's uplink and downlink traffic. However, if the De-Allocate QoS Resources (D) flag in the received Quality of Service option is set to a value of (1), then the QoS resources previously allocated have to be de-allocated.
- o If the mobile access gateway cannot support the requested QoS profile(s) in entirety, then it MUST reject 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 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 and for any of QoS profile, then the Quality of Service option MUST NOT be 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/revised QoS profile(s) or for a partial set of QoS profiles, then Quality of Service option(s) MUST be included in the Update Notification Acknowledgement message. The contents of each of the option (including the QoS attributes) MUST reflect the QoS profile that the mobile access gateway can support for that mobile node. This serves as an indication to the local mobility anchor to resend the Update Notification message with the

revised QoS profile(s).

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 message MUST be constructed as specified in [[RFC5213](#)] and can include the required mobility options.
 - * The message MUST additionally include the Quality of Service option(s) with the QoS attributes reflecting the requested QoS profile. Each instance of the Quality of Service option represents a specific QoS profile.
 - * If the request is for updating any of the parameters of an existing, negotiated QoS service request, the mobile access gateway MUST set the profile identifier to the identifier value allocated for that QoS service request. The Quality of Service option SHOULD have the updated values for the attributes. The (M) flag in the request MUST be set to a value of (1).
 - * If the request is for withdrawal of a currently negotiated QoS service request, the Quality of Service option MUST include the QoS parameters, DSCP value and the profile identifier matching that service request. The (D) flag in the request MUST be set to a value of (1). The (D) flag in the request MUST be set to a value of (1).
- o The response to the Proxy Binding Update message for the QoS service request has to be handled as follows.
 - * If the received Proxy Binding Acknowledgement message has the status code field set to a value of (0), the mobile access gateway MUST enforce the Quality of Service rules for the negotiated QoS profile(s) 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 MUST apply the following considerations.
 - + The denial for QoS service request MUST NOT result in removal of any existing Binding Update list entry for that mobile node.

- + If the message did not include any Quality of Service option(s), then it is indication from the local mobility anchor that QoS services are not enabled for the mobile node.
- + If the message includes one or more instances of the Quality of Service option, but the option contents reflect a downgraded/revised QoS profile, then the mobile access gateway MAY choose to agree to proposed QoS profile(s) by resending a new Proxy Binding Update message with the revised QoS profile(s). If any of the proposed QoS profile(s) are not acceptable to the mobile access gateway, then there is no further action needed.

General Considerations:

- o 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 MUST to be released.

6. QoS Services in Integrated WLAN-3GPP Networks

6.1. Technical Scope and Procedure

The QoS option specified in this document supports the setup of states on the local mobility anchor and the mobile access gateway to allow enforcement of QoS policies for packet differentiation on the network path between the local mobility anchor and the mobile access gateway providing non-cellular access to the mobile operator network. QoS differentiation is typically enabled in the mobile operator's network using Differentiated Services techniques in the IP transport network, whereas radio access specific QoS differentiation depends on the radio technology in use. Whereas accurate and fine granular traffic classes are specified for the cellular radio access, the IP transport network only supports enforcement of few Differentiated Services classes according to well-known Differentiated Services Code Points (DSCP) [[GSMA.IR.34](#)].

The QoS option specified in this document enables exchange of QoS policies, which have been setup for a mobile node's IP flows on the cellular network, between the local mobility anchor and a new mobile access gateway during handover from the cellular access network to the non-cellular access network. Furthermore, the QoS option can be used to exchange QoS policies for new IP flows, which are initiated while the mobile node is attached to the non-cellular mobile access gateway. 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 5 illustrates a generalized architecture where the QoS option can be used. 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 5. 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 case of a distributed access network architecture as per option (II), the mobile access gateway and the access technology specific control function (e.g. the WLC) must provide some protocol for QoS inter-

Based on the architecture illustrated in Figure 5, two key use cases can be supported by the QoS option. Use case A assumes a mobile node is attached to the network through cellular access and its local mobility anchor has QoS policy rules for the mobile node's data flows available. This specification does not depend on the approach how the cellular specific QoS policies have been configured on the local mobility anchor. During its handover, the available QoS policies are established on the target mobile access gateway, which serves the non-cellular access network. Use case B assumes that new policies need to be established for a mobile node as a new IP flow is initiated while the mobile node is attached to the network through the non-cellular network. These use cases are described in more detail in the [Appendix A.1](#) and [Appendix A.2](#) respectively. [Appendix A.3](#) describes a use case where established QoS policies are updated.

6.2. Relevant QoS Attributes

The QoS Option shall at least contain a DSCP value being associated with IP flows of a mobility session. Optional QoS information could also be added. For instance, in order to comply with 3GPP networks QoS, at minimum there is a need to convey the following additional QoS parameters for each PMIPv6 mobility session:

1. Per Mobile Node Aggregate Maximum Bit Rate to both uplink and downlink directions.
2. Per Mobility Session Aggregate Maximum Bit Rate to both uplink and downlink directions.

The following attributes represent a useful set of QoS parameters to negotiate during the session setup:

1. Allocation and Retention Priority (ARP).
2. Guaranteed Bit Rate
3. Maximum Bit Rate

For some optional QoS attributes the signaling can differentiate enforcement per mobility session and per IP flow. For the latter, 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.

Informational Note: If DSCP values follow the 3GPP specification and deployment, the code point can carry intrinsically additional attributes according to a pre-defined mapping table:

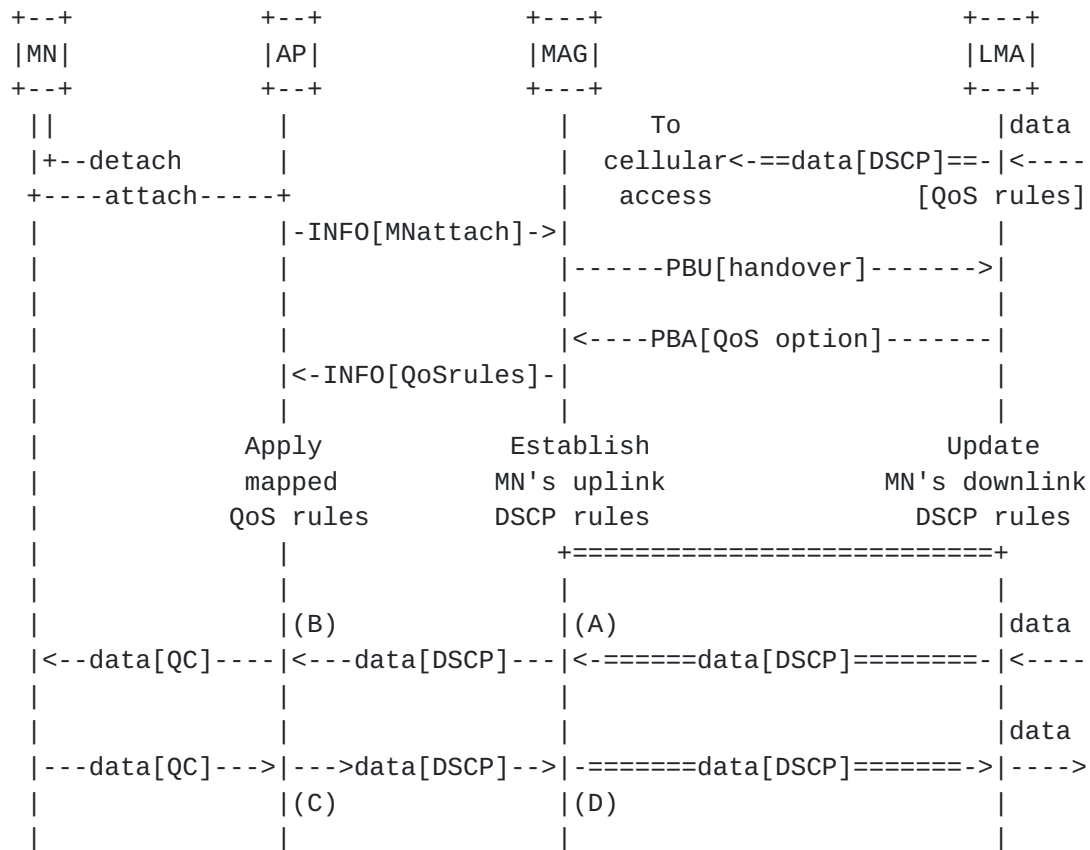
QCI	Traffic Class	DiffServ Per-Hop-Behavior	DSCP
1	Conversational	EF	101110
2	Conversational	EF	101110
3	Conversational	EF	101110
4	Streaming	AF41	100010
5	Interactive	AF31	011010
6	Interactive	AF32	011100
7	Interactive	AF21	010010
8	Interactive	AF11	001010
9	Background	BE	000000

QCI/DSCP Mapping Table

6.3. Protocol Operation

This section provides message flow charts for two scenarios where the QoS option extensions will apply. Correlated to use cases A and B as described in ([Section 6.1](#)), the protocol operation for 'Handover of existing QoS rules' is depicted in (Figure 6), whereas message sequences for 'Establishment of QoS rules' is shown in (Figure 7).

6.3.1. Handover of existing QoS rules



(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 rule

(D): Validate received DSCP and apply DSCP according to rule

Figure 6: Handover of QoS rules

6.3.2. Establishment of QoS rules



(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 7: Adding new QoS profile for MN initiated flow

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 sub-option format, Quality of Service attribute. The format of this mobility sub-option is described in [Section 4.2](#). This sub-option can be carried in Quality of Service mobility option. The type values for this sub-option needs to be managed by IANA, under the Registry, Quality of Service Attribute Registry. This registry should be created under "Mobile IPv6 Parameters" registry at <<http://www.iana.org/assignments/mobility-parameters>>. This specification reserves the following type values. Approval of new Quality of Service Attribute type values are to be made through IANA Expert Review.

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 [[I-D.ietf-netext-update-notifications](#)] 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 [[I-D.ietf-netext-update-notifications](#)] 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. 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 [[I-D.ietf-netext-update-notifications](#)] identify the security considerations for these signaling messages. The quality of service option when included in these signaling messages does not require additional security considerations.

9. Acknowledgements

The authors of this document thank the 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, John Kaippallimalil, Rajesh Pazhyannur, Carlos Jesus Bernardos Cano, Michal Hoefft, Dirk Von Hugo, Ryuji Wakikawa and Seil Jeon for their valuable comments and suggestions to improve this specification.

10. References

10.1. Normative References

- [I-D.ietf-netext-update-notifications]
Krishnan, S., Gundavelli, S., Liebsch, M., Yokota, H., and J. Korhonen, "Update Notifications for Proxy Mobile IPv6", [draft-ietf-netext-update-notifications-12](#) (work in progress), October 2013.
- [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.
- [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.

10.2. 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.
- [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.

- [RFC5845] Muhanna, A., Khalil, M., Gundavelli, S., and K. Leung, "Generic Routing Encapsulation (GRE) Key Option for Proxy Mobile IPv6", [RFC 5845](#), June 2010.
- [RFC5846] Muhanna, A., Khalil, M., Gundavelli, S., Chowdhury, K., and P. Yegani, "Binding Revocation for IPv6 Mobility", [RFC 5846](#), June 2010.
- [SMI] IANA, "PRIVATE ENTERPRISE NUMBERS", SMI Network Management Private Enterprise Codes, February 2011.
- [TS23.402] 3GPP, "Architecture enhancements for non-3GPP accesses", 2010.

Appendix A. General Use Cases

A.1. Use Case A -- Handover of Available QoS Context

The MN is first connected to the cellular network, e.g. an 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 WiFi AP (e.g., at home or in a cafe) and switches to it provided that WiFi access has a higher priority when available. Not only is the session continued, but also the QoS is maintained after moving to the WiFi access. In order for that to happen, the LMA delivers the QoS parameters to the MAG on the WLC via the PMIPv6 signaling and the equivalent QoS treatment is provided toward the MN on the WiFi link.

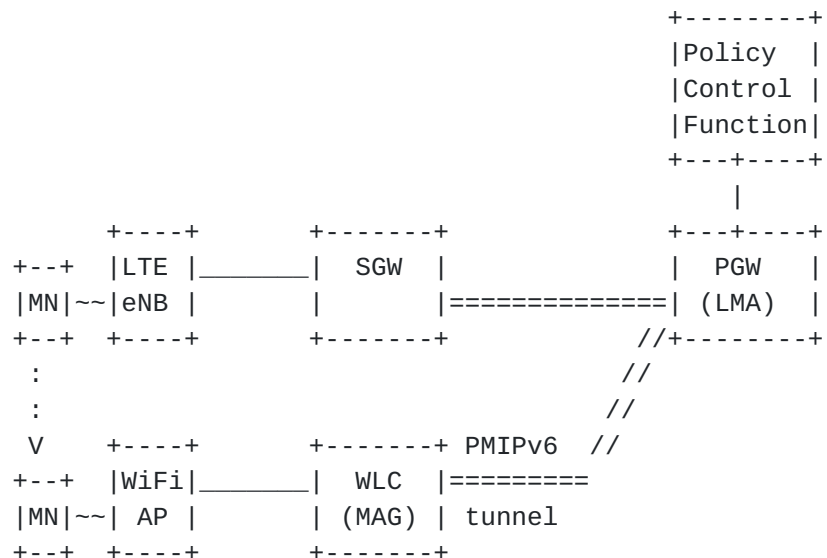


Figure 8: Handover Scenario

A.2. Use Case B -- Establishment of new QoS Context in non-cellular Access

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 9. 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 in an extension to 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 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 as part of the PMIP signaling and the equivalent QoS treatment is provided towards the MN on the WiFi link.




Figure 9: QoS policy provisioning

A.3. Use Case C -- 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 [[I-D.ietf-netext-update-notifications](#)] 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. 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 new flow, the application network function again notifies the PCF function for removing the established bearers. 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". MAG on receiving this message Update Notification Acknowledgement and completes the PBU/PBA signaling for obtaining the new QoS parameters. MAG provisions the newly obtained QoS parameters on the access network to ensure the dedicated network resources are now removed.

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 10 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	
+-----+-----+		

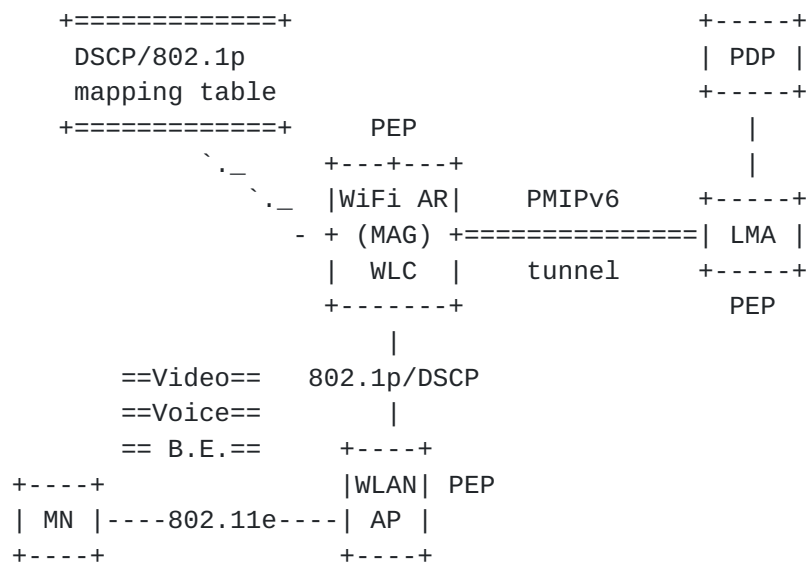


Figure 10: End-to-end QoS management with 802.11e

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 [Section 6.3](#) 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 11 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 renegotiate 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 11: 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., WiFi 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 12: 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 WiFi, 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

