

Mobile IP Working Group
Internet-Draft
Expires: January 11, 2002

A. Festag
X. Fu
H. Karl
G. Schaefer
Technical University Berlin
C. Fan
C. Kappler
M. Schramm
Siemens AG
July 13, 2001

QoS-Conditionalized Binding Update in Mobile IPv6
draft-tnkn-mobileip-qosbinding-mip6-00.txt

Status of this Memo

This document is an Internet-Draft and is in full conformance with all provisions of [Section 10 of RFC2026](#).

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

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

The list of current Internet-Drafts can be accessed at <http://www.ietf.org/ietf/1id-abstracts.txt>.

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

This Internet-Draft will expire on January 11, 2002.

Copyright Notice

Copyright (C) The Internet Society (2001). All Rights Reserved.

Abstract

This draft presents a scheme for QoS support in Mobile IPv6 based on the architecture of Hierarchical Mobile IPv6. A QoS option, a hop-by-hop header extension option first presented by Chaskar et al. [6], is carried in the message containing the Binding Update (BU) option to the mobility anchor point (MAP). Each node between the

mobile node (MN) and the switching MAP (including the MAP itself) which is concerned with managing QoS resources will forward the QoS requirement contained in the QoS option to its internal QoS mechanisms and check resource availability (essentially, performs admission control). Depending on this check, it will either drop the message and send back a message carrying a negative acknowledgement to the MN, or forward the message (possibly after certain modifications to the QoS option) to the next hop. Only when the QoS-conditionalized BU arrives at the MAP where old and new path meet (not necessarily the top-level MAP, but also potentially an intermediate MAP) and the check for resource availability in MAP is also successful, the MAP can make the final binding decision and reply with a BA. Hence, handoffs are conditionalized upon the availability of sufficient resources on the route between MN and MAP to meet QoS requirements. When sufficient resources are not available and there is more than one new access router, the procedure can be reiterated. By way of this scheme, local handoffs are managed locally, transparently to correspondent nodes (CNs) and proper QoS forwarding treatment is established in the new data path, while in the worst case (global mobility) it is managed with Mobile IPv6. Our scheme is flexible (vertical handoff support and several levels of hierarchy can be used), scalable, and potentially supports interworking with multiple QoS mechanisms.

Table of Contents

1.	Introduction	4
1.1	State of the art	4
1.2	Overview of our scheme	4
2.	Terminology	7
3.	Goals and assumptions	9
4.	Detailed description of our scheme	10
4.1	Format of QoS option	10
4.2	Detailed description of QoS-conditionalized binding procedure	12
5.	Comparison of our scheme with the requirements draft	17
5.1	Performance requirements	17
5.2	Interoperability requirements	17
5.3	Exception condition requirements	17
5.4	Miscellaneous requirements	18
5.5	Obvious requirements	18
6.	Further discussion	20
6.1	QoS control in entities unaware of the BU/BA options	20
6.2	Signaling downstream QoS requirements	20
6.3	Upgrading the level of QoS	21
6.4	Possibility of changing from a hop-by-hop option to destination option	21
6.5	Handling intermediate reservations	22
6.5.1	Optimistic reservation	22
6.5.2	Postponed reservation	23
6.6	Handling large signaling packets over the wireless link	23
7.	Security considerations	25
	References	26
	Authors' Addresses	27
	Full Copyright Statement	29

1. Introduction

1.1 State of the art

With the advent of various radio access technologies and increasing deployment of sophisticated applications in mobile end systems, IPv6-based networks will increasingly have to support Quality of Service (QoS) in mobile environments. Mobile IPv6 ensures correct routing of packets to a mobile node when the mobile node changes its point of attachment with the IPv6 network. However, it is also required to provide proper QoS forwarding treatment to the mobile node's packet streams at the changed route in the network due to node mobility in a fast, flexible, and scalable way, so that QoS-sensitive IP services can be supported over Mobile IPv6 [2]. A QoS scheme for Mobile IPv6 should (i) be able to localize the QoS (re-) establishment to the affected parts of the packet path in the network, and (ii) in cases where more than one access technology or access router (AR) is available, it may be desirable for the MN to choose an appropriate AR that can satisfy its QoS requirements among several potential new ARs when the MN moves into such a region (especially since in vertical handoff scenarios, choosing a "good" access router might be more important than the mere speed of reestablishing a QoS path). In these cases, a handoff should not be performed if the MN's QoS requirement is not met; yet if the QoS can be met, handoff should be performed as quickly as possible.

In reference [6] a new IPv6 option called "QoS option" is introduced. One or more QoS objects are included as a hop-by-hop option in IPv6 packets carrying Binding Update (BU) and Binding Acknowledgement (BA) messages. When one packet for this purpose traverses different network domains in the end-to-end path, the QoS option is examined at these intermediate network domains to trigger QoS support for the MN's data packets.

1.2 Overview of our scheme

The mechanism described in reference [6] outperforms RSVP [11][7] in that its signaling overhead is decreased. However, it does not allow to check whether the QoS requirements are satisfied along the new route before performing the handoff. We therefore introduce a QoS-conditionalized binding update. The node at which old and new paths diverge ("switching router") makes the final decision whether or not to update the binding, depending on the result of QoS checks. A binding update will only take place (in the sense of modifying the route) if all nodes along the route between the AR and the switching router are capable of complying with the QoS request, otherwise, the old route will still be used and a negative acknowledgement will be returned to the MN.

Our scheme is based on the architecture of Hierarchical Mobile IPv6 (HMIPv6) [5] to localize the QoS-conditionalized bindings. In HMIPv6, a new entity, the Mobility Anchor Point (MAP), is introduced and a MN only needs to perform one local BU through MAP when changing its layer 3 access point within the MAP domain. HMIPv6 is not able to express QoS requirements, let alone to provide feedback regarding the success of such request. We built on the work described in reference [6] to overcome these limitations.

In our scheme, a QoS hop-by-hop option is carried in the message containing the BU option to the MAP - this message is called BU+QoS message. Each node concerned with QoS management between the MN and the MAP (including the MAP) will pass the QoS requirement represented by the QoS option to internal QoS mechanisms and check its resource availability. If resources are available locally, they are reserved and the message will be forwarded along its route. If specified in the BU+QoS message, reservations covering less than the desired amount of resources are also possible; the request in the BU+QoS message is then updated accordingly. If resources are not available, negative feedback will be provided to the MN. Upon receiving the BU+QoS message, the MAP also checks resource availability and, if successful, will update the binding status and respond with a positive BA+QoS message, including the actual amount of reserved resources, if different from the requested amount. Otherwise, no binding update is performed and a negative BA+QoS message is returned to the MN.

By way of this scheme, QoS (re-)establishment due to local handoffs is managed locally and transparently to the CNs, while in the worst case (global mobility) it is managed with Mobile IPv6 and [6]. Only if all routers along the new path find that sufficient resources are available will a handoff (switching from old to new path) take place. In this sense, the handoff process is conditional on the availability of QoS resources and our scheme can take advantage of HMIPv6. The additional advantage, however, is that mobile terminals will only perform a handoff to an AR that can fulfill the QoS requirement (if there are multiple ARs to choose from; in case there is only a single AR able to serve the mobile terminal, even best-effort service would likely be acceptable, however, this is an application-level concern).

The rest of this document provides a detailed description of the QoS-conditionalized binding update procedure(s) for Mobile IPv6. The document is organized as follows:

- o [Section 2](#) gives the terminology used in the document.
- o [Section 3](#) describes our goals and assumptions.

- o [Section 4](#) gives a detailed description of the scheme, including a description of the rules/considerations for processing and forwarding messages containing Binding Update and Binding Acknowledgement options and QoS option at MNs, MAPs and intermediate routers.
- o [Section 5](#) compares our scheme with the requirements for a QoS solution for Mobile IP as described in [2].
- o [Section 6](#) presents a few important issues brought up by our scheme and gives the reasons for choosing a particular solution among different possibilities within our scheme.
- o [Section 7](#) addresses the security considerations.

2. Terminology

This document uses the following terms in addition to those defined in the Mobile IPv6 protocol [4] and Hierarchical Mobile IPv6 protocol [5]:

QoS option: A Hop-by-Hop option introduced in reference [6]. A QoS option contains zero or more QoS objects in Type/Length/Value (TLV) format.

QoS object: An object introduced in reference [6]. Essentially, QoS OBJECT is an extension of RSVP QoS and FILTER_SPEC objects, and contains certain parameters representing QoS requirements and traffic characteristics for a QoS flow.

QoS entity: An entity responsible for QoS negotiation and establishment. Examples are RSVP daemons in RSVP/IntServ, a Bandwidth Broker in a DiffServ domain, or a Label Edge Router in an MPLS domain. From the perspective of MIPv6, QoS (re)establishment is treated transparently in QoS-capable routers or hosts; the IPv6 nodes MAY ask QoS entities to check the QoS requirements included in the QoS option, and afterwards the latter SHOULD perform such a reservation and respond with an acknowledgement possibly along with (possibly modified) QoS parameters.

Switching router/MAP: The node at which old and new paths diverge

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](#) [1].

Besides, the following acronyms and abbreviations are used in this document:

MIP/MIPv6/HMIPv6: Mobile IP/Mobile IPv6/Hierarchical Mobile IPv6

MAP: Mobility Anchor Point

MN: Mobile Node

CN: Correspondent Node

QoS: Quality-of-Service

CoA: Care-of-Address

RCoA/LCoA: Regional/On-Link CoA

HoA: Home Address

AR: Access Router

BU/BA: Binding Update/Binding Acknowledgement

ER: Edge Router of network domain

IR: Interior Router of network domain

MPLS: Multi-Protocol Label Switching

LSP: Label Switched Path

DiffServ: Differentiated Services

IntServ: Integrated Services

RSVP: Resource ReSerVation Protocol

Upstream(UP)/Downstream(DW) direction: From MN/Towards MN

3. Goals and assumptions

The QoS-conditionalized binding update procedure for Mobile IPv6 is based on the following goals and assumptions:

1. We assume that most handoffs are local. As a result, a QoS solution minimizing the time for QoS (re)establishment may take the advantage of the regional mobility solutions. Furthermore, HMIPv6 model is assumed to be the regional mobility solution within our work.
2. In future wireless communication systems, it is likely that MNs can select among different ARs (possibly implementing different technologies or belonging to different administrative domains). There is hence a requirement of selecting among ARs when performing a handoff and that a handoff to a certain AR should only be performed if already established QoS guarantees can be maintained via the new AR.
3. The QoS entities in the route between (hierarchically the highest) MAP and MN are assumed to be capable of determining whether a given flow toward/from the MN can be admitted and, if so, are capable of (initiation of) reserving and releasing appropriate resources. We do not define any traffic control and resource management solutions.
4. Any QoS (re-)negotiation beyond the highest-level MAP (between this MAP and other network domains) is an administration concern and out of our scope. As we are mostly concerned with local mobility, end-to-end negotiation is not necessary and the QoS negotiation scheme described here therefore is only used in the part of the network between hierarchically highest MAP and MN; the usage of a range of different QoS mechanisms is conceivable here.
5. Support QoS-aware (vertical) handoffs over multiple access technologies collaborating on the IP layer.
6. QoS should be supported for both uplink (from MN) and downlink (to MN) traffic, provided that both uplink and downlink data travels along the same route. (Otherwise, only uplink traffic can directly be supported, support for downlink traffic in this case is in principle feasible yet complicated).
7. Support the specification of both "Acceptable QoS" and "Desired QoS" if so desired by the MN.

4. Detailed description of our scheme

4.1 Format of QoS option

1. The format of the QoS object (see Figure 1) follows [6].

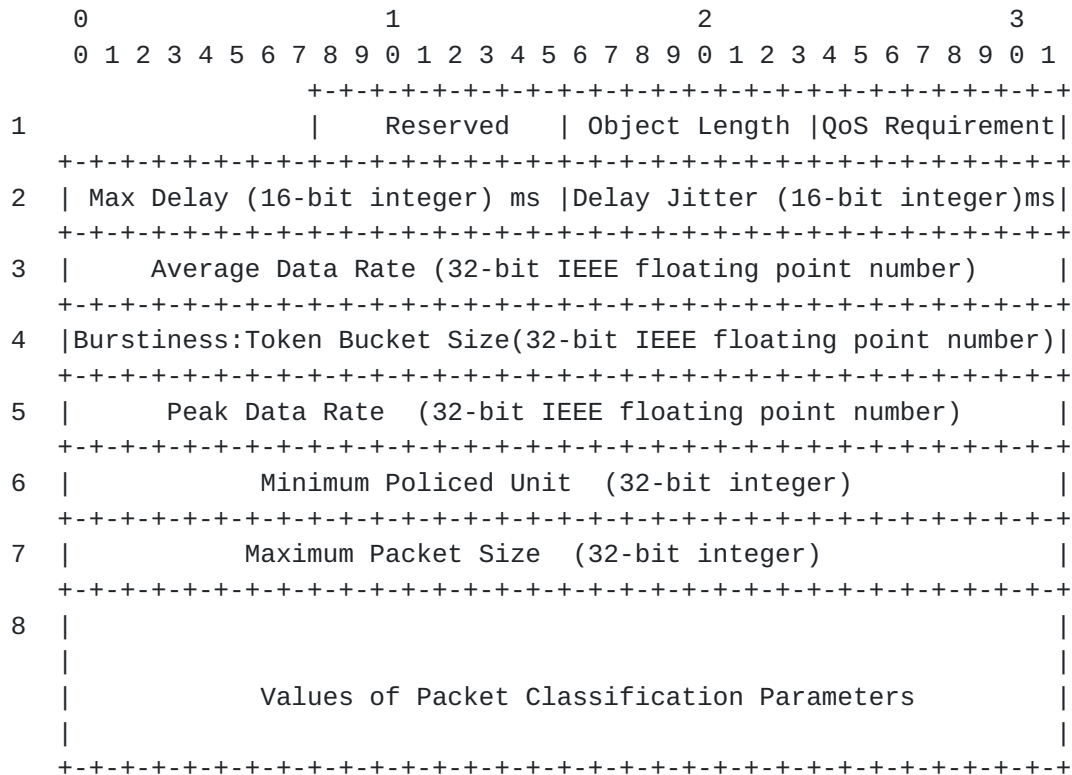


Figure 1 - Composition of a QoS OBJECT

2. Format of QoS Option - follows [6], except that 3 bits of "Reserved" bits are used to specify whether QoS requirement indicated by this option can be met, how to include acceptable and/or desired QoS and up- and/or downstream QoS. (see Figure 2):
 1. If the "F" bit is set, this means "QoS can not be met", otherwise "(up to current node) QoS can be met".
 2. If the "D" bit is set, this means "both upstream QoS and downstream QoS are specified separately", otherwise "upstream and downstream QoS are specified to be the same in both directions" (hence only one QoS object is required). Alternatively, the interpretation of a set "D" bit could be to indicate that only one direction (preferably downlink which is important e.g. for streaming media) is specified. This is yet to be decided.

4.2 Detailed description of QoS-conditionalized binding procedure

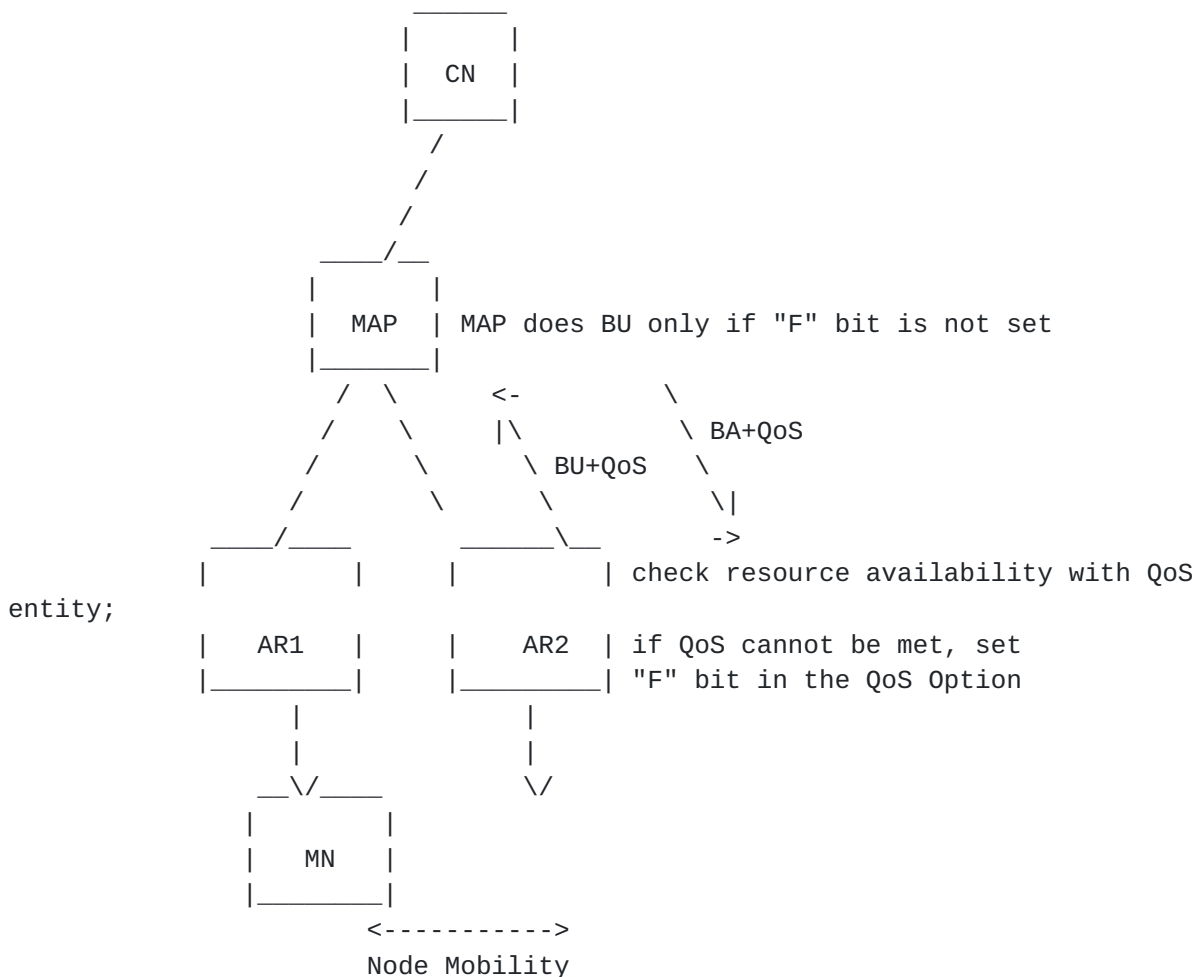


Figure 3 - An Example of a QoS-Conditionalized Binding Procedure

Figure 3 shows an example of a QoS-conditionalized binding update in a MAP domain. In this figure, the MAP is the switching router, the AR and the MAP are the only nodes concerned with QoS control, and IRs are not shown. In general, the path between the switching router and the AR may contain several MAPs, as well as DiffServ/MPLS domains and/or IntServ nodes, or combinations of both. QoS entities in such nodes or domains should make admission control decisions based on the QoS option. The QoS Option is a hop-by-hop header extension option and treated as described below. (As an optimization, the new AR could obtain QoS information from the old AR via context transfer protocols in order to save wireless bandwidth - see discussion in [Section 6.6](#).)

In HMIPv6, outside the MAP domain, destination address or source address of any packet to and from MN is marked as the MAP's IP

address or MN's RCoA in the MAP domain, not the HoA or LCoA of the

Festag, et. al.

Expires January 11, 2002

[Page 12]

MN. Hence, the CN is oblivious of the BA, and a QoS (re-) establishment procedure due to handoff only happens inside the MAP domain. here we only discuss the case of the basic mode of HMIPv6 and the treatment in the extended mode of HMIPv6 needs more investigation.

The QoS-conditionalized binding procedure works as follows.

1. MN gets a new LCoA and composes a message including a Binding Update option (where the so-called "A" bit is set to indicate that an acknowledgement should be generated once the binding update has taken place) and a QoS option. The QoS option may contain up to four QoS objects: it may contain just one QoS object describing the least acceptable upstream QoS; or two QoS objects, additionally describing the desirable upstream QoS; further it may contain both aforementioned objects for the downstream direction. This is discussed in detail in [Section 6.2](#). Note it is possible (cf. [Section 4.1](#)) to skip these last two objects by specifying upstream and downstream QoS as being identical. The now composed message is called "BU+QoS".
 2. MN sends the BU+QoS message to the new AR (towards the MAP). (As stated above, the new AR could obtain QoS information from the old AR via context transfer protocols in order to save wireless bandwidth).
- * IRs forward the BU+QoS message as a normal IPv6 packet. However, each router concerned with QoS control (IntServ node and ER) between the new AR and MAP (including the AR), before passing on the BU+QoS message, SHOULD check whether the "F" bit in the QoS option has been set by a previous router concerned with QoS control. This can be the case if this previous router was unable to generate a BA+QoS message.
 - * If the "F" Bit has been set, it generates a BA+QoS message stating the reason for the fail (Status 131 in BA option - "insufficient resources"), and returns it to the MN, if it is capable of doing so. Thereby no QoS object is returned. If the node is incapable of generating BA+QoS messages, it just passes the message on to the next upstream router. Unless the BA+QoS encounters a router capable of generating BAs, it continues up to the MAP. The capabilities of routers regarding the interpretation of QoS objects, BAs and BUs are discussed in detail in [Section 6.1](#).
 - * If the "F" Bit is not set, it inspects the QoS option, checking whether it (or the QoS domain) can provide this level of QoS by requesting it from the related QoS entity by, e.g.,

asking a Bandwidth Broker of a DiffServ domain to setup the QoS (UP and/or DW).

2.1 If yes, this amount of resources will be reserved, and the BU+QoS will be forwarded to the upstream router.

2.2 If no, if the router is not able to provide the resources, its reaction depends on whether there is just one or two QoS objects.

2.2.1 (a single QoS object is included in the QoS option): the "F" bit in the QoS option will be set. Then the entity will either construct a negative BA+QoS (as described in 2) and return it to the MN if this node is capable to do so. Or else (now the "F" bit in the QoS option is set) the BU+QoS will be forwarded to the next upstream router where (eventually) a negative BA+QoS message will be forwarded to the next upstream router (at least a MAP must be capable of generating negative BA+QoS messages, other intermediate routers may pass on such messages to be handled upstream).

2.2.2 (both desired and acceptable QoS object are included in the QoS option): if the router can provide at least the acceptable QoS, it can reserve whatever capacity it deems appropriate (at or above the acceptable level), write what it reserved in the "desired QoS" object, and then propagate the BU+QoS further upstream. If not even the acceptable QoS can be provided, then this case is treated as 2.2.1.

3. The MAP will perform almost the same steps as in (2), except in (2.1) and (2.2.2). In both cases, after the MAP was able to reserve resources (after possibly adapting them), it stores the MN's LCoA information in its Binding Cache (a binding update is performed) and generates a positive BA+QoS message, containing at most two QoS objects: "desired QoS", for both UP and DW, if available. "Acceptable QoS" objects are dropped, since a positive BA means the acceptable QoS could be met.

4. IRs forward the BA+QoS message as a normal IPv6 packet. However, if this is a router concerned with QoS control (IntServ node and ER) between the new AR and MAP (including the AR), it should inspect the QoS option:

* In case of a negative BA ("F" bit set, and Status 131 of BA option "insufficient resources"), release all resources reserved for this flow. (Other possibilities how to handle

such only semi-valid, intermediate reservations are discussed later in [Section 6.5.](#))

- * Otherwise, in case of a positive BA (no "F" bit set) check whether the QoS resources described are identical to those requested previously. (If no QoS object is present, this means only "acceptable QoS" was specified, which remained unchanged, hence no reservation needs to be changed). If the QoS resources have changed, change the reservation to adapt to this new QoS. Then the BA+QoS is forwarded to the next downstream router.
 - * Note that these routers MUST NOT interpret these QoS options as request for new resources. Rather, these QoS options are interpreted as providing more up-to-date information about a flow for which reservations have already been made.
5. Upon receiving the BA+QoS message, the MN should do the following:
- * If the BU succeeds and the QoS requirement has been met, it starts QoS-guaranteed transmission.
 - * Otherwise, there are essentially three options:
 - + If there is another AR/LCoA available, initiate another BU+QoS message procedure, possibly with different QoS requirements, e.g., the desired level of QoS could be reduced. There are several possibilities of how the number of available access routers could influence the setting of lowest acceptable QoS. E.g., acceptable QoS could be a function of the number of available ARs and/or the MN's speed.
 - + If there is no other LCoA available, try again with this new AR, but reduce QoS requirements, possibly down to best-effort (it should have done so in the first place).
 - + Stay attached to the old AR if this is feasible (particularly in vertical handoff scenarios, a handoff is not necessarily time-critical and connectivity to an old AR can be maintained).

Note that in order to correctly process the BA+QoS message, all routers concerned with QoS management, such as MAPs, ARs, and possibly DiffServ and MPLS ERs, as well as IntServ nodes need to maintain state for each flow. However we believe this is not an additional burden to these entities as they need to maintain this

same state anyways: MAPs must maintain the binding cache, and also the AR has to keep a information, including QoS information, for each MN. DiffServ and MPLS ERs typically act as aggregation routers, i.e. they (as opposed to IRs) still know individual flows, just as do IntServ nodes. Nevertheless, this constitutes an argument in favor of restricting QoS control to AR and MAP.

Note also that the QoS reservation as well as Binding Update option may be maintained as soft state, where QoS option should be refreshed periodically. The timer of QoS option may differ from that for the BU option and the procedure of refreshing QoS options needs further investigation. One possibility would be to periodically resend a packet containing a QoS option with the level of QoS that the MN has actually received. An intermediate router would then have to check whether it already has a QoS reservation for a given flow. If not, this would represent a new flow (and should potentially be accompanied by a BU option in the same packet, if this new flow has to do with mobility). If information about the flow already exists, this QoS option is interpreted as a refresh message, similar to the way the QoS option in the BA message is interpreted.

There are two alternative approaches for our scheme to release the QoS status in the old path if a handoff is successful. One is timing-out, i.e., if no QoS option is received in a certain period of time, the correspondent QoS status in a QoS domain will be cleared. Another way is to release explicitly by sending a QoS option with "F" bit set along the old path towards the MN after a handoff is performed. The selection of using which approach depends on their performance evaluation in different network scenarios and also needs further investigation.

5. Comparison of our scheme with the requirements draft

In [2], a number of requirements are listed which a QoS solution for Mobile IP has to satisfy. The following sections discuss how the conditionalized binding update presented in this draft compares with these requirements.

5.1 Performance requirements

A QoS solution

- o MUST minimize the interruption in QoS at the time of handoff - our scheme minimizes this interruption, because it provides the possibility to check for and reserve resources simultaneously with the binding update, and also allows for negotiating with several ARs to find one that can meet the QoS required.
- o MUST localize the QoS (re)programming to the affected parts of the packet path in the network - satisfied with HMIPv6.
- o MUST provide means to release any QoS state along the old path that is not required after handoff - one possibility is to let the MAP initiate the release request for the old path; the other is timing-out: as BUs time out, the QoS state along the old path will be released.

5.2 Interoperability requirements

A QoS solution

- o MUST be interoperable with other mobility protocols related to mobile IP. This is an open issue, however, the scheme as such could be applicable to other mobility protocols as well.
- o MUST be interoperable with heterogeneous QoS paradigms. As discussed in [Section 4.2](#) above, our scheme interoperates with DiffServ, IntServ and MPLS. Since its task is just carrying QoS information which is then used by QoS entities to do whatever the QoS paradigm requires, it should in fact interwork with any QoS paradigm.

5.3 Exception condition requirements

A QoS solution

- o MUST provide means to handle a situation in which the old QoS

agreement cannot be supported after handoff - our scheme informs the MN the old QoS requirements cannot be met via a negative BA. The MN may initiate another BU with another AR or the same AR with lowered QoS requirements or stay attached to the old AR.

5.4 Miscellaneous requirements

A QoS solution

- o SHOULD be able to support QoS along different potential paths, such as route-optimized path between the MN and the CN, triangle route via HA, temporary tunnels between old and new access router. This is an open issue and requires additional investigation.
- o MAY provide information to link layer to support required QoS, such as acceptable IP packet loss ratio for wireless link. Not supported, extensions are conceivable.

5.5 Obvious requirements

A QoS solution MUST satisfy

- o scalability: the major scalability concern in this scheme is the need to maintain state in intermediate entities. However, as most of the are MAP and hence must maintain binding update mappings, they do keep state on a per-flow level anyway. Hence, this scheme does not introduce any new scalability problems.
- o security - see [Section 7](#)
- o conservation of wireless bandwidth - apart from obtaining a new LCoA address from a new basestation/access router, wireless bandwidth is used only to send BU+QoS and to receive BA+QoS. It can, however, be decreased further by transferring context from old AR to new AR as described in [3] and as discussed later in [Section 6.6](#)
- o low processing overhead on mobile terminals - MN need to insert QoS object into BU and must be able to interpret negative BAs (but compare discussion about the use of context transfer in [Section 6.6](#)).
- o hooks for authorization and accounting - needs further investigation
- o robustness against failure of any MIP-specific QoS component in

the network - since we use the QoS option in a context of HMIP, if (one) MAP fails, the QoS option will be delivered further without any treatment for QoS option (esp. if a destination option for QoS option is used). This needs further investigations.

6. Further discussion

6.1 QoS control in entities unaware of the BU/BA options

In our discussion, we distinguish two kinds of QoS-controlling entities. Both of them are able to interpret the QoS object. While one kind is capable of recognizing the BU/BA options in order to decide what kind of message arrived, and are also capable of generating (negative) BAs, the other kind of QoS-controlling entity do not know about BUs and BAs. Such an entity bases its behavior only on the QoS option along with the QoS objects, but cannot use the BU/BA option to decide how to handle a message. In particular, it must be able to distinguish a QoS option for a flow that has not yet established any reservations at this particular entity from a flow that already has a reservation.

As described in detail in [Section 4.2](#), our mechanism works correctly with both kinds of QoS-controlling entities.

6.2 Signaling downstream QoS requirements

One concern is how to include QoS requirement for downstream traffic into a message carrying Binding options. In an end-to-end signaling scenario, e.g., when using standard Mobile IP, the QoS information for the downstream traffic can easily be provided by the CN. When using a hierarchical architecture, however, the downstream traffic information must still be available for the new path between the MAP and the MN. Requesting this information from the CN would defeat the purpose of using hierarchical mobility schemes in the first place. On the other hand, making this information available in the router might be feasible with some QoS paradigms that provide per-flow QoS, yet QoS schemes that only work on aggregated traffic schemes should not burden intermediate nodes with maintaining information about individual traffic flows (rendering the entire idea of aggregate flow treatment useless) - and this information would have to be present in every router that would potentially be a MAP.

Hence, the downstream QoS description cannot be obtained from the CN, neither can intermediate routers be expected to store this information for every flow. Rather, the downstream traffic QoS requirements should be provided along with the upstream requirements in the BU+QoS message. The MN could know this information (e.g., from some application-level negotiation of QoS) but how to get it is out of the scope of this document.

The main disadvantage of this approach is that the BU packets become larger. While this should not pose much of a problem in the wired backbone network, it could be considered a serious drawback when the

BU+QoS message has to be communicated over the wireless link. There are some possible remedies to this problem which will be discussed later.

Therefore, it is reasonable to assume that the MN also specifies the downstream QoS requirements in the BU+QoS (the MN should be capable of providing this information, e.g., derived from application-level negotiation protocols). While this does increase the amount of protocol data of the solution proposed here, the possibility to reduce state information in the network appears to be the outweighing concern - mechanisms like transferring context from old to new AR (e.g., [3], see discussions later) can additionally reduce wireless bandwidth requirements. The treatment of up- and downstream QoS information in the routers directly follows [6].

6.3 Upgrading the level of QoS

Another concern is which level of QoS requirements is appropriate for a MIPv6 QoS solution. When the MN requests (in preparation of a handoff) a QoS along the new path that is larger than the one used on the old path, the switching router alone can no longer decide whether or not this request should be accepted (assuming that it would be possible to provide this level of QoS on the new path). This inability is partly caused by the need to contact the CN to check whether it agrees as well, whether the CN's access network can provide such an increased capacity (otherwise, upgrading the MN's local reservation would make little sense), and it may also involve inter-domain QoS (re-)negotiation out of the (highest) MAP domain. Therefore, we suggest to consider during a handoff only the problem of maintaining the currently used QoS (and possibly specifying an acceptable lower limit) and to treat the problem of upgrading to higher service levels separately (the main points involved here would be authorization/charging, providing indication of the availability of more resources to the application, and application-level QoS renegotiation).

6.4 Possibility of changing from a hop-by-hop option to destination option

The feature of hop-by-hop option for the QoS option obviously will be a drawback for a fast handoff. Hence, a solution trying to use a destination option may be favorable. If the AR is also a MAP, the MN may specify a destination for the QoS option (destined to the AR) and the AR may relay it as the destination option (destined to the next hop MAP) again, and so forth. Then the QoS option can be carried as a destination option in the whole QoS-conditionalized binding procedure.

Using such a destination option is straightforward if the MAPs are the only entities concerned with QoS control. Typically, at least the AR would also perform QoS control without necessarily being a MAP as well. An important case would be when the AR is the only QoS control entity besides the MAPs. Here, the QoS option can be transmitted from the MN to the (switching) MAP in a hop-by-hop way, but it would be possible for the AR to change the QoS option from a hop-by-hop option to a destination option, destined to the next upstream MAP. Upon receiving this destination option and necessary work regarding QoS management, a MAP between AR and the highest MAP may encapsulate the BU message again to destine it to the hierarchically next higher MAP. When the highest MAP finally receives the BU+QoS message, it will issue a BA+QoS message and follow a reverse procedure (from destination option to a hop-by-hop option).

6.5 Handling intermediate reservations

As the process of accepting a binding update is a distributed one in which several routers can participate, it is necessary to further specify in detail how this decision process should take place. Specifying such distributed processing is further complicated by the fact that multiple binding updates from different MNs could be processed at the same routers with only small temporal offset.

The main issue is how a router handles a BU+QoS message that it could serve, but that also has to be passed upstream onto other routes in order to check whether they can also provide the requested QoS. In general, this is a distributed commit problem and can be solved with well-known techniques, requiring a number of message exchanges e.g., [9]. Here we are interested in faster approaches that should ideally work using only one round trip from MN to switching router and back; sacrificing some optimality aspects is unavoidable in such schemes. Two main schemes are conceivable: optimistic or postponed reservation.

6.5.1 Optimistic reservation

An intermediate router considers the requested QoS as actually being reserved, optimistically assuming that all other routers along the way can also grant the request. Reserving the capacity is the correct decision if all upstream routers can also grant the request. If any upstream router cannot comply with the request, a NACK is returned and the "lower" routers have to release the spuriously made reservation. This optimistic reservation approach can be problematic if a lower router made a reservation that will later be denied and has had to reject other reservation requests that could have been granted upstream.

However, if the round-trip time for BUs is short (which is reasonable in an access network using HMIPv6) and if there is less traffic (relative to the available capacity) towards the core than there is at the edges of the network, this situation should be rather improbable and might hence be regarded as an acceptable risk (in typical networks, the bottlenecks are likely to be closer to the edge than towards the core).

6.5.2 Postponed reservation

In order to circumvent the problems of optimistic reservations, one possibility would be to postpone the actual reservation: when receiving a BU, a router only checks the instantaneous availability of resources, without actually reserving anything when forwarding the BU. Actual reservation only takes place when positive acknowledgements are returned from upstream routers.

The problem with such postponed reservations is that a BA+QoS might not be able to actually reserve capacity because of overlapping BU/BA messages from different MNs. In such a case, the switching MAP has incorrectly reserved capacity and, even worse, performed a handoff to a path that is not QoS-guaranteed. This is a rather serious drawback, and we hence propose to use an optimistic reservation scheme.

6.6 Handling large signaling packets over the wireless link

At the beginning of an application, QoS information needs to be transported over the wireless link in order to enable end-to-end application-level negotiation of QoS requirements. However, as both wireless communication capacity and processing power on mobile terminals are precious resources, once this QoS information has been established, it would be desirable to minimize the amount of QoS information traversing the wireless link and the amount of processing the MN has to perform. A number of different approaches exist for this problem in general: compression schemes, moving protocol functionality away from MNs onto proxies that reside in the wired network; in the present context, context transfer schemes appear to be particularly useful [3].

In particular, it should be feasible to assume that the old AR has the QoS requirement information for each of its MNs. When an MN wishes to associate itself with a new AR, it could simply inform the new AR of the old AR's identity as well as of its own address (permanent and temporary address should work). The new AR then fetches the QoS requirement description from the old AR and initiates the BU process on behalf of the MN; acknowledgements would still have to be provided eventually to the MN.

Alternatively, the binding update process could also be initiated by the old AR. Here, the MN (or even the new AR) becomes aware of a new address it wants to use. The MN asks the old AR to initiate a binding update procedure for this new address. The old AR contacts the corresponding new AR, providing QoS requirements, and the new AR constructs a BU message to be sent in the usual fashion. As soon as the BA arrives, the new AR informs the MN that the new address is no valid and that this new link should now be used. Negative feedback should be provided via the old AR. This scheme is particularly attractive if the MN is not capable of maintaining two different link layer bindings (i.e., communicate with both old and new AR simultaneously).

Choosing between directly transmitting BU/QoS information and transferring context from another AR depends on a number of factors (delay, bandwidth and cost of both the wired and the wireless link and the respective weights assigned to them). Additionally, applying context transfer is orthogonal to different ways of initiating the actual handoff (controlled by the MN, the old or the new AR). However, this question requires additional investigations.

7. Security considerations

The QoS scheme described in this document raises the following threats, mainly concerning the integrity of BU/BA and QoS options:

- o An attacker might possibly try to forge or replay BU messages with specific QoS options in the name of another entity in order to either just harm that entity or to even gain economic benefits as QoS reservations may imply some form of billing consequences.
- o An attacker might try to delay, delete, or modify passing BU+QoS messages (especially, the QoS options), e.g. in order to reduce the desired QoS specification of another entity which might possibly affect its own QoS requests or the QoS requests of a third entity it wants to support in a positive manner.

The above threats SHOULD be averted by protecting the integrity of BU+QoS messages with some kind of cryptographic signature, e.g. as it is done with Mobile IP registration messages. However, this requires the availability of appropriate key material in the signing and the checking entities. It is out of the scope of this specification and for further study if this can be realized with a hop-by-hop approach, that is every intermediate node that processes BU+QoS messages or just the QoS options checks their integrity and signs the outgoing BU+QoS / QoS options, or an end-to-end approach which could, for example, require the last MAP to check the integrity of the mobile node's original BU+QoS message and its QoS option(s).

References

- [1] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", RFC (Best Current Practice) 2119, March 1997.
- [2] Chaskar(Ed.), H., "Requirements of a QoS Solution for Mobile IP", Internet Draft, work in progress, [draft-ietf-mobileip-qos-requirements-00.txt](#), June 2001.
- [3] Koodli, R. and C. Perkins, "Context Transfer Framework for Seamless Mobility", Internet Draft, work in progress, [draft-koodli-seamoby-ctv6-00.txt](#), February 2001.
- [4] Johnson, D. and C. Perkins, "Mobility Support in IPv6", Internet Draft, work in progress, [draft-ietf-mobileip-ipv6-13.txt](#), November 2000.
- [5] Soliman, H., Castelluccia, C., El-Malki, K. and L. Bellier, "Hierarchical MIPv6 mobility management", Internet Draft, work in progress, [draft-ietf-mobileip-hmipv6-03.txt](#), April 2001.
- [6] Chaskar, H. and R. Koodli, "A Framework for QoS Support in Mobile IPv6", Internet Draft, work in progress, [draft-chaskar-mobileip-qos-01.txt](#), March 2001.
- [7] Thomas, M., "Analysis of Mobile IP and RSVP Interactions", Internet Draft, work in progress, [draft-thomas-seamoby-rsvp-analysis-00.txt](#), February 2001.
- [8] Zhang, L., Berson, S., Herzog, S. and S. Jamin, "Resource ReSerVation Protocol (RSVP) -- Version 1 Functional Specification", [RFC 2205](#), September 1997.
- [9] Lyon, J., Evans, K. and J. Klein, "Transaction Internet Protocol Version 3.0", [RFC 2371](#), July 1998.
- [10] Blake, S., Black, D., Carlson, M., Davies, E., Wang, Z. and W. Weiss, "An Architecture for Differentiated Services", [RFC 2475](#), December 1998.
- [11] Braden, B., Clark, D. and S. Shenker, "Integrated Services in the Internet Architecture: an Overview", [RFC 1633](#), June 1994.
- [12] Rosen, E., Viswanathan, A. and R. Callon, "Multiprotocol Label Switching Architecture", [RFC 3031](#), January 2001.

Authors' Addresses

Andreas Festag
Technical University Berlin
Skr. 5-2, Einsteinufer 25
Berlin 10587
Germany

Phone: +49-30-314-79171
EMail: festag@ee.tu-berlin.de

Xiaoming Fu (corresponding author)
Technical University Berlin
Skr. 5-2, Einsteinufer 25
Berlin 10587
Germany

Phone: +49-30-314-23827
EMail: fu@ee.tu-berlin.de

Holger Karl
Technical University Berlin
Skr. 5-2, Einsteinufer 25
Berlin 10587
Germany

Phone: +49-30-314-23826
EMail: karl@ee.tu-berlin.de

Guenter Schaefer
Technical University Berlin
Skr. 5-2, Einsteinufer 25
Berlin 10587
Germany

Phone: +49-30-314-23836
EMail: schaefer@ee.tu-berlin.de

Changpeng Fan
Siemens AG
ICM N MC ST3
Berlin 13623
Germany

Phone: +49-30-386-36361
EMail: changpeng.fan@icn.siemens.de

Cornelia Kappler
Siemens AG
ICM N MC ST3
Berlin 13623
Germany

Phone: +49-30-386-32894
EMail: cornelia.kappler@icn.siemens.de

Mirko Schramm
Siemens AG
ICM N MC ST3
Berlin 13623
Germany

Phone: +49-30-386-25068
EMail: mirko.schramm@icn.siemens.de

Full Copyright Statement

Copyright (C) The Internet Society (2001). All Rights Reserved.

This document and translations of it may be copied and furnished to others, and derivative works that comment on or otherwise explain it or assist in its implementation may be prepared, copied, published and distributed, in whole or in part, without restriction of any kind, provided that the above copyright notice and this paragraph are included on all such copies and derivative works. However, this document itself may not be modified in any way, such as by removing the copyright notice or references to the Internet Society or other Internet organizations, except as needed for the purpose of developing Internet standards in which case the procedures for copyrights defined in the Internet Standards process must be followed, or as required to translate it into languages other than English.

The limited permissions granted above are perpetual and will not be revoked by the Internet Society or its successors or assigns.

This document and the information contained herein is provided on an "AS IS" basis and THE INTERNET SOCIETY AND THE INTERNET ENGINEERING TASK FORCE DISCLAIMS ALL WARRANTIES, EXPRESS OR IMPLIED, INCLUDING BUT NOT LIMITED TO ANY WARRANTY THAT THE USE OF THE INFORMATION HEREIN WILL NOT INFRINGE ANY RIGHTS OR ANY IMPLIED WARRANTIES OF MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE.

Acknowledgement

Funding for the RFC Editor function is currently provided by the Internet Society.

