Mobile Networks and Applications 9, 235–241, 2004

Ongoing Research on QoS Policy Control Schemes in Mobile Networks

HAIHONG ZHENG and MARC GREIS
Nokia Research Center, 6000 Connection Drive, Mail Drop 2-700, Irving, TX 75039, USA

Abstract. This paper describes the ongoing research in IETF on the QoS policy control schemes in the Internet, in particular for Quality of Service (QoS) in IP based mobile networks. The paper gives a general introduction to policy control and then discusses special requirements for policy control in mobile networks. A policy control framework, which matches these requirements, is presented. A detailed analysis and description of the protocols used in the policy control framework is given, followed by some usage examples.

Keywords: QoS, policy control, AAA, mobile network

1. Introduction

Whenever network operators provide enhanced Quality of Service to subscribers, it is necessary to provide the operators with a policy control framework [6], which allows them to control access to this enhanced Quality of Service. The most basic functionality of this policy control is to give the operator the possibility to restrict access to better Quality of Service to users who are willing to pay for it. Policy control enables operators to regulate which users, applications or host should have access to which resources and services and under which conditions (in the telecom world these regulations are often referred to as “subscription plans”). Also, policies can aid the operator in managing the network, for example by restricting certain services to certain times to avoid overloading the network.

It is expected that mobile networks will become more and more important in the near future, both through the growing popularity of technologies like WLAN and Bluetooth as well as the convergence of cellular technologies (e.g., GPRS, UMTS) towards the Internet. However, there are important aspects related to policy control in mobile networks which have not been taken into account within the Internet community so far, most specifically issues related to policy control across different administrative domains.

The purpose of this paper is to give a summary of different QoS control frameworks proposed in IETF, especially the new work of the policy control for the mobile internet, with a special focus being on the usage of the Diameter protocol [1] for transfer of policy rules as well as policy requests and policy decisions between different administrative domains. A short overview over the generic policy architecture as defined by the IETF is given first. We then describe issues related specifically to policy control in mobile networks, after which we give a more detailed view of how Diameter can accommodate the specific requirements for policy control in mobile networks.

2. Intra-domain policy control framework in IETF

A generalized policy management architecture as suggested by the IETF Policy Framework and RAP working group includes a policy management service, a dedicated policy repository, at least one Policy Decision Point (PDP), and at least one Policy Enforcement Point (PEP). The different policy components are shown in figure 1 and a description of their functions follows:

- The policy management service supports the specification, editing, and administration of policy, through a user interface as well as programmatically.
- The policy repository is a model abstraction representing an administratively defined, logical container for reusable policy conditions and policy actions. More specifically, it is a data store that holds policy rules, their conditions and actions. It is also a logical container representing the administrative scope and naming of policy rules, their conditions and actions.

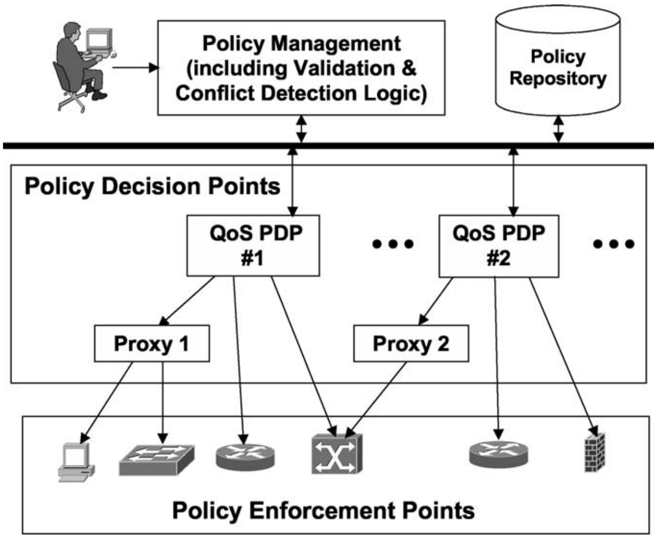
Policy Management (including Validation & Conflict Detection Logic)
Policy Repository
Policy Decision Points
QoS PDP #1
QoS PDP #2
Proxy 1
Proxy 2
Policy Enforcement Points

Figure 1. Intra-domain policy control architecture.

- The PDP is responsible for handling events and making decisions based on those events and updating the PEP configuration appropriately. The policy decisions it makes are based on the policy rules stored in the policy repository. An example protocol used between PDP and policy repository is LDAP [5]. Additionally, it may be responsible for providing the initial configuration of the PEP.
- The PEP enforces policy based on the policy rule sets it receives from the PDP. The policy information can be communicated between PDP and PEP through a variety of protocols, such as COPS [3].
- A Proxy may be used between PDP and PEP to translate information contained in the protocols used between PDP and PEP to the forms that the devices can consume (e.g., command line interface commands or SNMP sets).

The policy control framework defined above focuses only on the intra-domain policy definition and administration for a heterogeneous set of Policy Decision and Enforcement Points. The "intra-domain" refers to policy components that are all under the same administrative control. As mentioned before, the major protocols used in this intra-domain policy framework include COPS and LDAP. However, since roaming and mobility are not considered for the policy control when these IETF working groups are established, the protocol proposed for the intra-domain framework can not be directly used for the inter-domain policy control functions. Section 3 gives more details of requirement of the inter-domain policy control and the reasons of why these inter-domain policy control protocols cannot meet the requirements.

3. Policy control framework in mobile networks

As mobility and roaming are brought into the discussion by the work done in IETF working groups such as Mobile IP, inter-domain policy control becomes a relevant issue and should be considered in an appropriate policy framework. Among various issues, inter-domain policy control needs to allow enforcement of policies as a combination of the roaming user-specific policies defined by the home network and the local network policies. This section first identifies the requirement of the inter-domain policy control brought by the mobility, and then describes a new policy control framework that takes inter-domain and mobility into consideration.

3.1. New requirement of policy control in mobile networks

Roaming and user mobility impose a number of new requirements on policy control. A mobile user may be restricted by policy rules coming from different sources.

- The user's home network operator may regulate the user's access to certain levels of Quality of Service based on the user's subscription. Based on their "subscription plan", certain users may not be allowed to initiate certain sessions under certain circumstances (application type, day of week, time of day, location, etc.). It would usually be desirable for network operators that these regulations are not changed when the user roams to another network. In addition, the home network may also wish to use a different set of user specific policies (i.e., profile parameters) for this particular user based on the roaming relation with the foreign network where the user is visiting.
- When the user is roaming in another network, the operator of this visited network may also want to regulate the user in certain ways. For example, the user's home operator may allow the user to reserve resources for video conferencing sessions, while the visited operator may not want to allow such sessions. Especially in wireless environments where the visited network would provide the expensive radio resources, it would be in the interest of the visited network to enforce certain policies on top of the policies, which have been imposed upon the user by the home network operator. Again, these policy rules can depend upon the aforementioned circumstances like the time of day, etc.

Based on the reasons above, there are certain requirements for policy control arising from user roaming and mobility: it must be possible to base policy control for a roaming user on policy rules both from the user's home network and from the network which the user is currently visiting. This implies either that there needs to be a way to transfer policy decision requests and subsequently policy decisions from the visited to the home network and vice versa or that there is a way to transfer the policy rules for a user from the home network to the visited network so that the policy decision for both networks can be made in the visited network. This in turn implies that there needs to be a standardized protocol for transfer of policy decisions and/or policy rules, which needs to contain a well-defined standardized representation of policy requests, policy rules and policy decisions. Details of these requirements also have been proposed in [4]. The existing protocols used in intra-domain policy frameworks such as COPS and LDAP cannot satisfy such requirements, as described in section 3.3.

3.2. General inter-domain policy control framework

The new inter-domain policy control framework as shown in figure 2 is an extension of the existing intra-domain policy framework.

The policy server (i.e. the PDP) is an authority that controls and admits policy requests based on policies. As mentioned before, the policy framework defined in the policy working group successfully addressed the issues related to "local" policy control (i.e. for users which are currently located in their home network), but not the issues related to "remote" policy control (i.e. for users which are roaming in a visited network). In a mobile network, the remote policy control includes, but is not restricted to, subscription policy control. To simplify the discussion, we use the term "network policy server" to refer to the policy server that controls local network specific policy, and the term "subscription policy server" to refer to the policy server that controls subscription related policy. Also, we

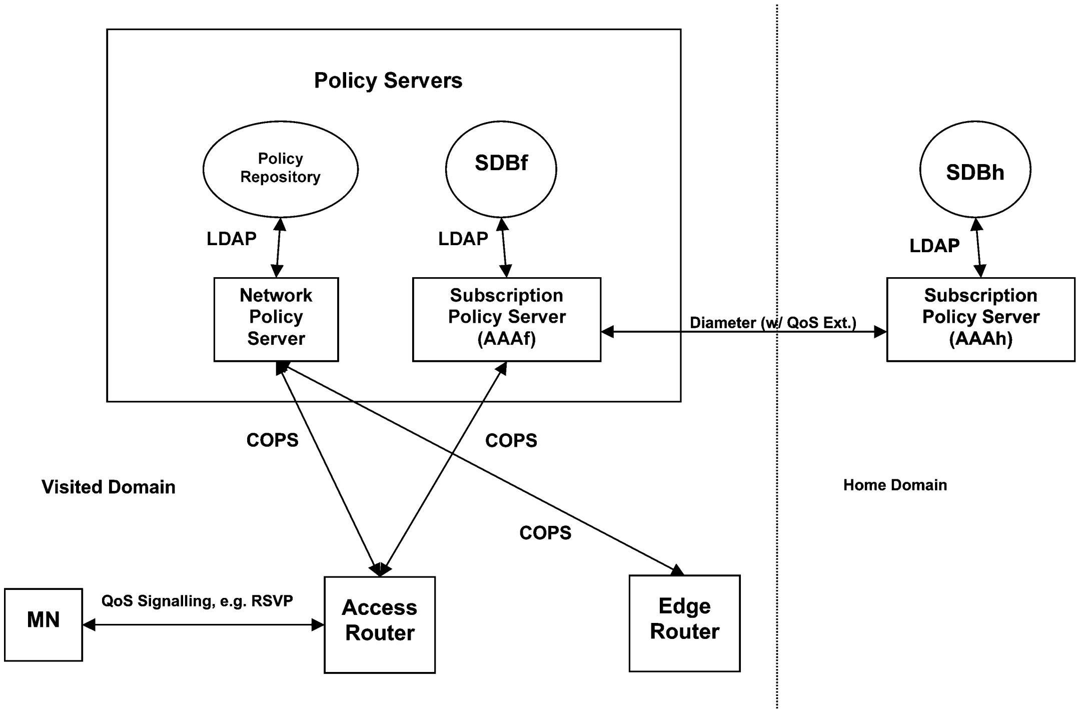

Figure 2. Inter-domain policy control framework.

use the abbreviation "SDB" for the "subscriber database" and "SDBh" and "SDBf" ("f" representing "foreign") for the subscriber databases in the user's home network and the visited network, respectively.

There are two different models for handling subscription policies between a user's home network and the visited network:

- The subscription policies are stored only in the home network of the subscriber. Thus, the policy checking is performed in the home network upon receipt of a request from the subscription policy server in the visited network. The subscription policy server in the home network then sends a policy decision back to the subscription policy server in the visited network. This model provides subscription profile privacy for the user, i.e. only the home network knows the subscription profiles for the user, and the subscription profiles do not have to be transferred between the networks, thus eliminating a possibility for hackers to intercept this transfer. However, this model creates an extra delay in the policy process, as it is always necessary to involve the user's home network in every policy decision.
- To reduce the policy checking delay, a downloading model can be used. The basic idea is to let the subscriber policy server in the visited network download the necessary portion of the subscription profile for the specific subscriber from the subscription database maintained by the subscriber policy server in the home network (i.e. the SDBh). The downloading could be performed for example during the mobile IP registration. The subscription profile can be stored in the subscription database maintained by the subscription policy server in the visited domain (i.e. the SDBv). When a policy request is received by the subscription policy server in the visited network, the policy decisions can be made by directly consulting the subscription database in the visited network. Of course the subscription policy server in the home network should also be able to update the SDBv whenever needed. When the user leaves the domain, the corresponding subscription profile should be removed from the SDBv. Faccin et al. [4] specified such subscription downloading model and analysed the issues to be considered in this model.

After the policy server makes the policy decision, it configures the result into the PEPs such as access router and edge router. A push model and a pull model can be used for the policy transfer both between policy servers as well as between a policy server and PEPs. With the pull model, when a PEP needs a policy decision from the PDPs due to the occurrence of a certain event, it performs policy control by sending policy requests to the relevant policy servers (network policy server and subscription policy server). One example of an event that can trigger the policy control could be the arrival of a QoS request from a mobile node to an access router. The policy servers (including network policy server and subscription policy server) make decisions based on the policies stored in the policy repository (including both network policy repository and SDB) and send reply to the PEP that originated the request. Note that if the user subscription profile cannot be downloaded to the SDBv in the visited network, the subscription policy server in the visited network needs to contact the subscription policy server in the home network to get the policy decision. The PEP then combines the decisions and enforces the policy decision by appropriately accepting or denying the request that triggers the policy control. Examples of PEPs are access routers and edge routers of administrative domains. With the push model, whenever a policy server decides to configure or reconfigure the policies enforced in the router, it sends the policy update request to the PEPs (e.g., access router and edge router). An event that triggers such a new decision could for example be a change in the subscription profile.

3.3. Protocols in the framework

COPS can be used at the interface between various routers (e.g., access router and edge router) and the local network policy server. It specifies a simple client-server protocol that can be used to exchange information between PDPs and PEPs. The client-server model used in COPS requires PDPs and PEPs to know the identity such as the hostname or the IP address of each other. This can be easily implemented by static configuration or dynamic service discovery of the policy servers in a single domain network, however it is not practical in the inter-domain case. This is because it is not a scalable solution to let all the PEPs in the foreign network know a PDP in the home network of a roaming user. Therefore, it is not a scalable and practical solution to use COPS to perform user specific policy evaluation across different domains.

LDAP, a directory access protocol, can be used at the interface between PDPs and policy repositories. Similar to COPS, LDAP uses a simple client–server model. A LDAP client transmits a protocol request describing the operation to be performed to a LDAP server. The server is then responsible for performing the necessary operations in the directory, and returns a response containing any results or errors to the requesting client. It is difficult for a LDAP client in a foreign network where a user roams, to access a LDAP server in the user's home network, because LDAP clients and servers need to know each other's identity to set up a communication and it is not a scalable solution to let every PDP (LDAP client) in the foreign network know a policy repository (LDAP server) in every possible foreign network that may be the home network of a roaming user. Thus, it is also not scalable and practical to use LDAP to retrieve user specific policies from user home network to the network the user is visiting. Therefore, the existing protocols used in intra-domain policy framework cannot satisfy the requirement for the inter-domain case, and a new protocol is needed for this purpose.

The Authentication Authorization and Accounting (AAA) infrastructure has been introduced by the IETF, among other purposes, in order to support inter-domain authentication and authorization. As the AAA protocol, Diameter is the best and the only candidate to be used over the interface between the subscription policy server in the home and foreign networks. The inter-domain functionality defined in Diameter enables an AAA node in the foreign network to communicate with AAA servers in the home network (AAAh) without knowing its hostname or IP address. All the intermediate AAA servers or proxies between the visited and home network can route the AAA message based on the destination realm. This AAA inter-domain routing capability enables AAA node in the foreign domain to locate AAAh only based on the realm of the user NAI, while COPS and LDAP protocols lack this kind of functionality. This feature of AAA can enable the subscription policy server in the foreign network to easily send messages to the subscription policy server in the home network without knowing its hostname or IP address. The signalling messages could be the request for downloading the subscription profile for a certain user, or the request for the subscription policy server in the home network to make a policy decision.

3.4. Policy control functions

This section summarizes the major policy control functions used in the proposed inter-domain policy control framework. These functions include, but are not limited to policy transfer, policy checking, policy update, policy removal and policy execution.

Policy transfer. Policies (e.g., subscription profile) can be transferred from home domain to the visited domain. Two models can be used for the policy transfer, a "push model" and a "pull model".

- In the pull model, a subscription policy server in the visited network sends a request for user-related policy information to the subscription policy server in the home network. The request contains information that is necessary to identify the user whose policy information is to be transferred. This may happen for example if a user appears in a new administrative domain, and the subscription policy server in the new domain attempts to request the policy rules for the new user from the user's home domain. Note that if Diameter is used for the policy transfer, the visited subscription policy server don't need to know the address of the one in the home network, but rely on the inter-domain router capability to route the request to the right home subscription policy server.
- In the push model, a home subscription policy server sends policy rules to a visited subscription policy server without a prior request. The visited subscription policy server may accept or reject the policy rules based on the roaming agreement between two domains. The trigger of pushing the policy into the visited domain from the home could be a registration message, such as Diameter AMR message defined in [2].

Policy checking. The base assumption for a policy check is that a client needs to contact a server to check a user's authorization for requesting a certain service, including QoS. This implies that at least user identification information and a Quality of Service profile are transferred from a client to a server that performs the policy checking. Additionally, more detailed information about the traffic which is supposed to receive the requested Quality of Service should be signalled. This includes any typical flow identification information like addresses, port numbers, etc. It should generally be possible to request authorization for more than one QoS request at a time, as the user may be using an application as well as QoS signalling mechanisms, which allow for multiple QoS requests (e.g., a typical SIP application may give a user the possibility to choose from different media types which would result in different QoS requirements, so the user may attempt to request authorization for all of them).

Upon receiving the service authorization request from the client, the policy server checks the related policy rules and makes policy decision. The local network policy decision is made by the network policy server, and the subscription policy decision is made by the subscription policy server. In the latter case, if the SDBf already contains the subscription profile for the specific user, the policy rules are evaluated locally; otherwise, the visited subscription policy server needs to send the request to the home subscription policy and obtains the decision. The policy servers in the visited network combine all the decision together and issue a reply to the client. The simplest way to accommodate this reply would be a yes/no answer, i.e. the request is either granted or denied. However, to cater for more general cases, it should be possible for the server to reply with a downgraded QoS profile that is within the range allowed by the policies applying to the user.

Policy execution. Both a push model and a pull model can be used for the policy transfer between policy servers and PEPs.

With the pull model, when a PEP needs a policy decision from the PDPs due to the occurrence of a certain event, it performs policy control by sending policy requests to the relevant policy servers (network policy server and subscription policy server). The policy servers make decisions based on the policies stored in the policy repository, or the SDBf, or even the SDBh, and then send reply to the PEP which originated the request. The PEP then combines the decisions and enforces the policy decision by appropriately accepting or denying the request that triggers the policy control. With the push model, whenever a policy server decides to configure or reconfigure the policies enforced in the router, it sends the policy update request to the PEPs (e.g., access router and edge router). An event that triggers such a new decision could for example be a change in the subscription profile.

Policy update and removal. It is necessary for a home subscription policy server to update or remove either all of the policy rules for a given user from the visited subscription policy server or to selectively update or remove only single rules. The latter case is important for example when the home server has configured a temporary policy rule into the client which coexists with permanent policy rules only for a certain time, e.g., the lifetime of an application session.

To update or remove the policy rules from the visited subscription server, the home subscription policy server sends a update or removal request, indicating the user for whom policy rules are to be updated or removed. The message contains a set of policy rules which are to be update or removed. When the home subscription policy server receives an error message, it should re-initialise the policy rules in the target server, i.e. it should delete all rules in the target server and then retransmit them. The visited network may also intend to remove the subscription profile from the SDBf upon certain event triggering. Such event could be receiving deregistration message from a user. After removing the profile from the SDBf, the visited subscription server may need to update the home subscription policy server by sending a removal indication.

3.5. Example policy interaction

This section gives some signalling flows for possible policy interactions to illustrate some policy control mechanisms described above. Note that only some examples are given here and they don't cover all the mechanisms described before.

Figure 3 shows a policy download procedure from a user's AAAh to the AAAf in the user's visited domain, which is triggered by a Binding Update sent from the Mobile Node (MN).

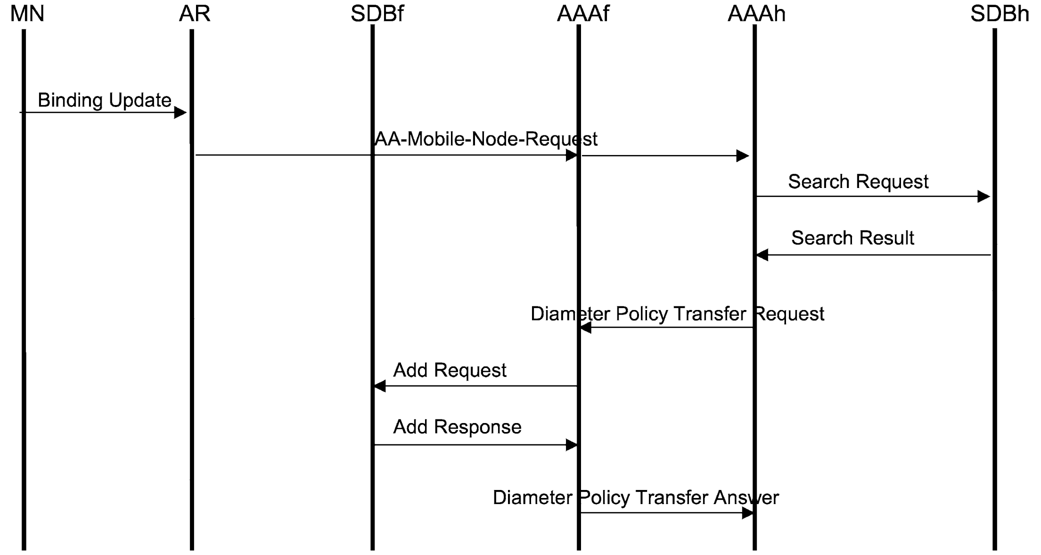

Figure 3. Policy download procedure during registration.

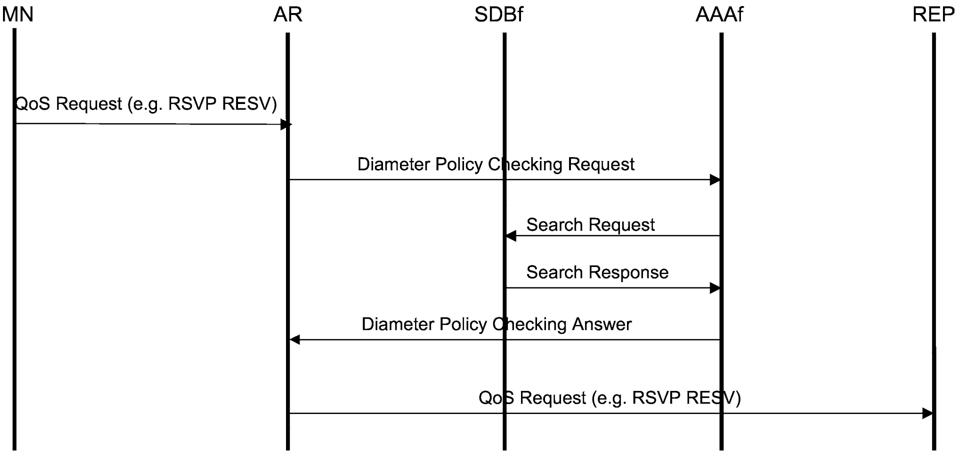
MN
AR
SDBf
AAAf
REP
QoS Request (e.g. RSVP RESV)
Diameter Policy Checking Request
Search Request
Search Response
Diameter Policy Checking Answer
QoS Request (e.g. RSVP RESV)

Figure 4. QoS setup procedure.

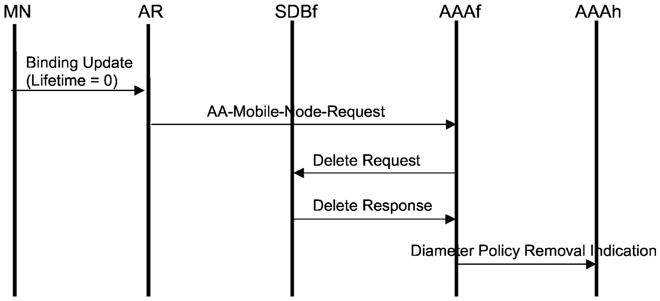

Figure 5. Policy removal procedure.

The following AA-Mobile-Node-Request (see [2]) from the Access Router (AR) to the AAAf is forwarded to the AAAh, which looks up the user's subscription data from the SDBh. The policies which are relevant for the user are transferred to the AAAf using a policy transfer request message and are then added to the SDBf by the AAAf. The AAAf subsequently send a policy transfer answer back to the AAAh to report the status.

Figure 4 shows a policy check during a QoS setup procedure. First, the MN sends a QoS request to the AR. This may be for example an RSVP RESV message, which would have to follow an RSVP PATH message from the remote endpoint (REP). Upon receipt of this request, the AR sends a policy checking request to the AAAf to determine of the user is authorized to request this QoS. The AAAf obtains the user's policy information from the SDBf, and then evaluate the request, after which it returns a policy checking answer to the AR. After receiving a positive authorization for the requested QoS, the AR forwards the QoS request (in this case the RSVP RESV message).

Figure 5 presents a policy removal procedure initiated in the visited network. To deregister, the MN may send a binding update with the lifetime set to zero, upon which the AR would send a AA-Mobile-Node-Request to the AAAf. The AAAf would now delete the policy information for the user from the SDBf and indicate this to the AAAh in a policy removal indication message.

4. QoS policy representation

Current policy representations as defined by IETF working groups do not allow for the simple representation of user-/subscription-related policies, as the current efforts within the IETF are mostly directed towards network management, device configuration and traffic policing. Figure 3 gives an example for a potential representation of policy rules. In this example, a number of filter profiles is defined for a user. At most one of these profiles can be empty, which means that the QoS profile associated with it is the default QoS profile. For each filter profile, one QoS profile is defined (several filter profiles can share the same QoS profile, but not vice versa). Note that the example presented here is very simple, as it does not take into account issues like, e.g., time-based policies (a user is allowed to request a certain QoS only at a certain time).

The QoS profiles in the rules are always seen as the highest QoS that a user can request for a certain flow. This may be a bandwidth specification, e.g., a user may be allowed to request a maximum of 20 kbps for any flow, or a DiffServ codepoint, e.g., a user may be restricted to use AF DiffServ codepoints.

In the simplest case, the policy rules for a user would consist only of a single QoS profile without any filter profiles, which implies that the QoS profile describes the maximum QoS which the user can request for any flow. It is also possible that the policy rules for a user are empty (i.e. there is not even a single QoS profile), which implies that the user can request any QoS for any flow.

There are several scenarios for possible policy checks resulting from the representation in figure 3:

- A user requests QoS for no particular traffic flow (i.e. the user did not specify a filter profile for a request), and there is a default QoS profile for the user: The requested QoS profile is matched against the user's default QoS profile. If the requested QoS profile is not greater than the default QoS profile, the request is granted, otherwise it is denied.
- A user requests QoS for no particular traffic flow (i.e. the user did not specify a filter profile for a request), and there

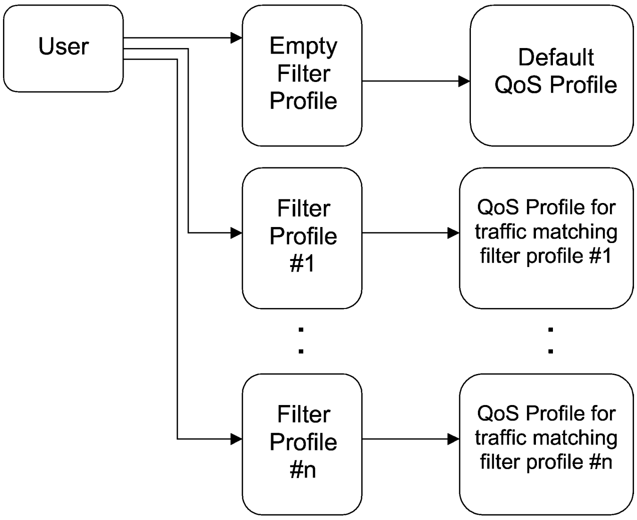
User
Empty Filter Profile
Default QoS Profile
Filter Profile #1
QoS Profile for traffic matching filter profile #1
Filter Profile #n
QoS Profile for traffic matching filter profile #n

Figure 6. A basis for simple policy rules.

is no default QoS profile for the user: The request is denied.

- A user requests QoS for a particular traffic flow (i.e. the user specifies a filter profile for a request) which matches existing filter profiles: The requested QoS profile is matched against all QoS profiles assigned to the matching filter profiles. If the requested QoS profile is greater than any of the QoS profiles, the request is denied, otherwise it is granted.
- A user requests QoS for a particular traffic flow (i.e. the user specifies a filter profile for a request) which does not match existing filter profiles: The requested QoS profile is matched against the user's default QoS profile. If the requested QoS profile is not greater than the default QoS profile, the request is granted, otherwise it is denied. If there is no default QoS profile for the user, the request is denied.

5. Summary

In this paper, we have shown the ongoing work of QoS policy control defined in IETF and the specific requirements for QoS policy control in mobile networks. Based on these requirements, we have described a general framework for inter-domain QoS policy control in mobile networks, and we have also given a more detailed description how the Diameter protocol can fit into this framework. However, the framework in this paper just summarized the beginning of a standardization process with the purpose of creating a commercially viable system for inter-domain policy control. The steps which will have to be taken in the future are:

- The IETF will need to commit to the extension of Diameter to support QoS policy control.
- The mechanisms for transferring, checking and removing policies have to be standardized, the mechanisms described in this paper being one example.
- The message contents for the policy interactions need to be defined, i.e. the Diameter AVPs which contain both the QoS policy rules as well as the QoS policy requests and answers will have to be standardized to allow for inter-domain interoperability between AAA servers.

We believe that the roaming procedures in wireless networks of the future can only match or even exceed the simplicity of roaming procedures for example of current cellular networks (e.g., GSM) if such a QoS policy control framework is standardized, which means that we see this as one of the most important standardization tasks within the IETF in the near future.

References

[1] P. Calhoun et al., Diameter Base Protocol, draft-ietf-aaa-diameter-08.txt, IETF Internet draft, work in Progress (May 2001).

[2] P. Calhoun and C. Perkins, Diameter mobile IPv4 extensions, draft-ietf-aaa-diameter-mobileip-08, IETF Internet draft, work in Progress (May 2001).

[3] D. Durham et al., The COPS (Common Open Policy Service) Protocol, IETF RFC 2748 (January 2000).

[4] S. Faccin, F. Le, S. Sreemanthula and H. Zheng, Profile management framework and Diameter profile management application, draft-zheng-diameter-pma-00.txt, IETF Internet draft, work in progress (November 2001).

[5] M. Wahl, T. Howes and S. Kille, Lightweight Directory Access Protocol (v3), RFC 2251 (December 1997).

[6] R. Yavatkar, D. Pendarakis and R. Guerin, A framework for policy-based admission control, IETF RFC 2753 (January 2000).

Haihong Zheng received her M.S. in computer science and engineering in 1996 with a specialization in Computer Networks from University of Electronic Science and Technology, China. She completed her Ph.D. in computer science and engineering in 2002 from Arizona State University specializing in mobile network systems. She joined Nokia Research Center in 1999 and has since been working on application of IP technology to mobile networks. She actively participated in IETF Robust Header Compression standardization and co-authored RFC 3095. She is now working on QoS aspects in IP based cellular systems and multi access networks. Her main interests include radio resource management, Quality of Service, policy control, AAA, mobility, MPLS and load balancing.
E-mail: haihong.zheng@nokia.com

Marc Greis received his degree in computer science from the University of Bonn, Germany in 1999. After working for the Nokia Research Center in Helsinki, Finland, he moved to the Nokia Research Center in Dallas in 2001. He attended 3GPP for Nokia in 1999 and 2000, as the editor of the QoS framework document TS 23.107 and as chairman of the QoS drafting group in 3GPP TSG SA WG2. Now, he is actively participating in the work of the NSIS working group in the IETF. His main interest lies in any issues related to the convergence of IP networks and mobile telephony networks, with a focus on Quality of Service.
E-mail: marc.greis@nokia.com