

Internet Engineering Task Force (IETF)
Request for Comments: 5866
Category: Standards Track
ISSN: 2070-1721

D. Sun, Ed.
Alcatel-Lucent
P. McCann
Motorola Labs
H. Tschofenig
Nokia Siemens Networks
T. Tsou
Huawei
A. Doria
Lulea University of Technology
G. Zorn, Ed.
Network Zen
May 2010

Diameter Quality-of-Service Application

Abstract

This document describes the framework, messages, and procedures for the Diameter Quality-of-Service (QoS) application. The Diameter QoS application allows network elements to interact with Diameter servers when allocating QoS resources in the network. In particular, two modes of operation, namely "Pull" and "Push", are defined.

Status of This Memo

This is an Internet Standards Track document.

This document is a product of the Internet Engineering Task Force (IETF). It represents the consensus of the IETF community. It has received public review and has been approved for publication by the Internet Engineering Steering Group (IESG). Further information on Internet Standards is available in Section 2 of RFC 5741.

Information about the current status of this document, any errata, and how to provide feedback on it may be obtained at <http://www.rfc-editor.org/info/rfc5866>.

Copyright Notice

Copyright (c) 2010 IETF Trust and the persons identified as the document authors. All rights reserved.

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

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

Table of Contents

1.	Introduction	3
2.	Terminology	3
3.	Framework	5
3.1.	Network Element Functional Model	7
3.2.	Implications of Endpoint QoS Capabilities	8
3.2.1.	Endpoint Categories	8
3.2.2.	Interaction Modes between the Authorizing Entity and Network Element	9
3.3.	Authorization Schemes	10
3.3.1.	Pull Mode Schemes	10
3.3.2.	Push Mode Schemes	13
3.4.	QoS Application Requirements	14
4.	QoS Application Session Establishment and Management	17
4.1.	Parties Involved	17
4.2.	Session Establishment	18
4.2.1.	Session Establishment for Pull Mode	18
4.2.2.	Session Establishment for Push Mode	21
4.2.3.	Discovery and Selection of Peer Diameter QoS Application Node	24
4.3.	Session Re-Authorization	24
4.3.1.	Client-Side Initiated Re-Authorization	25
4.3.2.	Server-Side Initiated Re-Authorization	26
4.4.	Session Termination	28
4.4.1.	Client-Side Initiated Session Termination	28
4.4.2.	Server-Side Initiated Session Termination	28
5.	QoS Application Messages	29
5.1.	QoS-Authorization Request (QAR)	30
5.2.	QoS-Authorization-Answer (QAA)	31
5.3.	QoS-Install Request (QIR)	32
5.4.	QoS-Install Answer (QIA)	32
5.5.	Re-Auth-Request (RAR)	33
5.6.	Re-Auth-Answer (RAA)	34
6.	QoS Application State Machine	34
6.1.	Supplemented States for Push Mode	34
7.	QoS Application AVPs	35
7.1.	Reused Base Protocol AVPs	36
7.2.	QoS Application-Defined AVPs	36
8.	Accounting	37

- 9. Examples 38
 - 9.1. Example Call Flow for Pull Mode (Success Case) 38
 - 9.2. Example Call Flow for Pull Mode (Failure Case) 40
 - 9.3. Example Call Flow for Push Mode 43
- 10. IANA Considerations 45
 - 10.1. AVP Codes 45
 - 10.2. Application IDs 45
 - 10.3. Command Codes 46
- 11. Security Considerations 46
- 12. Acknowledgements 47
- 13. Contributors 47
- 14. References 48
 - 14.1. Normative References 48
 - 14.2. Informative References 48

1. Introduction

This document describes the framework, messages, and procedures for the Diameter [RFC3588] Quality-of-Service (QoS) application. The Diameter QoS application allows Network Elements (NEs) to interact with Diameter servers when allocating QoS resources in the network.

Two modes of operation are defined. In the first, called "Pull" mode, the network element requests QoS authorization from the Diameter server based on some trigger (such as a QoS signaling protocol) that arrives along the data path. In the second, called "Push" mode, the Diameter server proactively sends a command to the network element(s) to install QoS authorization state. This could be triggered, for instance, by off-path signaling, such as Session Initiation Protocol (SIP) [RFC3261] call control.

A set of command codes is specified that allows a single Diameter QoS application server to support both Pull and Push modes based on the requirements of network technologies, deployment scenarios, and end-host capabilities. In conjunction with Diameter Attribute Value Pairs (AVPs) defined in [RFC5777] and in [RFC5624], this document depicts basic call-flow procedures used to establish, modify, and terminate a Diameter QoS application session.

This document defines a number of Diameter-encoded AVPs, which are described using a modified version of the Augmented Backus-Naur Form (ABNF), see [RFC3588].

2. Terminology

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

The following terms are used in this document:

AAA Cloud

An infrastructure of Authentication, Authorization, and Accounting (AAA) entities (clients, agents, servers) communicating via a AAA protocol over trusted, secure connections. It offers authentication, authorization, and accounting services to applications in local and roaming scenarios. Diameter and RADIUS [RFC2865] are both widely deployed AAA protocols.

Application Endpoint (AppE)

An Application Endpoint is an entity in an end-user device that exchanges signaling messages with Application Servers or directly with other Application Endpoints. Based on the result of this signaling, the endpoint may make a request for QoS from the network. For example, a SIP User Agent is one kind of Application Endpoint.

Application Server (AppS)

An Application Server is an entity that exchanges signaling messages with an Application Endpoint (see above). It may be a source of authorization for QoS-enhanced application flows. For example, a SIP server is one kind of Application Server.

Authorizing Entity (AE)

The Authorizing Entity is a Diameter server that supports the QoS application. It is responsible for authorizing QoS requests for a particular application flow or aggregate. The Authorizing Entity may be a standalone entity or may be integrated with an Application Server and may be co-located with a subscriber database. This entity corresponds to the Policy Decision Point (PDP) [RFC2753].

Network Element (NE)

A QoS-aware router that acts as a Diameter client for the QoS application. This entity triggers the protocol interaction for Pull mode, and it is the recipient of QoS information in Push mode. The Diameter client at a Network Element corresponds to the Policy Enforcement Point (PEP) [RFC2753].

Pull Mode

In this mode, the QoS authorization process is invoked by the QoS reservation request received from the Application Endpoint. The Network Element then requests the QoS authorization decision from the Authorizing Entity.

Push Mode

In this mode, the QoS authorization process is invoked by the request from the Application Server or local policies in the Authorizing Entity. The Authorizing Entity then installs the QoS authorization decision to the Network Element directly.

Resource Requesting Entity (RRE)

A Resource Requesting Entity is a logical entity that supports the protocol interaction for QoS resources. The RRE resides in the end-host and is able to communicate with peer logical entities in an Authorizing Entity or a Network Element to trigger the QoS authorization process.

3. Framework

The Diameter QoS application runs between an NE (acting as a Diameter client) and the resource AE (acting as a Diameter server). A high-level picture of the resulting architecture is shown in Figure 1.

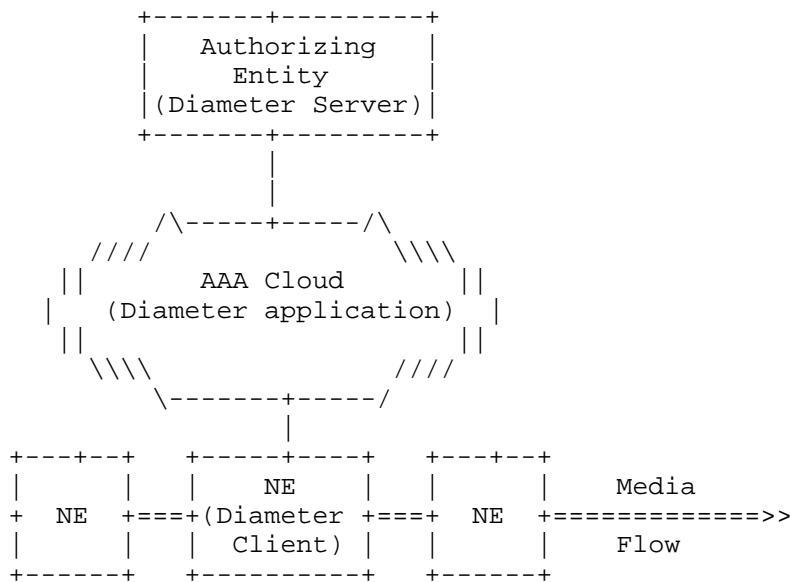


Figure 1: An Architecture Supporting QoS-AAA

Figure 1 depicts NEs through which media flows need to pass, a cloud of AAA servers, and an AE. Note that there may be more than one router that needs to interact with the AAA cloud along the path of a given application flow, although the figure only depicts one for clarity.

In some deployment scenarios, NEs may request authorization through the AAA cloud based on an incoming QoS reservation request. The NE will route the request to a designated AE. The AE will return the result of the authorization decision. In other deployment scenarios, the authorization will be initiated upon dynamic application state, so that the request must be authenticated and authorized based on information from one or more AppSs. After receiving the authorization request from the AppS or the NE, the AE decides the appropriate mode (i.e., Push or Pull). The usage of Push or Pull mode can be determined by the Authorizing Entity either statically or dynamically. Static determination might be based on a configurable defined policy in the Authorizing Entity, while dynamic determination might be based on information received from an application server. For Push mode, the Authorizing Entity needs to identify the appropriate NE(s) to which QoS authorization information needs to be pushed. It might determine this based on information received from the AppS, such as the IP addresses of media flows.

In some deployment scenarios, there is a mapping between access network type and the service logic (e.g., selection of Push or Pull mode and other differentiated handling of the resource admission and control). The access network type might be derived from the authorization request from the AppS or the NE, and in this case, the Authorizing Entity can identify the corresponding service logic based on the mapping.

If the interface between the NEs and the AAA cloud is identical regardless of whether or not the AE communicates with an AppS, routers are insulated from the details of particular applications and need not know that Application Servers are involved. Also, the AAA cloud may also encompass business relationships such as those between network operators and third-party application providers. This enables flexible intra- or inter-domain authorization, accounting, and settlement.

The processing of incoming QoS reservation requests includes three actions: admission control, authorization, and resource reservation.

The admission control function provides information about available resources and determines whether there are enough resources to fulfill the request. Authorization is performed by the Diameter client, which involves contacting an authorization entity through the AAA cloud shown in Section 3. If both checks are successful, the authorized QoS parameters are set in the packet classifier and the packet scheduler. Note that the parameters passed to the Traffic Control function may be different from the ones that requested QoS (depending on the authorization decision). Once the requested resource is granted, the Resource Management function provides accounting information to the AE via the Diameter client.

3.2. Implications of Endpoint QoS Capabilities

3.2.1. Endpoint Categories

The QoS capabilities of Application Endpoints are varied, and can be categorized as follows:

Category 1

A Category 1 Application Endpoint has no QoS capability at either the application or the network level. This type of AppE may set up a connection through application signaling, but it is incapable of specifying resource/QoS requirements through either application- or network-level signaling.

Category 2

A Category 2 Application Endpoint only has QoS capability at the application level. This type of AppE is able to set up a connection through application signaling with certain resource/QoS requirements (e.g., application attributes), but it is unable to signal any resource/QoS requirements at the network level.

Category 3

A Category 3 Application Endpoint has QoS capability at the network level. This type of AppE may set up a connection through application signaling, translate service characteristics into network resource/QoS requirements (e.g., network QoS class) locally, and request the resources through network signaling, e.g., Resource ReSerVation Protocol (RSVP) [RFC2205] or Next Steps in Signaling (NSIS) [NSIS-QOS].

3.2.2. Interaction Modes between the Authorizing Entity and Network Element

Different QoS mechanisms are employed in packet networks. Those QoS mechanisms can be categorized into two schemes: IntServ [RFC2211] [RFC2212] and Diffserv [RFC2474]. In the IntServ scheme, network signaling (e.g., RSVP, NSIS, or link-specific signaling) is commonly used to initiate a request from an AppE for the desired QoS resource. In the Diffserv scheme, QoS resources are provisioned based upon some predefined QoS service classes rather than AppE-initiated, flow-based QoS requests.

It is obvious that the eligible QoS scheme is correlated to the AppE's capability in the context of QoS authorization. Since Category 1 and 2 AppEs cannot initiate the QoS resource requests by means of network signaling, using the current mechanism of the IntServ model to signal QoS information across the network is not applicable to them in general. Depending on network technology and operator requirements, a Category 3 AppE may either make use of network signaling for resource requests or not.

The diversity of QoS capabilities of endpoints and QoS schemes of network technology leads to the distinction on the interaction mode between the QoS authorization system and underlying NEs. When the IntServ scheme is employed by a Category 3 endpoint, the authorization process is typically initiated by an NE when a trigger is received from the endpoint such as network QoS signaling. In the Diffserv scheme, since the NE is unable to request the resource authorization on its own initiative, the authorization process is typically triggered by either the request of AppSs or policies defined by the operator.

As a consequence, two interaction modes are needed in support of different combinations of QoS schemes and endpoint's QoS capabilities: Push mode and Pull mode.

Push mode

The QoS authorization process is triggered by AppSs or local network conditions (e.g., time of day on resource usage and QoS classes), and the authorization decisions are installed by the AE to the network element on its own initiative without explicit request. In order to support Push mode, the AE (i.e., Diameter server) should be able to initiate a Diameter authorization session to communicate with the NE (i.e., Diameter client) without any preestablished connection from the network element.

Pull mode

The QoS authorization process is triggered by the network signaling received from end-user equipment or by a local event in the NE according to pre-configured policies, and authorization decisions are produced upon the request of the NE. In order to support Pull mode, the NE (i.e., Diameter client) will initiate a Diameter authorization session to communicate with the Authorizing Entity (i.e., Diameter server).

For Category 1 and 2 Application Endpoints, Push mode is REQUIRED. For a Category 3 AppE, either Push mode or Pull mode MAY be used.

Push mode is applicable to certain networks, for example, Cable network, DSL, Ethernet, and Diffserv-enabled IP/MPLS. Pull mode is more appropriate to IntServ-enabled IP networks or certain wireless networks such as the General Packet Radio Service (GPRS) networks defined by the Third Generation Partnership Project (3GPP). Some networks (for example, Worldwide Interoperability for Microwave Access (WiMAX)) may require both Push and Pull modes.

3.3. Authorization Schemes

3.3.1. Pull Mode Schemes

Three types of basic authorization schemes for Pull mode exist: one type of two-party scheme and two types of three-party schemes. The notation adopted here is in respect to the entity that performs the QoS authorization (QoS Authz). The authentication of the QoS requesting entity might be done at the NE as part of the QoS signaling protocol, or by an off-path protocol (on the application layer or for network access authentication) or the AE might be contacted with a request for authentication and authorization of the QoS requesting entity. From the Diameter QoS application's point of view, these schemes differ in type of information that need to be carried. Here we focus on the "Basic Three-Party Scheme" (see Figure 3) and the "Token-Based Three-Party Scheme" (see Figure 4). In the "Two-Party Scheme", the QoS RRE is authenticated by the NE and the authorization decision is made either locally at the NE itself or offloaded to a trusted entity (most likely within the same administrative domain). In the two-party case, no Diameter QoS protocol interaction is required.




Figure 3: Three-Party Scheme

In the "Basic Three-Party Scheme", a QoS reservation request that arrives at the NE is forwarded to the Authorizing Entity (e.g., in the user's home network), where the authorization decision is made. As shown, financial settlement -- a business relationship, such as a roaming agreement -- between the visited network and the home network ensures that the visited network is compensated for the resources consumed by the user via the home network.

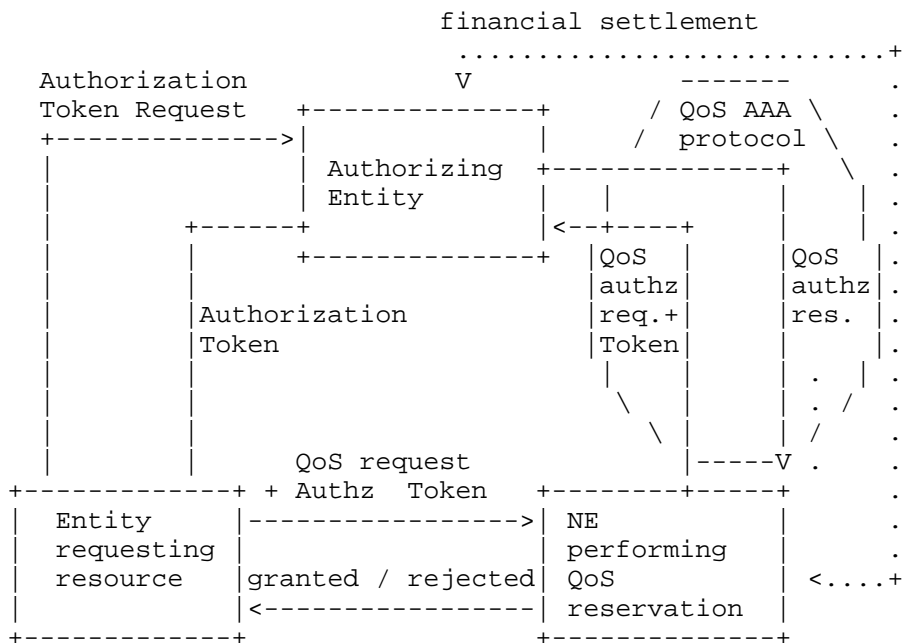


Figure 4: Token-Based Three-Party Scheme

The "Token-Based Three-Party Scheme" is applicable to environments where a previous protocol interaction is used to request authorization tokens to assist the authorization process at the NE or the AE [RFC3521].

The QoS RRE may be involved in an application-layer protocol interaction, for example, using SIP [RFC3313], with the AE. As part of this interaction, authentication and authorization at the application layer might take place. As a result of a successful authorization decision, which might involve the user's home AAA server, an authorization token is generated by the AE (e.g., the SIP proxy and an entity trusted by the SIP proxy) and returned to the end-host for inclusion into the QoS signaling protocol. The authorization token will be used by an NE that receives the QoS signaling message to authorize the QoS request. Alternatively, the Diameter QoS application will be used to forward the authorization token to the user's home network. The authorization token allows for the authorization decision performed at the application layer to be associated with a corresponding QoS signaling session. Note that the authorization token might either refer to established state concerning the authorization decision or the token might itself carry the authorized parameters (protected by a digital signature or a keyed message digest to prevent tampering). In the latter case, the

authorization token may contain several pieces of information pertaining to the authorized application session, but at minimum it should contain:

- o An identifier for the AE (for example, an AppS) that issued the authorization token;
- o An identifier referring to a specific application protocol session for which the token was issued; and
- o A keyed message digest or digital signature protecting the content of the authorization token.

A possible structure for the authorization token and the policy element carrying it are proposed in the context of RSVP [RFC3520].

In the scenario mentioned above, where the QoS resource requesting entity is involved in an application-layer protocol interaction with the AE, it may be worthwhile to consider a token-less binding mechanism also. The application-layer protocol interaction may have indicated the transport port numbers at the QoS RRE where it might receive media streams (for example, in SIP/SDP [RFC4566] signaling, these port numbers are advertised). The QoS RRE may also use these port numbers in some IP filter indications to the NE performing QoS reservation so that it may properly tunnel the inbound packets. The NE performing QoS reservation will forward the QoS resource requesting entity's IP address and the IP filter indications to the AE in the QoS authorization request. The AE will use the QoS RRE's IP address and the port numbers in the IP filter indication, which will match the port numbers advertised in the earlier application-layer protocol interaction, to identify the right piece of policy information to be sent to the NE performing the QoS reservation in the QoS Authorization response.

3.3.2. Push Mode Schemes

Push mode can be further divided into two types: endpoint-initiated and network-initiated. In the former case, the authorization process is triggered by AppS in response to an explicit QoS request from an endpoint through application signaling, e.g., SIP; in the latter case, the authorization process is triggered by the AppS without an explicit QoS request from an endpoint.

In the endpoint-initiated scheme, the QoS RRE (i.e., the AppE) determines the required application-level QoS and sends a QoS request through an application signaling message. The AppS will extract application-level QoS information and trigger the authorization process to the AE. In the network-initiated scheme, the AE and/or

AppS should derive and determine the QoS requirements according to application attribute, subscription, and endpoint capability when the endpoint does not explicitly indicate the QoS attributes. The AE makes an authorization decision based on application-level QoS information, network policies, end-user subscription, network resource availability, etc., and installs the decision to the NE directly.

A Category 1 AppE requires network-initiated Push mode and a Category 2 AppE may use either type of Push Mode.

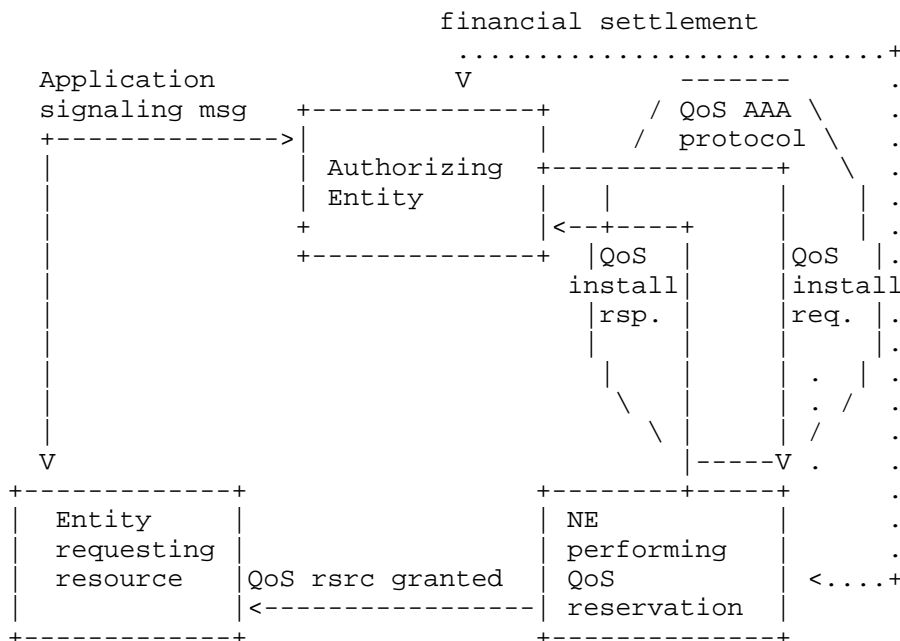


Figure 5: Scheme for Push Mode

3.4. QoS Application Requirements

A QoS application must meet a number of requirements applicable to a diverse set of networking environments and services. It should be compatible with different deployment scenarios having specific QoS signaling models and security issues. Satisfying the requirements listed below while interworking with QoS signaling protocols, a Diameter QoS application should accommodate the capabilities of the QoS signaling protocols rather than introduce functional requirements on them. A list of requirements for a QoS authorization application is provided here:

Identity-based Routing

The Diameter QoS application MUST route AAA requests to the Authorizing Entity, based on the provided identity of the QoS requesting entity or the identity of the AE encoded in the provided authorization token.

Flexible Authentication Support

The Diameter QoS application MUST support a variety of different authentication protocols for verification of authentication information present in QoS signaling messages. The support for these protocols MAY be provided indirectly by tying the signaling communication for QoS to a previous authentication protocol exchange (e.g., using network access authentication).

Making an Authorization Decision

The Diameter QoS application MUST exchange sufficient information between the AE and the enforcing entity (and vice versa) to compute an authorization decision and to execute this decision.

Triggering an Authorization Process

The Diameter QoS application MUST allow periodic and event-triggered execution of the authorization process, originated at the enforcing entity or even at the AE.

Associating QoS Reservations and Application State

The Diameter QoS application MUST carry information sufficient for an AppS to identify the appropriate application session and associate it with a particular QoS reservation.

Dynamic Authorization

It MUST be possible for the Diameter QoS application to push updates towards the NE(s) from Authorizing Entities.

Bearer Gating

The Diameter QoS application MUST allow the AE to gate (i.e., enable/disable) authorized application flows based on, e.g., application state transitions.

Accounting Records

The Diameter QoS application MAY define QoS accounting records containing duration, volume (byte count) usage information, and a description of the QoS attributes (e.g., bandwidth, delay, loss rate) that were supported for the flow.

Sending Accounting Records

The NE SHOULD be able to send accounting records for a particular QoS reservation state to an accounting entity.

Failure Notification

The Diameter QoS application MUST allow the NE to report failures, such as loss of connectivity due to movement of a mobile node or other reasons for packet loss, to the Authorizing Entity.

Accounting Correlation

The Diameter QoS application MAY support the exchange of sufficient information to allow for correlation between accounting records generated by the NEs and accounting records generated by an AppS.

Interaction with Other AAA Applications

Interaction with other AAA applications, such as the Diameter Network Access Server Application [RFC4005], may be required for exchange of authorization, authentication, and accounting information.

In deployment scenarios where authentication of the QoS reservation requesting entity (e.g., the user) is done by means outside the Diameter QoS application protocol interaction, the AE is contacted only with a request for QoS authorization. Authentication might have taken place already via the interaction with the Diameter application [RFC4005] or as part of the QoS signaling protocol (e.g., Transport Layer Security (TLS) [RFC5246] in the General Internet Signaling Transport (GIST) protocol [NSIS-NTLP]).

Authentication of the QoS reservation requesting entity to the AE is necessary if a particular Diameter QoS application protocol cannot be related (or if there is no intention to relate it) to a prior authentication. In this case, the AE MUST authenticate the QoS reservation requesting entity in order to authorize the QoS request as part of the Diameter QoS protocol interaction.

This document refers to three types of sessions that need to be properly correlated.

QoS Signaling Session

The time period during which a QoS signaling protocol establishes, maintains, and deletes a QoS reservation state at the QoS network element is referred to as a QoS signaling session. Different QoS signaling protocols use different ways to identify QoS signaling sessions. The same applies to different usage environments. Currently, this document supports three types of QoS session identifiers, namely a signaling session id (e.g., the Session Identifier used by the NSIS protocol suite), a flow id (e.g., identifier assigned by an application to a certain flow as used in the 3GPP), and a flow description based on the IP parameters of the flow's endpoints.

Diameter Authorization Session

The time period for which a Diameter server authorizes a requested service (i.e., QoS resource reservation) is referred to as a Diameter authorization session. It is identified by a Session-Id included in all Diameter messages used for management of the authorized service (initial authorization, re-authorization, termination), see [RFC3588].

Application-Layer Session

The application-layer session identifies the duration of an application-layer service that requires provision of a certain QoS. An application-layer session identifier is provided by the QoS requesting entity in the QoS signaling messages, for example as part of the authorization token. In general, the application session identifier is opaque to the QoS-aware NEs. It is included in the authorization request message sent to the AE and helps it to correlate the QoS authorization request to the application session state information.

Correlating these sessions is done at each of the three involved entities: The QoS requesting entity correlates the application with the QoS signaling sessions. The QoS NE correlates the QoS signaling session with the Diameter authorization sessions. The AE SHOULD bind the information about the three sessions together. Note that in certain scenarios, not all of the sessions are present. For example, the application session might not be visible to the QoS signaling protocol directly if there is no binding between the application session and the QoS requesting entity using the QoS signaling protocol.

4. QoS Application Session Establishment and Management

4.1. Parties Involved

Authorization models supported by this application include three parties:

- o Resource Requesting Entity
- o Network Elements (Diameter QoS application (DQA) client)
- o Authorizing Entity (Diameter QoS application (DQA) server)

Note that the QoS RRE is only indirectly involved in the message exchange. This entity provides the trigger to initiate the Diameter QoS protocol interaction by transmitting QoS signaling messages. The Diameter QoS application is only executed between the Network Element (i.e., DQA client) and the Authorizing Entity (i.e., DQA server).

The QoS RRE may communicate with the AE using application-layer signaling for the negotiation of service parameters. As part of this application-layer protocol interaction, for example using SIP, authentication and authorization might take place. This message exchange is, however, outside the scope of this document. The protocol communication between the QoS resource requesting entity and the QoS NE might be accomplished using the NSIS protocol suite, RSVP, or a link-layer signaling protocol. A description of these protocols is also outside the scope of this document.

4.2. Session Establishment

Pull and Push modes use a different set of command codes for session establishment. For other operations, such as session modification and termination, they use the same set of command codes.

The selection of Pull mode or Push mode operation is based on the trigger of the QoS authorization session. When a QoS-Authorization-Request (QAR, see Section 5.1) message with a new Session-Id is received, the AE operates in Pull mode; when other triggers are received, the AE operates in Push mode. Similarly, when a QoS-Install-Request (QIR, see Section 5.3) with a new Session-Id is received, the NE operates in Push mode; when other triggers are received, the NE operates in Pull mode.

The QoS authorization session is typically established per subscriber base (i.e., all requests with the same User-ID), but it is also possible to be established on a per node or per request base. The concurrent sessions between an NE and an AE are identified by different Session-Ids.

4.2.1. Session Establishment for Pull Mode

A request for a QoS reservation or local events received by an NE can trigger the initiation of a Diameter QoS authorization session. The NE converts the required objects from the QoS signaling message to Diameter AVPs and generates a QAR message.

Figure 6 shows the protocol interaction between a Resource Requesting Entity, a Network Element, and the Authorizing Entity.

The AE's identity, information about the application session and/or identity and credentials of the QoS RRE, requested QoS parameters, and the signaling session identifier and/or QoS-enabled data flows identifiers MAY be encapsulated into respective Diameter AVPs and included in the Diameter message sent to the AE. The QAR is sent to a Diameter server that can be either the home server of the QoS requesting entity or an AppS.

QoS-Specific Input Data	Diameter AVPs
Authorizing Entity ID (e.g., Destination-Host taken from authorization token, Destination-Realm, or derived from the Network Access Identifier (NAI) of the QoS requesting entity) Authorization Token Credentials of the QoS requesting entity QoS-Resources (including QoS parameters)	Destination-Host Destination-Realm QoS-Authorization-Data User-Name

Table 1: Mapping Input Data to QoS AVPs -- Pull Mode

Authorization processing starts at the Diameter QoS server when it receives the QAR. Based on the information in the QoS-Authentication-Data, User-Name, and QoS-Resources AVPs, the server determines the authorized QoS resources and flow state (enabled/disabled) from locally available information (e.g., policy information that may be previously established as part of an application-layer signaling exchange or the user's subscription profile). The QoS-Resources AVP is defined in [RFC5777]. The authorization decision is then reflected in the response returned to the Diameter client with the QoS-Authorization-Answer (QAA) message.

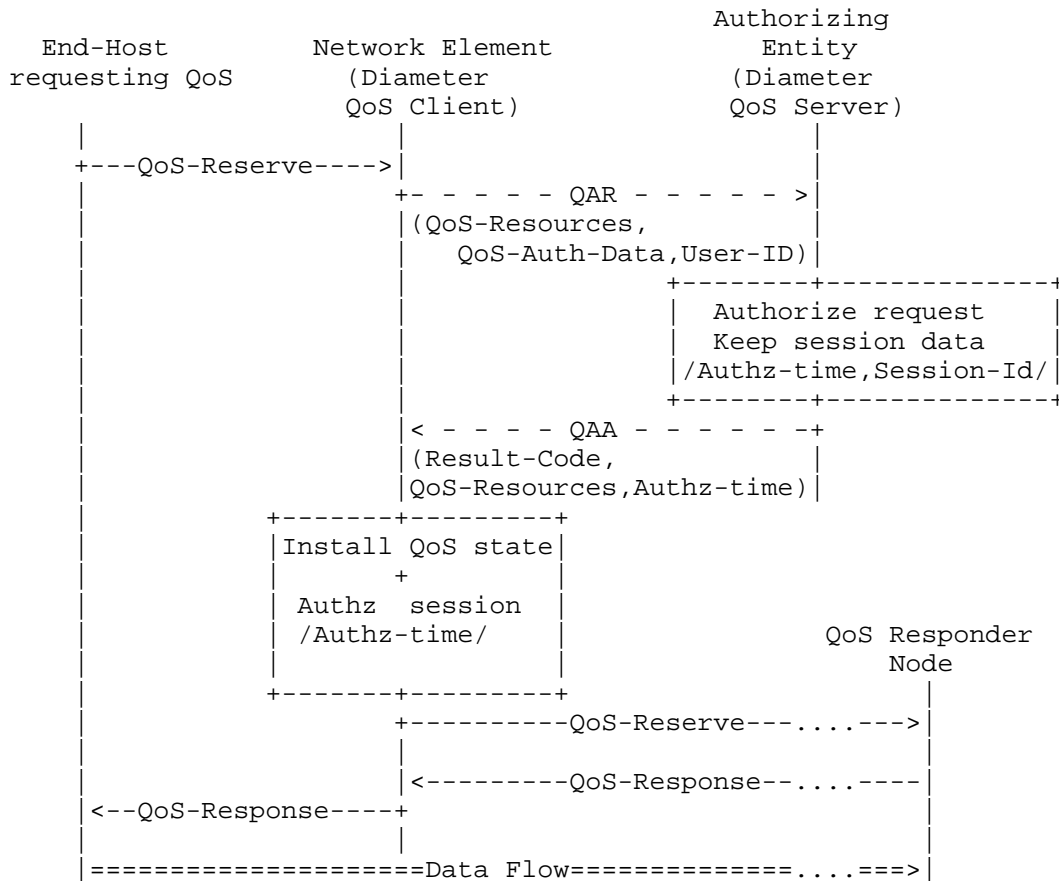


Figure 6: Initial QoS Request Authorization for Pull Mode

The Authorizing Entity keeps authorization session state and SHOULD save additional information for management of the session (e.g., Signaling-Session-Id, authentication data) as part of the session state information.

The final result of the authorization request is provided in the Result-Code AVP of the QAA message sent by the Authorizing Entity. In the case of successful authorization (i.e., Result-Code = DIAMETER_LIMITED_SUCCESS (see Section 7.1)), information about the authorized QoS resources and the status of the authorized flow (enabled/disabled) is provided in the QoS-Resources AVP of the QAA message. The QoS information provided via the QAA is installed by the QoS Traffic Control function of the NE. The value

DIAMETER_LIMITED_SUCCESS indicates that the AE expects confirmation via another QAR message for successful QoS resource reservation and for final reserved QoS resources (see below).

One important piece of information returned from the Authorizing Entity is the authorization lifetime (carried inside the QAA). The authorization lifetime allows the NE to determine how long the authorization decision is valid for this particular QoS reservation. A number of factors may influence the authorized session duration, such as the user's subscription plan or the currently available credits at the user's account (see Section 8). The authorization duration is time-based, as specified in [RFC3588]. For an extension of the authorization period, a new QoS-Authorization-Request/Answer message exchange SHOULD be initiated. Further aspects of QoS authorization session maintenance are discussed in Sections 4.3, 4.4, and 8.

The indication of a successful QoS reservation and activation of the data flow is provided by the transmission of a QAR message, which reports the parameters of the established QoS state: reserved resources, duration of the reservation, and identification of the QoS enabled flow/QoS signaling session. The Diameter QoS server acknowledges the reserved QoS resources with the QA Answer (QAA) message where the Result-Code is set to 'DIAMETER_SUCCESS'. Note that the reserved QoS resources reported in this QAR message MAY be different than those authorized with the initial QAA message, due to the QoS-signaling-specific behavior (e.g., receiver-initiated reservations with One-Path-With-Advertisements) or specific process of QoS negotiation along the data path.

4.2.2. Session Establishment for Push Mode

The Diameter QoS server in the AE initiates a Diameter QoS authorization session upon the request for a QoS reservation triggered by application-layer signaling or by local events, and generates a QoS-Install-Request (QIR) message to the Diameter QoS client in the NE in which it maps required objects to Diameter payload objects.

Figure 7 shows the protocol interaction between the AE, a Network Element, and an RRE.

The NE's identity, information about the application session and/or identity and credentials of the QoS resource requesting entity, requested QoS parameters, and signaling session identifier and/or QoS enabled data flows identifiers MAY be encapsulated into respective Diameter AVPs and included in the Diameter message sent from a Diameter QoS server in the Authorizing Entity to a Diameter QoS

client in the NE. This requires that the AE has knowledge of specific information for allocating and identifying the NE that should be contacted and the data flow for which the QoS reservation should be established. This information can be statically configured or dynamically discovered, see Section 4.2.3 for details.

QoS-Specific Input Data	Diameter AVPs
Network Element ID Authorization Token Credentials of the QoS requesting entity QoS-Resources (including QoS parameters)	Destination-Host Destination-Realm QoS-Authorization-Data User-Name

Table 2: Mapping Input Data to QoS AVPs -- Push Mode

Authorization processing starts at the Diameter QoS server when it receives a request from an RRE through an AppS (e.g., SIP Invite) or is triggered by a local event (e.g., a pre-configured timer). Based on the received information, the server determines the authorized QoS resources and flow state (enabled/disabled) from locally available information (e.g., policy information that may be previously established as part of an application-layer signaling exchange, or the user's subscription profile). The authorization decision is then reflected in the QoS-Install-Request (QIR) message to the Diameter QoS client.

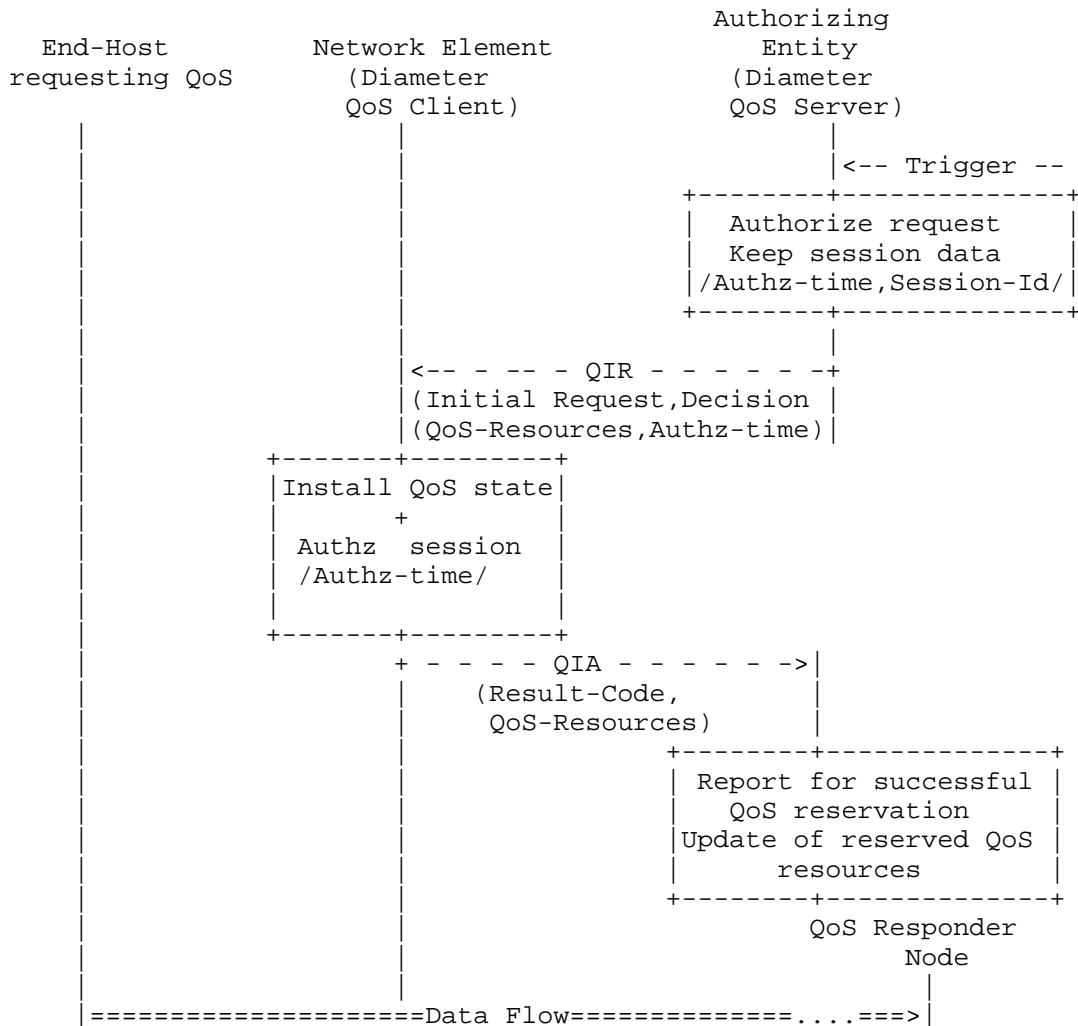


Figure 7: Initial QoS Request Authorization for Push Mode

The AE keeps authorization session state and SHOULD save additional information for management of the session (e.g., Signaling-Session-Id, authentication data) as part of the session state information.

The final result of the authorization decision is provided in the QoS-Resources AVP of the QIR message sent by the AE. The QoS information provided via the QIR is installed by the QoS Traffic Control function of the NE.

One important piece of information from the AE is the authorization lifetime (carried inside the QIR). The authorization lifetime allows the NE to determine how long the authorization decision is valid for this particular QoS reservation. A number of factors may influence the authorized session duration, such as the user's subscription plan or the currently available credits at the user's account (see Section 8). The authorization duration is time-based as specified in [RFC3588]. For an extension of the authorization period, a new QoS-Install-Request/Answer message or QoS-Authorization-Request/Answer message exchange SHOULD be initiated. Further aspects of QoS authorization session maintenance are discussed in Sections 4.3, 4.4, and 8.

The indication of QoS reservation and activation of the data flow can be provided by the QoS-Install-Answer message immediately. In the case of successful enforcement, the Result-Code (= DIAMETER_SUCCESS, (see Section 7.1)) information is provided in the QIA message. Note that the reserved QoS resources reported in the QIA message may be different than those initially authorized with the QIR message, due to the QoS signaling-specific behavior (e.g., receiver-initiated reservations with One-Path-With-Advertisements) or specific process of QoS negotiation along the data path. In the case that Multiple AEs control the same NE, the NE should make the selection on the authorization decision to be enforced based on the priority of the request.

4.2.3. Discovery and Selection of Peer Diameter QoS Application Node

The Diameter QoS application node may obtain information of its peer nodes (e.g., Fully-Qualified Domain Name (FQDN), IP address) through static configuration or dynamic discovery as described in Section 5.2 of [RFC3588]. In particular, the NE shall perform the relevant operation for Pull mode; the AE shall perform the relevant operations for Push mode.

Upon receipt of a trigger to initiate a new Diameter QoS authorization session, the Diameter QoS application node selects and retrieves the location information of the peer node that is associated with the affected user based on some index information provided by the RRE. For instance, it can be the Authorization Entity's ID stored in the authorization token, the end-user identity (e.g., NAI [RFC4282]), or a globally routable IP address.

4.3. Session Re-Authorization

Client- and server-side initiated re-authorizations are considered in the design of the Diameter QoS application. Whether the re-authorization events are transparent for the resource requesting

entity or result in specific actions in the QoS signaling protocol is outside the scope of the Diameter QoS application. It is directly dependent on the capabilities of the QoS signaling protocol.

There are a number of options for policy rules according to which the NE (AAA client) contacts the AE for re-authorization. These rules depend on the semantics and contents of the QAA message sent by the AE:

- a. The QAA message contains the authorized parameters of the flow and its QoS and sets their limits (presumably upper). With these parameters, the AE specifies the services that the NE can provide and for which it will be financially compensated. Therefore, any change or request for change of the parameters of the flow and its QoS that do not conform to the authorized limits requires contacting the AE for authorization.
- b. The QAA message contains authorized parameters of the flow and its QoS. The rules that determine whether parameters' changes require re-authorization are agreed out of band, based on a Service Level Agreement (SLA) between the domains of the NE and the AE.
- c. The QAA message contains the authorized parameters of the flow and its QoS. Any change or request for change of these parameters requires contacting the AE for re-authorization.
- d. In addition to the authorized parameters of the flow and its QoS, the QAA message contains policy rules that determine the NEs actions in case of a change or a request for change in authorized parameters.

Provided options are not exhaustive. Elaborating on any of the listed approaches is deployment/solution specific and is not considered in the current document.

In addition, the AE may use an RAR (Re-Authorization-Request) to perform re-authorization with the authorized parameters directly when the re-authorization is triggered by service request or local events/policy rules.

4.3.1. Client-Side Initiated Re-Authorization

The AE provides the duration of the authorization session as part of the QoS-Authorization-Answer (QAA) message. At any time before the expiration of this period, a new QoS-Authorization-Request (QAR) message MAY be sent to the AE. The transmission of the QAR MAY be triggered when the NE receives a QoS signaling message that requires

modification of the authorized parameters of an ongoing QoS session, or authorization lifetime expires.

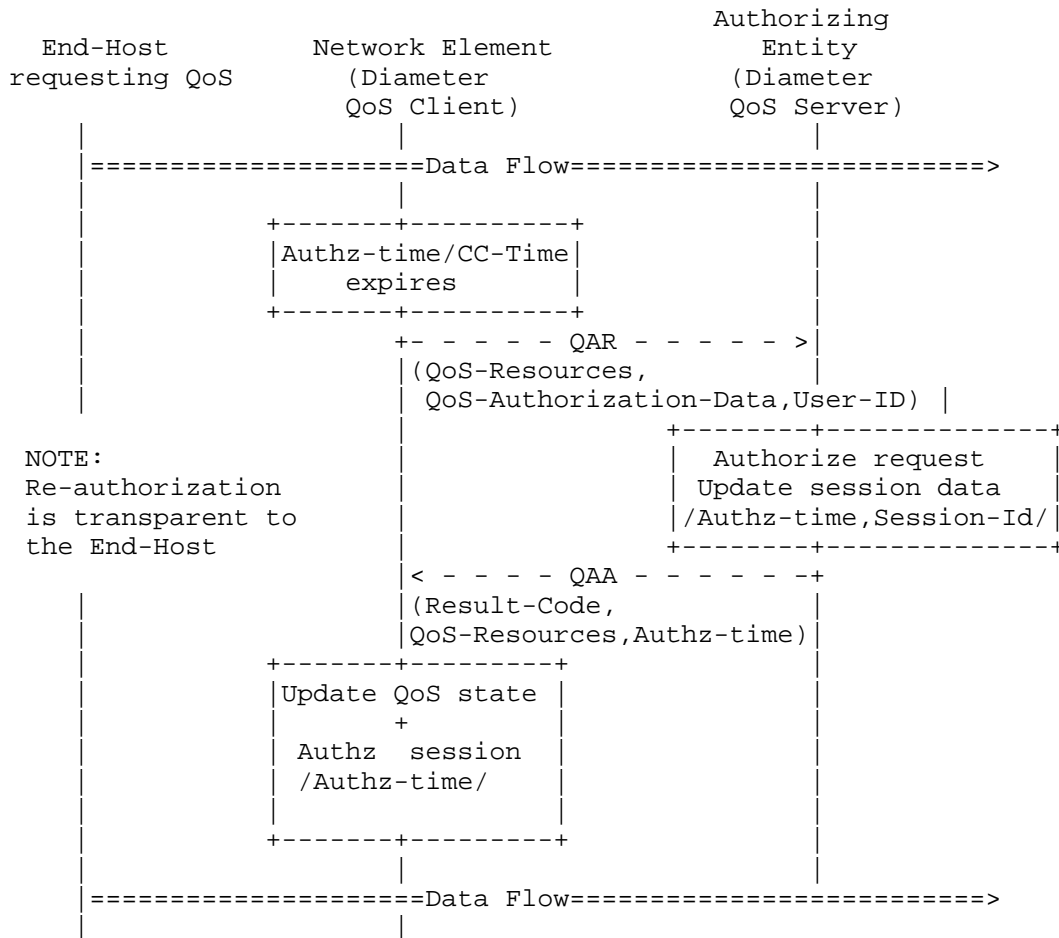


Figure 8: Client-side Initiated QoS Re-Authorization

4.3.2. Server-Side Initiated Re-Authorization

The AE MAY initiate a QoS re-authorization by issuing a Re-Authorization-Request (RAR) message as defined in the Diameter base protocol [RFC3588], which may include the parameters of the re-authorized QoS state: reserved resources, duration of the reservation, identification of the QoS-enabled flow/QoS signaling session for re-installation of the resource state by the QoS Traffic Control function of the NE.

An NE that receives such an RAR message with Session-Id matching a currently active QoS session acknowledges the request by sending the Re-Auth-Answer (RAA) message towards the AE.

If the RAR does not include any parameters of the re-authorized QoS state, the NE MUST initiate a QoS re-authorization by sending a QoS-Authorization-Request (QAR) message towards the AE.

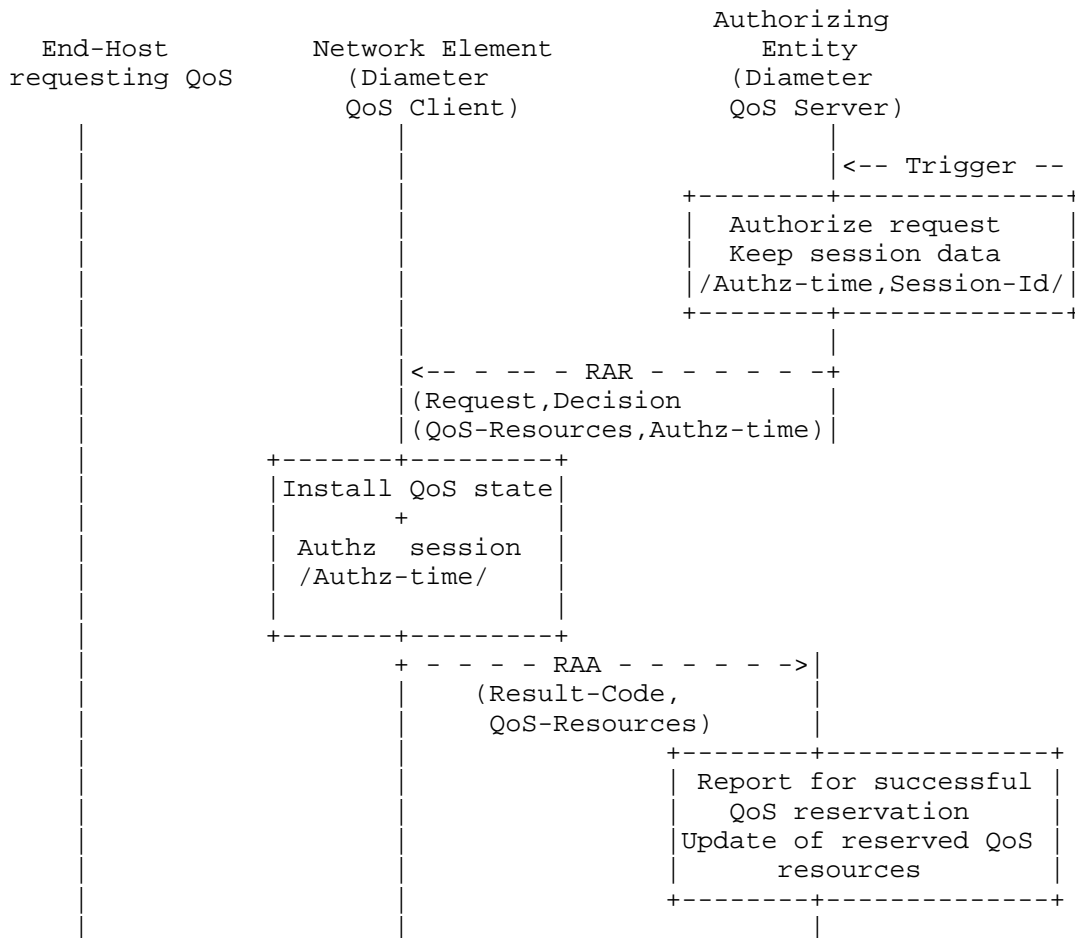


Figure 9: Server-Side Initiated QoS Re-Authorization

4.4. Session Termination

4.4.1. Client-Side Initiated Session Termination

The authorization session for an installed QoS reservation state MAY be terminated by the Diameter client by sending a Session-Termination-Request (STR) message to the Diameter server with a response Session-Termination-Acknowledgement (STA) message. This is a Diameter base protocol function and it is defined in [RFC3588]. Session termination can be caused by a QoS signaling message requesting deletion of the existing QoS reservation state, or it can be caused as a result of a soft-state expiration of the QoS reservation state.




Figure 10: Client-Side Initiated Session Termination

4.4.2. Server-Side Initiated Session Termination

At any time during a session, the AE MAY send an Abort-Session-Request (ASR) message to the NE. This is a Diameter base protocol function and it is defined in [RFC3588]. Possible reasons for initiating the ASR message to the NE are insufficient credits or session termination at the application layer. The ASR message results in termination of the authorized session, release of the

reserved resources at the NE, and transmission of an appropriate QoS signaling message indicating a notification to other Network Elements aware of the signaling session.

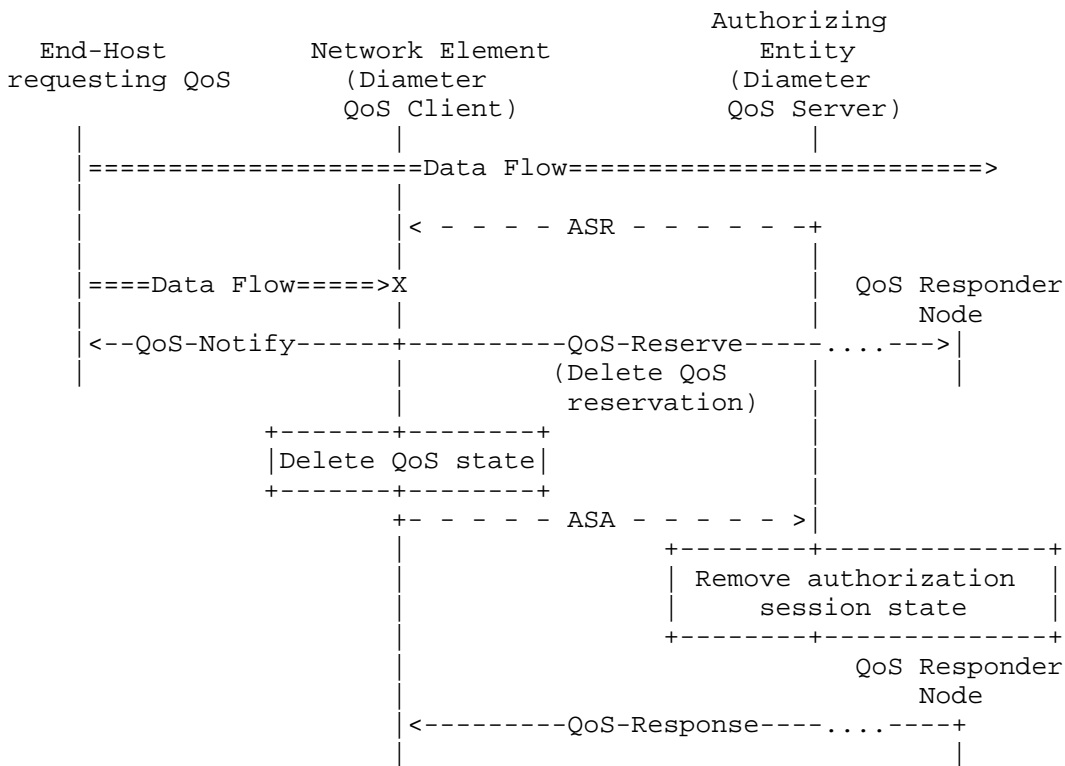


Figure 11: Server-Side Initiated Session Termination

5. QoS Application Messages

The Diameter QoS application requires the definition of new mandatory AVPs and Command-Codes (see Section 3 of [RFC3588]). Four new Diameter messages are defined along with Command-Codes whose values MUST be supported by all Diameter implementations that conform to this specification.

Command Name	Abbrev.	Code	Reference
QoS-Authorization-Request	QAR	326	Section 5.1
QoS-Authorization-Answer	QAA	326	Section 5.2
QoS-Install-Request	QIR	327	Section 5.3
QoS-Install-Answer	QIA	327	Section 5.4

Table 3: Diameter QoS Commands

In addition, the following Diameter base protocol messages are used in the Diameter QoS application:

Command-Name	Abbrev.	Code	Reference
Re-Auth-Request	RAR	258	[RFC3588]
Re-Auth-Answer	RAA	258	[RFC3588]
Abort-Session-Request	ASR	274	[RFC3588]
Abort-Session-Answer	ASA	274	[RFC3588]
Session-Term-Request	STR	275	[RFC3588]
Session-Term-Answer	STA	275	[RFC3588]

Table 4: Diameter Base Commands

Diameter nodes conforming to this specification MAY advertise support for the Diameter QoS application by including the value of 9 in the Auth-Application-Id or the Acct-Application-Id AVP of the Capabilities-Exchange-Request and Capabilities-Exchange-Answer commands, see [RFC3588].

The value of 9 MUST be used as the Application-Id in all QAR/QAA and QIR/QIA commands.

The value of zero (0) SHOULD be used as the Application-Id in all STR/STA, ASR/ASA, and RAR/RAA commands.

5.1. QoS-Authorization Request (QAR)

The QoS-Authorization-Request (QAR) message, indicated by the Command-Code field (see Section 3 of [RFC3588]) being set to 326 and the 'R' bit being set in the Command Flags field, is used by NEs to request quality of service related resource authorization for a given flow.

The QAR message MUST carry information for signaling session identification, AE identification, information about the requested QoS, and the identity of the QoS requesting entity. In addition, depending on the deployment scenario, an authorization token and credentials of the QoS requesting entity SHOULD be included.

The message format is defined as follows:

```
<QoS-Authorization-Request> ::= < Diameter Header: 326, REQ, PXY >
    < Session-Id >
    { Auth-Application-Id }
    { Origin-Host }
    { Origin-Realm }
    { Destination-Realm }
    { Auth-Request-Type }
    [ Destination-Host ]
    [ User-Name ]
    * [ QoS-Resources ]
    [ QoS-Authorization-Data ]
    [ Bound-Auth-Session-Id ]
    * [ AVP ]
```

5.2. QoS-Authorization-Answer (QAA)

The QoS-Authorization-Answer (QAA) message, indicated by the Command-Code field being set to 326 and the 'R' bit being cleared in the Command Flags field, is sent in response to the QoS-Authorization-Request (QAR) message. If the QoS authorization request is successfully authorized, the response will include the AVPs to allow authorization of the QoS resources and transport plane gating information.

The message format is defined as follows:

```
<QoS-Authorization-Answer> ::= < Diameter Header: 326, PXY >
    < Session-Id >
    { Auth-Application-Id }
    { Auth-Request-Type }
    { Result-Code }
    { Origin-Host }
    { Origin-Realm }
    * [ QoS-Resources ]
    [ Acct-Multisession-Id ]
    [ Session-Timeout ]
    [ Authorization-Session-Lifetime ]
    [ Authorization-Grace-Period ]
    * [ AVP ]
```

5.3. QoS-Install Request (QIR)

The QoS-Install Request (QIR) message, indicated by the Command-Code field being set to 327 and the 'R' bit being set in the Command Flags field, is used by the AE to install or update the QoS parameters and the flow state of an authorized flow at the transport plane element.

The message MUST carry information for signaling-session identification or identification of the flow to which the provided QoS rules apply, identity of the transport plane element, description of provided QoS parameters, flow state, and duration of the provided authorization.

The message format is defined as follows:

```
<QoS-Install-Request> ::= < Diameter Header: 327, REQ, PXY >
                          < Session-Id >
                          { Auth-Application-Id }
                          { Origin-Host }
                          { Origin-Realm }
                          { Destination-Realm }
                          { Auth-Request-Type }
                          [ Destination-Host ]
                          * [ QoS-Resources ]
                          [ Session-Timeout ]
                          [ Authorization-Session-Lifetime ]
                          [ Authorization-Grace-Period ]
                          [ Authorization-Session-Volume ]
                          * [ AVP ]
```

5.4. QoS-Install Answer (QIA)

The QoS-Install Answer (QIA) message, indicated by the Command-Code field being set to 327 and the 'R' bit being cleared in the Command Flags, field is sent in response to the QoS-Install Request (QIR) message for confirmation of the result of the installation of the provided QoS reservation instructions.

The message format is defined as follows:

```
<QoS-Install-Answer> ::= < Diameter Header: 327, PXY >
                          < Session-Id >
                          { Auth-Application-Id }
                          { Origin-Host }
                          { Origin-Realm }
                          { Result-Code }
                          * [ QoS-Resources ]
                          * [ AVP ]
```


5.5. Re-Auth-Request (RAR)

The Re-Auth-Request (RAR) message, indicated by the Command-Code field being set to 258 and the 'R' bit being set in the Command Flags field, is sent by the AE to the NE in order to initiate the QoS re-authorization from the DQA server side.

If the RAR command is received by the NE without any parameters of the re-authorized QoS state, the NE MUST initiate a QoS re-authorization by sending a QoS-Authorization-Request (QAR) message towards the AE.

The message format is defined as follows:

```
<RAR> ::= < Diameter Header: 258, REQ, PXY >
         < Session-Id >
         { Origin-Host }
         { Origin-Realm }
         { Destination-Realm }
         { Destination-Host }
         { Auth-Application-Id }
         { Re-Auth-Request-Type }
         [ User-Name ]
         [ Origin-State-Id ]
         * [ Proxy-Info ]
         * [ Route-Record ]
         * [ QoS-Resources ]
         [ Session-Timeout ]
         [ Authorization-Session-Lifetime ]
         [ Authorization-Grace-Period ]
         [ Authorization-Session-Volume ]
         * [ AVP ]
```

5.6. Re-Auth-Answer (RAA)

The Re-Auth-Answer (RAA) message, indicated by the Command-Code field being set to 258 and the 'R' bit being cleared in the Command Flags field, is sent by the NE to the AE in response to the RAR command.

The message format is defined as follows:

```
<RAA> ::= < Diameter Header: 258, PXY >
          < Session-Id >
          { Result-Code }
          { Origin-Host }
          { Origin-Realm }
          [ User-Name ]
          [ Origin-State-Id ]
          [ Error-Message ]
          [ Error-Reporting-Host ]
          * [ Failed-AVP ]
          * [ Redirect-Host ]
          [ Redirect-Host-Usage ]
          [ Redirect-Host-Max-Cache-Time ]
          * [ Proxy-Info ]
          * [ QoS-Resources ]
          * [ AVP ]
```

6. QoS Application State Machine

The QoS application defines its own state machine that is based on the authorization state machine defined in Section 8.1 of the Diameter base protocol ([RFC3588]). The QoS state machine uses its own messages, as defined in Section 5, and QoS AVPs, as defined in Section 7.

6.1. Supplemented States for Push Mode

Using the Diameter base protocol state machine as a basis, the following states are supplemented to the first two state machines in which the session state is maintained on the server. These MUST be supported in any QoS application implementations in support of server-initiated Push mode (see Section 4.2.2).

The following states are supplemented to the state machine on the server when state is maintained on the client, as defined in Section 8.1 of the Diameter base protocol[RFC3588]:

SERVER, STATEFUL			
State	Event	Action	New State
Idle	An application or local event triggers an initial QoS request to the server	Send QIR initial request	Pending
Pending	Received QIA with a failed Result-Code	Clean up	Idle
Pending	Received QIA with Result-Code = SUCCESS	Update session	Open
Pending	Error in processing received QIA with Result-Code = SUCCESS	Send ASR	Discon

The following states are supplemented to the state machine on the client when state is maintained on the server, as defined in Section 8.1 of the Diameter base protocol [RFC3588]:

CLIENT, STATEFUL			
State	Event	Action	New State
Idle	QIR initial request received and successfully processed	Send QIA initial answer, reserve resources	Open
Idle	QIR initial request received but not successfully processed	Send QIA initial answer with Result-Code != SUCCESS	Idle

7. QoS Application AVPs

Each of the AVPs identified in the QoS-Authorization-Request/Answer and QoS-Install-Request/Answer messages and the assignment of their value(s) is given in this section.

7.1. Reused Base Protocol AVPs

The QoS application uses a number of session management AVPs, defined in the base protocol ([RFC3588]).

Attribute Name	AVP Code	Reference [RFC3588]
Origin-Host	264	Section 6.3
Origin-Realm	296	Section 6.4
Destination-Host	293	Section 6.5
Destination-Realm	283	Section 6.6
Auth-Application-Id	258	Section 6.8
Result-Code	268	Section 7.1
Auth-Request-Type	274	Section 8.7
Session-Id	263	Section 8.8
Authorization-Lifetime	291	Section 8.9
Auth-Grace-Period	276	Section 8.10
Session-Timeout	27	Section 8.13
User-Name	1	Section 8.14

The Auth-Application-Id AVP (AVP Code 258) is assigned by IANA to Diameter applications. The value of the Auth-Application-Id for the Diameter QoS application is 9.

7.2. QoS Application-Defined AVPs

This document reuses the AVPs defined in Section 4 of [RFC5777].

This section lists the AVPs that are introduced specifically for the QoS application. The following new AVPs are defined: Bound-Auth-Session-Id and the QoS-Authorization-Data AVP.

The following table describes the Diameter AVPs newly defined in this document for use with the QoS Application, their AVP code values, types, possible flag values, and to determine whether the AVP may be encrypted.

				AVP Flag rules		
Attribute Name	AVP Code	Section Defined	Data Type	MUST	SHLD NOT	MUST NOT
QoS-Authorization-Data	579	7.2	OctetString	M		V
Bound-Auth-Session-Id	580	7.2	UTF8String	M		V

M - Mandatory bit. An AVP with the "M" bit set and its value MUST be supported and recognized by a Diameter entity in order for the message, which carries this AVP, to be accepted.
V - Vendor-specific bit that indicates whether the AVP belongs to an address space.

QoS-Authorization-Data

The QoS-Authorization-Data AVP (AVP Code 579) is of type OctetString. It is a container that carries application-session or user-specific data that has to be supplied to the AE as input to the computation of the authorization decision.

Bound-Authentication-Session-Id

The Bound-Authentication-Session AVP (AVP Code 580) is of type UTF8String. It carries the ID of the Diameter authentication session that is used for the network access [RFC4005]. It is used to tie the QoS authorization request to a prior authentication of the end-host done by a co-located application for network access authentication ([RFC4005]) at the QoS NE.

8. Accounting

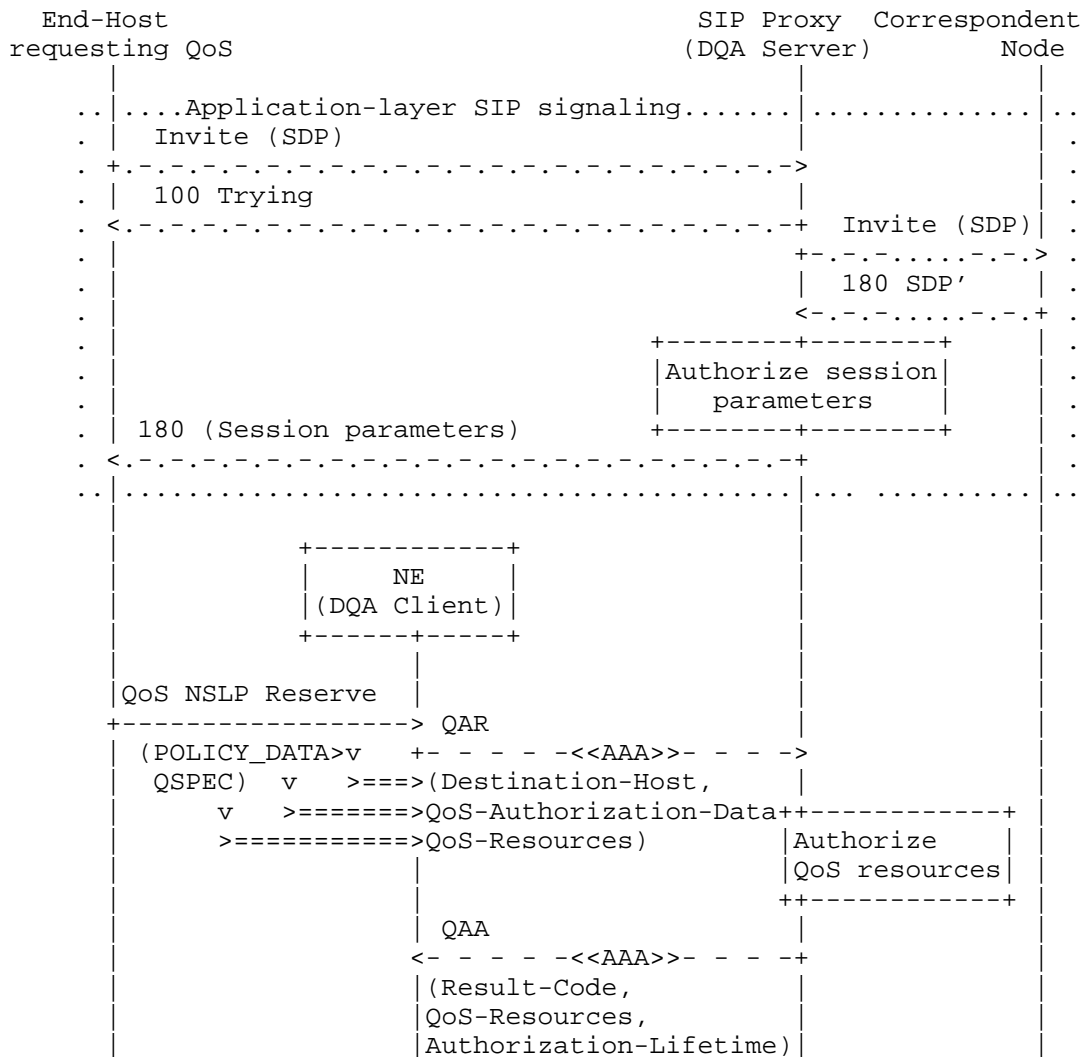
An NE MAY start an accounting session by sending an Accounting-Request (ACR) message after successful QoS reservation and activation of the data flow (see Figures 6 and 7). After every successful re-authorization procedure (see Figures 8 and 9), the NE MAY initiate an interim accounting message exchange. After successful session termination (see Figures 10 and 11), the NE may initiate a final exchange of accounting messages for the termination of the accounting session and report final records for the use of the QoS resources reserved. It should be noted that the two sessions (authorization and accounting) have independent management by the Diameter base protocol, which allows for finalizing the accounting session after the end of the authorization session.

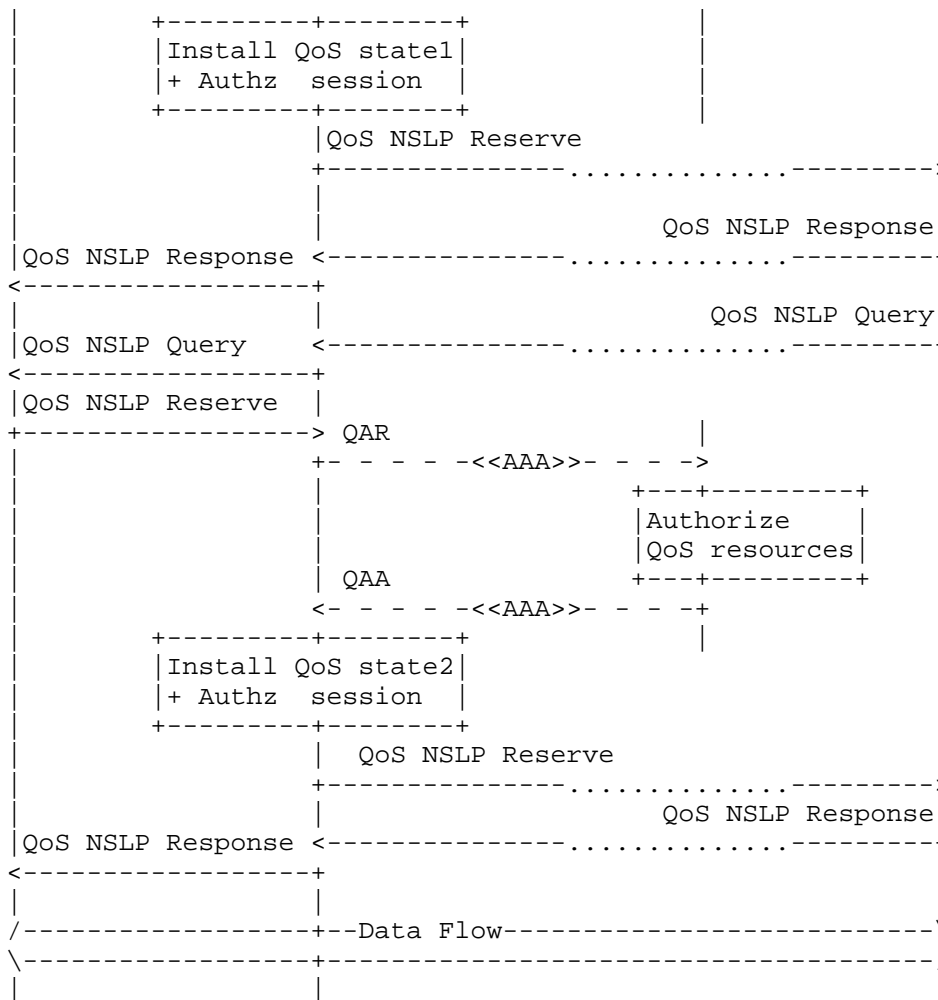
The detailed QoS accounting procedures are out of scope in this document.

9. Examples

9.1. Example Call Flow for Pull Mode (Success Case)

This section presents an example of the interaction between the end-host and Diameter QoS application entities using Pull mode. The application-layer signaling is, in this example, provided using SIP. Signaling for a QoS resource reservation is done using the QoS NSIS Signaling Layer Protocol (NSLP). The authorization of the QoS reservation request is done by the Diameter QoS application (DQA).





.-.-.-.-. SIP signaling
 ----- QoS NSLP signaling
 - - - - - Diameter QoS Application messages
 ===== Mapping of objects between QoS and AAA protocol

Figure 12: QoS Authorization Example - Pull Mode

The communication starts with SIP signaling between the two endpoints and the SIP proxy for negotiation and authorization of the requested service and its parameters (see Figure 12). As a part of the process, the SIP proxy verifies whether the user at Host A is authorized to use the requested service (and potentially the ability to be charged for the service usage). Negotiated session parameters are provided to the end-host.

Subsequently, Host A initiates a QoS signaling message towards Host B. It sends a QoS NSLP Reserve message, in which it includes description of the required QoS (QSPEC object) and authorization data for negotiated service session (part of the POLICY_DATA object). Authorization data includes, as a minimum, the identity of the AE (e.g., the SIP proxy) and an identifier of the application-service session for which QoS resources are requested.

A QoS NSLP reserve message is intercepted and processed by the first QoS-aware Network Element. The NE uses the Diameter QoS application to request authorization for the received QoS reservation request. The identity of the AE (in this case, the SIP server that is co-located with a Diameter server) is put into the Destination-Host AVP, any additional session authorization data is encapsulated into the QoS-Authorization-Data AVP, and the description of the QoS resources is included into the QoS-Resources AVP. These AVPs are included into a QoS Authorization Request message, which is sent to the AE.

A QAR message will be routed through the AAA network to the AE. The AE verifies the requested QoS against the QoS resources negotiated for the service session and replies with a QoS-Authorization-Answer (QAA) message. It carries the authorization result (Result-Code AVP) and the description of the authorized QoS parameters (QoS-Resources AVP), as well as duration of the authorization session (Authorization-Lifetime AVP).

The NE interacts with the Traffic Control function and installs the authorized QoS resources and forwards the QoS NSLP reserve message farther along the data path. Moreover, the NE may serve as a signaling proxy and process the QoS signaling (e.g., initiation or termination of QoS signaling) based on the QoS decision received from the Authorizing Entity.

9.2. Example Call Flow for Pull Mode (Failure Case)

This section repeats the scenario outlined in Section 9.1; however, in this case, we show a session authorization failure instead of success. Failures can occur in various steps throughout the protocol execution, and in this example, we assume that the Diameter QAR request processed by the Diameter server leads to an unsuccessful

result. The QAA message responds, in this example, with a permanent error "DIAMETER_AUTHORIZATION_REJECTED" (5003) set in the Result-Code AVP. When the NE receives this response, it discontinues the QoS reservation signaling downstream and provides an error message back to the end-host that initiated the QoS signaling request. The QoS NSLP response signaling message would in this case carry an INFO_SPEC object indicating the permanent failure as "Authorization failure" (0x02).

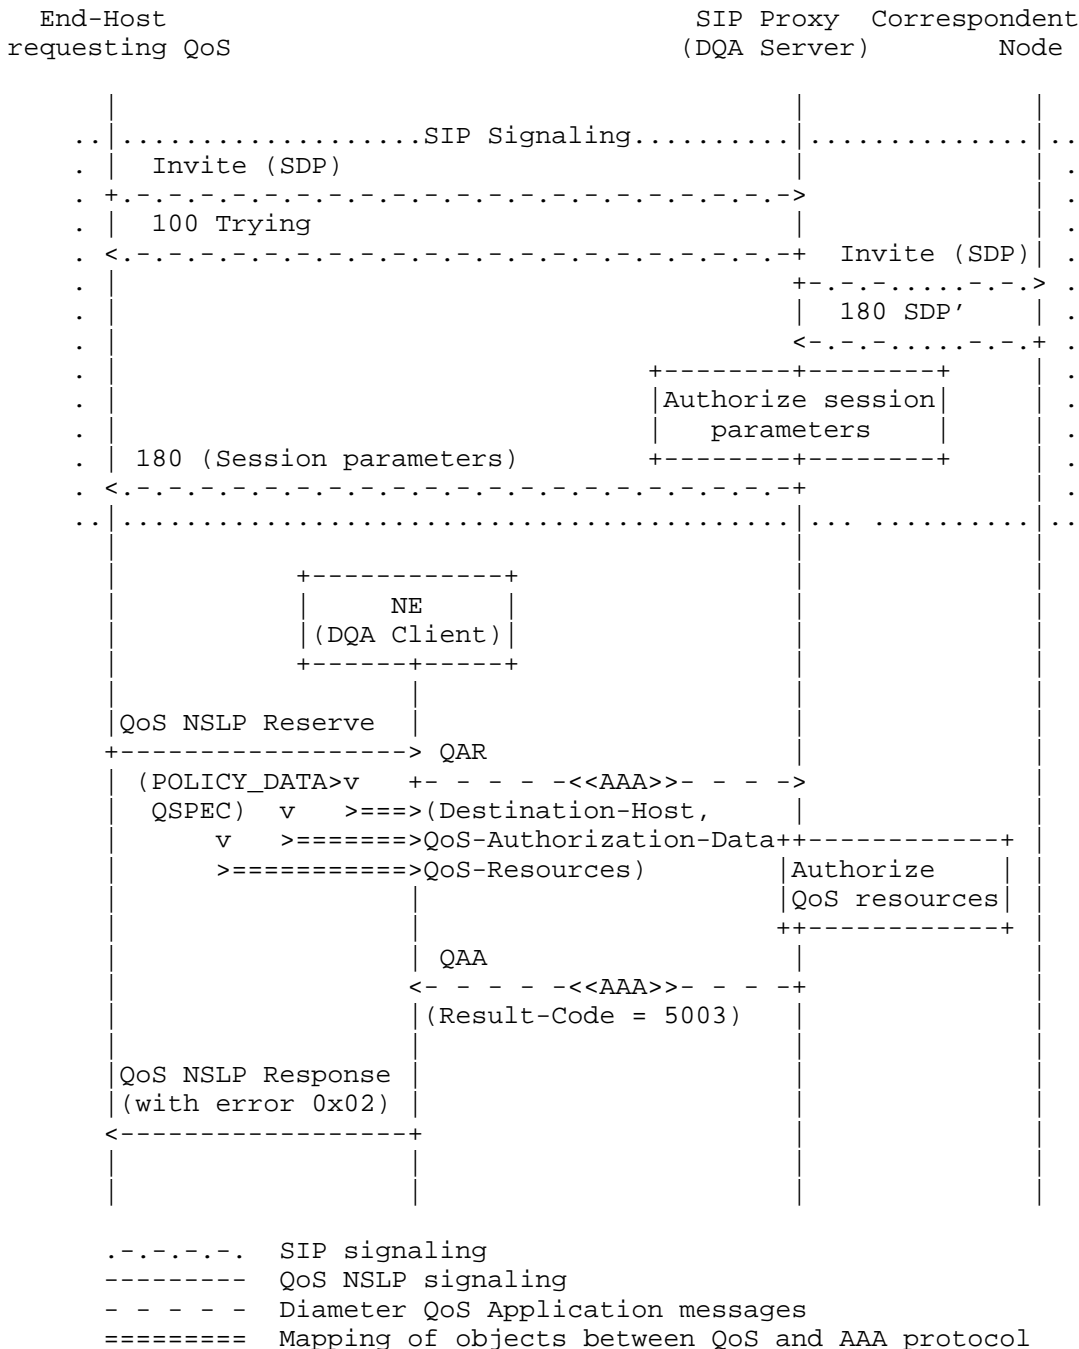
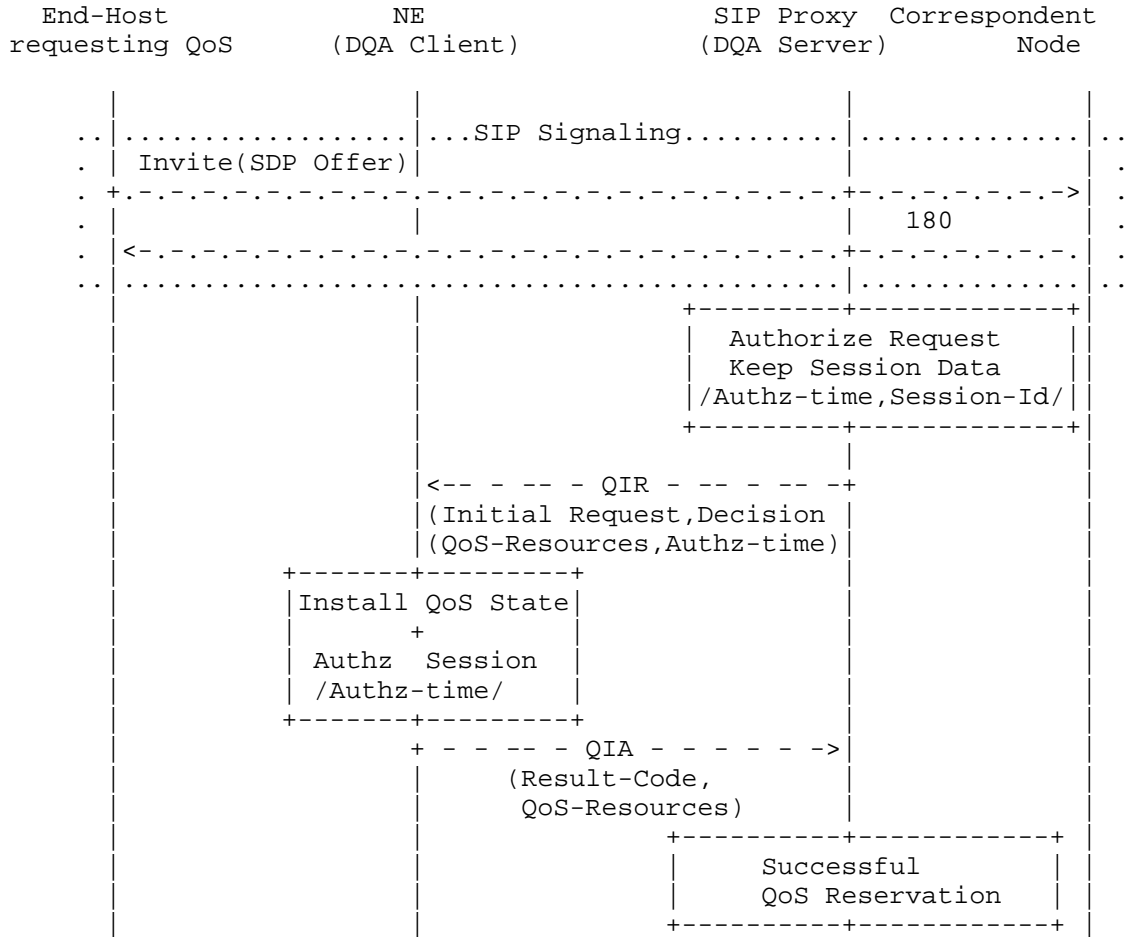


Figure 13: QoS Authorization Example - Pull Mode (Failure Case)

9.3. Example Call Flow for Push Mode

This section presents an example of the interaction between the end-host and Diameter QoS application entities using Push mode. The application-layer signaling is, in this example, provided using SIP. Signaling for a QoS resource reservation is done using the QoS NSLP. The authorization of the QoS reservation request is done by the Diameter QoS application (DQA).



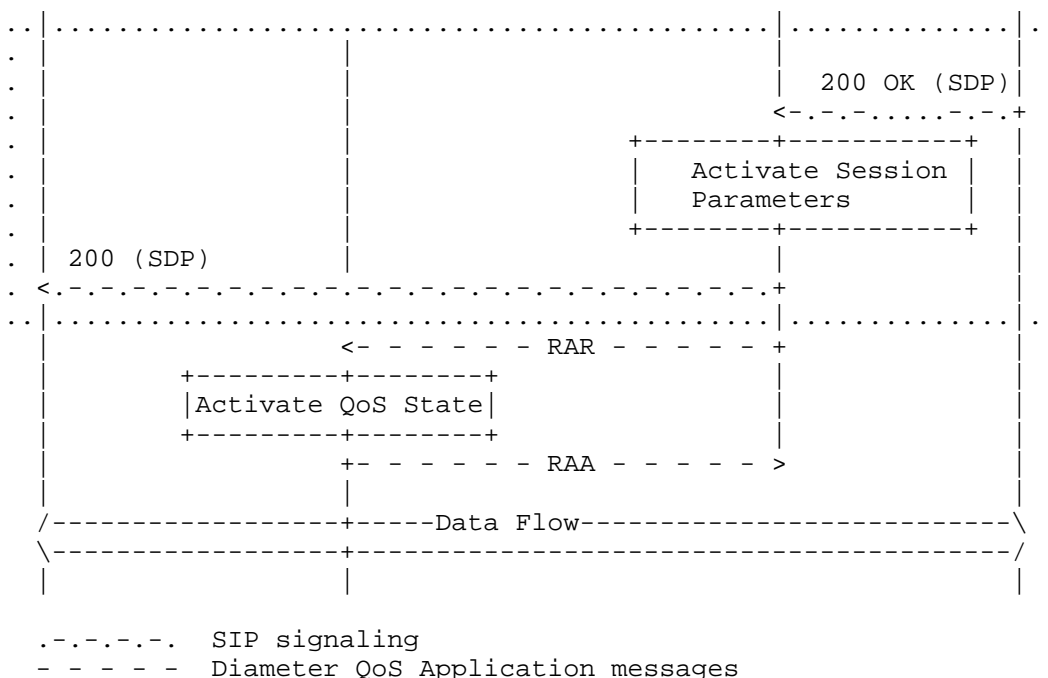


Figure 14: QoS Authorization Example - Push Mode

The communication starts with SIP signaling between the two endpoints and the SIP proxy for negotiation and authorization of the requested service and its parameters (see Figure 14). As a part of the process, the SIP proxy verifies whether the user at Host A is authorized to use the requested service (and potentially the ability to be charged for the service usage).

A few implementation choices exist regarding the decision about when to initiate the QoS reservation. [MMUSIC-MEDIA] discusses this aspect with a focus on firewalling. In the example above, the DQA server is triggered to authorize the QoS request based on session parameters from the Session Description Protocol (SDP) payload. It will use a QIR message to do so. For this example message flow, we assume a two-stage commit, i.e., the SIP proxy interacts with the NE twice. First, it only prepares the QoS reservation, and then, with the arrival of the 200 OK, the QoS reservation is activated.

This example does not describe how the DQA server learns which DQA client to contact. We assume pre-configuration in this example. In any case, the address of the DQA client is put into the Destination-Host AVP, the description of the QoS resources is included into the

QoS-Resources AVP, and the duration of the authorization session is carried in the Authorization-Lifetime AVP.

When the DQA client receives the QIR, it interacts with the Traffic Control function and reserves the authorized QoS resources accordingly. At this point in time, the QoS reservation is not yet activated.

When a 200 OK is returned, the DQA server may verify the accepted QoS against the pre-authorized QoS resources and send a Diameter RAR message to the DQA client in the NE for activating the installed policies and commit the resource allocation.

10. IANA Considerations

This section contains the namespaces that have either been created in this specification or had their values assigned to existing namespaces managed by IANA.

10.1. AVP Codes

IANA has allocated two AVP codes to the registry defined in [RFC3588]:

Registry:

AVP Code	AVP Name	Reference
579	QoS-Authorization-Data	Section 7.2
580	Bound-Auth-Session-Id	Section 7.2

10.2. Application IDs

IANA has allocated the following application ID from the registry defined in [RFC3588] (using the next available value from the 7-16777215 range).

Registry:

ID values	Name	Reference
9	Diameter QoS application	Section 5

10.3. Command Codes

IANA has allocated command code values from the registry defined in [RFC3588].

Registry:

Code Value	Name	Reference
326	QoS-Authorization-Request (QAR)	Section 5.1
326	QoS-Authorization-Answer (QAA)	Section 5.2
327	QoS-Install-Request (QIR)	Section 5.3
327	QoS-Install-Answer (QIA)	Section 5.4

11. Security Considerations

This document describes a mechanism for performing authorization of a QoS reservation at a third-party entity. The Authorizing Entity needs sufficient information to make such an authorization decision and this information may come from various sources, including the application-layer signaling, the Diameter protocol (with its security mechanisms), policy information stored available with a AAA server, and a QoS signaling protocol.

Below there is a discussion about considerations for the Diameter QoS interaction between an Authorizing Entity and a Network Element. Security between the Authorizing Entity and the Network Element has a number of components: authentication, authorization, integrity, and confidentiality.

Authentication refers to confirming the identity of an originator for all datagrams received from the originator. Lack of authentication of Diameter messages between the Authorizing Entity and the Network Element can seriously jeopardize the fundamental service rendered by the Network Element. A consequence of not authenticating the message sender by the Network Element would be that an attacker could spoof the identity of a "legitimate" Authorizing Entity in order to allocate resources, change resource assignments, or free resources. The adversary can also manipulate the state at the Network Element in such a way that it leads to a denial-of-service attack by, for example, setting the allowed bandwidth to zero or allocating the entire bandwidth available to a single flow.

A consequence of not authenticating the Network Element to an Authorizing Entity is that an attacker could impact the policy-based admission control procedure operated by the Authorizing Entity that provides a wrong view of the resources used in the network. Failing to provide the required credentials should be subject to logging.

Authorization refers to whether a particular Authorizing Entity is authorized to signal a Network Element with requests for one or more applications, adhering to a certain policy profile. Failing the authorization process might indicate a resource theft attempt or failure due to administrative and/or credential deficiencies. In either case, the Network Element should take the proper measures to log such attempts.

Integrity is required to ensure that a Diameter message has not been maliciously altered. The result of a lack of data integrity enforcement in an untrusted environment could be that an imposter will alter the messages exchanged between a Network Entity and an Authorizing Entity potentially causing a denial of service.

Confidentiality protection of Diameter messages ensures that the signaling data is accessible only to the authorized entities. When signaling messages from the Application Server (via the Authorizing Entity towards the Network Element) traverse untrusted networks, lack of confidentiality will allow eavesdropping and traffic analysis. Additionally, Diameter QoS messages may carry authorization tokens that require confidentiality protection.

Diameter offers security mechanisms to deal with the functionality demanded in the paragraphs above. In particular, Diameter offers communication security between neighboring Diameter peers using Transport Layer Security (TLS) or IPsec. Authorization capabilities are application specific and part of the overall implementation.

12. Acknowledgements

The authors would like to thank John Loughney and Allison Mankin for their input to this document. In September 2005, Robert Hancock, Jukka Manner, Cornelia Kappler, Xiaoming Fu, Georgios Karagiannis, and Elwyn Davies provided a detailed review. Robert also provided us with good feedback earlier in 2005. Jerry Ash provided us review comments in late 2005/early 2006. Rajith R provided some inputs to the document in early 2007.

We would also like to thanks Alexey Melnikov, Adrian Farrel, and Robert Sparks for their IESG reviews.

13. Contributors

The authors would like to thank Tseno Tsenov and Frank Alfano for starting the Diameter Quality of Service work within the IETF, for their significant contributions and for being the driving force for the first few draft versions.

14. References

14.1. Normative References

- [RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, March 1997.
- [RFC3588] Calhoun, P., Loughney, J., Guttman, E., Zorn, G., and J. Arkko, "Diameter Base Protocol", RFC 3588, September 2003.
- [RFC4005] Calhoun, P., Zorn, G., Spence, D., and D. Mitton, "Diameter Network Access Server Application", RFC 4005, August 2005.
- [RFC5624] Korhonen, J., Tschofenig, H., and E. Davies, "Quality of Service Parameters for Usage with Diameter", RFC 5624, August 2009.
- [RFC5777] Korhonen, J., Tschofenig, H., Arumaithurai, M., Jones, M., and A. Lior, "Traffic Classification and Quality of Service (QoS) Attributes for Diameter", RFC 5777, February 2010.

14.2. Informative References

- [MMUSIC-MEDIA] Stucker, B. and H. Tschofenig, "Analysis of Middlebox Interactions for Signaling Protocol Communication along the Media Path", Work in Progress, March 2009.
- [NSIS-NTLP] Schulzrinne, H. and M. Stiemerling, "GIST: General Internet Signalling Transport", Work in Progress, June 2009.
- [NSIS-QOS] Manner, J., Karagiannis, G., and A. McDonald, "NSLP for Quality-of-Service Signaling", Work in Progress, January 2010.
- [RFC2205] Braden, B., Zhang, L., Berson, S., Herzog, S., and S. Jamin, "Resource ReSerVation Protocol (RSVP) -- Version 1 Functional Specification", RFC 2205, September 1997.
- [RFC2211] Wroclawski, J., "Specification of the Controlled-Load Network Element Service", RFC 2211, September 1997.

- [RFC2212] Shenker, S., Partridge, C., and R. Guerin, "Specification of Guaranteed Quality of Service", RFC 2212, September 1997.
- [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.
- [RFC2753] Yavatkar, R., Pendarakis, D., and R. Guerin, "A Framework for Policy-based Admission Control", RFC 2753, January 2000.
- [RFC2865] Rigney, C., Willens, S., Rubens, A., and W. Simpson, "Remote Authentication Dial In User Service (RADIUS)", RFC 2865, June 2000.
- [RFC3261] Rosenberg, J., Schulzrinne, H., Camarillo, G., Johnston, A., Peterson, J., Sparks, R., Handley, M., and E. Schooler, "SIP: Session Initiation Protocol", RFC 3261, June 2002.
- [RFC3313] Marshall, W., "Private Session Initiation Protocol (SIP) Extensions for Media Authorization", RFC 3313, January 2003.
- [RFC3520] Hamer, L-N., Gage, B., Kosinski, B., and H. Shieh, "Session Authorization Policy Element", RFC 3520, April 2003.
- [RFC3521] Hamer, L-N., Gage, B., and H. Shieh, "Framework for Session Set-up with Media Authorization", RFC 3521, April 2003.
- [RFC4282] Aboba, B., Beadles, M., Arkko, J., and P. Eronen, "The Network Access Identifier", RFC 4282, December 2005.
- [RFC4566] Handley, M., Jacobson, V., and C. Perkins, "SDP: Session Description Protocol", RFC 4566, July 2006.
- [RFC5246] Dierks, T. and E. Rescorla, "The Transport Layer Security (TLS) Protocol Version 1.2", RFC 5246, August 2008.

Authors' Addresses

Dong Sun (editor)
Alcatel-Lucent
600 Mountain Ave
Murray Hill, NJ 07974
USA

Phone: +1 908 582 2617
EMail: d.sun@alcatel-lucent.com

Peter J. McCann
Motorola Labs
1301 E. Algonquin Rd
Schaumburg, IL 60196
USA

Phone: +1 847 576 3440
EMail: pete.mccann@motorola.com

Hannes Tschofenig
Nokia Siemens Networks
Linnoitustie 6
Espoo 02600
Finland

Phone: +358 (50) 4871445
EMail: Hannes.Tschofenig@gmx.net
URI: <http://www.tschofenig.priv.at>

Tina Tsou
Huawei
Shenzhen,
P.R.C

EMail: tena@huawei.com

Avri Doria
Lulea University of Technology
Arbetsvetenskap
Lulea, SE-97187
Sweden

EMail: avri@ltu.se

Glen Zorn (editor)
Network Zen
1310 East Thomas Street
#306
Seattle, Washington 98102
USA

Phone: +1 (206) 377-9035
EMail: gwz@net-zen.net