International Telecommunication Union

ITU-T

TELECOMMUNICATION
STANDARDIZATION SECTOR
OF ITU

Q.3309

(10/2009)

SERIES Q: SWITCHING AND SIGNALLING

Signalling requirements and protocols for the NGN – Resource control protocols

---

QoS coordination protocol

Recommendation ITU-T Q.3309

ITU-T

ITU-T Q-SERIES RECOMMENDATIONS

SWITCHING AND SIGNALLING

SIGNALLING IN THE INTERNATIONAL MANUAL SERVICE	Q.1–Q.3
INTERNATIONAL AUTOMATIC AND SEMI-AUTOMATIC WORKING	Q.4–Q.59
FUNCTIONS AND INFORMATION FLOWS FOR SERVICES IN THE ISDN	Q.60–Q.99
CLAUSES APPLICABLE TO ITU-T STANDARD SYSTEMS	Q.100–Q.119
SPECIFICATIONS OF SIGNALLING SYSTEMS No. 4, 5, 6, R1 AND R2	Q.120–Q.499
DIGITAL EXCHANGES	Q.500–Q.599
INTERWORKING OF SIGNALLING SYSTEMS	Q.600–Q.699
SPECIFICATIONS OF SIGNALLING SYSTEM No. 7	Q.700–Q.799
Q3 INTERFACE	Q.800–Q.849
DIGITAL SUBSCRIBER SIGNALLING SYSTEM No. 1	Q.850–Q.999
PUBLIC LAND MOBILE NETWORK	Q.1000–Q.1099
INTERWORKING WITH SATELLITE MOBILE SYSTEMS	Q.1100–Q.1199
INTELLIGENT NETWORK	Q.1200–Q.1699
SIGNALLING REQUIREMENTS AND PROTOCOLS FOR IMT-2000	Q.1700–Q.1799
SPECIFICATIONS OF SIGNALLING RELATED TO BEARER INDEPENDENT CALL CONTROL (BICC)	Q.1900–Q.1999
BROADBAND ISDN	Q.2000–Q.2999
SIGNALLING REQUIREMENTS AND PROTOCOLS FOR THE NGN	Q.3000–Q.3999
General	Q.3000–Q.3029
Network signalling and control functional architecture	Q.3030–Q.3099
Network data organization within the NGN	Q.3100–Q.3129
Bearer control signalling	Q.3130–Q.3179
Signalling and control requirements and protocols to support attachment in NGN environments	Q.3200–Q.3249
Resource control protocols	**Q.3300–Q.3369**
Service and session control protocols	Q.3400–Q.3499
Service and session control protocols – supplementary services	Q.3600–Q.3649
NGN applications	Q.3700–Q.3849
Testing for NGN networks	Q.3900–Q.3999

For further details, please refer to the list of ITU-T Recommendations.

Recommendation ITU-T Q.3309

QoS coordination protocol

Summary

Recommendation ITU-T Q.3309 defines an admission control coordination protocol for NGN and includes the definition of interfaces between the admission control coordination layer and higher-layer signalling systems, and between the admission control coordination layer and lower-layer transport networks.

Source

Recommendation ITU-T Q.3309 was approved on 29 October 2009 by ITU-T Study Group 11 (2009-2012) under Recommendation ITU-T A.8 procedures.

Keywords

QoS coordination, RSVP.

FOREWORD

The International Telecommunication Union (ITU) is the United Nations specialized agency in the field of telecommunications, information and communication technologies (ICTs). The ITU Telecommunication Standardization Sector (ITU-T) is a permanent organ of ITU. ITU-T is responsible for studying technical, operating and tariff questions and issuing Recommendations on them with a view to standardizing telecommunications on a worldwide basis.

The World Telecommunication Standardization Assembly (WTSA), which meets every four years, establishes the topics for study by the ITU-T study groups which, in turn, produce Recommendations on these topics.

The approval of ITU-T Recommendations is covered by the procedure laid down in WTSA Resolution 1.

In some areas of information technology which fall within ITU-T's purview, the necessary standards are prepared on a collaborative basis with ISO and IEC.

NOTE

In this Recommendation, the expression "Administration" is used for conciseness to indicate both a telecommunication administration and a recognized operating agency.

Compliance with this Recommendation is voluntary. However, the Recommendation may contain certain mandatory provisions (to ensure e.g., interoperability or applicability) and compliance with the Recommendation is achieved when all of these mandatory provisions are met. The words "shall" or some other obligatory language such as "must" and the negative equivalents are used to express requirements. The use of such words does not suggest that compliance with the Recommendation is required of any party.

INTELLECTUAL PROPERTY RIGHTS

ITU draws attention to the possibility that the practice or implementation of this Recommendation may involve the use of a claimed Intellectual Property Right. ITU takes no position concerning the evidence, validity or applicability of claimed Intellectual Property Rights, whether asserted by ITU members or others outside of the Recommendation development process.

As of the date of approval of this Recommendation, ITU had not received notice of intellectual property, protected by patents, which may be required to implement this Recommendation. However, implementers are cautioned that this may not represent the latest information and are therefore strongly urged to consult the TSB patent database at http://www.itu.int/ITU-T/ipr/.

CONTENTS

Recommendation ITU-T Q.3309

QoS coordination protocol

1 Scope

This Recommendation defines an admission control coordination protocol. Major design aspects of the protocol considered include the definition of interfaces between the admission control coordination layer and higher-layer signalling systems, and the admission control coordination layer and lower-layer transport networks. Protocol semantics are also included in the definition.

NOTE – This Recommendation is a specification of protocol requirements; it can functionally be part of different architectures.

2 References

The following ITU-T Recommendations and other references contain provisions which, through reference in this text, constitute provisions of this Recommendation. At the time of publication, the editions indicated were valid. All Recommendations and other references are subject to revision; users of this Recommendation are therefore encouraged to investigate the possibility of applying the most recent edition of the Recommendations and other references listed below. A list of the currently valid ITU-T Recommendations is regularly published. The reference to a document within this Recommendation does not give it, as a stand-alone document, the status of a Recommendation.

[IETF RFC 1633] IETF RFC 1633 (1994), *Integrated Services in the Internet Architecture: an Overview*.

[IETF RFC 2205] IETF RFC 2205 (1997), *Resource ReSerVation Protocol (RSVP0), Version 1 Functional Specification*.

3 Definitions

3.1 Terms defined elsewhere

This Recommendation uses the following terms defined elsewhere:

3.1.1 Adspec [IETF RFC 2205]: A Path message may carry a package of OPWA advertising information, known as an "Adspec".

NOTE – An Adspec received in a Path message is passed to the local traffic control, which returns an updated Adspec; the updated version is then forwarded in Path messages sent downstream.

3.1.2 flowspec [IETF RFC 2205]: Defines the QoS to be provided for a flow. The flowspec is used to set parameters in the packet scheduling function to provide the requested quality of service. A flowspec is carried in a FLOWSPEC object. The flowspec format is opaque to RSVP and is defined by the Integrated Services Working Group of the IETF.

3.1.3 Rspec [IETF RFC 2205]: The component of a flowspec that defines a desired QoS.

NOTE – The Rspec format is opaque to RSVP and is defined by the Integrated Services Working Group of the IETF.

3.1.4 Tspec [IETF RFC 2205]: A traffic parameter set that describes a flow.

NOTE – The format of a Tspec is opaque to RSVP and is defined by the Integrated Service Working Group of the IETF.

3.2 Terms defined in this Recommendation

This Recommendation defines the following terms:

3.2.1 forwarder: A node that is responsible – within a domain – to receive end-to-end QoS coordination requests, dispatch them to the admission control layer through the admission controller interface, process them and forward them to the next domain on the end-to-end path.

3.2.2 last forwarder: The forwarder of the last domain along the end-to-end path.

3.2.3 QoS requester: The node where the session control function or the end-user equipment requests QoS treatment to the network.

4 Abbreviations and acronyms

This Recommendation uses the following abbreviations and acronyms:

OPWA One-Pass With Advertising

QCP QoS Coordination Protocol

RSVP Resource ReSerVation Protocol

YESSIR YEt another Sender Session Internet Reservations

5 Conventions

None.

6 High-level description

In order to define the admission control coordination protocol, four layers are identified as illustrated in Figure 6-1: session control layer, coordination layer, admission control layer and QoS enforcement layer.

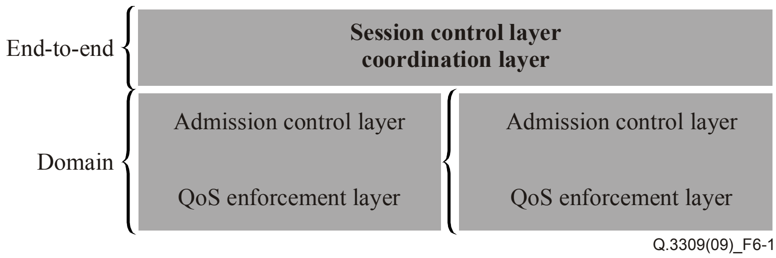

Figure 6-1 – Layer sketch of end-to-end QoS architecture

The QoS enforcement layer is instantiated at each admission controlled domain. It sends the necessary information to the upper admission control layer, so that when the admission control layer accepts to deliver a particular service to a flow or a user, the domain is able to provide it.

The admission control layer is instantiated at each admission controlled domain; it lies on top of the QoS enforcement layer and uses the information the QoS enforcement layer provides to act as a decision point for each admission controlled domain. Its task is to offer an interface to perform end-to-end requests for QoS treatment. The admission control layer interprets these requests for end-to-end QoS treatment, and supplies an answer on behalf of the whole domain whether or not the request can be accepted and the requested service can be delivered.

The coordination layer lies on top of the admission control layer. This common layer links together the admission controlled domains, giving end-to-end control to a set of local controllers.

The session control layer lies on top of the coordination layer and uses its services to enhance a session with the assurance of a certain type of service from the network. A session is made by a set of participants who engage themselves in a communication. Identifying these participants as well as transporting the information about the agreement on the QoS requirements among all the participants are two of the tasks of this layer. After this, QoS requirements are communicated to the coordination layer that triggers the mechanisms to create the end-to-end service.

6.1 Admission control coordination models

The basic building block of the network is the admission controlled domain, which is an interconnection of network elements that provides an admission controller interface. A subject (flow, packet, etc.) can request QoS treatments and receive a positive or negative response to their request. The admission control coordination protocol acts as a bridge between a single end-to-end request and the heterogeneity of the admission control mechanisms that are already deployed in the network; it provides all the means to locate, contact, query and coordinate the responses of the admission controllers on the path. The coordination protocol interacts with an actor called the coordination daemon.

There are many ways to design the coordination protocol. It can be:

a) Path-coupled vs. path-decoupled: A coordination protocol is path-coupled if it is part of the data path, path-decoupled otherwise; there are no assumptions on the location of the coordination daemon; if it is tightly coupled with its admission controller, it will be a matter of inter-process communication, whereas if they are loosely-coupled, an additional message will be sent across the network.

b) Stateless vs. stateful, with respect to the QoS request information: In a stateful scenario, the coordination daemon keeps a list of all the subject-treatment couples for the subjects that use part of the resources in that coordination daemon's scope; whereas in a stateless scenario, the coordination daemon does not keep a list of any subject.

The overall scheme can be proactive or reactive: this applies in the presence of failures, re-routes, mobility or other exceptional events. In these cases, subjects may have changed their path; hence the coordination phase may have to be performed again, as well as the admission control phase. If this is done proactively, the coordination mechanism is triggered periodically to refresh the state and to react to exceptional events. In a reactive scenario, the coordination daemon is aware of changes that cause exceptional events, and it reacts by re-triggering the coordination mechanism.

From such a characterization, we can sketch different possible architectures, summarized in Figure 6-2, where each oval represents an admission controlled domain; a spot represents an admission controller interface and a square represents a coordination daemon (note that they might overlap and be instantiated at the same node).

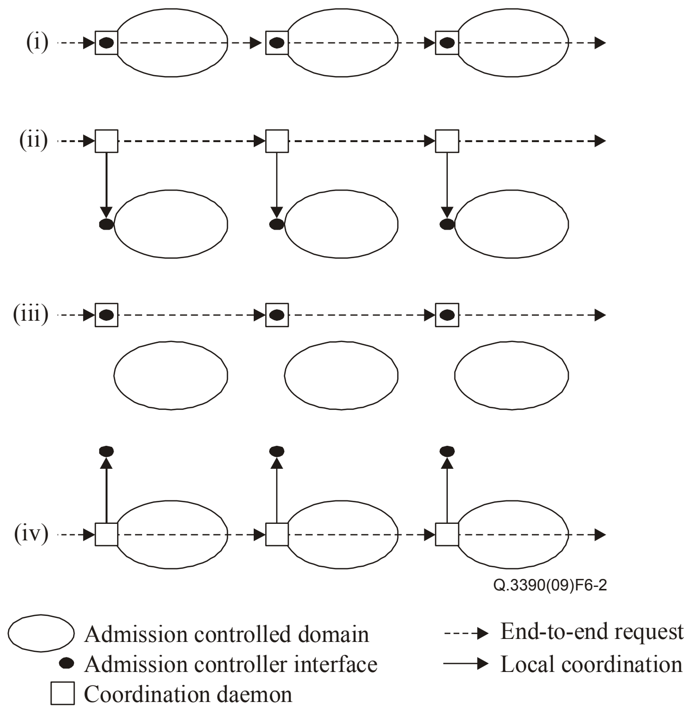

Figure 6-2 – Admission control coordination models

Among all the possible solutions, we might have (i) a scenario where the coordination daemon and the admission controller are both path-coupled: the coordination daemon receives an end-to-end request, in turn, it both propagates on to the next coordination daemon, and pushes down to the admission controllers; (ii) a situation in which the coordination daemon is path-decoupled and the admission controller is path-coupled; (iii) a scenario where the coordination daemon and the admission controller are both path-decoupled; and finally, (iv) a path-coupled coordination daemon and path-decoupled admission controller solution.

This model could be further extended to introduce the concept of hierarchy. The four models presented above could in turn be seen as basic building blocks. The most general approaches from Figure 6-2 are (ii) and (iii). (i) is the same as (ii) with the introduction of hierarchy; (iv) is the same as (ii), since no assumptions can be made on whether the admission controller and the coordination daemon are tightly or loosely coupled.

7 Protocol description

Clause 6 details the main models for arranging the admission control coordination. In this clause, a set of design guidelines for an admission control coordination protocol for the Internet is formulated. RSVP is extended in order to make it suitable as an admission control coordination protocol.

7.1 Design guidelines and basic protocol operation

An instantiation of such a coordination protocol consists of a protocol initiator, a terminator, and a sequence of coordination daemons. From initiator to terminator, the protocol follows the data path, communicating one by one with every coordination daemon along this path. Both a path-coupled approach and path-decoupled approach are allowed. The use of RSVP means that a separate routing mechanism is not necessary. Coordination daemons first trigger their local resource management mechanism (within their admission control domain) and then propagate the outcome of such a process until every coordination daemon along the path has provided its outcome. This communication occurs via the admission controller interface. The coordination protocol allows the various outcomes to be merged together into a single response to the QoS requester.

7.1.1 Interfaces

In general, the protocol daemon can have up to four interfaces: the application interface, the routing interface, the admission controller interface, and the coordination interface.

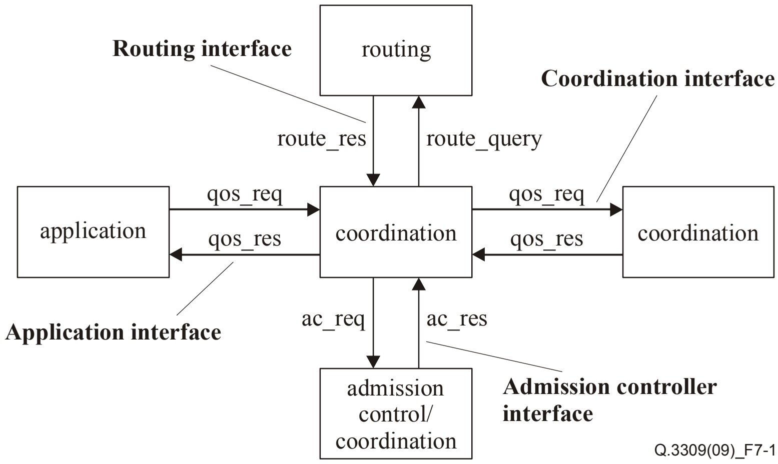

Figure 7-1 – Protocol interfaces

The admission controller interface is invoked by the coordination function and is implemented by either another coordination function (i.e., recursively) or by an admission control function. When it is implemented by an admission control function, this will allocate the necessary resources, if available. A successful allocation returns a positive admission control outcome, whereas a failure returns a negative outcome.

On the other hand, when the admission control function is implemented by a coordination function, the interface recursively looks for all the lower layer admission control functions, queries them and coordinates an outcome in one domain.

The routing interface uses RSVP to determine the path to forward on. It also receives route change notifications. The coordination interface is used to communicate to other coordination daemons or to the application or call agent at the end of the protocol chain.

7.1.2 Recursion

Since the invocation of the admission controller interface is transparent, the coordination protocol is unaware of the nature of the layer below. In particular, the layer below can be a protocol implementing either an admission control function or another coordination function. If an admission control function is implemented, then the recursion ceases; otherwise, if a further coordination function is implemented, then the current domain is further divided and coordinated within the divided domains. This process of recursively invoking the admission controller interface iterates until each domain configures its local elements' traffic control and reports the outcome of the admission control. Thus an admission controlled domain could be divided into sub-domains, where admission control is decided at the lower level. Optionally, the original semantics of RSVP could be maintained in a domain or sub-domain so that an RSVP reservation is made on routers, which is out of scope of this Recommendation.

7.1.3 Protocol semantics

Provided a coordination protocol shows the properties discussed in the previous clauses, the semantics of the message exchange are similar to those of resource reservation protocols such as RSVP [IETF RFC 1633] and YESSIR. Two models are supported:

- A two-stage commit-reserve: A coordination daemon receives a QoS request which is passed to the resource management system with the aim to reserve resources without committing them.
- A one-stage commit: A coordination daemon receives a QoS request which is passed to the resource management system with the aim of committing resources.

The difference between both cases is when resources are committed, i.e., when the admission control decision has been made. With respect to this, the protocol can operate in two ways, depending on whether a QoS request is in (non-)blocking mode for a coordination daemon.

- Blocking mode: A coordination daemon passes a commit request to the resource management system, waits for a response, and then forwards on both the commit request and the response.
- Non-blocking mode: A coordination daemon passes a commit request to the resource management system, forwards the commit request on, and waits for an upstream response to come back.

7.2 RSVP extensions to support coordination mode

RSVP [IETF RFC 2205] provides QoS signalling for application data streams. QoS requesters can use RSVP to request a specific QoS from the network for particular application flows. The admission control domain uses RSVP to deliver QoS requests to all nodes along the data path. RSVP can also maintain and refresh states for a requested QoS application flow.

The design of RSVP is distinguished in a number of fundamental ways, including soft state management, two-stage reserve-commit message exchanges, and separation of signalling from routing.

RSVP carries QoS signalling messages through the network, visiting each node along the data path, while following normal IP routing. To make a resource reservation at a node, the RSVP daemon communicates with two local decision modules: admission control and policy control. Admission control determines whether the node has sufficient available resources to provide the requested QoS. Policy control provides authorization for the QoS request. If either check fails, the RSVP module returns an error notification to the application process that originated the request. If both checks succeed, the RSVP module sets parameters in a packet classifier and packet scheduler to obtain the desired QoS.

RSVP is suitable to support admission control coordination because it shows most of the properties the QoS coordination protocol needs:

1) RSVP can be path-coupled, and
2) can be used recursively.

Nevertheless, there are a number of substantial features that RSVP is not designed for, including:

– Domain-awareness – RSVP is not domain aware.
– Admission controller interface – RSVP daemons support an admission controller interface. Although, while in coordination mode, this interface would likely be different and more generic.
– Protocol semantics – RSVP message exchange is based on a 2-stage reserve-commit model. A 1-stage exchange model is needed to support a coordination mode.

– Mode field – RSVP needs to support a message header field in order to distinguish when it is used in reservation mode from when it is used in coordination mode. An admission controlled domain could use an extended RSVP by handling coordination mode in the same way as reservation mode. This would be an internal choice by the domain.
– Recursion field – RSVP could be used recursively; however, it does not have an explicit message header field explicitly stating the level of recursion. This would be required to handle recursive protocol instantiations.
– QoS pre-booking – This is when QoS is requested for a future time period rather than for immediate use. RSVP does not support pre-booking of QoS.
– Sender initiation – RSVP is receiver initiated. Sender initiation is needed as is the ability of initiation by proxies of the sender or receiver.

The extended version of RSVP combines both resource reservation and coordination functions. When used in reservation mode, RSVP operates in the traditional way. When used in coordination mode, RSVP operates as follows. The message exchange involves one node per admission control domain, i.e., the data traverses each ingress node. Upon receipt of a request, the ingress node initiates its local resource management mechanism, and forwards the request to the next ingress node along the data path. This process iterates until the request reaches the terminator. The terminator replies with a response that traverses the same ingress nodes, collecting and merging the outcomes of the various admission control processes (success/failure). Once at the initiator, the response carries the overall, initiator-to-terminator, admission control response

Ingress nodes represent the default point of contact for that domain by running an instance of the admission control coordination protocol. When it is not necessary for the admission controller to represent the default point of contact for that domain, the coordination daemon forwards (via the admission controller interface) the request to the appropriate, locally configured, entity.

7.3 QoS reservation types

There are different possible semantics of reservation: One-pass mechanism (commit-error), Two-pass mechanism (reserve-commit), and One-pass with advertising (OPWA). For the reservation type OPWA, the advertising messages from the QoS requester are delivered to each admission controller interface in the domains along the end-to-end path. When they reach the receiver, they contain all the Rspec metrics resulting in the reservation of the given subject (defined by its Tspec) under the different types of services that the network can deliver. The decision can then be taken, choosing among all the different services and knowing all the Rspec metrics.

OPWA is used in this Recommendation since this is the method that allows complete flexibility. The QoS coordination protocol is recommended to contain a third message flow to deal with the situation where, at the endpoint, advertising of network services is available (most likely the receiver) and the endpoint is not the point where the decision of which service to choose is made. In such a case, the entity that takes the decision is recommended to be contacted, given the information about the different service and asked which service to choose.

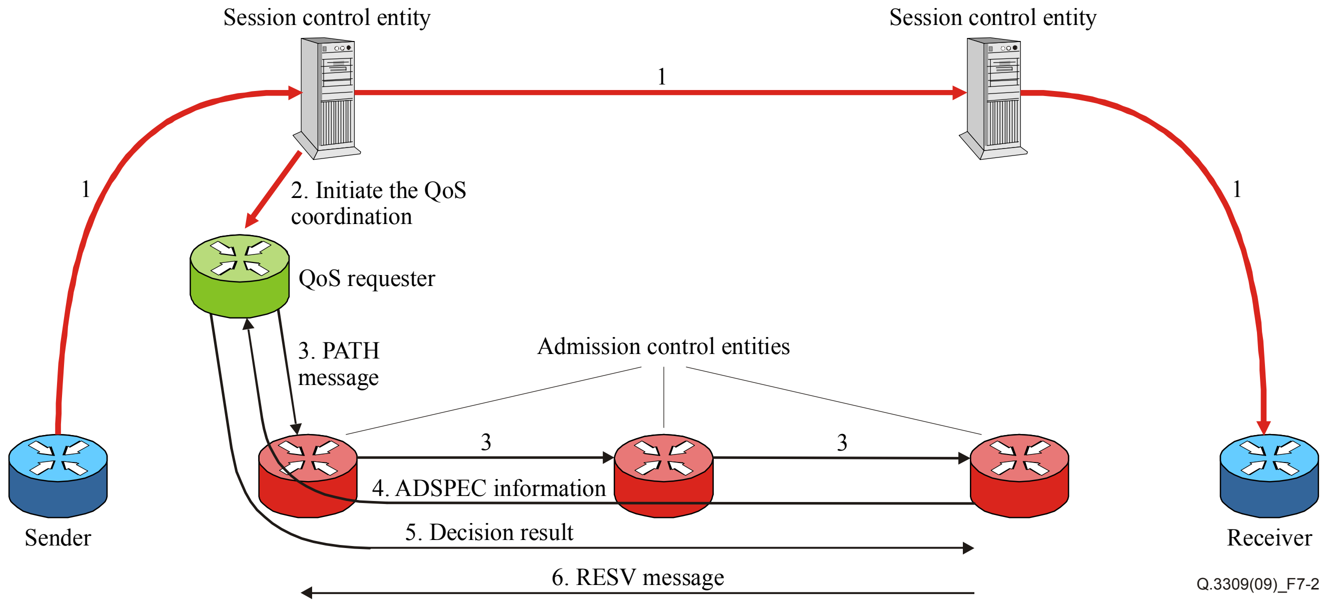

Figure 7-2 – Scenario of One-pass with advertising (OPWA)

A sketch of the above scenario is shown in Figure 7-2:

1) The sender contacts the session signalling overlay that differentiates the participants and negotiates the type of media.

2) The session control entity initiates the QoS coordination mechanisms via the application interface (as described in Figure 7-1), contacting the QoS requester.

3) If RSVP is assumed as the signalling protocol, the QoS requester sends a PATH message downstream (passing through all the admission control entities via the admission controller interfaces), with the Tspec of the flow. The PATH message collects advertising data in the ADSPEC object; the complete Adspec is present at the last forwarder (i.e., admission controller interface).

4) The Adspec information is then sent back to the QoS requester (assuming that the QoS requester is the suitable decision maker).

5) The QoS requester makes the decision and sends it to the last forwarder.

6) The last forwarder can then send back the RESV message to reserve the resource for the chosen service.

7.4 Aggregation

The model of coordination presented in clause 6 is explicitly per-flow and seeks end-to-end agreement on the treatment. This results in the scheduling of a single message exchange for each request.

Aggregation is a way to reduce the number of messages, the amount of information stored and the processing time for the coordination layer, when such tasks become excessive in the inner domains of the network. Generally, aggregation presumes an effort by the edges (i.e., where per-flow handling of QoS is possible) to aggregate flows so that they share a common path in the core and similar QoS requirements. This causes them to be treated together, reducing, therefore, the cost of individual treatment, without reducing the efficiency to obtain services.

In general, two different solutions can be identified to determine the boundaries of an aggregation area: static aggregation and dynamic aggregation.

Static aggregation is where an aggregation area is identified so that, within its boundaries, flows are treated as aggregates. In this solution, the scope of the domain is decided by network administrators so as to get the best out of the aggregation process. This may be complicated by the presence of multiple domains.

Dynamic aggregation is used dynamically, depending on different factors such as the direction of the flow, the number of flows and the position in the network of the domains handling a particular flow. Dynamic aggregation is out of scope of this Recommendation since it is complex to realize in practice.

7.4.1 Marking method

In the case when MPLS is used as a tunnelling technique, different label switched paths can be used to differentiate the QoS treatment among flows that follow the same aggregated path. A second encapsulated header is used to take care of the multiplexing of different QoS treatments among an aggregate. For example, compliant with the FILTER_SPEC classification of packets, a UDP header could be used for the aggregated traffic; the UDP source port field could be used to multiplex different packet treatments, the UDP destination port could be used for other purposes. This can be seen in Figure 7-3.

Encapsulated IP header	Encapsulated UDP header	IP header	Transport layer packet

Q.3309(09)_F7-3

Figure 7-3 – Packets for aggregated traffic

Appendix I

(This appendix does not form an integral part of this Recommendation)

I.1 List of different QoS reservation types

Table I.1 gives informational explanation of different QoS reservation types to help understand the texts in clause 7.3.

Table I.1 – QoS reservation types

QoS reservation types	Definitions
One-pass mechanism (commit-error)	Reservation requests from the sender are delivered to each admission controller interface in the domains along the end-to-end path. The reservation is either admitted, in which case the process goes on, or it is not, in which case a cascade of errors is propagated upstream to tear down the already admitted reservations. In this scenario though, there is no reasonable way, in which at the end of the reservation path, the coordination layer can know metrics such as end-to-end delay or jitter bound, for instance; hence, with this mechanism, the coordination layer can just instruct all the domains about the desired per-domain service.
Two-pass mechanism (reserve-commit)	Reservation requests from the sender are delivered to each admission controller interface in the domains along the end-to-end path in two passes. With the first pass, the sender inserts the Rspec, the network reserves the service that fits the request with the tightest metrics. At the receiver end, the global reservation is received; if the service characteristics that the network can offer are less than the requested ones, the reservation is rejected and torn down; in case they are even tighter, the service characteristics can be relaxed containing information about the excess resources. This scenario allows the specification of all the Rspec metrics.
One-pass with advertising (OPWA)	See clause 7.3.

I.2 Dynamic aggregation

In the dynamic aggregation, aggregation is used dynamically, depending on different factors such as the direction of the flow, the number of flows and the position in the network of the domains handling a particular flow. Dynamic aggregation area creation is not a straightforward task, and can be carried out as a result of coordination between local decisions of the admission controlled domains. Each domain in the path decides whether to use aggregation based on local policies. Consequently, the local decisions are shared with the help of the coordination layer and aggregation areas are dynamically built. This does not pose a problem in the intra-domain situation, but may pose a problem in the inter-domain situation between neighbouring nodes where less information may be available. In addition, this approach may lead to a bigger number of messages exchanged and to a more expensive computation for the communication of the local decisions and the coordination of different domains.

Bibliography

[b-ITU-T H.360] Recommendation ITU-T H.360 (2004), *An architecture for end-to-end QoS control and signalling*.

[b-ETSI RACS] ETSI RACS Release 1 (2005), *NGN functional architecture; Resource and Admission Control Subsystem (RACS)*.

[b-IETF RFC 2638] IETF RFC 2638 (1999), *A Two-bit Differentiated Services Architecture for the Internet*.

[b-IETF RFC 2702] IETF RFC 2702 (1999), *Requirements for Traffic Engineering Over MPLS*.

[b-IETF RFC 2814] IETF RFC 2814 (2000), *A Protocol for RSVP-based Admission Control over IEEE 802-style networks*.

[b-IETF RFC 3312] IETF RFC 3312 (2002), *Integration of Resource Management and Session Initiation Protocol (SIP)*.

[b-IETF WG Charter] IETF WG Charter (1996), *Integrated Services over Specific Link Layers, IETF Working Group charter*, <http://www.ietf.org/proceedings/37/charters/issll-charter.html>.

[b-IETF WG Charter 1] IETF WG Charter (2009), *Next Steps In Signalling, IETF Working Group charter*, <http://www.ietf.org/html.charters/nsis-charter.html>.

[b-3GPP TR23.802] 3GPP TR23.802 (2007), *Architectural enhancements for end-to-end Quality of Service*.

[b-ACM Tussle] ACM SIGCOMM Tussle in Cyberspace (2002), *Tussle in Cyberspace: Defining Tomorrow's Internet*.

[b-ATM forum TM V4.0] The ATM Forum Technical Committee (1996), *Traffic Management Specification Version 4.0*.

[b-BTTJSE Guaranteed QoS synthesis] Hovell, P., Briscoe, R., Corliano, G., (2005), *Guaranteed QoS synthesis – An example of a scalable core IP quality of service solution,* British Telecommunications Technical Journal Special Edition on IP Quality of Service.

[b-IFIP TC6 EuQoS] IFIP TC6 Conference EuQoS (2005), *End to End Quality of Service over Heterogeneous Networks (EuQoS), In Proc. Network Control and Engineering for QoS, Security and Mobility, IFIP TC6 Conference, NetCon'05*, Lannion, France.

[b-MSF-TR-ARCH-005-FINAL] Gallon, C., Schelén, O., (2005), *Bandwidth management in next generation packet networks*, *MSF Technical Report ARCH-005-FINAL*.

[b-PacketCable 050812] pkt-sp-dqos-I12-050812 (1999), *PacketCable Dynamic Quality of Service Specification.*

[b-QBBB] QBBB, Internet2 Bandwidth Broker Working Group, QBone Bandwidth Broker Architecture. <http://qbone.internet2.edu/bb/>.

[b-YESSIR] YESSIR (1999), YESSIR: *A Simple Reservation Mechanism for the Internet,* Computer Communication Review, vol. 29.

SERIES OF ITU-T RECOMMENDATIONS

Series A	Organization of the work of ITU-T
Series D	General tariff principles
Series E	Overall network operation, telephone service, service operation and human factors
Series F	Non-telephone telecommunication services
Series G	Transmission systems and media, digital systems and networks
Series H	Audiovisual and multimedia systems
Series I	Integrated services digital network
Series J	Cable networks and transmission of television, sound programme and other multimedia signals
Series K	Protection against interference
Series L	Construction, installation and protection of cables and other elements of outside plant
Series M	Telecommunication management, including TMN and network maintenance
Series N	Maintenance: international sound programme and television transmission circuits
Series O	Specifications of measuring equipment
Series P	Terminals and subjective and objective assessment methods
Series Q	**Switching and signalling**
Series R	Telegraph transmission
Series S	Telegraph services terminal equipment
Series T	Terminals for telematic services
Series U	Telegraph switching
Series V	Data communication over the telephone network
Series X	Data networks, open system communications and security
Series Y	Global information infrastructure, Internet protocol aspects and next-generation networks
Series Z	Languages and general software aspects for telecommunication systems

Printed in Switzerland
Geneva, 2010