

PROQOS—Dynamic SLA Management in DiffServ Space Links

Manuel Pessoa,¹ António Alves,¹ Gonçalo Quadros,¹ Fernando Boavida,² Michael Henke,³ Milva Natcheva,³ Patrick Halke,³ Peter Maurutschek,⁴ Zenon Huskic,⁴ Kurt Wagner,⁴ Frank Zeppenfeldt,⁵ and Roberto Donadio⁵

Several base elements for the provision of quality of service guarantees have been developed in the recent past. Of these, the Differentiated Services (DiffServ) architecture stands out as the most promising. In spite of this, various issues remain, especially when multidomain DiffServ services are concerned. In this case, some forms of distributed management of Service Level Agreements that allow the specification, exchange, enforcement and monitoring of quality of service data must be in place. Although, again, some isolated solutions exist for each of these problems, considerable effort is necessary to make them work together. The project presented in this paper tried to assess the feasibility of providing differentiated quality of service in satellite IP networks, by developing a dynamic Service Level Agreement management solution for an IP over Digital Video Broadcast Satellite system. The functionality of the implemented system comprises system configuration, dynamic SLA negotiation, QoS monitoring and metering, SLA conformance checking, and QoS reporting to customers.

KEY WORDS: Policy-based management; service level agreements trading; quality of service; DVB-S.

1. INTRODUCTION

The Internet Standard “Differentiated Services” [1, 2] has the potential to support scalable and flexible provisioning of quality of service (QoS) in terrestrial networks. Nevertheless, its use in satellite network environments is still incipient and places a set of technological challenges.

¹Critical Software SA, Coimbra, Portugal.

²University of Coimbra, Coimbra, Portugal.

³VCS Aktiengesellschaft, Bochum, Germany.

⁴Salzburg University, Salzburg, Austria.

⁵European Space Agency, Noordwijk, The Netherlands.

⁶To whom correspondence should be addressed at University of Coimbra, Coimbra, Portugal. E-mail: boavida@dei.uc.pt

“Quality of service” is used in different branches and domains with sometimes completely different interpretations. In the scope of packet-switched networks, QoS refers to the following parameters:

- Throughput, as the sum of all transmitted packet sizes within a given time interval;
- Packet loss rate, as the ratio between all lost packets and all injected packets;
- Latency, as the time difference between entry and departure times of packets into/out of a network;
- Jitter, or interpacket delay variance, as the absolute difference of successive packet latencies of a single packet stream.

In a wider sense, the availability of the transmission system can also be used as a QoS parameter, characterized as the following ratio: $MTBF/(MTBF+MTTR)$, where MTBF is the Mean Time Between Failures and MTTR the Mean Time To Repair.

This paper presents an overview the PROQOSL project, which was an extension to the work contracted by the European Space Agency (ESA) in the scope of the project “PROQOS—Development of Protocols for Quality of Service”—ESA ITT A0/1-3765/00/NL/US. The work developed in this context addressed the following general areas: QoS measuring, evaluation of Service Level Agreement (SLA) conformance, and automatic SLA negotiation and configuration.

One of the main objectives of the project was the implementation, assessment and demonstration of the applicability of DiffServ QoS mechanisms to IP/DVB-S (Internet Protocol/Digital Video Broadcast Satellite) environments. A second major objective was the development of mechanisms for a flexible and dynamic negotiation and management of SLAs between Space Link Providers (SLP) and their customers (namely, Internet Service Providers, ISP).

There is a specific business model in place between SLPs and ISPs which is based on an essentially static relation. Thus, allowing automatic and dynamic negotiation of SLAs between them will result in strong improvement of the effectiveness and efficiency of satellite capacity use.

The assumed baseline scenario is shown in Fig. 1. In order to transmit IP traffic over DVB-S, a content-providing organization has, in most cases, to use the services of a terrestrial network service provider (NSP) and of a satellite/teleport operator (space-link provider) for the transmission to the uplink hub.

If end-to-end data transfers require certain quality of service guarantees, then SLAs specifying these guarantees are required between content providers and network service providers on one side, and between network service providers and the space link provider on the other, thus requiring a distributed SLA management approach.

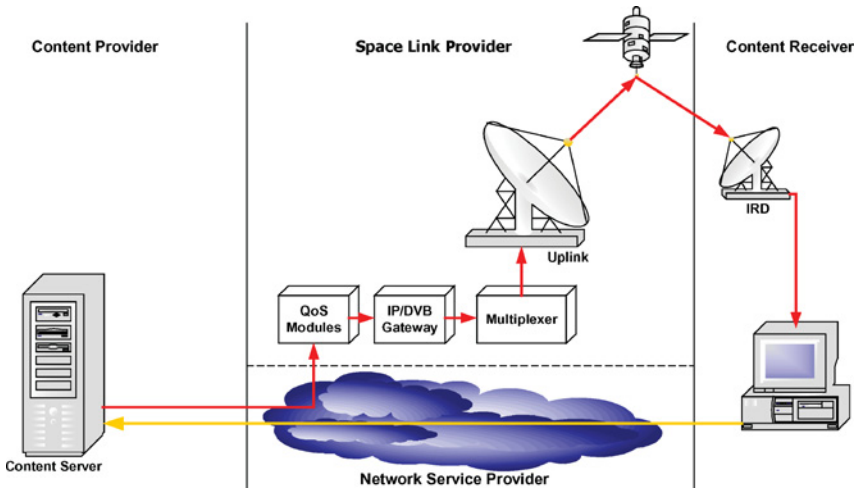


Fig. 1. PROQOS baseline scenario.

Currently, service level agreements are of a highly static, centralized nature. Typically, for IP/DVB-S environments, the contractual duration of SLAs nowadays is 1–3 years and they usually specify bandwidth and availability of the transmission system only. In some cases statements with respect to average latencies and packet loss rates are also included.

Within the PROQOS project, a more dynamic approach to SLAs was implemented. The base idea is that in “paper-SLAs” only global parameters, e.g. availability, are specified. Furthermore, these “paper-SLAs” define various service classes with corresponding QoS parameters. The customer identifies per service class minimum and maximum bandwidth values. During the contract period the customer may dynamically request bandwidth for a specific class and receive a corresponding proposal from the link provider, which he may accept or reject. This communication is performed automatically via specific machines (SLA Traders), as shown in the Fig. 2.

The advantage on the customer side is that only the actual required bandwidth has to be requested and paid for. For the link provider, this opens the possibility to offer available bandwidth to other customers in a more flexible way. In a scenario with multiple customers, there is even the possibility to offer more bandwidth than actually available in the uplink (overbooking) because, in general, customers do not use all of the requested bandwidth. If, however, in overbooking situations all customers use the requested bandwidth, the provider will probably not be able to fulfill his QoS commitments to all customers for all traffic classes. In this case a charge-back mechanism comes into action, in order to reimburse the concerned customers.

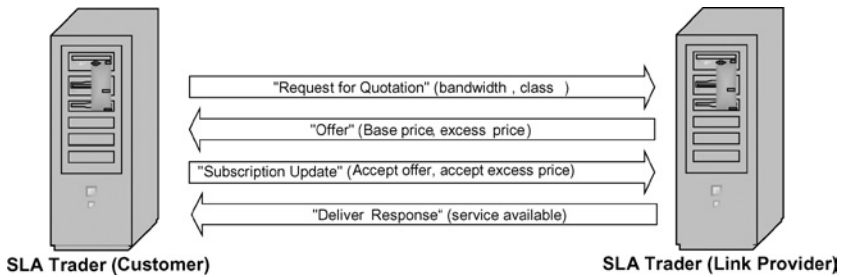


Fig. 2. Communication between customer and provider trader nodes.

By means of a dynamic SLA negotiation, it is possible to optimize the long-term usage of the available transponder capacity. For this to be effective, though, there is the need for QoS monitoring, evaluation of SLA conformance, and automatic SLA configuration, qualities which are at the base of the PROQOS platform presented in the following sections of this paper. Section 2 briefly identifies the related work. In Section 3 the overall functionality and architecture of PROQOS are presented. Section 4 describes the main issues pertaining to PROQOS SLAs, whereas Section 5 details the core element of PROQOS: the trader node. Some of the tests carried out on the developed prototype are presented in Section 6. Section 7 summarizes the contributions and identifies guidelines for further work.

2. RELATED WORK

This section presents the most relevant research work, related to the PROQOS project, being carried out in the scope of different research projects.

Under the umbrella of the Internet2 [3] initiative, a number of projects have been funded for advancing network, middleware and applications for the next generation Internet [4].

As part of the Internet2 initiative, the QBONE working group established an interdomain testbed for QoS-enhanced IP services, in order to provide the higher-education community with end-to-end services in support of emerging advanced networked applications. The technical approach focuses on supporting multidomain DiffServ services. QBONE addresses policy exchange and QoS negotiations, by means of dynamic SLAs.

Each QBONE network peers with neighboring QBONE DiffServ domains via specific policy servers called bandwidth brokers. Bilateral service level specifications (SLSs) exist between adjacent QBONE DiffServ domains [5]. These SLSs specify how traffic is classified, policed, and forwarded by Differentiated Services (DS) boundary nodes.

In the scope of the QBONE project there is a group specifically working on Bandwidth Broker (BB) issues (BBWG) [6]. In this project's context, a BB refers to an abstraction that automates the admission control decisions for QoS requests in a network domain. This means that it is responsible for keeping track of the current allocation of reserved traffic, it is configured with policies that define which traffic flows belong to which traffic classes, and it interprets new requests in light of these policies and the current bandwidth usage.

A group of European IST projects [7] of action line *IV.2.3: Network integration, interoperability, and interworking*, address QoS issues in IP networks. These projects, of which the TEQUILA project and the CADENUS project will be addressed below, started around the year 2000, and have a typical duration of 3 years.

The objective of the TEQUILA project [8] is to study, specify, implement and validate a set of service definitions and traffic engineering tools to obtain quantitative end-to-end Quality of Service (QoS) guarantees through careful planning, dimensioning and dynamic control of scaleable and simple qualitative traffic management techniques within DiffServ based networks. It addresses the specification of static and dynamic, intra- and interdomain SLSs (Service Level Specification), protocols and mechanisms for negotiating, monitoring and enforcing SLSs, and intra- and interdomain traffic engineering schemes to ensure that the network can cope with the contracted SLSs—within domains, and in the Internet at large.

One of the major contributions of this project to this research field is the specification of the basic information to be handled by SLSs when considering the deployment of value-added IP service offerings over the Internet. This specification was published in February 2002 as an Internet Engineering Task Force (IETF) Internet Draft [9] and proposes a standard formalism for listing the set of basic parameters which will actually compose an SLS. The document explicitly states that the requirements for SLS negotiation between service providers and between service providers and customers is subject to future study.

A monitoring and measurement architecture is also being investigated within the context of the TEQUILA project. The technological issues of the architecture have not been addressed yet, but the Simple Network Management Protocol (SNMP) is given as an example protocol for monitoring data passing from network devices to node monitors.

Another relevant IST project is the CADENUS project [10]. The primary goal of CADENUS is to develop, implement, validate and demonstrate a framework for the configuration and provisioning of end-user services with QoS guarantees in Premium IP networks.

The CADENUS framework proposes the creation and delivery of end-user services in premium IP with QoS guarantees by means of an SLA management

framework composed of three functional blocks at the user–provider interface, within the service provider domain and between the service provider and the network provider: the Access Mediator (AM); the Service Mediator (SM); and the Resource Mediator (RM). These functional blocks are responsible for the management of user access to the service, for presenting the available services and for interacting and managing the QoS-aware network elements available in the underlying network infrastructure [11]. This project adopted the SLS specification model defined within the TEQUILA project.

On the technology level, the ebXML architecture [12] is being considered because it is a standard proposal coming from the electronic business research community that aims at creating a single global electronic marketplace where enterprises of any size and in any geographical location can meet and conduct business with each other through the exchange of eXtensible Markup Language (XML) based messages.

In The PROQOS project a very pragmatic approach was adopted. The idea was to make use of proven concepts, technologies and mechanisms to the maximum possible extent and, in that way, to reduce project risks and accelerate the achievement of the project's goals in a short, 6-month time window.

Given the above, in addition to the natural adoption of the DiffServ proposals, the project followed the TEQUILA proposals for SLS specification, as it was considered the broader, most mature, and best accepted work in the field. The ebXML architecture was not used because of the lack of stability of the proposal, having been decided to use the Internet Open Trading Protocol (IOTP) [13] architecture instead.

Although several projects in the SLA/SLS arena have been reported in the literature [14, 15], of which the ones mentioned in this section are the most relevant to the problem at hands, it was not possible to inherit more than general knowledge and lessons learnt from them. In most cases, the PROQOS goals are targeted to a particular, different scenario—a specific business model. In other cases the solutions worked out in those projects are not fully documented and/or available, which prevented an evaluation of their adequateness/applicability.

In contrast, the PROQOS project shows how it is possible to support QoS and dynamic SLA negotiation using well-known, standardized technologies, and makes available a simple and pragmatic platform for supporting hands-on assessment and testing in this domain.

3. PROQOS OVERALL FUNCTIONALITY AND ARCHITECTURE

The developed system is a prototype for demonstrating the general benefits of applying DiffServ on IP/DVB satellite links. To ease integration between all

components, the complete development was performed using the Linux operating system.

The developed system implements DiffServ mechanisms in the environment of an IP/DVB Uplink Hub [1, 2] . As such, the system implements the following overall functionality:

- Interfacing to one or more Network Service Providers, which feed IP traffic streams into the system (these traffic streams consist of DiffServ marked packets);
- traffic classification of the incoming streams according to their Differentiated Services CodePoints (DSCPs) and source/destination address pairs;
- handling of unicast, multicast and broadcast flows. As one of the strengths of satellite communications is the inherent support of broadcasting and multicasting, the PROQOS system is capable of forwarding multicast and broadcast traffic, in addition to unicast traffic. However, for multicast traffic the PROQOS system does not act as a multicast router, being transparent within the datapath and assuming multicast routers both in the ISP nodes and in the networks behind the customer landing points. Traffic with multicast destination addresses is forwarded unchecked and unchanged by PROQOS to multicast routers on the receiver side, which are responsible for maintaining the tree for delivery to the receivers within a multicast group;
- traffic shaping of the streams;
- multiplexing of the streams into one IP/DVB uplink stream;
- uplink of the streams;
- automated bandwidth and SLA trading with the interfaced network providers, in a distributed environment;
- support of standard IP security (IPSec) mechanisms for tunneling and encryption. For satellite transmission it must certainly be assumed that in most cases the transported payload is secured against unauthorized access by means of encryption. Hence, systems like ProQos must be able to transport both encrypted and unencrypted traffic.

Figure 3 presents the PROQOS system architecture.

As users of the system, organizations called Space Link Providers are assumed. An SLP owns or has leased a certain uplink capacity on a communication satellite and operates an uplink hub. An SLP may be identical or distinct from a satellite operator. It is assumed that the interest of SLPs is to provide services to different Network Service Providers (NSP).

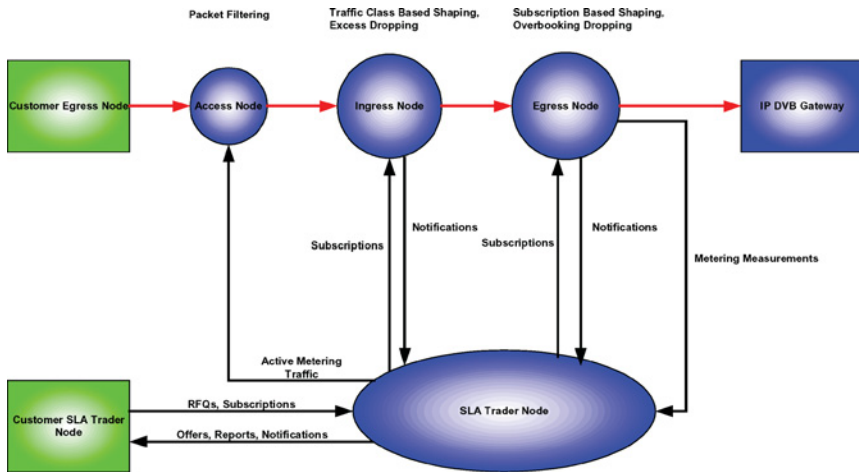


Fig. 3. General system architecture.

The PROQOS system performs its functions in a fully automated way, with only infrequent operator supervision/intervention. The system interfaces to the following external entities:

- Per connected NSP, the system shall connect to
- one NSP egress node; this node feeds the payload data of the NSP into the system;
- one NSP SLA trader node; the NSP SLA trader node connects to the system's SLA trader node; both systems perform automated negotiation of bandwidth and other service guarantees based on actual system workload;
- The IP/DVB Gateway interfaces to a DVB modulator, which is responsible for generating the uplink signal.

The traffic is injected into the system by several Customer Egress Nodes through an Access Node. This node is in charge of rejecting possible packets from unidentified and/or unauthorized sources. All authorized traffic is forwarded to the Ingress Node. Here, prioritization is performed according to the service classes and nonpaid excess traffic is dropped.

The Egress Node performs prioritization according to the actual ratio of customer subscriptions and handles overbooking situations applying a least-cost dropping approach, before it forwards the aggregated traffic to an IP/DVB Gateway for encapsulation.

The SLA Trader Node is responsible for communication with its counterpart on the customer side, in order to accept and answer requests for quotation. The status of subscriptions is forwarded to the ingress and egress nodes, to allow them

to perform decisions for prioritization and packet dropping. Furthermore, at the SLA Trader Node all metering information is aggregated, in order to enable the generation of notifications and reports to the customer.

This architecture follows the provisioning model of the Policy-Based Network Management (PBNM) approach proposed by the IETF's Policy Framework Working Group. The Policy Decision Point (PDP) is embodied by the SLA Trader Node, and the Access Node, Ingress Node and Egress Node are Policy Enforcement Points (PEP).

The PBNM model has some problems that are well known to the research community—mapping policies to network resources and their corresponding configurations is, perhaps, the most important. The ProQoS system architecture solves this problem by limiting the number of network nodes and restricting their traffic handling functions. As such, system nodes need only to support the services that can be provided by the predefined SLA template.

The SLAs established between the system owner and its customers define four traffic classes, each of them having fixed guarantees for packet drop ratio and packet latency. The SLA parameters for each traffic class are, then, restricted to

- a nominal bandwidth specification,
- a price for the nominal bandwidth, and
- a price for traffic exceeding the nominal bandwidth (excess traffic).

The network has a fixed number of nodes. Each node has a well-defined traffic handling function (see Fig. 3). To perform each of these functions on the traffic of each class of each customer, a fixed set of DiffServ mechanisms is needed. The configuration of these mechanisms is directly derived from the three contract parameters listed above.

4. SERVICE LEVEL AGREEMENTS

The PROQOS system is to be used by an SLP that provides network access to an IP/DVB Gateway. The SLP domain is able to offer differentiated treatment to traffic classes that are charged according to the degree of quality received within the SLP DiffServ domain. The typical customers of such a service are ISPs, and we will use the terms ISP and customer interchangeably.

The agreement for service provisioning always starts with the signature of an SLA by two entities—the ISP and the SLP. After this step the SLP operator will enter the signed SLA data into the system. From this point onwards, the ISP is able to dynamically and automatically negotiate with the SLP some of the SLA's parameters.

While an SLA is active, the SLP must monitor the QoS provided in its context. This monitoring data is used for reporting the system usage to the system operator, through the Reporting Interface, and for sending QoS summary reports and SLA




Fig. 4. Interfaces within the PROQOS system.

Conformance/Nonconformance notifications to the corresponding ISP through the SLA Trading Interface.

Some of the monitoring data is collected by the Access, Ingress, and Egress nodes of the system, and therefore an interface is needed to report this information to the SLA Trader Node. On the other hand, to configure the SLP DiffServ nodes according to the SLA information in the system (either entered by the SLP operator or negotiated by the ISP) an interface between the SLA Trader Node and the remaining system nodes is also needed. The SLA Exchange Interface is defined for these purposes. Figure 4 shows a context diagram of the various interfaces within the PROQOS system.

The SLA Exchange Interface is a scaled-down Common Open Policy Service—Provisioning model (COPS-PR) [16, 17] communication interface, used to exchange SLA configuration and monitoring data between the SLA Trader Node and the other nodes of the system. As mentioned in the previous section, in the SLA Exchange Interface the SLA Trader Node plays the policy decision point (PDP) role. The Access Node, Ingress Node and Egress Node play the policy enforcement point (PEP) role.

A full COPS-PR protocol implementation was used to support the SLA Exchange Interface. The usage of the protocol by PDP and PEPs is said to be scaled-down because some of the protocol mechanisms and facilities are not used:

- The security mechanism is not used,
- The secondary connection verification mechanism (Keep Alive) is not used,
- Only one Client-Type (session) is used over each connection,
- The state synchronization facility is not used; if a synchronization problem arises the connection is closed and a new connection is established,
- The COPS-PR Accounting Timer object is not used.

The SLA Exchange Interface provides a means of exchanging the active SLAs' data between the PDP and the PEPs so that, at any time, the active SLAs data store in the SLA Trader Node and the PEPs' configuration are consistent. It also provides a means for PEPs to report network usage data to the PDP. Figure 5 represents the SLA Policy Information Base (PIB) tree.

SLA agreements cover the provision of an amount of bandwidth for one or more traffic classes with QoS guarantees, which subscribers are permitted to impose on the ingress node of the SLP, for uplink via an IP/DVB-S stream within the transponder capacity available to the SLP on the telecommunication spacecraft (i.e., satellite).

The QoS parameters specified in the SLA are guaranteed by the provider between the Access Node and the IP/DVB Gateway performing the encapsulation of IP Packets in the DVB uplink stream. The SLP further warrants that beyond this gateway no multiplexing, neither on IP nor on DVB level, with data streams other than originating at the SLP egress node is performed which may hamper the Quality of Service in the uplink.

The SLA framework covers the following QoS parameters fixed for the duration of the agreement:

- MTTR: 4 h
- MTBF: 2160 h
- Resulting availability: 99.82%
- Resulting maximum monthly downtime: 1h and 20 min

The SLP provider offers the following traffic classes, each of which has a set of QoS parameters inherently contained

- Streaming: identified by the Expedited Forwarding (EF) per-hop behavior (PHB) DiffServ Code Point (DSCP)
- Gold: identified by the Assured Forwarding AF1x PHB DSCP
- Silver: identified by the AF2x or AF3x PHB DSCP
- Best effort: identified by the Best Effort (BE) PHB DSCP

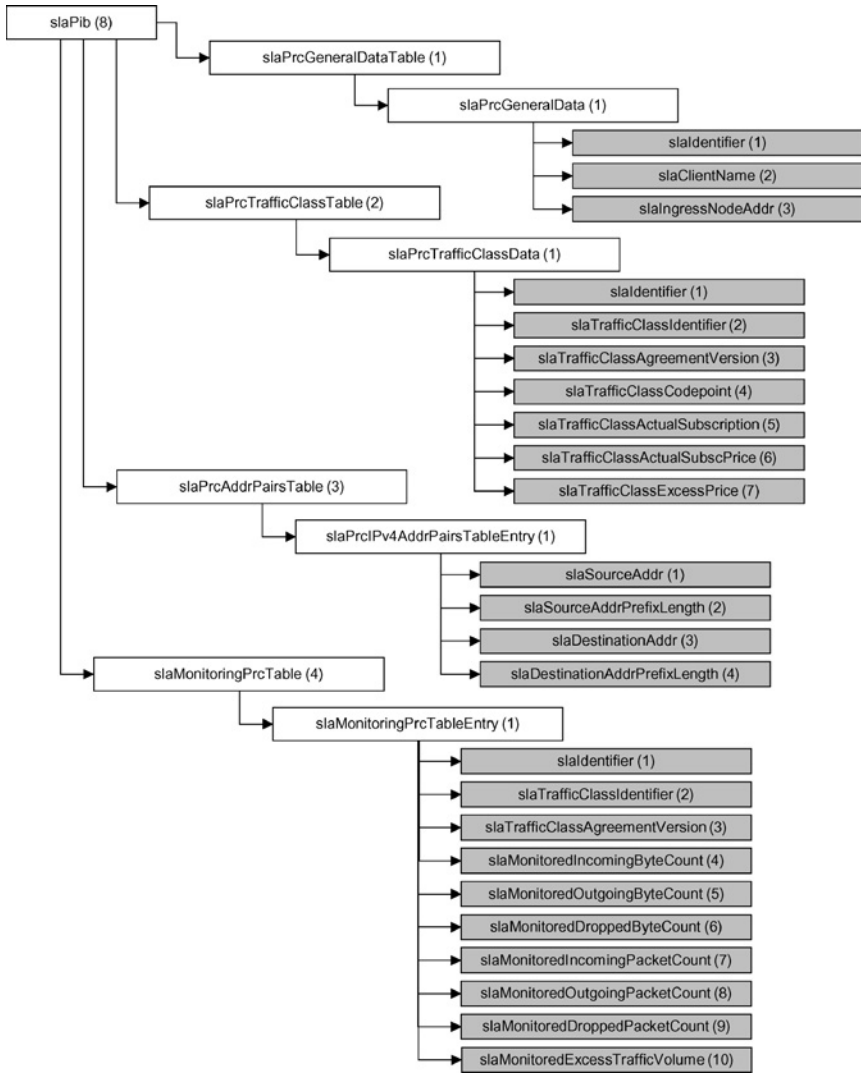


Fig. 5. SLA PIB tree.

For each traffic class, subscriber and provider agree on minimum and maximum bandwidth subscriptions. Within the range between minimum and maximum subscription, subscribers may at any time request a quotation for a new “actual subscription,” by means of an SLA trading node located at the subscriber’s site, which connects to a corresponding SLA trading node at the provider’s site. The

provider's SLA trading node will reply with a quotation for the requested bandwidth in the requested traffic class. Upon confirmation of this quotation by the subscriber's SLA trading node, the new "actual subscription" is sent to the Access, Ingress and Egress Nodes and a Deliver Response message is sent back to the subscriber's trading node notifying it that the new subscription became effective. This message exchange sequence is illustrated in Fig. 2.

Whenever a quotation for a new "actual subscription" is requested, not only the price for the "actual subscription" will be quoted, but also a price for an excess rate per kbit/s and minute will be quoted. In the confirmation response, subscribers may confirm this excess rate, which results in the price for the excess rate being charged in addition to the normal subscription rate whenever excess rate conditions occur. Alternatively, subscribers may reject such an excess rate pricing in the confirmation response. In this case, the IP Packet drop ratio associated with the traffic class is not guaranteed for excess rate traffic.

Whenever the provider is not capable of fulfilling the agreed QoS level, either in terms of static QoS parameters or parameters associated with traffic classes (namely, drop ratio and average maximum latency), charge-back mechanisms apply.

5. THE PROQOS TRADER NODE

As can be inferred from the previous sections, the Trader Node is the central building block of the PROQOS system, playing the role of the system coordinator in the Space Link Provider DiffServ network. In this section the main aspects of its operation will be highlighted.

Figure 6 presents the Trader Node block diagram. This node performs the following categories of functions:

- System configuration, including SLA management;
- SLA negotiation;
- QoS monitoring and metering;
- SLA conformance checking; and
- Periodic QoS reporting to customers.

The system's manager configuration interface is a web interface for Trader Node configuration and SLA data introduction. This interface allows for the start/stop of the Trader Node, the insertion, listing, updating and removal of SLAs, addition of traffic classes, changes of subscription data for a defined traffic class, removal of traffic classes, generation of unsolicited offers, and configuration of test traffic parameters. Figure 7 presents a screenshot of this interface.

The Trader Node uses the SLA Exchange Interface to update the configuration of all the other system nodes (Access, Ingress and Egress nodes) reflecting the changes introduced in the Active SLAs data store (maintained by the Trader

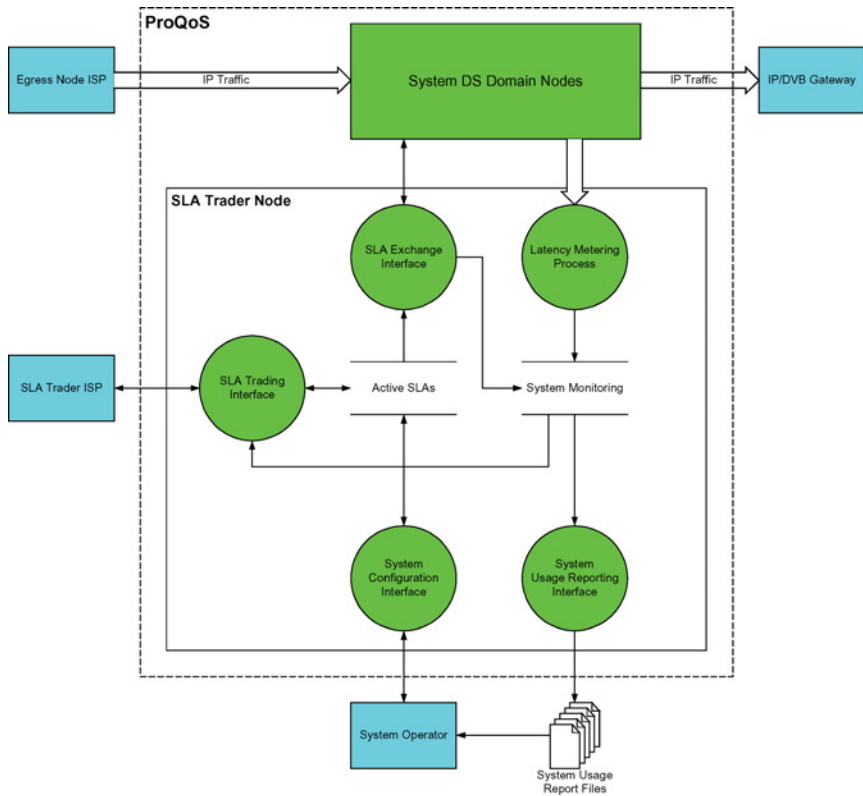


Fig. 6. Trader Node block diagram.

Node), triggered either by the system operator or by ISP request acceptance. As mentioned in section 4, this interface is a COPS-PR communication interface for the exchange of SLA's configuration and monitoring data.

The SLA Trader Node allows an ISP to dynamically request a new service, in the scope of an SLA previously agreed upon with the SLP. The system negotiates an active SLA with its subscriber ISP through the SLA Trading Interface. The procedure is initiated either by an ISP request or by an unsolicited bandwidth offer operation.

The SLA Trading Interface is based on a scaled-down IOTP (Internet Open Trading Protocol [13, 18]) Version 1.0 implementation, which implements only the Purchase transaction involving only two trading roles—the Merchant and the Consumer. To accomplish this, the following message types are used: Request for Quotation (reception), Offer (sending), Subscription update (reception), Delivery

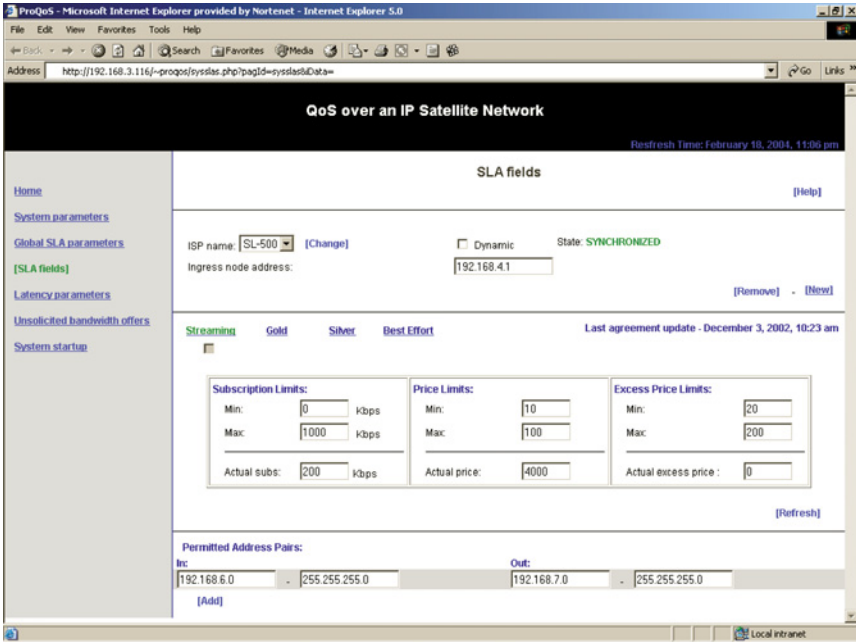


Fig. 7. Screenshot of the system management configuration interface.

Response (sending) and Unsolicited Bandwidth Offer (sending). Moreover, the SLA trader node provides the following additional message types which are out of the scope of any IOTP transaction: Out-of-Qos Notification (sending), Back-in-Qos Notification (sending), QoS Warranty summary report (sending). These message types are sent asynchronously by the Provider SLA trader node in accordance with the system status and configuration. Messages exchanged between the SLA trader node and subscribers are XML documents. Distinction between documents is made based on the XML document prolog.

The SLA trader node maintains a data store to keep the monitoring data collected by the system. The System Monitoring data store is composed of records with bandwidth usage information, latency metering information and availability metering information. The SLA trader node collects bandwidth usage metering data from the system nodes using the SLA Exchange Interface.

The SLA Trader Node also measures average packet latency per active traffic class between the Access Node entry interface and the Egress Node output interface. Measurement of average packet latency is made by injecting packets with a timestamp at the input interface segment of the Access Node, with characteristics that correspond to all current subscribed traffic classes per customer ISP. The

IPTables module of the Linux kernel is used in order to have the injected packets routed back to the originating node, ensuring that the same time source is used in calculations. The packet's arrival time is compared with the initial timestamp. The results are used to calculate the Average Packet Latency over one minute time intervals, for each traffic class, per ISP.

Another Trader Node function is SLA conformance evaluation. Whenever new data is added to the System Monitoring data store, the SLA Trader Node evaluates the system conformance to the corresponding SLA. The conformance evaluation is based on the data in the System Monitoring data store and on the parameters stored in the Active SLAs' data store. Additionally, the SLA Trader Node keeps track of the state associated with each active SLA. This state shall be In-QoS or Out-of-QoS. After every execution of an SLA conformance evaluation operation, the SLA Trader Node updates the state associated with each SLA to In-QoS or to Out-of-QoS, according to the outcome of the operation.

Every time that the state associated with one active SLA changes, the SLA Trader Node sends a notification message to the ISP that subscribes that SLA, through the SLA Trading Interface. In addition, the SLA Trader Node sends a QoS Warranty Summary Report to each ISP every QoS Reporting Period, through the same interface.

Finally, the SLA Trader Node reports the data in the System Monitoring data store to the system operator through the System Usage Reporting Interface. This interface consists of human readable text files that can also be parsed by external billing software. The files include all the data required to charge each customer according to the corresponding SLA.

6. TESTS

Both quantitative and perceptual tests were conducted to validate the performance of the PROQOS system. The tests were performed both using a simulated space link (via the Open Source Software "nistnet") and 2 Mbit/s capacity on SESAT provided by ESA.

Quantitative end-to-end tests were conducted using the Open Source software "tg" as traffic sources/sinks on the same platform. A set of seven different traffic patterns reflecting nominal and nonnominal load conditions were injected into the system. Figure 8 shows the test setup for quantitative tests over SESAT.

As example, Figs. 9 and 10 show the test results for the bandwidth injected and received during a dynamic SLA negotiation process.

The traffic is initially subscribed to 200 kbit/s, no payment of excess traffic is agreed. At $t_0 + 30$ s the sender exceeds the agreed bandwidth, traffic is shaped down by the PROQOS system to the agreed rate. At $t_0 + 60$ s the subscription is updated to 750 kbit/s and is shaped to the new bandwidth. At $t_0 + 90$ s the subscription is again updated to include payment of excess traffic. Now all traffic

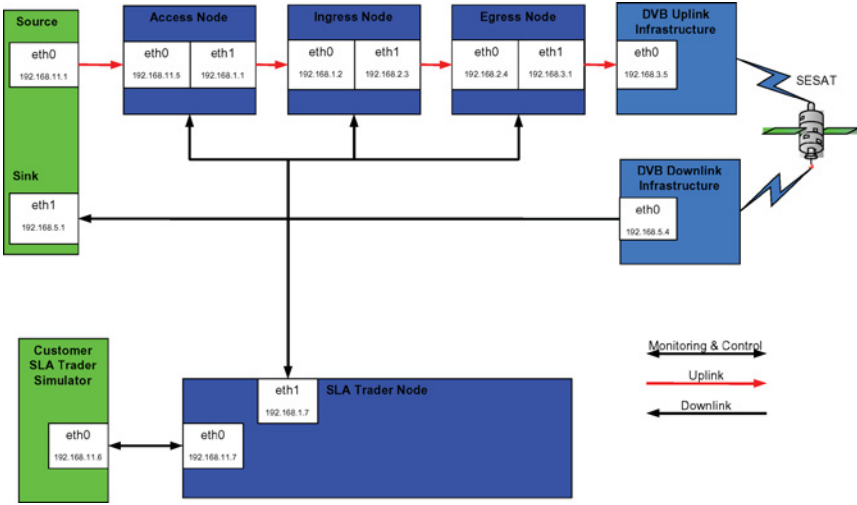


Fig. 8. Setup for quantitative tests.




Fig. 9. Sender rate.

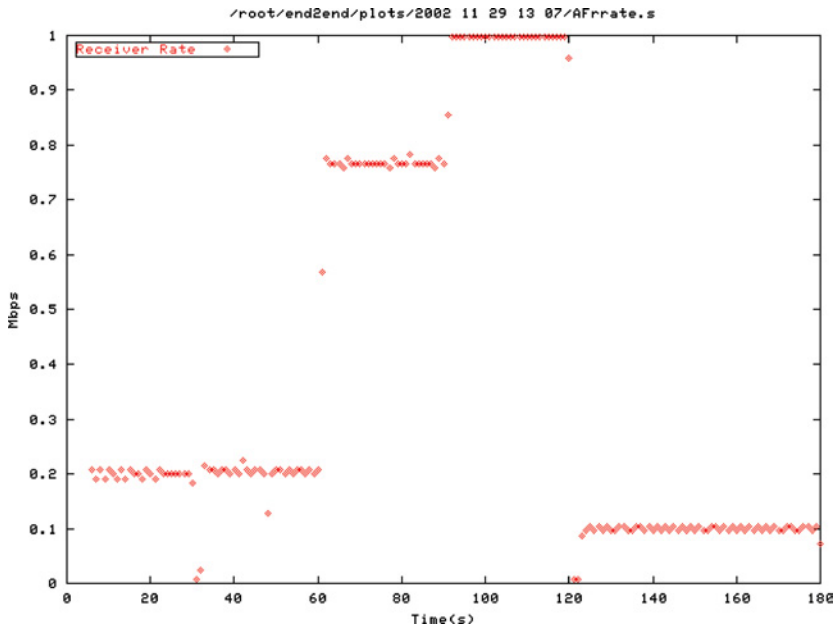


Fig. 10. Received rate.

is passed through the PROQOS system to the receiver's site. At $t_0 + 120$ s the injected traffic is reduced to 100 kbit/s with the latest subscription update still valid.

In order to test and visualize the benefits of QoS for application traffic, perceptual end-to-end tests were conducted using streamed video in one traffic class and observing the occurrence or nonoccurrence of visible and audible artifacts on the receiver side with the following test setup.

1. A streaming server with a test video of 69 seconds length was set up. This video was requested from a streaming client via RTSP (Real Time Streaming Protocol, <http://www.rtsp.org/>) with minimum buffering. The bandwidth occupied by the test video was almost constant 1450 kbit/s. The available space link bandwidth (incl. encapsulation overhead) was 2048 kbit/s.
2. This video was received at the client without visible or audible degradation both via local loop and via space link.
3. In the next step on all nodes in the uplink chain, queuing disciplines were removed. Hence the nodes acted as simple best effort routers. Then again the test video plus additional distortion traffic was injected into the chain.

Serious visible and audible degradation occurred in the streaming client side, which in some cases led to the crash of the streaming client software.

4. Finally the test as of step 3 was re-run with PROQOS queuing disciplines installed so that the distortion traffic had a lower subscription price than the streaming traffic. The resulting display at the streaming client was identical to that of step 2, although the space link had been overbooked by almost 75%.

Summarizing, the test proved the feasibility of providing QoS in a IP/DVB-S environment including the dynamic negotiation of QoS parameters. However, two particular areas requiring further research were identified during the tests.

1. The encapsulation overhead induced by the IP/DVB gateway is not negligible. Hence the development of more predictable encapsulation methods and/or the integration of IP/DVB gateways into the overall PROQOS monitoring & control, e.g. by means of a feedback channel from the gateway to the egress node, is desirable.
2. With the used queuing mechanisms Class Based Queueing (CBQ) and Hierarchical Token Bucket (HTB), the accuracy in limiting outgoing bandwidth is dependent on the packet size.

7. CONCLUSIONS AND FUTURE WORK

The project presented in this paper demonstrated the applicability of Diff-Serv QoS mechanisms to IP/DVB-S systems and showed that, although with considerable development effort, it is possible to provide mechanisms for flexible and dynamic negotiation and management of Service Level Agreements in this traditionally static SLP-ISP environment. Distribution of various aspects of SLA management such as configuration, provisioning, negotiation, monitoring, accounting, and conformance checking was achieved, with clear benefits in terms of functionality, reliability, scalability and efficiency, when compared to the solutions in use in today's IP-over-satellite systems. It is also noteworthy to mention that the developed prototype used standard solutions as a basis, namely COPS and IOTP, although these have been scaled down in order to simplify the implementation.

In addition to the research topics identified during the tests, further work will address the deployment of the system in a real environment, for which there are already some plans. The challenges of this deployment include charging issues, the refinement of the reporting facilities of the system, including client service history, and the automation of the generation of unsolicited offers.

Another challenge is the development of strategies for integrating DVB and QoS negotiation and enforcement technologies (which has not been explored in this project). It is obvious that the use of integrated mechanisms for doing

opportunistic multiplexing in DVB links can strongly improve the effectiveness and efficiency of the presented solution.

Last but not least, it is worthwhile to mention that the presented prototype has characteristics that make it very useful not only in space environments but also in terrestrial environments. It proves that it is possible to address the challenges of dynamic SLA negotiation, using simple approaches and well-known technologies.

REFERENCES

1. K. Nichols, S. Blake, F. Baker, and D. Black. RFC 2474—Definition of the Differentiated Services Field (DS Field) in the IPv4 and IPv6 Headers. December 1998.
2. S. Blake, D. Black, M. Carlson, E. Davies, Z. Wang, and W. Weiss. RFC 2475—An Architecture for Differentiated Service. December 1998.
3. Internet2 URL: www.internet2.edu
4. INTAP, Survey on Policy-Based Networking—Addressing Issues, Technological Trends, Future Prospects of Policy Exchange Methods in Multi-Domain Scenarios, INTAP, June 2001. Available at <http://www.net.intap.or.jp/INTAP/>
5. Qbone architecture (v1.0), Internet2 QoS Working Group Draft, August, 1999, Editor: Ben Teitelbaum. Available at <http://www.internet2.edu/qos/wg/papers/qbArch/1.0/draft-i2-qbone-arch-1.0.html>
6. Ben Teitelbaum and Phil Chimento, QBone Bandwidth Broker Architecture, Work in Progress. Available at <http://qbone.internet2.edu/bb/bboutline2.html>
7. Information Society Technologies Programme (IST), European Union's Fifth RTD Framework Programme 1998–2002, IST KA4/E1 projects: Computing, Communications and Networks. Available at <http://www.cordis.lu/ist/ka4/ipcn/projects.htm>
8. TEQUILA URL: <http://www.ist-tequila.org/>
9. D. Goderis, S. van den Bosch, Y. T'joens, O. Poupel, C. Jacquenet, G. Memenios, G. Pavlou, R. Egan, D. Griffin, P. Georgatsos, L. Georgiadis, and P. van Heuven, Service level specification semantics and parameters, Internet Draft document, February 2002. draft-tequila-sls-02.txt
10. CADENUS URL: <http://www.cadenus.org>
11. S. D'Antonio, M. Esposito, M. Gargiulo, S. P. Romano, and G. Ventre, Resource usage monitoring in SLA networks, April 2002—CADENUS deliverable D2.2.
12. ebXML Technical Architecture Project Team, ebXML Technical Architecture Specification v.1.0.4, Technical Specification—ebXML_TA_v1.0.4.doc. Available at <http://www.ebxml.org>
13. Burdett, D., RFC 2801—Internet Open Trading Protocol—IOTP Version 1.0. April 2000.
14. A. Keller and H. Ludwig, The WSLA framework: Specifying and monitoring service level agreements for web services, *Journal of Network and Systems Management*, Vol. 11, No. 1, 2003.
15. MESCAL project, Specification of Business Models and a Functional Architecture for Inter-domain QoS Delivery, Deliverable 1.1, IST-2001-37961, June 2003.
16. J. Boyle, R. Cohen, D. Durham, S. Herzog, R. Raja, and A. Sastry, "The COPS (Common Open Policy Service) Protocol," RFC 2748, January 2000.
17. K. Chan, J. Seligson, D. Durham, S. Gai, K. McCloghrie, S. Herzog, F. Reichmeyer, R. Yavatkar, and A. Smith, "COPS Usage for Policy Provisioning (COPSPR)," RFC 3084, March 2001.
18. D. Burdett, *Internet Open Trading Protocol*, 1st edn., McGraw-Hill, 2000; ISBN: 0071355014.

Manuel Pessoa worked as Senior Engineer in the Telecom and Networking unit at Critical Software, Portugal, where he participated in a considerable number of national and international

projects in the networking area. He is currently a teaching assistant at Polytechnic Institute of Castelo Branco and a PhD student at the University of Coimbra.

Antonio Alves works as Senior Engineer in the Telecom and Networking unit at Critical Software, Portugal. He has more than 10 years of experience in R&D projects such as ForeThought ATM platform for FreeBSD, New Service model for Internet Communications, Software Implemented Fault Injector for JPL/NASA Testbed, Development of Protocols for Quality of Service (ESA), and Edge Device.

Goncalo Quadros graduated in 1987 in Electrical Engineer and has worked in industry and R&D prior to obtaining his PhD in Computer Science from the University of Coimbra in 2002. He is cofounder of Critical Software in mid 1998. Currently he is Critical's Vice-President with administration and business development functions.

Fernando Boavida received his PhD in Informatics Engineering in 1990, and currently is an Associate Professor at the Department of Informatics Engineering (DEI) of the Faculty of Sciences and Technology of the University of Coimbra. He is the leader of the Laboratory of Communications and Telematics of DEI and the Strategic Director for Communications and Information Technology of the University of Coimbra.

Michael Henke works as Senior Project Manager in the Division of Information and Communication Solutions for Space at VCS Aktiengesellschaft, Bochum, Germany. He received his MSc in Computer Science from the University of Applied Sciences, Dortmund, Germany in 1994. He has performed project management and engineering tasks for several major European missions, such as the Columbus Control Centre and the future Galileo Ground Mission Segment.

Milva Natcheva works as a Software Engineer in the Division of Information and Communication Solutions for Space at VCS Aktiengesellschaft, Bochum, Germany. She received her MSc in Computer Science from the Sofia University "St. Kliment Ohridsky" Bulgaria in 1999. She worked at several European missions, e.g., the Columbus Control Centre, as a Software Element Manager.

Patrick Halke works as Software Engineer in the Division of Information and Communication Solutions for Space at VCS Aktiengesellschaft, Bochum, Germany. He has been working as a consultant for several software companies for about 12 years and beside ProQoS has also taken part in the development of the Data Services Subsystem for the Columbus Control Centre.

Peter Maurutschek is a project assistant at the Institute of Computer Science at the University of Salzburg, working in the field of computer networks and distributed systems. His research interests concentrate on IP data broadcasting over a digital carrier using DVB.

Zenon Huskic and Kurt Wagner are Professors at the Institute of Computer Science at the University of Salzburg, working in the field of computer networks and distributed systems.

Frank Zeppenfeldt works at the Telecommunications Department of ESA as a communications engineer. Before joining ESA, he worked at a NATO Research Laboratory where he was responsible for the technical implementation of a research network for secure collaborative experiments between NATO partners.

Roberto Donadio has been with ESA since 1989, first at the ESA Control Centre in Darmstadt, Germany, and is currently a Senior Applications and Standardization Engineer at the Research and Development Centre of ESA in Noordwijk, The Netherlands. He has been actively involved in standardization matters in the ITU and ETSI and he currently represents ESA in the DVB-RCS technical committee.