

The use of COPS and NSIS in the EuQoS Project

E. Monteiro, F. Boavida, P. Simões, J. Sá Silva, L. Cordeiro, R. Eugénio, M. Carmo

University of Coimbra
Laboratory of Communications and Telematics
CISUC-DEI
Pólo II, 3030-290 Coimbra – Portugal
edmund,boavida,psimoes,sasilva,cordeiro,romulo,maxweel}@dei.uc.pt

Abstract. This paper discusses the use of the Policy Based Network paradigm, COPS (Common Open Policy Service) and NSIS (Next Steps In Signaling) frameworks being developed by the IETF, in the context of the European Project EuQoS (End-to-End Quality of Service over Heterogeneous Networks). An overview of the EuQoS architecture is provided followed by the analysis of the use of PBN/COPS and NSIS to support the Signaling and Service Negotiation (SSN) function. The implementation status and validations aspects are also discussed in the paper and some preliminary results are included.

1 Introduction

The motivation of the European Integrated Project EuQoS (End-to-end Quality of Service support over heterogeneous networks) [1] is to solve the outstanding design issues presently associated with the delivery of end to end QoS service across heterogeneous networks. It is necessary to resolve these issues and accordingly upgrade the infrastructures so that new applications can be supported by the Internet and new service packages can be offered by operators and service providers.

The key objective of EuQoS is to research, integrate, test, validate and demonstrate end-to-end QoS technologies to support the infrastructure upgrade for advanced QoS-aware applications over multiple, heterogeneous network domains, belonging to research, scientific and industrial communities.

The project will deliver the EuQoS system which will support the delivery of end to end QoS. As QoS is primarily a challenge for the access network, the EuQoS system will be developed and trialed on various types of research access networks with the GEANT [2] core providing Pan European support. This heterogeneous infrastructure, which models the production networks of the future, requires a QoS technical solution that has not been synthesized to date.

The EuQoS project will propose and develop new QoS mechanisms which build upon the state of the art and incorporate the following components: Monitoring and Measurements, Admission Control, Failure Management, Signaling & Service Negotiation, Security and AAA, Charging and Traffic Engineering and Resource Optimization. EuQoS will integrate state of the art protocols and technologies with new

mechanisms to be developed in the project in order to build the above specified functions

This paper addresses the use of the Policy Based Network (PBN) paradigm [3] and its support protocol, the Common Open Policy Service (COPS) [4] together with the NSIS (Next Steps in Signaling) IETF (Internet Engineering Task Force) framework [5] to support some of the EuQoS functionalities.

The remaining of the paper is structured as follows. In Section 2 the EuQoS architecture is briefly described. Sections 3 and 4 will provide some background about PBN and NSIS, two important EuQoS building blocks for the EuQoS architecture. Section 5 will describe the proposed solution. Section 6 will discuss validation issues. Finally, Section 7 will conclude the paper and point out some directions for future (already ongoing) work in the scope of the EuQoS project.

2 EuQoS Architecture

The EuQoS project aims at creating a flexible and secure QoS Assurance System (the EuQoS System) by developing new QoS mechanisms which build upon the state-of-the-art. The EuQoS System consists of two major research components: User and QoS aware Control Plane and QoS Protocol Stack. The QoS Protocol Stack will provide a new API over existing and new transport protocols that will provide variable levels of order and reliability.

The Control Plane will include a set of functions that might be supported in network elements such as routers and in end systems. To integrate the Control Plane of the EuQoS architecture six main functions were identified:

Function 1 – Signaling and Service Negotiation (SSN); Function 2 – Connection Admission Control (CAC); Function 3 – Monitoring Measurement, Fault Management (MMFM); Function 4 – Traffic Engineering and Resource Optimization (TERO); Function 5 – Security and AAA (SAAA); Function 6 – Charging (CHAR).

Figure 1 shows a simplified view of the EuQoS system.

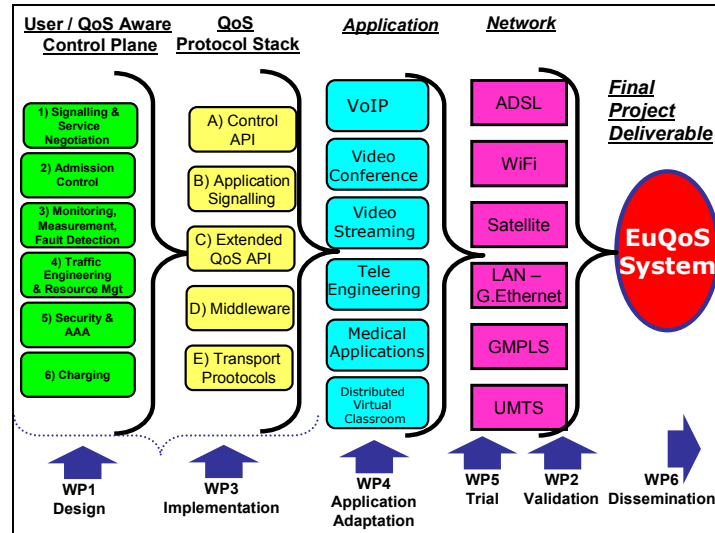


Fig. 1. The EuQoS system

The Signaling and Service Negotiation (SSN) function is responsible for the dynamic configuration of the communication system, in order to provide the requested level of QoS to applications. The SSN function covers the application connection setup, data transfer and teardown phases and is triggered by applications located at end systems, by application proxies or by network elements.

The EuQoS SSN function will be built over state-of-art signaling and service negotiation mechanisms complemented with new functionalities to be researched and developed in the EuQoS project. Among other proposals the PBN, COPS and NSIS all being developed at IETF will play an important role in EuQoS as will be described in the following sections.

3 PBN and COPS in EuQoS architecture

This section discusses the use of the Policy-Based Networking and the Common Open Policy Service in the EuQoS project context.

3.1 Policy-Based Networking Concepts

The PBN concept and architecture were developed by IETF's Policy Framework group [3]. The objective of PBN is to enable centralized network administration by supporting the definition of abstract rules, that is, by specifying what the network should do in a high-level language instead of specifying how to do it. This is done in a network-element-independent way, through a policy specification language.

The PBN architecture comprises four basic components: Management Console, Policy Repository, Policy Decision Point (PDP) and Policy Enforcement Point (PEP). These components use two kinds of protocols – repository access protocols and policy transfer protocols (Figure 2).

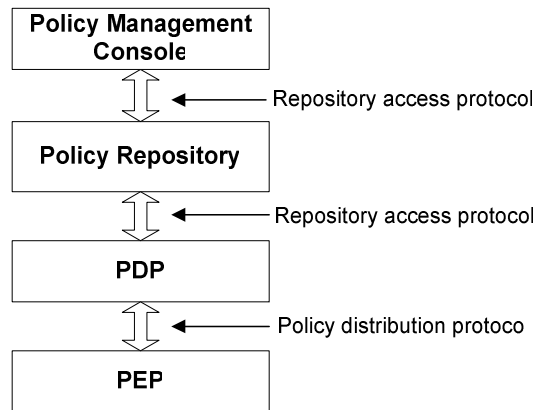


Fig. 2. Components of the PBN architecture

The Management Console provides an interface between the administrator and the policy management system. This interface allows for the specification, edition, translation and validation of the policies stored in the policy repository. Typically, this Console has a Web-based interface.

The PR stores the policies that are applicable to a given administration. Every policy database is structured as a Policy Information Base (PIB) – a common structure which resembles SNMP Management Information Bases – independently of the specific implementation approaches (plain text files; centralised or distributed relational databases; LDAP-schemas; etc.).

PDPs, also called Policy Servers, translate high-level policies stored in the policy repository into lower level policy rules. By combining this information with other network information (e.g. information regarding the status of the network), PDPs can produce the precise rules to be implemented in the Policy Enforcement Points.

PEPs receive rules from PDPs and deal with network elements accordingly. PEPs that are integrated with network elements can perform actions such as packet filtering, packet marking in addition to other resource management operations. Besides enforcing the rules established by PDPs, PEPs can send information to PDPs as, for instance, happens when physical configurations change.

Figure 3 shows a typical PBN configuration. Communication between PDPs and PEPs is achieved using a policy transfer protocol such as COPS protocol. The information to be transmitted by COPS is stored in a Policy Information Base (PIB).

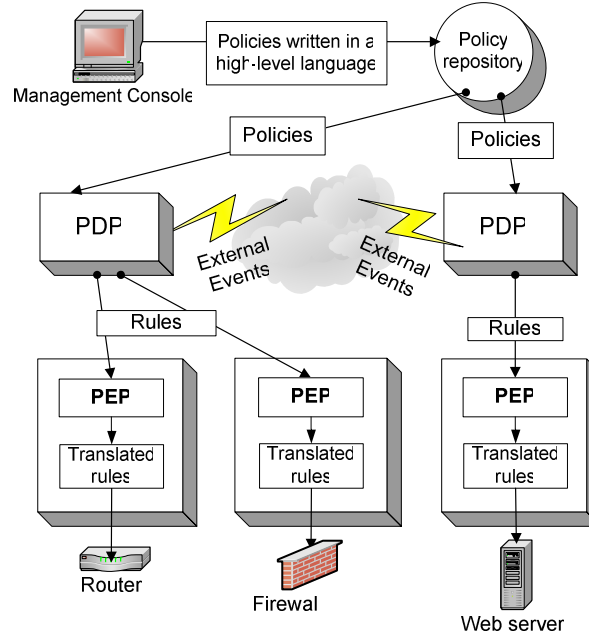


Fig. 3. Typical PBN configuration

The original version of COPS [4] corresponds to the Outsourcing Model, in which PEPs contact PDPs each time they need information on how to treat a particular event (i.e., there is an one-to-one relation between requests and decisions). PDPs answer with the appropriate rules to be used by PEPs. COPS uses TCP for reliable communication, supports security mechanisms (authentication, replay protection and integrity) and provides a statefull protocol: PDPs keep track of installed rules and react to requests from PEPs in accordance with these rules.

COPS soon evolved to fulfil the requirements of the Provisioning Model, leading to the appearance of COPS-PR (COPS for Policy Provisioning) [6]. According to this model, PDPs send all the information required by PEPs at start up. All the relevant rules are stored in PEPs and all events are treated using this information, without additional need for PEP-PDP communication.

3.2 PBN in the Context of EuQoS

PBN can play an important role in the EuQoS framework, since it provides an elegant and technology independent schema to map high-level QoS domain policies into low-level network equipment configuration, coping both with the required autonomy of QoS-domain management boundaries and the need to establish a network technology independent sub-layer that aggregates the resource managers of the various QoS domains.

Figure 4 illustrates how COPS fits into the EuQoS reference architecture. Each administrative domain (QoS Domain) maintains its own Policy Repository and its own Policy Decision Point. The Policy Repository stores domain-specific policies according to an LDAP (Lightweight Directory Access Protocol) schema. Those policies are then used at two distinct levels:

- To define the technology-independent behaviour of the Resource Manager (which QoS requests should be satisfied, under what circumstances);
- And to translate QoS requirements into specific network equipment configuration, using COPS-PR for the communication between the PDP (usually located in the Resource Manager) and the managed network nodes.

In this way, PBN plays an important role in the definition of the QoS Domain as a whole (policies regarding interaction with peer QoS domains) and in the translation of high-level policies into equipment specific configurations – using COPS-PR in the communication between the PDP and the various Policy Enforcement Points.

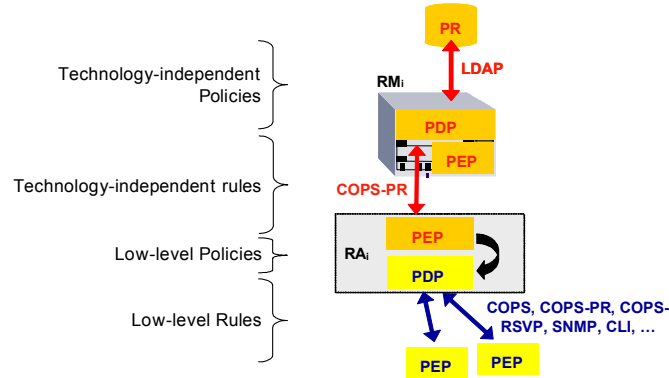


Fig. 4. PBN in the EuQoS context

4 NSIS in the EuQoS Context

This section discusses the use of the NSIS framework and protocols in the EuQoS project context.

4.1 NSIS Concepts

The Next Steps in Signaling suite of protocols is being developed by the IETF NSIS working group [5]. The main objective of the NSIS protocols is to support various signaling applications that need to install or manipulate state in the network.

The NSIS framework aims at providing signaling elements for the global Internet environment, to answer different needs (e.g. resource reservation, firewall transversal), without requiring a complete end-to-end deployment. Signaling protocol mes-

sages can take either the same path as the data path between endpoints or a different path (on-path and off-path signaling).

The NSIS protocol stack is divided in two layers: a generic lower layer for signaling transportation and an upper layer specific to each signaling application. For the lower layer a General Internet Messaging Protocol for Signaling (GIMPS) [7] is currently being standardized. For signaling data transport, GIMPS can use TCP or UDP (or other transport protocols being standardized). For the upper layer the NSIS working group is currently working in two protocols: the QoS NSIS Signaling Layer Protocol (QoS-NSLP) [8] for resource reservation and QoS negotiation and the NAT/Firewall NSIS Signaling Layer Protocol (NAT/FW-NSLP) for firewall transversal [9].

4.2 NSIS in the context of EuQoS

The NSIS framework is well fitted for the EuQoS project architecture. GIMPS provides the signaling transport mechanism that can be used between Resource Managers (off-path signaling) and also between routers and other network equipment (on-path signaling). This protocol also provides an abstraction layer that can be used for all the high-level signaling functions needed in the EuQoS framework. QoS-NSLP provides mechanisms to establish and maintain state at nodes along the path of a data flow in an end-to-end environment.

QoS-NSLP is similar to RSVP [RFC2205]. It is based in soft-state refresh messages for state installation and refreshment. However the signaling is performed between pairs of adjacent nodes rather than in an end-to-end fashion along the complete signaling path like RSVP. Also unlike RSVP, QoS-NSLP supports sender or receiver-initiated reservations, bi-directional reservation and reservations between arbitrary nodes (e.g. edge-to-edge, end-to-access). Nevertheless, QoS-NSLP does not support IP multicast as happens with RSVP.

In the context of EuQoS, QoS-NSLP can be used for resource reservation between RMs, across network domains.

The main limitation for the adoption of the complete NSIS framework in the EuQoS context is its relative immaturity which affects primarily the application signaling protocols like QoS-NSLP. To overcome this limitation, the use of NSIS in EuQoS will start by GIMPS and will progressively include QoS-NSLP mechanisms. In the meanwhile, a simplified version of QoS-NSLP will be developed and deployed.

5 Putting it all together: the SSN Function

The Signalling and Service Negotiation (SSN) function is responsible for the dynamic configuration of the communication system, in order to provide the requested level of QoS to applications. The SSN function covers application connection setup, data transfer and teardown phases and can be triggered by applications located at end systems, by application proxies or by network elements.

5.1 General Overview of the SSN Function

There are four different “signalling levels” in the EuQoS architecture, as illustrated in Figure 5 by the various red arrows (Level 1 to Level 4).

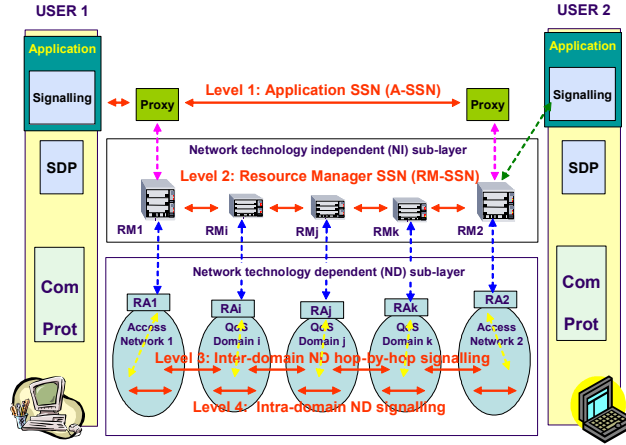


Fig. 5. Signaling levels in the EuQoS architecture

Besides horizontal signalling interactions between peer entities (i.e., signalling entities of the same signalling level), there are other, complementary signalling interactions that take place vertically between adjacent levels. These interactions are named “cross-layer signalling” and are identified by the vertical arrows in Figure 5. Cross-layer signalling interactions will also be discussed in this section.

5.2 Level 1: A-SNN

The first signalling level pertains to applications and provides Application QoS-based Signaling and Service Negotiation (A-SSN). In order to establish, maintain and release sessions with the necessary QoS levels, applications must express their needs and interact with the communication system. The objectives of this interactions is to guarantee that the needs of the application will be fulfilled, and that it will adapt to network conditions under which the available resources don’t allow the requested level of QoS. Furthermore, the interaction between applications and the communication system will be responsible for releasing the resources when sessions end.

From the applications and communication system points of view, the main requirements of Level 1 signalling functions are:

- Identification of users, accounting and billing functions (AAA functions);
- Negotiation and definition of session characteristics between user terminals including the possibility of choice of QoS characteristics by the user;
- Prioritization of data-flows, which allows the communication system to favour higher importance data flows at the expense of lower priority data flows;

- Identification of QoS requirements for the communication system;
- Verification of the possibility to setup connections with the required (or available) quality requirements (by interacting with CAC functions);
- Set-up connections between user terminals with the required (or available) quality level (by interacting with RMs);
- Maintain the quality level during sessions or adapt to quality variation in the communication system (by interacting with RMs);
- Release communication system resources at the end of sessions;
- Provide information to users about the quality level and relevant session characteristics during session set-up, session life and at session tear-down.

Signalling interactions between applications and the communication system can be explicitly made in the control plane by using a signalling protocol like SIP [10], or they can be implicitly initiated by the inspection of the data path and detection of session activation and termination.

When explicit signalling is used, Level 1 signaling can be direct end-to-end between the applications or mediated by proxy entities as illustrated in Figure 1 (proxies can also be transparent, thus preserving the end-to-end signalling characteristic between applications).

If implicit signalling is used, the quality requirements of applications have to be defined a priori by some sort of mechanism (e.g. policy mechanisms defined at application and/or user level).

To enable interaction between Level 1 signalling functions and Resource Managers (RMs), vertical interactions are needed. Four different cases can occur as described below.

Explicit Level 1 signalling and EuQoS-aware applications: in this case, Level 1 signalling flows directly between applications, and a mechanism in the control plane of the EuQoS-aware end-system will intercept signalling information and interact directly with the RM in the Access Network (cross-layer signalling identified by the green arrows in Figure 5). These interactions can be bidirectional and the RM can also notify ESs of changes in network conditions.

Explicit Level 1 signalling with EuQoS-aware proxies: if EuQoS-aware application proxies are used, these entities can interact with RMs for CAC and trigger Level 2 signalling for resource reservation (cross-layer signalling identified by the pink arrows in Figure 5). These interactions can be bidirectional and the RM can also notify proxies of changes in network conditions.

Explicit Level 1 signalling without EuQoS-aware applications or proxies: in this case Level 1 signalling must be intercepted somewhere (e.g. at the Edge Routers of the Access Networks) and Level 1 signalling information must arrive at RMs in order to trigger CAC and resource reservation functions. Since RMs do not have direct access to network domains, the information obtained from the Level 1 signalling inspection must be conveyed through Resource Allocators (blue arrow in Figure 5). In this case the information provided to the user will be limited.

Implicit Level 1 signalling without EuQoS-aware applications or proxies: if no explicit signalling is used at the application layer some mechanisms are needed somewhere in the data path (e.g. at the Edge Router of the Access Network) in order to trigger CAC and resource reservation functions at RMs. Since RMs do not have direct

access to network domains, the information obtained from the Level 1 signalling inspection must be conveyed through Resource Allocators (blue arrow in Figure 1). In this case the information provided to the user will be limited.

Of all the cases described above, the use of Level 1 signalling with EuQoS-aware proxies (case 2) is the less intrusive. Case 1, with EuQoS-aware applications, is the one that will enable more functionality to be available at the ESs, at the expense of a high deployment cost. Cases 3 and 4 are needed for backward compatibility reasons.

As said before, SIP is the natural candidate to support Level 1 signalling in the EuQoS architecture. SDP will also be used to support session characterization. For the purpose of EuQoS, SIP QoS extensions will be used and other possible extensions can also be proposed to meet specific requirements. In the context of this document, the EuQoS SIP QoS extensions protocol is named EQ-SIP.

The vertical interactions between RMs and RAs, and between RAs and network equipment (e.g. network domains Edge Routers), can be supported by the COPS protocol, by SNMP or by any proprietary solution available inside network domains.

5.3 Level 2: RM-SSN

The Resource Managers Signaling and Service Negotiation (RM-SSN) is the most important and the most complex to be developed in the EuQoS architecture.

The main objective of Level 2 signaling functions is to support resource reservation and management along the data path across the various network domains. To enable domain administrative independence, Level 2 signalling functions will operate on a hop-by-hop basis, between Resource Managers (RMs) of adjacent administrative domains. The main requirements of Level 2 signalling functions are:

- Activation of SLSs with adjacent domains (chosen by the TERO function) in reaction to local application needs (triggered by Level 1 signalling);
- Renegotiation of SLSs with adjacent domains, in reaction to local application needs (triggered by Level 1 signalling) and quality level variations;
- Termination of SLSs with adjacent domains when the former are no longer needed (triggered by Level 1 signalling);
- Reception of SLS requests from adjacent domains;
- Verification of the availability of resources to support the requested SLAs with the specified (or available) quality requirements (by interacting with CAC functions);
- Activation of SLSs at the request of adjacent domains;
- Maintenance of SLSs with adjacent domains with the specified quality (and renegotiation of the SLAs when needed);
- Termination of SLSs with adjacent domains and release of local communication system resources;
- Support SLS Monitoring and Measurement (interface with MMFM function);
- Support SLS optimization (interface with TERO function).

To support inter-domain service negotiation, Level 2 signalling functions will perform hop-by-hop negotiation between RMs in the path starting in the local RM at the origin access network (Access Network 1 domain in Figure 5). Each RM interacts with the Traffic Engineering and Resource Optimization (TERO) function to find the

next hop domain, install provisional reservations if resources are available (using the CAC function that, in turn, will use cross-layer signalling identified by the blue arrow in Figure 5 to interact with the RA of the domain), and then will start signalling with the next domain's RM, which will repeat the process until the last RM in the path is reached (Access Network 2 domain in Figure 5). When the remote RM is reached, signalling is sent in the opposite direction to confirm the provisional reservations made in the downstream direction. Alternatively to the use of the TERO function to find next hops between domains, this can be done in a static way at each RM.

The choices to support Level 2 signalling are not as clear as for Level 1 signalling. The options are:

- Develop a specific, simple signalling solution for the EuQoS architecture from scratch;
- Use an existing solution (e.g. SIIBS) as starting point for the development of the RM signalling solution;
- Use the NSIS framework being developed at the IETF. NSIS provides path-decoupled signalling mechanisms between control entities (e.g., Bandwidth Brokers Resource Managers); despite its relatively immature state, the basic functionality is defined, which allows the development of a simplified version and the subsequent adoption in the EuQoS architecture; in the context of this document, this simplified version of NSIS developed for the purpose of EuQoS QoS extensions is named EQ-NSIS;
- Use the SLS negotiation mechanism developed in the scope of the European IST project Mescal [11];
- Use the COPS extension for SLS negotiation (COPS-SLS) [12].

Since this is a critical issue in the EuQoS architecture it is recommended that two alternative options from the above list are developed and evaluated.

Figure 6 illustrates option 3 (use of EQ-NSIS).

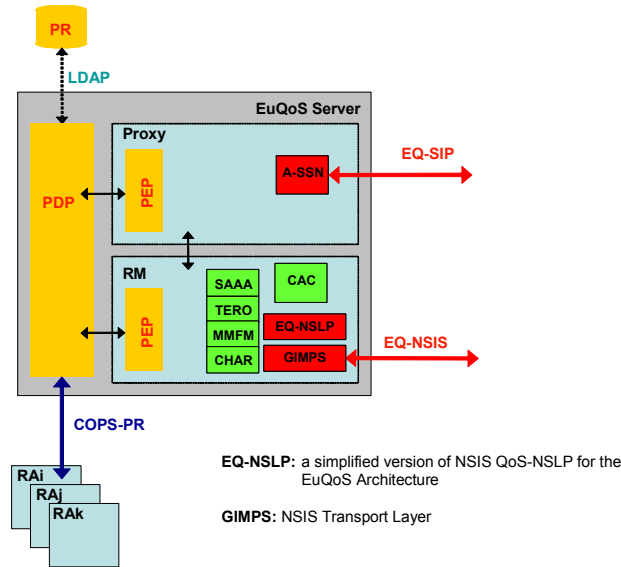


Fig. 6. EuQoS signaling at application and RM levels

5.4 Level 3: Inter-domain ND signaling

The third level of signaling in the EuQoS Signaling and Service Negotiation (SSN) Function is the network-technology-dependent hop-by-hop inter-domain signaling. At this level BGP (Border Gateway Protocol) is the most common solution currently used. BGP provides mechanisms for interdomain traffic routing and enables the use of routing policies to control the exchange of routing information between different administrative domains (Autonomous Systems).

The BGP version currently in use supports only Best Effort traffic, but extensions are being studied to deal with different traffic classes with different QoS requirements. The idea is to enable multiple paths to be created at interdomain level to support different traffic classes allowing traffic to be routed according to its QoS requirements. These extensions are normally named BGP+.

The use of BGP+ as inter-domain network-dependent hop-by-hop signaling mechanism (in the domains where it is available) can provide an extra level of QoS guarantees and enable high level RMs to rely on network level mechanisms for QoS support and network resilience.

The control of BGP+ routing policies can be made by RAs (and, indirectly, by RMs using cross-layer signaling interactions illustrated by the blue arrows in Figure 1) to route specific traffic classes via specific intermediate domains according to QoS and business requirements. Traffic Engineering and Resource Optimization (TERO) and Monitoring Measurement, Fault Management (MMFM) functions can provide useful information about the state of the network to support the policy control decisions.

Figure 7 illustrates the use of BGP+ inter-domain signaling, as well as intra-domain signaling interactions (discussed below).

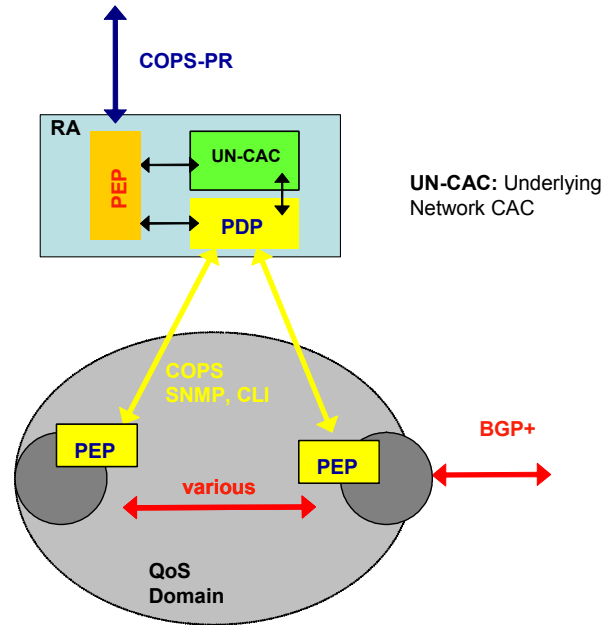


Fig. 7. Inter-domain ND hop-by-hop signaling in the EuQoS architecture

5.5 Level 4: Intra-domain ND signaling

The definition of signaling mechanisms to be used at intra-domain level in network-technology-dependent domains is also in the scope of the EuQoS project. The knowledge about specific network-dependent signaling mechanisms needs to be included in the domain's Resource Allocator (RA) in order to enable intra-domain resource management and QoS control.

Examples of intra-domain ND signaling that can be used by RAs to configure network domain's resources include RSVP and specific access network mechanism like ADSL bandwidth management, 3GPP, and Ethernet and WIFI priorities. NSIS path-coupled can also be used for this level of signaling.

When these signaling mechanisms are not available, the RAs must have a detailed knowledge of the domain configuration and topology so they can act directly in the configuration of each network element along the path subject to QoS requirements, using cross-layer signaling interactions illustrated by the yellow arrows in Figure 5.

Figure 8 presents the overall EuQoS signaling picture, according to the description presented above.

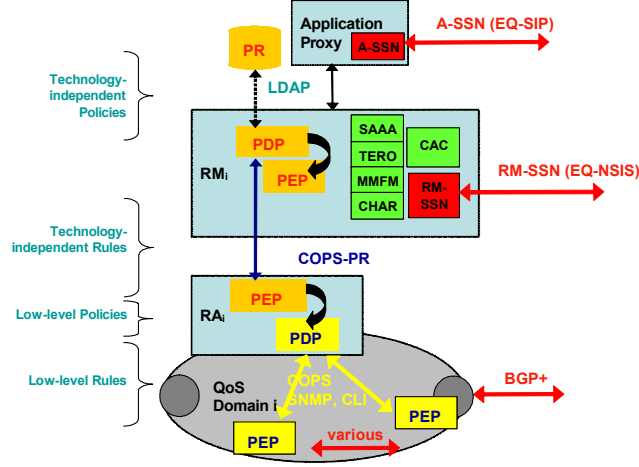


Fig. 8. Overall in-layer and cross-layer signaling protocols in the EuQoS architecture

6 Validation

The EuQoS system will be validated in two different ways: by simulation and by trial evaluation in wide test scenarios including different user communities, trial applications and access technologies. A separate EuQoS Work Package will deal with the overall validation of the QoS provided by the EuQoS system in heterogeneous network scenarios. The objectives of this WP are twofold. First, we evaluate the capabilities of the EuQoS system for delivering end-to-end QoS by using simulation tools (with main focus on these QoS mechanisms and algorithms that will be developed by the project). The second objective is to develop monitoring and measurement systems with test probes deployed in the whole EuQoS test bed. This system will be used in the trials.

In what concerns COPS and NSIS validation, the two parallel methods are also being conducted. Using the NS2 framework, a simulation model of NSIS is being developed at the University of Coimbra to integrate the EuQoS simulation package. In parallel the same project team is also developing one of the first NSIS implementations (GIMPS and QoS-NSLP). This implementation is being developed in Java and will be simplified in the first phase and focused in the EuQoS architecture (EQ-NSIS). In the second phase the implementation will include the complete feature set and will be made available to the Open Source community.

6.2 Simulation

In this section we present two of the studies that we made to evaluate the NSIS advantages in network scenarios, particularly by the use of the GIMPS layer specified in the NSIS framework, when it is necessary to implement a QoS signaling mechanism. For

this purpose, we have developed a set of new modules for NS-2: 802.1p/Q and a NSIS package.

A simple NSLP have also been deployed to work over Ethernet switch. Our NSLP interacts with the switch in order to configure prioritization levels, using the IEEE 802.1p and Q standards.

The next figure presents the Small Office Home Office (SOHO) scenario used in the first study.

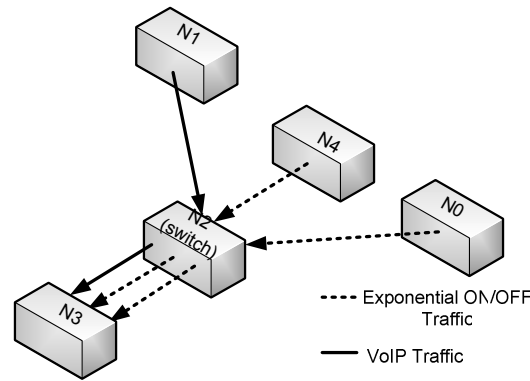


Fig. 9. Scenario #1

For foreground traffic we decided to use VoIP G.711 sources without Voice Activity Detection (VAD). The details about these source parameters are described in table 1.

Table 1. Source parameters

Traffic Type	Description
VoIP	G711 without VAD Packet size 200 B Inter-arrival 0.02 s
Traffic1	Packet Size 800 Burst_time 0.03 Idle_time 0.03 Rate 1000k Exponential traffic
Traffic2	PacketSize 200 Burst_time 0.02 Idle_time 0.02 Rate 600k Exponential traffic

In the scenario #1, a NSIS agent was configured in each node. The purpose of this study was to evaluate the NSIS effect on the traffic and on the congestion level. In

these simulations, the traffic between the nodes N1 and N3 was modeled as G.711 VoIP, requiring a high prioritization level. The background traffic (from N0 and N4) was modeled as exponential ON/OFF sources in order to overload the links.

Figure 10 shows the one way delay (OWD) experienced by each of the traffic flows, when no prioritization was performed in the switch. In this case, the switch had just two queues. As expected, all the traffic presented similar behavior, experiencing high delays in the face of network congestion.

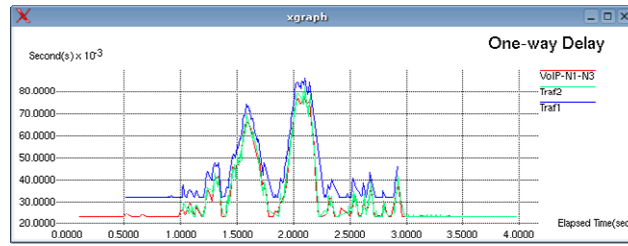


Fig. 10. Packet delay without traffic prioritization

Figure 11 shows the OWD when prioritization was introduced in the switch, activated by a NSIS signaling message at 1.3 seconds. The switch was initially configured with just 2 queues, and no prioritization mechanism was performed. In the elapsed time 1.3 sec, the NSIS agent, located in N1, sent a path-message to node N3. The payload of this NSIS message contained the command to reconfigure the switch (N1 =>N2=>N3) with 7 virtual queues. Each NSIS agent initialized the peer discovery negotiations while NSIS (NSLP function) reconfigured each switch along the data path. After that, the prioritization traffic was set between N1 and N3. As can be observed in figure 11 the VoIP delay remained low during all the simulation.

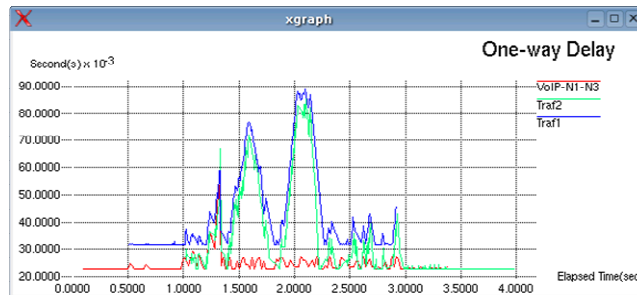


Fig. 11. Packet delay with traffic prioritization after the NSIS Message

In the scenario of figure 12, we simulated three NSIS sessions for three different VoIP flows. Each session was initialized at different elapsed moments, in order to observe the interaction of the NSIS agents with the switches and nodes. The traffic foreground was modeled as VoIP G.711 sources. Three sources were added from N1 to N10, N8 to N11 and N9 to N5. Two background flows were initialized, in order to

overload the network. In figure 13, we can observe OWD of this scenario without any prioritization.

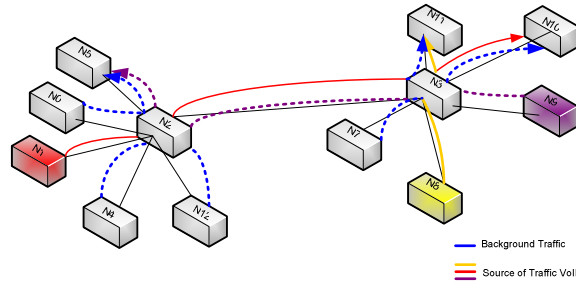


Fig. 12. Scenario #2

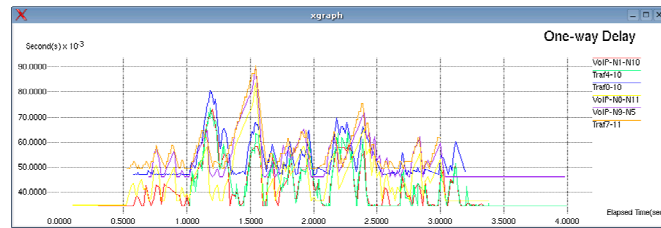


Fig. 13. Scenario #2 Ethernet without prioritization

A NSIS session was setup between N1 and N10 at 1.3s. A path message was sent from the node N1 to N10, in order to reconfigure the switch along the data path with a higher prioritization level. As a consequence, it can be observed, that the corresponding OWD was not affected (Figure 14). However, the others VoIP flows remained with critical OWDs until their corresponding NSIS sessions were initiated.

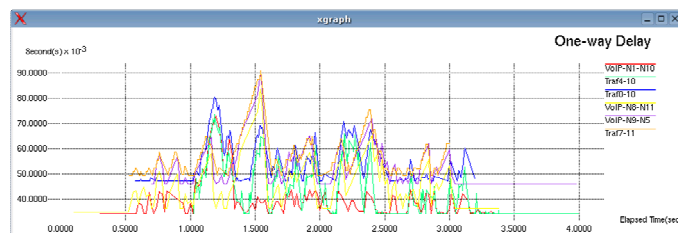


Fig. 14. Ethernet with traffic prioritization for the VoIP traffic VoIP-N1-N10

The NSIS agent in N8 sent a path message to N11 at 2s, and, at the same time, the agent NSIS in N9 sent a similar path message.

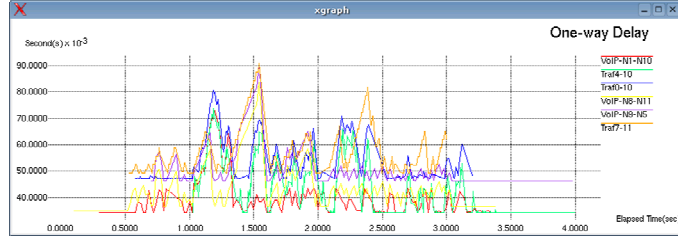


Fig. 15. OWD remain low for the VoIP-N8-N11 and VoIP-N9-N5 after prioritization traffics in elapse time 2s

These simulation tests demonstrated the NSIS advantages when it is necessary to implement different classes of traffic. The NSIS agents, along the simulation studies, reconfigured the switches with prioritization levels, controlling the OWD of the foreground flows. Different NSIS mechanisms can be built, extending the NSLP functionalities used in these examples.

7 Conclusion

This paper discussed the use of the Policy Based Network paradigm, COPS (Common Open Policy Service) and NSIS (Next Steps In Signaling) frameworks being developed by IETF, in the context of the European Project EuQoS (End-to-End Quality of Service over Heterogeneous Networks).

The analysis of the use of PBN/COPS and NSIS to support the Signaling and Service Negotiation (SSN) function revealed that these protocols can be used to provide a basis to support the desired functionality and a framework for further enhancements.

The implementation is being carried out in two parallel tracks: simulations and prototype that will support the validation tasks planned in the EuQoS project. Preliminary simulation results show that the integration of NSIS and COPS enable control of the underlying network devices and QoS provision to applications.

Acknowledgements

This work was partially funded by the European Union 6th Framework Programme under contract IST FP6 IP 004503 EuQoS Integrated Project. The authors also acknowledge the various comments and valuable suggestions received from the EuQoS project team.

References

1. EuQoS Project consortium, Integrated Project EuQoS End-to-end Quality of Service support over heterogeneous networks Annex I - Description of Work, May 2004.

2. The GÉANT project – <http://www.geant.net>
3. IETF Policy Framework Working Group (IETF-PG) – <http://www.ietf.org/html.charters/policy-charter.html>
4. Durham, Ed., J. Boyle, R. Cohen, S. Herzog, R. Rajan, A. Sastry, The COPS (Common Open Policy Service) Protocol, RFC 2748, IETF, January 2000.
5. R. Hancock G. Karagiannis J. Loughney S. van den Bosch, Next Steps in Signaling: Framework, Internet-Draft, IETF, 2005.
6. K. Chan, J. Seligson, D. Durham, S. Gai, K. McCloghrie, S. Herzog, F. Reichmeyer, R. Yavatkar, A. Smith, COPS Usage for Policy Provisioning (COPS-PR), RFC 3084, IETF, March 2001.
7. Schulzrinne, R. Hancock, GIMPS: General Internet Messaging Protocol for Signaling, Internet-Draft, IETF, 2004.
8. S. Van den Bosch, G. Karagiannis, A. McDonald, NSLP for Quality-of-Service signaling, Internet-Draft, IETF, 2004.
9. M. Stiernerling H. Tschofenig M. Martin, C. Aoun, NAT/Firewall NSIS Signaling Layer Protocol (NSLP), Internet-Draft, IETF, 2004.
10. Handley, H. Schulzrinne, E. Schooler, J. Rosenberg, SIP: Session Initiation Protocol, MRFC 2543, March 1999
11. MESCAL Project, Management of End-to-end Quality of Service Across the Internet at Large - <http://www.mescal.org/>
12. T.M.T. Nguyen, G. Pujolle, N. Boukhatem, COPS Usage for SLS negotiation (COPS-SLS), Internet-Draft, IETF, June 2001.