Roberto Mameli

The Strong Need for Policy in the QoS Environment

Quality of service (QoS) refers to the 'classification of packets for the purpose of treating certain classes or flows of packets in a particular way compared to other packets'. It is the idea that transmission rates, error rates, and other characteristics can be measured, improved, and, to some extent, guaranteed in advance. Ideally, it means making the data delivery service of the otherwise unpredictable best-effort Internet protocol (IP) network, predictable.
The QoS requires the use of protocols such as RSVP, that affects the allocation of network resources on a per-flow basis, and DiffServ for differentiated services, that provides marks within the IP packet headers to allow prioritisation of traffic aggregates.

Since enabling QoS on an IP network effectively means that some users will get better network service than others, it creates some incentive to steal. Some users inevitably want the better service, but they do not want to have to pay the (likely) higher costs involved. Hence, there is a need to authenticate those that request the better service levels, and to do so requires verifying the identity of traffic 'owners' on a per-packet basis.

Since there are varying circumstances in which traffic owners—be they end-users, applications, Internet hosts, company entities or whatever—are entitled to the services they request, there is a need for rules, a need for 'police' to enforce these rules, and a need for 'judges' to decide when they apply. The rules, the police and the judges all comprise a policy system that is an essential component of a QoS-enabled network.

This paper will show policy solutions applied to some generally accepted QoS scenarios.

Introduction

Today the common opinion of analysts is that the telecommunication and networking worlds are converging together. This phenomenon is justified by several technological and commercial reasons. The Internet Protocol (IP) is playing a key role in this process. Designed a couple of decades ago, it became the universal infrastructure for data communications, taking the place of more traditional solutions, such as X.25 or asynchronous transfer mode (ATM). The reason for this huge success is very simple. Originally designed for academic and public use, IP was developed to be simple, flexible, scalable and efficient. It was based on the 'end-to-end' principle, where the network performs minimal operations and the 'intelligence' is concentrated on terminals. This means that the network simply provides delivery of data packets, without any guarantees, while complex tasks—for example, retransmission, reliability mechanisms and congestion control—are performed by source and destination hosts. This model is usually referred to as *best-effort* service.

Due to its explosive success, IP is now gradually becoming the transport infrastructure used in converged networks for both voice and data. However, the same reasons that contributed to its success are now showing their limitations. In fact, the best-effort service offered by IP is considerably poorer in terms of quality of service (QoS) than the service offered by traditional networks designed for real-time traffic; for example, the telephone network. As a consequence, from the point of view of applications, QoS has become a key issue. New applications, such as voice over IP (VoIP), video streaming, conference calling and other multimedia applications, heavily depend on the QoS provided by the underlying network.

At the same time, there are many actors attracted by the use of IP for new services. Among them Internet service providers (ISPs) that aim to provide value-added services to differentiate from competitors, emerging operators that try to increase their market penetration by offering multimedia and other specialised services at low prices, and finally traditional telcos that are obviously interested in new technologies. The efforts and the investments they spend to provide the end-user with the desired level of QoS must be returned by adequate profits. The availability of better services is an incentive for malicious users to steal. As a consequence there is a need for rules, a need for 'police' to enforce these rules, and a need for 'judges' to decide when they apply. The

Roberto Mameli:
Ericsson Telecomunicazioni S.p.A.
Tel: +39 06 20410038
Fax: +39 06 20410037
E-mail: mameli@coritel.it

rules, the police and the judges all comprise a policy system that is an essential component of a QoS-enabled network.

This paper is organised as follows: the next section illustrates the main concepts of a policy-based architecture and explains the related terminology. The following section presents some QoS scenarios in which such an architecture could be applied and finally conclusions are drawn.

Policy Actors and Terminology

In a broad sense, *policy* refers to the aim of aligning business goals with network resources. A complete policy system allows the creation of rules that specify how traffic will be dealt within the network. These rules are usually decided by the network administrator with various purposes; for example, to select which users have access to which network resources, to prioritise critical applications or to deliver service differentiation upon user needs. They are enforced by network elements in order to provide some assurances that if network resources, such as bandwidth, become scarce, then the key applications will experience the least disruption.

In a standard policy-based network, policy rules consist of two components:

- A set of *conditions* under which the policy apply. This might include parameters such as user name, addresses, protocols and application types.
- A set of *actions* that apply as a consequence of satisfying (or not satisfying) the conditions including bandwidth guarantees, access control, service load balancing and so on.

Note that these rules do not specify how to achieve the desired objective, but rather what should be the desired outcome. A comprehensive policy-based system allows the network manager to define in a human readable fashion policy rules that will automatically take effect on specific equipment in the network environment. The set of architectural components needed to accept input in the form of high-level policy rules and translate them into specific configuration commands to individual network devices is usually referred to as policy-based networking (PBN) architecture.

A simple model of PBN architecture includes a *policy manager*, which is the central policy administration and directory repository point, and a *policy enforcer*, which consists of remote active management components that make up the local policy decision and enforcement points throughout the local wide area networks. Figure 1 shows the main components of a PBN architecture.

Figure 1 Elements of PBN architecture

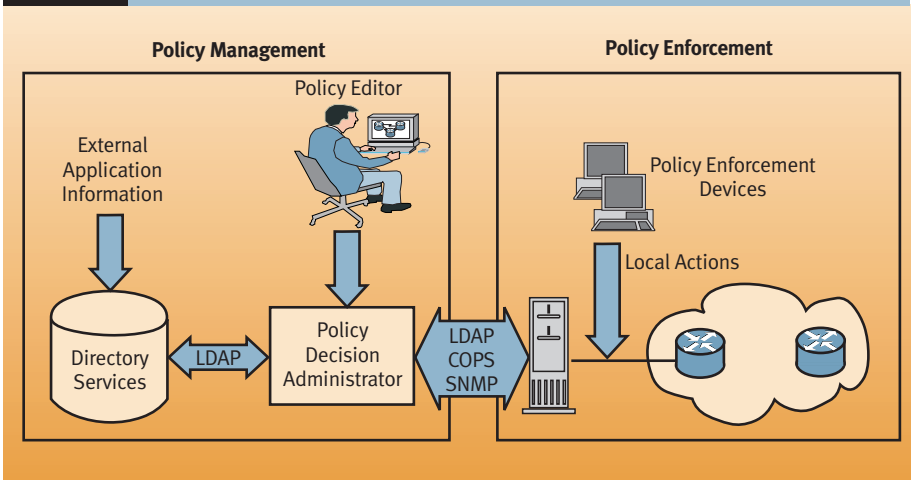

The policy management functions are responsible for coordinating administrator input with other external and internal policy information and translating this information into network terminology. The policy manager consists of various logical pieces: the policy editor, the interface to the directory database and the policy decision administrator:

- *Policy editor*: Provides the network manager with the capability of configuring rule-based policies in a centralised way. The policy editor can be accessed directly by network managers via a Java/web interface or can be part of a centralised management structure. The editor usually stores information in an lightweight directory access protocol (LDAP) based database directory.
- *Interface to the directory database*: The directory database server provides the central repository for policy information. It is responsible for the storage of a wide range of information including user login and network specific policy information. The directory can be grouped in a distributed and hierarchical fashion for sharing data. It will interface to the policy decision administrator in order to create higher-level dynamic policies. An administrator will be able to define a customer as *gold level* and then setup a specific policy that says that all gold-level customers will receive a specific priority throughout the network.
- *Policy decision administrator*: coordinates policies in the network. It performs many functions including creating policies based on static information in the directory and other databases as well as transient information based on the status of the network. This information is then relayed to various policy enforcement devices for specific network actions.

The policy enforcement functions are performed by network elements in response to policy management decisions. The policy enforcer could be a simple router that applies actions based on a field in the packet or, alternatively, it might reside in a particular piece of equipment that locally analyses traffic flows and network conditions. Typically, the policy enforcer is placed at the edge of a policy domain. It communicates with the policy manager to understand overall network policies and then performs various enforcement functions including traffic shaping/conditioning, policing and signal provisioning.

An issue strictly related to end-to-end QoS provisioning throughout the Internet is the concept of peering contracts between neighboring domains, also called *service level agreements* (SLAs). Bandwidth brokering is one of the mechanisms for implementing bilateral SLAs. As its name suggest, the bandwidth broker (BB) is responsible of allocating network resources on demand. More specifically, the BB performs both intra-domain and inter-domain functionality. Within the boundary of a specific policy domain, the BB manages resources on links and devices. Intra-domain functions may include monitoring resources, and gathering routing and topology information, useful for policy-based and resource-based admission control. In contrast, inter-domain functionality consist mainly in the communication with peering BBs for purposes of coordinating service level agreements and bandwidth

brokering between the different domains (see Figure 2).

Figure 2 Policy enforcers and bandwidth brokers

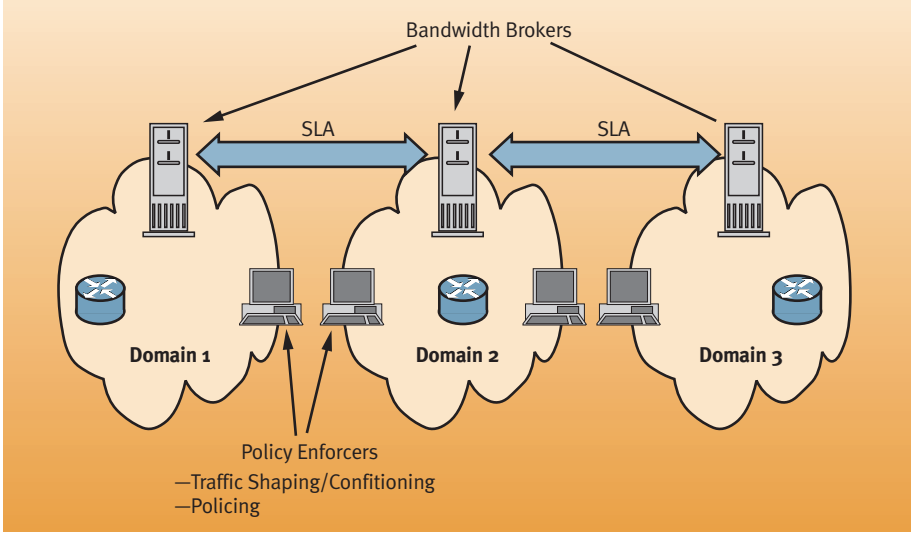

QoS and Policy Scenarios

The first step in the deployment of policy-based QoS systems will probably consist in the provision of the so-called *Olympic service* over enterprise backbones or ISP networks. In this model, the network is provisioned to support multiple levels of service, in descending ranking: Gold, Silver, and Bronze. In general terms, the higher classes are assumed to have lower delay, and lower drop probability with respect to lower classes. It is relatively easy to deploy these services onto networks supporting either differentiated services or legacy IP precedence. Two independent mechanisms are needed to determine the transmission service. First, traffic is classified at the edges and shaped (rate controlled); its header is marked with the appropriate precedence value. Second, core nodes are configured to provide a certain class of service to packets marked by a certain priority value. The first mechanism (edge) is rather dynamic and can be changed frequently. The second mechanism (core configuration) is relatively heavy, and should be changed rather infrequently.

An example of such a service could be implemented in a differentiated services network by means of three out of the four assured forwarding (AF) classes. The following set of configurations and actions must be installed to enable the Olympic service.

- DiffServ routers must be configured to allocate a certain bandwidth over every appropriate link for each class;
- edge interfaces must be configured to classify inbound traffic into one of these three classes (based on SLA and policy); and
- packets must be marked as a combination of their class and rate (whether they are in-profile).

Figure 3 depicts the structure of a policy-based QoS network for Olympic service provisioning. The network structure comprises a centralised QoS policy control—that is, a single policy administrator—and a fully distributed execution by means of multiple policy servers. This architecture has excellent scaling properties, and can virtually scale to any size network.

The next phase in the realisation of a complete policy-based QoS system would be the possibility for the customer to obtain dynamic on-demand QoS allocation. In fact, the service described above is incapable of handling congestion that arises when multiple sources compete within a single class of service. Furthermore, it lacks the ability to adapt to the changing status of the network, or perform automated capacity planning and provisioning. The policy system controls the traffic mapped to these Olympic classes by controlling the edge marking. This control is critical to assure that classes are not over-populated, and the higher-priority services actually get better service. To achieve this, several prominent researchers in the field introduced a process of negotiations and admission control by means of the previously mentioned bandwidth broker. This concept is developed as part of the Internet2 programme (and the QBone project specifically). From the customer perspective, the ability to request bandwidth and QoS on demand allows them to address elastic usage in their networks. Rather than overpay some of the time, and be short on resources other times, bandwidth on-demand provides customers with as much resources as they need, when they need it, and at a reasonably economical price.

There are many possible scenarios in which a complete policy-based system as described above could be employed. This section considers a couple of examples, but it is worth observing that many others are likely. The first refers to a virtual private network (VPN) environment, while the second foresees a voice-over-IP (VoIP) application.

Figure 3 Policy-based QoS system

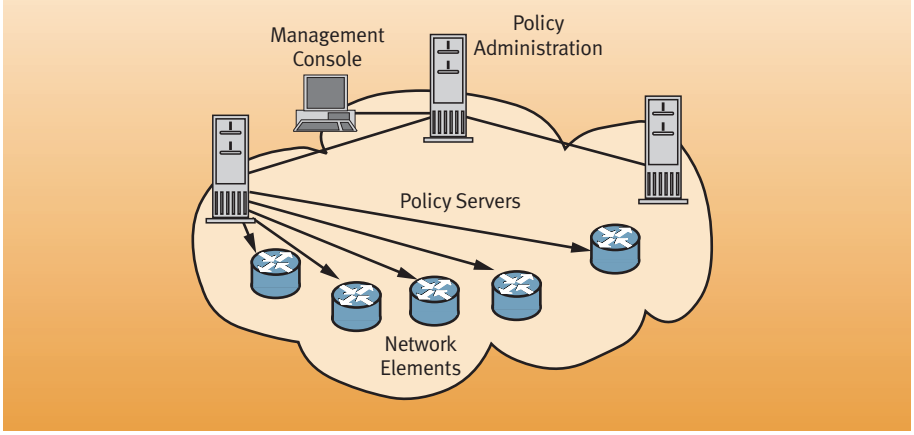

Virtual private network scenario

As its name suggests, a virtual private network (VPN) is a 'private' network that relies on a public shared transport infrastructure. For example, a VPN could be used by an enterprise to connect the headquarters to remote sites avoiding the expense of a leased line. VPN security mechanisms protect the confidentiality of information as it flows over the public network. A VPN basic function is to provide cost-effective secure Internet communications for the corporate enterprise. Due to its ubiquity, IP is the best candidate as a viable, cost-effective alternative to dedicated lines for implementing VPNs. Corporations that use VPNs become more efficient organisations, strengthened by

their increased productivity and reduced operating expenses. On the other side, Internet service providers (ISPs) can capitalise on the benefits that VPN technology affords the corporate enterprise, by offering secure outsourced VPN services to corporations. The role of end-to-end quality of service, and especially of policy-based QoS solutions, is very important, since it enables organisations to implement an intelligent, high-performance VPN customised to meet the specific information needs of their business. The result is a business-oriented network strategically designed to speed information flow throughout the network based on criteria defined by the organisation. Figure 4 shows a basic VPN scenario without QoS support.

Figure 4 Basic VPN configuration without QoS support

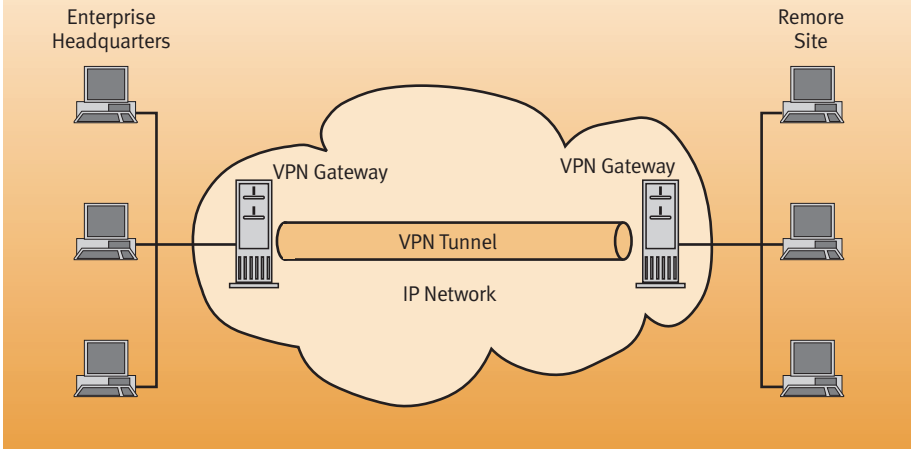

As confidence in VPN security grows, organisations are increasingly taking advantage of low-cost Internet services for sending business information over the enterprise wide are network (WAN). This information includes ordinary e-mail messages, mission-critical financial data, as well as bandwidth-intensive web graphics, IP telephony applications, and real-time multimedia. When an organisation's mission-critical and time-sensitive business information is travelling over the VPN, IP best-effort service is just not good enough. Although businesses can significantly reduce their costs by implementing VPN technology, to optimise VPN performance, the enterprise needs more control over how the network handles different types of traffic. QoS technology provides these capabilities by enabling organisations to prioritise traffic flow and to guarantee bandwidth on the VPN. To improve VPN performance, QoS solutions let network administrators use business-based criteria to prioritise network traffic. Packets of information are prioritised based on who in the organisation sent the packet (the president or a staff engineer), the type of business the packet contains (quarterly financial results or ordinary e-mail), or the time of day the packet is sent (noon time or 2:00 am). The assigned QoS determines the order in which different traffic types are forwarded over the VPN. By assigning a high-priority QoS to SAP R/3 traffic, for example, the network administrator can ensure that SAP data is forwarded over the VPN before ordinary e-mail and other low-priority traffic. If congestion occurs, the low-priority traffic is delayed to ensure that the high-priority SAP R/3 traffic receives expedited service. And if congestion causes packets to be dropped, the low-priority packets are dropped first. For time-critical VPN traffic that cannot tolerate network latency, more sophisticated QoS solutions are needed. In addition to simply prioritising traffic, QoS solutions offer policy-based dynamic bandwidth management capabilities, which can automatically reserve bandwidth completely through the network to ensure an application receives a steady rate of throughput. For example, the network administrator can define a QoS policy that automatically allocates 256 kbit/s for the weekly executive videoconference that uses VPN technology to link a remote branch office with the corporate headquarters. By reserving bandwidth, QoS technology guarantees the smooth delivery of real-time applications, minimising gaps, out-of-sequence data, and other problems that could affect transmission quality. Adding QoS traffic management features to the standard tunnelling capabilities of a VPN can significantly enhance performance by ensuring that traffic flows through the network according to business-driven priorities rather than the random, best-effort service provided by the underlying IP network infrastructure.

Voice-over-IP scenario

Another interesting scenario of applicability of the policy-based QoS architecture concerns emerging real-time and multimedia services and applications over IP. Among them a key technology is certainly represented by voice over IP (VoIP), which is gradually becoming a valuable solution for delivering cost-effective and business-grade voice applications. The issues behind this growth include new voice-compression techniques, VoIP gateways that compress and packetise voice traffic for transmission over IP networks, and end-to-end quality of service (QoS) for dynamic allocation of bandwidth and service guarantees. VoIP works by transmitting voice traffic over a packet-based corporate intranet or the global Internet rather than using the higher-priced circuit-switched public network. A typical enterprise VoIP solution includes VoIP gateways that enable long-distance calls between international corporate sites to bypass the PSTN and, thus, avoid the associated telephone-company charges. With this type of application, a VoIP gateway is installed between the PBX at each branch site and the router at the edge of the branch LAN (see Figure 5).

The PBX at the calling site switches the long-distance calls to the VoIP gateway,

Figure 5 VoIP configuration

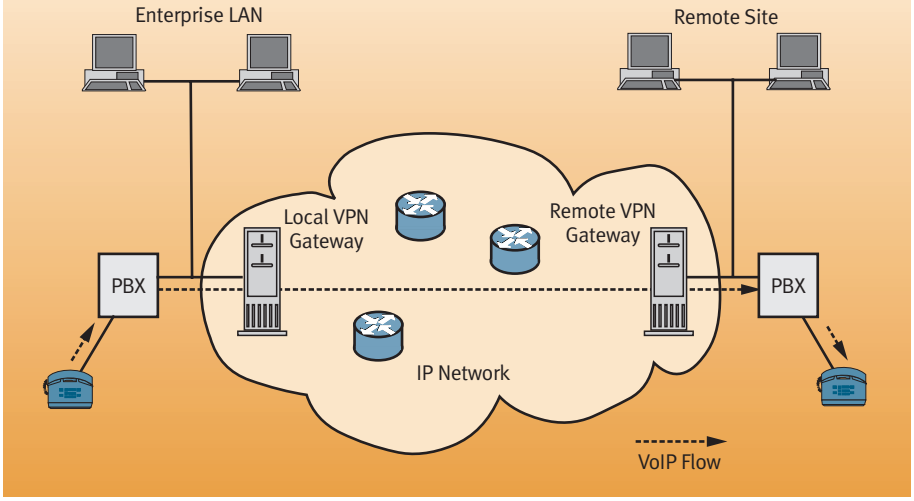

which in turn compresses and packetises the voice traffic, determines the IP address of the VoIP gateway on the remote branch and sends the VoIP packets to the edge router for transmission onto the IP network. On the other side, the destination VoIP gateway receives the voice traffic, depacketises and decompresses it, then send it to the PBX to be switched onto the local branch telephone lines. The entire process is transparent to the telephone call participants. New compression algorithms have significantly improved VoIP application performance. By reducing voice bandwidth requirements to under 5 kbit/s, compression techniques have increased the voice-carrying capacity of an IP network of a magnitude order. However, the biggest problem for VoIP application performance remains the inevitable delay associated with the connectionless packet-based IP network, which implies degradation of the audio quality. Moreover, network congestion causes inconsistent voice quality or *jitter* during peak traffic loads. When network latency exceeds a few hundreds milliseconds, audio quality becomes unacceptable. Since IP best-effort service cannot differentiate between time-sensitive voice traffic and less-time-critical traffic, such as e-mail messages, VoIP applications deliver poor voice quality when latency reaches high levels. Fortunately, emerging QoS solutions now provide the ability to deliver more predictable, reliable, and guaranteed audio-transmission services on IP networks. By assigning a high priority to VoIP traffic, the network administrator ensures that VoIP data is forwarded before ordinary e-mail and other low-priority traffic. If congestion occurs, the low-priority traffic is delayed to ensure that the high-priority VoIP traffic receives expedited service. And if congestion causes packets to be dropped, the low-priority packets are dropped first. IP precedence can be used to expedite traffic across the Internet backbone, where it is difficult to guarantee bandwidth without establishing a virtual private network managed by a single ISP or corporation. But, time-sensitive two-way voice applications generally require more sophisticated QoS solutions, which supports policy-based dynamic bandwidth management capabilities that automatically reserve bandwidth completely through the network to ensure the VoIP application receives a steady rate of throughput. For example, the network administrator can define a policy that automatically allocates 10 kbit/s for international long-distance VoIP calls between the remote sales office and the corporate headquarters. By reserving bandwidth end-to-end across the entire enterprise network, QoS technology can guarantee the smooth delivery of real-time VoIP applications, minimising gaps, out-of-sequence data, and other problems that could impact audio transmission quality.

Conclusions

In recent years we have been observing a gradual process of convergence between the telecomunications and datacommunications realities. Due to its characteristics and its ubiquity, IP certainly represents a key actor in this scenario. However, the classical best-effort service offered by IP is not well suited to emerging real-time and multimedia applications. This observation justifies the effort spent in the development of QoS solutions. On the other hand, the ability to provide better service to some users in spite of others leads to the introduction of policybased architectures. The purpose of such architectures is to find an elegant solution to the problem of congested networks. From the point of view of both enterprises and service providers they offer a sophisticated and cost-effective way to tune their networks to the needs of their business. At the same time, they allow the end-user to benefit from all the advantages related to QoS support. As a consequence, the interest in policy-based architectures certainly represent a key issue in the development of QoS over IP technologies for future multimedia applications.

Bibliography

1 YAVATKAR, R.; PENDARAKIS, D.; and GUERIN, R., A Framework for Policy Based Admission Control. IETF RFC 2753, Jan. 2000.

2 BLAKE, S.; BLACK, D.; CARLSON, M.; DAVIES, E.; WANG, Z.; and WEISS, W. An Architecture for Differentiated Services. IETF RFC 2475, Dec. 1998.

3 NICHOLS, K.; JACOBSON, V.; and ZHANG, L. A Two-bit Differentiated Services Architecture for the Internet. RFC 2638, July 1999.

4 NEILSON, R.; J. WHEELER, J.; REICHMEYER, F.; and HARES, S. (Editors) A Discussion of Bandwidth Broker Requirements for Internet2 Qbone Deployment. Version 0.7.

5 DURHAM, D., Ed.; BOYLE, J.; COHEN, R.; HERZOG, S.; RAJAN, R.; and SASTRY, A. The COPS (Common Open Policy Service) Protocol. IETF RFC 2748, Jan. 2000.

6 BLAKE, S.; BLACK, D.; CARLSON, M.; DAVIES, E.; WANG, Z.; and WEISS, W. An Architecture for Differentiated Services. IETF RFC 2475, Dec. 1998.

Glossary

AF Assured forwarding
BB Bandwidth broker
COPS Common object policy server
EF Expedited forwarding
IP Internet protocol
ISP Internet service provider
LAN Local area network
LDAP Lightweight directory access protocol
PBN Policy-based networking
PBX Private branch exchange
QoS Quality of service
RSVP Reservation protocol
SLA Service level agreement
SNMP Simple network management protocol
VoIP Voice over IP
VPN Virtual private network

Biography

Roberto Mameli
Ericsson
Telecomunicazioni S.p.A.

Roberto Mameli received his degree in Telecommunication Engineering from the University of Rome 'La Sapienza' in 1997. Since September 1998 he has worked in CoRiTeL, where he takes part in the research activity of the SOFIA group (SOlution for Full IP Access). Within this project he is mainly interested in the investigation of problems related to QoS and aspects concerning wireless access techniques and mobile IP issues.