Advanced QoS Provisioning in IP Networks: The European Premium IP Projects

Silvia Giordano, LCA — EPFL
Stefano Salsano, DIE — University of Rome "Tor Vergata
Steven Van den Berghe, IMEC
Giorgio Ventre, University of Naples Federico II
Dimitrios Giannakopoulos, National Technical University of Athens

ABSTRACT

This article describes the current evolution of QoS architectures, mechanisms, and protocols in the Internet, as it is ongoing in the framework of the European Union funded research projects on premium IP networks. A short review of the proposed standard approaches to QoS (e.g., differentiated services, integrated services, and label switching technologies) is given. Then we focus on the state-of-the-art architectures, mainly based on DiffServ concepts. Several issues arise when trying to implement these architectures in the real world: QoS aspects, network monitoring of the offered QoS, and end-user control of received QoS. The article then discusses the existing results and the current direction of European research and development in these areas.

INTRODUCTION

The traffic generated by Web-based and multimedia applications produce a great amount of burstiness, which is difficult to describe by a static set of traffic parameters. The dynamic and efficient usage of resources is one of the fundamental aspects of multimedia networks: the traffic specification should first reflect the real traffic demand, but optimize, at the same time, the resources assigned to data flows.

The straightforward and simplest answer to the applications' traffic demand is to add resources (i.e., capacity) on the links. This is known as *overprovisioning* and it represents an alternative to the IP quality of service (QoS) architecture and mechanisms we will discuss in this article. However, there are some shortcomings of the overprovisioning approach that justify the research for more "complex" QoS mechanisms. First of all, the current trend is that link technologies used in IP networks are becoming more heterogeneous, ranging from fiber optic backbone links to various kinds of wired and wireless link technologies at the edges, with these latter characterized by different transport capacity. In addition, the user community is looking forward to getting both application and service heterogeneity (linking service expectations and willingness to pay for the service). All these trends point toward the Internet becoming a ubiquitous multiservice network, where different actors (network operators, service providers, content suppliers) will compete and interoperate to offer novel applications that are dynamically created and managed. The presence of different actors of course makes it more difficult to find a common strategy for engineering the capacity on the links. Finally, a fast overprovisioning solution is not always available for technical and economical motivations. For that reason, overprovisioning per se cannot really be an effective solution on an end-to-end basis. Consequently, there are strong commercial reasons for network operators and equipment providers to offer QoS differentiation in IP networks [1].

Currently, some network operators are actively designing a new service model that includes flexible service creation. Manufacturers of network equipment continually introduce new solutions and products characterized by conformance to existing or proposed standards and recommendations concerning micro-QoS issues, but at the same time adopting specific and proprietary solutions for provisioning macro-QoS functionality. Thus, the network operators are willing to open up their network resources to innovative new service providers, which must include mechanisms for supporting end-to-end QoS guarantees (across multiple domains), and for the flexible and dynamic creation of new services.

For the last 20 years QoS has been one of the major research topics in the networking community. First in academia, then in industry, the issues related to the provision of guarantees in the performance achievable in the provisioning of communication services has been subject to an intense debate that is still continuing in the various fora (e.g., the Internet Engineering Task Force, IETF). In particular, we will present the ongoing debate and the most recent developments originating from a very specific research community, that active around the so-called premium IP network projects. These projects (funded by the European Commission in the framework of the IST research program [2]) are involving a quite large number of research and industry partners around the common problem of paving the way to the effective deployment of QoS-capable IP networks.

In this article, we first give a short review of the proposed standard approaches to QoS: differentiated services (DiffServ), integrated services (IntServ), and label switching technologies (see also [3]). Then we focus on state-of-the-art architectures, mainly based on DiffServ concepts. Several issues arise when trying to implement these architectures in the real world: definition, dynamic creation and configuration of such services, traffic engineering tools to obtain quantitative end-to-end QoS guarantees, dynamic service invocation, interdomain QoS aspects, network monitoring of the offered QoS, and end-user control of received QoS. Finally, the article discusses the existing results and the current direction of research and development currently underway in Europe in these areas, and these will be introduced in detail in the accompanying articles composing this Feature Topic.

QoS Architectures, Mechanisms, and Protocols in the Internet

As an alternative or additive strategy to overprovisioning to accommodate the requirements arising from current services and those expected to arrive in the near future, a network and service management architecture will be needed to optimize resource allocation in networks. The term *QoS provisioning* usually indicates the set of technologies for managing delay, jitter, and congestion events throughout a network via traffic policing and resource control. QoS technologies are generally handled by the network administrators and comprise both traffic handling mechanisms, and provisioning and configuration mechanisms [4]. In the research and standardization bodies the development of QoS support has experienced a significant evolution during these years. However, the evolution has mainly involved the backbone. To realize true QoS, its architecture must be applied end to end, not just at the edge or at selected network devices.

RSVP, DiffServ and MPLS

The first recognized tool for QoS provisioning is Resource Reservation Protocol (RSVP). RSVP is the signaling protocol defined in IETF that can be used by routers supporting IntServ along the path to set up per-flow QoS. According to the taxonomy presented in [4], RSVP can be considered a mechanism for configuring traffic handling mechanisms in network devices for per-flow-based traffic. The IntServ architecture foresees per-flow traffic classification and queue servicing algorithms for handling the traffic of each flow. Each router along the path performs per-flow admission control and then guarantees the service to each traffic flow in strict isolation from other traffic flows. This model was considered insecure, too complex, and not scalable enough to be used in the backbone of the Internet.

The DiffServ architecture was designed in IETF to overcome the scalability problems of per-flow QoS management in routers. According to the taxonomy presented in [4], DiffServ is also a traffic handling mechanism, but in this case it works with aggregate traffic. The objective is to provide scalable QoS support by avoiding per-flow state in routers. The basic idea is that packet headers include a field called the DiffServ codepoint (DSCP). The DSCP allows packets to be classified and identifies the specific queuing behavior (per-hop behavior, PHB). Edge routers are configured with a large number of per-flow policing policies, while core routers are configured with a few forwarding classes. The traffic from many flows having similar QoS requirements is marked with the same DSCP, so the flows are aggregated to a common queue or scheduling behavior. The architecture relies on packet markers and policing functions at the boundaries of the network to ensure that the intended services are provided.

Orthogonal to these node-level congestion avoidance mechanisms, path-level avoidance can bypass the shortest path routing for a part of the end-to-end traffic in order to obtain a more balanced network load. Multiprotocol label switching (MPLS) provides a mechanism for engineering network traffic patterns that is not subject to the limitations of different routing protocols, transport layers, and addressing schemes. MPLS was also designed in IETF, with the objective of increasing the efficiency of data throughput by optimizing packet processing overhead in the IP network. For this purpose, MPLS assigns short a layer 2 label (applied to the IP frame) to network packets that describe how to forward them through the network. This label corresponds to an established (configured/signaled) path through the network. The analysis of the IP packet header is performed just once, when a packet enters the MPLS cloud. Therefore, MPLS makes it possible to switch traffic through IP routers that, historically, had to interrogate each IP header before forwarding to the next hop. Additionally, at the time a label is applied to the flow, predefined traffic engineering parameters can be programmed into the forwarding hardware to guarantee levels of traffic bandwidth, delay variation, and congestion control. Alternatively, a DSCP can be encoded in the MPLS header allowing the DiffServ mechanisms to be applied to MPLS-encapsulated IP traffic.

The Current Status of DiffServ-Based Architectures

Although it is not clear if and when operators will implement a "standard-compliant" DiffServ archi-

The first recognized tool for QoS provisioning is RSVP, the signaling protocol defined in IETF that can be used by routers supporting IntServ along the path to set up per-flow quality of service.

Despite the IntServ and DiffServ work, the deployment of end-to-end QoS is negligible in the Internet. Starting from this consideration, the IETF's goal is to enhance QoS signaling in the Internet in order to deploy it in the real world.

tecture, the basic DiffServ concept (i.e., the differentiation of packets in traffic classes) is currently being exploited. The main driver is the support of voice traffic over IP networks. In both operator backbones and bottleneck links of corporate networks, voice traffic is carried with some differentiation with respect to other traffic.

Because DiffServ is simply a traffic handling mechanism, it needs to be configured in order to be able to provide useful services. According to the IETF DiffServ working group, the DiffServ network is configured to provide a certain service level as specified in a formal document defined as a *service level agreement* (SLA). An SLA is defined at the point where the customer submits traffic to the DiffServ network's ingress router, and represents a mutual commitment between the customer and the network: the first will abide to the traffic description he/she has provided to the network; the second will offer a communication service satisfying the performance requirements specified in the agreement.

The provisioning and configuration of DiffServ networks are based on the configuration of the traffic handling mechanisms. According to [4], this top-down provisioning is able to offer limited degrees of traffic prioritization, and needs to learn classification criteria and anticipate traffic patterns and volumes. Moreover, it requires centralization and coordination of the information.

Therefore, while DiffServ architecture solves the scalability problems of QoS provisioning, it fails to be the solution for end-to-end provisioning. The combination of RSVP signaling with aggregate traffic handling mechanisms, as discussed in [4], could address the deficiencies of the exclusively top-down provisioned approach of DiffServ without incurring the scalability problems of classical RSVP/IntServ usage. Anyway, this combination is not yet realized in a real-life network.

The Next Steps in Signaling Activity

Despite the IntServ and DiffServ work, the deployment of end-to-end QoS is negligible in the Internet. Starting from this consideration, the IETF working group Next Steps in Signaling (NSIS) has chartered its activity. The goal is to check how it is possible to enhance QoS signaling in the Internet in order to deploy it in the real world. For example, an issue is how to deal with interworking between several administrative domains where different QoS solutions exist. The NSIS WG also takes into account requirements coming from user mobility. The existing work (e.g., RSVP) is evaluated as a possible starting point. The NSIS is currently analyzing the architectural elements that should be involved in QoS signaling and where such signaling has to take place (end-to-end, end-to-edge, end-to-proxy, edge-to-edge, etc.). A document containing the requirements for the QoS signaling protocol was produced recently and is being discussed [5], while a document on the framework for NSIS signaling is at an earlier stage of preparation.

Policy-Based Networking

Policy-based network management leads to a standard and consistent way of network configuration, independent of the underlying architecture and QoS provisioning model assumptions. The event-driven paradigm, well established in the general-purpose programmers' world, through the policy-based approach begins to play its role also in the field of network management. A policy is a set of rules or methods, representing an object behavior or a decision strategy to be applied in order to ultimately reach a particular goal. Policy-based network management is the application of these organizational policies in order to manage networks. With this approach, the role of network management moves from passive network monitoring to active QoS and network SLA provisioning. The network view of policy is an intuitive high-level perspective of topology, connectivity, end-to-end performance objectives, and dynamic state of the network. The network view is composed of different nodal views, which correspond to the policy objectives and requirements at various network nodes. These, in turn, are composed of policy rules, which may be regarded as atomic injunctions through which various network nodes are controlled. Each network node has vendor-specific resource allocation mechanisms. Hence, packet forwarding paths and nodal views need to be ultimately translated into device-specific instructions [6].

However, while this technology is powerful and alluring, it is also generally untested and unproven. Worse, this area still suffers from a lack of standards and a lack of ad hoc use of existing ones. There are two key issues that are not yet totally addressed: first, how the vendors will access and control their hardware, and second, how these systems glean information about an organization's users and resources.

Issues for DiffServ-Based Architectures in the Real World

DiffServ-based architectures have received the greatest favor of operators and industry. However, despite this large interest, they are not yet implemented in the real world. The following are the major missing pieces needed for deploying QoS in the Internet. We also provide reference to the ongoing work in these areas.

Service Definition, Service Level Agreements

A key issue in packet networks is how to define the service offered by a network. This involves both technical and standardization aspects. In fact, the service definition should be simple and clear, and at the same time useful for allocating resources in the network. When the traffic is bursty, there is the known problem that using deterministic definition of services can be inefficient, while statistical service definition can be too difficult to understand, measure, and verify. However, this is also a standardization issue, since it is necessary to have a common agreement among service providers and users to formalize this service definition. An SLA (i.e., the contract between the provider and the users) needs to be based on this formal service definition. An effort in this direction is being pursued with the proposition of a service level specification (SLS) for IP QoS [7].

Dynamic Service Creation and Service Configuration Tools

These tools would allow an operator to define the services in an abstract way and then have an automatic mapping into the configuration of network devices. When moving QoS from laboratory trials to large-scale IP networks, these automatic management/configuration tools become essential, since it is not feasible to manually configure a large number of devices in a consistent way. Note that there is dependence on the issue of SLA/SLS, since a formalization of the QoS characteristics of the service could be useful in the proper configuration of the network. The idea of dynamic service creation is not new in the telecommunication systems area. However, in premium IP networks this idea assumes a different, more general and challenging meaning, since it implies the capability of a network to reconfigure itself in order to host new applications and services involving traffic patterns with completely new characteristics and requirements.

Traffic Engineering Tools

The DiffServ architecture basically specifies the PHB mechanisms, which apply "locally" on a router/link. How to manage the resources globally within a network backbone is out of the scope of the DiffServ WG, but it could be an important point for a provider. The process of managing the resources in a network backbone is referred to as *traffic engineering* and is the focus of the IETF TE WG. MPLS architecture is currently seen as a possible answer to TE needs. Both the MPLS WG [8] and TE WG [9] are working on the mapping of DiffServ over a MPLS backbone.

Dynamic Service Invocation

The first step in the introduction of IP QoS can consider a static association of QoS resources to a class of customers and/or applications. The further step is to make this association dynamic. The rationale behind this is that a more efficient use of resources is possible. Dynamic service invocation could add considerable value to QoS-enabled networks. For example, the allocation of bandwidth on demand can be of great interest for applications like video communication and videoconferencing, where a large amount of resources is needed for a limited period of time. The needed components to have dynamic IP QoS are a dynamic resource management architecture and an access signaling mechanism. The former is needed within the provider network and may include dynamic signaling between network elements, and the latter allows the user to signal its QoS needs to the provider. As discussed above, the NSIS work is addressing the aspects of QoS signaling in IP networks.

Monitoring

Monitoring of the provided QoS level is a critical component in the QoS framework. The IP Performance Metrics (IPPM) IETF Working Group [10] is currently developing a set of standard metrics that can be applied to the quality, performance, and reliability of Internet data delivery services.

In terms of measurement technologies, QoS provisioning needs to rely on a wide range of measurement tools, both active (analysis of network quality impact between two points on packets injected into the network by a measurement infrastructure) and passive (nonintrusive analysis of observed packet streams and router behavior at one or multiple points in the network). Hence, a measurement management architecture is required to unify the automated configuration, reporting, and analysis over these heterogeneous measurement tools.

QoS monitoring comprises two different aspects. The first aspect is "internal" to the service provider, which needs to monitor its own network. Traditionally the monitoring in packet networks is done offline with respect to network operation. Usage data are collected and used to build statistical reports for planning network upgrades (i.e., at the timescale of weeks or months). Rather, the monitoring of a QoS network could be "interlaced" with network operation since it can drive dynamic adjustments of network configuration at the timescale of seconds. Since traffic behavior is more difficult to foresee, monitoring becomes a part of network operation. An example of this concept is measurement-based admission control, where the current measured QoS level is used to dynamically control the access of new flows.

The other side of network monitoring is end-user control of received QoS. The user would like to have a clear understanding of whether the network is providing QoS or not. It is not always possible or enough to send probes, even though this is one of the possible techniques. There are some problems with elastic applications where the throughput depends on several factors besides the network QoS.

Interdomain QoS Aspects

The provisioning of IP QoS in a single domain requires proper solutions to most of the above issues. The standardization of these solutions facilitates interworking between equipment, but some nonstandard solutions (or "patches") can be accepted. If we want to really achieve Internet QoS rather than IP QoS, it should be possible to receive QoS across different domains. In this case standardization can become mandatory, and this is the first reason interdomain IP QoS is an even more challenging task.

In general, we can see that IP QoS issues should be solved in the single-domain case before tackling interdomain QoS. Then the new dimensions of multidomain QoS must be taken into account: what happens at the interface between different providers must be clarified; the combination of single-domain QoS levels into the whole end-to-end QoS is to be faced. For example, the standardization of service definition and SLS for the customer-network interface for the single domain can be the base to define SLSs for the end-to-end QoS and for provider-to-provider QoS.

Again, interdomain QoS can be on a static or dynamic level. Probably as in the intradomain case, the approach will be to start from static associations of users and applications to QoS, since these do not require interdomain signaling. Later, the use of signaling can increase the efficiency.

In terms of measurement technologies, QoS provisioning needs to rely on a wide range of measurement tools, both active and passive. Hence, a measurement management architecture is required to unify the automated configuration, reporting, and analysis over these heterogeneous measurement tools.

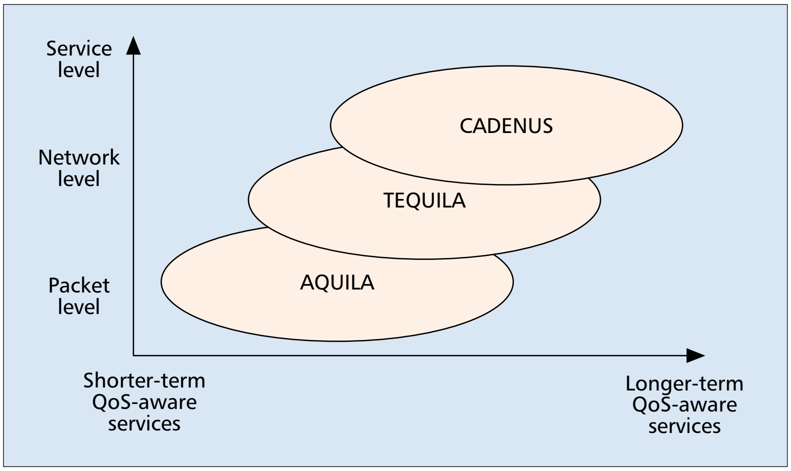

■ **Figure 1.** *QoS focus of the AQUILA, TEQUILA, and CADENUS projects.*

The Proposed Solutions of European Projects

In the context of the IST framework program [2], the EU has founded three research projects in the area of QoS support in large IP networks: AQUILA, CADENUS, and TEQUILA (Fig. 1). They have the same high-level objective of developing solutions for provisioning IP premium services (i.e., IP QoS) as a fundamental step toward the next-generation networks (NGNs). The three projects have focused on different specific issues, but of course with some overlaps in the problem space and the solution space. This was an opportunity to collaborate and progress in parallel. Each project also benefited from the possibility to integrate and compare its own results with other projects' work and results (Fig. 2).

Advanced service offerings require a sound and well-structured business model in which to operate, combined with technical possibilities to specify and manage the services in an automated way. Once a service is specified, it can be used as an input to the previously mentioned QoS provisioning architectures in order to actually implement the services on the wire. In the following paragraphs, the service and network management focus of the abovementioned projects will be outlined.

Service Management and Service Definition

The specification of a business process is the main activity required when creating a new service. The CADENUS project considers the current business processes involved, especially including the definition, and increased flexibility and dynamic nature, of SLAs and SLSs. A fundamental aspect of the CADENUS architecture is the separation of service providers from network (resource) providers. In this way, the resources aspects are hidden to the service providers, which have only an abstract view of them. The architecture also includes other levels in order to simplify and add robustness. CADENUS maps its architecture to the Telecommunications Management Forum (TMF) business processes. It considers the overall environment (access, service, and wholesale), including residential and business end systems.

AQUILA and TEQUILA assume an already specified service to be given as an input to the system. AQUILA has therefore developed the End-user Application Toolkit (EAT) that aims to provide access for end-user applications to QoS features. The EAT is a middleware between the end-user applications and the network infrastructure. The EAT supports two major kinds of (Internet) applications: legacy applications (QoS-unaware) and QoS-aware applications. The AQUILA network offers different network services with different predefined QoS characteristics to the customers of the network and implements them internally by different traffic classes. The network services are stored as XML data (based on a common Document Type Definition) on a central directory server. The QoS Management Tool provides network operators access to network services.

TEQUILA makes abstractions by considering "customer-provider" relationships. The provider is a network provider and the customer may be a company, a residential user, an application provider, another provider, or any other legal entity subscribing to a (network) service. The TEQUILA network architecture expects to activate service requirements in two (possible coinciding) epochs: a subscription and an invocation action. In the subscription action, a service will be specified in an SLS, and will be transported from the user to the TEQUILA architecture for which the generic session-oriented Service Negotiation Protocol (used to establish, modify, and terminate service contracts) was developed. The actual activation or invocation of the service is done either implicitly (i.e.,, an always-on service) or explicitly by notifying the provider's edge (e.g., using RSVP).

Network Resource Management and Traffic Engineering

Once application requirements are translated into services, their resource requirements must be translated into a network configuration. This builds on technologies like DiffServ and MPLS available in the network, management technologies to configure them, and measurements to validate the provisioning and optionally adapt the management strategy.

The three projects deal with QoS-provisioning using traffic engineering from both the long term centralized and short-term distributed aspects. CADENUS and TEQUILA explicitly base their resource management on a policy-based approach. CADENUS introduces the policies at different levels, starting from the translation of the SLA down to the set of commands for setting the network devices. The policies are stored in resource repositories (databases), which are accessible from the entity that handles the resources.

In TEQUILA, the use of high-level policies is the key element in driving the traffic engineering algorithms. In particular, this is used for network dimensioning, which allows calculating the long-term network configuration based on network status information, and current and forecast service subscriptions. TEQUILA proposed a set of service definition and traffic engineering tools to

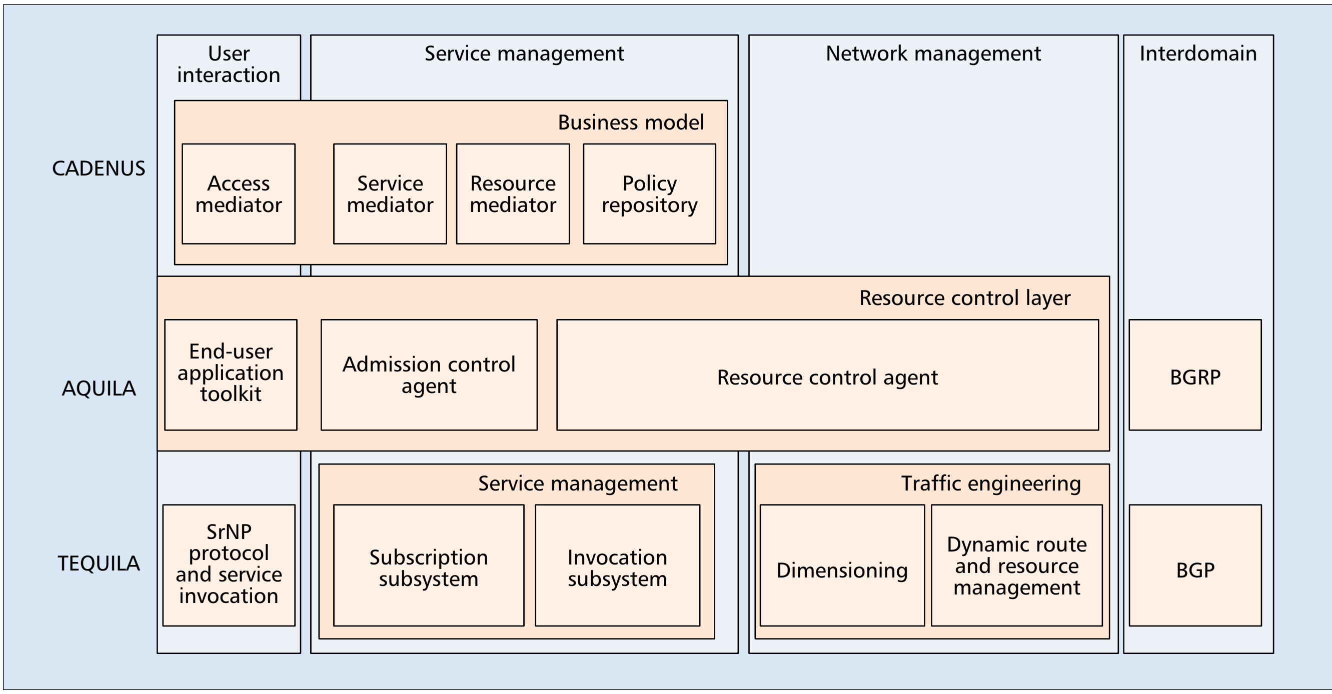

Figure 2. *Architectural view of the AQUILA, TEQUILA, and CADENUS projects.*

obtain quantitative end-to-end QoS guarantees through careful planning, dimensioning, and dynamic control of scalable and simple qualitative traffic management techniques within the Internet (i.e., DiffServ over MPLS).

AQUILA does not use policies, but performs global management of the resources via the resource control agent (RCA), which monitors, controls, and distributes the resources in the network.

Moreover, in addition to the approaches to long-term management, the three projects propose approaches for handling the backbone resources in the short term in a dynamic way. This two-level traffic engineering approach (delegating short-term traffic engineering decisions to components closer to the wire, as opposed to a single monolithic bandwidth broker) is required to perform QoS provisioning in a scalable and flexible way. TEQUILA performs dynamic admission control, as well as dynamic route and resource management that is operated within boundary conditions set by the network dimensioning. Communication between the entities that handle the services and resources allows dynamic admission control at the time of registration and invocation, as appropriate for the service. Finally, AQUILA dynamically handles the resources with a hierarchical structure with resource pools at the network edges and a per-flow-based admission control independent of resource management. The admission control agent (ACA) controls access to the network by performing policy control and admission control. EAT, ACA, and RCA together make up the resource control layer (RCL), an overlay network on top of the DiffServ core network. This way the RCL provides an abstraction of the underlying layers.

MONITORING

Since the projects aim at provisioning advanced QoS services in a heterogeneous flexible environment, network monitoring plays an important role.

In CADENUS monitoring and measurement is performed at all of the interfaces between the mediators. TEQUILA investigates a monitoring architecture for: assisting traffic engineering in allocating resources efficiently and dimensioning the network for any short- or long-term changes; in-service verification of the traffic and (QoS) performance characteristics by monitoring customer-specific SLSs. AQUILA defines and implements a distributed measurement architecture that is used for two main tasks. First is support of network operation and resource control performed through the network operator. This is used to enable measurement-based admission control and to give the operator a view of the current situation within the network. The second task is to validate the implemented QoS architecture including evaluation of the end-to-end QoS of network services and validation of the admission control. This is used to support traffic engineering in the design of the algorithms and their parameters.

INTERDOMAIN QOS

End-to-end QoS provisioning must also reach across administrative boundaries. AQUILA investigated the extensions of its QoS model to interdomain starting from the BGRP framework proposed in [11]. Dynamic reservations spanning multiple domains are supported. In order to address the scalability issue, the Border Gateway Reservation Protocol (BGRP) aggregates reservations that span multiple domains considering the destination domain, limiting the state information that must be maintained. Some mechanisms are also defined in order to avoid the signaling for each reservation having to travel all along the path from the source to the destination, limiting the signaling load. The BGRP relies on routing information provided by the

In order to address the scalability issue, the BGRP aggregates reservations that span multiple domains considering the destination domain, limiting the state information that must be maintained.

traditional interdomain routing protocol (BGP): routing is not affected, as BGRP is basically an admission control mechanism.

TEQUILA takes a complementary approach by enhancing the BGP protocol in order to convey QoS-related information, such as the delay or loss experienced by packets for a given destination on a route. This information can be used to modify the routing according to the QoS requirements, with the aim of building interdomain end-to-end QoS paths.

CHALLENGES AND FUTURE DIRECTIONS

In the last decades a powerful set of technologies has emerged from worldwide research to support QoS in today's networks. The IST Premium IP projects have built on this in order to provide the required infrastructure to create and offer advanced QoS-aware flexible services. By automating all layers involved in management and provisioning, this research has made "zero-effort" network management combined with value-added service offering possible. This also opens up the way to make services currently considered hard to implement — demanding multimedia, gaming, global grid computing, content distribution, and so on — as common as basic connectivity.

The current efforts spent on bringing people together, with backgrounds ranging from business definition, service creation, and traffic engineering to low-level monitoring and forwarding differentiation, should be a starting point to further enhance cooperation between all these different networking areas. Better and standardized synchronization between the network layers, and of the interconnection between players is a logical continuation of this.

REFERENCES

[1] Operax, "The Challenge of Enabling QoS in IP Networks," Oct. 2001, http://www.operax.com/docs/whitepaper-QoS-challenge-C.pdf

[2] The EU IST Programme, http://www.cordis.lu/ist

[3] X. Xiao and L. M. Ni "Internet QoS: A Big Picture," IEEE Networks, Mar. 1999.

[4] Y. Bernet, " The Complementary Roles of RSVP and Differentiated Services in the Full-Service QoS Network," *IEEE Commun. Mag.*, Feb. 2000.

[5] M. Brunner, Ed., "Requirements for Signalling Protocols," <draft-ietf-nsis-req-04.txt>, Aug. 2002, work in progress.

[6] R. Rajan *et al.*, "A Policy Framework for Integrated and Differentiated Service in the Internet," http://www.allot.com/html/white_products_technology_pbn.shtm

[7] D. Goderis *et al.*, "Service Level Specification Semantics, Parameters and negotiation requirements," <draft-tequila-sls-01.txt>, June 2001, work in progress, available at http://www.ist-tequila.org/standard

[8] F. Le Faucheur *et al.*, Ed. "Multi-Protocol Label Switching (MPLS) Support of Differentiated Services," IETF RFC 3270, May 2002.

[9] F. Le Faucheur and W. Lai, "Requirements for Support of Diff-Serv-aware MPLS Traffic Engineering," <draft-ietf-tewg-diff-te-reqts-06.txt>, Sept. 2002, work in progress.

[10] IPPM WG, http://ww.ietf.org/html.charters/ippm-charter.html

[11] P. P. Pan, E. L. Hahne, and H. G. Schulzrinne "BGRP: Sink-Tree-Based Aggregation for Inter-Domain Reservations," KICS 2000.

BIOGRAPHIES

SILVIA GIORDANO [M] (silvia.giordano@epfl.ch) has a Ph.D. from EPFL, Lausanne, Switzerland. She is currently working as senior researcher at the ICA institute at EPFL and at SUPSI (Ticino), and, since October 2001, as scientific collaborator at CNUCE, Pisa. She has published extensively in the areas of quality of services, traffic control, and mobile ad hoc networks. She has participated in several European ACTS/IST projects and European Science Foundation (ESF) activities. Since 1999 she is a Technical Editor of *IEEE Communication Magazine*. She was already co-editor of some special issues of *IEEE Communication Magazine* and B*altzer MONET and Cluster Networks* on mobile ad hoc networking and QoS networking. She was on the executive committee and TCP of several international conferences, and has served as reviewer on transactions and journals, as well as for several important conferences. She is a member of IFIP WG 6.8. Her current research interests include traffic control and mobile ad hoc networks.

STEFANO SALSANO (stefano.salsano@uniroma2.it) received his Laurea degree in 1994 from the University of Rome "Tor Vergata" and his PhD in 1998 from the University of Rome "La Sapienza". Since November 2000 he has been an assistant professor at the University of Rome "Tor Vergata." From 1997 to 2000 he was been with CoRiTeL, a research institute on telecommunications, coordinating IP-related research activities. His current research interests include QoS and traffic engineering in IP networks, IP telephony, MPLS, and IP over optics.

STEVEN VAN DEN BERGHE graduated in computer science from Ghent University in 1999. In July 1999 he joined the Broadband Communications Networks Group of the Department of Information Technology, Ghent University, where he is working toward a Ph.D. In January 2001 he was granted an IWT scholarship. His research interests are mainly in the area of quality of service and traffic engineering in IP. He is focusing on measurement-based traffic engineering in a Diff-Serv/MPLS/multipath environment. He was active in the IST TEQUILA project, development of DiffServ support for MPLS in the Linux community, and has made several contributions to conferences and standardization efforts on the use of measurement architectures in traffic engineered IP networks.

GIORGIO VENTRE is associate professor of computer networks in the Department of Computer Engineering and Systems of the University of Napoli Federico II. His has a Laurea degree in electronic engineering and a Ph.D. in computer engineering, both from the University of Napoli Federico II.. From 1989 to 1991 he worked at CPS, Center for Research on Parallel Computers of the Italian National Research Council (CNR), doing research in the area of system support for distributed memory computers. From 1991 to 1993 he was with the Tenet Group at the International Computer Science Institute and the University of California at Berkeley, working in the area of real-time protocols for multimedia applications. Since 1993 he has been at the University of Napoli Federico II, where he is co-leader of the Computer for Interaction and Communications (COMICS) team, a research initiative in the areas of networking and multimedia communications. Recently he was appointed director of ITEM, a research laboratory on multimedia application founded by CINI, the Italian University Consortium for Informatics. ITEM is located in Napoli and hosts a state-of-the--art technological infrastructure in the area of telematics and multimedia systems. He is a member of the IEEE Computer Society and the ACM. He has served on the Program Committees of international conferences and workshops, and is a reviewer for several international journals and conferences in the area of distributed systems and communication networks.

DIMITRIOS GIANNAKOPOULOS received his diploma in electrical and computer engineering from National Technical University of Athens in 1998 and has continued his postgraduate studies in the same department. Currently he works as a research associate at the Telecommunications Laboratory of the National Technical University of Athens dealing with software development for networking and monitoring traffic, and is involved in the TEQUILA project in the area of service management.