INTERNATIONAL JOURNAL OF COMMUNICATION SYSTEMS
Int. J. Commun. Syst. 2004; **17**:(in press) (DOI: 10.1002/dac.645)

Providing and verifying advanced IP services in hierarchical DiffServ networks—the case of GEANT

Athanassios Liakopoulos[1,2,‡], Basil Maglaris[1,2,§], Christos Bouras[3,4,*,†] and Afrodite Sevasti[2,3,¶]

[1] *Department of Electrical & Computer Engineering, National Technical University of Athens (NTUA), 15780 Zografou, Athens, Greece*
[2] *Greek Research and Technology Network (GRNET), 56 Mesogion Ave., 11574, Athens, Greece*
[3] *Department of Computer Engineering and Informatics, University of Patras, 26500 Rion, Patras, Greece*
[4] *Computer Technology Institute, 61 Riga Feraiou Str., 262 21 Patras, Greece*

SUMMARY

The differentiated services (DiffServ) framework is widely proposed as an efficient method for providing advanced IP services to large-scale networks, with QoS requirements. However, the provisioning of such services in production networks has proved to be more difficult than initially expected, in defining, setting and verifying appropriate Service Level Agreements (SLAs). GEANT, the Gigabit core pan-European research network, on a pilot basis introduced 'Premium IP' service, offering bounded delay and negligible packet loss to the European National Research & Education Networks (NRENs) that it interconnects. However, large scale provisioning of this new service requires the definition of efficient interaction procedures between administrative domains involved and methods for SLA monitoring. This paper focuses on these issues and presents the experience acquired from the early experiments in GEANT, as an example of hierarchical Gigabit multi-domain environment, enabled with QoS provisioning to its constituent NRENs. This model scales more efficiently than the common peering Internet Service provider (ISP) commercial paradigm. Finally, we outline other options that promise QoS, such as Layer 2 VPNs in MPLS backbones, with non-standard (yet) mechanisms. Copyright © 2004 John Wiley & Sons, Ltd.

KEY WORDS: SLAs; service provisioning; monitoring; differentiated services; premium IP

1. INTRODUCTION

Nowadays, the co-existence of diverse networked applications has rendered the traditional best-effort service model of the Internet inadequate. New applications, such as IP telephony, require

*Correspondence to: Christos Bouras, Research Academic Computer Technology Institute, Riga Feraiou 61 str, GR 26221 Patras, Greece.
† E-mail: bouras@cti.gr
‡ E-mail: aliako@grnet.gr
§ E-mail: maglaris@mail.ntua.gr
¶ E-mail: sevasti@grnet.gr

Contract/grant sponsor: GEANT and SEQUIN; Contract/grant sponsor: GRNET

Received March 2003
Revised December 2003
Accepted January 2004

QoS and virtual private network (VPN) services to be supported across multiple administrative domains. Today, ATM and newer standardized or proprietary Layer 2—multi-protocol label switching (MPLS) IP technologies are used for partially fulfilling the afore-mentioned requirements as they can separate traffic into different categories, e.g. data vs voice, offer service guarantees and provide many-to-many connectivity to a limited number of sites. However, the increase of voice and other time-sensitive traffic over packet-switched networks gives rise to new business models for IP networks that can provide QoS and VPNs in an efficient and scalable way.

Among several proposals, the DiffServ framework specified by IETF [1] stands out for providing service differentiation to traffic in a scalable manner. It suggests the aggregation of individual application flows with similar quality needs and it introduces the definition of different service classes to which such aggregates are appointed. The DiffServ framework operates on the basis of marking the packets of individual flows that belong to a certain QoS class with a single differentiated services codepoint (DSCP) value (or a group of DSCPs). Marking is achieved by setting the DSCP field, namely the six most significant bits of the IP packet header ToS field. In the interior of each DiffServ-enabled domain, queuing and scheduling of packets is performed according to their DSCP value and not to the flow to which they belong. In other words, in the interior of a DiffServ domain, all packets belonging to the same QoS class are indistinguishable and thus receiving the same treatment or per-hop-behaviour (PHB) [1] according to the DiffServ terminology.

Although the DiffServ framework has been initially received in a positive way, due to its scalability and straightforward implementation, it was soon realized that the introduction of advanced IP services with QoS requirements in large scale networks was not an easy task [2]. Actually, the provisioning of DiffServ-based IP services networks has not been widely deployed with success in production networks.

Generally, state-of-the-art core routers adequately support the required QoS functionality as application specific integrated circuits (ASICs) enable multiple complex classification, marking, queuing and scheduling mechanisms at gigabit speeds. Furthermore, several proof-of-concept demonstrations of advanced IP services in lab environments were successful in the past, which also proves that needed functionality is already available in the market. However, as already stated, service providers with DiffServ-enabled backbone networks hesitate to provide their customers with advanced IP services even if their core routers support the necessary QoS functionality. The main reason for these rather 'contradictory' facts is that there is no efficient provisioning model for the new IP services. A provisioning model defines the methodology for accomplishing service level agreements (SLAs), i.e. contracts that specify service guarantees, between the service provider and its customers and defines the procedures for setting-up and validating the appropriate QoS mechanisms in the core routers.

GEANT[‖], the Next Generation pan-European Research Network, is a high-speed backbone network that connects more than 30 European National Research and Education Networks (NRENs) (Figure 1). Its predecessor, TEN-155, offered different QoS guarantees to traffic based on ATM technology, especially through the managed bandwidth service (MBS) that allowed the provision of VPNs with committed bandwidth between the NRENs. However, as GEANT does not support ATM services any more, its advanced IP services can only be supported via DiffServ

[‖] http://www.dante.net/geant/

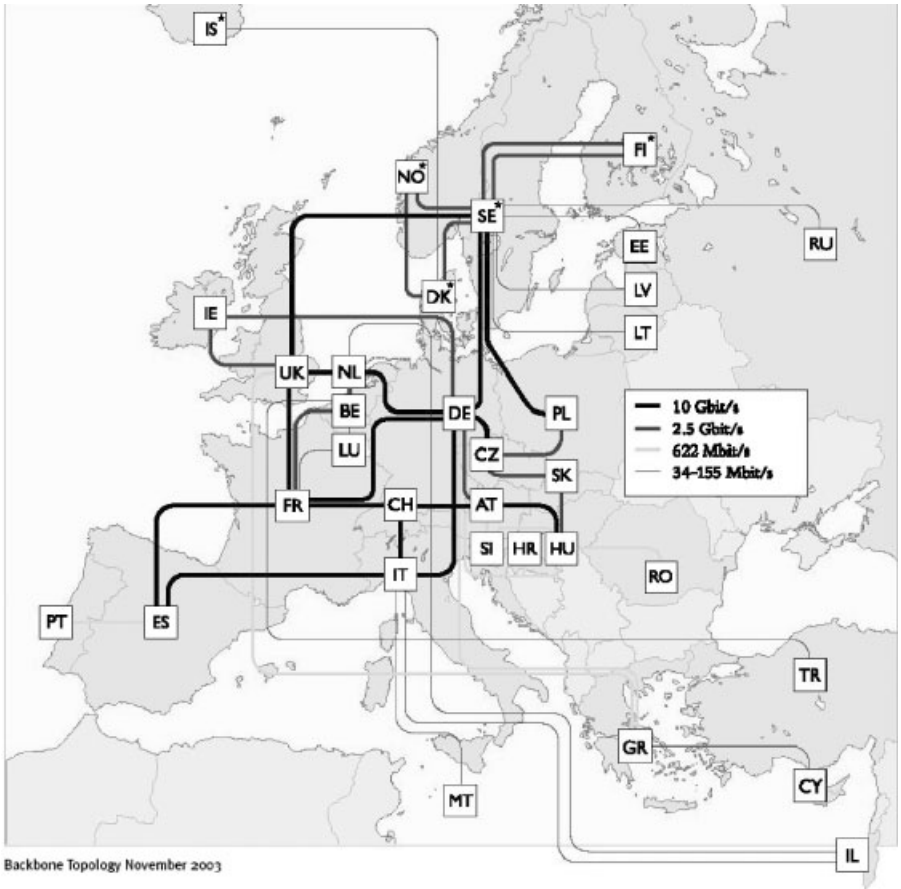

Figure 1. GEANT logical topology map.

mechanisms. SEQUIN [3], a European Commission Information Society Technologies (IST) research project, in cooperation with DANTE, the coordinator of GEANT, specified the 'Premium IP' service implementation model based on the DiffServ Expedited Forwarding PHB (EF-PHB) [4]. Premium IP offers explicit QoS guarantees, such as bounded delay and delay variation, and negligible packet loss to conforming traffic. In simplified terms, Premium IP could be defined as 'virtual leased line' service.

SEQUIN addressed problems related to end-to-end QoS provisioning across multi-domain networks, determined the procedures for establishing SLAs between different administrative domains and investigated methods for deploying a QoS monitoring infrastructure in a large-scale network. Premium IP service was initially deployed in a limited part of the GEANT production network during the proof-of-concept trials and the service was provided to multiple research groups in Europe.

In this paper, we present our experience gained from our participation in the SEQUIN project. Initially, we present the different phases in provisioning DiffServ-based QoS-aware IP services in backbone networks. We describe the agreements between the service providers and

their customers with respect to the QoS provided and outline the principles for monitoring the compliance of the service provider. Finally, we focus on Premium IP service and suggest how the afore-mentioned issues could be addressed in a hierarchical environment, such as GEANT that interconnects European NRENs.

2. SERVICE PROVISIONING PHASES

An advanced IP service in a DiffServ-enabled network is realized in two consecutive steps. In the first step, the IP service implementation model is carefully designed and deployed. Multiple per-hop behaviours (PHBs) are carefully defined in the network and appropriate QoS mechanisms, such as queuing and scheduling mechanisms, are enabled at the network elements. Quality guarantees may then be provided to portions of traffic. In the second phase, a feasible service-provisioning model is defined. This facilitates the service request by the customer and the service provisioning by the service provider. In addition, an appropriate infrastructure is required to assess the provided services with respect to what the service is promised to deliver.

The provision of advanced IP services between two users has to be established through a number of phases, as depicted in Figure 2. At the beginning, a *negotiation phase* should clarify the entities involved, the purpose for which advanced IP service will be provided, the feasibility of provision etc.

During the *service set-up phase*, all service provision details have to be collected, the necessary service level agreements (SLAs) have to be signed and detailed configuration of the equipment involved must be performed. Each SLA can be considered as a formal definition of the quality of the service provided, setting the performance objectives that the service provider must achieve [5]. The SLA subset that contains technical details of the contract is referred to as the service level specification (SLS). After the definition of an SLA, several SLSs have to be produced, one for each of the several peering domains involved in an end-to-end path. Each SLS contains the

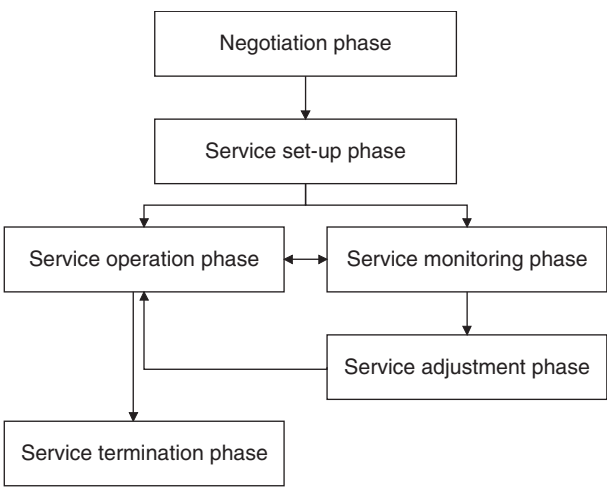

Figure 2. Phases for advance IP service provision.

parameters and values that describe the transport service that a specified flow should receive in one direction of the end-to-end path. Bi-directional SLAs across the network are possible by combing two SLSs (one in each direction) when negotiating a service.

During the *service operation phase*, no specific activities have to be performed unless indications of service degradation or failure occur. In such a case, measures have to be taken so that the service operation is restored. In parallel to the service operation phase, the *monitoring phase* should take place, comprising of constant measurement activities with the purpose of verifying the service quality. In case that the service performance deviates from the desired one the *service adjustment phase* will have to be initiated, which involves adjustments to the router configuration along the service provision path. A service adjustment phase always results in new service operation and monitoring phases that run in parallel, until the service provision time frame expires and the service termination phase is introduced.

The following two sections will mainly deal with the establishment of SLAs in the service set-up phase and the specification of the monitoring infrastructure** required in the monitoring phase.

3. INTERCONNECTION MODELS AND SLAS ESTABLISHMENT

There are two fundamental models for building interconnection networks that serve numerous communities across large geographical areas; the peering and the hierarchical. The peering model is based on multiple interconnection agreements among providers and customers or between peering providers in large tele-houses or Internet Exchange (IX) points. Traditionally, Internet was based on this model, which still prevails in liberalized commercial markets. However, large communities with common objectives, such as the education, research or public administration communities are suitable cases for the hierarchical model. GEANT and the European NRENs are characteristic examples of networks developed in the last decade according to this model. In fact, this model allows for scaleable multi-party SLA management, as we will demonstrate below.

The set-up phase of a new service includes the SLA negotiation between the involved entities. SLAs are contracts between customers and providers that specify the agreed upon service guarantees. In many cases, they are bilateral contracts between peering providers that provide inter-domain IP services to their customers. QoS-aware SLA specification for DiffServ-enabled networks aims at providing quality guarantees and setting out the limits of the services provided. In analogy to ATM traffic contracts [6], QoS aware SLAs enhance traditional SLAs by adding to availability, security and quantity of allocated resources metrics, appropriate quality parameters.

The establishment of QoS-aware SLAs in networks with multiple peering agreements is rarely adopted by large service providers, as it imposes many technical and operational problems. The complexity of supporting QoS-aware SLAs becomes a prohibiting factor admitting new service requests; these should carefully be handled in accordance to the numerous peering agreements of the provider(s) involved and take into account possible inter-domain routing changes.

**We define as "*monitoring infrastructure*" the group of network nodes, which perform the QoS performance measurement tests, analyse the collected data and exhibit it through a user-friendly interface.

In addition, it is difficult to standardize the provisioning procedures, as there is a large number of peering scenarios already in service today.

On the contrary, the establishment of QoS-aware SLAs in hierarchical networks can be more easily realized due to the well-established network topology and the existence of a common controlling body responsible for a common policy across multiple administrative domains. We will deal in the sequel with hierarchical networks, as the title of the paper suggests. The establishment of QoS aware SLAs in hierarchical networks is presently a rather static and labour intensive task. The current non-uniform nature of SLA establishment procedures in each domain, coupled with QoS interoperability problems across multi-domain IP networks are factors that limit the number of end-to-end service requests. The proposed approach is one of establishing bilateral SLAs between neighbouring domains and using them as a basis for forming an end-to-end SLA along the service provisioning path.

The DiffServ framework provides a set of guidelines rather than strictly defined service models. This is one of the advantages of the DiffServ approach while at the same time has been a limitation for its wide deployment. Even in the case of offering a limited set of supported service classes, the implementation details for each service across consecutive domains differ (due to different policies, different functionality supported by the available equipment, etc.) and therefore SLA parameters are actually variable and negotiable over a range. This emphasizes the need for end-to-end SLA establishment, which focuses on masking differences of individual SLA and achieving a uniform end-to-end service level.

Standardization efforts of SLA procedures will allow for a highly developed level of automation and dynamic negotiation of the contract technical details, the service level specifications (SLS). This automation may prove helpful in providing customers (as well as providers) with the technical means for the dynamic provisioning of QoS guaranteed transport services across multiple domains.

4. MONITORING: SERVICE VERIFICATION

During the service-monitoring phase, the service provider is committed to constantly assess the performance of his/her network and verify that the performance agreed upon with his/her customers is met. This is a complicated and extremely cumbersome task, as the appropriate monitoring infrastructure should be deployed in the network and multiple performance measurements should be obtained. However, service verification is a strict requirement imposed by the customers, who demand well-defined levels of service.

4.1. Monitoring scope

The exact location of the monitoring 'nodes' in the network and the measurements performed between them defines the monitoring scope of the service.

The usefulness of QoS measurements is affected by the location of the monitoring nodes. Service providers that wish to evaluate the end-to-end network performance install nodes as close as possible to their users. For increasing scalability, the monitoring nodes may be located at central PoPs of each administrative domain, as in Figure 3(a). In this case, the distance, in terms of number of hops, between the users and the monitoring nodes is typically short and thus measurements provide a good approximation of the end-to-end network performance.

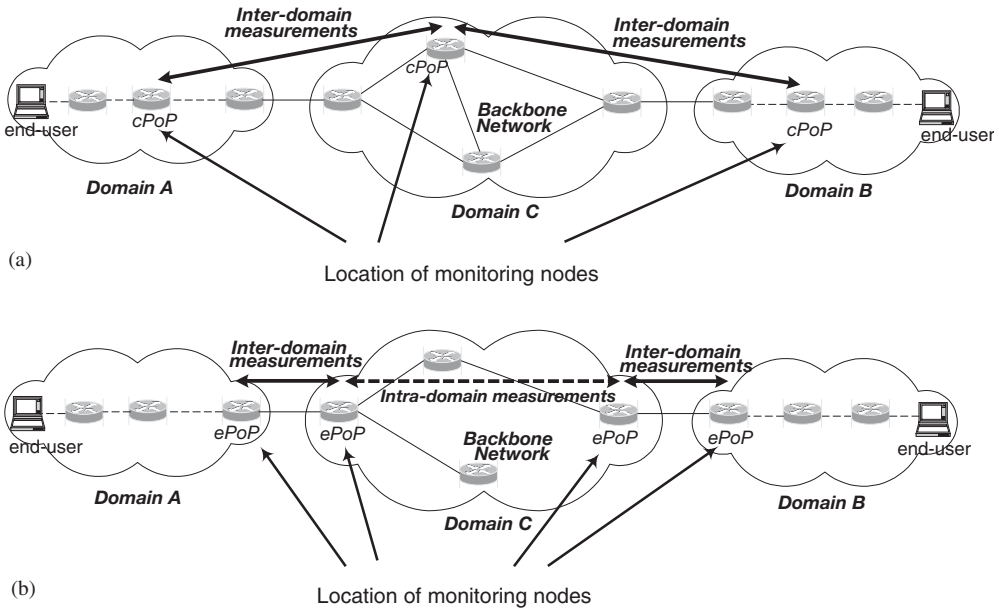

Figure 3. Monitoring scope.

However, it is difficult to locate which part of the network is responsible for the performance deviation. Alternatively, the monitoring nodes are located at the edge PoPs of each administrative domain, as in Figure 3(b). Intra- and inter-domain measurements are performed that can identify which domain is responsible for the SLAs violations in case the end-users perceive low service performance.

4.2. Active vs passive methods

Monitoring methods are classified as either *active* (*intrusive*) or *passive* (*non-intrusive*). In active monitoring methods, synthetic test traffic is generated and injected into the network by monitoring nodes. Test traffic shares the same network resources with the real traffic and therefore it encounters the same queuing delays and packet losses in congested networks. Active method experiments are easily controlled and repeated by the service provider who should, however, generate the appropriate test traffic patterns. Passive monitoring methods are based on information collected by network nodes, usually dedicated packet capturing nodes or backbone routers. In a typical configuration, an optical splitter is used to divert real traffic to the packet capturing (monitoring) node, which timestamps each received packet, classifies it and forwards collected information (packet ID and timestamp) to external servers for post-analysis. Passive monitoring methods are usually deployed easily and they do not burden the network under consideration with additional traffic. However, the lack of accurate network synchronization between the network nodes reduces the achieved accuracy of the measurements.

Monitoring in best-effort networks is often used for diagnostic purposes, such as locating network connectivity problems. On the contrary, in QoS-enabled networks a monitoring

infrastructure should provide accurate time measurements that can record minimum delays of time-sensitive traffic, e.g. for Voice over IP (VoIP) and Video Conference.

Accurate synchronization between the clocks of the monitoring nodes is obviously required for accurate time-related measurements in backbone networks that expand in wide areas and consist of diverse equipment and transmission links.

4.3. Clock synchronization

Synchronization of network nodes in distant locations is commonly achieved with external time sources that assist each node to adjust and correct its internal clock. Some commonly deployed methods are based on global position systems (GPS) or radio receivers coupled with network time protocol (NTP). It should be noted that there is a trade-off between accuracy, complexity and required cost for each of these methods.

In contemporary networks, NTP is widely used but the achieved synchronization usually suffers from approximately ± 10 ms error. When precise synchronization is a strict requirement, GPS receivers are often used.

4.4. Data presentation and analysis

Service providers should make available to their customers a user interface to access the network measurement data. Graphical views of raw or statistically analysed data are strongly desirable.

SLAs usually determine the graphical presentation of monitoring data required to verify the agreement. If, for example, a qualitative performance metric is evaluated every five minutes, then the monitoring graphs should be accordingly generated. In addition, SLA contracts are influenced by the ability of the service provider to collect accurate performance measurement data and publish them to its customers. For example, if the time accuracy of collected monitoring data is 20 ms then the granularity of time-related values of the service provider's SLAs should be at least of the same order.

4.5. Scalability

A monitoring infrastructure should be able to scale efficiently according to the network size (in terms of number of nodes and number of links), the transmission speed of the links (in terms of bandwidth) and the number of diverse QoS services provided.

Data collected during the measurement tests at the monitoring nodes could either be processed locally (distributed model) or forwarded to a central server (centralized model) for analysis and archiving. For QoS-enabled networks, the distributed model is preferable as it scales more efficiently and thus the monitoring infrastructures can be expanded across multiple administrative domains and come close to the end-users.

4.6. Security and data integrity

A monitoring infrastructure is necessary to support security features. There are two hazards that should be considered. Firstly, QoS measurements may be tampered by malicious users who are exploiting security breaches and take control of monitoring nodes. In this case, corrupted information may be sent to other monitoring nodes and be accepted as legitimate traffic. Secondly, in cases where active measurements are performed, a misbehaving monitoring node may inject large amounts of traffic in the network and become a potential harm for supported

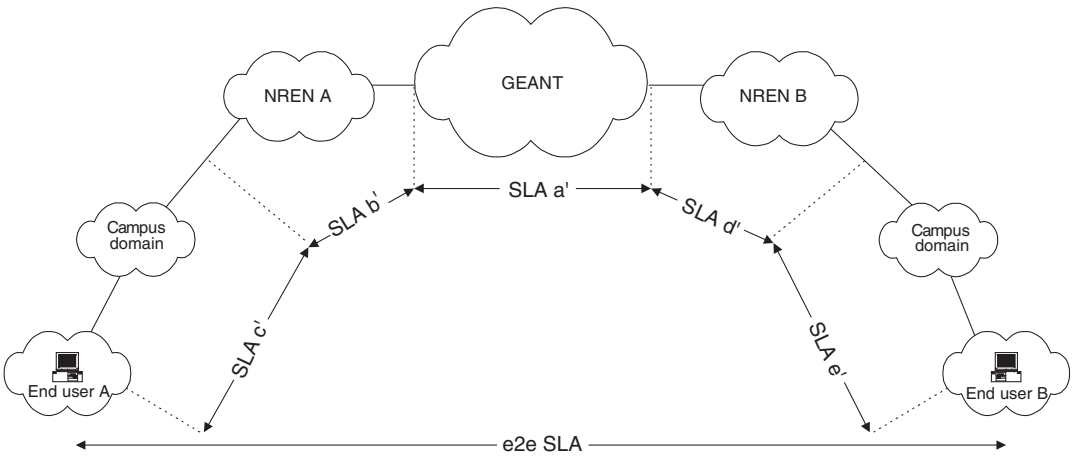

Figure 4. Premium IP provisioning: end-to-end SLA establishment.

SLAs. This could be considered as a denial of service (DoS) attack. Apparently, the monitoring infrastructure should support for security purposes authentication and authorization features.

5. PREMIUM IP SERVICE PROVISIONING IN GEANT

A typical example of SLA-based QoS provisioning in hierarchical multi-domain networks is the Premium IP Service within the GEANT–NREN community; its main purpose is the support of large multinational R&D Projects involving QoS sensitive experiments (e.g. tele-teaching, GRIDs and *e-Science* applications). Hence, Premium IP has to be delivered across multiple domains (end-user sites, NRENs and GEANT as shown in Figure 4).

A number of entities should be appointed and involved in the different service provision phases of Premium IP. For the co-ordination of the negotiation, set-up and operation phases, it is recommended that the end-users (NRENs or large R&D Projects) appoint a common representative towards the core interconnection network GEANT, referred to as the *Premium IP Service Provision Coordinator*. His/hers duties involve the mediation between GEANT and the end-users, the coordination of the service provision, the establishment procedures as well as any tasks required during the operation phase.

A technical person is also appointed as responsible for the service provision and implementation for each of the end-users. As depicted in Figure 5, *Technical Contacts A & B* (TC A & TC B) should be responsible for the service set-up and maintenance from end-users A and B, up to NREN A and NREN B domains, respectively. Their responsibilities include technical assistance to all *Site Administrators* (SA)[††] in their domain (NREN).

Similarly, GEANT has to appoint a technical person for the Premium IP service provision and maintenance. As depicted in Figure 5, the GEANT *Technical Contact* (GEANT TC) is

[††] Sites in NRENs may be University or Research Centre Local Area Networks.

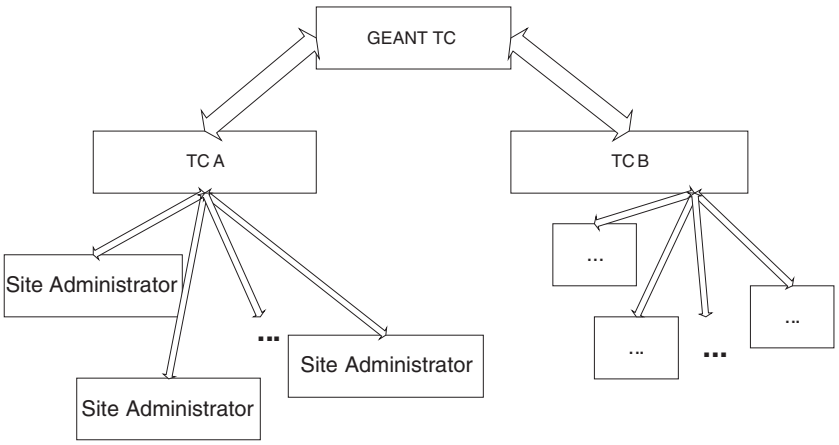

Figure 5. Hierarchical communication of technical contacts involved in Premium IP provisioning.

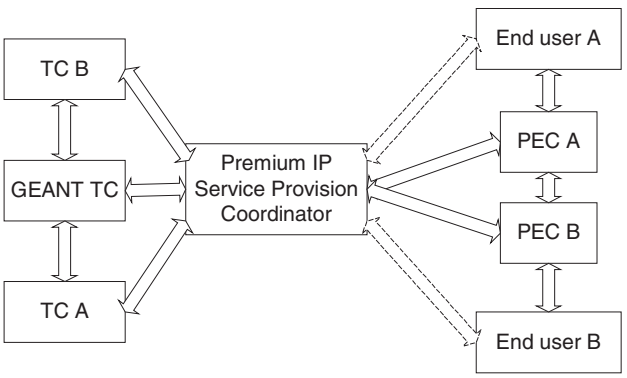

Figure 6. Entities involved in the provisioning of the Premium IP service.

responsible for the Premium IP service provisioning in GEANT network while at the same time providing any feedback required to the TCs from each end-user's side.

According to the previous definitions, Premium IP provisioning is performed in a hierarchical manner. The GEANT TC coordinates end-user domain TCs, while the latter coordinates its site SAs (Figure 5).

Apart from the technical responsibility, each end-user domain (NREN) appoints a *Performance Evaluation Contact* (PEC). These are responsible for checking whether the Premium IP implementation delivers to the end-users the quality they requested and if not, for advising adjustments to the SLA/SLSs. Their recommendations for adjustments should then be delivered to the TCs of each side via the SPC in order to be translated into re-configuration actions in the equipment involved.

Figure 6 depicts a possible way for the proposed entities' communication, with the *Premium IP Service Provision Coordinator* acting as an intermediate between TCs, PECs and the user

sides. Alternatively, in order to reduce communication overhead, end-users can avoid direct contact with the SPC and communicate any information via the PEC of each end-user side.

5.1. Premium IP SLA/SLSs

The Premium IP SLA specification between GEANT and a European NREN domain comprises of two parts: the administrative/legal part and the SLS part [7].

The administrative/legal part defines the procedure and the framework for the provision of the service of which the SLA is established for. The proposed fields are:

- Administrative/technical parties involved: It specifies the administrative/technical contacts for each of the two sides participating in the SLA.
- Duration in time: It contains the period in which the service agreement is valid.
- Availability guarantees: It defines the calculation of the service's availability figures and how these will be derived.
- Monitoring: It specifies how will the provided service be monitored and where monitoring nodes should be installed.
- Response times: It defines the response times guaranteed by the provider in cases of client requests for adjustment of the SLA.
- Fault handling: It specifies the actions taken by the provider when quality guarantees are violated.
- Quality and performance of support and helpdesk: It specifies the contracted service's support infrastructure.
- Pricing of the contracted service: It determines the cost of the service.
- Description of the service: It is a general description of the service and its qualitative characteristics.

The SLS part contains most of the technical details of the agreement and consists of the following fields:

(a) Scope: It defines the topological region to which the service is provided.
(b) Flow description: It indicates which IP packets should receive the QoS guarantees.
(c) Performance guarantees: It depicts the guarantees that the network offers to the customer. Performance parameters for in-profile Premium IP traffic are proposed to be:
    - Maximum one-way delay for 99% of the service provisioning time.
    - Maximum inter-packet delay variation (IPDV) for 95% of the service provisioning time.
    - One-way packet loss for in-profile traffic.
    - Capacity guaranteed at egress points of the domain.
    - Maximum Transmission Unit (MTU) size for packets guaranteed not to be fragmented.
(d) Traffic Envelope and Traffic Conformance: It describes how the stream of packets should look like in order to get the guarantees indicated at the SLS. Traffic Conformance algorithms identifies packets as either 'in-profile' or 'out-of-profile'.
(e) Excess treatment: It specifies how out-of-profile traffic is treated, e.g. dropping.
(f) Service schedule: It indicates a time period, in months, for which the service is provided.
(g) Reliability: It defines the allowed mean downtime per year (MDT) and the maximum allowed time to repair (TTR) in case of breakdown.

6. MONITORING METHODOLOGY IN GEANT

As GEANT interconnects European NRENs, most of its traffic is transit, i.e. network connections are not originated or terminated in GEANT. For this reason, the monitoring infrastructure for the GEANT domain is planned to comply with the topology of Figure 3(b). Therefore, GEANT will be able to evaluate the network performance inside its boundaries and verify the compliance of the network behaviour to the signed SLAs. In addition, monitoring nodes must be installed next to each *Border Router* of the NRENs. These nodes will evaluate the performance of the NRENs access links. Obviously, the performance perceived by the end-users will not be explicitly measured or monitored in GEANT.

An open-source tool, *rude*/*crude*[‡‡], is suggested for building an open-architecture monitoring infrastructure in GEANT. Therefore, the monitoring infrastructure may be modified in the future in order to support measurements for new IP services or adapted to other network topologies. In addition, it can be extended to a point closer to the end-users, allowing end-to-end performance measurements.

We further suggest that monitoring nodes generate variable bit rate traffic consisting of two different packet sizes of 64 and 1500 bytes. A single packet size is considered adequate for measuring delay, jitter and packet loss in a network and thus evaluating Premium IP performance. However, for troubleshooting purposes two packet sizes may be required, since packet-forwarding mechanisms in routers may occasionally destroy packet ordering, favouring short packets versus longer ones. Thus, packet-reordering fraction (which, although not explicitly defined in SLAs, may cause service degradation) is more reliably measured by tests that inject at least two packet sizes.

It is difficult to define a generic guideline that specifies how often monitoring packets should be injected to the network so that the network performance is accurately assessed while the network is not overwhelmed by monitoring bandwidth overhead. Experiments in GEANT indicate that, one monitoring packet per second (a mix made from 75% 64-byte and 25% 1500-byte packets) should be injected into the network for each Mbps of bandwidth of Premium IP service delivered between two PoPs. According to GEANT Premium IP specifications, such service may occupy up to 10% of the 10 Gbps core network links [4]. Therefore, the overall overhead in each network link is expected to be less than 1% of the link capacity.

6.1. *Synchronization*

The time-related metrics for Premium IP service are one-way-delay and inter-packet delay variation or *jitter* [8]. However, only one-way-delay needs accurate network synchronization among the monitoring nodes located in distant, in terms of number of hops, PoPs. The reason is that jitter derives from the relative delay between consecutive packets as they are received by a monitoring node. As only the clock of the receiving monitoring node is used, jitter measurements are considered adequately accurate.

One-way-delay consists of three terms: switching delay, queuing—transmission delay and propagation delay. Switching delays in GEANT state-of the-art gigabit routers are almost negligible, approximately 60–200 ns. Also, as most of the backbone consists of 2.5–10 Gbps optical links, with utilization less than 30% (usually around 10–15%) the transmission and

[‡‡] http://rude.sourceforge.net/

queuing delays are also negligible, e.g. 1.6 ns is required for a 512-byte packet to be transmitted over a 2.5Gbps link. Therefore, the most significant term in one-way-delay measurements is the propagation delay, approximately 1 ms per 150 km.[§§]

GEANT plans to synchronize its monitoring infrastructure using the NTP protocol, which is a simple, inexpensive but not very accurate method. For improving the synchronization accuracy, a small number of primary NTP servers scattered in main PoPs of GEANT are suggested to use GPS receivers as an external clock. The NTP traffic is to be given high priority in order to improve synchronization accuracy. This will minimize the asymmetric delays that NTP traffic encounters traversing from the client to the server and backwards due to congestion–delay asymmetries in forward and reverse end-to-end paths.

In the case of GEANT, the distance between two end-users from different countries is usually hundreds of kilometers and the transmission speed decreases as we move closer to the end-users (the access lines are usually the bottleneck of the path). Therefore, the typical end-to-end delay that packets are encountering in the network is several tens of milliseconds, as also confirmed by tests in GEANT [9]. In such cases, a synchronization accuracy of 10 ms is adequate for monitoring the Premium IP service. However, there are exceptional cases where end-users from different (neighbouring) countries are connected to GEANT through high-speed access links while the distance between their countries' PoPs is few hundreds of kilometers. If this is the case, one-way-delay between the PoPs is expected to be only few milliseconds making compulsory the use of GPS receivers for the NTP servers. Based on the above constrains, the number and locations of the primary NTP servers with a GPS receiver have to be selected to satisfy the aforementioned accuracy. This activity is currently in progress by the GEANT community.

6.2. Service verification

We propose that data collected from the monitoring infrastructure are presented through dynamically generated graphs. The open-source *Multicast Beacon*[¶¶] tool is recommended as the basis for building the monitoring web-interface for Premium IP. Obviously, simple tools that calculate QoS performance metrics, such as mean or confidence intervals, should be integrated to the afore-mentioned tool.

The monitoring infrastructure is suggested to support alarm-triggering functionality. If an SLA contract is violated, an appropriate monitoring node should send notifications to the GEANT network operation centre and to the involved customer. In addition, the monitoring infrastructure may continuously compare the most recent network measured values with the long-term values, e.g. compare the last-five-minute and the last-one-hour mean values. If the evaluated results are outside an expected range, a warning notification should be sent to the corresponding operation centre.

In parallel to the GEANT monitoring infrastructure, end-users (R&D initiatives and/or local site administrators using IP Premium) may build their own end-to-end (non-mesh) monitoring infrastructure across multiple administrative domains, using similar open-source tools to GEANT.

[§§] The propagation delay between distant, in terms of optical path length, PoPs in GEANT is estimated in several tens of milliseconds (e.g. 60 ms round trip delay between Athens and London).
[¶¶] http://dast.nlanr.net/Projects/Beacon

Finally, a restricted access to the GEANT monitoring infrastructure is suggested to be provided to the end-users through the service verification interface. End-users should be allowed to initiate measurements between monitoring nodes in the backbone network in order to estimate the network performance. This feature of monitoring infrastructure will assist end-users in requesting a Premium IP SLA from GEANT that fits their needs.

7. CONCLUSIONS

The provision of QoS-enabled IP services in gigabit multi-domain networks has proved to be more difficult than initially expected. In peering interconnection models, where multiple agreements are established among ISPs, the provision of QoS-aware SLAs is inhibited by technical and operational problems of extreme complexity. In hierarchical networks, though, the provisioning of QoS-aware SLAs is more easily realized, due to the well-established network topology and the common policy that is enforced across the dominant backbone network and its constituent domains. Such networks are usually deployed to serve large communities with common objectives, e.g. education, research and public administration communities.

To illustrate potential problems in QoS provisioning and monitoring, even for hierarchical multi-domain networks, we reported on the experience of the pan-European Research & Educational gigabit network GEANT. We concluded that wide scale provision of Premium IP service (a 'leased line' type service) among requesting NRENs or large R&D multinational projects is feasible, but requires well-defined procedures to regulate establishment and verification of QoS-aware SLAs. Involved entities and their interaction need to be explicitly defined and the SLA/SLS contract to be standardized. Consequently, the success and further deployment of Premium IP in European NRENs and GEANT depends on the provisioning methodology followed, that should allow for efficient and reliable service provision to the end-users.

It is worth noting, that due to the unforeseen complexities in QoS-aware SLA management, vendors and providers invest on Layer 2 VPN technologies and protocols. These protocols, so far proprietary, try to re-introduce in IP networks the old ATM-based provisioning concept by using MPLS tunnels. Such mechanisms include the CISCO AToM® (Any Transport over MPLS) and the Juniper CCC® (Circuit Cross Connect). They plan to provide some QoS guarantees, on top of VPN addressing-security–billing isolation; in spite of their limitations and interoperability problems, they increase their presence in the rapidly growing VPN market. The GEANT–NREN community is actively investigating provision of Layer 2 VPNs as an alternative to IP Premium service.

ACKNOWLEDGEMENTS

This work greatly benefited from the IST research project SEQUIN. The authors would like to thank all the project participants from the NREN community and DANTE that contributed to the successful completion of the project. This paper was partially supported by the GEANT and SEQUIN European Commission Projects and the Greek Research & Technology Network GRNET. The views expressed by the authors do not necessarily reflect the opinions of the GEANT–SEQUIN community.

REFERENCES

1. Blake S *et al*. An Architecture for Differentiated Services, *RFC 2475*, December 1998.
2. Teitelbaum B, Shalunov S. Why Premium IP has not deployed (and probably never will). *Internet2 Qos Working Group Informational Document*, May 2002.
3. Roth R *et al*. IP QoS across multiple management domains: practical experiences from pan-European experiments. *IEEE Communications Magazine*, January 2003.
4. Campanella M. Implementation Architecture specification for the Premium IP service. *Deliverable D2.1-Addendum 1*, SEQUIN Project (IST-1999-20841).
5. Verma D. Supporting service level agreements on IP networks. *MacMillan Technology Series*, 1999.
6. ATM Forum. ATM Traffic Management Specifications Version 4.1, *af-tm-0121.000*, March 1999.
7. Bouras C, Campanella M, Sevasti A. SLA definition for the provision of an EF-based service. *Proceedings of the 16th International Workshop on Communications Quality & Reliability (CQR 2002)*, Okinawa, Japan, 2002; 17–21.
8. Campanella M *et al*. Quality of Service Definition. *Deliverable D2.1, SEQUIN Project (IST-1999-20841)*, March 2001.
9. Bouras C, Campanella M, Przybylski M, Sevasti A. QoS and SLA aspects across multiple management domains: the SEQUIN approach. *Future Generation Computer Systems*, vol. 939. Elsevier: Amsterdam, 2002; 1–15.

AUTHORS' BIOGRAPHIES

Athanassios Ch. Liakopoulos received the Dipl-Ing degree in Electrical and Computer Engineering from the National Technical University of Athens (NTUA), in 1996 and MSc with Distinction in Telematics (Telecommunications & Computer Engineering) from the Electrical Engineering Department in University of Surrey (UniS) in 1998. Since 2000, he joined the Greek Research and Technology Network S.A. (GRNET S.A.) and participated in several national and European research projects. He is also pursuing a PhD from NTUA in the field of QoS provisioning at high-speed networks. He has awarded for his performance during his academic studies and has published articles in recognized technical journals. His major areas of expertise are: IP networking, QoS provisioning and monitoring, routing protocols and optical networks.

Professor Basil Maglaris is a Professor of the National Technical University of Athens (NTUA) and the Chairman of the Board of the Greek National Research & Education Network GRNET. He received the Diploma in Mechanical & Electrical Engineering from the National Technical University of Athens (NTUA), Greece in 1974, the MSc in Electrical Engineering from the Polytechnic Institute of Brooklyn (now Polytechnic University), Brooklyn, New York in 1975 and the PhD degree in Electrical Engineering & Computer Science from Columbia University, New York in 1979. From 1979 to 1981 he was a research engineer at the Network Analysis Corp., New York, a leading firm in designing the ARPANET (the predecessor of Internet). From 1981 to 1989 he was with the faculty of Electrical & Computer Engineering of the Polytechnic Institute of Brooklyn, involved in teaching and research on computer networks. Since 1989 he is with the Department of Electrical & Computer Engineering at NTUA. Prof. Maglaris served in various academic and professional boards; from 1993 to 1995 he was the Managing Director of the National Hellenic Research Foundation (NHRF) and from 1996 until now he is the Chairman of the Board of GRNET. Since 1995 he serves as a Commissioner of the Greek National Regulatory Authority for Telecommunications. He participated in several R&D projects in the USA and in Europe, supervised nine graduate students that obtained their Doctoral Degree, authored more than sixty research papers and gave numerous talks in scientific conferences and other fora. Since 1992, he leads at NTUA a major development effort that resulted in a state-of-the-art integrated high-speed campus-wide Local Area Network and in 1996 he planned the development of GRNET.

Associate Professor Christos Bouras obtained his Diploma and PhD from the Computer Science and Engineering Department of Patras University (Greece). He is currently an Associate Professor in the above department. Also he is a scientific advisor of Research Unit 6 in Research Academic Computer Technology Institute (CTI), Patras, Greece. His research interests include Analysis of Performance of Networking and Computer Systems, Computer Networks and Protocols, Telematics and New Services, QoS and Pricing for Networks and Services, e-learning, Networked Virtual Environments and WWW Issues. He has extended professional experience in Design and Analysis of Networks, Protocols, Telematics and New Services. He has published 150 papers in various well-known refereed conferences and journals. He is a co-author of five books in Greek. He has been a PC member and referee in various international journals and conferences. He has participated in R&D projects such as RACE, ESPRIT, TELEMATICS, EDUCATIONAL MULTIMEDIA, ISPO, EMPLOYMENT, ADAPT, STRIDE, EUROFORM, IST, GROWTH and others. Also he is member of, experts in the Greek Research and Technology Network (GRNET), Advisory Committee Member to the World Wide Web Consortium (W3C), Task Force for Broadband Access in Greece, ACM, IEEE, EDEN, AACE and New York Academy of Sciences.

Afrodite Sevasti obtained her Diploma and her Master's degree (MSc) from the Computer Engineering and Informatics Department of Patras University in Greece. She is a PhD candidate at the same department of Patras' University. She has worked as an R&D Computer Engineer at the RA Computer Technology Institute (Greece) and she is currently with the Greek Research and Technology Network (GRNET) S.A. Her main interests and expertise lie in the fields of Computer Networks, Telematics, Distributed Systems and especially in technologies and architectures of high performance networks, in traffic and network resources' management, in Managed Bandwidth Services, provisioning of Quality of Service (QoS), SLAs and pricing/billing of next generation networks. She has published 20 papers in well-known refereed conferences and journals. She has participated in several R&D projects.