

Dynamic SLA-Based Management of Virtual Private Networks

M. D'Arienzo, M. Esposito, S.P. Romano, and G. Ventre

University of Napoli "Federico II", Via Claudio 21 – 80125
{maudarie, mesposit, spromano, giorgio}@unina.it

Abstract. This paper aims at shading some light on the concept of *Service Level Agreements* (SLAs) and on their usefulness in the context of the so-called Premium IP Networks. Such networks provide users with a portfolio of services, thanks to their intrinsic capability to perform a service creation process while relying on a QoS-enabled infrastructure. We will introduce a definition of SLAs and we will then focus on an actual example, namely the negotiation and management of QoS-aware Virtual Private Networks (VPNs). We believe that VPNs are a significant application due both to their importance in corporate scenarios and to the high revenues they guarantee to service providers. We will discuss the issues related to the effective SLA-based management of resources in those cases where the need arises for an entity that is capable of optimizing resource utilization in the presence of network infrastructures shared by a community of users. Finally, a novel component, named *SLA Manager*, that accomplishes these tasks, will be presented.

1 Introduction

QoS has been in the last years 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 when offering communication services have been subject to an intense debate that is still continuing in various fora (as for example the IETF).

It should be noted, however, that most of such activity has been focused on the technological aspects of the assurance of QoS within network elements, nodes and terminals. Most of the work has indeed been performed in the area of the mechanisms and architectures required for assuring that in packet switched networks data belonging to certain flows can be differentiated from others. This analytic approach has led to the definition of basic mechanisms and standards to be adopted within and across network elements. We might define such results as the basic elements for the provisioning of QoS "at a low level", or also *micro-QoS*.

In our opinion, however, this path towards technological development is missing a critical evolution factor, i.e. the definition of technologies for a system-wide approach to the provision of QoS-aware communication services. These technologies should be the ones responsible for the provisioning of QoS "at a high level" across complex network infrastructures with a real process and business model

oriented philosophy. We will define them as technologies for *macro-QoS* provisioning.

Currently, from a Network Operator point of view the task of creating a QoS-aware communication infrastructure is obliged to be simply that of assembling new, advanced network components and adding such new infrastructure to the existing ones, trying to maintain and re-use the existing business models and architectures. Manufacturers of network equipment introduce continually new solutions and products characterized by conformance to the existing or proposed standards and recommendations for what concerns the micro-QoS issues, but at the same time adopting specific and proprietary solutions for the provisioning of macro-QoS functionality.

The aim of this paper is to bring theoretical and practical contributions exactly in this area, with the goal of allowing the definition of tools, procedures and processes for the offering of advanced communication services in Premium IP networks across global infrastructures which might have a high degree of complexity in terms not only of scale, but also of the number of operators and level of technological heterogeneity.

The paper is organized in six sections. The reference framework where this work has to be positioned is presented in section 2. Section 3 discusses the main issues related to the definition, negotiation and activation of Service Level Agreements in Premium IP Networks. An actual example of the applicability of these new concepts is shown in section 4, where we will expand on the issues related to the negotiation and management of QoS-aware Virtual Private Networks (VPNs). This will allow us to introduce a model for a novel network component, named the *SLA Manager (SLAM)*, whose main objective is the dynamic, SLA-based management of corporate traffic (section 5). Finally, section 6 provides some concluding remarks, together with some information concerning our future work in this field.

2 Reference Framework

This section introduces the general architecture proposed for the dynamic creation and provisioning of QoS based communication services on top of Premium IP networks [1]. Such an architecture includes key functional blocks at the user-provider interface, within the service provider domain and between the service provider and the network provider. The combined role of these blocks is to manage user's access to the service, to present the portfolio of available services and to appropriately configure and manage the QoS-aware network elements available in the underlying network infrastructure. Their internal operations comprise activities such as authentication, aggregation and a mediation procedure that includes the mapping of user-requested QoS to the appropriate service/network resources, taking into account existing business processes.

In our view, network architectures are expected to be highly heterogeneous in terms of variety of systems and nodes, owing to the fact that they should be able to support dynamic service creation and service configuration on top of

generic QoS-aware IP networks. Closely related to this activity is the management of those resources in the underlying networks that are reserved at registration/subscription, as well as those that are used – and maybe subsequently modified – when the service is invoked/configured. Associated with the reservation and usage of resources is the automated production and presentation of the corresponding SLAs to the user and the translation from the SLA to the corresponding Service Level Specification(s) (SLS) [2].

To the purpose, we have introduced three major components [8] (figure 1) that we believe are needed to supervise the dynamic service creation and service configuration process:

- Access Mediator (AM);
- Service Mediator (SM);
- Resource Mediator (RM).




Fig. 1. The reference framework for Premium IP Networks.

The Access Mediator is the device into which users input their requests to the system. It adds value for the user, in terms of presenting a wider selection of services, ensuring the lowest cost, and offering a harmonised interface: the Access Mediator presents to the user the currently available services. The source of the services is a so-called “Service Directory” database, but the Access Mediator performs processing of the raw information. For example, it can select the cheapest offer if a movie is available from more than one service provider, and it can notify the user as soon as a new movie becomes available that matches the

stored user's profile. Its main role thus consists in assisting and easing the service selection process. This functionality may be under the control of a trusted third party and appears to offer excellent novel opportunities for a value-added service provider. The usage of a service generally involves two business processes: registration to be a user of the service and invocation of the service at the moment when it is used (any modification of the service parameters during a session can be considered as a new invocation). The following sequence of events is broadly applicable to both processes:

- after authentication, the user requirements are captured, and the Access Mediator sends the information to Service Mediators (which in turn employ the Resource Mediators) to map the requested and (subsequently) selected service into the deployed physical network;
- once the service selection has been agreed with all parties, the SLA is “signed” between the user and the Service Mediator;
- records of usage and the associated SLAs are stored in the Access Mediator for future reference.

A graphical user interface associated with the Access Mediator is expected to provide a harmonised interface to the user for all the available service offers.

The Access Mediator may form associations with one or more Service Mediators to which requests are issued. Generally off-line, the Service Mediator will supervise the incorporation of new services, their presentation in the “Service Directory” and the management of the physical access to these services via the appropriate underlying network, using the Resource Mediator(s). It is the task of the Service Mediator to prepare the SLA for the user to sign, and subsequently map the SLA from the Access Mediator into the associated SLS(s) to be instantiated in cooperation with the Resource Mediator(s).

The Service Mediator has an important role, as this is the place where services are created and from where the impacts of service reconfigurations are communicated to the network resource management. The Service Mediator has to inform the Access Mediators (usually via the Service Directory) of all new service offerings, so to allow them to present the updated portfolio to their users. It also has to check that the addition of a new service, or invocation of an existing one, will not affect the services that are currently operational.

In this scenario, a policy based approach is a possible solution to ensure the correct operation of the network. Subsequent to the service creation, a policy extension could be applied to the network to ensure that all services can be managed correctly. The system would have a global view of the configuration of the devices (including an accounting system) and of the policy rules to be applied. In such a case, it would be the function of the Service Mediator to update the service level management system with new rules and configuration as required, in conjunction with the Resource Mediators.

The communication between the Service Mediator and the Resource Mediator should be generic (i.e. independent of the technology employed by the underlying network). According to our design, it is the Resource Mediator that

will hold the current end-to-end view of the network QoS, by communicating with all the appropriate underlying network management systems. A network provider wishing to offer its resources should support an interface capable of handling messages defining an SLS, from its network management system to one or more Resource Mediator(s). The SLS templates we envision are in line with the descriptions in [2,3].

Since there can be more than one Resource Mediator, a Service Mediator can issue identical requests for information about network resource availability to several Resource Mediators. The Resource Mediators will either act on their own image of the network, or explicitly enquire to the individual network management systems, before returning an answer to the Service Mediator. The Service Mediator will accept the best offer, on the basis of the current policy decisions.

In order for the Resource Mediator to maintain and update its end-to-end network view of the current QoS availability, it may use a set of policy rules that are agreed with the underlying network management systems [4].

A common feature of the communication process surrounding the Access Mediator, Service Mediator and Resource Mediator components is a “one-to-many, search-and-selection” mechanism. In particular:

- the Access Mediator is responsible for selecting the appropriate Service Mediator(s), according to the user’s request;
- the Service Mediator is responsible for finding – and, in some cases, building from individual elements – the service, requesting information from (and then selecting) the appropriate Resource Mediator(s);
- the Resource Mediators are responsible for selecting the appropriate network capabilities, given several available options.

The potential for developing a common protocol for all of these similar actions is described in [7].

3 Service Level Agreements

In this section we address the issues related to the definition of the Service Level Agreements [6] suitable for the services envisaged in the framework of Premium IP (PIP) Networks. As we saw in the previous section, PIP networks are capable of delivering new services to the end-users. Such services are characterized by different levels of Quality of Service (QoS). In this scenario, the definition of a service creation framework [5] plays a major role, given its aim to dynamically create application-level services by appropriately combining and configuring pre-existing service components. Based on these considerations, we will provide definitions for SLA and a number of scenarios describing their relationship with the service creation and resource management framework. This set of definitions is part of a comprehensive conceptual model (SLA modeling framework), extensible in nature and capable to include sections specific to the different technologies or service architectures and cover SLA modeling for both transport services and end-user services.

3.1 Service

A service is a concept which may be modeled from different perspectives. Whether acquiring services in a retail or whole-sale mode (see below) a service provisioning/creation life cycle starts with a service description, which then requires a proper degree of formalization. Emphasis is on the different types of information that need to be modeled and on the needed ability to model relations between them.

3.2 SLA

An SLA is a contract between the customer and the provider of a specified service. Such a contract is signed upon subscription to the service itself. An SLA is prepared from templates specifically conceived for the available services.

SLA templates are used during customer negotiation to define the required level of service quality. The production of an SLA template is an intrinsic part of service development. These SLA templates may either relate to standard product/service offerings, where they are used “as-is” to define the required level of service quality, or provide a baseline for custom negotiation (either automated or human-assisted). SLAs are defined on something perceived by the customer (i.e. explicitly subscribed to), that is the service elements composing the service product offering.

3.3 Retail vs Whole-Sale SLA

The retail SLA refers to the agreement between an end-user and a service provider. The end-user might be either a single person or a user organisation (e.g. a corporate or a public institution). Such an end-user could be induced to establish a SLA with his provider in order to support different kinds of applications. Some of the applications which we deem worth investigating are, for example:

- Adaptive Multimedia Applications (e.g. Video on Demand, Video Conference, etc.);
- Voice over IP (VoIP);
- Virtual Private Networks (VPNs).

SLAs trigger the negotiation of hierarchical agreements between different contractors. In the case of multi-domain scenarios, service providers may need to create inter-network agreements in order to support their end-user SLAs. We call whole-sale SLAs these inter-provider contracts. A whole-sale SLA takes into account traffic aggregates flowing from one domain to another. In general, there is no direct connection between r-SLAs and w-SLAs. In particular, w-SLAs might not be based on parameters related to a single service but might focus on statistical indicators related to the Grade of Service of the entire bundle provided by one provider to one of its neighbors. The focus of this paper is mainly on retail SLAs.

3.4 Static vs Dynamic SLA

As already stated, an SLA is a contract between two parties. To date, the general trend has been to consider only static SLAs: the contracts are instantiated after negotiations by human agents and their terms cannot be modified during their lifetime. We do believe that dynamic features are needed in order to better match the requirements of real-world operational scenarios. We envision at least two different flavors of dynamic behavior:

- time-varying user requirements (with different time-scales of time variability induced by specific application characteristics);
- time-varying network conditions (of which the user is made aware via feedback signals raised by the network itself).

Both flavors may induce the end-user to change over time the terms of a pre-established SLA. Typical usage scenarios are shown below.

1. No time-varying user requirements, No time-varying network conditions:

in this situation the end-user establishes a static SLA, with no feedback. This implies that, once successfully terminated the negotiation phase, no modifications to the contract are allowed. In case of an admission control failure, the user is provided with no information concerning the reasons behind it. Thus, he has no clues on how to better re-formulate his request.

2. Time-varying user requirements, Time-varying network conditions:

this case refers to the most complex possible scenario, which requires re-negotiable SLAs. During the negotiation and usage phases, the network may provide the user with useful information for tuning his request. The contract may be re-negotiated at any time.

3. No time-varying user requirements, Time-varying network conditions:

in this scenario the network is capable of keeping users informed about its state, even if the users themselves cannot re-negotiate contracts on the fly. To exploit network hints they are compelled to tear down pre-existing contracts and re-formulate their requests from scratch.

4. Time-varying user requirements, No time-varying network conditions:

here, users feel free to change their contracts, according to their new requirements. It is obvious that in this case, users' requests may incur admission control failures, due to the absence of specific data concerning the current state of network resources.

Figure 2 summarizes the aforementioned scenarios.

		Variability of user requirements	
		No	Yes
Usage of feedback signals	No	Wholly static SLA	Pro-active multimedia multiparty applications
	Yes	Long distance IP telephony/ Adaptive Applications	“Pro-active/ Adaptive” applications

Fig. 2. Dynamic SLAs and related applications.

3.5 Content-Based SLAs

The service creation framework we depicted in the previous section envisages a scenario where users contact an Access Mediator (AM) in order to gain access to a number of value-added services, by means of negotiation of specific Service Level Agreements. The AM, in turn, needs to interact with one or more Service Mediators, each providing a certain set of services, to retrieve information about the characteristics of the services themselves. Afterwards, it organizes these information in order to let the user choose the service that most appropriately fits his needs. Once a specific service has been chosen, the involved Service Mediator(s) is (are) in charge of interacting with one or more Resource Mediators which, eventually, configure network elements so to efficiently satisfy the negotiated requests.

The process described foresees the generation of a number of documents (Service Level Agreement, Service Level Specification, policy rules), each describing the same instance of the service at a different level of abstraction and thus requiring creation/interpretation by the modules (Access Mediator, Service Mediator, Resource Mediator) belonging to the corresponding level of the overall architecture.

Digging into the details of such mechanisms, we can see in figure 3 that the Service Level Agreement is a contract between the end-user and the Service Mediator, negotiated via mediation of the Access Mediator. Once this contract has been signed, the Service Mediator is in charge of translating it into an appropriate Service Level Specification, containing a technical description of the service itself. This translation is a uni-directional process, requiring some additional in-

formation on the SM's side in order to retrieve, where necessary, service-specific data.

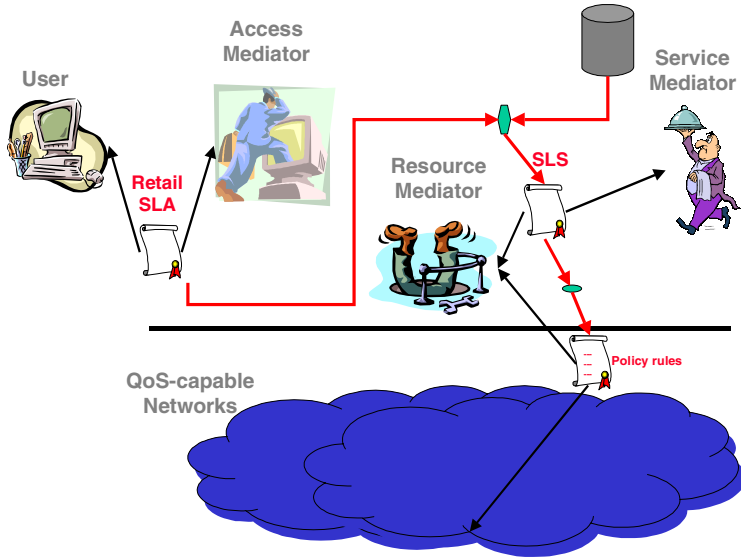


Fig. 3. SLAs, SLSs, policy rules.

The SLS is in turn given to the Resource Mediator, which translates it into a format that is the most appropriate for the QoS-capable network it manages. For example, it might build a list of policy rules, needed inside Policy Decision Points (PDP) in order to configure the underlying network elements (or Policy Enforcement Points - PEPs) via a policy protocol like COPS [4] .

3.6 Issues Related to the Definition of Service Level Agreements

As a general remark, an SLA should give the user the possibility to negotiate a certain type of service, among those offered by the network operator. We expect that most users will simply ignore the details of the service they expect from the network (especially those concerning the traffic characterization), either because such information is not available at all, or because they lack the motivation or the necessary technical skills required to understand their semantics. To ease the process of filling the contract template, a number of different SLA models might prove useful: the contract would become easier to understand, being focused on the actual needs expressed by the user. These SLAs may be considered as formed by two different parts, one containing information that does not depend on the particular application and the other containing application-specific data. The first might, for example, include:

- service level;
- user authentication module;
- information concerning availability/reliability of the service;
- encryption services;
- pricing and billing policies;
- options(enabling, for example, contract re-negotiation in case of unavailability of the required resources).

The second part of the agreement is analyzed in further detail in the next sub-section, devoted to examples on possible operational scenarios.

What we want to point out here is that from the network perspective the need arises to unambiguously specify all of the details and the characteristics of the service the user is willing to receive. SLAs should thus be translated into related Service Level Specifications (SLSs) containing all of the technicalities associated to the service itself. SLSs should be independent from both the high-level applications they stem from and the low-level network infrastructures on which they operate. Work is in full swing in the Internet community to define the main aspects related to SLS definition.

3.7 Example Operational Scenarios for SLAs

Interactive Multimedia Applications

Such applications include audio/video transmissions where a user connects to a video-server archive containing a number of movies that can be sent, in a streaming fashion, to a client host. In the same category we can also put those applications, like Video Conference and Tele-medicine, where video and audio data are generated from live sessions. For these applications, a mechanism is required to grant access to either the movie list or the session directory, in order to let the user choose the file/event he is interested in. The user is required to indicate the service he is willing to perceive, optionally specifying service lifetime. After defining these parameters, the translation module has to retrieve the traffic characterization associated to the specified files/sessions, in order to insert it inside the network SLS.

Virtual Private Networks

In this section we will consider only issues related to the creation of an SLA for a VPN service with respect to the problems linked to the provision of Quality of Service guarantees. We will therefore not cope with any aspect related to security or fault tolerance. We envision a scenario where a company, or in general an organization based on multiple facilities or sites, asks for the provisioning of a Virtual Private Network service as a way to efficiently interconnect its networking infrastructures. We expect, therefore, to see the VPN service forward traffic generated by a variety of users, network infrastructures, services, and applications. For example, we might consider a situation where two or more sites

of a company are connected via a VPN service so to have both data-like connectivity (LAN-to-LAN) and voice-like connectivity (VoIP interconnections). The SLA will therefore be related to the provision of services to a mix of traffic, with different requirements in terms of bandwidth and QoS.

In the case of VPNs, the r-SLA negotiated with the Service Mediator might not be as fine-grained as needed to accommodate the internal needs of the company/institution for which the VPN is being set up. Thus, the customer may apply further traffic management on its own premises. We will focus on these issues in the next two sections: more precisely, in section 4 we will propose an SLA template for VPNs, while in section 5 we will show how effective management of the VPN resources may be achieved thanks to the introduction of a novel component, called SLA Manager (SLAM).

4 SLAs for Virtual Private Networks

While designing a template for a Service Level Agreement for VPNs, we kept in mind the fact that users are looking for simple solutions to complex problems: that is to say, when buying an enhanced VPN service, they would just like to express their needs at a high level of abstraction, with no need to bother with all the underlying technicalities. Thus, we exploited the capabilities of the aforementioned mediation components in order to fill the semantic gap between the user's and the provider's perspectives on the same service.




Fig. 4. The “Drag&Drop” VPN.

We provide the user with a friendly graphical interface (figure 4) enabling him to design his own private network by simply drawing a graph representing the

sites he wishes to interconnect, together with the related tunnels (represented by the edges in the graph). Furthermore, for each tunnel, we envisaged the possibility to choose its specific features:

- type: mono or bi-directional;
- required bandwidth;
- desired quality;
- time schedule: permanent or time-dependent (either periodic or related to a specified time frame);
- authentication procedure;
- encryption algorithm.

This graphical representation of the VPN is then passed, in the form of an incidence-like matrix, to the Access Mediator, which in turn builds a formal description (e.g. in a language like XML) of the agreement the user is willing to negotiate (figure 5).

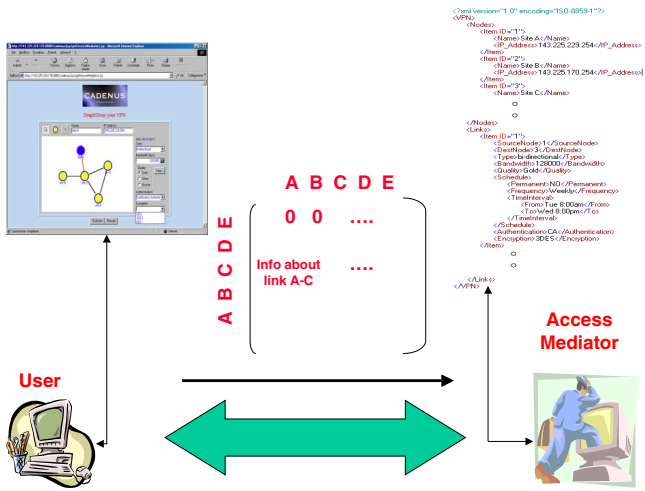


Fig. 5. User to AM communication of the VPN features.

This template SLA is the grounds on which lays the negotiation process involving all of the mediators and described in the previous sections. In particular, the Service Mediator, upon reception of this information, is in charge of performing the following tasks:

- retrieve information about links;
- for each link:
 - create an SLS, based on the Virtual Leased Line concept [2];

- compute the cost associated to this VLL, by contacting the involved Resource Mediators (this might require “SLS/policy rules” translation).
 - compute overall cost (e.g. sum of all link costs);
 - send cost back to the AM.

Figure 6 gives an idea of how the SLS for a single Virtual Leased Line might look like.

```
-Scope: one-to-one (143.225.229.254,143.225.170.254)

-Flow description: (143.225.229.0,143.225.170.0) DSCP=EF

-Traffic Conditioning: token bucket (b,r): r=128kbps

-Excess Treatment: dropping: only in-profile packets allowed

-Delay guarantee: qualitative (e.g. delay="low")

-Loss guarantee: p = 0 (implying a throughput guarantee R = r)

-Service Schedule: Weekly, 8am on Tuesday to 8pm on Wednesday

-Reliability: May be specified

-Options: Authentication (from CA),3DES encryption
```

Fig. 6. SLS for a Virtual Leased Line of the VPN.

After retrieving overall service costs, the AM is able to build the candidate SLAs and send them to the user: once the user has chosen a specific offer, it can store the associated SLA inside the user’s repository and notify SMs about the final decision, so to let them appropriately configure network resources.

It should be noticed that in this scenario, one SLA corresponds to multiple SLSs and only at the SLA level of abstraction the concept of a VPN does exist: the only component who speaks “Dragged&Dropped VPNs” is the AM. Furthermore, the SLSs are both service-independent and network-independent.

5 A Model for an SLA Manager

One of the issues to be considered in SLA based Premium IP networks is related to the possible complexity of the interactions between the users and the network entities responsible for the presentation and negotiation of SLAs. An SLA is a complex set of data, pertaining to different aspects of the provision of a network service or application: as already mentioned, in the case of a video delivery service, an SLA can include content dependent elements. In this business scenario,

the user is requested simply to accept the purchase of a service (for example, the delivery of a movie or of a multimedia document) and to ask the network (via the Access Mediator) for the provision of a communication service suited to the performance requirements of that specific content/service. Even though novel services could introduce new, more complex business scenarios, we expect that such kind of interaction between an end-user and the Access Mediator will remain quite simple, and perfectly manageable directly by the final user himself, for example via a web based interface.

However, we believe that there exist cases where such interactions could become much more complex, due to the presence of multiple technical and commercial aspects related to the nature of the offered service and to the specific needs of what we define the end-user. This is the case, for example, of the provision of a VPN service for the interconnection of multiple network infrastructures via public Premium IP networks.

Virtual Private Networks are being considered as the real killer applications for future networks. From the network provider's standpoint this is primarily due to the complexity of their setup and management. Users, on the other side, are mostly interested in the possibility of exploiting an infrastructure capable to provide services with guaranteed QoS.

In a QoS-capable network, a VPN service will be based on the idea of provisioning, on top of purely virtual links, a tailored communication service that could be differentiated from existing ones in terms of:

- performance and QoS;
- monitoring and accounting;
- security and privacy.

As far as the first issue, we expect that a VPN should offer a bouquet of services corresponding to the possible different QoS requirements that the data flowing across the VPN itself might demand, according to policy decisions that are — and should be kept — free to be defined and set by the user. Indeed, since VPNs are more and more linked to the need of creating and operating the so called “network companies”, one of the major issues will be dynamic management of the virtual communication infrastructure.

We introduced two differently-grained SLA types: (retail) r-SLA and (wholesale) w-SLA. The former is negotiated between an end-user and a service provider on a relatively small time scale, involving generally small resource amounts. The latter, instead, is negotiated between two different network domains less frequently than r-SLA, in order to create a pipe which merges different flows relating to r-SLAs already instantiated.

In the case of a users' organization, for example, an organisation representative signs an SLA with the access provider, reserving a large amount of resources. Then, single organisation members can apply for a smaller part of these resources up to the bulk available. Further requests have to be declined or a new contract has to be (re)negotiated by the organization.

In such a scenario the necessity arises for an entity that manages bulk resources, assigns sub-portions of them and is capable, in case of need, to (re)negot-




Fig. 7. The Service Level Agreement Manager.

iate larger quantities of resources. It can be considered as a mediator between a single user and a larger generic service provider (reached via an Access Mediator). This entity is in charge of negotiating with the Service Mediator a retail SLA that applies to the users’ organization as a whole. We call *SLA Manager (SLAM)* this entity (figure 7). A SLAM should have the following major functions:

- towards the single users:
 - AAA (Authentication, Authorization & Accounting);
 - internal negotiation of the r-SLA, in the form of what we call “ r^2 -SLA” (figure 8);
 - providing a friendly GUI;
 - enabling a user preference profile to be “bookmarked” for future use.
- towards the Access Mediator:
 - (re)negotiation of the r-SLA.
- from a more global perspective:
 - policing and/or shaping of single user’s flows;
 - admission-control;
 - providing a fair sharing of total resources;
 - triggering local network configuration.

The SLAM is therefore a new entity, placed between an existing Access Mediator and the end-user, but surely pertaining to the user’s domain.

We are currently prototyping a SLAM entity for VPNs, based on a model where QoS is taken into account at different levels of granularity inside the VPN




Fig. 8. The role of the SLAM.

tunnels. More precisely, a small and fixed number of traffic classes are considered in the core of the network (as delimited by the LAN gateways) and share the available bandwidth in a controlled way (exploiting, for example, a Priority Queuing algorithm [10]). Inside each class, a further level of discrimination is applied, by identifying the single micro-flows and deservng to each of them an ad hoc treatment via an additional scheduling (e.g. Weighted Fair Queuing [10]). Figure 9 gives an idea of how this scenario is provided.

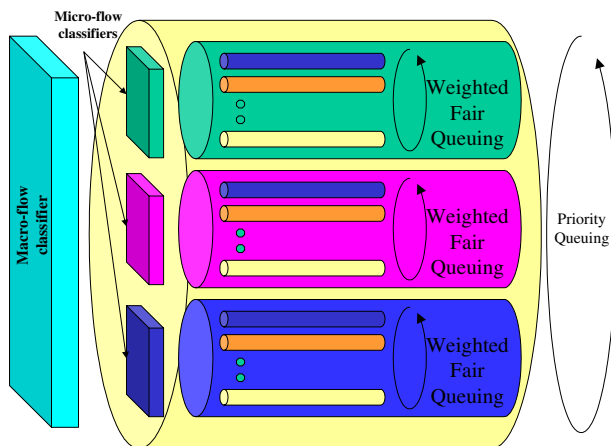


Fig. 9. Scheduling flows inside VPN tunnels

We implemented such scheduling algorithms on a programmable networks platform [9], where they may be combined in a structured fashion, thus bringing to complex configurations, ranging from simple series/parallel structures to nested ones. The situation depicted in the figure might relate, for example, to the case where all IP telephony flows are given priority over those generated by other applications (as guaranteed by the external scheduler) and are further discriminated among each other, by appropriately configuring the internal WFQ

scheduler (that acts as a “gate keeper” element). In this case, the SLAM is responsible for negotiating the r-SLA related to the bulk resources assigned to the VPN tunnel, while re-distributing them in a controlled fashion in the inner local network.

6 Conclusions and Future Work

In this paper we presented a novel approach to the design and development of a global architecture for the effective deployment of value-added Internet services upon Premium IP networks. As we saw, this is an ambitious goal, requiring a comprehensive understanding of all the procedures involved, from user-to-network interaction all the way through appropriate configuration of network devices, passing by a formal description of the service. Based on these considerations, we pointed out the need for a thorough definition of the concept of Service Level Agreements and associated Service Level Specifications. We then presented a model based on a modular decomposition of tasks involved in the deployment process, exploiting at their best the concepts of “mediation” and of recursive group communication. Finally, to give an idea of how this concepts apply in a real-world operational scenario, we focused on an actual example related to dynamic, SLA-based management of Virtual Private Networks.

This definitely represents the research field we are mostly interested in further investigating, due to the challenging topics it proposes. We are firmly convinced that the approach we presented may prove extremely useful for next generation Value Added Service Providers (VASP).

Acknowledgements

Research outlined in this paper is partially funded by the IST project CADENUS IST-1999-11017 “Creation and Deployment of End-User Services in Premium IP Networks”, <http://www.cadenus.org/>

References

1. IST project CADENUS (IST-1999-11017) “Creation and Deployment of End-User Services in Premium IP Networks”, <http://www.cadenus.org/>
2. D. Goderis et al, “Service Level Specification Semantics and Parameters”, `draft-tequila-sls-00.txt`, November 2000, expires March 2001.
3. Y. T’Joens et al, “Service Level Specification and Usage Framework”, `draft-manyfolks-sls-framework-00.txt`, October 2000, expires April 2001.
4. J. Boyle, R. Cohen, D. Durham, S. Herzog, R. Rajan, A. Sastry, “The COPS (Common Open Policy Service) Protocol”. IETF Proposed Standard RFC2748, January 2000.
5. M. Smirnov, “Service Creation in SLA Networks”, work in progress, Internet Draft `draft-cadenus-slan-screation-00.txt`, November 2000, <http://www.cadenus.org/papers/>

6. S. P. Romano, M. Esposito, G. Ventre, G. Cortese, “Service Level Agreements for Premium IP Networks”, work in progress, IETF draft `draft-cadenus-sla-00.txt`, November 2000, <http://www.cadenus.org/papers/>
7. M. Smirnov (Ed.), “SLA Networks in Premium IP”, CADENUS Deliverable D1.1, March 2001.
8. G. Cortese (Ed.), “Mediation Components Release 1 — Requirements and Architecture”, CADENUS Deliverable D3.1, March 2001.
9. R. Maresca, M. D’Arienzo, M. Esposito, S. P. Romano and G. Ventre, “An Active Network approach to QoS-enabled Virtual Private Networks”, submitted to 2nd Workshop on Quality of Future Internet Services (QofIS2001), 24-26 September 2001, Coimbra, Portugal.
10. S. Keshav, “An Engineering Approach to Computer Networking. ATM Networks, the Internet, and the Telephone Network”, *Addison Wesley Professional Computing Series*, 1996.