SLA definition for the provision of an EF-based service

Christos Bouras
Computer Technology Institute,
61 Riga Feraiou Str., 262 21 Patras,
Greece
and
Department of Computer Engineering and Informatics, University of Patras,
26500 Rion, Patras,
Greece

Mauro Campanella
INFN-GARR,
Via Celoria 16, 20133 Milano,
Italy

Afrodite Sevasti
Greek Research and Technology Network,
56 Mesogion Av., 11574, Athens,
Greece
and
Computer Technology Institute,
61 Riga Feraiou Str., 262 21 Patras,
Greece

***Abstract*-The evolution of mechanisms for providing Quality-of-Service (QoS) over the contemporary network infrastructures has introduced the need for regulation and management of the emerging QoS services with the use of Service Level Agreements (SLAs). SLAs for QoS-enabled networks move one step forward in the direction of traditional ones, in the sense that they do not only have to specify availability, security, quantity of allocated resources and a number of other quantitative values but also have to specify the values of appropriate quality parameters. This paper deals with the particular cases of introduction of QoS mechanisms to large transport networks according to the DiffServ architecture. In these cases, the extensive level of aggregation of flows and the connectionless nature of QoS services' provisioning makes the definition of QoS parameters and the engineering of QoS metrics in the traditional SLA specification a demanding task. Due to the fact that strictly bounded deterministic guarantees are not realistic, usually only upper bounds for the relevant parameters can be defined and the corresponding SLAs have to be defined accordingly.**

I. INTRODUCTION

There exist several efforts towards the standardization of definition of SLAs and their instantiation in QoS enabled networks ([4], [10], [9], [1]). This paper outlines a suggested template for the SLAs used to provide a QoS service called 'IP Premium' ([3]) from a backbone network (GEANT- the Next Generation of pan-European Research Network, [12]) to all its peering domains (the National Research Networks-NRENs). However, this template is general enough to be applied between any peering DiffServ-enabled domains in order for the regulation of QoS services' provision.

The IP Premium service itself was based on the Expedited Forwarding Per Hop Behavior (EF-PHB, [6]) of the DiffServ architecture and was defined in the framework of the GEANT and SEQUIN IST projects and more specifically in [3] and [8]. The implementation architecture for the Premium IP service aims at offering the equivalent of an end-to-end Virtual Leased Line (VLL) service at the IP layer across multiple domains.

SLA specification for QoS enabled networks aims at providing positive quality guarantees and setting out the limits of the services provided. In networks where QoS is inherently supported (such as ATM) the provision of SLAs comes as a natural delimitation of the relevant parameters. However, in IP networks where best-effort traffic has no quality guarantees, the introduction of QoS and associated services requires a thorough and accurate engineering of QoS metrics in the SLA specification on top of the guarantees for availability and characteristics of the transport medium, security, fault handling etc.

The analytical computation of such metrics is extremely complex taking into consideration the extensive level of aggregation and more generally the nature of traffic flowing in large interconnection domains. Usually only upper bounds for the relevant parameters can be defined. Therefore, SLA specification for QoS enabled networks becomes a process where intensive testing and probing of the available infrastructure has to take place, before being able to quantify the QoS offering and include concrete parameters and values in the agreement. Also, during the operation of the service monitoring its behavior is crucial.

In section II of this work, we are proposing an SLA structure and a series of fields that support the provision of IP Premium service over a DiffServ enabled set of interconnected domains. In subsection A, the administrative/legal part of the SLA is described, while in subsection B the technical SLS part is outlined. Section III concludes this work and specifies out future steps. The proposed model is currently being tested in order for its verification across GEANT and several NRENs in Europe.

II. SLA STRUCTURE

The 'IP Premium' SLA specification between GEANT and an NREN X is proposed to comprise of two parts:

- The administrative/legal part
- The Service Level Specification (SLS) part, defining the set of parameters and their values, for the provision of IP Premium service to a traffic aggregate by a DiffServ domain.

After the definition of the SLA specification, several instantiations of it will have to be produced, one for each of the several peering AS couples involved in the GEANT-NRENs architecture. Next steps will be to define the mechanisms for SLA negotiation and, of course, for the establishment of end-to-end services based in the individual SLAs. Each instantiation of a SLS will comprise a so-called Service Level Object (SLO) and will contain the parameters and their values that describe the transport service a specified flow is to receive over the transport domain.

Bi-directional services will also be possible by the combination of two Service Level Objects that will be taken atomically when negotiating a service pertain to two flows, one at each direction. These SLOs will comprise a Transport Service, which will be part one SLA defined between the two AS, among which the bi-directional service is established. Fig. 1 displays an SLA template and two instantiations of it, bringing the aforementioned individual SLA components together. The SLA instantiation on the left is an example of a bi-directional SLA containing two uni-directional SLOs.

A. *The administrative/legal part of the SLA*

The administrative/legal part of the SLA is suggested to comprise of a number of fields that will define the procedures and framework for the provision of the service that the SLA is established for. Proposed fields are:

- Administrative and technical parties involved: This section should contain at least one administrative and one technical contact from each of the two sides participating in the SLA.
- Duration in time: This section should contain the period for which the SLA is valid. This period can differ from the period defined at the 'service schedule' field of the SLS part of the SLA, but the value of the 'service schedule' field has to be a period within the period defined at this section of the SLA. The 'service schedule' is a set of time periods for which the service is active, while the SLA duration is a time period for which the SLA for the service's provision is valid.

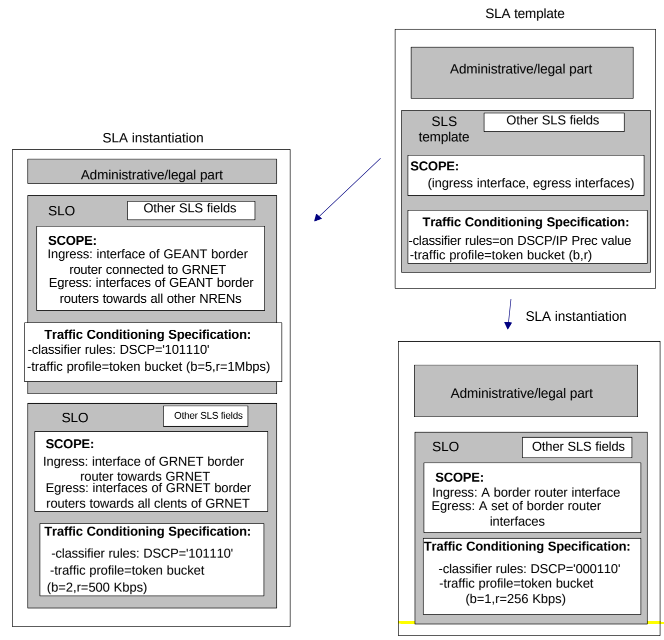

Fig. 1. SLA template, SLA instantiations, SLS and SLOs

- Availability guarantees: This section should define the calculation of the service's availability figures and how these will be derived (e.g. from the trouble ticketing system). The section should also provide a service availability ratio according to the SLA's duration in time in comparison, an Unavailable Time Limit (UTL) and formulas for the calculation of compensation for unavailability
- Monitoring: This section should specify how and when (constantly vs. periodically) will the SLA be monitored. It should specify the points of network topology where monitoring equipment is installed or where measurements are retrieved from. It should also specify the SLS metrics that will be visible to the client and how the client will have access to this monitoring data.
- Response times: This section concerns the overall response times guaranteed by the provider in cases of client requests for adjustment of the SLA (and/or SLS) and for necessary configuration of the relevant devices
- Fault handling-trouble ticket: This section should specify the actions taken by the provider when faults concerning the delivery of the service defined in the SLS occur and the reaction times.
- Quality and performance of support and helpdesk: This section should thoroughly specify the contracted service's support infrastructure.

- Pricing of the contracted service: Pricing of the service provided is a crucial part of a SLA between a client and a provider of network services. In order for a differentiated services pricing scheme that will efficiently reflect the service value and will maximize or meet several criteria (client revenue, efficient resource allocation, accepted service requests etc.) a very thorough and interdependent with the SLA monitoring and accounting infrastructure has to be used.

- Description of the service: A general description of the provided service, describing qualitatively its characteristics (in terms of e.g. delay, packet loss, throughput) and operation has to be provided here.

B. The SLS part

The proposed SLA template applies directly to the case where a transport domain establishes agreements for the provision of connectivity services with its customers in a uni-directional manner. Based on this assumption, the SLS part of the SLA is proposed to contain the following fields:

(i) Scope: The scope field should define the topological region to which the IP Premium service defined at the SLS will be provided. This field, according to [11], must specify where packets conforming to the SLS are entering and exiting a DiffServ domain. The recommended field is:

(ingress interface of upstream domain, set of ingress interfaces of downstream domains)

where the 'ingress interface of upstream domain' will specify the interface of a GEANT border router to which NREN X is connected and the 'set of ingress interfaces of downstream domains' will specify the set of ingress interfaces where packets injected to the GEANT from X can enter other customers.

(ii) Flow description: The flow description field will indicate for which IP packets the QoS guarantees of the specific SLS is to be enforced or in other words, which packets will receive the PHB treatment resulting in the QoS guarantees of the SLS. The flow descriptor is suggested to be the DSCP (Differentiated Services CodePoint) or IP Precedence value for the provision of the IP Premium service to aggregates injected by NRENs to GEANT. The DSCP or IP precedence value can uniquely identify the packets to receive IP Premium treatment among all packets injected from the NREN to GEANT, provided that the packets have already been through control admission in the NREN. However, additional information, already present in the packets or derived from the network topology, can optionally be included in the flow description field.

The IP Premium definition under consideration supports aggregation policing according to the packet’s destination NREN domain, and therefore classification of IP Premium packets must be extended to further granularity among different policers. Thus, the destination NRENs’ identities are recommended as part of the flow identification field, which becomes:

(QoS tag attribute,[source attribute], [destination attribute])

The source attribute is a unique value representing NREN X and the destination attribute is a set of values representing destination NRENs. The ‘QoS tag attribute’ is strongly recommended to be a global DSCP for all SLSs between GEANT and attached NRENs.

As pointed out in [11], the IP routing scheme MAY put restrictions on combining scope and flow description within an SLS. It might not be possible for a flow described by the flow description field to be serviced within the scope defined in the scope field e.g. traffic from NREN a to NREN b is routed via some specific egress interfaces of GEANT that must be included in the scope otherwise support of IP Premium traffic between the two NRENs will not be possible.

(iii) Performance guarantees: The performance guarantees field depicts the guarantees that the network offers to the customer for the packet stream described by the flow descriptor over the topological extent given by the scope value. In the case of SLS between an NREN and GEANT, this field provides seamless quality guarantees for the aggregate IP Premium flow injected from the NREN to GEANT.

The suggested performance parameters for in-profile traffic in the case of IP Premium and their respective values in the case of GEANT (see [13] for more details) are:

- One-way delay: It is suggested to be guaranteed as the maximum packet transfer delay between the scope-defined points measured. Indicative values are (26,.., 40) ms. The distance delay can be roughly computed using a signal speed of about 7 us/Km. A quintile could also be optionally defined to specify the delay guarantee in 99% of the cases, since users might find the worst-case figure misleading.

- IPDV: It is suggested to be guaranteed as the maximum packet transfer delay variation measured between the scope-defined points. Indicative values are equal or less than 2.351ms. Again a quintile is also defined to specify the IPDV guarantee in the majority (95%) of cases.

- One-way packet loss: It is suggested to be guaranteed as the ratio of lost in-profile packets between the scope endpoints and the injected in-profile packets at the ingress defined by the scope field. Indicative value is 10^-4 for 95% of the cases. It is suggested that the appropriate numbers will be based on the actual contracted values for the transmission lines and modified (increased) to take into account the service-induced figures.

- Capacity: It is defined as the rate measured at the set of egress points (defined by the scope field) of all packets identified by the flow descriptor. A suggested value for the IP Premium aggregate is 5% of ingress capacity. For the NREN-GEANT SLS it is suggested that this capacity is distributed to a guaranteed throughput vector of values corresponding to traffic from the NREN under consideration towards all other NRENs. The value of each vector item is the throughput guaranteed by the SLS for traffic originating at the NREN under consideration and terminating at a different NREN each time, through the interconnecting network. Each one of these values is used to calculate the rate parameter of each macro-aggregate policer at the ingress interface (see next SLS field).

- MTU (Maximum Transmission Unit): It is the largest physical packet size in bytes that the SLS guarantees to be transmitted without being fragmented. The suggested value for a WAN is 4470 bytes.

(iv) Traffic Envelope and Traffic Conformance: The traffic envelope is a set of traffic conformance (TC) parameters describing how the stream of traffic from NREN X towards GEANT should look like in order to get the guarantees indicated by the performance parameters of the SLS. The traffic conformance algorithm itself is part of the SLS, describes how is traffic examined against the targeted/contracted behavior and has as its input the traffic conformance parameters. It is possible to have either a binary-based or a multi-level based TC algorithm, but in the case of IP Premium, a binary-based algorithm identifying packets as either 'in-profile' or 'out-of-profile' is appropriate.

The IP Premium service aims at offering the equivalent of an end-to-end virtual leased line at the IP layer. Therefore, the conformance parameters are conformance to a shape and a limit of throughput/capacity. The traffic conformance algorithm adopted is that of token bucket with b as the depth and r as the capacity parameters.

Within the SLS framework, the following specification for the traffic envelope and conformance field is proposed:

- Conformance parameters = *(b, r)*
- Conformance algorithm = the *(b, r)* token bucket

In the particular case considered here, that is SLSs between an NREN and GEANT, the following values are suggested:

b = f(number of router interfaces on the same router that are part of the service, distance from the source)

$$r = \{1.2 .. 1.5\}* r_C$$

where r_C is the contracted capacity as defined in the 'Performance guarantees' field of the SLS.

However, as it has already been mentioned, policing of traffic in an NREN-GEANT SLS is not performed uniformly to the IP Premium flow aggregate injected by NREN X to the GEANT ingress. Instead, the entire flow aggregate is divided into several macro-flow aggregates according to the NREN the packets are heading towards and each macro-flow is policed by a different policer obeying to the conformance parameters and algorithm already specified, but with a different r_C value. Actually the r_C of each individual policer is derived from the throughput vector defined at the 'Performance guarantees' field of the SLS.

The suggested policing is thus for the macro-flow of each (NREN X, egress NREN) pair using a capacity value between 1.2 and 2 times larger than the contracted value between them and a token bucket with depth of at least 5 packets or more. Experimental evidence ([13]) has demonstrated that such a configuration results in minimum packet loss for IP Premium traffic.

(v) Excess treatment: This attribute specifies how excess traffic (or out-of-profile traffic, according to the profile described by the traffic envelope and traffic conformance field) is treated. For the purposes of IP Premium dropping of out-of-profile packets is suggested.

(vi) Service schedule: This field indicates the start time and end time of the period for which the service is provided. It is suggested to be of month range, either a single month or a group of sequential months.

(vii) Reliability: Reliability should define:

- allowed mean downtime per year (MDT)
- maximum allowed time to repair (TTR) in case of breakdown

for the provision of the service described by the SLS. The values of these guarantee parameters must be compliant with the guarantees provided via the administrative part of the SLA.

(viii) User visible SLS metrics: According to the SLA, NREN X must be aware of the available to him IP Premium bandwidth, one-way packet loss, burstiness of one-way packet loss, IPDV and one-way delay.

C. End-to-end SLAs

SLA definition between two peers is the structural unit for the establishment of end-to-end services. Provided that SLAs are properly defined all the way from the desired origin to the desired destination, proper mechanisms (such as the Bandwidth Brokers, see also [4], [5], [2], [7]) can evaluate all connections between consecutive peers and determine the resources (according to the SLAs) that are available for servicing requests for the specific service. This procedure can successively conclude with a valid outcome on whether the end-to-end service can be provided or not, based on the individual SLAs, and which are the specific quality features of the service provided.

III. CONCLUSIONS

This work contributes to the direction of defining end-to-end EF-based, VLL-like services over a DiffServ-based architecture of interconnected domains by specifying one ring of the end-to end chain: the SLA template for bilateral management of such a service between a backbone and its peering domains. Our future work on SLA definition for the IP Premium service will focus on further elaboration of different scenarios for the establishment of end-to-end SLAs in operating conditions. We also intend to focus on testing and calibration of the parameters that comprise the ‘Performance guarantees’ and ‘Traffic Envelope and Traffic Conformance’ fields of the proposed SLA template.

REFERENCES

[1] S. Salsano et al., “Definition and usage of SLSs in the AQUILA consortium”, draft-salsano-aquila-sls-00.txt, Internet Draft, November 2000.

[2] G. Fankhauser, and B. Plattner, “Diffserv Bandwidth Brokers as Mini-Markets”, Workshop on Internet Service Quality Economics, MIT, US, Dec. 2-3, 1999.

[3] M. Campanella, “Implementation architecture specification for an Premium IP service”, Deliverable D9.1 - Addendum 1, GEANT Project (IST-2000-26417), July 2001.

[4] “QBone Bandwidth Broker Architecture”, Work in progress, accessible at: http://qbone.internet2.edu/bb/bboutline2.html

[5] R. Neilson, J. Wheeler, F. Reichmeyer, and S. Hares (editors), “A Discussion of Bandwidth Broker Requirements for Internet2 Qbone Deployment”, Internet2 Qbone BB Advisory Council, Version 0.7, August 1999.

[6] V. Jacobson, K. Nichols, and K. Poduri, “An Expedited Forwarding PHB”, RFC 2598, June 1999.

[7] K. Nichols, V. Jacobson, and L. Zhang, “A two-bit differentiated services architecture for the Internet”, RFC 2638, July 1999.

[8] M. Campanella, P. Chivalier, A. Sevasti, and N. Simar, “Quality of Service Definition”, Deliverable D2.1, SEQUIN Project (IST-1999-20841), 31 March 2001.

[9] R. Rajan, E. Celenti, and S. Dutta, “Service Level Specification for Inter-domain QoS Negotiation”, draft-somefolks-sls-00.txt, Internet Draft, November 2000.

[10] D. Goderis et al., “D1.1: Functional Architecture Definition and Top Level Design”, TEQUILA Project: Traffic Engineering for Quality of Service in the Internet, at Large Scale, (IST-1999-11253), 11 September 2000.

[11] S. Blake et al., “An Architecture for Differentiated Services”, RFC 2475, December 1998.enough to be applied between any peering DiffServ-enabled domains in order for the regulation of QoS services’ provision.

[12] GEANT: The pan-European Gigabit Research Network, more info at: http://www.dante.net/geant/index.html

[13] A. Sevasti, ‘SEQUIN results on QoS’, TERENA Networking Conference 2002, Limerick, Ireland, June 3-6, 2002 (to appear).