A Distributed Admission Control Model for Class-based Networks using Edge-to-Edge QoS and SLS Monitoring

Solange Lima, Paulo Carvalho, Alexandre Santos, Vasco Freitas
University of Minho, Department of Informatics, 4710-059 Braga, Portugal, e-mail: {solange,paulo,alex,vf}@uminho.pt

Abstract - The advent of class-based networks has brought new needs for network traffic control in order to assure a certain QoS level. Despite the existing proposals, achieving a generic admission control (AC) strategy for traffic entering these networks is still an open issue. This paper provides new insights on how AC shall be accomplished proposing an encompassing AC model for multi-service class-based networks, which covers both intra-domain and end-to-end operation, without requiring changes in the network core and complex AC signaling. For each service type, AC is distributed and based on both on-line edge-to-edge monitoring of relevant QoS parameters and SLSs utilization. Service monitoring, performed at egress nodes, provides adequate metrics to ingress nodes which take implicit or explicit AC decisions based on service-dependent criteria. Although being oriented to flow AC, the model can easily be applied to SLS AC. SLS auditing and SLS traffic conditioning are tasks also covered.

Keywords - QoS architectures, Network Traffic Control, Admission Control, SLA/SLS, QoS Monitoring.

I. Introduction

The integration of present and emerging applications and services in the Internet, with different quality of service (QoS) requirements, has been fostering the evolution of the network architecture and protocols. Class of service (CoS) networks, where Diffserv architecture [1] is a reference model as regards QoS provision in the Internet, will soon be widely deployed. Although for few ISPs overprovisioning is an attainable solution, some form of network traffic control will have to be in place so that QoS requirements can be honored. In this context, the standardization of Service Level Specification (SLS) between domains and the definition of Admission Control (AC) strategies assume a relevant role. Despite the existing proposals, achieving a generic, yet feasible and light, AC model supporting different service types and covering both intra and inter domain AC is still an open issue.

This paper provides new insights on how AC shall be accomplished without adding significant complexity to the network control plane. The underlying idea is to take advantage of the consensual need for QoS and SLS monitoring in CoS networks and use the resulting information to perform AC. The AC model also takes into account (i) the network services to be provided; (ii) the level of assurance of each service; (iii) the overhead and scalability of the control strategy; and (iv) the easiness of integration in the Internet.

Resorting to edge-to-edge monitoring of relevant QoS parameters for each service type, the proposed model considers not only the service availability in a domain but also the sharing of SLSs between domains. AC is distributed and performed at domain ingress nodes, using service-dependent AC criteria. The QoS and SLS monitoring, which provides metrics for AC, are performed at egress nodes. Therefore, the network core is treated as a black box, without any changes. The model covers intra and inter-domain operation (end-to-end), while reducing AC latency, inter-domain signaling and state information. Although being detailed to flow AC, this model can be adopted for SLS AC. The additional tasks required at the provider side for SLS Auditing (SLS Monitoring and Conformance Verification), and SLS Traffic Conditioning are also described.

Considering that, from an end-to-end perspective, different QoS solutions are expected to be in place, the proposed model can be easily applied to distinct QoS scenarios and adjusted to technological, service and application evolution.

This paper is structured as follows: current AC approaches are reviewed in section II, while SLS definition issues are discussed in section III. The proposed AC model, which includes explicit and implicit AC, is detailed in section IV. SLS management and monitoring issues are discussed in sections V and VI, respectively. Although implementation aspects have been matter of concern when defining the AC architecture, specific implementation details are not covered here.

II. AC Approaches: a service oriented overview

Either in flow-based or class-based QoS architectures controlling the admission of traffic entering the network allows to: (i) avoid over-allocation of existing network resources; (ii) avoid new flows from impairing flows already accepted; (iii) fulfill service level agreements; (iv) prevent instability and assure QoS. Despite its need, the complexity introduced by AC has to be carefully assessed as Internet traffic is highly dynamic and not every application has strict QoS requirements. A lesson learned from the AC model used in Intserv [2] is that keeping resource reservations per flow in all network nodes, although allowing high QoS guarantees, is not a scalable solution. Aggregating these reservations [3] reduces the problem but does not solve it. Associated with CoS-based architectures, such as Diffserv, new AC approaches have been defined, avoiding per-flow state information in the core. Some proposals suggest the use of central entities for AC and resource management (bandwidth brokers) [4], [5]. However, the well-known problems of centralization led to many decentralized approaches to AC.

Generically, either using centralized or decentralized AC, the level of guarantee to be provided determines the complexity of the underlying traffic control strategy. For instance, to provide quantitative service guarantees (e.g. for hard real time traffic) current AC proposals need to control the state and the load of traffic aggregates in the core nodes [5], [6], [7], or even perform AC in these nodes [6], [7]. These solutions tend to require significant network state information and, in many cases, changes in all network nodes. Furthermore, as they are closely tied to network topology and routing, their complexity increases with the network dynamics.

Providing qualitative service guarantees (e.g. for soft real time) leads to reduced control information and overhead, but to eventual QoS degradation. To obtain a good compromise between efficient resource utilization and guarantee or predictability of QoS is a major challenge. In this context, measurement-based AC (MBAC) solutions have deserved special attention. Initially performed in all network nodes [8], recent studies suggest that AC should be carried out only at the edges (end-systems or edge routers), using either active (EMBAC) or passive measurement strategies of network load and/or QoS parameters [9], [10], [11], in order to keep the core simple. Despite not requiring changes in the network, EMBAC increases the initial latency and network load as probing is carried out on a per application basis. The use of routing protocols to propagate QoS metrics to edges is also a solution [12].

The need to control elastic traffic, for more efficient network utilization, has also been discussed and implicit AC strategies have been defined [13]. This means that no explicit signaling between the application and the network is needed. Conversely, AC approaches for streaming applications usually assume a form of signaling between the application and the network, where upon a traffic profile and QoS objectives description the network sends an explicit acceptance/rejection message.

In any AC strategy the admission criterion plays a crucial role as regards service guarantees and network efficiency. There are more or less conservative proposals [14], [8], which consider the estimation and control of parameters such as available bandwidth, delay, loss or ECN marks. Most of AC approaches only control the available bandwidth. Although being simple for a single link or node-by-node AC, controlling it in the full path is not straightforward. Methodologies and tools for estimating the available path capacity and available bandwidth are [15], [16]. The accuracy and on-line utilization relevance of these tools have been discussed in [17].

A complete survey comparing the main features and limitations of current AC strategies is available in [18].

III. Defining a Service Level Specification

A Service Level Agreement (SLA) is defined as a contract between a customer and a service provider or between service providers, specifying the expected service level. The technical part of an SLA - called Service-Level Specification (SLS) - describes the appropriate QoS-related parameters.

The definition of SLSs is a key aspect for QoS provisioning. A standardized set of SLS parameters and semantics is crucial for end-to-end QoS delivery and for simplifying SLS negotiations. Several working groups are committed to SLS definition [19], [20]. Taking these works into account, a possible SLA template including relevant QoS-related parameters, their typical contents and general meaning is defined in Table I. Although a large combination of performance and reliability parameters is possible, it is expected that service providers will offer a limited number of services. Instantiate the SLS template in quantitative and qualitative standard services adapted to different application types is, in fact, the major objective. To fulfill this, substantial work has been done on identifying the relevant QoS parameters and of the perceived quantitative quality of applications [21], [22]. Table II summarizes the QoS parameters upper bounds for common applications and services.

ITU-T work on QoS in IP networks and particularly the IETF IPPM working group have defined a set of standard metrics for QoS and performance measures and proposed measuring methodologies for them [22]. Several tools have also been developed and tested to measure SLS metrics [23], [17].

These inputs will be considered to identify a complete set of services, QoS parameters and measurement methodologies to be used in the proposed AC model. This process will also take into account Diffserv PDB definitions [24].

IV. Proposed Admission Control Traffic Model

A. Admission Control Perspectives

When AC takes an SLS as reference, two AC perspectives can be considered: (i) flow AC, when the admitted flows share an SLS; or (ii) SLS AC, when the admitted SLSs share a service (see Fig. 1). Although these are two distinct levels of AC, they use a similar principle to handle different objects. Whereas flow AC is based on the traffic profile and QoS objectives of a flow, SLS AC is based on the aggregate traffic profile and QoS objectives of the SLS. In fact, the semantic of the process is equivalent, only the granularity upon which the decision is taken changes. Therefore, the proposed model can be oriented either to flow AC or SLS AC, with minor adjustments.

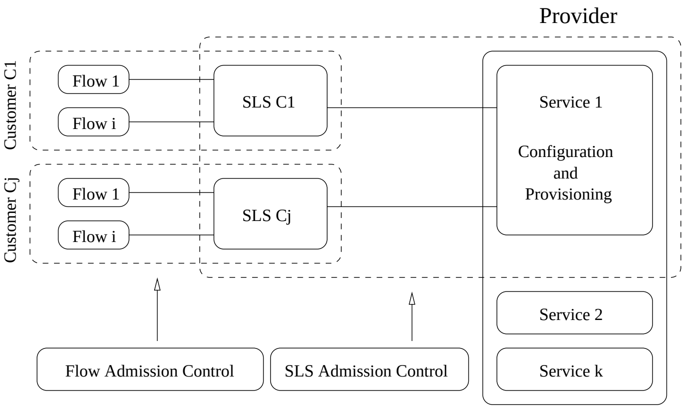

Fig. 1. Flow AC and SLS AC

Administrative Information		
Administrative entities		Contractual parties involved
Description of service		Description of service behavior
Validity		Contract validity period
Pricing/Tariffs		Pricing and tariffs of the service
Helpdesk info/Trouble tickets		Customer support actions
Monitoring/Accounting reports		Monitoring/accounting rules
Response time to changes		Time for enforcement of changes
Other		Other rules: e.g. provisioning
SLS		
Scope	- Ingress interfaces - Egress interfaces	Boundaries of the region over which the service will be enforced
Classifying rules	- MultiField criterion - DSCP or ToS Precedence	Packet fields used to identify a traffic flow or aggregate
Traffic Cond. rules	- Conformance algorithm - Conformance parameters - Treatment on excess	Information used to identify in-profile and out-of-profile traffic and corresponding treatment
QoS/Perf. parameters	- Delay, jitter, loss,... - Qualitative objectives - Quantitative objectives	Expected QoS of the conforming traffic stream in the Scope region
Reliability	- Mean downtime - Time to repair,...	Expected service reliability
Service scheduling	- Start/End time	Service time availability
Others	- Route, security, ...	Left for future study

TABLE I

Service Level Agreement (SLA) template.

ITU-T classes	Class 0	Class 1	Class 2	Class U	
Applications	Real-time	VoIP / Interact.	Non-Interac.	WWW / Free Serv.	Stream. video (VHS)
IPTD	150 ms	400 ms	1 s	Undefined	400 ms
IPDV	50 ms	50 ms	1 s	Undefined	17 ms
IPLR	10^{-3}	10^{-3}	10^{-3}	Undefined	10^{-5}
IPER	10^{-4}	10^{-4}	10^{-4}	Undefined	10^{-4}

TABLE II

Upper bound on QoS parameters for some applications

B. *Model Description*

As mentioned, the proposed AC model considers:

- intra-domain and inter-domain aspects, allowing: (i) the control of the available resources in a domain; (ii) the sharing of the existing SLS between domains; (iii) the ability to provide end-to-end QoS;
- different application QoS requirements and traffic profiles;
- the support for distinct network services;
- scalability, efficiency and deployment aspects.

The model is based both on edge-to-edge QoS monitoring of relevant QoS parameters for each service type and on the corresponding SLS control. A monitoring module, present at each egress router, measures the QoS parameters of each service taking into account the origin ingress router (Ingress/Egress Service Matrix), and also measures the egress SLSs occupancy. The resulting metrics, which reflect the domain service availability, are then used for AC at the corresponding ingress routers. The AC module operates based on service-dependent AC equations and proper parameters threshold intervals. The decision process can be implicit or explicit depending on the service characteristics, candidate application types and QoS guarantees.

Fig. 2 presents an overview of the required tasks to fulfill our first objective: assuring intra-domain QoS, by controlling the

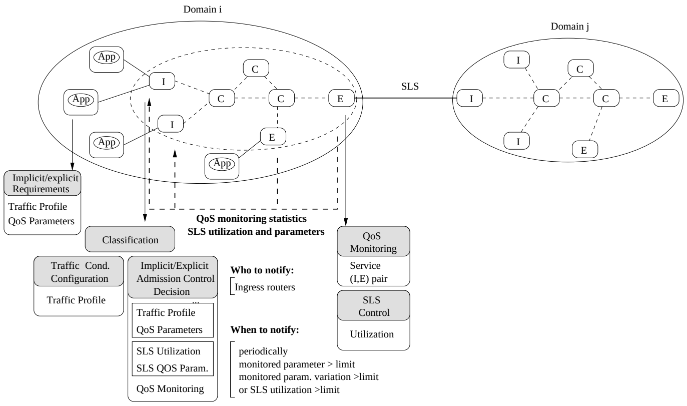

Fig. 2. Domain Activities location

QoS parameters and the fair sharing of the existing SLS.

From an end-to-end perspective and for flows requiring more guarantees, QoS performance metrics have to be included in the admission requests (e.g. using RSVP), to inform the downstream domain of the available service in the current and downstream domains. Using the incoming and its own metrics each domain ingress router decides if the flow can be accepted. The last acceptance/rejection decision is taken at the receiver (see Fig. 3). This solution leads to a generic AC model, which can be applied both to source and transit domains.

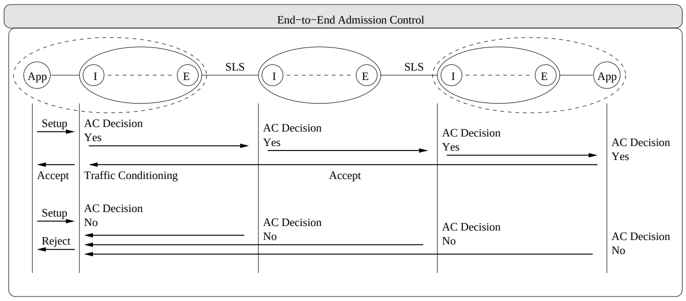

Fig. 3. End-to-end Admission Control Procedure

Ingress routers, in addition to classification and marking, deal with AC and TC. AC can be explicit or implicit, depending on the CoS and application type. When it is explicit, TC is parameterized based on the flow's traffic profile[1]. An explicit AC decision, described below, takes into account application requirements, domain QoS monitoring information for the corresponding CoS and, eventually, SLS information. Implicit AC may use one of the known methods for implicit detection of new flows, accepting them or not [13].

Egress routers deal with on-line QoS monitoring and SLS control. QoS monitoring measures relevant parameters for each service, using appropriate time-scales and methodologies (see section VI). The resulting metrics reflect the available service from each ingress[2]. The use of SLSs at each egress is also controlled to assure that traffic to other domains does not exceed the negotiated profiles. The QoS metrics and SLS information are then sent to the corresponding ingress routers (see Fig. 2), to update the ingress/egress service matrix and to be used for distributed AC. This notification can be carried out periodically; when a metric or metric variation exceeds a limit; or when the SLS utilization exceeds a threshold.

C. The criterion for AC

Explicit AC - the AC decision process for explicit AC in a domain is illustrated in Fig. 4. As shown, flow AC requires two initial verifications:

- SLS utilization control: the SLS can accommodate the new flow traffic profile.
- QoS control: for a particular egress node and service, the domain QoS metrics and the SLS QoS parameters[3] must fulfill the application's QoS requirements;

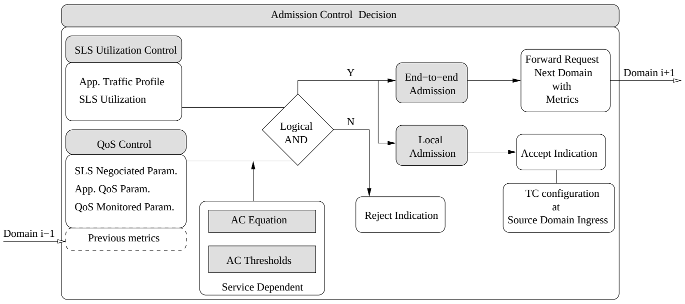

Fig. 4. Admission Control Criterion

Each AC decision is based on a service dependent AC equation, more or less conservative, depending on the service guarantee to be provided. In general, a conservative criterion will take the worst-case working scenario (e.g. peak rates, concurrent traffic, optimistic metrics, etc.). For each class, admission thresholds must be stricter than the class QoS objectives, which in turn, must be stricter than the requirements of all accepted flows. These thresholds and tolerances need to be carefully established and tuned in order to achieve an efficient AC[4].

In the admission process, if one of the tested conditions fails, the flow request is rejected, and the application notified. When the flow is accepted in the domain, the notification may be generated either locally (local admission) or remotely (end-to-end admission). The latter case occurs when an end-to-end availability check is required. In this case, the request including the

[1] Note that per-flow TC is only performed at source domain; in transit domains, TC is based on SLS traffic profile.

[2] For downstream domains, this QoS monitoring is also used for SLS monitoring and conformance verification.

[3] While domain QoS metrics are always verified, SLS verification (QoS and utilization) is undertaken when the destination of the flow request is outside the domain. When the destination is inside the domain, SLS verification is not mandatory, however internal SLS may be in place (intra-domain SLSs). This will turn the AC process generic and independent of the destination location.

[4] This need is stressed as QoS metrics result from a measurement based process which may introduce inaccuracy. In addition, concurrent AC requests may take place at other ingress routers, therefore defining such tolerances shall take this into account.

QoS metrics is propagated across domains up to the destination, and the notification is sent back to the source. Figure 3 illustrates this process.

Implicit AC - Implicit AC may be oriented to applications which do not use signaling, and in particular to elastic applications. This type of AC, likely to be implemented only in the source domain, will be restricted to SLS information and QoS monitoring. Two possible implicit reject actions are (i) SYN packets discarding or (ii) simply packet discarding based on flow accept/reject tables [13].

D. Model features

The proposed model has important features such as (i) the requests are only processed at ingress nodes; (ii) the state information per flow is only kept at the ingress router of the source domain; (iii) the other domains in the path maintain the TC based on the SLS traffic profile, as usual; (iv) the use of AC signaling does not imply the soft-state behavior and symmetric routing paths usually associated with a signaling process[5]; and (v) the network core is treated as a black-box, without changes. Despite the simplicity of the model, a lighter AC solution could be adopted if the transit ingress nodes do not implement AC, and just use application signaling for reporting QoS metrics to the destination.

Performing AC using on-line QoS monitoring avoids extra control mechanisms and simplifies the network control plane. When performed in a systematic way, measurements can be intrinsically auto-corrective and can detect short-term or long-term traffic fluctuations, depending on the measuring time unit or interval. Additionally, the effect of cross traffic and other internally generated traffic (e.g. routing, management and multicast traffic) is implicitly taken into account. Furthermore, the initial latency is reduced (as the metrics are available for on-line decisions) and per-application intrusive traffic to obtain the metrics is avoided.

V. Additional SLS deployment issues

For a service provider, the service monitoring process, and in particular, SLS auditing needs to be extended in order to include SLS monitoring and SLS conformance verification tasks (see Fig. 5). SLS monitoring will be carried out resorting to the proposed QoS monitoring module. Service generic information, including the service parameters and monitoring rules, and eventually specific service information (customer dependent) are used in this module to provide input for SLS conformance verification. The SLS conformance verification output reports whether the negotiated service is being provided or not, and allows resource reconfiguration and/or tuning actions.

Recall that as the QoS monitoring module has an ingress-egress view of the service QoS, it is particularly suited to SLS auditing. Other important aspect to consider is the space and time granularities in which monitoring is accomplished. Although off-line monitoring is a common approach, current studies highlight that it should be performed on-line [25].

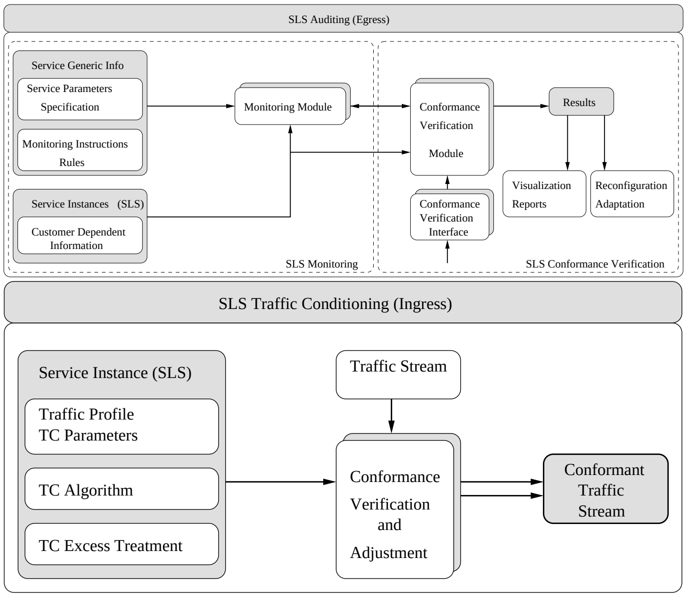
SLS Auditing (Egress)
Service Generic Info
Service Parameters Specification
Monitoring Instructions Rules
Service Instances (SLS)
Customer Dependent Information
Monitoring Module
Conformance Verification Module
Conformance Verification Interface
Results
Visualization Reports
Reconfiguration Adaptation
SLS Monitoring
SLS Conformance Verification
SLS Traffic Conditioning (Ingress)
Service Instance (SLS)
Traffic Profile TC Parameters
TC Algorithm
TC Excess Treatment
Traffic Stream
Conformance Verification and Adjustment
Conformant Traffic Stream

Fig. 5. SLS Related Tasks

In opposition to egress nodes, where the service to be provided is assessed, ingress nodes control incoming traffic aggregates. This consists of SLSs traffic profile conditioning based on negotiated TC parameters, TC algorithm and out-of-profile traffic treatment.

VI. Monitoring Issues

In the proposed AC model the monitoring module and the AC module are independent. Thus, the monitoring methodologies and corresponding implementation details are hidden from the AC module. This level of abstraction provides flexibility and portability to the model. Despite this, and due to the relevance of monitoring for AC, this section highlights major monitoring issues to be considered. The problematic of monitoring [26] involves the definition of metrics, measurement methodologies, and timing issues. The definition of these aspects will lead to an efficient and feasible monitoring solution.

Metric definition issues - The definition of metrics requires identifying of relevant parameters for each service type and corresponding statistics. As referred in section III, IPPM aims to develop a set of standard metrics providing unbiased quantitative measures of quality, performance and reliability of operational Internet data delivery services [22]. Defining a metric, identifying its type (analytical or empirical), its composition (in spatial and temporal terms) and its corresponding instances (singleton or sample metric) are topics to be addressed for every single parameter to be defined [27].

[5] There is no guarantee that the path used for the flow data is the same used for the flow request. This may not be particularly problematic as long as the new path is established maintaining the same QoS characteristics. In this case, the metrics of the new path will soon reflect the load increase and the corresponding AC will control it accordingly.

Measurement methodologies issues - The definition of a measurement strategy or methodology can use either passive or active measurements, or combinations thereof. Passive measurements use existing traffic for computing metrics without resorting to intrusive traffic. In this method, the amount of data gathered can be substantial, especially if information from all packets is required. In some cases, sampling techniques may be needed [28]. Active measurements allow a wide range of emulation scenarios by injecting extra traffic into the network. In some cases, small traffic volumes are enough to obtain meaningful measurements. Due to their characteristics, both mechanisms are important for QoS monitoring and SLS auditing. Apart from these techniques, the metrics propagated by routing protocols can provide useful QoS information about links and paths to the edges.

Timing issues - Depending on the measured parameter, timing issues can be particularly significant. For instance, one-way delay requires clock synchronization between measurement points. The temporal validity of metrics (short-term, long-term) needs to be defined according to the purpose of its use. For AC, the metrics do not require a fine granularity, i.e. seconds to minutes are the common operating time scales for AC mechanisms [14]. Although computing on-line metrics can be a difficult task, especially in high-speed networks due to large traffic volumes and reduced packet processing time, they are feasible for the time granularity the AC model requires.

VII. Acknowlegments

This work has been partially funded by FCT under the Project QoS II, POSI EEI/10168/98.

VIII. Conclusions

This paper proposes a generic AC traffic model for CoS architectures based on edge-to-edge QoS monitoring of each service type and on the control of SLS between domains. The intra-domain AC is distributed and performed at ingress routers using a service-dependent criterion, QoS metrics and SLS utilization. The QoS metrics are obtained through on-line monitoring performed at egress nodes, in an ingress-egress service basis. As the network core is treated as a black-box, neither state information (per-flow or aggregate) nor core changes are required. The model considers both implicit and explicit AC, and reduces AC latency as the metrics are available on-line. From an end-to-end AC perspective, only domain ingress routers are involved, but no per-flow state information is required. This information is only kept at source domain ingress routers for TC; the other domains in the path maintain SLS traffic profile conditioning. Besides SLS definition, additional SLS deployment issues such as SLS Monitoring and Conformance Verification, and SLS Traffic Conditioning, are identified and described. As the model is monitoring based, the definition of metrics, measurement methodologies and timing issues, need to be carefully assessed. Current work includes testing the model operation and efficiency both intra and inter-domain, for a limited set of services.

References

[1] S. Blake, D. Black, M. Carlson, E. Davies, Z. Wang, and W. Weiss. An Architecture for Differentiated Services, IETF RFC2475, 1998.

[2] R. Braden, D. Clark, and S. Shenker. Integrated Services in the Internet Architecture: an Overview, IETF RFC1633, 1994.

[3] F. Baker, C.Iturralda, F. Le Faucheur, and B. Davie. Aggregation of RSVP for IPv4 and IPv6 Reservations, IETF RFC3175, September 2001.

[4] B. Teitelbaum, S. Hares, L. Dunn, R. Neilson V. Narayan, and F. Reichmeyer. Internet2 QBone: building a testbed for differentiated services. *IEEE Network*, 13(5):8...16, September/October 1999.

[5] Zhi-Li Zhang et al. Decoupling QoS Control from Core Routers: A Novel Bandwith Broker Architecture for Scalable Support of Guaranteed Services. In *SIGCOMM'00*, 2000.

[6] I. Stoica and Hui Zhang. Providing Guaranteed Services Without Per Flow Management. In *ACM SIGCOMM'99*, October 1999.

[7] L.Westberg et al. Resource Management in Diffserv (RMD) Framework, February 2002.

[8] L. Breslau and Sugih Jamin. Comments on the Performance of Measurement-Based Admission Control Algorithms. In *IEEE INFOCOM'00*, March 2000.

[9] L. Breslau et al. Endpoint Admission Control: Architectural Issues and Performance. In *ACM SIGCOMM'00*, 2000.

[10] C. Cetinkaya, V. Kanodia, and E. Knightly. Scalable Services via Egress Admission Control. *IEEE Transactions on Multimedia*, 3(1):69...81, March 2001.

[11] V. Elek, G. Karlsson, and R.Rnngren. Admission Control Based on End-to-End Measurements. In *IEEE INFOCOM'00*, 2000.

[12] M.Gerla, S. Lee, G. Reali, and D. Sorte. Performance of Different Call Admission Schemes in a QoS Diffserv Domain, 2001. http://www.cs.ucla.edu/ñrl/hpi/papers/2001-milcom-0.ps.gz.

[13] R. Mortier and I. Pratt and C. Clark and S. Crosby. Implicit Admission Control. *IEEE Journal on Selected Areas in Communications*, 18(12):2629...2639, December 2000.

[14] A. Bak, W.Burakowski, F. Ricciato, S. Salsano, and H. Tarasiuk. Traffic Handling in AQUILA QoS IP Networks. In *QoFIS'01*, September 2001.

[15] V.Reijs. QoS Monitoring and SLS Auditing, July 2001. http://www.heanet.ie/Heanet/projects/nat_infrastruct/qosmonitoringtf-ngn.html.

[16] C. Dovrolis and M. Jain. 'End-to-End Available Bandwidth: Measurement methodology, Dynamics, and Relation with TCP Throughput. In *ACM SIGCOMM'02*, August 2002.

[17] M. Przybylski, S. Trocha, and Poznan. Network Measurement Tools Test, 2001. http://qos.man.poznan.pl/files/measurement_full.pdf.

[18] S. Lima. A Comparative Analysis of Admission Control Strategies, March 2002. Technical Report TR0102l.

[19] Goderis et al. Service Level Specification Semantics and Paramters, February 2002. IETF draft: draft-tequila-sls-02.txt (work in progress).

[20] A. Sevasti and M. Campanella. Service Level Agreements Specification for IP Premium Service, October 2001. Geant and Sequin Projects.

[21] V.Reijs. Perceived quantitative quality of applications, July 2001. http://www.heanet.ie/Heanet/projects/nat_infrastruct/perceived.html.

[22] Tijani Chahed. IP QOS Parameters. *TF-NGN*, November 2000.

[23] V.Reijs. Tools for measuring the SLS metric, July 2002. http://www.heanet.ie/Heanet/projects/nat_infrastruct/nettools.html.

[24] K. Nichols and B. Carpenter. Definition of Diff. Services Per Domain Behaviors and Rules for their Specification, IETF RFC3086, 2001.

[25] P. Bhoj, S. Singhal, and S. Chutani. SLA Management in Federated Environments. *Computer Networks*, 35(1), January 2001.

[26] Yuming Jiang, Chen-Khong Tham, and Chi-Chung Ko. Challenges and approaches in providing QoS monitoring. *International Journal of Network Management*, 10(6):323...334, November 2000.

[27] V. Paxson, G. Almes, J. Mahadavi, and M. Mathis. Framework for IP Performance Metrics, IETF RFC2330, 1998.

[28] I. Cozzani and S. Giordano. Traffic Sampling Methods for End-to-End QoS Evaluation in Large Heterogeneous Networks. *Computer Networks and ISDN Systems*, September 1998.