SERVICE LEVEL AGREEMENT PROVISIONING AND MONITORING FOR END TO END QOS

Vassiliki Pouli*, Stella Kafetzoglou*, Mary Grammatikou*, Chrysostomos Tziouvaras, Symeon Papavassiliou*,Vassilis Maglaris***

* {vpouli, skafetzo, mary, papavass, maglaris}@netmode.ntua.gr
Network Management & Optimal Design Lab. (NETMODE)
National Technical University of Athens, Iroon Polytexniou 9, 15780 Zografou, Greece

** tziou@grnet.gr
Greek Research & Technology Network - GRNET S.A.
56, Mesogion Av., Ampelokipi, 11527 Athens, Greece

Abstract
Grid applications require computing, storage and network resources offered from different Resource Centers (RCs). For the exchange of resources and the data transfers between two RCs an end-to-end path should be created. To provide QoS assurances to this end-to-end path, the complete path should be reserved and configured to offer that particular service and the global network setup should be compliant and aware of the service offered. In this paper we propose a coupled procedure for the establishment of end to end Service Level Agreements (SLAs) along with the service provisioning in order to assure that the service and network providers comply with the level of QoS that the applications expect along the path. These SLAs stipulate the appropriate requirements for service availability and performance providing guarantees for the delay, inter packet delay variation, packet loss and capacity.

Keywords
end-to-end Service Level Agreement , quality of service, monitoring, service provisioning, resource center

1. Introduction
A Service Level Agreement (SLA) is a contract in which a certain level of service is agreed between a service provider and a service consumer. It may specify the levels of availability, serviceability, performance and operation conditions [7], [11]. It may also define the procedures and the reports that must be provided to track and ensure compliance with the SLA or describe other attributes of the service like billing or penalties in the case of violation of the SLA [14].

In the framework of the EGEE-II project [5] we have defined an end-to-end SLA implementation procedure that is applicable to the GÉANT2 network, the interconnected National Research & Education Networks (NRENs) and the EGEE-II Resource Centers. In this procedure we establish end-to-end Service Level Agreements (SLAs) along with the provisioning of a service in order to define the required availability, reliability and performance guarantees for the requested service along the end-to-end path traversing the GÉANT2/NRENs networks.

The GÉANT2 network forms together with the NREN backbones a hierarchical IP network infrastructure upon which EGEE is based in order to achieve interconnection of the involved computing infrastructures into a pan-European Grid infrastructure. It forms together with the NRENs the major part of the end-to-end path between two EGEE Resource Centres (RCs) that provide certain computing and storage resources. It is obvious that the end-to-end path for data transport between any two EGEE RCs is not fully provided by the GÉANT2/NRENs path, as quite often EGEE RCs are connected to the NREN of their country via an NREN client domain (MAN, campus or institution domain), as shown in figure 1.

In grid environments [12] where the applications or users require resources from different service providers the issue of acquiring a specified level of QoS is of great importance [9]. A model that provides scalable QoS is the DiffServ [15] which polices and marks the DiffServ flows in edge routers according to SLAs, thus offering a variety of end-to-end services across multiple, separately administrated domains. Based on the DiffServ architecture, the Premium IP service offered today by the GÉANT2 network is a service, as implemented in the SEQUIN project [16], that offers network priority over other traffic on GÉANT2 [13], namely the Best Effort (BE) and Less Than Best Effort (LBE) traffic [17]. The Best Effort IP service is a service which does not provide any performance guarantees but is useful for non-real-time applications such as FTP. The Less than Best Effort service utilises the capacity that is not used by Best Effort and Premium IP traffic thus is targeted to flows of data which do not require particular quality assurances; for example transfer vast amounts of (non real-time) data for applications such as GRID computing without adversely affecting other traffic. On the contrary, the Premium IP service provides a service similar to that of a virtual leased line. Data packets that are sent using the Premium IP service will experience no congestion in the network

regardless of the load of the other traffic classes. As a result, delay and packet loss are kept to a minimum which can be particularly useful for real-time applications, such as Voice over IP (VoIP) and video conferencing.

The rest of this paper is organized as follows: in section 2 we classify the participating in the SLA establishment procedure domains in PIP compliant, supportive and indifferent depending on whether they comply with, simply support or are indifferent to the PIP IP service. In section 3 we present the SLA establishment procedure, in section 4 we propose the SLA monitoring procedure while in section 5 we conclude the paper and present future work items.

2. Participating domains' classification

In the SLA establishment procedure that we define whenever a user/application requests a distinct or aggregate flow of EGEE traffic between two distant EGEE RCs, travelling through the GÉANT2 and NREN networks and requiring special constraints to be met by the network, the involved domains between the two EGEE RCs have to be identified and a decision has to be taken whether the request can be satisfied. If the request can be accepted, the service path is reserved and an SLA instance is created in order to define the requested level of QoS, by providing all the technical and administrative details for maintaining, monitoring and troubleshooting this data flow.

Depending on what class of service a domain is able to provide, domains are classified as: PIP compliant, PIP supportive and PIP indifferent. A PIP compliant domain implements the PIP service and has the following features: classification of PIP packets, remarking of traffic at the ingress, policing to control exceeding PIP traffic, admission control on PIP service requests and a well-known QoS domain profile. A PIP supportive domain supports the following: preserving marking of PIP packets, operating in an over-provisioned manner, monitoring of performance and performing admission control on each new PIP request. Finally, a PIP indifferent domain cannot be classified as a PIP compliant or PIP supportive.

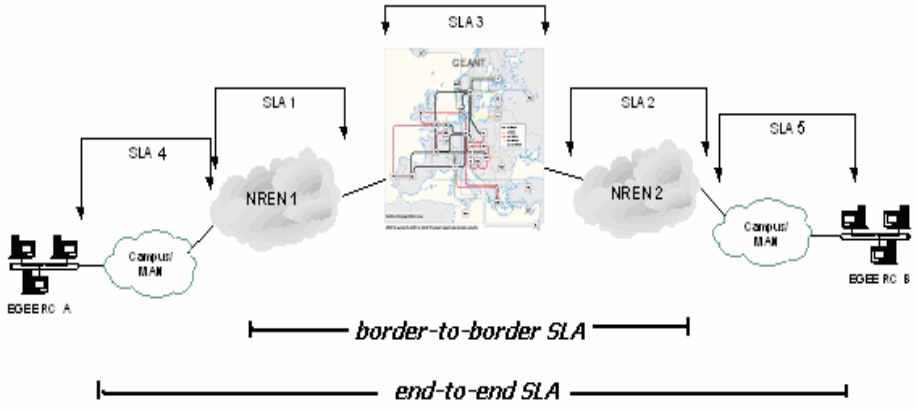

Figure 1: End-to-End network path between two EGEE Resource Centers (RCs)

Due to the fact that Premium IP service is the responsibility of each administrative domain, the network services provisioning model is based on the establishment of per-domain agreements (SLAs) for Premium IP service. Each such SLA defines the administrative and technical details, the availability, quality and performance guarantees, along with the monitoring and troubleshooting process.

As shown in figure 2, depending on whether a domain is fully PIP compliant, simply PIP supportive or PIP indifferent, a different SLA type is provided. When combined the per-domain SLAs along the path that the packets of a data flow follow, can form an end-to-end SLA that characterizes the treatment of this flow [1]. Each such SLA corresponds to one direction. If the data flow is bidirectional then two SLAs should be defined, one for each direction. The contracted values for these two SLAs might be different (asymmetric capacity for example).

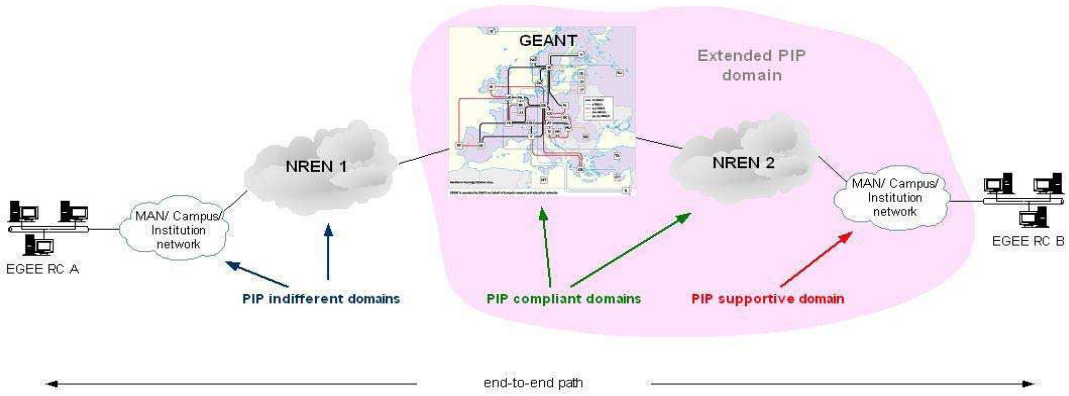

Figure 2: end to end network path between two EGEE RCs, comprising of PIP indifferent, compliant and supportive domains

3. SLA Establishment Procedure

The SLA establishment is a procedure that accompanies the service request and provisioning procedure and is implemented in two stages [2]. The rationale behind this two-stage process is that Grid applications require that the end-to-end service path be available in time, meaning in the order of a few minutes or hours when the application requests it. The reservation of the service path is foreseen to be established automatically in the near future through the Advanced Multi-domain Provisioning System (AMPS) of GÉANT2, but currently the process is done manually. Due to the fact that the reservation of the extended QoS network made by the AMPS (Advanced Multi-domain provisioning System) [4], implies a configuration lead time (currently this is two working days), we concluded that the end-to-end SLA establishment procedure should be implemented as a two-stage process where in the first stage, the service reservation in the extended QoS network should take place and in the second stage the activation of the service should be realized (see figure 3). Accordingly, stage 1 is defined as Service Reservation (SR) and stage 2 is defined as Service Activation (SA).

More specifically during the Service Reservation stage (first stage), a PIP reservation is made along the QoS extended domain which comprises the NRENs and the GÉANT2 network and a border-to-border SLA (figure 1) is established for the given reservation derived from the SLAs of the domains belonging to this extended QoS domain. It should be noted that the extended QoS domain can be expanded up to the sites if some intermediate domains implement the PIP service. During the Service Activation stage (second stage), configuration of the routers in the last mile network takes place and an end-to-end SLA (figure 1) is produced to characterize the treatment of the service flow that is instantiated. The end-to-end SLA extends the pre-established border-to-border SLA with the SLAs from the site networks, i.e. last mile domains.

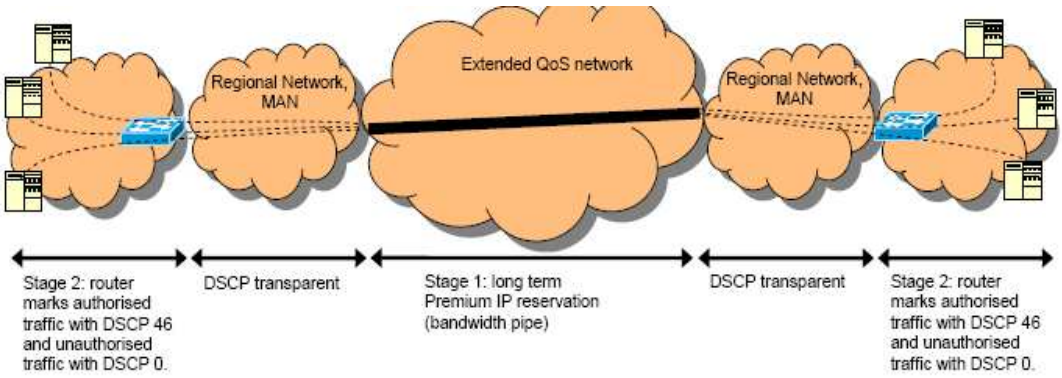

Figure 3: Two-stage provisioning process

In order to proceed with the two phase SLA establishment procedure we introduce the EGEE Network Operations Center (ENOC) which is the counterpart of a NREN NOC for EGEE; it is in charge of the "EGEE network" operations and maintains a database, called Network Operational Database (NOD) which stores topology information along with the maximum bandwidth offered by every domain participating in the EGEE-II project. [3]

In the following section the two phase model is presented in detail.

3.1. Phase1: Service Reservation

During this phase, the PIP reservation along with the SLA establishment in the QoS extended network is made. The reservation of the PIP request is made by the AMPS server [4] which is the GN2's automatic provisioning system. At time of writing the GN2-SA3 activity works on the deployment of this service called Advanced Multi-domain Provisioning System (AMPS), which accepts PIP service requests and is able to identify the involved domains and reserve the necessary network resources. AMPS servers should be installed in every Premium IP compliant and supportive NREN and a user will be able to address the AMPS server at his NREN, in order to perform its reservation. It should be noted here that this reservation is made between two sites that are part of the extended QoS network (GEANT/NRENs), where AMPS servers are actually deployed, and do not include MAN/regional networks. If a MAN/regional network decides to support QoS and deploy an AMPS server, it will then become part of the extended QoS network. In this phase a border-to-border SLA is also produced as depicted in Fig. 1 that will stipulate the requirements of the service reservation.

In the case of EGEE users, the ENOC entity will be responsible for making the request for a service reservation, and if the reservation is accepted, an SLA will also be produced to describe the expected quality of the user's reservation. The SLA that is provided in this stage is called border-to-border SLA and is produced from the merging of the individual SLAs reported by the domains belonging to the extended QoS network, namely GÉANT2 and NRENs. The rules for the merging of the individual SLAs are described in [1]. The reason for the ENOC's proxying between the user, Service Request actor, and the AMPS is two-fold: first, the user may not have sufficient knowledge about the reservation procedure or where to address , and second the AMPS server may not be able to authenticate/authorize each individual user but only authenticate/authorize EGEE as a whole through the ENOC. Thus the authentication/authorization of each EGEE individual user is the responsibility of the ENOC.

The steps that should be followed for the service reservation are shown below:

• *Action 1:* The SR actor (EGEE VO/RC/application) asks the ENOC for a service. The SR actor should provide the ENOC with the source and destination sites or hosts of the service flow, the requested bandwidth and the duration of the requested service (start and end time). Additional information should be provided about the purpose of the request and the credentials of the SR, so that the ENOC can authenticate/authorize the SR actor. The user requests to the ENOC can be made either by email or, alternatively, the SR actor can fill a form provided at the ENOC's web site.

• *Action 2, 3:* The ENOC retrieves from its database (NOD) the appropriate information in order to identify the involved domains and discover the maximum bandwidth allowed at each involved domain for EGEE traffic.

• *Action 4:* The ENOC checks whether the service request is within the "agreement" boundaries of each involved domain.

• *Action 5(a):* If the service request cannot be supported by at least one involved domain the ENOC informs the SR actor and the flow stops here.

• *Action 5(b):* Otherwise, if the service request can be satisfied by all participating domains, the ENOC forwards the request to the first AMPS server of the involved networks on the path which has a web interface that accepts user requests for a service reservation, which then communicates with its peers (AMPS servers) on the path to assess the feasibility of the request given the reserved bandwidth from other requests in every domain and emails the administrators of involved NOCs to effect the configuration.

• *Action 6(a):* The AMPS reply is negative, i.e. the service request cannot be satisfied. The ENOC informs the SR actor and the flow stops here.

• *Action 6(b)*: Affirmative reply returns to the ENOC.

• *Action 7:* The request can be satisfied within the Extended QoS Network (GÉANT2/NRENs), so the ENOC asks from the involved domains the completion of the SLA templates.

• *Action 8:* The completed SLA templates are returned to the ENOC.

• *Action 9:* The ENOC merges the returned SLAs into a border-to-border SLA and stores it to its database (NOD).

• *Action 10:* The service reservation ID is returned to the SR actor.

The flow diagram below shows the described procedure:

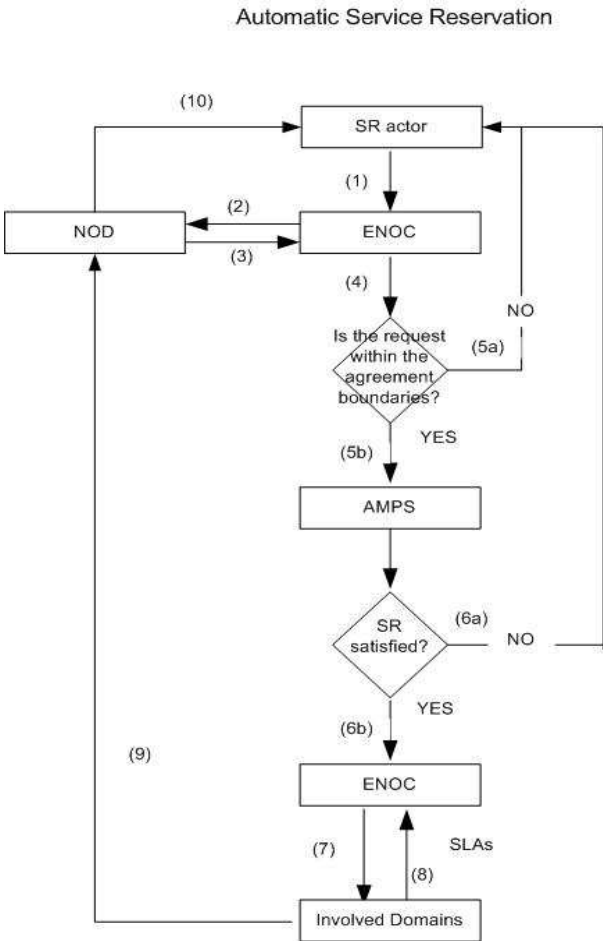

Figure 4: Automatic service reservation and border-to-border SLA establishment procedure

It should be noted that in cases when domains have not deployed an AMPS server (or a similar service) or in situations where PIP indifferent domains are involved in the end-to-end path of the user's request as shown in Figure 2 -where, as mentioned before, no AMPS server will be deployed in these domains- the reservation of the user's service request has to be performed manually, i.e. each domain has to be asked separately whether it can reserve the necessary network resources in order to satisfy the user's request. As before, the user addresses his request to the ENOC and if it is satisfied, the path is reserved and a border-to-border SLA is created.

3.2. Phase 2: Service Activation

It should be noticed that after the service is reserved with the above procedure, service activation is required. This activation takes place when a grid job is about to start. As long as the reservation in the extended QoS domain is still in effect, the only configuration needed is the configuration of the routers in the last mile networks.

Along with the configuration of the routers, the NOCs of the two last mile networks provide the SLAs of their domains to ENOC. ENOC merges them with the border-to-border SLA provided from phase 1 and produces an end-to-end SLA as shown in Fig. 1. The SLA produced assures the end-to-end QoS for the specific service instance.

The Service Activation phase is shown in more detail below:

• *Action 1, 2:* The ENOC retrieves from its database (NOD) the service reservation information in the extended QoS network.
• *Action 3:* The ENOC checks whether the reservation is still effective.
• *Action 4(a):* If the reservation is not effective, the ENOC informs the SR actor and the flow stops here.
• *Action 4(b):* If the reservation is still effective, the ENOC forwards the request to the site networks and asks whether the service can be satisfied given the traffic in every domain.
• *Action 5(a):* If at least one domain's reply is negative, then the service request cannot be satisfied. The ENOC informs the SR actor and the flow stops here.

• Action 5(b): Affirmative reply returns to the ENOC.
• *Action 6:* The request can be satisfied within the NREN client domains, so the ENOC asks from them to complete the SLA templates.
• *Action 7:* The NOCs of the NREN client domains send the completed SLA templates.
• *Action 8:* The ENOC merges the returned SLAs with the border-to-border SLA and the end-to-end SLA is produced and stored to the NOD.

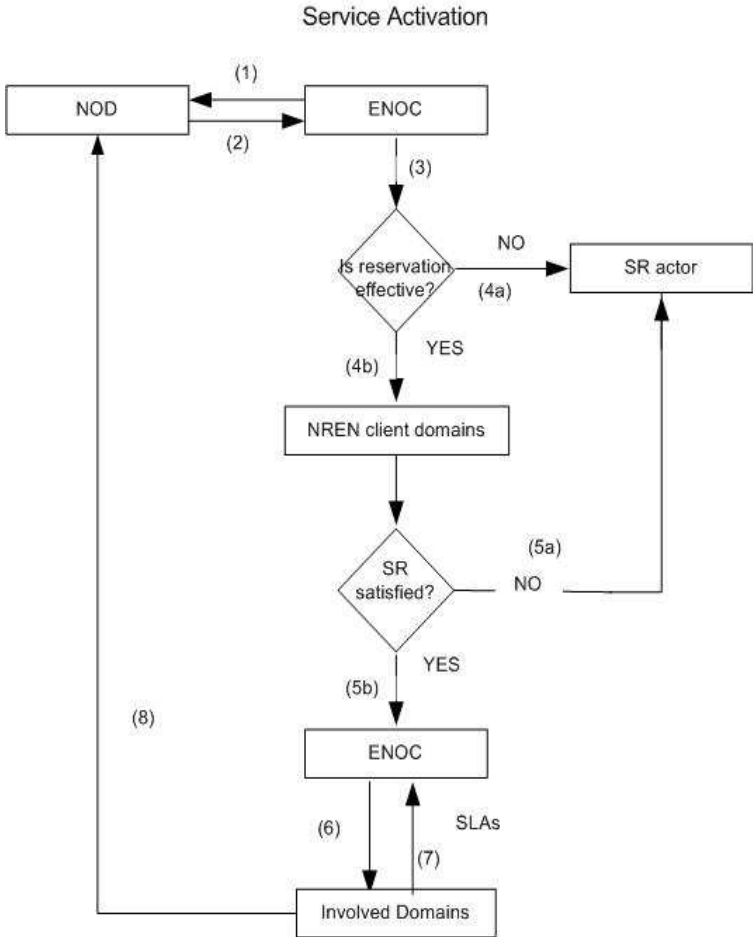

Figure 5: Service activation and end-to-end SLA establishment procedure

4. SLA Monitoring

After the end-to-end SLA for a certain traffic instance is established ENOC is responsible for monitoring if the SLA guarantees are honored. It should be noted that the SLA monitoring ought not to impact network performance [8].

The SLA monitoring and troubleshooting procedures are described below:
• In order for the ENOC to monitor an end-to-end SLA it should access the NPM DT (Network Performance Monitoring Diagnostic Tool) and provide all the required parameters to this tool (e.g. time period over which the network monitoring for the SLA applies, source and destination IP addresses, metrics that need to be measured, such as OWD, achievable bandwidth etc). The NPM DT is a monitoring framework developed as part of the EGEE project and provides end-to-end measurement results. The measured data provided by this tool can be viewed in various ways (data table, matrix, time plot, histogram) so the ENOC can choose the way in which the resultant data is displayed. It should be mentioned that the ENOC must inform the NPM DT Administrator, in advance, about the sites that it wants to measure and what kind of measurements it wants to make in order for the available configurations from the NPM Diagnostic Tool and the involved sites (installation of appropriate software) to take place.
• From the results, the ENOC should try to identify if the SLA guarantees are honoured. In order for the problem identification to be easier it would be desirable for the NPM DT to provide the ability to set thresholds on the metrics that the ENOC wants to measure and accordingly generate alarms in case of threshold violations.

• In case an SLA metric is violated due to a problem (e.g. bandwidth falls below the specified value in the SLA template) then the ENOC should try to identify in which domain the problem occurred by taking intermediate measurements if needed and accordingly notify the Technical Contact(s) (provided in the administrative part of the end-to-end SLA) of one or more of the domains along the path in order to solve the problem. Those involved in a service path, such as VOs/RCs, can install and/or run their own application oriented tests in order to verify the performance of their service and if it is degraded. They can also access the NPM DT (Network Performance Monitoring Diagnostic Tool) as long as they are authorized by the NPM DT administrator and make the same measurements as ENOC would do in order to monitor their SLA service. In case of faults, perceived by them, the problem must then be reported to GGUS [9] and GGUS in turn redirects it to the ENOC. It should be mentioned that the ENOC must refer to the specific SLA instance for which the problem exists when reporting a problem to a domain, so that the service details, as registered in the corresponding SLA instance, are taken into consideration for troubleshooting.
• When the location of the problem cannot be located then the GN2 PERT [6] team is informed in order to troubleshoot.
• In each domain, the fault-handling and trouble ticket procedures specified in the corresponding SLA must be followed. Normally, the domain(s) in which the problem occurs will be able to troubleshoot it.
• After troubleshooting they should report back to the ENOC what were the causes of the problem and if normal operation has been restored.
• Finally, actual failure indication response times (as declared by a domain in the corresponding SLA) must be crosschecked by the ENOC against the promised values.

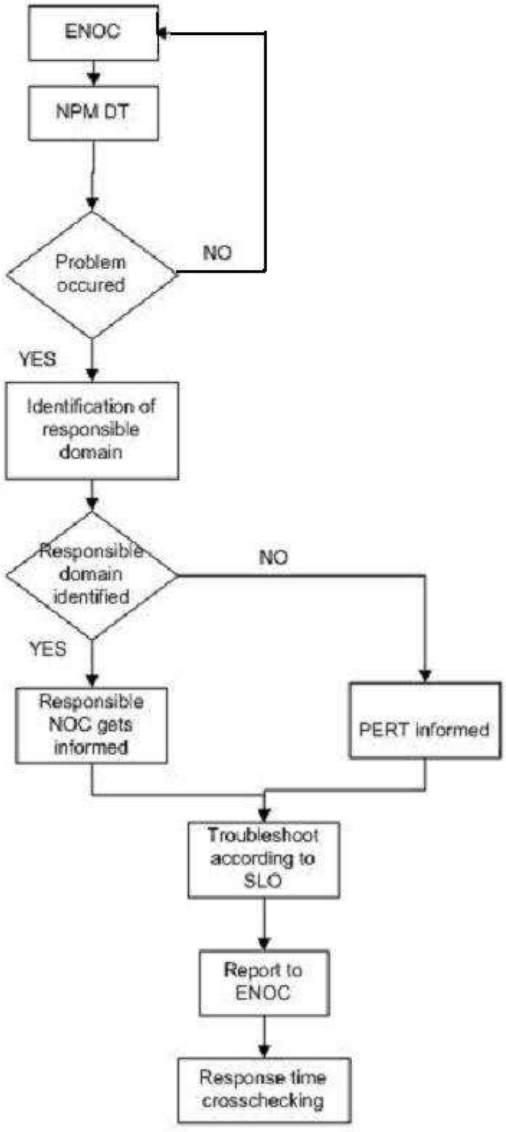

Figure 6: SLA monitoring and troubleshooting procedure

5. Conclusion And Future Work

In this paper we described a coupled procedure for the establishment in the IP layer of an end-to-end SLA along with the provisioning of a service thus defining the level of service corresponding to the specific requirements of the user request concerning the availability, duration, reliability and performance guarantees. The monitoring and troubleshooting procedures were also presented. Integrating the end-to-end SLA service in the provisioning model of networks which will start offering dedicated lightpath services with network capacity up to 10Gbps is for future work.

6. Acknowledgements

. This work is supported by the EU project EGEE[5], sponsored by the European Union under contract number INFSO RI 031688.

7. References

[1]. EGEE DSA 2.2: Institution of SLAs and Appropriate Policies, https://edms.cern.ch/document/565447
[2]. EGEE-L-NSAP Functional Specification, https://edms.cern.ch/document/635851/1
[3]. EGEE-MSA2.3: Operational Interface Between EGEE and GEANT/NRENs, https://edms.cern.ch/document/565449 https://edms.cern.ch/document/694911/1
[4]. AMPS, http://www.geant2.net/upload/pdf/GN2-04-153v4AMPS.pdf
[5]. EGEE project, http://public.eu-egee.org/
[6]. PERT (performance Enhancement Response Team): http://www.geant.net/server/show/nav.1235
[7]. J. Martin, A. Nilsson: "On Service Level Agreements for IP Networks", IEEE INFOCOM 2002
[8]. C. Molina-Jimenez, S. Shrivastava, J. Crowcroft, P. Gevros: "On the monitoring of contractual service level agreements", Proc. 1st IEEE International Workshop on Electronic Contracting, 2004
[9]. GGUS (Global Grid User Support): http://gus.fzk.de/
[10]. P. V-B Primet, F. Chanussot: "End To End Network Quality Of Service In Grid Environments: The Qosinus Approach", Workshop on Networks for Grid Applications, California, USA, 2004
[11]. D.C. Verma: "Service Level Agreements on IP Networks", Proceedings of the IEEE, 2004
[12]. I. Foster and C. Kesselman, "The grid : Blueprint for a new computing infrastructure," Morgan Kaufmann Publishers Inc., 1998.
[13]. Premium IP Service in GÉANT2: http://www.geant2.net/server/show/nav.1527
[14]. C. Courcoubetis and V. Siris: "Managing and Pricing Service Level Agreements for Differentiated Services", Proc. of IEEE/IFIP IWQoS'99,UCL, London, UK, May 31-June 4, 1999
[15]. K. Nichols, B. Carpenter, "Definition of Differentiated Services Behavior Aggregates and Rules for their Specification", IETF Draft, "draft-ietf-diffserv-ba-def-00.txt', Feb 2000.
[16]. SEQUIN project: http://www.dante.net/sequin/
[17]. Quality of Service in GÉANT2: http://www.dante.net/server/show/nav.00700a009