Available online at www.sciencedirect.com
SCIENCE DIRECT

Journal of Network and
Computer Applications 28 (2005) 285–302

Journal of NETWORK and COMPUTER APPLICATIONS
www.elsevier.com/locate/jnca

Service level agreements for DiffServ-based services' provisioning

Christos Bouras[a,b,*], Afrodite Sevasti[b,c]

[a]RA Computer Technology Institute-RACTI, Kolokotroni 3, Patras 26221, Greece
[b]Department of Computer Engineering and Informatics, University of Patras, Rion, Patras 26500, Greece
[c]Greek Research and Technology Network-GRNET, 56 Mesogion Av., Athens 11574, Greece

Received 13 January 2004; received in revised form 24 June 2004; accepted 6 July 2004

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 define the qualitative and quantitative characteristics of the services provided from a network provider to peering networks or customers. In this work, we define a template for the SLA structure to support the provision of a QoS service between two peering domains and then we proceed with the definition of an end-to-end SLA across consecutive domains, based on the bilateral ones. We also propose a model for the service provisioning procedures.

Keywords: Service Level Agreement; DiffServ-based service; Provisioning; End-to-end SLA; QoS

1. Introduction

A Service Level Agreement (SLA) is an explicit statement of the expectations and obligations that exist in a business relationship between two entities: a service provider and a customer (Rajan et al., 2000). Bilateral SLAs can also be defined among organizations that have a symbiotic relationship, with each being a customer of the other's

* Corresponding author. Address: Computer Technology Institute, 61 Riga Feraiou Str, Patras 26221, Greece. Tel.: +30-2610-960375; fax: +30-2610-960358.

E-mail addresses: sevasti@grnet.gr (C. Bouras), sevasti@grnet.gr (A. Sevasti).

1084-8045/$ - see front matter
doi:10.1016/j.jnca.2004.07.001

services. The SLA provides a means of defining the service. It specifies what the customer wants and what the supplier is committing to provide. It defines the standards for the quality of the service provided, setting performance objectives that the supplier must achieve. It also defines the procedure and the reports that must be provided to track and ensure compliance with the SLA. In the field of telecommunications networking, SLAs play a significant role, reinforced by the latest advances in differentiated services' provisioning.

The availability of high-speed transmission media and networking equipment, as well as the evolution of quality-demanding applications has focused research interest on the provision of Quality-of-Service (QoS) in addition to the traditional best-effort model of the Internet. A number of alternatives for service differentiation and QoS provision have been proposed and standardized, but in the case of IP-based backbone networks the Differentiated Services (DiffServ) architecture has prevailed, due to its scalability and deployment feasibility. The provisioning of IP services according to the DiffServ framework has introduced complexity in the corresponding business model and has raised the requirements for controlled resource allocation and management, definition, monitoring and verification of the quality provided. At this point, the appropriate definition of SLAs between customers and service providers is envisaged to provide the controlled environment required. In this framework, SLAs will act as mediators for mutual service provisioning between peering domains.

The DiffServ framework stands out for attempting to provide service differentiation to traffic in a scalable manner, by suggesting the aggregation of individual application flows with similar quality needs. It introduces the definition of different service classes to which such aggregates are appointed and the implementation of mechanisms for differential treatment by network elements (Per-Hop-Behaviour, PHB) of the packets belonging to each service class. A PHB is thus describing the treatment of aggregated traffic in a manner that ensures the quality guarantees provided by the corresponding service class.

Although, DiffServ has been initially confronted with a positive attitude, due to its scalability, the DiffServ framework mechanisms have proved difficult to deploy and monitor in a large scale in production networks. Based on the DiffServ framework, a number of service models constructed by a combination of DiffServ mechanisms have been proposed and experimentally evaluated up to our days. However, real-world implementations of DiffServ-based services in production networks have not successfully operated in large-scale yet. Missing DiffServ functionality from IP routers, translation to last-mile (end user access links) QoS, considerable operational and economical paradigm shifts required from operators, lack of a flexible business model, inadequacy of service verification infrastructure, inadequate standardization and architectural gaps are some of the major deployment problems analysed in (Paxson et al., 1998).

We believe that the DiffServ framework and DiffServ-based services do have a significant potential in upgrading the best-effort service model in today's Internet. However, due to its probabilistic rather than deterministic differentiation mechanisms, the provisioning model of DiffServ has to be thoroughly specified and standardized as soon as possible. Among the deployment problems already mentioned, a flexible business model for intra-domain development and peering agreements according to the existent agreements are considered crucial for the successful deployment of DiffServ-based