MANAGEMENT OF DIFFERENTIATED SERVICES WITH ACTIVE POLICIES

Paulo Rogério Pereira
Instituto Superior Técnico,
Universidade Técnica de Lisboa.
INESC, Rua Alves Redol, 9.
1000-029 Lisboa, Portugal.
Voice: +351-213100345.
Fax: +351-213145843.
prbp@inesc.pt

Elionildo Menezes, Djamel Sadok
Universidade Federal de Pernambuco.
Centro de Informática.
Caixa Postal 7851.
CEP 50732–970. Recife, PE, Brazil.
Voice: +55-81-2718430.
Fax: +55-81-2718438.
{esm, jamel}@di.ufpe.br

Paulo da Fonseca Pinto
Faculdade de Ciências e Tecnologia,
Universidade Nova de Lisboa.
UNINOVA, Quinta da Torre.
2825 Monte de Caparica, Portugal.
Voice: +351-212948300.
Fax: +351-212948532.
pfp@uninova.pt

1. Introduction

The explosive growth of the Internet and the use of new services such as e-business, VoIP (Voice over IP) and multimedia applications raises the need to support Quality of Service (QoS) requirements and to accommodate different service levels leading to differentiated user pricing policies.

The Differentiated Services architecture [1] aggregates traffic with similar QoS requirement in traffic classes that share the same per-hop behavior (PHB) throughout the network. The border nodes implement packet classification and traffic conditioning functions including metering, marking, shaping, and policing.

The availability of different service policies on such large networks demands for a semi-automatic management architecture which can decide the best configuration with little intervention from the manager. This paper describes an Internet Service Provider (ISP) scenario that offers three DiffServ based commercial services:

- *ENTERPRISE*: this service has the best performance. It offers rigid bounded delay guaranties. Hence, it is ideal for delay sensitive applications such as videoconferencing. It is normally implemented using the DiffServ Expedited Forwarding (EF) PHB.
- *STANDARD*: ideal for clients looking for a service that performs better than LIGHT, but cannot, or would not, pay for the limited and more expensive ENTERPRISE service. This service offers minimum QoS guaranties, whereby the network seems lightly loaded. This service is normally implemented using the DiffServ Assured Forwarding (AF) PHB.
- *LIGHT*: characterized by its occupation of whatever network resources are left. This service is implemented using the DiffServ Best Effort (BE) PHB.

Policies are used to classify the user traffic into the available PHBs according to the selected commercial service and the user application. Furthermore, the current number of users, the network load and failures are continuously monitored to prevent degrading the QoS. To attain these objectives, active management policies based on [2] were used. Decisions are made at each of several stacked planes describing directives at a decreasing abstraction level.

All results, presented in this paper, were obtained with a 15-node network with 15 points-of-presence (PoP) simulated using the Berkeley network simulator [3]. The results are presented in the next section showing that the use of active policies outperforms the use of more classic approaches.

2. Architecture and Simulation Results

Figure 1 shows the system architecture, and the functions performed at each of the four planes. At the Business Plane level, four different user admission policies where used:

a. Users are accepted until a fixed maximum number of users is reached for each PoP.
b. Users are accepted until a fixed maximum number of users is reached for each class, (Enterprise, Standard), after which the users get a lower available class, possibly Light.
c. Similar to **a**, but the maximum number of users depends on the current PoP load.

d. Combination of **b** and **c**: all maximums depend on current PoP load.

At the Service Plane level (the second), the traffic is mapped into one of the available DiffServ PHBs. Four different policies may be used to restrict this mapping according to the user application involved, as shown in table 1.

Policy	telnet/TCP	CBR/UDP	OnOff/UDP	http/TCP	ftp/TCP
1	any	any	any	any	any
2	any	any	any	AF BE	BE
3	EF AF	EF AF	EF AF	AF BE	BE
4	EF AF	EF AF	EF AF	BE	drop

Table 1. Policies to map application traffic to PHBs.

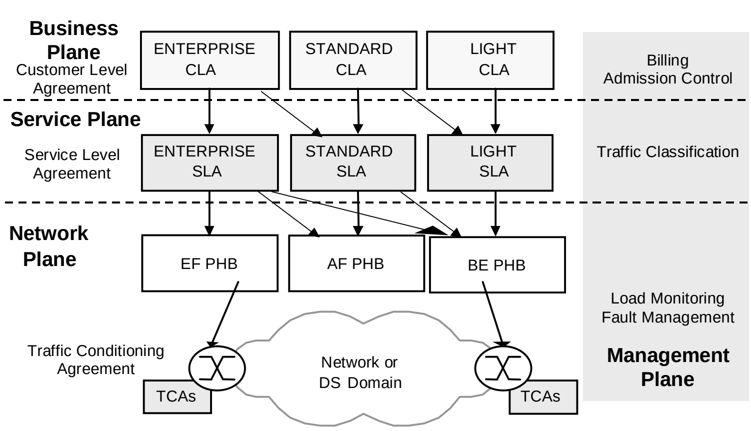

Figure 1: System Architecture.

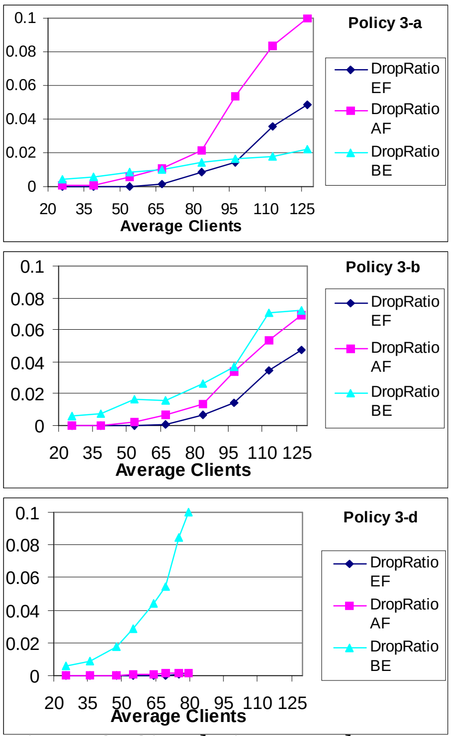

Figure 2: Simulation Results.

Simulations show that if too much traffic is allowed to use the EF PHB, as in policy **1**, the router schedulers will be overloaded, causing an unacceptable drop ratio, delay and jitter for the EF PHB and as a consequence also for the other lower priority PHBs. Policies **2**, **3** and **4** restrict the traffic of bandwidth intensive applications to the lower priority PHBs to avoid this problem. Policy **4** is used to further reduce traffic whenever a core link fails.

However, if too many users are allowed to forward traffic to the same PoP, the same effect occurs. Policy **b** is the simplest policy to overcome this problem, by monitoring the number of users in each class, and placing users that exceed their class quota in the next lower available class. A limit of 1/3rd of the users is allowed for classes Enterprise, Standard and Light. The downgraded users get worse service than what they have contracted, but in general, most users get better service. Figure 2 shows the simulation results. Policy **d** is an active policy that works by monitoring the packet drops in each PHB, and by dynamically adjusting the maximum number of users for each PoP, and the maximum number of users allowed in each class. Whenever a certain number of packet drops is detected for a class, the corresponding maximum number of users is set to the current number of users minus 3 (10 for all classes). Current users remain active, but new users are refused access or downgraded, to prevent further network overload. When there are no more drops, limits are increased again. The results demonstrate that the use of policies, especially active ones, improves the service offered to users, adapting to the network state.

References

[1] S. Blake et al., “An Architecture for Differentiated Services”, RFC 2475, December 1998.

[2] Paulo Pereira, Paulo Pinto, “Algorithms and Contracts for Network and Systems Management”. IEEE Latin American Network Operations and Management Symposium, Rio de Janeiro, Brazil, 3-5 December 1999, pp.385-396. ISBN: 85-900382-3-8.

[3] VINT Network Simulator (version 2). http://www-mash-cs.berkeley.edu/ns