Available online at www.sciencedirect.com
SCIENCE DIRECT®
Future Generation Computer Systems 21 (2005) 515–524
FGCS
Future Generation Computer Systems
www.elsevier.com/locate/fgcs

# Experiments with equivalent differentiated services in a grid context

Pascale Vicat-Blanc Primet*, François Echantillac, Mathieu Goutelle

*INRIA-RESO, Laboratoire de l'Informatique du Parallélisme, Ecole Normale Supérieure de Lyon, 46, allée d'Italie, 69 007 Lyon, France*

Available online 7 January 2005

**Abstract**

Grids generally rely on a complex interconnection of Internet Protocol (IP) domains that offer heterogeneous services and unpredictable performance characteristics, particularly at the local area network/wide area network boundary. The total lack of end-to-end resource control in IP networks is responsible for performance problems that may affect the whole Grid environment. An end-to-end service differentiation architecture that controls heterogeneous communication performance is thus needed. We propose the *Equivalent Differentiated Services* (EDS) architecture, based on a layer-4 service differentiation solution exploiting a new layer-3 relative DiffServ model. In this paper, we present the *EDS* packet forwarding principles, the router mechanisms and two adaptive packet marking algorithms. As a proof of concept, we have implemented the *EDS* architecture in Linux and performed experiments on a transoceanic testbed.


*Keywords:* Grid networking; End-to-end performance; Service differentiation; Quality of service; Adaptive packet marking; EDS

## 1. Introduction

Grid traffic can be decomposed into two categories of flows: bulk file transfers (e.g., replications of large databases) and time-sensitive transfers (e.g., real-time applications or interactive web traffic). However, the underlying communication infrastructure of these large scale distributed environments is a complex interconnection of multi-Internet Protocol (IP) domains that present bottlenecks and changing performance characteristics, particularly at the local area network/wide area network (LAN/WAN) boundary. Consequently, end-to-end flows experience varying latency (due to queues on network interfaces or traffic loss that can damage dramatically the end-to-end throughput. The lack of end to end resource control in IP networks is responsible for performance problems that may affect the whole Grid environment.

There are basically two solutions to control and differentiate the communication performance [1]: *end-based QoS control* consists in masking the variability of network performance by using compensation mechanisms in the applications to meet specific QoS objectives; *network-based QoS control* consists in adding control mechanisms inside the network in order to avoid performance variation and offer some guarantees for the transfer.

* Corresponding author.
*E-mail address:* pascale.primet@ens-lyon.fr (P.V.-B. Primet).

0167-739X/$ – see front matter © 2004 Elsevier B.V. All rights reserved.
doi:10.1016/j.future.2004.10.009

The advantages of the first approach are two-fold: the application monitors the experimented QoS, detects variation and reacts appropriately; and the network is not modified. The drawbacks of this approach are three-fold: in some cases, the adaptation is not sufficient and the functionality of the application itself may be incorrect; the end-to-end applications have to take care of network details and be network-aware, which makes their programming and debugging more complex; and the context-dependent adaptive mechanisms are difficult to reuse.

Consequently, a large effort to provide QoS support at the network level has been done during the last 10 years. The advantages of network-based QoS approaches such as IntServ [2] or DiffServ [3] are two-fold: the network offers strict or statistical guarantees, and the application is not modified. The drawbacks are three-fold: resource reservation and appropriate packet scheduling mechanisms have to be integrated and configured in routers [4]; the offered guarantees may be higher than what applications really need, making the overall architecture unnecessarily complex; and the approaches face deployment problems as they impose per-domain resource provisioning, admission control and interdomain interoperability. The result is that network-based QoS architectures are not yet widely and homogeneously deployed in IP networks.

This paper explores an alternate hybrid approach that offers a solution for end-to-end service differentiation at the transport level. This model is based on relative service differentiation at the IP level (soft network QoS control) merged with performance adaptation at the transport level (end-to-end adaptation). This model avoids classical IP QoS approach drawbacks as it does not impose strict performance objectives. *Equivalent Differentiated Services*, (EDS) relies on a relative and asymmetric packet level service differentiation that allows end-to-end performance adaptation.

The rest of this paper is organized as follows. Section 2 develops the EDS principle, architecture and layer-3 mechanisms. Section 3 describes a set of adaptive packet marking algorithms exploiting the EDS network layer to provide specific layer-4 service differentiation. Section 4 analyzes the results obtained in experimentations in the DataTAG testbed. Section 5 investigates related work. Finally, conclusions and perspectives are given in Section 6.

## 2. Equivalent differentiated services

This section presents the EDS model and the layer-3 mechanisms.

### 2.1. EDS goals and principles

Our *Equivalent Differentiated Services* scheme [5] has two main goals:

- Improving the best-effort service provided by the network by adding suitable differentiation mechanisms where congestion occurs. We assume that, in Grids, the main bottlenecks are located at the LAN/WAN boundary.
- Enhancing the global performance of a mix of traffic and not privileging any type of flow.

Best-effort packet forwarding is replaced by a new service discipline and queue management mechanism that is able to accelerate or assure packet forwarding *locally*. In *EDS*, unlike DiffServ, no absolute bounded performance from one end to another is guaranteed by a continuous concatenation of particular forwarding treatments called Per Hop Behaviors (PHB). Consequently, the *EDS* architecture avoids the bounded domain concept and control plane features such as resource reservation and admission control required by the traditional DiffServ architecture.

To provide predictable end-to-end performance, transport protocols adapt end-to-end performance to the specific needs of applications and the actual load of the network, by allocating appropriate packet code point, a feature known as *adaptive packet marking*. Different types of adaptive packet marking algorithms, integrated in standard transport protocols, are designed to fully exploit the network differentiation.

### 2.2. EDS PHB model

The *EDS* scheme provides a range of different but equivalent packet forwarding classes that make a trade-off between delay and loss rate for each packet in each *EDS* network node. In this model, each class experiences a specific relative and local per-hop behavior, which includes the differential treatment that an individual packet receives. The *EDS* PHB model defines an arbitrary number ($N$) of equivalent packet forwarding