

Business-driven Management of Policies in DiffServ Networks

Antonio Astorga and Javier Rubio-Loyola

Network Management Group, Departament de Teoria del Senyal i Comunicacions.
Universitat Politècnica de Catalunya.
{aastorga, jrubio}@nmg.upc.edu

Abstract. This paper presents a framework to control the life cycle of enforceable policies aligned to business directives. The framework relies on three main aspects: the formalization of measurable business indicators underpinning the management cycle of policies, a holistic approach to conceal Quality of Service delivery with business alignment, and the use of policy-based management as the vehicle to control system behavior accordingly. The core contribution of this research lies in the network management area whose most widespread solutions for QoS delivery have been systematically decoupled from business value, although the research community recognizes business profit as the main motivation for any management solution. The ultimate goal of this research is to develop a management framework that allows exploiting business value in telecommunications infrastructures. As each application domain may have intrinsic peculiarities we propose to validate our approach in the context of DiffServ networks. Simulations will be conducted to evaluate and to optimize the performance of diverse business directives under different patterns of service invocations and patterns of inter-domain traffic exchange between autonomous systems. To the best of our knowledge no other approach has dealt with the above research area in such a holistic view.

Keywords: business driven management, business indicators, policy life cycle, network incidents, profit maximization.

1 Introduction

In networks supporting different classes of services it is expected that users will enjoy different levels of service for the traffic injected into the network. On the one hand, users should be encouraged by service providers to exploit the managed network resources. On the other hand, the allocation of resources to assess the arranged quality of service may need to be adjusted accordingly. The same behavioural pattern applies with inter-domain traffic exchange.

It is widely accepted that without control of any kind like price, admission control, resources limitation, etc, traffic injection would produce resources starvation and consequently network congestion [1]. Statistical phenomena like congestion may eventually cause service degradation which in turn would have some impact on the

business value. From an economical view, SLA violation would have some impact on the profit, reward, and refund strategies of the administrative parties.

In the network management area, the research community has directed some efforts developing mechanisms to deliver end-to-end Quality of Service in the Internet. Mechanisms for network congestion prevention and solving, control of service subscriptions and invocations, and dynamic traffic engineering functions have been the centre of study in potential intra-domain [2] and inter-domain [3] network management solutions. Moreover, although these solutions have been proved to be efficient to guarantee QoS delivery, the requirements, implications and the incremental efforts to elevate their business value have remained almost unexplored.

The ability to carry out business- and QoS-oriented network management introduces several challenging problems addressed in this work. Initially, business strategies must be properly modelled with appropriate business indicators, pivotal for the management of policies. Second, business indicators should be monitored and modelled as functions of measurable parameters of the managed systems. Third, the dynamicity of events occurring in the managed network should be constantly evaluated as to define proactive and corrective management actions enforced through policy. The creation, deployment and modification of policies in runtime should be devoted to optimise business value, all in all under QoS delivery constraints. These problems make this research highly challenging, mainly when we consider a holistic approach to optimise business value under different patterns of resources utilization, patterns of traffic exchange between administrative domains and diverse network topologies. We intend to evaluate the framework in the DiffServ Network Management application domain.

2 Application Domain – particular objectives of our research

We strongly believe that each application domain should have inherent peculiarities to both maximize business profit and to provide QoS. Developing a bi-directional framework of this kind, valid for any application domain may be very difficult if not impossible. This research lies in the DiffServ Network Management application domain. In principle we propose a framework that exploits the mechanisms of the TEQUILA architecture [2] – Traffic Engineering for Quality of Service for the Internet at Large Scale – which to the best of our knowledge is the only validated network management approach that brings together Service Management and Traffic Engineering functionalities for QoS support in next generation IP Networks. TEQUILA uses policy based management as the key enabler for programmability.

The Service Management part of the TEQUILA architecture has two objectives: to control the traffic entering the network and to commit with the service provider's QoS guarantees. This research deals with the critical nature of achieving these two objectives relying on the dynamic evaluation of business indicators affected by the amount of traffic entering the network. In this sense we propose an overlay approach to dynamically enforce the most optimal admission control settings for service subscriptions and service invocations. Similarly, our proposed overlay approach dynamically enforces the most optimal preventive and corrective actions due to

statistical network conditions affecting the business indicators, for example resource starving users (or other service providers), or potential states of network congestion.

The Traffic Engineering functionality of the TEQUILA architecture is concerned with the management of physical network resources. An off-line dimensioning process is responsible for mapping the predicted traffic demand to the physical network resources. In addition, real-time operations are implemented as the means to first, balance the load amongst the established Label Switched Paths (LSPs) in the network, and second, to ensure that link capacities are appropriately distributed among the different Per-Hop-Behaviours (PHBs) sharing each link in the core network. In this functionality our proposed approach dynamically triggers the most suitable policies for the centralized off-line dimensioning process devoted to maximize business value. In addition, our overlay approach dynamically evaluates the business indicators affected by statistical fluctuations of traffic in the core network and enforces real-time operations accordingly.

So far, the policy refinement process in this particular application domain [4] has been carried out taking into account pure QoS-oriented aspects, isolated completely from any business considerations. In summary this research work proposes to introduce a business optimization overlay that dynamically activates the most appropriate set of service management and traffic engineering enforceable policies as a result of a dynamic evaluation of the business indicators.

3 Business indicators and management of policies – Foreseen technical approach

The chosen approach to address the research problems of this work is inspired in game theory principles which in turn have enabled recent business-oriented advances in IT Service Management [5], [6]. The proposed solution is an overlay environment in which business strategies are formalised through measurable business indicators. The approach is devoted to prioritise active changes in the managed network in order to prevent negative effects on the business indicators due to statistical changes of the network. The active changes are enforced through policies. More concrete, the framework basically consists in a multi-thread environment that *collects* information from the managed network, *analyses* it against the business indicators and *enforces* changes in the active policies as to maximise the business value due to actual states of the network threatening business value.

Consider for instance a business indicator linked to “the profit generated from a certain type of customers”. The overlay framework will produce appropriate sets of policies devoted to maximize the profit obtained from that type of customers under statistical changes in the network utilisation, taking into account Quality of Service constraints. The approach is intended to handle several business indicators.

A critical issue in this research is the modelling that formalises the relationships between the business indicators and pure technological aspects of the DiffServ management domain. Initially, the overlay framework considers the identification of potential *incidents* affecting the business. In the DiffServ domain incidents may be simple ones like a “*service rate threshold crossing*”, or composed like “*a service rate*

threshold crossing under network congestion state". A model establishes formal relationships to link incidents both, simple or composed, with business value threatens. For example, the fact that the network is congested or in normal state may have a measurable impact on service level degradation. Also, traffic injection threshold crossings may imply that users may be starving the network, under-utilising it or injecting traffic according to the pre-signed SLA. In this context the modelling formalises the effects of an incident up to the SLAs. For instance, an incident may affect a set of Per-Hop-Behaviours (PHBs) in a network link. Each PHB affected by an incident may in turn affect a set of services and finally, a service may have some effect on a given SLA. SLAs may eventually be linked with user information or with other service provider's information. Incident effect trees are the source of information to carry out the *analysis* and consequently this information should be modelled properly.

The complexity of the *analysis* stage of the framework lies in the generation of policies aimed at minimising the negative effect of the above incidents on the business indicators throughout system execution. The central part of the *analysis* deals with the formalisation of the effect, from now on referred to as *impact*, of incidents over the *business indicators*. The total impact of an incident i should consider all the business indicators affected by such incident. As an incident may have different impact on different business indicators, the chosen approach considers the formalisation of this situation by means of weights [6], [7]. Weights represent the degree of importance ω that an incident i has on a business indicator j .

The analysis stage evaluates the total impact of incidents occurring on the network affecting all the business indicators and defines priority actions *enforced* with QoS policies. An automatic policy refinement approach [4] deploys QoS-oriented policies for this last *enforcement* step.

The evaluation of priority actions is by no means a trivial task. The DiffServ application domain is a multi-service environment with shared resources. Priority actions committed to some user's (or administrative domains) may have some influence on other business indicators, most probably linked to other users of the network. A trade-off processing step committed to stability is mandatory at this stage of this work. As priority actions are enforced through policy, stability passes through the management cycle of policies. An event handler enables the communication amongst the monitoring sub-systems of the managed network and the overlay system itself. This way the incidents affecting our business indicators are correlated dynamically and in case of further priority actions are needed those can be effectively scheduled.

4 Concluding Remarks

This paper has described our ongoing research work towards controlling the life cycle of enforceable policies aligned to business directives for the DiffServ network management area. We rely on game theory principles applied to the business area in which business indicators underpin the prioritization of corrective actions through policy once threatening incidents affecting the business are detected. The approach is

to be validated with an overlay framework embedded to the network management functions of the described application domain.

So far we have built a strong knowledge in the DiffServ network management application domain, in particular the QoS-oriented policy refinement process [4] with a validated methodology and appropriate tools. We are currently defining the model to bring QoS DiffServ Network Management and business alignment into a unified framework.

Our immediate future work will be the definition of a holistic scenario in which we demonstrate our approach with realistic business indicators. Implementing our overlay framework and integrate it with our existing functional refinement solution is also part of our future work. Our intentions are to validate and refine our solution through OPNET simulations. We intend to execute scenarios to analyze different patterns of user's utilization and patterns of traffic exchange between different administrative domains. The ultimate objective of our research is to define patterns of policy management procedures to minimize profit losses and possibly to maximize profit due to patterns of behavior from users and eventually from other service providers e.g. for inter-domain issues. The results of this research can be further validated in real life scenarios in which traffic traces collected from real-life capturing process may be used to prove the effectiveness of the patterns induced in this holistic study.

We strongly believe that the management of policies constrained to business perspectives is an issue that is tightly coupled to application domains. With this regard, our contribution lies in the network management area in which to the best of our knowledge, not a single approach to manage policy in this context has been proposed so far as agreed by the research community.

Acknowledgments. This paper is supported in part by the IST-EMANICS Network of Excellence (#26854).

References

1. Chang X, Petr D.W., A survey of pricing for integrated service networks. Elsevier Computer Communications, Volume 24, Number 18, 1 December 2001
2. EU IST TEQUILA – Traffic Engineering for the Internet at Large Scale www.ist-tequila.org
3. EU IST MESCAL – Management of End-to-End Quality of Service Across the Internet at Large Scale www.mescal.org
4. J. R. Loyola, et al. "A Methodological Approach towards the Refinement Problem in Policy-based Management Systems" IEEE Communications Magazine, October 2006
5. Issam Aib. A business-driven approach to policy optimization. University Pierre & Marie Curie July 2007
6. Bartolini, C.; Salle, M.; "Business Driven Prioritization of Service Incidents", DSOM 2004
7. Buco M. J., Chang R. N., Luan L. Z., Ward C., Wolf J. L., Yu P. S., "Utility computing SLA management based upon business objectives", IBM Systems Journal, Vol 43, No 1, 2004