End-to-End QoS Frameworks for Heterogeneous Networks - A Survey

B. Iancu, V. Dadarlat, A. Peculea
Technical University of Cluj-Napoca, Computer Science Department
26-28 Gh. Baritiu St., 400027 Cluj-Napoca, Romania
{Bogdan.Iancu, Vasile.Dadarlat, Adrian.Peculea}@cs.utcluj.ro

Abstract

One of the main concerns of the current Internet, in terms of Quality of Service, is guarantying end-to-end delivery of services over heterogeneous networks, where more than one management entity exists. Due to the complex, underdeveloped business relations among ISPs and a low level of trust, new QoS frameworks are being designed to guarantee reachability, availability and network performance. This survey is an effort to briefly discuss some of the developments in the QoS for heterogeneous networks research area. The survey is part of the ongoing QAF research project, which aims to develop an end-to-end QoS aware framework for networks and middleware.

Keywords – *end-to-end QoS, heterogeneous networks, survey.*

1. Introduction

Several definitions for Quality of Services (QoS) have been identified throughout the years in different standards and bibliographical references, without any unique and exhaustive formal definition. The most ilustrative definitions of the concept are: [ISO 8402/1986] states that quality is the totality of features and characteristics of a product or service that bears on its ability to satisfy stated or implied needs. The product, respectively the service, represents the result of the activities or processes within the system; [ITU-T, 1994] defines the collective effect of services' performance, which determines the level of satisfaction of the service user; [ISO/IEC X641] defines the qualities that refer to the way an object or a group of objects (components) collectively works, or reflect the qualitative performance of the service offered through the network

At network level, QoS represents the network capability to deliver better services for the selected flows over different technologies. The main goal of QoS is to provide priority including dedicated bandwidth, controlled latency and jitter, and improved loss characteristics, which represent the main QoS parameters.

The concept of QoS has emerged with the developing of new services, such as VoIP or multimedia applications that requires bandwidth availability, controlled latency and jitter and improved loss characteristics. Best-effort networks no longer meet their QoS requirements. Also Internet Service Providers and their clients require means for rating and pricing the quality of the offered/received services. MESCAL project points out that the issue of provisioning end-to-end QoS in the Internet is currently being investigated by both research and standardization communities

The definition of heterogeneous network gathers many characteristic and aspects. A heterogeneous network can be seen as collection of computers and other devices from different manufacturers, all working together as a single unit. It can also be defined as a hybrid network with different transmission technology over different links. A comprehensive definition of heterogeneous networks is found in [1]: network portions may be managed by different Service Providers may use different transmission means such as cable, satellite, radio and may implement different protocols such as ATM, IP and MPLS; a network may also be heterogeneous from the point of view of users, who can require different services and have a different availability to pay for them.

Providing end-to-end QoS in heterogeneous networks is a challenge. Each portion of the network may implement different QoS techniques, guarantying the required performances at their level. But, because of the poor communication and the lack of guaranties between management entities, of different portions of the heterogeneous network, the fulfillment of end-to-end requirements have no guarantees. The paper is organized as follows: Section II provides background information related to QoS technologies. Section III provides an overview of the frameworks to be analyzed and compared. Sections IV to VI present and compare the business models, the implementation approaches and the functional architectures of the two. Section VII concludes the paper.

2. QoS end-to-end frameworks

The main frameworks and technologies used for implementing QoS are Integrated Services, Differentiated Services, Integrated Services over Differentiated Services, Multiprotocol Label Switching, traffic engineering. These technologies are detailed described in [1], [2].

Integrated Services have developed a new architecture for resource allocation in order to satisfy the requirements of the real time applications. The main idea is resource reservation for each flow. A flow is identified by five fields in packet headers: source and destination IP addresses, protocol ID and source and destination ports. These five fields are often referred as five-tuple. The aim of the Integrated Services is the preservation of the IP networks datagram based model and in the same time the reservation of the resources for real time applications. In Integrated Services architecture a set of mechanisms and protocols is used for explicit reservation of the resources. Before packet transmission, the applications reserve the necessary resources along the path. The reservation setting begins with the description of the flow characteristics and necessary resources. The network can accept the new application only if there are available resources. After the reservation establishment, the application can send packets along the reserved path. Integrated Services architecture assumes that delay is the main QoS parameter guaranteed by the network.

Differentiated Services Architecture provides several service levels by classifying the traffic into a small number of forwarding classes and allocating resources based on these classes. The individual forwarding class represents the aggregated traffic and is encoded in the IP packet header. Edge nodes classify the packets and condition the traffic. Interior nodes forward the packets based on the forwarding classes in the packet header. The forwarding treatment is described by per-hop behavior (PHB), each PHB being represented by a 6 bit value named Differentiated Services Codepoint (DSCP). All the packets with the same DSCP are referred as behavior aggregate. IP packet header has an 8 bit field named IP TOS. This field is composed of 3 bits precedence, 3 bits type of service and 2 unused bits. Differentiated Services standard redefines IP TOS field in order to indicate per-hop behaviors. The replacing field, named DS, replaces IPv4 TOS and IPv6 traffic class bytes definitions. First 6 bits of DS field are used as DSCP in order to encode the PHB and the other 2 bits are not used.

Multiprotocol Label Switching (MPLS) presents a series of advantages as QoS support through connection orientation, traffic engineering support, VPN support or multiprotocol support.

Traffic engineering optimizes the performance of the network by reducing the congestion and improving resource utilization through traffic distribution management.

Recent QoS end-to-end frameworks, such as [3], [4] try to resolve the problems of quality of service across heterogeneous networks. This paper briefly discusses the developments of the two frameworks, pointing out, in a comparative manner, the novelties of the approaches, the problems solved and also the drawbacks.

3. Frameworks' overview

MESCAL [3] (Management of End-to-end Quality of Service Across the Internet at Large) project was partially funded by the EU as part of the IST Programme, within the Fifth Framework Programme. The project started in 2002 and ended in April 2005. MESCAL scope was to deliver end-to-end Quality of Service across the Internet, by focusing on supporting QoS across multiple domains administered by different organizations.

AGAVE [4] (A liGhtweight Approach for Viable End-to-end IP-based QoS Services) project was partially funded by the EU as part of the IST Programme, within the Sixth Framework Programme. The project started in 2005 and will end in May 2008. AGAVE objective is to develop an approach for deploying, in a lightweight way, open end-to-end service provisioning with QoS in order to support added-value IP-based services. [5] MESCAL can be consider the precursor of AGAVE, in terms of heterogeneous network decomposition and the separation between service providers and network providers, and their roles.

The projects are motivated by the underdeveloped business relations, agreements and specific guarantees on reachability, availability or network performance between peered ISPs. The business relationships between ISPs must be considered in order to provide end-to-end QoS across the Internet. Also explicit QoS information must be part of the agreements between ISPs.

MESCAL research points out the need of new techniques required to propagate QoS-based agreements among the set of providers involved in the chain of inter-domain service delivery. They approach the problem of how current agreements between ISPs should be enhanced to propagate QoS information between domains, and, in the absence of any form of central control, how these agreements may be used together to guarantee end-to-end QoS levels across all involved domains of control/ownership.

AGAVE approaches the problem of end-to-end QoS by studying, developing, and validating an inter-domain architecture based on the novel concept of Network Planes, which will allow multiple IP Network Providers to build and provide Parallel Internets tailored to end-to-end service requirements [6].

The AGAVE functional architecture is described in [7]. It introduces the concepts of Network planes and Parallel Internets to ensure service differentiation with both intra- and inter-domain scope.

4. Business Models

The Internet can be described as a collection of Internet Service Providers (ISPs), viewed as Autonomous Systems (ASes), which must interconnect to allow users of different ISPs to exchange traffic. Usually, ISPs use peering and transit agreements to interconnect and to exchange information about how to reach destinations in the Internet according to their interconnection strategies [8].

Fig.1 shows how different types of ISPs are interrelated. In each case, a subscriber — whether an end user or a lower-level service provider — connects to a higher-level service provider at that ISP's Point of Presence (POP). A POP is just a nearby router to which the subscriber can connect via dialup or a dedicated local loop. At the highest level, the network service providers interconnect via network access points (NAPs). A NAP is a LAN or switch—typically Ethernet, FDDI, or ATM—across which different providers can exchange routes and data traffic [9].

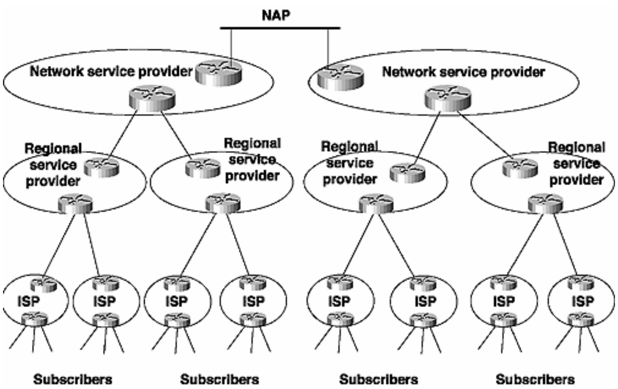

Figure 1. ISPs relations [9]

MESCAL staring point is the two exiting forms of distinct relations between ISPs for traffic exchange, that underline the business agreements: *peer-to-peer* and *transit* (customer-provider) relations between ISP's at different levels of the three-tier model.

The Business Model Stakeholders, which enables inter-domain Quality of Service (QoS) across the Internet, described in [10] and shown in Fig. 2 are:

A. *Service Providers* that interact with IP Network Providers following a customer-provider paradigm on the basis of respective agreements (SLAs).

B. The focus of the MESCAL project is on the interactions between the set of *IP Network Providers* (INPs) involved in the end-to-end delivery of QoS-based IP services, i.e. across multiple domains. The INPs are responsible for their own network domains and for reaching peering agreements with other INPs to extend the scope of the QoS services that they can offer to their customers and other INPs.

C. *Physical Network Providers* provides physical (up to the link layer) connectivity services between protocol-compatible equipment in determined locations.

D. *Customers (Users)* are the beneficiaries of the services provided by Service Providers

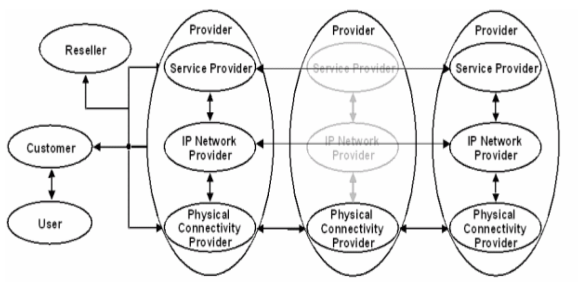

Figure 2. MESCAL Business Model [11]

AGAVE, assumes the previous business model and distinguishes between the service provider (SP) and the IP network provider (INP) business roles (Fig. 3). In terms of business, there is no mapping one-to-one, such that a business entity may implement more than one role.

The relationships between IPNs, and their role, are described as following: INPs offer IP connectivity to SPs but they do not offer their services directly to end customers. INPs interact with each other on a one-to-one relationship basis regulated by INP interconnection agreements (NIAs), thus extending the scope of their IP connectivity. An NIA defines the QoS and availability performance of the traffic exchanged between the INPs, the scope and the profile of the traffic entitled to the agreed performance, and identifiers to capture distinct flows for providing differentiated treat.

SPs are the entities which offer IP-based services to end customers. SPs establish connectivity provisioning agreements (CPAs) with underlying INPs in order to fulfill the IP connectivity aspects of their services. Similarly to NIAs, CPAs specify the performance, constraints, and identifiers of the service traffic entering the INP network from the SP sites. In terms of relationship between SPs and INPs, beyond the connectivity mechanism, the INP offers to the SP the means to control the connectivity provisioning. SPs can set policing and routing rules and also, they receive feedback reports. The specific provisioning rules and required feedback also are agreed upon during the CPA negotiation.

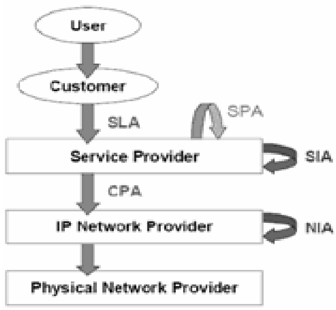

Figure 3. AGAVE business relationships

Another important business relationship is the one between SPs (SP to SP) – SP interconnection agreement (SIA). The content of an SIA is service-specific, for example, a VoIP SIA may include telephony performance metrics, such as average success rate or simultaneous calls capacity [7]. The SP-to-customer relationship is defined by SLAs.

A novelty of MESCAL is the introduction of the *QoS-class (QC)* concept. QC denotes a basic network-wide QoS transfer capability of a Provider domain [10], as a set of attribute-value pairs (e.g. packet loss – upper bound). This concept is further refined as following: in a single domain (only one provider) MESCAL defines l*ocal-QoS-class (l-QC)* and that is the QoS transfer capability provided by the provider itself. To extend this concept to heterogeneous networks (accros multiple domains) Mescal defines the *extended-QoS-class (e-QC)*, shown in Fig. 4, by combining l-QCs or e-QCs.

The relationships between an INP and a customers/a peer INP are set as SLS (where SLS denotes the technical characteristics of a given service in the context of an SLA) and are used to define QoS-based services through multiple domains. Mescal defines several QoS classes that an INP can offer: a) SLS between INPs are called peer SLS *(pSLS)* for inter provider relationship; b)

SLS between IPNs and customers are called customer SLS *(cSLS)* for end-customer – provider relationship. INPs establish pSLS with their peers using the further described cascades model, though maintaining the loosely coupled Internet model. In accordance with the agreed QoS levels, IPNs then provision and configure the network resources. Also the introduction of pSLSs, service agreements, implies legal obligations, thus providing an economic relationship between ISPs and also adding a certain level of trust.

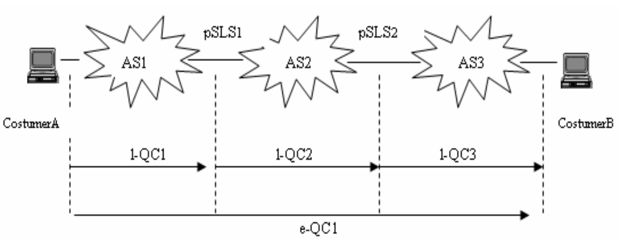

Figure 4. QoS classes

Also, a *QoS Proxy* (upstream-QoS-proxy) concept is proposed to provide stronger business relationships. The QoS-proxy relationship assumes that an ISPs may agree with another neighboring ISP to provide a transit QoS-based connectivity service to (a subset of) the destinations it can reach with a certain QoS level. Agreements may be established independently in either direction, as each ISP wishes. The ISP offering the transit QoS service would have built its QoS reach capabilities based on similar agreements with its directly attached ISPs and so on. In this way MESCAL provisions harder QoS guarantees to specific destinations. This type of business relationship is thought as being the QoS Internet counterpart of call termination agreements in the PSTN or VoIP business world.

5. Implementation approaches

Mescal presents three QoS Peering Approaches [12]: a source-based approach, a cascaded approach, and also a hybrid approach (a combination of the first two approaches).

The source-based approach disassociates pSLS negotiations from the existing BGP peering arrangements. The originating domain must have knowledge of the end-to-end topology of the Internet and then it establishes pSLSs with a set of adjacent and distant domains. The purpose is to reach a set of destinations, with a particular QoS. Fig. 5 [12] shows the fact that the originating domain (AS1) has the responsibility for managing the overall requested QoS service/connection. The provider (AS1) directly requests peering agreement (pSLS1 and pSLS2) with providers AS2 and AS3 or with any other network provider involved in order to create an e-QC (from AS1 to AS3).

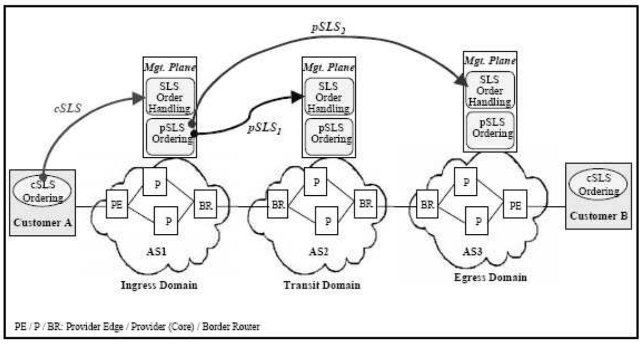

Figure 5. Source-Based Approach

The *cascaded approach* keeps the loosely coupled structure of the internet, because each ISP established pSLSs only with adjacent ISPs (ISPs with whom there are existing BGP peering relationships).

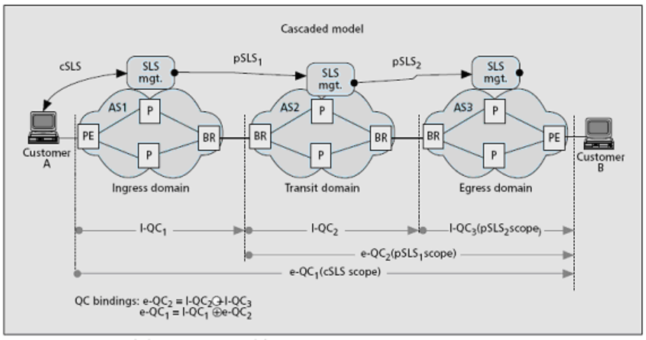

Figure 6. Cascaded-Based Approach

Fig. 6 [11] gives an overview of the operations in this approach. [11] describes the scenario: domain AS3 supports an intra-domain QoS capability (l-QC1). AS2 supports an intra-domain QoS capability (l-QC2) and is a BGP peer of AS3. AS2 and AS3 negotiate a contract (*pSLS2*) that enables customers of AS2 to reach destinations in AS3 with a QoS (e-QC1). This process can be repeated recursively to enable AS1 to also reach destinations in AS2 and AS3, but at no point do AS1 and AS3 negotiate directly. Because there are no explicit end-to-end agreements in the cascaded approach, this solution is more scalable. Every domain builds its own e-QCs to the required destinations based upon the capabilities of adjacent downstream ASs to. This recursive approach results in an approximation of end-to-end agreements.

The above described have also some limitations. For the source-based approach an up-to-date topology of the Internet is necessary, with information about the existence and operational status of every physical link between ASs. The creation, maintenance and update of such a topology represent a drawback from the scalability point of view. Also, in both the source-based approach and the cascaded approach, inter-domain routing is pSLS constrained, i.e. traffic will only pass through ASs where pSLS agreements are already in place.

AGAVE approach is based on realization of two novel concepts: network planes and parallel internets.

The *network planes* concept is used to differentiate the delivery behaviors experienced by IP flows when crossing an IP domain managed by a single INP. The NP notion is internal to INPs, and its engineering can be carried out before or after the formulation of service requirements as expressed by SPs. AGAVE proposes the use of a mechanism that can route traffic with different QoS requirements through distinct paths, thus being able to satisfy the specific performance demands of the traffic. In this case traffic flows may follow different paths to reach the same destination.

One of the INP function is to plan, select, and engineer its NPs to meet the SP requirements. A given NP can be used to convey service traffic managed by the same or distinct SPs in an aggregate fashion, resembling to the DiffServ paradigm.

INPs must engineer corresponding NPs within their own network, to accomplish the service requirements specified in the CPA. A characteristic of INPs is that they can select the most appropriate combination of mechanisms, from the routing, forwarding and resource management dimensions in order to implement specific NPs according to the service requirements, through the combination of several processes.

- *Routing dimension*: individual NPs have to implement different paths in order to provide support for heterogeneous services. The routing dimension has the capability to provide pre-configured backup paths/topologies inside high availability NPs, to permit fast reroute.
- *Forwarding dimension* an INP can deploy different queuing and scheduling mechanisms, e.g. DiffServ PHBs, for traffic belonging to different NPs [6].
- *Resource management dimension*: an IP packet is subject of admission control, traffic shaping and policing mechanisms associated with a NP and its allocated resources (dedicated/shared network capacity).

The concept of parallel Internets is introduced in the sense of horizontal QoS, as classified by [1], and is the concatenation of NPs from across different INPs. This is the key concept in enabling end-to-end QoS over heterogeneous networks, by providing the means to service differentiation across INPs. INPs must negotiate and establish NIAs between each other to bind NPs with similar service characteristics and apply specific mechanisms in the control/data plane to enforce the realization of individual PIs [7].

A novelty aspect of this proposed concept is that each instance of PI is not necessarily implemented in the same way across multiple INPs. This gives a great degree of flexibility allowing each INP to make local decision in binding NPs to the PI. Parallel Internets can be viewed as inter-domain extensions of Network Planes. [5]

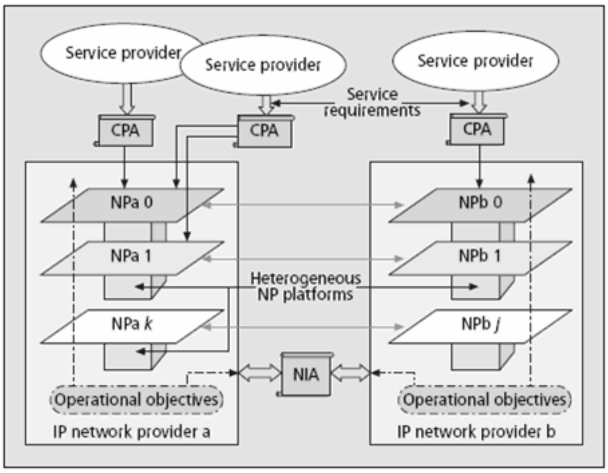

Figure 7. Network planes and Parallel Internet realization

Mescal and AGAVE frameworks are compared from the point of view of end-to-end QoS realization (Table 1), putting emphasis on the specific intra-domain and inter-domain techniques.

MECAL	AGAVE
Intra-domain QoS capability	*Inter-domain QoS capability*
l-QC	e-QC – concatenation of l-QCs
NP	PI – concatenation of NPs

Table 1. End-to-End QoS realization

The Table 2 provides a summarized overview of the implementations discussed in this section. The two frameworks are compared in terms of scalability capabilities, inter-domain routing and reliability.

MESCAL		AGAVE
Scalability		
Source-based approach	Cascaded approach	Network planes/Parallel Internets
-the up-to-date topology of the Internet represents a drawback	-scalable because of the adjacent peering relation	-scalable because of concatenation of NPs (local decision in binding NPs to the PI)
Inter-domain routing		
Source-based approach	Cascaded approach	Network planes/Parallel Internets
-pSLS constrained	-pSLS constrained	-NIA constrained
Reliability		
Source-based approach	Cascaded approach	Network planes/Parallel Internets
-not specified	-not specified	-pre-configured backup paths/topologies inside high availability NPs
Reliability		

Table 2. Comparison of implementations

6. Functional Architectures

The MESCAL architecture comprehends the full set of functions required in the management (service and resource), control and data planes for the provision of end-to-end QoS-based IP connectivity services.

A detailed functional architecture is defined in [11] and Fig. 8, from the perspective of a single provider, showing the main planes, designed to provide inter-domain QoS-based services: 1) management, 2) control and 3) data plane.

First, the Management Plane is composed of the Service Planning and QoS Capabilities Exchange that has the role of defining the services that a provider can offer. Prior to the use of a service, a provider must first learn about the potential services. The QoS Capabilities Discovery is used to discover those services and also their costs. Also, the provider has the capability to advertise its own services using the QoS Capabilities Advertisement function.

The Monitoring and Assurance function is responsible for node and network monitoring and communication with other functional blocks. SLS Assurance is used to monitor SLS and to ensure that the services are provided accordingly to their contracted SLS. The Monitoring and Assurance component has two distinct motivations: the first function is for network management and monitoring, and the second is for financial purposes, assuring a feasible network architecture, both from the Service providers' point of view, and also from customers' point of view.

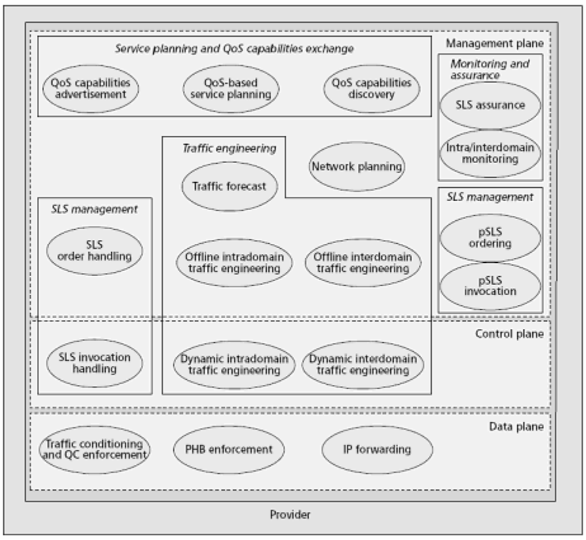

Figure 8. MESCAL functional architecture

[10] defines Network planning and provisioning as off-line processes that are responsible for determining the type, quantity and geographical location of the physical resources required by an IP Network Provider and for ensuring that the physical resources are deployed as planned and with the appropriate physical configuration. The implementation of planning decisions is not an automated process, because it requires manual installation or configuration of physical equipment.

Traffic engineering functional block it is used to estimate traffic demand, and this estimation is in turn used to provision inter and intra-domain traffic. The offline inter-domain and intra-domain traffic engineering functional blocks are used closely together to construct potential e-QC that meet the service requirements and optimize the e-QC. The interaction between the intra and inter TE is construct in such a manner to avoid e-QC that meet the inter-domain requirements, but not the intra-domain requirements, avoiding thus potential bottlenecks or overloads.

The SLS management functional blocks at the management plane are pSLS order handling (performs subscription-level admission control.), pSLS ordering (receives and negotiates requests from offline interdomain TE for new pSLSs) and pSLS invocation (handles pSLS invocations).

The SLS management functional blocks at the Control Plane receives requests from customers or peer providers for pSLS or cSLS invocations and also performs admission control. The Dynamic inter-domain TE uses real-time monitoring information for inter-domain routing and load balancing, using information from Offline inter-domain TE block. Similary, the Dynamic intra-domain TE uses information from Offline intra-domain TE block and performs, in real-time, intra-domain routing, load balancing, and dynamic bandwidth assignment.

Traffic conditioning and QC enforcement, in the Data Plane is responsible for packet classification, policing, traffic shaping, and marking accordingly to the agreed SLSs, in the same manner as DiffServ. PHB enforcement functional block is used for traffic queuing and scheduling in order to realise the different PHBs.

An overview of the AGAVE functional architecture is presented in [7], [6], from the perspective of the interfaces that are to be supported by the business roles.

The functional blocks are summarized from [6] and shown in Fig. 9. The components are presented from the Service Providers perspective and from the Internet Network Providers perspective.

The SP functional blocks are grouped in three main modules: the Customer Interface, the SIA Interface and the CPA Interface with the INP.

The Customer Interface provides the Customer Service Advertisement function that is responsible to advertise the service provider's capabilities and service offerings, as generated by the *Service Planning and Engineering* function, to potential customers. The *SLA Order Handling* function is to negotiate SLA set up or modification requests with the customers. The *SLA Assurance* function assumes the actions related to SLA assurance: reports or notifications in relation to the fulfillment of the effective service to the SLA terms.

The CPA interface uses *Network Services Discovery* function to discover the INPs' capabilities and send them to the the *Service Planning and Engineering* function. The *CPA Ordering* function is used for the negotiation of requests for set up, execution of actions, modification, explicit verification and teardown, originated from the *Service Planning and Engineering* function. *CPA Verification* function has the task of processing the received CPA verification reports and notifications sent by the INPs

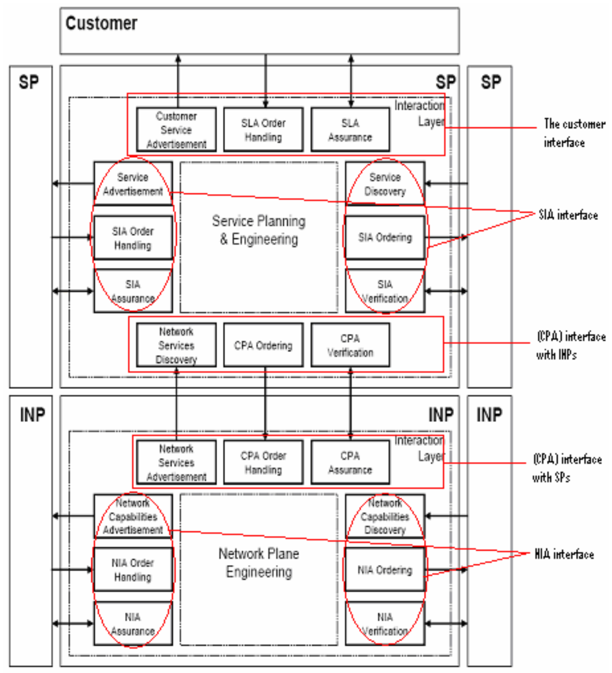

Figure 9. AGAVE functional architecture [6]

The SIA interface has the *Service Advertisement* function to advertise the *Service Planning and Engineering* available services and the *Service Discovery* function to discover the corresponding capabilities of other service providers. SIA orders issued by peer service providers are managed by the *SIA Order Handling* function. The *Service Planning and Engineering* function performs also the task of implementing the requests. The SIA order requests generated by the *Service Planning and Engineering* function are negotiated with the implicated peer service providers through the *SIA Ordering* function. The SIA assurance reports and notifications generated by the *Service Planning and Engineering* function are sent through the *SIA Assurance* function.

SIA Verification function has a similar task as CPA Verification, that being the task of processing the received SIA verification reports and notifications sent by the peer service providers.

The *Service Planning and Engineering* functional block is also responsible to set guidelines for the Interaction Layer to manage SIAs with other service providers and CPAs with INPs.

The INPs functional blocks are similar with the SP functional blocks. It has the CPA Interface to interact with the SPs and the NIA Interface for the interaction with other INPS, thus extending the scope of their IP connectivity.

The CPA interface uses *Network Services Advertisement* function to announce the INPs' capabilities, as formulated by the *NP Engineering* functions, to potentially interested service providers. For the establishment or modification of CPA with the service providers the *CPA Order Handling* function is used. *NP Engineering* functions have the responsibility of the fulfillment of CPAs. The *CPA Assurance* function is used in a similar way as the previous described SLA assurance function.

The NIA interface announces the interconnection capabilities of the INP using the *Network Capabilities Advertisement* function and discovers the capabilities of other INPs through the *Network Capabilities Discovery* function, sending them to the *NP Engineering* functions. The *NIA Order Handling* function receives NIA order requests for establishment, execution of permissible actions, modification or teardown issued by peer IP network providers, but the actual implementation of the received request is the responsibility of *NP Engineering* functions. *NIA Ordering* function is used for the negotiation with the implicated peer IP network providers, and it supervised by the *NP Engineering* functions by the use of *NIA Assurance* function. A *NIA Verification* function is also used.

Several Mechanisms to implement Network Planes and Parallel Internets have also been investigated [13]. We can classify the specific mechanisms based on their scope – intra or inter domain:

a. Intra-domain scope

- Multipath Routing with Dynamic Variance: used for resource optimization, such as load balancing or congestion delay, by computing and selecting alternative paths to be used to route the different types of traffic according to their QoS requirements
- Multi-Topology Routing: used for achieving service differentiation through diverse paths implemented with dedicated IP routing topologies for individual Network Planes and for dynamic traffic control across topologies within an NP.

b. Inter-domain scope

- q-BGP [3]: used for QoS-based route advertisement and route selection, supporting multiple QoS planes
- Resilience Aware BGP TE: used for intra/inter-domain traffic optimization under normal conditions (no failure) or link failure.
- Overlay routing: used for fast recovery from failures and network congestions, using pre-computed paths.

7. Conclusions

Providing end-to-end communication and service guarantees in heterogeneous networks is a challenge. Both MESCAL and AGAVE had defined a series of abstractions to hide the differences of heterogeneous internetworking technologies and to enable applications to operate and to provide services across them. A high level description of the architecture of the two frameworks was presented.

The MESCAL architecture is a generic one, which allows a variety of different performance guarantees to be provided. The model defines the functional blocks an INP needs to deploy in order to support QoS over multiple providers, thus MESCAL is suited for providing QoS over heterogeneous networks. MESCAL proposes suitable business models to provide both loose, qualitative and statistical, quantitative end-to-end QoS guarantees.

AGAVE project proposes an approach to ease the introduction of differentiated services. A clear separation among services and network aspects was achieved through the introduction of two new concepts: network planes and parallel Internets. Quality of service, reliability and availability will be enabled through different Network Planes within domains, whose interconnection will form global Parallel Internets that will support the needs of well-defined services, as stated in [6]. The proposed business model and the functional architecture suppose a decoupling between the service provider and IP network provider roles, the complexity being pushed to the INPs, achieving a lightweight approach for the SPs. The separation of service and network concerns will allow the deployment, by the SPs, of added-value services in addition to plain Internet access. Several techniques and mechanisms have also been investigated to implement network planes and parallel internets. The main goal stated by the AGAVE (A liGhtweight Approach for Viable End-to-end IP-based QoS Services) project is to solve technical problems related to end-to-end provisioning of QoS-aware services over IP networks.

Acknowledgments

This work was supported by the “QAF - Quality of Service Aware Frameworks for Networks and Middleware” research project within the framework “National Research, Development and Innovation Programme” initiated by The National University Research Council – Romania (CNCSIS).

References

[1] M. Marchese, “QoS Over Heterogeneous Networks”, Morgan Kaufman, 2007.

[2] Z. Wang, “Internet QoS: Architectures and Mechanisms for Quality of Service”, Morgan Kaufmann, San Francisco, 2001.

[3] www.mescal.org

[4] www.ist-agave.org/

[5]ftp://ftp.cordis.europa.eu/pub/fp7/ict/docs/netmedia/20080416-17-agave_en.pdf, 2008.

[6] D1.1: Parallel Internets Framework, http://www.ist-agave.org/results/D1.1-final-public.pdf, 2006.

[7] Boucadair, M. Levis, P. Griffin, D. Wang, N. Howarth, M. Pavlou, G. Mykoniati, E. Georgatsos, P. Quoitin, B. Rodriguez Sanchez, J. Garcia-Osma, M.L., “A Framework for End-to-End Service Differentiation: Network Planes and Parallel Internets”, Communications Magazine, IEEE**,** Volume: 45, Issue: 9, pg. 134-143, 2007.

[8] P. Ferreira, M. Sirbu, ”Inefficiency in provisioning interconnected communication networks”, Proceedings of the 6th ACM conference on Electronic commerce, pg. 137 – 146, 2005.

[9] J. Doyle, J. DeHaven, “Routing TCP/IP, Volume II (CCIE Professional Development)”, Cisco PressPub, 2001.

[10] MESCAL deliverable D1.3: “Final specification of protocols and algorithms for inter-domain SLS management and traffic engineering for QoS-based IP service delivery” Chapter 2, (http://www.mescal.org/deliverables/d1.3_finalv2.pdf), 2005.

[11] Howarth, M.P.; Flegkas, P.; Pavlou, G.; Ning Wang; Trimintzios, P.; Griffin, D.; Griem, J.; Boucadair, M.; Morand, P.; Asgari, A.; Georgatsos, P., “Provisioning for interdomain quality of service: the MESCAL approach”, Communications Magazine, IEEE Volume 43, Issue 6, pg. 129 – 137, 2005.

[12] P. Georgatsos, J. Spencer, D. Griffin, P. Damilatis, H. Asgari, J. Griem, G. Pavlou, P. Morand, "Provider-level Service Agreements for Inter-domain QoS delivery", Proc. of Fourth Int. Workshop on Advanced Internet Charging and QoS Technologies (ICQT04), Springer, September 2004.

[13] D4.1: Test Specification and Experimentation Plan, http://www.ist-agave.org/results/D4.1-final.pdf, 2007.