Abstract: *Currently more and more applications work distributed and use Internet Protocol based transport services on a best effort network base. A large number of applications has special requirements concerning required traffic characteristics and parameters like bandwidth, delay, jitter or loss of information. The applications do not need the highest level of quality for all these parameters at the same time. Some applications are mostly concerned with bandwidth, others with delay and so on. This leads to the approach to support a limited number of specific Network Services (NS), where each NS is optimized for a group of applications, like loss critical or streaming applications.*

The AQUILA project proposes a light DiffServ model for IP core networks with four quality specialized NSs, beside the standard best effort, named Premium Constant Bit Rate (PCBR) and Premium Variable Bit Rate (PVBR) for UDP/RTP based applications, and Premium Mission Critical (PMC) and Premium Multi-Media (PMM) for TCP based applications. A set of parameters is collected and stored in application profiles, which are built per application generically. These sets are formalized and structured in templates as Service Level Specifications (SLS). This intention is fixed in an IETF draft prepared by the European project consortium. Real applications were chosen to prove the NS approach in a trial.

1 Quality of Service in the IP World

The popularity of the Internet is still rapidly growing, and in particular, there is a clear trend towards serious and critical business transactions over the Internet. Examples for this kind of applications are online banking and brokerage, electronic commerce and multimedia communication. However, the only kind of network service, which is currently supported in the Internet, is the “best effort” model. The network is not yet able to properly adjust itself to the concurrent traffic flows of different types. For the customer, the consequence is that all transactions over the Internet still have to be considered as unreliable. So it is currently necessary to always provide a “traditional” (non-Internet) way to perform the transactions, using for instance the public telephone network. For a next wave of Internet applications, customers will demand Quality of Service (QoS) guarantees from the network operators.

In this paper, the view of a network operator is taken and therefore specifically network-based techniques for QoS are discussed. There exist various application-based quality adaptation methods, which are not discussed here; however, also these methods ultimately rely on a basic quality guarantee from the underlying network. So it is assumed that the customer establishes a so-called Service Level Agreement (SLA) with an Internet Service Provider (ISP) or a network operator. The SLA is a contract defining clear responsibilities for the operator. The most important part of the SLA is a precise specification of the kind of network service required from the customer. The technical parameters are fixed in the Service Level Specification (SLS). This paper describes a concrete approach to specify Network Services (NS) and to check for their correct realization. The information given here was developed in the European IST research project [IST] AQUILA [AQUILA]. In AQUILA, a layered resource control architecture for the Internet is developed and concrete trials for proof of QoS are carried out. The architecture and implementation are described in [Wint00]. Here it should be sufficient to state that it is essentially based on a Differentiated Services (DiffServ) approach, enhanced with sophisticated control and management layers [Hust00].

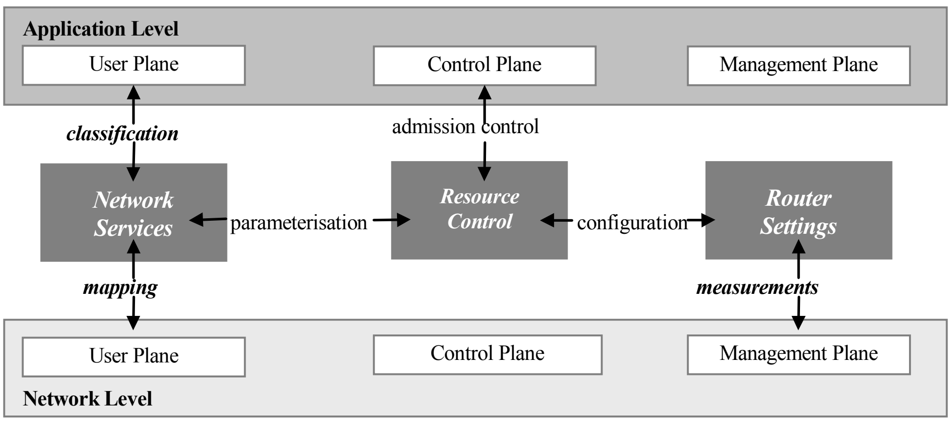

Figure 1: *Network services and measurement support as intermediate between application and network level.*

This paper is structured as follows: In section 2, the requirements for SLSs from a customer and application point of view are analyzed. It is argued that a few NSs are sufficient for the majority of SLSs. Section 3 explains the defined set of NSs on real networks. The aspects discussed in these two sections are brought together in section 4, where the upper and lower peers of NSs are described.

2 Application and End-user Requirements

The default best-effort service, the standard offer in the Internet, is adequate for many accepted applications [Krau00] like e-mail, basic web surfing, and text-based live communication using web chat lines, and text oriented games, for example Multi-User Dungeons (MUDs). However, the best-effort service is inadequate to support real-time applications such as voice over IP and video conferencing. In the case of a congested network the end-user may not be able to use an application. The application's execution is not like it is expected for a normal usage. For example if the packet delay and jitter get too big, a phone call can be impossible: the delivered message is not anymore comprehensible for a human. In the case of an online "real-time" shooting game if the players do not get the information in "real-time" and synchronized playing can be impossible. This leads to a relatively low acceptance and usage [GVU98] of such applications by end-users.

Applications are at the edge between the network and the end-user, they are used by the latter and feed the network with traffic that influences their behavior.

At *network level*, they, on the one hand, produce complex traffic flows, and on the other hand they are influenced by the behavior of the network. The complexity of the traffic flows depends on the codecs used, the application itself, the implementation etc. The traffic flows are a function of many factors at different levels (traffic type, transport protocol, connection rate, living time, bit rate, micro-flows, and packet size) that can vary from stream traffic to elastic traffic, or from UDP traffic to TCP traffic. Application's behavior is influenced by the network QoS parameters: throughput, delay, jitter, packet errors and loss ratio affect differently, as depicted in Table 1, the application performance. ** indicate a strong dependency. * indicates importance with a limit of flexibility.

Application Group	Application Type	Throughput	Delay	Jitter	Error	Loss
1	IP telephony		**	**		*
	MPEG 1/2 A/V	*	**	**		*
2	MPEG 4/7 A/V	**	**	**	*	**
3	Audio on demand	*		*		*
	Video on demand	**		*		*
	Internet TV	**		*		*
4	Online banking		*		**	**
	Online games		**		**	**
	Online trading		**		*	*
5	Chat		*			
	Web browsing		*			
	Telnet		*			*
	classical e-mail					

Table 1: QoS parameters of popular, potential profitable applications and their similarities.

An analysis of Table 1 enables a grouping of Internet applications into five main groups, characterized by similar requirements as depicted in Table 2. It is not a mandatory list of characteristics for all imaginable applications of one group. A compromise is needed to keep the number of groups low.

Application types group	QoS parameters	Traffic characteristics
1 Audio-(video) based conversational, and interactive multimedia applications where audio (video) data are exchanged and retrieved in soft real-time with constant data flows	• delay ≤150ms • loss: low • jitter: very low	• Stream (time integrity preservation) • Low rate • Constant bit rate • Long live • UDP
2 Video-(audio) based conversational, and interactive multimedia applications where video (audio) data (large amount of data) are exchanged and retrieved in soft real-time with variable data flows	• delay ≤150ms • loss: low • jitter: very low	• Stream (time integrity preservation) • High rate • Variable bit rate • Long live • UDP
3 Streaming multimedia applications where mul-	• jitter: low	• Elastic (loose time requirements)

Application types group	QoS parameters	Traffic characteristics
timedia data are downloaded by the client, buffered and then retrieved, but where the time relation between the entities making up the stream has to be preserved	• loss: low	• High rate • Variable bit rate • Long live • TCP
4 Mission critical, secure, retrieval, and games applications where discrete, real-time data are reliably exchanged	• delay: low • loss: very low	• Elastic • TCP • Low rate • Short live
5 Non-time critical basic web applications for distribution or human interface based actions	• none	• none

Table 2: Application types groups with QoS and traffic characteristics.

It becomes obvious that a solution, at network (and end-user) level, consists in offering to the application a number of transport options supporting the different QoS requirements. A limited but powerful set of network offers, called *Network Services (NS)* has to be pre-defined, as discussed in section 3 below.

At *end-user level*, it is important for the network operator, within the scope of an SLA negotiation with a customer, to use pre-defined information using the end-user vocabulary. User-friendly descriptions for QoS correspond to a universal apprehension of applications and ideally make reference to well-known similar "services" from everyday life like: tv, video recorder, hi-fi, or phone etc. The "video conferencing" type, for example, is described by the service component types: "video", "voice" and "data". The service component "video" again is described by attributes like "picture size", "standard tv quality" etc. This information takes the form of files listing common application types and possible quality levels the so called *application profile*. The information is associated hierarchical to application types, service component types in order to increase the universality of the description. A converting mechanism introduced in section 4.1 then maps the concrete customer requirements to the required NSs.

3 AQUILA Network Services Design

In the AQUILA project [AQUILA], a number of transport options for user IP traffic was defined, so called Network Services (NS). The NS and their characteristics are defined by the network operator. The idea of these services, as described in the former chapters, is to provide a few specific offerings from the network operator to the customer, which are relatively easy to understand, and are suitable for a specific group of applications, and which can be practically realized in large networks.

3.1 Definition Set

A Network Service

- describes how customer traffic is technically handled while passing the network,
- is implemented by one (or more) Traffic Classes,
- aggregates similar packet flows.

In the AQUILA project five NS have been identified and implemented, as listed in Table 3 to satisfy the variety of requirements [RiSa00] similar to the application groups (see Table 2). The set consists of four premium services with different focus and the classical best-effort service called Standard (STD). The concept is open to add an additional Custom Service if required to meet special requests of future applications. The network view using Traffic Classes is described in section 4.2.

A balanced mixture of NS load on real networks is required as the DiffServ approach does not allow to give absolute guarantees. The first approach defines fixed limits per premium class, which some does not exceed the link capacities. Per NS performance measurement and monitoring ensures the expected quality. Within the second specification phase measurement driven control loops will replace the static limits.

Technically the NSs are coded in the Type of Service (ToS) field of each IP packet, namely in the first 3 bits, which are known as precedence bits. Routers act autonomously for handling packets in a service specific way.

NS	Description	Application Group	Traffic Class	ToS Coding
PCBR	Premium Constant Bit Rate	1	TCL-1	"110 000 xx"
PVBR	Premium Variable Bit Rate	2	TCL-2	"101 000 xx"
PMM	Premium Multi-Media	3	TCL-3	"100 000 xx"
PMC	Premium Mission Critical	4	TCL-4	"010 000 xx"
STD	Standard	5	TCL-STD	"000 000 xx"

Table 3: AQUILA Network Services and their associated Traffic Classes.

Within the first trial period successful co-existence of these classes could be proved [Kope01]. Additional profiling of packets within one NS (in-profile vs. out-of-profile), which was defined in the specification [Wint00], was dropped from the implementation in order to limit the needed router QoS features to gain performance speed-up.

3.2 Characteristics at Network and Traffic Flow Level

The definition of The AQUILA NS focuses on both, network level QoS and traffic flow characteristics. Additional criteria like elasticity and greediness complete the behavior description.

Network level QoS characteristics are the measurable parameters loss, delay and jitter (delay variation). Figure 2 shows this three dimensions and places the defined NSs with their primary and secondary focus in the map. A close neighborhood to the arrowhead indicates a high sensitivity against the criteria. For example the Premium Mission Critical Network Service focuses firstly on low loss and secondly on low delay. The Standard Network Service is placed in the middle, because it does not take care of any of those three parameters.

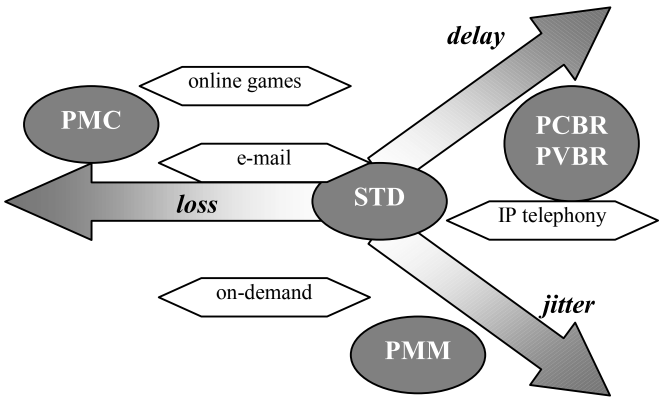

Figure 2: Network level QoS characteristics: loss, delay, jitter (delay variation).

Traffic flow characteristics are mainly the bandwidth, burstiness and the MTU (packet size). Figure 3 shows the placement of the NSs against these characteristics. While for network level characteristics PCBR and PVBR focuses on the same aim here they are on the opposite ends for burstiness.

There is a close dependency between bandwidth and MTU. Therefore the two arrows are in parallel. The requirements against burstiness are completely independent from bandwidth and MTU. Again typical applications with their characteristics, respectively requirements are placed on the map.

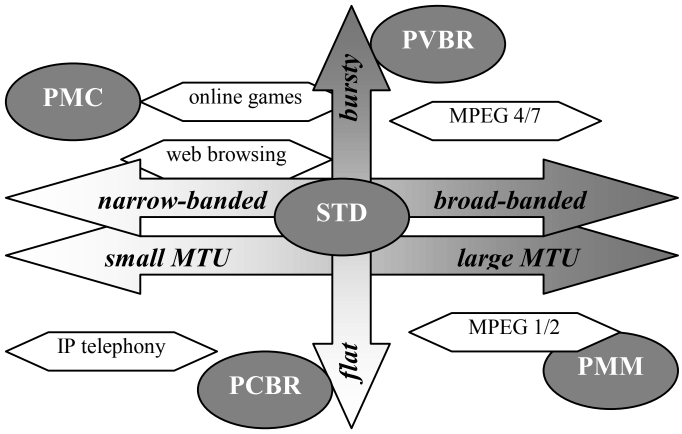

Figure 3: Traffic flow characteristics: bandwidth, burstiness and MTU (packet size).

As a third group of characteristics transport protocol support and elasticity influence the grouping of flows to NSs. The typical IP transport protocols are typically TCP on the one hand side and UDP/RTP on the other side. Figure 4 illustrates the orientation of the defined NSs.

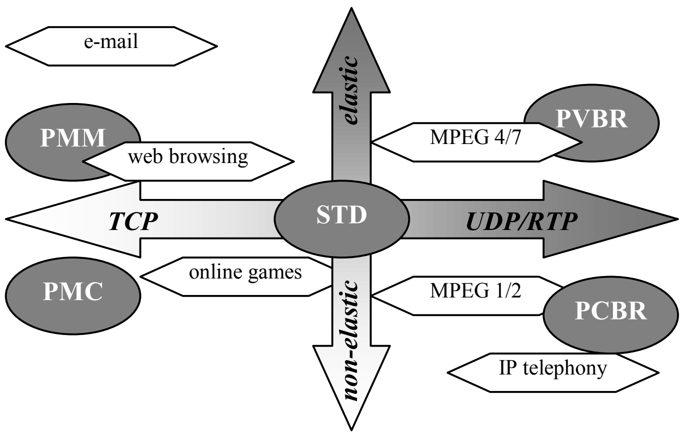

Figure 4: Protocol support and elasticity characteristics.

From administrative point of view additional characteristics like lifetime of flows or availability of NSs on demand are important. The Service Level Specification (see section 4.3) allows to control additional conditions. A detailed description of the NSs including parameters and algorithms for their appliance is to be found in [RiSa00].

3.3 Network Services Landscape

As there are multiple dimensions of characteristics it is difficult to define a limited set of NSs, which exactly meets all requirements. Figure 5 summarizes the application and traffic oriented views concerning their characteristics. The five AQUILA NSs form a well suited intermediate between both.

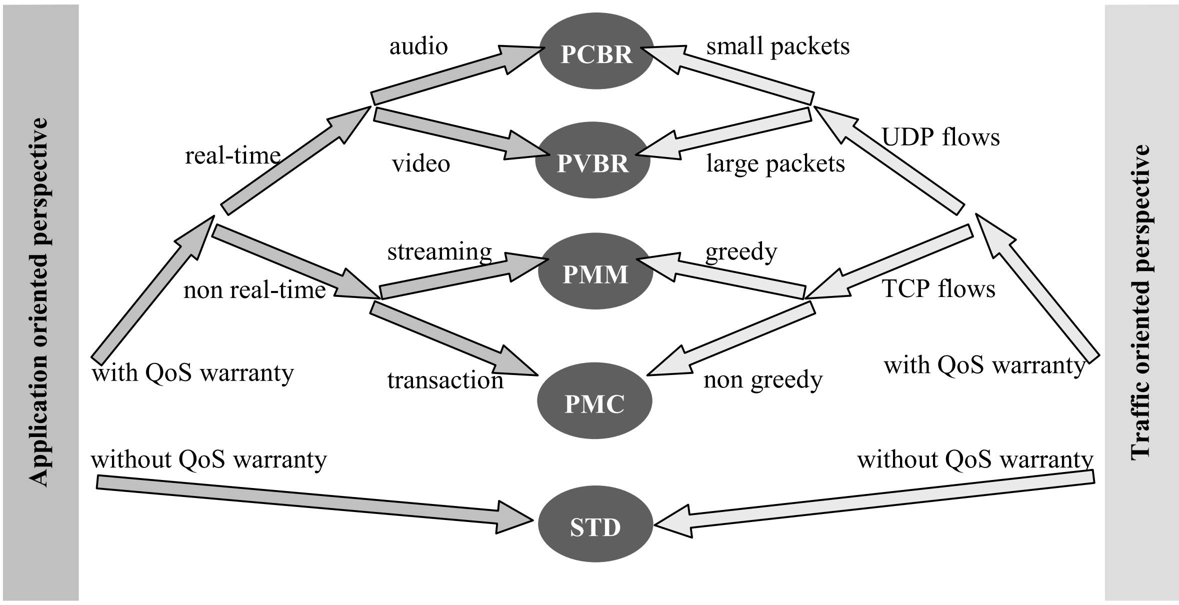

Figure 5: Landscape of Network Services from application and traffic oriented perspective.

4 Embedding of Network Services

The NSs form a central abstraction layer as an intermediate between application and network view. By knowing the applications requirements or on user requests the End-User Application Toolkit is asking for the appropriate NS or a mixture of NSs as upper peer. Technically speaking NSs rely on the underlying Traffic Classes, which are associated with router parameterization for scheduling and queuing.

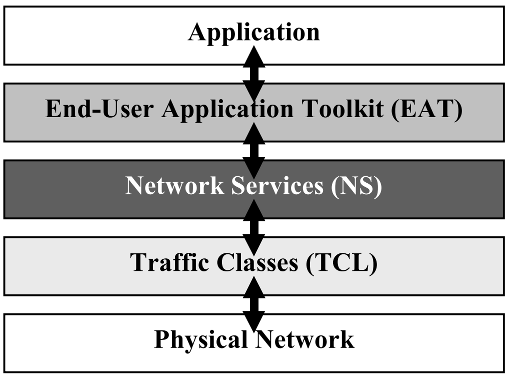

Figure 6: Network Services as intermediate between application and network view.

4.1 The Application View: End-User Application Toolkit

One of the main shortcomings of the DiffServ architecture in comparison to IntServ is the lack of a mechanism for the generation of resource reservation requests by end-users. The End-user Application Toolkit (EAT) aims to fill the gap between the end-user and the network, thus providing true end-to-end QoS support.

The EAT is a middleware architecture for the transparent support of QoS. The main objective of the EAT is the provision of a scalable and efficient approach for transparent QoS support for a wide range of applications including, in a first step, legacy applications. (With the term legacy, we refer to commercial, non QoS-aware applications.)

The EAT takes care of the end-user and reservation management. It provides an authentication mechanism for the authorized use of the network resources, retrieves the available network services from the RCL, presents them into available QoS options in a user friendly way to the customer (for the selection of the appropriate quality parameters for an application), after the end-user's selection it translates and forwards the QoS requests to the RCL. Different mechanisms to request QoS are supported. Proxies are provided for special applications, which are natively not QoS-aware. All accepted requests are mapped to the appropriate network services.

To present in a well-understandable manner the QoS options the EAT co-operates with a so called converter. For the supported legacy applications, the converter maps back and forth the application's information stored in so-called application profiles. The information kept resume on the one hand end-user specific wordings and on the other hand the corresponding technical parameters needed for the mapping into NSs [Fünf01].

4.2 The Network View: Traffic Classes

In the AQUILA approach the set of NSs as a whole represents a kind of QoS interface between the applications and the network. In the above sections the role of NSs was presented from the applications point of view: by choosing the appropriate NS the application – or better the EAT – chooses a combination of service elements, expressed in terms of an SLS (see section 4.3). These combinations should be exhaustive in the range of the application requirements, but necessarily small in order to keep the implementation simple. This section discusses the role of the NSs from the network point of view. The job of providing QoS in a multi-service network embeds the following tasks:

- providing QoS performances assurances,
- providing QoS differentiation,
- providing traffic separation (e.g. between TCP and UDP flows).

At flow level, the first task is accomplished by means of Traffic Admission Control and Traffic Conditioning functions, enforced at the network edge with the appropriate algorithms. The tools for the subsequent tasks are those mechanisms that differentiate the access traffic to the network resources at the packet level (bandwidth, buffer): Scheduling and Queue Management algorithms, enforced at each network interface. Each traffic flow must access to a combination of such algorithms. The more combinations are present in the system, the more complex will be the implementation and the operation of the system itself, which weakens its predict-ability. All that pushes towards the implementation of a reduced set of combinations, called *Traffic Classes (TCL)* in AQUILA. To date the set of TCLs defined in AQUILA map 1:1 the set of NSs (see Table 3), i.e. each NS is associated to a different TCL. This choice was done to allow a simple NS->TCL mapping in the initial phase of the project. Nevertheless, NSs and TCLs should be regarded as different objects, and there is the possibility in the future to introduce new NSs upon the same set of TCLs.

Each AQUILA TCL – numbered from 1 to 5 - is associated a different queue in the router output interface. All the queues except the first one are served by a Weighted Fair Queuing (WFQ) scheduler, thus are associated to a WFQ weight. The queue dedicated to TCL-1 is served with strict priority over the others. Figure 7 shows the complete scheduling scheme at each router output port.

TCL-1 and TCL-2 are intended to support UDP traffic with stringent QoS requirements. In particular TCL-1 will support the PCBR service, characterized by very high QoS performance (very low delay and very low losses), accomplished by a conservative AC scheme based only on the declaration of the peak-rate only. Typically, TCL-1 queues will be entered by flows with small to medium packet size (< 256 B) and moderate peak-rate, as typically originated by real-time streaming applications like voice over IP, etc. TCL-2 will instead support the PVBR service. It will deliver a lower QoS level (low delay and low losses) to those streaming application with high emission rate variability and/or large packets: the AC for TCL-2 scheme allows for some degree of statistical multiplexing, thus the mean rate of the flows must be taken into account in the AC algorithm along with the peak-rate. For both this classes a simple FIFO policy is used in the queues, while the packets exceeding the declared profile are dropped.

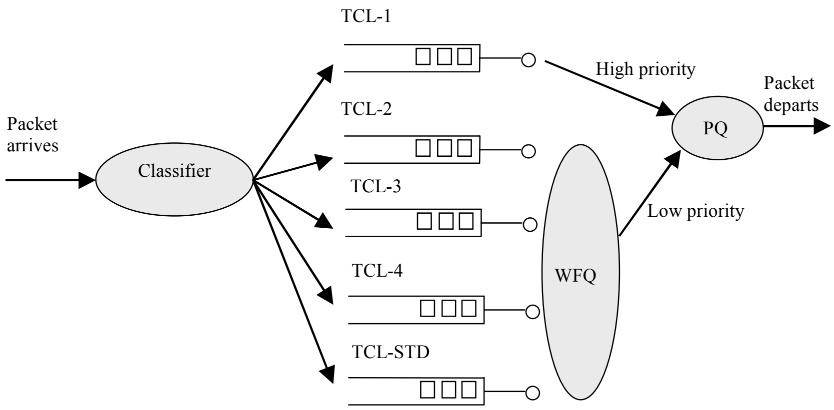

Figure 7: Design of router output port with queues per Traffic Class.

TCL-3 and TCL-4 are dedicated to reactive flows (TCP and TCP-like). In particular, TCL-3 will support PMM service and will take long-lived TCP connections (long file transfers) or other adaptive application flows (audio/video download, adaptive video). These flows are typically greedy, as they continue to expand the emission rate until congestion is reached. TCL-4 instead will support PMC service and will receive non-greedy elastic flows, typically short-lived TCP connections originated by some critical transaction application (e.g. finance) or interactive games. Note that separating long-lived and short-lived TCP connections into different classes, thus avoiding direct competition on the same resources (the WFQ scheduler acts as a sort of arbiter for the bandwidth) prevents the former from starving the latter. The queues for TCL-3 and TCL-4 uses RIO queuing management scheme. Flows in TCL-3 only declare the expected average rate, and excess packets are simply marked as out-of-profile. Flows in TCL-4 declares both average and peak rate. Finally, TCL-5 is dedicated to the STD NS, and is built with a separate FIFO queue associated to a small WFQ weight to give it lower service priority. Further details on the TCL definition can be found in [RiSa00].

4.3 Service Level Specification

The Service Level Specification (SLS) represents the technical description of the object of a Service Level Agreement. The SLS should unambiguously define the scope of the QoS contract: which flows should receive QoS, which are their traffic characteristics, which is the expected QoS level. It has been recognized that a common standardized way to express the semantic of the SLS could be useful in an open multi-provider environment. Some work has been produced in this direction [Tjoe00], [TjRe01] and related discussion within IETF is ongoing. If a “static” QoS provisioning approach is envisaged the agreement is negotiated “off-line” between the user and the provider, may be involving human intervention. A formal SLS can be useful to have a clear and commonly understood picture of the service and of the required QoS. If a dynamic approach is used, where the user application can automatically send QoS requests to the network, the SLS should also be mapped into signaling in-formation exchanged by QoS aware elements. This is the approach followed by the AQUILA project, where the EAT sends its reservation request messages to the RCL, specifying an SLS. The semantic definition of the AQUILA SLS is described in [Sals00]. The three most important components of the reservation requests are: the scope and amount of reservation (where the reservation applies and how much is the bandwidth); the type of requested service (possibly including a set of QoS parameters); the flow identification (i.e. to which IP flow or aggregate of flows the reservation apply). The reservation request messages, as all the AQUILA control messages, are transported using CORBA.

It is important to discuss what is the capability offered by the network operator to the QoS user with this SLS mechanism. We believe that the more generic definition SLS should be standardized, trying to capture all the possible service offerings that can be provided over a DiffServ network. A quite complex set of parameters will be contained in this generic SLS definition. An instance of the generic SLS, containing concrete values for the parameters represents an "external" service that is provided to a customer by the DiffServ network operator. The DiffServ network operator will use the SLS parameters to map the user requirements into internal mechanisms (e.g. DiffServ QoS classes) [CISCO]. The mapping process between the generic SLS and the concrete QoS mechanisms can be very complex if the user can freely select and combine the parameters. Inefficiencies can also be added. Moreover, as we have discussed above, it is commonly believed that only a few QoS classes will be handled in core DiffServ networks.

To solve these problems a drastic approach is indeed to standardize the external services offered to the customer, possibly having a one-to-one mapping with the internal mechanisms in the DiffServ network. The idea of having "Globally Well known Services" is for example used in [TeCh01]. The approach that we have proposed in [Sals00] is more generic. The idea is to have a powerful generic SLS description template and to define external services as "predefined SLS type", in terms of the generic SLS template. A "predefined SLS type" fixes values (or range of values) for a subset of the parameters in the generic SLS. It may also fix some relationships or dependencies between some parameters. If this mechanism is used, the SLS that is invoked by the customer will carry an indication of the "predefined SLS type" and it will contain the actual values for the parameters that are not fixed. The AQUILA NSs can be seen as example of these SLS types, where the user can customize some parameters (for example the amount of requested bandwidth), while other are fixed with the Net-work Service (for example the delay cannot be negotiated in the QoS request).

We think that the set of NSs defined in the AQUILA context can fulfil the requirements of the applications and the needs of the user, both the current one and those that can be envisaged in the near future. It also represents a balance trade-off between complexity and flexibility in service offering for a DiffServ net-work. The definition of standardized SLS should open the door for the evolution of this service model. Additional new NSs will be defined giving user control over a larger set of parameters.

5 Conclusions and Further Work

In this paper the set of NSs defined in the AQUILA project is presented. Such set was designed to support a broad range of applications by clustering their requirements in a reduced number of combinations. We discussed the implementation of the NS concept within the AQUILA architecture, both on the user side (EAT) and on the network side (TCLs).

At the time of writing an implementation of such NSs has been produced and is currently under test in a experimental trial with some real applications. In particular among the trial objectives a special focus is given to the matching between the actual QoS performances and the expected ones. Preliminary available results are conformant. Nevertheless a review of the NSs and of the supporting TCLs will be carried on based on the trial results. For a second trial phase complex Internet Services are taken into consideration, relying on a adequate mix of NSs.

Further work is due to clarify the service levels that can be provided to those users connected to the network through low bandwidth links, which likely impose some restrictions on the range of services that can be provided. Encoding such restrictions is a matter that touches the NS definition and the SLS.

An important aspect of the NS defined till now is the a-priori declaration of few emission profile values, in particular the average rate and the burstiness, in order to perform traffic regulation at the network edge. Since the first trial experiments we realized that in most practical cases it is very problematic to gain the a-priori knowledge of such values at the application side. Moreover, it became clear that even in case of perfect knowledge of such parameters it is difficult at the network side to exploit such an information in a fully effective way. Based on these considerations, a further step will be taken to simplify the traffic declaration for some NS – mainly PVBR and PMC – while moving the traffic regulation mechanisms (mainly Admission Control) towards some measurement based algorithms. This will likely have an impact on the pricing mechanisms. In facts in case the emission profile declaration is loose or even null, there is no way to limit greedy behaviors by means of packet level conditioning (dropping/marking of excess packets). In this case appropriate pricing policies can be used to penalize such misbehaviors. All these issues represents interesting material for further works in the NS for the AQUILA project.

6 Acknowledgements: The AQUILA Project

The AQUILA (Adaptive Resource Control for QoS using an IP-based Layered Architecture) project [AQUIILA] is an European research project partially funded by the IST programme [IST], placed in the "Next Generation Networks" cluster. Well-structured, three workpackage groups, each consisting of several work-packages, bundle the activities of requirement analysis, specification and simulation, design and implementation as well as integration, trial and exploitation following a defined milestone plan. Close co-operation with the IST projects CADENUS [CADENUS] and TEQUILA [TEQUILA], which cover other QoS aspects, is performed.

The AQUILA project supports an interdisciplinary research and implementation teamwork testing modern software engineering methods. The partners are: Siemens (D), Q-Systems (GR) as manufacturers; Telekom Austria (A), Telekomunikacja Polska (PL), Elisa Communications (FIN), T-Nova Deutsche Telekom (D) from the ISP and operator side; Bertelsmann mediaSystems (D) as content provider, and from the universities and research institutes: Dresden University of Technology (D), National Technical University of Athens (GR), Salzburg Research (A), Warsaw University of Technology (PL) and CoRiTeL (I).

The AQUILA project lasts until end of 2002. It will have a second trial phase following the spiral approach to verify the RCL results in 2002.

7 Paper specific Abbreviations

AQUILA	Adaptive resource control for QoS Using an IP-based Layered Architecture
CADENUS	Creation And Deployment of ENd-User Services in premium IP networks
EAT	End-User Application Toolkit
NS	Network Service
PCBR	Premium Constant Bit Rate
PMC	Premium Mission Critical
PMM	Premium Multi-Media
PVBR	Premium Variable Bit Rate
RCL	Resource Control Layer
SLS	Service Level Specification
TCL	Traffic CLass
TEQUILA	Traffic Engineering for QUality of service in the Internet, at LArge scale

8 List of References

[AQUILA] AQUILA consortium. *Adaptive Resource Control for QoS Using an IP-based Layered Architecture.* AQUILA homepage, URL: http://www-st.inf.tu-dresden.de/aquila/, IST project 2000-2002.

[CADENUS] CADENUS consortium. *Creation and Deployment of End-User Services in Premium IP Networks.* CADENUS homepage, URL: http://www.cadenus.org/.

[CISCO] CISCO White Paper. *Technical Benefits of Using a Service Provider with a Cisco Powered Network.* URL: http://www.tradespeak.com/htmldocs/1245.html .

[Fünf01] F. Fünfstück et al. *A QoS Middleware between Users, Applications and the Network.* ComCon 8, June 2001, Crete.

[GVU98] GVU's WWW Surveying Team. *GVU's 10th WWW User Survey.* Graphics, Visualization & Usability (GVU) Center at Georgia Tech, Atlanta 1998 http://www.cc.gatech.edu/gvu/user_surveys/survey-1998-10/graphs/use/q38.htm.

[Gode00] D. Goderis et al. *Service Level Specification Semantics and Parameters.* URL: http://www.ietf.org/internet-drafts/draft-tequila-sls-00.txt.

[Hust00] G. Huston; *Next Steps for the IP QoS Architecture*. IETF RFC2990, Nov. 2000

[IST] European Commission. *Information Society Technologies Programme (IST).* IST homepage, URL: http://www.cordis.lu/ist/.

[Kope01] Z. Kopertowski et al. *First trial report.* AQUILA deliverable D3201, Warsaw, June 2001.

[Krau00] R. Kraut et al. *Internet Paradox; A Social Technology that reduces Social Involvement and Psychological Well-Being?* American Psychologist, December 2000, Vol. 55, No. 12. http://www.apa.org/journals/amp/amp5391017.html.

[RiSa00] F. Ricciato, S. Salsano et al. *Specification of traffic handling for the first trial.* AQUILA deliverable D1301, Rome, June 2000.

[Sals00] S. Salsano et al. IETF draft on *Definition and usage of SLSs in the AQUILA consortium*. URL: http://www.ietf.org/internet-drafts/draft-salsano-aquila-sls-00.txt, November 2000.

[TEQUILA] TEQUILA consortium. *Traffic Engineering for Quality of Service in the Internet, at Large Scale.* TEQUILA homepage, URL: http://www.ist-tequila.org/.

[TjRa01] Y. T'Joens, R.Rajan. *Service Level Specification and Usage (SLSU) WG.* Proposed Chapter for SLSU WG (version 1) February 2001; URL: http://www.ist-tequila.org/slsuwgv1.txt .

[Tjoe00] Y. T'Joens et al. *Service Level Specification and Usage Framework.* IETF draft: draft-manyfolks-sls-framework-00.txt , October 2000;

[TeCh01] B. Teitelbaum, P. Chimento. *QBone Bandwidth Broker Architecture.* Work in Progress, June 2001; URL: http://qbone.internet2.edu/bb/bboutline2.html .

[Wint00] M. Winter et al. *System architecture and specification for the first trial.* AQUILA deliverable D1201, Munich, June 2000.