ELISA: European Linkage between Internet Integrated and Differentiated Services over ATM

T. V. Do (α), G. Eichler (β), H. Hussmann (γ), B. Koch (δ), G.Mamais (ε), C. Prehofer (δ), S. Salsano (ζ), J.-J. Tchouto, C. Tittel, P. Todorova (η), I.S. Venieris (ε)

(α) Technical University of Budapest
(β) *Deutsche Telekom AG*
(γ)Dresden University of Technology, Department of Computer Science
(δ) Siemens AG
(ε) National Technical University of Athens
(ζ) Consorzio di Ricerca sulle Telecomunicazioni (CoRiTeL)
(η) GMD-Fokus

Abstract

Nowadays we are witnessing the growth of Internet at an unprecedented rate. However, the current Internet does not guarantee Quality of Service for user flows, which prevents the provision of good quality for multimedia streaming applications over long distance. This led to the development of the Integrated Services architecture and recently to the Differentiated Services architectures. On the other hand, ATM networks are being installed for the public infrastructure of B-ISDN throughout the world with advanced QoS capabilities to offer various service classes to customers.

Considering these trends it seems important to network operators to integrate Internet protocols and the mature ATM QoS technology. This paper presents a specific approach for such an integration which is taken in the international project ELISA (European Experiment on the Linkage between Internet Integrated Services and ATM), funded by the EU through the ACTS programme.

The approach is based on a combination of the Integrated Services architecture and the Differentiated Services architecture. Moreover, the approach serves as a linkage between those architectures and the ATM technology. A trial will demonstrate the feasibility of the proposed approach.

1. Introduction

The Internet protocol (IP) family is rapidly becoming the network layer technology of choice for building packet networks and it is enabling the growth of the worldwide Internet. One explanation is that there are more and more IP based applications, which fulfil user needs. However, the Internet itself does not guarantee any Quality of Service for user flows, which prevents the provision of good quality for multimedia applications such as video and voice over wide area. This led to the development of the IETF Integrated Service (IntServ) architecture (based on the use of the RSVP protocol) and recently the Differentiated Service (DiffServ) architecture to provide complementary approaches to the problem of providing QoS for Internet applications.

On the other hand, ATM networks are being installed, mainly as backbone networks. As the deployment of current IP-over-ATM technologies (such as LAN Emulation and MPOA) continues and new technologies for operating IP-over-ATM emerge, the rationale behind IP-over-ATM will increase more dramatically. Considering these trends, network operators can find important revenue sources by means of the integration of Internet protocols and the mature ATM QoS technology.

The main objective of the ELISA project is to offer a practical approach for the provisioning of Integrated Services (with the help of RSVP) and Differentiated Services architectures based on Internet IP technology. Moreover, the ELISA project brings together such an advanced IP service approach with the capabilities offered by an ATM-core network (e.g. guaranteed QoS through ATM switched connections).

In order to achieve that goal the ELISA architecture focuses on an Edge Device which acts as an edge router as well as an Integrated Services gateway. The AC310 ELISA consortium is formed by public network operators (Deutsche Telekom AG, Telefonica I+D), equipment manufacturers (Siemens AG, Siemens Telefongyár Kft, Topologix GmbH), as well as research institutes and universities (GMD-Fokus, Coritel, Dresden University of

Technology, National Technical University of Athens, Technical University of Budapest). It will provide a transnational testbed to demonstrate the feasibility of the architectural approach.

The paper is organised as follows. Section 2 describes ELISA reference configuration and gives details about protocols and interfaces to be considered. Section 3 provides further details on the elements of the architecture such as the Edge Device and the end-user terminals. The applications and scenarios which will demonstrate the feasibility of the proposed solution are discussed in Section 4. Section 5 describes some business scenarios where the Edge Device can be utilised. Section 6 provides an overview on the trial and demonstration plan of the ELISA project.

2. ELISA architecture

One of the main goals of the ELISA architecture is to enable operators of wide-area ATM networks to offer the advanced features of ATM technology to end-users which are using IP-based applications. The architecture assumes as the standard case that end-users are not connected directly to the ATM wide-area network but are connected through the infrastructure of an Internet Service Provider or an intranet, which may use, for instance, various LAN technologies, ISDN dial-up access or xDSL access networks. The goal of the proposed architecture is to make QoS available to such end-users in a way which is based on the features of the ATM core network, but which does not force the end-users to leave the world of Internet protocols. The ELISA architecture is flexible enough to be applied to configurations where the core network is not ATM based, but this is out of focus of the prototype currently being developed.

In order to bring QoS to Internet-based applications, it is necessary that QoS is somehow accessible in the Internet world, i.e. that applications or end systems can define their specific QoS requirements. There are two main approaches to QoS which are currently discussed in the Internet community:

- The *Integrated Services* (IntServ) architecture [1] is based on individual resource reservations issued by the applications using reservation protocols, in particular the Resource Reservation Protocol (RSVP) [2]. The routers in the core network have to reserve resources for the individual flows.
- The *Differentiated Services* (DiffServ) architecture [3, 4] takes a much simpler approach and simply assumes that IP packets are marked as belonging to one out of a number of different traffic classes. Core routers simply have to follow the general rules for the respective traffic class and do not deal with individual QoS requests.

Both approaches have their advantages and disadvantages. The main problem with the Integrated Services architecture is that it does not scale well to large wide-area networks, since in this case reservation processing on central core routers becomes a bottleneck. The problem with the Differentiated Services approach is that it can give true service guarantees only if it is combined with rigorous policing of the access to the high-quality service classes. And a problem with both approaches is how to interwork with already available ATM infrastructure. ELISA tries to solve these weak points.

Several aspects related to the interworking of IntServ and DiffServ approaches are considered in [5]. The definition of the ELISA architecture has progressed in parallel with the work reported in this IETF draft. The ELISA architecture focuses on one important network element, which is particularly relevant for QoS. This network element is called *Edge Device* (ED) within ELISA, but it may also be called an extended access router or a QoS gateway. The Edge Device prototype developed in ELISA connects an access network to a wide-area core network. The interconnection of EDs realises an overlay QoS network for the IP-based applications. ELISA targets only end-systems that are connected through an Edge Device. The access network is typically based on LAN technologies or on POTS, ISDN or xDSL. The core network can be a plain ATM network, which is the optimal network base for ELISA. But the ELISA Edge Device also works in more general configurations where the core network is an arbitrary DiffServ-based IP network, which possibly resides upon an ATM layer. The ELISA network architecture does not require any changes or upgrades to any of the involved sub-networks.

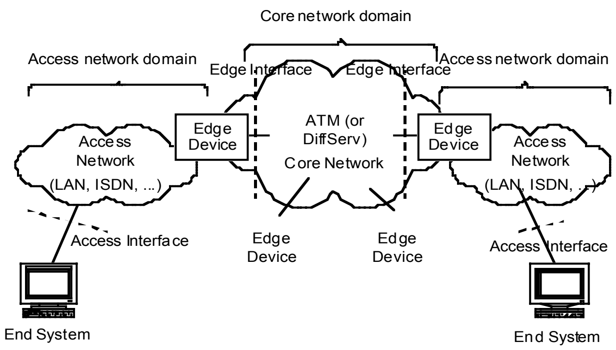

Figure 1: ELISA reference architecture

Figure 1 shows the reference network architecture of ELISA. Details are shown only for a pair of end-systems, but the general assumption is to have many end-systems connected to each Edge Device and many Edge Devices to the core network.

The end-systems connected to the access network use only IP-based protocols and can specify their QoS requirements in two different ways: either they issue reservation requests using IntServ approach or their traffic belongs to a specific DiffServ traffic class. The Edge Device deals with both kinds of requirements.

Reservations are analyzed, policed and mapped onto appropriate core network traffic classes (Figure 2). This core network traffic classes can be carried over a dedicated ATM Switched Virtual Channel (SVC) to another Edge Device, or over a portion of a pre-established ATM Permanent Virtual Channel (PVC) using a specific DiffServ class. Note that the end-system just uses the reservation protocol without any knowledge of the mechanism chosen by the Edge Device. Note further that the reservation processing is done only in the Edge Device and not in the core network. The flexible mapping into ATM VCs distinguishes the Edge Device from other proposals for the IntServ/DiffServ integration known to the authors.

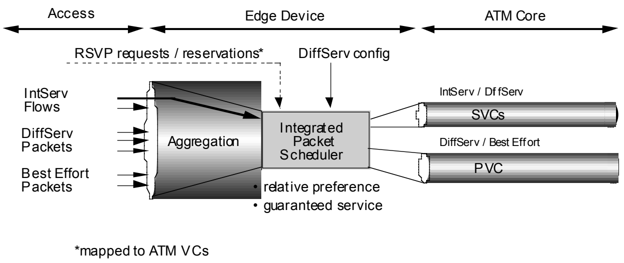

Figure 2: ELISA Edge Device – an approach for integrating IntServ/DiffServ over ATM

DiffServ traffic is analyzed and policed in the Edge Device. The Edge Device is able to carry out DiffServ marking for end-systems that do not themselves support DiffServ marking. In this case, the traffic class is derived from the IP address or other selectors (e.g. TCP port number) based on information, that is kept on the Edge Device in a configurable way. The ED is able to do policing based on volume limits for traffic in a specific DiffServ class, and therefore enables QoS guarantees for reservations following the DiffServ approach.

As an additional function, the ED is able to collect information for usage-based charging for QoS usage, for explicit reservations as well as for DiffServ class usage.

The advantages of the ELISA architecture are that it flexibly combines the Integrated Services and Differentiated Service model by mapping reservation requests onto DiffServ classes where possible and that it removes some of the scalability problems for explicit QoS reservations in wide-area networks. Of course, this raises the question how the scalability of the ELISA architecture itself is ensured. Aspects related to the scalability of an architecture combining IntServ and DiffServ are considered in [6].

In the ELISA architecture, two different kinds of interfaces must be considered. Figure 1 shows these two interfaces, which will be referred to as *Access Interface* and *Edge Interface*. The Access Interface allows the host to interact with the network and in particular to be connected to the Edge Device. The Edge Interface allows the Edge Device to access the core network. This distinction is consistent with the choice to partition the network in an access domain and a core domain. The IP technology is the common denominator between the two domains and is a basic component of both interfaces.

Let us briefly analyse the protocol stacks used at both interfaces to transport IP packets (user plane) and the control mechanism needed to signal QoS requirements (control plane). On the Access Interface, the IP packets are carried over LAN (Ethernet) or ISDN (using PPP). At the control plane level, the RSVP protocol is the mechanism which is used by the applications to signal their QoS requirements on the Access Interface, according to the Internet Integrated Service model. The ELISA RSVP protocol messages are fully compliant with the relevant RFCs issued by the IETF [2]. In the ELISA architecture, the RSVP can be used within the access network domain to allocate resources up to the Edge Device. RSVP is used in the ED to allocate resources and to map the user requests onto the mechanisms used in the core network.

The Edge Interface is based on ATM, under the hypothesis that a network operator provides ATM semi-permanent and switched connection. It is a UNI interface; for the ELISA project a common subset of ATM Forum 3.1 and ITU-T 2931 is used. IP packets are encapsulated on ATM according to RFC1483. Within the core network domain, the IP layer is enhanced with Differentiated-Services concepts, allowing different classes of packets to receive a specific handling. The different classes of packets are discriminated against a special byte in the IP header. The Differential Service classes are selected by the set of classes under standardisation in the IETF Differentiated Service WG [7]. In particular the Expedited Forwarding (EF) class, will be used to support applications with real-time QoS requirements.

The control plane of the Edge Interface is based on two concepts. The first concept is that the RSVP is used with "end-to-end" significance through the core network domain. This means that RSVP messages are interpreted only in the access network domain and in the Edge Device, but not in the core network domain. The second concept is that the RSVP reservations can be mapped either onto DiffServ classes or directly onto a dedicated ATM connection that will be set-up on demand.

In the scenarios described so far, the user application is always supposed to signal its QoS requirements with RSVP. A possible extension is to allow the application to directly mark the IP packets with a given DiffServ class, and leave to the ED the task to police the user behaviour.

3. Elements of the architecture

3.1 Edge Device

The previous section imposes a set of requirements that must be fulfilled by the ELISA Edge Device. At first, the ELISA ED has to be capable to support QoS guarantees to selected IP flows handled by it. Therefore, functions such as flow admission control, policing, shaping, and bandwidth management must be provided. In the second place, the ELISA ED is also the gateway that interconnects users attached to the low cost access network with the core ATM network. Consequently, bridging functions for both user plane and control plane are accommodated, too. Moreover, the efficient utilisation of the ATM network resources is a matter of significant importance. Many of the existing IP based applications generate many short living connections. For instance, a WWW browser application may establish many TCP connections with multiple different servers within a small time period. Each connection is used for the transfer of text or an image contained within an HTML page and then the TCP connection is released. If the ED allocates resources from the ATM network every time that a request is generated in the access network, this would result in a lot of signalling traffic within the ATM network and low utilisation. Therefore, the ELISA ED performs bandwidth management algorithms that decide whether a new request will be fulfilled by opening a new SVC, dedicated to this request, or not. Requests that are not handled by opening a new dedicated SVC are classified and then aggregated with other traffic according to the Differentiated Services model. Hence, the ED must provide functions for classification and marking of the IP packets. Finally, since the ELISA ED offers multiple services with different QoS parameters, an appropriate mechanism for recording the utilisation of the network per customer must be provided, too.

The hardware platform of the ELISA ED is a SUN Sparc workstation together with SUN ATM-adapters. Solaris 2.5.1 or 2.6 will be used as operating system. All functional blocks will be implemented in C/C++ and a CORBA based communication platform will be used for the inter-edge-device communication.

3.2 Network services

The following network service capabilities will be implemented in the Edge Devices:

- IntServ, where resources are explicitly requested. Resources will be provided by a dedicated SVC or aggregation into an existing link.
- DiffServ, where resources are provided according to the marking of the packets. The following classes are available: Best Effort (BE), Priority (P), High Priority (HP), Expedited Forwarding (EF).

Therefore, the network operators can use those capabilities to support PNO (Public Network Operator) services, as detailed in Section 5 on business scenarios.

3.3 Terminals

One of the intentions of the ELISA project is to provide the advantages of ATM networks for the users with low cost multimedia terminals connected to the Edge Device through various access technologies such as LANs, ISDN, xDSL. Most of today's PCs and workstations used in the cooperate environment, education and research institutes and at home are equipped with audio and video capabilities. To support audio features sampling rates of at least CD quality and a full duplex operating mode are required, while to support video frame capturing and rendering of PAL/NTSC quality is expected. Moreover, hardware support of video encoding and decoding is recommended to achieve good quality video. The terminal equipment with various operating systems (Unix, Linux, Windows,…) is able to run multimedia resource aware (RSVP) IP applications.

In the ELISA project two platforms have been selected to represent various terminal options. The first platform is a Sun Sparc Ultra workstation which will be used as a high-end terminal for business users. The second platform is a PC running Linux operating system. Linux is a low cost, powerful and stable operating system, with growing acceptance. Especially developers make use of the availability of the source code, the documentation and the rapidly increasing capabilities. The same argument is true for resource reservation features required by ELISA applications.

Note that these two platforms are only representatives, other kind of terminals and workstations with other operation systems (e.g: Windows from Microsoft enhanced with RSVP capability) can also access the ELISA Edge Device, which proves the versatility of the ELISA approach.

4. Applications and scenarios

4.1 Multimedia service scenarios

The ELISA project will provide a wide range of capabilities which will be offered to users in form of services:

- Multimedia conferencing services,
- File retrieval services and
- WWW navigation services.

Selected applications represent typical use of today's Internet and also allow demonstration of the network services delivered by the ELISA architecture.

- **Multimedia conferencing services**

Multimedia conferencing services provide resources to enable two or more parties to communicate with each other. The quality of the conferencing services will be defined by the packet loss probability, the end to end delay and delay variation. For the trial, three services have been selected to represent each service group:

1. Multimedia conferencing: This service provides its users with the full capabilities of a video conference but no guarantees are given for the audio and video quality.
2. Video telephony: Video telephony provides the resource to connect two parties with ISDN like links.

Only two users may participate, but the ISDN quality (64 Kbits/s) is guaranteed for audio/video transfer.

3. Premium multimedia conferencing: The premium service within the group of conferencing services provides the highest flexibility and quality. It allows to connect two or more parties and guarantees the desired audio/video quality depending on the used coding available to all participants.

- **File retrieval services**

File retrieval services provide the user with the possibility to get data from a certain destination. After disconnecting, the user is able to process the data offline. The quality of the file retrieval services depends on the time to download error-free remote data, i.e. on the throughput and the packet loss. Two services have been selected to be demonstrated:

1. File transfer: The file transfer service offers the transmission of remote data to a local host without guaranteed quality.
2. On demand retrieval: On demand retrieval provides the same functionality as the file transfer, but allows to reserve a certain bandwidth for the duration of the connection. The bandwidth of the service may be selected by Internet service provider (ISP) or by the user, in the later case the ISP must define an allowed range and consider the selected bandwidth when charging the user.

- **WWW navigation services**

Navigation services allow the on-line access to remote Internet data. The quality of the services depends on the end to end delay and variation, the throughput, the packet loss probability, as any form of data might be accessed during the use of the service, i.e. live audio or video, text documents, pictures etc. Two services have been selected:

1. WWW navigation: This service allows on line WWW navigation without providing any guaranteed quality.
2. Premium WWW navigation. The premium service offers the same functionality as the WWW navigation service, but provides and guarantees network resources on request. As an example, "normal" navigation receives the same quality as above, but resources for audio/video data are reserved, if a movie is requested.

4.2 Mapping of end services

The end services described in the previous section have been selected to demonstrate the projects network services. The mapping of the end services to the network services is summarized in Table 1.

Table 1: Service mapping

End service	*Service components*	*Service profile*	*ELISA mapping*
Multimedia conferencing	Video Audio	Regular IP video conference.	BE, P, HP
Video telephony		2 party ISDN like video conference	EF, Dedicated SVC
Premium Multimedia Conferencing		Regular video conference with QoS	EF, Dedicated SVC
File Transfer	Data	Regular FTP	BE
On demand retrieval		FTP / Guaranteed bandwidth	EF, Dedicated SVC
WWW navigation	Video, Audio Data	Regular WWW	BE
Premium WWW navigation		Regular WWW with QoS	P, HP

The low quality services of each group (i.e., multimedia conferencing, File Transfer and WWW navigation) are mapped to a best effort DiffServ class.

End services, which request QoS, like video telephony, premium multimedia conferencing, on demand retrieval, and premium www navigation, are mapped to a guaranteed flow, either dedicated or aggregated or will be assigned to one the DiffServ classes exploited by the project.

4.3 Internal architecture of end-user applications

The software architecture for RSVP aware applications consists of up to three modules as shown in Figure 3.

- The applications themselves, which implement the core functionality;
- The resource reservation modules containing the resource reservation daemons and the resource reservation manager;
- The graphical user interface (GUI).

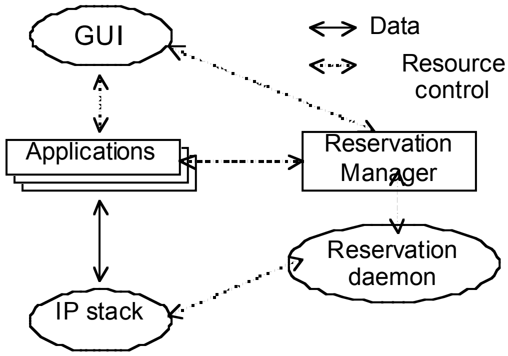

Figure 3: Terminal architecture

The application core module may make use of the reservation manager by directly accessing it or may be controlled by a graphical user interface, which communicates with the reservation module.

The resource reservation module for the terminal prototype is based on the RSVP implementation from the USC Information Science Institute [8]. A resource reservation manager is under development on top of the RSVP daemon. The overall software architecture of the resource reservation module is shown in Figure 4.

This manager is designed to provide a simple interface to RSVP for unidirectional or symmetric bi-directional connections. The interfaces are event-driven and include a library interface and a text based interface.

Moreover, in an environment where the user is required to pay for the reserved network resource, the RSVP manager will also support the option of automatically closing QoS sessions in case of an application software failure.

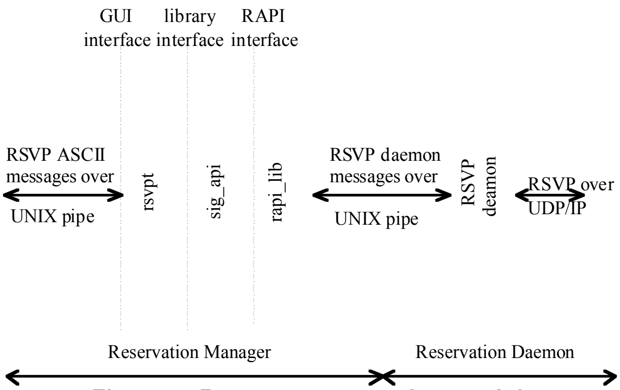

Figure 4: Resource reservation module

The applications to implement the end services include RSVP aware conferencing, integrated file transfer and integrated web navigation.

An RSVP aware conferencing application is designed to operate within a multimedia toolkit [9]. The MASH toolkit is a scaleable multimedia architecture for distributed multimedia collaboration in heterogeneous environments and was developed by the research group at the University of California, Berkeley. It allows the specification of programs using core objects derived from various MBONE implementations [10,11]. The conferencing application:

- implements audio and video functionality;
- simplifies the usage by providing an unified user interface, and
- enables resource reservation by using the reservation module.

The prototype implementation uses unicast connections to build a fully meshed network between the participants to allow different resources allocation between each pair of participants. After the start-up of the conferencing software, the user can specify a list of destinations. The conferencing connections will be established for each of the destinations. If a service class is assigned to a certain destination, the conference connection data are copied to the reservation manager and the reservation manager informs the application after the resource request about the success of the operation. Best-effort traffic is used, if no service class is assigned to a destination or the resources were denied.

5. Business scenarios

In this section two possible business scenarios based on the network architecture developed in ELISA are discussed. Depending on who owns which part of the network, different revenue models are considered. In addition, each player can optimize the network usage in order to optimize his part of the network structure and to optimize revenues. There are two main scenarios, the public access and the corporate access scenario, which are discussed below.

The public access scenario is shown in Figure 5. Users connect to Edge Devices owned by ISPs, e.g. by ISDN, xDSL or any other access technology. The ISPs connect their Edge Devices via an ATM network. This ATM network is owned by an ATM core network provider. Clearly, a possible special case is that both ISP and ATM operator coincide.

In this scenario, we have at least two business relations. The end-user will be charged by the ISP for the individual user services or by flat rate models, which may include access to some services. In addition, it is possible that the service is initiated and paid by a content provider, who in turn charges the user, e.g. by credit card. Another option is that the ISP charges for application level services, which provide QoS.

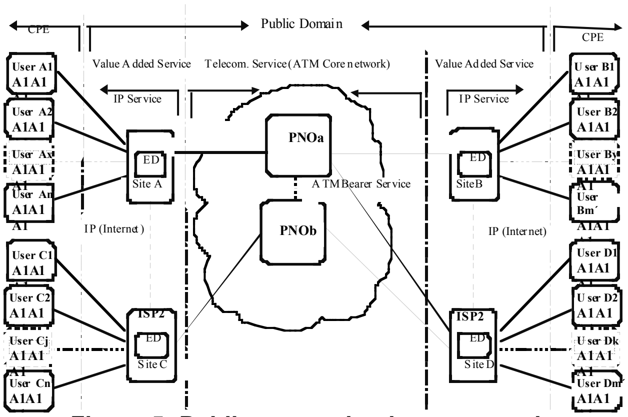

Figure 5: Public access business scenario for the ELISA architecture

In turn, the ISP has to lease the ATM lines from the ATM operator. In the case of dedicated ATM connections for one user QoS request, there is a direct mapping of the cost for the ATM connection to the user service. In the other cases, traffic is aggregated in service classes. The ISP has to lease a number of statically provisioned ATM connections (fixed costs) which are supported by on-demand switched VCs (variable costs). Although one user request may trigger an extra VC, there is no direct relation between the ATM connection and the user service. An option for small networks is to use only switched VCs, which can be setup by means of a management tool.

From a technical point of view, bandwidth management can take place on the IP-level (DiffServ scheduling in the ISP Edge Device) and on the ATM layer via ATM VCs. With the former, DiffServ classes can borrow unused bandwidth of other DiffServ classes in order to utilize the available link layer bandwidth. On the other hand, we can map DiffServ classes to different ATM VCs providing different traffic capabilities. In this case, the ATM layer has to assure that bandwidth is not wasted. However, in the first case, the ISP can use extra bandwidth for best effort and other traffic, while in the other case, the ATM operator may use the extra bandwidth. Hence the ELISA Edge Device design provides flexible mechanisms to utilize bandwidth (and in turn service charges) for both ISP and ATM operator.

In the corporate scenario in Figure 6, the Edge device is operated by a cooperation inside its local area network. In this case, the users are typically connected via LAN infrastructure. The CPE equipment and the Edge Devices are administered by the same corporate network operator. (In case the Edge Device is leased from the another operator, the model is similar to the above one.) For the core network, it is convenient for the cooperation to lease ATM connections which serve as a virtual private network (VPN). In this case, other issues like security and service availability can play a more important role than in the above mentioned case.

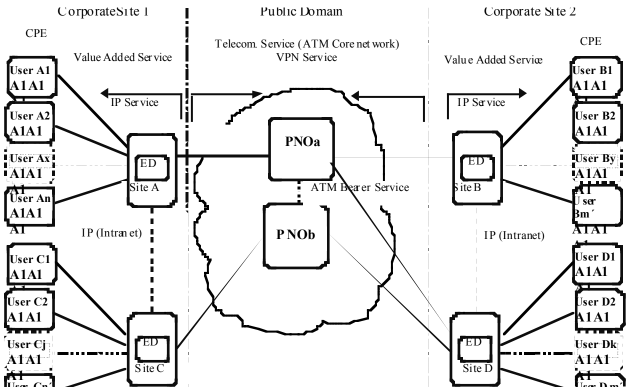

Figure 6: Corporate business scenario for the ELISA architecture

5. Trial

The final target of the project is to evaluate the DiffServ and IntServ features mapped onto ATM SVC connections and also the charging capabilities of the ELISA platform related with the provision of services. For these purpose, an Edge Device and Customer Premises Equipments are being implemented. The feasibility of the ELISA approach will be tested and demonstrated in an international testbed. The testbed will involve Edge Devices located in four core sites (Munich, Darmstadt, Berlin, Budapest) and end systems deployed to several remote sites in different European countries (e.g. Germany, Spain, Italy, Greece). The remote sites and local termials will be connected to EDs using selected access technologies.

The experiments focus on:

- the functional verification of interworking functions between IntServ, DiffServ and ATM, signaling, admission control, traffic policing and traffic shaping as well as charging modules,
- traffic related performance measures affecting Quality of Service,
- experiments assessing end service aspects and taking into account both, the user perception and the service provider perception.

6. Conclusion

The ELISA project provides a practical approach for the provision of Internet Integrated and Differentiated Services using the key advantages of ATM networks (Quality of Services guarantees). The proposed approach supports end systems which are not directly connected to an ATM network but via other technologies to get the advantage of the ATM technology.

For this purpose, the prototype of the ELISA Edge Device and a set of applications are being developed and will be tested to several trial sites in different European countries where they will be connected to typical access technologies for business and residential users. Each access network will be connected to the public ATM core network via Edge Devices acting as gateways. Within the testbed the decisions taken by the project concerning the IntServ/Diffserv mapping and the feasibility of the architectural approach will be evaluated.

Acknowledgements

This work was supported in part by the European Commission under the AC310 ELISA project. The authors are solely responsible for the content of this paper.

References

[1] P.White, “RSVP and Integrated Services in the Internet: A Tutorial”, IEEE Communications Magazine, May 1997

[2] Braden et al., “Resource Reservation Protocol (RSVP)” - IETF RFC 2205, September 1997.

[3] Blake et al., “An Architecture for Differentiated Services”, IETF Draft, October1998.

[4] Blake et al., “A Framework for Differentiated Services”, IETF Draft October 1998.

[5] Bernet et al., “A Framework for Use of RSVP with DiffServ Networks”, IETF Draft, June 1998.

[6] A. Detti et al. “Supporting RSVP in a Differentiated

Service Domain: an Architectural Framework and a Scalability Analysis", ICC'99, Vancouver, Canada.
[7] http://www.ietf.org/ html.charters/diffserv-charter.html
[8] Braden, R., Hoffman, D., "RSVP RAPI Version 5", draft-ietf-rsvp-rapi-00.ps, June 1997
[9] http://www-mash.cs.berkeley.edu/mash
[10] Vat - http://www-nrg.ee.lbl.gov/vat
[11] Vic - http://www-nrg.ee.lbl.gov/vic