

QoS and Multiprotocol Label Switching Experiments for the Design of an ATM-based National Network

Tiziana FERRARI (ferrari@cnaif.infn.it)^{1,2}

Antonia GHISELLI (ghiselli@cnaif.infn.it)¹

Cristina VISTOLI (vistoli@cnaif.infn.it)¹

¹Italian National Institute for Nuclear Physics, National Networking Center (CNAIF)

²Dep. Elettronica Informatica e Sistemistica (DEIS), University of Bologna

Abstract

During the last few years evolution in networking technologies has been a continuous fast process driven by the development of new applications. Nevertheless, the fast increase in the number of users, the huge cost of high speed links and the growth of commercial networks are preventing the use of these advanced applications in a production environment. The traditional best-effort approach is not enough to guarantee acceptable performance and a fair sharing of resources: more complex services and efficient data forwarding are needed. ATM, RSVP and MPLS, some of the approaches devised by the scientific community to address these requirements, are presented. Advantages, performance results and use of these solutions are discussed and compared. The paper focuses on the design of a high speed ATM-based national network and shows how ATM, MPLS and RSVP can be used and integrated in the same infrastructure. ATM, deployed as high speed technology for the backbone, can be integrated with MPLS to enhance data forwarding and achieve simplicity and scalability. On the other hand, RSVP can support per flow quality of service to a limited set of users and/or applications when end-to-end dedicated ATM connections cannot be provided.

Keywords: Internet, Quality of Service, Classes of Service, ATM, Multiprotocol Label Switching, Tag Switching, performance, GARR-B.

Table of Contents

1. [Introduction](#)
2. [QoS and CoS](#)
3. [ATM](#)
4. [RSVP](#)
5. [Multiprotocol label switching](#)
6. [Architecture design for an advanced backbone network: GARR-B](#)
7. [Conclusions](#)
8. [Acknowledgments](#)
9. [Bibliography](#)

1 Introduction

During the last few years evolution in networking technologies has been a continuous fast process driven by the development of new applications. In their turn more advanced and demanding applications, like videoconferencing, computer supported collaborative work, video on demand, tele-teaching, distributed computing etc., have been devised thanks to the availability of enhanced network features. The result of this is a circular process which has recently reached a crisis.

The fast increase in the number of users, the huge cost of high speed links, the growth of commercial networks and in some countries the lack of networking infrastructures dedicated to the research community, are preventing the use of these advanced applications in a production environment.

In almost every country national and international network infrastructures, especially the backbones, have increased their link capacity. TEN-34 [1] is an example in Europe: A new European backbone for interconnection of National Research Networks has been designed and implemented to provide the European research community with interconnection link capacities from 2 Mbps to 34 Mbps. The project for a further increase to 155 Mbps is under development.

Nevertheless, even more complex services are needed. The traditional best-effort approach, in which packets are stored and forwarded without guaranteed services, is not enough to provide acceptable performance and a fair sharing of resources. Thus, an increase in bandwidth is simply not enough to provide any kind of application with the needed performance. New protocols,

architectures and networking technologies need to be developed by the scientific community for quality of service provisioning.

This article supplies a short overview of some of the approaches devised to cope with the problem of quality of service, and of enhanced packet forwarding techniques. Section 2 states the problem of service provisioning and defines the concept of Quality of Service and Classes of Service. Section 3 and 4 describe, compare and analyze *ATM* and *RSVP*, two different approaches for quality provisioning, and presents the performance results obtained during a test program carried out by INFN-CNAF. Section 5 focuses on *Multiprotocol Label Switching*, a technique for enhanced packet forwarding, very suitable for ATM infrastructures, which might also provide support for classes of service. In particular, we describe the testbed set up to test *tag switching* in collaboration with the European task force *tf-ten*, and we compare the performance and advantages of tag switching to the traditional IP routing scheme. Section 6 describes the project which aims at developing a new ATM-based network for the Italian national research community: *GARR-B*. We analyze how the network architecture can be integrated with RSVP and tag switching. In section 7 we draw the conclusions.

2 QoS and CoS

Quality of service (QoS) provisioning is important to distinguish traffic in the network according to specific criteria, for example in terms of packet drop priority, queuing delay, delay variation tolerance, average and peak data rate. In this way resources allocation can be tailored to meet the user's needs.

ATM and RSVP are two different approaches which both address the QoS problem by providing per flow reservations at the cost of high complexity: both of them deploy an end-to-end signalling protocol and a mechanism for traffic flow specification. Classes of service (CoS) is a different solution for coarser grained quality of service provisioning. In CoS the concept of flow is substituted by "the class". Packets belonging to different streams are associated to a given class at the network border according to the packet preference. The core routers queue and forward each class with different priority.

The following sections show different architectures for QoS and CoS provisioning and compare them in terms of features, performance and utilization.

3 ATM

Asynchronous Transfer Mode (ATM), originally defined to support B-ISDN (Broadband Integrated Services Digital Network [2]), is one of the first technologies standardized to face the problem of service differentiation. Five classes of service were devised to transmit different types of information (data, voice and video), namely: CBR (Constant Bit Rate), ABR (Available Bit Rate), VBR-rt (Variable Bit Rate real time), VBR-nrt (VBR non real time) and UBR (Unspecified Bit Rate) [3]. ATM has an advantage in comparison with the current single service class architecture: the capacity of setting up an end-to-end virtual circuit with specific service features through a signalling protocol. In this way each flow is defined by a traffic profile and makes use of a dedicated connection. ATM connections can be set up and released dynamically, in this manner thanks to statistical multiplexing, bandwidth utilization is more efficient.

3.1 Features and performance

The advantages of ATM are manifold. ATM is a flexible technology which can be applied in the local area as well as in the metropolitan and wide area. It provides high bandwidth for the desktop when applied in the LAN and a large range of data rates up to 622 Mbps and more in the trunk. vBNS is an example of network backbone making use of high speed ATM links [4].

Test results of TCP/IP and UDP/IP over ATM show that in an ATM network (without IP routers) full link capacity utilization is achieved. A necessary condition for this is the tuning of the transport and application layer parameters like: socket buffer sizes at the sending and receiving end-systems (to act on the TCP window size [5]), the TCP_NODELAY option to control the Nagle's algorithm (introduced to avoid the *silly window syndrome*[6]); and the application message size. ATM performance measures were collected by means of benchmarking programs using TCP/IP over ATM [7]. The driving factors for this choice are the more and more widespread use of TCP and UDP/IP to support applications and their predominance over native ATM applications.

Figure 1 shows the relationship between throughput achieved by a memory-to memory data transfer and the send socket buffer size of the sender. A TCP stream was generated in an ATM LAN between two end-systems (DIGITAL Alpha station 3000/600 and 3000/400 running DIGITAL Unix) connected to a DIGITAL GIGAswitch/ATM through OC-3c network interfaces.



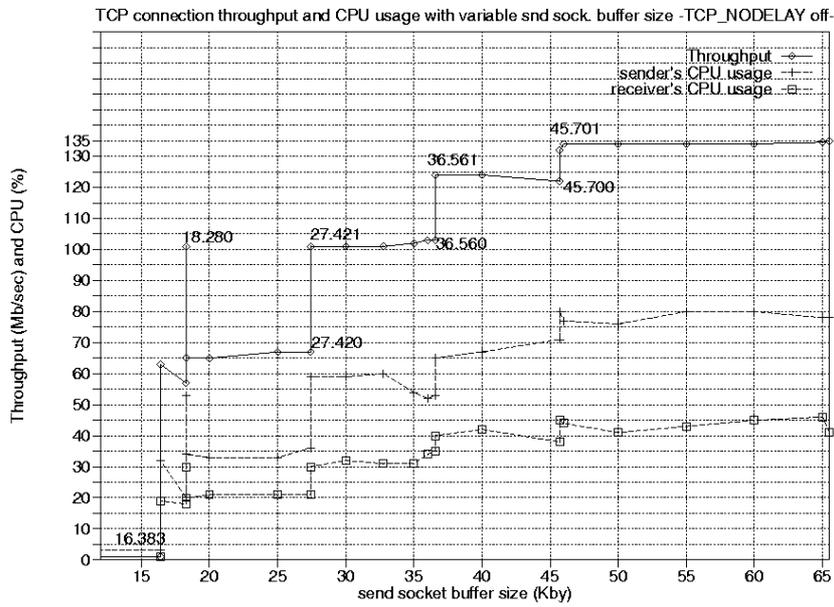


Figure 1: full link capacity utilization in an ATM local network

The maximum throughput, about 135 Mbps, is obtained when the local send socket buffer size is greater than 40700 bytes. As the graphic shows, the send socket buffer size is a key parameter. Efficiency of TCP/IP over ATM varies a lot depending on the way TCP/IP is implemented in the operating system. In the case of DIGITAL Unix, the throughput curve is not linear because the operating systems takes the socket size value set by the user and rounds it up to the minimum integer multiple of the MSS (Maximum Segment Size), which is 9140 bytes in the case of IP over ATM. The graphic expresses also the percentage of CPU utilization both at the sending and receiving node. CPU utilization is strongly dependent on the efficiency of the operating system TCP-UDP/IP implementation.

In a pure ATM network (where "pure" means that no routers are used, i.e. just ATM switches and ATM interfaces), TCP/IP performance is very good and close to the theoretical throughput. In this case the correct setting of the TCP/IP parameters, especially of the socket sizes, is fundamental because connections are characterized by long round trip times (RTT). If the TCP window is not large enough to cope with the long RTT, the sending process behaves in a *stop and wait* manner with a consequent inefficient use of bandwidth.

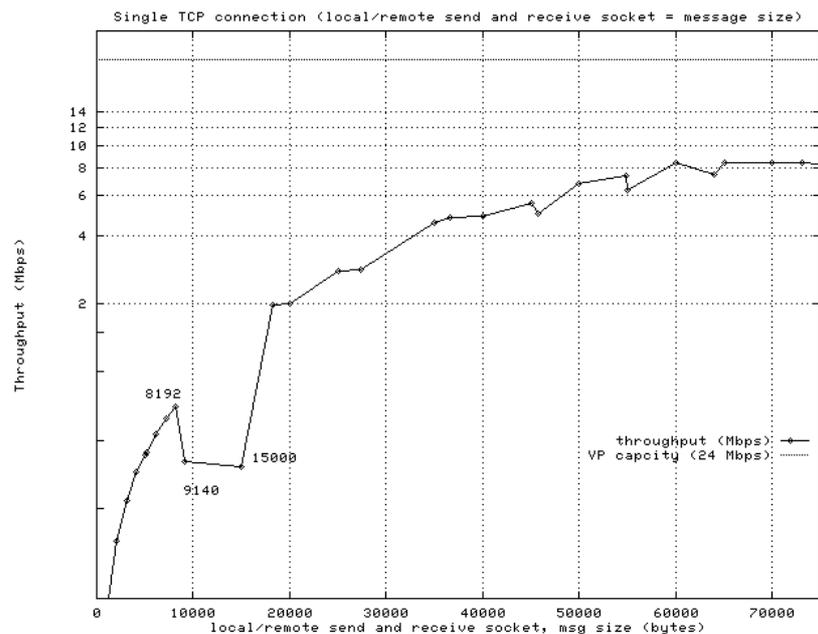


Figure 2: partial utilization of bandwidth on a WAN ATM connection because of large RTT and small TCP window

Figure 2 shows the throughput of a TCP connection over a geographic 24 Mbps CBR ATM PVC (Permanent Virtual Circuit) connecting two workstations (a Sparc Station 5 and an Indigo) in Italy and Sweden. The ATM connection was provided by JAMES, an ATM network for the support of European research projects [8]. The test was carried out in collaboration with the task

force *tf-ten* of TERENA [9].

The test was performed by modifying at the same time the send and receive socket buffer sizes of both the sending and receiving host, and the message size. These parameters were set to the same value.

In this test the RTT of a maximum size IP packet (9180 bytes) was 50 msec. End-systems were not deploying the window scaling option, which enables hosts to use window sizes larger than 65535 bytes. As a result the throughput increases until the socket size is 65535 bytes and then it keeps constant. The maximum achieved is only about 9 Mbps, while the capacity is 24 Mbps. This is the effect of the stop and wait behavior. Tests show that the socket buffer size in the end-systems is not enough, full utilization of bandwidth can be achieved by increasing the number of concurrent TCP streams. By running multiple TCP connections between two senders and one receiver in the same testbed, the aggregate utilization, i.e. the sum of the throughput achieved by each single stream, reaches 21 Mbps. For a complete description of the test program results refer to [9].

Nevertheless, according to our test results the performance achieved with "pure" ATM network is better than with networks including also IP routers, especially in presence of multivendor equipment. In the latter case bandwidth utilization achieved by TCP streams is not full, not only with single half duplex connections but also with multiple half duplex and/or full duplex concurrent TCP flows. The problem could be due to bad interoperability, to the queue management in the routers or to the reassembling and segmentation of IP datagrams in the router. This analysis is argument of current work.

3.2 ATM signalling

A major problem in the use of ATM is the limited range of ATM services offered by the carriers. For example, in several European countries signalling is not supported: only permanent connections can be configured. This implies that allocation of bandwidth must be static with a consequent inefficient use of resources. Thus, in the wide area ATM signalling can be deployed only by tunnelling the signalling cells into a PVC and by forcing new connections to belong to this PVC. If it is a CBR PVC, then SVCs (Switched Virtual Circuits) must be properly shaped in order to satisfy the policing adopted by the carrier.

European tests done in collaboration with the task force *tf-ten* show that SVC set-up times can be high [10]. Considering a local area ATM network made of end-systems and a single ATM switch, with UNI 3.1 set-up times vary in the range [20..30] msec. Set-up times depend on the end-points of the SVC and on their interoperability.

In the wide area set-up times from an end-system in Italy to end-systems and/or routers in Europe were more than twice or thrice the normal IP datagram round trip time measured in the same network set-up, as shown by the following table.

destinations:	Austria (Cisco 7010)	Switzerland (Sun)	Germany (Sun)	Germany (Cisco 7010)	Belgium (Cisco)
IP RTT	18	7	47	47	74
set-up time	66	44	82	86	142

table 1: ATM connection set-up times

The importance of dynamic ATM connections is fundamental to achieve an efficient use of resources and end-to-end quality of service. To make a widespread use of SVCs, an improvement in signalling performance, an NSAP addressing scheme and a plan for E.164-NSAP address translation must be devised. Nevertheless, because of the high cost of ATM equipment, of the dominance of TCP-UDP/IP applications and of the emergence of new technologies for the high speed in the LAN (like Fast Ethernet and Gigabit Ethernet), ATM will be more deployed in the WAN.

4. RSVP

RSVP, the Resource Reservation Setup Protocol, is the receiver driven signalling protocol devised by IETF [11]. Given a flow specification generated by the sending application, it's the receiver which selects the reservation and specifies its profile by means of a message which is processed by the intermediate nodes on the path to the source. The reservation state needs to be periodically refreshed by the receiver. RSVP is interesting because it can provide end-to-end per flow reservations independently of the layer 2 technology.

4.1 RSVP testbed



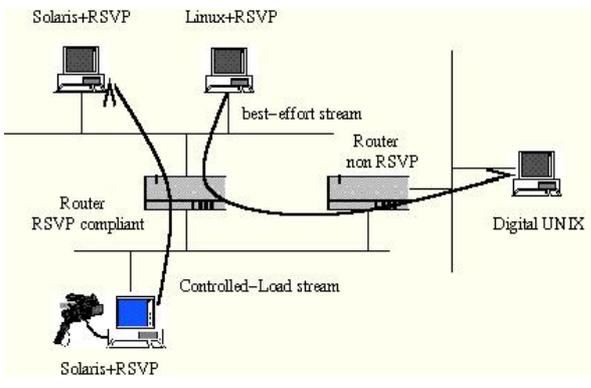


Figure 3: network set-up for RSVP tests in the local area

Local area RSVP tests were carried out at INFN-CNAF to analyze the functionality, the efficacy and the scalability of resource reservation. The network set-up consists of a CISCO 7505 RSVP capable router, of a traditional router and of four workstations (two of them running RSVP capable applications and the other best-effort data streams). Equipment was connected through Ethernet networks as shown in picture 3.

Mgen3.1 [12], *Vic* (a tool for videoconferencing), *Netperf2.1* [13] and *ftp* are the applications used: *Mgen3.1* and *Vic* to set up streams with reservation

guarantees, *Netperf2.1* and *ftp* for best-effort traffic.

4.2 RSVP performance

Both *Mgen* and *Vic* implement only two of the three types of RSVP service devised by IETF: the best-effort and the controlled-load [14] service. As a result, no performance measures were collected for the guaranteed service [15].

Both the mechanism for PATH and RSVP message exchange and for admission control work fine. Moreover, an admission control failure is correctly generated by a router on the path if a reservation exceeds the resource reservation threshold.

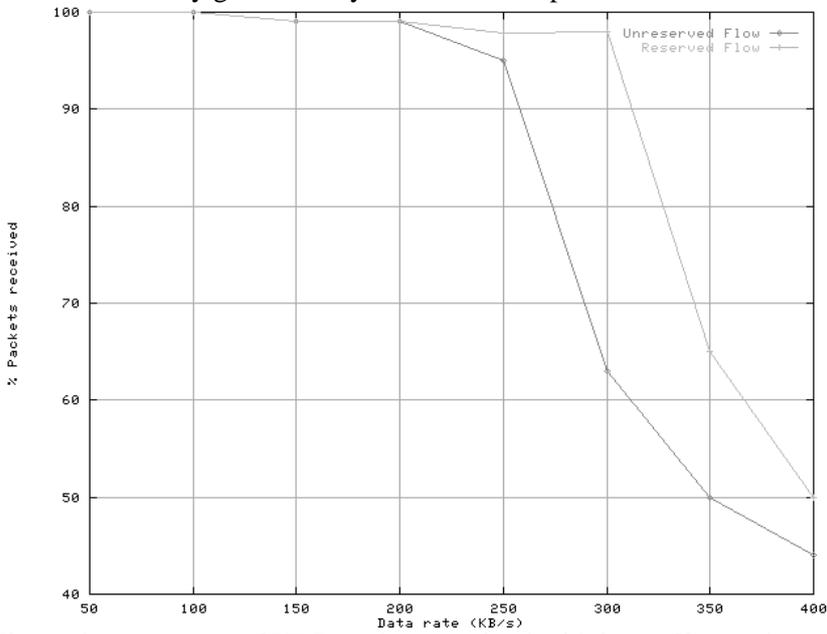
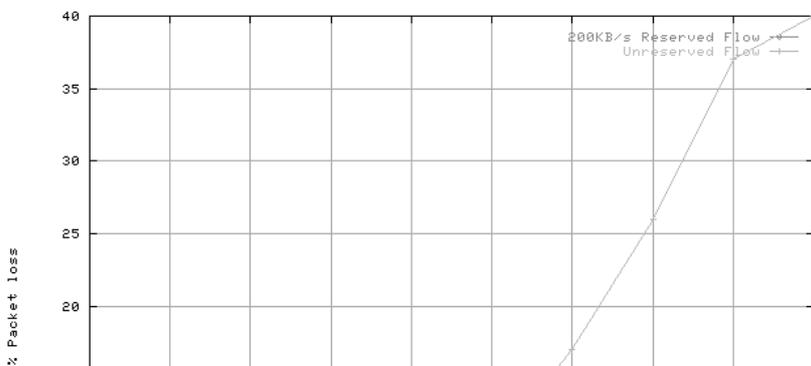


Figure 4: percentage of UDP packets received with best-effort and controlled-load service

The controlled-load service provided by RSVP is good as shown by the comparison of the packet loss percentage achieved with the best-effort and with the controlled-load service.

UDP *Mgen* streams with different data rates were generated between two hosts of the network. The reservation profile of a controlled-load stream was defined by two main parameters: *average* and *peak* bandwidth of the reservation. Figure 4 shows the results obtained for a reservation with peak and average equal to 200 and 300 Kbytes/sec respectively. The received packet percentage decreases dramatically for data rates bigger than 200 Kbytes/sec. On the other hand, the performance of the controlled-load stream is better: a relevant number of packets is dropped for data rates bigger than 300 Kbytes/sec. In this case packet loss is due to the presence of extra traffic on the two Ethernet networks (the LAN was not dedicated to the test).

When a best-effort and a controlled-load stream run concurrently, the RSVP quality of service is preserved independently of the data rate of the best-effort flow. As the data rate increases, also the controlled-load packet loss percentage increases, but it still keeps below 4% for a best-effort data rate equal to 4 Mbps.



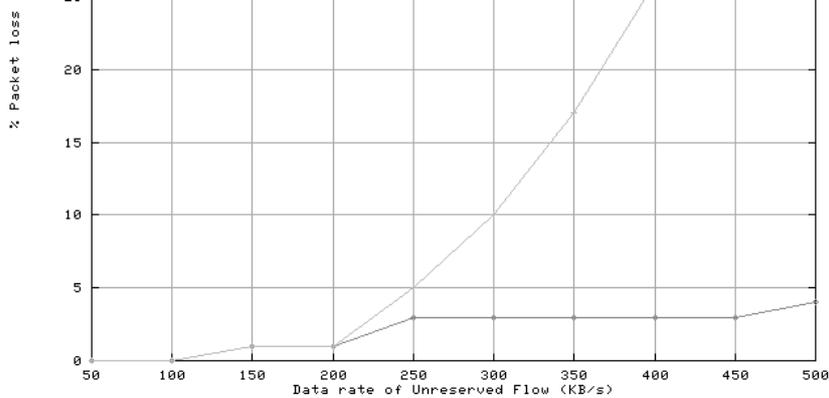


Figure 5: comparison traffic impact on a controlled-load and a best-effort stream

Figure 5 compares the results obtained for a reserved stream of 1.6 Mbps (200 Kbytes/sec) and a best-effort stream for several best-effort data rates.

Different results were obtained when the best-effort load is generated by several streams instead of just one as in the previous example. Against any expectation in this case the impact of best-effort traffic was much more relevant: the packet loss percentage of the controlled load stream duplicates, even if the aggregate best-effort load was the same. This problem is not clear, but it's probably due to the implementations of RSVP and/or CBQ tested in our environment.

RSVP scalability was also measured by generating up to 55 reserved streams (55 was a limit of the end-systems). In these conditions CPU utilization of the CISCO 7505 jumps from 0% to 2%. This increase is relevant given the small number of reserved streams. Further studies of scalability with more reserved streams in an extended LAN topology are argument of current work.

Major problems of the RSVP implementation tested are the lack of control on the amount of host resources needed to handle the reservation (for example the amount of CPU cycles) and of tools to manage the reservation when the local area network is congested. If the sending machine is overloaded, then also the throughput of the controlled-load stream is impacted. Moreover, the packet loss rate of the controlled-load stream increases significantly in presence of high traffic over the LAN connection.

5. Multiprotocol label switching

Quality of service provisioning is one of the most important requirement for the enhancement of both networks and applications, nevertheless, to make it work, fast and reliable connections should be available. The improvement and "simplification" of packet forwarding capabilities in the equipment is thus one the means to achieve these requirements. Multiprotocol Label Switching (MPLS) under development at IETF [16] and *IP switching* [17] are some of the possible techniques for packet forwarding enhancement. In this paper we focus on MPLS and in particular on *tag switching*, an implementation of MPLS by CISCO. MPLS is interesting for the design and implementation of backbones because it's two-fold: enhanced IP routing is combined with the flexibility of label switching, in particular with ATM. For this reason it could be a mean to integrate two different worlds and have benefit of the best features of both. MPLS can also offer an easy support for classes of service.

With MPLS routing tables are integrated with a label information base (named "Tag information Base" in tag switching jargon), so that packets to the same destination network/host/application can be identified by one unique label (or "tag"). Labels can be encapsulated in different ways: into the IP datagram itself (in the case of IPv6), in the layer 2 header (for example in the ATM cell header) or between the layer 2 and layer 3 header. Once the packet is associated to a label, in the next hops forwarding is implemented by just looking at the label. Intermediate systems (routers and/or switches) can perform label swapping, i.e. they replace the input label value with a different one in output. With ATM thanks to its small length and to the fixed position into the cell, switches implement label switching at wire speed.

Labels must be consistent throughout the network. This is achieved by the *Label Distribution Protocol* (LDP), which is "a set of procedures by which one Label Switching Router informs another of the label/stream mappings it has made".

This paper aims at showing the benefits of MPLS, its application for the implementation of a production network and the results obtained in a European test bed. For an extensive description of the protocol refer to the document repository maintained by the Multiprotocol Label Switching working group at IETF [16].

5.1 MPLS deployment

MPLS, and the implementation tested, *tag switching* [18], are technologies suitable for the design and implementation of an ATM-based wide area network for the following reasons. MPLS runs both on routers and ATM switches, thus the label switching forwarding function can be easily mapped into ATM cell switching by carrying it into the VCI and/or VPI field of the ATM cell. In this way we have a gain in performance and speed of packet forwarding is obtained in this way.

With MPLS the IP infrastructure can be more scalable because the number of IP connections between routers could be reduced: Let's consider a network environment made of a core ATM network and border IP routers. With the traditional IP scheme to achieve IP connectivity with only two hops router to destination, a full mesh of ATM connections must be configured (traditional ATM switches has no IP routing capabilities). With MPLS ATM switches support IP routing functions, for this reason each router has only to configure a PVC with the switch it is directly connected to.

Another advantage is the dynamic configuration of ATM connectivity in the MPLS core network: "Tag Virtual Circuits" (TVCs), in tag switching jargon, do not need to be configured manually. LDP deploys a dedicated VCI and establishes tag virtual circuits. VPI and VCI ranges deployed for standard ATM communication and for tag switching, are kept separated.

In IP over ATM, in order to connect hosts belonging to different subnetworks crossing only two routers a complete mesh of VPIs/VCI between the subnetworks is necessary and must be manually configured. On the other hand, with MPLS a complete mesh of TVCs is configured between networks and the number of router hops is not increased.

Finally, as already mentioned, MPLS is a promising technique for the support of classes of services, since it provides a straightforward mean for packet classification. At the ingress point of the label switching cloud a given label is associated to each packet depending on the network or application layer information carried in its header. Through the binding between labels and classes of services, the packet can be queued consistently throughout the label switching cloud.

5.2 A case study: tag switching

5.2.1 Tag switching testbed

Tag switching tests were performed by the task force *tf-ten* using the European ATM infrastructure JAMES already mentioned in section 3. Participants (Austria, France, Germany, Italy, Spain and Switzerland) were interconnected through a set of ATM CBR permanent virtual circuits. The link bandwidth was 4515 cells/sec in two cases: on the France-Spain and on the France-Germany link, and 4750 cells/sec in the rest of the international links. A loop between Austria, Italy and Switzerland was configured to verify the use of the metric for a correct route computation between the three countries. The *tag cloud* (corresponding to the area inside the dashed line in figure 6) was made of LightStream1010 switches and CISCO routers of the 7500 and 7200 series, all running tag switching. For a more detailed description of the test see [19].

Hierarchical IP routing configuration was implemented. Switches and Routers in the core run interior routing (OSPF), belonging to the same backbone area 0. Border routers (*192.168.11.1*, *192.168.21.1*, *192.168.31.1*, *192.168.41.1*, *192.168.51.1* and *192.168.61.1*) were adjacent, since they were directly connected through TVCs, set up by TDP thanks to the OSPF protocol that runs both in routers and switches. Border routers run both interior and exterior BGP sessions. The external session was run by peering with routers outside the tag cloud. IP datagrams with same destination prefix were associated to the same tag. In this way per flow TVCs are not necessary and the number of ATM connections is reduced.

Three workstation (two Sun Ultra and a Sparc Station 10) were deployed to generate end-to-end TCP streams. The goal was the comparison between the performance achieved in the same network set-up with and without tag switching. Tag switching control protocol was tunnelled into the international PVCs, so tag switching was completely transparent for the ATM equipment on the public network side.

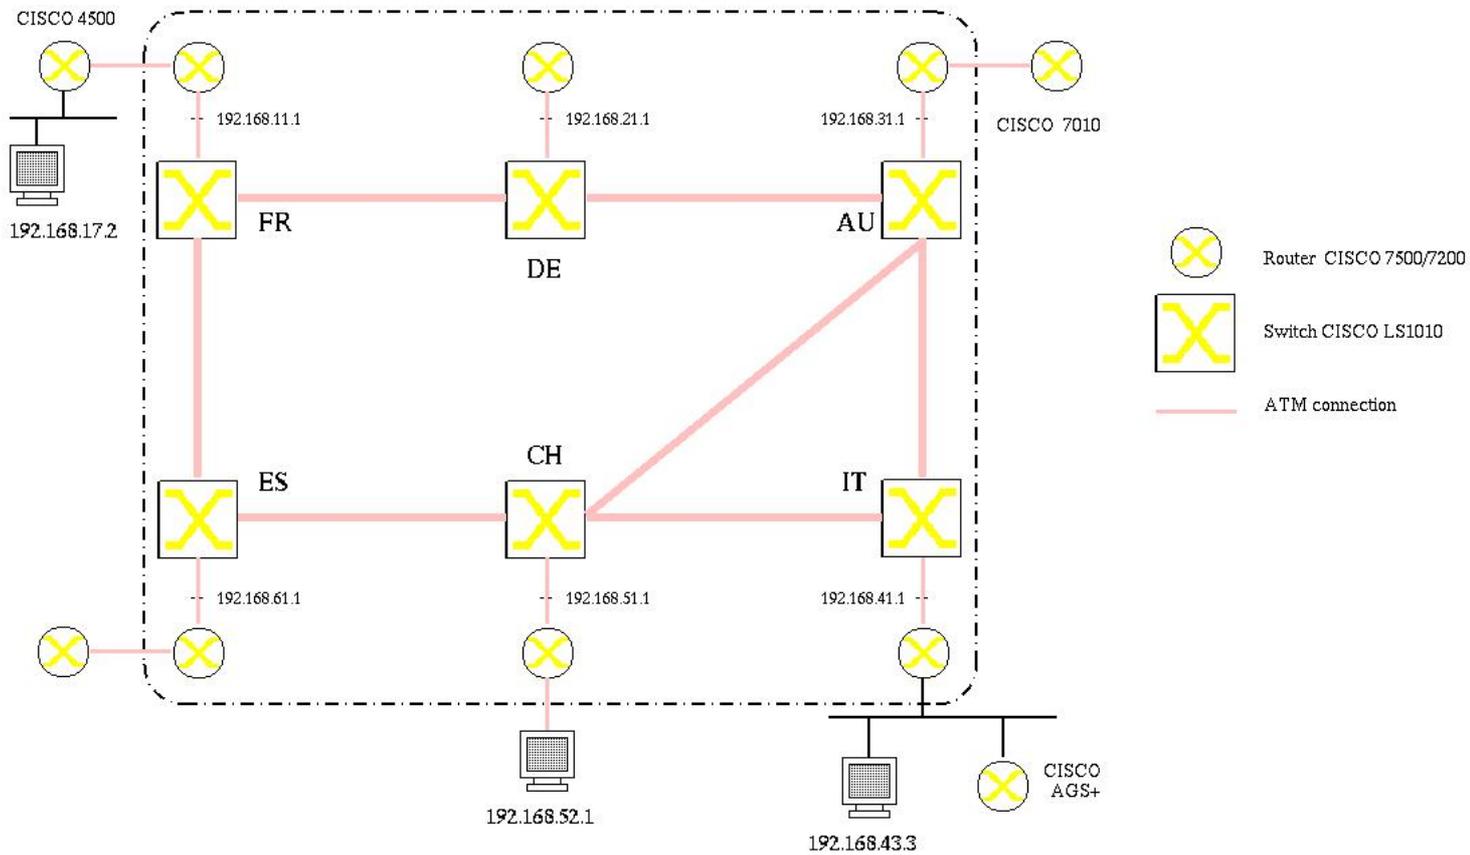


Figure 6: tag switching test bed

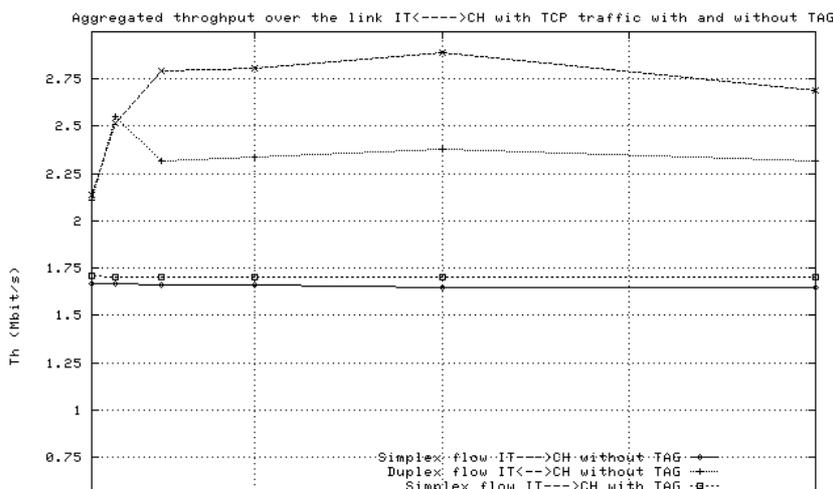
5.2.2 Tag switching functionality

Tag switching showed good functionality. Automatic tag virtual circuits (TVCs) were generated by each node for each network prefix and neighbor in the routing table. The following is an extract of the binding table of router 192.168.41.1, in which IP networks and hosts are associated to the corresponding TVC VPI/VCI. 192.168.10.0 is the IP network of the French ATM switch

Destination	Headend Router	hops	VPI/VCI	Status
Destination: 192.168.40.1/32	Headend Router ATM2/0.10	(1 hop)	2/50	Active, VCD=23
Destination: 192.168.10.0/24	Headend Router ATM2/0.10	(4 hops)	2/33	Active, VCD=7
Destination: 192.168.11.1/32	Headend Router ATM2/0.10	(5 hops)	2/34	Active, VCD=8

We also tested *traffic engineering* i.e. the capability of configuring preferential routes for traffic selected through filters, and it worked. With traffic engineering a set of streams can be treated differently from the rest of the IP datagrams by overriding the standard routing table. The goal of this approach is traffic tailoring to maximize bandwidth utilization.

5.2.3 Tag switching performance



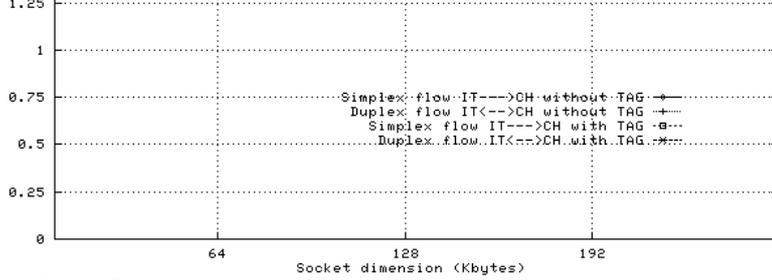


Figure 7: performance gain with tag switching

Full PVC capacity utilization was achieved during the tests with tag switching. Picture 7 shows the application layer throughput achieved by half and full duplex TCP connections between Italy and Switzerland and compares the results measured with and without tag switching

With tag switching TCP performance is slightly better than with classical IP routing independently of the socket buffer sizes (send and receive buffer sizes were set to the same value), this is explained by the different AAL5 encapsulation schemes used by tag switching.

TVCs use AAL5 VC based multiplexing encapsulation, while ATM PVCs deploy AAL5 LLC-SNAP encapsulation [20]. With LLC-SNAP 8 bytes (LLC header plus SNAP header) are added to the IP PDU when it's encapsulated into the AAL5 CPCS PDU payload. On the other hand, with VC based multiplexing no overhead is added at all with a consequent performance gain which depends on the IP PDU size distribution, i.e. on the number of padding bytes added in the AAL5 CPCS PDU.

With half duplex connection and tag switching the maximum throughput was 1.70 Mbps, about the maximum achievable by an application. In fact, considering the maximum size of an IP PDU (1500 bytes -two end-systems had Ethernet interfaces), the resulting size of the AAL5 CPCS PDU is 1536 bytes, i.e. 32 cells. Since the ATM link rate between Italy and Switzerland was 4750 cells/sec, this corresponds to 148 MTU-long IP PDUs/sec, i.e. 1.73 Mbps of application layer throughput.

6 Architecture design for an advanced backbone network: GARR-B

GARR-B is the project for the implementation of a new infrastructure for universities and national research institutes in Italy. The name GARR-B also identifies the network infrastructure (figure 8). GARR-B addresses the need of an increased network capacity up to the user access node, of higher network reliability and of a wider range of quality of services to support specific applications like videoconferencing, remote X-sessions, access to remote file systems and distributed computing.

Scalability, flexibility and simplicity of management are three fundamental requirements of this infrastructure. With scalability we mean the capability of an architecture to cope with an increase in the backbone capacity and in the user community size. On the other hand, flexibility is the ability of a network to adapt to new requirements and new technologies.

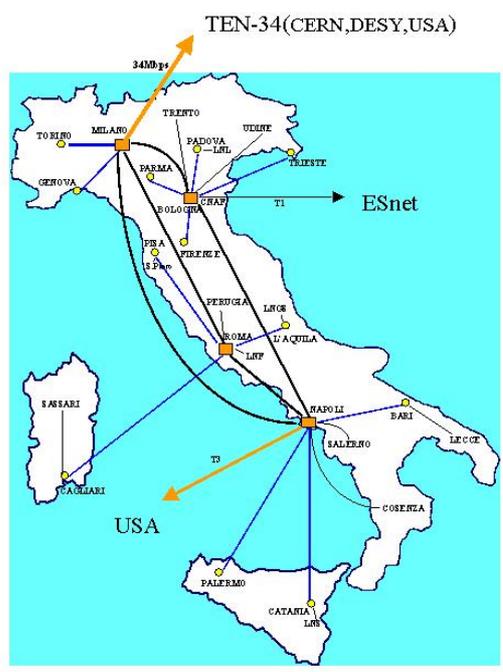


Figure 8: plan of GARR-B topology

6.1 Architecture design

To satisfy the requirements listed above a hierarchical network infrastructure has been designed. Three are the main components: the core ATM *transport network*, the *access network* and the *user sites* (see figure 9). User sites will access the network with speed in the range [2..34] Mbps. The number of user sites will be from 200 to 300.

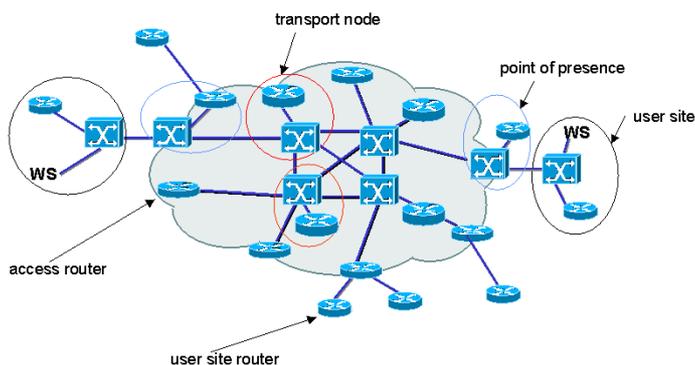


Figure 9: hierarchy in GARR-B: transport network, access network and user sites

The transport network is made of transport nodes, i.e. routers connected to ATM switches, which are interconnected together through a full mesh of CBR ATM permanent virtual circuits. The user equipment (routers and/or switches) is connected to the transport infrastructure through the *Point of Presence* (PoP). Routers can be deployed for classical IP connectivity. On the other hand, ATM switches also supply direct ATM connectivity to end-systems. In this way special purpose applications have the benefit of direct end-to-end reserved connections and private connections can be built on top of GARR-B.

6.2 RSVP and MPLS in the GARR-B production environment

6.2.1 GARR-B and RSVP

The design of GARR-B has several drawbacks, both from the ATM and from the IP point of view.

First of all, the core ATM network is completely static. This means that the bandwidth is statically provisioned per PVC and no statistical multiplexing gain is achieved. Moreover, ATM connections for special purpose applications have to be manually configured and managed, and per flow quality of service cannot be supplied. Nevertheless, QoS can be provisioned by enabling a selected set of end-systems to establish reservations through RSVP. Even if tests showed that some of the current protocol implementations are not totally stable and still incomplete as compared with its specifications, RSVP can be viable in a carefully controlled environment.

The well known poor scalability of RSVP in a wide area network does not cause major problems if RSVP is used to support a limited number of special purpose applications.

6.2.2 GARR-B and MPLS

At the IP layer, the routing hierarchy implies that in many cases traffic going through the shortest path between two end-systems has to cross several routers. This disadvantage could be overcome through the use of label switching techniques in the following way.

- The role of transport node (router plus switch) can be played by a single ATM switch running label switching. In this manner the ATM core infrastructure will be made of TVCs with the advantages that the inner IP topology of the core is transparent to the rest of the network and the TVC infrastructure (i.e. the ATM core) is dynamically set up.
- The role of access router can be played by a label switching router. In this way the access router is directly connected to the rest of the access routers through TVCs. The hierarchy is simplified by removing a level and the number of routers in the end-to-end path is reduced. The resulting network is still manageable and scalable.

In the future classes of service could be easily supported by label switching. Fine grained per flow quality of service could be also supported through the integration of RSVP and ATM, which is under study at IETF [21]. ATM signalling and PNNI [22] will be fundamental to make this interoperability viable.

Picture 10 compares the current GARR-B infrastructure and the architecture modified to support label switching.

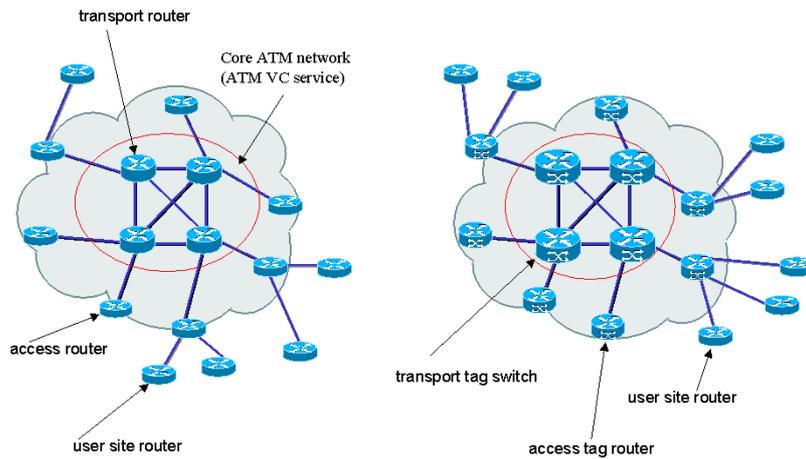


Figure 10: GARR-B architecture: current model and enhanced model to support label switching

7 Conclusions

Two were the goals of the paper: the focus on the importance of enhanced network features like quality of service and efficient packet forwarding, and the analysis of how these capabilities can be used and integrated in a production environment. Sections 3, 4 and 5 show the details of ATM, RSVP and MPLS, and present the results of tests carried out in local and wide area testbeds. ATM, RSVP and MPLS are important techniques which can be used to address different goals. ATM is an efficient technology for the implementation of high speed networks in a flexible and scalable way. If end-to-end ATM signalling is not possible, RSVP can be deployed to dynamically establish per flow reservations for selected applications and users. The presence of an ATM core network gives also the possibility to improve the packet forwarding capabilities of the network. The application of MPLS in the design of a national research ATM-based network like GARR-B, is useful to achieve simplicity and scalability.

8 Acknowledgments

The ATM test program was developed at INFN in the framework of the ATM Project (Gruppo V). We thank Simon Leinen, Jean-Marc Uzé for coordinating the RSVP and the tag switching tests in the wide area network, and all the partners in the task force tf-ten for their valuable collaboration. We thank also CISCO Systems Europe for the provisioning of test equipment and beta software for the tag switching test program. The contribution of Andrea Chierici and Alessandro Canzian was fundamental during the RSVP and tag switching test program.

9 Bibliography

1. TEN-34, *the Information Superhighway for European R&D*; <http://www.dante.net/ten-34/>
2. International Telecommunication Union; *B-ISDN Protocol Reference Model and its Application*; Recommendation I.321.
3. ATM Forum; <http://www.atmforum.com/>
4. vBNS engineering; *vBNS Quality of Service Architecture*; http://www.vbns.net/presentations/vbnsqos/vbnsqos_public.html
5. Clark, David; *Window and Acknowledgment Strategy in TCP*; RFC 813.
6. Braden, Robert; *Requirements for the Internet Hosts*; RFC 1122.
7. Laubach, Mark; *Classical IP and ARP over ATM*; RFC 1577.
8. JAMES, *Joint ATM Experiment on European Services*; <http://www.labs.bt.com/profsoc/james/>
9. Task Force tf-ten, TERENA (<http://www.dante.net/tf-ten/>); *Deliverable 11.3: Results of Phase 1 Test Programme*; 30 April 1997; <http://www.dante.net/ten-34/DELIVERABLES/D11.3/>
10. Alborghetti, Stefania; Campanella, Mauro; Colombo, Diego; Maggi, Simone; *Study of SVC set-up times*; July 1997.
11. *Resource Reservation Setup Protocol*; <http://www.ietf.org/html.charters/rsvp-charter.html>
12. *Mgen 3.1*; <http://tonnant.itd.nrl.navy.mil/ipresearch/mgen.html>
13. *Netperf*; <http://www.cup.hp.com/netperf/Netperpage.html>

14. Wroclawski, John; Partridge, Craig; Guerin, Roch; *Specification of the Controlled-Load Network Element Service*; RFC 2211.
15. Shenker, Scott; *Specification of Guaranteed Quality of Service*; RFC 2212.
16. *Multiprotocol Label Switching working group*; <http://www.ietf.org/html.charters/mpls-charter.html>
17. *IP switching*; <http://www.ipsilon.com/about/technology/>
18. *Tag switching*; http://www.cisco.com/warp/public/732/tag/tag_resources.html
19. Task Force tf-ten; *Label-based Switching Experiment*; <http://www.urec.fr/Renater/jmu/jameslbs.html>
20. Heinanen, Juha; *Multiprotocol Encapsulation over ATM Adaptation Layer 5*, RFC 1483.
21. Crawley, Eric; Berger, Lou; Berson, Fred; Borden, Marty; Krawczyk, John; *A framework for Integrated Services and RSVP over ATM*; draft-ietf-issll-atm-framework-01.txt.
22. ATM Forum; *P-NNI v1.0*; <http://www.atmforum.com/atmforum/specs/approved.html>