#37

TEN-155: High Speed Networking for European Research

Roberto Sabatino

Jan Novak

This paper will be presented by Roberto Sabatino at the INET Conference in San Jose, California on 23 June 1999.

The Internet Society
ISBN#: 1-891562-06-1
INET'99 The Internet Global Summit
(Proceedings - CD-Rom)

DANTE IN PRINT is a track record of papers and articles published by, or on behalf of DANTE. HTML and Postscript versions are available from: http: / / www.dante.net / pubs / dip

For more information about DANTE or *DANTE IN PRINT* please contact:

DANTE
Francis House
112 Hills Road
Cambridge CB2 1PQ
United Kingdom

Tel: +44 1223 302992
Fax: +44 1223 303005
E-mail: dante@dante.org.uk

TEN-155: High Speed Networking for European Research

Roberto Sabatino

Jan Novak

Abstract

This paper addresses the implementation of a pan-European network to support co-operative research amongst European researchers: TEN-155. Besides providing a high speed IP service, the purpose of TEN-155 is also to support research in networking by providing an international test bed for advanced networking technologies (the Quantum Test Programme) and by providing VPNs, with dedicated and guaranteed bandwidth, for specific research projects in countries connected to TEN-155. Predecessors to TEN-155 were TEN-34 and EuropaNet. TEN-155 supersedes these two networks not only in terms of capacity offered but also in terms of the Managed Bandwidth Service (MBS) offered.

This paper will illustrate the use of ATM technology to support the MBS and the Quantum Test Programme in co-existence with the standard best efforts IP service.

KEYWORDS: TEN-155, ATM, Managed Bandwidth, Quantum Test Programme

Introduction

The success of EuropaNet[1] and TEN-34[2] has demonstrated that a dedicated networking infrastructure to the European research community is essential for successful co-operation amongst European researchers. Both these networks have had a short lifetime, mainly due to the excessive cost of international capacity. Nevertheless, with the help of funding from the European Commission (EC) it has been possible to deploy them and demonstrate their vital importance to the research community. EC funding for TEN-34 ended in July 1998, and as a consequence effort has been put into the Quantum project, which led to the deployment of a replacement network for TEN-34: TEN-155.

The paper will outline the Quantum project and describe in detail the resulting TEN-155 network which supersedes TEN-34 in terms of available bandwidth and above all in the ability to offer a Managed Bandwidth Service to guarantee end-to-end Quality of Service. TEN-155 makes use of ATM, considered the most effective technology to offer guaranteed capacity end-to-end.

A section is dedicated to the Quantum Test programme, the purpose of which is to evaluate new and emerging technologies and investigate the possibility of their deployment on the production network.

Finally, a description of the Managed Bandwidth Service (MBS), in terms of its organisation and development phase will be provided.

2. The Quantum Project

The Quantum project (Quality Network Technology for User-Oriented Multi-media, http://www.dante.net/quantum) foresees the exploration and implementation of providing Quality of Service across a pan-European network of very high speed. The Quantum project also calls for experimentation of new IP and ATM technology using a Wide Area and international test network.

TEN-155 is the operational network built as

Jan Novak and Roberto Sabatino are Network Engineers at DANTE. Their email addresses are Jan.Novak@dante.org.uk and Roberto.Sabatino@dante.org.uk .

a result of the Quantum project.

A group of 16 national Research Networks and one regional network, co-ordinated by DANTE, are responsible for the Quantum Project which is co-funded under a joint initiative by DGXIII (Telematics for Applications, Esprit and ACTS) of the European Commission.

DANTE is a not-for-profit company set up in 1993 by European National Research Network organisations. DANTE plans, builds and manages advanced networking services for the European research community.

3. TEN-155 physical topology

The physical topology of TEN-155 was dictated by a combination of the following factors:

- cost of international circuits;
- traffic flow requirements, derived from TEN-34 traffic analysis;
- traffic growth prediction;
- what the tenderers were able to offer;
- overall cost of the network;

Following the issue of an open tender, several offers from the tenderers were evaluated, and in August 1998 a contract was awarded to Unisource Belgium for the provision of a number of 155Mbps SDH circuits, the supply and management of an ATM service in all TEN-155 countries, and for facilities management in most of them. Unisource Belgium relies on KPN (Netherlands) for the implementation and technical support of these services. Fig.1 illustrates the physical toplogy of TEN-155.

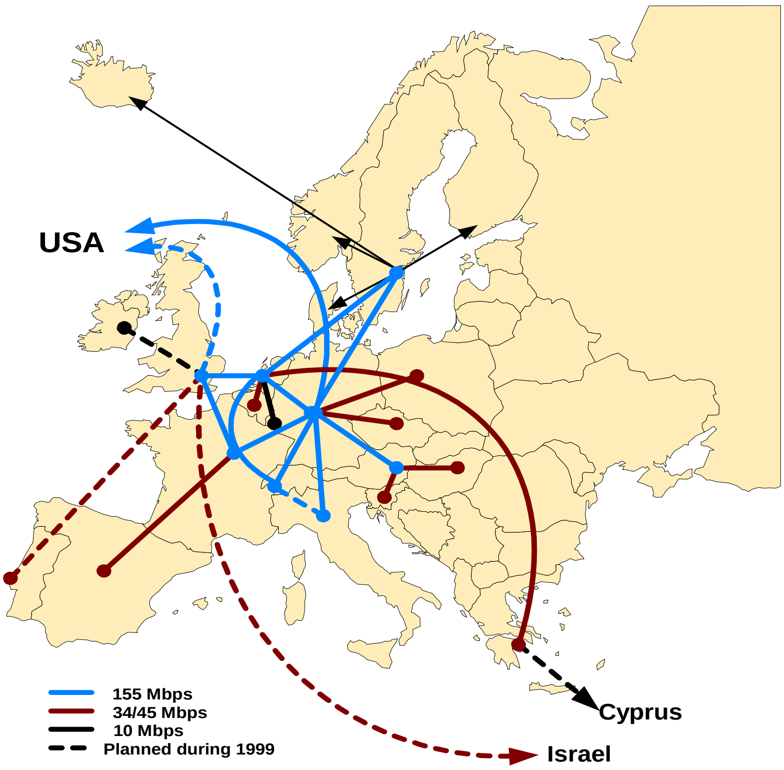

Fig. 1 TEN-155 physical topology

The figure outlines the existence of transit nodes (AT CH DE FR IT NL SE UK), interconnected via non protected SDH STM-1 circuits, to which NRNs are directly connected in addition to international circuits to peripheral sites (HU SI GR ES CZ LU BE PT PL IE) or other transit sites.

When comparing TEN-155 to TEN-34, a huge increase in available capacity is immediately noticeable. On TEN-34 the highest bandwidth available was a 34Mbps leased line between Germany and Switzerland, now it is 155Mbps on many circuits. In addition, even in countries where capacity still has relatively high prices (Greece, Portugal, Slovenia) it has been possible to deploy 34Mbps circuits. As a result of the 1998 liberalisation of European telcoms services, the overall cost of TEN-155 is similar to that of TEN-34, but the amount of available bandwidth is much higher.

4. The Engineering of TEN-155

The engineering of the TEN-155 network took into account that:

- initially a best efforts IP service is to be provided, to ensure continuity of service from TEN-34;
- a MBS service is to be provided. Developments in the telecommunications market indicated ATM as the only reliable way to provide this end-to-end and across network management domains;
- ATM could also be used as a bandwidth management tool (for example, to distribute IP traffic across available circuits).
- it must be possible to setup, using the MBS, a VPN dedicated to the evaluation of networking technology. This evaluation is a component of the Quantum project and is implemented in the Quantum Test Programme;
- it must be possible to implement value-added services such as advanced monitoring and Mbone.

These requirements and conditions led to the following design decisions:

- as ATM was indicated as the most reliable way to provide MBS, the international SDH circuits are terminated on ATM switches dedicated to TEN-155;
- to provide an IP service to the NRNs, IP routers managed by TEN-155 are installed in the transit nodes of TEN-155. NRNs connect to TEN-155 by setting up BGP sessions between their access routers and the nearest TEN-155 router;
- to support at the same time best efforts IP traffic on TEN-155 and guaranteed traffic for MBS, UBR-like VCs between routers to carry the IP traffic are established whilst CBR VCs are established between the end points that require MBS. ATM switches that enable co-existence of the two ATM Traffic Classes (ATCs) in a dynamic manner have been deployed: best effort IP traffic (UBR) may take up the whole bandwidth if no guaranteed traffic (CBR) is flowing, whilst in cases of congestion the switch will drop the UBR-like traffic and let through the CBR or guaranteed traffic;
- to minimise the required capacity between switch and router (which in turn means minimising the number of STM-1 interfaces on switches and routers) a full mesh of UBR-like PVCs is set up between the TEN-155 routers in the transit nodes. In addition the full mesh allows to load balance the bandwidth usage independently of IP routing and provides more stability at IP level;
- Workstations are installed in the transit PoPs to allow implementation of advanced monitoring tools and deployment of Mbone.

Fig.2 illustrates an example PoP setup. The example outlines the benefits of a full mesh of UBR-like PVCs to carry the IP traffic: the capacity (hence number of interfaces) required between switch and router is equivalent to the capacity of the connected NRNs, in this case DFN (155Mbps) and Cesnet (34Mbps). Therefore 3 STM-1 connections between switch and router are required. In the event of hop-by-hop PVCs, the capacity required between switch and router will have to accommodate all the transit traffic, i.e. 9 STM-1 interfaces will be necessary.

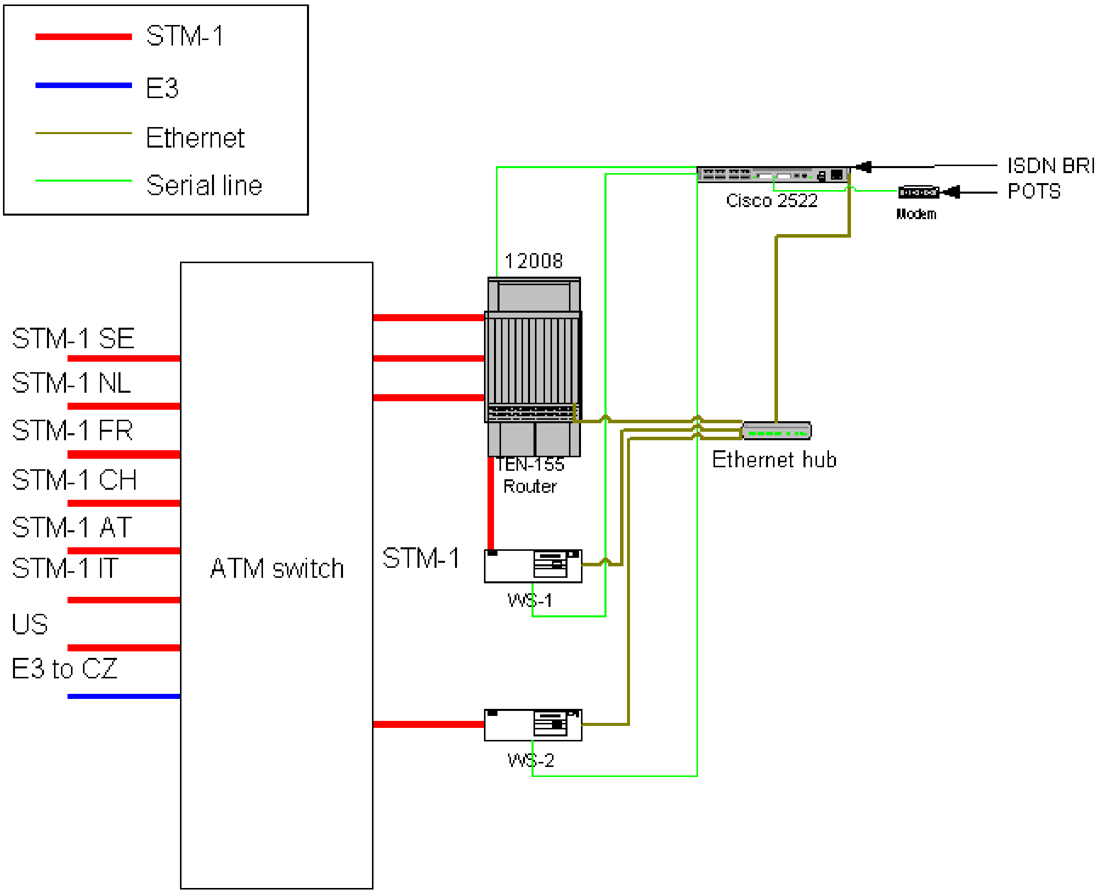

Fig2. TEN-155 DE PoP setup

The figure also outlines that 2 PoP workstations are available: one is to host operational services such as Mbone and advanced monitoring, the other for experimentation within the Quantum Test Programme.

The UBR-like PVCs are implemented by SBR3 (VBR, with SCR=10cps , PCR=line rate). Cell tagging is enabled on these PVCs, so all cells which exceed the contract (virtually all cells, given SCR=10cps) are tagged with CLP=1. Cell tagging is not enabled on the CBR PVCs, hence the cells are transmitted with CLP=0. In cases of congestion, the switch will discard CLP=1 cells first, and Early Packet discard (EPD) is implemented. This ensures that complete AAL-5 frames (i.e whole IP packets) are discarded - which has the beneficial effect of removing entirely from the network cells that no longer serve any useful purpose. In fact, IP packets are mapped onto AAL-5 frames, and if a single cell is dropped, the whole corresponding IP packet is invalid and will result in a CRC error detected on the router and in turn this represents a serious waste of bandwidth. By removing a complete AAL-5 frame the network is able to entirely accommodate the next incoming packet.

This mechanism has been thoroughly tested in a laboratory environment and in cases of 200% congestion, the result was to have 100% utilisation on the congested link and more than 95% goodput. Without the EPD mechanism, situations of congestion as low as 110% may result in approximately 15% goodput on some switches. Several switches have been tested, and Fig.3 illustrates a simple setup that was used for testing the EPD mechanism.

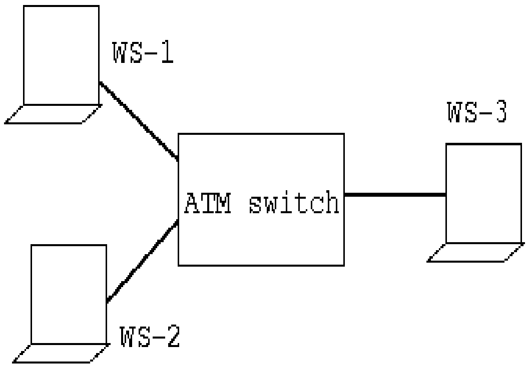

Fig.3 EDP test setup

Ttcp, with UDP, sessions between WS-1 and WS-3 and between WS-2 and WS-3 were run at the same time, therefore the bottleneck was the physical STM-1 connection between the switch and WS-3. The goodput was measured on WS-3 and was calculated by adding the results reported by WS-3 for the two ttcp sessions. The ATM software on the workstation was also able to report the number of CRC errors occurring when EPD was not enabled.

It is not possible to publish the detailed results due to NDAs with the switch suppliers.

On IP level, for the implementation of the best effort IP service, TEN-155 is configured as one AS with the necessary full mesh of iBGP peerings using loop backs. This avoids loss of connectivity in the event of an ATM PVC outage. TEN-155 uses OSPF as internal routing protocol with OSPF cost of ATM PVCs configured to reflect the underlying physical and ATM topology. As ATM level re-routing of PVCs is not configured, this set-up ensures:

- the selection of shortest path on physical / ATM level for IP traffic in operational status;
- the selection of shortest path on physical / ATM level for IP traffic when a failure of a main trunk occurs.

The rule applied for calculating the OSPF cost between two nodes (see Fig. 4) is:

```
100 for the first physical line, 20
for each subsequent ATM level hop.

Or 100 + 20 * (number_of_hops -1)
```

In other words, two IP level hops have always a longer path than one direct ATM VC between any two nodes as the highest ATM hop count in the backbone does not xceed 3. Therefore 2 IP level hops between two nodes would have cost 200, whilst the maximum cost for a direct ATM PVC is 140.

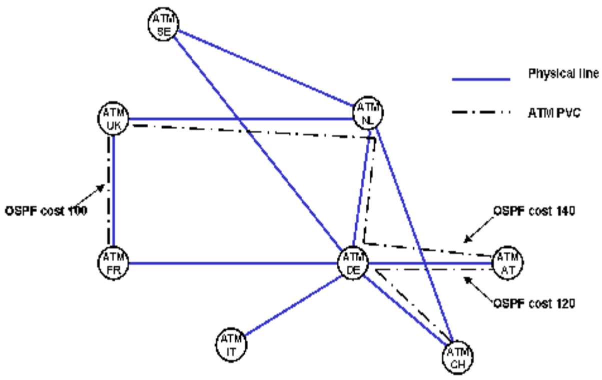

Fig. 4 TEN-155 OSPF setup

5. Impact of TEN-155 on inter-NRN traffic

One of the reasons for deploying TEN-155 was to provide more bandwidth for European research traffic, therefore an analysis of the immediate impact of providing more capacity was carried out.

Table 1 shows an extract (a full matrix would not be easily readable) of inter-NRN traffic measurements on TEN-34 in November 1998 (the last month of full operation of TEN-34) and on TEN-155 in January 1999. The values shown are monthly daytime averages of traffic between NRNs expressed in Mbps and are derived from DANTE's inter-NRN traffic statistics package [3], initially deployed on TEN-34.

From the table it is clear that the total traffic on TEN-34 was much higher. This is due to mainly two reasons:

- a considerable amount of NRNs had still not migrated to TEN-155 (IT,AT,HU, SI) and most of these were accessing their Global Internet service via TEN-34;
- on TEN-34 many commercial ISPs were indirectly interconnected to TEN-34, thus generating a significant amount of traffic between an NRN and a commercial organisation.

On the other hand, a significant increase of traffic between the NRNs that did migrate to TEN-155 is noticeable. This is partly due to the natural increase of bandwidth usage on the Internet, but also to the availability of more bandwidth, which removes bottlenecks from the network. On TEN-34, traffic between Nordunet(SE), SURFnet(NL), DFN(DE) and Janet(UK) was significantly affected by these bottlenecks.

the US. However it is still more economic for some countries to organise the setting up of shared transatlantic capacity. DANTE to this end has organised the procurement of a SDH / STM-1 circuit from Frankfurt to New York. The transatlantic SDH / STM-1 terminates on ATM switches (the TEN-155 switch in Frankfurt, and on another switch managed by DANTE in New York). The service offered to the subscribing NRNs is ATM directly into the NRN, so it is as if the NRN had their own transatlantic link on IP level. The advantages of this are both economic and technical. From the economic point of view it is cheaper to purchase SDH / STM-1 rather than separately purchasing capacities in the order of 15-40Mbps whilst from the technical point of view the set-up allows guaranteed capacity to the subscribing NRNs and simplifies the IP level setup for TEN-155.

TEN-34, November 1998					-	TEN-155, January 1999			
src/dst	DE	NL	SE	UK	-	DE	NL	SE	UK
DE	-	1.04	1.89	2.51	-	-	2.38	3.06	2.5
NL	0.96	-	1.18	0.96	-	2.74	-	1.87	1.34
SE	3.04	2.5	-	2.09	-	5.74	4.1	-	2.77
UK	4.81	1.87	2.24	-	-	5.04	2.4	2.5	-
Total of extract: 25.09					-	Total of extract: 36.44			
Total of TEN-34: 192.27					-	Total of TEN-155: 114.19			

Table 1. Inter-NRN traffic statistics on TEN-34 and TEN-155

The table emphasises that inter-NRN traffic has increased by almost 50% which outlines both the benefits and need for more bandwidth to support the European research community. It is expected that once all NRNs have migrated to TEN-155 the increase of inter-NRN traffic will be even more significant.

6. TEN-155 US service and connections to other networks

Prices of intercontinental circuits are falling making it more feasible to obtain high capacity connections from European cities to

Another add-on to the TEN-155 network is the connection of Israel and Cyprus, as a result of the EC approved Q-MED project. DANTE is co-ordinating the connection of these countries to TEN-155 which n the case of Israel will be implemented by a E3 circuit from London to Tel Aviv. The connection is engineered in such a way to offer to Israel the same services available to the NRNs taking part in the Quantum project.Planning of the connection to Cyprus is still in progress.

7. The Quantum Test Programme

As previously mentioned, the Quantum Test Programme (QTP) is a substantial component of the Quantum project. The main objectives of QTP are to carry out evaluation of advanced networking technology and network related technology on a dedicated test bed and migrate where possible the technology to the production service. The dedicated test bed is obtained by setting up a set of PVCs to create a VPN. QTP activities are carried out by a joint DANTE-TERENA task force, TF-TANT, which carries out evaluation of technologies that are also relevant to the TERENA working group on lower layer technologies (WG-LLT). The activities relevant to QTP are as follows:

- Multicasting (IP and ATM)
- IP QoS (diff-serv, RSVP, RSVP to ATM SVC mapping)
- IP over ATM
- ATM SVCs
- IPv6
- MPLS
- Route monitoring
- QoS and Flow based monitoring

At the time of writing this paper, the QTP activity has just started so there are no results published in this paper. However these will be available at the DANTE web site (http://www.dante.net/quantum/qtp)

Particular priority has been assigned to native IP multicast, in order to gain experience with PIM[4,5] and MBGP[6] and transfer this technology to the production network. Some NRNs already deploy PIM and MBGP in their own domain, but in a multi-domain environment there is still the need for experimentation before it can be deployed without the risk of disrupting production service. Therefore DANTE and QTP will commit efforts and resources to the experimentation of PIM and MBGP using workstations and GateD initially to gain experience with the concepts. The scheme of experimental set-up is on Fig. 5 below. In parallel, the procurement of test routers, of the same type as the production routers, to deploy in the dedicated VPN will be organised.

Testing of PIM and MBGP will then resume on this infrastructure. It is estimated that by the end of 1999 TEN-155 will support a native IP multicast backbone.

Currently multicast is supported on TEN-155 by deployment of a European Mbone, described in detail in [7] for TEN-34 and migrated to TEN-155 (http://www.dante.net/mbone). This Mbone is implemented by setting up DVMRP tunnels between TEN-155 PoP workstations and multicast capable routers in the NRNs.

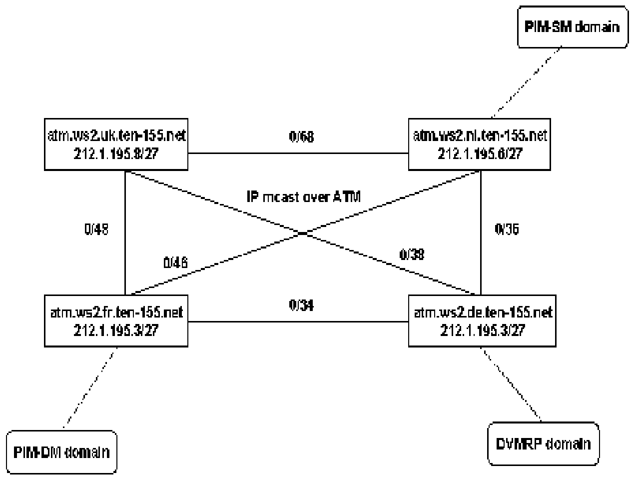

Fig.5 Experimental IP multicast setup

IP over ATM furthermore is viewed as high priority because the intention is to make use of the most suitable ATM traffic classes (ATCs) for carrying IP traffic and guaranteed traffic. These have been initially identified as SBR-3 with SCR=~0 and DBR, but also the use of SBR-3 with SCR0, SBR-2 and ABR needs to be thoroughly evaluated, in co-operation with the supplier of the ATM service (KPN).

In July 1999 there is a contractual obligation of the supplier of the service (KPN) to provide ATM signalling - but it is not yet defined how it is possible to make use of it in TEN-155, especially for the MBS service which would require end-to-end signalling, hence across management domains. Therefore the QTP will dedicate resources to the investigation of the deployment of SVCs.

ATM is not viewed as the only way to provide guaranteed capacity to users. In fact the networking community is considering both ATM and IP mechanisms. The performance and behaviour of ATM QoS is well known (if the most simple ATM technology is used), whilst the IP QoS / CoS mechanisms are still under development - however there is strong belief in their potential. Consequently the Quantum Test Programme will investigate the developments of diff-serv, RSVP and RSVP to ATM signalling mapping. One issue which will need immediate investigation is the co-existence of a fully meshed ATM backbone and the IP QoS mechanisms and techniques for their interoperability.

8. The TEN-155 Managed Bandwidth Service

The Managed Bandwidth Service (MBS) (http: / / www.dante.net / mbs) is one of the major components of the Quantum Project. It aims at enabling end-to-end guaranteed capacity between sets of hosts or networks across Europe in order to create VPNs with dedicated bandwidth in the support of pan-European research projects. The capability for establishing VCs / VPNs and closing them down at short notice or according to a predefined timetable is an essential component of the MBS.

Setting up operational and management procedures and tools is the main challenge of this acitivity, as from the technical point of view it is possible to create VPNs using a set of ATM VCs between the participating nodes. Before the MBS becomes an operational service, it will undergo a pilot or alpha test phase with selected users. These users have been identified within ERCIM[8], whilst the project to make use of the MBS in the alpha phase is the MECCANO[9] project. The aim of the alpha test phase is to identify the exact operational procedure requirements and support tools necessary. This phase is expected to last from January through March 1999. Once the alpha phase has been considered a success, a second phase (beta) will follow for two months, expanding the number of projects and National Research Networks involved.

The purpose of the beta phase is to prepare the service for production by verifying procedures and reducing time consumed in critical tasks. Production service is expected before June, 1999. The Quantum Test Programme will make use of the MBS to setup the VPNs necessary to carry out evaluation of networking technologies.

Conclusions

The Quantum project has been successful in deploying a pan-European network for research, providing a huge increase in capacity when compared to its predecessors and at a similar price. This is mainly due to developments in the European telecommunications market. From the technical point of view, TEN-155 has been designed to make efficient use of the available bandwidth and fair sharing of the bandwidth in situations of congestion. These targets have been met by combining different ATCs, and the results obtained in test laboratories confirmed the theoretical expectations. Rollout of the MBS and the work carried out within the Quantum Test Programme is expected to enhance even further the capabilities of the network, with the deployment of native multicast and developments in the area of the co-existence of differentiated and integrated services.

References

1.EuropaNet, http: / / www.dante.net / europanet.html
2.TEN-34, The European Information Superhighway for European R&D, Feb. 97-Dec.98 http: / / www.dante.net / ten-34
3.Roberto Sabatino, Traffic Accounting using Netflow and Cflowd, 4th International Symposium on Interworking in Ottawa,

Canada, July 1998, DANTE in Print N. 32, http://www.dante.net/pubs/dip/32/32.html
4.PIM-SM RFC2362 - Protocol Independent Multicast-Sparse Mode (PIM-SM): Protocol Specification
5.PIM-DM - draft draft-ietf-pim-v2-dm-01.txt, Protocol Independent Multicast Version 2 Dense Mode Specification;
6.MBGP - RC2283 - Multiprotocol Extensions for BGP-4
7.Jan Novak, Roberto Sabatino, The TEN-34 Mbone pilot, DANTE in Print N. 36, October 1998
http://www.dante.net/pubs/dip/36/36.html
8.ERCIM - European Research Consortium for Informatics and Mathematics.
http://www.ercim.org/
9.MECCANO - Multimedia Education and Conferencing Collaboration over ATM networks and Others
http://www-mice.cs.ucl.ac.uk/multimedia/projects/meccano/

Acknowledgements

The authors would like to thank the European Comission for the co-funding of the Quantum Project. Numerous technical experts from the European National Reserach Networks have provided advice on the engineering of the TEN-155 network, and also carried out physical work beyond any expectation for the deployment of the network. Members of TF-TANT, the task force that carries out the technical evaluation of networking technology in the Quantum Test Programme, who provide continuous input to the introduction of new services on the operational network.