#28

The Implementation of TEN-34

Michael Behringer

This paper was presented at JENC8, the 8th annual Joint European Networking Conference, Edinburgh, May 12-15 1997.

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 302 992
Fax: +44 1223 303 005
E-mail: dante@dante.org.uk

The Implementation of TEN-34

Michael Behringer

Abstract

In 1995 the TEN-34 project [1] was launched with the goal of providing an advanced high speed networking platform for the European Research community, by interconnecting the national research networks (NRNs). TEN-34 is carried out by a Consortium of European National Research Networks, with DANTE as Coordinating Partner. It took more than two years from the first meeting to the backbone finally becoming available. The reason for this is that standard PNO (public network operator) services in Europe could not fulfil the requirements of the R&D community in Europe, and to some extent they still cannot.

This paper describes the technical implementation of the TEN-34 backbone, which is shaped by a series of non-technical influences such as non-availability of required public services and cost factors. The backbone is a hybrid network consisting of ATM VPs and traditional leased lines. The design is explained in detail and the routing and other IP related topics are discussed.

Introduction: The TEN-34 Hurdles

The usual way of designing a new network is to assess the requirements of the users, to translate this into requirements for the network and to define how the network needs to be engineered to fulfil these requirements. Then one would buy or lease the components and plug it together. In the case of the pan-European TEN-34 network however the solution was determined by what is available in terms of components. Our initial plan was to build a network based on 34 Mbit / s leased lines. However, we found a great number of non-technical problems on the way. It is important to understand these problems to appreciate why the TEN-34 network was built the way it is.

Michael Behringer works as Senior Network Engineer for DANTE. His e-mail address is: Michael Behringer@dante.org.uk

The first problem is that in Europe international 34 Mbit / s lines are either not available at all, or they do not provide much economy of scale over 2 Mbit / s prices, making them far too expensive. Only in a few countries can international 34 Mbit / s half-circuits be leased easily, and then it still depends whether the PNO in the other country offers an equivalent and matching circuit, which is often not the case. We were reduced to using ATM VPs in some cases, which are new services at an international level. They therefore required a lot of attention from our side to ensure that the VPs would be configured consistently over the network.

Also the mere fact that the network has multiple suppliers leads to potential operational problems as well as difficulties of receiving satisfactory service from the suppliers, because most suppliers contribute only a minor part to the network.

The TEN-34 project is partly funded by the European Commission, which is currently the only way of financing this network at all, since it costs 40 MECU per year. But this contribution comes at the price of added delays and bureaucratic hurdles in the process. TEN-34 also currently has a planned service period of 18 months only, which is rather alarming given that it took two years to plan the network. Due to the short service period some significant compromises were made in terms of the configuration, as national preferences influenced the shape of the network to some real extent.

These non-technical problems resulted in a “bottom-up” engineering process, in which we were presented with a - very limited - set of technical options. But these options were not necessarily the ones we would have chosen. The engineering challenge was to check basic suitability of available components and to find ways of plugging together quite different technologies, rather than starting from the requirements point of view. This unnatural way of working explains why the topology sometimes looks rather odd.

Designing the Network

A. From Dreams to Reality

Ideally one would wish to have one transparent network using one basic transmission technology throughout. Our original goal was to start with a traditional network based on E3 leased lines, because this is a well known solution and would present the least problems for a fast implementation [2]. In parallel we would experiment with new technologies and bring those into the network at a later stage, once they have been proven to work reliably. This testing activity is being carried out by TERENA Task Force TEN [3]. This ideal set-up proved to be impossible due to the non-technical problems mentioned above.

From an early stage on it became clear that the TEN-34 network would have to consist of more than one logical part. There was an offer from Unisource to provide services in their countries (NL, CH, SE, ES) with interconnections to the UK and DE. This was designed as a standard high speed IP service. It was not possible due to political and market reasons to have one single PNO provide the whole network. Thus we ended up with a separation of the network in a "managed IP sub-network", provided by Unisource, and a collection of mostly ATM VPs and leased lines, provided by the PNOs in the other countries.

B. Calculating the Bandwidth Requirements

Given the cost of the lines involved it was important to make most efficient use of the resources. Therefore it was necessary to anticipate the bandwidth requirements on the backbone, to make sure that single backbone lines are not under-engineered, but also not too big to make efficient use of them.

The required bandwidths for the lines of the backbone were extrapolated from the traffic figures of existing backbones. Each NRN presented a table with its bandwidth requirements to other NRNs, based on existing traffic measurements. We combined these individual tables into one common table with country to country traffic requirements, where the bandwidth required between a pair of countries was averaged from the two estimates.

The next step was to create a routing policy for a given topology. With this routing policy we could then project the traffic estimates for each country to country connection onto the backbone. Adding the requirements up for each line gave us an impression of whether there is sufficient bandwidth available. In some cases we had to change the routing policy to "re-direct" traffic to less used lines.

Whilst this is the only projection we had, it is unclear how accurate it is. Firstly, the existing traffic figures are dependent on the existing infrastructure, as existing bottlenecks would limit traffic now, but would not necessarily exist in the new backbone. Therefore traffic figures that are strongly restrained by an existing bottleneck are not indicative for real traffic demands. Secondly, the traffic figures were sometimes old, as the gathering of reliable statistics is not possible in all cases. Thirdly, the figures we based this calculation on were projections of up to three years ahead. There are far too many influential factors, such as growth of the national network itself, to give precise estimations of demand. Also new applications that might come up in the future can change the pattern significantly.

All these "whiteboard" calculations and projections had a high degree of uncertainty in them. But given that these are the only figures available, we did not have an alternative. Other estimation methodologies such as extrapolating from the population of a country were dismissed as too speculative.

At the time of writing this paper the TEN-34 network is just starting operation, and it was not possible yet to judge whether our calculations were correct.

C. The Managed IP Subnetwork

One part of the TEN-34 network is a managed IP network, initially based on 34 Mbit/s lines. This network is provided by Unisource, and it covers the Unisource countries (NL, CH, SE, ES), with extensions to the UK and DE. It is built on E3 leased lines, with a topology as shown in figure 1. This network provides a full Internet Service with interconnections to other European backbones. The NRNs interface the network through defined access ports on an access router. As the full 34 Mbit/s are not available from the beginning, an access speed limitation is implemented between the Unisource access routers and the core routers of the network. This limitation is implemented by using ATM between the access and core routers, with a defined bandwidth

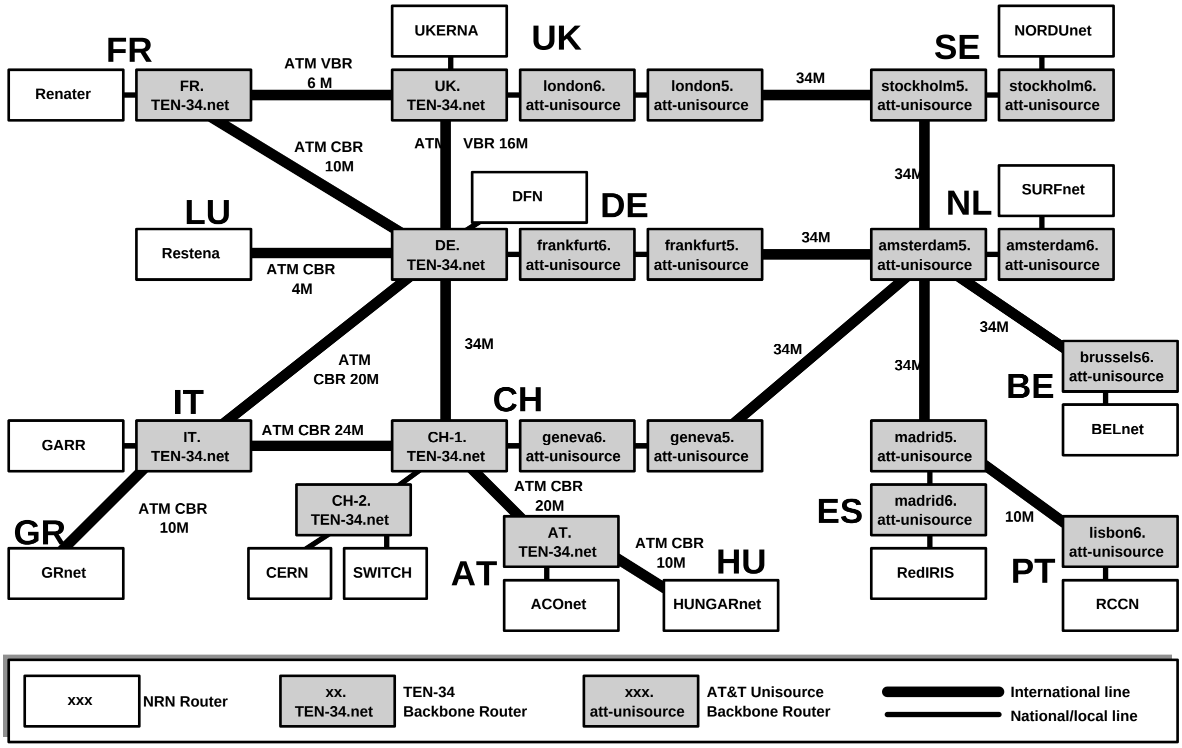

Figure 1: The IP topology of the TEN-34 network

(see table 1).

At a later stage it is envisaged to migrate the international E3 infrastructure of this IP network to an ATM based platform. However, the topology on the IP layer will not change, as the leased lines between the routers will be replaced with parallel ATM services. Therefore this migration will remain transparent for the TEN-34 user community.

D. The FUDI Side

Lacking an appropriate name of the other side of the network, we named it after the initials of the countries who originally started this side of the network: France, the UK, Germany (D) and Italy: FUDI. In the meantime more countries have joined this sub-network, but the name has so far outlived the changes.

The first choice for the FUDI part of the network was to have leased lines at E3 speed between those countries. However, either international E3 services were not available at all, or only at very high prices. Pushed for a solution, the PNOs proposed ATM, as both CBR services and VBR services.

CBR services are in their characteristics close to a leased line and therefore usable for IP services without problems. VBR provides a sustainable cell rate (SCR), which can be increased to a peak cell rate (PCR) for a short time. Whilst this sounds ideal for IP traffic, measurements of IP throughput over VBR services have actually shown that IP throughput can decrease below the SCR if the PCR is higher than the SCR [4, 5]. We therefore insisted on VBR services with SCR = PCR for the initial phase of TEN-34. From an IP perspective both, CBR and VBR with SCR = PCR are equally suitable.

In some cases we had a choice between CBR and VBR ATM services between countries. This choice was then mainly driven by cost factors. In some countries VBR was cheaper then CBR, in others it was the opposite. On one particular VP CBR was cheaper in one country, and VBR in the other country. In this case we chose the overall cheapest solution. ATM services were not available everywhere either, for example it was not possible to get ATM services between the UK and Italy.

ATM would in theory enable us to create a full mesh between routers, so that each country has a direct connection on IP level to each other country. We decided against a full or even partial mesh, and use ATM virtual circuits (VCs) only as a substitution for leased lines between routers, with the VC occupying the full bandwidth available on this VP. The reason for this decision is that with a full mesh between a number of countries the bandwidth out of each country must be pre-allocated to each other country, when CBR or VBR services are used. But this preallocation would only be useful if there is either very high bandwidth available - which was not the case -, or if more sophisticated ATM traffic classes such as ABR were available. Thus our approach was to have VCs only along the physical lines and allocate the full VP speed to this one connection.

As on this side of the network only "layer 2" services will be provided by the PNOs, it is the task of the TEN-34 NRNs to interconnect these lines and ATM VPs on the IP layer, and to provide a uniform IP service. For this reason we deployed dedicated TEN-34 routers to interconnect the international services and to interface the NRNs in the countries. This set of TEN-34 routers is managed by the TEN-34 NOC and provides a transparent IP service to the NRNs.

E. The Interface to the NRN and the Connection of the two Sub-Networks

The interface to the NRNs is slightly different on the two sides of the network. On the Unisource side the interface is provided by the Unisource access router, which is usually co-located with the NRN router that connects to it. These connections are usually over FDDI. On the FUDI side the interface is provided by the TEN-34 router, which is in all cases co-located with the connecting NRN router.

Country	Network	Access Speed
AT	ACOnet	10 Mbit/s
BE	BELnet	22 Mbit/s
CH	SWITCH, CERN	22 Mbit/s
DE	DFN	45 Mbit/s
ES	RedIRIS	22 Mbit/s
FR	Renater	16 Mbit/s
GR	GRnet	10 Mbit/s
HU	HUNGARnet	10 Mbit/s
IT	GARR	34 Mbit/s
LU	Restena	4 Mbit/s
NL	SURFnet	22 Mbit/s
PT	RCCN	10 Mbit/s
SE	NORDUnet	22 Mbit/s
UK	UKERNA	34 Mbit/s

Table 1: TEN-34 access speeds of NRNs[1]

[1] Not all NRNs were connected at the time of writing this paper. Not all details were final.

The two sub-networks, the FUDI network and the Unisource network, are co-located in three countries: UK, DE and CH. The connection between the two sub-networks is made by connecting the TEN-34 router in those countries to the Unisource access router, instead of the NRN router. The NRNs in those countries interface the TEN-34 network at the TEN-34 router. With this set-up the two sub-networks can be treated as one contiguous IP network, as from any country to any other country only backbone routers from either sub-network will be traversed, and no other NRN routers.

Some NRNs are in topologically advantageous positions, where significant backbone bandwidth is going to other countries, enabling them in theory to make full use of that bandwidth. However, this bandwidth is also used for transit traffic between other NRNs, so that an NRN accessing the network could compete with other NRNs using this country as transit. Therefore it was considered optimal to limit the access speed of each NRN to a defined amount, so that each NRN would get its fair share. This can then also be used for cost distribution. Table 1 shows the access speeds of the NRNs connecting to the TEN-34 backbone.

Obviously the backbone must be designed in such a way that there is sufficient bandwidth in all transit nodes to support both the NRN in a particular country and other NRNs' traffic passing the node. This bandwidth limitation was usually implemented by configuring an ATM VC at the appropriate speed. In some countries this was done between the TEN-34 router and the NRN router, in the Unisource network this ATM speed limitation is between the Unisource access router and the core router. In some countries there is no limitation necessary, as the access speed is equal to the bandwidth out of the country.

Implementing Routing

A. Routing Set-Up

A basic principle for the set-up of the TEN-34 network was to provide one contiguous network, so that the internal structure remains transparent to the NRNs connecting to the network. With the topology described above this could be achieved. On a routing level it was important not to have any NRN router in the path between two other NRNs, as the policy of the intermediate NRN might then influence the backbone routing. Therefore separate autonomous systems (AS)

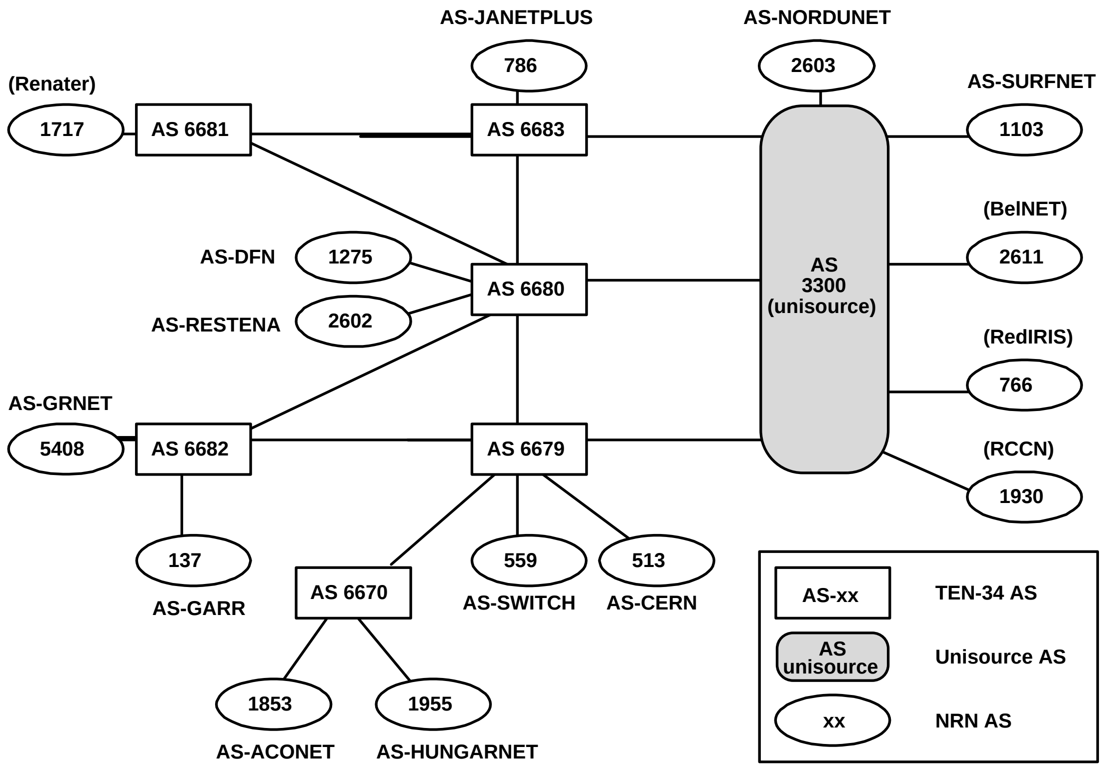

Figure 2: TEN-34 routing set-up.

are used for the TEN-34 routers. Figure 2 shows the routing set-up of the TEN-34 network in terms of autonomous systems.

The whole of the Unisource subnetwork resides in one single AS, so that routing between NRNs over the Unisource network is straightforward. On the FUDI side the plan was originally to use one AS for all TEN-34 routers as well. However, to make routing as flexible and manageable as possible it was decided to use one AS for each of the TEN-34 routers. This way traffic from a FUDI NRN to a NRN on the Unisource part can be routed differently - within limits - from other FUDI traffic to the same Unisource NRN, as there is not only one common routing policy on the FUDI network. This makes it possible to distribute the load more finely to the three interconnections of the sub-networks. This was considered important as these interconnections are very expensive.

The interface to the NRNs is the BGP-4 routing protocol. Internally, in the FUDI part, OSPF is being used to route between the TEN-34 routers. Unisource interfaces TEN-34 and the Unisource NRNs via BGP-4.

B. Filtering and Dampening

To achieve maximum stability of the network, and to ensure correct routing, several filtering poli

cies operate on the network. The TEN-34 network deploys two kind of filtering: Inbound (from an NRN) we operate AS filters and route filters, based on the RIPE DB, where all official routes must be registered. We also operate filters for the “private address space” according to RFC 1918 [6], so that bogus announcements do not propagate through the backbone. There are no plans for filtering on prefix length.

Route dampening is applied to all NRNs and peers on the network to ensure a stable routing set-up.

C. Interconnections to other Backbones and Acceptable Use

The global peering policy of the TEN-34 network was not precisely defined at the time of writing this paper. Interconnections with EuropaNET and Ebone are considered important especially for the transition phase of NRNs to TEN-34. Other arrangements will be considered on a case-by-case basis.

The Unisource network is being used by TEN-34 as well as commercial organisations. As the same physical network is being shared between TEN-34 and others it is basically impossible to separate traffic between those, so that interconnections of the Unisource network will also be usable for TEN-34. On the other hand Unisource commercial customers will get automatic access to the TEN-34 NRNs.

As there is public money involved in this network, there is an Acceptable Use Policy (AUP) for the network. Without going into the details of the definitions, the intention of this AUP is to prevent purely commercial traffic on the network. Obviously there will be traffic between R&D sites and commercial sites through the peerings with commercial networks. There is no way of limiting this due to the way IP works. But this kind of traffic is not considered commercial, as one side is an R&D institution.

The NRNs connecting to the TEN-34 network all have a national AUP, which is similar to the AUP of the TEN-34 network. Therefore the NRNs can be considered non-commercial in this context, and the only thing that needs to be prevented is transit traffic between commercial peers of the TEN-34 network. This kind of filtering can be done with simple AS filters, which makes the enforcement of this AUP is relatively straightforward.

D. US connectivity

Unfortunately the cost of the TEN-34 backbone is very high, so that it is economically not feasible for big NRNs to make use of US connectivity through the backbone. This is because the cost of traversing the backbone to reach a trans-Atlantic line are so high that it is cheaper to buy a direct line into the US. There is also the additional complication of total absence of US co-funding for trans-Atlantic lines, which makes this very expensive for Europeans.

The only cases where the consolidation of trans-Atlantic bandwidth is economically feasible is for smaller countries with demand for US bandwidth in the range of a few Mbit/s. For those countries trans-Atlantic connectivity is being procured. A 45 Mbit/s line is planned from Germany to New York, where Internet services will be purchased from a US provider. The connection point is ideally located in Germany, as this is a central location in the TEN-34 network. Technically this set-up is not much different from existing low-speed arrangements.

Operational Pitfalls

The TEN-34 network was not organised centrally by a call for tender for the whole service, but it was negotiated in small parts between NRNs and national PNOs. As a consequence there are now approximately 30 contracts in place for different parts of the network. The responsibility for those parts is also highly distributed, and there are a large number of boundaries between suppliers.
A contract for a network operations centre was awarded to UKERNA. This TEN-34 NOC is responsible for the day-to-day operations of the network. This includes configuring, monitoring, troubleshooting and reporting.

It was already a big challenge to bring this network up with so many contracts that need to be in line. On the technical side there were many differences in the initial drafts that needed to be brought in line. For example the calculation of a cell rate from the contractual Mbit/s values for the VPs was done differently in different countries. Some PNOs included the ATM headers in the calculation, others did not, and one PNO even calculated with a 20% safety margin. And these differences could be seen even on two ends of the same VP, which is not strengthening the view of a good co-ordination between PNOs.

Whilst these problems could be tackled for the installation of the network, now, that the network is operational, a lot of potential problems might arise from the fact that there are a large number of service boundaries, on each of which there can be misunderstandings about responsibilities. There is the additional risk of international ATM services being deployed. This is a new service for the PNOs, hence teething problems cannot be ruled out in the initial stages, especially in the operation of the service.

At the time of writing this paper no operational experience could be reported yet, as the network was just becoming available. The TEN-34 NRNs hope to keep control of the network by defining precise operational procedures between all parties involved. Whilst this is a difficult task given the complex set-up we are confident of providing a stable service.

What Next?

With the extremely short service period of the TEN-34 network of 18 months, the planning of the next phase had to start right after the TEN-34 network became operational. At the end of the TEN-34 project the traffic requirements will most likely exceed the capacity again. Therefore the next phase has to cater for 155 Mbit/s connections at least between the major networking countries. During the next months more NRNs, also from central and eastern Europe, are expected to join the TEN-34 project.

For the TEN-34 project itself a number of enhancements are planned for the near future, to advance the network to a “service” rather than a pure “bit forwarding mechanism”. DANTE’s plans include an Mbone service that is part of the network service, a number of statistical monitoring systems for network performance and utilisation, and a set of network management tools such as a “looking glass”, traceroute and ping servers, and servers for the multicast tools. Most of these features can be deployed fairly soon, as the technology is available today.

There are also plans to make more use of the ATM parts of the network to provide dedicated bandwidth for certain applications. These are currently still under investigation. At a later stage it is envisaged to provide direct access to ATM for the NRNs, and not only an IP service. This also needs further planning.

Summary

The way the TEN-34 network was being procured as a collection of partly small individual components was driven by supply: Public network services at the speeds required by TEN-34 are very difficult to obtain in Europe, and there is a wide gap in service offerings between various countries. Therefore the engineering of the network had to make use of whatever was available, rather than defining an ideal network and buying the components for it.

This “bottom-up” approach has complicated the design of the network both in terms of making it work and of operating this diverse environment. We believe to have solved this problem satisfactorily - we made it work, although it is not the ideal solution in technical and operational terms. The TEN-34 network now provides a well defined and stable service to the NRNs connecting to it, and for the first time there is a network that covers all western-European countries, with more in central and eastern Europe to come.

The TEN-34 project also caters for new services, and research for those is being carried out in parallel [4]. It is envisaged that more sophisticated networking services will be available on the TEN-34 network at a later stage.

References

[1] TEN-34 stands for “Trans-European Network Interconnect at 34 Mbit/s”; for more information see: http://www.dante.net/ten-34.

[2] Behringer, M. H.: “Towards a European High-Speed Backbone”; Proceedings INET’95, Page 569; 1995 (also available on-line as http://www.dante.net/pubs/dip/11/)

[3] TF-TEN stands for Task Force on Trans-European Networking; see http://www.dante.net/ten-34/tf-ten/

[4] Michael H. Behringer (ed): “Experiments for Advanced Backbone Services”; Proceedings JENC8, paper 332; 1997 (also on-line on http://www.dante.net/pubs/dip/29/)

[5] Reijs, Victor: “VBR and IP”; http://www.nic.surfnet.nl/surfnet/persons/reijs/sn4/pcr.htm

[6] RFC1918 (BC): Y. Rekhter, R. Moskowitz, D. Karrenberg, G. de Groot, E. Lear, “Address Allocation for Private Internets”, 02/29/1996.

Acronyms

ABR	available bit rate
ATM	asynchronous transfer mode
AUP	acceptable use policy
BGP	border gateway protocol (routing protocol)
CBR	continuous bit rate (ATM traffic class)
E3	transmission speed: 34 Mbit/s
FDDI	fiber distributed data interface (local area network standard)
FUDI	TEN-34 subnetwork, started by FR, UK, DE and IT
IP	Internet protocol
NOC	network operations centre
NRN	national research network
OSPF	open shortest path first (routing protocol)
PCR	peak cell rate (ATM traffic parameter)
PNO	public network operator
SCR	sustainable cell rate (ATM traffic parameter)
TEN-34	trans-European interconnect at 34 Mbit/s
TF-TEN	task force for trans-European networking
VBR	variable bit rate (ATM traffic class)
VC	virtual circuit
VP	virtual path