#10

Technical Options for a European High-Speed Backbone

Michael Behringer

This paper was presented by Michael Behringer at JENC, the annual conference of TERENA (RARE), held in in May 1995 in Tel Aviv, Israel.

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

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

Technical Options for a European High-Speed Backbone

Michael Behringer

Abstract

As many European national research networks have already deployed high-speed backbones nationally, there is a growing demand for a European backbone supporting line speeds of 34 Mbit/s and more. Unlike in the USA, in Europe it is not possible yet to buy high-speed connectivity as a standard product from the PNOs, so that the R&D community is forced to seek solutions on its own.

This paper is based on the results of a survey carried out by DANTE in October 1994, in which all national research networks in Europe and some international user groups were asked to provide information about their current and future requirements with respect to a pan- European infrastructure. The results lead to the definition of services needed on a high-speed backbone. Based on this analysis, technical options are discussed, which are suitable to meet these service requirements. A possible implementation strategy is outlined which takes into account the feasibility of new technologies such as ATM. Several technical possibilities to physically implement the backbone are discussed on their overall benefit to the users of the backbone.

Keywords: High-Speed Backbone, IP, ATM, Technical Implementation

1. Requirements

In October 1994, DANTE carried out a survey with the national research networks, enquiring about current and future demands on a pan-European high-speed backbone. The results of this are being used in planning the new infrastructure. Apart from non-technical topics, which are not discussed in this paper, the survey covered the view of line speeds required, theprotocols which ought to be supported, and the technology that should be used for a pan-European backbone. On the question of line speeds, there was unanimous agreement that at least 34 Mbit/s or multiples are required immediately. The same consensus could be seen on the need for future expansion to higher speeds, beginning with 155 Mbit/s. There was also general agreement that IP is absolutely needed as a service on the backbone. Most countries took the view that an ATM service is at least desirable in a future infrastructure, while some considered ATM being essential. CLNS was considered to be required by approximately half the countries. A service in this context means certain availability, throughput and maximum delay guarantees on a backbone that is continuously monitored. A pilot on the other hand does not provide those strong service level guarantees. On the question of which technology should be used for delivering the service, there was no general consensus. Five countries are open in respect to this issue, while two stated leased lines as the technology to be used. As there was a general agreement of the importance of an ATM service, it is obvious that ATM is also seen as the technology that should be used on the backbone.

The need for new applications has been specified as well. In this area there is a general tendency towards multimedia applications such as video conferencing. This kind of application of course puts special requirements on the provision of a high-speed backbone. Primarily this puts forward the need for constant bitrate services, which cannot be delivered on today's R&D infrastructure to the extent that is required.

A problem that arose during the work on the survey was the lack of clarity of the border between a technology and a service. For this article, a service is defined as the 'towards the national network visible interface between the pan-European backbone and the regional network'.

Michael Behringer is responsible for network planning and development in DANTE. His e-mail address is M.H.Behringer@dante.org.uk

The technology used is the implementation on the backbone that allows the delivery of the service; therefore, it is not necessarily visible to the regional (a regional is defined here as a national research network connecting to the backbone). In the case of ATM however, this might coincide, as it then represents the technology on the backbone as well as the service delivered to the regional network.

To summarise the result of the survey: the national research networks consider an IP service on line speeds of 34 Mbit / s as an immediate requirement. An ATM service would at least be desirable. Based on these results, the following analysis is investigating which technical options there are to deliver these services to the national research networks.

2. Constraints on Technical Options

The requirements given in the last sections specify the users+ needs on a high-speed backbone, where 'user' in this context means a user of the backbone, i.e. a national research network. Although not expressed explicitly, there are some other constraints on the procurement of the backbone, which are dealt with in this section.

2.1 Reliability of the Backbone

The pan-European backbone that is currently used by nearly all the R&D community in Europe is EuropaNET, which delivers a service of IP, CLNS and X.25 to the national networks. On this backbone there are certain service level guaranties, which ensure a very high quality of service on the backbone. A future high-speed backbone will have to meet similar quality parameters, in order to be accepted by the R&D community. Therefore a new high-speed backbone needs to run at a high service level. This demand constrains the technical options to the extent that the technology used must be robust and fault-tolerant.

2.2 Cost Effectiveness

In theory, all requirements of the R&D community can be met by installing enough hardware, to provide all services needed. There are of course financial constraints, which suggest that there are limits to the number of leased lines and switches / routers that can be set up.

2.3 Interconnection to Existing Low-Speed Backbones

The state of the development of high-speed infrastructure within the single European countries is very diverse. Apart from a small number of countries, which deploy a 34 Mbit / s network today there are a considerable number of countries which are not yet in a position to deploy such an infrastructure nationally for the next few years. Therefore a pan-European high-speed backbone cannot include all European countries from day one. This makes gateways to the existing lower speed backbones necessary, which need to be considered as an essential part of the new backbone as well.

2.4 Availability of PNO Services

This analysis has to be based on the PNO services currently available on a pan-European scale. At the moment, there are no higher layer services (such as e.g. SMDS) available, so this analysis will not take for example PNO ATM services into consideration. Yet one requirement of the new backbone is flexibility to the extent that a migration to PNO services which might become available in the future must be easily accomplishable.

3. Technical Options Analysis

The question to be answered in this section is which technology can deliver an IP service for line speeds of 34 Mbit / s and more immediately. It would be desirable if an ATM service could be provided on the same infrastructure now or possibly at a later stage.

3.1 IP over leased lines

The most straightforward implementation option with respect to delivering an IP service is the deployment of leased lines, with IP routers attached. This technology is well understood, and it is widely known. Although there are some performance issues with IP at higher speeds, these do either not have a big influence on the overall performance (as opposed to the performance of one single connection on the backbone) or they seem to be manageable. US network providers have proved that IP over 45 Mbit / s is working and performing, and higher speeds like 155 Mbit / s are started to being used for IP as well.

Multimedia applications are being deployed on the Internet today. For example Mbone provides users with multicast video and audio sessions over the Internet. Although the quality of especially the video transmission is far from perfect, it is the general attitude of the R&D community that a non-optimal service for many users is better than a perfect service available to only a few users. Thus the 'best effort' paradigm of IP seems to meet the demands of the community. Nevertheless with new applications available it will become necessary to deploy services with a constant bitrate. This is not possible with IP.

To summarise, IP over leased lines is able to provide the services needed for the immediate future. Of course it is not possible to interconnect ATM networks through an IP network, so that an ATM service cannot be provided on this infrastructure. As IP is a best effort protocol, constant bitrate services can not be delivered.

3.2 IP over ATM

Some European countries have already deployed ATM networks, and are starting to use them or use them for a production service. There is a clear demand to interconnect those national ATM networks on an international basis.

The ATM Forum [1] is working on a global standard on ATM. The User- Network Interface (UNI) has been specified already [ATM-UNI94], and currently the Forum is working on the Network-Network Interface (NNI). Unless those specifications have been made and are implemented by all vendors of ATM equipment, the full advantages of ATM can not be used internationally, as there are already different vendor's equipment in use today.

Even without the uniform specifications of the NNI, ATM can be used to the extent of Virtual Paths (VP) between two switches today. This basically amounts to a set of virtual leased lines between a set of end points. Therefore, a network could be set up on ATM basis, which would add some flexibility to the network with respect to easier changes of the virtual network topology. At the moment though it is not possible to make use of all features of ATM such as Switched Virtual Circuits (SVC), as the signalling system is vendor specific. For the provision of an IP service on top of ATM at this stage, there are no major benefits in using ATM compared to using leased lines. The set-up of a European backbone on ATM basis would still have the advantage of being able to provide ATM services at least to the extent of VP's to the regionals, or at a later stage full ATM service. An IP service can be provided on top of the ATM infrastructure, as has been proposed by the Internet community [RFC 1483, RFC 1577]. Other services could be offered on top of ATM as well. The major advantage though is the easy possibility to use the full ATM features like SVC's and use of multimedia applications as soon as they become available. This could then be done by updating the code on the switches already deployed.

The major issue with ATM today is the lack of experience in the deployment. Although there are ATM networks being deployed today, most of them are not considered as production services that have to meet certain service level guarantees. Generally speaking, the development of ATM has not come to an end yet, leaving a lot of unknowns for the moment [Laubach94]. This imposes a certain risk to the deployment of a European backbone. In the worst case the ATM Forum might not come to an agreement on a common signalling standard, with the possible consequence of difficult or impossible interoperability in the future. The standardisation process might also take longer than expected, possibly delaying the deployment of an ATM backbone.

3.3 IP over Other Protocols

IP services can be implemented on top of several other protocols. Examples include Frame Relay, SMDS or DQDB. Given that there are no such services available on a pan-European scale, a technology like SMDS would have to be implemented and maintained by the operator of the backbone. This imposes higher costs, higher protocol overhead and a higher fault risk. Therefore, without going into the technical details of those protocols, the assumption is made here that it is not feasible to provide IP over such a technique in this case. Additionally, none of those techniques has been considered to be required by the national research networks as such; due to that fact, there must be an internal benefit of providing IP over e.g. SMDS, where 'internal' means a benefit in operational, financial or technical terms to the operating entity of the backbone. The regionals would be able to see this benefit only indirectly, as they do not require the service underneath the IP layer.

An internal benefit for the usage of an underlying protocol could be seen if there were public PNO services on that technique available, so that for example the protocol overhead is hidden from the customer. As outlined above, at the moment there are no higher layer services than leased lines available on a pan-European scale, at least not for the speeds required here. For the operating entity of such a backbone to deploy an underlying protocol has several disadvantages. Firstly, the overhead of the underlying protocol is real bandwidth loss as compared to native IP on leased lines or buying a correspondent service from a PNO; secondly, another protocol adds complexity to the stack, which leads to other complications, like higher fault likelihood and possible interworking problems between the layers.

3.4 Conclusion

From the possible candidate technologies, two seem to be suitable for the provision of a European high-speed R&D backbone. Other services than IP and possibly ATM are not required by the research community in Europe, and will therefore not be considered further in this paper. A pure IP service is straightforward to implement. It is the easiest and most secure way to provide the backbone. IP over ATM at the moment (i.e. with PVC+s only) does not show any major advantages compared to native IP, but as this is a new technology it bears certain risks in operational aspects, which make it difficult to provide a service with QoS guarantees. However, as ATM is most likely going to be needed in the future, there should be a pilot network based on ATM technology.

To provide an IP service on top of another protocol like FR or SMDS imposes additional overhead and complexity which can not be justified unless such services were available as a PNO service, where such disadvantages would be hidden to the regionals. There are still the purely technical disadvantages, which have been discussed above, and which still apply, even if a PNO service were available.

4. Implementation of the Backbone

4.1 Implementation Strategy

An Implementation of a European R&D backbone should take into consideration all requirements specified by the users of the backbone, which are the national research networks in Europe, plus some international groups like the High-Energy Physicists. The approach of this paper is to highlight the principal possibilities of providing first of all a production high-speed IP network, as well as at the beginning an ATM pilot network, which can be upgraded to a production level as soon as seems feasible. Then the IP service could be offered on top of ATM, so that there is only one basic technology on the backbone. Therefore the project 'pan-European R&D high-speed backbone' should consist of the following phases:

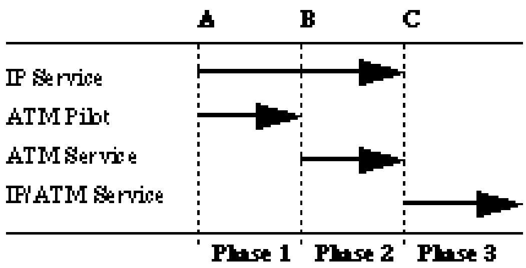

Phase 1 starts at the beginning of the project (A) with an IP service and a parallel ATM pilot network. This fulfils the immediate requirements, and furthermore provides the possibility to test ATM for its suitability for the provision of a service network. Of course the start of the ATM pilot network does not need to coincide exactly with the start of the IP service network. The outcome of the ATM trials might be that ATM is not suitable for the purpose needed here, or there might be general problems in the global provision of ATM services in a multivendor environment. In this case, the ATM pilot network can be ceased, and the backbone will consist of an IP service only. In case ATM proves useful to the community, the next step can be undertaken (B)

In phase 2 the ATM pilot network will be operational and suitable for the provision of a service, with the same or similar service level guarantees as the IP network (see above). The timing of (B) depends on external influences, such as the progress on the ATM development with respect to signalling, as well as on internal testing of suitability for the special requirements in this case. During this phase, testing of IP over ATM can be carried out with the possibility to fall back to the IP service network in case of problems. This provides an ideal testing environment for IP over ATM. At a certain point in time it will then become possible to provide the same service level guarantees for IP over ATM as for native IP (C). In phase 3 there is a backbone, based on ATM technology, which offers both ATM and IP as a full service to the regional networks. It will be provided on one physical backbone, which offers the best value for money.

In the case that during the ATM pilot phase or at a later stage ATM should, for whatever reason, not prove to be a suitable backbone technology, there is always the option to fall back to the native IP backbone. This ensures that investments into the backbone are not lost completely.

4.2 Technical Implementation Options

On the basis of the implementation strategy described above, there are several ways to implement a backbone that fulfils all the requirements. Important features of this strategy are the flexibility that is needed for testing a new technology, and financial aspects. In order for the ATM pilot network not to influence the service IP network, there should be sufficient independence between the different solutions.

4.2.1 Provision of two Physically Separated Backbones

This is the easiest way to ensure that the ATM pilot will not impact badly on the IP production network. It is easy to engineer initially, as the two parts do not influence one another. Once IP is being used on the ATM infrastructure as well, there might be routing problems as there will be several paths to one destination, but those are not problems related to ATM. As the techniques do not interfere in this model, this is a very secure option in terms of reliability of the IP service. A disadvantage is that at time (C), when the IP service is starting to be provided on the ATM network, one network has to be taken down, which might impose extra costs and effort, although the equipment can probably be reused.

The drawback of this solution is the costs. It is not sure at the moment, whether individual countries can afford a connection to a separate high-speed test network.

4.2.2 Lab Testing

The interworking of equipment from different vendors could also be tested in a local environment, which would bring down the costs as no leased lines are required. However, recent test of this kind have shown that it is very difficult to simulate real traffic within a laboratory. The major disadvantages are:

- The delay on leased lines has a big impact on the size of buffers in the switch. It is therefore very difficult to extrapolate results from a local environment to a production network.
- The ATM network will be used for IP traffic primarily, which is by its nature bursty. This shape of traffic is difficult to simulate in a test environment. But the impact of this on the buffer sizes in the switching equipment is a very important topic.

It is certainly possible to test the interoperability of different switches in a local environment to a certain extent. But to be able to provide a production service on that equipment requires testing in a real user environment, albeit at a later stage. Therefore lab testing can provide valuable input, but it will not be able to replace field tests completely.

4.2.3 Deployment of a Backbone with IP and ATM links

To save costs in the above model, it would be a possibility to create a network with native IP links and ATM links, which could back each other up in case of a failure on one of them. The assumption here is that neither of the two different backbones provides the resilience to be able to guarantee a very high service level, but taken together they can. Furthermore it is assumed that for real testing of ATM it is not sufficient just to use routers with ATM interfaces, but to deploy at least one ATM switch.

Of course this model can only be deployed when there is reasonable confidence that the ATM part of the network is sufficiently stable. Based on current experience with the configuration of ATM networks that use only PVCs, it seems possible to do that.

The savings in terms of lines are dependent on the topology deployed and on the level of risk one is willing to take. ATM networks only utilising PVCs seem to be reasonably stable, which makes it acceptable to take that risk.

4.2.4 Multiplexing of High-Speed Lines

Another possibility is to deploy one physical high-speed backbone, but to multiplex the links. A 34 Mbit/s link could be split into two 17 Mbit/s links for example, using normal TDM. This could be done by the operator of the network, but the high-speed link could also be obtained by the PNO in two smaller portions, which add up to the purchased bandwidth. The multiplexing of physical links offers several advantages in comparison with the other options:

- The pilot backbone can be kept virtually physically separate from the production network, providing maximum security for the production service.
- The bandwidth of 34 Mbit/s will not be used immediately by IP traffic. Therefore it is reasonable to split up some bandwidth for testing purposes. When the testing period is finished, the full bandwidth can be made available for service traffic.
- The costs are only to be paid for one infrastructure. This might lead to possibly high savings compared with the other solutions. The additional costs of TDM equipment should be comparatively small in relation to the overall benefit.
- Given that a TDM splitting of a line can be changed easily, there is big flexibility in this approach. So it would be possible to start with 10 Mbit/s for IP and 24 Mbit/s for the ATM pilot; when the IP service bandwidth does not suffice any longer, the splitting could be adjusted.

4.3 Implementation Summary

ATM is not sufficiently mature yet to provide an IP service from the beginning. Therefore the implementation has to be done in several phases, starting with a separate ATM pilot network, and aiming at providing IP over ATM, with both protocols being delivered to the regionals as services. Building two distinct backbones for service and testing is the best and most secure option, but expensive. Local test within laboratories cannot replace a real testing environment. The provision of one backbone with IP and ATM links which back up each other, is cheaper, but bears a slight risk. Another possibility is to provide one backbone and multiplex the lines, so that there are two distinct backbones, one for service, and one for testing. This is a cheap and flexible approach, which might be used if a second backbone for testing ATM turns out to be not feasible.

5. Summary

A survey on the requirements of national research networks in Europe with regard to a high-speed R&D backbone showed that there is an immediate need for a high-speed IP network. In the future, most countries consider ATM as necessary.

Technically, in the absence of PNO services, there are two ways to provide such a combined service. One is to start with an IP network, and develop ATM as a parallel infrastructure. The other possibility is to deploy an ATM network on which IP is carried as a service. As ATM does not seem to be sufficiently mature yet to provide such a service, the first option is proposed with the final goal to migrate to an ATM structure with IP on top.

Possibilities to provide the combined production and testing environment have been outlined. The best but most expensive option is to provide two physically distinct backbones, one for the production network and one for testing ATM. Other solutions were described which are slightly less optimal but could be chosen in case the first proposal might not be feasible.

References

[ATM-UNI94]
ATM Forum: 'User Network Interface (UNI) Specification Version 3.0', Prentice Hall, ISBN 0-13-225863-3, 12/1993

[Laubach94]
Mark Laubach: 'IP over ATM and the construction of High- Speed Subnet Backbones', Connexions Volume 8, No 7, July 1994

[RFC 1483]
J. Heinanen, 'Multiprotocol Encapsulation over ATM Adaptation Layer 5', 07/20/1993.

[RFC 1577]
M. Laubach, 'Classical IP and ARP over ATM', 01/20/1994.

Note

[1] For detailed information on the ATM Forum see: http://www.atmforum.com/