#11

Towards a European High-Speed Backbone

Michael Behringer

This paper was presented by Michael Behringer at INET, the annual conference of the Internet Society, in Honolulu, Hawaii, June 1995.

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

Towards a European High-Speed Backbone

Michael H. Behringer

Abstract

Unlike the USA, in Europe it is not yet possible to buy high-speed (i.e. 34 Mbit/s and above) international connectivity as a standard product from the PNOs. As about half of the Western European national research networks have already deployed high-speed pan-European backbones nationally, there is a growing demand for a backbone supporting line speeds of 34 Mbit/s and more on a European scale. This paper describes the requirements of the European R&E community for a high-speed backbone, as well as the plans for the development; it also explains the difficulties in getting such lines or services.

The first section of this paper outlines the technical issues. It describes the requirements of the national research networks in Europe for a high-speed European infrastructure, as indicated in a survey carried out by DANTE in October 1994. The available technical options for implementing the backbone are discussed and an overview of the proposed implementation strategy is given.

The second section covers non-technical issues in relation to the deployment of a high-speed pan-European backbone. The major problem there is the unavailability of high-speed lines from the PNOs, which, in most of the European countries, still have a monopoly position. The EuroCAIRN project was initiated by the national governments to find a solution for these problems. This section also gives an overview about the EuroCAIRN project, its objectives, strategy and goals.

Keywords:

High-Speed Backbone, IP, ATM, Technical Implementation, EuroCAIRN

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

1 Technical Issues

1.1 Requirements

In October 1994, DANTE carried out a survey of the national research networks, enquiring about current and future demands for a pan-European high-speed backbone [EuroCAIRN95]. These results are being used in planning new infrastructure. The survey covered the view of line speeds required, the protocols 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 on the absolute need for IP as a service[1] 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. While CLNS was considered to be required by approximately half the countries at the time of the survey, in April 1995 none of the regional networks (a regional is defined here as a national research network connecting to the backbone) subscribed to a CLNS service on the successor of EuropaNET.

On the question of which technology should be used for delivering the service, there was no 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 the technology to 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 puts special requirements on the provision of a high-speed backbone. Primarily this requires constant bit rate 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'. 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. In the case of ATM however, this might coincide, as it then represents the technology on the backbone as well as a possible service delivered to the regional network.

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

1.2 Constraints on Technical Options

The requirements given in the last sections specify the users' needs[2] for a high-speed backbone. Although not expressed explicitly, there are some other constraints on the procurement of the backbone, which are dealt with below.

- 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 [Møller93]. There are certain service level guaranties on this backbone 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 provide a high service level. This demand constrains the technical options to the extent that the technology used must be robust and fault-tolerant.

- 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 implies that there are limits to the number of leased lines and switches / routers that can be set up.

- Interconnection to Existing Low-Speed Backbones

The state of the development of high-speed infrastructure within the individual European countries varies widely. 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 an immediate position to deploy such an infrastructure nationally. 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.

- Availability of PNO Services

The analysis in this paper 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 cannot 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 should be easy to accomplish.

1.3 Technical Options Analysis

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

1.3.1 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 and 622 Mbit / s are starting 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 bit rate. This is not possible with IP.

To summarise, IP over leased lines is able to provide the services needed for the immediate future. However, 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 bit rate services can not be delivered.

1.3.2 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[3] is working on a global standard on ATM. The User-Network Interface (UNI) has already been specified [ATM-UNI94], and currently the Forum is working on the Network-Network Interface (NNI). Unless those specifications have been completed and are implemented by all vendors of ATM equipment, the full potential 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 amounts to a set of virtual leased lines between a set of end points. Therefore a network could be set up with ATM switching technology, which would add some flexibility to the network with respect to easier changes of the virtual network topology. At the moment, however, 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 VPs 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 main advantage though is the possibility using the full ATM features such as SVCs and use of multimedia applications easily as soon as they become available. This could then be done by updating the software on the switches already deployed.

The major issue with ATM today is the lack of experience in its 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 isn't yet finished, leaving a lot of unknowns for the moment [Laubach94]. This imposes a certain risk on 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.

1.3.3 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.

None of the national research networks see a need for these techniques. This is because there needs to be an intrinsic benefit for providing IP over e.g. SMDS. In this context 'intrinsic' means benefit in operational, financial or technical terms to the operating entity of the backbone. The national research networks would only be able to see the benefit indirectly, as they have no need of and cannot exploit the service underneath the IP layer.

Such an intrinsic benefit from 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.

1.3.4 Conclusion

Currently IP is provided primarily on top of leased lines. This is a well-known set-up, and it is certain to work, even for high line speeds. IP could also be implemented on top of ATM, but there are major concerns about the maturity of ATM for a production service, especially since a very high QoS is required. Other underlying protocols like Frame Relay or SMDS only seem to be suitable if such a service could be bought directly from the PNOs on a European level. As this is not the case today, this option does not seem feasible.

From the possible candidate technologies, two seem to be suitable for the provision of IP services on 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 PVCs 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 will most likely be needed in the future, there should be a pilot network based on ATM technology.

1.4 Implementation of the Backbone

1.4.1 Implementation Strategy

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 possibilities of providing first of all a production high-speed IP network, as well as 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:

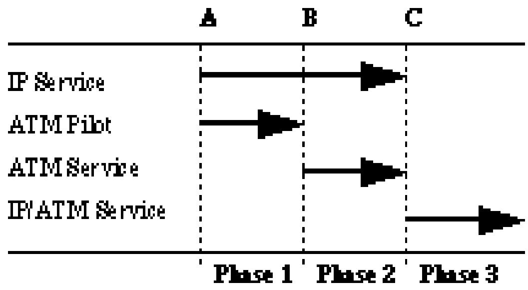

Phase 1 starts at the beginning of the project (A) with an IP service and a parallel ATM pilot network. This fulfils the immediate requirements, and furthermore provides the possibility to test the suitability of ATM for the provision of a service network. The start of the ATM pilot network does not 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 (B) can be undertaken.

In phase 2 the ATM pilot network would 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 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 of falling 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 serv-

ice 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.

1.4.2 Technical Implementation Options

On the basis of the implementation strategy described above, there are several ways to implement a backbone that fulfil 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.

Provision of two Physically Separated Backbones
This is the easiest way to ensure that the ATM pilot will not have a negative impact on the IP production network. It is easy to engineer initially, as the two parts do not influence each 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 ceased, which might impose extra costs and effort, although the equipment can probably be re-used.
The drawback of this solution is the costs. It is not any certain, whether individual countries can afford a connection to a separate high-speed test network.

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.

Deployment of a Backbone with IP and ATM links
To save costs in the above model, it would be possible 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 necessary 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.

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.

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.
- Costs only have 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 considerable 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. A possible problem here is that some ATM switches cannot operate on random line speeds, but only on certain pre-defined ones, which constrains the flexibility here.

1.4.3 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 because of the immaturity of ATM. 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. The aim however must be to work on the best environment possible, being two physically distinct backbones, one for production and a possibly smaller one for testing.

2 Non-Technical Issues

2.1 Availability of Lines

In the USA high-speed lines, i.e. lines with a bandwidth of more than 2 Mbps, are commonly available. There, the situation is demand-driven, and primarily limited by the technical feasibility of the transmission technology. This is possible because of the highly competitive market in long-distance infrastructure and telecommunications services. In the short-distance telecommunications market there are still monopolies in place to ensure that also areas with a lower profit margin are being served.

In most European countries PNOs still operate on a monopoly basis. There are only a few exceptions such as the United Kingdom and Finland, where the telecommunications market is liberalised. Thus in most European countries only one provider owns the telecommunications infrastructure and holds a licence to install new lines. This makes the provision of international leased lines through Europe complicated, since in most cases several different PNOs are involved in providing such a line.

The major concern of European PNOs with regard to providing high speed leased lines is the reselling of this bandwidth for carrying voice circuits. Up to bandwidths of 2 Mbps the cost per voice channel is sufficiently unattractive for reselling; 34 Mbps lines though can carry app. 1000 voice circuits, which makes reselling attractive if the cost are considerably less than 17 times the cost of a 2 Mbps circuit. Since international voice telephony is a highly profitable business, the PNOs try to keep possible competitors out of their country either by not selling the bandwidth required or not at reasonable rates. Their monopoly status enables them to do so.

The liberalisation of the European telecommunications market is planned for 1998. By then, the PNO monopolies will have to make the transition to a competitive environment. This change is already taking place in a few countries like the Netherlands, with the effect of lowering the prices for PNO services. The same effect was observed in the United Kingdom when British Telecom lost its status as a monopoly. But most of the countries are only slowly moving in that direction and prices are still very high, compared to the USA. Unless the telecommunications market is fully liberalised it is not certain if high-speed lines will become available under normal conditions.

2.2 Possible Solutions for the R&E community

The R&E community in Europe does not aim to resell high-speed bandwidth for voice transmission, but bandwidth is needed purely for data transmission. Although most of the PNOs are aware of this they are still not willing to deliver the required bandwidth on normal business terms to the R&E community. This is because they might then be forced to sell the same connectivity to all customers, including competitors.

If the liberalisation does not take effect considerably earlier than 1998, the R&E community in Europe will have to look for other means of obtaining high-speed bandwidth. A possibility would be to agree with PNOs not to resell this bandwidth, and to commit to only use it for research purposes. This would allow the PNOs to offer more favourable conditions since reselling for voice would not be possible. Public funding associated with this might lead to an Acceptable Use Policy (AUP) on the infrastructure, dedicating the lines to research and development traffic for the academic community.

In the USA the AUP model was used for the NSFnet backbone service, which created an artificial boundary on the otherwise seamless Internet and led to a complicated administrative overhead. As the market developed Internet services became generally available, these offered the sort of performance now sought in Europe. An AUP was therefore no longer necessary since there were at least three service providers. The cessation of this AUP service was generally appreciated, even though it is widely acknowledged that the AUP was needed in the beginning. Given the current difficulties in obtaining lines in Europe there might be no other possibility than the provision of another AUP service here, with the well-known problems seen on NSFnet and also on EuropaNET when the AUP was still in place there.

2.3 The EuroCAIRN Project

To strengthen the position of the European research community with respect to the PNOs and to provide a framework for organising the pooling of national funding resources, EuroCAIRN (European Co-operation for Academic and Industrial Research Networking) was launched in 1994 as a Eureka project. The goal of the EuroCAIRN project is to 'improve network facilities for European researchers'.

Members of the EuroCAIRN committee were primarily the ministries of research and education of most of the European countries. These bodies fund the international networking facilities in their countries. Given the difficulties in technical as well as political and financial terms such a initiative was deemed necessary.

The key elements of the project were split into three areas:

- Organisation: Work towards a common European policy with regard to R&D networking in the high-speed range. This required a collaboration between all European countries.
- Technical: Promote leading-edge technologies and support the upgrade of current technologies in European networking.
- Industry: Mobilise Europe's IT industry to collaborate with the research community. This extended specifically to the telecommunications market.

The main challenges the project faced were the availability of suitable telecommunications equipment and to get the funding for new infrastructure organised. Both issues need to be addressed on a European level rather than on a country by country basis.

DANTE was subcontracted by the EuroCAIRN committee to produce a recommendation on how to address these issues. DANTE was the logical partner for this activity since it was set up by the national research networks in Europe to organise their international networking. The report produced by DANTE included a broad survey of the requirements of the national networks, technical recommendations as to the suitability of different technologies and a roll out plan for the implementation of a high-speed backbone [EuroCAIRN95]. Financial matters and connectivity issues to the rest of the Internet were part of the survey as well. The general goal of the report was the immediate procurement of a European high-speed backbone.

2.4 Next Steps

The EuroCAIRN report prepared by DANTE represents the blueprint for the development of research networking in Europe. In parallel the European Commission issued a call for tender under its Fourth Framework Programme which invited proposals for the development of research networking very much along the lines of the EuroCAIRN report. The deadline for the submission of proposals was in March 1995. The proposals were being reviewed at the time this paper was submitted.

The Telematics framework provides the basis for resolving the financial issues. However, it is unclear when and where high-speed lines will be made available by the European PNOs. This is still the major issue for establishing a high-speed backbone for the European R&E community.

References

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

[EuroCAIRN95]
EuroCAIRN, European Research Information Highways, Volume 1 and 2 (Vol 2 = Trans-European High Performance Interconnect for Academic and Industrial Research Networks - A study prepared by DANTE for EuroCAIRN), 1995.

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

[Moeller93]
Svend Moeller Nielsen, 'EuropaNET - contemporary high speed networking', Computer Networks and ISDN Systems (Suppl. 1), September 1993

[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.

Notes

[1] For this document a service is defined as providing 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.

[2] 'User' in this context means a user of the backbone, i.e. a national research network. This paper does specifically not address the end users, but only the perspective of the regionals of the backbone.

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

Autor Information

Michael Behringer is Senior Network Engineer at DANTE, a company founded by the national research networks in Europe to organise international networking. He is working on network development and routing planning.