Project Number: RE 1009
Project Title: TEN-34

Deliverable D14.3
Summary Results of the Test Programme

Deliverable Type:	PU - Public
Contractual Date:	30 June 1998
Actual Date:	-
Work Package:	14 - Phase 2 Test Programme
Nature of Deliverable:	RE - Report

Authors:

Michael Behringer, DANTE, UK
Simon Leinen, SWITCH, CH

Abstract:

The TEN-34 project has carried out tests of new ATM WAN features and other advanced internetworking technologies for more than two years, in the framework of the TERENA Task Force TEN. The detailed results of this extensive experimentation programme have been published as deliverables. Now, towards the end of the TEN-34 network, a new European backbone network is being developed: TEN-155. This paper looks at the results of the TEN-34 test programme, and examines the tested technologies on their practical applicability in the design of a new generation network.

This paper does not make a final recommendation on most technology questions for the TEN-155 network. It outlines the applicability, advantages and disadvantages of various options, based on the experience of the TEN-34 test programme.

Keywords:

TEN-34, TF-TEN, TEN-155, QUANTUM, ATM experiments WAN networking, Internet backbone design

---

Executive Summary

The TEN-34 project included an extensive test programme in parallel to the deployment of the production network. This test programme, carried out by the TERENA Task Force TEN, investigated new networking technologies on their suitability for a European Research backbone such as TEN-34.

The results have shown that ATM technology and advanced IP services are becoming increasingly stable and reliable for production services. The use of static VCs and VPs has been reliable for some time, and new services such as SVCs in connection with dynamic ATM routing through PNNI are becoming generally usable, although with some restrictions. Call admission control and the bandwidth management of CBR/VBR SVCs are still problematic. There is a better understanding today of new traffic classes such as ABR and UBR, although also here certain restrictions apply, mainly in the area of interoperability. Also RSVP provides interesting new methods for guaranteeing bandwidth on IP level, although the mapping of RSVP parameters to ATM SVCs is not standardised yet.

In general the experiments have revealed interesting new possibilities for networking. Within one administrative domain and one vendor's equipment there are a number of new technologies that can now be used reliably. The results of these experiments can be used for the planning of the TEN-34 successor network TEN-155.

---

1. Introduction

The TEN-34 project contains a production network and a test component. The test programme was designed with the view of providing new services on the production network, once these are sufficiently stable and useful in an operational backbone network. Due to the immaturity of most of the new technologies tested (see D11.3 and D14.2 for details) it was not possible during the lifetime of TEN-34 to implement new services on the production backbone.

The test programme has however provided the necessary information and knowledge to implement new services once they can be considered mature and stable, and once they are required. The results of the test programme outline for each technology tested how it can be used, and how it can be implemented in a backbone network.

This deliverable summarises the results of the TEN-34 test programme with a view of applicability of the various technologies in a production backbone. Since the TEN-34 network will be closed down in the second half of 1998, the results of the test programme are described in the view of usability on the successor network, TEN-155.

2. The Approach to the Successor of TEN-34

Although none of the experiments of the TEN-34 test programme addressed directly the questions raised with regard to the general backbone design, the experience from the programme is invaluable for the initial design of the successor backbone.

In the following sections the various technologies examined in the test programme are taken into account to list the possible design criteria for the successor network, TEN-155, together with pros and cons of the options. The TEN-155 network is being designed to support two basic services to the National Research Networks in Europe:

- A best-efforts IP service, like the one provided on TEN-34.
- A Managed Bandwidth Service (MBS) for the provisioning of virtual private networks (VPNs).

From the work with various backbone technologies carried out in the test programme it is obvious that the underlying technology for a network supporting these services has to be ATM at the current state of development, to enable the MBS for the creation of virtual private networks (VPNs). New directions of providing VPNs, such as

IP differentiated services, are currently being developed, but are not at a maturity level yet for the provisioning of a managed bandwidth service.

2.1 Ways of Providing a Best Efforts IP service

During the TEN-34 test programme several ways of mapping IP traffic onto ATM have been examined. These divide into dynamic and more sophisticated technologies such as LAN Emulation or MPOA, and more static ones such as PVCs. During the test programme it was discovered that protocols like MPOA are more designed for the use in LANs and have potential problems in the WAN environment. These problems relate mainly to timing issues and resilience of the various servers required.

For the use in a backbone network such as TEN-34 it is furthermore not obvious what advantages the more sophisticated services would bring over a mainly static set-up with PVCs. The backbone topology is quite static compared to a LAN environment, and does not contain a large number of IP devices to interconnect. Thus one of the results of the test programme is that the use of the simpler and therefore more reliable static PVCs is to be preferred for backbone networks such as TEN-34.

One decision to be made in a backbone network is whether to provide resilience (i.e., re-routing) on the IP level only or also in the ATM layer. Re-routing on IP level has some disadvantages such as potential routing instabilities, which can be avoided if the ATM layer (or the SDH layer) automatically switches to a back-up path in case of line failures. The advantage of this is that the IP layer remains untouched by the outage if the ATM switch-over happens fast enough. The drawback is the added complexity in the ATM bandwidth management, which would need to take into account all back-up paths considered for bandwidth planning.

Most of today's ATM switches provide automatic re-routing of VCs. This facility relies on a routing protocol on ATM level, and signalling between the switches. As shown in the PNNI experiment, ATM routing works stable and reliable so that this mechanism can be considered in a production network, although the properties of proprietary switches' protocols should be investigated before the deployment.

Some ATM switches also offer "soft PVCs", which means that they appear to be PVCs from outside the ATM network, but are in fact based on a signalling protocol with a re-routing capability. This can be assumed reliable, although it has not directly been tested in the TEN-34 test programme. From the SVC experiments can be concluded that SVCs within the backbone network can also be deployed in a stable fashion, but there is no obvious advantage to soft PVCs.

The traffic class for the IP service can be either CBR or VBR (see the VBR experiment in D11.3), or UBR. More advanced traffic classes such as ABR have shown too many complications in the production environment and are therefore not yet reliable to be deployed (see experiment 8 in D14.2).

CBR and VBR services are easy to deploy, but have the disadvantage of firm bandwidth allocations on the backbone. This means that the bandwidth needs to be assigned to one VC and cannot easily be re-used. A full mesh between routers is practically not possible in this model since the average bandwidth per VC would be too small to allow for fluctuations in traffic between different destinations. Therefore a set-up with CBR or VBR VCs would likely result in a topology with one VC per link, or possibly a very small number of additional connections that pass more than one link. This model however is closest to the traditional leased-line IP network and thus very reliable and provides optimal bandwidth usage.

UBR on the other hand offers the advantage of not having to specify bandwidths for single VCs, making a full or partial mesh possible. The problem is that the routers on the edge do not see congestion any longer, and in the case of congestion cells will be dropped in the ATM cloud. This set-up only works if the ATM network discards on packet (AAL5 SDU) level, otherwise the performance has been shown to be unusable in the case of congestion.

Another advantage of the UBR mesh is that traffic flows do not have to be passed to each router on the way through the network, but will stay in the ATM network. This allows the bandwidth between the routers and the ATM network to be smaller than with the un-meshed CBR/VBR network, saving interface bandwidth.

The TEN-34 test programme has undertaken extensive research in the area of network management, which will be invaluable for future deployments. The use of SNMP for the management of ATM networks in particular has been researched extensively, and OAM based monitoring has also been investigated. See experiment 12 in D14.2 for more details.

2.2 Possibilities for a Managed Bandwidth Service

The simplest way of providing a MBS is with PVCs using either VPC or VCC. This way of working was used extensively within the TEN-34 test programme over the JAMES network (see D10.1 and experiment 1 of D11.3). The experiments on high-speed ATM performance (see D11.2 and D11.3) have proven that individual transcontinental host-to-host connections can be supported with bandwidths up to at least 30 Mbits/s with ATM PVCs, and current workstation hardware and TCP/IP implementations can fully exploit this bandwidth with some tuning.

From a technical point of view this PVC approach is easy to implement, but since most configuration work has to be done by hand, it is expensive in terms of operational overhead. VPNs based on VPs have the additional advantage of allowing tunneling of ATM services that might not be supported on the ATM backbone.

A more sophisticated way of operation is the usage of soft PVCs in the backbone network, as described in section 2.1 for the IP service. This eases the management in the backbone, but still requires additional interaction between the backbone network and the connected national research networks. As mentioned above, soft PVCs are generally stable enough to be considered in production networks, and they have the advantage of easier set-up over PVCs.

The ideal way to set up VPNs however would be through end-to-end signaling, so that the cost for set-up and maintenance of the VPN can be kept at a minimum. Whilst experiment 6 in D14.2 has shown that the signalling works reliably, there are still open issues in the call admission control for such networks. Therefore SVCs are only suitable for a network like TEN-34 today in a restricted environment, where the bandwidth consumption can be controlled by other means, e.g., administratively.

A further problem with SVCs is that bandwidth management for CBR and VBR SVCs is still not fully implemented on some equipment. Before the implementation of an SVC service more testing is required with a special emphasis on CBR and VBR SVCs. The use of UBR SVCs is possible, but since most users of a MBS require bandwidth guarantees, is not useful.

2.3 Other Technologies and Services

Several of the services and protocols investigated in the testing programme will not be deployed on TEN-155 from the start. In general, candidate techniques must meet several conditions in order to promise successful deployment: implementations that have proven themselves at least in a testbed setting, sufficient theoretical understanding and operational experience to make configuration and network management tractable, predictably limited impact on existing services, and tolerance against common wide-area communication problems such as line failures. Another frequent problem is that, although a protocol may in principle be mature for deployment, it hasn't been integrated into the software mainstream yet, and the test versions of the software that include the new functionality are not suitable for production use because of support issues or lack of other, more important, functions. All this has to be weighed against the potential benefits to the connected networks. In many cases, new backbone services would only benefit those NRNs that can provide an equivalent service to their respective connectees, so there is also a chicken-and-egg problem for incentive towards new services.

Classical IP with ARP (Address Resolution Protocol) over ATM could have been used between routers in the IP backbone. It was decided not to do so because the backbone is simple enough to make a statically configured address mapping quite feasible, and an address resolution protocol would introduce new potential sources of configuration errors, interoperability problems, and single points of failure without providing significant functionality that cannot be obtained otherwise.

NHRP (Next-Hop Resolution Protocol) is an alternative that distributes address resolution over the network, thus eliminating single points of failure. However, it seems too complex for such a simple topology.

PNNI (Private Network-Network Interface), if supported by the equipment used in TEN-155, might be used internally to support ATM routing for Soft PVCs even in the initial configuration of the IP backbone. However, it is not expected that NRNs will be able to connect their own ATM networks to TEN-155 via PNNI in the early phase. Besides the policy issues mentioned above, the interoperation of different vendors' PNNI implementations are still seen as somewhat problematic.

MPLS (Multi-Protocol Label Switching) is an IETF effort to standardise label-based switching methods for IP networks, of which Cisco's Tag Switching is an example that we have investigated in the framework of the testing programme. MPLS represents an attractive way to leverage ATM's cell-switching capabilities for an IP backbone, while not being restricted to ATM as a lower layer. It offers interesting possibilities for traffic engineering and IP VPNs. However, MPLS is still in the early stages of standardisation, and the use of a proprietary scheme such as Tag Switching would severely limit the choice of platforms for the network. Also, the management of the ATM and IP parts of the TEN-155 would have to be tightly coupled if label-based IP switching were used, which would impose additional constraints on the division of labour in the set-up and operation of the network.

3. Potential New Pilot Services

There are a number of new potential services, which can also partly be based on the results of the TEN-34 test programme. Whilst all of them require some degree of testing and tuning, the basic principles are well understood and believed to work. Also it is anticipated that these services are useful on a pan-European research network, such as TEN-34 or its successor network TEN-155.

3.1 Native Multicast

The experiment 5 in D14.2 gives details about the ATM point-to-multipoint facility. This is a relatively new technology, and will require some more work before it can be used in a production environment. Especially the mapping between IP based multicast and ATM point-to-multipoint still presents some open issues.

As an interim step, the current Mbone infrastructure should be moved to a native service. This involves the installation of new IOS versions on routers, and the migration of the current tunneled Mbone infrastructure to the backbone.

3.2 Routing Monitoring Tools

With the Internet growing and the dynamic nature of it there is an increasing demand for monitoring of routing information, to detect abnormalities. These tools should monitor route advertisements and withdrawals and run some sanity check on the data. Abnormal behaviour such as multiple withdrawals of the same route, route flaps and similar pathological behaviour should be brought to the attention of the network management centre.

3.3 A Native IPv6 Backbone

Currently the protocols for IPv6 are being standardised. To pilot a native IPv6 network, this needs to be trialled on a test infrastructure. Tests should include the new routing protocols, and applications. The goal is to provide a native IPv6 backbone, to which national IPv6 test beds can connect. Routing should be hierarchical, i.e., use an exterior gateway protocol such as BGP 4+.

3.4 SVC Pilot Service

Building on the experience of the TEN-34 tests on signalling, an SVC pilot service should be developed, allowing SVCs to work end to end. By making the pilot service available only to pilot users, and by limiting the bandwidth by other means (e.g., tunneling over a limited VP) the current shortcomings of SVCs described above can be overcome. A similar test bed is currently in use in France.

3.5 Extended Network Monitoring

The TEN-34 experiments have provided insight into the monitoring of ATM equipment. For the production network further traffic monitoring tools on IP level are required. There has been some work on flow based accounting methods, but this needs to be further refined.

3.6 IPv4 Differentiated Services

The IETF is currently standardising the way of making use of the type of service (ToS) field in IPv4 datagrams. This can be used to differentiate between different classes of service (CoS) on IP backbones. A model for a European backbone needs to be developed, and implemented on a European backbone such as TEN-34/155.

4. Additional Experiments

In some other areas, more experimental work is needed to assess the potential benefits and deployment scenarios for TEN-155 or its successors.

4.1 RSVP (Resource ReSerVation Protocol)

RSVP is the IETF standard signaling protocol for ``microflow'' (i.e. per individual transport-layer connection) Quality of Service guarantees in the Internet. End-system support for RSVP is expected to become quite widespread in the near future, and router implementations become more available and more mature, too. It is widely recognised, however, that RSVP has some open problems which currently impede its deployment on the greater Internet. The problems of *policy-based admission control*, and in particular of service-level agreements (and their enforcement) that would permit providers to accept reservations from their peers, are still largely unsolved.

Future research is required to:

- define an end-to-end policy model for networks such as TEN-155, the participating NRNs, and their connected organisations
- investigate implementations of lower-layer mappings of RSVP/IntServ for ATM-based networks
- follow the integration of RSVP with emerging lower-layer techniques such as MPLS and Differentiated Services
- measure router performance with RSVP signaling and resource management at high best-effort and

reserved traffic loads.

4.2 MPLS (Multi-Protocol Label Switching)

MPLS is currently being standardised by the IETF (Internet Engineering Task Force) should be further investigated for applicability in a backbone context. In particular, work needs to be done to

- track standardisation efforts
- validate multi-vendor interoperability
- test performance of implementations under load
- study interaction with emerging techniques such as Integrated Services/Differentiated Service, as well as coexistence with a native ATM service with QoS.

4.3 IP VPN Provision and Firm QoS

Virtual Private Networks (VPNs) can be seen as a logical extension of a Managed Bandwidth Service, and thus as an interesting potential service of TEN-155 and its successors. IP VPNs are an area in which fast progress is made, and several techniques to realise them are either already available or being developed. This includes techniques based on IP tunneling, RSVP (Integrated Services), Differentiated Service, and MPLS. The available approaches should be investigated for their respective potential in terms of e.g. quality-of-service assurance and scaling. Some interesting techniques should then be implemented in test setups, to help understanding the configuration and operational issues.

4.4 Managed Multicast Service

The goal of this work is to propose network mechanisms in support of real-time applications involving multiple parties. Examples include videoconferencing and other synchronous Computer-Supported Cooperative Work (CSCW) applications. The network requirements of such applications will have to be determined in a collaborative effort with experts from the application area, such as the MECCANO group. Available lower-layer mechanisms such as a dedicated and tightly controlled IP multicast infrastructure, RSVP/Integrated Services, Differentiated Services, ATM SVCs, or MPLS traffic engineering, can then be evaluated for their suitability for those applications. Application designers could benefit from experiments at the network layer by having access to a high-performance infrastructure, and by providing feedback, could improve the understanding of their requirements.

5. Conclusions and Outlook

The work conducted by the TERENA Task-Force TEN on the TEN-34 experiments is an extremely useful source of information for the development of networking services for future backbone networks. This work provided valuable input in the design of TEN-155, the successor network of TEN-34.

Moreover, the programme fostered collaboration between researchers and network engineers from about a dozen countries. Besides allowing them to conduct experiments that most countries would have had difficulties to set up on their own, it also created a vivid exchange of ideas, technical understanding, and operational experience.

Some of the techniques that could not be introduced on a European backbone have found or will find their way into National Research Networks or campus networks. Experience with such installations is then efficiently disseminated throughout the participating organisations via task force meetings and mailing list discussions. This cross-fertilisation has become an important driver for technical progress in the European R&E network community.

Most of the new technical features tested in this programme are becoming reliable and stable, and can be used at least under some constraints as pilot services. For the general case, which includes multi-vendor interoperability and end-to-end functionality there are still some restrictions with regard to the applicability. Generally, the step from a test environment to a production network has been shown to be very difficult.

The test programme has also shown that the way of testing new features carried out by the TEN-34 community through a European ATM overlay network proved to be efficient and produced applicable results for a production network, although the tunneling of VCs over an ATM VP network introduced additional complications.

Continuation of the experimental work done by the task force is not only necessary for the development of new services on TEN-155, but also in the interest of the national R&E networking communities in Europe. We believe that a similar approach should be taken towards future testing activities within the TEN-155 project, and that the testbed facilities offered by the new backbone will allow for more interesting experiments.

The recent past has brought a tremendous acceleration of development in the transmission and internetworking areas. New technologies such as packet over SDH, Wavelength-Division Multiplexing (WDM), or all-optical networking will probably be deployed in high-speed wide-area networks over the coming years, and schemes like xDSL (Digital Subscriber Line), Cable Modems or Universal Mobile Telecommunications Services (UMTS) hold some promise for solving the ``last mile'' problem.

The European R&E networking community will certainly have to deal with these new technologies in the medium term, and it would be very helpful to have similarly fertile international cooperation. Whether these activities are best pursued in the same framework is an open issue, however. The challenge for the future of this work is to remain focused while being able to integrate new developments, and to maintain the productivity of the social ``network'' established during TF-TEN's work while still being open to contributions from new participants.

Annex A: Glossary

```
IETF     Internet Engineering Task Force (http://www.ietf.org/)
         The Internet protocol standardization body

IP       Internet Protocol

JAMES    Joint ATM Experiment on European Services
         Experimental ATM network jointly operated by 18 European PNOs

MBS      Managed Bandwidth Service; also:
         Maximum Burst Size (ATM Forum: traffic parameter)

MPLS     Multi-Protocol Label Switching

NHRP     Next Hop Resolution Protocol
         (ftp://ds.internic.net/internet-drafts/draft-ietf-rolc-nhrp-11.txt)

NRN      National Research Network

OAM      Operations And Maintenance
```

P-NNI	Private Network to Network Interface
PNO	Public Network Operator
PVC	Permanent Virtual Circuit
RSVP	Resource ReSerVation Protocol Version 1 Functional Specification - Internet draft; http://www.internic.net/internet-drafts/draft-ietf-rsvp-spec-12.txt
SNMP	Simple Network Management Protocol
SVC	Switched Virtual Circuit
VBR	Variable BitRate (ATM Forum: traffic class)
VC	Virtual Circuit, or Virtual Channel
VCC	Virtual Channel Connection
VP	Virtual Path
VPC	Virtual Path Connection
VPN	Virtual Private Network