#29

Experiments for Advanced Backbone Services

Michael Behringer

This paper was compiled for, and 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

Experiments for Advanced Backbone Services

Michael H. Behringer

Abstract

The TEN-34 network delivers today stable high-speed networking services to the European R&D community. At the same time research into more advanced backbone technologies is being carried out with the goal of making more sophisticated services such as bandwidth on demand available on the TEN-34 network. This paper describes the first phase of the experiments on advanced backbone technologies.

A set of 10 experiments was defined and carried out for this purpose by the TERENA Task Force TEN. These experiments deal mostly with aspects of ATM, such as signalling, but also with new IP related technologies such as RSVP. The results were partly reassuring, and partly quite disappointing. Overall these experiments have shown that more time is needed to bring the advanced features of ATM to a mature state.

This paper contains after a short introduction into the goals of the project the results of each of the experiments, presented by the leader(s) of the respective experiment in the TF-TEN. In a summary an attempt is made to assert the current status of backbone technology with regard to new services.

Introduction

When the TEN-34 project [0.1] was launched in 1995, it was recognised from the beginning that more advanced features will be needed on a European R&D network than there are available today. At the same time higher bandwidths were required urgently, so that the TEN-34 network was initially set up with higher speeds, but well known technology. Therefore an experimentation programme was launched in parallel to the installation of the production network. The goal of these experiments was to trial new technologies, and test them on their suitability for production services. The experimentation programme was then carried out by TERENA Task Force TEN [0.2], where mostly ATM experts from across Europe came together for these tests.

As a platform for testing we were using the JAMES network [0.3]. This network provided us with a set of mostly CBR ATM services across Europe, and linked the test sites of the TF-TEN members together. This set-up would be very similar to a set-up over the TEN-34 network at a later stage [0.4, 0.5].

The experiments described here cover a wide range of aspect of ATM, from network management to signalling, as well as performance related issues and in depth examinations of ATM technology related issues such as cell delay variation. Unfortunately the results are far from encouraging. In several areas of research we came to the conclusion that too many pieces of the jigsaw are still missing. Especially the use of SVCs does not seem to be feasible over a WAN yet.

Quite a few buzzwords of ATM, such as ABR and PNNI do not appear in this paper. The reason is not that we do not consider them, but implementations of them were either not available, or we didn't have time to do experiments on them. The second phase of the TF-TEN tests, which starts in May 1997, will cover these areas, and continue some of the tests described here.

Results of the Experiments

In the following sections, the results of each of the experiments carried out by TF-TEN will be described in detail. More information on all experiments can be found on the TF-TEN web page [0.2]. Each section lists the leader of the experiment. However, there were many more people from the task force involved in the work of each experiment, partly to a significant extent, and partly in the writing of these results. The full list of people involved in the experiments can be found at the end of the paper.

Michael Behringer works as Senior Network Engineer for DANTE. He is Chairman of the Task Force TEN and compiled this paper. His e-mail address is Michael.Behringer@dante.org.uk

1. TCP / UDP Performance over ATM Constant Bit Rate

Mauro Campanella <campanella@mi.infn.it>
Tiziana Ferrari <ferrari@infn.it>
Vegard Engen <vegard.engen@uninett.no>
Celestino Tomas <ctomas@chico.rediris.es>
Magnus Danielson <magda@it.kth.se>

Introduction

The wide area ATM infrastructure operated by JAMES offered the opportunity to analyse the impact of the TCP / IP flow control mechanism on the performance of applications on long baseline, high-speed ATM links. The efficiency of the windowing flow control mechanism was tested varying the setting of the socket options which directly size the dimension of the window: the send socket buffer size and the receive socket buffer size. Also the impact of the application message size, i.e. of the amount of data written in the kernel memory through a single system call write(), on the throughput was measured.

1.2 Set-up

All the tests were done by generating a real data stream between two or more workstations equipped with ATM interfaces connected by CBR circuits. Different and complex stream topologies were configured in order to stress the switches and to analyse the TCP / IP flow control efficiency. The public domain benchmarking application Netperf developed at Hewlett Packard was used. The experiment consists of a single phase during July and august 1996, divided into two test sessions, each run on a different network topology configuration and by a different set of partners. The VPs used in the tests are:

- Italy-Sweden: on a 56600 cell per second CBR circuit (24 Mb / s) of 43 ms round trip time
- Norway-Spain: on a 36000 cells per second CBR circuit (13.8 Mb / s) of 63 ms round trip time

For more details on the hardware configuration see reference [0.6].

1.3 Results

Tests gave a straightforward proof of the round trip time impact on the achievable throughput of a one-way TCP / IP connection over an ATM CBR VP. The window sizes should be large enough to allow the sender to generate one packet and receive back the corresponding acknowledgement without stopping the sending process in the meanwhile.

A rough estimate of the lower bound of the window size necessary to prevent the stop-and-wait syndrome gives for the link:

- Norway-Spain: Window = (63 msec * 13.824 Mbps) / 8) or about 109 Kbytes
- Italy-Sweden: Window = (48 msec * 24 Mbps) / 8 or about 120 Kbytes

These window sizes are larger than the allowed upper limit (64 KB) of most operating systems, then the window scaling option, which permits larger window sizes, must be implemented in the operating system. Tests show (see figure 1.1) that with the proper operating system set-up and applying a patch to allow for the window scaling option, the total bandwidth reserved on the CBR VP on the JAMES infrastructure is available to an application running on top of TCP / IP.

In figure 1.1 the send socket buffer (ssb) and the receive socket buffer (rsb) and the message size are variable with ssb=rsb=message, the maximum throughput is achieved when the parameters sizes are about 120 KB, according to our rough estimate. Up to that value the throughput increases linearly.

The shape of the function strictly depends on the operating system running on the sending machine: if it runs Solaris 2.5, the throughput increases regularly up to the maximum value, than it stays constant.

On the opposite, in the test session IT-SE, where the standard versions of Solaris 2.5 and IRIX 5.3 were used, which feature a maximum buffer size of 64 KB, the maximum throughput achieved by one connection was only 8.5 Mbps, which is 35% of the available bandwidth.

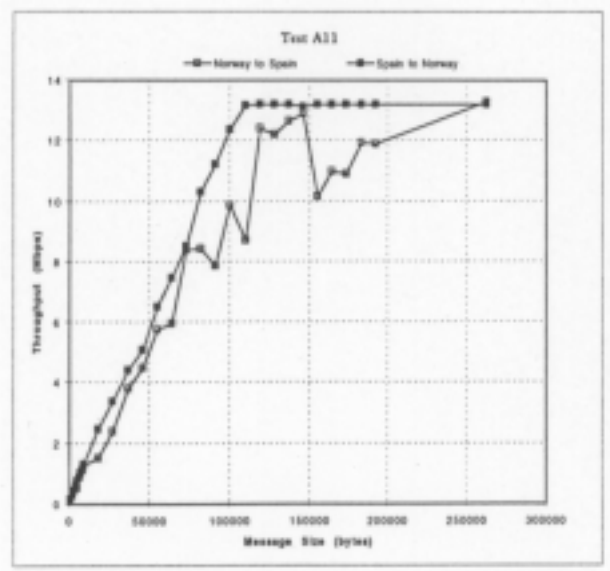

Figure 1.1: Test session Norway-Spain: Throughput for a one-way TCP connection with variable local ssb/rsb sizes, remote ssb/rsb sizes and message size (ssb = rsb = msg)

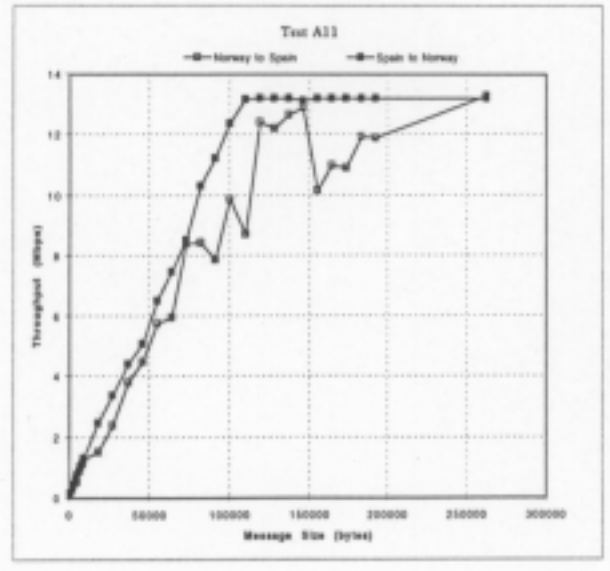

Figure 1.2: Test session Italy-Sweden: Aggregated throughput for a bunch of one-way TCP connections between 1 pair of end-nodes

Throughput for a one-way TCP connection with aggregated throughput for a bunch of one-way variable local ssb / rsb sizes, remote ssb / rsb sizes TCP connections between 1 pair of end-nodes and message size (ssb = rsb = msg).

In addition it can be stated that the optimal "ssb and rsb sizes combination'' does not exist, since it strongly relies on the TCP / IP stack implemented in each operating systems.

In any case, as we could expect, a symmetrical configuration, where both sizes are set to the maximum allowed value, makes the throughput as high and stable as possible.

With VP bandwidth in the range [0..30] Mbps, the size of the message does not impact the throughput at all. In fact, even with messages smaller than 10 byte, the CPU power of the sending host is still enough to guarantee the maximum throughput. A small message size makes the application generate an higher number of system calls, that is, more software interrupts and as a consequence some overhead for their management. If the amount of CPU cycles used by the sending process is not high — this is the case if the VP bandwidth is "low'' — this additional overhead is negligible.

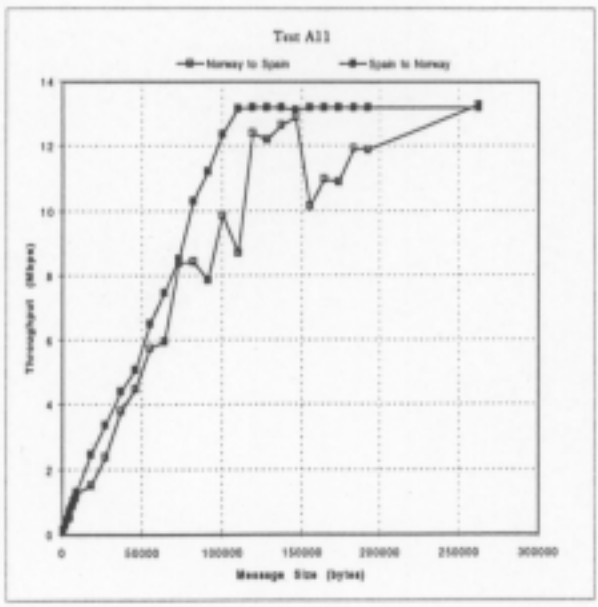

Figure 1.3: Test session Norway-Spain: Test of aggregate throughput for a bunch of two-way TCP connections between two end-nodes

When hosts have a limited TCP window size the global bandwidth utilisation can increase only if more TCP connections run in parallel (see figure 1.2).The improvement of the performance with more concurrent TCP connections is a positive result, because this model is much more similar to the real Internet traffic patterns, in which typically many users share the same circuit

When the traffic is not on-way, but full-duplex, i.e. it is generated by two data streams in both VP

directions, the aggregate throughput increases, but the maximum value measured is still lower than the total amount of bandwidth allocated (i.e. VP_capacity * 2) as shown in figure 1.3. When concurrent connections are activated, the aggregated throughput is about 75% of the maximum achievable throughput, in particular, 35.7 Mbps on the IT-SE VP with 24 Mbps bandwidth in each direction, and 20 Mbps on the NO-SP VP with about 13.8 Mbps, again, in both ways.

As far as UDP / IP is concerned, tests show that for appropriate datagram sizes almost the total available capacity of the VP can be used to successfully transmit UDP datagrams. In the traffic patterns tested the cell drop rate has no impact on the throughput measured for UDP streams. Figure 1.4 shows the throughput function shape for different values of Peak Cell Rate (PCR) assigned to the CBR VC when the size of the datagram increases. The interesting range of datagram size is [1..9152] bytes. If the UDP datagram is longer, it does not fit in one ATM MTU (Maximum Transfer Unit) any longer and because of the No_fragment option enabled, the receiver can't assemble the original packet.

As last remark, during all the tests the ATM service available in the JAMES infrastructure was good, continuous and reliable.

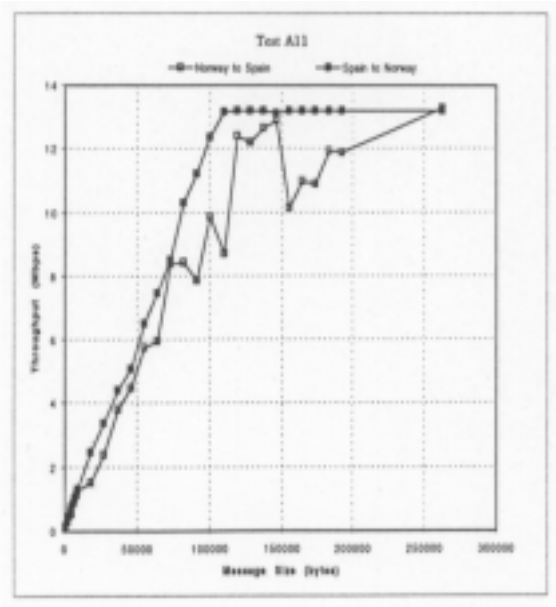

Figure 1.4: Relationship between variable peak cell rate (PCR) configured for the ATM VC and throughput achieved by a single one-way TCP connection on top of the VC itself

2. SVC in the WAN

Christoph Graf <Christoph.Graf@dante.org.uk>

2.1 Main Goals

- Gaining experience in interconnecting ATM end systems with SVCs in a multivendor environment.
- Prove the suitability of SVC over long distance PVPs.
- Providing an infrastructure for further tests based on SVCs and JAMES services, namely NHRP and ATMARP.

2.2 LAN vs. WAN

SVCs are being used successfully in local area networks already. Extension to the WAN is not easily possible as the LAN and WAN environments differ in some key aspects.

Bandwidth in the LAN is relatively abundant and cheap, while it is scarce and expensive in the WAN. The ATM services used today in the LAN and WAN differ for that reason, even though the same physical infrastructure is used to carry both: While UBR SVCs are perfectly suitable for a WAN environment, they do not work over policed WAN VPs, where less than line speed is available for cost reasons. Furthermore, signalling is not currently supported by the JAMES infrastructure available for our tests. On the other hand, the CBR or VBR services used over WAN links play a less important role in the LAN due to the configuration overhead involved. We are left for the time being with two somewhat incompatible worlds. In our experiment, we try to expand the typical LAN service SVC to the WAN.

2.3 SVC extension to the WAN

The most straight forward way of extending the SVC network to the WAN would be by establishing signalling relationships between NRN and PNO switches. Due to lack of signalling support on JAMES switches, this was currently not feasible.

2.4 SVC tunnelling

We decided therefore to follow another path and deploy SVC tunnelling instead.
Signalling information is not exchanged between neighbouring switches using the reserved VPI = 0, but between two switches linked together with a standard PVP connection with an arbitrary VPI. The same PVPC is used to carry the SVCs established between those two switches.

Besides the necessity in our case to use tunnelling due to the lack of signalling support on the JAMES network there may be other reasons to choose tunnelling instead of direct signalling between adjacent switches. It might prove very difficult to define and enforce an access policy for potentially very expensive SVCs, while tunnelling relies on a fixed bandwidth with fixed costs. Furthermore, it might help to cope with incompatible signalling protocols and / or addressing schemes in the provider and user network.

2.5 Experiments and findings

1. International IP network based on SVCs
JAMES provided PVPs were used to link the ATM test networks of all participating NRNs. Some technical data about this network:

- Bandwidth of all JAMES provided PVPCs: 2Mbps
- Signalling: UNI3.1 signalling was used where supported, UNI3.0 otherwise.
- ATM addressing scheme: NSAP addressing with static IP to NSAP address mapping (no ATMARP server) on each host.
- ATM routing protocol: IISP (a static NSAP prefix based routing protocol) was used on all links.
- IP related information: all ATM end systems were placed into the same / 24 IP LIS (logical IP subnetwork) without routing protocol and no transit traffic.

Performance of this network was extremely poor and could only be used to verify proper functioning of the signalling protocol and to test set-up and tear down of SVCs. Usage on the IP layer was limited to the exchange of small sized packets to test connectivity.

- ATM hosts use normally UBR traffic class to communicate among each other over SVCs, either without or with proprietary flow control and congestion avoidance. This works well in a possibly single vendor LAN environment, but fails miserably as soon as policed PVPCs are used, since no mechanism is provided to adjust the sender's cell rate to match the contractual value of the PVPC. Small packets of only a few cells length and signalling messages are normally acceptable due to the CDVT of the PVPC.

- Tests using CBR traffic class instead of UBR failed, as our switches could not handle those signalling requests.

- Some ATM end systems offer the possibility to limit the bandwidth of each or the sum of all SVCs to a configurable value. While this might look promising, it does really not solve the problem: Since it does not get feedback from intermediate switches about the behaviour of other SVCs from possibly other ATM end systems, it is not known how much bandwidth is available. Limiting the bandwidth of each SVC in such a way that the sum of all possible SVCs will not exceed the contractual bandwidth will help, but only at the cost of possibly very poor bandwidth utilisation and throughput on a single SVC.
- All equipment available for our tests was able to establish signalling sessions among each other.
- Problems were encountered in the way VCIs are chosen when establishing SVCs: The VCI range on some equipment is hard coded but it has to match the range of its signalling partner. Otherwise, intermittent SVC set-up failures will result, whenever a VCI outside the configured range is chosen. We could detect that some equipment even tried to use VCI values outside the configured range of VCIs, which were consequently rejected and resulted in set-up failures.
- PVPC will generally be assigned different VPIs on both ends. But on some of our switches this caused signalling problems. One of the two switches on the ends of a signalling tunnel has to be configured as "user" switch, the other one as "network" switch. The latter will assign VPI / VCI values to new SVCs. The "user" switch must use the VPI value assigned to this end of the PVPC instead of the signalled value. This does not work on some of our switches and we had therefore to ask JAMES to assign the same VPI values on both ends of the PVPCs.
- High SVC set-up failure rates were observed, specially when multiple long distance VPs were involve. The reasons could not be determined properly due to the lack of proper ATM analysers. Some of them might have been caused by packet loss due to improper traffic shaping.
- In some cases SVCs were established, but not used by the end systems, most probably due to bugs in the IP stack of those systems. SVC set-up times very often exceed our expectations by far, further investigation is required on that point.
- Probably due to bugs in the IP stack, establishment of multiple SVCs between pairs of sites were observed.

2. SVC network with reshaping

Most, but not all of the switches used in our SVC network are capable of doing some form of traffic reshaping. This is required to ensure that the traffic contract on policed links is not violated. With sufficiently large buffers this can be used to shape reasonably well behaved UBR traffic into CBR VPs. No cells will be lost any longer due to policing, but excessive cells might have to be removed from the output queue. EPD, when available, makes sure that no fragments of packets get transmitted. This should work quite well with TCP / IP traffic, as the source will dynamically react to packet loss by adjusting the bandwidth. Excessive packet loss is thus prevented.

We therefore eliminated or upgraded the switches without support for reshaping from our SVC network and enabled reshaping on all VPs towards policed VPs. Observations:

- Traffic policing is performed on a per VP basis as should shaping. But some of our equipment is only capable of shaping on physical interfaces. (With two spare interfaces per VP and a cable to interconnect them, per VP shaping can be emulated).
- Reshaping could cure neither the high failure rate to establish SVCs across the network nor the too high set-up times.
- Once established, SVCs proved to be suitable to carry general purpose TCP / IP traffic with good bandwidth utilisation.
- Very little documentation is available in general about the way reshaping is realised on our equipment and how it copes with different traffic classes. We suspect, however, that only one queue is used for all traffic classes on our equipment.

2.6 Further work

- The high SVC set-up failure rate and the too high set-up times will be investigated in more detail.
- Products supporting ABR or similar traffic classes offering best effort services with congestion avoidance and flow control might become available in the not too far future. The suitability of new traffic classes in an SVC environment including long distance connections should be tested.

3. Classical IP and ARP Over ATM

Simon Leinen <Leinen@switch.ch>
Ramin Najmabadi Kia <najmabadi@helios.iihe.ac.be>

3.1 Introduction

In Classical IP over ATM as defined in [3.1], a specialised variant of ARP server is used to resolve layer-three (IP) addresses to layer-two (ATM NSAP or E.164) addresses. The main difference to traditional ARP is that - because ATM lacks a broadcast facility - there is a single designated ATMARP server whose layer-two address has to be configured statically in each client.

3.2 Protocol Operation

Classical IP (CLIP) is based on ATM Switched Virtual Circuits (SVCs). It is only defined within a Logical IP Subnet (LIS). When a CLIP node wants to send an IP packet to another CLIP node on the same LIS, and no SVC between the two nodes has been established yet, the sending node has to request an ATM SVC to the receiver. For this purpose, it needs to know the receiver's ATM address. Unless the mapping is already in the cache, it queries the ATMARP server.

Likewise, when a CLIP node receives an SVC connection request from another node, it uses an Inverse ARP (InARP) request to the ATMARP server to find the protocol address of the sender. Communication between CLIP nodes and the ATMARP server is performed using AAL5 / SNAP over a regular SVC, and the ATM address of the ATMARP server has to be configured statically in each node on the LIS. This SVC can also be used to carry IP traffic between a node and the node running the ARP server.

3.3 Experiment Set-up

Building on the configuration for the SVC tunnelling experiments, another range of network addresses (193.203.225.0 / 24) was reserved for this experiment. The participants had to configure an additional ATM sub-interface in a CLIP / ATMARP configuration on their nodes. An ATMARP server was configured on a Cisco router at the University of Linz in Austria, which was used by all participants. The only addresses that had to be configured on each participating interface were:

- the local IP address
- the ESI of the local NSAP address
- the NSAP address of the ATMARP server

Compare this with the set-up for the SVC tunnelling experiment (see section 2.), where every participant needed a complete table of IP / NSAP mappings for all other interfaces.

3.4 Observations

Using an ATMARP server didn't introduce any new instabilities for the participants. However, problems with the SVC tunnelling network could prevent potential participants from contacting the ATMARP server, which would make all communication within the LIS impossible, even though some destinations would be reachable on the ATM level. On the other hand, static IP-to-NSAP address mappings aren't necessary when ATMARP is used, removing another common source of errors and maintenance effort.

The dependency on a single ATMARP server per LIS is a severe drawback, in particular in a WAN setting. Most ATM-level connectivity problems cause the ATMARP server to be reachable for some parties, making all communication impossible. The NBMA Next Hop Resolution Protocol [3.3] alleviates the problem by allowing smaller LISes and permitting layer-2 connectivity outside the LIS. Having multiple redundant address resolution servers necessitates a synchronisation protocol such as SCSP used in MARS [3.2].

3.5 Timing Results

The following table compares response times for ICMP echo requests ("pings") within the same LIS, on the one hand using an ATMARP server, on the other hand using statically configured IP-NSAP address mappings. The first packet takes a bit longer to respond to using ATMARP, because the ATM server has to be contacted by the sender (ARP request) and / or responder (InARP request). For subsequent packets, the response time was the same in both set-ups, except for differences due to the SVC tunnel topology that had changed between both tests.

The SVC between an ATMARP client and the server is usually kept active permanently, so calls to the ATMARP server are not included in the timings below. All measurements were carried out from Switzerland (unit is msec).

To country	To host	with ATM-ARP 1st	with ATM-ARP nth	without ATM-ARP 1st	without ATM-ARP nth
CH	CH1	NR	1	NR	1
(local)	CH2	0	0	0	0
	CH3	NR	1	11	1
AT	AT1	13	13	55	13
NO	NO1	386	84	242	86
	NO2	619	78	299	90
UK	UK1	931	125	305	136

3.6 Conclusion

ATMARP works quite well as an address resolution protocol mapping IP to ATM NSAP addresses. Its use yields an extremely simple configuration for an IP subnetwork over an ATM SVC infrastructure. Protocol overhead is very small and only noticeable on new connections.

The most important drawback is the dependency on a single address resolution server. This may be tolerated on a LAN, where the server can be run on a system whose functioning is vital to the network anyway (such as the single ATM switch). In a WAN setting however, lower-layer communication problems involving paths to the ATMARP server have fatal effects for its clients. More redundancy seems to be needed here.

4. IP routing over ATM with NHRP

Olav Kvittem <Olav.Kvittem@uninett.no>
Vegard Engen <vegard.engen@uninett.no>

4.1 Background

An IP-system at the edge of an ATM-network needs to find for a destination IP-address the ATM-address for the optimal next hop over the ATM-network so that it can set up a call there. A partial solution to this problem is the ATM ARP in RFC 1577 (Classical IP over ATM) which solves the problem for one IP subnet. This does not scale to large multi-organisation networks. The Next Hop Resolution Protocol (NHRP) proposes a solution for shortcutting subnet based routing so that one can minimise the number of hops through the same ATM cloud.

Given an European academic ATM-based backbone with possibly more than 40-50 nodes, NHRP might be the way a pan-European academic IP-network could be practical to set up. With statically set up connections the network would be tedious to maintain and lead to incomplete direct connectivity and thus inefficient use

of network resources. With NHRP one could hope for automatic set-up of connections to new nodes with a traffic interest. The same problem exists perhaps to an even bigger scale in national academic networks.

4.2 Clients and servers

The current status of the development of NHRP at IETF is that the protocol is under consideration by the IESG as a proposed standard. This means that the protocol is fairly stable.

There is however still few implementations available. There has been an implementation for Cisco routers available for a while. This was chosen for the tests. There is also one for a workstation, but for an older incompatible version of the protocol.

4.3 SVC infrastructure

The NHRP operation is dependent on having a ubiquitous SVC-connectivity among the participants forming a logical NHRP cloud over the TF-TEN ATM VP overlay network. Such an infrastructure has been prepared by the ATM SVC-project described in section 2. However the NHRP project copied that set-up putting in its own VC's in order not to interfere with the other experiments like SVC and ATM-ARP.

4.4 NHRP operation

The routers at the edge of the ATM-cloud will act as NHRP servers. There need to be a initial connectivity between routers / hosts on the IP level so that NHRP can work. This can be a slow indirect path. The initial pre-defined VC-connections defined are that each country connects to one of two interconnected centres in Germany and Austria.

The NHRP servers will have the same network-id, that will tell them that they are on the same ATM-cloud when receiving an NHRP-request, and may return info about their ATM-address to the requester. A NHRP request will be sent when a pre-defined amount of packets has been sent towards a destination. The request will be passed along to NHS servers on the ATM-cloud until no further downstream NHS-servers are available. The egress router from the ATM-cloud will then return his ATM-address to the ingress router.

4.5 Tests

A demonstration of the basic behaviour of the Cisco implementation is shown in the following simple test:

- A and B has a IP-link with the ATM address of each other as well as B and C, but A and C does not know how to contact each other.
- A sends echo packets towards C via the default route to B.
- A brings a SVC to B to serve that traffic.
- B sends the packet on to C and brings up an SVC to do that.
- C returns the packets to A via B
- A sends a NHRP request to C via B after some packets. The NHRP packet contains A's ATM-address and C tries to set-up a direct connection via the ATM-address from the request but fails due to SVC-problems
- C responds to A via B with it's ATM-address
- B receives the reply and sets up a SVC to A

This experiment was performed between Austria, Switzerland and Norway (ABC) and the round-trip time with the A-B-C path was 108 milliseconds, while it was about 76 ms with the A-C VC.

4.6 Conclusions and Further studies

This simple experiment has demonstrated that and how the NHRP basic functions works. The implementation is still largely untested and we experienced router crashes, routing tables flushes and looping SVC-control processes during testing. There were also some problems on the ATM SVC-level that are mentioned in the SVC-experiment

There is also some functionality missing in the ATM implementation of the router, like queuing up packets while waiting for a call to be set up. As it is now, packets coming in are lost until a call is active.

This version of NHRP only supports lookup of addresses directly reachable from the egress router. This means that to make transit traffic between the networks behind the respective connections flow on the NHRP connection one must use normal routing on top. NHRP would be more useful in a backbone with such and extension (NHRP-R2R).

The present ATM network (JAMES + NRNs) does not support any means of resource control in the network besides static allocations. Due to inherent properties of ATM the packet loss can be disastrous when a link is saturated using Unspecified Bit Rate (UBR). Setting up a large number of unrelated NHRP UBR VC's in a not controlled resource environment is not recommended. NHRP does not have any resource res-

ervation mechanisms, so one would have look to ATM mechanisms like Packet Discard, Available Bit Rate Services and resource reservation, or to higher level like RSVP.

It is therefore highly recommended that this project continues with the targets of advancing on the above mentioned issues like a larger scale pilot, transit routing and resource control.

5. European ATM Addressing

Kevin Meynell <K.Meynell@terena.nl>

5.1 Introduction

The forthcoming introduction of ATM signalling means it will be necessary to devise an ATM addressing scheme for European NRNs. The aim was to investigate a scheme that would allow experiments with UNI signalling and routing services. It was also hoped a universal scheme would allow the scope of the JAMES experiments to be easily expanded, and avoid a lot of re-configuration work in the future.

5.2 Findings

Most NRNs have indicated they would prefer to use NSAP addressing as this provides the fine address resolution they are likely to require [5.1, 5.2, 5.3]. As various NSAP formats are well defined, it is really only necessary for each NRN to obtain an NSAP prefix from the ISO National Member Authority for their country (e.g. the British Standards Institute in the UK). The NRN may then allocate the undefined octets in a manner that suits it is topology / organisational structure. JANET (UK) and RENATER (France) have already devised DCC-format NSAP schemes that could possibly be adapted by other NRNs.

Most of the European PNOs however, have indicated they will be using E.164 addressing, the ITU standard relating to international ISDN numbering. Consequently, this means there must be a method for NSAP addresses to traverse the PNO-provided network.

ATM Forum standards state that where a call originates from, and is destined for, networks supporting NSAP addresses, the NSAP address may be carried in the E.164 sub-address field over an E.164 network. The E.164 address (Called Party Number) required for transit must be derived from the NSAP address at the gateway between the two networks. Where a call originates from a network supporting NSAP addresses and is destined for a network only supporting E.164, the Called Party Number will be coded as an NSAP-formatted E.164 address.

Unfortunately, there are not any standards for this and translation appears to have been left to the switch suppliers to implement. The only switch supplier known by the author to be working on a solution is Cisco and this is proprietary. Another problem is the differences in field length between E.164 and NSAP addresses, and the fact that some telecommunications switch manufacturers may not support the full E.164 field length. This could conceivably mean that parts of an NSAP address would be discarded when entering an network only supporting E.164. Indeed, it appears the PNOs themselves are not yet sure how to proceed on these issues.

5.3 Future Progress

The ATM Forum and ITU are currently working to define some standards in these areas, but nothing firm has been published. Until this happens, which is unlikely to be until next year, further progress will be inevitably restricted.

Nevertheless, it is not currently an issue for the SVC experiments over JAMES as they are being tunnelled over VPs. The NRNs should still also be able to determine their NSAP addressing schemes to use, which would allow real values to be assigned to their equipment (as JANET as done). Indeed, it would be of benefit to their own internal ATM networks.

5.4 Summary

The table below is a summary of the known address schemes that will be used by European NRNs and the JAMES partners.

Country	NRN	PNO
Austria	NSAP DCC	E.164
France	NSAP DCC*	E.164
Germany	NSAP DCC	E.164
Italy	NSAP DCC	No decision
Netherlands	NSAP DCC	E.164
Norway	NSAP ICD*	No decision
Spain	NSAP DCC	E.164
UK	NSAP DCC*	E.164 NSAP (interim)

* denotes an official scheme has been published

6. ATM Network Management

Zlatica Cekro <cekro@helios.iihe.ac.be>

6.1 Goals

The principal goals in this experiment were to gain experience of the management aspects on ATM management service offered by JAMES as a public ATM network. The management service covers the domain of system interoperability at the management plane in the aspects of configuration, performance, accounting, fault and security. The idea was to test the Customer Network Management for ATM Public Network Service i.e. M3 interface in ATM Forum classification / Xuser interface in the ITU-TS classification. According to the M3 specification "read only" management service (Class I of requirements) is mandatory if the service provider offers any management service. The service includes: General UNI Protocol Stack Information, General ATM Level Performance Information, ATM Level Virtual Path / Virtual Channel (VP / VC) Link Configuration and Status Information, Traffic Characterisation Information, Event Notifications from the Public Network Provider.

6.2 Infrastructure

A Management Platform for Network monitoring and statistics collection based on the SunNet Manager-SunNet Domain Manager version 2.3 on Solaris 2.4 was used. It was connected to the Overlay ATM Network and Internet network for the tests. The Platform was remotely accessible to all participants in the tests through X-window terminal. Some of collected statistics were presented in the WWW page.

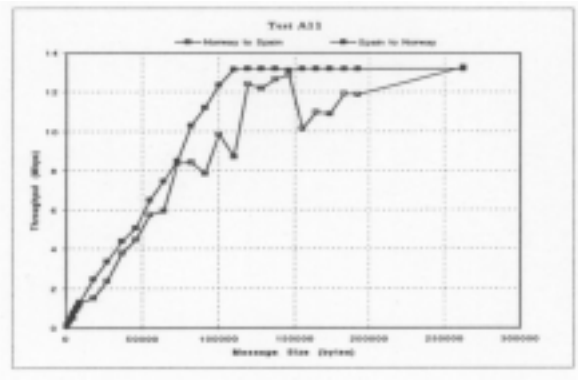

6.1: Infrastructure used for ATM Network Management tests

Tested MIBs:

- LS1010: ATM-MIB, PNNI-MIB, Cisco-ATM-Addr-MIB, Cisco-ATM-conn-MIB, Cisco-ATM-if-MIB, Cisco-ATM-Phys-MIB, Cisco-ATM-RM-MIB, Cisco-ATM-switch-addr-MIB, Cisco-traffic-MIB, ATM-forum-MIB
- UB GeoSwitch 155: ATM-forum-MIB, ATM-forum-addr-reg-MIB, ATM-MIB
- FORE ASX200: ATM-forum-MIB, ATM-forum-addr-reg-MIB, FORE-switch-MIB
- Cisco LS100: LS100-MIB

6.3 Summary of results

The initial scenario required an active role both of the NRNs and of the service provider - JAMES. The JAMES for this moment does not offer any management service to the end-users. Because of that we decided to test the similar functionality on the M2 / M3 interface only at the NRNs side as described above. The interface M2, a management interface needed to manage a private ATM network, has not been standardised and in practice it has the same functionality as M3 interface. The test results could be summarised as follows:

- SNMPv1 or SNMPv2 based agents are widely implemented in the NRN edge devices which we tested: CISCO LS1010, CISCO LS100, FORE ASX200, UB GeoSwitch, CISCO routers with ATM interfaces.
- ATM based standard MIBs like ATM MIB (RFC 1695) and ATM FORUM UNI MIB are supported widely by tested ATM switches.
- Very rich proprietary ATM specific MIBs were tested at CISCO LS1010 and FORE ASX200.
- OAM F4 and F5 Loopback end-to-end flows (ITU-TS I.610) are supported at tested ATM switches.

7. CDVT Experiments

Victor Reijs <Victor.Reijs@SURFnet.nl>
Phil Chimento <chimento@cs.utwente.nl>

7.1 Introduction

The CDVT experiments were created to study experimentally the behaviour of CBR cell streams in an ATM network. CDVT can be thought of as the deviation of the cell stream from a perfect constant cell rate. CDVT can be introduced into a cell stream by a variety of causes, including the slotted nature of ATM, switch fabric service disciplines within particular switches, and queuing and output port contention within switches. The

experiments performed so far have focused mostly on CBR streams which are important in a number of applications, including: PDH circuit emulation, 64 kbit/s voice support and CBR video applications. These applications demand very tight bounds on the CDVT in order to operate correctly.

7.2 Configuration and Experiments

Our approach was to look at specially generated cell streams as they traversed an ATM network. We set up an international VPC from the University of Twente in the Netherlands to the University of Stuttgart in Germany. There were a total of 5 measurement points set up along the path of the cell stream, although not all measurement points were active for each experiment. Two measurement points were at the University of Twente, one point at KPN Research in Leidschendam, one point at Deutsche Telekom in Cologne, and one point at the Univeristy of Stuttgart.

SURFnet bv and the partners just mentioned composed the team that performed the experiments. The people involved are: Victor Reijs (SURFnet), Harrie van de Vlag (KPN Research), Dirk Hetzer and Mr. Schurillis (Deutsche Telekom), Robert Stoy (University of Stuttgart) and Edward Meewis and Phil Chimento (University of Twente).

The VPC used was a permanent VPC which was established at a CBR circuit with a maximum speed of 4750 cells per second. The actual path that the cell streams took (for most of the experiments) went through nine different switches located in the local university networks, the Dutch and German national networks and the JAMES network. Some switches were traversed twice, making a total of eleven times that the cell stream went through a switch. The links were 155 Mb/s except in the Dutch national network and JAMES, where they were 34 Mb/s.

The measurement equipment consisted of an HP 5200A, two HP75000s and a W&G DA30c. All the HPs have a time resolution of 100 ns and the W&G has a time resolution of 10 microseconds. The streams consisted of 110,000 cells each and were completely captured by the measurement equipment. No cells were lost during the experiments.

In December 1996 and January 1997, we ran a total of 17 international trials. Most of these concentrated on CBR streams generated by a SUN Ultra 1 workstation with a SUN ATM card and software. The CBR streams were low speed (500 kbit/s, 1.1 Mb/s, 1.6 Mb/s) and, as measured at the tailgate of the SUN, quite stable. We measured on a total of 5 different days, around mid-afternoon each of the days. The traces themselves can be found at URL: http://wwwtios.cs.utwente.nl/~chimento/tf10exp/tracelist-1.html.

7.3 Results

Figure Figure 1 shows an analysis of the GCRA parameters of a 1.1 Mb/s cell stream that has been captured at each of the measurement points. This figure shows the CDVT required for the cell stream to pass a policing GCRA with zero violations (that is, no dropping or tagging due to policing) as a function of the inverse of the PCR. In this case, the cells are sent at a nominal rate of 2624.67191 cells per second, or in other words, an inter-cell time of .381 milliseconds. The interesting part of this graph lies to the left of the steep rise which occurs at the nominal inter-cell time. This portion of the graph shows the effect of the spread of the cell stream as it moves through the international network. The measurement point “Enschede1” is the tailgate of the SUN Ultra generating the cell stream. The CDVT required for zero violation at the nominal inter-cell time is about 2.1 microseconds which accords with the spread of the inter-cell time distribution shown

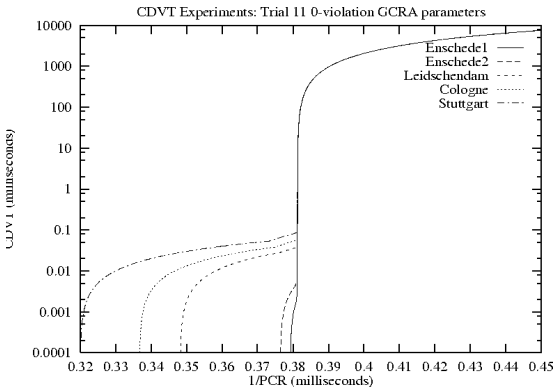

It is clear that with each successive measurement point, the CDVT increases considerably. “Enschede 2” is the measurement point just after the first switch in the network. At this point, the CDVT required for zero violation has increased to 5.4 microseconds. At Leidschendam, the CDVT required is 37.6 microseconds, and at Cologne it has risen to 55.9 microseconds. Finally when the cell stream reaches Stuttgart, the CDVT required is 86.0 microseconds. If we look

at the minimum and maximum inter-cell times at each measurement point, we find that the difference in these inter-cell times grows also at each measurement point. For this trace, the difference between the min and max at "Enschede 1" is 3.3 microseconds (a little more than 1 cell time at 155 Mb/s). After the first switch it is 8.8 microseconds, at Leidschendam, the difference is 65.3 microseconds, at Cologne it is 110.9 microseconds, and at Stuttgart it is about 130 microseconds (+ or - 9 microseconds).

We are continuing to investigate these results with a view to drawing more information from them. Initially it looks like this spread is introduced by the switches themselves, but there are other possible explanations which we will investigate further.

8. IP over VBR

Olivier Martin <omartin@dxcoms.cern.ch>

The goal of the VBR experiment is to assess the technical advantages of using VBR for IP traffic, as opposed to CBR. At first glance it looks like VBR provides the same sort of flexibility as Frame Relay does. However, on a VBR service the peak cell rate (PCR) can only be used for very short periods of time over the guaranteed average rate, the SCR (sustainable cell rate). The burst tolerance (BT) specifies the duration of the peak, and is usually in the range of a few cells. After a peak the transmission has to slow down for some time, to reach on average the SCR again. Thus a VBR service does not provide "additional bandwidth" on top of the SCR for IP traffic.

The main question of the experiment is whether there is any benefit at all for IP traffic on a VBR service. Experiments on VBR services were carried out in the Netherlands [8.1] and in Switzerland, in both cases nationally. The main results from the tests were:

- The IP throughput is mainly determined by the SCR. It will not go significantly beyond the SCR, and certainly not to the PCR.
- If the BT on the router is higher than the burst tolerance on the ATM network, the "goodput", i.e., the usable IP throughput is close to zero. This is not surprising, as the policing of the ATM network would discard cells that do not correspond to the traffic contract, and as any cell loss leads to severe IP packet losses.
- Higher BT (within the traffic contract) leads to slightly higher throughput. This is also not surprising.
- Measurements with PCR = 2 * SCR showed slightly lower throughput than with PCR = SCR. This is surprising, since an increase in PCR should increase the performance, rather than decrease it. Up to the time of writing this paper no explanation for this behaviour could be found.

For more information on these results see [8.1]. A confirmation of these results over an international VP was being carried out at the time of writing this paper. The results shown here demonstrate that VBR services can be used, but they do not offer any advantages over CBR services for IP traffic, except that on the ATM network itself it is cheaper to provide VBR services. Configurations should always be with PCR = SCR, and the BT should be as big as possible.

The VBR class of service is meant to be used by, so called, on/off sources in such a way that the average traffic is at best equal to the SCR. This kind of traffic profile which is well adapted to special applications with periodic bursts, is obviously not well adapted to aggregated IP traffic. However, it would seem that non real time VBR can also be used to provide a frame relay like service, provided suitable bursts are accepted by the ATM network. It could therefore be that there are significant price performance advantages in using nrt-VBR rather than CBR services. Still not clear are the parameters that will be offered beyond the PCR and SCR values, i.e. CDVT and MBS.

9. IP resource reservation over ATM

Olav Kvittem <Olav.Kvittem@uninett.no>
Sabine Kuehn <kuehn@ibdr.inf.tu-dresden.de>
Frank Breiter <breiter@ibdr.inf.tu-dresden.de>
Steinar Haug <sthaug@nethelp.no>

9.1 Background

The purpose of the experiments is to evaluate the use and support of RSVP in the network like a European or national academic backbone, as opposed to what and how services could be provided over RSVP (the latter is e.g. done by the MERCI-project). The focus is therefore on how the routers and ATM would work to support RSVP.

9.2 Summary of results

University of Dresden has developed RSVP over IP over ATM for DEC Workstations so that this functionality could be tested in a local ATM environment using an own performance tool with an integrated graphical RSVP user interface. Moreover, the University of Dresden is developing a video conferencing system testing RSVP over ATM by a more practical-relevant example.

IETF is working on how to map RSVP to ATM (ISSL working-group), and we present and implements some ideas on how to do this. Since RSVP is potentially interesting for IP multicast applications, we are also looking at MARS (Multicast Address Resolution Server), a way of keeping track of point-to-multipoint VC's for IP multicast.

RSVP is still experimental, and is only on link in building Integrated Services in the Internet. The understanding and specification of policy related issues like who's paying and how to tell the router ones policy still has to mature both technically and organisationally for some years. RSVP could at the moment soon be put to use in confined environments.

Some simple experiments with RSVP over IP through RSVP-capable routers are being performed over IP / ATM over JAMES.

9.3 Conclusions and future work

We have explored and got experience with some aspects of RSVP use in routers and implementations in workstations and routers. The next step would be to follow the development of RSVP to ATM mapping more closely and also to look at mapping RSVP for IP multicast over ATM.

10. Security in ATM Networks

Paulo Neves <pneves@rccn.net>
Roberto Canada <canada@rccn.net>
José Vilela <vilela@rccn.net>

10.1 Introduction

This workpackage aims to study security issues on ATM WAN networks. This is a recent field of study. Standardisation work at the ATM Forum is under way regarding the future shape of ATM Security infrastructure. This infrastructure considers the use of special signalling procedures to allow for negotiation of security parameters between communicating parties. There is also some academic research on this area [10.1].

We intend to analyse equipment and services to find their vulnerabilities in the JAMES context. Our goal is not to produce new specifications or to propose changes to the existing protocols, but to try to define a set of recommendations on how to configure the equipment in order to improve security.

10.2 General Security Requirements

Some basic behaviours concerning security are expected from any communication network. We find that at least availability (in the sense of not denying available resources), secure communication channels and accurate auditing information are essential requirements. Aspects like user authentication and non-repudiation of contents (of user messages) should not be expected from the network as an entity, although they might be supported by other means.

10.3 Possible Threats

We identify three classical attacks and their consequences on each ATM flow:

- Data or traffic flow confidentiality loss due to an intruder eavesdropping the network and deducing user data content or user traffic features;
- Data integrity loss caused by accidental or malicious injection / removal / modification of cells / signalling messages in transfer;
- Overloading problems following a mass-injection of cells / signalling messages;

1. User Data Flows
Confidentiality and integrity losses are particularly damaging when applied to user data flows since an intruder eavesdropping at a point on the network can retrieve all the cells belonging to one connection. An intruder may also disrupt the network by injecting, modifying or removing user cells. Most often these cells are removed at the receiving entity causing retransmission of upper-layer frames, and overloading the network.

2. Signalling Flows
Signalling flows vulnerability is message type dependent. Since SET UP messages for establishing point-to-point connection are the only ones bearing the sensitive information - called and calling end-entities addresses, they appear as the most vulnerable messages to eavesdropping attacks. Also overloading the network with SET UP messages is damaging since this causes mass connec-

tions set ups and therefore end or network entities overload and consequently legitimate connections rejections. Other messages such as RELEASE and RELEASE COMPLETE are vulnerable to integrity attacks because their injection immediately causes a connection release, which can also be viewed as a DoS attack.

3. Management Flows
An intruder eavesdropping performance management cells can infer the number of user cells transmitted over one connection. Also an intruder realising an attack on integrity may cause line errors to remain undetected, a connection release whereas the connection is still operational or a wrong line problem location.

10.4 Security Services

In order to avoid the threats mentioned above security services need to be introduced within ATM planes. These are summarised in the table below.

10.5 Considerations

We consider the availability of some of these services (namely to the Control and Management Planes) is essential for the robustness of the network itself. In fact, we find that the integrity of the network depends on the existence of means to avoid some forms of attack (Denial of Service, Masquerade, Spoofing and Repudiation), on signalling and management protocols, even if user security services could be performed at higher layers.

10.6 Future Work

In the JAMES framework we are confined to user data channels, running through PVCs, without any means to directly contact intermediate ATM switches, for connection negotiation or management. In order for us to test the most interesting issues, some control and management functions would have to be present. In the meantime we will try to develop some experimental work over the SVC infrastructure.

10.7 Acknowledgements

During our work we had a valuable contribution from Maryline Laurent and Pierre Rolin [10.1], on which this section is partly based on. We also are in contact with Martin Moore in order to define some experiments in co-operation with JAMES.

Summary

The work carried out in this framework shows that most of the advanced features of ATM and the new IP protocols are not yet in a state where they can be used safely for operational services. The problem seems to be in most cases that the development of hard- and software is not mature enough. The results of the experiments do however show how to make best use of the existing services (CBR, VBR), and give a good insight into the problems that arise with new technologies.

All the experiments listed here are part of the first phase of the TEN-34 testing programme, which was not finished at the time of writing this paper. More work is clearly needed to fully understand the capabilities of ATM networks and of comparable IP services. In some of the areas described above new questions arose during the tests.

There are also a number of technologies which were not yet examined in phase one. Phase two of the project, starting in May 1997, will also investigate into other technologies, such as ATM routing and new traffic classes such as ABR. The focus of the tests carried out here is to make experimental services available on the production TEN-34 network. Although the more interesting features of ATM seem to be not in a state yet for an operational service, we will keep on following the developments in ATM and IP related activities. The latest information on our experiments can always be found on the TF-TEN home page [0.2].

References

[0.1] TEN-34: Project to establish a pan-European high-speed research backbone (Trans-European Networking on 34 Mbit/s). See http://www.dante.net/ten-34/

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

[0.3] For more information on JAMES see http://www.labs.bt.com/profsoc/james/

[0.4] Behringer, Michael: “TEN-34 and JAMES: Technical Plans”; Proceedings JENC7, paper 131, 1996. Online: http://www.dante.net/pubs/dip/23/

[0.5] Behringer, Michael: “The Implementation of TEN-34”; Proceedings JENC8, paper 331, 1997. Online: http://www.dante.net/pubs/dip/28/

[0.6] Project TEN-34, Deliverable D11.2, Interim Results of Phase 1 Test Programme. See http://www.dante.net/ten-34/DELIVERABLES.html

[1.1] Permanent virtual circuits configuration and TCP-UDP / IP performances in a local ATM network; C.Battista, M.Campanella, T.Ferrari, A.Ghiselli, C.Vistoli. INFN Internal Note n. 1069, July 1995

[1.2] Performance evaluation of TCP(UDP) / IP over ATM networks; S.Dharanikota, K.Maly, C.M.Overstreet, Computer Science Dep., Old Dominion University, Norfolk VA

[1.3] A Performance Analysis of TCP / IP and UDP / IP Networking Software for the DECstation 5000; J.Kay, J.Pasquale; Computer Systems Laboratory, Dep. of Computer Science and Engineering, University of California, San Diego

[1.4] High Performance TCP in ANSNET; C.Villamizar, C.Song

[1.5] High-performance TCP / IP and UDP / IP Networking in DEC OSF / 1 for Alpha AXP; Digital Technical Journal, vol. 5, n. 1, win 1993

[1.6] How a large ATM MTU causes deadlocks in TCP data transfers; K. Moldeklev, P. Gunninberg (Norwegian Telecom Research and Swedish Institute of Computer Science).

[3.1] RFC1577; M. Laubach, Hewlett-Packard Laboratories, "Classical IP and ARP over ATM", January 1994

[3.2] SCSP: James V. Luciani, Grenville Armitage, Joel Halpern, "Server Cache Synchronization Protocol (SCSP)", November 1996; draft-ietf-ion-scsp-00.txt (work in progress)

[3.3] NHRP: James V. Luciani, Dave Katz, David Piscitello, Bruce Cole, "NBMA Next Hop Resolution Protocol (NHRP)", March 1997; draft-ietf-rolc-nhrp-11.txt (work in progress)

[5.1] George Howat - University of Edinburgh. The JANET ATM Addressing Scheme (http://www.ed.ac.uk/~george/ukac-index.html).

[5.2] Kjetil Olsen - University of Oslo. The Uninett ATM Addressing Scheme (http://www.uninett.no/info/nett/supernett/atm-adresser.html)

[5.3] Victor Reijs - SURFNET. ATM Addressing (http://www.nicsurfnet.nl/surfnet/projects/atm/atmaddr.htm)

[6.1] ATM Forum, Customer Network Management (CNM) for ATM Public Network Service (M3 Specification), Rev. 1.05, January 1996

[6.2] ITU-T, I.610, Integrated Services Digital Network (ISDN), Maintance Principles, B-ISDN Operation and Maintenance Principles and Functions, November 1995

[6.3] ITU-T, I.751, Integrated Services Digital Network (ISDN), B-ISDN Equipment Aspects, Asynchronous Transfer Mode, Management of Network Element View, March 1996

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

[10.1] Maryline Laurent and Pierre Rolin, "Securite ATM : une analyse de flux menee sur quatre architectures de reseaux", GRES'95, Paris, September 1995. Online: ftp//ftp.rennes.enst-bretagne.fr/pub/security/ml_GRES95.ps.gz.

Acronyms

ABR	Available Bit Rate
ARP	Address Resolution Protocol
BT	Burst Tolerance
CBR	Continuous Bit Rate (ATM traffic class)
CDV(T)	Cell Delay Variation (Tolerance)
CLIP	Classical IP
E.164	(ITU-T recommendation on addressing)
EPD	Early Packet Discard
IESG	Internet Engineering Steering Group. Manages the working groups and standardisation process in IETF
IETF	Internet Engineering Task Force (http://www.ietf.org); The Internet protocol standardisation body
IISP	Interim Inter-Switch Signalling Protocol
JAMES	A European experimental ATM-network. See [0.3]
LAN	Local Area Network
LIS	Logical IP Subnetwork
MARS	Multicast Address Resolution Server
MBS	Maximum Burst Size
MIB	Management Information Base
NBMA	Non-Broadcast Multiple Access
NHRP	Next Hop Resolution Protocol
NHRP-R2R	NHRP for Destinations off the NBMA Subnetwork
NHS	Next Hop Server
NRN	National Research Network
NSAP	(OSI addressing standard)
OAM	Operations, Administration and Maintenance
PCR	Peak Cell Rate (ATM traffic parameter)
PNNI	Private Network-to-Network Interface
PNO	Public Network Operator
PVC	Permanent Virtual Circuit
PVP(C)	Permanent Virtual Path (Connection)
RSVP	Resource ReSerVation Protocol
SCR	Sustainable Cell Rate (ATM traffic parameter)
SCSP	Server Cache Synchronisation Protocol
SNMP	Simple Network Management Protocol
SVC	Switched Virtual Circuit
TEN-34	Trans-European Network interconnect at 34 Mbit/s
TF-TEN	Task Force for Trans-European Networking
UBR	Unspecified Bit Rate (ATM traffic class)
UNI	User-Network Interface
VBR	Variable Bit Rate (ATM traffic class)
VC(I)	Virtual Circuit (Identifier)
VP(I)	Virtual Path (Identifier)
WAN	Wide Area Network

(for more acronyms, see: http://www.fore.com/atm-edu/acronyms.html)

Information about TF-TEN

The work described here was carried out and described by the participants of the TERENA Task Force TEN (Trans-European Networking), which is being chaired by Michael Behringer. More information on the work of the TF-TEN can be found on the TF-TEN home page:

http://www.dante.net/ten-34/tf-ten/

The members of the Task Force were actively involved in one or more of the experiments. The work described here is the work of all Task Force members. These are in particular:

Frank Breiter <breiter@ibdr.inf.tu-dresden.de>
Mauro Campanella <campanella@mi.infn.it>
Roberto Canada <canada@rccn.net>
Zlatica Cekro <cekro@helios.iihe.ac.be>
Phil Chimento <chimento@cs.utwente.nl>
Magnus Danielson <magda@it.kth.se>
Vegard Engen <vegard.engen@uninett.no>
Peter Feil <peter.feil@rus.uni-stuttgart.de>
Tiziana Ferrari <ferrari@infn.it>
Joaonuno Ferreira <ferreira@rccn.net>
Alain Frieden <frieden@restena.lu>
Christoph Graf <christoph.graf@dante.org.uk>
Steinar Haug <sthaug@nethelp.no>
Ramin Najmabadi Kia <najmabadi@helios.iihe.ac.be>
Sabine Kuehn <sabine_kuehn@ibdr.inf.tu-dresden.de>
Olav Kvittem <olav.kvittem@uninett.no>
Cees de Laat <C.T.A.M.deLaat@fys.ruu.nl>
Simon Leinen <simon@switch.ch>
Simone Maggi <smaggi@net4u.it>
Olivier Martin <omartin@dxcoms.cern.ch>
Kevin Meynell <k.meynell@terena.nl>
Paulo Neves <pneves@rccn.net>
Victor Reijs <victor.reijs@surfnet.nl>
Albert Schindler <albert.schindler@seinf.unige.ch>
Guenther Schmittner <Schmittner@edvz.uni-linz.ac.at>
Robert Stoy <stoy@rus.uni-stuttgart.de>
Celestino Tomas <ctomas@chico.rediris.es >
Jean-Marc Uze <jean-marc.uze@renater.fr>
José Vilela <vilela@rccn.net>
Baoyu Wang <b.wang@ukerna.ac.uk>