TERENA/DANTE TASK FORCE FOR TESTING ADVANCED NETWORKING TECHNOLOGIES

Minutes of the 6th TF-TANT meeting held on the 25th and 26th of November 1999 at ETSIT, Madrid, Spain.

Kevin Meynell - Issue 2

PRESENT

Name	Organisation	Country
Hamad el Allali	U.Twente	Netherlands
Werner Almesberger	EPFL	Switzerland
Miguel Angel Sotos	RedIRIS	Spain
Michael Behringer	Cisco	Spain
Graca Caivallio	Cisco	Spain
Massimo Carboni	GARR	Italy
Valentino Cavalli	TERENA	-
Zlatica Cekro	VUB/ULB	Belgium
Nicole Chiminelli	CSELT	Italy
Tryfon Chiotis	GRNET	Greece
Maribel Cosin	RedIRIS	Spain
Luca Dell'Agnello	INFN/GARR	Italy
Tiziana Ferrari	INFN Bologna	Italy
Avgust Jauk	ARNES	Slovenia
Joop Joosten	CERN	Switzerland
Cees de Laat	U.Utrecht	The Netherlands
Simon Leinen	SWITCH	Switzerland
Lladislav Lhotka	CESNET	Czech Republic
Pedro Lorga	FCCN	Portugal
Kevin Meynell (Sec)	TERENA	-
Harald Michl	U.Vienna/ACOnet	Austria
Antonio Pinizzotto	CNR	Italy
Alex van der Plas	Ericsson Telebit	Denmark
Agnes Pouele	DANTE	-
Herve Prigent	RENATER/CRIHAN	France
Juergen Rauschenbach	DFN	Germany
Filipe Rebelo	FCCN	Portugal
Esther Robles	RedIRIS	Spain
Roberto Sabatino (Chair)	DANTE	-
Rina Samani	UKERNA	United Kingdom
Robert Stoy	RUS/DFN	Germany
Celestino Tomas	RedIRIS	Spain
Jean-Marc Uze	RENATER	France
Bert Wijnen	IBM	The Netherlands
Wilfried Woeber	ACOnet	Austria

Apologies were received from:

Leon Gommans	U.Utrecht/Cabletron	The Netherlands
Dimitrios Kalogeras	GRNET	Greece
Olav Kvittem	Uninett	Norway
Jan Novak	DANTE	-
Simon Nybroe	Ericsson Telebit	Denmark
Victor Reijs	SURFnet	The Netherlands

1. APPROVAL OF MINUTES

The minutes of the TF-TANT meeting held on the 30th of September 1999 were approved.

2. STATUS OF QUANTUM & TEN-155

Roberto reported the connections from Belgium to the Netherlands and France had been upgraded to 45 Mbps. The links were operational, but the TEN-155 router was not yet installed. Belnet was multihomed to TEN-155, and routing was configured to use both circuits. The triangle between Spain-France-Switzerland-Spain was due to be upgraded to 45 Mbps, whilst an additional 155 Mbps circuit would be installed between Italy and Switzerland. A 60 Mbps ATM service was being considered between the London PoP and Greece.

Cees had recently discovered some serious problems with the Germany to Netherlands connection. It was only possible to transmit a few Kbps during peak hours, and a few Mbps during off-peak hours. It transpired that the Ascend switches had been badly configured by KPN and were discarding traffic unnecessarily. A temporary workaround had been implemented and it was now possible to transit up to 50 Mbps during off-peak hours. Peak hour performance had improved slightly to a few megabits per second, but further investigation was clearly required.

A number of connections were also showing CRC errors. These had been traced to certain port adapters on the Cisco routers, and seemed to be resolved by upgrading these to Revision 1.

Wilfried asked whether the cell discard problem was due to the switch architecture or the configuration. Roberto replied it was more likely to be the latter.

Wilfried then asked whether the configuration details could be made available to the group. Roberto thought this would be difficult, but he would talk to him off-line about this.

ACTION 6.1 - Roberto Sabatino

Howard reported that due to a price review and as compensation for delayed and/or badly performing circuits, a number of connections were being substantially upgraded. The ring between UK-Netherlands-Germany-France-UK would be upgraded to STM-4. Belgium would also be incorporated into the ring as it already traverses Brussels. In order to improve resilience, an additional 155 Mbps connection to Austria, and an additional 45 Mbps to Hungary would be installed. Existing connections to Italy, Sweden and Switzerland would be duplicated.

Even with these upgrades, the total cost of TEN-155 would be 10% less than projected. The European Commission had agreed the remaining money could be used to fund QUANTUM until October 2000, in order to cover the period between the end of the Fourth Framework, and the start of the Fifth Framework Programmes.

Around 160 million euros had been allocated in the Fifth Framework for research networking. This was split into RN1 - for production networks, and RN2 - for testbeds. The European Commission was having difficulty deciding how to allocate this money, as it was felt research networking was a special case (i.e. DANTE represents all EU/EEA NRNs), but existing regulations required an open call for proposals.

As a consequence, an independent Requirements Advisory Committee (RAG) had been established by the European Commission. As the RN1 funding was clearly intended for DANTE, the Commission officials hoped to get a decision from the Commissioners that NRNs should not have to bid for their own money. Instead, the RAG would produce a specification for the network, and would issue the Call for Tender. The membership of about 12 members included David Williams (CERN & TERENA), Tomaz Kalin (ARNES), Enzo Valente (INFN), Peter Villemoes (NORDUnet), Martin Wilhelm (DFN) and 2 representatives from ETNO (European Public Telecommunications Network Operators Association).

Cees asked whether the successor to TEN-155 actually required EU funding as they only covered about 30% of the costs. The commercial market moved quickly, and it may be possible to obtain better deals if NRNs did not have to involve themselves with the European Commission.

Howard replied the shortfall in funding would not be a problem to some NRNs, but others only received funding from their national governments if they were part of a EU project. Nevertheless, the NRNs were considering this option.

3. REVIEW OF EXPERIMENTS

3.1 Differentiated Services

Antonio presented the interim results of the DiffServ experiment (URL?). Tiziana added there had been problems with the SmartBits equipment, so they had tried to synchronise their test workstations using NTP. This had not proved possible as the workstation clocks were not accurate enough. They had therefore attempted to set-up connections to and from the same point in order to obtain accurate measurements. Unfortunately, this configuration had taken some time to debug as there were also problems with the Cisco and IBM routers, in addition to the underlying ATM infrastructure.

Tiziana also outlined the next phases of the experiment. The aim was to implement a pre-production service, and validate whether it could deliver bandwidth according to specified parameters. The experiment should now be extended beyond the initial three sites, with additional DiffServ implementations from other vendors (e.g. Nortel and Torrent).

Werner was interested in testing his implementation of DiffServ for Linux. He agreed to produce a plan for testing this on the overlay network.

ACTION 6.2 - Werner Almesberger

Roberto asked when the SmartBits equipment had to be returned. Tiziana was not sure, but she would contact Netcom to find out.

ACTION 6.3 - Tiziana Ferrari

Tiziana mentioned that a DiffServ BoF had been held during IETF46 where the results of the TF-TANT experiment had been presented. As a result, Rob Nielsen of BCIT and Joe Mambretti of iCAIR had expressed some interest in working with TF-TANT. Such collaboration might be useful as BCIT were working on a bandwidth broker, whilst iCAIR were developing digital video services.

3.2 Policy Control

Leon Gommans was not present at the meeting, but Cees mentioned he had been participating in the AAA Policy Working Group at the IETF. Some work had also been conducted with IP Highway, but this was subject to a non-disclosure agreement. Nevertheless, it might be possible to establish a private sub-group if TF-TANT members were interested in testing the IP Highway implementation.

3.3 IP over ATM

Roberto reported the previous two months had been spent fixing performance problems on TEN-155. The buffers on the Ascend switches had been misconfigured, with the default value of 80 cells proving insufficient. This value had been increased to 500 cells, and throughput had improved. It was hoped all the problems would be fixed by the end of the year.

Kevin thought it was unsatisfactory that KPN did not know how to configure their own equipment properly. TEN-155 would almost be finished by the configuration had been sorted out, and if this was likely to happen again in the future, DANTE (who had done a good job tracing the problems) would be better off managing the equipment themselves. Roberto agreed that DANTE could have done a better job, but this did not necessarily mean they should do it.

3.4 MPLS

Jean-Marc outlined the next phases of this experiment. The first activity was to investigate DiffServ mapping on MPLS, followed by evaluation of fast restoral. Finally, some interoperability tests should be conducted between equipment from Cisco, Nortel Networks and Juniper. Ericsson Telebit had offered their laboratory facilities for this.

Unfortunately, Jean-Marc would be leaving RENATER at the end of the year, so a new experiment leader would need to be found. In addition, the loan period for the Cisco equipment had expired.

ACTION 6.4 - Roberto Sabatino

Jean-Marc asked whether the loan period could be extended until April 2000. Michael was not sure this would be possible, but he agreed to try and arrange an extension.

ACTION 6.5 - Michael Behringer

Ladislav mentioned they had tried to run MPLS on their backbone network. This had not been successful as a number of routers had crashed, and routing information was not advertised outside of the MPLS cloud. It had been necessary to remove it after a short period of operation.

4. EXPERIMENT DISCUSSIONS (Parallel Sessions)

4.1 IP Version 6

Alex reported that ACOnet had upgraded their router software which had created routing instabilities. These problems disappeared after

downgrading. SWITCH had also experienced routing instabilities, and these appeared to be related to certain IOS versions. More involvement with Cisco was required to help solve these problems.

ACOnet, SURFnet and the University of Southampton were participating in this experiment using IPv6 over ATM. SWITCH and RedIRIS were using IPv4 tunnelling. JOIN (Cisco/Telebit/3Com) would be connecting shortly. Nevertheless, the experiment was not progressing as quickly as hoped, and it needed more participation.

In order to progress this experiment, ACOnet volunteered 1.5 FTEs for 4 months, DFN said 2 FTEs were being employed for 2 years, CESNET thought 2 FTEs could be made available until May 2000, RedIRIS volunteered 0.33 of an FTE until May 2000, INFN volunteered 0.25 of an FTE until May 2000, whilst DANTE thought they could spare a few hours per week.

It was desirable to test as many different IPv6 router implementations as possible. Routers from Telebit and Cisco were already in use, and 3Com would be shortly. MRTd and Zebra were also suggested. It would also be desirable to conduct tests across networks other than TEN-155, but one or more NRNs would need to provide transit to do this.

Wilfried thought some further investigation into multihoming problems was necessary. These were extremely difficult to track down because of the nature of the 6Bone, but he was looking for volunteers to help him with this. A request would be sent to the mailing list.

ACTION 6.6 - Wilfried Woeber

Wilfried also reported he had been unable to get DNSv6 to use anything other than IPv4 for transport. There were also other issues to resolve before DNSv6 could be used in a production environment, and he would summarise these on the mailing list.

ACTION 6.7 - Wilfried Woeber

Roberto asked whether a connection to the 6REN should be established. Kevin however, thought it was necessary to peer with the 6REN at the 6TAP (located at STARTAP), to which TEN-155 did not have a connection.

Lladislav asked whether IPSEC existed for IPv6. Alex replied this would be supported by Telebit by the end of Q1. Juergen added that some IPv6 implementations for the PC already supported this.

Juergen mentioned he was compiling a list of IPv6 applications, with

a view to running them between participants. He asked the group for input.

4.2 Flow-based Monitoring

Simon said there were some problems with some flow analysis tools. The type of data that could be extracted with cflowd was limited as it was only possible to produce two-dimensional tables. It was also difficult to measure traffic that used non-specified port numbers. Furthermore, the high-end Cisco routers only supported aggregated NetFlow statistics.

Roberto thought it might be worth investigating OCXmon. He asked whether anyone had any experience of this.

Simon asked who was interested in using accounting for volume-based charging. Avgust replied ARNES were considering this.

Roberto said that DANTE had to produce statistics for the four AUCS interconnects to ensure load balancing was working properly. This was being carried out by David Harmelin, but he would welcome input.

4.3 Multicasting (notes from Valentino Cavalli)

Robert presented a multicast route monitoring architecture. This was currently being implemented by Cisco, and was based on a Manager-Agent scheme. Agents monitoring multicast sessions ran at the edge of the network, and reported to the Manager. The software for Solaris 7.5 was available from http://imj.ucsb.edu/mrm/, whilst the protocol specification was available from http://www.ietf.org/internet-drafts/draft-ietf-mboned-mrm-00.txt

Robert proposed to test this architecture on TEN-155, and asked which NRNs would be interested in participating. CESNET, INFN, RENATER, RedIRIS volunteered for this, while Agnes would check to see whether DANTE could provide a workstation in London. DANTE were also asked to establish a mailing list for this activity.

ACTION 6.8 - DANTE

The Point-to-Multipoint SVC tests were discussed. Such SVC tests should ideally be initiated between the PoPs, but the Ascend switches currently only supported one SVC per port. It was therefore proposed to initiate them on local switches at the NRNs, and use tunnelling over TEN-155.

ACTION 6.9 - Robert Stoy

Robert also discussed the development of BGMP/MASC. The latest

implementations were available from http://www.gated.org/ and http://netweb.usc.edu/masc/mascd/. The IETF draft specification was available from: http://netweb.usc.edu/masc/draft-ietf-malloc-masc-04.txt

Robert would download and compile the latest version of BMGP/MASC, and produce a test description.

ACTION 6.10 - Robert Stoy

5. DATE OF NEXT MEETING

The next meeting will be held on the 3rd and 4th of February 2000. The venue would probably be CERN, Switzerland.

6. ANY OTHER BUSINESS

Cees mentioned that one of their professors had recently won the Nobel Prize for Physics. They would therefore like to take a high-quality video stream from the awards ceremony in Stockholm on Wednesday, 8 December. He asked whether it was possible to obtain guaranteed bandwidth from TEN-155 for this.

Howard replied bandwidth could be reserved via the MBS procedure, and Cees should contact Telia. This could usually be set-up very quickly, but the connection from the TEN-155 PoP to the Nobel Institute would also need to be investigated.

Roberto reported that as Kevin was leaving TERENA, Valentino Cavalli would take-over as Secretary of TF-TANT.

Roberto also thanked Celes, Esther and RedIRIS for hosting the meeting.

7. ACTIONS FROM LAST MEETING

5.1 DANTE to provide input to IPv6 report.
- Done.

5.2 DiffServ participants to provide input to Internet draft by 22/10/99.
- Done.

4.3 All IPv6 experiment participants to supply information about their available equipment, bandwidth and manpower.
- Ongoing.

3.8 Victor Reijs to draft document expressing the concerns of the research community about STM-4c.
- Status unknown. Victor was not at the meeting.

OPEN ACTIONS

6.1 Roberto Sabatino to discuss the TEN-155 switch configuration with Wilfried Woeber.

6.2 Werner Almesberger to produce a plan for testing DiffServ for Linux on the overlay network.

6.3 Tiziana Ferrari to contact Netcom to find out when the SmartBits equipment should be returned.

6.4 Roberto Sabatino to find a new leader for the MPLS experiment.

6.5 Michael Behringer to try and arrange an extension for the Cisco equipment loan.

6.6 Wilfried Woeber to ask on the mailing list for assistance with IPv6 multihoming issues.

6.7 Wilfried Woeber to summarise DNSv6 issues on the mailing list.

6.8 DANTE to establish a mailing list for the multicast experiment.

6.9 Robert Stoy to produce proposal for tunnelling point-to-multipoint SVCs over TEN-155.

6.10 Robert Stoy to produce test description for BMGP/MASC.

4.3 All IPv6 experiment participants to supply information about their available equipment, bandwidth and manpower.

3.8 Victor Reijs to draft document expressing the concerns of the research community about STM-4c.