



# Draft minutes of the 4th TF-TANT meeting (issue 2)



---

TERENA/DANTE TASK FORCE FOR TESTING ADVANCED NETWORKING TECHNOLOGIES

Minutes of the 4th TF-TANT meeting held on the 17th and 18th of June 1999 at INFN-CNAF, Bologna, Italy.

Kevin Meynell - Issue 2

PRESENT

Name	Organisation	Country
----	-----	-----
Werner Almesberger	EPFL	Switzerland
Michael Behringer	Cisco	United Kingdom
Marc van den Bergh	KPN Research	The Netherlands
Marjan Bogatinovski	KPN Research	The Netherlands
Zlatica Cekro	VUB/ULB	Belgium
Phil Chimento	Uni. Twente	The Netherlands
Howard Davies	DANTE	-
Tiziana Ferrari	INFN Bologna	Italy
Silvia Giordano	EPFL	Switzerland
Leon Gommans	U.Utrecht/Cabletron	The Netherlands
Ivano Guardini	CSELT	Italy
Christoph Graf (Chair)	DANTE	-
Mark Jansen	Uni. Utrecht	The Netherlands
Joop Joosten	CERN	Switzerland
Tom Kosnar	CESNET	Czech Republic
Simon Leinen	SWITCH	Switzerland
Silvia Matteoni	Uni. Bologna	Italy
Kevin Meynell (Sec)	TERENA	-
Paolo Moroni	CERN	Switzerland
Jan Novak	DANTE	-
Vaclav Novak	CESNET	Czech Republic
Simon Nybroe	Telebit	Denmark
Giovanni Pau	Uni. Bologna	Italy

Antonio Pinizzotto	IAT-CNR	Italy
Esther Robles	RedIRIS	Spain
Gianni Rossi	CSELT	Italy
Roberto Sabatino	DANTE	-
Guenther Schmittner	JKU/ACOnet	Austria
Jeremy Sharp	UKERNA	United Kingdom
Robert Stoy	RUS/DFN	Germany
Jean-Marc Uze	RENATER	France
Christina Vistoli	INFN Bologna	Italy
Bert Wijnen	IBM	The Netherlands
Wilfried Woeber	ACOnet	Austria

Apologies were received from:

Mauro Campanella	INFN Milano	Italy
Olav Kvittem	Uninett	Norway
Cees de Laat	U.Utrecht	The Netherlands
Ladislav Lhotka	CESNET/USB	Czech Republic
Vassilis Merikoulias	GRNET	Greece
Juergen Rauschenbach	DFN	Germany
Victor Reijs	SURFnet	The Netherlands

## 1. APPROVAL OF MINUTES

The minutes of the TF-TANT meeting held on the 29th and 30th of March 1999 were approved.

## 2. STATUS OF QUANTUM & TEN-155

Howard reported the connection to Portugal was currently being tested and delivery was planned for the 24th of June. There were however, some problems with the connections to Ireland and Luxembourg. The line to Ireland had been installed, but the ATM switch would not be installed for another four weeks. A temporary solution was being planned in the meantime. There were also equipment problems in Luxembourg which meant that connection had been delayed.

A contract had still not been signed with the Slovenian PTO. Negotiations had broken down and alternative suppliers were being approached. Whilst there was a telecommunications monopoly in Slovenia, this did not extend to data provision.

The TEN-155 Managed Bandwidth Service (MBS) was currently in the beta-test phase which was scheduled until the end of June. The SUSIE (Germany/Switzerland), EDISON (France/Germany), METHODIS (France/Germany), ENCart (Austria/Hungary) and TF-TANT (MPLS

experiment) projects were piloting this. This had raised a number of issues about the management chain, which involved up to three separate organisations (DANTE, local NRNs and local institutions). It was clear that unified management procedures would be required in future.

As a result of the upheavals in the European Commission, it was likely that the Fifth Framework Call for Proposals would be postponed from September 1999 until January 2000. The consequence of this may be problems with continuity of funding between QUANTUM and the next generation network.

The QUANTUM Policy Committee (QPC) had met on the 1st of June and had expressed a lot of interest in the QUANTUM Test Programme (QTP). They had decided that QTP deliverables should be peer reviewed before being published in order to validate the quality of the results. In addition, they had asked whether a joint workshop could be held with TF-TANT in October, so they could evaluate progress.

Wilfried said he did not have a problem with peer reviews themselves, but he wanted reviewers that would provide the group with constructive guidance. During the TEN-34 project, criticisms were only voiced after the test programme had been completed.

Kevin asked whether the QPC would find the peer reviewers, or whether the TF-TANT group was expected to do this. Howard believed this was the responsibility of the QPC.

### 3. FORMAT OF FUTURE MEETINGS

Christoph said the group had agreed to have two parallel sessions during the current meeting. There were too many experiments for them all to be discussed during a plenary session, and the parallel sessions would allow more time for thrashing out technical details. The group however, still needed to decide the optimum format of future meetings. One option was to have separate experiment meetings on one day, with a plenary meeting on the second day. Alternatively, separate experiment meetings could be held in a common location, or even in different physical locations.

Tiziana proposed that a general meeting could still be held, but individual experiment meetings could be arranged as necessary. Wilfried however, said there would then be very little incentive for participants to come to the general meeting. In the long-term, this might affect the work programme as the current format provided opportunities for receiving feedback and finding participants.

It was agreed that future meetings would have one day for separate

experiment discussions, followed by a joint meeting the next day for inter-experiment coordination and progress reporting. This would hopefully strike a good balance between the need to discuss technical details, and allowing participants to gain an overview of all the experiments.

#### 4. EXPERIMENT DISCUSSIONS (Parallel Sessions)

##### 4.1 QoS Monitoring

Philip discussed the different approaches to monitoring QoS. This could be conducted actively by inserting test traffic into the network (e.g. Surveyor equipment), or passively using techniques such as NetFlow. The parameters that should be measured included one-way packet loss, one-way delay variation, link utilisation, link bandwidth, EF commitments and EF reservation load. It was also suggested that interface discards, one-way packet delay and burst throughput should be measured. The goal was to define a meaningful measurement architecture that produced consistent traffic reports.

Tiziana asked what techniques could be used to conduct these measurements. Philip replied it might be possible to utilise the Surveyor or RIPE test traffic equipment. Another alternative was to use NTP, although this was not very precise.

Tiziana went on to ask whether the Surveyor equipment could generate TCP streams. Philip thought this was possible, but they not yet managed to get their GPS antenna to work. It was also possible to shape traffic although the equipment was not intended to be a traffic generator.

It was thought the Surveyor equipment would be more suitable for this experiment because the RIPE Test Traffic Project was aimed at commercial networks where confidentiality was important. Kevin said he had some contacts at Advanced Networks (who developed Surveyor) and would ask about the possibilities of loaning some equipment.

ACTION 4.1 - Kevin Meynell

Jean-Marc said it might be possible to loan more Smartbits equipment from Netcom. They had been quite helpful during the MPLS experiment, although they had just started a new activity. He would approach them about this.

ACTION 4.2 - Jean-Marc Uze

Roberto mentioned the University of Cambridge also had a measurement project. They had some prototype equipment they would like to trial.

It was agreed that a permanent DiffServ network should be established after the initial tests. This should not be too complicated and should not involve more than about five sites. A topology had already been specified in the experiment proposal.

Tiziana mentioned she was looking for someone to help her coordinate this experiment. Roberto volunteered for this.

#### 4.2 IP Version 6 (notes from Christoph Graf)

Simon N reported that Telebit had shipped one of their routers to the Amsterdam PoP, which was currently awaiting installation by DANTE (<http://www.tbit.dk/quantum/slides/tftantbologna/index.htm>). A pTLA had been obtained (3ffe:8030:/28, QTPVSIX) and plans for the MBS overlay were being finalised. They would start connecting participants shortly. There were also plans to connect to the 6REN in the future.

Simon N asked whether end-to-end tests should be conducted, and whether additional services should be tested in a coordinated manner. The group was in favour.

Wilfried proposed the experiment also cover address allocation, peering and deployment of IPv6 in existing infrastructures, tunnelling, DNS and reverse DNS.

It was agreed the first phase of the experiment should connect Austria and Switzerland to the 6Bone via TEN-155 during the first phase. 512 Kbps to 1 Mbps bandwidth would be adequate, but this should be on a permanent basis if possible. The remaining participants would be connected during the second phase of the experiment.

The following address usage was agreed: pTLA = 28 bits, NLA = 6 bits, SLA = 14 bits and subSLA = 16 bits. An AUP was not considered necessary as this was only a pilot service.

Simon N asked each participant to supply information about their available equipment, bandwidth and manpower. DANTE was also asked to create a mailing list for this activity.

ACTION 4.3 - All IPv6 experiment participants

ACTION 4.4 - DANTE

#### 4.3 Differentiated Services

Christina gave a presentation on the MONARC project. This was a distributed computing application for processing high-energy physics

data, and could potentially benefit from DiffServ.

Silvia asked whether the MONARC application was DiffServ enabled. Tiziana replied it was not, but DiffServ could be enabled on the routers.

Tiziana thought testing an application that was only used by a specialist user community would not be of much interest to most people. Leon however, said this was an extremely important project and worth considering.

Mark Jansen gave a presentation on the DiffServ tests conducted at the University of Utrecht (<http://www.phys.uu.nl/~jansen/diffserv.ppt>). Their results showed that bi-directional TCP streams from multiple source to multiple destination provided the best performance.

Werner Almesberger gave a presentation on the Linux DiffServ implementation. The latest version (<http://icawww1.epfl.ch/linux-diffserv/dist.html>) had a modular framework and supported traffic control, classification, scheduling and policing. Policy management, ingress policing and improved classification was still being worked on.

Tiziana asked about the scripts necessary to use Linux as a DiffServ router. Werner replied the scripting language was a bit messy at the moment, but help could be obtained from the mailing list (see <http://lrcwww.epfl.ch/linux-diffserv/> for details).

Tiziana then presented her proposal for the differentiated services experiment. The first phase was scheduled from the 21st of June to the 28th of August. There would be a single domain consisting of INFN, Stuttgart and Utrecht running Cisco and Cabletron implementations of DiffServ. Basic QoS measurements using Smartbits equipment would be conducted during this phase.

The second phase would involve multiple domains running different DiffServ implementations (IBM and Linux). This would investigate SLS definitions, LDAP policy servers, traffic shaping and policing, remarking and heterogeneous platforms. The third phase would test real applications over DiffServ (e.g. MONARC and video-streaming). The dates for these phases however, still had to be decided.

The requirements for each participant were a Smartbits box and a test workstation with NetPerf, mgen, ttcp and NetXpert installed. It was also necessary to configure NTP and a common account for SSH access. At least two sites would require a GPS antenna for the Smartbits boxes.

Silvia thought the proposal was a bit ambitious. Interoperability testing was very time consuming and needed a lot of support from the vendors. She suggested the number of domains should be reduced.

Simon L commented that SWITCH would not have a native ATM service until the end of August.

Philip mentioned that Torrent had just been acquired by Ericsson and were keen to participate in the experiment. Their DiffServ implementation was due in October.

Werner asked why addresses were allocated in blocks of ten. Tiziana replied this was for aggregation purposes. Simon L thought blocks of eight would be better in this case, especially if static routes were used.

Paolo said he did not want to use reserved addresses because he had researchers who wished to access the network. The network should be routeable even if it was still physically isolated.

Werner suggested the ISABEL video-conferencing application developed by the University of Madrid may be an interesting application to test over a DiffServ network. It had been ported to Linux and would be useful to most people.

Tiziana asked about participants unavailability during the first phase of the experiment:

CERN	Available all weeks
DANTE	Available all weeks
INFN	Weeks 29-30, 33-34
RedIRIS	Weeks 33-34
Uni. Stuttgart	Week 34
SWITCH	Week 29, 31
Uni. Twente	Week 30, 31
Uni. Utrecht	Available all weeks

Tiziana said further activities might include the introduction of other vendors (e.g. Nortel, Telebit and Torrent) into the DiffServ cloud, interworking trials between QoS-capable MPLS clouds and DiffServ clouds; testing of bandwidth brokers; and interconnection with other testbeds such as the QBone.

#### 4.4 Multicasting (notes from Christoph Graf)

Jan gave a presentation on the proposed TEN-155 multicast service (<http://www.dante.net/mbone/mcast99/mphase2.html>). The PIM Sparse Mode cloud would contain nodes in the Netherlands, Denmark and Sweden. NRNs were being encouraged to migrate to MBGP as parallel

DVMRP caused forwarding loops which could not be maintained in the longer term.

There were some interoperability problems between Cisco IOS 12.0 and 11.1.25cc. Participants were therefore requested to run IOS 12.0 (Wilfried thought 12.0.48 was okay as well). Other problems were that NRNs were not supplying much information to generate MBGP routes, DVMRP interaction did not work well, and multicast running over an ATM full mesh did not make efficient use of the underlying ATM infrastructure.

The group discussed the last problem, but were not very happy with any of the current solutions. Nevertheless, a proposal by Steffen Baur looked promising and would be tested.

[Note: the subsequent results of this test can be found at:  
<http://www.dante.org.uk/mbone/mcast99/mphase2.html>]

## 5. REVIEW OF EXPERIMENTS

### 5.1 ATM Signalling

Jan reported the service description was now available on the WWW (<http://www.dante.net/nep/dante-kpn/docs/1999/DK-99-21.html>). KPN had agreed to most requirements, and there was a lot of freedom to specify parameters. One problem though, was a lack of support for SBR1. This meant that bandwidth usage could not be restricted unless CBR was used.

The proposal for the acceptance tests was also available on the WWW (<http://www.dante.net/nep/dante-kpn/docs/1999/DK-99-16.html>). Some dates still needed to be specified, and the acceptance criteria still had to be agreed with KPN.

DANTE required some feedback from the experiment participants. Did they plan to use production interfaces on their switches; did they have a workstation with an ATM interface that supported LANE; and did they plan to use their own addressing scheme?

Jean-Marc and Guenther said they planned to use their own addressing schemes. Jan was also asked to include the KPN addressing plan (NSAP/AESA E.164) in the experiment proposal.

ACTION 4.5 - Jan Novak

### 5.2 Policy Control

Leon reported the research phase of this activity was about to



finish. The next stage was to develop plans for testing equipment using LDAP and policy management applications. A number of vendors had LDAP-compliant products, including Cabletron, Cisco, IBM, Nortel Networks, Telebit and Torrent. Any suggestions of other equipment that could be tested were welcomed.

### 5.3 Multicasting (IP and ATM)

Robert reported the TEN-155 multicast service was being migrated to MBGP. They hoped to move to a fully operational service in time for the Oslo IETF.

The ATM Point-to-Multipoint tests would be scheduled during the TEN-155 acceptance tests in September. The participants would be DFN/RUS, AConet, CESNET, INFN and RedIRIS. An experiment description was currently being drafted, and this could be discussed at the next meeting. The requirements at each site were a Cisco router with an ATM interface (AIP or ATM-PA3), an ATM switch terminating a VP to TEN-155, and a test workstation.

### 5.4 IP over ATM

Roberto said there was little to report since the last meeting. KPN had been informed by Ascend that a fix was now available for SBR/SCR and Y2K problems in their switches. Implementation of this fix would therefore need to be coordinated by the NOC.

### 5.5 Flow-based Monitoring

Simon L reported that DANTE had opened a test account on a workstation located at the TEN-155 PoP in Switzerland. This received a copy of the NetFlow accounting data from the router also located there. The software installed to analyse the data included cflowd, Cisco FlowCollector/FlowAnalyser and Fluxoscope (A NetFlow accounting tool developed by SWITCH).

The aim was to investigate how accurate traffic statistics could be produced for exchange points. This could help facilitate volume-based charging (such as the JANET scheme), detection of network abuse (e.g. smurf attacks), identify long-term trends (e.g. application mix, new applications, interesting source/destination networks), and detection of routing anomalies. The available tools would be described, compared and deployed where appropriate.

Simon N asked why multicast was not being used to collect the data. Simon L replied they had been unable to get it to work.

### 5.6 Route Monitoring

Simon L reported that IRRd (routing registry mirror server), mrt (multi-threaded routing toolkit) and the RAtools software had been installed on a workstation located at the TEN-155 PoP in Switzerland. This had been difficult in some cases where the programs had required root privileges. The next stage would be to set-up routing registry mirror servers so the actual TEN-155 BGP routing table could be compared with these.

## 5.7 IP Version 6

Simon N summarised the outcome of the earlier IPv6 session (<http://www.tbit.dk/quantum/slides/tftantbologna-plenum/index.htm>). The aim was to conduct end-to-end tests over 0.5 Mbps VCs, and to run some native services such as DNS. These activities would continue until the end of the QUANTUM project. It was hoped that other router vendors would also participate.

Bert mentioned that IBM had routers available for loan that were IPv6-compliant. CERN were responsible for distributing these.

Jean-Marc asked whether RENATER could connect using their own IPv6 addressing scheme. Simon N replied this would not be a problem.

## 5.8 RSVP to ATM Mapping

Tiziana said this activity would commence when ATM signalling was available on TEN-155.

## 5.9 MPLS

Jean-Marc gave a presentation on the MPLS trials conducted during the previous three months (<http://www.renater.fr/jmu/QTP/MPLS-tftant-180699/>). Configuration was straightforward and the stability of MPLS had improved since the previous tests. Performance had been tested with low bandwidth links using Netcom Smartbits 2000, ttcp, mgen, NetPerf and ping, but had consistently proved to be unacceptable. This required further investigation.

The next stage was to conduct more performance tests, and investigate PVCC configuration, Tag QoS, and management issues. This would probably happen in the Autumn.

Jean-Marc thanked everyone who had participated in this experiment, particularly Victor Reijs.

## 6. TF-TANT OVERLAY NETWORK

Roberto said the previous meeting had postponed the decision on

whether a permanent overlay network should be established for the TF-TANT experiments. There was a discussion of the problems experienced during the MPLS experiment, and whether a permanent overlay network would help future experiments. A possible topology was also presented.

Wilfried pointed out that logistical problems had recently cost them a lot of time. A switch had been misconfigured to send and receive cells on different interfaces. This may be due to the fact that different switch vendors had their own interface numbering schemes. It would also be useful to have a common agreement for loopback tests at management boundaries.

Robert asked about the response time for setting-up VPs on the Managed Bandwidth Service. Howard replied they aimed to fulfil requests within two hours. They did not however, have any control over the response times of local networks.

Wilfried said he would require more information about the MBS and the duration of the TF-TANT programme in order to put management procedures in place in Austria. Christoph suggested that questions be sent to the MBS Operational Mailing List.

It was concluded that the establishment of a permanent overlay network would not really help. Many problems were local to the NRNs, permanent connections sometimes incurred costs, and only one ATM traffic class would be available.

## 7. DATE OF NEXT MEETING

The next meeting was provisionally arranged for the 30th of September and the 1st of October to coincide with the next QPC meeting. The first day would be reserved for specific experiment discussions, whilst the second day would be a joint session with the QPC. The venue would probably be Amsterdam.

Kevin said that TERENA would be unable to host this event as their meeting room could only hold a maximum of twenty-five people. Schiphol Airport was suggested as an alternative.

## 8. ANY OTHER BUSINESS

Christoph thanked Tiziana and INFN-CNAF for hosting the meeting.

## 9. ACTIONS FROM LAST MEETING

- 3.1 Christoph Graf to circulate EU funding notice/disclaimer on the mailing list.
  - Done.
- 3.2 Christoph Graf to circulate a questionnaire about MBS requirements on the mailing list.
  - Superseded.
- 3.3 Tiziana Ferrari to incorporate bandwidth brokers into the Differentiated Services experiment proposal.
  - Ongoing, but nearly completed.
- 3.4 Roberto Sabatino to describe laboratory tests in the IP over ATM experiment proposal.
  - Closed. Unable to publicly disclose these.
- 3.5 Victor Reijs to send URL of document discussing the use DiffServ in conjunction with RSVP to the mailing list.
  - Done.
- 3.6 Roberto Sabatino to obtain approval for the TEN-155 multicast network from the QUANTUM Policy Committee.
  - Done.
- 3.7 Guenther Schmittner to ask KPN whether they will support PNNI.
  - Done. KPN would not be supporting this.
- 3.8 Victor Reijs to draft document expressing the concerns of the research community about STM-4c.
  - Ongoing.
- 3.9 All MPLS experiment participants to arrange ATM connectivity to TEN-155, complete MBS questionnaire, put diagram of local Tag architecture on the Web, specify delivery address for loan equipment, and sign Non-Disclosure Agreement.
  - Done.
- 3.10 Jean-Marc Uze to propose backbone infrastructure and VPN set-up, circulate the configuration document, and organise Smartbits tutorial.
  - Done.
- 3.11 Cisco to check the hardware configurations, determine software availability, and provide feedback on test plan.
  - Done.
- 2.1 Simon Leinen to contact Tiziana Ferrari to determine what activities could be included in the RSVP to ATM SVC Mapping Experiment.

- Done. No conclusion was reached.
- 2.2 Robert Stoy and Jan Novak to specify the TEN-155 multicast facilities.
  - Done.
- 2.12 Experiment Leaders to specify equipment requirements in their proposals as soon as possible.
  - Done.

#### OPEN ACTIONS

- 4.1 Kevin Meynell to contact Advanced Networks about the possibility of loaning some Surveyor equipment.
- 4.2 Jean-Marc Uze to approach Netcom about the possibility of loaning more Smartbits equipment.
- 4.3 All IPv6 experiment participants to supply information about their available equipment, bandwidth and manpower.
- 4.4 DANTE to create a mailing list for the IPv6 experiment.
- 4.5 Jan Novak to include the KPN addressing plan in the ATM signalling experiment proposal.
- 3.3 Tiziana Ferrari to incorporate bandwidth brokers into the Differentiated Services experiment proposal.
- 3.8 Victor Reijs to draft document expressing the concerns of the research community about STM-4c.

Contact: [nep@dante.org.uk](mailto:nep@dante.org.uk)

[DANTE](#) / [TERENA](#)

*[September 1999]*