Minutes of the 7th TF-TANT meeting held on the 3rd and 4th of February 2000 at CERN, Geneva, Switzerland.

Valentino Cavalli - Issue 2

PRESENT

Name	Organisation	Country
Werner Almesberger	EPFL	Switzerland
Michael Behringer	CISCO	Spain
Wim Biemolt	SURFnet	The Netherlands
Massimo Carboni	GARR	Italy
Valentino Cavalli (Secr)	TERENA	-
Phil Chimento	University of Twente	The Netherlands
Nicola Chiminelli	CSELT	Italy
Tryfon Chiotis	GRNET	Greece
Ehsan Chirazi	University of Brussels	Belgium
Tim Chown	University of Southampton	United Kingdom
Howard Davies	DANTE	-
Stefaan De Cnodder	Alcatel	Belgium
Larry Dunn	CISCO	USA
Francis Dupont	ENST Bretagne	France
John Dyer	TERENA	-
Hans Joachim Einsiedler	T-Nova Berkom	Germany
Tiziana Ferrari	INFN Bologna	Italy
Leon Gommans	U.Utrecht/Cabletron	The Netherlands
David Harmelin	DANTE	-
Avgust Jauk	ARNES	Slovenia
Joop Joosten	CERN	Switzerland
Dimitrios Kalogeras	GRNET	Greece
Tom Kosnar	CESNET	Czech Republic
Olav Kvittem	Uninett	Norway
Simon Leinen	SWITCH	Switzerland
Ladislav Lhotka	CESNET	Czech Republic
Jiri Navratil	CERN	Switzerland
Simon Nybroe	Ericsson Telebit	Denmark
Olivier Martin	CERN	Switzerland
Octavio Medina	ENST Bretagne	France
Paolo Moroni	CERN	Switzerland
Christian Mueller Boehm	Join/ University of Muenster	Germany
Antonio Pinizzotto	CNR	Italy
Agnes Pouele	DANTE	-
Herve Prigent	RENATER/CRIHAN	France
Victor Reijs	SURFnet	The Netherlands & HEAnet Ireland
Esther Robles	RedIRIS	Spain
Roberto Sabatino (Chair)	DANTE	-
Rina Samani	UKERNA	United Kingdom
Trond Skjesol	Uninett	Norway
Robert Stoy	RUS/DFN	Germany
Bernard Tuy	RENATER	France
Paul Van Binst	University of Brussels	Belgium

David Williams	CERN	Switzerland
Wilfried Woeber	ACOnet	Austria

Apologies were received from:

Bert Wijnen, Juergen Raus and Alex van der Plas.

1. APPROVAL OF MINUTES

The minutes of the TF-TANT meeting held on the 25th and 26th of November 1999 were approved.

2. UPDATE ON TEN-155

Roberto Sabatino reported that the T3 link from Vienna to Slovenia had been activated before the end of December 1999, whilst the T3 triangle between Spain France and Switzerland had been delayed until mid February 2000. The

AUCS interconnection was completed with 127 Mbps total capacity. DANTE had been working on the CRC errors and FCP discards. The problems were definitely solved before the end of December 1999. Roberto said that the transatlantic connection to Abilene is running now at 50 Mbps. DANTE is currently negotiating with canarie for a interconnection with Canet3, and a interconnection with ESNET is due really soon now.

The plans for the TEN-155 regard the STM-4 ring linking London, Amsterdam, Brussels, Frankfurt and Paris planned for July 2000; doubling the capacity out of Sweden, Switzerland and Italy is planned for October 2000. GX550 switches will be installed in those 5 cities, connected to CBX50 switches for low speed connections (<=45Mbps). Upgrades are planned to take place mid March for the links between Vienna and Hungary and between Vienna and Switzerland.

3. UPDATE ON QUANTUM /FP5

Howard Davies reported the position of the Commission with respect to the continuation of QUANTUM in the 5th Framework Programme for the years 1998-2002 (FP5). The Commission had to face the problem of insuring continuity of infrastructure in the total discontinuity between the FP4 and FP5. The situation had been clarified in a series of documents available on the DANTE site at http://www.dante.org.uk/geant/. The latest paper (http://www.dante.org.uk/geant/GEA-00-002.pdf dated 15 January 2000) is the report from the Requirement Advisory Group (RAG) defining the general requirements for RN1, the first action line of the IST programme for provision of infrastructure to Research Networks. Howard anticipated that the main requirements are to provide 2.5 Gbps as soon as possible in the core, and to ensure the upgrade to tens of Gbps by the end of FP5 in 2002. The new project should ensure continuity with QUANTUM /TEN-155, and, in addition, cover all countries included in the FP5. Interconnection should be provided to those countries that directly benefit from Commission funding (CH, CY, CZ, PO, HU, SI), as well to those countries that do not

get direct funding but are eligible to participate in the programme (BG, EE, LT, LV, RO, SK). Besides operational service, the new infrastructure should be used to provide capacity for "experiments with disruptive technologies". Howard said that these tests would be carried out on a separate VPN. The next steps for implementing the GEANT proposal are delivery of detailed specifications end of March 2000, extending the QUANTUM and Q-MED projects from May to October 2000, and signing the contract between the Commission and the new consortium by November 2000. Immediately after the contract signature open tenders will be issued for the creation of the 2.5 Gbps core and for links to the new countries. Howard reported that according to the latest draft of the IST work programme for the year 2000 four action lines for Research Networks projects have been outlined: RN1 to support operational service (GEANT), and RN2 to RN4 for network developments and testbeds. RN2 will deal with middleware and end-to-end experiments, RN3 will support deep technology developments, mainly regarding access and wireless networks, RN4 will

support experiments with alternative/disruptive network technologies such as optical networks, etc. Howard handed out a copy of the relevant section of the draft IST work programme to the participants. He anticipated that the call for proposal should be published by mid-February 2000 with estimated deadline in May 2000. He invited the TF-TANT group to start thinking about a set of coordinated experiments addressing the RN2 call. Each experiment should be proposed by a group of NRNs with a leading organization for each of them. If the QTP will be extended, plans would also be needed to justify the possible continuation of testing activity.

4. UPDATE ON EC /RAG

David Williams is a member of the Requirement Advisory Group for RN1. He had contributed to work out the Commission document dated 15 January 2000, already anticipated by Howard. David said that a formal decision on the content of the document, and therefore on the formal steps necessary to go ahead with RN1 was expected by May 2000.

5. REPORT ON ACTIVITIES SINCE LAST TF-TANT

Summaries of the main results of experiments were anticipated in the plenary about QoS, Policy Control and IPv6. In addition, David Harmelin reported that FCP discards experienced before Christmas were due to a standard configuration problem of the CBX500 FCP Module in TEN-155. All parameters like Available Cell Rate (ACR), Initial Cell Rate (ICR), RIF/RDF, local buffer congestion and discard threshold, were set to default values. The solution adopted to prevent FCP discards consisted in changing the buffer and discard profiles on the switches.

Victor Rijs would be interested in investigating what the NRN are doing at the national level in relation to WDM. He mentioned the availability of an international link with 60 Km of dark fibers between Belgium and the Netherlands. He said he had started a survey, but did not yet received enough responses from the NRN to be able to present a complete view at the meeting. Victor invited the group to get involved.

6. IPV6 TEST PLANS

Simon Nybroe briefed on the status of the IPv6 experiment. An IPv6 network with native ATM PVC is now operational, and the group is currently collecting routing information. A to do list was distributed recently but a few replies were received.

The network is stable enough to start the experiments. According to the plan outlined by Alex van der Plas IPv6 experiments will be carried out in four areas: interoperability, DNS, multihoming and applications. Wilfried presented the plan in detail. Four groups would be created with participants active in at least one area, taking responsibility for one deliverable or report. He also invited the group to work in their own environment and to report to others their experiences.

Wilfried said that a message by Harald Michel providing information about the maximum IPv6 "ping size" had been posted to the qtp-v6@dante.org.uk list. Then, he reported about recent successful interoperability testing, like load sharing with two serial lines. Various Host based routing implementations were tested: both MRTD and Zebra worked fine with BGP4 and RiPng, but Simon Nybroe observed that one needs to better investigate into

peer OSPFng. The comparison between MRTD and Zebra shows that they have similar configuration language (CISCO style), but Zebra uses a modular configuration file approach. The GATE-D implementation will not be tested because it is too expensive.

There was not much progress on the DNS area since the last TF-TANT meeting, mainly because the group is still waiting for a new record-types release of Bind a.x. Investigations are planned regarding IPv6 capability in resolvers (including transparent resolvers), forwarding and recursion, tree-walk, and coupling of stateful and stateless auto-configuration.

Wilfried said that the group is looking into the deployment of both PTLA and STLA addresses. He mentioned a draft presented by jessica Ju at a meeting in Japan. Simon said that there were many useful presentations and that the whole documentation of the meeting is interesting for the group. The proposed plan for the following phase of multihoming testing envisages to use QTPv6 prefixes from existing PTLA specification to/from AS8933, and to use STLA for NRN to NRN peering. The group should start building a pilot by initially setting up the routing cloud, investigating into the routing registry aspects and discuss transit to provide global connectivity. Many applications are available for IPv6, among these Ping, Traceroute, Telnet, SSH, Free BSD like Mozilla, Apache, Qmail. Quake was recently announced too. It was not clear if any such application was available for Linux, however they work fine with Solaris 8.

An IPv6 Policy document had been distributed on the mailing list Ipv6-wg@ripe.net. The document reviews the rules for assignment of STLA and TLA by way of the RIRs. Updates regarding DNS included RevDNS, update

http://www.isi.edu/~sekiya/IPv6/DNS.html, and the release of IPv6 enabled Linux.

Wilfried concluded his presentation by mentioning the need to develop coexistence and migration flow from v6 to v4 by working on NLA distribution and registry, tunneling technology and NAT protocol translators. Technical activity is also needed to upgrade IOS versions and install filters. He said he had presented a report at the IPv6 Forum Conference in Berlin in December 1999, describing the framework of the IPv6 test both in QUANTUM and ACOnet. A short paper will be submitted for the TNC 2000.

ACTION 7.1 Simon and Wilfried to submit a short paper on IPv6 for the TNC 2000.

Christian Mueller Boehm, from JOIN IPv6 lab said that they have IPv6 router from Telebit TBC2000 connected to 6Bone and to QTPv6. They plan to start interoperability tests with other routers, particularly CISCO 4500, 3COM Pathbuilder and NetbuilderII and the Nortel BLN. The test should be carried out on various operating systems. They invited the TF-TANT IPv6 group to participate in the tests. Simon Nybroe reported a conversation he had with Tiziana about the IBM routers loaned for the Diffserv experiment. They do not perform properly for the Diffserv experiment, but are IPv6 capable, this means that they might be used in the IPv6 experiments. He added that he had sent out a questionnaire for people to participate in the tests, but

so far he had received answers from ACOnet, Ericsson and JOIN only. A matrix would be filled in for all participants to chose their favoured area.

ACTION 7.2 All, except GRNET, INFN, RedIRIS to respond to the questionnaire.

7. MULTICAST

Robert Stoy reported that "places all over TEN-155" is announced via sdr and that CENSET, INFN, REDIRIS, DANTE and RUS are transmitting on this session.

Mtrace statistics are updated every ten minutes, showing varying levels of packet-loss from zero to ten percent and up to 20 percent. More information is available at http://www-ks.rus.uni-stuttgart.de/TF-TANT/. It was observed that MRTG could also be used to monitor packet loss, but Robert remarked that it would monitor a single router and therefore would not provide information on the multicast topology.

Results about the activity based on Multicast Routing Monitor (MRM) are not available yet. A beta release of MRM is being used at RUS in a configuration linking a MRM manager to several MRM testers running on different platforms. The MRM code is only available for Solaris 2.7. MRM test-senders send TS-requests to the network about parameters like packet ID, packet size, multicast group, port number. MRM test-receivers produce TR-reports providing timestamp, TS-address and TR-address, DIL. Robert said that one test session only is possible for each multicast group; to be able

to monitor multiple MRM sessions one have to have multiple multicast groups. He was asked to investigate about MRM test on routers, especially in relation to MRM manager producing loss graphs. Routers to be used in the experiments can be provided by Dante, GRNET and CESNET, Renater and RedIRIS. Discussion about the experiment plan will be carried out in the mailing list qtp-mc@dante.org.uk

ACTION 7.3 Robert Stoy to coordinate MRM test on routers and publish results

A third activity of the Multicast group is related to testing BGMP/Masc implementations. Robert said that the latest version of Masc demon has not been tested yet. The group would be interested in investigating point to multi-point SVC.

8. WRED

Nicola Chiminelli presented the WRED experiment carried out at CSELT with the collaboration of INFN. They have been using CISCO 7500 at CSELT and 7200 at INFN in a PVC at 3.5 Mbps. Traffic was generated by 2 Sun Solaris workstations and collected by a Linux machine. The testbed was intended to evaluate Per-Hop Behaviour Assured Forwarding with different priorities in a test environment with maximum number of 26 TCP streams and background UDP stream (best effort) at constant rate of 1.5 Mbps. The results are divided into two types of traffic: balanced traffic and non-balanced traffic.

ACTION Nicola to publish his presentation. Presentations are available at

http://carmen.cselt.it/papers/index.html

Nicola remarked that there is a need to reduce the complexity of the problem. The discussion was resumed later in the meeting and Nicola said that the next thing to be verified is the RTT dependency. The current RTT with INFN is 8 milliseconds. The plan is to have a new testbed with one new destination for 100 milliseconds to generate TCP flows towards each

destination.

9. MPLS TEST PLANS

Herve Prigent is the new leader of the MPLS group taking over from Jean Marc Uze. Herve listed the test areas for the new set of experiments: a) fast restoral with new version of IOS, b) traffic engineering, c) Diffserv mapping on MPLS network, d) interoperability of software from different vendors, e) investigating new software functionality of VPNs, f) MPLS over non-ATM networks. The plan is to prepare the test during February 2000, setting the MPLS network in March and carrying out the experiments in April. For the experiments, four Smartbit testers are available until May 2000 at INFN, University of Twente, University of Utrecht and Renater. It was remarked that according to the plan the MPLS experiment will overlap with the Diffserv experiment, and therefore the detailed plan should very carefully take into account the availability of routers. Herve needs to

fill in a matrix of NRN interested in working in the specific tests.

ACTION 7.4 Herve' to chase matrix completion.

The discussion helped in the initial definition of people interests. CERN, INFN, University of Twente and GRNET expressed their interest, RUS would be interested but need to investigate their equipment, and CESNET would like to test multicast over MPLS. Joop recommended to take into account in the new plan those tests that were not finished by Jean marc, and to address in particular those which had failed.

Ladislav reported that CESNET continues the attempts to migrate their production backbone to MPLS. They experience problems both in the stability of the core switches and routers and in the interface with metropolitan area networks, especially regarding IP addressing. Concerns about the rationale of continuing to focus on MPLS over ATM were raised, since MPLS should be transparently used and ATM switches are not being used as LSR. Herve considered the opportunity to test MPLS functionality without ATM.

10. EFFECT OF DIFFERENT PACKET SIZE ON RED PERFORMANCE

Stefaan De Cnodder presented results of Random Early Detect (RED) experiments carried out at Alcatel. The queue management algorithm is based on the specifications from Floyd and Jacobsen (1993). The behaviour of RED depends on the implemented algorithm. Stefaan compared the IST variant, were drop probability is independent of packet size, and the ZND variant, were drops depend on the packet size. Due to this feature ZND is useful for telnet packets. However this version can give very severe throughput collapse. Stefaan said that most vendors had implemented uniform drop probability RED versions, but they are likely independent of the packet size. More information is provided at:

http://www.alcatel.com/crc/publi/topic_7.htm.

11. FLOW MEASUREMENT (from draft Notes by John Dyer)

The goals of the meeting are mutual update on Local Activities and agreement on the next steps, basically access to test environment and information collection.

Simon Leinen said that SWITCH have got two workstations co-located at La Poste in Geneve, one used for production activity and the other for test. The first workstation sends one copy of CflowD to purgatory and another to

the second workstation (which can be fanned out even further if necessary). They want to find people to commit to doing some serious tests, provided that they have access to the workstations. Simon said that in agreement with DANTE, SWITCH can provide SSH remote access. People that need access should contact him. Tom said he will think about installing his software on one workstation. Jiri said he will try to install his software on the test

one, because he is interested in seeing how a system designed for a large end-user site hold up on a backbone.

Jiri reported the activity at CERN. They started to work with Netflow by converting data into the form that was used before and run old statistics, however that was found to be too heavy. Now they use Netflow directly with an update interval of 5 minutes. Results about analysis of packets, bytes and flows by source and destination are accessible via the web. On-line statistics of IP traffic at CERN are not yet in production, but will be made available at http://sunstats.cern.ch/.

Avgust said he thought that CISCO would not be supporting full Netflow data on interfaces running at 2.5 Gbps and above.

Olav reported about results from sampling comparable to the netflow full results. In those cases, a random distribution of the packets permits quite a good representation of the data. The most interesting traffic is seen in the US line.

David said that DANTE have been monitoring Netflow packets at regular intervals to get an indication of full traffic. A 50% more or less correct picture resulted out of one hour monitoring (10% error possible). However the situation gets worse as the sample interval gets longer. Sampling one packet out of 100 will never get accurate picture independently of the monitoring time.

Tom Kosnar gave a demonstration of his package. CESNET wants to analyse traffic on each border router from each customer into the cloud. They can choose time interval (time or calendar), and many other parameters with various levels of precision. Traffic generated by several cache engine farms appears to be about 18% of the whole traffic being monitored.

Simon said that they want to identify the common requirements and problems and feed them back to CISCO. At the moment they are working in close collaboration with CAIDA.

David said that with the Netflow Statistics Visualiser DANTE can look at last hour traffic matrix. But they have main problems when there are some asymmetry and no routing data are available for a particular AS. It would be possible to receive data from another AS and not knowing whom to account it to.

Simon said that experiment planning would be left to the mailing list. As head of the list, he would get the applications specific enough to map into the configuration of the packages. Some of the applications are simple to solve with CflowD. The web site contains descriptions of the applications, lists of RFC's pointers to useful software and a list of projects.

David agreed to set up a mailing list at DANTE to discuss these issues.

ACTION 7.5 David Harmelin to set up a mailing list for flow measurement.

ACTION 7.6 Leon and Simon to get information on contacts in Cabletron on LFAP

12. POLICY CONTROL TEST PLANS (from notes by John Dyer)

Leon Gommans leads the activity of the Policy Control Group. This mainly regards configuring QOS parameters and provision of QOS over a Diffserv infrastructure. The group is composed of the University of Bologna, INFN, CTIT at the University of Twente, WFI and the University of Utrecht. On 19th January the group held a first meeting to discuss proposals about deployment of an Open Policy System (V2.0 IP Highway), a Vendor independent Policy Management Solutions for QoS provisioning, involving 3com, CISCO and others.

The GUI talks to the policy administrator and keep policy in a Directory (such as LDAP). The Policy server talks to the policy enforcers (the routers); it may also have backup policy servers. The components can either be implemented on a single machine or different machines. The policy definitions are made through the GUI, which can even group policies based on interface or device types. GUI works on NT Win95/98, Admin NT, Server NT or Solaris.

The short Term Objectives of the Test Plan are to investigate policy based management solutions which will ultimately help the group to take some of the complexity of managing a Diffserv infrastructure. A prerequisite is a Diffserv capable router infrastructure. The Long Term Objective is to look at the work going on in AAA based policy control. Thinking about having something like a session key that gives access to all the things that the policy allows and then pushes this authority out to the equipment on the network. A lot of this will be administered on the edges of the network. Local authorisation at the edge might give you authorisation to get the aggregated access to the backbone resources.

Howard Davies expressed his concern about the actual implementation in a real service environment. In reality a single end to end connection involves 5 management responsibilities (Uni1, NRN1, DANTE, NRN2 Uni2), therefore special attention should be paid to Multi-domain aspects. Leon said that he will need to negotiate access to the Diffserv routers in order to do the testing. Simon and Olav added that Switch and Uninett would undertake some testing.

13. DIFF-SERV TESTING RESULTS AND PLANS

Hans Joachim Einsiedler presented the DISCMAN project (DIfferentiated Services network Configuration and MANagement), funded by Eurescom

http://www.eurescom.de. They want to cooperate with TF-TANT by sharing information about results of experiments, in particular relating to the project tasks task 2 (service models) and 3. The latter is about testing single providers scenarios, testing commercial platforms and open-source platforms. The task also foresees a second phase with the provision of a

more sophisticated testbed. Hans invited people from the TF-TANT group to attend the next project meeting that will be held on 27-29 march 2000. More info will be provided to the TF-TANT list.

ACTION 7.7 Hans to provide information about the next DISCMAN project meeting.

Olav Kvittem said that Uninett have 4-5 routers (CISCO 70200 and 4700) in three different locations in Norway and a 2 Mbps link between Oslo and Stockholm. The testbed is not yet configured, and no particular test has been planned yet. In addition, Olav said that they should set up a policy server later in the year 2000 to carry out Diffserv and Policy control testing.

Tiziana reported her discussion with Roberto and Victor about the use of diffserv for replacing missing ATM connections at SURFnet. They need to test the 622 MBPS interfaces towards the DANTE production network before it can be used. The test would be carried out at the Universities of Utrecht and Twente. They would start as soon as possible with the aim of being ready on the 1st of June 2000.

Werner Almersberger reported results of diffserv measurement on Linux. The test result regarded TCP, UDP with EF and TCP, UDP with no EF. Details are available at

ftp://icaftp.EPFL.ch/pub/people/almersberger/slides/quage00.ps.gz. EPFL are interested in things happening at the lower level, especially the interaction between flows.

Tiziana reported that in the last two months her group had been tuning parameters in the routers, in particular by concentrating on transmission queue, service rate and queue size for Expedited Forwarding queuing. On CISCO 70200 they experienced improvements with regards to both delay and jittering. They have been testing scheduling algorithms for priority queuing on both the INFN ATM-based MAN and the INFN-University of Twente WAN in different network scenarios. Tiziana said that she plans to finish the ongoing testing with EF very soon and to start testing with real applications like video conferencing. In the future they should focus more on AF, but that work needs to be done from scratch. Octavio Medina remarked that some test still needs to be done with WRED. Tiziana said the interoperability test is critical, and expressed her concern about continuing with the IBM equipment.

Experiments were progressing since the last TF-TANT meeting. In Italy tests had been carried out with two Sun Ultra (Solaris 2.7), Smartbit router, CISCO C7200 and C7500 routers, and ATM switches. The experiment compared one way delay versus packet size. Problems were detected with background Best Effort traffic. Tiziana observed that with priority queuing it is expected to observe less delay. According to the experiments, without Best Effort traffic there is no difference, but with Best Effort traffic a huge increase of delay is experienced.

The EF test behaviour had been carried out in a real WAN architecture with

multiple aggregation points and multiple congestion source at INFN, CERN, GRNET and the University of Utecht, The EF was tested with different input rates. EF should not have packet loss, however, loss of packets was experienced in the test. CISCO is currently looking into the problem. Additionally, increase in the number of EF streams (aggregate) and the aggregate bandwidth were detected in the experiment.

Tiziana mentioned some International collaboration activities: a proposal with the British Columbia University, and other with ESNET, in particular

about a Diffserv project. She announced that she would participate to the Internet2 QOS workshop in Houston, Texas on 9-10 February 2000.

14. DATE OF NEXT MEETING

The next meeting will be held on the 17th and 18th of April 2000. The venue would be ARNES, Slovenia.

15. ANY OTHER BUSINESS

No other business.

16. ACTIONS FROM LAST MEETINGS

6.5 Michael Behringer to try and arrange an extension for the CISCO equipment loan.
- ongoing

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

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

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

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.
- Ongoing.

OPEN ACTIONS

7.1 Simon and Wilfried to submit a short paper on IPv6 for the TNC 2000.

7.2 All, except GRNET, INFN, RedIRIS to respond to the questionnaire about IPv6 experiments.

7.3 Robert Stoy to coordinate MRM test on routers and publish results.

7.4 Herve' to chase matrix completion.

7.5 David Harmelin to set up a mailing list for flow measurement.

7.6 Leon and Simon to get information on contacts in Cabletron on LFAP.

7.7 Hans to provide information about the next DISCMAN project meeting.