

TERENA/DANTE TASK FORCE FOR TESTING ADVANCED NETWORKING TECHNOLOGIES

Minutes of the 8th TF-TANT meeting held on the 17th and 18th of April 2000 at ARNES, Ljubljana, Slovenia.

Valentino Cavalli - Issue 2

PRESENT

Name	Organisation	Country
-----	-----	-----
Werner Almesberger	EPFL	Switzerland
Michael Behringer	CISCO	Spain
Valentino Cavalli (Secr)	TERENA	-
Nicola Chiminelli	CSELT	Italy
Ehsan Chirazi	University of Brussels	Belgium
Tim Chown	University of Southampton	United Kingdom
David Crochemore	RENATER	France
Howard Davies	DANTE	-
Hamad el Allali	Univ. of Twente	The Netherlands
Tiziana Ferrari	INFN Bologna	Italy
Leon Gommans	U.Utrecht/Cabletron	The Netherlands
Mark Janssen	U.Utrecht	The Netherlands
Avgust Jauk	ARNES	Slovenia
Joop Joosten	CERN	Switzerland
Shai Herzog	IPHighwais	USA
Dimitrios Kalogeras	GRNET	Greece
Vadim Kotov	JINR DUBNA	Russia
Tom Kosnar	CESNET	Czech Republic
Olav Kvittem	Uninett	Norway
Simon Leinen	SWITCH	Switzerland
Ladislav Lhotka	CESNET	Czech Republic
Octavio Medina	ENST Bretagne	France
Paolo Moroni	CERN	Switzerland
Antonio Pinizzotto	CNR	Italy
Agnes Pouele	DANTE	-
Juergen Rauschenbach	DFN-verein	Germany
Victor Reijs	SURFnet	The Netherlands & HEAnet Ireland
Roberto Sabatino (Chair)	DANTE	-
Rina Samani	UKERNA	United Kingdom
Christian Schild	JOIN Projekt (Uni-Ms)	Germany
Robert Stoy	RUS/DFN	Germany
Bernard Tuy	RENATER	France
Jean-Marc Uze	Air France	France
Alex van der Plas	Ericsson Telebit	Denmark
Strg Venas	Uninett	Norway
Wilfried Woeber	ACOnet	Austria

Apologies were received from:

Esther Robles	RedIRIS	Spain
Herve Prigent	Crihan	France

1. APPROVAL OF MINUTES

The minutes of the 7th TF-TANT meeting held on the 3rd and 4th of February 2000 were approved.

2. FLOW MEASUREMENT AND ANALYSIS

Simon expressed an interest in knowing the plans of other NRNs and would have liked a presentation session. He reported that the LFAP protocol (RFC 2124) differs from the Cisco Netflow because LFAP is TCP based, the accounting information is sent through intermediate periodic updates, and has its own accounting software. Simon said that he is looking for solutions to provide accounting information to the universities in Switzerland, and invited the participants to get in contact with him.

ACTION 8.1 Simon to check who are using Cabletron routers in their networks.

Simon reported that Netflow data are used to determine the type of application traffic. He said that traffic on the US line amounts to 70% of all traffic, and remarked that passive-mode ftp, used by Netscape and other applications, is not detected, whereas traditional active-mode ftp is detected. H.323 traffic and streaming media are also detected, in addition to non-better identified UDP flows. It is possible to detect a huge ICMP traffic in connection to Denial of Service Attacks and the originating network can be detected, but not the actual source.

Simon mentioned the FlowScan tool, developed by Dave Plonka from the University of Wisconsin, which integrates a cflowd patch and is able to detect NAPSTER and other types of applications. A pointer is available at <http://www.switch.ch/tf-tant/floma/software.html>.

Simon reported that Cisco had developed a new architecture based on Forwarding Information Base (FIB) and that he had implemented it on one of the SWITCH routers connected to an Akamai server in monitoring the amount of traffic towards specific NRNs.

Finally, Simon said that SWITCH have a lot of Flow Monitoring software installed in Switzerland and they are now looking for people to participate into their usage experiments.

3. IPV6

Alex reported that 14 NRNs are involved in the IPv6 test. Updates of activities carried out in the last two months were the following. SURFnet reported that they are investigating with Cisco about some BGP problems. RENATER deployed native IPv6 inside the ATM backbone. GRnet intend to transfer all services from IPv4 to IPv6 but there will be no progress before the end of the year. They still do not have any group working on IPv6. ACONET reported that they spent lot of effort on the DNS issue and finally decided to upgrade the IOS version to 12.0 based. Uninett are carrying out interoperability testing with Cisco router and Linux

workstations as well as in DNS testing; they plan to move to multihoming soon. SWITCH are using a slow old Cisco router but want to connect to 6Bone; their BGP connection is stable and small parts of their network can do multi-protocol IPv4 - IPv6; if there is time available they plan to experiment 624 encapsulation and prefix. CESNET are currently disconnected from the IPv6 network due to some problems with the "thanksgiving" release; they have been carrying out interoperability test with PC demons Zebra MRT on Linux versus Cisco routers and would like to coordinate the tests with other people. DANTE activity was focused mainly on DNS tests carried out on Solaris workstations in the Netherlands and on IPv6 enabled applications that have been tested with different Operating Systems in the offices at Dante. JOIN have their labs connected via native IPv6, they support multihoming with both 6Bone and QTP prefixes and have large DNS server for 6Bone; finally, they have been testing applications on behalf of universities requesting to migrate to IPv6. At CERN they have both connection to QTP6 and, via tunneling to a router in Chicago, to 6tap.

Christian said that JOIN is working on interoperability. He reported some problems in dealing with large packets and mentioned that some Cisco machine went down. He said that similar problems were also detected by SURFnet. Uninett and CESNET have gained some experience in that and should contact him to coordinate. In general, known problems should be reported by all participants to JOIN.

ACTION 8.2 all, report known problems on IPv6 interoperability to Christian Schild.

With regard to DNS activity Wilfried reported that DANTE had been working on testing Bind9 beta-1 and beta-2 basic features and that AConet and the University of Southampton had plans to collaborate with them in reverse DNS lookup experiment. SWITCH had offered to provide transit for the QTPv6 network, but there was a need for an official authorization. He also said that the IETF had been discussing about changing the IP6.int domain for reverse lookup, but there was no formal decision yet.

Ongoing IPv6 activity at AConet are related to working with Bind9 beta-2 and set up of DNS at the University of Vienna. Progress in working with bind are regularly reported. Simon remarked that a beta-3 release of bind would be available in a few months, even if it was not been announced yet. Furthermore, AConet is integrating IPv6 into (pre) production status and is now involved in issues still related to bind, including connectivity and transit. In the near future they will focus on multihoming and routing registry development. Wilfried also mentioned a project to deploy ADSL technology to deliver cable TV in the Vienna area to 20-30,000 users, and their plan to use IPv6.

Tim mentioned the "6INIT" EU project that the University of Southampton is carrying out together with BT, Telenord and Telia. They have 4 clusters in the UK, Germany, France and Scandinavia for testing QoS and Voice over IP. More information is available at <http://www.6init.org>. The University is also involved in multihoming and testing of different Operating Systems and

routers. Tim said that the use of freeBsd routers in deployment of those homes proved to be cheap and effective. AConet only provided input on the experiments, Tim asked for contribution by other partners. Updates about Applications are available at <http://phoebe.urec.fr/G6/#LABEL5>.

It was discussed whether Tunnel Brokers were used in Europe and Bernard Tuy asked if there was an opportunity to have them in the IPv6 NRN backbones. Among the participants only Uninett had one TB, and Wilfried reported that maybe SURFnet had one too. Simon remarked that there would be drawbacks in having more TB because this generates a weird topology and therefore it would be just a nice solution for groups of users needing temporary connection. Bernard did not think this remark is relevant since Tunnel Brokers provide IPv6 connectivity to "leaf end users". So there is no risk related to routing problems. Moreover, TB is designed to provide temporary connectivity (as IPv4 dialup is today) and therefore it does not impact the core backbone topology. Bernard said students developing IPv6 projects very often request Tunnel Brokers. They typically need IPv6 connectivity to check that what they did is Ok on a larger scale than the LAN (supporting IPv6), but they do not need a permanent connection to an IPv6 ISP.

Juergen said that in the near future there will be a need for 3.2 Millions IP addresses to support GPS and this is related to support for address planning. He mentioned a draft IPv6 distribution policy document by the IPNG Working Group and addressing issues that have been discussed at RIPE NCC.

4. TEN155 UPDATE

Roberto said that the T3 ring between Spain, France and Switzerland was operational since mid March, the STM-1 between Switzerland and Austria since end of March and the E3 for Ireland since 11th April. On the transatlantic side, connectivity with ESNET was upgraded to 50 Mbps, connectivity with Canarie to 30 Mbps since end March. He reported that additional STM-1 was still to be provided for DFN and SURFnet, as well as upgrades to 68 Mbps for Hungary and 155 Mbps for Greece. AUCS was planned to be upgraded to a monthly set at 127 Mbps. The 622 Mbps 5 city ring was planned for July 2000 and the 9 city setup, involving doubling the capacity to Stockholm, Milan and Geneva (with 2 STM1) was planned for October 2000. The full ATM mesh originally planned was offered by KPN Qwest, but it proved not to work effectively, therefore the 622 Mbit link will be delivered as four STM-1, out of which one will be used for ATM and three will be used for POS (IP). Due to this solution in the 9 city setup the POS ring will bypass Belgium. The NRNs should not move away from ATM, Roberto said that DANTE will continue to support the Managed Bandwidth Service and make it available to all network nodes. It is the sole responsibility of the NRN to decide whether to use the service or not.

Dante had evaluated both the Cisco 12K and the Juniper M40. Both are judged suitable for TEN-155, although the M40 is more performant and \ the 12k has more functionality. Other issues such as cost, maintenance, support and NOC integration would make the difference for a final choice. The evaluation of

these points was overtaken by contractual negotiations with KPNQ which forced KPNQ to provide and purchase the routers for TEN-155. KPNQ had already chosen Juniper for their own network and therefore these are also supplied to TEN-155.

No decision had been made yet about the routers for the upgrades to Stockholm, Milan and Geneva. As these are DANTE's responsibility, opportunities are still open for both Cisco and Juniper.

5. RN1 UPDATE

Howard briefed about the status of the RN1 call for proposal by the European Commission. According to the Fp5 it should be an open call for proposal with an independent evaluation, but on the other hand the EC wants the proposal to be published in May to be open to NRN only. Howard said that the position of the EC was still unclear and that a risk was that the evaluation might find that the GEANT proposal does not contain enough R&D. Actually this would be awkward because the GEANT proposal is about provision of infrastructure and not about R&D. In order to ensure continuity with Quantum Dante planned to issue a call for Expression of Interest in view of the future tender at the same time as the EC published the call for proposal, estimated on the 2nd May 2000 and actually published on 13th May 2000.

6. MULTICAST

Robert Stoy reported that MRM is being tested by Kevin Amelroth at the University of Southern California. In Europe the tests are in phase 1, test topology are being carried out by Rediris, DANTE, Cesnet, Janet and DFN (at the University of Stuttgart). A first test showed 10-20% packet loss. In the second test which was carried out in April between Madrid and Stuttgart the packet loss rate dropped down to 5-10%, however the statistics showed ambiguities because no report was done for a while. The test should be repeated.

Robert said that the workstations used in the experiment are available to other NRN. Simon said that SWITCH would be interested in using them. Who wants to join send message to Robert.

Regarding Interdomain Multicast Routing and BGMP Masc Robert said that the IETF proceeds slowly and that the Cisco MBGP version is implemented on TEN155. Roberto suggested not to pursue PIM -SM mapping to ATM point to multipoint SVCs as developments of TEN-155 suggest that the backbone is moving away from an ATM mesh. However NRNs may still find this topic interesting and if so they should say so and encourage to continue the work.

7. POLICY CONTROL

Shai Herzog founder and CTO of IPHighway was invited to hold a lecture on Policy Based Networking for QoS. The motivation for Policy Based Networking is that various classifications of services are made possible by Diffserv, Priority Queuing etc. This raises a need for policy management. There are different standardisation efforts being carried out at IETF, leading to protocols like RAP (Resource Allocation Protocol), COPS-Base, COPS-RSVP (RFC2784, . . . ,52) and COPS-PR, all of them dealing with the communication mechanism, PIB (Policy Information Base). Shai described SNMP configuration and IP Security Policy, then discussed various policy architectures and introduced the Integrated Policy Architecture based on a distributed Open Policy System. A copy of Shai's presentation is available at:

<http://www.phys.uu.nl/~lgommans/herzog/ljubljana/index.htm>.

8. DIFFSERV TEST

Tiziana asked Michael if it was possible to further extend the Cisco loan, but Michael said that an extension was already agreed in November and that a further one was not possible. A simple extension would be difficult also because the loans were separately managed by the NRN at the local level, with the exception of the equipment used in the MPLS experiment. The team should make a new coordinated request to Cisco, listing all the equipment on loan which has been used in the TF-TANT experiments and specifying the rationale for a new loan on the basis of concrete new testing plans for the next months.

ACTION 8.3 Michael to provide a full list of the equipment lend by the TF-TANT participants.

ACTION 8.4 Roberto to coordinate with Tiziana and to write a single proposal to Cisco for a new loan allowing the NRN to keep the same equipment for new planned experiments until the end of TF-TANT in October 2000.

Tiziana presented the update on the last months, starting with a repeated comparison between Priority Queuing and Waited Fair Queuing by increasing the number of Best Effort streams (Experiments 5b and 7b). Results show that one-way delay is different but jitter remains almost the same. Test performance of Expedited Forwarding with aggregation was carried out to test how burstiness changes. Tests on WRED were focused on the behaviors with different packet size. Comparison of burstiness with WFQ and PQ was done too. Antonio informed the participants that he had started to make experiments with three-color markers on Cisco routers to check rate/burstiness used in the output with WRED in a single network topology with two sources mtu 250 and mtu 1500 and 2 Mbps connection between the routers.

The next plans by Tiziana regard the validation of Expedited Forwarding with ICAIR, and possibly to try the configuration used in the past experiments in a real application configuration. A possibility is to work

with CERN to inject video streams to the Diffserv network using EF packets. Other experiments are planned with WFQ and WRED for Assured Forwarding, as well as about color-marking mechanisms and deployment of AF in production environment.

9. GRNET BW MANAGEMENT PROPOSAL

Dimitrios briefly summarized a proposal he had earlier submitted to the quantum mailing list and discussed it with the participants. The objective is to guarantee minimum bandwidth per group of customers and at the same time enforce maximum bandwidth in order to reverse usual "unfairness" towards small customers. There was some objection that enforcing maximum bandwidth when lot of bandwidth is available is unfair too. Dimitrios said that the project would be developed in different phases. In phase one 7 classes would be implemented plus one level for best effort using Class Based WFQ by Cisco. In phase two WRED might be implemented if needed. The following table summarises the envisaged steps.

	Step 1	Step 2
Signaling	BGP community	BGP community
Marking	BGP QOS	BGP QOS
Policing	--	CAR
Queuing	CBWFQ	CBWFQ
Dropping	--	WRED

There were questions about signaling mainly related to the fact that customers should not interface directly with the network and remarks about the lack of a monitoring mechanism. Statistics for the queuing per VC should be collected too. Other issues were related to QoS BGP accounting. Simon suggested to use CAR in marking instead than in Policing, however he would recommend to mark individual customers instead of groups of customers. Roberto remarked that the proposal should be refined further with the collaboration of people from other NRNs, including Simon Leinen. On the other hand the proposal scheme can be considered as a roadmap for other NRNs which might need to implement solutions to meet objectives similar to GRnet. Howard saw a lot of problems with the operational aspects, but Dimitrios remarked that many mechanisms like BGP QoS had been implemented already some year ago and have passed all testing, so there should not be any problem. Other NRNs use part of them already, even if Simon, who had been using CAR in the past, said that it has to be further tested before going to production. The final recommendation from Roberto was to use a three step approach: 1) to refine the technical proposal by incorporating the monitor mechanisms etc. 2) to implement the scheme in a test environment and 3) to work out a plan for deployment in a production environment.

ACTION 8.5 Dimitrios to work with Tiziana, Simon and Roberto on refining the proposal in the first week of May 2000.

10. MPLS

Agnes presented a report about the work done so far. Documentation for the test preparation is available at <http://www.crihan.fr/mplsdoc.html>, which contains IETF and Cisco documentation as well as references to MPLS networks. The first test was carried out as a tutorial at the Cisco lab in Paris about fast rerouting behaviour, BGP VPN configuration, traffic engineering and MPLS related to Diffserv. The fast rerouting test consisted in generating a link failure and checking the time needed to reroute the packets to an alternative path. The test showed a 40 seconds recovery time.

The test platform for the VPN/Traffic engineering was based on active routers running tag switching OSPF based on a physical infrastructure involving the UK, Greece, the Netherlands, Switzerland, Italy, the Czech Republic and France. The fast rerouting test was carried out by cutting links one by one in the nodes. Results were same as in the lab. It takes the same time also when you put up the link again. The VPN test scenario contains experiments on CE reachability between customers and PE connectivity as well as set up of green, red and blue VPNs.

Agnes ToDo schedule was to setup VPN green in week 16, 17, start VPN testing in weeks 17, 18 and then undertake Traffic Engineering and MPLS and Diffserv tests. With respect to the latter activity to be carried out in June 2000 Agnes asked support in finding out further documentation about Cisco 12-11a/t

ACTION 8.6 Tiziana to check if Cisco 12-11a/t version is Ok for the MPLS and Diffserv experiment.

ACTION 8.7 Michael to help in finding documentation about Cisco 12-11a/t version.

Lada said that MPLS had been running on the CESNET backbone for three weeks but did not make use of advanced features like VPN and Traffic Engineering. One year ago they had problems for 8 months, not specifically related to MPLS, but in that case they worked together with Cisco tech in Brussels and sent bug fixes and releases etc. Nowadays there are 10 towns in CZ connected. All unicast traffic is now transparent through this MPLS network, but some problems might be expected when going to encapsulated MPLS and back. They run route reflectors for BGP. They use transparent web caching in two sites and plan to use traffic engineering and classes of services. However VPN does not seem to have high priority.

11. DATE OF NEXT MEETING

The next meeting will be held on the 13th and 14th of July 2000. The venue would be Dublin, Ireland.

12. ANY OTHER BUSINESS

No other business.

13. ACTIONS FROM LAST MEETINGS

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

- superseded by new ACTIONS 8.3 and 8.4

6.9 Robert Stoy to produce proposal for tunneling 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

- DONE.

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

- DONE.

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

- DONE

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

- DONE

7.4 Herve' to chase matrix completion.

- DONE

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

- DONE

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

- DONE

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

- DONE

OPEN ACTIONS

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

- ongoing

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

- ongoing

- 8.1 Simon to check who are using Cabletron routers in their networks.
- 8.2 all, report known problems on IPv6 interoperability to Christian Schild.
- 8.3 Michael to provide a full list of the equipment lend by the TF-TANT participants.
- 8.4 Roberto to coordinate with Tiziana and to write a single proposal to Cisco for a new loan allowing the NRN to keep the same equipment for new planned experiments until the end of TF-TANT in October 2000.
- 8.5 Dimitrios to work with Tiziana, Simon and Roberto on refining the proposal in the first week of May 2000.
- 8.6 Tiziana to check if Cisco 12-11a/t version is Ok for the MPLS and Diffserv experiment.
- 8.7 Michael to help in finding documentation about Cisco 12-11a/t version.