<pre>
TERENA/DANTE TASK FORCE FOR TESTING ADVANCED NETWORKING TECHNOLOGIES

Minutes of the 9th TF-TANT meeting held on the 13th and 14th of July 2000
at University College Dublin, Dublin, Ireland.

Valentino Cavalli - Issue 1

PRESENT

Name	Organisation	Country
Werner Almesberger	EPFL	Switzerland
Kostas Anagnostakis	TERENA	-
Alain Bidaud	Crihan	France
Michael Behringer	CISCO	Spain
Marc Berenschot	Univ. of Twente	The Netherlands
Valentino Cavalli (Secr)	TERENA	-
Phil Chimento	Univ. of Twente	The Netherlands
Nicola Chiminelli	CSELT	Italy
Tim Chown	University of Southampton	United Kingdom
Joel Corral	ENST-France	France
Raffaele D`Albenzio	CSELT/Telecom Italia	Italy
Howard Davies	DANTE	-
Larry Dunn	CISCO	USA
Tiziana Ferrari	INFN Bologna	Italy
Ruediger Geib	Deutsche Telekom T-Nova	Germany
Leon Gommans	U.Utrecht/Cabletron	The Netherlands
David Harmelin	DANTE	-
Joop Joosten	CERN	Switzerland
Tom Kosnar	CESNET	Czech Republic
Olav Kvittem	Uninett	Norway
Simon Leinen	SWITCH	Switzerland
Ladislav Lhotka	CESNET	Czech Republic
Octavio Medina	ENST Bretagne	France
Christian Mueller-Boem	JOIN-DFN	Germany
Mike Norris	HEAnet	Ireland
Jan Novak	DANTE	-
Antonio Pinizzotto	IAT-CNR	Italy
Herve Prigent	Crihan	France
Victor Reijs	SURFnet	The Netherlands & HEAnet Ireland
Esther Robles	RedIRIS	Spain
Roberto Sabatino (Chair)	DANTE	-
Robert Stoy	RUS/DFN	Germany
Alex van der Plas	Ericsson Telebit	Denmark
Franz Widhofner	Univ. of Linz/ACOnet	Austria
Wilfried Woeber	ACOnet	Austria
Mirjana Zafirovic-Vukotic	Univ. of Twente	The Netherlands

Apologies were received from:

Cees de Laat	Univ. of Utrecht	The Netherlands
Juergen Rauschenbach	DFN-verein	Germany
Rina Samani	UKERNA	United Kingdom
Bernard Tuy	RENATER	France

1. APPROVAL OF MINUTES

The minutes of the 8th TF-TANT meeting held on the 17th and 18th of April 2000 were approved.

Michael reported the status of the Cisco loan. At the previous TF-TANT meeting it was suggested to coordinate the different loans at a single unified European level, but Michael said that the plan was changed. The loans are still managed at the level of individual NRN and individual country Cisco representatives. The loan is globally extended until the end of TF-TANT, however the individual account managers are responsible for it and possible requests to buy the equipment after the end of the experiments should be addressed to them.

Roberto announced that DANTE will discontinue Managed Bandwidth Service to the US PoP due to staff shortage and therefore they have decided to focus on straightforward IP commodity and interconnections to North American Reserach networks.

The Dutch universities will be disconnected from the Diffserv network at the end of July. Tiziana asked people to start planning how to reuse the two Smartbit boxes used by the University of Twente and the University of Utrecht. People from Twente said they will keep their router and reuse it in the MPLS test. The router at Utrecht is available for other organisations. In addition, Tiziana said the IBM boxes are IPv6 capable and might also be reused in other experiments. However it was remarked that IPv6 support is limited to Ethernet only. Also, the IBM router is now unsupported so it doesn't make sense to invest further efforts in this.

ACTION 9.1 To ship the IBM boxes back to CERN.

Roberto said the Quantum deliverable was sent to the peer reviewers, after which no change was needed. The document was available only on the Quantum web site
so far, but he announced it was going to be made available on the open TF-TANT web site soon.

2. TEN 155 update

Roberto reported that Hungary was upgraded to 68 Mbps and that the AUCS interconnection was upgraded to 155 Mbps, but the upgrade of the Greece circuit to 155 Mbps was still pending. A new 155 Mbps ring was being order linking Spain, France and Switzerland.

Phase 1 of the 5 cities ring at 622 Mbps was due on July 1st. Roberto said the capacity was in place, but part of the ring failed showing that there

were still some problems with testing IP over ATM. The migration did not start yet. The following phase aiming at doubling the capacity to Switzerland, Sweden and Italy was unlikely to be delivered on time.

3. Activities after October 2000

Howard, explained the formal difference between the GEANT infrastructure and GN1, the 4 years project that DANTE and the NRN were negotiating with the European Commission. The project timetable was:
31 may - proposal to EC
7 June - Call for evaluation and parallel call for expression of interest from DANTE for the service procurement
20 June - IST committee decision
30 June - Negotiation of contract with EC and deadline for Expression of Interest (93 responses were received by DANTE)
7 July - DANTE issued the Invitation to tender
12 July - EC agreement with minor changes on the Technical Annex
? - EC contract

29 September - Deadline for ITT
1 November - EC funding start

Howard said the new network should technically start on 1st November, but capacity cannot be expected to be in place before 2001. The current network will not be stopped, the GEANT network will be based on continue upgrades to the existing one.

Howard said GEANT is meant to provide operational services, but due to EC rules its technical programme contains some research issues. In the negotiation of the GN1 contract the test programme was defined not with the aim of doing research "per se", but with the overall intention to take pieces of new technologies and integrating them into the operational service. Three phases were defined: technology equipment testing, specification of the service that GEANT wants to provide and finally piloting services. The technical annex defines topics for year one and contains a Technology roadmap for services to be introduced in the next years. The roadmap will be redefined every year. The initial topics are:
- MPLS / guaranteed capacity / VPN
- Premium IP service
- Improved multicast
- IPv6
- Traffic measurement

DANTE would like to have working groups on each of these areas, with open membership, cooperation with equipment suppliers and combined meetings of the working groups 3-4 times a year. The timeline for the areas to be covered in year one is flexible, but MPLS/VPN is expected to be operational as a pilot service in November 2001. For IPv6, Traffic measurement and multicast there is no schedule. Premium IP should be delivered "as soon as possible".

Howard remarked that optical switching is not in the list but the tender for the service procurement issued by DANTE leaves the bidders free to offer it: he said the topic would be included in the programme depending on the result of the tender.

Olav remarked that there are issues for NRNs that are beyond the small list of technical areas specified by DANTE for GEANT, but Howard said the things that need to be done as part of GEANT have to be specific and not "blue sky R&D". He suggested that people get a draft copy of the GN1 proposal. Phil asked how the network would look like in the future. Specifically, would migration from ATM to POS be the case? Howard acknowledged and said that the intention is to have 2.5Gbps in a number of countries initially, and that the technology choice was up to the tenderers. ATM is not considered feasible at those high speeds, so it is likely to be at least POS initially.

Olav asked how the tasks in the experiments should be divided between operational people and the TF-TANT group. Others remarked that there is a need to define exactly what kind of people is needed. Howard said that people in the TF-TANT group have the right mix and need to be involved. Others said the TF-TANT community is not only made of the people who attended the meeting, and the technical programme should be able to motivate other people beyond this group to participate. Mike Norris, supported by other people, said that having a roadmap and clear deliverables is the right approach, but the community also wants to see some "blue sky" research. Michael said that the group always had long term research on the side to more operational work and Olav added that the topics identified are too much short term: a common broader view of what to work on in the next years was needed. Howard accepted the comments and said he regrets that there was not enough time to discuss all issues with the EC. Victor remarked the good scheme of TF-TANT, which was based on deliverables and reports but also on long term more open research. Another question was how the group was going to cooperate in defining the topics that are not covered by the GEANT project. Howard answered that that would need to be discussed and to get support by the NRNs.

4. Interdomain issues for Multicast

Jan reported on a problem/bug with multicast that causes inconsistency in multicast group memberships: the joins keep coming from some users even after the sender stops to send data and there are no more Session Announcement messages. It is a known problem which everybody should be aware of in case they or their customers have problems with multicast. A solution by Cisco requires upgrade to IOS 12.0(8.6)S1 and higher, but this cannot be used in the backbone. People are reluctant to upgrade, both for

technical and financial reasons. However it was remarked that only the IOS on the sender (the closest to the source) needs to be upgraded (i.e. the one on the LAN responsible for sending registers). The detailed description and resolution of the problem are available at: http://www.dante.net/mbone/nop/sm_non.html

5. Flow Measurement (NOTES FROM KOSTAS ANAGNOSTAKIS)

No agenda was set for this meeting, everybody agreed it should proceed as an open discussion. Simon asked whether everybody got the QTP final report. Some people were not able to access it.

ACTION 9.2 Simon to put the report on the flo.ma Web page (DONE)

David reported on Purgatorio, which is based on MRTG for statistics per NRN, with 1 hour matrixes. However, only integers were supported by MRTG which causes problems with the statistics presentation. The URL to see the statistics is http://stats.dante.org.uk/. Some scalability issues appear, esp. lack of resources to store data. The package uses cflowd.

Tom Kosnar, CESNET, mentioned that Cisco GSR netflow works although some problems with timeouts etc. occurred. David said that the Cisco 700 with Cisco Express Forwarding operates at 70% load when there are 10k fps.

Olav Kvittem, UNINETT, observed that this is a potential threat for DoS attacks.

Simon mentioned that due to problems with charging at SWITCH and discrepancies that have appeared, AS-based lookups, port number and network number were kept in file.

Olav reported on an HTML interface that was built at UNINET that displays raw cflowd information. It involves data collection only and not any kind of post-processing. He asked whether it is possible for cflowd to reduce data by summarizing. Simon replied that the rrdtools do have that feature of information being made coarser as it becomes older.

ACTION 9.3 Olav to post URL (the password is the same as for QTP)

Simon reported that at SWITCH raw data is kept in archive for billing issues. The list has been pretty quiet. However, two interesting papers from AT&T Research on measurements for traffic engineering. Specifically, a tool called "netdb" is described there, which works on aspects of how to perceive and characterize traffic. Another paper describes a T-Matrix method that is useful for network design and planning.

Herve Prigent mentioned another similar tool called "netnet".

ACTION 9.4 Herve Prigent to provide information/URL on this tool

Olav Kvittem asked whether there is a way of using flow measurement for intrusion detection. David replied with his experience on reporting potential DoS attacks through an e-mail notification system. However, too many false alarms were produced. "Abnormal" traffic between two AS'es is used to trigger these alarms. However, sometimes this is natural.

Simon reported on a catalog of DoS signatures, like ICMP, TCP SYN flood and port scanning. Also, a new service called Gnutella with thousands of servers was identified in recent accounting information. How do changes affect specification of flows especially the perception of things. Netdb seems to provide answers to this kind of questions

Tom Kosnar asked whether it should be suggested to Cisco to have AS number based accounting. Simon said that in Netflow you can select 3 different types of AS related information: 1) number, origin and neighbor. The alternative to this is policy based
accounting.

Tom Kosnar asked whether all this scales to Gbit/s.

Simon asked whether somebody sued netflow on the catalyst. 6509 v6 was not supported, only sampling.

David said that the Juniper flow monitoring, based on the Internet Processor II, that uses 1 out of X sampling stores netflow data locally, except for AS information.

Simon observed that there is an overall tendency to adopt netflow (cisco, juniper and cabletron). Building export capability into flexoscope could be useful. Also, a front end on a web page would be interesting.

FUTURE ACTIVITIES
Simon requested participants to share their ideas on how the group should proceed and what future activities they would be interested in.

Leon Gommans asked whether the flow measurement work should incorporate the overlapping diffserv activities.

Tiziana added that different metrics are interesting when monitoring quality of service. People are interested in doing something advanced, not necessarily with the existing tools.

Simon reported on the availability of an SNMP MIB for SLA accounting

Phil reported that he tried to use the Cisco SNMP implementation but the system crashed. This was Cisco IOS 12.0.17

Simon asked what kind of stats would be interesting in this case (for RED etc)

Tiziana replied that monitoring on production is required; QoS is blamed for degradation so some way is required to identify the source of performance loss or prove the contrary.

Tom mentioned that there is a new URL for the MIB design and tool creation resources.

ACTION 9.5 Tom Kosnar to send new URL on MIB design and tool creation

Phil asked where the limits should be set for this activity Leon replied that making sure it's meaningful

Olav proposed that per-hop monitoring in forms of a histogram , would be required, rather than an end-to-end performance analysis.

Simon said that flow bandwidth would be necessary to detect aggresive flows, compute maximum delay. Vern Paxson's paper on the impact of packet loss on TCP performance was mentioned. It is obvious that different flows observe different rates on the same link. An empirical analysis is needed, which is a good subject for a paper, i.e. monitor TCP window size per IP address. It was asked whether a different mailing list should be set-up, but everybody agreed that the current form is fine. QoS Validation would be the new activity of the flo.ma. group.

Phil asked whether accounting-performance aggregation is what we are talking about.

Simon said that this is a different problem.

Phil put the question whether sampling is ok in this case.

Simon suggested that any thoughts should be posted to the floma list.

Tiziana asked about when these new activities should begin

Simon proposed that people should just see how it works and that there is no need to wait until October etc.

Olav requested that the tasks should be distributed among the participants. However, things are not concrete yet and more brainstorming is required.

6. Multicast

It was discussed why many end users do not use multicast and what can be done to improve usability and deployment of multicast. Lada said that ideally to improve end-to-end multicast NICs should filter more multicast addresses at level 2, LAN switches should use IGMP snooping, and leaf routers should communicate via PIM SM and/or BGMP. However there are problems because most Ethernet cards provide limited hardware support for multicast MAC addresses and often islands still exist that use dense mode protocols. Moreover there is insufficient demand for native multicast applications (unicast workarounds are often more reliable than native multicast) and finally most of the existing software is difficult to use and poorly documented.

Lada outlined the following possible activities to improve multicast in the context of GEANT:
1 use existing monitoring tools (MRM, mtrace) regularly,
2 prepare a test suite for end users able to inspect the user environment, perform tests and generate reports,
3 support existing products,

4 perform and publish tests of equipment,
5 cooperate with equipment vendors,

6 apply results of QOS, policy and security research groups,
7 set up web site, FAQ, platform for discussion,
8 workarounds: reflectors, tunnels, etc.

ACTION 9.6 Lada to distribute his action plan for the next month

Lada and Victor discussed the need of killer applications being able to support multicast deployment. Such applications are there but they do not enforce multicast mainly because unicast is still an option whenever multicast is not available, for instance at home. Victor said Mbone tools are still driving the demand for multicast so HEAnet need to provide it, however he wonders if this is enough for stimulating vendors to provide good multicast equipment. There are still many problems to be solved. It was proposed to run an open mailing list, but the decision was to watch out the existing Mbone list. In addition the actions below were agreed.

ACTION 9.7 Ladislav to co-ordinate collection and publishing on dante's web of
known problems/solutions.

Wilfried asked a number of questions. Who might benefit from deploying multicast? What sort of community would be able to source multicast? Is any-cast necessary at all? Traffic engineering replication in a tree-like scheme for content delivery is not enough? A backbone which is multicast capable would really be cost effective? It was remarked that in the universities there is a real interest in applications requesting many-to-many multicast features, however Robert pointed out that there are ways to simplify the current multicast model when there is only one sender, like it was done by Real Player.

Robert talked about latest developments of MRM and asked about BGMP implementations on the vendor routers. Then, he described the components of MALLOC Dynamic address allocation:
- MADCAP multicast address dynamic client allocation protocol (RFC 2730),
- AAP multicast address allocation protocol ID-04-June-2000,
- MASC multicast address set claim protocol ID-06-July-2000.
He also said that implementation of experimental maskd is available but not tested yet. The MRM test configuration included DANTE, Rus, CESnet and RedIRIS, however RedIRIS was not reachable because they had IOS 11xx on Designated router. Rus was also not reachable because MSDP messages did not reach DFN.

Lada was interested in the possibilities to use multicast over tag switching and wondered about any plans to support multicast over MPLS. Larry said he would ask at Cisco for information about any such plan.

7. IPv6 (NOTES FROM KOSTAS ANAGNOSTAKIS)

AGENDA

- interoperability,
- applications,
- multihoming issues,
- DNS,
- Geant,
- addressing

INTEROPERABILITY
Christian said that there is not too much progress, actually nothing to report. A Cisco router is now available so that tests can be done with more heterogeneous configurations and not just telebit boxes. Collecting configuration information for IPv6 over ATM as example configurations for interoperability of cisco and telebit and 3com / zebra could be appealing.

Wilfried said that besides MRTD to Cisco IOS tests it could be interesting of getting hold of a juniper box when they actually implement it. It was asked whether anyone tested IOS's new DNS implementation esp. wrt reverse resolution.

It was discussed whether there would be problems with the NDA. As long as it was stated that any results were obtained on a pre-production system, there should not be any problem.

Alex asked Wilfried whether the 3com-mrtd interoperability tests can be done.

Wilfried agreed.

Alex announced that a new release of the Telebit software is out.

Simon noted that some problem might arise by MTU mismatch between ethernet and ATM which might light to exsessive drops. This problem was confirmed by CERN.

Christian mentioned OSPF for IPv6 is implemented in gated, zebra and mrtd, but not cisco.

CERN also reported that the new cisco release returns erroneous DNS names.

Simon suggested that the configuration should be verified.

David observed that everyone can be BGP with telebit boxes and asked whether anyone has done OSPF.

Wilfried replied that they are waiting for the new Cisco implementation. This would include IS-IS.

APPLICATIONS
Tim said that there is need to implement IPv6 applications and port existing applications to IPv6. However, no IPv6 multicast is supported yet. UCL is expanding it's multicast tools: a multicast mp3 jukebox has been developed. Also, an IPv6 implementation of sendmail, of web servers, and IPv6 plugins work. There is a problem with clients though, only mozilla can

be made to work. Uninett however, is running an IPv4/IPv6 proxy.

Alex reported that Quake/IPv6 works fine.

Simon observed that for deployment , implementation on a widely used OS like Microsoft Windows NT 5.0 would be required.

Christian also reported that DFN's IRC network is IPv6-based.

MULTIHOMING
Tim reported that not too much progress has been made by the IETF. New drafts have been published. On for cross-over v4/v6 tunnels. The other 3 drafts are about address selection. However, because of the address allocation issues that arise, any progress has been stalled, so there are no direct results.

Wilfried noted that 3 addresses are currently held and used.

Tim requested that the group should define what tests would be required, in which environment (FreeBSD was used as an example). In general low priority should be put on that as it will come up later.

DNS
Wilfried reported that there is not too much to talk about in this subject. Bind9 beta 4 runs smooth with no problems. Reverse DNS delegation has not been tried though yet.

David wants to push for delegations to NRR's to simulate root.

Wilfried said that all that is necessary is to send a request to Mill Mahne using pgp. At the last IETF no final decision was made to move from .int o .arpa. New stuff in .arpa is still a moving target. Also, a problem that arises with Cisco IOS's DNS is that it is required for the router to be able to use DNS for the local IPv6 applications i.e. traceroute etc. A new person was hired by ACOnet to work on IPv6 and bind.

Tim reported that Viagene would run as rot but there are obvious problems for example running the .com for IPv6.

Simon reported that IPv6 has been added to the .ch registry's Web interface.

David would like everybody who has a prefix to be delegated. Commitment from all qtp6 members would be required.

Tim reported that there are three kind of prefixes given by the registry:
- /29 long term
- /35 slow rollout
- /48 end site

Wilfried reported that /64 , /48 and /56 prefixes were discussed, and the IAB interim suggested that address allocation guidelines should be more generous although they should try to avoid wasting address space. He also

reported that they are trying to put ADSL in student residencies which would result in an immediate requirement for 17000 addresses, a clear need for IPv6 deployment.

Tim added that 5 million users are expected to join JANET

GEANT AND FUTURE ACTIVITIES
Roberto reported that in the Geant activities, IPv6 is clearly put forward as a working item. Testing of IPv6 will continue as well as following the developments. From an organisational perspective, it would be preferred for a Geant partner to lead any IPv6 activity, so UKERNA would take over from Telebit. Alex will continue to be the activity group until October 2000. The network would also continue to be in place until further arrangement.

Wilfried expected no major bottleneck unless traffic engineering technology is extensively used and changes the general picture.

Simon mentioned that IPv6 can be used over an MPLS ore just like it was over ATM.

Tim reported on UKERNA's collaboration with I2/ISI/NASA, potential 6tap interconnection, UCL QoS testing link, and the availability of NTT's commercial IPv6 secvice.

8. MPLS

Alain talked about the VPN connection model and said that different sites in different VPNs can use the same address space. There are three main players: VRF (VPN Route Forwarding), RD, a Route Distinguisher used by MBGP in making the IPv4 route globally unique, and RT (Route Target), an extended community attribute. He said VPNs are defined on the basis of routing table sharing. There are possible different strategies to import and export routes between routing tables. In one case several sites are grouped in the same colour VPN, in the other, each site has a different colour on the local backbone (its own routing table, own VPN). With the second strategy every combination is possible but this makes configuration more difficult, especially in large networks.

Alain reported some problems in managing VPNs related to assignment of VRF to more VPNs. The routes are local to the PE and are not exported to other PE. Victor remarked that VRF is a local name and the same name should not be used for both VRF and VPN. A second issue is related to providing Internet access on a MPLS network though different ISPs. One should use different VPNs for each ISP and set default route on client side. Alain said VPN across different administrations will be supported in the next release. Other problems are related to filters: SNMP requests on PE routers do not work when you are part of a MPLS VPN, the SNMP monitoring devices are part of the backbone. Problems also occur with Netflow because it does not support yet tag switch enabled interfaces.

Alain closed his presentation talking about MPLS Traffic Engineering. He

said Uni-directional MPLS tunnels need to be set up for a specific path. This can be done either in three different ways: by specifying origin and destination and use IGP computed path, by specifying each single node and finally by specifying the conditions to enforce the selection of a path modifying the IGP computed path.

Herve listed the open issues to be looked after in the future. The first issue is about traffic engineering and guaranteed bandwidth. Hard point-to-point QOS constrains in MPLS should be implemented, moreover there is a need to emulate ATM-like services. Finally the current beta version from Cisco does not work with 72xx routers: 75xx only are supported now, with POS and Ethernet interface. Herve said training on that had already

started with Cisco. The second open issue regards Inter-domain MPLS testing, particularly in propagating services across carrier boundaries. A test is being carried out using Juniper routers by France Telecom (vTHD) and by Renater (PlaGE) on 2.5 Gbps POS links with Gigabit Ethernet client interface. The last open issue regards interoperability between Cisco and Juniper routers.

Herve closed his talk saying Crihan is operating the network at the moment, but France Telecom will take over in a few weeks. He added that the MAN in Rouen has dark fibers. Current network deployment includes the ongoing expansion in Normandy: starting in 2001 the IP service currently based on ATM will be based on MPLS.

Lada opened a discussion about the need for MPLS VPNs saying that CESnet does not have any request for VPN. He was wondering if there is a demand for it in other NRNs. Howard replied that projects need MPLS VPNs for specific purposes, either as defined VPNs or with some guaranteed bandwidth. Victor said SURFnet and HEAnet have such requirements and are providing some basic services right now. Herve remarked that to provide VPN service across countries one need to know who are the end users.

Tiziana asked for a roadmap about Traffic Engineering and Guaranteed Bandwidth. Herve said discussion with Cisco started and he would organise a meeting with them soon in order to set up a timetable. They would also need to ensure that they have the appropriate hardware to be able to carry out the experiments, but some test with basic TE functionality would be carried out anyway. Herve asked people who are interested in TE to provide input for those tests. Overall the tests should study the interaction between MPLS, TE, and Diffserv. Phil observed that this kind of testing is very useful, especially because the ATM infrastructure will be disconnected very soon.

9. Policy Control

Leon said that Mark Janssen from the University of Utrecht had problems in working on the Policy Control stuff so far and therefore the activity is behind schedule. However they have now set up a new deadline for end of August to implement 1-1 virtual leased lines between INFN and Twente via Utrecht using Open Policy System (OPS) and the Cisco QoS Policy Manager

(QPM). Michael announced that now there is a COPS version of QPM. Leon said the current tasks are mostly related to write documentation and comparing OPS and QPM Full detail will be available on the web at http://www.phys.uu.nl/~mjanssen/.

Next activities include intra domain service definition, Bandwidth Brokers allowing different service domain to inter-operate and intelligence of edge devices, controlling user access to defined services.

10. Diffserv

Tiziana briefed about deliverables and presentations of work at international fora, including IWQoS in Pittsburgh, TERENA TNC 2000 and QoFIS in Berlin. All documents are on the Diffserv Web site http://www.cnaf.infn.it/~ferrari/tfng/ds/.

Tiziana said the Diffserv network topology was updated to include Norway and IRISA in France, but needed to be change again because of the disconnection of Utrecht and Twente. Tiziana proposed to connect GRnet to SWITCH and setting up additional connection from INFN to SWITCH. Roberto said Greece would be connected to London very soon and therefore he suggested to connect GRnet to RedIRIS instead of SWITCH, but Tiziana chose to connect it to Switzerland. After Twente and Utrecht will be disconnected they can still be part of the network by tunneling to INFN and SWITCH. Tiziana proposed to tunneling Utrecht with INFN

ACTION 9.8 Leon to define what has to be tested and then to decide whether he wants to connect Utrecht to INFN or to SWITCH

Test reports started with the tests with eight queues on the MAN inside Italy with the University of Bologna and INFN. Independently on the packet size, delay, jitter, etc. with PQ was always better than WFQ

Then Tiziana reported router performance tests based on CPU utilisation as a function of traffic rate. She said without CAR on 7500 routers 90% usage of CPU is reached with 24000 small (64 bytes) packets per second, but with CAR saturation is reached sooner. The maximum packet rate with 10 CAR is about 18000 packets per second. Delay and jitter were computed by focusing on particular access list (ACL) configuration. Delay and Jitter were also tested at low speed. Details of all experiments are provided at:

http://www.cnaf.infn.it/~ferrari/tfng/ds/. Roberto expressed an interest in knowing the impact of Delay and Jitter at higher speed. Tiziana said the test would be repeated. Larry suggested to try also with triple CAR or triple ACL, but Tiziana said that this test would not show any useful result. Tiziana asked people which other features would be useful to test. Simon Leinen said he had tested RED on a transatlantic line and found out that the CPU usage was increased by a factor of 2. It might be useful to repeat that test. There were also proposals to test RED and PQ at different rates. Finally Tiziana reminded that demos of EF testing in the wide area are available on the web.

Larry outlined a number of Cisco updates on QOS. He said the functionality supported by the IOS before version 12.1(2)T were, on the ingress side, CAR, MQC(Modular QoS CLI)for metering, policing, marking; dWRED (d standing for distributed) was supported in the in-box, whilst on the egress (d)CBWFQ and LLQ (low latency queuing = PQ/CBWFQ) were
supported. The new features introduced from IOS 12.1(2)T are generic traffic shaping in modular CLI (MQC), CBWFQ with GTS and DSCP marking using MQC. Starting from IOS 12.1(4)T a number of new features would be introduced, like WRED enhancement for DSCP, CP QOS MIB support (low latency queue on VIP 12.0(5)XE6, 12.1(4)T) and Hierarchical policy maps (on 12.1(2)E and 12.1(4)T).

Then Larry made some remarks about general features in the main IOS "TRAINS": IOS 12.0( )T supports new features for particular platforms, whilst 12.0 ( )XE provides them on 75xx routers. The same hold for 12.1( )T and 12.1( )E. IOS 12.0( )S is ISP oriented and supports multicast MSDP/MBGP/PIM-SM. 12.0( )ST has limited availability and supports GSR MPLS/VPN development. Latest updates regard development of tools at Cisco. The role of RSVP at edge and the role of MPLS, MPLS/TE, MPLS/VPN are being explored, whist tools related to bandwidth broker, admission control and provisioning are under development now. There are still open issues on validation of the cloud model as opposite to pipe model services by integrating different components that have been tested separately so far. Another remark was that many people use features like WRED or CAR, but Larry does not know of many people using some sort of "Diffserv packaged service", probably because they still do not see a ROI. Michael said he knew about one deployment in Finland, that however did not provide guarantee, but only priority.

Phil Chimento presented the results of Diffserv Aggregation Measurement, which was based on the architecture used by Tiziana and described in previous meetings, aggregating EF and BE streams at different nodes. The network topology had three bottlenecks: at INFN, CERN and GRNET. The measurement of loss per router (aggregation point) both with PQ and WFQ increased with aggregation: load was less at earlier aggregation points, and loss was found to increases with EF load. The measurement of loss ratio per source proved that the loss is affected by adding BE traffic as well as by the number of hops. Overall, the experiments seemed to show a problem:
that the loss depend on the length of the path-number of hops traversed.

ACTION 9.9 Phil to mail his presentation to Tiziana.

ACTION 9.10 Tiziana to put presentation by Phil on the Diffserv web site.

Mirjana presented the results collected by the Smartbit boxes. She said PQ was better in loss with large queue limit and in delay range, whereas WFQ was better in IPDV, Cdf and density loss, and sensitivity to the number of flows. Other results showed that IPDV gives an insight into the delay series and is sensitive to the number of flows, whilst loss and delay are sensitive to the queue limit. Also, periodicity was found in the IPDV and

less empirical density in the cascade of routers.

Tiziana remarked the need of developing the analysis further and write a report. Phil said they will draw conclusions from the experiments and will write recommendations based on the analysis.

ACTION 9.11 University of Twente to write recommendations based on experiments reported by Mirjana.

Antonio briefed about new tests with RED and WRED he had carried out since the previous TF-TANT meeting. He said this time he tested UDP instead of TCP flows. The first three tests were done with RED only. The first with 2 UDP flows and same packet rate, the second with 2 UDP flows and same bit rate and the third with 3 flows and same bit rate. A fourth experiment was carried out by testing WRED with two flows and same packet rate.

Tiziana closed the session about Diffserv experiments updates by listing the planned activities. The groups should agree on a demo to be carried out on wide area networks, and further shaping test to be carried out on 75xx routers. Finally to carry out further test with regard to AF. In this respect Tiziana said discussion was started with Octavio and people from Northwest University.

11. DATE OF NEXT MEETING

The next meeting will be held on the 5th and 6th of October 2000. The venue will be Vienna, Austria.

12. ANY OTHER BUSINESS

Phil Chimento presented activities related to the development of Simple Inter-domain Bandwidth Broker. It is being developed in a workgroup including the University of Twente, Siemens, and the Argonne lab. The goal of the workgroup is to define Bandwidth Broker model to be used in Qbone, produce recommendation for the deployment phase and specify inter-domain protocol.

The Inter-domain Communication Protocol is based on Simple Inter-domain Bandwidth Broker Signaling (SIBBS). The Intra-domain Communication Protocol makes use of tunnels to create virtual adjacency between brokers, making communication point to point rather than per-hop.

Phil said the are some open research questions regarding service definition, mapping services to Diffserv PHBs, resource allocation, tunnel triggering and aggregation mechanisms. All participants in the workgroup are bringing good ideas, but need to work on more detail on the protocols. Phil expects early in fall 2000 good design so that people start to implement the model.

13. ACTIONS FROM LAST MEETINGS

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

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.
- Ongoing

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

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

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.
- Done

8.5 Dimitrios to work with Tiziana, Simon and Roberto on refining the proposal in the first week of May 2000.
- Dropped due to limited resources and plan change

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

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

OPEN ACTIONS

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.
- Ongoing

9.1 Those who have the IBM boxes to ship them back to CERN.

9.2 Simon to put the report on the flo.ma Web page (DONE)

9.3 Olav to post URL (the password is the same as for QTP)

9.4 Herve Prigent to send information/URL about "netnet" tool

9.5 Tom Kosnar to send new URL on MIB design and tool creation

9.6 Lada to distribute his action plan for the next month

9.7 Ladislav to co-ordinate collection and publishing on dante's web of known problems/solutions.

9.8 Leon to define what has to be tested and then to decide whether he wants to connect Utrecht to IN

9.9 Phil to mail his presentation to Tiziana.

9.10 Tiziana to put presentation by Phil on the Diffserv web site.

9.11 University of Twente to write recommendations based on experiments reported by Mirjana.

</pre>