Project Number: RE 1009
Project Title: TEN-34

Deliverable D2.3

Final Programme Conference Paper

Deliverable Type:	LI - Limited
Contractual Date:	30 November 1998
Actual Date:	29 December 1998
Work Package:	2 - External Relations
Nature of Deliverable:	RE - Report

Author: Roberto Sabatino, DANTE
Cathrin Stöver, DANTE
Maria Pallares, DANTE

Abstract:

This deliverable describes the presentations and papers given regarding TEN-34 during its lifetime. It replaces a final programme conference paper which was not produced as the event for which it was planned was rescheduled as the Fifth Framework Programme launch.

Keywords:

TEN-34, Presentations and Papers, demonstrations

Table of Contents

EXECUTIVE SUMMARY

The TEN-34 work programme included the presentation of TEN-34 at a final programme conference. The Project Officer had proposed that a presentation on TEN-34 should be given at a conference planned for February 1999. As this conference is now planned as the launch of the Fifth Framework Programme, TEN-34 will not be presented. Nonetheless, TEN-34 has been presented at many conferences and exhibitions throughout the lifetime of the project. As agreed with the Project Officer, this deliverable is based on earlier presentations of TEN-34. It also provides an overview of the promotional activities undertaken in the project.

1. INTRODUCTION

TEN-34 has been presented at many conferences and events during its lifetime. Various papers have been published giving an overview of the project, highlighting technical activities, and focussing on policy issues. Although part of the project work programme, a final programme conference paper has not been produced since the event foreseen for the presentation will now be used as the Fifth Framework Programme launch event. This deliverable therefore provides an example of an overview paper and a technical paper, and reviews the promotional and dissemination activities undertaken by the project.

2. SAMPLE PAPERS

Appendix 1 provides a sample of a paper giving an overview of TEN-34 and the future of European research networking. This article was written by Cathrin Stöver of DANTE for ***axis***, The UCISA Journal of Academic Computing and Information Systems, and appeared in Volume 4, Number 4 1997.

Appendix 2 provides a sample of a technical paper documenting the traffic measurement and monitoring facilities implemented in TEN-34. The paper was presented at the Interworking 98 Conference in Ottawa in July 98, and again at the CCIRN Measurement Working Group meeting which was held during INET '98 in Geneva (21-24 July 1998).

3. TEN-34 PROMOTIONAL ACTIVITIES

Since its inception the TEN-34 project and the TEN-34 network have been presented at many conferences and events and in a number of press articles; papers on TEN-34 have been included in conference proceedings as well as in *DANTE in Print*. Various brochures and leaflets have been produced for distribution at exhibitions where TEN-34 was represented at a stand, and in many cases demonstrations were also given. A list of the TEN-34 promotional and dissemination activities is given below.

Brochures/leaflets:

- TEN-34 - The Information Superhighway for European Researchers, 4 pages A-4, A-2 size map on the back, September 1996.
- The TEN-34 Network, two page A-4 leaflet describing the implementation of the network, November 1996.
- TEN-34 TF-TEN: Testing of Advanced ATM Technology, one page A-4 leaflet listing the planned ATM tests as part of the TEN-34 Project, November 1996.
- TF-TEN: Advanced ATM Testing on a European Scale, a two page leaflet summarising the results of the experiments conducted, May 1997.
- TEN-34 - Creating the information highway for European researchers, 1 page A4, November 1997, handout for EITC Conference, Brussels.
- TEN-34 – Looking Ahead…, 1 page A4, November 1997, handout for EITC Conference, Brussels.
- TEN-34 – Creating the information highway for European researchers, 1 page A4, February 1998, handout for TAP Conference, Barcelona
- TEN-34 – Transition Ahead…, 1 page A4, February 1998, handout for TAP Conference, Barcelona
- TEN-34 – Remote data lenses for visualisation, 2 pages A4 with images, February 1998, handout for TAP Conference, Barcelona

- TEN-34 – TF-TEN: ATM Experiments, 2 pages A4, February 1998, handout for TAP Conference, Barcelona

Press releases:

(some URLs may no longer be valid)

- TEN-34 Consortium: towards a real Superhighway for European Research, 16 May 1995 http://www.dante.net/ten-34/ten34-pr.html
- EC Contract for TEN-34 Breakthrough in European Networking, 17 May 1996 http://www.dante.net/ten-34/contract/english.htm
- Launch TEN-34 Network: A Milestone in European Research Networking, 20 May 1997 http://www.dante.net/ten-34/launch/launch-pr.html
- Leading Six launch Pan-European ATM Academic Network, BT (on behalf of FUDI), 21 May 1997, http://www.dante.net/ten-34/launch/BT-pr.html
- Unisource launches the first Pan-European Internet services platform, Unisource, 20 May 1997 http://www.dante.net/ten-34/launch/Unisource-pr

Articles in the Press:

- 'Creating Europe's R&D Virtual Laboratory' in the Information Society Newsletter (News from the EC's Information Society Project Office) http://www.ispo.cec.be/ispo/newsletter/ISPOJUNE97/ISPOJUNE05.html
- 'Internet2: Traffic moves into the fast lane' Financial Times, 9 July 1997 by Richard Poynder, p.10 http://www.ft.com
- 'USA should fund half transatlantic links' by Alison Abbott & Colin Macilwain, Nature, 17 July 1997 http://www.nature.com
- 'TEN-34 – Building the information highway for European researchers' by Howard Davies in How to benefit from the Information Society, Publication of EC DGXIII, October 1997
- 'TEN-34 and the commercial realities of the next generation Research Internet', Dai Davies, ERCIM Newsletter, October 1997
- 'TEN-34: Implementing the Network'; Michael Behringer, ERCIM News magazine, October 1997
- 'TEN-34, the new high speed pan-European research backbone', Michael Behringer and Josefien Bersee, submitted to ConneXions, the Interoperability Report
- 'TEN-34 and the Future of European Research Networking' by Cathrin Stöver in ***axis*** – The UCISA Journal of Academic Computing and Information Systems, Volume 4, Number 4 1997
- 'Schnelleres Netz für Europa' by Cathrin Stöver in DFN Mitteilungen 46, March 1998
- 'TEN-34' Entry in PROSOMA Book (ESPRIT Programme), June 1998
- 'TEN-34 – Trans-European Network at 34 Mbps' by Cathrin Stöver in Towards the Knowledge Society, Publication of EC DGXIII, August 1998

Papers and presentations on TEN-34

- ***TEN-34: Ein Breitbandnetz für Europäische Forschung,*** Dai Davies, CeBIT, Hannover, March 1996
- ***TEN-34: A 34 Mbit/s Infrastructure for European Research**** Dai Davies, ATM Europe 96, a conference organised by ATM Perspectives, 19-21 March 1996 in Paris, France
- ***TEN-34, Progress on High Speed Networking in Europe,*** Howard Davies, 7th Joint European Networking Conference, 13-16 May 1996 in Budapest, Hungary
- ***TEN-34 and JAMES: Technical Plans*,*** Michael Behringer, Prepared (not selected for presentation) for 7th Joint European Networking Conference, 13-16 May 1996 in Budapest, Hungary
- ***TEN-34 - Progress on High Speed Networking in Europe,*** Dai Davies, June 1996, ESnet Conference, Elba
- ***ATM and the Internet - A Service Provider Perspective*,*** Michael Behringer, ATM - Developing the Broadband Future 1996 Conference, 2-4 July 1996, Paris, France.
- ***Trans-European Network Interconnect at 34 Mbps,*** David Hartley, Meeting of the European Networking Policy Group, Stockholm, September 1996
- ***TEN-34 and DANTE,*** Maria Pallares, INFN Network Workshop, 14 October 1996, Catania, Italy
- ***The Challenges of building the pan-European Optoroute for Research*,*** Dai Davies, Optoroute'96, Grenoble, 16 October 1996
- ***Building the pan-European information superhighway for research,*** Dai Davies, Asia-Pacific Advanced Network (APAN) Workshop, Tokyo, Japan, 8 November 1996
- ***Broadband Telematics Infrastructure,*** David Hartley, First Annual Concertation Meeting of the Telematics Application Programme, 2-3 December 1996
- ***The Implementation of TEN-34****, Michael Behringer, JENC 8, 12-15 May 1997, Edinburgh, Scotland
- ***TEN-34***, Howard Davies, JENC 8, 12-15 May 1997, Edinburgh, Scotland
- ***Experiments for Advanced Backbone Services****, Michael Behringer, JENC 8, 12-15 May 1997, Edinburgh, Scotland
- ***ATM Experiments for Advanced Backbone Services****, Michael Behringer, INET'97, June 1997, Kuala Lumpur, Malaysia
- ***TEN-34: The Trans-European IP Backbone***; M.H.Behringer; 1997, Proceedings 39th IETF, Munich, August 1997
- ***TEN-34 and the Future of European Research Networking***; M.H.Behringer, EuroInfo Conference, Warsaw, PL, November 1997
- ***TEN-34 – Trans European Broadband Networks***, Dai Davies, Trans-European Broadband Networks Conference, London, December 1997
- ***TEN-34 – The European Telecommunications Challenge***, Dai Davies, The Global Telecoms R&D Summit, London, December 1997
- ***TEN-34, TEN-155 and QUANTUM***, Howard Davies, Athens, January 1998
- ***Trans-European Research Networks***, Dai Davies, Navigation and Communication Services, TAP Conference, Barcelona, February 1998
- ***Policy and Planning of International Internet Connections****, Howard Davies, EuroMedNet, Nikosia, Cyprus, March 1998, DANTE in PRINT No.31, Cambridge, March 1998
- ***Inter NRN traffic statistics on TEN-34***, Roberto Sabatino - DANTE, in proceedings of Networkshop 26, University of Aberdeen (UK), March/April 1998-12-16
- ***The European Network Providers in the Era of Liberalisation***, Dai Davies, European Networking Policy Group, Vienna, April 1998

- ***The TEN-34 Mbone Pilot****, Jan Novak, Roberto Sabatino, DANTE in PRINT No. 36 Cambridge, October 1998
- ***Traffic Accounting Using Netflow and Cflowd****, Roberto Sabatino, at 4th International Symposium on Interworking, Ottawa, Canada, July 1998, and at the CCIRN Measurement Working Group held during INET'98, Geneva, July 1998.

* Titles marked with an asterisk are available in *DANTE in Print*

TEN-34 materials were distributed at the following events

- EITC Conference, Brussels, November 1997
- TAP Conference, Barcelona, February 1998
- INET'98 exhibition, Geneva, July 1998
- TERENA Networking Conference, Dresden, October 1998, with demo of the TEN-34 real-time network map.
- UK Information Day for UK Business in EC projects (Telematics, Esprit, ACTS), London, October 1998.

TEN-34 was used to support the participants' Terminal Room and the exhibitors at the INET '98 conference and the preceding workshops (14-23 July 1998).

Updates on TEN-34 appeared regularly on the TEN-34 web pages and in The Works of DANTE, a bi-monthly e-mail newsletter.

General References

• TEN-34 home page:	http://www.dante.net/ten-34
• TF-TEN home page:	http://www.dante.net/ten-34/tf-ten
• TEN-34 public deliverables:	http://www.dante.net/ten-34/DELIVERABLES/
• TEN-34 progress:	http://www.dante.net/ten-34/progress.html
• The Works of DANTE:	http://www.dante.net/pubs/works.html
• *DANTE IN PRINT:*	http://www.dante.net/pubs/dip.html
• DANTE Annual Reports:	http://www.dante.net/pubs/ar.html

APPENDIX 1

TEN-34 and the future of European Research Networking

The last decade has witnessed an increasing demand for on-line communication between members of the European research community. In the Western European nations, national research networks have been set up to satisfy the demand for rapid exchange of information nationally. However, high-bandwidth connectivity on a European level was missing for a long time.

In 1993 DANTE was established by a number of national research networks to co-ordinate and manage the creation of a high-speed pan-European connectivity for research purposes. The first step in this process was the building of EuropaNET (1992-1997), which connected 18 European networks and offered overseas connectivity to Japan and the USA.

Since Spring 1997, the National Research Networks have been migrating to the EuropaNET successor, TEN-34, the **T**rans **E**uropean **N**etwork at **34** Mbps, which was officially launched in May 1997. Today TEN-34 interconnects the National Research Networks of all the EU Member States and also of the Czech Republic, Hungary, Norway, Slovenia and Switzerland. TEN-34 also offers researchers direct connectivity to the United States. DANTE manages the TEN-34 network and is co-ordinating partner in the TEN-34 consortium.

The TEN-34 network topology (see network map: http://www.dante.net/ten-34/ten34net.gif) presents itself consisting of two subnetworks: one an IP (Internet Protocol) subnetwork run by Unisource connecting Germany, the Netherlands, the Nordic countries, Spain, Switzerland and the UK at 22 Mbps. The second subnetwork is based on ATM VP (Asynchronous Transfer Mode Virtual Path) technology and links Austria, the Czech Republic, France, Germany, Greece, Hungary, Italy, Luxembourg, Slovenia, Switzerland and the UK. The two networks are united at interconnection points in Geneva, Frankfurt and London. A Transatlantic line from Frankfurt connects TEN-34 to the United States.

In the past three months, the performance of TEN-34 in terms of availability has been very good. With one or two exceptions, access port availability and line availability has been greater than 99%. Line loads on the ATM VP side of the network range between 20-70% during the busiest periods except between Paris and London where line loads are approaching 90%. However, with these loads, packet loss is very low and TEN-34 adds very little performance overhead to end-to-end connections between users.

Therefore the benefits of a high-speed European research backbone are obvious. For the first time in European history, European researchers have the means for proper communication. They are no longer dependent on the public Internet and all its downsides starting from having to share lines with thousands of private users to lack of Quality of Service. For the first time scientists can transfer large files or databases from one site or University to the other, even during the day without any major packet loss or delay. For the first time they can follow experiments at partner sites from their own laboratory or they can run interactive remote applications. In short, they can communicate.

Economic realities

So, if we take a look at the achievements of TEN-34 today, we see that it has provided a very good pan-European coverage to support the European research community. Accesses to TEN-34 are running at up to 45 Mbps. However, there remains a significant bandwidth gap between what is available to individual institutions that are connected to the National Research Networks, typically at 155 Mbps and the maximum access speed to TEN-34 at 45 Mbps. Thus TEN-34 is undersized for a true pan-European network to support, for example, multi-media applications. This bandwidth gap between what is available nationally and what is achievable on a pan-European basis, represents a significant inhibitor to a homogeneous trans-European broadband network.

From a technology point of view, bandwidth should not be a significant problem. The current capacity of the fibre optic systems and predictable developments in the area of wavelength multiplexing, all imply that this building block for broadband networks should not act as an inhibitor. The current monopolistic environment in the European telecommunications market builds certainly commercial and perhaps to an extent, regulatory roadblocks that prevent customers from obtaining such bandwidths.

TEN-34 is an expensive project. The annual cost of the network is about 40 MECU. For this reason DANTE has made limited commitments until July 1998. We are currently working on a successor activity under the project name QUANTUM. We believe that within QUANTUM we will be able to obtain much more cost-based prices. However, there will be a considerable variation in the deregulation activities of each European nation which might lead to an even greater distortion of the cost-structure for national connectivity in Europe. Eventually the liberalisation of the European telecommunications market will lead to an improvement of the commercial position, but this improvement will be gradual and patchy in terms of geographic coverage.

A choice of two technologies

The Internet is becoming a persuasive part of the telecommunications world. There is, however, much confusion between Internet routing technology and the services that can be built using this technology. Internet routing technology has historically been supplied by two dominant US suppliers and their dominance of the market has meant that de-facto standardisation has been quick and effective in what is, in telecommunications terms, a relatively primitive technology.

Internet services represent a much more complicated picture. The origins of Internet services lie in the research community. Unlike most telecommunications development which has derived from the slow technical co-operation of monopolists, the Internet has demonstrated that it is possible to be technically innovative at a pace associated with the computer industry. The Internet has had, in its short life, two major advantages over traditional telecommunications services. It has never had to face economic reality in the sense that the development has been paid by continuous government funding. Secondly, it has never had to provide Quality of Service. Now with huge growth, continuing and augmenting demand and increasing complaints about quality, the issue of cost can no longer be ignored.

In contrast ATM (Asynchronous Transfer Mode) is a technology of the traditional telecommunications' world. ATM developments have been underway for at least fifteen years. Much effort has been devoted to standardisation in the traditional monopolist fora and via more informal routes such as the ATM forum but, even in 1997, TEN-34 has demonstrated that ATM switched connections with pre-defined Quality of Service remain an elusive dream. Again there is a significant difference between technology and service. ATM technology is provided by a much larger group of suppliers, who do not have the advantage of a quasi-monopoly power of the Internet technology suppliers. ATM suppliers service a world built on the tradition of defining and providing service in a monopoly environment. The ability of those developing ATM technology to understand the service requirements of the diverse group of service providers will be of prime importance for the future.

Internet technology may be a very cost effective way of providing teleshopping and low grade entertainment offering the networking equivalent of commercial television, but if it is to form the basis of a serious next generation networking technology supporting real time applications in the commercial marketplace, it will have to move away from the lowest common denominator service quality to offering real and predictable Quality of Service. The challenge from ATM which has lacked the technical innovativeness and speed of development of the Internet will become more apparent as Internet routing technology seeks to move from a primitive but cost effective data technology to mainstream telecommunications service technology. It is a battle between monopoly providers of service and monopoly providers of technology. The European research community will play an important role in determining who will win this battle. Interestingly the battle is truly global, since, whilst the centre of gravity of Internet technology is the USA, ATM technology is world-wide.

Implications for the next generation of European research networking

The growth of Internet technology and its increasing importance as a mainstream telecommunications technology will lead to a much more commercial approach to the provision of Internet services. The research community has to be very careful to ensure that a research network, particularly an international research network remains at the forefront of available technology. Interchange with the commercial worlds is indeed necessary but this must not be allowed to dilute the focus of research networking. Research networking is about exploiting technology in advance of that which the marketplace can offer. Research networks must play the vital role in acting as the bridge between development and the market place. This vital role precludes a simple customer-supplier relationship for service. Thus we can provide the best service for the European research community, while developing the network technology itself towards the fulfilment of European researchers' requirements.

For further information, please contact:

DANTE
Francis House
112 Hills Road
Cambridge
CB2 1PQ
UK
Tel: +44 1223 302992
Fax: +44 1223 303005
e-mail: dante@dante.org.uk
http://www.dante.net
http://www.dante.net/ten-34

APPENDIX 2

Traffic Accounting Using Netflow and Cflowd

ROBERTO SABATINO

ABSTRACT

The availability of traffic statistics serves several purposes amongst which accounting and charging. From a purely technical point of view their availability is an excellent prerequisite for network planning. Network topologies or routing plans can be readapted to match the actual traffic flows only if these are well known and understood.

In this paper the author discusses what has been done in the TEN-34 network (the pan-European research network) to obtain traffic flows statistics using netflow and cflowd. A set of web tools developed by the author of this paper, of simple design and straightforward understanding, provides access to the statistics produced.

1. INTRODUCTION

The TEN-34 consortium was set up by a set of European National Research Networks (NRNs) in order to provide a pan-European network infrastructure for research. DANTE is co-ordinating partner of the consortium and as such has the tasks, amongst others, of network planning and proposing routing strategies. In order to carry out these tasks efficiently it is necessary to have access to traffic statistics. This paper describes in detail the project carried out by DANTE to produce such statistics.

The interest is mainly in traffic between management domains, in the TEN-34 specific community between the NRNs and commercial peers of the TEN-34 network. For example, it is important to know how much traffic is exchanged between DFN (DE) and NORDUNET (SE). In addition to this it would be very useful to know what kind of traffic i.e. what applications (web, mail, telnet, news...) are being used between two NRNs. This information may be less useful for planning and routing purposes but it can be very useful for charging purposes. It is technically possible to obtain this information, but the efforts were focused on obtaining traffic volume statistics which are equally valid for charging and network planning.

The building blocks of the project are netflow and cflowd [2] that have been complemented by a set of tools and programs developed by the author of this paper.

1.1 NETFLOW

Netflow was initially developed by Cisco in its QoS program. It is a switching method that allows more efficient switching of packets according to the type of packet.

The concept of a flow is introduced. A flow is an IP communication between a pair identified by an IP address and a port number. When a packet between a pair is noticed for the first time on a router a flow structure is initialised and contains all the information regarding that flow, i.e. number of bytes exchanged, IP addresses, port numbers, AS numbers etc. of both the source and the destination of a packet. Each packet of that flow will contribute to the byte count and packet count of the flow structure.

Cisco added the functionality to export this data, and software has been developed to collect, store and analyse this information. Netflow data in production routers is exported continuously at a rate of up to 800 Kbps in the TEN-34 network. The data of a flow structure is exported once it expires, or after a time-out of 30 minutes.

1.2 CFLOWD

Cflowd was developed by ANS to collect and analyse the information available from netflow. It allows the user to store the information and enables several views on the data. It produces port matrices, AS matrices, network matrices and pure flow structures. The amount of data stored depends on the configuration of cflowd and varies from a few hundred Kbytes to hundreds of Mbytes in one day. For the purposes of the project the choice was to store AS matrix information which also turns out to be the smallest consumer of disk space. AS matrices simply contain for each pair of ASes the number of packets and bytes exchanged.

1.3 POST-PROCESSING

As previously mentioned, the AS matrices as produced by cflowd need to be post-processed in order to produce readable data. Programs and scripts that read an AS matrix produced by cflowd have been developed for this purpose. The output of these programs allows the user to produce NRN to NRN traffic matrices or graphs.

The matrices produced can contain the volume of traffic between pairs of NRNs on a daily average basis, hourly average or daytime average (10 working hours). They can also be produced on monthly daytime average or monthly daily average. The graphs show the changes in traffic patterns between NRNs on an hourly, daytime or daily basis.

The matrices and the graphs have been extremely useful in detecting traffic changes and have helped the TEN-34 partners plan upgrades and routing modifications.

2. DESIGN

The objective of the project was to provide inter-NRN traffic statistics, which no technology or software alone is capable of providing today and the challenge was to design the project in such a way that it would make the best possible use of currently available technology and public domain software and to complement these with a set of in-house developed tools or programs.

As previously mentioned, the implementation of the system is based on netflow v5 and cflowd. This choice was dictated from the fact that the TEN-34 network is entirely composed of Cisco routers and netflow v5 is a well established technology, proprietary to Cisco. Moreover, cflowd has been written by ANS specifically to collect, process and store netflow data in various formats, amongst which AS matrix format. This lead to the following top level design decisions:

- enable netflow on all TEN-34 border routers;
- collect netflow data on workstations directly connected to the routers with cflowd;
- perform centralised post processing of the data.

Figure. 1 PoP set-up (AT)

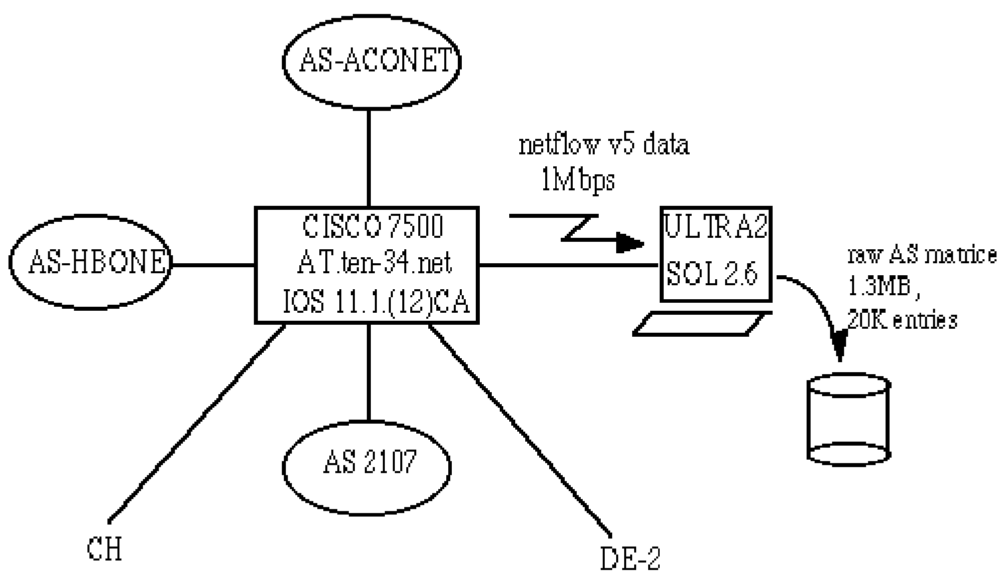

Figure 1. illustrates an example PoP set-up. Netflow is enabled on all of the TEN-34 router's interfaces. The netflow data is exported to the locally connected workstation. The amount of traffic generated by netflow is proportional to the amount of traffic that flows through the router

and varies from 400 Kbps to 1.3 Mbps. The cflowd process running on the workstation is configured to produce AS matrices. Depending on the set-up of netflow, which can generate origin AS or peer AS information, the size of the AS matrix will vary. Given the AS set-up of TEN-34 there was only the option to configure netflow to generate origin AS information.

Both origin and peer AS information have pros and cons. Peer AS information results in an AS matrix with very few entries but doesn't allow an in depth analysis of the traffic, whereas the opposite is true with origin AS information. It has to be said, though, that the system that has been developed does not currently exploit the more in depth analysis capabilities which are provided by having origin AS information, this is a topic for further development.

As mentioned, the choice of origin AS was dictated by the AS set-up of TEN-34. If peer AS information was chosen, from each TEN-34 router the only information extracted would be the AS of the next-hop TEN-34 router, which is obviously useless.

The size of an AS matrix generated by cflowd is approximately 1.2 Mbytes and contains approximately 20K entries.

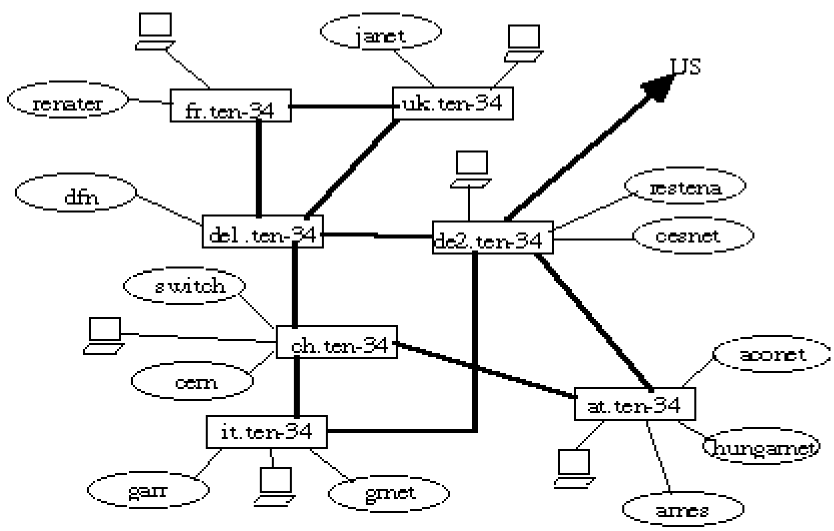

Figure 2. illustrates the set-up of TEN-34 with regard to the netflow data collection. Each router (rectangular boxes) sends data to the workstation directly connected to it and the AS matrix generated by the cflowd process running on the workstation will contain information regarding all traffic transiting through the corresponding router. One key design decision was to take into account only traffic that is to or from one of the NRNs directly connected to the border router being analysed. With this approach, for each pair of NRNs connected to two different routers there will be two sources of information, for instance traffic from JANET (UK) to ACONET (AT) will be accounted for on the uk.ten-34.net router and on the at.ten-34.net router. The same information will be available on all transit routers the traffic flows through (DE, CH), but this will not be taken into account. Moreover, ideally all sources of information of the same traffic should deliver the same results, but in practice we see differences in the order of 1%, which is viewed as perfectly acceptable. For the final calculation of the traffic between the pair of NRNs, the greater value of the two is used.

Accounting of traffic between pairs of NRNs takes place on a central processing workstation. The AS matrices are down loaded form the PoP workstations and an algorithm that converts the AS matrix into an NRN matrix is applied. The algorithm relies on the consistency of the RIPE database, where AS-macros and AS-lists are stored. The effect of the algorithm is to produce a file in an intermediate format suitable for long term storage and easy to be processed by the viewing tools (purgatorio, see section 3.)

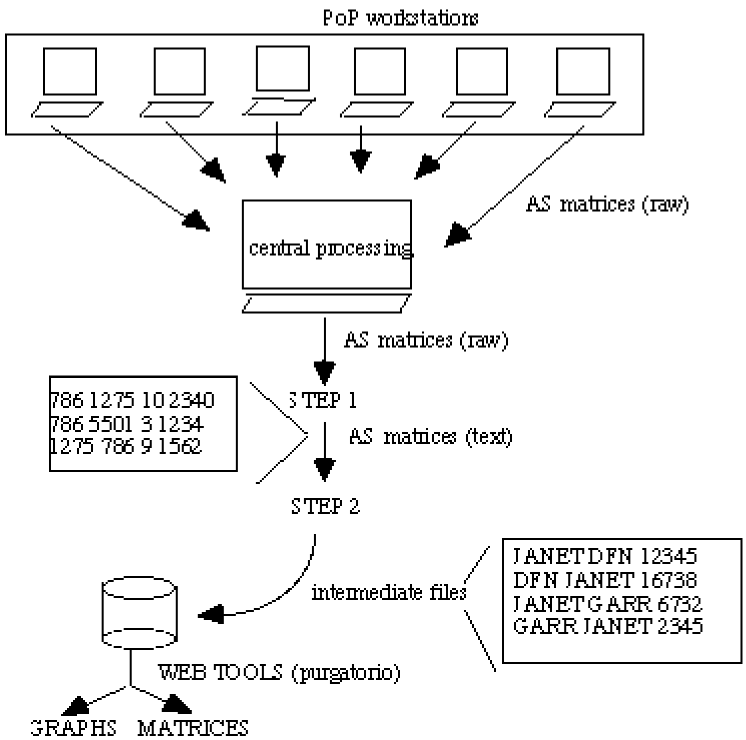

Figure 3. Data flow

Figure 3. shows how the data is processed and the path from AS matrices on the PoP workstations to intermediate traffic files read by the purgatorio system.

3. PURGATORIO: THE USER INTERFACE

Purgatorio is the collection of web pages and cgi scripts that interfaces the user to the data and allows different views of the data. The first question that arises is *why did we call it purgatorio*?. Recall that the name of our company is DANTE, and the first statistical package we developed was called *inferno* - hence there was only one choice for the name of the next generation statistics package.

Two different views of the data are possible:

- traffic matrices (full or per-NRN);
- traffic graphs.

When choosing the view via a traffic matrix the user may select:
- daytime (business hours) average of a given day
- 24 h average of a given day
- monthly (business or 24 hour) average
- hourly average

whereas when selecting the view via a graph the user may select:
- daytime (business hours) average over a selected range of days;
- 24 h average over a selected range of days;
- hourly average of a selected day.

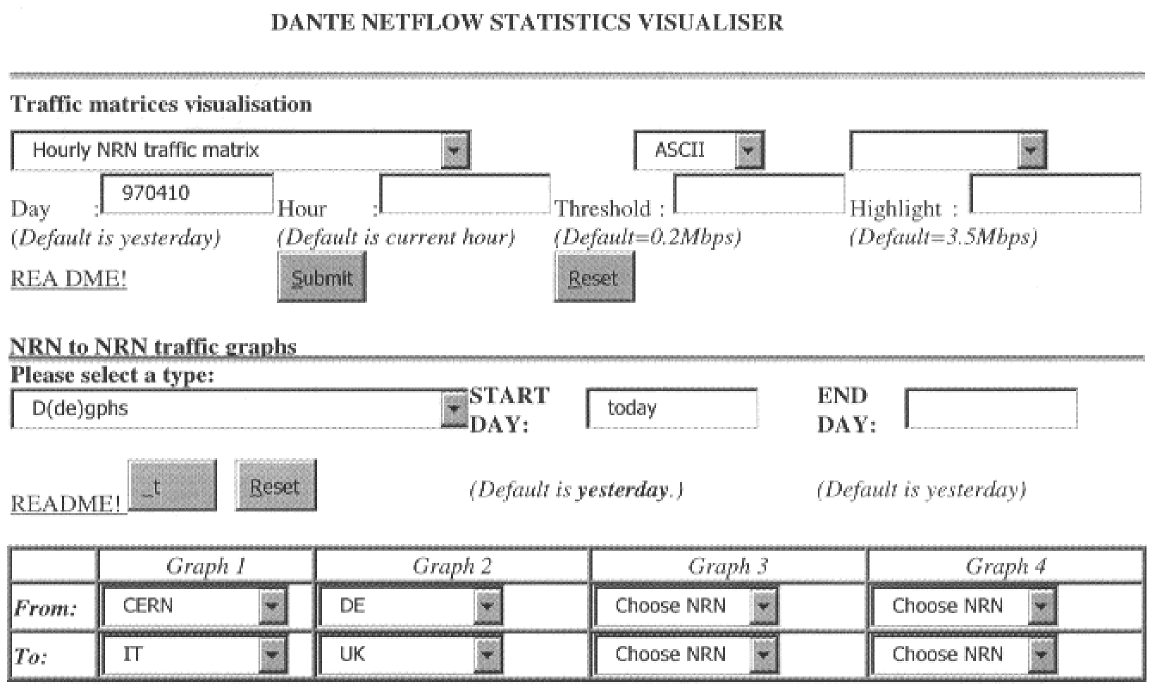

Figure 4. shows the web form, composed of two sub-forms. One sub-form is used to generate the traffic matrices, the other is used to generate the traffic graphs.

3.1 TRAFFIC MATRIX GENERATION

When generating a traffic matrix, the user chooses which kind of matrix to produce (monthly, daytime, hourly). The user may choose which day (and hour) and may specify two tuning parameters (threshold and highlight fields) that help the reading of a large matrix. The threshold field specifies to show only the cells that have a value greater than the value specified, while cells with a value greater than the value specified in the highlight field will be highlighted. For a matrix with more than 10 rows and columns these two parameters are very useful. The values in the cells show the average rate (in Mbps) of traffic between the correspondent NRNs in the row (source) and column (destination).

Together with the matrix, another table is generated that shows the effect of the two tuning parameters: the percentage of shown traffic and the percentage of highlighted traffic. For example with a particular threshold value many cells in the final matrix may be empty, but the percentage of traffic could still be close to 100%, which means that the main sources of traffic have been identified and therefore the matrix may be interpreted reliably and network planning decisions can be made with sufficient confidence of correctness. Reducing the threshold value could result in a matrix with fewer empty cells (which leads to a less readable matrix) but with little gain in terms of the total traffic shown. Figure 5. shows an example matrix.

Monthly (24h average)) traffic matrix of 04 / 98

NRN Matrix: Values are Mbps

src/dst	AT	CERN	CH	CZ	DE	IT	UK	??	*totals*	dst/src
AT	-	-	0.51	-	0.67	-	-	3.68	5.01	AT
CERN	-	-	-	-	0.62	1.04	0.39	0.78	2.85	CERN
CH	0.64	-	-	-	0.31	0.26	-	3.04	4.31	CH
CZ	-	-	-	-	0.52	-	-	5.57	6.35	CZ
DE	0.61	0.29	0.33	0.35	-	0.74	1.62	15.68	19.62	DE
IT	-	0.20	-	-	0.73	-	-	3.55	4.71	IT
UK	-	-	0.20	-	2.09	0.36	-	6.85	9.82	UK
??	2.06	0.66	1.76	7.87	12.61	8.66	5.08	-	38.70	??
totals	3.51	1.30	2.86	8.40	17.55	11.13	7.47	39.14		

Total traffic	Threshold value	Highlight value	Shown Traffic	Percentage of Shown Traffic	Highlighted Traffic	Percentage of Highlighted Traffic
91.36 Mbps	0.20 Mbps	3.50 Mbps	90.31 Mbps	98.85	69.53 Mbps	76.11

Software written by Dante NMC <nmc@dante.org.uk>

Figure 5.

3.2 TRAFFIC GRAPH GENERATION

Traffic graphs are mainly used to show the evolution of traffic between pairs of NRNs over given periods. The hourly graphs show the evolution of traffic over the hours of a given day, while the daily graphs (business or 24 hour) show the evolution of traffic between pairs of NRNs over a selected range of days. Currently there is a limitation of 4 pairs of NRNs that may be selected. This is for reasons to do with the layout of the web page and the readability of the resulting graph. Figure 6. illustrates an example output.

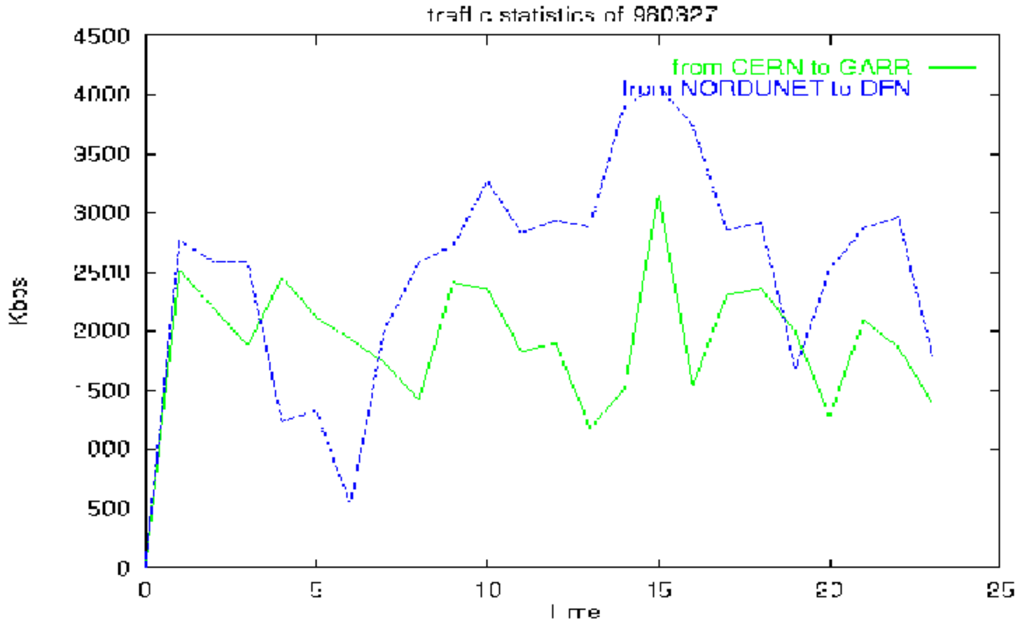

Figure 6

3.3 CONFIGURATION

The system has been developed and used on the TEN-34 network. It is possible to use the system on any network that resembles TEN-34, i.e. a set of ASes interconnected by a backbone network. The requirement is to deploy Cisco routers with an IOS version that supports netflow v5 as border routers and to have workstations directly connected to the routers.

The system has been developed with maintenance and extendibility requirements well in mind, therefore a set of configuration scripts that read a set of configuration files has been developed. These tools generate all the necessary configuration files for the PoP workstations, the central processing scripts, web pages and cgi scripts. Four configuration files are to be maintained:

- **nrn.config:** This file specifies for each network appearing in the traffic matrix or in the traffic graphs the correspondent AS-macro or number, the router from which the information is to be extracted from and the string that identifies the NRN or network in the traffic matrix. The format is:

NRN Country-code AS-macro/number
router_name

example:
JANET UK AS-JANETEURO UK
ARNES SI 2107 AT

The first line says that the NRN JANET will appear in the matrix with the country-code UK, the associated AS-macro is AS JANETEURO and the information for JANET will be extracted from the UK router.

- **router.config**: This file contains the information that the downloading scripts need to know. The format is the following:

host login router_name

example:
uk-ws.ten-34.net dante UK
de-ws.ten-34.net ten-34 DE1
de-ws.ten-34.net ten-34 DE2

This instructs the downloading scripts to download the UK router data from the PoP workstation uk-ws.ten-34.net, while both the DE1 and DE2 router data will be extracted from de-ws.ten-34.net.

- **site.config**: Contains the set of site specific definitions, i.e. home directories, location of tools and additional programs etc.

- **PoP workstation configuration files**. For each PoP workstation there is a configuration file which is read by a script that generates all the cflowd-specific files for the PoP workstations. The format of the file is the following:

WS login home-dir. dataport collectport TZ
router_name ip_address

Example traffic graph

example:
de-ws ten-34 /opt/ten-34 8765 8764 CET
DE1 193.203.227.25
DE2 193.203.254.5

This instructs the script to generate files for de-ws, where the login name is ten-34, the home-dir is /opt/ten-34 and the two port numbers used by cflowd are 8765 and 8764. The next lines tell which router data is to be accepted and stored by the cflowd process running on the workstation.

4. CONCLUSIONS

A system that has proved to be very useful has been developed and is currently being used on a daily basis by the European NRNs that are part of the TEN-34 consortium and by DANTE, the co-ordinating partner of the consortium. Similar data was available via IP accounting on routers a few years ago, but this mechanism is no longer feasible given the growth of the Internet in Europe. The system developed is the first in Europe on this scale we are aware of and has been met with interest and enthusiasm by the NRNs.

There is, of course, a lot of scope for improvement and extension of the system starting from application statistics (web, sendmail, telnet etc....) and top-10 users statistics, both of which are important for planning and charging. More work can be done to provide a larger choice of views of the data, aggregating the information on a configurable periodic basis, real-time analysis of traffic patterns and we welcome input from the users of this service in this respect.

ACKNOWLEDGEMENTS

I wish to thank my colleagues Steven Bakker, Michael Behringer, Howard Davies and Christoph Graf for the valuable advice and feedback they provided.

REFERENCES

[1] TEN-34: Trans-European Network with access ports of 34 Mbps
http://www.dante.net/ten-34

[2] cflowd: a Cisco flow-export collector
Daniel W. McRobb, John Hawkinson
http://figaro.ans.net/cflowd