Computer Networks and ISDN Systems 25 (1992) 546–553
North-Holland

Status reports on networks

The regionalization of EARN

Daniele Bovio

EARN Office, CIRCE BP 167, F-91403 Orsay Cedex, France

Abstract

Bovio, D., The regionalization of EARN, Computer Networks and ISDN Systems 25 (1992) 546–553.

During 1991 the EARN Network decided to reorganize its NJE routing following a plan issued by the EARN Routing Project Group. The EARN-RPG plan (also EARN2) is based on the separation of the network into regions, mainly following the model already adopted by BITNET for the definition of the BITNET II network. The criteria followed for the division are based on the density of nodes, traffic, and need of services in the area. The regions are connected through a set of EARN sites (called "core sites") located at strategic positions from the point of view of traffic patterns and connectivity, with the ability to have virtual NJE links among them and in some cases NJE/IP links with selected partners on the BITNET II core.

The main goal of EARN2 is to improve the performance of the whole network reducing the load on the key backbone sites in Europe and exploiting several existing IP lines between the USA and European network organizations like NORDUNET, JANET and others. These two aims were achieved through the second half of 1991. Thanks to EARN2, better services are now available for the users. The statistical data collected by EARN showed in fact a fair reduction of the load on the key backbone sites and a general improvement of the network performance. The EARN2 goal for 1992 is to implement the BITNET II approach with full connectivity among the core sites. A full mesh configuration is the logical further step of the regionalization process. It is highly desirable because, besides the overall improvement of services already achieved, it facilitates network management at the routing level both inside Europe and with BITNET and the other Cooperating Networks.

Keywords: EARN; NJE/IP; routing; network performance evaluation; network management.

A little bit of history

The purpose of the EARN Regionalization plan was to implement the solutions to a number of problems which were affecting the EARN network during 1990.

There were a number of serious bottlenecks which arose as a result of increasing traffic and the network topology EARN had at that time (almost all the physical links were connected to two sites in Europe, namely CNUSC, Montpellier (FR), and CERN, Geneva (CH)). In the same period a set of IP based lines became available connecting several EARN sites both inside Europe and toward the USA. The use of these resources required a radical reconsideration of the structure of EARN and in November 1990 the EARN Routing Project Group discussed and afterward proposed a working plan.

The proposed restructuring was based on the concept of "regionalization", the separation of the network into areas or regions not necessarily bound to national borders. The criteria for the division among regions were based on the density of nodes, traffic, and need of services in the area. The RPG also proposed to select among the EARN international sites a set of sites (called in the following chapters "core sites") located at

Correspondence to: Dr. D. Bovio, EARN Office, CIRCE BP 167, F-91403 Orsay Cedex, France. Tel. (+33) 1 6982 3973, Fax (+33) 1 6928 5273, E-mail: hi@frors12.bitnet.

REGION:	CORE SITE:	N of Nodes:
FR, BE, ES, PT, EG, GR, CY, TR, TN	FRMOP11	281
DE	DEARN	197
IT	ICNUCEVM	132
CH, IN	CEARN	59
NL, LU, IE	HEARN	101
UK	UKACRL	2 (+JANET)
IL	TAUNIVM	47
AT, BG, YU, HU, CS	AEARN	37
SE, NO, FI, DK, PL, CIS	SEARN	88

Fig. 1. Current regional subdivision of EARN.

strategic positions within each region from the point of view of traffic patterns and connectivity. It was essential for these core sites to have the ability to establish virtual NJE links amongst themselves and, in some cases, NJE/IP links with selected partners on the BITNET network. One of the main purposes of the plan, in fact, was to drastically increase the connectivity between Europe and the USA, exploiting the various existing transatlantic lines never used to carry EARN traffic before.

It is possible to carry the Network Job Entry (NJE) protocol on different network layer protocols like BSC, SNA(X.25) and IP. For the latter, particular thanks go to Princeton University where VMNET (the IP application that allows the transport of NJE over TCP/IP protocol (NJE/IP)) was developed. With VMNET, it is also possible to have virtual NJE links over sets of physically interconnected IP lines. The RPG proposed to exploit this IP connectivity being mindful of the continuing evolution of the network services and the conviction that EARN must have the ability to adapt to them. From this point of view the RPG thought that the solution proposed was the fastest and cheapest one available at that time.

A short plan overview

EARN is subdivided into regions. The subdivision, as of today, is shown in Fig. 1. The ISO country codes of the countries are shown as well as the NJE name of the core site and the number of nodes belonging to the region. Each European core site shall be connected to all the other core sites in order to form a full mesh. This reduces the number of store & forward operations inside the backbone and at the same time drastically reduces the complexity of the routing since no routing is needed in a full mesh network because

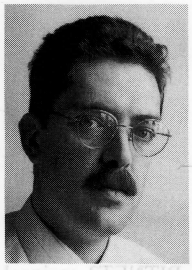

Daniele Bovio, graduated in physics from the University of Milan in 1984, is chairman of the EARN Network Performance Evaluation Project Group and is currently coordinating the work of the EARN permanent technical staff which he joined in March 1990. He worked for the Computer Center of the Institute of Cosmic Physics of the Italian National Research Council (CNR) between 1985 and 1990, initially responsible for the Network Department and then as Technical manager of the Computer Center. During this period he represented CNR in two technical advisory committees, concerning e-mail and IBM communication protocols, of the Italian Academic Network (GARR). In the same period he was involved in the Italian management of EARN as deputy Network Country Coordinator in 1988 and as a member of the Italian Executive of EARN from 1989.

every node has a direct link to all the others. All the other EARN International nodes (not core sites) have a connection to the core site of their region and can thus reach any other International node in 3 hops. The connectivity to BITNET is secured from each EARN region through that region's core site to a BITNET2 regional core site (the concept is schematically shown in Fig. 2). This allows all the EARN International nodes to reach the US in 2 hops.

The implementation of the plan

The implementation of the plan started in March 91 with the definition of the first virtual link between EARN and BITNET core sites. Following the steps defined by the EARN RPG, all the planned links have been added gradually. All the links were tested and monitored before being put into production, including a verification of the consequences to routing prior to the addition of the link. The time scale of the changes to the EARN international topology is shown in Fig. 3.

Thanks to the invaluable collaboration of Michael R. Gettes (coordinator of the BITNET2 activities and lately member of the EARN-RPG) a subset of US core sites was selected to be coupled to their European counterparts. Before the implementation of the plan, only two links/lines were used to carry the traffic between Europe and the USA: the 64K EARN funded line between Montpellier and New York, and the T1 EASINET line between CERN and Cornell. Thanks to the implementation of the plan, it was possible to exploit other existing transatlantic lines and open 9 VMNET links. In particular the Janet/joint funded 512Kb line between ULCC, London and FIX-EAST; the NORDUNET 256Kb line between KTH, Stockholm and Cornell; the Israeli University Consortium 64Kb line between the Weizmann Institute of Science and PSI; and the DFN 128Kb line between GMD, Birlinghoven and the Princeton Plasma Physics Laboratory are now in use.

The overall situation (both inside Europe and to the US) of the links defined is shown in Fig. 4. The dotted lines represent links under test or planned for the near future. The 9 big spots on the right part of the figure are the European core sites, the 9 on the left side are the coupled US core sites. All the other spots represent the international EARN nodes. The traffic patterns among European regions are shown in Fig. 5.

For each region, the amount of traffic exchanged with the rest of the world from October 91 through January 92 is shown. The traffic flow between the US and Europe has been regulated through a careful but uncomplex routing paradigm: each European region is sending through its link to the US only the traffic coming from inside its region. The routing is symmetrical so that returning traffic (from the US to Europe) first negotiates the US backbone and then crosses

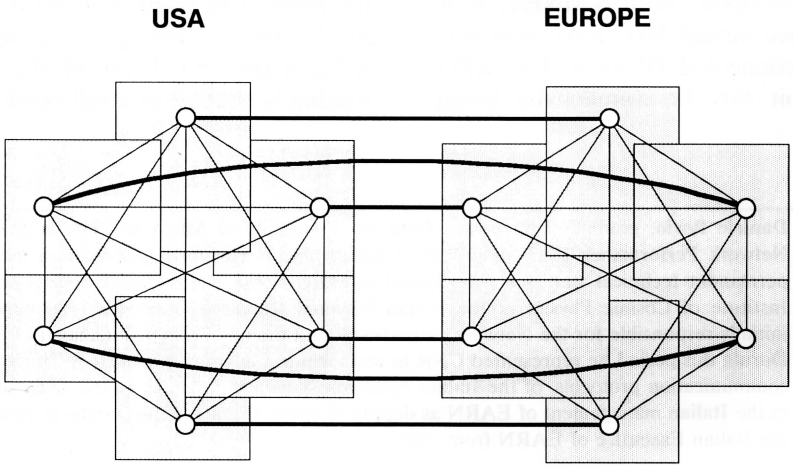

Fig. 2. Connectivity to BITNET is secured from each EARN region through that region's core site to a BITNET2 regional core site.

the Atlantic on the link connected to the European region containing its final destination. The proportion of traffic that each region sent to the US through its link in respect to the total traffic exchanged during January 92 is shown in Fig. 6.

After having implemented the full US-European transatlantic link grid, new traffic patterns showed that it was possible to phase out the EARN funded line to CUNY in New York and to reallocate the money to other cooperative initiatives leading to a better connectivity to the USA. Eventually the EARN line was in fact removed from production last February, cancelled at the end of March, and the money reallocated to the EBONE'92 initiative.

Looking at Fig. 4 it is clear that a high level of connectivity among core sites has been achieved. A few of the European core sites have already achieved full mesh connectivity status (HEARN, ICNUCEVM), and others will do so in the near future.

Results

Besides the consensus between European core and international sites that EARN services have progressively improved during the past 12 months, it is possible to empirically confirm this improvement.

Since October 1990 the EARN Office has collected data about the status of the network, monitoring the up/down time periods of the links (at the NJE application layer) and the number of

LINK		LOCAL TEST	TRAFFIC	MONITOR	BITEARN NODES
TO US					
CEARN	PUNFSV2	3/9/90	9003	9010	9010
ICNUCEVM	UICVM	10/7/91	9108	9108	9108
TAUNIVM	RICEVM1	18/3/91	9103	9104	9104
SEARN	CORNELLC	15/5/91	9105	9105	9106
AEARN	UGA	4/6/91	9106	9106	9107
DEARN	YALEVM	7/6/91	9107	9107	9108
UKACRL	VTBIT	18/7/91	9108	9108	9108
FRMOP11	MITVMA	15/4/91	9105	9105	9112
HEARN	UIUCVMD	12/91	9112	9201	9201
EUROPE					
ICNUCEVM	CEARN	28/2/91	9103	9103	9107
ICNUCEVM	SEARN	30/1/91	9102	9103	9107
ICNUCEVM	AEARN	10/5/91	9105	9105	9107
ICNUCEVM	HEARN	27/4/91	9105	9105	9105
HEARN	SEARN	21/4/91	9105	9105	9105
HEARN	CEARN	22/4/91	9105	9105	9105
HEARN	AEARN	10/5/91	9105	9106	9107
HEARN	DEARN	15/5/91	9108	9108	9108
GBGBOX	FRMOP22	02/9/91	9103	9104	9106
ITGBOX	GBGBOX	06/9/91	9107	9109	9109
DEARN	GBGBOX	08/91	9108	9110	9110
HEARN	FRMOP11	15/7/91	9107	9108	9108
ICNUCEVM	FRMOP11	30/8/91	9109	9109	9110
FRMOP11	AEARN	15/7/91	9108	9108	9108
HEARN	TAUNIVM	08/91	9109	9110	9112
CEARN	TAUNIVM	09/91	9110	9112	9112
ICNUCEVM	TAUNIVM	18/11/91	9112	9112	9112
DKEARN	SUEARN2	10/91	9111	9111	9111
FRMOP11	TAUNIVM	11/91	9111	9112	9112
DEARN	TAUNIVM	11/91	9111	9111	9112
AEARN	TAUNIVM	11/91	9111	9101	9202
ICNUCEVM	DEARN	13/11/91	9112	9112	9112
UKACRL	HEARN	11/12/91	9201	9201	9201
AEARN	DEARN	27/1/92	9202	9202	9203/4
AEARN	UKACRL	11/12/91	9201	9201	9201
UKACRL	TAUNIVM	15/2/92	9202	9203	9207?
UKACRL	SEARN	15/2/92	9202	9203	9206?

Fig. 3. Time scale of changes to the EARN international topology.

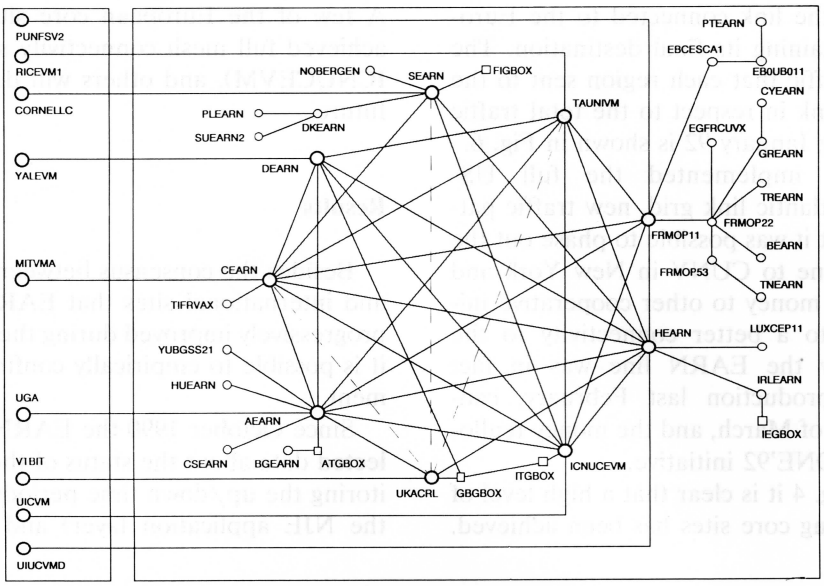

Fig. 4. Overall situation (both inside Europe and to the US) of defined links.

files queued at all the international EARN nodes. The number of files queued is indirectly a measure of the quality of the service offered. A high number of files queued means delay in the transmission and delivery of the traffic. A situation where files are queued constantly on a link means that the users will invariably receive their mail/files later than they should. By these means, we consider the sharp decrease in the number of files queued on the network as showing a tangible improvement (service being inversely proportional to queue size). It must be noted that we do not collect detailed information about each discreet queued file (traffic figures are collected at a macro level). However, we have no reason to believe that the composite characteristics of each queue changes sharply over a monthly interval and thus consider our data intrinsically sound.

Figure 7 shows a breakdown by threshold of the number of files queued (on a monthly average) at all the international links at any time. The thresholds are fixed at values of greater than 50, 100, 200, 500, 1000, and 3000 files queued. As we previously stated, implementation of the plan started at the beginning of March 91, with the definition of the first VMNET links. It is quite clear from the figure that we have experienced a sharp decrease in the size and the frequency of the queues, especially for the low thresholds (50, 100, 200, 500); this indicates an overall improvement and a change in the behavior of the network. As can be seen from the last quarter (1Q92), the percentage of time in which we find more

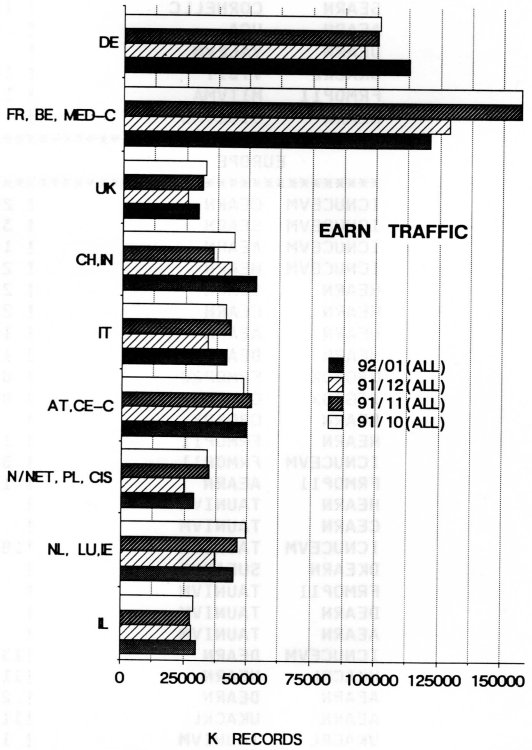

Fig. 5. Traffic patterns among European regions; the amount of traffic exchanged with the rest of the world from October 91 through January 92.

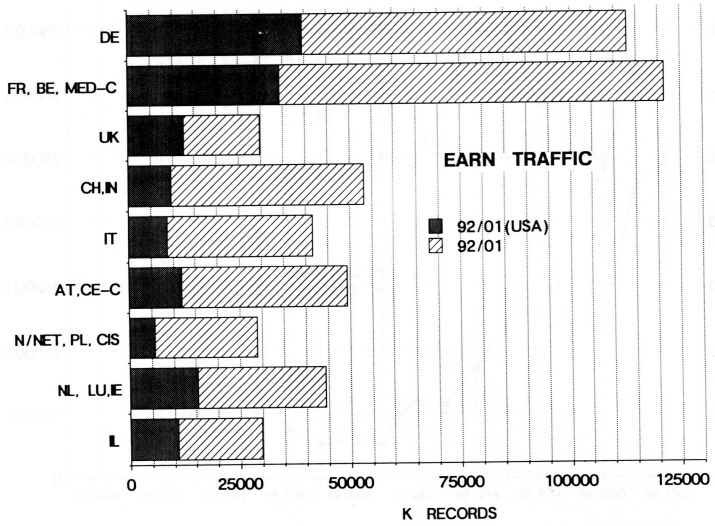

Fig. 6. Proportion of traffic that each region sent to the US through its link in respect to the total traffic exchanged during January 92.

than 50 files queued went down approximately 75% (from 12% in February 91 to 3% in February 92). On the other hand the probability of finding large queues is still high, and comparable with the same period of 1991. This means that the queue formation/resolution is today much faster than in the past. In other words, it is not related to the load of a particular link but rather to the up/downtime periods experienced by that link.

By comparing the monthly average of the queued files and the traffic data (Fig. 8), it is also clear that the reduction of the queues is not bound to the reduction of the traffic volume. The two curves are, in fact, following a very different path: the traffic is increasing (except the traditional lower values for the Summer and over Christmas) while the average of the queued files is decreasing.

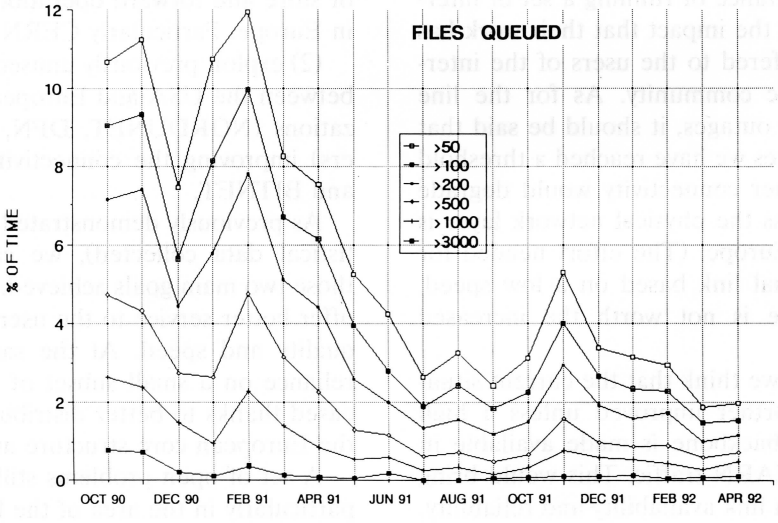

Fig. 7. Breakdown by threshold of the number of files queued (on a monthly average) at all the international links at any time.

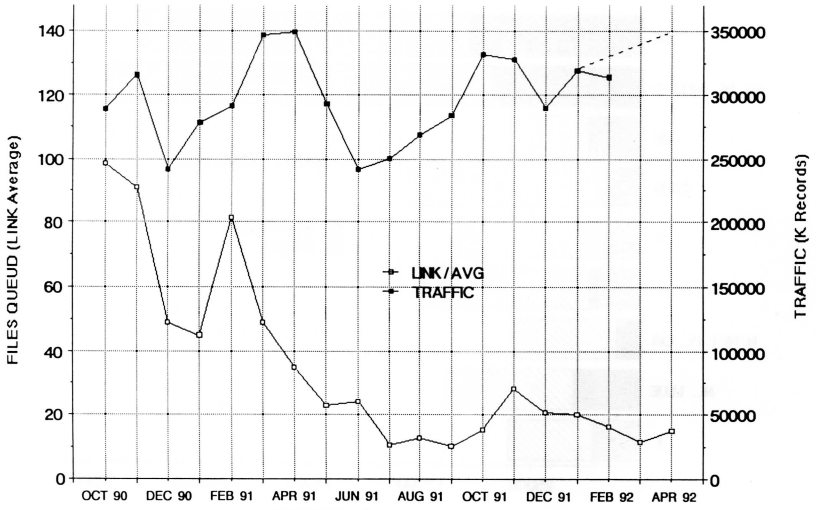

Fig. 8. Monthly average of queued files compared with traffic data.

The future

The immediate future goal of EARN is quite clear: to reach a full mesh configuration inside Europe. It is also quite obvious, however, that this is not sufficient to achieve global improvements in terms of management and services. As we earlier stated, one of the major problems degrading network performance is link downtime period. This is a double sided problem: on the one hand, the availability of links at the application layer and on the other, the physical line performances and outages. The link downtime periods could be reduced by offering a set of NJE monitoring tools in order to facilitate alerting them to the importance of running a set of international links and the impact that their work has on the services offered to the users of the international Academic community. As for the line performances and outages, it should be said that for several core sites we have reached a threshold above which further connectivity would degrade performance unless the physical network layer is improved inside Europe. (The effort needed for monitoring a virtual link based on a low speed, low reliability line is not worth the increased connectivity).

In other words we think that the current situation cannot be further improved unless a high speed redundant backbone is made available in Europe to carry EARN traffic. This would bring an improvement in link availability and reliability. The EBONE, as well as other similar initiatives, could be regarded as a first step in this direction. EARN is currently cooperating with all the technical groups active in this field and the cooperation will be continued to achieve this desired result. A realistic time scale to reach a full mesh configuration seems to be in this framework the last quarter of 1992.

Conclusions

The two main goals of the EARN Regionalization plan can be summarized as follows:

(1) improve the performance of the whole network by reducing the traffic load (i.e. the number of store and forward operations) on the key sites in Europe. Particularly CERN and CNUSC.

(2) exploit previously unused transatlantic lines between the USA and European network organizations (NORDUNET, DFN, JANET, and others) improving the connectivity between EARN and BITNET.

As previously demonstrated (thanks to the statistical data collected), we can today consider those two main goals achieved. As a result we can offer better service to the users, both in terms of quality and speed. At the same time, EARN's reliance on a small subset of key sites has been eased thanks to better distribution of traffic over the European core structure and to the USA.

A set of open problems still stands, of course, particularly in the area of the link uptime/downtime periods. It is our understanding that this can

be addressed and our service to users be further improved. This will be our challenge for the coming year.

Acknowledgement

I am pleased to thank the EARN Routing Project Group, which discussed and proposed this successful plan and in particular the chairman of the group H. U. Giese. I also would like to particularly thank M.R. Gettes, from Princeton University and coordinator of BITNET II. He has been extremely helpful and patient in helping "this side of the pond". Thanks also to my colleague G. Lloyd who helped me review this document.

References

[1] P.A. Olenick, RSCS Over TCP/IP—VMNET, A Technical Overview and Definition of the BITNET II Protocol, Princeton University, 26 April 1989.
[2] H. Nussbacher, EARN-IP Project Group NJE-IP Policy and Procedures, EARN Document (EXEC14 90), March 1990.
[3] D. Bovio, U.Giese and H.Nussbacher, Proposal for the Regionalization of EARN, EARN Document (BOD7 91), April 1991.
[4] M. Gettes, L. Varian and P. Olenick, BITNET2 Proposal, 21 September 1990 (revised version).
[5] D. Bovio, Definition of the EARN2 Regions, EARN Document (EXEC83 91), August 1991.
[6] D. Bovio, Backup for Transatlantic IP Links, EARN Document (EXEC70 91), August 1991.
[7] D. Bovio and U. Giese, Regionalization Plan Progress Report, EARN Document (BOD31 91), October 1991.