Distribution via the Internet DNS of RFC 1327 mapping rules and the INFNet distributed gateway system

Pierluigi Bonetti <Bonetti@infn.it>

Claudio Allocchio <Allocchio@elettra.trieste.it>

Antonia Ghiselli <Ghiselli@infn.it>

Abstract

There is a costant increase in the traffic load between RFC822/SMTP and X.400 mail protocols, thus new and easy to use interconnection tools are required. In particular, the duplication and decentralization of gateways between the above protocols requires that they should be usable also by non experts. The new gateways should be maintenance free. One of the major problem is RFC1327 mapping table maintenence. A mechanism for automatic distribution of mapping rules by means of Internet Domain Name System (DNS) has been implemented and tested; the positive results are presented here. In particular, we specified an Application Programming Interface (API) and developed a library which allows a gateway to interact with DNS in order to obtain mapping rules, without the need of statical tables. Using this basic tools, we developed also a new RFC 1327 gateway which does not require local maintenance. This new gateway was distributed over our INFNet backbone, resulting in an efficient and fault tolerant two protocol (X.400 and RFC 822/SMTP) E-mail distribution backbone.

I. Introduction

Electronic mail is one of the most popular network services among users. The set of rules by which two applications can transfer messages is called *mail protocol*: this term includes both trasmission rules (E-mail transport protocol) and message format description (format protocol). On the other hand, the computer programs transmitting messages along the network are often called *Message Transfer Agents* or *Mailers*, while the user interface is called *User Agent*.

There are many E-mail protocols, which differ for transport type, message format and offered additional services (such as positive delivery reports, address and reacheability check, etc.). However, just a few among these E-mail protocols will probably play a first rate role in the next years.

The E-mail transport protocol mainly based upon network protocol TCP/IP and used over the Internet is called SMTP (Simple Mail Transfer Protocol) [1]. It uses an addressing scheme organized in "domains" and described in RFC 822 [2]. Since its wide diffusion, it represents a *de facto* reference point for E-mail.

Its main features are the simpleness of configuration and the very low maintenance activity. Moreover, the domain-based addressing schema is rather practical and it is becoming familiar among large user groups. However, one of the RFC 822 major limits is the restriction to use US-ASCII 7-bit character set for the header and the body of messages.

MIME (Multipurpose Internet Mail Extensions) [3] tries to bypass this strong format restriction, leaving the SMTP transport system unchanged.

The E-mail official standard, issued from international organizations (CCITT and ISO), is called X.400 [4, 5]. It describes both a transport and a format for messages and exists in 1984 and 1988 versions. Due to the fact that it's an OSI (Open Systems Interconnection) standard and to the completeness of services offered, it's becoming one of the most used E-mail protocol over the large national and international communication infrastructures.

In X.400 a message content can be of different types, allowing thus natively multimedia transmission. Moreover the addressing schema is usually based upon well known logical-geographical informations only, leaving physical network infrastructure and host names out of consideration. Complexity of installation and configuration of an X.400 Message Handling System (MHS) and a still large lack of User Agents, however, makes this protocol almost unknown to the end users.

It is likely that during the next years most E-mail users will probably continue to use the tools already in use, i.e. mainly a domain-based addressing scheme (RFC 822) and their current User Agents. However the E-mail transport system over large infrastructures will be more and more based upon X.400, especially within non R&D communities.

While it is desirable a unique E-mail protocol, in the meantime we thus need to use gateways. In particular, it is absolutely needed to have gateways between RFC 822 and X.400 able to support the already intense traffic in the two directions.

II. The mail gateways

Studying the characteristic of the most used multiprotocol[1] mail gateways, such as MX [6], PMDF [7], PP [8], GIVEME [9], some common aspects arise.

All of them are designed for centralised architectures, i.e. for sites connected with all the supported protocols.

They are complex objects, which heavily depend on network connectivity status and on directory services of involved networks. They need thus a continuous monitoring and maintenance.

Some multiprotocol gateways act a direct conversion between all the possible couples of supported protocols, while some others use an intermediate internal format. Moreover, these architectures often do not have specific tools to keep into account the different amount of traffic through the various protocols. Often, a First-in First-out criterion in queue management, with sequential message processing, can produce a dangerous bottleneck whenever there are problems in sending a message.

Typically, major gateways are just a few ones in each national network, located in strategical positions for the traffic. A problem in one of these gateways causes large scale effects; this critical role of the centralized gateways is one of the greatest limit of today E-mail distribution systems.

Finally, most gateways doesn't yet support X.400-88 nor multimedia conversion.

II.A. Possible evolutions

Assuming that the protocol diffusion could evolve as described, let's try to analyze a plausible scenario.

Multiprotocol gateways are, without doubt, still adequate to guarantee communications among less used protocols, such as Mail-11 (over DECnet), NJE (over BITNET/EARN), UUCP (dial-up), etc. and from them towards X.400 and RFC 822/SMTP. They will probably continue to be used for this purpose.

Today's amount of traffic between RFC 822 and X.400, however, is so high to make existing installations inadequate for fast real time processing, resulting in the introduction of delays.

Increasing the number of multiprotocol gateways does not appear to be the best solution, since the biggest need for conversion concerns only two of the supported protocols and these gateways are too complex to be mantained by non-experts.

The hypothesis to build "simple" gateways, i.e. only between RFC 822 and X.400 protocols seems to be a better solution in order to reduce the heavy load on multiprotocol gateways.

Similar objects could then be installed in multiple distributed copies and not necessarily in a centralized way, as for the multiprotocol gateways.

Decentralization of gateways, and therefore local format conversion of messages, has the positive effect to reduce network traffic and to increase fault tolerance and conversion efficiency. But, on the other hand, new coordination and management problems arise. We will discuss them shortly.

Firstly, new gateways should be compatible with both versions of X.400 (84 and 88).

Moreover, they should be compliant with RFC 1327 [10] (which just describes the mapping between X.400/84 - 88 and RFC 822), and with RFC 1494, 1495 and 1496 [11, 12, 13] (which describe the conversion between X.400 and MIME). Compliance to RFC 1405 [14] is required if gateways have to support also the Mail-11 protocol.

Anyhow, the true feature of the new type of gateway consists on its capability to be easily distributed in many sites.

III. Distributed RFC 822 - X.400 gateways

A "replicable" gateway needs the following features:

- easy installation and configuration;

- low resource (disk space and CPU) consumption in order to not require a dedicated host;

- do not require human intervention for maintenance.

The last two requirements are correlated: one of the major maintenance activity for current gateways consists of the periodic update of statical address mapping tables installed on the gateway itself.

III.A. Mapping tables update

At the time of writing, the RFC 1327 mapping tables are maintained by an International Coordination Service: an up to date version is periodically issued and made available via ftp and E-mail.

Each gateway manager must get those tables and make them available to his own gateway. This retrieval can be easily automated, satisfying thus our third requirement.

[1] Here "multiprotocol" should be intended as "between three or more protocols."

But, in addition, if we totally suppress all static mapping tables, we will comply also with the second requirement, regarding to disk space requirements.

In this work we consider the last approach.

In order to overcome the mechanism of manual distribution of static tables (usually through ftp or E-mail), there are two alternatives: to move around information through X.500 or through the Internet DNS.

The current problem in using the X.500 approach is the lack of available servers. The few existing ones will be exposed to an excessive amount of query, as soon as the number of gateways grows.

The other approach, parallel and supplementary to X.500, is described into the document "Using the Internet DNS to distribute RFC 1327 address mapping tables" [15] from IETF X.400 Ops Working Group and the RARE WG-MSG Group.

In the next section we describe the first implementation of the latter solution.

IV. Distribution of mapping rules between RFC 822 and X.400 through Internet DNS

At first let's summarize briefly the concepts explained into [15].

The Internet Domain Name Service (DNS) [16, 17] is a well tested mechanism to distribute useful informations over Internet for various network services.

There are many advantages in using a distributed system to propagate mapping rules:

- It avoids fetching and storing of entire mapping tables by every host that wishes to run an RFC 1327 gateway or needs them for other purpouses;
- Mapping rules updates are propagated more rapidly than using static mapping tables (currently issued monthly). Moreover, updates are done in a decentralized way;
- It is possible to determine the mapping used by a remote gateway querying the DNS (remote debugging);
- Other tools, like address converter and User Agents, can take advantage of the real-time availability of RFC 1327 mapping rules.

The proposed mechanism to encode information into DNS uses a new Resource Record, called "PX" (Pointer to X.400/RFC822 mapping information). It will contain mapping rules in the format:

<rfc822-domain> <x400-in-domain-syntax>

The format is the same for rules going from RFC 822 to X.400 and in the opposite direction.

Mapping rules from RFC 822 to X.400 use as key an Internet domain name, already present in the DNS tree. The new Resource Record will contain the corresponding rule.

The key to search a mapping rule from X.400 to RFC 822 is an O/R address: to store these rules we thus need a new subtree into DNS. Since O/R address structure is country-oriented and ISO two letter country code are the same used for the RFC 822 country top-level domains, a subtree under each country top-level is created.

To avoid collisions with the existing tree, a unique reserved domain, called "X42D" (X.400 To Domain), is placed under each country top-level domain. The subtree under X42D is labelled with O/R addresses, and mapping rules from X.400 to RFC 822 are stored into the PX Resource Record.

An example of some new subtrees of DNS is shown in Figure 1.

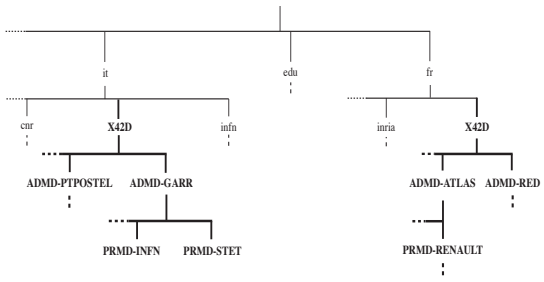

Figure 1: New DNS structure with RFC 1327 mapping rules

The update of mapping rules can take place in a totally distributed way, as for any other informations within DNS. Initially, a single entry point for each country will be defined, then the authority will be gradually distributed.

The syntax used for O/R addresses and RFC 1327 mapping rules is not consistant with Internet domain name syntax. DNS syntax, defined in RFC 1034, limit the character set for valid domain names to alphanumeric plus the symbol "-" (hyphen). Let's call this set <alphanumhyphen>.

RFC 1327 mapping rules contains characters from a set called <IA5-PrintableString>[2] and <alphanumhyphen> is a proper subset. A mechanism to map characters in <IA5-PrintableString> and not in

[2] <IA5-PrintableString> is constituted by the following characters: A..Z a..z 0..9 <blank> ' + , - . / : = ?

<alphanumhyphen>, into symbols of the latter set is thus also defined.

IV.A. Implementation of the system

The distribution of mapping rules via DNS is intended to provide an easy to use tool for obtaining these data. RFC 1327 rules must be used in gateways, but also a number of other applications could provide better services by using them. To make life easier for implementors, we have specified an Application Programming Interface (API) describing tools to perform the encoding and the decoding needed to store and retrieve RFC 1327 mapping rules into the Internet DNS. Moreover, we also implemented the application library following the API specifications.

In order to minimize changes in existing gateways and enable them to use the new automatic distribution service, we designed the interface to be compliant with the format of rules retrieved in static tables.

The query to a nameserver asks for the best matching rule, using an RFC 822 domain or an X.400 O/R domain as the key. Possible results are the same as for any other DNS query: the mapping rule is found, then it's used; the rule is not found, then use the other RFC 1327 specifications to convert the address; time-out, no answer from the nameserver. In this last case, not foreseen in RFC 1327, the gateway or conversion operation should be delayed and then retried at a later time. The latter case should anyhow be avoided with a carefully planned DNS zone duplication.

The DNS query described mechanism is shown schematically in Figure 2.

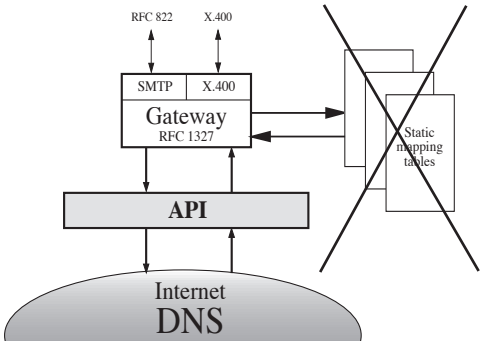

Figure 2: Interaction between a gateway and the Internet DNS

Since the implementation of the described mechanism requires changes in the nameserver software, it will need some times before it become globally available. A transition plan is thus needed to allow the use of the new DNS service since now.

At first the mapping rules will be stored centrally, i.e. in a single master point of DNS, and the application library will ensure a transparent access to the information.

Then, as soon as new nameserver implementations spread around, the authority will be distributed on a per country base, i.e. each countries will take in charge their own DNS subtree for mapping rules. At this stage we will be in a similar situation as the current mechanism for mapping rules creation: a national mapping authority prepares the national mapping tables, but instead of submitting it to the coordination center for redistribution, the national authority will insert it into DNS, just under to national top level domain.

At last the authority can furthermore be distributed within the national tree.

In the future, with the development of X.500 servers, there will also be growth of mapping data available via X.500. Therefore, coordination of the DNS distribution of mapping rules with the equivalent project via X.500 is needed. The ideal situation being a global unique dynamic distribution system where applications can find a mapping rule transparently, we must think of appropriate exchange mechanisms to pass information between DNS and X.500. Dumb data duplication in both systems should be avoided. On the contrary, transparent gateways between DNS and X.500 should be implemented in the mean time. Some works in this field has already been done with encouraging results [18].

To immediately test the system, we created a temporary subtree into DNS, under the domain "X400.IT" and we stored in it all mapping rules, using the "PTR" Resource Record, already available in current DNS implementations. Two branches were created, one labelled with RFC 822 domains and one with DNS-encoded O/R addresses.

We developped some tools, useful to create a unique SOA file, starting from the three files containing the International Mapping Tables.

We also developed a first application library to be used within the "X400.IT" subtree. All this software proved to be functional, and is thus now being converted into the final version in accord with [15][3] .

Using the library[4] with calls specified into the API, it is possible since now to obtain RFC 1327 mapping

[3]The final version of the library and tools will be available as public domain software: contact the Italian National Institute for Nuclear Physics (INFN) at mailing@infn.it.

[4]The library is written in C and functions are callable also from applications written in other languages.

rules, querying the Internet DNS. These rules can be used from any gateway, address converter, User Agent and everything that needs mapping rules between X.400 and RFC 822 to operate.

As a first immediate result of the availability via DNS of mapping rules, we realized an RFC 1327 address converter, called "POSTINO", which uses the library following the API: POSTINO just calls a library function when it needs a mapping rule and then performs the appropriate conversion. It can thus run on any host with DNS access, without any other requirements, and it was also ported and tested on various operating systems and platforms, including both Unix and VMS.

V. The INFNet RFC 1327 distributed gateway system

The first positive tests with the address converter, made us go further, developing an RFC 1327 gateway based on our API and library. This new gateway, which was named GIVEME2, is based on DEC MR and MRX[5] for the X.400(84) side and on a standard TCP/IP socket interface on the SMTP side.

As GIVEME2 uses DNS to retrieve its mapping rules, its very simple installation can convert any MR/MRX MTA into an RFC 1327 gateway which does not require any specific and complex additional maintenance to the usual MTA one. In fact GIVEME2 does not have any additional queueing system, as it uses the DEC MR usual queues, nor any extra routing system, as it uses DNS on the RFC 822/SMTP side and the usual MR/MRX routing data on the X.400 one.

Actually an additional automatic X.400 MR/MRX routing tables update mechanism mantained by a central coordination service also removes from the local manager the X.400 routing database maintenance, making possible for GIVEME2 installations to be totally maintenance free.

Figure 3 shows the logical configuration of an INFNet site.

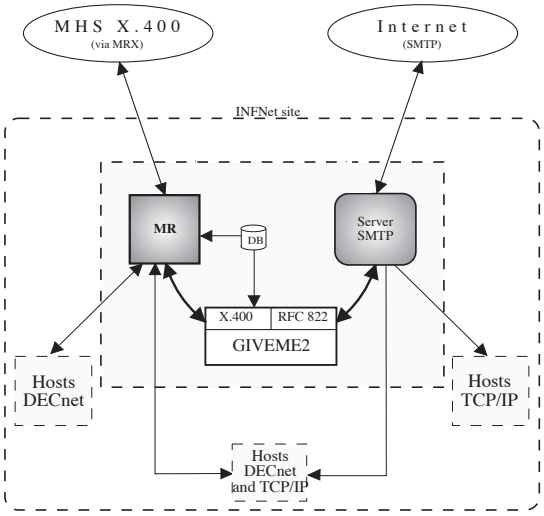

Figure 3: A tipical INFNet site configuration

An additional feature of the GIVEME2 gateway is the RFC 822/SMTP local delivery service, which in conjunction with other MR tools, allows to implement a full local distribution system, reaching users nearly on any of the existing platforms. This also allows a full implementation of purely logical addresses both in RFC 822 and in X.400, i.e. addresses where there is no reference to the actual final host where the end user prefers to receive his/her mail.

The INFNet E-mail system was previously based on an X.400 transport backbone, implemented using MR and MRX, and one central RFC 1327/RFC 1405 multiprotocol gateway allowing message interchange among X.400, RFC 822, NJE and Mail-11. An additional multiprotocol gateway was used as a backup of the main one.

Using GIVEME2 we turned any X.400 MTA serving a local domain on INFNet into an RFC 1327 gateway. This set of gateways has allowed us to strongly reduce the heavy load on the central gateways: in fact if a message needs to be delivered via RFC 822, it can immediately be converted locally and delivered directly to its final destination. Moreover incoming traffic from RFC 822 is not forced into the bottleneck of the central gateways, and can enter the X.400 backbone locally.

Figure 4 schematically describes the two transport system with a gateway in each site.

[5]GIVEME2, like its multiprotocol ancestor GIVEME, does not need DEC DDS to register X.400 subscribers.

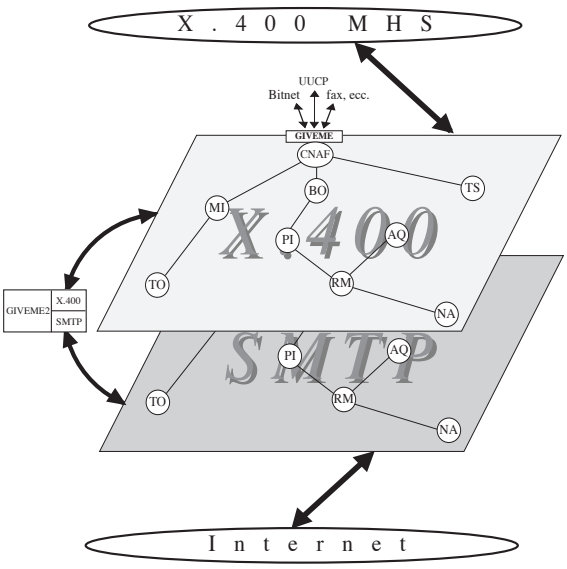

Figure 4: The two parallel backbone E-mail architecture of INFNet

Last, but not the least, such a number of equivalent RFC 1327 gateway let us implement also a fully redundant conversion system. In fact, even if for any other protocol different from X.400 and RFC 822 we still need to cross the central multiprotocol gateways, any local GIVEME2 gateway can act as a backup for any other local gateway.

As an overall result, the whole E-mail distribution system over INFNet dramatically increased its performance and reliability. Moreover, reducing the single gateway load to local traffic only, we can start to seriously plan for a migration to X.400(88) and hence for support of Multimedia local gateway conforming to RFC 1494, RFC 1495 and RFC 1496.

Additional development plans, regarding User Agents being able to transparently accept RFC 822 or X.400 addresses are also being considered.

VI. Conclusions

There is a need to facilitate the conversion between the mail protocols RFC 822/SMTP and X.400.

We implemented a dynamic mechanism to distribute address mapping rules, in an efficient and fast way, to the various objects that needs them, firstly mail gateways.

It's then possible to duplicate and distribute maintenance-free gateways, as regard to RFC 1327 tables.

Some additional tools allowed us to implement a fully distributed and redundant RFC 1327 gateway system over our national network E-mail system. Moreover, a purely logical addressing schema is fully possible both in RFC 822 and X.400 syntax.

The reduced traffic load also open us serious chances to implement local multimedia gateways.

VII. References

[1] Jon B.Postel. "Simple Mail Transfer Protocol". Request for Comments 821, DDN Network Information Center, SRI International, August 1982.

[2] David H.Crocker. "Standard for the Format of ARPA Internet Text Messages". Request for Comments 822, DDN Network Information Center, SRI International, August 1982.

[3] Nathaniel S.Borenstein, Ned Freed. "MIME - Multipurpose Internet Mail Extensions". Request for Comments 1521, DDN Network Information Center, SRI International, September 1993.

[4] "Message Handling: System and Service Overview". International Telegraph and Telephone Consultative Committee, 1988. Recommendation X.400.

[5] ISO/IEC International Standard 10021". Dallas, October 1988.

[6] "Message Exchange Guides V.3.1". Engineering Computing Services, Rensselaer Polytechnic Institute, Troy, New York, March 1992.

[7] Ned Freed, Mark Vasoll. "System manager's guide to PMDF-822 V.3.0". May 1988.

[8] J.P.Onions, P.Cowen, A.da Cruz, S.Kille. "PP - Vol.1: Installation and operation (version 6.0)". X-Tel Service Ltd., University College London, December 1991

[9] Claudio Allocchio, Antonia Ghiselli. "The INFN GIVEME 987 gateway". Computer Networks and ISDN systems, vol.19 (1990), pp. 255-260.

[10] Steve Hardcastle-Kille. Mapping between X.400 (1988) / ISO 10021 and RFC 822. Request for Comments 1327, DDN Network Information Center, SRI International, May 1992.

[11] Harald T.Alvestrand, Steven J.Thompson. "Equivalences between 1988 X.400 and RFC-822 Message Bodies". Request for Comments 1494, DDN Network Information Center, SRI International, August 1993.

[12] Harald T.Alvestrand, S.Kille, Robert S.Miles, Marshall T.Rose, Steven J.Thompson. "Mapping between X.400 and RFC-822 Message Bodies". Request for Comments 1495, DDN Network Information Center, SRI International, August 1993.

[13] Harald T.Alvestrand, James A.Romaguera, Kevin E.Jordan. "Rules for Downgrading Messages

from X.400/88 to X.400/84 When MIME Content-Types are Present in the Messages". Request for Comments 1496, DDN Network Information Center, SRI International, August 1993.

[14] Claudio Allocchio. "Mapping between X.400 (1984/1988) and MAIL-11 (DECnet mail)". Request for Comments 1405, DDN Network Information Center, SRI International, January 1993.

[15] Claudio Allocchio, Antonio Blasco Bonito, Bruce Cole, Silvia Giordano, Robert Hagens. "Using the Internet DNS to distribute RFC1327 Address Mapping Tables". IETF X.400 Ops Working Group, Internet-Draft v5.2, DDN Network Information Center, SRI International, April 1994.

[16] Paul Mockapetris. "Domain Names - Concepts and Facilities". Request for Comments 1034, DDN Network Information Center, SRI International, November 1987.

[17] Paul Mockapetris. "Domain Names - Implementation and Specification". Request for Comments 1035, DDN Network Information Center, SRI International, November 1987.

[18] Costa, A., Macedo, J., Freitas. V. "Accessing and Managing DNS information in the X.500 Directory". Proceeding of the 4th Joint Networking Conference, Throndheim, NO, May 1993.

Author Information

Pierluigi Bonetti was graduated in Computer Science at the University of Bologna with a dissertation on E-mail developed at the National Networking Center (CNAF) of INFN, the italian National Institute for Nuclear Physics. Now he works at CNAF with a fellowship. His research focuses on E-mail protocols and architectures, Directory Services, multimedia and network information systems. He is author of the "POSTINO" RFC 1327 address converter and co-author of the "GIVEME2" mail gateway. He partecipates also to the X.500 deployment for INFN.

Claudio Allocchio is the responsible for electronic mail and network services at Elettra Synchrotron Radiation facility in Trieste and the italian national responsible for e-mail services at INFN. He is also the chairman of the italian National Research Network (GARR) Working Group on E-mail, and coordinates the italian ISO registration authority. Since 1986 he partecipated to the development of the INFN national plan for E-mail services, including the logical E-mail naming scheme. He is also an active member of the HEPnet Technical E-mail committee, of RARE WG-MSG and since 1990 of the Internet Engeneering Task Force (IETF). He is also the developer of the first GIVEME 987 gateway and one of the developers of the current GIVEME2 1327 gateway.

Antonia Ghiselli, graduated in Physics from the University of Bologna, is a researcher at INFN, in Bologna. Since 1980 she has been working in setting up, managing and software developing for the INFN network. In addition, from March 1986 till June 1987 she has been at CERN, as a scientific associate, working in the on-line group of the LEP ALEPH experiment. Now she is network manager of INFNet and partecipates to the coordination and planning of the GARR network.