INTERNATIONAL ATOMIC ENERGY AGENCY
UNITED NATIONS EDUCATIONAL, SCIENTIFIC AND CULTURAL ORGANIZATION
INTERNATIONAL CENTRE FOR THEORETICAL PHYSICS
I.C.T.P., P.O. BOX 586, 34100 TRIESTE, ITALY, CABLE: CENTRATOM TRIESTE

H4.SMR/585-11

FIRST INTERNATIONAL SCHOOL ON COMPUTER NETWORK ANALYSIS AND MANAGEMENT

(3 - 14 December 1990)

The INFN GIVEME 987 Gateway

Claudio Allocchio

Area di Ricerca
Sincrotrone Trieste S.p.A.
Padriciano, 99
Trieste

The INFN GIVEME 987 Gateway

Claudio ALLOCCHIO
Istituto Nazionale di Fisica Nucleare, Via Castagneto 47, I - 34127 Trieste, Italy (Tel: =39 40 3758523, Fax +39 40 226338, E.mail: ALLOCCHIO@elettra-ts.infn.it)

Antonia GHISELLI
Istituto Nazionale di Fisica Nucleare, CNAF, Bologna, Italy

Abstract: The major aim of the INFN mailing system project is to provide a user address scheme based on logical fields, completely independent from the network communication infrastructure, and in agreement with the X.400 recommendations. Full connectivity with other existing mailing systems must be maintained. The GIVEME 987 gateway has now been implemented with address mapping and operations procedures as defined by the RARE MHS Project. An X.400 user interface, satisfying typical research environment requirements, has been agreed and developed with DEC in order to smooth the user transition from RFC822-style addresses to X.400 Standard Attribute Addresses.

1. Introduction

"How should I send an electronic mail to my colleague in Princeton?". Similar questions have come into the offices of user support organisations over the years. The confusion arises over the different addressing schemes in use across the networks. These schemes, which in the early days of networking developed to give a degree of independence from the network "hardware" addressing were hierarchic, postal style "logical" addresses indicating the geographic and organisational location of the addressee. Various schemes evolved. The Internet "domain" style, being prevalent in the USA, became the North American de facto standard formalised as RFC822. European standardisation led to the wide acceptance of the CCITT X.400 specification based on the OSI mailing model and the subsequent implementation by the RARE R&D MHS Pilot Project.

This "two-world" scenario of X.400 and RFC822 is the environment and motivation for the development of conversion gateways. The RFC987 specifications and subsequent enhancements provide the future development route.

2. Addressing Scheme

A user-friendly mailing system uses attributes to identify a sender/recipient in a tree structure corresponding to country, organisation, site/departmental unit and name. The branching levels can vary depending on the internal structure of each organisation. The network name of a computer, workstation or terminal is not required.

For example, we could refer to Mr. Tizio, an Italian researcher, by attributes in the following way:

he is "italian" => country = IT
he works with INFN => organisation = INFN
he works in Roma => unit = Roma
his name is Tizio => surname = Tizio

This method uses "labels" to identify the various attributes of the address, making the order in which we write them irrelevant. The other possible approach is to define a positional order for the attributes, thus avoiding the need to label them individually:

Tizio Roma INFN IT

To be more precise, the first approach has been adopted by X.400 specifications defining the Standard Attribute Address (SAA) format, in which our user's address becomes:

C = it; ADMD = garr; PRMD = INFN;
OU = Roma S = Tizio;

The latter approach is typical of RFC822 addressing scheme:

Tizio@Roma.infn.it

Both schemes can be used to implement our purely logical naming scheme; in fact our address prototype is:

C = country
ADMD = network administration;
PRMD = organisation;
OU = site;
S = surname;
G = givenname;
or
givenname.surname@site.infn.it

where "givenname" and "surname" identify the person, possibly using other attributes to solve ambiguities, and "site" is the name of the city or the Laboratory where our correspondent works, equivalent to the Organisation Unit (OU). There are advantages and disadvantages in both address formats: RFC822 is shorter and easier to type, but its element are strictly positional. It is easy for the sender to misplace one element, especially if he is unfamiliar with the structure of the recipient organization. On the other hand X.400 SAA is more "position tolerant" but introduces some elements like ADMD and PRMD - Administration Management Domain and Private Management Domain - which are useful for the "post-office" people (the MHS at transport level), not for the final users. They were introduced by the various national PTT bodies when specifying the X.400 recommendations: they apply to a world where all the national "private" networks connect only to their national public service, and international connectivity was provided by ADMD-to-ADMD connections. In reality X.400 protocols started to run in an efficient way on private networks much more than on public services. The topic of address formats has been well treated in the literature and at previous RARE networkshops [1,2].

Our fundamental goal is to establish an OSI mailing system, thus the primary choice is X.400 SAA as an addressing scheme, but we included the RFC822 address format in our mailing project with the same dignity as X.400; it is up to the user to choose, in a transparent way, which format he prefers to use. In a later section we will describe how this was implemented.

3. Network infrastructure

INFNet, the INFN network, consists of many different computers, from different manufactures like DEC, IBM, APOLLO, SUN, Apple and so forth, and its communication infrastructure is integrated into the Italian GARR backbone, based on multiple protocols like DECnet, X.25, IP and SNA. Moreover it is inerconnected with ITAPAC, the Italian public X.25 network, and with other private international research networks like EARN/BITNET, HEPNET, EUnet and INTERNET. Thus the main idea during the transition to a full OSI system was to start using X.400 protocols at upper levels of the OSI stack, using currently available low-level protocols as transport, in a fully transparent way for the users. In this way they won't notice at all the low-level transitions, being able already to use all features available in X.400 MHS.

The Message Router (MR) system, from DEC, has been chosen as a basic mailing transport system, for the following reasons:

- It uses DECnet protocols in an efficient way, which are still the widest used protocols within the HEP community, and gives the possibility to distribute the routing, avoiding bottlenecks.
- Several MR gateways already exist giving access to the most important mailing systems like MRX/X400.SNA/PROFFS, and UMC/TCP-IP.
- It will migrate to full OSI X.400 with DECnet phase V.
- It allows an easy user interface to develop applications.

Using the simple routing rules available in MR we then developed a transport system where the routing scheme is based on C, ADMD, PRMD, 0, OU, obtaining an efficient Message Transfer Service within INFNet.

Our MHS must, however, communicate efficiently with the other non-X.400 communities in the world, and with any computer connected onto our LANs, including those not yet supporting X.400 protocols. Hence, the absolute need to develop a multi-protocol flexible gateway with a wide range of configuration capabilities, from minimal LAN application up to the international service level.

4. The General Gateway Problem

Two-protocol and multi-protocol gateways can be distinguished. In the first case a gateway is nothing else than a conversion table. In the latter case the mail format and protocol which should be used to send it out are determined after the incoming message has already been decoded. Moreover, a message can be sent to multiple destinations on different networks thus requiring different output formats. The solution is to have an "internal format" for our object from which we convert from and to any of required actual external formats.

One of the easiest ways to obtain a flexible architecture is to use modular strategy. In hard-ware terms there are one or more buses, with their own microprocessor controlling the traffic among the various cards which are plugged in as needed. The same architecture can be successfully used for software: the communication bus becomes the internal format for our object (the message), the microprocessor is our gateway engine, the micro-code are the routing and conversion tables and the cards are our interfaces to the needed external protocols.

Let's now look at the world a gateway must talk to: there are essentially two kinds of external agents to connect with: User Agents (UA) to deliver the messages directly to the final users and other Message Transfer Agents (MTA) which can be other transport elements or other gateways used for relaying purposes. In our bus architecture thus there will be two classes of software cards: one for UA's, the other for MTAs. The OSI stack level for these two kinds of cards is different, but they are both driven by the same central engine.

5. The GIVEME 987 Implementation

The GIVEME design objective is to ensure connectivity between the RFC822 and the X.400 MHS worlds. Thus it must handle two routing schemes and two naming conventions. However, in the RFC822 part of the world there can be many different RFC822-based protocols, i.e. many different software cards. On the X.400 side we can have many cards as well (X.400 implemented on all OSI levels, X.400 on top of TCP/IP or on top of DECnet and so forth).

We have a two-world scenario. GIVEME is thus structured as a two-bus system, the first one with internal RFC822 routing rules and addressing scheme, the second one X.400-based. The two buses must be linked via a fast channel (a mailbox, an ethernet link or at least a fast remote link) to exchange messages. The kernel of the RFC987 conversion rules then sits on top of the conversion channel itself. Both buses can hold as many software cards as needed, and the translation to/from X.400 and RFC822 is performed only if a "change of bus" action is needed.

For the implementation, it was decided to adopt the DEC Message Router bus (also called Mail-Bus) for the X.400 side, ignoring the different configuration approach, and to write our own bus on the RFC822 side.

For the MailBus, we designed and developed a completely new addressing and routing scheme, avoiding the DECnet node and mailbox names as suggested by the standard DEC configuration, basing everything on the standard X.400 attributes. As a consequence, the addressing and routing rules on the MailBus became practically identical, obviating the mapping between MRX and MailBus address format.

On the RFC822 bus, instead, it was decided not to create another specific format for RFC822 routing tables, and the DOMAIN NAMES table syntax was adopted, expanding it in order to support the much complex cases present in GIVEME. New keywords were added, allowing address rewrite rules, the use of different protocols like X.400, SMTP, DECnet, VAX PSI and BSMTP, and a more complex address resolution algorithm, in order to allow the maximum flexibility to the routing configuration. Compatibility with standard DOMAIN NAMES table was also ensured.

A conversion is, of course, needed when a message leaves the MailBus to enter the RFC822 world (or vice versa). The RFC987 specifications fully describe this case; thus we have adopt the mapping procedures and tables according to the RARE MHS project, ensuring the proper connectivity.

Like on the MailBus side, where any available standard module can be plugged in, on the RFC822 bus there are a number of software modules available: VMSmail, BSMTP, RSCS/SNA, PSImail, SMTP over DECnet, SMTP over TCP/IP (UCX, Wollongong, CMU implementations), SMTP over X.25. These modules ensure full connectivity both to end users and to other mailer or gateways, like Columbia Mailer, Unix Sendmail, PMDF and any other gateway implementation using mail-11 foreign protocol.

An interesting feature of the two bus architecture is the possibility to run them separately, on two different CPUs; on the same physical machine. As already said, the only need is to have a fast link between them. Thus the gateway load can be easily shared to increase performance. Moreover it is not needed to have all the mailing protocols implemented awt the same time on the same machine: one machine can implement the OSI networking protocols and the other the Internet, DECnet, etc. ones, increasing even more efficiency and performance. An experiment was performed with the OSI bus (the Mail-Bus half) sitting in Europe and the Internet bus (RFC822) sitting across the Atlantic in the USA , joined by a 64K bps link. The result was connectivity between Internet, HEPnet, SPAN and RARE R&D MHS. More generally, when all the possible cards are plugged into the two buses the GIVEME 987 provides connectivity between most of the currently available "messaging systems", including the traditional ones, like Telex. A typical minimum implementation includes the SMTP on TCP.IP and the MRX cards only, connecting Unix machines on a LAN WITH THE X.400 MHS.

6. RFC822 Compatibility

Users usually do not like to re-format addresses, especially when they are complex as in the RFC987 case: if they read on a business card an E-mail address they want to type it into their user agents "as it is" without translation. Thus we asked DEC to include into the deign of their new X.400 user agent (named MAIL400) the compatibility with RFC822 addresses ,i.e. to enable it to accept them, passing them as they are to the transport system.

Many possible options to handle these RFC822 addresses are now available: forward the message to the closest GIVEME 987 gateway on the MHS; implement a parallel routing in RFC822 style on top of our X.400 routing scheme; or a mixture of these solutions. The most efficient choice is the last one. If the link to the candidate GIVEME implementations fast and reliable then we will forward the message but a local routing for RFC822 addresses on a LAN is highly recommended. No matter which configuration is adopted, the end user simply types the address he knows, X.400 SAA, RFC822, or even a mixed list of addresses and his mail is forwarded towards the final destination.

7. Implementation shortcoming

During the implementation phase it was possible to verify the functionalities of all parts of the system and to compare the MRX.400 MTA with other X.400 MHS products. Thanks to this test phase, the product has proved its functionality in the MTA-to-MTA communication aspects, undergoing a wide interoperability test. Connectivity to all the X.400 MTAs presently used within the RARE MHS Project was fully proven and many problems were solved. Currently the INFN/MTA is the only X.400 entry point for Italy accepted by the RARE MHS project as operational and conformant. Nevertheless we identified other aspects of the MRX functionality that should be improved, concerning "subscriber registration", and "address format".

The present version of MRX requires subscriber registration as a basis for the authorization and routing mechanism. Unfortunately does not use the "X.400 recipient address", contained in the incoming message, as a routing address, but uses an intermediate mapping of subscriber to/from MR-address on the mailbus. With our routing and addressing scheme implemented on the MailBus, this mapping is now superfluous, but can be retained to solve special cases like Telex addresses. Since the research network is the facto an open system security-wise any strict authorization mechanism is not necessary, at least at single-user level.

It has been concluded that subscriber registration is not required, moreover its implementation represents a problem. The problem has been submitted to DEC engineering, and we are now working together to modify the MRX implementation in order to solve it.

8. Conclusion

The standardization of our mail services is aimed at OSI. We based our Message Handling System on DEC products because there is a commitment from DEC to migrate to OSI/MHS with the next release of DECnet phase V or DECnet/OSI. Integrating these products with our own developments has resulted in a powerful and extremely flexible GIVEME 987 mail gateway, connecting in a fully transparent way with nearly any messaging system available in the research world. Our users are now experiencing its services and capabilities, without the hard shock of changing drastically from the RFC822 style to SAA. Meanwhile the low levels of the OSI stack migrate, without impact on users, towards OSI. The last details and corrections of our implementation have been prepared, taking into account one of the most important features needed to obtain an "happy user" in front of his terminal: user feedback.

References

[1]: A. Hansen, RARE MHS Project Review, Comput. Networks ISDN Syst. 16 (1988) 8 - 12.

[2]: R. Grimm and D. Heagerty, Recommendation for a Shorthand X.400 Address Notation, Comput. Networks ISDN Syst. 17 (1989) 263 - 267.

[3]: Alf Hansen and Havard Kvernely, RARE MHS Project Report, RARE, 1989.