

INTERMAIL AND COMMERCIAL MAIL RELAY SERVICES

STATUS OF THIS MEMO

This RFC discusses the history and evolution of the Intermail and Commercial mail systems. The problems encountered in operating a store-and-forward mail relay between commercial systems such as Telemail, MCI Mail and Dialcom are also discussed. This RFC provides information for the Internet community, and does not specify any standard. Distribution of this memo is unlimited.

INTRODUCTION

The evolution of large electronic mail systems testifies to the increasing importance of electronic mail as a means of communication and coordination throughout the scientific research community.

This paper is a summary of the development of, and a status report on, an experiment in protocol interoperation between mail systems of different design. USC/Information Sciences Institute (ISI) began work on this experiment in 1981 and over the years has provided an evolving demonstration service for users to exchange mail between the Internet and a few commercial mail systems.

Recently other organizations have begun to provide similar services, demonstrating the ongoing need for interoperation of the Internet and the commercial mail systems. We believe that ISI's pioneering work in this area has promoted this expansion of service.

These systems include the Internet mail system, the US Sprint Telemail system, the MCI Mail system, and the Dialcom systems. All of the systems were designed to operate autonomously, with no convenient mechanism to allow users of one system to send electronic mail to users on another system.

The Intermail and Commercial Mail Relay (CMR) services described in this paper were developed to provide a means for sending mail between the Internet and these commercial mail systems.

The Internet is an interconnected system of networks using the SMTP mail protocol, which includes the ARPANET, MILNET, NSFNET, and about 700 other networks; mail relays allow the exchange of mail with BITNET, CSNET, and the UUCP networks as well. To the users, this Internet looks like one large mail system with at least 100,000 computers and at least 400,000 users. Figure 1 illustrates the path of a message sent by a user on one Internet host to a user on another Internet host. For more details on the Internet and connected networks (see Appendix A).

As commercial mail systems came into popular use, it became clear that a mail link between the Internet and the commercial mail systems was necessary (see Appendix B). More and more commercial and research entities needed to communicate with the Internet research community, and many of these organizations (for one reason or another) were inappropriate candidates for Internet sites. The Intermail and CMR services allow these groups to communicate with Internet users by purchasing electronic mail services from commercial companies.

INTERMAIL

Intermail is an experimental mail forwarding system that allows users to send electronic mail across mail system boundaries. The use of Intermail is nearly transparent, in that users on each system are able to use their usual mail programs to prepare, send, and receive messages. No modifications to any of the mail programs on any of the systems are required. However, users must put some extra addressing information at the beginning of the body of their messages.




Figure 1 – Internet to Internet Mail

The earliest version of Intermail was developed in 1981, by Jon Postel, Danny Cohen, Lee Richardson, and Joel Goldberg [1]. It ran on the TOPS-20 operating system and was used to forward VLSI chip specifications for the MOSIS project between the ARPANET and the

Telemail system. The original addressing model used in this system was called “Source Route Forwarding”. It was developed to handle situations in which a message might travel multiple hops before reaching its destination.

Later, in 1983, Annette DeSchon converted Intermail into a more general-purpose mail-forwarding system, supporting forwarding between the Internet mail system and three commercial mail systems: Telemail, MCI Mail, and Dialcom [3,4].

As it became apparent that the level of generality of Source Route Forwarding was not needed, and as Intermail gained acceptance among users, an easier approach to addressing was developed. The new addressing model is called “Simple Forwarding”. This form of addressing, like Source Route Forwarding, appears at the beginning of the text of each message. It can be used to include various Internet mail header fields in addition to the standard “To” and “Cc” address fields. This format also allows the use of special address formats, such as U.S. postal addresses and TELEX addresses, which are supported by the MCI Mail system. The Intermail system performed partially automated error handling. Error messages were created by the Intermail program and were then either approved or corrected by a human postmaster.

Figure 2 illustrates the pathways between the user mailboxes in the commercial mail systems and the user mailboxes in the Internet via the Intermail accounts and program modules. Figure 3 shows the Intermail processing in more detail.

COMMERCIAL MAIL RELAY

In 1988, the Commercial Mail Relay (CMR) was developed to run on a dedicated UNIX system, replacing the TOPS-20-based Intermail system. The CMR is a store-and-forward mail link between the Internet and two commercial systems, Telemail and Dialcom. The only remaining forwarding performed by the TOPS-20 Intermail system is in support of the MCI Mail system. (This is planned for conversion to the CMR.) The CMR supports relay-style addressing in the “Internet to commercial system” direction, as well as Simple Forwarding in both directions. One advantage of relay-style addressing is that users from different commercial systems can appear on Internet mailing lists. Another advantage is that the reply features of most Internet user applications can be used by Internet users to respond to mail that originated on a commercial system. Unfortunately, since we do not have access to the address-parsing software on the commercial systems, it is not possible for users of the commercial systems to enter addresses directly into the message header, and they must continue to use Simple Forwarding.

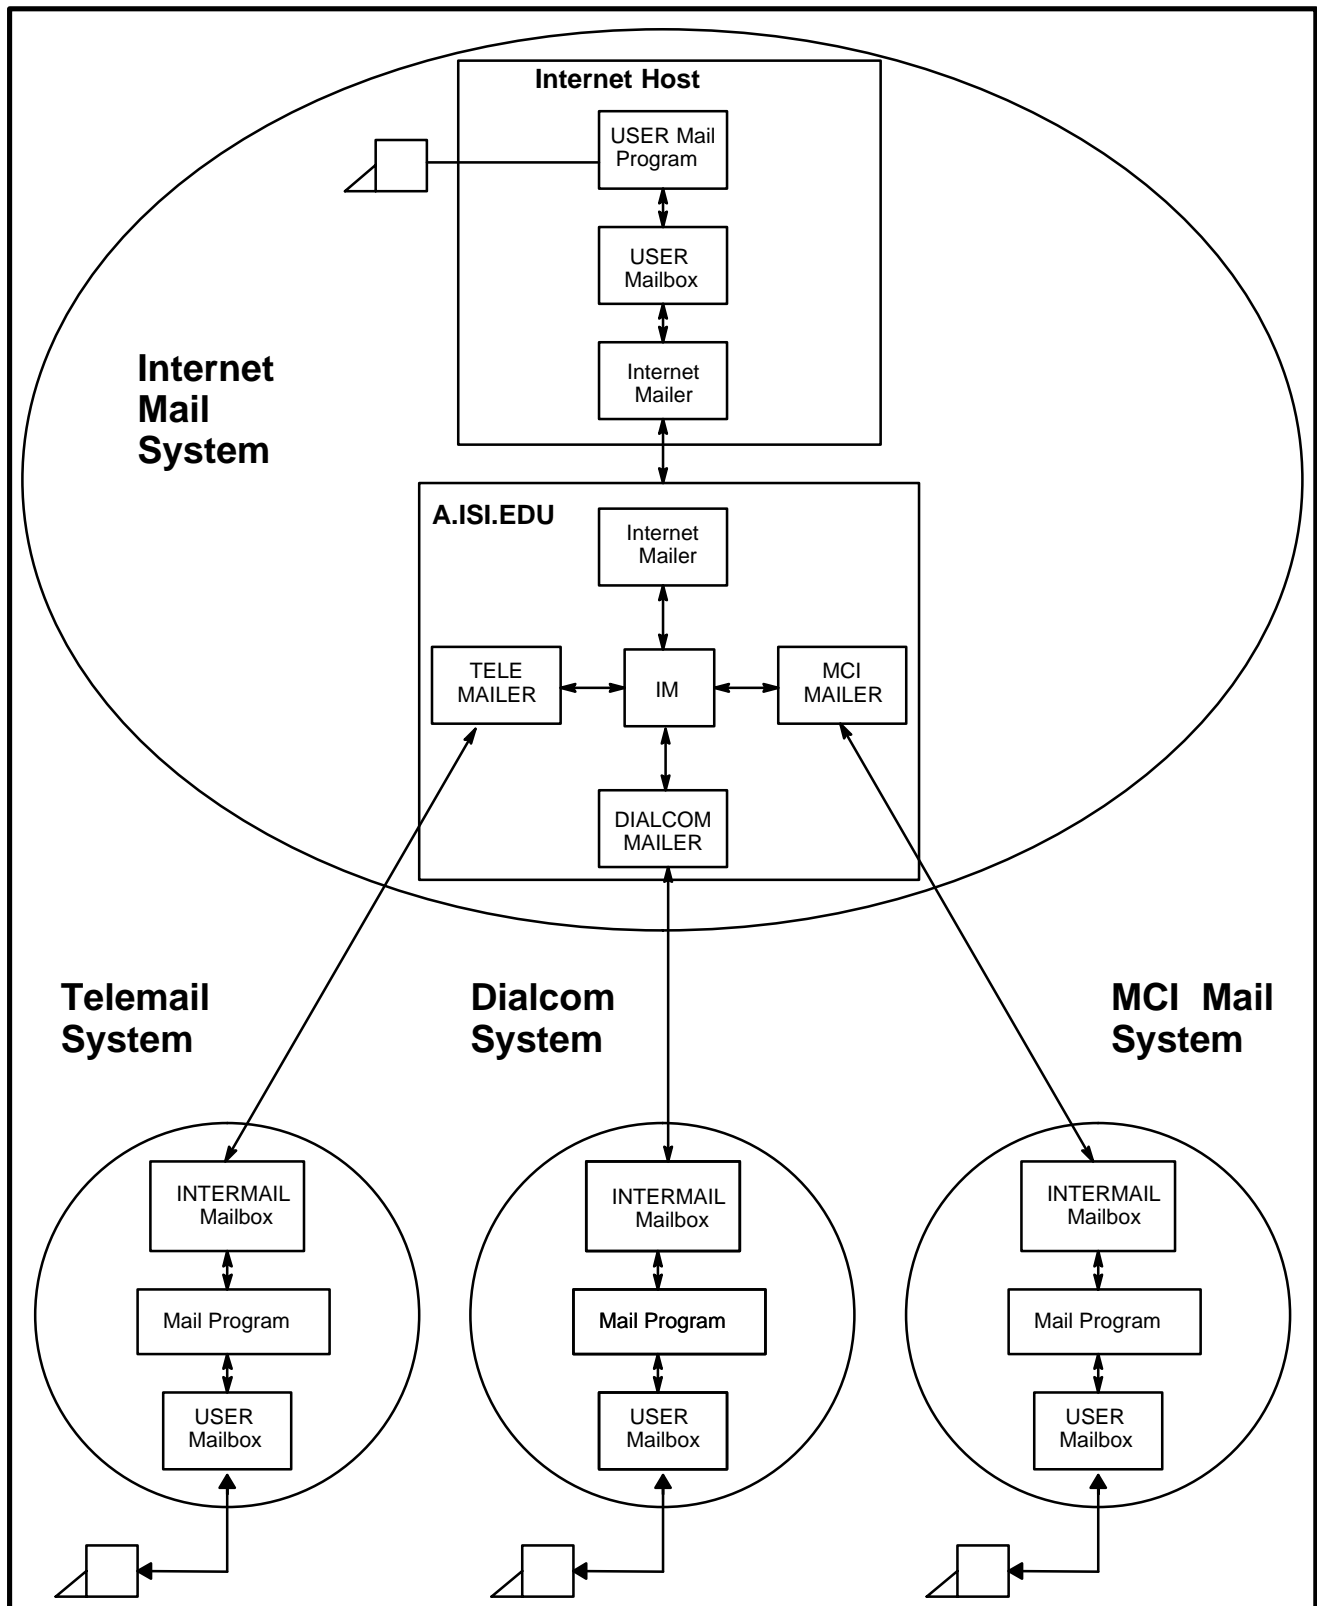


Figure 2 - Commercial Mail to Intermail

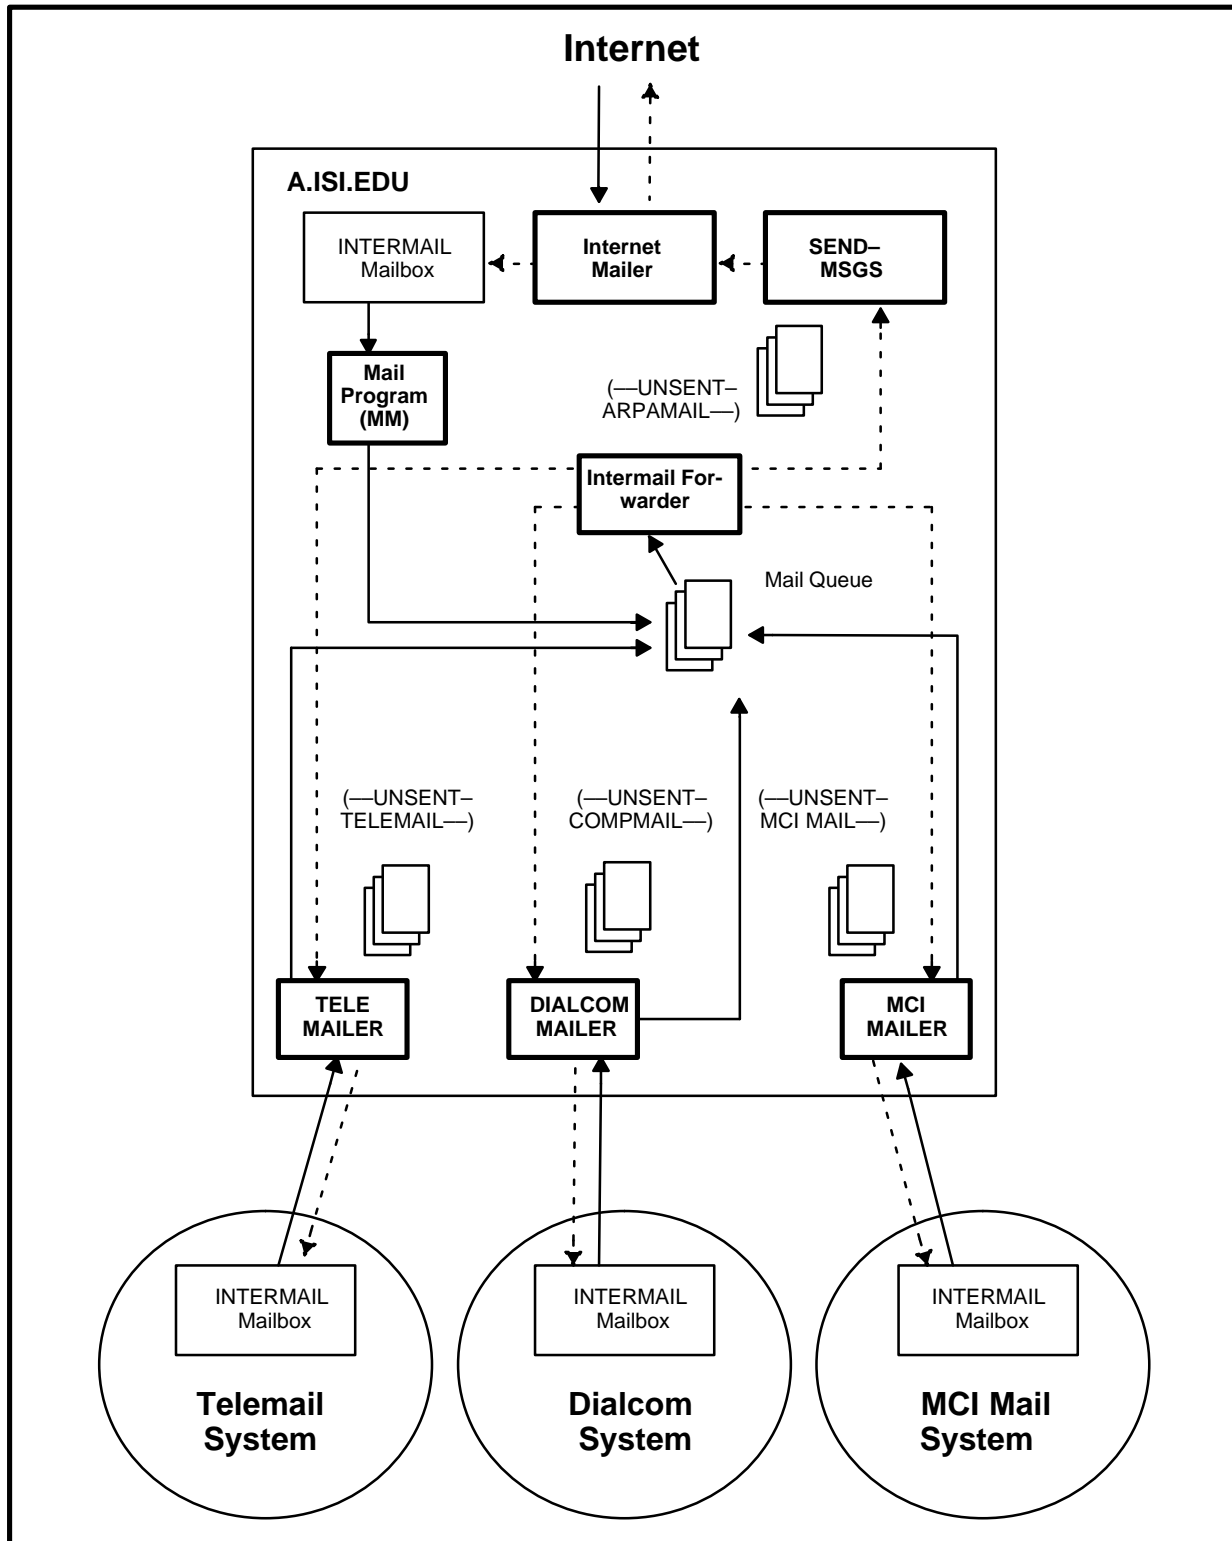


Figure 3 - Intermail Processing

The CMR supports automated error handling, which enables the system to provide faster turnaround on messages containing addressing errors, and requires much less intervention from a human postmaster.

DESCRIPTION OF THE CMR SYSTEM

The Multi-channel Memo Distribution Facility (MMDF) is used as the system mail software because of its notion of separating the mail queue into separate channels [5]. This makes it easy to dedicate a channel/queue combination to each commercial system. Internet mail comes in over the standard SMTP port, and the system parses the destination address, queuing the message in the proper outgoing queue. A tag can be added to outgoing traffic so that replies can be made without any special processing at the destination site.

The CMR uses a relay mailbox on each commercial system. Commercial users send mail to this mailbox with a Simple Forwarding Header (SFH) at the head of their message text. Each channel, in addition to sending outgoing mail into the commercial system, reads all messages in the relay mailbox and places them in a spool file in the CMR host computer.

The processing of the spool file is performed by a single daemon. It parses each commercial system message header to find the sender and subject, then it searches for and processes the SFH. The SFH contains the destination Internet addresses. Figure 4a illustrates the path of mail from the Internet to the commercial systems. Figure 4b illustrates the path from the commercial systems to the Internet. Note: MCI Mail is not yet implemented.

The CMR employs a simple accounting mechanism: a shell script counts the number of times a string marker occurs in the MMDF logs. At the end of the month, another script uses an “awk” program to total the number of messages sent and received with each commercial system. The Commercial Mail Relay is being developed by Craig E. Ward. Ann Westine served as the Postmaster for both Intermail and the CMR until March 1989. Currently, our Action Office serves as Postmaster. Questions may be sent to “Intermail-Request@ISI.EDU”.

COMMERCIAL SYSTEMS SERVED

The CMR provides mail relay service between the Internet and two commercial electronic mail systems: the US Sprint Telemail system and the Dialcom system. A CMR connection to MCI Mail is under development. MCI Mail is currently served by the TOPS-20 Intermail system. See Appendix C for recent traffic data.

The Internet to Commercial Systems:

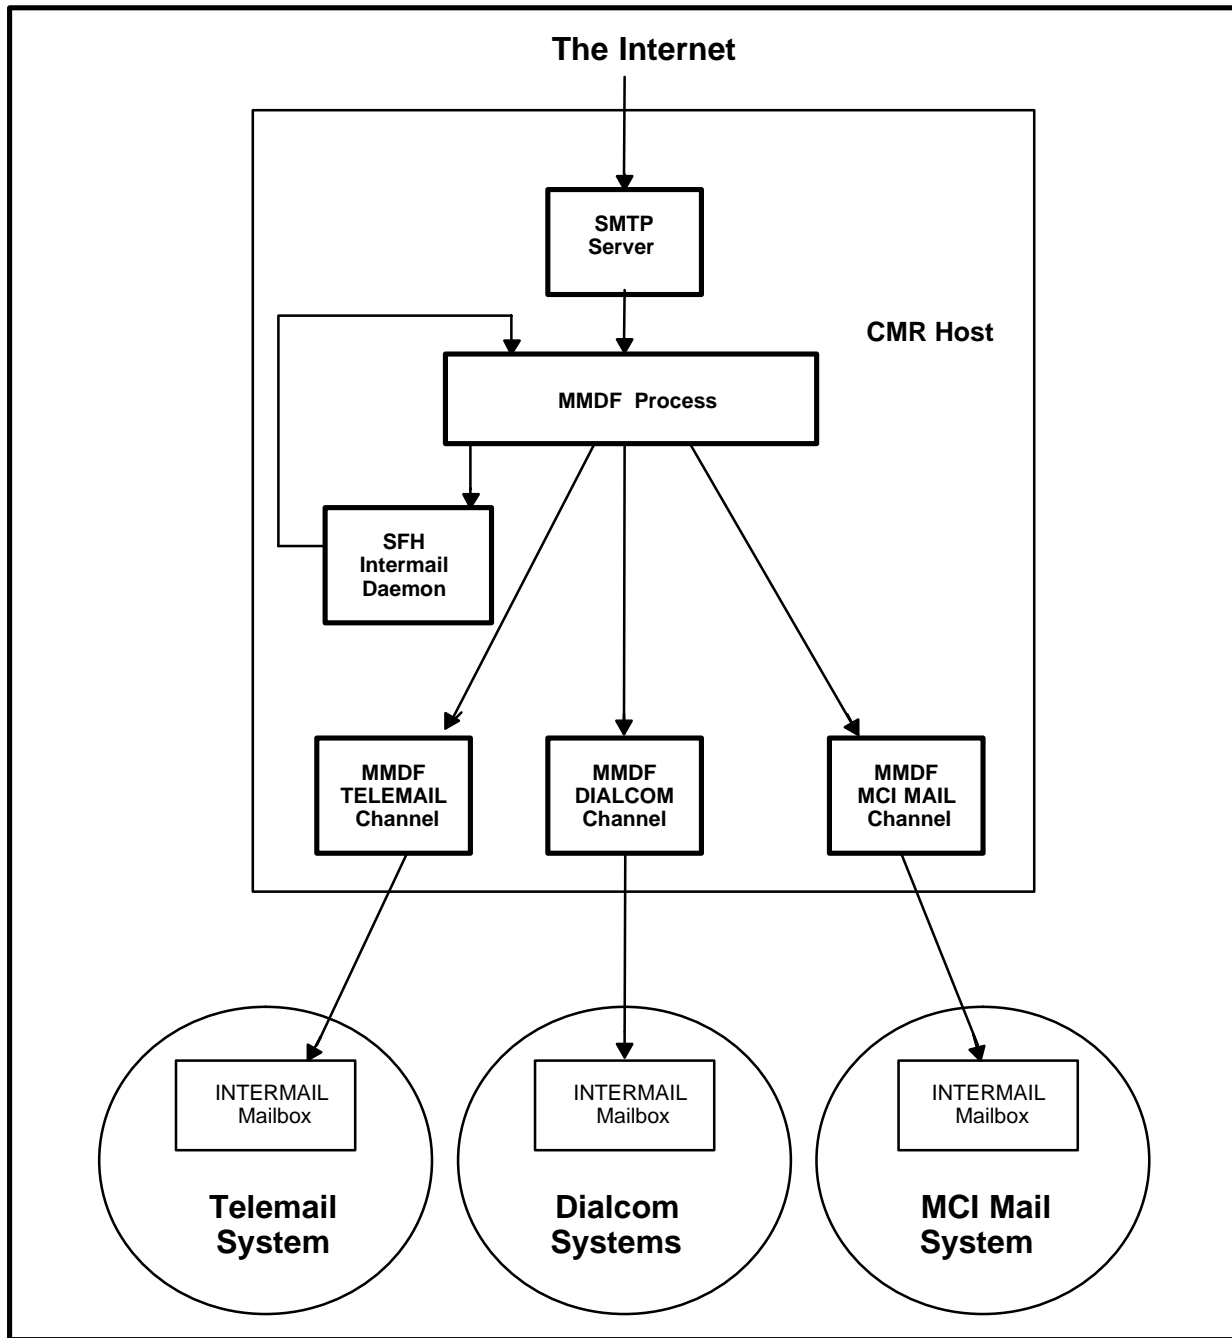


Figure 4a

Commercial Systems to the Internet:

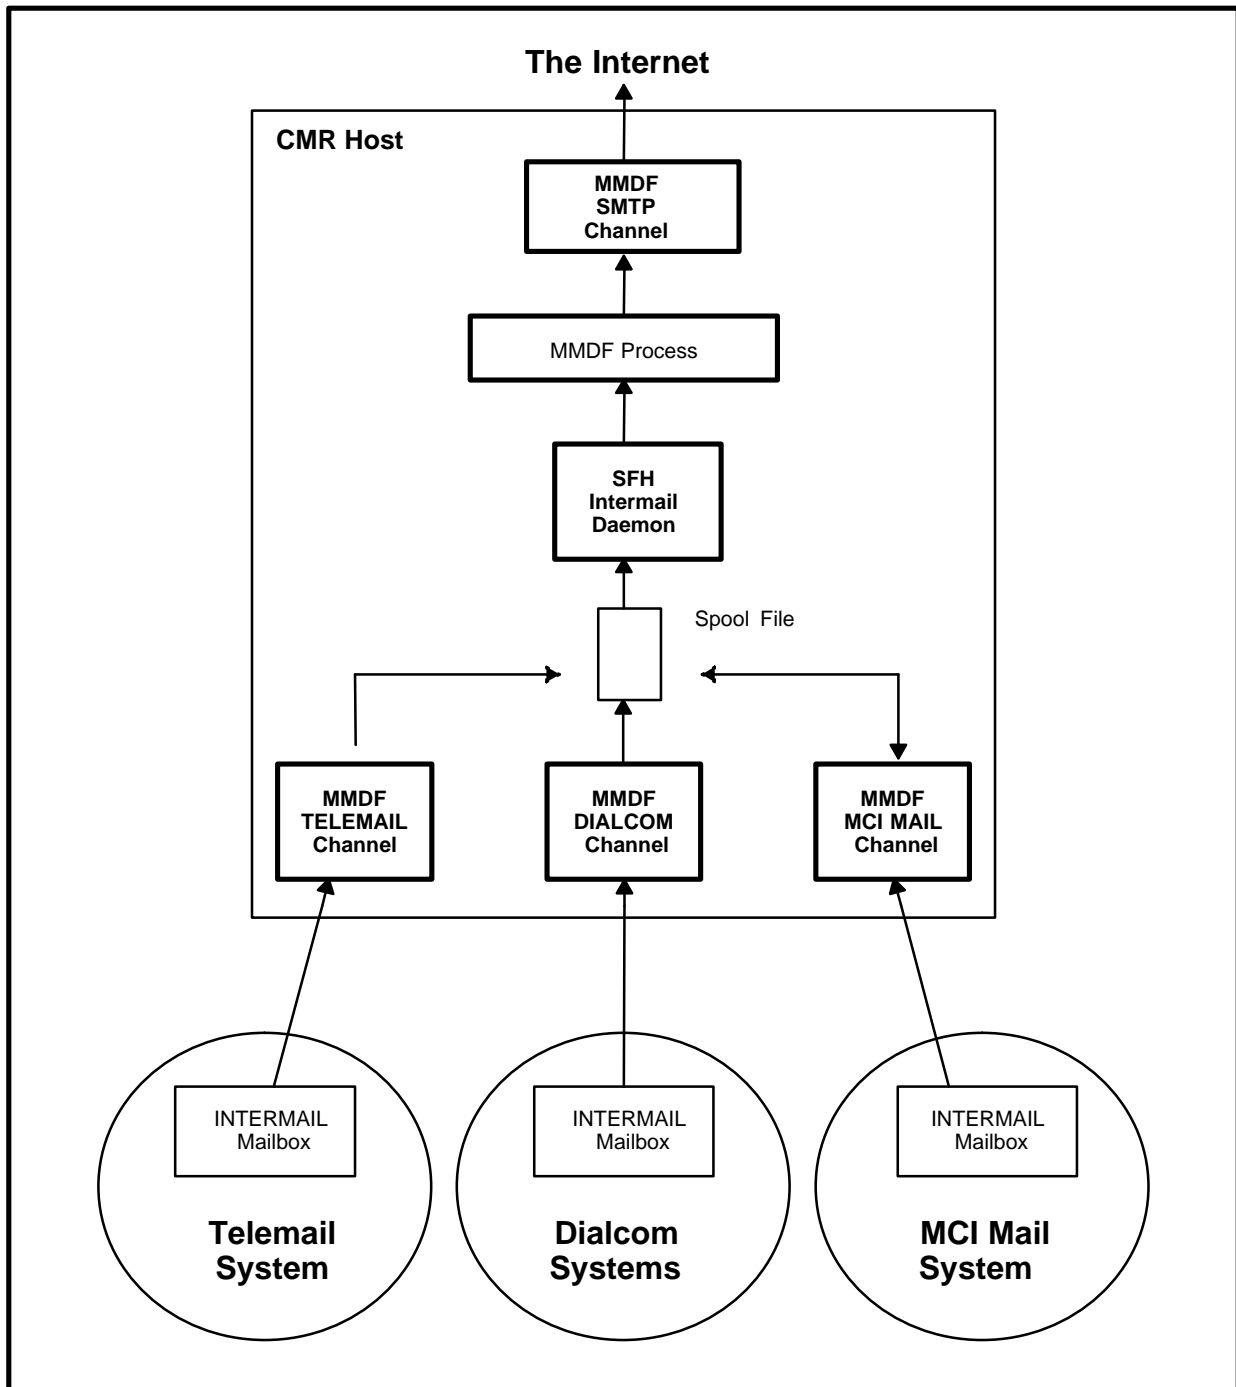


Figure 4b

Telemail is an international commercial service. Some of the Telemail systems served by the CMR include MAIL/USA, NASAMAIL/USA, and GSFC/USA. Some government agencies, such as NASA and the Environmental Protection Agency, have dedicated Telemail systems. Companies also exist that purchase bulk services from Telemail and resell the service to individuals. Omnet's Sciencenet is a very popular example of this type of service.

Dialcom is a commercial service similar to Telemail in that it has facilities for allowing groups to purchase tailored services, and some government agencies (such as the National Science Foundation and the U.S. Department of Agriculture) have special group-access plans. The IEEE Computer Society also has a dedicated group service, called IEEE Compmail, which is open to members of the IEEE Computer Society.

MCI Mail is operated by MCI and is marketed to large companies as well as individual users.

Specific examples of the users of Intermail and the CMR are as follows:

1) Scientists in Oceanography, Astronomy, Geology, and Agriculture use Intermail and the CMR to communicate with colleagues. Many of these scientists have accounts on "Sciencenet", which is actually part of a Telemail system administered by Omnet.

(2) The IEEE Computer Society's publication editors use the Dialcom system "Compmail" to manage the papers being prepared for their numerous publications. Many of the authors are in university departments with access to the Internet. Intermail and the CMR support a significant exchange of large messages containing manuscripts.

(3) NASA uses Telemail systems for its own work and has extensive exchanges through its own relay service, as well as Intermail and the CMR, for communicating with university scientists on the Internet.

Other developments to interoperate between the Internet and Commercial mail systems are:

- (1) The Merit gateway to Sprintmail and IEEE Compmail
- (2) The CNRI gateway to MCI Mail
- (3) The Ohio State University gateway to Compuserve, and,
- (4) NASA-Ames gateway to Telemail

ACCEPTABLE USE POLICY FOR INTERMAIL AND THE CMR

The Internet is composed of many networks sponsored by many organizations. However, all the Internet's long-haul networks are provided by U.S. government agencies. Each of these

agencies limits the use of the facilities it provides in some way. In general, the statement by an agency about how its facilities may be used is called an “Acceptable Use Policy”.

The various agencies involved in the Internet are currently preparing their Acceptable Use Policy statements. Most of these are in draft form and have not been released as official agency statements as yet. None of these policies are currently available as online documents.

In the least restrictive case, all bona fide researchers and scholars, public and private, from the United States and foreign countries (unless denied access by national policy) are allowed access.

The Intermail and Commercial Mail Relay (CMR) systems at ISI are resources provided by the Defense Advanced Research Projects Agency (DARPA) for computing and communication. Use of these resources must be limited to DARPA-sponsored work or other approved government business (or must otherwise meet the acceptable use policy of DARPA), only.

However, DARPA, as a member of the Federal Research Internet Coordinating Committee (FRICC), has agreed to the FRICC draft policy for communication networks, which provides in part that: “The member agencies of the FRICC agree to carry all traffic that meets the Acceptable Use Policy of the originating member agency”.

Thus, e-mail messages (i.e., “traffic”) that meet the Acceptable Use Policy of an agency and pass through some facility of that agency (i.e., “the originating member”) on the way to Intermail or CMR are allowed.

The current member agencies of the FRICC are DARPA, NSF, DOE, NASA, and NIH.

BITNET and UUCP (and other) networks are interconnected to the Internet via mail relays. It is the responsibility of the managers of these mail relays to ensure that the e-mail messages (“traffic”) that enter the Internet via their mail relays meet the Acceptable Use Policy of the member agency providing the Internet access.

In addition, we cannot allow CMR or Intermail to be used simply as a bridge between two commercial systems, even though CMR has this technical capability. At least one end of the communication must be related to FRICC acceptable use.

DETAILS OF CMR SYSTEM USE

The CMR host computer is Internet host INTERMAIL.ISI.EDU (128.9.2.203).

The users of the commercial system are required to know the proper gateways between the Internet and other networks such as BITNET, CSNET, or UUCP. Users on networks interconnected to the Internet likewise need to know how to reach the Internet to send mail through INTERMAIL.ISI.EDU to a commercial system.

The relay connection to Telemail is through their host TELEMAIL/USA. The general syntax for Telemail addresses is “[USER/ORGANIZATION]HOST/COUNTRY”, making the full address for the relay mailbox:

[INTERMAIL/USCISI]TELEMAIL/USA

Users across the entire Telemail service can send mail to this address. Users on the TELEMAIL host need only send to INTERMAIL.

Internet users can use the basic Telemail format, append a “%TELEMAIL” to it, and mail to the resulting address as if it really existed on INTERMAIL.ISI.EDU, e.g.:

[CWARD/USCISI]TELEMAIL/USA%TELEMAIL@INTERMAIL.ISI.EDU

Note that the CMR system will accept anything before the “%TELEMAIL”, that is, the CMR does not validate Telemail addresses before transmitting them to Telemail.

The CMR handles Dialcom mail delivery in a similar way, but this system has what might be called “virtual hosts”. Groups can be set up with an alias system to allow easier intra-group access. For example, both NSF and USDA share the same Dialcom host (157); but, while both groups send relay messages to Intermail, their actual fully qualified Dialcom mailboxes are different. For example, NSF’s mailbox is NSF153, and USDA’s mailbox is AGS9999.

Mail going in either direction may use an embedded Simple Forwarding Header. An SFH *must* be the first part of the message text. It starts with a “Forward:” field followed by a “To:” field. “Cc:”, “Subject:”, and other fields may follow the “To:” fields. The SFH is terminated by a blank line.

This is a template of an SFH:

```
Forward: Destination–Network
To: User@host1, User@host2,
    User2@host2
Cc: User@host1
Subject: This subject supercedes the subject in the host net header
<Blank–Line>
```

Dialcom syntax is “Host–ID:User–ID”, for example, 134:ABC1234. This format will work from any Dialcom host; but users in the same group as ABC would be able to use the user name, for example, JSMITH.

Using the SFH format, mail to a Dialcom system could be sent as follows:

To: Intermail@ISI.EDU
Subject: Test Message

Forward: Compmail
To: 134:ABC1234

Here is the text of the message.

Proper destination network names include ARPA, Telemail, Compmail, NSF-Mail, and USDA-Mail.

It is possible for a user to make mistakes at many points in the process. Errors are handled as automatically as possible by the CMR. Many errors are caught in the standard Internet mail traffic, and users receive the usual error messages from the system. Messages with incorrect commercial system addresses or faulty SFHs are also automatically returned to sender. Messages that the software cannot handle are sent to the CMR's user-service mailbox, Intermail-Request@ISI.EDU. This mailbox has been set up to take care of user problems and to be a central distribution point for user instructions.

PROBLEMS

Several problems arise from the store-and-forward nature of the CMR. One of the biggest is that almost all of the commercial systems lack a machine-to-machine interface — the CMR software must mimic a human user of the commercial system. Another problem is that the Internet and a commercial system have different forms (or syntax) for electronic mail addresses. A major goal of the CMR project is to make the link between networks as transparent as possible, allowing Internet users to use off-the-shelf mail programs. Making commercial address formats fit the Internet standard is a major task [2].

Compatibility with Internet addressing standards is also a concern. The commercial accounts are not able to take advantage of the transparency features of the Domain Name System (DNS) (see Appendix D); and some commercial addresses are incompatible with the Internet syntax—this requires Internet users to continue using the older methods.

Another general problem to be solved is to reduce the amount of time needed to maintain the system. Because most commercial systems force our software to mimic a human user, automatic error detection and handling are quite complex. The Intermail system requires human intervention in processing failed mail. A goal of the CMR is to fully automate these processes.

A related problem facing the CMR, as well as its predecessor Intermail, is the frequency with which commercial systems change their software. The changes are usually minor and do not bother most human users; however, the CMR depends on being able to recognize certain strings. To avoid the necessity of rebuilding the whole CMR when these strings change, most of the string markers are stored in ASCII files that are read at run time.

The translation of commercial system addresses has created a new set of problems, most of which are caused by the use of “special” characters by the commercial systems.

Telemail uses square brackets (“[” and “]”) around user names. While these characters are not special by Internet standards when found in the local part of an address, many (perhaps most) Internet mailers refuse to accept these characters unless they are quoted. MMDF was modified locally to correct this.

The square bracket problem is even worse for users of IBM mainframe machines, many of which are used on BITNET. The square bracket is not a printable character on many BITNET IBM hosts, and all kinds of strange addresses can result from its use.

The colon is another example. Dialcom uses it as the delimiter between host and mailbox. However, the colon is a special character in the Internet mail standard [2]. Users can avoid this problem by using the SFH and placing the Dialcom address at the beginning of the message text. Although the CMR can accept addresses with colons, many Internet hosts and relays are unable to accept addresses that contain colons. Mail with colons in the address fields is often rejected by Internet hosts and is returned to the Intermail–Request mailbox for error processing. This can cause significant delays.

Problems have also been caused by confusion about which hosts are mail relays between the Internet and other systems compatible with the Internet mail standard [2]. (e.g., BITNET, UUCP, and CSNET). When the CMR was implemented, a decision was made that the CMR would not keep track of these mail relays. When a relay is changed, as the BITNET mail relays were in 1988, mail may be rejected because the host either no longer exists or refuses the mail.

The mail relay problem is a subset of the larger problem of communicating information about new features and changes to the user community. Virtually none of the users of the CMR are local. Many are hidden behind the veil of the commercial system. (Dealing with commercial system customer support people has proven to be frustrating — few of them seem to understand the concept of machine-to-machine exchanges.) Enhancements to commercial software that necessitate minor changes can disrupt some CMR users for days.

Another problem that has not been adequately solved is validation of commercial system addresses and processing of failed commercial system mail. The Telemail system will not

validate a user/host combination until after the full text of the message has been transmitted. If a long message is sent to an invalid address, it can be very expensive in terms of wasted time and connect charges.

Telemail also gives inadequate information when the host is correct but the user name is not. The failed mail notice received from Telemail is of little use to either a human reader or the CMR software. The only information that Telemail returns is the message ID number — it provides no subject, and no text to distinguish the message from the numerous others that pass through the mailbox.

Dialcom does a better job of validating addresses. If an address is not recognized, the system immediately prompts for a correction. A simple <RETURN> will delete the invalid address from the list.

The commercial systems are geared for paying customers to send and receive mail to other paying customers. They are not equipped to handle reverse billing, or “collect calls.” ISI is currently charged for connect time needed to transmit and receive mail to and from other Internet sites. A possible solution to this problem would be to extend the CMR. to include accounting and billing procedures that would pass the costs of CMR to its users.

What had been GTE Telemail became Sprint SprintMail, Telenet became Sprintnet, and the host TELEMAIL/USA became SM66/USA.

In April 1990, Sprint installed its X.400 implementation. For the time being, the old-style Interconnect syntax will work. The CMR telemail channel and the Simple Forwarding Header (SFH) processor, were modified to accept either format in the SprintMail “From” field.

Sprint uses the following syntax for X.400:

```
(O:USCISI,UN:INTERMAIL,TS:SM66)
```

The SFH processor will “translate” this into:

```
/O=USCISI/UN=INTERMAIL/TS=SM66/%TELEMAIL
```

The channel program will reverse the process. In the translation, parentheses become slashes, colons become equal signs and commas become slashes and vice versa.

Unfortunately, the translation algorithm is not foolproof. A Sprint/Internet relay did not use the same field names and values as those in SprintMail. Consequently, a CMR translated address can not be sent unmodified to Sprint’s relay, Sprint.COM, and Sprint.COM processed addresses cannot be sent unmodified to the CMR.

From experimentation, the modifications necessary to a CMR processed address to make it acceptable to Sprint.COM are (1) take the “non-standard” X.400 fields of “UN” and “TS” and prepend “DD.” to them, (2) add the country field and code (C:US) and (3) add the Telemail administrative domain name (ADMD:Telemail). The above example would become:

```
/O=USCISI/DD.UN=INTERMAIL/DD.TS=SM66/ADMD=TELEMAIL/C=US/@Sprint.COM
```

The country code must be changed from “US” to “USA.” The CMR queue name must also be appended: “%TELEMAIL.”

The situation is further complicated by Sprint’s decision to only relay mail to and from its own administrative domain. Other X.400 ADMDs may be added in the future if payment problems can be overcome.

SprintMail encoded Internet addresses are not parsed correctly by the SFH processor, but that should not be a major problem — who on the Internet is going to send to the commercial side of the relay?

When the NSF decided to terminate NSFMAIL, it became clear that the CMR Project needed a way to get news out to the commercial users. The CMR channel programs now are able to append a news file to the end of messages going into the commercial networks. After transmitting a message, each channel checks for a news file with the channel name and if present, sends it.

The biggest costs of the CMR are the connect times to the Sprintnet X.25 network and the commercial machines. Making the CMR transmit faster is the current number one problem.

Three strategies are being pursued:

- Improve the implementation of the current method
- Change the method to take advantage of changes in the commercial software
- Upgrade the modems and increase the number of phone lines

For a list of known problems or bugs in the CMR software, see the Appendix of the program logic manual [6].

FUTURE DIRECTIONS

No software project is ever completed, and the CMR is no exception. There are many possible extensions, some more difficult than others.

One addition that will be made to the CMR is a channel for interacting with MCI Mail. MCI Mail is one of the original TOPS-20 commercial systems that were serviced by Intermail; the CMR will need to replace this function before all of the TOPS-20 machines are removed from service on the Internet.

The adaptability of the CMR is such that adding new commercial systems should not be a major problem. Additional commercial systems under consideration include General Electric's GENIE, Western Union's EasyLink, and Compuserve.

One possible addition to the CMR system could be maintenance of a list of gateways. This would allow commercial system users to incorporate the native address formats of other networks into the SFHs. An advantage of this would be that users could simply tell the CMR to forward a message to BITNET, for example, and the CMR would find the gateway and properly format the address for that gateway.

To increase the ease of use to Internet users, the system might treat each commercial system as an Internet host and create DNS database records for them. This would allow users to send mail to a non-Internet user at an Internet-style domain name.

Another improvement would be the possibility of accepting X.400-style addressing. The current system rejects them.

In order to further reduce the hazards of string changes in the commercial system software, an AI component could be added to the commercial system interfaces. Such an AI component might be able to "figure out" what marker a new prompt represents and to remember it.

ACCESS INFORMATION

For instructions on how to use Intermail and CMR contact Intermail-Request@ISI.EDU.

REFERENCES

- [1] Cohen, D., "A Suggestion for Internet Message Forwarding for MOSIS", IEN-180, USC/Information Sciences Institute, March 1981.
- [2] Crocker, D., "Standard for the Format of ARPA Internet Text Messages", RFC-822, University of Delaware, August 1982.
- [3] DeSchon, A. L., "MCI Mail/ARPA Mail Forwarding", USC/Information Sciences Institute, ISI Research Report, RR-84-141, August 1984.
- [4] DeSchon, A. L., "INTERMAIL, An Experimental Mail Forwarding System", USC/Information Sciences Institute, ISI Research Report, RR-85-158, September 1985.
- [5] Kingston, D., "MMDF II: A Technical Review", Usenix Conference, Salt Lake City, August 1984.
- [6] Ward, C. E., "The Commercial Mail Relay Project: Intermail on UNIX", USC/Information Sciences Institute, 1988.

APPENDIX A

The Internet and Connected Networks

The Internet is a network of networks interconnected by gateways or routers. The common element is the TCP/IP protocol suite. The Internet now includes approximately 800 networks and 100,000 host computers. The Internet is made up of local area networks in research institutes and university campuses, regional networks, and long-haul networks. These resources are supported by the using organizations and by several US government agencies (including DARPA, NSF, NASA, DOE, and NIH). The long-haul networks in the Internet are the ARPANET, the MILNET, the NSFNET Backbone, the NASA Science Internet (NSI), and the DOE Energy Science Network (ESNET).

Other systems using TCP/IP or other protocols may be networks of networks or “internets” with a lower case “i”. The capital “I” Internet is the one described above.

There are other networks with (semi-) compatible electronic mail systems. These include BITNET (and EARN and NETNORTH), UUCP (and EUNET), CSNET, ACSNET, and JANET. Users of electronic mail may not necessarily be aware of the boundaries between these systems and the Internet.

The Domain Name System (DNS) is a mechanism used in the Internet for translating names of host computers into addresses. The DNS also allows host computers not directly on the Internet to have registered names in the same style.

BITNET (Because It's Time NETWORK)

BITNET has about 2,500 host computers, primarily at universities, in many countries. It is managed by EDUCOM, which provides administrative support and information services. There are three main constituents of the network: BITNET in the United States and Mexico, NETNORTH in Canada, and EARN in Europe. There are also AsiaNet, in Japan, and connections in South America. Gateways exist between BITNET and the Internet. The most common gateway used is CUNYVM.CUNY.EDU.

CSNET (The Computer + Science Network)

CSNET has 180 member organizations, primarily computer science research laboratories at universities and research institutes, including international affiliates in more than a dozen countries. CSNET has adopted DNS-style names for all its host computers. It is administered by the University Corporation for Atmospheric Research (UCAR) and provides administrative support and information services via the CSNET Information Center (CIC) at Bolt Beranek and Newman (BBN). The gateway between CSNET and the Internet is RELAY.CS.NET.

Note: CSNET and BITNET have officially merged into a single organization as of October 1, 1989.

UUCP (UNIX to UNIX Copy)

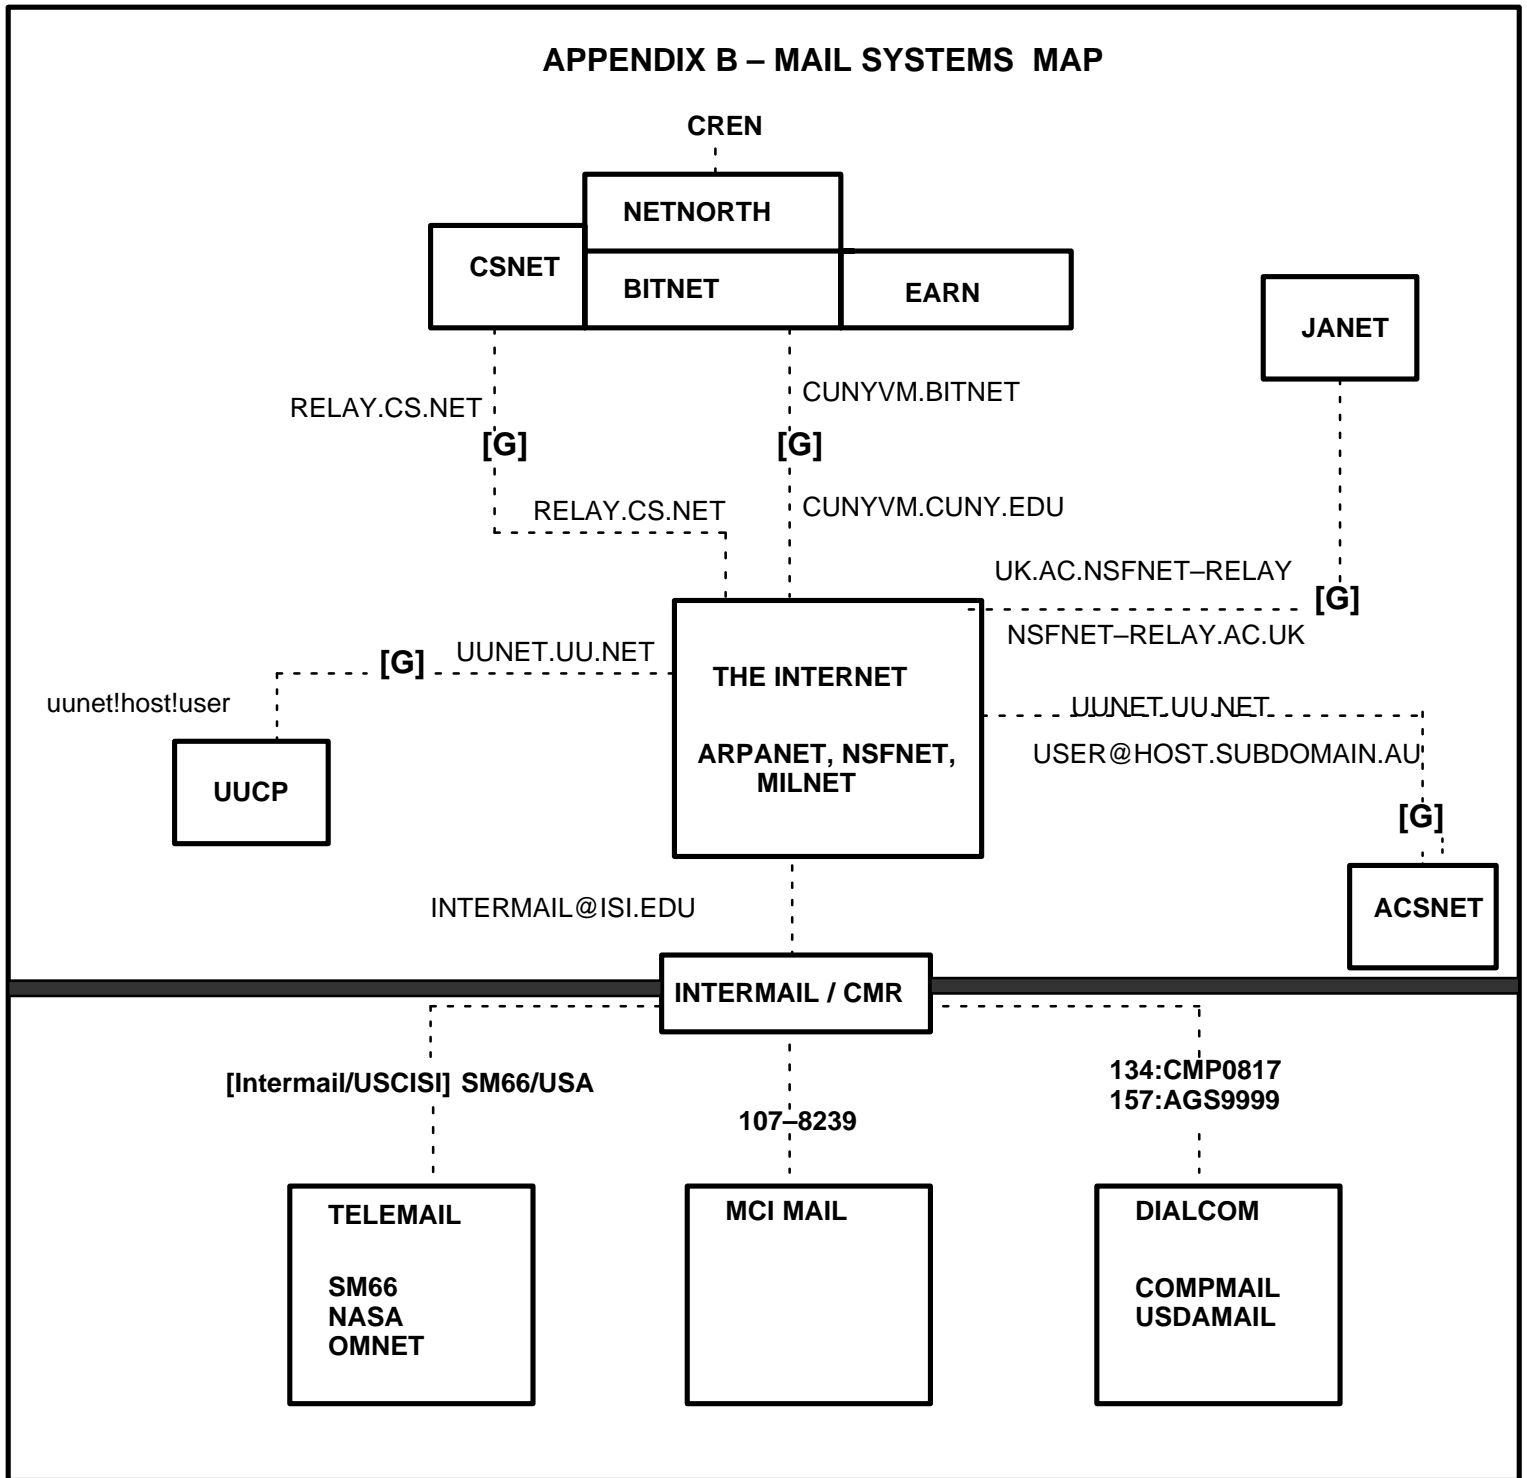
UUCP is a protocol, a set of files, and a set of commands for copying data files from one UNIX machine to another. These procedures are widely used to implement a hop-by-hop electronic mail system. This simple mechanism allows any UNIX host computer to join the system by arranging a connection (dial-up or permanent) with any UNIX host already in the system. In the basic UUCP system, mail is source routed by the sending user through a path of connected hosts to the destination. Currently, there are databases of connection information (UUCP maps) and programs (pathalias) that aid in determining routes. There is some use of DNS-style names by UUCP hosts. EUNET is a UUCP-based network in Europe, and JUNET is a similar net in Japan. These international branches of UUCP use DNS-style names as well. There are many hosts that may relay mail between UUCP and the Internet. One prominent gateway is UUNET.UU.NET.

JANET (Joint Academic NETwork)

JANET is the primary academic network in the United Kingdom, linking about 1,000 computers at about 100 universities and research institutes. JANET has a domain name system similar to that of the Internet, but the order of the domain name parts is opposite (with the top-level domain on the left). The protocols used in JANET are the UK "Coloured Books". The primary gateway between JANET and the Internet is NSFNET-RELAY.AC.UK.

ACSNET (Australian Computer Science Network)

ACSNET is the principal electronic mail system for the computer science and academic research community in Australia, connecting about 300 sites. It works similarly to UUCP. ACSNET has a domain naming syntax similar to that for Internet domains. The gateways between ACSNET and the Internet are MUNNARI.OZ.AU and UUNET.UU.NET.



APPENDIX C

March 1990	sent	read	total
Telemail	1291	768	2059
MCI Mail	56	44	100
Compmail	634	306	940
NSF-Mail	370	243	613
CGnet	171	82	253
USDA Mail	6	1	7
BellSouth	6	10	16
Other	0	0	0

Total	2534	1454	3988
Days in Month	31		
Messages per Day	128.65		

April 1990	sent	read	total
Telemail	1363	696	2059
MCI Mail	40	39	79
Compmail	771	329	1100
CGnet	361	191	552
USDA Mail	28	31	59
BellSouth	98	73	17
Other	0	0	0

Total	2661	1361	4022
Days in Month	30		
Messages per Day	134.07		

May 1990	sent	read	total
Telemail	1007	561	1568
MCI Mail	23	12	35
Compmail	815	359	1174
CGnet	406	210	616
USDA Mail	12	5	17
BellSouth	167	93	260
Other	0	0	0
Total	2430	1240	3670
Days in Month	31		
Messages per Day	118.39		

APPENDIX D

The Domain Name System

The Domain Name System (DNS) provides for the translation between host names and addresses. Within the Internet, this means translating from a name, such as “ABC.ISI.EDU”, to an IP address such as “128.9.0.123”. The DNS is a set of protocols and databases. The protocols define the syntax and semantics for a query language to ask questions about information located by DNS-style names. The databases are distributed and replicated. There is no dependence on a single central server, and each part of the database is provided in at least two servers.

In addition to translating names to addresses for hosts that are in the Internet, the DNS provides for registering DNS-style names for other hosts reachable (via electronic mail) through gateways or mail relays. The records for such name registration point to an Internet host (one with an IP address) that acts as a mail forwarder for the registered host. For example, the Australian host “YARRA.OZ.AU” is registered in the DNS with a pointer to the mail relay “UUNET.UU.NET”. This gives electronic mail users a uniform mail addressing syntax and avoids making them aware of the underlying network boundaries.

SECURITY CONSIDERATIONS

Security issues are not discussed in this memo.

AUTHORS' ADDRESSES**Ann Westine**

USC/Information Sciences Institute
4676 Admiralty Way
Marina del Rey, CA 90292-6695

Phone: (213) 822-1511
EMail: Westine@ISI.EDU

Annette DeSchon

USC/Information Sciences Institute
4676 Admiralty Way
Marina del Rey, CA 90292-6695

Phone: (213) 822-1511
EMail: DeSchon@ISI.EDU

Jon Postel

USC/Information Sciences Institute
4676 Admiralty Way
Marina del Rey, CA 90292-6695

Phone: (213) 822-1511
EMail: Postel@ISI.EDU

Craig E. Ward

USC/Information Sciences Institute
4676 Admiralty Way
Marina del Rey, CA 90292-6695

Phone: (213) 822-1511
EMail: Ward@ISI.EDU

Printed on: Mon Jul 16 15:49:36 1990

Last saved on: Mon Jul 16 15:51:09 1990

Document: cmr90-ps

For: westine