


AARNet



Engineering
a
Connection
to
AARNet

Geoff Huston
Network Technical Manager
AARNet

G.Huston@aarnet.edu.au

GPO Box 1142
Canberra ACT 2601

Tel: (06) 249 3385
Fax: (06) 249 1369

Connecting To AARNet

Table of Contents

1. Background	1
1.1 What is AARNet?.....	1
1.2 Who owns, operates and funds AARNet ?.....	2
1.3. What organisations are members of AARNet and who can join ?.....	2
1.4. What services are supported by AARNet ?	2
2. Affiliate Membership.....	3
2.1. Acceptable Usage.....	3
2.2. Fee Structure.....	4
2.3. Mail Affiliate Membership	4
3. Network Affiliate Membership.....	5
3.1. The Physical Link.....	5
3.1.1. Router - Router Links	6
3.1.2. SLIP/PPP Links	6
3.1.3. Other Connections.....	7
3.2. Network Addressing and Routing	7
3.2.1. An Introduction to IP Addressing.....	7
3.2.2. Obtaining an IP Network Number.....	8
3.2.3. Routing.....	8
3.3. Domain Name System.....	9
3.3.1. Why ?	9
3.4.2. Structure.....	9
3.4.3. Domain Name Servers	10
3.4.4. Setting up a Domain Name Server.....	11
3.4.5. Making applications use the DNS.....	15
3.5. Electronic Mail.....	16
3.6. Security	17
3.6.1. Welcome to the Internet.....	17
3.6.2. Security Measures.....	17
3.6.3. Known UNIX System Configuration Problems.....	18
3.7. Network News.....	20
Appendix A. AARNet Members	21
Appendix B. AARNet User and Resource Guides.....	27
Appendix C. Mail Gateway Services and Providers	28
Appendix D. AARNet Regional Hub Contacts	31

Attachments

1. AARNet Mail Affiliate Membership Application
2. AARNet Network Affiliate Membership Application
3. AUUG Mail Service Application
4. Internet Number Application Form
5. *Introduction to Administration of Internet-based Local Networks*, Charles Hedrick, Rutgers University

CONNECTING TO AARNET

1. Background

This tutorial describes how to establish a connection to the Australian Academic and Research Network (AARNet).

The intended audience for this document are technical staff who are about to undertake, or are evaluating, a connection to AARNet. In particular emphasis is placed on establishing a network level Internet (TCP/IP) link, rather than an electronic mail link.

1.1 What is AARNet?

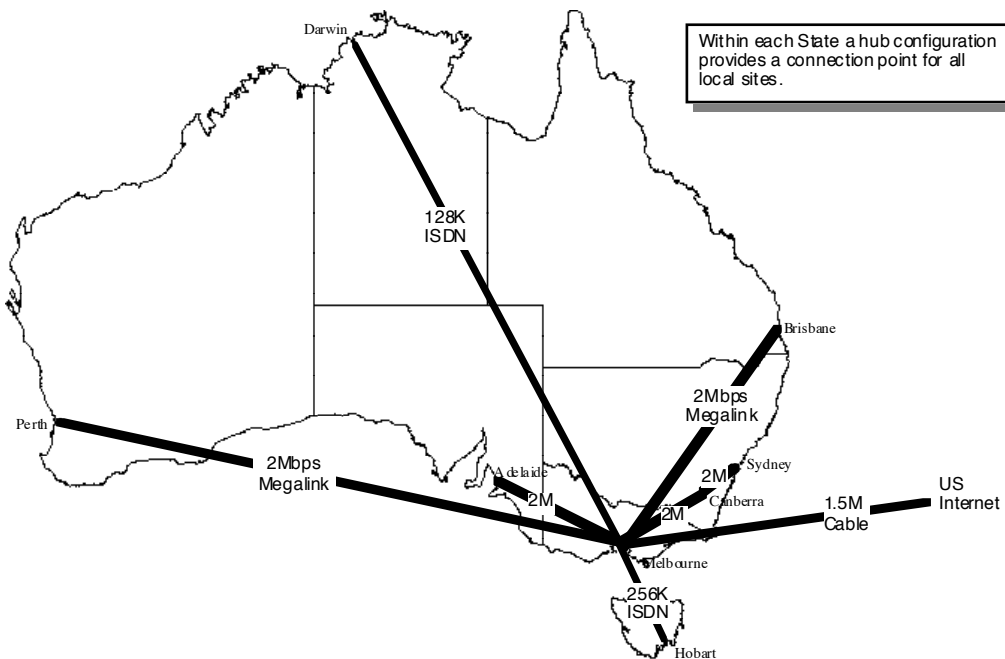
AARNet is a private data network that provides dedicated telecommunications services in support of its members' research, academic, service and operational activities. The network links all Australian Universities, CSIRO Divisions and a large number of other organisations under the Affiliate Membership program described below.

AARNet was commissioned in May 1990, and over three years of operation it has experienced an approximate twentyfold increase in usage (as measured by traffic volume), and in October 1993 there were an estimated 85,000 connected computers connected to AARNet, and more than 550 organisations receiving communications services from the network.

AARNet provides interconnections between the local area networks (LANs) of its client members and other networks within Australia and overseas. AARNet itself does not provide services to end-users; it enables users of systems that are connected to AARNet to access resources and services made available by other entities connected to the network. Thus, there is no AARNet "userid" and "password" - access to AARNet is obtained by providing some form of communications link between the computing facilities as a member site and either a mail gateway system (for Mail Affiliate membership) or an AARNet network gateway (for Network Affiliate membership).

AARNet is a member network of the Internet, an interconnection of many networks that use the Internet Protocol (IP) to provide direct access to tens of thousands of organisations throughout the world. The Internet is the world's largest electronic network, with some 1,500,000 connected computer systems and an estimated 20 million individual

**The Australian Academic and Research Network
Trunk Configuration - September 1993**



CONNECTING TO AARNET

users. Any system connected to the Internet at the network level (eg. through AARNet Affiliate Membership) can both access, and be access by, any other system that is also connected to the Internet (assuming that they have been granted the authority to do so). The Internet also interconnects with many other networks, making it possible to exchange electronic mail with users of networks such as:

BITNET	BTGold	Dialog	UUCP	USENET
Keylink	Dialcom	CompuServe	AT&T Mail	MCI Mail

AARNet is architected around medium (192Kbps and 256Kbps) and high speed (2M and 10Mbps) leased lines that link multiprotocol routers in a hierarchical star structure centred on Melbourne. AARNet operates a regional hub located within a major university in every capital city of Australia.

AARNet operates a regional access hub in every capital city of Australia, located within a major metropolitan university.

1.2 Who owns, operates and funds AARNet ?

AARNet is owned and operated by the Australian Vice-Chancellors' Committee (AVCC) as a service to the participating member institutions of the AVCC and the Commonwealth Scientific and Industrial Research Organisation (CSIRO). Operational staff, based in Canberra, are responsible for the engineering and management of the network and the regional hub sites are contracted to provide additional regional operational support.

The network is funded through contributions from the AVCC members and CSIRO and fees from Affiliate Members.

1.3. What organisations are members of AARNet and who can join ?

As noted above, all Australian Higher Education Institutions (HEIs) are members of AARNet and are connected to the AARNet backbone infrastructure at speeds of between 48Kbps and 10Mbps (microwave). Most CSIRO divisions also have connections to AARNet, which depending on the division's size and communications needs, are operated at speeds of between 9.6Kbps and 10Mbps.

Under it's general brief, AARNet also provides telecommunications services to other organisations. In this and other AARNet documents the member institutions of the funding bodies (AVCC and CSIRO) are termed *Members* of AARNet. Other organisations and bodies connecting to AARNet are termed *Affiliate Members* of AARNet.

In providing AARNet access to Affiliate Members the objectives of AARNet are to:

- further the opportunities available to Higher Education Institutions (HEIs) and CSIRO through assisting collaborative activities via a common network infrastructure;
- enable vendors to improve the quality and cost-effectiveness of their support to the HEIs and CSIRO;
- facilitate the dissemination of science, technology and humanities information through Australia in the direct support of the role of the HEIs and CSIRO within this country;
- improve access to, and the quality of, information services of interest to the HEIs and CSIRO; and
- encourage broad Australian participation in the development of a common network infrastructure realising benefits to the national academic and research endeavour through economies of scale.

A complete list of all AARNet members as at July 1993 is attached to this document.

1.4. What services are supported by AARNet ?

As noted above, AARNet is simply a network of networks. The services supported by AARNet are therefore simply those supported by the computing systems connected to AARNet's member networks. Details of these services can be found in a variety of on-line directories, manuals and user guides, but a brief summary of the basic functions supported by AARNet (and hence the Internet) are presented below.

CONNECTING TO AARNET

- **Electronic Mail** The ability to transfer electronic messages is the most useful and pervasive facility provided by AARNet. As mentioned above mail-gateways allow electronic mail to be exchanged between AARNet connected organisations and literally millions of other users throughout the world.
- **Remote Interactive Access** This service allows users with access to an AARNet-connected PC, Macintosh or mainframe system to establish interactive terminal sessions with remote hosts. This facility is supported by the IP application *telnet* and is used for a wide variety of purposes, including remote access to library catalogues, databases, directories, information servers and supercomputing systems.
- **File Transfer** File transfer facilities enable the sharing of resources and information between users, and also provide a mechanism for the rapid distribution of software, applications and data. The basic transfer mechanism is based on the *File Transfer Protocol (FTP)* that supports information exchange on a client-server basis between dissimilar computer hardware and software platforms.
- **News** As well as mail transfer, the other major messaging mechanism is the exchange of public notes between users on a network of host systems, creating a global network-wide bulletin board. The USENET news network currently delivers 3,000 new messages per day to a global readership estimated to be of the order of 2,500,000 readers. NEWS is used to distribute software and software updates, provide technical assistance on a peer basis on a wide variety of subjects, and to allow the interchange of ideas on a wide range of subjects of particular relevance to the research and academic community.
- **Others** There are many other network applications not mentioned above, including a whole range of distributed applications to support X.500-based directories, wide-area file systems and information servers. Recent advances in the development and deployment of information services that incorporate intuitive, graphical interfaces to make searching multiple databases for keywords, documents or data are making it possible to harness the enormous volume of widely distributed information stored throughout the Internet. This "new wave" of applications include the Wide Area Information Servers (WAIS), World Wide Web (WWW), Gopher and Archie.

2. Affiliate Membership

There are two basic connection services offered by AARNet to Affiliate Members;

- **Mail Affiliate Membership**, which enables an affiliate organisation to exchange electronic mail with AARNet. Other services that may be layered above electronic mail transfer or similar store-forward messaging-based connections, such as network NEWS, can also be used.
- **Network Affiliate Membership**, which provides full TCP/IP network-layer connectivity to AARNet and the peer networks that make up the Internet. The full range of services supportable by TCP/IP are available to network affiliate members, including remote login, file transfer, electronic mail and access to on-line information servers just to name a few.

For users that wish to use electronic mail only occasionally, or that have little existing computing infrastructure or funding to support a full scale AARNet mail or network affiliate connection, it is possible to send electronic mail to/from AARNet, and connected networks, through Telecom Australia's Keylink electronic messaging service or via the commercial messaging service provider, Pegasus Networks. These organisations can be contacted as follows:

Keylink	Pegasus
Telecom Plus	Pegasus Networks
Tel: 008 032 333	Tel: 07 257 1111

The services offered by these organisations are dial-up computing services, where the service provider maintains a mailbox, and an interactive computing environment on the service host, supported by remote terminal access via dial-up.

The applications forms for Affiliate Membership that detail the procedures, terms and conditions for connections to AARNet are included with these tutorial notes.

2.1. Acceptable Usage

Both Mail and Network Affiliate members are required to comply with the AARNet acceptable use policy included in the AARNet affiliate membership applications. In addition, network affiliate members must comply with a number of conditions, described in the Network Affiliate membership document.

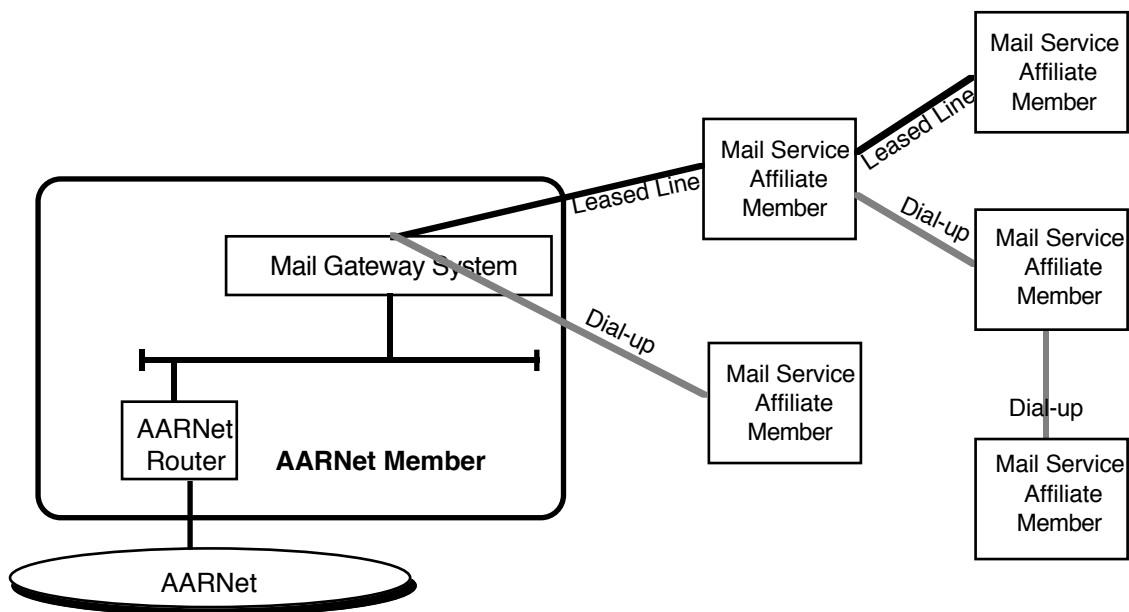
CONNECTING TO AARNET

2.2. Fee Structure

It is not the intent of AARNet to provide telecommunications services to Affiliate Members at a level of charges where AARNet is subsidising such services, nor is it intended that the total costs of the network project be funded through the Affiliate Membership program. The underlying principle of the charge schedules relating to Affiliate Membership of AARNet is to reflect the actual costs associated with the service as far as is feasible. In so doing, the charging model adopted by AARNet is that of fixed annual charges to Affiliate Members of AARNet - there are no usage based charges. Details of these charges are described in the AARNet membership documents.

2.3. Mail Affiliate Membership

Mail Affiliate membership allows organisations to exchange electronic mail over AARNet. This is done by firstly establishing a link (usually via periodic dial-up) between the mail affiliate organisation and a *mail gateway* system either directly, or indirectly via another mail (or network) affiliate member, as shown below. The mail gateway stores electronic mail destined for the affiliate member, and when a link is established passes it on towards its destination. At the same time, any outgoing mail from the affiliate members is passed to the mail gateway for delivery. This is commonly referred to as a *store and forward* mail system.



Electronic Mail Connections to AARNet

The protocol used to operate this link is at the discretion of the provider of the mail gateway service and the client affiliate mail affiliate member, but is typically either MHSnet (the commercial implementation of the SUN III software that supports the Australian Computer Science Network - ACSnet) or the Unix to Unix Copy Program (UUCP). An in-depth discussion of the configuration of either is beyond the scope of this tutorial, but it should be noted that MHSnet is available for UNIX platforms only, whereas there are implementations of UUCP available for a wide range of systems including public domain versions for DOS-based PC's, Macintoshes, and, VAX/VMS systems.

The second step is the installation of a mail exchange (MX) record within the Internet domain name system. The MX record registers the mail affiliate organisation's electronic mail address along with the the name of the appropriate gateway to which Internet electronic mail should be sent in order to reach its ultimate destination. The domain name system is a distributed database that is consulted by Internet mailing systems during the process of mail transmission - absence of an MX record implies that the mail address is invalid and transmission attempts should not proceed. Because the domain name system is global, this messaging service extends to all peer international academic and research networks as described above.

AARNet uses the installation of an MX record as a mechanism to provide mail affiliate membership, for which there is an annual charge of \$1,000 (no volume charges apply). Note that this charge does not include network news, which is not charged for by AARNet. This also does not cover the costs associated with the provision of the mail

CONNECTING TO AARNET

gateway service, which is available from any of the mail gateway providers listed in Appendix C; typical gateway service charges vary between \$100-\$2,000 (which may include a full or partial network news feed).

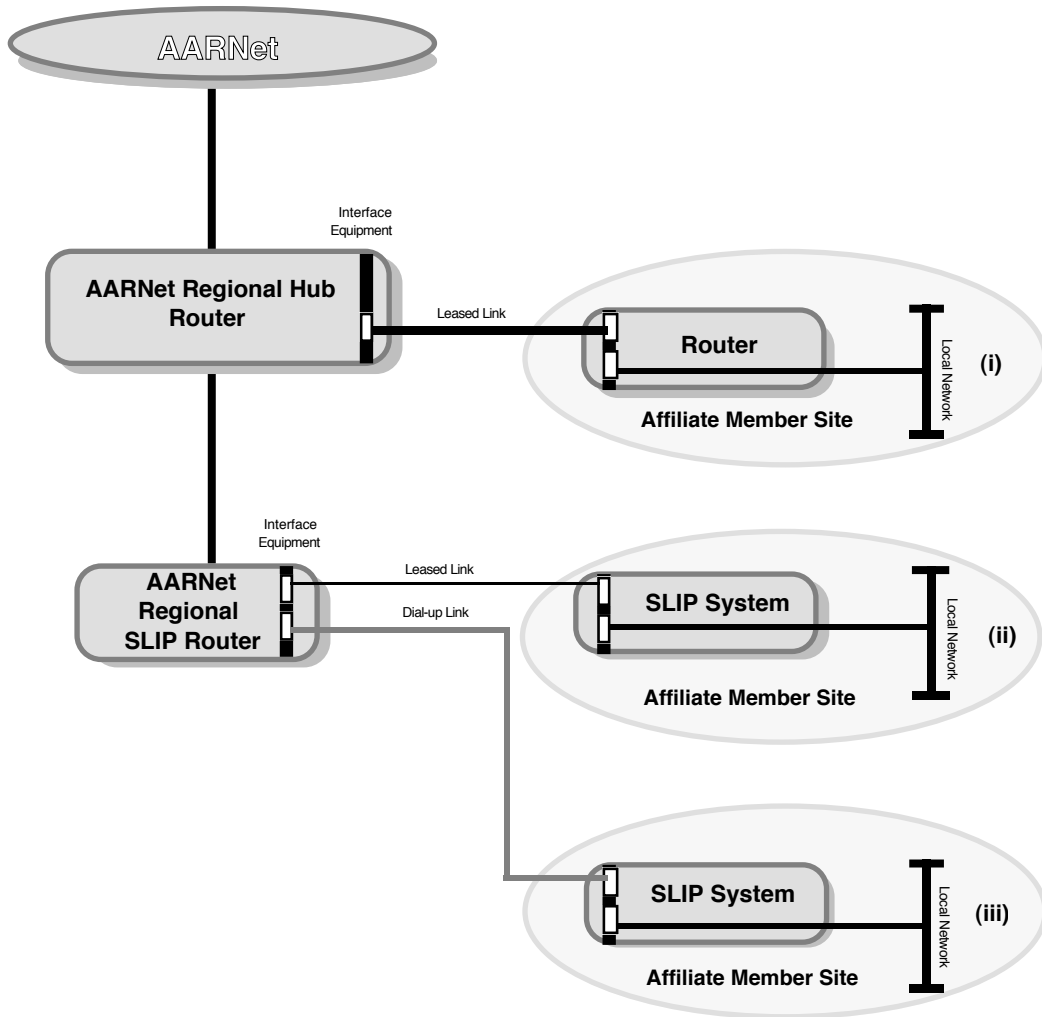
In addition to the \$1,000 per annum mail affiliate program, AARNet has negotiated an arrangement with the Australian UNIX Users Group (AUUG) to allow for a transition and formalisation of the ACSnet network for sites that only exchanged small volumes of electronic mail. This program allows members of AUUG to obtain AARNet Mail affiliate membership, through the AUUG, for \$250 per year (see attached application).

3. Network Affiliate Membership

AARNet Network Affiliate Membership involves an organisation establishing a link to AARNet using the Internet Protocol (IP), often referred to as TCP/IP (Transmission Control Protocol/Internet Protocol). In so doing the connecting organisation enters AARNet as a peer entity and has access to any and all services that can be provided over the TCP/IP connection. Because AARNet is a member of the Internet, network affiliate organisations also become Internet members.

3.1. The Physical Link

There are a number of means of interconnection to AARNet as shown in the following figure:



Preferred Connection Methods to AARNet

CONNECTING TO AARNET

The preferred method of connection to AARNet is via a leased line direct to one of the AARNet Regional Hubs located in each state/territory capital city.

Where a link to the Regional Hub is impractical (far North Queensland, for example), a link to a local AARNet AVCC member or CSIRO site may be considered by AARNet. However as this implies that the Affiliate Member is sharing common capacity with the host AARNet Member to connect to the Regional Hub, this method of link may not be possible in all cases. Again in such a case a router is required on the Affiliate Member site, or SLIP servers may be a practical alternative. This method of interconnection requires the specific approval of AARNet as well as the member institution involved in hosting the connection.

Where the connection is made to AARNet equipment there is a once-off connection fee as described in the AARNet Affiliate Membership document. In addition, a annual membership fee is payable, based on the effective bandwidth of the connection to AARNet. No other charges are levied by the AVCC for the provision of network affiliate membership. Where the connection is made to equipment owned and operated by a university or CSIRO division, that university or CSIRO division may charge for operational support services with respect to their end of the connection.

Note that network affiliate member is responsible for the ordering, installation, operation and maintenance of any equipment at their site, and the Telecom services that connect it to the AARNet Regional Hub. AARNet purchases, maintains and operates the hub itself, including the interface that terminates the link to the network affiliate member.

3.1.1. Router - Router Links

The "standard" form of AARNet connection is a leased line between the AARNet Regional Hub and the affiliate member's site. The leased line service may be any form of synchronous digital transmission service, including the Telecom DDS, DMS, Datel, or ISDN services. Terminating the leased line is a gateway server, or router, which must be interoperable with the AARNet Cisco Systems router used at all Regional Hubs. Alternatives to a cisco router that AARNet has confirmed are acceptable include Wellfleet, NAT B/280 and Proteon routers. This method of interconnection is shown as the first connection in the preceding figure (i).

The benefits associated with using a router-router connection, rather than a SLIP connection (described below) include:

- better utilisation of link bandwidth;
- no host processing overhead;
- simpler and more reliable;
- efficient, manageable network wide access to systems connected to the the affiliate site's LAN;
- incorporation into an affiliate member's WAN strategy;
- network management controls through both standard (SNMP) and vendor-specific mechanisms; and
- the ability to enforce strict security controls with a fine degree of granularity.

The major drawback with using routers at both ends is the capital cost of the connection. A low end router (such as a NAT B/280) retails for some \$5,000, but a full function router (such as a cisco, Wellfleet or Proteon) may cost \$6,000-\$10,000. The AARNet fee for connection to the AARNet Regional Hub router is \$3,500.

3.1.2. SLIP/PPP Links

To support low speed, low cost asynchronous links, connections to an AARNet Regional Hub SLIP server are available. Such connections normally take the form of a 9.6Kbps leased line (shown as (ii) in the preceding figure), or dial-up service (iii), using the Serial Line IP Protocol (SLIP) or Point-To-Point Protocol (PPP). In this case a router is not required at the affiliate member's site but either a system which supports SLIP/PPP, or a terminal server with SLIP/PPP is required.

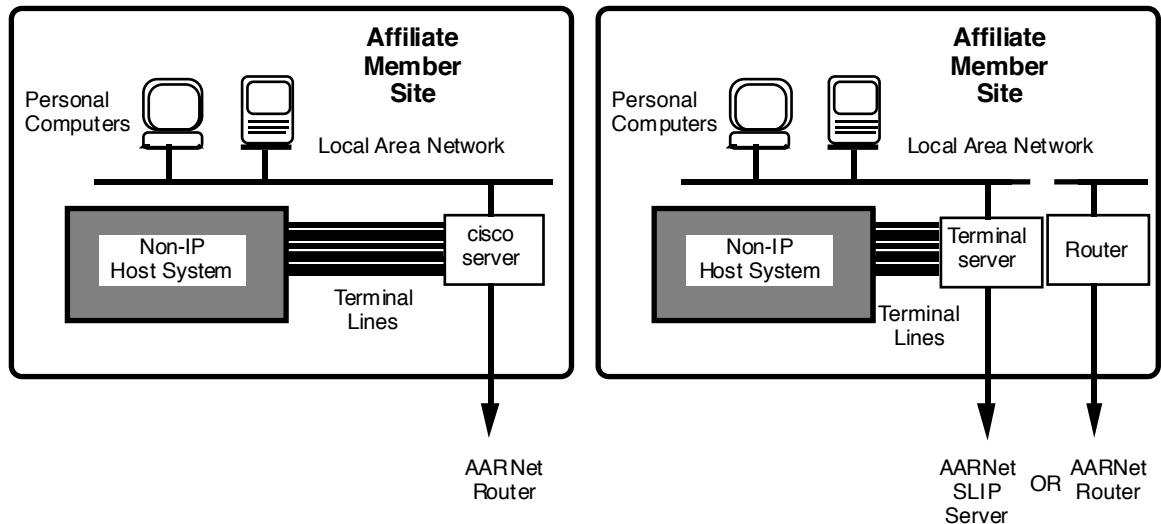
SLIP is available for most systems that support the Internet Protocol suite, and for UNIX, PC and Macintosh systems there are implementations in the public domain. SLIP connections are appropriate when the number of systems to be connected is small, the volume and quality of service required by the affiliate member are not high, and/or the budget is severely constrained.

In situations where the affiliate member wishes to provide modems for connection to a Telecom service, the modems provided must be acceptable to the AARNet Regional Hub. Note that the dial-up access is to a dedicated port on the AARNet SLIP server, and thus it is still necessary for the affiliate member to purchase a Telecom service, even if it is simply a standard phone line. The AARNet fee for connection to an AARNet Regional SLIP/PPP server is \$1,000.

CONNECTING TO AARNET

3.1.3. Other Connections

In some circumstances, neither of the above techniques will service an organisation's requirements for connectivity to AARNet. The most common of these is when access needs to be provided to a host that does not support IP natively. In such cases hybrid solutions, terminal servers can be used to "front-end" a non-IP host. The following indicates potential configurations for such a case.



The diagram on the left in the preceding figure shows a cisco combination terminal server and router, whereas the example on the right is a more standard arrangement that can support either a SLIP connection, or a router-router connection. In either case the personal computers are not required to support the AARNet connection to the non-IP hosts, but allow outgoing access to AARNet from the affiliate member.

If such a link is considered appropriate, the proposed configuration should be discussed with AARNet.

3.2. Network Addressing and Routing

In this section some of the basic issues of IP network addressing and routing are discussed. The bulk of this section is attached as a separate paper written by Charles Hedrick of Rutgers University, New Jersey and made available to the public over the Internet.

3.2.1. An Introduction to IP Addressing

IP network addresses are 32 bit binary numbers that uniquely identify a connection point to an IP network. A system with a serial interface and an ethernet interface, or two ethernet interfaces, will therefore have two IP addresses. All IP packets contain a source and destination IP address; the former indicates which network connection generated the packet, and the destination is the address to which the packet should be routed.


IP addresses are composed of two parts; a network part and a host part (an extension to this two level hierarchy is described below), and are typically written as four dot-separated decimal bytes, eg. 192.83.123.129. The network part is assigned to an organisation by a central numbering authority, and the host part is assigned by the owner of the network number part to their individual IP connections. This mechanism ensures that all IP addresses are globally unique and delegates the authority for allocating addresses to organisational entities.

IP addresses are categorised as follows:

Class	Network Bits	Host Bits	Number of Networks (of this class)	Max Number of Connections (for each network number in this class)
A	8	24	128	16 Million
B	16	16	16,000	32,768
C	24	8	2 Million	256

CONNECTING TO AARNET

Normally, a unique network number is assigned to every physical network, be it an Ethernet segment, Token Ring LAN, FDDI ring or serial link. Clearly if this were the case the available network numbers would be very quickly exhausted, so a middle level of hierarchy is used called *subnetting* to allow the same network number to be used for multiple interconnected physical networks within the same organisation. Subnetting allows a configurable portion of the host part of a network number to be designated by the number's owner as a subnet network number, thus allowing many physical segments within the same organisation to be assigned unique subnet numbers rather than network numbers. For example, the host part of a class B network may be divided in half (or at any other bit position in the Host Part) to provide subnetting:



All hosts on the same physical network therefore share the same network (or subnet) number, but each has a unique host number. The fact that a network is subnetted is visible only to those systems connected to the subnetted network. This means that all subnets of a subnetted network must be interconnected.

3.2.2. Obtaining an IP Network Number

It is a requirement that any organisation wishing to join AARNet as a network affiliate member must have a number allocated by the Internet Network Information Centre (NIC) in the United States. If your organisation already has a NIC-allocated number, AARNet must be told what it is so that AARNet can arrange for it to be known to the rest of the Internet and configured into the AARNet routers (this does not automatically happen when a network number is allocated). If you do not yet have a proper network number, then AARNet can assign you one (for which there is no charge, by either AARNet or the NIC).

In fact, any organisation that is contemplating establishing an IP network should apply for a unique network number from the NIC (through AARNet), even if they have no immediate plans to establish a network level connection to AARNet. The effort required to renumber host and network systems if at some stage the organisation does become an AARNet member, or if the organisation wishes to connect to another organisation that by chance is using the same randomly chosen network number, is something that should be avoided if at all possible.

Because there are so few Class A networks defined, they are only ever assigned to very large organisations (eg. IBM, Hewlett-Packard, Stanford University, etc.). On the other hand, Class C networks are relatively small in terms of maximum number of assignable hosts and bits available for subnetting, and as a result are useful to only small sites. Class B networks are therefore the most popular even though it is unusual for an organisation to allocate even a moderate fraction of the 32,768 addresses that are available. This has led to a situation whereby class B network numbers are in relatively short supply compared to class C network numbers, and the NIC is now allocating to some organisations bundles of class C networks rather than a class B network.

The only problem with using bundles of class C networks instead of a single B is that this imposes routing overhead on the Internet, since instead of an organisation appearing as just one routing entry, it may appear as several. This does not scale well in a network as large as the Internet. To alleviate this problem the NIC limits the number of class C networks it will bundle to between 12 and 16, depending on an organisation's likely requirements (which must be described in the application for the network number). It is important in applying for a number that the number of physical LANs, serial lines, network routers and connected host systems (including personal computers) be identified, along with the expected growth over a three year period. From this information the NIC will make a decision on whether to provide a bunch of C's, or a B that can be subnetted.

The network number application form, and some guidelines on how best to complete it are included with these notes.

3.2.3. Routing

Charles Hedrick's paper *Introduction to Administration of an Internet-based Local Network* provides an excellent overview of the IP routing process, and how it can be manipulated within a UNIX environment. It also discusses the

CONNECTING TO AARNET

differences between bridging and routing in some detail, which is recommended to anyone contemplating the installation of anything other than a trivial LAN.

Currently AARNet supports routing interchanges with affiliate members using the Routing Information Protocol (RIP), and supplies a default route (0.0.0.0) to the affiliate network.

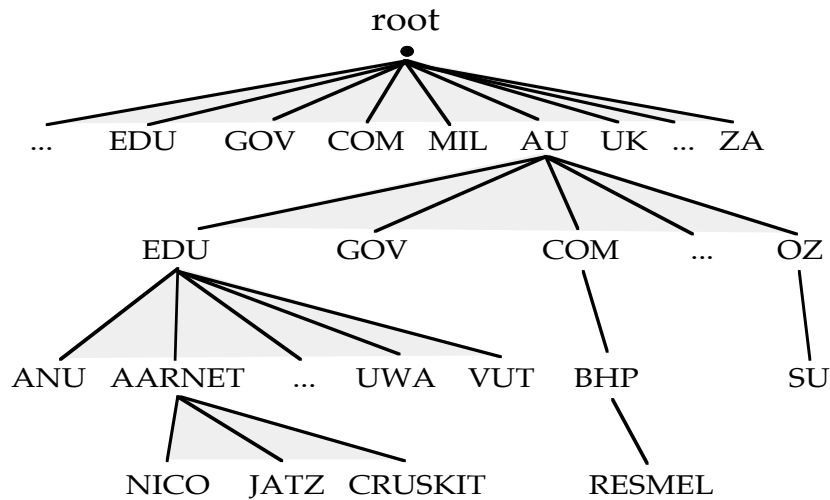
3.3. Domain Name System

3.3.1. Why ?

In the early days of the Internet, all host names and their associated IP addresses were recorded in a single file called HOSTS.TXT, maintained by the Network Information Centre in the USA. Not surprisingly, as the Internet grew so did HOSTS.TXT, so that in about 1985 it had become impractically large to distribute to all systems over the network, and impossible to keep up to date. The Internet Domain Name System (DNS) was developed as a distributed database to solve this problem. It's primary goal is to allow the allocation of host names to be distributed amongst multiple naming authorities, rather than centralised at a single point.

3.4.2. Structure

The DNS is arranged as a hierarchy, both from the perspective of the structure of the names maintained within the DNS, and in terms of the delegation of naming authorities. At the top of the hierarchy is the root domain "." which is owned by the US NIC. Ownership of the root domain gives the NIC the authority to allocate domains beneath the root, as shown in the diagram below:



The process of assigning a domain to an organisational entity is called delegating, and involves the owner of a domain (eg. the NIC) creating a sub-domain (AU) within an existing domain they own (the Root) and assigning the ownership and authority for allocating sub-domains of the new domain to another organisation (the University of Melbourne for AU). A complete DNS name, known as a full-qualified domain name, is obtained by writing the simple names obtained by tracing the DNS hierarchy from the leaf nodes to the root, from left to right, separating each name with a stop ".", eg.

JATZ . AARNET . EDU . AU .

is the name of a host system (JATZ) within the AARNET organisation, an educational (EDU) institution within Australia (AU), and

METRO.UCC.SU.OZ.AU.

is the name of a host system (METRO) within the University Computing Centre (UCC) of Sydney University (SU), that is retaining it's original ACSnet name (OZ), within Australia (AU).

CONNECTING TO AARNET

The sub-domains of the root are known as the top-level domains, and include the EDU (educational), GOV (government), MIL (military) and COM (commercial) domains. Although an organisation anywhere in the world can register beneath these three-character top level domains, the vast majority that have are located within, or have parent companies based in, the United States. The top-level domains represented by the ISO two-character country codes are used in most other countries, thus organisations in Australia are registered beneath AU.

The majority of country domains are sub-divided into organisational-type sub-domains. In some countries two character sub-domains are created (eg. AC.NZ for New Zealand academic organisations), and in others three character sub-domains are used (eg. COM.AU for Australian commercial organisations). Regardless of the standard adopted each domain may be delegated to a separate authority. The owners of the AU sub-domains are listed below:

EDU.AU	Academic and Educational Institutions	Geoff.Huston@AARNet.EDU.AU
GOV.AU	Governmental Agencies	Geoff.Huston@AARNet.EDU.AU
COM.AU	Commercial Entities	Warwick.Jackson@Praxa.COM.AU
CSIRO.AU	CSIRO	Bob.Smart@mel.dit.CSIRO.AU
OZ.AU	ACSnet	kre@munnari.OZ.AU (Robert Elz)
ORG.AU	Other Organisations	kre@munnari.OZ.AU (Robert Elz)

Organisations that wish to register a domain name, even if they do not plan to establish an AARNet connection in the immediate short term, should contact the owner of the domain which most closely describes their activities. If you require assistance, contact AARNet. Note that registering a domain within the AU domain does not require the involvement of the US NIC.

Even though the DNS supports many levels of sub-domains, delegations should only be made where there is a requirement for an organisation or organisational sub-division to manage their own name space. Any sub-domain owner must also demonstrate they have the technical competence to operate a domain name server (described below), or arrange for another organisation to do so on their behalf.

3.4.3. Domain Name Servers

The DNS is implemented as collection of inter-communicating nameservers. For each domain, there is one primary nameserver, and one or more secondary nameservers (see exceptions discussed below). At any given level of the DNS hierarchy, a nameserver for a domain has knowledge of all the immediate sub-domains of that domain.

Information is extracted from the DNS using resolver routines that direct queries to nameservers that forward the query through the DNS hierarchy to locate the data requested. For example, if a request was passed to the resolver routines, by the `gethostbyname` function perhaps, to return the IP address of `JATZ.AARNET.EDU.AU` the following (much simplified) steps would be executed by the resolver routines:

1. A (system specific) default nameserver is requested to provide the nameservers for the domain `AARNET.EDU.AU`, since this is the domain that contains information about `JATZ.AARNET.EDU.AU`.
2. If the nameserver is unable to answer this query, it will forward the request to one of the root nameservers.
3. The root nameservers are unlikely to be able to provide the nameservers for `AARNET.EDU.AU`, but instead return the nameservers of the "best match domain" that they have knowledge of, in this case `AU`.
4. One of the `AU` name servers is queried for the nameservers for `AARNET.EDU.AU`.
5. If the `AU` nameserver queried is unable to answer, it also provides it's "best effort", which for this server would be the nameservers for the `EDU.AU` domain.
6. An `EDU.AU` nameserver is queried to provide the nameservers for `AARNET.EDU.AU`.
7. One of which is queried for the IP address of `JATZ.AARNET.EDU.AU`.
8. The address is returned to the application that issued the `gethostbyname`.

Clearly if this simple process was repeated for every query the DNS would be impractically slow and inefficient. Most nameservers implementations therefore cache the results of queries for later use, and can also be configured to forward queries to other nameservers acting as local "stores" of answers to common queries.

CONNECTING TO AARNET

Secondary nameservers provide backup to the primary nameserver when it is not operational, and further improve the overall performance of the DNS, since the nameservers of a domain that respond to queries most quickly are used in preference to any others. Thus, in addition to having a primary nameserver on site, each organisation should have at least one secondary on site, and another elsewhere on AARNet (or the Internet), preferably well connected (ie. on or near the AARNet backbone). AARNet can assist you on locating an off-site secondary, usually at a regional hub site. This is particularly important for affiliate members with slow speed AARNet connections (eg. 9.6K) to reduce use of their link to support the DNS.

3.4.4. Setting up a Domain Name Server

Nameservers running on UNIX systems are based on the Berkeley Internet Name Daemon (BIND), which usually runs as the daemon `named`. Configuring BIND is a simple process that is full of large pit falls. Because of the way nameservers closely interact with each other, and the fact that the majority of BIND implementations are rather more delicate than we would like, mistakes are easily propagated throughout the DNS and often quite difficult to resolve.

To start with, there is the `named.boot` file that `named` reads at startup. It can usually be found in `/etc` (or `/usr/local/etc`) of a Berkeley based UNIX system. This is a simple example:

```
a directory    /usr/local/domain
;
b ;type        domain          source host/file  backup file
;
c cache        .               root.cache
;
d primary      EDU.AU           zone/EDU.AU
primary        aarnet.EDU.AU    zone/aarnet.EDU.AU
primary        avcc.EDU.AU      zone/avcc.EDU.AU
e primary      130.139.IN-ADDR.ARPA  rev/130.139.rev
primary        123.83.192.IN-ADDR.ARPA  rev/123/83.192.rev
primary        0.0.127.IN-ADDR.ARPA    rev/loopback
;
f secondary    AU               128.250.1.21     bak/AU
g secondary    latrobe.EDU.AU    132.147.1.2      bak/latrobe.EDU.AU
;
h forwarders   139.130.4.4
```

Where:

- a** Specifies the directory in which other files specified in the `named.boot` file can be found.
- b** Comments are indicated with a leading semi-colon.
- c** This says that the file `root.cache` (in `/usr/local/domain`) has information about root nameservers that can get the nameserver started. This information is replaced as soon as possible with "real" information obtained from the root nameservers.
- d** The `primary` keyword identifies a domain for which this nameserver shall act as a primary nameserver. The definition of the domain itself is called a *zone file* and in this case is the file `aarnet.EDU.AU` in the directory `/usr/local/domain/zone`. A nameserver can be primary for several domains, including sub-domains of domains it is primary for, eg. `aarnet.EDU.AU` and `EDU.AU`.
- e** The DNS contains more than just domain-name to IP address mappings. This line and the next indicate that this nameserver will act as a primary nameserver for the so-called reverse mappings of IP address to domain-name for the Class B network 139.130.0.0 and the Class C network 192.83.123.0. This feature is described below.
- f** Configuring a nameserver to be a secondary nameserver for a domain is done using the `secondary` keyword and providing the name of the domain to be secondaried, the IP address of the primary name server from which the corresponding zone file will be loaded (by this `named` communicating with the `named` on the primary

CONNECTING TO AARNET

nameserver), and the name of a file in which this zone file will be stored. Secondary zone transfers occur when named starts, and at intervals defined within the zone file.

- g** This nameserver is secondary for another domain. As noted above, it's good to have at least one secondary "somewhere else", in case your organisation's link to AARNet is down for an extended period.
- h** The `forwarder` line specifies the IP address of one or more nameservers to which this nameserver should pass all queries and responses. These "forwarders" act as caches of frequently requested DNS questions and answers, potentially improving the speed of DNS lookups, reducing the load on nameservers higher up in the DNS hierarchy (particularly the root servers), and limiting the amount of DNS traffic carried over major trunk links, eg. from Australian to the the rest of the Internet. AARNet operates a system located at the topological centre of the network to support this function, and recommends that all Australian sites include the IP address of this system, 139.130.4.4, in the `forwarder` line of any `named.boot` configuration file.

The current root nameservers are included in the following `root.cache` file:

```
;  
; Initial cache data for root domain servers.  
;  
.           99999999   IN      NS      NS.NASA.GOV.  
.           99999999   IN      NS      A.ISI.EDU.  
.           99999999   IN      NS      C.NYSER.NET.  
.           99999999   IN      NS      TERP.UMD.EDU.  
.           99999999   IN      NS      NS.NIC.DDN.MIL.  
.           99999999   IN      NS      AOS.BRL.MIL.  
.           99999999   IN      NS      NIC.NORDU.NET.  
;  
; Prep the cache (hotwire the addresses). Order does not matter  
;  
NS.NIC.DDN.MIL.  99999999   IN      A       192.112.36.5  
A.ISI.EDU.AU.    99999999   IN      A       128.9.0.107  
                99999999   IN      A       26.3.0.103  
AOS.BRL.MIL.    99999999   IN      A       128.120.1.2  
                99999999   IN      A       192.5.25.82  
C.NYSER.NET.    99999999   IN      A       192.33.4.12  
NIC.NORDU.NET.  99999999   IN      A       192.36.148.17  
NS.NASA.GOV.    99999999   IN      A       192.52.195.10  
                99999999   IN      A       128.102.16.10  
TERP.UMD.EDU.   99999999   IN      A       128.8.10.90
```

This shows the basic format of a *zone file*. It contains Internet (IN), nameserver (NS), and IP address (A) resource records. The NS records specify, on the right-hand-side, the names of the systems that are nameservers for the root domain, as indicated in the label field (a ".") on the left-hand-side. For each nameserver so specified there are A records that provide the IP address(es) of the corresponding systems. Note that if there is no left-hand-side label, then the previous label is used. The value 99999999 is the Time-To-Live value for the resource record; normally the default value for the current domain is used (see below). None of the root nameservers are located in Australia.

Let's now look at the start of zone file `zone/aarnet.EDU.AU`:

```
; Authoritative data for aarnet.EDU.AU (ORIGIN assumed aarnet.EDU.AU)  
;  
; NB: format of serial number: yyyymmddxx  
;  
@           IN      SOA      jatz.aarnet.EDU.AU.  mit900.jatz.aarnet.EDU.AU.  (  
                1993092801      ; Serial  
                10800           ; Refresh - 3 hours  
                1800            ; Retry - 30 minutes  
                3600000          ; Expire - 1000 hours (41.6 Days)  
                43200           ; Minimum - 12 hours
```

The "origin" for this zone file is `aarnet.EDU.AU` since this is the domain name specified with the `primary` keyword in the `named.boot` file.

CONNECTING TO AARNET

BEWARE! Any name or label that appears on either the left or right-hand-side of a resource record that does not have a terminating full stop will have the origin added to the name/label. Missing full stops are one of the most common causes of error in DNS zone files.

- The origin (or current domain) can be explicitly referred to using the "@" symbol. This is used in this example as the label for the Start Of Authority (SOA) record that must appear at the beginning of any zone file defining a domain. The first two fields of the SOA record are the name of the primary nameserver, and the electronic mail address of the person responsible for this domain's nameserver (with the traditional "@" replaced by a "."). This is followed by a number of important numbers.
- The **serial number** identifies the current version of the zone file. The format shown above is not mandatory, but is recommended; the N.M format for example is interpreted by BIND as $(N*10000 + M)$ - it's not a decimal number. The only requirement is that every time any change is made to the zone file (other than to just comments) the serial number is incremented in a monotonically increasing fashion. If you do not change the serial number, then any changes made will not be propagated to the secondary nameservers for the domain, and the DNS becomes inconsistent.

All the other values are timers, measured in seconds.

- The **refresh** rate, controls how frequently secondary nameservers check to see if their copy of the zone file is the same as the primary's (ie. has the same serial number). This should be set to a value appropriate for the communications links between the various servers. If they are very close, or on a high bandwidth network, then an hour or so is probably OK; if they're further apart, then several hours would be better. In any case, it should be many times less than the expire counter.
- Next is **retry**. This determines how frequently a secondary should attempt to communicate with the primary (for refresh purposes) if its first attempt fails. Again, this should be set based on the link types, but also based on common sense. If the primary takes about 1/2 an hour to boot, setting refresh to anything less is pointless.

The values for refresh and retry put extra load on the nameservers (especially the secondaries) and on the network if they're too low, but conversely allow nameserver inconsistencies for longer than is really desirable if they're too high. A value of 2 or 3 hours (7200 or 10800 seconds) for refresh, and 30 minutes (1800 seconds) for retry is suggested.

- The **expire** number controls how long a secondary is permitted to give out information for the domain even though it has been unable to contact the primary. Typically should this be long enough so that the secondary can keep working through long network outages, system hardware crashes, or anything which may affect operation of the primary nameserver. 1000 hours (3600000 seconds) is the most usual value here.
- Finally **minimum** sets the minimum cache validity period for any records in the zone. This is effectively a Time To Live (TTL) value that specifies how long remote servers will remember DNS data sent to them before asking again. Its possible to assign individual TTL values to each resource record, but all are required to be at least as large as this minimum. Minimum value should be set reasonably high, a value of a day is generally regarded as acceptable, 12 hours is perhaps OK, anything less is too low.

Note that the value of minimum field effectively controls how quickly you can make changes to the zone file since it determines how long it will take between when you decide that you want to make a change, and when you can confidently say that the change has propagated everywhere (by virtue of the changed record's TTL expiring).

Immediately after the SOA record are a list of NS records that specify the nameservers (primary and all secondaries) for the domain (remember that resource records without a label assume the label of the preceding record, ie. the origin, "@", or aarnet.EDU.AU.):

```
IN  NS  jatz.aaarnet.EDU.AU.
IN  NS  cruskit.aarnet.EDU.AU.
IN  NS  munnari.OZ.AU.
```

Each system listed in these NS records **must** either be the primary nameserver, or be configured as a secondary (ie. be running named and have an appropriate `secondary` statement in that nameserver's `named.boot` file). Just naming a system with an NS record as shown above will not make the named system a secondary nameserver. Note

CONNECTING TO AARNET

that the NS records that appear in this list must be the same as the NS records that appear in the zone file of the parent domain (in this example EDU.AU).

The remainder of the aarnet.EDU.AU zone file contains information about host systems with the domain:

```
jatz          IN      A           139.130.204.4
              IN      MX           10          jatz
sao           IN      A           139.130.204.5
              IN      HINFO        "VaxStation" "VMS"
              IN      MX           10          sao
scotch-finger IN      A           139.130.204.8
              IN      MX           10          jatz
```

- A records were mentioned above; they define the IP address(es) associated with the label on the left-hand-side.
- The HINFO record is optional and supplies two information strings, a description of the hardware platform, and the name of the operating system associated with this name. These strings are meant to be selected from a defined list (the Internet Assigned Numbers RFC), but this requirement is not strictly adhered to.
- The MX records provide Mail Exchange information. The right-hand-side of an MX record specifies the name of the system to which electronic mail for the domain name specified on the left-hand-side should be forwarded.

Thus mail for `jatz.aarnet.EDU.AU` should be forwarded to `jatz`. This would be the normal case where a system is configured to accept mail from the Internet (described below). For systems that are generally not capable of exchanging electronic mail directly, the MX record will point at some form of mail server, eg. as shown above for the Macintosh system `scotch-finger.aarnet.EDU.AU`. The decimal value in the MX record is a preference that selects one MX record over another, for example

```
@            IN      MX           10          jatz
              IN      MX           20          munnari.OZ.AU.
              IN      MX           100         relay1.UU.NET.
              IN      MX           100         relay2.UU.NET.
```

This defines a series of MX records that indicate (for this domain) that mail for `aarnet.EDU.AU` should be directed to `jatz`, but if `jatz` is not reachable, then try `munnari.OZ.AU`, and failing that either `relay1.UU.NET` or `relay2.UU.NET` with equal preference. Note that any system that is the target of an MX record should be configured to accept or forward mail for the domain name associated with the MX.

At the end of this example zone file the following resource records are found:

```
;
; Standard domain names
;
mailhost      IN      CNAME    jatz
localhost    IN      A         127.0.0.1
loghost      IN      CNAME    localhost
;
; Glue records
;
munnari.OZ.AU. IN      A         128.250.1.21
              IN      A         192.43.207.1
```

The CNAME record allows aliases to be defined within the DNS. Since several applications assume the existence of the hostnames `mailhost` and `loghost`, these are provided as aliases for actual hosts. Because they are just aliases, the left-hand-side of a CNAME record cannot be used anywhere else on the right-hand-side of a resource record.

The name `localhost` is configured to be the internal loopback IP address which is defined to be `127.0.0.1`.

The glue records provide the IP addresses of the nameservers listed with the NS records under the SOA that are not defined within this zone file. Because this information is statically defined, if a secondary nameserver is to have its IP address changed, the owner of the domain for which it is being a secondary for must be informed, since the glue records will have to be updated.

CONNECTING TO AARNET

The format of the zone files that support the reverse mapping of IP address to domain name use different resource records but are in all other respects formatted in the same way. For our small example domain the following zone might be used;

```
; Authoritative data for 130.139.in-addr.arpa
;
@           IN   SOA   jatz.aarnet.EDU.AU. mit900.jatz.aarnet.EDU.AU. (
                                1993092001   ; Serial
                                10800        ; Refresh - 3 hours
                                1800        ; Retry - 30 minutes
                                3600000     ; Expire - 1000 hours
                                43200      ) ; Minimum - 12 hours
           IN   NS    jatz.aarnet.EDU.AU.
           IN   NS    anu.anu.EDU.AU.
;
4.204      IN   PTR   jatz.aarnet.EDU.AU.
5.204      IN   PTR   sao.aarnet.EDU.AU.
8.204      IN   PTR   scotch-finger.aarnet.EDU.AU.
;
; Glue
;
anu.anu.EDU.AU. IN  A    130.56.4.1
```

The records in this domain are of the form $x4.x3.x2.x1.IN-ADDR.ARPA$ representing the IP address $x1.x2.x3.x4$. The right-hand-side of the Pointer (PTR) record indicates the host associated with this address. The reason the address bytes are backwards is because IP address have the most significant component on the left, but domain names have the most significant component on the right. The in-addr.arpa domains are used by the `gethostbyaddr` to return the name of a system given its IP address.

Since the parent domain of `130.139.in-addr.arpa` is `in-addr.arpa`, you need to contact the owner of `in-addr.arpa`, the US NIC, to get your sub-domain of `in-addr.arpa` delegated to your nameservers. Once you have established an AARNet connection, a form to arrange the delegation of an `in-addr.arpa` sub-domain can be retrieved from `aarnet.edu.au:pub/netinfo/in-addr-template`, filled in and sent via electronic mail to `hostmaster@nic.ddn.mil`. An `in-addr.arpa` domain is also required for the loopback address described above (this is `rev/loopback`):

```
@           IN   SOA   jatz.aarnet.EDU.AU. mit900.jatz.aarnet.EDU.AU. (
                                1992092001   ; Serial
                                10800        ; Refresh - 3 hours
                                1800        ; Retry - 30 minutes
                                3600000     ; Expire - 1000 hours
                                43200      ) ; Minimum - 12 hours
           IN   NS    jatz.aarnet.EDU.AU.
1          IN   PTR   localhost.
```

Because the loopback domain is just that, the information in this zone file is not passed to any other server and thus the values assigned to the various timers is not critical.

3.4.5. Making applications use the DNS

Having configured BIND by creating the appropriate boot and zone files, and started the nameserver(s) by launching `named`, applications that make use of the network and the associated name to IP address functions must somehow be made to direct queries to the DNS through the DNS resolver, in addition too, or instead of using a `hosts` file or a Network Information Service (NIS), formerly Yellow Pages (YP), server. The method by which this is done varies from one UNIX implementation to another. In all cases however, the resolver will consult the file `/etc/resolv.conf` to obtain the name of the domain in which this machine has been configured, and the addresses of nameservers to be used to initiate DNS queries, eg. `resolv.conf` on a system within the `aarnet.EDU.AU` domain might contain:

CONNECTING TO AARNET

```
domain aarnet.EDU.AU
nameserver 139.130.204.4
nameserver 139.130.204.2
nameserver 130.56.4.1
```

Note that the nameservers specified do not need to be primary or even secondary for the corresponding domain, any nameserver will be able to respond to DNS queries (but there are obvious efficiencies in including at least one server for the local domain).

Under SunOS, use of the DNS can be enabled by either building the NIS host maps using the `-b` flag to `makedbm` or by replacing the `libc` run-time shareable library with a version that only uses the DNS for resolving host names and addresses. In either case, because dynamically linked shareable libraries are used, applications do not need to be recompiled or relinked. The DNS-capable library is available over the network, or from Sun Microsystems.

Under ULTRIX there is a file called `/etc/svc.conf` that specifies which methods (DNS, hosts, NIS) should be used to resolve names (and other things), and the order they are applied in. Like SunOS, this is implemented through shareable libraries, so again no application changes are required. The `svcsetup` command can be used to modify `svc.conf` to enable use of BIND.

A similar facility is provided under System V Release 4 implementations, but in this case the relevant file is `/etc/netconfig`.

Some UNIX systems come with resolver routines that detect the presence of `resolv.conf` and automatically use the DNS. Others will include one set of network utilities (telnet, ftp, etc.) that use the hosts file (or NIS), and another set that can be used instead of those that use the DNS. In all cases, the vendor's supplied documentation should be consulted.

3.5. Electronic Mail

Within an IP network, mail is generally transported using the Simple Mail Transfer Protocol (SMTP). Systems that wish to send and receive electronic mail therefore need to run an SMTP-capable daemon to accept incoming electronic mail connections. This daemon is usually also responsible for delivering outgoing mail, and thus is known as the Mail Transport Agent (MTA). The functions of reading and composing mail are incorporated into a separate Mail User Agent (MUA) that may be different for every user. Examples of MUA's include the (in)famous Berkeley mail program, `mh` and its X-windows incarnation `xmh`, `elm`, and vendor supplied tools such as Sun's Open Look compliant `MailTool`.

A wide variety of MTA's are also available, but the most widely used within the Internet is `sendmail`. Because of the size and apparent complexity of the `sendmail` configuration file many administrators are reluctant to delve too deeply into its inner workings. The problem is further complicated by the fact that vendor supplied configurations incorporate many bells and whistles, many of which are not necessary for a simple IP connection.

One alternative to basic `sendmail` is the IDA `sendmail` package which allows a `sendmail` configuration file to be generated from a series of simple `m4` macro definitions. It also provides a number of other useful features that improve the functionality and management aspects of the mail environment such as the ability to modify `From:` addresses on a per user basis and the use of `dbm` files to store additional configuration information. It is available from `munnari.oz.au:pub/sendmail-5.65c+IDA-1.4.4.1.tar.Z`.

Whatever path is chosen in configuring your mail system some basic pointers should be observed:

- Make sure `sendmail` is configured to recognise and accept mail directed to the system it is running on. If mail for the hosts system's domain (rather than the host itself), or even other domains is directed to the `sendmail` MTA as well (via `MX` records for example), it must be setup to accept such mail as well, and either treat it as local or pass it to some other mail system for final delivery as appropriate.
- The mail aliases file should have an entry for `postmaster`. By convention, the address `postmaster@<hostname>` is expected to reach someone responsible for the system, to enable systems administrators to communicate problems and advice. It is also recommended that the domain alias `postmaster@<domain-name>` is also supported to assist people in trying to contact users within your organisation. This latter form will usually require an appropriate `MX` record, as well as `sendmail` changes.
- Electronic messages that leave your domain or organisation should have their return address (specified in the `From:` field of the message header) *fully qualified*. This means that mail from `gih@jatz` should have the

CONNECTING TO AARNET

From: address set to `gih@jatz.aarnet.edu.au` (or `gih@aarnet.edu.au`) - if this is not done then the recipient will be unable to reply to the message since their (the recipient's) default domain will be appended to `pte@jatz`.

- Always attempt to deliver mail using the MX records of the DNS. Some systems (SunOS is one) come with two versions of sendmail, one that does not consult MX records and one that does (`/usr/lib/sendmail.mx`).
- The only exceptions to the above rule are for the pseudo domains UUCP and BITNET. These do not actually exist within the DNS, and therefore need to be explicitly dealt with within the sendmail configuration. This involves forwarding messages of the form `user@host.UUCP`, and `user@host.BITNET` to a system that knows how to handle them. Within Australia, messages of either form should be directed to `munnari.oz.au`. Many sendmail configurations include a definition of "smart" or "gateway" mailer" which specifies the name of a system that will handle mail this sendmail cannot, and this is how UUCP and BITNET addresses are often passed on to a smarter mailer.

3.6. Security

3.6.1. Welcome to the Internet

Once connected to AARNet, your organisation has access to all the (considerable) resources of the Internet. Unfortunately, all of the Internet also has access to you. There continues to be a lot of publicity regarding the success of "hackers" in obtaining unauthorised access to systems connected to the Internet, yet with relatively small amounts of effort it is possible to implement and maintain a level of security that is adequate for most installations. For organisations that have very high security requirements, such as the military and some commercial companies, there are tools and techniques that permit very secure access to the Internet.

At the same time that the network is a source of security concerns, it is also a source of help and assistance. There are many publicly available papers on how to make host systems (particularly UNIX systems) more secure, and a broad assortment of utilities, system enhancements and monitors to make the system managers life easier. In addition the Computer Emergency Response Team (CERT), based at Carnegie Mellon University acts as a focal point for the collection of security related material, and issues security alerts whenever bugs or problems are identified with vendor supplied software. Such alerts usually include workaround solutions to problems, and/or directions on how to obtain a fix or patch for the problem over the network. They also provide a 24 hour international help desk to assist you if you believe that your system has been compromised. The CERT can be contacted as follows:

Internet E-mail: `cert@cert.org`
Telephone: `+1 412- 268 7090` (24-hour hotline)
CERT/CC personnel answer 7:30a.m. to 6:00p.m. EST(GMT-5)/EDT(GMT-4), and are on call for emergencies during other hours.

Computer Emergency Response Team/Coordination Center (CERT/CC)
Software Engineering Institute
Carnegie Mellon University
Pittsburgh, PA 15213-3890

The CERT/CC issues CERT Advisories, which warn you about problems and inform you about preventive techniques. The CERT/CC maintains a CERT Advisory mailing list, which is also distributed via the USENET newsgroup `comp.security.announce`. If you are unable to receive this newsgroup, send mail to: `cert-advisory-request@cert.org` to be added to the Security Advisory mailing list.

Past advisories and other computer security related information are available for anonymous FTP from the `cert.org` (192.88.209.5) system.

3.6.2. Security Measures

Security measures fall into essentially two categories. The first involves making the end systems themselves secure so that even if they were open to unlimited access, it would not be possible to obtain unauthenticated access to resources of any kind maintained on that system. The second category operates at the network level, and is implemented using the packet filtering capabilities of modern routers. Controls at this level are very secure, yet at

CONNECTING TO AARNET

the same time can be very finely tuned, eg. permit incoming electronic mail connections to hosts A and B, and file transfers to host Z, but deny incoming access, of any kind, to all other hosts.

It should be noted that because some features of the UNIX system, such as the Berkeley "r"- commands, allow transparent and semi-transparent access to other computing systems, the level of security maintained by an organisation as a whole is generally no higher than the level of the most insecure system at that site.

The level of access controls chosen by any organisation will be a function of resources and requirements, but the following very fundamental measures should be undertaken in all cases:

- No account should ever have a null password.
- Never provide guest accounts that do not require a password to obtain access.
- Ensure that your users use "good" passwords (ie. use combinations of letters and digits, mixed upper and lower case, special symbols and avoid simple words).
- Be sure you understand the implications of offering network based services that might allow remote systems to obtain sensitive data such as password files
- Read the sections on security in the manuals provided by your computer vendor, and ensure you have installed all appropriate software bug fixes and patches.

3.6.3. Known UNIX System Configuration Problems

The following summary is extracted from a CERT advisory.

1. Weak passwords

Intruders often use `finger` or `ruser` to discover account names and then try simple passwords. Encourage your users to choose passwords that are difficult to guess (for example, words that are not contained in any dictionary of words of any language; no proper nouns, including names of "famous" real or fictitious characters; no acronyms that are common to computer professionals; no simple variations of first or last names. Furthermore, inform your users not to leave any clear text username/password information in files on any system.

A good heuristic for choosing a password is to choose an easy-to-remember phrase, such as "By The Dawn's Early Light", and take the first letters to form a password. Insert in some punctuation or mixed case letters as well. For the phrase above, one example password might be: `bT}DeL{.` (DO NOT use this sample phrase for your password!)

If intruders can get a password file, they will usually take it to another machine and run password guessing programs on it. These programs involve large dictionary searches and run quickly even on slow machines. The experience of many sites is that most systems that do not put any controls on the types of passwords used probably have at least one password that can be easily guessed.

If you believe that your password file may have been taken, change all the passwords on the system. At the very least, you should always change all system passwords because an intruder may concentrate on those and may be able to guess even a reasonably 'good' password.

2. Use of TFTP (Trivial File Transfer Protocol) to steal password files

To test your system for this vulnerability, connect to your system using `tftp` and try `'get /etc/passwd'`. If you can do this, anyone else on the network can probably get your password file. To avoid this problem, either disable `tftpd` if you don't require it or ensure that it is configured with restricted access.

If you believe your password file may have been taken, the safest course is to change all passwords in the system.

CONNECTING TO AARNET

3. Accounts without passwords or known passwords (accounts with vendor-supplied default passwords are favourites)

Intruders often exploit system default passwords that have not been changed since installation. Be sure to change all default passwords when the software is installed. Also, be aware that product upgrades can quietly change account passwords to a new default. It is best to change the passwords of default accounts after updates are applied.

Scan your password file for extra UID 0 accounts, accounts with no password, or new entries in the password file. Do not allow any accounts without passwords. Remove entries for unused accounts from the password file. To disable an account, change the password field in the `/etc/passwd` file to an asterisk '*', and change the login shell to `/bin/false` to ensure that an intruder cannot login to the account from a trusted system on the network.

4. Holes in sendmail

Make sure that you are running the latest sendmail from your vendor. BSD 5.65 fixes all known holes. To establish which version of sendmail that you are running, use telnet to connect to the SMTP port (25) on your system: `telnet localhost 25`

5. Old versions of FTP; misconfigured anonymous FTP

Make sure that you are running the most recent version of `ftpd`, which is the Berkeley version 5.60 of July 22, 1990. Check with your vendor for information on configuration upgrades. Also check your anonymous FTP configuration. It is important to follow the instructions provided with the operating system to properly configure the files and directories available through anonymous FTP (for example, file and directory permissions, ownership and group). Note that you should not use your system's standard password file or group file as the password file or group file for FTP. The anonymous FTP root directory and its two subdirectories, `etc` and `bin`, should not be owned by `ftp`.

6. Fingerd hole used by the Morris Internet worm

Make sure that you're running a version of `finger` that is more recent than November 1988. Numerous Berkeley-derived versions of UNIX were vulnerable.

7. Inappropriate network configuration file entries

Several vendors supply `/etc/hosts.equiv` files with a '+' (plus sign) entry. The '+' entry should be removed from this file because it means that your system will trust all other systems. Other files that should not contain a '+' entry include `/etc/hosts.lpd` and all `~/ .rhost` files on the system. These files should not be world-writable. We recommend that you do not support the following services in your `/etc/inetd.conf` file unless you specifically require them:

```
port 11 - systat
port 69 - tftp
port 87 - link
```

8. Misconfiguration of UUCP

If your machine supports UUCP, check the `L.cmds` (Permissions) file and ensure that only the commands you require are included. This file should be owned by root (not by `uucp!`) and world-readable. The `L.sys` (Systems) file should be owned by `uucp` and protected (600) so that only programs running `setuid uucp` can access it.

9. Inappropriate 'secure' settings in /etc/ttys and /etc/ttytab

Check the file `/etc/ttys` or `/etc/ttytab` depending on the release of UNIX being used. The default setting should be that no terminal lines, pseudo terminals or network terminals are set secure except for the console.

CONNECTING TO AARNET

10. Inappropriate entries in /usr/lib/aliases

Examine the `/usr/lib/aliases` (mail alias) file for inappropriate entries. Some alias files include an alias named 'uudecode' or just 'decode'. If this alias exists on your system, and you are not explicitly using it, then it should be removed.

11. Inappropriate file and directory protections

Check your system documentation to establish the correct file and directory protections and ownership for system files and directories. In particular, check the `'/'` (root) and `'/etc'` directories, and all system and network configuration files. Examine file and directory protections before and after installing software or running verification utilities. Such procedures can cause file and directory protections to change.

12. Old versions of system software

Older versions of operating systems often have security vulnerabilities that are well known to intruders. To minimize your vulnerability to attacks, keep the version of your operating system up-to-date and apply security patches appropriate to your system(s) as soon as they become available.

(Section 3.6.3 Copyright © 1991, 1992 Carnegie Mellon University)

3.7. Network News

There are three basic software components associated with setting up a network news system over an IP network such as AARNet.

- The news software itself. The most widely used is called C-news. It handles the reception, posting, expiration and general management of news articles. INN is also suitable for high volume news sites.
- The Network News Transport Protocol (NNTP) software. This client and server program to enable transmission and reception of news articles both to and from your site, and from a news server at you site to news reading clients.
- A news reader that allows users to read, reply to and post news articles. There are a many readers available for UNIX including `vn`, `rn`, `trn`, `tin` and `nn`. An X-windows reader, called `xrn`, and a wide range of PC and Macintosh based readers are also available.

All this software is freely available over the network. A collection of tailored distributions and a short paper on how to proceed with the installation of the various pieces can be found in `wolfen.cc.uow.edu.au:pub/news`.

Unless you have made other arrangements, the AARNet regional hub site will be able to provide a network news feed via NNTP to your site. All news is brought into Australia by a single system (`munnar1.oz.au`) which then redistributes complete copies to each AARNet hub site. This minimises the volume of traffic carried over the AARNet backbone, so arbitrary feeds from arbitrary locations are discouraged (unless they are low in volume).

Appendix A. AARNet Members

A.1 AARNET MEMBERS

Australian Academic and Research Network aarnet.edu.au		Swinburne Institute of Technology swin.oz.au
Australian Catholic University acu.edu.au		University College of Northern Victoria ucnv.edu.au
Australian Defence Force Academy adfa.oz.au		University of Adelaide adelaide.edu.au
Australian National University anu.edu.au		University of Canberra canberra.edu.au
Australian Vice Chancellors' Committee avcc.edu.au		University of Central Queensland ucq.edu.au
Ballarat University College buc.edu.au		University of New England une.oz.au
Bond University bu.oz.au		University of New South Wales unsw.oz.au
Charles Sturt University csu.edu.au		University of Newcastle newcastle.edu.au
Curtin University of Technology curtin.edu.au		University of Queensland uq.oz.au
Deakin University deakin.oz.au		University of South Australia unisa.edu.au
Edith Cowan University cowan.edu.au		University of Southern Queensland usq.edu.au
Flinders University flinders.edu.au		University of Sydney su.oz.au
Griffith University gu.edu.au		University of Tasmania utas.edu.au
James Cook University jcu.edu.au		University of Technology, Sydney uts.edu.au
La Trobe University latrobe.edu.au		University of Western Australia uwa.oz.au
Macquarie University mq.edu.au		University of Western Sydney uws.edu.au
Melbourne University unimelb.edu.au		University of Wollongong uow.edu.au
Monash University monash.edu.au		Victoria College viccol.edu.au
Murdoch University murdoch.edu.au		Victoria University of Technology vut.edu.au
Northern Territory University ntu.edu.au		Victorian College of Agriculture and Horticulture vcah.edu.au
Queensland University of Technology qut.edu.au		
Royal Melbourne Institute of Technology rmit.oz.au		

A.2 CSIRO

Audit Group audit.csiro.au		Division of Manufacturing Technology dmt.csiro.au
Australia Telescope National Facility atnf.csiro.au		Division of Material Science and Technology mst.csiro.au
Australian Animal Health Laboratory aahl.dah.csiro.au		Division of Mathematics and Statistics dms.csiro.au
Biometrics Group biom.csiro.au		Division of Mineral Process and Engineering dmpe.csiro.au
Centre for Environmental Mechanics enmech.csiro.au		Division of Mineral Products dmp.csiro.au
Division of Animal Health (Parkville) mel.dah.csiro.au		Division of Plant Industry pi.csiro.au
Division of Animal Production prospect.anprod.csiro.au		Division of Radio Physics rp.csiro.au
Division of Applied Physics dap.csiro.au		Division of Soils soils.csiro.au
Division of Atmospheric Research dar.csiro.au		Division of Tropical Animal Production tap.csiro.au
Division of Biomolecular Engineering dbe.csiro.au		Division of Tropical Crops and Pastures tcp.csiro.au
Division of Building Construction and Engineering dbce.csiro.au		Division of Water Resources dwr.csiro.au
Division of Chemicals and Polymers chem.csiro.au		Division of Wildlife and Ecology dwe.csiro.au
Division of Coal and Energy Technology dcet.csiro.au		Division of Wool Technology dwt.csiro.au
Division of Entomology ento.csiro.au		Information Services Branch isb.csiro.au
Division of Exploration Geoscience deg.csiro.au		Institute of Information Sciences and Engineering iise.csiro.au
Division of Food Processing dfp.csiro.au		Joint Supercomputer Facility jsf.csiro.au
Division of Geomechanics geomechanics.csiro.au		Management Information Systems mis.csiro.au
Division of Human Nutrition dhn.csiro.au		Marine Laboratories (Oceanography + Fisheries) ml.csiro.au
Division of Information Technology dit.csiro.au		Tropical Forest Research Centre tfrc.csiro.au

CONNECTING TO AARNET

A.3 NETWORK AFFILIATE MEMBERS

Aeronautical Research Laboratory	arl.gov.au	Department of Prime Minister and Cabinet	
AIDAB - Australian International Development Assistance Bureau	aidab.gov.au		pmc.gov.au
Anglo Australian Observatory	aao.gov.au	Department of the Parliamentary Library	aph.gov.au
Australia Computing & Communications Institute	acci.com.au	Department of Transport and Communications	dotc.gov.au
Australian Antarctic Division	antdiv.gov.au	Earth Resource Mapping	erm.oz.au
Australian Artificial Intelligence Institute	aaii.oz.au	Edex	edex.edu.au
Australian Bibliographic Network, National Library of Australia	abn.nla.gov.au	Environmental Protection Authority	epa.gov.au
Australian Broadcasting Commission, R&D Laboratory	abc.gov.au	Environmental Resources Information Network (ERIN)	erin.gov.au
Australian Bureau of Agriculture and Resource Economics	abare.gov.au	Exicom Pty Ltd	exicom.oz.au
Australian Electoral Commission	aec.gov.au	Fluid Thinking Pty Ltd	fluid.com.au
Australian Heritage Commission	ahc.gov.au	Geelong and District Water Board	gdwb.oz.au
Australian Institute of Marine Science	aims.gov.au	Great Barrier Reef Marine Park Authority	gbrmpa.gov.au
Australian Maritime College	amc.edu.au	Internode Systems	internode.com.au
Australian Numerical Simulation and Modelling Services (ANSAMS)	ansto.gov.au	Ipex Information Technology Group	ipex.com.au
Australian Oceanographic Data Centre, RAN Hydrographic Office	aodc.gov.au	IPS Radio and Space Services	ips.oz.au
Australian Road Research Board	arrb.gov.au	James Bennett Pty Ltd	bennett.com.au
Australian Seismological Centre	ausseis.gov.au	LISWA - Library and Information Service of WA	liswa.gov.au
Australian Sports Commission	ausport.gov.au	Melbourne Information Technology Services	mits.com.au
Australian Surveying and Land Information Group	auslig.gov.au	Museum of Victoria	mus.vic.gov.au
AWA Defence Industries	awadi.com.au	National Centre for Studies in Travel and Tourism	ncstt.edu.au
BHP Research Laboratories	bhp.com.au	National Library of Australia	nla.gov.au
Biznet.com.au Pty Ltd	biznet.com.au	National Resource Information Centre	nric.gov.au
Bureau of Meteorology	bom.gov.au	Northern Territory Department of Mines and Energy	dme.nt.gov.au
Bureau of Mineral Resources	bmr.gov.au	NSW Agriculture	agric.nsw.gov.au
Bureau of Resource Sciences	brs.gov.au	NSW TAFE	tafensw.edu.au
Canberra Deep Space Communication Complex (Tidbinbilla)	cdscc-net.jpl.nasa.gov	Pacific Power	powernet.gov.au
CLSI	clsi.com.au	Pegasus Networks	pegasus.oz.au
Collaborative Information Technology Research Institute	citri.edu.au	Premier and Cabinet Information Technology, (Tas)	tas.gov.au
Comalco Research Centre	cra.com.au	Pyramid Technology Corporation	pyramid.com.au
Computer Education Centre	cec.edu.au	Queensland Institute of Medical Research	qimr.edu.au
Computer Power Software	cpg.oz.au	Rural Water Corporation, Victoria	rwc.org.au
Concept Personnel Services	concept.com.au	South Australian Department of Employment and TAFE	tafe.sa.edu.au
connect.com.au	connect.com.au	South Australian Department of Environment & Planning	dep.sa.gov.au
Cooperative Library Association Network of NSW	clann.edu.au	South Australian Dept of Lands	lands.sa.gov.au
Corinthian Engineering	ci.com.au	South Australian Dept of Roads	roads.sa.gov.au
CRA - Advanced Technical Development	atd.craatd.telememo.au	State Library of NSW	slnsw.gov.au
Cray Communications Limited	craycomm.com.au	State Library of South Australia	slsa.sa.gov.au
CSR Limited	csr.com.au	State Library of Victoria	statelib.vic.gov.au
DA Books and Journals	dadirect.com.au	State Systems (South Australian Computing Service)	systems.sa.gov.au
Defence Science and Technology Organisation	dsto.gov.au	Sugar Research Institute	sri.org.au
Department of Arts, Sport, Environment & Territories	daset.gov.au	Sydney Cricket and Sports Ground Trust	scgt.oz.au
Department of Ed, Vocational Education and Training	deet.gov.au	Technology Development Corporation	tpa.com.au
Department of Primary Industry and Energy	dpie.gov.au	Technology One Pty Ltd	techone.com.au
		Techway Networks	techway.com.au
		Telecom - Network and Workforce Systems	telstra.com.au
		Telecom Australia Research Laboratories	trl.oz.au, telecom.com.au
		The Miden Group	miden.com.au

CONNECTING TO AARNET

UniSuper	unisuper.com.au	West Australian Petroluem Pty Ltd	wapet.com.au
University Co-operative Bookshop	coop-bookshop.com.au	Western Australia Tertiary Admissions Centre	tisc.edu.au
VIC Roads	vicroads.gov.au	Western Australian Department of Agriculture	agriv.wa.gov.au
WAEPA - WA Environmental Protection Authority	waepa.gov.au	Walter Eliza Hall Institute	wehi.edu.au
Wellfleet Communications (Aus)	wellfleet.com.au		

A.4 MAIL AFFILIATE MEMBERS

21C Magazine	21c.com.au	Australian Public Access Network Association (NSW)	apana.org.au
3COM ANZA	anza.3com.com.au	Australian Public Access Network Association (SA)	apana.org.au
Aberfoyle Resources Limited	aberfoyle.oz.au	Australian Public Access Network Association (Vic)	apana.org.au
ACT Electricity and Water	actew.oz.au	Australian Radiation Laboratory	arl.oz.au
ACT Government Computing Service	agcs.act.gov.au	Australian Software Innovations	asi.oz.au
Adacel Pty Ltd	adacel.com.au	Australian Technology Resources	atr.com.au
Adept Software	adept.oz.au	Australian Tourism Commission	atc.gov.au
Alcatel Australia	alcatel.oz.au	Australian Unix Users Group (AUUG) Inc	auug.oz.au
All Graphics Industries	agi.oz.au	AUUG Inc (Canberra Chapter)	canb.auug.org.au
Amalgamated Televisions Services (ATN Channel 7 Sydney)	atn7.oz.au	AUUG Inc.	auug.oz.au
Amcor Research and Technology Centre	amcor.com.au	Avid Systems	avid.oz.au
AMP Investments Australia	amp.com.au	Bay Technologies Pty Ltd	baytech.com.au
Ampac Life Ltd	ampac.com.au	Beta Communications Pty Ltd	betacomm.com.au
Andersen Consulting	andersen.com.au	BHA Computer	bha.oz.au
Angle Park Computing Centre	magill.unisa.edu.au	BHP CPD Research and Technology Centre	kembal.oz.au
Animal Logic Research P/l (formerly Discreet Logic)	dl.oz.au	BHP Info Tech - Newcastle	bhpese.oz.au
ANL Limited	anl.oz.au anl.com.au	BICC Communications	bicc.oz.au
Anti-Cancer Council of Victoria	accv.oz.au	Blackwells	blackwell.com.au
ANZ Banking Group - Global Technical Services	anz.com.au	Boral Elevators - Research and Development	jpl.oz.au
Apple Computer Australia	apple.oz.au	Box Hill College of TAFE	bhtafe.edu.au
Apple Users Society of Melbourne	ausom.oz.au	Brisbane Public Access Network	brisnet.org.au
Applied Micro Systems Australia Pty Ltd	ams.oz.au	British Aerospace Australia	baea.com.au
Aquasoft Pty Ltd	aquasoft.com.au	Burdett Buckeridge and Young	bby.oz.au
Assorted C Software	asstdc.oz.au	Byrne and Davidson	bnd.oz.au
Atlantek	atlantek.oz.au	Byrne and Davidson (B&D Australia)	bnd.oz.au
Ausonics	ausonics.oz.au	Bytecraft	bytecraft.oz.au
Auspex Systems	auspex.oz.au	Cadcom Solutions Pty Ltd	cad.com.au
Austek Microsystems	austek.oz.au	Cadence Design Systems	cadence.oz.au
Austel	austel.gov.au	Canon Information Systems Research Australia	canon.oz.au
Australia Post	auspost.com.au	CEA Technologies Pty Ltd	cea.com.au
Australian Airlines	ausair.com.au	Chancery Lane Computing Services	clcs.com.au
Australian Centre for Unisys Software	syacus.acus.oz.au	Chiron Mimotopes	chiron.com.au
Australian Computer Society	acs.org.au	Cisco Systems	cisco.com.au
Australian Council for Educational Research	acer.edu.au	CITEC	citec.oz.au
Australian Defence Industries	adied.oz.au	Classified Computers	ccpl.oz.au
Australian Information Processing Centre	aipc.oz.au	Clegg Driscoll Consultants Pty Ltd	clegg.oz.au
Australian Institute of Criminology	crime.oz.au	Co-Cam Computer Group	cocam.oz.au
Australian Institute of Family Studies	aifs.org.au	Coal & Allied Operations	canda.com.au
Australian Mineral Industries Research Association	amira.com.au	Codex Software Development	codex.oz.au
Australian National Botanical Gardens	anbg.gov.au	Coles Myer Ltd	colesmyer.oz.au
Australian National Parks and Wildlife Service	anpws.gov.au	Coles Supermarkets	coles.storesys.oz.au
		Colonial Mutual	cmutual.com.au
		Colonial Mutual (cancelled 92aug)	cmutual.com.au
		Com Net Solutions	comnet.oz.au
		Com Tech Communications	ct.oz.au

CONNECTING TO AARNET

Commercial Dynamics	cdsyd.oz.au	ETX Consultants	etx.oz.au
Common Sense Computing Pty Ltd	csc.com.au	Expert Solutions	bacchus.oz.au
Communica Software Consultants	communica.oz.au	FAI Life Insurance	failife.com.au
Communications Development and Research	cdr.oz.au	Faxon Australian Pty Ltd	faxon.com.au
Comperex Systems	cx.oz.au	Federated Australian University Staff Association	fausa.org.au
Computechnics	ctechnics.oz.au	FGH Decision Support Systems	fgh.oz.au
Computer Law Corporation	lawnet.com.au	Fire Fighting Services (James Hardie Building Service & Technology Pty Ltd)	ffe.com.au
Computer Sciences of Australia	csasyd.oz.au	First State Computing	firststate.com.au
Computer Systems (Australia)	compsys.com.au	Frank McCulloch	xacom.oz.au
Convergent Systems Inc	conv.com.au	Fujitsu Australia	faloz.oz.au
Copper Refineries Pty Ltd	crl.oz.au	Fulcrum Consulting Group	fulcrum.oz.au
Corinthian Engineering	ci.oz.au	Functional Software	fs.com.au
Corporate WorkGroup Resources	cwr.oz.au	GCS Pty Ltd	gcs.oz.au
County Natwest Australia	county.oz.au	GEC Alsthom Australia Ltd	gecaa.com.au
Cray Communications Limited	craycomm.com.au	GEC Marconi Systems Ltd	gecms.com.au
CSL Limited	csl.com.au	Gemco (Incorporating Aldetec)	aldetec.oz.au
Cyberdyne Systems Corporation	cyberdyne.oz.au	Genasys II	genasys.com.au
Cybersource Pty Ltd	cyber.com.au	General Automation Pty Ltd	pdev.ga.oz.au
Dames & Moore	dm.com.au	Gestetner Pty Ltd	gestetner.oz.au
Data General Australia	dg.oz.au	Graphics Computer Systems	gcs.oz.au
Datacraft Technologies	datacraft.oz.au	Gunnedah Abattoir	gsc.oz.au
David Burren	eyrie.oz.au	Halcyon Communications Pty Ltd	halcyon.com.au
David Nugent	csource.oz.au	Haltek	haltek.com.au
Dawn Technologies/Auspex Systems	auspex.oz.au	HCS Australia	hcsaust.com.au
DB Bain Group Pty Ltd	bain.oz.au	Headquarters, Australian Defence Force	hqadf.defence.gov.au
DECUS Australia	decus.com.au	Health Computing Services Australia	hcsaust.com.au
Defence Housing Authority	dha.gov.au	Hedland College	hedland.edu.au
Defence Service Homes	dsh.oz.au	Highland Logic	highland.oz.au
Department of Administrative Services (ITSB)	itsb.das.gov.au	Honeywell	honeywell.oz.au
Department of Agriculture and Rural Affairs, Victoria	agvic.gov.au	Hong Kong Jockey Club Systems	hkjcs.oz.au
Department of Defence - Canberra	defcen.gov.au	Hydrographic Sciences Australia Pty Ltd	hsa.com.au
Department of Employment Education and Training (HED)	hed.deet.gov.au	Iconix	iconix.oz.au
Department of Energy and Minerals	mines.vic.gov.au	Ideas International Pty Ltd	ideas.oz.au
Department of Health, Housing and Community Services	hhcs.gov.au	Industry Commission	ic.gov.au
Department of of Primary Industries (Qld)	qdpi.oz.au	Information Technology Branch, Department of Education and Arts	deet.gov.au
Department of Planning and Housing	dph.vic.gov.au	Information Technology Consultants	itc.oz.au
Dept Conservation and Natural Resources	dce.vic.gov.au	Informed Technology	it.com.au
Dept of Conservation and Environment - Victoria	mssl.oz.au	Insession	signum.oz.au
Dept of Premier and Cabinet - Victoria	vicgov.oz.au	Integration Design Pty Ltd	id.oz.au
Derpartment of Industry, Technology and Regional Development	dit.gov.au	International Development Program	idp.edu.au
Desmond Fitzgerald and Associates	dfa.com.au	International Railroad Systems (Elec)	augean.ua.oz.au
Dialix	dialix.oz.au	Internode Systems	internode.com.au
Digital Arts Research	dafat.com.au	Ipec Metro Distribution	main.com.au
Disability Information and Resource Centre	dirc.sa.org.au	Irvine & Associates	irvine.com.au
DNA Pty Ltd	dna.com.au	Ivanhoe Grammar School	igs.edu.au
Duesburys Information Technology	dues.oz.au	J. A. Almgren	jna.com.au
Duke Australia	duke.com.au	Jeffrey Coleman	mana.oz.au
EASAMS (Australia)	can.easams.com.au	Jeffrey Smartt	smartt.oz.au
Eastek	eastek.oz.au	JTEC Pty Ltd	jtec.com.au
Electronic Financial Services	efs.com.au	Kel Aerospace Pty Ltd	kel.com.au
Emergent Technology	emergent.com.au	Knowledge Engineering	ke.oz.au
Empire Ridge Pty Ltd	empire.com.au	KPMG Solutions	kpmgsol.oz.au
Engineering Computer Services :Pty Ltd	ecsaus.oz.au	Labtam Pty Ltd	labtam.oz.au
Equinet	equinet.oz.au	Lancorp	lancorp.oz.au
ESRI Australia	esri.oz.au	Land and Water Resources R&D Corporation	lwrrdc.lwdc.gov.au
		Land Information Centre	mswlands.oz.au
		Leeds and Northrup Australia	lna.oz.au
		Liquor Administration Board of NSW	lab.oz.au

CONNECTING TO AARNET

Lochard Environment Systems	lochard.com.au	Prospect Electricity	pcc.oz.au
Lyons Computer Pty Ltd	lyons.oz.au	Public Works Department, NSW	pwd.nsw.gov.au
Mayne Nickless Courier Systems	mayne.oz.au	QANTAS Information Technology	qantek.com.au
Media Solutions	mediasolutions.com.au	QH Tours	qhtours.oz.au
Megadata Pty Ltd	megadata.mega.oz.au	QH Tours	qhtours.oz.au
Memtec Limited	memtec.com.au	QPSX Communications	qpsx.oz.au
Menhennitt Consultants	mencon.oz.au	Quality Bakers Australia	qba.oz.au
Mentor Technologies	mg.oz.au	Quality By Design	qbd.oz.au
Mercurius Secretariat Pty Ltd	mercurius.com.au	Queensland Department of Education	qed.qld.gov.au
Metal Trades Industry Association of Australia	mtia.oz.au	Recognition Equipment	plexus.oz.au
		Release 4	release4.oz.au
MHS Systems Pty Ltd	mhs.oz.au	Rinbina Pty Ltd	esi.com.au
Michael O'Brien	step.oz.au	Road User Research P/L	roaduser.com.au
MINCOM Pty Ltd	mincom.oz.au	Roads and Traffic Authority - NSW	rta.oz.au
Mitsubishi Motors Australia Ltd	mmal.oz.au	Rocket Media Pty Ltd	rocket.com.au
Mitsui Computer	mitsui.oz.au	Royal Automobile Club of Victoria (RACV)	racv.com.au
MLC Computer P/L	mlc.com.au		
Moldflow Pty Ltd	moldflow.oz.au	RUNX Unix Timeshare	runx.oz.au
Moss Products Pty Ltd	moss.oz.au	Saki Computer Services Pty Ltd	saki.com.au
MPA International Pty Ltd	mpa.oz.au	Santos Pty Ltd	santos.com.au
Multibase	yartou.oz.au	Satellite Telemetry Pty Ltd	sattel.com.au
Multiline BBS	multiline.com.au	SBC Dominguez Barry	dbsm.oz.au
Multiuser Solutions Pty Ltd	mus.com.au	Scitec Communications Systems	scitec.com.au
Murray-Darling Basin Commission	mdbc.gov.au	Silicon Graphics Pty Ltd	sgi.oz.au
National Languages and Literacy Institute of Australia	nllia.oz.au	Snowy Mountains Authority	smhea.sma.oz.au
		Softpac	softpac.com.au
Naval Support Command	navy.gov.au	Softsmiths Pty Ltd	softsmiths.oz.au
NEC Australia	nec.com.au	Softway Pty Ltd	sw.oz.au
NEC Information Systems Australia	necisa.oz.au	Solbourne Computer Pty Ltd	solbourne.com.au
Network Solutions	nsg.com.au	Sony Australia	sony.oz.au
New South Wales Fisheries	fisheries.nsw.gov.au	Sony Technology Centre of Australia	tca.oz.au
Niche Publishing Pty Ltd	niche.com.au	South East Queensland Electricity Board	seqeb.gov.au
Nick Andrew	zeta.org.au fido.oz.au	Southern Sky Magazine	southsky.com.au
NorTel Technology Centre	nortel.com.au	St Gregory's Armenian School	nareg.oz.au
NSW Dept of Agriculture	nswda.oz.au	Stallion Technologies	stallion.oz.au
Objectif Pty Ltd	objectif.oz.au	State Bank of NSW	state.com.au
Ochre Development	ochre.oz.au	State Electricity Commission of Victoria	secvic.oz.au
Office of the Supervising Scientist	oss.gov.au	Sterling Software	sdipl.sdi.oz.au
Opal Communications Systems Pty Ltd	opal.com.au	Stockdale Prospecting	spl.oz.au
Open Connexion Pty Ltd	connexion.com.au	Stone Microsystems	stone.oz.au
Open Software Associates	osa.oz.au	Systech	systech.com.au
Open Vision Systems	ovs.com.au	System Builder Development	sbdev.oz.au
OpSys Can Unit - Telecom	canocc.telecom.oz.au	System Services Pty Ltd	syserv.com.au
Optimization Software Engineering	optimization.com.au	Tandem Australia	tandem.com.au
Osix	osix.oz.au	Tasmanian Department of Education and the Arts	tased.oz.au
OTC Australia Research & Development	otca.oz.au		
OZ-E-MAIL	ozemail.com.au	Tech Pacific	techpac.oz.au
Ozware Developments Pty Ltd	ozware.com.au	Technix Consulting Services	technix.oz.au
P.T. Freeport Indonesia Company	fmi.com.au	Telecom - Customer Network Management Branch	cnm.telecom_cnm.oz.au
Pacific Semiconductors Pty Ltd	pacsemi.oz.au		
Pacific Star Communications	pacstar.com.au	Telecom - Network Management and Protocol Development	nmpd.oz.au
Patrol Software Pty Ltd	patrol.com.au		
Paul Fitchett & Associates	pfa.com.au	Telecom - Network Management Systems Development	nimrod.ta.oz.au
Paul Williams	css.oz.au		
Petrosys Pty Ltd	petrosys.oz.au	Telecom - Network Management Systems Development	nmmdmelb.telecom.oz.au
Philips PTS	philips.oz.au		
Phillips Australian Oil Company	paoc.com.au	Telecom Australia - Commercial and Consumer	qld.cnc.oz.au
Platform Technologies	platform.oz.au		
Port of Melbourne Authority	pma.vic.gov.au	Telecom Australia - Integrated Network Products Unit	card.inpu.oz.au
Praxa Ltd	praxa.com.au		
Prentice-Hall	prenhall.com.au	Telecom Australia, WA	telecomwa.oz.au
Pro Medicus	promed.com.au	Telecom Network Engineering and Support Services	tne.oz.au
Process Software Solutions	process.oz.au		

CONNECTING TO AARNET

Telecom Network Operations, Systems Development	mnsd.oz.au
Telectronics Pty Ltd	tpl.oz.au
Telerate (Sydney)	telerate.com.au
The Armidale School	as.edu.au
The Australian Museum	austmus.oz.au
The Preston Group	tpg.oz.au
The Southport School	tss.oz.au
TimeKeeping (SA) Pty Ltd	tksa.com.au
Timekeeping Victoria Pty Ltd	tkv.com.au
Toshiba International	csdc.toshiba.oz.au
Tower Technology	towers.oz.au
TurboSoft	abccomp.oz.au
Tusc Computer Systems	tusc.oz.au
Tyndall Australia Limited	tyndall.com.au
UBI Consulting Pty Ltd	ubi.com.au ubi.com.au
Unidata	unidata.oz.au
UniLink Data Systems Pty Ltd	unilink.com.au
Uninet Consulting Pty Ltd	uninet.com.au
Uniq Professional Services	uniq.com.au
University of Melbourne Student Union Inc	mac.unimelb.edu.au
Unixpac	unixpac.oz.au
Vanoco	van.oz.au
Vibro Acoustic Sciences	vibro.oz.au
VME Systems	vme.oz.au
VSM Software Services	vsm.com.au
Wacher Pty Ltd	wacher.oz.au
Webster Computer Corporation	wcc.oz.au
Welkin Software	welkin.com.au
Wesley College	wesley.oz.au
Western Australia Police Department	wapol.gov.au
Whitesmiths Australia Pty Ltd	wsa.oz.au
Wizard Software Company	web.com.au
Wollongong City Council	wcc.nsw.gov.au
Woodside Petroleum	woodside.com.au
WorkNet	worknet.com.au
Workstations Plus	plus.oz.au
Wormald Electronics	wormald.com.au
Zircon Systems	zircon.oz.au
Zircon Systems	zircon.oz.au

Appendix B. AARNet User and Resource Guides

Appendix B.1. User Guide - Booklet

AARNET USER GUIDE ORDER FORM

Copies of the AARNet User Guide, *Getting the Most out of AARNet*, are available from the Queensland University of Technology at a cost of \$7.50 each including postage and packing. The guide, a 50 page booklet, comes in two versions; one tailored for VAX/VMS, the other for UNIX systems. Orders should indicate the number of each versions required, and payment must be included. Make cheques payable to "Computing Services, QUT".

AARNet User Guide
C/- Computing Services,
(Ground floor, V Block)
Queensland University of Technology
GPO Box 2434
Brisbane QLD 4001

Please send me _____ copies of the Vax/VMS version.

Please send me _____ copies of the Unix version.

I enclose a cheque / money order for \$_____

(\$7.50 per copy including postage and packing.)

Please deliver to:

Appendix B.2. User Guide - On-Line

The AARNet User Guide is available on-line via anonymous FTP from system `aarnet.edu.au` in text (.txt), PostScript (.ps) or WordPerfect (.wp) format. The files are in the directory `/pub/user-guide`.

```
aarnet_user_guide_unix_v1.ps      (or .txt or .wp)
aarnet_user_guide_updates.txt     (or .wp)
aarnet_user_guide_vms_v1.ps      (or .txt or .wp)
```

Note that WordPerfect files should be retrieved in FTP binary mode.

Files can also be obtained via electronic mail (uuencoded and split into small parts). Send a message to `ftpmail@cs.uow.edu.au` with (for example) the following commands in the body of the message:

```
HOST aarnet.edu.au
GET pub/user-guide/aarnet_user_guide_unix_v1.ps
```

Appendix C. Mail Gateway Services and Providers

C.1. Mail Gateway Services

MHSnet

MHSnet is the commercial implementation of the networking protocols used to implement ACSnet. It runs only on UNIX systems and is widely deployed within Australia. MHSnet provides a number of other facilities in addition to electronic mail, including file transfer and the transmission of network news between other systems running MHSnet. Available from:

Message Handling Systems Pty. Ltd.
1st Floor
2 King Street
Newtown
NSW
Tel: (02) 550 4448
Fax: (02) 519 2551

UUCP

The UNIX to UNIX Copy Program (UUCP) is part of every UNIX system and versions are also available for just about every other system including PCs and Macintoshes. Some commercial E-mail packages also support UUCP connections to gateways and servers.

PHONEnet

PHONEnet is a dial-up protocol for use with Digital's VAX/VMS systems. It is distributed with the excellent PMDF mail package which is available from:

The PMDF Project
Information Technology Services
Thomas Cherry Building
University of Melbourne
Information Technology Services
Parkville
VIC
Tel: (03) 3447045
Fax: (03) 347 4803

SLIP (Serial Line TCP/IP Protocol)

By using TCP/IP to establish a point to point connection to the gateway system, a link can be engineered at relatively low cost since TCP/IP is incorporated into all UNIX systems and is available from the public domain for PC and Macintosh systems. Many commercial E-mail packages support links to TCP/IP networks via TCP/IP's Simple Mail Transfer Protocol (SMTP).

Dial-Up Terminals

As an alternative to transferring electronic mail to the mail affiliate's site, some gateway service providers offer clients the option of establishing a dial-up terminal session with the gateway system itself to read and compose messages.

Other Services

Some gateway service providers provide, or may be able to provide, links to proprietary electronic mail systems such as QuickMail.

CONNECTING TO AARNET

C.2. Mail Gateway Service Providers

National

connect.com.au
Hugh Irvine
Tel: 03 528 2239
Fax: 03 528 5887
Services: UUCP, MHSnet, SLIP/PPP

Victoria

University of Melbourne
Information Technology Services
Chris Chaundy
Tel: 03 344 7045
Fax: 03 347 4803
Services: Phonenet, SLIP

Deakin University
Computing & Communications Services
Craig Warren
Tel: (052) 27 1163
Fax: (052) 27 2010
Services: UUCP, ACSnet, SLIP, PPP

Royal Melbourne Institute of Technology
Val Freda
Tel: 03 660 3816
Fax: 03 660 2499

Latrobe University
Computing Services Department
Phil County,
Tel: 03 479 1500
Fax: 03 479 1999

New South Wales

University of Sydney
University Computer Centre
David Phipps
Tel: (02) 692 3491
Fax: (02) 660 6557
Services: MHSnet, SLIP

University of Wollongong
Department of Computer Science
Tel: 042 21 3810
Fax: 042 21 3262

University of New South Wales
Computing Services Department
Phil Cohen
Tel: 02 697 2926
Fax: 02 662 8665

University of Newcastle
David Morrison
Tel: 049 215397
Fax: 049 216910

University of Technology, Sydney
Communication Services Branch
Anne Mikita or Patrick Herlihy
Tel: 02 330 1990
Fax: 02 330 1994
Services: MHSnet, SLIP, Quickmail

Message Handling Systems Pty Ltd
Elaine Pensabene
Tel: (02) 550 4448
Fax: (02) 519 2551
Services: MHSnet, UUCP

Queensland

University of Queensland
Prentice Centre
Wilber Williams
Tel: 07 365 4232
Fax: 07 365 4477
Services: MHSnet, SLIP, UUCP

Griffith University
Information Technology Services
Greg Watson
Tel: 07 875 7468
Fax: 07 875 7877
Services: SLIP/PPP, Dial-up

Queensland University of Technology
Computing Services
Alan Agnew
Tel: 07 864 1703
Fax: 07 864 1808
Services: Phonenet

Bond University
Computer Services
Chris Christoff
Tel: 075 95 1424
Fax: 075 95 1456
Services: MHSnet

CONNECTING TO AARNET

South Australia

University of Adelaide
Mark Prior
Tel: 08 228 5680
Fax: 08 223 6245
Services: MHSnet, Phonenet, SLIP/PPP, Dial-up

University of South Australia
Rollo Ross
Tel: 08 302 3158
Fax: 08 302 3385
Services: Phonenet, Dial-up

Western Australia

University of Western Australia
Winthrop Technology
Julie Stout
Tel: 09 380 3505
Fax: 09 382 1688
Services: MHSnet, UUCP, SLIP, Dial-up

Tasmania

University of Tasmania
Computer Centre
Ray Jones
Tel: 002 202792
Fax: 002 231772

Australian Capital Territory

Australian Defence Force Academy
Computer Centre
Geoff Collin
Tel: 06 268 8142
Fax: 06 268 8150
Services: MHSnet, UUCP

Appendix D. AARNet Regional Hub Contacts

Victoria

Melbourne University
Chris Chaundy
Technical Manager, Networks
Ph: (03) 344 7045
FAX: (03) 347 4803

New South Wales

University of Sydney
David Phipps
Communications Manager
Ph: (02) 692 3491
FAX: (02) 660 6557

Queensland

University of Queensland
Max Norris
Network Manager
Ph: (07) 365 3942
FAX: (07) 365 4477

Western Australia

University of Western Australia
Toivo Pedaste
Ph: (09) 380 2605
FAX: (09) 380 1014

Australian Capital Territory

Australian National University
David Jones
Head, Network and Technical Services
Ph: (06) 249 3436
FAX: (06) 247 3425

South Australia

University of Adelaide
Geoff Cleave
Ph: (08) 228 5099
FAX: (08) 223 6245

Tasmania

University of Tasmania
Rod Bilson
Networks Supervisor
Ph: (002) 20 2811
FAX: (002) 23 1772

Northern Territory

Northern Territory University
Terry Anstey
Deputy Director, Computer Services
Ph: (089) 46 6108
FAX: (089) 27 0612

Attachment 1

AARNet Mail Affiliate Membership Application

Organisations wishing to establish an electronic mail service to AARNet should complete this application and return it to the Australian Vice-Chancellors Committee at the address shown below. If you wish to be invoiced for the AARNet Mail Affiliate Membership fee, please indicate this on your application.

Note that there may be other charges associated with an electronic mail service as well as the AARNet Mail Affiliate Membership fee. These may include any charges levied by the organisation which provides the actual gateway service between AARNet and the mail system you are using, and charges payable to a telecommunications carrier for your physical link to the gateway service provider. Such charges are not within the discretion of AARNet.

Provision of the Mail Service to the Affiliate Member is conditional on the continued use of the service as per the Acceptable Use Policy of AARNet, contained in this document.

How to Apply for Mail Affiliate Membership of AARNet

To register as a Mail Affiliate Member of AARNet, please complete a printed copy of this form, and return the form to:

**Customer Services,
AARNet,
PO Box 1142,
Canberra ACT 2601.**

You will be notified by mail when your application has been processed and the required Mail Exchange records for your organisation have been entered into the domain name system.

The Australian Academic and Research Network
1-5 Geils Court, Deakin, ACT
GPO Box 1142, CANBERRA, ACT 2601
E-mail: AARNet@AARNet.edu.au
Telephone: (06) 249 3385
Fax: (06) 249 1369

Australian Vice-Chancellors Committee
ACN 008 502 930

AARNet Acceptable Use Policy

1. *[AARNet]*

The Australian Academic and Research Network (AARNet) is a private telecommunications network owned and operated by the Australian Vice-Chancellors' Committee (AVCC) as a common service provided to the member institutions of the AVCC, the Commonwealth Scientific and Industrial Research Organisation (CSIRO) and other organisations who have entered into an AARNet Affiliate Membership Agreement with the AVCC.
2. *[Function of AARNet]*

The primary function of AARNet is to support the open academic and research enterprise of Australia.
3. *[Inclusion of Support Services as Acceptable Use]*

The use of AARNet in connection with the administrative activities of the member organisations of AARNet, use for institutional support services, and use in support of the various infrastructural service support activities associated with the academic and research enterprise, are considered to be acceptable use of AARNet.
4. *[Affiliate Membership]*

The use of AARNet by Australian organisations other than AVCC member institutions and CSIRO divisions requires the execution of an AARNet Affiliate Membership Agreement between the organisation and the AVCC. The AVCC reserves the right to reject or approve each application for Affiliate Membership, which it will consider on a case by case basis.
5. *[Academic and Research and Commercial Traffic as Acceptable]*

Traffic relating to the national academic and research enterprise and related commercial traffic is considered acceptable use of AARNet.
6. *[Peer Traffic as Acceptable]*

Traffic originating from peer academic and research networks which conforms to the originating networks' acceptable use policies is in general considered acceptable traffic to be carried over AARNet.
7. *[Illegal Use as Unacceptable]*

AARNet facilities cannot be used by any individual or group of persons for any activities of an illegal or fraudulent nature, including any activities prohibited under the Australian Commonwealth Government Telecommunications Act 1989, or under other applicable Australian State and Commonwealth laws.

Additional Notes

- A.1 The AVCC, in conjunction with CSIRO, and through the AARNet Advisory Board, will be the sole arbiter of which categories of material are considered acceptable or unacceptable for transmission over AARNet;
- A.2 Notwithstanding Note A.1, permission may be granted for a member to use AARNet to transmit or receive specific materials not normally considered acceptable if the request to transmit such materials is supported by the Chief Executive Officer of the institution requesting such materials; and
- A.3 If material which the AVCC has resolved should not be transmitted is found to have been transmitted over AARNet, then the membership of the network member responsible for that action may be terminated by the AVCC.

AARNet Mail Affiliate Membership Application

On behalf of the organisation listed below I wish to apply to be a Mail Affiliate Member of AARNet, and accordingly request that the Australian Vice-Chancellors' Committee maintain on my behalf an electronic mail delivery record in the Australian Academic and Research Network (AARNet) to allow my organisation to send and receive electronic mail carried across AARNet.

I understand that I will be invoiced for this service for subsequent years unless the AVCC receives my organisation's written advice to terminate the Affiliate Membership of AARNet. I understand that the AVCC has the right to obtain mail logs maintained by the organisation providing my organisation's mail gateway service to confirm that the volume of mail transmitted/received by my organisation is within the limits specified in the schedule or charges below.

I understand that the AVCC maintains the right to vary the Mail Affiliate Membership charges from year to year, and maintains the right to cease offering this service to my organisation at the start of any year, at their discretion. I understand that in the event of any variation of the Mail Affiliate Membership of AARNet, my organisation will be advised in writing by the AVCC to the address below.

I understand that in consideration of the AARNet Mail Affiliate Membership charge, AARNet will undertake to maintain a mail directory entry which will direct incoming electronic mail to the gateway system(s) which I have nominated below. Furthermore I accept that there is no other undertaking made by AARNet in terms of reliability of mail delivery or any other form of undertaking by AARNet or the AVCC in consideration of the payment to AARNet for the maintenance of the mail directory entry on AARNet.

My organisation undertakes to use AARNet facilities within the terms and conditions stated in the AARNet Acceptable Use Policy. I accept the right of the AVCC to immediately terminate this service at their discretion if these undertakings are abused by my organisation (where the AVCC retains the right to determine what constitutes such abuse).

Date: _____

Name of Organisation: _____

Signed: _____

Name: _____ Position: _____

on behalf of the organisation named above.

Address: _____

Administrative/Billing

Contact Name: _____

Address: _____

E-Mail: _____ Phone: _____ Fax: _____

AARNet Technical Contact

Technical Contact: _____

E-Mail: _____ Phone: _____ Fax: _____

Mail Delivery Information to be entered in AARNet (see note A)

Preferred Electronic Mail Address: _____

Mail Gateway(s) in order of preference: _____

Membership Fee

Please indicate which category of membership you are applying for and whether you wish to be invoiced for for the membership fee.

	Category	AARNet Charge
<input type="checkbox"/>	1	\$1,000.00 per annum
<input type="checkbox"/>	2	\$4,000.00 per annum

- Payment is attached

- Please send invoice

••••

Note A. Mail Delivery Information

Two items of information are required: firstly your organisation's preferred electronic mail address in Internet domain format, and secondly the name (or names) of gateway systems who will accept electronic mail from AARNet (and connected overseas networks) on your behalf and forward it to you.

Example Mail Delivery information request:

```
Mail addresses required:      acme.oz.au, *.acme.oz.au
Mail Gateways (primary)      gw.somewhere.edu.au
                             (secondary) munnari.oz.au
```

Note that the entry *.acme.oz.au acts as a wildcard entry to direct mail for <any-name>.acme.oz.au to your mail gateway site. It is necessary only if there are many hosts or many electronic mail subdomains within your organisation to which mail needs to be individually addressed.

Electronic mail addresses are generally assigned from within the following Australian domains:

- EDU.AU for educational organisations,
- GOV.AU for government organisations,
- COM.AU for commercial organisations,
- ORG.AU for organisations that do not fit the above categories, and
- OZ.AU for organisations have MHSnet (or ACSnet) connections

It will be necessary to consult your mail gateway service provider to confirm that your preferred electronic mail address is acceptable, and to obtain the name of the gateway system(s).

Attachment 2

AARNet Network Affiliate Membership Application

To establish a Network Affiliate connection to AARNet you will need to complete this form and return it to the AVCC. The AVCC will provide advice as to the available options for a network connection, and liaise with you regarding the termination of the physical link to AARNet.

You will be invoiced by the AVCC for the Affiliate Membership charge, which is of the form of an annual charge, and for the once-off connection fee to AARNet if applicable. Note that there may be other charges associated with the AARNet network connection, including applicable Telecom transmission charges and communications equipment purchase and maintenance charges.

The provision of Network Affiliate Membership is conditional on the continued use of the service as per the Acceptable Use Policy of AARNet reproduced overleaf.

••••

How to Apply for Network Affiliate Membership of AARNet

To register as a Network Affiliate Member of AARNet, please complete a printed copy of this form, and return the form to:

AARNet
PO Box 1142
Canberra ACT 2601

The AVCC will, in the normal course of events, respond with approval to connect, and invoice you for the appropriate charges for connection and the first year's Affiliate Membership fee.

In the following pages are:

AARNet Acceptable Use Policy

Terms and Conditions of AARNet Network Affiliate Membership

Network Affiliate Membership Application

The Australian Academic and Research Network
1-5 Geils Court, Deakin, ACT
GPO Box 1142, CANBERRA, ACT 2601
E-mail: AARNet@AARNet.edu.au
Telephone: (06) 249 3385
Fax: (06) 249 1369

Australian Vice-Chancellors Committee
ACN 008 502 930

AARNet Acceptable Use Policy

1. *[AARNet]*

The Australian Academic and Research Network (AARNet) is a private telecommunications network owned and operated by the Australian Vice-Chancellors' Committee (AVCC) as a common service provided to the member institutions of the AVCC, the Commonwealth Scientific and Industrial Research Organisation (CSIRO) and other organisations who have entered into an AARNet Affiliate Membership Agreement with the AVCC.
2. *[Function of AARNet]*

The primary function of AARNet is to support the open academic and research enterprise of Australia.
3. *[Inclusion of Support Services as Acceptable Use]*

The use of AARNet in connection with the administrative activities of the member organisations of AARNet, use for institutional support services, and use in support of the various infrastructural service support activities associated with the academic and research enterprise, are considered to be acceptable use of AARNet.
4. *[Affiliate Membership]*

The use of AARNet by Australian organisations other than AVCC member institutions and CSIRO divisions requires the execution of an AARNet Affiliate Membership Agreement between the organisation and the AVCC. The AVCC reserves the right to reject or approve each application for Affiliate Membership, which it will consider on a case by case basis.
5. *[Academic and Research and Commercial Traffic as Acceptable]*

Traffic relating to the national academic and research enterprise and related commercial traffic is considered acceptable use of AARNet.
6. *[Peer Traffic as Acceptable]*

Traffic originating from peer academic and research networks which conforms to the originating networks' acceptable use policies is in general considered acceptable traffic to be carried over AARNet.
7. *[Illegal Use as Unacceptable]*

AARNet facilities cannot be used by any individual or group of persons for any activities of an illegal or fraudulent nature, including any activities prohibited under the Australian Commonwealth Government Telecommunications Act 1989, or under other applicable Australian State and Commonwealth laws.

Additional Notes

- A.1 The AVCC, in conjunction with CSIRO, and through the AARNet Advisory Board, will be the sole arbiter of which categories of material are considered acceptable or unacceptable for transmission over AARNet;
- A.2 Notwithstanding Note A.1, permission may be granted for a member to use AARNet to transmit or receive specific materials not normally considered acceptable if the request to transmit such materials is supported by the Chief Executive Officer of the institution requesting such materials; and
- A.3 If material which the AVCC has resolved should not be transmitted is found to have been transmitted over AARNet, then the membership of the network member responsible for that action may be terminated by the AVCC.

Terms and Conditions of AARNet Network Affiliate Membership

Conditions of Use:

1. Use of AARNet must conform to the AARNet Acceptable Use Policy.
2. The AVCC reserves the right to disconnect the Affiliate Member in the event of serious operational or security incidents arising through the Affiliate Member's connection to AARNet, or in the event of continued abuse of these Conditions of Use of AARNet (where the AVCC retains the right to determine what constitutes such abuse).

Network Access:

3. The Affiliate Member cannot, under these terms of AARNet Network Affiliate Membership, allow third party organisations or individuals to gain access to AARNet via the Affiliate Member's facilities. This includes the explicit prevention of access to AARNet by either direct network layer connectivity or indirect gateway services,

Network Connection:

4. The physical connection from the Affiliate Member to AARNet is to be terminated within AARNet at a location which will be nominated by AARNet. In general the Regional Hub sites of AARNet will be nominated as the appropriate connection point to AARNet. The type and characteristics of the physical connection must be approved by AARNet prior to installation of the facility.
5. The physical link must be terminated on the Affiliate member site with equipment which has the prior approval of AARNet as being technically compatible with AARNet equipment.
6. The responsibility for operational support and maintenance of the link, and the equipment associated with the physical link on the Affiliate Member's site is that of the Affiliate Member. The responsibility for the maintenance and operation of equipment on the AARNet site is that of AARNet or its agent.

Disclaimer of Liability:

7. The AVCC, CSIRO, and member institutions of these bodies accept no responsibility for any loss or damage which may be incurred by the Affiliate Member through the connection to AARNet. Affiliate Members connect to AARNet at their own risk.

Affiliate Membership

8. To become an Affiliate Member of AARNet an application must be made in writing to the AVCC. The application must indicate the undertaking of the organisation to abide by these Terms and Conditions of Affiliate Membership.
9. The AVCC reserve the right to refuse a Network Affiliate Membership application at their discretion.
10. The Affiliate Membership shall operate on a year-to-year basis, and requires the advance payment of an annual Affiliate Membership fee to the AVCC.
11. The AVCC maintains the right to vary the Network Affiliate Membership fee from year to year, and maintains the right to cease offering this service at the start of any year at their discretion.
12. The AVCC reserve the right to vary the Affiliate Membership fee for an organisation at their discretion on a case by case basis.

Termination of Affiliate Membership

13. Affiliate membership will be terminated upon the receipt of notice to terminate the Affiliate membership by the AVCC.
14. The AVCC reserves the right not to renew the annual Affiliate Membership for an organisation at their discretion. In such a case notice of intention to terminate Affiliate Membership as of the following year will be delivered to the organisation in writing.
15. Affiliate Membership shall be considered to have terminated in the event of non-payment of the annual Affiliate Membership fees to the AVCC by the invoice due date.
16. In the event of termination of Affiliate Membership of AARNet all connection services will be withdrawn by AARNet.

AARNet Network Affiliate Membership Application

On behalf of the organisation listed below I wish to apply to be a Network Affiliate Member of AARNet.

I have read and understood the Terms and Conditions of Connection to AARNet as stated in this document. I undertake that my organisation will abide by the Terms and Conditions of Affiliate Membership of AARNet.

I understand that following initial approval as an Affiliate Member of AARNet, AARNet will nominate a connection point for a physical connection to AARNet, and provide advice as to the method of connection. A specific proposal will be lodged by the Affiliate Member with AARNet for endorsement before proceeding with the implementation of the connection.

I understand that following evaluation of the connection proposal I will be charged an annual fee based on the effective bandwidth of the connection to AARNet as per the attached charge Schedule, and will also be charged an initial connection fee covering the expenditure of the AVCC in establishing the connection. I understand that I will be invoiced for this service for subsequent years unless the AVCC receives my organisation's written advice to terminate the Affiliate Membership of AARNet.

I understand that the AVCC maintains the right to vary the Network Affiliate Membership charges from year to year, and maintains the right to cease offering this service to my organisation at the start of any year, at their discretion. I understand that in the event of any variation of the Network Affiliate Membership of AARNet, my organisation will be advised in writing by the AVCC to the address below.

I understand that in consideration of the AARNet Network Affiliate Membership charge, AARNet will undertake to maintain a data communications link between AARNet and my organisation. Furthermore I accept that there is no other undertaking made by AARNet in terms of reliability of the network connection service, nor any other form of undertaking by AARNet or the AVCC in consideration of the payment to AARNet for the maintenance of the network connection service to AARNet.

My organisation undertakes to use AARNet facilities within the terms and conditions stated in the AARNet Acceptable Use Policy. I accept the right of the AVCC to immediately terminate this service at their discretion if these undertakings are abused by my organisation (where the AVCC retains the right to determine what constitutes such abuse).

Date: _____

Name of Organisation: _____

Signed: _____

Name: _____ Position: _____

on behalf of the organisation named above.

Address: _____

Administrative/Billing

Contact Name: _____

Address: _____

E-Mail: _____ Phone: _____ Fax: _____

AARNet Technical Contact

Technical Contact: _____

E-Mail: _____ Phone: _____ Fax: _____

AARNet Charges for Network Affiliate Members

The schedule of annual charges is based on effective bandwidth of the physical connection into AARNet. There are no further incremental charges for use of AARNet. The charges for 1993 are as follows:

Annual Access Fees

\$5,000 p.a.	dial in at 9.6Kbps or lower
\$8,000 p.a.	9.6Kbps capacity - leased line or permanent call
\$15,000 p.a.	19.2Kbps capacity - leased line or permanent call
\$25,000 p.a.	48Kbps & 64Kbps capacity
\$30,000 p.a.	128Kbps capacity
\$35,000 p.a.	256Kbps capacity
\$40,000 p.a.	512Kbps capacity
\$45,000 p.a.	1Mbps capacity
\$50,000 p.a.	2Mbps capacity
\$60,000 p.a.	10Mb capacity

Connection Fee

\$3,500	Connection to a single AARNet cisco router interface
\$1,000	Connection to a single AARNet SLIP/PPP server interface

AARNet Connection Information

Preferred AARNet Connection Point: _____

Please indicate your desired connection bandwidth:

- Link Capacity
- 9.6K or less (dial-up)
 - 9.6K
 - 19.2K
 - 48K/64K
 - 128K
 - 256K
 - 512K
 - 1M
 - 2M
 - 10M

Please indicate your preferred connection type:

- Connection AARNet SLIP Server
- AARNet Router
- Other

Please provide details:

- Payment is attached

- Please send invoice

Configuration Details :

Internet network numbers in use by your organisation:

Prefered Internet Domain Name: _____

Attachment 3

AUUG Mail Service Application

Dear Site Administrator,

As you may be aware, the arrangements for mailing to addresses outside Australia (and also to AARNet sites) changed in May 1991. Since then, the University of Melbourne are no longer managing the the administrative details associated with maintaining this service. The AARNet (Australian Academic and Research Network) management has taken over administering the service, and are requiring all ACSnet and similar sites to register with AARNet and pay a fee for continued access to Internet mail services. AARNet have set this fee as \$1000 per annum for most sites, with larger sites paying more (you know who you are).

The fee is intended to cover use of AARNet bandwidth for your network traffic. Registration with AARNet, however, provides ONLY the registration of your address in worldwide address tables - your site will be unreachable without this registration. The fee does NOT cover the costs involved in obtaining a connection to AARNet or ACSnet NOR does it include a guarantee that you can be connected or even to help you find a connection point. See Note B for some information about connection services.

AUUG as a service to its members has negotiated with AARNet to achieve a lower price for this basic address registration service. The lower price is based on the reduction in paperwork for the AARNet management authorities. The AUUG/AARNet fee is dependent on the membership status of the owner of the machine(s)/domain involved, and is currently \$250 for members and \$600 for non-members. As such it is a substantial discount on the AARNet fee, but only applies to sites in the AARNet \$1000 category. Larger sites will need to negotiate directly with AARNet.

The address registration is for one AUUG membership year. Membership years start on the 1st of January or July, whichever is nearest to receipt of your application. Sites which do not renew their AUUG/AARNet registration annually with their AUUG membership each year will be removed from the Internet tables and will no longer be able to communicate with international and AARNet hosts. Reminders/invoices will be sent along with your membership renewal.

The required initial registration form is attached below. It should be completed and forwarded to AUUG's (postal) mailing address at the bottom of the form or faxed to (02) 332 4066. If you have any queries on the AUUG/AARNet arrangements please direct them to Liz Fraumann at the AUUG office on (02) 361 5994 (eaf@swift.sw.oz.au) or myself (chris@softway.sw.oz.au).

Regards,
Chris Maltby
Vice-President
AUUG Inc.

AUUG - AARNet Mail Service Affiliate Membership Application

On behalf of the organisation listed below I wish to apply to be a Mail Service Affiliate Member of AARNet, and accordingly request that AUUG Incorporated arrange for the Australian Vice-Chancellors' Committee (AVCC) to maintain on my behalf an electronic mail delivery record in the Australian Academic and Research Network (AARNet) to allow my organisation to send and receive electronic mail carried across AARNet.

I understand that the AVCC may consult the recorded logs of my organisation's usage of AARNet facilities, and determine that I am ineligible for registration under the terms of the agreement between AVCC and AUUG Inc. I understand that AUUG Inc will invoice my organisation for this service for the current calendar year and for subsequent years unless it receives my organisation's written advice to terminate the Affiliate Membership of AARNet.

I understand that the AVCC and AUUG Inc maintain the right to vary the Mail Service Affiliate Membership charges from year to year, and maintains the right to cease offering this service to my organisation at the start of any year, at their discretion. I understand that in the event of any variation of the Mail Service Affiliate Membership of AARNet, my organisation will be advised in writing by the AVCC or AUUG Inc to the address below.

I understand that in consideration of the AARNet Mail Service Affiliate Membership charge, AARNet will undertake to maintain a mail directory entry which will direct incoming electronic mail to the AARNet gateway system(s) which I have nominated below. Furthermore I accept that there is no other undertaking made by AARNet in terms of reliability of mail delivery, maintenance of connection or any other form of undertaking by AARNet or the AVCC in consideration of the payment to AARNet for the maintenance of the mail directory entry on AARNet.

I undertake that my organisation's use of the mail delivery services over AARNet will not be used as a common commercial carrier service between my organisation and other organisations receiving similar services from AARNet, nor will it be used as a commercial carrier service between branches of my organisation. Furthermore my organisation undertakes to use AARNet facilities within the terms and conditions stated in the AARNet Acceptable Use Policy. I accept the right of the AVCC or AUUG Inc to immediately terminate this service at their discretion if these undertakings are abused by my organisation (where the AVCC retains the right to determine what constitutes such abuse).

I understand that a fee is payable with this application: of \$250 if the host/hosts covered are owned by a member of AUUG Incorporated, or \$600 if the host/hosts covered are not owned by an AUUG member. Corporation host owners may only claim the member price if the corporation is an Institutional member of AUUG Inc. My cheque payment of either \$250 or \$600 as appropriate is enclosed with this application.

AUUG - AARNet Mail Service Affiliate Membership Application

Date: _____ PLEASE P_R_I_N_T CLEARLY!

Name of Organisation/Owner: _____

Signed: _____ AUUG Membership No (if known): _____

Name: _____ Title: _____

on behalf of the organisation named above.

Address: _____

_____ Postcode: _____

Administrative Contact: _____ Title: _____

E-Mail: _____ Phone: (____) _____ - _____

Fax: (____) _____ - _____

Technical Contact: _____ Title: _____

E-Mail: _____ Phone: (____) _____ - _____

Fax: (____) _____ - _____

Mail Delivery Information to be entered in AARNet tables (see Note A)

Domain Names Requested: _____

Gateway Addresses: _____

Expected Link Protocol: UUCP SL/IP MHSnet Other: _____

* * * *

Send this page only to:

AUUG Incorporated

PO Box 366 Phone: +61 2 361 5994
Kensington NSW 2033 Fax: +61 2 332 4066

Note A. Mail Delivery Information

Two items of information are required: firstly the preferred name of your mail host (or the domain name(s) of a group of hosts) in Internet domain name system format, and secondly the name (or names) or AARNet gateway systems who will accept electronic mail over AARNet (and connected overseas networks) on your behalf and forward it to you. The primary requirement for an AARNet gateway is its ability to recognise your host/domain addresses and perform the necessary mail header rewriting reliably.

Please check with the postmaster at your preferred AARNet gateway host site before citing them as a gateway for AARNet mail delivery. For ACSnet addresses (*.oz.au), the host "munnari.oz.au" (Melbourne University) is a recommended gateway. Other possible sites include "metro.ucc.su.oz.au" (Sydney University), sirius.ucs.adelaide.edu.au (University of Adelaide), uniwa.uwa.oz.au (University of WA) and bunyip.cc.uq.oz.au (University of Qld). Note that all gateway addresses must be fully domain qualified.

Example Mail Directory information request:

```
Mail addresses required:  acme.oz.au, *.acme.oz.au

Mail Gateways (primary)  gw.somewhere.edu.au
                        (secondary) munnari.oz.au
                        (secondary)  uunet.uu.net
```

The addressability of your site and the willingness of your nominated gateways to act in that capacity will be determined before registration proceeds. Processing will be made faster if you contact the postmaster at your nominated gateways in advance to inform them of your intentions. Your nominated technical contact will be notified by email when registration is complete.

Note B. Getting Connected

New sites will need to find an existing AARNet or ACSnet site who will accept their site as a connection, and also select a protocol for transferring data over their mutual link. Although the UUCP package is a standard inclusion with UNIX, it is little used in Australia due to its relatively poor performance. Other possible choices for your link protocol include SLIP (TCP/IP) and MHSnet.

Among a number of organisations who provide connection services, Message Handling Systems Pty Ltd have announced a special offer on both their link software and connect time for AUUG members. For more details on this offer, contact Message Handling Systems on (02) 550 4448 or elaine@mhs.oz.au.

Attachment 4

IP Network Number Application

Australian Academic and Research Network Network Information Centre

Please complete the following 6 questions and return this form to the AARNet NIC (please type the responses - our handwriting recognition software is not all that good!).

Our contact details are:

Email: hostmaster@aarnet.edu.au

Fax: 06 249 1369

Post: AARNet
GPO Box 1142
Canberra ACT 2601

Phone: 06 249 3385

Network Number Application Form

1. Applicant Details

1a. Organisation Name:

1b. Postal Address:

2. Technical Point of Contact

2a. Name (Lastname, Firstname):

2b. Title:

2c. Postal Address (is as per 1b then enter "as above"):

2d. Email (if known):

2e. Phone:

2f. Fax:

2g. NIC Handle (if known):

3. Network Name for the network (up to 12 characters)

(This is the name that will be used as an identifier for the network in the Network Information Centre database. This is usually a compressed organisation name or similar. This name will have the suffix '-AU' appended to it to avoid clashes with names assigned by other Network Information Centres.)

3a. Network Name:

4. Size of requested Network Address Space

Network number space is being allocated on the basis of demonstrated need. The allocation guidelines applied across the Internet are described in RFC 1466: "Guidelines for Management of IP Address Space", Gerich, May 1993.

These guidelines (in summary) are:

Class A network numbers (24 bits host address space)

No longer allocated

Class B network numbers (16 bits host address space)

Organizations applying for a Class B network number should fulfill the following criteria:

- 1) the organization presents a subnetting plan which documents more than 32 subnets within its organizational network

AND

- 2) the organization has more than 4096 hosts

Organizations applying for a Class B network number must submit an engineering plan that documents its need for a Class B network number. This document must demonstrate that it is unreasonable to engineer its network with a block of class C network numbers. The engineering plan must include how many hosts the network will have within the next 24 months and how many hosts per subnet within the next 24 months.

The submitted engineering plans will be held in strict confidence by the Internet registries and will only be used to judge whether an application is justified. If it is deemed that the applicant's engineering plan, including the number of hosts and subnets, does not warrant a Class B assignment, the applicant will be allocated a block of Class C addresses.

Class C network numbers (units of 8 bits host address space)

By default, if an organization requires more than a single Class C, it will be assigned a bit-wise contiguous block from the Class C space allocated for its geographic region.

The number of allocated contiguous Class C networks is based on the subscriber's 24 month projection of required end system addresses according to the following criteria:

Organization	Assignment
1) requires fewer than 256 addresses	1 class C network
2) requires fewer than 512 addresses	2 contiguous class C networks
3) requires fewer than 1024 addresses	4 contiguous class C networks
4) requires fewer than 2048 addresses	8 contiguous class C networks
5) requires fewer than 4096 addresses	16 contiguous class C networks
6) requires fewer than 8192 addresses	32 contiguous class C networks
7) requires fewer than 16384 addresses	64 contiguous class C networks

Exceptions from the above stated criteria will be determined on a case-by-case basis.

The following questions allow the NIC to determine the most suitable size of address block to allocate to your organisation.

4a. Estimate the number of hosts that will be connected to this network

Initially:

Within 1 year:

Within 2 years:

Within 5 years:

4b. Requested address block (and reasons):

5. Connecting to AARNet

When networks are connected to AARNet we have to arrange for their IP network numbers to be added to the global routing tables. This can take some time. If we know that your network is to be connected in the near future we can begin this process now.

5a. Will the allocated IP address(es) be connected to AARNet

Within 3 months (Y/N):

Within 1 year (Y/N):

6. NIC Profile Information

The following is used for statistical purposes only

Networks are characterized as being either Research, Military, Government, Educational or Commercial. Which type is this network?

6a. Type of network:

(AARNet IP Template \$Date: 1993/09/21 05:44:38 \$)

Attachment 5

Introduction to Administration of an Internet-based Local Network

C R
C S
Computer Science Facilities Group
C I
L S

RUTGERS
The State University of New Jersey
Center for Computers and Information Services
Laboratory for Computer Science Research

3 October 1988

This is an introduction for people who intend to set up or administer a network based on the Internet networking protocols (TCP/IP).

Copyright (C) 1988, Charles L. Hedrick. Anyone may reproduce this document, in whole or in part, provided that: (1) any copy or republication of the entire document must show Rutgers University as the source, and must include this notice; and (2) any other use of this material must reference this manual and Rutgers University, and the fact that the material is copyright by Charles Hedrick and is used by permission.

Unix is a trademark of AT&T Technologies, Inc.

Table of Contents

1. The problem.....	1
1.1 Some comments about terminology	1
2. Routing and Addressing	2
3. Choosing an addressing structure	3
3.1 Should you subdivide your address space?	4
3.2 Subnets vs. multiple network numbers	4
3.3 How to allocate subnet or network numbers.....	5
3.4 Dealing with multiple "virtual" subnets on one network.....	5
3.4.1 Dealing with Multiple Subnets by Turning off Subnetting.....	6
3.4.2 Multiple Subnets: Implications for Broadcasting.....	6
3.5 Choosing an address class	7
3.6 Dialup IP and Micro gateways: Dynamically assigned addresses	7
3.6.1 Dialup IP.....	7
3.6.2 Micro gateways	8
4. Network-wide Services, Naming.....	9
5. Setting up routing for an individual computer.....	11
5.1 How datagrams are routed	12
5.2 Fixed routes.....	14
5.3 Routing redirects.....	14
5.4 Other ways for hosts to find routes	15
5.4.1 Spying on Routing.....	16
5.4.2 Proxy ARP	16
5.4.3 Moving to New Routes After Failures.....	19
6. Bridges and Gateways	21
6.1 Alternative Designs.....	21
6.1.1 A mesh of point to point lines.....	21
6.1.2 Circuit switching technology.....	22
6.1.3 Single-level networks.....	22
6.1.4 Mixed designs.....	23
6.2 An introduction to alternative switching technologies	23
6.2.1 Repeaters	23
6.2.2 Bridges and gateways.....	24
6.2.3 More about bridges	25
6.2.4 More about gateways.....	26
6.3 Comparing the switching technologies.....	26
6.3.1 Isolation	26
6.3.2 Performance.....	27
6.3.3 Routing.....	28
6.3.4 Network management.....	29
6.3.5 A final evaluation.....	30
7. Configuring Gateways.....	31
7.1 Configuring routing for gateways.....	32

This document is intended to help people who are planning to set up a new network based on the Internet protocols, or to administer an existing one. It assumes a basic familiarity with the TCP/IP protocols, particularly the structure of Internet addresses. A companion paper, "Introduction to the Internet Protocols", may provide a convenient introduction. This document does not attempt to replace technical documentation for your specific TCP/IP implementation. Rather, it attempts to give overall background that is not specific to any particular implementation. It is directed specifically at networks of "medium" complexity. That is, it is probably appropriate for a network involving several dozen buildings. Those planning to manage larger networks will need more preparation than you can get by reading this document.

In a number of cases, commands and output from Berkeley Unix are shown. Most computer systems have commands that are similar in function to these. It seemed more useful to give some actual examples than to limit myself to general talk, even if the actual output you see is slightly different.

1. The problem

This document will emphasize primarily "logical" network architecture. There are many documents and articles in the trade press that discuss actual network media, such as Ethernet, Token Ring, etc. What is generally not made clear in these articles is that the choice of network media is generally not all that critical for the overall design of a network. What can be done by the network is generally determined more by the network protocols supported, and the quality of the implementations. In practice, media are normally chosen based on purely pragmatic grounds: what media are supported by the particular types of computer that you have to connect, the distance you have to go, and the logistics of installing various kinds of cable. Generally this means that Ethernet is used for medium-scale systems, Ethernet or a network based on twisted-pair wiring for micro networks, and specialized high-speed networks (typically token ring) for campus-wide backbones, and for local networks involving super-computer and other very high-performance applications.

Thus this document assumes that you have chosen and installed individual networks such as Ethernet or token ring, and your vendor has helped you connect your computers to these network. You are now faced with the interrelated problems of

- configuring the software on your computers
- finding a way to connect individual Ethernets, token rings, etc., to form a single coherent network
- connecting your networks to the outside world

My primary thesis in this document is that these decisions require a bit of advance thought. In fact, most networks need an "architecture". This consists of a way of assigning addresses, a way of doing routing, and various choices about how hosts interact with the network. These decisions need to be made for the entire network, preferably when it is first being installed.

1.1 Some comments about terminology


I am going to use the term "IP" throughout this document to refer to networks designed to carry TCP/IP. IP is the network-level protocol from the Internet (TCP/IP) family of protocols. Thus it is common practice to use the term "IP" when referring to addresses, routing, and other network-layer items. In fact the distinction is not always very clear. So in practice the terms Internet, TCP/IP, and IP may appear to be almost interchangeable.

The terms "packet" and "datagram" are also almost interchangeable. Ideally, "packet" is used for the lowest-level physical unit, whereas "datagram" refers to a unit of data at the level of IP. However these are identical for most media, so people have nearly stopped making the distinction. I have tried to use the terms correctly, even though these days it may sound a bit pedantic. The term "packet" seems to be winning out in common speech. For example, gateway speeds are generally given in "packets per second." I have used the more technically accurate "datagrams per second," since it is really datagrams that are being counted.

I use the term "gateway" where some other authors use "router." "Gateway" is the original Internet term. Unfortunately, the ISO community has begun using the same word with a rather different meaning. People have started using "router" because it doesn't have this ambiguity. I am continuing to use "gateway" because, like the companion Introduction to the Internet Protocols, this document is intended to help you make sense of the Internet specifications. Those specifications use "gateway."

2. Routing and Addressing

Many of the decisions that you need to make in setting up an IP network depend upon routing, so it will be best to give a bit of background on that topic now. I will return to routing in a later section when discussing gateways and bridges. In general, IP datagrams pass through many networks as they are going between the source and destination. Here's a typical example. (Addresses used in the examples are taken from Rutgers University.)



This diagram shows three normal computer systems, two gateways, and three networks. The networks might be Ethernets, token rings, or any other sort. Network 2 could even be a single point to point line connecting gateways R and S.

Note that computer A can send datagrams to computer B directly, using network 1. However it can't reach computer C directly, since they aren't on the same network. There are several ways to connect separate networks. This diagram assumes that gateways are used. (In a later section, we'll look at alternatives.) In this case, datagrams going between A and C must be sent through gateway R, network 2, and gateway S. Every computer that uses TCP/IP needs appropriate information and algorithms to allow it to know when datagrams must be sent through a gateway, and to choose an appropriate gateway.

Routing is very closely tied to the choice of addresses. Note that the address of each computer begins with the number of the network that it's attached to. Thus 128.6.4.2 and 128.6.4.3 are both on network 128.6.4. Next, notice that gateways, whose job is to connect networks, have an address on each of those networks. For example, gateway R connects networks 128.6.4 and 128.6.21. Its connection to network 128.6.4 has the address 128.6.4.1. Its connection to network 128.6.21 has the address 128.6.21.2.

Because of this association between addresses and networks, routing decisions can be based strictly on the network number of the destination. Here's what the routing information for computer A might look like:

network	gateway	metric
128.6.4	none	0
128.6.21	128.6.4.1	1
128.121	128.6.4.1	2

From this table, computer A can tell that datagrams for computers on network 128.6.4 can be sent directly, and datagrams for computers on networks 128.6.21 and 128.121 need to be sent to gateway R for forwarding. The "metric" is used by some routing algorithms as a measure of how far away the destination is. In this case, the metric simply indicates how many gateways the datagram has to go through. (This is often referred to as a "hop count".)

When computer A is ready to send a datagram, it examines the destination address. It gets the network number from the beginning of the address, and then looks in the routing table. The table entry indicates whether the datagram should be sent directly to the destination or to a gateway.

Note that a gateway is simply a computer that is connected to two different networks, and is prepared to forward datagrams between them. In many cases it is most efficient to use special-purpose equipment that are designed as gateways. However it is perfectly possible to use ordinary computers, as long as they have more than one network interface, and their software is prepared to forward datagrams. Most major TCP/IP implementations (even for microcomputers) are designed to let you use your computer as a gateway. However some of this software has limitations that can cause trouble for your network.

Note that a gateway has several addresses -- one for each network that it's attached to. This is a difference between IP and some other network protocols: each interface from a computer to a network has an address. With some other protocols, each computer has only one address, which applies to all of its interfaces. A gateway between networks

128.6.4 and 128.6.21 will have an address that begins with 128.6.4 (for example, 128.6.4.1). This address refers to its connection to network 128.6.4. It will also have an address that begins with 128.6.21 (for example, 128.6.21.2). This refers to its connection to network 128.6.21.

The term "network" probably makes you think of things like Ethernet, which can have many machines attached. However it also applies to point to point lines. In the diagram above, networks 1 and 3 could be in different cities. Then network 2 could be a serial line, satellite link, or other long-distance point to point connection between the two locations. A point to point line is treated as a network that just happens to have only two computers on it. As with any other network, the point to point line has a network number (in this case 128.6.21). The systems connected by the line (gateways R and S) have addresses on that network (in this case 128.6.21.1 and 128.6.21.2).

It is possible to design routing software that does not require a separate network number for each point to point line. In that case, the interface between the gateway and the point to point line doesn't have an address. This can be useful if your network is so large that you are in danger of running out of network numbers. However such "anonymous interfaces" can make network management somewhat more difficult. If there is no address, network management software may have no way to refer to the interface. Thus you may not be able to get data on throughput and errors for that interface.

3. Choosing an addressing structure

The first comment to make about addresses is a warning: Before you start using an IP network, you must get one or more official network numbers. IP addresses look like this: 128.6.4.3. This address is used by one computer at Rutgers University. The first part of it, 128.6, is a network number, allocated to Rutgers by a central authority. Before you start allocating addresses to your computers, you must get an official network number. Unfortunately, some people set up networks using either a randomly-chosen number, or a generic number supplied by the vendor. While this may work in the short run, it is a very bad idea for the long run. Eventually, you will want to connect your network to some other organization's network. Even if your organization is highly secret and very concerned about security, somewhere in your organization there is going to be a research computer that ends up being connected to a nearby university. That university will probably be connected to a large-scale national network. As soon as one of your datagrams escapes your local network, the organization you are talking to is going to become very confused, because the addresses that appear in your datagrams are probably officially allocated to someone else.

The solution to this is simple: get your own network number from the beginning. It costs nothing. If you delay it, then sometime years from now you are going to be faced with the job of changing every address on a large network. Network numbers are currently assigned by the DDN Network Information Center, SRI International, 333 Ravenswood Avenue, Menlo Park, California 94025 (telephone: 800-235-3155). You can get a network number no matter what your network is being used for. You do not need authorization to connect to the Defense Data Network in order to get a number. The main piece of information that will be needed when you apply for a network number is the address class that you want. See below for a discussion of this.

In many ways, the most important decision you have to make in setting up a network is how you will assign IP addresses to your computers. This choice should be made with a view of how your network is likely to grow. Otherwise, you will find that you have to change addresses. When you have several hundred computers, address changes can be nearly impossible.

Addresses are critical because IP datagrams are routed on the basis of their address. For example, addresses at Rutgers University have a 2-level structure. A typical address is 128.6.4.3. 128.6 is assigned to Rutgers University. As far as the outside world is concerned, 128.6 is a single network. Other universities send any datagram whose address begins with 128.6 to the nearest Rutgers gateway. However within Rutgers, we divide up our address space into "subnets". We use the next 8 bits of address to indicate which subnet a computer belongs to. 128.6.4.3 belongs to subnet 128.6.4. Generally subnets correspond to physical networks, e.g. separate Ethernets, although we will see some exceptions later. Systems inside Rutgers, unlike those outside, contain information about the Rutgers subnet structure. So once a datagram for 128.6.4.3 arrives at Rutgers, the Rutgers network will route it to the departmental Ethernet, token ring, or whatever, that has been assigned subnet number 128.6.4.

When you start a network, there are several addressing decisions that face you:

- Do you subdivide your address space?
- If so, do you use subnets or class C addresses?
- How big an address space do you need?

3.1 Should you subdivide your address space?

It is not necessary to use subnets at all. There are network technologies that allow an entire campus or company to act as a single large logical Ethernet, so that no internal routing is necessary. If you use this technology, then you do not need to subdivide your address space. In that case, the only decision you have to make is what class of address to apply for. However we recommend using either a subnet approach or some other method of subdividing your address space in most networks:

- In section 6.2 we will argue that internal gateways are desirable for all networks beyond the very simplest.
- Even if you do not need gateways now, you may find later that you need to use them. Thus it probably makes sense to assign addresses as if each Ethernet or token ring were going to be a separate subnet. This will allow for conversion to real subnets later if it proves necessary.
- For network maintenance purposes, it is convenient to have addresses whose structure corresponds to the structure of the network. For example, when you see a stray datagram from system 128.6.4.3, it is nice to know that all addresses beginning with 128.6.4 are in a particular building.

3.2 Subnets vs. multiple network numbers

Suppose that you have been convinced that it's a good idea to impose some structure on your addresses. The next question is what that structure should be. There are two basic approaches. One is subnets. The other is multiple network numbers.

The Internet standards specify the format of an address. For addresses beginning with 128 through 191 (the most common numbers these days), the first two octets form the network number. E.g. in 140.3.50.1, 140.3 is the network number. Network numbers are assigned to a particular organization. What you do with the next two octets is up to you. You could choose to make the next octet be a subnet number, or you could use some other scheme entirely. Gateways within your organization must be set up to know the subnetting scheme that you are using. However outside your organization, no one will know that 140.3.50 is one subnet and 140.3.51 is another. They will simply know that 140.3 is your organization. Unfortunately, this ability to add additional structure to the address via subnets was not present in the original IP specifications. Thus some older software is incapable of being told about subnets.

If enough of the software that you are using has this problem, it may be impractical for you to use subnets. Some organizations have used a different approach. It is possible for an organization to apply for several network numbers. Instead of dividing a single network number, say 140.3, into several subnets, e.g. 140.3.1 through 140.3.10, you could apply for 10 different network numbers. Thus you might be assigned the range 140.3 through 140.12. All IP software will know that these are different network numbers.

While using separate network numbers will work just fine within your organization, it has two very serious disadvantages. The first, and less serious, is that it wastes address space. There are only about 16,000 possible class B addresses. We cannot afford to waste 10 of them on your organization, unless it is very large. This objection is less serious because you would normally ask for class C addresses for this purpose, and there are about 2 million possible class C addresses.

The more serious problem with using several network numbers rather than subnets is that it overloads the routing tables in the rest of the Internet. As mentioned above, when you divide your network number into subnets, this division is known within your organization, but not outside it. Thus systems outside your organization need only one entry in their tables in order to be able to reach you. E.g. other universities have entries in their routing tables for 128.6, which is the Rutgers network number. If you use a range of network numbers instead of subnets, that division will be visible to the entire Internet. If we used 128.6 through 128.16 instead of subdividing 128.6, other universities would need entries for each of those network numbers in their routing tables. As of this writing the routing tables in many of the national networks are exceeding the size of the current routing technology. It would be considered extremely unfriendly for any organization to use more than one network number. This may not be a problem if your network is going to be completely self-contained, or if only one small piece of it will be connected to the outside world. Nevertheless, most TCP/IP experts strongly recommend that you use subnets rather than multiple networks. The only reason for considering multiple networks is to deal with software that cannot handle subnets. This was a problem a few years ago, but is currently less serious. As long as your gateways can handle subnets, you can deal with a few individual computers that cannot by using "proxy ARP" (see below).

One warning about subnets: Your subnets must all be "adjacent". That is, you can't have a configuration where you get from subnet 128.6.4 to subnet 128.6.5 by going through some other network entirely, e.g. 128.121. For example, Rutgers has campuses in New Brunswick and Newark. It is perfectly OK for the networks in both cities to be subnets of 128.6. However in that case, the lines between New Brunswick and Newark must also be part of 128.6. Suppose we decided to use a regional network such as JvNCnet to talk between our two campuses, instead of providing our own lines. Since JvNCnet is 128.121, the gateways and serial lines that they provide would have addresses that begin with 128.121. This violates the rules. It is not allowable to have gateways or lines that are part of 128.121 connecting two parts of 128.6. So if we wanted to use JvNCnet between our two campuses, we'd have to get different network numbers for the two campuses. (This rule is a result of limitations in routing technology. Eventually gateway software will probably be developed that can deal with configurations whose networks are not contiguous.)

3.3 How to allocate subnet or network numbers

Now that you have decided to use subnets or multiple network numbers, you have to decide how to allocate them. Normally this is fairly easy. Each physical network, e.g. Ethernet or token ring, is assigned a separate subnet or network number. However you do have some options.

In some cases it may make sense to assign several subnet numbers to a single physical network. At Rutgers we have a single Ethernet that spans three buildings, using repeaters. It is very clear to us that as computers are added to this Ethernet, it is going to have to be split into several separate Ethernets. In order to avoid having to change addresses when this is done, we have allocated three different subnet numbers to this Ethernet, one per building. (This would be handy even if we didn't plan to split the Ethernet, just to help us keep track of where computers are.) However before doing this, make very sure that the software on all of your computers can handle a network that has three different network numbers on it. This issue is discussed in more detail in section 3.4.

You also have to choose a "subnet mask". This is used by the software on your systems to separate the subnet from the rest of the address. So far we have always assumed that the first two octets are the network number, and the next octet is the subnet number. For class B addresses, the standards specify that the first two octets are the network number. However we are free to choose the boundary between the subnet number and the rest of the address. It's very common to have a one-octet subnet number, but that's not the only possible choice. Let's look again at a class B address, e.g. 128.6.4.3. It is easy to see that if the third octet is used for a subnet number, there are 256 possible subnets and within each subnet there are 256 possible addresses. (Actually, the numbers are more like 254, since it is generally a bad idea to use 0 or 255 for subnet numbers or addresses.) Suppose you know that you will never have more than 128 computers on a given subnet, but you are afraid you might need more than 256 subnets. (For example, you might have a campus with lots of small buildings.) In that case, you could define 9 bits for the subnet number, leaving 7 bits for addresses within each subnet. This choice is expressed by a bit mask, using ones for the bits used by the network and subnet number, and 0's for the bits used for individual addresses. Our normal subnet mask choice is given as 255.255.255.0. If we chose 9 bit subnet numbers and 7 bit addresses, the subnet mask would be 255.255.255.128.

Generally it is possible to specify the subnet mask for each computer as part of configuring its IP software. The IP protocols also allow for computers to send a query asking what the subnet mask is. If your network supports broadcast queries, and there is at least one computer or gateway on the network that knows the subnet mask, it may be unnecessary to set it on the other computers. However this capability brings with it a whole new set of possible problems. One well-known TCP/IP implementation would answer with the wrong subnet mask when queried, thus leading causing every other computer on the network to be misconfigured. Thus it may be safest to set the subnet mask explicitly on each system.

3.4 Dealing with multiple "virtual" subnets on one network

Most software is written under the assumption that every computer on the local network has the same subnet number. When traffic is being sent to a machine with a different subnet number, the software will generally expect to find a gateway to handle forwarding to that subnet. Let's look at the implications. Suppose subnets 128.6.19 and 128.6.20 are on the same Ethernet. Consider the way things look from the point of view of a computer with address 128.6.19.3. It will have no problem sending to other machines with addresses 128.6.19.x. They are on the same subnet, and so our computer will know to send directly to them on the local Ethernet. However suppose it is asked to send a datagram to 128.6.20.2. Since this is a different subnet, most software will expect to find a gateway that handles forwarding between the two subnets. Of course there isn't a gateway between subnets 128.6.19 and 128.6.20, since they are on the same Ethernet. Thus you have to find a way to tell your software that 128.6.20 is actually on the same Ethernet.

Most common TCP/IP implementations can deal with more than one subnet on a network. For example, Berkeley Unix lets you use a slight modification of the command used to define gateways. Suppose that you get from subnet 128.6.19 to subnet 128.6.4 using a gateway whose address is 128.6.19.1. You would use the command

```
route add 128.6.4.0 128.6.19.1 1
```

This says that to reach subnet 128.6.4, traffic should be sent via the gateway at 128.6.19.1, and that the route only has to go through one gateway. The "1" is referred to as the "routing metric". If you use a metric of 0, you are saying that the destination subnet is on the same network, and no gateway is needed. In our example, on system 128.6.19.3, you would use

```
route add 128.6.20.0 128.6.19.1 0
```

The actual address used in place of 128.6.19.1 is irrelevant. The metric of 0 says that no gateway is actually going to be used, so the gateway address is not used. However it must be a legal address on the local network.

Note that the commands in this section are simply examples. You should look in the documentation for your particular implementation to see how to configure your routing.

3.4.1 Dealing with Multiple Subnets by Turning off Subnetting

There is another way to handle several subnets on one physical network. This method involves intentionally misconfiguring your hosts, so it is potentially dangerous if you don't watch what you are doing. However it may be easier to deal with when you have lots of subnets on one physical network. An example of this is a site that uses bridges, and uses subnets simply for administrative convenience. The trick is to configure the software on your hosts as if you were not using subnets at all. In this case your hosts will not make any distinction between the subnets, and they'll have no trouble dealing with all of them. Now the only problem is how to talk to subnets that are not on this multi-subnet network. However if your gateways handle proxy ARP, they will solve that problem for you. This approach is likely to be convenient when the same network is carrying many subnets, particularly if additional ones are likely to be added later. However it has two problems:

If you have any hosts with multiple interfaces, you will have to be very careful. First, only one interface should be on the multi-subnet network. For example, suppose you have a "network" that is made up of several Ethernets connected by bridges. You can't have a machine with interfaces on two of those Ethernets. However you can have a system with one interface on the multi-subnet network and another on some totally separate subnet. Second, any machine with multiple interfaces will have to know the real subnet mask, and will need to be told explicitly which subnets are on the multi-subnet network. These restrictions come about because a system with multiple interfaces has to know which interface to use in any given case.

You will have to be careful about the ICMP subnet mask facility. This is a facility that allows systems to broadcast a query asking what the subnet mask is. If most of your hosts think the network is not subnetted, but your gateways and multi-interface hosts think it is, you've got a potential for confusion. If a gateway or multi-interface host happens to send an ICMP subnet mask reply giving the real subnet mask, some of your other hosts may pick it up. The reverse is possible as well. This means that you will either have to

- disable ICMP subnet mask replies on all of the systems that know the real subnet mask. (This may be easy if only gateways know it.)
- make sure that your hosts ignore ICMP replies

According to the most recent documents, as long as you set the subnet mask explicitly, hosts are supposed to ignore the ICMP subnet mask mechanism. So you should be able to set different masks on different hosts without causing any problem, as long as you set the mask explicitly for all of them. However we have noticed that some IP implementations will change their subnet mask when they see an ICMP subnet mask reply.

3.4.2 Multiple Subnets: Implications for Broadcasting

When you have more than one subnet on the same physical network, you need to give some thought to broadcast addresses. According to the latest standards, there are two different ways for a host on subnet 128.6.20 to send a broadcast on the local network. One is to use address 128.6.20.255. The other is to use address 255.255.255.255. 128.6.20.255 says explicitly "all hosts on subnet 128.6.20". 255.255.255.255 says "all hosts on my local network". Normally these have the same effect. However they do not when there are several subnets on one physical network. If subnet 128.6.19 is on the same Ethernet, it is also going to receive messages sent to 255.255.255.255. However hosts with numbers 128.6.19.x will not listen to broadcasts to 128.6.20.255. The

result is that the two different forms of broadcast address will have somewhat different meanings. This means that you will have to exercise some care in configuring software on networks such as this, to make sure that broadcasts go where you intend them to go.

3.5 Choosing an address class

When you apply for an official network number, you will be asked what class of network number you need. The possible answers are A, B, and C. This affects how large an address space you will be allocated. Class A addresses are one octet long, class B addresses are 2 octets, and class C addresses are 3 octets. This represents a tradeoff: there are a lot more class C addresses than class A addresses, but the class C addresses don't allow as many hosts. The idea was that there would be a few very large networks, a moderate number of medium-size ones, and a lot of mom-and-pop stores with small networks. Here is a table showing the distinction:

class	range of first octet	network	rest	possible addresses
A	1 - 126	p	q.r.s	16777214
B	128 - 191	p.q	r.s	65534
C	192 - 223	p.q.r	s	254

For example network 10, a class A network, has addresses between 10.0.0.1 and 10.255.255.254. So it allows 254×3 , or about 16 million possible addresses. (Actually, network 10 has allocated addresses where some of the octets are zero, so there are a few more addresses possible.) Network 192.12.88, a class C network has hosts between 192.12.88.1 and 192.12.88.254, i.e. 254 possible hosts.

In general, you will be expected to choose the lowest class that will provide you with enough addresses to handle your growth over the next few years. Organizations that have computers in many buildings will probably need and be able to get a class B address, assuming that they are going to use subnetting. (If you are going to use many separate network numbers, you would ask for a number of class C addresses.) Class A addresses are normally used only for large public networks and for a few very large corporate networks.

3.6 Dialup IP and Micro gateways: Dynamically assigned addresses

In most cases, each of your computers will have its own permanent IP address. However there are a few situations where it makes more sense to allocate addresses dynamically. The most common cases involve dialup IP, and gateways intended primarily for microcomputers.

3.6.1 Dialup IP

It is possible to run IP over dialup lines. The protocol for doing so is called SLIP ("serial line IP"). SLIP is useful in at least two different circumstances:

- As a low-cost alternative to permanent point to point lines, for cases where there isn't enough traffic to justify dedicated lines.
- As a way to connect individual PC's into a network when they are located in buildings that don't have Ethernets or other LAN technology.

I am going to use the term "SLIP server" to refer to a computer system that has modems attached, which other systems can connect to using SLIP. Such a system will provide a gateway into your network for PC users or for other networks that connect using SLIP.

If you have a number of individual PC's dialing up with SLIP, it is often not practical to assign each PC its own IP address. For one thing, there may just not be enough addresses. In order to keep the routing straight, the dialup systems have to get addresses on the same subnet as the SLIP server. Generally there are only 256 or so addresses available on each subnet. If you have more PC's than that, you can't give each one its own address. If you have SLIP servers on more than one subnet, this will make permanent addresses even more difficult. If a user wanted to be able to call both servers, his PC would need two addresses, one for each subnet.

In order to avoid these problems, many SLIP implementations assign addresses dynamically. When a PC first connects to the SLIP server, the server finds an unused IP address and assigns it to the PC. The simplest way to manage this is to give each SLIP server a range of IP addresses that it keeps track of and can assign.

When you use such a scheme, your SLIP software has to include some way for the server to tell the PC what address to use. If each PC has a permanent address, you have the reverse problem: when a PC connects to a server, there has

to be a way for the PC to tell the server what its address is. Some care is needed. Otherwise someone could have his PC claim to be yours and steal all your files.

Unfortunately, there is no standard way to manage these addressing issues with SLIP. There are several SLIP implementations that handle them, but there isn't a single standard yet. Until such a standard is developed, you need to check out SLIP software carefully. Make sure that it assigns addresses the way you want, and that your SLIP server and your PC's agree on how to figure out the PC's address.

I recommend giving the PC's permanent addresses in cases where other computers have to be able to tell which PC they are talking to. This would be the case if the PC is going to receive private computer mail, or engage in other sensitive transactions. I recommend using dynamic addresses where you have a lot of PC's, and where the applications that they access over the network do their own security checking.

When you are using SLIP to connect two networks, you have three choices for handling addressing (although not all SLIP software can handle all three choices):

- Treat SLIP connections like point to point lines that just don't happen to be up all the time. If you call more than one computer, each pair of computers that talks has a separate network number which they use only when they talk to each other.
- Use routing software that allows anonymous interfaces. In that case no address is needed at all.
- Assign addresses dynamically when the connection is opened, just as you would for a PC that is dialing up.

If you make connections only to one or two other systems, it is quite reasonable to use a network number for each connection. This method makes it easy to keep usage and error statistics.

If you have many different connections, it is probably best to use anonymous interfaces. You would probably use dynamic address allocation only if your routing technology did not support anonymous interfaces.

3.6.2 Micro gateways

It is perfectly possible for microcomputers to participate in an IP network. However there seems to be a tendency for micros to use somewhat different network technology than larger systems. This is because many micro users start with specialized network software whose design is tailored specifically to the needs of micros, or even some particular type of micro. Micro users quite naturally want to be able to start using TCP/IP without having to abandon any special micro network that they are already using. For that reason there is a growing number of gateway products that allow PC's to access both some micro-oriented network product and TCP/IP.

In this section, Apple's AppleTalk is used as an example. This is because gateways for it have existed for some time, and are in widespread use. However similar products exist for several other micro network technologies. Note that the term AppleTalk refers to the Apple network protocols, whereas LocalTalk refers to the specific twisted-pair technology on which AppleTalk was initially implemented. Thus AppleTalk is analogous to the TCP/IP protocols, whereas LocalTalk is analogous to the Ethernet medium.

Several vendors supply gateways to connect AppleTalk running over a LocalTalk network with IP running over Ethernet. Although there are several products of this kind, most of them supply the following services:

- TCP/IP applications on the PC can connect to TCP/IP systems on the Ethernet. Special facilities are defined to allow IP datagrams to be carried over LocalTalk between the PC and the gateway. TCP/IP applications on the PC have to be written using a special library that uses a mixture of AppleTalk and TCP/IP. The AppleTalk facilities are needed to get the datagrams to the gateway, where they are transformed into pure TCP/IP before being put out onto the Ethernet. Thus the TCP/IP systems on the Ethernet don't know they are talking to micros.
- AppleTalk applications can be written for larger systems, so that PC's can use them as servers. These applications are written using a special library that is more or less the reverse of the one just described. Again, it uses a mixture of AppleTalk and TCP/IP. But this time TCP/IP facilities are needed to get the datagrams to the gateway, where they are transformed into pure AppleTalk before being put onto the LocalTalk network to communicate with the PC's. Thus the PC's can access applications on the larger systems, without knowing that they are on the Ethernet rather than an Apple network.

- A campus or corporate IP network can be used to connect AppleTalk networks at different locations. Gateways at each location wrap up AppleTalk datagrams inside IP datagrams, and send them over the main IP network.

In addition, some newer gateways will translate at the application level. For example one gateway will translate between the Apple filing protocol and Sun's Network File System. This allows a PC to access a Unix file system, with the PC using the Apple filing protocol, and the final access to the Unix system being done using Sun's Network File System.

Unfortunately the flexibility of products like this also means that they are complex. Addressing issues are particularly complicated. For the same reasons as SLIP, these gateways often use dynamic IP address allocation. A range of IP addresses is assigned to each gateway. When a PC attempts to open its first TCP/IP connection, the gateway picks a free IP address and assigns it to the PC. As with SLIP, you will often need to choose whether you want addresses to be assigned this way, or you want each PC to have its own address. Again, this depends upon how many PC's you have and whether you have applications which must be able to use the IP address to identify the particular PC that is talking to it.

Addressing is further complicated by the fact that AppleTalk has its own addressing structure. So you must define a mapping between AppleTalk and IP network numbers. There must also be a mapping between individual IP addresses and AppleTalk addresses, but this mapping is maintained dynamically by the gateways.

4. Network-wide Services, Naming

If you are going to have a TCP/IP network, there are certain things that you are going to have to do centrally. Some of them are simply administrative. The most important is that you will have a central registry of names and IP addresses. The DDN Network Information Center performs this role for the Internet network as a whole. If you are connected to the international Internet, your administrator will need to register with the DDN Network Information Center, so that queries from other institutions about your hosts are forwarded to your servers.

You will want to maintain a database containing information about each system on your network. At a minimum, you need to have the host name and IP address for each system. Probably the central registry will assign IP addresses. If your network is subnetted, or if you use multiple class C network numbers, the registry will probably assign network numbers to new networks or subnets. Most commonly, individual host administrators will be allowed to choose their own host names. However the registry must at least verify that there are no duplicate names. If you have a very large network, you may choose to delegate some of these tasks to subregistries, possibly one for each department.

We suggest that you assign numbers in the simplest way: starting from 1. Thus if your network is 128.6, you would assign 128.6.1 as your first subnet, 128.6.2 as the second, etc. IP addresses for individual hosts should probably start at 2. This allows you to reserve 1 on each subnet for use by a gateway. Thus the first host on subnet 128.6.4 would be 128.6.4.2, the next 128.6.4.3, etc. There is a specific reason for keeping addresses as small as possible. If you have a large organization, you may run out of subnet numbers. If you do, and if your host numbers are small, you can assign another bit for the subnet. For example, we use the entire third octet as a subnet number. As long as all of our host numbers are less than 128, we will be able to expand to 9-bit subnet numbers. For example, subnet 128.6.4 would be split into two separate subnets, 128.6.4.0 and 128.6.4.128. If we had assigned host numbers above 128, this split would be impossible.

Host names need not be so systematic. They can start with almost any word made up of letters numbers, and hyphens. It is safest for the first character to be a letter. It will be easier for users if the name is fairly short. (We have seen software that has trouble dealing with names longer than 16 characters.) Many times departments or projects choose a theme, and pick names that are consistent with them. For example, the machines used by computer science graduate students at Rutgers are named after rock bands: STEELEYE, BAND, TREX, DEVO, etc. Our math department uses famous mathematicians: GAUSS, FERMAT, etc. If your institution does not have any connection with the outside world, such one-word names are all you need.

If you are connected to with the international Internet, your organization will need to get a "domain name." This is assigned to you by the DDN Network Information Center, just as your network number is. Unlike the network number, you can get along without one if your network is isolated. If you find later that you need one, it is easy to add a domain name. (We recommend that you start with an official network number from the beginning because changing network numbers later can be traumatic.) Domain names normally end in .EDU for educational institutions, .COM for companies, etc. For example, Rutgers University has a domain name of .RUTGERS.EDU A full domain-style host name consists of your one-word internal name followed by your organization's domain name. For example, the computer I normally use is known internally as ATHOS. It's full name is

ATHOS.RUTGERS.EDU If you have a large organization, it is possible to have sub-domains. For example, you might have a subdomain for each department. This adds another period to your names. For example, the computer science department might have decided to create a subdomain. In this case, my computer would probably be called ATHOS.CS.RUTGERS.EDU Once you get a domain name assigned to you, it is wise to change all of your configuration files so that the full form of name is used. However your software can be set up so that the one-word versions are accepted as nicknames. That way your users don't have to type out the long form.

If you have more than one or two systems, you are going to need some way to keep host information up to date on all of your systems. TCP/IP software needs to be able to translate host names into IP addresses. When a user tries to connect to another system, he wants to be able to refer to it by name. The software has to translate the name into the IP address in order to open the connection. Most software provides two ways to do this translation: a static table or a name server. The table approach is probably easier for small organizations, as long as they are not connected to any other network. You simply create a file that lists the names and addresses of all your hosts. Here's part of our host table:

```
HOST: 128.6.4.2, 128.6.25.2 : ARAMIS.RUTGERS.EDU,ARAMIS : SUN-3-28
HOST: 128.6.4.3 : GAUSS.RUTGERS.EDU,GAUSS : SUN-3-180 : UNIX ::
HOST: 128.6.4.4, 128.6.25.4 : ATHOS.RUTGERS.EDU,ATHOS : SUN-4-280
```

This format has one line for each system, and lists its addresses, names, and other information about it. Note that aramis and athos are both on two networks, so they have two addresses. They have both primary names, e.g. ARAMIS.RUTGERS.EDU, and nicknames, e.g. ARAMIS. Since we are attached to the Internet, our primary name is a full domain name. We supply brief nicknames to make it easier for our users. There is one other commonly-used format for the host table. Here's an example of that format:

```
128.6.4.2   aramis.rutgers.edu aramis
128.6.25.2  aramis.rutgers.edu aramis
128.5.4.3   gauss.rutgers.edu  gauss
128.6.4.4   athos.rutgers.edu  gauss
128.6.25.4  athos.rutgers.edu  gauss
```

In this format, each line represents a single IP address. If a system has two interfaces, there are two lines in the table for it. You should try to put the address first that is likely to be used more often. The documentation for your systems should indicate what format they want the host information to use.

In the simplest setup, every computer has its own copy of the host table. If you choose to use the setup, you will want to set up procedures to make sure that systems get updated copies of the host table regularly.

Larger sites, and all sites that are connected to the Internet, should use name servers instead of individual host tables. A name server is a program that you run on a few of your systems to keep track of names. When a program needs to look up a name, instead of looking for a copy of the host table, it sends a network query to the name server. This approach has two advantages:

- For a large site, it is easier to keep tables up to date on a few name servers than on every system.
- If your site is connected to the Internet, your name server will be able to talk to name servers at other organizations, and look up names elsewhere.

Using a name server is the only way to have access to complete host information about the rest of the Internet.

It is important to understand the difference between a name server and a resolver. A name server is a program that accesses a host database, and answers queries from other programs. A resolver is a set of subroutines that can be loaded with your program. It generates queries to the name server, and processes the responses. Every system should use the resolver. (Actually, the resolver is generally loaded with each program that uses the network, since it's simply a set of subroutines.) You only need a few name servers. Many people confuse these two concepts, and come to believe that every computer needs to run a name server.

In order to use a resolver, each computer will need a configuration file or other option that specifies the address of a name server where queries should be sent. Generally you should specify several name servers, in case one of them is down. If your system cannot reach any name server, much of your software is likely to misbehave. Thus you should be very careful to have enough name servers around that every system can always reach at least one name server.

Name servers generally have a number of configuration options. Rather than giving advice here on setting up a name server, I am going to refer you to two official Internet standards documents. Both are available from the DDN Network Information Center, SRI International, 333 Ravenswood Avenue, Menlo Park, California 94025 (telephone:

800-235-3155). RFC 1032 contains instructions for getting a domain name from the Network Information Center, including the necessary forms. RFC 1033 contains instructions on how to set up a name server. Like this document, these documents are conceptual. You will also need documentation for the specific name server software that you are going to use. [This paragraph is a cop-out. Future editions of this document will contain some advice on setting up a name server. However RFC 1033 is almost unique in that it is directed at administrators rather than networking experts. Thus it is reasonable to direct people there for the moment.]

In some cases you may need to use both fixed tables and name servers. If you have some TCP/IP implementations that do not include resolvers, then you will have to have host tables for those systems. If your network is connected to the international Internet, you are going to have problems with systems that don't have resolvers. The Internet is too big for there to be a host table that lists all of its hosts. Thus you will have to put together a host table that lists those hosts that your users tend to use. The DDN Network Information Center maintains a host table that will be a good starting point. However it is by no means complete. So you will have to add your users' favorite hosts to it. Systems that use a resolver will not have this problem, since the name servers are able to translate any legal host name.

Host name and number allocation is the only facility that has to be done centrally. However there are other things that you may prefer to do centrally. It is very common to have one or two computers that handle all computer mail. If are on the Internet, it is easy for every one of your computers to talk directly to any other computer on the Internet. However most institutions want to communicate with systems on other networks, such as Bitnet and Usenet. There are gateways between the various networks. But choosing the right gateway, and transforming computer mail addresses correctly is a rather specialized business. Thus many sites set up the appropriate software only one place and direct all external mail (or all external mail to hosts that are not on the Internet) through this system.

5. Setting up routing for an individual computer

All TCP/IP implementations require some configuration for each host. In some cases this is done during "system generation". In other cases, various startup and configuration files must be set up on the system. Still other systems get configuration information across the network from a "server". While the details differ, the same kind of information needs to be supplied for most implementations. This includes

- parameters describing the specific machine, such as its IP address.
- parameters describing the network, such as the subnet mask (if any)
- routing software and the tables that drive it
- startup of various programs needed to handle network tasks

Before a machine is installed on your network, a coordinator should assign it a host name and IP address, as described above. Once you have name and address, you are ready to start configuring your computer. Often you have to put the address and name into a configuration file on the computer. However some computers (particularly those without permanent disks on which configuration information could be stored) get this information over the network. When such a system starts, it broadcasts a request over the network. In effect, this request says "who am I?" If you have any computers like this, you will have to make sure that some system on your network is ready to answer these questions. The obvious issue is: how can another system tell who you are? Generally this is done based on Ethernet address (or the analogous address for other types of network). Ethernet addresses are assigned by the computer manufacturer. It is guaranteed that only one machine in the entire world has any particular Ethernet address. The address is normally stored in ROM somewhere in the machine. The machine may not know its IP address, but it does know its Ethernet address. Thus the "who am I" request includes the Ethernet address. Systems that are set up to answer such requests have a table that lists Ethernet addresses and the corresponding IP address. This lets them know how to answer. Unfortunately, you have to set this table up manually. Generally you know the IP address, because your address coordinator has assigned it. The only problem in constructing the table will be finding out the Ethernet address for each computer. Generally, computers are designed so that they print the Ethernet address on the console shortly after being turned on. However in some cases you may have to find a way to bring the computer up and then type a command that displays information about the Ethernet interface.

Generally the subnet mask should be specified in a configuration file associated with the computer. (For Unix systems, the "ifconfig" command is used to specify both the Internet address and subnet mask.) However there are provisions in the IP protocols for a computer to broadcast a request asking for the subnet mask. The subnet mask is an attribute of the network. It is the same for all computers on a given subnet. Thus there is no separate subnet

table corresponding to the Ethernet/Internet address mapping table used to answer address queries. Ideally, only a few authoritative computers will answer queries about the subnet mask. However many TCP/IP implementations are set up so that any machine on the network that believes it knows the subnet mask will answer. If your TCP/IP is like this, an incorrect subnet mask setting on one machine can cause confusion throughout the network.

Normally the startup files do roughly the following things:

- load any special device drivers that may be necessary. (This is particularly common with PC's, where network access is likely to depend upon add-on controller cards and software that is not part of the original operating system.)
- enable each of the network interfaces (Ethernet interface, serial lines, etc.) Normally this involves specifying an Internet address and subnet mask for each, as well as other options that will be described in your vendor's documentation.
- establish network routing information, either by commands that add fixed routes, or by starting a program that obtains them dynamically.
- turn on the domain system (used for looking up names and finding the corresponding Internet address -- see the section on the domain system in the Introduction to TCP/IP). Note that the details of this will depend upon how the domain system is configured. In most cases only a few hosts actually run domain name servers that must be started. Other hosts simply need configuration files that specify where the nearest name server is located.
- set various other information needed by the system software, such as the name of the system itself.
- start various "daemons". These are programs that provide network services to other systems on the network, and to users on this system. In the case of PC's, which often cannot run multiple processes, similar facilities may be provided by so-called "TSR"s, or they may be built into the device drivers.

It is not practical to document these steps in detail, since they differ for each vendor. This section will concentrate on a few issues where your choice will depend upon overall decisions about how your network is to operate. These overall network policy decisions are often not as well documented by the vendors as the details of how to start specific programs. Note that some care will be necessary to integrate commands that you add for routing, etc., into the startup sequence at the right point. Some of the most mysterious problems occur when network routing is not set up before a program needs to make a network query, or when a program attempts to look up a host name before the name server has finished loading all of the names from a master name server.

5.1 How datagrams are routed

If your system consists of a single Ethernet or similar medium, you do not need to give routing much attention. However for more complex systems, each of your machines needs a routing table that lists a gateway and interface to use for every possible destination network. A simple example of this was given at the beginning of this section. However it is now necessary to describe the way routing works in a bit more detail. On most systems, the routing table looks something like the following. (This example was taken from a system running Berkeley Unix, using the command "netstat -n -r". Some columns containing statistical information have been omitted.)

Destination	Gateway	Flags	Interface
128.6.5.3	128.6.7.1	UHGD	il0
128.6.5.21	128.6.7.1	UHGD	il0
127.0.0.1	127.0.0.1	UH	lo0
128.6.4	128.6.4.61	U	pe0
128.6.6	128.6.7.26	U	il0
128.6.7	128.6.7.26	U	il0
128.6.2	128.6.7.1	UG	il0
10	128.6.4.27	UG	pe0
128.121	128.6.4.27	UG	pe0
default	128.6.4.27	UG	pe0

The example system is connected to two Ethernets:

controller	network	address	other networks
i10	128.6.7	128.6.7.26	128.6.6
pe0	128.6.4	128.6.4.61	none

The first column shows the name for the Ethernet interface. The second column is the network number for that Ethernet. The third column is this computer's Internet address on that network. The last column shows other subnets that share the same physical network.

Now let's look at the routing table. For the moment, let us ignore the first 3 lines. The majority of the table consists of a set of entries describing networks. For each network, the other three columns show where to send datagrams destined for that network. If the "G" flag is present in the third column, datagrams for that network must be sent through a gateway. The second column shows the address of the gateway to be used. If the "G" flag is not present, the computer is directly connected to the network in question. So datagrams for that network should be sent using the controller shown in the third column. The "U" flag in the third column simply indicates that the route specified by that line is up. (Generally a route is assumed to be up unless attempts to use it consistently result in errors.)

The first 3 lines show "host routes", indicated by the "H" flag in column three. Routing tables normally have entries for entire networks or subnets. For example, the entry

```
128.6.2 128.6.7.1 UG i10
```

indicates that datagrams for any computer on network 128.6.2 (i.e. addresses 128.6.2.1 through 128.6.2.254) should be sent to gateway 128.6.7.1 for forwarding. However sometimes routes apply only to a specific computer, rather than to a whole network. In that case, a host route is used. The first column then shows a complete address, and the "H" flag is present in column 3. E.g. the entry

```
128.6.5.21 128.6.7.1 UHGD i10
```

indicates that datagrams for the specific address 128.6.5.21 should be sent to the gateway 128.6.7.1. As with network routes, the "G" flag is used for routes that involve a gateway. The "D" flag indicates that the route was added dynamically, based on an ICMP redirect message from a gateway. (See below for details.)

The following route is special:

```
127.0.0.1 127.0.0.1 UH lo0
```

127.0.0.1 is the address of the "loopback device". This is a dummy software module. Any datagram sent out through that "device" appears immediately as input. It can be used for testing. The loopback address can also be handy for talking to applications that are on your own computer. (Why bother to use your network to talk to a program that is on the same machine you are?)

Finally, there are "default" routes, e.g.

```
default 128.6.4.27 UG pe0
```

This route is used for datagrams that don't match any other entry. In this case, they are sent to a gateway with address 128.6.4.27.

In most systems, datagrams are routed by looking up the destination address in a table such as the one just described. If the address matches a specific host route, then that is used. Otherwise, if it matches a network route, that is used. If no other route works, the default is used. If there is no default, the user should get an error message such as "network is unreachable".

The following sections will describe several ways of setting up these routing tables. Generally, the actual operation of sending datagrams doesn't depend upon which method you use to set up the routes. When a datagram is to be sent, its destination is looked up in the table. The different routing methods are simply more and less sophisticated ways of setting up and maintaining the tables.

5.2 Fixed routes

The simplest way to set up routing is to use fixed commands. Your startup files contain commands to set up the routing table. If any changes are needed, you make them manually, using commands that add and delete entries in the routing table. (When you make such a change, don't forget to update the startup files also.) This method is practical for relatively small networks, particularly if they don't change very often.

Most computers automatically set up some routing entries for you. Unix will add an entry for the networks to which you are directly connected. For example, your startup file might contain the commands

```
ifconfig ie0 128.6.4.4 netmask 255.255.255.0
ifconfig ie1 128.6.5.35 netmask 255.255.255.0
```

These specify that there are two network interfaces, and your addresses on them. The system will automatically create routing table entries

```
128.6.4 128.6.4.4 U ie0
128.6.5 128.6.5.35 U ie1
```

These specify that datagrams for the local subnets, 128.6.4 and 128.6.5, should be sent out the corresponding interface.

In addition to these, your startup files would contain commands to define routes to whatever other networks you wanted to reach. For example,

```
route add 128.6.2.0 128.6.4.1 1
route add 128.6.6.0 128.6.5.35 0
```

These commands specify that in order to reach network 128.6.2, a gateway at address 128.6.4.1 should be used, and that network 128.6.6 is actually an additional network number for the physical network connected to interface 128.6.5.35. Some other software might use different commands for these cases. Unix differentiates them by the "metric", which is the number at the end of the command. The metric indicates how many gateways the datagram will have to go through to get to the destination. Routes with metrics of 1 or greater specify the address of the first gateway on the path. Routes with metrics of 0 indicate that no gateway is involved -- this is an additional network number for the local network.

Finally, you might define a default route, to be used for destinations not listed explicitly. This would normally show the address of a gateway that has enough information to handle all possible destinations.

If your network has only one gateway attached to it, then of course all you need is a single entry pointing to it as a default. In that case, you need not worry further about setting up routing on your hosts. (The gateway itself needs more attention, as we will see.) The following sections are intended to provide help for setting up networks where there are several different gateways.

5.3 Routing redirects

Most TCP/IP experts recommend leaving routing decisions to the gateways. That is, it is probably a bad idea to have large fixed routing tables on each computer. The problem is that when something on the network changes, you have to go around to many computers and update the tables. If changes happen because a line goes down, service may not be restored until someone has a chance to notice the problem and change all the routing tables.

The simplest way to keep routes up to date is to depend upon a single gateway to update your routing tables. This gateway should be set as your default. (On Unix, this would mean a command such as "route add default 128.6.4.27 1", where 128.6.4.27 is the address of the gateway.) As described above, your system will send all datagrams to the default when it doesn't have any better route. At first, this strategy does not sound very good if you have more than one gateway. After all, if all you have is a single default entry, how will you ever use the other gateways in the cases where they are better? The answer is that most gateways are able to send "redirects" when they get datagrams for which there is a better route. A redirect is a specific kind of message using the ICMP (Internet Control Message Protocol). It contains information that generally translates to "In the future, to get to address XXXXX, please use gateway YYYYYY instead of me". Correct TCP/IP implementations use these redirects to add entries to their routing table. Suppose your routing table starts out as follows:

Destination	Gateway	Flags	Interface
127.0.0.1	127.0.0.1	UH	lo0
128.6.4	128.6.4.61	U	pe0
default	128.6.4.27	UG	pe0

This contains an entry for the local network, 128.6.4, and a default pointing to the gateway 128.6.4.27. Suppose there is also a gateway 128.6.4.30, which is the best way to get to network 128.6.7. How do you find it? Suppose you have datagrams to send to 128.6.7.23. The first datagram will go to the default gateway, since that's the only thing in the routing table. However the default gateway, 128.6.4.27, will notice that 128.6.4.30 would really be a better route. (How it does that is up to the gateway. However there are some fairly simple methods for a gateway to determine that you would be better off using a different one.) Thus 128.6.4.27 will send back a redirect specifying that datagrams for 128.6.7.23 should be sent via 128.6.4.30. Your TCP/IP software will add a routing entry

```
128.6.7.23 128.6.4.30 UDHG pe0
```

Any future datagrams for 128.6.7.23 will be sent directly to the appropriate gateway.

This strategy would be a complete solution, if it weren't for three problems:

- It requires each computer to have the address of one gateway "hardwired" into its startup files, as the initial default.
- If a gateway goes down, routing table entries using it may not be removed.
- If your network uses subnets, and your TCP/IP implementation does not handle them, this strategy will not work.

How serious the first problem is depends upon your situation. For small networks, there is no problem modifying startup files whenever something changes. But some organizations can find it very painful. If network topology changes, and a gateway is removed, any systems that have that gateway as their default must be adjusted. This is particularly serious if the people who maintain the network are not the same as those maintaining the individual systems. One simple approach is to make sure that the default address never changes. For example, you might adopt the convention that address 1 on each subnet is the default gateway for that subnet. For example, on subnet 128.6.7, the default gateway would always be 128.6.7.1. If that gateway is ever removed, some other gateway is given that address. (There must always be at least one gateway left to give it to. If there isn't, you are completely cut off anyway.)

The biggest problem with the description given so far is that it tells you how to add routes but not how to get rid of them. What happens if a gateway goes down? You want traffic to be redirected back to a gateway that is up. Unfortunately, a gateway that has crashed is not going to issue Redirects. One solution is to choose very reliable gateways. If they crash very seldom, this may not be a problem. Note that Redirects can be used to handle some kinds of network failure. If something fails in a distant part of the network, your current route may no longer be a good one. As long as the gateway to which you are talking is still up and talking to you, it can simply issue a Redirect to the gateway that is now the best one. However you still need a way to detect failure of one of the gateways that you are talking to directly.

The best approach for handling failed gateways is for your TCP/IP implementation to detect routes that have failed. TCP maintains various timers that allow the software to detect when a connection has broken. When this happens, one good approach is to mark the route down, and go back to the default gateway. A similar approach can also be used to handle failures in the default gateway. If you have marked two gateways as default, then the software should be capable of switching when connections using one of them start failing. Unfortunately, some common TCP/IP implementations do not mark routes as down and change to new ones. In particular, Berkeley 4.2 Unix does not. However Berkeley 4.3 Unix does do this, and as other vendors begin to base products on 4.3 rather than 4.2, this ability is expected to become more common.

5.4 Other ways for hosts to find routes

As long as your TCP/IP implementations handle failing connections properly, establishing one or more default routes in the configuration file is likely to be the simplest way to handle routing. However there are two other routing approaches that are worth considering for special situations:

- spying on the routing protocol
- using proxy ARP

5.4.1 Spying on Routing

Gateways generally have a special protocol that they use among themselves. Note that redirects cannot be used by gateways. Redirects are simply ways for gateways to tell "dumb" hosts to use a different gateway. The gateways themselves must have a complete picture of the network, and a way to compute the optimal route to each subnet. Generally they maintain this picture by exchanging information among themselves. There are several different routing protocols in use for this purpose. One way for a computer to keep track of gateways is for it to listen to the gateways' messages among themselves. There is software available for this purpose for most of the common routing protocols. When you run this software, your computer will maintain a complete picture of the network, just as the gateways do. The software is generally designed to maintain your computer's routing tables dynamically, so that datagrams are always sent to the proper gateway. In effect, the routing software issues the equivalent of the Unix "route add" and "route delete" commands as the network topology changes. Generally this results in a complete routing table, rather than one that depends upon default routes. (This assumes that the gateways themselves maintain a complete table. Sometimes gateways keep track of your campus network completely, but use a default route for all off-campus networks, etc.)

Running routing software on each host does in some sense "solve" the routing problem. However there are several reasons why this is not normally recommended except as a last resort. The most serious problem is that this reintroduces configuration options that must be kept up to date on each host. Any computer that wants to participate in the protocol among the gateways will need to configure its software compatibly with the gateways. Modern gateways often have configuration options that are complex compared with those of an individual host. It is undesirable to spread these to every host.

There is a somewhat more specialized problem that applies only to diskless computers. By its very nature, a diskless computer depends upon the network and file servers to load programs and to do swapping. It is dangerous for diskless computers to run any software that listens to network broadcasts. Routing software generally depends upon broadcasts. For example, each gateway on the network might broadcast its routing tables every 30 seconds. The problem with diskless nodes is that the software to listen to these broadcasts must be loaded over the network. On a busy computer, programs that are not used for a few seconds will be swapped or paged out. When they are activated again, they must be swapped or paged in. Whenever a broadcast is sent, every computer on the network needs to activate the routing software in order to process the broadcast. This means that many diskless computers will be doing swapping or paging at the same time. This is likely to cause a temporary overload of the network. Thus it is very unwise for diskless machines to run any software that requires them to listen to broadcasts.

5.4.2 Proxy ARP

Proxy ARP is an alternative technique for letting gateways make all the routing decisions. It is applicable to any broadcast network that uses ARP or a similar technique for mapping Internet addresses into network-specific addresses such as Ethernet addresses. This presentation will assume Ethernet. Other network types can be accommodated if you replace "Ethernet address" with the appropriate network-specific address, and ARP with the protocol used for address mapping by that network type.

In many ways proxy ARP it is similar to using a default route and redirects, however it uses a different mechanism to communicate routes to the host. With redirects, a full routing table is used. At any given moment, the host knows what gateways it is routing datagrams to. With proxy ARP, you dispense with explicit routing tables, and do everything at the level of Ethernet addresses. Proxy ARP can be used for all destinations, only for destinations within your network, or in various combinations. It will be simplest to explain it as used for all addresses. To do this, you instruct the host to pretend that every computer in the world is attached directly to your local Ethernet. On Unix, this would be done using a command

```
route add default 128.6.4.2 0
```

where 128.6.4.2 is assumed to be the IP address of your host. As explained above, the metric of 0 causes everything that matches this route to be sent directly on the local Ethernet. Alternatively, some systems will allow you to get the same effect by setting a subnet mask of 0. If you do this, you may have to take precautions to make sure that it isn't reset by an ICMP subnet mask broadcast by a system that knows the real subnet mask.

When a datagram is to be sent to a local Ethernet destination, your computer needs to know the Ethernet address of the destination. In order to find that, it uses something generally called the ARP table. This is simply a mapping from Internet address to Ethernet address. Here's a typical ARP table. (On our system, it is displayed using the command "arp -a".)

```

FOKKER.RUTGERS.EDU (128.6.5.16) at 8:0:20:0:8:22 temporary
CROSBY.RUTGERS.EDU (128.6.5.48) at 2:60:8c:49:50:63 temporary
CAIP.RUTGERS.EDU (128.6.4.16) at 8:0:8b:0:1:6f temporary
DUDE.RUTGERS.EDU (128.6.20.16) at 2:7:1:0:eb:cd temporary
W2ONS.MIT.EDU (18.70.0.160) at 2:7:1:0:eb:cd temporary
OBERON.USC.EDU (128.125.1.1) at 2:7:1:2:18:ee temporary
gatech.edu (128.61.1.1) at 2:7:1:0:eb:cd temporary
DARTAGNAN.RUTGERS.EDU (128.6.5.65) at 8:0:20:0:15:a9 temporary


```

Note that it is simply a list of IP addresses and the corresponding Ethernet address. The "temporary" indicates that the entry was added dynamically using ARP, rather than being put into the table manually.

If there is an entry for the address in the ARP table, the datagram is simply put on the Ethernet with the corresponding Ethernet address. If not, an "ARP request" is broadcast, asking for the destination host to identify itself. This request is in effect a question "will the host with Internet address 128.6.4.194 please tell me what your Ethernet address is?". When a response comes back, it is added to the ARP table, and future datagrams for that destination can be sent without delay.

This mechanism was originally designed only for use with hosts attached directly to a single Ethernet. If you need to talk to a host on a different Ethernet, it was assumed that your routing table would direct you to a gateway. The gateway would of course have one interface on your Ethernet. Your computer would then end up looking up the address of that gateway using ARP. It would generally be useless to expect ARP to work directly with a computer on a distant network. Since it isn't on the same Ethernet, there's no Ethernet address you can use to send datagrams to it. And when you send an ARP request for it, there's nobody to answer the request.

Proxy ARP is based on the concept that the gateways will act as proxies for distant hosts. Suppose you have a host on network 128.6.5, with address 128.6.5.2. (computer A in diagram below) It wants to send a datagram to host 128.6.4.194, which is on a different Ethernet (subnet 128.6.4). (computer C in diagram below) There is a gateway connecting the two subnets, with address 128.6.5.1 (gateway R):



Now suppose computer A sends an ARP request for computer C. C isn't able to answer for itself. It's on a different network, and never even sees the ARP request. However gateway R can act on its behalf. In effect, your computer asks "will the host with Internet address 128.6.4.194 please tell me what your Ethernet address is?", and the gateway says "here I am, 128.6.4.194 is 2:7:1:0:eb:cd", where 2:7:1:0:eb:cd is actually the Ethernet address of the gateway. This bit of illusion works just fine. Your host now thinks that 128.6.4.194 is attached to the local Ethernet with address 2:7:1:0:eb:cd. Of course it isn't. But it works anyway. Whenever there's a datagram to be sent to 128.6.4.194, your host sends it to the specified Ethernet address. Since that's the address of a gateway R, the gateway gets the datagram. It then forwards it to the destination.

Note that the net effect is exactly the same as having an entry in the routing table saying to route destination 128.6.4.194 to gateway 128.6.5.1:

```
128.6.4.194 128.6.5.1 UGH pe0
```

except that instead of having the routing done at the level of the routing table, it is done at the level of the ARP table.

Generally it's better to use the routing table. That's what it's there for. However here are some cases where proxy ARP makes sense:

- when you have a host that does not implement subnets
- when you have a host that does not respond properly to redirects
- when you do not want to have to choose a specific default gateway
- when your software is unable to recover from a failed route

The technique was first designed to handle hosts that do not support subnets. Suppose that you have a subnetted network. For example, you have chosen to break network 128.6 into subnets, so that 128.6.4 and 128.6.5 are separate. Suppose you have a computer that does not understand subnets. It will assume that all of 128.6 is a single network. Thus it will be difficult to establish routing table entries to handle the configuration above. You can't tell it about the gateway explicitly using "route add 128.6.4.0 128.6.5.1 1". Since it thinks all of 128.6 is a single network, it can't understand that you are trying to tell it where to send one subnet. It will instead interpret this command as an attempt to set up a host route to a host whose address is 128.6.4.0. The only thing that would work would be to establish explicit host routes for every individual host on every other subnet. You can't depend upon default gateways and redirects in this situation either. Suppose you said "route add default 128.6.5.1 1". This would establish the gateway 128.6.5.1 as a default. However the system wouldn't use it to send datagrams to other subnets. Suppose the host is 128.6.5.2, and wants to send a datagram to 128.6.4.194. Since the destination is part of 128.6, your computer considers it to be on the same network as itself, and doesn't bother to look for a gateway.

Proxy ARP solves this problem by making the world look the way the defective implementation expects it to look. Since the host thinks all other subnets are part of its own network, it will simply issue ARP requests for them. It expects to get back an Ethernet address that can be used to establish direct communications. If the gateway is practicing proxy ARP, it will respond with the gateway's Ethernet address. Thus datagrams are sent to the gateway, and everything works.

As you can see, no specific configuration is needed to use proxy ARP with a host that doesn't understand subnets. All you need is for your gateways to implement proxy ARP. In order to use it for other purposes, you must explicitly set up the routing table to cause ARP to be used. By default, TCP/IP implementations will expect to find a gateway for any destination that is on a different network. In order to make them issue ARP's, you must explicitly install a route with metric 0, as in the example "route add default 128.6.5.2 0", or you must set a subnet mask of 0.

It is obvious that proxy ARP is reasonable in situations where you have hosts that don't understand subnets. Some comments may be needed on the other situations. Generally TCP/IP implementations do handle ICMP redirects properly. Thus it is normally practical to set up a default route to some gateway, and depend upon the gateway to issue redirects for destinations that should use a different gateway. However in case you ever run into an implementation that does not obey redirects, or cannot be configured to have a default gateway, you may be able to make things work by depending upon proxy ARP. Of course this requires that you be able to configure the host to issue ARP's for all destinations. You will need to read the documentation carefully to see exactly what routing features your implementation has.


Sometimes you may choose to depend upon proxy ARP for convenience. The problem with routing tables is that you have to configure them. The simplest configuration is simply to establish a default route, but even there you have to supply some equivalent to the Unix command "route add default ...". Should you change the addresses of your gateways, you have to modify this command on all of your hosts, so that they point to the new default gateway. If you set up a default route that depends upon proxy ARP (i.e. has metric 0), you won't have to change your configuration files when gateways change. With proxy ARP, no gateway addresses are given explicitly. Any gateway can respond to the ARP request, no matter what its address.

In order to save you from having to do configuration, some TCP/IP implementations default to using ARP when they have no other route. The most flexible implementations allow you to mix strategies. That is, if you have specified a route for a particular network, or a default route, they will use that route. But if there is no route for a destination, they will treat it as local, and issue an ARP request. As long as your gateways support proxy ARP, this allows such hosts to reach any destination without any need for routing tables.


Finally, you may choose to use proxy ARP because it provides better recovery from failure. This choice is very much dependent upon your implementation. The next section will discuss the tradeoffs in more detail.

In situations where there are several gateways attached to your network, you may wonder how proxy ARP allows you to choose the best one. As described above, your computer simply sends a broadcast asking for the Ethernet address for a destination. We assumed that the gateways would be set up to respond to this broadcast. If there is more than one gateway, this requires coordination among them. Ideally, the gateways will have a complete picture of the network topology. Thus they are able to determine the best route from your host to any destination. If the gateways coordinate among themselves, it should be possible for the best gateway to respond to your ARP request. In practice, it may not always be possible for this to happen. It is fairly easy to design algorithms to prevent very bad routes.

For example, consider the following situation:



1, 2, and 3 are networks. A and B are gateways, connecting network 2 to 1 or 3. If a host on network 2 wants to talk to a host on network 1, it is fairly easy for gateway A to decide to answer, and for gateway B to decide not to. Here's how: if gateway B accepted a datagram for network 1, it would have to forward it to gateway A for delivery. This would mean that it would take a datagram from network 2 and send it right back out on network 2. It is very easy to test for routes that involve this sort of circularity. It is much harder to deal with a situation such as the following:



Suppose a computer on network 1 wants to send a datagram to one on network 2. The route via A and D is probably better, because it goes through only one intermediate network (3). It is also possible to go via B, C, and E, but that path is probably slightly slower. Now suppose the computer on network 1 sends an ARP request for a destination on 2. It is likely that A and B will both respond to that request. B is not quite as good a route as A. However it is not so bad as the case above. B won't have to send the datagram right back out onto network 1. It is unable to determine there is a better alternative route without doing a significant amount of global analysis on the network. This may not be practical in the amount of time available to process an ARP request.

5.4.3 Moving to New Routes After Failures

In principle, IP routing is capable of handling line failures and gateway crashes. There are various mechanisms to adjust routing tables and ARP tables to keep them up to date. Unfortunately, many major implementations of TCP/IP have not implemented all of these mechanisms. The net result is that you have to look carefully at the documentation for your implementation, and consider what kinds of failures are most likely. You then have to choose a strategy that will work best for your site. The basic choices for finding routes have all been listed above: spying on the gateways' routing protocol, setting up a default route and depending upon redirects, and using proxy ARP. These methods all have their own limitations in dealing with a changing network.

Spying on the gateways' routing protocol is theoretically the cleanest solution. Assuming that the gateways use good routing technology, the tables that they broadcast contain enough information to maintain optimal routes to all destinations. Should something in the network change (a line or a gateway goes down), this information will be reflected in the tables, and the routing software will be able to update the hosts' routing tables appropriately. The disadvantages are entirely practical. However in some situations the robustness of this approach may outweigh the disadvantages. To summarize the discussion above, the disadvantages are:

- If the gateways are using sophisticated routing protocols, configuration may be fairly complex. Thus you will be faced with setting up and maintaining configuration files on every host.
- Some gateways use proprietary routing protocols. In this case, you may not be able to find software for your hosts that understands them.
- If your hosts are diskless, there can be very serious performance problems associated with listening to routing broadcasts.

Some gateways may be able to convert from their internal routing protocol to a simpler one for use by your hosts. This could largely bypass the first two disadvantages. Currently there is no known way to get around the third one.

The problems with default routes/redirects and with proxy ARP are similar: they both have trouble dealing with situations where their table entries no longer apply. The only real difference is that different tables are involved. Suppose a gateway goes down. If any of your current routes are using that gateway, you may be in trouble. If you

are depending upon the routing table, the major mechanism for adjusting routes is the redirect. This works fine in two situations:

- where the default gateway is not the best route. The default gateway can direct you to a better gateway
- where a distant line or gateway fails. If this changes the best route, the current gateway can redirect you to the gateway that is now best

The case it does not protect you against is where the gateway that you are currently sending your datagrams to crashes. Since it is down, it is unable to redirect you to another gateway. In many cases, you are also unprotected if your default gateway goes down, since routing starts by sending to the default gateway.

The situation with proxy ARP is similar. If the gateways coordinate themselves properly, the right one will respond initially. If something elsewhere in the network changes, the gateway you are currently issuing can issue a redirect to a new gateway that is better. (It is usually possible to use redirects to override routes established by proxy ARP.) Again, the case you are not protected against is where the gateway you are currently using crashes. There is no equivalent to failure of a default gateway, since any gateway can respond to the ARP request.

So the big problem is that failure of a gateway you are using is hard to recover from. It's hard because the main mechanism for changing routes is the redirect, and a gateway that is down can't issue redirects. Ideally, this problem should be handled by your TCP/IP implementation, using timeouts. If a computer stops getting responses, it should cancel the existing route, and try to establish a new one. Where you are using a default route, this means that the TCP/IP implementation must be able to declare a route as down based on a timeout. If you have been redirected to a non-default gateway, and that route is declared down, traffic will return to the default. The default gateway can then begin handling the traffic, or redirect it to a different gateway. To handle failure of a default gateway, it should be possible to have more than one default. If one is declared down, another will be used. Together, these mechanisms should take care of any failure.

Similar mechanisms can be used by systems that depend upon proxy ARP. If a connection is timing out, the ARP table entry that it uses should be cleared. This will cause a new ARP request, which can be handled by a gateway that is still up. A simpler mechanism would simply be to time out all ARP entries after some period. Since making a new ARP request has a very low overhead, there's no problem with removing an ARP entry even if it is still good. The next time a datagram is to be sent, a new request will be made. The response is normally fast enough that users will not even notice the delay.

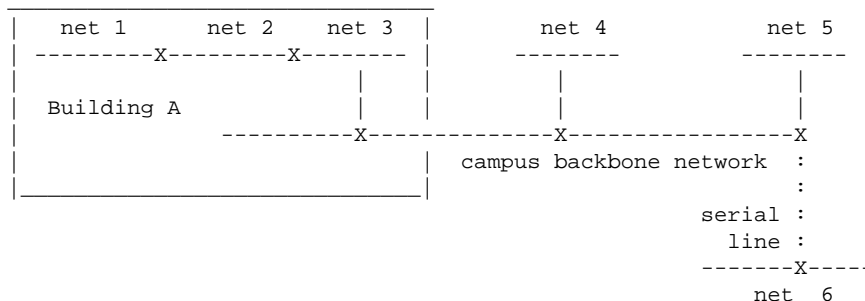
Unfortunately, many common implementations do not use these strategies. In Berkeley 4.2, there is no automatic way of getting rid of any kind of entry, either routing or ARP. They do not invalidate routes or ARP entries based on failures. If gateway crashes are a significant problem, there may be no choice but to run software that listens to the routing protocol. In Berkeley 4.3, routing entries are removed when TCP connections are failing. ARP entries are still not removed. This makes the default route strategy more attractive for 4.3 than proxy ARP. Having more than one default route may also allow for recovery from failure of a default gateway. Note however that 4.3 only handles timeout for connections using TCP. If a route is being used only by services based on UDP, it will not recover from gateway failure. While the "traditional" TCP/IP services use TCP, network file systems generally do not. Thus 4.3-based systems still may not always be able to recover from failure.

In general, you should examine your implementation in detail to determine what sort of error recovery strategy it uses. We hope that the discussion in this section will then help you choose the best way of dealing with routing.

There is one more strategy that some older implementations use. It is strongly discouraged, but we mention it here so you can recognize it if you see it. Some implementations detect gateway failure by taking active measure to see what gateways are up. The best version of this is based on a list of all gateways that are currently in use. (This can be determined from the routing table.) Every minute or so, an echo request datagram is sent to each such gateway. If a gateway stops responding to echo requests, it is declared down, and all routes using it revert to the default. With such an implementation, you normally supply more than one default gateway. If the current default stops responding, an alternate is chosen. In some cases, it is not even necessary to choose an explicit default gateway. The software will randomly choose any gateway that is responding. This implementation is very flexible and recovers well from failures. However a large network full of such implementations will waste a lot of bandwidth on the echo datagrams that are used to test whether gateways are up. This is the reason that this strategy is discouraged.

6. Bridges and Gateways

This section will deal in more detail with the technology used to construct larger networks. It will focus particularly on how to connect together multiple Ethernets, token rings, etc. These days most networks are hierarchical. Individual hosts attach to local-area networks such as Ethernet or token ring. Then those local networks are connected via some combination of backbone networks and point to point links. A university might have a network that looks in part like this:



Nets 1, 2 and 3 are in one building. Nets 4 and 5 are in different buildings on the same campus. Net 6 is in a somewhat more distant location. The diagram above shows nets 1, 2, and 3 being connected directly, with switches that handle the connections being labelled as "X". Building A is connected to the other buildings on the same campus by a backbone network. Note that traffic from net 1 to net 5 takes the following path:

- from 1 to 2 via the direct connection between those networks
- from 2 to 3 via another direct connection
- from 3 to the backbone network
- across the backbone network from building A to the building in which net 5 is housed
- from the backbone network to net 5

Traffic for net 6 would additionally pass over a serial line. With the setup as shown, the same switch is being used to connect the backbone network to net 5 and to the serial line. Thus traffic from net 5 to net 6 would not need to go through the backbone, since there is a direct connection from net 5 to the serial line.


This section is largely about what goes in those "X"s.

6.1 Alternative Designs

Note that there are alternatives to the sort of design shown above. One is to use point to point lines or switched lines directly to each host. Another is to use a single-level of network technology that is capable of handling both local and long-haul networking.

6.1.1 A mesh of point to point lines

Rather than connecting hosts to a local network such as Ethernet, and then interconnecting the Ethernets, it is possible to connect long-haul serial lines directly to the individual computers. If your network consists primarily of individual computers at distant locations, this might make sense. Here would be a small design of that type.



In the design shown earlier, the task of routing datagrams around the network is handled by special-purpose switching units shown as "X"s. If you run lines directly between pairs of hosts, your hosts will be doing this sort of routing and switching, as well as their normal computing. Unless you run lines directly between every pair of computers, some systems will end up handling traffic for others. For example, in this design, traffic from 1 to 3 will go

through 4, 5 and 6. This is certainly possible, since most TCP/IP implementations are capable of forwarding datagrams. If your network is of this type, you should think of your hosts as also acting as gateways. Much of the discussion below on configuring gateways will apply to the routing software that you run on your hosts. This sort of configuration is not as common as it used to be, for two reasons:

- Most large networks have more than one computer per location. In this case it is less expensive to set up a local network at each location than to run point to point lines to each computer.
- Special-purpose switching units have become less expensive. It often makes sense to offload the routing and communications tasks to a switch rather than handling it on the hosts.

It is of course possible to have a network that mixes the two kinds of technology. In this case, locations with more equipment would be handled by a hierarchical system, with local-area networks connected by switches. Remote locations with a single computer would be handled by point to point lines going directly to those computers. In this case the routing software used on the remote computers would have to be compatible with that used by the switches, or there would need to be a gateway between the two parts of the network.

Design decisions of this type are typically made after an assessment of the level of network traffic, the complexity of the network, the quality of routing software available for the hosts, and the ability of the hosts to handle extra network traffic.

6.1.2 Circuit switching technology

Another alternative to the hierarchical LAN/backbone approach is to use circuit switches connected to each individual computer. This is really a variant of the point to point line technique, where the circuit switch allows each system to have what amounts to a direct line to every other system. This technology is not widely used within the TCP/IP community, largely because the TCP/IP protocols assume that the lowest level handles isolated datagrams. When a continuous connection is needed, higher network layers implement it using datagrams. This datagram-oriented technology does not match a circuit-oriented environment very closely. In order to use circuit switching technology, the IP software must be modified to be able to build and tear down virtual circuits as appropriate. When there is a datagram for a given destination, a virtual circuit must be opened to it. The virtual circuit would be closed when there has been no traffic to that destination for some time. The major use of this technology is for the DDN (Defense Data Network). The primary interface to the DDN is based on X.25. This network appears to the outside as a distributed X.25 network. TCP/IP software intended for use with the DDN must do precisely the virtual circuit management just described. Similar techniques could be used with other circuit-switching technologies, e.g. ATT's DataKit, although there is almost no software currently available to support this.

6.1.3 Single-level networks

In some cases new developments in wide-area networks can eliminate the need for hierarchical networks. Early hierarchical networks were set up because the only convenient network technology was Ethernet or other LAN's, and those could not span distances large enough to cover an entire campus. Thus it was necessary to use serial lines to connect LAN's in various locations. It is now possible to find network technology whose characteristics are similar to Ethernet, but where a single network can span a campus. Thus it is possible to think of using a single large network, with no hierarchical structure.

The primary limitations of a large single-level network are performance and reliability considerations. If a single network is used for the entire campus, it is very easy to overload it. Hierarchical networks can handle a larger traffic volume than single-level networks if traffic patterns have a reasonable amount of locality. That is, in many applications, traffic within an individual department tends to be greater than traffic among departments.

Let's look at a concrete example. Suppose there are 10 departments, each of which generates 1 Mbit/sec of traffic. Suppose further than 90% of that traffic is to other systems within the department, and only 10% is to other departments. If each department has its own network, that network only needs to handle 1 Mbit/sec. The backbone network connecting the department also only needs 1 Mbit/sec capacity, since it is handling 10% of 1 Mbit from each department. In order to handle this situation with a single wide-area network, that network would have to be able to handle the simultaneous load from all 10 departments, which would be 10 Mbit/sec.

However this example was carefully constructed to be favorable to the hierarchical design. If more of the traffic in the department is going to other departments, then the backbone will need a higher bandwidth. For example, suppose that a campus has a few centralized resources, e.g. mainframes and other large systems in a computing center. If most of the network traffic is from small systems attempting to get to the central system, then the argument above does not work. In this case a hierarchy may still be useful. However it doesn't reduce the bandwidth

required for the long-haul network. In the example above, if all 10 departments communicated primarily with systems at the computer center, the backbone would have to be able to carry all of their traffic, 10Mbits per second. The computer center would either attach its systems directly to the backbone, or it would have a "departmental" network with a capacity of 10Mbits per second rather than the 1Mbits per second needed by the other departments.

The second limitation on single-level networks is reliability, maintainability and security. Wide-area networks are more difficult to diagnose and maintain than local-area networks, because problems can be introduced from any building to which the network is connected. They also make traffic visible in all locations. For these reasons, it is often sensible to handle local traffic locally, and use the wide-area network only for traffic that actually must go between buildings. However if you have a situation where each location has only one or two computers, it may not make sense to set up a local network at each location, and a single-level network may make sense.

6.1.4 Mixed designs

In practice, few large networks have the luxury of adopting a theoretically pure design.

It is very unlikely that any large network will be able to avoid using a hierarchical design. Suppose we set out to use a single-level network. Even if most buildings have only one or two computers, there will be some location where there are enough that a local-area network is justified. The result is a mixture of a single-level network and a hierarchical network. Most buildings have their computers connected directly to the wide-area network, as with a single-level network. However in one building there is a local-area network which uses the wide-area network as a backbone, connecting to it via a switching unit.

On the other side of the story, even network designers with a strong commitment to hierarchical networks are likely to find some parts of the network where it simply doesn't make economic sense to install a local-area network. So a host is put directly onto the backbone network, or tied directly to a serial line.

However you should think carefully before making ad hoc departures from your design philosophy in order to save a few dollars. In the long run, network maintainability is going to depend upon your ability to make sense of what is going on in the network. The more consistent your technology is, the more likely you are to be able to maintain the network.

6.2 An introduction to alternative switching technologies

This section will discuss the characteristics of various technologies used to switch datagrams between networks. In effect, we are trying to fill in some details about the black boxes assumed in previous sections. There are three basic types of switches, generally referred to as repeaters, bridges, and gateways, or alternatively as level 1, 2 and 3 switches (based on the level of the OSI model at which they operate). Note however that there are systems that combine features of more than one of these, particularly bridges and gateways.

The most important dimensions on which switches vary are isolation, performance, routing and network management facilities. These will be discussed below.

The most serious difference is between repeaters and the other two types of switch. Until recently, gateways provided very different services from bridges. However these two technologies are now coming closer together. Gateways are beginning to adopt the special-purpose hardware that has characterized bridges in the past. Bridges are beginning to adopt more sophisticated routing, isolation features, and network management, which have characterized gateways in the past. There are also systems that can function as both bridge and gateway. This means that at the moment, the crucial decision may not be to decide whether to use a bridge or a gateway, but to decide what features you want in a switch and how it fits into your overall network design.

6.2.1 Repeaters

A repeater is a piece of equipment that connects two networks that use the same technology. It receives every data packet on each network, and retransmits it onto the other network. The net result is that the two networks have exactly the same set of packets on them. For Ethernet or IEEE 802.3 networks there are actually two different kinds of repeater. (Other network technologies may not need to make this distinction.)

A simple repeater operates at a very low level indeed. Its primary purpose is to get around limitations in cable length caused by signal loss or timing dispersion. It allows you to construct somewhat larger networks than you would otherwise be able to construct. It can be thought of as simply a two-way amplifier. It passes on individual bits in the signal, without doing any processing at the packet level. It even passes on collisions. That is, if a collision is generated on one of the networks connected to it, the repeater generates a collision on the other network.

There is a limit to the number of repeaters that you can use in a network. The basic Ethernet design requires that signals must be able to get from one end of the network to the other within a specified amount of time. This determines a maximum allowable length. Putting repeaters in the path does not get around this limit. (Indeed each repeater adds some delay, so in some ways a repeater makes things worse.) Thus the Ethernet configuration rules limit the number of repeaters that can be in any path.


A "buffered repeater" operates at the level of whole data packets. Rather than passing on signals a bit at a time, it receives an entire packet from one network into an internal buffer and then retransmits it onto the other network. It does not pass on collisions. Because such low-level features as collisions are not repeated, the two networks continue to be separate as far as the Ethernet specifications are concerned. Thus there are no restrictions on the number of buffered repeaters that can be used. Indeed there is no requirement that both of the networks be of the same type. However the two networks must be sufficiently similar that they have the same packet format. Generally this means that buffered repeaters can be used between two networks of the IEEE 802.x family (assuming that they have chosen the same address length and maximum packet size), or two networks of some other related family. A pair of buffered repeaters can be used to connect two networks via a serial line.

Buffered repeaters share with simple repeaters the most basic feature: they repeat every data packet that they receive from one network onto the other. Thus the two networks end up with exactly the same set of packets on them.

6.2.2 Bridges and gateways

A bridge differs from a buffered repeater primarily in the fact that it exercises some selectivity as to what datagrams it forwards between networks. Generally the goal is to increase the capacity of the system by keeping local traffic confined to the network on which it originates. Only traffic intended for other networks goes through the bridge. So far this description would also apply to a gateway. Bridges and gateways differ in the way they determine what datagrams to forward. A bridge uses only the OSI level 2 address. In the case of Ethernet or IEEE 802.x networks, this is the 6-byte Ethernet or MAC-level address. (The term "MAC-level address" is more general. However for the sake of concreteness, examples in this section will assume that Ethernet is being used. You may generally replace the term "Ethernet address" with the equivalent MAC-level address for other similar technologies.) A bridge does not examine the datagram itself, so it does not use the IP address or its equivalent for routing decisions. In contrast, a gateway bases its decisions on the IP address, or its equivalent for other protocols.

There are several reasons why it matters which kind of address is used for decisions. The most basic is that it affects the relationship between the switch and the upper layers of the protocol. If forwarding is done at the level of the MAC-level address (bridge), the switch will be invisible to the protocols. If it is done at the IP level, the switch will be visible. Let's give an example. Here are two networks connected by a bridge:




Note that the bridge does not have an IP address. As far as computers A, B, and C are concerned, there is a single Ethernet (or other network) to which they are all attached. This means that the routing tables must be set up so that computers on both networks treat both networks as local. When computer A opens a connection to computer B, it first broadcasts an ARP request asking for computer B's Ethernet address. The bridge must pass this broadcast from network 1 to network 2. (In general, bridges must pass all broadcasts.) Once the two computers know each other's Ethernet addresses, communications use the Ethernet address as the destination. At that point, the bridge can start exerting some selectivity. It will only pass datagrams whose Ethernet destination address is for a machine on the other network. Thus a datagram from B to A will be passed from network 2 to 1, but a datagram from B to C will be ignored.

In order to make this selection, the bridge needs to know which network each machine is on. Most modern bridges build up a table for each network to which they are connected, listing the Ethernet addresses of machines known to be on that network. They do this by watching all of the datagrams on each network. When a datagram first appears on network 1, it is reasonable to conclude that the Ethernet source address corresponds to a machine on network 1.

Note that a bridge must look at every datagram on the Ethernet, for two different reasons. First, it may use the source address to learn which machines are on which network. Second, it must look at the destination address in order to decide whether it needs to forward the datagram to the other network.

As mentioned above, generally bridges must pass broadcasts from one network to the other. Broadcasts are often used to locate a resource. The ARP request is a typical example of this. Since the bridge has no way of knowing what host is going to answer the broadcast, it must pass it on to the other network. Some bridges have user-selectable filters. With them, it is possible to block some broadcasts and allow others. You might allow ARP broadcasts (which are essential for IP to function), but confine less essential broadcasts to one network. For example, you might choose not to pass rwhod broadcasts, which some systems use to keep track of every user logged into every other system. You might decide that it is sufficient for rwhod to know about the systems on a single segment of the network.

Now let's take a look at two networks connected by a gateway:



Note that the gateway has IP addresses assigned to each interface. The computers' routing tables are set up to forward through appropriate address. For example, computer A has a routing entry saying that it should use the gateway 128.6.5.1 to get to subnet 128.6.4.

Because the computers know about the gateway, the gateway does not need to scan all the packets on the Ethernet. The computers will send datagrams to it when appropriate. For example, suppose computer A needs to send a message to computer B. Its routing table will tell it to use gateway 128.6.5.1. It will issue an ARP request for that address. The gateway will respond to the ARP request, just as any host would. From then on, datagrams destined for B will be sent with the gateway's Ethernet address.

6.2.3 More about bridges

There are several advantages to using the MAC-level address, as a bridge does. First, every packet on an Ethernet or IEEE network has such an address. The address is in the same place for every packet, whether it is IP, DECnet, or some other protocol. Thus it is relatively fast to get the address from the packet. A gateway must decode the entire IP header, and if it is to support protocols other than IP, it must have software for each such protocol. This means that a bridge automatically supports every possible protocol, whereas a gateway requires specific provisions for each protocol it is to support.

However there are also disadvantages. The one that is intrinsic to the design of a bridge is

- A bridge must look at every packet on the network, not just those addressed to it. Thus it is possible to overload a bridge by putting it on a very busy network, even if very little traffic is actually going through the bridge.

However there is another disadvantage that is based on the way bridges are usually built. It is possible in principle to design bridges that do not have this disadvantage, but I don't know of any plans to do so. It stems from the fact that bridges do not have a complete routing table that describes the entire system of networks. They simply have a list of the Ethernet addresses that lie on each of its networks. This means

- Networks that use bridges cannot have loops in them. If there were a loop, some bridges would see traffic from the same Ethernet address coming from both directions, and would be unable to decide which table to put that address in. Note that any parallel paths to the same destination constitute a loop. This means that multiple paths cannot be used for purposes of splitting the load or providing redundancy.

There are some ways of getting around the problem of loops. Many bridges allow configurations with redundant connections, but turn off links until there are no loops left. Should a link fail, one of the disabled ones is then brought back into service. Thus redundant links can still buy you extra reliability. But they can't be used to provide extra capacity. It is also possible to build a bridge that will make use of parallel point to point lines, in the one special case where those lines go between a single pair of bridges. The bridges would treat the two lines as a single virtual line, and use them alternately in round-robin fashion.

The process of disabling redundant connections until there are no loops left is called a "spanning tree algorithm". This name comes from the fact that a tree is defined as a pattern of connections with no loops. Thus one wants to disable connections until the connections that are left form a tree that "spans" (includes) all of the networks in the system. In order to do this, all of the bridges in a network system must communicate among themselves. There is an IEEE proposal to standardize the protocol for doing this, and for constructing the spanning tree.

Note that there is a tendency for the resulting spanning tree to result in high network loads on certain parts of the system. The networks near the "top of the tree" handle all traffic between distant parts of the network. In a network that uses gateways, it would be possible to put in an extra link between parts of the network that have heavy traffic between them. However such extra links cannot be used by a set of bridges.

6.2.4 More about gateways

Gateways have their own advantages and disadvantages. In general a gateway is more complex to design and to administer than a bridge. A gateway must participate in all of the protocols that it is designed to forward. For example, an IP gateway must respond to ARP requests. The IP standards also require it to completely process the IP header, decrementing the time to live field and obeying any IP options.

Gateways are designed to handle more complex network topologies than bridges. As such, they have a different (and more complex) set of decisions to make. In general a bridge has fairly simple decision to make: should it forward a datagram, and if so which interface should it send the datagram out? When a gateway forwards a datagram, it must decide what host or gateway to send the datagram to next. If the gateway sends a datagram back onto the same network it came from, it should also issue a redirect to the source of the datagram telling it to use a better route. Many gateways can also handle parallel paths. If there are several equally good paths to a destination, the gateway will alternate among them in round-robin fashion. (This is done by some bridges also, though it is less common there. In both cases, there are some issues raised by round-robin alternation. It tends to lead to datagrams arriving in an order different than the order in which they were sent. This can complicate processing by the destination host. Some older TCP/IP implementations have bugs in handling out of order datagrams.)

In order to handle these decisions, a gateway will typically have a routing table that looks very much like a host's. As with host routing tables, a gateway's table contains an entry for each possible network number. For each network, there is either an entry saying that that network is connected directly to the gateway, or there is an entry saying that traffic for that network should be forwarded through some other gateway or gateways. We will describe the "routing protocols" used to build up this information later, in the discussion on how to configure a gateway.

6.3 Comparing the switching technologies

Repeaters, buffered repeaters, bridges, and gateways form a spectrum. Those devices near the beginning of the list are best for smaller networks. They are less expensive, and easier to set up, but less general. Those near the end of the list are suitable for building more complex networks. Many networks will contain a mixture of switch types, with repeaters being used to connect a few nearby network segments, bridges used for somewhat larger areas, and gateways used for long-distance links.

Note that this document so far has assumed that only gateways are being used. The section on setting up a host described how to set up a routing table listing the gateways to use to get to various networks. Repeaters and bridges are invisible to IP. So as far as previous sections are concerned, networks connected by them are to be considered a single network. Section 3.4 describes how to configure a host in the case where several subnets are carried on a single physical network. The same configuration should be used when several subnets are connected by repeaters or bridges.

As mentioned above, the most important dimensions on which switches vary are isolation, performance, routing, network management.

6.3.1 Isolation

Generally people use switches to connect networks to each other. So they are normally thinking of gaining connectivity, not providing isolation. However isolation is worth thinking about. If you connect two networks and provide no isolation at all, then any network problems on other networks suddenly appear on yours as well. Also, the two networks together may have enough traffic to overwhelm your network. Thus it is well to think of choosing an appropriate level of protection.

Isolation comes in two kinds: isolation against malfunctions and traffic isolation. In order to discuss isolation of malfunctions, we have to have a taxonomy of malfunctions. Here are the major classes of malfunctions, and which switches can isolate them:

- Electrical faults, e.g. a short in the cable or some sort of fault that distorts the signal. All types of switch will confine this to one side of the switch: repeater, buffered repeater, bridge, gateway. These are worth protecting against, although their frequency depends upon how often your cables are changed or disturbed. It is rare for this sort of fault to occur without some disturbance of the cable.
- Transceiver and controller problems that generate signals that are valid electrically but nevertheless incorrect (e.g. a continuous, infinitely long packet, spurious collisions, never dropping carrier). All except the simple repeater will confine this: buffered repeater, bridge, gateway. (Such problems are not very common.)
- Software malfunctions that lead to excessive traffic between particular hosts (i.e. not broadcasts). Bridges and gateways will isolate these. (This type of failure is fairly rare. Most software and protocol problems generate broadcasts.)
- Software malfunctions that lead to excessive broadcast traffic. Gateways will isolate these. Generally bridges will not, because they must pass broadcasts. Bridges with user-settable filtering can protect against some broadcast malfunctions. However in general bridges must pass ARP, and most broadcast malfunctions involve ARP. This problem is not severe on single-vendor networks where software is under careful control. However sites with complex network environments or experimental network software may see problems of this sort regularly.

Traffic isolation is provided by bridges and gateways. The most basic decision is how many computers can be put onto a network without overloading its capacity. This requires knowledge of the capacity of the network, but also how the hosts will use it. For example, an Ethernet may support hundreds of systems if all the network is used for remote logins and an occasional file transfer. However if the computers are diskless, and use the network for swapping, an Ethernet will support between 10 and 40, depending upon their speeds and I/O rates.

When you have to put more computers onto a network than it can handle, you split it into several networks and put some sort of switch between them. If you do the split correctly, most of the traffic will be between machines on the same piece. This means putting clients on the same network as their servers, putting terminal servers on the same network as the hosts that they access most commonly, etc.

Bridges and gateways generally provide similar degrees of traffic isolation. In both cases, only traffic bound for hosts on the other side of the switch is passed. However see the discussion on routing.

6.3.2 Performance

Absolute performance limits are becoming less of an issue as time goes on, since the switching technology is improving. Generally repeaters can handle the full bandwidth of the network. (By their very nature, a simple repeater must be able to do so.) Bridges and gateways often have performance limitations of various sorts. Bridges have two numbers of interest: packet scanning rate and throughput. As explained above, a bridge must look at every packet on the network, even ones that it does not forward. The number of packets per second that it can scan in this way is the packet scanning rate. Throughput applies to both bridges and gateways. This is the rate at which they can forward traffic. Generally this depends upon datagram size.

Normally the number of datagrams per second that a unit can handle will be greater for short datagrams than long ones. Early models of bridge varied from a few hundred datagrams per second to around 7000. The higher speeds are for equipment that uses special-purpose hardware to speed up the process of scanning packets. First-generation gateways varied from a few hundred datagrams per second to 1000 or more. However second-generation gateways are now available, using special-purpose hardware of the same sophistication as that used by bridges. They can handle on the order of 10000 datagrams per second. Thus at the moment high-performance bridges and gateways can switch most of the bandwidth of an Ethernet. This means that performance should no longer be a basis for choosing between types of switch. However within a given type of switch, there are still specific models with higher or lower capacity. And there may still be differences in price/performance. This is particularly true at the low end. The least expensive bridges are currently less than half the price of the least expensive gateway.

Unfortunately there is no single number on which you can base performance estimates. The figure most commonly quoted is packets per second. Be aware that most vendors count a datagram only once as it goes through a gateway, but that one prominent vendor counts datagrams twice. Thus their switching rates must be deflated by a factor of 2.

Also, when comparing numbers make sure that they are for datagrams of the same size. A simple performance model is

$$\text{processing time} = \text{switching time} + \text{datagram size} * \text{time per byte}$$

That is, the time to switch a datagram is normally a constant switching time, representing interrupt latency, header processing, routing table lookup, etc., plus a component proportional to datagram size, representing the time needed to do any datagram copying. One reasonable approach to reporting performance is to give datagrams per second for minimum and maximum size datagrams. Another is to report limiting switching speed in datagrams per second and throughput in bytes per second, i.e. the two terms of the equation above.

6.3.3 Routing

Routing refers to the technology used to decide where to send a datagram next. Of course for a repeater this is not an issue, since repeaters forward every packet.

The routing strategy for a bridge turns into two decisions: (1) enabling and disabling links in order to maintain the spanning tree, and (2) deciding whether it should forward any particular packet, and out what interface (if the bridge is capable of handling more than two interfaces). The second decision is usually based on a table of MAC-level addresses. As described above, this is built up by scanning traffic visible from each interface. The goal is to forward those packets whose destination is on the other side of the bridge. This algorithm requires that the network configuration have no loops or redundant lines. Less sophisticated bridges leave this up to the system designer. With these bridges, you must set up your network so that there are no loops in it. More sophisticated bridges allow arbitrary topology, but disable links until no loops remain. This provides extra reliability. If a link fails, an alternative link will be turned on automatically. Bridges that work this way have a protocol that allows them to detect when a unit must be disabled or reenabled, so that at any instant the set of active links forms a "spanning tree". If you require the extra reliability of redundant links, make sure that the bridges you use can disable and enable themselves in this way. There is currently no official standard for the protocol used among bridges, although there is a standard in the proposal stage. If you buy bridges from more than one vendor, make sure that their spanning-tree protocols will interoperate.

Gateways generally allow arbitrary network topologies, including loops and redundant links. Because of their more general routing algorithms, gateways must maintain a model of the entire network topology. Different routing techniques maintain models of greater or lesser complexity, and use the data with varying degrees of sophistication. Gateways that handle IP should generally support the two Internet standard routing protocols: RIP (Routing Information Protocol) and EGP (External Gateway Protocol). EGP is a special-purpose protocol for use in networks where there is a backbone under a separate administration. It allows exchange of reachability information with the backbone in a controlled way. If you are a member of such a network, your gateway must support EGP. This is becoming common enough that it is probably a good idea to make sure that all gateways support EGP.

RIP is a protocol designed to handle routing within small to moderate size networks, where line speeds do not differ radically. Its primary limitations are:

- It cannot be used with networks where any path goes through more than 15 gateways. This range may be further reduced if you use an optional feature for giving a slow line a weight larger than one.
- It cannot share traffic between parallel lines (although some implementations allow this if the lines are between the same pair of gateways).
- It cannot adapt to changes in network load.
- It is not well suited to situations where there are alternative routes through lines of very different speeds.
- It may not be stable in networks where lines or gateways change a lot.

Some vendors supply proprietary modifications to RIP that improve its operation with EGP or increase the maximum path length beyond 15, but do not otherwise modify it very much. If you expect your network to involve gateways from more than one vendor, you should generally require that all of them support RIP, since this is the only routing protocol that is generally available. If you expect to use a more sophisticated protocol in addition, it may be useful for the gateways to translate between their own protocol and RIP. However for very large or complex networks, there may be no choice but to use some other protocol throughout.

More sophisticated routing protocols are possible. The primary ones being considered today are cisco System's IGRP, and protocols based on the SPF (shortest-path first) algorithms. In general these protocols are designed for larger or more complex networks. They are in general stable under a wider variety of conditions, and they can handle arbitrary combinations of line type and speed. Some of them allow you to split traffic among parallel paths, to get better overall throughput. Some newer technologies may allow the network to adjust to take into account paths that are overloaded. However at the moment I do not know of any commercial gateway that does this. (There are very serious problems with maintaining stable routing when this is done.) There are enough variations among routing technology, and it is changing rapidly enough, that you should discuss your proposed network topology in detail with all of the vendors that you are considering. Make sure that their technology can handle your topology, and can support any special requirements that you have for sharing traffic among parallel lines, and for adjusting topology to take into account failures. In the long run, we expect one or more of these newer routing protocols to attain the status of a standard, at least on a de facto basis. However at the moment, there is no generally implemented routing technology other than RIP.

One additional routing topic to consider is policy-based routing. In general routing protocols are designed to find the shortest or fastest possible path for every datagram. In some cases, this is not desired. For reasons of security, cost accountability, etc., you may wish to limit certain paths to certain uses. Most gateways now have some ability to control the spread of routing information so as to give you some administrative control over the way routes are used. Different gateways vary in the degree of control that they support. Make sure that you discuss any requirements that you have for control with all prospective gateway vendors.

6.3.4 Network management

Network management covers a wide variety of topics. In general it includes gathering statistical data and status information about parts of your network, and taking action as necessary to deal with failures and other changes. The most primitive technique for network monitoring is periodic "pinging" of critical hosts. Pinging is a monitoring technique that depends on an "echo" datagram. This is a specific type of datagram that requests an immediate reply. Most TCP/IP implementations contain a program (usually called "ping") that sends an echo to a specified host. If you get a reply, you know that the host is up, and that the network connection to the host works. If you don't get a reply, you know that something is wrong with one of the other. By pinging a reasonable sample of hosts, you can normally tell what is going on. If all the hosts on a network suddenly stop returning pings, it is reasonable to conclude that the connection to that network has gone bad. If one host stops returning pings, but other hosts on the same network still do, then it is reasonable to conclude that the host has crashed.

More sophisticated network monitoring requires the ability to get specific status and statistical information from various devices on the network. These should include various sorts of datagram counts, as well as counts of errors of various kinds. This data is likely to be most detailed in a gateway, since the gateway classifies datagrams using the protocols, and may even respond to certain types of datagram itself. However bridges and even buffered repeaters can certainly have counts of datagrams forwarded, interface errors, etc. It should be possible to collect this data from a central monitoring point.

There is now an official TCP/IP approach to network monitoring. The first stages use a related set of protocols, SGMP and SNMP. Both of these protocols are designed to allow you to collect information and to make changes in configuration parameters for gateways and other entities on your network. You can run the corresponding interface programs on any host in your network. SGMP is now available for several commercial gateways, as well as for Unix systems that are acting as gateways. There is a limited set of information which any SGMP implementation is required to supply, as well as a uniform mechanism for vendors to add information of their own. By late 1988, the second generation of this protocol, SNMP, should be in service. This is a slightly more sophisticated protocol. It has with it a more complete set of information that can be monitored, called the MIB (Management Information Base). Unlike the somewhat ad hoc collection of SGMP variables, the MIB is the result of numerous committee deliberations involving a number of vendors and users. Eventually it is expected that there will be a TCP/IP equivalent of CMIS, the ISO network monitoring service. However CMIS, and its protocols, CMIP, are not yet official ISO standards, so they are still in the experimental stages.

In general terms all of these protocols accomplish the same thing: They allow you to collect critical information in a uniform way from all vendors' equipment. You send commands as UDP datagrams from a network management program running on some host in your network. Generally the interaction is fairly simple, with a single pair of datagrams exchanged: a command and a response. At the moment security is fairly simple. It is possible to require what amounts to a password in the command. (In SGMP it is referred to as a "session name", rather than a password.) More elaborate, encryption-based security is being developed.

You will probably want to configure the network management tools at your disposal to do several different things. For short-term network monitoring, you will want to keep track of switches crashing or being taken down for

maintenance, and of failure of communications lines and other hardware. It is possible to configure SGMP and SNMP to issue "traps" (unsolicited messages) to a specified host or list of hosts when some of these critical events occur (e.g. lines up and down). However it is unrealistic to expect a switch to notify you when it crashes. It is also possible for trap messages to be lost due to network failure or overload. Thus you can't depend completely on traps. You should also poll your switches regularly to gather information. Various displays are available, including a map of your network where items change color as their status changes, and running "strip charts" that show datagram rates and other items through selected switches. This software is still in its early stages, so you should expect to see a lot of change here. However at the very least you should expect to be notified in some way of failures. You may also want to be able to take actions to reconfigure the system in response to failures, although security issues make some managers nervous about doing that through the existing management protocols.

The second type of monitoring you are likely to want to do is to collect information for use in periodic reports on network utilization and performance. For this, you need to sample each switch periodically, and retrieve numbers of interest. At Rutgers we sample hourly, and get the number of datagrams forwarded for IP and DECnet, a count of reloads, and various error counts. These are reported daily in some detail. Monthly summaries are produced giving traffic through each gateway, and a few key error rates chosen to indicate a gateway that is being overloaded (datagrams dropped in input and output).

It should be possible to use monitoring techniques of this kind with most types of switch. At the moment, simple repeaters do not report any statistics. Since they do not generally have processors in them, doing so would cause a major increase in their cost. However it should be possible to put network management software in buffered repeaters, bridges, and gateways. Gateways are the most likely to contain sophisticated network management software. Most gateway vendors that handle IP are expected to implement the monitoring protocols described above. Many bridge vendors make some provisions for collecting performance data. Since bridges are not protocol-specific, most of them do not have the software necessary to implement TCP/IP-based network management protocols. In some cases, monitoring can be done only by typing commands to a directly-attached console. (We have seen one case where it is necessary to take the bridge out of service to gather this data.) In other cases, it is possible to gather data via the network, but the monitoring protocol is ad hoc or even proprietary.

Except for very small networks, you should probably insist that any switch more complex than a simple repeater should collect statistics and provide some way of querying them remotely. Portions of the network that do not support such operations can be monitored by pinging. However ping simply detects gross failures. It does not allow you to look at the noise level of a serial line and other quantities needed to do high-quality maintenance. In the long run, you can expect the most software to be available for standard protocols such as SGMP/SNMP and CMIS. However proprietary monitoring tools may be sufficient as long as they work with all of the equipment that you have.

6.3.5 A final evaluation

Here is a summary of the places where each kind of switch technology is normally used:

- Repeaters are normally confined to a single building. Since they provide no traffic isolation, you must make sure that the entire set of networks connected by repeaters can carry the traffic from all of the computers on it. Since they generally provide no network monitoring tools, you will not want to use repeaters for a link that is likely to fail.
- Bridges and gateways should be placed sufficiently frequently to break your network into pieces for which the traffic volume is manageable. You may want to place bridges or gateways even in places where traffic level alone would not require them for network monitoring reasons.
- Because bridges must pass broadcast datagrams, there is a limit to the size network you can construct using them. It is probably a good idea to limit the network connected by bridges to a hundred systems or so. This number can be increased somewhat for bridges with good facilities for filtering.
- Because certain kinds of network misbehavior will be passed, bridges should be used only among portions of the network where a single group is responsible for diagnosing problems. You have to be crazy to use a bridge between networks owned by different organizations. Portions of your network where experiments are being done in network technology should always be isolated from the rest of the network by gateways.
- For many applications it is more important to choose a product with the right combination of performance, network management tools, and other features than to make the decision between bridges and gateways.

7. Configuring Gateways

This section deals with configuration issues that are specific to gateways. Gateways that handle IP are themselves Internet hosts. Thus the discussions above on configuring addresses and routing information apply to gateways as well as to hosts. The exact way you configure a gateway will depend upon the vendor. In some cases, you edit files stored on a disk in the gateway itself. However for reliability reasons most gateways do not have disks of their own. For them, configuration information is stored in non-volatile memory or in configuration files that are uploaded from one or more hosts on the network.

At a minimum, configuration involves specifying the IP address and address mask for each interface, and enabling an appropriate routing protocol. However generally a few other options are desirable. There are often parameters in addition to the IP address that you should set for each interface.

One important parameter is the broadcast address. As explained above, older software may react badly when broadcasts are sent using the new standard broadcast address. For this reason, some vendors allow you to choose a broadcast address to be used on each interface. It should be set using your knowledge of what computers are on each of the networks. In general if the computers follow current standards, a broadcast address of 255.255.255.255 should be used. However older implementations may behave better with other addresses, particularly the address that uses zeros for the host number. (For the network 128.6 this would be 128.6.0.0. For compatibility with software that does not implement subnets, you would use 128.6.0.0 as the broadcast address even for a subnet such as 128.6.4.) You should watch your network with a network monitor and see the results of several different broadcast address choices. If you make a bad choice, every time the gateway sends a routing update broadcast, many machines on your network will respond with ARP's or ICMP errors. Note that when you change the broadcast address in the gateway, you may need to change it on the individual computers as well. Generally the idea is to change the address on the systems that you can configure to give behavior that is compatible with systems that you can't configure.

Other interface parameters may be necessary to deal with peculiarities of the network it is connected to. For example, many gateways test Ethernet interfaces to make sure that the cable is connected and the transceiver is working correctly. Some of these tests will not work properly with the older Ethernet version 1 transceivers. If you are using such a transceiver, you would have to disable this keepalive testing. Similarly, gateways connected by a serial line normally do regular testing to make sure that the line is still working. There can be situations where this needs to be disabled.

Often you will have to enable features of the software that you want to use. For example, it is often necessary to turn on the network management protocol explicitly, and to give it the name or address of a host that is running software to accept traps (error messages).

Most gateways have options that relate to security. At a minimum, this may include setting password for making changes remotely (and the "session name" for SGMP). If you need to control access to certain parts of your network, you will also need to define access control lists or whatever other mechanism your gateway uses.

Gateways that load configuration information over the network present special issues. When such a gateway boots, it sends broadcast requests of various kinds, attempting to find its Internet address and then to load configuration information. Thus it is necessary to make sure that there is some computer that is prepared to respond to these requests. In some cases, this is a dedicated micro running special software. In other cases, generic software is available that can run on a variety of machines. You should consult your vendor to make sure that this can be arranged. For reliability reasons, you should make sure that there is more than one host with the information and programs that your gateways need. In some cases you will have to maintain several different files. For example, the gateways used at Rutgers use a program called "bootp" to supply their Internet address, and they then load the code and configuration information using TFTP. This means that we have to maintain a file for bootp that contains Ethernet and Internet addresses for each gateway, and a set of files containing other configuration information for each gateway. If your network is large, it is worth taking some trouble to make sure that this information remains consistent. We keep master copies of all of the configuration information on a single computer, and distribute it to other systems when it changes, using the Unix utilities `make` and `rdist`. If your gateway has an option to store configuration information in non-volatile memory, you will eliminate some of these logistical headaches. However this presents its own problems. The contents of non-volatile memory should be backed up in some central location. It will also be harder for network management personnel to review configuration information if it is distributed among the gateways.

Starting a gateway is particularly challenging if it loads configuration information from a distant portion of the network. Gateways that expect to take configuration information from the network generally issue broadcast requests on all of the networks to which they are connected. If there is a computer on one of those networks that is prepared to respond to the request, things are straightforward. However some gateways may be in remote locations where

there are no nearby computer systems that can support the necessary protocols. In this case, it is necessary to arrange for the requests to be routed back to the network where there are appropriate computers. This requires what is strictly speaking a violation of the basic design philosophy for gateways. Generally a gateway should not allow broadcasts from one network to pass through to an adjacent network. In order to allow a gateway to get information from a computer on a different network, at least one of the gateways in between will have to be configured to pass the particular class of broadcasts used to retrieve this information. If you have this sort of configuration, you should test the loading process regularly. It is not unusual to find that gateways do not come up after a power failure because someone changed the configuration of another gateway and made it impossible to load some necessary information.

7.1 Configuring routing for gateways

The final topic to be considered is configuring routing. This is more complex for a gateway than for a normal host. Most TCP/IP experts recommend that routing be left to the gateways. Thus hosts may simply have a default route that points to the nearest gateway. Of course the gateways themselves can't get by with this. They need to have complete routing tables.

In order to understand how to configure a gateway, we have to look in a bit more detail at how gateways communicate routes. When you first turn on a gateway, the only networks it knows about are the ones that are directly connected to it. (They are specified by the configuration information.) In order to find out how to get to more distant parts of the network, it engages in some sort of "routing protocol". A routing protocol is simply a protocol that allows each gateway to advertise which networks it can get to, and to spread that information from one gateway to the next. Eventually every gateway should know how to get to every network. There are different styles of routing protocol. In one common type, gateways talk only to nearby gateways. In another type, every gateway builds up a database describing every other gateway in the system. However all of the protocols have some way for each gateway in the system to find out how to get to every destination.

A metric is some number or set of numbers that can be used to compare routes. The routing table is constructed by gathering information from other gateways. If two other gateways claim to be able to get to the same destination, there must be some way of deciding which one to use. The metric is used to make that decision. Metrics all indicate in some general sense the "cost" of a route. This may be a cost in dollars of sending datagrams over that route, the delay in milliseconds, or some other measure. The simplest metric is just a count of the number of gateways along the path. This is referred to as a "hop count". Generally this metric information is set in the gateway configuration files, or is derived from information appearing there.

At a minimum, routing configuration is likely to consist of a command to enable the routing protocol that you want to use. Most vendors will have a preferred routing protocol. Unless you have some reason to choose another, you should use that. The normal reason for choosing another protocol is for compatibility with other kinds of gateway. For example, your network may be connected to a national backbone network that requires you to use EGP (exterior gateway protocol) to communicate routes with it. EGP is only appropriate for that specific case. You should not use EGP within your own network, but you may need to use it in addition to your regular routing protocol to communicate with a national network. If your own network has several different types of gateway, then you may need to pick a routing protocol that all of them support. At the moment, this is likely to be RIP (Routing Information Protocol). Depending upon the complexity of your network, you could use RIP throughout it, or use a more sophisticated protocol among the gateways that support it, and use RIP only at the boundary between gateways from different vendors.

Assuming that you have chosen a routing protocol and turned it on, there are some additional decisions that you may need to make. One of the more basic configuration options has to do with supplying metric information. As indicated above, metrics are numbers which are used to decide which route is the best. Unsophisticated routing protocols, e.g. RIP, normally just count hops. So a route that passes through 2 gateways would be considered better than one that passes through 3. Of course if the latter route used 1.5Mbps lines and the former 9600 bps lines, this would be the wrong decision. Thus most routing protocols allow you to set parameters to take this sort of thing into account. With RIP, you would arrange to treat the 9600 bps line as if it were several hops. You would increase the effective hop count until the better route was chosen. More sophisticated protocols may take the bit rate of the line into account automatically. However you should be on the lookout for configuration parameters that need to be set. Generally these parameters will be associated with the particular interface. For example, with RIP you would have to set a metric value for the interface connected to the 9600 bps line. With protocols that are based on bit rate, you might need to specify the speed of each line (if the gateway cannot figure it out automatically).

Most routing protocols are designed to let each gateway learn the topology of the entire network, and to choose the best possible route for each datagram. In some cases you may not want to use the "best" route. You may want traffic to stay out of a certain portion of the network for security or cost reasons. One way to institute such controls

is by specifying routing options. These options are likely to be different for different vendors. But the basic strategy is that if the rest of the network doesn't know about a route, it won't be used. So controls normally take the form of limiting the spread of information about routes whose use you want to control.

Note that there are ways for the user to override the routing decisions made by your gateways. If you really need to control access to a certain network, you will have to do two separate things:

- Use routing controls to make sure that the gateways use only the routes you want them to.
- Use access control lists on the gateways that are adjacent to the sensitive networks.

These two mechanisms act at different levels. The routing controls affect what happens to most datagrams: those where the user has not specified routing manually. Your routing mechanism must be set up to choose an acceptable route for them. The access control list provides an additional limitation which prevents users from supplying their own routing and bypassing your controls.

For reliability and security reasons, there may also be controls to allow you to list the gateways from which you will accept information. It may also be possible to rank gateways by priority. For example, you might decide to listen to routes from within your own organization before routes from other organizations or other parts of the organization. This would have the effect of having traffic use internal routes in preference to external ones, even if the external ones appear to be better.

If you use several different routing protocols, you will probably have some decisions to make regarding how much information to pass among them. Since multiple routing protocols are often associated with multiple organizations, you must be sure to make these decisions in consultation with management of all of the relevant networks. Decisions that you make may have consequences for the other network which are not immediately obvious. You might think it would be best to configure the gateway so that everything it knows is passed on by all routing protocols. However here are some reasons why you may not want to do so:

- The metrics used by different routing protocols may not be comparable. If you are connected to two different external networks, you want to specify that one should always be used in preference to the other, or that the nearest one should be used, rather than attempting to compare metric information received from the two networks to see which has the better route.
- EGP is particularly sensitive, because the EGP protocol cannot handle loops. Thus there are strict rules governing what information may be communicated to a backbone that uses EGP. In situations where EGP is being used, management of the backbone network should help you configure your routing.
- If you have slow lines in your network (9600 bps or slower), you may prefer not to send a complete routing table throughout the network. If you are connected to an external network, you may prefer to treat it as a default route, rather than to inject all of its routing information into your routing protocol.