Vikas Aggarwal
JvNCnet Computer Network
Request for Comments: 1291
December 1991

Mid-Level Networks
Potential Technical Services

Status of this Memo

This FYI RFC provides information for the Internet community. It does not specify an Internet standard. Distribution of this memo is unlimited.

Abstract

This document proposes a set of technical services that each Internet mid-level network can offer within the mid-level network itself and and to its peer networks. The term *"mid-level"* is used as a generic term to represent all regional and similar networks, which, due to continuous evolutions and transitions, can no longer be termed 'regional'. [MAN]. It discusses the pros and cons of offering these services, as well as areas in which mid-level networks can work together.

A large portion of the ideas stem from discussions at the IETF Operational Statistics (OPstat), User Connectivity Problems (UCP) and Network Joint Management (NJM) working groups.

Contents

1 Introduction **3**

2 The Generic Model **3**

3 Technical Services **4**

3.1 Domain Name Service 4

3.2 Public Domain Software 5

3.3 Network Time 6

3.4 Network News 6

3.5 Mailing Lists 7

4 Experimental Testbeds **7**

5 Network Information Services **8**

6 Network Operations **8**

7 References **9**

8 Security Considerations **10**

9 Author's Address **10**

A Appendix A - Mailing Lists **12**

B Appendix B - DNS Architecture Strategy **12**

1 Introduction

Over the past few years, the Internet has grown to be a very large entity and its dependability is critical to its users. Furthermore, due to the size and nature of the network, the trend has been to decentralize as many network functions (such as domain name-service, whois, *etc.*) as possible. Efforts are being made in resource discovery [SHHH90] so that the work of researchers is not lost in the volumes of data that is available on the Internet.

A side result of this growth has been the logical structure imposed in the Internet of networks classified by *function*. Tangible examples in the present state are the NSFnet national backbone, the mid-level/regional networks and campus networks. Each of these can be viewed as hierarchies within an organization, each serving a slightly different function than the other (campus LANs providing access to local resources, mid-level networks providing access to remote resources, etc.). The functions of each hierarchy then become the *'services'* offered to the organizational layer below it, who in turn depend on these services.

This document proposes a set of basic technical services that could be offered by a mid-level network. These services would not only increase the robustness of the mid-level network itself, but would also serve to structure the distribution of *resources* and *services* within the Internet. It also proposes a uniform naming convention for locating the hosts offering these services.

2 The Generic Model

The Internet model that is used as the basis for this document is a graph of mid-level networks connected to one another, each in turn connecting the campus/organization networks and with the end users attached to the campus networks. The model assumes that the mid-level networks constitute the highest level of functional division within the Internet hierarchy described above (this could change in the unforeseen future). With this model in perspective, this document addresses the objectives of minimizing unnecessary traffic within the Internet as well as making the entire structure as robust as possible.

The proposed structure is a derived extension of organizational LANs where certain services are offered within the organizational LAN itself, such as nameservice, mail, shared files, single or hierarchical points of contact for problems, etc.

The following are the services that are discussed as possible functions of a mid-level network:

- **Technical services**
- **Experimental sites** for testing and dissemination of new software and technology to end sites on the network

In addition, the following services are mentioned briefly which are discussed in detail elsewhere [SSM91, ML91]:

- **Network Operation** services and the interaction between different mid-level networks in this area
- **Network Information services**

3 Technical Services

The Internet has grown to be an essential entity because of the services that it offers to its end users. The list of services is long and growing, but some services are more widely used and deployed than others. This section attempts to list and discuss those technical services that could help a mid-level network provide robust and improved services to its end sites.

3.1 Domain Name Service

According to the NSFnet traffic statistics collected for May 1991, about 7% of the packets on the NSFnet backbone were domain nameserver (DNS) packets. This is a significant amount of traffic, and since most of the other network applications depend on this service, a robust DNS service is critical to any Internet site.

Proper location of secondary nameservers so that they are located on different physical networks can increase the reliability of this service to a large extent [MOC87a, MOC87b]. However, the nature of the service requires that the nameservers for the next highest level be available in order to resolve names outside of one's domain. Thus, for 'foo.princeton.edu' to resolve 'a.mid.net', the *root* nameservers which point to mid.net's nameservers have to be reachable.

To make the service more reliable, the mid-level network could have at least one nameserver that is able to resolve nameserver queries for all domains directly connected to it. Thus, in the event that the entire mid-level network becomes isolated from the rest of the Internet, applications can still resolve queries for sites directly connected to the mid-level network. Without this functionality, there is no way of

resolving a name if the root (or higher level) nameservers become unreachable, even if the query is for a site that is directly connected and reachable.

Strategies for implementing this architecture are discussed in appendix B.

To locate such a "meta-domain" server within a mid-level network, it is proposed that a nameserver entry for 'meta-dns' exist within the mid-level network's domain.

3.2 Public Domain Software

File transfer traffic constituted 23% of the NSFnet backbone traffic for May 1991. Public shareware is a very valuable resource within the Internet and a considerable amount of effort is being put into developing applications to track all available resources in the public archives[SHHH90].

It would be difficult, if not impossible to create an up-to-date repository for every public domain package available on the Internet, simply because of the volume of software and the rate at which new software is being developed every day. Some hosts have gained popularity as good public archives (such as uunet.uu.net, sumex-aim.stanford.edu, wuarchive.wustl.edu) and new developers tend to distribute the software to these sites as distribution points. The economics of maintaining centralized archives is another deterrent to centralization (the UUnet archives at uunet.uu.net take up roughly 1GB of disk storage).

Recently however, a number of methods for resource discovery have been developed and are available on the Internet ('ftp-list' file compiled by John Granose-*odin@pilot.njin.net*, Archie at *archie.cs.mcgill.ca* and Prospero [NEU]).

It is desirable that the mid-level networks be able to provide up-to-date *pointers* to the distribution hosts for available public software archives. Coordinating the distribution of a static list is difficult (though not impossible) and the use of automated resource discovery mechanisms such as Archie and Prospero is recommended. Under ideal conditions, any software that is popular and significant (*e.g.* X11, TeX, RFC's) could be archived and distributed within the mid-level network, but measuring "popularity" and "significance" are debatable and left for further evaluation. Furthermore, a nameserver entry for host 'sw-dist' within the domain can provide information on the various available alternatives for software distribution and discovery (static file location, pointers to Archie servers, *etc.*)– this nameserver entry can be an alias for a CNAME or a TXT entry.

3.3 Network Time

An important feature of any computer network providing distributed services is the capability to synchronize the local clocks on the various systems in the network. Ideally, the clocks of all the reference sources would be synchronized to national standards by wire or radio. The importance and immense popularity of this service makes Network Time a very useful potential service that can be provided by a mid-level network. No specific *protocol* for maintaining time is proposed, and any available protocol that maintains time with reasonable accuracy could be used.

Network Time Protocol (NTP) traffic constituted 1% of the NSFnet traffic during May 1991. The traffic might seem insignificant, but there have been instances where a particular stratum-1 timeserver (*e.g.* one of the stratum-1 servers at University of Delaware) has reached a point of overload with too many different sites trying to peer with it.

It is proposed that at least one stratum-1 and two stratum-2 servers be located *within* a mid-level network (the selection of three servers is based on the NTP standards documentation [MIL89]). Note that the servers can be located at any of the directly connected sites in the network as long as they are publicly accessible. All sites connected to the mid-level network can then coordinate their system times with the servers within the mid-level network itself. Besides increasing the reliability of the timekeeping network, this approach would also limit the load on each timeserver.

For locating the network time servers within a domain, nameserver entries for 'timekeeper-*x*' (where x= 1,2,3..) can be made within the domain. The servers are numbered in order of preference and accuracy. Thus, 'timekeeper-1.foo.net' would be the primary timekeeper and 'timekeeper-2.foo.net' would be additional (possibly secondary) timekeepers within domain 'foo.net'. If such hosts are not available within a domain, a TXT entry pointing to other recommended time-servers could be provided instead.

3.4 Network News

Network News (or Usenet News) constituted 14% of the NSFnet traffic in May 1991. Netnews is an expensive service, both in terms of disk and CPU power, as well as network bandwidth consumed.

The present structure of Network News consists of several hub sites which are distributed over the Internet. End sites get news feeds from other sites, and an article gets injected into the news stream by sending it to the nearest "upstream" site, which then forwards it to its connected news sites, and so on. There is no preset norm for finding a site willing to provide a news feed, and it usually ends up being a site

with whom the site administrator happens to be acquainted. However, this could easily result in some sites not being able to get an economical news feed from within the mid-level network and actually having to derive the feed from a site located on another mid-level network.

A mid-level network could alleviate such occurrences by being able to provide a newsfeed to any or all of its directly connected end sites. Though an expensive resource, some of the costs can be moderated by acting as a *transit* news feeder so that the news needn't be stored for a long time on disk. The software for providing the news feed is not specific and depends entirely on the newsfeed provider.

3.5 Mailing Lists

Internet mailing lists are another popular source of information in parallel to Network News. However, like public software, there is no central repository of all the possible mailing lists available on the Internet, and it would require considerable effort to compile one (at the time of writing this document, a fairly comprehensive list is available on the Internet and mentioned in appendix A.

At this time, there is no clear strategy for distributing or maintaining mailing lists. However, it can be very expensive for a site to distribute mail to all individual end users directly, and if a clear strategy for maintaining a list of mailing-lists can be devised, then mail exploders can be set up at the mid-level networks, each of which forwards the mail to exploders at the end sites. This mechanism would reduce the load on the originating systems, and provides a clean path for tracking down mailer problems. Also, in order to prevent bounced mail from propagating back to the originator of the message, the mailing lists should be set up in a way so that bounced mail goes to the the "owner" of the list and not to the originator of the mail message.

A list of major mailing lists for the services discussed in this document are listed in appendix A.

4 Experimental Testbeds

Due to the working relationships that they have with their end sites and peer networks, the mid-level networks are very good media for distribution of new ideas and technology. Examples of this function are the *White Pages pilot project* [RS90] established by NYSERnet, the NSAP routing schema for OSI transitioning [CGC91], etc.

The mid-level networks could establish cooperative experimental testbeds for testing and deployment of new technologies similar to the ones mentioned above. Besides deployment and testing of new technology, this could also serve to provide a "help" service to the end-sites and to get them started with the new software.

The exact interaction *between* the mid-level networks in this area is not very clear. It is complicated by competition for members between the mid-level networks and needs to be discussed further.

5 Network Information Services

There are a variety of new and useful user services available on the Internet that are difficult to document and provide a comprehensive list of. Some attempt has been made at documenting such resources [NNS] and a mid-level network can be the initial point of contact for distribution of such information on a wide basis. The information can be disseminated in a more controlled and complete manner using this hierarchical approach if each mid-level network maintains up-to-date information about its directly connected sites. Network Information services (NIC) also make the network easier and more attractive to end users. Examples of these services are:

- provide information resources
  - security advisory messages
  - list of library catalogs [GL91]
  - geographical information servers
  - password generators
- resolve end user problems (user support)

These services are NIC related and discussed in detail elsewhere [SSM91]. For accessibility information, an entry for 'nic' could exist in the DNS for the domain (this could be a TXT entry listing email or phone number information for users or other NIC's).

6 Network Operations

The Network Operation Center's (NOC's) at the mid-level networks need to cooperate with each other to resolve network problems. In the event of a network problem

between two mid-level networks or if an end-site has trouble getting to any host, the mid-level network NOCs can serve to be the initial point of contact. The procedures for interaction among NOCs and the formats for exchange of trouble-tickets between the NOCs are described elsewhere [JOH91, ML91].

It is important for cooperating NOCs to have contact information for their directly connected campus/organizational sites and also about their peer mid-level networks. A distributed mechanism for maintaining contact information could be implemented by using a nameserver TXT entry for 'noc' or by maintaining *'finger'* information for user 'noc@domain' or 'noc@noc.domain'. A NOC 'phonebook' listing the contact information for the various NOCs can be used as a static non-distributed mechanism (it is understood that the phonebook can contain outdated information, but the distributed mechanisms can provide correct and updated NOC information provided that the hosts are reachable at the desired time). If it is undesirable to publish the phone number or email address of the NOC for any reason, an entry saying *"unpublished"* (or words to that effect) could exist in the nameserver or *'finger'* entry instead.

7 References

[BOG] K.J Dunlap, M.J. Karels *"Nameserver Operations Guide for Bind Release 4.8"*, CSRG, Department of Electrical Engineering and Computer Sciences, University of California, Berkeley, California.

[CCI88] CCITT Blue Book *"X.500 Series Recommendations"*, ITU, 1989.

[CGC91] R. Collela, E. Gardner, R.Callon *"Guidelines for OSI NSAP Allocation in the Internet"*, RFC 1237, USC Information Sciences Institute, July 1991.

[SSM91] D. Sitzler, P. Smith, A. Marine *"Building a Network Information Services Infrastructure"*, RFC in preparation.

[GL91] A.S. George, R. Larsen *"Internet Accessible Library Catalogs & Databases"*, Aug 1991.
Available via anonymous FTP from 'ariel.unm.edu'.

[JOH91] D.S. Johnson *"NOC TT Requirements"*, RFC in preparation.

[MAN] R. Mandelbaum, P.A. Mandelbaum *"The Strategic Future of the Mid-Level Networks"*, University of Rochester, NY, 1991.

[MOC87a] P. Mockapetris *"Domain Names - Implementation and Specification"*, RFC 1035, USC Information Sciences Institute, Nov 1987.

[MOC87b] P.Mockapetris *"Domain Names - Concepts and Facilities"*, RFC 1034, USC Information Sciences Institute, Nov 1987.

[MIL89] D. Mills *"Network Time Protocol"*, RFC 1129, USC Information Sciences Institute, Oct 1989.

[ML91] M. Mathis, D. Long *"User Connectivity Problems Working Group"*, RFC in preparation.

[NEU] B.C. Neuman *"The Virtual System Model: A Scalable Approach to Organizing Large Systems"*, Department of Computer Science, University of Washington, FR-35, Seattle, WA, May 1990.

[NNS] NSF Network Service Center *"Internet Resource Guide"*, Cambridge, MA.
Available via anonymous FTP from 'nnsc.nsf.net'.

[RS90] M.T. Rose, M. Schoffstall *"The NYSERnet White Pages Pilot Project"*, NYSERnet, Inc., Mar 1990.

[SHHH90] M.F. Schwartz, D.R. Hardy, W.K. Heinzman, G.C. Hirschowitz *"Supporting Resource Discovery Among Public Internet Archives"*, Department of Computer Science, University of Colorado, Boulder, CO. Sep 1990.

8 Security Considerations

Security issues are not discussed in this memo.

9 Author's Address

Vikas Aggarwal
JvNCnet
6 von Neumann Hall
Princeton University
Princeton, NJ 08544

Phone: +1-609-258-2403
Email: vikas@jvnc.net

A Appendix A - Mailing Lists

The following is a list of popular mailing lists for the services listed in this document. To subscribe to a particular mailing list, send a request to *'mailing-list*-request' (do *not* send a request to the entire mailing list).

- ietf@isi.edu: The general mailing list for the Internet Engineering Task Force. This group is concerned with the evolution and development of Internet related protocols and standards. Old mail is archived at 'venera.isi.edu' in directory ftp/irg/ietf.
- ntp@trantor.umd.edu: For discussions on the Network Time Protocol (NTP).
- namedroppers@nic.ddn.mil: Mailing list for discussions on DNS topics. Old mail is archived at 'nic.ddn.mil'.

At the time of writing this document, a list of mailing lists on the Internet is available via anonymous FTP from host 'ftp.nisc.sri.com' in the file 'netinfo/interest-groups'.

B Appendix B - DNS Architecture Strategy

This section discusses practical strategies for implementing a nameserver architecture within a mid-level network, so that it can resolve nameserver queries for all domains directly attached to it.

In order to resolve queries for all directly connected networks, a host that is authoritative for all directly attached domains will need to exist within the mid-level network. Nameservers at the end sites would then treat this "group-of-domains" nameserver as a *forwarding* server to resolve all non-local queries. This can be done by adding a line to the named.boot file on the end site nameservers such as:

```
forwarders 128.121.50.7 128.32.0.4
```

This method has the added advantage that the forwarding server builds up a very rich cache of data [BOG] and acts like a meta-cache that all hosts can benefit from. Note that the forwarding server is queried only if the end-site server cannot service a query locally– hence the "meta-domain" server is not overloaded with queries for *all* nameserver lookups.