IEN 162

W-NOTE 16

TRANSPORT, ADDRESSING, AND ROUTING IN THE WIDEBAND NET

Wideband Net Working Note #16

John A. Pershing Jr.

October 1980

Bolt Beranek and Newman Inc.
50 Moulton Street
Cambridge, Mass. Cambridge, Mass. 02238

TRANSPORT, ADDRESSING, AND ROUTING IN THE WIDEBAND NET

This note proposes a model for addressing and routing in the DARPA Wideband Satellite Experiment. The purpose of the organization described herein is twofold: (1) to hide the physical structure of the Wideband Net from The Internet and its gateways; and (2) to unify the transport and routing functions performed within the Wideband Net.

Certain terms, defined in a glossary at the end of this memo, are used with specific, somewhat non-standard meanings in order to avoid ambiguities.

1. Current Organization

The Wideband Satellite Experiment involves the development of the PSAT satellite network, several local networks, and connections to several existing networks such as the ARPANET. Currently, no plan exists for the organization of these parts into a unified communication medium. So far, the Wideband Net is a collection of independent, sovereign networks, each with its own transport protocol, addressing, and routing schemes. This sovereignty of the constituent networks is a feature which should be preserved as much as possible, so that the development of local and satellite network technologies can proceed without artificial constraints.

Figure 1 illustrates a hypothetical organization sometime in the future, after the network has expanded somewhat. Each site has one PSAT, and perhaps a Voice Funnel and/or mini-concentrator; additionally, sites have a number of local networks with various interconnections. Hosts may be connected directly to a PSAT, a Voice Funnel, or to a local network; some hosts may have connections to two or more nets, etc.

A uniform plan for communication in the Wideband Net will avoid ad hoc schemes involving specialized interface machines which

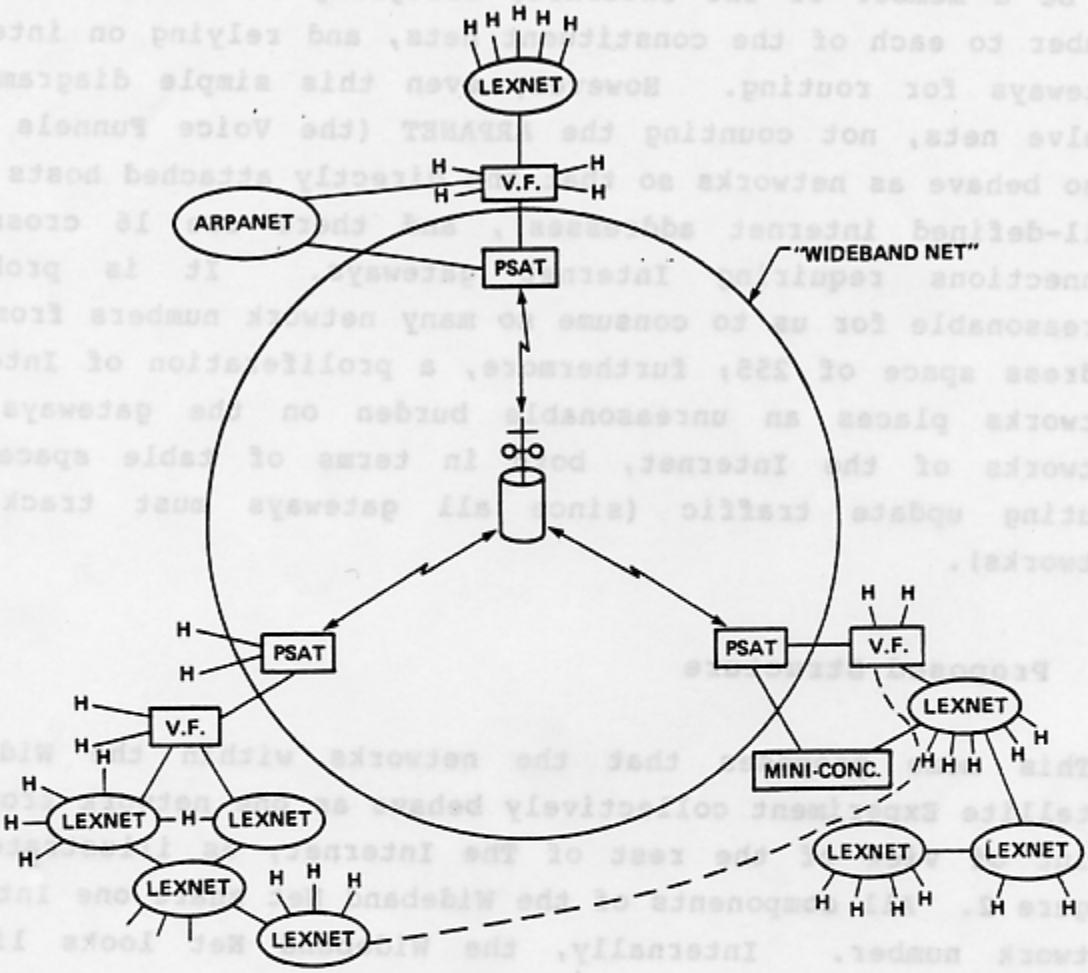

Figure 1: Current Organization of the Wideband Net

transform one net's local protocol and addressing into that of another net. Using specialized gateways in combination with source routing will not take full advantage of the topology and dynamics of the situation. (This is particularly evident if the link indicated by the dotted line exists, since only the Voice Funnel at the right side of the figure knows the status of both the dotted link and the satellite link, and only that Voice Funnel is able to choose the appropriate link.) Such ad hoc schemes are inflexible, inefficient in terms of manpower (since large amounts of special-purpose code must be implemented), and

do not allow the Wideband Net to be readily integrated into the Internet.

A "simple" approach is to consider every component in Figure 1 to be a member of The Internet, assigning an Internet Network Number to each of the constituent nets, and relying on internet gateways for routing. However, even this simple diagram has twelve nets, not counting the ARPANET (the Voice Funnels must also behave as networks so that the directly attached hosts have well-defined internet addresses), and there are 16 cross-net connections requiring Internet gateways. It is probably unreasonable for us to consume so many network numbers from the address space of 255; furthermore, a proliferation of Internet Networks places an unreasonable burden on the gateways and networks of the Internet, both in terms of table space and routing update traffic (since all gateways must track all networks).

2. Proposed Structure

This memo proposes that the networks within the Wideband Satellite Experiment collectively behave as one network from the point of view of the rest of The Internet, as illustrated in Figure 2. All components of the Wideband Net share one Internet network number. Internally, the Wideband Net looks like a catenet; this structure was suggested by Vint Cerf in IEN-48 [1]. The various subnets are interconnected by gateways; each subnet maintains its own autonomy, and hosts that are only involved in local communication can ignore the catenet (and Internet) aspects of the arrangement.

We will adopt an internal addressing and routing scheme which is transparent to the Internet Protocol (IP) [2], so that all hosts on the Wideband Net will have well-defined Internet addresses; we adopt IP as the Wideband Net's "catenet transport protocol", and superimpose a fine-structure on the 24-bit local

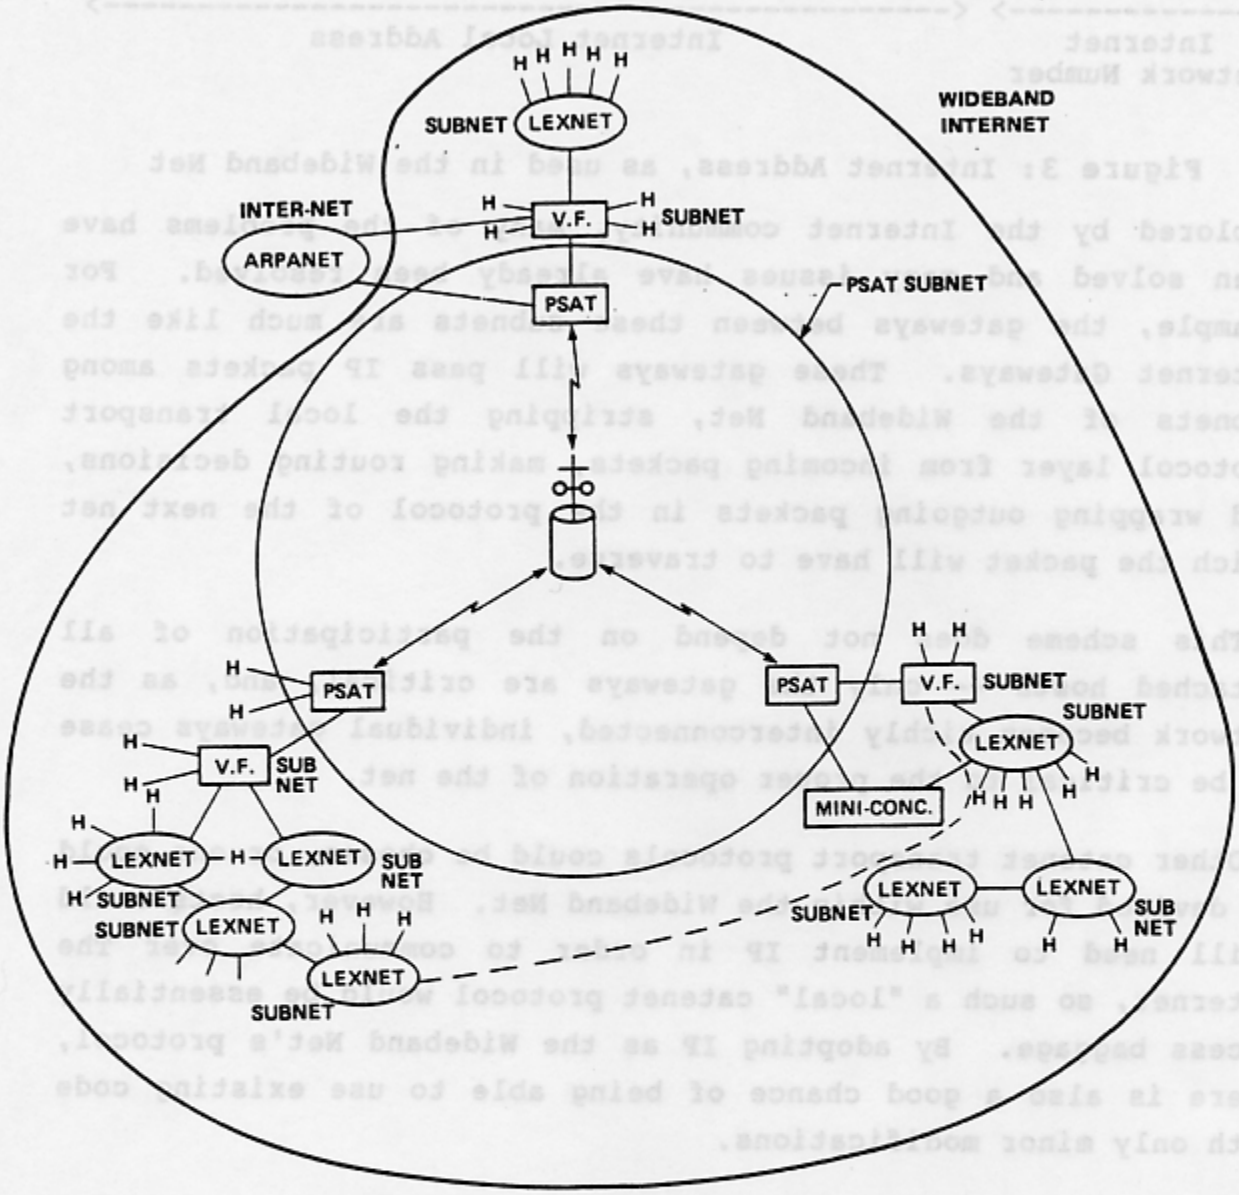

Figure 2: Proposed Organization of the Wideband Net

address part of the (32-bit) Internet address (see Figure 3). The high-order 8 bits specify a subnet of the Wideband Net (e.g., a specific LexNet), and the remaining 16 bits specify the local address on that net.

Since the model we are adopting is one that has been thoroughly

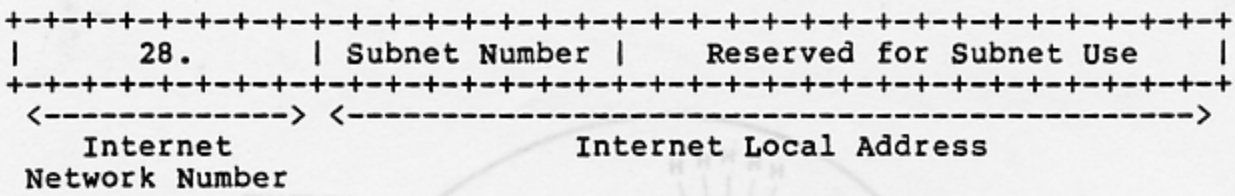

Figure 3: Internet Address, as used in the Wideband Net

explored by the Internet community, many of the problems have been solved and many issues have already been resolved. For example, the gateways between these subnets are much like the Internet Gateways. These gateways will pass IP packets among subnets of the Wideband Net, stripping the local transport protocol layer from incoming packets, making routing decisions, and wrapping outgoing packets in the protocol of the next net which the packet will have to traverse.

This scheme does not depend on the participation of all attached hosts -- only the gateways are critical, and, as the network becomes richly interconnected, individual gateways cease to be critical to the proper operation of the net.

Other catenet transport protocols could be chosen, or one could be devised for use within the Wideband Net. However, hosts would still need to implement IP in order to communicate over The Internet, so such a "local" catenet protocol would be essentially excess baggage. By adopting IP as the Wideband Net's protocol, there is also a good chance of being able to use existing code with only minor modifications.

The issue of what is an Internet network (with its own assigned Internet network number) and what is a subnet is more of a managerial problem than a technical one. The relevant issue is the partitioning of The Internet, in terms of name space and the burden of the routing algorithm on the gateways, as well as in terms of managerial responsibility. The proposed scheme for the Wideband Net is quite flexible in this regard. If one of the subnets becomes an Internet network, then the gateways will

perform somewhat different routing with respect to that network, and the other hosts of the Wideband Net will be mostly indifferent to the change (except that the addresses of the hosts on the "promoted" subnet will change). Presumably, any Wideband Net gateways which were connected to the net will be "promoted" to Internet gateways.

While we believe that this proposed organization should be adopted, it has several problems. However, these problems also exist in The Internet and are treated more fully in various Internet Experiment Notes; they are only briefly mentioned here.

The name space is becoming crowded; using 8 bits for the subnet number and 16 bits for the local host address is perhaps the wrong partition. However, this choice seems to be a good tradeoff between the potential size of the Wideband Net (measured in subnets) and the addressing requirements of any individual subnet. Eventually, the issue of name-space size will have to be addressed by the Internet community as a whole.

Since the fine structure of the Wideband Net is not known to The Internet, it is possible that Internet Gateways will make non-optimal routing decisions with respect to the Wideband Net. This is the penalty which must be paid for trying to minimize the number of Internet Networks. If the Internet Gateways which are connected to the Wideband Net are also Wideband Net gateways (that is, they participate in Wideband Net routing as well as Internet routing), then they may be able to fine-tune the routing of Internet packets through the use of advisory messages exchanged with the other Internet Gateways.

3. Transport Protocol Layers

For purposes of discussion, we define the notion of catenet adjacency. Two hosts on some catenet are adjacent if they are connected to the same constituent network. Conversely, two hosts are non-adjacent if they are connected to different constituent

networks of some catenet. Note that hosts with interfaces on more than one network (such as a gateway) may be both adjacent and non-adjacent to a given host; in fact, such a host is non-adjacent to itself by this definition.

The transport, addressing, and routing schemes suggested in this memo are intended to be used by non-adjacent hosts on the Wideband Net. Of course this does not prohibit adjacent hosts from communicating with these protocols; however, such hosts have the option of communicating using protocols which are local to the network to which they are attached (and, for efficiency reasons, such hosts will probably exercise this option).

Datagrams which are to be transmitted through the Wideband Net are wrapped in a layer of catenet protocol, which is common to the entire Wideband Net, followed by a (possibly null) layer of protocol which is dictated by the particular network that the datagram is traversing. The catenet protocol header is considered to be part of the datagram, and is preserved (with only minor changes) as the datagram traverses the various constituent networks. The layer of local protocol is volatile, and will be discarded as soon as the datagram exits the network defining that particular local protocol.

For compatibility with The Internet, the catenet protocol used is the DoD Standard Internet Protocol (IP) [2]. This permits hosts of the Wideband Net to communicate with hosts on other Internet Networks without resorting to yet another layer of protocol. Initially, no internet gateways will be provided on the Wideband Net, so that a restricted subset of IP may be implemented by the various hosts. The restrictions of the initial implementation follow:

- Fragmentation and Reassembly: Not implemented. Total length may not exceed 576 octets; we believe that all participating hosts, gateways, and networks will support Internet datagrams of this length without fragmentation. The "Flags" field will always be set to '010' (binary) to inhibit fragmentation; the "Fragment Offset" field will always be zero.

- Source/Destination Addresses: The Internet Network Number (first 8 bits) will always be '28' (decimal). The Source/Destination Local Address (remaining 24 bits) is further interpreted as being 8 bits of network number and 16 bits of local address on that network.
- Options: Strictly optional. Any options which are present must be accounted for by the Internet Header Length (IHL) and Header Checksum fields; however, no host is required to interpret any options.

The local protocol is dependent on the particular network(s) to which a host is attached -- there is potentially a different local protocol for every constituent network. In addition, a single network may have more than one link-level protocols, depending on the particular type of port to which the host is attached. Issues of local protocols are of no concern to the purposes of this memo; the implementor is referred to the (possibly nonexistent) documentation on the local protocol of the network(s) to which a particular host is to be interfaced.

4. Addressing

All hosts on the Wideband Net have at least one unique Internet address. Since the 8-bit Internet Network Number is already specified by the Internet protocols (the Wideband Net is number 28 decimal), this leaves a 24-bit address space for use within the Wideband Net.

Since the Wideband Net is organized as a catenet, the "network number / local host number" strategy employed by The Internet is also employed within the Wideband Net. Each distinct, constituent network is assigned a unique, 8-bit network number; and each host is assigned a unique, 16-bit local address on each network to which it is attached. These two numbers are concatenated to produce the 24-bit Wideband Net Address.

5. Routing and Gateways

Routing in the Wideband Net is implemented in the same manner as routing in The Internet [4]. The constituent networks are connected by gateways; these gateways implement the routing function. One of the primary functions of the Voice Funnels is to implement this gateway function for each network to which it is attached. If a network is to be attached to the Wideband Net at some place besides a Voice Funnel, then it is the responsibility of that network to provide a gateway.

The IP server in each host needs to know very little about routing in order to function properly. It is assumed that each host knows its own address (that is, its network number and its address on that network) and the address of at least one adjacent gateway. Also, the IP server must be able to produce a local network protocol header from a 24-bit Wideband Net address which specifies this local network.

When a host's IP server is called on to deliver a datagram to an adjacent host it wraps the datagram first in an IP header, then in a local header which is addressed to that host, and sends the result directly to its destination. When an IP server is called on to deliver to a non-adjacent host, it wraps the datagram first in an IP header, then in a local header which is addressed to an adjacent gateway and delivers it. That gateway, in cooperation with the other gateways of the Wideband Net, will deliver the datagram to its destination.

In addition to transmitting outgoing datagrams, IP servers will receive datagrams from their network(s). These should be routed internally as appropriate; this is probably dependent on the "Protocol" field of the IP header (but note that a protocol of '3' indicates an advisory message from a gateway directed at the IP server itself).

If an IP server is to function optimally, it must keep track of

ALL adjacent gateways which are up. Additionally, it must maintain a cache of those non-adjacent hosts to which it has recently sent datagrams, along with the gateway through which these datagrams were forwarded. When it is about to send a datagram, the IP server first checks the cache. If the addressee is found in the cache and if the gateway named in the cache is still up, then the datagram is forwarded to that same gateway. Otherwise, a gateway is chosen arbitrarily through which to forward the datagram.

Occasionally, a routing advisory message will be sent to the IP server by a gateway. Currently, one of two advisories may be received: "destination unreachable", meaning that it will be futile to send more datagrams to some host for a while (e.g. a couple of minutes); and "redirect", meaning that a non-optimal gateway was used, and that further datagrams should be forwarded through the gateway specified in the redirect message. When an IP server receives a routing advisory, it should update its cache, and perhaps notify one or more processes as appropriate.

6. Unresolved Issues

Since the currently planned constituents of the Wideband Net (the PSAT Net, LexNets, and Voice Funnels) are all capable of supporting broadcasting, it might be worthwhile to consider some form of broadcasting as a basic Wideband Net service. This would probably take two forms: group addressing, essentially extending the group concept of the PSAT Net to allow broadcasting to a designated group of hosts of the Wideband Net; and general broadcasting, directed at all hosts on a particular subnet or at all hosts of the Wideband Net.

The ST protocol should also be supported directly by the Wideband Net's gateways. ST will allow the capacity of the various subnets to be more fully utilized. It should also be made to take advantage of the broadcast nature of the various subnets.

7. Glossary

In order to avert the ambiguities inherent in this two-level structure, this memo uses a two-level nomenclature, defined as follows:

network: A physical communication service in which all attached hosts communicate with all other hosts using a uniform, local set of link, addressing, and transport protocols; i.e. the "usual" meaning of the word. To quote Cerf [1], "the term 'local' is used in a loose sense here, since it means 'peculiar to the particular network' rather than 'a network of limited geographical extent'. A satellite-based network, such as the ARPA packet satellite network, therefore has 'local' characteristics (e.g. broadcast operation) even though it spans many thousands of square miles geographically speaking." A network must appear to be homogeneous from the "outside looking in"; however, this does not necessarily preclude an internal structure.

Internet Network: A communicating system of hosts and/or networks belonging to The Internet, and which can be uniquely identified by an 8-bit "network number" assigned by the number czar [3]. This may correspond to part of a network, one network, or a concatenation of many networks. Note that the Wideband Net is an Internet Network.

catenet: A collection of two or more networks, arbitrarily interconnected by gateways, in which the communicating hosts have agreed, a-priori, on some canonical "catenet protocol" which is used for datagram transport.

The Internet: The collection of Internet Networks, along with their Internet Gateways. Hosts on The Internet communicate using version 4 of the DoD Standard Internet Protocol [2].

gateway: A logical host which is connected to two or more networks, and which can forward "catenet protocol" datagrams arriving from any of these networks to the appropriate outgoing network. Many adjectives may be applied to "gateway" as needed. A hidden gateway is one whose presence is generally unknown by the hosts attached to the networks adjacent to the gateway. A non-routing (or static-routing) gateway makes its routing

decisions based on a-priori information; a routing gateway exchanges information with other routing gateways in order to be able to make dynamic adjustments to its routing information as the conditions of the catenet change.

Internet Gateway:
A gateway connected to two or more Internet Networks, which can forward internet datagrams.

REFERENCES

[1] Cerf.
The Catenet Model for Internetworking.
Internet Experiment Note 48, July, 1978.

[2] Postel, ed.
The DoD Standard Internet Protocol.
Internet Experiment Note 128, January, 1980.

[3] Postel, ed.
Assigned Numbers.
Internet Experiment Note 127, January, 1980.

[4] Strasisar.
How to Build a Gateway.
Internet Experiment Note 109, August, 1979.