

Request for Comments: 851
Obsoletes RFC: 802

The ARPANET 1822L Host Access Protocol

RFC 851

Andrew G. Malis
ARPANET Mail: malis@bbn-unix

Bolt Beranek and Newman Inc.
50 Moulton St.
Cambridge, MA 02238

April 1983

This RFC specifies the ARPANET 1822L Host Access Protocol, which is a successor to the existing 1822 Host Access Protocol. 1822L allows ARPANET hosts to use logical names as well as 1822's physical port locations to address each other. The RFC is also being presented as a solicitation of comments on 1822L, especially from host network software implementers and maintainers.

Table of Contents

1	INTRODUCTION.....	1
2	THE ARPANET 1822L HOST ACCESS PROTOCOL.....	4
2.1	Addresses and Names.....	6
2.2	Name Translations.....	8
2.2.1	Authorization and Effectiveness.....	8
2.2.2	Translation Policies.....	11
2.2.3	Reporting Destination Host Downs.....	13
2.2.4	1822L and 1822 Interoperability.....	16
2.3	Uncontrolled Packets.....	18
2.4	Establishing Host-IMP Communications.....	20
2.5	Counting RFMS When Using 1822L.....	22
2.6	1822L Name Server.....	24
3	1822L LEADER FORMATS.....	27
3.1	Host-to-IMP 1822L Leader Format.....	28
3.2	IMP-to-Host 1822L Leader Format.....	35
4	REFERENCES.....	43

FIGURES

1822 Address Format.....	6
1822L Name Format.....	7
1822L Address Format.....	7
Communications between different host types.....	17
Host-to-IMP 1822L Leader Format.....	28
NDM Message Format.....	31
IMP-to-Host 1822L Leader Format.....	35
Name Server Reply Format.....	39

1 INTRODUCTION

This RFC specifies the ARPANET 1822L Host Access Protocol, which will allow hosts to use logical addressing (i.e., host names that are independent of their physical location on the ARPANET) to communicate with each other. This new host access protocol is known as the ARPANET 1822L (for Logical) Host Access Protocol, and is a successor to the current ARPANET 1822 Host Access Protocol, which is described in sections 3.3 and 3.4 of BBN Report 1822 [1]. Although the 1822L protocol uses different Host-IMP leaders than the 1822 protocol, the IMPs will continue to support the 1822 protocol, and hosts using either protocol can readily communicate with each other (the IMPs will handle the translation automatically).

There is one major restriction to the new 1822L protocol: it will be implemented in C/30 IMPs only, and will therefore only be usable by hosts connected to C/30 IMPs, as Honeywell and Pluribus IMPs do not have sufficient memory to hold the new programs and tables. This restriction also means that logical addressing cannot be used to identify a host on a non-C/30 IMP. While this is not a problem on the ARPANET, which only has C/30 IMPs, the restriction will apply if logical addressing is used on any network that mixes C/30 and non-C/30 IMPs.

The RFC's terminology is consistent with that used in Report 1822, and any new terms will be defined when they are first used. Familiarity with Report 1822 (section 3 in particular) is assumed. As could be expected, the RFC makes many references to Report 1822. As a result, it uses, as a convenient abbreviation, "see 1822(x)" instead of "please refer to Report 1822, section x, for further details".

This RFC updates, and obsoletes, RFC 802. The changes from that RFC include:

- o The Short Blocking Feature, which had also been described in RFC 802, now has its own RFC, RFC 852 [2]. It was moved to its own RFC, since it is completely independent of logical addressing.
- o In section 2.2, descriptions of the three address selection policies and of host error handling have been added.
- o In section 2.3, the IMP's uncontrolled packet service has been further improved. This applies to hosts using 1822 as well as 1822L.
- o Pointers on using RFLM counting with 1822L have been added as section 2.5.

- o Section 2.6 describes the new "1822L name server" in the IMP, which makes use of two new Host-to-IMP messages to allow hosts to do their own name-to-address mapping.
- o In section 3.2, the subtypes for the type 15 (1822L Name or Address Error) IMP-to-Host message have been changed.

2 THE ARPANET 1822L HOST ACCESS PROTOCOL

The ARPANET 1822L Host Access Protocol allows a host to use logical addressing to communicate with other hosts on the ARPANET. Basically, logical addressing allows hosts to refer to each other using an 1822L name (see section 2.1) which is independent of a host's physical location in the network. IEN 183 (also published as BBN Report 4473) [3] gives the use of logical addressing considerable justification. Among the advantages it cites are:

- o The ability to refer to each host on the network by a name independent of its location on the network.
- o Allowing different hosts to share the same host port on a time-division basis.
- o Allowing a host to use multi-homing (where a single host uses more than one port to communicate with the network).
- o Allowing several hosts that provide the same service to share the same name.

The main differences between the 1822 and 1822L protocols are the format of the leaders that are used to introduce messages between a host and an IMP, and the specification in those leaders of the source and/or destination host(s). Hosts have the choice of

using the 1822 or the 1822L protocol. When a host comes up on an IMP, it declares itself to be an 1822 host or an 1822L host by the type of NOP message (see section 3.1) it uses. Once up, hosts can switch from one protocol to the other by issuing an appropriate NOP. Hosts that do not use the 1822L protocol will still be addressable by and can communicate with hosts that do, and vice-versa.

Another difference between the two protocols is that the 1822 leaders are symmetric, while the 1822L leaders are not. The term symmetric means that in the 1822 protocol, the exact same leader format is used for messages in both directions between the hosts and IMPs. For example, a leader sent from a host over a cable that was looped back onto itself (via a looping plug or faulty hardware) would arrive back at the host and appear to be a legal message from a real host (the destination host of the original message). In contrast, the 1822L headers are not symmetric, and a host can detect if the connection to its IMP is looped by receiving a message with the wrong leader format. This allows the host to take appropriate action upon detection of the loop.

2.1 Addresses and Names

The 1822 protocol defines one form of host specification, and the 1822L protocol defines two additional ways to identify network hosts. These three forms are 1822 addresses, 1822L names, and 1822L addresses.

1822 addresses are the 24-bit host addresses found in 1822 leaders. They have the following format:

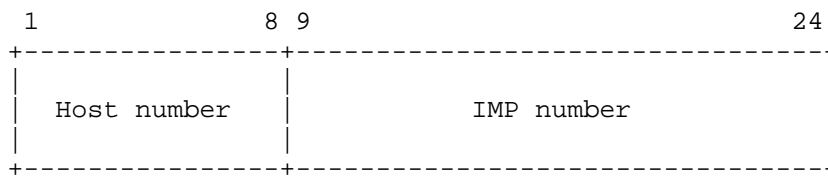


Figure 1. 1822 Address Format

These fields are quite large, and the ARPANET will never use more than a fraction of the available address space. 1822 addresses are used in 1822 leaders only.

1822L names are 16-bit unsigned numbers that serve as a logical identifier for one or more hosts. 1822L names have a much simpler format:

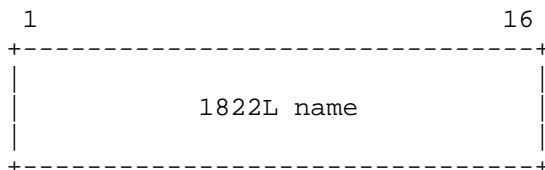


Figure 2. 1822L Name Format

The 1822L names are just 16-bit unsigned numbers, except that bits 1 and 2 are not both zeros (see below). This allows over 49,000 hosts to be specified.

1822 addresses cannot be used in 1822L leaders, but there may be a requirement for an 1822L host to be able to address a specific physical host port or IMP fake host. 1822L addresses are used for this function. 1822L addresses form a subset of the 1822L name space, and have both bits 1 and 2 off.

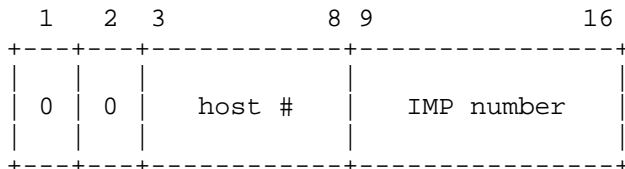


Figure 3. 1822L Address Format

This format allows 1822L hosts to directly address hosts 0-63 at IMPs 1-255 (IMP 0 does not exist). Note that the highest host

numbers are reserved for addressing the IMP's internal fake hosts. At this writing, the IMP has seven fake hosts, so host numbers 57-63 address the IMP fake hosts, while host numbers 0-56 address real hosts external to the IMP. As the number of IMP fake hosts changes, this boundary point will also change.

2.2 Name Translations

There are a number of factors that determine how an 1822L name is translated by the IMP into a physical address on the network. These factors include which translations are legal; in what order different translations for the same name should be attempted; which legal translations shouldn't be attempted because a particular host port is down; and the interoperability between 1822 and 1822L hosts. These issues are discussed in the following sections.

2.2.1 Authorization and Effectiveness

Every host on a C/30 IMP, regardless of whether it is using the 1822 or 1822L protocol to access the network, can have one or more 1822L names (logical addresses). Hosts using 1822L can then use these names to address the hosts in the network independent of their physical locations. Because of the implementation

constraints mentioned in the introduction, hosts on non-C/30 IMPs cannot be assigned 1822L names. To circumvent this restriction, however, 1822L hosts can also use 1822L addresses to access all of the other hosts.

At this point, several questions arise: How are these names assigned, how do they become known to the IMPs (so that translations to physical addresses can be made), and how do the IMPs know which host is currently using a shared port? To answer each question in order:

Names are assigned by a central network administrator. When each name is created, it is assigned to a host (or a group of hosts) at one or more specific host ports. The host(s) are allowed to reside at those specific host ports, and nowhere else. If a host moves, it will keep the same name, but the administrator has to update the central database to reflect the new host port. Changes to this database are distributed to the IMPs by the Network Operations Center (NOC). For a while, the host may be allowed to reside at either of (or both) the new and old ports. Once the correspondence between a name and one or more hosts ports where it may be used has been made official by the administrator, that name is said to be authorized. 1822L addresses, which actually refer to physical host ports, are always authorized in this sense.

Once a host has been assigned one or more names, it has to let the IMPs know where it is and what name(s) it is using. There are two cases to consider, one for 1822L hosts and another for 1822 hosts. The following discussion only pertains to hosts on C/30 IMPs.

When an IMP sees an 1822L host come up on a host port, the IMP has no way of knowing which host has just come up (several hosts may share the same port, or one host may prefer to be known by different names at different times). This requires the host to declare itself to the IMP before it can actually send and receive messages. This function is performed by a new host-to-IMP message, the Name Declaration Message (NDM), which lists the names that the host would like to be known by. The IMP checks its tables to see if each of the names is authorized, and sends an NDM Reply to the host saying which names were actually authorized and can now be used for sending and receiving messages (i.e., which names are effective). A host can also use an NDM message to change its list of effective names (it can add to and delete from the list) at any time. The only constraint on the host is that any names it wishes to use can become effective only if they are authorized.

In the second case, if a host comes up on a C/30 IMP using the 1822 protocol, the IMP automatically makes the first name the IMP

finds in its tables for that host become effective. Thus, even though the host is using the 1822 protocol, it can still receive messages from 1822L hosts via its 1822L name. Of course, it can also receive messages from an 1822L host via its 1822L address as well. (Remember, the distinction between 1822L names and addresses is that the addresses correspond to physical locations on the network, while the names are strictly logical identifiers). The IMPs translate between the different leaders and send the proper leader in each case (see section 2.2.4).

The third question above has by now already been answered. When an 1822L host comes up, it uses the NDM message to tell the IMP which host it is (which names it is known by). Even if this is a shared port, the IMP knows which host is currently connected.

Whenever a host goes down, its names automatically become non-effective. When it comes back up, it has to make them effective again.

2.2.2 Translation Policies

Several hosts can share the same 1822L name. If more than one of these hosts is up at the same time, any messages sent to that 1822L name will be delivered to just one of the hosts sharing that name, and a RFNM will be returned as usual. However, the

sending host will not receive any indication of which host received the message, and subsequent messages to that name are not guaranteed to be sent to the same host. Typically, hosts providing exactly the same service could share the same 1822L name in this manner.

Similarly, when a host is multi-homed, the same 1822L name may refer to more than one host port (all connected to the same host). If the host is up on only one of those ports, that port will be used for all messages addressed to the host. However, if the host were up on more than one port, the message would be delivered over just one of those ports, and the subnet would choose which port to use. This port selection could change from message to message. If a host wanted to insure that certain messages were delivered to it on specific ports, these messages could use either the port's 1822L address or a specific 1822L name that referred to that port alone.

Three different address selection policies are available for the name mapping process. When translated, each name uses one of the three policies (the policy is pre-determined on a per-name basis). The three policies are:

- o Attempt each translation in the order in which the physical addresses are listed in the IMP's translation tables, to find

the first reachable physical host address. This list is always searched from the top whenever an uncontrolled packet is to be sent or an end-to-end connection has to be created. This is the most commonly used policy.

- o Selection of the closest physical address, which uses the IMP's routing tables to find the translation to the destination IMP with the least delay path.
- o Use load leveling. This is similar to the second policy, but differs in that searching the address list for a valid translation starts at the address following where the previous translation search ended. This attempts to spread out the load from any one IMP's hosts to the various host ports associated with a particular name. Note that this is NOT network-wide load leveling, which would require a distributed algorithm and tables.

2.2.3 Reporting Destination Host Downs

As was explained in report 1822, and as will be discussed in greater detail in section 2.5, whenever regular messages are sent by a host, the IMP opens a subnetwork connection to each destination host from the source host. A connection will stay open at least as long as there are any outstanding (un-RFNMed)

messages using it and both the source and destination hosts stay up.

However, the destination host may go down for some reason during the lifetime of a connection. If the host goes down while there are no outstanding messages to it in the network, then the connection is closed and no other action is taken until the source host submits the next message for that destination. At that time, ONE of the following events will occur:

- A1. If 1822 or an 1822L address is being used to specify the destination host, then the source host will receive a type 7 (Destination Host Dead) message from the IMP.
- A2. If an 1822L name is being used to specify the destination host, and the name maps to only one authorized host port, then a type 7 message will also be sent to the source host.
- A3. If an 1822L name is being used to specify the destination host, and the name maps to more than one authorized host port, then the IMP attempts to open a connection to another authorized and effective host port for that name. If no such connection can be made, the host will receive a type 15 (1822L Name or Address Error), subtype 5 (no effective translations) message (see section 3.2). Note that a type 7 message cannot be returned to the source host, since type 7 messages refer to a particular destination host port, and

the name maps to more than one destination port.

Things get a bit more complicated if there are any outstanding messages on the connection when the destination host goes down. The connection will be closed, and one of the following will occur:

- B1. If 1822 or an 1822L address is being used to specify the destination host, then the source host will receive a type 7 message for each outstanding message.
- B2. If an 1822L name is being used to specify the destination host, then the source host will receive a type 9 (Incomplete Transmission), subtype 3 (message lost due to network failure) message for each outstanding message. The next time the source host submits another message for that same destination name, the previous algorithm will be used (either step A2 or step A3).

The above two algorithms also apply when a host stays up, but declares the destination name for an existing connection to no longer be effective. In this case, however, the type 7 messages above will be replaced by type 15, subtype 3 (name not effective) messages.

Section 2.3 discusses how destination host downs are handled for uncontrolled packets.

2.2.4 1822L and 1822 Interoperability

As has been previously stated, 1822 and 1822L hosts can intercommunicate, and the IMPs will automatically handle any necessary leader and address format conversions. However, not every combination of 1822 and 1822L hosts allows full interoperability with regard to the use of 1822L names.

The following figure illustrates how these addressing combinations are handled, showing how each type of host can access every other type of host. There are three types of hosts: "1822 on C/30" signifies an 1822 host that is on a C/30 IMP, "1822L" signifies an 1822L host (on a C/30 IMP), and "1822 on non-C/30" signifies a host on a non-C/30 IMP (which cannot support the 1822L protocol). The table entry shows the protocol and host address format(s) that the source host can use to reach the destination host.

Source Host	Destination Host		
	1822 on C/30	1822L	1822 on non-C/30
1822 on C/30	1822	1822 (note 1)	1822
1822L	1822L, using 1822L name or address (note 2)	1822L, using 1822L name or address	1822L, using 1822L address only (note 2)
1822 on non-C/30	1822	1822 (note 1)	1822

Note 1: The message is presented to the destination host with an 1822L leader containing the 1822L addresses of the source and destination hosts. If either address cannot be encoded as an 1822L address, then the message is not delivered and an error message is sent to the source host.

Note 2: The message is presented to the destination host with an 1822 leader containing the 1822 address of the source host.

Figure 4. Communications between different host types

2.3 Uncontrolled Packets

Uncontrolled packets (see 1822(3.6)) present a unique problem for the 1822L protocol. Uncontrolled packets use none of the normal ordering and error-control mechanisms in the IMP, and do not use the normal subnetwork connection facilities. As a result, uncontrolled packets need to carry all of their overhead with them, including source and destination names. If 1822L names are used when sending an uncontrolled packet, additional information is now required by the subnetwork when the packet is transferred to the destination IMP. This means that less host-to-host data can be contained in the packet than is possible between 1822 hosts.

Uncontrolled packets that are sent between 1822 hosts may contain not more than 991 bits of data. Uncontrolled packets that are sent to and/or from 1822L hosts are limited to 32 bits less, or not more than 959 bits. Packets that exceed this length will result in an error indication to the host, and the packet will not be sent. This error indication represents an enhancement to the previous level of service provided by the IMP, which would simply discard an overly long uncontrolled packet without notification.

Other enhancements that are provided for uncontrolled packet service are a notification to the host of any errors that are detected by the host's IMP when it receives the packet. A host will be notified if an uncontrolled packet contains an error in the 1822L name specification, such as if the name is not authorized or effective, if the remote host is unreachable (which is indicated by none of its names being effective), if network congestion control throttled the packet before it left the source IMP, or for any other reason the source IMP was not able to send the packet on its way.

In most cases, the host will not be notified if the uncontrolled packet was lost once it was transmitted by the source IMP. However, the IMP will attempt to notify the source host if a logically-addressed uncontrolled packet was mistakenly sent to a host that the source IMP thought was effective, but which turned out to be dead or non-effective at the destination IMP. This non-delivery notice is sent back to the source IMP as an uncontrolled packet from the destination IMP, so the source host is not guaranteed to receive this indication.

If the source IMP successfully receives the non-delivery notice, then the source host will receive a type 15 (1822L Name or Address Error), subtype 6 (down or non-effective port) message. If the packet is resubmitted or another packet is sent to the

same destination name, and there are no available effective translations, then the source host will receive a type 15, subtype 5 (no effective translations) message if the destination name has more than one mapping; or will receive either a type 7 (Destination Host Dead) or a type 15, subtype 3 (name not effective) message if the destination name has a single translation.

Those enhancements to the uncontrolled packet service that are not specific to logical addressing will be available to hosts using 1822 as well as 1822L. However, logically-addressed uncontrolled packets must be used in order to receive any indication that the packet was lost once it has left the source IMP.

2.4 Establishing Host-IMP Communications

When a host comes up on an IMP, or after there has been a break in the communications between the host and its IMP (see 1822(3.2)), the orderly flow of messages between the host and the IMP needs to be properly (re)established. This allows the IMP and host to recover from most any failure in the other or in their communications path, including a break in mid-message.

The first messages that a host should send to its IMP are three NOP messages. Three messages are required to insure that at least one message will be properly read by the IMP (the first NOP could be concatenated to a previous message if communications had been broken in mid-stream, and the third provides redundancy for the second). These NOPS serve several functions: they synchronize the IMP with the host, they tell the IMP how much padding the host requires between the message leader and its body, and they also tell the IMP whether the host will be using 1822 or 1822L leaders.

Similarly, the IMP will send three NOPS to the host when it detects that the host has come up. Actually, the IMP will send six NOPS, alternating three 1822 NOPS with three 1822L NOPS. Thus, the host will see three NOPS no matter which protocol it is using. The NOPS will be followed by two Interface Reset messages, one of each style. If the IMP receives a NOP from the host while the above sequence is occurring, the IMP will only send the remainder of the NOPS and the Interface Reset in the proper style. The 1822 NOPS will contain the 1822 address of the host interface, and the 1822L NOPS will contain the corresponding 1822L address.

Once the IMP and the host have sent each other the above messages, regular communications can commence. See 1822(3.2) for

further details concerning the ready line, host tardiness, and other issues.

2.5 Counting RFMS When Using 1822L

When a host submits a regular message using an 1822 leader, the IMP checks for an existing simplex virtual circuit connection from the source host to the destination host. If such a connection already exists, it is used. Otherwise, a new connection from the source host port to the destination host port is opened. In either case, there may be at most eight messages outstanding on that connection at any one time. If a host submits a ninth message on that connection before it receives a reply for the first message, then the host will be blocked until the reply is sent for the first message.

Such connections can stay open for some time, but are timed out after three minutes of no activity, or can be closed if there is contention for the connection blocks in either the source or destination IMP. However, a connection will never be closed as long as there are any outstanding messages on it. This allows a source host to count the number of replies it has received for messages to each destination host address in order to avoid being blocked by submitting a ninth outstanding message on any

connection.

When a host submits a regular message using an 1822L leader, a similar process occurs, except that in this case, connections are distinguished by the source name/destination name combination. When the message is received from a host, the IMP first looks for an open connection for that same source name/destination name pair. If such a connection is found, then it is used, and no further name translation is performed. If, however, no open connection was found, then the destination name is translated, and a connection opened to the physical host port. As long as there are any outstanding messages on the connection it will stay open, and it will have the same restriction that only eight messages may be outstanding at any one time. Thus, a source host can still count replies to avoid being blocked, but they must be counted on a source name/destination name pair basis, instead of just by destination host address as before.

Since connections are based on the source name as well as the destination name, this implies that there may be more than one open connection from physical host port A to physical host port B, which would allow more than 8 outstanding messages simultaneously from the first to the second port. However, for this to occur, either the source or destination names, or both, must differ from one connection to the next. For example, if the

names "543" and "677" both translate to physical port 3 on IMP 51, then the host on that port could open four connections to itself by sending messages from "543" to "543", from "543" to "677", from "677" to "543", and from "677" to "677".

As has already been stated, the destination names in regular messages are only translated when connections are first opened. Once a connection is open, that connection, and its destination physical host port, will continue to be used until it is closed. If, in the meantime, a "better" destination host port belonging to the same destination name became available, it would not be used until the next time a new connection is opened to that destination name.

2.6 1822L Name Server

There may be times when a host wants to perform its own translations, or might need the full list of physical addresses to which a particular name maps. For example, a connection-based host-to-host protocol may require that the same physical host port on a multi-homed host be used for all messages using that host-to-host connection, and the host does not wish to trust the IMP to always deliver messages using a destination name to the same host port.

In these cases, the host can submit a type 11 (Name Server Request) message to the IMP, which requests the IMP to translate the destination 1822L name and return a list of the addresses to which it maps. The IMP will respond with a type 11 (Name Server Reply) message, which contains the selection policy in use for that name, the number of addresses to which the name maps, the addresses themselves, and for each address, whether it is effective and its routing distance from the IMP. See section 3.2 for a complete description of the message's contents.

Using this information, the source host can make an informed decision on which of the physical host ports corresponding to an 1822L name to use, and can subsequently send the messages to that port, rather than to the name.

The IMP also supports a different type of name service. A host needs to issue a Name Declaration Message to the IMP in order to make its names effective, but it may not wish to keep its names in some table or file in the host. In this case, it can ask the IMP to tell it which names it is authorized to use.

In this case, the host submits a type 12 (Port List Request) message to the IMP, and the IMP replies with a type 12 (Port List Reply) message. It contains, for the host port over which the IMP received the request and sent the reply, the number of names

that map to the port, the list of names, and whether or not each name is effective. The host can then use this information in order to issue the Name Declaration Message. Section 3.2 contains a complete description of the reply's contents.

3 1822L LEADER FORMATS

The following sections describe the formats of the leaders that precede messages between an 1822L host and its IMP. They were designed to be as compatible with the 1822 leaders as possible. The second, fifth, and sixth words are identical in the two leaders, and all of the existing functionality of the 1822 leaders has been retained. In the first word, the 1822 New Format Flag is now also used to identify the two types of 1822L leaders, and the Handling Type has been moved to the second byte. The third and fourth words contain the Source and Destination 1822L Name, respectively.

3.1 Host-to-IMP 1822L Leader Format

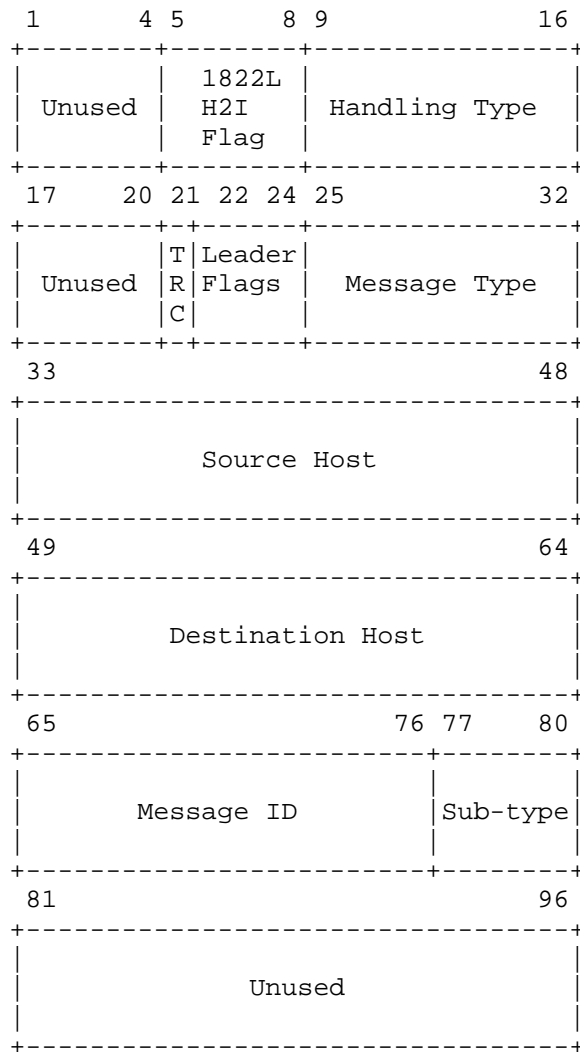


Figure 5. Host-to-IMP 1822L Leader Format

Bits 1-4: Unused, must be set to zero.

Bits 5-8: 1822L Host-to-IMP Flag:

This field is set to decimal 13 (1101 in binary).

Bits 9-16: Handling Type:

This field is bit-coded to indicate the transmission characteristics of the connection desired by the host. See 1822(3.3).

Bit 9: Priority Bit:

Messages with this bit on will be treated as priority messages.

Bits 10-16: Unused, must be zero.

Bits 17-20: Unused, must be zero.

Bit 21: Trace Bit:

If equal to one, this message is designated for tracing as it proceeds through the network. See 1822(5.5).

Bits 22-24: Leader Flags:

Bit 22: A flag available for use by the destination host.

See 1822(3.3) for a description of its use by the IMP's TTY Fake Host.

Bits 23-24: Reserved for future use, must be zero.

Bits 25-32: Message Type:

Type 0: Regular Message - All host-to-host communication occurs via regular messages, which have several sub-types, found in bits 77-80. These sub-types are:

0: Standard - The IMP uses its full message and error control facilities, and host blocking may occur.

3: Uncontrolled Packet - The IMP will perform no message-control functions for this type of message, and network flow and congestion control may cause loss of the packet. Also see 1822(3.6) and section 2.3.

4-15: Unassigned.

Type 1: Error Without Message ID - See 1822(3.3).

Type 2: Host Going Down - see 1822(3.3).

Type 3: Name Declaration Message (NDM) - This message is used by the host to declare which of its 1822L names is or is not effective (see section 2.2.1), or to make all of its names non-effective. The first 16 bits of the data portion of the NDM message, following the leader and any leader padding, contains the number of 1822L names contained in the message. This is followed by the 1822L name entries, each 32 bits long, of which the first 16 bits is a 1822L name and the second 16 bits contains either of the integers zero or one. Zero

indicates that the name should not be effective, and one indicates that the name should be effective. The IMP will reply with a NDM Reply message (see section 3.2) indicating which of the names are now effective and which are not. Pictorially, a NDM message has the following format (including the leader, which is printed in hexadecimal):

1	16 17	32 33	48
+-----+-----+-----+-----+			
0D00	0003	0000	
+-----+-----+-----+-----+			
49	64 65	80 81	96
+-----+-----+-----+-----+			
0000	0000	0000	
+-----+-----+-----+-----+			
97	112 113	128 129	144
+-----+-----+-----+-----+			
# of entries	1822L name #1	0 or 1	
+-----+-----+-----+-----+			
145	160 161	176	
+-----+-----+-----+-----+			
1822L name #2	0 or 1	etc.	
+-----+-----+-----+-----+			

Figure 6. NDM Message Format

An NDM with zero entries will cause all current effective names for the host to become non-effective.

Type 4: NOP - This allows the IMP to know which style of leader the host wishes to use. A 1822L NOP signifies that the host wishes to use 1822L leaders, and an 1822 NOP signifies that the host wishes to use 1822 leaders. All of the other remarks concerning the NOP message in 1822(3.3) still hold. The host should always issue NOPs in groups of three to insure proper reception by the IMP. Also see section 2.4 for a further discussion on the use of the NOP message.

Type 8: Error with Message ID - see 1822(3.3).

Type 11: Name Server Request - This allows the host to use the IMP's logical addressing tables as a name server. The destination name in the 1822L leader is translated, and the IMP replies with a Name Server Reply message, which lists the physical host addresses to which the destination name maps.

Type 12: Port List Request - This allows the physical host to request the list of names that map to the host port over which this request was received by the IMP. The IMP replies with a Port List Reply message, which lists the names that map to the port.

Types 5-7,9-10,13-255: Unassigned.

Bits 33-48: Source Host:

This field contains one of the source host's 1822L names (or, alternatively, the 1822L address of the host port the message is being sent over). This field is not automatically filled in by the IMP, as in the 1822 protocol, because the host may be known by several names and may wish to use a particular name as the source of this message. All messages from the same host need not use the same name in this field. Each source name, when used, is checked for authorization, effectiveness, and actually belonging to this host. Messages using names that do not satisfy all of these requirements will not be delivered, and will instead result in an error message being sent back into the source host. If the host places its 1822L address in this field, the address is checked to insure that it actually represents the host port where the message originated. If the message is destined for an 1822 host on a non-C/30 IMP, this field MUST contain the source host's 1822L address (see figure 4 in section 2.2.4).

Bits 49-64: Destination Host:

This field contains the 1822L name or address of the destination host. If it contains a name, the name will be checked for effectiveness, with an error message returned to

the source host if the name is not effective. If the message is destined for an 1822 host on a non-C/30 IMP, this field MUST contain the destination host's 1822L address (see figure 4 in section 2.2.4).

Bits 65-76: Message ID:

This is a host-specified identification used in all type 0 and type 8 messages, and is also used in type 2 messages. When used in type 0 messages, bits 65-72 are also known as the Link Field, and should contain values specified in Assigned Numbers [4] appropriate for the host-to-host protocol being used.

Bits 77-80: Sub-type:

This field is used as a modifier by message types 0, 2, 4, and 8.

Bits 81-96: Unused, must be zero.

3.2 IMP-to-Host 1822L Leader Format

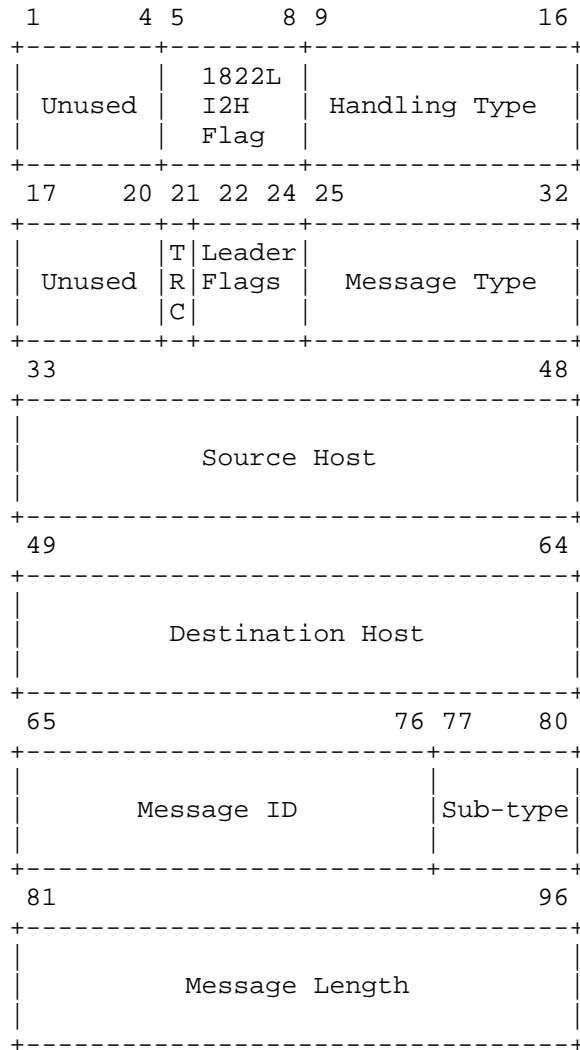


Figure 7. IMP-to-Host 1822L Leader Format

Bits 1-4: Unused and set to zero.

Bits 5-8: 1822L IMP-to-Host Flag:

This field is set to decimal 14 (1110 in binary).

Bits 9-16: Handling Type:

This has the value assigned by the source host (see section 3.1). This field is only used in message types 0, 5-9, and 15.

Bits 17-20: Unused and set to zero.

Bit 21: Trace Bit:

If equal to one, the source host designated this message for tracing as it proceeds through the network. See 1822(5.5).

Bits 22-24: Leader Flags:

Bit 22: Available as a destination host flag.

Bits 23-24: Reserved for future use, set to zero.

Bits 25-32: Message Type:

Type 0: Regular Message - All host-to-host communication occurs via regular messages, which have several sub-types. The sub-type field (bits 77-80) is the same as sent in the host-to-IMP leader (see section 3.1).

Type 1: Error in Leader - See 1822(3.4).

Type 2: IMP Going Down - See 1822(3.4).

Type 3: NDM Reply - This is a reply to the NDM host-to-IMP message (see section 3.1). It will have the same number of entries as the NDM message that is being replying to, and each listed 1822L name will be accompanied by a zero or a one (see figure 6). A zero signifies that the name is not effective, and a one means that the name is now effective.

Type 4: NOP - The host should discard this message. It is used during initialization of the IMP/host communication. The Destination Host field will contain the 1822L Address of the host port over which the NOP is being sent. All other fields are unused.

Type 5: Ready for Next Message (RFNM) - See 1822(3.4).

Type 6: Dead Host Status - See 1822(3.4).

Type 7: Destination Host or IMP Dead (or unknown) - See 1822(3.4).

Type 8: Error in Data - See 1822(3.4).

Type 9: Incomplete Transmission - See 1822(3.4).

Type 10: Interface Reset - See 1822(3.4).

Type 11: Name Server Reply - This reply to the Name Server Request host-to-IMP message contains a word with the selection policy and the number of physical addresses to which the destination name maps, followed by two words per physical address: the first word contains an

1822L address, and the second word contains a bit signifying whether or not that particular translation is effective and the routing distance (in 6.4 ms units) to the address's IMP. In figure 8, EFF is 1 for effective and 0 for non-effective, and POL is a two-bit number indicating the selection policy for the name (see section 2.2.2):

- 0: First reachable.
- 1: Closest physical address.
- 2: Load leveling.
- 3: Unused.

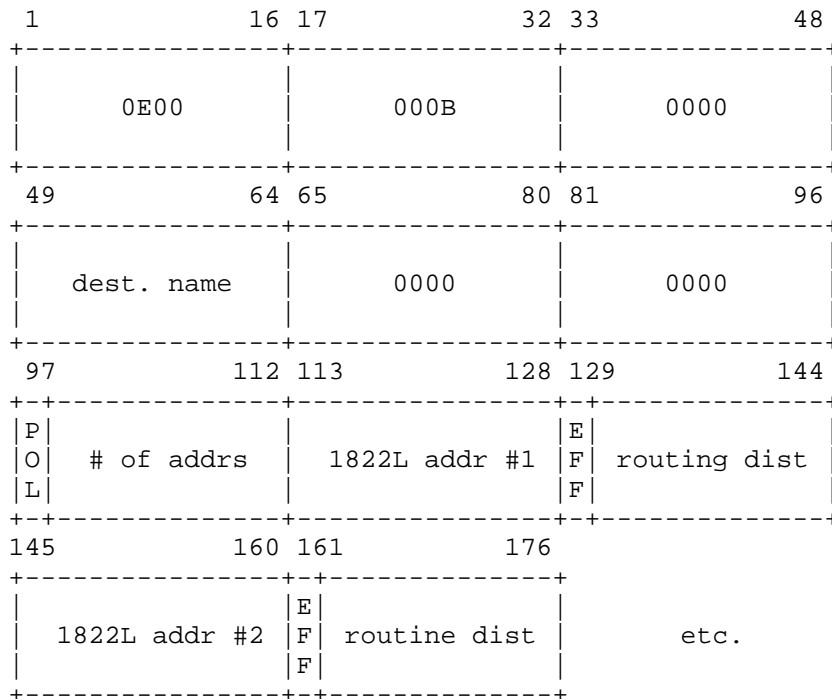


Figure 8. Name Server Reply Format

Type 12: Port List Reply - This is the reply to the Port List Request host-to-IMP message. It contains the number of names that map to this physical host port, followed by two words per name: the first word contains an 1822L name that maps to this port, and the second contains either a zero or a one, signifying whether or not that particular translation is effective. The format is identical to the type 3 NDM Reply message

(see figure 6).

Type 15: 1822L Name or Address Error - This message is sent in response to a type 0 message from a host that contained an erroneous Source Host or Destination Host field. Its sub-types are:

- 0: The Source Host 1822L name is not authorized or not effective.
- 1: The Source Host 1822L address does not match the host port used to send the message.
- 2: The Destination Host 1822L name is not authorized.
- 3: The physical host to which this singly-homed Destination Host name translated is authorized and up, but not effective. If the host was actually down, a type 7 message would be returned, not a type 15.
- 4: The Source or Destination Host field contains a 1822L name, but the host being addressed is on a non-C/30 IMP (see figure 4 in section 2.2.4).
- 5: The multi-homed Destination Host name is authorized, but has no available effective translations.
- 6: A logically-addressed uncontrolled packet was sent to a dead or non-effective host port. However, if it is resubmitted, there may be another effective host port to which the IMP may be able to attempt

to send the packet.

7: Logical addressing is not in use in this network.

8-15: Unassigned.

Types 13-14,16-255: Unassigned.

Bits 33-48: Source Host:

For type 0 messages, this field contains the 1822L name or address of the host that originated the message. All replies to the message should be sent to the host specified herein. For message types 5-9 and 15, this field contains the source host field used in a previous type 0 message sent by this host.

Bits 49-64: Destination Host:

For type 0 messages, this field contains the 1822L name or address that the message was sent to. This allows the destination host to detect how it was specified by the source host. For message types 5-9 and 15, this field contains the destination host field used in a previous type 0 message sent by this host.

Bits 65-76: Message ID:

For message types 0, 5, 7-9, and 15, this is the value assigned by the source host to identify the message (see section 3.1). This field is also used by message types 2

and 6.

Bits 77-80: Sub-type:

This field is used as a modifier by message types 0-2, 5-7, 9, and 15.

Bits 81-96: Message Length:

This field is contained in type 0, 3, 11, and 12 messages only, and is the actual length in bits of the message (exclusive of leader, leader padding, and hardware padding) as computed by the IMP.

4 REFERENCES

- [1] Specifications for the Interconnection of a Host and an IMP,
BBN Report 1822, December 1981 Revision.
- [2] A. Malis, The ARPANET Short Blocking Feature, Request For
Comments 852, April 1983.
- [3] E. C. Rosen et. al., ARPANET Routing Algorithm Improvements,
Internet Experimenter's Note 183 (also published as BBN
Report 4473, Vol. 1), August 1980, pp. 55-107.
- [4] J. Postel, Assigned Numbers, Request For Comments 820,
January 1983, p. 11.

INDEX

1822.....	4
1822 address.....	6
1822 host.....	5
1822L.....	4
1822L address.....	7
1822L host.....	5
1822L name.....	6
address selection policy.....	12
authorized.....	9
blocking.....	22
closest physical address.....	13
connection.....	22
destination host.....	33, 41
effective.....	10
first reachable.....	12
handing type.....	29, 36
host downs.....	13
leader flags.....	29, 36
link field.....	34
load leveling.....	13
logical addressing.....	4
message ID.....	34, 41
message length.....	42
message type.....	30, 36
multi-homing.....	4
name server.....	24, 32, 37
NDM.....	10, 30
NDM reply.....	10, 37
NOC.....	9
NOP.....	5, 20, 32, 37
priority bit.....	29
regular message.....	30, 36
RFNM.....	22, 37
source host.....	33, 41
standard message.....	30
sub-type.....	34, 42
symmetric.....	5
trace bit.....	29, 36
uncontrolled packet.....	18, 30

