

The ARPANET AHIP-E Host Access Protocol (Enhanced AHIP)

1. Status of this Memo

This RFC is a proposed specification for the encoding of Class A IP addresses for use on ARPANET-style networks such as the Milnet and Arpanet, and for enhancements to the ARPANET AHIP Host Access Protocol (AHIP; formerly known as 1822). These enhancements increase the size of the PSN field, allow ARPANET hosts to use logical names to address each other, allow for the communication of type-of-service information from the host to the PSN and enable the PSN to provide congestion feedback to the host on a connection basis. Distribution of this memo is unlimited. Comments on this RFC should be sent to the netmail address "ahipe@bbn.com".

Table of Contents

1	INTRODUCTION.....	4
2	IP ISSUES.....	6
2.1	Current Interpretation of Class A IP Address Fields	6
2.2	Requirements and Constraints Affecting New Class A Mapping	7
2.3	New Interpretation of IP Address Fields.....	8
2.4	Discussion of the New Mapping.....	10
2.5	Interoperability between Current AHIP and AHIP-E	11
3	LOGICAL ADDRESSING.....	13
3.1	Addresses and Names.....	13
3.2	Name Translations.....	14
3.2.1	Authorization and Effectiveness.....	15
3.2.2	Translation Policies.....	16
3.2.3	Reporting Destination Host Downs.....	17
3.3	Establishing Host-PSN Communications.....	18
3.4	Name Server.....	19
4	OTHER CHANGES.....	20
4.1	Type-of-Service Specification.....	20
4.2	Subnet Congestion Feedback.....	21
4.3	Precedence Level Information.....	21
5	FORMATS FOR NEW AHIP-E MESSAGES.....	23
5.1	Host-to-PSN AHIP-E Leader Format.....	23
5.2	PSN-to-Host AHIP-E Leader Format.....	27
6	AHIP-E VERSIONS.....	33
7	REFERENCES.....	34

FIGURES

2.1	IP Class A Mapping.....	6
2.2	New Class A IP Address Interpretation.....	8
2.3	AHIP-E Address and Name.....	9
3.1	Current AHIP Address Format.....	13
3.2	AHIP-E Address Format.....	14
3.3	Logical Name Format.....	14
5.1	Host-to-PSN AHIP-E Leader Format.....	23
5.2	NDM Message Format.....	25
5.3	PSN-to-Host AHIP-E Leader Format.....	27
5.4	Name Server Reply Format.....	30

1 INTRODUCTION

This RFC is a proposed specification for the encoding of Class A IP addresses for use on ARPANET-style networks such as the Milnet and Arpanet, and for enhancements to the AHIP Protocol (AHIP is the preferred term for what has previously been known as the 1822 protocol). These enhancements and modifications are partially motivated by a need to overcome the current address limitation of 256 PSNs per network and by a desire to allow hosts to take advantage of logical addressing with minimal change to their AHIP software. This enhanced AHIP protocol will be referred to as "AHIP-E". These enhancements will:

1. Increase the size of the PSN field to 10 bits.
2. Allow hosts to use logical names (i.e., host names that are independent of physical location on the network) in addition to physical port addresses to communicate with each other.
3. Enable the host to specify a type-of-service to the PSN.
4. Provide a mechanism for the PSN to communicate subnetwork congestion information to the host on a destination host basis. This will give the host an opportunity to selectively reduce its congesting flows, thus preventing all of its flows from being blocked by the network. Currently, a host has no way of knowing which of its flows is experiencing congestion; consequently, it is possible that one congesting flow can result in the blocking of all the host's flows .
5. Enable the PSN to inform the host about changes in precedence cutoff levels and about precedence level violations.

A host can take advantage of the extended and logical addressing capabilities without making substantial changes to its AHIP implementation. In particular, the specification provides three versions of AHIP-E: version 0 is current AHIP with no changes; version 1 allows use of logical and extended addressing with minimal change to code; version 2 constitutes full-fledged AHIP-E. This is described in further detail in chapter 6.

This RFC's terminology is consistent with that used in BBN Report 1822 [1], and any new terms are 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".

The rest of this RFC is organized as follows. Chapter 2 describes the new mapping between IP class A addresses and subnetwork hosts. Chapter 3 discusses logical addressing. Chapter 4 describes the enhancements related to type-of-service and reliability specification and to congestion and precedence feedback. Chapter 5 includes a specification of the new message types and their formats. Finally, chapter 6 describes the AHIP-E version numbering scheme.

2 IP ISSUES


This section discusses the changes to the mapping between Class A IP addresses [5] and subnet addresses. These changes are made necessary by:

1. The introduction of logical names.
2. The expansion of the PSN-number field.

Note that this RFC does not affect Class B and C mappings [5].

2.1 Current Interpretation of Class A IP Address Fields

Class A IP addresses are 32 bits in length, with 8 bits devoted to network number and 24 to the local address. In particular, they are of the form n.h.l.i, where n,h,l and i are decimal integers less than 256. AHIP addresses are 24 bits in length. The current ARPANET-style class A mapping is as follows (from RFC 796):



IP Class A Mapping
Figure 2.1

The LH (logical host) field is used by the hosts only and is not passed to the network.

2.2 Requirements and Constraints Affecting New Class A Mapping

This section discusses some of the requirements and constraints that were considered significant in determining the new address mapping.

1. Address Mapping Stability Requirement:

Any current IP physical address with 1 (logical host) = 0 should remain unchanged under the new design. For example, the binary string corresponding to 10.0.0.51 should continue to refer to sri-nic.arpa (assuming, of course, that sri-nic continues to reside on psn 51, port 0). This requirement is motivated by a desire to avoid a network-wide address switchover.

2. Existing implementation compatibility:

Existing compliant implementations of AHIP should continue to function for destinations with addresses fitting the restrictions in 1. In other words, such addresses should continue to refer to their original destinations, not only with the AHIP-E implementation (which is the condition in 1), but also with current ones.

3. Compatibility between X.25's IP address to subnet host mapping and AHIP's IP address to subnet host mapping:

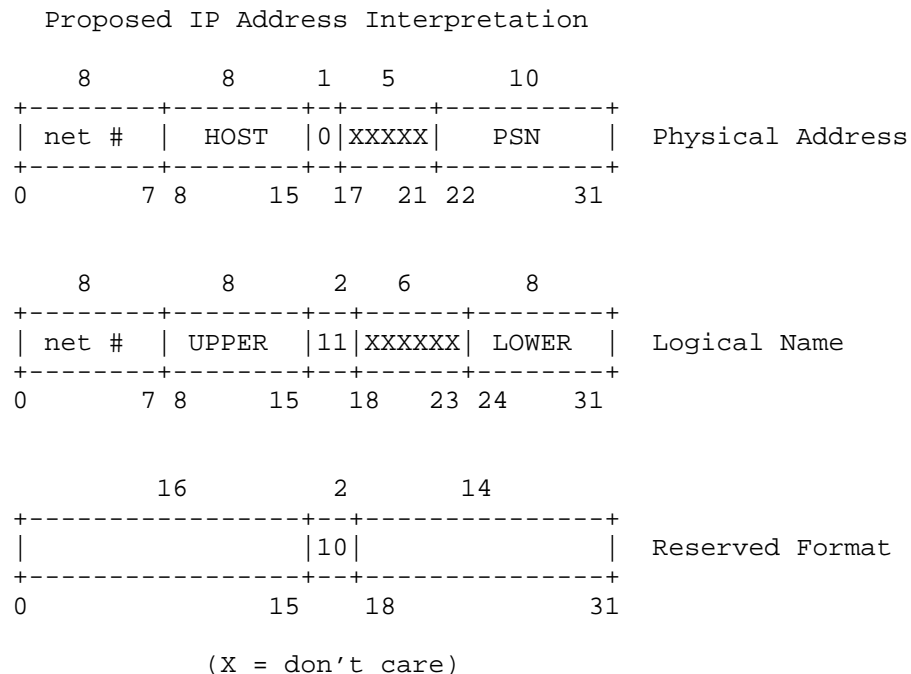
The AHIP-E IP to host mapping should be able to co-exist in some sense with the IP to host mapping specified by the DDN X.25 Specification [6]. In particular, restricted use of the revised IP to DDN host mapping should produce addresses that are consistent with the current X.25 mapping. In other words, there should be a set that includes "sufficiently many" logical names and physical addresses, with the property that each address/name in the set maps onto the same host under both the AHIP and X.25 mappings.

4. Maximum number of PSNs that can be supported:

The new design should support a maximum of more than 256 PSNs per network.

2.3 New Interpretation of IP Address Fields

The following is the new interpretation of the IP address field, in the context of ARPANET-style networks:



New Class A IP Address Interpretation
Figure 2.2

The fields have the following meanings:




HOST = host-number

PSN = 10 bit PSN-number field

UPPER = upper 8 bits of the 16-bit logical name

LOWER = lower 8 bits of the 16-bit logical name

AHIP-E physical addresses and logical names have the following formats:



AHIP-E Address and Name
Figure 2.3

The reserved address format is currently undefined and will be rejected by the PSN, which will return an error message (message type 6, subtype 3) to the host.

```
-----
|This design does not require the AHIP-E host to do any processing|
|of the address -- the host need only copy bits 8-31 of the IP    |
|address into bits 41-64 of the AHIP leader. The host no longer  |
|needs to zero out bits 49-56 of the AHIP leader. The PSN will   |
|take care of the AHIP to subnet address conversion. In other   |
|words, bits 8-31 of the IP address field should be passed      |
|unchanged to the PSN, which interprets them exactly as shown in |
|figure 2.3.                                                     |
-----
```

2.4 Discussion of the New Mapping

This section presents an evaluation of the design in terms of the requirements in section 2.2

1. Address mapping stability requirement:

Current physical IP addresses will not have to be changed, as long as they have been following the convention of setting LH = 0. This ensures that bit 16 is set to 0, indicating that the address is physical, and that the PSN number comes out right.

2. Existing implementation compatibility:

The design meets this requirement, as the address that gets to the PSN has its second octet = 0, which results in its correct interpretation as a physical address.

3. Compatibility with the current X.25 IP address to DDN host mapping:

The current X.25 IP to HOST mapping [6] is as follows: If $h < 64$, the address is considered physical, i.e., it refers to host h on PSN i . If $h \geq 64$, the address is considered logical, i.e., it refers to the host whose logical name is h concatenated with i .

The design is compatible in a limited sense with the current X.25 logical addressing implementation, as long as logical names are assigned such that host-number > 63 (also PSN-number < 256 which is automatic, given the 16-bit size of the logical name field) and physical addresses are in the range host-number < 64 and PSN-number < 256 , with the appropriate setting of bits 16 and 17 of the IP address field. This works because the X.25 mapping ignores the value of the 1 field, i.e., the third IP address octet.

Given the desire to be able to address more than 64 hosts physically and for PSN numbers > 255 , this address assignment restriction should not be considered permanent, but rather as an interim compromise until the hosts' X.25 implementations are revised to incorporate the new mapping between IP and DDN addresses.

4. Maximum number of PSNs that can be supported:

The design allows addressing of up to 1024 PSNs per network.

2.5 Interoperability between Current AHIP and AHIP-E

This section discusses the interoperability between hosts using current AHIP and AHIP-E. It also discusses the general issue of current AHIP host operation in the AHIP-E addressing environment.

The proposed modifications to AHIP have been designed with backward compatibility in mind. However, note that bits 41-64 of the PSN-to-host leader (see 1822(3.4)) will always contain the physical address of the source host. This means that an error could occur when a host on a PSN numbered greater than 255 attempts to send a message to a host running a current AHIP implementation, which interprets the address of the source host as one with PSN-number < 256.

There are other possibilities for errors, caused by incorrect address translation between IP and current AHIP:

1. A host running current AHIP cannot physically address any host on a PSN numbered greater than 255 (see Figure 3.1). Consequently, an error will result if the host attempts to use an address from the NIC host table that has PSN-number > 255.
2. If a host running current AHIP attempts to use a logical name that it might have in its host table, an error will occur. This is because the logical name flag bits 16 and 17 of the IP address, bits 49 and 50 of the AHIP leader. Recall that bits 49 - 56 of the AHIP leader get set to zero with current AHIP (see figure 2.1).

Since these errors cannot be detected by the subnetwork, it is essential that all hosts implement at least version 1 AHIP-E (see chapter 6) before PSN numbers over 255 and logical names are assigned.

Another aspect of interoperability has to do with the IP LH field, which is currently used by a handful of Arpanet hosts to demultiplex a single host port. The 5 don't-care bits of the physical IP address (bits 17-21) and the 6 don't-care bits of the IP logical name (bits 18-23) can be used for this purpose -- in particular, the use of these bits is divided between the network and external devices, based on administrative agreement. At the very least, the IP addresses of such hosts will have to change to reflect the changed position of the LH field. However, the preferred way to demultiplex a single host port is via the mechanism of logical names. The only change this involves is to get the port expander implementation to look at the entire IP address, rather than just the LH field.

3 LOGICAL ADDRESSING


The modifications to AHIP allow a host to use logical addressing to communicate with other hosts on the network. Basically, logical addressing allows hosts to refer to each other using a logical name (see section 3.1) which is independent of a host's physical location in the network. IEN 183 (also published as BBN Report 4473) [2] 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 in the network (especially important if the host has to move to another physical port).
- 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.
- o Allowing a host to provide services that have their own unique names.

3.1 Addresses and Names


The AHIP-E protocol allows two forms of host specification. The first is a slightly modified version of the form used by the current AHIP protocol, the physical address. The second form is the logical name (the terms "name", "logical name" and "logical address" are used interchangeably in this document).

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




Current AHIP Address Format
Figure 3.1

AHIP-E addresses have the following format:



AHIP-E Address Format
Figure 3.2

Logical names are 16-bit unsigned numbers that serve as a logical identifier for one or more hosts. A logical name is the concatenation of two separate octets in the AHIP leader, bits 41-48 (Upper 8) and 57-64 (Lower 8) in particular.



Logical Name Format
Figure 3.3

3.2 Name Translations

There are a number of factors that determine how a logical name is translated by the PSN 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; and which legal translations should not be attempted because a particular host port is down. These issues are discussed in the following sections.

3.2.1 Authorization and Effectiveness

Every host on a PSN, regardless of whether it is using the AHIP or AHIP-E protocol to access the network, can have one or more logical names. Hosts using AHIP-E can then use these names to address the hosts in the network independent of their physical locations.

At this point, several questions arise: How are these names assigned, how do they become known to the PSNs (so that translations to physical addresses can be made), and how do the PSNs 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 PSNs by the Monitoring Center (MC). 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. Physical addresses, which actually refer to physical host ports, are always authorized in this sense.

When the PSN detects that a host has come up on one of its ports, it makes effective the default name(s), if any, for that host. This default action is specified in the configuration table for that host, and can be one of the following: Enable All Names, Enable No Names, Enable One Particular Name. In the case of an AHIP-E host, the default name might not be the one that the host desires to be known as (recall that several hosts may share the same port, or one host may prefer to be known by different names at different times). This requires that an AHIP-E host be able to declare its name to the PSN. This function is performed by a new host-to-PSN message, the Name Declaration Message (NDM), which lists the names that the host would like to be known by. The PSN 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.

If a host is using the current AHIP protocol, it can still receive messages from hosts via its logical name. Of course, it can also receive messages from a current AHIP host via its physical address as well. (Remember, the distinction between logical names and physical

addresses is that the addresses correspond to physical locations on the network, while the names are strictly logical identifiers).

The third question above has by now already been answered. An AHIP-E host can use the NDM message to tell the PSN which host it is (which names it is known by). Thus, even if this is a shared port, the PSN knows which host is currently connected.

WHENEVER A HOST GOES DOWN, ITS NAMES AUTOMATICALLY BECOME NON-EFFECTIVE. When it comes back up, the default action (from the host's configuration) is taken. If the host wishes to be known by a name other than the default, it will have to issue a NDM. It will also have to do this upon receipt of reset NOPS from the PSN.

3.2.2 Translation Policies

Several hosts can share the same logical name. If more than one of these hosts is up at the same time, any messages sent to that logical name will be delivered to just one of the hosts sharing that name, and a RFNMM 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 logical name in this manner.

Similarly, when a host is multi-homed, the same logical 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 physical address or a specific logical 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 administratively 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 PSN's translation tables, to find the first reachable physical host address. This list is always searched from the top whenever a new virtual circuit connection has to be created. This is the most commonly used policy.

- o Selection of the closest physical address, which uses the PSN's internal routing tables to find the translation to the destination PSN with the least cost path for the particular type-of-service whenever a new virtual circuit connection has to be created.
- o Use load leveling. This is similar to the first policy, but differs in that searching the address list for a valid translation starts at the address following where the previous translation search ended whenever a new virtual circuit connection has to be created. This attempts to spread out the load from any one PSN's hosts to the various host ports associated with a particular name. Note that this is NOT network-wide load leveling, which would require knowledge about flows throughout the network.

3.2.3 Reporting Destination Host Downs

As is explained in Report 1822, whenever regular messages are sent by a host, the PSN opens a virtual circuit connection to each destination host from the source host. A new connection is opened for each new source-address/destination-name (or address, as the case might be)/handling-type/type-of-service combination. 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. Connections are also closed after a period of inactivity.

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 a physical address is being used to specify the destination host, then the source host will receive a type 7, subtype 0 (Destination Host Dead) message from the PSN.
- A2. If a logical name is being used to specify the destination host, and the name maps to only one authorized host port, then a type 7, subtype 0 message will be sent to the source host.
- A3. If a logical name is being used to specify the destination host, and the name maps to more than one authorized host port, then the PSN 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 (AHIP Name or Address Error), subtype 5 (no effective translations) message (see section 5.2). 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. However, in the case of a version 0 or 1 host, a type 7, subtype 0 message will be returned for each outstanding message. See chapter 6 for further details on version numbers.

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 a physical 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 a logical name is being used to specify the destination host, then the source host will receive a type 9 (Incomplete Transmission), subtype 6 (message lost due to logically addressed host going down) 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). However, in the case of a version 0 or 1 host, a type 7, subtype 0 message will be returned for each outstanding message. See chapter 6 for further details on version numbers.

3.3 Establishing Host-PSN Communications

When a host comes up on a PSN, or after there has been a break in the communications between the host and its PSN (see 1822 (3.2)), the orderly flow of messages between the host and the PSN needs to be properly (re-)established. This allows the PSN and host to recover from almost 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 PSN are three NOPs. Three messages are required to ensure that at least one message will be properly read by the PSN (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 to synchronize the PSN with the host, to inform the PSN about how much padding the host requires between the message leader and its body and to specify the host's AHIP-E version number to the PSN (see chapter 6).

Similarly, the PSN will send three NOPs to the host when it detects that

the host has come up. The NOPs will be followed by an Interface Reset message. These NOPs will contain the physical address of the host interface.

Once the PSN 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.

3.4 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 PSN 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 PSN, which requests the PSN to translate the destination name and return a list of the addresses to which it maps. The PSN 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 (for the particular type-of-service specified in the Name Server Request message) from the PSN. See section 5.2 for a complete description of these messages' contents.

Using this information, the source host could make an informed decision on which of the physical host ports corresponding to a logical name to use and then send the messages to that port, rather than to the name.

The PSN also supports a different type of name service. A host needs to issue a Name Declaration Message to the PSN in order to change its effective names, but it may not wish to keep its names in some table or file in the host. In this case, it can ask the PSN 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 PSN, and the PSN replies with a type 12 (Port List Reply) message. It contains, for the host port over which the PSN 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 5.2 contains a complete description of the reply's contents.

4 OTHER CHANGES

This section describes the enhancements to the AHIP protocol involving type-of-service specification, subnet congestion feedback and network precedence level feedback. Note that only version 2 hosts will receive the congestion and precedence messages described in this section.

4.1 Type-of-Service Specification

Bits 9 and 10 of the AHIP leader, currently unused, will be used by the host to specify desired delay and throughput characteristics to the PSN. Bit 11, also currently unused, will be used to specify reliability. The bits have the following meaning:

Bit 9: delay bit

0 -- normal delay
1 -- low delay

Bit 10: throughput bit

0 -- normal throughput
1 -- high throughput

Bit 11: reliability bit

0 -- normal reliability
1 -- high reliability

The values of these bits are consistent with those of IP, and bits 11, 12 and 13 of the IP header can be copied directly into bits 9, 10 and 11 of the AHIP leader.

The type-of-service bits should be considered as extensions of the "Handling Type" field (bits 33-40 of the AHIP leader -- see 1822 (3.3)). Messages from host A to host B using the same destination name and of the same handling type and type-of-service will use the same connection, while those that differ in either type-of-service, destination name or handling type will use separate connections. In other words, for a given source host and destination name pair, a new connection will be established whenever a message with a new handling-type/type-of-service combination is received.

4.2 Subnet Congestion Feedback

This section describes the new messages that are part of the mechanism used by the PSN to communicate subnetwork congestion information to the host. Note that a host will be blocked by the PSN when its share of buffers in the PSN is used up. Thus, this information, which is communicated on a connection basis, will give the host an opportunity to selectively reduce its congesting flows, thus preventing all of its flows from getting blocked. Currently, a host has no way of knowing which of its flows is experiencing congestion; consequently, it is possible that one congesting flow can result in the blocking of all the host's flows.

Three new PSN-to-host messages have been created. These messages are:

1. STOP: Blocking Imminent -- Stop Sending on this Connection (Message type 13)
2. SLOW: Subnet Congestion -- Send at Slow Rate on this Connection (Message type 14) -- Maintain Window Size of 1, i.e., do not send a new message to this destination host with this type-of-service and handling type until all previous messages have been acknowledged by RFNMs.
3. GO: Congestion Subsided -- Send at Regular Rate on this Connection (Message type 16) -- Maintain Window Size of 8

These messages may be sent in any order and correspond to states, not transitions. A participating host should support three states with effective windows of 8, 1 and 0. The format of these messages can be found in section 5.2.

4.3 Precedence Level Information

Two new messages have been created:

1. Network Not Accepting Messages at this Precedence Level (Message type 9, subtype 7).
2. Network Precedence Level Cutoff Change (Message type 17).

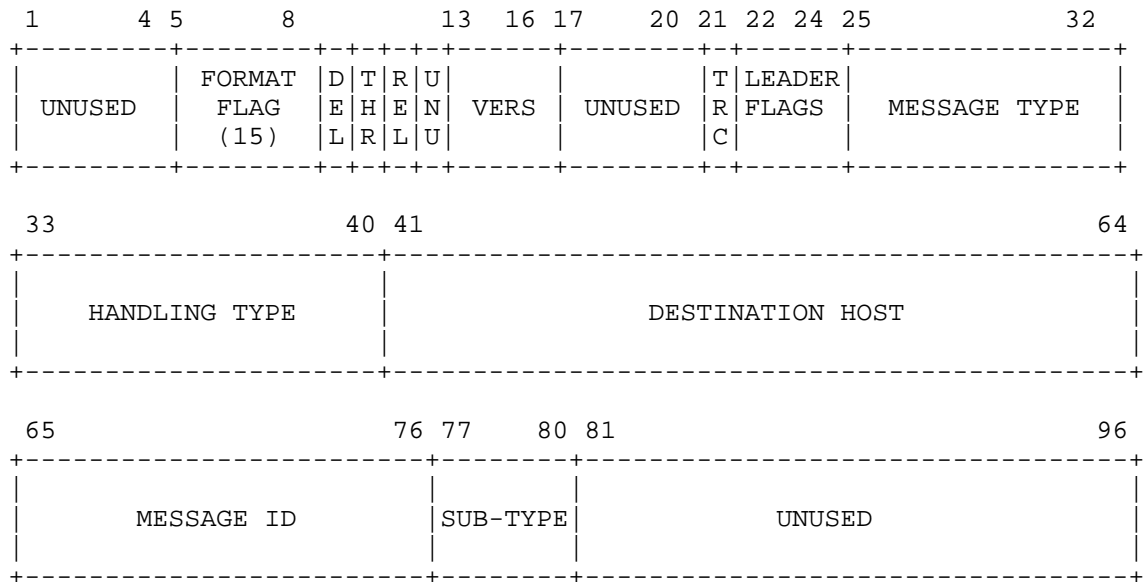
The first message will be generated whenever the host attempts to send a message at a precedence level lower than the cutoff. The cutoff represents a precedence level below which no traffic may be submitted

into the subnetwork; note that a cutoff set to the lowest possible precedence level implies that no precedence restrictions are in effect. If the host has chosen not to receive the new AHIP-E messages, then the PSN will send a type 7, sub-type 3 message (communication with the destination host is administratively prohibited) instead. The second message will be generated whenever the network precedence level cutoff changes. Both messages contain the network precedence cutoff value. The format of these messages can be found in section 5.2.

5 FORMATS FOR NEW AHIP-E MESSAGES

The following sections describe the formats of the leaders that precede messages between an AHIP-E host and its PSN. The formats are almost identical to those of AHIP (see 1822(3.3) and 1822(3.4)). New message types are marked by margin bars (as shown | here).

5.1 Host-to-PSN AHIP-E Leader Format



Host-to-PSN AHIP-E Leader Format
Figure 5.1

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

Bits 5-8: Format Flag
This field is set to decimal 15 (1111 in binary).

Bits 9-11: Type-of-Service

- Bit 9: Delay Bit:
 - 0 -- normal delay
 - 1 -- low delay
- Bit 10: Throughput Bit:
 - 0 -- normal throughput
 - 1 -- high throughput
- Bit 11: Reliability Bit:

- 0 -- normal reliability
- 1 -- high reliability

Bit 12: Unused, must be set to zero.

Bits 13-16: AHIP-E Version number

Ignored by the PSN except in the case of a NOP -- see chapter 6.

Bits 17-20: Unused, must be set to 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 AHIP(3.3) for a description of its use by the PSN'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 PSN uses its full message and error control facilities, and host blocking may occur.
- 3: Uncontrolled Packet - The PSN 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). 1-2,4-15: 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 logical names is or is not effective (see section 3.2.1), 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 logical names contained in the message. This is followed by the logical name entries, each 32 bits

long, of which the first 16 bits is a logical 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. Note that only the names explicitly in the NDM will remain enabled after the NDM is processed (assuming that they are authorized). The PSN will reply with a NDM Reply message (see section 5.2) 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, and without any leader padding):



NDM Message Format
Figure 5.2

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

Type 4: NOP -- see 1822(3.3). Bits 13-16 of the NOP leader are used to determine the host's AHIP-E version -- see chapter 6.

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

Type 11: Name Server Request - This allows the host to use the PSN's logical addressing tables as a name server. The destination name in the AHIP-E leader is translated, and the PSN replies with a Name Server Reply message, which lists the physical host addresses to which the destination name maps. The type-of-service bits (bits 9-11) should be set correctly by the host, as the Name Server Reply message contains information about characteristics of the subnetwork route(s) to that destination, which will depend on the type-of-service.

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 PSN. The PSN 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 33-40: Handling Type:

The top two bits (33 and 34) specify the precedence of the connection. There are 4 precedence levels, level 3 being the highest and level 0 the lowest. Bits 35-40 are used to specify up to 64 separate connections at a particular precedence level and type-of-service.

Bits 41-64: Destination Host:

This field contains the name or address of the destination host, as described in figures 3.3 and 3.2 respectively. 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.

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

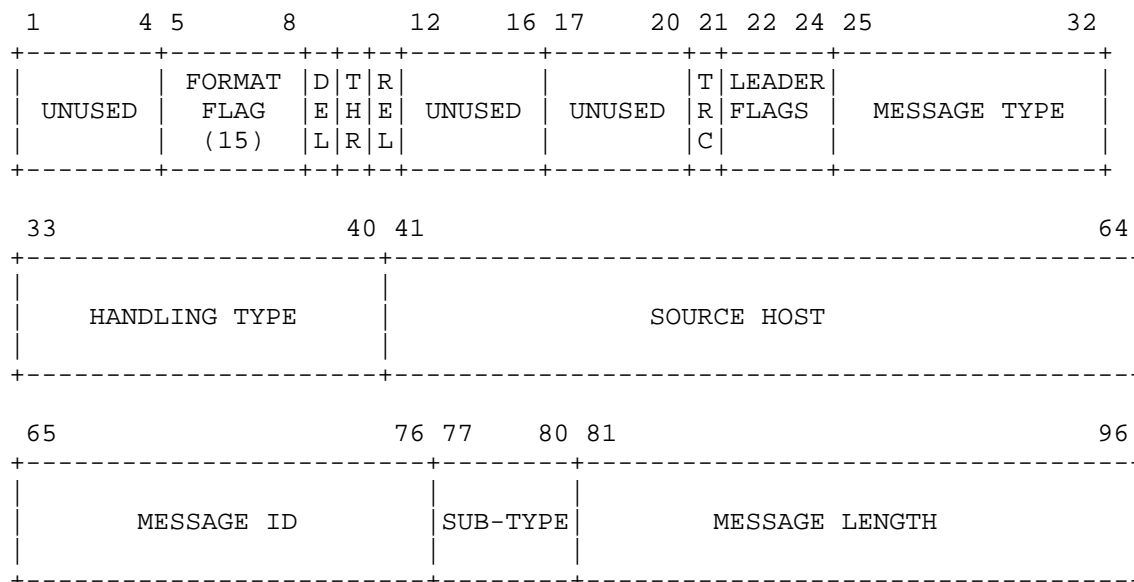
Assigned Numbers [3] appropriate for the host-to-host protocol being used.

Bits 77-80: Sub-type:

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

Bits 81-96: Unused

5.2 PSN-to-Host AHIP-E Leader Format



PSN-to-Host AHIP-E Leader Format
Figure 5.3

Bits 1-4: Unused and set to zero.

Bits 5-8: Format Flag

This field is set to decimal 15 (1111 in binary).

Bits 9-11: Type-of-Service

Specified by the source host (see section 5.1).

Bits 12-20: Unused, must be set to zero.

Bit 21: Trace Bit:

If equal to one, the source host has 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 that sent in the host-to-PSN leader (see section 5.1).

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

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

Type 3: NDM Reply - This is a reply to the NDM host-to-PSN message (see section 5.1). It has the same number of entries as the NDM message to which it replies, and each listed name is accompanied by a zero or a one (see figure 5.2). 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 PSN/host communication. The Destination Host field will contain the physical 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 PSN Dead (or unknown) - See 1822(3.4).

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

Type 9: Incomplete Transmission - See 1822(3.4). In addition to its already defined sub-types, this message has two new sub-types:
6: Logically Addressed Host Went Down - A logically

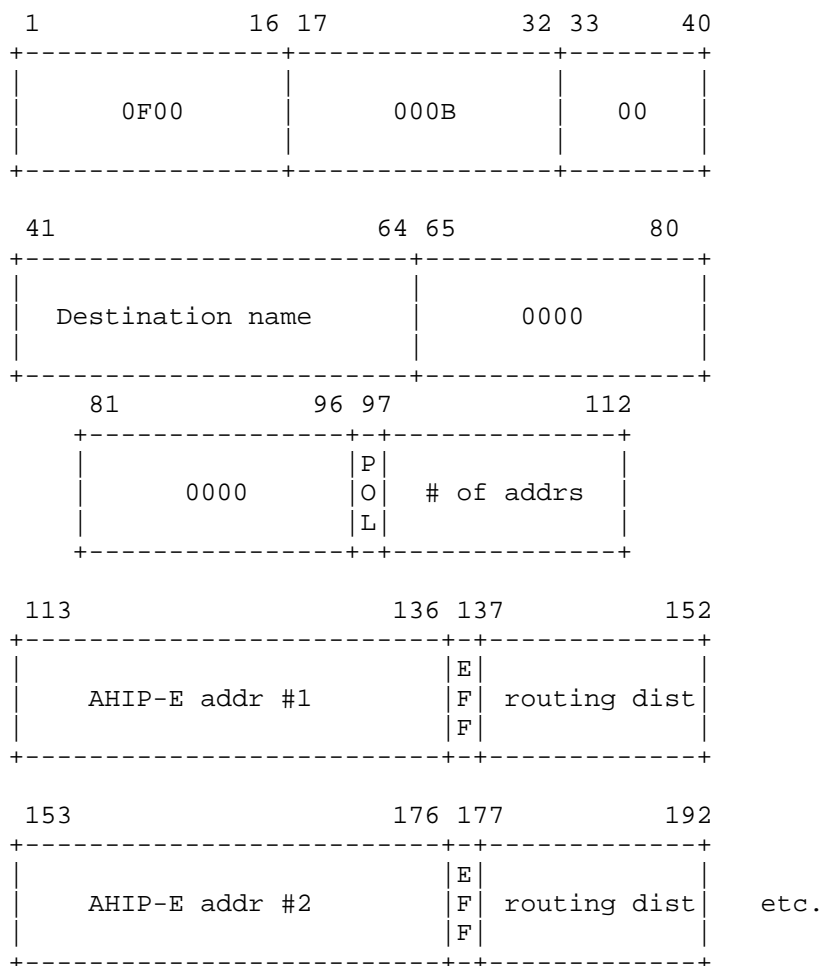
addressed message was lost in the network because the destination host to which it was being delivered went down. The message should be resubmitted by the source host, since there may be another effective host port to which the message could be delivered (see section 2.2.3).

7: Network Not Accepting Messages at this Precedence Level - bits 33 and 34 encode the minimum precedence level currently being accepted by the network. See section 4.3.

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

Type 11: Name Server Reply - This reply to the Name Server Request host-to-PSN message contains, following the leader and any leader padding, a word with the selection policy and the number of physical addresses to which the destination name maps, followed by five octets per physical address: the first three octets contain an AHIP-E address, and the last two contain a bit signifying whether or not that particular translation is effective and the routing distance (expected network transmission delay, in 6.4 ms units) to the address's PSN for the type-of-service specified in the Name Server Request being replied to. This type-of-service will be included in the Name Server Reply leader. In figure 5.4, which includes the leader without any leader padding and has type-of-service set to 000, EFF is 1 for effective and 0 for non-effective, the destination name is in the format of figure 3.3, and POL is a two-bit number indicating the selection policy for the name (see section 3.2.2):

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



Name Server Reply Format
Figure 5.4

Type 12: Port List Reply - This is the reply to the Port List Request host-to-PSN message. It contains the number of names that map to this physical host port, followed by two words per name: the first word contains a logical 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 5.2).

Type 13: STOP -- Stop Sending on this Connection. See

section 4.2.

Type 14: SLOW -- maintain window size of 1 on this connection. See section 4.2.

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

2: The Destination Host 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.

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 PSN may be able to attempt to send the packet.

7: Logical addressing is not in use.
The PSN has no table of mappings from logical addresses to physical host ports.

0, 1, 4, 8-15: Unassigned

Type 16: GO -- maintain window size of 8 on this connection. See section 4.2.

Type 17: Network Precedence Level Cutoff Change -- bits 33 and 34 encode the minimum precedence level currently being accepted by the network. See section 4.3.

Types 18-255: Unassigned.

Bits 33-40: Handling Type:

This has the value assigned by the source host (see 1822(3.1)). This field is only used in message types 0, 5-9, and 13-16.

Bits 41-64: Source Host:

See 1882(3.4). For type 0 messages this contains the physical address of the source host, in the format detailed in figure 3.2. For type 4 messages, this contains the physical address of the local host. For messages of type 5-9, 11 and 13-16 which are responses to messages from the

local host, this contains the destination name as specified in the message from the local 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 5.1). 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 messages only, and is the actual length in bits of the message (exclusive of leader, leader padding, and hardware padding) as computed by the PSN.

6 AHIP-E VERSIONS

This specification provides three versions of AHIP-E and allows a host to specify its version in bits 13-16 of the leader of the NOP. The PSN will set the version of a host based on the value contained in the most recent NOP that it has received from the host. Thus, a host can change the PSN's idea of its version by issuing a NOP containing a different version value. Note that the version field in all other host-to-PSN messages will be ignored by the PSN.

Version 0:

A host that doesn't change its current AHIP implementation will presumably have the version bits in the AHIP leader set to zero. Version 0, thus, is nothing but current AHIP.

A version 0 host will not receive any of the new AHIP-E messages from the PSN, nor will the PSN expect any of the new host-to-PSN message types from the host. The type-of-service bits will always be set to zero in the PSN-to-host leader.

Version 1:

A version 1 host will be able to use logical names to address other hosts, will be able to use the 10-bit PSN field, will be able to specify desired type-of-service to the PSN, but will not receive any of the new AHIP-E messages from the PSN. The PSN will not expect any of the new host-to-PSN message types from the host either.

To implement version 1, a host need only make the following changes to its AHIP implementation:

1. Set the version number field to 1 when sending type 4 messages (NOPs).
2. When sending type 0 messages, copy IP address bits 8-31 into bits 41-64 of the AHIP leader.
3. When sending type 0 messages, copy IP header bits 11-13 to AHIP leader bits 9-11.

Version 2:

A version 2 host is one that is fully compliant with the AHIP-E protocol as described in this document. In addition to being able to take advantage of the features described under version 1 above, it should be able to send and receive all the new AHIP-E messages described in this document.

7 REFERENCES

- [1] "Specifications for the Interconnection of a Host and an PSN", BBN Report 1822, as found in "DDN Protocol Handbook", December 1985, vol. 3, section 3.10.
- [2] 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.
- [3] J. Reynolds and J. Postel, "Assigned Numbers", Request For Comments 990, November 1986.
- [4] J. Postel, ed., "Internet Protocol -- DARPA Internet Program Protocol Specification", Request for Comments 791, September 1981.
- [5] J. Postel, "Address Mappings", Request for Comments 796, September 1981, as found in "DDN Protocol Handbook", vol. 3, section 3.4.
- [6] "Defense Data Network X.25 Host Interface Specification", pp. 497-498, DDN protocol handbook, vol. 1, December 1985.