

Internet Area
Internet Draft
Expires May 1993

Tom Evans
Webster Computer
Christopher Ranch
Novell, Inc.
November 1992

A Method for the Transmission of Internet Packets Over AppleTalk Networks [MacIP]

Status of this Memo

This document is an Internet Draft. Internet Drafts are working documents of the Internet Engineering Task Force (IETF), its Areas, and its Working Groups. Note that other groups may also distribute working documents as Internet Drafts.

Internet Drafts are draft documents valid for a maximum of six months. Internet Drafts may be updated, replaced, or obsoleted by other documents at any time. It is not appropriate to use Internet Drafts as reference material or to cite them other than as a ``working draft'' or ``work in progress.'' Please check the `1id-abstracts.txt` listing contained in the `internet-drafts` Shadow Directories on `nic.ddn.mil`, `nnsf.nsf.net`, `nic.nordu.net`, `ftp.nisc.sri.com`, or `munnari.oz.au` to learn the current status of any Internet Draft.

Abstract

This Internet Draft describes an existing protocol for transporting IP packets over AppleTalk networks. It is intended to specify, standardize, and enhance existing implementations. Distribution of this memo is unlimited.

This Internet Draft is a product of the Apple-IP Working Group of the Internet Engineering Task Force (IETF).

Table of Contents

(To be generated and inserted later)

Overview

There is much confusion over what MacIP really is, as it has grown by many small increments designed and implemented by at least as many people. There has been no central or formal design document, so each of the protocol additions and their effects are not well understood. We hope to dissect the anatomy of the current state of the art, and recommend a functional subset of the protocol features and a new

gateway architecture that will help provide services to multiple zones more reliably.

The MacIP protocol is used to transport IP packets across AppleTalk networks. There are many existing host and gateway implementations of this protocol, and most of them have many slight differences.

We'll call the functional subset MacIP-1. It is our intention that it is compatible with as many of the current implementations as is practically possible. It also describes a set of existing features that will not work except in specific topologies. A future document, MacIP-2, will describe a protocol that provides the intended features, but does so for all AppleTalk topologies.

The new proposed gateway architecture describes an Internal Virtual MacIP (VIM) , which requires a gateway to maintain an internal virtual network that has all zones that are supported in its zones list. This does not require any modifications on existing MacIP host implementations, and existing gateways will have a simple single-point modification.

1. Introduction

The goal of the MacIP architecture is to provide TCP/IP services and connectivity to computers that are not directly connected to an IP network, but are connected indirectly via an AppleTalk network. Typically these are Apple Macintosh computers, but the use of MacIP is not restricted to them.

IP hosts must be connected directly to appropriate media in order to communicate with other IP hosts or gateways. All Macintosh computers come equipped with LocalTalk, Apple Computer's medium speed network media. LocalTalk is not capable of carrying IP directly. Thus, MacIP gateways have been developed to provide IP front-end forwarding agents on behalf of IP hosts embedded in AppleTalk networks.

It is recommended that hosts use this protocol only if they do not have a direct connection; that is their network media does not support IP. The details IP hosts directly connected to IP networks are covered in other well known RFCs.

1.1. Terminology and Topology

In this document, the term "MacIP" refers to the encapsulation protocol and associated services. Specific parts of the protocol are given more specific names as appropriate. In a "MacIP internet" there are two distinct types of devices.

The first are termed "MacIP hosts", and are usually Apple Macintosh computers. They are running applications over TCP/IP, and can use MacIP to communicate between themselves within the confines of their AppleTalk internet.

When communication is desired with common IP devices on IP supported networks (hereafter called "IP hosts"), the services of the second MacIP device, a "MacIP gateway", is required. A MacIP gateway forwards IP datagrams between IP supported networks and MacIP devices on AppleTalk networks, as well as provides other server-type functions.

The term "MacIP Packets" specifically refers to IP datagrams encapsulated in AppleTalk packets that are sent between the forwarding modules in MacIP hosts and gateways.

The term "MacIP Range" refers to the set of IP addresses configured into a MacIP gateway that are considered to belong to the MacIP hosts being serviced by that MacIP gateway.

Other terms will be defined as they are used, and most appear in the Glossary.

The AppleTalk protocols used by MacIP are detailed in [section 8](#) which describes DDP (Datagram Delivery Protocol), ATP (AppleTalk Transaction Protocol) and NBP (Name Binding Protocol). If you are not familiar with AppleTalk, then please read the relevant sections in Inside AppleTalk [[1](#)].

[1.2.](#) Intended Audience

It is expected that there are five different groups of people likely to be reading this document: MacIP host implementors, MacIP gateway implementors, system managers in trouble, the terminally curious, and Jon Postel.

[1.3.](#) Assumptions

This document assumes the reader is familiar with the AppleTalk suite of protocols. AppleTalk is documented in "Inside AppleTalk" [[1](#)]. It is also assumed that the reader is familiar with the IP suite of protocols, particularly IP [[2](#)], IP over Ethernet [[4](#)], Ethernet ARP [[3](#)], RIP [[9](#)], subnetting and routing [[6](#)], IP host requirements [[10](#)], and as documented in other various RFC's.

This document is purposefully confined to the description of the two port gateway, where one port is connected to an AppleTalk network and one port is connected to a Ethernet IP network. This is for

Evans and Ranch

November 11, 1992

Page 3

MacIP

November 1992

descriptive simplicity and should not restrict implementations in any way.

The described AppleTalk port of the gateway and of the host does not need to be restricted to LocalTalk as AppleTalk may be transmitted over other media, such as Ethernet (EtherTalk), Token Ring (TokenTalk), a Serial connection (ARAP), or anything else.

1.4. Structure of this Document

MacIP is a complex protocol; there are multiple modules in both the gateway and the host. Some of them operate in a client/server mode. Others operate peer-to-peer and/or peer-to-proxy.

The original protocol specification, MacIP-1 is described in sections 2 and 3. It begins with a description of the whole system, moves on to detailed descriptions of the individual modules, then describes the protocols used.

[Section 4](#) discusses MacIP limitation, and recommends a subset of features and the Virtual Internal MacIP architecture. These limitations and recommendations are for informational purposes.

[Section 5](#) provides a brief synopsis of AppleTalk, and discusses MacIP required variations of the AppleTalk Name Binding Protocol, NBP.

Sections [6](#) and [7](#) provide protocol constant definitions and a glossary.

[Section 8](#) provides implementation notes on many of the previous sections. Its section numbering is discontinuous, and follows which previous section the note applies to. These notes are for informational purposes.

2. MacIP Protocol Overview

MacIP must satisfy requirements imposed by IP, AppleTalk (except where noted), the Macintosh, and the MacIP gateway.

2.1. Required Functionality

2.1.1. Basic Requirements

IP connectivity is required between MacIP hosts embedded in an AppleTalk network, and IP hosts elsewhere on an IP internet.

Evans and Ranch

November 11, 1992

Page 4

MacIP

November 1992




Figure 1. Example MacIP Internet

There are four different cases to consider in the internet shown in Figure 1.

1. Between a MacIP host and an IP Host (H1 and IP3 via GW1).
2. Between MacIP hosts on different AppleTalk networks via intervening MacIP gateways (H1 and H4 via GW1 and GW2).
3. Between MacIP hosts in the same zone connected to the same MacIP gateway (H1 and H2 via GW1).
4. Between MacIP hosts in the same zone without a MacIP gateway being present (H1 and H2 directly).

Cases 1 and 2 are very similar. Given that MacIP to IP connectivity is established, there is no difficulty supporting MacIP to IP to MacIP. Cases 3 and 4 may appear identical, but are not. Traditional implementations require that MacIP packets must follow the most "efficient" route, and not go through the gateway when the hosts are

on the same network. The following is a simple requirement matrix.

From/To	Local with No Gateway					
Local	Local	Remote	IP Host	Local with No Gateway		
Local	Yes	Yes	Yes	Yes		
Remote	Yes	Yes	Yes	-		
IP Host	Yes	Yes	-	-		

Table 1. MacIP-1 Requirement Matrix

2.1.2. IP Requirements

In order to function as an IP host, a minimum of two things are required:

1. An IP Address for the host.
2. A means by which to forward packets to the host.

With these, an IP host can function on a point-to-point link. In order to function on a broadcast network (such as Ethernet or LocalTalk) the following is also required:

3. An Address Resolution scheme (ARP) to resolve IP addresses to native network addresses.
4. A subnet mask to allow local and remote routing decisions to be made.
5. Gateway address to send off-subnet packets to (more sophisticated routing is sometimes desirable, but not essential).

The MacIP protocol can provide all five, although the last two add significant complexity. A variation on ARP provides this functionality.

2.1.3. Macintosh Considerations

Macintoshes are mobile, and are often moved from one network to another. The AppleTalk protocol handles this automatically without requiring any user reconfiguration of the Macintosh. It would be inconvenient to have to reconfigure the IP Protocol stack under these

circumstances. The MacIP protocol provides an automatic IP Address Assignment protocol that can be used to circumvent the requirement for reconfiguration when moved. IP addresses so assigned are called "Dynamic" Addresses.

Because Macintoshes often are connected directly to Ethernet, the IP protocol stacks usually support both direct Ethernet and MacIP modes of connection. The MacIP protocol is modular and designed to attach simply to an existing Ethernet implementation.

2.2. Protocol Relationships

In order to provide the required functions, MacIP has the following relationship to the other protocols in MacIP hosts and gateways:

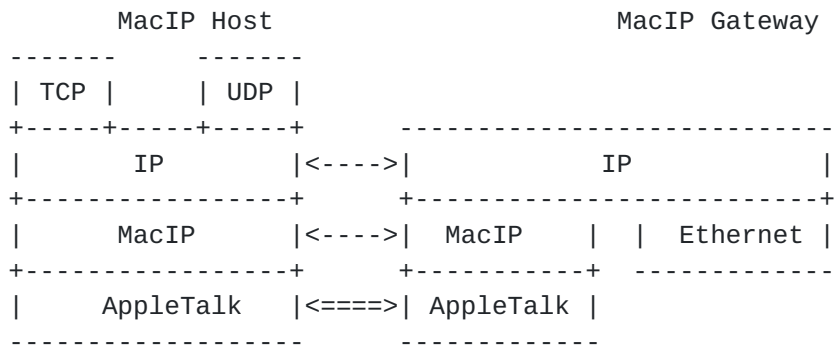


Figure 2. MacIP Host and Gateway Connections

MacIP acts as a Link-layer protocol to IP, while acting as a client of all the session, transport and network layers of AppleTalk. This should serve to warn readers of the complexities ahead:

Layer Name	Protocol

4 - Transport	TCP UDP
	+-----+-----+
3 - Network	IP
	+-----+
2 - Link to IP
5 - Session to	MacIP
AppleTalk	+-----+-----.




Figure 3. Relation Between IP and MacIP

[Phil doesn't like the above diagram. Shall it go in the appendix?]

2.3. Protocol Mapping

MacIP uses the zone-wide protocol NBP to perform the ARP function. This makes for a very simple ARP implementation on a Macintosh as most of the code is already built in to the operating system. It does lead to the following unfortunate restrictions:

1. There has to be one MacIP gateway per zone,

Evans and Ranch

November 11, 1992

Page 7

MacIP

November 1992

2. There can't be more than one MacIP gateway per zone, and
3. MacIP hosts can't use a MacIP gateway in a "remote" zone.

There have been many attempts to overcome the above restrictions, but they are all "outside" of MacIP-1 and cause problems of their own.

2.4. MacIP Functions and Services

MacIP provides address assignment, address resolution and packet transport services.

2.4.1. Address Assignment

The MacIP gateway contains an Address-Assignment module which is configured with a set of IP addresses to assign to MacIP hosts. The module advertises its presence on the network with NBP registration, lookup, and confirmation. The MacIP gateway is discovered by a MacIP host during the initialization of the MacIP host's protocol stack, then an IP address can be requested and granted. MacIP also allows for a MacIP host to be assigned a fixed or "Static" IP address within a range of addresses known to the MacIP gateway.

2.4.2. Address Resolution

Ethernet-connected IP hosts use the Ethernet Address Resolution Protocol (ARP) to discover the hardware address corresponding to a required IP address. The AppleTalk NBP protocol provides similar capabilities and is used to implement the address resolution function in MacIP. This is referred to as NBP ARP.

When any device supporting MacIP acquires an IP address, it registers it through its local NBP process. It is then visible to NBP ARP requests from other MacIP devices. There is the added advantage of possibly discovering configuration errors caused by duplicate IP addresses, as NBP guarantees unique registration within the local zone.

With Ethernet ARP, the "working range" of an ARP request is the IP subnet, which corresponds to the "reach" of the Ethernet broadcast packet. With NBP ARP, the "broadcast reach" corresponds to an AppleTalk construct called a "Zone". A zone consists of one or more AppleTalk networks, the actual topology of which is controlled by the network administrator.

The zone that the MacIP gateway and the hosts are in is referred to as the "MacIP zone". The zone that a particular device is in is

referred to as its "local zone".

Therefore there is a direct correspondence between an IP Subnet and an AppleTalk Zone for NBP ARP. Unfortunately the same does not apply for the delivery of any other sort of AppleTalk or MacIP packet, as the "broadcast reach" corresponds to a single AppleTalk network.

2.4.3. Transport

Transport of IP datagrams over LocalTalk is achieved by encapsulating them in Datagram Delivery Protocol (DDP) packets and sending them over an AppleTalk internet. The destination device can be another Macintosh computer supporting MacIP or a gateway. The latter can either be explicitly selected or discovered through a proxy-based process.

3. MacIP Protocol Specifics

This section describes the MacIP protocol as originally implemented and documented in the Stanford KIP gateway code. This forms the simplest possible version of MacIP, and one which should be supported by all host and gateway implementations.

3.1. Gateway Addressing Styles

There are two alternative approaches to integrating an AppleTalk network into an IP network. One approach involves treating the AppleTalk network as an IP subnet, with the MacIP gateway assuming the role of an IP router. The alternative is to allocate to the AppleTalk network a small range of addresses "stolen" from the Ethernet IP subnetwork that the MacIP gateway is connected to. In this case, the MacIP gateway forwards IP packets to and from AppleTalk.

The forwarding method is conceptually easier, and thus easier to configure. No large range of subnet addresses needs to be calculated and allocated to the AppleTalk network, and no changes need to be made to the rest of the network.

The routing method is more difficult conceptually and, hence, harder for an administrator to configure. It is, however, more consistent with the requirements of many large sites, and can be more practical in complicated networks. This is especially true if the MacIP gateway will emit Routing Information Protocol (RIP, [\[9\]](#)) packets to inform the Ethernet network of the MacIP AppleTalk subnet.

MacIP gateways conforming to MacIP-1 MAY implement either or both of these styles. This doesn't affect MacIP hosts as they should not be

able to tell the difference.

3.1.1. MacIP Forwarding

When forwarding with the MacIP architecture, the AppleTalk network is treated as an extension of the Ethernet IP network. This is done by situating the "MacIP Range" within the range of IP addresses defined by the Ethernet IP network. When a host on the Ethernet ARPs for an IP address which is in the MacIP range, the gateway will answer, performing the proxy ARP function [\[8\]](#).

For example, if the Ethernet has the IP subnet "192.9.200.0", then the MacIP gateway might be configured to assign the addresses "192.9.200.100" through "192.9.200.150" to MacIP hosts. The gateway

will respond to ARP requests on Ethernet to all these addresses, plus its own IP address.

3.1.2. MacIP Routing

Routing via the MacIP protocol is straightforward from the perspective of IP routing. The gateway is configured with two IP addresses and subnet masks, one for the Ethernet and one for the AppleTalk networks.

As the MacIP gateway is acting as an IP Gateway (and is thus performing IP routing), it is necessary for the TCP/IP hosts on the Ethernet side of the gateway be informed of the existence of the subnet corresponding to the MacIP Range, and that the MacIP gateway is the gateway to this subnet. This can be done via static routing tables or via the RIP protocol. MacIP gateways MAY provide either a full or conservative (the latter only advertises the MacIP subnet) RIP implementation in the MacIP gateway.

3.2. MacIP Initialization

3.2.1. Configuration Required For MacIP Hosts

The MacIP host requires an IP address for configuration. "Dynamic" or "Static" addresses refer to the method by which the address is acquired. A dynamic address is assigned from the MacIP gateway's address range. A static address is assigned at the MacIP host, then confirmed through the MacIP gateway.

3.2.2. MacIP Host Initialization

The initialization code is responsible for finding the MacIP gateway's Address Assignment server using NBP, requesting "server" information, acquiring an IP address, either from the address

assignment service or from a statically-configured address, and registering the MacIP host's IP address with NBP.

3.2.2.1. Locating the MacIP gateway's Address Assignment Server

The Address Assignment Service in the MacIP gateway [Server] is assumed to have registered itself with NBP with a type of "IPGATEWAY". The MacIP host initialization process uses NBP to search for "=:IPGATEWAY@*", which performs a search for all Servers in the same zone that the MacIP host is in. Under MacIP-1 the

implementation assumes that it will only receive one response from one gateway. Multiple gateways in one zone are not covered in MacIP-1.

3.2.2.2. Requesting Server Information

The MacIP host and the Server exchange information using a protocol called "MacIPGP", described later. The MacIP host can optionally send a MacIPGP SERVER request to the Server, and SHOULD then receive a response packet. The information in the response packet is mainly obsolete and not very useful, although the returned IP Broadcast address might be usable from some gateways.

3.2.2.3. Requesting a Dynamic IP Address

If the MacIP host is not configured to use a Static IP address, it sends a MacIPGP ASSIGN request to the Server. It will either respond with an appropriate IP address or an error status and optional message which should be displayed to the user. Errors at this point are non-recoverable.

3.2.2.4. Registering the IP Address

The MacIP host registers its IP address through NBP. It MUST use its IP address in dotted decimal notation as its NBP Name. This representation is the four bytes of the IP address, in network order, in decimal with no leading zeros and separated by periods. For its NBP Type, the string "IPADDRESS" MUST be used. For example, to register the IP address 131.161.1.2, the NBP registration would be "131.161.1.2:IPADDRESS@*".

If the registration fails then it is an indication of a duplicate IP address, a gateway, or a network misconfiguration. The failed address MUST NOT be used. This SHOULD be clearly reported to the user.

The MacIP host MUST register on DDP socket 72. Some current MacIP gateway implementations assume that the MacIP host is using socket 72

whether it is or not, so using anything else is not advisable.

3.2.2.5 Closing Down

When the MacIP protocol stack closes down, it must remove its IP address registration from NBP.

3.2.3. Configuration Required For MacIP Gateways

A MacIP gateway has to be configured with an IP address, a subnet mask and (possibly) a default gateway. It also needs to be configured with the "range" of IP addresses that are to be allocated to the MacIP hosts so that the gateway can provide address assignment and forwarding service to its MacIP hosts.. The "union of all IP addresses that can be assigned to MacIP hosts and that the MacIP gateway is required to forward packets to" is called the "MacIP Range". Historically, this is a single contiguous range, but implementations are not confined to this.

This range contains a "Dynamic" and a "Static" range, either of which may be empty. The "Dynamic" range consists of IP addresses that can be allocated to MacIP hosts on request. The "Static" range consists of IP addresses that can be configured into MacIP hosts that require the same IP addresses all the time. The MacIP gateway will not forward packets to a MacIP host that has an IP address outside of this MacIP range.

3.2.4. MacIP Gateway Initialization

The initialization code in a MacIP gateway is responsible for setting up certain data structures used by other modules, registering the Address Assignment server and certain IP addresses with NBP, and performing an initial search for already- registered MacIP hosts.

3.2.4.1. Proxy ARP Initialization

If the MacIP gateway is configured in "forwarding" mode , then the ARP module is initialized to respond to the "MacIP range" of IP addresses in addition to other MacIP gateway addresses.

3.2.4.2. Registration of Address Assignment Server

The MacIP gateway MUST register itself through NBP, using the type "IPGATEWAY", on DDP socket 72. It is necessary for the registration to be unique in the zone, and early implementations guaranteed this by using as the NBP name field (which has to be unique), the dotted decimal representation of their IP address.

3.2.4.3. Registration of IP Addresses

The IP addresses that the MacIP gateway has that are within the MacIP Range SHOULD be registered with the NBP protocol on the gateway in the same way that IP addresses are registered on MacIP hosts. This guarantees that MacIP hosts will not succeed in registering the same address in the same zone. Also, this makes the IP addresses visible to both MacIP hosts (that may wish to send datagrams to these IP addresses) and to network management software. If the registration fails then MacIP MAY not be able to function on the gateway, and appropriate actions should be taken. "Passive Registration" ([section 5.3.2.1](#)) SHOULD not be used.

3.2.4.4. Reregistration Function

The MacIP gateway is responsible for assigning unique IP addresses to MacIP hosts. If the gateway has been running, has assigned addresses and is then restarted (or crashes), it is in danger of reassigning the same addresses to other MacIP hosts. In order to recover the previous assignments, the MacIP gateway uses NBP to search for "=:IPADDRESS*", which will locate all IP addresses registered with NBP ARP in the zone. The NBP responses are directed back to the Initialization module. Those discovered IP addresses that are within the dynamic range assignable by the gateway can be used to initialize the assignment table.

It is important that the MacIP gateway attempts to use the same AppleTalk node address after a restart that it had before, as the MacIP hosts will have the node address of the gateway stored in NBP ARP tables and elsewhere. The gateway should not use a "random" node address on restart.

3.3. Proxy ARP

As mentioned in 4.1.1, when configured to act as a "Forwarding Gateway", the MacIP gateway must perform Proxy ARP for the MacIP Range of addresses. This is simply implemented by adding the MacIP Range to the IP Address(es) that the MacIP gateway's Ethernet ARP process will respond to. All IP addresses in the MacIP range are proxied for, whether there are MacIP hosts using these IP addresses or not as this simplifies the implementation.

Proxy ARP MUST be disabled when the MacIP gateway is configured to act as a "Routing Gateway".

3.4. NBP ARP - Address Resolution

Any MacIP device (host or gateway) can resolve an IP address to an

AppleTalk address by using NBP to search for the device that has that IP address registered. The name is the requested IP address in "dotted decimal" notation and the type is "IPADDRESS". NBP requires repeat and delay specifications (the number of retries and the delay between them). These should be set particularly leniently, especially considering that MacIP may be running over WANs and/or ARAP (Apple Remote Access Protocol). Recommended values are given in [section 10](#), "Definitions".

NBP ARP is functionally very similar to Ethernet ARP, and can often be implemented by using the host or gateway Ethernet ARP code and data structures. Both the MacIP host and gateway are assumed to implement a cache for the addresses resolved by NBP ARP.

With Ethernet ARP, the "working range" of an ARP request is the IP subnet, which corresponds to the "reach" of the Ethernet broadcast packet. With NBP ARP, the "reach" corresponds to an AppleTalk construct called a "Zone". A Zone consists of one or more AppleTalk networks, the actual topology of which is controlled by the network administrator.

There is therefore a direct correspondence between an IP Subnet and an AppleTalk Zone, but ONLY when performing NBP ARP, and not when routing certain packets such as those to an IP broadcast address, such as might be used by RIP or RWHO.

[3.4.1.](#) NBP ARP - Details

The NBP ARP module is passed an IP address to resolve by the delivery module. Resolution is first attempted by searching for a matching IP address in the local ARP cache. A successful match should reset any usage timers. If a no match is found, then the address has to be searched for.

The search on the network is made by using NBP to send an NBP LookUp to all devices in the MacIP zone. The entity-name in the LookUp is the dotted-decimal representation of the IP address (see [section 6.2.8](#)). The entity type is "IPADDRESS". The entity zone is the MacIP zone. The retry count and time is as specified in [section 10](#).

The NBP ARP response carries the AppleTalk address of the destination. The full AppleTalk address (net, node and socket - socket 72 is not to be assumed) must be stored in the ARP cache together with the IP address.

[3.5.](#) Delivery

IP packets including the full IP header MUST be encapsulated in DDP

packets of type 22 (decimal). The source and destination sockets of the packet are 72 (decimal) by convention.

The AppleTalk DDP protocol limits the data size of a DDP packet to 586 bytes, which is then the maximum possible Maximum Transmission Unit (MTU) of the AppleTalk network when transporting IP datagrams. The MTU of 576 is more commonly used so as to conform to the minimum IP MTU.

The Ethernet network has an MTU size of 1486 bytes. The smaller MTU size of AppleTalk requires that gateways must fragment Ethernet packets bound for AppleTalk which are larger than the AppleTalk MTU [3].

Packets received by the MacIP host MUST be dropped (not processed) if the destination IP address does not match the IP address of the MacIP host. No "IP Broadcast" destination addresses are allowed. This function can be performed either by the Delivery or IP modules.

3.6. MacIP Routing Decisions

MacIP's design imposes restrictions that renders MacIP hosts incapable of correctly performing IP Broadcasts and of correctly interpreting ICMP Redirects. Both of these features are required in IP implementations. However, current MacIP host implementations do attempt this, so the feature described in the following section was developed.

3.7. Gateway NBP Proxy ARP

In order to support the simple routing method used by the MacIP client, it is necessary for there to be a service in the MacIP gateway that performs the equivalent of a "Proxy ARP" service, but for all of its MacIP clients. The NBP Proxy ARP service must respond to all NBP ARP requests for all IP addresses EXCEPT for the ones in the MacIP range (those allocated to the MacIP clients).

NBP Proxy ARP MUST respond to NBP requests with the IP address being requested, the type "IPADDRESS" and the AppleTalk net, node and socket (72) address of the "Delivery" module. The NBP Proxy ARP module MUST NOT respond to "wildcard" lookups for type "IPADDRESS". NBP Proxy ARP has to be implemented on a "non-standard" NBP implementation described in [section 5.3](#) as "Conditional NBP" (CNBP)..

3.8. MacIP Gateway Protocol

MacIPGP is a simple request-response protocol based on ATP (AppleTalk Transaction Protocol) ALO (at-least-once) transactions.

Evans and Ranch

November 11, 1992

Page 15

MacIP

November 1992

The MacIP host sends an ATP TREQ (ATP Request) packet to the AppleTalk address of the Address Assignment Server, and the MacIP gateway responds with an ATP TRSP (ATP Response) packet.

There are two defined functions which MacIP-1 uses. These are "get server info" (SERVER) and "assign IP address" (ASSIGN).

3.8.1. SERVER Function

The SERVER function returns a list of common server IP addresses, such as name server, file server and the broadcast address. These are usually defined as part of the gateway's configuration and simply passed on via the protocol without interpretation, so their specific contents is not part of the MacIP protocol. Most of these can be considered "obsolete". Their "usual" designation is detailed below.

3.8.2. ASSIGN Function

The ASSIGN function returns an IP address which can be used by the MacIP host to communicate with other IP hosts. The following is the description of the implementation in the December 1986 KIP code in the "ip.at" document.

The gateway is configured with a "Dynamic Range" range of IP addresses and maintains a table of these containing the fields:

IP address; timer; flags; AppleTalk address (including socket);

When the MacIP Client sends an ASSIGN request to the gateway, the gateway searches the table described above. The service tries to reassign the same IP address to the same AppleTalk address if possible. Otherwise any free IP address is used. If an IP address is available, it is sent in an ASSIGN Reply packet, and the timer field of that table entry will be started.

Thereafter, every "Confirm Period" (60 seconds if not configurable), an "echo command"(NBP ARP Confirm) is sent to the client and the timer bumped. Echo replies received will restart the timer. If 5 periods pass with no replies received, that table entry will be available for potential reassignment. In making "new" table assignments the timer field is used to locate the oldest unused table entry. This increases the chances of a given MacIP host to keep

It is important to note that IP addresses assigned via the gateway ATP protocol must be registered and resolved using the NBP ARP technique.

The MacIPGP protocol request return an error if a request other than ASSIGN or SERVER is made. The ASSIGN response may return an error if there are no Dynamic addresses available.

1. Data-link Header,
2. DDP Header,
3. ATP Header,
4. MacIPGP Header,
5. MacIPGP Data Packet (only on MacIPGP Response packets).

[illegible]




Figure 6. MacIPGP Packet

3.8.5. ATP Control Fields

The Control Info, Sequence Number and Transaction ID fields are documented in Inside AppleTalk. The MacIP gateway returns these fields as-is to the MacIP host except that the Control Info field is changed from a TREQ to a TRSP.

3.8.6. ATP User Bytes

These four bytes are unused in MacIP-1 and can be expected to contain random data.

3.8.7. MacIPGP Request and Response

If the ATP Control info is TREQ, then this is a packet from the MacIP host to the gateway. The MacIPGP Data field is empty (zero length). The defined values of the MacIPGP Request field are:

Val.	Name	Meaning
1	ASSIGN	Request assignment of Dynamic address.
3	SERVER	Return Server Information.
-1	ERROR	Only valid in MacIPGP Response packet.

ASSIGN is a request for the MacIP gateway to assign an IP address from its configured "Dynamic Range" to the MacIP host. SERVER is a request for the "Server Information" to be returned.

The MacIP gateway normally returns the packet as-is with the MacIPGP Response field in the returned packet set to the MacIPGP Request field in the received packet. If an error occurred, then the Response field is set to "-1" (hex FFFFFFFF) with an optional zero-terminated

error string returned in the "Error Message" field. The length of the data field of the returned ATP packet is 64 bytes plus the length of the terminated error string. If the string is empty then 65 bytes are returned.

3.8.8. MacIPGP Data Field

The MacIPGP Data Field is returned in MacIPGP Response packets. The format is:

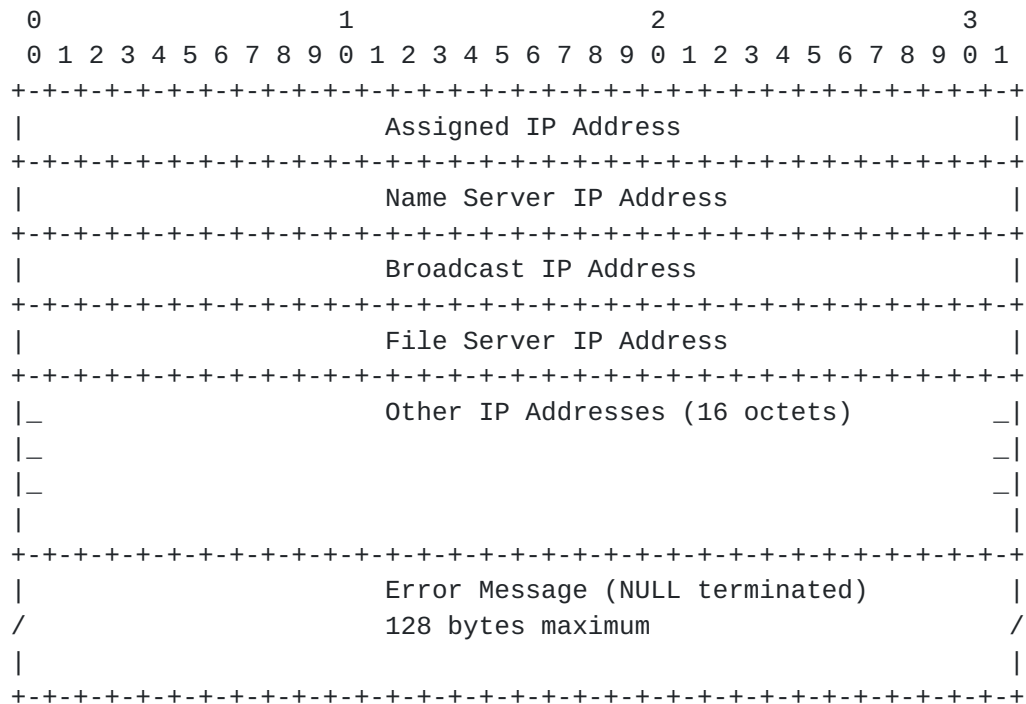


Figure 7. MacIPGP Data Field

Originally the Name-server, Broadcast, File-server and Other fields were simply copies of configuration data transferred to the MacIP gateway by the AppleTalk Administrative Daemon (atalkad). Their meaning was thus a "private matter" between the administrator and the MacIP host code in use, and did not involve the MacIP gateway at all.

It did not interpret or generate the data (except for the Broadcast field which the MacIP gateway does use), it only transferred it.

Some manufacturers have changed this simple relationship by using the atalkatab fields for gateway configuration and/or having the gateway generate the data transferred to the MacIP host. Thus, the meaning is unclear, and thus for MacIP-1, undefined.

3.8.8.1. Assigned IP Address

This is the IP address assigned by the MacIP gateway to the MacIP host. It is only returned in ASSIGN response packets. This field is only valid in ASSIGN response packets.

3.8.8.2. Name Server IP Address

This historically has been used to contain the IP address of an IEN-116 Domain Name Server.

Evans and Ranch

November 11, 1992

Page 19

MacIP

November 1992

3.8.8.3. Broadcast IP Address

This may be the Ethernet IP broadcast address which is only "appropriate" for the MacIP hosts if the MacIP gateway is configured in "Forwarding" mode. If it is in "Routing" mode, this field could be anything. MacIP hosts MUST not rely on this data.

3.8.8.4. File Server IP Address

Originally the address of EFS (Electronic File Server) included with CAP v4 (Columbia AppleTalk Package). Obsoleted by AUFS.

3.8.8.5. Other IP Addresses

Available for "private" use between consenting MacIP hosts and MacIP gateway configurations. These may contain the "IPOTHER" fields from the "atalkatab" file, but they may not.

3.8.8.6. Error Message

If the returned MacIPGP Response code is -1, this may contain a 128-byte maximum null-terminated error message. Otherwise it is a zero-length null-terminated string (one byte of null).

3.8.9. Delivery Packets

If the Address Assignment Service does not have the same AppleTalk

address as the Delivery Module (as returned by NBP Proxy ARP), then it may still receive delivery packets (IP packets of DDP type 22) sent by MacIP hosts disobeying MacIP-1. For backward compatibility, these packets SHOULD be forwarded to the Delivery module.

4. MacIP Limitations and Recommendations

The specification of MacIP imposes certain basic restrictions on operations of both hosts and gateways, particularly on the allowable network topology. Most MacIP gateway vendors have found different, incompatible, and incomplete methods ways to work around these basic protocol limitations.

This specification recommends methods for managing the restrictions on both hosts and gateways. The host recommendations focus on the routing decisions hosts are attempting, and the gateway recommends that MacIP gateways reside on a virtual internal network. This network's zones list contains all zones supported by the MacIP gateway, and provides a more reliable mechanism for NBP ARP. This gateway and topology architecture is referred to as Virtual Internal MacIP (VIM).

Evans and Ranch

November 11, 1992

Page 20

MacIP

November 1992

4.1. MacIP Routing Decisions

MacIP 1 recommends that the MacIP host perform no IP routing decisions whatsoever. There are many advantages to this requirement:

1. It simplifies the MacIP host implementation,
2. It minimizes the network-specific information that has to be sent to or configured into the MacIP host,
3. It allows for MacIP to be extended so that Hosts can use the services of a MacIP gateway from "out of zone".

Because of restrictions imposed by the design of MacIP, MacIP hosts are incapable of performing IP Broadcasts and of correctly interpreting ICMP Redirects, both of which would be expected of a full IP implementation.

To be specific, it is mandatory that MacIP 1 Hosts should not do any routing-related operations and should obey the following:

1. They must use NBP ARP to resolve all IP addresses, no matter what they are, even IP addresses that may be, or

are manifestly in another IP subnet or network, and includes any possible broadcast IP addresses.

2. They must ignore all subnet masks, network and subnet numbers and all possible "default gateway" addresses.
3. They should not send IP Datagrams to either the AppleTalk or IP address of the Address Assignment Server (received as the "name" field of the NBP LookUp).
4. They should ignore all ICMP Redirect messages and RIP packets that they may receive.

The simplest way to disable all off-host routing decisions in an existing IP (or MacIP) implementation is to set the subnet mask to "0.0.0.0", and to disable any other "special-case" or error-checking code that may defeat the intention of this.

The MacIP host code should not assume that the Address Assignment server and the NBP Proxy ARP module have the same socket address, or even the same AppleTalk address, as it is permissible for them to be implemented on different hardware and even on different networks.

4.2. Multiple Gateways in One Zone Limitation

Evans and Ranch

November 11, 1992

Page 21

MacIP

November 1992

As detailed in [section 3.7](#) (Gateway NBP Proxy ARP), having multiple MacIP gateways in the same zone will completely prevent MacIP hosts in that zone from working. The following hacks have been invented to get around this:

4.2.1. NBP Proxy ARP Directed Port

With this method, the MacIP gateway restricts the operation of NBP Proxy ARP to requests received from networks connected via particular physical ports, (usually the LocalTalk port, where the majority of MacIP hosts usually are). This allows multiple MacIP gateways to exist in the same zone as long as the network topology guarantees that there is never more than one MacIP gateway's directed port on one AppleTalk network. This approach also requires NBP (CNBP) in the gateway to not answer NBP LookUps for "IPGATEWAY" unless they originate from the same port that NBP Proxy ARP is enabled on. This asymmetry plays havoc with network management; devices will be visible depending on where in the AppleTalk topology devices are looking from.

4.2.2. NBP Proxy ARP Hop Limit

The DDP Hop Count is available to the NBP Proxy ARP process. It can refuse to answer any requests from MacIP hosts that aren't zero hops away (directly connected). This solves some problems at the expense of restricting the topology. It also will not work if two gateways share a common network.

4.2.3. Friendly Requests - Dynamic

In this method, the NBP Proxy ARP module will only answer requests from the SAME MacIP hosts that the Address Assignment module knows about. The source AppleTalk address of the incoming NBP ARP is compared against all entries in the Assigned Address table. If the AppleTalk address is present (and the MacIP host is not registering its own address), then the NBP Proxy ARP module responds.

This fails completely with MacIP hosts that have Static IP addresses, as they aren't in the in the table.

4.2.4. Friendly Requests - Static

Various methods have been used to force friendly requests to work with static hosts. In all cases the gateway has to somehow record the AppleTalk addresses of MacIP hosts with static addresses. This table has to be filled with periodic "Reregistration" calls (3.3.4.4), or "directed wildcard NBP ARP" calls have to be issued to "suspected hosts". VIM provides a superior solution.

Evans and Ranch

November 11, 1992

Page 22

MacIP

November 1992

4.2.5. NBP Proxy ARP Not Required

If the MacIP host disobeys MacIP-1 and sends all packets to the MacIP gateway then NBP Proxy ARP may not be required, but MacTCP (Apple's MacIP driver for the Macintosh) doesn't allow this.

4.2.6. Other Multiple Gateway Problems

Given a choice of multiple MacIP gateways, MacIP hosts will invariably pick the worst possible one - the one "furthest away". To date, MacIP hosts cannot select which MacIP gateway to attach themselves to. Thus, it is a race condition for NBP responses to return to the MacIP host. Generally, the first response received is used. See section y.y.y "Out-of-order NBP" for a discussion.

4.3. Out-of-Zone Limitation

NBP ARP can only work between devices that are in the same zone. In spite of this, most MacIP implementations try to allow the MacIP hosts to be in a zone different to that of the MacIP gateway (the MacIP zone). The problems for the MacIP host in the wrong zone are:

1. Its NBP ARP Registration won't find duplicate IP addresses,
2. It can't answer NBP ARPs from other MacIP hosts,
3. It can't answer NBP ARPs from the MacIP gateway,
4. The gateway "reregistration" function can't find this host.

The following "Out-of-Zone-Hacks" have been used in existing implementations. We none of them, and suggest VIM instead (described in the next sub-section).

4.3.1. Ask Address Assignment - Dynamic

MacIP hosts that have requested Dynamic IP addresses have their AppleTalk and IP addresses in the Address Assignment Module's table. NBP ARP can be modified to check this table. This is a very incomplete solution as it only solves problem (3) above, while ignoring (1), (2) and (4). It can't work with Statically-addressed hosts at all.

4.3.2. NBP ARP Reverse Resolution

To allow Static-IP address MacIP hosts to operate out-of-zone, it is

possible for the MacIP gateway to "listen" for NBP ARPs (as it does in 4.1.2 Friendly Requests). These may be the Static hosts attempting to "register" their IP addresses in the zone of the MacIP gateway (they should do this in order to discover duplicate registration, but none do). It may be the Static hosts performing NBP ARP in the MacIP gateway's zone. In either case, if the MacIP gateway doesn't find the source AppleTalk address in the table, it then sends a "directed wildcard NBP ARP" to the source AppleTalk address. Any response will be directed by CNBP to the "reregistration" part of the Initialization Module, and will serve to register the address for the next time.

This fails to work for statically addressed hosts that are running an IP service application. IP service applications tend not to send unsolicited packets, and thus no address mapping exists, preventing address resolution. It still doesn't solve problems (1) and (2) either.

4.3.3. Glean From MacIP Packets

Some Static hosts may default to sending MacIP packets to the MacIP gateway. The IP to AppleTalk address mapping can be "gleaned" from these packets. There are two problems here. Firstly it is computationally expensive to "glean" every MacIP packet. Secondly it relies on the host sending the packet to the gateway, and not all do. It doesn't work with hosts running a server application, as they don't send any gleanable MacIP packets.

Gleaning can partly solve the problem that reregistration can't work with out-of- zone hosts - the next MacIP packet from the MacIP host after a gateway restart will update the table.

4.4. Virtual Internal MacIP Recommendation

We recommend that MacIP gateways implement a virtual internal network that has no physical port associated with it, but has a network range and zones list. The zone list contains all the zones that MacIP hosts are going to reside. Then, the MacIP gateway is made visible in all of those zones by registering itself through NBP to all MacIP hosts in those zones. If the MacIP Host is not in the same zone as the MacIP Gateway, then you don't get in.

This simplifies everything back to "Classic 1986 MacIP" which we all have code for. All existing MacIP host code works, both "in" and "out-of" zone (because there aren't any out-of-zones). Pretty much all of the existing MacIP Gateway code should be able to remain "as-is" as well, and should be straight-forward to document and instrument.

It is possible to make the "VIM zone list acquisition" automatic. When a new "MacIP Zone" is found (by a MacIP host trying to register), the gateway uses RTMP to "poison" the old VIM network range (let's say it was AppleTalk network 10-10), creates a new VIM network that overlaps the old one (say network 10-11) with all the old zones and the new one in it, and then advertises that with RTMP. Disabling "automatic zone list acquisition" counts as a marketable

"security feature".

We could then optionally add Phil Budne's idea of having MacIP Gateways search for other MacIP Gateways (using NBP to look for "=:IPGATEWAY@zzz"), and then send RIP packets to them - we've just solved all of the tricky "RIP-in-MacIP" problems too.

For the "efficiency purists" who demand "minimum path" delivery between MacIP hosts that are in different zones, that's now easy to support too. In this case, NBP Proxy ARP would look in its mapping table for the IP address in question, and use the mapped AppleTalk address in the NBP Response's Entity address fields. This effectively provides an address translation service, or could be considered NBP Proxy ARP redirect.

5. AppleTalk Basics and Variations

For those not familiar with AppleTalk, this section gives a very brief summary of the parts of AppleTalk used by MacIP. The reference for AppleTalk is "Inside AppleTalk, Second Edition", published by Addison Wesley.

5.1. AppleTalk Addresses

Devices on an AppleTalk Internet are uniquely identified by a 16-bit network number combined with an 8-bit node number. These addresses are handled very similarly to the net and host part of a Class C IP address. MacIP has no special requirements on AppleTalk addresses.

5.2. AppleTalk Zones

AppleTalk networks are grouped together into named collections called Zones. This is implemented in the routers responsible for the network numbers by associating a Zone Name (or list of zones for extended networks) with each network. Zone names form the user-visible topology of AppleTalk Internets, hiding the actual network-level topology.

5.3. Name Binding Protocol

NBP provides registration and location services for named entities

within zones. A service entity begins by attempting to register a "name" and a "type" with the local NBP software. NBP places the name:type combination into a local registry, so their nodes may

locate it, and then (with the cooperation of the local router) broadcasts "LookUp" packets on every network that is associated with the host's zone. If the name:type is already registered in that or some other host, a "LookUp Reply" packet will be received by NBP and the registration attempt fails. Generally, the service will modify the its name, and re-attempt registration. If no LookUp Replies are received, then the registration is considered successful.

Once a name:type is registered, NBP will answer searches made by other devices for that name:type, and will supply the AppleTalk address (net:node:socket) of the registered entity. Wildcard LookUps are permitted on both the name and type fields.

5.3.1. Strict NBP

The previously-described NBP is the type implemented by the Macintosh Computer OS. It is characterized by "Active NBP Registration" where NBP always performing a "uniqueness-search" on registration, and "Unconditional NBP Replies" where NBP unconditionally answers all NBP LookUps that match registered entities.

5.3.2. Loose NBP

There are a lot of ways to abuse NBP and to step outside the purposes for which it was written. The following are some of the ways that have been found to abuse it, collectively called "Loose NBP". The problems caused are mainly those of confusion, both to network managers and to network management software.

5.3.2.1. Passive NBP Registration

Some NBP implementations perform "Passive NBP Registration" by skipping the "uniqueness search". For the case of registering single unique entities this can only be described as poor practice, as it will fail to detect duplicate registration. It may also confuse network managers who are monitoring the device to see if it is working properly.

In the case of implementing NBP Proxy ARP it is impractical to "Actively Register" the 4,261,412,864 IP addresses that the MacIP gateway is proxying for.

5.3.2.2. Conditional NBP Replies

NBP can also be abused to allow "Conditional NBP Replies" (CNBP)

where the generation of the NBP LookUp Reply does not depend solely on the contents of the local NBP Registry, but depends on the requesting device, or the particulars on what is being asked for. This isn't possible on a Macintosh as NBP is built into the OS. Again NBP Proxy ARP requires this.

5.3.2.3. Out-of-order NBP

A packet monitor observing an NBP transaction would expect to see the following operations in the following order, and may in fact depend on the order to correctly report the operation:

1. MacIP host wishing to search for a particular named entity within its own zone sends NBP BrRq (Broadcast Request) to a local AppleTalk router,
2. The router rebroadcasts the NBP Query as an NBP LkUp (LookUp) on all networks in the requested zone,
3. One or more LkUp Replies are sent from the searched-for entity to the requesting host.

In the case of a MacIP host searching for a MacIP gateway there exists a serious problem. The router that the MacIP host sent the NBP BrRq to is likely to be the most appropriate MacIP gateway. However while this router/gateway is busy performing step (2) above, other MacIP gateways in the same zone are likely to be on step (3). Current MacIP host code will use the first LkUp Reply and will thus select the inappropriately "remote" gateway.

If the first router reverses steps (2) and (3) above, replying to the request before rebroadcasting it, then the MacIP host will select it. It looks weird on a packet monitor though.

5.3.2.4. Port Hopping

NBP is intended to allow a client to search for a service by name, and to discover the network address which will be used for the delivery of subsequent datagrams. In the case of a multi-port service-provider, it is sensible to return the "nearest" AppleTalk address to the client, in this case the address of the port that the NBP Reply was returned to the client through. Unfortunately this port (and address) may not be in the zone that the request was originally made in. This causes no problems apart from confusion to network managers again, and should probably be avoided for this reason.

5.4. DDP

MacIP

November 1992

The Datagram Delivery Protocol is close to the AppleTalk equivalent of UDP/IP. It implements an "unreliable" Datagram Delivery service, with the destination address being a socket at a specific net:node address. DDP packets also have a "type" field which permits another level of multiplexing on top of the socket multiplexor, but is often used redundantly with the socket.

5.5. ATP

The AppleTalk Transaction Protocol adds a request/response/timeout and retry mechanism on top of DDP packet delivery. Transactions commence with an ATP Request packet (TREQ) and must be answered by an ATP Response (TRSP) or the requestor will retry and eventually time out and report back an error to the client.

6. Definitions

6.1. AppleTalk Protocol constants

MacIP MTU	576 bytes
DDP constants	
MacIP packet type	22 (decimal) DDP
ARP packet type	23 (decimal)
DDP ARP constants	
AppleTalk address type	3 (decimal)
NBP constants	
gateway object type	IPGATEWAY registered IP
address object type	IPADDRESS LookUp retransmit
count	4 tries LookUp retransmit interval
seconds	5

6.2. Gateway ATP Protocol Constants

ATP Protocol Constants	
ATP retransmit count	4 tries ATP retransmit
interval	5 seconds
ATP request command codes	
ASSIGN assign IP address	1 NAME name
server	2 (obsolete) SERVER get server
info	3
ATP response codes	
SUCCESS	same as request code

7. Glossary

MacIP

The encapsulation protocol and associated services.

MacIP-1

The original MacIP specification as implemented in the KIP code.

MacIP-2

A future version of MacIP that will allow out-of-zone and multiple-gateway operation.

MacIP Host

A device implementing the host-side of the MacIP protocol, usually an Apple Macintosh computer.

MacIP Gateway

A device implementing the gateway-functions of the MacIP protocol. It converts between the MacIP transport-layer and those used by the other IP hosts, as well as provides other server-type functions.

IP Host

An IP host on an IP internet, usually not using MacIP, but with which MacIP hosts can communicate.

MacIP Packets

IP datagrams encapsulated in AppleTalk packets that are sent between the "Delivery" modules in MacIP hosts and gateways.

MacIP Network

The collection of MacIP hosts and the MacIP gateway that are in direct communication with each other - similar to the concept of an IP Subnet.

MacIP Range

The range of IP addresses configured into a MacIP gateway that are considered to belong to the MacIP hosts.

MacIP Dynamic Range

The IP addresses in the MacIP range that are available for automatic assignment to MacIP hosts by the gateway MacIPGP module.

Dynamic Address

An IP address assigned by a MacIP gateway from its Dynamic range for the use of a MacIP host.

Evans and Ranch

November 11, 1992

Page 29

MacIP

November 1992

MacIP Static Range

The IP addresses in the MacIP range that are available for permanent assignment to MacIP hosts by the system administrator.

Static Address

A fixed IP address assigned by the system administrator to a MacIP host. This address must be in the MacIP Static range of the host's MacIP gateway.

MacIP Zone

The AppleTalk zone that the MacIP gateway is situated in, and which corresponds to the "reach" of the NBP ARP function.

Local Zone

The AppleTalk zone that a MacIP host is in.

MacIP Routing Mode

The MacIP gateway configuration where the MacIP network IP addresses are in a different subnet to that of the IP backbone, and where the MacIP gateway is acting as an IP router.

MacIP Forwarding Mode

The MacIP gateway configuration where the MacIP network IP addresses are in the same subnet as the IP backbone, and where the MacIP gateway is performing Proxy ARP for the MacIP hosts.

In-Zone Mode

The MacIP host configuration where the host is in the same zone as the MacIP gateway (the MacIP zone is the same as the Local zone).

Out-of-Zone Mode

The MacIP host configuration where the host is in a different zone to the MacIP gateway (the MacIP zone is not the same as the Local zone).

MacIPGP

The MacIP Gateway Protocol. Used to implement the address assignment service, and to forward other information.

NBP ARP

Method for resolving an IP address into an AppleTalk address.
Equivalent to Ethernet ARP.

NBP Proxy ARP

Method by which the MacIP gateway answers NBP ARP requests for IP addresses of IP hosts.

Evans and Ranch

November 11, 1992

Page 30

MacIP

November 1992

NBP ARP Forwarding

Method by which the MacIP gateway forward NBP ARP requests to MacIP hosts that ore located outside of the MacIP zone.

Proxy ARP

Method by which the MacIP gateway responds to Ethernet ARP requests from IP hosts for IP addresses in the MacIP range.

VIM

Virtual Internal MacIP. This refers to the proposed new MacIP gateway architcture. The gateway implements an internal, virtual network, and associates itself with it. If the internal network's zones list contain more than one zone name, the MacIP gateway registers itself on all of them, providing gateway services to all MacIP hosts in those zones.

8. Implementation Notes

This section provides notes and insights to the reader. Its sections numbers are organized to follow the preceding section numbers.

8.1.2 Intended Audience

It should be noted that members from the various groups mentioned (host implementors, gateway implementors, system managers, and the curious) have different preconceived notions about what comprises a MacIP protocol system. Also, what is acceptable functionality in undocumented limitations will differ between these groups, and between members of the groups themselves. Careful attention must be paid to the fact that TCP/IP is not AppleTalk, and AppleTalk is not TCP/IP. There are some similarities that are close enough to conceal the important differences. NBP is nothing like DNS. RIP is different from RTMP. UDP isn't DDP. Mapping one protocol system to another presents a significant challenge, and that is what we're attempting with MacIP. It is very easy for things to go very wrong.

8.2.1.1 Basic Requirements

Implementing MacIP so as to fulfill this requirement as well as the local with no gateway requirement (which is very rarely used in practice) is one of the things that makes the protocol so complex.

8.2.3. Protocol Mapping

MacIP follows NBP ARP. The following are two obvious ways to "map" IP onto AppleTalk at the transport level, but are not followed.

NBP ARP is used instead. This combines some of the worst features of

Evans and Ranch

November 11, 1992

Page 31

MacIP

November 1992

"Direct Mapping" and "Server Client" without sufficient advantages to offset it. It also complicates implementation, documentation and installation. But we're stuck with it.

8.2.3.1. Direct Mapping

MacIP provides very similar capabilities to Ethernet and Ethernet ARP, so one possible protocol mapping would be to treat each AppleTalk network as an IP subnet, with IP addresses within each subnet being resolved by an equivalent network-wide broadcast protocol to Ethernet ARP. All routing would be performed by the MacIP gateways, although the hosts would be required to perform "this- subnet" type routing decisions.

This was actually implemented early on in the history of MacIP. It has many advantages. It satisfies the "requirement matrix" of 2.1.1 and its similarity to IP makes it easily understandable and documentable. The downside is that it requires every AppleTalk router in an Internet to also be an IP router. This complicates the configuration of an Internet and also consumes IP address space at an alarming rate. This isn't MacIP-1.

8.2.3.2. Client-Server

An approach leading to a very simple implementation would be a strict client- server one, such as the protocols that are already used to implement DECnet, LAT and SNA services over AppleTalk. MacIP hosts would use NBP to find the gateway(s) and then establish a "session" which would allow the gateway to record the AppleTalk address of the host so it would know how to route a packet back to them.

Advantages would be the ease of implementation, documentation and debugging. The MacIP hosts wouldn't have to know anything about routing - they would send all packets to the gateway. It would also allow MacIP hosts that are anywhere on a large and complex AppleTalk Internet to use a gateway no matter where it was - there would only need to be one gateway. It would allow operation across AppleTalk zones. The downside is that all traffic between two MacIP hosts on the same network would have to pass through the gateway, and that the gateway would be required for the protocol. This isn't MacIP-1 either, much the pity.

8.2.4.4. MacIP Routing Decisions

8.2.4.4.1. Routing in Standard IP Hosts

An IP host on a single point-to-point link only has one simple routing decision to make when presented with an IP packet to forward:

Evans and Ranch

November 11, 1992

Page 32

MacIP

November 1992

1. Is the destination IP address equal to my IP address?

If it is, the packet is sent "inwards". If not, it is sent out the link. An IP host connected to single broadcast medium has another two questions to answer:

2. Is the destination address within "this" network/subnet?

If it is, the packet can be sent "directly" to the destination. If it isn't, then the packet has to be sent via a gateway, often a single "default gateway", the IP address of which is known.

3. Is the destination address a Broadcast address?

In all cases ARP is used to resolve the selected IP address to a hardware address. In case (3), ARP will substitute the appropriate configured hardware broadcast address.

8.2.4.4.2. Routing in MacIP Hosts

This "broadcast medium behavior" is insufficient for MacIP, as the Address Assignment protocol sensibly and deliberately provides neither the subnet mask nor default gateway information. This behavior should not be defeated (by allowing manual entry of the data) as it defeats the Macintosh "plug and play" requirement.

Even if Address Assignment did provide this information, the

disparity between the "reach" of NBP ARP and normal DDP broadcast datagrams prevents a MacIP host from broadcasting an IP datagram to all hosts in the zone. This is strictly not "required" by IP, but it is certainly "expected" by some applications.

MacIP hosts SHOULD never perform any routing. This includes anything to do with subnet masks, broadcast addresses, default gateways, RIP or ICMP redirects.

8.3.2. MacIP Modules - Introduction

MacIP can be best understood and described if the various functions are categorized into functional modules.

The following gives the relationship between the IP , MacIP and AppleTalk modules in the MacIP host:

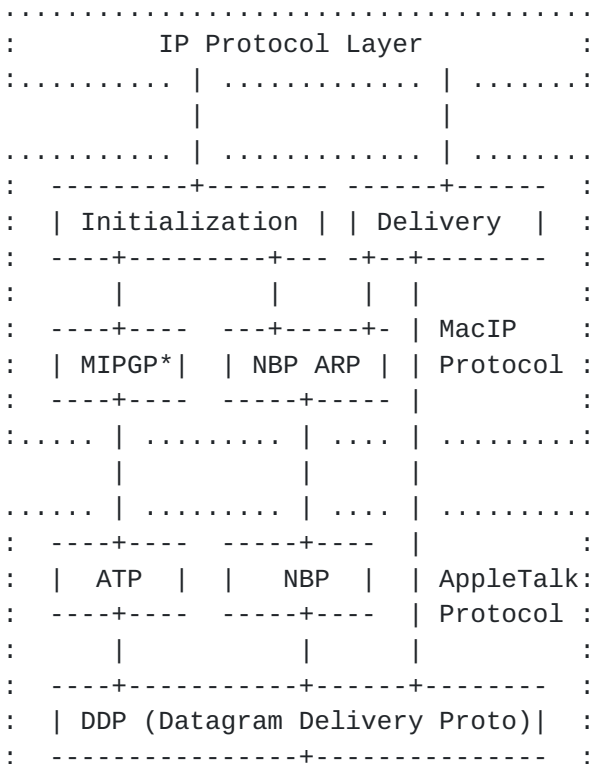
Evans and Ranch

November 11, 1992

Page 33

MacIP

November 1992







Figure 4. MacIP Host Implementation

The following gives the relationship between the IP , MacIP and AppleTalk modules in the MacIP gateway:



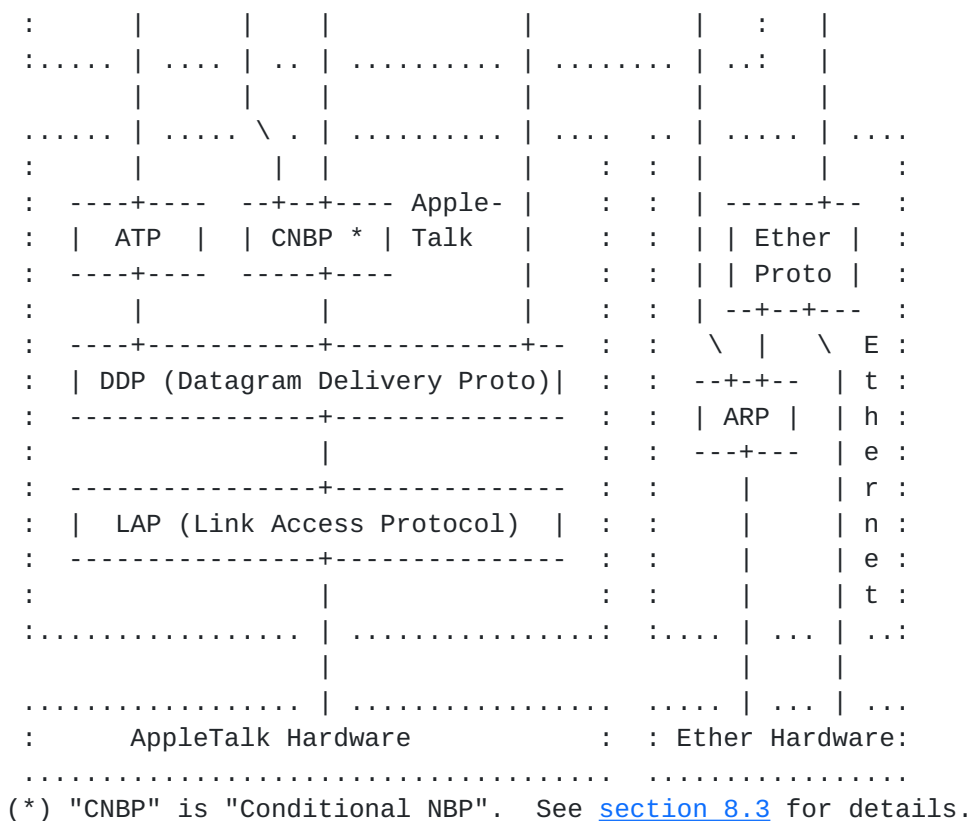


Figure 5. MacIP Gateway Implementation

[8.3.2.1.](#) Configuration Required For MacIP Hosts

Evans and Ranch

November 11, 1992

Page 35

MacIP

November 1992

The terms static and dynamic SHOULD be used consistently in the user interface.

If the host has multiple MacIP and/or IP interfaces, then a mechanism is required to allow selection between them. It is important to distinguish plainly between MacIP- based and "Native" IP connections.

[8.3.2.2.1.](#) Locating the MacIP gateway's Address Assignment Server

The implementation SHOULD not take any notice of the returned NVE name field in the NBP response. However, it should be noted that many implementations expect to find the IP address corresponding to the Server in dotted decimal format.

The implementation SHOULD record the full AppleTalk address returned

in the NBP response (including the returned socket) and use it in subsequent transactions with the Server. It SHOULD not simply assume that the Server is on DDP socket 72. However, there are many existing implementations that make this assumption.

8.3.2.4.2. Registration of Address Assignment Server

Unfortunately many MacIP host implementations wrongly rely on the name being an IP address, so for compatibility with these the use of the IP address as the name is "recommended". A simple two-port MacIP gateway configured in "forwarding" mode may only have one IP address to use for this purpose, but a multi-IP-port MacIP gateway may suffer an embarrassment of choices, so the question arises as to which IP address to use for this purpose. Fortunately this choice can be narrowed by the existence of bugs in some MacIP host implementations that can require the name of the Server to be an IP address in the same "subnet" as the IP address of the MacIP host. Good luck!

8.3.4.1. NBP ARP - Details

The aging of ARP cache entries is required. The mobility of Macintoshes makes this particularly important. Any algorithm may be used as long it is at least as good as the following.

Each use of the ARP cache entry should reset a "usage" timer. When a new entry is to be put into the cache, it might be full (or the bucket associated with the hashing function might be full). The entry with the oldest "usage" timer should be the one chosen to be replaced.

ARP cache entries should be "confirmed" by sending a single NBP Confirm directly to the AppleTalk address in the cache entry once a minute. If no response is received after five requests the entry

should be deleted.

The NBP Cache should be capable of being flushed on request by other modules.

8.3.5. Delivery

If the "Delivery" module experiences a hard error (the packet transport code cannot transmit the packet, and returns an error) when attempting to transmit a MacIP packet, then it is recommended that the NBP ARP cache entry for the destination IP address should be

flushed and the transmit retried. This is to recover from MacIP gateway restarts where the gateway has picked a different node address to the one it previously had.

8.3.6. MacIP Routing Decisions

The requirement in [section 3.6](#) that MacIP hosts should always use NBP ARP for all destination IP addresses is widely disobeyed in actual practice. Most MacIP host implementations disobey this fundamental specification, as the requirements of being an IP host were not well understood at the time.

It is thus expected that some MacIP-1 hosts will send some or all packets directly to the AppleTalk address of the discovered MacIP gateway for delivery. However it should be noted that if the MacIP gateway restarts, it may acquire a different AppleTalk node address to the one it had prior to the restart. If the host is sending MacIP packets to the old gateway address, they will not be delivered. The MacIP host SHOULD take measures to recover from this by following address cache aging and updating algorithms as discussed in the Host Requirements RFC [[10](#)]. This applies to NBP ARP caches as well.

8.3.7. Gateway NBP Proxy ARP

It is the implementation of this service that causes the most problems in MacIP. If the NBP Proxy ARP module wrongly responds to an IP addresses that is assigned to a MacIP host, then the response from the gateway will look to the host attempting to register its address as a duplicate registration.

This results in the restriction that with MacIP-1 there must never be multiple MacIP gateways in the one zone configured with different MacIP Ranges, as they will defeat registration attempts by MacIP hosts that have been assigned addresses by another gateway. Then again, two MacIP gateways can't be configured with the same MacIP range, as normal IP routing (Routing case) and Proxy ARP (Forwarding case) would fail to route packets under these circumstances. Therefore there can't be more than one MacIP gateway in any one zone.

This is the "Multiple Gateways in the One Zone" problems, and will be addressed later.

8.5.3. Name Binding Protocol

NBP can be thought of implementing a form of "zone-wide broadcast". The only thing that can be broadcast is a search for a named entity -

it is not possible to broadcast data-containing packets. This is important when considering mapping one network protocol (such as IP) onto an AppleTalk Internet.

It is not possible to discriminate a "registration attempt" NBP LookUp from a "searching for" one. This makes the implementation of NBP Proxy ARP far more difficult than you would first suspect.

9. Issues

9.1 IP Routing Issues

Phil pointed out that we didn't discuss IP hop count or checksum issues. Since we are subject to [rfc1122](#) (naturally), should we still discuss it? Furthermore, what do y'all think of bumping [or not bumping] the hop count if we are just in 'forwarding mode'?

Also, Phil asked if forwarding gateways need to intercept ICMP redirects. IP routers must ignore all ICMP redirects, as they're meant to have their routing tables updated by a "more reliable" method. Thus, all ICMP redirects MUST be ignored by the gateway, and MacIP hosts SHOULD ignore all ICMP redirects too. Should this go into this document? Where?

9.2 More Hacks

Jonathan described the following additional 'hacks'. Should they be in this document, and if so, where?

(1) The EtherGate has one ethernet port and 2 localtalk ports. We do client IP address assignment out of a single range of addresses for both localtalk ports. This means that besides it's usual weirdnesses, we had to make NBP-proxy-ARP work across both localtalk ports if they're in different zones. When a Mac tries to confirm that it can use its IPADDRESS the EtherGate forwards the NBP Lookup to the other port and changes the zone name. Yuck.

(2) When the EtherGate receives an NBP-ARP for an IPADDRESS not in its client range and on the same (sub)net as the EtherGate itself, it IP ARPs on the ethernet side to see if it can reach the IP host before responding on the NBP side.

Would someone with a good familiarity of Proxy ARP routers like to contribute - there may be some good practice in existing Proxy ARP routers that we can adopt, copy, steal documentation from etc.

One last note, I assume the implementation notes will include hints for the host (Mac client) implementor as well as the gateway implementor? We should clearly point out pitfalls and suggestions for both!

Any host-implementors like to contribute a few paragraphs?

9.3 MacIP Host Recommendations - Another Idea

From: Tom Evans

The MacIP _HOST_ doesn't need a "session" with anything.

Q. What is this "session" actually FOR?

A. It is established at MacIP Host startup for the sole purpose of getting a Dynamic IP Address assigned.

Q. Does the MacIP Host have to MAINTAIN the "session" with the Address Assignment Service?

A. NO.

Q. What is the ADDRESS of the IPGATEWAY?

A. I've been reading Radia's "Interconnections" too. By her definition in [section 2.3](#), an AppleTalk "address" doesn't fit her definition of an "address". The node-part of an AppleTalk address CHANGES.

So the "address" of the IPGATEWAY starts out as its NAME, which is initially "=:IPGATEWAY@*" ("any gateway"). The MacIP Host picks a particular one (which is now "aa.bb.cc.dd:IPGATEWAY@*"), which fits the definition of a NAME (but has an IP ADDRESS embedded in it). This is the only real "fixed" representation of the "address". The AppleTalk address MAY CHANGE, so it shouldn't be stored. If the IPGATEWAY is required for some reason, it should probably be searched for again at the time that it is needed.

The IP Address ("aa.bb.cc.dd" above) is a better "fixed address" than the AppleTalk address is. It can be resolved to an AppleTalk address when required by existing (NBP ARP) code.

Q. "My code treats IPGATEWAY as a "default IP gateway", so I have to

keep its address".

A. XXXXXXXXXXXX (removed by censors). Nothing else that I know of in a Mac stores AppleTalk addresses - they're too volatile. The LaserWriter driver stores the NAME of the printer. The AppleTalk address "A_Router" stored by any Mac is never more than 40 seconds old. ASP sessions store the AppleTalk address, but they're continuously maintained sessions that tell you when they break.

The logical thing to do is to store the MacIP Gateway's IP Address and NOT its AppleTalk Address. This will make a lot more sense in the host's existing IP Routing table after all. No "special case MacIP code" required in the IP layer.

Use NBP ARP to resolve ALL IP addresses passed down from the IP layer, including the one for the Default Gateway. The ARP/NBP-ARP code should have timeouts and confirmation. This will recover nicely if the IPGATEWAY changes its AppleTalk address, which it is "allowed" to do.

Q. How does the MacIP Gateway keep ITS session information (required for NBP Proxy ARP)?

A. By REREGISTRATION at gateway startup, by answering MacIPGP requests, by "tickling" them thereafter and by "probing" suspected "clients" that it receives NBP ARP requests from. It isn't the MacIP Host's problem.

So the MacIP Host shouldn't keep any information that it doesn't need to, and thus it shouldn't be in the MIB. If it is, then it should be the IP address that is stored and not the AppleTalk address.

Of course the above seems to be contradicted in the MacIP doc in sections [5.2.5](#) and [6.2](#) where it says that the MacIP Host should be able to receive multiple IPGATEWAY responses and then choose the "best" one. However, once the MacIP Host HAS an IP address, it doesn't need the Address Assignment Server anymore, and it should probably throw all the information away. So if the multiple gateway's addresses WERE accessible via SNMP, then they'd only be there for less than a second.

10. Acknowledgments

Bill Croft, for the initial implementation of this protocol in the SEAGATE gateway at Stanford University and for the documents provided with the KIP code.

Gaige Paulson and Tim K., for the NCSA Telnet source code.

MacIP

November 1992

John Romkey, for the original PC/IP source code.

Tim Maroney and ?, for the port of PC/IP to Macintosh.

Jeannine Smith for TCP and UDP in MacTCP, and John Veizades and his team for the rest.

Brad Parker and Josh Littlefield. This document is directly based on theirs written in February 1990.

John Veizades, for a version of this document and for being the Chair of the Apple- IP working group.

11. References

- [1] Sidhu, G., Andrews, R., and Oppenheimer, A., Apple Computer, "Inside AppleTalk, 2nd. Edition" Addison-Wesley Publishing Company, Inc., Reading, MA, May, 1990.
- [2] Postel J., "Internet Protocol", [RFC-791](#), USC Information Sciences Institute, September 1981.
- [3] Plummer, D., "An Ethernet Address Resolution Protocol", [RFC-826](#), Symbolics, September 1982.
- [4] Hornig, C., "Standard for the transmission of IP datagrams over Ethernet", [RFC-894](#), Symbolics, April 1984.
- [5] Mogul, J., "Internet Subnets", [RFC-917](#), Stanford University, October 1984.
- [6] Mogul, J., and Postel, J., "Internet Standard Subnetting Procedure", [RFC-950](#), Stanford University, August 1985.
- [7] Brandon, R., and Postel, J., "Requirements for Internet Gateways", [RFC-1009](#), USC Information Sciences Institute, June 1987.
- [8] Carl-Mitchell, S., Quarterman, J.S., "Using ARP to Implement Transparent Subnet Gateways", [RFC-1027](#), October 1987.
- [9] Hendrick, C., "Routing Information Protocol", [RFC-1058](#), June 1988.
- [10] Braden, R.T., ed, "Requirements for Internet Hosts", [RFC-1122](#), October 1989.

12. Expiration Date

This Internet Draft will expire in May 1993.

13. Security

This document does not discuss security in any manner.

14. Contact Points

The Apple-IP working group can be contacted by emailing the following address:

apple-ip@cayman.com

The Apple-IP Working Group Chairperson is:

John Veizades
Apple Computer
Cupertino, CA
(408)974-2672

EMail: veizades@apple.com

The author's addresses are:

Tom Evans
Webster Computer Corporation
1270 Ferntree Gully Rd
Scoresby Melbourne
3179 Victoria, Australia
61-3-764-1100

EMail: tom@wcc.oz.au

Christopher S. Ranch
Novell, Inc.
2180 Fortune Drive
San Jose, CA 95131
(408)473-8667

EMail: cranch@novell.com

