Simple Internet Protocol Plus White Paper

Status of this Memo

Abstract

   This document was submitted to the IETF IPng area in response to RFC
   1550. Publication of this document does not imply acceptance by the
   IPng area of any ideas expressed within. Comments should be
   submitted to the author and/or the sipp@sunroof.eng.sun.com mailing
   list.

1. Introduction

   This white paper presents an overview of the Simple Internet Protocol
   plus (SIPP) which is one of the candidates being considered in the
   Internet Engineering Task Force (IETF) for the next version of the
   Internet Protocol (the current version is usually referred to as
   IPv4). This white paper is not intended to be a detailed
   presentation of all of the features and motivation for SIPP, but is
   intended to give the reader an overview of the proposal. It is also
   not intended that this be an implementation specification, but given
   the simplicity of the central core of SIPP, an implementor familiar
   with IPv4 could probably construct a basic working SIPP
   implementation from reading this overview.

   SIPP is a new version of IP which is designed to be an evolutionary
   step from IPv4. It is a natural increment to IPv4. It can be
   installed as a normal software upgrade in internet devices and is
   interoperable with the current IPv4. Its deployment strategy was
   designed to not have any "flag" days. SIPP is designed to run well
   on high performance networks (e.g., ATM) and at the same time is
   still efficient for low bandwidth networks (e.g., wireless). In
   addition, it provides a platform for new internet functionality that
   will be required in the near future.

   This white paper describes the work of IETF SIPP working group.
   Several individuals deserve specific recognition. These include
   Steve Deering, Paul Francis, Dave Crocker, Bob Gilligan, Bill

Simpson, Ran Atkinson, Bill Fink, Erik Nordmark, Christian Huitema,
Sue Thompson, and Ramesh Govindan.

2. Key Issues for the Next Generation of IP

   There are several key issues that should be used in the evaluation of
   any next generation internet protocol. Some are very
   straightforward. For example the new protocol must be able to
   support large global internetworks. Others are less obvious. There
   must be a clear way to transition the current installed base of IP
   systems. It doesn't matter how good a new protocol is if there isn't
   a practical way to transition the current operational systems running
   IPv4 to the new protocol.

2.1 Growth

   Growth is the basic issue which caused there to be a need for a next
   generation IP. If anything is to be learned from our experience with
   IPv4 it is that the addressing and routing must be capable of
   handling reasonable scenarios of future growth. It is important that
   we have an understanding of the past growth and where the future
   growth will come from.

   Currently IPv4 serves what could be called the computer market. The
   computer market has been the driver of the growth of the Internet.
   It comprises the current Internet and countless other smaller
   internets which are not connected to the Internet. Its focus is to
   connect computers together in the large business, government, and
   university education markets. This market has been growing at an
   exponential rate. One measure of this is that the number of networks
   in current Internet (23,494 as of 1/28/94) is doubling approximately
   every 12 months. The computers which are used at the endpoints of
   internet communications range from PC's to Supercomputers. Most are
   attached to Local Area Networks (LANs) and the vast majority are not
   mobile.

   The next phase of growth will probably not be driven by the computer
   market. While the computer market will continue to grow at
   significant rates due to expansion into other areas such as schools
   (elementary through high school) and small businesses, it is doubtful
   it will continue to grow at an exponential rate. What is likely to
   happen is that other kinds of markets will develop. These markets
   will fall into several areas. They all have the characteristic that
   they are extremely large. They also bring with them a new set of
   requirements which were not as evident in the early stages of IPv4
   deployment. The new markets are also likely to happen in parallel
   with other. It may turn out that we will look back on the last ten
   years of Internet growth as the time when the Internet was small and

only doubling every year. The challenge for an IPng is to provide a
solution which solves todays problems and is attractive in these
emerging markets.

Nomadic personal computing devices seem certain to become ubiquitous
as their prices drop and their capabilities increase. A key
capability is that they will be networked. Unlike the majority of
todays networked computers they will support a variety of types of
network attachments. When disconnected they will use RF wireless
networks, when used in networked facilities they will use infrared
attachment, and when docked they will use physical wires. This makes
them an ideal candidate for internetworking technology as they will
need a common protocol which can work over a variety of physical
networks. These types of devices will become consumer devices and
will replace the current generation of cellular phones, pagers, and
personal digital assistants. In addition to the obvious requirement
of an internet protocol which can support large scale routing and
addressing, they will require an internet protocol which imposes a
low overhead and supports auto configuration and mobility as a basic
element. The nature of nomadic computing requires an internet
protocol to have built in authentication and confidentiality. It
also goes without saying that these devices will need to communicate
with the current generation of computers. The requirement for low
overhead comes from the wireless media. Unlike LAN's which will be
very high speed, the wireless media will be several orders of
magnitude slower due to constraints on available frequencies,
spectrum allocation, and power consumption.

Another market is networked entertainment. The first signs of this
emerging market are the proposals being discussed for 500 channels of
television, video on demand, etc. This is clearly a consumer market.
The possibility is that every television set will become an Internet
host. As the world of digital high definition television approaches,
the differences between a computer and a television will diminish.
As in the previous market, this market will require an Internet
protocol which supports large scale routing and addressing, and auto
configuration. This market also requires a protocol suite which
imposes the minimum overhead to get the job done. Cost will be the
major factor in the selection of a technology to use.

Another market which could use the next generation IP is device
control. This consists of the control of everyday devices such as
lighting equipment, heating and cooling equipment, motors, and other
types of equipment which are currently controlled via analog switches
and in aggregate consume considerable amounts of power. The size of
this market is enormous and requires solutions which are simple,
robust, easy to use, and very low cost.

The challenge for the IETF in the selection of an IPng is to pick a
protocol which meets today's requirements and also matches the
requirements of these emerging markets. These markets will happen
with or without an IETF IPng. If the IETF IPng is a good match for
these new markets it is likely to be used. If not, these markets
will develop something else. They will not wait for an IETF
solution. If this should happen it is probable that because of the
size and scale of the new markets the IETF protocol would be
supplanted. If the IETF IPng is not appropriate for use in these
markets, it is also probable that they will each develop their own
protocols, perhaps proprietary. These new protocols would not
interoperate with each other. The opportunity for the IETF is to
select an IPng which has a reasonable chance to be used in these
emerging markets. This would have the very desirable outcome of
creating an immense, interoperable, world-wide information
infrastructure created with open protocols. The alternative is a
world of disjoint networks with protocols controlled by individual
vendors.

2.2. Transition

At some point in the next three to seven years the Internet will
require a deployed new version of the Internet protocol. Two factors
are driving this: routing and addressing. Global internet routing
based on the on 32-bit addresses of IPv4 is becoming increasingly
strained. IPv4 address do not provide enough flexibility to
construct efficient hierarchies which can be aggregated. The
deployment of Classless Inter-Domain Routing [CIDR] is extending the
life time of IPv4 routing routing by a number of years, the effort to
manage the routing will continue to increase. Even if the IPv4
routing can be scaled to support a full IPv4 Internet, the Internet
will eventually run out of network numbers. There is no question
that an IPng is needed, but only a question of when.

The challenge for an IPng is for its transition to be complete before
IPv4 routing and addressing break. The transition will be much
easier if IPv4 address are still globally unique. The two transition
requirements which are the most important are flexibility of
deployment and the ability for IPv4 hosts to communicate with IPng
hosts. There will be IPng-only hosts, just as there will be IPv4-
only hosts. The capability must exist for IPng-only hosts to
communicate with IPv4-only hosts globally while IPv4 addresses are
globally unique.

The deployment strategy for an IPng must be as flexible as possible.
The Internet is too large for any kind of controlled rollout to be
successful. The importance of flexibility in an IPng and the need
for interoperability between IPv4 and IPng was well stated in a

message to the sipp mailing list by Bill Fink, who is responsible for
a portion of NASA's operational internet. In his message he said:

   "Being a network manager and thereby representing the interests of
   a significant number of users, from my perspective it's safe to
   say that the transition and interoperation aspects of any IPng is
   the key first element, without which any other significant
   advantages won't be able to be integrated into the user's network
   environment. I also don't think it wise to think of the
   transition as just a painful phase we'll have to endure en route
   to a pure IPng environment, since the transition/coexistence
   period undoubtedly will last at least a decade and may very well
   continue for the entire lifetime of IPng, until it's replaced with
   IPngng and a new transition. I might wish it was otherwise but I
   fear they are facts of life given the immense installed base.

   "Given this situation, and the reality that it won't be feasible
   to coordinate all the infrastructure changes even at the national
   and regional levels, it is imperative that the transition
   capabilities support the ability to deploy the IPng in the
   piecemeal fashion... with no requirement to need to coordinate
   local changes with other changes elsewhere in the Internet...

   "I realize that support for the transition and coexistence
   capabilities may be a major part of the IPng effort and may cause
   some headaches for the designers and developers, but I think it is
   a duty that can't be shirked and the necessary price that must be
   paid to provide as seamless an environment as possible to the end
   user and his basic network services such as e-mail, ftp, gopher,
   X-Window clients, etc...

   "The bottom line for me is that we must have interoperability
   during the extended transition period for the base IPv4
   functionality..."

Another way to think about the requirement for compatibility with
IPv4 is to look at other product areas. In the product world,
backwards compatability is very important. Vendors who do not
provide backward compatibility for their customers usually find they
do not have many customers left. For example, chip makers put
considerable effort into making sure that new versions of their
processor always run all of the software that ran on the previous
model. It is unlikely that Intel would develop a new processor in
the X86 family that did not run DOS and the tens of thousands of
applications which run on the current versions of X86's.

Operating system vendors go to great lengths to make sure new
versions of their operating systems are binary compatible with their

old version. For example the labels on most PC or MAC software
usually indicate that they require OS version XX or greater. It
would be foolish for Microsoft come out with a new version of Windows
which did not run the applications which ran on the previous version.
Microsoft even provides the ability for windows applications to run
on their new OS NT. This is an important feature. They understand
that it was very important to make sure that the applications which
run on Windows also run on NT.

The same requirement is also true for IPng. The Internet has a large
installed base. Features need to be designed into an IPng to make
the transition as easy as possible. As with processors and operating
systems, it must be backwards compatible with IPv4. Other protocols
have tried to replace TCP/IP, for example XTP and OSI. One element
in their failure to reach widespread acceptance was that neither had
any transition strategy other than running in parallel (sometimes
called dual stack). New features alone are not adequate to motivate
users to deploy new protocols. IPng must have a great transition
strategy and new features.

3. History of the SIPP Effort

   The SIPP working group represents the evolution of three different
   IETF working groups focused on developing an IPng. The first was
   called IP Address Encapsulation (IPAE) and was chaired by Dave
   Crocker and Robert Hinden. It proposed extensions to IPv4 which
   would carry larger addresses. Much of its work was focused on
   developing transition mechanisms.

   Somewhat later Steve Deering proposed a new protocol evolved from
   IPv4 called the Simple Internet Protocol (SIP). A working group was
   formed to work on this proposal which was chaired by Steve Deering
   and Christian Huitema. SIP had 64-bit addresses, a simplified
   header, and options in separate extension headers. After lengthly
   interaction between the two working groups and the realization that
   IPAE and SIP had a number of common elements and the transition
   mechanisms developed for IPAE would apply to SIP, the groups decided
   to merge and concentrate their efforts. The chairs of the new SIP
   working group were Steve Deering and Robert Hinden.

   In parallel to SIP, Paul Francis (formerly Paul Tsuchiya) had founded
   a working group to develop the "P" Internet Protocol (Pip). Pip was
   a new internet protocol based on a new architecture. The motivation
   behind Pip was that the opportunity for introducing a new internet
   protocol does not come very often and given that opportunity
   important new features should be introduced. Pip supported variable
   length addressing in 16-bit units, separation of addresses from
   identifiers, support for provider selection, mobility, and efficient

forwarding. It included a transition scheme similar to IPAE.

After considerable discussion among the leaders of the Pip and SIP
working groups, they came to realize that the advanced features in
Pip could be accomplished in SIP without changing the base SIP
protocol as well as keeping the IPAE transition mechanisms. In
essence it was possible to keep the best features of each protocol.
Based on this the groups decided to merge their efforts. The new
protocol was called Simple Internet Protocol Plus (SIPP). The chairs
of the merged working group are Steve Deering, Paul Francis, and
Robert Hinden.

4. SIPP Overview

SIPP is a new version of the Internet Protocol, designed as a
successor to IP version 4 [IPV4]. SIPP is assigned IP version number
6.

SIPP was designed to take an evolutionary step from IPv4. It was not
a design goal to take a radical step away from IPv4. Functions which
work in IPv4 were kept in SIPP. Functions which didn't work were
removed. The changes from IPv4 to SIPP fall primarily into the
following categories:

   o Expanded Routing and Addressing Capabilities

      SIPP increases the IP address size from 32 bits to 64 bits, to
      support more levels of addressing hierarchy and a much greater
      number of addressable nodes. SIPP addressing can be further
      extended, in units of 64 bits, by a facility equivalent to
      IPv4's Loose Source and Record Route option, in combination
      with a new address type called "cluster addresses" which
      identify topological regions rather than individual nodes.
      The scaleability of multicast routing is improved by adding
      a "scope" field to multicast addresses.

   o Header Format Simplification

      Some IPv4 header fields have been dropped or made optional, to
      reduce the common-case processing cost of packet handling and to
      keep the bandwidth cost of the SIPP header almost as low as that
      of IPv4, despite the increased size of the addresses. The basic
      SIPP header is only four bytes longer than IPv4.

o Improved Support for Options

Changes in the way IP header options are encoded allows for more
efficient forwarding, less stringent limits on the length of
options, and greater flexibility for introducing new options in
the future.

o Quality-of-Service Capabilities

A new capability is added to enable the labeling of packets
belonging to particular traffic "flows" for which the sender
requests special handling, such as non-default quality of
service or "real-time" service.

o Authentication and Privacy Capabilities

SIPP includes the definition of extensions which provide support
for authentication, data integrity, and confidentiality. This
is included as a basic element of SIPP.

The SIPP protocol consists of two parts, the basic SIPP header and
SIPP Options.

4.1 SIPP Header Format

```
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
|Version|                    Flow Label                         |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
|        Payload Length         | Payload Type  |   Hop Limit   |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
|                                                               |
+                                                               +
|                                                               |
+                        Source Address                         +
|                                                               |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
|                                                               |
+                                                               +
|                     Destination Address                       |
+                                                               +
|                                                               |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
```

Version   4-bit Internet Protocol version number = 6.

Flow Label   28-bit field. See SIPP Quality of Service
                     section.

Payload Length   16-bit unsigned integer. Length of payload,
                     i.e., the rest of the packet following the
                     SIPP header, in octets.

        Payload Type   8-bit selector. Identifies the type of
                               header immediately following the SIPP
                               header. Uses the same values as the IPv4
                               Protocol field [STD 2, RFC 1700].

        Hop Limit   8-bit unsigned integer. Decremented by 1
                               by each node that forwards the packet.
                               The packet is discarded if Hop Limit is
                               decremented to zero.

        Source Address   64 bits. An address of the initial sender of
                               the packet. See [ROUT] for details.

        Destination Address   64 bits. An address of the intended
                               recipient of the packet (possibly not the
                               ultimate recipient, if an optional Routing
                               Header is present).

4.2 SIPP Options

   SIPP includes an improved option mechanism over IPv4. SIPP options
   are placed in separate headers that are located between the SIPP
   header and the transport-layer header in a packet. Most SIPP option
   headers are not examined or processed by any router along a packet's
   delivery path until it arrives at its final destination. This
   facilitates a major improvement in router performance for packets
   containing options. In IPv4 the presence of any options requires the
   router to examine all options. The other improvement is that unlike
   IPv4, SIPP options can be of arbitrary length and the total amount of
   options carried in a packet is not limited to 40 bytes. This feature
   plus the manner in which they are processed, permits SIPP options to
   be used for functions which were not practical in IPv4. A good
   example of this is the SIPP Authentication and Security Encapsulation
   options.

   In order to improve the performance when handling subsequent option
   headers and the transport protocol which follows, SIPP options are
   always an integer multiple of 8 octets long, in order to retain this
   alignment for subsequent headers.

The SIPP option headers which are currently defined are:

```
   Option                        Function
   ---------------               ----------------------------------------
   Routing                       Extended Routing (like IPv4 loose source
                                 route)
   Fragmentation                 Fragmentation and Reassembly
   Authentication                Integrity and Authentication
   Security Encapsulation        Confidentiality
   Hop-by-Hop Option             Special options which require hop by hop
                                 processing
```

4.3 SIPP Addressing

   SIPP addresses are 64-bits long and are identifiers for individual
   nodes and sets of nodes. There are three types of SIPP addresses.
   These are unicast, cluster, and multicast. Unicast addresses
   identify a single node. Cluster addresses identify a group of nodes,
   that share a common address prefix, such that a packet sent to a
   cluster address will be delivered to one member of the group.
   Multicast addresses identify a group of nodes, such that a packet
   sent to a multicast address is delivered to all of the nodes in the
   group.

   SIPP supports addresses which are twice the number of bits as IPv4
   addresses. These addresses support an address space which is four
   billion (2^{32}) times the size of IPv4 addresses (2^{32}). Another
   way to say this is that SIPP supports four billion internets each the
   size of the maximum IPv4 internet. That is enough to allow each
   person on the planet to have their own internet. Even with several
   layers of hierarchy (with assignment utilization similar to IPv4)
   this would allow for each person on the planet to have their own
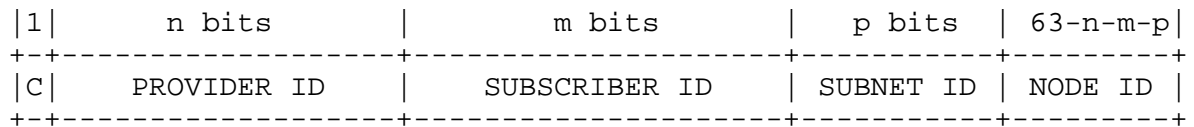   internet each holding several thousand hosts.

   In addition, SIPP supports extended addresses using the routing
   option. This capability allows the address space to grow to 128-
   bits, 192-bits (or even larger) while still keeping the address units
   in manageable 64-bit units. This permits the addresses to grow while
   keeping the routing algorithms efficient because they continue to
   operate using 64- bit units.

4.3.1 Unicast Addresses

   There are several forms of unicast address assignment in SIPP. These
   are global hierarchical unicast addresses, local-use addresses, and
   IPv4- only host addresses. The assignment plan for unicast addresses
   is described in [ADDR].

4.3.1.1 Global Unicast Addresses

   Global unicast addresses are used for global communication. They are
   the most common SIPP address and are similar in function to IPv4
   addresses. Their format is:

```
   |1|      n bits       |        m bits       |  p bits  | 63-n-m-p|
   +-+-------------------+---------------------+----------+---------+
   |C|     PROVIDER ID   |     SUBSCRIBER ID   | SUBNET ID | NODE ID |
   +-+-------------------+---------------------+----------+---------+
```

   The first bit is the IPv4 compatibility bit, or C-bit. It indicates
   whether the node represented by the address is IPv4 or SIPP. SIPP
   addresses are provider-oriented. That is, the high-order part of the
   address is assigned to internet service providers, which then assign
   portions of the address space to subscribers, etc. This usage is
   similar to assignment of IP addresses under CIDR. The SUBSCRIBER ID
   distinguishes among multiple subscribers attached to the provider
   identified by the PROVIDER ID. The SUBNET ID identifies a
   topologically connected group of nodes within the subscriber network
   identified by the subscriber prefix. The NODE ID identifies a single
   node among the group of nodes identified by the subnet prefix.

4.3.1.2 Local-Use Address

   A local-use address is a unicast address that has only local
   routability scope (within the subnet or within a subscriber network),
   and may have local or global uniqueness scope. They are intended for
   use inside of a site for "plug and play" local communication, for
   bootstrapping up to a single global addresses, and as part of an
   address sequence for global communication. Their format is:

```
   | 4  |
   |bits|    12 bits    |                   48 bits                  |
   +----+---------------+--------------------------------------------+
   |0110|    SUBNET ID  |                   NODE ID                  |
   +----+---------------+--------------------------------------------+
```

   The NODE ID is an identifier which much be unique in the domain in
   which it is being used. In most cases these will use a node's IEEE-
   802 48bit address. The SUBNET ID identifies a specific subnet in a
   site. The combination of the SUBNET ID and the NODE ID to form a
   local use address allows a large private internet to be constructed
   without any other address allocation.

   Local-use addresses have two primary benefits. First, for sites or
   organizations that are not (yet) connected to the global Internet,
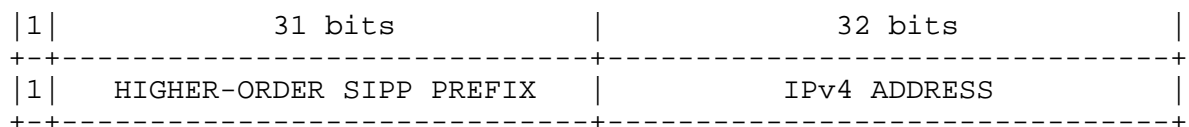   there is no need to request an address prefix from the global

Internet address space. Local-use addresses can be used instead. If
the organization connects to the global Internet, it can use it's
local use addresses to communicate with a server (e.g., using the
Dynamic Host Configuration Protocol [DHCP]) to have a global address
automatically assigned.

The second benefit of local-use addresses is that they can hold much
larger NODE IDs, which makes possible a very simple form of auto-
configuration of addresses. In particular, a node may discover a
SUBNET ID by listening to a Router Advertisement messages on its
attached link(s), and then fabricating a SIPP address for itself by
using its link-level address as the NODE ID on that subnet.

An auto-configured local-use address may be used by a node as its own
identification for communication within the local domain, possibly
including communication with a local address server to obtain a
global SIPP address. The details of host auto-configuration are
described in [DHCP].

4.3.1.3 IPv4-Only Addresses

SIPP unicast addresses are assigned to IPv4-only hosts as part of the
IPAE scheme for transition from IPv4 to SIPP. Such addresses have
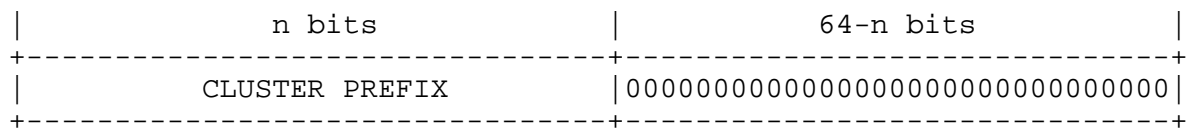the following form:

```
 |1|          31 bits           |           32 bits            |
 +-+----------------------------+------------------------------+
 |1|   HIGHER-ORDER SIPP PREFIX  |         IPv4 ADDRESS         |
 +-+----------------------------+------------------------------+
```

The highest-order bit of a SIPP address is called the IPv4
compatibility bit or the C bit. A C bit value of 1 identifies an
address as belonging to an IPv4-only node.

The IPv4 node's 32-bit IPv4 address is carried in the low-order 32
bits of the SIPP address. The remaining 31 bits are used to carry
HIGHER- ORDER SIPP PREFIX, such as a service-provider ID.
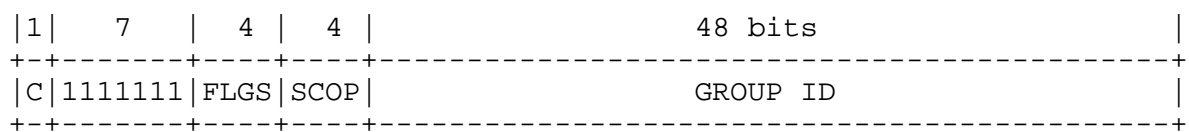
4.3.2 Cluster Addresses

Cluster addresses are unicast addresses that are used to reach the
"nearest" one (according to unicast routing's notion of nearest) of
the set of boundary routers of a cluster of nodes identified by a
common prefix in the SIPP unicast routing hierarchy. These are used
to identify a set of nodes. The cluster address, when used as part
of an address sequence, permits a node to select which of several
providers it wants to carry its traffic. A cluster address can only
be used as a destination address. In this example there would be a

cluster address for each provider. This capability is sometimes
called "source selected policies". Cluster addresses have the
general form:

```
|              n bits          |          64-n bits           |
 +------------------------------+------------------------------+
|         CLUSTER PREFIX       |00000000000000000000000000000000|
 +------------------------------+------------------------------+
```

4.3.3 Multicast Addresses

A SIPP multicast address is an identifier for a group of nodes. A
node may belong to any number of multicast groups. Multicast
addresses have the following format:

```
|1|   7   | 4 | 4 |                48 bits                     |
+-+-------+----+----+------------------------------------------+
|C|1111111|FLGS|SCOP|                GROUP ID                   |
+-+-------+----+----+------------------------------------------+
```

Where:

C = IPv4 compatibility bit.

1111111 in the rest of the first octet identifies the address as
being a multicast address.

```
                           +-+-+-+-+
FLGS is a set of 4 flags:  |0|0|0|T|
                           +-+-+-+-+
```

The high-order 3 flags are reserved, and must be initialized to 0.

T = 0 indicates a permanently-assigned ("well-known") multicast
        address, assigned by the global internet numbering authority.

T = 1 indicates a non-permanently-assigned ("transient") multicast
        address.

SCOP is a 4-bit multicast scope value used to limit the scope of
the multicast group. The values are:

```
    0   reserved                 8   intra-organization scope
    1   intra-node scope         9   (unassigned)
    2   intra-link scope         10  (unassigned)
    3   (unassigned)             11  intra-community scope
    4   (unassigned)             12  (unassigned)
```

            5   intra-site scope   13   (unassigned)
            6   (unassigned)   14   global scope
            7   (unassigned)   15   reserved

     GROUP ID identifies the multicast group, either permanent or
     transient, within the given scope.

4.4 SIPP Routing

     Routing in SIPP is almost identical to IPv4 routing under CIDR except
     that the addresses are 64-bit SIPP addresses instead of 32-bit IPv4
     addresses. This is true even when extended addresses are being used.
     With very straightforward extensions, all of IPv4's routing
     algorithms (OSPF, BGP, RIP, IDRP, etc.) can used to route SIPP [OSPF]
     [RIP2] [IDRP].

     SIPP also includes simple routing extensions which support powerful
     new routing functionality. These capabilities include:

          Provider Selection (based on policy, performance, cost, etc.)
          Host Mobility (route to current location)
          Auto-Readdressing (route to new address)
          Extended Addressing (route to "sub-cloud")

     The new routing functionality is obtained by creating sequences of
     SIPP addresses using the SIPP Routing option. The routing option is
     used by a SIPP source to list one or more intermediate nodes (or
     topological clusters) to be "visited" on the way to a packet's
     destination. This function is very similar in function to IPv4's
     Loose Source and Record Route option. A node would publish its
     address sequence in the Domain Name System [DNS].

     The identification of a specific transport connection is done by only
     using the first (source) and last (destination) address in the
     sequence. These identifying addresses (i.e., first and last
     addresses of a route sequence) are required to be unique within the
     scope over which they are used. This permits the middle addresses in
     the address sequence to change (in the cases of mobility, provider
     changes, site readdressing, etc.) without disrupting the transport
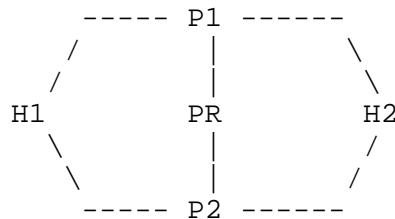   connection.

     In order to make address sequences a general function, SIPP hosts are
     required to reverse routes in a packet it receives containing address
     sequences in order to return the packet to its originator. This
     approach is taken to make SIPP host implementations from the start
     support the handling and reversal of source routes. This is the key
     for allowing them to work with hosts which implement the new features
     such as provider selection or extended addresses.

Three examples show how the extended addressing can be used. In
these examples, address sequences are shown by a list of individual
addresses separated by commas. For example:

     SRC, I1, I2, I3, DST

Where the first address is the source address, the last address is
the destination address, and the middle addresses are intermediate
addresses.

For these examples assume that two hosts, H1 and H2 wish to
communicate. Assume that H1 and H2's sites are both connected to
providers P1 and P2. A third wireless provider, PR, is connected to
both providers P1 and P2.

```
                    ----- P1 ------
                   /       |       \
                  /        |        \
               H1         PR         H2
                  \        |        /
                   \       |       /
                    ----- P2 ------
```

The simplest case (no use of address sequences) is when H1 wants to
send a packet to H2 containing the addresses:

        H1, H2

When H2 replied it would reverse the addresses and construct a packet
containing the addresses:

        H2, H1

In this example either provider could be used, and H1 and H2 would
not be able to select which provider traffic would be sent to and
received from.

If H1 decides that it wants to enforce a policy that all
communication to/from H2 can only use provider P1, it would construct
a packet containing the address sequence:

        H1, P1, H2

This ensures that when H2 replies to H1, it will reverse the route
and the reply it would also travel over P1. The addresses in H2's
reply would look like:

        H2, P1, H1

If H1 became mobile and moved to provider PR, it could maintain (not
breaking any transport connections) communication with H2, by sending
packets that contain the address sequence:

        H1, PR, P1, H2

This would ensure that when H2 replied it would enforce H1's policy
of exclusive use of provider P1 and send the packet to H1 new
location on provider PR. The reversed address sequence would be:

        H2, P1, PR, H1

The address extension facility of SIPP can be used for provider
selection, mobility, readdressing, and extended addressing. It is a
simple but powerful capability.

4.5 SIPP Quality-of-Service Capabilities

The Flow Label field in the SIPP header may be used by a host to
label those packets for which it requests special handling by SIPP
routers, such as non-default quality of service or "real-time"
service. This labeling is important in order to support applications
which require some degree of consistent throughput, delay, and/or
jitter. The Flow Label is a 28-bit field, internally structured into
three subfields as follows:

```
   +-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
   |R|  DP |                   Flow ID                     |
   +-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
```

   R (Reserved)   1-bit subfield. Initialized to zero for
                       transmission; Ignored on reception.

   DP (Drop Priority) 3-bit unsigned integer. Specifies the
                       priority of the packet, relative to other
                       packets from the same source, for being
                       discarded by a router under conditions of
                       congestion. Larger values indicates a
                       greater willingness by the sender to allow
                       the packet to be discarded.

   Flow ID   24-bit subfield used to identify a
                       specific flow.

A flow is a sequence of packets sent from a particular source to a
particular (unicast or multicast) destination for which the source
desires special handling by the intervening routers. There may be
multiple active flows from a source to a destination, as well as

traffic that is not associated with any flow. A flow is identified
by the combination of a Source Address and a non-zero Flow ID.
Packets that do not belong to a flow carry a Flow ID of zero.

A Flow ID is assigned to a flow by the flow's source node. New Flow
IDs must be chosen (pseudo-)randomly and uniformly from the range 1
to FFFFFF hex. The purpose of the random allocation is to make any
set of bits within the Flow ID suitable for use as a hash key by the
routers, for looking up the special-handling state associated with
the flow. A Flow ID must not be re-used by a source for a new flow
while any state associated with the previous usage still exists in
any router.

The Drop Priority subfield provides a means separate from the Flow ID
for distinguishing among packets from the same source, to allow a
source to specify which of its packets are to be discarded in
preference to others when a router cannot forward them all. This is
useful for applications like video where it is preferable to drop
packets carrying screen updates rather than the packets carrying the
video synchronization information.

4.6 SIPP Security

The current Internet has a number of security problems and lacks
effective privacy and authentication mechanisms below the application
layer. SIPP remedies these shortcomings by having two integrated
options that provide security services. These two options may be
used singly or together to provide differing levels of security to
different users. This is very important because different user
communities have different security needs.

The first mechanism, called the "SIPP Authentication Header", is an
option which provides authentication and integrity (without
confidentiality) to SIPP datagrams. While the option is algorithm-
independent and will support many different authentication
techniques, the use of keyed MD5 is proposed to help ensure
interoperability within the worldwide Internet. This can be used to
eliminate a significant class of network attacks, including host
masquerading attacks. The use of the SIPP Authentication Header is
particularly important when source routing is used with SIPP because
of the known risks in IP source routing. Its placement at the
internet layer can help provide host origin authentication to those
upper layer protocols and services that currently lack meaningful
protections. This mechanism should be exportable by vendors in the
United States and other countries with similar export restrictions
because it only provides authentication and integrity, and
specifically does not provide confidentiality. The exportability of
the SIPP Authentication Header encourages its widespread

implementation and use.

The second security option provided with SIPP is the "SIPP
Encapsulating Security Header". This mechanism provides integrity
and confidentiality to SIPP datagrams. It is simpler than some
similar security protocols (e.g., SP3D, ISO NLSP) but remains
flexible and algorithm-independent. To achieve interoperability
within the global Internet, the use of DES CBC is proposed as the
standard algorithm for use with the SIPP Encapsulating Security
Header.

5. SIPP Transition Mechanisms

The two key motivations in the SIPP transition mechanisms are to
provide direct interoperability between IPv4 and SIPP hosts and to
allow the user population to adopt SIPP in an a highly diffuse
fashion. The transition must be incremental, with few or no critical
interdependencies, if it is to succeed. The SIPP transition allows
the users to upgrade their hosts to SIPP, and the network operators
to deploy SIPP in routers, with very little coordination between the
two.

The mechanisms and policies of the SIPP transition are called "IPAE".
Having a separate term serves to highlight those features designed
specifically for transition. Once an acronym for an encapsulation
technique to facilitate transition, the term "IPAE" now is mostly
historical.

The IPAE transition is based on five key elements:

  1) A 64-bit SIPP addressing plan that encompasses the existing
     32-bit IPv4 addressing plan. The 64-bit plan will be used to
     assign addresses for both SIPP and IPv4 nodes at the beginning
     of the transition. Existing IPv4 nodes will not need to change
     their addresses, and IPv4 hosts being upgraded to SIPP keep their
     existing IPv4 addresses as the low-order 32 bits of their SIPP
     addresses. Since the SIPP addressing plan is a superset of the
     existing IPv4 plan, SIPP hosts are assigned only a single 64-bit
     address, which can be used to communicate with both SIPP and IPv4
     hosts.

  2) A mechanism for encapsulating SIPP traffic within IPv4 packets so
     that the IPv4 infrastructure can be leveraged early in the
     transition. Most of the "SIPP within IPv4 tunnels" can be
     automatically configured.

   3) Algorithms in SIPP hosts that allow them to directly interoperate
      with IPv4 hosts located on the same subnet and elsewhere in the
      Internet.

   4) A mechanism for translating between IPv4 and SIPP headers to
      allow SIPP-only hosts to communicate with IPv4-only hosts and to
      facilitate IPv4 hosts communicating over over a SIPP-only
      backbone.

   5) An optional mechanism for mapping IPv4 addresses to SIPP address
      to allow improved scaling of IPv4 routing. At the present time
      given the success of CIDR, this does not look like it will be
      needed in a transition to SIPP. If Internet growth should
      continue beyond what CIDR can handle, it is available as an
      optional mechanism.

   IPAE ensures that SIPP hosts can interoperate with IPv4 hosts
   anywhere in the Internet up until the time when IPv4 addresses run
   out, and afterward allows SIPP and IPv4 hosts within a limited scope
   to interoperate indefinitely. This feature protects for a very long
   time the huge investment users have made in IPv4. Hosts that need
   only a limited connectivity range (e.g., printers) need never be
   upgraded to SIPP. This feature also allows SIPP-only hosts to
   interoperate with IPv4-only hosts.

   The incremental upgrade features of IPAE allow the host and router
   vendors to integrate SIPP into their product lines at their own pace,
   and allows the end users and network operators to deploy SIPP on
   their own schedules.

   The interoperability between SIPP and IPv4 provided by IPAE also has
   the benefit of extending the lifetime of IPv4 hosts. Given the large
   installed base of IPv4, changes to IPv4 in hosts are nearly
   impossible. Once an IPng is chosen, most of the new feature
   development will be done on IPng. New features in IPng will increase
   the incentives to adopt and deploy it.

6. Why SIPP?

   There are a number of reasons why SIPP should be selected as the
   IETF's IPng. It solves the Internet scaling problem, provides a
   flexible transition mechanism for the current Internet, and was
   designed to meet the needs of new markets such as nomadic personal
   computing devices, networked entertainment, and device control. It
   does this in a evolutionary way which reduces the risk of
   architectural problems.

Ease of transition is a key point in the design of SIPP. It is not
something was was added in at the end. SIPP is designed to
interoperate with IPv4. Specific mechanisms (C-bit, embedded IPv4
addresses, etc.) were built into SIPP to support transition and
compatability with IPv4. It was designed to permit a gradual and
piecemeal deployment without any dependencies.

SIPP supports large hierarchical addresses which will allow the
Internet to continue to grow and provide new routing capabilities not
built into IPv4. It has cluster addresses which can be used for
policy route selection and has scoped multicast addresses which
provide improved scaleability over IPv4 multicast. It also has local
use addresses which provide the ability for "plug and play"
installation.

SIPP is designed to have performance better than IPv4 and work well
in low bandwidth applications like wireless. Its headers are less
expensive to process than IPv4 and its 64-bit addresses are chosen to
be well matched to the new generation of 64bit processors. Its
compact header minimizes bandwidth overhead which makes it ideal for
wireless use.

SIPP provides a platform for new Internet functionality. This
includes support for real-time flows, provider selection, host
mobility, end-to- end security, auto-configuration, and auto-
reconfiguration.

In summary, SIPP is a new version of IP. It can be installed as a
normal software upgrade in internet devices. It is interoperable
with the current IPv4. Its deployment strategy was designed to not
have any "flag" days. SIPP is designed to run well on high
performance networks (e.g., ATM) and at the same time is still
efficient for low bandwidth networks (e.g., wireless). In addition,
it provides a platform for new internet functionality that will be
required in the near future.

7. Status of SIPP Effort

   There are many active participants in the SIPP working group. Groups
   making active contributions include:

   Group   Activity
   ---------------------   ---
   Beame & Whiteside   Implementation (PC)
   Bellcore   Implementation (SunOS), DNS and ICMP specs.
   Digital Equipment Corp. Implementation (Alpha/OSF, Open VMS)
   INRIA   Implementation (BSD, BIND), DNS & OSPF specs.
   INESC   Implementation (BSD/Mach/x-kernel)
   Intercon   Implementation (MAC)
   MCI   Phone Conferences
   Merit   IDRP for SIPP Specification
   Naval Research Lab.   Implementation (BSD) Security Design
   Network General   Implementation (Sniffer)
   SGI   Implementation (IRIX, NetVisulizer)
   Sun   Implementation (Solaris 2.x, Snoop)
   TGV   Implementation (Open VMS)
   Xerox PARC   Protocol Design
   Bill Simpson   Implementation (KA9Q)

   As of the time this paper was written there were a number of SIPP and
   IPAE implementations. These include:

   Implementation   Status
   --------------   -------------------------------------
   BSD/Mach   Completed (telnet, NFS, AFS, UDP)
   BSD/Net/2   In Progress
   Bind   Code done
   DOS &Windows   Completed (telnet, ftp, tftp, ping)
   IRIX   In progress (ping)
   KA9Q   In progress (ping, TCP)
   Mac OS   Completed (telnet, ftp, finger, ping)
   NetVisualizer   Completed (SIP & IPAE)
   Open VMS   Completed (telnet, ftp), In Progress
   OSF/1   In Progress (ping, ICMP)
   Sniffer   Completed (SIP & IPAE)
   Snoop   Completed (SIP & IPAE)
   Solaris   Completed (telnet, ftp, tftp, ping)
   Sun OS   In Progress

8. Where to Get Additional Information

   The documentation listed in the reference sections can be found in
   one of the IETF internet draft directories or in the archive site for
   the SIPP working group. This is located at:

           ftp.parc.xerox.com   in the /pub/sipp   directory.

   In addition other material relating to SIPP (such as postscript
   versions of presentations on SIPP) can also be found in the SIPP
   working group archive.

   To join the SIPP working group, send electronic mail to

           sipp-request@sunroof.eng.sun.com

   An archive of mail sent to this mailing list can be found in the IETF
   directories at cnri.reston.va.us.

9. Security Considerations

   Security issues are discussed in section 4.6.

10. Author's Address

   Robert M. Hinden
   Manager, Internet Engineering
   Sun Microsystems, Inc.
   MS MTV5-44
   2550 Garcia Ave.
   Mt. View, CA 94303

   Phone: (415) 336-2082
   Fax: (415) 336-6016
   EMail: hinden@eng.sun.com

11. References

   [ADDR]   Francis, P., "Simple Internet Protocol Plus (SIPP): Unicast
            Hierarchical Address Assignment", Work in Progress, January
            1994.

   [AUTH]   Atkinson, R., "SIPP Authentication Payload",
            Work in Progress, January, 1994.

   [CIDR]   Fuller, V., Li, T., Yu, J., and K. Varadhan, "Supernetting:
            an Address Assignment and Aggregation Strategy", RFC 1338,
            BARRNet, cisco, Merit, OARnet, June 1992.

   [DISC]   Simpson, W., "SIPP Neighbor Discovery", Work in Progress,
             March 1994.

   [DIS2]   Simpson, W., "SIPP Neighbor Discovery -- ICMP Message
             Formats", Work in Progress, March 1994.

   [DHCP]   Thomson, S., "Simple Internet Protocol Plus (SIPP): Automatic
             Host Address Assignment", Work in Progress, March 1994.

   [DNS]   Thomson, S., and C. Huitema, "DNS Extensions to Support
             Simple Internet Protocol Plus (SIPP)", Work in Progress,
             March 1994.

   [ICMP]   Govindan, R., and S. Deering, "ICMP and IGMP for the Simple
             Internet Protocol Plus (SIPP)", Work in Progress, March 1994.

   [IDRP]   Hares, S., "IDRP for SIP", Work in Progress, November 1993.

   [IPAE]   Gilligan, R., et al, "IPAE: The SIPP Interoperability and
             Transition Mechanism", Work in Progress, March 1994.

   [IPV4]   Postel, J., "Internet Protocol- DARPA Internet Program
             Protocol Specification", STD 5, RFC 791, DARPA,
             September 1981.

   [OSPF]   Francis, P., "OSPF for SIPP", Work in Progress, February
             1994.

   [RIP2]   Malkin, G., and C. Huitema, "SIP-RIP", Work in Progress,
             March 1993.

   [ROUT]   Deering, S., et al, "Simple Internet Protocol Plus (SIPP):
             Routing and Addressing", Work in Progress, February 1994.

   [SARC]   Atkinson, R., "SIPP Security Architecture", Work in Progress,
             January 1994.

   [SECR]   Atkinson, R., "SIPP Encapsulating Security Payload (ESP)",
             Work in Progress, January 1994.

   [SIPP]   Deering, S., "Simple Internet Protocol Plus (SIPP)
             Specification", Work in Progress, February 1994.