            The Recommendation for the IP Next Generation Protocol

Status of this Memo

Abstract

   This document presents the recommendation of the IPng Area Directors
   on what should be used to replace the current version of the Internet
   Protocol. This recommendation was accepted by the Internet
   Engineering Steering Group (IESG).

Table of Contents

1. Summary

   The IETF started its effort to select a successor to IPv4 in late
   1990 when projections indicated that the Internet address space would
   become an increasingly limiting resource. Several parallel efforts
   then started exploring ways to resolve these address limitations
   while at the same time providing additional functionality. The IETF
   formed the IPng Area in late 1993 to investigate the various
   proposals and recommend how to proceed. We developed an IPng
   technical criteria document and evaluated the various proposals
   against it. All were found wanting to some degree. After this
   evaluation, a revised proposal was offered by one of the working

groups that resolved many of the problems in the previous proposals.
The IPng Area Directors recommend that the IETF designate this
revised proposal as the IPng and focus its energy on bringing a set
of documents defining the IPng to Proposed Standard status with all
deliberate speed.

This protocol recommendation includes a simplified header with a
hierarchical address structure that permits rigorous route
aggregation and is also large enough to meet the needs of the
Internet for the foreseeable future. The protocol also includes
packet-level authentication and encryption along with plug and play
autoconfiguration. The design changes the way IP header options are
encoded to increase the flexibility of introducing new options in the
future while improving performance. It also includes the ability to
label traffic flows.

Specific recommendations include:

* current address assignment policies are adequate
* there is no current need to reclaim underutilized assigned network
  numbers
* there is no current need to renumber major portions of the Internet
* CIDR-style assignments of parts of unassigned Class A address space
  should be considered
* "Simple Internet Protocol Plus (SIPP) Spec. (128 bit ver)"
  [Deering94b] be adopted as the basis for IPng
* the documents listed in Appendix C be the foundation of the IPng
  effort
* an IPng Working Group be formed, chaired by Steve Deering and Ross
  Callon
* Robert Hinden be the document editor for the IPng effort
* an IPng Reviewer be appointed and that Dave Clark be the reviewer
* an Address Autoconfiguration Working Group be formed, chaired by
  Dave Katz and Sue Thomson
* an IPng Transition Working Group be formed, chaired by Bob Gilligan
  and TBA
* the Transition and Coexistence Including Testing Working Group be
  chartered
* recommendations about the use of non-IPv6 addresses in IPv6
  environments and IPv6 addresses in non-IPv6 environments be
  developed
* the IESG commission a review of all IETF standards documents for
  IPng implications
* the IESG task current IETF working groups to take IPng into account
* the IESG charter new working groups where needed to revise old
  standards documents
* Informational RFCs be solicited or developed describing a few
  specific IPng APIs

   * the IPng Area and Area Directorate continue until main documents
     are offered as Proposed Standards in late 1994
   * support for the Authentication Header be required
   * support for a specific authentication algorithm be required
   * support for the Privacy Header be required
   * support for a specific privacy algorithm be required
   * an "IPng framework for firewalls" be developed

2. Background

   Even the most farseeing of the developers of TCP/IP in the early
   1980s did not imagine the dilemma of scale that the Internet faces
   today. 1987 estimates projected a need to address as many as 100,000
   networks at some vague point in the future. [Callon87] We will reach
   that mark by 1996. There are many realistic projections of many
   millions of interconnected networks in the not too distant future.
   [Vecchi94, Taylor94]

   Further, even though the current 32 bit IPv4 address structure can
   enumerate over 4 billion hosts on as many as 16.7 million networks,
   the actual address assignment efficiency is far less than that, even
   on a theoretical basis. [Huitema94] This inefficiency is exacerbated
   by the granularity of assignments using Class A, B and C addresses.

   In August 1990 during the Vancouver IETF meeting, Frank Solensky,
   Phill Gross and Sue Hares projected that the current rate of
   assignment would exhaust the Class B space by March of 1994.

   The then obvious remedy of assigning multiple Class C addresses in
   place of Class B addresses introduced its own problem by further
   expanding the size of the routing tables in the backbone routers
   already growing at an alarming rate.

   We faced the dilemma of choosing between accepting either limiting
   the rate of growth and ultimate size of the Internet, or disrupting
   the network by changing to new techniques or technologies.

   The IETF formed the Routing and Addressing (ROAD) group in November
   1991 at the Santa Fe IETF meeting to explore this dilemma and guide
   the IETF on the issues. The ROAD group reported their work in March
   1992 at the San Diego IETF meeting. [Gross92] The impact of the
   recommendations ranged from "immediate" to "long term" and included
   adopting the CIDR route aggregation proposal [Fuller93] for reducing
   the rate of routing table growth and recommending a call for
   proposals "to form working groups to explore separate approaches for
   bigger Internet addresses."

In the late spring of 1992 the IAB issued "IP version 7" [IAB92], concurring in the ROAD group's endorsement of CIDR and also recommending "an immediate IETF effort to prepare a detailed and organizational plan for using CLNP as the basis for IPv7." After spirited discussion, the IETF decided to reject the IAB's recommendation and issue the call for proposals recommended by the ROAD group. This call was issued in July 1992 at the Boston IETF meeting and a number of working groups were formed in response

During the July 1993 Amsterdam IETF meeting an IPng (IP Next Generation) Decision Process (ipdecide) BOF was held. This BOF "was intended to help re-focus attention on the very important topic of making a decision between the candidates for IPng. The BOF focused on the issues of who should take the lead in making the recommendation to the community and what criteria should be used to reach the recommendation." [Carpen93]

3. A Direction for IPng

   In September 1993 Phill Gross, chair of the IESG issued "A Direction for IPng". [Gross94] In this memo he summarized the results of the ipdecide BOF and open IESG plenary in Amsterdam.

   * The IETF needs to move toward closure on IPng.
   * The IESG has the responsibility for developing an IPng
     recommendation for the Internet community.
   * The procedures of the recommendation-making process should be open
     and published well in advance by the IESG.
   * As part of this process, the IPng WGs may be given new milestones
     and other guidance to aid the IESG.
   * There should be ample opportunity for community comment prior to
     final IESG recommendation.

   The memo also announced "a temporary, ad hoc, 'area' to deal specifically with IPng issues." Phill asked two of the current IESG members, Allison Mankin (Transport Services Area) and Scott Bradner (Operational Requirements Area), to act as Directors for the new area. The Area Directors were given a specific charge on how to investigate the various IPng proposals and how to base their recommendation to the IETF. It was also requested that a specific recommendation be made.

   * Establish an IPng directorate.
   * Ensure that a completely open process is followed.
   * Develop an understanding of the level of urgency and the time
     constraints imposed by the rate of address assignment and rate of
     growth in the routing tables.
   * Recommend the adoption of assignment policy changes if warranted.

      * Define the scope of the IPng effort based on the understanding of
        the time constraints.
      * Develop a clear and concise set of technical requirements and
        decision criteria for IPng.
      * Develop a recommendation about which of the current IPng candidates
        to accept, if any.

4. IPng Area

   After the IPng Area was formed, we recruited a directorate. (Appendix
   B) The members of the directorate were chosen both for their general
   and specific technical expertise. The individuals were then asked to
   have their management authorize this participation in the process and
   confirm that they understood the IETF process.

   We took great care to ensure the inclusion of a wide spectrum of
   knowledge. The directors are experts in security, routing, the needs
   of large users, end system manufacturers, Unix and non-Unix
   platforms, router manufacturers, theoretical researchers, protocol
   architecture, and the operating regional, national, and international
   networks. Additionally, several members of the directorate were
   deeply involved in each of the IPng proposal working groups.

   The directorate functions as a direction-setting and preliminary
   review body as requested by the charge to the area. The directorate
   engages in biweekly conference calls, participates in an internal
   mailing list and corresponds actively on the Big-Internet mailing
   list. The directorate held open meetings during the March 1994
   Seattle and July 1994 Toronto IETF meetings as well as two additional
   multi-day retreats. To ensure that the IPng process was as open as
   possible, we took minutes during these meetings and then published
   them. Additionally, we placed the archives of the internal IPng
   mailing list on an anonymous ftp site. (Hsdndev.harvard.edu:
   pub/ipng.)

5. ALE Working Group

   We needed a reasonable estimate of the time remaining before we
   exhausted the IPv4 address space in order to determine the scope of
   the IPng effort. If the time remaining was about the same needed to
   deploy a replacement, then we would have select the IPng which would
   only fix the address limitations since we would not have enough time
   to develop any other features. If more time seemed available, we
   could consider additional improvements.

   The IETF formed an Address Lifetime Expectations (ALE) Working Group
   in 1993 "to develop an estimate for the remaining lifetime of the
   IPv4 address space based on currently known and available

technologies." [Solens93a] Tony Li of Cisco Systems and Frank
Solensky of FTP Software are the co-chairs. The IETF also charged
the working group to consider if developing more stringent address
allocation and utilization policies might provide more time for the
transition.

5.1 ALE Projections

   The ALE Working Group met during the November 1993 Houston,
   [Solens93b] March 1994 Seattle [Bos93] and July 1994 Toronto
   [Solens94] IETF meetings. They projected at the Seattle meeting,
   later confirmed at the Toronto meeting that, using the current
   allocation statistics, the Internet would exhaust the IPv4 address
   space between 2005 and 2011.

   Some members of the ipv4-ale and big-internet mailing lists called
   into question the reliability of this projection. It has been
   criticized as both too optimistic and as too pessimistic.

   Some people pointed out that this type of projection makes an
   assumption of no paradigm shifts in IP usage. If someone were to
   develop a new 'killer application', (for example cable-TV set top
   boxes.) The resultant rise in the demand for IP addresses could make
   this an over-estimate of the time available.

   There may also be a problem with the data used to make the
   projection. The InterNIC allocates IP addresses in large chunks to
   regional Network Information Centers (NICs) and network providers.
   The NICs and the providers then re-allocate addresses to their
   customers. The ALE projections used the InterNIC assignments without
   regard to the actual rate of assignment of addresses to the end
   users. They did the projection this way since the accuracy of the
   data seems quite a bit higher. While using this once-removed data
   may add a level of over-estimation since it assumes the rate of large
   block allocation will continue, this may not be the case.

   These factors reduce the reliability of the ALE estimates but, in
   general, they seem to indicate enough time remaining in the IPv4
   address space to consider adding features in an IPng besides just
   expanding the address size even when considering time required for
   development, testing, and deployment.

5.2 Routing Table Size

   Another issue in Internet scaling is the increasing size of the
   routing tables required in the backbone routers. Adopting the CIDR
   block address assignment and aggregating routes reduced the size of
   the tables for awhile but they are now expanding again. Providers now

need to more aggressively advertise their routes only in aggregates.
Providers must also advise their new customers to renumber their
networks in the best interest of the entire Internet community.

The problem of exhausting the IPv4 address space may be moot if this
issue is ignored and if routers cannot be found that can keep up with
the table size growth. Before implementing CIDR the backbone routing
table was growing at a rate about 1.5 times as fast as memory
technology.

We should also note that even though IPng addresses are designed with
aggregation in mind switching to IPng will not solve the routing
table size problem unless the addresses are assigned rigorously to
maximize the affect of such aggregation. This efficient advertising
of routes can be maintained since IPng includes address
autoconfiguration mechanisms to allow easy renumbering if a customer
decides to switch providers. Customers who receive service from more
than one provider may limit the ultimate efficiency of any route
aggregation. [Rekhter94]

5.3 Address Assignment Policy Recommendations

The IESG Chair charged the IPng Area to consider recommending more
stringent assignment policies, reclaiming some addresses already
assigned, or making a serious effort to renumber significant portions
of the Internet. [Gross94]

The IPng Area Directors endorse the current address assignment
policies in view of the ALE projections. We do not feel that anyone
should take specific efforts to reclaim underutilized addresses
already assigned or to renumber forcefully major portions of the
Internet. We do however feel that we should all encourage network
service providers to assist new customers in renumbering their
networks to conform to the provider's CIDR assignments.

The ALE Working Group recommends that we consider assigning CIDR-type
address blocks out of the unassigned Class A address space. The IPng
Area Directors concur with this recommendation.

6. IPng Technical Requirements

The IESG provided an outline in RFC 1380 [Gross92] of the type of
criteria we should use to determine the suitability of an IPng
proposal. The IETF further refined this understanding of the
appropriate criteria with the recommendations of a Selection Criteria
BOF held during the November 1992 IETF meeting in Washington D.C.
[Almqu92] We felt we needed to get additional input for determining
the requirements and issued a call for white papers. [Bradner93] This

call, issued as RFC 1550, intended to reach both inside and outside
the traditional IETF constituency to get the broadest possible
understanding of the requirements for a data networking protocol with
the broadest possible application.

We received twenty one white papers in response to the RFC 1550
solicitation. (Appendix E) We received responses from the
industries that many feel will be the major providers of data
networking services in the future; the cable TV industry [Vecchi94],
the cellular industry [Taylor94], and the electric power industry
[Skelton94]. In addition, we received papers that dealt with
military applications [Adam94, Syming94, Green94], ATM [Brazd94],
mobility [Simpson94], accounting [Brown94], routing [Estrin94a,
Chiappa94], security [Adam94, Bell94b, Brit94, Green94, Vecchi94,
Flei94], large corporate networking [Britt94, Fleisch94], transition
[Carpen94a, Heager94], market acceptance [Curran94, Britt94], host
implementations [Bound94], as well as a number of other issues.
[Bello94a, Clark94, Ghisel94]

These white papers, a Next Generation Requirements (ngreq) BOF
(chaired by Jon Crowcroft and Frank Kastenholz) held during the March
1994 Seattle IETF meeting, discussions within the IPng Area
Directorate and considerable discussion on the big-internet mailing
list were all used by Frank Kastenholz and Craig Partridge in
revising their earlier criteria draft [Kasten92] to produce
"Technical Criteria for Choosing IP The Next Generation (IPng)."
[Kasten94] This document is the "clear and concise set of technical
requirements and decision criteria for IPng" called for in the charge
from the IESG Chair. We used this document as the basic guideline
while evaluating the IPng proposals.

6.1 The IPng Technical Criteria document

The criteria described in this document include: (from Kasten94)

* complete specification - The proposal must completely describe the
  proposed protocol. We must select an IPng by referencing specific
  documents, not to future work.
* architectural simplicity - The IP-layer protocol should be as
  simple as possible with functions located elsewhere that are more
  appropriately performed at protocol layers other than the IP layer.
* scale - The IPng Protocol must allow identifying and addressing at
  least 10**9 leaf-networks (and preferably much more)
* topological flexibility - The routing architecture and protocols
  ofIPng must allow for many different network topologies. They must
  not assume that the network's physical structure is a tree.
* performance - A state of the art, commercial grade router must be
  able to process and forward IPng traffic at speeds capable of fully

     utilizing common, commercially available, high-speed media at the
     time.
   * robust service - The network service and its associated routing and
     control protocols must be robust.
   * transition - The protocol must have a straightforward transition
     plan from IPv4.
   * media independence - The protocol must work across an internetwork
     of many different LAN, MAN, and WAN media, with individual link
     speeds ranging from a ones-of-bits per second to hundreds of
     gigabits per second.
   * datagram service - The protocol must support an unreliable datagram
     delivery service.
   * configuration ease - The protocol must permit easy and largely
     distributed configuration and operation. Automatic configuration of
     hosts and routers is required.
   * security - IPng must provide a secure network layer.
   * unique names - IPng must assign unique names to all IP-Layer
     objects in the global, ubiquitous, Internet. These names may or
     may not have any location, topology, or routing significance.
   * access to standards - The protocols that define IPng and its
     associated protocols should be as freely available and
     redistributable as the IPv4 and related RFCs. There must be no
     specification-related licensing fees for implementing or selling
     IPng software.
   * multicast support - The protocol must support both unicast and
     multicast packet transmission. Dynamic and automatic routing of
     multicasts is also required.
   * extensibility - The protocol must be extensible; it must be able
     to evolve to meet the future service needs of the Internet. This
     evolution must be achievable without requiring network-wide
     software upgrades.
   * service classes - The protocol must allow network devices to
     associate packets with particular service classes and provide them
     with the services specified by those classes.
   * mobility - The protocol must support mobile hosts, networks and
     internetworks.
   * control protocol - The protocol must include elementary support for
     testing and debugging networks. (e.g., ping and traceroute)
   * tunneling support - IPng must allow users to build private
     internetworks on top of the basic Internet Infrastructure. Both
     private IP-based internetworks and private non-IP-based (e.g., CLNP
     or AppleTalk) internetworks must be supported.

7. IPng Proposals

   By the time that the IPng Area was formed, the IETF had already aimed
   a considerable amount of IETF effort at solving the addressing and
   routing problems of the Internet. Several proposals had been made
   and some of these reached the level of having a working group
   chartered. A number of these groups subsequently merged forming
   groups with a larger consensus. These efforts represented different
   views on the issues which confront us and sought to optimize
   different aspects of the possible solutions.

   By February 1992 the Internet community developed four separate
   proposals for IPng [Gross92], "CNAT" [Callon92a], "IP Encaps"
   [Hinden92a], "Nimrod" [Chiappa91], and "Simple CLNP" [Callon92b]. By
   December 1992 three more proposals followed; "The P Internet
   Protocol" (PIP) [Tsuchiya92], "The Simple Internet Protocol" (SIP)
   [Deering92] and "TP/IX" [Ullmann93]. After the March 1992 San Diego
   IETF meeting "Simple CLNP" evolved into "TCP and UDP with Bigger
   Addresses" (TUBA) [Callon92c] and "IP Encaps" evolved into "IP
   Address Encapsulation" (IPAE) [Hinden92b].

   By November 1993, IPAE merged with SIP while still maintaining the
   name SIP. This group then merged with PIP and the resulting working
   group called themselves "Simple Internet Protocol Plus" (SIPP). At
   the same time the TP/IX Working Group changed its name to "Common
   Architecture for the Internet" (CATNIP).

   None of these proposals were wrong nor were others right. All of the
   proposals would work in some ways providing a path to overcome the
   obstacles we face as the Internet expands. The task of the IPng Area
   was to ensure that the IETF understand the offered proposals, learn
   from the proposals and provide a recommendation on what path best
   resolves the basic issues while providing the best foundation upon
   which to build for the future.

   The IPng Area evaluated three IPng proposals as they were described
   in their RFC 1550 white papers: CATNIP [McGovern94] , SIPP
   [Hinden94a] and TUBA. [Ford94a]. The IESG viewed Nimrod as too much
   of a research project for consideration as an IPng candidate. Since
   Nimrod represents one possible future Internet routing strategy we
   solicited a paper describing any requirements Nimrod would put on an
   IPng to add to the requirements process. [Chiappa94]

7.1 CATNIP

   "Common Architecture for the Internet (CATNIP) was conceived as a
   convergence protocol. CATNIP integrates CLNP, IP, and IPX. The CATNIP
   design provides for any of the transport layer protocols in use, for

example TP4, CLTP, TCP, UDP, IPX and SPX, to run over any of the
network layer protocol formats: CLNP, IP (version 4), IPX, and
CATNIP. With some attention paid to details, it is possible for a
transport layer protocol (such as TCP) to operate properly with one
end system using one network layer (e.g., IP version 4) and the other
using some other network protocol, such as CLNP." [McGovern94]

"The objective is to provide common ground between the Internet, OSI,
and the Novell protocols, as well as to advance the Internet
technology to the scale and performance of the next generation of
internetwork technology."

"CATNIP supports OSI Network Service Access Point (NSAP) format
addresses. It also uses cache handles to provide both rapid
identification of the next hop in high performance routing as well as
abbreviation of the network header by permitting the addresses to be
omitted when a valid cache handle is available. The fixed part of the
network layer header carries the cache handles." [Sukonnik94]

7.2 SIPP

"Simple Internet Protocol Plus (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 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. 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." [Hinden94b]

"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."

"SIPP 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. 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." [Hinden94a]

7.3 TUBA

   "The TCP/UDP Over CLNP-Addressed Networks (TUBA) proposal seeks to
   minimize the risk associated with migration to a new IP address
   space. In addition, this proposal is motivated by the requirement to
   allow the Internet to scale, which implies use of Internet
   applications in a very large ubiquitous worldwide Internet. It is
   therefore proposed that existing Internet transport and application
   protocols continue to operate unchanged, except for the replacement
   of 32-bit IP addresses with larger addresses. TUBA does not mean
   having to move over to OSI completely. It would mean only replacing
   IP with CLNP. TCP, UDP, and the traditional TCP/IP applications would
   run on top of CLNP." [Callon92c]

   "The TUBA effort will expand the ability to route Internet packets by
   using addresses which support more hierarchy than the current
   Internet Protocol (IP) address space. TUBA specifies the continued
   use of Internet transport protocols, in particular TCP and UDP, but
   specifies their encapsulation in ISO 8473 (CLNP) packets. This will
   allow the continued use of Internet application protocols such as
   FTP, SMTP, TELNET, etc.   TUBA seeks to upgrade the current system by
   a transition from the use of IPv4 to ISO/IEC 8473 (CLNP) and the
   corresponding large Network Service Access Point (NSAP) address
   space." [Knopper94]

   "The TUBA proposal makes use of a simple long-term migration proposal
   based on a gradual update of Internet Hosts (to run Internet
   applications over CLNP) and DNS servers (to return larger addresses).
   This proposal requires routers to be updated to support forwarding of
   CLNP (in addition to IP). However, this proposal does not require
   encapsulation nor translation of packets nor address mapping. IP
   addresses and NSAP addresses may be assigned and used independently
   during the migration period. Routing and forwarding of IP and CLNP
   packets may be done independently." ([Callon92c]

8. IPng Proposal Reviews

   The IPng Directorate discussed and reviewed the candidate proposals
   during its biweekly teleconferences and through its mailing list. In
   addition, members of the Big-Internet mailing list discussed many of
   the aspects of the proposals, particularly when the Area Directors
   posted several specific questions to stimulate discussion. [Big]

   The directorate members were requested to each evaluate the proposals
   in preparation for a two day retreat held near Chicago on May 19th
   and 20th 1994. The retreat opened with a roundtable airing of the

views of each of the participants, including the Area Directors, the
Directorate and a number of guests invited by the working group
chairs for each for the proposals. [Knopper94b] We will publish
these reviews as well as a more detailed compendium review of each of
the proposals as companion memos.

The following table summarizes each of the three proposals reviewed
against the requirements in the IPng Criteria document. They do not
necessarily reflect the views of the Area Directors. "Yes" means the
reviewers mainly felt the proposal met the specific criterion. "No"
means the reviewers mainly felt the proposal did not meet the
criterion. "Mixed" means that the reviewers had mixed reviews with
none dominating. "Unknown" means that the reviewers mainly felt the
documentation did not address the criterion.

	CATNIP	SIPP	TUBA
complete spec	no	yes	mostly
simplicity	no	no	no
scale	yes	yes	yes
topological flex	yes	yes	yes
performance	mixed	mixed	mixed
robust service	mixed	mixed	yes
transition	mixed	no	mixed
media indepdnt	yes	yes	yes
datagram	yes	yes	yes
config. ease	unknown	mixed	mixed
security	unknown	mixed	mixed
unique names	mixed	mixed	mixed
access to stds	yes	yes	mixed
multicast	unknown	yes	mixed
extensibility	unknown	mixed	mixed
service classes	unknown	yes	mixed
mobility	unknown	mixed	mixed
control proto	unknown	yes	mixed
tunneling	unknown	yes	mixed

8.1 CATNIP Reviews

All the reviewers felt that CATNIP is not completely specified.
However, many of the ideas in CATNIP are innovative and a number of
reviewers felt CATNIP shows the best vision of all of the proposals.
The use of Network Service Attachment Point Addresses (NSAPs) is well
thought out and the routing handles are innovative.

While the goal of uniting three major protocol families, IP, ISO-CLNP
and Novell IPX is laudable our consensus was that the developers had
not developed detailed enough plans to support realizing that goal.

The plans they do describe suffer from the complexity of trying to be the union of a number of existing network protocols. Some reviewers felt that CATNIP is basically maps IPv4, IPX, and SIPP addresses into NSAPs and, as such, does not deal with the routing problems of the current and future Internet.

Additionally the reviewers felt that CATNIP has poor support for multicasting and mobility and does not specifically deal with such important topics as security and autoconfiguration.

8.2 SIPP Reviews

Most of the reviewers, including those predisposed to other proposals, felt as one reviewer put it, that SIPP is an "aesthetically beautiful protocol well tailored to compactly satisfy today's known network requirements." The SIPP Working Group has been the most dynamic over the last year, producing a myriad of documentation detailing almost all of the aspects necessary to produce a complete protocol description.

The biggest problem the reviewers had with SIPP was with IPAE, SIPP's transition plan. The overwhelming feeling was that IPAE is fatally flawed and could not be made to work reliably in an operational Internet.

There was significant disagreement about the adequacy of the SIPP 64 bit address size. Although you can enumerate 10**15 end nodes in 64 bits people have different views about how much inefficiency real-world routing plans introduce. [Huitema94] The majority felt that 64 bit addresses do not provide adequate space for the hierarchy required to meet the needs of the future Internet. In addition since no one has any experience with extended addressing and routing concepts of the type proposed in SIPP, the reviewers generally felt quite uncomfortable with this methodology. The reviewers also felt that the design introduces some significant security issues.

A number of reviewers felt that SIPP did not address the routing issue in any useful way. In particular, there has been no serious attempt made at developing ways to abstract topology information or to aggregate information about areas of the network.

Finally, most of the reviewers questioned the level of complexity in the SIPP autoconfiguration plans as well as in SIPP in general, other than the header itself.

8.3 TUBA Reviews

   The reviewers generally felt that the most important thing that TUBA
   has offers is that it is based on CLNP and there is significant
   deployment of CLNP-capable routers throughout the Internet. There
   was considerably less agreement that there was significant deployment
   of CLNP-capable hosts or actual networks running CLNP. Another
   strong positive for TUBA is the potential for convergence of ISO and
   IETF networking standards. A number of reviewers pointed out that,
   if TUBA were to be based on a changed CLNP then the advantage of an
   existing deployed infrastructure would be lost and that the
   convergence potential would be reduced.

   A number of aspects of CLNP were felt to be a problem by the
   reviewers including the inefficiencies introduced by the lack of any
   particular word alignment of the header fields, CLNP source route,
   the lack of a flow ID field, the lack of a protocol ID field, and the
   use of CLNP error messages in TUBA. The CLNP packet format or
   procedures would have to be modified to resolve at least some of
   these issues.

   There seems to be a profound disagreement within the TUBA community
   over the question of the ability of the IETF to modify the CLNP
   standards. In our presentation in Houston we said that we felt that
   "clone and run" was a legitimate process. This is also what the IAB
   proposed in "IP version 7". [IAB92] The TUBA community has not
   reached consensus that this view is reasonable. While many,
   including a number of the CLNP document authors, are adamant that
   this is not an issue and the IETF can make modifications to the base
   standards, many others are just as adamant that the standards can
   only be changed through the ISO standards process. Since the
   overwhelming feeling within the IETF is that the IETF must 'own' the
   standards on which it is basing its future, this disagreement within
   the TUBA community was disquieting.

   For a number of reasons, unfortunately including prejudice in a few
   cases, the reviews of the TUBA proposals were much more mixed than
   for SIPP or CATNIP. Clearly TUBA meets the requirements for the
   ability to scale to large numbers of hosts, supports flexible
   topologies, is media independent and is a datagram protocol. To the
   reviewers, it was less clear that TUBA met the other IPng
   requirements and these views varied widely.

   There was also disagreement over the advisability of using NSAPs for
   routing given the wide variety of NSAP allocation plans. The
   Internet would have to restrict the use of NSAPs to those which are
   allocated with the actual underlying network topology in mind if the
   required degree of aggregation of routing information is to be

achieved.

8.4 Summary of Proposal Reviews

To summarize, significant problems were seen in all three of the
proposals. The feeling was that, to one degree or another, both SIPP
and TUBA would work in the Internet context but each exhibited its
own problems. Some of these problems would have to be rectified
before either one would be ready to replace IPv4, much less be the
vehicle to carry the Internet into the future. Other problems could
be addressed over time. CATNIP was felt to be too incomplete to be
considered.

9. A Revised Proposal

As mentioned above, there was considerable discussion of the
strengths and weaknesses of the various IPng proposals during the
IPng 'BigTen' retreat on May 19th and 20th 1994. [Knopper94b] After
this retreat Steve Deering and Paul Francis, two of the co-chairs of
the SIPP Working Group, sent a message to the sipp mailing list
detailing the discussions at the retreat and proposing some changes
in SIPP. [Deering94a]

The message noted "The recurring (and unsurprising) concerns about
SIPP were:

(1) complexity/manageability/feasibility of IPAE, and

(2) adequacy/correctness/limitations of SIPP's addressing and routing
    model, especially the use of loose source routing to accomplish
    'extended addressing'".

They "proposed to address these concerns by changing SIPP as follows:

* Change address size from 8 bytes to 16 bytes (fixed-length).

* Specify optional use of serverless autoconfiguration of the 16-byte
  address by using IEEE 802 address as the low-order ("node ID")
  part.

* For higher-layer protocols that use internet-layer addresses as
  part of connection identifiers (e.g., TCP), require that they use
  the entire 16-byte addresses.

* Do *not* use Route Header for extended addressing."

After considerable discussion on the sipp and big-internet mailing
lists about these proposed changes, the SIPP working group published
a revised version of SIPP [Deering94b], a new addressing architecture
[Francis94], and a simplified transition mechanism [Gillig94a].
These were submitted to the IPng Directorate for their consideration.

This proposal represents a synthesis of multiple IETF efforts with
much of the basic protocol coming from the SIPP effort, the
autoconfiguration and transition portions influenced by TUBA, the
addressing structure is based on the CIDR work and the routing header
evolving out of the SDRP deliberations.

10. Assumptions

10.1 Criteria Document and Timing of Recommendation

In making the following recommendations we are making two assumptions
of community consensus; that the IPng criteria document represents
the reasonable set of requirements for an IPng, and that a specific
recommendation should be made now and that from this point on the
IETF should proceed with a single IPng effort.

As described above, the IPng Technical Criteria document [Kasten94]
was developed in a open manner and was the topic of extensive
discussions on a number of mailing lists. We believe that there is a
strong consensus that this document accurately reflects the
community's set of technical requirements which an IPng should be
able to meet.

A prime topic of discussion on the big-internet mailing list this
spring as well as during the open IPng directorate meeting in
Seattle, was the need to make a specific IPng recommendation at this
time. Some people felt that additional research would help resolve
some of the issues that are currently unresolved. While others
argued that selecting a single protocol to work on would clarify the
picture for the community, focus the resources of the IETF on
finalizing its details, and, since the argument that there were open
research items could be made at any point in history, there might
never be a 'right' time.

Our reading of the community is that there is a consensus that a
specific recommendation should be made now. This is consistent with
the views expressed during the ipdecide BOF in Amsterdam [Gross94]
and in some of the RFC 1550 white papers [Carpen94a].

There is no particular reason to think that the basic recommendation
would be significantly different if we waited for another six months
or a year. Clearly some details which are currently unresolved could

be filled in if the recommendation were to be delayed, but the
current fragmentation of the IETF's energies limits the efficiency of
this type of detail resolution. Concentrating the resources of the
IETF behind a single effort seems to us to be a more efficient way to
proceed.

10.2 Address Length

One of the most hotly discussed aspects of the IPng design
possibilities was address size and format. During the IPng process
four distinct views were expressed about these issues:

1. The view that 8 bytes of address are enough to meet the current
   and future needs of the Internet (squaring the size of the IP
   address space). More would waste bandwidth, promote inefficient
   assignment, and cause problems in some networks (such as mobiles
   and other low speed links).

2. The view that 16 bytes is about right. That length supports easy
   auto-configuration as well as organizations with complex internal
   routing topologies in conjunction with the global routing topology
   now and well into the future.

3. The view that 20 byte OSI NSAPs should be used in the interests of
   global harmonization.

4. The view that variable length addresses which might be smaller or
   larger than 16 bytes should be used to embrace all the above
   options and more, so that the size of the address could be
   adjusted to the demands of the particular environment, and to
   ensure the ability to meet any future networking requirements.

Good technical and engineering arguments were made for and against
all of these views. Unanimity was not achieved, but we feel that a
clear majority view emerged that the use of 16 byte fixed length
addresses was the best compromise between efficiency, functionality,
flexibility, and global applicability. [Mankin94]

11. IPng Recommendation

After a great deal of discussion in many forums and with the
consensus of the IPng Directorate, we recommend that the protocol
described in "Simple Internet Protocol Plus (SIPP) Spec. (128 bit
ver)" [Deering94b] be adopted as the basis for IPng, the next
generation of the Internet Protocol. We also recommend that the
other documents listed in Appendix C be adopted as the basis of
specific features of this protocol.

This proposal resolves most of the perceived problems, particularly
in the areas of addressing, routing, transition and address
autoconfiguration. It includes the broad base of the SIPP proposal
effort, flexible address autoconfiguration features, and a merged
transition strategy. We believe that it meets the requirements
outlined in the IPng Criteria document and provides the framework to
fully meet the needs of the greater Internet community for the
foreseeable future.

11.1 IPng Criteria Document and IPng

A detailed review of how IPng meets the requirements set down in the
IPng Criteria document [Kasten94] will soon be published. Following
is our feelings about the extent to which IPng is responsive to the
criteria.

   * complete specification - the base specifications for IPng are
     complete but transition and address autoconfiguration do remain to
     be finalized
   * architectural simplicity - the protocol is simple, easy to explain
     and uses well established paradigms
   * scale - an address size of 128 bits easily meets the need to
     address 10**9 networks even in the face of the inherent
     inefficiency of address allocation for efficient routing
   * topological flexibility - the IPng design places no constraints on
     network topology except for the limit of 255 hops
   * performance - the simplicity of processing, the alignment of the
     fields in the headers, and the elimination of the header checksum
     will allow for high performance handling of IPng data streams
   * robust service - IPng includes no inhibitors to robust service and
     the addition of packet-level authentication allows the securing of
     control and routing protocols without having to have separate
     procedures
   * transition - the IPng transition plan is simple and realistically
     covers the transition methods that will be present in the
     marketplace
   * media independence - IPng retains IPv4's media independence, it may
     be possible to make use of IPng's Flow Label in some connection-
     oriented media such as ATM
   * datagram service - IPng preserves datagram service as its basic
     operational mode, it is possible that the use of path MTU discovery
     will complicate the use of datagrams in some cases
   * configuration ease - IPng will have easy and flexible address
     autoconfiguration which will support a wide variety of environments
     from nodes on an isolated network to nodes deep in a complex
     internet
   * security - IPng includes specific mechanisms for authentication and
     encryption at the internetwork layer; the security features do rely

on the presence of a yet to be defined key management system
   * unique names - IPng addresses may be used as globally unique names
     although they do have topological significance
   * access to standards - all of the IPng standards will be published
     as RFCs with unlimited distribution
   * multicast support - IPng specifically includes multicast support
   * extensibility - the use of extension headers and an expandable
     header option feature will allow the introduction of new features
     into IPng when needed in a way that minimizes the disruption of the
     existing network
   * service classes - the IPng header includes a Flow Label which may
     be used to differentiate requested service classes
   * mobility - the proposed IPv4 mobility functions will work with IPng
   * control protocol - IPng includes the familiar IPv4 control protocol
     features
   * tunneling support - encapsulation of IPng or other protocols within
     IPng is a basic capability described in the IPng specifications

11.2 IPv6

   The IANA has assigned version number 6 to IPng. The protocol itself
   will be called IPv6.

   The remainder of this memo is used to describe IPv6 and its features.
   This description is an overview snapshot. The standards documents
   themselves should be referenced for definitive specifications. We
   also make a number of specific recommendations concerning the details
   of the proposed protocol, the procedures required to complete the
   definition of the protocol, and the IETF working groups we feel are
   necessary to accomplish the task.

12. IPv6 Overview

   IPv6 is a new version of the Internet Protocol, it has been designed
   as an evolutionary, rather than revolutionary, step from IPv4.
   Functions which are generally seen as working in IPv4 were kept in
   IPv6. Functions which don't work or are infrequently used were
   removed or made optional. A few new features were added where the
   functionality was felt to be necessary.

   The important features of IPv6 include: [Hinden94c]

   * expanded addressing and routing capabilities - The IP address size
     is increased from 32 bits to 128 bits providing support for a much
     greater number of addressable nodes, more levels of addressing
     hierarchy, and simpler auto-configuration of addresses.

     The scaleability of multicast routing is improved by adding a
     "scope" field to multicast addresses.

     A new type of address, called a "cluster address" is defined to
     identify topological regions rather than individual nodes. The use
     of cluster addresses in conjunction with the IPv6 source route
     capability allows nodes additional control over the path their
     traffic takes.

   * simplified header format - 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 overhead of the IPv6
     header as low as possible in spite of the increased size of the
     addresses. Even though the IPv6 addresses are four time longer
     than the IPv4 addresses, the IPv6 header is only twice the size of
     the IPv4 header.

   * support for extension headers and options - IPv6 options are placed
     in separate headers that are located in the packet between the IPv6
     header and the transport-layer header. Since most IPv6 option
     headers are not examined or processed by any router along a
     packet's delivery path until it arrives at its final destination,
     this organization facilitates a major improvement in router
     performance for packets containing options. Another improvement is
     that unlike IPv4, IPv6 options can be of arbitrary length and not
     limited to 40 bytes. This feature plus the manner in which they are
     processed, permits IPv6 options to be used for functions which were
     not practical in IPv4.

     A key extensibility feature of IPv6 is the ability to encode,
     within an option, the action which a router or host should perform
     if the option is unknown. This permits the incremental deployment
     of additional functionality into an operational network with a
     minimal danger of disruption.

   * support for authentication and privacy - IPv6 includes the
     definition of an extension which provides support for
     authentication and data integrity. This extension is included as a
     basic element of IPv6 and support for it will be required in all
     implementations.

     IPv6 also includes the definition of an extension to support
     confidentiality by means of encryption. Support for this extension
     will be strongly encouraged in all implementations.

   * support for autoconfiguration - IPv6 supports multiple forms of
     autoconfiguration, from "plug and play" configuration of node
     addresses on an isolated network to the full-featured facilities
     offered by DHCP.

   * support for source routes - IPv6 includes an extended function
     source routing header designed to support the Source Demand Routing
     Protocol (SDRP). The purpose of SDRP is to support source-initiated
     selection of routes to complement the route selection provided by
     existing routing protocols for both inter-domain and intra-domain
     routes. [Estrin94b]

   * simple and flexible transition from IPv4 - The IPv6 transition plan
     is aimed at meeting four basic requirements: [Gillig94a]

     - Incremental upgrade. Existing installed IPv4 hosts and routers
       may be upgraded to IPv6 at any time without being dependent on
       any other hosts or routers being upgraded.
     - Incremental deployment. New IPv6 hosts and routers can be
       installed at any time without any prerequisites.
     - Easy Addressing. When existing installed IPv4 hosts or routers
       are upgraded to IPv6, they may continue to use their existing
       address. They do not need to be assigned new addresses.
     - Low start-up costs. Little or no preparation work is needed in
       order to upgrade existing IPv4 systems to IPv6, or to deploy new
       IPv6 systems.

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

12.1 IPv6 Header Format

   The IPv6 header, although longer than the IPv4 header, is
   considerably simplified. A number of functions that were in the IPv4
   header have been relocated in extension headers or dropped.
   [Deering94b]

```
    +-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
    |Version|                    Flow Label                         |
    +-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
    |         Payload Length        |  Next Header  |   Hop Limit   |
    +-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
    |                                                               |
    +                                                               +
    |                                                               |
    +                        Source Address                         +
    |                                                               |
    +                                                               +
    |                                                               |
    +-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
    |                                                               |
    +                                                               +
    |                                                               |
    +                      Destination Address                      +
    |                                                               |
    +                                                               +
    |                                                               |
    +-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
```
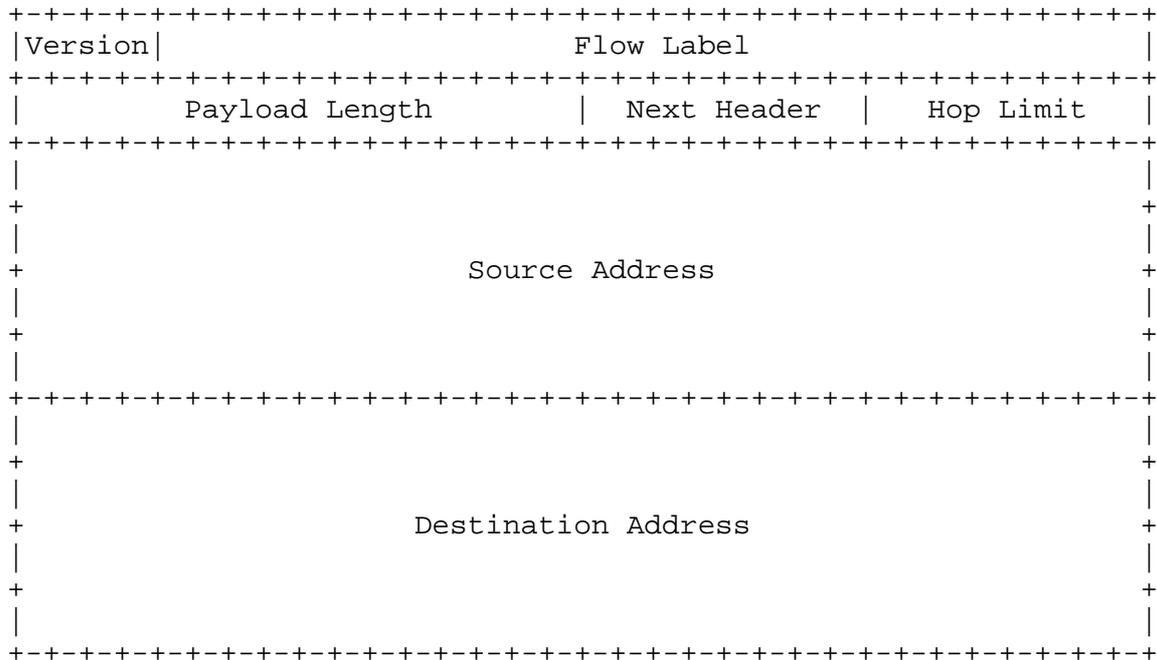
   * Version - Internet Protocol version number. IPng has been assigned
     version number 6. (4-bit field)

   * Flow Label - This field may be used by a host to label those
     packets for which it is requesting special handling by routers
     within a network, such as non-default quality of service or "real-
     time" service. (28-bit field)

   * Payload Length - Length of the remainder of the packet following
     the IPv6 header, in octets. To permit payloads of greater than 64K
     bytes, if the value in this field is 0 the actual packet length
     will be found in an Hop-by-Hop option. (16-bit unsigned integer)

   * Next Header - Identifies the type of header immediately following
     the IPv6 header. The Next Header field uses the same values as the
   IPv4 Protocol field (8-bit selector field)

   * Hop Limit - Used to limit the impact of routing loops. The Hop
     Limit field is decremented by 1 by each node that forwards the
     packet. The packet is discarded if Hop Limit is decremented to
     zero. (8-bit unsigned integer)

   * Source Address - An address of the initial sender of the packet.
     (128 bit field)

   * Destination Address - An address of the intended recipient of the
     packet (possibly not the ultimate recipient, if an optional Routing
     Header is present). (128 bit field)

12.2 Extension Headers

   In IPv6, optional internet-layer information is encoded in separate
   headers that may be placed between the IPv6 header and the
   transport-layer header in a packet. There are a small number of such
   extension headers, each identified by a distinct Next Header value.
   [From a number of the documents listed in Appendix C.]

12.2.1 Hop-by-Hop Option Header

   The Hop-by-Hop Options header is used to carry optional
   information that must be examined by every node along a packet's
   delivery path. The Hop-by-Hop Options header is identified by a
   Next Header value of 0 in the IPv6 header, and has the following
   format:

```
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
|   Next Header  |  Hdr Ext Len  |                             |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+                             +
|                                                             |
.                                                             .
.                           Options                           .
.                                                             .
|                                                             |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
```
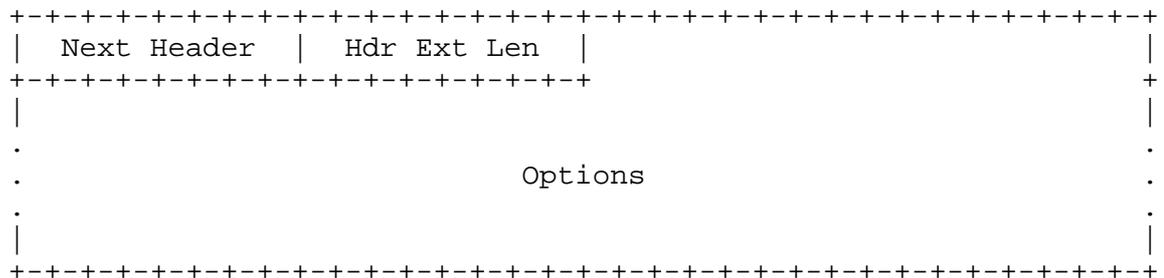
   * Next Header - Identifies the type of header immediately
     following the Hop-by-Hop Options header. Uses the same values
     as the IPv4 Protocol field. (8-bit selector)

   * Hdr Ext Len - Length of the Hop-by-Hop Options header in 8-octet
     units, not including the first 8 octets. (8-bit unsigned
     integer)

   * Options - Contains one or more TLV-encoded options. (Variable-
     length field, of length such that the complete Hop-by-Hop
     Options header is an integer multiple of 8 octets long.)

   12.2.2 IPv6 Header Options

   Two of the currently-defined extension headers -- the Hop-by-Hop
   Options header and the End-to-End Options header -- may carry a
   variable number of Type-Length-Value (TLV) encoded "options", of
   the following format:

   +-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+- - - - - - - - -
   | Option Type | Opt Data Len | Option Data
   +-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+- - - - - - - - -
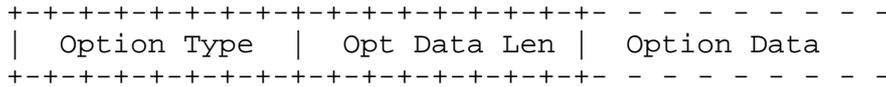
   * Option Type - identifier of the type of option (8-bit field)

   * Opt Data Len - Length of the Option Data field of this option,
     in octets. (8-bit unsigned integer)

   * Option Data - Option-Type-specific data. (Variable-length field)

   The Option Type identifiers are internally encoded such that their
   highest-order two bits specify the action that must be taken if
   the processing IPv6 node does not recognize the Option Type:

   00 - skip over this option and continue processing the header
   01 - discard the packet
   10 - discard the packet and send an ICMP Unrecognized Type message
          to the packet's Source Address, pointing to the unrecognized
          Option Type
   11 - undefined.

   In the case of Hop-by-Hop options only, the third-highest-order
   bit of the Option Type specifies whether or not the Option Data of
   this option shall be included in the integrity assurance
   computation performed when an Authentication header is present.
   Option data that changes en route must be excluded from that
   computation.

12.2.3 Routing Header

   The Routing header is used by an IPv6 source to list one or more
   intermediate nodes (or topological clusters) to be "visited" on
   the way to a packet's destination. This particular form of the
   Routing Header is designed to support SDRP. [Estrin94b]

```
 0                   1                   2                   3
 0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
| Next Header   |Routing Type=1 |M|F| Reserved    | SrcRouteLen |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
| NextHopPtr    |              Strict/Loose Bit Mask            |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
|                                                               |
.                                                               .
.                         Source Route                          .
.                                                               .
|                                                               |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
```
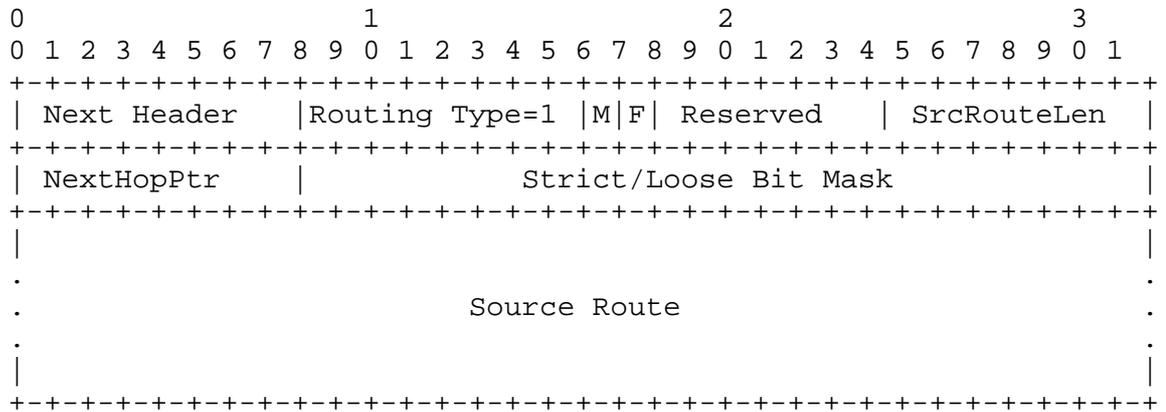
     * Next Header - Identifies the type of header immediately
       following the Routing Header. Uses the same values as the IPv4
       Protocol field. (8-bit selector)

     * Routing Type - Indicates the type of routing supported by this
       header. Value must be 1.

     * MRE flag - Must Report Errors. If this bit is set to 1, and a
       router can not further forward a packet (with an incompletely
       traversed source route), as specified in the Source Route, the
       router must generate an ICMP error message. If this bit is set
       to 0, and a router can not further forward a packet (with an
       incompletely traversed source route), as specified in the Source
       Route, the router should not generate an ICMP error message.

     * F flag - Failure of Source Route Behavior. If this bit it set
       to 1, it indicates that if a router can not further forward a
       packet (with an incompletely traversed source route), as
       specified in the Source Route, the router must set the value of
       the Next Hop Pointer field to the value of the Source Route
       Length field, so that the subsequent forwarding will be based
       solely on the destination address. If this bit is set to 0, it
       indicates that if a router can not further forward a packet
       (with an incompletely traversed source route), as specified in
       the Source Route, the router must discard the packet.

        * Reserved - Initialized to zero for transmission; ignored on
          reception.

        * SrcRouteLen - Source Route Length - Number of source route
          elements/hops in the SDRP Routing header. Length of SDRP
          routing header can be calculated from this value (length =
          SrcRouteLen * 16 + 8) This field may not exceed a value of 24.
          (8 bit unsigned integer)

        * NextHopPtr - Next Hop Pointer- Index of next element/hop to be
          processed; initialized to 0 to point to first element/hop in the
          source route. When Next Hop Pointer is equal to Source Route
          Length then the Source Route is completed. (8 bit unsigned
          integer)

        * Strict/Loose Bit Mask - The Strict/Loose Bit Mask is used when
          making a forwarding decision. If the value of the Next Hop
          Pointer field is N, and the N-th bit in the Strict/Loose Bit
          Mask field is set to 1, it indicates that the next hop is a
          Strict Source Route Hop. If this bit is set to 0, it indicates
          that the next hop is a Loose Source Route Hop. (24 bit
          bitpattern)

        * Source Route - A list of IPv6 addresses indicating the path that
          this packet should follow. A Source Route can contain an
          arbitrary intermix of unicast and cluster addresses. (integral
          multiple of 128 bits)

   12.2.4 Fragment Header

      The Fragment header is used by an IPv6 source to send payloads
      larger than would fit in the path MTU to their destinations.
      (Note: unlike IPv4, fragmentation in IPv6 is performed only by
      source nodes, not by routers along a packet's delivery path) The
      Fragment header is identified by a Next Header value of 44 in the
      immediately preceding header, and has the following format:

        +-+
        | Next Header | Reserved | Fragment Offset |Res|M|
        +-+
        | Identification |
        +-+
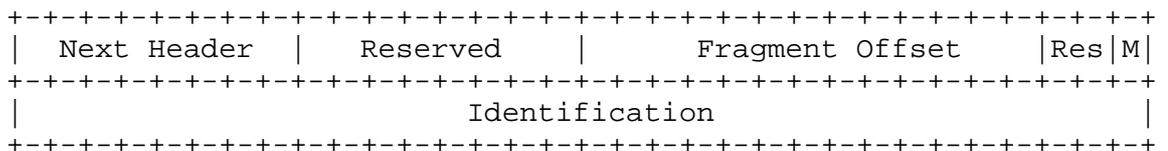
        * Next Header - Identifies the type of header immediately
          following the Fragment header. Uses the same values as the IPv4
          Protocol field. (8 bit selector)

    * Reserved, Res - Initialized to zero for transmission; ignored on
      reception.

    * Fragment Offset - The offset, in 8-octet units, of the following
      payload, relative to the start of the original, unfragmented
      payload. (13-bit unsigned integer)

    * M flag - 1 = more fragments; 0 = last fragment.

    * Identification - A value assigned to the original payload that
      is different than that of any other fragmented payload sent
      recently with the same IPv6 Source Address, IPv6 Destination
      Address, and Fragment Next Header value. (If a Routing header
      is present, the IPv6 Destination Address is that of the final
      destination.) The Identification value is carried in the
      Fragment header of all of the original payload's fragments, and
      is used by the destination to identify all fragments belonging
      to the same original payload. (32 bit field)

  12.2.5 Authentication Header

    The Authentication header is used to provide authentication and
    integrity assurance for IPv6 packets. Non-repudiation may be
    provided by an authentication algorithm used with the
    Authentication header, but it is not provided with all
    authentication algorithms that might be used with this header.
    The Authentication header is identified by a Next Header value of
    51 in the immediately preceding header, and has the following
    format:

```
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
| Next Header  | Auth Data Len |             Reserved           |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
|                     Security Association ID                   |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
|                                                              |
.                                                              .
.                      Authentication Data                     .
.                                                              .
|                                                              |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
```
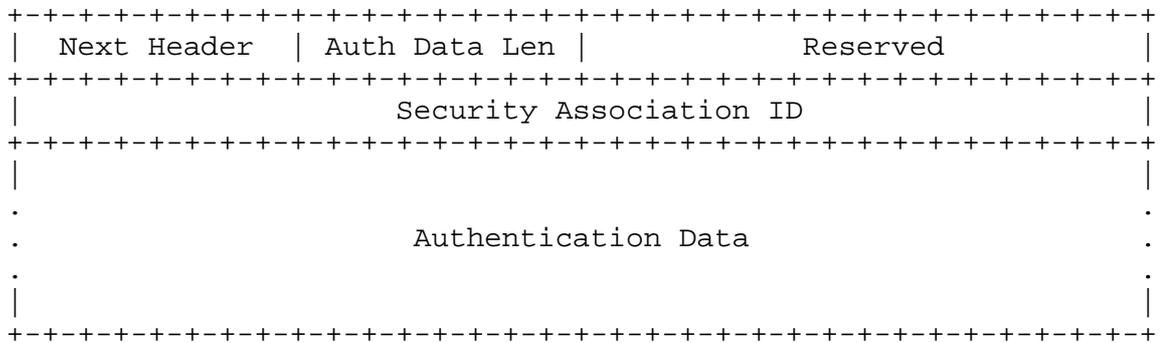
    * Next Header - Identifies the type of header immediately
      following the Authentication header. Uses the same values as
      the IPv4 Protocol field. (8-bit selector)

    * Auth Data Len - Length of the Authentication Data field in 8-
      octet units. (8-bit unsigned integer)

    * Reserved - Initialized to zero for transmission; ignored on
      reception.

    * Security Assoc. ID - When combined with the IPv6 Source Address,
      identifies to the receiver(s) the pre-established security
      association to which this packet belongs. (32 bit field)

    * Authentication Data - Algorithm-specific information required
      to authenticate the source of the packet and assure its
      integrity, as specified for the pre-established security
      association. (Variable-length field, an integer multiple of 8
      octets long.)

   12.2.6 Privacy Header

   The Privacy Header seeks to provide confidentiality and integrity
   by encrypting data to be protected and placing the encrypted data
   in the data portion of the Privacy Header. Either a transport-
   layer (e.g., UDP or TCP) frame may be encrypted or an entire IPv6
   datagram may be encrypted, depending on the user's security
   requirements. This encapsulating approach is necessary to provide
   confidentiality for the entire original datagram. If present, the
   Privacy Header is always the last non-encrypted field in a packet.

   The Privacy Header works between hosts, between a host and a
   security gateway, or between security gateways. This support for
   security gateways permits trustworthy networks to exist without
   the performance and monetary costs of security, while providing
   security for traffic transiting untrustworthy network segments.
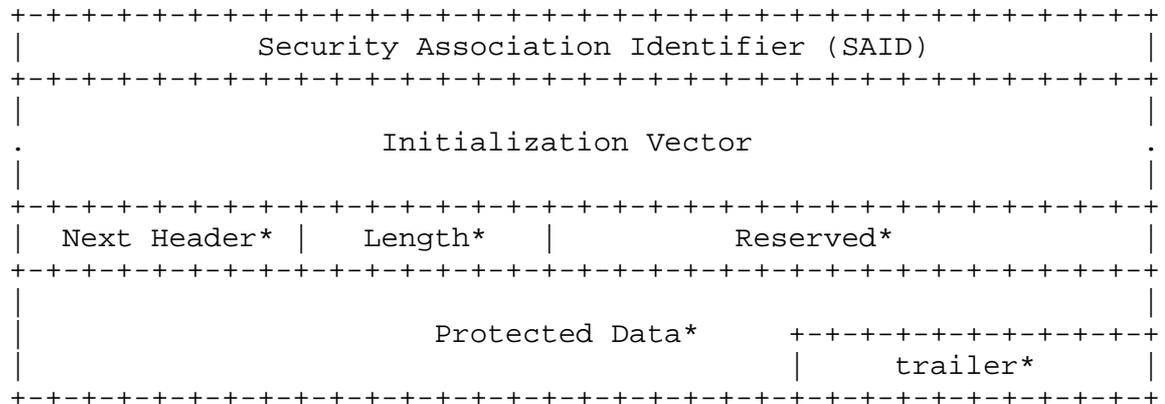
```
    +-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
    |            Security Association Identifier (SAID)             |
    +-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
    |                                                              |
    .                   Initialization Vector                      .
    |                                                              |
    +-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
    | Next Header* |  Length*   |          Reserved*               |
    +-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
    |                                                              |
    |              Protected Data*      +-+-+-+-+-+-+-+-+-+-+       |
    |                                   |       trailer*     |     |
    +-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+

                                                       *encrypted
```

   * Security Association Identifier (SAID) - Identifies the security
     association for this datagram. If no security association has
     been established, the value of this field shall be 0x0000. A
     security association is normally one-way. An authenticated
     communications session between two hosts will normally have two
     SAIDs in use (one in each direction). The receiving host uses
     the combination of SAID value and originating address to
     distinguish the correct association. (32 bit value)

   * Initialization Vector - This field is optional and its value
     depends on the SAID in use. For example, the field may contain
     cryptographic synchronization data for a block oriented
     encryption algorithm. It may also be used to contain a
     cryptographic initialization vector. A Privacy Header
     implementation will normally use the SAID value to determine
     whether this field is present and, if it is, the field's size
     and use. (presence and length dependent on SAID)

   * Next Header - encrypted - Identifies the type of header
     immediately following the Privacy header. Uses the same values
     as the IPv4 Protocol field. (8 bit selector)

   * Reserved - encrypted - Ignored on reception.

   * Length - encrypted - Length of the Privacy Header in 8-octet
     units, not including the first 8 octets. (8-bit unsigned
     integer)

   * Protected Data - encrypted - This field may contain an entire
     encapsulated IPv6 datagram, including the IPv6 header, a
     sequence of zero or more IPv6 options, and a transport-layer
     payload, or it may just be a sequence of zero or more IPv6
     options followed by a transport-layer payload. (variable
     length)

   * trailer (Algorithm-dependent Trailer) - encrypted - A field
     present to support some algorithms which need to have padding
     (e.g., to a full cryptographic block size for block-oriented
     encryption algorithms) or for storage of authentication data for
     use with a encryption algorithm that provides confidentiality
     without authentication. It is present only when the algorithm
     in use requires such a field. (presence and length dependent on
     SAID)

   12.2.7 End-to-End Option Header

      The End-to-End Options header is used to carry optional
      information that needs to be examined only by a packet's
      destination node(s). The End-to-End Options header is identified
      by a Next Header value of TBD in the immediately preceding header,
      and has the same format as the Hop-by-Hop Option Header except for
      the ability to exclude an option from the authentication integrity
      assurance computation.

13. IPng Working Group

   We recommend that a new IPng Working Group be formed to produce
   specifications for the core functionality of the IPv6 protocol suite.
   The working group will carry out the recommendations of the IPng Area
   Directors as outlined at the July 1994 IETF and in this memo. We
   recommend that this working group be chaired by Steve Deering of
   Xerox PARC and Ross Callon of Wellfleet.

   The primary task of the working group is to produce a set of
   documents that define the basic functions, interactions, assumptions,
   and packet formats for IPv6. We recommend that Robert Hinden of Sun
   Microsystems be the editor for these documents. The documents listed
   in Appendix C will be used by the working group to form the basis of
   the final document set.

   The work of the IPng Working Group includes:

   * complete the IPv6 overview document
   * complete the IPv6 detailed operational specification
   * complete the IPv6 Addressing Architecture specification
   * produce specifications for IPv6 encapsulations over various media
   * complete specifications for the support of packets larger than 64KB
   * complete specifications of the DNS enhancements required to support
     IPv6
   * complete specification of ICMP, IGMP and router discovery for
     support of IPv6.
   * complete specification of path MTU discovery for IPv6
   * complete specifications of IPv6 in IPv6 tunneling
   * complete the suggested address format and assignment plan
   * coordinate with the Address Autoconfiguration Working Group
   * coordinate with the NGTRANS and TACIT Working Groups
   * complete specifications of authentication and privacy support
     headers

The working group should also consider a few selected enhancements
including:

* consider ways to compress the IPv6 header in the contexts of native
  IPv6, multiple IPv6 packets in a flow, and encapsulated IPv6
* consider specifying support for a larger minimum MTU

14. IPng Reviewer

   Currently it is the task of the IPng Area Directors, the IPng
   Directorate and the chairs of the proposed ipng working group to
   coordinate the activities of the many parallel efforts currently
   directed towards different aspects of IPng. While this is possible
   and currently seems to be working well it can not be maintained over
   the long run because, among other reasons, the IPng Area will be
   dissolved eventually and its Directorate disbanded. It will also
   become much more difficult as IPng related activities start up in
   other IETF areas.

   We recommend that an IPng Reviewer be appointed to be specifically
   responsible for ensuring that a consistent view of IPv6 is maintained
   across the related working groups. We feel that this function is
   required due to the complex nature of the interactions between the
   parts of the IPng effort and due to the distribution of the IPng
   related work amongst a number of IETF areas. We recommend that Dave
   Clark of MIT be offered this appointment.

   This would be a long-term task involving the review of on-going
   activities. The aim is not for the IPng Reviewer to make
   architectural decisions since that is the work of the various working
   groups, the IAB, and the IETF as a whole.. The aim is to spot gaps or
   misunderstandings before they reach the point where functionality or
   interworkability is threatened.

15. Address Autoconfiguration

   As data networks become more complex the need to be able to bypass at
   least some of the complexity and move towards "plug and play" becomes
   ever more acute. The user can not be expected to be able to
   understand the details of the network architecture or know how to
   configure the network software in their host. In the ideal case, a
   user should be able to unpack a new computer, plug it into the local
   network and "just" have it work without requiring the entering of any
   special information. Security concerns may restrict the ability to
   offer this level of transparent address autoconfiguration in some
   environments but the mechanisms must be in place to support whatever
   level of automation which the local environment feels comfortable
   with.

The basic requirement of "plug and play" operation is that a host
must be able to acquire an address dynamically, either when attaching
to a network for the first time or when the host needs to be
readdressed because the host moved or because the identity of the
network has changed. There are many other functions required to
support a full "plug and play" environment. [Berk94] Most of these
must be addressed outside of the IPv6 Area but a focused effort to
define a host address autoconfiguration protocol is part of the IPv6
process.

We recommend that a new Address Autoconfiguration Working Group
(addrconf) be formed with Dave Katz of Cisco Systems and Sue Thomson
of Bellcore as co-chairs. The purpose of this working group is to
design and specify a protocol for allocating addresses dynamically to
IPv6 hosts. The address configuration protocol must be suitable for
a wide range of network topologies, from a simple isolated network to
a sophisticated globally connected network. It should also allow for
varying levels of administrative control, from completely automated
operation to very tight oversight.

The scope of this working group is to propose a host address
autoconfiguration protocol which supports the full range of
topological and administrative environments in which IPv6 will be
used. It is the intention that, together with IPv6 system discovery,
the address autoconfiguration protocol will provide the minimal
bootstrapping information necessary to enable hosts to acquire
further configuration information (such as that provided by DHCP in
IPv4). The scope does not include router configuration or any other
host configuration functions. However, it is within the scope of the
working group to investigate and document the interactions between
this work and related functions including system discovery, DNS
autoregistration, service discovery, and broader host configuration
issues, to facilitate the smooth integration of these functions.
[Katz94a]

The working group is expected to complete its work around the end of
1994 and disband at that time. The group will use "IPv6 Address
Autoconfiguration Architecture" [Katz94b] draft document as the basis
of their work.

16. Transition

The transition of the Internet from IPv4 to IPv6 has to meet two
separate needs. There is a short term need to define specific
technologies and methods to transition IPv4 networks, including the
Internet, into IPv6 networks and an IPv6 Internet. There is also a
long term need to do broad-based operational planning for transition,
including developing methods to allow decentralized migration

strategies, understanding the ramifications of a long period of
coexistence when both protocols are part of the basic infrastructure,
developing an understanding of the type and scope of architectural
and interoperability testing that will be required to ensure a
reliable and manageable Internet in the future.

16.1 Transition - Short Term

   Any IPng transition plan must take into account the realities of what
   types of devices vendors will build and network managers will deploy.
   The IPng transition plan must define the procedures required to
   successfully implement the functions which vendors will be likely to
   include in their devices. This is the case even if there are good
   arguments to recommend against a particular function, header
   translation for example. If products will exist it is better to have
   them interoperate than not.

   We recommend that a new IPng Transition (NGTRANS) Working Group be
   formed with Bob Gilligan of Sun Microsystems and xxx of yyy as co-
   chairs to design the mechanisms and procedures to support the
   transition of the Internet from IPv4 to IPv6 and to give advice on
   what procedures and techniques are preferred.

   The work of the group will fall into three areas:

   * Define the processes by which the Internet will make the transition
     from IPv4 to IPv6. As part of this effort, the group will produce
     a document explaining to the general Internet community what
     mechanisms will be employed in the transition, how the transition
     will work, the assumptions about infrastructure deployment inherent
     in the operation of these mechanisms, and the types of
     functionality that applications developers will be able to assume
     as the protocol mix changes over time.
   * Define and specify the mandatory and optional mechanisms that
     vendors should implement in hosts, routers, and other components of
     the Internet in order for the transition to be carried out. Dual-
     stack, encapsulation and header translation mechanisms must all be
     defined, as well as the interaction between hosts using different
     combinations of these mechanisms. The specifications produced will
     be used by people implementing these IPv6 systems.
   * Articulate a concrete operational plan for the Internet to make the
     transition from IPv4 to IPv6. The result of this work will be a
     transition plan for the Internet that network operators and
     Internet subscribers can execute.
                                                            [Gillig94c]

The working group is expected to complete its work around the end of
1994 and disband at that time. The group will use the "Simple SIPP
Transition (SST)" [Gilig94a] overview document as the starting point
for its work.

16.2 Transition - Long Term

There are a number of transition related topics in addition to
defining the specific IPv4 to IPv6 mechanisms and their deployment,
operation and interaction. The ramifications and procedures of
migrating to a new technology or to a new version of an existing
technology must be fully understood.

We recommend that the Transition and Coexistence Including Testing
(TACIT) Working Group, which was started a few months ago, explore
some of the basic issues associated with the deployment of new
technology into an established Internet. The TACIT Working Group
will focus on the generic issues of transition and will not limit
itself to the upcoming transition to IPv6 because, over time,
enhancements to IPv6 (IPv6ng) will be developed and accepted. At
that point they will need to be deployed into the then existing
Internet. The TACIT Working Group will be more operationally
oriented than the NGTRANS Working Group and will continue well into
the actual IPv6 transition.

The main areas of exploration are:

* Make the transition from a currently deployed protocol to a new
  protocol while accommodating heterogeneity and decentralized
  management.
* Since it is often difficult or impossible to replace all legacy
  systems or software, it is important to understand the
  characteristics and operation of a long period of coexistence
  between a new protocol and the existing protocol.
* The Internet must now be considered a utility. We are far removed
  from a time when a new technology could be deployed to see if it
  would work in large scale situations. Rigorous architectural and
  interoperability testing must be part of the predeployment phase of
  any proposed software for the Internet. Testing the scaling up
  behaviors and robustness of a new protocol will offer particular
  challenges. The WG should determine if there are lessons to be
  learned from: OSPF, BGP4 and CIDR Deployment, the AppleTalk 1 to 2
  transition, DECnet Phase 4 to Phase 5 planning and transition,
  among others.

The TACIT Working Group will explore each of these facets of the
deployment of new technology and develop a number of documents to
help guide users and managers of affected data networks and provide
to the IETF:

* Detailed descriptions of problem areas in transition and
  coexistence, both predicted, based on lessons learned, and observed
  as the IPv6 process progresses.
* Recommendations for specific testing procedures.
* Recommendations for coexistence operations procedures
* Recommendations for the smoothing of decentralized transition
  planning.

                                                        [Huston94]

17. Other Address Families

   There are many environments in which there are one or more network
   protocols already deployed or where a significant planning effort has
   been undertaken to create a comprehensive network addressing plan. In
   such cases there may be a temptation to integrate IPv6 into the
   environment by making use of an existing addressing plan to define
   all or part of the IPv6 addresses. The advantage of doing this is
   that it permits unified management of address space among multiple
   protocol families. The use of common addresses can help facilitate
   transition from other protocols to IPv6.

   If the existing addresses are globally unique and assigned with
   regard to network topology this may be a reasonable idea. The IETF
   should work with other organizations to develop algorithms that could
   be used to map addresses between IPv6 and other environments. The
   goal for any such mapping must be to provide an unambiguous 1 to 1
   map between individual addresses.

   Suggestions have been made to develop mapping algorithms for Novell
   IPX addresses, some types of OSI NSAPs, E164 addresses and SNA
   addresses. Each of these possibilities should be carefully examined
   to ensure that use of such an algorithm solves more problems than it
   creates. In some cases it may be better to recommend either that a
   native IPng addressing plan be developed instead, or that an IPv6
   address be used within the non-IP environment. [Carpen94b]

   We recommend that, in conjunction with other organizations,
   recommendations about the use of non-IPv6 addresses in IPv6
   environments and IPv6 addresses in non-IPv6 environments be
   developed.

18. Impact on Other IETF Standards

   Many current IETF standards are affected by IPv6. At least 27 of the
   current 51 full Internet Standards must be revised for IPv6, along
   with at least 6 of the 20 Draft Standards and at least 25 of the 130
   Proposed Standards. [Postel94]

   In some cases the revisions consist of simple changes to the text,
   for example, in a number of RFCs an IP address is referred to in
   passing as a "32 bit IP address" even though IP addresses are not
   directly used in the protocol being defined. All of the standards
   track documents will have to be checked to see if they contain such
   references.

   In most of the rest of the cases revisions to the protocols,
   including packet formats, will be required. In many of these cases
   the address is just being carried as a data element and a revised
   format with a larger field for the address will have no effect on the
   functional paradigm.

   In the remaining cases some facet of the operation of the protocol
   will be changed as a result of IPv6. For example, the security and
   source route mechanisms are fundamentally changed from IPv4 with
   IPv6. Protocols and applications that relied on the IPv4
   functionality will have to be redesigned or rethought to use the
   equivalent function in IPv6.

   In a few cases this opportunity should be used to determine if some
   of the RFCs should be moved to historic, for example EGP [Mills84]
   and IP over ARCNET. [Provan91]

   The base IPng Working Group will address some of these, existing IETF
   working groups can work on others, while new working groups must be
   formed to deal with a few of them. The IPng Working Group will be
   responsible for defining new versions of ICMP, ARP/RARP, and UDP. It
   will also review RFC 1639, "FTP Operation Over Big Address Records
   (FOOBAR)" [Piscit94] and RFC 1191 "Path MTU Discovery" [Mogul90]

   Existing working groups will examine revisions for some of the
   routing protocols: RIPv2, IS-IS, IDRP and SDRP. A new working group
   may be required for OSPF.

   The existing DHCP Working Group may be able to revise DHCP and
   examine BOOTP.

A TCPng Working Group will be formed soon, and new working groups
will have to be formed to deal with standards such as SNMP, DNS, NTP,
NETbios, OSI over TCP, Host Requirements, and Kerberos as well as
reviewing most of the RFCs that define IP usage over various media.

In addition to the standards track RFCs mentioned above there are
many Informational and Experimental RFCs which would be affected as
well as numerous Internet Drafts (and those standards track RFCs that
we missed).

We recommend that the IESG commission a review of all standards track
RFCs to ensure that a full list of affected documents is compiled. We
recommend that the IESG charge current IETF working groups with the
task of understanding the impact of IPv6 on their proposals and,
where appropriate, revise the documents to include support for IPv6.

We recommend that the IESG charter new working groups where required
to revise other standards RFCs.

19. Impact on non-IETF standards and on products

Many products and user applications which rely on the size or
structure of IPv4 addresses will need to be modified to work with
IPv6. While the IETF can facilitate an investigation of the impacts
of IPv6 on non-IETF standards and products, the primary
responsibility for doing so resides in the other standards bodies and
the vendors.

Examples of non-IETF standards that are effected by IPv6 include the
POSIX standards, Open Software Foundation's DCE and DME, X-Open, Sun
ONC, the Andrew File System and MIT's Kerberos. Most products that
provide specialized network security including firewall-type devices
are among those that must be extended to support IPv6.

20. APIs

It is traditional to state that the IETF does not "do" APIs. While
there are many reasons for this, the one most commonly referenced is
that there are too many environments where TCP/IP is used, too many
different operating systems, programming languages, and platforms.
The feeling is that the IETF should not get involved in attempting to
define a language and operating system independent interface in the
face of such complexity.

We feel that this historical tendency for the IETF to avoid dealing
with APIs should be reexamined in the case of IPv6. We feel that in
a few specific cases the prevalence of a particular type of API is
such that a single common solution for the modifications made

necessary by IPv6 should be documented.

We recommend that Informational RFCs be solicited or developed for
these few cases. In particular, the Berkeley-style sockets
interface, the UNIX TLI and XTI interfaces, and the WINSOCK
interfaces should be targeted. A draft document exists which could
be developed into the sockets API description. [Gillig94b]

21. Future of the IPng Area and Working Groups

In our presentation at the Houston IETF meeting we stated that the
existing IPng proposal working groups would not be forced to close
down after the recommendation was made. Each of them has been
working on technologies that may have applications in addition to
their IPng proposal and these technologies should not be lost.

Since the Toronto IETF meeting the existing IPng working groups have
been returned to the Internet Area. The group members may decide to
close down the working groups or to continue some of their efforts.
The charters of the working groups must be revised if they choose to
continue since they would no longer be proposing an IPng candidate.

In Toronto the chairs of the SIPP Working Group requested that the
SIPP Working Group be concluded. The chairs of the TUBA Working
Group requested that the TUBA working group be understood to be in
hiatus until a number of the documents in process were completed, at
which time they would request that the working group be concluded.

We recommend that the IPng Area and its Directorate continue until
the basic documents have entered the standards track in late 1994 or
early 1995 and that after such time the area be dissolved and those
IPng Area working groups still active be moved to their normal IETF
areas.

22. Security Considerations

The security of the Internet has long been questioned. It has been
the topic of much press coverage, many conferences and workshops.
Almost all of this attention has been negative, pointing out the many
places where the level of possible security is far less than that
deemed necessary for the current and future uses of the Internet. A
number of the RFC 1550 White Papers specifically pointed out the
requirement to improve the level of security available [Adam94,
Bell94b, Brit94, Green94, Vecchi94, Flei94] as does "Realizing the
Information Future". [Nat94]

In February of 1994, the IAB convened a workshop on security in the
Internet architecture. The report of this workshop [IAB94] includes
an exploration of many of the security problem areas and makes a
number of recommendations to improve the level of security that the
Internet offers its users.

We feel that an improvement in the basic level of security in the
Internet is vital to its continued success. Users must be able to
assume that their exchanges are safe from tampering, diversion and
exposure. Organizations that wish to use the Internet to conduct
business must be able to have a high level of confidence in the
identity of their correspondents and in the security of their
communications. The goal is to provide strong protection as a matter
of course throughout the Internet.

As the IAB report points out, many of the necessary tools are not a
function of the internetworking layer of the protocol. These higher
level tools could make use of strong security features in the
internetworking layer if they were present. While we expect that
there will be a number of special high-level security packages
available for specific Internet constituencies, support for basic
packet-level authentication will provide for the adoption of a much
needed, widespread, security infrastructure throughout the Internet.

It is best to separate the support for authentication from the
support for encryption. One should be able to use the two functions
independently. There are some applications in which authentication
of a corespondent is sufficient and others where the data exchanged
must be kept private.

It is our recommendation that IPv6 support packet authentication as a
basic and required function. Applications should be able to rely on
support for this feature in every IPv6 implementation. Support for a
specific authentication algorithm should also be mandated while
support for additional algorithms should be optional.

Thus we recommend that support for the Authentication Header be
required in all compliant IPv6 implementations.

We recommend that support for a specific authentication algorithm be
required. The specific algorithm should be determined by the time
the IPv6 documents are offered as Proposed Standards.

We recommend that support for the Privacy Header be required in IPv6
implementations.

We recommend that support for a privacy authentication algorithm be
required. The specific algorithm should be determined by the time the
IPv6 documents are offered as Proposed Standards.

Clearly, a key management infrastructure will be required in order to
enable the use of the authentication and encryption headers.
However, defining such an infrastructure is outside the scope of the
IPv6 effort. We do note that there are on-going IETF activities in
this area. The IPv6 transition working groups must coordinate with
these activities.

Just as clearly, the use of authentication and encryption may add to
the cost and impact the performance of systems but the more secure
infrastructure is worth the penalty. Whatever penalty there is
should also decrease in time with improved software and hardware
assistance.

The use of firewalls is increasing on the Internet. We hope that the
presence of the authentication and privacy features in IPv6 will
reduce the need for firewalls, but we do understand that they will
continue to be used for the foreseeable future. In this light, we
feel that clear guidance should be given to the developers of
firewalls on the best ways to design and configure them when working
in an IPv6 environment.

We recommend that an "IPv6 framework for firewalls" be developed.
This framework should explore the ways in which the Authentication
Header can be used to strengthen firewall technology and detail how
the IPv6 packet should be analyzed by a firewall.

Some aspects of security require additional study. For example, it
has been pointed out [Vecchi94] that, even in non-military
situations, there are places where procedures to thwart traffic
analysis will be required. This could be done by the use of
encrypted encapsulation, but this and other similar requirements must
be addressed on an on-going basis by the Security Area of the IETF.
The design of IPv6 must be flexible enough to support the later
addition of such security features.

We believe that IPv6 with its inherent security features will provide
the foundation upon which the Internet can continue to expand its
functionality and user base.

23. Authors' Addresses

   Scott Bradner
   Harvard University
   10 Ware St.
   Cambridge, MA 02138

   Phone: +1 617 495 3864
   EMail: sob@harvard.edu

   Allison Mankin
   USC/Information Sciences Institute
   4350 North Fairfax Drive, Suite 400
   Arlington, VA 22303

   Phone: +1 703-807-0132
   EMail: mankin@isi.edu

Appendix A - Summary of Recommendations

    5.3 Address Assignment Policy Recommendations
       changes in address assignment policies are not recommended
       reclamation of underutilized assigned addresses is not currently
          recommended
       efforts to renumber significant portions of the Internet are not
          currently recommended
       recommend consideration of assigning CIDR-type address blocks out
          of unassigned Class A addressees
    11. IPng Recommendation
       recommend that "Simple Internet Protocol Plus (SIPP) Spec. (128
          bit ver)" [Deering94b] be adopted as the basis for IPng
       recommend that the documents listed in Appendix C be the basis of
          IPng
    13. IPng Working Group
       recommend that an IPng Working Group be formed, chaired by Steve
          Deering and Ross Callon
       recommend that Robert Hinden be the document editor for the IPng
          effort
    14. IPng Reviewer
       recommend that an IPng Reviewer be appointed and that Dave Clark
          be that reviewer
    15. Address Autoconfiguration
       recommend that an Address Autoconfiguration Working Group be
          formed, chaired by Dave Katz and Sue Thomson
    16.1 Transition - Short Term
       recommend that an IPng Transition Working Group be formed, chaired
          by Bob Gilligan and TBA
    16.2 Transition - Long Term
       recommend that the Transition and Coexistence Including Testing
          Working Group be chartered
    17. Other Address Families
       recommend that recommendations about the use of non-IPv6 addresses
          in IPv6 environments and IPv6 addresses in non-IPv6
          environments be developed
    18. Impact on Other IETF Standards
       recommend the IESG commission a review of all standards track RFCs
       recommend the IESG charge current IETF working groups with the
          task of understanding the impact of IPng on their proposals
          and, where appropriate, revise the documents to include support
          for IPng
       recommend the IESG charter new working groups where required to
          revise other standards RFCs
    20. APIs
       recommend that Informational RFCs be developed or solicited for a
          few of the common APIs

   21. Future of the IPng Area and Working Groups
      recommend that the IPng Area and Area Directorate continue until
         main documents are offered as Proposed Standards in late 1994
   22. Security Considerations
      recommend that support for the Authentication Header be required
      recommend that support for a specific authentication algorithm be
         required
      recommend that support for the Privacy Header be required
      recommend that support for a specific privacy algorithm be
         required
      recommend that an "IPng framework for firewalls" be developed

Appendix B - IPng Area Directorate

   J. Allard - Microsoft   <jallard@microsoft.com>
   Steve Bellovin - AT&T   <smb@research.att.com>
   Jim Bound - Digital   <bound@zk3.dec.com>
   Ross Callon - Wellfleet   <rcallon@wellfleet.com>
   Brian Carpenter - CERN   <brian.carpenter@cern.ch>
   Dave Clark - MIT   <ddc@lcs.mit.edu >
   John Curran - NEARNET   <curran@nic.near.net>
   Steve Deering - Xerox   <deering@parc.xerox.com>
   Dino Farinacci - Cisco   <dino@cisco.com>
   Paul Francis - NTT   <francis@slab.ntt.jp>
   Eric Fleischmann - Boeing   <ericf@atc.boeing.com>
   Mark Knopper - Ameritech   <mak@aads.com>
   Greg Minshall - Novell   <minshall@wc.novell.com>
   Rob Ullmann - Lotus   <ariel@world.std.com>
   Lixia Zhang - Xerox   <lixia@parc.xerox.com>

   Daniel Karrenberg of RIPE joined the Directorate when it was formed
   but had to withdraw due to the demands of his day job.

   Since the Toronto IETF meeting Paul Francis has resigned from the
   Directorate to pursue other interests. Robert Hinden of Sun
   Microsystems and Yakov Rekhter of IBM joined.

Appendix C - Documents Referred to the IPng Working Groups

   [Deering94b] Deering, S., "Simple Internet Protocol Plus (SIPP) Spec.
      (128 bit ver)", Work in Progress.

   [Francis94] Francis, P., "SIPP Addressing Architecture", Work in
      Progress.

   [Rekhter94] Rekhter, Y., and T. Li, "An Architecture for IPv6 Unicast
      Address Allocation", Work in Progress.

   [Gillig94a] Gilligan, R., "Simple SIPP Transition (SST) Overview",
      Work in Progress.

   [Gillig94b] Gilligan, R., Govindan, R., Thomson, S., and J. Bound,
      "SIPP Program Interfaces for BSD Systems", Work in Progress.

   [Atkins94a] Atkinson, R., "SIPP Security Architecture", Work in
      Progress.

   [Atkins94b] Atkinson, R., "SIPP Authentication Header", Work in
      Progress.

   [Ford94b] Ford, P., Li, T., and Y. Rekhter, "SDRP Routing Header for
      SIPP-16", Work in Progress.

   [Hinden94c] Hinden, R., "IP Next Generation Overview", Work in
      Progress.

Appendix D - IPng Proposal Overviews

   [Ford94a] Ford, P., and M. Knopper, "TUBA as IPng: A White Paper",
      Work in Progress.

   [Hinden94a] Hinden, R., "Simple Internet Protocol Plus White Paper",
      RFC 1710, Sun Microsystems, October 1994.

   [McGovern94] McGovern, M., and R. Ullmann, "CATNIP: Common
      Architecture for the Internet", RFC 1707, Sunspot Graphics, Lotus
      Development Corp., October 1994.

Appendix E - RFC 1550 White Papers

   [Adam94] Adamson, B., "Tactical Radio Frequency Communication
      Requirements for IPng", RFC 1677, NRL, August 1994.

   [Bello94a] Bellovin, S., "On Many Addresses per Host", RFC 1681, AT&T
      Bell Laboratories, August 1994.

   [Bello94b] Bellovin, S., "Security Concerns for IPng", RFC 1675, AT&T
      Bell Laboratories, August 1994.

   [Bound94] Bound, J., "IPng BSD Host Implementation Analysis", RFC
      1682, Digital Equipment Corporation, August 1994.

   [Brazd94] Brazdziunas, C., "IPng Support for ATM Services", RFC 1680,
      Bellcore, August 1994.

   [Britt94] Britton, E., and J. Tavs, "IPng Requirements of Large
      Corporate Networks", RFC 1678, IBM, August 1994.

   [Brownl94] Brownlee, J., "Accounting Requirements for IPng", RFC
      1672, University of Auckland, August 1994.

   [Carpen94a] Carpenter, B., "IPng White Paper on Transition and Other
      Considerations", RFC 1671, CERN, August 1994.

   [Chiappa94] Chiappa, N., "IPng Technical Requirements Of the Nimrod
      Routing and Addressing Architecture", RFC 1753, December 1994.

   [Clark94] Clark, R., Ammar, M., and K. Calvert, "Multiprotocol
      Interoperability In IPng", RFC 1683, Georgia Institute of
      Technology, August 1994.

   [Curran94] Curran, J., "Market Viability as a IPng Criteria", RFC
      1669, BBN, August 1994.

   [Estrin94a] Estrin, D., Li, T., and Y. Rekhter, "Unified Routing
      Requirements for IPng", RFC 1668, USC, cisco Systems, IBM, August
      1994.

   [Fleisch94] Fleischman, E., "A Large Corporate User's View of IPng",
      RFC 1687, Boeing Computer Services, August 1994.

   [Green94] Green, D., Irey, P., Marlow, D., and K. O'Donoghue, "HPN
      Working Group Input to the IPng Requirements Solicitation", RFC
      1679, NSWC-DD, August 1994.

   [Ghisel94] Ghiselli, A., Salomoni, D., and C. Vistoli, "INFN
      Requirements for an IPng", RFC 1676, INFN/CNAF, August 1994.

   [Heager94] Heagerty, D., "Input to IPng Engineering Considerations",
      RFC 1670, CERN, August 1994.

   [Simpson94] Simpson, W. "IPng Mobility Considerations", RFC 1688,
      Daydreamer, August 1994.

   [Skelton94] Skelton, R., "Electric Power Research Institute Comments
      on IPng", RFC 1673, EPRI, August 1994.

   [Syming94] Symington, S., Wood, D., and J. Pullen, "Modeling and
      Simulation Requirements for IPng", RFC 1667, MITRE, George Mason
      University, August 1994.

   [Taylor94] Taylor, M., "A Cellular Industry View of IPng", RFC 1674,
      CDPD Consortium, August 1994.

   [Vecchi94] Vecchi, M., "IPng Requirements: A Cable Television
      Industry Viewpoint", RFC 1686, Time Warner Cable, August 1994.

Appendix F - Additional References

   [Almqu92] Almquist, P., "Minutes of the Selection Criteria BOF",
      Washington DC IETF, November 1992, (ietf/nov92/select-minutes-
      92nov.txt).

   [Berkow94] Berkowitz, H., "IPng and Related Plug-and-Play Issues and
      Requirements", Work in Progress, September 1994.

   [Bos94] Bos, E. J., "Minutes of the Address Lifetime Expectations BOF
      (ALE)", Seattle IETF, March 1994, (ietf/ale/ale-minutes-
      94mar.txt).

   [Big] Archives of the big-internet mailing list, on munnari.oz.au in
      big-internet/list-archives.

   [Bradner93] Bradner, S., and A. Mankin, "IP: Next Generation (IPng)
      White Paper Solicitation", RFC 1550, Harvard University, NRL,
      December 1993.

   [Callon87] Callon, R., "A Proposal for a Next Generation Internet
      Protocol", Proposal to X3S3, December 1987.

   [Callon92a] Callon, R., "CNAT", Work in Progress.

   [Callon92b] Callon, R., "Simple CLNP", Work in Progress.

[Callon92c] Callon, R., "TCP and UDP with Bigger Addresses (TUBA), A
     Simple Proposal for Internet Addressing and Routing", RFC 1347,
     DEC, June 1992.

[Carpen93] Carpenter, B. and T. Dixon, "Minutes of the IPng Decision
     Process BOF (IPDECIDE)", /ietf/93jul/ipdecide-minutes-93jul.txt,
     August 1993.

[Carpen94b] Carpenter, B, and J. Bound, "Recommendations for OSI NSAP
     usage in IPv6", Work in Progress.

[Chiappa91] Chiappa, J., "A New IP Routing and Addressing
     Architecture", Work in Progress.

[Clark91] Clark, D., Chapin, L., Cerf, V., Braden, R., and R. Hobby,
     "Towards the Future Internet Architecture", RFC 1287, MIT, BBN,
     CNRI, ISI, UC Davis, December 1991.

[Deering92] Deering, S., "The Simple Internet Protocol", Big-Internet
     mailing list, 22 Sept. 1992.

[Deering94a] Deering, S., and P. Francis, Message to sipp mailing
     list, 31 May 1994.

[Estrin94b] Estrin, D., Zappala, D., Li, T., Rekhter, Y., and K.
     Varadhan, "Source Demand Routing: Packet Format and Forwarding
     Specification (Version 1)" Work in Progress.

[Fuller93] Fuller, V., Li, T., Yu, J., and K. Varadhan, "Classless
     Inter-Domain Routing (CIDR): an Address Assignment and Aggregation
     Strategy", RFC 1519, BARRNet, cisco Systems, MERIT, OARnet,
     September 1993.

[Gillig94c] Gilligan, B., "IPng Transition (ngtrans)", Work in
     Progress.

[Gross92} Gross, P., and P. Almquist, "IESG Deliberations on Routing
     and Addressing", RFC 1380, ANS, Stanford University, November
     1992.

[Gross94] Gross, P. "A Direction for IPng", RFC 1719, MCI, December
     1994.

[Hinden92a] Hinden, R., "New Scheme for Internet Routing and
     Addressing (ENCAPS)", Work in Progress.

[Hinden94b] Hinden, R., Deering, S., and P. Francis, "Simple Internet
     Protocol Plus", Working Group Charter, April 1994.

[Hinden92b] Hinden, R., and D. Crocker, "A Proposal for IP Address
    Encapsulation (IPAE): A Compatible version of IP with Large
    Addresses", Work in Progress.

[Huston94] Huston, G., and A. Bansal, draft charter for the
    "Transition and Coexistence Including Testing (TACIT) Working
    Group, June 1994.

[Huitema93] Huitema, C., "IAB Recommendations for an Intermediate
    Strategy to Address the Issue of Scaling", RFC 1481, INRIA, July
    1993.

[Huitema94] Huitema, C., "The H ratio for address assignment
    efficiency", RFC 1715, INRIA, October 1994.

[IAB92] Internet Architecture Board, "IP Version 7", Work in
    Progress.

[IAB94] Braden, R., Clark, D., Crocker, S., and C. Huitema, "Report
    of IAB Workshop on Security in the Internet Architecture -
    February 8-10, 1994", RFC 1636, USC/Informaiton Sciences
    Institute, MIT Laboratory for Computer Science, Trusted
    Information Systems, Inc., INRIA, IAB Chair, June 1994.

[Kasten92] Kastenholz, F, and C. Partridge, "IPv7 Technical
    Criteria", Work in Progress.

[Kasten94] Partridge, C., and F. Kastenholz, "Technical Criteria for
    Choosing IP: The Next Generation (IPng)", RFC 1726, BBN Systems
    and Technologies, FTP Software, December 1994.

[Knopper94a] Knopper, M., and P. Ford, "TCP/UDP Over CLNP-Addressed
    Networks (TUBA)", Working Group Charter, January 1994.

[Knopper94b] Knopper, M., and D. Piscitello, "Minutes of the BigTen
    IPng Retreat, May 19 & 20 1994".

[Leiner93] Leiner, B., and Y. Rekhter, "The MultiProtocol Internet",
    RFC 1560, USRA, IBM, December 1993.

[Mankin94] Mankin, A., and S. Bradner, message to big-internet, tuba,
    sipp, catnip and ietf mailing lists, 7 July 1994.

[Mills84] Mills, D. "Exterior Gateway Protocol Formal Specification",
    RFC 904, UDEL, April 1984.

[Mogul90] Mogul, J., and S. Deering, "Path MTU Discovery", RFC 1191,
    DECWRL, Stanford University, November 1990.

   [Nat94] National Research Council, "Realizing the Information Future:
       The Internet and Beyond", National Academy Press, 1994.

   [Piscit94] Piscitello, D., "FTP Operation Over Big Address Records
       (FOOBAR)", RFC 1639, Core Competence, June 1994.

   [Provan91] Provan, D., "Transmitting IP Traffic over ARCNET
       Networks", RFC 1051, Novell, February 1991.

   [Postel94] Postel, J., Editor, "Internet Official Protocol
       Standards", RFC 1720, USC/Information Sciences Institute, November
       1994.

   [Solens93a] Solensky, F., and T. Li, "Charter for the Address
       Lifetime Expectations Working Group", FTP Software, Cisco Systems,
       November 1993.

   [Solens93b] Solensky, F., "Minutes of the Address Lifetime
       Expectations BOF (ALE)", Houston IETF, November 1993,
       (ietf/ale/ale-minutes-93nov.txt).

   [Solens94] Solensky, F., "Minutes of the Address Lifetime
       Expectations BOF (ALE)", Toronto IETF, July 1994, (ietf/ale/ale-
       minutes-94jul.txt).

   [Sukonnik94] Sukonnik, V., "Common Architecture for Next-Generation
       IP (catnip), Working Group Charter, April 1994.

   [Tsuchiya92] Tsuchiya, P., "The 'P' Internet Protocol", Work in
       Progress.

   [Ullmann93] Ullmann, R., "TP/IX: The Next Internet", RFC 1475,
       Process Software Corporation, June 1993.

Appendix G - Acknowledgments