

IP: The Next Generation

Written by Scott Phillips.

This page is a general overview of the IPng protocol.

Table of Contents

1. [What is IPng?](#) A quick overview of what IPng is and why it is being developed.
2. [IPng goals and development](#) Some background on the issues involved in the IPng design.
 1. [Problems with IPv4.](#)
 2. [Transition from IPv4 to IPng.](#)
 1. [Simple Internet Transition \(SIT\).](#)
3. [Specification](#) Click here to see an overview of the IPng specification.
 1. [Addressing](#)
 2. [Routing](#)
 3. [Headers and Options](#)
 1. [IPng header format](#)
 4. [Flow Control](#)
 5. [Security](#)
4. [Who is doing it?](#) Click here to see a list of the folks that are actually doing the work.
5. [References](#) See the words of those who have gone before. . .

1. What is IPng?

IPng stands for Internet Protocol: The Next Generation. Its official name is IPv6, and it is intended to replace the IP that is currently used in the Internet today (version 4, or IPv4).

Over the past few years, as the Internet has grown at an exponential rate, the deficiencies of IP have quickly become apparent. The number of allowable addresses in IPv4 isn't in line with the vast number of nodes connecting to the Internet and there simply aren't enough addresses to last. Although valiant attempts at prolonging the life of IP are currently under way, they are only delaying the inevitable. A more drastic solution is required.

IPng was designed to correct all of IPv4's deficiencies and to implement some new functionality as well. Please examine the presented IPng specification for more details.

[Back to the Table of Contents](#)

2. IPng: The Design Issues

The major issue is to solve the [problems with IPv4](#), such as scalability and routing.

For this reason, IPng supports large hierarchical addresses. Look in the specification section for more detailed information on IPng addressing.

How do we get there from here?

The [transition from IPv4 to IPv6](#) is obviously a big deal. IPv6 will have to be gradually phased in and, as a result, will have to completely interact with IPv4.

As if that weren't enough. . .

They've included support for real-time flows, provider selection, host mobility, end-to-end security, auto-configuration and auto-reconfiguration.

How does it handle?

IPng is designed to run on high performance networks such as ATM as well as in low bandwidth applications such as wireless communication. Its headers are less expensive to process, and the 128 bit address was chosen to match the new generation of 64 bit processors.

[Back to the Table of Contents](#)

2.1 IPv4 Problems

The number of users on the Internet is growing at a ridiculously fast rate and hence, the number of available IPv4 addresses is quickly dwindling. Also, as the world becomes more and more involved in networking, new technologies and applications that are currently very difficult, if not impossible to support must be enabled.

The major problem:

- In IP you have only four classes of nets: Class A, B, C and D (a fifth class, class E, is only for research purposes). These classes differ in the number of nets and hosts:
 - Class A 125 nets with 16 million hosts per net
 - Class B 16382 nets with 65534 hosts per net
 - Class C 2 million nets with 254 hosts per net
 - Class D multicast network class

Unsupported features:

- Provider selection.
You can choose special providers for routing the packet. This is necessary for commercial usage of the internet, making it possible to choose only special providers i.e. trusted providers.
- Scalable [multicast](#).
Multicast in IP is only possible in subnets. For Multimedia-applications it should be possible to address different hosts in different subnets.
- Mobility. "plug-and-play"
Mobility means that you can plug in a host to the net with no need to configure it by yourself and others can address it and reach it. This would mean that your host would have to get an address from a provider.
- Real-time flow.
That feature is important for real time services like video conferences etc...

[Back to the Table of Contents](#)

2.2 Transition from IPv4 to IPv6

Getting there from here. . .

The problem is how to convert the Internet into IPv6, without disrupting the operation of the existing IPv4 network. The transition is planned to be processed in two phases. At the end of phase 1, there will be both IPv4 and IPv6 hosts and routers. At the end of phase 2, there will only be IPv6 hosts and routers. This means that the SIT (**S**imple **I**IPv6 **T**ransition) should at least ensure the following:

1. IPv6 and IPv4 hosts can interoperate
2. IPv6 routers and hosts can be deployed in the Internet in a highly diffuse and incremental fashion, with few interdependencies
3. The transition should be as easy as possible for end-users, system administrators, and network operators to understand and carry out.

The SIT provides a number of features, including:

- 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.
- Minimal upgrade dependencies.
The only prerequisite to upgrading hosts to IPv6 is that the DNS server must first be upgraded to handle IPv6 address records. There are no prerequisites to upgrading routers.

- 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.

The following mechanisms are employed in SIT to realize the above features:

- Use of the dual IP layer (IPv4 and IPv6) technique in hosts and routers for direct interoperability with nodes implementing both protocols.
- Two IPv6 addressing structures that embed an IPv4 addresses within IPv6 addresses.
- A mechanism for tunneling IPv6 packets over IPv4 routing infrastructures. This technique uses the embedded IPv4 address structure, which eliminates the need for tunnel configuration in most cases.
- An optional mechanism for translating headers of IPv4 packets into IPv6, and the headers of IPv6 packets into IPv4. This technique allows nodes that implement only IPv6 to interoperate with nodes that implement only IPv4.

[Back to the Table of Contents](#)

2.2.1 Details on SIT

Types of Hosts and Routers

To understand the Transition Model, it is necessary to know the various kinds of hosts and routers. In the model there exists 4 types:

1. IPv4-only-nodes
These are host and routers that only understand IPv4.
2. IPv6/IPv4-nodes
The routers and hosts of this category have both the IPv4 and the IPv6 protocol stacks. In addition to that they have mechanisms such as IPv6-over-IPv4 tunneling. These nodes can directly interoperate with both IPv4 and IPv6 nodes, but for communication with IPv4-only-nodes they have to be configured with an IPv4-compatible IPv6 address.
3. IPv6-only-nodes
That are hosts and routers that only understand IPv6.
4. IPv6/IPv4-header-translating-router
These routers translate IPv6 packets into IPv4 packets and vice-versa.

IPv6-over-IPv4 Tunneling

Tunneling is used to carry IPv6 packets across IPv4 routed network areas. One of the requirements for tunneling is that the begin and endpoints of the tunnel are IPv6/IPv4-nodes with IPv4-compatible IPv6 addresses.

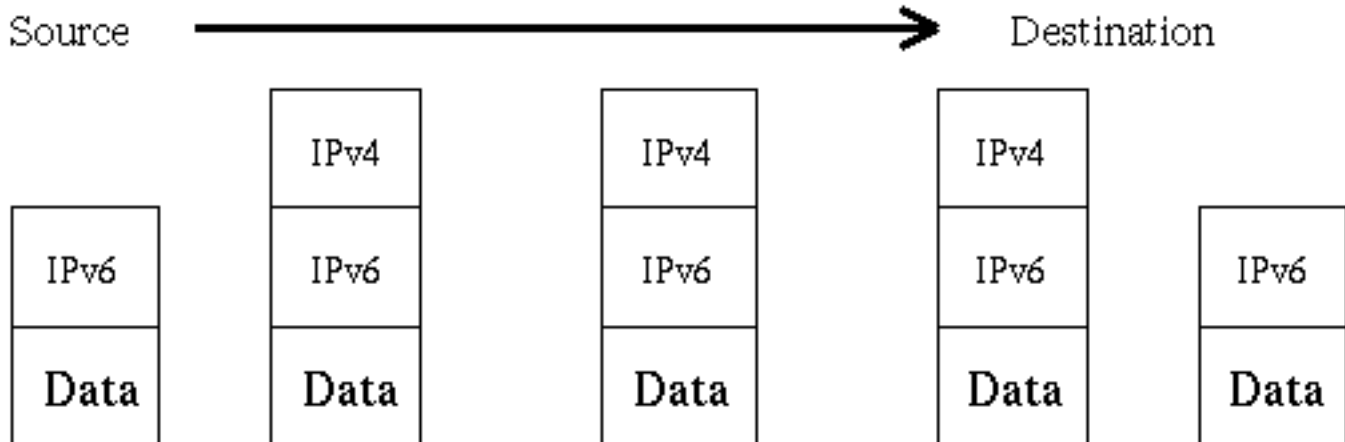

Tunneling means that the whole IPv6 packet is mapped into a body of an IPv4 packet and sent across the IPv4 network area. The endpoint of the tunnel has to be either a IPv6/IPv4-header-translating-router or a IPv6/IPv4-node to de-encapsulate the packet. The destination address of the new IPv4 packet is the address of the node representing the tunnel endpoint.

There are two types of tunneling: automatic tunneling and configured tunneling.

Automatic Tunneling

Automatic tunneling is used between two IPv6/IPv4-hosts. It is "end-to-end". It can also be used if a router is going to send an IPv6 packet to an IPv6/IPv4-host that is connected to the same IPv4 network area. It is important that the endpoint of the tunnel is the destination host.

Net-Structure

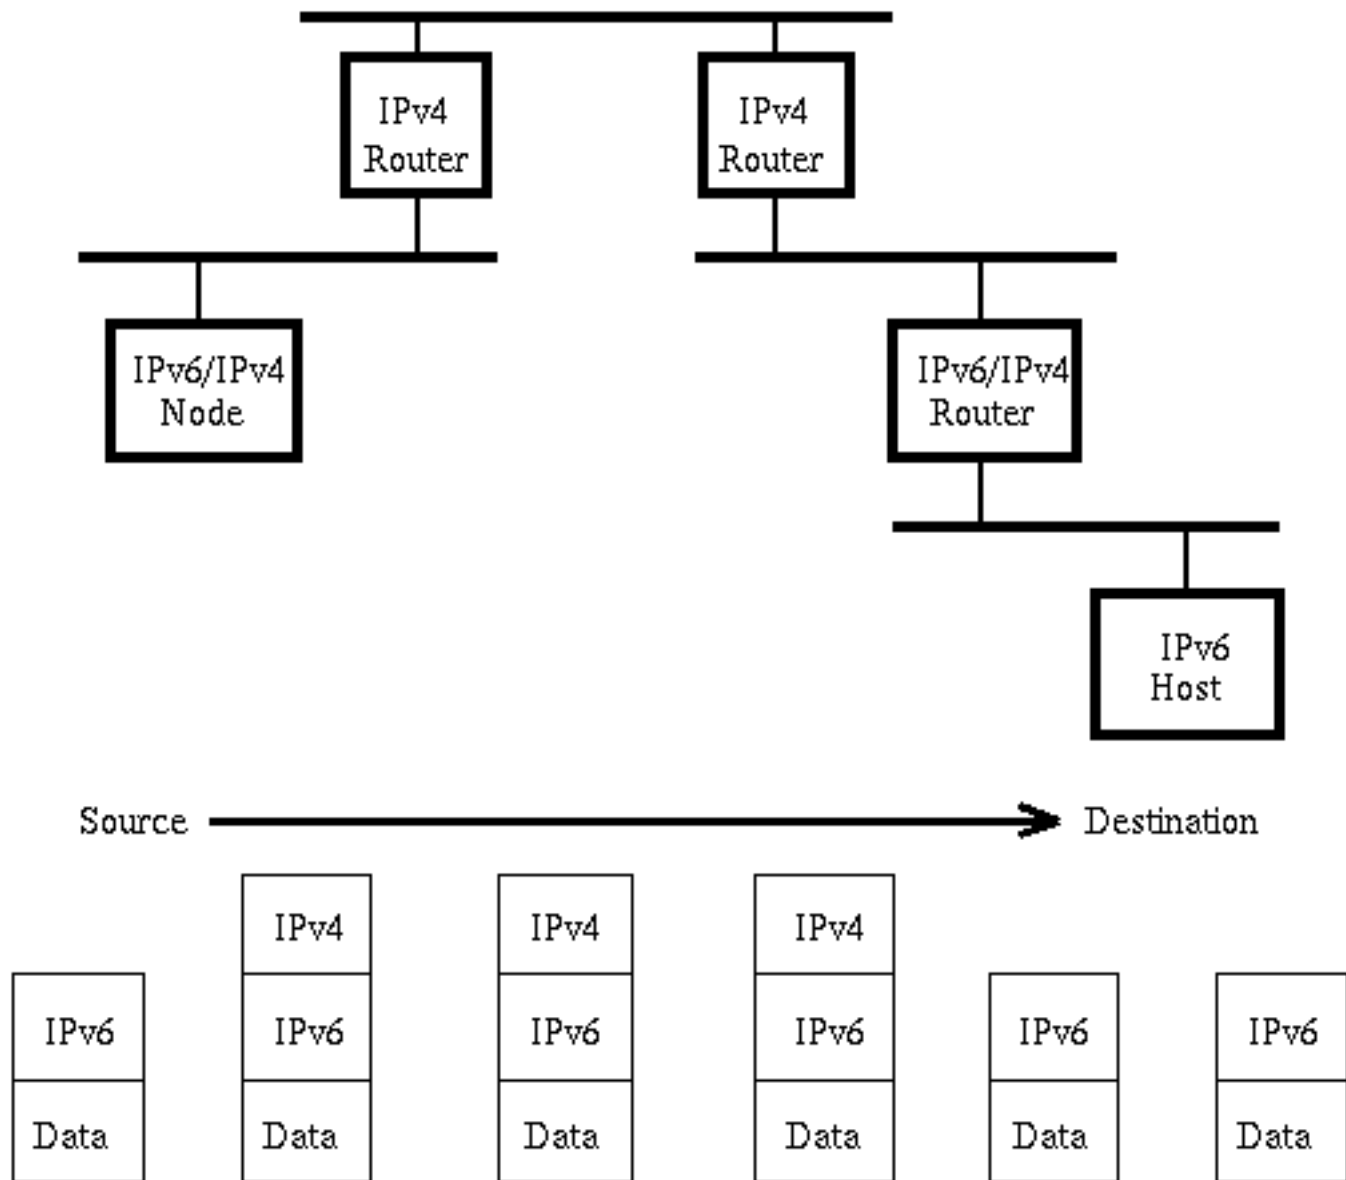


Packet-Structure

Configured Tunneling

Configured tunneling is used if the destination host is different from the endpoint of the tunnel. In this case, the destination address for the IPv4 header, ie the address of the endpoint of the tunnel, could not be simply mapped from the IPv6 destination address. The endpoint of the tunnel has to be configured in the IPv6/IPv4-node.

Net-Structure



Packet-Structure

[Back to the Table of Contents](#)

3. The Specification

So what's new?

Expanded Routing and Addressing

The IPng address size increases from 32 bits to 128 bits. This allows for more levels of hierarchy, a greater number of addressable nodes, and easier auto-configuration.

Multicast addresses get an added "scope" field to enhance routing scalability.

A Cluster Address is defined to identify topological regions rather than individual nodes. When used in source addresses, cluster addressing allows nodes to control the path their traffic will take.

Header Format Simplification

Some redundant IPv4 header fields have been dropped or made optional. This reduces the processing cost of packet handling and header bandwidth. Although IPng addresses are four times larger than those in IPv4, the header is only two times larger.

Improved Option support

Option lengths have less stringent limits and there is much greater flexibility for the introduction of new options in the future.

Quality of Service

Packets are allowed to be labeled as part of a flow. This enables real-time service.

[Back to the Table of Contents](#)

3.1 IPng Addressing

IPng addresses are 128 bits long. They can identify individual nodes or sets of nodes. Three types of IPng addresses exist, namely:

- [unicast](#) - a single node.
- [cluster](#) - a group of nodes that share a common prefix such that a packet sent to the address will be delivered to one of the nodes.
- [multicast](#) - a group of such that a packet sent to the address will be delivered to all nodes in the group.

IPng has 665,570,793,348,866,943,898,599 addresses per square meter of the Earth. Assuming the most pessimistic hierarchical division possible it is estimated that there is still 1564 addresses per square meter.

The address space is divided into NSAP, IPX, Provider based unicast, Geographic, Local use and multicast addresses. These take up 15% of the total. The rest is reserved for future use.

Unicast

There are several forms of unicast in IPng, these are global provider hierarchical, geographical hierarchical, NSAP hierarchical, IPX hierarchical, local use addresses, IP-only host addresses. Additional ones can be defined in the future.


Provider Based Unicast Addresses



Provider based unicast address

These are used for global communication. The first 3 bits identify it as of this type. A provider ID is supplied to the providers, who can assign parts of it to it's subscribers.


Local Use Unicast Addresses



Local use address

A local address is one that only has scope within its own subnet. It may have local or global uniqueness. They are intended as "plug and play" addresses for bootstrapping to a fixed address. They are not yet connected to the global internet, without the need to request an address prefix.


IPv4 only Unicast Addresses



IPv4 only address

These are assigned to IP only hosts as part of the transition scheme.


Cluster Addresses



The cluster address allows a node to select which of several providers that it wants to use. A cluster address can only be used as a destination.

[Back to the Table of Contents](#)

Multicasting with IPng



IPng multicast is an identifier for a group of nodes. A node can belong to any number of multicast groups.

FLGS is set to 000X, where the 000 is reserved, and

- X=0 indicates a permanently assigned or well known address, assigned by a naming authority.
- X=1 indicates a transient address.

SCOP indicates scope, for which the values are

- 0 reserved
- 1 intra-node scope
- 2 intra-link scope
- 3 (unassigned)
- 4 (unassigned)
- 5 intra-site scope
- 6 (unassigned)
- 7 (unassigned)
- 8 intra-organisation scope
- 9 (unassigned)

- A (unassigned)
- B intra community scope
- C (unassigned)
- D (unassigned)
- E global scope
- F reserved

[Back to the Table of Contents](#)

3.2 Routing

Every node on the Internet has a unique IP address. In order to communicate with other nodes, each must keep a routing table with information about other nodes on the network and how to get to them. As the size of the Internet balloons at such an alarming rate, these tables can become exceptionally large and hence, very memory intensive and inefficient.

Since the internet is so large, the new scheme must be compatible with the old in order for it to be feasible. For this reason, IPng routing is nearly identical to IPv4 routing. A simple extension to the IPv4 routing algorithms is all that is required for them to work with IPv6.

Differences

- 128-bit addresses instead of IPv4's 32-bit addresses. This reduces the size of routing tables by adding hierarchical levels and eliminates the current shortage of IP addresses.
- Routing extensions
 - Provider selection: This option allows a host to specify the route (intermediate nodes) used to reach a destination. It may be determined by security (policy), performance, cost, or all of the above.
 - Host-mobility: This is also called plug-and-play. This will allow the addition of a host to the network with no manual configuration. When a station is added to the network, its IP address can be automatically assigned and all of the appropriate tables updated accordingly.
 - Auto re-addressing: A destination host can reply to the source by merely reversing the address sequence, thereby eliminating the routing process.

[Back to the Table of Contents](#)

3.3 IPng Headers and Options

Headers

In IPng, the [headers](#) have been greatly simplified with regard to IPv4. Many of the fields have been dropped, or made optional. The reason for this was to make the cost of processing packets as low as possible, despite a greatly increased address size. Even though the addresses in IPng are four times


longer than in IP , the headers are only twice as big as the IPv4 header.

Options

- Routing
- Fragmentation
- Authentication
- Security Encapsulation
- End to End options

[Back to the Table of Contents](#)

3.3.1 IPng Header Format



IPng Header

Version

4-bit Internet Protocol version number. This is 6 for IPng.

Flow Label

28 bits. See [quality of service](#).

Payload Length

16-bit field that measures the length of the rest of the packet following the header, measured in octets.

Next Header.

This uses the same values as the IP Protocol field. It specifies the type of the header immediately following the header, such as TCP or UDP. It has been renamed (renamed from SIP which was renamed from IPv4) to avoid confusion as to what is being referred to as the IP protocol - the protocol field or IP itself.

Hop Limit

The Hop Limit is set to some nonzero value, and decremented by one by each system that forwards the packet. The packet is discarded if the hop limit reaches zero. This is to prevent the packet getting stuck in a forwarding loop.

Other uses include limiting the propagation of multicast packets, and it can also be used for diagnostic purposes. The "time to live" field in IP provided the same function, plus one extra one. This was to limit the amount of time that a packet spent in transit. This was discarded because it proved too costly to implement, and in some cases impossible to implement, for example in large subnets whose transit time is unpredictable. In practice many IP routers implemented time to live as hop limit, SIP legitimised this. Any higher level functions that cannot tolerate delivery delays, must provide their own method of recognising old packets.

Source Address

This is the 128 bit address of the sender.


Destination Address

This is the 128 bit address of the initial destination. This may not be the ultimate destination, if an optional routing header is used.

[Back to the Table of Contents](#)

3.4 Flow Control

Flow control is provided for applications which require consistent throughput, delay, and/or jitter.



The Flow Label Field in the IPv6 header may be used by a host to label packets that require special handling by routers, such as non-default quality of service or "real-time" service.

The nature of this handling might be conveyed to the routers by either a control protocol (such as RRP, resource reservation protocol) or by information within the packets themselves.

A flow is identified by a source address and a non-zero FLOW ID. (Packets that do not belong to a flow carry a FLOW ID of 0.)

FLOW ID

- Assigned to a flow by the source node
- Randomly chosen from 0x1 to 0xFFFFFFFF to make any set of bits within the FLOW ID suitable for use as a hash key by the routers (for looking up the special handling state associated with the flow)
- Must not be used again for a new flow while any state associated with the previous usage still exists in any router

TCLASS

- Used to identify the desired delivery priority of its packets, relative to other packets from the same source
- Two ranges of values:
 - 0 to 7:
 - used to label **flow-controlled packets**, e.g. packets that belong to a TCP connection.
 - Particular applications:
 - filler traffic (netnews...), interactive traffic (telnet, X...), internet control traffic (SNMP, routing protocol ...)
 - 8 to 15:
 - used to label **non-flow-controlled packets**, e.g. "real-time" packets being sent without any flow-control feedback from the receivers.
 - Particular applications:
 - for those packets that the sender is most/least willing to have discarded under conditions of congestion (e.g. high/low fidelity video traffic),

[Back to the Table of Contents](#)

3.5 Security

Below the application layer, IPv4 lacks privacy and authentication methods. Today's IPv4 has one or two problems with security. IPng offers to cure this by providing two integrated security options. These two options can be used separately or in conjunction with each other depending on the user's needs.

Two Security Options

IPng Authentication Header

This option will provide authentication and integrity but no confidentiality. The option will be algorithm-independent and will support various authentication techniques. The purpose for providing all this without the confidentiality is that this mechanism needs to be exportable by vendors in countries that restrict the export of confidentiality algorithms, such as the United States, for example.

In order to help ensure interoperability within the Internet, the use of keyed MD5 has been proposed. This will also eliminate a number of network attacks, (including host masquerading attacks).

This internet layer protection will provide the upper layers with the host origin authentication that they currently lack.

Encapsulating Security Header

This option will provide the integrity and confidentiality missing from the IPng Authentication Header option. It is both flexible and algorithm-independent.

The DES algorithm has been proposed as the standard, again with the aim of achieving interoperability within the worldwide Internet. This mechanism, however, probably won't be as exportable as the Authentication Header, but the use of DES as a standard should help.

[Back to the Table of Contents](#)

5. References and Pointers to More Information

R. Hinden, IP Next Generation Overview, URL, <http://playground.sun.com/pub/ipng/html/INET-IPng-Paper.html>, May 1995.

J. Bound, [Dynamic Host Configuration Protocol for IPv6](#), Internet Draft, draft-ietf-dhc-dhcpv6-14.txt, February 1999.

S. Bradner, A. Mankin, [The Recommendation for the IP Next Generation Protocol](#), RFC 1752, January 1995.

Internet Engineering Steering Group, Protocol Action: The Recommendation for the IP Next Generation Protocol to Proposed Standard, November 19, 1994.

Last modified 8/24/95.

[Other reports on Recent Advances in Networking 1995](#)

[Back to Raj Jain's Home Page](#)