

IPv6: The Essentials You Must Know¹

Ibrahim Haddad

Introduction

IPv6 is the next generation protocol designed by IETF to replace the current version of the Internet Protocol, IPv4. Most of today's Internet uses IPv4, which has been remarkably resilient in spite of its age, but it is beginning to have problems. Most importantly, there is a growing shortage of IPv4 addresses, which are needed by all new devices connecting to the Internet. As a result, IETF defined IPv6 to fix the problems in IPv4 and to add many improvements to cater for the future Internet. These improvements come in different areas such as routing, autoconfiguration, security, QoS, and mobility.

In this article, we address the problems in IPv4 that lead to the design of IPv6 and present the capabilities of IPv6 that have been developed in direct response to critical business requirements for scalable network architectures. We also briefly touch on IPv6 implementations and the migration from IPv4 to IPv6.

The article intends to provide readers with a global view over IPv6 and will be followed up by technical articles that will guide readers on how to support IPv6 on their Linux servers and provide basic services over IPv6.

“32 bits should be enough address space for the Internet”

– Dr. Vint Cerf

In 1977, Dr. Vint Cerf – Senior Vice President of Internet Architecture and Engineering at MCI WorldCom, and an Internet pioneer credited as being the father of the Internet – confidently asserted that “32 bits should be enough address space for the Internet”. Today, as honorary chairman of the IPv6 Forum, Dr. Cerf’s strong voice calls for the immediate adoption of IPv6 in order to “take the Internet where no other network has gone before”.

IPv4: An Aging Protocol

IPv4 has proven to be robust, easily implemented and interoperable, and has stood the test of scaling to be the size of today's Internet using different mechanisms such as NAT. However, The initial design of IPv4 did not take into consideration several issues that are of importance today, such as a large address space, mobility, security, autoconfiguration, and QoS.

¹ To be published in Linux User & Developer – Last updated 2003-04-28 9:15 AM

In this section, we present the shortcomings of IPv4, and in the following section, we examine how IPv6 solved these problems and added new features to help us design and build the future Internet.

IPv4 Address Space Limitations

IPv4 has a 32-bit address field, which in theory can support as many as four billion unique addresses. However, the actual allocation² of space has locked up nearly 75% of these addresses. Consequently, any organization (schools, university, carriers, ISPs, ASPs, etc.) applying today for IP addresses will be assigned with a fraction of the remaining Class-C addresses.

Inefficient Routing

The explosive growth of the Internet has affected the ability of the Internet backbone routers to maintain routing tables and provide fast routing. This is due to different reasons, such as:

- The design of the IPv4 address structure.³
- The fact that new networks within the same organization are assigned new prefixes because the ones previously assigned to the same organization were all used. Therefore, routers in the DFZ need to have several entries for the same corporate or ISP network.

As a result, routing was on the list of features to be improved in IPv6.

Security

Long considered an issue to be addressed at the higher network layers, security has emerged as an area where the next version of IP could provide some useful functions. One shortcoming of IPv4 packets is that it is not natively secure. This has been somehow corrected by RFC 2401, Security Architecture for the Internet Protocol, but the method remains cumbersome (IPsec will not work over NAT), as this was not an original design feature.

Lack of Autoconfiguration

Configuring IPv4-nodes has always been complex. Network administrators as well as users would prefer to be able to plug a computer into the network and start using it. This issue has been addressed in IPv6 (explained in a later section).

Lack of Flow Label

A routing issue that is lacking in IPv4 and being addressed by IPv6 is the lack of flow control. IPv4 is a connectionless, best effort network. There is no guarantee when and if the packet will arrive at its destination. IPv6 provides flow labels that can be used to provide QoS. The flow labels, identifying the packets as belonging to a flow, can be used in conjunction with a hop-by-hop routing extension header (allowing predefined routes) and the priority field (allowing for QoS).

² In the early days, a limited number of entities were each allocated Class-A IP addresses. Today, each controls more than 16 million addresses.

³ The way it is divided into a network and a host portion.

Performance

Although IP performs remarkably well, some of the design decisions made 20 and more years ago in retrospect could stand improvement. IPv6 improved the network performance by introducing changes into the address structure with the hierarchical address-scheme, changing the header format, and introducing extension headers (explained more in a later section).

Poor Mobility Support

IPv4 has some difficulties managing mobile computers or nodes. Illustrated are some case scenarios:

1. Mobile computers need to use a forwarding address at each new point of attachment to the Internet. With IPv4, getting this address is not always easy.
2. Good authentication facilities, which are not commonly deployed in IPv4 nodes, are required to inform any agent in the routing infrastructure about the new location of the mobile node.
3. It can be difficult for mobile nodes to determine whether they are attached to the same network or not.
4. Mobile nodes usually cannot inform their communication partners about a change in location.

A redesign for mobility support was needed and it is reflected with the mobility schemes supported in IPv6.

IPv4: Fix or Replace?

Given all the problems we were facing with IPv4, the Internet community had to take action. There were two options available:

1. Fix IPv4 and risk the continuation of the degradation of the Internet model:

Taking this option would lead to more complex and volatile network services, lower performance, less robust, less secure, and less manageable networks. In addition, it will add a significant barrier to innovations in new applications and services with all the limitations imposed by the protocol.

2. Replace IPv4 and restore the Internet model with a new Internet protocol:

Following this path would lead to simple and stable network services, higher performance, and more robust, secure, and manageable networks. It will also enable anyone to provide new innovative applications and services allowing rapid innovation and growth.

The choice was clear, and the main concern was to avoid doing small changes now and then repeat the same exercise a few years later on. Therefore, the decision was to proceed forward and design a new protocol.

From IPv4 to IPv6 – What about IPv5?

IPv5 (defined in RFC 1190) was an experimental connection-oriented alternative protocol to IPv4 for resource reservation intended to provide Quality of Service. However, before discussions started to go too far, IPv6 standards emerged and started to be implemented, and other protocols provided the proposed functionality of IPv5.

IPv6 Philosophies

The rapid increase in the number of Internet users, combined with the expected growth in the number of wireless Internet devices, and new applications, require a scalable and flexible IP technology, which is not provided by IPv4. In response, IETF produced a set of specifications that defines IPv6. IPv6 architecture and design include a number of attractive features that make it a very suitable component for IP-based next generation networks.

IPv6 was designed, following two main philosophies:

- **Design Philosophy:** IPv6 must be scalable, designed with a futuristic mind to provide a large address space with a simple structure, an original end-to-end environment, a NAT free network, fast processing, and many features needed by current and future applications (security, QoS, mobility, etc.).
- **Operation Philosophy:** The process of migration from IPv4 to IPv6 and IPv6 deployment should not be expensive. IPv6 should inter-operate with IPv4 and provide tools and mechanisms needed by hosts running different IP versions to communicate with each other and to enable applications to work with both IP versions.

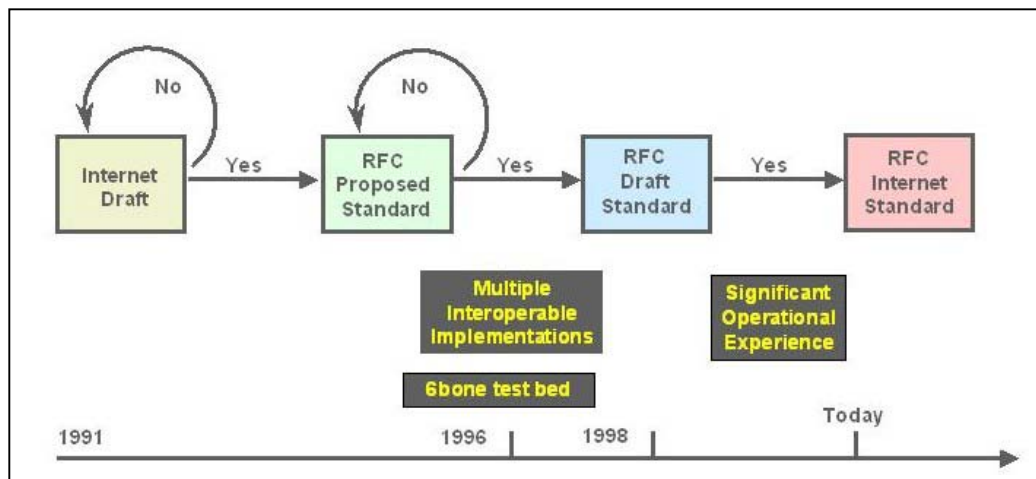


Figure 1: IPv6 - A Decade of Design and Testing

Let There be ... IPv6!

IPv6 has been designed to enable high-performance, scalable networks to remain viable well into the next century and to avoid running into similar problems in the future. A large part of this design process involved correcting the inadequacies of IPv4. Some of the qualities of IPv6 are found in obviously enhanced features, such as the larger address space and streamlined packet design. Other qualities are less tangible and relate to the fresh start that IPv6 gives to those who build and administer networks.

IPv6 represents a big package of capabilities of which addressing is the most visible component.

The addressing issue gets a lot of attention, however, it is only one of many important issues that IPv6 designers have tackled. Other IPv6 capabilities include improved security and data integrity, integrated QoS, automatic configuration, mobile computing, data multicasting, and more efficient network route aggregation at the global backbone level.

The following sub-sections provide an overview of the improvements and features that IPv6 brings to networking and the Internet.

128 Bit Addresses

IPv6 solves the lack of addresses with its 128-bit address field, providing a big improvement from IPv4's 32-bit addresses. This extended address space is very essential, as IP addresses will be assigned to mobile phones, home appliances, motor vehicles and other equipment. In addition, with such a huge address space, we can create multi-level hierarchies of addresses, which will simplify the problem of routing – simpler routing algorithms and less space needed for routing tables.

New Types of Addresses

IPv6 defines three types of addresses: unicast (global, link local, site local), multicast, and anycast.

- **IPv6 unicast address** identifies a single interface. A packet sent to a unicast address is delivered to the interface identified by that address.
 - *Global unicast address*, used for point-to-point communication.
 - *Link local unicast address*, used to let packets traverse on only one link or segment. Routers will not forward packets with link local unicast addresses.
 - *Site local unicast address*, used to limit the packet delivery scope to your intranet. The edge router connecting your internal network to the external network will not forward packets with site local unicast addresses to the external network.
- **IPv6 multicast address** delivers copies of one source packet to recipients. In the IPv6 multicast address, you can specify multicast scope, which can be node-local, link-local, site-local, or global. In IPv6, multicasting to all nodes in your organization replaces the broadcasting capability in IPv4.
- **IPv6 anycast address** identifies a set of interfaces typically belonging to different nodes. A packet sent to an anycast address is delivered to one of the interfaces identified by that address (the nearest one, according to the routing protocol's measure of distance). Anycast differs from multicast in that it delivers a message to any one of the nodes in a group. When one node, often the nearest node in the group, receives the message, anycast is finished. You can group routers in an anycast group, and a host can send a query to the anycast group to find the nearest router.

Addresses Autoconfiguration

TCP/IP designers have recognized the difficulty of installing and configuring TCP/IP networks and have tried over the years to come up with solutions that would overcome these pitfalls. One of IPv6's useful features is its ability to automatically configure itself without the use of a stateful configuration protocol, such as DHCPv6. By using router discovery, a host can also determine the addresses of routers, additional addresses, and other configuration parameters. Included in the router advertisement message is an indication of whether a stateful address configuration protocol should be used.

There are two types of autoconfiguration:

- **Stateless:** Configuration of addresses is based on the receipt of router advertisement messages. These messages include stateless address prefixes and require that hosts not use a stateful address configuration protocol.
- **Stateful:** Configuration is based on the use of a stateful address configuration protocol, such as DHCPv6, to obtain addresses and other configuration options. A host uses stateful address configuration when it receives router advertisement messages that do not include address prefixes and require that the host use a stateful address configuration protocol. A host will also use a stateful address configuration protocol when there are no routers present on the local link.

By default, an IPv6 host can configure a link-local address for each interface.

The main idea with autoconfiguration is that IPv6 has significant features that enhance the ability of a host to configure itself. The aim of the designers with this feature was that a host should be able to discover automatically all the information it needs to connect to the Internet, without any human intervention.

Streamlined Header Format

The IP headers were modified with IPv6 by allowing headers to be chained together – optimized for efficient processing. IPv6 packet headers contain many of the fields found in IPv4 packet headers; some of these fields have been modified from IPv4. The 40-byte IPv6 header consists of the following eight fields:

- Version (4 bits): Identifies the version of the Internet Protocol.
- Traffic Class (8 bits): Identifies different classes or priorities.
- Flow label (20 bits): Used by a source node to identify packets that belong to the same flow (source address + flow label = uniquely identify flows).
- Payload length (16 bits): Length of the IPv6 payload.
- Next header (8 bits): Indicates the encapsulated protocol – the next extension header to examine.
- Hop limit: Indicates the maximum number of hops allowed.
- Source address (128 bits): Address of the source node sending the packet.
- Destination address (128 bits): Final destination node address for the packet.



Figure 2: IPv6 Header Format

In addition, IPv6 is much more flexible in its support of options through extension headers. Extension headers are used to encode optional Internet-layer information. They are placed between the IPv6 header and the upper layer header in a packet. Extension headers are chained together using the next header field in the IPv6 header. There are six different extension headers: Hop-by-hop Options header, Destination Options header, Routing header, Fragment header, Authentication header, and Encapsulated Security header.

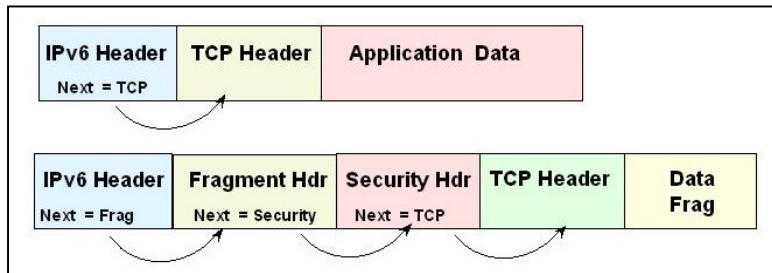


Figure 3: IPv6 Extension Headers

The next header field indicates to the router which extension header to expect next. If there are no more extension headers, the next header field indicates the upper layer header (TCP header, UDP header, ICMPv6 header, an encapsulated IP packet, or other items).

In addition, header compression in IPv6 improves interactive response time, allows using small packets for bulk data with good line efficiency, allows using small packets for delay sensitive low data-rate traffic, decreases header overhead, and reduces packet loss rate over lousy links.

Network renumbering

One new interesting feature of IPv6 is the renumbering of a network, which makes it simpler, for instance, to move a whole network to a new ISP by reconfiguring the router with the new routing prefix from the new ISP. The new ISP will then propagate its prefix from its router to the customer routers, who in turn will advertise the prefix to all hosts in the network. The new addresses will replace the old addresses when the hosts receive the new address prefix through the router advertisements.

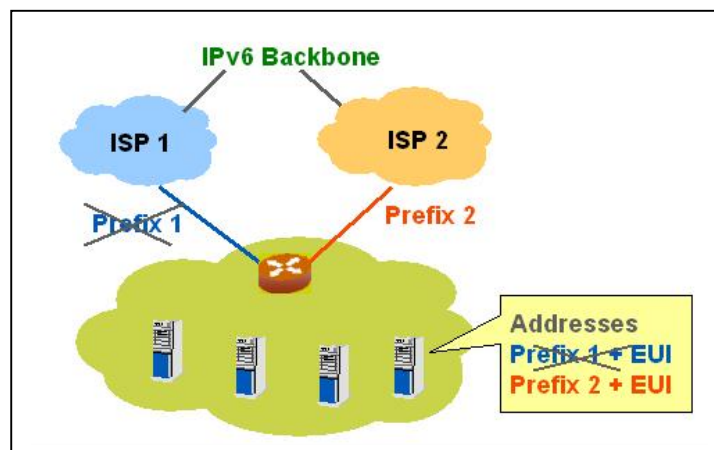


Figure 4: Network renumbering

Mobility Support

Each mobile node (such as mobile phones, PDAs, etc.) is always identified by its home address, regardless of its current point of attachment to the Internet. While situated away from its home, a mobile node is also associated with a care-of address, which provides information about the mobile node's current location. IPv6 packets addressed to a mobile node's home address are transparently routed to its care-of address.

The protocol enables IPv6 nodes to cache the binding of a mobile node's home address with its care-of address, and to then send packets destined for the mobile node directly to it at this care-of address. The mobility advantage of IPv6 can be further emphasized by the addition of flow-label management, which gives mobile nodes an even better quality of service.

Mandated Support for IPsec

Security has always been an important topic on the Internet. To increase Internet security, the IETF has specified IP-layer security (IPsec) in parallel with IPv6 and mandated its support with IPv6 so it will not be an optional extension as was the case with IPv4. As a result, IPsec support is a standard feature with IPv6.

QoS and the Flow Label Field

Today, the Internet works on a best efforts basis and QoS is mostly outside the users' control. Commercial users of the Internet are increasingly demanding guaranteed levels of service for all of their traffic or for selected parts of their traffic. Without improved and selectable quality, services such as Internet telephony and video conferencing, that demand minimum service levels, will not be feasible.

IETF specified two approaches (integrated services and differentiated services) to provide guaranteed and selectable QoS over the Internet. In addition, IPv6 provides flow labels that can be used to provide QoS. The flow labels, identifying the packets as belonging to a flow, can be used in conjunction with a hop-by-hop routing extension header (allowing predefined routes) and the priority field (allowing for QoS). The flow label also serves as a key in the router cache to reduce the amount of processing. When a datagram comes to a router the first time, it can save the flow label in the cache so that the next time a datagram arrives from the same flow (with the same flow label) the router will recognize the flow label in its cache table and finds the next hop without having to look in the routing table. This technique reduces the processing time in the router considerably.

As a result, IPv6 will make it easier to build and deploy applications requiring particular or selectable QoS over the Internet.

Better Performance

In IPv6 networks, we experience better performance, thanks to the hierarchical address-scheme, which makes the routing tables smaller which makes the table lookup in forwarding faster, and the

streamlined header design which makes it much easier to build in fast hardware support for parsing the IP header because the options is put in the extension headers.

IPv6 Implementations

IPv6 implementations started to emerge way before the specifications matured or became complete. As a result, the definition of specifications and developing implementation were going in parallel with the implementation trailing a bit behind the specs. Currently, most operating systems (and routers) provide an implementation to support IPv6, however, the status of each implementation is different being anywhere from experimental to commercial releases.

For information on available IPv6 implementations, visit <http://www.ipv6.org/impl/>. This site provides an extensive list of ipv6 implementations sorted by operating systems and routers' vendors.

For the purpose of this article, we will not go into further detail on the IPv6 implementation, as this will be the subject of an upcoming article in Linux U&D.

Migrating to IPv6

It was clear from the beginning that a complete rollout of IPv6 would not happen overnight. Therefore, it was very important for IPv6 to be able to inter-operate with IPv4 by providing practical transition mechanisms. The two most important transition requirements are flexibility of deployment and the ability for IPv4 hosts to communicate with IPv6 hosts.

When designing IPv6, IETF had transition in mind and it has developed many transition and co-existence mechanisms to offer a toolbox with special tools for special cases.⁴ As a result, IPv6 came out transition rich with many defined transition mechanisms and techniques that should be applied in the right context.

These techniques fall into three categories:

- **Dual-stack techniques:** IPv4 and IPv6 co-exist in the same devices and on the same network. This approach requires hosts and routers to implement both IPv4 and IPv6 protocols. This technique enables networks to support both IPv4 and IPv6 services and applications during the transition period in which IPv6 services emerge and IPv6 applications become available.
- **Tunneling techniques:** This approach enables the interconnection of IP clouds. Separate IPv6 networks can be interconnected through a native IPv4 service by means of a tunnel. IPv6 packets are encapsulated by a border router before transportation across an IPv4 network and de-capsulated at the border of the receiving IPv6 network.
- **Translation techniques:** This technique allows IPv6-only devices to communicate with IPv4-only devices. Translation is a simple extension to NAT, and is used to translate header formats

⁴ Since deployment depends on a specific situation and no single mechanism applies to all situations.

as well as addresses. IPv6 nodes behind such a translator get full IPv6 functionality when talking to other IPv6 nodes located anywhere; they also get normal NAT functionality when talking to IPv4 devices.

Conclusion

IPv6 is addressing the future by providing addresses for new devices, new applications, and new users, restoring the Internet model optimized for performance, robustness, security, and manageability, and enabling rapid innovation for next-generation applications.

Although IPv6 today has all the functional capabilities of IPv4, implementations are not yet as advanced and the deployment is at its early stages. Many efforts are still needed to move towards full IPv6 connectivity.

People often think that they still have a lot of time left before they start considering IPv6. However, it is not too early to begin planning, deploying, and testing IPv6 networks. By preparing for the transition now, instead of later, we can build a solid knowledge base and be prepared for an easy migration.

IPv6 is a key technology and a long-term solution to build scalable, reliable, manageable, secure, and high-performance IP networks. The technology you have been waiting for is here. Start deploying today!

References

Internet Engineering Task Force	www.ietf.org
IPv6 Forum	www.ipv6forum.com
IPv6 Information Page	www.ipv6.org

Abbreviations

DFZ	Default Free Zone
DHCPv6	Dynamic Host Configuration Protocol for IPv6
DNS	Domain Name System
ICMPv6	Internet Control Message Protocol for IPv6
IETF	Internet Engineering Task Force
IPsec	Internet Protocol Security
IPng	Internet Protocol next generation
IPv4	Internet Protocol Version 4
IPv5	Internet Protocol Version 5
IPv6	Internet Protocol Version 6
ISP	Internet Service Provider
MAC	Media Access Control
NAT	Network Address Translation
TCP	Transmission Control Protocol
UDP	User Datagram Protocol
QoS	Quality of Service

Acknowledgment

The author would like to acknowledge Reiner Ludwig, Mats Näslund, and Hesham Soliman for their expert reviews.

About the author



Ibrahim Haddad (Ibrahim.Haddad@Ericsson.com) is a Researcher at the Open System Lab – member of the IP Network branch at Ericsson Research.