> Learn About Differences in Addressing Between IPv4 and IPv6

IPv6 is the most recent generation of the Internet Protocol (IP) defined by the Internet Engineering Task Force (IETF). Initially defined in a number of RFCs in 1995, which have been obsoleted by RFC 2460, IPv6 has repeatedly been enhanced and modified as scalability and security have taken off in modern networks. While IPv6 is intended to eventually replace IPv4, they are tightly mingled right now. It's an excellent time to understand the differences between these protocols.

The IP layer of the TCP/IP protocol stack is the most crucial piece of the whole Internet architecture. Common figures and illustrations often show an "hourglass" shape with many options and supported technologies for frame and bit-level transport (the bottom of the hourglass), and the many choices for end-to-end messages and the applications that send and receive them (the top of the hourglass), but in the narrow waist of the Internet, there is only IP. The first stable version of IP became IPv4 (IP version 4), and today, whenever the term "IP" appears by itself, it is, without qualification, taken to mean IPv4.

However, within ten years of IP going mainstream in the early 1980s, the limitations of IPv4 in terms of scalability *("Will everyone in the world need an IP address? Never!")* and capability (IPv4 requires several awkward add-ons like ICMP and ARP to function) became painfully obvious. By the mid-1990s, a replacement scheme was developed (IPv5 was taken) but the explosion of the Internet made a simple transition impossible at the time (and some say it will be even more difficult going forward).

This *Learn About* provides a working comparison of IPv4 and IPv6, and an overview of how Juniper Networks implements the growing standard. Rather than throw the switch, and transition completely from IPv4 to IPv6, most engineers are running IPv4 and IPv6 together. So it's critical to understand the differences (and similarities) now, until at some point in the near future when the sheer size of billions of new devices will throw the IPv6 switch.

> Differences Between IPv4 and IPv6

It's important to understand that IPv6 is much more than an extension of IPv4 addressing. IPv6, first defined in RFC 2460, is a complete implementation of the network layer of the TCP/IP protocol stack and it covers a lot more than simple address space extension from 32 to 128 bits (the mechanism that increases IPv6's ability to allocate almost unlimited addresses to all the devices in the world for years to come).

IPv6 offers many improvements over IPv4, and Table 1 compares IPv4 and IPv6 operation at a glance.

- More efficient routing. IPv6 routers no longer have to fragment packets, an overhead-intensive process that just slows a network down.
- Quality of service (QoS) built-in. IPv4 has no way to distinguish delay-sensitive packets from bulk data transfers, requiring extensive workarounds, but IPv6 does.
- Elimination of NAT to extend address spaces. IPv6 increases the IPv4 address size from 32 bits (about 4 billion) to 128 bits (enough for every molecule in the solar system).
- Network layer security built-in (IPsec). Security, always a challenge in IPv4, is an integral part of IPv6.
- Stateless address autoconfiguration for easier network administration. Many IPv4 installs were complicated by manual default router and address assignment. IPv6 handles this in an automated fashion.
- Improved header structure with less processing overhead. Many of the fields in the IPv4 header were optional and used infrequently. IPv6 eliminates these fields (options are handled differently).

> Table 1 Key IPv4 and IPv6 Comparisons <

IPv4	IPv6
32-bit (4 byte) address supporting 4,294,967,296 address (although many were lost to special purposes, like 10.0.0.0 and 127.0.0.0)	128-bit (16 byte) address supporting 2^{28} (about 3.4 x 1038) addresses
NAT can be used to extend address limitations	No NAT support (by design)
IP addresses assigned to hosts by DHCP or static configuration	IP addresses self-assigned to hosts with stateless address auto-configuration or DHCPv6
IPSec support optional	IPSec support required
Options integrated in header fields	Options supported with extensions headers (simpler header format)

Figures 1 and 2 compare the header of a IPv4 packet and an IPv6 packet. Even if you don't study packet header fields, you can see the difference between the IPv4 header (fragmentation fields are *italicized*) and the "streamlined" IPv6 header. Note the reduction in header fields in the IPv6 packet that routers need to process or examine.

IPv6 packets have their own frame Ethertype value, `0x86dd`, making it easy for receivers that must handle both IPv4 and IPv6 to distinguish the frame content on the same interface. The IPv6 header is comprised of the following fields:

- **Version:** A four-bit field for the IP version number (0x06).
- **Traffic Class**: An 8-bit field that identifies the major class of the packet content (for example, voice or video packets). The default value is 0, meaning it is ordinary bulk data (such as FTP) and requires no special handling.
- **Flow Label**: A 20-bit field used to label packets belonging to the same flow (those with the same values in several TCP/IP header parameters). The flow label is normally 0 (flows are detected in other ways).

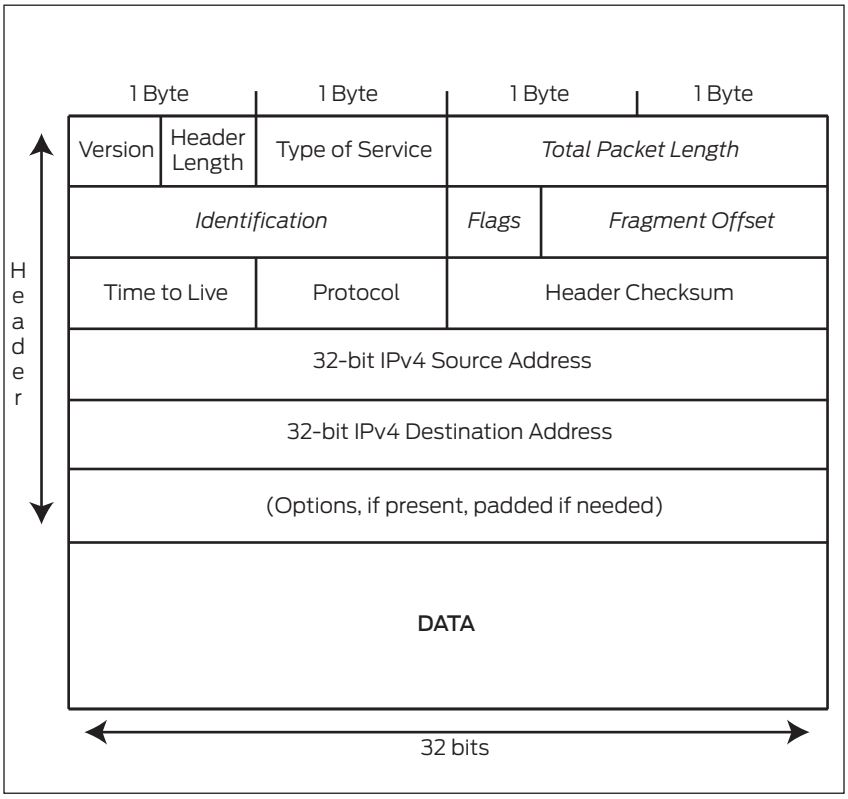

> Figure 1 IPv4 Header <

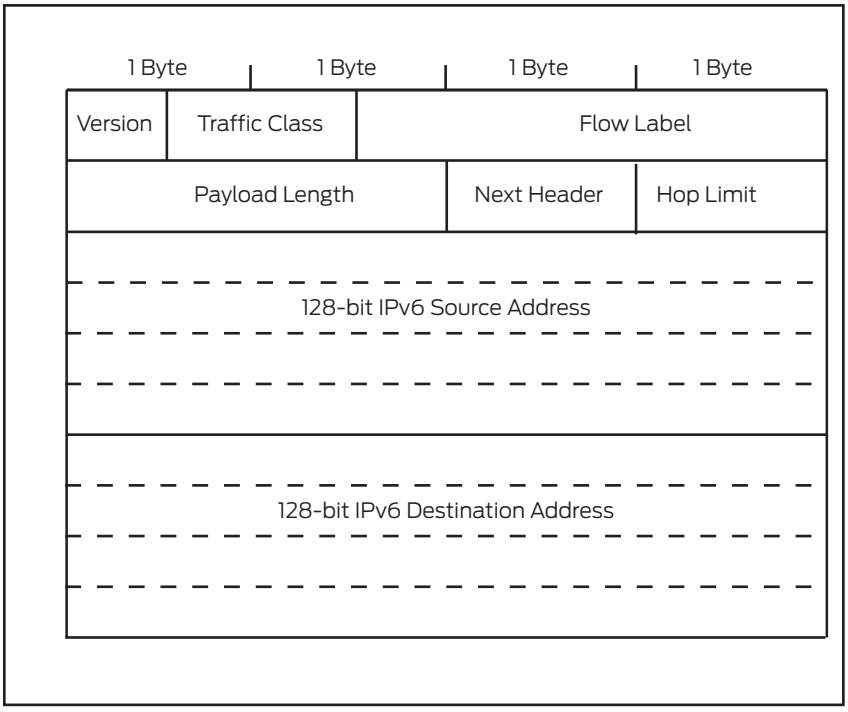

> Figure 2 IPv6 Header <

- **Payload Length**: A 16-bit field giving the length of the packet in bytes, excluding the IPv6 header.
- **Next Header**: An 8-bit field giving the type of header immediately following the IPv6 header (this serves the same function as the Protocol field in IPv4).
- **Hop Limit**: An 8-bit field set by the source host and decremented by 1 at each router. Packets are discarded if Hop Limit is decremented to zero (this replaces the IPv4 Time To Live field). Generally, implementers choose the default to use, but values such as 64 or 128 are common.

> IPv6 Address Notation

Those headers on IP packets are important, but often invisible and network administrators are not always concerned about the value of a certain header field. Since most readers are familiar with IPv4 address notation (172.23.22.45), let's explore how IPv6 addresses are written down and discussed.

The dotted decimal address notations of IPv4, shown in the last paragraph, should be familiar to everyone. IPv6 addresses are written very differently than IPv4 notation. Remember there are now 128-bits in the IPv6 address, and a notation system was devised to represent this enormous address pool potential. So, in IPv6, the 128-bits in the IPv6 address:

- Are written as eight 16-bit hexadecimal blocks separated by colons (not case sensitive)
- Use abbreviations to simplify the notation
- Omit (optionally) leading zeroes
- Use double colons (::) to replace consecutive zeros (or leading or trailing zero strings), but never more than once per address

So an address like:
`2dfc:0000:0000:0000:0217:cbff:fe8c:0000`

Can be written:
`2dfc:0:0:0:0217:cbff:fe8c:0`

Or:
`2dfc::0217:cbff:fe8c:0`

But *not*:
`2dfc::0217:cbff:fe8c::`

Why? The double use of :: makes it unclear how many zeros were in each 0 string originally.

Juniper Networks complies with RFC 5952 in the standard assignment and display rules for IPv6 addresses. These rules mean that:

- Devices must accept all methods of address entry
- Leading zeros are always suppressed
- The double colon (::) is always used when applicable (if there are two 0 strings, the longer string always gets the ::)
- Hexadecimal characters (a, b, c, d, e, f) must be represented in lower case

> IPv6 Prefix Notation

Routers seldom have to worry about complete ("host") addresses because their routing tables (and the forwarding tables they are based on) usually employ a prefix and examine only the number of bits that match the longest entry in the table. This is the "longest match" rule, and ensures that a 64-bit prefix (if present) is preferred over a 32-bit prefix; when the first 32 bits of a packet's destination address are the same in both table entries (longest match wins).

Prefixes used for routing are defined in IPv6 by RFC 4291. So the IPv6 address:

`2bfc:0000:0000:0000:0217:cbff:fe8c:5c85/64`
…Has a 64-bit prefix of 2bfc:0000:0000:0000:

RFCs The rules and format for using literal IPv6 addresses in URLs are detailed in RFC 2732. All Juniper Networks devices comply with both RFC 4291 and RFC 2732.

> IPv4 and IPv6 Address Usage Compared

The expanded address space is certainly one of the most distinctive features of IPv6. However, there is more to addressing than assigning an end device or network node (router or switch) an address. IPv6 eliminates some of the most awkward applications of IPv4 addresses and has some notable enhancements.

For example, the IPv4 address space was divided into "classes" with Class A networks (a few huge networks with thousands of end user devices), Class C networks (thousands of small networks with only a handful of devices), and Class B networks that are somewhere in between. But today's world operates on millions of networks, and many of them are enormous, especially on the service provider side. So IPv6 does away with the rigid class concept (it was already obsolete in IPv4) and uses subnetting (more on that later) to adjust network sizes with a given address space assignment.

Another oddity of IPv4 was the use of a class-type address space for multicast use (224.0.0.0/4). Devices need both unicast and multicast addresses to receive multicast traffic. IPv6 uses a more integrated address space for multicast, at FF00::/8.

IPv4 allowed for "broadcast" addresses that forced each and every device to stop and look at packets, even if they were useless to them (these were very popular in denial of service attacks). IPv6 does away with broadcast and uses multicast groups for everything.

IPv4 uses 0.0.0.0 as an unspecified address (for example, when a device does not yet know its IP address), and uses another whole class-type address (127.0.0.1) for loopback (testing). IPv6 uses :: and ::1 as unspecified and loopback address respectively.

Finally, IPv4 uses globally unique public addresses for normal traffic and special "private" addresses that should not appear on the public Internet (such as 10.0.0.0/8). Instead, IPv6 uses globally unique unicast addresses and local addresses (FD00::/8).

Table 2 shows some of the major differences between IPv4 address types and their use in IPv6.

> Table 2 IPv4 and IPv4 Addressing Concepts <

IPv4	IPv6
Mulitcast address space at 224.0.0.0/4	Multicast address space at FF00::/8
Has broadcast addresses for all devices	No such concept in IPv6 (uses multicast groups)
Uses 0.0.0.0 as unspecified address	Uses :: as unspecified address
Uses 127.0.0.1 as loopback address	Uses ::1 as loopback address
Supports globally unique "public" addresses	Supports globally unique unicast addresses
Uses 10.0.0.0/8, 172.16.0.0/16, and 192.168.0.0/16 as "private" addresses	Uses FD00::/8 as unique local addresses

> IPv6 Address Allocation and Types

In the same way as IPv4, organizations and end users get IPv6 address allocations from their ISPs. The process is defined in RFC 6177, which is supported by Juniper Networks devices:

- Any subnet with a prefix of /64 or shorter (that is, less than 64 bits), is supported. In other words, an ISP should not be giving out addresses with /96 prefixes, or any other value greater than /64.
- Initial address assignments should allow for a site to establish multiple subnets. Subnetting in IPv6 is discussed a little later in this *Learn About*.
- Users should allow a user or site to easily obtain additional IPv6 address space for additional subnets. Usually, a simple form request is all that is needed.

IPv6 defines three distinct types of addresses:

- Unicast addresses: Addresses used to send packets to one specific address.
- Multicast addresses: Addresses used to send to every member of a specific group. Group membership is either mandatory (a device *must* belong to certain multicast groups) or optional (a device can decide whether to join the multicast group or not).
- Anycast addresses: Addresses used to send to any *one* member of a specific group (such as the "closest" in terms of routing metrics). For example, a device can use the same anycast address to access a video stream located in the same country, or halfway around the world. The routers and other network nodes decide which source is "closest."

> Notable IPv6 Addresses

Certain IPv6 addresses are seen over and over, and it is a good idea to know what types of IPv6 address they represent. Note that it is common for end devices to have many IPv6 addresses, and for routing tables to contain many entries reflecting these types of IPv6 addresses.

> Table 3 Notable IPv6 Addresses <

6bone (experimental) (Currently unallocated per RFC 3701)	5F00::/8 (RFC 1897) 3FFE::/16 (RFC 2471)
6to4 (Must be filtered and some are not used)	2002::/16 (RFC 3056, but see RFC 3964)
Default route (See IANA registry)	For experimental purposes (RFC 4773)
Documentation prefix for user manuals, etc. (Should never be advertised publicly)	2001:DB8::/32 (RFC 3849)
IPv4-mapped Addresses (Should never be advertised publicly)	::FFFF:0:0/96 (RFC 4291)
IPv4-compatible Addresses (Deprecated)	::<ipv4-address>/96 (RFC 4291)
Link-scoped Unicast (Should never be advertised publicly)	FE80::/10 (RFC 4291)
Multicast (Must not appear in unicast routing tables)	FF00::/8 (RFC 4291)
Node-scope Unicast (Should never be advertised publicly)	::1/128 is loopback (RFC 4291) ::/128 is unspecified (RFC 4291)
ORCHID (Not IP routable) (Should never be advertised publicly)	2001:10::/28 (RFC 4843)
Teredo (Advertised for Teredo service)	2001::/32 (RFC 4380)
Unique-local Addresses (Should never be advertised publicly)	FC00::/7 (RFC 4193)

> IPv6 Host Addressing

IPv4 hosts are fairly easy to configure: usually, the network interface has one IPv4 address. When coupled with the default router – if there is a way off the subnet – the host has everything it needs to decide where things go. But IPv6 does much more.

In contrast to IPv4 hosts, IPv6 hosts (end devices) normally have multiple addresses on each interface. But these multiple addresses greatly simplify the operation of the IPv6 network layer (finding network neighbors, routers, and so on).

- Unicast addresses
  - Link-local address on each interface (Unique Local Addresses, ULA, beginning with FE00::/7 can use www.sixxs.net/tools/grh/ula/ to generate and register site local prefixes based on RFC 4193)
  - Additional unicast addresses for each interface, which can be multiple global addresses or unique local addresses
  - Loopback address (::1)
- Multicast addresses
  - FF01::1 – The interface-local scope all-nodes multicast address
  - FF02::1 – The link-local scope all-nodes multicast address
  - The solicited-node address for each assigned unicast address
  - The multicast addresses of any groups the host has joined

> IPv6 Router Addressing

Remember, packets sent with unicast IPv6 address are sent to only one destination, and packets sent to anycast IPv6 addresses are destined for any one member of a defined group. Routers in IPv6 have different requirements for unicast and anycast addresses. The physical interfaces on an IPv6 router are assigned the following IPv6 addresses:

- Unicast addresses
  - Link-local address on each interface
  - Additional unicast addresses for each interface, which can be multiple global addresses or unique local addresses
- Anycast addresses
  - Subnet-Router anycast address for each subnet established
  - Optionally, additional anycast addresses

All IPv6 router interfaces also listen for traffic on the following multicast addresses:

- FF01::1 – The interface-local scope all-nodes multicast address
- FF01::2 – The interface-local scope all-routers multicast address
- FF02::1 – The link-local scope all-nodes multicast address
- FF02::2 – The link-local scope all-routers multicast address
- FF05::2 – The site-local scope all-routers multicast address
- The solicited-node address for each assigned unicast address
- The multicast addresses of any groups the router has joined

> IPv6 Protocols

There is more to IPv6 than addressing, of course. IPv6 bundles several adjunct protocols that were completely separate in IPv4 into the main IPv6 operation. IPv6 does not use ARP to "map" network layer IPv4 addresses to link level (frame) addresses. Also, IPv4's ICMP (the Internet Control Message Protocol) has been redone as ICMPv6.

With the assistance of ICMPv6, IPv6 can discover neighbors and routers through special neighbor discovery (ND) and router advertisement (RA) messages. This simplifies host configuration on LANs and lets devices find default routers on their own. ICMPv6 can automatically determine path MTU limitations, and notify the sender if this is not respected (IPv6 routers do not fragment packets —all fragmentation must be done at the sending host).

IPv6 also does away with the IPv4 IGMP protocol for multicast group membership because the group membership function is bundled with ICMPv6.

> IPv6 Subnetting

Subnetting is a way for users to take an assigned IP address space and partition it to meet their specific needs. For example, if an organization has two office locations, one large and one small, the IP address space can be subnetted so that the hosts get

the addresses they need, and traffic can be efficiently handled internally without concerning the global public Internet. Subnetting today in IPv4 and IPv6 is done by prefixes, but again IPv6 has its own rules.

As an example of how an IPv6 address would be subnetted, assume that the ISP has provided a 48-bit prefix to the user. Consider the address and prefix *2001:0867:5309/48* for instance. Because each four hexadecimal digits are 16 bits, and the last 64 bits are usually the interface identifier (the "host" portion of the address in IPv4), this leaves 16 bits for subnetting.

The 16 subnet bits allow for 2^8 or 256 subnets. In each subnet, all address possibilities are allowed except for all zeros, which is reserved in IPv6 for subnet router anycast messages. Assume the subnet bits, assigned locally, for a certain subnet are *9abc*.

This subnet would allow 2^{64}-1 addresses in the range:

2001:0867:53099abc:0000:0000:0000:0001 to 2001:0867:53099abc:ffff:ffff:ffff:ffff

So, when subnetted, an IPv6 address consists of three parts:

- The global routing prefix (*2001:0867:5309/48* in this example)
- The subnet identifier (*9abc* in this example)
- The interface identifier or "host address" (the remaining 64 bits in this example)

Things can get a little tricky when subnets are not established on the 16 bit "boundaries." However, the principles of subnetting are almost the same as in IPv4.

> Junos Address Aware

Junos Address Aware is an addressing and tunneling software portfolio for the MX Series routers that helps network operators conserve and extend their IPv4 address pool, ensure IPv4/IPv6 coexistence, and pragmatically transition to IPv6 in a cost-effective and low risk manner. It's available as a software license for MX Series MS-MPC and MS-MIC Service Cards:

- Offered as licensed software that protects investments in new or existing MX Series 3D Universal Edge Routers.
- MS-MPC and MS-MIC service cards provide very high address translation performance and density, enabling cost-effective service scale without impacting routing or forwarding functions.
- One software license supports many technologies, including IPv4/ IPv6 dual stack, NAT44, and NAT64, among others, providing deployment flexibility and non-disruptive address scheme evolution.
- Low risk service adoption is based on mature, field-proven products and protocols deployed on a carrier-grade router via software update.
- Rich application-level gateway (ALG) support protects traffic and revenue streams by ensuring compatibility with popular applications that cannot be translated.

Junos Address Aware has several datasheets and more information on Juniper.net.

> Junos Pulse

Junos Pulse 4.0 and Pulse Secure Access Service 7.4 have been enhanced so that today, end users are now able to access IPv6 resources – along with IPv4 resources – from an IPv4 network, by simply using Junos Pulse to access Junos Pulse Secure Access Service 7.4.

> Juniper.net/IPv6

So, while the move to IPv6 is necessary to accommodate the continuing explosion of Internet space requirements, current Internet technology profile mandates that access via IPv4 and access via IPv6 have to coexist now, and for the foreseeable future. Read more about Juniper's Address depletion strategies, here at Juniper.net/IPv6.

> Interworking Mechanisms

RFC 2893 includes the following well-known interworking mechanisms for IPv6 and IPv4:

- Dual IP Layer (Dual stack): Hosts and routers implement a complete suite for both IPv6 and IPv4
- Configured Tunneling of IPv6 over IPv4: Point-to-point tunnels are used to encapsulate IPv6 packets with IPv4 headers to carry them over IPv4 routing networks
- Automatic Tunneling of IPv6 over IPv4: Uses IPv4-compatible addresses to automatically tunnel IPv6 over IPv4 routing infrastructures
- IPv4 Multicast Tunneling: A form of IPv6-over-IPv4 tunneling for multicast where the IPv4 tunnel endpoint address is determined using neighbor discovery which, unlike configured tunneling, does not require address configuration and, unlike automatic tunneling, does not require the use of IPv4-compatible addresses.

> References and Further Reading

- The following RFCs can be examined at: http://tools.ietf.org/
  RFC 2460, RFC 5952. RFC 2732, RFC 6177, RFC 4291, RFC 4193, RFC 3849, RFC 3056, RFC 3964, RFC 4380, RFC 3701, RFC 1897, RFC 2471, RFC 4843, RFC 4773, RFC2893,
- *The Illustrated Network*, by Walter Goralski
- *Day One: Exploring IPv6*
- *Junos 13.1 IPv6 Implementation Guide*, and other documentation.

> Learn About Differences in Addressing Between IPv4 and IPv6

by Walter Goralski

This *Learn About* reviews the key addressing differences between IPv4 and IPv6 that you can use as a backdrop for understanding transition technologies to bridge the IPv4 / IPv6 world.

Walter Goralski is a Technical Lead at Juniper Networks, where he has worked for the past 12 years. He has published a half dozen books about networking and the telecommunications industry, including the best-selling The Illustrated Network.

Version History: First Edition, January 2014 2 3 4 5 6 7 8 9

> For more *see*:
juniper.net/documentation

ISBN 978-1936781904

9 781936 781904 50500