



Introduction to IPv6 -- Part 2

Peter J. Welcher

Introduction

This month we continue our discussion of IPv6. Our focus will be on introducing the technical basics of IPv6.

In case you missed it, last month's article introduced IPv6, with an eye towards answering the very important question "and why do I care about IPv6?" The article can be found at the link <http://www.netcraftsmen.net/welcher/papers/ipv6part01.html>.

In that article, we saw that there is no clear "killer application" yet for IPv6. (This article is being written in November of 2006).

We also saw that U.S. government agencies are making their network backbones IPv6 ready as equipment is replaced. And we noted that both the Pacific Rim and Europe are pushing into IPv6. New applications that originate from those parts of the world could have impressive features requiring IPv6. Do you want to be the person who has to tell the CIO that you didn't think your equipment needed to be IPv6 capable, and so you cannot run IPv6 until the next hardware refresh in your network, say in 3-5 years? In short, if you ignore IPv6, you're betting you won't need to run it during the lifetime of your next equipment purchase.

We also talked about device-specific support for IPv6 at a high level.

At the [end of this article](#) you will find some of the best Cisco links I've found to date. Prior to that, I've also included [IPv6 security issues links](#), links to good discussions about what's different with security in IPv6. One reason I'm doing this is because I have little to add on that topic, and because I suspect that some other topics will have broader appeal. I do plan to discuss the IPv6 migration mechanisms supported by Cisco in subsequent article.

IPv6 Address Basics

IPv6 addresses are 128 bits long. This provides 2^{128} or approximately 10^{43} addresses. One intent is to address running out of address space with IPv4. Last month's article mentioned some of the interesting ideas for how to use this plethora of addresses.

The addresses are written as 8 blocks of hexadecimal digits separated by colons (:). Thus 1234:ABCD:ef01:0000:0000:0000:0987:4321 might be an IPv6 address. Leading zeros in any of the 8 blocks can be omitted, so the "0987" could be written as just "987". And "0000" as "0". Furthermore, any **one** block of consecutive zeros can be replaced by double colon (::). So the previous address could be shortened to 1234:abcd:ef01::987:4321. The addresses are not case sensitive.

The shortest IPv6 address is the all-zero address ::0 or ::. It is used for an unspecified address, as when a node is requesting an address be assigned by DHCPv6.

Yes, they could be a bit hard to remember (or 128 bits hard to remember).

The reason that double colon (::) can only be used once is that the parser needs to be able to figure out how many zeros got omitted.

When an IPv4 address is embedded in an IPv6 address, it can be written in the usual dotted form. In such a case, it is preceded by a prefix written in IPv6 form.

IPv6 address prefix lengths are specified as with IPv4, using slash (/) notation. Thus /24 would be prefix, the first 24 bits of which are significant. A /24 would be a much bigger block of addresses in IPv6 than in IPv4.

The official reference on IPv6 addressing is the "IPv6 Addressing Architecture", [RFC4291](#).

IPv6 Address Types

IPv6 uses three type of addresses: unicast, anycast, and multicast. These are used a bit differently than in IPv4.

The unicast address belongs to a single interface on a single device. **It may be one of several kinds, discussed below. A link-local unicast address (see below) cannot be the source address in a packet. [Part in red revised 8/24/2007].**

An anycast address identifies any of several interfaces, generally one interface on each of several devices. A packet with an anycast destination is delivered to the closest interface with that anycast address. One use would be for a service such as DNS, where a reply from any of several servers would suffice. IPv4 anycast is common already for some of the top level DNS domain servers. IPv6 just makes this practice potentially a more common one.

A multicast address is assigned to interfaces on different devices. A packet sent to the multicast address is to be delivered to all the interfaces with that multicast address. IPv6 multicast addresses have scopes, which reflects the size of the portion of the network where that address is valid. For example, site scope would mean multicast destinations throughout a company or site.

IPv6 Unicast Address Information

IPv6 addresses are assigned to interfaces, as in IPv4. Unlike IPv4, which assigns a unique address to an interface, an IPv6 interface **must** have multiple addresses. Any one of the assigned unicast addresses can be used to communicate with a device.

Some special cases:

- Multiple physical interfaces may have a single unicast address when used for load sharing.
- Routers can use unnumbered interfaces on point-to-point links.

The following sorts of unicast addresses are available for use:

Type of IPv6 address	Prefix
Global unicast addresses	2000::/3
Site-local unicast addresses (deprecated)	FEC0::/10
Unique local unicast addresses (replaces site-local)	FD00::/8
Link-local unicast addresses	FE80::/10
IPv4-mapped IPv6 address (96 0's then the hex form of the IPv4 address)	::pqrs:tuvw, also written as ::M.N.P.Q
IPv4 compatible IPv6 address (somewhat deprecated) (80 0's then FFFF then the IPv4 address)	0:0:0:0:FFFF:pqrs:tuvw, also written as ::FFFF:M.N.P.Q
Unspecified address	0:0:0:0:0:0:0:0 or 0::0 or ::/128
Loopback	0:0:0:0:0:0:0:1 or ::1

Further details:

IPv6 global unicast addresses are often considered to be in the form [subnet prefix, interface ID (identifier)]. As noted above, we use notation such as /64 to indicate the number of subnet bits. All unicast addresses other than those starting with binary "000" are required to have 64 bit interface IDs in Modified EUI-64 format.

A **Modified EUI-64 interface identifier** can be universal (global scope) or local. One way they are created is by taking a 48 bit MAC address and inserting 0xFFFE in the middle. The next to low order bit for global/local scope is toggled from 0 to 1 as well. Thus 00:11:22:33:44:55 becomes 02:11:22:FF:FE:33:44:55. (The bit toggling is where the "modified" comes in).

See the RFC for further details concerning Modified EUI-64. .

The **site-local unicast** approach is deprecated. Instead, use **unique local unicast addresses**. See [RFC4193](#). The idea is that unique local unicast addresses are not routed on the Internet, and generally filtered inbound. They are Internet Service Provider independent, suitable for use at sites not connected to the Internet. If leaked accidentally, there is no conflict with assigned IPv6 prefixes.

The key idea with unique local unicast addresses is to follow the specified prefix in the table with a 40 bit pseudo-random "global ID". This greatly reduces the chance that two users of unique local addresses might find themselves using the same prefixes. That is, what happens now when two companies using network 10.0.0.0 /8 have to communicate with each other -- forcing one-way to two-way NAT. Unique local greatly mitigates this "network merge" issue. It also provides suitable addressing for VPNs between sites.

FC00::/8 is reserved for future use.

Link-local unicast addresses are for use on a single link for the following uses:

- automatic address configuration
- neighbor discovery
- when no routers are present

Routers should not forward packets with link-local source or destination addresses to other links. That is, the point to link local addresses is that they are for communication strictly on one link.

The **unspecified address** is for use when an address is absent, for example before an initializing host determines its own address, as occurs with DHCP.

The **loopback address** is used by a node sending packets to itself. It must not be a source address for packets transmitted outside a single node.

Anycast Address Information

The first bits of an anycast address are used for scope. Zero bits of scope would mean an Internet-wide anycast address. Since each anycast address consumes a routing table entry, there are expected to be very few Internet-wide anycast addresses. However, [RFC4291](#) goes on to note that one expected use would be to identify the routers belonging to an Internet Service Provider. (Doesn't that have security implications?)

The **Subnet-Router anycast address** must be used (is required). The format is simple: [subnet prefix, zeros]. Packets sent to this address are delivered to a router on the subnet.

IPv6 Multicast Address Information

IPv6 multicast addresses begin with the prefix FF00::/8. That means the first byte is all one bits. The next two nibbles (4 bit half-bytes) are used for flags and scope.

The format as specified in [RFC4291](#):





Fun Fact: IPv6 multicast flag bits:

- 0: reserved
- T bit: 0 = permanently-assigned "well known" multicast address, 1 = transient (non-permanent)
- P bit: P = 0 is a "normal" multicast address, P = 1 requires T = 1, and indicates a multicast address based on network prefix. This allows Internet Service Provider (ISP) - assigned multicast addresses without some method for allocating multicast addresses across ISPs. For details, see [RFC3306](#).
- R bit: R = 1 indicates the address embeds the address of the Rendezvous Point (RP). This in turn forces T = 1. Another way of looking at this: the prefix FF70::/12 indicates embedded RP information. For details see [RFC3956](#).

The 4 bits for scope govern how broadly the multicast is to be forwarded. The assigned values per [RFC4291](#):

Value	Multicast Scope	Usage
0	reserved	
1	Interface-Local scope	multicast on a single interface, useful for loopback multicast
2	Link-Local scope	multicast on a single link
3	reserved	
4	Admin-Local scope	must be administratively configured
5	Site-Local scope	multicast within a single site
8	Organization-Local scope	multicast within multiple sites belonging to one organization
E	Global scope	
F	reserved	

The unassigned scopes are for "administrative control of additional multicast regions".

Routers are required to not forward multicast packets beyond the scope indicated in the destination multicast address.

The multicast addresses FF00 - FF0F:0:0:0:0:0:0 are reserved, not for assignment.

Multicast addresses of interest:

- All Nodes Addresses: FF01:0:0:0:0:0:1 (interface-local scope) and FF02:0:0:0:0:0:1 (link-local scope)
- All Routers Addresses: FF01:0:0:0:0:0:2 (interface-local), FF02:0:0:0:0:0:2 (link-local), FF05:0:0:0:0:0:2 (site-local)
- Solicited-Node Address: FF02:0:0:0:0:1:FFXX:XXXX

A **Solicited-Node multicast address** is formed by taking the low-order 24 bits of an address (unicast or anycast) and appending those bits to the prefix FF02:0:0:0:0:1:FF00::/104 resulting in a multicast address in the range FF02:0:0:0:0:1:FF00:0000 to FF02:0:0:0:0:1:FFFF:FFFF.

For each interface it has, a node must determine (compute) and join the solicited-node multicast address for every unicast and anycast address on that interface.

For a list of assigned IPv6 multicast addresses, see <http://www.iana.org/assignments/ipv6-multicast-addresses>.

Required Addresses

Now that we know something about all the IPv6 address types, you are probably wondering "just how many addresses does IPv6 use?"

The answer is, several. You may have noticed from the brief discussion of solicited-node multicast above that "we're not in IPv4 anymore".

Here is the official list of all the addresses required of a node:

- The link-local address for each interface
- Additional unicast and anycast addresses configured on interfaces (manually or automatically)
- Loopback
- All-nodes multicast addresses as above
- Solicited-node multicast addresses for each unicast and multicast address
- Multicast addresses for the multicast groups the node belongs to (joins)

Routers must recognize all the above addresses, plus more:

- The Subnet-Router anycast addresses for all interfaces on which routing is enabled
- All other anycast addresses the router has been configured with
- The All-Routers multicast address (see above)

Routing IPv6

As we noted in the previous article, the various IPv4 routing protocols have been extended into IPv6 routing protocols. Details vary. The good news, there seems to be far less novelty here than there is with the addressing, above.

Here is a summary providing a little more detail:

IPv6 Routing Protocol	Comments	First Cisco Support
RIPng	RFC2080 Uses the multicast address FF02::9 for RIP updates.	12.2(2)T, 12.0(22)S, etc.
OSPFv3	RFC2740 Uses the multicast addresses FF02::5 and 6 (AllSPFRouters, AllDRouters).	12.0(24)S, 12.2(15)T, etc.
EIGRP support for IPv6	Cisco proprietary (of course). Uses multicast FF02::A.	12.4(6)T
IS-IS support for IPv6	Being submitted to the IESG as Proposed Standard, see http://www.ietf.org/internet-drafts/draft-ietf-isis-ipv6-06.txt or http://www.ietf.org/html.charters/isis-charter.html . Basically the same IS-IS but with new TLV's, similar to the way IPv4 is supported.	12.0(22)S, 12.2(8)T
MBGP support for IPv6	RFC2545	12.0(22)S, 12.2(2)T

In general, see the [support reference link](#) below, for which Cisco IOS Release various IPv6 features were first supported in.

There is another aspect to IPv6 Routing you may not think of at first. A vital part of modern routing is high performance multi-layer switching. The Cisco documents on the topic refer to this as "hardware support for IPv6 forwarding". Please note that the older platforms (including MSFC2) do not support this. If you can live with the performance implications of software (process) switching of IPv6 traffic, say due to low volumes of IPv6 traffic, then this may not concern you. If you

plan to support high volumes of IPv6 traffic, then you should look at the Cisco document at the previous link. The specific sub-section is http://www.cisco.com/en/US/products/sw/iosswrel/ps5187/products_configuration_guide_chapter_09186a00801d65ed.html#wp1076187.

Operation of IPv6

The IETF designers changed or improved on IPv4 behavior in some areas that affect how the protocols works.

IPv6 uses a Neighbor Discovery Protocol, also Router Discovery.

IPv6 also makes more substantial use of ICMPv6, changes the rules on who may fragment a packet, and incorporates Path MTU Discovery for IPv6.

IPv6 supports stateless autoconfiguration, also DHCPv6. Stateless autoconfiguration uses a prefix learned from a router plus the Modified EUI-64 MAC address to build an address. DHCPv6 provides an alternative. There is also work on privacy, which I read as in effect working with a random EUI-64 address instead of one based on the MAC address.

Neighbor Discovery lets nodes and routers find the link-layer (MAC) address of a neighbor on the same link, find neighboring routers, and track neighbors. It uses ICMPv6 and solicited-node multicasts to determine the link (MAC) address or verify reachability. You should be thinking "ARP replacement, with some enhancements". The two aspects of this are Neighbor Solicitation and Neighbor Advertisement.

Router Discovery allows IPv6 nodes to discover routers on the local link. It is similar to Router Discovery (IRDP) in IPv4, [RFC1256](#). It is based on Router Solicitation and Router Advertisement.

Concerning ICMPv6 and Path MTU Discovery, routers are not allowed to perform fragmentation, which is both undesirable anyway but potentially a real performance impact. That means Path MTU Discovery (P-MTU-D) must be allowed to operate, or there will be no remedy. Blocking "MTU exceeded" error messages defeats P-MTU-D and has been a persistent nuisance with IPv4.

Personally, I'll grant that fragmentation is extremely undesirable. Having said that, I would like to see a real solution to P-MTU-D issues caused by naive firewall settings. Assuming administrators are all well-versed and will configure firewalls correctly seems to be proven to be overly optimistic. Are we seeing denial that P-MTU-D has some real operational problems? Or am I missing something? (Cisco has various techniques to address the problem, but none completely solves the problem).

One conclusion: security administrators will need to be trained, and must be very careful not to block several types of ICMP packets with IPv6, as the protocol makes a more fundamental use of ICMP than IPv6 does.

IPv6 Security

The way that IPv6 works changes some of the things one must consider concerning security. Rather than attempting to summarize the present state of ongoing discussions, I prefer to suggest that the reader look at the following interesting links.

Title	Link
IPv6 Security Technology Paper Author/Editor: Merike Kaeo, Contributing Authors: David Green, Jim Bound, Yanick Pouffary	http://www.ipv6forum.org/modules.php?op=modload&name=News&file=article&sid=32&mode=thread&order=0&thold=0
IPv6 Security Links Sean Convery	http://www.seanconvery.com/ipv6.html
IPv6 and IPv4 Threat Comparison	http://www.seanconvery.com/v6-v4-threats.pdf

and Best-Practice Evaluation, March 2004, S. Convery, D. Miller
--

There are a lot more articles on the topic. The discussions about Shimv6 and general Multihoming for IPv6 have a security component to them.

Conclusion

The RFCs mentioned in this article are (as usual) carefully written and fairly short and readable.

Here are some IPv6 links to sites with links that I've found useful:

Title	Link
North American IPv6 Task Force Library	http://www.nav6tf.org/html/nav6tf_library.html
IPv6 Forum	http://www.ipv6forum.org/
IPv6 Portal	http://www.ipv6tf.org/
IPv6 Security page, Adrian Portelli	http://www.stindustries.net/IPv6/index.html
IPv6 Summit, 6Sense newsletter	http://www.usipv6.com/publications.html
ARIN IPv6 Policies	http://www.arin.net/policy/nrpm.html#ipv6
RIPE IPv6 Addressing Policy	http://www.ripe.net/ripe/docs/ipv6policy.html

For convenience, I'm going to repeat the previous article's table of the best Cisco links I've found concerning IPv6:

Title	Link
Cisco IPv6 Technology Page	http://www.cisco.com/en/US/products/ps6553/products_ios_technology_home.html
Cisco IOS IPv6 Configuration Library (12.4)	http://www.cisco.com/en/US/products/sw/iosswrel/ps5187/products_configuration_guide_chapter09186a00801d65ed.html
Cisco IOS Release Specifics for IPv6 Features	http://www.cisco.com/en/US/products/sw/iosswrel/ps5187/products_configuration_guide_chapter09186a00801d65ed.html

Your comments, questions, and suggestions for future articles are of course welcome! See below to decipher my email address.

Dr. Peter J. Welcher (CCIE #1773, CCSI #94014, CCIP) is a Senior Consultant with Chesapeake NetCraftsmen. NetCraftsmen is a high-end consulting firm and Cisco Premier Partner with multiple specializations, dedicated to quality consulting and knowledge transfer. NetCraftsmen has eight CCIE's, with expertise including large network high-availability routing/switching and design, IP Telephony, QoS, MPLS, IPsec VPN, wireless LAN and bridging, network management, security, IP multicast, and other areas. See <http://www.netcraftsmen.net> for more information about NetCraftsmen. Pete's links start at <http://www.netcraftsmen.net/welcher>. New articles will be posted under the Articles link. Questions, suggestions for articles, etc. can be sent to pjw <at> netcraftsmen <dot> net.

11/11/2006 Updated 8/24/2007
Copyright (C) 2006 Peter J. Welcher