

# Scotti-BYTE Enterprise Consulting Services

## The Joys of Ipv6

We are all familiar with the Ipv4 dotted address format of a.b.c.d where each octet is an eight bit number from 0 to 255. We also know about subnet masks and network address translation (NAT). RFC 1918 (NAT) provides a means to replace the private address in packets sent by hosts behind a router on an internal network with its own address and the reverse for packets coming into the private network from the outside.

There are two reasons for NAT. One is so that multiple computers on the local network can share one public internet address. The other is to provide a mechanism for firewall protection to the LAN client nodes. The more practical reason is that the Internet Assigned Numbers Authority (IANA) has been running out of assigned numbers blocks in the Ipv4 address range to grant. Ipv4 has a total of 4,294,967,296 unique addresses. This pool of numbers is  $2^{32}$  bits in size and was thought at one time to be unlimited.

IPv6 addresses are made up of eight 16 bit numbers (each 0-65,535). This address range provides a pool of numbers that is  $2^{128}$  bits in size with a total of 340,282,366,920,938,463,463,374,607,431,768,211,456 IPv6 addresses. IPv6 addresses are represented as eight of these sixteen bit numbers separated by colons and normally each number is represented as a 16 bit hexadecimal number having four hexadecimal digits.

When the Internet Assigned Numbers Authority (IANA) first assigned addresses in the early days of the Internet, some organizations got excessively large address blocks. For instance, IBM, Xerox, HP, DEC, Apple and MIT all received "class A" address blocks of nearly 17 million addresses. (So HP, which acquired DEC, has more than 33 million addresses.) Today we are basically out of addresses. Although IPv6 is in wide use on the public internet, most home users do not make use of the new address scheme. Eventually, Ipv4 will be retired.

For right now, it is very viable to run a local area network with both IPv4 and IPv6 concurrently. Linux has supported IPv6 since 1996. Windows started supporting IPv6 back in 2009 with Windows Vista. Some older devices and IoT may not support IPv6 or may not support it very well. That's a good reason to run IPv4 and IPv6 together,

The typical way that IPv4 addresses are granted is with Dynamic Host Configuration Protocol (DHCP) or if a client node is statically addressed. Newer routers provide DHCP address reservations which is a way to provide static addresses and yet let the DHCP server manage them. DHCP works well if there's a single DHCP server, but not so much when there's more than one and they supply conflicting information. It can also be hard to get a system to have the same

address across reboots with DHCP if not properly managed.

IPv6 addresses can be compressed and rewritten as follows.

## IPv6 Address Compression

- Leading zeros can be omitted (trailing zeros cannot)
- A single instance of continuous zeros can be replaced with a double colon ::
- When decompressing remember there are supposed to be 8 hextets of 4 hexadecimal characters each

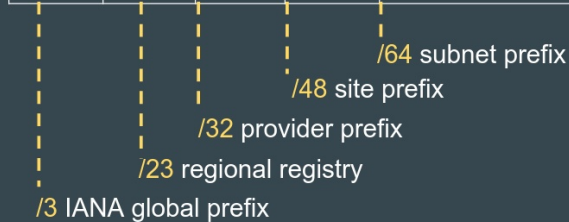
global address	2001:0EB8:00C1:2200:0001:0000:0000:0331 /64
leading zeros	2001:0EB8:00C1:2200:0001:0000:0000:0331 /64
continuous zeros	2001:0EB8:00C1:2200:0001:0000:0000:0331 /64
compressed	2001:EB8:C1:2200:1::331 /64

With IPv6, each host interface has a link-local address and multiple routable addresses. These addresses are the global unicast address which is routable on the Internet, the Unique local address which is similar to a IPv4 private address behind a NAT and an Anycast address which is the same unicast address on multiple devices.

## IPv6 Global Unicast Address

- The network portion of the global unicast address is hierarchically structured

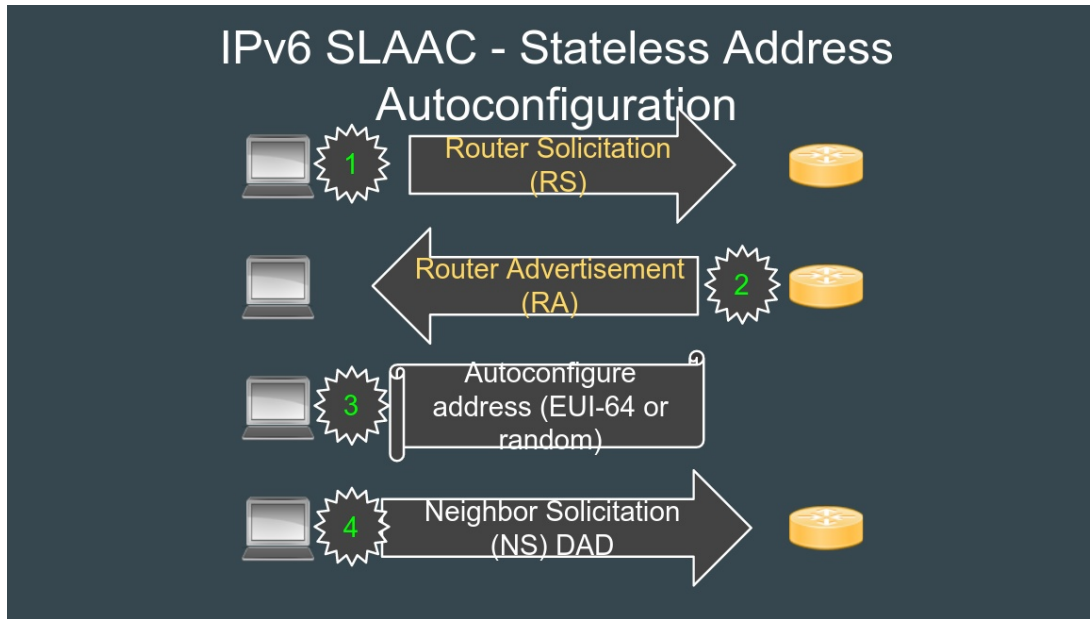
Global Routing Prefix		Subnet	Interface ID		prefix
2001	00A1	2233	0001	0800:27FF:FE00:0008	/64



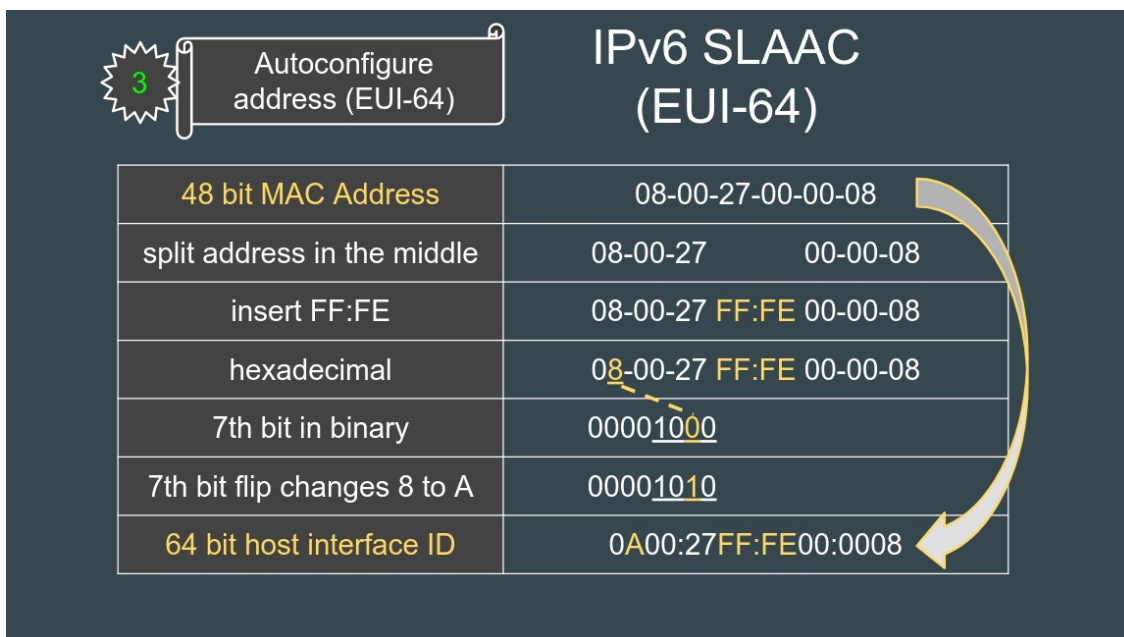
- /64 subnet prefix
- /48 site prefix
- /32 provider prefix
- /23 regional registry
- /3 IANA global prefix

With IPv6, DHCP is unnecessary because of stateless autoconfiguration. This is a mechanism whereby routers send out "router advertisements" (RAs) that contain the upper 64 bits of an IPv6 address, and hosts generate the lower 64 bits themselves in order to form a complete address.

The lower 64 bits of an IPv6 address is comprised of an Extended Unique Identifier by Stateless Address Auto-Configuration (SLAAC) using the MAC address as defined below.



The bottom 64 bits of an IPv6 address are generated from a MAC address by flipping a bit and adding the bits ff:fe in the middle. So the Ethernet MAC address 08-00-27-00-00-08 results in 0A00:27FF:FE00:0008 as the lower 64 bits of an IPv6 address, called the "interface identifier" in IPv6.



SLAAC also uses neighbor solicitation and neighbor advertisement to request and provide information about client node reachability similar to IPv4 ARP.

# IPv6 SLAAC

RFC 4861 Neighbor Discovery - SLAAC - ICMPv6			
Message	Source address	Destination address	Type
<b>Neighbor Solicitation (NS)</b> <ul style="list-style-type: none"> <li>Similar to ARP in IPv6</li> <li>check host availability</li> <li>check for DAD</li> </ul>	link-local or unspecified address :: /128 if duplicate address detection (DAD)	all-solicited nodes multicast FF02::1:FFxx:xxxx	135
<b>Neighbor Advertisement (NA)</b> <ul style="list-style-type: none"> <li>response to (NS),</li> <li>used to announce a link-layer address change</li> </ul>	link-local address FE80::x	link-layer address, or all-nodes multicast FF02::1	136
R=router flag-neighbor unreachable S=solicited flag-response to NS O=override flag-update address			

Here is a summary of IPv6 address prefixes. Note that any address beginning with FE80 is your link-local address which is the logical equivalent to the IPv4 private address behind a NAT. A good take-away from this is to understand that if your ISP is supporting IPv6 and your router is configured for IPv6, then you will see an FE80 link local address on each device on your network as well as a Global Unicast address and an Anycast address.

## IPv6 Addresses

IPv6 Address	Purpose
::/0	all networks - used for default route
::1/128	loopback - similar to 127.0.0.1
::/128	unspecified address - no address yet
FE80::/10	link-local
FF02::1	all-nodes multicast
FF02::2	all-routers multicast
FF02::1:FFxx:xxxx	All-solicited nodes multicast - autoconfiguration and neighbor discovery (similar to ARP)

An IPv4 client behind a NAT had to have a port forwarding rule on the router to allow a connection from the public internet to connect to a client system with a private address. With IPv6, every client system on your LAN network already has its own global address that can be accessed from the Internet. This is possible because most ISPs assign /48 network prefixes to subscribers' sites (the End Users' networks). Because all IPv6 networks have /64 prefixes, a /48 network prefix allows 65,536 LANs in an end user's site.

My cable provider grants me an IPv6 network size that is typical for an ISP and they provide a /56 which is a prefix that allows for 256 LAN segments. To get an idea of what this means, my ISP provides me with  $2^{56}$  IPv6 addresses or a total of 4,722,366,482,869,645,213,696 addresses. If this seems extreme, realize that there are a total of 340,282,366,920,938,463,463,374,607,431,768,211,456 IPv6 addresses in the total 128 bit IPv6 address space. Here is my WAN port configuration:

If it seems wasteful that my ISP grants me a /56 prefix, consider that many ISP's grant a /64 and a single /64 can have 18 quintillion nodes, but a customer network may only have less than a hundred computers connected.

On most modern networks, there are nodes that are capable of both IPv4 and IPv6 communications. In many cases, the nodes have IPv4 addresses either statically assigned or leased with DHCP, while the IPv6 nodes have only a link-local address. If there isn't an IPv6-enabled router active on the LAN, then the IPv6 nodes do not receive a global unicast address. However, increasingly, LANs are actively using IPv6 and the end nodes have both an IPv4 and an IPv6 address and operate in dual-protocol mode.

Because IPv4 is still required for reachability to much of the Internet, it is unlikely that you will have an IPv6-only network today. You can add IPv6 to a LAN, but you must keep IPv4 active because most services still use only IPv4. Therefore, dual-protocol operations are preferred, but the network would have all the IPv4 ARP traffic on the LAN, in addition to the IPv6 traffic.

IPv6 operates differently than IPv4 on a LAN. IPv6 does not use broadcast message delivery. IPv6 strives for increased efficiency using only unicast, multicast, or anycast.

Link local addresses are used to communicate over a single physical or logical subnetwork, such as an Ethernet. These addresses start with fe80 and are extensively used for IPv6's internal house keeping.

Site Local addresses are the IPv6 equivalent of the RFC 1918 private address space in IPv4. However, the IETF found the situation where different organizations use the same address space undesirable, so they created "unique site local" addresses where everyone takes a randomly selected block out of the IPv6 address space starting with FC.

A multicast address is a group address, so every packet sent to a multicast address is received by all members of the group. Multicast addresses start with FF and can be used for applications where

several hosts must receive the same information at the same time, such as live video broadcasts and also for autoconfiguration and discovery.

An IPv6 anycast address is an address that is assigned to more than one interface. Typically, the address belongs to different nodes. A packet that is sent to an anycast address is routed to the nearest interface that has that address. Anycast

**IPv6 Address Types**

Type	Purpose	Prefix
Global Unicast	routable on the Internet	2000:: /3
Link Local	local link/network (only)	FE80:: /10
Multicast	sending to groups (no broadcast)	FF00:: /8
Unique Local	routable on a LAN (private address)	FC00::/8 FD00::/8
Modified EUI-64	converts 48 bit MAC to 64 bit host ID	7th bit flipped + FF:FE inserted
Autoconfiguration	stateless address autoconfiguration (SLAAC)	RS RA NS NA (ICMPv6)
Anycast	shared address, used to send to nearest available device interface	syntactically identical to unicast (identified as anycast) 2001:DB8:1::/64 (all zeros ID)



addresses can be used as part of a route sequence.

So, what about IPv6 security? The idea was to give IPv6 security a big push by making IPsec support mandatory. IPsec encrypts each individual packet, so it can be applied to all IP traffic, unlike the widely used SSL, which only works on top of TCP. However, for a number of reasons, it's very difficult to build IPsec support into applications, so it never gained much real-world use except as a mechanism to implement VPNs. And despite the fact that IPsec was developed for IPv6 or at least with IPv6 in mind, it also works with IPv4. All in all, IPsec can't be considered a security advantage for IPv6.

With IPv4, there will generally be a NAT device that functions as a simple firewall by blocking incoming sessions (although there are ways to trick NATs into allowing them). Since there are more than enough public addresses to go around in IPv6, there is no NAT. This might be scary.

The good news is that because the IPv6 address space is so large, randomly scanning for systems that are vulnerable is completely infeasible. The story goes that at the height of the self-propagating malware explosion a few years ago, an unpatched Windows system would be infected faster than it could download the necessary security updates. With IPv6, that is simply impossible: even with a billion infected hosts each scanning a billion IPv6 addresses per second, it takes more than a hundred million years to scan just the IPv6 address space that's given out to ISPs right now, which is about 0.01 percent of what's available.

Many software firewalls that run on the to-be-firewalled host itself only support IPv4 and don't get in the way of IPv6 packets at all. The Windows and Mac OS built-in firewalls don't have this problem, but if you're doing any firewalling on Linux or BSD (or command line firewalling with Mac OS X), make sure that your services are firewalled over IPv6, too. Of course, systems that use domain filters, GeoIP filters and IPS/IDS end up blocking a lot of undesirable traffic to begin with.

Systems with IPv6 connectivity decide whether to use IPv4 or IPv6 to reach a destination by consulting the DNS. Communication over the Internet requires addresses, but we generally work with domain names. The DNS takes care of the difference by having one or more A (address) records that contain an IPv4 address associated with a given name. If a system also has an IPv6 address, this is added to the DNS with an AAAA (quad-A) record. Hosts that only have IPv4 connectivity ignore the AAAA records, but dual stack hosts ask the DNS for both the A and AAAA records. They will then generally prefer to connect to a destination over IPv6 if possible, and use IPv4 if there's no AAAA record in the DNS or connecting over IPv6 doesn't work.

So, in summary, with IPv6 NAT appears to be obsolete. In addition with IPv6 every system on VLANs where you have configured IPv6, all have global IPv6 addresses. As for DDNS, you would have to find a DDNS that would support an AAAA record since that is what IPv6 uses. I am reading where ISPs have DHCP for their IPv4 addresses, but their IPv6 addresses are static. I don't have enough experience to comment on this.