


# IPv6 Briefing – A Spitfire Review



Graham Lewis  
Director of IP Engineering  
Briefing Paper – Rev 1  
Spitfire Network Services  
February 2014

[www.spitfire.co.uk](http://www.spitfire.co.uk)

**SPITFIRE**<sup>®</sup>  
VOICE • INTERNET • DATA

## IPv6 briefing

Whilst much of the attention given to IPv6 is the extended address, size it is important to understand that it is a totally different protocol to IPv4. The two will happily run alongside each other on a LAN and over WAN circuits but the packets travel as though “ships in the night” never interacting or being aware of each other.

Thus devices with only IPv4 addresses cannot send or receive packets to IPv6 hosts, and vice versa, they are simply not listening out for packets from the other protocol. While some invention has been made to create IPv4 to IPv6 “gateways”, these are often inadequate, incomplete and unsatisfactory. Almost all successful implementations to date, are made by “dual stacking”, whereby hosts run both IPv4 and IPv6 software at the network layer.

Most implementations of Unix/Linux have had IPv6 enabled for many years and it has been in Windows since Windows 7 and in OS X since v10.1

## Protocol version 6

IPv6 is a whole new protocol with lots of features that have been born out of experience with IPv4, however the main reason for people adopting IPv6 right now is the much larger address size that IPv6 has. An IPv4 address is 32 bits long which gives  $2^{32}$  possible values for IP addresses and all those values have now been assigned or squandered. An IPv6 address is 128 bits long (16 bytes) which gives  $2^{128}$  possible values for IP addresses.

## How big?!

Everyone has heard the superlatives - that is over 54 million IP addresses for each square meter of the earth's surface.

10 million trillion times the number of grains of sand on all the beaches in the World etc



## Writing it down

For the benefit of humans the 128 bits are represented in colon-hexadecimal which consists of eight 16-bit hexadecimal numbers separated by colons like this:

**2A02:0B90:DC01:1777:0000:0000:0000:1001**

IPv6 address on one of our routers.

It's still a big number so the rules say that to make it easier to read you can remove leading zeros and take out a consecutive block of zeros (only once) to give you this:

**2A02:B90:DC01:1777::1001**

Using the same rules the loopback address

**0000:0000:0000:0000:0000:0000:0000:0001**

is abbreviated to

**::1**

You be pleased to know that subnet masks are not used instead the slash notation from IPv4 CIDR is used

e.g. **21cd:53::/64**

and

**2A02:B90:DC01:1777::1001/128**

*Never  
heard of  
hexadecimal ?*

It is a way of representing 4 bits in a single alpha-numeric character from 0 to the letter F.

Binary	Hex
0000	0
0001	1
0010	2
0011	3
0100	4
0101	5
0110	6
0111	7
1000	8
1001	9
1010	A
1011	B
1100	C
1101	D
1110	E
1111	F

## Address types

There are no broadcast addresses in IPv6 only Unicast, Multicast and Anycast.

## Link-Local IPv6 Unicast Addresses

Link-Local addresses are for use on a single link (LAN). Link-Local addresses have the following format:



On a local network, IPv6 systems assign themselves an address automatically called a link-local address. All link local addresses have the same network identifier, a 10-bit prefix of 11111110 010 followed by 54 zeros which gives the network address fe80:0000:0000:0000 /64 or fe80::/64

You will find that all modern operating systems have already assigned you an IPv6 link local address – go on have a look run “ipconfig” on Windows or “ifconfig” on a Unix/Linux/OSX machine and you will probably see a link-local IPv6 address already configured.

```
Z:\>ipconfig
Windows IP Configuration
Ethernet adapter Local Area Connection:
    Link-local IPv6 Address . . . . . : fe80::9123:a597:ca7d:d5b7%22
    IPv4 Address. . . . . : 192.168.56.1
    Subnet Mask . . . . . : 255.255.255.0
    Default Gateway . . . . . :
```


If you see this on your PC, it means is that the machine has IPv6 enabled but hasn't heard from an IPv6 enabled router yet, so it only has a Link-local address.

Devices on the local LAN could still use IPv6 to communicate with each other but just would not get any traffic off the local network.


## IPv6 Global Unicast address

To get onto the Internet using IPv6 you will need a Global Unicast address.

The general format for IPv6 Global Unicast addresses is as follows:



The recommendation from RIPE, is that customers are assigned a /48 IPv6 address. So when you apply for an IPv6 allocation this is what your IPv6 addresses will look like:



## The global routing prefix

The global routing prefix will be assigned by the ISP. Spitfire have been assigned **2A02:B90/32**

## The subnet ID

The subnet ID allows a customer 65,536 /64's or 65,536 separate LANs on a site. So if you were assigned **2a02:b90:5003/48** you can create subnets of

2a02:b90:5003:0001/64  
2a02:b90:5003:0002/64  
2a02:b90:5003:0003/64 etc up to 2a02:b90:5003:FFFF/64

The subnet ID is under your control and you can make it whatever makes sense to you in your addressing scheme. Often people will use DC01 etc to designate different data centres, put VLAN ID's in, or companies have branch office codes that make sense in organising a consistent addressing scheme.

## The Interface ID

[www.spitfire.co.uk](http://www.spitfire.co.uk)

IPv6 is structured so that on a LAN segment there are 64 bits of network prefix (a /64) known as the Interface ID or host bits. Each LAN segment therefore has a possible 18,446,744,073,709,551,616 IPv6 addresses. So you could change your IP address every second for 584 billion years (that's a lot of typing).

It may seem a waste to allow the Interface ID to be 64 bits long but it allows for easy configuration because the device can create a Network ID based upon the 48 bits of its MAC address in the Network ID called Stateless Address Auto Configuration (SLAAC).

## Stateless Address Auto Configuration (SLAAC)

In most cases on a LAN, it is expected that customers will use SLAAC where the interface ID is generated based on the MAC address of the interface card or by a pseudo-random generator. It is called stateless, because unlike DHCP, you don't have to run a special server that keeps track of which device has which IP address and you don't have to do any DHCP style administration to deploy or maintain it. In SLAAC, devices pick an IP address, check that it is not already in use and then contact a router to get their global routing prefix. It really is very easy to deploy IPv6 on the LAN using SLAAC.

The method of using the MAC address of the interface card to create the interface ID is a known as creating an EUI-64 (Extended Unique Identifier). There are three steps to creating an EUI-64

1. Split the 6 byte MAC address into two halves of 6 hex digits each.
2. Insert "FFFE" between the two
3. Invert the 7<sup>th</sup> bit of the interface ID.

So MAC address **0015.1234.ABCD** becomes EUID-64 **021512FFFE34ABCD**

Obviously the network devices will do this for you automatically, just recognise that if you see FFFE in the middle of your interface ID this is almost certainly an EUI-64.

Note that in particular Microsoft OSs by default use a randomly generated interface ID and regularly change the interface ID (and hence your IPv6 address). This is regarded as a privacy measure.

The way that IPv6 addressing has been designed, is that that an IPv6 address is essentially in two 64 bit halves – a network ID and a host ID. If you use a MAC derived EUI-64 identifier as the interface ID then this remains the same when you move networks, only the first portion of your IPv6 address will change.

A device may roam across different networks but the interface ID will remain the same and hence is trackable and gives away more information than you may want to allow.

This may not be desirable and hence RFC3041 defines these privacy extensions to SLAAC, where the interface ID is randomised. If you run a Windows network, you may want to have static IPv6 addresses, so you need to turn this privacy behaviour off. Other OS have the option but it is disabled by default.

Once a device creates its link local address, it uses IPv6 Neighbour Discovery (ND) protocol to check that it is unique (checking that its network ID is unique). If the link local address is unique, it assigns it to its interface and starts sending Router Solicitation messages. Any routers on the LAN advertise themselves using Router Advertisement protocol, using multicast address FF02::2. When it receives a reply from a router on the LAN, it learns the network prefix for the LAN from the router reply. It will prepend the network prefix it learns from the router to its interface ID to create its global unicast address.


Therefore enabling IPv6 on a LAN, for most people, is as simple as configuring an IPv6 address on their router and enabling IPv6 routing. Most modern OS will then use SLAAC to join the IPv6 network.

## Manually created interface IDs

You don't have to use SLAAC. You might want to manually configure IP addresses in which case you can use any whacky scheme you like, one of them is to map existing IPv4 addresses onto an IPv6 address.

## IPv4-Mapped IPv6 Address

A type of IPv6 address that holds an embedded IPv4 address is defined. This address type is used to represent the addresses of IPv4 nodes as IPv6 addresses. The format of the "IPv4-mapped IPv6 address" is as follows:



Note when you are doing this, you are not turning an IPv6 address into an IPv4 address or vice-versa, you are just giving an aide memoir in the IPv6 address. It was designed to help people remember the where they had left their IPv6 addresses, in relation to their IPv4 addresses, during transition phase from IPv6 to IPv4.

## Get creative with IPv6 addresses

Within the part of the address that is under your control, you can use the letters to make meaningful or vanity addresses. Facebook amusingly use FACE:B00C in their IPv6 addresses.

```
;; QUESTION SECTION:
;facebook.com.      IN      AAAA

;; ANSWER SECTION:
facebook.com. 900 IN AAAA 2a03:2880:2110:df07:face:b00c:0:1
```

## IPv6 and DNS

IPv6 addresses have the DNS resource record type AAAA, so once you have put your servers on IPv6 addresses you need to get DNS AAAA records added and the associated reverse IPv6 records (ipv6.arpa).

## Test tools

Most tools, like ping and traceroute, have their IPv6 equivalents. Some OS have renamed them as ping6.

## What you can do with it?

You need to keep this picture in your mind, of there being an IPv4 internet and an IPv6 Internet and the two cannot directly exchange data with each other, except through hosts that are dual stacked. So if you have IPv6 only hosts, you will find that there is not a lot out there (less than 2% of Internet hosts have IPv6 addresses). Other hosts on IPv6 are limited at the moment, while there is little commercial drive to force people to adopt IPv6.

A lot of Google services are available on IPv6, including Google search at [ipv6.google.com](http://ipv6.google.com) and they run a handy IPv6 test tool. You can stream films from Netflix over IPv6 and keep in contact with friends on Facebook over IPv6. However many business and applications are not on IPv6 yet until they are forced to adopt it by demand or lack of IPv4 addresses.



## IPv6 “brokenness”

Many applications are still not IPv6 ready and some applications can be expected to have IPv6 bugs. Other than individual applications having bugs the biggest cause of problems is once you have IPv6 enabled most OS prefer to use IPv6 and will try to use an IPv6 address before failing over to IPv4.

If your IPv6 connectivity is incomplete you will experience that as slow to login to things and slow responses from servers etc.



## About the author



Graham Lewis MSc MIET MBCS is Director of IP Engineering at Spitfire Network Services Ltd a leading UK ITSP.