All about IP addresses.

Introduction
Yo' Momma is the Plain Old Telephone System.
Standing In a Broken Phone Booth, Tuppence in My Hand...
Pensylvania 128.91.65.000
  Let's pretend the world is a happy convenient place:
  I can take this… Show me the numbers!
Back to reality…
  Class-A
  Class-B
  A quick interlude to explain those slashes some more
  Class-C
What happens in the 1990s stays in the 1990s.
What works, what doesn't work and what just seems to work?
Where the nice IPs at?
Some notes on private IPs and IPv6

Introduction

As more of our admin folks have access to Superadmin, I thought I would clear up a few issues about IP addresses, and answer some of the questions I have been asked more than once.

This document won't win any Institute of Electrical and Electronics Engineers awards for accuracy but that isn't the intent; if I end up with a readable document that explains the landscape as we see it, I'll be content. If I am in the mood I may flag things with a little (*) that are actually inaccurate but convenient enough explanations, and would get you laughed at if you ever went to a party with a gaggle of Cisco engineers. That being said, Cisco engineers never go to parties; so you are probably safe.

If you are lucky, I will avoid using pictures, and if I do I will steal clip-art from other places - if you are unlucky, I will try and recreate it myself in Paint.

Yo' Momma is the Plain Old Telephone System.

Traditionally, IP addresses have been explained in much the same way as phone numbers. This is reasonable, so let's go with that for now, and then tinker with it later.

Let's look at a phone number. In the long, dark and distant past, phone numbers were quoted as "town/exchange-name, number" - So "Pennsylvania 6-5000" or in my case, the first home number I remember as a kid was, "Fleetwood 3735".

What happened in the real world is that most calls were made within the local exchange, so in this case, Penn Station in NYC, or Fleetwood. If you wanted to call a local number you didn't quote the town, you just dialled 65000 or 3735 and it would connect you. Once operators became obsolete and everything was done by mechanical dials, the number of numbers in a number (points for that phrase, please!) became important. In the Penn Station exchange vicinity you dialled 5 digits, allowing 99999 possible numbers, in Fleetwood you had just 4, allowing 9999 (*) - This was called Local Exchange dialing because everything was handled by your local telephone exchange (or groups of interconnected exchanges).

If you wanted to dial another exchange, you would dial extra numbers and these calls would be routed over trunk-lines which were long distance cables connecting interchange-exchanges to other ones. they made these into 3 letter mnemonics using the letters on the telephone dial so Penn Station becomes been PEN (736), and Fleetwood FLE (353) - It didn't actually work like that except in places like London and NYC - But we'll stick with the theory because it's useful.

So now we have a telephone number made up of two parts - The local exchange (PEN or 736), and the phone number. To call me back in 1975 you would have dialed 353-3735 and Glenn Miller's phone number would be 736-65000 - Except it wouldn't because Glenn Miller is a filthy liar and lyrical purposes he added the 6, and it's actually 736-5000 (which is actually the Hotel Pennsylvania).

Standing In a Broken Phone Booth, Tuppence in My Hand...

As time went on and more and more people got phones this caused problems because there simply weren't enough numbers in some areas. Area codes got split, and with the changeover from mechanical to digital routing in the 90s, it became much easier to have numbers all over the place. Voice over IP providers would buy large allocations of "spare numbers" anywhere in the world and sell them on to their subscribers. With VoIP it was perfectly possible for somebody in Fleetwood to to have a Penn Station telephone number even if they'd never even heard of the place. With the addition of vanity numbers, mobile phones, telemarketing operations and large-scale VoIP providers like Skype and Google Voice the whole thing became a mess. Somebody having a number in Fleetwood no longer comes with any certainty that are even in the United Kingdom.

Pensylvania 128.91.65.000

For the rest of this document, we are going to be using **IPv4** terminology. This is because IPv6 is horrible, and for some reason IPv5 was stolen by aliens. The evolution of IP addresses is much the same as the evolution of the pre-digital phone system, which we will now call POTS (the Plain Old Telephone System) because you are now experts, and are therefore entitled to use abbreviations to sound cooler. When people talk about IP (which, incidentally, stands for Internet Protocol) they are pretty much always talking about IPv4.

Let’s pretend the world is a happy convenient place:

It’s convenient to think of an IP address in very similar ways to phone numbers because they kind of look the same. As we all know (although some people know it’s wrong), an IP address looks like A.B.C.D - Four dotted numbers from 0 to 255.

If we are making a rather geeky simile, we could think of this as Country.State.City.Number - It’s not true but it’s worth pretending and thinking about it that way for a while.

This would make Pennsylvania 6-5000 something like US.NYC.PEN.5000

In our lovely convenient fairyland, you can see how nice and easy it would be to ban everybody living around Penn Station by banning US.NYC.PEN.* (In computing terms, * usually means “everything”) - We would call that a subnet and in fairytale IP world, subnets are separated by dots. In the real world, nothing could possibly be so simple.

I can take this… Show me the numbers!

In reality an IP address isn’t A.B.C.D at all, it is actually a string of 32 binary digits - If you don’t know binary feel free to skip this, and be sure to tut and shake your head as you do it.

A typical IP address, for instance 194.72.6.51 actually looks like this:

11000010010010000000011000110011

Because we humans like to group things together in convenient chunks, we tend to write it as:

```
11000010.01001000.00000110.00110011
```

That being said, we could just as well write it as

```
1100.0010.0100.1000.0000.0110.0011.0011
```

Which could translate back to decimal as: 12.2.4.8.0.6.3.3

Obviously, we can divide these 32 bits into chunks (subnets) of any size we want. We don't have to use 8 bits within each dot just because that is how people prefer to write them down; and these days, we rarely do.

Back to reality…

In the olden days before IP addresses became a somewhat scarce commodity, allocations made some degree of sense.

(This explanation about classes is a fib, it's not true but it's true enough for this document. That being said, don't quote it in a pub quiz or some bright spark will spend an hour telling you how wrong you are. I will switch to the much better /nn representation quite quickly to get around that.)

IPs were handed out as Class-As, Class-Bs, Class-Cs and Host-Addresses.

Class-A

A Class-A was a number in the first block, and then the owner was free to use the rest of it as they wanted - For example, some famous Class-As are:

3.0.0.0	General Electric.
9.0.0.0	IBM
17.0.0.0	Apple
18.0.0.0	MIT

There's a list at http://en.wikipedia.org/wiki/List_of_assigned_/8_IPv4_address_blocks if you are curious, it provides a strange insight into the early Internet.

You will see these things called /8's - So somebody may say that "IBM has a slash-eight" - This is shorthand for the first 8 binary digits of the address, which is much easier when we stop subdividing into 8 bits - It’s easier to say a “/10” than “A Class-A and a couple of extra bits”.

In the early days of the Internet, they handed these things out like social science degrees; you can see from the Wikipedia page that the US DoD has 11 of the things. People who have them hold onto them for dear life (except apparently for Stanford Uni who made a big deal of giving lots of theirs back to the community) and there is no means to forcibly take them back.

Class-B

A large number of Class-As exist simply to be split into 265 Class-Bs which can then be allocated to smaller organisations. If we go to the state of Pennsylvania as opposed to the station, we can see that The University of Pennsylvania (UPenn) has the following B’s (or /16s):

128.91.0.0
130.91.0.0
158.130.0.0
165.123.0.0
170.212.0.0
120.166.0.0

To add some extra terminology here, we actually represent these as (for example) 128.91.0.0/16 - Which means that everything after the first 16 bits are for the owner to do what they want with - This means that they can theoretically use any address from

128.91.0.0 to 128.91.255.255 - This gives 265 x 256 address, which is 65,536 individual addresses.

This means that UPenn has 6 allocations of 65,536 addresses, giving them 393,216 addresses. Given that they have 25,000 students you could well wonder why they need 10 times more addresses than students, even if they allocated each an individual address (which they don’t). it’s all down to hoarding and bad housekeeping - Something you can smugly point out when people whine about the addresses running out.

Back to a real-world situation, if we wanted to ban the entire UPenn from a site, we would have to ban the 6 individual /16s.

A quick interlude to explain those slashes some more

I am throwing around /16s and /8s (and soon /24s) a lot - The reason for the terminology is all to do with binary logic and binary netmasks but you don't need to know any of that, you just need to know it refers to how many bits of an IP address we can use when it's written in binary.

The number after the slash basically says "We have to block off that many bits" so let's say we have our previous IP address:

```
11000010010010000000011000110011
```

Writing this as:

```
11000010010010000000011000110011  /8
```

(or 194.72.6.51/8) means we can use anything after the first 8 bits - Those 8 will ALWAYS stay the same… So we can have addresses like:

```
1100001  0000000000000000000000000
1100001  0111111111111111111111111
1100001  0101010101010101010101010
1100001  1111000011110000111110000
1100001  1110010110110110111001101
```

Moving up to a Class-B or /16 we'd write this as:

```
11000010010010000000011000110011  /16
```

(or 194.72.6.51/16) and that means we can use anything after the first 16 bits - Again those 16 will be fixed… So we can have things like:

```
1100001001001000  0000000000000000
1100001001001000  1111111111111111
1100001001001000  1010101010101010
1100001001001000  1110011110101000
1100001001001000  0001110001110001
```

Now let's pick a /29 - Which is a perfectly legitimate allocation - This would be written as 194.72.6.51/29 or:

```
11000010010010000000011000110011 /29
```

This allows us only the following 8 combinations:

```
11000010010010000000011000110 000 (194.72.6.48)
11000010010010000000011000110 001 (194.72.6.49)
11000010010010000000011000110 010 (194.72.6.50)
11000010010010000000011000110 011 (194.72.6.51)
11000010010010000000011000110 100 (194.72.6.52)
11000010010010000000011000110 101 (194.72.6.53)
11000010010010000000011000110 110 (194.72.6.54)
11000010010010000000011000110 111 (194.72.6.55)
```

Class-C

Typically Class-Bs were allocated to large companies, large telecom companies and Internet Service Providers (ISPs) and occasionally, people who used to collect them to win bets. Class-B's tended to be split into organisational units or customers and in the earlier days of the Internet the most convenient way of doing this was to split it off at a Class-C or a /24 - An ISP backbone provider like British Telecom, which mostly sold to smaller ISPs would historically have split some of its customer Class-B space into 256 C's and when they signed up a new customer they would re-sell them as many as they needed.

It's this historic snapshot that creates a lot of the terminology we use today and a lot of the misconceptions about how the Internet works in terms of addressing.

What happens in the 1990s stays in the 1990s!

A lot of networking stuff was written back in the 90s and a lot of people learned a lot of bad habits. I wrote my own guide to network management back then that was meant to parody a lot of the common beliefs; the odd thing was that although it did, it was also mostly accurate.

Back in the 90s not many people ever considered that IP addresses would become scarce, so not many people worried about allocating addresses in blocks of /24s. To give some perspective, at one point I had about 6 machines using up 4 Class-C's, which was 1024 addresses. These days people pay about $10 a month per single extra address from ISPs and hosting companies.

We talk about things like “subnets” and we have tools to look at and ban subnets. When people talk about subnets they generally assume a Class-C and tools are written to work with these. This can sometimes accidentally work out quite well but can also be a disaster.

What works, what doesn’t work and what just seems to work?

Biggish ISPs and backbone providers still tend to think in terms of /24s and will allocate /24s to smaller ISPs. The smaller ISPs will then resell them in as small an allocation as they can get away with.

A cybercafe with 12 terminals doesn’t need 12 IP addresses these days since pretty much all of the user terminals will be hidden behind a connection sharing system that will present them all as a single IP host-address to the outside world. In reality customers tend to try and get as many IP addresses as they can, so a cybercafe will probably ask for an extra IP for a webserver, and another for remote management.

So let’s look a real world example - Let’s say that we see a whole bunch of fake scammer users using our website from a Nigerian cybercafe on a certain IP address. From everything we have said, we cannot safely assume that the whole /24 (or subnet, as we call them) is tainted by association. At most we can probably assume a /30 (just a few addresses). But in practice we will nearly almost find that the entire /24 is full of scammers. Why?

The reality is more social than technical. An ISP that will sell IP space to a cybercafe known to host scammers will likely as not sell IP space to other dodgy businesses. They may well have some good customers in the same area as well, a garage or a couple of churches; but these aren’t likely to have legitimate users using an American dating site and anyone using the IPs belonging to the legitimate companies is likely to be using them illicitly.

In this case banning the /24 works purely by accident - But it still works. Whether this is safe to rely on is questionable. The easy answer is “Hell no!” but in practice is it almost always “Why not, if it works?”

Where the nice IPs at?

I'll close this by adding to the confusion and mess. Basically, we go back to the idea of a nice neat hierarchically-structured IP system like: US.NYC.PEN.65000 - And then we take a huge great crap all over it.

It started off quite well-meaning. Back in 1981, IBM had their /8 which was 9.0.0.0 and then they started using it - They allocated some to partners, some to subdivisions, some to customers, research groups, anyone who needed them and within a few years this block had gone all over the world. Bits of IBM were sold off, research groups moved around and there became very little in the way of sense of geography to the blocks. 9.9.9.1 may be in the USA, and 9.9.10.1 may be in China. When I worked for Hughes, we had a bunch of IPs in space, which adds an extra level of geographical complexity because there's not actually a controlling authority for that - But as ever, I digress.

IBM is amazingly well managed and they are a mess so think what has happened to the other blocks. As IP space is running out, people have started trading them. ISPs will try and hang onto their allocations, but often legal issues mean that they actually have to formally re-allocate chunks of their IP space, eventually fragmenting it all over the place again. Whilst it's not strictly legal, trade in IP space has also become a big thing. Since no laws can stop you selling a company along with its assets, this is usually how IP space ends up getting traded. The last /16 I was involved in selling went for about $250,000 for some perspective on how large a business this is now. Spam companies burn up thousands of addresses a week and since most security mechanisms still work in terms of subnets based on a /24 policing and moderating this is not keeping up.

This means that you simply can't put vast amounts of faith in where databases think IP addresses are, and you certainly can't put much faith into assuming that two similar blocks may be related to each other.

Hosting companies, Virtual IPs, Botnets, TOR exits and anonymizers all complicate this even more, but that's another story.

Some notes on private IPs and IPv6

IPv6 is an abomination against the gods, and that is all you need to know. It’s only good for fridges and freaks.

Private IPs allow people to hide a huge amount of IP addresses behind a router. So if we look back at our Nigerian cybercafe which has 12 terminals, even though there is only one external IP address those machines need to be able to talk to the router, the printer, and the hundreds of cellphones that are connecting to the WiFi hotspot.

This is achieved by having the local network (LAN / Local Area Network) on private addresses which are specifically not (and cannot be) routed by the internet as a whole. The ones you will see most often are the 10.0.0.0/8 and 192.168 .0.0/16. There is also 172.16.0.0/20 which nobody ever uses because hardly anybody understands netblocks that arent /8, /16 or /24 and a special one, 169.254.0.0/16 which we just don't talk about in polite company. You may ask why there is so much private IP space available when it’s all private; that’s probably a good question too.

As a rule of thumb, big companies use the 10 block whereas home users and small businesses use the 192.168 block. I remain convinced that this is purely a historical side-effect of the default IP addresses of commercial vs. consumer routers. If you pay half a million dollars for a router, you want it to have much more addresses.

It’s worth noting as a side-observation that most router companies can’t decide whether people’s home networks should be on router default address 192.168.0.1 or 192.168.1.1 Belkin and SMC use 192.168.2.1 because they are even more of a pain than most.