IP address management

P A Roberts and S Challinor

Anything connected to the Internet needs, at some point, to pick up a globally unique address to identify it among everything else on the Internet. In the case of today's Internet, these addresses are Internet protocol version 4 (IPv4) addresses and they are limited in number. Responsibility for management of this limited resource is delegated from a global level through regional organisations and ultimately to individual Internet service providers. In order to gain maximum use out of the available number of addresses against the background of explosive growth in the size of the Internet, this valuable resource has to be carefully managed to ensure that there is enough to go round. This paper discusses what IP addresses are, the global organisations used to manage IP addresses, and how BT manages its IP addresses.

1. Introduction

Internet protocol (IP) addresses are the unique numbers used to identify individual connections to the Internet; they are also a limited resource and the careful management of these IP addresses is essential to the running of the Internet. BT must be able to efficiently manage its own and its customers' IP addresses and it must be able to demonstrate good management of IP addresses in order to qualify for more IP address space. Without more new IP addresses BT would not be able to expand its IP network or launch new IP services. It is therefore of central importance to BT that it manages its IP addresses well.

This paper starts with a description of what IP addresses are before going on to outline the global IP address environment and how the global address space is managed through regional registries and organisations such as Internet service providers (ISPs) for distribution to end users. This is followed by an examination of the process for obtaining addresses for an end user which highlights some of the rules which have to be followed by the ISP as part of this process. The penalties for abuse of this system are punitive and could seriously affect the ability of an ISP to conduct its business.

2. IP addresses

An IP address consists of 32 bits, written as four 8-bit numbers (thus in the range 0 to 255) separated by a decimal point. Every IP address is invisibly divided into two parts, the first part identifies a particular network, and the second part identifies a particular machine on that network (called a 'host'). The point at which the first part, the network part, of the address ends and the second part, the host part, of the address begins depends upon the size of the network. If the network is very large, with many hosts, the division point is nearer the beginning of the address than it would be in a smaller network with fewer hosts. For instance, with the address 132.146.100.100, the first two numbers (132.146) may refer to a single large network, leaving address space from 0.0 through to 255.255 (i.e. approximately 65 000 addresses). Whereas the address 194.72.10.10 may be the address of a host on a smaller network represented by the numbers 194.72.10 which leaves only 255 addresses on that network (0 to 255).

The exact position of this division point is denoted by the 'subnet mask', which indicates which part of the address represents the network, and which represents the host. The 132.146.100.100 address would have a subnet mask of '255.255.0.0', the 255s indicating which numbers represent network, and the zeros indicating which numbers represent the hosts on that network, the 194.72.10.10 example would have the subnet mask '255.255.255.0'. What is in fact happening here is that every bit in the 32 bits of the address that represents the network is being masked with a '1', and every bit that represents a host is being masked with a zero. Thus the subnet mask for the first example is actually 11111111.11111111.00000000.00000000, but for convenience this is converted to decimal and written as 255.255.0.0.

The first address on a network cannot be given to a host on that network, since it is used to represent the network as a whole. For example, the address 194.72.10.0, with a subnet mask of 255.255.255.0 is the first address on that network, and is used only to represent that entire network. This is called a 'network' address. Similarly the last address on a network cannot be given to the host, since this is used as a broadcast address. Any traffic sent to a broadcast address is effectively addressed to all the hosts on that network. The address 194.72.10.255 with subnet mask 255.255.255.0 is a broadcast address.

The original Internet plan was to break the whole address space into different sections to allow for the creation of these different sized networks; this plan is no longer used, but the terminology from it still is — it is presented below:

Class A: 0-127.X.X.X
126 Class As exist (0 and 127 are reserved)
16 777 214 hosts on each Class A

Class B: 128-191.X.X.X
16 384 Class Bs exist
65 532 host on each Class B

Class C: 192-223.X.X.X
2097 152 Class Cs exist
254 hosts on each Class C

Class D: 224-247.X.X.X
Class Ds are multicast addresses

Class E: 248-255.X.X.X
Reserved for experimental use

The above scheme proved very wasteful of address space, since the smallest network that could be created had 255 addresses. Networks with less than 254 machines would lock up address space, preventing any other network from using it. The scheme of classless interdomain routeing (CIDR) or ‘variable subnetting’ was devised to solve this. With a variable subnet, the network part of the address is extended into the last 8 bits of the address, creating space for more networks, but with less hosts since there are less bits left with which to address hosts. For example if a network address were represented by the first 26-bits of the 32 bits available, that would leave only 6 bits for the machines on that network.

If the subnet mask for a 26-bit network was written out in binary, it would look like this:

11111111.11111111.11111111.11000000

Converting each of the 8 binary digits to decimal gives the subnet mask for a 26-bit subnet as:

255.255.255.192

Under the original addressing plan, the Class C network 194.72.10.0 would have 255 hosts, but there would only be one network, and any unwanted addresses would be wasted. If this is converted to a 26-bit network, the one class C network is swapped for 4 smaller networks. To see how this is done, only the last 8 bits of the address will be dealt with for the moment. The first 2 bits have been borrowed for use by the network, and with 2 bits only four different combinations can be represented: 00, 01, 10 and 11. Now since these two bits are actually the first two bits in an 8-bit binary number, they actually represent the numbers 128 and 64. Thus the four combinations 00, 01, 10 and 11 are actually written as 0, 64, 128 and 192. The hosts may be written with 6 binary digits, giving 64 combinations, representing the decimal numbers from 0 through to 63. Since the last two bits of the 26-bit network description and the 6 bits of the host’s description are to be written as one number, the host’s value (0 to 63) must be added on to the decimal value of the network (0, 64, 128 or 192), giving the following four networks to use:

Network 1: 194.72.10.0 through to 194.72.10.63

Network 2: 194.72.10.64 through to 194.72.10.127

Network 3: 194.72.10.128 through to 194.72.10.191

Network 4: 194.72.120.192 through to 194.72.10.255

All have the subnet mask 255.255.255.192.

Since the first address of any network is the network address, and the last is always the broadcast address, there are only 62 addresses available for use by hosts on each of the above networks. A Class C network capable of supporting 254 hosts has been exchanged for four smaller networks, each capable of supporting 62 hosts. It can be seen that the subnet mask is critical in understanding what an IP address is actually representing, and even with the subnet mask it is far from intuitive to determine whether the address is a machine address, a network address or a broadcast address.

The boundary between the network part of the address can be moved further into the last 8 bits, allowing for more and more networks, each with fewer and fewer hosts. For example, a 30-bit network would have only 2 bits available for addressing the machine. Since 2 bits only allows four combinations, and two of those have to be used for network and broadcast addresses, there can only be two machines on this network — which is perfect for addressing a point-to-point link. Thus 30-bit networks are very common within an IP network.

Table 1 shows the number of addresses available with a given size network, where the network size is specified in both netmask notation and the ‘/’ (slash) notation which identifies the number of bits belonging to the network part of the address.

3. Reserved ranges of IP addresses

The following ranges have been reserved for use on private networks and are known as private addresses or RFC 1918 addresses, these ranges are not routeable across the Internet:

Table 1 Comparison of netmask and / (slash) notations.

Mask	/	Hosts	Networks	Host bits	Network bits	Addresses
255.255.255.255	32	1	N/A	0	0	1
255.255.255.254	31	0	2147483648	1	31	2
255.255.255.252	30	2	1073741824	2	30	4
255.255.255.248	29	6	536870912	3	29	8
255.255.255.240	28	14	268435456	4	28	16
255.255.255.224	27	30	134217728	5	27	32
255.255.255.192	26	62	67108864	6	26	64
255.255.255.128	25	126	33554432	7	25	128
255.255.255.0	24	254	16777216	8	24	256
255.255.254.0	23	510	8388608	9	23	512
255.255.252.0	22	1022	4194304	10	22	1024
255.255.248.0	21	2046	2097152	11	21	2048
255.255.240.0	20	4094	1048576	12	20	4096
255.255.224.0	19	8190	524288	13	19	8192
255.255.192.0	18	16382	262144	14	18	16384
255.255.128.0	17	32766	131072	15	17	32768
255.255.0.0	16	65534	65536	16	16	65536
255.254.0.0	15	131070	32768	17	15	131072
252.252.0.0	14	262142	16384	18	14	262144
255.248.0.0	13	524268	8192	19	13	524288
255.240.0.0	12	1048574	4096	20	12	1048576
255.224.0.0	11	2097150	2048	21	11	2097152
252.192.0.0	10	4194302	1024	22	10	4194304
255.128.0.0	9	8388606	512	23	9	8388608
255.0.0.0	8	16777214	256	24	8	16777216
254.0.0.0	7	33554430	128	25	7	33554432
252.0.0.0	6	67108862	64	26	6	67108864
248.0.0.0	5	134217726	32	27	5	134217728
240.0.0.0	4	268435454	16	28	4	268435456
224.0.0.0	3	536870910	8	29	3	536870912
192.0.0.0	2	1073741822	4	30	2	1073741824
128.0.0.0	1	2147483646	2	31	1	21474836484

10.0.0.0/8 or 10.0.0.0 — 10.255.255.255
172.16.0.0/12 or 172.16.0.0 — 172.31.255.255
192.168.0.0 or 192.168.0.0 — 192.168.255.255

4. IP address co-ordination

4.1 Global co-ordination

The growth of the global Internet from its beginnings as the US Government's experimental ARPAnet soon identified a need to control the allocation of addresses to the many networks being connected. This task fell to an organisation called the Internet Assigned Numbers Authority (IANA) which was set up and funded by the US government. IANA controlled the management of global IP address allocation by delegating the responsibility to regional Internet registries (RIRs), each responsible for a geographic area:

- the American Registry for Internet Numbers (ARIN) [1] covers North and South America, the Caribbean and some parts of Africa — ARIN took over this role from Network Solutions in December 1997,

- Réseaux IP Européens (RIPE) [2] covers Europe, North Africa and parts of the Middle East and was formed in 1989 — a number of working groups exist within RIPE, operating in areas of common interest to the Internet community which are run by volunteers from participating local Internet registries,

- the Asia Pacific Network Information Centre (APNIC) [3] covers the Asia Pacific region.

There have been proposals to set up regional registries for Africa and Latin America which would possibly be known as AFRINIC [4] and LATNIC [5].

IANA was financed by the US government which, due to the global nature of the network, felt the Internet community should take responsibility for financing a top-level management organisation. As a result of this, the Internet Corporation for Assigned Names and Numbers (ICANN) was formed in October 1998. ICANN [6] is a non-profit corporation that not only has responsibility for IP address space allocation, but also looks after protocol parameter assignment, domain name system management, and root server system management functions. ICANN is already operational and it is in the process of assuming the responsibility for the technical management and policy development functions that have been assigned to it. The transition to ICANN from the US Government and IANA is scheduled to be completed by September of 2000. The current structure of ICANN (see Fig 1) is composed of a board of nineteen directors consisting of nine 'at-large' directors, nine nominated by ICANN's three supporting organisations, and the President/CEO. The nine current 'at-large' directors are serving initial terms and will be succeeded by at-large directors selected by an at-large membership organisation which will seek to represent the broadest possible spectrum of the Internet community. The three ICANN supporting organisations consist of the:

- Protocol Supporting Organisation (PSO) which is made up of representatives of the Internet Engineering Task Force (IETF), World Wide Web Consortium (W3C), International Telecommunications Union (ITU) and European Telecommunications Standards Institute (ETSI) — the PSO will advise the ICANN board with respect to matters relating to the assignment of parameters for Internet protocols and technical standards,

- Address Supporting Organisation (ASO) which consists of representatives from the ARIN, RIPE and APNIC RIRs — the ASO will advise the ICANN board with respect to policy regarding IP addresses,

- Domain Name System Supporting Organisation (DNSO) which consists of representatives from top-level domain registries, commercial and business entities, gTLD registries, ISPs and connectivity providers, non-commercial domain name holders, registrars, and trademark, other intellectual property and anti-counterfeiting interests — the DNSO will advise the ICANN board with respect to policy issues relating to the domain name system.

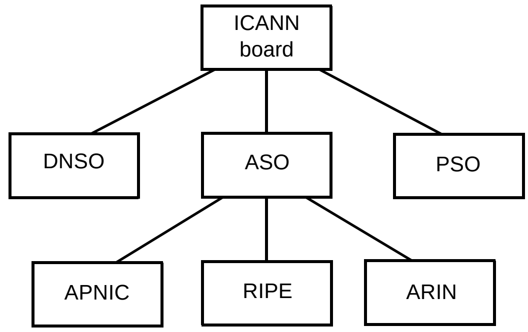

Fig 1 The structure of ICANN with details of the ASO part.

4.2 Regional Internet registries

As outlined in the above section there are currently three RIRs, each of which has responsibility for a different area of the world, each RIR works towards the following goals:

- fair distribution of address space to ensure there is enough address space to go round, and ensuring that end users are considering the use of private address space and that assignments to users are only for the short term needs of their network,

- conservation to prevent the practice of stockpiling of addresses by ensuring that each application for an addresses is genuine and that assignments to users are only for the immediate and short-term needs of their network which are actually going to happen, rather than for long-term projects or growth which may not come to fruition,

- aggregation to create a hierarchical distribution of globally unique address space which will permit the simplification of routeing information — this is achieved by the RIRs being allocated very large blocks of address space, typically /8s, and in turn allocating large blocks of addresses to the LIRs for onward assignment to their customers, which results in a relatively small number of blocks representing large areas of the world and each LIR will only announce the larger blocks to other ISPs, thus helping to reduce the size of the global routeing tables,

- registration, including the provision of public registry or database, to ensure uniqueness and assist troubleshooting — details of all addresses allocated to an LIR are entered into a public database as well as the addresses assigned to end users by an LIR (contact information for these address ranges is also included and this information may be accessed using the 'Whois' tool [7]).

All of the above registries are run on a non-profit basis, charging the participating local Internet registries to cover their administrative costs. Regional registries are also responsible for issuing Autonomous System (AS) numbers and ‘in-addr-arpa’ reverse DNS delegations.

4.3 Local Internet registries

Most LIRs are Internet service providers who require addresses to assign to customers connected to their network. Some organisations that have very large networks or that connect to multiple ISPs can become enterprise registries which, like an ISP, will have an IP address allocation but it is only for use within the organisation. The LIR will announce the block of addresses allocated to it to the Internet using the BGP exterior routeing protocol. BT runs an enterprise registry for its internal address space and a public registry for its ISP BTnet and a number of registries in the joint ventures and Concert.

5. IP assignment requests

When an LIR first starts up it is given a limited number of addresses as its first allocation and the first assignments from this block have to be sent to the RIR for approval. As the LIR demonstrates competence in this activity, then it will be given more power to assign address space independently. In order to achieve this, APNIC and RIPE allocate each LIR an assignment window; this is a representation of the number of IP addresses the registry may assign to an end user without first seeking approval from the RIR. The assignment window is usually expressed as a network mask in ‘/’ notation. The use of the terms ‘assign’ and ‘allocate’ are important and should not be confused — an RIR allocates addresses to an LIR for onward assignment to end users.

An LIR has to follow the rules laid down by the RIR when they assign IP addresses to an end user. Only an LIR can assign IP addresses to an end user; an LIR cannot assign an IP address to a third party for onward assignment to the end user. Failure to abide by the rules could result in the LIR having its assignment window reduced to zero, meaning that all requests for assignment would have to be processed through the RIR, resulting in delays to the customer’s IP address application process.

An end user applying for address space from an LIR will be expected to provide the following information prior to an assignment being made:

- overview of end-user organisation describing the structure of the organisation requesting the addresses and in which part of that organisation they would be used, including information on:

  — how the addresses will be distributed among the various departments,

  — the geographical set-up of the organisation,

  and, if the organisation is a subsidiary of another organisation or it has any subsidiaries, address usage of all parts of the organisation will have to be considered as part of the application,

- contact details for a technical representative which will be entered into the RIR’s public database,

- contact details for an administrative representative, who must be someone on the site where the IP addresses are used — this data will also be entered into the RIR’s public database,

- an addressing plan showing how the organisation is using all public address space that has been assigned to it in the past and what its future plans are for that space (see Fig 2),

- an addressing plan showing how the requested IP addresses will be used over the next two years (see Fig 3) — RIPE expects that 80% or more of the IP addresses requested will be used within two years of assignment while ARIN expects 50% or more of addresses requested will be used within one year of assignment.

The LIR is expected to keep records of this information in case of an audit by the RIR and forward to the RIR if the request is greater than their assignment window. In Europe this would be submitted to RIPE using the RIPE-141 form, an example of which is shown in Fig 4; the majority of the fields can be filled in from the information provided by the end user. The form is intended to present enough information to the RIPE hostmaster to make an informed decision about the request.

Once the form is completed it is e-mailed to RIPE via a syntax checking robot and is allocated a ticket number which is mailed back to the requester. The request is then placed in a wait queue before being passed to a member of the RIPE staff known as a hostmaster, who may then enter into an e-mail conversation with the requester if they require any further details. It is in the interest of the requester and the end user to ensure all information presented is correct and clear in order to avoid any unnecessary delays in this stage of the application. The hostmaster has to consider the request in terms of the goals of the RIR/LIR system. Once the request is authorised the LIR will assign the address space from their existing block. Before the end user starts using the addresses the assignment should be entered into the RIR’s public database.

[CURRENT ADDRESS SPACE USAGE TEMPLATE]#

			Addresses Used			
Prefix	Subnet Mask	Size	Current	1-yr	2-yr	Description
194.72.193.0	255.255.255.192	64	28	34	60	Department A
194.72.193.64	255.255.255.224	32	10	21	28	Department B
194.72.193.96	255.255.255.224	32	8	13	27	Department C
194.72.193.128	255,255,255,128	128	0	0	0	Unused
	255.255.254	512	317	429	480	Department D
		640	363	497	595	Totals

Fig 2 The ‘Current Address Space Usage Template’ is used to show how addresses are currently used within the organisation.

#[ADDRESSING PLAN TEMPLATE]#

Relative		Addresses Used				
Prefix#	Subnet mask	Size	Immediate	1yr	2yr	Description
0.0.0.0	255.255.255.128	128	80	95	100	Department A
0.0.0.128	255.255.255.224	32	12	17	25	Department B
0.0.0.160	255.255.255.224	32	0	15	28	Department C
0.0.0.192	255.255.255.240	16	7	8	10	Department D
0.0.0.208	255.255.255.240	16	10	14	14	Department E
0.0.0.224	255.255.255.240	16	11	12	12	Department F
		240	120	161	189	Totals

An entry is made for each physically separate subnet in the network, subnets are considered to be physically separate if there is an IP router between them.

Prefix: or Relative Prefix: The network address of the subnet; if this Addressing Plan was being used to request addresses, then a relative prefix would be used where the first entry would be 0.0.0.0 and all subsequent entries would be sequential from that depending on the size of each subnet. For example, 0.0.0.0 with a mask of 255.255.255.0 would be followed by a prefix of 0.0.1.0, as shown above.
Subnet mask: The subnet mask of each entry.
Size: The size of the subnet based on the subnet mask.
Addresses Used: In these three fields, show the current and estimated usage for the next two years.
Current: The number of network interfaces currently used in the subnet.
Immediate: The number of addresses the organisation will use immediately.
1-yr: The number of addresses to be used in one year.
2-yr: The number of addresses to be used in two years.
Description: Used to specify a short description of the use of the subnet in the user’s organisation.

Fig 3 The ‘Addressing Plan Template’ is used to show how the organisation proposes to use the requested addresses.

6. LIR’s address space

How the assignments are made from the block of addresses allocated is left to the LIR to control. This allows some scope for network summarisation and simplification of the routeing within the LIR’s network. The ideal situation would be to allocate customers connected to each access node from specific ranges of addresses and maybe to have a large dial network from another block.

For example, say the LIR has the following block allocated to it — 194.72.0.0/16; this could be broken down:

194.72.0.0/18 = Customers in the north

194.72.64.0/18 = Customers in the midlands

194.72.128.0/18 = Customers in the south

194.72.192.0/19 = Dial network

194.72.224.0/20 = Infrastructure addresses

194.72.240.0/20 = SPARE

Every customer assignment made to customers connected to the LIR’s network in the south would therefore come from the 194.72.128.0/18 block.

X-NCC-RegID: *each LIR has a unique RegistryID, this is used to identify who has sent the RIPE141 form in.*

I. General Information
#[OVERVIEW OF ORGAISATION TEMPLATE]#
As provided by the end user and modified or added to by the person from the LIR making the request.

#[REQUESTER TEMPLATE]#
Details of the person from the LIR making the request.

name:
organisation
country:
phone:
fax-no: (optional)
e-mail:

#[USER TEMPLATE]#
Details of the person from the end user's organisation requesting the addresses.

name:
organisation:
country:
phone:
fax-no: (optional)
e-mail:

#[CURRENT ADDRESS SPACE USAGE TEMPLATE]#
Details of all current public address usage by the end user organisation — see the Addressing Plan diagram (Fig 3) for further details.

			Addresses Used			
Prefix	Subnet Mask	Size	Current	1-yr	2-yr	Description

II. The Request
#[REQUEST OVERVIEW TEMPLATE]#
Summary information about the request

request-size: *information from the request addressing plan*
addresses-immediate: *information from the request addressing plan*
address-year-1: *information from the request addressing plan*
address-year-2: *information from the request addressing plan*
subnets-immediate: *information from the request addressing plan*
subnets-year-1: *information from the request addressing plan*
subnets-year-2: *information from the request addressing plan*
inet-connect: *whether the organisation is currently connected to the Internet and if so how*
country-net: *the ISO 3166 two letter country code where the addresses will be used*
private-considered: *whether the organisation has considered the use of private addresses for the requested networks*
request-refused: *whether the organisation has ever been refused an IP address request in the past and if so who by and why*
PI-requested: *if the request is for Provider Independent address space*
address-space-returned: *whether the organisation is planning to return address space as part of this request*

#[ADDRESSING PLAN TEMPLATE]#
Details of the proposed use of the requested addresses — see the Addressing Plan diagram (Fig 3) for further details.

Relative			Addresses Used			
Prefix#	Subnet Mask	Size	Current	1-yr	2-yr	Description

III. Database Information
#[NETWORK TEMPLATE]#
Information that will be entered in the RIPE database about the organisation requesting the address space.

inetnum:
netname:
descr:
descr:
country:
admin-c:
tech-c:
status:
mnt-by: (optional)
notify: (optional)
changed:
source: RIPE

#[PERSON TEMPLATE]#
If contact information for the network administration persons entered in the Network Template is already in the RIPE database then this should be checked to see if it is up to date. For each person specified in the network template for which there is no entry in the RIPE, one can be created using the following template.

person:
address:
address:
address:
e-mail:
phone:
fax-no: (optional)
mnt-by: (optional)
notify: (optional)
nic-hdl:
changed:
source: RIPE

#[TEMPLATES END]#
IV. Optional Information

Any other information which may help the hostmaster with the request can be added here.

Fig 4 Example of RIPE 141 form, with comments added in italics.

However, the blocks will not be used evenly and eventually one will run out or there will no longer be a large enough free range of addresses to satisfy an assignment. The spare addresses will therefore have to be used for the region which has run out of addresses:

194.72.0.0/18 = Customers in the north
194.72.64.0/18 = Customers in the midlands
194.72.128.0/18 = Customers in the south
194.72.192.0/19 = Dial network
194.72.224.0/20 = Infrastructure addresses
194.72.240.0/20 = More customers in the south

This is not a problem until one of the other blocks fills up and the LIR has to request more addresses from the RIR. In Europe, this is done by sending a list of all assignments from the allocation to RIPE who will audit them and ask for further details of some of them as evidence that the LIR is maintaining the address space properly. RIPE will not entertain a request for address space until 80% of the current allocation is used which in our example can cause further problems if one of the blocks is under-utilised. In order to fully use the block, the LIR may choose to spread one block over two regions:

194.72.0.0/18 = Customers in the north
194.72.64.0/18 = Customers in the midlands and some customers from the north
194.72.128.0/18 = Customers in the south
194.72.192.0/19 = Dial network
194.72.224.0/20 = Infrastructure addresses
194.72.240.0/20 = More customers in the south

As this progresses the advantages of network summarisation begin to erode; when the LIR is then allocated a further block of addresses, the process will continue. While it is not easy to strictly maintain the ideal model of address management across the network it should be considered as the network grows.

7. Address management tools

In order to keep track of their address space LIRs need to keep a top-level record of how their allocations have been split up — this is needed in the event of an audit by the RIR. A second layer of information is required which consists of detailed information about the assignments from the blocks of addresses, that have been used. Returning to the previous example the top-level record would appear

194.72.0.0 Allocation 1
  194.72.0.0/18 Block for customers in the north
  194.72.64.0/18 Block for customers in the midlands
  194.72.128.0/18 Block for customers in the south
  194.72.192.0/19 Assigned to the dial network on RIR ticket number NCC#1234567
  194.72.224.0/20 Block for infrastructure addresses on RIR ticket number NCC#1234567
  194.72.240.0/20 Block 2 for customers in the south

The 194.72.192.0/19 range of addresses is directly assigned to the dial network and therefore does not have any underlying records of addresses — all of the other blocks contain underlying records:

194.72.0.0 Allocation 1
  194.72.0.0/18 Block for customers in the north
    194.72.0.0/25 Manchester Metalworkers Co RIR ticket number NCC#123444
    194.72.0.128/25 Scottish Tartan Ltd RIR ticket number NCC#127811
    etc
  194.72.64.0/18 Block for customers in the Midlands
    194.72.0.0/27 Birmingham Tyres
    194.72.1.0/24 Shakespeare Hotels RIR ticket number NCC#122222
    194.72.3.0/29 Little Company
    etc
  194.72.128.0/18 Block for customers in the south
    194.72.0.0/23 London Bridge Co RIR ticket number NCC#124545
    194.72.3.0/24 Tower Bridge Co RIR ticket number NCC#126544
    194.72.22.0/28 Brighton Beach
    etc
  194.72.192.0/19 Assigned to the dial network on RIR ticket number NCC#1234567
  194.72.224.0/20 Block for infrastructure addresses on RIR ticket number NCC#1234999
    194.72.224.0/24 Block for connecting customers in the North
      194.72.224.1 Northern customers access router
      194.72.224.2 First customer on the northern customers access router
      194.72.224.254 Last customer on the northern customers access router
    194.72.225.0/24 Block for connecting customers in the midlands
      194.72.225.1 Midland customers access router
      194.72.225.2 First customer on the midland customers access router
      194.72.225.254 Last customer on the midland customers access router
    194.72.225.0/24 Block for connecting customers in the south
      194.72.225.1 Midland customers access router
      194.72.225.2 First customer on the southern customers access router
      194.72.225.254 Last customer on the southern customers access router
    194.72.226.0/24 Block for special customers connections
      194.72.226.0/30 Very fast customer
      194.72.226.4/30 BGP customer
      etc
  194.72.240.0 /20 Block 2 for customers in the south
    194.72.240.0/27 South coast hotels

Apart from the structure of the records there are a few things to note about the above records. Only the larger blocks have RIR assignment ticket numbers associated with them; this is because they are larger than the LIR's assignment window and have therefore had to have passed through the RIR's approval process. The assignments from within the blocks are not sequential — there is no requirement for this from the RIR. The assignment of the infrastructure top-level block has been approved by the RIR not the assignments from it — each sub-block of the infrastructure block would have been represented as a separate line in the infrastructure requests addressing plan. An up-to-date copy of the addressing plan for the LIR's network infrastructure and other addresses used by the organisation, with current utilisation levels, should also be kept for when the organisation to which the LIR belongs requires more addresses. In the above example, the dial network addresses would fall into this category.

While there are a few commercial and freeware IP address record applications for aggregating address space, the maintenance of top-level records and addressing plans can adequately be carried out using spreadsheets.

All assignments and associated records should be registered in the RIR's public database and the tool for querying it is called 'Whois'. It can be accessed via Web pages [7] or it can be downloaded as an application and run on a local machine.

8. Address conservation

A number of methods have been employed in order to extend the use of IP version 4 address space, two of which are examined below.

8.1 Network address translation

Using software commonly available in routers and firewalls, it is possible to effectively hide whole networks using private addresses behind one or a pool of public addresses. There are two forms of network address translation (NAT). Port address translation (PAT), which is also known as 'NAT overload', works by utilising the port numbers available in the TCP/IP packet, and allocating a different port on the outside address to an internal address. Port numbers are represented by a 16-bit field in the IP packet header giving a total of 65 536 different ports available; however, many of these are reserved for use by specific applications. The number of inside addresses, which can be mapped to a single outside address, is reduced to a few thousand depending on the manufacturer's software. Extensive use of this method could go a long way towards the conservation of IPv4 address space. Unfortunately not all applications will work through PAT; for example, an address for an e-mail server which requires a globally unique address will be sent to the single outside address. To overcome this the more basic form of NAT allows a one-to-one mapping of an inside address to an outside address allowing the implementation of these services.

8.2 Dynamic allocation of addresses

Another method of saving public addresses is to dynamically allocate an address to the end user for the duration of their session connected to the network. This method is extensively used by an ISP's dial-up customers and allows a large number of users to use a relatively small pool of address space. Another example of this is dynamic host configuration protocol (DHCP), which allows end-user hosts to be configured automatically with an address for the duration of their use. When a DHCP client starts up it searches for a DHCP server and obtains set-up information, including an IP address; this can be used in combination with NAT and has the advantage of allowing changes to network configuration, which affect all clients, to be carried out centrally at the server.

9. Futures

Given the explosive growth of the Internet it is conceivable that a global IP address shortage could become a reality. Work is well under way towards the deployment of the next generation of the Internet based on IP version 6. IPv6 extends the size of addresses to 128 bits or 2^{128} addresses which equates to 340,282,366,920,938, 463,463,374,607,431,768,211,456 addresses. IP version 6 also adds a number of features including enhanced methods of dynamic address allocation, class-of-service-based routeing and enhanced security features. The problem this creates is when and how to migrate or whether the existing addresses can be extended indefinitely.

10. Conclusions

This paper has shown some of the issues BT has to face in order to successfully manage its vitally important IP address space within the overall global framework. The consequences for getting this activity wrong could adversely affect BT's ability to operate in the new global IP environment.

References

1 http://www.arin.net

2 http://www.ripe.net

3 http://www.apnic.net

4 http://www.afrinic.org

5 http://www.latnic.org

6 http://www.icann.org

7 http://www.ripe.net/cgi-bin/whois

Paul Roberts joined BT as an apprentice on leaving school in 1980, after which he spent a number of years working on exchange construction before moving on to wideband provision.

He has spent the last three years working in Technical Services at the Internet Operations Centre in St Albans where among other support work he manages the RIPE registry for BTnet.

He graduated with an Upper Second Class Honours degree in Electronics and Computing from the Open University in 1994.

After graduating from Nottingham University in 1993 with a 2.1 Honours BSc in Computer Science, Simon Challinor joined the Adastral Park team which was embarking on BT's first Internet Service Provider venture — BTnet. Working as part of BT's core IP design team, he has played a leading role in designing BT's UK Internet backbone, most recently leading the IP design activity for the Colossus project which provided BT with the largest IP backbone network in the country. With backbone IP network growth being crucial to BT's intentions to be a major player in the broadband, mobile, leased line and dial IP service arena, he is currently designing next year's network developments to keep ahead of the exponential growth in IP traffic.