Understanding IP Address Prefixes

Objective of this Note

This Note is targeted at NRENs that are moving to establish their unique identities on the net so as to be recognised as bona fide members of the global research and education community. Such NRENs must acquire their autonomous system numbers (ASNs) as well as independent IP addresses. IP addresses permit the unique identification of the device that is sending data or is the target recipient of data by other devices on the Internet. Application for the autonomous system numbers and independent IP address space requires a prior understanding of the rationale and meaning of IP addresses: this note provides a quick technical primer to the subject.

The Note has been developed by the UbuntuNet Alliance for Research and Education Networking as part of a set of tools to support the growth of REN activity in the continent with the intent of empowering and creating global equity for Africa-based educationists and researchers.[1]

1. Historical note

Prior to 1994 IP addresses were allocated to organisations in one of three sizes, identified as Class A, Class B and Class C. Class A addresses were characterized by having the first bit set to zero; Class B by having the first two bits set to 10; and Class C by having the first three bits set to 110. For Class B allocations, the first two octets (i.e. the first 16 bits), were reserved for specifying the network address and the last two octets for specifying the host address. For Class C allocations, the first three octets (i.e. the first 24 bits), are reserved for specifying the network address and the last octet for specifying the host address. Many older institutions in Africa and elsewhere have one or more Class B and/or Class C allocations that were made during this early era in the history of the Internet.

In September 1993, RFCs 1518 and 1519 appeared and established an entirely new policy and practice for the allocation of IP address blocks and routing practice. This was done to control and limit the growth of the routing tables upon which routers base their routing decisions. The new paradigm was called Classless Inter-Domain Routing (CIDR), and is also sometimes referred to as supernetting.

Prior to the advent of CIDR, blocks of IP addresses were allocated without regard to the disposition of adjacent blocks. Under the CIDR practice, allocations were structured hierarchically and in such a way that assignments by any ISP to its customers would be made by the ISP from a large allocation to the ISP, and in such a way that these assignments could be aggregated within the ISP's network, reducing the number of independent entries in the routing tables.

This Note describes one key aspect of the CIDR paradigm, viz the specially selected address blocks, which are called prefixes, in which IP addresses are allocated to organisations. The

[1] We acknowledge with thanks and appreciation the role that has been played by Dr Duncan Martin, Dr Lishan Adam, Mr Victor Kyalo, and Dr F F Tusubira in the development of these notes. Note UA 152 was developed by Dr Duncan Martin.

exposition is limited to IPv4, as the same concepts and concomitant practices apply to the distribution of IPv6 address blocks.

2. Prefixes

Internet Registries are organizations assigned the responsibility of managing the allocation of IP addresses. Regional Internet Registries (RIR's) allocate chunks of address space as contiguous ranges of addresses, each of which constitutes what is called a "prefix". Here are the formal definitions.

An ***IPv4 address*** is a number between 0 and 2^{32}.

An ***IP prefix*** is a contiguous range of addresses containing a number of addresses that is a power of 2 and with the smallest address being a multiple of that same power of 2. Consequently the largest address is one less than the next multiple of the same power of 2. This definition applies to both IPv4 and IPv6.

Example: The $2^5 = 32$ numbers between (24756×2^5) and $(24757 \times 2^5 - 1)$ inclusive constitute a prefix. In binary representation the low end and high ends of the prefix are respectively:

00000000 00001100 00010110 10000000 and
00000000 00001100 00010110 10011111.

In quad-decimal representation (translating each binary octet block into the equivalent decimal number and separating blocks by a "."), they are:

0.12.22.128 and
0.12.22.159

The adjacent chunk of 32 numbers comprises those between (24757×2^5) and $(24758 \times 2^5 - 1)$ inclusive, and also comprise a prefix. What about these 64 numbers viewed together as a single chunk? Would that chunk be a prefix? The low end number is (24756×2^5), which can be written as (12378×2^6), while the high end number is $(24758 \times 2^5 - 1)$, which can be written as $(12379 \times 2^6 - 1)$. So the aggregate of these two prefixes is itself a prefix. Note that this would not have been the case if we had started with an odd number, say 24755, instead of 24756.

3. The CIDR or "slash" notation for prefixes

When written out as 32-bit binary numbers, the 32 addresses in our example prefix differ from each other only in last five bits, and together cover the entire range from 00000 to 11111 of these last five bits. All 32 addresses have the same bit values for the leading 27 bits (viz 00000000 00001100 00010110 100). This provides another way of thinking about prefixes. An ***"slash n" IPv4 prefix*** is a contiguous range of $2^{(32-n)}$ addresses, all of which, when represented as binary numbers, have the same bit values for the leading n bits, and cover all possible distinct values for the trailing 32-n bits. Note that the number n, which is called the ***mask length*** and can range from 0 to 32, specifies the number of leading bits that are the same, while $2^{(32-n)}$ is the number of addresses contained in the prefix. The longer the mask, the smaller the prefix!

One can specify a prefix by specifying the address at the low end of the range and also specifying the mask length. The standard notation is to specify the low end address followed by the mask length, separating the two by a slash character, as in 0.12.22.128/27. The mask length is read as “slash 27”.

4. Subdividing prefixes

Once an institution has been allocated a prefix, say a /18 prefix (which has 2^{14} = 16,384 addresses) the network planner will want to implement his/her numbering plan, and will subdivide this prefix into several sub-prefixes for numbering campus and building subnets.

The principle is: *Any prefix can be split into two equal prefixes*. This applies not only to the allocated /18, but to the each of the sub-prefixes produced by such splits. For example, the /18 can be split into two /19s. One of these could be split into two /20s for use at the institution’s two satellite campuses, while the other /19 is for the Main Campus. On each campus, the campus prefix can be split and re-split to produce prefixes that are suitably sized for the various building subnets. Generally a small prefix is reserved for network interfaces.

5. Why prefixes? The aggregation imperative

The reason prefixes were introduced is that they enable IP numbering resources to be distributed in a way that controls and limits the growth of the Internet routing tables. In the routing tables, all the addresses within a single prefix are represented by a single entry. An Internet Service Provider (ISP), having been authorised by AfriNIC[2] to act as a Local Internet Registry (LIR), assigns to its various customers sub-prefixes of the ISPs own prefix. Within its network, the ISP can route traffic to each customer according to the customers’ sub-prefixes, but the ISP’s entire network can still be represented by the ISP’s single prefix in the routing tables. This is referred to as aggregation.

The minimum allocation size is /22

AfriNIC will not normally allocate IPv4 prefixes that are smaller than one with a mask length of 22. Such a prefix has $2^{(32-22)} = 2^{10}$ = 1,024 addresses. Organisations that cannot justify an allocation of a /22 or larger address block prefix have no option but to use a prefix assigned by their ISP or other LIR. (For IPv6, the smallest allocation size to an end-user organisation is a /48).

/24 is the minimum prefix size for inclusion in the routing tables

When ISPs announce routes to each other, they will not normally accept prefixes longer than /24. Consequently, if a subnet might have to be announced separately – e.g. in countering a DOS attack - that subnet should not be assigned a longer mask than a /24.

[2] AfriNIC is the RIR for Africa and the Indian Ocean region. See http://www.afrinic.net.