IPv6 Address Space Management

Paul Wilson
Raymond Plzak
Axel Pawlik

Document ID: ripe-261
Date: 31 October 2002

Abstract

This document provides the management process for IPv6 global unicast address space whereby address allocations are made from a single global pool according to a "sparse allocation" algorithm. This allocation process will maximise aggregation of address space, ensuring that most ISPs retain a single prefix as they grow, and avoiding the address space fragmentation which results from the current IPv4 allocation technique. This document also describes the registration process and the administration of the IP6.ARPA domain.

Table of Contents

- 1.0 Overview
- 2.0 Common Address Pool
- 3.0 Sparse Allocation Algorithm
- 4.0 Allocation Request Process
- 5.0 Avoiding Fragmentation
- 6.0 CAP Contention
- 7.0 Review of Allocation Process
- 8.0 Common Registration Service
- 9.0 Reverse DNS (ip6.arpa) Requirements
- 10. IANA Considerations

- Author's Addresses

1.0 Overview

Under the current system of management of global IP address space, Regional Internet Registries (RIRs) are responsible for allocation of address space to organisations within their respective geographic regions.

RIRs receive address space periodically from IANA, and then manage that address space as a "pool" from which subsequent allocation are made to organisations within their regions. RIRs individually use various allocation techniques within their respective pools of address space (including sparse allocation techniques), in an attempt to ensure that subsequent allocations to ISPs are able to be aggregated.

Under this current system, aggregation of allocated address space is limited by several factors. These factors include the requirement for RIRs to utilise their existing pool of addresses to a level of 80% prior to requesting more address space from IANA, as well as the relatively small size of the address pools held by RIRs at any one time. Over a period of several years, a single large ISP may receive many discontiguous allocations from its RIR.

This document provides a management system for IPv6 which avoids these problems by delegating to the RIRs collectively a single large pool of IPv6 address space which will be managed under a new allocation system. Under this system, an RIR wishing to make an allocation will receive that allocation from the common pool, rather than from a regional pool already held by that RIR. Furthermore, the allocation it receives from that pool will be selected so as to maximise the room for future expansion of the allocation as a single aggregatable block.

2.0 Common Address Pool

For the purposes of this document, the source of address space described above is referred to as the Common Address Pool (or CAP). The CAP will be established by the IANA as a single allocation of address space for management by the RIRs under this proposed framework. See IANA considerations below.

Allocations from the CAP will be selected according to a Sparse Allocation (or "binary-chop") algorithm, designed to maximise aggregation of address blocks allocated. This algorithm is described in the next section.

The administration of the CAP will be conducted jointly by the RIRs, through a "registration service" which is described below. For the purposes of this document, the term CRS is used to refer to the Common (Address Pool) Registration Service.

3.0 Sparse Allocation Algorithm

An allocation of IPv6 address space is defined uniquely by a start address (expressed as an IPv6 address) and an address block prefix (expressed as the integer number of bits in the network mask, between 0 and 64). Examples are:

```
2000::/3
2001::/16
2002:200::/23
2002:298::/35
```

Under the sparse allocation system, the start addresses for successive allocations are generated according to a simple algorithm, as illustrated below.

For example, within a 6-bit address space, the first 16 start addresses would be as follows.

Seq#	Address	Decimal
1	000000	00
2	100000	32
3	010000	16
4	110000	38
5	001000	08
6	101000	40
7	011000	24
8	111000	56
9	000100	04
10	100100	36
11	010100	20
12	110100	52
13	001100	12
14	101100	44
15	011100	28
16	111100	60

The following illustration shows this 6-bit address space (comprising 64 locations), and the location of the first 16 allocations to be made (according to their sequence number), according to the above list.

```
|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|
1   |   |   |   |   |   |   |       |   |   |   |   |   |   |   |
    |   |   |       |   |   |   2   |   |   |       |   |   |
    |       |   3   |       |       |       |   4   |       |
        5               7               6               8
    9       13      11      15      10      14      12      16
```

The effect of the sparse allocation algorithm is to successively subdivide each remaining block of free address space into 2 equal parts, the first being left to accommodate growth of an existing allocation, and the second being made as a new allocation.

In the first instance (perhaps for the first 1000-2000 allocations, or as otherwise agreed), successive blocks can be selected according to a predefined list as described above. Subsequently however, it will be desirable to select blocks of free space for subdivision according to their size and the rate of growth demonstrated by the immediately preceding allocation. See "Avoiding Fragmentation" below.

4.0 Allocation Request Process

When a new address block is required by an RIR, a request will be sent to the CRS for a block of the required size. The CRS will apply the sparse allocation algorithm (as described above) to determine the start address of the next free block of address space which at of at least the size required by the RIR. This block will be registered by the CRS in an internal database, and then returned to the requesting RIR.

When a subsequent allocation is required by an RIR (that is, in order to expand an address block already allocated by that RIR), a request must be sent to the CRS for the expansion of an existing allocation to the required size. The CRS will then examine the current state of the Common Address Pool to ensure that the required additional space is available, and return a confirmation of the allocation (and updating internal records as above).

5.0 Avoiding Fragmentation

Under the sparse allocation system described here, the "distance" between neighbouring allocations is initially very large, so it is not expected that any fragmentation of address space will occur for some time.

However, depending on the rate of allocation, and the rate of growth of individual allocations, the avoidance of fragmentation will need to be addressed eventually. A technique is proposed here which avoids fragmentation of address blocks which exhibit rapid growth.

Before allocating a start address for any new allocation, the CAP should be examined to determine whether the immediately preceding allocation is likely to exhaust its available address space in the foreseeable future. If that preceding allocation has grown to occupy more than an agreed percentage of the address space potentially available to it, the selected start address should not be allocated.

In the case below, an existing allocation A has grown to occupy 25% of the space potentially available to it. In this case, the candidate address B would not be allocated, and another address would be selected instead.

```
   A----------- space available to A ----------->
---|xxxxxxxxxx-----------|---------------------|-------
                         ^
                         B-- new allocation --->
```

The question of which alternative address should be selected in this case could be addressed in a variety of ways, but these are not examined in this document.

6.0 CAP Contention

Since the CAP will be accessed by multiple RIRs, a mechanism will be required to manage possible contention by simultaneous requests. This mechanism is beyond the scope of this document.

It is also recognised that a central administrative agency, even a very lightweight one, may for various reasons be unavailable or unreachable for a period of time exceeding the required response time on allocation requests. For this reason it will be necessary for the CRS to make provisional allocations to RIRs of multiple address blocks, so that each RIR always has a "reserve" to be used in the exceptional case that the CRS is unavailable. To avoid duplicate allocations, this "reserve" for each RIR would be updated by the CRS with every allocation made to that RIR.

7.0 Review of Allocation Process

The initial set of allocations made under this policy should be limited, on the basis that allocation system is reviewed before the limit is reached, and adjusted in light of experience. Such a review would take place within the communities of the RIRs, through the regional Open Policy Meeting processes.

The initial set of allocations will have 2048 entries, with a review to be commenced immediately after the 1024th allocation is made from that set.

8.0 Common Registration Service

A registration service entity will be formed to act as the custodian for the global space. The operation of this entity will rotate amongst the RIRs and will be the Secretariat for the entity. Each RIR will operate a database server, the server that is located at the Secretariat will be the master, the servers at the other RIR locations will be mirrors of the master. This data base will NOT be a whois data base. It will contain only the information necessary to identify the RIR that made the allocation and the information necessary to generate the reverse delegation DNS zone file. All information required for the RIR to manage the address space and to provide whois information will be resident at the respective RIRs and subject to their specific processes and procedures. Allocations will be made from the master server by each RIR using AAA. The mirror data base servers at the other RIR locations server provide robustness to the system by being redundant to the master.

9.0 Reverse DNS (ip6.arpa) Requirements

Administration

The ip6.arpa domain and any third level domains will be administered by the registration service entity. Technical administration will be performed by the RIR that is the Secretariat. Each RIR will operate a DNS server, the server that is located at the Secretariat will be the silent master for these domains, the servers at the other RIR locations will be mirrors of the master. The zone files will be created from the master registration data base server that is located at the secretariat location. Each RIR will operate a suite of servers that will be the secondary servers for these domains. The SOA records for these domains will be transparent.

Delegation

All delegations will be made on a nibble boundary regardless of allocation boundary. Thus a /32 allocation could also have a corresponding delegation, whereas a /35 would have two (2) /36 delegations.

The initial delegation will consist of the two (2) /4 prefixes that comprise the unicast address space. Allocations made by the RIRs will be delegated from the appropriate /4 delegation and will be done on a nibble boundary as described above. As the number of delegations in the /4 domain grow, intermediary delegations can be made to move the flat space into a hierarchical tree.

10. IANA Considerations

This proposal is intended to provide a deterministic means of allocating address blocks. Once agreed by the relevant communities, this process can be carried out by any party.

As discussed above, the algorithm used to generate start addresses will be subject to revision in the light of experience, and both the algorithm itself and the broader allocation policy will be subject to regular review by the addressing community.

It is proposed that in order to maximise the benefits of the system, the entire FP001 be allocated by IANA to this system, for management by the RIRs.

Author's Addresses

Paul Wilson
APNIC
Level 1, 33 Park Road
Milton QLD
AUSTRALIA

Phone: +61 7 3367 0490
Email: pwilson@apnic.net

Raymond Plzak
ARIN
3635 Concorde Parkway, Suite 200
Chantilly, Virginia
United States

Phone: +1 703 227 9850
Email: Plzak@arin.net

Axel Pawlik
RIPE NCC
Singel 258
1016 AB Amsterdam
The Netherlands

Phone: +31 20 535 4444
Email: axel@ripe.net