

IPv6 Resource Management - RIR Proposal


Background and Motivation

- IANA-RIR allocation system
 - Unchanged in 10+ years
 - Major IPv4 address space fragmentation
 - Many ISPs have many separate prefixes
 - IPv6 should not go the same way
- Proposal for new system for IPv6
 - Designed to minimise fragmentation
 - Most ISPs will have 1 prefix for many years
- Document development
 - Document jointly authored by RIRs
 - Published as *ripe-261*

Current Allocation System

- IANA allocates to RIR
 - RIR maintains a pool of addresses
 - Attempts to maximise aggregation within pool
 - Short-term reservations
 - Sparse allocation
- RIRs allocate to LIRs/ISPs
 - When pool runs low, RIR receives more from IANA
 - Subsequent allocations to existing ISPs cannot be aggregated

Current Allocation System (v4)



ISP has 2 prefixes after 3 requests!

Current Allocation System

- IPv4
 - IANA to RIR allocation unit: /8
 - RIR to LIR/ISP: /20... /10...
 - Many ISPs have multiple prefixes
- IPv6
 - IANA to RIR allocation unit: /23 (64 x /29)
 - RIR to LIR/ISP: /32 minimum
 - IPv6 swamp is being created already
 - Maximum reservation per ISP is /29

Proposal

- “Sparse Allocation” system
 - Maximise “distance” between separate portable allocations
 - Maximise chance of aggregation of subsequent allocations
 - Implemented as list of address prefixes to be allocated in order
- For example...



Proposal

- Sparse allocation system will maximise aggregation
 - Simple system, easily understood
 - Otherwise known as “binary chop”
 - Used in practice by RIRs already (IPv4)
 - Within large address blocks (e.g. /8)
 - Used in other allocation systems
 - e.g. dynamic memory allocation

Proposal

- Benefits increase as address pool increases
 - System breaks down in “overflow condition”
 - i.e. where pool becomes too crowded or full, and another pool must be allocated
 - Therefore RIRs propose to share a single global pool
 - Known as Common Address Pool (CAP)
 - Managed by RIRs jointly, under “Common Registry Service” (CRS)

Proposal

- CAP needs to be as large as possible
 - to ensure long life of single pool
 - to avoid unaggregatable allocations
- So...
 - IANA to allocate 2000::/3 (FP001) for CAP
 - For management by CRS
 - This address space already designated by IETF as Global Unicast, for allocation by RIRs

Allocation Request Process

1. First IPv6 allocation to ISP

- RIR sends request to CRS for new block of specified size
- CRS allocates next entry from list of start addresses

2. Subsequent allocation to ISP

- RIR sends request to CRS for expansion of existing allocation for that ISP (to certain specified size)
- CRS provides extension of existing allocation
 - If extension is not available, new prefix must be allocated

Avoiding Fragmentation

- Distance between neighboring allocations is initially very large
 - “Dumb” algorithm can be used initially
- However, some ISP allocations will grow faster
 - Threatening to “collide” with neighbour
- “Smarter” algorithm for new allocations
 - e.g. If existing preceding allocation has grown to occupy more than a certain % of address space available to it, select next start address from the list

Avoiding Fragmentation

- “Smarter” algorithm...



However note that this is a far future scenario...

Other Details

- Review of allocation process
 - Initial set of allocations limited to 2048
 - Providing each ISP with up to /14 (!)
 - Commence review after 1024th entry (2-3 years?)
- Common Registry Service (CRS)
 - Function to rotate between RIRs
 - ‘Master’ server at one RIR
 - Mirror servers elsewhere
- Reverse DNS requirements (ip6.arpa)
 - CRS administers master DNS server
 - Other RIRs will be mirrors of master



Disadvantages

- Requires single large allocation
 - Maybe “Putting all our eggs in one basket”
 - RIR proposal is to utilise very large block, only one-eighth of IPv6 address space
- Not possible to identify specific blocks allocated to specific RIRs/regions
 - e.g. for filtering purposes
 - RIRs note that this is not possible in IPv4 due to historical allocations

Further information

- Document available from
 - <http://www.ripe.net/ripe/docs/ipv6-sparse.html>
- APNIC IPv6 SIG
 - <http://www.apnic.net/meetings>
 - <http://www.apnic.net/lists>