

IPv6 Global Routing

John Barlow

Some IPv6 design goals


- Reduce routing table growth via address allocation policy.
- Basic Principle of IPv6: Routers along the way should do minimal processing.

Address Allocation Policy

- There were many discussions of how to interpret the address space when IPv6 was being developed.
- Suggestions included:
 - Provider Independent
 - Essentially what v4 does
 - Provider Based
 - Geographical
- Ostensibly provider-based addressing was selected.
 - It is important to understand the difference between allocation and assignment.

Provider Based Unicast Addresses

- Aggregatable global unicast address



Types of Unicast Addresses

- Aggregatable global unicast address
 - Used in production IPv6 networks
 - Goal: minimize global routing table size
 - From range 2000::/3
 - Three fields in /64 prefix
 - 16-bit Top Level Aggregator (TLA)
 - 8-bit reserved
 - 24-bit Next Level Aggregator (NLA)
 - 16-bit Site Level Aggregator (SLA)

Unicast Address Terminology

- TLA, NLA, SLA no longer used in RFCs
- Instead we have
 - Global routing prefix
 - Subnet identifier
- Doesn't affect basic ideas

Top-Level Aggregators

- Allocated by RIRs to transit providers
 - They in turn allocate to customers.
- In practice, RIRs have adopted “slow-start” strategy
 - Start by allocating /32s
 - Expand to /29s when sufficient use in /32
 - Eventually move to /16s

Getting IPv6 address space

- Claim you will hand out 200 /48 networks in the next 2 years ...
- Get an allocation from an ISP (you can not take it with you ...)

AARNet Allocation

- Allocated 2001:388::/32
- This leaves 32 bits of network space available.
- We will see later how this is to be used.

NLAs and SLAs

- NLAs used by providers for subnetting
 - Allocate blocks to customers
 - Can be multiple levels of hierarchy
- SLAs used by customers for subnetting
 - Analogous to campus subnets
 - Also can be hierarchical
 - Minimum size is /48

Current Practice and Aggregation

- In fact the use of terms like TLA and NLA is not longer in use.
- However the intent of Provider based addressing is still the same.
- The goal here is aggregation.
 - As you move up the provider chain many addresses get aggregated into larger blocks.
 - If implemented completely the result would be a default free zone with a very small number of prefixes.

Other Unicast Addresses

- Original provider-based
- Original geographic-based
- GSE (8+8)
- Tony Hain's Internet Draft for provider-independent (geographically-based) addressing

Routers & BGP

- You can start cheap with a PC running FreeBSD or Redhat (zebra for BGP, RADVD for auto-configuration)
- If using a Cisco router you should update Cisco IOS to new syntax
 - `conf t`
 - `bgp upgrade-cli`
 - requires 12.0(22)S or 12.0(14)ST or 12.2(15)T ...

Routers & IGP's

- Limited options for IGP with IPv6, but updates being released (ISIS seems to be popular with Cisco, OSPF out soon ?) – expect to be at the bleeding edge of releases for a while ...
- RIPv6 is available.

Global Routing Table

TLAs – Top Level Aggregators

- AARNet has 2001:388::/32, and can not advertise smaller blocks than this – no longer “small allocations” to sites, but large chunks to “aggregators”.
- You can have multiple addresses, which provides the same as multi-homing.
- `[[show global routing table]]`