

IPv6 DEPLOYMENT

GLOBAL TRANSIT COMMUNICATIONS

Presented by Mark Tinka
Chief Network Architect
Global Transit
Kuala Lumpur, Malaysia



IPv6: The Interest

IPv6: The Interest

- Main reason; readiness for effects of IPv4 exhaustion.
- Understand the technology so we aren't caught flat-footed.
- To be ahead of the game.
- Dual-stack IP network deployed December, 2008.

IPv6: The Interest

- Knowledge primarily acquired from the “6NET IPv6 Deployment Guide” of 2005.
- Although not very current at the time, this 442-page document served as an excellent base.
- It became easier to build upon this information with present-day concepts and designs related to IPv6 (as well as IPv4).

Design Choice

Design Choice

- For our IPv6 deployment, a dual-stack design was the obvious choice.
- A simple, scalable, clean IPv6 design was our strongest motivation for a dual-stack deployment.
- 6PE promised IPv4 forwarding rates, but this was not a concern at the time; our kit still had a lot of “juice” left in it.

Pre-IPv6 Network

Pre-IPv6 Network

- ❑ Multi-vendor network (Cisco & Juniper).
- ❑ Fully MPLS-enabled network.
- ❑ IS-IS carrying Loopback & infrastructure IPv4 addresses.
- ❑ iBGP carrying customer prefixes.
- ❑ MP-iBGP carrying customer VPNv4 prefixes.
- ❑ BGP-free core; core routing 100% MPLS-based.
- ❑ Routers were a combination of software- & hardware-based platforms.
- ❑ Production code was current at the time, IOS 12.2(33)SRC2, 12.2(33)SXH3 & JunOS 9.2R2.

First things first

First things first

- We needed IPv6 addresses (obviously).
- APNIC requirement at the time of acquisition was that LIR's should indicate the capability to make 200 assignments.
- This policy has since been updated as at 4th August, 2008.

Deployment

Deployment

- ❑ Verified IPv6 support on installed kit.
- ❑ Nothing special to be done at Layer 2; this was all transparent to the IPv6 protocol.
- ❑ IOS & JunOS code on the routers supported all the IPv6 features we planned to deploy.
- ❑ IPv6-signaled MPLS core not supported; means IPv6 core would be BGP-aware.

Deployment

- ❑ Servers and stations running Mac OS X 10.5 & openSuSE-11.1 had native IPv6 support, enabled by default.
- ❑ Servers running FreeBSD-7.0 had native IPv6 support, enabled with 3 lines in ``/etc/rc.conf'`.

Addressing

Addressing

- Yes, IPv6 provides millions of addresses and can satisfy even the most insatiable and wasteful of networks...
- But we did not see the need to “waste” address space if we did not have to... did we not once think 4.2 billion IPv4 addresses were more than enough :-).
- A debate for another day, perhaps...

Addressing

- We assigned our address space like so:
 - /48 per PoP; included both infrastructure + customer point-to-point WAN links.
 - Each /48 broken down further to individual /64's.
 - /64 assigned to backbone-wide Loopbacks (taken from main PoP's /48 assignment, for ease of administration).
 - /64 assigned to routing & switch infrastructure.
 - /64 assigned to servers/services.
 - /64 assigned to WAN point-to-point links

Addressing

- On more specific addressing:
 - /128 used for Loopback addresses.
 - /112 used for routing & switching infrastructure addressing.
 - /112 used for server/services subnets.
 - /126 used for WAN point-to-point links.
- For customers, our plan was to assign addresses, per the APNIC policy, but we are reviewing this for internal purposes.

Configuration

Configuration

- ❑ Dual-stack configuration had to be done delicately, as IS-IS implements IPv4 and IPv6 in the same IGP.
- ❑ Support for IPv6 routing in the switches was available as well. However, some of the edge switches didn't support IS-IS, while others required licenses to do so ☹️.
- ❑ This problem was solved by redistributing the switch's Loopback address via the attached edge router, into IS-IS.

Configuration

- A few tips on IPv6 Cisco IOS configuration:
 - ipv6 cef - enable IPv6 CEF switching; it's not enabled by default (I wonder why).
 - ipv6 unicast-routing - also not enabled by default... hmmmh... could this hurt deployment :-\ ?.
 - no ipv6 source-route - disable source routing for IPv6.
 - no ipv6 redirects - disable IPv6 ICMP Redirects.
 - ipv6 nd supress-ra - disable Neighbor Discovery Router Advertisements.
 - multi-topology - enable IPv6 MT in order to prevent IPv4 outages during transition (details follow).

Configuration

- A few tips on IPv6 Juniper JunOS configuration:
 - IS-IS support for IPv4 & IPv6 is enabled by default when IS-IS is turned on.
 - topologies ipv6-unicast – enable IPv6 MT in order to prevent IPv4 outages during transition (details follow).

Configuration

□ IS-IS (IOS)

- Enabled at the interface level.
- Other IS-IS parameters configured at the IS-IS process level.
- Supports an IPv6 address family for values specific to IPv6.

```
interface GigabitEthernet0/1
  ipv6 address 2001:4498:0::1/112
  ipv6 router isis 1
  isis ipv6 metric 10 level-1
```

```
router isis 1
  net 49.0001.1234.5678.9012.00
  passive-interface Loopback0
  !
  address-family ipv6
    multi-topology
```

Configuration

□ IS-IS (JunOS)

- Enabled at the interface level (defining of “iso” family).
- Other IS-IS parameters configured at the IS-IS protocol level.

```
ge-0/0/0 {  
  unit 0 {  
    family iso;  
    family inet6 {  
      address 2001:4498:0::1/112;  
    }  
  }  
}
```

Configuration

```
lo0 {  
  unit 0 {  
    family iso {  
      address 49.0001.1234.5678.9012.00;  
    }  
    family inet6 {  
      address 2001:4498::2/128;  
    }  
  }  
}
```

```
topologies ipv6-unicast;  
interface ge-0/0/0.0 {  
  level 2 disable;  
  level 1 {  
    metric 10;  
    ipv6-unicast-metric 10;  
  }  
}
```


Configuration

```
interface lo0.0 {  
    passive;  
}
```

Configuration

□ BGP (IOS)

- Requires MP (Multi-protocol) BGP.

```
router bgp 24218
  no bgp default ipv4-unicast
  !
  address-family ipv6
```

- Configuration similar to IPv4 BGP setup.
- Only difference is 'network' statements are written using CIDR notation.

```
router bgp 24218
  network 2001:4498::/32
```

Configuration

□ BGP (JunOS)

- Requires MP (Multi-protocol) BGP.

```
[edit protocols bgp]
```

```
group rr-peers6 {
```

```
  family inet6 {
```

```
    unicast;
```

```
  }
```

```
  export BGP-OUTBOUND-IPV6-POLICY;
```

```
}
```

- Configuration similar to IPv4 BGP setup.

Configuration

- ❑ Static routing, ACL's and prefix lists also utilize CIDR notation when describing IPv6 addresses (IOS).
- ❑ No more inverse mask calculations :-).

```
ipv6 route 2001:4498::2:0/112 Null0
ipv6 route ::/0 Serial4/0
!
ipv6 prefix-list upstreams-in6 seq 10 deny 3FFE::/16 le 128
ipv6 prefix-list upstreams-in6 seq 300 permit ::/0 le 48
!
ipv6 access-list filter-incoming6
deny ipv6 3FFE::/16 any
permit ipv6 any 2001:4498::/32
```

Configuration

- Static routing, firewall filters and route-filters also utilize the same basic syntax as IPv4 (JunOS).

```
rib inet6.0 {  
  static {  
    route 2001:db8::1/128 discard;  
  }  
}
```

```
policy-statement upstreams-in6 {  
  term 10 {  
    from {  
      route-filter 3ffe::/16 upto /128;  
      route-filter 2001:db8::/32 upto /128;  
      route-filter fe00::/9 upto /128;  
      route-filter ff00::/8 upto /128;  
    }  
    then reject;  
  }  
}
```

Configuration

```
family inet6 {
  filter filter-incoming6 {
    term 10 {
      from {
        source-address {
          3ffe::/16;
          2001:db8::/32;
          fe00::/9;
          ff00::/8;
        }
        destination-address {
          ::/0;
        }
      }
      then {
        discard;
      }
    }
  }
}
```

Configuration

- openSuSE-11.1
 - Compiled by default within Linux kernel
 - Enabled by default after install.
 - IPv6 address and routing added via the ``ifconfig`` and ``route`` commands.

```
ifconfig eth0 inet6 add 2001:42b0::6
```

```
route -A inet6 add default gw 2001:4498::1
```

Configuration

□ Mac OS X 10.5.6

- Supports IPv6 out-of-the-box.
- We experienced a bug in OS X that prevents static IPv6 addresses from being assigned via the GUI. The workaround is to assign the address via the CLI:

```
ifconfig en0 inet6 2001:4498::1/112
```

```
route add -inet6 -prefixlen 0 default 2001:4498::2
```

- Automatically assigned IPv6 addresses are unaffected by this bug.

Configuration

□ FreeBSD-7.0

- Compiled by default within kernel.
- Enabled via `‘/etc/rc.conf’` with 3 lines.

```
ipv6_enable="YES"
```

```
ipv6_defaultrouter="2001:4498:0000:0000:0000:0000:0000:FF7F"
```

```
ipv6_ifconfig_dc1="2001:4498:0000:0000:0000:0000:0000:0006/112"
```

DNS

DNS

- Dual-stack hosts had no problem resolving AAAA and PTR records via an IPv4-only DNS server.
- BIND-9 running on FreeBSD-7.0 had no problems listening on an IPv6 interface.
- No recompilation needed, simply add the following to `/etc/named.conf`:

```
options {  
    listen-on-v6 { any; }  
}
```

DNS

- DNS name resolution using an IPv6-only resolver on our openSuSE-11.1 & Mac OS X stations worked fine too (watch out, Windows XP users!!!).
- Creation of fully functional AAAA forwarding records worked fine.
- The 'ip6.arpa' reverse records worked with no problems, for both our ::1 (localhost) and 2001:4498::/32 zones.
- We used 'sipcalc' to create full 'ip6.arpa' domains (command was 'sipcalc -r')

IPv6 Transit

IPv6 Transit

- With our core network functional as a dual-stack platform, it was time to talk to the outside world.
- We were able to establish native IPv6 peering with some of our in-country peers.
- To the rest of the world, we were able to acquire native and tunneled IPv6 transit access from some of our upstreams that had current support.

IPv6 Transit

- 'aut-num' and 'route' objects were successfully created in the RIPE WHOIS database.
- At the time, we received ~1,403 IPv6 prefixes on a full BGP feed with our IPv6 transit upstreams.

IPv6 Transit

route6: 2001:4498::/32
descr: Global Transit Communications, Malaysia, Networks - 2001:4498::/32
descr: In case of abuse, please contact abuse@globaltransit.net
origin: AS24218
mnt-by: GTI-MY-MNT
source: RIPE # Filtered

Basic security

Basic security

- Cisco IOS & Juniper JunOS
 - SSH over IPv6 was supported.
 - ACL's for IPv6 are supported.
 - The equivalent of BCP-38 and RFC 3330 for IPv6 at the time included filtering of the following addresses:
 - 3FFE::/16
 - 2001:DB8::/32
 - FE00::/9
 - FF00::/8

Basic security

- For IOS, application of the ACL's to an interface uses a slightly different command than that of IPv4.

```
interface GigabitEthernet0/1
  ipv6 traffic-filter acl-name in|out
```

- For JunOS, application of the firewall filters to an interface is similar to doing the same in IPv4.

```
lo0 {
  unit 0 {
    family inet6 {
      filter {
        input | output firewall-filter-name;
      }
    }
  }
}
```

Basic security

- openSuSE-11.1
 - 'ip6tables' included within this distribution's Linux kernel.
 - 'ip6tables' application ships with the distribution.
 - SSH over IPv6 worked with no problems.

Basic security

- Mac OS X 10.5.6
 - 'IPfw' is the system's default packet filter.
 - Tested with custom scripts not that dissimilar from FreeBSD's IPfw implementation.
 - SSH over IPv6 worked with no problems.

Basic security

- FreeBSD-7.0
 - 'IPfw' and 'Pf' are supported.
 - SSH over IPv6 worked with no problems.

Issues

Issues

□ Cisco

- Newer-generation “desktop” switches require the SDM (Switch Database Management) template to be changed to support ‘dual-ipv4-and-ipv6’ template before the switch can support IPv6 (both forwarding & control planes). SDM templates only take effect after a system reboot; problem!
- Low-to-mid range “desktop” switches do not support IS-IS. Support is rumoured to appear this year, but only for IPv4 (and depending on customer demand).

Issues

□ Cisco

- IPv6 BGP sessions do not support the 'fall-over' feature. Sessions configured with this feature will not form.
- Loopback interfaces already configured as 'passive' under IS-IS (IPv4) will not be installed in the IS-IS IPv6 database after an IPv6 address is added to them. Workaround was to ensure the Loopback interface is made passive in IS-IS "after" an IPv6 address was added – this state survives a reboot as well. This bug has already been fixed in 12.2(33)SRC3.

Issues

□ Cisco

- IPv6 traceroute output displays AT&T's ASN, AS2686 at each router hop. This is a problem particular to the 12.2SR* train. This is only a cosmetic problem and does not affect ongoing network operations – but it might be bad for business depending on who's watching 😊.
- In some cases, the IPv6 'neighbor' statements appear under the IPv4 address family. A bug has already been filed for this issue. Workaround was to disable the neighbor with the 'no' form of command.

Issues

- Juniper
 - New EX-series “desktop” switches require a license to run and support IS-IS.

Issues

- IS-IS Multi-topologies
 - Dual-stack deployment of IPv6 in IS-IS requires that topologies be congruent.
 - However, the actual fact is that during transition of router interfaces from a single to a dual IP stack, topologies are not congruent.
 - After IPv6 addresses are configured on a dual-stack interface with IS-IS for IPv6 enabled, the router will lose IS-IS adjacencies.
 - The solution was to ensure IS-IS multi-topology support was enabled prior to configuring IPv6 addresses on the router interfaces.

Issues

- IS-IS Multi-topologies
 - Enabling multi-topology support for IS-IS on the production routers was a safe procedure, although it is still recommended that any changes be made during maintenance windows.

Issues

- openSuSE-11.1
 - 'ip6tables' contributed scripts not as rife as those for IPv4 ('iptables') – build your own 😊.
 - We found the 'ping6' and 'traceroute6' commands rather annoying (same goes for FreeBSD & Mac OS X).

Issues

- Security (router):
 - This applies to router ACL security
 - Cisco and Juniper both have different ways they treat the filtering of special IPv6 addresses, e.g., link-local, multicast, e.t.c.
 - Cisco will NOT filter link-local addresses (used for next-hop information to global IPv6 addresses).
 - Juniper WILL filter these, which makes sense.
 - We are currently investigating this “phenomenon” with both vendors.

Note

Note

- ❑ Should you ever need to configure Cisco IOS with IPv6 only, you still need a 32-bit Router-ID.
- ❑ Essentially, an IPv4 address on a Loopback interface.
- ❑ The IPv6 routing protocols on an IPv6-only IOS router require an IPv4 Router-ID.

Concerns

Concerns

- ❑ A scalable transition mechanism for IPv4-only and/or IPv6-only sites.
- ❑ Router forwarding rates in native IPv6 deployments, especially in software-based platforms.
- ❑ The readiness of essential ISP ancillary software, e.g., billing, provisioning, management, e.t.c.
- ❑ Current IPv6 address allocation and assignment policies.

Helpful resources

Helpful Resources

<http://www.civil-tongue.net/6and4>

http://www.getipv6.info/index.php/Main_Page

<http://www.cisco.com/en/US/products/sw/iossw>

END

Thank you!

Q&A

mtinka@globaltransit.net