

Summer V6cation

Pete Templin

TexLink Communications,
A PacWest Company

Background

- Chose to roll out IPv6 to break the ice
- Single state, four city network
 - Now larger after being acquired...
- Easier to campaign for related resources if this piece was complete

Architecture

- Two core routers/POP
 - Mostly GSR
- Dedicated upstream edge routers
 - All GSR
- 7206/7507/GSR/6509 customer attach
 - 7507s going away

Deployment Strategy

- IPv4/IPv6 Dual Stack deployment
 - Every GSR
 - Every Sup720
- Based on equipment lifecycle
 - All upstream
 - Most of core (skipped one POP)
 - Ethernet in two largest POPs, T3s+ in 1 POP

Deployment tactics

- No effective testing lab at hand
- 2N redundancy at core layer
 - Move traffic away from one core
 - “set-overload-bit” or “max-metric router-lsa”
 - “Reload in 20:, apply commands, cross fingers
 - Check, “reload cancel”

Deployment tactics

- VPN access from San Antonio
 - Deploy outward from San Antonio to Dallas
 - Work through Dallas (and secondary Dallas site)
 - Work to/through Houston
 - Finish through San Antonio
- 1 router/day, then 2/day, then 3, etc.

Deployment tactics

- Phase 1:
 - Interface Addressing
 - OSPFv3
- Phase 2:
 - iBGP

Addressing

- 2607:F1B8::/32 from ARIN
- Initially:
 - 2607:F1B8:FFFF::/48 carved for internal ops
 - Note to self: use :0000: instead...those Fs get old!
- Current:
 - 2607:F1B8::/48 for production, corporate ops

Internal Addressing

- Loopbacks:
 - 2607:F1B8:POP:node/128
- Intra-POP links:
 - 2607:F1B8::POP:link-ID:endpoint/112
- Inter-POP links:
 - 2607:F1B8::POP#1:POP#2:link-ID:endpoint/112

Lessons learned

- Define “endpoint-ID” carefully before beginning:
 - Tried to get too fancy with endpoint Ids
 - Inconsistent relevance
 - Easier to use :1 and :2 across the board.
- Too many endpoint ID collisions until I got this squared away.

Customer assignments

- Plan:
 - /112 WAN/interconnect link
 - /56 for customer, automatically
 - /48 for customer, on request
 - Larger assignments ICB
- Reality:
 -

Customer assignments

- Divide & Conquer method
 - :8000:, then :4000:, then :C000:, then :2000:, etc.
 - Maximize the chance that any customer can expand their assignment.
- Spreadsheet started to track customer assignments
 - Not many customers, we'll deal with scaling later. For us, it's worked for v4...hush!

IGP History

- Originally OSPF, based on operational experience at the time.
- Migrated to ISIS (thanks for the hints, Vijay/ATDN!) ~3 years ago.
 - Our OSPF mesh had auto-costing (and even that had issues) – migration provided a safer path, along with other *perceived* benefits.

IGP Routing

- IPv4 rolled back to OSPF from ISIS for consistency
 - Had issues with multi-topology ISIS
 - Had other equipment reasons
- OSPF once again
- iBGP ‘hack’ as in IPv4: only loopbacks, intra-POP links, inter-POP links in IGP, all others in BGP.

Internal Services

Upstream Transit

- Cogent
 - “Not yet, no roadmap”
- AboveNet
 - “Not yet, no roadmap”
- NTT
 - “Gladly, for a small fee”

Interim Transit

- Hurricane Electric tunnel
 - Allowed operational experience
- See “Internal Services”...

State of the packets

- So far, so good
- No customer demand yet
- Full expectation that sales will sell IPv6 to a customer on kit that won't do v6.
- Verified that equipment roadmap is restricted to v6-capable gear.

Moral of the story

- Once we got the ball rolling, the roll-out was easy.
- Anyone who said there isn't revenue to be made with IPv6 hasn't picked up the phone.
- Corollary: Availability of IPv6 (and price) will likely influence future transit decisions.

Phase one config snippet

- `ipv6 unicast-routing`
- `ipv6 cef dist`
- `!`
- `ipv6 ospf name-lookup`
- `!`
- `interface Loopback0`
- `ipv6 addr 2607:F1B8::2:2/128`
- `ipv6 ospf 1 area 0`
- `ipv6 ospf cost 1`
- `!`
- `interface Serial4/0`
- `ipv6 addr 2607:F1B8::2:3:3:1/112`
- `ipv6 ospf 1 area 0`
- `ipv6 ospf cost 38`
- `ipv6 ospf hello 1`

Phase two config snippet

- `router bgp 11457`
- `neig dlls-core-v6 peer`
- `neig dlls-core-v6 remote-as 11457`
- `neig dlls-core-v6 upd lo0`
- `neig 2607:f1b8::2:1 peer dlls-core-v6`
- `!`
- `address-family ipv4`
- `no neig 2607:f1b8::2:1 activate`
- `!`
- `address-family ipv6`
- `neig 2607:f1b8::2:1 peer dlls-core-v6`
- `network 2607:f1b8::/32 route-map bgpnet-bgp-v6`
- `!`
- `route-map bgpnet-bgp-v6 permit 100`
- `set local-preference 400`
- `set community 11457:30200`
- `!`
- `ipv6 route 2607:F1B8::/32 Null0`

Closing

- I'll gladly share full config deltas as applied to our routers.
- Mail me: templin@templin.org