

All this garbage...

What sort of IPv6 BGP filters are useful?

Gert Döring, SpaceNet AG

May 3rd, 2005

RIPE 50, Stockholm, SE - Routing WG

Historic Overview

- in the 6bone, nobody filtered anything
- all sort of garbage in the BGP table
 - typos (3FFE:xx:: \Leftrightarrow 3FF3:xx::)
 - accidental more-specific leaks (50 x /48, etc.)
 - router vendor bugs (Cisco 1000::... 6PE leaks)
 - ...
- \Rightarrow early BGP filter recommendation on <http://www.space.net/~gert/RIPE/ipv6-filters.html>
- so people started filtering, and all was good...

but...

- Initial filtering recommendation looked like this:

```
ipv6 prefix-list ipv6-strict permit 2001::/16 ge 35 le 35
ipv6 prefix-list ipv6-strict permit 2001::/16 ge 21 le 32
ipv6 prefix-list ipv6-strict deny 0::/0 le 128
```

- “everything in 2001:: should be between /21 and /32” (+/35)
- then came 2001:2000::/20 (Telia)
- and the well-meant filters broke their announcements :-)

updates...

- filtering recommendations updated (of course)
- but broke again when 2001:5000::/21 was allocated
- and *again* for 2003::/19
- and *AGAIN* for 2600:: and 2A00:: allocations
- lots of problems for large network operators due to other participants that don't update their filters in time
- so it seems the current model ("permit those things that we know about") is just not appropriate for the backbone

So how to tackle this?

- filtering towards BGP customers is a *GOOD* thing.
 - you *really* want and should do this
 - but this sort of filters should not be done in a generic “one-size-fits-all” matter, but the classic way, querying the RIPE DB for route6 objects, and building filters from that
- but what about peering / upstream BGP sessions?
 - *what threads are you protecting yourself?*
 - are there commonly-agreed “good” and “bad” routes?
 - we should get some consensus on that and base recommendations on it

BGP table threads?

- very long more-specifics (/64, /127, ...)
 - fairly agreed-upon that this is not what we want
- intermediate more-specifics (/48, /40, ...)
 - some say “useful”, other say “don’t encourage that!!”
 - watch out for /48 microallocations
- hijacked space (someone using 3000:1234::/32)?
 - you’d need filters against 2001:609::/32 as well...!
- router overload, *make sure only /32 and shorter are allowed* ?
 - there are 65536 /32s inside 2001::/16 alone
- fallback to max-prefix on peer/upstream links? Secure-BGP?