

Some Initial Measurements of Prefix Length Philtres

Jennifer Rexford <jrex@research.att.com>

Steve Bellovin <smb@research.att.com>

Randy Bush <randy@research.att.com>

NANOG / Scottsdale

2001.05.21

What was Measured

- Prefix length filters applied to some well-known BGP tables
 - Telstra internal of 2001.05.15
 - route-views Univ of Oregon 2001.03.16
 - RIPE multi-view of 2001.03.01
 - Verio customer of 2001.05.16
- Applying various filters
 - /24
 - RIRs' smallest allocation sizes in /8s

The Increasingly Satanic Philtres

[applied sequentially]

- Clean data by removing Martians: 1918, IXs, ...
This gives the **Baseline Reachability**
- Remove all prefixes longer than /24
- Remove all prefixes in **A** space longer than RIRs allocate
- Remove all prefixes in **B** space longer than RIRs allocate
- Remove all prefixes in **C** space ≥ 206 longer than RIRs allocate

Legend

- Number of prefixes removed at this stage
- Prefix count after this filter is applied
- Count of baseline address space /32s which have become unreachable
- Portion of baseline address space /32s which was reachable in baseline but is no longer reachable after this filter is applied

Telstra

	Prefixes Removed	Prefixes Left	/32s Lost	How Much Unreachable
Telstra clean	31	108,215		
/24	5,960	102,255	1,835	0.000002
A	13,846	94,369	366,123	0.000328
B	20,881	87,334	1,834,795	0.001643
C	36,303	71,912	3,845,675	0.003443

route-views (Univ of OR)

	Prefixes Removed	Prefixes Left	/32s Lost	How Much Unreachable
route-views				
clean	50	116,097		
/24	12,938	103,159	3,456	0.000003
A	20,086	96,011	497,536	0.000488
B	27,599	88,498	1,910,144	0.001720
C	42,974	73,123	3,836,544	0.003455

Verio Customer

	Prefixes Removed	Prefixes Left	/32s Lost	How Much Unreachable
Verio				
clean	31	86,966		
/24	80	86,886	0	0.000000
A	741	86,225	15,104	0.000014
B	1,748	85,218	1,352,960	0.001221
C	17,024	69,942	3,345,920	0.003021

RIPE

	Prefixes Removed	Prefixes Left	/32s Lost	How Much Unreachable
RIPE				
clean	44	121,014		
/24	5,347	115,667	2,591	0.000002
A	14,315	106,699	438,815	0.000394
B	22,182	98,832	1,695,007	0.001522
C	41,363	79,651	3,544,095	0.003183

Warnings

- RIR published boundaries can be wrong
I.e. RIPE is allocating /20s in 62/8 while RIPE-211 says /19s
- APNIC does not document at all
- These results are fresh out of the oven.
There could be bugs

Warning

- If we assume multi-homing is a major cause of redundant announcements
- I.e. redundant announcements are folk trying to engineer how I send traffic to them
- Then the described filtering techniques maintain reachability but defeat those desires for redundancy and traffic engineering

But why should the rest of us allow you to consume extra resources on our routers when you are not paying us for that service?

[credit to smd]

Is There a Sweet Spot?

- The majority of the benefit comes from the A and B filters
- The majority of the unreachable prefixes comes from the C filter
- Might an intermediate stance be useful?

Where to Go from Here?

- For discussion of the taxonomy of prefix coverage, filtering, ... the ptomaine BOF of the IETF, <ptomaine@shrubbery.net>
- Searching for ways to multi-home without routing table bloat, the IETF multi6 WG, <multi6@ops.ietf.org>
- In the meantime, philtre aggressively and often!

References

List of Prefixes Lost

<http://research.att.com/~jrex/nanog/lost.html>

ARIN Allocation Boundaries

<http://arin.net/regserv/IPStats.html#cidr>

RIPE Allocation Boundaries

<http://www.ripe.net/ripe/docs/ripe-211.html>

This Presentation

<http://research.att.com/~jrex/nanog/lost.html>

<http://psg.com/~randy/010521.nanog/>