

BGP table fragmentation: what & who?

Julien Gamba

Romain Fontugne, Cristel Pelsser, Randy Bush, Emile Aben

University of Strasbourg

Why looking at fragmentation?

647 000

Why looking at fragmentation?

647 000

Amount of IP prefixes announced on the Internet (and counting)

Why looking at fragmentation?

647 000

Amount of IP prefixes announced on the Internet (and counting)

- Large routing tables consume memory
- Routers memory (TCAM) is expensive, so they do not have a lot of it (especially the older ones)
- Too many routes means routers slowing down or shutting down

Why looking at fragmentation?

- Some of these prefixes are fragmented and could be aggregated

Why looking at fragmentation?

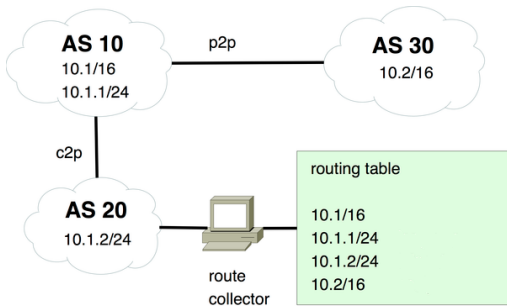
- Some of these prefixes are fragmented and could be aggregated
- Some of these prefixes are redundant and could be removed from the routing table

Why looking at fragmentation?

- Some of these prefixes are fragmented and could be aggregated
- Some of these prefixes are redundant and could be removed from the routing table

- How bad is the fragmentation?
- How many prefixes are redundant?
- What are the causes of all this?


IP prefixes classification



- **Top**: is covering some smaller prefix(es);
- **Deaggregated**: is covered by another prefix which is originated by the same AS;
- **Delegated**: is covered by another prefix which is originated by another AS;
- **Lonely**: does not overlap with any other prefix.

Luca Cittadini et al. (2010). "Evolution of Internet Address Space Deaggregation: Myths and Reality". In: *IEEE journal on selected areas in communications*


IP prefixes classification



- **Top**: is covering some smaller prefix(es);
- **Deaggregated**: is covered by another prefix which is originated by the same AS;
- **Delegated**: is covered by another prefix which is originated by another AS;
- **Lonely**: does not overlap with any other prefix.

Luca Cittadini et al. (2010). "Evolution of Internet Address Space Deaggregation: Myths and Reality". In: *IEEE journal on selected areas in communications*


IP prefixes classification



- **Top:** is covering some smaller prefix(es);
- **Deaggregated:** is covered by another prefix which is originated by the same AS;
- **Delegated:** is covered by another prefix which is originated by another AS;
- **Lonely:** does not overlap with any other prefix.

Luca Cittadini et al. (2010). "Evolution of Internet Address Space Deaggregation: Myths and Reality". In: *IEEE journal on selected areas in communications*

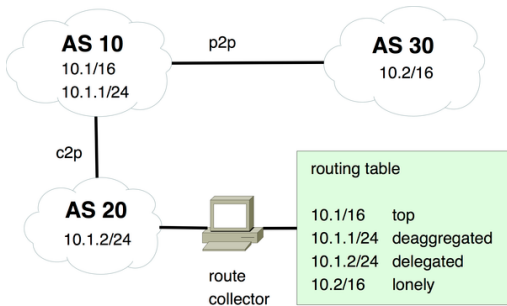
IP prefixes classification



- **Top:** is covering some smaller prefix(es);
- **Deaggregated:** is covered by another prefix which is originated by the same AS;
- **Delegated:** is covered by another prefix which is originated by another AS;
- **Lonely:** does not overlap with any other prefix.

Luca Cittadini et al. (2010). "Evolution of Internet Address Space Deaggregation: Myths and Reality". In: *IEEE journal on selected areas in communications*


IP prefixes classification



- **Top**: is covering some smaller prefix(es);
- **Deaggregated**: is covered by another prefix which is originated by the same AS;
- **Delegated**: is covered by another prefix which is originated by another AS;
- **Lonely**: does not overlap with any other prefix.

Luca Cittadini et al. (2010). "Evolution of Internet Address Space Deaggregation: Myths and Reality". In: *IEEE journal on selected areas in communications*

Deaggregation evolution — prefixes




- routeviews data, as seen by AS 3356 (Level3), counting prefixes
- Proportion of deaggregated is increasing over time
- Combined fraction of deaggregated and delegated prefixes is constant

AS10	10.1/16	Top
AS10	10.1.1/24	Deaggregated
AS20	10.1.2/24	Delegated
AS30	10.2/16	Lonely

Deaggregation evolution — addresses

Evolution of IPv4 addresses classification



- Deaggregated addresses are still increasing, top addresses are increasing too
- Lonely addresses are decreasing

routeviews data, as seen by AS 3356 (Level3)

AS10	10.1/16	Top
AS10	10.1.1/24	Deaggregated
AS20	10.1.2/24	Delegated
AS30	10.2/16	Lonely

Deaggregation evolution

In other words

- There are more and more deaggregated prefixes
- Combined fraction of deaggregated and delegated prefixes is constant
- The proportion of lonely addresses is decreasing, while top and deaggregated increase

Deaggregation evolution

In other words

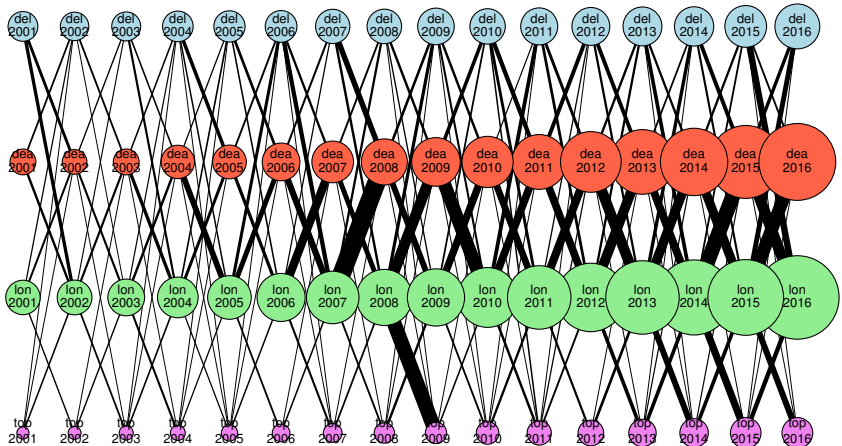
- There are more and more deaggregated prefixes
- Combined fraction of deaggregated and delegated prefixes is constant
- The proportion of lonely addresses is decreasing, while top and deaggregated increase

What is going on?

- Are lonely prefixes becoming top? Deaggregated?
- Do they disappear? Is there just less new lonely prefixes?
- What are the movements between the categories? What is the stability of the prefixes?

Prefixes dynamics

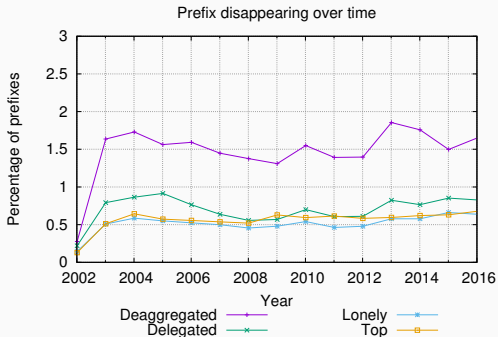
Movements between categories



Most movements appear between lonely and deaggregated prefixes, both ways

Lonely prefixes become deaggregated when the covering prefix is announced

Prefixes disappearances



- Counting how many prefixes disappear from the RIB between two months
- Graph show average per year
- Same trend across categories, although deaggregated prefixes disappear more

routeviews data, as seen by AS 3356 (Level3)

AS10	10.1/16	Top
AS10	10.1.1/24	Deaggregated
AS20	10.1.2/24	Delegated
AS30	10.2/16	Lonely


Deaggregated prefixes seems to be the most volatile kind:

- Why is that?
- Is it traffic engineering?
- Is it something else?

Detection of traffic engineering

AS path prepending

Proportion of prefixes using path prepending

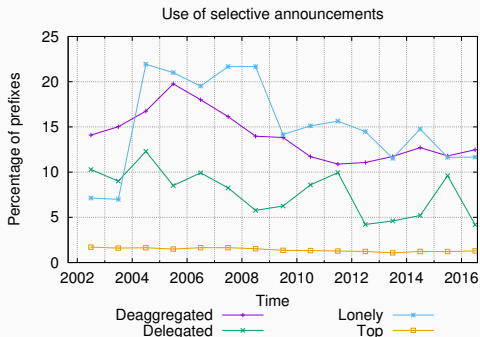


- Graph shows average per year
- Showing the proportion of prefixes announced using path prepending in each category
- Path prepending usage is slowly increasing
- No clear trend between categories

routeviews data, as seen by AS 3356 (Level3)

AS10	10.1/16	Top
AS10	10.1.1/24	Deaggregated
AS20	10.1.2/24	Delegated
AS30	10.2/16	Lonely

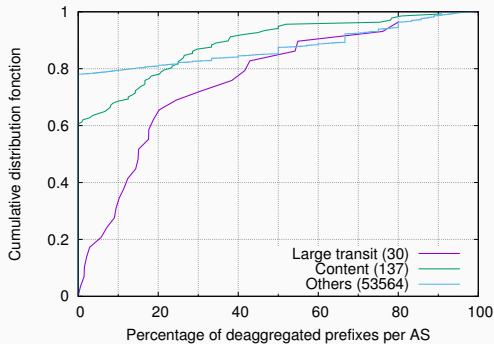
Selective announcements



- Graph shows average per year
- Showing the proportion of prefixes seen by less than half the peers by category
- Lonely and deaggregated prefixes are the categories that use selective announcements the most

AS10	10.1/16	Top
AS10	10.1.1/24	Deaggregated
AS20	10.1.2/24	Delegated
AS30	10.2/16	Lonely

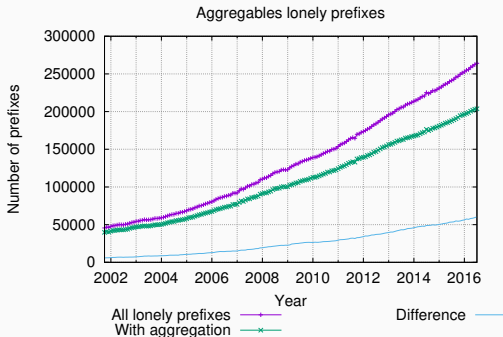
Who is engineering traffic?



- Graph shows proportion of deaggregated prefixes by AS business type
- Large transit AS list comes from CAIDA's classification
- Large transit providers may split their address space to do traffic engineering

AS10	10.1/16	Top
AS10	10.1.1/24	Deaggregated
AS20	10.1.2/24	Delegated
AS30	10.2/16	Lonely

Aggregability of lonely prefixes



Two prefixes are aggregables if:

- they have the same AS origin
- they are consecutive
- the aggregate falls on a power of two boundary

routeviews data, as seen by AS 3356 (Level3)

AS10	10.1/16	Top
AS10	10.1.1/24	Deaggregated
AS20	10.1.2/24	Delegated
AS30	10.2/16	Lonely

In conclusion

- Deaggregation has been increasing for the last fifteen years
- Combined fraction of deaggregated and delegated prefixes is constant
- Some large transit AS split heavily their address space

In conclusion

- Deaggregation has been increasing for the last fifteen years
- Combined fraction of deaggregated and delegated prefixes is constant
- Some large transit AS split heavily their address space

We still don't know:

- why the heavy splitting?
- for traffic engineering?
- for security?

Questions?