

# /48 Considered Harmful

On the Interaction of Strict IPv6 Prefix Filtering  
and the Needs of Enterprise LIRs

Dr. Roland Langner, [roland.langner@burda.com](mailto:roland.langner@burda.com)  
Nicolas Schätzle, [nicolas.schaetzle@burda.com](mailto:nicolas.schaetzle@burda.com)  
Enno Rey, [erey@ernw.de](mailto:erey@ernw.de)



## Who We Are

---



- Dr. Roland Langner
  - Global Network Architect  
@ Hubert Digital Systems (BDS)
  - Chemist
    - Which leads to a certain scientific look at the world ;-)
- Nicolas Schätzle
  - Network wizard at BDS
- Enno Rey
  - Involved in operations in carrier space in the past,  
mostly enterprise space today.
  - IPv6 since 1999.

# Agenda

---

- ▢ Problem Statement
- ▢ Research Approach & Results
- ▢ Conclusions & Proposition



## Problem Statement

---

"There are occasionally requirements for the advertisement of more specific routes from within an allocation. With a few ISPs currently filtering at the minimum PA allocation (/32) within the relevant address ranges, this can cause significant difficulties for some networks wishing to deploy IPv6."

*[RIPE-532]*



## Let's Have a Closer Look at this Statement

---

- What (nature) could those “occasional requirements” be?
- Who's the “few” ISPs filtering?
  - Is it just “a few”?
  - Does the number change over time?
  - If so, how (& why)?



## Quick Refresher

---



- RIPE members / LIRs
  - Types & their motivations
- “Strict Filtering”
  - Origin & practice

# RIPE Members

Let's put them into two main categories



## → “Transit-LIRs”

- “Traditional ISPs”
- Business model centered around connectivity & traffic transport
- Open to many customers



## → “Enterprise-LIRs”


- Large enterprises mostly running their own networks.
- Often main business is not IT-related.
- Have specific requirements...

## Numbers & Trends (I)

Do you think this is all  
“Transit-LIRs”?

Source:  
<https://labs.ripe.net/statistics/lirs-with-and-without-ipv6>

LIRs With and Without IPv6





## Why Do They Do This? Can't they just apply for PI space?

Here's what Deutsche Telekom responded to an organization willing to apply for a single /48 PI space through them (as *sponsoring LIR*).

“vielen Dank für Ihren Auftrag zur Beantragung eines /48 Provider Independent (PI) IPv6-Adressraums beim RIPE NCC. Zunächst einige Informationen zur gewünschten Beantragung.

[...]

Wir weisen ausdrücklich darauf hin, dass die internationale Konnektivität unter Umständen instabiler sein kann als die eines PA-Blocks der Deutschen Telekom.

Das hängt damit zusammen, dass nicht nur IP-Adressraum knapp ist, sondern auch die globale Routingtabelle (von der alle “border gateways” weltweit eine Kopie halten müssen) bereits eine kritische Grösse erreicht hat. Daher gibt es Provider, die strikte prefix-length Filterung in ihrer Routing-Policy einsetzen. Als Folge hiervon kann es sein, dass Ihr Netz von diesen Providern aus nicht erreichbar ist.“

“Die Gebühren für die Beantragung betragen ca. 1600 EUR“



## Strict Filtering



- Initially described by Gert Döring in 2002

- <http://www.space.net/~gert/RIPE/ipv6-filters.html>

- Relevant parts (RIPE space)

```
ipv6 prefix-list ipv6-ebgp-strict permit 2001::/16 ge 35 le 35
ipv6 prefix-list ipv6-ebgp-strict permit 2001::/16 ge 19 le 32
ipv6 prefix-list ipv6-ebgp-strict permit 2001:0678::/29 le 48
(PI-Assignments)
ipv6 prefix-list ipv6-ebgp-strict permit 2003::/16 ge 19 le 32
ipv6 prefix-list ipv6-ebgp-strict permit 2a00::/12 ge 19 le 32
```

## Reasons for Filtering

---

- Save TCAM memory
  - → protect hardware investments.
- Keep routing tables tidy (→ Hierarchical routing)
  - Preserve RFC 1518 spirit.
    - Has this ever worked anyway?



## Does this really Happen?

---



- Well, unfortunately (from Enterprise-LIR perspective): yes.
- Two cases studies from 2014.

## Case Study (I)

---

Burda (Media)

### **Hubert Burda Media**



- 2.6 Billion EUR revenue in 2013
- +10K employees
- Business activities in Germany (~ 70%) and rest of the world

## Case Study (I)

---

Burda (Media)

### Hubert Burda Media



- Three main (data centers|POPs| network hubs) across Germany
  - All dual-homed, full BGP peering, each with an own ASN.
  - Stateful firewalls at each of those ...

## IPv6 @ HBM

---



- Given their industry sector “they are very interested in IPv6”
  - Which for probably everybody here in the room is good news.
- Started related activities back in 2011.
  - And, quite quickly, faced the crucial question: “To LIR or Not?”

## The Dilemma

---



- Current (network) topology *requires* splitting \$ALLOCATED\_SPACE into several pieces
  - One for each POP (e.g. a /34)
  - Did I already mention they have stateful firewalls?
- Some of you will already see where this is going.



## What Happened

---



- Announcement/propagation of that specific POP's /34 at one of their main sites
  - Before you ask: yes, appropriate `route6` object was created in advance.
- Boom! One of two uplink providers (regional SP with 500K broadband subscribers & ~ 5K business customers. kind-of traditional city carrier) filtered announcement, expecting announcement aligned with `inet6num` object.
  - Could be solved by a phone call.

## What Else?

Did that phone call solve the problem?



- Once /34 announcement accepted by \$UPLINK, they closely monitored route propagation via RIPEstats and some looking glasses.
- Took several days until stable situation.
  - Right now everything looks good.
  - Problem solved?
  - Not really, this leaves bad after taste.
    - Can this happen again in the future, for other/future announcements?
    - Probably “business” won’t like this too much...

## Case Study (II)

Evonik Industries



- Chemical
- ~ 14 bn EUR yearly revenue, 33K employees
- RIPE LIR since 2005, IPv6 /32 allocated in 2011.

## Case Study (II)

### Network Details



- Two main European data centers, at two sites.
- Both sites multi-homed, with large German carrier + Top EU Tier-1.
- IPv6 address concept with GUAs only (no ULAs → good!)
  - Planned sec architecture: replace sec-benefit of RFC 1918 addresses by proper filtering and “selective propagation of address space”

## Case Study (II)



- Just one problem with this approach...
- Both carriers expected covering aggregate in addition, they wouldn't route /48 only.
- Which led to somewhat cumbersome 6VPE config on MPLS network between data centers.

But that's only Anecdotal  
Evidence, Isn't It?

---



- That exactly was our question.
- Research question:
  - What is the amount of *more-specific* IPv6 routes from RIPE PA space out there?
  - Read: how much harm does an ISP performing strict-filtering potentially cause? Is this common practice?

## Research Approach

---



- Look at RIS data from 2010–2015
  - <http://www.ripe.net/data-tools/stats/ris>
- Analyze length of prefixes announced at several IXs.

# Analyzed RRCs

Netnod / SE – Stockholm

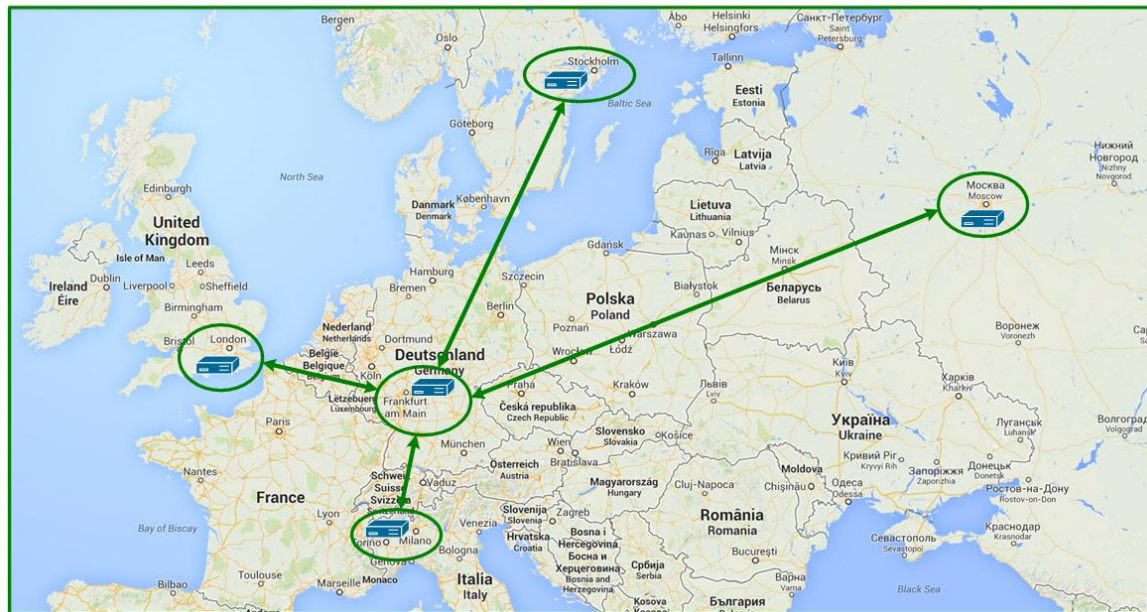
MIX / IT – Milan

LINX / UK – London

MSK-IX / RU – Moscow

DE-CIX / DE – Frankfurt

## Analyzed RRCs of the RIPE RIS Project





## Some more Details

---

- One bview per quarter
  - 20+ overall
- Filter
  - Only IPv6 routes
  - Only PA scopes (**not** PI)
  - Filter on neighboring AS (peering point)
  - Calculate number of prefixes per prefix length (/12 - /48)
- Analyze some parameters

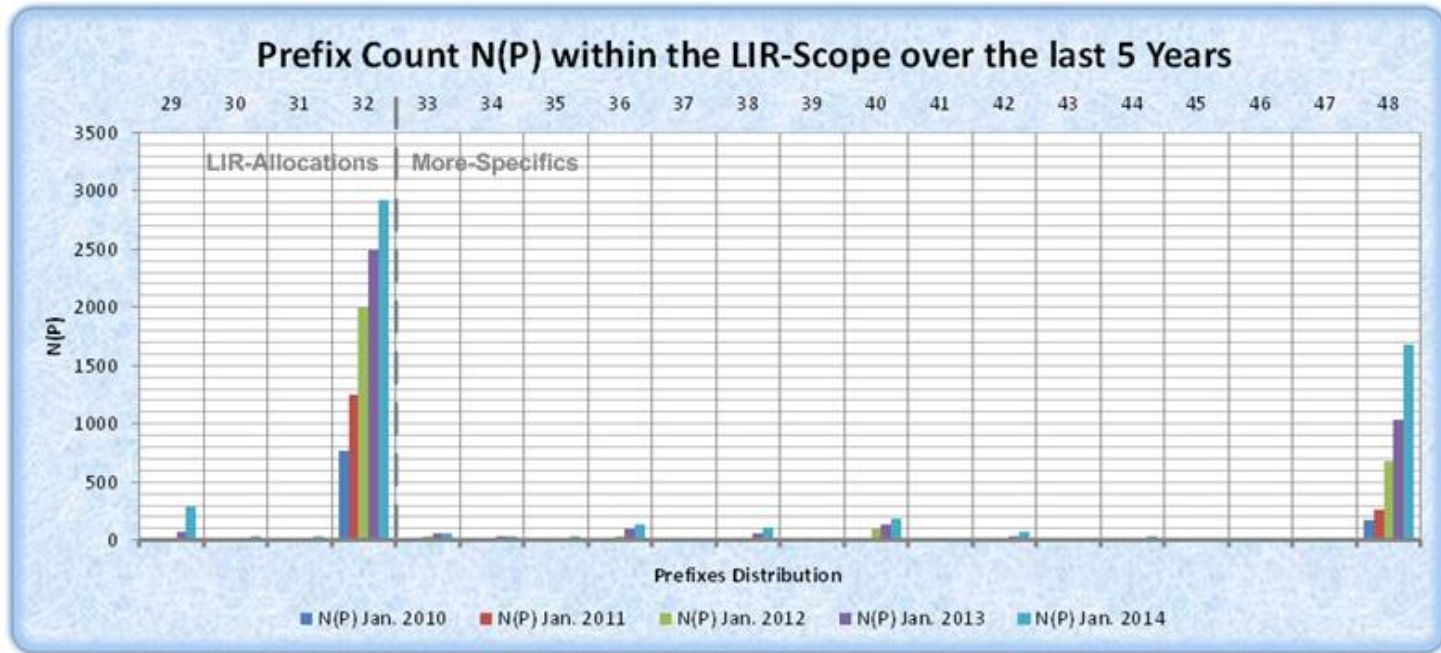


# Some Results

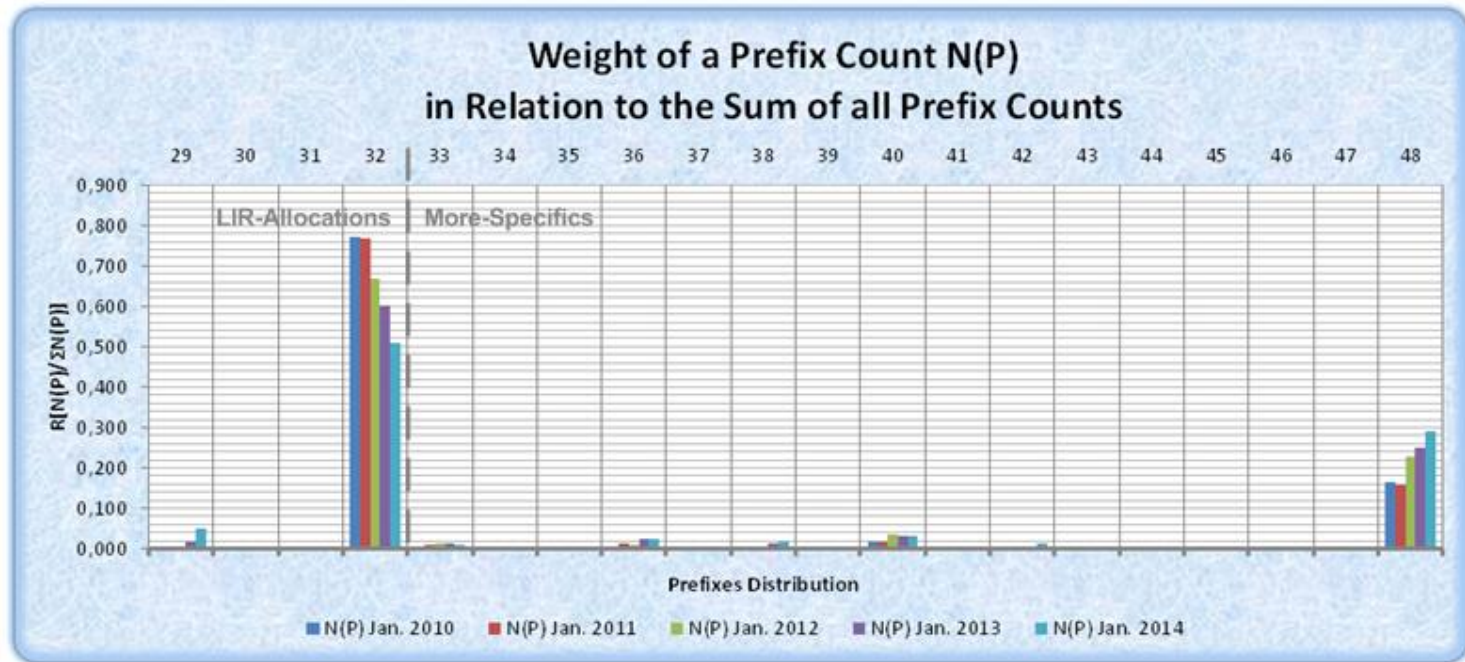
---



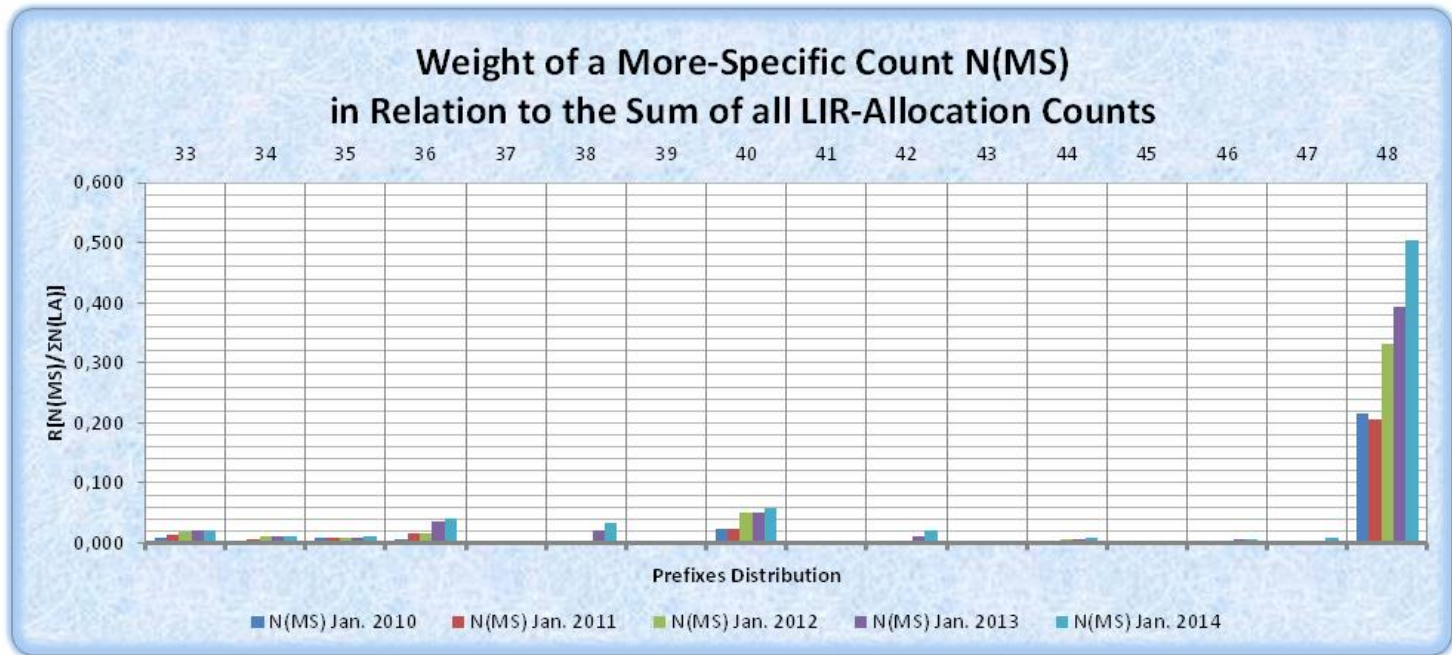
## Overall Numbers



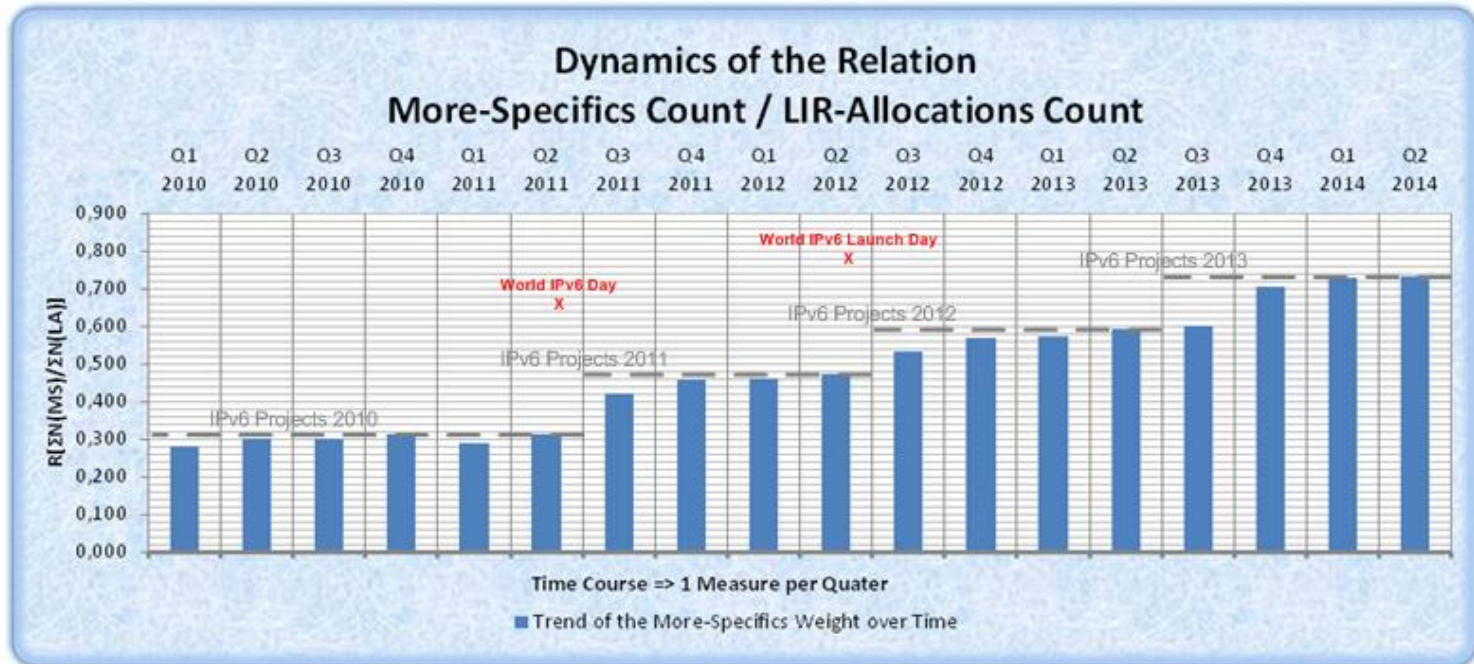
# Shares



## Weight of More-Specifics

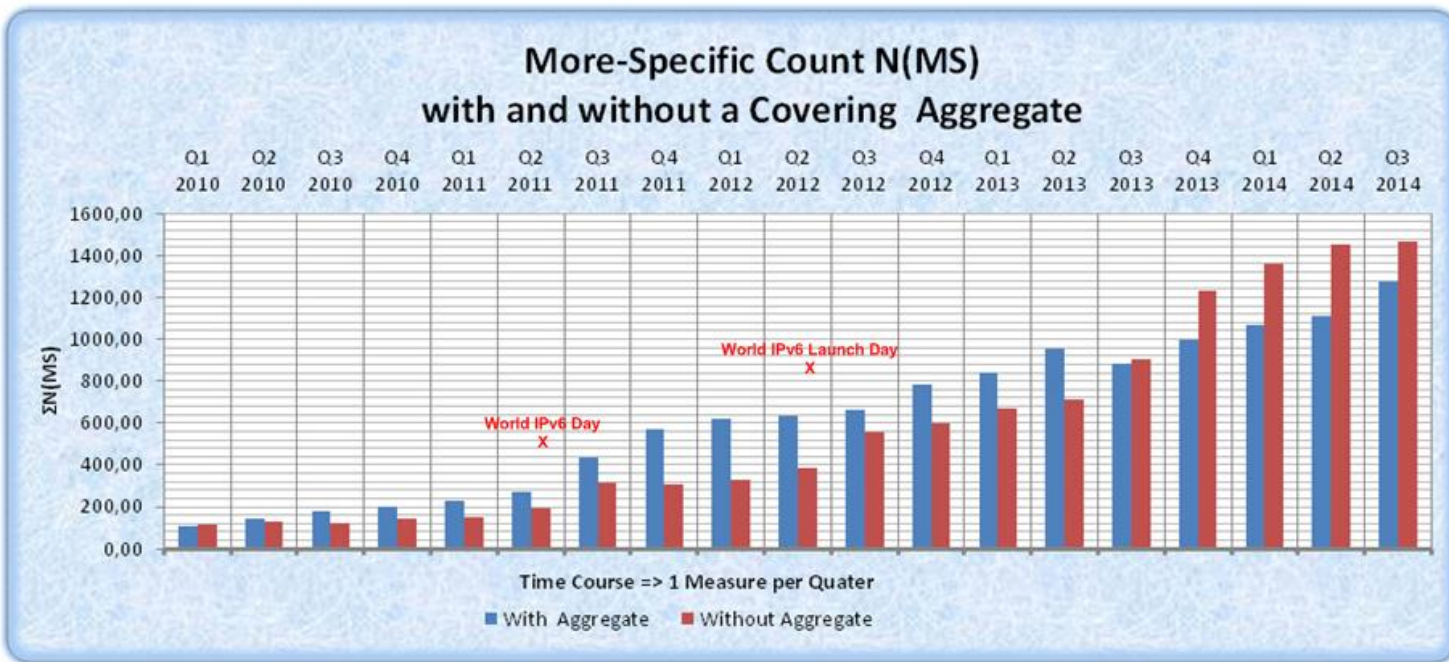


# Dynamics



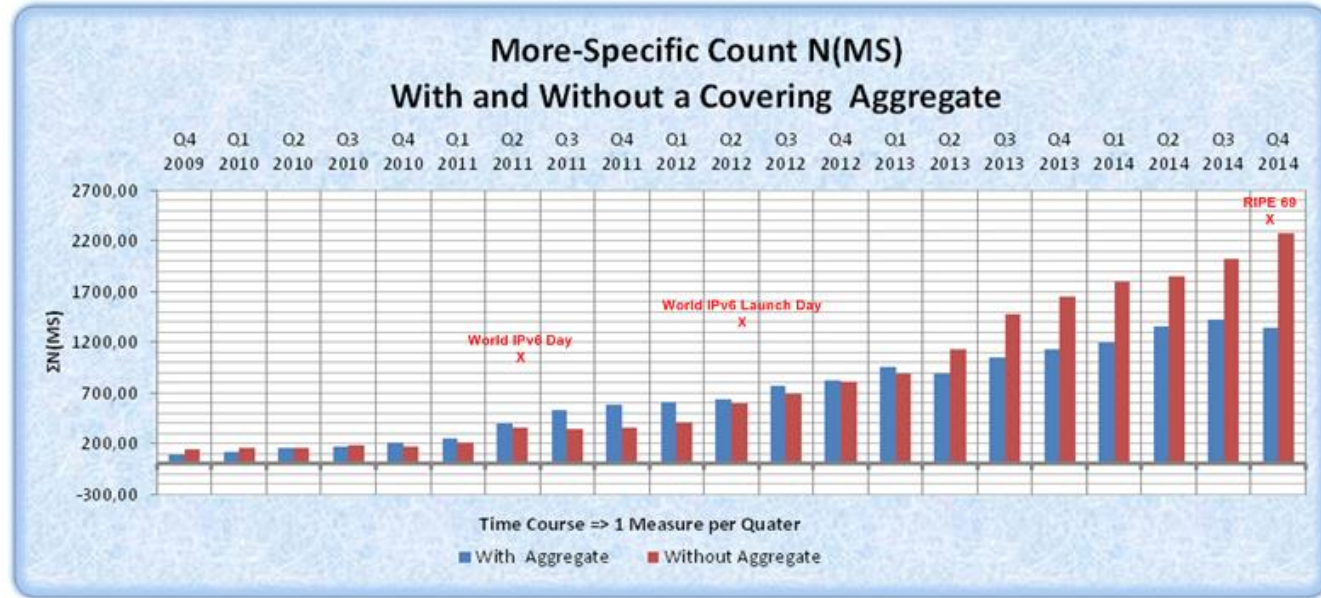


## With/-out Covering Aggregate



# Update

Very interesting numbers!  
Probably Geoff Huston is right:  
“the market will fix it”.





## Interim Conclusion



- Share of *more-specifics* in overall prefixes (in RIPE PA space) is growing.
  - /48 most prominent
  - Many *without* covering aggregate.
- Continuous trend since 2011.
- Over time second law of thermodynamics might realize.

# Strict Filtering Anyone? (I)

Global Routing Table of a Strict Filtering LIR in last 5 Years

Date	/19	/20	/21	/22	/23	/24	/25	/26	/27	/28	/29	/30	/31	/32	/33	/34	/35	/36	/37	/38	/39	/40	/41	/42	/43	/44	/45	/46	/47	/48
01.01.2010	1	2	2	1	0	1	1	2	1	1	0	0	2	780	0	0	5	0	0	0	0	0	0	0	0	0	0	0	0	91
01.04.2010	1	2	2	1	0	2	1	3	1	2	0	1	2	881	1	0	7	0	0	0	0	0	0	0	0	0	0	0	0	97
01.07.2010	1	2	2	1	0	2	1	3	2	3	1	1	2	990	1	0	7	0	0	0	0	0	0	0	0	0	0	0	0	96
01.10.2010	1	2	2	1	0	2	1	3	2	4	1	1	2	1090	1	0	7	0	0	0	0	0	0	0	0	0	0	0	0	99
01.01.2011	1	2	2	1	1	2	1	3	2	4	1	1	2	1261	0	0	7	0	0	0	0	0	0	0	0	0	0	0	0	104
01.04.2011	1	2	2	1	1	3	2	3	3	4	1	2	2	1514	2	0	6	0	0	0	0	1	0	0	0	0	0	0	0	118
01.07.2011	1	2	2	1	2	3	3	4	3	4	2	4	3	1760	1	0	6	0	0	0	0	0	0	0	0	0	0	0	0	177
01.10.2011	1	2	2	1	2	3	4	5	3	5	4	5	5	1894	1	0	6	0	0	0	0	1	0	0	0	0	0	0	0	138
01.01.2012	1	2	2	1	2	3	4	4	5	7	4	6	6	2019	1	0	6	0	0	0	0	1	0	0	0	0	0	0	0	142
01.04.2012	1	2	2	1	3	3	4	5	5	10	6	6	5	2111	2	0	10	0	0	0	0	1	0	0	0	0	0	0	0	146
01.07.2012	1	2	2	1	3	3	4	5	7	13	9	9	5	2246	0	0	10	0	0	0	0	0	0	0	0	0	0	0	0	151
01.10.2012	1	2	2	1	3	3	4	6	9	13	38	10	5	2361	0	0	10	0	0	0	0	0	0	0	0	0	0	0	0	152
01.01.2013	1	2	2	1	4	3	4	6	9	13	73	15	8	2514	0	0	10	0	0	0	0	0	0	0	0	0	0	0	0	153
01.04.2013	1	2	2	1	3	3	4	7	9	14	126	20	12	2640	0	0	9	0	0	0	0	0	0	0	0	0	0	0	0	154
01.07.2013	1	2	2	1	4	4	4	7	11	15	173	22	15	2740	0	0	9	0	0	0	0	0	0	0	0	0	0	0	0	155
01.10.2013	1	2	2	1	3	4	4	7	11	16	237	32	34	2839	0	0	9	0	0	0	0	0	0	0	0	0	0	0	0	155
01.01.2014	1	2	2	1	4	4	4	7	11	16	293	39	34	2940	0	0	9	0	0	0	0	0	0	0	0	0	0	0	0	159
01.04.2014	1	2	2	1	4	4	4	8	11	15	349	50	37	3037	0	0	9	0	0	0	0	0	0	0	0	0	0	0	0	162
01.07.2014	1	2	2	1	4	4	4	8	11	15	398	57	37	3125	75	43	39	193	12	88	7	215	8	69	1	151	17	21	24	1861

# Strict Filtering Anyone? (II)

Global Routing Table of a Strict Filtering LIR

Date	Total	/19	/20	/21	/22	/23	/24	/25	/26	/27	/28	/29	/30	/31	/32	/33	/34	/35	/36	/37	/38	/39	/40	/41	/42	/43	/44	/45	/46	/47	/48
01.01.2010	974	1	2	2	1	0	1	1	2	1	1	0	0	2	769	7	4	8	6	0	0	2	20	1	4	0	1	1	0	3	134
01.04.2010	1125	1	2	2	1	0	1	1	3	1	2	0	1	2	859	9	9	11	22	0	0	2	15	1	4	0	2	1	0	3	170
01.07.2010	1304	1	2	2	1	0	1	1	3	2	3	1	1	2	986	13	8	11	23	0	0	2	25	0	4	0	2	1	0	3	199
01.10.2010	1441	1	2	2	1	0	1	1	3	2	4	1	1	2	1080	17	7	12	24	1	0	2	29	0	6	0	3	1	3	3	232
01.01.2011	1649	1	2	2	1	1	1	1	3	2	4	1	1	2	1254	18	9	11	26	1	1	2	31	1	5	1	2	1	3	4	256
01.04.2011	2012	1	2	2	1	1	2	2	3	2	4	1	2	2	1505	26	10	11	18	1	6	2	47	2	6	0	4	1	8	6	334
01.07.2011	2563	1	2	2	1	2	3	3	4	3	4	2	4	2	1756	36	14	14	21	1	6	2	70	4	5	3	6	2	9	6	569
01.10.2011	n.a.																														
01.01.2012	3004	1	2	2	1	2	3	4	4	4	7	4	6	5	2001	39	23	16	34	1	7	3	97	4	5	3	13	3	9	7	682
01.04.2012	3188	1	2	2	1	3	3	4	5	4	10	6	6	5	2090	46	30	21	43	1	8	2	113	4	2	27	15	2	10	8	698
01.07.2012	3552	1	2	2	1	3	3	4	5	7	13	9	9	5	2224	45	31	22	68	2	40	3	124	4	32	25	17	1	12	9	807
01.10.2012	3840	1	2	2	1	3	3	4	6	9	13	35	10	5	2343	46	32	21	78	1	45	2	136	4	32	29	14	2	13	8	936
01.01.2013	4158	1	2	2	1	4	3	4	6	9	13	69	15	8	2494	53	31	21	94	1	55	3	140	4	32	0	15	2	15	9	1048
01.04.2013	4492	1	2	2	1	3	3	4	7	9	13	123	20	12	2612	53	32	20	111	1	44	2	150	7	42	35	16	2	14	10	1137
01.07.2013	4776	1	2	2	1	4	4	4	7	11	14	170	22	15	2720	57	30	20	110	1	79	3	159	9	22	37	16	2	15	14	1221
01.10.2013	5381	1	2	2	1	3	4	4	7	11	15	231	32	32	2810	63	38	26	132	2	110	3	181	8	67	0	22	3	18	24	1523
01.01.2014	5735	1	2	2	1	4	4	4	7	11	15	289	39	33	2917	63	34	37	135	3	109	4	196	8	74	0	33	3	18	24	1661
01.04.2014	6034	1	2	2	1	4	4	4	8	11	15	340	49	35	3013	65	40	39	165	4	105	7	201	8	74	1	62	17	20	24	1710
01.07.2014	6365	1	2	2	1	4	4	4	8	11	15	391	56	36	3104	67	43	39	190	11	87	7	213	8	69	1	65	16	20	24	1819
01.10.2014	6715	1	2	2	1	4	4	4	8	12	15	448	61	36	3175	71	42	38	208	10	89	9	227	8	68	1	66	16	53	25	1960
01.01.2015	4093	1	2	2	1	4	4	4	8	12	15	506	63	38	3260	0	0	9	0	0	0	0	0	0	0	0	0	0	0	0	164

Now, what does all this mean?

---





- While those numbers indicate certain directions, right now there's a huge amount of unsettledness amongst Enterprise-LIRs, leading to
  - Operational effort.
  - Taking decisions which create *tragedy of the commons*.
  - Postponing IPv6 deployment.
- All these are disadvantageous for the whole community
  - Not least because Enterprise-LIRs are usually customers of several Transit-LIRs.

## The Underlying Problem

---



- Dilemma / tragedy of the commons as of above.
- How much deaggregation are we willing to accept and what does this mean for filter policies?

# Ways to Address The Dilemma

---



- Discuss the problem and make involved parties aware of the needs of each other.
  - This is exactly why I stand here right now ;-)
  - => “compromise between practitioners”
- Create predictability & transparency by policy
  - Has not worked in the past.
  - Different players have different agendas.
    - What normative value has \$RIPE\_POLICY?
- Solve commercially
  - Extra fee for extra announcements?
  - Hence \$ENTERPRISES (which have plenty of money anyway, right?) could “pay their due share”.
- Wait (and hope problem goes away).
- Others
  - <https://tools.ietf.org/id/draft-van-beijnum-grow-controlled-deagg-00.txt>

## Compromise Between Practitioners

---

What Could this Look Like?



- Proposal: /40
- Maximum deaggregation  $2^8$
- Overall weight of /40 growing.



## Conclusion



- If we don't solve this issue, everybody loses!
- The earlier we tackle it, the better for us, as a community
  - And for IPv6 deployment in enterprise space, which is a non-negligible part of the Internet out there...

There's never enough time...

**THANK YOU...**



**...for yours!**

Slides:

<https://www.insinuator.net>

## References

---



- RIPE Routing Working Group  
Recommendations on IPv6 Route Aggregation
  - <http://www.ripe.net/ripe/docs/ripe-532>
- [Atlas2012]
  - <https://labs.ripe.net/Members/emileaben/ripe-atlas-a-case-study-of-ipv6-48-filtering>
- [Bayer2010]
  - <https://labs.ripe.net/Members/dbayer/visibility-of-prefix-lengths>