Friedwart Kuhn

Digital unterschrieben von Friedwart Kuhn
DN: cn=Friedwart Kuhn, givenName=Friedwart Nikodemus Michael, sn=Kuhn, serialNumber=DSV0000026276
Datum: 2014.11.05 19:20:37 +01'00'

ERNW Newsletter 44 / November 2014

Dynamics of IPv6 Prefixes within the LIR Scope in the RIPE NCC Region

Version:	1.0
Date:	11/5/2014
Author(s):	Dr. Roland Langner & Nicolas Schätzle (Burda Digital Systems), Enno Rey (ERNW)

1 Abstract

Just recently there have been constant discussions within the Internet community which impact a strict BGP filter of transit ISPs (LIRs = local Internet registries) might have on the actual reachability of address space. Which organizations would be affected by strict filtering and what options are available to the players involved?

The objective of the present study was to analyze the IPv6 prefix dynamics of the default-free zone (DFZ) within the PA/LIR scopes of the RIPE NCC region. For this purpose, data of the routing information service (RIS) project from 2010 to 2014 was evaluated and discussed. It could be argued that the IPv6 prefix dynamics within the DFZ space of the LIR scopes are developing steadily and with increasing weight towards more specifics, i.e. in the direction of IPv6 prefixes between /33 and /48. Apparently there is already a significant effort within the LIR community to further deaggregate LIR allocations according to individual policies.

This general trend towards more-specifics and subsequently to more equal distribution in the DFZ phase space seems to be a spontaneous process under the current conditions of the LIR community and sometimes exhibits chaotic behavior within its fine structure. For further discussion and interpretation of the results it has been postulated in a first approximation that increasingly two columns are establishing within the LIR community which represent two communities of interest in their respective environment – the pillar of "transit LIRs" (traditional ISPs) as well as, like we call them, "enterprise LIRs". Both groups operate LIR infrastructures that are tailored to their own business models which in turn might cause conflicts of interest.

The results of this study raise the question on how strict BGP filter on allocated LIR scope could violate individual LIR policies and how this might affect the overall deployment of IPv6. As part of discussion of the results we will specifically address this question and provide an outlook on possible compromises.

Table of Content

List of Figures

List of Tables

2 INTRODUCTION

The chief scientists of regional Internet registries (RIRs) such as the RIPE NCC or the APNIC are periodically publishing analyses of the DFZ development in their administrative regions and thereby give interested parties insight into core functions of the global Internet routing. Publications of this type provide information about the current state of BGP routing updates and convergence times, respectively, the availability of address space or various qualitative and quantitative aspects of the prefix distribution within the DFZ. In particular, the IPv6 prefixes of the global routing table can be assigned to different scopes, which are primarily managed by the IANA. So the IANA initially assigns IPv6 address space to RIRs, which in turn then assign address space to LIRs (companies with a LIR contract) or end users (companies without a LIR contract). In this way LIRs receive provider aggregated (PA) scope, the so-called LIR allocations that serve especially for further distribution of address space, while end users can also directly get provider independent (PI) scope for their own use. The DFZ represents the entirety of all scopes propagated by means of BGP.

2.1 Default-free Zone (DFZ):

The default-free zone is commonly understood as the ensemble of all via BGP connected autonomous systems (ASs) on the Internet in which any destination can be reached without a default route. Accordingly, DFZ routers have a so called "full BGP table" available which is currently about 500.000 IPv4 and 18.000 IPv6 routes (status in July 2014) and is known as "Internet routing table" or "global routing table".

Based on experiences from the IPv4 past as well as current discussions on the IPv6 deaggregation potential, it is a strong desire in the LIR community to keep the global routing table as compact and small as possible. Background of this policy is not least the investment protection of hardware resources (memory/TCAM and CPU), the stabilization and minimization of BGP convergence times or, more generally, the practice of the simplicity principle (RFC 3439) [1] to name only a few aspects. A commendable overview of a more detailed range of aspects is provided by an IAB workshop from 2007, which was documented as RFC 4984 [2]. Protagonists of a prefix conserving policy are found especially in the fraction of traditional Internet service providers (ISPs), since an exponential growth of the DFZ would take an immediate impact on their business and above all on their cost-model.

Assuming that a hierarchical routing across RIR regions in IPv6 address space would be feasible a thought experiment is to compare the order of magnitude of the IPv4 and IPv6 deaggregation potentials. Hierarchical routing would demand a strict aggregation on the border between the RIR domains. The address space of another RIR domain would thus only be seen as a conglomerate of some aggregated prefixes. Table 1 and Table 2 compare the theoretically possible deaggregation potential of IPv4 and IPv6 [3] address space.

Status Quo of the Default-free Zone	Ca. 500.000 Prefixes	July 2014
Maximum Deaggregation Count	Ca. 16x 10^6 Prefixes	theoretically possible

Table 1: Potential of Deaggregation in the entire IPv4 Address Space up to a /24 Prefix Size

Status Quo of the Default-free Zone	Ca. 18.000 Prefixes	July 2014
Maximum Deaggregation Count	Ca. 70x 10^9 Prefixes	theoretically possible

Table 2: Potential of Deaggregation in the (per IANA as of mid 2014) assigned RIPE IPv6 Address Space (PI & PA), up to /48 Prefix Size

Just the IANA assigned address pool 2001:0600::/23 from which the RIPE NCC currently generates PI scope (2001:678::/29) has already a larger deaggregation potential with approximately 32x 10^6 than the entire IPv4 address space. Furthermore, one should keep in mind that PI scope basically removes the possibility of aggregation, as it is directly assigned by the RIPE NCC to the end user. Hence, prefixes of PI scope will always appear as independent routes in the DFZ without covering aggregate. Looking at this magnitude it becomes clear that the LIR community has a legitimate interest in avoiding a further deaggregation of address space within the PA scope. Thus, while PI scope is not subject to a BGP filter policy < /48, as a rule, filtering of PA scope (allocations) reflects the desire to counteract the risk of deaggregation. However, exactly this addresses the core problem of an uncertainty of planning for companies with LIR contract which rely on deaggregation of PA scope because of their individual networks and policies. Table 3 shows a summary of the address space that has been allocated by the IANA to the RIPE NCC and based there on the PA allocations and PI assignments of the RIPE NCC with regard to their prefix length [4] (status in July 2014).

RIPE designated Scope		Prefixes Size Distribution of the RIPE NCC allocated and assigned IPv6 Address Space in July 2014																													
Prefix	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
2001:0600::/23	01.07.1999	---	---	---	---	---	---	---	---	---	---	16	1	0	48	---	---	---	---	---	---	---	1	0	1	1	2	7	12	32	1.766
2001:0800::/23	02.05.2002	---	---	---	---	---	---	---	---	---	---	16	3	2	40	---	---	---	---	---	---	---	---	---	---	---	---	---	---	---	---
2001:0a00::/23	02.11.2002	---	---	---	---	---	---	---	---	---	---	11	0	0	50	---	---	---	---	---	---	---	---	---	---	---	---	---	---	---	---
2001:1400::/23	01.02.2003	---	---	---	---	---	---	---	---	1	0	6	1	0	55	---	---	---	---	---	---	---	---	---	---	---	---	---	---	---	---
2001:1600::/23	01.07.2003	---	---	---	---	---	---	---	---	1	0	6	1	0	24	---	---	---	---	---	---	---	---	---	---	---	---	---	---	---	---
2001:1a00::/23	01.01.2004	---	---	---	---	---	---	---	---	---	---	12	1	0	51	---	---	---	---	---	---	---	---	---	---	---	---	---	---	---	---
2001:1c00::/22	04.05.2001	---	---	---	---	1	---	---	---	---	---	---	---	---	---	---	---	---	---	---	---	---	---	---	---	---	---	---	---	---	---
2001:2000::/20	04.05.2001	---	1	---	---	---	---	---	---	---	---	---	---	---	---	---	---	---	---	---	---	---	---	---	---	---	---	---	---	---	---
2001:3000::/21	04.05.2001	---	---	---	---	---	---	---	---	---	---	---	---	---	---	---	---	---	---	---	---	---	---	---	---	---	---	---	---	---	---
2001:3800::/22	04.05.2001	---	---	---	---	---	---	---	---	---	---	---	---	---	---	---	---	---	---	---	---	---	---	---	---	---	---	---	---	---	---
2001:4000::/23	11.06.2004	---	---	---	---	---	---	---	---	---	---	7	0	0	53	---	---	---	---	---	---	---	---	---	---	---	---	---	---	---	---
2001:4600::/23	17.08.2004	---	---	---	---	---	---	---	---	---	---	---	---	---	---	---	---	---	---	---	---	---	---	---	---	---	---	---	---	---	---
2001:4c00::/23	15.10.2004	---	---	---	---	---	---	---	---	1	0	7	0	0	25	---	---	---	---	---	---	---	---	---	---	---	---	---	---	---	---
2001:5000::/20	10.09.2004	---	---	1	---	---	---	---	---	---	---	---	---	---	---	---	---	---	---	---	---	---	---	---	---	---	---	---	---	---	---
2003:0000::/18	12.01.2005	1	---	---	---	---	---	---	---	---	---	---	---	---	---	---	---	---	---	---	---	---	---	---	---	---	---	---	---	---	---
2a00:0000::/12	03.10.2006	1	1	1	2	3	3	3	7	8	5	1.502	56	37	5.531	---	---	---	---	---	---	---	---	---	---	---	---	---	---	---	---

Table 3: IANA - IPv6 Global Unicast Address Assignments

2.2 Strict IPv6 BGP Filtering

As early as 2002 Gert Döring from SpaceNet[1] published a document "IPv6 BGP Filter Recommendations" [5] that presented the LIR community some BGP filter practice. Among other things this document discusses the possibility of strict filtering on PA scope respectively on LIR allocations assigned by the RIPE NCC. Although the current version of the document warns of strict filtering and draws attention to the risks involved, there seem to be LIRs on the Internet that actually employ strict filtering. Anecdotal cases [6] lead to the conclusion that strict filters reportedly led to non-reachability of certain scopes. To make matters worse, even an extension of the allocation from /32 to /29 would not change the fundamental problem of strict filtering since required filter information could be dynamically acquired from the RIPE database. Once a RIPE allocation generated more-specifics in terms of deaggregation regardless of the prefix length /32 or /29, those could be eliminated by strict BGP filters. In these cases one has to rely on a covering aggregate.

An excerpt from the strict filter set which includes the RIPE scopes confirms that almost no more-specifics > /32 would be permitted with the exception of PI scope and some of /35 prefixes. Accordingly, for example, a company that has been assigned a /32 allocation from the range 2A00::/12 upon conclusion of a LIR contract, should definitely propagate the /32 prefix everywhere as covering-aggregate.

```
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
```

2.2.1 Postulate of the Two LIR Columns

The next two pictures show two statistics published by the RIPE NCC [7] [8], which reflects the dynamics of LIR-contracts and LIR-allocations, respectively. As can be seen, LIR allocations have significant growth rates since 2012 making IPv6 a driving force of new LIR activities. It is hard to believe that this significant increase in LIR members is only due to traditional transit ISPs. It can be rather assumed that a second fraction is increasingly established in the LIR community - the conventional non-transit company that wants to secure a larger IPv6 address space for future use.

Figure 1: LIRs With and Without IPv6

1 https://www.space.net/

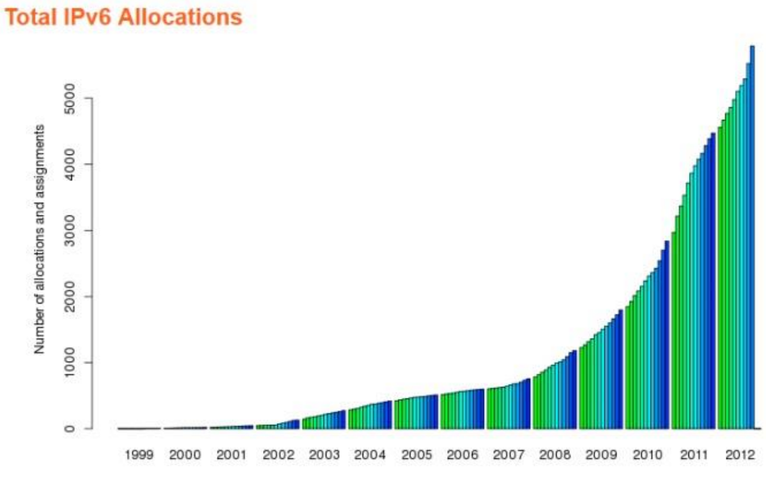

Figure 2: Total IPv6 Allocations

For interpretation of the results both columns were distinguished within the LIR community, the transit LIRs which include Internet carriers and traditional service providers as well as enterprise LIRs whose LIR business activity is usually limited to their own profit centers, divisions and departments. Such companies operate LIR infrastructures that vary greatly from those of traditional transit LIRs and thus, have different attributes, features and characteristics. The non-transit LIR organizations often work with stateful firewalls at their peering points and hence are dependent on bidirectional symmetrical forwarding. Send and return paths of communication must be done through the same Internet breakout / firewall. For traditional ISPs this is not true in most cases.

In the present work, we now venture a detailed look into the dynamics of the IPv6 prefix distribution of LIR scopes the RIPE NCC has allocated to the LIR community. For this purpose the archived RIS data (so called bview files) from a total of five remote route collectors (RRCs) were evaluated and analyzed from 2010 to 2014. As shown in Figure 3 the investigated RRCs form a cross through Europe to deliver a representative average of the interesting observables.

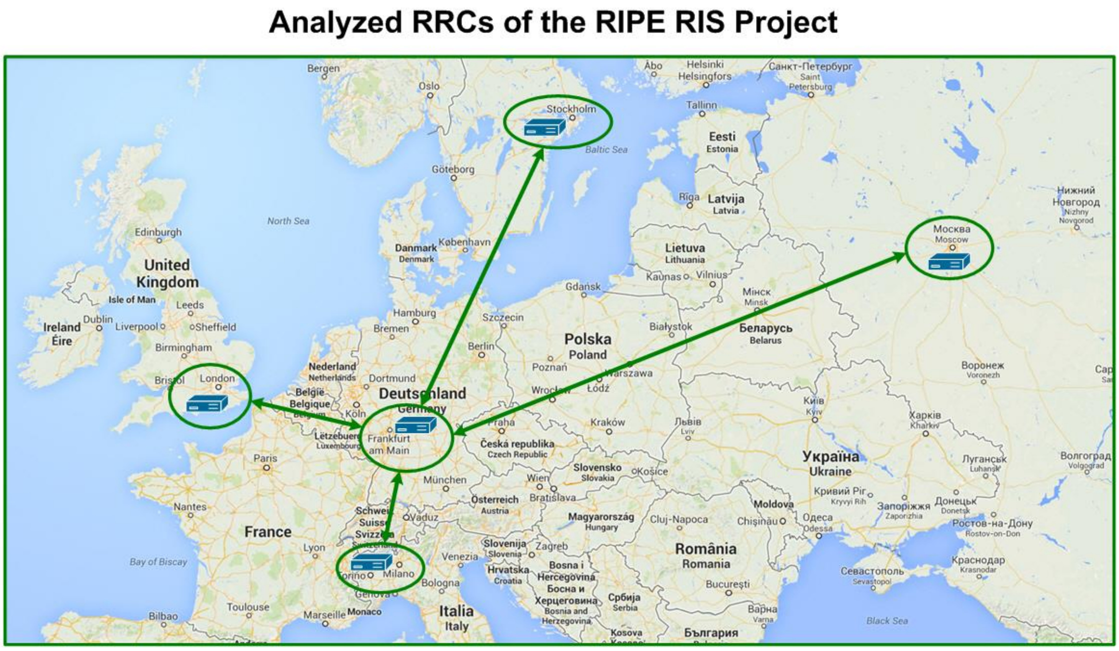

Figure 3: Analyzed RRCs of the RIPE RIS Project

3 EXPERIMENTAL

For analysis of the DFZ prefix distribution, the bview files of the following Internet exchange points (IXPs) respectively its RRCs were examined:

- Netnod / Schweden – Stockholm
- MIX / Italy – Milan
- LINX / England – London
- MSK-IX / Russia – Moscow
- DE-CIX / Germany – Frankfurt

In each case the first bview file per quarter was evaluated, i.e. for each RRC peering point we analyzed a total of 19 images of the global routing table from 2010 - 2014 (01. Jan / 01. May / 01. Jul / 01. Oct). However, only those peering points have been considered that have reached at least 90% of the maximum /32 prefix count per quarter evaluation. Accordingly, peering points with a significantly incomplete global routing table were rejected.

In the following the individual working steps are described by which the bview files were sequentially filtered and further processed to gain access to the relevant raw data:

- Step 1: Setting a filter on all IPv6 entries
- Step 2: Setting a positive filter on all PA scopes (s. Tab. 2)
- Step 3: Setting a filter on the neighboring ASN (peering point)
- Step 4: Calculation of the prefix number per prefix length (/12 - /48)

The raw data thus obtained gave the following observables per peering point (global routing table):

3.1 Description of the Observable N(P)

1. N (P) indicates the prefix count for a particular prefix length.

2. ΔN(P)/Δt shows the time variation of the prefix count per quarter. For this purpose, the difference between the prefix values to the following quarter ($Q_{n+1} - Q_n$) was calculated. The size is a measure of the rate of change of the observable.

3.2 Description of the Observable R[N(P)/ΣN(P)]

1. R[N(P)/ΣN(P)] indicates the ratio of a prefix N(P) to the sum of all prefixes. The observable reflects the weight of a prefix in relation to the total number of all prefixes between /12 and /48.

2. ΔR[N(P)/ΣN(P)]/Δt shows the time variation of the ratio R[N(P)/ΣN(P)] per quarter. For this purpose, the difference between the ratio values to the following quarter ($Q_{n+1} - Q_n$) was calculated. The size is a measure of the rate of change of the observable.

3.3 Description of the Observable R[N(MS)/ΣN(LA)]

1. R[N(MS)/ΣN(LA)] indicates the ratio of a more-specific N(MS) between /33 and /48 to the sum of all LIR allocations between /12 and /32. The observable reflects the weight of a more-specific in relation to the total number of LIR allocations.

2. ΔR[N(MS)/ΣN(LA)]/Δt shows the time variation of the ratio R[N(MS)/ΣN(LA)] per quarter. For this purpose, the difference between the ratio values to the following quarter ($Q_{n+1} – Q_n$) was calculated. The size is a measure of the rate of change of the observable.

3.4 Description of the Observable R[ΣN(MS)/ΣN(LA)]

1. R[ΣN(MS)/ΣN(LA)] indicates the ratio of the sum of all more-specifics between /33 and /48 to the sum of all LIR allocation between /12 and /32. A value of 1 corresponds to an equal distribution of both scopes. Increasing values indicates a shift in favor of the more-specifics.

2. ΔR[ΣN(MS)/ΣN(LA)]/Δt shows the time variation of the ratio R(ΣN(MS)/ΣN(LA)) per quarter. For this purpose, the difference between the ratio values to the following quarter ($Q_{n+1} – Q_n$) was calculated. The size is a measure of the rate of change of the observable.

As part of the evaluations the mean values per IXP/RRC were initially calculated over all peering points and analyzed. Since the analyzes doesn't show significant DIFFs, the mean values over all examined IXPs/RRCs were calculated to get most representative values of an observable in the administration region of the RIPE NCC.

In order to represent the dynamics of the observables, i.e. the temporal changes of observables per quarter, as detailed as possible so-called phase portraits were created requiring a more precise explanation. More information about the terms phase space and phase portrait can be found e.g. in Wikipedia [9] [10].

3.5 What is a Phase Portrait?

Phase portraits go back to the French mathematician Henri Poincare and generally serve for the analysis of dynamic systems. Here, the phase portrait reflects graphically the so-called phase space that represents all possible states of a physical system in terms of an observable. Each state of the system corresponds to a point in phase space. The migration of a time-dependent state through the phase space is called trajectory. The trajectory thus describes the path/line of the physical system (observable) through the phase space. Phase portraits were e.g. known in the context of chaos theory to represent attractors and repellors, or only to project the dynamics of an oscillating pendulum. In the case of the pendulum, the speed of the pendulum ($v=\Delta x(t)/\Delta t$) is plotted against their displacement ($x(t)$), i.e. the time variation of the observable is plotted against the observable itself in a Cartesian coordinate system. Analogously, various prefix based observables were presented and interpreted in DFZ phase space in the present study.

3.6 Interpretation of Phase Portraits

To be able to read a phase portrait correctly, we turn our gaze to a harmonically oscillating pendulum. The following figure represents the phase space of a harmonic oscillator.

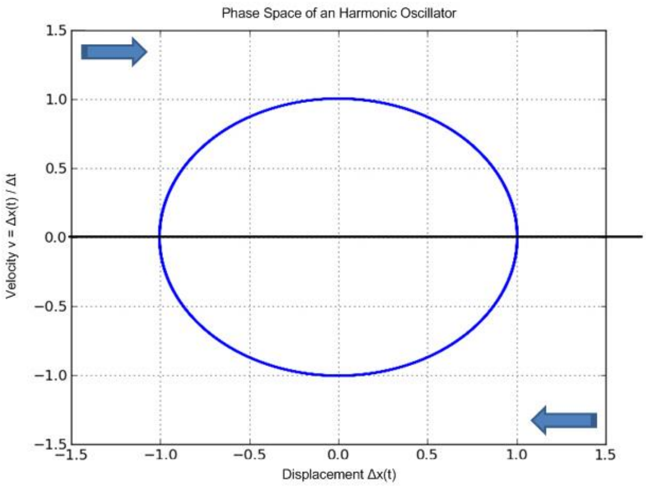

Figure 4: Phase Space of a Harmonic Oscillator

One of the special features of phase portraits is represented by the x axis itself respectively the abscissa which is to be understood in the sense of a border. All states above the abscissa have a positive velocity and in terms of their trajectory, i.e. the displayed line that reflects the time change of states, are to be read from left to right. The pendulum oscillates with positive speed in the positive direction of the displacement. All states below the abscissa then have logically a negative velocity, and the trajectory has to be interpreted from right to left. The pendulum oscillates with negative speed in the negative direction of the displacement. In this way, the oscillatory motion of a pendulum can be represented very simply in a plane of the phase space. The blue arrows indicate the time course and thus the reading direction of the trajectory. In other words, the trajectory with positive ordinate values is read in terms of an increase of the observable and negative ordinate values vice versa. In an analogous manner, other observables can be displayed and interpreted in a phase portrait. Phase portraits are still playing an important role in the chapter "Results and Discussion".

4 Results and Discussions

4.1 Dynamics of the Prefix Distribution within the LIR Scope

As already indicated in the introduction the evaluations of the bview files focused on the dynamics of the prefix distribution in the LIR scope. The study was to show how LIR allocations and their more-specifics behaved over the period of 2010 - 2014 and whether we can detect a trend. For this purpose, the prefix counts were compared in the range between /29 and /48 in Figure 5. Each prefix shows in dependence on its numerical value up to five bars, each corresponding to January of the years 2010 to 2014. Because prefix lengths less than /29 yielded no visible results, representation has been waived. Likewise prefix lengths greater than /48 are not included as they are generally not routed in the DFZ.

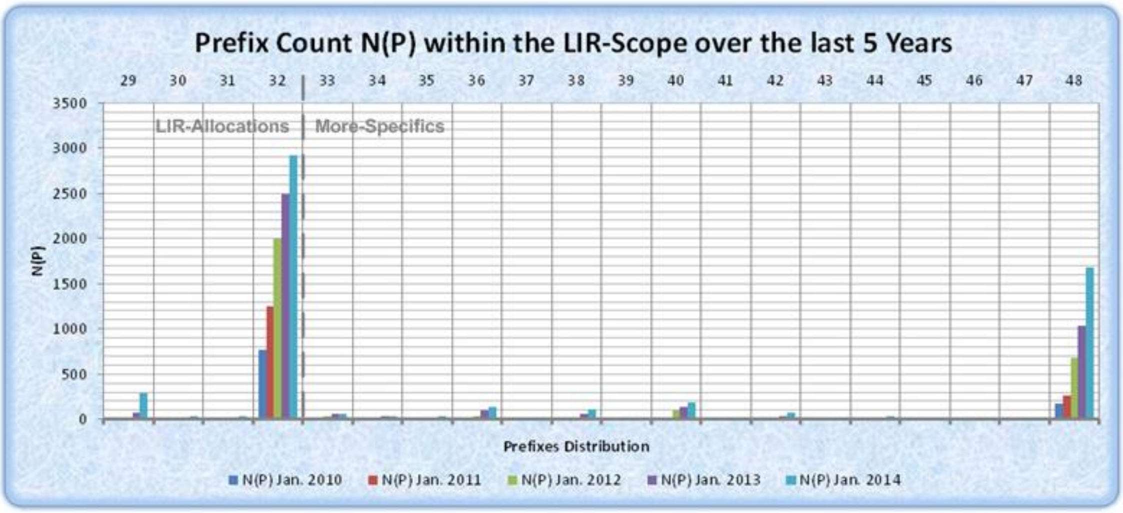

Figure 5: Prefix Count N(P) within the LIR-Scope over the last 5 Years

As expected and compliant with the RIPE policies, the /32 prefix and thus the smallest possible LIR allocation is the most represented. Over the last five years a steady increase of the prefix can be recognized. Interestingly, also the /29 LIR allocation seems to increase significantly in 2014. Within the scope of the more-specifics the /48 prefix establishes as the strongest representative and reaches well over 50% of /32 LIR-allocations. This /48 prefix is not to be confused with the RIPE assignment from the provider independent scope that was not subject of the present investigation (s. Experimental). The /48 prefix from the LIR scope should be in most cases rather attributed to the desire of LIRs to only announce those more-specifics which are absolutely necessary for their productive environments. A potentially underlying security aspect was already commented on by the RIPE NCC in the Policy RIPE-399 [11] as follows.

»The result is that some providers work around their concerns about the relative lack of routing system security by simply announcing the smallest acceptable prefix. This means that no other autonomous system can announce more specific versions of the same prefixes thereby causing a denial of service on the legitimate user of the address space. However, the authors are aware of very few such incidents being recorded, so deaggregation for security reasons seems a somewhat overly unfriendly activity compared with the potential risk.«

In addition to the /48 prefixes some smaller peaks at /40, /36 and /38 are seen especially in the last two years. These prefix lengths seem to correspond more to a moderate deaggregation that may have different causes according to individual environments within the LIR community which we will discuss in more detail later.

If we now disregard the absolute numbers and rather consider the relative weights of the individual prefixes in relation to the total number of prefixes, then there seems to be a similar picture at first glance. Upon closer examination of Figure 6 , however, show interesting deviations.

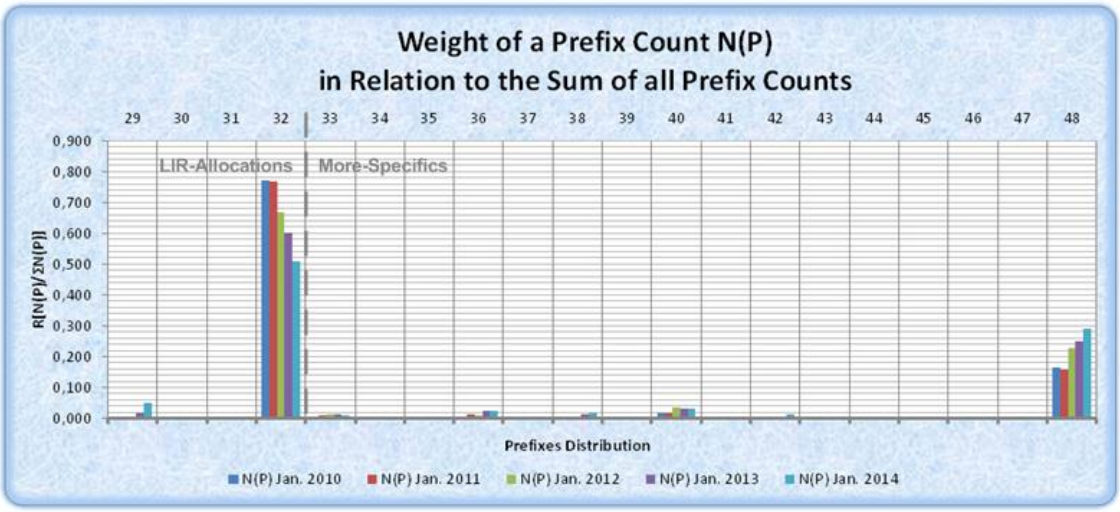

Figure 6: Weight of a Prefix Count N(P) in Relation to the Sum of all Prefix Counts

As the chart shows, the weighting of the /32 prefix in the years 2010 and 2011 still remains relatively stable while significantly breaking down in the three following years. It seems evident that from 2012 onwards a marked decrease in /32 weighting occurs. In contrast, the weighting of the /29 allocation seems unimpressed to rise. This result may suggest that new LIR members of the RIPE NCC increasingly take a /29 allocation in claim to miss the more-specifics discussion and the impact of strict filter. As already pointed out in the introduction such logic might pose certain risks. The RIPE NCC itself stated in the policy RIPE-532 [12] as follows on the subject:

»There is a valid need for some LIRs to advertise more than one IPv6 PA prefix. As either obtaining more address space and advertising more /32 prefixes, or advertising more specific prefixes within an already allocated /32 have the same impact on the routing table, it is suggested that the latter approach is taken to prevent address space wastage.«

While the weighting of the /32 allocation decreases sharply from 2012, that of /48 prefixes seems to increase over the same period of time. Only in the years 2010 and 2011 there was a minor slump which is possibly due to aggregation of /48 prefixes during this period. The smaller peaks around /40, /36 and /38 remain unimpressed.

As we have seen from the results in Figure 5, all prefixes occurring significantly in the LIR scope seem generally to increase with time. Now let's take a more detailed look at the scope of more-specifics. Figure 7 shows the relative weighting of a more-specific in relation to the total number of LIR allocations. Consequently, if a more-specific should increase faster than the total of LIR allocations one would expect a curve with a positive gradient in this chart. Accordingly, for more-specifics that increase with the same speed one would expect a steady course with a gradient of zero and for those that develop more slowly a decreasing trend with a negative gradient.

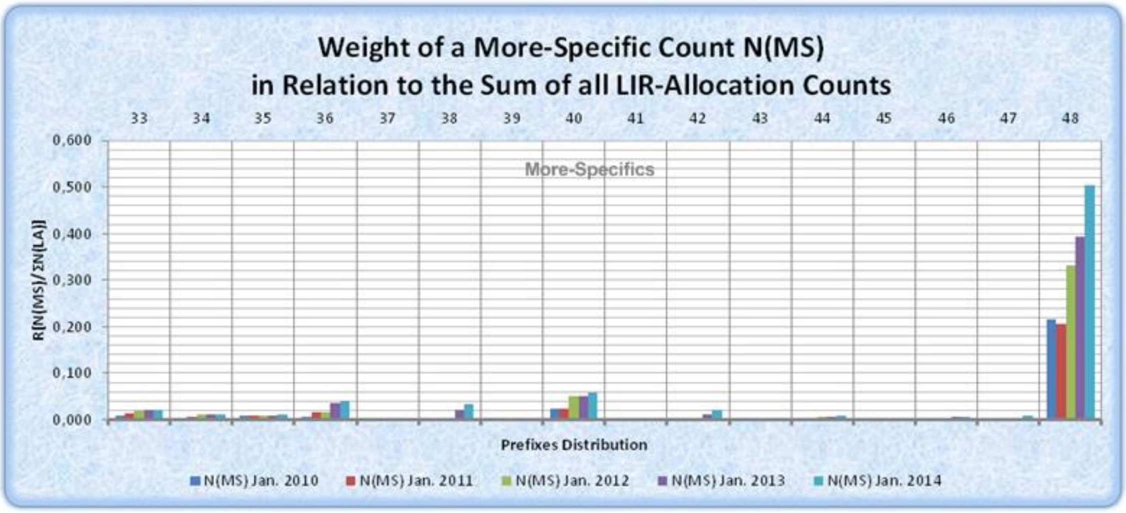

Figure 7: Weight of a More-Specific Count N(MS) in Relation to the Sum of all LIR-Allocation Counts

While the more-specifics /33, /34 and /35 are developing almost in line with the LIR allocations in the last three years and therefore have the same growth rates as the LIR allocation itself, in particular the /48 prefix seems to increase much faster. Although we recognize in 2010 - 2011 a period in which a larger quantity of /48 prefixes has apparently been aggregated or filtered, the prefix count significantly increases again in the three subsequent years. In addition, four other scopes are visible which gain in weight over the last few years the more-specifics /40, /36, /38 and /42. All other prefixes are currently negligible yet.

In Figure 7 two approaches of deaggregation can be observed. The first one starts at the right edge of the graph and corresponds, as already mentioned, to the policy of only announcing that scope which is essential for the productive IPv6 environment of the company. In this way you generate a maximum of prefixes in the LIR scope over time which ultimately cannot be objective of the LIR community. Each LIR with a /32 allocation would thus have a deaggregation potential of 2^16 or theoretically 65.536 prefixes to enrich the DFZ. The second approach starts at the left edge of the graph and keeps the focus to a minimum deaggregation by propagating only larger IPv6 blocks into the DFZ as needed. And this demand will arise quickly in the LIR community, as currently composed. The RIPE NCC has already addressed the issue in the policy RIPE-532 [12]:

```
»However, as understood in the RIPE Routing Working Group's Recommendations on Route
Aggregation [RIPE-399], 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.

Some reasons for wanting to advertise multiple prefixes from a PA allocation could be:

    1. The LIR has several networks that are not interconnected.
    2. Traffic engineering: A single prefix that covers an LIR's entire customer base
       may attract too much traffic over a single peering link.«
```

In addition to the reasons above, there are further situations in which an LIR is dependent on a deaggregation of its allocation.

- The LIR operates several networks / peerings that are interconnected with insufficient bandwidth.
- The LIR operates several networks / peerings behind stateful firewalls that allow only symmetric bidirectional communication.
- The LIR operates several networks / peerings via different ASNs that each announce their individual IPv6 scope.

- For the purpose of DDOS protection the LIR would like to announce a more-specific scope via an appropriate service provider.
- For the purpose of RTT optimization the LIR would like to announce a more-specific.

This list could certainly be still continued and clarified but it already suggests that for affected companies' deaggregation of allocated scope must be classified not only as a wish but as absolutely obligatory. As we explained in the introduction, such companies come increasingly from the column of the enterprise LIRs and less from those of traditional transit LIRs. Just the common use of stateful firewalls at the peering points that only allow bidirectional symmetrical communication will encounter only in a subordinate way in the pillar of the transit LIRs. Furthermore enterprise LIRs usually operate no comparable ISP transit networks between their locations but are rather subject to limited degrees of freedom according to their corporate infrastructures and their interconnections between affected networks. Effects and problems on the consistent announcing of the LIR allocation are therefore inevitable. The deaggregation of the LIR allocation establishes in these cases as a need in order to perceive the business interests of these companies with regard to the new IPv6 world. A strict filter policy undermines the IPv6 planning reliability of these enterprise LIRs by exposing the accessibility of deaggregated IPv6 scopes an uncertainty. This risk is implicitly made aware in document RIPE-532 [12].

Even if companies in the LIR community face no need for a deaggregation today, this requirement could arise spontaneously in future. In a world of increasing globalization and dynamics other companies are daily acquired and integrated in their own corporate infrastructure. It is not excluded that in such situations the need for deaggregation of LIR scopes emerges spontaneously. Then, in order to fulfill the tasks of a LIR and to provide its own customers respectively divisions and profit centers with IPv6 address space, deaggregation might be required. The RIPE NCC formulates a recommendation on strict filtering in the document RIPE-532 [12] but leaves the decision in the last instance up to the LIR-operators.

»It is suggested that prefix filters allow for prudent subdivision of an IPv6 allocation. The operator community will ultimately decide what degree of subdivision is supportable, but the majority of ISPs accept prefixes up to a length of /48 within PA space.

Advertisement of more specific prefixes should not be used unless absolutely necessary and, where sensible, a covering aggregate should also be advertised.«

In a similar manner recurrently formulations can be found in the RIPE policies with regard to aspiring goals like the protection of the global routing table which, however, assign the LIR members themselves the final responsibility and decision making authority for routing and announcing policies. The current policy RIPE-589 [13] is expressed as follows on the subject.

»Wherever possible, address space should be distributed in a hierarchical manner, according to the topology of network infrastructure. This is necessary to permit the aggregation of routing information by ISPs and to limit the expansion of Internet routing tables.

This goal is particularly important in IPv6 addressing, where the size of the total address pool creates significant implications for both internal and external routing.

IPv6 address policies should seek to avoid fragmentation of address ranges.

There is no guarantee that any address allocation or assignment will be globally routable.

However, RIRs must apply procedures that reduce the possibility of fragmented address space which may lead to a loss of routability.«

In contrast, however, we also find formulations in the policies RIPE-555 [14] and RIPE-399 [11], which clearly express the decision-making authority with respect to the announcing policy.

»Routing decisions are the responsibility of network operators.
The RIRs allocate address space to ISPs in blocks, with the expectation that these blocks are announced to the Internet unaltered. It should be noted that the RIRs have no rules about how this address space should be announced to the Internet.«

Summa summarum, a goal conflict arises between the theoretical desire to keep the DFZ routing table as small and compact as possible, as it is also represented by nearly all ISOC organizations (IETF, IAB, etc.), and the practical requirement profiles and laws of the free market. How well theoretical approaches work in the DFZ, can be guessed for example by the lofty goal of a hierarchical routing between RIR domains. A look inside the global routing table indicates not everywhere only aggregated address blocks from other RIR domains. Rather, the free market seems to write its own laws and the LIR community must ultimately find a solution itself. So especially the representatives of both LIR pillars are in the responsibility to harmonize their interests and work out compromises in solidarity. Both the community of interest of the transit LIRs as well as that of the enterprise LIRs is prompted to make a reasonable announcing policy, on the one hand to protect the global routing table, on the other hand to imply a certain degree of planning reliability for companies that depend on a moderate deaggregation of their LIR allocation without an covering aggregate. Otherwise an uncertainty in terms of accessibility of deaggregated IPv6 scope would arise specifically in the area of enterprise LIRs which counteracts the conveying of IPv6. From a holistic point of view this may not be the objective of the LIR community because finally the enterprise LIRs turn out again as customers of transit LIRs. Or as Randy Bush has expressed in the v6ops mailing list:

»Maybe if we start filtering now. But we know how well that went in IPv4 when their suits called our suits and said "we pay you to let us contact <deaggregator>"«

The irony of fate provides a system with strong feedback at this point.

Hence, the evaluation of the RIS data reveals a spontaneously occurring process which increasingly generates more-specifics in the LIR scope. This trend towards more equal distribution and deaggregation in the scope of LIR allocations is probably due to the different requirement profiles and business environments within the LIR community, and should be also represented graphically as a quotient. Figure 8 shows the ratio of the sum of all the more-specifics to the sum of all LIR allocations.

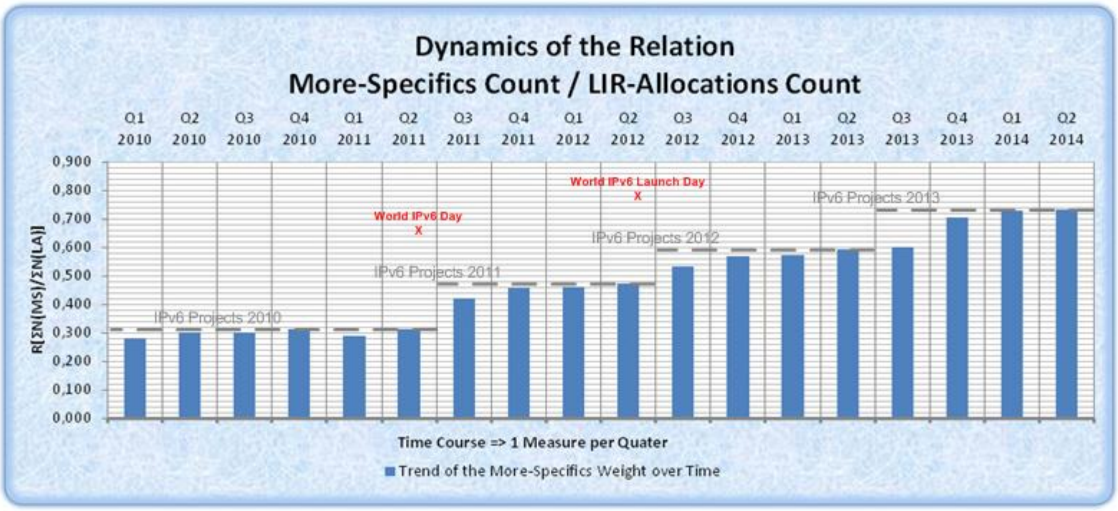

Figure 8: Dynamics of the Relation More-Specific Count / LIR-Allocations Count

Interestingly, the ratio seems to be developing in year levels. Whereas the first two years hardly have dynamics, the increase of the more-specifics in the LIR scope starts shortly after the World IPv6 Day in June 2011. The next level is achieved shortly after the IPv6 World Launch Day in June 2012, followed by a further step about one year later.

Apparently the World IPv6 Day has brought this process in motion, continuing significantly in subsequent years. As can be seen in Figure 8, also in 2012 and 2013 a sudden increase of the more-specifics towards the middle of the year is occurring, which is probably attributable to the implementation of IPv6 projects.

4.2 Phase Portraits of the Prefix Distribution within the LIR Scope

Figure 5 - Figure 7 illustrate the dynamics of the prefix distribution in annual steps. If you want to increase the temporal resolution in the next step, and present the data in quarterly increments this form of representation is no longer suitable. To this end, we rely on so-called phase portraits which were already described in the experimental part. As an example and to introduce the new form of presentation, we use the results from Figure 8 and generate a phase portrait.

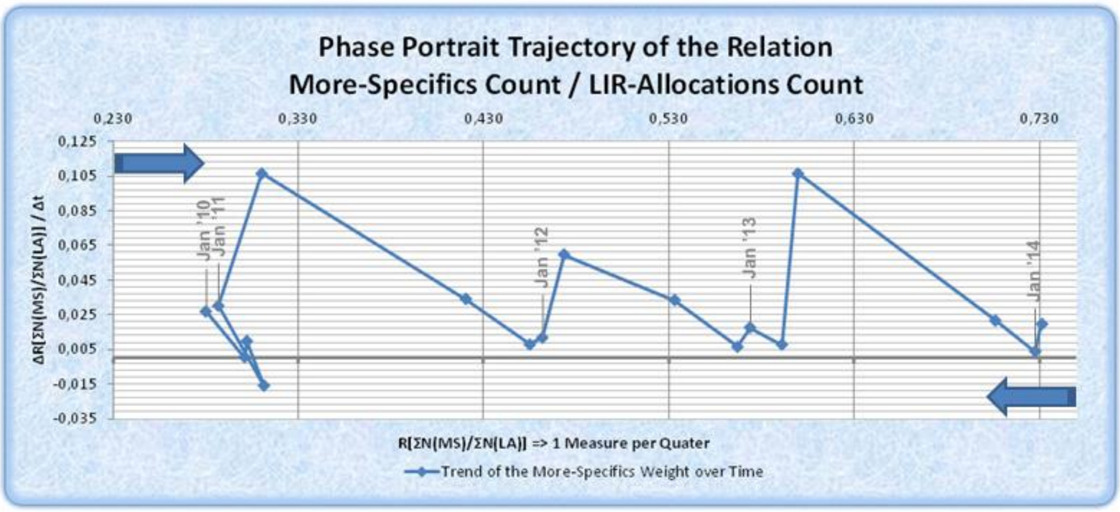

Figure 9: Phase Portrait Trajectory of the Relation More-Specific Count / LIR-Allocations Count

The abscissa of the graph represents the quotient of the sum of all more-specifics over the sum of all LIR allocations; the ordinate indicates the change of this value over time. Thus, we no longer apply the observable directly against time, but rather their time variation against the observable itself. The advantage of this form of representation is the fact that time is an implicit component of the variable itself. For better orientation, the measurement points in January of each year were selected.

The graph shows a total of 18 measurement points and represents the trajectory of the observable through the DFZ phase space of the LIR scopes (space of all possible states). As already explained in the experimental part, any point above the zero line (abscissa) indicates a positive speed of the observable, i.e. the observable moves from left to right which is symbolized by the blue arrows. The results indicate that apart from an initial phase in 2010 the trajectory of the observable moved consistently and steadily towards the more-specifics. Furthermore, it can be seen that the time course as already discussed in Figure 8 takes place in three stages in which first the speed increases (ordinate) then to complete a larger jump to the right (abscissa) and finally loses speed again. However, the results also show that in the last four years, not a single phase with a negative speed, i.e. a regression in favor of LIR allocations (aggregates), has occurred.

If we now transmit the results of Figure 6 in a phase portrait and increase the temporal resolution of one year on a quarter again, we can take a look on the fine structure of the complete prefix distribution in the DFZ phase space. Figure 10 presents the dynamics of all prefixes between /29 and /48, each based on 18 measurement points. The trajectories give an impression of the time course of the weight of a prefix in relation to the total number of prefixes.

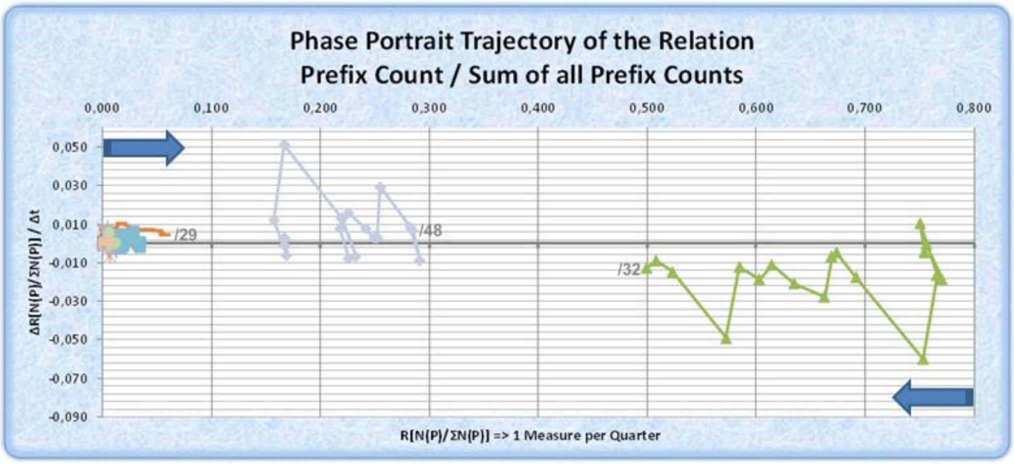

Figure 10: Phase Portrait Trajectory of the Relation Prefix Count / Sum of all Prefix Counts

The trajectories do not show above all a linear behavior with time, they do not even show a simple non-linear behavior but rather a general trend with a chaotic fine structure. Ultimately, this reflects the heterogeneity of the LIR community whose DFZ policies can apparently develop the prefix distribution in a certain direction. The combination of both LIR columns the transit and the enterprise LIRs, their individual requirement profiles, their business environments and resulting announcing policies finally lead to a chaotic fine structure, but revealed a long-term trend.

After an initial stable phase in 2010 the /32 trajectory indicates a steady decline in their weight, while at the same time the /29 trajectory reveals an increase. Although the rate of this increase drops slowly down over time, it can still illustrate the trend to a /29 allocation. The /48 trajectory indicates mostly positive rates and thus long-term development from left to right for more weightage. The fine structure reveals, however, that this trend is broken repeatedly by shorter phases in which data points fall in the negative range of the speed and thus signal a short-term regression of weight. These phases are apparently attributable to filtering or aggregation phases in which a larger number of /48 prefixes were removed from the DFZ. The remaining trajectories are currently still at lower weights and require a more detailed analysis.

Next we consider the results of Figure 7 as a phase portrait and thus the weighting of the more-specifics in relation to the total number of allocations. Figure 11 shows a zoomed section of the prefixes between /33 and /47. The trajectory of the /48 prefix is similar to that in Figure 10. A representation was waived.

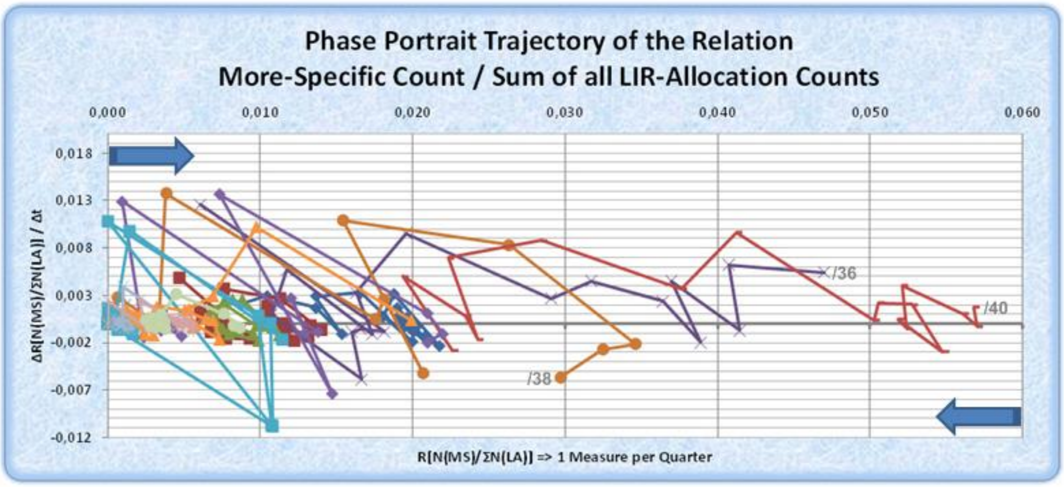

Figure 11: Phase Portrait Trajectory of the Relation More-Specific Count / Sum of all LIR-Allocation Counts

The figure shows that most prefixes develop at approximately the same growth rate as the LIR allocations in the DFZ, i.e. the values jump chaotically in the range of the zero-line (abscissa). In addition, however, we also find three exceptions which we want to dwell. At first the /38 trajectory seems to develop with an increasing growth rate above the abscissa but again drifts in the negative area below the abscissa in the last three measurement points. In this phase of the trajectory the growth rate in relation to that of the LIR allocations reveals to regress again and the prefix loses presence. In contrast, according to the general trend the prefixes /36 and /40 appear to be more dominant over time. On average the growth rates are higher than those of LIR-allocations. The trajectories reveal a highly nonlinear and chaotic character respectively on their path through the DFZ-phase space. In particular, the /40 trajectory reveals even a small loop in the last eight measuring points. From chaos theory it is known that such phenomena may be caused by introducing nonlinearities.

All in all at least two more-specifics in addition to the /48 prefix appear to develop more dominant than the rest - the prefixes /36 and /40. However, accurately these more-specifics reflect the deaggregation in the context of 4-bit nibbles (/32 - /36 - /40) and could represent the basis for a moderate and responsible announcing policy. Assuming the LIR community would in the long run remember it to build the deaggregation of LIR scopes on absolutely obligate requirement profiles and not to operate the deaggregation beyond a prefix length of /40 (> /40) then it could create a breeding ground for harmonizing divergent interests. An agreement on this basis would both protect the global routing table by limiting the deaggregation potential of a /32 LIR scope to a maximum of 2^8 = 256 and as well give the LIR the freedom to perform a moderate deaggregation in compelling cases without running the risk of falling into a strict filter policy. An agreement including a collective commitment can only be achieved in solidarity of all LIR operators who register an interest to overcome this existing uncertainty. In view of the higher goal of promoting IPv6 across all boundaries and restrictions to come, it must be formulated at least as a rational desire in the LIR community.

Continuous monitoring of the DFZ phase space of all prefixes affected could be updated daily by an automated evaluation of the RIS project bview files. This would improve the temporal resolution of the statistics on one day and nevertheless obtain a deep insight into the actual processes within the DFZ phase space. An online monitoring of this kind would even be equivalent to a simple indication system that could establish correlations between RIPE policies and the development of the prefix distribution in a timely manner. The same applies to all political decisions in the IPv6 context respectively current IPv6 events and happenings.

4.3 Strict Filtering Reloaded?

But what about the "Yes, but ..." fraction, which could represent the conviction that this discussion is unnecessary provided that any more-specific would have a covering-aggregate and therefore the accessibility of deaggregated scope would be guaranteed. Thus, the resulting question can be formulated as follows:

How many more-specifics in the LIR scope do not have a covering-aggregate?

There should be agreement that isolated more-specifics without any covering-aggregate would be directly affected by strict filtering policies, since they operate without a backup prefix in the DFZ. Whether these isolated scopes can be subsequently provided with a covering-aggregate will heavily depend on the individual business environments of the LIRs. As already discussed, especially in the column of the enterprise LIRs there are numerous restrictions that just do not allow the announcements of an aggregate. Figure 12 compares both more-specific variants.

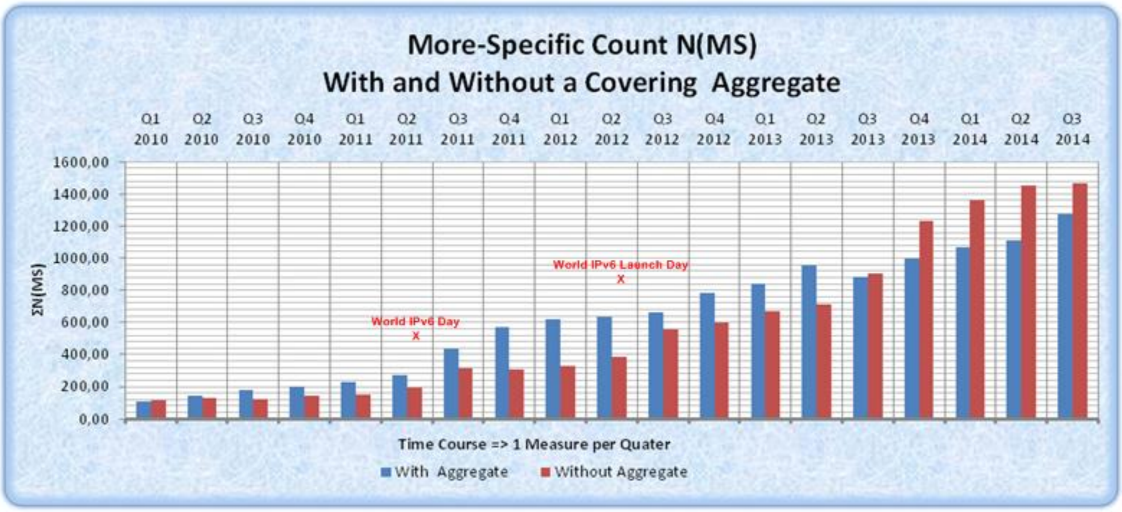

Figure 12: More-Specific Count N(MS) With and Without a Covering Aggregate

The chart illustrates that over the past five years, a significant proportion of isolated more-specifics was constantly recorded. Starting with about 50% in 2010 the proportion developed initially below 50% until it finally rose again. Until the third quarter of 2013 the proportion of the more-specifics without covering-aggregates increases even well above 50%. This development will not defuse the issue but rather force the debate over strict filtering. Any LIR which carries with the idea to use strict filtering is aware of the risk to filter out larger scopes of isolated more-specifics from its global routing table and thus deprive its customers. A strict filtering policy should therefore be critically reviewed also from a customer perspective respectively from the perspective of an end user (RIPE jargon) and, where appropriate, be included in the appointment decision of Internet services.

Next the question arises: Are there LIRs to the current state of the discussion at all that employ strict filter? How can you imagine their global routing table? In Table 3, the anonymized Internet routing table of an affected LIR was shown.

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	22	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

Table 4: Global Routing Table of a Strict Filtering LIR in last 5 Years

Table 4 illustrates that more-specifics between /33 and /48 largely fall victim to the filter policy. Only a few prefixes are managed in this range up to Q2 2014. How the Internet routing table should look like without strict filter can be observed in the last line. Apparently, the affected LIR has abandoned its strict filter policy since Q3 2014 and now maintains a complete DFZ.

To conclude the debate there still remains the issue for further clues in the RIS data that underpin an obligate deaggregation policy. Are there any LIR infrastructures that make a deaggregation required and no longer allow the announcing of a covering-aggregate meaningful? While bview files are giving no detailed insight into individual LIR infrastructure environments or even LIR policies, however, using BGP path information the origin ASN can be evaluated and compared with the LIR-allocations. LIRs that propagate their scope via at least two ASN cannot do this without deaggregation of their allocation, because you can assign an ASN one covering aggregate only. In combination with insulated networks, too little bandwidth between the networks or the use of stateful firewalls at the peering points will become a show-stopper for the covering aggregate. Figure 13 shows the development of such LIR environments with at least two independent ASNs over the past 5 years.

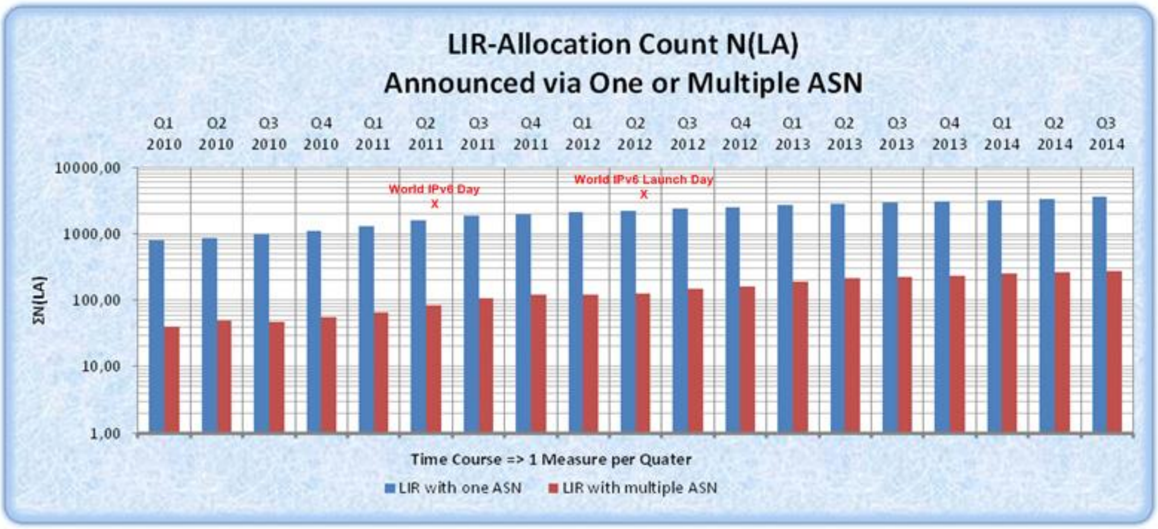

Figure 13: LIR-Allocation Count N(LA) Announced via One or Multiple ASN

The chart shows that a quite perceivable proportion of the LIR community propagates sub-allocations via at least two ASNs. This share increases slightly in the past three years and achieved in Q3 2014 at least 7% of all allocations. In other words, at least 7% of all LIRs must already deaggregate their allocation via multiple ASNs. That is a fact.

4.4 Via Rationis – The Rational Path

The present evaluations have shown that there should be a legitimate interest within the LIR community to discuss the development of LIR scopes openly. Where should lead the way and how can we control it responsibly. As already indicated, at least two columns are establishing in the field of the LIR community for quite some time - the traditional transit LIRs and the enterprise LIRs. The history of RIPE policies indicates that organizations are actually encouraged in the context of IPv6 to conclude a LIR contract with the RIPE NCC. The last RIPE Meeting (RIPE-68) presented figures of new LIR contracts and the subsequent debate concerning the reduction of the LIR fees speak their own language here. On the other hand, almost no one would welcome a return to the antiquated conceptions of NAT in the IPv6 community. NAT stands figuratively for "No Address Transparency" and thus counteracts the IPv6 philosophy of transparent end-to-end transmission and implementing the principle of simplicity [1], respectively. The path of the IPv6 protocol should lead predominantly and long term via public address space for all participating units ... up to the client device. A more sophisticated discussion about IPv6, however, would fill another publication.

Consequently, if a profit organization applies for an IPv6 PI scope, exactly one /48 prefix according to the current policy RIPE-589 [13] is assigned per Internet breakout. As of today a larger block for more IPv6 future security and freedom in planning of the end user infrastructure is not provided. More /48 prefixes are obtained only under detailed justification of current demand. In contrast, however, each end user is open to the conclusion of a LIR contract, where initial hurdles such as the presentation of redundant Internet breakouts via at least two transit LIRs have disappeared from the RIPE policies over time. All this not only promotes the IPv6 path but also the LIR path of non-transit organizations. Here the RIPE NCC pursues its own strategy which is quite different from those of other RIR domains such as the ARIN for example. Ultimately, divergent interests could arise in the LIR community and infrastructures which are characterized by differing features as well as production environments that are subject to entirely different restrictions. Task will be to accept this in the first instance and then to work out a reasonable path for the future.

Based on the differentiation of both LIR pillars you can also characterize their underlying network infrastructure. Accordingly, we consider on the one hand the globally networked Internet infrastructure of the carrier and transit LIRs and on the other hand the infrastructure environments of different organizations and corporations with LIR contract whose main business activity is beyond IP transit services - the enterprise LIRs. While the global Internet transit backbone is characterized by rather moderate growth and manageable dynamics the ongoing development of the enterprise LIR environments is characterized by flexible sometimes chaotic growth and ever increasing dynamics of their locations. Not least paid to the trend of globalization, the fraction of the enterprise LIRs thus exerts a noticeable influence on the quantitative development of the DFZ. The group of the enterprise LIRs brings especially more dynamics and degrees of freedom into the infrastructure environments of the LIR community.

One effect of these newly gained degrees of freedom could be already reflected in the chaotic fine structure of the more-specifics trajectories. Another consequence of these new developments will manifest itself in a progressive trend that inevitably leads to a more equal distribution of prefixes over time. This trend is fueled by further liberalization of degrees of freedom in an open system with only a few constraints and always strives against the most probable state. The law of nature which is behind this general trend is the much cited second law of thermodynamics. The affected open system corresponds to the DFZ phase space in the LIR scope and the most likely state is an increasing equal distribution of more-specifics within this scope if one does not counteract the process. And right here could set a responsible announcing policy that corresponds figuratively to a kind of conditioning or the establishing of certain boundary conditions.

Of course new boundary conditions could be implemented successfully in the solidarity of all LIRs only, i.e. both the representatives of the transit LIRs as well as those of the enterprise LIRs must be convinced of the need and the advantages of a common announcing policy and support it openly. A clear commitment of both communities of interest in favor of a moderate deaggregation up to a prefix length of /40 could trigger the necessary impetus. Harmonization on the basis of a compromise could also pave the way in the DFZ phase space to establish two attractors that would win more dominance in the half-byte Nibbles /36 and /40 over time. The desire and the prospect of a self-healing phenomenon in the LIR scope will only strengthen a law of nature and thus the increase of equal distribution over the entire LIR scope in favor of high-bit prefixes like /48. Exactly this battle was already lost in the IPv4 DFZ which is reflected by the persistent and uncontrollable reality of deaggregation in IPv4 space. This process will continue to move forward, because it is driven by an open system with minimal feedback and boundary conditions.

5 CONCLUSIONS

The analysis of the IPv6 RIS data could point out that a general trend in the LIR scope is establishing towards more deaggregation. The growth rates of the more-specifics seem to develop sometimes above average in relation to the LIR allocations itself. This trend is particularly noticeable in the last two years and continues to progress. Within the more-specific scope foremost the /48 prefix reached dominance, thus generating a maximum potential of equal distribution similar to the /24 prefix in the IPv4 address space or, more generally, high-bit prefixes of any address space. Exactly these high-bit prefixes sooner or later encourage a maximum degree of equal distribution and thus satisfy the underlying law of nature in an open system with very limited boundary conditions. The system evolves in the direction of the most probable state. It may be assumed that this trend is increasingly fueled by the pillar of enterprise LIRs, commonly faced with greater restrictions on their own infrastructure environments. In particular, the use of stateful firewalls that allow only bidirectional symmetric communication via peering point not only promotes the deaggregation of LIR allocations but makes it even an obligate requirement profile. In this situation affected organizations sometimes have no way to propagate a covering aggregate into the DFZ and, as a result, are defenselessly exposed to a strict filter policy. With view to IPv6 migration projects the resulting fuzziness of planning counteracts the promotion of IPv6 and leads at best to more uncertainty as well as open questions. It is exactly this uncertainty that we must tackle and eliminate within the LIR community.

If one doesn't accept this topic, hazard potentials for legitimate business interests of enterprise LIRs may arise. In this sense, any business model that builds on globally functional routing policies regarding to the own deaggregated LIR scope is possibly undermined by strict filtering. Therefore, only a common approach within the LIR community leads to success. It should be an objective on the one hand to further push IPv6 as well as to protect the global routing table but also on the other hand to provide all LIR organizations that rely on a moderate deaggregation of their own allocation, a legitimate and protected by the LIR community phase space. Ultimately, the LIR operators who have admirably placed their own role in the document RIPE-399 [11] bear the responsibility for the envisaged harmonization of announcing policies.

```
»The industry considers it improper for the RIRs to tell ISPs how to announce address
space; in the same way that libraries won't tell their readers how to read the books it
lends.«
```

That is right, however, in an appropriate context the books should be treated as such that no part of chapter is ripped out after reading and hence is no longer available for the rest of the reader community. To meet this responsibility is certainly an arduous path in the community of LIR operators but you should always keep the cash flow in mind. Because in the end enterprise LIRs are the customers of the transit LIRs and as mankind at least since Isaac Newton know: Each actio usually results in a reactio. This is very wisely not to be understood as a threat but due to the experience that complex systems increasingly use their feedback loops in such situations. Durable solutions are usually achieved through harmonization and compromise that means achieving a kind of equilibrium in their feedback loops.

As a possible solution a reduction of the deaggregation potential to 8 bits and thus a maximum of 256 prefixes per /32 LIR allocation was proposed. This newly defined phase space for more-specifics in the LIR scope would include all prefix lengths between /32 to a maximum of /40. The resulting announcing policies would have to be accepted by the community of LIR operators and implemented long term. For more-specifics greater than /40 there would be no actual guarantee on unfiltered scopes. Finally, the proposed compromise would also eliminate the currently existing uncertainty in IPv6 address planning of affected LIR organizations. Coupled with an online monitoring of the DFZ phase space we would have an instrument which is able to recognize developments in the DFZ at an early stage and consequently offers the opportunity to adapt IPv6 announcing policies responsibly. However, for this purpose some rethinking with the focus on changing LIR infrastructure environments is needed, which demands certain openness to change and a holistic basic adjustment. Or to express it in the words of a scientist and Nobel laureate (translated by the author of this publication):

»We must not in favor of one organ - the rational analysis - stunt all the others; it is much more important to catch the reality with all organs that are given to us …«

Werner Heisenberg, ›steps across borders‹

6 APPENDIX

6.1 References

[1] *"Some Internet Architectural Guidelines and Philosophy", Randy Bush, David Meyer, RFC 3439*, December 2002,
https://www.ietf.org/rfc/rfc3439.txt

[2] *"Report from the IAB Workshop on Routing and Addressing", David Meyer, Lixia Zhang and Kevin Fall, RFC 4984, September 2007*
https://www.ietf.org/rfc/rfc4984.txt

[3] *"IPv6 Global Unicast Address Assignments", IANA*
http://www.iana.org/assignments/ipv6-unicast-address-assignments/ipv6-unicast-address-assignments.txt

[4] *"RIR Prefix Size Distribution", RIPE NCC*
https://stat.ripe.net/widget/rir-prefix-size-distribution

[5] *"IPv6 BGP Filter Recommendations", Gert Döring, SpaceNet*
http://www.space.net/~gert/RIPE/ipv6-filters.html

[6] *"RIPE Atlas – A Case Study of IPv6 /48 Filtering", RIPE NCC*
https://labs.ripe.net/Members/emileaben/ripe-atlas-a-case-study-of-ipv6-48-filtering

[7] *"LIRs With and Without IPv6", RIPE NCC*
https://labs.ripe.net/statistics/lirs-with-and-without-ipv6

[8] *"Total IPv6 Allocations", RIPE NCC*
https://labs.ripe.net/statistics/total-ipv6-allocations

[9] *"Phase Space", Wikipedia*
http://en.wikipedia.org/wiki/Phase_space

[10] *"Phase Portrait" Wikipedia*
http://en.wikipedia.org/wiki/Phase_portrait

[11] *"RIPE Routing Working Group – Recommendations on Route Aggregation", Philip Smith, Rob Evans and Mike Hughes, RIPE-399, December 2006*
http://www.ripe.net/ripe/docs/ripe-399

[12] *"RIPE Routing Working Group – Recommendations on IPv6 Route Aggregation", Rob Evans and Philip Smith, RIPE-532, November 2011*
http://www.ripe.net/ripe/docs/ripe-532

[13] *IPv6 Address Allocation and Assignment Policy", Address Policy Working Group, RIPE-589, May 2013*
http://www.ripe.net/ripe/docs/ripe-589

[14] *"Address Space Managed by the RIPE NCC", RIPE NCC, July 2012*
http://www.ripe.net/ripe/docs/ripe-555

6.2 Disclaimer

All products, company names, brand names, trademarks and logos are the property of their respective owners.