

ASN Missing In Action

A Comparison of RIR Statistics and RIS Reality

René Wilhelm, Henk Uijterwaal
RIPE NCC

Date: October 13, 2005

Document: RIPE-353.ps

Abstract

In this paper, we look at Autonomous System Number (ASN) assignments by the Regional Internet Registries (RIRs) and the number of unique ASNs seen by routers on the Internet.

On 1 August 2005, 33000 ASNs were assigned by the RIRs and the RIRs currently assign 284 ± 43 ASN/month. If this were to continue, the pool of available ASNs could run out somewhere between 2013 and 2016.

At the same time, the routing table is only growing by 183 ± 44 ASN/month. There seem to be two reasons for this disparity:

1. ASNs are being requested based on future expansion plans, but are never actually used;
2. ASNs are being retained that are no longer in use.

If we could recover all of these unused ASNs, then the available pool could last until somewhere between 2025 and 2030. If we cannot recover them, then Internet Service Providers (ISPs) should start thinking about deploying 4-byte ASNs.

We make recommendations for ASN policy changes, that we feel would allow the RIRs to better allocate ASNs to End Users.

Contents

1	Introduction	2
2	Theoretical Background	3
2.1	Date format	3
2.2	Definition of an AS	3
2.3	AS Numbers	4
2.4	4 byte AS Numbers	4
2.5	When does one need an AS?	4
3	The RIR and LIR System	4
4	The Routing Information Service (RIS)	5

5	Data Sources and Selection.	5
5.1	RIS collected routing data	5
5.1.1	Pitfalls	6
5.2	RIR allocation statistics	6
5.2.1	Problems with Stats Files	8
5.3	CIDR Report	9
5.4	Availability of the Data Sets	9
6	ASN Missing In Action.	9
7	ASN Missing in Statistics	10
8	The total number of ASN seen.	11
8.1	The total number	11
8.2	Per registry	11
8.3	Growth rates of RIS and RIR ASN counts	12
8.4	Growth rates per registry	14
9	Fraction of ASNs seen	16
10	ASNs no longer seen	16
11	Age of ASNs in RIS on 31 March 2005	19
12	Activation delay of ASNs	20
12.1	Yearly data	20
12.2	By registry	22
12.3	Discussion	22
13	When will the Internet run out of 16 bit ASNs?	22
14	ASN Assignment Policy Changes	23
15	Conclusions	23

1 Introduction

The global Internet can be seen as a network of networks. Individual organisations first build a local network connecting the machines that they operate. These individual networks then are connected to each other in order to create the global Internet.

For machines to send data across this “network of networks”, they need a “map” of the Internet and the capability to decide which path data has to take. This process is called routing and the computers that perform this task are called routers. How data is routed between networks, depends on policies set by the operators of the networks. These policies will vary from one network to another. A network with its own distinct routing policies is called an Autonomous System (AS). A formal definition of an AS can be found in [1].

The design of the routing system is beyond the scope of this paper. However, it should be clear that for a routing system to work, each AS (as well as each individual computer) should have a unique identifier. In the case of an AS, this is known as its AS Number (ASN). To guarantee uniqueness, the RIRs allocate ASNs to networks in a hierarchical way. This is discussed further in section 3. The RIRs publish statistics showing which ASNs have been allocated.

Each AS should be aware how to reach other ASs. If not, then parts of the Internet will not be visible. This information is stored by the routers in the Routing Information Base (RIB). RIBs are maintained dynamically. Several projects exist that collect the RIBs for debugging and research purposes, one such project is described in section 4. By analysing the RIBs, you can see which ASNs are in active use on the Internet.

By early 2003 [2], there were around 14000 unique ASNs in the RIBs. This total was growing by around 200 each month. At the same time, the RIRs allocated about 300 new ASN each month. This meant that around 100 ASNs were allocated but not used, or appeared to be “Missing In Action”. In this paper we analyse what happened to these ASNs.

The outline of the remainder of this paper is as follows. In section 5, we discuss the data sources used for this analysis. In section 6, we give an exact definition of “Missing In Action”. Our analysis will show that there are not only ASNs Missing In Action, but also ASNs in use that are not registered. This is discussed in section 7.

The total number of ASNs seen is discussed in section 8. ASNs appear to be missing for two reasons, this is discussed in section 11. As we will show in section 2, only 65000 unique numbers are available and half of those numbers have already been used. In section 13, we will discuss when the Internet might run out of ASNs and what can be done to avoid this.

This paper is a write-up of all our results, in order to document our work and generate further discussion. Some background information has been added to improve readability. Parts of this paper have been presented at various meetings [2,3].

2 Theoretical Background

2.1 Date format

Dates are written in standard European order, rather than the US format, either in text (1 September 2005) or as DD/MM/YYYY with leading zeros dropped.

2.2 Definition of an AS

There are many definitions of an AS. The best one can be found in the definition of the Border Gateway Protocol (BGP), the protocol used to distribute routing information between ASs [1]. A more informal definition would be that an AS is a network (usually) maintained by a single organisation, with its own distinct routing policies.

From these definitions, it should be clear that a site only needs an AS when its network has the need to specify its own set of routing policies, these policies will be different from the policies of its upstream provider. This is (usually) the case for networks that have two or more upstream providers.

2.3 AS Numbers

Each AS has a unique label, its ASN. The current version of the BGP specifications [1] assigns 16 bits for this number. This gives 65536 possible values.

The first and last ASNs (0 and 65535) have been reserved and cannot be used to identify networks. The upper 1023 ASNs (64512 through 65534) have been reserved for private use. Anybody can use these without having to register them. These ASNs should never appear on the global Internet. The most common case for the use of a private ASN is where you want to set up sub-domains inside a domain.

The remaining numbers (1 through 64511) are available for Internet routing. The ASNs are assigned sequentially, without any structure in the numbers. Regional and Local Internet Registries assign these to End Users. They guarantee uniqueness and maintain a register of all ASNs in use. Details on this process can be found in section 3.

2.4 4 byte AS Numbers

While 64510 ASNs seems a large number, if the Internet continues to grow at its current rate, there will come a point where all ASNs have been assigned. The industry has started to discuss what to do next.

One solution would be to increase the number of bits used to express an ASN. This is not as simple as it sounds. ASNs are embedded in BGP protocol messages and the routers expect messages formatted according to this protocol. There is a proposal [4] to change this and use 4-byte ASNs. It has received limited attention and to date only one vendor has implemented it. We are unaware of any deployment plans.

2.5 When does one need an AS?

Not every domain needs to have a unique ASN. An AS is a domain with its own routing policies and not every domain does want to do that. A very common case is a sub-allocation, where a site splits off some of its IP address space to customers. The routing policies for the sub-allocation will still be the same as for the larger block and no AS is needed. Even if a separate AS is preferred for practical reasons, this could be done with a private ASN.

3 The RIR and LIR System

ASNs have to be globally unique. In order to ensure uniqueness, they are distributed in a hierarchical fashion: a global organisation (IANA [5]) administrates the pool of all available ASNs and hands out blocks to the five RIRs. They distribute them further to End Users, usually through a Local Internet Registry (LIR). Both IANA and the RIRs maintain registers of the ASNs that they have assigned.

In some countries in the Asian-Pacific region, there is another layer between the RIR and the End Users. A National Internet Registry (NIR) deals with resources for a particular country. It applies for a block from APNIC, then reassigns the ASNs in this block further inside that country. This has some effects on the statistics discussed further on in this paper.

Handing out ASNs from the RIRs to the LIRs is based on policies set by the RIRs [6–10]. These policies differ slightly between regions but are all based on the guidelines described in [11]. The guidelines in this RFC are based on “demonstrated need”. A network must show

that it needs an ASN before one is assigned. This is only checked *before* the ASN is assigned. None of the RIRs check that the plans used to justify the assignment actually become reality.

A common factor across all RIR policies is that you have to show that a network that needs an unique global identifier will be deployed in near future, of the order of 1 to 3 months.

All RIR policies insist that the network should tell the RIR when it no longer needs the AS. The RIR can then reassign the AS. However, there are no (financial) incentives to return an unused ASN, and a large number of unused ASNs are never returned.

While the theoretical model for ASN assignment is based on temporary assignment to meet a demonstrated need, in practice the model is closer to allocation for once and forever.

At the moment, there are 5 RIRs: RIPE NCC, ARIN, APNIC, LACNIC and AfriNIC. LACNIC was only active from 2002 onwards. AfriNIC started its operations in February 2005. As this study focuses on data for the period 1 January 2001 until 1 August 2005, AfriNIC is not included as a separate category when statistics are split out over RIRs.

4 The Routing Information Service (RIS)

The Routing Information Service (RIS) [12] started as a RIPE NCC project in July 1999. It collects and stores Border Gateway Protocol (BGP) routing information and makes it publicly available.

To collect this information, 14 route collectors have been installed world-wide and peering sessions have been set up with approximately 450 ASs, both IPv4 and IPv6.

The information that is collected consists of dumps of the routing tables taken three times a day and all BGP updates sent by each peer.

The data is stored in a SQL database for three months. This “RIS” database and the tools built on top of this can help network operators troubleshoot routing problems by providing eBGP data collected over time without being limited to a single BGP view. All data is also archived and all raw RIS data since September 1999 is available for research purposes.

5 Data Sources and Selection.

5.1 RIS collected routing data

Our study started with the analysis of the RIS data. We analysed the data taken between 18 August 2000 and 1 August 2005. Earlier data, from the first 11 months of RIS operation, were skipped for technical reasons. Since the majority of ASNs have been assigned after August 2000, we will see this subset of the RIS data is rich enough for the purposes of this study.

For each day in the data set, we processed all RIS data files, both the full table dumps and the updates. For each recorded RIB entry or BGP update message, we broke down the AS path attribute into its components. This gave us a list of ASNs that were seen at least once, as origin AS, as transit AS or both. We stored the results in a data base, giving each AS the following fields:

- `First_seen`: date the ASN was first seen in the RIS
- `Last_seen`: date the ASN was last seen in the RIS

- `Days_seen`: total number of days the ASN was seen in the RIS

We also stored the total number of ASNs seen on each day in a separate table. On 31 July 2005 we thus found 20422 ASNs in the routing system. Note that the count excludes private ASNs which might have popped up in RIS data. These are not included in our analysis, as they are not assigned by RIRs.

5.1.1 Pitfalls

While in general the routing tables reflect reality, there are two potential issues when comparing RIS data with RIR assignments:

1. It is always possible for AS numbers to show up in RIS data without having been assigned by a registry. Usually, this happens by mistake and the corresponding BGP announcements do not last very long. To enable a better comparison to RIR assignment statistics, we filtered the RIS data based on the number of days the ASN was seen in the RIS. Any ASN seen for less than a total of 7 days will be discarded from detailed analysis. From all 26757 ASNs detected by RIS in the period August 2000 to August 2005, only 449 or 1.7% match this rule and were excluded.
2. The absence of an ASN in RIS data does not necessarily mean the ASN is not used at all. In some cases, assigned AS numbers might not be visible in the global routing tables. This includes those used in private interconnects and confederations using public ASNs. In this study we do not judge what is good or bad usage of an ASN, we only focus on what is seen in the RIS.

It is obviously difficult to estimate the number of ASNs used in these situations, however, we believe that it is small.

5.2 RIR allocation statistics

Since 2002, the RIRs have published “RIR Stats Files” containing details of all assigned resources (IP addresses and AS numbers) [13]. For ASNs, each record of in the Stats files lists:

- the ASN (or block of ASNs)
- the assignment date
- the assigning registry
- the country code of the recipient.

In theory, it should be possible to obtain a historic count of assigned ASNs as a function of time from the latest Stats files: start from the total number of assigned ASNs and count back day by day, using the dates of all ASN assignments. However, since Stats Files only provide the status of a given day and do not contain information on returned or recycled resources, we have to look at each individual Stats File from past years to obtain a more accurate view of the number of ASNs assigned at those dates. For this analysis, we processed all Stats Files published until 1 August 2005. For each ASN appearing in the Stats Files, we store the following information in a database:




Figure 1: RIRs ASN Count, for the individual statistics files (+, solid line) and a count-back from 1 September 2005 (x, dashed line).




Figure 2: Absolute time in days between first and last assignment. 6913 ASNs have their assignment date changed, 67% have exactly 1 month difference, 23% exceeds 1 year.

- `first_assign` – assignment date of the ASN as listed in the first (oldest) Stats File containing the ASN
- `last_assign` – assignment date of the ASN as listed in the last (latest) Stats File containing the ASN
- `first_stats` – date of the first (oldest) Stats File containing the ASN
- `last_stats` – date of the last (latest) Stats File containing the ASN
- `first_rir` – registry first seen listing the ASN in a Stats File
- `last_rir` – registry last seen listing the ASN in a Stats File
- `first_country` – country of origin of ASN holder as in the first (oldest) Stats File
- `last_country` – country of origin of ASN holder as in the last (latest) Stats File

In addition we keep track of a total per registry count of ASNs for each date.

On 1 August 2005, the total RIR ASN count according to Stats Files was 32805, whereas the accumulated statistics from the past 3 years show an additional 879 ASNs had once been assigned, but no longer appear in the latest Stats Files. Figure 1 shows how the countback from 1 August 2005 starts to diverge from the “per Stats File ASN count” as we go further back in time. Two effects contribute to this divergence:

1. ASNs that had been assigned but were returned at a later date are missing from the projected ASN count from the day they were first assigned, until the day they were withdrawn from the Stats Files.
2. ASNs that had been assigned in the past, were returned and later reissued appear in the latest Stats Files with a later date than that of first assignment. These ASNs are missing from the projected ASN count for the days between first assignment and reassignment.

5.2.1 Problems with Stats Files

Though processing past Stats Files provides more information, it also brings problems to the analysis.

- *Duplicates*: In 2002 and early 2003, we found overlaps of ≈ 500 ASNs in the assignment statistics, with some ASNs appearing in the data from more than 1 RIR. When we counted the total number of ASNs assigned on each date by the combined RIRs, we removed these duplicates.
- *Last assign before first assign*. 104 ASNs appear with an earlier date in later Stats Files.
- *No assignment date*. 193 ASNs always had an assignment date of "0000-00-00", these are all in ARIN Stats Files.
- *First and last assignment date not the same*. As shown in Figure 2, 6913 ASNs have different first and last assignment dates:

- 191 were assigned after 1 January 2002
- 4709 have first and last assignment dates that differ by no more than 1 calendar month. (Only 20 of these were assigned after 1 January 2002).
- 3351 of the 6913 were still seen by RIS on 1 August 2005.

We have assumed that the ASNs with assignment date differences of less than 33 days have not been reassigned by the RIRs. The date changes appear to be administrative corrections.

5.3 CIDR Report

The CIDR report [14] is an analysis of the routing table as seen from AS4637. It is generated once a week on Friday afternoon. The first CIDR report was published in 1994, making it one of the longest running analysis of the Internet routing system. One of the items reported is the total number of ASN seen on the Internet. This number was extracted from all reports and stored in a database.

5.4 Availability of the Data Sets

All data sets discussed in this section are available on request, in order to verify our studies or for further analysis.

Results

6 ASN Missing In Action.

We now have two sets of data:

- The combined data from the RIS and CIDR report showing which ASNs are seen on the Internet, and
- The data from the RIR statistics showing which ASNs have been assigned.

If we compare first data set against the second one, there are three possibilities:

1. An ASN is seen in both sources. This is the normal case. An ASN is allocated by the registries and is used by the organisation to which it has been assigned.
2. An ASN appears on the net but not in the registry statistics. This should never happen as you should not use an ASN without an assignment. Though, as we will see in section 7, it does happen.
3. An ASN has been registered but does not appear in the RIS data. This is called an *ASN Missing In Action*, or ASNMIA. The ASN should have been in use or should have been returned to RIRs when the need for it ended.




Figure 3: Total ASN Count since 1995, RIS ASN count (+), CIDR Report ASN count (x), RIR ASN count (*) and extrapolated RIR ASN count (□).

7 ASN Missing in Statistics

During the period covered by this study, we found 436 ASNs that were not seen in the statistics files but were seen at least once in the RIS. 256 of these were seen for at least 30 days and 247 for at least a year.

Looking at the ASNs, we find 62 out of the 436 to come from reserved blocks that were inappropriately used. The remaining 374 are ASNs that were seen for at least 7 days in the RIS tables, without ever being listed in the Stats Files of any of the registries.

On 31 July 2005, 255 of these 436 were still seen in the RIS. Examining the numbers in detail shows that of those 255, 40 are in the ranges administrated by ARIN and 215 in the ranges by the RIPE NCC.

Further research showed that the 215 ASNs from the ranges administrated by the RIPE NCC are all old assignments, pre-dating the current system of registry files. For 214 of them, other records were available showing to whom the ASN was assigned. The 40 ASNs in ARIN's range were reported to ARIN.

Besides these 436 ASNs, we also found 7 ASNs where it is unclear which RIR is responsible for them. These are all old assignments that were not properly included in the ERX project [16]. This is being corrected.

These numbers allow us to put an upper limit on the accuracy of the registry system. On 31 July 2005, approximately 33000 ASN were registered by the RIRs. For at most 41, no records were available, or about 0.12%. In practice, the number will be even lower. For the RIPE NCC, the number is 1 ASN out of approximately 10000.

8 The total number of ASN seen.

8.1 The total number

Figure 3 shows the ASN count in RIS routing tables, the combined (and corrected for duplicates) ASN count as in the RIR Stats Files and the RIR ASN count obtained by counting backwards from the June 2002 RIR Stats Files. This is a lower limit of the actual ASN count, because returned, reclaimed historic ASN assignments do not appear in the 2002 Stats Files. We have made two corrections to the data obtained from the Stats Files.

1. 678 ASNs that were part of the ERX project [16], were (accidentally) removed from the statistics files on 6 September 2002 and reintroduced on 10 March 2005. These ASNs have been included for the entire period.
2. 223 ASNs were transferred from the RIPE NCC to AfriNIC in February 2005, but continued to be included in the Stats Files from both registries. We removed these ASNs from the totals.

These corrections were also made in other graphs involving the number of ASNs extracted from the Stats Files.

In addition, we plotted the ASN count as listed in the weekly CIDR reports, starting at the end of 1994. We used the individual RIR stats files in order to get the first assignment date for ASNs that have later been reassigned, see section 5.2.

Several observations can be made about this graph. First of all, the RIR assignment count has two phases. There is a slow growth up to 1999, then an increase and growth at a much higher rate. The change in pace occurred around 1999, at the start of the Internet bubble. At this time, the Internet became a mission critical service for many companies. Many of them became multi-homed in order to have redundancy on their networks. This meant that they needed their own ASNs. The RIS and CIDR counts also reflect this behavior. However, at the end of the Internet bubble (late 2001), the RIS/CIDR curve returned to a slower growth rate, while the RIR curve continued to grow at a higher rate. In other words, LIRs still made plans to deploy ASNs (and RIRs considered them solid enough to assign ASNs), but the plans were never turned into reality and the ASN did not show up on the Internet.

Then, the graph shows that the RIS and CIDR ASN counts are consistent. The RIS always sees a fraction more. This could be due to the larger coverage of the RIS, or due to the fact that the RIS ASN count also looks at BGP updates to find ASNs.

Finally, until about 1999, the RIR and RIS/CIDR curves have about the same slope. In other words, there was a pool of about 500 ASNs that was not in use or “Missing In Action (MIA)”. After 2002, the curves started to diverge and the fraction of ASNs MIA started to grow.

8.2 Per registry

Figure 4 shows the number of ASNs allocated by the registries. The density of the points varies due to the fact that until the middle of 2003, some of the registries only published their statistics on a monthly basis. The ARIN count shows the most interesting pattern:

- August–September 2002, LACNIC ASNs removed from Stats Files




Figure 4: Allocated ASN count per registry. ARIN (+), RIPE NCC (x), APNIC (*), LACNIC (□) and AfrinIC (■).

- October 2002, Early Registration Transfer (ERX) of ASNs. About 680 ASNs that were transferred to RIPE NCC did not make it to the RIPE NCC Stats Files. This was corrected in March 2005.
- January 2003, LACNIC ASNs make a one time reappearance in ARIN stats file.
- June–July 2004, a slowdown, even decrease in ASN count as ARIN chases bad debtors, terminating registration services for some assigned ASNs [15].

8.3 Growth rates of RIS and RIR ASN counts

Looking at the curves in Figure 3, it appears that over the last 2 to 3 years, the net growth of ASNs remained constant. Other reports [17] suggest an exponential behavior.

In order to prove our hypothesis, we fitted the total number of ASNs assigned by the registries over time (t) to two equations, first a polynomial:

$$\#ASN(t) = \sum_{n=1}^N a_n t^n \quad (1)$$

for various values of N up to 4, and then to an exponential curve:

$$\#ASN(t) = p_o e^{\alpha t} \quad (2)$$

and the fitting parameter χ^2 was calculated. The lowest value of χ^2 was obtained for (1) with $N = 1$ or a simple linear curve. This is further illustrated in Figure 5, where we plot the data since 1 January 2003 and a fit to a linear and an exponential curve. The plots shows




Figure 5: Linear (solid line) and exponential fit (dashed line) to the number of ASNs assigned by the registries.

that the linear curve fits the data over the entire period, whereas the exponential curve fails to describe the data near the start and the end of the period. Polynomials of order 2 and above also fail to describe the data.

Another indication that the behavior is linear can be seen when (2) is written as a power-series:

$$\#ASN(t) = p_0 \sum_{n=0}^{\infty} \frac{(\alpha t)^n}{n!} \quad (3)$$

This is, of course, a polynomial and when fitting to (1) you might expect the coefficients a_n to behave as $\frac{a_{n+1}}{a_n} = \frac{\alpha}{n+1}$. This is not the case, for a polynomial of order 2 or higher, the 2nd and higher coefficients do not show this behavior. As a final test, we plotted the number of ASNs assigned per day, or the first derivative of (3), over time and repeated our fits. If the behavior is linear, then the first derivative will be constant, whereas for an exponential behavior the first derivative will still be an exponential curve. The result can be seen in Figure 6 and no trend other than constant is visible.

In short, we do not see any indication that the net growth rate behaves other than linear.

In Figure 7, we plot the number of ASNs assigned each month by the RIRs over the last years, as well as the number of ASNs that disappeared from the statistics during that period. Since January 2002 the rate of first assignments is more or less constant at about 284 ± 43 ASN/month. The number that disappears varies more, but is about 40 ± 31 ASN/month for the last 2 years.

The figure also shows the number of ASNs that appeared in the Stats Files for the first time, and the number that was assigned after having been assigned earlier.

Until 2005, the first and last assignment graphs more or less follow the same pattern, which indicates that most ASN assignments didn't change assignment date in a later Stats File, or if they did change, the new date (last assigned) is within a few months from the first date.




Figure 6: Linear (solid line) and exponential (dashed line) fit to the first derivative of the number of ASNs assigned by the registries.

After 1 January 2005, however, we see a widening gap between last assigned and first assigned lines. This means that a sizable fraction of all 2005 assignments were ASNs that had been assigned in years before but are now recycled, reassigned to a different entity. In numbers: from the 1809 ASNs assigned in 2005, 521 first had an assignment date at least 1 year earlier. The majority, 81%, of these are in ARIN.

It is also interesting to note that as of June 2004, there is a significant number of ASNs disappearing from the Stats Files. These ASNs have been revoked or were returned to a RIR when no longer needed. Looking at ASNs that have not been in RIR stats for at least one month, we find ARIN to again have the majority. 307 out of 358 ASNs leaving the RIR stats files in 2005 were last assigned by ARIN.

You can draw a similar figure for the data seen by the RIS, this is Figure 8. Over the last years, the number of new ASNs seen in the RIS grew by about 270 ± 25 ASN/month. At the same time, about 105 ± 15 ASN/month disappeared from the RIS.

8.4 Growth rates per registry

If we split out the growth rates over the individual RIRs, as shown in Figure 9, we see remarkable differences.

The ARIN rates for assigning ASNs that never have been allocated before have been steadily decreasing, while in the same period those of the RIPE NCC have been increasing. So far, these effects have compensated each other, resulting in linear growth as the best fit to overall consumption of ASNs from the unallocated pool, but if the trend of the past 3 years in RIPE NCC assignment rates continues, the overall global consumption rate will increase faster than linear.




Figure 7: Monthly RIR assignments and RIR stats changes. (Re-) assigned by the RIRs (\times), new in RIR stats ($*$), leaving RIR stats (\square).







Figure 8: Monthly RIR assignments and RIS observations. (Re-) assigned by the RIRs ($+$), new arrivals in RIS (\times), disappearing from RIS ($*$). The horizontal lines indicate the averages of 284 and 105.




(a) APNIC



(b) ARIN



(c) LACNIC



(d) RIPE NCC

Figure 9: ASN Assignment rates for the 4 RIRs. In all figures: ASN first assigned from the pool of unused ASN (+), ASN first assigned or reassigned (x) and ASN no longer assigned (*).

9 Fraction of ASNs seen

In Figure 10, the RIS ASN count is divided by the extrapolated (before June 2002) and actual (after June 2002) RIR ASN count. Since March 2001, the fraction of ASNs seen in RIS, as opposed to those assigned by RIRs, fluctuates between 0.6 and 0.63.

10 ASNs no longer seen

On 31 July 2005, the RIRs had assigned about 12700 more ASNs than were seen in the RIS tables. These can be divided into two categories:

1. ASNs that were once in use but have since disappeared from the routing tables,
2. ASNs that have been assigned but have not been used, see section 12.




Figure 10: Fraction of ASNs seen in the routing system over time. RIS ASN count over RIR stats (+), RIS ASN count over RIR countback (×), CIDR report over RIR stats (*) and CIDR report over RIR countback (□).




Figure 11: Lifetime distribution of all ASNs no longer seen in the RIS. 5163 ASNs have been allocated but have not been seen for at least 3 months.




Figure 12: Lifetime in RIS of ASNs assigned and seen after 18 August 2000, in months (30 day periods). *Last_seen - first_seen (months)* (+), *total_days_seen (30 day units)* (x) and *last_seen - last_assign (months)* (*).

From the 32802 ASNs in the combined RIR Stats Files of 1 August 2005, only 20101 showed up in the RIS on that day. 5046 were at least once in the routing table but were not seen in the 3 months before 1 August 2005.

We also see 7037 ASNs that have been allocated but were never seen in the RIS on any day before 1 August 2005. This includes 4338 ASNs that were allocated in the previous century, before 1 January 2001. Though we expect a few percent of the ASN assignments from 2004 to appear in the RIS at a later stage, the majority of them are unlikely to ever appear in the global routing tables.

To get an idea of how long ASNs had been in use before they disappeared from the routing system, we calculate for each ASN the number of days between the last day in RIS and the last assignment date, divide by 30 to get an (approximate) lifetime in months and then plotted the resulting distribution.

Figure 11 shows a majority of ASNs have a lifetime of just a few years (< 50 months). When looking at the data, it is important to remember that the majority of these ASNs were assigned in the last few years. The peak at 160 months is when US Airforce Head Quarter System Center stopped using (or advertising to the public Internet) a large number of their ASNs. The “noise” with negative values are from ASNs that were seen before their alleged first assignment. This could be either an error in the Stats Files or ASNs which were used inappropriately.

Figure 12 shows for all ASNs assigned after 18 August 2000 (first day of RIS data) how long they were seen in the RIS. This includes both ASNs that are active and ASNs that have gone missing. On 31 July 2005, 12743 out of 15032 were seen in the RIS. The graph has three lines. The first looks at the lifetime in RIS assuming 100% visibility between last and first day seen. The second plots the lifetime using the actually recorded number of days it was




Figure 13: Fraction of ASNs seen in the RIS on 31 July 2005 as a function of their assignment date.

seen in RIS. The last takes lifetime as the difference between last_seen and last_assign dates.

No significant pattern differences are seen between the (last-first) seen and (days_seen) graph, which indicates most ASNs seen are seen for the full time between their first and last appearance in the routing system. The distributions are slightly skewed, more shorter than longer living ASNs. That could be a side effect of startup effects, where it takes time for ASNs to make their way into the routing system after initial assignment by a RIR.

The distribution of lifetime as (last_seen - last_assign) is more even. 50 month “old” ASNs being as abundant as 5 month old ASNs. Compared to the other distributions, there are significantly less ASNs that are just a few months old. This reflects the “slow start” in ASNs appearing in the routing system (see section 12).

11 Age of ASNs in RIS on 31 March 2005

Figure 13 shows the fraction of ASNs seen in the RIS on 31 July 2005 as a function of their assignment date. The oldest ASN in our data set is AS4, assigned 22 February 1984 to ISI/University of Southern California.

Looking at Figure 13, you will see that the fraction of visible ASNs peaks at about 80% a year after the assignment of the ASN. Then, as the assignments age, the fraction that is visible starts to drop.

We suspect that this is caused by two effects.

- First of all, sites will go out of business. In this case, the network is (usually) shut down and ASN will no longer appear on the network. As there is no incentive to return it, it will disappear from the pool.




Figure 14: Activation Delay (time between assignment and first appearance) for ASNs assigned by the RIRs during the last 4 years (2001 through 2004), for the 4 RIRs APNIC (+), ARIN (x), RIPE NCC (*) and LACNIC (□).

- Smaller sites often merge with others or are bought by larger network operators. When this happens, the networks are often merged into one bigger network and the need for one ASN (usually) disappears. Again there is no incentive to return it and it will again disappear from the pool.

This explains how the fraction of ASN in use can decrease over time.

12 Activation delay of ASNs

12.1 Yearly data

Figure 14 shows how long it takes assigned ASNs to appear in the RIS. Each graph plots as a function of elapsed time since first assignment, the cumulative fraction of ASNs that have




Figure 15: Activation Delay (time between assignment and first appearance) for ASNs assigned by the RIRs during the last 4 years (2001 (+), 2002 (x), 2003 (*) and 2004 (□)).

been seen in the RIS for at least 7 days¹.

Interestingly, even in 2004, some fraction of ARIN and RIPE NCC ASNs were seen in the RIS well before they got assigned. For ARIN, these are mostly ASNs from the US Department of Defence, but others are there too. Given the low value of the ASNs (< 8000), it is likely that the exact date of the assignment is wrong. The assignment is done by one of the the entities predating the RIR system and the data has been transferred multiple times.

The CDF (Cumulative Distribution Function) patterns for 2004 assignments seem to indicate the ceiling with maximum visibility is reached much earlier (and lower). However, these effects are likely to be artificial. The RIS data used ends on 31 July 2005, which means the ASNs assigned in August 2004 and later, had less than a full year to make it to the routing tables.

¹Note, the graphs do not show the total fraction of ASNs seen at a certain date, i.e. ASNs appearing and disappearing; this is not easy to obtain from the current statistics extracted from RIS raw data.

12.2 By registry

Figure 15 shows the same data as in section 12, but now grouped by RIR, plotting the activation delay of the 4 years in one graph. For RIPE NCC, the curves almost overlap in the first part, indicating a 80% statistical probability for an ASN to make its appearance in RIS within ≈ 6 months (180 days) after being first assigned.

12.3 Discussion

ARIN [7] has the policy that the requestor of an AS must have plans to use it within 30 days of assignment. The RIPE NCC does not have a policy, but a public discussion [18] showed that people thought that it would be reasonable for an AS to appear within 3 months after assignment.

Looking at the curves in Figure 14 or 15, you will see that in practice, the situation is different. In the case of ARIN only about 40% of the ASNs appear within 1 month, the time specified in their policies. In the case of the RIPE NCC, 1 out of 3 ASNs does not appear within 3 months.

You will also see the curves flatten off at about 80% after one year. In other words, some 20% of the assigned ASNs are never used, even though a “demonstrated need” was shown at the time of the assignment.

Modeling of the Data.

13 When will the Internet run out of 16 bit ASNs?

In section 8.3, we presented several statistics on the growth rates of RIS and RIR ASN counts. Each of these touches upon different aspects of the ASN assignment and deployment processes. When it comes to the question how long we have before the internet runs out of 16 bit ASN, it is in our opinion best to look at the rate of first assignment. The last, most recent, assignment dates for the ASNs includes a fraction of reissued ASN and thus would provide a biased view on the depletion of the pool of unallocated, never used, AS numbers.

Figure 8 showed the first assignment rate to fluctuate around an average of 284 since 1 January 2002. Extrapolating this into the future, we can estimate when the last available, unallocated 16 bit ASN will be issued. On 31 July 2005, the RIRs had used 33681 of the 64511 available ASN. This leaves us with $30830/284 = 108.5$ months or 9 years for the remaining ASN to be assigned. Taking the error estimate on the average assignment rate into account, this statistic projects 16 bit ASN to run out in 8 to 11 years.

If a more aggressive recovery approach was being used and all ASN disappearing from RIS after 31 July 2005 were to be reclaimed, the net growth rate would be $(284 - 105) = 179 \pm 53$ ASN/month. In this case the pool would last for 15 ± 4 years.

Finally, if all ASN not seen in the routing table at present and disappearing from this table in future could be reclaimed, the net growth of unallocated ASNs would become equal to the net growth rate of the routing table, i.e. 160 ± 40 ASN/month. On 31 July 2005, 20400 ASN were seen in the routing table, leaving some 44000 to be (re-)assigned. With the present growth that would take 23 ± 5 years.

The first case (9 ± 1 years) is slightly worrying. A proposal to increase the number of bits for an ASNs to 32 exists. If this proposal would be implemented, the number of ASNs will increase to 4.2×10^9 ASNs and the pool would last for a very long time.

However, 9 ± 1 years means that the Internet community only has some 8 years to standardize, implement and deploy 32 bit ASN. This is relatively short, considering the complexity of the problem.

All these expectations assume that the RIR policies will remain unchanged.

14 ASN Assignment Policy Changes

Besides increasing the number of bits, another way to make the pool of ASNs last longer, would be to change the ASN assignment policies. From the data in this paper, we can suggest several possibilities.

The RIRs could reconsider the “demonstrated need” criteria currently used. 20% of the recently assigned ASNs never appear on the Internet and only 63% of all ASNs are actively being used. This suggests that it is too easy to “demonstrate need” even if you do not have an actual need. Stricter criteria could be put in place. The implementation of this idea is left to the policy working groups of the RIRs.

Another approach would be to actively recover unused resources. So far, there has been no good mechanism for this. However, APNIC and RIPE NCC are now looking into certification of resources [19,20]. When this is implemented, each ASN assigned would be accompanied by an electronic document showing who can use the ASN and during which time period. With certification, recovery would be easy. If a certificate expires, the resource is no longer needed and can be reassigned.

From the data in this paper, it appears reasonable to initially assign ASN for a period of 1 year. If the ASN shows up in the routing tables during that year, the certificate can be extended. If not, the certificate can be revoked and the ASN reassigned. This ensures that the 20% of the ASNs assigned but never used, would be recovered and the growth rate would slow down by about 20%.

For ASNs that appear in the routing system, the certificate can be extended, for example for another period of 1 year, allowing the site to continue to use the number without interruptions or the need to renumber to a different ASN. If a site stops using an ASN, the certificate will expire and the resource can be reassigned.

A combination of RIS and registry data can be used to automatically warn sites that a certificate is about to expire and even automatically extend certificates. Only for ASN that are intentionally not visible, would LIRs have to take action.

This mechanism does require that sites peering with an AS will have to check the validity of the accompanying certificate before accepting routes from it. However, as securing the routing system is likely to become standard practice, this will not create much overhead.

15 Conclusions

In this paper, we studied ASN assignments by the RIRs and the number of unique ASNs seen by routers on the Internet.

The rate of ASNs assigned by the RIRs is about 284 ± 43 ASN/month. At this rate, all available ASNs will have been assigned somewhere between 2013 and 2016.

By comparison, the routing table only grows by about 183 ± 44 ASN/month. We have shown that are two reasons for this:

1. ASNs that are assigned, based on future plans, but are never actually used, or used in such a way that they are not visible on the global Internet. However, we believe that the impact of the latter is small.
2. ASNs that were once seen but are no longer in use.

If all these ASNs could be recovered, the pool of ASNs would last until somewhere between 2025 to 2030. One way to accomplish this would be through the use of certification of ASNs.

It can take considerable time for ASNs to appear in the RIS after being assigned. About 10% of the ASNs assigned in the RIPE region in recent years never show up in the routing tables, with higher numbers for the other regions. It could be that the criteria used to decide whether to assign an ASN are too weak.

If the present situation continues, the pool of ASNs will soon run out. LIRs need to start to considering deployment plans for 4-byte ASNs.

Acknowledgments

The authors would like to thank Alex Le Heux, for investigating what happened to the ASN missing from the RIPE NCC statistics, Tony Bates (CISCO), Philip Smith (CISCO) and Geoff Huston (APNIC) for running the CIDR report, Cathy Murphy (ARIN) for the explanation of the ARIN ASN numbers, Daniel Karrenberg, Leo Vegoda and Filiz Yilmaz for their feedback on this paper, all colleagues from RIPE NCC involved in the RIS project, all staff from the 5 RIRs involved in maintaining the statistics files, and finally Adrian Bedford for his comments on the English.

This paper was written in a large part while waiting for various modes of transportation. HU would like to thank the Nederlandse Spoorwegen (Dutch Railways) and Delta Airlines for all train and plane delays, without them, he would not have found the time to finish this paper.

References

- [1] Y. Rekhter, T. Li and S. Hares (editors), “A Border Gateway Protocol 4 (BGP-4)”, `draft-ietf-idr-bgp4-26.txt` (work in progress).
- [2] H. Uijterwaal, Results presented in the RIPE Routing WG, RIPE44, <http://www.ripe.net/ripe/meetings/ripe-44/presentations/#routing>, January 2003.
- [3] R. Wilhelm, Results presented at the RIPE 50 meeting, <http://www.ripe.net/ripe/meetings/ripe-50/presentations>, May 2005,
- [4] Q. Vohra, E. Chen, “BGP support for four-octet AS number space”, `draft-ietf-idr-as4bytes-10.txt` (work in progress).

- [5] IANA AS-Number Assignments, <http://www.iana.org/assignments/as-numbers>,
- [6] M. Kühne, N. Nimpuno, S. Wilmot, “Autonomous System (AS) Number Assignment Policies and Procedures”, RIPE-263, <http://www.ripe.net/ripe/docs/asn-assignment.html>, January 2003.
- [7] ARIN, “ARIN Number Resource Policy Manual”, <http://www.arin.net/policy/nrpm.html>, January 2005.
- [8] APNIC, “Policies for Autonomous System number management in the Asia Pacific region”, <http://www.apnic.net/docs/policy/asn-policy.html>, February 2004.
- [9] LACNIC, “Internet Resource Management Policies in Latin America and the Caribbean”, http://lacnic.net/en/politicas/2002-11-num_sis.html, November 2002.
- [10] A. Akplogan, E. Byaruhanga, “Autonomous System Numbers Management in the AfriNIC region”, <http://www.afrinic.net/docs/policies/afpol-as200407-000.htm>, March 2004.
- [11] J. Hawkinson, T. Bates, “ Guidelines for creation, selection, and registration of an Autonomous System (AS)”, RFC 1930, March 1996, <ftp://ftp.rfc-editor.org/in-notes/rfc1930.txt>.
- [12] RIPE NCC Routing Information Service (RIS), <http://www.ripe.net/ris>.
- [13] Available by anonymous ftp at: <ftp://ftp.<registry>.net/pub/stats/>, with registry one of afrinic, apnic, arin, lacnic or ripenc.
- [14] T. Bates, Ph. Smith and G. Huston, “The CIDR report”, <http://www.cidr-report.org/>.
- [15] C. Murphy, ARIN, private communication.
- [16] ERX project, <http://www.ripe.net/projext/erx>.
- [17] G. Huston, “Exploring Autonomous System Numbers”, in “The ISP Column”, August 2005
- [18] See <http://www.ripe.net/ripe/maillists/archives/lir-wg/2002/msg00780.html> for the archives of this discussion.
- [19] G. Michaelson, “APNIC Resource Certification”, <http://www.apnic.net/meetings/20/programme/sigs/routing.html>, September 2005.
- [20] RIPE NCC, “Activity Plan 2006”, Unpublished at the time of writing.