

# The ISP Column

*An occasional column on things Internet*

---

August 2005

Geoff Huston

## Exploring Autonomous System Numbers

So what's an "Autonomous System Number", and what role do these numbers play in the technology of the Internet? In this article I'd like to explore the Autonomous System space and the associated Autonomous System Number pool in some detail.

### Part 1 - Using AS Numbers

To set the stage for this examination of Autonomous Systems Numbers it will be necessary to set the stage with a brief overview of the Internet's routing architecture.

#### Internet Routing

The Internet's routing architecture is structured as a two-level hierarchy. The environment is firstly partitioned into "domains", with each domain using an internal routing environment. These network domains use an interior routing protocol (commonly referred to as an "IGP", after the term "*interior gateway protocol*"), which maintains a complete mapping set for the current internal topology of the domain, together with the set of "best paths" between any two points within the network domain. While this approach of automatically maintaining a comprehensive view of the current topology of a network through the actions of a distribution computation in the form of a routing protocol can be made to work within even quite large routing domains, such an approach does not scale to the size of the entire Internet, within the capabilities of any known routing technology, nor does it even make much sense for fine-grained topology information within one part of the Internet to be passed to all other parts of the network. Fine-grained topology information is only useful in "local" situations, and is best omitted when forming a larger view of the network.

The approach used within the Internet's routing architecture to partition the extent to which this fine-grained topology information is propagated across the Internet is to call upon a second level of routing hierarchy. This "*inter-domain*" routing environment maintains a somewhat different map of the network. The inter-domain routing environment supports a description of how each of the local routing domains are connected to each other, but avoids the task of also maintaining the internal topology of each domain. In the inter-domain space a path to an address is described as a sequence of domains that must be transited to reach the domain that originates the particular address prefix. This inter-domain space is maintained these days using version 4 of the Border Gateway Protocol (BGP4).

Each routing domain is a single administrative domain, operated within a uniform set of routing policies, and is operated independently from any other domain. The domain is in effect an autonomous unit in the overall routing architecture, and is termed an "Autonomous System" (AS). Each of these

---

AS's are uniquely identified using an Autonomous System Number (ASN), which brings us to the topic of consideration of this ASN number pool.

In this article we'll first explore how the ASN number space is structured, how ASN's are used in the inter-domain routing environment, and then look at the consumption rate of these numbers and finally examine our options once we get to the point of likely ASN number pool exhaustion.

## What's an Autonomous System?

One of the best definitions of an Autonomous System can be found in an IETF working document that describes the Border Gateway Protocol:

```
The classic definition of an Autonomous System is a set of routers under a single technical administration, using an interior gateway protocol (IGP) and common metrics to determine how to route packets within the AS, and using an inter-AS routing protocol to determine how to route packets to other ASs. Since this classic definition was developed, it has become common for a single AS to use several IGPs and sometimes several sets of metrics within an AS. The use of the term Autonomous System here stresses the fact that, even when multiple IGPs and metrics are used, the administration of an AS appears to other ASs to have a single coherent interior routing plan and presents a consistent picture of what destinations are reachable through it.
```

Internet Draft (Work In Progress): [\*draft-ietf-idr-bgp4-26.txt\*](#)

## The AS Number Pool

AS numbers are drawn from a 16 bit number field, allowing for 65,536 possible values.

AS **0** is reserved, and may be used to identify non-routed networks. The largest value, AS **65,535** is also reserved. The block of AS numbers from **64,512** through to **65,534** is designated for private use. ASN **23,456** is reserved for use in AS number pool transition.

The remainder of the values, from **1** through to **64,511** (less **23,456**), are available for use in Internet routing. The number space is unstructured, there are no internal fields in the number structure, nor is there any aggregation or summarization capability for AS numbers. Autonomous Systems are simply numbered sequentially, allowing for a potential of 100% utilization of the number space.

## How AS Numbers are used in BGP

The inter-domain routing space is constructed using two components: address prefixes and AS numbers, which are used as domain identifiers. Every prefix has an originating domain, known as the "origin AS" from which reachability for the prefix is propagated across the inter-domain space.

As the routing advertisement is propagated across the inter-domain space each prefix accumulates an associated "AS Path". As a prefix advertisement transits each domain, the domain effectively "signs" the prefix advertisement by having its AS number prepended to the AS Path associated with the address prefix. At any point in the network the AS path describes a sequence of connected domains which forms a path from the current point to the originating domain. This is shown in Figure 1, where AS 1 originates an advertisement for the address prefix 192.0.2.0/24. At AS 5, the AS will receive two BGP advertisements for this prefix. One will have the AS Path (4,2,1), while the other will have the AS Path (3,1). The AS Path is a sequence of AS number values. The left-most number is the AS number

of the adjacent Autonomous System from which the address prefix advertisement was received. The sequence of numbers indicates the sequence of Autonomous Systems through which this update was propagated. The right-most, or final AS number is the AS number of the Autonomous System that originated the address prefix advertisement, or "Origin AS".




Figure 1 – AS Path Generation in BGP

AS numbers sequences into AS Paths serve two purposes in inter-domain routing: that of a path length metric and a loop detection mechanism.

The AS Path is used as a path metric in BGP's path selection algorithm. The default route selection metric in BGP is that of selection of the minimal length AS Path, with each AS in the path counting as a single unit of 'cost'. When a BGP speaker receives two or more advertisements for the same address prefix, the default selection mechanism is to prefer the advertisement with the minimal AS Path length. In the case of the example network in Figure 1, AS 5 will prefer to use the path via AS3 to reach the originating AS1, in preference to the longer path of AS4 then AS2.

While enumerating the AS Path vector within the routing protocol is one way of passing the path cost through the routing domain, it may appear in the first instance that the best path selection function could just as easily be supported by carrying a simple path cost metric of a domain transit counter, as is undertaken by other distance vector protocols, such as RIPv2 or EIGRP. However, the issue with distance vector protocols is the "count to infinity" problem.

To illustrate the need for explicit AS Path enumeration in BGP, let's look at what happens when the AS Path vector is replaced by a simple path cost metric. In the configuration shown in Figure 2, AS1 originates a routing advertisement towards AS2, AS2, AS3, and AS4 are interconnected in a simple loop configuration. When AS2 receives AS1's advertisement with a path cost of 1, it will pass the advertisement on to both AS3 and AS4, with a path cost of 2. Both AS3 and AS4 will select as their best path this advertisement from AS2 with a path metric 2, corresponding to the AS path (2, 1). Now if the connection between AS1 and AS2 is broken, then AS2 will no longer see AS1, and will withdraw its best path to the prefix via AS1. AS2 will then stop advertising a path to AS3 and AS4. But AS3 is already advertising a path to AS4, with a metric of 3, corresponding to the AS path (3, 2, 1). Upon the withdrawal of the advertisement from AS2, AS4 will then select this as its next best path, with a path cost of 3. AS4 will then advertise this prefix to AS2 with a path cost of 4, corresponding to the AS path (4, 3, 2, 1). At this point, without the explicit AS path in the advertisement, AS2 cannot deduce that this advertisement is in fact a loop. Accordingly, AS2 will accept this path with a metric of 4 as its best path. AS2 will then advertise this to AS3 with a metric of 5, corresponding to the AS path (2, 4, 3, 2, 1). AS3 will update its best path to AS1 with this new metric and then send an update to AS4, and so on. This process will continue around the loop until the path cost metric reaches some defined maximal value. The higher the maximal value for the path cost metric the longer the time taken to detect the loop condition. The smaller the maximal path cost metric the smaller the span of network that the protocol can encompass. This requires some considerable care in setting the maximal path cost parameter, and the operation of the protocol can be extremely slow to converge in terms of loop detection.




Figure 2 – Loop Formation in Distance Vector Protocols

By replacing the path cost counter with a fully enumerate AS sequence this form of loop can be averted. Continuing the example in Figure 2, when AS2 withdraws its route to AS3 and AS4, AS4 will still select the other route it has heard, but this time the selected prefix will have the path (3, 2, 1). When AS4 attempts to pass this advertisement to AS2, AS2 will see its own value in the associated AS Path and reject the advertisement. At the same time AS3 will withdraw its advertisement to AS4, and at that point the prefix will be dropped from the entire routing system.

The use of AS numbers and AS Path vectors in BGP provides an effective solution to this classic problem of loop detection.

### Who needs an AS Number?

Not every network needs to have its own AS Number. The guiding principle is that AS numbers are used to express distinct inter-domain routing policies, and not every network has the requirement to express its own unique set of routing policies.

In the case where a network has a single upstream connection then the routing policies of the network are precisely the same as those of its upstream service provider, and there would normally be no need for the network to use a distinct AS number. Even if the network domain uses BGP for its upstream connection, the originating domain can use a private AS number (from the number range 64512 – 65534) to support the BGP session to the upstream network. The upstream network will strip off the private AS number when it readvertises the prefix, and the upstream network would appear to the rest of the Internet as the originating AS. Even if the AS has “downstream” networks it can still use a private AS, even when the downstream AS’s are using public AS numbers. The stripping of the private AS only removes the instances of the private AS from the AS path, not the public AS numbers. (Figure 3)




Figure 3 Use of Private AS Numbers

---

In the case where a network has two or more upstream transit connections then it is more likely that the network will use its own unique AS number. It is not always the case that a distinct AS number is required here, and the distinguishing factor is that of the network wanting to express particular routing policies. Where the network has no particular preference as to which of the upstream services should be used for incoming traffic, the network can also use a private AS number for each of its routing sessions. In such a case the external view is that the prefix appears to be originated from multiple AS's.

In the case where there are multiple paths to reach the network, and where these paths need to be distinguished in the routing system by different AS Paths that have the same originating AS (i.e. there is a need to express a routing policy), then the network needs to use a unique AS number within the inter-domain routing system.

## Can you split an AS Number across separated sub-domains?

There are many cases of dispersed networks that exist in multiple locations. If these locations are all administered by a single entity it may be desired to use a single AS number across all these domains. This is possible, but considerable care needs to be exercised when designing the routing configuration.

In the following example there are two distinct sub-domains of AS1, and they are not interconnected internally.




Figure 4 – Split AS

AS1 (A) advertises the prefix 192.0.2.0/25 to AS2, and this advertisement is propagated to AS2, AS3 and AS4. When AS4 passes this advertisement to the other segment of AS1 (B), this router will reject the advertisement because the associated AS path (4,3,2,1) indicates that the route has already passed through AS1. Similarly, the first segment of AS1 (A), will reject the advertisement of 192.0.2.128/25 from AS2, as its path (4,3,2,1) also indicates that a loop has formed. To restore complete connectivity between the distinct parts of AS1, AS1 needs to configure static routes at its edges. If AS1 (A) configures a static route to 192.0.2.128/25 pointing towards AS2, and AS1 (B) similarly configures a route to 192.0.2.0/25 via AS4, then the configuration enables full connectivity.

In more complex configurations where each of the segments of the network are multiply connected the static route configuration becomes more complex. However, with very careful configuration, a single AS number can be distributed across multiple distinct networks.

## How are AS numbers used to express Route Policies?

The basic mechanism of path preference in BGP is that of the AS Path length. Where there are two advertised paths to reach a particular address prefix then the default selection algorithm in BGP is to prefer the advertisement with the shorter AS Path length.

A multi-homed domain may wish to have other domains prefer one particular path over another to reach it. This may be due to the local domain wishing to optimise its traffic costs between the multiple upstream providers, or wishing to balance the traffic load across multiple paths, or set up various forms of primary and backup relationships across the multiple provider upstream paths.

While it is often the case that such policy preferences are set up using BGP communities, BGP community signalling requires the cooperation of multiple parties in consistent interpretation of the community values. A more coarse form of expressing such policy preferences can be achieved through AS Path prepending. AS Path prepending is a technique of deliberately extending the AS Path length of a prefix advertisement by adding additional AS numbers into the AS Path of an advertised prefix. Normally the form of AS Path prepending uses the local AS number to perform the prepending.

In the example in Figure 5, AS1 wants to express the policy to prefer incoming traffic via AS2, and only use the link to AS3 as a backup. To achieve this with AS Path prepending AS1 prepends itself twice in the AP path of the advertisement passed to AS3, in order to artificially lengthen the AS3 transit path. AS5 would've normally used the shorted AS path via AS3 to reach AS1. As a result of AS1 artificially lengthening its path to AS3, AS5 will now select the transit path via AS4 and AS2 to reach AS1.




Figure 5 – AS Path Prepending

Of course AS Path prepending is a very imprecise technique, and can often produce surprising results in real world situations. A more deterministic method of traffic engineering uses additional signals attached to address prefix advertisements, via BGP communities.

A more subtle, and more controversial, prepending technique is that of so-called AS Path Poisoning, where an AS uses some other value to prepend in the AS Path. In the example below AS1 wants to express the policy that under no circumstances should AS5 use the transit via AS3 to reach AS1. In this case AS1 could use AS5 as the prepending value in its advertisement to AS3. When AS5 receives this advertisement, the presence of its own AS number in the AS Path means that it will not accept this advertisement, and will prefer the transit path via AS4 and AS2. The difference between these two examples is that in the case where the connection between AS1 and AS2 is broken none of AS2, AS4 or AS5 can possibly reach AS1 when this AS Path Poisoning technique is being used.




Figure 6 - AS Path Prepending with Poison AS

In the next part of this article we'll look at how quickly the AS Number pool is being consumed and make some predictions as to the most likely point of exhaustion of this space.

---

---

## Part 2 - AS Number Consumption

Of the 64,510 available AS numbers, as of July 2005 we have already allocated some 39,000 AS numbers, or well over half of the number pool. This raises two immediate questions – how long before the number pool is completely exhausted, and what are our options for an expanded number pool that can encompass a larger inter-domain routing environment?

### The Drivers for AS Number Consumption

Before looking at these two questions in further detail, it would be useful to understand the driving factors behind AS number consumption.

From one perspective it would appear to be counter-intuitive to assume that the Internet will evolve from tens of thousands of distinct routing domains to one of hundreds of thousands or even millions of distinct routing domains. Individual routing domains are essentially equivalent to individual Internet Service Providers (ISPs), and in the first instance it would appear that there is a reasonable level of correlation between the number of active ISPs in the Internet and the number of advertised AS numbers. If forecasting a future demand for hundreds of thousands or even millions of AS numbers it would appear that we are forecasting continued fragmentation of the service provider industry with large numbers of small enterprises who collectively compose the Internet. This does not appear to be a likely scenario. The ISP industry has been one that has had a continuing aspect of relatively intense competition between providers, and in many, if not all, market segments there has been an underlying factor of economies of scale. Larger ISPs have access to more efficient use of their resources and are more capable of sustaining a market share at competitive prices with reasonable operating margins due to these economies of scale. Smaller providers tend to service niche markets, and in general are highly susceptible to pricing pressures in the competitive supply market. The overall result is strong pressure for continued aggregation in the service provider market, tending to aggregate to a smaller number of larger providers.

From this perspective, if the number of AS's in use is roughly commensurate to the number of service providers, then this view of the market dynamics would lead to a view that the service provider population is either in a state of equilibrium where the entrance of new niche-oriented players is roughly the same as the rate at which smaller players are aggregated into larger providers, or one of relatively small growth based on the larger dynamics of continued expansion of the Internet on a global basis.

In practice this has not been the case, and we see a continuous rate of consumption of new AS numbers. This rate appears to be some 3,500 AS Numbers per year, or an average of 9.8 AS numbers per day, and this consumption rate appears to have been steady since early 2003. Does this signify a continuous rate of new entrants into the ISP industry at a global level, or are other factors at play here? Certainly there do appear to be some additional factors that have some bearing on AS number consumption rates.

One of these factors is the practice of multi-homing at the edge of the network. Many end site networks have business-critical needs for assured Internet connectivity, and a common way to achieve this is by using the services of two or more upstream providers. In such situations the end site may want to express different routing policies to each upstream provider, and it does so through using its own AS number and expressing these routing policies using BGP to each of its upstreams.

AS numbers are also used in other contexts. In MPLS L3 networks one form of generating the Route Distinguisher value for a VPN client network is through the use of concatenating a VPN host AS number concatenated with a serial number (RFC 2547). To what extent this semi-private use of AS



numbers in a VPN context contributes to the consumption rate of AS numbers if difficult to assess, simply because these numbers may not appear in the public Internet.

With the public Internet there are other contributory factors to AS Number consumption. ISP's with diverse product portfolios may wish to express different routing policies for various product families. Again this can be achieved through the use of distinct AS numbers of each routing policy set.

Also there is little incentive for AS Number return and recycling. With the current framework there is no cost to maintain an AS number allocation, and the overall characteristic of AS number allocation appears to be a "once and forever" allocation model. Once AS numbers are no longer required in general AS numbers generally do not return to the unallocated pool for subsequent recycling.

Taken together these factors lead to the conclusion that continued AS number consumption is based on a larger set of considerations than the dynamics of the service provider industry. Accordingly we can be a little more confident in making the assumption that the drivers for AS number consumption in the recent past will continue to be drivers in the near term future. This leads to some further confidence in a predictive technique that uses recent consumption data to generate trends that can made predictive forecasts of future demands.

We'll apply this technique to AS number consumption data to make some forecasts of the time by which the current AS number pool will be exhausted.

## AS Number Pool Status

There are some 65,536 AS numbers. As noted already some 1026 numbers are reserved and unable to be used in the public Internet. This leaves 64,510 for use in the Public Internet.

The pool of AS numbers is administered by the Internet Assigned Numbers Authority (IANA), and blocks of 1024 numbers are allocated to the Regional Internet Registries (RIRs) periodically when the RIR's pool drops below a threshold level.

As of July 2005 the IANA is holding 25,600 numbers, while 38,910 numbers have been allocated to the RIRs. The overall picture of the AS number space is shown in Figure 1.




Figure 1 – AS Pool Status

The pool of 38,910 allocated numbers is further divided up to each of the RIRs, and the complete picture of AS number allocations is shown in Figure 2. The large number of allocations listed to the ARIN registry includes a component of historical allocations undertaken prior to the establishment of the RIR system, so this is not an accurate correlation of the distribution of AS numbers across the various global regions.

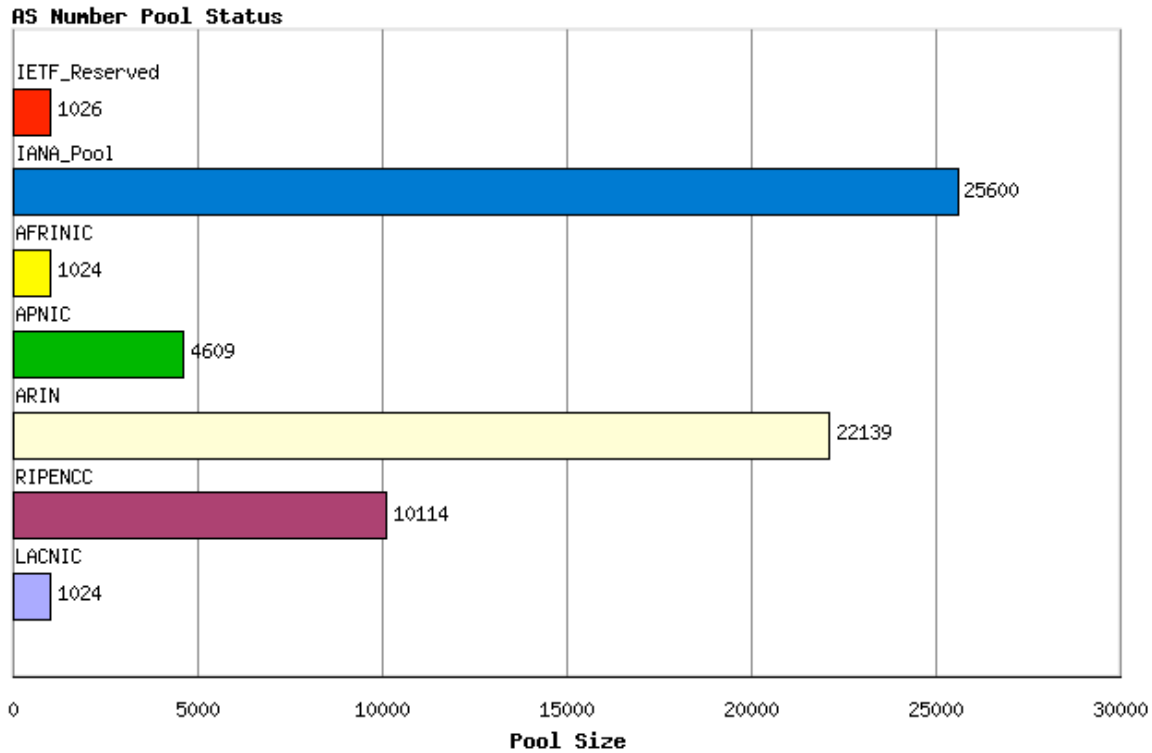


Figure 2 – AS allocations to RIRs

While some 38,910 AS numbers have been allocated by IANA, there is a further classification of AS numbers. A working pool of numbers is held by the RIR for subsequent assignment to ISPs. Of the assigned AS numbers some are visibly used in the inter-domain routing table of the public Internet, while others are not visible in the Internet. The breakdown of AS numbers into the RIR pool, assigned but not advertised, and assigned and advertised, as of July 2005, is shown in Figure 3. Of the 32,557 assigned AS numbers, some 19,859 are advertised, while 12,698 have been allocated in the past, but are not currently advertised in the BGP routing table.

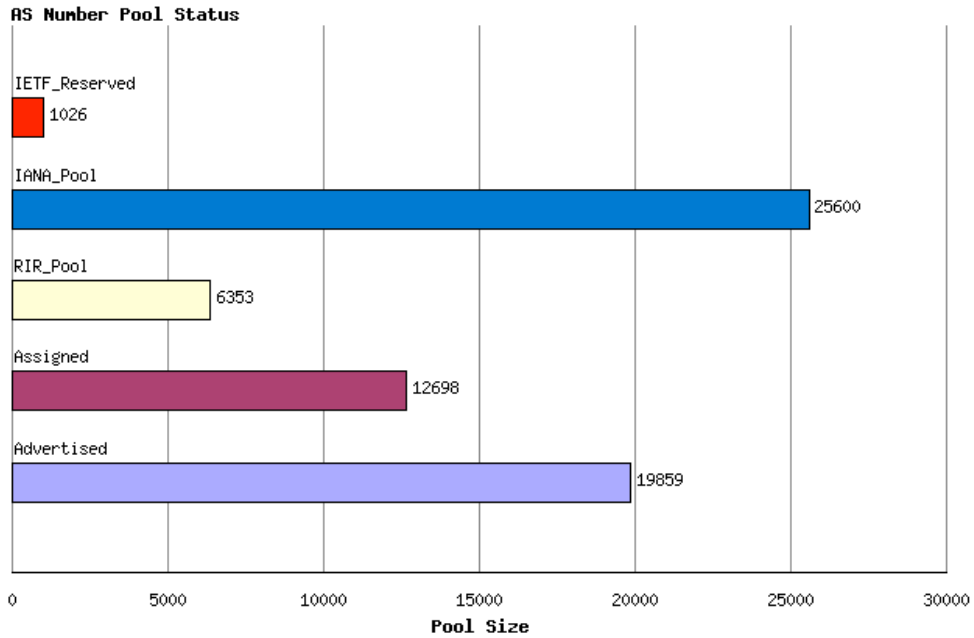


Figure 3 – AS number status of Advertised, Unadvertised and Unallocated pools

Again, this can be further broken down to each RIR, as shown in Figure 4. It is evident from this view that the ratio of unadvertised to advertised AS's varies across the RIR system. The highest ratio is in the ARIN region, which also coincides with the earliest allocations. There appears to be some relationship between the probability of an AS being visible in the routing system and the allocation lifetime of an AS.




Figure 4 – AS status pools by RIR

## IANA AS Allocations

It is also possible to construct a number of time series of the AS allocation data, showing the growth in the use of AS numbers over the past decade or so. The first of these series is the allocation of blocks of AS numbers from IANA to the RIRs, shown in Figure 5.




Figure 5 – IANA AS Allocation running total

This is not a smooth series as IANA makes periodic allocations of blocks of 1024 AS numbers to each RIR. This can be further classified into per-RIR allocations, as shown in Figure 6.

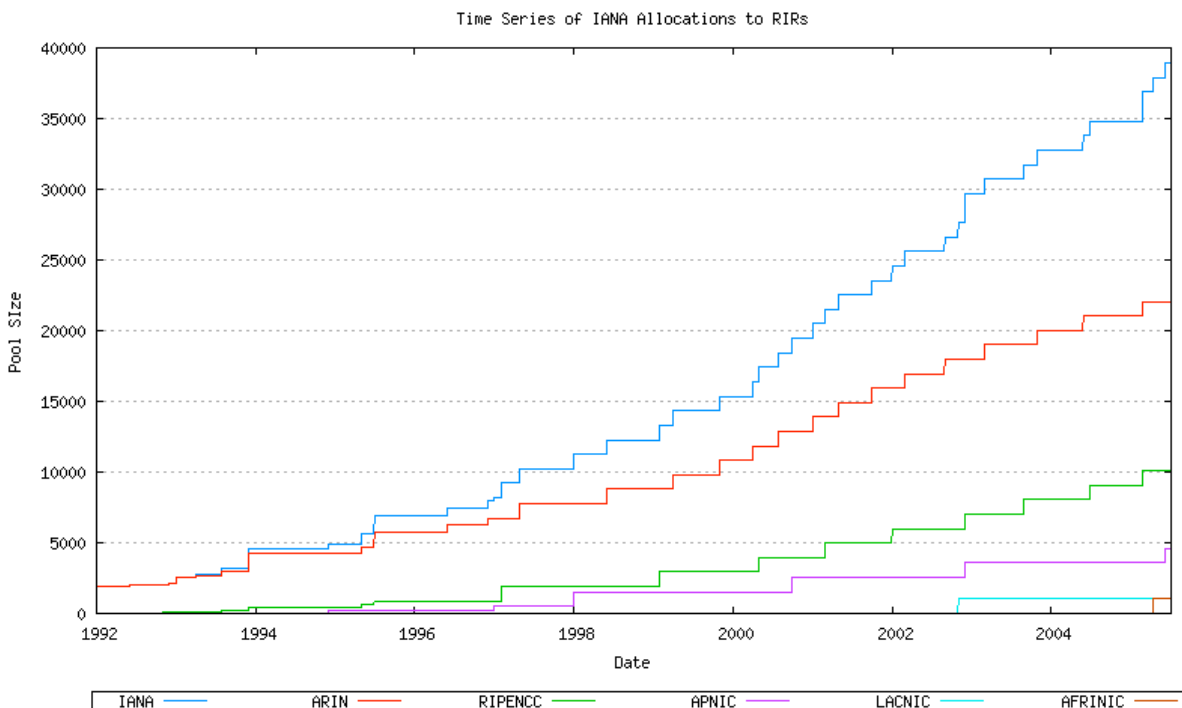


Figure 6 – IANA AS allocations per RIR

It is evident that the rate of AS allocation was highest in the ARIN region across the Internet boom of 1999 – 2002, while more recent allocation activity is evident in the European region of the RIPE NCC.

## RIR AS Allocations

These allocated AS's are then assigned to ISPs by the RIRs, and a second time series can be generated, showing the rate of AS assignment to ISPs and Local Internet Registries. This allocation time series is shown in Figure 7. Again the series shows the effects of the Internet boom across the period from 1999 through to late 2001. The subsequent market correction is also evident, which was visible in the AS allocation rate change by early 2002.




Figure 7 – RIR Allocations

The data can also be presented on a per-RIR basis, as shown in Figure 8. Overall the picture corresponds to the earlier IANA allocation data, where the rate of allocations in the ARIN region appears to be slowing down, while the trend in Europe in the RIPE NCC-serviced regions shows rising levels of activity in recent years.

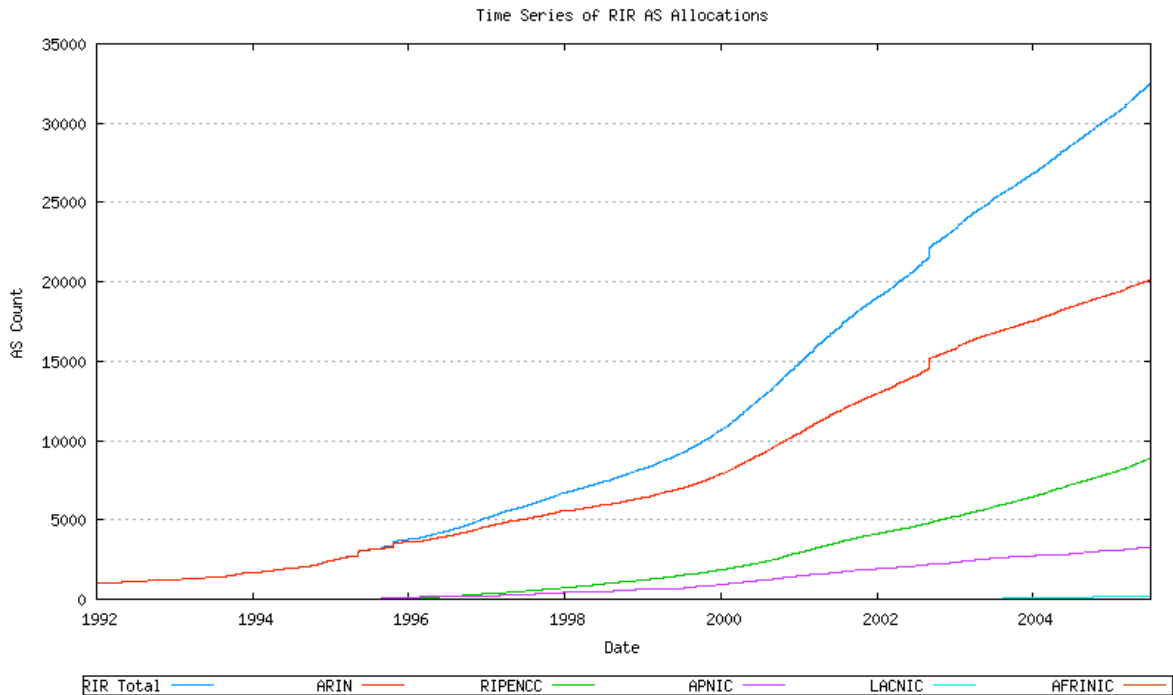


Figure 8 – RIR allocations

## BGP AS Advertisements

As well as looking at allocation rates, a further source of AS number data is the inter-domain routing table. The number of distinct Autonomous Systems advertised in the inter-domain routing space of the public Internet has been measured on a regular basis since 1997. The time series of this count since that date is shown in Figure 9.


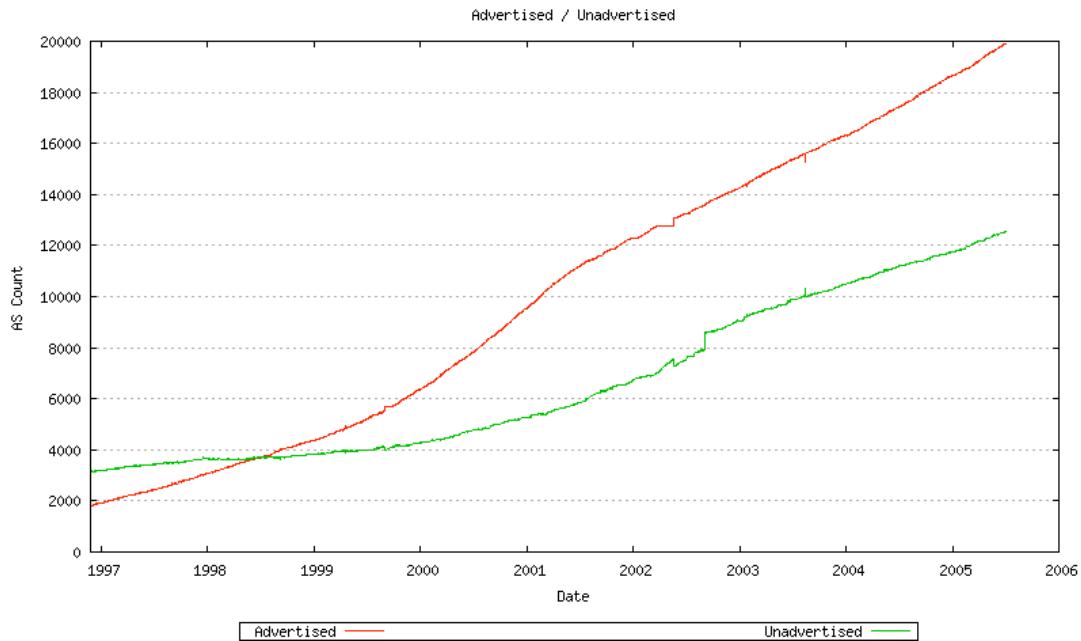


Figure 9 – Advertised AS Numbers

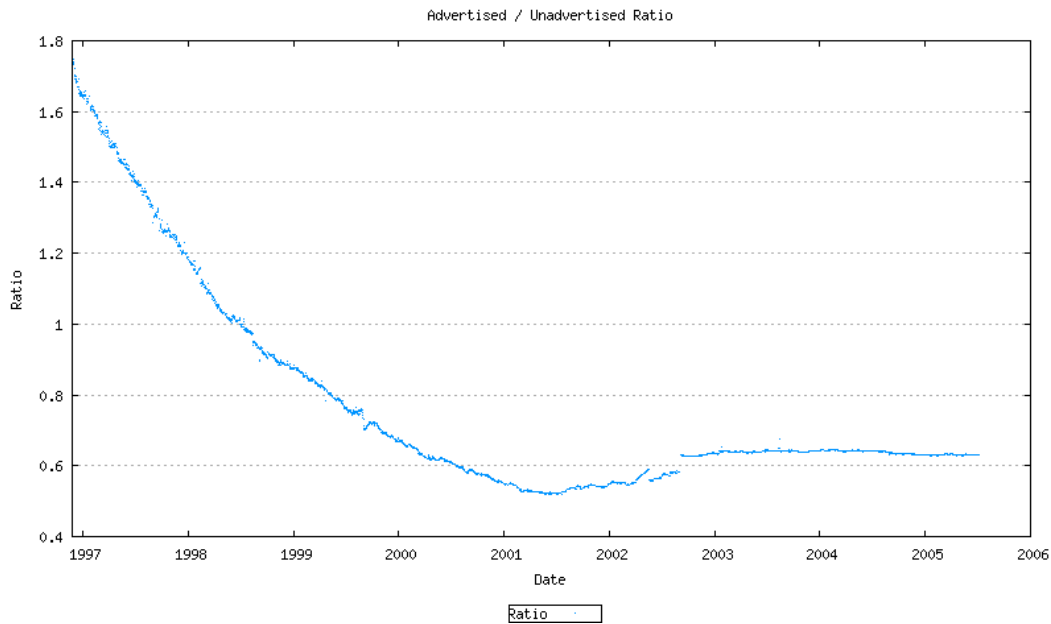
The current count of advertised AS numbers is some 20,000, as of July 2005. This time series has a visible correlation to the recent Internet business cycles, with a sharp upward trend across the Internet boom period of 1999 – 2001, and a marked change in the trend in early 2001. Since that date the rate of increase of AS numbers has been far slower.

There are also some 12,600 AS numbers that have been allocated, but are not advertised in the BGP table. In terms of generating projections of AS number consumption it is also necessary to factor in the trends in this unadvertised AS number pool. This is shown in Figure 10.



*Figure 10 – Advertised and UnAdvertised AS Numbers*

It is evident that the number of unadvertised ASs is rising over time. The consequent question is whether this rate of increase is greater or less than the number of advertised ASs? One way to look at this is to look at the ratio of unadvertised to advertised ASs over time. This is shown in Figure 11.




*Figure 11 – Time Series Ratio of Unadvertised to Advertised AS’s*

This figure indicates that since 2003 the rate of growth of unadvertised AS’s is slightly lower than the growth rate of advertised ASs.

## AS Disposition by AS Number Block

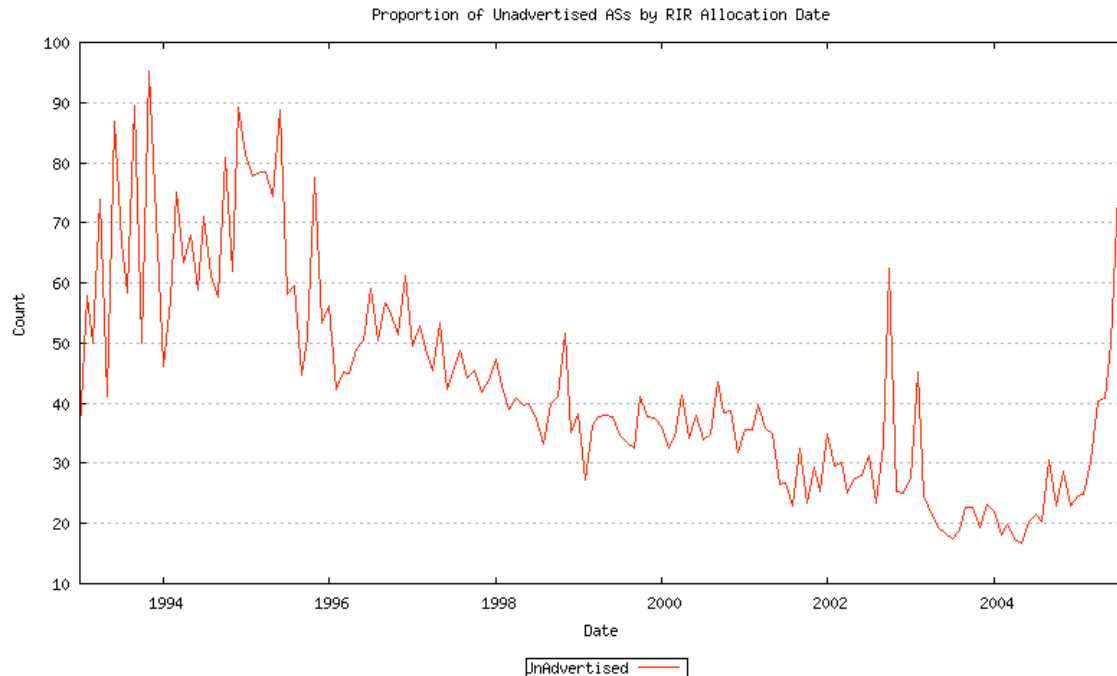
Another way to look at the AS number space is to look at the AS Numbers in blocks of 256 AS numbers. The AS number pool has been one that to a large extent has been consumed in numerical sequence. Smaller AS numbers were allocated first, and larger AS numbers were allocated more recently. Figure 12 shows the breakdown of each block of 256 AS numbers, looking at how many numbers from each of these smaller pools is currently advertised, how much is not advertised, and how much is in the RIR pre-allocation pools.





*Figure 12 – AS Number Status by Block*

Interestingly the older AS number blocks appear to have a much higher amount of unadvertised ASs than more recent AS number blocks. Another way to look at this is to examine the age distribution of Unadvertised AS's. As shown in Figure 13, the age distribution of unadvertised AS numbers is one where the probability that an AS is not visible in the routing table appears to be directly related to the period that has elapsed since allocation. The figure also shows a peak for very recent allocations, and it would appear that there is a delay of up to three months between the allocation action and the appearance of the AS in the BGP routing tables. What this table indicates is that there is a constant rate of change in the Internet environment, and without any major reclamation of unused AS numbers then once an AS number is no longer required it does not get handed back to the unallocated pool for reuse, but is placed into a dormant state where the AS is considered to be allocated, but is unadvertised.



*Figure 13 Age Distribution of Unadvertised AS Numbers*

## AS Consumption Projections

At this point it is now possible to make some projections on AS number consumption. The technique here is to use the available data from a starting point of the start of 2003 and using a least squares best fit to the data generate trend data series. This trend data is then projected forward in time to forecast the point in time when the resource reaches a certain threshold point.

This technique is perhaps best shown in Figure 14. Here the IANA allocation data is used as the base data series, and two projection models are used. The first model is to apply a linear best fit to the data. A linear model assumes that the underlying drivers for resource use are constant and the rate of growth over time remains a constant. The second model is an exponential growth model, where the underlying drivers increase over time. The model used here is an exponential model, where the growth rate doubles over a certain time interval. The definition of the doubling time interval is based on a least squares best fit of an exponential growth curve to the available data.

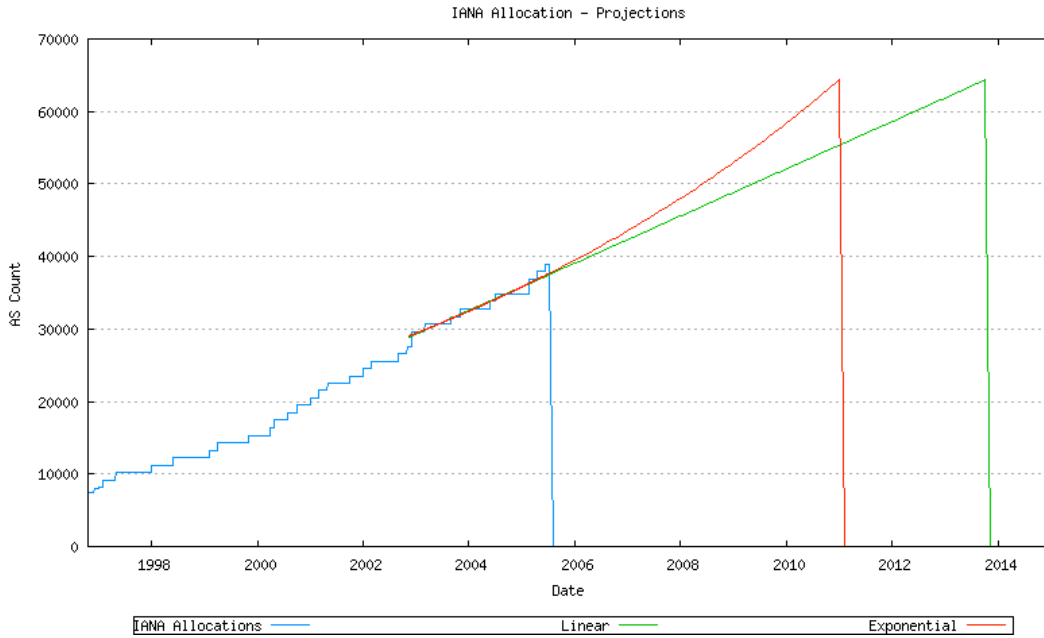


Figure 14 – Projections of IANA AS Allocations

This projection shows that a best fit to a linear (constant growth) model of AS number allocation from the IANA pool would see the final AS block being allocated in the final quarter of 2013. An exponential (increasing growth) model of AS number allocation from the IANA pool would see the final AS block being allocated in early 2011.

Do the other data series provide comparable outcomes?

The next place to look is in the RIR allocation data. The true point of AS number exhaustion is not when there are no further AS's for IANA to allocate to the RIRs, but the point at which the RIRs have no further AS numbers to allocate. A similar technique can be applied to the RIR allocation data. This is indicate in Figure 15.

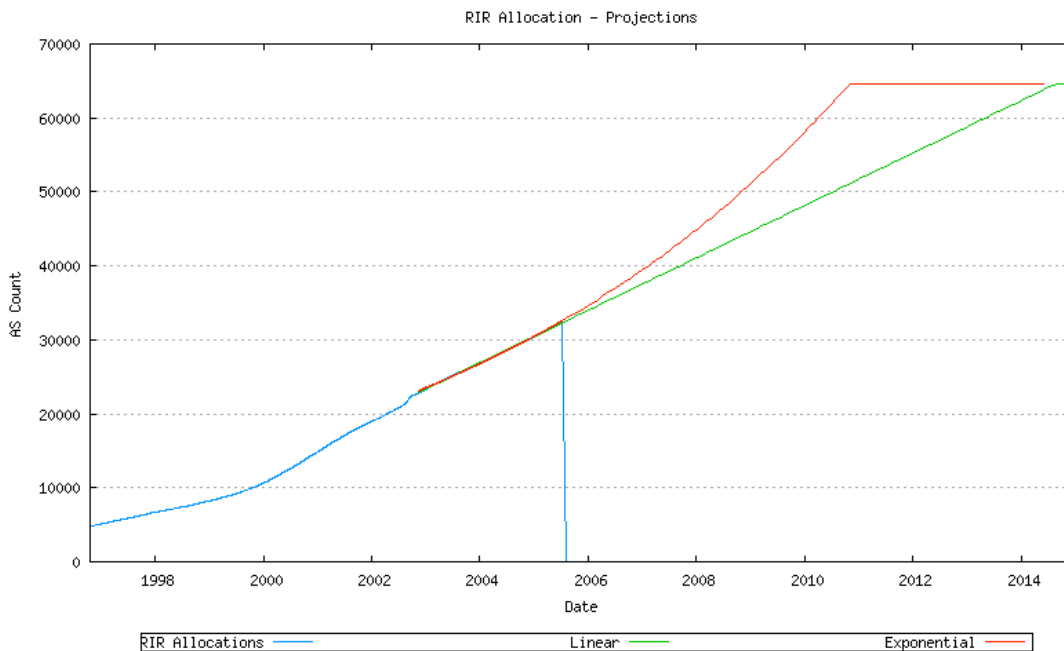


Figure 15 – Projections of RIR AS Allocations

This projection shows a similar picture. A linear model projects exhaustion of the RIR AS number pool in mid 2014, and an exponential growth model predicts exhaustion by the final quarter of 2010. This is certainly reasonably aligned to the IANA data.

So which model is the best reflection of the underlying drivers of AS number consumption – linear or exponential?

One way to delve a little deeper into the data is to look at the first order differential of allocation rates, to see if there is a discernable trend in that data series. The daily rate of AS allocations is shown in Figure 16




Figure 16 – First order differential of the time series RIR Allocations

This appears to indicate that the daily rate of AS allocation has increased over time. Between mid 2003 and the recent past the daily allocation rate appears to have increased from a mean of 9 to a mean of 12 AS's per day. This would indicate that a linear model appears to be an overly conservative predictive model, and does not accurately reflect the recent history of AS allocations.

The other sanity check of the most appropriate projection model is to look at the first order differential of the logarithm of the allocation rate. If this is a constant value overtime then the exponential growth model would be a best fit to the data.




Figure 17 – First order differential of the logarithm of the time series RIR Allocations

Here the time series appears to be relatively constant since mid-2003, with a mean value of 0.0035.

The next data sequence to examine is that of the advertised AS data series. A projection of Linear and Trend (exponential) models onto the data is shown below.




Figure 18 – Projections of AS Advertisements

This projection is flawed, in that it assumes that all AS numbers will, eventually, be reclaimed and used in the Internet's routing space. A slightly more relevant model of consumption needs to take into account that unadvertised AS numbers will not be reclaimed and reused in the routing system.

The model used here is to project forward the number of AS's found in the Internet's BGP table (advertised AS's), and also model the best fit projection of the ratio of unadvertised to advertised AS Numbers. These two series will allow the projection of the total assigned AS numbers. Together with a model of the behavior of the RIR AS number pool, a complete AS number consumption model can be generated. The modelling uses the most recent 800 days as the baseline data for this exercise.

The next data series is the projection of the unadvertised AS count / advertised AS count ratio. This is shown in the following figure, with a linear best fit superimposed on the measured data.



Figure 19 -Unadvertised to Advertised AS Number Ratio

The next step is to model the advertised AS number count. The previous section looked at both linear and exponential trend models. This is shown in the figure below.




Figure 20 -Fit of Linear and Trend models to advertised data

The most recent data and the correlation to the linear and best fit models is shown in the figure below.




Figure 21 -Fit of Linear and Trend models to advertised data, recent data

The error between the original data series and the linear and exponential trend models is shown below. The closer the values in these series are to zero, the better the fit.



Figure 22 -Linear and Trend model error

A best fit to the data can be established by looking at the first order differential of the data. If the first order differential is constant, then a linear trend model would be the best fit to the data. This is shown in the following figure.

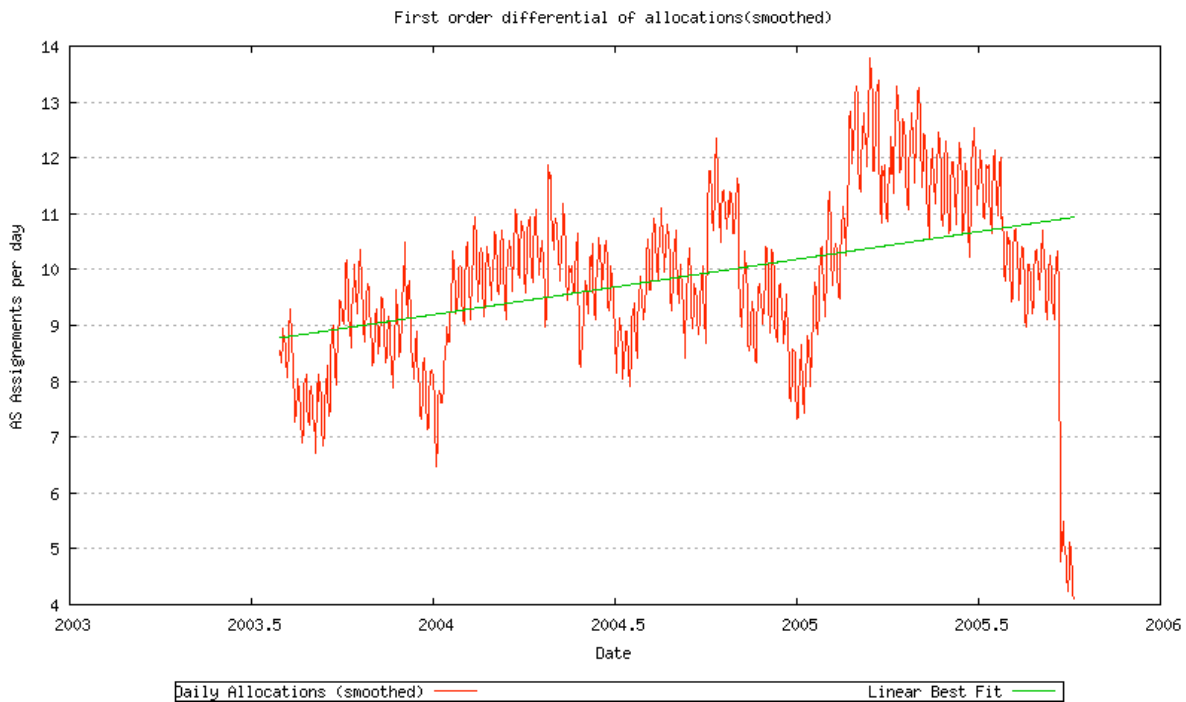



Figure 23 -First order differential of the AS advertisement rate


There is a rising trend in this data, indicating that an exponential model may be a better fit. Here the first order differential of the log of the data is shown in the following figure.





*Figure 24 – First order differential of the log of the AS advertisement rate*

Using this exponential trend model of advertised AS number growth and a linear trend model of the ratio of unadvertised to advertised AS numbers, it is now possible to model each RIR's projected AS allocation rate and then look at the predicted behaviour of the RIR pool size for each RIR. This is shown below.



*Figure 25 – Model of RIR AS Pool Behaviour*

Using this model of each RIR's relative rate of allocation, it is possible to generate a model of AS number consumption. Here the end point is the date where the first RIR has exhausted its available pool of 2-byte AS numbers, and no further numbers are available in the IANA unallocated pool to replenish the RIR's pool. The data available suggests a best fit predictive model where this will occur on October 2010.

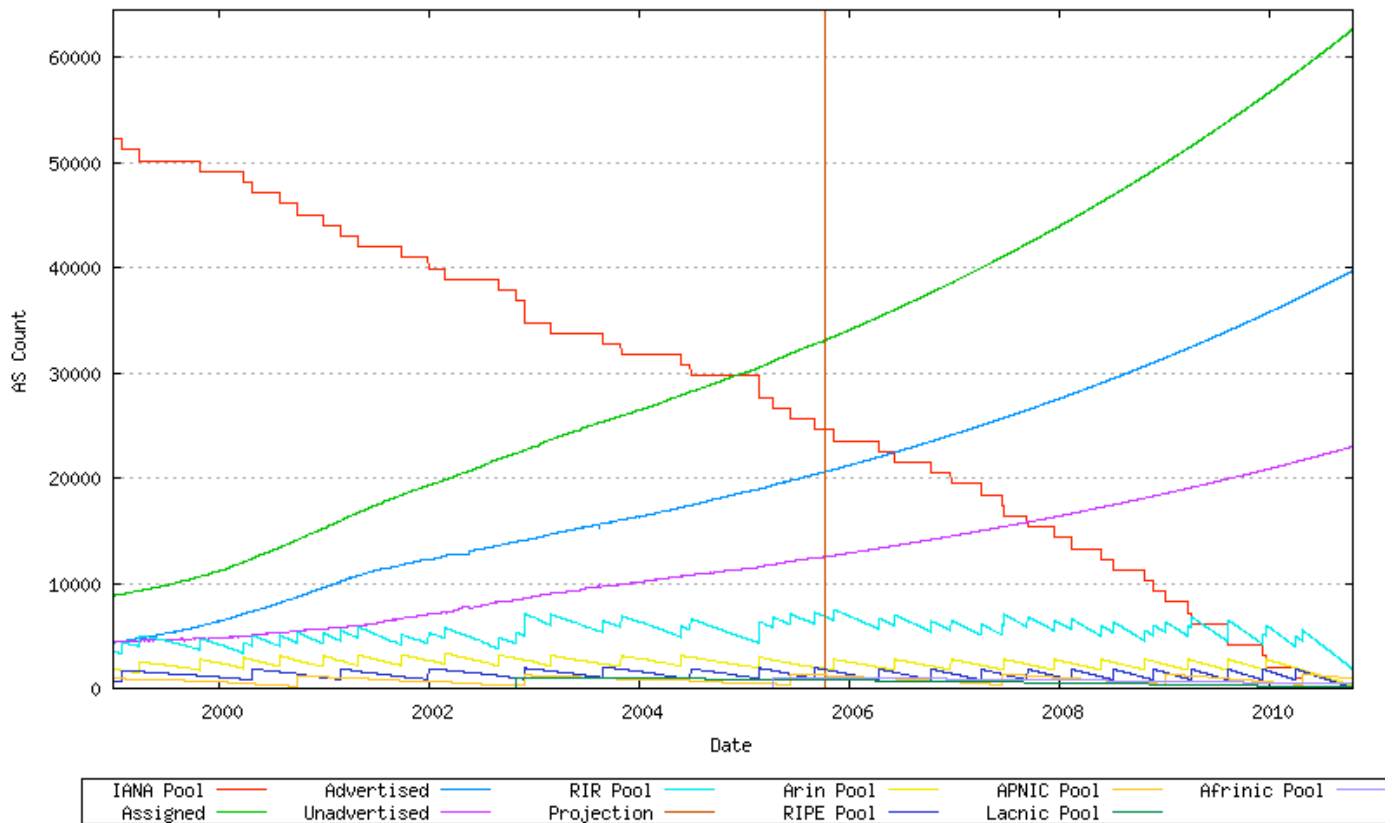


Figure 26 - Predictive Model of 2-Byte AS Number Consumption

So it would appear that we are looking at a best fit model of accelerating consumption of AS numbers, and a projected exhaustion date of late 2010 of available AS numbers to allocate to ISPs, or some 5 years from the time of writing this article.

## What does an AS Pool exhaustion date of 2010 actually mean?

The implication of these projections is that by late 2010 the Internet should be using a new protocol for inter-domain routing that does not rely on Autonomous Systems numbers, or, more likely, that the Internet should be using a version of BGP that supports the use of larger AS numbers that are drawn from a number pool significantly larger than 16 bits.

The first option appears to be somewhat unrealistic, to say the least.

And the second option, although simpler and very much the preferred path, is still going to take some years to deploy, particularly considering the growing size of the Internet's inter-domain space and the diversity of these component domains.

When contemplating a transition to a larger AS Number pool, it should be remembered that every day there are more networks that will need to undertake a transition to a longer AS number field in their deployed instances of the BGP protocol.

The steps in this transition path appear to include:

- the completion of the relevant protocol standards for a larger AS number field in BGP,
- the production of code in available implementations of BGP that support this protocol standard,
- various forms of testing of this code, both in terms of its correct operation and interoperability, and in terms of the correctness and viability of the relevant transition steps,

- 
- developing the necessary infrastructural support system to manage the distribution of this new number pool, and
  - a process of deployment of this protocol so that the deployment of larger AS numbers can commence well before the point at which the existing AS number pool is exhausted.

Even an aggressive schedule of transition across such a large and diverse network as is the Internet will take a number of years to reach the final step. It also appears that a prudent course of action would see us reach that position not by 2010, but by 2008 at the latest, allowing us a margin of some two years (and some 10,000 remaining AS numbers) to complete the task.

The implication is that we have reached the end of the period when the agenda for transition can be further deferred without undue risk. We now have some 3 years to undertake these tasks if we want to avoid an uncomfortable phase of potential disruption in the Internet's inter-domain routing space.

In the next part we will explore the aspects of the proposed 4-byte AS number space and the manner by which the Internet's transition across to this space can be performed.

This data analysis and associated predictive trend generation has been scripted, and the latest daily update of the data sets, and the associated generated predictive model is available at <http://www.potaroo.net/tools/asns/index.html>.

---

## Part 3 - 4-Byte AS Numbers

In the previous parts of this article on the topic of AS numbers we've looked at the aspects of the use of AS numbers in the Internet's inter-domain routing space, and then looked at the range of projections that are derived from the existing AS number consumption data. The most likely consumption trend that is derived from this data is a model of an increasing AS number consumption rate in the coming years.

This model has a best fit projection that forecasts exhaustion of the unallocated pool of AS numbers in late 2010. In working backward from this date to the necessary steps to facilitate the associated industry transition, it would appear that the time to commence the transition is in the coming months rather than in the coming years. To recap the steps involved in this transition, we appear to be looking at a sequence of actions that includes:

- the completion of the relevant protocol standards for a larger AS number field in BGP,
- the production of code in available implementations of BGP that support this protocol standard,
- various forms of testing of this code, both in terms of its correct operation and interoperability, and in terms of the correctness and viability of the relevant transition steps,
- developing the necessary infrastructural support system to manage the distribution of this new number pool, and
- a process of deployment of this protocol so that the deployment of larger AS numbers can commence well before the point at which the existing AS number pool is exhausted.

One of the very first steps is to look at what is being proposed to address this forecast exhaustion is standardize the use of a larger AS number pool within the BGP protocol, and also understand the implications of an associated transition plan.

In this part of the article we'll look at the current proposal for a larger AS number pool.

As of July 2005 the document defining this proposal is an IETF Internet Draft : draft-ietf-idr-as4bytes-10.txt.

The approach proposed in this work-in-progress document is to expand the size of the AS number pool space from 16 bits to 32 bits. In number terms this expands the number space from a pool of 65,536 numbers to 4,294,967,296 billion numbers. In terms of the current use of AS numbers, the current scaling properties of the BGP routing protocol, and the use of AS's in the context of inter-domain routing, a pool of 4.4 billion numbers would easily encompass a network environment of significantly greater levels of domains, and inter-domain interconnection density. Such a pool size would exceed some current guesses of the scaling capabilities of the BGP protocol by up to a further two orders of magnitude.

Its also proposed to preserve the first block of 4-Byte AS Numbers to align with the allocations of the 2-Byte numbers.

Let's use a new form of terminology here for 4-Byte AS number values, where the first 65,536 AS numbers are numbers use the form "0:0" through to "0:65535". The second set of 65536 numbers would be written as 1:0 through to 1:65535, and so on. So we'll be using a number format of <upper16 bits>:<lower 16 bits>.

What is the inventory of issues that need to specifically addressed here?

Obviously there is a need for some changes to the routing protocol, and this change needs to be able to accommodate a number of situations. An ordered inter-domain transition will be unrealistic to expect. More reasonable on the expectation scale is the piecemeal transition of domains, where individual domains will shift to supporting 4-byte AS's in their own time. Domains that are currently using 2-byte AS's may have less reason to undergo an early transition to 4-byte AS support, while those domains who are assigned a non-mappable 4-byte AS number will find that they have to support 4-byte AS numbers from the outset. The following diagram hopefully illustrates most of the potential issues associated with this kind of piecemeal transition. Here "OLD" refers to 2-byte AS number domains and "NEW" refers to 4-byte AS number domains.

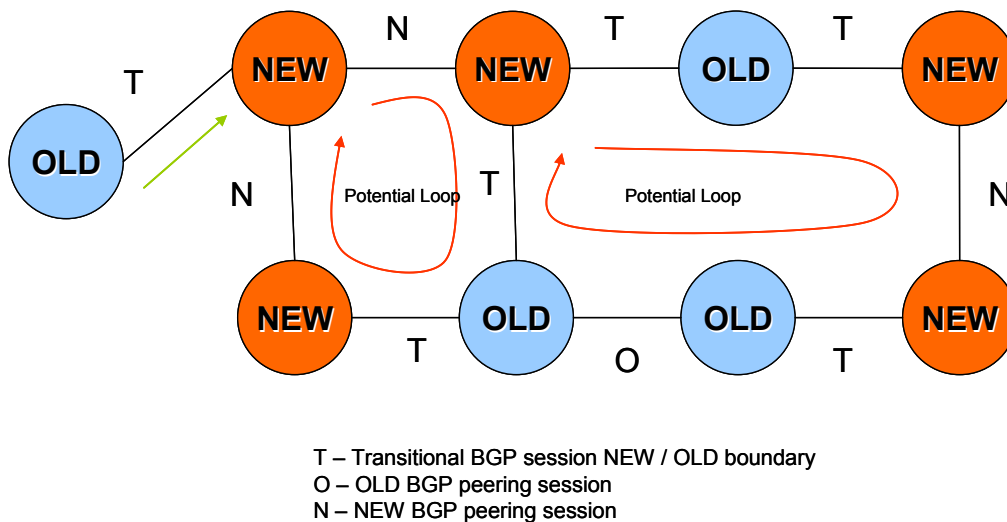


Figure 1 – BGP Transition Cases

Figure 1 describes an original prefix injection from an OLD BGP originating AS and a number of transitions from OLD to NEW BGP AS's and NEW to OLD BGP AS's, as well as OLD and NEW BGP peerings. The diagram also contains a number of potential loops, encompassing both OLD and NEW transits. Any proposed solution should be able to loop detect in this scenario without having to alter the behaviour of the OLD BGP speakers.

## Changes to the BGP protocol

BGP has two major parts within its protocol: opening up a BGP conversation with a peer BGP speaker, and then transfer of protocol objects that describe reachability of address prefixes and associated attributes of these address prefixes. Both parts include AS Number components, and in considering

---

changes to the current protocol, both parts of the protocol require some change. The message objects that need to be considered here are therefore the BGP OPEN message and the BGP UPDATE message.

The changes to the BGP protocol create a "NEW" BGP implementation that is capable of supporting a 4-byte AS number environment. The essential task of the changes is to define mechanisms that all NEW BGP speakers to speak to each other and pass all AS number values in 4 byte fields. However the Internet is way too large to set up a "flag day" at which point the entire collection of BGP speakers will undertake a switch from "OLD" BGP to NEW BGP. Accordingly, its also necessary to define protocol interactions in NEW BGP where the transition in the Internet will be gradual and essentially uncoordinated. NEW BGP speakers will have to set up sessions with OLD BGP speakers, and of course OLD BGP speakers will also be peering with other OLD BGP speakers. The information associated with 4-Byte AS paths must be passed across sections of the network that normally support only 2-Byte AS Paths. In other words 4-Byte AS information needs to be passed to OLD BGP speakers and between OLD BGP speakers.

## Opening a BGP session

BGP carries its own AS number in the "My Autonomous System" field of the BGP OPEN message.

The proposed approach is to initiate a NEW BGP session in a mode that is compatible with the OLD BGP protocol, and also inform the remote peer of its capability to conduct a NEW BGP conversation if the remote peer is also a NEW BGP speaker. This capability advertisement is part of OLD BGP, and OLD BGP speakers who open a peer session with a NEW BGP speaker will simply ignore the NEW capability and operate in OLD mode. A NEW BGP peer will respond positively to the NEW capability, and the BGP session can then operate in NEW mode.

The BGP OPEN message includes a fixed length 2-byte "My AS field" (as shown in Figure 2) as well as potentially containing a capability query as part of the Optional Parameters section. In order to ensure that NEW and OLD speakers can communicate, then this 2 byte My AS field needs to be preserved in NEW BGP even when the Optional Parameters section encompasses the capability to undertake a NEW peering session. This may appear contradictory in the first instance, as the OPEN message then contains both a 2-byte Autonomous System number and a 4-byte AS Capabilities Query.

The mechanism proposed for the OPEN Message varies according to whether the NEW speaker is using a mappable AS number drawn from the original pool (i.e. with a My AS number in the range 0:0 through to 0:65535), or its using a number drawn from a higher-numbered 4-byte number block. In the first case the OPEN message would use the 2-byte mapped value in the My AS field (dropping out the zero-valued high order 16 bits of the AS value), while in the second case the BGP speaker would use for My AS a special 2-byte value that is reserved for this purpose (AS 23456). In both cases the Optional Parameter section would include a capability code to indicate that the local BGP speaker can support 4-byte AS Numbers (Capability Code 65).

The side effect is that in the OLD BGP domains AS 23456 may appear to be connected to the 2-byte BGP realm in many different locations, and advertising a collection of different address prefixes in different locations. From the OLD BGP realm this does not present a protocol problem, although, as always, there is the potential here that this repeated use of AS 23456 as a 4-byte AS substitution token may create a somewhat confusing BGP-view of the Internet from the perspective of the OLD BGP world!




Figure 2- BGP Open Protocol Message – From "draft-ietf-idr-bgp4-26.txt"

The capability exchange uses a protocol described in RFC3392. The NEW BGP speaker adds an optional capability field to the OPEN message. The 4-byte AS capability code 65 carries as its capability value the local 4-byte local AS number value. For a NEW peer this capability value is to be interpreted as the actual AS of the remote side, on the basis that the My AS field in the body of the OPEN is either a truncation of the local 4-byte AS value (in the case of mappable 4-byte AS values), or the special value of AS 23456.

One response from the remote BGP speaker is to accept the capabilities announcement with a comparable OPEN message, in which case the remote side is also a NEW BGP speaker, and the session may proceed using 4-byte AS values.

If the session is being opened with an OLD BGP peer, the OLD BGP peer may respond with a NOTIFICATION message indicating that the 4-byte capability is an Unsupported Optional Capability parameter. In response to this unsupported notification the NEW BGP speaker will re-establish the connection by resending the OPEN message, and this time drop the 4-byte capability option from the message. The NEW BGP speaker will then assume that it is peering with an OLD BGP peer.

The "Unsupported" response to a capabilities parameter was not included in the original specification. Older versions of BGP allowed a BGP speaker to optionally send a NOTIFICATION message and terminate the peer session. If the NEW BGP speaker sees a session termination in response to its OPEN message it may need to re-open the TCP session, and this time omitting the 4-byte capability advertisement in the initial BGP OPEN message. Once again, the NEW BGP speaker will then assume that it is peering with an OLD BGP peer.

In general, however, a BGP implementation should not send a NOTIFICATION when a capability parameter is unrecognized because the Capabilities Optional Parameter is still optional. With such general implementations, the OLD speaker would just pick up the 2-byte AS (23456) in the OPEN received from the NEW speaker. As the OLD speaker does not advertise the 4-byte AS Capability in its OPEN, the NEW speaker has to use the 2-byte AS it advertised in the OPEN (that is, the AS\_TRAN - 23456) for peering. A NOTIFICATION is not involved in this scenario.

## The BGP UPDATE Message

For a NEW BGP session (4-byte peering with 4-byte) the changes to the protocol are the use of 4-byte AS numbers in the AS\_PATH attribute of UPDATE messages. All 2-byte AS values are padded with a zero high order 16 bits. If the AGGREGATOR attribute is used it is similarly carried as a 4-byte value. So in the 4-byte peering, all 2-byte information is carried in mapped 4-byte AS numbers.

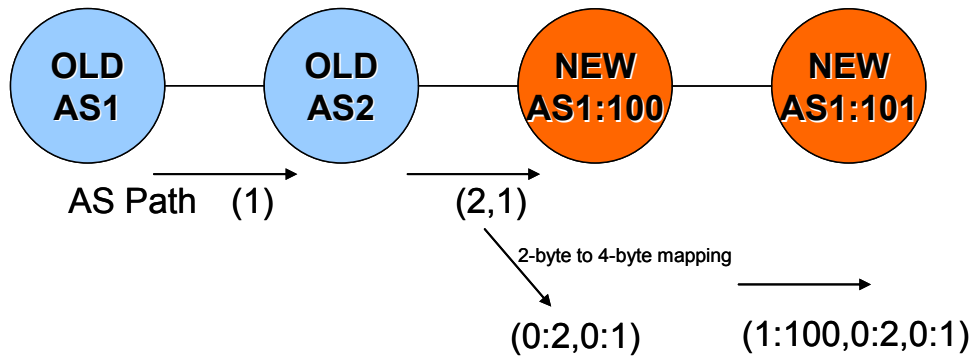


Figure 3 – OLD to NEW BGP AS Path Mapping

In this way AS Path length is preserved without change when translating 2-byte AS information into the 4-byte domain.

The next case is where an OLD BGP peers with a NEW BGP. We've already seen the simple case where the information is coming from a 2-byte path and there is no additional 4-byte information, and in this case the 2-byte values are simply mapped into 4-byte values. What about the reverse case where 4-byte information is being passed back into the 2-byte world?

There are two parts to this case: firstly creating an equivalent 2-byte AS Path and secondly packing up the 4-byte AS Path information in such a way that it transits across the 2-byte domain in such a way that it can be reassembled in any subsequent transition into a 4-byte domain. In the first case the equivalent path information is constructed by either stripping off the high order 2-bytes off the AS value, as long as this part is all zeros. Where this is not possible the transition AS number, 23456, is substituted in its place. This is indicated in Figure 4.

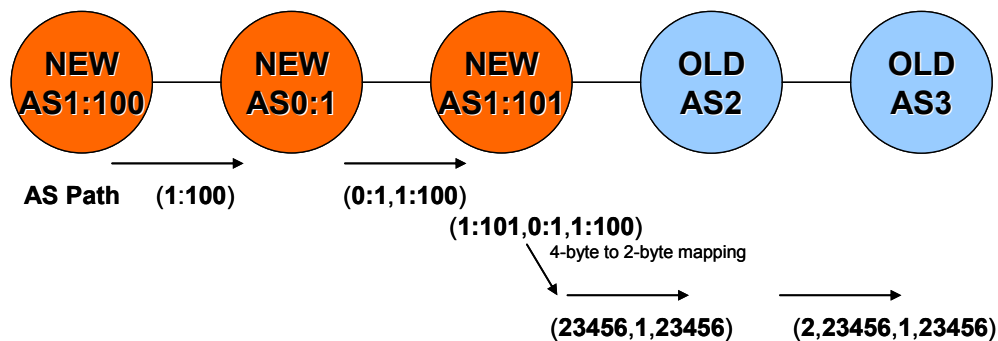


Figure 4 – NEW to OLD BGP AS Path Mapping

In this way the AS Path length metric is preserved, and the prevention of count-to-infinity loops in the 2-byte domain is avoided.



The second part to this case is packaging up the 4-byte path into the OLD BGP session in such a way that it can be unpacked at any subsequent boundary into a 4-byte realm. Here the proposal calls for new transitive community attributes to be supported for OLD BGP. These attributes are defined as transitive attributes, and should be passed through the OLD BGP peering sessions without alteration. It should be noted that this is not a protocol change per se, but it does require the explicit support within OLD BGP implementations of this attribute as a transitive community. The proposed mechanism is an extended community attribute called "NEW\_AS\_PATH". When a NEW BGP speaker is speaking to an OLD BGP, the NEW BGP prepends its own AS value to the AS PATH and copies this information into the NEW\_AS\_PATH. It then translates the 4-byte AS Path into a 2-byte equivalent AS Path. The translation is straightforward, in that where the 4-byte AS has all zeros in the high order 2 bytes, the translation truncates the AS value to a 2-byte value, and where the high order 2-bytes are non-zero the translation substitutes the reserved 2-byte value AS 23456 in its place.

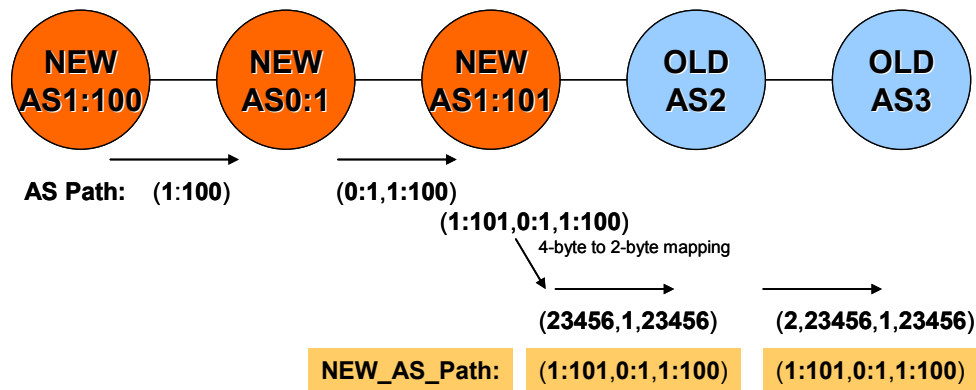


Figure 5 – NEW to OLD BGP AS Path Mapping

The transit across the OLD BGP domains leaves the NEW\_AS\_PATH untouched, and prepends 2-byte AS values to the AS\_PATH.

The next transition is one from the OLD to the NEW domain, as shown by a continuation of the previous example

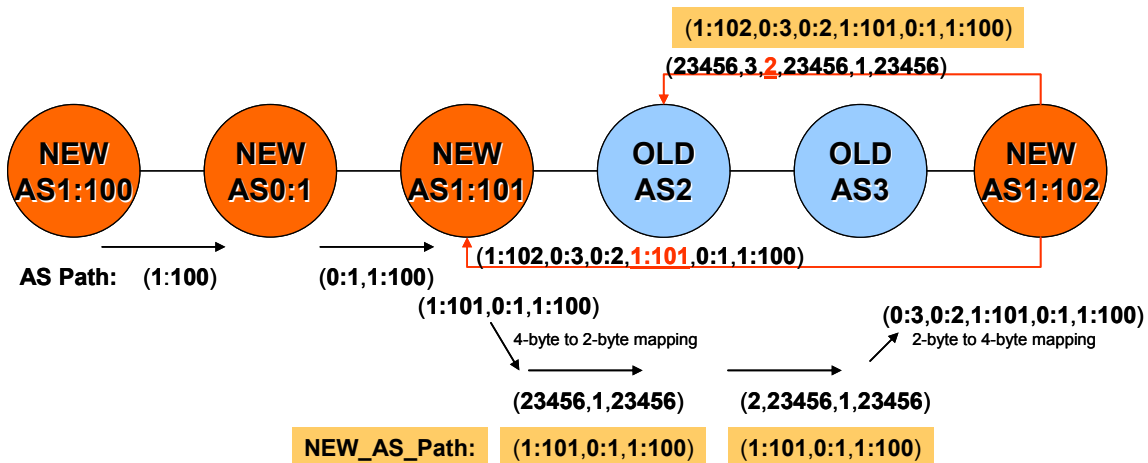


Figure 6 – NEW-> OLD-> NEW Transition with Potential Routing Loops

The Figure 6 there is a further OLD to NEW transition. In this case the NEW BGP speaker takes the AS Path as presented by the OLD BGP speaker and converts the 2-byte values to 4-byte values by adding 2-bytes of zero padding to each entry, and then overwrites the trailing entries with the values specified by the NEW\_AS\_PATH attribute. The net result is that the 4-byte path that entered the 2-byte sequence is prepended with the 2-byte transit AS sequence. The NEW\_AS\_PATH is then removed, leaving an intact 4-byte path as the AS\_PATH attribute.

This ensures that the resultant BGP environment can detect loops in both the NEW 4-byte and OLD 2-byte realms.

Further extending this example, we can construct a potential loop in the 4-byte world by adding a path back to AS 1:101. The restoration of the original 4-byte AS Path at the OLD-to-NEW transition ensures that the potential loop is discarded even when the loop needs to traverse one or more 2-byte OLD BGP AS's. A similar form of loop can be constructed for a 2-byte OLD BGP AS, that traverses a 4-byte NEW BGP AS. Again the transition mapping ensures that the potential routing loop is detected by BGP.

The ability to perform AS Path Prepending is also unaltered in this mixed NEW and OLD BGP environment, The AS simply prepends its local AS value to the AS\_PATH as normal. In the case of prepending on a NEW-to-OLD boundary the prepended AS Path is mapped into the NEW\_AS\_Path attribute as described above.

The earlier article of use of AS numbers in routing also noted the less common use of AS PATH poisoning, where the prepending uses a different AS number value in order to ensure that the particular advertisement is not learned by a remote AS. For NEW BGP speakers there is no change to this capability. For OLD BGP speakers the AS Path Poisoning can only be directed towards 2-byte AS's, as the OLD BGP speaker has no knowledge of the structure or content of the NEW AS\_PATH attribute. From the perspective of the OLD BGP speaker, the NEW\_AS\_PATH attribute is an opaque data block.

The same translation technique applies to the AGGREGATOR attribute. In a NEW-to-OLD transition the AGGREGATOR may be a mappable AS number, in which case the value is truncated to 2-bytes and no further action is required. Otherwise the 4-byte AGGREGATOR value is rewritten to the NEW\_AGGREGATOR attribute and the transition 2-byte value, AS 2356 is placed into the AGGREGATOR attribute. On an OLD-to-NEW transition the NEW\_AGGREGATOR attribute is copied back into the AGGREGATOR attribute, if defined, otherwise the AGGREGATOR is padded out with leading zeros.

The general approach adopted for transition is preserve AS Path length information across the OLD and NEW BGP boundaries, while recognising that some 4-byte AS information cannot be cleanly mapped into a 2-byte AS Path. In order to preserve 4-byte information, which is necessary to prevent loop formation for 4-byte AS's, the 4-byte information is preserved across OLD transit paths and restored upon re-entry into NEW BGP realms.




Figure 7 – 2-byte and 4-byte AS Realms

---

## BGP Communities

BGP communities require some additional consideration. If the high order 16 bits of the community attribute are neither all zeros or all ones, then it is assumed to contain a 2-byte AS value. Where it is necessary to specify a 4-byte AS number in the community attribute it is necessary to turn to the extended community attribute to support this.

This extended communities feature is documented in the Internet draft: draft-ramachandra-bgp-ext-communities-10.txt, now on the RFC publication track as a Proposed Standard.

## Transition

Transition in this environment is relatively straightforward. NEW BGP speakers can be deployed within the network in a piecemeal fashion without any major concerns. The size of BGP UPDATE messages is slightly longer due to the extended length of the AS PATH attribute in NEW BGP and the NEW\_AS\_PATH attribute that has been added in the OLD BGP environment, but it should not prove to be a major factor.

BGP loop prevention appears to be adequately addressed in all commonly encountered situations, and there appear to be no other significant transition considerations.

There does appear to be one precondition for the use of 4-byte AS numbers, and that is for a routing domain to actually be numbered with a non-mappable 4-byte AS number, all the BGP speakers in the domain should be NEW BGP speakers. Aside from that consideration there do not appear to be any further constraints associated with this transition.

## Alternative Approaches

It is certainly a challenging task to contemplate an environment when we would exhaust a 4-byte AS number space, but one would suppose that same consideration was in the minds of the original BGP protocol designers when they opted to use 2-byte AS numbers. Of course a 32 bit number pool is not double the pool size of a 16 bit number – its 65,536 times larger. That does appear to lead one to believe that this time it will be far more challenging task to exhaust this number pool.

This approach of simply extending the number space appears to offer a path of minimal disruption and minimal change in terms of operational configuration, storage, message size and processing overheads for BGP. Nothing much has changed here except the range of the number space, and some ancillary considerations relating to transitional arrangements.

Of course, other labelling spaces remain possibilities, and a shift to a different labelling scheme could well use the same transitional approach. There is no significance whatsoever in the AS Number apart from its uniqueness, and any other form of namespace would function equally well in terms of its role in BGP. One could use strings such as domain names, URIs, fixed length hashes of public keys, the public keys themselves, or even IPv6 addresses as a distinguishing AS identifier. There is no direct requirement for summarization of AS number ranges within the protocol use, no requirement within the protocol to continue to use number identifiers, and no direct requirement to stick with values that are encoded in a fixed length field.

---

However, such approaches would add to the size of BGP UPDATE messages, increase the storage requirements, and, perhaps marginally, increase processing overheads for BGP. The more complex the identity space the more complex the basic task of BGP configuration and the higher the possibility of mistakes. "Borrowing" AS identifiers from another namespace, such as domain names, or derived URIs, has the associated issue that the uniqueness of the space is derived from the inherent stability and uniqueness of the name space upon which the identifiers are derived. It's a definite possibility that at times this trust is misplaced.

Numbers are certainly simple neutral identifiers. The decision to simply extend the number pool space appears to be another instance of a design trade-off between the size of the number space and the additional BGP overheads and ease of transition. In sticking with numeric labels this approach represents minimal change to the installed base of BGP speakers, and there is no requirement for an existing routing domain using a 2-byte AS number and OLD BGP to make any changes to its routing environment at all.

The 4-byte transition appears to offer flexibility, orderly transition and minimal disruptions to existing operational practices.

We are certainly running out of the 2-byte AS number pool, and an industry of this size needs to have a considerable period of advance warning of change in order to be able to integrate such changes into various operational cycles of testing and transitional deployment prior to integration into production environments. The first steps that need to happen are the completion of the technical specification of this approach in the form of an IETF Standard and the subsequent production and distribution of NEW BGP implementations from the existing sources of BGP implementations. It would be preferable to get this underway now, while there is still time to complete this transition well before we exhaust the current 2-Byte AS Number space.

---

## Acknowledgement

Thanks to Enke Chen, one of the authors of the 4-Byte AS working document, for some clarification regarding the OPEN behaviour between OLD and NEW BGP implementations.

## Disclaimer

The views expressed are the author's and not those of APNIC, unless APNIC is specifically identified as the author of the communication. APNIC will not be legally responsible in contract, tort or otherwise for any statement made in this publication.

## About the Author

*Geoff Huston* B.Sc., M.Sc., has been closely involved with the development of the Internet for many years, particularly within Australia, where he was responsible for the initial build of the Internet within the Australian academic and research sector. He is author of a number of Internet-related books, and has been active in the Internet Engineering Task Force for many years.

[www.potaroo.net](http://www.potaroo.net)