Internet Society

The ISP Column
An occasional column on things Internet

August 2005

Exploring Autonomous System Numbers

Geoff Huston

So what's an "Autonomous System Number", and what role do these numbers play in the technology of the internet?

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 of 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 maintaining a comprehensive view of the current topology of a network 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 routing technology, nor does it make much sense for fine-grained topology information within one part of the Internet to be passed to all other parts of the network.

To partition the extent to which this fine-grained topology information is propagated across the Internet, the approach used within the Internet's routing architecture is to call upon a second level of routing hierarchy. This "inter-domain" routing environment maintains a somewhat different map of the network. This inter-domain routing environment supports a description of how each of the 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 a set of three articles 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 clearest definitions of an Autonomous System can be found in an IETF working document that describes the forthcoming new IETF Standard specification of 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 tohave a single coherent interior routing plan and
presents a consistent picture of what destinations are reachable through it.

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 are AS Numbers 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. Ast 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**). In general the AS Path reflects the sequence of AS's through which the prefix advertisement has traversed to reach the current AS. And the general intention is that the AS Path reflects the sequence of transit AS's that a packet will traverse to reach the destination prefix.

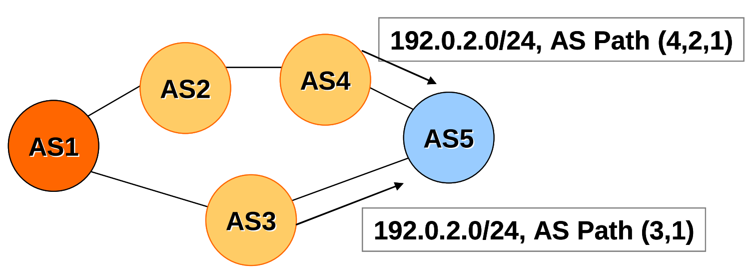

Figure 1 – AS Path Generation in BGP

Holding AS numbers sequences in AS Paths in the BGP protocol 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, 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 initially appear 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 RIP or EIGRP. The issue with all distance vector protocols, including BGP, is the "count to infinity" problem.

Lets 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. Both AS3 and AS4 will select a best path of metric 2, corresponding to the AS path (2, 1). If the connection between AS1 and AS2 is broken AS2 will 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. According 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 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.

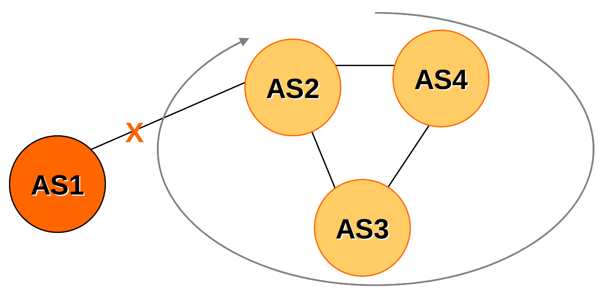

Figure 2 – Loop Formation in Distance Vector Protocols

By replacing the AS transit counter with a full path metric, this form of loop can be averted. 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 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 the 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 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 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 would appear to the rest of the Internet as the originating AS.

In the case where a network has two or more upstream transit connections then it is more likely that the network will use an 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, then the network needs to be assigned a unique AS number.

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.

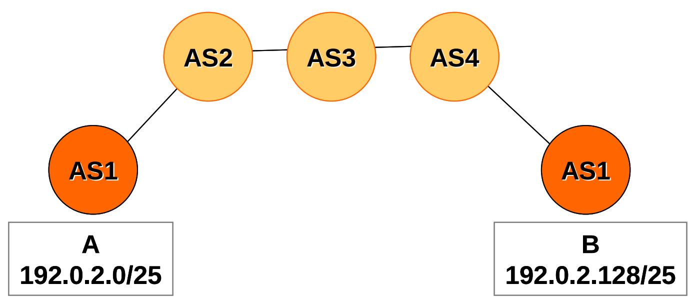

Figure 3 – Split AS

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 (2,3,4,1) 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 primariy 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 4, 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.

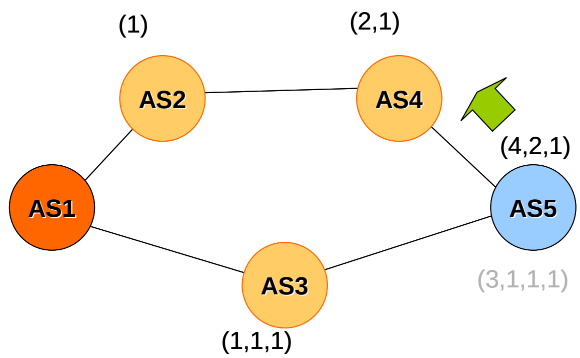

Figure 4 – AS Path Prepending

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 reach AS1.

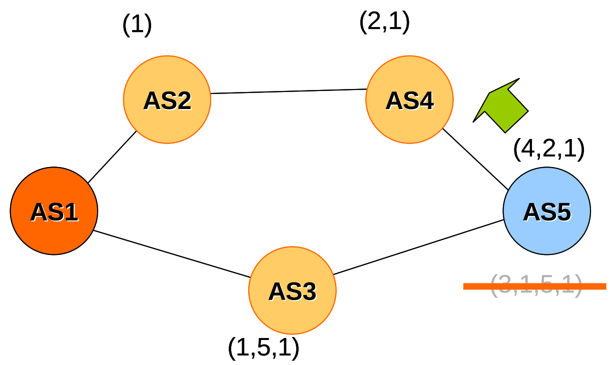

Figure 5 – AS Path Prepending with Poison AS

Summary

AS numbers are a basic component of inter-domain routing in the Internet. They form the means of identifying the distinct lower level units in the two-level hierarchy of the Internet's routing architecture. AS numbers are used to uniquely identify each domain.

The AS numbers are drawn from a 16 bit field, and there are 64,510 useable AS numbers.

As of July 2005 we've already consumed some 39,000 of these AS numbers.

- Are we running out of AS numbers?
- How long before we exhaust this number pool,?
- What are our options to expand this number space to encompass a much larger domain population of the future Internet?

We'll explore these questions in some detail in the next parts of this look at AS numbers in the Internet.

Disclaimer

If were are any views expressed in this article, they do not necessarily represent the views or positions of the Asia Pacific Network Information Centre, nor those of the Internet Society.

About the Author

GEOFF HUSTON holds a B.Sc. and a M.Sc. from the Australian National University. He 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 is currently the Senior Internet Researcher at APNIC, the Regional Internet Registry serving the Asia Pacific region. He has served as a Trustee of the Internet Society, and also as a member of the Internet Architecture Board of the Internet Engineering Task Force.