

4-Byte AS Numbers



The view from the old BGP world

Geoff Huston
October 2006
APNIC



4 Byte AS Numbers

- We are running into the exhaustion of the 2 Byte AS Number pool
 - Current estimate: 2200 UTC 14 October 2010
 - See <http://www.potaroo.net/tools/asns>
- From **1 January 2007** some / all of the RIRs will be allocating 4 Byte AS numbers upon specific request
- From **1 January 2009** some/ all of the RIRs will be allocating 4 Byte AS numbers by default



The 4-Byte ASN Approach

- Objectives
 - Change as little as possible in the BGP spec
 - Be 'backward compatible' with 2-Byte BGP implementations
 - Preserve AS semantics
 - Preserve loop detection capability
 - Preserve AS Path length metric
 - No 'flag day'
 - Allow 2-Byte implementations to continue to operate indefinitely in a mixed 2 / 4-Byte AS world



What does this imply?

- If you are a 2 Byte AS
 - as all of you are today

and you don't want to upgrade all your instances of BGP

- something you probably want to avoid

then you don't have to do anything at all!



Well, **almost** nothing

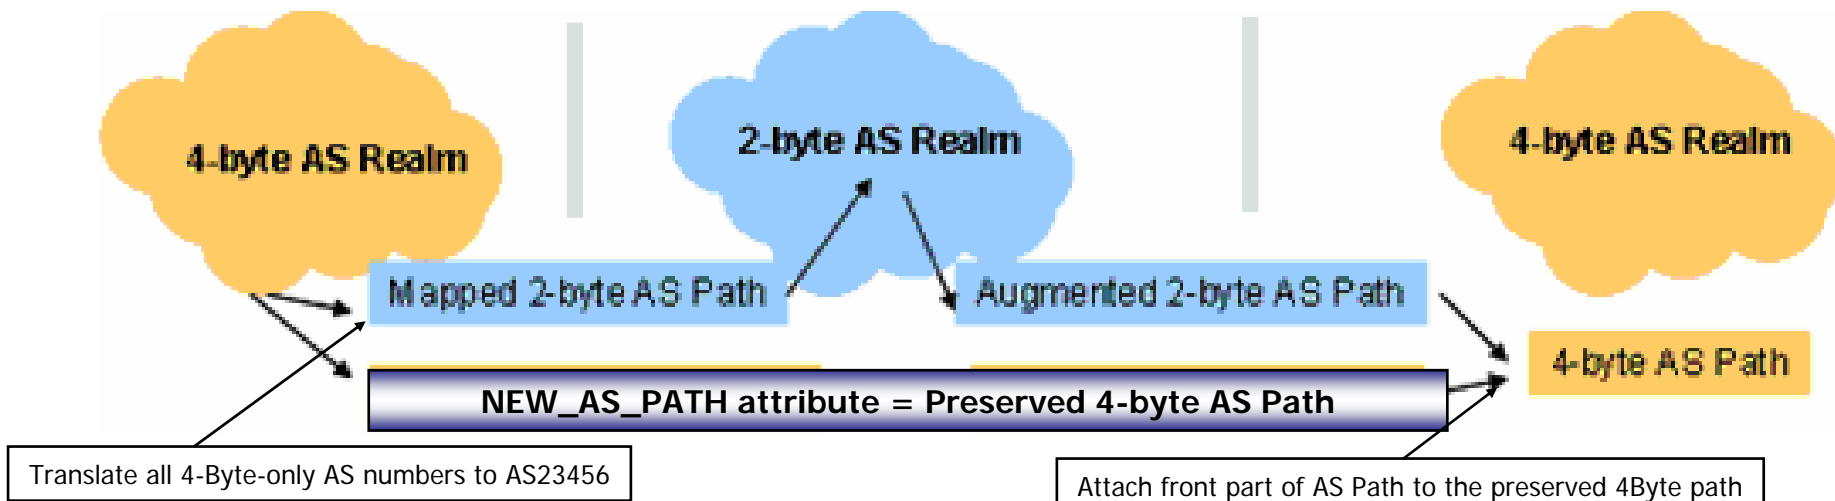


AS Path Semantics

- It's a path metric where the LENGTH is used as in path selection
- It's a loop detector where the presence of your own AS in a PATH is an indicator of a distance-vector-I'm-counting-to-infinity-unless-you-stop-me-loop
- You don't have to have an entirely accurate AS path – but at a minimum you do have to have a path-metric and loop-detecting properties

4-Byte AS Transition

- Think about this space as a set of NEW / OLD boundaries
- Define the NEW / OLD and the OLD / NEW transitions
- Preserve all BGP information at the transition interfaces
 - Translate 4-Byte AS Path information into a 2-Byte representation
 - Tunnel 4-Byte AS Path information through 2-Byte AS domain as an update attribute





AS Path approach

- In the 2-byte world we 'lie' about the 4-byte path
 - 4-Byte ASs appear as AS 23456 in the 2-Byte world
 - As long as you preserve the integrity of Path length and don't change 2-byte values in the 2 byte world, then BGP "works" in terms of path metric and loop detection
- In the 4-byte world we preserve 4-byte values of the entire AS Path

- Length integrity is preserved in all cases
- Loops entirely in the 2-Byte world are detected
- Loops entirely in the 4-Byte world are detected
- In a mixed 2-byte 4-byte potential loop then make sure that at least one party in the loop can see its own AS



Implications for 2-Byte BGP

- BGP speakers in 2-Byte AS domains should support NEW_AS_PATH as a transitive opaque community attribute
 - because that's where the 4-byte path is hiding
 - That's a "SHOULD" not a "MUST", by the way
 - Its better if you do, but nothing fatally breaks if you don't
 - Mixed 2 / 4 Byte loops will get detected in the 2-Byte world as a fallback



Implications for 2-Byte BGP

- AS 23456 is going to appear in 2-Byte AS paths – both origin and transit
- This implies that what's in the route registries and what your customers tell you about their AS and what's in your OSS and your routing system will differ:
 - E.g.: AS 1.2 gets translated into AS 23456 in a number of places, including in your OSS
 - You may need to peer with AS 23456, transit across AS 23456, and have multiple customers on AS 23456. Your OSS might get terminally confused



Implications

- If you want to explicitly signal to a 4-Byte AS using communities you will need to explicitly signal the 4-Byte AS using BGP Extended Communities

See:

- [RFC4630](#)
- [draft-rekhter-as4octet-ext-community-01.txt](#)



Implications

- BGP memory requirements will increase
 - 4-Byte BGP speakers will need twice the memory used to hold AS paths¹
 - 2-Byte BGP speakers will need up to three times the memory used to hold AS paths plus NEW_AS_PATH extended community attribute²

1 - Not "twice the memory" but "twice the memory used for AS Path storage"

2 - Not "three times the memory", but "three times the memory used for AS Path Storage"



Implications

- BGP bandwidth requirements will increase
 - 4-Byte BGP speakers will need twice the size used to carry AS paths
 - 2-Byte BGP speakers will need up to three times the size used to carry AS paths (factoring in the NEW_AS_PATH extended community attribute)



Implications

- BGP convergence times may increase in some cases
 - As any instance of 2-Byte BGP world destruction of the tunnelled NEW_AS_PATH attribute implies extended times on loop detection in order to fully complete prefix withdrawal
 - Its not that the withdrawal will loop forever, its that the loop will take some additional AS hops before it is detected



Implications

- If you proxy aggregate in the 2-Byte world then make sure that the aggregate is strictly larger than the components
 - Or loop detection may be harder than otherwise
 - As the AS Set object generated in the 2-Byte world as a result of this proxy aggregation is not cleanly translatable into the 4-Byte world



Implications

- No dynamic capability for 2/4-Byte ASN support
- You cannot flick from “2-Byte OLD” to “4-Byte NEW” mode within an active BGP session on the fly
 - A single BGP speaker could, in theory, simultaneously be a NEW and an OLD speaker in different sessions, but this is not required in the specification



Implications

- In a complex iBGP AS that wants to transition to using a 4-Byte “home” AS then you are going to have to think about the transition VERY carefully



The AS Number Report:

<http://www.potaroo.net/tools/asns/>

Paper:

<http://www.potaroo.net/ispcol/2005-08/>