

4 Byte ASN with Cisco IOS Software



4 Byte AS

- RFC 4271 defines an AS number as 2-bytes
- Private AS Numbers = 64512 through 65535
- Public AS Numbers = 1 through 64511
 - 39000+ have already been allocated
 - We will eventually run out of AS numbers
- Need to expand AS size from 2-bytes to 4-bytes
 - 4,294,967,295 AS numbers
 - Cannot have a “flag day” solution
 - On Jan 1, 2010 - all BGP speakers must support feature FOO
 - Solution must support a gradual deployment

4 Byte AS

- RFC4893 – “BGP Support for Four-octet AS Number Space”
Provides 4-byte AS support without a flag day
- RFC5396 – “Textual Representation of Autonomous System (AS) Numbers “

ASDOT

- Representation is based upon the existing 2-Byte AS representation
 - The full binary 4-byte AS number is split two words of 16 bits each
 - Notation:
<higher2bytes in decimal>.<lower2bytes in decimal>
For example: AS 65546 is represented as “1.10”
 - Easy to read, however hard for regular expressions
- Note: If the higher order 16 bits represent the value of a decimal zero, then the 4-Byte AS can be represented in as the traditionally well known 2-Byte AS format

ASPLAIN

- IETF preferred notation
- Continuation on how a 2-Byte AS number has been represented historically
- Notation: The 32 bit binary AS number is translated into a Single decimal value
Example: AS 65546

4-byte AS

- 4-byte AS support is advertised via BGP capability negotiation

Speakers who support 4-byte AS are known as NEW BGP speakers

Those who do not are known as OLD BGP speakers

- New Reserved AS#

AS_TRANS = AS #23456

2-byte placeholder for a 4-byte AS number

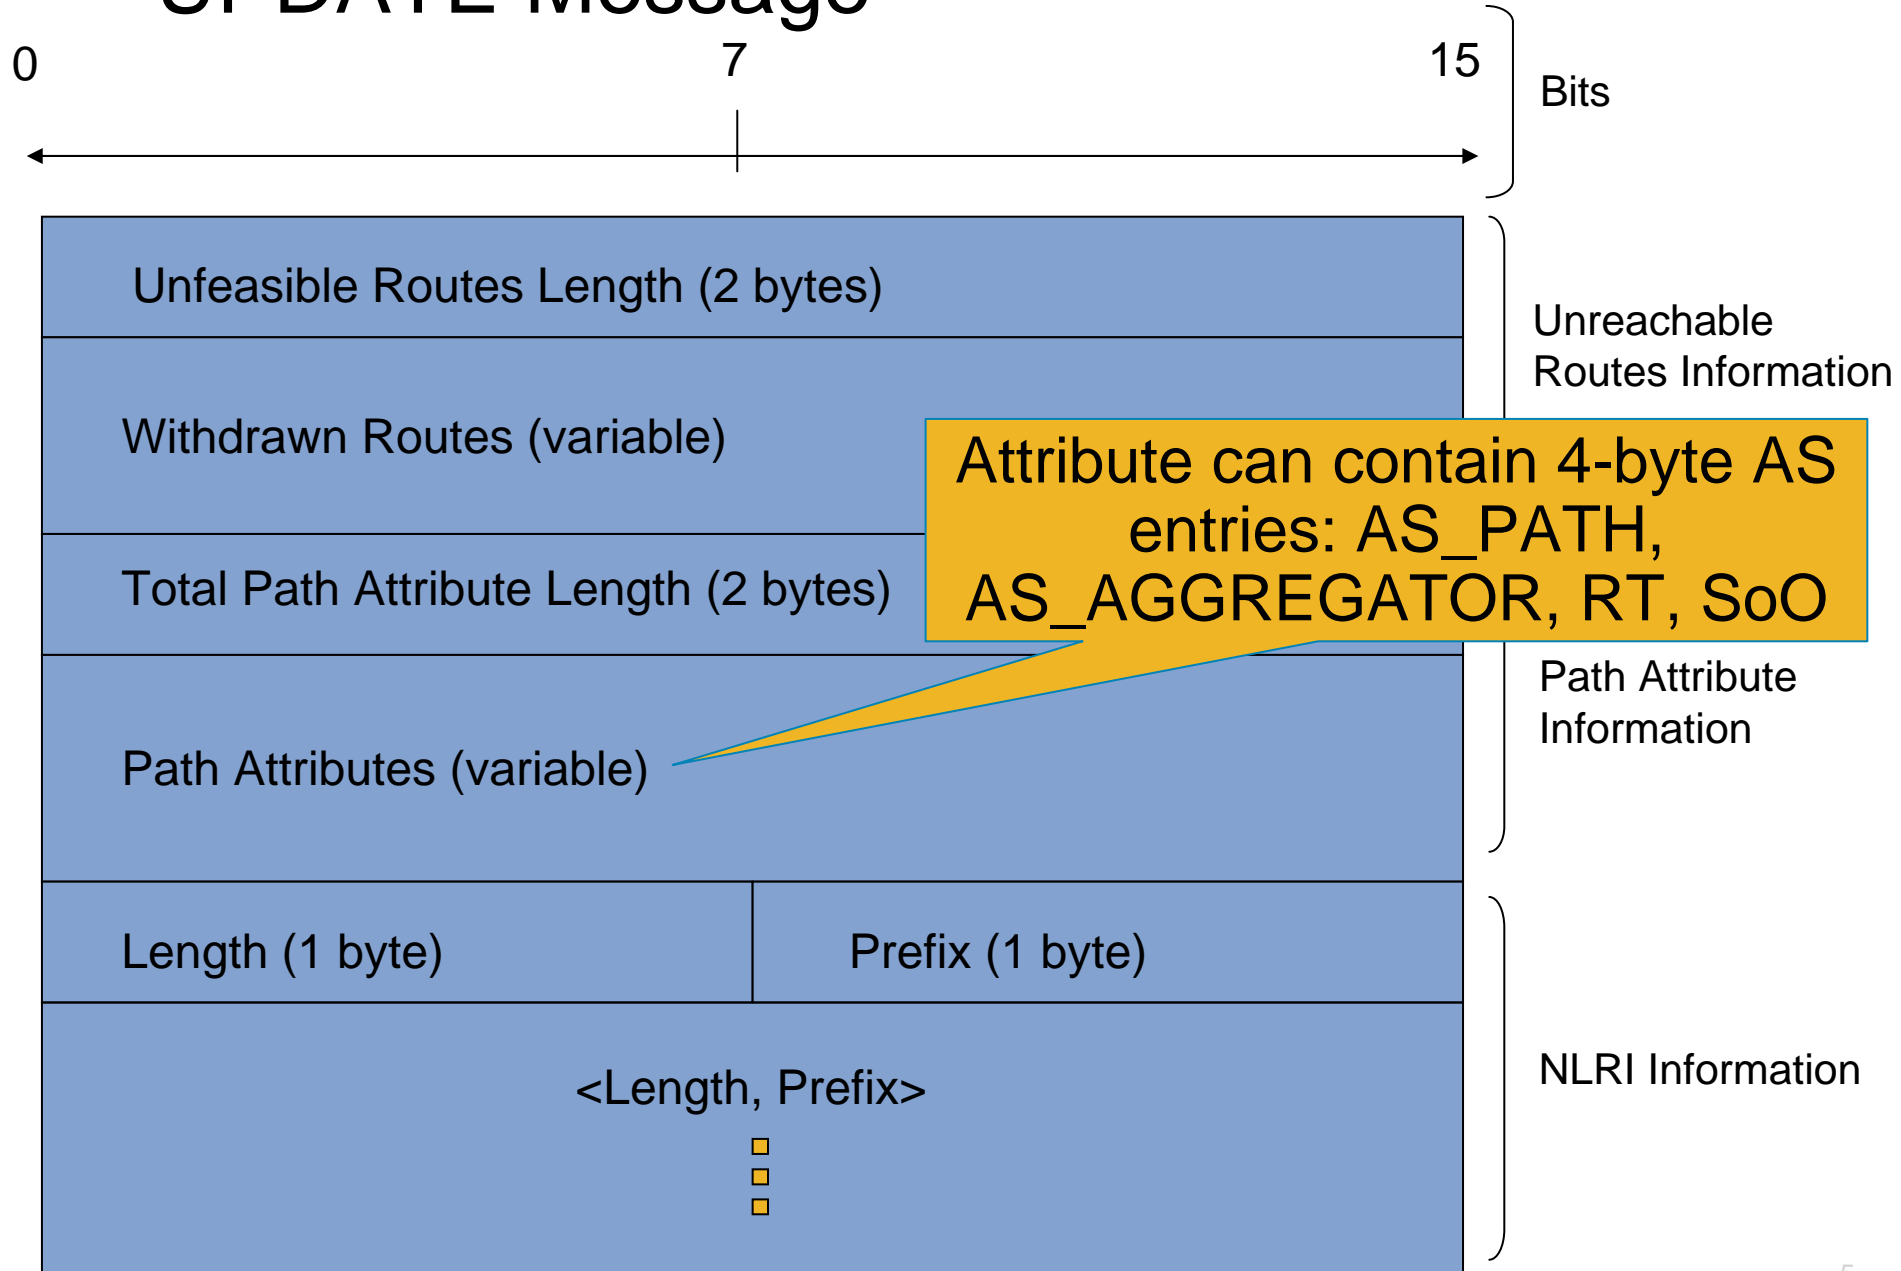
Used for backward compatibility between OLD and NEW BGP speakers

- Two new attributes, both are “optional transitive”

AS4_AGGREGATOR

AS4_PATH

UPDATE Message



4 Byte AS

- Formatting UPDATES for a NEW speaker

Encode each AS number within the AS_PATH in 4-bytes

AS_PATH and AGGREGATOR attributes are affected

For VPN Route-Target (RT) and Site-of-Origin (SoO) are affected also

- Formatting UPDATES for an OLD speaker

If the AGGREGATOR/ASPATH does not contain a non-2-byte mappable 4-byte AS we are fine

If it does, substitute AS_TRANS (AS #23456) for each 4-byte AS

AS4_AGGREGATOR and/or AS4_ASPATH will contain a 4-byte encoded copy of the attribute if needed

OLD speaker will blindly pass along NEW_AGGREGATOR and NEW_ASPATH attributes

4 Byte AS

- Receiving UPDATES from a NEW speaker
Decode each AS number as 4-bytes
AS_PATH and AGGREGATOR are effected
- Receiving UPDATES from an OLD speaker
AS4_AGGREGATOR will override AGGREGATOR
AS4_PATH and ASPATH must be merged to form the correct as-path
- Merging NEW_ASPATH and ASPATH

```
ASPATH –          275  250  225  23456  23456  200  23456  175
NEW_ASPATH –                100.1  100.2  200  100.3  175
Merged as-path – 275  250  225  100.1  100.2  200  100.3  175
```

Mappable Autonomous System Numbers

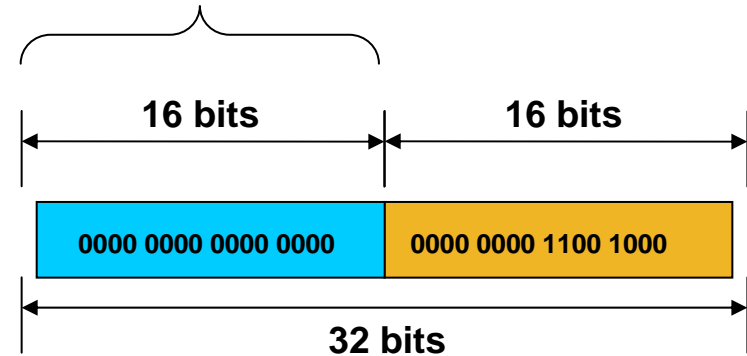

2 Byte Autonomous System

Mappable AS

4 Byte Autonomous System



Autonomous System # 200

All "ZERO"




Autonomous System # 50000

Translate in Binary

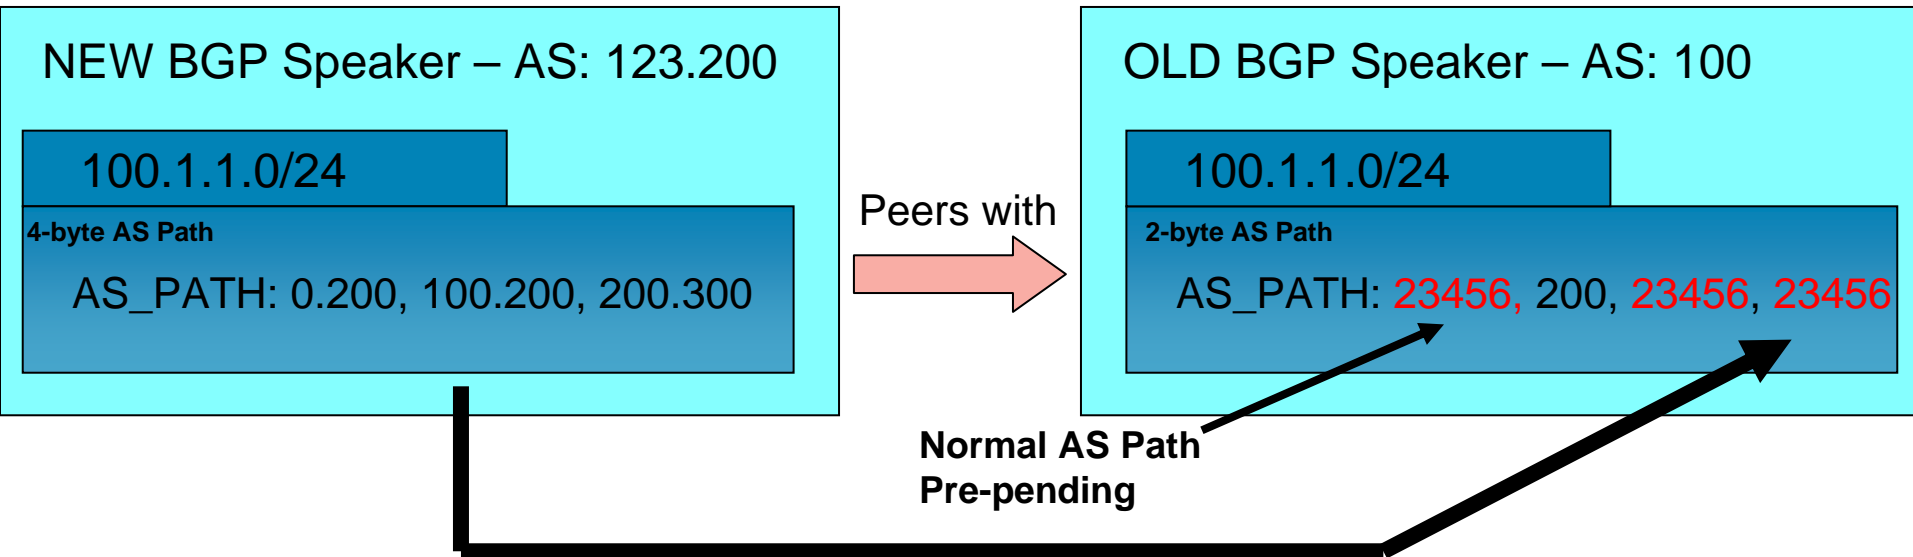


Backward Compatibility Mappable AS Numbers



Backward Compatibility

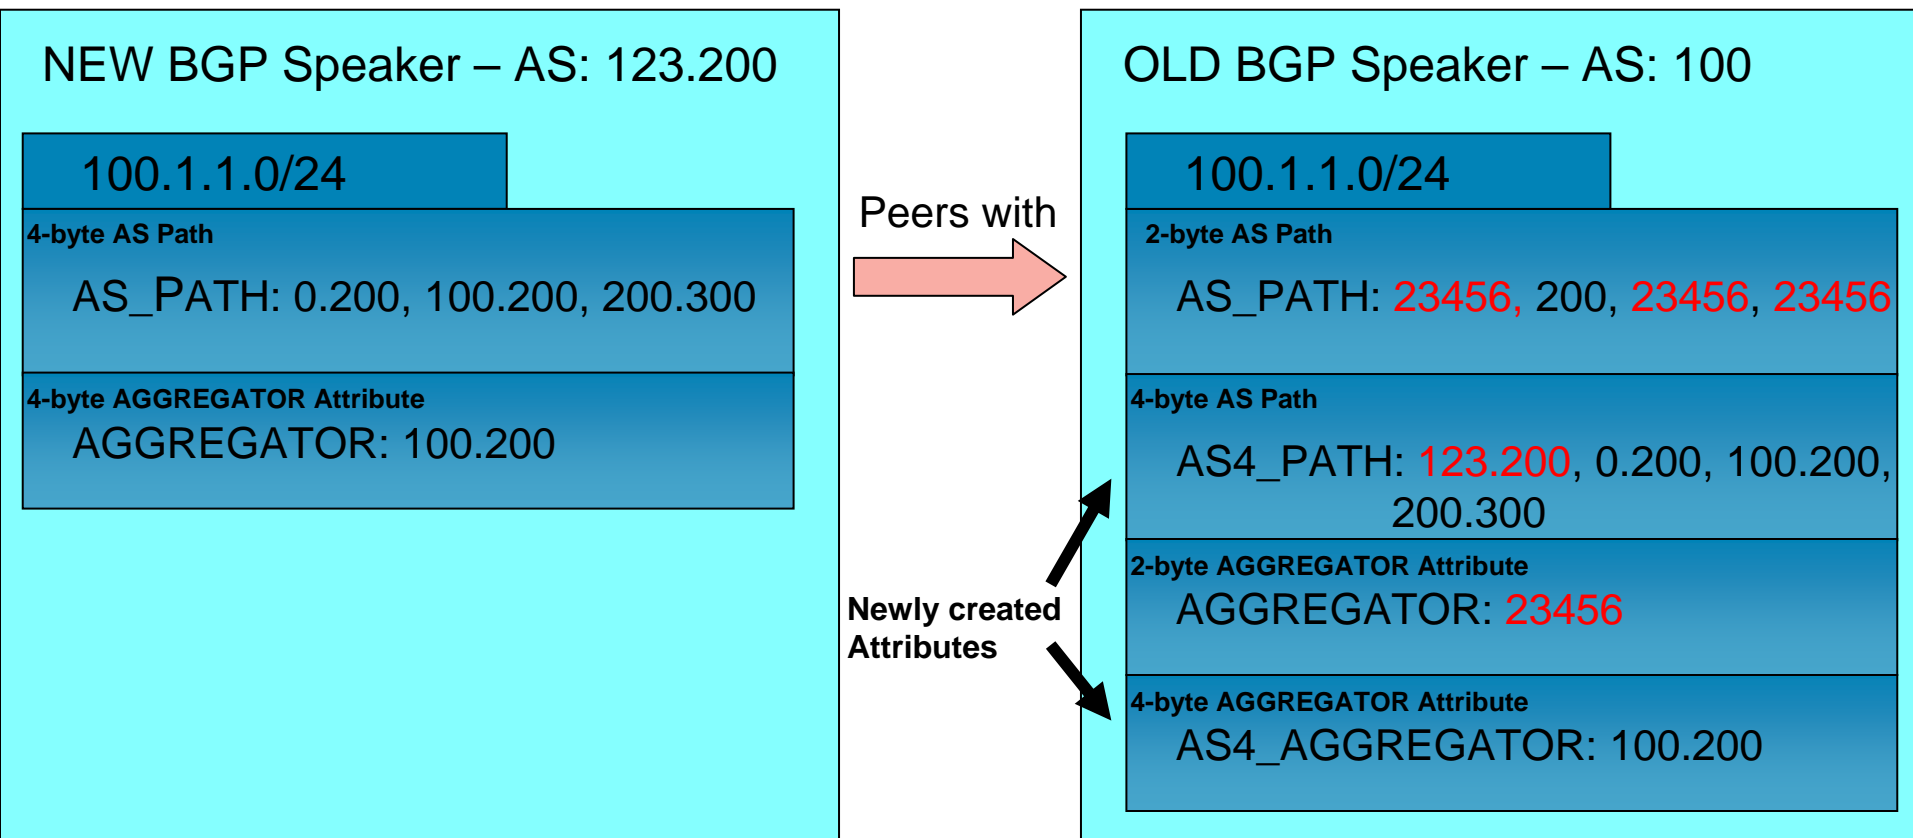
Non-mappable AS Numbers



Each full non-Mappable AS entry will be swapped with well known AS_TRANS (23456) Autonomous Number

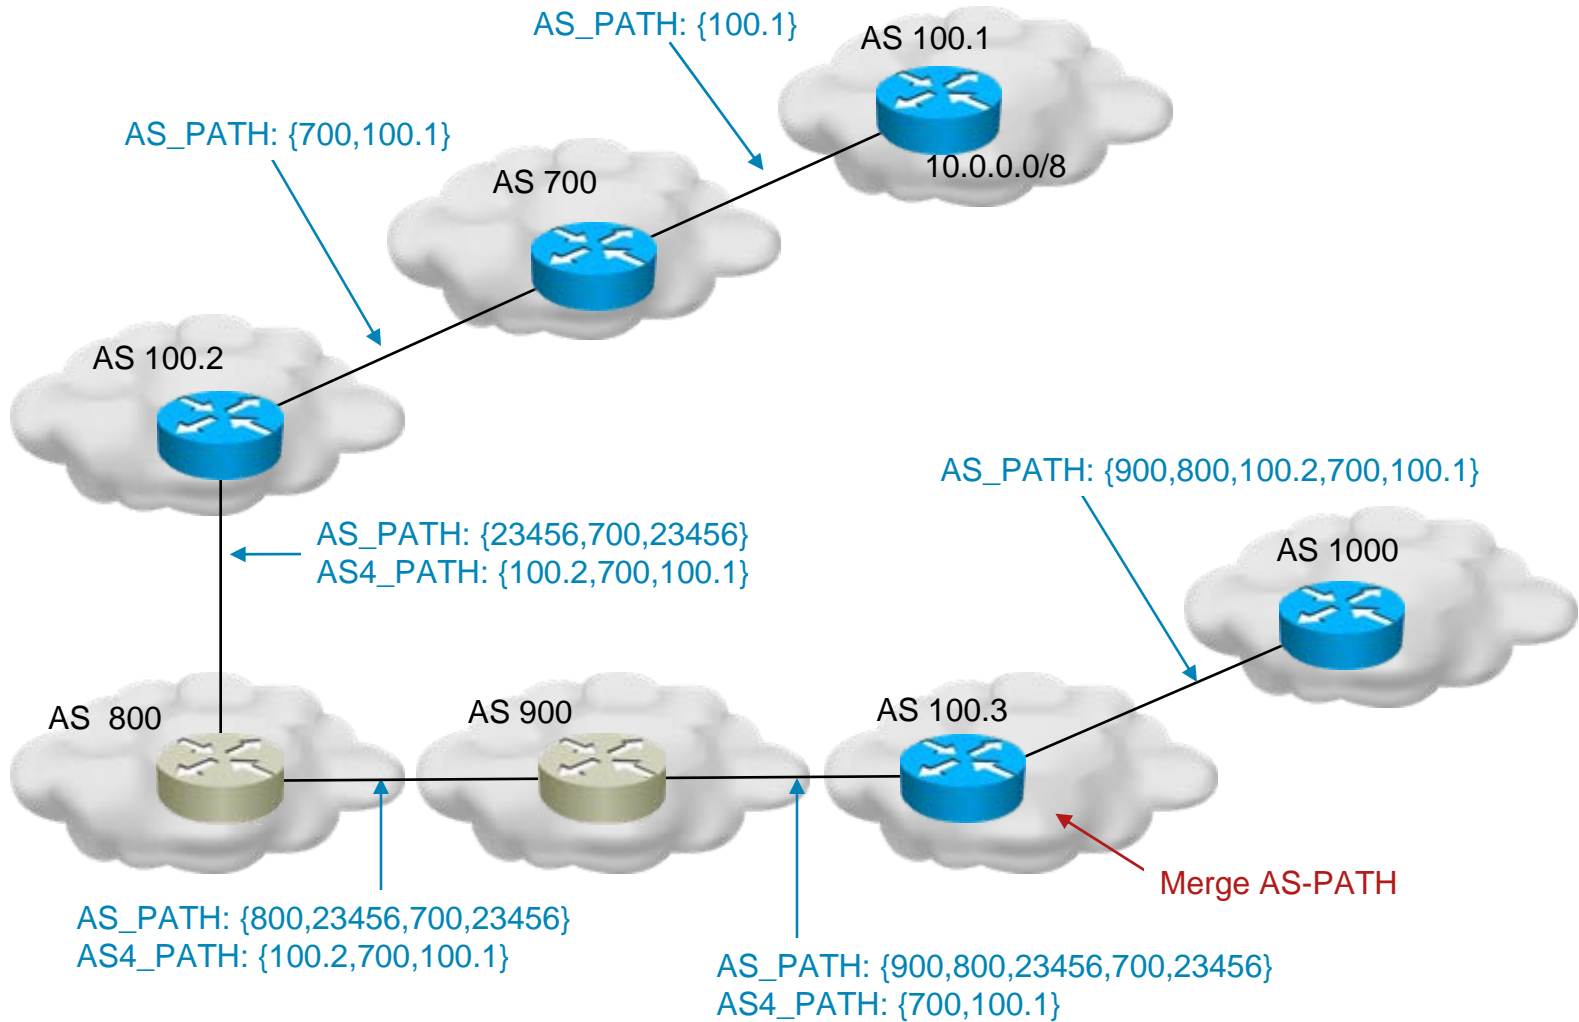
Backward Compatibility

Non-mappable AS Numbers (Cont.)




4 Byte AS

Operation Example



4 Byte AS

Aggregation Example



- AS 100.3 creates 10.1.1.0/24 aggregate

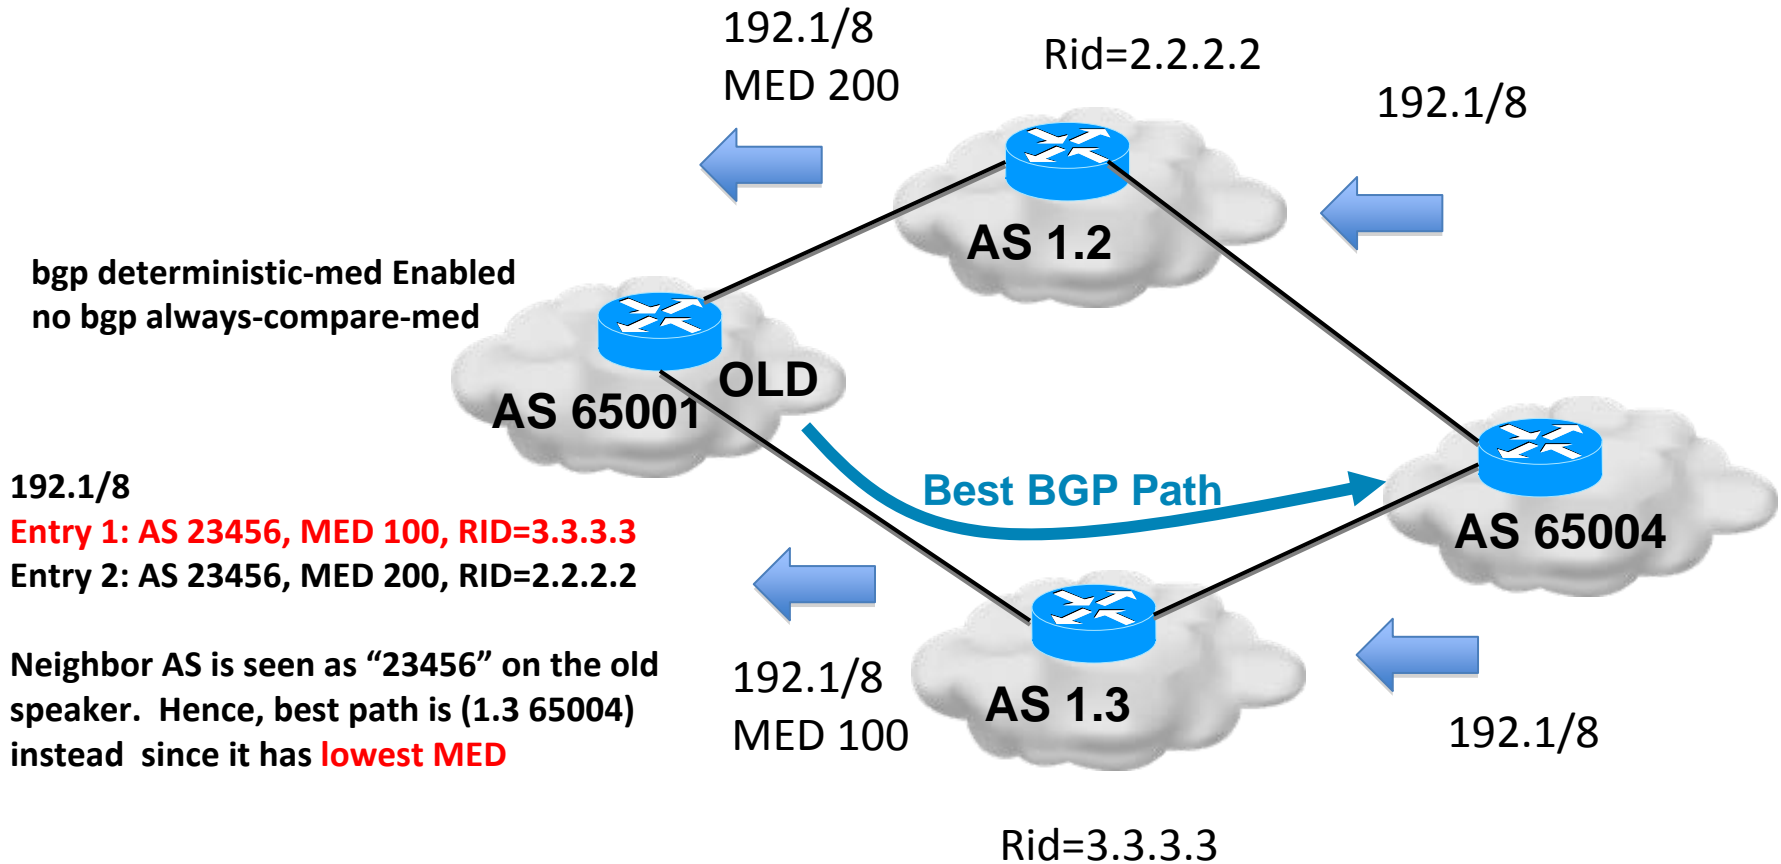
Considerations When My BGP Autonomous System Does Not Support 4-byte AS

- Filtering based on 4-byte AS Numbers is impossible on a OLD BGP speaker
- It is illegal to use the well known AS_TRANS as a BGP Autonomous System Number
- 4-byte AS Numbers can experience additional BGP memory utilization on OLD BGP speakers due to usage of AS4_PATH and AS4_AGGREGATOR attributes
- Due to AS_TRANS usage, the NetFlow v9 created traffic matrix may be gradually more and more incorrect when 4-byte AS numbers are really allocated to users on an OLD BGP speaker
- BGP route aggregation on an OLD BGP speaker may create routing BGP loops under certain conditions (ref. RFC4893)
- Upgrading an OLD BGP speaker peering with a non-mappable 4-byte neighbor AS will need a new neighbor configuration when being upgraded from an OLD BGP speaker to a NEW BGP speaker (swap AS “23456” with the real 4-byte ASN within the BGP neighbor statement)
- Due to the usage of AS_TRANS, this could result in the wrong usage of the MED metrics during BGP path selection (see next slides)

MED - Old Speaker with 4b AS Transit




For your reference



MED - New Speaker with 4b AS Transit



For your reference





For your
reference

Configuration

```
router bgp 4.4  
  bgp log-neighbor-changes  
  neighbor 134.0.0.3 remote-as 3.3
```

Neighbor Configuration

R4#sh ip bgp 1.1.1.0

```
BGP routing table entry for 1.1.1.0/24, version 2  
Paths: (1 available, best #1, table default)  
Flag: 0x820  
  Not advertised to any peer
```

3.3 2 1.1

```
  134.0.0.3 from 134.0.0.3 (134.0.0.3)  
    Origin IGP, localpref 100, valid, external, best
```

R4#sh ip bgp sum

```
BGP router identifier 134.0.0.4, local AS number 4.4  
BGP table version is 2, main routing table version 2  
1 network entries using 124 bytes of memory  
1 path entries using 52 bytes of memory  
2/1 BGP path/bestpath attribute entries using 184 bytes of memory  
1 BGP AS-PATH entries using 40 bytes of memory  
0 BGP route-map cache entries using 0 bytes of memory  
0 BGP filter-list cache entries using 0 bytes of memory  
BGP using 400 total bytes of memory  
BGP activity 1/0 prefixes, 1/0 paths, scan interval 60 secs  
Neighbor      V  AS MsgRcvd MsgSent  TblVer  InQ OutQ Up/Down State/PfxRcd  
134.0.0.3    4  3.3   28    27     2    0  0 00:25:33    1
```

BGP Show Command



For your
reference

Configuration

```
R3#sh ip rout | include B
```

Codes: L - local, C - connected, S - static, R - RIP, M - mobile, B - BGP

```
B      2.2.2.0 [20/0] via 123.0.0.2, 00:11:01
```

```
B 192.0.0.0/24 [20/0] via 123.0.0.2, 00:11:01
```

```
R3#sh ip route 192.0.0.0
```

Routing entry for 192.0.0.0/24

Known via "**bgp 3.3**", distance 20, metric 0

Tag 2, type external

Redistributing via ospf 1

Advertised by ospf 1

Last update from 123.0.0.2 00:12:14 ago

Routing Descriptor Blocks:

* 123.0.0.2, from 123.0.0.2, 00:11:09 ago

Route metric is 0, traffic share count is 1

AS Hops 1

Route tag 2

Routing Show Commands

```
ip as-path access-list 1 permit ^1\.4$  
router bgp 1  
neighbor 4.4.4.4 remote-as 1.4  
neighbor 4.4.4.4 route-map foo in
```

Note that the "." must be
escaped from the
regular expression with a "\"

AS-PATH Filter for ASDOT notation


```
route-map foo permit 10  
match as-path 1
```

References

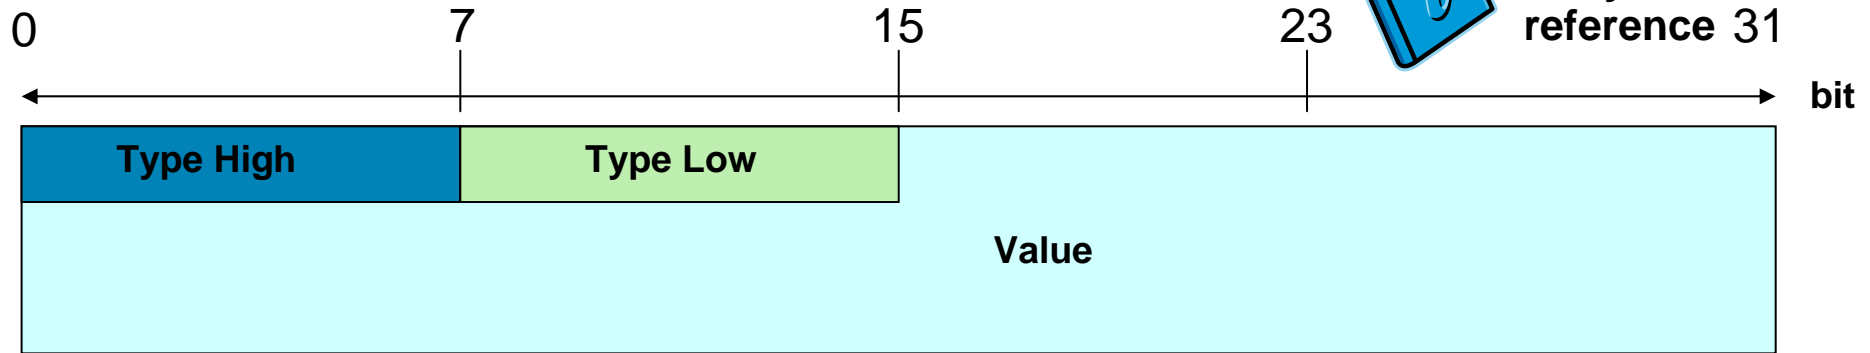
- RFC4893 – “BGP Support for Four-octet AS Number Space”
- RFC5396 – “Textual Representation of Autonomous System (AS) Numbers “
- RFC2842 – “Capabilities Advertisement with BGP-4 ”
- 16-bit AS Number Report
<http://www.potaroo.net/tools/asn16/>
- ARIN, AS Number Change on 1 January 2009
<http://www.arin.net/announcements/07242008.html>
- RIPE NCC, AS Number change could affect Internet routing from 1 January 2009
<http://www.ripe.net/news/asn-32-pr2008.html>
- APNIC, AS number change could affect Internet routing from 1 January 2009
<http://www.apnic.net/news/2008/0725.html>



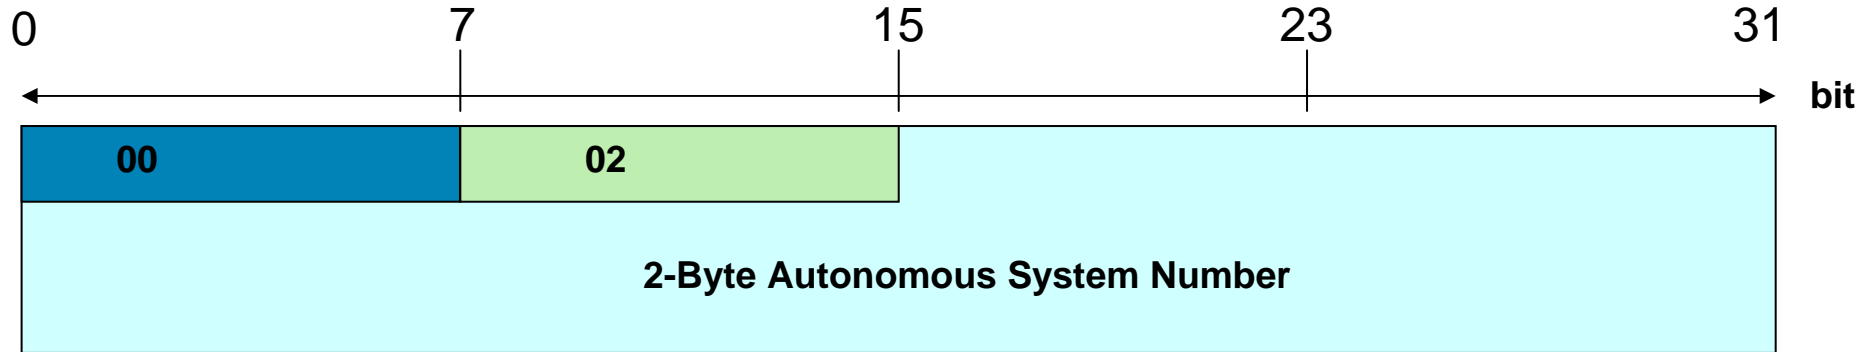
RFC4360: BGP Extended Communities Attribute



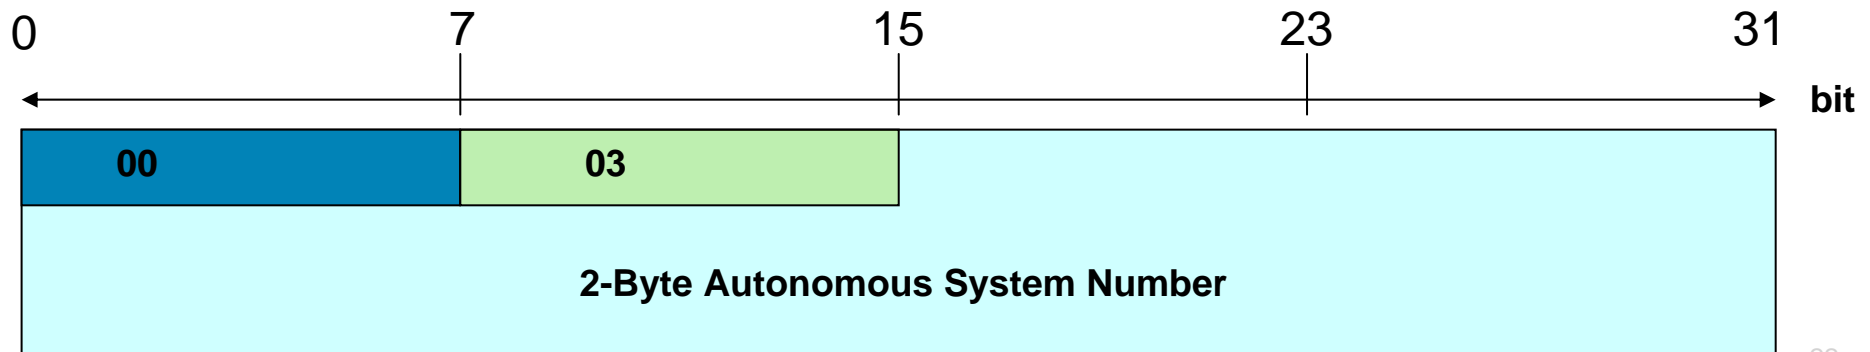
For your reference 31



RFC4360: Route-Target Extended Communities Attribute – 2 Byte AS



RFC4360: SOO Extended Communities Attribute - 2 Byte AS

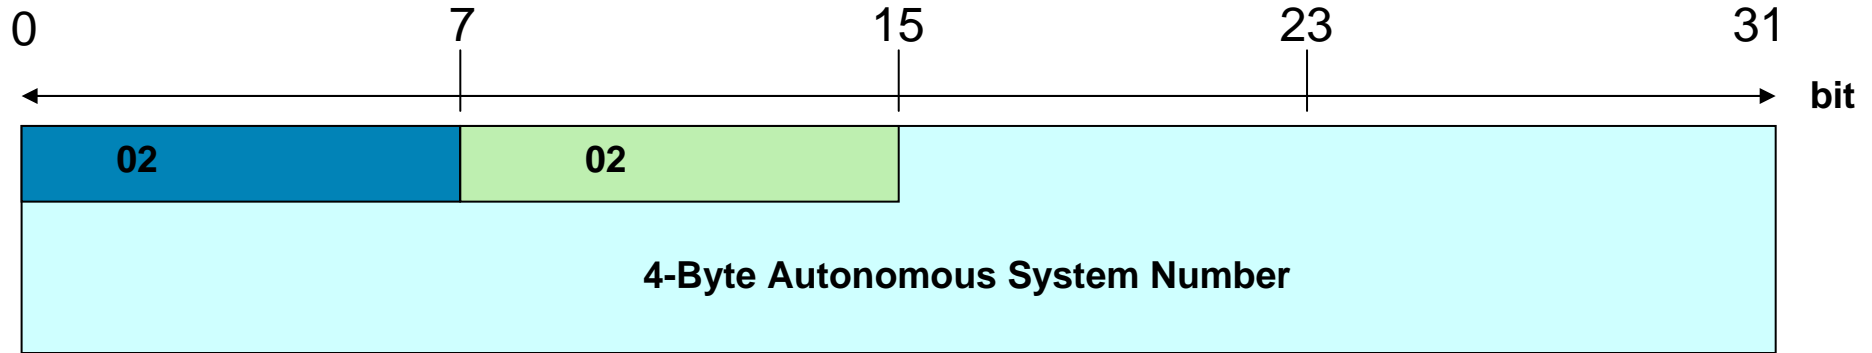


Four-octet AS Specific BGP Extended Community (ietf-l3vpn-as4octet-ext-community-02.txt)



For your
reference

Route-Target Extended Communities Attribute – 4 Byte AS



SOO Extended Communities Attribute - 4 Byte AS

