



4-Bytes ASN




PLNOG 3 2009-09-11

Klaudiusz Staniek

kstaniek@cisco.com



Autonomous System (AS)



- Collection of networks with same routing policy
- Single routing protocol
- Usually under single ownership, trust and administrative control
- Identified by a unique number

Autonomous System Number

- An ASN is a 16-bit integer
 - 1-64511 are assigned by the Regional Internet Registries
 - 64512-65534 are private ASNs and should never be used on the Internet
 - 0 and 65535 are reserved
- Current allocations up to 56318 have been made to the RIRs


ASN status

- The pool of 16-bit ASNs will soon be exhausted

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

Current estimates are that the 16-bit ASN pool will be exhausted in 2011

- Work started in 2001 to extend the ASN pool to 32-bits



Source: <http://www.potaroo.net/tools/asns/fig28.png>

32-bit ASNs

- Standards documents

Description of 32-bit ASNs

www.rfc-editor.org/rfc/rfc4893.txt

Textual representation

www.rfc-editor.org/rfc/rfc5396.txt

New extended community

<http://tools.ietf.org/html/draft-ietf-idr-as4octet-extcomm-generic-subtype-00>

<http://tools.ietf.org/html/draft-rekhter-as4octet-ext-community-03>

<http://tools.ietf.org/html/draft-ietf-l3vpn-as4octet-ext-community-03>

- AS 23456 is reserved as interface between 16-bit and 32-bit ASN world

Getting a 32-bit ASN

- Sample RIR policy (RIPE)
<http://www.ripe.net/ripe/docs/asn-assignment.html>
From 1st January 2007
32-bit ASNs available on request
- From 1st January 2009
32-bit ASNs assigned by default
16-bit ASNs only available on request
- From 1st January 2010
No distinction – ASNs assigned from 32-bit pool

Representation

- 32-bit ASNs extend the pool:
0-65535 extended to 0-4294967295
- Representation of 65536-4294967295 range

Most operators favour traditional format (asplain)

A few prefer dot notation (X.Y):

asdot for 65536-4294967295, e.g 2.4

asdot+ for 0-4294967295, e.g 0.64513

But regular expressions will have to be completely rewritten for asdot and asdot+ !!!

asplain vs asdot(+)

- Problem:

$^{[0-9]^+}$ matches any ASN (16-bit and asplain)

This and equivalents extensively used in BGP multihoming configurations for traffic engineering

- Equivalent regexp for asdot is

$^{([0-9]^+)|([0-9]^+\.[0-9]^+)}$

- Equivalent regexp for asdot+ is

$^{[0-9]^+\.[0-9]^+}$

- \Rightarrow BGP policy regular expressions will need to be rewritten

IANA Assignments

- Using dot notation for readability
- 0.0 - 0.65535 16-bit ASN block
- 2.0 - 2.1023 APNIC
- 3.0 - 3.1023 RIPE NCC
- 4.0 - 4.1023 LACNIC
- 5.0 - 5.1023 AfriNIC
- 6.0 - 6.1023 ARIN
- Remainder are reserved or held by IANA

IANA Assignments (Special)

- 0 & 65535 Reserved
- 23456 32-bit ASN transition
- 64512 - 65534 Private ASNs
- 64496 - 64511 Documentation (16-bit)
- 65536 - 65551 Documentation (32-bit)

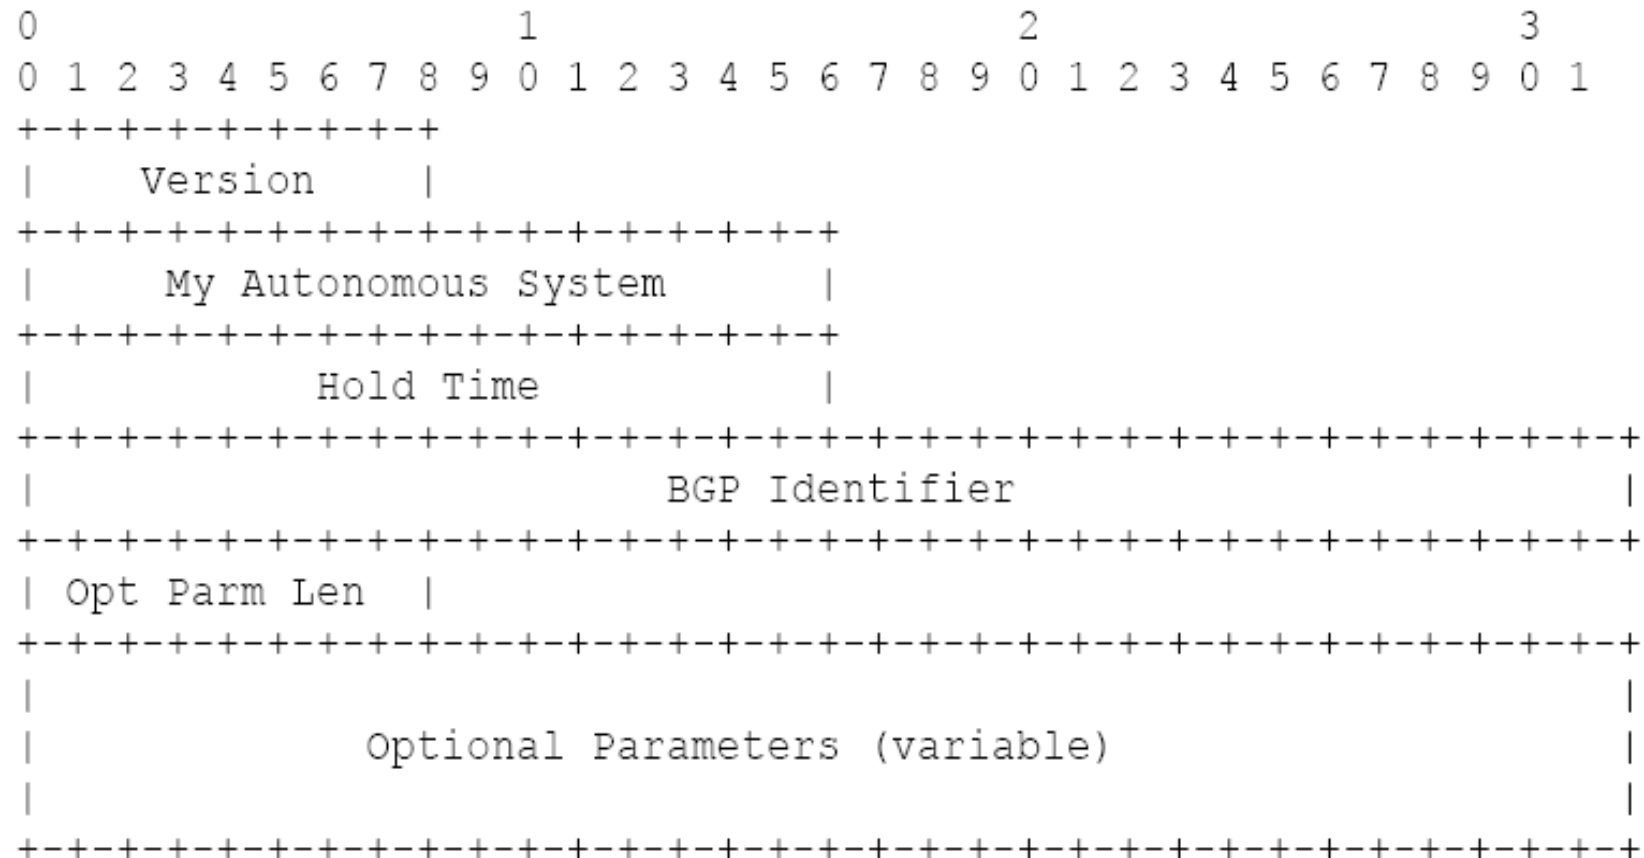
Changes (1)

- 32-bit ASNs are backwardly compatible with 16-bit ASNs
- There is no flag day
- You do NOT need to:
 - Throw out your old routers
 - Replace your 16-bit ASN with a 32-bit ASN

Changes (2)

- You do need to be aware that:
 - Your customers will come with 32-bit ASNs
 - ASN 23456 is not a bogon!
 - You will need a router supporting 32-bit ASNs to use a 32-bit ASN
- If you have a proper BGP implementation, 32-bit ASNs will be transported silently across your network

“BGP OPEN protocol message format”



How does it work (1)?

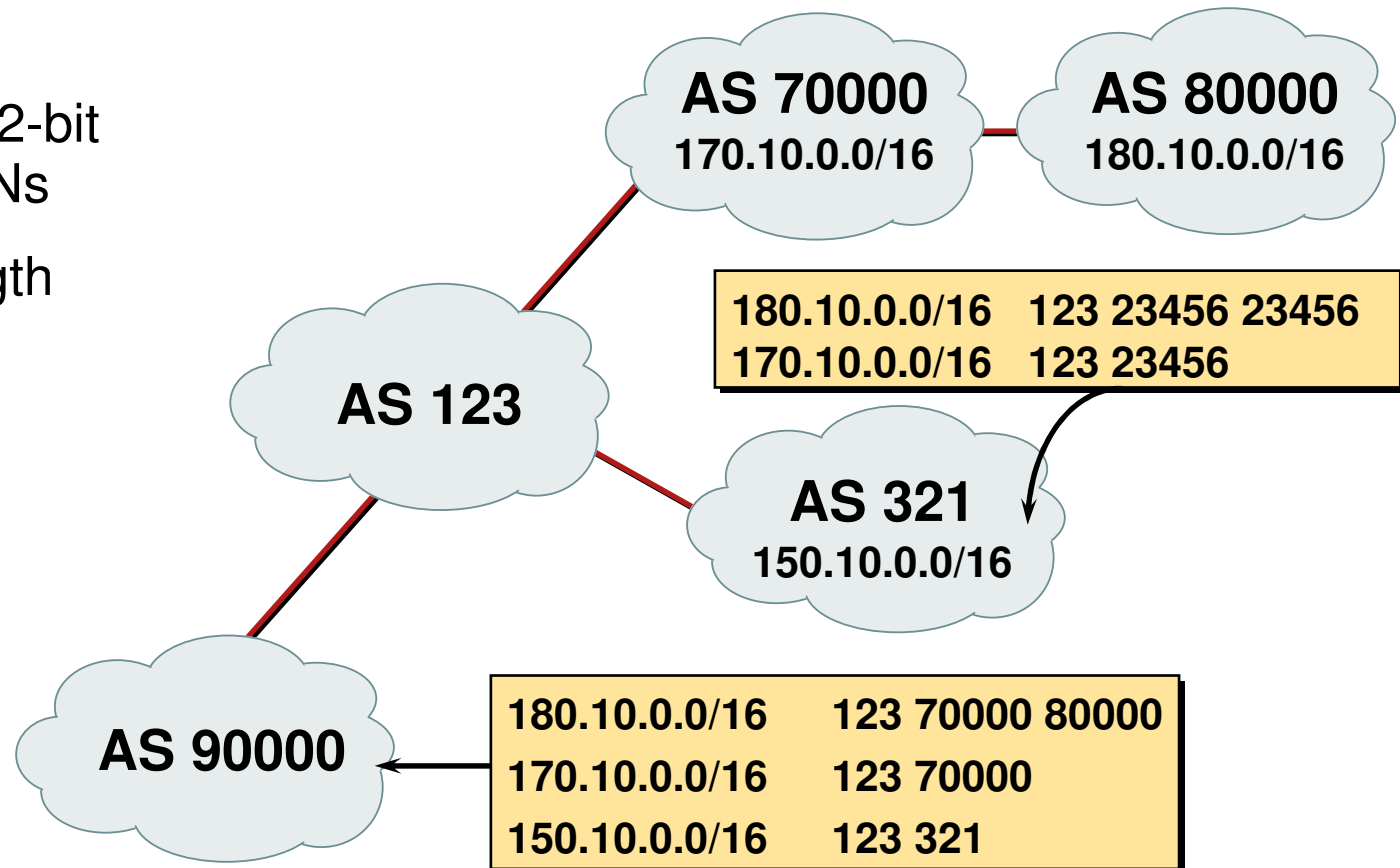
- Local router only supports 16-bit ASN
- Remote router uses 32-bit ASN
- BGP peering initiated:
 - Remote asks local if 32-bit supported (BGP capability negotiation)
 - When local says “no”, remote then presents AS23456
 - Local needs to be configured to peer with remote using AS23456

How does it work (2)?

- BGP peering initiated (cont):
 - BGP session established using AS23456
 - 32-bit ASN included in a new BGP attribute called AS4_PATH
 - (as opposed to AS_PATH for 16-bit ASNs)
- Result:
 - 16-bit ASN world sees 16-bit ASNs and 23456 standing in for 32-bit ASNs
 - 32-bit ASN world sees 16 and 32-bit ASNs

Example:

- Internet with 32-bit and 16-bit ASNs
- AS-PATH length maintained



Configuration Example (1)

- AS70000 and AS80000 border routers:

Configuration on AS80000:

```
router bgp 80000
```

```
neighbor 1.1.1.6 remote-as 70000
```

Configuraton on AS70000:

```
router bgp 70000
```

```
neighbor 1.1.1.5 remote-as 80000
```

Both routers have to support 32-bit ASNs

Configuration Example (2)

- AS123 and AS70000 border routers:

Configuration on AS123:

```
router bgp 123
```

```
neighbor 1.1.1.2 remote-as 23456
```

Configuraton on AS70000:

```
router bgp 70000
```

```
neighbor 1.1.1.1 remote-as 123
```

AS70000 router supports 32-bit ASNs

AS123 router does **not** support 32-bit ASNs

What has changed?

- Two new BGP attributes:

AS4_PATH

Carries 32-bit ASN path info

AS4_AGGREGATOR

Carries 32-bit ASN aggregator info

Well-behaved BGP implementations will simply pass these along if they don't understand them

- AS23456 (AS_TRANS)

What do they look like?

- IPv4 prefix originated by AS196613

```
as4-7200#sh ip bgp 145.125.0.0/20
BGP routing table entry for 145.125.0.0/20, version 58734
Paths: (1 available, best #1, table default)
```

**asplain
format**

```
131072 12654 196613
204.69.200.25 from 204.69.200.25 (204.69.200.25)
Origin IGP, localpref 100, valid, internal, best
```

- IPv4 prefix originated by AS3.5

```
as4-7200#sh ip bgp 145.125.0.0/20
BGP routing table entry for 145.125.0.0/20, version 58734
Paths: (1 available, best #1, table default)
```

**asdot
format**

```
2.0 12654 3.5
204.69.200.25 from 204.69.200.25 (204.69.200.25)
Origin IGP, localpref 100, valid, internal, best
```

What do they look like?

- IPv4 prefix originated by AS196613

But 16-bit AS world view:

```
BGP-view1>sh ip bgp 145.125.0.0/20
```

```
BGP routing table entry for 145.125.0.0/20, version 113382
```

```
Paths: (1 available, best #1, table Default-IP-Routing-Table)
```

```
23456 12654 23456
```

```
204.69.200.25 from 204.69.200.25 (204.69.200.25)
```

```
Origin IGP, localpref 100, valid, external, best
```


**Transition
AS**

What do they look like?

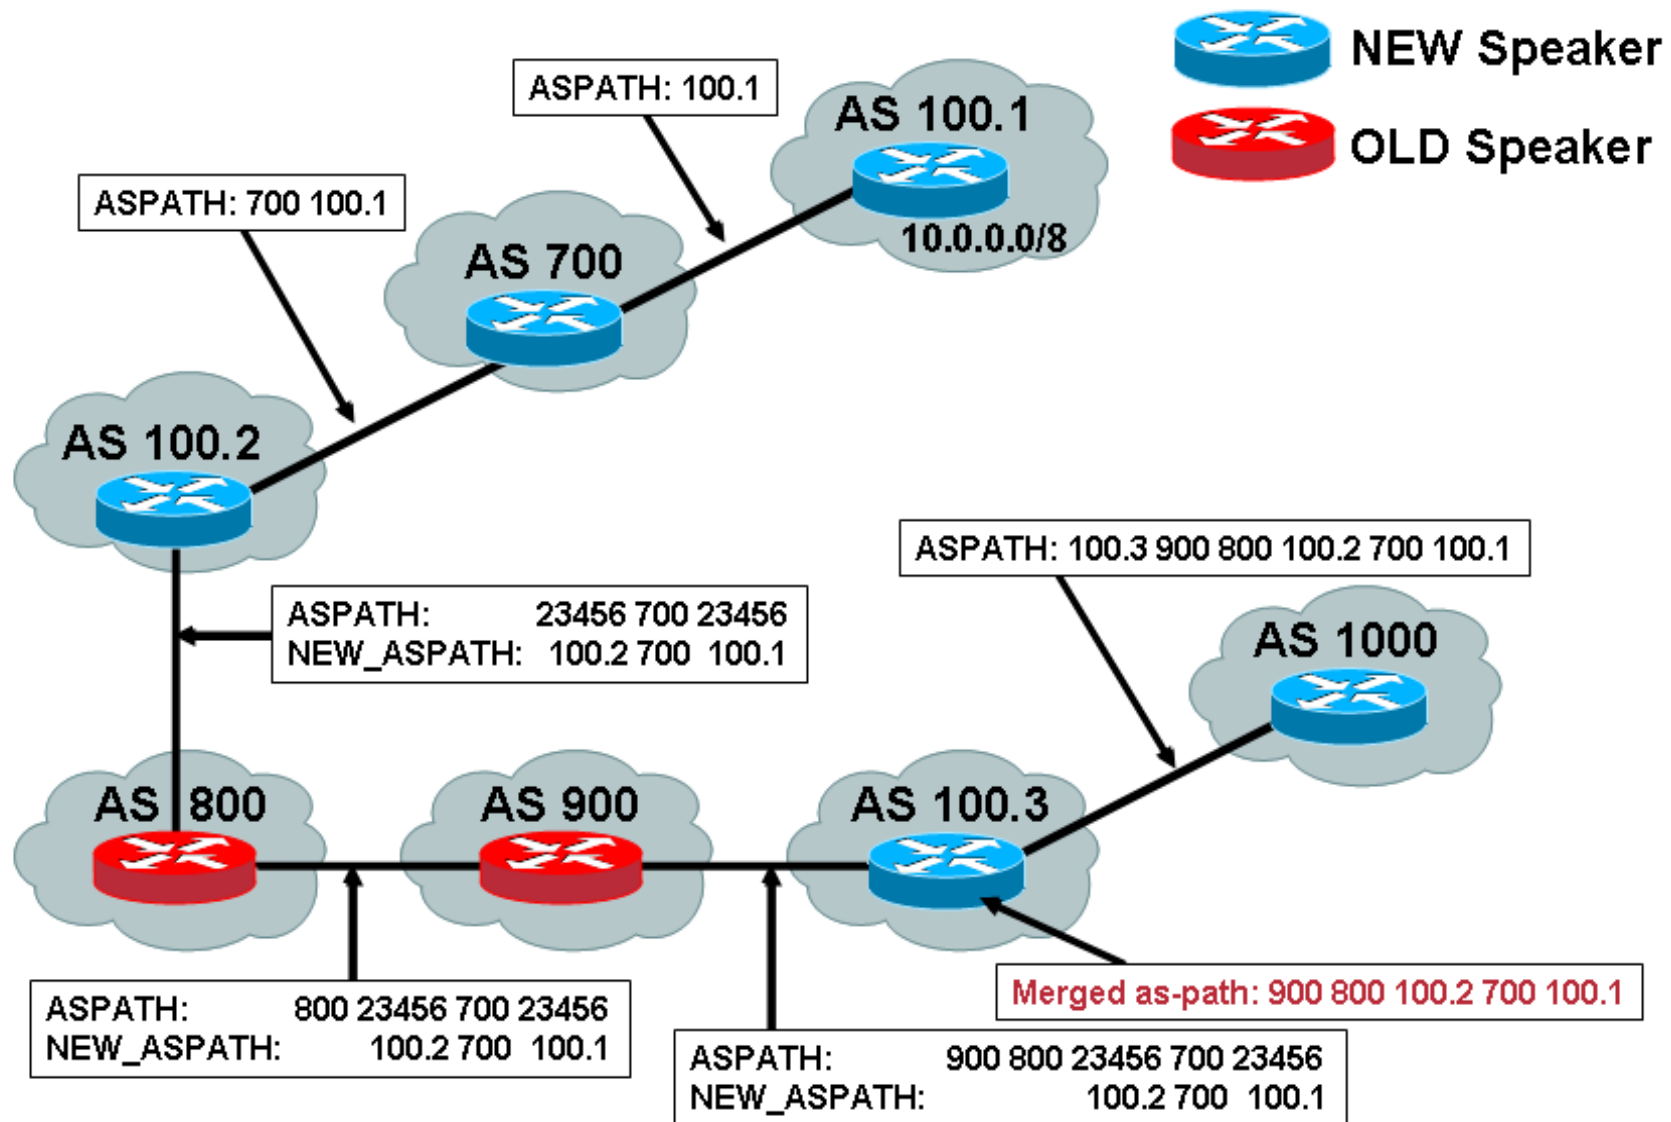
- IPv6 prefix originated by AS 2.9

```
RP/0/0/CPU0:as4byte#show bgp ipv6 unicast 2403:2000::/32
BGP routing table entry for 2403:2000::/32
Versions:
  Process          bRIB/RIB  SendTblVer
  Speaker          93        93
Paths: (1 available, best #1)
  Not advertised to any peer
  Path #1: Received by speaker 0
    109 6175 2497 2500 18146 2.9
      2001:420:0:8001::1 from 2001:420:0:8001::1 (204.69.200.22)
      Origin IGP, localpref 100, valid, external, best
```

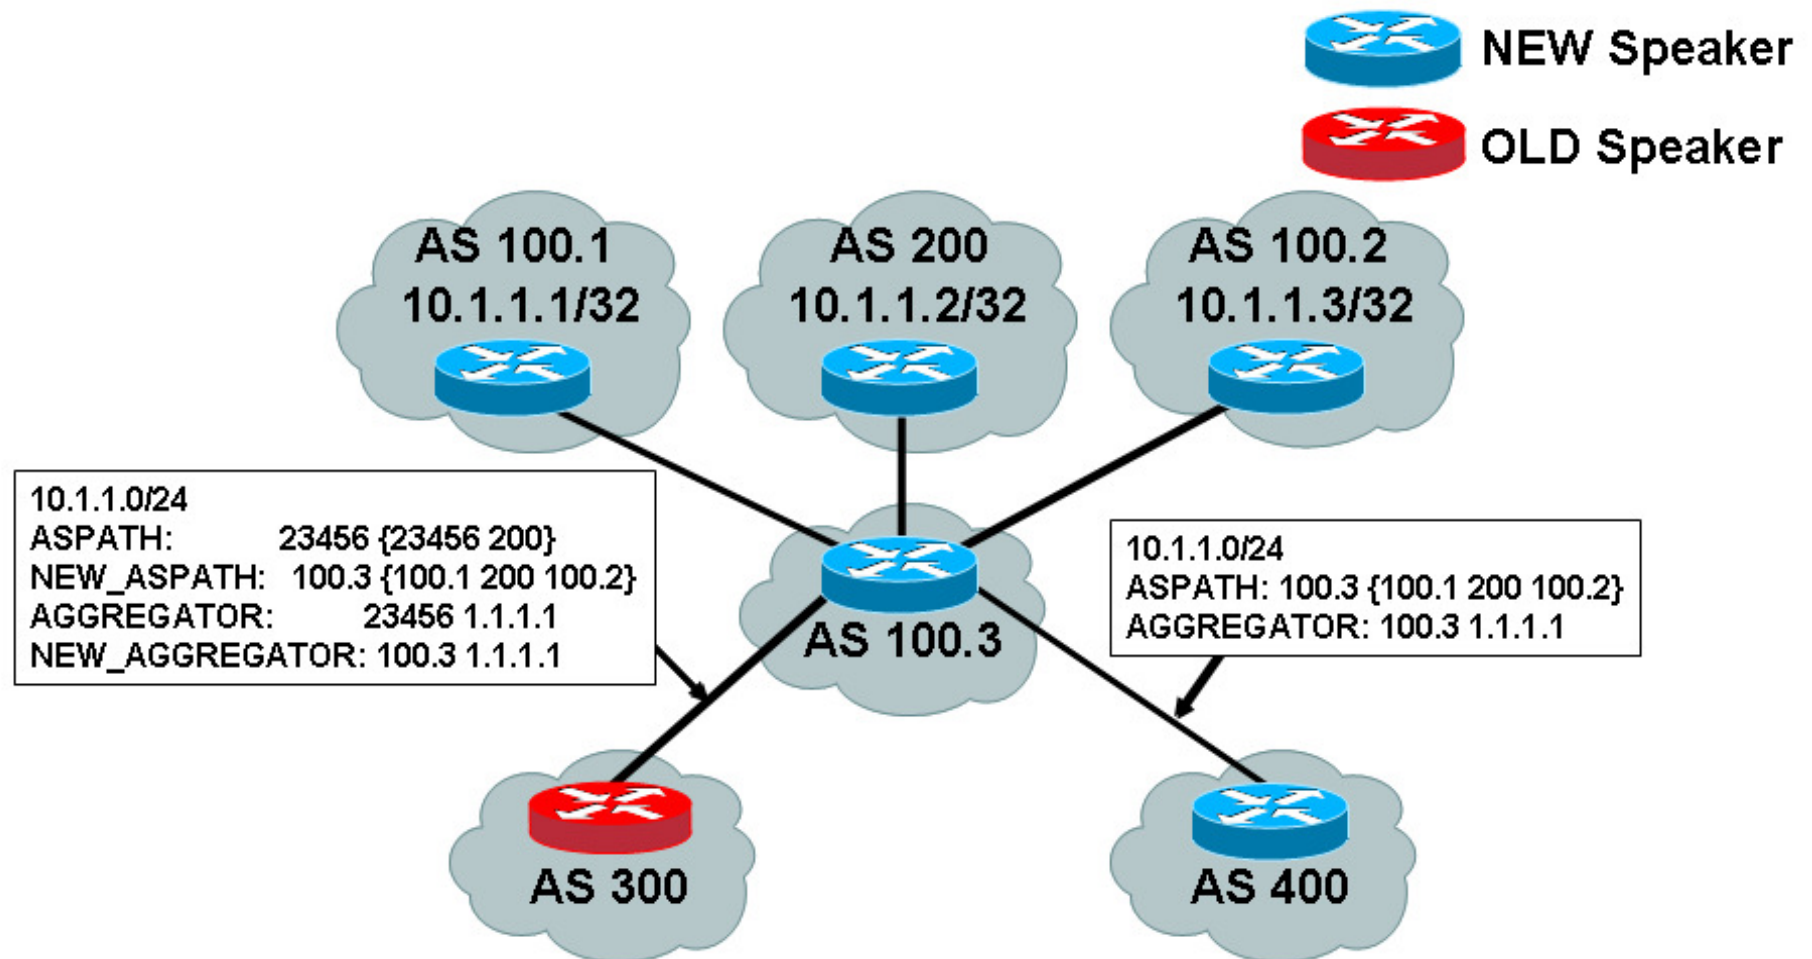
**32-bit
ASN**



4-byte AS – ASPATH & NEW_AS_PATH in a mixed environment



4-byte AS – aggregation



- AS 100.3 creates 10.1.1.0/24 aggregate

Considerations

- AS4_PATH attribute are aggregated by an OLD BGP speaker

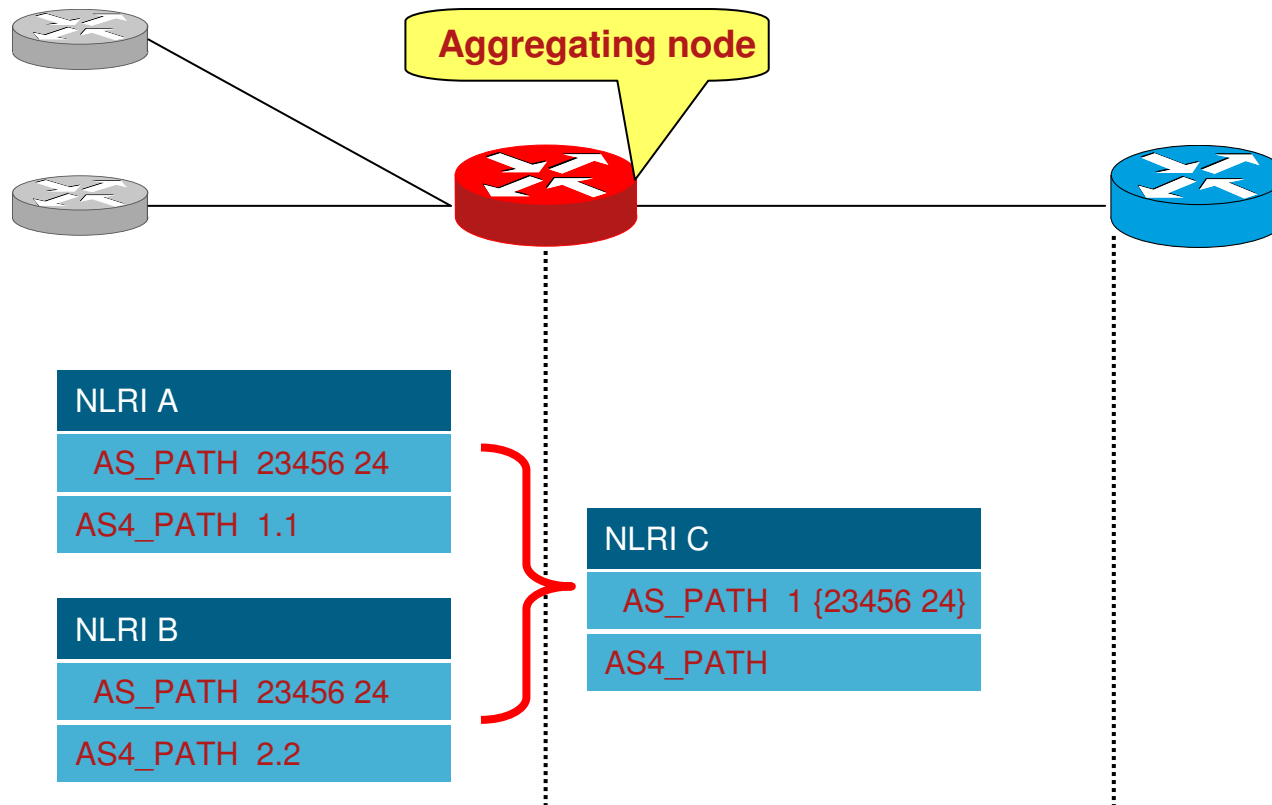
When two or more routes that carry the AS4_PATH attribute are aggregated by an OLD BGP speaker, and the AS4_PATH attribute of at least one of these routes carries at least one 4-octet AS number. It may not be possible to reconstruct the entire AS path information from the AS_PATH and the AS4_PATH attributes of a route.

When such aggregation results in creating a route that is less specific than any of the component routes (route whose Network Layer Reachability Information (NLRI) covers NLRI of all the component routes), loss of the AS path information does not create a risk of a routing loop. In all other cases, loss of the AS path information does create a risk of a routing loop.

Considerations

OLD Speaker use as AGGREGATION need care for Loop detection

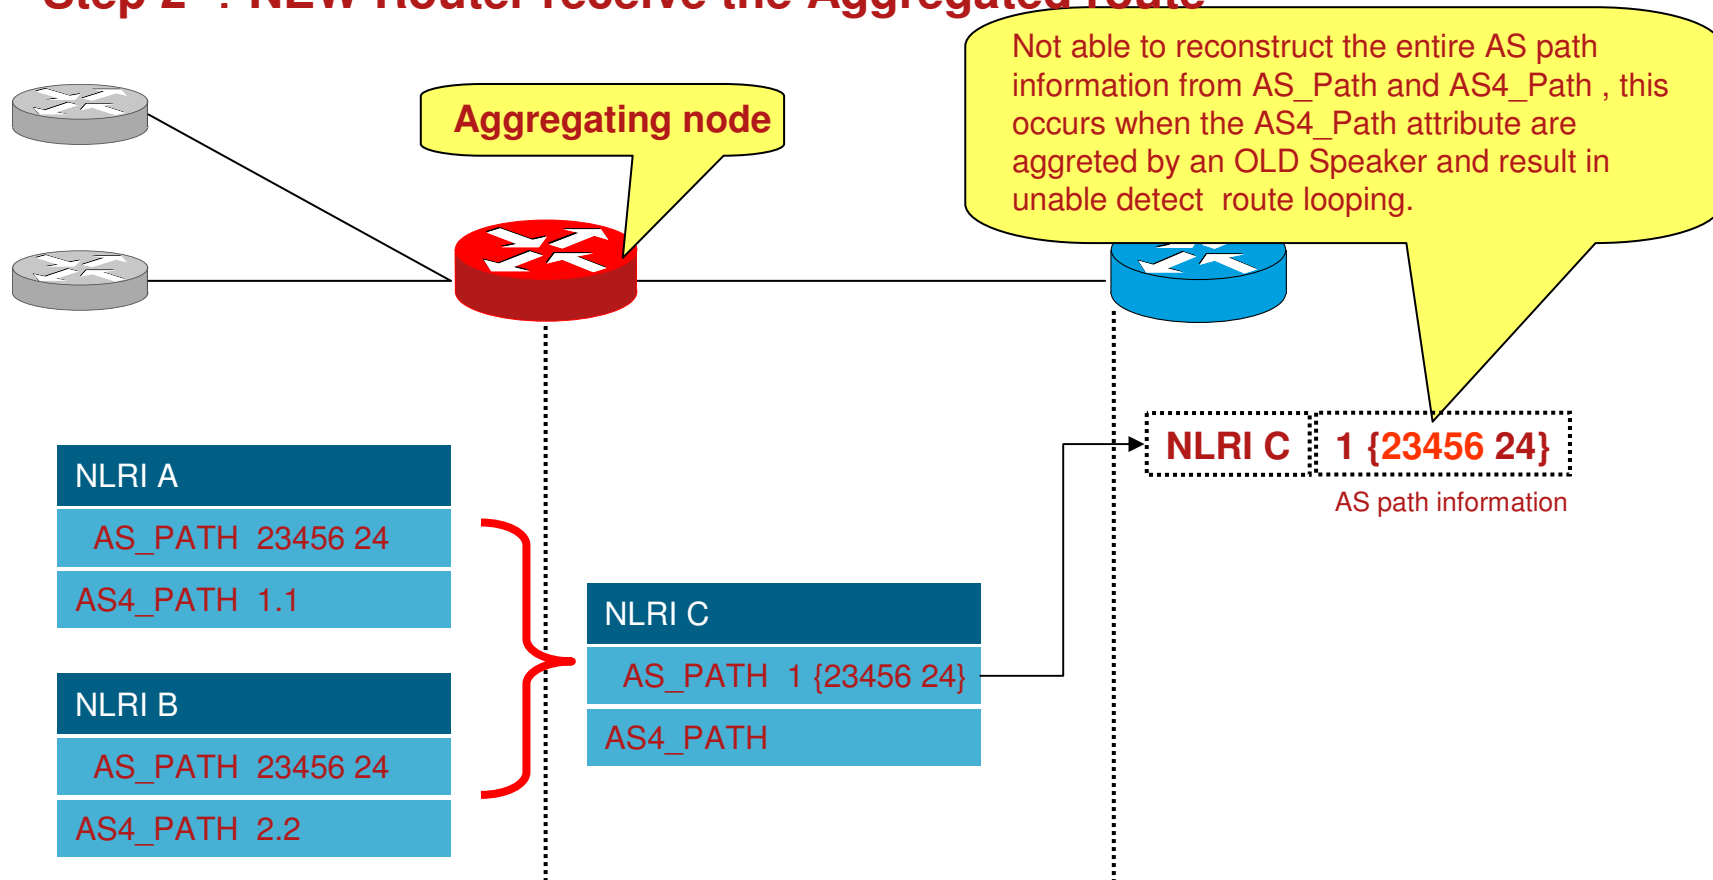
Step 1 : As-set option set on this Aggregating node



Considerations

OLD Speaker use as AGGREGATION need care for Loop detection

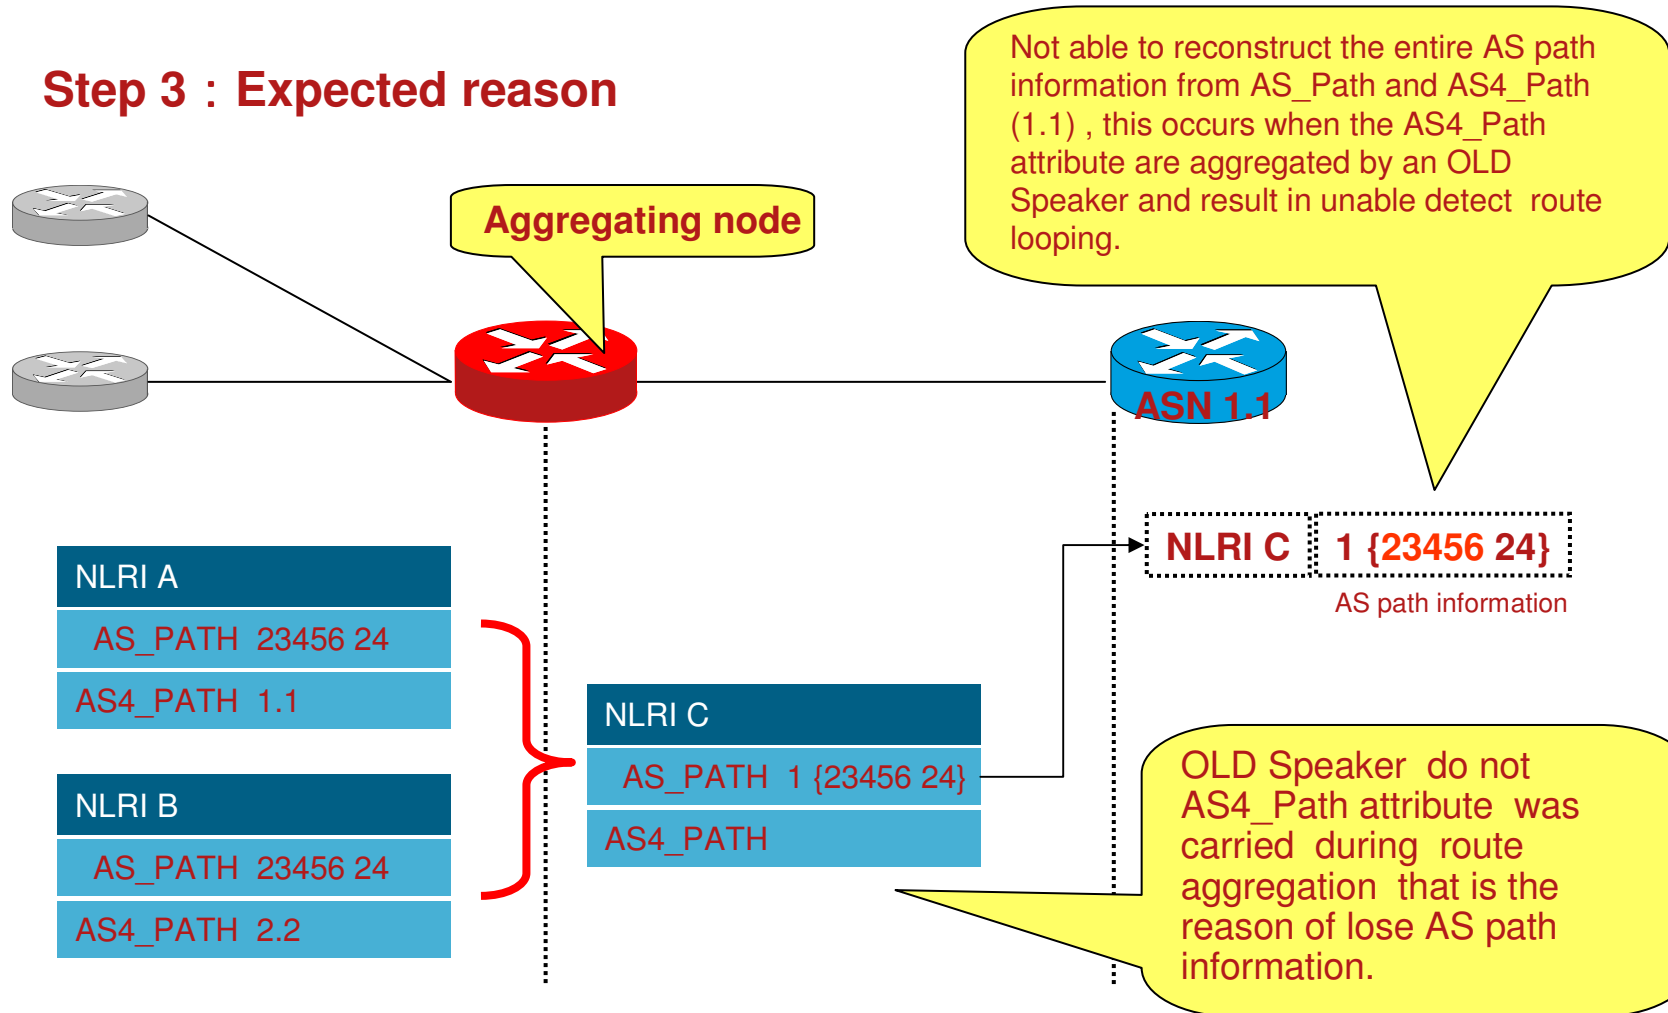
Step 2 : NEW Router receive the Aggregated route



Considerations

OLD Speaker use as AGGREGATION need care for Loop detection

Step 3 : Expected reason



Considerations (Cont)

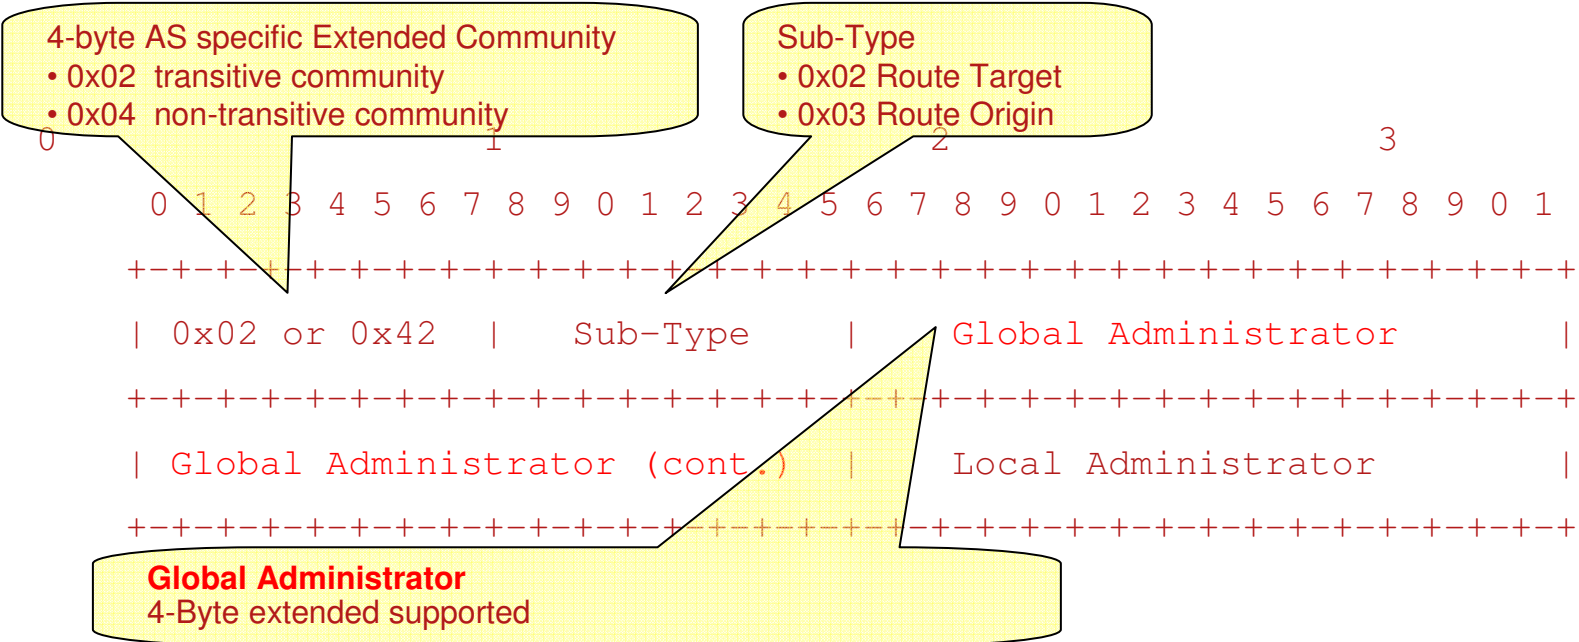
- Handling BGP Communities

As specified in [RFC1997], when the high-order two-octets of the community attribute is neither 0x0000 nor 0xffff, these two octets encode the Autonomous System number. **Quite clearly this would not work for BGP speakers that use 4-octets Autonomous System numbers.** Such BGP speakers should use the Four-octet AS Specific Extended Communities [AS-EXT-COM] instead.

Considerations

Handling BGP Extended Communities

- [RFC1997], when the high-order 2-octets of the community attribute is neither 0x0000 nor 0xffff, these two octets encode the Autonomous System number .
- NEW BGP Router, use the 4-octet AS Specific Extended Communities instead.
- **How OLD router to processing when received Extended Community .**

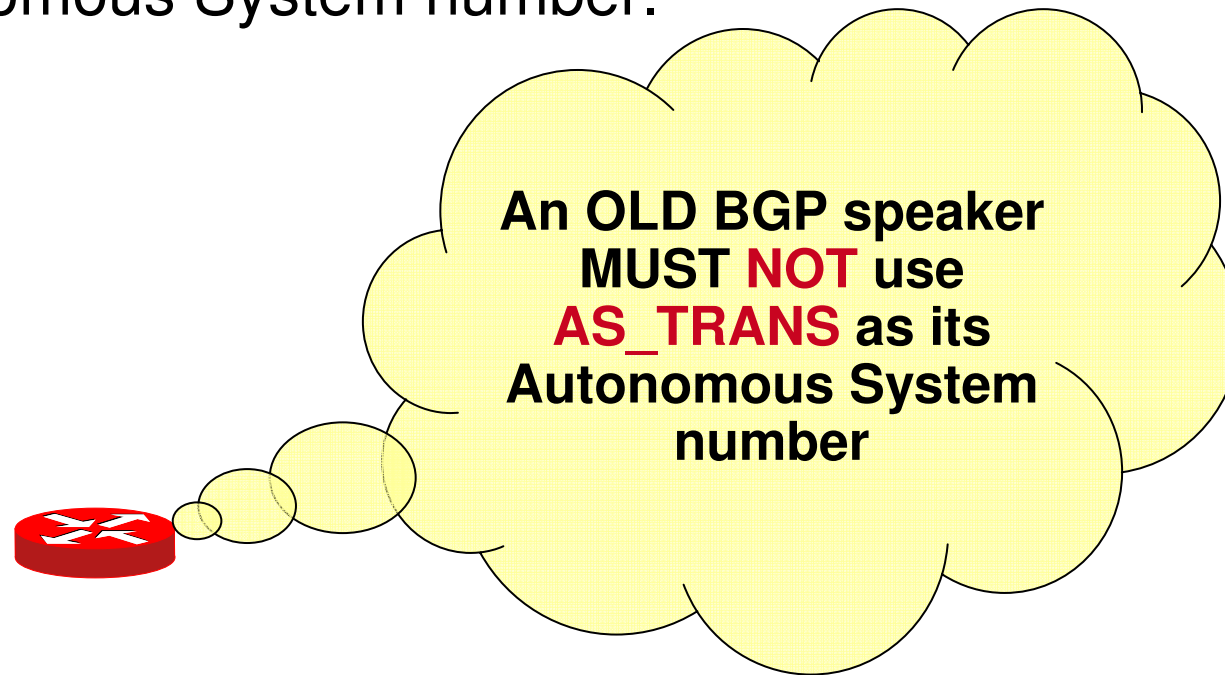


Considerations

AS_TRANS

- AS_TRANS

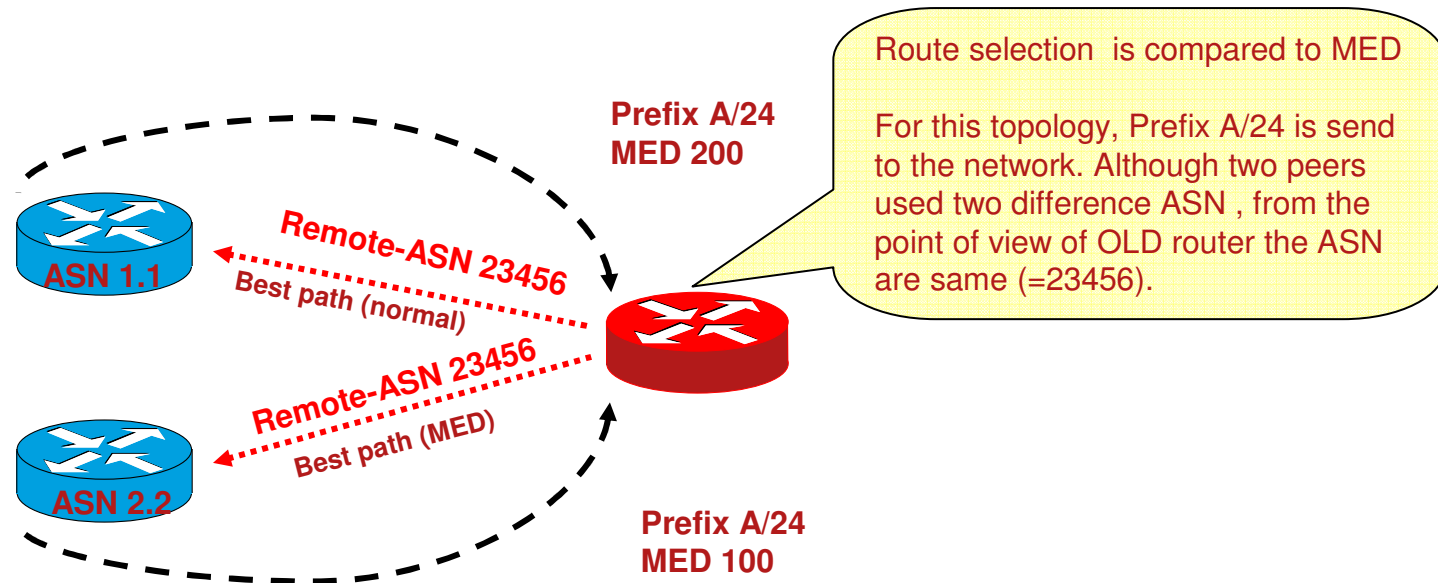
An OLD BGP speaker **MUST NOT** use AS_TRANS as its Autonomous System number.



Considerations

Multi-Exit Discriminators

- An Autonomous System that has OLD BGP speakers peers with two or more Autonomous Systems that have NEW BGP speakers and use AS_TRANS, use of Multi-Exit Discriminators by the Autonomous System with the OLD speakers may result in a situation where **Multi-Exit Discriminator will influence route selection among the routes that were received from different neighboring Autonomous Systems.**



32-bit ASN not supported:

- Inability to distinguish between peer ASes using 32-bit ASNs

They will all be represented by AS23456

Could be problematic for transit provider's policy

- Inability to distinguish prefix's origin AS

How to tell whether origin is real or fake?

The real and fake both represented by AS23456

(There should be a better solution here!)

32-bit ASN not supported:

- Incorrect NetFlow summaries:

Prefixes from 32-bit ASNs will all be summarised under AS23456

Traffic statistics need to be measured per prefix and aggregated

Makes it hard to determine peerability of a neighbouring network

Implementations (Sept 09)

- Cisco IOS-XR 3.4 onwards
- Cisco IOS-XE 2.3 onwards
- Cisco IOS 12.0(32)S12 & 12.4(24)T
- Cisco NX-OS 4.0(1)
- Quagga (patches for 0.99.6)
- OpenBGPD (patches for 3.9 & 4.0)
- Juniper JunOSe 4.1.0 & JunOS 9.1
- Redback SEOS
- Force10 FTOS7.7.1 onwards

http://as4.cluepon.net/index.php/Software_Support

What next?

- Pester your router vendors for 32-bit ASN support

Do you really want to run beta software in your core network?

Depletion of the 16-bit pool is not far away

Stable software, deployment cycles

Insist your vendors support “asplain”

Otherwise prepare to rewrite all your regular expressions!!

Conclusion

- The Internet will not break
- Your network will not break
- If you have an ASN today:
 - You don't need to change anything
 - 32-bit ASNs appear as AS 23456
- If you have no ASN today:
 - Your routers will need 32-bit ASN support after 1st January 2009
 - (Or you will need to ask RIRs for a 16-bit ASN)

FOOD FOR THOUGHT !!!

- Do we ever again need to expand BGP ASN from 32-bits to 64 or even 128-bits ?????

FOOD FOR THOUGHT !!! (CORRECT ANSWER)

POSSIBLE !!! If the world can become global village , moon recreational flights have almost begun, life on Mars seem possible then why not !!! Our solar system's planets and moons in the galaxy can be interconnected several decades later & AT THAT TIME, we eventually need more AS Numbers.

