TWNOG

Taiwan Network Operators Group

11th TWNIC IP Open Policy Meeting
2008/11/12, Taipei

4-Byte AS Number Migration Suggestion

Agenda

- Quickly Recap
- Assumptions before implementation
- Common BGP deployments
- 4-byte AS number implementation
- Operational issues
- Summary

Quickly Recap

- Changes in BGP 4-byte AS:
  - New BGP capability advertisement
    - Capability code: 65
    - Capability length: 4
  - Attribute update
    - AS_PATH attribute carry 4 byte AS
    - AGGREGATOR attribute carry 4 byte AS
  - New attribute
    - Add AS4_PATH attribute
      - optional, transitive
      - Construct from AS_PATH attribute by 4-byte AS enabled router when face to 2-byte AS only router
    - Add AS4_AGGREGATOR attribute
      - optional, transitive
      - Construct from AGGREGATOR attribute by 4-byte AS enabled router when face to 2-byte AS only router

Quickly Recap

- Limitations:
  - NEW BGP speaker need a 2 byte AS to peer with the OLD BGP speaker.
    - No-mappable 4 byte AS can use “23456” as 2 byte AS
  - AS4_PATH is not compatible with:
    - AS_CONFED_SEQUENCE
    - AS_CONFED_SET
    - Must use mappable 4 byte member ASN during the migration
      - i.e. 0.64512 ~ 0.65535
- Other considerations
  - BGP community handling
    - Current AS encoding method in BGP community should be updated
    - ref: draft-rekhter-as4octet-ext-community-03.txt
  - Netflow statistics
    - Netflow version 9 support 4-byte AS number

Quickly Recap

» *From "4-Byte AS Numbers, The view from the Old BGP world, p21" by Geoff Huston, APNIC*

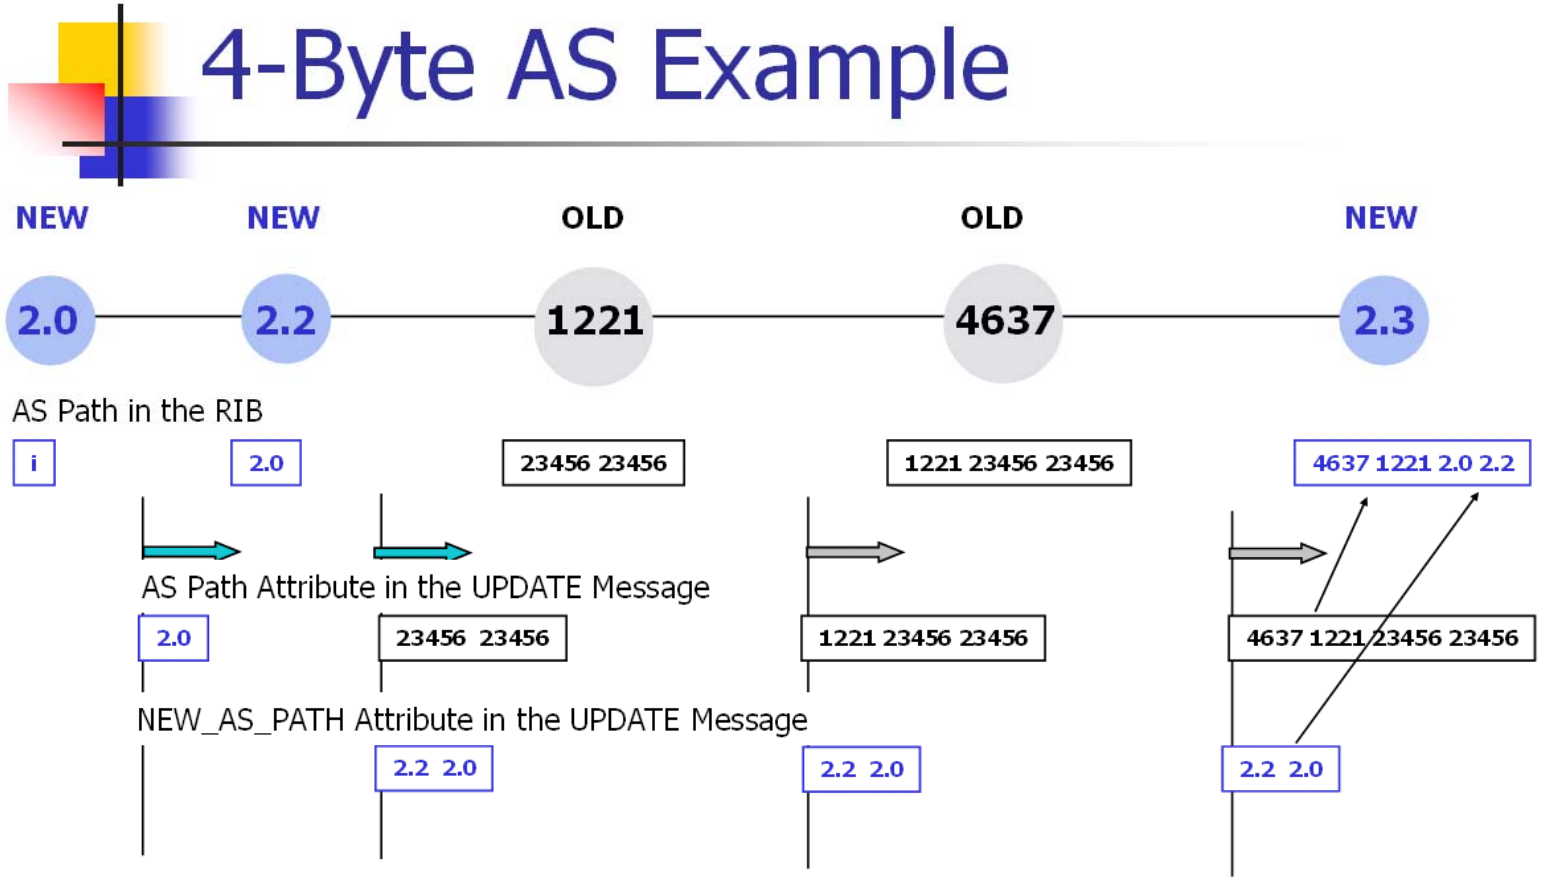

Quickly Recap

- AS4_PATH example
  - 193.5.68.0/23

```
AS path: 6830 8758 23456 I (Originator) Cluster list:  0.0.0.110
AS path:  Originator ID: 193.251.245.13
AS path:  Unrecognized Attributes: 17 bytes
AS path:  Attr flags e0 code 11: 02 03 00 00 1a ae 00 00 22 36 00 03 00 0d
```

  - Attribute Type/Length/Value
    - Attribute Type (2 octets)
      - 0xe0 => *optional, transitive*
      - 0x11 => *AS4_PATH (TYPE CODE 17)*
    - Attribute Length (Variable)
    - AS Path Segments (Variable), which includes
      - Segment Type (1 octet)
        - » 0x02 => *AS Sequence*
      - Segment Length (1 octet)
        - » 0x03 => *AS numbers in the AS Sequence*
      - Value (Variable)
        - » 0x00 00 1a ae => 6830 *(The AS who generated the AS4_PATH)*
        - » 0x00 00 22 36 => 8758
        - » 0x00 03 00 0d => 3.13

Quickly Recap

- AS4_PATH example
  - 195.47.195.0/24

```
AS path: 3257 8495 23456 I Unrecognized Attributes: 9 bytes
AS path:  Attr flags e0 code 11: 02 01 00 03 00 10
```

  - Attribute Type/Length/Value
    - Attribute Type (2 octets)
      - 0xe0 => *optional, transitive*
      - 0x11 => *AS4_PATH (TYPE CODE 17)*
    - Attribute Length (Variable)
    - AS Path Segments (Variable), which includes
      - Segment Type (1 octet)
        - » 0x02 => *AS Sequence*
      - Segment Length (1 octet)
        - » 0x01 => *AS numbers in the AS Sequence*
      - Value (Variable)
        - » 0x00 03 00 10 => 3.16 *(The AS who generated the AS4_PATH)*

Quickly Recap

- AS4_PATH example
  - 2001:7fb:ff00::/48

```
AS path: 3549 1103 1125 23456 12654 I (Originator) Cluster list:  193.251.245.70
AS path:  Originator ID: 193.251.151.105
AS path:  Unrecognized Attributes: 17 bytes
AS path:  Attr flags e0 code 11: 02 03 00 00 04 65 00 03 00 05 00 00 31 6e
```

  - Attribute Type/Length/Value
    - Attribute Type (2 octets)
      - 0xe0 => *optional, transitive*
      - 0x11 => *AS4_PATH (TYPE CODE 17)*
    - Attribute Length (Variable)
    - AS Path Segments (Variable), which includes
      - Segment Type (1 octet)
        - » 0x02 => *AS Sequence*
      - Segment Length (1 octet)
        - » 0x03 => *AS numbers in the AS Sequence*
      - Value (Variable)
        - » 0x00 00 04 65 => 1125 *(The AS who generated the AS4_PATH)*
        - » 0x00 03 00 05 => 3.5
        - » 0x00 00 31 6e => 12654

Quickly Recap

» *From "4-Byte AS Numbers, The view from the Old BGP world, p23" by Geoff Huston, APNIC*

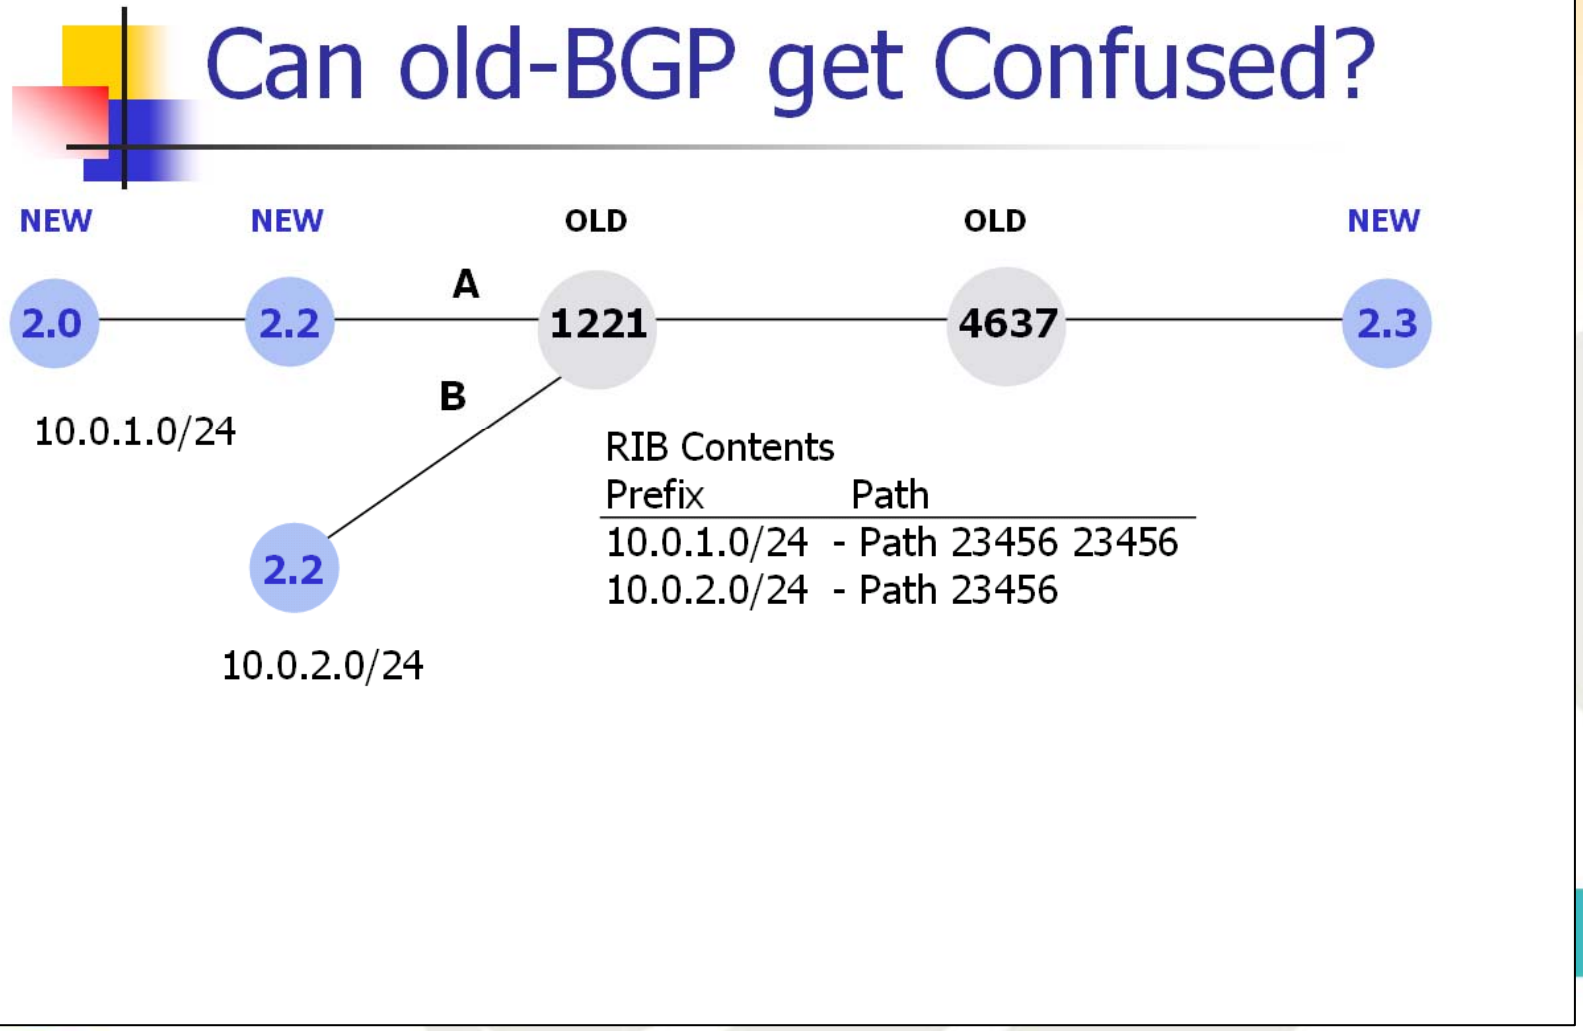

Quickly Recap

» *From "4-Byte AS Numbers, The view from the Old BGP world, p24" by Geoff Huston, APNIC*

NO! BGP Nexthop is the key!

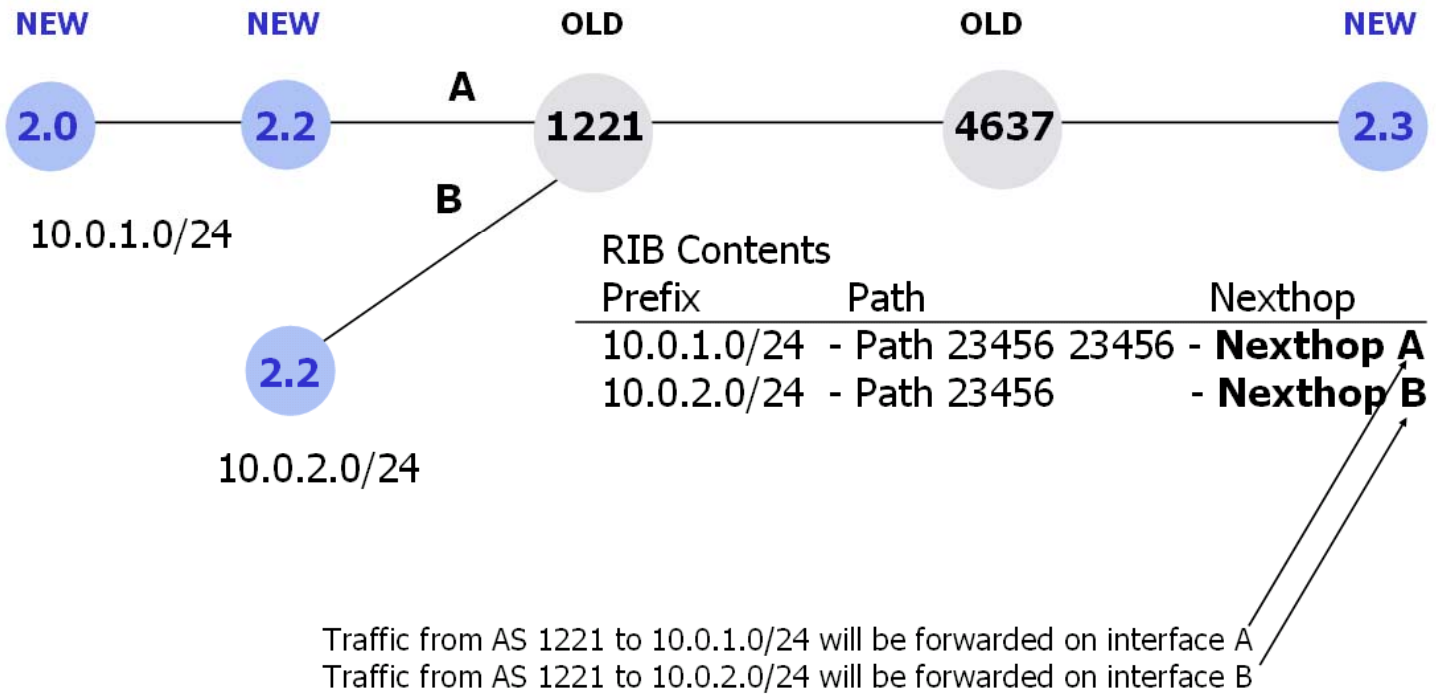

This is standard BGP behaviour – nothing changes here for BGP as it is used today

Assumptions before implementation

- Everyone (uplink ISP, peering ISP, transit customers and yourself)
  - enable BGP
  - has at least one 2 bytes ASN before 2 byte ASN exhausting
- Following common BGP deployments are considered:
  - Only border routers enable BGP
  - All routers enable BGP and full-mesh with each other
  - All routers enable BGP and implement Route-Reflector
  - All routers enable BGP with BGP confederation
  - There are more than one AS in the same ISP domain

Common BGP deployment

- Only border routers enable BGP

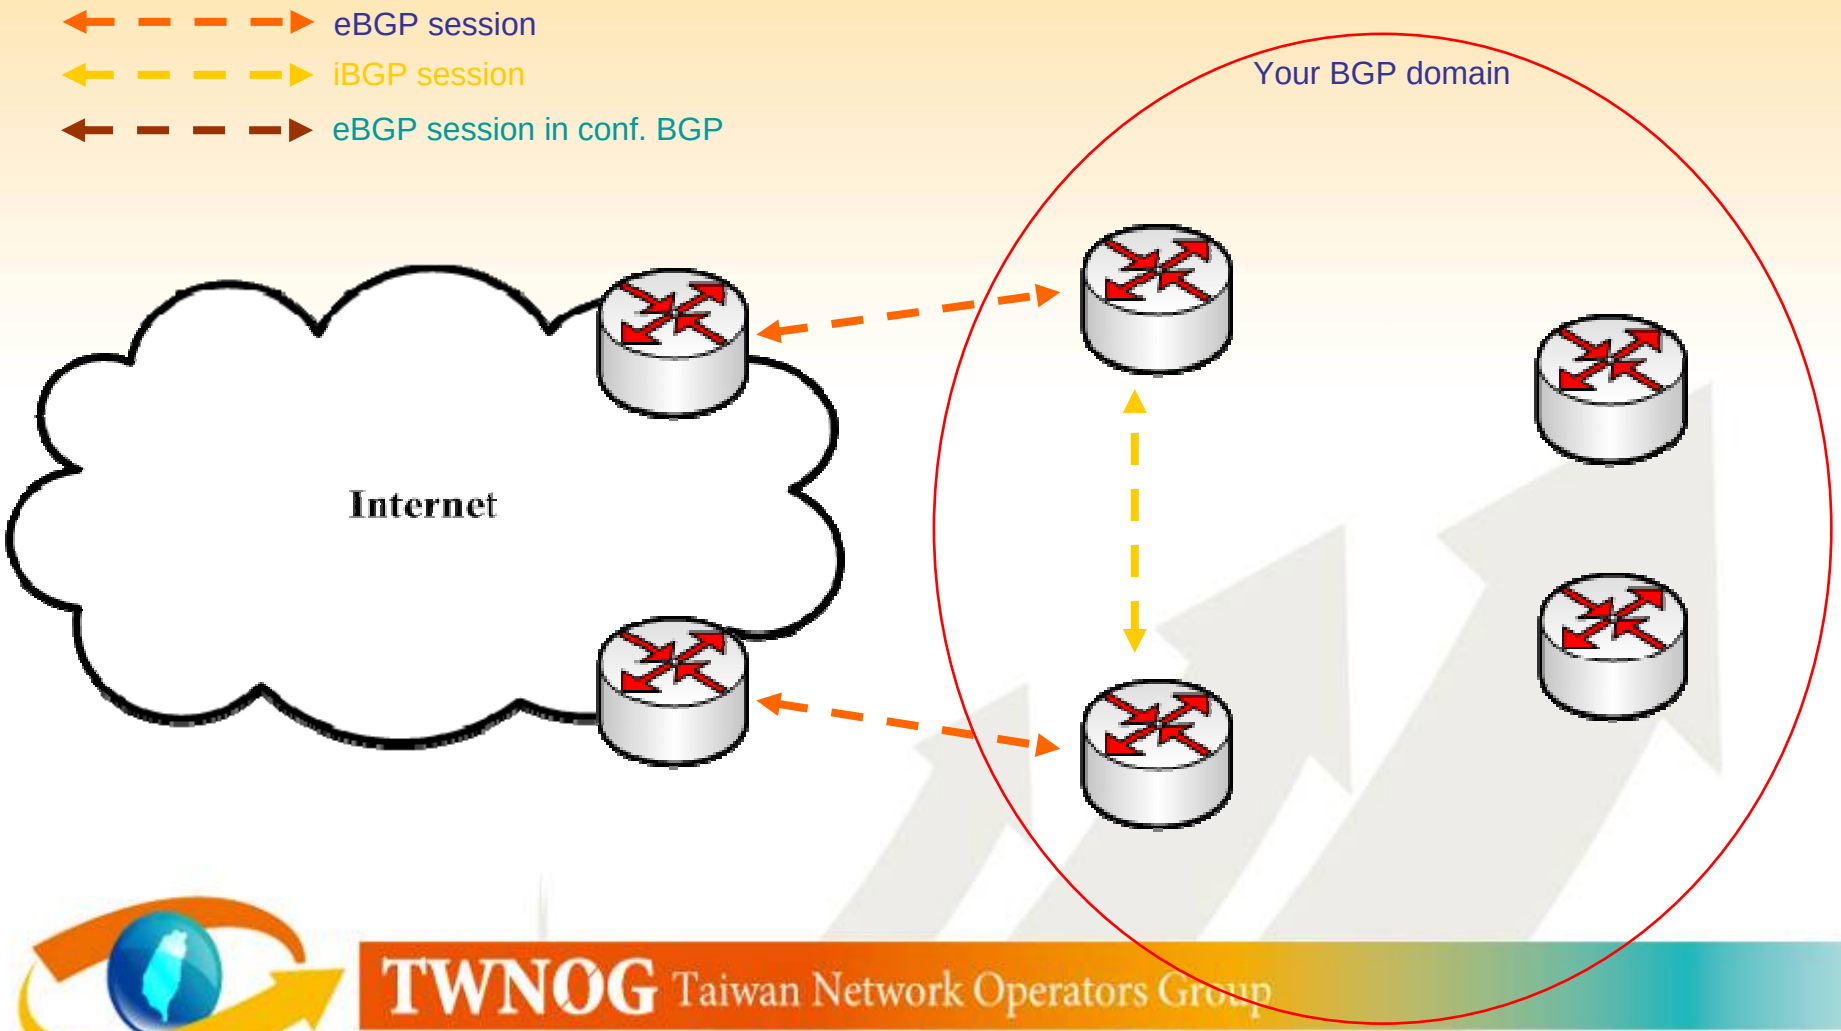

Common BGP deployment

- All routers enable BGP and full-mesh with each other

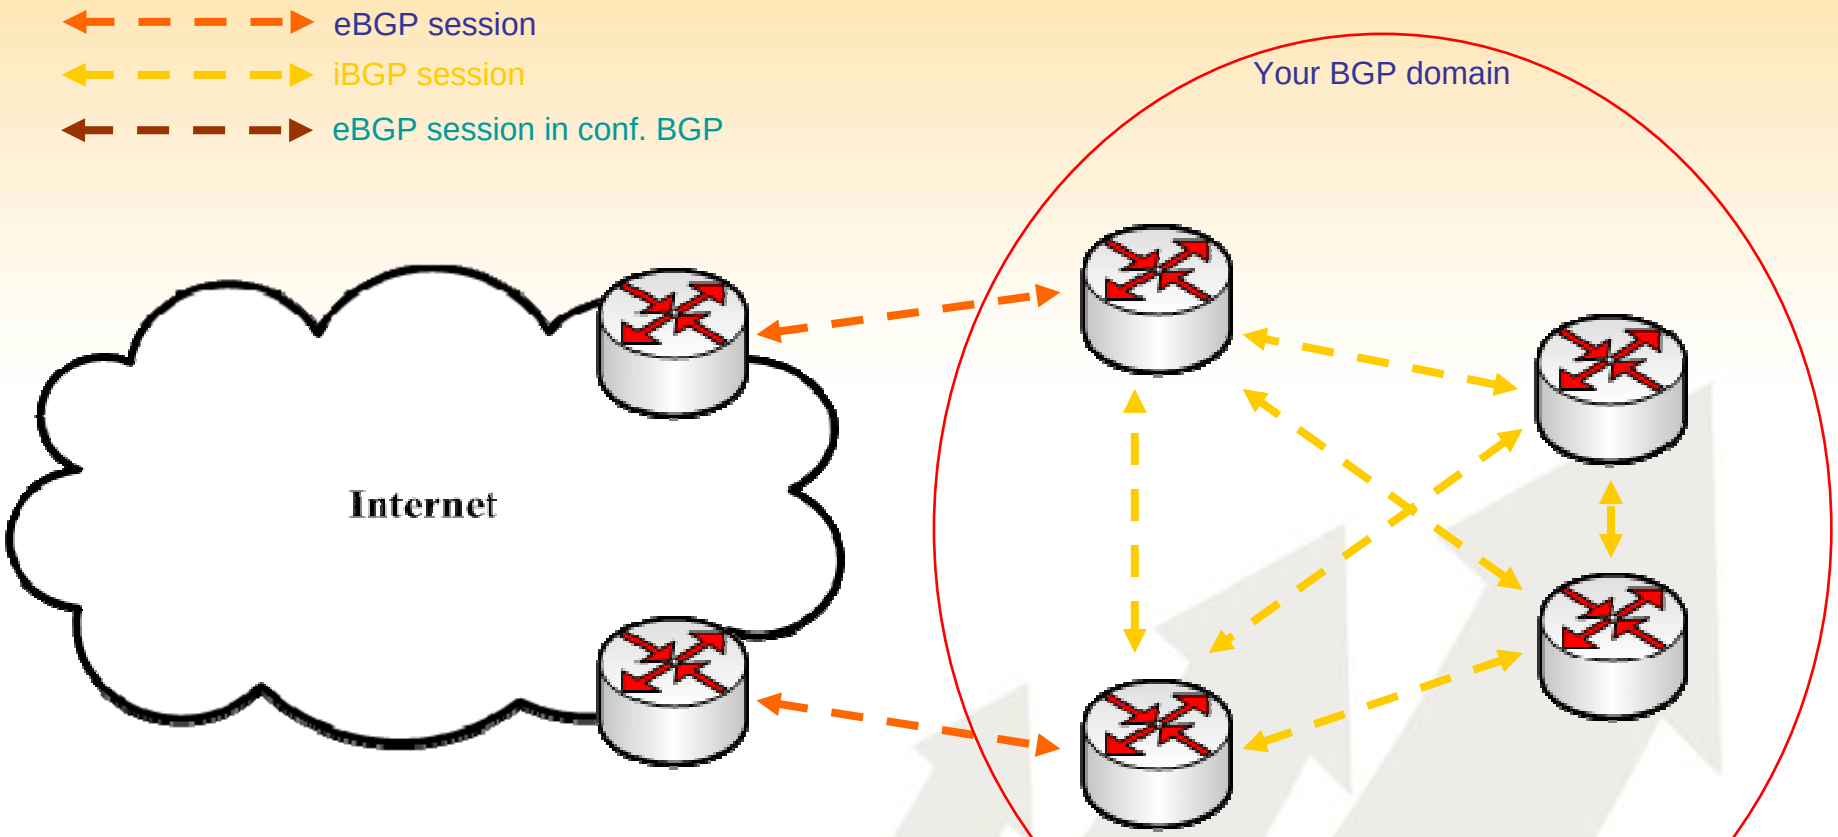

Common BGP deployment

- All routers enable BGP and implement Route-Reflector

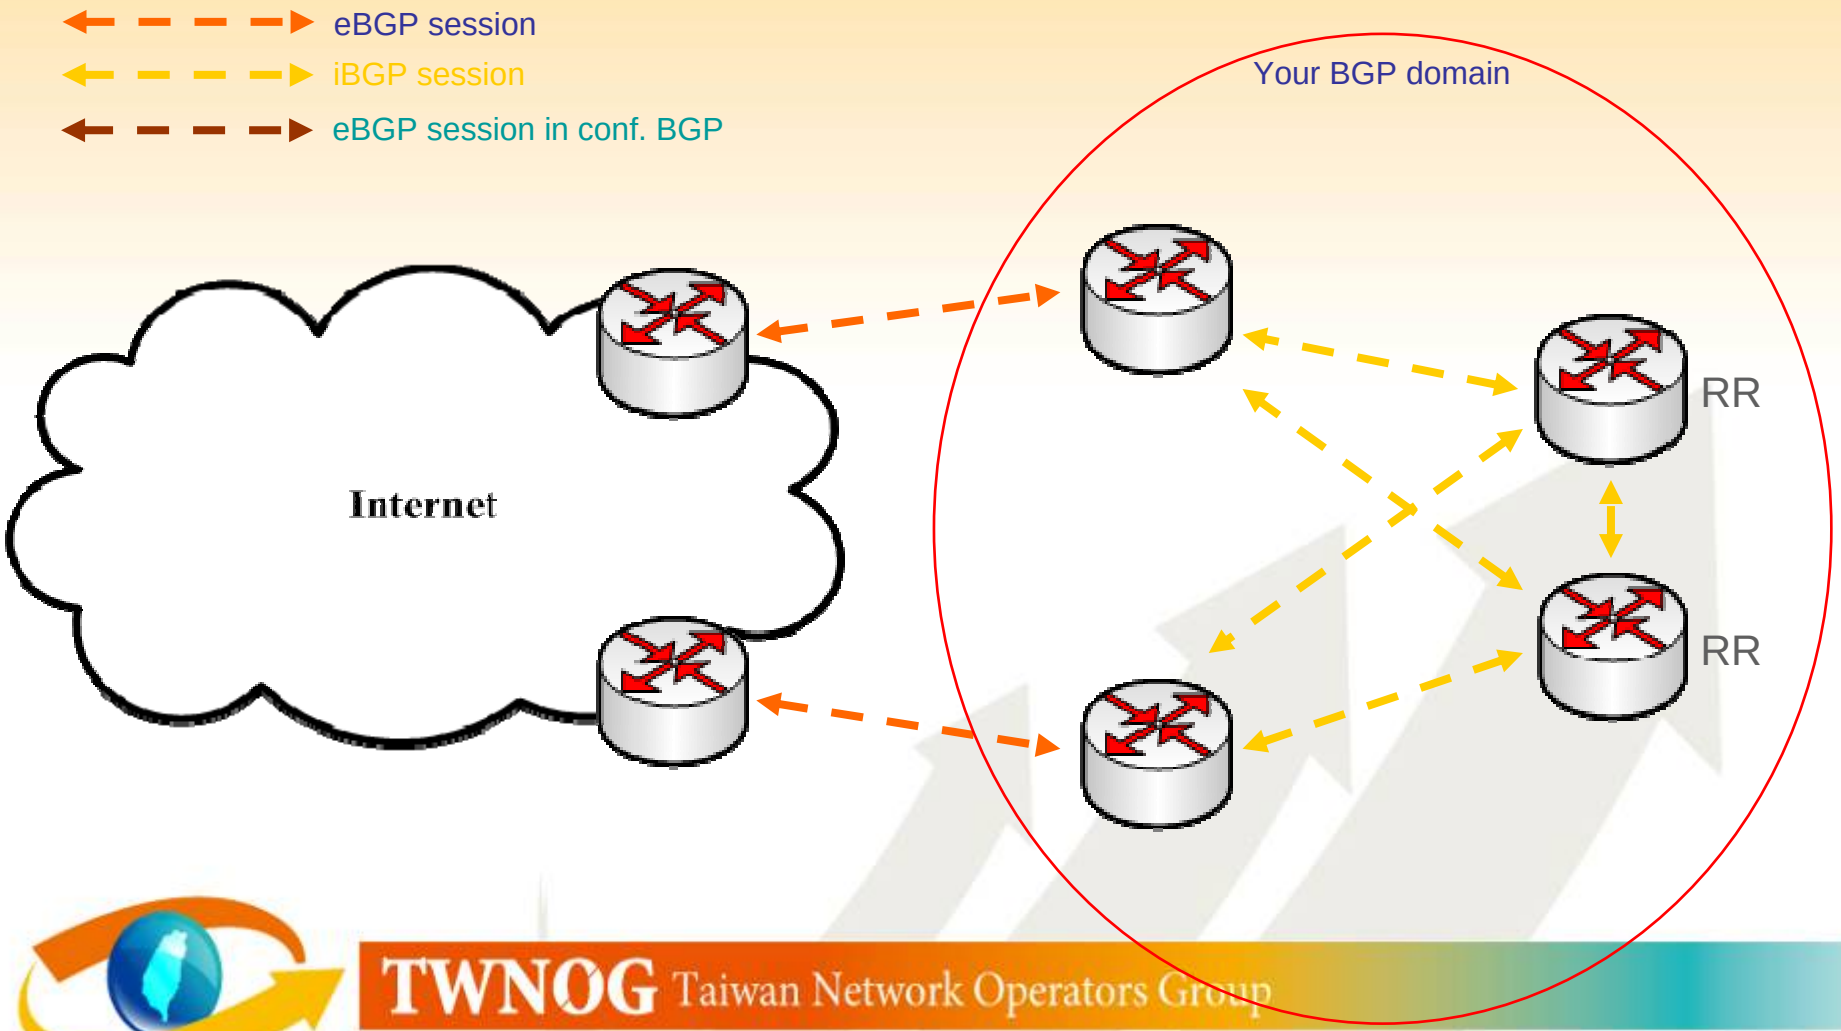

Common BGP deployment

- All routers enable BGP with BGP confederation

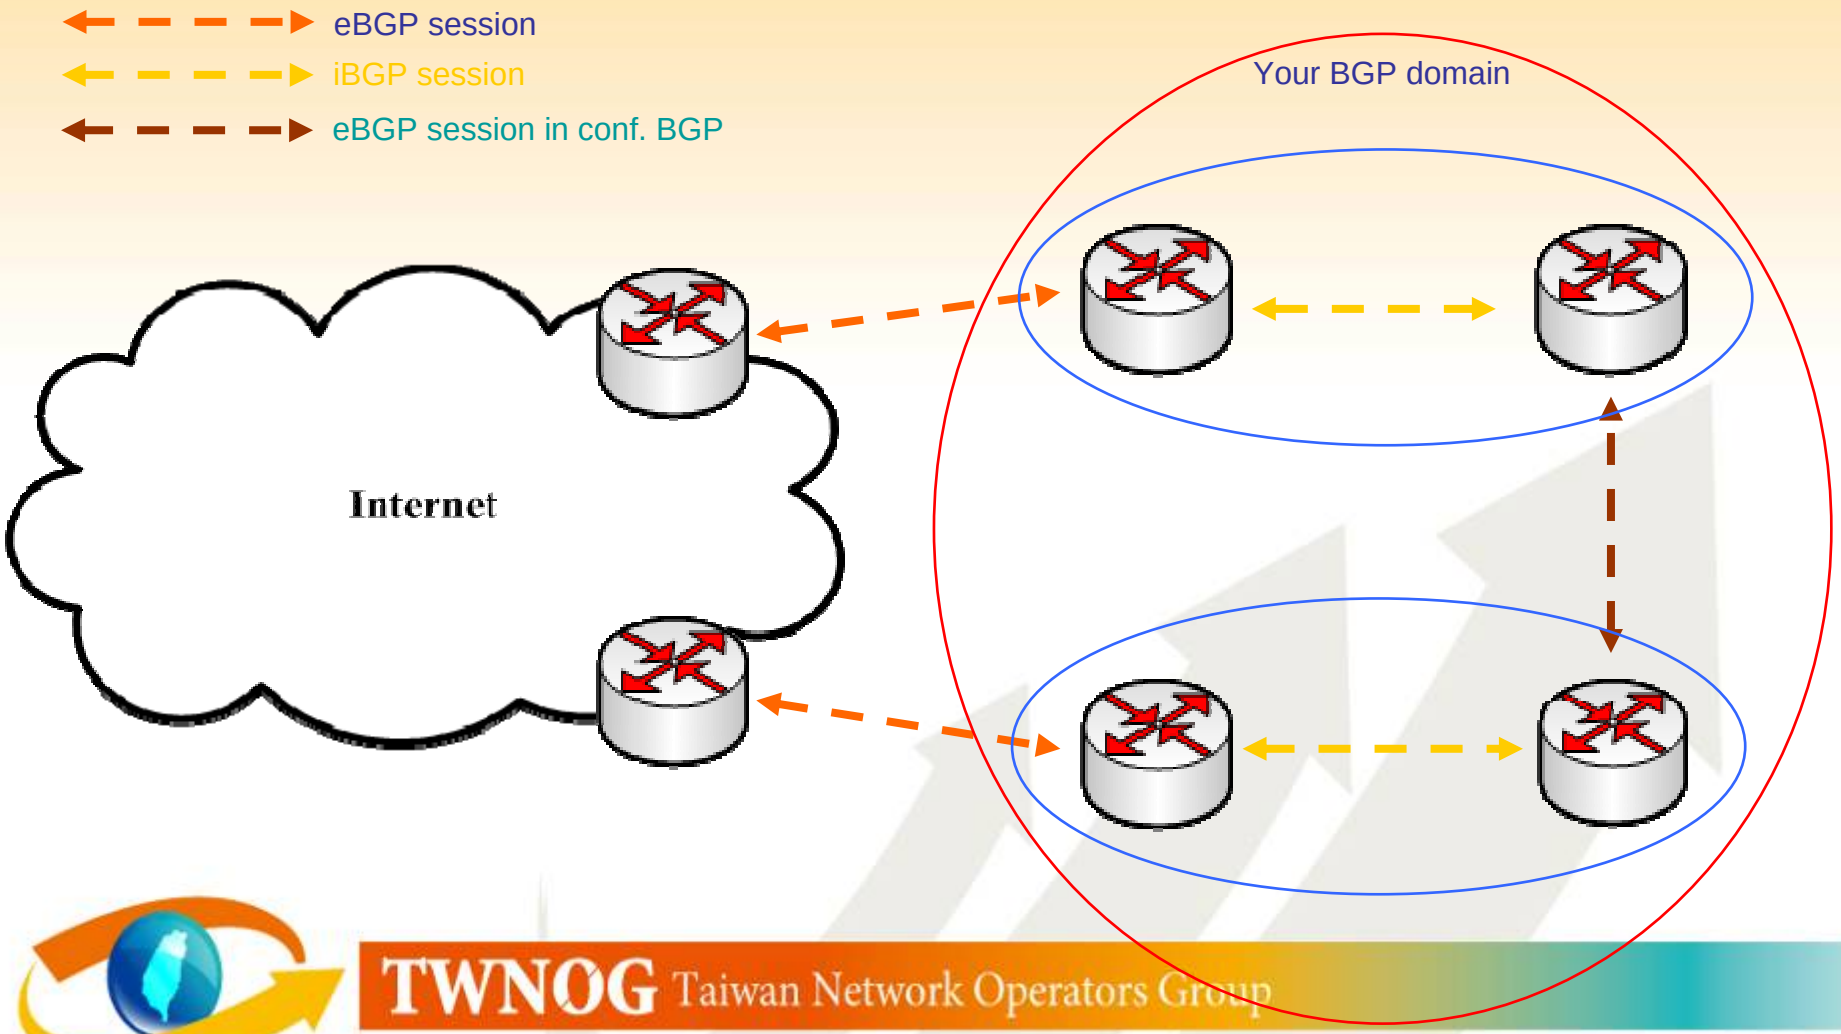

Common BGP deployment

- There are more than one AS in the same ISP domain

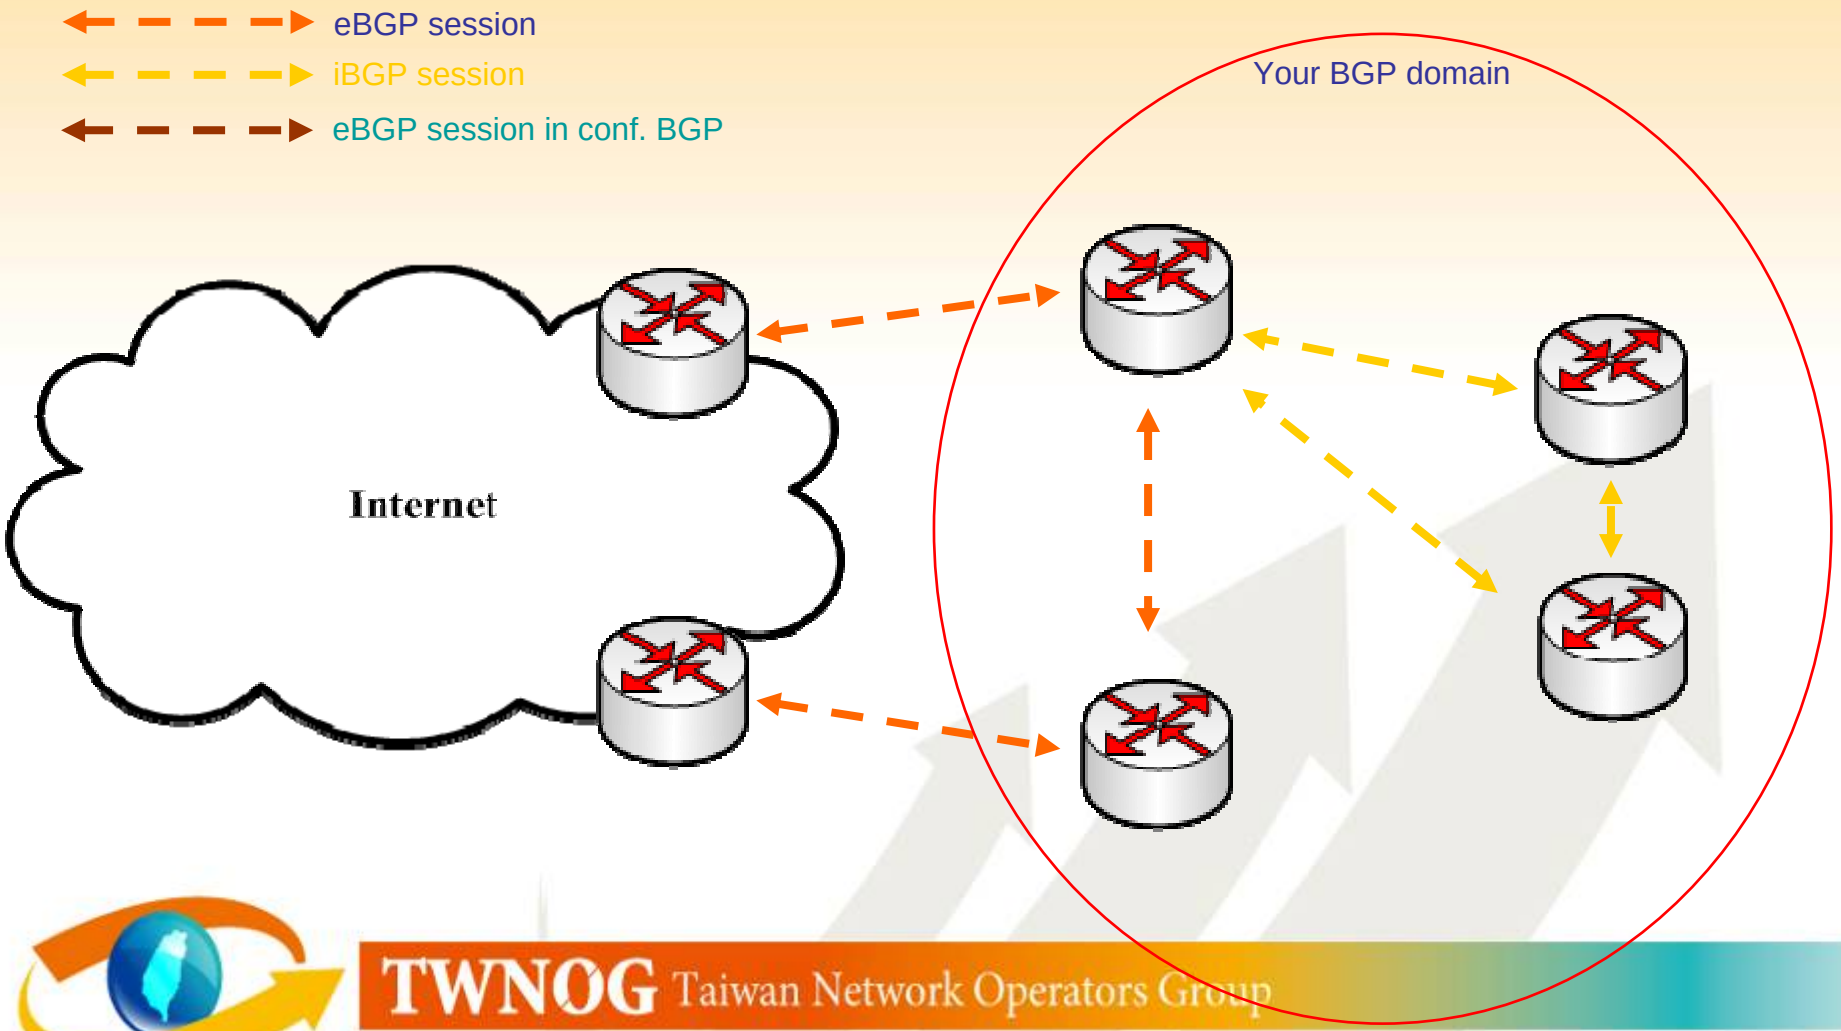

4-byte AS number implementation

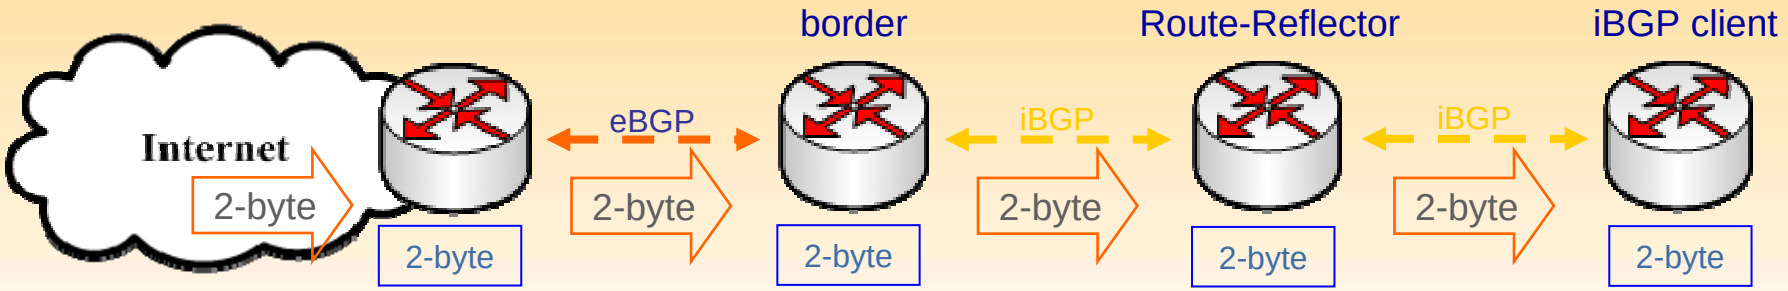

- What we have to do?
  - Arrange a perfect plan first
  - Upgrade the router operating system (ROS) then
- What is the safe implementation approach?
  - from border router?
  - from Route-Reflector?
    - Route-Reflector is so important, it seems not a good choice to upgrade Route-Reflector first
  - from iBGP client?

4-byte AS number implementation

- Example 1-1-1
  - Your partner is NOT 4-byte AS ready yet
    - 1st: You upgrade your border router ROS

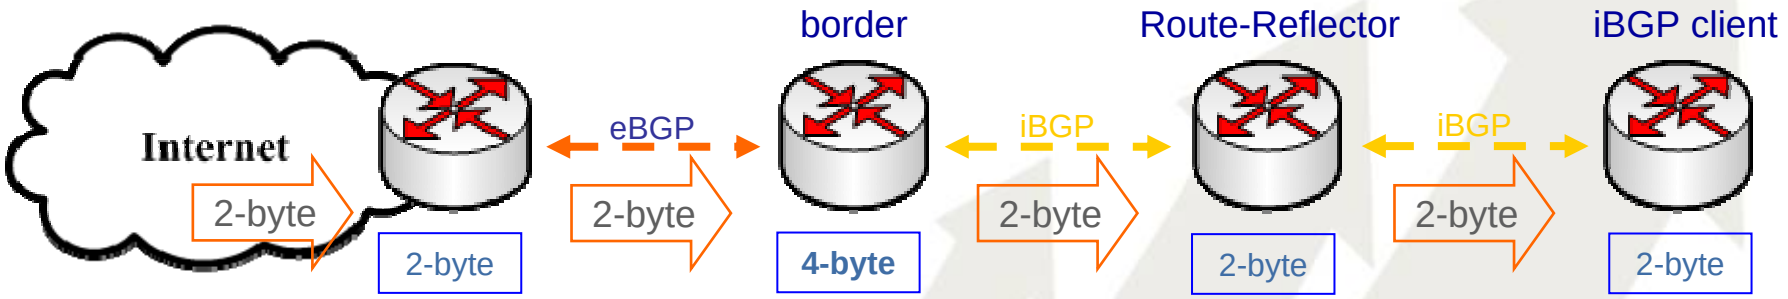

4-byte AS number implementation

- Example 1-1-2
  - Your partner is NOT 4-byte AS ready yet
    - 1st: You upgrade your border router ROS
    - 2nd: You upgrade your iBGP client ROS

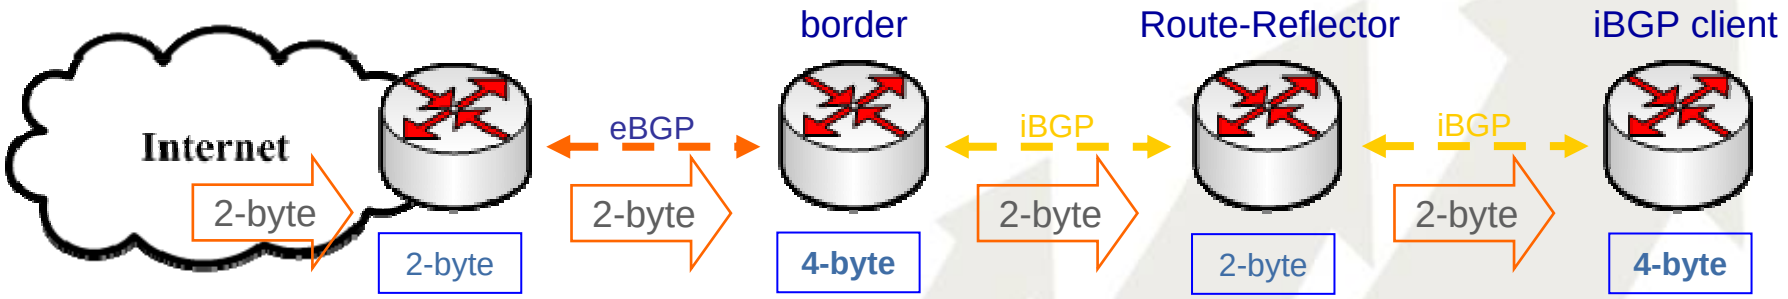

4-byte AS number implementation

- Example 1-1-3
  - Your partner is NOT 4-byte AS ready yet
    - 1st: You upgrade your border router ROS
    - 2nd: You upgrade your iBGP client ROS
    - 3rd: You upgrade your Route-Reflector ROS

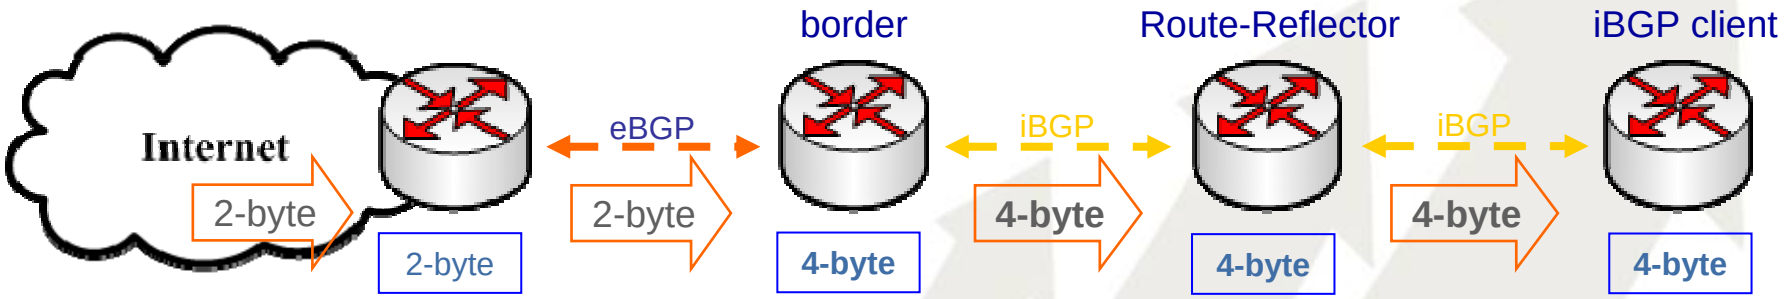

4-byte AS number implementation

- Example 1-2-1
  - Your partner is NOT 4-byte AS ready yet
    - 1st: You upgrade your iBGP client ROS

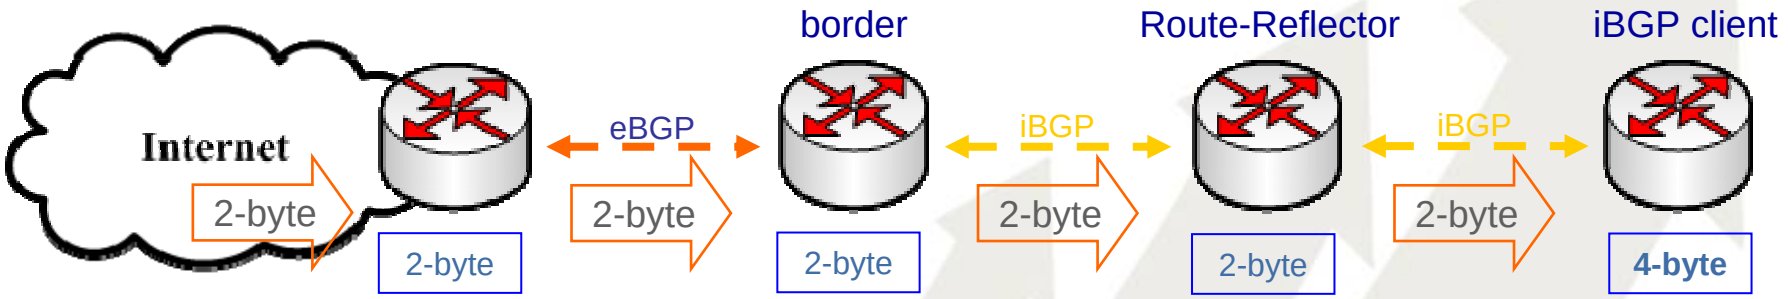

4-byte AS number implementation

- Example 1-2-2
  - Your partner is NOT 4-byte AS ready yet
    - 1st: You upgrade your iBGP client ROS
    - 2nd: You upgrade your border router ROS

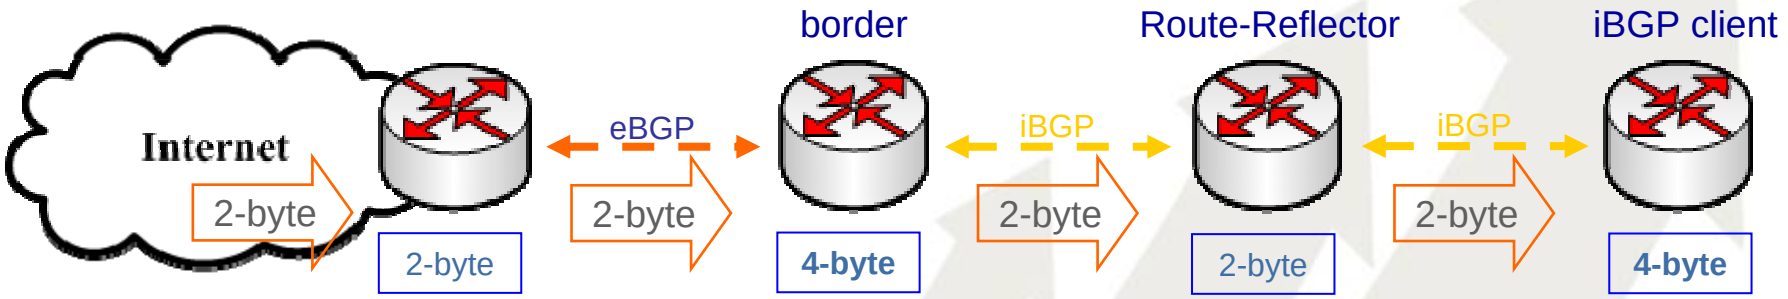

4-byte AS number implementation

- Example 1-2-3
  - Your partner is NOT 4-byte AS ready yet
    - 1st: You upgrade your iBGP client ROS
    - 2nd: You upgrade your border router ROS
    - 3rd: You upgrade your Route-Reflector ROS

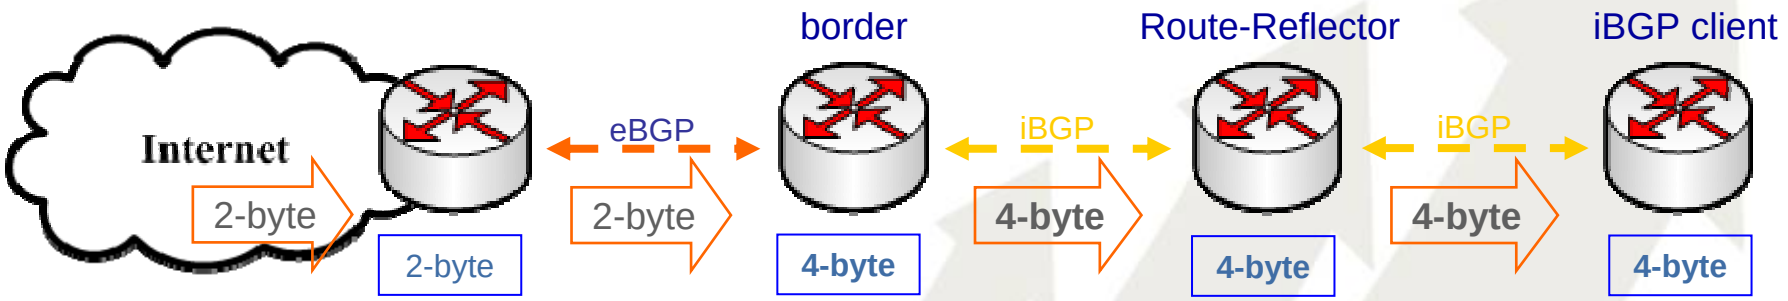

4-byte AS number implementation

- Example 2-1-1
  - Your partner is 4-byte AS ready
    - 1st: You upgrade your border router ROS

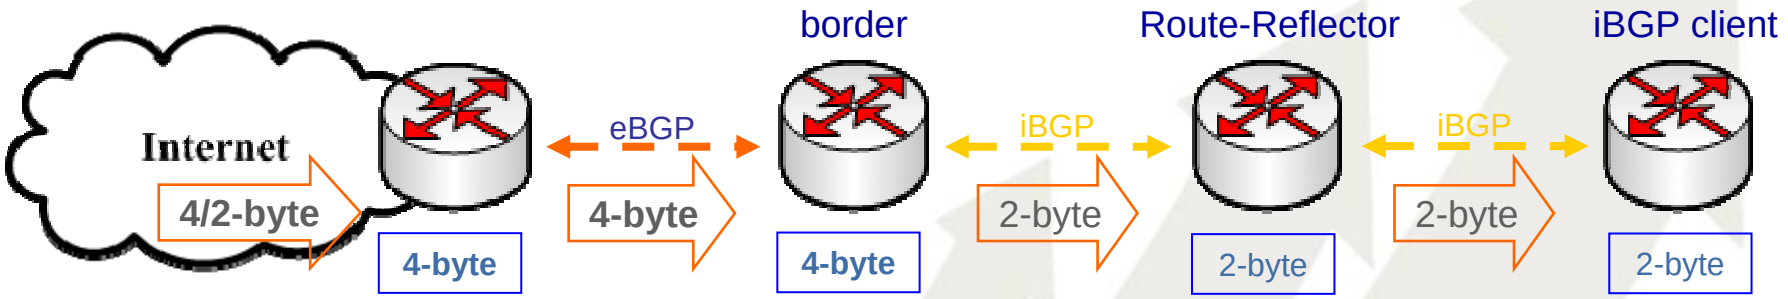

4-byte AS number implementation

- Example 2-1-2
  - Your partner is 4-byte AS ready
    - 1st: You upgrade your border router ROS
    - 2nd: You upgrade your iBGP client ROS

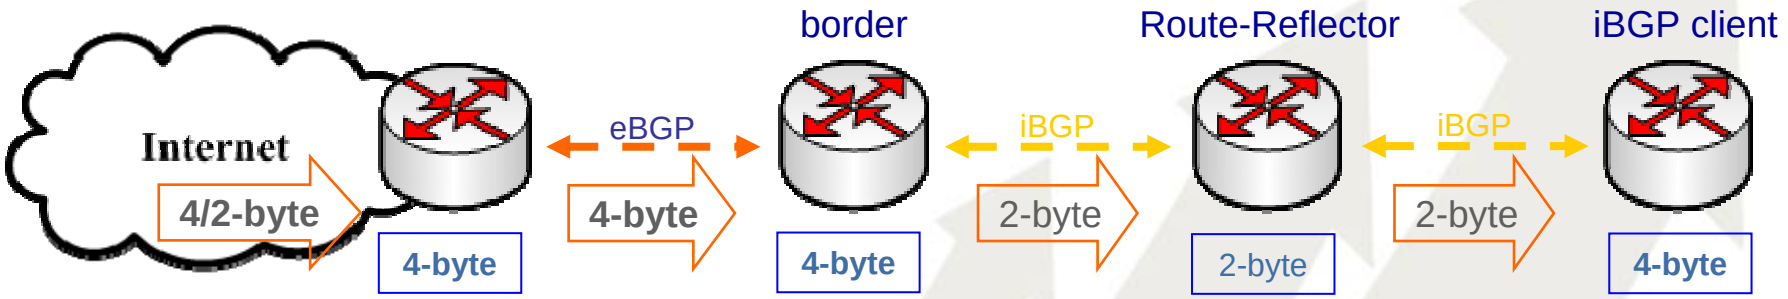

4-byte AS number implementation

- Example 2-1-3
  - Your partner is 4-byte AS ready
    - 1st: You upgrade your border router ROS
    - 2nd: You upgrade your iBGP client ROS
    - 3rd: You upgrade your Route-Reflector ROS

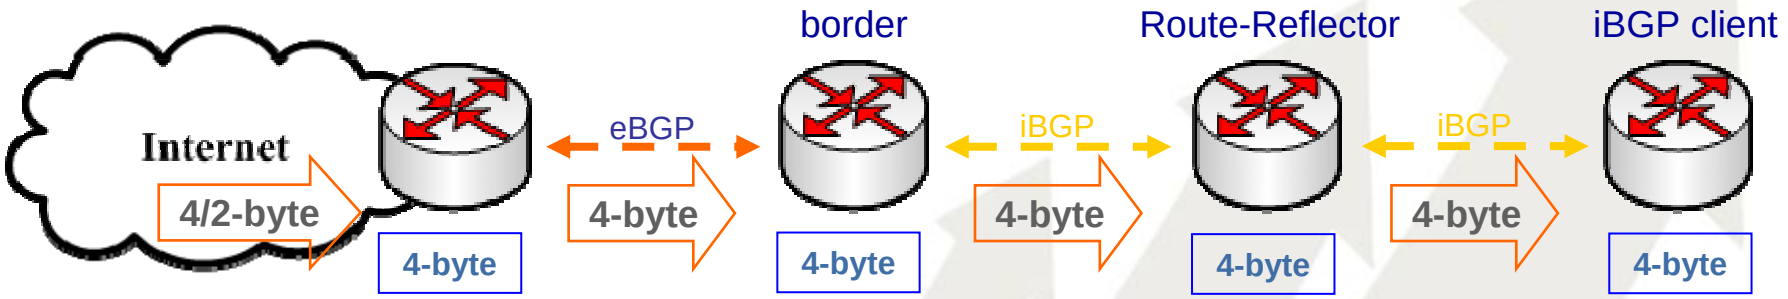

4-byte AS number implementation

- Example 2-2-1
  - Your partner is 4-byte AS ready
    - 1st: You upgrade your iBGP client ROS

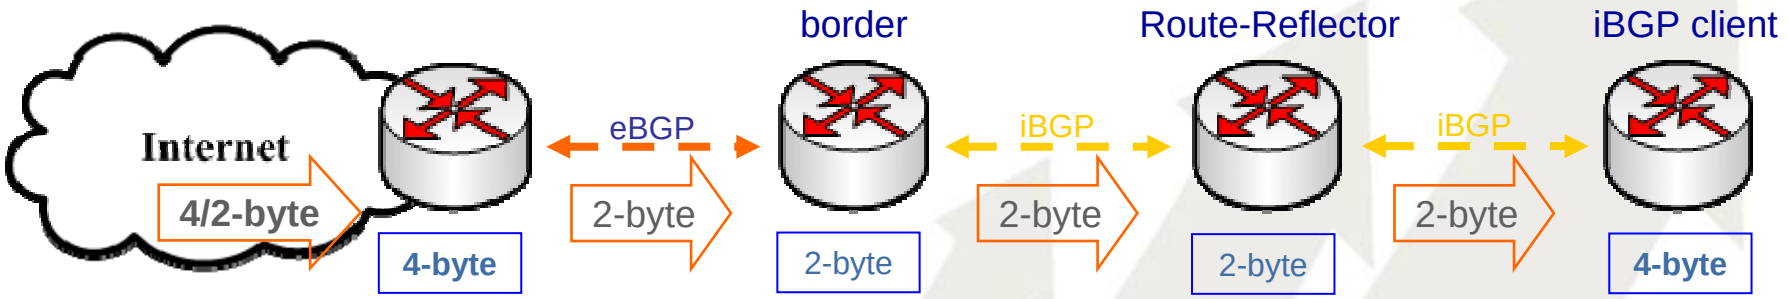

4-byte AS number implementation

- Example 2-2-2
  - Your partner is 4-byte AS ready
    - 1st: You upgrade your iBGP client ROS
    - 2nd: You upgrade your border router ROS

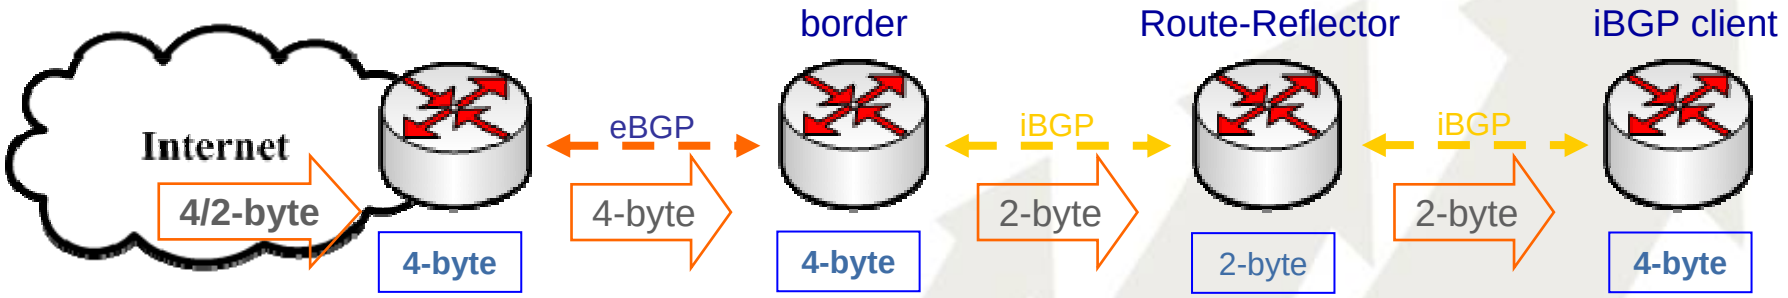

4-byte AS number implementation

- Example 2-2-3
  - Your partner is 4-byte AS ready
    - 1st: You upgrade your iBGP client ROS
    - 2nd: You upgrade your border router ROS
    - 3rd: You upgrade your Route-Reflector ROS

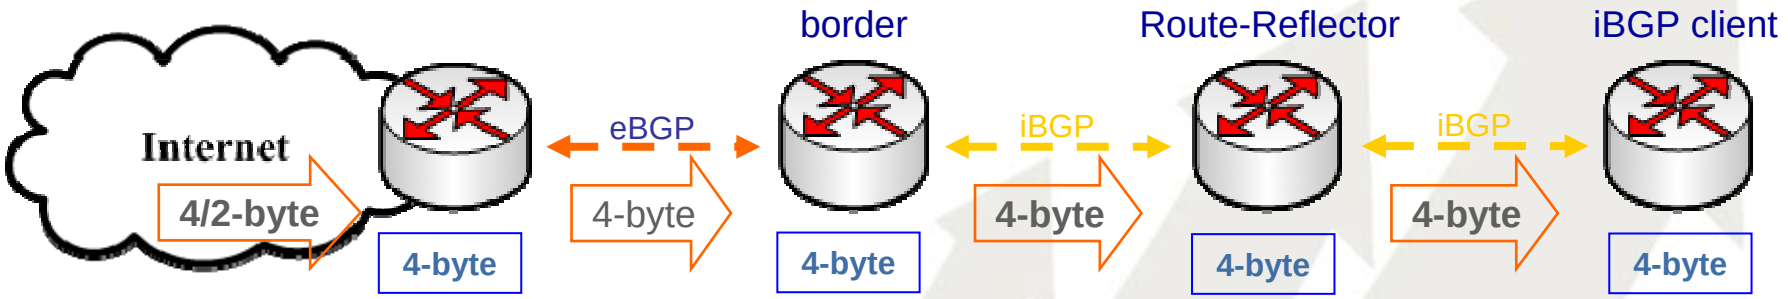

Summary

- It is fine to upgrade ROS no matter
  - from border router, iBGP client then Route-Reflector
  - from iBGP client, border router then Route-Reflector
- Thank you
- But…..

Are they still good approaches when consider operational issues?

Operational issues

- AS-PATH issues on border router
  - AS-PATH is a very important attribute for BGP policy design
    - Allow or deny BGP routes
    - Set local preference
    - Set BGP community
  - If the border router did not support 4-byte AS number, we can NOT handle 4-byte AS BGP routes by AS-PATH attribute
    - Can NOT permit/deny BGP routes by AS path information
      - Old ROS can not recognize AS4_PATH attribute either
    - Can NOT set BGP local preference on border router by AS path information
    - Can NOT set BGP community on border router by AS path information

Operational issues

- Upgrade Router Operating System issues
  - It is not a easy work
    - Heavy loading in testing the new ROS
    - Much different from patch/upgrade personal computer OS
  - Upgrade ROS then reboot router impact the network and SLA
    - Especially in those network environments without appropriate redundant design
    - Higher SLA requirement is a challenge: 99.99% even 99.999%
  - With a serious, detail plan, upgrade all router ROS will spend many months even more than one year.
    - Upgrade processes are risky
    - Any un-conditional network event will postpone the scheduled upgrade process
  - Upgrade ROS guide in operation:
    - DON'T TOUCH YOUR ROS UNLESS YOU HAVE TO!!!!!

Summary

- If we want to consider the operational issues, to upgrade border router ROS first is a better approach
  - The AS-PATH handling will not be a problem anymore
  - It is fine to upgrade border router only
    - Old BGP router will NOT be confused by duplicate AS23456
    - We can reboot the router for serious ROS bug-fix patch only rather than just 4-byte AS upgrade
  - After the border router, the upgrade sequence could be considered by
    - (option) other eBGP router in the same autonomous domain
    - (option) CONFED border router
    - (option) iBGP client
    - (option) Route-Reflector

Reference

- TWNIC OPM
  - 8th OPM
    - 4 byte ASN的現況
    - 4-byte ASNs Test Scenarios
  - 6th OPM
    - BGP Support for Four-octet AS Number Space
- RFC
  - RFC4271 - “A Border Gateway Protocol 4 (BGP-4)”
  - RFC4893 - “BGP Support for Four-octet AS Number Space”
- NANOG
  - NANOG 39; “4-Byte AS Numbers, The view from the Old BGP world” by Geoff Huston, APNIC
- IETF draft
  - Four-octet AS Specific BGP Extended Community

Thank you

- Questions?