Asia Pacific Network Information Centre

APNIC

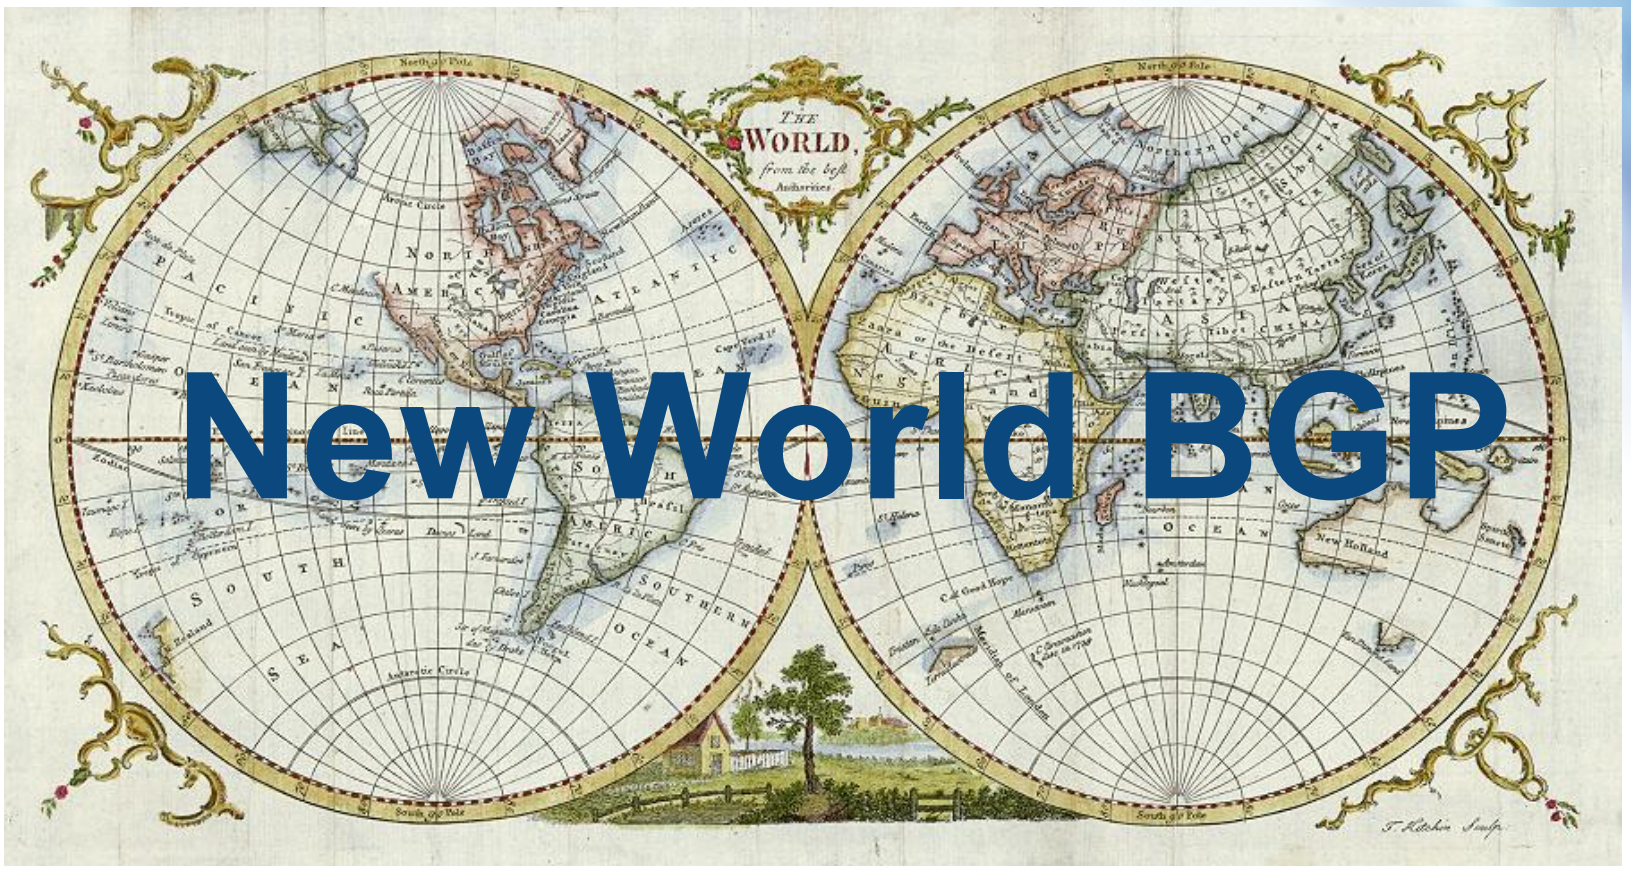

New World BGP

Geoff Huston
January2010
APNIC

Introduction

- Presenters
  - Nurul Islam Roman
    - nurul@apnic.net

16-bit AS Number Map

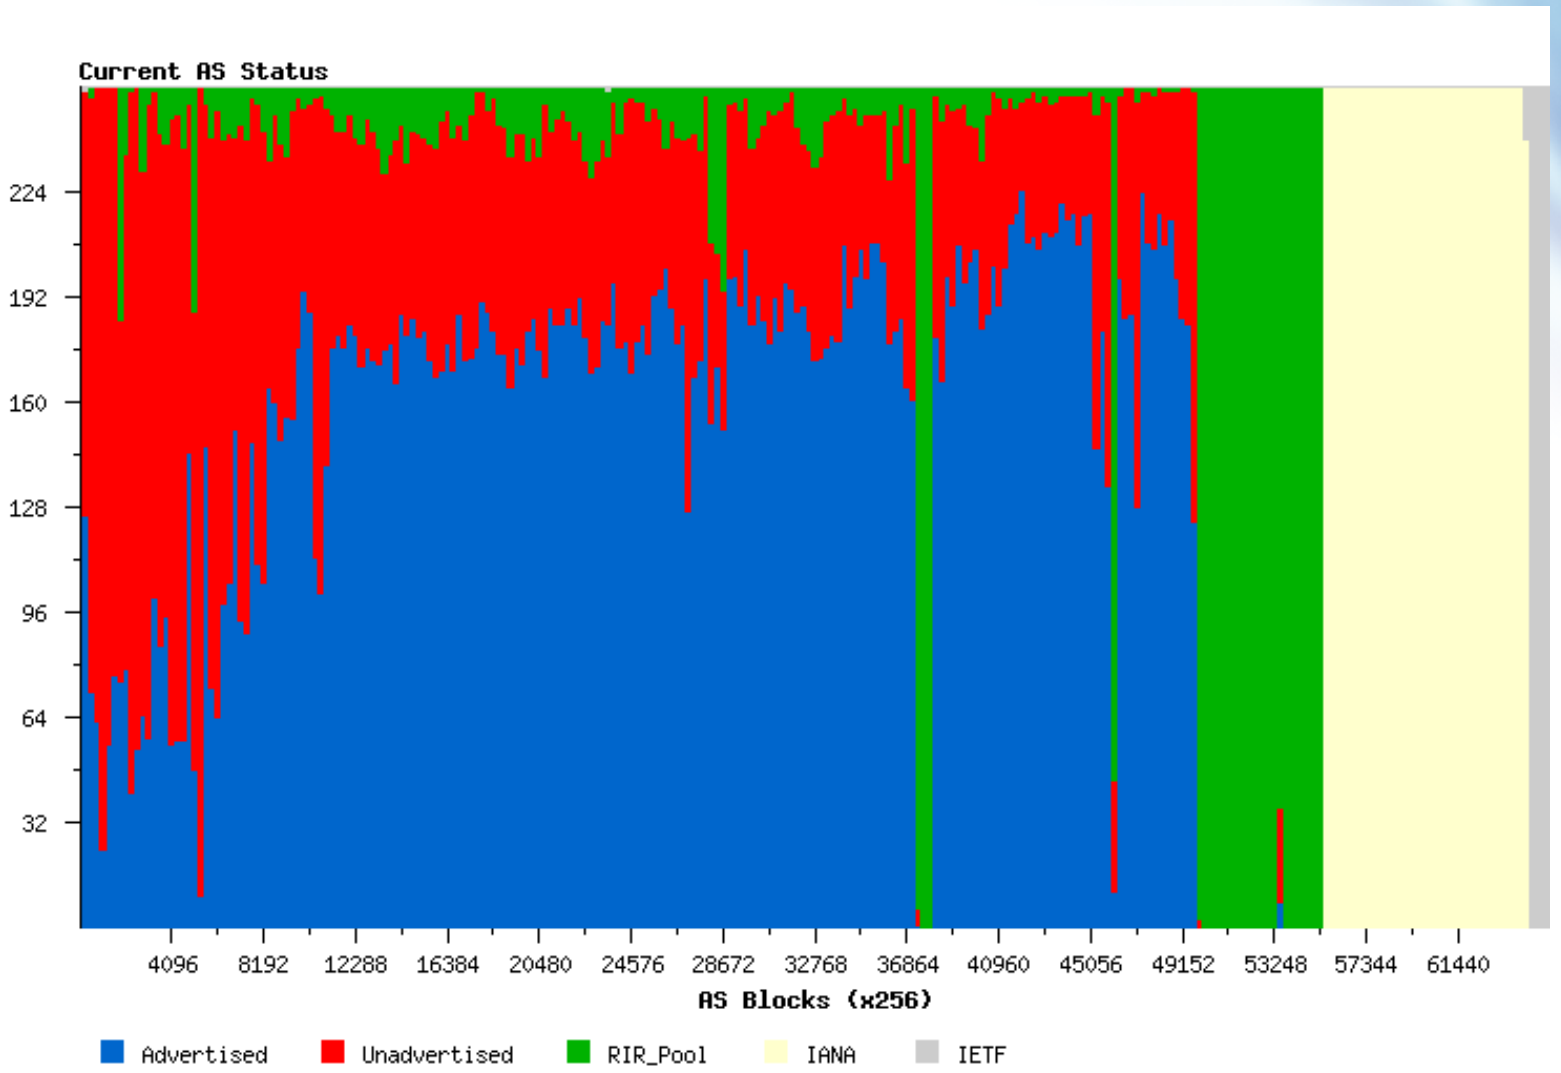

16-bit AS Number Map

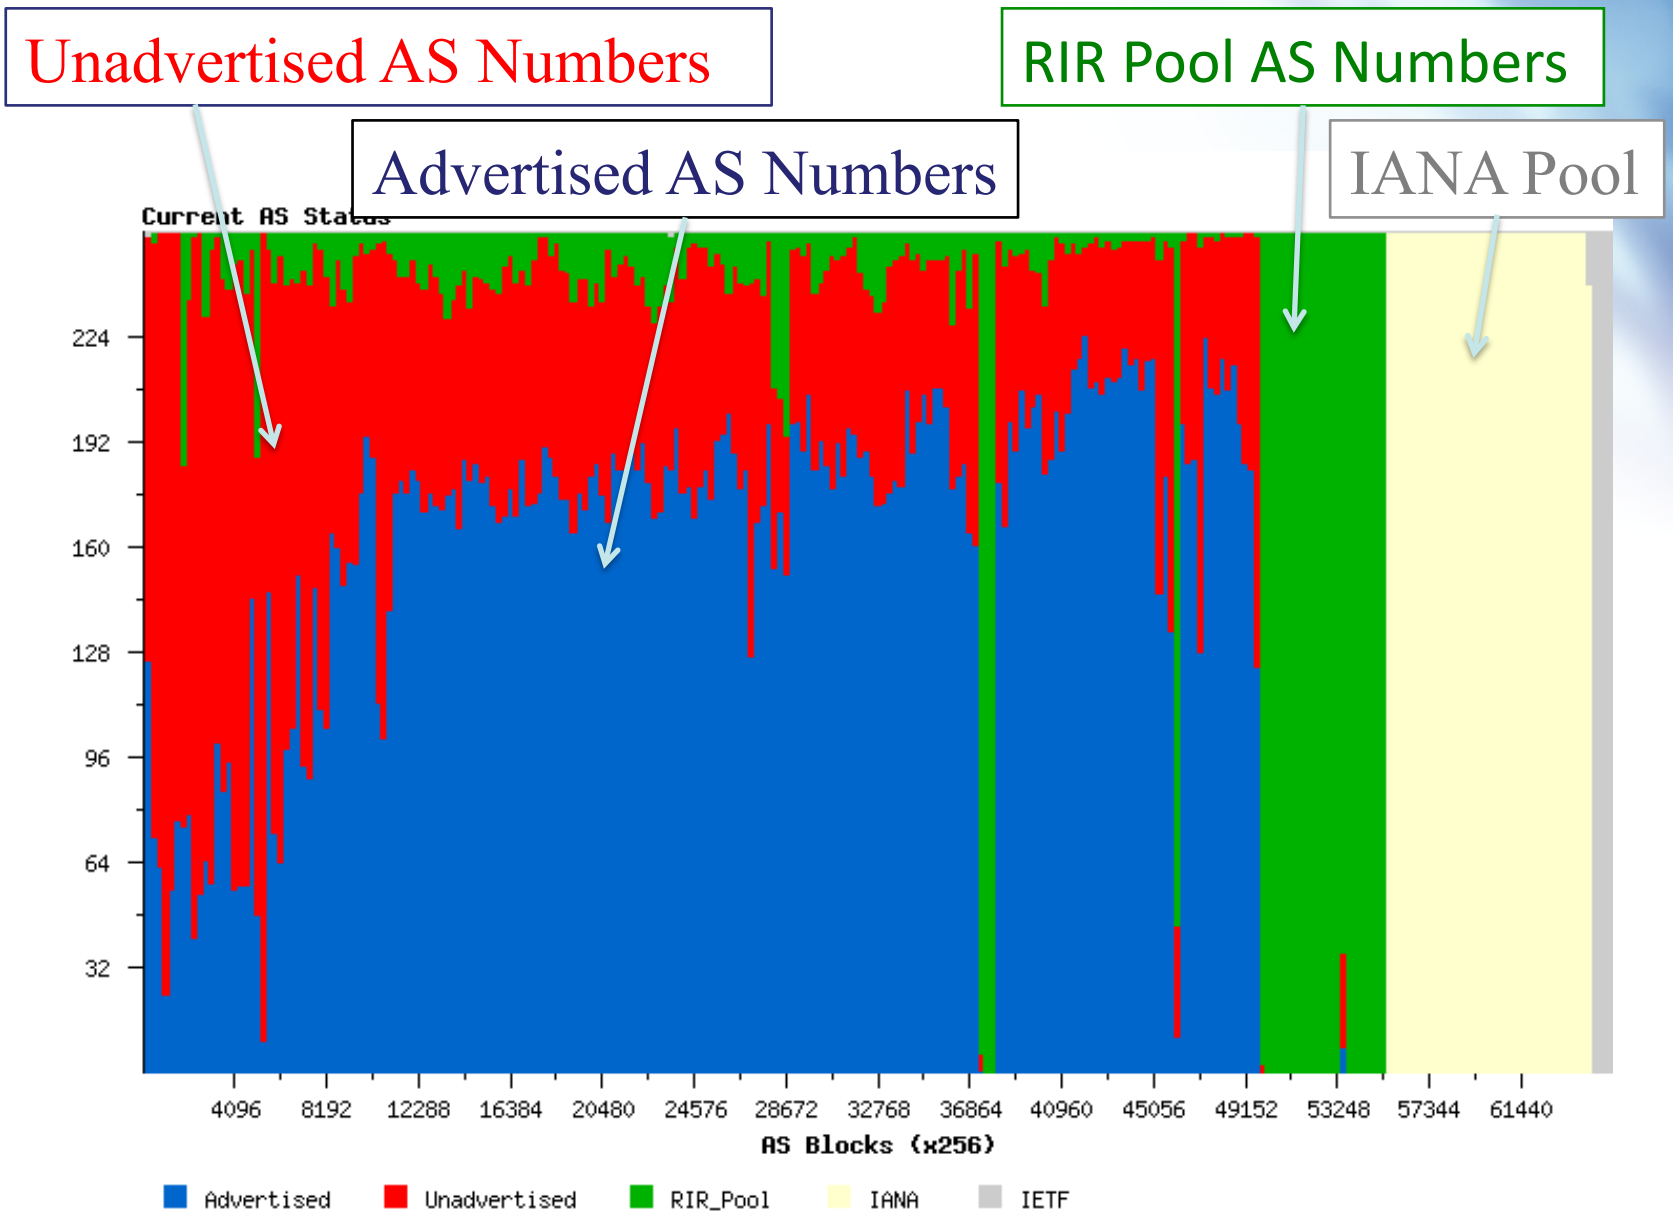

16-bit AS Number Map

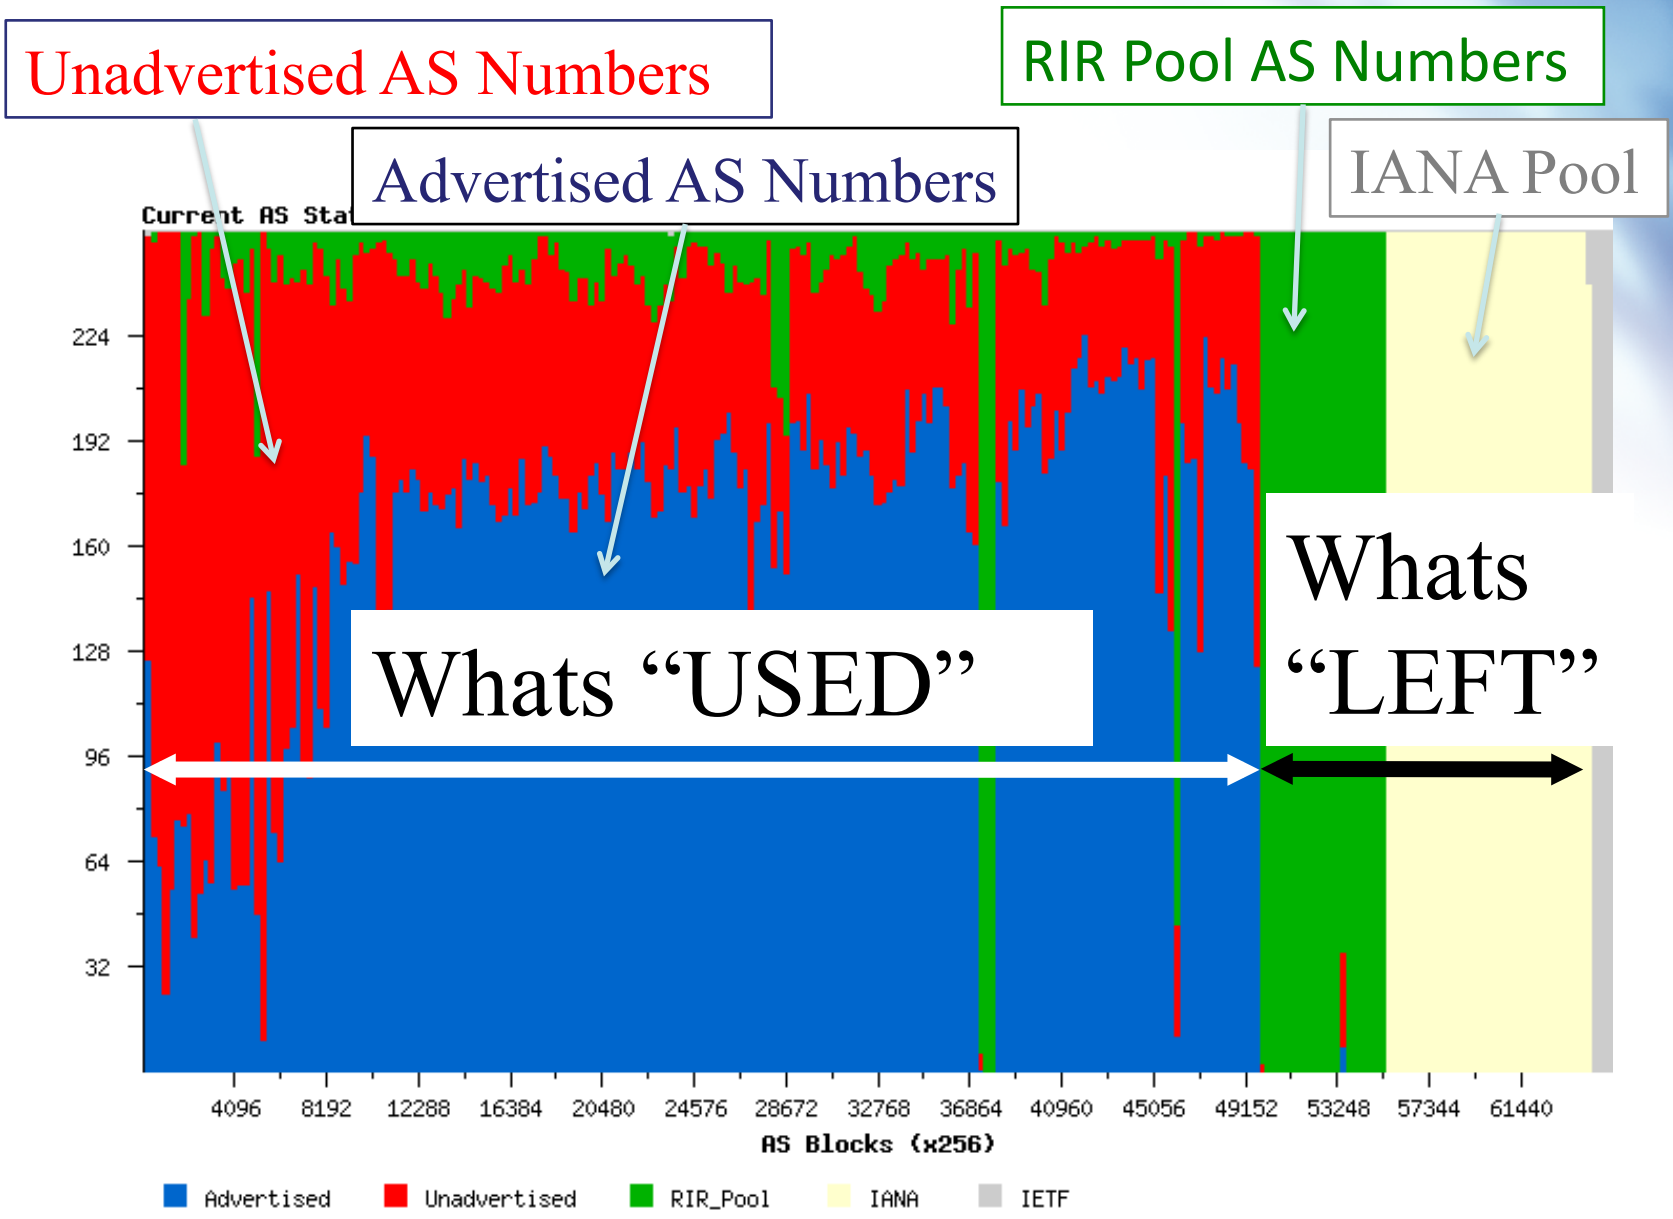

Consumption Rate

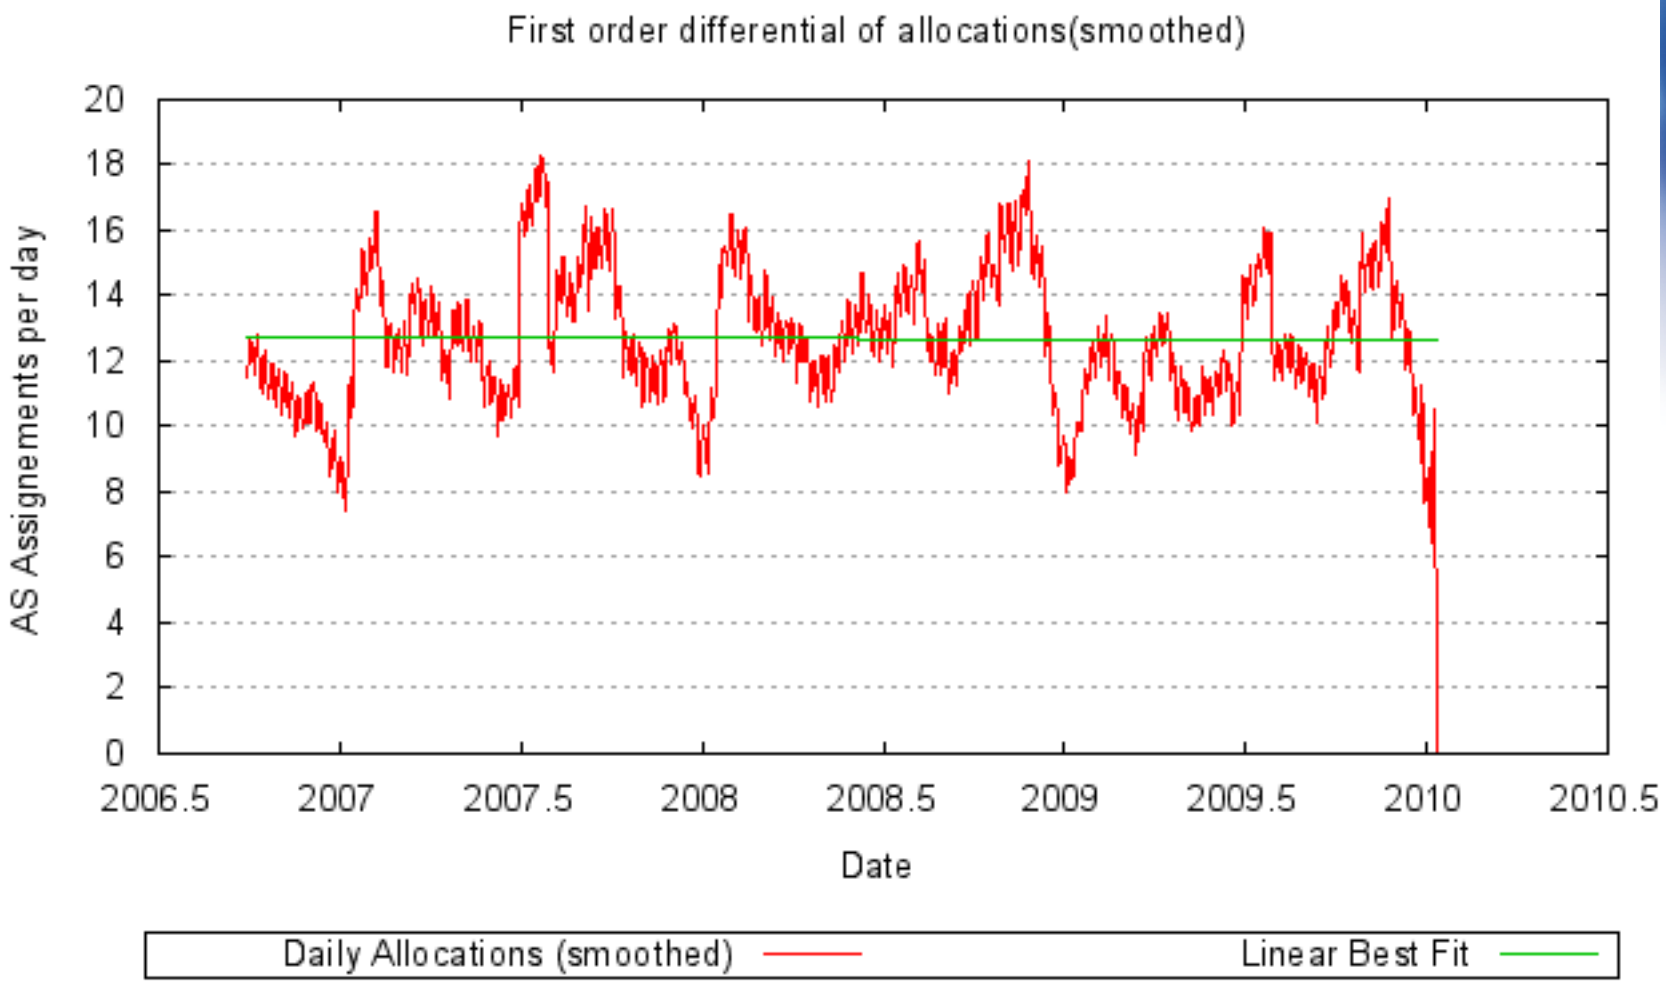

Consumption Rate

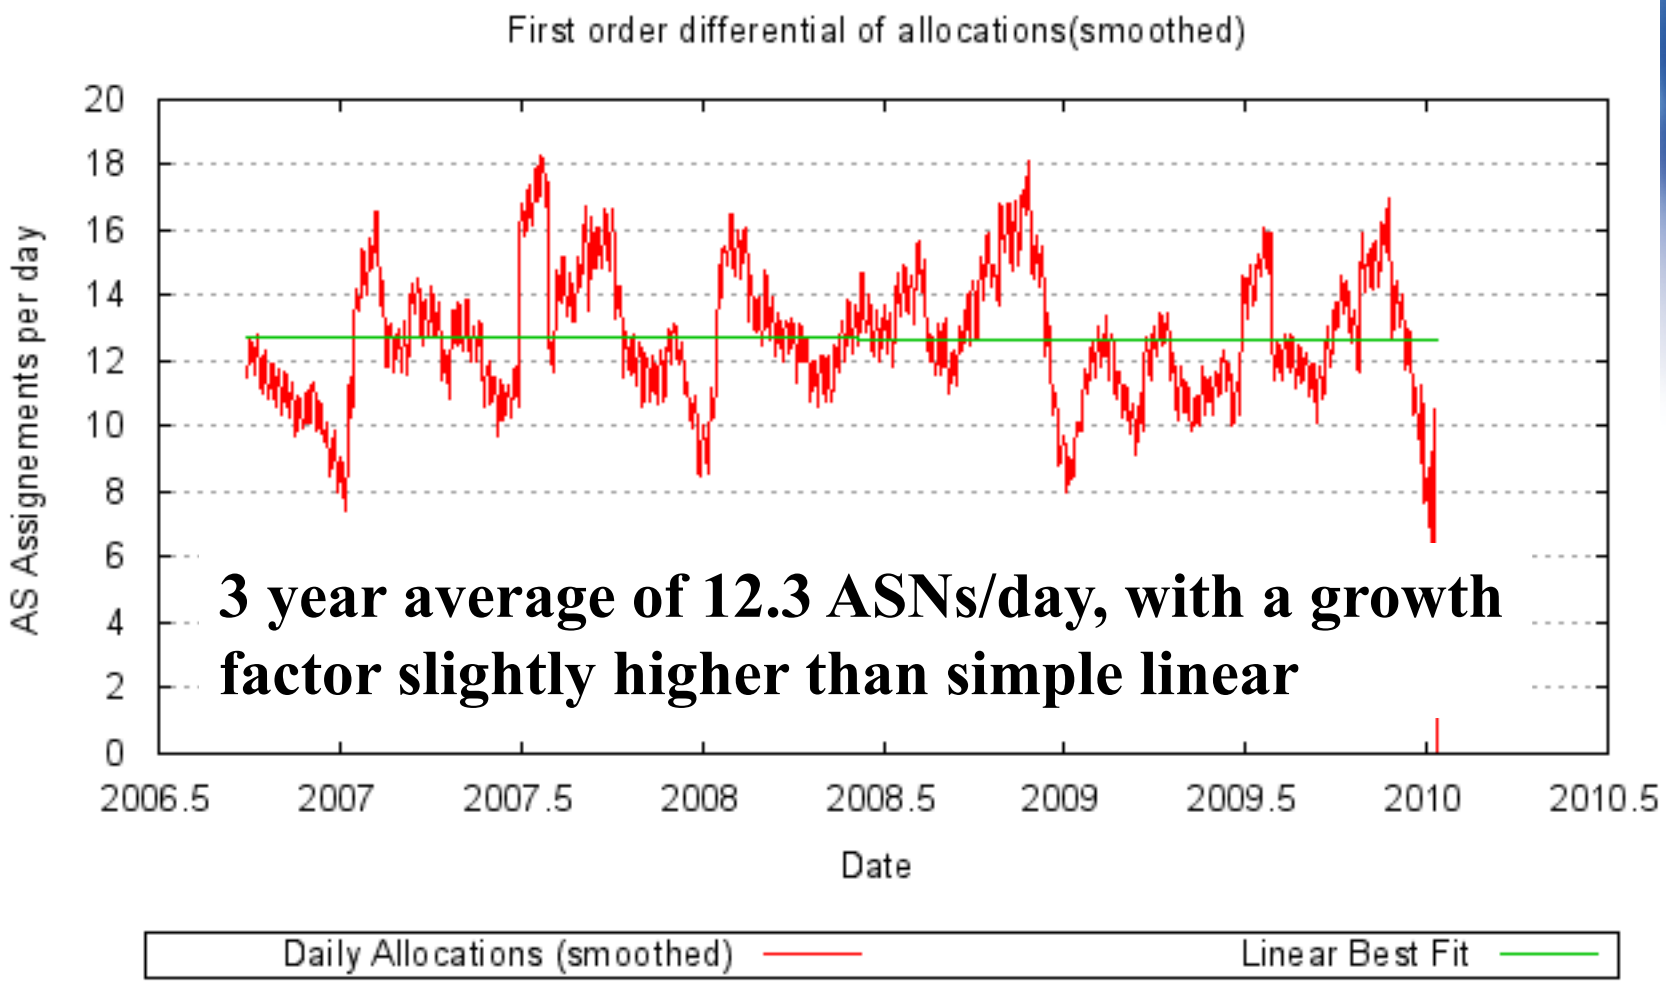

How long have we got?

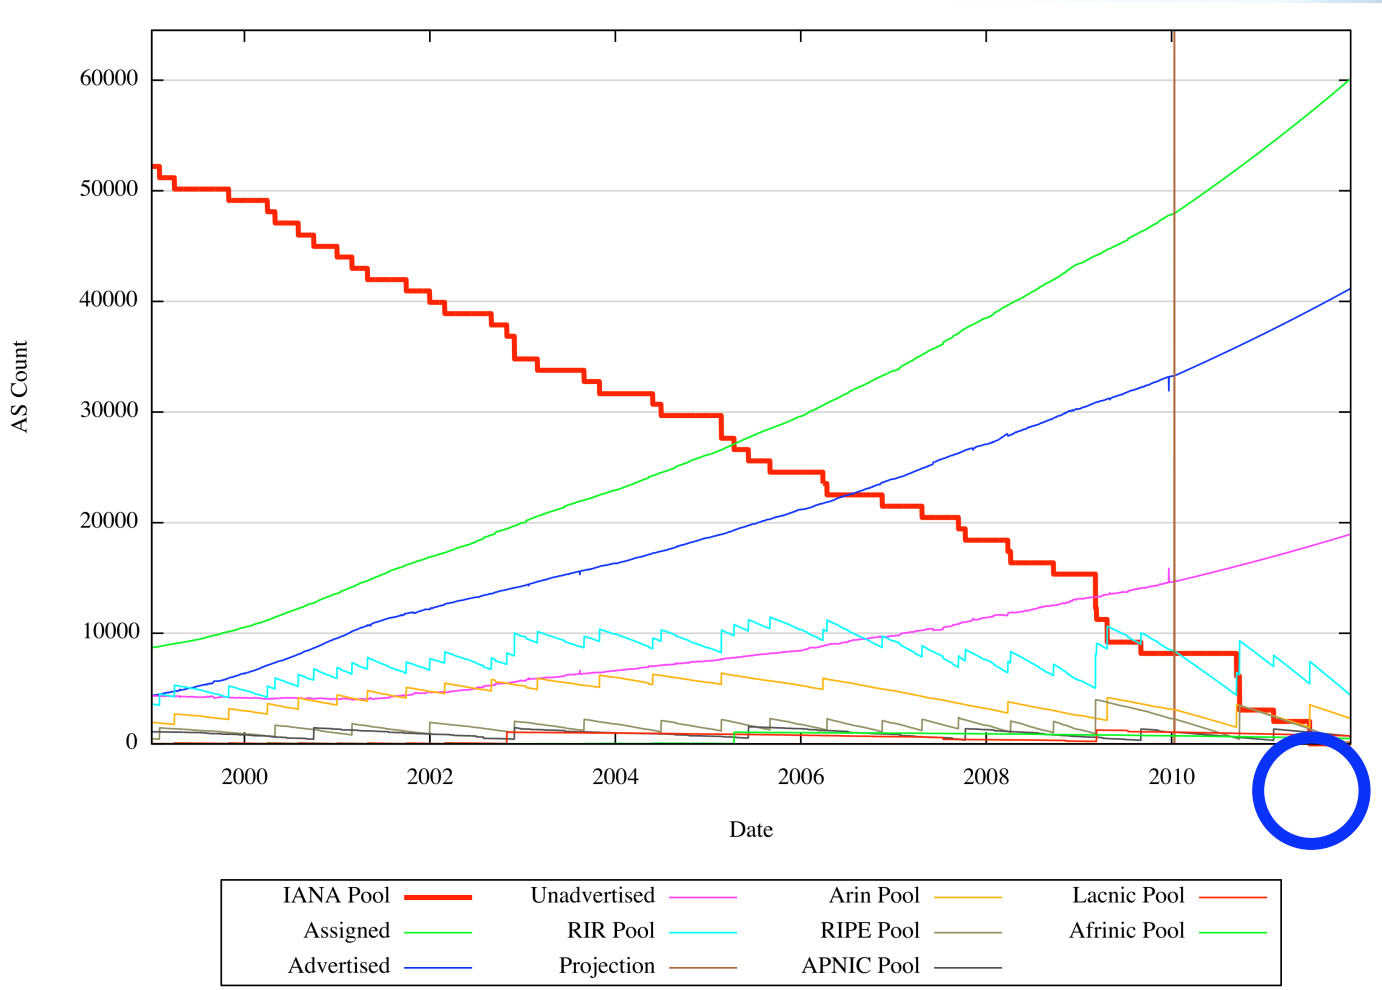

How long have we got?

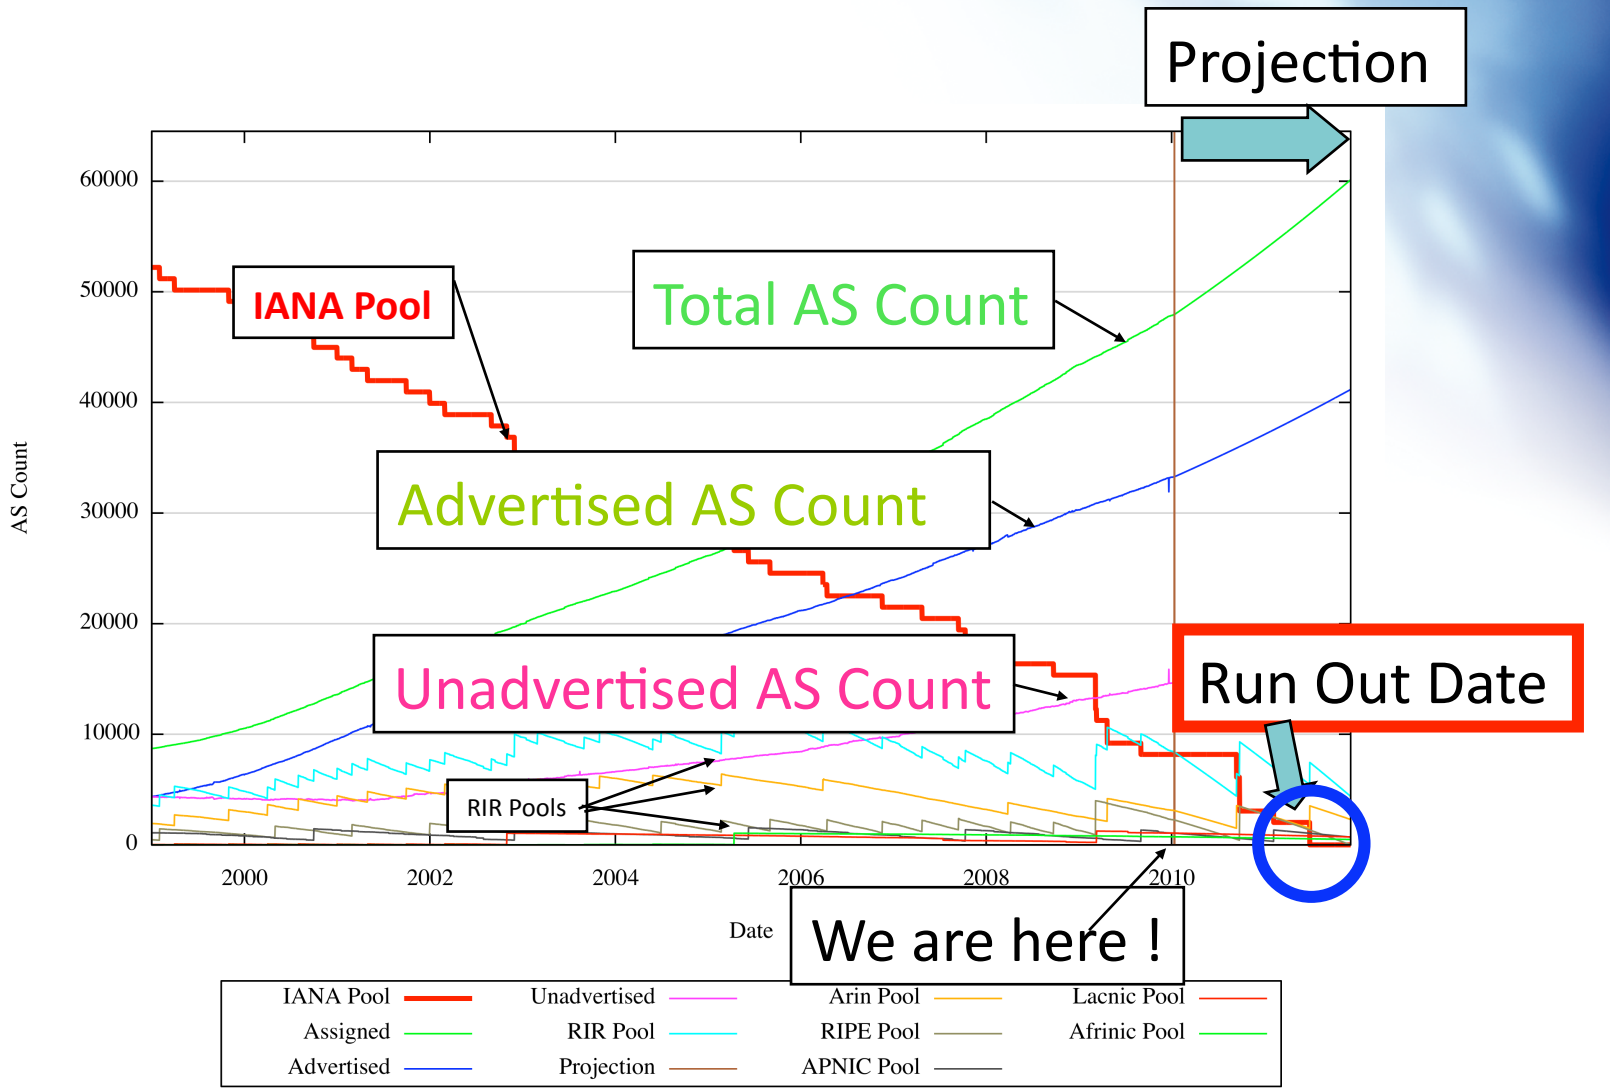

Asia Pacific Network Information Centre

APNIC

16-bit AS Number Exhaustion

- We are exhausting the 16-bit AS Number pool
  - IANA will allocate its last AS number block in **March 2011**
  - RIPE will exhaust its 16 bit AS Number pool in **December 2011**

See http://www.potaroo.net/tools/asns

This is not exactly news!

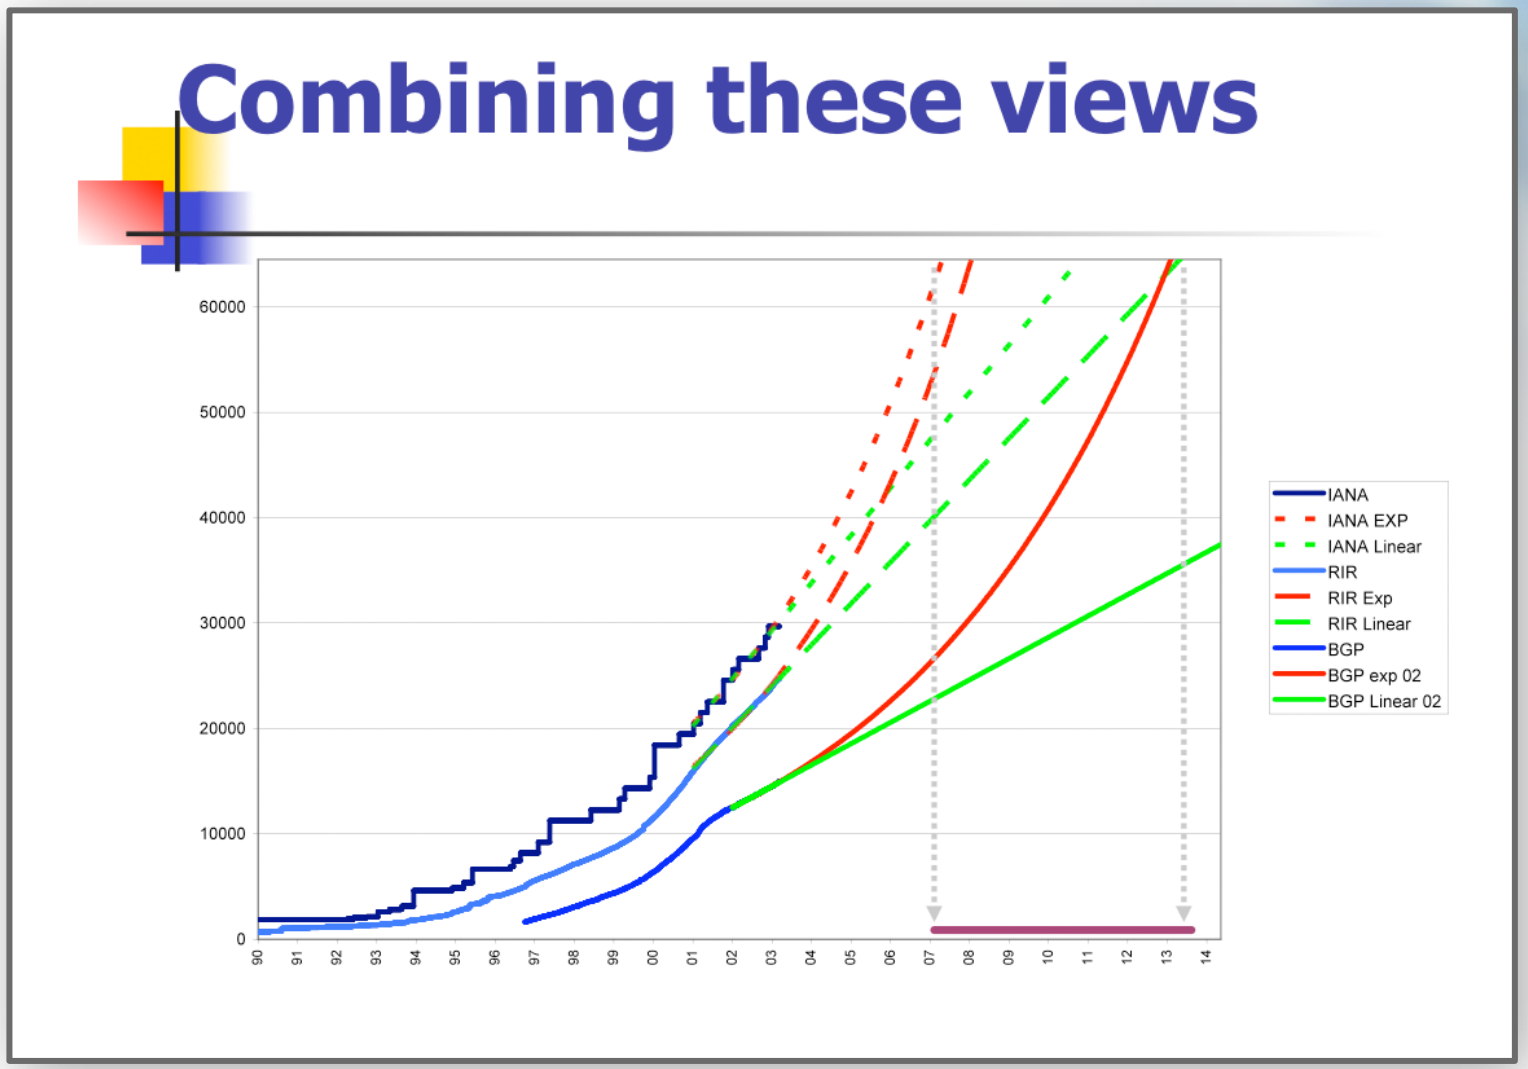

AS Number Exhaustion – March 2003

2003 Projection

Current AS Forecast

- The available AS number pool will exhaust in the timeframe of 2009-2011 if current AS use trends continue

2009

- no significant reclamation in old AS number space
- No coordinated effort to increase utilization density of AS numbers

2011

- reclamation and increased deployment efficiency

The Agenda for AS Transition

Developed in 2004 as a 4 step process:

1. IETF to complete BGP Standards to support transition mechanisms to 32-bit AS numbers

   ~2 years

2. RIRs to start making 32-bit AS numbers available

   ~½ year

3. Vendors to provide 32-bit AS number capable BGP implementations

   ~1 year

4. BGP networks to commence deployment

   - **ready for deployment by 2008!**

1. IETF Standards Activity

- 4-Byte AS Specification
  - Initial draft prepared in Feb 2001
    - Change BGP Attribute Definitions to extend AS components from 16 to 32 bits
    - Change BGP OPEN message to include capability negotiation for peer 4 byte support
    - Carry 32-bit AS path across 16-bit AS domains using new opaque transitive attribute (AS4_PATH)
    - Transition mechanism via translation and tunneling that allows piecemeal introduction of 4-byte AS numbers into the Internet
  - Specification ready for publication in late 2005
- IANA 32 bit AS number registry created in November 2006
- **RFC 4893** published in May 2007

The Agenda for AS Transition

1. IETF to complete BGP Standards to support transition mechanisms to 32-bit AS numbers
   RFC 4893 – May 2007
2. RIRs to start making 32-bit AS numbers available
3. Vendors to provide 32-bit AS number capable BGP implementations
4. BGP networks to commence deployment

2. RIR ASN Allocation Policy

- Globally coordinated policy proposal 2005 / 2006
- Intended to avoid surprises and disappointment during the run-out of the 16-bit AS number space
- State clear milestones for vendors, ISPs and network admins for 32-bit ASN uptake
- Phased transition to the 32-bit AS number pool:
  - 2007 – 32 bit ASNs available upon request
  - 2009 – 32 bit ASNs available by default
  - 2010 – transition projected to be complete

The Agenda for AS Transition

1. IETF to complete BGP Standards to support transition mechanisms to 32-bit AS numbers
2. RIRs to start making 32-bit AS numbers available
3. Vendors to provide 32-bit AS number capable BGP implementations
4. BGP networks to commence deployment

RFC 4893 – May 2007

32 bit AS Policy – 2006

3. Vendor Support in BGP

Name	Version	Notation
Alcatel-Lucent SR OS	7.0	asplain
Arbor Peakflow SP	5.5	asplain
BIRD	1.0.12	asplain
Brocade (Foundry) IronWare	4.0.00 for the NetIron MLX and XMR, 2.8.00 for the BigIron RX	asdot, asdot+, asplain
Cisco IOS	12.0(32)S12, 12.0(32)SY8, 12.2(33)SXI1, 12.4(24)T	asdot (asplain planned for future)
Cisco IOS XE	2.3	asplain (asdot optional)
Cisco IOS XR	3.4(1)	asdot (asplain planned for 3.9)
Cisco NX-OS	4.0(1)	asdot (asplain planned for 4.1(3))
ExtremeXOS	Need Information	Need Information
Juniper JUNOS	9.1R1	asplain (asdot optional)
Juniper JUNOSe	4.1.0	asplain
Force10 FTOS	7.7.1.0	asplain (asdot, asdot+ optional)
OpenBGPD	4.2, patches for 3.9 and 4.0	asdot
Quagga	0.99.10, patches for 0.99.6 and other versions	asplain
Redback SEOS	2.0	ascolon (asplain planned for end of 2009)

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

The Agenda for AS Transition

1. IETF to complete BGP Standards to support transition mechanisms to 32-bit AS numbers
2. RIRs to start making 32-bit AS numbers available
3. Vendors to provide 32-bit AS number capable BGP implementations
4. BGP networks to commence deployment

RFC 4893 – May 2007

32 bit AS Policy – 2006

Currently gathering pace:

2010

RIR Allocation Data of 32-bit AS's

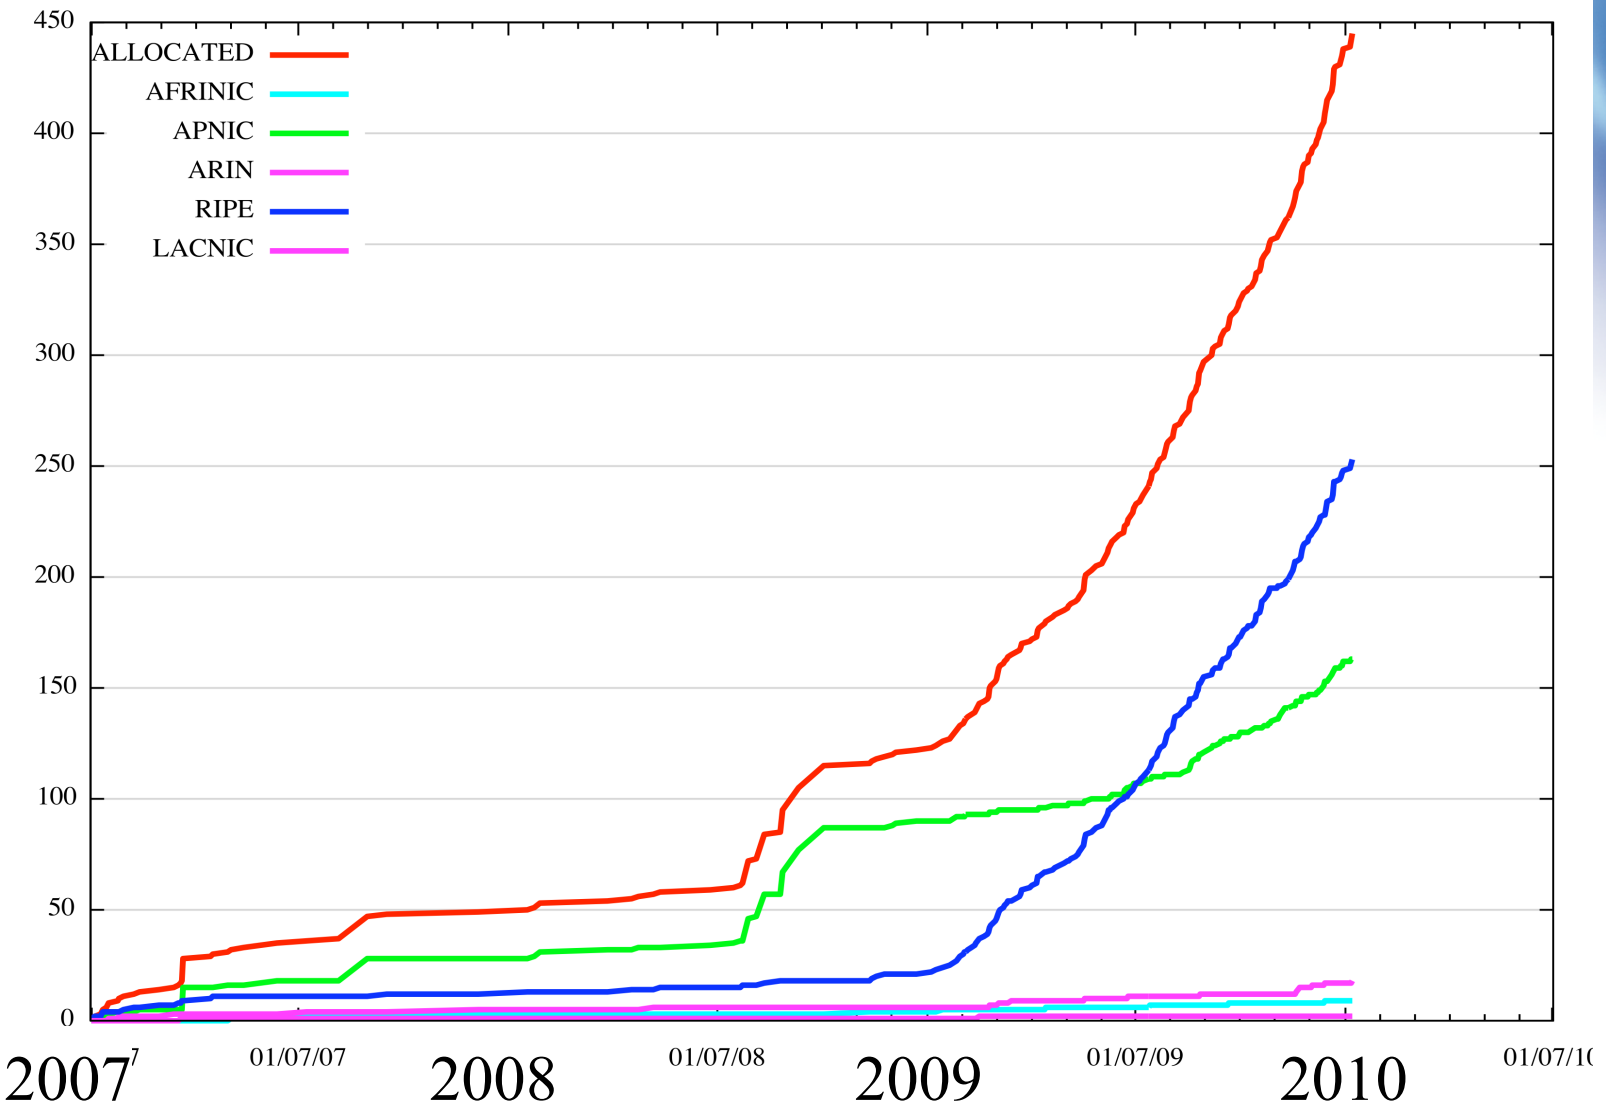

32-bit ASNs in BGP

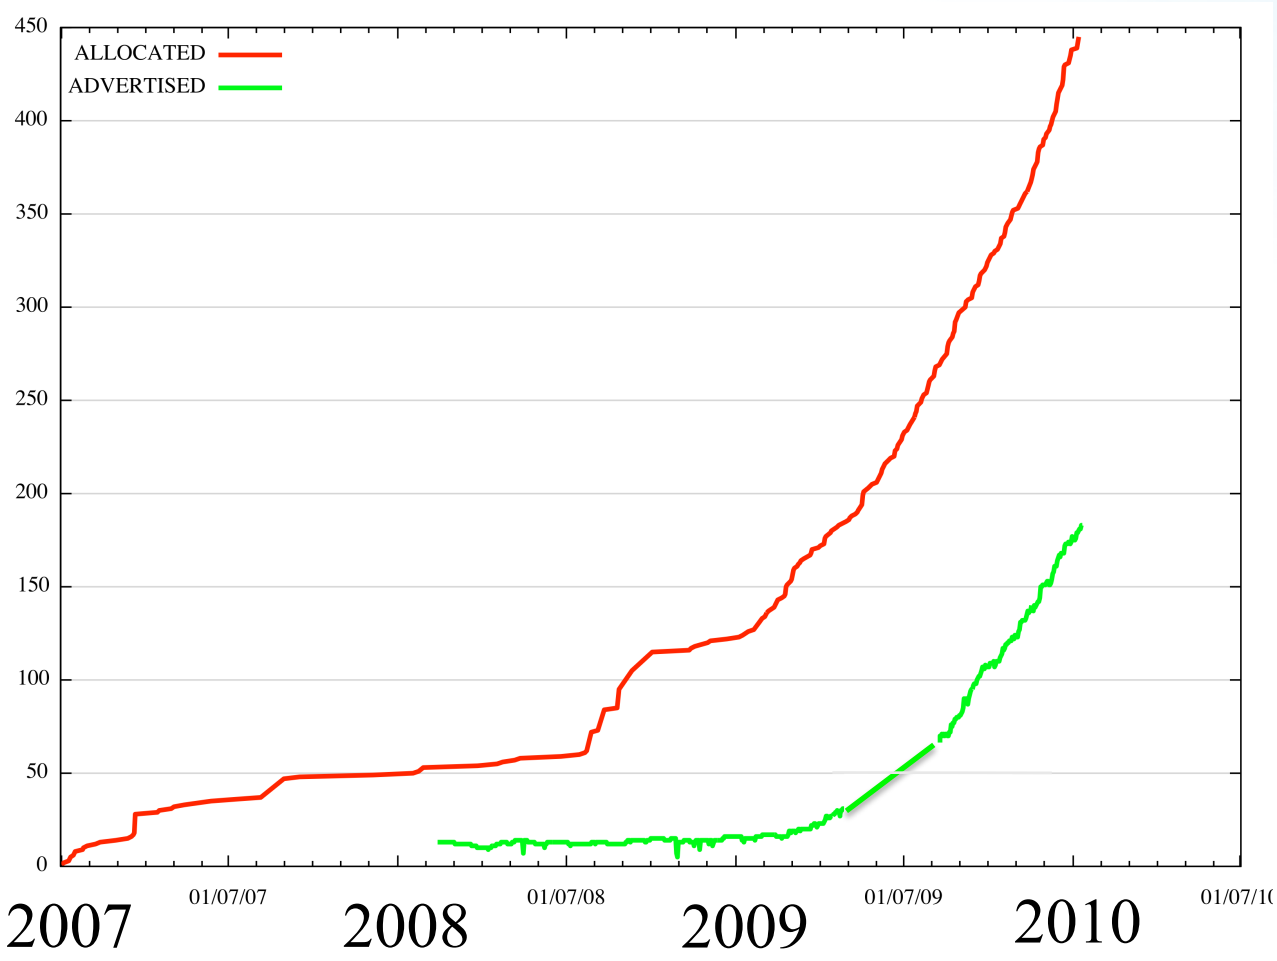

32-bit ASN Deployment

- Allocation status as of January 2010:
  - Advertised: 183
  - Unadvertised: 262

- In 2009 the RIRs allocated 4,761 ASNs
  - 4,445 were 16-bit ASNs
  - 316 were 32 bit ASNs

The Agenda for Transition

1. IETF to complete BGP Standards to support transition mechanisms to 32-bit AS numbers
2. RIRs to start making 32-bit AS numbers available
3. Vendors to provide 32-bit AS number capable BGP implementations
4. BGP networks to commence deployment

RFC 4893 – May 2007

32 bit AS Policy – 2006

Currently gathering pace 2009 - 2010

Lagging – badly!

How can we assist with 32-bit AS deployment?

- Information and education
  - Keep the community informed
  - Address some common misunderstandings about 4 byte AS numbers
- Supply chain pressure
  - Add 4 byte AS support to your “mandatory to support” in your next BGP purchase

AS Path Semantics in BGP

- It's a *path metric* where the length of the AS Path 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-going-to-loop-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 path-metric and loop-detecting properties for BGP to function correctly

32-bit 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** 32-bit AS Path information into a 16-bit representation
  - **Tunnel** 32-bit AS Path information through 16-bit AS domain as an update attribute

32-bit 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** 32-bit AS Path information into a 16-bit representation
  - **Tunnel** 32-bit AS Path information through 16-bit AS domain as an update attribute

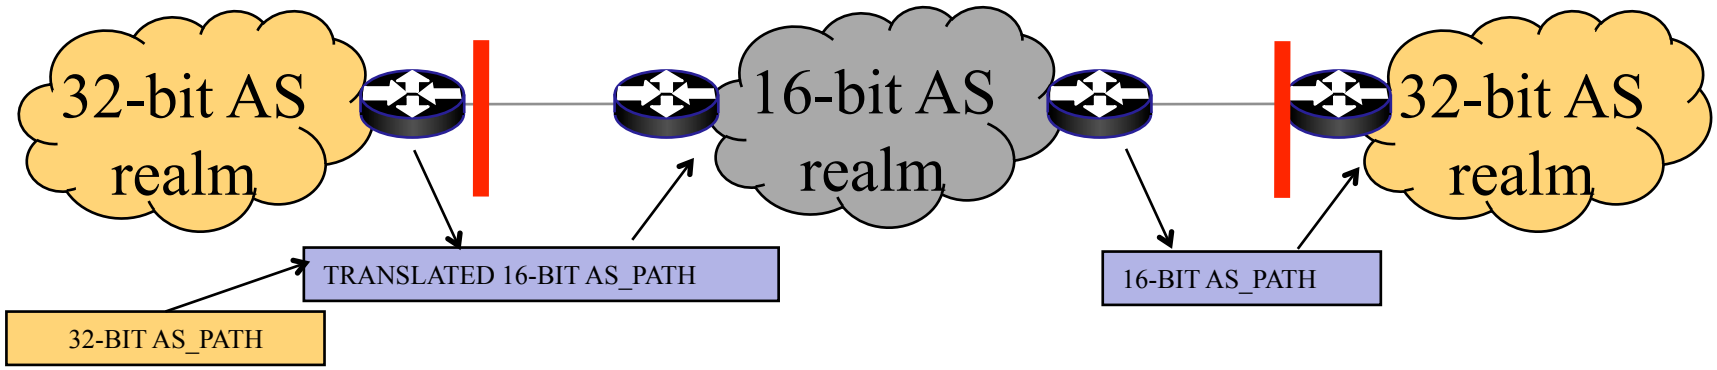

TRANSLATE all 32-bit-only AS numbers to **AS23456**

32-bit 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** 32-bit AS Path information into a 16-bit representation
  - **Tunnel** 32-bit AS Path information through 16-bit AS domain as an update attribute

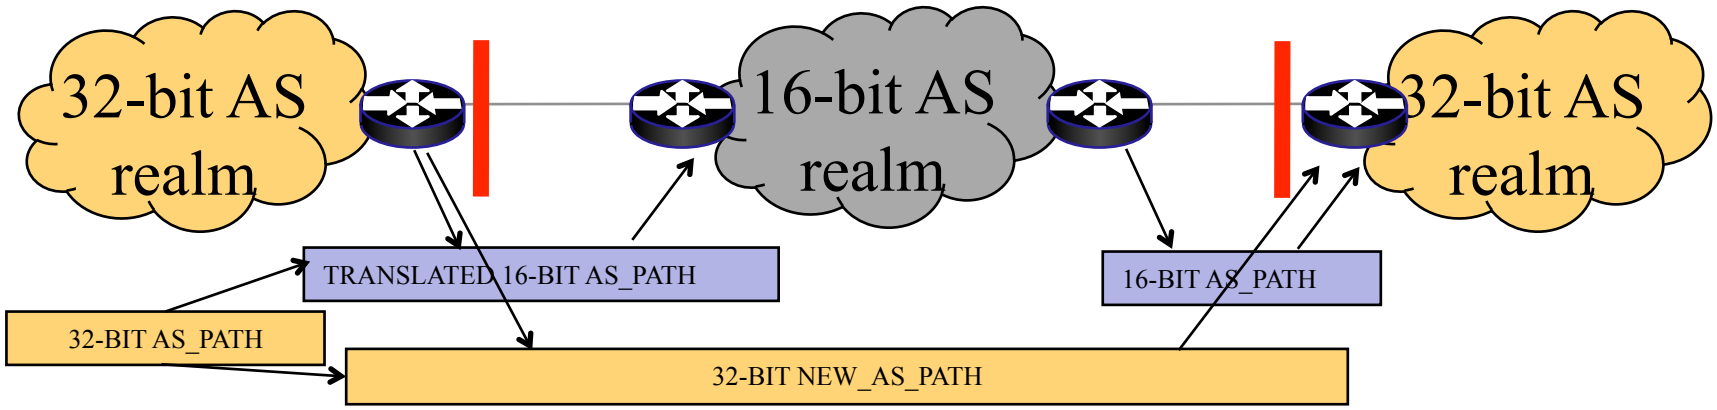

COPY 32-bit AS_PATH to NEW_AS_PATH

32-bit 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** 32-bit AS Path information into a 16-bit representation
  - **Tunnel** 32-bit AS Path information through 16-bit AS domain as an update attribute

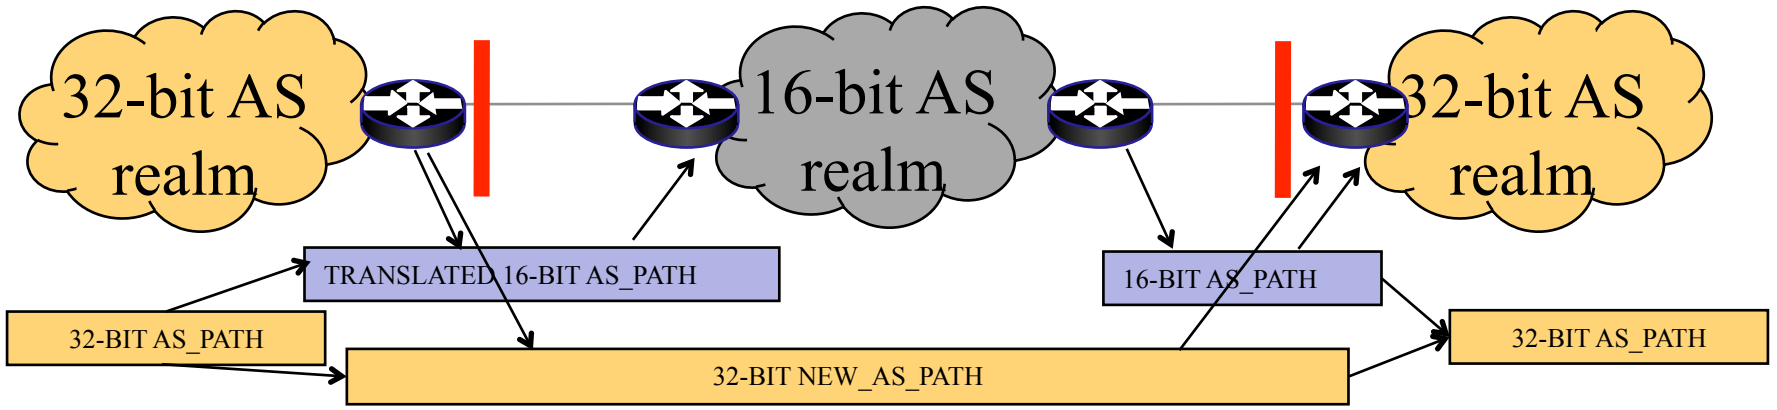

REASSEMBLE 32-bit AS_PATH

32-bit / 16-bit BGP Example

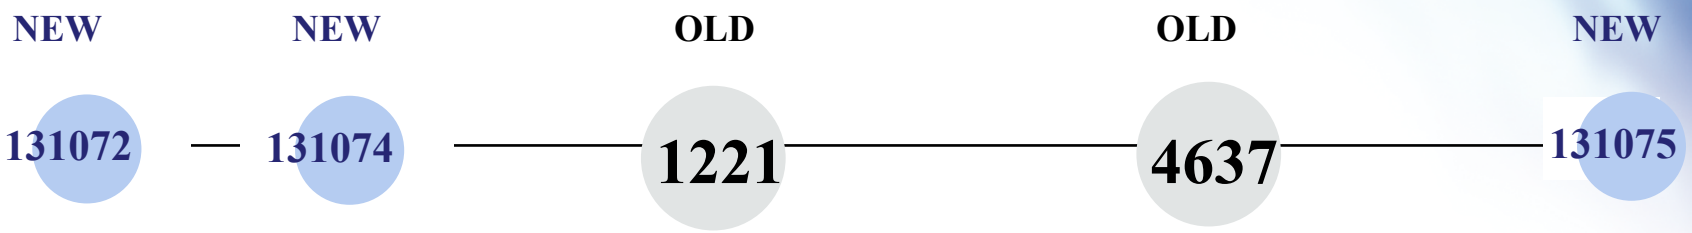

AS Path in the RIB

i

32-bit / 16-bit BGP Example

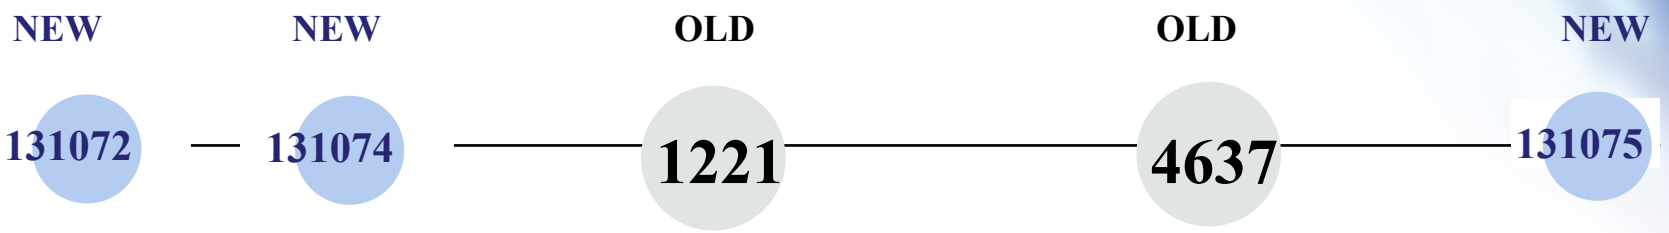

AS Path in the RIB

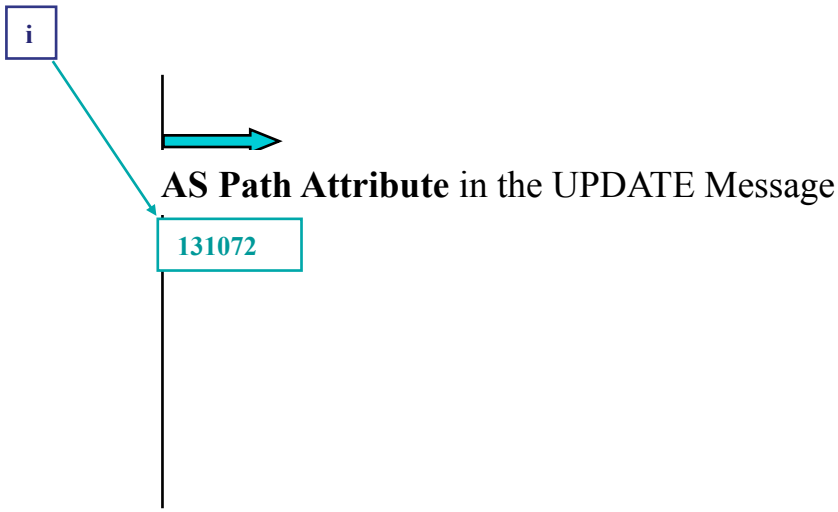

32-bit / 16-bit BGP Example

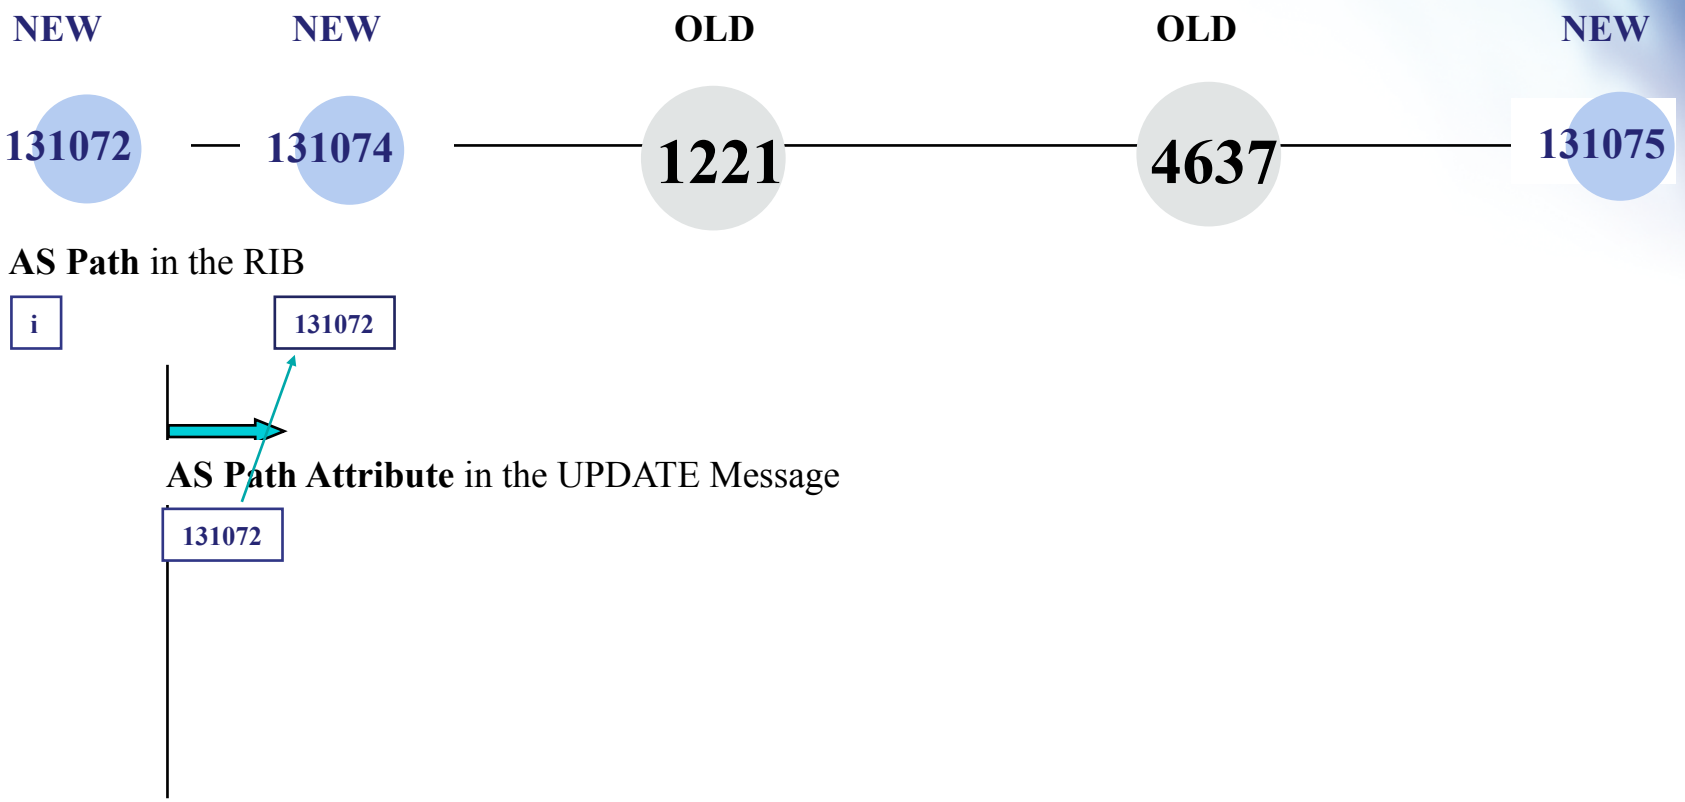

32-bit / 16-bit BGP Example

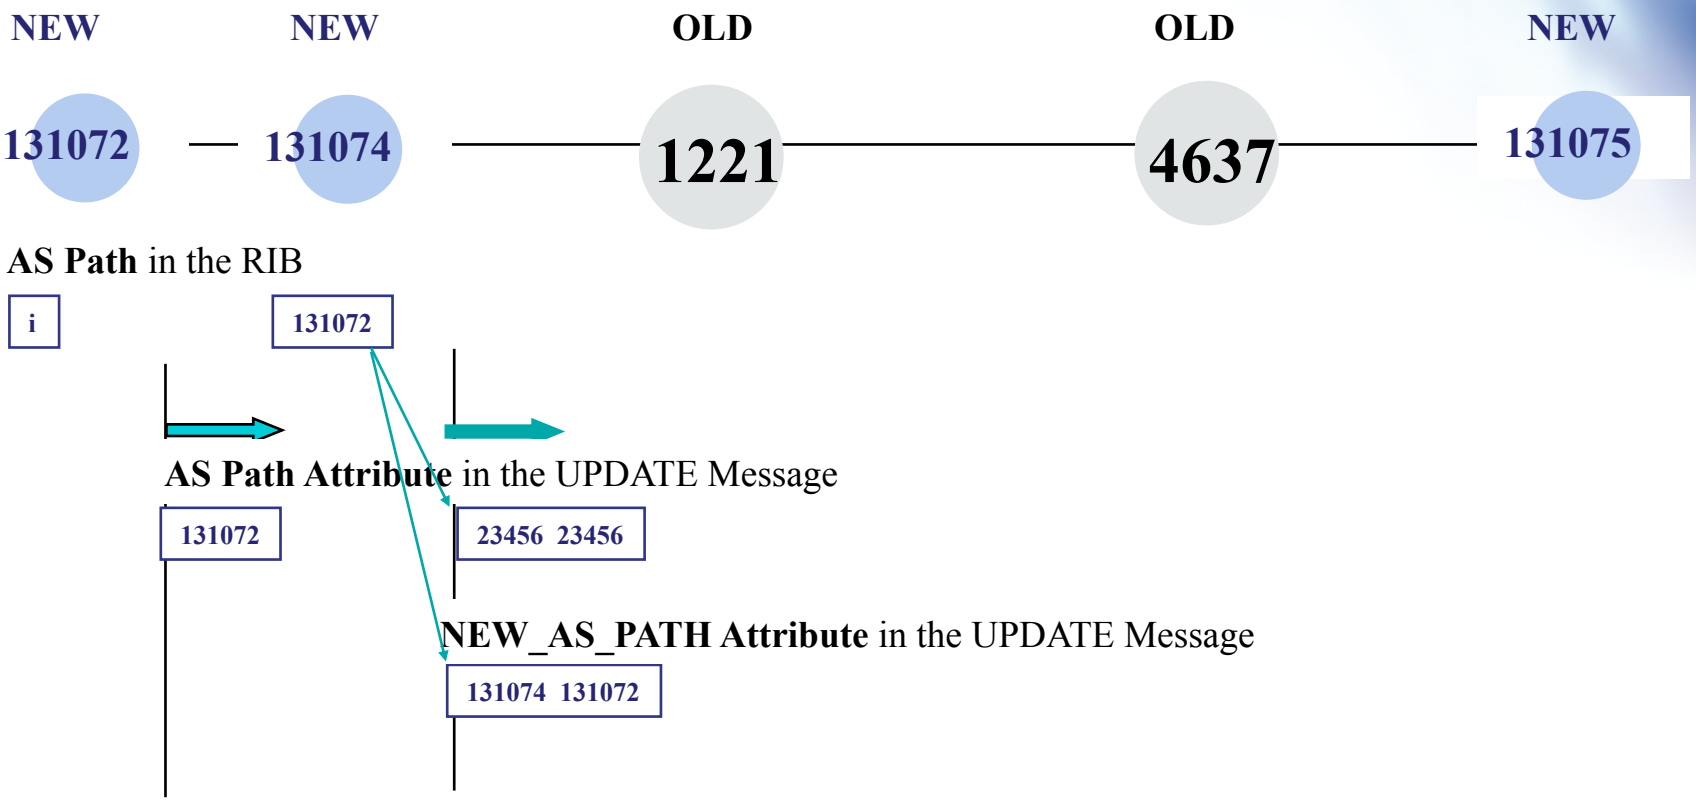

32-bit / 16-bit BGP Example

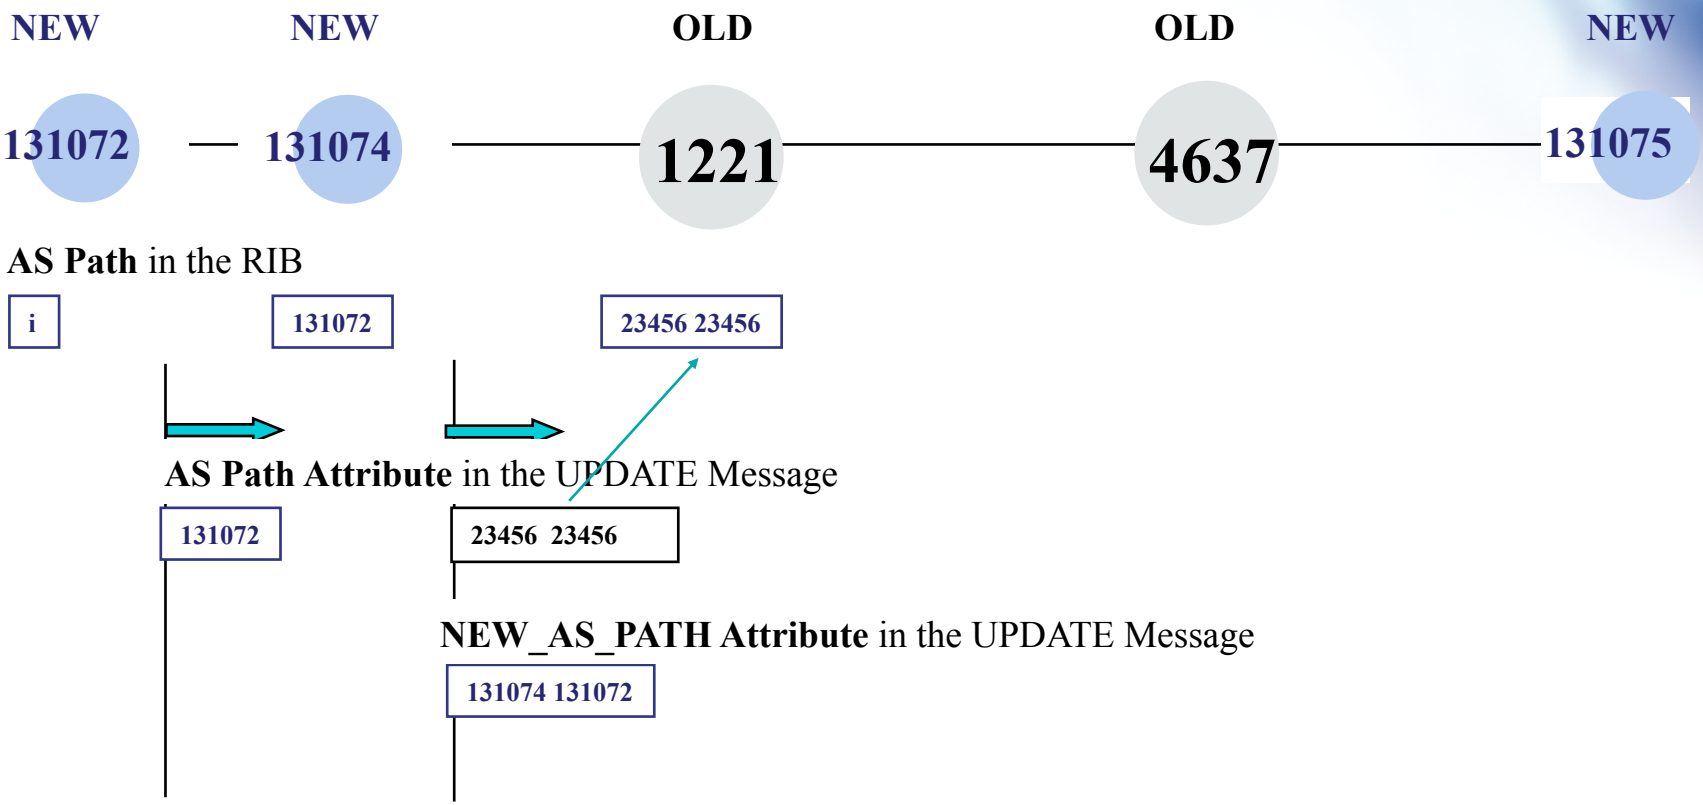

32-bit / 16-bit BGP Example

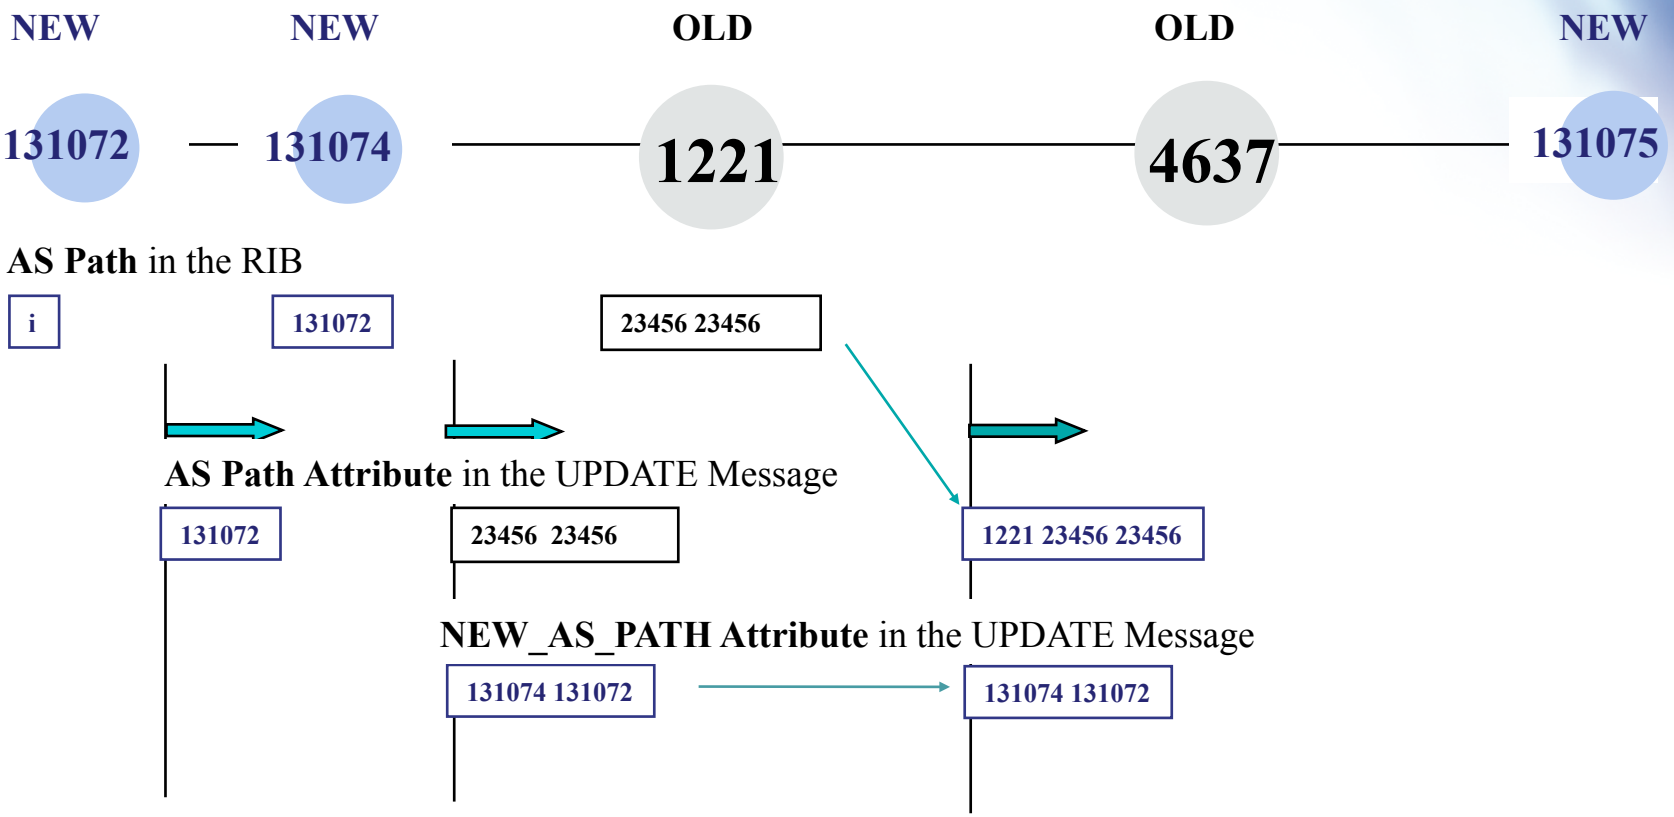

32-bit / 16-bit BGP Example

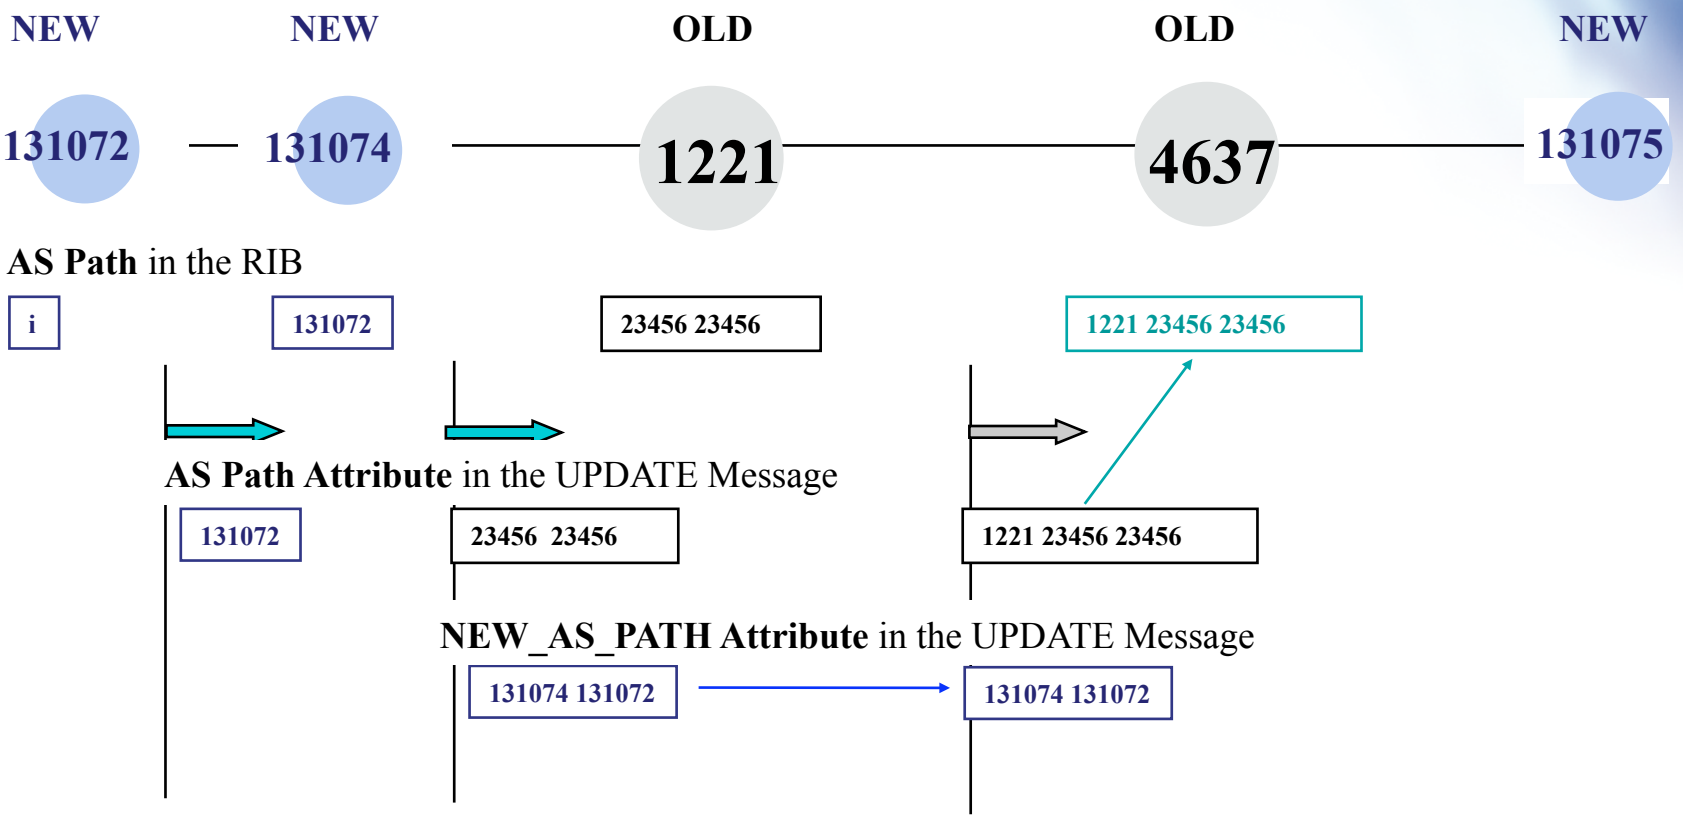

32-bit / 16-bit BGP Example

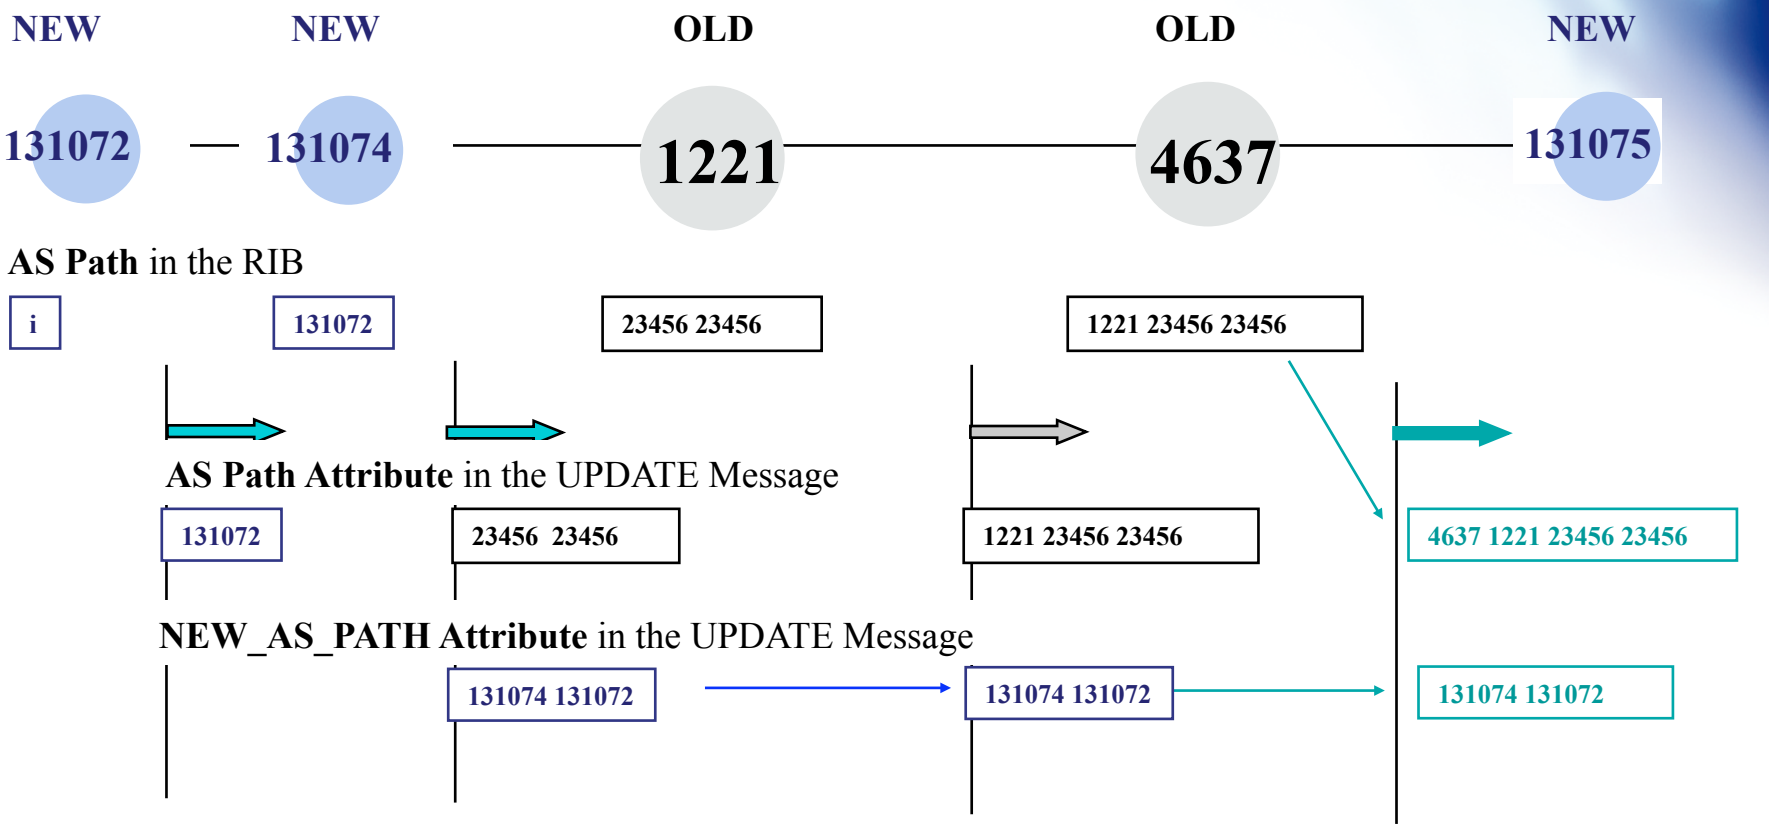

32-bit / 16-bit BGP Example

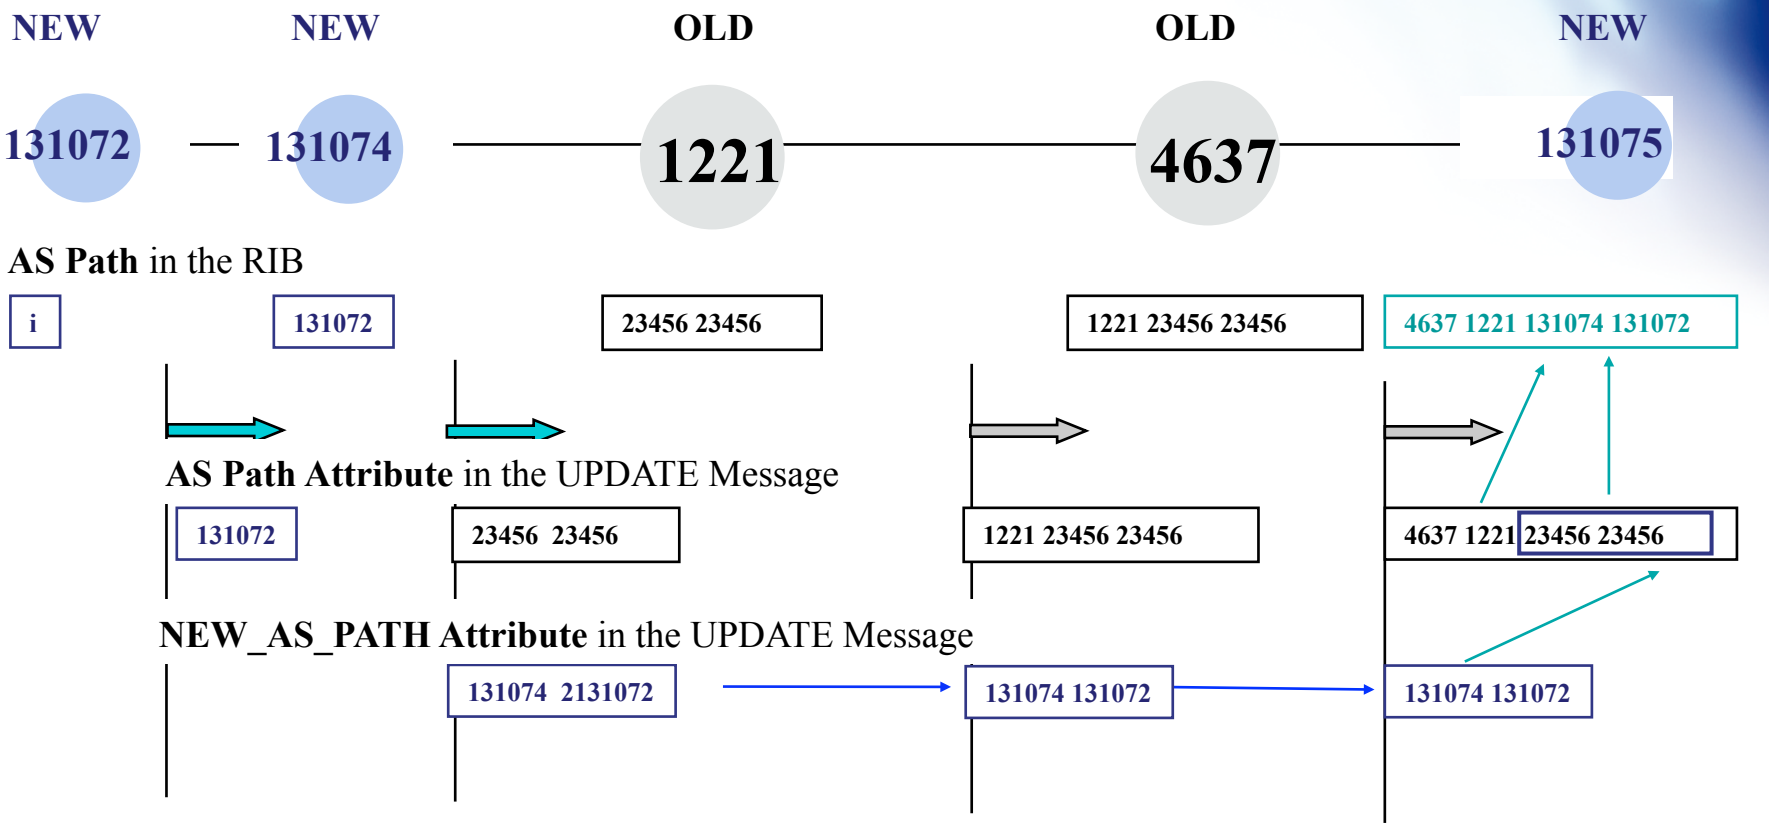

32-bit / 16-bit BGP Example

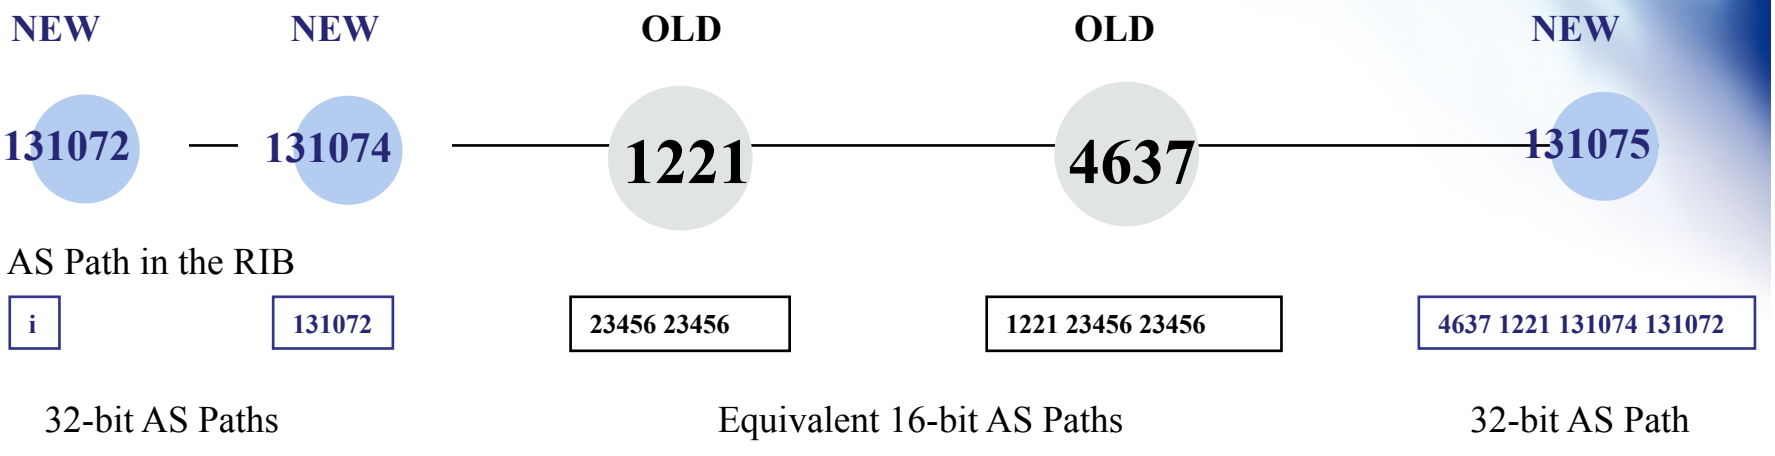

Can old-BGP get Confused?

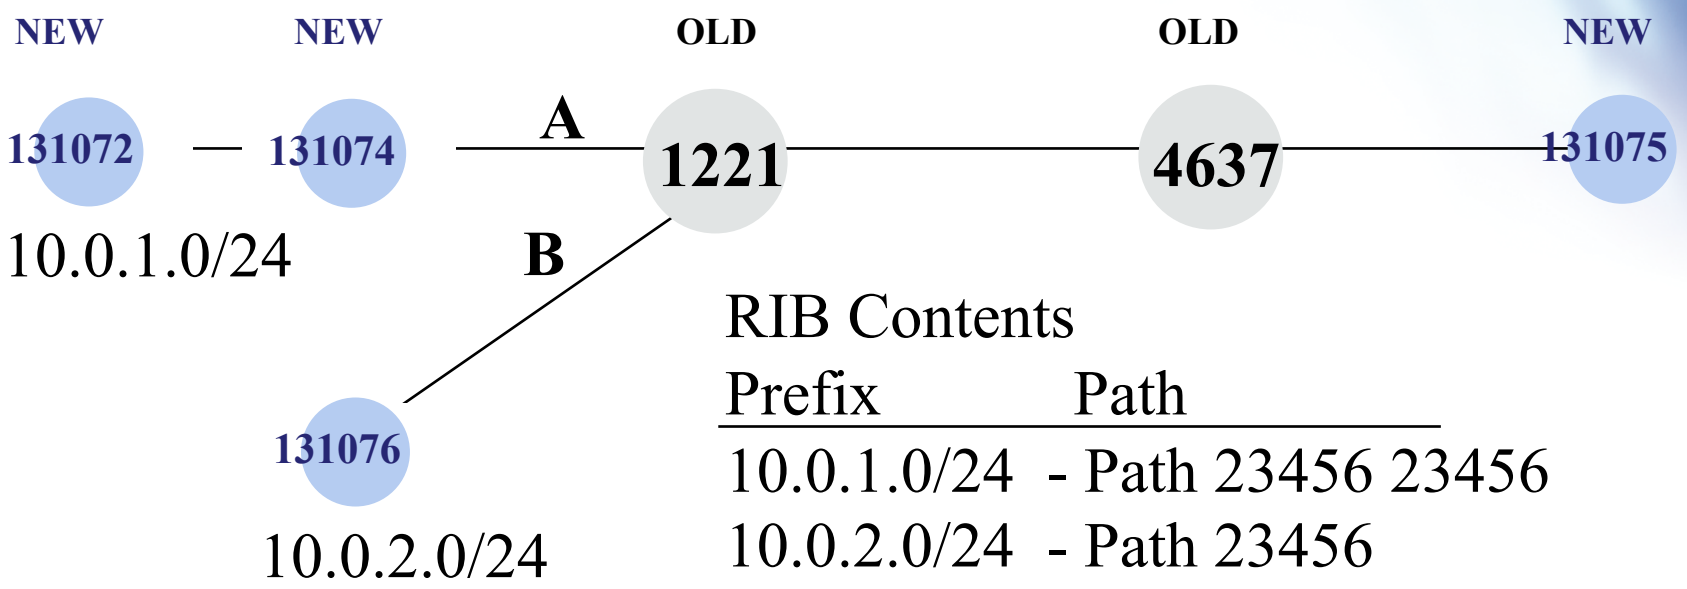

RIB Contents

Prefix	Path
10.0.1.0/24	- Path 23456 23456
10.0.2.0/24	- Path 23456

NO! BGP Nexthop is the key!

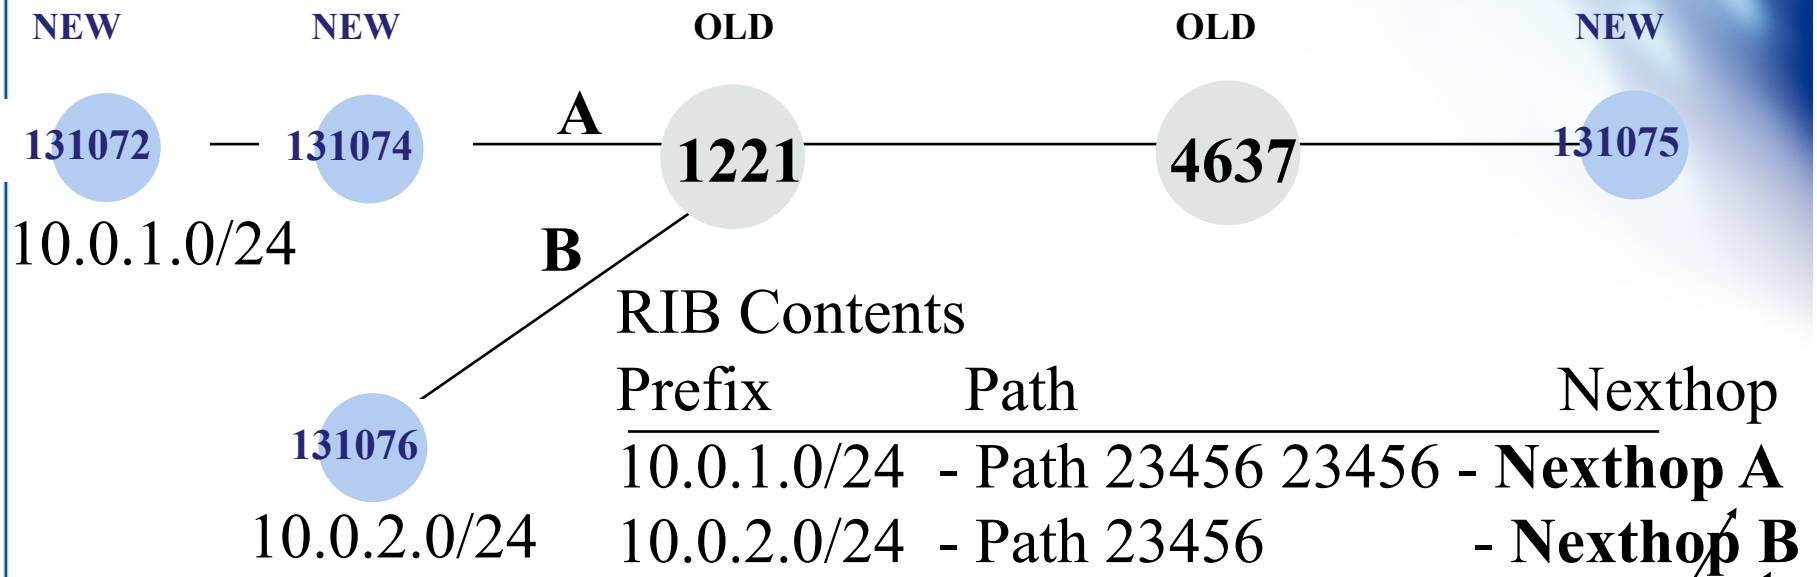

RIB Contents

Prefix	Path	Nexthop
10.0.1.0/24	- Path 23456 23456	- **Nexthop A**
10.0.2.0/24	- Path 23456	- **Nexthop B**

Traffic from AS 1221 to 10.0.1.0/24 will be forwarded on interface A
Traffic from AS 1221 to 10.0.2.0/24 will be forwarded on interface B

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

NEW_AS_PATH Attribute

- BGP speakers in 16-bit AS domains should support **NEW_AS_PATH** as a transitive optional attribute in UPDATE messages
  - because that's where the 32-bit 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 16-bit world as a fallback

Default BGP configurations will do the right thing here

NEW_AGGREGATOR Attribute

- BGP speakers in 16-bit AS domains should support **NEW_AGGREGATOR** as a transitive optional attribute in UPDATE messages
  - because that's where the 32-bit Aggregator AS 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

Default BGP configurations should do the right thing here

AS 23456

- **AS 23456** is going to appear in many 16-bit AS paths – both origin and transit

This is not an error – it's a 16-bit token holder for a 32-bit AS number

Netflow and Sflow

- Netflow analyzers may need to be reviewed
  - Netflow version 9 supports 32-bit AS numbers
    - But may not report the 32-bit ASN unless the netflow collector is a 32-bit BGP
    - Does your analyzer support 32-bit AS numbers?
  - Netflow version 8 and earlier are 16-bit AS constrained
    - Which implies that you'll be seeing AS 23456 more than you may want!
- Sflow
  - Appears to define a source and dest AS using a 32 bit field
  - So it *should* be ok!

BGP Communities

- If you want to explicitly signal to a 32-bit AS using communities in BGP then you will need to explicitly signal the 32-bit AS using **BGP Extended Communities**
  - Attempting to use AS 23456 in this context will have unintended consequences!

  See:
  - RFC 4630
  - RFC 5568

BGP Memory requirements

- BGP memory requirements will increase
  - 32-bit BGP speakers will need twice the memory used to hold AS paths[1]
  - 16-bit BGP speakers will need up to three times the memory used to hold AS paths plus NEW_AS_PATH extended community attribute[2]
    - 30,000 unique AS paths with an average length of 4 implies an additional memory requirement of 240Kb for 32-bit BGP and up to a further 480Kb for 16-bit BGP

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”

Operational Support Systems

What happens when you have a customer / transit / peer with a 32-bit AS Number?

- 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 65538** needs to be auto-translated into **AS 23456** in a number of places, including in your OSS
  - Your BGP routers may need to peer with **AS 23456**, transit across **AS 23456**, and have multiple customers on **AS 23456** at the same time, while also understanding that these refer to different external parties

Related Systems

- Anything that wants to manipulate AS numbers, including your local support systems, scripts and databases

Common Questions

1. *Someone out there is using 4 byte AS numbers. Do I have to upgrade my BGP to support 4-byte AS numbers in order to reach the prefixes that they are announcing?*

Common Questions

1. *Someone out there is using 4 byte AS numbers. Do I have to upgrade my BGP to support 4-byte AS numbers in order to reach the prefixes that they are announcing?*

NO!

Common Questions

2. My customers / peers/ upstreams are using 4-byte AS numbers. Do I have to upgrade my BGP to support 4-byte AS numbers?

Common Questions

2. My customers / peers/ upstreams are using 4-byte AS numbers. Do I have to upgrade my BGP to support 4-byte AS numbers?

NO!

Common Questions

3. Can I use communities for 4-byte ASNs?

Common Questions

3. Can I use communities for 4-byte ASNs?

YES and NO

Common Questions

3. Can I use communities for 4-byte ASNs?

- **NO** - if your BGP does not support RFC5668
  - because there is only a 2 byte field for the ASN in the conventional BGP community
  - You need to use a BGP extended community to define a set of communities for 4-byte origin and target AS values
  - This is specified in RFC5668
  - Ask your vendor when they will be supporting BGP extended communities with 4-byte ASNs – RFC5668
- **YES** – if your BGP supports RFC5668

Common Questions

4. Is Internet about to crash and die?

Common Questions

4. I see AS 23456 in a 4-byte AS path – Is the Internet about the crash and die?

Calm down!

Common Questions

4. I see AS 23456 in a 4-byte AS path – Is the Internet about the crash and die?

It may be abnormal, but its not fatal

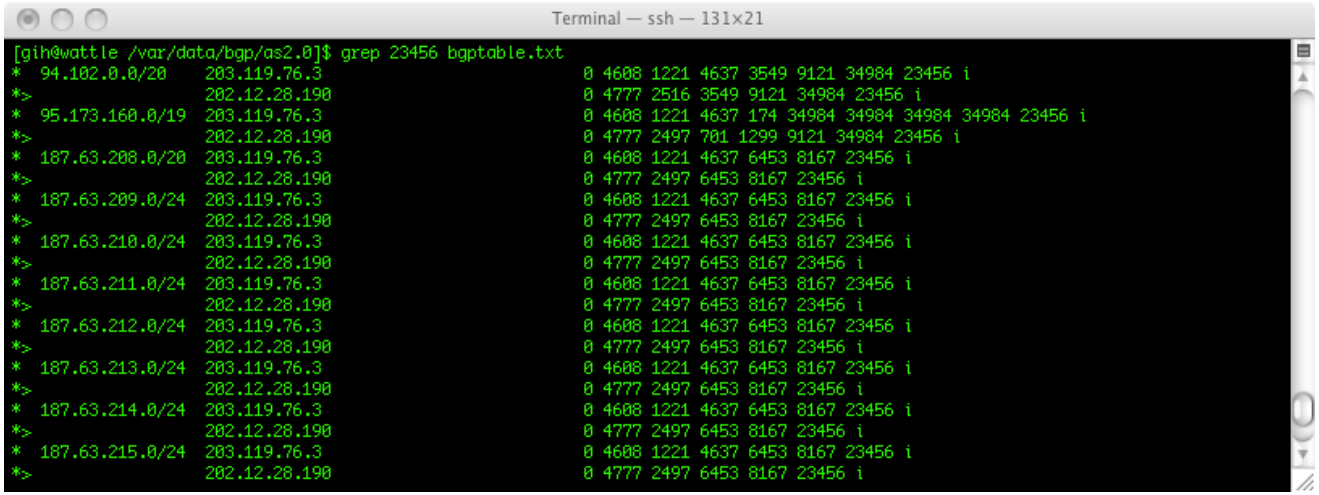

Common Questions

4. I see AS 23456 in a 4-byte AS path – Is the Internet about the crash and die?

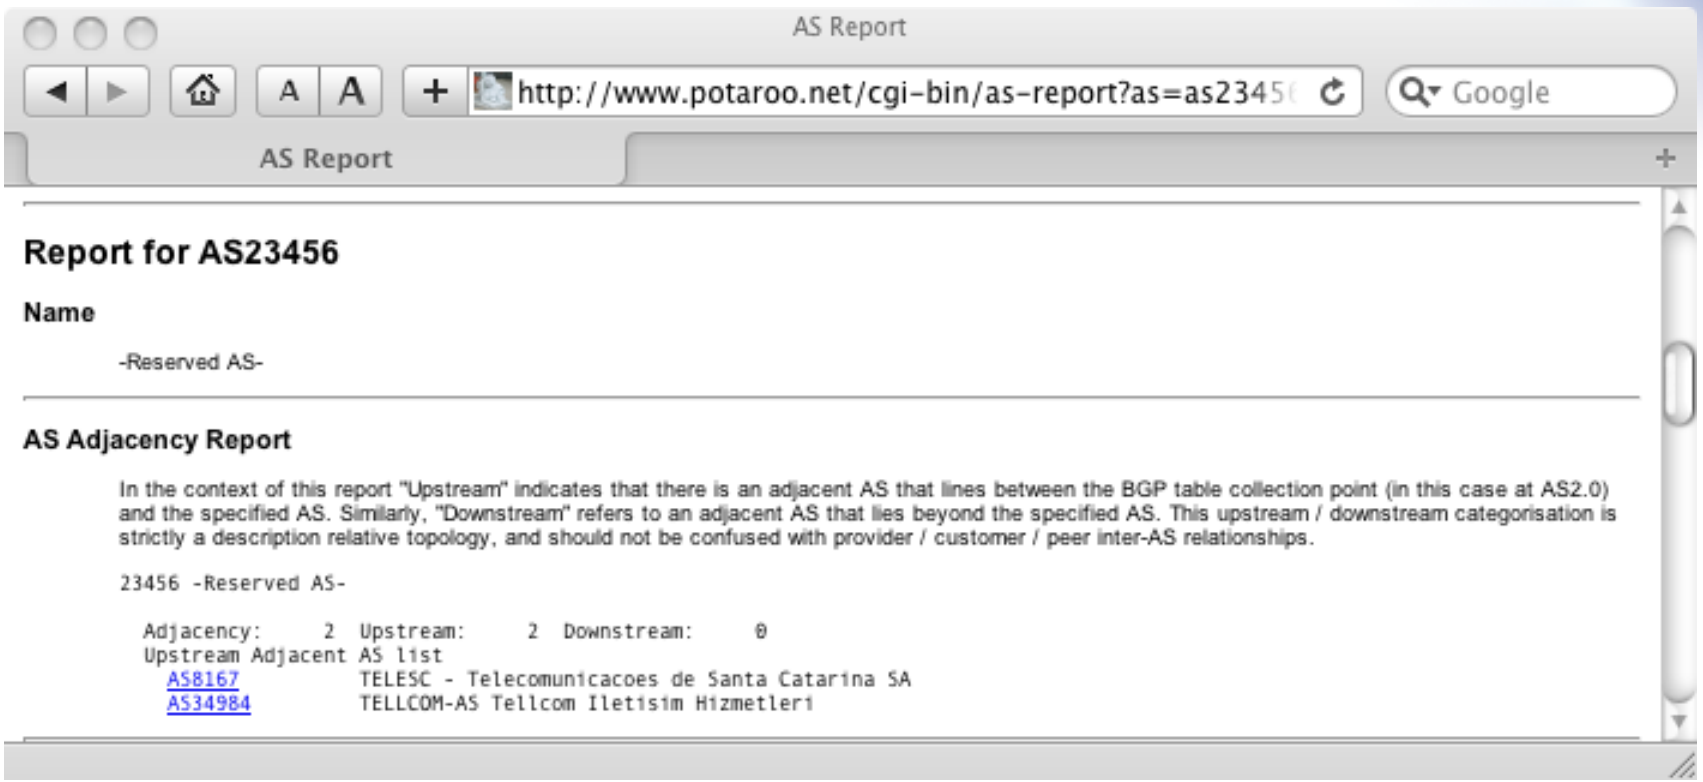

32-bit ASN Resources

- **IETF Specifications**

  RFC4893 – the 4-byte AS specification

  draft-ietf-idr-rfc4893bis – working document that adds some further clarity and error handling to the specification

- **Documentation**

  Exploring AS Numbers – Internet Protocol Journal, Vol 9, No 1
  (http://www.cisco.com/web/about/ac123/ac147/archived_issues/ipj_9-1/autonomous_system_numbers.html)

- **Reports and Resources**

  The AS Reports

  http://www.potaroo.net/tools/asn16/

  http://www.potaroo.net/tools/asn32/

  ISP Resource Wiki for ASNs

  http://as4.cluepon.net

Questions?

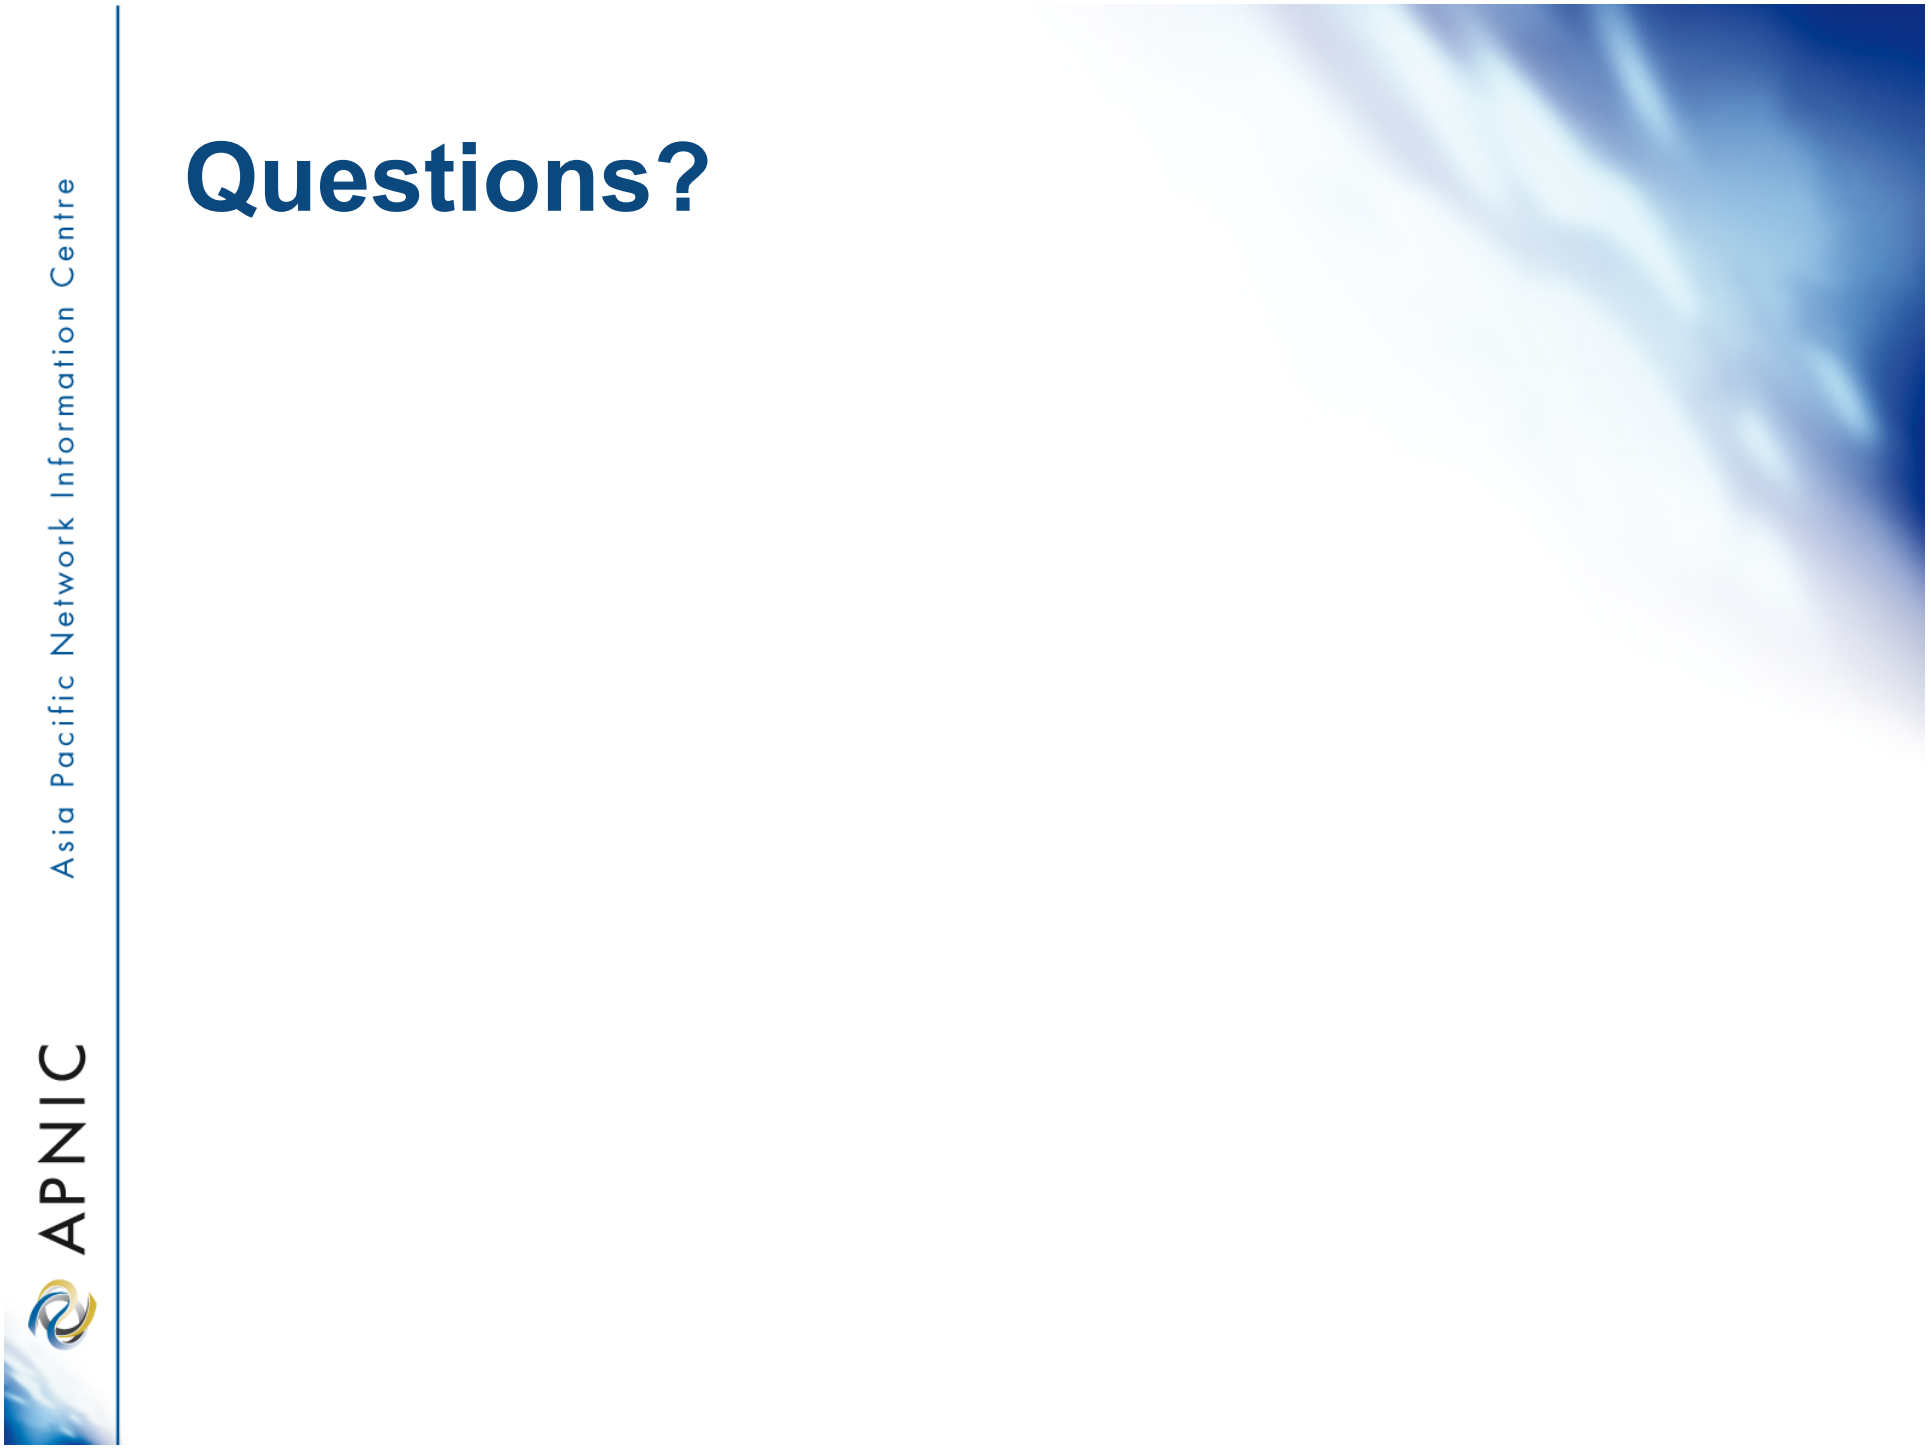

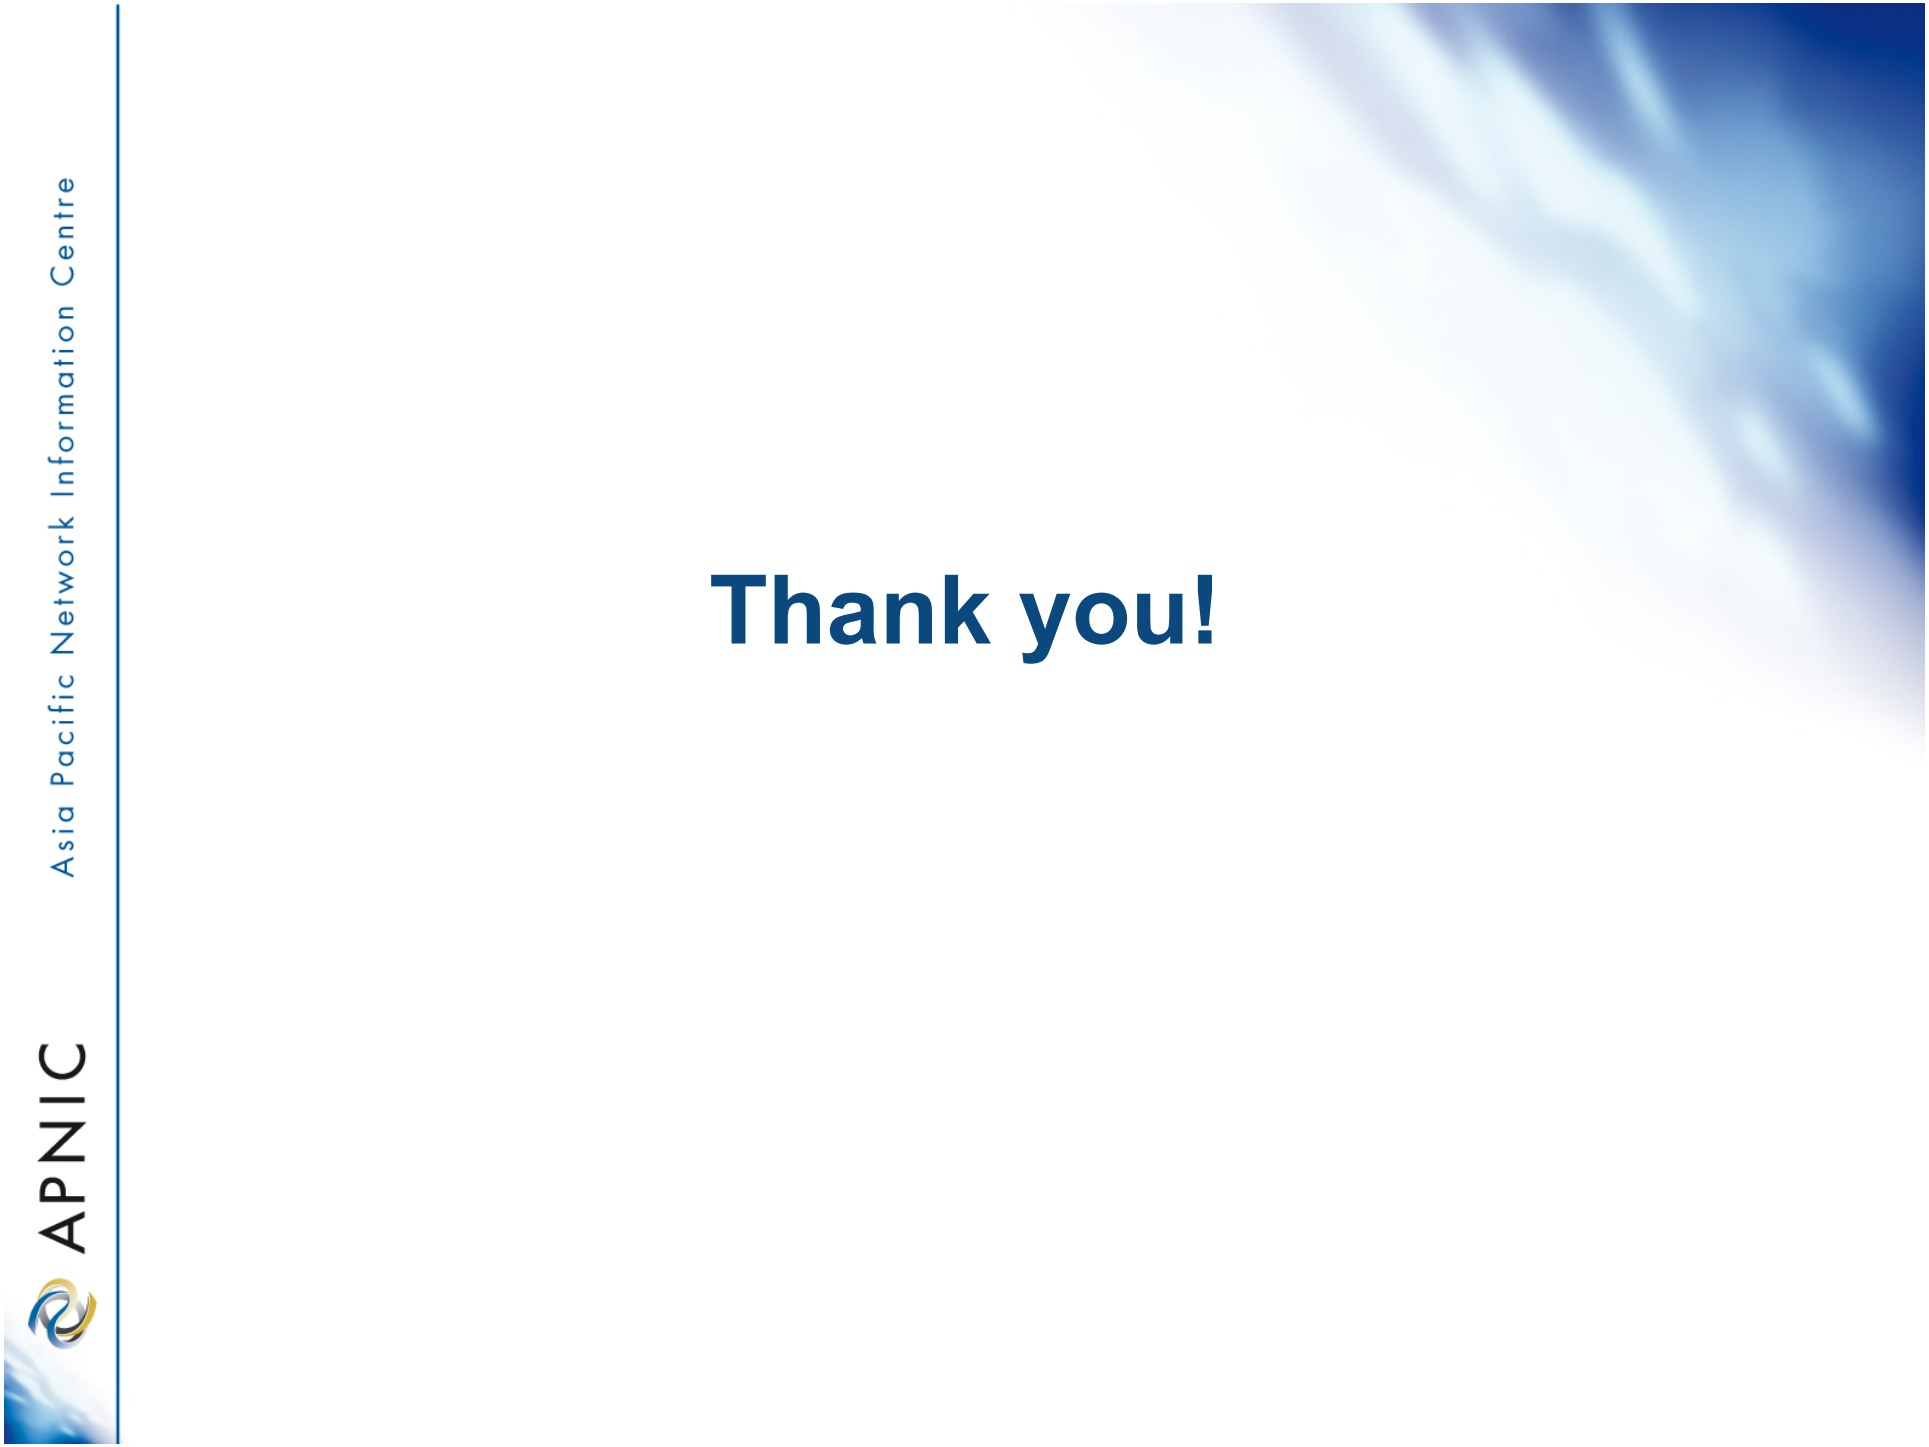
Asia Pacific Network Information Centre
APNIC
Thank you!