



# IPv4 and Two-byte ASNs running out How to craft the Internet beyond?


SANOG 12  
Kathmandu, Nepal  
13 August 2008

Srinivas (Sunny) Chendi

# Acknowledgements



Geoff Huston  
Chief Scientist  
APNIC




Intec NetCore, Inc.

<http://www.potaroo.net>

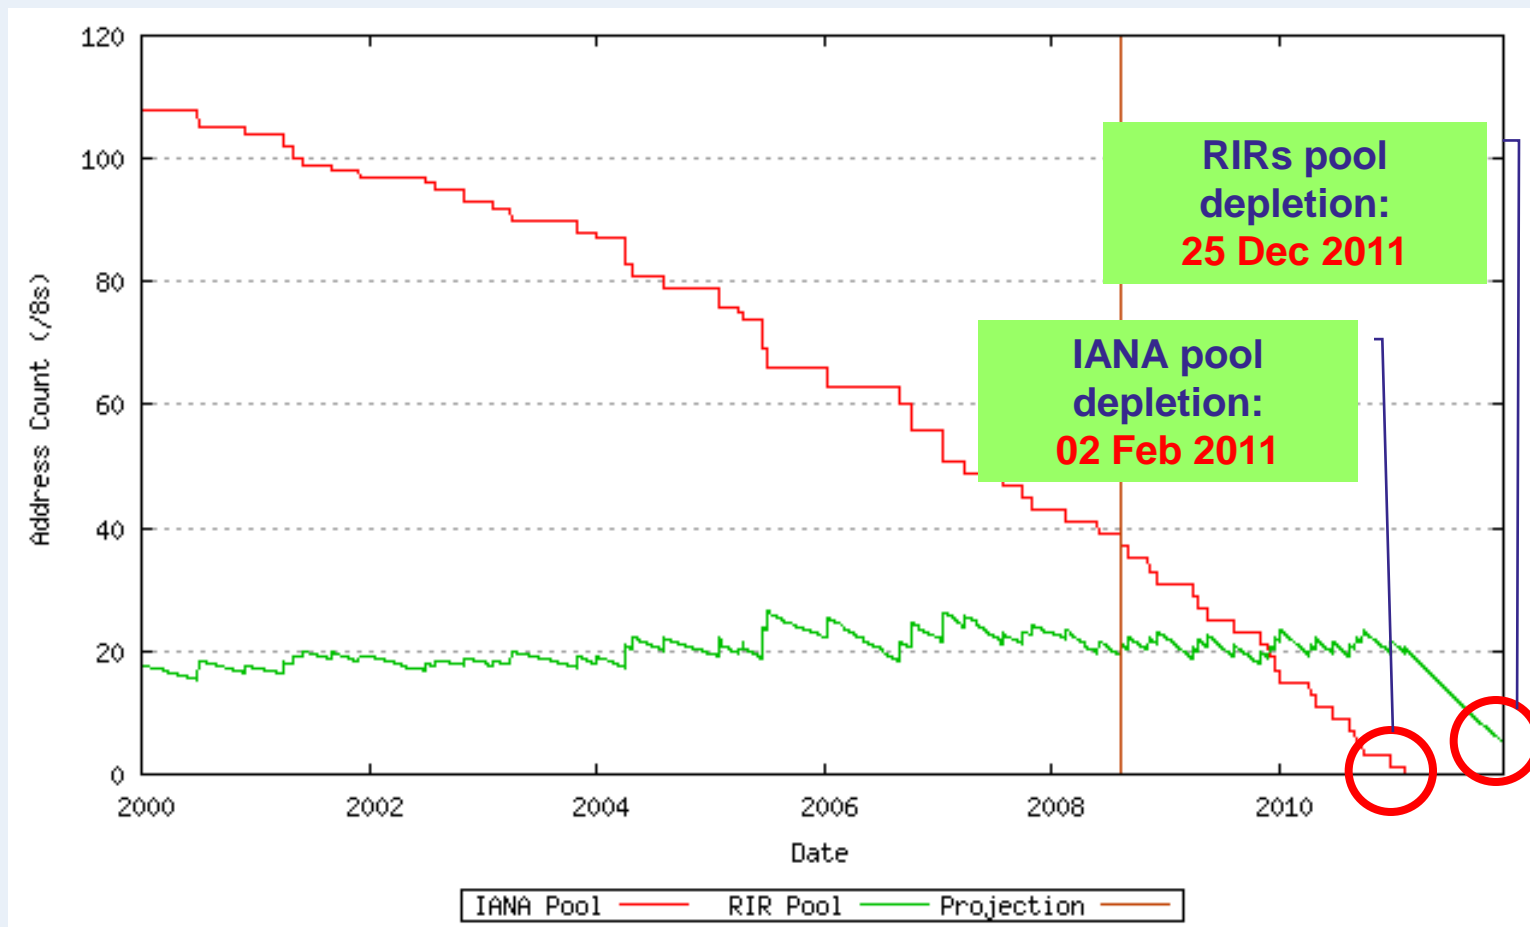
# IPv4 address distribution

# Current distribution of the whole IPv4 address space



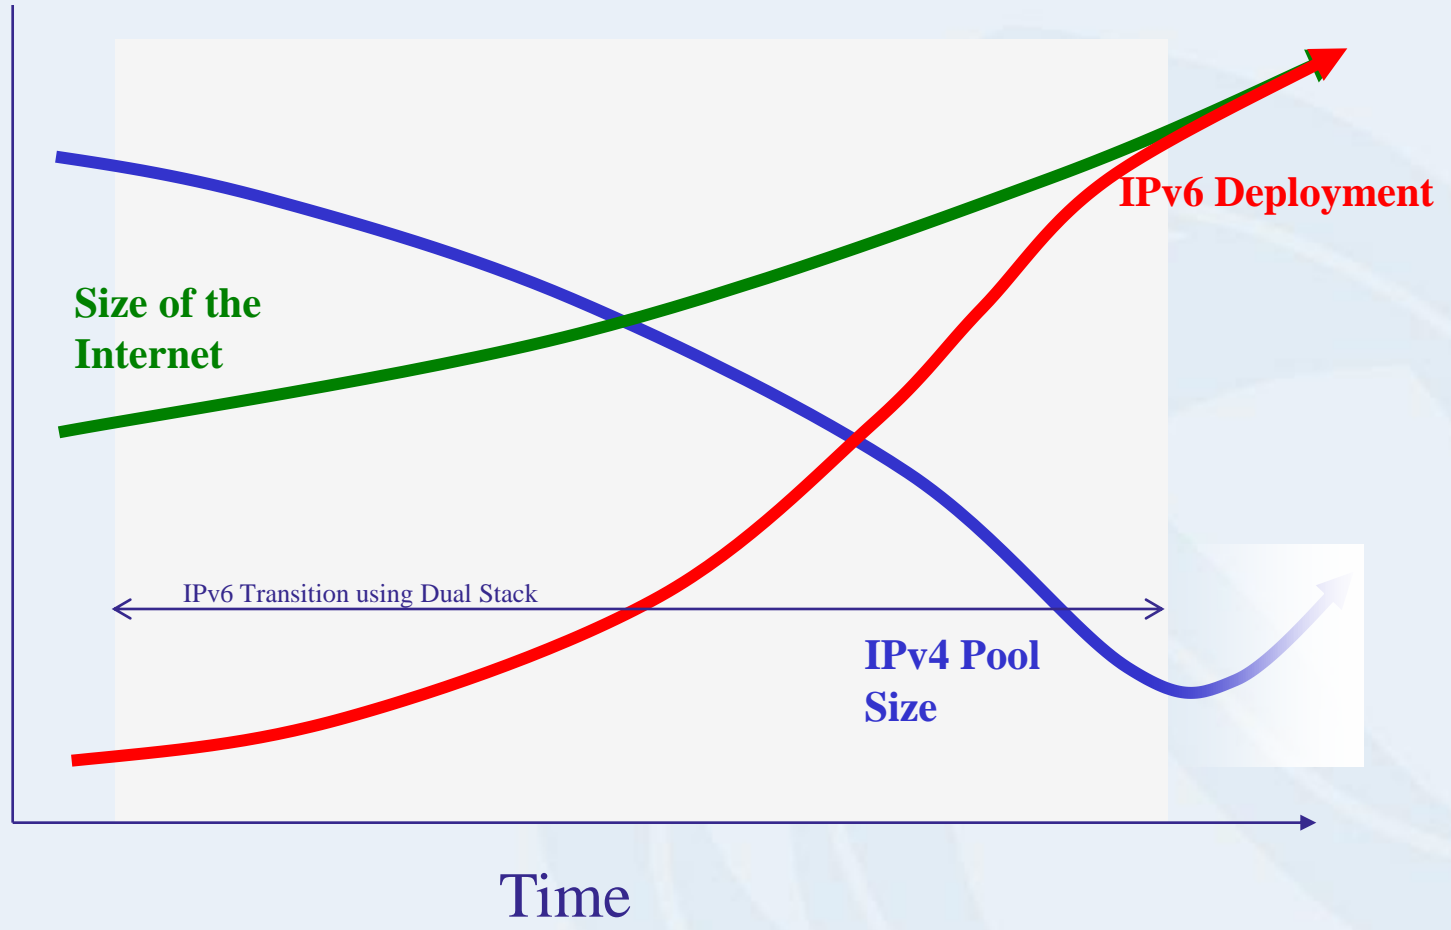
Source: <http://www.potaroo.net> (As of this date)

# Projected lifetime of remaining IPv4 addresses

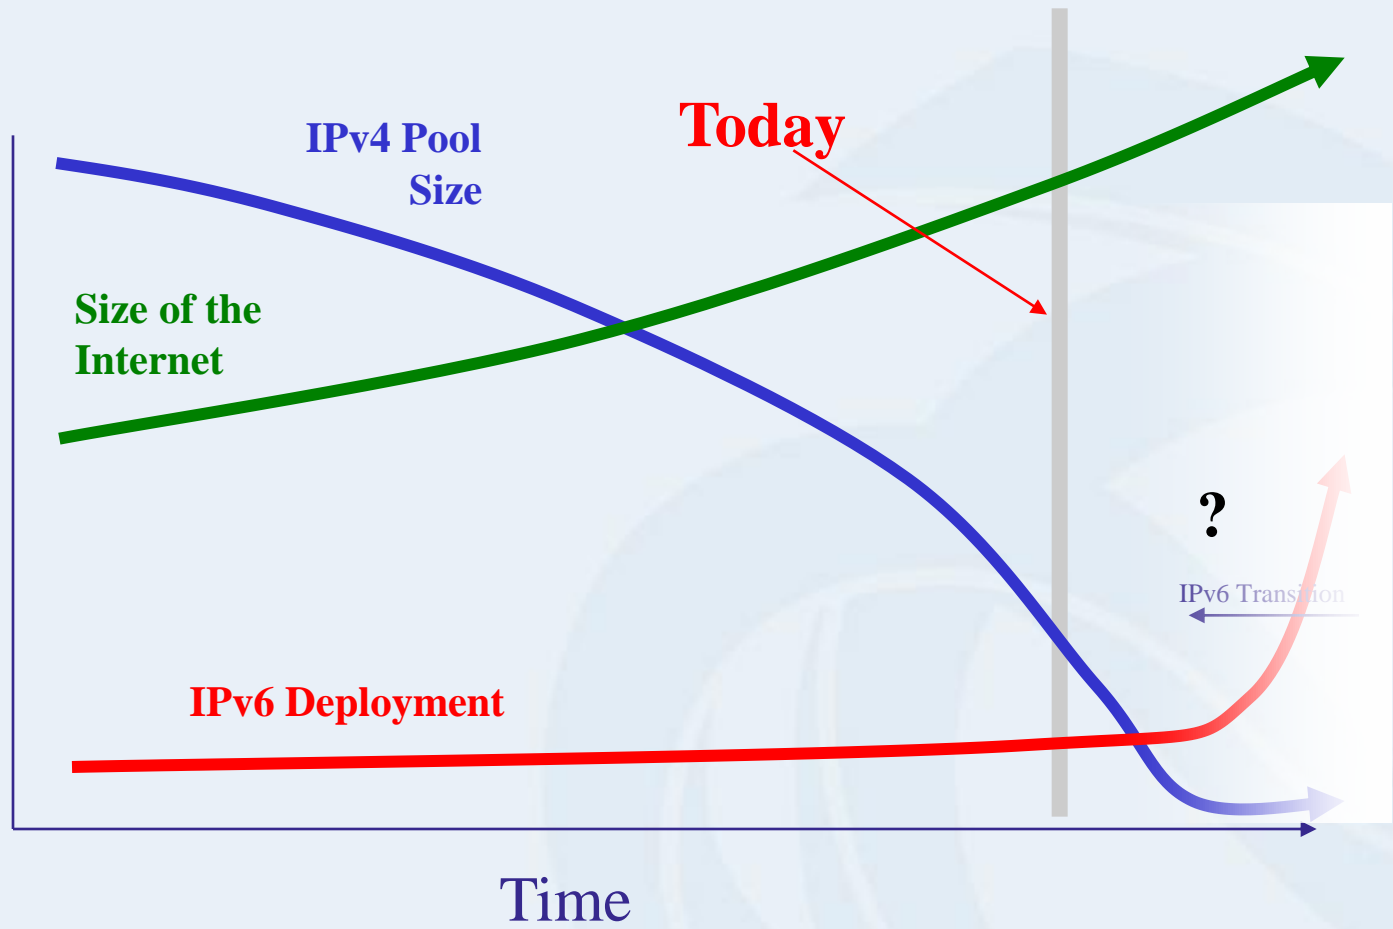


Source: <http://www.potaroo.net> (As of this date)

# We had a plan ...



# What's the revised plan?


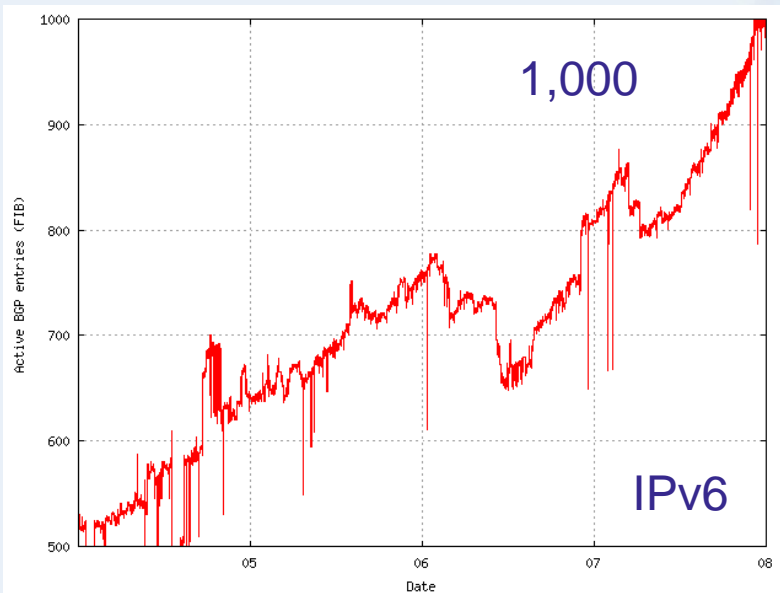
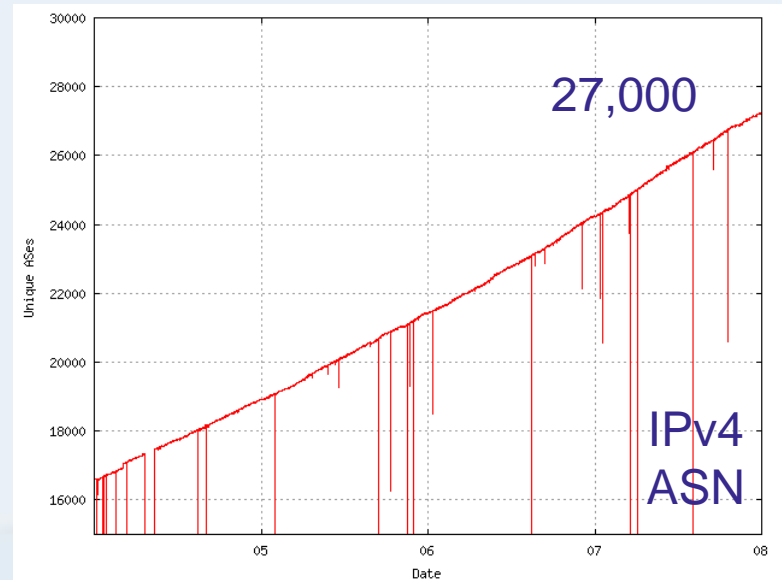
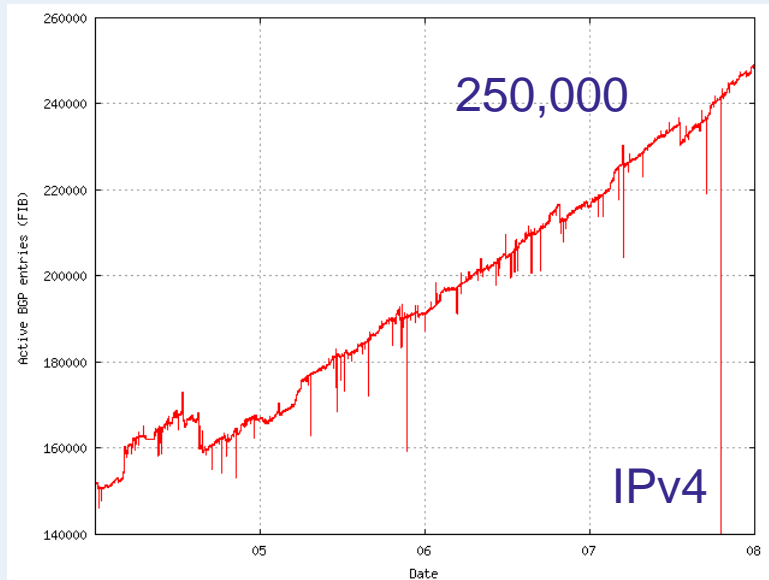


Source: <http://www.potaroo.net> (As of this date)

Is IPv6 actually in use?



# Yes, it is, far less than IPv4 but growing!



Source: <http://www.potaroo.net> (As of this date)

What are beyond the *depletion*?

# How can we expand the Internet after the IPv4 address depletion?

- Procuring global IPv4 address *by any means*
- Deploying IPv6 for new users
- Using NAT not to use global IPv4 addresses

Is IPv4 address any longer available?

## Not so longer, not always

- The current free pool is being depleted in 2010 - 2011
- Re-circulated IPv4 address will not always be supplied
  - Returning unused IPv4 address DOES COST. Available space by reclamation will be QUITE LIMITED.
  - A market for second-hand IPv4 address *might* emerge, but the supply is NOT COMMITTED.

Then, don't we need to deploy IPv6?

Yes, we do.

- Why?
- Simply, servers connected via NATs cannot be reached to meet end-to-end connectivity
  - Internet users benefit from cool services on servers. Not from the network itself.

## Frequently heard but questionable arguments – 1&2

- *IPv4 address depletion? I don't care since I'll make much more use of NAT*
- *IPv4 address depletion? I don't care since I've already got more than sufficient IPv4 address space.*
- You must care. Your customers will have more and more destinations which they cannot get through.



## Frequently heard but questionable arguments – 3

- *IPv6? Yet no one uses. Why and for whom should we deploy it? The cost will never be justified.*
- It is not a brand-new service only to extend your business. IPv4 address depletion is a CRISIS, and IPv6 is the only sustainable countermeasure.

## An IPv6 revolution...

- “Internet for Everything” instead of *Everyone*
- Serving the communications requirements of a device-dense world
- Device population some 2–3 orders of magnitude larger than today’s Internet
- Service costs must be cheaper by 2-3 orders of magnitude – per packet

# IPv6 – From PC to IPOD to iPOT...

- A world of billions of chattering devices



- Or even trillions...


# Two-byte ASNs

IPv4 not the only protocol running out of numbers

## 2-byte ASN and 4-byte ASN format


- Two-byte ASN (16-bit)  
0 ~ 65535
- Four-byte ASN (32-bit)  
0.0 ~ 65535.65535
- APNIC four-byte ASN range  
2.0 ~ 2.1023

# Current distribution of two-byte ASNs



Source: <http://www.potaroo.net> (As of this date)

# Projected lifetime of remaining two-byte ASNs



Source: <http://www.potaroo.net> (As of this date)

## RIRs and 32-bit AS Numbers

- From **1 January 2007** the RIRs are allocating 32-bit AS numbers (upon specific request)
- From **1 January 2009** the RIRs will be allocating 32-bit AS numbers by default (leaving some 16-bit AS numbers available upon specific request)



# What does this imply?

If you are using 16-bit AS  
as most (all) of you are today

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

something you probably want to avoid (or at least defer!)

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

***NOTHING*** changes!

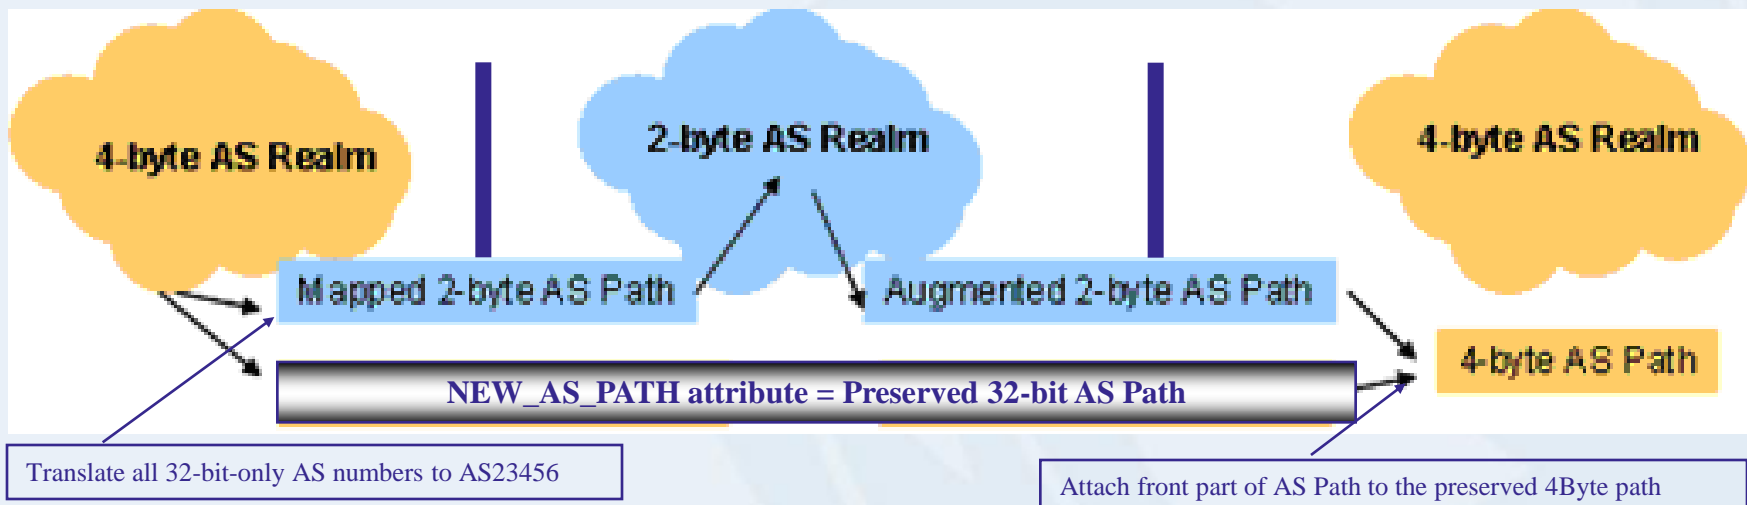
Well, **almost** nothing!

## What's changed?

- BGP Update messages in the 16-bit world
  - May contain “lies” in parts of the AS Path
  - May be larger in size due to tunneled additional information
- But prefix reachability information is still communicated between 16-bit and 32-bit BGP “realms”

# 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 / 16-bit BGP Example...




**AS Path** in the RIB (Routing Information Base)

i

# 32-bit / 16-bit BGP Example...




AS Path in the RIB



# 32-bit / 16-bit BGP Example...



AS Path in the RIB



AS Path Attribute in the UPDATE Message



# 32-bit / 16-bit BGP Example...



AS Path in the RIB

i

2.0



AS Path Attribute in the UPDATE Message

2.0

23456 23456

NEW\_AS\_PATH Attribute in the UPDATE Message


2.2 2.0




# 32-bit / 16-bit BGP Example...



AS Path in the RIB



AS Path Attribute in the UPDATE Message



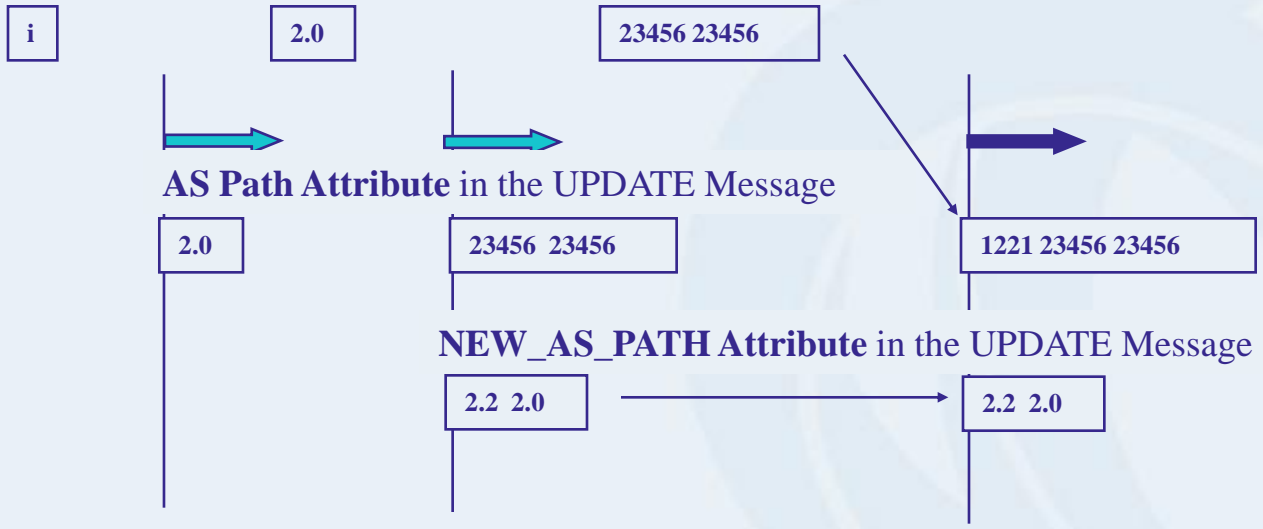
NEW\_AS\_PATH Attribute in the UPDATE Message



# 32-bit / 16-bit BGP Example...



AS Path in the RIB



# 32-bit / 16-bit BGP Example...




AS Path in the RIB




AS Path Attribute in the UPDATE Message



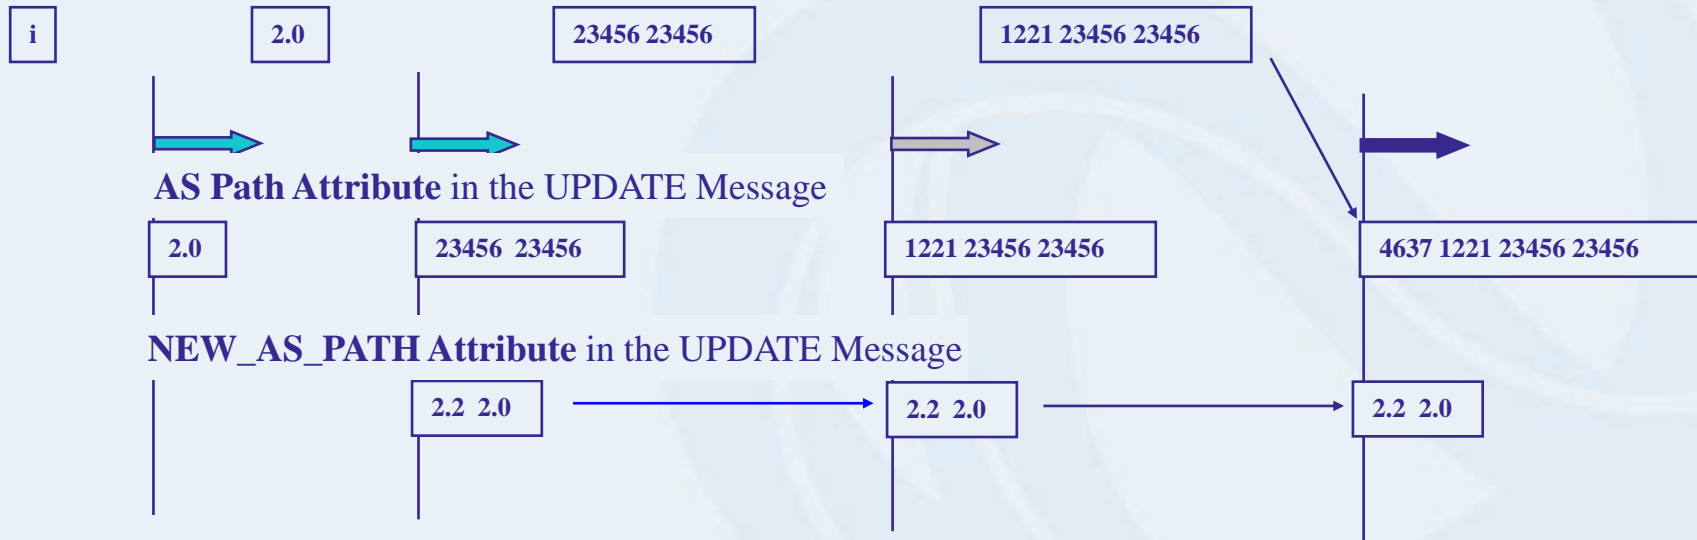
NEW\_AS\_PATH Attribute in the UPDATE Message



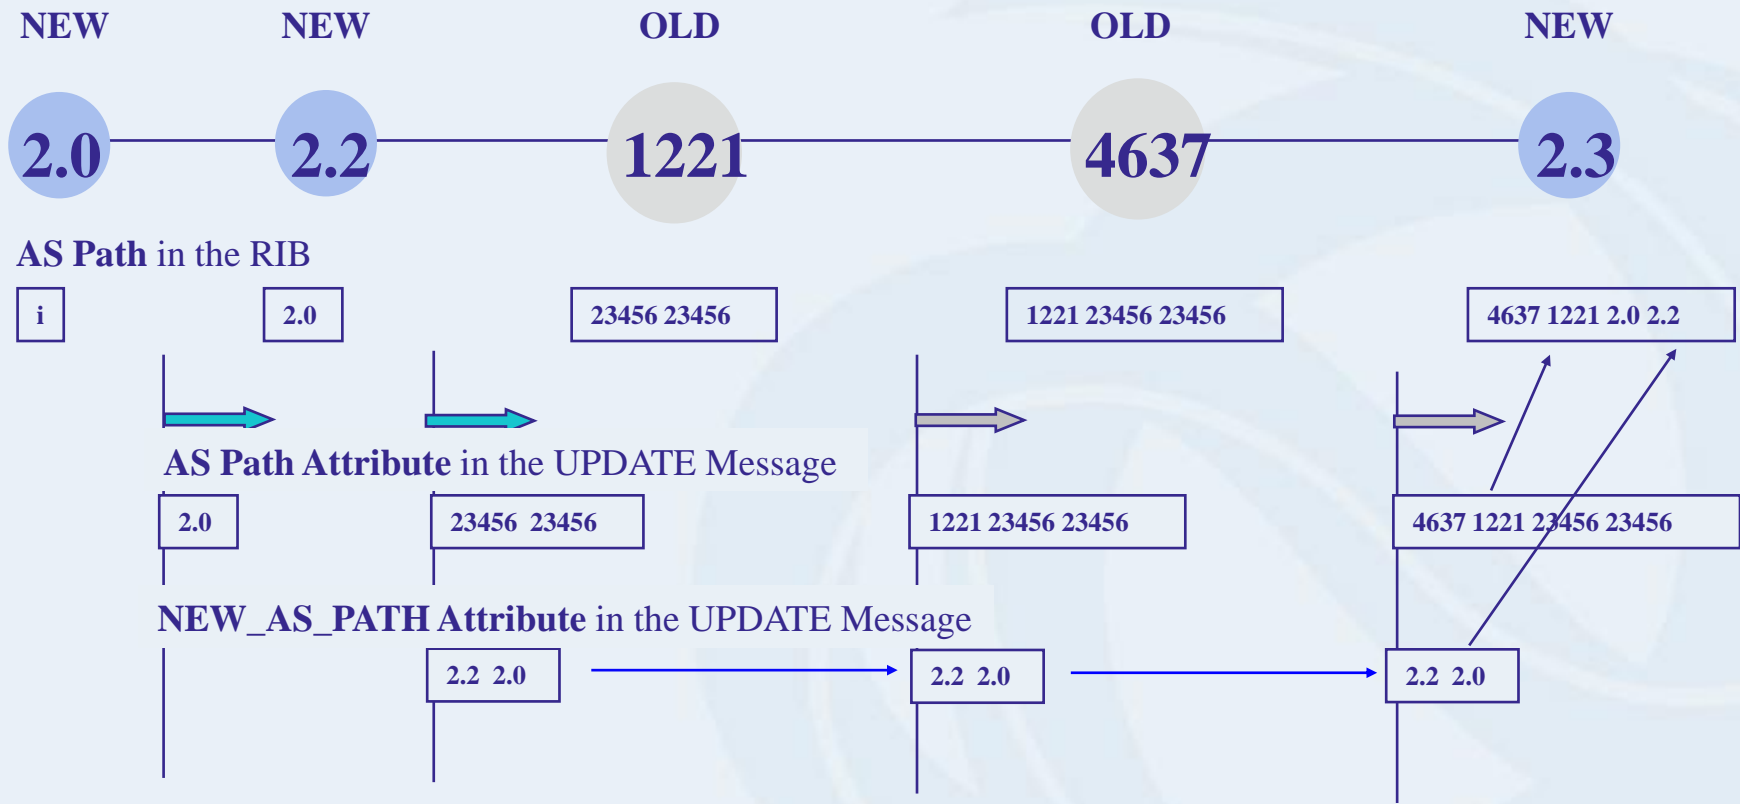
# 32-bit / 16-bit BGP Example...




AS Path in the RIB



# 32-bit / 16-bit BGP Example...



# 32-bit / 16-bit BGP Example



# 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*

# The Route-Views View

```

route-views.oregon-ix.net>show ip bgp 203.10.62.0/24
BGP routing table entry for 203.10.62.0/24, version 177310093
Paths: (43 available, best #39, table Default-IP-Routing-Table)
Not advertised to any peer
3277 3216 3549 4637 1221 23456
  194.85.4.55 from 194.85.4.55 (194.85.4.16)
    Origin IGP, localpref 100, valid, external
    Community: 3216:3000 3216:3004 3277:3216 3549:2141 3549:30840
7500 2497 4637 1221 23456
  202.249.2.86 from 202.249.2.86 (203.178.133.115)
    Origin IGP, localpref 100, valid, external
2493 3602 812 812 4637 1221 23456
  206.186.255.223 from 206.186.255.223 (206.186.255.223)
    Origin IGP, localpref 100, valid, external
2905 701 1239 4637 4637 4637 4637 4637 4637 1221 23456
  196.7.106.245 from 196.7.106.245 (196.7.106.245)
    Origin IGP, metric 0, localpref 100, valid, external
...
  
```



# 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 1.2** 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
  - Your OSS might get terminally confused!

## Mixed environments

- No dynamic capability for 16/32-bit ASN mode shift
  - You cannot flick from “16-bit OLD” to “32-bit NEW” mode within an active BGP session
  - You need to clear the session and then perform a clean start to trigger the initial capability exchange

## 4 Byte AS Testing

- Tests have been undertaken using closed BGP networks, and over the public Internet
- Tests of 16-bit/32-bit transition boundaries in various permutations of transits and loops
- Current announcement of 203.10.62.0/24 originating from AS 2.2 to assist others in local testing of 32-bit BGP

# 32-bit Path Reconstruction

```
srv0# bgpctl show rib 203.10.62.0/24
flags: * = Valid, > = Selected, I = via IBGP, A = Announced
origin: i = IGP, e = EGP, ? = Incomplete

flags destination gateway lpref med aspath origin
*> 203.10.62.0/24 147.28.0.1 100 0 0.3130 0.1239 0.4637
0.4637 0.4637 0.4637
0.4637 0.4637 0.1221
1.202 i
```

*Experiment performed on January 11 2007, with the assistance of Randy Bush and George Michaelson, using OpenBGPD 3.9 with 4Byte AS support patches as the origin and the observer points.*

# Resources

- IETF Specification
  - [RFC4893](#)
- OpenBGPD patches
  - <http://www.potaroo.net/tools/bgpd>
- Quagga patches
  - <http://quagga.ncc.eurodata.de/>

In conclusion...

## Possible steps YOU could take

- Staff training
  - Send staff to events like SANOG, APNIC and APRICOT to participate in training
  - Request APNIC to conduct workshop in your economy
- Request for IPv6 & 4-byte ASN from APNIC
  - IPv4 and IPv6 networks can co-exist
    - Most IPv4 software and hardware are IPv6 capable
  - No extra fees
    - Existing APNIC members with IPv4 space
- Start now
  - Transition takes time

## Possible steps YOU could take (cont)

- Join mailing lists to keep up to date on developments
  - APNIC mailing lists
    - <http://www.apnic.net/community/lists/index.html>
  - IPv6 global operator forum
    - <http://lists.cluonet.de/mailman/listinfo/ipv6-ops>
- Access relevant websites
  - Internet Community of Online Networking Specialists (ICONS)
    - Keep up to date on operational matters
    - <http://icons.apnic.net>
  - Global IPv6 forum
    - Latest events and information on IPv6 development
    - <http://www.ipv6forum.org>





Thank you!

[sunny@apnic.net](mailto:sunny@apnic.net)