



Beijing, China 24-28 August 2009

APNIC 28

AS Numbers - Again

Geoff Huston
APNIC

August 2009

research@APNIC.net



Current Status

- 2-byte AS Numbers
 - 31,750 AS numbers visible in BGP
 - 14,983 not advertised in BGP
 - 9,200 AS numbers left with IANA
 - 8,561 AS numbers in the RIR pools



AS Number Map




AS Number Map

Unadvertised AS Numbers

RIR Pool AS Numbers


Advertised AS Numbers

IANA Pool AS Numbers






Consumption Rate






Consumption Rate

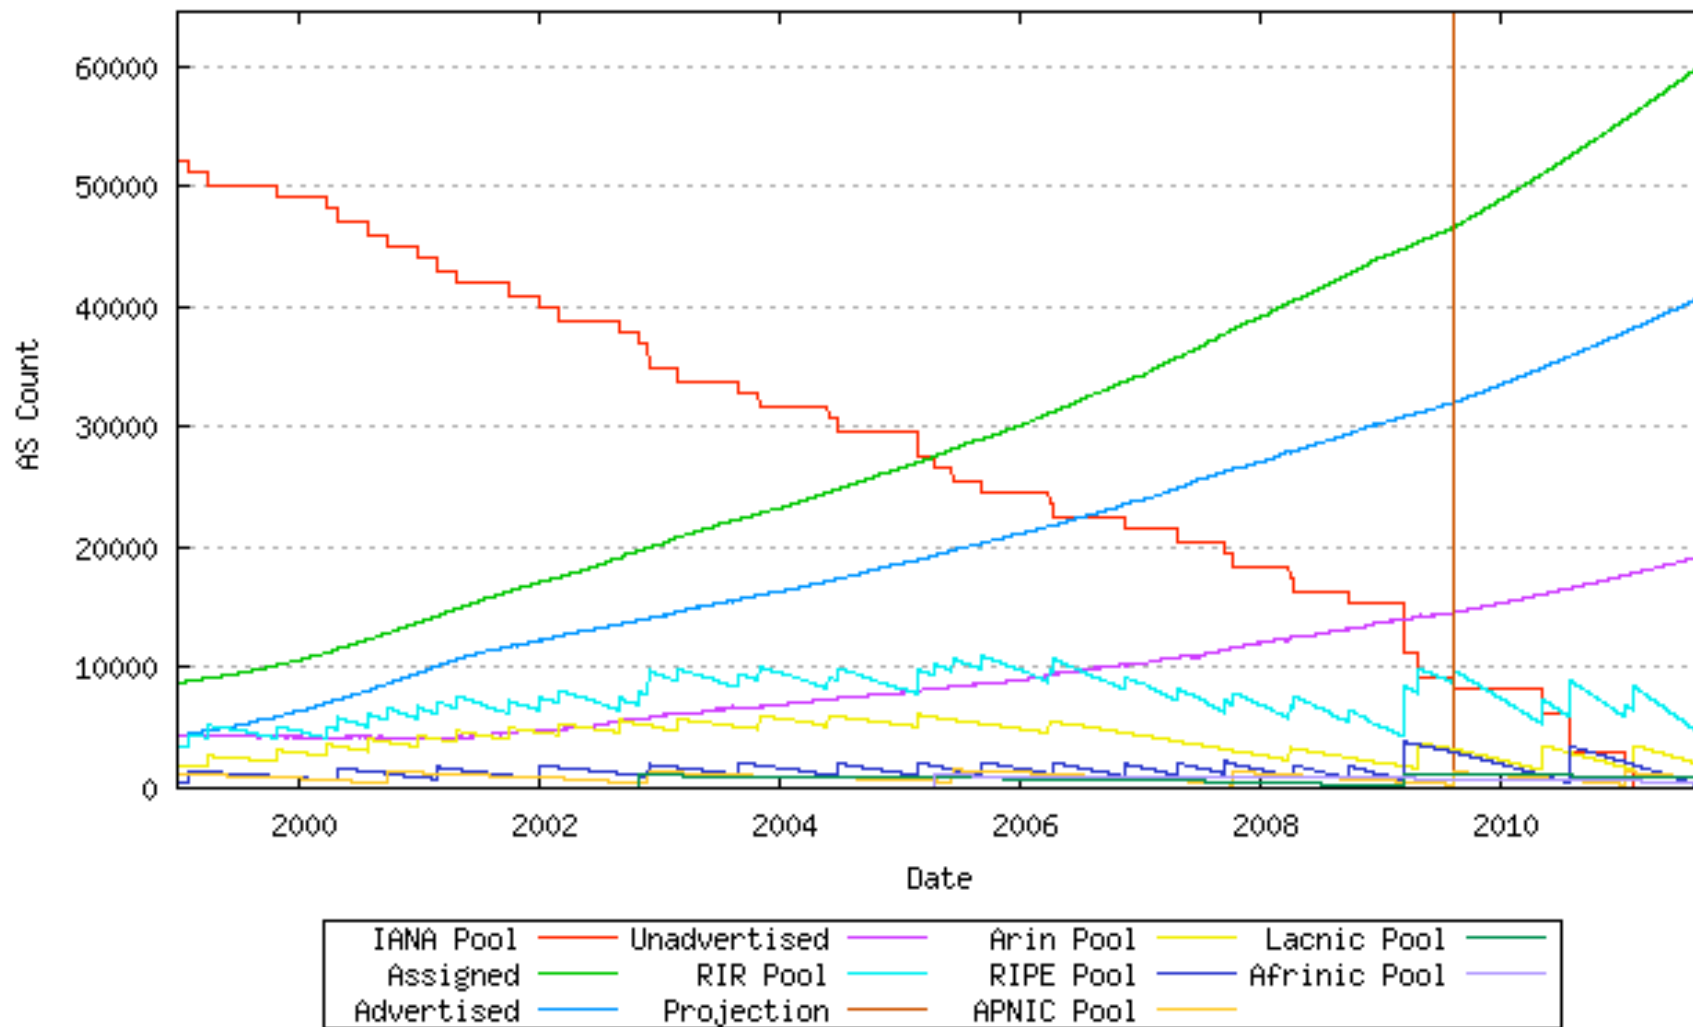





Consumption Rate



How long have we got?



How long have we got?






How long have we got?

- IANA will allocate its last 16-bit AS number block in **February 2011**
- RIPE will exhaust its 16-bit AS Number pool in **September 2011**

This is not exactly news!

Combining these views



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
~1/2 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
RFC 4893 May 2007
2. RIRs to start making 32-bit AS numbers available
32 bit AS Policy -2006
3. Vendors to provide 32-bit AS number capable BGP implementations
4. BGP networks to commence deployment



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)SX11, 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
May 2007
RFC 4893
2. RIRs to start making 32-bit AS numbers available
32 bit AS Policy -2006
3. Vendors to provide 32-bit AS number capable BGP implementations
*Currently gathering pace:
2009-2010*
4. BGP networks to commence deployment

RIR Allocation Data of 32-bit AS's





32-bit ASes in BGP



32-bit ASN Deployment

- Allocation status as of August 2009:
 - Advertised: 70
 - Unadvertised: 221
- In 2009 the RIRs allocated 2,683 ASNs up to the 12th August
 - 2,514 were 16-bit ASNs
 - 169 were 32 bit ASNs



The Agenda for Transition

1. IETF to complete BGP Standards to support transition mechanisms to 32-bit AS numbers
May 2007
RFC 4893
2. RIRs to start making 32-bit AS numbers available
32 bit AS Policy -2006
3. Vendors to provide 32-bit AS number capable BGP implementations
currently gathering pace
2009 - 2010
4. BGP networks to commence deployment
Lagging - badly!

How can we assist with 4-byte 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



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

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?*

- BGP uses a translation approach to mapping 4-byte AS numbers into a 2-byte AS number
- The 4 byte BGP speaker does all the translation work, so the existing BGP world will not need to upgrade to “see” these additional networks that lie within 4-byte ASNs in the routing space
- All that you will see is:
 - AS 23456 appearing in many AS paths
 - A very minor increase in memory use by BGP - associated with the storage of the additional AS4_PATH attribute
 - which contains the 4-byte AS path
 - but its an opaque transitive attribute to you, so you don't care about its contents

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

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

- You need to do **nothing!**
- The new 4-byte BGP speaker figures out its talking to your old 2-byte BGP speaker and the 4-byte BGP speaker does **all** the work
 - it translates all instances of 4 byte AS numbers in the AS Path and Aggregator attributes to 23456 and stores the original 4-byte AS Path and Aggregator in new opaque transitive attributes (tunneling) before sending you the update
 - and restores the 4-byte information in any updates it received from you from the tunneled attribute information

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?*

- **But** you should've checked out your operational support system by now to make sure it can cope:
 - because you will need to support multiple peers / customers / upstreams who will have 4-byte AS numbers
 - and you will want to differentiate between them
 - but your routers' BGP configs will be peering with AS 23456 for each instance
 - so your support system better be able to work this all out and not get confused!

Common Questions

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

Common Questions

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

NO – not yet

Common Questions

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

- NO, 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 draft-ietf-l3vpn-as4octet-ext-community
 - not yet an RFC - currently in IESG review
- Ask your vendor when they will be supporting BGP extended communities with 4-byte ASNs

Common Questions

4. If I upgrade BGP, will BGP crash?

Common Questions

4. If I upgrade BGP, will BGP crash?

MAYBE!

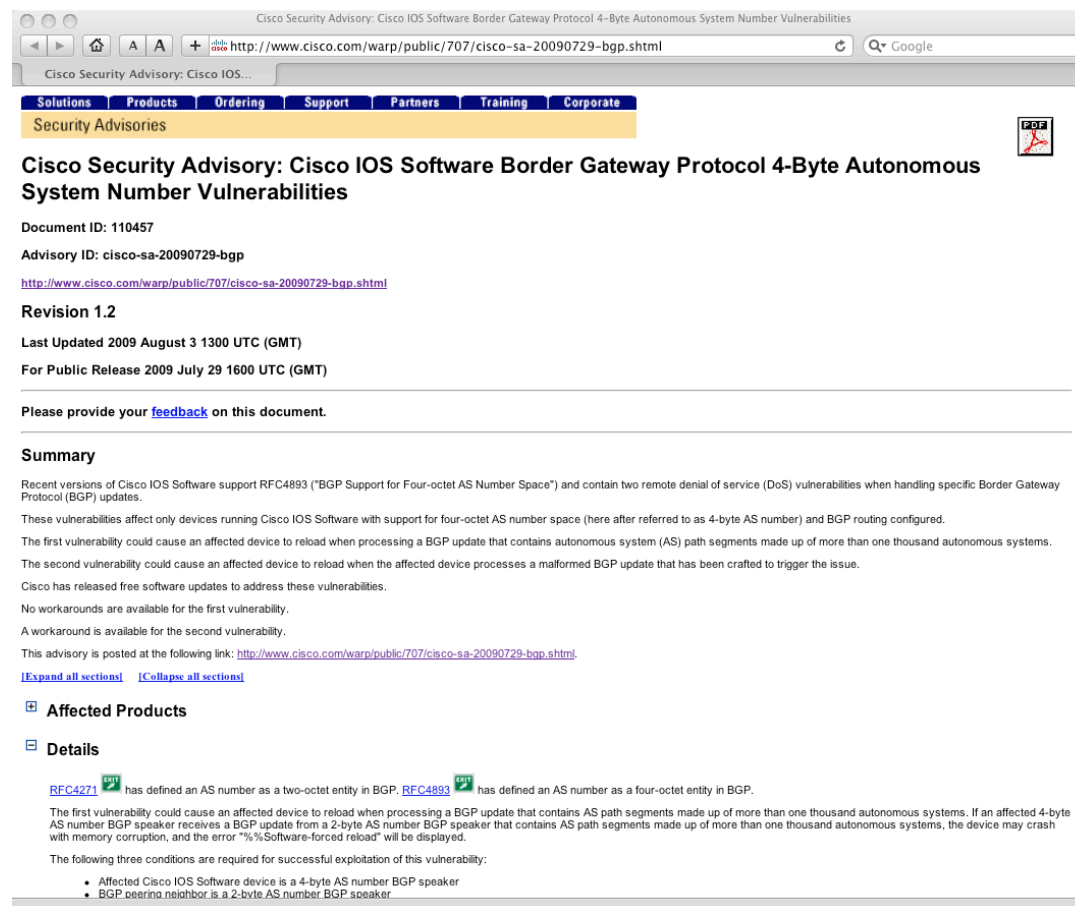
Common Questions

4. If I upgrade BGP, will BGP crash?

- Some Cisco implementations of BGP with 4-byte ASN support get unhappy when the number of elements in the AS path gets to over 1,000
- The maxas-limit setting is your friend

Common Questions

4. If I upgrade BGP, will BGP crash?



Cisco Security Advisory: Cisco IOS Software Border Gateway Protocol 4-Byte Autonomous System Number Vulnerabilities

Document ID: 110457
Advisory ID: cisco-sa-20090729-bgp
<http://www.cisco.com/warp/public/707/cisco-sa-20090729-bgp.shtml>

Revision 1.2
Last Updated 2009 August 3 1300 UTC (GMT)
For Public Release 2009 July 29 1600 UTC (GMT)

Please provide your [feedback](#) on this document.

Summary

Recent versions of Cisco IOS Software support RFC4893 ("BGP Support for Four-octet AS Number Space") and contain two remote denial of service (DoS) vulnerabilities when handling specific Border Gateway Protocol (BGP) updates.

These vulnerabilities affect only devices running Cisco IOS Software with support for four-octet AS number space (here after referred to as 4-byte AS number) and BGP routing configured.

The first vulnerability could cause an affected device to reload when processing a BGP update that contains autonomous system (AS) path segments made up of more than one thousand autonomous systems.

The second vulnerability could cause an affected device to reload when the affected device processes a malformed BGP update that has been crafted to trigger the issue.

Cisco has released free software updates to address these vulnerabilities.

No workarounds are available for the first vulnerability.

A workaround is available for the second vulnerability.

This advisory is posted at the following link: <http://www.cisco.com/warp/public/707/cisco-sa-20090729-bgp.shtml>

[Expand all sections](#) | [Collapse all sections](#)

Affected Products

Details

[RFC4271](#) has defined an AS number as a two-octet entity in BGP. [RFC4893](#) has defined an AS number as a four-octet entity in BGP.

The first vulnerability could cause an affected device to reload when processing a BGP update that contains AS path segments made up of more than one thousand autonomous systems. If an affected 4-byte AS number BGP speaker receives a BGP update from a 2-byte AS number BGP speaker that contains AS path segments made up of more than one thousand autonomous systems, the device may crash with memory corruption, and the error "%Software-forced reload" will be displayed.

The following three conditions are required for successful exploitation of this vulnerability:

- Affected Cisco IOS Software device is a 4-byte AS number BGP speaker
- BGP peer/neighbor is a 2-byte AS number BGP speaker

Common Questions

4. If I upgrade BGP, will BGP crash?

Also, there is the issue of the “standard” method for handling invalid components in the the AS4_PATH attribute

- AS Confederation path segments are declared invalid in the AS4_PATH attribute (RFC4893)
- If an optional attribute in an UPDATE is recognised then it must be checked, and if it is detected as invalid then a NOTIFICATION message must be sent and the BGP session is closed (RFC4271)
- A literal implementation of 4-byte AS BGP will be triggered to repeatedly tear down the local BGP session if AS Confederation elements are added into the AS PATH by a 4-byte AS BGP speaker, and then immediately propagated to a 2-byte AS BGP peer



Common Questions

4. If I upgrade BGP, will BGP crash?

The “safest” option is for the 4-byte BGP speaker to remove the offending element and reconstruct the AS Path as best it can, and log the error

- Which appears to be what many BGP implementations now do
- And this consideration of “soft handling” of update errors applies to any BGP update, not only those with the AS4_PATH attribute, such as the use of AS0 in an AS Path
- The IETF is working on refining the BGP specification to treat such BGP update attribute errors with some circumspection, rather than a rather brutal “just drop the session” response!

Common Questions

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

Common Questions

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

Calm down!



Common Questions

5.1 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

```
Terminal — ssh — 105x10
[gih@wattle /var/data/bgp/as2.0]$ grep 23456 bgptable.txt
* 193.169.148.0/23 203.119.76.3 0 4608 1221 4637 3356 20960 23456 i
*> 202.12.28.190 0 4777 2516 3356 20960 23456 i
*> 194.0.68.0/22 202.12.28.190 0 4777 2516 3491 20485 23456 i
* 203.119.76.3 0 4608 1221 4637 3491 20485 23456 i
*> 195.88.154.0/23 202.12.28.190 0 4777 2497 9002 44237 34267 23456 i
* 203.119.76.3 0 4608 9942 6453 1299 44237 34267 23456 i
* 203.217.144.0/22 203.119.76.3 0 4608 1221 4637 1273 9498 23456 i
*> 202.12.28.190 0 4777 2516 7473 9498 23456 i
[gih@wattle /var/data/bgp/as2.0]$
```

Common Questions

5.1 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

The screenshot shows a web browser window titled "AS Report" with the URL <http://www.potaroo.net/cgi-bin/as-report?as=as23456>. The page content includes:

- Report for AS23456**
- Name**: -Reserved AS-
- AS Adjacency Report**
- Text: "In the context of this report "Upstream" indicates that there is an adjacent AS that lines between the BGP table collection point (in this case at AS2.0) and the specified AS. Similarly, "Downstream" refers to an adjacent AS that lies beyond the specified AS. This upstream / downstream categorisation is strictly a description relative topology, and should not be confused with provider / customer / peer inter-AS relationships."
- Summary: "23456 -Reserved AS-"
- Adjacency statistics: "Adjacency: 4 Upstream: 4 Downstream: 0"
- Upstream Adjacent AS list:
 - [AS34267](#): DEBRYANSK-AS-1 JSC "CenterTelecom", Bryansk branch
 - [AS20960](#): TKTELEKOM-AS Telekomunikacja Kolejowa is an ISP operating in Poland
 - [AS20485](#): TRANSTELECOM JSC Company TransTeleCom
 - [AS9498](#): BBIL-AP BHARTI Airtel Ltd.

Common Questions

5.1 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

- The AS Path is used for loop detection and path metric
- Even when AS23456 appears in the AS path, routing loops cannot form in BGP
 - but such “hybrid” loops may take a few more AS hops to detect and kill

Common Questions

6. Are there AS Bogons in the 4-byte space?

Common Questions

6. Are there AS Bogons in the 4-byte space?

Yes!

6% of the 4-byte ASNs in BGP are bogons

Common Questions

6. Are there AS Bogons in the 4-byte space?

- Advertised 4-byte ASNs: 70
- Advertised Bogons: 4
 - 196636 advertised by AS 29608 – WAN2MANY
 - 262657 advertised by AS 12956 - Telefonica
 - 393392 advertised by AS 12874 - Fastweb
 - 2076901376 advertised by AS 43314 – DIANET

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>



APNIC 28 Beijing, China 24-28 August 2009

APNIC 28

Thank You

Questions?