APNIC Apster

*Apster* is the quarterly newsletter for APNIC members and the Internet community.

Issue 20 - January 2007

# IGF Athens: Democracy, philosophy, arranged marriages, and dialogue

The first Internet Governance Forum (IGF) was held in Athens, Greece, in November 2006. **Gerard Ross** describes the role played by the NRO in this historic meeting and reviews some of the perceptions that have been expressed about the experience.

It was unexpectedly described as resembling an arranged marriage. On the closing day of the first Internet Governance Forum (IGF), Nitin Desai, the Indian-born Chair of the Forum found a surprisingly appropriate metaphor to sum up the intensive four-day program.

"In my country," he said, "when people get married, we have arranged marriages, and usually the first meeting between the boy and the girl, they are scoping each other out... the conversation tends to cover everything, you see. And at the second and the third meeting they start talking about more specific things, what are your tastes in this area or that area. And it is some time before they actually start holding hands. So let's just treat this as a first meeting where people have just gotten to know one another and maybe it will lead to marriage".

No one had really known what to expect from the first IGF, the new forum created by former UN Secretary-General Kofi Annan to carry on the dialogue started in the World Summit on the Information Society (WSIS). Certainly no one had expected the many and varied stakeholders to be holding hands by the end of the first week – and, of course, they were not. Nevertheless, many were pleasantly surprised by the spirit of the dialogue which characterised much of the meeting.

The Number Resource Organization (NRO) noted in 2006 that it looked forward to the IGF "with interest and expectation", hoping it would become "a useful tool for dealing with those issues which constitute real problems for the community and for which there are no adequate governance mechanisms".

The NRO had also expressed the view, shared by many others, that the "IGF must be a multi-stakeholder forum without decision-making attributions". For the forum to succeed the "different stakeholders must participate on a level playing field, all of them having exactly the same privileges. The archives of the IGF meeting minutes and documents must be accessible to anyone without the need for accreditation".

On this measure, the inaugural IGF can be considered a success. While the conduct of most of the sessions would have been quite familiar to people accustomed to meetings held in the "Internet

▲ Athens, venue for the first IGF, inspired many speakers to draw parallels between the development of democracy and issues of Internet governance.

tradition" – free seating, open microphones, remote participants, and actual dialogue – it was clearly unfamiliar territory for many participants used to more rigid intergovernmental and diplomatic practices. This lack of familiarity did lead to some teething problems. Many people throughout the week noted that the size of some of the panels, while allowing broad representation of the many stakeholder groups, did not always create the most focused or coherent discussions.

Indeed, there was a feeling, particularly in the early part of the programme, that many of the people present simply did not know how to talk to each other. Nitin Desai drew attention to this in his closing remarks, noting that in this new mix of cultures and styles, government representatives may need to be more open, civil society may need to be less vigourous and forceful, and business may need to accept more discussion of principles and generalities. There were signs of these adjustments beginning to occur, especially in the context of the smaller workshop sessions, where the nature of the discussions brought people from across the stakeholder groups closer to understanding of – if not necessarily agreement with – the needs of other groups.

The NRO participated officially in two of the workshops. The first, the "Participation Workshop", focused on encouraging multi-stakeholder participation in the organisations and mechanisms responsible for the management, administration, and development of the Internet. Its purpose was to convey the importance of participating in the processes of those organisations and mechanisms. The NRO is perfectly suited to a panel

P 2

# IPv4 exhaustion projections: Then and now

At APNIC 22 in Kaohsiung, Taiwan, the opening plenary session featured an active panel discussion, entitled 'IPv4 exhaustion: what's the real story?', following a similar discussion at ARIN XVI earlier in the year. In this article, **Tina Bramley** provides an overview of the main themes of these sessions.

The IPv4 address pool is, by definition, finite and the consumption of the pool is of critical importance to the Internet community. The issue is not new, but it is now in the spotlight more than ever before. This article examines IPv4 consumption by reviewing how we saw the situation three years ago, and how subsequent developments in the IT and Internet landscapes have affected IPv4 exhaustion projections and proposed strategies.

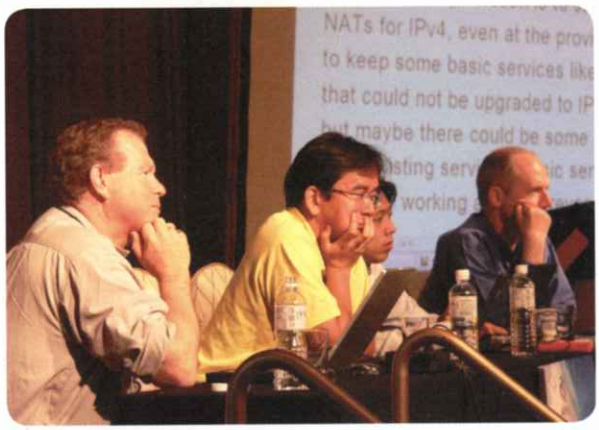


▲ Geoff Huston, Akinori Maemura, Tomoya Yoshida, and Paul Wilson on the APNIC 22 plenary panel "IPv4 exhaustion: What's the real story?"



## The view from 2003

### IPv4 consumption in the news

In August 2003, APNIC published the article 'IPv4 – How long have we got?' in *Apster* (issue 7). The article was written in response to mainstream media claims that exhaustion of the IPv4 address pool was imminent, particularly one sensationalist BBC news report that had indicated that IPv4 would run out in 2005.

At the time, another common, though mistaken, theme of media reports was that North America was immune to the exhaustion problem because of 'stockpiling', and that later-developing economies like China were facing difficulties getting the IPv4 address space they needed.

In issue 12 of *Apster*, APNIC published the article 'IP addressing in China' (December 2004) which debunked the myth that Asia Pacific nations, such as China, were unable to get IPv4 address space. The article explained how the myth grew from misinterpretations of the historical growth of the Internet, which began in North America and Europe. Claims that shortages could arise in any economy or region overlook the allocation policies of the RIRs, which are designed to ensure that everyone who needs IP addresses gets equal access to them. APNIC allocation statistics show that Asia Pacific economies were receiving IPv4 blocks faster than any other region. In 2003, APNIC was allocating as much IPv4 address space to members as was requested and that remains the case today; IPv4 space is still available to any organisation that satisfies the criteria set down in APNIC's allocation policies.

### Preliminary projections

Geoff Huston's paper 'IPv4 – How long have we got?' drew on three different data sources to examine consumption and usage patterns, and arrived at the following projections:

- IANA allocations of IPv4 address blocks to RIRs would end by 2019,
- RIR allocations of address blocks to LIRs and ISPs would end by 2026, and
- extrapolation of BGP routing table analysis indicated total exhaustion of the IPv4 address pool by 2029.

Using a modelling technique that combined these three basic forecasts and accounted for address space held in reserve, Huston nominated 2022 as the year in which IPv4 would be effectively exhausted.

However, Huston's projections are purely mathematical, combining real data, gathered up until 2003, with assumptions about the future. While Huston is confident about the historical data, he states clearly that we can only make informed guesses about future address usage.

For example, if someone had made a similar forecast in the early 1980s, they would likely have assumed that IPv4 usage would increase, but by how much? And would they have taken into account the future introduction of CIDR, DHCP, and NAT, and changes to allocation practices? Technical and policy developments such as these are difficult, if not impossible to predict. When making his 2003 projections, Huston warned about the uncertainties related to the following potential factors:

- increases in address rate consumption caused by personal mobile IP devices or by an inability of NATs to support emerging popular applications,
- disruptions with a social origin, such as the boom and bust cycles seen in the late 1990s,
- changes to the way in which IP addresses are distributed.

## The view from 2006

### Recent projections

The last three years has by no means brought about consensus on the question of an IPv4 exhaustion date. Research done in this area has continued to produce a variety of projections, spurring lively debate. Factors affecting these various projections include:

- researchers using different projection methodologies to calculate the exhaustion date (for a detailed examination of these methodologies refer to JPNIC's report 'Analysis and Recommendations on the Exhaustion of IPv4 Address Space'),
- researchers focusing on data collected over different periods (for example a 10 year sample compared to a 5 year sample),
- researchers including or discounting social factors and their impacts (such as policy changes),
- varying assumptions about the ultimate destination of the allocated IPv4 addresses (such as whether it will be routed in the public IPv4 Internet).

Cisco's Tony Hain argues that IPv4 exhaustion may happen much sooner than previously thought. One of the factors contributing to this, he suggests, is that there could be a 'last minute rush' as imminent exhaustion of the unallocated address pool becomes more visible. His 'Pragmatic Report on IPv4 Address Space Consumption' focuses on data from the previous five years and projects that the IANA and RIR pools will be exhausted between 2009 and 2016.

Huston has also adjusted his projections, taking into account the acceleration of IPv4 allocation rates in recent years. His figures, based on the frequently-updated 'IPv4 Report' (www.potaroo.net/tools/ipv4) now project that IANA's pool of unallocated address space will be exhausted in 2011 and the RIR's pool of unallocated address space will be exhausted in 2012.

## Policy implications

At the APNIC 22 IPv4 exhaustion panel discussion, held in September 2006, Geoff Huston noted that for regulators and policy-makers "phrasing clear objectives in a regulatory framework with clear and unambiguous signals to industry players will be an extremely difficult challenge".

One potential subject of policy change suggested by some relates to bringing more IPv4 addresses into the public pool. JPNIC's Akinori Maemura, who was also a panellist at APNIC 22, noted that "some think that IPv4 addresses would never be depleted if all class A addresses were returned to registries... But we know that that is not true". In addition to the fact that reclamation would only delay, rather than halt, the exhaustion process, it also appears to be widely accepted that reclamation would be a lengthy, expensive exercise, with limited effectiveness.

▲ Akinori Maemura of JPNIC urged caution on any policy approaches proposed in reaction to IPv4 consumption rates.

So what about the reserved IPv4 space? Hain notes that while some have called for the 16 /8s reserved in the experimental space to be used, there would be some serious hurdles in that many existing network configurations which have been tested would consider such ranges to be a configuration error and refuse to accept them. And the pain would be not be worth suffering because, according to Hain, "assuming the sustained growth trend in allocations continues... the entirety of the old Class E space would amount to about 6 months of run rate".

The RIR exhaustion point could be impacted by changes in policy. However, strategies such as reclaiming unused address space, releasing reserved space, and arbitrarily slowing down allocation rates can only delay the inevitable.

Slowing down consumption (conservation) is only one consideration for policy-makers. The five current goals of IPv4 allocation policy are uniqueness, registration, aggregation, conservation, and fairness. The JPNIC report 'Analysis and Recommendations on the Exhaustion of IPv4 Address Space' projects that a 'demand surge' could take place prior to exhaustion (possibly peaking years before the actual exhaustion point), raising a number of issues related to fairness.

The report argues that under current policies smaller organisations are disadvantaged because they may not be able to use their allocated resources as quickly as larger organisations. If a smaller organisation cannot comply with the 80% utilisation rule, the report argues, they may not be able to secure additional allocations for the future. Assuming that the RIR pools may not run out at exactly the same time, the opportunity for global organisations to go 'RIR-shopping' could also disadvantage smaller organisations. The JPNIC report suggests a range of policy responses, including allowing organisations to get allocations every six months, rather than annually, and reducing minimum allocation block sizes.

## IPv4 exhaustion scenarios

Because IPv4 address space is a finite resource, it can be argued that the crucial question is not when will it be exhausted, but, rather, what will be the response of the Internet community?

During a roundtable discussion at ARIN XVI in October 2005, entitled 'Future of IPv4', KC Claffy from the Cooperative Association for Internet Data Analysis made the point that the IPv4 exhaustion issue is not a data analysis problem, but rather a scenario planning problem. Along similar lines, Hain asked: "when the current RIR IPv4 allocations policies are no longer applicable, what are the most appropriate address management policy measures that will support the continued well-being of the global Internet and its users?"

The answer to this question requires some educated guesses about what is likely to happen during the pre-exhaustion and post-exhaustion periods. These scenarios can then be assessed to decide which would pose the greatest threat to the Internet's long-term future, and which would contribute to its stability.

## IPv4 trading markets

The JPNIC report observes that the concept of an IPv4 trading market would mean major changes for RIRs. The report notes that:

> *At present, Internet registries are prohibited from carrying out any type of address transfers such as selling or purchasing allocated IP address space. However... just after exhaustion, public institutions such as RIRs may need to create a suitable distribution structure, for example, for the securities market. If such a structure is not adequate... a black market may develop.*

It is widely speculated that an IPv4 trading market will appear when demand exceeds supply. Huston, in his APNIC 22 presentation, noted that scarcity is commonly expressed in terms of increased prices. This price rise could place IPv4 addresses out of reach for some industry players, which, in turn, may create additional costs for end-users. These are, he suggests, "probably undesirable outcomes that may not be equitable, efficient or effective".

Huston further explained that while IPv4 trading in the short term may extend the availability of IPv4, the implications for aggregation are not well understood. Trading fragmented IPv4 blocks could create potential issues for the viability of the routing system, meaning an IPv4 trading market would be unlikely to be viable in the long term. "An IPv4 address trading market can provide a short term incentive to expose unused addresses for reuse, and can provide incentives for high address utilisation efficiencies... (but) markets cannot make the finite infinite," he said.

## IPv4 and NATs

Is NAT a possible solution to IPv4 exhaustion? At APNIC 22 Geoff Huston's answer was yes – and no.

He argued that more NATs would be deployed, adding that one interesting unknown is what this would cost, because presently ISPs do not tend to pay for NATs; their customers do. NATs,



P 6

he said, were extensively used, and popular applications were NAT-agile. Therefore, he claimed, NATs were a reality in the short-term. He acknowledged, however, that many people in the industry are anti-NAT because of a lack of standardisation. This could be fixed, but the real long-term problem, he argues, is not standards, but complexity:

> *...you've got to make application-specific identity domains. You've got to do a whole new set of application-level technology. You've got to start making the NATs aware of the application, that is, traverse them. I think then about trying to deploy a new application if every NAT in the world needs a software application...it's...difficult to run reliably, let alone securely. Once you start threading your NATs...your brain will explode, as will the network. As far as I can see, the long-term future of NATs is relatively depressing.*

In his 'Pragmatic Report on IPv4 Address Space Consumption' virtual roundtable Hain put it equally bluntly: "NAT and CIDR did their jobs and bought the 10 years needed to get IPv6 standards and products developed. Now is the time to recognise the end to sustainable growth of the Internet has arrived and that it is time to move on."

## IPv6

The IPv6 discussion has been on the agenda for quite some time and public IPv6 addresses have been available since 1999. However, switching over to IPv6 is not a simple matter and organisations have responded to the issue in a variety of ways. Some economies, like Japan, have invested significant time and resources implementing IPv6-ready networks, whereas others have done comparatively little. Why is this the case?

In many quarters it has been argued that if it is left up to players in the marketplace, planning for the long-term future is not something the IT industry is good at. In most cases large-scale IPv6 network development has occurred where institutions like governments have taken an interest (and provided funding).



At ARIN XVI, Claffy noted that IPv6 development (in the USA, at least) was blocked on two fronts: capital and incentive. She argued that the focus on profits in a competitive market runs counter to a culture of innovation, and "the ones who need to innovate in the core don't have the capital".

Huston makes a similar point: "I'm not sure many industry players are spending money for problems that are going to happen the year after next. Competitive markets are tight. Discipline of competition leads to short-term focus. There is no long-term money. Almost all major telcos in the world have shut down their research labs. They are no longer investing in a future."

He noted that from a service provider perspective, IPv4 exhaustion:

> *...can be re-expressed as a problem relating to investment lifecycles. The ISP industry and the enterprise sector have already made considerable investments in IPv4-based infrastructure in equipment...and operational capability, and we are seeing some considerable reluctance to add to this with additional investment into IPv6 capability at this time. The direction of the use of various forms of NAT-based approaches and increasing use of application layer gateways...can be seen as an effort to extend the lifetime of the existing infrastructure investment. In a volume-based market with relatively low margins, this position certainly has some sound rationale from a business management perspective. But I agree with Tony (Hain)...that such business approaches are ultimately short-term in nature...The numbers all indicate that this is not a matter which can be deferred indefinitely. (The) call for some timely attention to the need to commence investment in IPv6-based service infrastructure is one that I hope the industry is listening to attentively.*

## Technical challenges – IPv4 and IPv6 co-existence

In the 'Pragmatic Report on IPv4 Address Space Consumption' virtual roundtable, Hain argued that IPv6 is technically ready for implementation, but that a more substantial hurdle is attitude. "When CIOs make firm decisions to deploy IPv6, the process is fairly straightforward. Staff will need to be trained, management tools will need to be enhanced, routers and operating systems will need to be updated, and IPv6-enabled versions applications will need to be deployed. All these steps will take time – in many cases multiple years," he said.

JPNIC's 'Analysis and Recommendations on the Exhaustion of IPv4 Address Space' includes an overview of the technical challenges created by the transition to an IPv6 Internet, noting that while IPv4 and IPv6 are able to run across the same network and share physical facilities the two protocols are "not compatible and cannot communicate directly with each other. Therefore, the Internet constructed with IPv4 and the Internet constructed with IPv6 should be considered as two independent networks, though they share some of the same facilities".

Dual-stacked networks that use both IPv4 and IPv6 are operating well across the Internet now. However, the JPNIC report details the challenges to be faced by network operators upon full IPv4 exhaustion, when it will be necessary to ensure that new single-stacked IPv6 networks can communicate with existing single-stacked IPv4 networks.

Solutions to these challenges will be critical to the long term viability of the IPv6 protocol.

## Conclusion

At APNIC 22, Geoff Huston commented on the current perception of the IPv4 exhaustion issue, noting that "the major key to looking at the future is actually understanding the past... What we're looking at is no different from many other crises... There will be disruption. It won't be seamless and it won't be costless".

Paul Wilson, APNIC Director-General agrees that this issue will certainly attract more attention in future, both inside and outside the RIR communities. "Through the policy processes of APNIC and the other RIRs, we can be sure that developments take full account of the operational feasibility and implications of any changes," says Wilson. "I have no doubt that collectively we will find the most realistic and workable solutions, to ensure ongoing stable growth of the Internet worldwide".

## Further reading

A full archive of the APNIC panel discussion, 'IPv4 exhaustion: what's the real story?', is available at:

www.apnic.net/meetings/22/program/panel.html

The ARIN roundtable, 'The Future of IPv4', is archived at:

www.arin.net/meetings/minutes/ARIN_XVI/ppm.html

Tony Hain's article, 'A pragmatic report on IPv4 address space consumption' was published in the Cisco Internet Protocol Journal at:

www.cisco.com/web/about/ac123/ac147/archived_issues/ipj_8-3/ipv4.html

JPNIC's report 'Analysis and Recommendations on the Exhaustion of IPv4 Address Space' is available here:

www.nic.ad.jp/en/research/IPv4exhaustion_trans-pub.pdf

# Seeking solutions for IPv6 multihoming

As IPv6 deployment continues, albeit slower than many would hope, a range of technical issues associated with its use are being discussed by organisations and individuals around the world. One such discussion addresses best practice for IPv6 multihoming facilitation. A recent paper compiled by Marla Azinger, Senior IP Engineer for Frontier Communications and member of the ARIN AC, is being used as one step toward finding a universal solution to this routing issue.

The paper, entitled 'IPv6 Multi-homing Solutions and their Pros and Cons', sets out a range of solutions proposed by the Internet community, including possible reasons for and against each strategy, and any questions that have been identified. The paper is a living document, as the author will update it as other pros and cons are identified, questions are answered, or further solutions are proposed. As Azinger notes in introducing the paper, the "ultimate solution may or may not be as written in this document."

This article provides a brief overview of some of the solutions discussed by Azinger. For a more in-depth examination of issue please see the original document, which can be found on the NRO website.

## Possible solutions

Azinger's document details a number of solutions, under the following categories:

- CIDR
- Metro/regional assignment of IP address space
- Community codes
- A published list of address blocks approved for multihoming
- Policy
- Maximum prefix
- Shim6
- 8+8/GSE
- PI only

Pros and cons are identified, and factors considered for each solution include:

- Economic impact
- Technical impact
- Policy impact (will the RIRs have to develop new policies?)
- Additional knowledge requirements
- Impact on routing table bloat
- Facility for traffic engineering

## CIDR

A solution requiring filters to be opened to a selected CIDR boundary would have the benefit of mirroring the current IPv4 multihoming method. This means a minimum of new knowledge is required, fast implementation is possible, and there is no drastic economic impact.

This solution presents a range of questions and potential problems, however. For example, which CIDR boundary should be adopted? Azinger's paper examines the pros and cons of adopting a CIDR boundary of /48, a CIDR boundary more specific than /48, and allowing a specific number of aggregations per AS.

## Metro/regional

This solution would involve assigning IP addresses to city or region representatives, such as a city council or regional authority, as opposed to large networks, ISPs, or end users. As described by Azinger:

> *The city then chooses a single or list of relevant providers to serve as the interchange. Each of the providers will advertise the region's prefix to the Internet. Based on a protocol or by a contract, these providers will accept more specific prefixes from subscribers that are within the regional/geographic location, and will then interchange traffic to the other relevant providers appropriately.*

Such a solution would mark a significant departure from the way that routing in the Internet has been done to date, and as such its economic and policy impact would be quite large. However, proponents of this view believe it could help control routing table bloat and "provide a multi-home solution for everyone".

## Community codes

Tagging multihoming prefixes with a BGP community attribute is another suggestion. This would mean a new community attribute, such as MULTIHOME, would be created, and that prefixes tagged with this attribute would be propagated by default.

This solution is economically efficient, and addresses both routing table bloat and the need for traffic engineering. However, it relies on all parties correctly tagging and propagating the new attribute, and is therefore highly susceptible to human error. It may also present an easy way to hijack address space.

## Published list



Another option is that the RIRs maintain and publish a list of approved multihoming blocks. This solution is economically lightweight and would solve numerous security concerns, such as address hijacking, and spam and virus propagation. However, it would involve a considerable amount of policy discussion and additional responsibilities for the RIRs. It is also unclear whether the requirement to open filters to the published list could be a mandate, or only a strong suggestion; this would have to be decided by the community, and would have significant impact on the effectiveness of the solution.

## Policy

This solution would involve a policy written to work around the PA multihoming problem, by allowing providers to request PI space for the purpose of multihoming. This would not allow for traffic engineering, but it would allow providers to offer multihoming to customers.

## Maximum prefix

This solution would limit each origin AS to a maximum number of prefixes. This would provide for both multihoming and traffic engineering, but its impact on routing table bloat is unclear. Should routing table bloat prove a problem, finding further solutions for this may prove expensive.

## Shim6

Shim6 is a protocol that could provide a solution for IPv6 multihoming, but it requires both provider ends to be running the protocol. Therefore it would require a lot of education and training, and may prove economically demanding.

P 8

### 8+8/GSE

Both 8+8 and GSE are also protocol-based solutions that function through identification manipulation. They would also require a large amount of education and training, but would address issues such as routing table bloat, and would allow both PI and PA space to be multihomed.

### PI only

PI only multihoming would only allow multihoming for a select few, and is therefore not regarded as an ideal solution.

### The next step

The document outlining all of these solutions in more detail, with pros and cons for each, is now open for development by any interested parties. Any suggestions, clarifications or new ideas for solutions can be emailed to marla.azinger@frontiercorp.com, and will be included in further updates to the document. In this way the community can move toward a final consensus decision.

'IPv6 Multi-homing Solutions and their Pros and Cons' version 1.0 is available at:

www.nro.net/documents/pdf/MultihomeIPv6procon.pdf

An archive of the discussions on this subject at the APNIC 22 meeting in Taiwan is available at:

www.apnic.net/meetings/22/program/sigs/ipv6.html

# ICANN ratifies new global policy for allocating IPv6 addresses to RIRs

At its September 7, 2006 Meeting the Board of ICANN ratified a global policy for the allocation of IPv6 addresses by the Internet Assigned Numbers Authority (IANA) to Regional Internet Registries.

This new global policy describes how IANA distributes IPv6 address ranges to the RIRs.

Under the policy, the unit of IPv6 allocation is now a /12, with IANA allocating sufficient address space for each RIR to support their respective registration needs for at least 18 months.

An RIR will be eligible to receive additional IPv6 address space from the IANA when either its available space is less than 50 percent of a /12 or its available space is less than its established necessary space for the following 9 months.

"This is an outcome which provides certainty to Internet Registries and their customers who include Internet Service Providers and users" said Dr Paul Twomey, President and CEO of ICANN.

The policy was developed in a bottom-up process, coordinated through the various RIR communities, before being submitted to the ASO AC by the Executive Council of the Number Resource Organization, Raúl Echeberría, in June 2006.

The Global Addressing Policy document is available from the ASO website at:

www.aso.icann.org/docs/aso-global-ipv6.pdf



# ISOC announces IETF fellowship programme

The Internet Society (ISOC) has launched a new programme aimed at helping technologists from developing countries become more involved in the IETF by attending a meeting in person. The program, called the 'ISOC Fellowship to the IETF', sponsors the cost of attending (meeting fees, airfare, hotel accommodation, and a small stipend) for up to five individuals per meeting.

The programme is aimed at individuals from developing countries that possess a solid level of technical education and enough knowledge about concrete areas of IETF work to follow and benefit from the meeting's technical discussions (and ideally they will have already been participating in one or more IETF mailing lists.)

Fellowships will be awarded through a competitive application process. ISOC currently is accepting fellowship applications for two IETF meetings: IETF 68 being held in Prague, Czech Republic on 18 - 23 March and IETF 69 being held in Chicago, USA on 22 - 27 July.

The Internet Society (ISOC) was founded in 1992 as an international, non-profit, membership organization whose mission is to assure the open development, evolution, and use of the Internet for the benefit of all people throughout the world.

ISOC is the organizational home for the groups that develop Internet infrastructure standards, including the IETF and the Internet Architecture Board (IAB). It provides funding and operational services to support the work of these bodies.

Full details of the fellowship programme, including how to apply, is located at

www.isoc.org/educpillar/fellowship

Interested parties are also encouraged to contact Karen Rose <rose@isoc.org> or Mirjam Kuehne <mir@isoc.org> for more information.

Applications are due on 2 February 2007.

# 4-byte Autonomous System Numbers: The view from the old BGP world

◀ 4-byte Autonomous System numbers are now being distributed by the 5 RIRs under a transition policy. **Geoff Huston** explains how 4-byte ASNs will coexist with the existing 2-byte ASNs and what ISPs will observe during this transition.

The IANA has now expanded the AS number registry from its original 16 bit range (AS numbers 0 through 65535) to a 32 bit range (AS numbers 0 through 4,294,967,295).

This is a timely action, in that we were running out of AS numbers in the 2-byte number range, and the rate of AS number consumption was such that by October 2010 the AS number range would have been completely exhausted. APNIC has adopted a policy that allows ISPs to transition to this extended number range without the need for last minute rushed changes in BGP.

From 1 January 2007 until 31 December 2008, ISPs may specifically request an AS number from the extended 4-byte number pool, but, by default, they will be assigned an AS Number from the original 2-byte number pool. From 1 January 2009, the allocation practice will be reversed, and unless specifically requested, AS numbers will be allocated from the extended 4-byte number pool.

What are the implications for ISPs with this AS Number allocation policy?

If an ISP wants to use an AS number that is greater than 65535 then it will need to deploy "new" BGP. That is, it will need to deploy a version of the BGP protocol in its routers that understands 4-byte AS numbers.

But what about everyone else? What about the existing "old" BGP world that uses 2-byte AS numbers? Even though they have a 2-byte AS number, will they need to upgrade their BGP to see these new extended AS numbers?

The approach in the 4-byte AS number transition has been carefully constructed to be backward compatible. The reassuring news is that if you have a 2-byte AS number and are running BGP then you need to change nothing at all. The Internet will still work and you will continue to see routes to all advertised networks, irrespective of the existence of 4-byte AS numbers in the network. You don't need to upgrade your version of BGP, nor make any router configuration changes in your network. Nothing need change.

Well, almost nothing! Some things might change, and in this article I'd like to highlight some of the things to think about if you are running an old BGP that supports only 2-byte AS numbers.

First, some background. In BGP, the AS Path attribute is used for two essential roles. It's a metric of path length where, by default, BGP will prefer a short AS path over a longer one for the same advertised prefix. It's also a loop detector, where each AS is capable of detecting a potential routing loop by seeing its own AS already in the AS path of received BGP advertisements. Strictly speaking, the AS path does not have to be entirely accurate, but it does need to have these capabilities of path metric and loop detection.

The transition mechanism of 4-byte BGP is a combination of translation and tunnelling.

When passing a routing update into the 2-byte old BGP world the 4-byte new BGP speaker converts all AS numbers in the AS path to 2-byte values. If the AS number was between 0 and 65535 then all it does is strip off the leading 16 zero bits of the AS number value to perform this conversion. If the AS number is greater than 65535 then it translates the AS number of the special 2-byte value of 23456. If any AS number is translated in this way the new BGP speaker also saves a copy of the 4-byte AS path in a new transitive opaque community attribute called "NEW_AS_PATH".

When passing a routing update from a 2-byte old BGP speaker to a 4-byte new BGP speaker, all the AS numbers in the AS Path attribute are expanded to the equivalent 4-byte values by adding the leading 16 zero bits to the AS number value. If there is a NEW_AS_PATH community attribute, then this AS string is substituted back into the AS PATH. If all goes well, the 4-byte BGP world sees an accurately re-constructed 4-byte AS PATH, preserving both AS path length metrics and the BGP loop detection capability.

But even if the NEW_AS_PATH attribute is not present, or cannot be substituted back into the 4-byte AS Path, it is not a fatal condition. Even without the substitution, the AS Path length metric is preserved, and loop detection still can be performed, although in a degraded fashion. Potential routing loops entirely within the 4-byte new BGP world are detected as normal, as are potential routing loops entirely within the 2-byte old BGP world. And in the case of a mixed 2-byte and 4-byte potential routing loop, the detection will still happen when the loop formation reaches the 2-byte old BGP world. So if the NEW_AS_PATH attribute is lost in the 2-byte old BGP world, then the only casualty is speed of routing convergence, where it may take a number of additional AS hops for a potential routing loop to be detected and removed.



The implications for old world BGP appear to include the following considerations:

- The first implication for the old BGP world is that it is preferred if the NEW_AS_PATH is carried as a transitive opaque community attribute when present. That's a "SHOULD", not a "MUST", by the way.
- The second implication is that the old 2-byte BGP world will see more and more instances of AS23456 as both an originator and as a transit provider. This is not a mistake, it is just the only way that the 2-byte world can carry a place-holder for a 4-byte AS value.
- The third implication is that an old BGP ISP may see routing peers, both as customers, peers, and possibly upstreams, using 4-byte AS numbers. But as your local BGP is an old world BGP, your routers will not be aware of these 4-byte AS values. From your routers' perspective, AS23456 is going to start popping up both as a diverse prefix originator and a ubiquitous transit provider. The ISP's operating support system (OSS) probably should be able to store the corresponding AS numbers of these routing peers as 4-byte number values, simply to avoid unnecessary confusion and potential ambiguity! But if you use the OSS to generate router configuration fragments, AS path filters and similar, then you may need to revise your OSS to transform the 4-byte AS number values into the 2-byte equivalent value of AS23456. The same situation occurs when using a Routing Registry to generate local configuration state for your 2-byte BGP. So, for example, your OSS may have a configuration relating to a BGP peer with AS1.2, but your old BGP router will need to be provided with a generated configuration fragment that refers to AS23456.

▶▶ P 10

- Many ISPs use directed community attributes to signal to a remote AS. A prefix that has explicit signalling to AS65505 may use a community attribute of "65505:123", for example. But this will not work as cleanly if the old BGP wishes to generate a signal to a 4-byte target AS. At the very least, your BGP version should support expanded community attributes (RFC4630) and also support the means of entering 4-byte AS numbers into these attributes (draft-rekhter-as4octet-ext-community-01.txt).
- You should also expect a modest increase on memory and bandwidth requirements for BGP. While nothing much is changing in your view of the routing world, you will be carrying these NEW_AS_PATH transitive community attributes along with the prefixes, and the memory and bandwidth required to hold AS Paths will triple for old world BGP routers. That's not saying that total BGP memory demands will triple, just that requirement relating to AS path storage.
- We might anticipate slightly poorer performance in routing. The specific cases where convergence times will extend are in those circumstances where the NEW_AS_PATH attribute is lost on transit through the old BGP 2-byte world. In such cases loop detection will take slightly longer, and this will have some level of impact on convergence times.
- There is no dynamic capability to support a change from 2-byte old BGP to 4-byte new BGP. When a routing domain wants to transition from a 2-byte to a 4-byte AS number, then the BGP session will need to be reset via a complete shutdown and restart. The transition from old BGP to new BGP within a domain includes a number of considerations with respect to iBGP as well as eBGP sessions, and the transition will need to be planned very carefully.



# APNIC launches 4-byte ASN service

In November 2006, the IANA issued APNIC with its first 4-byte assignment range: 2.0 - 2.1023

On 1 January 2007, APNIC began processing applications for 4-byte AS numbers. APNIC will continue to assign 2-byte AS numbers by default, unless otherwise requested.

AS number syntax in the APNIC Whois Database has been extended to accommodate 4-byte AS numbers. This may affect any automated tools you use to interact with the database. The syntax for 2-byte AS numbers will remain unchanged.

The 4-byte syntax changes affect the following database objects:

- aut-num
- as-set
- inet-rtr
- route
- route-set
- as-block
- filter-set
- peering-set
- route6
- rtr-set

For more information, visit:

www.apnic.net/services/asn_guide.html

# Call for nominations: APNIC Executive Council

Four positions on the APNIC Executive Council will be opened for election at APNIC Member Meeting in Bali, Indonesia, on Friday 2 March 2007.

The members whose positions are up for re-election are Kuo-Wei Wu, Moo-Ho Billy Cheon, Qian Hualin, and Ma Yan, who were elected at the APNIC Member Meeting in Kyoto on 25 February 2005.

Nominations are due by close of business Friday 16 February 2007 at the latest.

Only APNIC members may nominate and vote for candidates, but nominees do not have to be representatives of members.

Online voting will be available for this election, via the MyAPNIC web site. Online voting will open soon after the nomination period closes.

If you are not able to vote online or attend the AMM, you can still make your vote count by appointing a proxy to represent you at the meeting.

Nominees should note that positions on the APNIC Executive Council are voluntary and that APNIC may not be able to reimburse EC members all expenses associated with EC duties. Where possible, however, APNIC will reimburse actual expenses for attendance to APNIC meetings, providing that these fall within budget and cash flow constraints.

Nominations may be made using the online nomination form available at:

www.apnic.net/meetings/23/ec/nomination.html

For more information on nominations received, proxy registration, current EC members, and the role of the EC, please refer to:

www.apnic.net/meetings/23/ec

# Call for nominations: ICANN Board

The Address Supporting Organization Address Council (ASO AC) has issued a call for nominations to the ICANN Board to fill the ASO seat currently held by Raimundo Beca, whose term expires in June of 2007.

All candidates interested in selection to the ICANN board by the Address Council must meet the selection criteria and conflict of interest requirements as stated by ICANN in its Bylaws and other relevant documents. Candidates will also be required to sign a Letter of Certification attesting to compliance with the ICANN stated requirements as well as attesting to previous conduct and character.

Candidates may be required to travel one time during the month of March or April 2007 if the Address Council decides to conduct in-person interviews.

This nomination period will close on 21 February 2007.

For detailed information, please visit:

www.aso.icann.org

# Staff updates

## Finance

### Maggie Liu, Accounts Officer

Maggie joined APNIC in November 2006. She has a Master of Professional Accounting. Maggie has worked in a variety of accounts and co-ordination roles in the Australian and Chinese travel and tourism industries. She is fluent in Mandarin. Her responsibilities at APNIC include general accounts keeping, billing related enquiries, and other administrative tasks within the Finance unit.

## Marketing & External Relations

### James Lemon, Web Developer

James joined APNIC in December 2006. He has qualifications in Computer Science and Graphic Design. Originally from London, he has previously worked for the Law Society of England and Wales, updating content and developing mini-apps. He has also had similar positions within the Queensland government. His work involved developing consistent cross-browser/platform web templates and web publishing standards. He was also responsible for assisting in the roll-out of a content management system and providing user training.

## Member Services

### Nurul Islam (Roman), Internet Resource Analyst

Roman joined APNIC as an Internet Resource Analyst, or Hostmaster, in September 2006. He has a Bachelor of Computer Science and an MSC in Digital Communication Networks. Roman has worked for network communications companies in the UK and Bangladesh. Roman specialises in developing Internet infrastructure for service provider networks, particularly routing and switching. He is fluent in Bengali and Hindi. His responsibilites at APNIC include processing requests for IP address space and AS number allocations within the Asia Pacific region.

# New *Apster* subscription policy

You may recall that the previous issue of *Apster* included an article about the Eco-APNIC initiative. This staff-driven project aims to reduce our impact on the environment by doing many things, including reducing our paper usage.

So, from now on, *Apster* will now be mainly published in an electronic format. If you'd like to be notified when the next issue is published, you can subscribe to the *Apster* mailing list here:

mailman.apnic.net/mailman/listinfo/apster-subscribers

If you prefer to use RSS, then you can subscribe to this feed:

www.apnic.net/docs/apster/news.rss

We will still be printing some copies of *Apster*, which we will give away at meetings and other events. If you would like to keep receiving hard copies, then please subscribe using this form:

www.apnic.net/docs/apster

eco APNIC

## Training schedule

### 2007

**January**

| | |
|---|---|
| 19 | Colombo, Sri Lanka |
| 19 | Colombo, Sri Lanka |
| 21-22 | Colombo, Sri Lanka (with SANOG 9) |
| 31 | Palmerston North, New Zealand (with NZNOG) |

**February**

| | |
|---|---|
| 12-16 | Dhaka, Bangladesh |
| 26-27 | Bali, Indonesia (with APRICOT/APNIC23) |

**March**

| | |
|---|---|
| 4-6 | Australia (tba) |
| 16 | Kuala Lumpur, Malaysia (with IPv6 Summit) |
| 18 | Kuala Lumpur, Malaysia |

**April**

| | |
|---|---|
| 10-13 | Cebu, Philippines |
| 16-20 | Singapore |
| 23-25 | Kathmandu, Nepal |

**May**

| | |
|---|---|
| 1-5 | Pakistan (tba) |
| 7 | India (tba) |
| 9 | Bangalore, India (with IPv6 Summit) |
| 7-11 | Vietnam, Laos, Cambodia (tba) |
| 28-2 June | China (tba) |

**June**

| | |
|---|---|
| 4-7 | Bangkok, Thailand |
| 11-14 | Ulaanbaatar, Mongolia |
| 25-28 | Vanuatu (with PACNOG) |

The APNIC training schedule is subject to change. Please check the web site for regular updates at:

www.apnic.net/training

If your organisation is interested in sponsoring APNIC training sessions, please contact us at:

training@apnic.net

# Calendar

**Pacific Inslands Telecommuications Association Meeting**

*13 January 2007*
Honolulu, USA
http://www.pita.org.fj

**PTC '07 Conference and Exhibition**

*14-17 January 2007*
Honolulu, USA
http://www.ptc07.org

**SANOG 9**

*14-24 January 2007*
Colombo, Sri Lanka
http://www.sanog.org

**23rd APAN Meeting**

*22-26 January 2007*
*Manila*, Philippines
http://apan.net/meetings/manila2007

**NANOG '07**

*31 January - 2 February 2007*
Palmerston North, New Zealand
http://www.nznog.org

**Asia IPv6 Summit**

*20-21 February 2007*
Makati, Philippines
http://www.asiaipv6.com

**APRICOT 2007/APNIC 23**

*21 February - 2 March 2007*
Bali, Indonesia
http://www.apnic.net/meetings/23

**68th IETF**

*18-23 March 2007*
Prague, Czech Republic
http://www.ietf.org/meetings/meetings.html

**ICANN Meeting**

*26-30 March 2007*
Europe (venue TBA)
http://www.icann.org/meetings

**ARIN XIX**

*22-25 April 2007*
San Juan, Puerto Rico
http://arin.net/meetings

**AfriNIC 6**

*23 April - 4 May 2007*
Abuja, Nigeria
http://afrinic.net/meeting

**CeBIT Australia 2007**

*1-3 May 2007*
Sydney, Australia
http://www.cebit.com.au/main/about

**RIPE 54**

*7-11 May 2007*
Tallinn, Estonia
http://www.ripe.net/meetings/current.html

**LACNIC X**

*21-25 May 2007*
Isla Margarita, Venezuela
http://www.lacnic.net/en/eventos

**Interop Tokyo 2007**

*13-15 June 2007*
Tokyo, Japan
http://www.interop.jp

**ICANN Meeting**

*25-29 June 2007*
North America (venue TBA)
http://www.icann.org/meetings

## How to contact APNIC

| | |
|---|---|
| Street address | Level 1, 33 Park Road, Milton, Brisbane, QLD 4064, Australia |
| Postal address | PO Box 2131, Milton QLD 4064, Australia |
| Phone | +61-7-3858-3100 |
| SIP | info@voip.apnic.net |
| Fax | +61-7-3858-3199 |
| Web site | www.apnic.net |
| General enquiries | info@apnic.net |
| Hostmaster (filtered) | hostmaster@apnic.net |
| Helpdesk | helpdesk@apnic.net |
| Training | training@apnic.net |
| Webmaster | webmaster@apnic.net |
| *Apster* | apster@apnic.net |

## Member Services Helpdesk

The Member Services Helpdesk provides APNIC members and clients with direct access to APNIC Hostmasters.


www.apnic.net/helpdesk


helpdesk@voip.apnic.net


helpdesk@apnic.net


+61 7 3858 3188

**Helpdesk Hours:** 9:00 am to 7:00 pm (UTC + 10 hours) Monday - Friday

## Communicate with APNIC via MyAPNIC

APNIC members can use MyAPNIC to:

- view APNIC resources held by their organisation
- monitor the amount of address space assigned to customers
- view current and past membership payments
- view current tickets open in the APNIC email ticketing system
- view staff attendance at APNIC training and meetings
- vote online

For more information on MyAPNIC's features, see:

www.apnic.net/services/myapnic

eco APNIC
This issue of *Apster* is printed on ONYX recycled paper.