APNIC

Apster

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

Issue 23 - September 2007

The Internet in India

APNIC 24 will be held in conjunction with SANOG 10 from 27 August to 7 September 2007 in New Delhi, India, and will be hosted by the ISP Association of India (ISPAI). This article traces the history of the Internet in India, the current situation, and anticipated future developments.

With over one billion people, India is the second most populous country in the world. It is largely rural, with 70% of its population living outside the major metropolitan areas. India's geography is vast and diverse; its landmass covers 3.3 million square kilometres and features mountain ranges, valleys, desert regions, tropical rain forests, fertile plains and a dry plateau.

In addition to these demographic and geological factors, poverty and limited telecommunications infrastructure have posed a challenge for Internet development in India.

Early days

In 1986, the Indian government established a multi-protocol network called the Education and Research Network (ERNET). This project received technical and financial support from the UN Development Program. It aimed to set up a nationwide computer network for academic and research communities, conduct research and develop computer networking, and provide network training and consulting services.

Other developments in the late 1980s included:

- INDONET was India's first SNA network. By the end of the following decade, after several upgrades, INDONET connected eight Indian cities via 64 Kbps leased lines.
- NICNET was India's first nationwide VSAT network. This network provided data communications for government agencies. NICNET's services included email, remote database access, data broadcasting, electronic data interchange, and an emergency communication system.
- The National Centre for Software Technology (NCST) was the first Indian institution to establish an international connection to the global Internet. The 9.6 Kbps UUCP link between NCST and UUNet Technologies in the USA was launched in February 1989.

ISPs

VSNL, a company owned by the Indian government, operated the global packet switched service (GPSS) at speeds of up to 64 Kbps. The GPSS network consisted of three packet switching exchanges in Calcutta, Mumbai, and New Delhi. It provided connections to most foreign packet switched networks and served as a gateway to the Internet, allowing text communications. VSNL also operated the international Gateway Electronic Mail Service (GEMS 400), with nodes in Bangalore, Calcutta, Chennai, Mumbai, New Delhi, and Pune.

In 1998, VSNL signed a national data carrier agreement to market services developed by Global One. Managed network nodes were established in Bangalore, Mumbai, and New Delhi to provide high speed connections to Global One's 1400 points of presence around the world. While ISPs had not yet been officially licensed, VSNL is generally recognised as India's first ISP.

The regulatory and legislative framework that Indian ISPs operate within is unique and complex. The ISP Association of India (ISPAI) was established in 1998 to represent the interests of ISPs in India. The organisation also works closely with Indian government and industry associations in an advisory capacity.

P 2

24th APNIC Open Policy Meeting

29 August - 7 September 2007 New Delhi - INDIA

P 1

Recent developments and growth

In 2003, ISPAI established the National Internet Exchange of India (NIXI) with assistance from India's Department of Technology. This neutral Internet exchange enables ISPs to route Internet traffic locally in order to improve Internet performance and reduce costs.

APNIC has established MoUs with ISPAI and NIXI. In 2005, this partnership bore fruit when three root nameservers were installed in Chennai (F), Mumbai (I) and Delhi (K).

Around the same time, in 2004, the Indian Computer Emergency Response Team (CERT-In) was established. CERT-In responds to computer security incidents and assists the Indian IT community to implement proactive measures to reduce the risk of these incidents.

In 2005, it was estimated that around 60% of Internet users were accessing the Internet via a network of 200,000 cybercafes.

According to Peter Wolcott and Seymour Goodman in their Association for Information Systems report, one of the most common measurements of Internet penetration is known as pervasiveness: The fraction of the total population that uses the Internet regularly. In 2002, usage data rated India's Internet pervasiveness as 'established' and increased it to 'common' by the end of that year. Unlike the case in many countries where the Internet takes hold first within a single metropolitan area, India established points of presence in multiple states from the outset. Development was largely confined to the states where it started; therefore, in the context of the entire population, Internet usage remained low. However, growth rates in recent years have been impressive.

Internet subscriber growth in India

Year	Dialup subscribers (in lakhs)	Growth on previous year	Broadband subscribers (in lakhs)	Growth on previous year
Mar 1998	1.4			
Mar 1999	2.8	100%		
Mar 2000	9	221%		
Mar 2001	30	233%		
Mar 2002	32	7%		
Mar 2003	36	13%	0.08	
Mar 2004	45	25%	0.19	138%
Mar 2005	56.5	26%	1.8	847%
Mar 2006	69.4	23%	13.5	650%
Sep 2006	88	27%	18.2	35%
Dec 2006			21	15%

Source: ISPAI
Note: 1 lakh = 100,000

▲ The Taj Mahal is India's most famous tourist attraction. Located less than 200km from New Delhi, it is one of many impressive historical sites located in the area.

Looking forward

The Indian government, like that of many other countries, has taken a keen interest in expanding broadband services and preparing for IPv6. A Telecom Regulatory Authority of India (TRAI) consultation paper, published in 2005, outlines the issues related to India's transition from IPv4 to IPv6. Their final recommendations were released in January 2006.

Wolcott and Goodman note that in recent years private sector initiatives have greatly expanded the Internet infrastructure and services market. At the same time, government initiatives have promoted the expansion of the Internet in areas that are not served well by private ISPs. They note that "...time is an ally; the basic elements for continued growth of the Internet are largely in place."

Sources and further information:

Communications of the Association for Information Systems:

Global Diffusion of the Internet I - Peter Wolcott and Seymour Goodman, 2003

NIXI:	www.nixi.in
ISPAI:	www.ispai.in
CERT-In:	www.cert-in.org.in

Consultation Paper on Issues Relating to Transition of IPv4 to IPv6 in India – 26 Aug 2005:

http://www.trai.gov.in

India IPv6 forum: http://ipv6forum.in

APNIC 24 Open Policy Meeting

APNIC's 24th Open Policy Meeting is being held in conjunction with SANOG 10 in India's picturesque capital New Delhi. The event will run from 29 August to 7 September 2007 at the Intercontinental The Grand Hotel.

APNIC 24 is hosted by the ISP Association of India (ISPAI). It is the largest ever dual event of this nature to be held in South Asia.

Attendees have the unique opportunity to meet with Internet and networking experts from all over the world, with delegates from America, Europe, South America and the Asia Pacific present to share their wealth of expertise and experience.

Opportunities for key personnel from Internet organisations in developing economies to attend APNIC 24 were offered through the APNIC 24/SANOG 10 combined fellowship program.

Attending the Open Policy meeting is a great opportunity to participate in activities and processes that are critical for managing Internet resources both in the Asia Pacific region and globally. Decisions made at this meeting may have a direct effect on your organisation.

You can follow the events at the meeting as they happen, onsite or remotely, using our live features:

Live transcripts

Most meeting sessions are transcribed and broadcast via a chat client so that you can read along. Also, full transcripts of the proceedings will be posted to the website within 24 hours.

Chat rooms

You can actively participate in the meetings using the Jabber chat protocol, which allows you to:

- Read the meeting transcripts as they are typed
- Discuss issues with other participants
- Have your comments read out live at the meeting
- Voice your opinion on policy proposals as part of the consensus process

Video and audio streaming

Live video streaming will be available for selected sessions.

For more information please see:

http://www.apnic.net/meetings/remote

◀ Attendees at a recent APNIC Open Policy Meeting

index

Policy proposals under discussion in the APNIC community

prop-051 Global policy for the allocation of the remaining IPv4 address space

This proposal suggests that when an agreed minimum amount of available space remains, an identical number of IPv4 allocation units (/8s) will be allocated by IANA to each RIR.

Status in other RIR regions:

AfriNIC Posted to the Resource Policy Discussion List on 9 July 2007.

ARIN Awaiting review by the ARIN AC for possible inclusion as a formal policy proposal at ARIN XX.

LACNIC Reached consensus at LACNIC X and is awaiting Board approval.

RIPE NCC Open for discussion on the Address Policy WG mailing list until 27 August 2007.

prop-050 IPv4 resource transfer

This is a proposal to remove APNIC policy restrictions on the transfer of the registration of portable IPv4 address allocations and assignments between current APNIC account holders.

Status in other RIR regions:

This proposal has not been submitted in any other region.

prop-049 IANA policy for the allocation of ASN blocks to Regional Internet Registries

This proposes to have a global policy for the RIRs to receive blocks of Autonomous System Numbers (ASNs) from IANA.

Status in other RIR regions:

AfriNIC Not yet submitted.

ARIN Awaiting review by the ARIN AC for possible inclusion as a formal policy proposal at ARIN XX.

LACNIC Reached consensus at LACNIC X and is awaiting Board approval.

RIPE NCC Last call for comments ended 14 August 2007.

prop-048 Pv6 ULA-central

This proposes the assigning of IPv6 blocks within the 'Centrally Assigned Unique Local IPv6 Unicast Addresses' to organisations or individuals requiring it.

This proposal will be discussed at the APNIC 24 Policy SIG.

Status in other RIR regions:

AfriNIC Posted to the Resource Policy Discussion List on 1 April 2007.

ARIN Not submitted as a proposal.

LACNIC Did not reach consensus at LACNIC X.

RIPE NCC Last call for comments on the Address Policy WG mailing list ended 13 August 2007.

prop-047 eGLOP multicast address assignments

This is a proposal for RIRs to begin assigning multicast addresses from the range specified in RFC 3138.

This proposal was submitted after the deadline for policy proposals to be discussed at APNIC 23. Therefore, this proposal was presented as an informational proposal at APNIC 23, and the decision to adopt, modify or abandon the policy proposal deferred until a later meeting.

Status in other RIR regions:

AfriNIC Not submitted as a proposal.

ARIN Closed. The ARIN AC did not accept it as a formal proposal for ARIN XIX.

LACNIC Did not reach consensus at LACNIC X.

RIPE NCC Not submitted as a proposal.

prop-046 IPv4 countdown policy proposal (version 2)

This proposal focuses on measures that could be taken globally in the address management area to prepare for exhaustion of the free IPv4 pool.

Version 1 of this proposal was presented at APNIC 23; however, it did not reach consensus.

Status in other RIR regions:

AfriNIC Discussed informally at AfriNIC 6.

ARIN Version 1 did not reach consensus at ARIN XIX.

LACNIC Version 1 did not reach consensus at LACNIC X.

RIPE NCC The Version 1 comment period has ended. It is awaiting decision by the proposal authors.

prop-043 Proposal to remove reference to IPv6 policy document as an 'interim' policy document

This proposes to remove the reference to the "IPv6 Address Allocation and Assignment Policy" document as an 'interim' policy document .

The proposal was presented at APNIC 23, where it did not reach consensus. It was returned to the Policy SIG mailing list for further discussion.

Status in other RIR regions:

AfriNIC The 'interim' text does not appear in the policy.

ARIN To be implemented by 15 September 2007.

LACNIC Reached consensus at LACNIC X and is awaiting Board approval.

RIPE NCC Not submitted as a proposal.

prop-042 Proposal to change IPv6 initial allocation criteria

This is a proposal to remove the need to have "a plan to make 200 /48 assignments in two years" and replace it with "a plan to make a reasonable number of assignments in two years".

This proposal was presented at APNIC 23, where the proposer agreed to modify the proposed change to "a plan to make assignments within two years". However, the proposal did not reach consensus at APNIC 23 and was returned to the Policy SIG mailing list for further discussion.

Status in other RIR regions:

AfriNIC **The policy states:**
d) show a reasonable plan for making /48 IPv6 assignments to end sites in the AfriNIC region within twelve months. The LIR should also plan to announce the allocation as a single aggregated block in the inter-domain routing system within twelve months.

ARIN **Open for discussion. The proposal text differs:**
d. be an existing, known ISP in the ARIN region or have a plan for making at least 20 /48 assignments to other organizations within five years.

LACNIC **A version of this proposed change was implemented by LACNIC V in March 2004:**
d) Announce a single block in the Internet inter-domain routing system, aggregating the total IPv6 address allocation received, within a period not longer than 12 months.

e) Offer IPv6 services to clients physically located within the region covered by LACNIC within a period not longer than 24 months.

RIPE NCC **Implemented 30 July 2007:**
c. have a plan for making sub-allocations to other organisations and/or End Site assignments within two years.

How APNIC policies are developed

APNIC's policies are developed by the membership and broader Internet community. The major media for policy development are the face-to-face Open Policy Meetings, which are held twice each year, and mailing list discussions.

APNIC's policy development process is:

Open	**Anyone** can propose policies. **Everyone** can discuss policy proposals.
Transparent	APNIC publicly documents all policy discussions and decisions.
Bottom-up	The community drives policy development.

APNIC documents all of these discussions and decisions to provide complete transparency of the policy development process.

Before the meeting

You must submit your proposed policy or amendment to the APNIC Secretariat at least four weeks prior to the meeting at which the proposal will be considered.

After the SIG Chair accepts the proposal, it will be posted to the mailing list so that the community can discuss it. This allows anybody to discuss the proposal, and it is an important way for people who cannot attend the meeting to have their say. All discussion is taken into account when the proposal is discussed at the APNIC Open Policy Meeting (OPM).

At the meeting

At the OPM itself, the proposed policies are presented during the appropriate SIG session. This is your opportunity to present your proposal in person, or by other means if you are unable to attend. The community will use this opportunity to comment on the proposal.

If the proposal reaches consensus, the SIG Chair reports the decision at the APNIC Member Meeting (AMM) at the end of the week. The APNIC membership is then asked to endorse the SIG's decision.

After the meeting

Within a week of the proposal's endorsement at the APNIC Member Meeting (AMM), the proposal is sent back to the mailing list for an eight-week comment period. If any changes were made to the proposal during the APNIC meeting, this eight-week comment period gives the community the opportunity to comment on the modified proposal.

If the proposal is deemed to have reached consensus during the eight-week comment period, the SIG Chair will ask the APNIC Executive Council (EC) to endorse the proposal.

After the APNIC EC endorses the policy proposal, the APNIC Secretariat implements the policy. This usually occurs a minimum of three months after EC endorsement.

For more information please see:

http://www.apnic.net/docs/policy/dev

APNIC policy development process

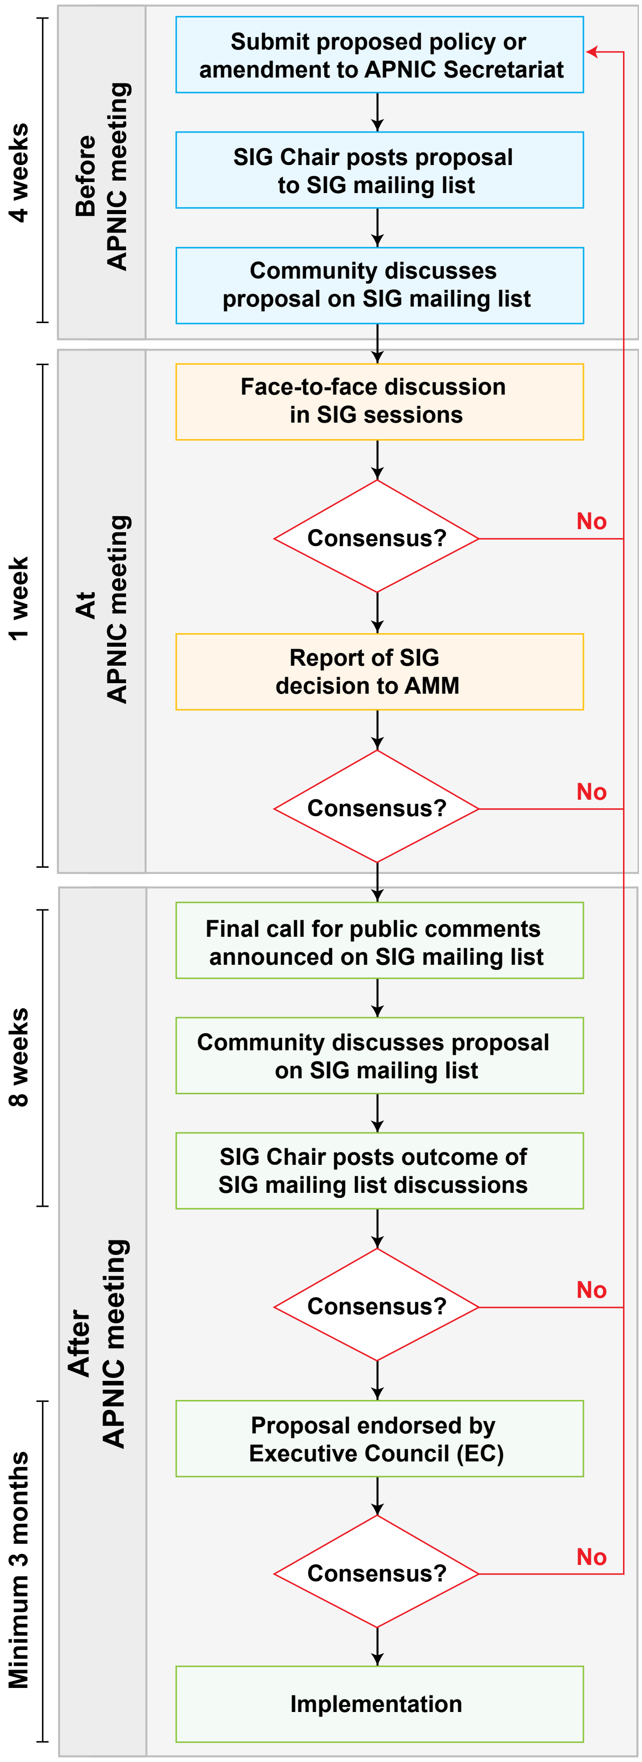

4 weeks
Before APNIC meeting
Submit proposed policy or amendment to APNIC Secretariat
SIG Chair posts proposal to SIG mailing list
Community discusses proposal on SIG mailing list
1 week
At APNIC meeting
Face-to-face discussion in SIG sessions
Consensus?
No
Report of SIG decision to AMM
Consensus?
No
8 weeks
Minimum 3 months
After APNIC meeting
Final call for public comments announced on SIG mailing list
Community discusses proposal on SIG mailing list
SIG Chair posts outcome of SIG mailing list discussions
Consensus?
No
Proposal endorsed by Executive Council (EC)
Consensus?
No
Implementation

History of SANOG

SANOG (South Asian Network Operators Group) was founded in January 2003 in Kathmandu. The first SANOG meeting was hosted in conjunction with the Computer Association of Nepal's Annual IT conference. Subsequent meetings have been held in Colombo, Bangalore, Dhaka, Thimphu, Mumbai, and Karachi. For the first time, SANOG and APNIC will be holding their meetings collaboratively in New Delhi from 29 August – 7 September 2007.

The Internet was introduced to the South Asian region in the early 1990s. Nepal, Bhutan, Bangladesh, and the Maldives were connected via satellite, and India, Pakistan, and Sri Lanka were connected through submarine cables. After a wave of deregulation swept through South Asia in the late 1990s, Internet service providers were established in large numbers throughout the region. As the Internet's importance grew and ISPs started to multihome and provide more services, the need arose for a forum to share their experiences.

As a result, SANOG was created to fulfill this need. SANOG enables ISPs in the South Asian region to communicate with each other, vendors, and engineers. Many customers also attend these meetings to keep abreast of recent industry developments.

The SANOG meeting structure includes workshops, tutorials, and conferences. Workshops are delivered in learning labs that enable hands-on experience with the latest best practice for service providers.

Popular past workshops included routing, multihoming, security, network management, DNSSec, and IP services. SANOG tutorials cover the latest advances in networking technology; and in the past they have covered topics such as security, MPLS, QoS, VoIP, and Internet exchange points.

▲ Thanks to Gaurab Raj Upadhaya from SANOG for providing information for this article.

SANOG collaborates closely with APNIC. They have hosted APNIC workshops and tutorials, featured APNIC presentations at SANOG meetings, and hosted APNIC OPM showcases.

SANOG would like to thank all of the organisations that have provided support over the years. The Internet Society (ISOC) has provided funding for the SANOG fellowship program since SANOG 3. The Network Resource Startup Center (NSRC) has also provided workshops and supported SANOG's activities. SANOG has also received assistance from other Internet related organisations, such as PCH, ISC, and Autonomica.

SANOG has been very well received in the South Asia region, and have already decided on the locations for the next three SANOG meetings.

SANOG 11: Dhaka, Bangladesh, 3-18 January, 2008

SANOG 12: Kathmandu, Nepal, 6-14 August, 2008

SANOG 13: Lahore, Pakistan, TBA

RIPE requests ICANN sign DNS root

The impact of not having a signed root is currently more apparent in the RIPE region than anywhere else. The RIPE NCC has begun using DNS Security Extensions (DNSSEC) 'in production' and so far has a number of signed zones.

Although they have been signing their reverse map zones, their parent zones 'in-addr.arpa' for IPv4 and 'ip6.arpa' for IPv6 are not signed. Thus, anyone relying on RIPE's DNSSEC activities have to configure multiple trusted keys.

In addition, the Swedish ccTLD is also signing their zones.

In May 2007 during the RIPE 54 Meeting in Tallinn, Estonia, the RIPE DNS Working Group agreed to formally request that the Internet Corporation for Assigned Names and Numbers (ICANN) sign the DNS root as soon as possible.

The working group formulated a statement and asked the RIPE meeting for their endorsement during the closing plenary. The statement was unanimously supported and a formal letter of request was sent to ICANN. The request was:

> *The lack of progress towards the deployment of DNSSEC is undermining the stability and security of the Internet. Operators and implementers are compelled to adopt ad-hoc, short-term solutions which will create long-term problems. The RIPE community urges ICANN to speed up and improve its efforts to get the root zone signed.*

This issue was discussed in July 2007 at the Internet Engineering and Planning Group (IEPG) meeting in Chicago, which was a lead-up event for the 69th Internet Engineering Task Force (IETF) meeting. The Internet Assigned Numbers Authority (IANA) gave an update on progress being made to implement signing of the root/infrastructure TLD zones.

They also demonstrated some experimental scripts that are being used to run their zone signing.

Notes from IEPG 2007:

http://www.potaroo.net/iepg/2007-07-ietf69/notes.txt

Example of signed zone put up by IANA for testing:

https://ns.iana.org/dnssec/status.html

DNSSEC links:

Theory	http://www.potaroo.net/ispcol/2006-08/dnssec.html
Practice	http://www.potaroo.net/ispcol/2006-09/dnssec2.html
Opinion	http://www.potaroo.net/ispcol/2006-10/dnssec3.html

Reuse of 240/4 address space for private use

A recent Internet draft, submitted by APNIC staff, proposes the redesignation of the IPv4 address block 240/4 from 'Future Use' to 'Limited Use for Large Private Internets'. This address space was originally designated by the IETF as 'Class E' address space and reserved for future use, but it is still unused today. As we approach the exhaustion of the IPv4 free pools held by IANA and the RIRs, it seems prudent to put this address space into use in a productive way.

The draft is referred to as "draft-wilson-class-e".

The motivation for this proposal is to service the demands of large networks that will be deployed behind NAT, which we believe will increase significantly as we undertake the transition to IPv6 through an extended period of dual stack IPv4/IPv6 networking. Such networks, large enough to exceed the existing private address space available under RFC 1918, are certainly likely in future. The use of public IPv4 address space for such purposes could be very wasteful, and could also consume a large proportion of the remaining IPv4 address space.

Private use address space and RFC 1918

RFC 1918 was written in 1996, and designates three separate blocks of IPv4 address space for private use. These blocks are:

10.0.0.0 - 10.255.255.255 (10/8 prefix)

172.16.0.0 - 172.31.255.255 (172.16/12 prefix)

192.168.0.0 - 192.168.255.255 (192.168/16 prefix)

Private address space can be used in many networks simultaneously, provided that it is not routed onto the global Internet. Such networks are therefore either not connected at all, or may be connected to the Internet through a NAT.

The problem with RFC 1918 address space today is that it is simply too small for some of the large private network applications that are being planned. Anyone wishing to build such a network is forced to use public space, and would be entitled under today's policies to receive that space, even if it is being used privately.

Are there any risks?

It seems from initial testing that most IP protocol implementations are correctly programmed to reject 240/4 address space and not to process or forward packets bearing those addresses. This makes it unlikely that anyone could put 240/4 addresses into use until equipment and software vendors make the necessary modifications. It also means that when such implementations are available, there may be a need for testing and verification, particularly in a heterogeneous environment.

It is for these reasons that the Internet draft specifically states that 240/4 is for 'limited use', specifically for large private networks. In other words, this address space should not be seen as equivalent to RFC 1918 address space or as a substitute for that space in general applications. It is intended specifically for those special purposes that require more space than is available under RFC 1918.

It should be noted that the current status of 240/4 makes it rather safe for it to be used in private networks since any 'leakage' of routes from this block onto the public Internet is unlikely to be propagated. However, any operator proposing to use the space in a private network deployment should also verify that behaviour.

Are there any alternatives?

It would be possible to allocate existing unicast IPv4 space for private use; however, this would remove precious /8 blocks from the existing pool. It would be very hard to arrive at a global consensus on the amount of space to be allocated because the cost benefit trade-off is not well defined. The other risk is that, in the case of route leakage, such space is likely to propagate through any networks without specific filters unless a global upgrade is performed.

Next steps

If there is clear support for this proposal, it would be passed through the formal IETF RFC publication procedures and adopted. At that time, it can be officially referenced by developers as a basis for the incorporation of the appropriate handling provisions into their equipment. The required changes are likely to be trivial in terms of the code and algorithms.

However, before publication of an RFC, prospective users of this address space are encouraged to contact their vendors to support this initiative and ensure that it is known and understood.

IANA update

Leo Vegoda, IANA

The first half of 2007 has been very busy for the IANA in general and the IPv4 registry in particular. Some of the updates to the IPv4 registry have been the result of legacy assignments being returned, while others have been new allocations to RIRs. APNIC received five /8s in January and the RIPE NCC was allocated two /8s in March and two /8s in July. In May, AfriNIC were listed as being responsible for 196.0.0.0/8. All five RIRs have old assignments registered in this /8, but only AfriNIC will be making assignments and allocations from it from now on.

We've also been busy working on updating the Public Data Network Numbers registry. This is the registry for 14.0.0.0/8, which was assigned to provide IPv4 addresses that could be mapped to the X.121 addresses used in X.25 networks. At the end of July all but one of these addresses had been returned to the IANA by the registrants. When the status of the last address is known we can prepare an update to RFC 3330 to make this /8 available for allocation to an RIR, possibly minus the first /24.

We have also been working with all five RIRs to improve the IPv4 registry's usefulness. It will soon be updated to include information about which RIR's whois server to consult for all of the 'Various Registries' /8s. The new format will also make it easier to distinguish special IPv4 reservations (such as private address space, multicast and the Class E space) from unicast space that has not yet been allocated.

This is only part of a larger set of improvements to all the IANA registries. We are working on converting them to an XML format that will allow us to vary the way we present them to different audiences. For instance, computers find it easier to parse a well-defined XML file, while people find it easier to read a web page. These formats will sit alongside the traditional plain text registries on our new web site. It is currently visible at http://beta.iana.org and will move to http://www.iana.org later on this year.

Responses to IPv4 address space consumption

Paul Wilson, APNIC

Introduction

IP address space is the Internet's fundamental numbering resource: Every device that is directly connected to the Internet requires an IP address. Today IPv4, the addressing standard that has been in place for the past 20 years, dominates the Internet. The total number of IPv4 addresses is strictly limited to a total of around 4 billion addresses; therefore, while IPv4 is dominant the size of the public Internet is also similarly limited. Consumption of the remaining unallocated IPv4 address space is accelerating, and based on current consumption rates it is projected that the remaining 'free pool' of addresses could be exhausted by the year 2011 (as shown in Figure 1, from http://ipv4.potaroo.net).

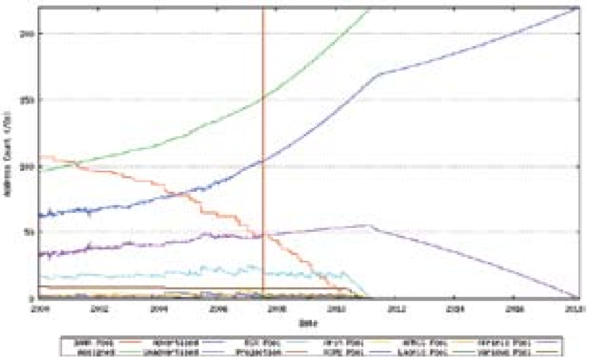

Figure 1 - IPv4 projection
Source from http://ipv4.potaroo.net

The fact that the IPv4 address space will be fully consumed in due course has always been known, and measures to ensure the ongoing growth of the Internet have been planned for many years. The best known measure is IPv6, which provides a much larger address space, as well as other useful features.

Until now, it seems that many of those who will be affected by the depletion of IPv4 addresses have been prepared to wait before taking action. However, as this critical period draws nearer, concerns are emerging about the readiness of the Internet to switch to IPv6 and about the consequences if it is not ready. At the same time, there is still great uncertainty about the timing and outcome of events, which themselves depend on how the Internet community responds to the situation. It is a classic 'Chicken and Egg' situation.

This article examines some possible responses to the current state of IPv4 address consumption, within the current IPv4 addressing environment. The aim of these responses would be a smooth transition from today's IPv4-dominated addressing system to an IPv6-dominated system, without risk to the essential attributes of the Internet. Specifically, we need to ensure that regardless of any change, the Internet remains functional, stable, coherent, and fully connected at all times.

Address policy environment

The management of Internet addresses is a matter of interest to many in the Internet community, particularly those involved in providing Internet services and infrastructure. Today a range of mechanisms exist that allow interested organisations and individuals to participate in the development of policies that govern the address management system. These mechanisms are provided by Regional Internet Registries (RIRs) and ICANN, and are generally well known within the Internet addressing and technical communities. Participants including Internet service providers, developers, users, researchers, governments, and members of society at large use them.

RIR policy process

Each RIR has a set of regional address management policies that govern its own regional processes for allocating and registering IP addresses. Each RIR also has well-defined mechanisms for developing (altering, extending or removing) those policies through Open Policy Meetings and associated activities and processes. Through these means, address management policies are continually developed and refined by the addressing community in direct response to emerging needs within the Internet itself.

It is worth noting that in each region the policy development process is itself defined by policy, which can be changed in the same open manner.

Global policy co-ordination

Among the five RIRs, there are currently nine annual Open Policy Meetings, which provide regular opportunities for collective consideration of policy issues and proposals. While these meetings serve to reflect regional priorities and objectives, they also provide an avenue for global cross-fertilisation of policy discussions and policy changes, which help to maintain a level of consistency between the regional addressing communities.

Global policy process

The Address Council (AC) of the Address Supporting Organization (ASO) is a global body whose formal responsibility is the co-ordination and oversight of global address policy initiatives. The Internet addressing community in each region appoints its members through an open election process.

According to a bottom-up process of policy development, global addressing policies must be approved by each regional addressing community before being passed to the AC for endorsement. Due to the need for support from every region the development of global policies can be slow, requiring more than one cycle of review and update within each region before consensus is reached.

Procedural responses to address consumption

RIRs are responsible for implementing policies as determined by their regional communities in the manner described above. In general, it is the address policies that determine address consumption rates and patterns, and RIR actions have little or no impact on those rates and patterns.

However, RIRs undertake a range of supporting activities, and may undertake administrative or service initiatives that respond to the current Internet environment. In this section we look at the RIRs responses, which may address some concerns about address consumption.

Global policy activities

The current processes for the development of global policies strongly support regional autonomy and bottom-up consensus decision-making. However, at a time when global policy initiatives may become quite critical, there may be opportunities to improve the efficiency of the global policy process without sacrificing its essential features.

For instance, the ASO Address Council could become more actively involved in inter-RIR exchange of information about regional policy developments, particularly about proposed global policies, or policies that have the potential to become global. The AC could also establish communication mechanisms, such as mailing lists, where global policy proposals could be raised and discussed at the earliest possible stage. This could certainly facilitate discussion and understanding of global policy proposals before they are formally introduced into each of the regional open policy meetings.

IANA request process

Currently there is a global policy governing the allocation of IPv4 address space from IANA to the RIRs, which specifies the size of those allocations. Under this policy it is possible for RIRs to receive allocations as large as five or more /8 blocks, equivalent to over 10% of the remaining IPv4 pool at this time. During 2007, the RIRs proposed and agreed to limit these allocations to a maximum of two /8 blocks at a time, regardless of entitlements under the policy.

While this measure will not have any net impact on the consumption of IPv4 addresses, it is intended to support a finer pattern of IPv4 allocation to each RIR, resulting in a distribution outcome that is more predictable and more equitable across the regions.

Administration of legacy address space

During 2007, the RIRs also undertook a joint analysis of the so-called ‘legacy’ address space, which was directly assigned by IANA prior to the establishment of the RIR system. Currently, there are 91 /8 IPv4 blocks classed as ‘legacy’ address space, representing some 35% of the total IPv4 address pool, and still significantly more than the total address space allocated to the RIRs themselves (81 blocks or 32% of the total).

The RIR analysis identified a collection of unallocated IPv4 addresses equivalent to approximately seven /8 address blocks in size. Based on this analysis, the RIRs then determined an equitable distribution of administrative responsibility for specific legacy /8 blocks to each of the RIRs, ensuring an equal share of that free address space for each RIR. This distribution was ratified with IANA in July 2007 and is being implemented at the time of writing.

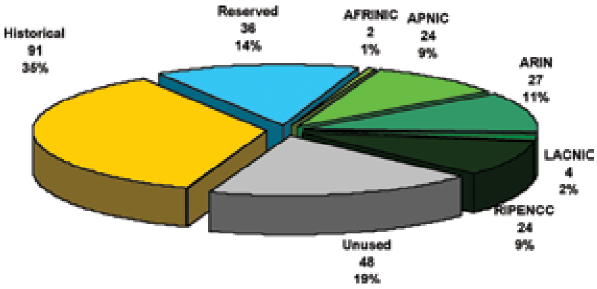

Figure 2 - Global distribution of IPv4 (July 2007)

Internet Resource Certification

With increasing interest in the efficient use of IPv4 address space, there is currently an increased focus on the importance of IPv4 registration information and on the role of RIRs as authoritative registries. At the same time, the security of registration information and the security and integrity of the Internet’s routing system are also becoming critical issues.

The RIRs, along with the IETF community, are currently examining options for the digital certification of Internet resource information as a means of addressing these needs, and in some cases conducting technical development and trials. APNIC is particularly active in this area and will deploy certificate management services and tools to its members beginning early in 2008.

Policy responses to address consumption

Numerous policy initiatives have been proposed in the past that have implications for IPv4 address space consumption, and in the current environment, more proposals can be expected in future. In this section, some past and current proposals are reviewed. However, there is no suggestion that any measure should be implemented without full reference to the RIR and/or ASO policy processes.

IPv6 considerations

Today, IPv6 standards are in a mature state, IPv6 products are available for many platforms and IPv6 connectivity is supported by many applications. IPv6 address space is actively being allocated, and some production IPv6 networks are being deployed. However, the current rate of this development is still not high enough for a smooth transition to IPv6 within the next two or three years.

The reasons for the lack of widespread IPv6 deployment in the public Internet to date have been debated widely in recent times, and they appear to be largely commercial rather than technical. It is clear that despite strong promotion efforts, IPv6 is yet to produce a business case for rapid deployment by a significant number of ISPs. However, as rates of deployment accelerate it is now becoming critical to enable a smooth transition when needed. Without sufficient time for this transition, it is likely that new Internet infrastructure will be delayed, and existing infrastructure may possibly be disrupted during a rushed and compromised transition process.

While IPv6 address policies will be revisited and may be revised at any time in the future, it seems that no one is suggesting that substantial IPv6 address policy changes are needed at this time. On the contrary, RIR communities have taken every measure available to ensure that there are no policy impediments to IPv6 address space availability for any viable application. However, all RIRs have become actively involved in recent years in the promotion of IPv6 within their communities, particularly in operator community education and training, and in government circles.

Address space reclamation

Of the IPv4 addresses that have ever been allocated, approximately 25% (or around 48 /8 blocks) do not appear in the routing system of the public Internet. This unrouted address space is certainly not used in the public Internet, and it is not known how much is in use in private networks, or how much is completely unused. In either case, the addresses could possibly be ‘reclaimed’ for use on the public network.

It has often been suggested that unused address space should be reclaimed for reuse. However, the lack of conditions on early allocations can make this process very difficult. Particularly in the USA, where most legacy allocations exist, reclamation processes are not only likely to be lengthy and difficult, but also legally expensive. At this time, only APNIC and LACNIC have active reclamation processes in place while in other regions address space returns are effectively voluntary. In either case, the amount of address space that can be recovered is relatively insignificant.

Private use address space: 240/4

Given the interest in large-scale, NGN-type projects, it seems that new network deployments on the drawing board are essentially private networks, but are too large to use existing private-use (RFC 1918) address space. Proposals have previously been made for allocation of additional IPv4 address blocks for private use, but these did not achieve consensus across the addressing communities. Without additional private space being made available, it is possible that such network deployments might consume large amounts of IPv4 addresses in future, contributing unnecessarily to address space depletion.

P 10

Another possibility under discussion is the redesignation of some or all the existing reserved '240/4' address space (comprising 16 /8 blocks of IPv4 address space) for this purpose. This space is still unused, and there are no other proposals for its use. It is also generally regarded as unsuitable for public unicast use as 'normal' Internet address space due to the need to upgrade many different devices across the Internet (similar to the challenge of transitioning to IPv6). However, within a private network setting, an operator would be able to assess, manage, and control the upgrade process according to their needs based on a full knowledge of the infrastructure in use and the costs and benefits of that upgrade.

IPv4 'flag date'

A recent policy proposal has suggested that a specific date should be established on which all RIRs will cease to make IPv4 address allocations, and that remaining addresses at that time should be reserved for future critical purposes. In addition, it is proposed that once the date is established, no further changes to address policy should be made in order to "ensure steady provision of IPv4 address space." The intention is to provide some certainty as to when alternative arrangements (namely IPv6) must be in place, and to avoid perceptions of unfairness between or within RIR regions.

To date, this proposal has not achieved consensus in any RIR region. It appears that communities are unwilling to voluntarily impose constraints on possible future action, as would be implied by the choice of a flag date and the prevention of future policy changes.

Allocation of remaining IPv4 address pool

Another recent global proposal has suggested that when the IANA IPv4 pool is reduced to a certain level (for instance to 20 or 15 /8 blocks), that the remaining address space should be divided equally among the five RIRs. The purpose of this proposal is to allow each individual RIR to autonomously determine its own chosen approach to the distribution of those blocks, and also to avoid a 'late run' on the central IANA pool, which would possibly favour some regions over others. The proposal was approved in the LACNIC region, and will be discussed in other regions during 2007.

Address space transfer and trade

It has been suggested that the immediate outcome of the exhaustion of the unallocated IPv4 address pools would be the emergence of a market for IPv4 address space in which addresses themselves, or licences for 'right-of-use' of address space, could be traded. Such a trading scenario could assume one of many possible forms, and could have a range of possible effects, both positive and negative.

The sudden emergence of a market for IP addresses, particularly after the consumption of existing IPv4 supplies, could create a significant disruption in the distribution of IP addresses and the integrity of the Internet itself. Possible risks include the escalation of prices for IP address space beyond the reach of many address users, the fragmentation of address blocks resulting in routing problems, illegal trading and fraudulent claims on address space, market distortions (such as hoarding, price speculation and attempts at seizure of control), the emergence of conflicting markets, and an Internet without a single authoritative system for address registration. On the other hand, it is also conceivable that a trading market could emerge in an orderly fashion: one that has the support of the Internet addressing community, in which the RIRs continue to play their crucial roles as registries of current address space holdings and of address 'right-of-use' licence transfers.

The implications of such a scenario deserve a more thorough and extended analysis than is possible here; however, the following observations may be made:

- IP address transfers are already undertaken by RIRs on a regular basis, with mergers and acquisitions of network providers or infrastructure.
- Today's address transfer policies could be relaxed to allow transfers to be recognised in a market environment, probably without major cost or administrative changes.
- Digital resource certificates, currently under consideration for use in routing security, may be adapted to represent address right-of-use 'licences' for use in transfer or trade.
- If transfer policies are adjusted before the exhaustion of IP address space, then address space users would have a choice in the source of addresses (but not in the choice of address registry).
- Existing address allocation services provided by the RIRs, along with fees charged for those services, would tend to moderate the market and impose a limit on price escalation.
- While address stockpiling would be possible, today's RIR policies of allocation only for demonstrated need would tend to prevent RIR allocations from being stockpiled.
- The speculative value of IPv4 addresses would likely be severely limited by the inevitable advent of IPv6 addressing, while any form of severe price escalation in an IPv4 market would only add further impetus to IPv6 deployment.
- A market for address space would provide strong incentives for unused IPv4 address space to be brought 'into circulation', and would relieve pressure on the remaining IANA address pool.
- The advent of a market of some kind may be inevitable after the exhaustion of the remaining address pool. If so, the structuring of such a market to avoid industry disruption and market distortion would be a critical role for RIRs.

Conclusion

This article has explored some current initiatives that may affect the management and consumption of the remaining IPv4 address space. Some of these are operational initiatives that can be easily implemented, while others involve deeper changes at the policy level that will certainly need further discussion by the Internet addressing community. The recent acceleration of IPv4 address space consumption, along with the ongoing slow (up take) of IPv6, add some extra urgency to these discussions and should be a cause for concern for everyone in this community.

Given the importance of addressing to the Internet and its development, it is no surprise that these issues are finding a much larger community of interest these days, or that they have been placed on the agenda of the Internet Governance Forum meeting that will be held later in 2007. It is the responsibility of all of us to become acquainted with and involved in these discussions, and to find workable solutions within the fairly near future.

Analysis of earthquake disruption of submarine cable system

Of the many potential threats to submarine cable systems, perhaps the two most prevalent are those caused by shipping activity and those caused by submarine earthquakes. Shipping incidents usually occur close to the cable landing points, where the cable lies in shallow water and can be snagged by anchors, or as a result of seabed trawling.

For this reason, the undersea cable that lies in shallow water is sheathed with an additional 5-8 cm of steel jacketing. However, deep-sea cable segments are not clad with this additional armour, and this leaves them particularly vulnerable to submarine earthquakes.

This is exactly what happened one evening in December 2006, 90 km south of Kaohsiung, Taiwan, in the Luzon Strait. A relatively powerful earthquake, it had a magnitude of 7.1, placing it squarely in the 'major' category. Seven aftershocks followed in the subsequent 48 hours, all with a Richter magnitude in excess of 5.2.

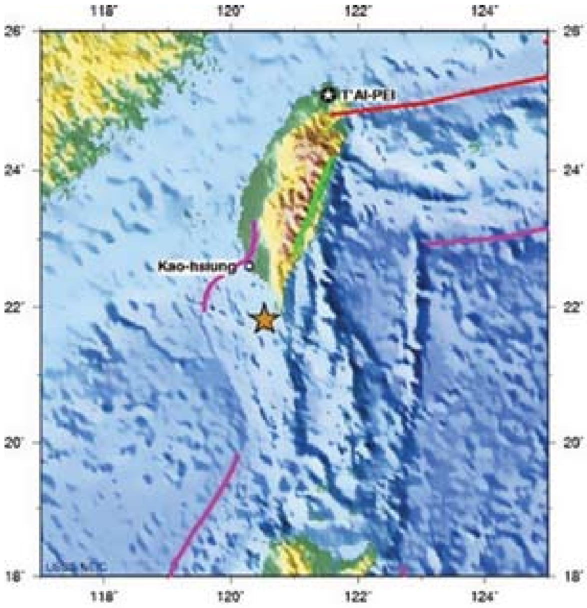

▲ **Figure 1** - Taiwan region
26 December 2006 12:26:21 UTC
21.82N 120.53E
Magnitude: 7.1
Source: earthquake.usgs.gov

Submarine earthquakes are not uncommon, but the Luzon Strait, which connects the South China Sea to the Philippine Sea, is particularly critical in terms of submarine cables. The Formosa Sea, which lies to the west of Taiwan, is too shallow for submarine cables, and the path to the south of the Philippines is impractically long.

As a result, all nine of the major submarine cable systems that connect Southeast Asia and the Indian subcontinent to North Asia and North America pass through the Luzon Strait. Even where the cable is configured as a self-healing ring, both segments of the ring pass through this same body of water.

The reason submarine earthquakes have such destructive potential is not just due to the physical disruption of the earthquake itself, but also the capacity of the earthquake to trigger oceanic landslides. Lasting from minutes to hours, these landslides travel across the sea floor at speeds ranging from tens to hundreds of kilometres per hour.

It appears that such a submarine landslide occurred in this case, as indicated by the cable outage log. The only cables left in service were Asia Netcom's EAC cable system and the Guam-Philippines Fiber Optic system. Obviously, the major impact here in terms of Internet communications was an extended loss of transit bandwidth for Southeast Asia, including Hong Kong, Singapore, Malaysia, Indonesia, and India.

Cable	Outage time (UTC)
SMW3 S1.7 & S1.8	12:25 26/12
China-US W2	12:27 26/12
RNAL Busan / TongFul	12:43 26/12
APCN2, Seg 7	16:06 26/12
APCN2, Seg 3	18:01 26/12
APCN Sys 1, Seg B17	18:15 26/12
China-US S1	18:59 26/12
RNAL HongKong / Toucheng	19:42 26/12
APCN Sys 2, Seg B5	20:44 26/12
FLAG FEA Sub-Sys B	20:56 26/12
China-US W1	02:07 27/12

Source: http://www.hardwarezone.com.au/news

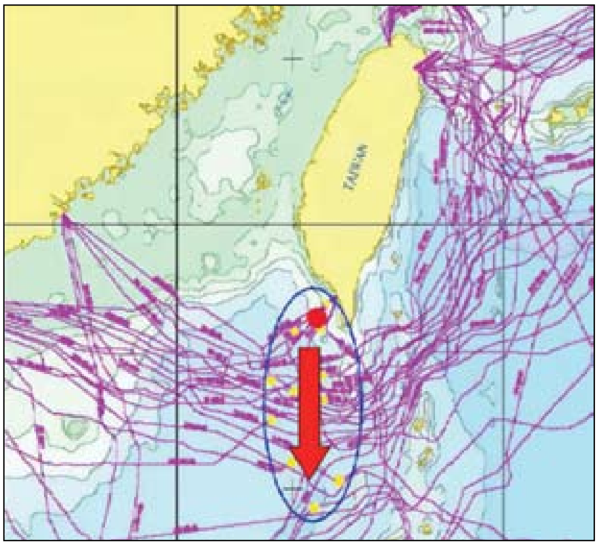

▲ **Figure 2**
Source: PCCW
The red dot indicates the epicentre of the earthquake, which triggered a landslide that travelled in the direction of the red arrow. The purple lines indicate submarine cables. The yellow dots indicate where breakages occurred on the submarine cables.

P 12

P 11

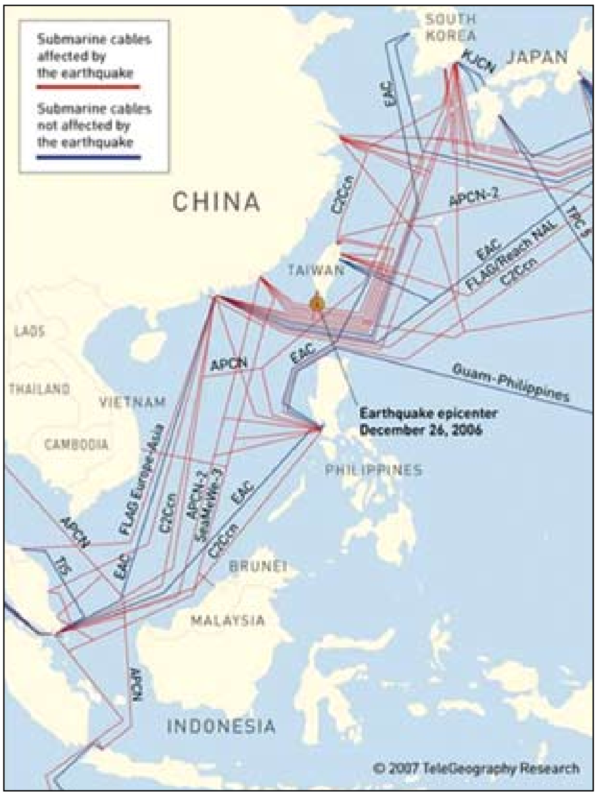

Figure 3

Source: Telegeography Research

The depth of the Luzon Strait is one of its major advantages as a cable location. Kilometres below the ocean's surface, there is no human activity that could snag the cable. However, when a cable break does occur, this depth becomes a serious challenge: It is beyond the operating depth of the cable repair Remotely-Operated Vehicles (ROVs) that would normally be used to locate the cable and bring it to the surface for repairs.

At these depths, the repair ship is forced to resort to an older method: Grapnels are dragged across the sea floor in an attempt to snag the cable, cut it, and recover it to the surface. This can be a tricky operation at the best of times, but when a submarine landslide has moved the cable and possibly buried it as well, the recovery operation can be extremely difficult and time consuming.

It took some seven to ten days to repair each cable, and a total of 18 separate cable repair operations were necessary in this case. Given the scale of the operation, it is remarkable that all cables were reported to be back in service by 14 February 2007.

These days, the bulk of what these systems relay is IP traffic, so it is helpful to look at these kinds of disruptive events from the perspective of the inter-domain routing system.

Renesys' analysis of the earthquake showed that the first set of cable outages, which occurred at 12:30 UTC on 26 December, triggered a shift of traffic onto backup paths (presumably the other half of the ring in many cases), and few routing outages were noted.

The first set of network outages occurred eight hours later, at around 19:30 UTC on the same day, which coincided with the second set of cable outages. At this time, both halves of the APCN ring were taken out of service. However, the major event occurred at 02:00 UTC on 27 December, when the second half of the China-US cable system ring was also rendered useless. At this point, the total number of affected route destinations reached a peak of 4,000 prefixes.

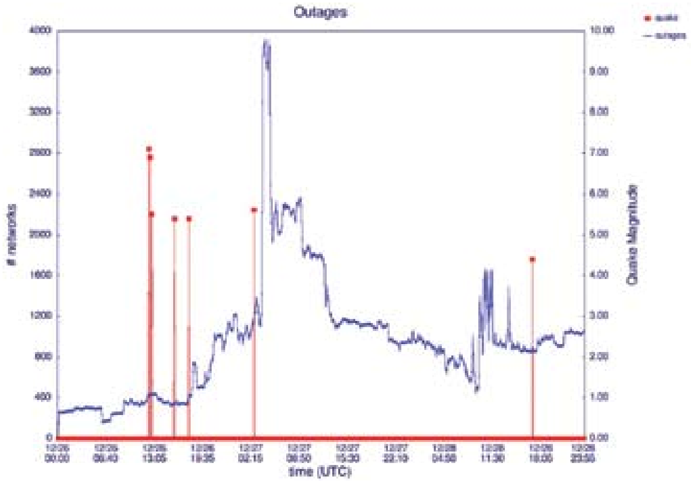

Figure 4

Source: Renesys

It was also clear by this time that the region's service providers' operations teams were fully engaged with the issue, and connectivity for 2,000 of these routes was restored within two hours of this final major outage.

What this tracking of routing connectivity does not show, however, was that while basic connectivity was generally restored within one to two days of the earthquake, the available bandwidth was severely diminished for a much longer period. Indeed, the impact of the earthquake on the end users, particularly those located in Southeast Asia, lasted for the next four to six weeks.

This earthquake revealed the underlying fragility of the infrastructure that supports global communications, and highlighted the potential that local events have to create a worldwide impact. The Luzon Strait is a critical point of vulnerability as it is a geologically active area located on part of the broader 'Ring of Fire' that encompasses the Pacific Rim.

One might expect the normal response to be the construction of alternative paths to provide some resilience against a recurrence of this event. The problem here is that the alternatives to passing through the Luzon Strait also represent high risk, high cost, or both. Alternative cable paths traverse the Indian Ocean, the Mediterranean Sea, and the Atlantic Ocean, but such paths are accompanied by much higher costs, and consequently are not heavily used.

One of the facts this event has revealed is that the designers of the Internet's infrastructure appear to naturally tend towards approaches that offer cost savings, even when such preferences represent a higher risk in terms of resilience in the face of disruptive geological phenomena.

Acknowledgements

Thanks to Todd Underwood and Renesys, and Sylvie La Perrière and Teleglobe VSNLI for their presentations on this topic at APRICOT 2007, from which parts of this article have been based.

Geoff Huston
Chief Scientist, APNIC

DNS measurement activities

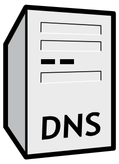

APNIC is currently bringing a new DNS statistics node online.

Currently, APNIC undertakes DNS operational logging and sampling measurement; however, these activities are unlikely to meet our future requirements. We now have a heavy-duty dual-CPU Dell 860 server to act as our new statistics node.

This new system is based on a completely different model to the one currently operating. Two 'passive aggregating' taps monitor network traffic flow. These are directly connected to our DNS servers and are continuously aware of the network state and their own internal states.

This is a 'fail safe' implementation: If a problem arises, the taps immediately shut themselves out of the circuit so that normal operation is not interrupted. This is quite an accomplishment at gigabit Ethernet speeds.

The statistics node collects all DNS traffic flowing to our two main servers (around 60Gb/day), and it has enough storage capacity to hold more than a week's worth of data online before having to be cycled.

The current statistical sampling method on this node will soon be superseded by a more powerful data analysis model that also allows the data to be held for much longer periods.

The new data analysis model was chosen after careful review of the available options. A product titled the DNS Statistics Collector (DSC), produced by the Operations, Analysis, and Research Center (OARC) was ultimately selected for our purposes.

APNIC is a member of the OARC, which is a research community consisting of DNS operators and other interested parties. OARC has a large system that acts as a central research 'repository' and collects/collates input from many different DNS servers worldwide.

We also intend to introduce reverse DNS measurements, using the OARC's data format.

In the event of anything unusual or interesting occurring on the network, with our week or more's worth of full packet capture available, we will be able to go back over the data to aid with reduction and analysis.

APNIC would like to express its gratitude to the OARC, ISC, and the WIDE community for their assistance with this new deployment.

New BGP research node in Japan

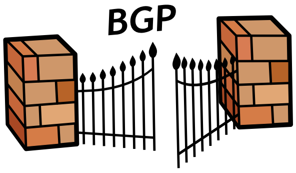

APNIC has recently brought a BGP (Border Gateway Protocol) research node online in Japan.

This node is a Unix host, running the Quagga routing software suite with a modification to support 4-byte AS and configured to act as a 4-byte AS BGP speaker.

This research node is peering with DIXIE, the Japanese exchange run by WIDE. It is also connected to another APNIC BGP research node located in Brisbane using the Exterior Border Gateway Protocol (EBGP). Soon, we expect to add IPv4 and IPv6 BGP route peering sessions both directly across DIXIE and via remote multihop BGP sessions.

As part of our routing research, APNIC will be announcing at least one IPv4 and one IPv6 prefix from this node. There will not be any traffic implications associated with this announcement because no services will be hosted on these addresses; also, no planned experiments will rely on connectivity via these experimental prefix announcements. Consequently, routing transit for these prefixes can be provided safely without risk of incurring data transit overheads.

The objectives of deploying this node are:

- To provide a neutral BGP peer in Japan, that other entities can connect to directly (and via EBGP multihop) to explore the interoperation of their routing systems with this experimental 4-byte AS routing peer. For example, KDDI and NTT have both announced their intention to introduce research 4-byte AS BGP peers in the near future. Therefore, this APNIC research facility at DIXIE should be useful as a visible 4-byte AS route collector to support this work.
- To provide a 4-byte AS 'beacon' that can be used to announce and withdraw prefixes in a predictable and well-understood manner. This will allow anyone to explore the visibility of 4-byte and 2-byte to 4-byte AS behaviours as routes are announced and withdrawn within the IPv4 and IPv6 routing realms.
- To provide another measurement and data collection point for ongoing research into routing behaviour and the scaling properties of the routing system.
- To demonstrate that 4-byte AS is a well-understood and safe technology that can be deployed at any time, by any ISP.

APNIC requires support for this research activity. We are seeking expressions of interest from ISPs who are in a position to offer full BGP route feeds in IPv4 and IPv6, and those who are willing to advertise data-less transit for a small set of APNIC research prefixes originating from this location.

If you can assist APNIC in this activity, please contact:

research@apnic.net

APNIC Resource Certification project

Resource Certification is a new service that APNIC will be providing its members in the near future. This article follows on from a previous Resource Certification article published in Apster in August 2006. We will continue by describing what Resource Certification is, why Resource Certification is useful, and how it works.

What is Resource Certification?

To answer this question, it would be useful to first look at the role of 'Public Key Certificates', in the context of public/private key cryptography.

Public/private key cryptography involves a pair of cipher keys with a unique property: Any object that is encrypted with one of the keys can only be decrypted with the corresponding key.

To use this technology, you first generate a key pair. You keep one of the keys – your 'private' key – a closely guarded secret, and publish the other key openly. When you want to send a message that is guaranteed to be authored by you, you simply sign the message with your private key. Only your public key can unlock this signature.

Conversely, if someone wants to send you a message that only you can read, they encrypt the message using your public key. Only your private key can unlock the message.

This is fine in theory; but if, for instance, I want to send you a message, how do I find out your public key, and how can I tell you my public key?

We could meet in person somewhere and exchange these public key values. That way, we could communicate securely and privately, with the knowledge that our communications were private for as long as our private keys were kept private. But, what if we could never meet? How could I trust that the public key purported to be yours is authentic, and doesn't belong to someone else masquerading as you?

Here's where Public Key Certificates play an important role. Certificates allow a third party, in the role of a trusted certificate authority, to attest that the public key that they are publishing in a certificate belongs to the party identified in the certificate, and they sign this attestation with their private key.

As long as you trust the integrity of this certificate authority, and you have a copy of their public key, then you do not have to meet in person to exchange public keys. As long as you are satisfied that the certificate authority has done its job correctly, then you can trust the certificate authority to vouch for the authenticity of my public key.

Public key certificates frequently mention identity or role authorities, such as "This party is identified as APNIC," or "The holder of this certificate is an APNIC hostmaster," or "This is APNIC's web site."

Resource Certificates extend the certification model to make a slightly different attestation: The holder of the corresponding private key is the current 'right-of-use' holder of a specific set of address and AS number resources, where the address blocks and AS numbers are listed in the Resource Certificate, along with the public key of the resource holder.

It is envisioned that Resource Certificates will be issued by the actual resource allocator, rather than just any Certificate Authority. Thus, a resource certificate issued by APNIC relating to member 'A' with resources 192.0.2.0/24 should only be issued by APNIC, and only if member 'A' is indeed the current holder of that address resource. The more general constraint is that the Resource Certificates can only be issued by the party that allocated the resources in the first place.

A fully populated resource certification authority hierarchy should follow the IP address and AS delegation hierarchy, with IANA certifying the RIRs, the RIRs certifying Local Internet Registries (LIRs), and LIRs certifying their downstream resource allocations.

Why do we need Resource Certification?

Routing security has always been a critical factor in successful Internet operation, and one of the ways to subvert or disrupt the Internet is to inject false information into the Internet's routing system.

The overall aim of routing security is to allow any party to be able to validate routing advertisements. This is to confirm that the information being passed through the routing system is indeed correct, and that it corresponds to the intentions of the address holder.

One of the key cornerstones of routing security is the ability to validate the implicit claims of the right to use an AS, or to route an address prefix, in an automated and efficient manner.

The objective of Resource Certification is to create a robust framework that allows the validation of assertions relating to IP addresses and ASNs and their use. It is also intended to make it easier for anyone to see if someone is misrepresenting their control over addresses and routing.

In other words, a secure routing environment should allow anyone to be able to answer the following questions:

- Does this routing information correspond to duly delegated address resources?
- Is the routing advertisement made with the explicit agreement of the current 'right-of-use' holder of the addresses being advertised?
- Does the network path represented in the routing advertisement correspond to a valid path through the network that will reach the advertised destination addresses?

IP addresses and AS numbers are used in many places. Obviously, the routing system is one such place; but there are also whois reports, Internet Routing Registry reports, Operating Support Systems used by ISPs, and so forth.

Anywhere that IP addresses and AS numbers are used, it is possible to sign the information with a digital signature and have the corresponding certificate for the signing key be a duly issued and validatable Resource Certificate. In this way, anyone using the signed information has the capability to ensure that the information they are using is exactly the same information that was originally entered by the resource's 'right of use' holder, and that the information has not been altered, truncated, or extended by anyone else.

Resource Certification does not prevent attempts to lie or misrepresent information about resource holdings and their use, but the use of digitally signed information and the associated resource certification public key infrastructure makes such misrepresentations readily detectable by anyone.

How does Resource Certification work?

It's all about signing and validation.

The relevant observation here is that 'signing' information with a digital signature essentially 'freezes' that information, and any effort to alter the signed information results in the signature being invalidated.

This property of signatures is independent of Resource Certification, which plays a part when it comes to validating this signature. The question that resource certification can answer is: “Was this signature generated by the current ‘right of use’ holder of a given address or AS number?”

Establishing a certificate’s validity involves assembling a chain of certificates that starts with a nominated trust anchor. This trust anchor issues a resource certificate to another entity, who in turn issues a resource certificate to another entity, and so on until we reach the certificate in question.

For example, if the nominated trust anchor was IANA, validating the certificate would entail the inspection of a certificate issued by IANA that describes the APNIC resource holdings. This is followed by inspection of an APNIC-issued certificate that describes the resource holdings of a Local Registry, which, in turn, has issued a certificate to the entity being validated.

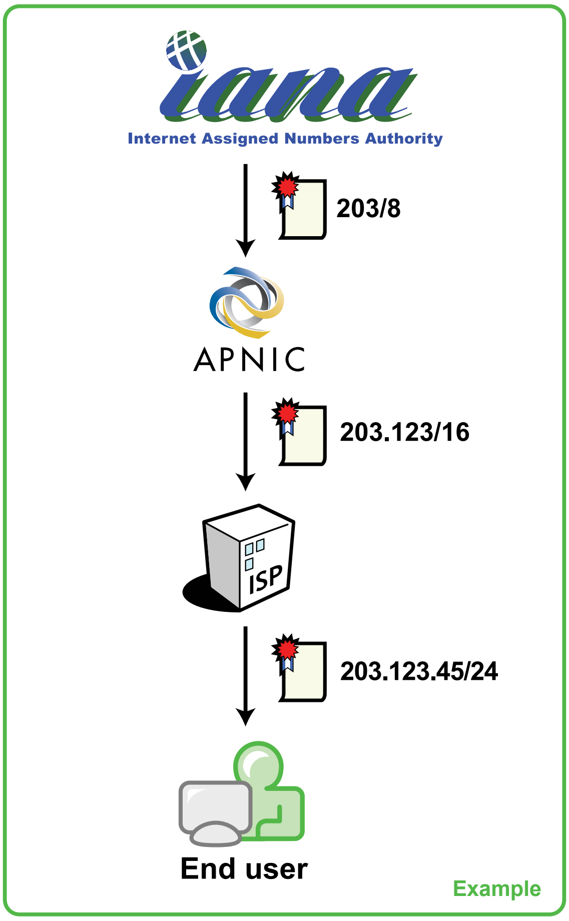

Resource certification is an explicit way of confirming the validity of resource allocations that underpin the Internet’s resource distribution system.

Potential APNIC services for Resource Certificates

The exact nature of the APNIC member services relating to Resource Certificates have yet to be fully specified; but as a general illustration of the way we anticipate Resource Certificates will be managed, here are some possible services that could support Resource Certificates:

Open software

APNIC is developing a software suite that includes modules for managing Resource Certificates, generating digital signatures using resource keys, managing repositories of Resource Certificates, and validating Resource Certificates.

The software is based on existing open source software, and uses this foundation to construct functions specific to Resource Certificates. APNIC will be making this software freely available as open source software, allowing others to use it to manage resource certificate functions.

Certificate management services

Some APNIC members may be entirely comfortable setting up local certificate management systems; however, we understand that many members would appreciate a member service option that passes responsibility for resource certificate management to APNIC.

At this stage, we envisage the extension of MyAPNIC to include functions to support Resource Certification, exposing the end use functionality of generation of digital signatures and validation of signed objects, while performing the specific tasks related to certificate management on behalf of the member as an automated internal function.

Resource Certificate repository services

As part of the certificate validation function, the end user needs to assemble a chain of certificates from a chosen trust anchor through to the certificate in question. Accessing these certificates using an on-demand access model could impose a significant overhead, particularly if validation is used in conjunction with an inter-domain routing protocol, such as BGP. We are looking at ways to manage a public aggregated repository cache, so that a member only needs to perform a daily synchronisation operation with such an aggregated cache in order to maintain their local copy of the resource certificate collection.

In summary

Resource Certification represents significant progress in managing Internet number resources in a secure and trustable fashion. The benefits, in terms of allowing anyone to make assertions about their right to use an IP address and allowing anyone else to validate such a claim in a secure and reliable fashion, represent a significant step forward in the larger effort of improving the security and utility of the Internet.

APNIC community interface

Southeast Asia liaison

APNIC Southeast Asia Liaison Officer Son Tran attended a number of APNIC training sessions in Southeast Asia in June and July. This was a valuable opportunity to find out about local IT professionals' interests and concerns.

Singapore

One of the issues raised by Singapore session attendees was the potential introduction of RIR-based anti-spam policies. This was a good opportunity to remind the attendees that participating in the APNIC Open Policy Meeting is the best opportunity to propose such policies.

Attendees were curious about how the concept of fairness was applied to resource distribution policies. They were also interested in how these policies were co-ordinated on a global scale.

APNIC would like to thank Ivan Wee from Republic Polytechnic for his invaluable help.

Bangkok, Thailand

In Bangkok, Son met with Suchok Ardhmad and Kamphol Boongsri from the Communication Authority of Thailand (CAT) and Chalermpol Charnsripinyo and Chatchai Chan from the National Electronics and Computer Technology Center (NECTEC).

Internet usage and broadband adoption in Thailand have increased markedly in the last few years, particularly as a result of a government initiative to significantly reduce broadband access costs. The Thai government also plans to migrate its entire network infrastructure to support IPv6 by 2014

NECTEC has had an IPv6 forum running since December 2004 to promote migration to IPv6. Meetings are held on a yearly basis; however, the forum faces many challenges in promoting IPv6 adoption. The organisation is currently providing IPv6 tunnelling for Thai ISPs, and is currently working to stimulate greater interest and use of this technology.

APNIC would like to thank the Asian Institute of Technology for their support.

Phnom Penh, Cambodia

Internet usage is increasing rapidly in Cambodia, and while the costs of Internet access are still considered to be high in relation to neighbouring economies, significant progress has been made in the past 10 years to enhance Internet service affordability. Home Internet services are still rare, but reasonably priced Internet cafés are becoming more widespread.

In past years most ISP operations were run by foreign interests. Today, the number of locally run ISPs in Cambodia is steadily increasing. Presently, there are around 10 ISPs operating in Cambodia; however, around 30 licences have been issued. Continued development is expected in this area as home Internet use becomes more common. In addition, Angkor Net recently introduced WiMAX, a broadband wireless service. This allows ISPs to offer broadband speeds without having to install any telecommunications infrastructure.

APNIC would like to thank Channda Sok from Anna Computer and Angkor Net, event sponsors, and the Sunway Hotel, for their assistance.

Hanoi, Vietnam

In Hanoi, Son met with Nguyen Chau Son and Cuong Nguyen from Vietnam Postal and Telecommunication (VNPT), which is currently the only ISP in Vietnam with IPv6 resources. VNPT are very eager to migrate their networks to IPv6 while minimising costs and disruption to their current services.

Son also met with Mr Hai Hong Pham from the Ministry of Posts and Telematics (MPT). Mr Hai Hong Pham said that he recognised that IPv6 deployment has become an urgent matter for Vietnam, and he will work closely with VNNIC to promote its adoption.

VNNIC are keen to work with the APNIC training team to conduct sessions about DNS delegation procedures and IPv6 development. They hope to offer these sessions in the near future.

APNIC would like to thank Mr Tan Minh Tran from VNNIC for his support.

PacNOG 3

PacNOG

PacNOG 3 was held in June this year at the Edgewater Resort on Rarotonga and was hosted by Telecom Cook Islands. About 35 people attended the event, and approximately half of the attendees were APNIC members. The vast majority of participants this year were attending PacNOG for the first time.

APNIC Internet Resource Analyst and Pacific Liaison Officer, Elly Tawhai, attended as a speaker and observer. This was a valuable opportunity to meet with APNIC members, provide information about APNIC activities and services, discuss issues, and hear their feedback firsthand.

IPv6 implementation and operational issues were high on the agenda in presentations and workshops. In addition, a wide variety of topics were covered at the meeting. These included BGP aggregation, spam, VoIP PBX using open source tools, and deploying an island-wide wireless network.

The next PacNOG meeting will be held in Vanuatu in June 2008.

http://www.pacnog.org/

▲ PacNOG 3 was held in Rarotonga from 16-22 June 2007.

NIC World Summit and CNNIC 10th birthday

In June this year, APNIC Resource Services Unit Manager and China Liaison Officer, Guangliang Pan, visited Beijing. He attended the Network Information Centre (NIC) World Summit, which coincided with 10th birthday celebrations for the China Internet Network Information Center (CNNIC).

The NIC World Summit was the venue where TLD registry representatives from CNNIC, DENIC, nominetUK, NIDA, JPRS, SGNIC, and NeuStar signed the 'Beijing declaration', stating their "desire and commitment to increase dialogue and sharing of best practice to contribute to building a harmonious information society." ICANN President/CEO Paul Twomey also attended and spoke at the event.

While onsite at CNNIC, Guangliang took the opportunity to view the I-root server, which has been operating since 2005. APNIC worked with CNNIC to install this server and now provides technical support. This project is an example of the fruitful partnership APNIC and CNNIC have enjoyed during the past ten years.

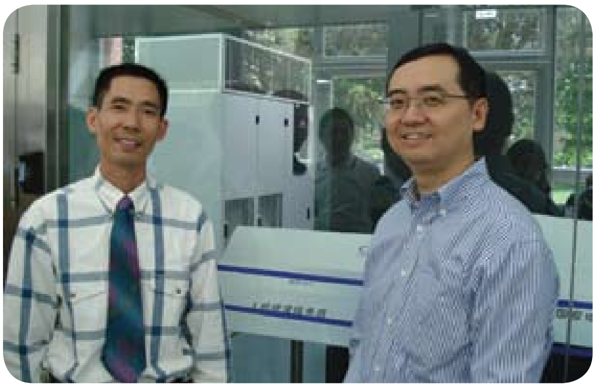

▲ APNIC's Guangliang Pan and CNNIC's Mao Wei at the Beijing I-root server installation.

JPNIC 12th Open Policy Meeting

JPNIC held their 12th Open Policy Meeting on 17 July. Our representatives Miwa Fujii and Guangliang Pan presented an APNIC update at the meeting, and also discussed a variety of topics in an informal information-sharing session with JPNIC staff.

The APNIC update presentation included an overview of topics to be discussed at APNIC 24. The JPNIC OPM participants showed particular interest in the IPv4 consumption issue.

The JPNIC OPM was attended by more than 100 people and featured a lot of lively discussion. "Attending this meeting gave us a great insight into our community's needs", Guangliang said.

▲ From left to right,
Back row: Hiroki Kawabata, Susumu Sato, Guangliang Pan (APNIC), Akinori Maemura, Taiji Kimura
Front row: Kanae Sato, Izumi Okutani

TWNIC 8th Open Policy Meeting

TWNIC held their 8th Open Policy Meeting on 6 July this year. APNIC Member Services Unit Manager George Kuo attended and presented an APNIC update at the event.

The TWNIC OPM agenda featured of variety of topics, ranging from IPv6 operations and security, through to 4-byte ASNs. During the meeting, TWNIC and NIDA signed an MoU to enhance co-operation and information sharing between the two NICs.

George also noted that 'TWNOG' featured in the TWNIC program for the first time. It appears that the TW community is very interested in having a NOG meeting, and is in the process of trying to make this happen. Stay tuned for more details.

▲ Lai Fei Pei (left) and Kwan-Ho Song (right) signed an MoU at the 8th TWNIC Open Policy Meeting.

NIR training

▲ Jin-Man Kim from KRNIC

In the interests of sharing skills and information, KRNIC Hostmaster Jin-Man Kim was a guest in APNIC's office from 6-31 August.

The training was conducted by the Resource Services Unit, and included an overview of many aspects of APNIC's day-to-day functions, including hostmaster, billing, technical, administration, and policy.

It is hoped that this type of training will further enhance our relationship with the NIRs, and enhance communication and service delivery in both organisations.

Secretariat spotlight - What's new at APNIC?

In response to member feedback, including the recent survey, the APNIC secretariat has made many changes to continually improve and streamline services and processes.

In this issue, Apster shines the spotlight on two internal units that have a significant impact on our ability to support our members: **Network Operations** and **Member Services**.

Network Operations update

In addition to providing the essential data communications services that support the APNIC secretariat, Network Operations has also recently undertaken many projects to benefit the wider Internet community.

New DNS generation system

APNIC has invested a significant amount of time and effort in the new APNIC DNS generation system. The new system is currently running in parallel with the existing system, and is set for release in the third quarter of 2007. From an end-user perspective the new system will behave the same way as the current sytem. Its features include:

- Updates in .arpa DNS in two minutes or less for APNIC zones
- Secure and immediate transactions with immediate feedback
- Compliance with a protocol so it can be implemented in members' automated systems
- Data transfer in well-formed XML
- A fully audited mechanism
- A roadmap for future automated interaction with APNIC IPv6 research

APNIC are currently considering options for providing more IPv6 services to the operator community. "IPv6 uptake appears to be slower than expected, and many people have reported issues of education and access to technology", said APNIC Network Operations Unit Manager, Terry Manderson. APNIC's IPv6 research will consider these issues and seek to address them in future IPv6 implementations.

Implementing ITIL practices

The Information Technology Infrastructure Library (ITIL) is a framework of best practice approaches and management procedures intended to support high quality, cost-effective IT service delivery. The procedures are supplier-independent and have been developed to provide guidance across the breadth of IT infrastructure, development, and operations.

Implementing this system benefits APNIC, and the flow-on effects, such as a formally structured approach to managing IT resources and issues, will also benefit APNIC members.

▲ Network Operations Unit

Member Services update

The APNIC Membership and Helpdesk teams recently united under the banner of Member Services. This unit was formed in order to create a single point of contact for all member enquiries.

Onsite helpdesk at APNIC 24

You can meet some of the Member Services team in person at the Services Lounge at APNIC 24. This will be located onsite at the Intercontinental The Grand Hotel. We encourage you to take the opportunity to meet our friendly staff and and get assistance with any APNIC issues you may have.

The helpdesk can assist with enquiries about:

- Membership
- Billing
- IP/ASN resource application forms
- Status of resource requests
- APNIC Whois Database
- Reverse DNS delegation
- APNIC digital certificates

Speedy service

One of the benefits of the new structure has been an improvement to request turnaround times. Member Services aims to respond to all requests by the next working day, but recently has been able to achieve same-day processing for urgent requests such as digital certificates and reverse DNS delegation.

English not your first language?

APNIC helpdesk staff speak a range of languages in addition to English, including Bengali, Cantonese, Filipino (Tagalog), Hindi, Mandarin, Tamil, Telugu, and Thai. To make an appointment to use our multilingual service, simply email:

helpdesk@apnic.net

Contact us

The Member Services helpdesk features extended operating hours (9:00am - 7:00pm UTC+10), with direct access to APNIC hostmasters and Member Services officers.

In addition to phone, fax, and email, you can also contact the helpdesk using VoIP and online chat. For more information, including VoIP and online chat contact details, please see:

http://www.apnic.net/helpdesk

▲ Member Services Unit

Staff updates

Communications

Jen Anderson
Resource Developer

Jen joined APNIC in July, bringing a wide variety of experience in knowledge management, interface design and IT project management to the Communications Area. She will work closely with staff from all areas of APNIC to enhance member service delivery. She will achieve this by developing various communications facilities, such as APNIC's content management systems, user interfaces for automated services, and intranet features.

Simon Nettle
Editor

Simon joined APNIC in June 2007. He recently moved to Brisbane after living in Japan for three years. Simon has previously worked as a freelance medical and scientific editor. He also has experience in software development and language education. Simon speaks conversational Japanese. Working within the APNIC Communications Area, Simon writes and edits publications across a range of media.

ecoAPNIC initiative

APNIC slashes lighting electricity consumption

The APNIC secretariat office recently replaced 200 fluorescent light tubes with a more advanced type that uses the same amount of electricity but emits a much brighter light. This means that we can now use one fluorescent tube where we previously used two while maintaining adequate lighting levels. Removing the need for a second tube in each lighting fixture has effectively halved the amount of electricity APNIC consumes for lighting.

How you can be eco-friendly at APNIC 24

At recent APNIC meetings, we have reduced the amount of paper products and other materials that we give out. Here are some simple ways you can reduce your ecological footprint at APNIC 24:

Reduce waste:

- Only take handouts if they are useful to you (many just end up in the bin)
- Think about how many copies you need to take back to your colleagues

Reduce energy:

- Turn the thermostat in your hotel room's air-conditioner up a few degrees
- Turn the light off when you leave your hotel room

Training schedule

2007

September

3-4	New Delhi, India, in conjunction with APNIC 24/ SANOG 10
24-28	China
24-28	Sri Lanka
24-28	Maldives

October

2-5	Suva, Fiji, in conjunction with USP
16-19	Bangkok, Thailand, in conjunction with intERLab
21-23	Bangkok, Thailand, in conjunction with intERLab & APTLD
24-27	Vientiane, Laos, in conjunction with NUOL
29-2 Nov	Vientiane, Laos, for NUOL, in conjunction with intERLab

November

3-7	Dhaka, Bangladesh, in conjunction with ISPAB
TBA	Hong Kong
TBA	India
TBA	New Zealand
TBA	Singapore
TBA	Malaysia
TBA	Pakistan

December

TBA	Bhutan
TBA	Vietnam
TBA	Indonesia

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

PacIP6/Pac INET

15-21 August 2007
Pacific Islands
http://www.ipv6forumpacific.org

AP* retreat

26 August 2007
Xi'an, China
http://www.apstar.org

9th APNG Camp

27-29 August 2007
Xi'an, China
http://www.apng.org

24th APAN Meeting

27-31 August 2007
Xi'an, China
http://www.apan.net/meetings/xian2007

APNIC 24/SANOG 10

3-7 September 2007
New Delhi, India
http://www.apnic.net/meetings

AfriNIC-7

24-28 September 2007
Durban, South Africa
http://www.afrinic.net/meeting

ARIN XX

17-19 October 2007
Albuquerque, USA
http://arin.net/ARIN-XX

RIPE 55

22-26 October 2007
Amsterdam, Netherlands
http://ripe.net/ripe/meetings/ripe-55

ICANN Meeting

29 October - 2 November 2007
Los Angels, USA
losangels.icann.org

IGF 2007

12-15 November 2007
Rio de Janeiro, Brazil
igfbrazil2007.br

AUSNOG

15-16 November 2007
Darling Harbour, Sydney, Australia
http://ausnog.net/

Australia IPv6 Summit

20-21 November 2007
Canberra, Australia
http://www.isoc-au.org.au/ipv6summit

70th IETF

2-7 December 2007
Vancouver, Canada
http://www.ietf.org/meetings/meetings.html

NZNOG '08

23-25 January 2008
Dunedin, New Zealand
http://2008.nznog.org/

APNIC 25/APRICOT 2008

20-29 February 2008
Taipei, Taiwan
http://www.apnic.net/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.

Chat
www.apnic.net/helpdesk

VoIP
helpdesk@voip.apnic.net

Email
helpdesk@apnic.net

Phone
+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.