APNIC

Apster26

Apster is the twice-yearly newsletter for APNIC members and the Internet community.

Issue 26 - August 2008

The Internet in New Zealand

As APNIC 26 will be held in the southern city of Christchurch, New Zealand, this article looks at some of the history and development of the Internet in New Zealand. Against a unique backdrop of economic deregulation that governed the New Zealand economy through the late 1980s and early 1990s, the country embarked on an ambitious mission to position the country's broadband infrastructure in the top quarter of the OECD's rankings.

New Zealand is a relatively small island nation with a population barely over 4 million people. Given its geographical isolation in the South Pacific, it is not surprising that New Zealanders have embraced the Internet.

From its early beginnings in 1986, the Internet in New Zealand has become, as in most developed countries, an indispensible and frequently used tool. According to the World Internet Project Interim Report (December 2007), the Internet in New Zealand now rates higher than all traditional media as an information source.

In 1986 Victoria University of Wellington was the first institution to introduce dial-up access to USENET. It wasn't long before New Zealand's first ISP, Actrix, was established in 1989. In 1995 a new public body was formed to help manage Internet infrastructure development with the formation of the Internet Society of New Zealand (ISOCNZ). At this time there were about 2000 domain names in the .nz registry.

In order to create more competition, the Shared Registry System (SRS) was formed in 2000. The SRS provides a single register for domain names and associated technical and administrative information. In 2001, ISOCNZ was rebranded as InternetNZ. By then there were over 100,000 domain names in the .nz register. This increased to well over 300,000 .nz domain names by mid-2007.

A not-for-profit organization that represents a wide membership, including Internet Service Providers (ISPs), web designers, academia, public information groups, and users, InternetNZ serves to promote and protect the Internet in New Zealand and strives to ensure it operates in an open and unfettered environment. As well as providing commentary and advice to various interest groups, InternetNZ also represents New Zealand in global Internet organizations having delegations for the .nz Country Code Top Level Domain (ccTLD), for example.

ICANN (Internet Corporation for Assigned Names and Numbers), the group responsible for the global co-ordination of identifiers such as country codes, is currently chaired by a New Zealander, Peter Dengate-Thrush. Peter is a member of InternetNZ's Public Policy Committee and has been involved with InternetNZ since its formation in 1995, chairing the society from 1999 to 2001.

As part of the delegation of .nz, InternetNZ operates the Domain Name Commission, which oversees the management of the .nz domain name space. It also owns New Zealand Domain Name Registry Ltd, trading as .nz Registry Services (NZRS), which operates the .nz register.

InternetNZ, along with its counterpart in Canada, is one of the few country-code managers where there is said to be "sensible" regulation. According to Keith Davidson, executive director of InternetNZ, "as long as there is open and transparent policy development for the .nz space" and "consensus in the community for our policies", InternetNZ is left alone by the government.

In 2004, the Maori were the first indigenous people in the world to have their own second level domain name, that is, .iwi.nz. InternetNZ is now considering a proposal that second-level domain names be available in both Maori and English languages to reflect the country's bilingual culture. Further, an idea that InternetNZ has put forward to ICANN is the internationalization of domain names, an increasingly urgent issue now that more than half of the Internet's 1.2 billion users do not speak English as a first language. According to Keith Davidson, approximately 80% of the world's population do not speak English, even though they are expected to use English when typing in a URL.

Deregulation

In recent years, Telecom and Clear have dominated the New Zealand telecommunications market.

As a result of poor economic performance during the late 1970s and early 1980s, the New Zealand economy went through a period of economic reform that focused on the removal of protection and the development of competitive markets. As a result, New Zealand was the first member of the

P 2

P 1

OECD to introduce full competition to all sectors of the telecommunications industry.

In 1987, the New Zealand Post Office (which at the time held a statutory monopoly over telecommunications services) was split into three differing areas, one of which was Telecom Corporation of New Zealand Ltd (Telecom). This was privatized in 1990, and in that year the Kiwi Shares Obligations were established. One of the requirements of this contractual agreement required Telecom to offer free residential local calls.

This deregulation and the growing use of telecommunications services resulted in the main network operators entering into interconnection agreements. The five-year interconnection agreement signed between Telecom and Clear in 1996 determined that the operators charged each other a certain sum per minute for terminating calls that originated on the other's network.

In May 1996, Telecom launched its own ISP, Xtra, and in November 1996, Clear launched ClearNet. ClearNet initially focused on large business customers and based its pricing strategy on peak-service demand times. During the period 1996-1999, there was a move away from time billing towards flat-rate Internet services. During this period, the price of Internet access to end-users decreased substantially. For example, user rates decreased from around NZD 150 per month to NZD 30 per month.

In 2000, a few ISPs emerged offering free Internet access, such as Freenet and i4free. The termination revenues received by the competing networks and assigned ISPs under the interconnection contracts encouraged a number of them to offer free Internet services, thereby attracting more customers. However, as soon as the interconnection contracts ended, the termination fees became no longer available. Consequently, most free ISPs could not continue offering free services. After termination of the interconnection agreements, a "bill and keep" arrangement was established whereby neither telco charged the other for calls terminated on its network. It is interesting to note that New Zealand was quite unique in the way free Internet services were offered. Free ISPs in countries like Australia and the USA were based on advertising revenue, which apparently proved to be unprofitable.

The growth of pay ISPs was not affected by these free ISPs. Many Internet users elected to have dual access; that is, existing pay Internet users kept their accounts because the free Internet providers usually only offered limited free Internet access. The advantage of the free ISPs was that they did, however, force down Internet call charges, resulting in the offering of unmetered Internet access packages. This trend is seen in a number of countries.

In 2000, Saturn Communications and Telstra New Zealand formed a 50:50 joint venture intending to invest more than NZD 1 billion over five years to build a broadband network. Telstra then purchased Clear in 2001, aiming to strengthen its competitive position relative to Telecom's. Both telcos intended on investing heavily into a broadband network.

Usage of the Internet in New Zealand

Through the World Internet Project Interim Report (December 2007), it is possible to monitor the developments and trends in Internet usage in New Zealand.

In this report, Internet access, usage, capability, and attitudes were all strongly graded by a New Zealander's age, income, and residential area. The younger, wealthier, and more urban people are, the more connected and confident they were likely to be online.

Of the 1,200 people sampled, 81% used the Internet. About 68% of those with an Internet connection used broadband, whilst 31% used dial-up. The percentage of those using broadband tended to be younger, more urban, and from higher household incomes. Broadband in New Zealand is dominated by ADSL, and in 2006, Telecom launched ADSL2+, which offered users much faster Internet access.

New Zealand has risen one place to 19th in the latest OECD broadband rankings. About 18.3 out of every 100 Kiwis have broadband, with approximately 757,132 broadband subscribers in New Zealand. This has been a significant increase from 8.1%, measured at the end of 2005, primarily because ISPs gained access to Telecom New Zealand's broadband. In 2006 the Telecommunications Act was amended to provide for the operational separation of Telecom New Zealand into retail, wholesale, and network access units, as well as the unbundling of the local telephone loop in order to allow greater competition among local ISPs. This has boosted competition and broadband uptake. New Zealand is now the sixth fastest growing OECD country in terms of broadband penetration, with a net increase of 4.37 subscribers per 100 people.

InternetNZ states, however, that broadband uptake would need to increase dramatically if the government's goal of reaching the OECD's top quarter by 2010 is to be achieved. For New Zealand to reach the level of broadband penetration expected in the OECD top quartile, a minimum of 1.17 million additional services needed to be connected from 2005. This represents a growth of 354% from 2005 levels, or a compound annual growth rate of 35.3%. Further, the standard speed will need to be increased from an estimated 5Mbit/s to 20Mbit/s (Ministry of Economic Development "The Broadband Divide").

The future of the Internet in New Zealand

Both the government (including the opposition) and the big telcos agree that investment in the broadband network is a priority for the future of the Internet in New Zealand. There is cross-party consensus that state investment alongside the private sector in fibre-optic cabling is necessary for economic and social progress.

The government is proposing a Broadband Investment Fund to provide NZD 325 million over five years to telcos, local authorities, and other organizations. Telecom plans to spend NZD 1.4 billion over the next three years deploying fibre-optic cabling closer to customers, which it says will give 80% of homes and businesses faster broadband. However, what consumers are expected to pay for super-fast broadband needs investigating.

The depletion of IPv4 is also of concern, and InternetNZ's current president, Peter Macaulay, suggests New Zealand follow the European Union's (EU) example and move towards IPv6. The EU has set a target for a quarter of EU businesses, public authorities, and households to use next-generation Internet addresses by 2010. Considering the advent of new technologies such as mobile Internet, which sees a large number of consumers accessing the Internet on the go, moves towards businesses implementing IPv6 become increasingly crucial as the free pool of IPv4 addresses depletes. InternetNZ believes the State Services Commission should already be encouraging the government to begin migration towards the new protocol. The Research and Education Advanced Network of New Zealand is already using IPv6. While some old routers and Internet switches will be rendered obsolete when IPv6 becomes more widespread, most modern hardware is already IPv6 compatible.

▲ Christchurch, New Zealand

Message from the Director General

Dear reader,

Welcome to the 26th edition of Apster, the regular newsletter of APNIC, the Asia Pacific Network Information Centre. In this issue you might notice a few changes.

Firstly, in response to reader feedback, we've decided to make Apster a more substantial but less-frequent publication, with two editions per year instead of four. With a longer lead time, we're looking forward to bringing you more in-depth content about topics of interest to the APNIC community.

As you know, APNIC hosts two open policy meetings per year, and a new Apster will now be published for release at each of those meetings. Coincidentally, this 26th edition of Apster is being released in time for the 26th APNIC meeting, which is being held in Christchurch, New Zealand. Likewise, Apster 27 will be released in time for APNIC 27 (in Manila, February 2009), and so on.

For all readers, we hope that Apster will be on interest. Our aim is to include topical news about Internet operations in the region and about global developments that are of interest, whether this be in the area of technical standards, infrastructure developments, applications, or policy.

We also hope that Apster will eventually bring you news that has been contributed by the community itself, by readers such as yourself, or by any of the various Asia Pacific focused organizations that are working around the region. If you have something to contribute, or simply some ideas about how Apster can be better for you, please do get in touch.

I hope you enjoy this 26th edition of Apster.

Paul Wilson

Future APNIC meetings

APNIC's forthcoming meetings will be held at the following locations:

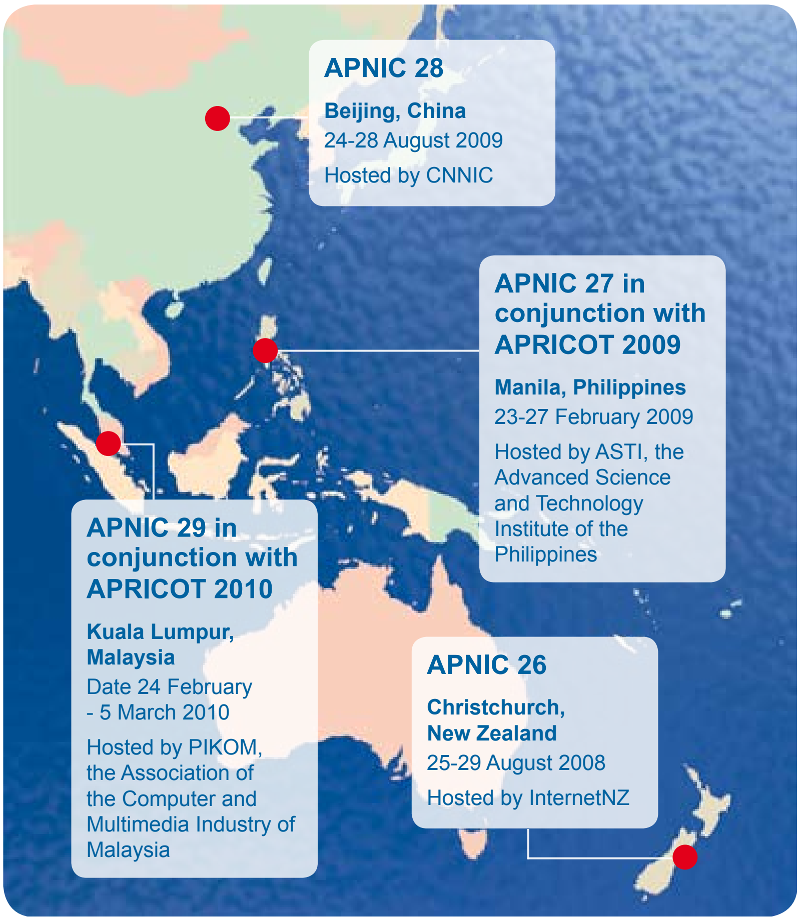

APNIC 28
Beijing, China
24-28 August 2009
Hosted by CNNIC

APNIC 27 in conjunction with APRICOT 2009
Manila, Philippines
23-27 February 2009
Hosted by ASTI, the Advanced Science and Technology Institute of the Philippines

APNIC 29 in conjunction with APRICOT 2010
Kuala Lumpur, Malaysia
Date 24 February - 5 March 2010
Hosted by PIKOM, the Association of the Computer and Multimedia Industry of Malaysia

APNIC 26
Christchurch, New Zealand
25-29 August 2008
Hosted by InternetNZ

index

Policy update

Recent policy implementations

On 4 August 2008, the following policy proposals were implemented:

prop-053 Changing minimum IPv4 allocation size to /22

This change makes it easier for small ISPs to receive direct allocations from APNIC. Under the newly implemented policy, two criteria that ISPs must meet to qualify for a /22 allocation have also been changed. An ISP must:

- Have used a /24 from their upstream provider or demonstrate an immediate need for a /24
- Demonstrate a detailed plan for use of a /23 within a year

prop-054 NIR operational policy document revision

This change updates “Operational policies for National Internet Registries in the APNIC region” to include a reference to IPv6 reverse delegations and aligns reverse DNS methods for NIRs with reverse DNS methods available to other APNIC members.

prop-057 Proposal to change IPv6 initial allocation criteria

This change makes it easier for more ISPs to receive IPv6 allocations from APNIC by enabling current LIRs with IPv4 allocations to receive IPv6 initial allocations without a plan for making 200 assignments. Instead, current LIRs can choose to meet one of the two following criteria:

- Have a plan for making at least 200 assignments to other organizations within two years OR
- Be an existing LIR with IPv4 allocations from an APNIC or an NIR, which will make IPv6 assignments or sub-allocations to other organizations and announce the allocation in the inter-domain routing system within two years

Policy proposals to be discussed at APNIC 26

prop-050 IPv4 address transfers

This is a proposal to remove APNIC policy restrictions on transferring registration of IPv4 addresses between current APNIC account holders. The first version of this proposal was presented at APNIC 24, where the author did not seek consensus. At APNIC 25, the author was asked to continue refining the proposal. The proposal will be presented again at APNIC 26.

Version three of the proposal includes a summary of discussions to date from the APNIC, ARIN, and RIPE communities about liberalizing IPv4 address transfers.

Proposals for address transfers similar to this proposal are currently being discussed at the following RIRs:

RIPE	2007-08: Enabling Methods for Reallocation of IPv4 Resources
ARIN	2008-2: IPv4 Transfer Policy Proposal

prop-055 Global policy for the allocation of the remaining IPv4 address space

This is a proposal for allocating the remaining IPv4 space from IANA to the RIRs. Under this proposal, IANA must reserve one /8 for each RIR as soon as this proposal is adopted. Later, when IANA receives a request for IPv4 address space that cannot be fulfilled using the remaining IANA IPv4 free pool, IANA will allocate each RIR a single /8 from the block reserved for this purpose. Any remaining /8s in IANA’s free pool will then be allocated to the RIR that makes the last request to IANA. The proposal was presented at APNIC 25, where it received majority support but not consensus.

This proposal has reached consensus at the AfriNIC 8 and LACNIC XI meetings and has been adopted by the ARIN Board of Trustees. It is currently under discussion in the RIPE region after having been deemed to reach consensus during the review phase.

prop-059 Using the Resource Public Key Infrastructure to construct validated IRR data

This is a proposal to introduce a new registry that augments Internet Routing Registry (IRR) data with the formally verifiable trust model of the Resource Public Key Infrastructure (RPKI) and provide ISPs with the tools to generate an overlay to the IRR, which is much more strongly trustable.

The proposal has been raised in the following RIRs:

ARIN	“ACSP suggestion 2008.14: Construct Validated IRR Data” has been submitted to the ARIN consultation and suggestion process.
RIPE	“2008:04: Using the Resource Public Key Infrastructure to Construct Validated IRR Data” is currently under discussion.

prop-060 Change in the criteria for the recognition of NIRs in the APNIC region

This is a proposal to update the criteria for recognizing new National Internet Registries (NIRs) in the APNIC region. Proposed changes include removing the need for government endorsement of a proposed NIR, requiring the APNIC membership to vote on whether to accept a proposed NIR, adding requirements for the composition of the proposed NIR's board structure and allowing economies with an APNIC operational office or branch office to be eligible to apply for an NIR. This proposal has not been submitted to any other RIR.

prop-061 32-bit ASNs for documentation purposes

This is a proposal to reserve a block of 32-bit Autonomous System (AS) numbers for the sole purpose of assisting the creation of Internet related documentation. This proposal has not been submitted to any other RIR, but the authors intend the AS number block to be available for use in documentation produced anywhere in the world. A similar proposal asking for an IPv6 documentation prefix was adopted by the community at APNIC 14 in 2002 and was later adopted as RFC 3849 by the IETF.

prop-062 Use of final /8

This proposal describes how APNIC should distribute the final /8 that would be allocated to it by IANA under a successful implementation of prop-055, ‘Global policy for the allocation of the remaining IPv4 address space’. Under this proposal, new and existing LIRs in the APNIC region would be able to receive a single /22 from the last /8 if they meet the current allocation criteria. In addition, it is proposed that a /16 be reserved from the final /8 for distribution for future, as yet unknown, technology requirements.

Proposals to dedicate part of the remaining RIR blocks similar to this proposal have been submitted to the following RIRs:

ARIN	"2008-5: Dedicated IPv4 block to facilitate IPv6 deployment" is currently under discussion.
LACNIC	"LAC-2008-04: Special IPv4 allocations/ assignments reserved for new members" reached consensus at the LACNIC XI meeting in May 2008.

prop-063 Reducing timeframe of IPv4 allocations from twelve to six months

This is a proposal to change the timeframe APNIC uses to make IPv4 allocations to meet LIRs' needs from twelve months to six months to help ensure the fairer distribution of the remaining unallocated free pool. This proposal has not been submitted in any other region.

prop-064 Change to assignment policy for AS numbers

This proposal seeks to create an awareness for the need to support 4-byte AS numbers by introducing another key date in the timeline for APNIC's move to assigning 4-byte AS numbers by default. The proposal suggests that from 1 July 2009, APNIC assigns 4-byte AS numbers by default unless the LIR can document that they cannot successfully announce a 4-byte AS number. This proposal has not been submitted to any other RIR.

prop-065 Format for delegation and recording of 4-byte AS numbers

This proposal recommends that APNIC change the format it uses to represent 4-byte AS numbers from the ASDOT format to the ASPLAIN format. This proposed format change would include AS number representation in all documentation and the APNIC Whois Database. The proposal has not been submitted to any other RIR to date, but may be submitted to other RIRs in future.

prop-066 Ensuring efficient use of historical IPv4 resources

This proposes including all historical address allocations when assessing a network's eligibility for more IPv4 addresses. This proposal has not been submitted to any other RIR. ARIN and LACNIC include historical addresses when assessing eligibility; however, RIPE and AfriNIC do not.

Sam Dickinson

New APNIC policies implemented in August broaden eligibility to APNIC resources

On 4 August 2008, APNIC implemented three policy proposals that reached consensus at APNIC 25 in Taipei, Taiwan in February 2008. The APNIC Executive Council (EC) endorsed the proposals during their May 2008 meeting. The three policy changes are:

- **Changing the minimum IPv4 allocation size to /22**
  This change makes it easier for small ISPs to receive direct allocations from APNIC.

- **Altering the IPv6 initial allocation criteria**
  This change makes it easier for more ISPs to receive IPv6 allocations from APNIC by enabling current LIRs with IPv4 allocations to receive IPv6 initial allocations without a plan for making 200 assignments

- **Revise the reverse zone delegation section of operational policies for National Internet Registries in the APNIC region**
  This change updates the document to include a reference to IPv6 reverse delegations and aligns reverse DNS methods for NIRs with reverse DNS methods available to other APNIC members.

To view the history of the policy proposals that reached consensus, see:

- **prop-053: Changing minimum IPv4 allocation size to /22**

- **prop-057: Proposal to change IPv6 initial allocation criteria**

- **prop-054: NIR operational policy document revision**

Draft document comment period

The draft policy documents that incorporate these policy changes were available for comment during a one-month comment period. All draft policies are subjected to this process, during which time interested parties may:

- **Object to the draft document on the grounds that it does not properly reflect the consensus decision reached in the Policy Review Process**

- **Suggest improvements of any aspect of the document**

- **Request that an additional call for comment be made to allow more consideration of substantial revisions**

What IPv6 address is that?

If you have enabled IPv6 on your computer and have browsed through the interface configuration information for IPv6 addresses you may have been surprised by the fact that there is not just one IPv6 address, but many. With IPv4, there was a single address that was bound to each interface, but with IPv6 it is not so clear, and an interface can have a number of addresses. It is also common to have automatic IPv6 over IPv4 tunnelling interfaces created. The result can be impressive in terms of the number of IPv6 addresses that are configured into a single host system. Here is an example collection of IPv6 addresses:

fe80::20e:7fff:feac:d687

2001:388:1000:4000:217:f2ff:fec9:1b10

2002:cb0a:3cdd:1::1

fc01:3db6:134a:4bb:1:217:f2ff:fec9

::1

What does each of these addresses mean? What IPv6 addresses are useable, and in what context?

The authoritative references for IPv6 addresses are RFC 4291, and the IPv6 address space registry, operated by the Internet Assigned Numbers Authority (IANA):

➡ http://www.iana.org/assignments/ipv6-address-space

Address formats

The major division within the IPv6 address architecture is between unicast and multicast addresses. The IPv6 addresses in the block FF::/8 are multicast addresses, while all other addresses are unicast addresses.

Currently, IPv6 unicast addresses are generically structured as a two-part address: a 64-bit topology part, used by routers to forward a packet to its intended destination network, and a 64-bit interface identifier, that identifies a particular end point. The general structure of an IPv6 unicast address is as follows:

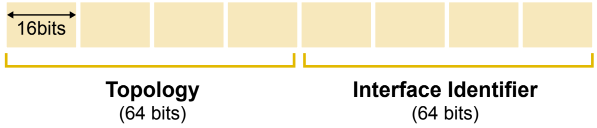

IPv6 multicast addresses are used to forward a packet to a group of end points, each of which is associated with a multicast group identity. The general structure of an IPv6 multicast address is as follows:

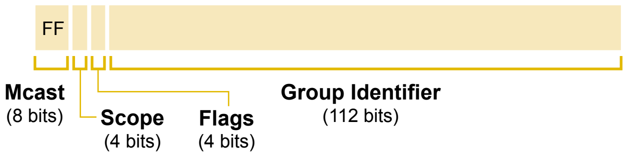

Unicast addresses

::/128 Unspecified

The address of all zeros (::) is the 'unspecified address'. This is never a valid destination address, but may be used as a source address by an initializing host before it has learned of its own address.

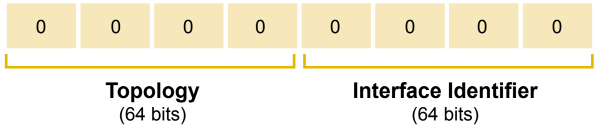

::1/128 Loopback

The address of a single 1 in bit 128 is the loopback address, and is the way for a host to address itself.

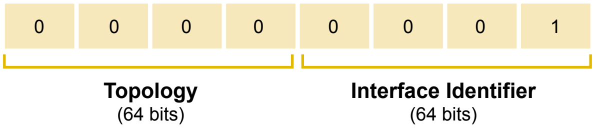

::FFFF/96 IPv4-mapped

IPv4-mapped IPv6 addresses are used to embed IPv4 addresses in an IPv6 address. One use for this is in a dual stack transition scenario where IPv4 addresses can be mapped into an IPv6 address. This is described further in RFC 4038.

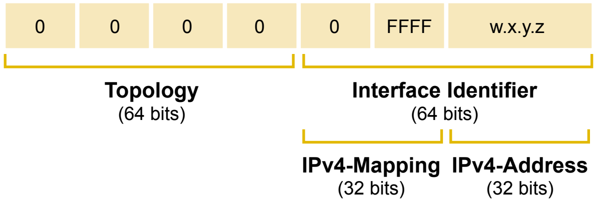

FC00::/7 Unique local addresses

This address block is analogous to the private address space in IPv4. This address space is intended to have a scope that equates to an enterprise environment, as distinct from global public address space. At this stage the address block FD00::/8 is defined, using a self-assigned global ID, where the local bit is set to 1. The global ID is not guaranteed to be unique, and there is no form of address registration. Packets with these addresses in the source or destination fields are not intended to be routed in the public Internet, but are intended to be routed in scoped contexts.

The address prefix FC00::/8 (where the local bit set to 0) is currently undefined.

Unique Local Addresses are described in RFC 4193.

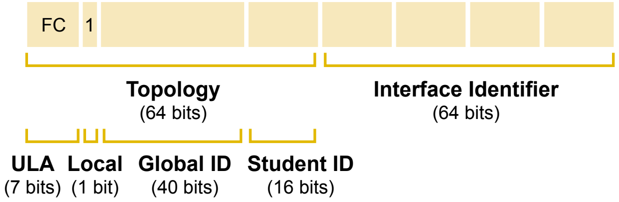

FE80::/10 Link-local addresses

IPv6 uses the concept of scoped addresses, and addresses that use the FE80::/10 prefix are scoped for use as a single link or a non-routed common access network, such as an Ethernet LAN. Their uniqueness is not required in contexts that are broader than this link scope.

While the prefix FE80::/10 is reserved for this purpose, the only defined prefix in this address block is FE80:0:0:0::/64 (that is, the following 54 bits are zero).

Link-local addresses may appear as the source or destination of an IPv6 packet, and are bound to interfaces. Routers must not forward IPv6 packets if the source or destination contains a link-local address.

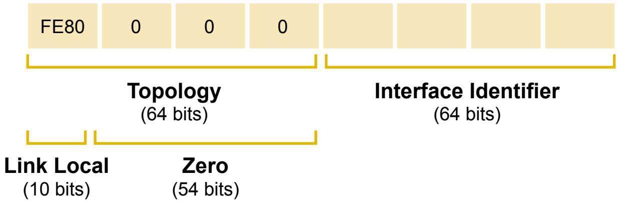

2000::/3 Global unicast

IPv6 global unicast addresses are assigned from the address prefix 2000::/3. These assignments are registered in the IPv6 global unicast address assignment registry:

www.iana.org/assignments/ipv6-unicast-address-assignments

Assigned prefixes are recorded in this registry. All other address prefixes are currently unallocated and should not be seen in the source or destination address of any IPv6 packet in the context of global routing.

All global unicast addresses are allocated via the RIR system with two exceptions:

Address prefixes in the range 2001:0000::/23 are allocated by IANA for special-use cases, as described in RFC 4773. To date, there have been three such assignments:

2001:0000::/32 Teredo
This is a mapped address to allow IPv6 tunnelling through network address tranlators (NATs), as described in RFC 4380. Packets with Teredo address prefixes in the source of destination field of the packet may be encountered in scoped or global routing contexts.

2001:0200::/48 Benchmarking
This is a non-routable address prefix to be used in the context of benchmarking tests (see RFC 5180 for details). This is not a routable address prefix, and packets with these addresses in the source or destination fields should not be seen on local scoped or global networks.

2001:0010::/28 Orchid
This is a fixed term experimental address allocation intended to support routable cryptographic hash identifiers. These identifiers are intended to be visible only on an end-to-end basis, and routers should not see packets with these addresses in the source or destination address fields of any packet, in either local or globally scoped contexts.

In addition to these special purpose allocations, there is a mapped address prefix used to support auto-tunnelling of IPv6 packets over IPv4 in non-NAT contexts.

2002::/16 6to4
A 6to4 gateway adds its IPv4 address to this 6to4 prefix, forming the 48-bit address prefix: 2002:w.x.y.z::/48. This prefix is used as the common prefix for all IPv6 client hosts that are serviced in IPv6 from this 6to4 gateway. This is described in RFC 3056.

2001:db8::/32 Documentation
This is the reserved documentation prefix. Packets should not carry addresses with this prefix in either the source or destination fields in the IPv6 packet header.

All others: reserved

All other addresses in the unicast range are reserved by the IETF for future definition. IPv6 packets should now not use addresses from this range in either the source or destination fields in private, scoped, or global contexts.

Multicast addresses

FF00::/8

This is the IPv6 multicast address prefix which is used to identify multicast groups. An interface may belong to a number of multicast groups. Multicast addresses must not be used as a source address in an IPv6 packet, only as a destination.

The following four items of this multicast address are flags, defined as follows:

0	must be zero
R	defined in RFC 3956 - embedded rendezvous point
P	defined in RFC 3306 - unicast prefix based multicast address
T	defined RFC 4291 - 0 = IANA assigned, 1 = transient

The next four items are a scope identifier, defined as follows:

0	reserved	**1**	Interface-local scope
2	Link-local scope	**3**	reserved
4	Admin-local scope	**5**	Site-local scope
6	(unassigned)	**7**	(unassigned)
8	Organization-local scope	**9**	(unassigned)
A	(unassigned)	**B**	(unassigned)
C	(unassigned)	**D**	(unassigned)
E	Global scope	**F**	reserved

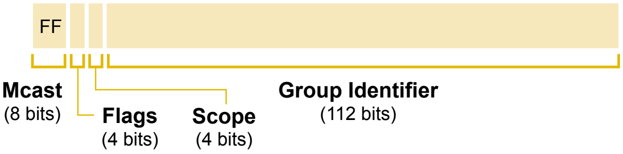

Deprecated IPv6 addresses

There are a number of address prefixes that were assigned in the past, but these assignments have been deprecated for one reason or another, and the addresses have reverted back to their original designation.

::w.x.y.z/128 IPv4 compatible IPv6 address

This form of mapped address, where an IPv4 address was padded out to 128 bits in length by the addition of leading zeros was originally thought to be a useful technique in the IPv4 to IPv6 transition. No technique has been devised to use this form of mapped IPv4 address, and the assignment has been deprecated.

0200::/7 NSAP

The intention of this prefix was to map certain OSI Network Service Access Point addresses into IPv6 addresses, as defined in RFC 1988. This approach was subsequently deprecated (RFC 4048) and the assignment of this address prefix was cancelled.

3FFE::/16 6bone

This prefix was used in the second operational phase of the IPv6 test network, the 6BPNE. This experimental network allocation was described in RFC 2471. It was shut down as of 6 June 2006 (RFC 3701).

5F00::/8 6bone

This address prefix was used in the first operational phase of the IPv6 test network, as described in RFC 1897. The assignment was deprecated with the second phase of the 6bone, RFC 2471.

FEC0::/10 Site local

In earlier versions of the IPv6 address plan, this prefix was intended for use in a scoped context that corresponded to a 'site' analogous to the IPv4 private-use address prefixes described in RFC 1918. This address assignment was deprecated by RFC 3879, and the function of scoped addresses has been replaced by unique local address prefixes, described in RFC 4193.

IPv6 registry lookup

In order to determine the end user of an IPv6 address, it is necessary to determine which registry holds the assignment information.

The registry that contains the IPv6 address plan is:

IPv6 address registry	IANA	http://www.iana.org/assignments/ipv6-address-space

P 8

P7

The unicast address assignments are contained in the global unicast registry:

Global unicast registry	IANA	http://www.iana.org/assignments/ipv6-unicast-address-assignments

The local and mapped IPv6 addresses are mapped as follows:

1::/128	local host
::FFFF:w.x.y.z	IPv4 mapped address - lookup IPv4 address w.x.y.z
FD00::/8	Unique Local Addresses. No registry is used for these addresses.
2001:0::/32	**Teredo address** • Bits 32-64 of the address contain the Teredo server address • Bits 76-91 contain the external IPv4 port address, XORed with 1s • Bits 92-128 contain the external IPv4 address XORed with 1s **For example, the Teredo IPv6 address:** 2001:0:cf2e:308c:0:323d:3fa1:c0b0 ...can be mapped as follows • Teredo server IPv4 address is cf.2e.30.8c = 207.46.48.140 • External obscured port of client is 323d = port 52674 • External obscured IP address of client is 3f.a1.c0.cb = 192.94.63.70
2002::/16	**6to4 address** Bits 16-48 contain the IPv4 address of the 6to4 gateway **For example, the 6to4 address** 2002:cb0a:3cdd:1::1 ...has the IPv4 gateway address of cb.0a.3c.dd = 203.10.60.221

Other addresses are allocated through the RIR system and the corresponding registry address location for the end user is of the form:

whois -h <whois *Registry*> <*ipv6 address*>

For example:

whois -h whois.apnic.net 2001:dc0:2001:10:20e:7fff:feac:d687

The following table lists the IPv6 address prefix ranges and the corresponding whois registry where its allocation details can be found.

Start prefix	End prefix	Whois registry
2001:0200	2001:03FF	whois.apnic.net
2001:0400	2001:05FF	whois.arin.net
2001:0600	2001:0BFF	whois.ripe.net
2001:0C00	2001:0FFF	whois.apnic.net
2001:1200	2001:13FF	whois.lacnic.net
2001:1400	2001:3BFF	whois.ripe.net
2001:4000	2001:5FFF	whois.ripe.net
2001:8000	2001:BFFF	whois.apnic.net
2003:0000	2003:3FFF	whois.ripe.net
2400:0000	24FF:FFFF	whois.apnic.net
2600:0000	26FF:FFFF	whois.arin.net
2610:0000	2610:01FF	whois.arin.net
2620:0000	2620:01FF	whois.arin.net
2800:0000	28FF:FFFF	whois.lacnic.net
2A00:0000	2AFF:FFFF	whois.ripe.net
2C00:0000	2CFF:FFFF	whois.afrinic.net

Geoff Huston

Is your IPv4 address really private?

The current predictions of IPv4 address consumption foresee the IANA unallocated address pool being completely depleted by January 2011 (http://ipv4.potaroo.net). This model also predicts that the RIRs will consume their available pools of IPv4 addresses some months thereafter. The implication is that we are going to see IPv4 addresses deployed in the public Internet that have not been used in this context before, and this could present some problems to private, corporate, or even home networks.

In the early 1990s, it was recognized that the pool of IPv4 addresses was considerably smaller than the expectations of the use of the Internet in the coming years. The Internet Engineering Task Force (IETF) initiated a number of efforts to address this problem. Some were short-term tactical measures intended to provide a 'once-off' benefit, others were medium-term approaches, and yet others were phrased to meet longer-term objectives.

The long-term measure was the specification of IPv6; which, even at the time, was recognized to be a major undertaking involving protracted periods in both specification and transition within the deployed Internet. The short and medium-term efforts included the removal of the classful structure of IPv4 addresses, allowing address deployment to be more precisely tailored to meet the exact requirement of each network, and as a result eliminating much of the inefficiency of address use. Also, the IETF specified a pool of addresses that were intended for private use. This specification, originally documented as RFC 1597 in March 1994 and subsequently in RFC 1918 in 1996, defined three address ranges that could be used in private contexts.

The addresses, the so-called 'RFC 1918 address blocks', are as follows:

```
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)
```

These addresses are not used in the public Internet, so any private network, whether it's a home or corporate enterprise network, can use these addresses with the knowledge that even if they connect their network to the public Internet via a NAT, these internal private addresses will not clash with any public address. These are not the same as the addresses used in IPv4 to identify oneself (0.0.0.0/8), or one's loopback interface (127.0.0.0/8). These RFC 1918 private-use addresses are intended to address devices whose connectivity is explicitly bounded and does not include direct visibility to the public Internet.

However, not all private-use contexts use addresses drawn solely from this RFC 1918 private number range. Some network equipment manufacturers and some Internet service providers used IPv4 addresses as private addresses, but chose to use other IPv4 addresses than those specified in RFC 1918. These addresses were unassigned addresses, or addresses that were registered in the IANA IPv4 address registry, but it was assumed that the current holder of the assigned address would never use the addresses in the context of the public Internet.

Some use-contexts extend back over decades, and their use of a particular address block is a legacy issue. It is also possible that some equipment vendors or software engineers were unaware of the details of the IPv4 address management framework and assumed that the term 'reserved' in the IANA IPv4 address registry meant 'reserved for all time', as distinct from 'not allocated at the moment, but will be allocated in the future.' Others used unallocated space because they had run out of RFC 1918 private use addresses and needed more, and were forced to use unallocated addresses, even with the knowledge that the address would be allocated for public use at a later date.

Irrespective of the reason as to why there is use of unallocated address space in private networks, the problem still remains that in about three years it is anticipated that all the remaining unallocated IPv4 address space will be allocated to end-users and will appear as routed address space in the public Internet.

For example, a common WiFi user authentication service uses addresses 1.1.1.1 and 2.2.2.2 to perform user authentication and connect the user to a NAT that is connected to the public Internet. The local user has a locally-assigned network address drawn from 1.0.0.0/8. As long as network 1.0.0.0/8 is unallocated, this approach works. But when IANA allocates 1.0.0.0/8 to an RIR, and the RIR allocates addresses to service providers from this block, problems will surface. The local users in this WiFi network will not be able to reach any public services that are addressed using addresses drawn from 1.0.0.0/8.

Well, so what? Interestingly, a significant proportion of the most popular services on the Internet today did not exist three years ago. In other words, these popular services have a tendency to use recently allocated IP addresses. The probability that a popular service will be addressed from network 1.0.0.0/8 or 2.0.0.0/8 is, therefore, extremely high. At this point, the decision to use 1.0.0.0/8 in a private network context becomes an extremely poor one. It's not that this private use of unallocated addresses will prevent anyone else from accessing services and using the Internet; the damage will occur in the context of the private network, where the ability to access public services that are numbered from the same address range as the local private network will be affected.

Leo Vegoda of IANA has undertaken a study of the use of unallocated IPv4 addresses in private contexts, and his results were written up in the Internet Protocol Journal in September 2007 (http://www.cisco.com/web/about/ac123/ac147/archived_issues/ipj_10-3/103_awkward.html) and in presentations at NANOG 41 and NANOG 42. Commonly used networks in this category include 1.0.0.0/8 (reported as widely used), 5.0.0.0/8 (reported as used by the Hamachi VPN service) and 42.0.0.0/8 (reported as used by the HP Procurve 700w appliance). There are also informal reports of the use of 7.0.0.0/8, 14.0.0.0/8 and 104.0.0.0/8 in private network contexts as well.

As Leo reports:

> Organizations using these address ranges in products or services may experience problems when more specific Internet routes attract traffic that was meant for internal hosts, or alternatively find themselves unable to reach the legitimate users of those addresses because those addresses are being used internally. The users of unregistered networks may also find problems with reverse Domain Name System (DNS) resolution, depending on how their DNS servers are configured. These problems are likely to result in additional calls to helpdesks and security desks at both enterprises and ISPs, with unexpected behaviour for end users that might be hard to diagnose. Users of unregistered address space may also experience problems with unexpected traffic being received at their site if they leak internal routes to the public Internet. Many ISPs have already had experience with this type of routing inconsistency as recent /8 allocations reach routing tables and bogon filters are updated.

So, if your local IP address is 7.1.20.1, or your ISP has used DHCP to provide your local NAT box with 14.1.0.20 so that you can use network 10.0.0.0/8 at home as a local private network without clashing with your ISP, then you can confidently expect to encounter connectivity issues in the near future.

In such situations, there is no choice but to renumber. One choice is to renumber into RFC 1918 private use address space. Another choice is to evaluate whether you are eligible for an allocation of unique public IPv4 address space from your service provider,

P 10

DNSSEC – Once more, with feeling!

Domain Name System Security Extensions (DNSSEC) have been in development for some ten years now, which makes it one of the more respectable Internet technologies, despite the fact that it has still not been effectively deployed! Why not? Is there something we can do now to get this DNSSEC ball rolling?

Background

DNSSEC is a way to add some 'security' to the DNS. The underlying motivation is that the DNS is a rather obvious gaping hole in the overall security of the Internet. With DNS as it stands today, attacks on the integrity of the DNS are virtually impossible to detect, particularly attacks by simple man-in-the-middle substitutions.

DNSSEC is a digital signature framework that allows a DNS client to check the integrity and completeness of a response to a DNS query. The approach used by DNSSEC is a relatively conventional application of public key cryptography, where DNS Resource Records (RRs) are digitally signed.

This digital signature can be added as additional information to a DNS response, and DNSSEC-aware resolvers can request that additional DNSSEC validation material be provided in addition to any DNS query. The resolver may then use this validation material to verify that the DNS response is genuine and has not been altered in any way whatsoever.

But, while the signed data may be valid, was the key used to sign the data a valid key? DNSSEC embeds a key hierarchy within the DNS itself. Each zone parent is given the role of signing over every delegated child's zone key, so that a zone's signing key can be verified by confirming the parent zone's signature over that key; which, in turn, can be verified by that zone parent's signature over this key, and so on until you reach a Trust Anchor or the root of the DNS, or, preferably, both at once.

All this results in the observation that the theoretical trust model of DNSSEC can be reduced to a reliance on trust in only one key, that key being at the top of this delegation hierarchy. In other words, DNSSEC was originally designed with a single trust anchor in mind, that being the public key used to sign the root zone of the DNS, and the deployment model was one of universal adoption across the entire public DNS.

So where are we on this DNSSEC deployment agenda? Is it within reach? Is it a bit of a stretch, but still plausible? Or maybe, it's so far out there that a manned mission to Pluto will happen first!

Root signing

In the case of DNSSEC, the task of generating a public/private key pair, signing the root zone with the private key, and publishing the public key as the material that "anchors" the DNSSEC signing hierarchy is an example of a task that seems simple, but actually poses an array of daunting problems.

The root zone of the DNS is quite small, and generating a key pair and creating a digital signature of the root zone is not actually the problem. Indeed, it's a straightforward task, and existing DNSSEC tools can achieve this outcome quite easily. An example of what such a signed zone might look like can be found at:

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

If the mechanics of actually signing the root are so trivial, and we know what the outcome will look like, then why hasn't it happened already on the actual DNS root? What's the problem here?

The general technical consideration when contemplating changes to the root zone of the DNS is whether the basic UDP response to a root zone-priming query will be able to be transmitted back to the query party. When a priming query is passed towards a root server with the DNSSEC OK bit set in the EDNS header of the DNS, the query should also use the EDNS 'sender's UDP payload' mechanism, because the inclusion of the DNSSEC information in the priming response will increase the size of the responding DNS packet to more than 576 bytes.

As long as everything between the DNSSEC and the EDNS-aware DNS entity performing the query is capable of doing the right thing with large UDP packets, then everything will work. And, if the packet is mangled by some reprehensible middleware that mangles DNS UDP packets in bizarre ways, then the query has to drop back out of DNSSEC and perform a non-DNSSEC priming query. It seems that the operational modes do not cause complete breakage here. So, as long as the key is correctly handled and the signatures are well formed, then from a technical perspective signing the root does not appear to present risks to the operation of the DNS.

Obviously, it's not a technical problem. So, if its not technical, then it must be yet another case of those endless political debates about the DNS, IANA, ICANN, and the roles and desires of the multitude of interested onlookers. The consideration of who has effective control of the private key of the DNS root zone appears to be the sticking point here.

While the idea that IANA should control the zone keys for the DNS root may appear obvious to many; for others, that response is politically naive. Somehow, the mechanical process of signing parts of the DNS root appears to have been confounded with the political process of the endless debate about the process of making changes to the DNS root.

The key that would allow us to bootstrap the DNSSEC validation hierarchy in a simple manner is now regarded as the key to defining which parties need to provide their permission before allowing any changes whatsoever to the DNS root zone itself. However, if the political process of DNS determination remains in the public debate over the governance of the DNS, the prospects of seeing a signed root zone in the DNS appear to be highly unlikely.

P 9 ...continued from page 9

your local Internet registry, or from the relevant regional Internet registry. You'll need to check with them to see if you are eligible for such an address allocation within their policy framework. Or, if you are renumbering your network anyway, you may want to consider if using IPv6, possibly with NAT-protocol translator interfaces, as being a step towards future-proofing your private network. The IPv6 address allocation policy framework is often somewhat different to IPv4, so while your private use network may not meet the IPv4 address allocation criteria, the same network may meet the IPv6 address allocation criteria. Or, you could consider the use of IPv6 unique local address space, which is analogous to private use address space in IPv6.

So, if you are using unallocated IPv4 addresses because they were not being used in the public Internet at the time, then you are strongly urged to check and renumber your network, as there is a very high probability that you will run into Internet connectivity problems in the very near future.

Geoff Huston

Signing everything else

There is one positive aspect of this longstanding political stalemate over signing the root zone of the DNS; namely, that the other rather improbable aspect of DNSSEC deployment can be conveniently ignored.

DNSSEC relies on the DNS structure itself to create the interlocking relationships that remove the need for additional verification mechanisms. For a DNS Resource Record (RR) to be verified by a DNSSEC-aware resolver, it is necessary to verify both the digital signature of the RR and confirm the validity of the signing key that generated the RRSIG (RR signature) resource record.

In DNSSEC, the parent signs across the child's public key, so the child's key can be verified by verifying the parent's digital signature of this key. This signature can be verified by confirming the digital signature and verifying the parent's public key. This key is verified by its parent, and so on. If the root zone key is the trust anchor of this verification chain, then every zone between the root zone and the zone of the DNS RR being verified must be DNSSEC signed.

Just how likely is universal DNSSEC deployment? It appears that the marginal additional overhead of adding DNSSEC to a zone is all borne by the zone administrator and the zone publisher. This involves additional zone administrative procedures, the introduction of DNSSEC key management functions, larger DNS zone files, interactions with delegated zone administrators and parent zone administrators, and the possibility of more potential points of service failure. The direct benefits, however, accrue to the DNS client.

In those situations where a large number of people bear the costs, and a smaller number of people stand to gain the benefit, and there is no compensatory flow of money to even the scorecard, then the entire proposition is not generally considered overly attractive from a business perspective. Worse still is the case where the costs far exceed the value of the benefit. Signing everything in the DNS just doesn't seem to be economically rational from the point of view of the DNS zone administrator.

Universal DNSSEC deployment, even with a signed root, has prospects that are even more dismal than the signed root proposition.

The great key hunt

So, we haven't yet managed to sign the root, and we haven't yet managed to achieve universal buy-in below the root, and the prospects for achieving both of these objectives look dim at the moment.

Are we ready to give up on DNSSEC yet? Of course not!

The design of secure systems often involves a set of design trade-offs – a set of compromises between security, scalability, and feasibility. If the total cost of deployment and the resultant benefit are not well aligned, can DNSSEC take some convenient shortcuts to ease this imbalance, even at the expense of the integrity of the security model? Is partial DNSSEC deployment possible?

Partial deployment of DNSSEC implies that only some zones are signed. A signed zone might have no DNSSEC parent, but may have DNSSEC children. These parentless DNSSEC 'orphans' become the apex of a local DNSSEC hierarchy. For a DNSSEC resolver to be able to verify as much as it possibly can, it has to load up all the zone keys that are at the apex of a local DNSSEC hierarchy. Unfortunately, there is no automated way to perform such a sweep of the DNS to expose these DNSSEC zones, so the process appears to be one that you perform by hand. Even then, you need to derive some form of trust in the authenticity of the key in this undefined process of DNSSEC zone key retrieval and maintenance.

Another way for the resolver to gather these zone keys up is to pick them up one-by-one on an as-required basis using the DNS itself. This gets rid of the entire process of trying to sweep the DNS for these DNSSEC local hierarchies just in case you might want to pose a query against the zones that you wanted to validate with DNSSEC. But, the entire DNSSEC effort is directed at providing some means of verifying DNS responses. If you use the DNS itself to deliver these keys that you are going to use as the basis of your trusted verification, then you have just introduced a fatal circularity of dependence. How can you trust the DNS to deliver you the correct key that you are then going to trust to validate DNS responses?

So, we are back once more to the basic problem of DNS integrity that DNSSEC was intended to solve. How can you pick up these local DNSSEC apex keys in a reliable and trusted manner without resorting to the DNS? Unfortunately, DNSSEC has no direct answer here. How do you know that a DNS zone has been legitimately signed and that the DNS response can be validated? In an environment of partial DNSSEC deployment, DNSSEC does not appear to be overly helpful.

Another response has been to contemplate DNSSEC Lookaside Vectors (DLVs), which attempt to aggregate a number of apex zone keys of local DNSSEC zone hierarchies into a synthesized DNSSEC zone that allows a single DLV zone key to substitute for the set of stored apex zone keys. If you change the DNSSEC-aware resolver to follow the Lookaside Vector, you can replace the collection of local apex keys with the single key of the Lookaside Vector zone, as long as all these DNSSEC orphans are prepared to live as the DNSSEC Lookaside orphanage. But, in this DLV model we've broken out of the interlocking DNS delegation model, and the question then arises as to the authenticity of the zone keys stored in this DLV zone.

While DLV is technically feasible, the validity of the outcome is no longer based on the elegant structure of interlocking DNSSEC keys that are precisely aligned to DNS zone delegation. Instead, the strength of the outcome is no stronger than the integrity and accuracy of the admission procedures that are undertaken by the DLV operator. It's hard to see how the DLV approach offers anything more secure than the current haphazard collection of DNS name certificates and the haphazard collection of certificate authorities who issue certificates for DNS names, while at the same time sit outside the name delegation hierarchy.

So, the alternatives for the DNSSEC client in a world of partial DNSSEC deployment is to locate and maintain a set of trust keys using undefined processes that presumably involve a considerable amount of human direction, or to outsource the problem to a trustworthy and reliable DLV operator.

Neither alternative sounds overly attractive if you are after true security rather than just the appearance of it.

The DNSSEC burden

DNSSEC is not free. The DNSSEC burden is spread across the zone administration, primary zone servers, secondary servers and resolvers, and the end clients of the DNS. In the late 1980s, the DNS traffic on the NSFNET contributed 20% of all packets on the network of the day. Things have improved considerably since then, and it's not anticipated that DNSSEC deployment is going to break either the DNS or the Internet at large; but, even so, DNSSEC is not free.

For the zone administrator, there're the additional tasks of key management, record signing, and managing zone updates. For the primary and secondary zone servers, there's the need to support incremental updates to the zone, the advisability of using a trusted channel for zone updates from the primary server to the secondaries, such as TSIG, the issues of larger zone sizes, the larger response sizes, and the corresponding increments in processor, memory, and bandwidth requirements for the servers. This, in turn, triggers reconfiguration of platform and network capacity for the server side of the DNS.

P 12

P 11

For resolvers, there is the issue of the larger response sizes, the need to ensure that resolvers and the local network infrastructure of firewalls, NAT boxes, and other inventive pieces of middleware can cope with EDNS0 and larger DNS response packets that are potentially fragmented, the additional query overhead associated with validation of responses, and that tricky question of what DNSSEC outcomes to cache and for how long.

One major consideration here is that the DNS is already relatively fragile and is the constant target of attack. One way to disrupt a service is to subject its name servers to a constant very high intensity load. In amongst all this noise traffic, genuine queries are lost and the servers effectively disappear off the network. For providers, this is a balancing act. While the DNS is essentially unfunded, if you are out of action, then the outage is highly visible. Given that the attack has been one of loading the DNS server with correctly formatted queries, there is no visible clue to the server as to which queries are genuine and which constitute the attack.

The common mitigation to this attack is firstly one of segmentation of the server domain through anycast of the DNS servers, and an effort to increase the capacity of the server to absorb the query load associated with an attack without falling over. In this environment, we are now seeing implementations of customized DNS servers and recursive servers that are engineered with a very high capacity. Here, DNSSEC applies more pressure through larger responses. A DNS server that is engineered for resiliency in a non-DNSSEC world may not necessarily be able to manage the increased load. From the perspective of a large, heavily-used DNS zone, DNSSEC necessitates reinvestment in server infrastructure as a consequence of DNSSEC signing the zone. This is not exactly delivering on the generic promise of all this technology as being better, faster, and cheaper.

What is the meaning of ‘failure'?

What should a resolver do in the event of a DNSSEC verification failure?

If the DNS response is a Resource Record (RR) and the associated DNSEC RR signature fails to verify, then what should the DNS resolver present to its client? Should the resolver synthesize an NXDOMAIN response in place of the suspect RR response on the basis that passing on a response that has failed verification is probably worse than a 'no such domain' response? What if the NSEC (or NSEC3) response has a DNSSEC RR signature that cannot be verified? Here the 'server failure' response is still relatively unhelpful, but a more assertive NXDOMAIN, or 'no such domain' response, is possibly an incorrect response; however, there is no better information at hand to substitute in its place.

But, what is failure anyway? And, more particularly, what is failure in a partially-deployed DNSSEC world? A digital signature that cannot be verified is a clear failure. The lack of an upward chain of parent signatures that leads to a trust anchor is also a failure, but of a different type. It could well be that in this case the DNS RR provided as the answer to the original query is perfectly good in every respect, including the DNSSEC RR signature, and it’s the resolver’s efforts to hunt down all of the apex zone keys for each and every local DNSSEC hierarchy. Or, the resolver could be the victim of a DNS attack. How can a resolver tell the difference? Is this a benign failure, or an instance of a failure that points to a directed attack via the DNS.

Of course, undermining all this is human behaviour. When the screen presents you with the message: “I cannot validate this certificate, do you want to proceed anyway: Yes or No?” we are all pretty much the same when we move the mouse to hover over that “yes” button. DNSSEC might be as paranoid as it wants, but we humans are all risk takers of one sort or another, and even when presented with a dire warning of the risks involved, there’s a certain air of digital bravado that creeps up upon your clicking finger when given the opportunity to take a risk! So all this security infrastructure can be undermined by an all too typical user behaviour mode.

So, tell me again…

Much of the public consideration of DNSSEC has been on the topic of the number and composition of the group of people whose fingers hold the pen that signs the root zone of the DNS. The politicization of this question of “who signs the root?” appears to ensure that the root remains unsigned, and as a consequence any hope of universal deployment of DNSSEC flies out the window. It appears that without a signed root, partial deployment of DNSSEC is the best one can hope for, and partial deployment of DNSSEC is, on the whole, an example of negative progress.

Without a signed root, DNS resolvers need to do the impossible and perform minor miracles in terms of trust key discovery and management. The impossible happens only with considerable pain and effort, and even then it happens rarely. The question is then raised as to why zone administrators and DNS operators should incur additional costs to locally deploy DNSSEC given that there are a very small number of DNSSEC queries to be answered from these few bold and adventurous DNSSEC-aware resolvers? From this perspective, it appears to be counter-productive to be fixated on the issue of how many potential signatories should be dancing on the head of this root zone pin.

Clearly, the answer to signing the DNS root is that with all zones, it is the role of the zone administrator to generate the keys and sign the zone file. In the case of the root zone of the DNS, it’s back to IANA to just do the job and let the rest of us move on.

But 'moving on' is not the same as solving all the issues of Internet insecurity though the scribbling of just one signature over a block of bits with a digital pen. It is apparent that the lack of a signed DNS root has become a convenient way to paper over the other issue of the lack of DNSSEC deployment across the remainder of the DNS and the poor prospects of ever changing that situation. It’s still a very hard problem to work out how to swing the balance around in the DNSSEC cost and benefit equation so that the benefits of deploying DNSSEC on the server side can motivate its deployment such that the client-side benefits of a more robust DNS can be realized.

And, even if we address that and move into a DNSSEC world, there is the observation that most of the technology failures we encounter in this area are outcomes of benign rather than hostile actions, and that failures are all too often a product of some failing of operational competence rather than intended malice, and this is no doubt going to continue. We have become habituated to seeing 'failure' as an accidental and unintended outcome, rather than a dire warning of potential danger. We have grown so trusting in the technology that when the technology offers us choices, then we believe that it really is a choice between equally viable options and that both options are equally 'safe' from the point of view of the underlying technical machinery. Otherwise, why would this machine be asking us for guidance in the first place?

The substantive issues for DNSSEC are much further down in the DNS hierarchy than at the root, but we’re never even going to have the opportunity to address them as anything more than hypothetical issues to be considered in the abstract for as long as the unsigned DNS root remains in our way.

Geoff Huston

Reflections on the OECD ministerial meeting on the future of the Internet economy

In June the Organization for Economic Cooperation and Development (OECD) hosted a Ministerial Meeting on the Future of the Internet Economy. It was attended by Ministers for communications from the 30 OECD member nations and some 15 other nations, who gathered to talk about the future of the Internet Economy.

The meeting was relatively unique in its class for a number of reasons. The OECD is a widely referenced, well-respected source of objective economic data and comparative studies of national economies. The issues discussed were more aligned to Internet public policies than limited to an interpretation of economic activity. Also, such OECD activities in the past have been instrumental in facilitating changes in governmental approaches to common issues that have broad economic and social dimensions.

It was interesting to note that the level of knowledge and insight into the current issues that face the Internet was generally very high. The discussion at the meeting was well-informed and thoughtfully focused. The various misconceptions that are all too often a part of such governmental discussions about the Internet were, on the whole, completely absent.

The Ministerial meeting was careful to wrap up the message of the Internet’s future in the catchy tag line – "Convergence, Creativity, and Confidence." This encompasses themes of technology evolution and the increasing reliance on IP as the universal substrate, the continual process of innovation in the range and scope of services, and the issues of security and integrity required from the network and its services.

On a deeper level than the choice of catchy theme phrase, this was a meeting that has made some important strides in the political landscape of the Internet. The two years of effort in preparing for this Ministerial meeting appears to have produced some very interesting outcomes. We've come a long way in the last decade when one can now see governments rejecting efforts to shoehorn the Internet back into the constrained box of regulatory initiatives as phrased though the inter-governmental treaty organizations, such as the ITU.

Signed by 39 national governments and the European Union “The Seoul Declaration for the Future of the Internet Economy”, which arose from this meeting, recognizes that the Internet is as much about robust economic well-being, cultural diversity, and social interaction as it is about the intricate technical task of shoveling large piles of bits into data pipes. It recognizes the engine that propels the Internet is not regulatory in nature. Instead, it is an engine fuelled by the cooperative open participation of many interests and communities. These are ignited by the deregulation of the communications sector, the introduction of intense competition at all levels of the network service environment, and the surge in innovation in user-visible services.

The Declaration recognizes that national economic performance is tightly bound to the prospects of the Internet itself. Further, it is stating that the overall cooperation and coordination effort required to support the Internet is far too broad and critical a task to leave to organizational structures that are bound by inter-governmental treaties, and by inference, to simply pass over the entire matter into the hands of the ITU-T. The task is far more than a conventional exercise in communications technology regulation. Its future will necessarily involve not only governments, but all forms of enterprise and individual actions as well.

Taking a more decentralized approach to policy formulation fits comfortably with the open, transparent, participatory processes that are the cornerstones of the self-regulatory policy development framework used by the Regional Internet Registries and by ICANN. The commitment by governmental signatories to this Declaration to work with business, civil society, and the technical community on maintaining a policy framework for the Internet that promotes competition, empowers and protects consumers, and expands Internet access and use worldwide would be nothing particularly novel, were it not for the explicit commitment to cooperate with these other stakeholders, and were it not for the particular emphasis on open competition, as distinct from the traditional recourse to imposed regulatory fiat.

The good news is that the various regulatory regimes have been successfully pressured to take a further step along the deregulatory path by recognizing there are other communities of interest who have a legitimate voice in the overall framework of "governance". It appears that the closed door position used by many governments when representing national interests at international levels is changing, and this ministerial declaration offers some evidence of this progress.

The major positive aspect of this Seoul Declaration is the recognition that in a deregulated, diverse activity sector serving a public communications utility there are many interests that sit alongside those of national governments, and this is indeed a welcome recognition.

Not everything about the Internet is solved, of course. There are important social, technical, and economic issues that require attention in the coming years. If the flowery rhetoric about the rosy future of a dramatically larger, more diverse, more ubiquitous Internet is to ever come even close to a reality, then such issues will demand some form of resolution. If the OECD fulfils its intention to meet again to evaluate progress, then the effectiveness of the Seoul Declaration should be gauged by the extent to which these business and economic issues have been addressed between now and then.

In any case, the Seoul Declaration makes one thing pretty clear even today: it’s not 'their' Internet and they're not 'their' issues, but it’s very much 'our' Internet, and its future is in 'our' hands. That overt recognition of a shared responsibility for the Internet is indeed a big shift in governmental perspective, and for me that's what made this particular meeting one of the more important meetings in 2008.

Geoff Huston

Plenary Session
Masashi NAKADO
Senior Vice Minister of Economy ... and Industry
Japan
OECD Ministerial Meeting
on the Future of the Internet Economy
Seoul, Korea, 17-18 June 2008

▲ Masashi Nakado

Day In The Life of the Internet (DITL)

This year, for the first time, APNIC participated in the 'Day in the Life' of the Internet (DITL) data collection event. From 18 to 19 March 2008, all traffic to and from the APNIC DNS nameservers in Tokyo, Hong Kong, and Brisbane was captured and stored. This raw data was then uploaded to the 'Cooperative Association for Internet Data Analysis' (CAIDA) repository:

⇨ http://www.caida.org

The data is available as a long-term analysis source, for registered researchers.

APNIC was able to provide the data from all its DNS nameservers by deploying new data collection servers, based on a 'passive tap' model, that is, highly reliable network tap devices were deployed at each location and inserted in front of the network connections of each APNIC nameserver. This allowed a new host to be connected and collect the data without any impact on the deployed service.

A total of over 297GB of data was collected by APNIC in the 48 hour period. This took another week to upload to the CAIDA repository, located in the USA. CAIDA collected over 1.9TB of data from over 50 sources, across the root, ccTLD, and other DNS operations communities.

Data collection exercises like this have previously been organized by CAIDA in 2006 and 2007. It is hoped that the continued data collection will provide a long-term snapshot of Internet activity in the future.

By arranging to capture the full dataflow (both queries and responses) for the DNS, it is hoped that information about worldwide trends in Internet usage can be measured. CAIDA's own website on the project provides a summary of the kinds of questions that researchers hope to be able to answer from analysis of the data. These cover areas such as the role of locality in Internet usage, workload and traffic/performance measures, questions about the DNS itself as a service, addressing and routing, and social issues.

⇨ http://www.caida.org/projects/ditl/questions/

Data is being held for use under a range of terms, depending on the policy of the data provider, and ranges from unrestricted anonymous versions of actual data capture to restricted rights and visibility of analysis outcomes.

Because APNIC provides both reverse DNS (in-addr.arpa) nameservers and also acts as a secondary for a range of ccTLDs, APNIC was in a position to provide data capture in both forward and reverse DNS contexts. APNIC is a 'secondary' nameserver for several of the other RIR's ranges and so provided data on DNS services for address ranges beyond the Asia-Pacific address management footprint.

Further research work on the DNS is planned by APNIC in 2008, exploiting the service deployed for the DITL data capture, which will be presented at our meetings and for publication.

George Michaelson

Increased demand for IPv4 driven by uptake of mobile telephony

The exponential growth of the Internet has naturally brought with it an increasing demand for IPv4 address space, particularly as the world's developing economies start to make use of the Internet and demand their share of these limited resources.

In addition, the strong uptake of mobile telephony is creating particularly strong demand for IPv4 address space as more mobile Internet connected devices proliferate and require their own unique public IP address.

The number of Internet-connected mobile phones is growing at a much faster rate than fixed installations, and this sector is responsible for the majority of new demand.

APNIC has measured a noticeable increase in requests for IPv4 address space in recent months.

Peer-to-peer applications, such as chatting, file sharing, and gaming, will increasingly dominate the mobile Internet landscape. These applications require publicly visible addresses, a need which is currently being fulfilled by IPv4.

Predictions regarding the longevity of the IPv4 address pool have not necessarily taken into account the popularity of third-generation (3G) mobile telephony; and at this increased rate of uptake, the available reserves may not last as long as expected.

The obvious solution to this problem is to use IPv6 exclusively for mobile connected devices, as this would immediately rectify the problem of address space scarcity.

Utilizing IPv6 for mobile telephony also simplifies network maintenance by removing the need for private address space and Network Address Translators (NATs).

The growth in 3G mobile telephony is likely to be one of the major factors driving the transition to IPv6 technology.

Support for IPv6 has been specified in various 3G standards for some time now. Therefore, it is important to ensure that the remaining network infrastructure is IPv6 compatible in order to facilitate the continued unrestricted deployment of mobile phones.

APNIC cooperates for DUMBO deployment in Myanmar

In the event of natural disasters, the telecommunications infrastructure that is so ubiquitous in everyday life is often compromised or completely destroyed, severly hampering the organization of rescue efforts in the affected region.

A case in point was the 2004 Indian Ocean earthquake and resultant tsunami. In the worst affected areas, the destruction of telecommunications systems caused a range of problems. All cell phones ceased to function, making it impossible to contact medical staff. Fortunately, many doctors made their way to the hospitals of their own initiative after discovering that the phones were not working. People away from the flooded areas were unaware what had happened, frustrating the doctors ability to treat the patients due to their lack of knowledge as to what had happened.

Medical information about how to treat these specific types of wounds, usually accessed via the Internet, was suddenly unavailable. It became impossible to request assistance or facilities from elsewhere, further retarding rescue and care efforts.

The major lesson from this event was the importance of having a network that can be set up at short notice and function independently of any existing infrastructure.

Project DUMBO

Project DUMBO (Digital Ubiquitous Mobile Broadband OLSR*) is an initiative that acts in response to emergency conditions, such as those immediately following a natural disaster, to deploy a mobile wireless network to assist in emergency communications.

This mobile ad-hoc network (MANET) consists of a collection of mobile nodes that automatically cooperate to allow wireless data transmission. The MANET does not rely on any external infrastructure and is lightweight, thereby allowing it to be installed quickly and easily.

With its portable nodes, MANET coverage can penetrate deep into areas inaccessible by road and into areas where all fixed telecommunications infrastructure has been destroyed.

DUMBO allows streaming video, VoIP, and short messages to be simultaneously transmitted from a number of mobile laptops to the central command centre or to the other rescuers at the same or different disaster sites.

Streaming video can be sent from each node to the central coordination unit, which can broadcast to all nodes on the network. This command centre can be located either in the disaster areas or indeed anywhere where there is Internet access.

The command centre uses a face recognition module to identify potential matches between photos of the faces of unidentified victims in the field with images stored at the command centre.

In addition, sensors can be deployed to measure environmental data, such as temperature and humidity. Data from the sensors can be sent to the command centre, where it can be analyzed or passed on to the other mobile nodes.

* Optimized Link State Routing Protocol

How is APNIC participating?

The Asian Institute of Technology (AIT), dotAsia, and APNIC have been cooperating for DUMBO deployment following the devastating cyclone Nargis that hit the delta regions of Myanmar, which caused catastrophic devastation and casualties to the five states of Yangon, Bago, Ayeyarwaddy, Kayin, and Mon.

Training sessions were held 21-24 May as part of the combined effort to assist in the establishment of a post-disaster recovery management infrastructure.

APNIC's role in the training involved:

- Delivering fundamental network concepts to the Burmese engineers
- Assisting in the mini-setup of the DUMBO system to allow the engineers to understand the setup process and the capabilities of the system
- Providing an avenue of technial assistance to AIT and the Myanmar engineers with support from the APNIC offices in Brisbane

APNIC contributed a hardware package to the project as well as substantial funding to aid the acquisition of additional resources and equipment for the deployment.

APNIC staff members Annuttara Tallents and John Tan were in charge of delivering training in networking fundamentals, providing the essential knowledge to assist the Burmese engineers in their understanding of the DUMBO project. This assistance was greatly appreciated, and APNIC also collaborated on discussions of topology design, which was challenging due to the long distances between rescue centres, ranging from 30 to 80 km.

Project DUMBO, while offering substantive support in post-disaster recovery, still remains in a proof-of-concept stage. More work is required, and through these instances of deployment, more is learnt about the best means of setting up effective mobile networks. Research continues to be undertaken and it is hoped that in future, Project DUMBO will continue to aid in providing the much-needed relief to disaster victims.

Sources:

➡ http://www.interlab.ait.ac.th/dumbo/index.php

32-bit AS numbers

From 1 January 2009, when you request an AS number from APNIC, you may receive more than you expected. The AS number registry has been expanded from its original 16-bit range (AS numbers 0 through 65535) to a 32-bit range (AS numbers 0 through 4,294,967,295). From that point on, APNIC will allocate 32-bit AS numbers by default, and will only allocate 16-bit values if explicitly requested.

Why the change?

This change is being made because AS numbers in the 16-bit AS number range are close to running out. The recent rate of AS number consumption has been such that the remaining pool would have been completely exhausted by March 2011. The Regional Internet Registries have each adopted policies that allow ISPs to transition their networks over the course of a few years to use this extended AS number range, without the need for the last-minute deployment of rushed changes to the BGP.

From 1 January 2007 until 31 December 2008, ISPs have been able to specifically request an AS number from the extended 32-bit number pool (that is, an AS number greater than 65535); but, by default they were assigned an AS number from the original 16-bit number pool (that is, a number less than 65536). From 1 January 2009, the allocation practice will be reversed, and unless a 16-bit AS number is specifically requested, AS numbers will be allocated from the extended 32-bit number pool.

Implications for ISPs

What are the implications for ISPs with this AS number allocation policy? What should an ISP have as a prerequisite to requesting a 32-bit AS number?

AS numbers are used in the context of inter-domain routing, particularly in the BGP protocol. If an ISP wants to use an AS number that is greater than 65535, it will need to deploy a 'new' version of BGP. That is, it will need to deploy a version of the BGP protocol in its routers that is compatible with 32-bit AS numbers, as most existing BGP implementations use 16-bit data structures.

The rest of us...

But what about everyone else? What about the existing BGP world that uses 16-bit AS numbers? Does the first public deployment of a 32-bit AS number force everyone else to also upgrade their versions of BGP?

Even though every other network already has a 16-bit AS number, will every network also need to upgrade their BGP to see these new extended-length AS numbers? Will everyone need to apply some form of upgrade to their equipment and operational support systems before the first extended-length AS number is used in the public Internet?

The short answer is "No!" Current networks using 16-bit AS numbers need not change anything!

Anything at all? Really?

Well, that's not quite true. Let's see why.

The issue here is actually one of 'transition', and the way in which this transition has been integrated into the specification of BGP properties to support 32-bit AS numbers. The transition to 32-bit AS numbers has been carefully constructed to be backward compatible, and the changes to BGP only affect 'new' 32-bit BGP implementations. The reassuring news is that if you have a 16-bit AS number and are running 16-bit AS BGP today, then you do not need to change anything at all in your routers. The Internet will still work, and you will continue to see routes to all advertised networks, irrespective of the existence of 32-bit AS numbers in the network. You will be able to send packets to those autonomous routing domains numbered from the 32-bit AS number space, and they will be able to send packets back to you. You don't need to upgrade your version of BGP, nor do you need to make any router configuration changes in your network. The Internet will work as intended without a break in connectivity.

BGP backwards compatibility

However, some things might change for you. To understand what is going on, it is useful to describe how BGP has managed to be backward compatible across this change.

A BGP session uses an initial handshake to determine the identity of its neighbour. To allow a 'new' version of BGP to speak to an 'old' version of BGP, it presents itself as the 16-bit AS 23456 in the initial handshake, and includes a 32-bit capability advertisement. If the neighbour is also a 'new' BGP, it will pick up this capability and use the extended length AS number and proceed as normal. If the neighbour is an 'old' BGP, the 'old' BGP speaker will believe it's speaking to AS 23456. In this case, the 'new' BGP will pick up that it has an 'old' neighbour and make some changes to the way BGP operates.

When the 'new' BGP speaker sends a BGP UPDATE to the 'old' BGP, it uses a combination of translation and tunnelling to transform the AS path from a sequence of 32-bit values to a sequence of 16-bit values. For translation, each AS is truncated to a 16-bit value. If the 32-bit AS value was less than 65536, then the leading zeros are stripped off and the equivalent 16-bit value is used. Otherwise, the 16-bit value 23456 is used in its place. This creates an equivalent 16-bit AS path of the same length as the 32-bit version. For tunnelling, the original 32-bit AS path is placed in an opaque community attribute.

This update will traverse the 'old' BGP world as normal. The 16-bit AS path will be augmented with each transit AS, and the opaque community attribute will be unaltered.

When passing a routing update from the 16-bit 'old' BGP world back into the 32-bit 'new' BGP world, the opposite transformation is applied. All the AS numbers in the AS path attribute are expanded to the equivalent 32-bit values by adding the leading 16 zero bits to the AS number value. If there is the appropriate opaque community attribute present, then all instances of AS 23456 can be converted back to their 32-bit values. If nothing untoward has happened, the 'new' 32-bit BGP world sees an accurately re-constructed 32-bit AS PATH, preserving both the AS path length metrics and BGP's routing loop detection capability.

This assumes that nothing unusual has happened as the BGP UPDATE message traversed the 16 bit 'old' BGP world. It may be the case that the opaque community attribute has been dropped, or some form of proxy aggregation in the 16-bit world has garbled the AS path so that the reverse substitution cannot be performed. But even this is not a fatal condition for BGP itself. Even without this substitution, the AS PATH length metric is preserved, and routing loop detection can still be performed, although in a degraded fashion. So, if BGP encounters something unexpected in the translation back from the 'old' to the 'new' 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.

So, even if you do absolutely nothing in your 16-bit 'old' BGP routing domain, Internet-wide routing will still work, and reachability information will still be propagated in useful forms. Nothing will 'break'. However, there are some things to check, and maybe alter, in the larger environment of your operational support framework.

Implications for existing BGP technology

The implications for 'old' world BGP routing domains include the following:

Do not strip the NEW_AS_PATH opaque community attribute from BGP Updates. The 32-bit injection will mark this attribute with the optional and transitive attribute flags. Sixteen-bit BGP speakers should avoid using local policy configurations that alter this attribute setting or remove this attribute from the prefix. If we are using standards-compliant language, then that's a 'SHOULD', not a 'MUST', by the way. Nothing will break if you don't, but we'll see faster routing convergence if this attribute is left intact. The default BGP action will handle this correctly.

Similarly, it's strongly preferred that the NEW_AGGREGATOR attribute also be carried as an optional transitive opaque community attribute when present in 16-bit BGP. Again, nothing will break if you don't, but we'll see faster routing convergence if this attribute is left intact. The default BGP action will handle this correctly.

The next implication is that the 'old' 16-bit BGP world will see more and more instances of AS 23456 as both an originator of prefixes, and as a transit provider. This is not a mistake; it's just the only way that the 16-bit world can carry a placeholder for a 32-bit AS value.

A word of caution

However, if you choose to not upgrade your BGP software, watch out for the following:

Many ISPs used directed community attributes to signal to a remote AS. A prefix that has explicit signalling to AS65505 uses a community attribute of '65505:123', for example. However, this will not work if you wish to generate a signal to a 32-bit target AS. At the very least, your BGP version should support extended community attributes (RFC 4630) and also support the means of entering 32-bit AS numbers into these attributes.

You should also expect a modest increase in the 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 BGP's total memory demands will triple, only those requirements relating to AS path storage.

We might anticipate slightly poorer performance in routing. The specific cases where convergence times will increase are in those circumstances where the NEW_AS_PATH attribute is lost in transit through the 'old' BGP 16-bit world. In such cases, loop detection will take slightly longer, and this will have some level of impact on convergence times.

So you don't need to upgrade the BGP software on your routers. However, you would be well advised to thoroughly audit the capabilities of your network management systems. In particular:

An 'old' BGP ISP may see routing peers, both as customers, peers, and possibly upstream transit providers, using 32-bit AS numbers. But as your local BGP is an 'old' world BGP, your routers will not be aware of these 32-bit AS values. From your router's perspective, AS 23456 is going to start popping up both as a diverse prefix originator and possibly as a ubiquitous transit provider. The ISP's operational support systems (OSSs) should be able to store the corresponding AS numbers of these BGP routing peers as 32-bit number values, simply to avoid unnecessary confusion and potential ambiguity. So, you should probably ensure that your OSS is 32-bit compliant for AS numbers, and is capable of storing and displaying the configuration state for AS numbers in 32-bit format, even if you are running a version of BGP that only supports 16-bit AS numbers. Depending on the way in which the OSS has been designed, implementing this requirement may vary from the trivial to the extensive.

If you use this OSS to generate router configuration fragments, AS path filters and similar, then you may need to ensure that your OSS is capable of generating both 16-bit BGP configuration fragments and 32-bit configuration fragments. In the case of the 16-bit version, the OSS will need to transform the locally held 32-bit AS number values into the 16-bit equivalent value of AS 23456.

Routing registries will also need to be updated, allowing the registry's clients the ability to deposit registry entries that refer to 32-bit AS numbers. When using a routing registry to generate local configuration fragments, the generated configuration entry will need to differ, depending on whether a 16-bit or 32-bit configuration fragment is required. For example, the routing registry may have an entry relating to a routing domain of AS 1.2, but your 'old' BGP router will need to be provided with a generated configuration fragment that refers to AS 23456.

If you filter based on AS numbers, then any filter generator code that you might use will need to translate the 32 bit AS numbers stored in your local routing policy database into instances of AS 23456.

If you do elect to perform a network upgrade, then there are also some further things to think about in the planning process:

There is no dynamic capability to support a change from 16-bit 'old' BGP to 32-bit 'new' BGP. When a routing domain wants to transition from a 16-bit to a 32-bit 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.

Further reading:

Exploring AS numbers:

➡ http://www.potaroo.net/ispcol/2005-08/as.html

Geoff Huston

New at APNIC 26: eTicketing

As part of APNIC's commitment to assisting our meeting attendees' registration needs, APNIC is introducing the eTicket. The eTicket has been designed to streamline the registration process, resulting in a smooth experience for both meeting attendees and the APNIC staff.

The eTicket contains a barcode that simplifies the process of identifying an attendee at the registration counter. The barcode contains the registration number, which links to the registrant's details in the APNIC events system.

These details include the registrant's:

- Name
- Organization
- Registration type
- Registered activities
- Charges

eTickets will be emailed to those registrants who have complimentary registrations or those registrants who have made online payments prior to meeting commencement.

Internet Governance Forum 2008

This year, the Internet Governance Forum (IGF) will be held in Hyderabad, India from 3 to 6 December and will be hosted by the Indian government. This will be the third IGF - the first two were held in Athens in 2006 and Rio de Janeiro in 2007, respectively. This is the first time the IGF will be held in an Asian country.

This year's IGF will see improvements in accessibility, with all the main sessions being webcast and translated into the official languages of the United Nations. Workshops will be audiocast, and online participation options will also be made available.

The IGF is principally a forum for discussion. No formal decisions are made; it exists solely for the purpose of policy dialogue regarding issues of Internet Governance.

The IGF involves what is now referred to as the Multistakeholder Advisory Group (MAG), whose purpose is to assist the Secretary-General convene the meeting. This group consists of members from governments, the commercial private sector, the technical community, and public civil society. All stakeholders appointed to the MAG participate as equals.

The proposed agenda for the upcoming IGF will include:

- Reaching the next billion
- Promoting cyber-security and trust
- Managing critical Internet resources
- Taking stock and the way forward
- Emerging issues - the Internet of tomorrow

The 'emerging issues' session, held at the end of the forum, is an important aspect of the IGF. This session recognizes the dynamic nature of Internet governance issues by allowing themes and topics that have arisen during the main sessions and the workshops to be discussed.

The overarching theme will be 'Internet for all', which includes and emphasizes the importance of ensuring that the world's 650 million people with disabilities have unencumbered Internet access as the focus shifts to connecting the next billion people yet to benefit from Internet access.

Issues regarding IP addressing issues are tentatively on the agenda, to be finalized in September 2008. These include IP addressing as a 'critical Internet resource'. This is one of the main themes at IGF 2008, and focuses on the transition from IPv4 to IPv6, IPv4 depletion, and the role of the RIRs in enabling routing security via Resource Public Key Infrastructure (RPKI).

It has been acknowledged by Professor Milton Mueller of Syracuse University that the RIRs will become a very important trust point in the Internet if RPKI flourishes. Professor Meuller has proposed a workshop at IGF 2008 to discuss this issue. Although this workshop proposal has not yet been officially accepted, it does indicate that an awareness of secure routing is reaching outside the small circle of the RIRs and into the wider community.

APNIC will be contributing to the IGF with a booth at the IGF village as part of the Number Resource Organization (NRO). It will also be contributing to the workshops and sessions relating to IP addressing, both as APNIC itself and as a component organization of the NRO.

APNIC has submitted a workshop proposal titled "Challenges facing Internet operators in developing countries", final approval for which will be announced in August.

In addition, APNIC will be assisting with providing web and audiocasts of the event.

Internet governance in the Asia Pacific

There are three key issues that dominate the landscape of Internet governance in the Asia Pacific region. These are:

1. Internationalized domain names
2. Internet access
3. IPv4 depletion and IPv6 adoption

The issue of internationalized domain names refers to the fact that many of the languages used by people in the Asia Pacific use non-ASCII scripts, and there has traditionally been no support in top-level domain names for non-ASCII text. Providing top-level domain name support for these languages is important to many people in the Asia Pacific.

While most people think of places such as Africa and South America as being the frontier of Internet provision, low Internet penetration is actually a problem in many areas of the Asia Pacific. This is particularly the case in the Pacific Islands, as well as economies in Southeast and South-central Asia. Poverty is a major barrier to Internet connectivity, so reducing costs and finding innovative methods to bring Internet access to these areas are critical points of focus.

IPv4 depletion and the lukewarm transition to IPv6 are making headlines around the world as the free pool of IPv4 addresses evaporates. This is particularly relevant in the Asia Pacific region, as it contains two of the largest and fastest developing economies in the world. China and India are together generating extremely strong demand for Internet resources as more and more of their populations get online, especially with mobile devices, which tend to use proportionally more address space. This awareness has led to economies in the Asia Pacific, such as Japan, China, and Korea, leading the way with IPv6 adoption.

How you can participate in the IGF:

If you would like to participate in the IGF, you can submit a paper outlining your opinions on the issues being discussed or on a completely different topic if you think there's an important Internet governance issue that needs discussing. The deadline for the submission of papers as an input for the Hyderabad meeting is 12 September 2008.

All papers submitted by that date will be reflected in a synthesis paper prepared by the Secretariat for the Hyderabad meeting. This is you and your community's chance to have a say on the issues to be discussed in Hyderabad.

See:

➡ http://www.intgovforum.org

Sam Dickinson

Training update

Over the past few months, the APNIC training team has held 34 informative training sessions throughout the Asia Pacific region. These included, for the first time, training in Japan, Guam, and Brunei. Trainers were also sent to Thailand after the devastating floods to assist AIT (Asia Institute of Technology) in the deployment of support to Myanmar.

▲ Thailand, June 2008

▲ Guam, June 2008

As part of APNIC's continuing development of its training courses, the training team is evaluating eLearning platforms for the interactive delivery of web classes. A selected group of members have been invited to participate and test the shortlisted products and provide their feedback. The testing will be done sequentially, and the course content centres on IPv6.

The training team is also implementing the results of the training survey that was held in December 2007. The purpose of the survey was to assess whether the provision of training was supportive and useful to our members. Overall, the results were very positive, and the team is now in the process of developing and improving the course content to align with the results of the survey and to meet current technical needs.

For example, the new topics being developed include:

- Security forensics
- IPv6 level 2: Deployment to the edges (customers, content providers, and corporations)
- Internet operations for corporations and content providers – multi-homing.

APNIC is also discussing the establishment of Memorandums of Understanding with a number of institutions, such as Team Cymru, VNNIC, ASTI, NAV6, ISPBD, NIXI, and AIT, to further assist in the development of training services for our members.

To further promote IPv6 in developing countries, APNIC has become involved with the "6Deploy" project, a worldwide initiative from the European Union to provide training and deployment support. APNIC is a sponsoring partner and recently assisted 6Deploy with their training lab in Kenya.

The training team has numerous upcoming training events. For schedules and further information, please see:

➡ www.apnic.net/training

▲ Indonesia, 2008

▲ Indonesia, 2008

▲ APTLD APNIC InterLab collaborative ccTLD workshop - Bangkok

▲ APJII Open Policy Meeting / APNIC training - Indonesia

MyAPNIC and login

APNIC is pleased to announce a significant upgrade to the MyAPNIC service and a number of other facilities we provide our members and the wider community. Currently, all APNIC members are using certificates to access MyAPNIC and manage their resources. Starting in the fourth quarter of 2008, members will be able to use a consistent username and password across MyAPNIC and all other web services APNIC offers. This upgrade will be deployed incrementally, so not all services will be enabled initially.

This change is in response to ongoing user feedback about the APNIC website and APNIC's strategic planning goals for continuous improvement to the Web and other related services. For all but a small number of high-privilege operations, it will now be possible to use a simple login process to access member services. This is expected to make the management process simpler and, in most cases, faster.

APNIC continues to support the use of client certificates for high-privilege applications. The username/password changes will complement the certificate, and be used as an additional check for certain functions. This follows an extensive security review conducted as part of APNIC's public key certificate practice statement (CPS) development. Use of both username/password and certificates for high-privilege services will allow us to combine low and high-privilege services in one integrated portal model.

For existing certificate holders, a simple re-enrollment process will be used online to register members into the system and create their username and password. Future (new) users can be managed directly from the account management screens.

MyAPNIC2 Beta will also offer enhanced contact and user-management screens for authorized users to manage access to APNIC services. We are now able to offer vastly simpler user-management processes, which will simplify staff changes, role changes, and access control for most account holders. In the future, a single authorized contact will be able to approve and manage all subsequent staff access into MyAPNIC, including the approval of certificate issuance. This change means that the turn-around time to be given an APNIC-issued certificate will be reduced from days to hours, with no need to present credentials to APNIC helpdesk/hostmaster staff. We also expect that future use of the services by independent consultants on behalf of the members will be improved so that APNIC members can consider outsourcing significant aspects of Internet number resource management if they so wish.

Existing information services, like ICONS, and registration for training and APNIC events will be modified over time to adopt the common username/password model. This means that single-time registration of your contact information, including email, will work for all our online services.

The same facilities will now be used to manage APNIC contact information, including mailout and corporate contacts. APNIC members will be able to use one management screen to control information regarding billing recipients, voting, membership, and resource management.

A future deployment (in the next 2-3 months) will see a new APNIC certification authority deployed to complement this new username/password model. This will permit new certificates to be issued directly by authorized member managers, and also allow for certificate re-issuance online. APNIC hopes that this will encourage wider adoption of the certificate, and access to MyAPNIC.

Amendments to the 'clients first' project have been integrated into these changes so that almost all initial company registration information is collected only once, and users will no longer be expected to re-enter basic contact details when they register for different APNIC services. In almost all cases, an existing username and password will 'just work'.

George Michaelson

ICONS gets a wiki

The Internet Community of Online Networking Specialists (ICONS) website is undergoing a migration.

ICONS, a community networking site run by APNIC, AfriNIC, APRICOT, and SANOG, was established as a space for the Internet community to share information addressing specific topics. Since its launch, ICONS has attracted nearly 600 members, but to make it even easier to contribute we are migrating the site to a wiki platform.

Initially, the new wiki area will focus on a couple of specific issues (particularly the introduction of 4-byte Autonomous System numbers and IPv6), but as the site grows we will provide more functionality and ultimately the wiki will provide a better environment for the community to interact, share ideas and information, or work cooperatively on issues faced by network users and builders.

Whether you are a vendor, ISP, user, or regulator, we encourage you to contribute anything interesting that you think may be of benefit to others. You can add comments or bookmarks as well as add and edit pages, or build an RSS feed of your favorite area to make sure you are kept up to date with any changes.

Please invite your friends and colleagues to join the ICONS community!

➡ http://icons.apnic.net

icons
The Internet Community of Online Networking Specialists
Sharing knowledge & experience
Want to know more? Please visit:
www.icons.apnic.net

Resource Certification

After the initial release of our resource certification engine earlier this year, APNIC is now working towards the deployment of a member portal (MyAPNIC) service for resource certification. The portal will be the beginning of full service delivery for resource certification in the Asia Pacific region.

This work is aimed at giving MyAPNIC users the ability to use APNIC as a one-stop shop to manage resource certificates, Route Origin Attestation (ROA), and other signed objects, all within the GUI we offer them for resource management. Users will be able to create, manage, apply, and destroy certificates over all their resources, or over subsets of these, and see them published in the worldwide resource certificate repository hierarchy at APNIC.

New screens are being added to MyAPNIC to manage collections of Internet number resources. These collections can then be used to nominate which certificates need to be used (or created, if need be) to sign with, so you can say things like:

"please route all my prefixes from AS <x>"

or

"please route all my customer prefixes as listed here from AS <y>"

…and have MyAPNIC manage the process.

This work is complemented by changes in the MyAPNIC access control model, which will free up non-resource management areas of MyAPNIC to a simple username/password access control model. This would also retain user certificate controls over the address management functions, including resource certification. This change is designed to ensure complete trust in the certification process over resources, while liberalizing access to the information, training, and other aspects of member management in MyAPNIC.

This new UI will be released at APNIC 26 in New Zealand.

The service has been implemented as a hosted certification engine, which is run separately to the existing APNIC service, allowing APNIC to meet its certification practice requirements for information management. The two services communicate using the protocols designed in collaboration with the RIR community, and deployed earlier in the year at APNIC.

APNIC has continued to test the inter-operation of its systems with other RIRs and code developers, which culminated in a series of informal tests carried out between the ARIN and APNIC codebases earlier this year.

APNIC has also recently submitted three new drafts to the IETF defining different aspects of the resource certification framework. This standardization activity aims to bring together both resource management and routing security outcomes in one consistent framework.

Resource certification continues to underpin emerging policy development around the globe, and there is an increasing interest in certification as the backbone of a worldwide resource transfer process, as well as securing routing. Recent court cases in the USA have made it clear that whois records and other information sources cannot be used as the basis of information management, which requires more rigorous proofs of responsibility over a resource to control its disposition.

Whois, of course, remains important for the publication of resource registrations, and APNIC is considering design issues to allow whois records to also carry resource certification records.

George Michaelson

PacNOG IV

In June 2008, APNIC participated in PacNOG IV in Port Vila, Vanuatu.

The Pacific Network Operators Group's 4th meeting was held in Vanuatu for the first time, enabling many locals to easily participate in the six days of workshops, lectures and tutorials offered at the meeting.

APNIC staff members Elly Tawhai, Champika Wijayatunga, Cecil Goldstein, and Sunny Chendi discussed Internet resource management, network security, APNIC services, and particularly how to prepare as it becomes necessary to transition to IPv6 and 4-byte AS numbers.

Islands Network Operators Group - PacNOG IV
29 June - 5 July 2008 Port Vila VANUATU
PacNOG
Last Row L to R Robertson Sarki, Junior Joseph, Eddie Sandy, Gerard Metsan, Sakaio Manoa, Teuila Fifita, Michael Seeng, Srinivan (Sunny) Chendi, Philip Smith, Jethro Webston, Mitchell Tutangata
3rd Row L to R Anacklet Kali, Sanjay Reddy, Philippe Gardon, Martin Mahe, Sam Sak, Gilbert Bourdon, Wallis Taura, Joseph Wong, Dennis Robert
2nd Row L to R Lynette Andre, Andrew Molivurae, Tamara Sikivou, Ralph Ray, Tearia Mateariki, Save Vocea, Ishaka Buaka
Front Row L to R Bernie Fuaoanee, Elly Tawhai, Champika Wijayatunga, Cecil Goldstein, Sebastien Le Guillou, Amond Melkaros, Paula Mataiiraga, Terry Rupeni, Fred Christopher

Ten years as Director General

– Paul Wilson's decade of service at APNIC

They say the Internet runs in 'dog years' – meaning that time flies much faster on the Internet – like, seven times as fast. It is also said that no one can forecast even three years into the future of the Internet, so it's obviously impossible to foresee an entire decade!

Thus, our Director General achieved something unforseeable this year.

Paul Wilson was appointed to the position of APNIC's Director General in 1998, and arrived at the tenth anniversary of his service as the DG on 3 August 2008. He is the most long-standing president among the five Regional Internet Registries.

My strongest memory from the early days of Paul's term was at the APNIC Member Meeting at APRICOT2000 in Seoul. There was a ceremony to launch the new APNIC logo, which has been in use since then. The original logo was lovely and actually very familiar to everyone, but the new one looked so neat that it seemed to envisage the successful forthcoming years of APNIC.

And this has, in fact, come true. The Secretariat's operation has grown ten times as big since 1998, both in terms of the headcount and of the annual budget. Not only has the fundamental service been stably run, but its services have extended to enhanced training courses, formulating the Number Resource Organization, handling Internet Governance issues, and so forth.

It goes without saying that the pioneers' devotion during this emerging period established the firm basis of APNIC's business today, but Paul's contribution stands apart and has been truly invaluable, allowing the entire APNIC community to benefit from APNIC's service today.

The Executive Council both very much appreciates and is extremely proud of Paul's professionalism as APNIC's Director General for the last decade, and we hope we can enjoy his excellence as long as possible.

Akinori Maemura

APNIC mourns loss of community champion

Dr. Masaki Hirabaru of Japan's NICT, the National Institute of Information and Communication Technology, passed away on Tuesday, 29 July 2008 at the young age of 48.

One of the initial staff members of the APNIC Secretariat, Dr. Hirabaru made a great contribution to the establishment of APNIC, which was originally APCCIRN's pilot project. He was also the founder of JNIC – the predecessor organization to JPNIC – and was the founding Committee Chair of JPNIC.

Born in Kitakyushu in 1960, Dr. Hirabaru studied at Kyushu University, where he gained a doctorate in engineering in 1988. Kyushu University benefited from Dr. Hirabaru's expertise as he worked as a lecturer in its engineering department.

During a distinguished career he worked at a number of Universities, including Tokyo University, the University of Michigan, and at the Nara Institute of Science and Technology.

At the National Institute of Information and Communications Technology, Dr. Hirabaru served as the Group Leader of the Network Architecture Group in the New Generation Network Research Centre. Prior to this, he was the Director of Internet Research at the Institute of Systems & Information Technologies in Fukuoka City for three years.

He also worked to launch a number of network research projects such as JAIN, TRAIN, QGPOP, and he worked on a new generation network architecture and led the AKARI Project.

The staff members at the APNIC Secretariat would like to express their sincere condolences.

May his soul rest in peace.

Contribute to Apster!

As part of our continued efforts to bring you a more interesting and relevant publication, the Apster team is now actively seeking contributions from all corners of the Asia Pacific networking community.

We're particularly interested in your experience in your part of the world. Do you have a certain speciality that you'd like to share? Apster is distributed to all our members throughout the region, so this is a great opportunity to spread your message and get your name out to our broad readership.

We understand that many potential contributers may feel intimidated by the task of writing in English, especially if it's their second language. There is no need to worry — our skillful editors will make sure that your message is communicated perfectly.

You will be given full credit for your contribution, so don't miss this opportunity to get your name up in lights and tell the world what it's really like in your neck of the woods, in your organization, or just being you.

Topics like:

- Interesting events you've attended
- Developments in network connectivity in your region
- Stories about how you coped with natural or man-made disasters - especially those impacting the network
- The social effects of bringing the Internet to your community

...or even just what it's like being a network engineer from your point of view.

Transitioning our network infrastructure to be IPv6 compatible is a hot topic and can be something of a challenge for various reasons. We'd love to hear about any surprises or difficulties you faced and overcame in upgrading to IPv6. Was it more or less difficult than you imagined? Do you have any advice for other network engineers in your position?

If you think you'd like to contribute any material to Apster, please contact us at: publication@apnic.net

Staff updates

German Valdez, Communications Area Manager

Even though German is new to APNIC, he has been part of the RIR community since 1998. Originally from México, he started 10 years ago as a Hostmaster at NIC México. He was part of the first LACNIC board of members who negotiated the recognition of the Latin American and Caribbean Regional Registry with ICANN. In 2003, he moved to Uruguay to be part of the recently created LACNIC. He spent the last five years working as the Policy and External Relations Manager at LACNIC before joining APNIC in February 2008. German is very pleased to be part of the Asia Pacific community and looks forward to contributing in his new position to the stability and growth of the Internet in the Asia Pacific.

Adam Gosling, Publications Unit Manager

Adam joined APNIC in May 2008 following a long career in computing and communications publishing for channel, trade and technical titles. He has held a variety of positions including Editorial Director and Publisher, working in Sydney, Singapore, and most recently Brisbane. Adam also has experience in IT PR representing some of the industry's biggest names in software and internetworking.

Bhadrika Magan, Senior Editor

Bhadrika joined APNIC in May 2008. She holds a Bachelor of Arts in Classical History and a Bachelor of Laws. She brings to APNIC several years' editorial and written experience, including copywriting and copyediting for a variety of publications. Bhadrika, along with a highly skilled team, is responsible for the accurate and timely delivery of all APNIC communications.

Alan Golding, Systems Administrator (Desktop Support)

Alan Golding originally comes from Ireland and has been working with APNIC for nearly 4 months as a Junior Systems Administrator. He recently graduated with an Honors Degree in IT Management, after which he wanted to travel. His plan was to travel around South East Asia, ending up in Australia to hopefully find some career employment. He feels extremely lucky to be a part of the diverse and friendly APNIC team.

Vivek Nigam, Internet Resource Analyst (Helpdesk)

Vivek joined APNIC in February 2008 and has a Bachelor's degree in Information Technology and a diploma of Internet Programming. Vivek has previously worked in Desktop Support and IT Helpdesk roles at the University of Queensland, as well as for various ISPs. He has extensive experience supporting members with broadband troubleshooting, configuring ADSL modems and network administration skills in Novel and Active directory.

Tanya Samuel, Internet Resource Analyst (Helpdesk)

Tanya is originally from Papua New Guinea. She graduated in 2005 with a Bachelor's degree in Engineering Technology, majoring in computer engineering, from the Manukau Institute of Technology in Auckland, New Zealand. She then returned to Papua New Guinea and worked for PNG's largest ISP, Datec PNG Ltd., for 2 years as an ISP engineer. Tanya commenced work for APNIC as an Internet Resource Analyst in May 2008.

Wita Laksono, Internet Resource Analyst (Helpdesk)

Wita joined APNIC in June 2008, relocating to Brisbane from Indonesia where he previously worked for ISPs, gaining extensive skills and experience in technical support and systems administration as well as network engineer roles. Wita completed his studies in computer science in Indonesia and will be a valuable asset to the Helpdesk. Wita speaks Bahasa Indonesia and will be available for member enquiries.

Gary Kennedy, Software Engineer

Gary comes from a technical background, but due to his imposing stature he got sick of crawling under desks to deal with cables and decided that *sitting* at desks would be easier on his knees. After a few short courses to see if his casual programming skills would pass professional muster, he felt confident enough to try his luck in the corporate marketplace. His first 'real' software job interview was with APNIC. Knowing genius when they saw it, they immediately snapped him up.

Training schedule

2008

August

25-29	Christchurch, New Zealand (APNIC 26)

September

1-5	Cook Islands
2-5	Fiji
8-12	New Caledonia

October

6-10	India
6-11	Colombo, Sri Lanka
13-15	Male, Maldives
11-16	Ulaanbaatar, Mongolia

November

10-12	Canberra, Australia
12-14	Samoa
17-19	Tonga

December

9-12	Hong Kong
9-12	South China, China
15-17	Macau

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

www.apnic.net/training

If your organization is interested in sponsoring APNIC training sessions, please contact us at training@apnic.net

▲ Training in Singapore, June 2008

Calendar

APNIC 26

25-29 August 2008
Christchurch, New Zealand
http://www.apnic.net/meetings

PacINET 2008

1-5 September 2008
Rarotonga, Cook Island
http://www.picisoc.org/tiki-index.php

ITU Telecom Asia 2008

2-5 September 2008
Bangkok, Thailand
http://www.itu.int/ASIA2008/marketing/email/asia2008.html

NANOG 44

12-14 October 2008
Los Angels, USA
http://www.nanog.org

ARIN XXII

15-17 October 2008
Los Angels, USA
http://www.arin.net/meetings/index.html

RIPE 57

26-30 October 2008
Dubai, UAE
http://www.ripe.net/ripe/meetings/ripe-57

ICANN 33rd International Public Meeting

2-7 November 2008
Cairo, Egypt
http://public.icann.org

73rd IETF

16-21 November 2008
Minneapolis, MN, USA
http://www.ietf.org/meetings/meetings.html

Australian IPv6 Summit

17-19 November 2008
Canberra, Australia
http://www.ipv6.org.au/summit

Asian Internet Engineering Conference (AINTEC 2008)

18-20 November 2008
Bangkok, Thaailand
http://www.interlab.ait.ac.th/aintec08

AfriNIC 9

22-28 November 2008
Addis Ababa, Ethiopia
http://www.afrinic.net/meeting/index.htm

Internet Governance Forum

3-6 December 2008
Hyderabad, India
http://www.intgovforum.org

APTLD members and board meeting

8 December 2008
New Delhi, India
http://www.aptld.org/newdelhidec2008.htm

APNIC 27/APRICOT 2009

18-27 February 2009
Manila, Philippines
http://www.apnic.net/meetings/index.html

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	+61-7-3858-3188
SIP	helpdesk@voip.apnic.net
Fax	+61-7-3858-3199
Website	www.apnic.net
Helpdesk	helpdesk@apnic.net
Training	training@apnic.net
Apster	publication@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: 09:00 - 19:00 (UTC + 10 hours) Monday - Friday

Are you using MyAPNIC?

APNIC members can use MyAPNIC to:

- View APNIC resources held by their organization
- 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.