APNIC

Apster

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

www.apnic.net

Issue 16 December 2005

APNIC - Addressing the challenge of responsible Internet resource distribution in the Asia Pacific region

The view from the Summit: Where to now for the Information Society?

As the dust settles and the North African city returns to normal, **Samantha Dickinson** examines the outcomes, achievements, and open questions of the World Summit on the Information Society.

When Phase II of the World Summit on the Information Society (WSIS) came to a close on 18 November 2005, much of the world's media announced that, particularly in the area of Internet governance, WSIS had changed nothing. "ICANN's rule unchanged", "USA still in charge of the Internet", the headlines read. But the real outcomes are more complex. This article examines what the WSIS recommendations on Internet governance mean for the technical Internet community.

▲ The Internet Pavilion at WSIS attracted attention from the delegates, the public, and international media.

The Internet governance scope widens

In the early days of WSIS, much discussion of Internet governance revolved around domain names, root DNS servers, and IP addresses. But one positive outcome from WSIS Phase II is the official acknowledgement that Internet governance covers more than these few issues. The hope that this will lead to more balanced Internet governance discussions in the future is reflected in paragraph 58 of the *Tunis Agenda*, which recognises "other significant public policy issues such as, inter alia, critical Internet resources, the security and safety of the Internet, and developmental aspects and issues pertaining to the use of the Internet."

However, this statement also implies that domain names and IP addressing fall into area of public policy. Earlier in the *Tunis Agenda*, it states "Policy authority for Internet-related public policy issues is the sovereign right of States". The document also makes specific statements about public policy development in relation to ccTLDs and gTLDs, but – critically for the addressing community – it is less clear about public policy issues related to IP addresses, calling for "the reinforcement of specialised regional Internet resource management institutions to guarantee the national interest and rights of countries in that particular region to manage their own Internet resources, while maintaining global coordination in this area."

The Number Resource Organization (NRO) interprets this as clear support for the current Regional Internet Registry (RIR) system with its established open processes. But where exactly do governments, post-WSIS, fit within those processes? The *Tunis Agenda* certainly strengthens governments' roles in developing future Internet public policy, and since the *Agenda* also considers IP addressing to be within the public policy sphere, it appears that some governments will be more actively involved in the RIRs in future. In words of the *Agenda*, there is a need to develop an "enhanced cooperation model"; however, at this stage it is far from clear just what that model will look like.

For its part, APNIC already does include a formal role for government in the endorsement of National Internet Registries in the Asia Pacific region, and some governments have been active in this area, while others have chosen not to be. In either case APNIC has indicated it will continue to strengthen relationships with governments in the Asia Pacific region and to encourage more dialogue on the technical issues associated with IP addressing policy.

P 4

21st APNIC Open Policy Meeting
27 February - 3 March 2006 Perth - AUSTRALIA

P 2

21st APNIC Open Policy Meeting
27 February - 3 March 2006 Perth - AUSTRALIA

21st APNIC Open Policy Meeting

APNIC invites anyone with an interest in Internet address policy to attend the 21st APNIC Open Policy Meeting (APNIC 21), which will be held in conjunction with APRICOT 2006 from 27 February - 3 March 2006. The venue for the meeting will be the Perth Convention and Exhibition Centre (PCEC) in Perth, Australia.

APNIC 21 will include tutorials, Special Interest Groups (SIGs), Birds of a Feather sessions (BoFs), hostmaster consultations, the APNIC Member Meeting (AMM), and a social event.

For the latest programme information, see:

http://www.apnic.net/meetings/21/programme

Remote participation

As with previous meetings, APNIC will provide a range of remote participation facilities for those unable to attend the meeting in person. Those with an interest will be able to follow events at APNIC 21 and participate in real time via video and audio streaming, online transcripts, and live chat rooms.

For more information on APNIC's remote participation facilities, and how they can enhance your meeting experience, see:

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

Become an APNIC 21 sponsor

Organisations throughout the region can play an important role in the APNIC meeting by becoming a sponsor. Sponsors will be presented with valuable opportunities to expose their organisation, products, and services to an international audience of Internet leaders, with approximately 200 delegates from the region and around the world expected to attend APNIC 21.

By becoming a sponsor, you help to:

- Reduce the financial burden on members attending the Open Policy Meeting through monetary contributions or in-kind support in areas such as meeting rooms, equipment, Internet connection, social activities, and meals
- Foster stronger, more supportive mutual relationships between members, as well as non-member organisations, and enhance the opportunities for effective communication and sharing of experience
- Provide opportunities for fellows to meet and network with their peers, gain valuable experience, and broaden their contacts with key people in the Asia Pacific Internet community.

For more information on how to become a sponsor, and the many benefits, see:

http://www.apnic.net/meetings/21/sponsors

Executive Council election

An important event at APNIC meetings is the election of community members to the APNIC Executive Council. Three positions on the APNIC Executive Council (EC) will become vacant in March 2006, and an open election to fill these vacancies will be held at APNIC 21 on Friday 3 March 2006.

Nominations for these positions are due by close of business Friday 17 February 2006. Nominees do not have to be representatives of APNIC members; however, only APNIC members may make nominations. Members are welcome to nominate a representative of their own organisation.

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

http://www.apnic.net/meetings/21/ec/nomination.html

If you are not able to attend the AMM, you can still make your vote count by appointing a proxy to represent you at the meeting. To do this, please print the proxy form and fax a signed copy to APNIC. Proxy forms must be received by Wednesday 1 March 2006, and can be downloaded from:

http://www.apnic.net/meetings/proxy.html

APNIC members can also cast their votes via the online voting facility in MyAPNIC. Online voting will close on Wednesday 1 March 2006, with further details to be announced shortly.

For information on current EC members and the role of the EC please refer to the Executive Council web page, at:

http://www.apnic.net/ec

More information

Regular meeting updates will be sent to the apnic-announce mailing list over the coming months.

Please send any meeting related enquiries to

meetings@apnic.net

▲ Perth, Western Australia is renowned for its beaches and outdoor lifestyles, as well as the fine wineries in the surrounding region.

New root server deployments for Bangladesh and Pakistan

Bangladesh and Pakistan have now joined the list of root server hosts, following two new deployments of F-root mirrors in December. The Dhaka server went live on 8 December and was followed on 14 December by the launch of the Karachi server.

These are the first root server deployments in each country and are expected to bring significant improvements in speed and reliability to Internet users in Bangladesh, Pakistan, and the surrounding region.

The number of root server mirrors in the Asia Pacific region has grown rapidly in recent years, to a large extent driven by APNIC's efforts in coordinating deployments with root server operators and local ISP representatives. These recent deployments bring the number of root DNS servers in the Asia Pacific region to 26, 18 of which have been made possible with APNIC's support.

The F-root is operated by Internet Systems Consortium (ISC), a nonprofit company based in California, which supports Internet infrastructure by developing and maintaining software, protocols, and operations. The F-root now has presence in 35 locations around the world, 12 of which are in the Asia Pacific.

APNIC coordinated the Bangladesh deployment with the Bangladesh Internet Exchange (BDIX), with support from SDNP Bangladesh and the Internet Service Provider Association of Bangladesh (ISPAB).

ISPAB President Md. Akhtaruzzaman Manju welcomed the deployment, explaining that as "root servers are the backbone of the Internet Domain Name System (DNS), it is vital that root servers are spread throughout the world to maintain high availability".

In Pakistan, the deployment was made possible with the support from Cyber Internet Services, the largest Internet and data communication network service provider in Pakistan, and the Internet Service Provider Association of Pakistan (ISPAK).

Root servers are a critical part of the Internet's domain name system (DNS), providing information about authoritative servers for the many top-level domains (such as ".COM", ".ORG", ".BD", and ".PK"). Computers need this DNS information to interpret URLs, email addresses, and perform many other types of Internet transactions.

▲ F-root deployments in Dhaka and Pakistan bring the number of root DNS servers in the Asia Pacific region to 26, 18 of which have been made possible with APNIC's support.

index

WSIS Internet governance timeline

1998

The ITU meets in Minneapolis and resolves to hold WSIS.

2003

Internet governance emerges as one of the major issues at WSIS Phase I, Geneva.

2004

Working Group on Internet Governance (WGIG) is established to help guide Internet governance discussions at WSIS.

2005

WGIG releases its report in mid-2005.

Late on the evening before WSIS Phase II begins, governments agree on watered-down Internet governance resolutions. The *Tunis Agenda* and the *Tunis Commitment* are subsequently endorsed by 174 States at WSIS Phase II.

2006

First Internet Governance Forum (IGF) to be held in Athens, Greece. The role and effectiveness of the IGF to be reviewed within five years of its creation.

Internet governance – A quick look at the future-

Next six months

No changes to current Internet governance systems.

One to five years

A non-binding, multistakeholder Internet Governance Forum (IGF) will be established

After five years

Discussions at the IGF and at other forums may lead to further changes in Internet governance systems and the long term development of an "enhanced cooperation model".

To read all WSIS recommendations on Internet governance, see paragraphs 29 to 82 of the Tunis Agenda for the Information Society at:

http://www.itu.int/wsis/docs2/tunis/off/6rev1.html

Internet Governance Forum

The most detailed Internet governance recommendation outlined in the *Tunis Agenda* is the formation of a multistakeholder Internet Governance Forum (IGF), to be convened by the UN Secretary-General before the middle of 2006. The IGF is to discuss public policy issues related to Internet governance and facilitate discussion of issues that have not yet found a home elsewhere.

While the detail remains to be seen, the IGF could be a productive way for governments, civil society, the private sector, and international organisations to make progress on Internet issues that cut across stakeholder boundaries. The *Tunis Agenda* is careful to state that existing structures and processes of Internet governance will be used by the IGF, not replaced. Since the WSIS process began, some of the Internet's established stakeholders have been concerned about dilution of existing bottom-up processes in future Internet governance systems. The *Agenda*'s assurance that the IGF is a non-threatening forum, in which all can contribute and grow, may encourage existing stakeholders to contribute openly.

A number of speakers in the programme of Parallel Events at WSIS made the observation that, between the two phases of WSIS, a greater dialogue had developed between the many stakeholders in Internet governance. This was demonstrated at WSIS Phase II, where government delegations actively sought out the opinions of non-government participants to gain a broader understanding of issues. For example, the Civil Society's Internet governance caucus was asked for its opinion by a number of governance delegations during last minute PrepCom-3 discussions. In addition, many speakers at the Parallel Events programme observed that there had been a substantial rise in the breadth and quality of understanding of Internet governance issues by various stakeholder groups. If the IGF is able to take advantage of better-informed and more active stakeholders, the forum may ultimately lead to truly responsive and cooperative Internet governance systems.

However, it is also possible that the IGF will not produce any positive concrete outcomes. The *Tunis Agenda* makes it clear that the IGF is to be an advisory body only, with no power to enforce any recommendations it makes. Since the roles of the stakeholders in the forum are not clear in the *Agenda*, it is possible that the forum will have a similar format to the PrepCom, where civil society and the private sector were often relegated to observer status and only official state delegations had real input into the drafting of the outcomes. If this is the case, the IGF may fall victim to wider international politics, preventing anything of real substance coming out of the forum.

Such nation-based politicking was evident at WSIS, where the Internet governance statements in the *Tunis Agenda* were hailed as a triumph, although no concrete targets for Internet governance were agreed upon. Instead, the difficulty of overcoming political differences between the 174 national delegations at WSIS meant it was an achievement simply to agree to keep discussing Internet governance.

The Tunis Agenda and technical Internet operations

The *Tunis Agenda* divides Internet governance into two main areas: 1) public policy, which is the main focus of the *Agenda* and the primary responsibility of governments, and 2) day-to-day technical operation of the Internet, which it leaves largely in the hands of the private and civil sectors. On first impressions, the Internet's technical community may see this as a sign that it can continue its operations in the knowledge it will not be hindered by government involvement. The *Agenda* certainly has been interpreted by many to mean that ICANN has now finally gained international approval for its role in the technical administration of the Internet.

However, on closer examination, it becomes evident that many public policy issues do have an impact on the day-to-day technical running of the Internet. For example, if the IGF were to make recommendations on stopping spam globally, and these recommendations were then passed as resolutions at the UN, this potentially could lead to pressure for changes in protocols such as email. Such protocol changes would need to be developed and standardised through bodies such as the IETF, then deployed throughout the Internet.

While the Internet traditionally has operated from a bottom-up technical development process, the *Tunis Agenda* could result in future technical development being driven, instead, by top-down public policy. This may have significant implications.

A complaint sometimes expressed about technical bodies such as ICANN, the IETF, and the RIRs is that these bodies have historically avoided becoming involved in finding solutions for major global problems, such as spam, by stating that such issues are outside the limited technical scope of the organisations. While participants of technical bodies understand there is a need to conserve the bodies' limited resources, it has been more difficult for the world's non-technical majority to understand why such bodies cannot solve problems that affect most Internet users.

A top-down public policy approach, as recommended by the *Tunis Agenda*, combined with a bottom-up technical implementation could perhaps result in a more robust

Internet; Internet protocols may be more effective at both the level of network administration and at the level of global security and usability. On the other hand, many fear that top-down enforcement of non-technical concerns could lead to a politicisation of Internet's core technologies, to the detriment of network health.

The role of the technical community in future Internet governance discussions

During the first phase of WSIS, some elements of the technical Internet community did not play a large role in the discussions. This was partly because it was not yet clear how prominent Internet governance discussions would be at WSIS and partly because the technical Internet community did not immediately foresee the full potential for WSIS discussions to impact on operational issues. Representatives from organisations such as APNIC and ICANN were present at WSIS Phase I, but attended more as observers than active participants. However, by the second phase of WSIS, there was greater participation by the technical community and, in fact, many from the community were registered as part of official government delegations. LACNIC CEO and NRO EC member, Raul Echeberría was an advisor to the Uruguay delegation. In addition, a number of Internet organisations, including the Number Resource Organization (NRO), the Internet Society (ISOC), and ICANN, worked together on the Internet Pavilion, a stand at the WSIS side event, the ICT4all exhibition. At this stand, members of the Internet's technical community were available throughout WSIS to answer technical questions from WSIS attendees. The Internet Pavilion was visited by a number of representatives from government delegations and the world's media, as well as from civil society and the private sector.

In the post-WSIS world, it is important that the technical Internet community continues to play an active role in Internet governance both the day-to-day Internet operations and in the development of public policy. However, to do this, the technical community needs to continue to learn the ways of diplomacy. Traditionally, the technical community has placed a lot of value on establishing the knowledge level and worthiness of new entrants to the community before engaging them in meaningful discussion. To a large degree, this attitude changed during the WSIS process as the technical community began to engage with less tech-savvy stakeholders in Internet governance. This newer, more inclusive approach continues to be essential in forums such as the IGF. The technical Internet community must continue to actively work to educate non-technical stakeholders about technical issues and to engage in public policy discussions. Otherwise, silence from the technical community may be mistaken for tacit approval. It is not the governments' responsibility to learn the intricacies of the technical operations of the Internet, but the technical Internet community's responsibility to help governments understand how their public policy interacts with the technical running of the Internet.

It is also important that the technical community understand the framework within which future Internet governance will develop. While the IGF may be a useful venue for airing important Internet governance issues, the IGF itself will be subject to the higher-level political intrigues. For example, since the US Government has made its distrust of the UN and the ITU well known, if the IGF were ever to recommend a move to a centralised UN-based Internet governance system, it would be very hard to achieve even if the rest of the world was in favour of it. To the technical Internet community, whose main desire is to get on with the job of keeping the Internet running, such high-level political wrangling may seem pointless. But it will not go away and it will be important for all stakeholders to build an understanding of how to work within this paradigm.

In summary, while WSIS has not resulted in any concrete changes to future Internet governance, it has signalled that the Internet governance discussions have finally matured, that the stakeholders now have a greater understanding of the issues, and that Internet governance issues are so complex, they cannot be resolved overnight. Internet governance discussions will continue, probably without any major structural changes to Internet governance bodies, for at least the next five years. The Internet governance changes that develop at the end of that time will be dependent on what the stakeholders contribute to forums such as the IGF. Therefore, it is important that technical Internet community continues to engage in Internet governance discussions. The technical community can do this on many levels: lobbying governments, participating as everyday citizens within civil society, working with the business community, as well as continuing to work within specific technical frameworks in organisations such as APNIC, ISOC, and the IETF.

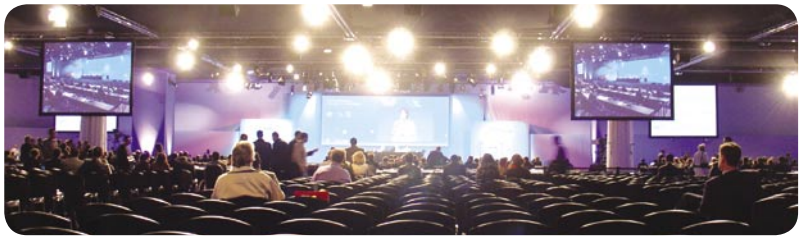

▲ WSIS Phase II was held in the Kram Palexpo conference centre in Tunis.

APNIC by numbers

Fast facts about APNIC and the region

19	APNIC staff speak a total of 19 different languages
100%	100% of the APNIC Hostmasters have a technical background and more than half have direct experience working at an ISP
56	The APNIC region covers 56 economies, with members in 47 of those economies
44	You can chat to APNIC Hostmasters via the online helpdesk chat 44 hours a week
18	APNIC has facilitated the establishment of 18 rootservers in the AP region
60	APNIC meetings are attended by people from more than 60 different nationalities
12	APNIC meetings have been hosted in 12 different locations in the region
21	The 21st APNIC meeting meeting will be held in Perth, Australia!

APNIC staff languages

Bahasa Indonesia	Malay
Bangla	Mandarin
Cantonese	Persian (Farsi)
Filipino (Tagalog)	Punjabi
Fijian	Singhalese
French	Tamil
Hindi	Thai
Japanese	Telugu
Korean	Vietnamese
Lao	

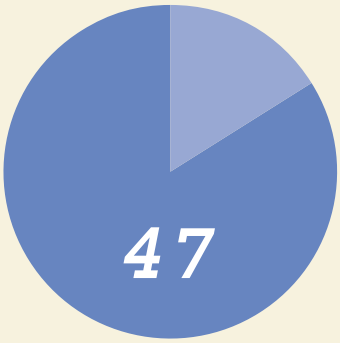

47 of the 56 economies in the region have APNIC members

The transition to 4-byte AS numbers

In the previous issue of *Apster*, **Geoff Huston** examined the consumption of 2-byte AS numbers, arguing that it would be prudent to begin a transition to larger AS number pool within the next three years. In this article, he describes the proposed 4-byte AS number space and suggests how the transition may take place.

Note: Readers may access the complete version of this article at http://potaroo.net/ispcol/2005-08/as.html

As discussed in the previous issue of Apster, current experience suggests that the unallocated 2-byte AS number pool could become exhausted by late 2010. Working backward from this date to the necessary steps that could ensure smooth transition to a new AS number pool, it would appear that we should start the transition in the coming months rather than in the coming years.

In this part of the article we'll look at the current proposal for a larger AS number pool within the BGP protocol and examine the implications of an associated transition plan.

The approach proposed in "draft-ietf-idr-as4bytes-10.txt" is to expand the size of the AS number pool space from 16 to 32 bits, expanding the number space from a pool of 65,536 to 4,294,967,296 billion numbers. In terms of the current use of AS numbers, the current scaling properties of the BGP routing protocol, and the use of ASs in the context of inter-domain routing, a pool of 4.4 billion numbers would easily encompass a network environment of significantly greater levels of domains and inter-domain interconnection density. Such a pool size would exceed some current guesses of the scaling capabilities of the BGP protocol by up to a further two orders of magnitude.

Its also proposed to preserve the first block of 4-byte AS numbers to align with the allocations of the 2-byte numbers.

We can use a new form of terminology here for 4-byte AS number values, where the first 65,536 AS numbers use the form "0:0" through to "0:65535". The second set of 65,536 numbers would be written as 1:0 through to 1:65535, and so on. So we'll be using a number format of <upper16 bits>:<lower 16 bits>.

So, what is the inventory of issues that need to be specifically addressed here?

Obviously there is a need for some changes to the routing protocol and this change needs to be able to accommodate a number of situations. It would be unrealistic to expect an ordered inter-domain transition. A more expectation is the piecemeal transition of domains, where individual domains will shift to supporting 4-byte ASs in their own time. Domains that are currently using 2-byte ASs may have less reason to undergo an early transition to 4-byte AS support, while those domains which are assigned a non-mappable 4-byte AS number will find that they have to support 4-byte AS numbers from the outset.

A piecemeal transition raises the potential of loops between "OLD" and "NEW" domains (see Figure 1). Any proposed solution should be able to detect such loops without having to alter the behaviour of the old BGP speakers.

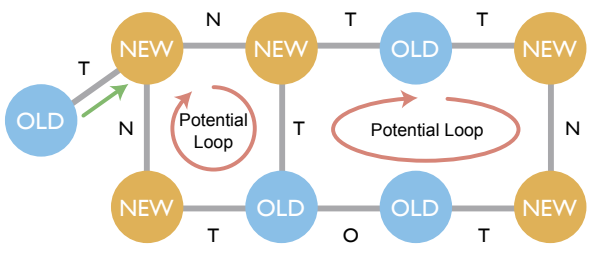

Figure 1 - BGP transition cases

Changes to the BGP protocol

BGP has two major parts within its protocol: opening a BGP conversation with a peer BGP speaker, then transfering protocol objects that describe reachability of address prefixes and associated attributes of these address prefixes. Both parts include AS number components and, in considering changes to the current protocol, both parts of the protocol require some change. The message objects that need to be considered are the BGP OPEN message and the BGP UPDATE message.

The changes to the BGP protocol create a new BGP implementation that is capable of supporting a 4-byte AS number environment. The essential task of the changes is to define mechanisms that all NEW BGP speakers use to speak to each other and pass AS number values in 4-byte fields. However the Internet is way too large to set up a "flag day", for all BGP speakers to switch from OLD BGP to NEW BGP. Accordingly, its also necessary to define protocol interactions in NEW BGP where the transition in the Internet will be gradual and essentially uncoordinated. NEW BGP speakers will have to set up sessions with OLD BGP speakers and, of course, OLD BGP speakers will also be peering with other OLD BGP speakers. The information associated with 4-byte AS paths must be passed across sections of the network that normally support only 2-byte AS paths. In other words, 4-Byte AS information needs to be passed to OLD BGP speakers and between OLD BGP speakers.

Opening a BGP session

BGP carries its own AS number in the "My Autonomous System" field of the BGP OPEN message.

The proposed approach is to initiate a NEW BGP session in a mode that is compatible with the OLD BGP protocol and also inform the remote peer of its capability to conduct a NEW BGP conversation if the remote peer is also a NEW BGP speaker. This capability advertisement is part of OLD BGP – OLD BGP speakers which open a peer session with a NEW BGP speaker will simply ignore the NEW capability and operate in OLD mode. A NEW BGP peer will respond positively to the NEW capability, and the BGP session can then operate in NEW mode.

The BGP OPEN message includes a fixed length 2-byte "My AS field" (as shown in Figure 2) as well as potentially containing a capability query as part of the Optional Parameters section. In order to ensure that NEW and OLD speakers can communicate, this 2-byte MyAS field needs to be preserved in NEW BGP even when the Optional Parameters section encompasses the capability to undertake a NEW peering session. This may appear contradictory in the first instance, as the OPEN message then contains both a 2-byte Autonomous System number and a 4-byte AS Capabilities Query.

The mechanism proposed for the OPEN Message varies according to whether the NEW speaker is using a mappable AS number drawn from the original pool (that is, with a My AS number in the range 0:0 through to 0:65535), or its using a number drawn from a higher-numbered 4-byte number block. In the first case the OPEN message would use the 2-byte mapped value in the My AS field (dropping out the zero-valued high order 16 bits of the AS value), while in the second case the BGP speaker would use for My AS a special 2-byte value that is reserved for this purpose (AS 23456). In both cases, the Optional Parameter section would include a capability code to indicate that the local BGP speaker can support 4-byte AS numbers (Capability Code 65).

The side effect is that in the OLD BGP domains AS 23456 may appear to be connected to the 2-byte BGP realm in many different locations, and advertising a collection of different address prefixes in different locations. From the OLD BGP realm this does not present a protocol problem; however, as always, there is the potential that this repeated use of AS 23456 as a 4-byte AS substitution token may create a somewhat confusing BGP-view of the Internet from the perspective of the OLD BGP world!

```
0                   1                   2                   3
0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1
+-+-+-+-+-+-+-+-+
|    Version    |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
|     My Autonomous System      |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
|           Hold Time           |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
|                         BGP Identifier                        |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
| Opt Parm Len  |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
|             Optional Parameters (variable)                    |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
```

Figure 2 - BGP Open Protocol Message – From "draft-ietf-idr-bgp4-26.txt"

The capability exchange uses a protocol described in RFC3392. The NEW BGP speaker adds an optional capability field to the OPEN message. The 4-byte AS capability code 65 carries as its capability value the local 4-byte local AS number value. For a NEW peer this capability value is to be interpreted as the actual AS of the remote side, on the basis that the MyAS field in the body of the OPEN is either a truncation of the local 4-byte AS value (in the case of mappable 4-byte AS values), or the special value of AS 23456.

One response from the remote BGP speaker is to accept the capabilities announcement with a comparable OPEN message – in which case the remote side is also a NEW BGP speaker – and the session may proceed using 4-byte AS values.

If the session is opened with an OLD BGP peer, the OLD BGP peer may respond with a NOTIFICATION message indicating that the 4-byte capability is an Unsupported Optional Capability parameter. In response to this unsupported notification the NEW BGP speaker will re-establish the connection by resending the OPEN message, and this time drop the 4-byte capability option from the message. The NEW BGP speaker will then assume that it is peering with an OLD BGP peer.

The "Unsupported" response to a capabilities parameter was not included in the original specification. Older versions of BGP allowed a BGP speaker to optionally send a NOTIFICATION message and terminate the peer session. If the NEW BGP speaker sees a session termination in response to its OPEN message it may need to re-open the TCP session, this time omitting the 4-byte capability advertisement in the initial BGP OPEN message. Once again, the NEW BGP speaker will then assume that it is peering with an OLD BGP peer.

In general, however, a BGP implementation should not send a NOTIFICATION when a capability parameter is unrecognised because the Capabilities Optional Parameter is still optional. With such general implementations, the OLD speaker would just pick up the 2-byte AS (23456) in the OPEN received form the NEW speaker. As the OLD speaker does not advertise the 4-byte AS Capability in its OPEN, the NEW speaker has to use the 2-byte AS it advertised in the OPEN (that is, the AS_TRAN - 23456) for peering. A NOTIFICATION is not involved in this scenario.

The BGP UPDATE Message

For a NEW BGP session (4-byte peering with 4-byte) the changes to the protocol are the use of 4-byte AS numbers in the AS_PATH attribute of UPDATE messages. All 2-byte AS values are padded with a zero high order 16 bits. If the AGGREGATOR attribute is used it is similarly carried as a 4-byte value. So in the 4-byte peering, all 2-byte information is carried in mapped 4-byte AS numbers (see Figure 3).

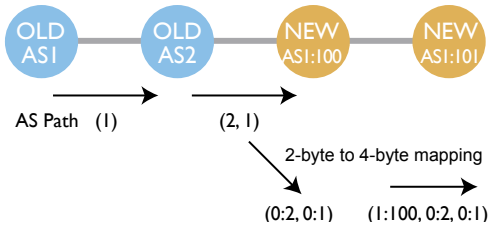

Figure 3 - OLD to NEW BGP AS path Mapping

In this way, AS path length is preserved without change when translating 2-byte AS information into the 4-byte domain.

The next case is where an OLD BGP peers with a NEW BGP. We've already seen the simple case where the information is coming from a 2-byte path and there is no additional 4-byte information, and in this case the 2-byte values are simply mapped into 4-byte values. What about the reverse case where 4-byte information is being passed back into the 2-byte world?

There are two parts to this case: first creating an equivalent 2-byte AS path and second packing up the 4-byte AS path information in such a way that it transits across the 2-byte domain in such a way that it can be reassembled in any subsequent transition into a 4-byte domain. In the first case, the equivalent path information is constructed by either stripping the high order 2-bytes from the AS value, as long as this part is all zeros. Where this is not possible, the transition AS number, 23456, is substituted in its place. This is indicated in Figure 4.

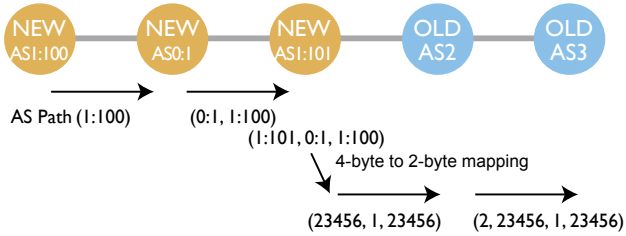

Figure 4 - NEW to OLD BGP AS path mapping

In this way, the AS path length metric is preserved, and the prevention of count-to-infinity loops in the 2-byte domain is avoided.

The second part to this case is packaging up the 4-byte path into the OLD BGP session in such a way that it can be unpacked at any subsequent boundary into a 4-byte realm. Here the proposal calls for new transitive community attributes to be supported for OLD BGP. These attributes are defined as transitive attributes, and should be passed through the OLD BGP peering sessions without alteration. It should be noted that this is not a protocol change, per se, but it does require the explicit support within

OLD BGP implementations of this attribute as a transitive community.

The proposed mechanism is an extended community attribute called “NEW_AS_PATH”. When a NEW BGP speaker is speaking to an OLD BGP, the NEW BGP prepends its own AS value to the AS path and copies this information into the NEW_AS_PATH. It then translates the 4-byte AS path into a 2-byte equivalent AS path. The translation is straightforward, in that where the 4-byte AS has all zeros in the high order 2 bytes, the translation truncates the AS value to a 2-byte value, and where the high order 2-bytes are non-zero the translation substitutes the reserved 2-byte value AS 23456 in its place (see Figure 5).

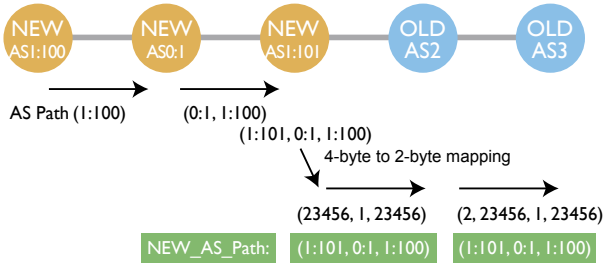

Figure 5 – NEW to OLD BGP AS path mapping

The transit across the OLD BGP domains leaves the NEW_AS_PATH untouched, and prepends 2-byte AS values to the AS_PATH.

The next transition is one from the OLD to the NEW domain, as shown by a continuation of the previous example (see Figure 6).

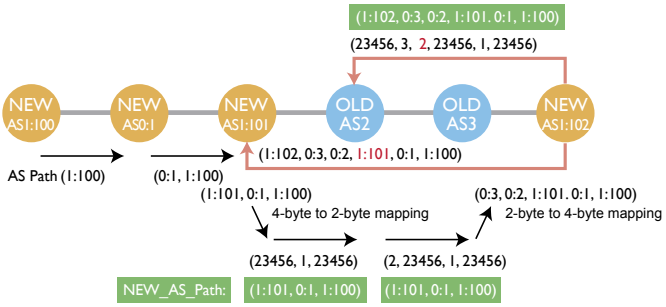

Figure 6 – NEW to OLD to NEW transition with potential routing loops

Figure 6 shows a further OLD to NEW transition. In this case the NEW BGP speaker takes the AS Path as presented by the OLD BGP speaker and converts the 2-byte values to 4-byte values by adding 2-bytes of zero padding to each entry, and then overwrites the trailing entries with the values specified by the NEW_AS_PATH attribute. The net result is that the 4-byte path that entered the 2-byte sequence is prepended with the 2-byte transit AS sequence. The NEW_AS_PATH is then removed, leaving an intact 4-byte path as the AS_PATH attribute.

This ensures that the resultant BGP environment can detect loops in both the NEW 4-byte and OLD 2-byte realms.

Further extending this example, we can construct a potential loop in the 4-byte world by adding a path back to AS 1:101. Restoring the original 4-byte AS path at the OLD-to-NEW transition ensures that the potential loop is discarded even when the loop needs to traverse one or more 2-byte OLD BGP ASs. A similar form of loop can be constructed for a 2-byte OLD BGP AS, that traverses a 4-byte NEW BGP AS. Again the transition mapping ensures that the potential routing loop is detected by BGP.

The ability to perform AS Path Prepending is also unaltered in this mixed NEW and OLD BGP environment. The AS simply prepends its local AS value to the AS_PATH as normal. In the case of prepending on a NEW-to-OLD boundary the prepended AS Path is mapped into the NEW_AS_Path attribute as described above.

In a less common use of AS PATH poisoning, the prepending uses a different AS number value in order to ensure that the particular advertisement is not learned by a remote AS. For NEW BGP speakers there is no change to this capability. For OLD BGP speakers the AS Path poisoning can only be directed towards 2-byte ASs, as the OLD BGP speaker has no knowledge of the structure or content of the NEW AS PATH attribute. From the perspective of the OLD BGP speaker, the NEW_AS_PATH attribute is an opaque data block.

The same translation technique applies to the AGGREGATOR attribute. In a NEW-to-OLD transition the AGGREGATOR may be a mappable AS number, in which case the value is truncated to 2-bytes and no further action is required. Otherwise, the 4-byte AGGREGATOR value is rewritten to the NEW_AGGREGATOR attribute and the transition 2-byte value, AS 2356 is placed into the AGGREGATOR attribute. On an OLD-to-NEW transition the NEW_AGGREGATOR attribute is copied back into the AGGREGATOR attribute, if defined, otherwise the AGGREGATOR is padded out with leading zeros.

The general approach adopted for transition is to preserve AS Path length information across the OLD and NEW BGP boundaries, while recognising that some 4-byte AS information cannot be cleanly mapped into a 2-byte AS Path. In order to preserve 4-byte information, which is necessary to prevent loop formation for 4-byte ASs, the 4-byte information is preserved across OLD transit paths and restored upon re-entry into NEW BGP realms.

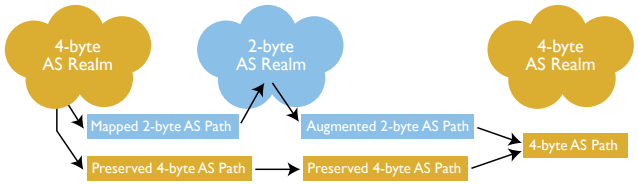

Figure 7 – 2-byte and 4-byte AS Realms

BGP communities

BGP communities require some additional consideration. If the high order 16 bits of the community attribute are neither all zeros or all ones, then it is assumed to contain a 2-byte AS value. Where it is necessary to specify a 4-byte AS number in the community attribute it is necessary to turn to the extended community attribute to support this.

This extended communities feature is documented in the Internet draft "draft-ramachandra-bgp-ext-communities-10.txt", now on the RFC publication track as a Proposed Standard.

Transition

Transition in this environment is relatively straightforward. NEW BGP speakers can be deployed within the network in a piecemeal fashion without any major concerns. The size of BGP UPDATE messages is slightly longer due to the extended length of the AS PATH attribute in NEW BGP and the NEW_AS_PATH attribute that has been added in the OLD BGP environment, but it should not prove to be a major factor.

BGP loop prevention appears to be adequately addressed in all commonly encountered situations and there appear to be no other significant transition considerations.

There does appear to be one precondition for the use of 4-byte AS numbers, and that is for a routing domain to actually be numbered with a non-mappable 4-byte AS number, all the BGP speakers in the domain should be NEW BGP speakers. Aside from that consideration there do not appear to be any further constraints associated with this transition.

We are certainly running out of the 2-byte AS number pool, and an industry of this size needs to have a considerable period of advance warning of change in order to be able to integrate such changes into various operational cycles of testing and transitional deployment prior to integration into production environments.

The 4-byte transition appears to offer flexibility, orderly transition and minimal disruptions to existing operational practices.

The first steps that need to happen are the completion of the technical specification of this approach in the form of an IETF Standard and the subsequent production and distribution of NEW BGP implementations from the existing sources of BGP implementations. It would be preferable to get this underway now, while there is still time to complete this transition well before we exhaust the current 2-byte AS number space.

Acknowledgement

Thanks to Enke Chen, one of the authors of the 4-Byte AS working document, for some clarification regarding the OPEN behaviour between OLD and NEW BGP implementations.

Certifying Internet address resources

Certification of IP addresses and AS numbers is just around the corner with the development of a new service aimed at extending X.509 certificates to provide new levels of security for Internet resources.

The early Internet emerged from an environment of implicit trust within a small community of like-minded researchers and engineers. Originally, many of the Internet's operations relied on that sense of common purpose and shared problems. As it expanded into a truly global and ubiquitous network, the sense of the Internet being a single community dissipated and the time has long since passed when anyone can take Internet security for granted. Nevertheless, in some critical aspects, relatively informal trust models still apply. Internet addressing is one key example where more modern security procedures are clearly needed.

The RIRs distribute IP addresses and AS numbers in a careful, responsible manner. The means by which the responsibility for these common resources can be delegated from one party to another is strictly defined by publicly-accepted policies. And, of course, the whois databases are a public record of resource status and custodianship. But when someone approaches an ISP and asserts their right to use a particular address range, how can that ISP really trust the assertion? By what easy, reliable mechanism can the ISP verify that the person is actually the legitimate custodian of those resources?

In this context, APNIC has begun work to establish an address resource certificate infrastructure that can provide a level of resource security that has not previously been possible. APNIC's resource certification service project is based on the model provided by RFC3779, which provides a mechanism for certifying IP addresses and AS numbers.

Under this project, APNIC will establish a service to issue RFC3779-compliant certificates to APNIC account holders, allowing them to make trusted assertions about their resources. To achieve this, APNIC will develop a policy and technical infrastructure to support the use of resource certificates, including a certification practice statement, a certificate repository, and a certificate revocation list.

In many respects, APNIC already has considerable experience in this field. For several years now, it has been issuing certificates to allow account holders to access the MyAPNIC secured web site. However, up until now, the only party which has needed to trust MyAPNIC certificates is APNIC itself; under this new service, it will be necessary to ensure that everyone is able to trust the resource certificates.

RFC3779 defines two extensions to the X.509 certificate format for IP addresses and AS numbers. The extensions allow a list of IP or AS resources to be incorporated. In practice, this can be used to verify that the holder of the certificate's private key has authority to use the listed resources.

The certificate model fits well within the established hierarchical resource allocation model, in which IANA delegates address management rights to RIRs, which themselves delegate responsibilities to ISPs, end entities, and other delegating bodies. This hierarchy of delegation allows for a certificate chain that is able to reflect address policy and create certainty of custodial rights.

Under the planned service, when APNIC delegates a resource, in addition to sending a confirmation email and registering the delegation in whois, APNIC will also send a digital certificate containing the latest list of resources that have been delegated to that account holder. The certificate will conform to the X.509 standard and will include the resource information extensions as defined in RFC3779.

The certificate format itself is neutral to any specific protocol. Each certificate will contain a simple statement that the certificate subject (who holds the private key corresponding to the certificate) has received custodianship of the listed resources. It is important to note that the scope of these certificates is restricted to resource and routing attestations. APNIC resource certificates are not intended to be used to certify websites, authenticate web clients, or serve any other general purpose.

Work is currently underway to develop protocols and tools for working with resource certificates. The tools will deal with procedures such as requesting certificates, issuing downstream certificates, and validating certificates. Any tools developed as part of this project will be open source and made available to the general community.

The first stage of APNIC's resource certification service project began in late 2005. This initial trial involves work with a small number of early adopters and consultations with software developers and router designers. This will be followed by a pilot program offered in the first quarter of 2006, which will use input from the trial to refine the service for wider deployment. It is expected that by the second quarter of 2006, all aspects of the certificate procedures should be stable enough for full service.

Progress on the resource certification service project will be reported in the Routing SIG at APNIC 21.

APNIC 20 project presentation archive is available at:

http://www.apnic.net/meetings/20/programme/sigs/routing.html

RFC3779, by C. Lynn, S. Kent, and K. Seo is available at:

http://www.ietf.org/rfc/rfc3779.txt

APNIC multimedia projects

Over the past year, APNIC Secretariat staff have created a range of multimedia presentations, including video and Flash animations. These presentations are designed to serve as educational tools, and focus on a variety of different subjects of relevance to APNIC members and the Internet community in general.

At this stage, the following presentations are available for viewing on the APNIC website:

Video

Inside the APNIC Open Policy Meeting
A video introduction to APNIC Open Policy Meetings, featuring interviews with prominent members of the community, SIG Chairs, and Secretariat staff.

Flash presentations

What is APNIC?

A general introduction to APNIC, its role in the Internet, and the services it offers.

The history of the RIRs

A brief history of Internet addressing and the development of the Regional Internet Registry (RIR) system. Discusses the role of the RIRs in the Internet community.

Policy development

An overview of the APNIC policy development process, looking at why policies are necessary, how they are agreed upon, and how individuals and organisations can become involved in the policy process.

The Number Resource Organization

A look at the Number Resource Organization, or NRO, a representative body of the five RIRs. Includes a discussion of the body's history and current role in the Internet community.

MyAPNIC demonstration

An overview of MyAPNIC, APNIC's online user interface system for members and account holders. Includes information on how to access MyAPNIC and the range of features available to users.

APNIC Outreach Coordinator Nurani Nimpuno described the project as "a significant addition to APNIC's information library".

"Issues that are complicated to understand in a basic text format can sometimes be made simpler by using graphics and animations," she said.

Training Manager John H'ng noted that some of the new multimedia resources "have been incorporated into APNIC training courses and presentations and the feedback we have received has been very encouraging. Multimedia materials allow you to be a little bit more creative in how you convey information and we are very excited about exploring these forms further."

As well as being used in training courses, some of the presentations have also been used in international forums, including the NRO display at the recent WSIS meeting in Tunis. Plans are also going ahead to expand the range of multimedia materials available.

"We are currently developing a suite of Flash animations that will be available on our website, covering the broader operations of APNIC," said Nurani. "We are also looking at what other areas would benefit from video, Flash, or other graphical formats. We have also begun to develop an APNIC multimedia CD as a resource to the community."

The full range of APNIC multimedia presentations can be accessed at:

http://www.apnic.net/multimedia

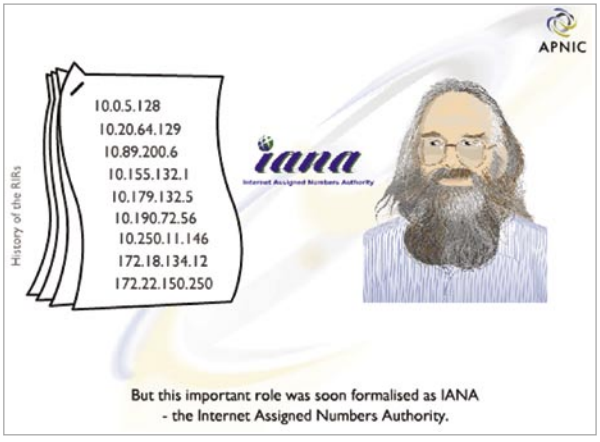

New staff

Communications Department

Holly Qi
Marketing Communications Officer

Holly Qi has recently joined the APNIC Outreach team as Marketing Communications Officer. Holly has over 10 years of experience working as a marketing specialist, and has worked for companies such as Dell and Nortel Networks in China, as well as with several companies in Australia, including The Hear and Say Centre (a not for profit organisation). She recently graduated with a Master of Business Adminstration from the University of Queensland.

In her role at APNIC, Holly is responsible for planning marketing and communications activities for the organisation.

Training Department

Sall'ee Ryman
E-learning Development / Training Officer

Sall'ee Ryman is the newest member of the APNIC Training team, filling the new role of E-learning Development/Training Officer. Sall'ee brings a wealth of experience to the position, including degrees in both Media Production and Education. Her previous work has seen her involved in a range of media organisations, as well as Education Queensland, during which time she received an Excellence in Teaching award for projects in e-learning.

Sall'ee is responsible for developing and delivering training to APNIC members and the wider Asia Pacific Internet community, and for helping to expand the range of training services offered by APNIC.

Technical Services Department

Robert Bailey
System Administrator

Robert Bailey joined the APNIC Technical team in November, having worked in various positions with the NSW Department of Education and Training. He is a graduate in Information Technology (Network Engineering) from the Macquarie Fields College of TAFE, and in his role at APNIC will be focused on internal customer support, as well as assisting with general IT infrastructure support and projects as required.

Visiting staff

Technical Services Department

Frank Nnebe
AfriNIC

AfriNIC's Senior Software Engineer, Frank Nnebe, visited the APNIC office at the beginning of November. He spent most of his time with the APNIC Technical team, learning about MyAPNIC and the overall APNIC resource management system. The visit was a great opportunity for both organisations to learn from each other and will provide a good base for software development collaboration in the future.

Training schedule

2006

January

- **16 - 25** Mumbai, India

February

- **1** Melbourne, Australia
- **22 - 3 March** Perth, Australia (In conjunction with APNIC 21 / APRICOT 06)

March

- **22 - 24** Wellington, New Zealand (In conjunction with NZNOG 06)
- **27 - 30** Manila, Philippines

April

- **7** Guam
- **25 - 26** China (Venue TBA)

May

- **2 - 5** Bangkok, Thailand
- **26** Brisbane, Australia

June

- **TBA** Jakarta, Indonesia
- **19 - 23** PACNOG 2
- **TBA** Japan (Venue TBA)

July

- **TBA** India
- **TBA** Islamabad, Pakistan
- **27 - Aug 4** Karachi, Pakistan (In conjunction with SANOG 8)

August

- **TBA** Mongolia
- **21 - 26** PICISOC (Venue TBA)

September

- **TBA** APNIC 22
- **TBA** Vietnam
- **TBA** Laos
- **TBA** Cambodia
- **TBA** Hong Kong

Octomber

- **9 -13** Bangkok, Thailand
- **16 - 20** Colombo, Sri Lanka

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

www.apnic.net/training

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

training@apnic.net

Calendar

PITA Meeting
14 January 2006
Honolulu, USA
http://www.pita.org.fj

PTC '06
15-18 January 2006
Honolulu, USA
http://www.ptc06.org

ICOIN: International Conference on Information Networking 2006
16-19 January 2006
Sendai, Japan
http://www.icoin.org

JANOG 17
16-19 January 2006
Sendai, Japan
http://www.janog.gr.jp

SANOG 7
16-24 January 2006
Mumbai, India
http://www.sanog.org

21st APAN Meeting
22-26 January 2006
Tokyo, Japan
http://apan.net/meetings/future.htm

NANOG 36
12-14 February 2006
Dallas, USA
http://www.nanog.org/future.html

APNIC 21 / APRICOT 2006
22 February - 3 March 2006
Perth, Australia
http://www.apnic.net/meetings

65th IETF
19-24 March 2006
Dallas, USA
http://www.ietf.org/meetings/meetings.htm

NZNOG 06
22-24 March 2006
Wellington, New Zealand
http://www.nznog.org

ICANN Meeting
27-31 March 2006
Wellington, New Zealand
http://www.icann.org/meetings

ARIN XVII
9-12 April 2006
Montreal, Canada
http://arin.net/meetings

RIPE 52
24-28 April 2006
Istanbul, Turkey
http://ripe.net/ripe/meetings/current.html

AfNOG
7-15 May 2006
Nairobi, Kenya
http://www.afnog.org/afnog2006

AfriNIC 4
16-17 May 2006
Nairobi, Kenya
http://www.afrinic.net/meeting

LACNIC IX
22-16 May 2006
TBD
http://lacnic.net/en/eventos

How to contact APNIC

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

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

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

Member Services Helpdesk

helpdesk@apnic.net

www.apnic.net/helpdesk

☎ +61 7 3858 3188

+61 7 3858 3199

Communicate with APNIC via MyAPNIC

APNIC members can use MyAPNIC to:

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

For more information on MyAPNIC's features, see:

www.apnic.net/services/myapnic