

August 2009

Geoff Huston

AS Numbers - Again

IPv6 is not the only number resource that is running out in the coming couple of years. The same fate awaits the pool of Autonomous System numbers, used to support the inter-domain routing protocol, BGP. In the original design of BGP. A previous article in August 2005 (http://www.potaroo.net/ispcol/2005-08/as.html) explored AS numbers in some detail, so in this article I'd like to update the situation and look at we are doing about exhaustion some four years later.

Current Status of the AS Number Pool

There are 65536 AS numbers in the 16 bit number pool. 1,041 numbers are reserved by the IETF for private use and documentation purposes, and one number, AS 23456, is reserved for protocol use. Of the remaining 64,494 AS numbers, As of the end of August 2009, some 56,318 AS numbers have been passed into the system managed by the Regional Internet Registries, and 8,176 AS numbers remain with the IANA.



Figure 1 - 16-bit AS Number Pool Status (August 2009)

If we look a little closer at these numbers, using the BGP routing table, as of the end of August 2009 some 31,969 AS numbers are advertised in BGP and 15,019 are allocated, but not visible as advertised, This data can be used to construct a map of the entire AS Number space.





Each column represents 256 AS Numbers. Each column shows the number of AS numbers visible in BGP (blue), the number allocated but not routed (red), the number held by an RIR in an unallocated pool (green), the number held by the IANA in the IANA unallocated pool (cream) and the number reserved by the IETF (grey).

16-bit AS Number Pool Exhaustion

So the question of the moment is that if we've managed to get through 46,988 AS numbers in August 2009, how long will the remaining pool of unallocated 16-bit AS numbers last?

Of course this is a question that we should never have had to ask. AS Number exhaustion was never planned to occur in this manner. The thinking in the early 1990s was that BGP was a simple short term approach to inter-domain routing that was more of an interim solution than a long term foundation for a massive Internet. The planned successor protocol to BGP, IDRP, or the Inter-Domain Routing Protocol, was a protocol that was being developed within an intended general framework for router that was intended to encompass more than just IP. The expectation at this time is perhaps best illustrated in RFC 1930, " Guidelines for creation, selection, and registration of an Autonomous System (AS)", published in March 1996.

9. AS Space exhaustion

The AS number space is a finite amount of address space. It is currently defined as a 16 bit integer and hence limited to 65535 unique AS numbers. At the time of writing some 5,100 ASes have been allocated and a little under 600 ASes are actively routed in the global Internet. It is clear that this growth needs to be continually monitored. However, if the criteria applied above are adhered to, then there is no immediate danger of AS space exhaustion. It is expected that IDRP will be deployed before this becomes an issue. IDRP does not have a fixed limit on the size of an RDI.

RFC 1930, "Guidelines for creation of an AS" J. Hawkinson, T. Bates, March 1996.

IDRP used protocol addresses rather than a synthetic numbers to undertake the role of domain identifiers. Given that IPv4 unallocated address pool exhaustion is occurring in a similar timeframe to 16-bit AS number exhaustion, then even if BGP had elected to use 32-bit IP address assigned from the network operator's address pool in a similar manner to IDRP, then we'd be facing exactly the same problem at pretty much the same time! Indeed, IDRP would have encountered the same problem around now if it had been deployed and had elected to use an IPv4 address as its domain identifier.

However we are on a different trajectory with inter-domain routing, and IDRP languished and then disappeared over the Internet's horizon at about the same time as this RFC was published, and since then its been BGP and only BGP, until we find ourselves in the situation of looking at the prospect of imminent AS Number pool exhaustion in BGP.

So the relevant question to ask is: How imminent is this event? When will the AS Number pool exhaust?

The starting point is to look at the history of AS Number allocation and derive a model of consumption that can be used to project into the future. The available data starts at 1992, and the time series of AS number consumption is shown in Figure 3.



Figure 3 - AS Number Allocation Time Series

The method I've used to provide one possible answer to this question is to construct a model of daily AS Number usage over the past three years, and use the change in average daily

consumption levels over that period to construct a mathematical model of future consumption. The critical component of the model is the first order differential of the AS Number consumption data, or the daily rate of allocations. This is shown in Figure 4.



First order differential of allocations(smoothed)

Figure 4 - AS Number Daily Allocation Rate

While there has been some significant short term variation (or "jitter') in the daily allocation rate of AS Numbers, from a low point of 8 ASNs per day to a high of 18 ASNs per day, the long term trend is an average of some 12.5 ASNs per day, with a slight increase over this 38 month period.

This data can then be used to construct a forward extrapolation of future demand. The technique is to construct a least squares best fit to the daily AS Number allocation rate, and then integrate this linear function into a second order polynomial to create a model of AS allocation. This model can be coupled with the prevailing policies of AS Number allocation from IANA to the RIRs in order to construct a complete model of projected AS number consumption, a graphical representation of which is shown in Figure 5.

This model predicts that IANA will allocate its final block of AS numbers in February 2011, to ARIN. The first RIR to completely exhaust its pool of 16 bit AS Numbers will be the RIPE NCC, and this is projected to occur in September 2011.



Figure 5 - AS Number Consumption Model

This is of course not new news. The expectation of exhaustion of the 16-bit AS Number pool has been with us for a decade or so. What is perhaps of more interest is that the rate of consumption of AS numbers has been so stable over many years. In March 2003 a similar predictive exercise on the longevity of the AS Number pool resulted in a projection of exhaustion of 2009, or possibly 2011 if there was some concerted efforts of reclamation and reuse of unused AS Numbers.



Figure 6 - March 2003 Prediction of AS Number Exhaustion

The Transition to 32-bit AS Numbers

Once it was evident that IDRP was not going to be the new inter-domain routing protocol for the Internet, and that we were going to be stuck with BGP for some time to come, then it was necessary for BGP itself to be modified to cope with this issue of exhaustion of the 16-bit AS Number pool.

The basic approach devised in the late 90's was simple: leave the AS number semantics alone, but simple change the length of the AS number field in BGP to use 32 bits rather than 16 bits, and do so in a fashion that was entirely backward compatible with the installed based of BGP speakers.

The transition process, developed informally in 2004, was one with four major milestones:

- 1. Standardization of the modifications to BGP
- 2. Modify the AS number distribution to include the distribution of 32-bit AS Numbers
- 3. Suppliers of BGP implementations to provide 32-bit AS BGP versions to their customers
- 4. BGP networks to commence deployment

The first step was estimated to take about 2 years, and the second and third steps could proceed in parallel once the standard was stable. The second step would take around six months and the third step was also estimated to take around 2 years. The intended outcome was that the transition to allow all new networks to use 32-bit AS numbers would be complete by the start of 2008, which was some 18 months earlier than the estimated time of 16-bit pool exhaustion

How well have we managed to stick to this timetable?

1. Standardization of the modifications to BGP

This standardization effort for 32-bit AS numbers commenced in the IETF some 10 years ago. The basic approach is quite simple: change the AS Path attribute to hold 32 bit AS Numbers rather than 16 bit AS numbers.

Most of the additional detail comes from a desire for this revised BGP to be fully interoperable with existing BGP implementations. For that reason the 32-bit AS BGP speaker (or "NEW BGP") will open a session with a remote BGP peer in a neutral fashion, and offer 32-bit AS number support, and its 32-bit AS number value, as a negotiated capability in the OPEN message exchange, while also offering a dummy 16-bit AS placeholder in the field of the OPEN message where "My AS" is recorded in the 16-bit AS BGP (or "OLD BGP"). If the remote BGP speak is also a NEW BGP speaker, the session will operate in an entirely conventional fashion, exchanging 32 bit AS numbers in the attributes of all BGP Update messages.

However, OLD BGP speakers will not recognise this 32-bit AS number capability, and in that case the NEW BGP speaker performs a number of additional steps to mimic the OLD BGP behaviour. Firstly, it presents itself as a "special" 16-bit AS value, namely AS 23456, in the My AS field of OPEN message. Secondly, it translates all the AS values in the AS Path attributes of Update messages that it sends to the OLD BGP speaker into 16-bit AS values. Those AS values between 0 and 65535 can be represented as 16 bit values by stripping off the leading 16 zeros. For larger values, the NEW BGP speaker maps these values to a single 16-bit token holder value of AS 23456. At the same time, the NEW BGP speaker copies the entire original 32-bit AS Path attribute into a transitive opaque attribute, the AS4_Path attribute. The OLD BGP speaker is not required to understand this AS4_Path attribute. As it is an opaque attribute, this attribute can be passed between OLD BGP speakers without generating any error condition. Thirdly, the NEW BGP speaker performs a reverse transformation on all Updates received from the OLD BGP speaker. All the AS's in the AS Path are expanded to 32-bit values. The AS4_Path attribute is then used to substitute original 32-bit values back into the AS Path in place of the AS 23456 place holders in the original AS Path. This combination of translation using AS 23456 and tunnelling using the AS4_Path attribute allows for piecemeal deployment of 32-bit ASes and NEW BGP speakers, both in the inter-domain environment, and even within an AS in the iBGP environment.

The 32-bit AS BGP specification was ready for publication as an RFC in late 2006. The IANA created the 32-bit AS number registry in November 2006, and the BGP specification, RFC 4893, was published in May 2007.

While this entire process took a little longer than originally anticipated, nevertheless the work was completed within a reasonable amount of time, allowing a further 2 - 3 years before the anticipated time when exhaustion of the 16 bit AS number pool would occur.

2. RIRs to start making 32-bit AS numbers available

In 2005 the Regional Internet Registries (RIRs) considered a coordinated policy framework to undertake the transition to 32-bit AS numbers. The framework that was adopted was one that attempted to provide a clear indication to vendors and network administrators as to when 32-bit AS-capable versions of BGP should be available for general use, namely by January 2009.

The adopted approach used a number of milestones:

• January 2007

32-bit AS numbers would be available upon request, while 16 bit AS numbers would be available as normal

- January 2009
 32-bit AS numbers would be available by default, while 16 bit AS numbers would be available upon specific request.
- January 2010 The transition was projected to be complete, and there was no need to continue to distinguish between 16-bit and 32-bit AS values.

3. Vendors to provide 32-bit AS number capable BGP implementations

Rather than provide a snapshot here that could well be outdated in a few months, perhaps its better to provide a reference to a resource page that lists a number of router equipment vendors and the versions of their software that include support for 32-bit AS Numbers: http://as4.cluepon.net/index.php/Software_Support

This aspect of the transition plan is certainly running later than was originally hoped, and this delay in availability of 32-bit AS BGP has had a significant impact on the level of uptake in the public Internet.

However, it should be remembered that there is no strict requirement for universal adoption of this revision of BGP at any particular time. Existing networks can continue to run 16-bit AS BGP without any requirement to switch over to a 32-bit AS capable version of BGP at any time. There is no universal deadline for existing deployment of BGP. It is new network deployments that are configured with a new AS number will be required to support 32-bit BGP, so while this aspect of the transition appears to be running a little behind the original expectation, this is not yet a critical issue.

0 0		Software Support – AS4							
 A 	A 🕂 😚 http://as4.cluepon.	net/index.php/Software_Support	RSS 🖒	Qr Google					
Software Sup	oport - AS4			+					
	nana discussion view so	urea Distory		& Log in / create account					
		dice marchy							
	Sonware Support								
	These versions of software support	32-bit ASNs.							
We	Software								
	Name	Version	Notation						
Main Page	Alcately ucent SB OS #	70	asolaio						
Community portal	Arbor Paskflow SP 4	55	asplain						
 Current events Becent changes 	Arbor Feakilow of B	5.5	aspian						
Random page	BIRD 6	1.0.12	asplain						
= Help	Brocade (Foundry) IronWare	4.0.00 for the NetIron MLX and XMR, 2.8.00 for the BigIron RX	asdot, asdot+, asplain						
search	Cisco IOS 🧬	12.0(32)S12, 12.0(32)SY8, 12.2(33)SXI1, 12.4(24)T	asdot (a	splain planned for future)					
Go (Search)	Cisco IOS XE 🛃	2.3	asplain (asdot optional)						
toolbox	Cisco IOS XR 💋	3.4(1)	asdot (a	asdot (asplain planned for 3.9)					
 Related changes 	Cisco NX-OS 🗗	4.0(1)	asdot (asplain planned for 4.1(3))						
Special pages Printable version	ExtremeXOS 🗗	Need Information	Need Information						
Permanent link	Juniper JUNOS	9.1R1	asplain (asdot optional)						
	Juniper JUNOSe 🚱	4.1.0	asplain						
	Force10 FTOS	7.7.1.0	asplain (asdot, asdot+ optional)						
	OpenBGPD ট	4.2, patches for 3.9 and 4.0	asdot						
	Quagga 🗗	0.99.10, patches for 0.99.6 and other versions	asplain						
	Redback SEOS	2.0	ascolon (asplain planned for end of 2009)						
	Tools								
	 RIR whois 								
	 IRRd accepts and indexes asdot 22 bit ASNo 	dot format AS numbers from mirrored databases, but the parser for submissions to local databases does not yet accept							
	NetFlow v9 and sFlow v5 both support 32-bit ASNs, but the router needs to fill in the 32-bit ASN and the collector needs to understand them								

Figure 7 - Support Table for 32-bit BGP

4. BGP networks to commence deployment

Obviously this aspect of the transition plan is running well behind the originally envisaged schedule. It was anticipated that by late 2009 the level of 32-bit AS Number support in networking products would be close to universal, and that all new network deployments would be able to use 32 bit AS numbers without any problems.

However a look at the RIRs' AS number allocation data paints a different picture. In 2009 the policy intention was that we would now be at the stage where 32 bit AS numbers were the default, and as the year progressed the number of 16 bit AS numbers being allocated was to dwindle off to a very small number. In the period 1 January to 12 August 2009 the RIRs allocated 2,683 AS numbers, of which 2,514 were 16-bit AS Numbers and only 169 were 32-bit AS Numbers. In total, since 1 January 2007, the RIRs have allocated 291 AS Numbers, of which only 70 can be seen in the BGP routing table. (Figure 8, Figure 9)





Figure 9 - 32 bit ASNs: Allocations and BGP Advertisements

It appears that in this phase of the AS transition we appear to be lagging badly, and rather than ensuring that we had completed the necessary preparatory activities in order to avoid the entire issue of exhaustion of 16-bit AS numbers, we may be in the rather unpleasant situation of running out of AS numbers before we have fully completed the BGP transition.

What could go wrong?

Mitigating the risks associated with any infrastructure transition lies in thorough preparation.

For network operators it may be preparing your own plan with your vendor and your suppliers to ensure that you can support customers, peers and upstreams using 32-bit AS numbers in the near future. And you may want to consider your own network and look at whether you want to prepare you own plan to migrate your AS routing domain to use a 32-bit BGP platform sometime in the coming months.

For vendors of network equipment it may be in preparing a wide variety of BGP upgrade paths for your customers to ensure that they don't have to transition to the absolute latest software revision just to get a 32-bit capable BGP. Many such upgrades are shunned by network operators, either because there are residual issues about software stability in such bleeding edge software releases, or because such recent software releases may entail other firmware or hardware changes that are not feasible for your customers' networks.

And part of the assistance lies in decent information about the situation. There appear to be a number of common misunderstandings and misapprehensions about using 32-bit AS Numbers. I'd like to look at some of the more common of these here.

Do I have to upgrade my BGP?

Someone out there is using 4 byte AS numbers. Do I have to upgrade my BGP to support 4-byte AS numbers in order to reach the prefixes that they are announcing?

NO!

BGP uses a translation and tunneling approach to mapping 32-bit AS numbers into 16-bit AS number The 32-bit AS BGP speaker does all the translation and tunneling work, so the

existing 16-bit BGP environment will not need to upgrade to "see" these additional networks that lie within 32-bit ASNs in the routing space All that you will see is AS 23456 appearing in many AS paths. But because your local AS is not AS23456 this has no impact whatsoever on routing and reacahability. All that 16-bit BGP implementations may notice is a very minor increase in memory use by BGP, associated with the storage of the additional AS4_PATH attribute that contains the 32-bit AS path. As this is an opaque transitive attribute to the 16-bit BGP speaker, this is no problem.

My customers and/or peers and/or upstreams are using 32-bit AS numbers. Do I have to upgrade my BGP to support 32-bit AS numbers?

NO!

You need to do nothing!

The new 32-bit AS BGP speaker figures out its talking to your o16-bit AS BGP speaker through the capability negotiation at the start of the BGP session, and the 32-bit AS BGP speaker then does all the necessary translation work. All these eBGP neighbors will appear to your eBGP routers as multiple instances of AS23456, all with distinct nexthops. To your local BGP environment this is perfectly normal and BGP routing will work just fine.

However, you should've checked out your operational support system by now to make sure it can cope with multiple external BGP adjacencies that all present themselves to your routers as AS23456. This may be an important consideration to you if, for example, you are using Internet Routing Registry (IRR) tools to load up AS path filters, for example, as you may need to compute the set of prefix and path filters using the 32-bit AS values contain in the IRR objects, but then be prepared to translate the instances of 32-bit AS numbers to AS23456 in the process of creating routing configuration fragments for your local 16-bit AS BGP routers. You may need to support multiple peers, customers and upstream providers who will have 32-bit AS numbers, and you will want to differentiate between them within your operational support environment, but ensure that if you autogenerate router configuration fragments, the correct mapping to AS23456 is taking place.

BGP Functionality

Can I use communities for 32-bit ASNs?

YES and NO

Using conventional communities, then this is not possible, as the field to specify the AS number is only 16 bits in length. To specify a 32-bit AS in a BGP community it is necessary to use the BGP extended community, and your BGP implementation should support the use of 32-bit AS numbers in extended communities. The specification is described at this stage in an Internet Draft document, draft-ietf-I3vpn-as4octet-ext-community, and while it is not published as an Proposed Standard RFC, it is in the final stages of pre-publication review, and many BGP implementations already include support for such extended communities in their latest versions of BGP. If not, ask your vendor when they will be supporting BGP extended communities with 32-bit ASNs.

Upgrading BGP

If I upgrade BGP, will my network crash?

NO!

Perhaps this should be a "MAYBE".

There are some recently encountered problems with BGP that affect both the commonly deployed 16-bit versions of BGP and some problems that have been encountered in the 32-bit AS versions of BGP. The original issue lies in the BGP specification, where the specification requires that if a BGP speaker receives an Update that has mal-formed fields, then the BGP speaker should send its peer a Notification message and reset the BGP session.

The manifestation of this problem that affects many 16-bit and 32-bit BGP speakers is where ASO is used in the AS Path. This is a reserved AS value that should not be used in any routing context, and in some implementations of BGP the peering session will be shutdown if AS 0 has been used anywhere in the AS Path. When the session reset and the peer once more attempts to pass across the same offending route object, the session will again be reset, and so on.

A more subtle manifestation of this problem can be exacerbated when AS0 is used in the AS4_Path attribute, but not in the AS Path attribute. The 32-bit AS BGP speaker, upon received such as update will send a notification to its peer and reset the session. But as this error is contained in an opaque transitive attribute, the 16-bit AS BGP peer has no knowledge of the problem.

There are also some slightly more subtle issues that are introduced here through the use of the AS4_Path attribute. The maximum BGP message size is 4096 octets. Making a very rough approximation that the remainder of the BGP Update message takes 96 octets, this would allow 4,000 octets for the AS Path, or a maximum of 2,000 AS Numbers in an AS Path, as each AS requires 2 octets. However, the maximum AS Path length under the same conditions when the update is being passed through a 16-bit BGP domain is smaller, as each AS is carried in the AS Path as a 16-bit value and also in the AS4-Path as a 32-bit quantity. The same 4,000 octets are capable of carrying a maximum AS Path length of 666 AS Numbers, as each AS now requires 6 octets in such a scenario. A maximum AS limit configuration setting in your BGP implementation is definitely worth using.

There is also the requirement that AS Confederations must not appear in a 32-bit AS configuration, or, to be a little more precise, should not be encoded into the AS4_Path attribute. Again, if a 32-bit AS BGP speaker receives an update with this in the AS4_Path, it is expected to send a Notification to its peer and shut down the BGP session.

The most "reasonable" behaviour here appears to be one that requires a deviation from the current BGP specification, and rather that shut down the BGP session when it receives a malformed attribute in an update message, it should elect to log all updates that contain such malformed attributes, then disregard the update from further local processing. This is similar to the manner in which local policy settings can be used to filter received BGP route objects such that they will not be used in the local RIB nor propagated to any other BGP peers. This course of action at least prevents the AS4_Path attribute from becoming a rather subtle form of DDOS weapon against BGP implementations. Some BGP implementations already perform this form of update processing, and there is some current work underway in the IDR Working Group of the IETF to equip the BGP specification with a "soft" notification which is a form of advisory message that informs the peer of a malformed Update message, but does not subsequently close the BGP session.

Can I upgrade my BGP routers one at a time?

Of course!

Assuming that the AS in which these BGP speakers are located is a 16-bit AS, then it is possible to upgrade the BGP speakers one at a time. The 16-bit / 32-bit transition behaviour is not reliant on eBGP or iBGP. However, many large iBGP environments make extensive use of route reflectors and local BGP "clusters," and in such situations it may be

appropriate to look at an upgrade to BGP that migrates each local iBGP cluster in turn, ensuring that the route reflectors remains synced with their set of iBGP clients across the transition.

32-bit AS transition

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

Calm down!

It may be abnormal, but its not fatal.

[gih@wattle /var/dat	ta/bgp/as2.0]\$ grep	23456 bgptable.txt								
* 193.169.148.0/23	203.119.76.3			4608	1221	4637	3356 20960	23456		
*>	202.12.28.190			4777	2516	3356	20960 23450			
*> 194.0.68.0/22	202.12.28.190			4777	2516	3491	20485 23450			
*	203.119.76.3			4608	1221	4637	3491 20485	23456		
*> 195.88.154.0/23	202.12.28.190			4777	2497	9002	44237 3426	7 23456		
*	203.119.76.3			4608	9942	6453	1299 44237	34267	23456	
* 203.217.144.0/22	203.119.76.3			4608	1221	4637	1273 9498 2	23456 i		
*>	202.12.28.190			4777	2516	7473	9498 23456			
[gih@wattle /var/data/bgp/as2.0]\$ [

It's either the result of some form of path prepending of AS23456 in the 32-bit AS world, or the result of AS Path manipulation in the 16-bit AS world, but in either case it should not cause any problems for BGP. The AS Path is used for loop detection and path metric. The presence of AS23456 in the AS Path is the same as the presence of any other AS in the AS Path, in that it adds 1 to the path metric. In terms of loop detection, each AS is looking for "itself" in the AS Path, so as long as you don't configure yourself as AS23456 (which would be a rather silly thing to do in the first place!) you will not interpret AS23456 in the AS Path as an instance of a BGP loop in any case.

More generally, it is possible for a 16-bit BGP speaker to drop the AS4_Path attribute, or attempt to corrupt it in various ways. The 32-bit BGP speaker that picks this up can either drop the update as an instance of a malformed attribute, or simply accept the 16-bit AS Path, including instances of AS23456, as a valid path. The only potential issue here if the update is accepted is that a routing loop that straddles the 16-bit AS and 32-bit AS domains may slip through the 32-bit AZS domain undetected. But when the corrupted update is passed back in to the 16-bit AS domain the loop will be detected. So in this case the loop detection may take a little longer, but the routing loop will be reliably detected in any case, even when the AS4_Path attribute is dropped from the update.

Are there AS Bogons in the 4-byte space?

Yes!

On the 12th August 2009 the following bogons were visible in the BGP routing table:

Advertised 4-byte ASNs: 70

Advertised Bogons: 4196636advertised by AS 29608 - WAN2MANY262657advertised by AS 12956 - Telefonica393392advertised by AS 12874 - Fastweb2076901376 advertised by AS 43314 - DIANET

When?

"Now" sounds like a good time to get moving with this transition.

At this stage we will have some remaining time while there are still some 16-bit AS numbers left to assist with the transition. But if we don't get serious about using 32-bit AS numbers in new network deployments we may well have to resort to quite extraordinary lengths of AS number sharing or similar measures to try and make a 16 bit number space encompass a larger routing domain. If we ever get to that point there are some very serious questions about routing integrity and security that will become quote difficult to answer.

Further Reading

32-bit ASN Resources

RFC4893 – the 32-bit AS specification draft-ietf-idr-rfc4893bis – working document: revision to the 32-bit AS specification draft-ietf-idr-optional-transitive - working document: BGP error handling for optional transitive BGP attributes

Documentation

Exploring AS Numbers – Internet Protocol Journal, Vol 9, No 1

Reports and Resources The AS Reports

http://www.potaroo.net/tools/asn16/ http://www.potaroo.net/tools/asn32/ ISP Resource Wiki for ASNs http://as4.cluepon.net

Disclaimer

The above views do not necessarily represent the views or positions of the Asia Pacific Network Information Centre, nor those of the Internet Society.

About the Author

GEOFF HUSTON is the Chief Scientist at APNIC, the Regional Internet Registry serving the Asia Pacific region. he graduated from the Australian National University with a B.Sc, and M.Sc. in Computer Science. He has been closely involved with the development of the Internet for many years, particularly within Australia, where he was responsible for the initial build of the Internet within the Australian academic and research sector. He is author of a number of Internet-related books, and was a member of the Internet Architecture Board from 1999 until 2005, and served on the Board of Trustees of the Internet Society from 1992 until 2001.

http://www.potaroo.net