The ISP Column
An occasional column on things Internet

November 2007

Geoff Huston

NANOGGING

There are many network operator group meetings being held these days. Even in the backwater of the South Pacific where I live there is now AUSNOG, and NZNOG is just next door in New Zealand. We now have MENOG in the Middle East and AFNOG in Africa. The original NOG was the North American Network Operators Group, NANOG, and they have the T-Shirts to prove it! NANOG meets three times a year, and I attended NANOG 41 in October 2007. NANOG meetings cover a broad variety of topics, from operational tools, measurement, and peering practices through to a commentary on the state of the Internet industry. Here's my impressions of the meeting.

A Perfect Storm

The first presentation in the NANOG meeting was a perspective on the state of the industry and an interpretation of the major trends by Dave Meyer, in a presentation on the Perfect Storm.

> The phrase "perfect storm" refers to the simultaneous occurrence of events which, taken individually, would be far less powerful than the result of their chance combination. Such occurrences are rare by their very nature, so that even a slight change in any one event contributing to the perfect storm would lessen its overall impact.
>
> The term is also used to describe a hypothetical hurricane that happens to hit at a region's most vulnerable area, resulting in the worst possible damage by a hurricane of its magnitude.
>
> [http://en.wikipedia.org/wiki/Perfect_storm]

The concept of a Perfect Storm is based on the observation that the following basic events could lead to a significant restructuring of our industry, the shift of the IP packet carriage business to a low margin commodity utility business, the weakening or even dismantling of monopolies in access, and the emergence of a set of peer-to-peer applications emerge that co-opt the incumbents' revenue streams.

The presentation explored some of the business implications of the end-to-end architectural model of the network, and in particular, noted that the application of this model tended to push not only functionality and utility to the edges of the network, but also pushed value and margin to the edges as well, leaving the network's "core" as a commodity packet-pushing enterprise. Here capital expenditure requirements can form a barrier to entry for players, because the activity is the subject of commodity-based competition with consequent margin erosion, making capital investment a proposition that has relatively low yields and increased risk. This is not a stable situation, and one of the factors that perturbs this situation is the tension between of economies of scale and technology creep. As long as incumbents can realize reduced marginal for

each additional customer, there is a tendency of the sector to aggregate and monopolies form. The balancing tension is technology creep, where a new entrant to the market can deploy new equipment and realize improved efficiencies due to superior price performance of the equipment.

Within this overall structure, however, the outlook is still one of a basic commodity utility enterprise. A more traditional view of the network operator as the critical central player in the telecommunications enterprise that exercises overarching control over all other parties and is able to reap the consequent rewards in terms of control over margins on capital and the maintenance of inherent barriers to entry that perpetuate monopolistic positions is challenged by this perspective of network carriage as a commodity. This is the first element of the perfect storm scenario.

The second element of this perfect storm is the proposition that monopolies over last mile monopolies are eroded or even destroyed. This part of the case offers alternate access technologies, including various flavours of wireless access and the potential for further regulatory involvement. It has often been the case that consumers are held captive to access providers and the defacto monopolies in access have tended to imply that in the absence of a persistent regulatory framework true competition for the user evaporates in exorbitant cross-carrier changes for last mile access. Weakening of this access monopoly, if this is indeed the case, is not a preferred outcome for an incumbent provider.

The final element for this scenario is the shift of services into edge-to-edge applications. This is not only evident in web-based applications but also in what were traditionally considered network-centric applications, such as voice and video applications.

The combination of these three elements, if all are played through, will place significant constraints on network providers, given that the previous incarnation of this industry, namely that of monopoly telcos, were entirely ignorant of commodity utility operations. There is a prospect of shift to a slim investment profile by carriage service providers in order to reduce their risk and in an attempt to improve margins, a move that in turn may exacerbate situations where further investment is being called for. Part of the reluctance of industry to make significant investments in IPv6 infrastructure at present could be attributed to commodity utility factors in this sector of the industry. There is a strong risk aversion factor in the industry today that may also impact the industry's ability to further scale the capacity of the network into terabit capacities.

IPFIX

No, its not a repair kit for IP! IPFIX is the name of an IETF standards activity that addresses the flexible export of IP flow data from routers, or other metering processes, to a flow collector. IP metering has had an extensive history, and its based on a desire to view the activity in an IP network as more than a collection of packets. The flow view of a network attempts to group together IP packets that are associated with a single TCP session, or with a single UDP transaction, or with any other grouping function that can relate packets together in a single unit and then look at metrics of the packet set, or flow, such as the amount of data transferred per flow, the number of packets per flow, and the duration of each flow. The IETF IPFIX work has started from the Cisco Netflow V9 specification and extended it. Now this work is not about how to perform flow accounting, but to define how to "package" flow information when it is being exported from a measurement device. This is related to sFlow, which is a packet sampling protocol which uses sampling as its basic data gathering rather than flow aggregation. There was one question raised at NANOG which I thought was an interesting commentary on the evolution of the standards work in the Internet space, which was: "Why is the IPFIX specification about the same size as that of TCP itself?" I'm not sure that the answer is at all reassuring if you are looking for lean and simple solutions in today's protocol standards.

Optical Networking

We've heard for some time that there was some divisions over where to go after 10G fibre optics. Some wanted to standardize 40G as being a readily achievable and understood technology. Others have argued that the pace of continuing traffic growth means that 40G is too little too late and that at the very least we should be standardizing 100G. The growth trend continues along a bandwidth demand trajectory that is steeper than linear and some commentators believe that they see a growth trend appears to be exponential with a doubling factor of every 2.5 years. The topic of high speed optical systems was discussed on a panel at this NANOG meeting.

If 10G per wavelength networks are feeling the strain today then not only do we need to pull up high data rates per optical wavelength, but ideally have direct coupling of IP to the Dense Wave Division Multiplexing (DWDM) system (or IPoDWDM if you are in to acronym concatenation!), removing local transponders from the physical plant. While some bandwidth can be gained by various forms of concatenation to simulate a larger system by a collection of smaller systems operating in parallel, the strong industry preference is delivery of single clocked bandwidth in as larger speed as possible. Here's where modulation enters the story. Optical systems up to 10G use a simple off / on encoding, where a "1" is a transition from "on" or "off" and a 0 is a transition from "off" to "on". A string of all zeros or all one involves twice the number of transitions between on and off states than the basic data rate. One way to increase the data rate is to increase the base frequency of the modulation. Another way is to use a number of discrete amplitude levels. Yet another is to alter the phase of the carrier signal, and with optical systems you can also alter the polarization of the signal. Taking a basic carrier signal at 10Ghz, it is possible to achieve 20Gbps using a modulation technique such as Quadrature Phase-Shift Keying (QPSK). Higher data rates can be achieved with "denser" modulation technique, such as QAM-16 or QAM-256. The capacity cab be further doubled by using "vertical" and "horizontal" polarization to place two signals on a single light carrier. In a perfectly lossless and coherent light transport system then 1Tbps would be readily achievable. When typical signal to noise ratios are added then a more realistic picture of using QPSK and polarization providing somewhere between 2 and 8 bits / Hz looks feasible. So it really does appear that as the 40G systems hit the market there are 100G systems being designed. The design objectives include an optical reach of at least 1,500km, support for existing 50Ghz DWDM channel spacing and existing DWDM equipment, operating within existing crosstalk parameters and chromatic dispersion tolerances and be bit level transparent for a 100GE client interface.

While this appears to be technically possible there is one additional parameter, namely that of cost. There is always a tradeoff between performance and cost and here the objective is to construct high speed systems with practical electronic speeds given that electrical signal processing is in general available at a lower cost than equivalent optical signal processing. Its also the case that digital signal processing is at an early stage in its development in optical communications. It seems likely that 100G systems can be constructed on a 25Gbaud optical system using QPSK to achieve 2 bits per Hz and then combine two of these streams using 90^{o} polarization, inferring that deploying 100Gbps systems using the existing 10Gbps optical infrastructure appears to feasible in many cases. Of more interest is the cost factor of such systems. Higher speeds implies a need for better signal to noise ratios which in turn implies a need for higher power coupling from the transducer into the fibre cable which in turn leads to an engineering issue of high power stable tuneable lasers constructed within an acceptable unit price structure. Whether 100G will achieve further unit cost efficiencies over the 10G and 40G systems remains an open question at this point in time.

But the Internet remains a voracious consumer of network capacity, and doubtless we will see the call for Tbps optical systems in the very near future.

The LISP Routing Experiment

The Computer Scientist David Wheeler is quoted as saying "Any problem in computer science can be solved with another layer of indirection."

Indirection as a tool to assist in routing scaling the routing system is being investigated in an experimental project termed 'LISP', presented to NANOG by Dino Farinacci. The term "LISP stands for locator ID separation, which for me is somewhat misleading as it is actually an instance of a relatively conventional form of tunnelling where the original IP datagram is encapsulated in a UDP datagram. Conventionally, tunnels are established by static configuration at both ends, but in LISP the tunnels are created dynamically. Ingress Tunnel Routers (ITR) take an unencapsulated packet and encapsulate it inside a UDP packet whose outer IP header uses a source address of the ITR and a destination address of the Egress Tunnel Router (ETR) associated with the destination address. The intended outcome is that the only packets passing across the "core" of the network are encapsulated IP packets and the routing load of the default free zone is essentially the routing space associated with the TRs, which, hopefully, will be a far smaller space than the host address space. The problem being addressed here is that of the pressures on the Internet through continued growth pressure being placed on the routing system. I must admit that I'm not exactly sure of the difference between the scale of the problem in 100G and faster IP transports and very large routing tables. In the case of the optical transports there is a continuing impression that solutions are there and that there will be customers for the solution, and that demand for ever higher speeds will continue. If the same argument is applied to inter-domain routing then there would be the expectation that the current routing solutions, which are one order of magnitude larger than the current routing table size, will continue to scale, and routing systems will continue to increase in size, and there will be customers for larger routing systems and that continued growth in the capacity of routing systems will be matched by continued demand for such systems. This does not appear to

be the case in routing, and there appears to be a widely held view that further growth of the routing system will encounter a hard limit, while the optical systems still have much more space to explore.

Back to LISP design, much of the challenge in this form of approach lies in distributing the mapping between ETR host addresses and the equivalence class of addresses that are “served” by that ETR. Approaches using ‘just-in-case’ push, ‘just-too-late’ pull and pre-loading the mappings are being explored.

A followup presentation by Dave Meyer explored “CONS” a control plane mapping system for LISP that relies on the approach of a large distributed database . Interestingly, while the basic LISP approach is quite conventional and simple, it’s the mapping systems that get very complex very quickly, and CONS is no exception to this generalization. The trade-offs in this kind of system repeat those that underlie the DNS, namely that speed, accuracy, state and update load are related. Systems that perform “push” escalate state and have to tradeoff accuracy and update load, but can operate at speed. Systems that perform “pull” can operate with less state, but are slower to return the mapping outcome for any particular query.

One of the problems encountered in tunnelling IPv6 over IPv4 has been the issue of MTU discovery and MTU-based traffic black holes. I suspect that LISP will also encounter this problem once it moves out of a more carefully controlled development environment into the broader Internet for the next set of tests.

As David Wheeler’s quote about indirection continues: “ But that usually will create another problem."

Lawyer Speak

It is unusual for a lawyer to be presenting at a network operator’s meeting. This presentation at NANOG was on the topic of a lawyer’s reflection on Internet Governance. Internet Governance has been a hot topic in recent years, starting with the preparations for the 2005 World Summit on the Information Society (WSIS) and the followup Internet Governance Forums in 2006 and the forthcoming one in November 2007. The underlying issues of international arrangements and the influence of individual governments and related political agendas being played out in the Internet has been a topic that has generated far more heat than light over the past few years. The perspective voiced in this session certainly was influenced by the stance of the US national position here, which is not unsurprising for a North American meeting. It started from the premise that the current structure, which is strongly based in concepts of deregulation of public telecommunications and the operation of markets in a strongly competitive framework, is far more appropriate than an alternative that asserts the primacy of national policies through direct regulatory control of incumbent operators.

Interestingly, while espousing the overall advantages of deregulation, market-based economies and the beneficial outcomes of competition the presentation was very emphatic that addresses should not be distributed in a market-based framework, and that addresses were not to be considered property. While there were ample IPv4 unallocated addresses left to fuel the growth of the Internet I’m not sure that this was a matter of any particular concern, but as the exhaustion of that pool is now in sight the matter of what happens in a few years time with IPv4 addresses is now far more prominent. Are addresses something that could be traded in a market-based framework? What are the policy dimensions of such a shift in perspective? And what should we be doing in the period leading up to the depletion of this particular address pool?

The State of IPv6

Given the prospect of change in the air over IPv4 and IPv6, the presentation on the IPv6 routing table was timely. Gert Döring has been presenting updates to the RIPE community for some years and at this NANOG his work was presented to the NANOG community. There is no doubt that IPv6 has a much smaller deployment footprint than IPv4. While the IPv4 routing table has some 240,000 entries, a growth of 40,000 entries in the past 12 months, the IPv6 routing table has some 920 entries, and has grown by 150 entries over the past 12 months. The table is still small enough to allow detailed analysis of each major routing event, but the overall picture is still somewhat surprising and concerning – with between 2 – to 3 years of a major disruptive event in IPv4 with the exhaustion of the unallocated address pool there is still scant evidence that industry has paid any attention to IPv6 , and the prospect of undertaking a manageable transition to IPv6 is fast disappearing.

TEOTWAWKI

I'm sure this will not be the last of "the end of the world, as we know it" sessions. Indeed the longer there is no sign of industry response the larger the looming problem will be and the more disruptive the situation. The lack of industry adoption of IPv6 is a serious concern, and it is becoming increasingly apparent that while there are many strong and compelling reasons for IPv6 deployment, the business case for deployment has not been anywhere near as convincing. IPv6 is, to put it bluntly, a business failure. (The quotes on each slide of this presentation are very aptly chosen!) So we are now facing two problems, not just one. The first is trying to address the business issues of IPv6 deployment and cast IPv6 in a more realistic perspective about what it can and cannot do. The second is to continue to support a growing IPv4 Internet after the current address distribution system runs to completion. At this stage we can anticipate far more intense use of NATs, possibly various forms of address markets and consequent pressure on rapid growth in the inter-domain routing space. Each of these changes in isolation would represent a significant challenge to a global deregulated industry which is still recovering from the last major upheaval of the Internet boom and bust at the start of this century. Taken together they represent a set of challenges to the industry that will cause significant disruption. When combined with the themes outlined in the perfect storm situation the scale tends to move from "significant disruption" to "major upheaval". With this prospect in mind it is possible to interpret the unfolding situation as being sufficiently complex, and sufficiently disastrous, that collective denial of the entire situation becomes the only possible response from the deregulated industry base!

Just the IPv6 Facts Thanks

IPv6 has been around for a very long time, and the marketing message has been inflated with more and more hyperbole in response to industry's continuing disinterest in the topic. Randy Bush gave a presentation that attempted to deflate some of the more excessive claims about the wonders of IPv6, such as the rather preposterous claim that IPv6 increases battery life in laptops! The premise behind the presentation is that the marketing hype is counter-productive and its not actually assisting in the effort to deploy IPv6, so its time to take a more grounded look at IPv6 and make an honest assessment of what needs to be done. The current situation was characterised as one with "no transition plan, [a] declared victory before the hard part started, no real long term plan, no realistic estimation of costs, no support for the folk on the front lines [and continual declaration that] victory will be next month" The problem is that IPv6 really is just IPv4 with more bits in the address space, and as a result it offers no real feature benefit but because there are extra bits in the address header the IPv6 packets are not backward compatible with the IPv4 installed base. Transition to IPv6 will be hard, and the longer we wait the larger the network; the longer we wait the fewer the number of IPv4 addresses left with which to fuel the transition; the longer we wait the longer the transition will take; the longer we wait the lower the commitment to make any transition actually work. This is not a pleasant position to be in.

The presentation advocated a level-headed assessment of the transition problem spaces. We've managed to build a rather complex network over the years and now wherever we may care to look there are instances of embedded IPv4 functionality where there are no equivalent IPv6 functions. From firewalls to load balancers, from mail filters to DLS modems, and from the DNS to DHCP we seem to have a lot of work to do. The transition to IPv6 will not be easy, nor quick. It will not eliminate NATs, nor will it alleviate routing load. The address plan is relatively wasteful and is already looking finite in size. There is no improvement in security. Incremental deployment is not possible and this process will be protracted and expensive.

The presentation looked at a number of pragmatic issues involving this transition, making the case that the time for "more features" in IPv6 was over and it was now time to move attention to the supply chains and ensure that IPv6 capability is well supported in products. Dual Stack should go as far as possible to the customer edge, but even so there will be significant use of Application Level Gateways.

The presentation was certainly a change from the more enthusiastic and overblown presentations that we've grown used to, and largely rejected as being completely oversold as part of some deluded marketing frenzy associated with IPv6. Hopefully this more critical assessment of the situation will gather a receptive audience. Either that or we need to get into IPv6 the Musical!

"Lightning" Talks

One of the common comments of attendees at technical conferences is that many talks are just too long! A typical technical presentation starts with some an extended introduction to the topic, some motivation as to why this is a relevant and interesting topic to study, identification of a problem in this space, and then a

description of the method taken to study the problem. Only then might the presentation touch upon some findings, and then the presentation will attempt to place these results into a broader context. All of this takes time. What if the speaker had to assume that the audience was thoroughly familiar with the topic and just wanted to hear the results? This has lead to the use of “lightning” presentations, where each presentation is limited to 10 minutes, and its just the results that matter!

This time at NANOG 41 there were 10 such presentations, which is a strong signal of the popularity of these “direct and straight to the point” presentations that appear to be a feature of many technical conferences these days. I found this particular set of lightening talks very much on the target for network operations.

Detecting Routing Leaks

Inter-domain routing involves large amounts of trust – trust that your routing peers are doing the right thing and trust that everyone else in the inter-domain routing environment is also doing the right thing. Filters and controls can help mitigate some of the exposures here, but not all – bad things happen and often its not clear what has happened until well after the event. One approach, presented by Jared Maunch, is to create a database of what is normally seen in the routing system in terms of AS paths, and then generate alarms if abnormalities are observed. The abnormalities involve an extended transit across the so-called tier-1 networks, or a AS path that appears contrary to generally observed inferred business relationships. The leak count since July 2007 is impressive, with some 360,000 anomalies detected using this approach. Weekends have less noise than weekdays, so operator finger trouble is a likely candidate reason for such route leaks. It’s a relatively lightweight but very interesting approach to this problem.

Unleashing Class E on an unsuspecting world

You know life, version 4, is getting desperate when we start looking in all those dim corners for more address space to consume. The tantalising prospect of making a further 268,435,454 addresses, or the equivalent of 16 /8 address blocks, available has tempted a number of folk, including Vince Fuller, who presented a proposal to bring this address block into play. Changing the registry file to change its designation is not the issue here – the real issue is that implementations of TCP/IP have opted to prevent the use of this address block in their code. While protecting users from themselves is often seen as a worthy end in and of itself, it can also have some unintended consequences. In this case it makes using 240/4 addresses a little more challenging. The presentation noted some efforts to patch existing code to support use of this address block and the likely problems that we face in trying to bring this block into a more general use. For those folk reading this on a MAC running MacOSX, or those of a FreeBSD persuasion, then the good news is that these addresses mostly work. For others then it’s a case of applying a set of patches to remove the “you must be crazy” protection lines in the IP code!

Routing Small Prefixes

From a routing perspective a routing advertisement is just another routing advertisement, irrespective of its size in terms of address span. Todd Underwood presented an analysis of the differences across a number of routing perspectives. The first observation is that spammers continue to inject /8 prefix advertisements into the routing system, particularly in Asia. The prefixes 63/8, 66/8 and 68/8 are being announced by AS2706, Pacific Internet Singapore, and have an announcement pattern that is strongly associated with spamming that were seen some time ago originating at AS 4678 Fine Canon in Hong Kong. There are also a number of advertisements of a /25 or smaller that have been noticed. The bulk of these advertisements originate with AS 23577, Korea Telecom, and appear to be selectively propagated across the Internet where minimum length prefix filters are not in place. The observation made in this presentation is that either there are more smaller prefixes being advertised today, or more transit providers are lifting their prefix length filters, or possibly a combination of both!

Passive DNS Data Collection

The typical cry from many folk researching the Internet is “send more data!” This presentation explores an approach to collecting data from the DNS by tapping into the responses from authoritative services and analysing them. It is possible to undertake analysis to profile the behaviour of the DNS using this approach, and to profile the behaviour of the inverse address to name mapping for use in activities such as botnet hunting. The presentation by Paul Vixie proposed an approach to instrument recursive DNS servers with DNS packet capture and to provide a central data collector as a means of facilitating real time access to live DNS data for researchers.

Who’s Been Squatting in My Address?

Not everyone has been studious in respecting the IPv4 address registry and keeping out of using the “reserved” address blocks. Leo Vegoda has investigated this and reported that network 1/8 has been used extensively for private use, 5/8 has been used by Hamachi, 42/8 by HP and 104/8 for Mesh networking. In the next couple of years these networks will be put into use and some local address clashes may result. File this under “you’ve now been warned.”

Is My Multicast Dead?

For a while back then multicast was all the rage, but these days it has retreated back into a relatively small sphere of use. Which is a shame in some ways as its quite an impressive forwarding technique. Stig Venaas reported on a tool he has developed, smpping, that tests multicast connectivity from an Source Specific Multicast (SSM) server by sending a unicast ping that requests the SMP server to generate both a unicast and multicast response. There’s also an Any Source Multicast (ASM) version of this tool. Now lets move onto the ever so slightly more difficult problem of actually deploying multicast!

BGP Anomalies

There has been a lot of work in the last couple of years to come to grips with various techniques to secure BGP. Josh Karlin’s presentation described a technique that has elements in common with the routing leak detector in that it attempts to detect anomalous situations in the routing system without resorting to the explicit labelling of “good” information via a Public Key Infrastructure. Some time back Lixin Gao published an algorithm that attempted to label every inter-AS relationship as being either a Provider / Customer relationship or a Peer relationship, exploiting the bimodal nature of inter-provider relationships in the Internet. The basic approach to classification was that every AS path was of the form of a sequence of none or more Customer -> Provider relationships, followed by at most one Peer relationships followed by none or more Provider -> Customer relationships. The basic concept here is that traffic does not transit a customer (or “every AS path is a mountain and there are no valleys”) At the time the research team noted that there were a number of anomalies to their classification technique. This work focuses in on these anomalies to see if they have the characteristics of a deliberate attempt to pervert the routing system, either by prefix hijacking or by AS path manipulation. NANOG has generally attracted researchers to the forum, as the opportunity to present their research concepts and ideas to an operational forum is a valuable opportunity to sanity test some concepts and gauge interest. This one does look interesting, although there are other approaches to this same problem of detection of routing attacks.

More Fun with Ping

If you operate an Internet exchange point it will become obvious pretty quickly that you need an address block for that exchange point to support the peering fabric. In theory this address space need not be announced outside the exchange point. Practice is different of course and John Kristoff has been pinging into the exchange point address space to see what is there. It appears possible to use this technique to construct peering maps that perform dynamic discovery of the who, where and how of peering.

Visualizations and Hilbert Curves

A long time ago in a different life I encountered Hilbert Curves while studying fractals, and I must admit that after heading off into networking I confidently expected never to encounter them ever again. After all, it is a member of that rather esoteric set of curves, the so-called space-filling curves.

> In 1890, Giuseppe Peano discovered a densely self-intersecting curve which passed through every point of the unit square. This was the first example of a space-filling curve. Peano's purpose was to construct a continuous mapping from the unit interval onto the unit square, motivated by Georg Cantor's earlier counterintuitive result that the infinite number of points in a unit interval is the same cardinality as the infinite number of points in any finite-dimensional manifold, such as the unit square.
>
> [http://en.wikipedia.org/wiki/Space-filling_curve]

So out pops a Hilbert curve map of the IPv4 address space, and on top of this curve one can colour it according to various properties of each address, such as the addresses of DNS open resolvers, maps of address space described in various address registries, or even measured round trip times for packet probes, as Duane Wessels has done in some quite remarkable visualizations of the Internet's address map. It's an intriguing visualization approach, although I'm still not sure why it needs to be done on a Hilbert Curve in particular. Perhaps next time we'll see the same data mapped onto a Heigway dragon curve, of mapped into a space-filling curve that is mapped into curved space or even a three dimensional space. Now that could really be quite neat!

And more!

No, this is not a complete report of NANOG 41, and there were more presentations, tutorials and sessions not covered here. The agendas and presentations of the NANOG41 meeting are all online, of course, as are all past NANOG meetings. But maybe you should just go there and experience it for yourself – its just three very filled days! NANOG 42 will be held on the 17th – 20th February 2008 at San Jose, California.

Disclaimer

The above views do not necessarily represent the views or positions of the Asia Pacific Network Information Centre.

About the Author

GEOFF HUSTON holds a B.Sc. and a M.Sc. from the Australian National University. 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 is currently the Chief Scientist at APNIC, the Regional Internet Registry serving the Asia Pacific region. He was a member of the Internet Architecture Board from 1999 until 2005, and served on the Board of the Internet Society from 1992 until 2001.

http://www.potaroo.net