

A Decade in the Life of the Internet

by Geoff Huston, APNIC

The evolutionary path of any technology can often take strange and unanticipated turns and twists. At some points simplicity and minimalism can be replaced by complexity and ornamentation, while at other times a dramatic cut-through exposes the core concepts of the technology and removes layers of superfluous additions. The technical evolution of the Internet appears to be no exception, and contains these same forms of unanticipated turns and twists.

This article presents a personal perspective of the evolution of the Internet over the last decade, highlighting my impressions of what has worked, what has not, and what has changed over this period. It has been an extraordinary decade for the Internet, encompassing a boom and a bust that would rate among history's best, a comprehensive restructuring of the communications industry, and a set of changes that have altered the way in which each of us now works and plays. And the Internet has even added a few new words to the language on the way.

Rather than offer a set of random observations, I will use the Internet Protocol model as a template, starting with the underlying transmission media, then looking at the internetwork layer, the transport layer, then applications and services, and, finally looking at the business of the Internet.

The Transmission Media Layer

It seems like it was in an entirely different lifetime, but the *Internet Service Provider* (ISP) business of 1998 was still centrally involved in the technology of dial-up modems. The state-of-the-art of modem speed had been continually refined from 9,600 bps to 14.4 kbps, to 28 kbps, to finally, 56 kbps, squeezing every last bit out the phase amplitude space contained in an analogue 3-KHz voice circuit. Modems were the bane of an ISP's life. They were capricious, constantly being superseded by the next technical refinement, unreliable, difficult for customers to use, and they were just slow. Almost everything else on the Internet was tailored to download reasonably quickly over a modem connection. Webpages were carefully tailored with compressed images, and plaintext was the dominant medium as a consequence.

Not all forms of Internet access were dial-up. ISDN was used in some places, but it was never cheap enough to take over as the ubiquitous access method. There were also access services based on *Frame Relay*, X.25, and various forms of digital data services. At the high end of the speed spectrum were T1 access circuits with 1.5-Mbps clocking, and T3 circuits clocked at 45 Mbps.

ISPs leased circuits from a telephony company (telco). In 1998 the ISP industry was undergoing a transition of its trunk IP infrastructure from T1 circuits to T3 circuits. It was not going to stop here, but squeezing even more capacity from the network was proving to be a challenge. Deployment of 622-Mbps IP circuits occurred, although many of these were constructed using 155-Mbps *Asynchronous Transfer Mode* (ATM) circuits using router load balancing to share the IP load over four of these circuits in parallel. Gigabit circuits were just beginning, and the initial tests of IP over 2.5-Gbps *Synchronous Digital Hierarchy* (SDH) circuits began in 1998.

In some ways 1998 was a pivotal year for IP transmission. Until this time IP was still just another application that was positioned as just another customer of the telco's switched-circuit infrastructure that was constructed primarily to support telephony. From the analogue voice circuits to the 64K digital circuit through to the trunk bearers, IP had been running on top of the voice network. By 1998 things were changing. The Internet had started to make ever larger demands on transmission capacity, and the factor accelerating further growth in the network was now not voice, but data. It made little sense to provision an ever larger voice-based switching infrastructure just to repackage it as IP, and by 1998 the industry was starting to consider just what an all-IP high-speed network would look like, from the photon all the way through to the application.

At the same time the fiber-optic systems were changing with the introduction of *Wavelength-Division Multiplexing* (WDM). Older fiber equipment with electro-optical repeaters and *Plesiochronous Digital Hierarchy* (PDH) multiplexers allowed a single fiber pair to carry around 560 Mbps of data. WDM allowed a fiber pair to carry multiple channels of data using different wavelengths, with each channel supporting a data rate of up to 10 Gbps. Channel capacity in a fiber strand is between 40 to 160 channels using *Dense WDM* (DWDM). Combined with the use of all-optical amplifiers, the most remarkable part of this entire evolution in fiber systems is that a Tbps cable system can be constructed today for much the same cost as a 560-Mbps cable system of the mid-1990s. The factor that accelerated deployment of these high-capacity fiber systems was never based on expansion of telephony, because the explosive growth of the industry was all about IP. So it came as no surprise that at the same time as the demand for IP transmission was increasing there was a shift in the transmission model, where instead of plugging routers into telco switching gear and using virtual point-to-point circuits for IP, we started to plug routers into wavelengths of the DWDM equipment and operate all-IP networks in the core of the Internet.

The evolution of access networks has seen a shift away from modems to numerous digital access methods, including DSL, cable modems, and high-speed wireless services. The copper pair of the telco network has proved surprisingly resilient, and DSL has achieved speeds of tens of megabits per second through this network, with the prospect of hundred-megabit systems appearing soon.

So, in terms of transmission, the last 10 years has seen the network migrate from an overlay system of kilobit-per-second access with multimegabit trunks operating as a customer of the telco switched network to a comprehensive IP network with access of megabits per second with multigigabit trunks, or a thousandfold increase in basic network capacity in that period.

The demand of the Internet for capacity continues, and we are now seeing work on standardizing 40- and 100-Gbps transmission systems in the IEEE; the prospect of terabit transmissions is now taking shape for the Internet.

The Internet Layer

If transmission has seen dramatic changes in the past decade, then what has happened at the IP layer over the same period?

The glib answer is “absolutely nothing!” But that answer would be ignoring a large amount of activity in this area. We have tried to change many parts of IP in the past decade, but, interestingly, none of the proposed changes has managed to gain any significant traction in the network, and IP today is largely no different from IP of a decade ago. *Mobility*^[1], *Multicast*^[2], and *IP Security* (IPSec)^[3] remain poised in the wings, still awaiting adoption by the Internet mainstream.

Quality of Service (QoS) was a “hot” topic in 1998, and it involved the search for a reasonable way for some packets to take the fast path while others took a more leisurely way through the network. We experimented with various forms of signaling, packet classifiers, queue-management algorithms, and interpretations of the *Type of Service* bits in the IPv4 packet header, and we explored the QoS architectures of *Integrated and Differentiated Services* in great detail. However, QoS never managed to achieve wide acceptance in mainstream Internet service environments. In this case the Internet took a simpler direction: In response to not enough network capacity, the alternate approach to installing additional mechanisms in the network—in the host protocol stack and even in the application in order to ration the capacity you have—is to simply expand the network to meet the total level of demand. So far the simple approach has prevailed in the network, and QoS remains largely unused^[4].

We have experimented with putting circuits back into the IP architecture in various ways, most notably with the *Multiprotocol Label Switching* (MPLS) technology^[5]. This technology used the label-swapping approach used in X.25, Frame Relay, and ATM virtual circuit switching systems; it created a collection of virtual paths from each network ingress to each network egress. The idea was that in the interior of the network you no longer needed to load up a complete routing table into each switching element, and instead of performing destination-address lookup you could perform a much smaller, and hopefully faster, label lookup.

This process did not eventuate, and switching packets using the 32-bit destination address continued to present much the same level of cost-efficiency at the hardware level as virtual circuit label switching. When you add the additional overhead of an additional level of indirection in terms of operational management of MPLS networks, MPLS became another technology that so far has not managed to achieve traction in mainstream Internet networks. However, MPLS is by no means a dormant technology, and one place where MPLS has enjoyed considerable deployment is in the corporate service sector where many *Virtual Private Networks*^[6] are constructed using MPLS as the core technology, steadily replacing a raft of traditional private data systems that used X.25, Frame Relay, ATM, *Switched Multimegabit Data Service* (SMDS), and switched Ethernet.

Of course one change at the IP level of the protocol stack that was intended in the past decade but has not occurred is *IP Version 6*^[7]. In 1998 we were forecasting that we would have consumed all the remaining unallocated IPv4 addresses by around 2008. We were saying at the time that, because we had completed the technical specification of IPv6, the next step was that of deployment and transition. There was no particular sense of urgency, and the comfortable expectation was that with a decade to go we did not need to raise any alarms. And this plan has worked, to some extent, in that today's popular desktop operating systems of Windows, MacOS, and UNIX all have IPv6 support. But other parts of this transition have been painfully slow. It was only a few months ago that the root of the *Domain Name System* (DNS) was able to answer queries using the IPv6 protocol as transport, and provide the IPv6 addresses of the root nameservers. Very few mainstream services are configured in a dual-stack fashion, and the prevailing view is still that the case for IPv6 deployment has not yet reached the necessary threshold. Usage measurements for IPv6 point to a level of deployment of around one-thousandth of the IPv4 network, and, perhaps more worrisome, this metric has not changed to any appreciable level in the past 4 years. So what about that projection of IPv4 unallocated pool exhaustion by 2008? How urgent is IPv6 now? The good news is that the *Internet Assigned Numbers Authority* (IANA) still has some 16 percent of the address space in its unallocated pool, so IPv4 address exhaustion is unlikely to occur this year. The bad news is that the global consumption rate of IP addresses is now at a level such that the remaining address pool can fuel the Internet for less than a further 3 years, and the exhaustion prediction is now sometime around 2010 to 2011.

So why have we not deployed IPv6 more seriously yet? And if we are not going to deploy IPv6, then what is the alternative? Of all the technical refinements to IP that have occurred, one that received little fanfare when it was first published has enjoyed massive deployment over the past decade, and that is the technology of *Network Address Translation* (NAT)^[8]. Today NAT devices are ubiquitous. It seems that every home access unit, every corporate firewall, every data center, and every service includes a NAT device.

One measure of the ubiquity of NATs is the transformation that has occurred in the application space. By 2008 applications have either adopted a strict client-server approach, where the client always initiates the network transaction, or were forced down a more complex path. Where there is some form of peer interaction, applications are now equipped with additional capabilities, including NAT behavior discovery, NAT binding management, application-level name spaces, and multiparty rendezvous mechanisms, all required to allow the application to function across NATs. So far we have managed to offload the problem of looming address scarcity in the Internet onto NATs, and the really significant change that has occurred in the past decade at the IP level is the default assumption about the semantics of an IP address. An IP address is no longer synonymous with the persistent identity of the remote party that anyone can use to initiate a communication, but a temporary token to allow a single transaction to complete. As a consequence, most Internet services have retreated into data centers and the business of hosting services has thrived. And the change that would have preserved the coherent end-to-end architecture of the Internet IP layer, namely IPv6, is still waiting for wide-scale deployment.

The next few years promise to be “interesting” in every form of meaning of the word. The exhaustion of the remaining IPv4 address pool is imminent, and if we are going to substitute IPv6 in place of IPv4, then we simply do not have enough time to achieve this substitution before the remaining IPv4 address pool is depleted. And although so far NATs have conveniently pushed the problem of increasing address scarcity off the network and over to the edge devices and onto applications, it is not clear that this approach can sustain an ever-growing Internet indefinitely. We have yet to understand just what a “carrier-grade NAT” might be, or whether it can even work in any useful manner. NATs were an accidental addition to the Internet, and their role in the coming years is unclear.

The early 1990s saw a flurry of activity in the routing space, and protocols were quickly developed and deployed. By 1998 the “standard” Internet environment involved the use of either *Intermediate System-to-Intermediate System* (IS-IS) or *Open Shortest Path First* (OSPF) as large-scale interior routing protocols and *Border Gateway Protocol 4* (BGP4) as the interdomain routing protocol^[9]. This picture has remained constant over the past decade. In some ways it is reassuring to see a technology that is capable of sustaining a quite dramatic growth rate, but perhaps that is not quite the complete picture.

We never quite completed the specification for the next interdomain routing protocol, and BGP4 is now showing signs of stress^[10]. The pool of *Autonomous System* (AS) numbers is forecast to run out early in 2011, and by then we need to have fielded a new variant of BGP that can operate with a much larger pool of AS numbers^[11].

Fortunately the technology development has been completed and an approach that allows incremental deployment has been devised, so this transition is not quite the traumatic transition that is associated with IPv6. But deployment is slow, and of the current level of adoption of the larger AS number set is, oddly enough, comparable to IPv6, at a level of around one-thousandth of the total AS number pool. The routing system has also been growing inexorably, and the capability of switching systems to cope with ever larger routing tables while at the same time offering continual improvements in cost-efficiencies is now looking less certain. So, once again we appear to be examining routing protocol theory and practice, and looking at alternate approaches to routing that can offer superior scaling properties to BGP for the future.

No listing of the major highlights in IP over the past decade would be complete without some mention of the perennial issue of *location* and *identity*.^[25] One of the original simplifications in the IP architecture was to place the semantics of identity, location, and forwarding into an IP address. Although that process has proved phenomenally effective in terms of simplicity of applications and simplicity of IP networks, it has posed some serious challenges with regard to mobility, routing, and network management. Each of these aspects of the Internet would benefit considerably if the Internet architecture allowed identity to be distinct from location. Numerous efforts have been directed at this problem over the past decade, particularly in IPv6, but so far we really have not arrived at an approach that feels truly comfortable in the context of IP.

So although it is possible to observe that not much has happened at the IP level in the past decade that is deployed in the Internet—and IP is still IP—there is still a considerable agenda to tackle at the Internet layer.

The Transport Layer

A decade ago, in 1998, the transport layer of the IP architecture consisted of the *User Datagram Protocol* (UDP) and TCP, and the network usage pattern was around 95-percent TCP and 5-percent UDP. Here, as well, not much has changed in the intervening 10 years.

We have developed two new transport protocols, the *Datagram Congestion Control Protocol* (DCCP) and the *Stream Control Transmission Protocol* (SCTP)^[12], which can be regarded as refinements of TCP to cover flow control for datagram streams in the case of DCCP and flow control over multiple reliable streams in the case of SCTP. However, in a world of transport-aware middleware that is the Internet today, the level of capability to actually deploy these new protocols in the public Internet is marginal at best.

TCP has proved to be remarkably resilient over the years, but as the capacity of the network increases the ability of TCP to continue to deliver ever faster data rates over distances that span the globe is becoming a significant concern. Recent times have seen much work to devise revised TCP flow-control algorithms that still share the network fairly with other concurrent TCP sessions, yet can ramp up to multigigabit-per-second data-transfer rates and sustain those rates over extended periods^[13]. At this stage much of this work is still in the area of research and experimentation, and TCP today as deployed on the Internet is much the same as TCP of a decade ago, with perhaps a couple of notable exceptions. The latest TCP stack from Microsoft in Vista uses dynamic tuning of the Receive window, and a larger inflation factor of the Send window in congestion avoidance where there is a large bandwidth delay product, and improved loss-recovery algorithms that are particularly useful in wireless environments. Linux now includes an implementation of *Binary Increase Congestion control* (BIC), which undertakes a binomial search to reestablish a sustainable send rate. Both of these approaches can improve the performance of TCP, particularly when sending the TCP session over long distances and trying to maintain high transfer speeds.

The Application and Service Layer

This area, unlike the transport layer, has seen quite profound changes over the past decade. A decade ago the Internet was on the cusp of portal mania, where *LookSmart* was the darling of the Internet boom and everyone were all trying to promote their own favorite “one stop shop” for all their Internet needs. We were still using various forms of hand-compiled directories, and navigation of the Internet was still the subject of various courses and books.

By 1998 *AltaVista* has made its debut, and change was already evident. This change, from directories and lists to active search, completely changed the Internet. These days we simply assume that we can type any query we have into a search engine and the search machinery will deliver a set of pointers to relevant documents. Each time this process occurs our expectations about the quality and utility of search engines are reinforced, and we have moved beyond swapping URLs as pointers and simply exchange search terms as an implicit reference to the material. Content is also changing as a result, because users no longer remain on a “site” and navigate around the site. Instead users are directing the search engines, and pulling the relevant page from the target site without reference to any other material.

Another area of profound change has been the rise of active collaboration over content, best typified in wikis. *Wikipedia* is perhaps the most cited example of user-created content, but almost every other aspect of content generation is also being introduced into the active user model, including *YouTube*, *Flickr*, *Joost*, and similar content.

Underlying these changes is another significant development, namely the changes in the content economy. In 1998 content providers and ISPs were competing for user revenue. Content providers were unable to make pay per view and other forms of direct financial relationship with users work in their favor, and were arguing that ISPs should fund content, because, after all, the only reason that users paid for Internet access was because of their perceived value of the content. ISPs, on the other hand, promoted the idea that content providers were enjoying a free ride across the ISP-funded infrastructure, and content providers should contribute to network costs. The model that has gained ascendancy as a result of this unresolved tension was that of advertised-funded content services, and this model has sustained a vastly richer, larger, and more compelling content environment.

At the same time the peer-to-peer network has emerged, and from its beginnings as a music-sharing subsystem, the distributed data model of content sharing now dominates the Internet with audio, video, and large data sets now using this form of content distribution and its associated highly effective transport architecture. Various measurements of Internet traffic have placed peer-to-peer content movement at between 40 and 80 percent of the overall traffic profile of the network.

In many ways applications and services have been the high frontier of innovation in the Internet in the past decade. An entire revolution in open interconnection of content elements is embraced under the generic term *Web 2.0*, and “content” is now a very malleable concept. It is no longer the case of “my computer, my applications, and my workspace” but an emerging model where not only the workspace for each user is held in the network, but where the applications themselves are part of the network, and all are accessed through a generic browser interface.

Any summary of the evolution of the application space over the last decade would not be complete without noting that whereas in 1998 the Internet was still an application that sat on top of the network infrastructure used to support the telephone network, by 2008 voice telephony was just another application layered on the infrastructure of the Internet, and the Internet had even managed to swallow the entire telephone number space into its DNS, using an approach called *ENUM*^[14].

The Business Layer

As much as the application environment of the Internet has been wildly erratic over the past decade, the business environment has been unpredictable as well, and the list of business winners and losers includes some of the historical giants of the telephone world as well as the Internet-bred new wave of entrants.

In 1998, despite the growing momentum of public awareness, the Internet was still largely a curiosity. It was an environment inhabited by geeks, game players, and academics, whose rites of initiation were quite arcane. As a part of the data networking sector, the Internet was just one further activity among many, and the level of attention from the mainstream telco sector was still relatively small. Most Internet users were customers of independent ISPs, and the business relationship between the ISP sector and the telco was tense and acrimonious. The ISPs were seen as opportunistic leeches on the telco industry; they ordered large banks of phone lines, but never made any calls; their customers did not hang up after 3 minutes, but kept their calls open for hours or even days at a time, and they kept ordering ever larger inventories of transmission capacity, yet had business plans that made the back of an envelope look professional by comparison. The telco was unwilling to make large long-term capital investments in additional infrastructure to pander to the extravagant demands of a wildcat set of Internet speculators and their fellow travelers. The telco, on the other hand was slow, expensive, inconsistent, ill-informed, and hostile to the ISP business. The telco wanted financial settlements and bit-level accounting, whereas the ISP industry appeared to manage quite well with a far simpler system of peering and tiering that avoided putting a value on individual packets or flows^[15]. This relationship was never going to last, and it resolved itself in ways that in retrospect were quite predictable. From the telco perspective it quickly became apparent that the only reason the telco was being pushed to install additional network capacity at ever increasing rates was the requirements of the ISP sector. From the ISP perspective the only way to grow at a rate that matched customer demand was to become one's own carrier and to take over infrastructure investment. And, in various ways, both outcomes occurred. Telcos bought ISPs, and ISPs became infrastructure carriers.

All this activity generated considerable investor interest, and the rapid value escalation of the ISP industry and then the entire Internet sector generated the levels of wild-eyed optimism that are associated only with an exceptional boom. By 2000 almost anything associated with the Internet, whether it was a simple portal, a new browser development, a search engine, or an ISP, attracted investor attention, and the valuations of Internet start-ups achieved dizzying heights. Of course one of the basic lessons of economic history is that every boom has an ensuing bust, and in 2001 the Internet bust happened. The bust was as inevitable and as brutal as the preceding boom was euphoric. But, like the railway boom and bust of the 1840s, when the wreckage was cleared away, what remained was a viable—and indeed a valuable—industry.

By 2003 the era of the independent retail ISP was effectively over. ISPs still exist, but those that are not competitive carriers tend to operate as IT business consultants that provide services to niche markets. Their earlier foray in to the mass market paved the way for the economies of scale that only the carrier industry could implement on the market.

But the grander aspirations of these larger players have not been met, and effective monopoly positions in many Internet access markets have not translated to effective control over the user's experience of the Internet, or anything even close to such control. The industry was already "unbundled," with intense competition occurring at every level of the market, including content, search, applications, and hosting. The efforts of the telco sector to translate their investment into mass-market Internet access into a more comprehensive control over content and its delivery in the Internet has been continually frustrated. The content world of the Internet has been reinvigorated by the successful introduction of advertiser-funded models of content generation and delivery, and this process has been coupled with the more recent innovations of turning back to the users themselves as the source of content, so that the content world is once again the focus of a second wave of optimism, bordering on euphoria.

And Now?

It has been a revolutionary decade for us all, and in the last 10 years the Internet has directly touched the lives of almost every person on this planet. Current estimates put the number of regular Internet users at 19 percent of the world's population.

Over this decade some of our expectations were achieved and then surpassed with apparent ease, whereas others remained elusive. And some things occurred that were entirely unanticipated. At the same time very little of the Internet we have today was confidently predicted in 1998, whereas many of the problems we saw in 1998 remain problems today.

What we have today is not the technical Internet we thought we were building a decade ago. It is not a coherent end-to-end network with clear signaling across commodity packet switching fabric, but a network that is replete with all forms of active middleware^[16], from NATs to firewalls^[17] and filters, including packet shapers, torrent detectors, *Voice over IP* (VoIP) blockers, and load balancers. It is neither a secure nor a safe network, but one that includes a continual barrage on end hosts in the form of more than a million different forms of viruses^[18], worms, and assorted malware^[19], as well as a barrage on users in the form of torrents of spam^[20]. The network is a host to a litany of hostile attacks, including gigabit traffic swamping attacks, redirection, inspection, passing off, and denial-of-service attacks^[21]. The attacks are directed at links, routers^[22], the routing protocols^[23, 24], hosts, and applications. Our ability to effectively defend the network and its connected hosts continues to be, on the whole, ineffectual. Our level of interest in paying a premium to support highly secure systems still remains slight. But somehow we are not deterred by this situation. Somehow each of us has found a way to make our Internet work for us.

I am not sure that the next decade will bring the same level of intensity of structural change to the global communications sector, and perhaps that is a good thing given the collection of other challenges that are confronting us all in the coming decades. At the same time I think it would be good to believe that the past decade of development of the Internet has completely rewritten what it means to communicate, rewritten the way in which we can share our experience and knowledge, and, hopefully, rewritten the ways in which we can work together on these challenges.

References

The Internet Protocol Journal (IPJ) has published articles on all the major aspects of the technical evolution of the Internet over the past decade. To illustrate the extraordinary breadth of these articles, I have included as references here only articles that have been published in the IPJ.

- [1] Stallings, W., "Mobile IP," *IPJ*, Volume 4, No. 2, June 2001.
- [2] Handley, M., and Crowcroft, J., "Internet Multicast Today," *IPJ*, Volume 2, No. 4, December 1999.
- [3] Stallings, W., "IP Security," *IPJ*, Volume 3, No. 1, March 2000.
- [4] Huston, G., "QoS — Fact or Fiction?" *IPJ*, Volume 3, No. 1, March 2000.
- [5] Stallings, W., "MPLS," *IPJ*, Volume 4, No. 3, September 2001
- [6] Ferguson, P., and Huston, G., "What is a VPN?" *IPJ*, Volume 1, No. 1 & No. 2, June & September 1998.
- [7] Fink, R., "IPv6," *IPJ*, Volume 2, No. 1, March 1999.
- [8] Huston, G., "Anatomy: Inside Network Address Translators," *IPJ*, Volume 7, No. 3, September 2004.
- [9] Huston, G., "The BGP Routing Table," *IPJ*, Volume 4, No. 1, March 2001.
- [10] Huston, G., "Scaling inter-Domain Routing," *IPJ*, Volume 4, No. 4, December 2001.
- [11] Huston, G., "Exploring Autonomous System Numbers," *IPJ*, Volume 9, No. 1, March 2006.
- [12] Huston, G., "The Future for TCP," *IPJ*, Volume 3, No. 3, September 2000.
- [13] Huston, G., "Gigabit TCP," *IPJ*, Volume 9, No. 2, June 2006.

- [14] Huston, G., "ENUM," *IPJ*, Volume 5, No. 2, June 2002.
- [15] Huston, G., "Peering and Settlements," *IPJ*, Volume 2, No. 1 & No. 2, March & June 1999.
- [16] Huston, G., "The Middleware Muddle," *IPJ*, Volume 4, No. 2, June 2001.
- [17] Avolio, F., "Firewalls and Internet Security," *IPJ*, Volume 2, No. 2, June 1999.
- [18] Fraser, B., Rogers, L., and Pesante, L., "Was the Melissa Virus So Different?" *IPJ*, Volume 2, No. 2, June 1999.
- [19] Chen, T., "Virus Trends," *IPJ*, Volume 6, No. 3, September 2003.
- [20] Crocker, D., "Challenges in Anti-Spam Efforts," *IPJ*, Volume 8, No. 4, December 2005.
- [21] Patrikakis, C., Masikos, M., and Zouraraki, O., "Distributed Denial of Service Attacks," *IPJ*, Volume 7, No. 4, December 2004.
- [22] Lonvick, C., "Securing the Infrastructure," *IPJ*, Volume 3, No. 3, September 2000.
- [23] Kent, S., "Securing BGP: S-BGP," *IPJ*, Volume 6, No. 3, September 2003.
- [24] White, R., "Securing BGP: soBGP," *IPJ*, Volume 6, No. 3, September 2003.
- [25] Meyer, D., "The Locator Identifier Separation Protocol (LISP)," *IPJ*, Volume 11, No. 1, March 2008.

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. The author of numerous Internet-related books, he 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. E-mail: gih@apnic.net