

---

*Author: Geoff Huston*

---

By any number of measures the Internet is a relatively mature technology. Whether you start the clock with the original deployment of the ARPANET in 1969, or move forward to 1973 with the initial TCP/IP specification activity, it's still some three decades on from these initial efforts, and for a communications technology thirty years is a pretty long time indeed.

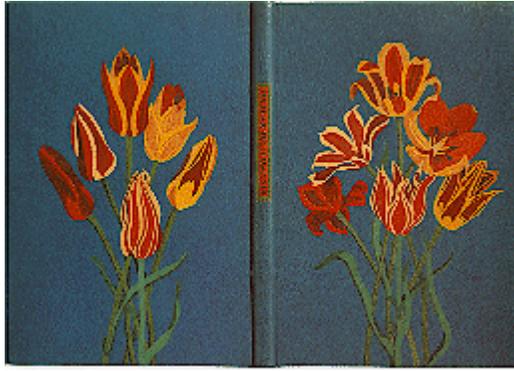
There are many histories of the Internet. I'd recommend <http://www.isoc.org/internet/history/brief.shtml> as a good write up of the early history of the Internet, written by those who were there at the time.

### Boom and Bust

Many technologies start life with a grand fanfare, and then settle down into a quieter mode as they assume an accepted mode of operation. And this may probably be the longer term outcome for the Internet, for the first three decades its worked in the opposite fashion. For most of this time the Internet has been a low key affair, of interest in the academic and research community as a useful testbed for experimentation into packet switched network architectures. The precise date may vary depending on the source, but it was around 1995 that the Internet started to truly shed its experimental trappings and seriously challenge the assumptions of the incumbent players in the communications industry. At this time the assertion was forming that the Internet could be a stable platform for commercial service provision, and, yes, there was a sustainable market for selling Internet access, email services, web-hosting services and Internet-based e-commerce solutions.

There were a few other factors at play as well, of course. The Internet was sufficiently different in its architecture and service profile that it was not just a case of changing the presentation format on existing data communications services. This was no change of the surface veneer, but a more basic change of the way in which the underlying common transmission and switching resource was to be shared between customers. This implied that incumbent operators were not in a position to react quickly enough to meet the rapidly escalating demands for Internet services. At the same time, in many parts of the world, the industry was being deregulated, and large markets were being opened up to private investment and competitive service provision. Further factors were at work, including a widespread acknowledgement that there was no clear understanding of just how large the Internet market could become, nor in what timescale. Not only had personal computers become so widely used that they were now being marketed in the same way as other items of consumer electronics, but the Internet was extending its reach into other areas of embedded devices and control operations. When you add to this a period of social wealth and a very large investment market in search of a new opportunity, then the conditions were set. A massive boom in the Internet was the only possible outcome.

Of course all this is nothing new. Anyone attempting to purchase a tulip in 1637 would've seen the same mania that exists at the height of a boom



There are some strong parallels between the Internet boom and the Tulip boom some four centuries ago. The site [itulip.com](http://itulip.com) explores some of the parallels in more detail. Persian and

Turkish tulips were imported into Holland in the latter half of the sixteenth century and were considered highly fashionable in the gardens of the rich. Breeding of special species took place, and due to a virus affecting bulbs, rare colours and streaks were created. One of the more popular and expensive varieties of tulip was *Semper Augustus*, and, at the height of the boom in February 1637 as much as 30,000 guilders was offered for three bulbs. To put this into some form of perspective, the average annual income of the time in Holland was 150 guilders, the most expensive houses in the centre of Amsterdam cost 10,000 guilders, and Rembrandt received less than 2,000 guilders for "The Night Watch". The social consequences of the consequent crash were severe.

Satirical engravings of Tulipmania are to be found at the Koninklijke Bibliotheek in Den Hague (<http://www.kb.nl>).

Less than a hundred years later in France the same cycle of boom and bust was played out with the Bank Royale and the infamous Scotsman John Law.

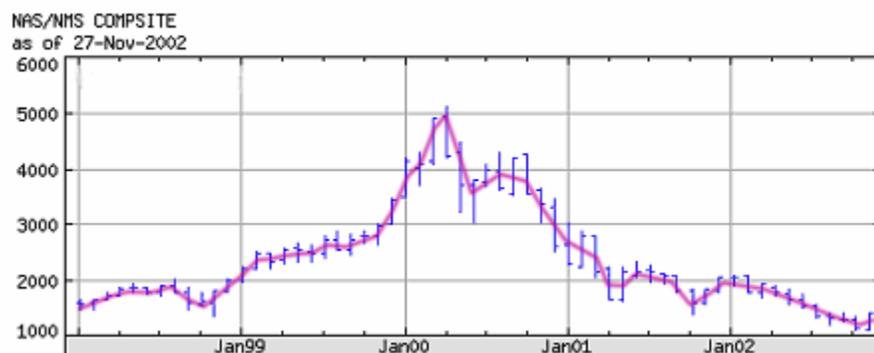


The annals of the beginnings of the banking system are filled with colourful characters. One of the more reckless, unbalanced fascinating geniuses of banking was John Law, who ingratiated himself into the French Court. His solution to the court's financial woes was truly inventive, and captured the imagination, and much of the money, of the French population. The term "millionaire" was first coined in 1719 to describe those who had invested a few thousands early in 1719 and were worth millions in a matter of weeks. The jam of people seeking to buy the stock was dense and the din was evidently deafening. The invested money was used to pay the expenses of the Government of France. In recognition of Law's financial genius he was made Comptroller General of France in early 1720. The bust happened in July that year, and the jam of people seeking to liquidate their investments was dense – so dense that 15 people were squeezed to death. The consequent crash so alienated the French from the term "banque" that "credit" has been used in France to this day to describe banking institutions.

Booms have occurred in land, gold, banking, currency speculation, oil, railways and in many others forms of goods and services. In each case the consequent bust is entirely inevitable. At some point in an inflated market the sobering view of reality intrudes uncomfortably into the wild-eyed flights of optimistic fantasy. And in the same way that the boom feeds on itself to reach ever higher levels of optimism, the consequent bust turns caution into panic as investors seek to leave the market, and the bust also feeds on itself. Indeed its not that the tendency of markets to operate in boom and bust cycles that should be so surprising, but that the evolution of regulatory-based control mechanisms over the centuries should continue to be so ineffectual in protecting our economy from the ravages of this erratic and disruptive mode of operation.

And it played through in a thoroughly predictable fashion in the Internet. The initial cautiously optimistic investment steps in the Internet were replaced with a stronger more assertive optimism. Optimism turned to mania with a rush of new services, operators, technologies and investors, all attempting to stake a claim in the future value of the Internet. In turn this was replaced by a sense of euphoria that the Internet sector would continue to rise indefinitely, until, in the final phase, the investment was not related to the business fundamentals of the industry, but more a case of attempting to squeeze additional growth out of the market before the inevitable crash. And, on the 6th March 2000, the edifice of investor optimism was simply too high to be matched by reality, and the first words of cautious doubt were amplified into a selling mania.

Perhaps the best boom and bust illustration of the Internet can be found in a 5 year tracking of the NASDAQ composite index. The sharp rise and fall at the start of 2000 is clearly evident



So is this the end of the Internet? Of course not! Tulips still blossom in spring across the world, railways still move freight and banks still conduct financial transactions. And, yes, the Internet will still be around for many years to come! But yes, some things have changed, and a combination of a post-bust business conservatism, couple with a renewed interest in the security of the Internet is making some quite profound changes to the Internet.

So in looking forward to a post-boom and bust Internet here are perhaps a number of observations we can make about likely futures.

## Dot-Post-Euphoria-Conservatism

The first of these is a likely swing to conservative business investment patterns, with a consequence that Internet expansion will no longer rocket ahead at rates of doubling in size every 12, 8 or even 6 months. Business emphasis on growth as the primary indicator is being replaced with an emphasis on return on investment and the current period is seen as one of consolidation of activities and their related revenue levels rather than an all out push to grow the total size of the market at any price.

Conservatism is also visible in a strong focus on security at present. There is a visible effort underway across the Internet to dismantle the worst excesses of a widely distributed trust model of Internet infrastructure and applications and replace it with a model of explicitly negotiated conditional trust. This is coupled with a renewed interest in the use of encryption and authentication within many levels of the IP protocol suite.

The related area of focus is a greater emphasis on service robustness rather than service differentiation. There is a growing appreciation that the Internet is now a major component of our communications infrastructure and disruptions to the service will have social and commercial impacts for both small and large communities. Within this examination of the level of resilience of the Internet's infrastructure there is bound to be a second debate hot on its heels: namely a debate as to whether a competitive, private-sector funded, diverse and largely deregulated environment is truly capable of making the necessary investment and technology decisions to achieve appropriate levels of infrastructure robustness, or whether some level of regulatory involvement is necessary in order to ensure that these social and strategic objectives are to be achieved for all. It promises to be an interesting debate across the many regulatory regimes in existence today.

## Goodbye to Convergence?

The increasing concentration on business fundamentals has also had a visible impact on the long-standing debate over convergence within the industry. For many years it was a cherished belief that it was possible to create a single service platform that would efficiently carry both the real-time load of voice conversations and video streams, and at the same time carry data traffic., and for many years the ATM architecture carried the bulk of this agenda of platform convergence. More recently this agenda was re-homed within IP and the converged service platform was seen as being related to the future of IP.

It appears that such optimism is, if not poorly based, at least premature, and its not yet time to throw away those time division multiplexors and time space switches. Creating real time outcomes from IP packet switched networks is hard and expensive work, and for the moment it appears to be too expensive for the industry to pursue it with a vengeance. There is no more talk about taking the entire communications service load and placing it on an IP foundation network. Instead there is a more visible emphasis for using a set of technologies within a service provider's suite of offerings, using real-time platforms for the core of the real time traffic load, and using a variety of data solutions, including ATM, Frame Relay, IP and even a resurgence of wide area Ethernet for the data market. With the advent of multiple wavelength multiplexing in both the short and long haul transmission markets, and the continual downward pressure on the cost of silicon-based switching, the use of an engineering base of multiple platforms with service specialty in each platform is now a viable approach to service the market.

Part of this shift is also due to a more sober view of the strengths and weaknesses of IP as a platform. IP's major strength is in supporting adaptable traffic sessions to operate highly efficiently over wired networks. IP's weaknesses become apparent when the load consists of a major component of real-time traffic to be mixed with adaptable traffic, or when the application is mobile high speed large coverage wireless, or in complex traffic engineering scenarios and in working with applications that require resource management and resource usage enforcement. While these weaknesses can be overcome with various engineering measures, the major issue is the efficiency of the resultant service platform. Unless IP represents immediate cost savings for an operator then its adaptation into any particular service offering is not a foregone conclusion.

## Areas of IP Development

It should not be assumed that a period of consolidation is entirely devoid of further development and refinement of the architecture of the service platform. There are a number of areas of further refinement of the IP architecture that appear likely in the near future.

There remains much work in the area of IPv6. While it's always been hard to separate the myth from the reality whenever IPv6 is concerned, there does appear to be some sound reasons why V6 is going to be useful. It's not that IPv6 offers superior security, superior QoS, superior auto-configuration or superior address structuring – all these are debating points where you will find points in favour and points against such suppositions. The main reason lies in the longer address field in IPv6. This allows for a large number of connected devices without the need for various forms of network middleware that attempt to use various assumptions about network application behaviour to map addresses from one realm into addresses from another realm on the fly. Yes, I'm talking about the problems of dynamic Network Address Translation units. Part of the reason for the success of the Internet architecture is that the Internet network itself makes very few assumptions as to the nature of the applications that use the network. The network's task is to maintain a consistent view of address reachability so that individual packets can be locally switched towards their destination. Attempts to extend this role invariably have to make limiting assumptions about application behaviour and such assumptions end up restricting what can be supported on the network, and this, in turn, restricts the potential utility of the network. So I'm of the view that there is value in an agenda that moves deployment of V6 beyond the experimental code sets and starts to engage in using the Internet for a broader world of devices and device-based communications.

The next area of likely activity is in the IP Virtual Private Network space. This has been a topic of considerable interest in the IP Service Provider world for some time, as it offers the service provider the ability to provide a higher level of service to the customer than simple commodity-based packet transit services. If it was an easy problem we'd all be drowning in VPNS today! The fact that we are some way away from that fate tends to suggest that it's quite a complex problem. Creating multi-layer routing domains, and the associated multi-layer signaling infrastructure, not to mention VPN-segmented support services such as DNS, access control, mail and web services all add further grist to this particular mill.

I also suspect that the network management picture is one that requires further attention. Much of the world's IP networks continue to operate on a base of the Simple Network Management Protocol, and there is a growing feeling of unease regarding this approach to network management. The protocol itself operates at a very basic level, and assumes a lowest common denominator approach to equipment and management operations. There is considerable potential to combine some of the experiences in the world of distributed object computing with the concept of intelligent objects, allowing the network management system to directly address the basic objective of network management: to alert the network operator for intervention on those events where the conditions are abnormal and automated recovery processes are not seen as being able to restore the network to a stable state.



The city of Oulu in Finland may have many claims to fame, but in IP circles it's the work of researchers at the University of Oulu that has done the trick. Their work on protocol conformance unearthed a set of serious errors in many implementations of the Simple Network Management Protocol (SNMP). This received worldwide

attention in February 2002, when details of their work were made public. The implementation errors were such that it had the potential to corrupt the memory map of the managed device, and exploitation of this could cause the device to head into a continuous cycle of rebooting, or allow the device to be 'captured'. Not only were the errors serious, they were also quite widespread across many deployed products.  
<http://www.ee.oulu.fi/research/ouspg/protos>

Sometimes using the work "simple" in a protocol name is not enough to master all the intricacies and complexities that can arise!

While it may be unrealistic to expect that the world's telephone traffic will switch to IP anytime soon, that does not mean that Voice over IP (VOIP) is going away. And if VOIP is not going away, the problem remains of how to make the various VOIP gateways aware of each other. I suspect that the use of the DNS to map telephone numbers into IP service addresses, or ENUM, will be a feature of the next few years of interworking between VOIP and the public switched telephone network. The ENUM solution is certainly elegant, but there are a number of rather thorny aspects that promise to keep us occupied for some years to come. Is a telephone number purely an addressing scheme for with within the telephone network? Or is it a universal device identifier for anything that can originate or terminate conversations irrespective of the form of last mile access network? Who "owns" a telephone number in such a multi-service platform world? And what regulatory measures apply to phone calls even when no part of the call transits any part of the public switched telephone network? What are the privacy implications of exposing part of the mapping from telephone number of IP address in the DNS? As you can appreciate there is ample fodder to fuel a public policy debate for some time to come/

No list of current activities in IP would be complete without mentioning the area of IP and mobility. There are two issues that are exposed here, firstly the issue of maintaining a constant identity while supporting a changing network location, and secondly the issues of IP over wireless subnets, where the issues of attempting to clearly delineate network congestion and wireless packet corruption come to the forefront. Also there is the consideration of the differences between nomadism and roaming mobility. With the massive investment in 3rd Generation wireless spectrum space in recent years, more notably in Europe, there is a considerable impetus to develop and deploy some form of technology that will support high speed and presumably high volume applications across wide coverage radio access systems.

## Where to from here?

Well, its safe to make one prediction - that the Internet is not scheduled to disappear anytime soon. The match between packet-based networks and computing applications is simply the right one. To revert to a time-space switching system, or to an approach that uses various forms of dynamically

created circuits as the fundamental component of network architecture is difficult to conceive. So boom or bust, its the Internet, and we're stuck with it.

However, there's no doubt that the landscape of the Internet will continue to change, but the uncritical acceptance of every bold new concept, whether in Internet marketing or Internet technology is over. Its back to an environment where business basics are reasserting their primacy, and investors are looking for enterprises that present a credible business proposition with minimal risk.

Meet the new economy. Same as the old economy.



The classic The Who song, written by Pete Townshend, *Won't Get Fooled Again* was first recorded as part of the aborted LifeHouse project in early 1971. It was re-recorded with a synthesizer track in April 1971 and released as a single and on the *Who's Next* album in August 1971. This song formed the climax of their stage set. This song is about the same age as the Internet.

---

## Disclaimer

The above views do not represent the views of the Internet Society, nor do they represent the views of the author's employer, the Telstra Corporation. They were possibly the opinions of the author at the time of writing this article, but things always change, including the author's opinions!

---

## 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 the past decade, particularly within Australia, where he was responsible for the initial build of the Internet within the Australian academic and research sector. Huston is currently the Chief Scientist in the Internet area for Telstra. He is also a member of the Internet Architecture Board, and is the Secretary of the APNIC Executive Committee. He was an inaugural Trustee of the Internet Society, and served as Secretary of the Board of Trustees from 1993 until 2001, with a term of service as chair of the Board of Trustees in 1999 – 2000. He is author of *The ISP Survival Guide*, ISBN 0-471-31499-4, *Internet Performance Survival Guide: QoS Strategies for Multiservice Networks*, ISBN 0471-378089, and coauthor of *Quality of Service: Delivering QoS on the Internet and in Corporate Networks*, ISBN 0-471-24358-2, a collaboration with Paul Ferguson. All three books are published by John Wiley & Sons.

E-mail: [gih@telstra.net](mailto:gih@telstra.net)