

June 2009

Geoff Huston

Twenty Years Later

As I write this, on the 23rd June 2009, I've been reminded that some 20 years ago, on the night of the 23rd June 1989, Robert Elz of the University of Melbourne and Torben Neilsen of the University of Hawaii completed the connection work that brought the Internet to Australia with a permanent connection via a 56kbps satellite circuit.

Since that day we've evidently connected some 56.8% of the local population, or 12,073,852 Australians, to the Internet (according to recent Internet user statistics published by the ITU-T). While that's an impressive outcome, I suppose I should say at the outset that when we started down this path in Australia some twenty years ago we had no intention of achieving this scale of outcome. Indeed we never thought that this type of data networking would ever cross the boundary from an esoteric tool to assist a select group of computer literate researchers and academics into the mainstream of society, and current concepts like twitter and social networking were completely foreign to us. In truth all we were trying to do was to save a bit of money for the universities and have some fun experimenting with some pretty novel technology on the way. At best this was all just an experiment and if there was ever going to be some mainstream commercial outcome, then that was the for the telephone companies to work on, and was certainly none of our business.

So how and why did all this get started?

The 1980's

The 1980's saw the end of one generation in computing and the start of the next generational shift. The era that was ending was the era of the mainframe computer. And the next generation was a combination of three quite revolutionary technologies: the personal computer, the Ethernet Local Area network, and Unix, an open source operating system.

Prior to the introduction of personal computer technology, computers were large and environmentally specialized systems that were located in protected environments with conditioned power, air, and access. The access network that reached out from such hosts was typically one of twisted pair wires supporting the RS-232 serial line protocol to allow the dedicated connection of data entry terminals. The PC introduced the wave of cheap personal computers on desktops, replacing this model of large central host computers and its terminals with a model of a distributed data environment. PCs were capable of exchanging data without a central common data repository. With this change in the landscape of computing came a change in the landscape of networking. Networking protocols were now the means by which these PCs were linked into the local environment. Networking protocols now had to provide the 'glue' for this distributed environment, which was originally hard wired into the single mainframe host computer. Network protocols now had to support access to shared storage and distributed file systems, as well as supporting access to common resources, such as printers and mass archival storage, as well as supporting the more traditional host-to-host models of remote terminal access, bulk data transfer, and electronic messaging.

Ethernet was to assist this change by introducing a matching change in the environment of the local area network. The original access networks were constructed using dedicated 2-pair copper wires, spanning out from the computer room or wiring hub to each workstation. The access system was designed to operate at very slow serial transmission speeds of 9600 baud or less. This speed was simply inadequate for the demands of a distributed environment of PCs. Ethernet changed all that in a very radical fashion. The early prototypes supported a common wiring plant operating as 3Mbps, while the production equipment was configured to operate at 10Mbps. Ethernet also changed the existing wiring topology, replacing the wiring hub closet and

radial spans with a single cable that passed every work station. This form of extremely high speed local networking spawned a diverse set of developments in networking.

Unix was the third component of this technology trio. The computing environment to that time typically encompassed a software environment which was customized to the hardware platform. Clients were locked into vendor-specific information technology environments, which presented few options for portability, and allowed vendors to lock in customers. Into this environment came Unix. For academic institutions, Unix was offered worldwide at nominal license fees, and allowed the client to migrate between an increasingly large range of hardware platforms, while preserving a software environment.

By the late eighties the computing environment for any large scale enterprise was hopelessly chaotic. There were a variety of proprietary network protocols and each protocol was linked to a specific set of underlying network types. Some protocols functioned only on Ethernet LANS, for example. And of course interoperability between these different protocol environments was non-existent.

Protocols, Protocols, Protocols

The mid 1980's saw a blossoming of all kinds of network protocol suites. There was Appletalk, SNA, IPX, DECnet, and X.25 to name but a few. At one point it appeared that each and every major vendor had developed its own proprietary network protocol.

What was being used in the academic and research environment at the time was much the same eclectic mix of protocols as you'd find in any IT shop. However one protocol had gained some prominence in that community, and that was DECnet. DECnet served as a common protocol to interconnect the range of DEC minicomputer and mainframe computers. DEC computers were widely used in the academic and research environments in the late 1970s and the 1980s, and in the wake of deployment of these computer systems, DECnet was widely deployed to interconnect these systems. In the academic and research world two of the largest DECnet networks at the time were the High Energy Physics Network (HEPnet) serving the international physics community, and the Space Physics Aeronautical Network (SPAN) in the international space aeronautics community. These networks used local Ethernet networks for on-site connectivity, and a mix of leased circuits and X.25 public switched data services for long-distance circuits. However, DECnet had some serious limitations in its internal addressing model, as it used a 16 bit address field. While a single network of tens of thousands of computers may have seemed like a fanciful notion in the era of mainframe computers, with the advent of PCs and with the accessibility of international digital circuits, such large scale networks were a practical reality, and by the late 1980's were hitting hard upon the basic address limitations of the DECnet address architecture.

As the deployment of PCs gathered in pace the need to provide an interoperable networking framework that could glue together a disparate multi-vendor networked environment became paramount, and the industry commenced along a path of developing an open vendor independent standard networking architecture. The study work undertaken within Technical Committee 97, Information Processing, of the International Organization for Standardization, published its first set of Open Systems Interconnection (OSI) standards, with the CCITT Red Book series of 1984. But this work was largely paperware for much of the 1980s. By the late eighties the effort was gathering some momentum, but the effort was hampered by a typical design-by-committee flaw that when confronted by choices, avoid any decision and instead describe all possible choices in the standard specification. On paper such an approach is easy, but when the choices result in mutually incompatible behaviours then its probably reasonable to start to question the value of the entire committee process. In OSI's case one of the most glaringly conspicuous problems was the inability to make a coherent choice between circuit-based connection services and datagram-based services, and the inclusion of both approaches into the OSI specifications was a particularly poor outcome.

However, there was an alternative open network architecture. Some years earlier the US Defense Advanced Research Project Agency (DARPA) had contracted the University of California, Berkeley, to develop an open source implementation of TCP/IP for Unix, and this implementation, complete with its source code, was freely available to the public. This was a critical step away from vendor-specific environments, into a world of working open source software. TCP/IP had made some very interesting design choices, including the adoption of a 32 bit address architecture, the use of dynamic fragmentation capability to allow for adaptation across disparate networks, and the omission of any form of hop-by-hop flow control or datagram integrity. This code spread quickly into a wide array of computers and by 1989 TCP/IP was available as either third party code or vendor-supported code on most vendors products from mainframes to personal computers and workstations.

The Rise of the Internet

TCP/IP was not just a network protocol that was used sporadically at the local network level. Throughout the 1970s and 1980's parts of the US research community had been involved in a national effort to experiment with a packet switched networking platform, and at the heart of this platform was the TCP/IP protocol suite. Even though the nascent Internet grew steadily throughout this decade, it was a commonly accepted view that it was an experiment, albeit a large scale experiment, and that this Internet effort would generate some valuable insights in terms of understanding of how such packet-switched networks operate and behave. This would become input to the evolving OSI standardization effort, that would be the source of the technology standards underpinning the next generation of public data networks. In the meantime the Internet was also addressing some of the immediate requirements of the research community. The unfolding evolution of technology in this space did not adhere to such an orderly planned regime. The OSI effort, which a number of Internet technologists described as a triumph of dysfunctional vaporware over the requirement to produce functional technological specifications, largely dissipated by the mid 1990's, with the overwhelming attention being focussed on the Internet.

But perhaps I'm getting ahead of myself here. In 1988 the next major transition of the evolving Internet commenced with the initial deployment of NSFnet, a networking project funded by the U.S. National Science Foundation (NSF). This project was perhaps one of the most startling success stories in this history. The project had an original objective of linking the five supercomputer centers with an Internet backbone network. However, it quickly grew beyond that original objective and encompassed a broader objective of servicing the research communications needs of the national academic and research community. Interestingly enough, a decision was made that the network would support only TCP/IP in the first instance, and would not support DECnet, then still in widespread use within the academic and research community. Perhaps two aspects of the 1988 NSFnet program which were truly long sighted and revolutionary. The first of these was to launch a project that had technical specifications, which could be confidently met by no existing services or equipment. The specification of 1.544Mbps transmission rates and high-speed gateways was certainly unique at that time, and provided a high-speed platform which defined new parameters for high speed long distance network performance that was on a par with local area Ethernet systems. The second quite novel aspect of the NSFnet program was the adoption of the organizational model of using the so called 'mid-level' networks as the means of local distribution of the backbone long-haul core network. The mid-level networks not only were a focus for local effort in the development of Internet infrastructure, but these entities had to quickly include a funding base that was more embracing than absolute reliance on federal funding, and, as such, provided an early template for the emerging ISP sector. The prevailing industry model of network operation through national telco monopolies could be summarized by the catch cry of "one network, one operator," while the NSFnet deliberately started with an operation model of "one network, many providers."

The NSFnet in the US certainly captured the attention of many academic and research communities across the world, including in Australia.

The Australian Experience

The Australian community had experimented with various forms of store / forward messaging networks, and by 1989 there was a quite sizeable community in the country that was connected through a locally developed protocol, the Australian Computer Science Network (ACSnet). We used this to support a national hookup that extended across most of the Australian academic and research institutions by the end of the 1980s. No doubt there were many issues with this network, but perhaps the most glaring was its success. The system relied heavily on diverse collection of low speed intermittent modem-based connections and as the level of usage of the network picked up, the volume of traffic started to overwhelm the network. At its worst the end to end delays in message transmission could be measured in days or even weeks. Normally such problems are amenable to the application of additional money, but this network was largely a volunteer effort which had no central funding or internal structure and the task of funding growth proved to be extremely challenging.

It fell to the Australian university sector to take the initiative, and in 1989 the university sector embarked on a centrally funded effort to link all Australian universities together with a cohesive network.

It seems incredible some twenty years later, but the first steps with the network were very tentative. The backbone of the national network was constructed using the digital service equivalent of a single voice circuit,

which in Australia was a 48Kbps circuit. The initial network that was completed in May 1990 had some 50 connected sites. Each site was connected to a single local hub, and each of these hubs connected to a single national hub, using a simple star hierarchy, all running at 48Kbps. There was no particular ulterior motive behind this topology other than simply one of working within a very constrained budget. There was a single international connection, initially provisioned as a 56Kbps satellite circuit connected to the University of Hawaii. This was before the web and was in a time when simple ascii text was the lingua franca of communications, so while by today's standards it sounds like a miniscule amount of capacity, it managed to service the needs of a community of a few thousand, where the basic mode of interaction was email.

The subsequent steps were ones that played out in a similar fashion in many other countries at the time. The take up in the academic and research community was relatively rapid, and the network's traffic volume was doubling every eight months, or an inexorable growth of 2% in total traffic volumes each and every week. The 56Kbps circuit quickly grew to a 128Kbps circuit, then 512K, followed by 768K and by 1993 a 1.5Mbps circuit was installed using the first of the undersea digital cable systems that reached Australia. Each time the additional capacity was consumed and each time more capacity was required. Not only was the community growing, but so were the demands in terms of the richness of the communication. From simple text-based information structuring, such as *gopher* ([http://en.wikipedia.org/wiki/Gopher_\(protocol\)](http://en.wikipedia.org/wiki/Gopher_(protocol))), we moved rapidly onto richer information structuring with *WAIS* (http://en.wikipedia.org/wiki/Wide_area_information_server), and then by the mid 90's its was entirely a world of web services and the communications requirements expanded exponentially. By then this was not all just academics and researchers. The Internet was embarking in a direction that followed the transition of the personal computer into the realm of consumer electronics, and by then the Internet had expanded in scope to encompass the mainstream communication service provider sector.

The Commercial Internet

As the academic and research Internet gathered pace, so did the pressure to construct a more commercial orientation of networking. The U.S. Federal government-Agency funded networks operated under an "Appropriate Use Policy", which effectively prevented their use for commercial purposes. Far from stifling commercial demand, such a restriction acted as a catalyst for commercial service providers to appear. UUNET in the US was one of the first commercial ISPs to meet this emerging market demand. The initial trial gathered momentum in the US in the early 1990s and the Internet's story then saw a rapid multiplication as commercial ISP operators made their debut. The interaction between these networks quickly became the area of prime focus. Rather than following a single thread of development, the framework of inter-provider interconnection became the major theme at this point in time. Indeed, this theme of interconnection, both as an aspect of network engineering, and as a commercial interaction between ISPs, has been an enduring aspect of the ISP domain.

By 1995 the Internet had attracted the interest of many of the incumbent telcos, and commercial offerings were appearing in the consumer market, based on low cost dial-up access, then moving on to DSL services by the end of the 1990's. Many of the initial academic and research networks, including AARNet in Australia changed from a focus on operating a dedicated special purpose IP network into a purchasing group, buying wholesale IP services from the commercial sector, and a rather unique period of the evolution of the Internet drew to a close.

In many ways by 1995, the Internet was firmly locked into place as the foundational technology for the next generation of mainstream telecommunications services. For the small group of network engineers in Australia where I was working at the time it was a fast paced 6 years from start to finish, but we had managed to do what we never expected to achieve at the start - we'd help to transformed a tiny research experiment into an entire industry. Twenty years later that still seems like a worthy achievement.



There's not much left of what was used to construct IP networks 20 years ago, but these two items are pretty representative. Many of the early milestones of the Internet were marked by T-shirts, and this is the T-shirt we produced to mark the start of AARNet. Underneath it is one of the very early Cisco router models, the Hybridge, which was equipped with one Ethernet port, one serial port and a single console interface. This particular one was used for some years in the early 1990's to connect the Australian Internet to the US Internet infrastructure.

Disclaimer

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

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>