New Technology Briefing

An overview of the Internet

Paul Eccleson

Date received: 22 March 1999

Paul Eccleson is manager of electronics business research at Hewlett-Packard's European research labs. He has been involved with the development of core World Wide Web technologies for over five years.

Keywords:

Internet, World Wide Web, history, marketing, personalisation

Abstract

This paper takes a broad look at the history and construction of the Internet, and tracks the evolution of the Internet from early concept as a military communications backbone through to its present commercial orientation. Specific references are drawn to the advent of marketing and personalisation on the Internet, and the possibilities for new business models and open competition.

How it all started

Imagine that you consider yourself to be the world's largest superpower. You begin to notice that the command-and-control capability for your military might is a series of point-to-point connections. Worse than this, the new technology of computing is proliferating across all your sites and doesn't communicate with any of the other sites very well. There's a real danger that a thoughtfully placed attack on a couple of key links in your communications could disable your military efforts very quickly. You're probably worried enough to spend a fair amount of taxpayers' money on fixing this.

This is roughly the situation that the US Air Force responded to in 1962. Paul Baran, of the RAND Corporation (a US government agency), was commissioned to find a way around this problem. His report recommended an extensive 'packet-switched' network. Messages would be cut into small chunks, each addressed with destination and sender information, and then sent from computer to computer in the linked network. If an individual computer were lost, the network could reroute packets via other machines. If packets were lost, the receiver could request a resend from the originator. The Internet was born.

The Defense Advanced Research Projects Agency (DARPA) began to fund an extensive development of the networking infrastructure across military and academic establishments in the US. As the research continues, the network establishes its own basic communications protocol (TCP/IP), basic services (e-mail), local area network and dial-up connectivity and a number of useful services for transferring data (eg the file transfer protocol and simple mail transport protocol). These are all based on the basic packet switching of the Internet and establish its use among academic institutions in the US.

Today, the Internet spans the globe, with most countries having their own Internet backbone, and with a large commercial underpinning to the funding and expansion of the network. The same basic protocols and architecture are still underpinning the Internet — although some are beginning to show signs of this global strain.

Packet-switched networks created to withstand attack

Paul Eccleson,
Hewlett-Packard Labs,
Filton Road,
Bristol BS34 8QZ
Phone: + 44 (0)117 3128240
Fax: + 44 (0)117 3128003
E-mail: pbe@hpl.hp.com

How is it built?

The structure of the Internet is best understood by considering your access point. Most readers will have come into contact with either Internet connections through a local area network, or through a direct dial-up account with an Internet service provider (ISP). In the case of the local area network link, someone on your site has created a local connection to the Internet through which all of your communications are routed. For the purposes of this discussion, only the dial-up account with an ISP will be considered.

How an Internet service is provided

Your ISP has a large-bandwidth link to the Internet. This is normally provided by a telecommunications company, (eg BT or TeleWest), or a reseller of telecommunications bandwidth (eg Pipex). We can consider this to be the 'outward' link to the Internet. The ISP needs to take many incoming phone lines and allow them to share access to this outward link. The ISP does this through a modem rack for the incoming calls, a cycling allocation of Internet addresses and shared access to the Net through a single computer attached to the outward link (Figure 1).

Figure 1: Dialling in to the Internet through an ISP from home

The ISP Internet link then connects to one of several 'network access points' (NAPs) around the country. These are large switches that are connected to the very-high-bandwidth telecommunications backbone in a single country. The network access points are also where the international interconnectivity occurs. The NAPs connect across satellite links to other NAPs in other countries. The operators of the NAPs, often telecommunications companies, need to cooperate with each other in order to provide uniform Internet access. This is not always the case, particularly in countries where deregulation has led to stern competition, so the best way to reach another NAP in the same, or a neighbouring, country might be to link to a US NAP and back again.

The computers connected to the Internet all require an Internet address, whether just temporarily, as is the case for the ISP dial-up, or permanently, as is the case for World Wide Web services for companies. These IP addresses (eg 192.4.55.104) are given some meaning via the domain naming system (DNS), where names associated with computers can be registered as belonging to an IP address. It is even possible to allocate several names to a single IP address, with the computer at the other end working out which account to route a single DNS access to.

So there we have it, connections to and from the Internet allow computers to do their stuff with access to all the other computers doing their stuff at the same time. The standard protocols mean that Apple, HP, Sun, Compaq, Dell and all the other computer manufacturers have machines that link to the Internet for both consuming and providing Internet services.

The growth of the Internet

For the first 20 years or so of its life, the Internet was the home of the academic. Its major components were written as summer student projects, and its traffic was mostly academic discussions. Computer science dominated because only computer scientists could put up with some of the hostile command-line-driven interactions that the Net demanded.

**The Internet
continues to grow**

Figure 2: Internet growth 1981–98 (Internet Domain Survey)¹

Figure 2 shows the Internet Domain Survey figures since 1981 (for ‘host’ read ‘computer’). This isn’t a particularly accurate translation, as it’s an underestimate of the number of computers attached. So what happened in 1990–92 that caused this explosive rise in the number of Internet connections? What has taken the Internet out of the realm of technical mystique and into common usage? One European invention — the World Wide Web.

The original proposal for the World Wide Web

The World Wide Web and accessibility

The original designers of the World Wide Web were aiming at common document and data retrieval by people without computer science degrees. They weren't really aiming at 'members of the general public' since the team worked at CERN, the nuclear physics research installation near Geneva, and the data they were sharing concerned sub-atomic particles. Still, 'no more command-line interfaces' seemed a good idea. Here's what the original authors of the Web (Tim Berners-Lee and Robert Cailliau) described as the objectives for the six-month project:

- to provide a common (simple) protocol for requesting human readable information stored at a remote system, using networks;
- to provide a protocol within which information can automatically be exchanged in a format common to the supplier and the consumer;
- to provide some method of reading at least text (if not graphics) using a large proportion of the computer screens in use at CERN;
- to provide and maintain at least one collection of documents, into which users may (but are not bound to) put their documents. This collection will include much existing data (this is partly to give us first-hand experience of use of the system, and partly because members of the project will already have documentation for which they are responsible);
- to provide a keyword search option, in addition to navigation by following references, using any new or existing indexes (such as the CERNVM FIND indexes). The result of a keyword search is simply a hypertext document consisting of a list of references to nodes which match the keywords;
- to allow private individually managed collections of documents to be linked to those in other collections;
- to use public domain software wherever possible, or interface to proprietary systems which already exist;
- to provide the software for the above free of charge to anyone.²

The idea was to provide common data formats to store and read textual information. There would be computers that stored the information (servers) to allow others to pull copies of that information across the network to be viewed by special pieces of software (browsers). Once again, the Internet academic ethos of 'free' appears in the list. This encouraged other adopters of the technology to extend its capabilities from text to graphics and other formats.

The Web not only established its own new protocols to run across the Internet, but also encompassed most of the important protocols already in existence. This is available globally within a hypertext and multimedia framework. Ever wondered why most WWW addresses start with 'http:'? Hypertext transfer protocol (http) is the communications protocol that the WWW is based on. You can also use 'ftp:' or 'gopher:' to tell the browser that you want to connect to that machine using a different protocol, but it only works if the server at the other end supports that protocol.

The WWW integrates Internet services

Security and the Web

The commercialisation of the Internet

The important step in opening up the Internet for commercial use was taken with the introduction of a secure encryption link between a server and a browser. A certification authority issues a server with a certificate of authenticity. When a browser contacts the server, it uses the encryption key presented in the certificate to establish another encryption key that will be used to scramble the packets in that interaction alone. The interactions will, therefore, be secret to the browser and the server, and be uniquely scrambled every time a new connection is made. This encryption scheme (secure sockets layer, SSL) was first proposed by Netscape and is the basis of secure WWW interactions in most browser-to-server links.

The addition of a guaranteed way to encrypt the channel between browsers and servers made the transmission of sensitive data somewhat safer. Lines could not be tapped and card numbers stolen. SSL does not add other types of security to the process (eg it does not ensure that your credit card details are never stored to hard disk in the clear), but it does add some extra protection against eavesdropping.

There is growing evidence that some classes of commerce do very nicely using this medium, despite the worries about its fundamental insecurity. Dell Computers (<http://www.dell.com>) have geared most of their business around their Website. For them, selling direct to customers allows the company to bypass a complex and costly dealer chain. Most efforts by supermarkets (eg <http://www.j-sainsbury.co.uk>) to allow home shopping are limited to larger stores, where a team of shelf pickers is dispatched to compile your order to be delivered or picked up. This still requires you to be around to handle the order and deal with the perishable or frozen items.

An intriguing variant of this idea is that from Streamline (<http://www.streamline.com>). They provide a 'Streamline Box' that is a hybrid fridge, freezer, wardrobe and bookshelf. This is installed in your garage and receives van deliveries of shopping items you order over the WWW. For \$40 a month, plus the cost of your shopping, they 'give you your weekend back'. Streamline isn't a supermarket, but they have similar guarantees of quality as supermarkets. They integrate across various types of purchase and deliver direct to your home. In doing so they offer a challenge to the postal service, video rental stores, dry cleaners, supermarkets and many other service industries who do not offer similar facilities. The Net is opening up new commercial opportunities in consumer practice that were not there before.

New communications technologies won't (necessarily) make it easier to personalise

It is often said that new technologies like the WWW and interactive television will open up a new world of direct customer interaction and allow 'integrated marketing' across channels. Communicating devices in people's work and home lives will give us unprecedented access to their behaviours and actions. Such claims should be treated with caution.

The Web opens up commercial opportunities

Personalisation on the Web

Personalisation and one-to-one marketing in any technology area is very hard and very specialised to that medium. Take a look at your business card right now. Has the advent of fax, e-mail, phone, mobile phone, telex, postal addresses, WWW given people more ways to contact you? Sure it has. Has the addition of any one of these channels integrated in any of the others? No it hasn't. Every new communication mechanism has meant another piece of technology to carry around, have on your desk or load on your computer. How many telephone numbers have you got? How many e-mail addresses might you have? As e-mail and interactive television take off, you'll be able to add a couple of more lines to that card.

Statements like the following will also be seen quite frequently:

'Web servers can observe the behaviour of millions of customers and immediately produce customised content, bundled with customised ads.'³

While this is undoubtedly true, it belies the technical difficulty and large cost involved. Take the Amazon site (<http://www.amazon.com>). When you express an interest in a particular book, the database also gives you a list of books that others who already bought that book included in their purchase ('People who bought that, also bought ...'). This is undoubtedly the kind of activity that the quote is referring to, but it does not do justice to the specialised software and skills needed to create this. Nor does it acknowledge that the service is specialised to the Web. If Amazon operated a call centre (which they don't), would a customer reasonably expect the same service over the phone? How about via post? Over interactive television? Via a fax-back service? Could they provide this? With some more effort, yes. Can we describe Amazon's one-to-one marketing as integrated? Possibly not.

The future for the Internet

Future challenges for the Internet

The 'free and open' tradition of the Internet is surviving through the uptake of commercial companies, mainly through standardisation efforts. The Internet Engineering Task Force is established to develop the basic protocols and structure of the Internet. The World Wide Web Consortium is there to see that WWW services and functions remain openly available as standard implementations. There are large technical challenges ahead for both of these bodies — how will the IP addressing system be expanded to cover all of the computers and devices that will be attached to the Internet in ten years' time? How will the domain naming system adapt to the increase in named computers? How will the merging of JavaScript, HTML, XML and Style Sheets result in a manageable tool-set for Web designers?

The installation of the Internet as part of so many people's everyday lives will undoubtedly provide sufficient justification for the resolution of these issues. There will be a growing commercial imperative to make the technology cope with this growth. In time, the Internet will become just another medium through which to reach

**Children do not see
the Internet as
technology**

consumers and business partners. We still tend to see this new medium as a technology, with technology drivers and the realm of IT departments and academic research. The early adopters in this market are not necessarily the ones that we would see as competitors to ourselves. (A Sainsbury's reader of this paper would probably not see Streamline as much of a threat.)

There is a real danger in treating the Internet as 'a new-fangled gizmo with no relevance to my business'. The medium is providing new players into the market, with new value propositions to consumers and businesses alike. The markets they enter may still be small, the consumers they acquire are early adopters of technology. In time, however, this consumer base will grow. If TV and telephones do not appear to be technological marvels to you, the Internet will not appear so to the next generation. They grew up with it in their schools, and possibly at home. Don Tapscott quotes a startling set of conclusions in *Growing Up Digital*.⁴ An American Demographics survey in 1997 concluded that:

- children are the prime motivators for PC purchases in the home;
- teenagers are the primary users of PCs in the home;
- children know more about the computers in their homes than adults do;
- children know more about computers than their teachers do.⁵

This new technology is in the hands of a generation who have mastered it, and we don't even know what they're up to with it.

With time, the 'small market' becomes the only market. It happens with most disruptive technology changes. Clayton Christensen in *Innovation Dilemma* covers this in far more detail than can possibly be attempted here, but he states:

'In a study of many disruptive technology changes, those companies dominant before the change went out of business within five years after it.'⁶

Invest or lose out

There are examples of this in disk drives, medicine, financial services and even in the manufacture of ice. Those who do not invest lose. The Internet is here to stay and is to be experimented with. There is a chance that one could get it completely (and expensively) wrong, but it is worth bearing this quote in mind:

'If we don't make utter fools of ourselves from time to time, we grow smug — that is, we do not grow at all.'⁷

Glossary of Terms

ftp — File Transfer Protocol — A basic Internet service for transmitting and receiving files over the Internet.

html — Hypertext Markup Language — The standard for defining basic WWW page content and basic formatting (eg tables). New features of the standard need to be added via a standards body, the World Wide Web consortium.

http — Hypertext Transfer Protocol — The format and protocol for exchanging WWW data over the Internet.

IP address — A unique number allocated to a computer when attached to the Internet. May be allocated temporarily for a dial-up session, or permanently for continuously connected computers.

JavaScript — A component of the WWW that allows small programs to be downloaded to other computers and run in Web browsers. Often used to add interactive features (eg calculators) to basic HTML.

Style sheets — A mechanism for controlling the display and layout of HTML pages. Adds control over fonts, page layout etc.

TCP/IP — Transmission Control Protocol/Internet Protocol — The format and protocol for exchanging packets of data over the Internet.

XML — Extensible Mark-up Language — A new mark-up proposal that allows page formatting similar to HTML, but allows definitions of new features by users. New features do not need to be added via a standards body.

References

1. Internet Domain Survey, (Network Wizards), <http://www.nw.com/>.
2. Berners-Lee, T. and Cailliau, R., 'WorldWideWeb: Proposal for a Hypertext Project', original CERN document available at <http://www.wz.org/proposal>.
3. Shapiro, C. and Varian, H. (1999) *Information Rules: A Strategic Guide to the Network Economy*, Harvard Business School Press, Boston.
4. Tapscott, D. (1997) *Growing Up Digital. The Rise of the Net Generation*, McGraw-Hill, New York.
5. American Demographics survey quoted in *ibid*
6. Christensen, Clayton (1997) *The Innovator's Dilemma: When New Technologies Cause Great Firms to Fail*, McGraw-Hill, New York.
7. David Glass, CEO of Wal-Mart, quoted in Peters, T. (1994) *The Tom Peters Seminar: Crazy Times Call for Crazy Organisations*, Macmillan, London, p. 278.