The Digital Imprimatur

How big brother and big media can put the Internet genie back in the bottle.

by John Walker

> **imprimatur 1.** The formula (= 'let it be printed'), signed by an official licenser of the press, authorizing the printing of a book; hence as *sb.* an official license to print.
>
> *The Oxford English Dictionary (2nd. ed.)*

Introduction

Over the last two years I have become deeply and increasingly pessimistic about the future of liberty and freedom of speech, particularly in regard to the Internet. This is a complete reversal of the almost unbounded optimism I felt during the 1994–1999 period when public access to the Internet burgeoned and innovative new forms of communication appeared in rapid succession. In that epoch I was firmly convinced that universal access to the Internet would provide a countervailing force against the centralisation and concentration in government and the mass media which act to constrain freedom of expression and unrestricted access to information. Further, the Internet, properly used, could actually *roll back* government and corporate encroachment on individual freedom by allowing information to flow past the barriers erected by totalitarian or authoritarian governments and around the gatekeepers of the mainstream media.

So convinced was I of the potential of the Internet as a means of global unregulated person-to-person communication that I spent the better part of three years developing *Speak Freely* for Unix and Windows, a free (public domain) Internet telephone with military-grade encryption. Why did I do it? Because I believed that a world in which anybody with Internet access could talk to anybody else so equipped in total privacy and at a fraction of the cost of a telephone call would be a better place to live than a world without such communication.

Computers and the Internet, like all technologies, are a double-edged sword: whether they improve or degrade the human condition depends on who controls them and how they're used. A large majority of computer-related science fiction from the 1950's through the dawn of the personal computer in the 1970's focused on the potential for centralised computer-administered societies to manifest forms of tyranny worse than any in human history, and the risk that computers and centralised databases, adopted with the best of intentions, might inadvertently lead to the emergence of just such a dystopia.

The advent of the personal computer turned these dark scenarios inside-out. With the relentless progression of Moore's Law doubling the power of computers at constant cost every two years or so, in a matter of a few years the vast majority of the computer power on Earth was in the hands of individuals. Indeed, the large organisations which previously had a near monopoly on computers often found themselves using antiquated equipment inferior in performance to systems used by teenagers to play games. In less than five years, computers became as decentralised as television sets.

But there's a big difference between a computer and a television set—the television can receive only what broadcasters choose to air, but the computer can be used to *create* content—programs, documents, images—media of any kind, which can be exchanged (once issues of file compatibility are sorted out, perhaps sometime in the next fifty centuries) with any other computer user, anywhere.

Personal computers, originally isolated, almost immediately began to self-organise into means of communication as well as computation—indeed it is the former, rather

than the latter, which is their principal destiny. Online services such as CompuServe and GEnie provided archives of files, access to data, and discussion fora where personal computer users with a subscription and modem could meet, communicate, and exchange files. Computer bulletin board systems, FidoNet, and UUCP/USENET store and forward mail and news systems decentralised communication among personal computer users, culminating in the explosive growth of individual Internet access in the latter part of the 1990's.

Finally the dream had become reality. Individuals, all over the globe, were empowered to create and exchange information of all kinds, spontaneously form virtual communities, and do so in a totally decentralised manner, free of any kind of restrictions or regulations (other than already-defined criminal activity, which is governed by the same laws whether committed with or without the aid of a computer). Indeed, the very design of the Internet seemed technologically proof against attempts to put the genie back in the bottle. "The Internet treats censorship like damage and routes around it."[1]. Certainly, authoritarian societies fearful of losing control over information reaching their populations could restrict or attempt to filter Internet access, but in doing so they would render themselves less competitive against open societies with unrestricted access to all the world's knowledge. In any case, the Internet, like banned books, videos, and satellite dishes, has a way of seeping into even the most repressive societies, at least at the top.

Without any doubt this explosive technological and social phenomenon discomfited many institutions who quite correctly saw it as reducing their existing control over the flow of information and the means of interaction among people. Suddenly freedom of the press wasn't just something which applied to those who owned one, but was now near-universal: media and messages which previously could be diffused only to a limited audience at great difficulty and expense could now be made available around the world at almost no cost, bypassing not only the mass media but also crossing borders without customs, censorship, or regulation.

To be sure, there were attempts by "the people in charge" to recover some of the authority they had so suddenly lost: attempts to restrict the distribution and/or use of encryption, key escrow and the Clipper chip fiasco, content regulation such as the Computer Decency Act, and the successful legal assault on Napster, but most of these initiatives either failed or proved ineffective because the Internet "routed around them"—found other means of accomplishing the same thing. Finally, the emergence of viable international OpenSource alternatives to commercial software seemed to guarantee that control over computers and Internet was beyond the reach of any government or software vendor—any attempt to mandate restrictions in commercial software would only make OpenSource alternatives more compelling and accelerate their general adoption.

This is how I saw things at the euphoric peak of my recent optimism. Like the transition between expansion and contraction in a universe with Ω greater than 1, evidence that the Big Bang was turning the corner toward a Big Crunch was slow to develop, but increasingly compelling as events played out. Earlier I believed there was no way to put the Internet genie back into the bottle. In this document I will provide a road map of precisely how I believe that could be done, potentially setting the stage for an authoritarian political and intellectual dark age global in scope and self-perpetuating, a disempowerment of the individual which extinguishes the very innovation and diversity of thought which have brought down so many tyrannies in the past.

One note as to the style of this document: as in my earlier *Unicard* paper, I will present many of the arguments using the same catch phrases, facile reasoning, and short-circuits to considered judgment which proponents of these schemes will undoubtedly use to peddle them to policy makers and the public. I use this language solely to demon-

[1]This observation is variously attributed to John Gilmore and John Nagle; I don't want to get into that debate here

strate how compelling the arguments can be made for each individual piece of the puzzle as it is put in place, without ever revealing the ultimate picture. As with *Unicard*, I will doubtless be attacked by prognathous pithecanthropoid knuckle-typers who snatch sentences out of context. So be it.

The Emerging Consumer Internet

The original design of the ARPANET, inherited by the Internet, was inherently peer to peer. I do not use the phrase "peer to peer" here as a euphemism for "file sharing" or other related activities, but in its original architectural sense, that all hosts on the network were logically equals. Certainly, Internet connections differed in bandwidth, latency, and reliability, but apart from those physical properties any machine connected to the Internet could act as a client, server, or neither—simply a *peer* of those with which it communicated. Any Internet host could provide any service to any other and access any service provided by them. New kinds of services could be invented as required, subject only to compatibility with the higher level transport protocols (such as TCP and UDP).

This architecture made the Internet something unprecedented in the human experience, the first *many-to-many mass medium*. Let me elaborate a bit on that. Technological innovations in communication dating back to the printing press tended to fall into two categories. The first, exemplified by publishing (newspapers, magazines, and books) and broadcasting (radio and television) was a one-to-many mass medium: the number of senders (publishers, radio and television stations) was minuscule compared to their audience, and the capital costs required to launch a new publication or broadcast station posed a formidable barrier to new entries. The second category, including postal mail, telegrams, and the telephone, is a one-to-one medium; you could (as the technology of each matured) communicate with almost anybody in the world where such service was available, but your communications were person to person—point to point. No communication medium prior to the Internet had the potential of permitting any individual to publish material to a global audience. (Certainly, if one creates a Web site which attracts a large audience, the bandwidth and/or hosting costs can be substantial, yet are still negligible compared to the capital required to launch a print publication or broadcast outlet with comparable reach.)

This had the effect of dismantling the traditional barriers to entry into the arena of ideas, leveling the playing field to such an extent that an individual could attract an audience for their own work, purely on the basis of merit and word of mouth, as large as those of corporate giants entrenched in earlier media. Beyond direct analogues to broadcasting, the peer to peer architecture of the Internet allowed creation of entirely new kinds of media—discussion boards, scientific preprint repositories, web logs with feedback from readers, collaborative open source software development, audio and video conferences, online auctions, music file sharing, open hypertext systems, and a multitude of other kinds of spontaneous human interaction.

A change this profound, taking place in less than a decade (for despite the ARPANET's dating to the early 1970s, it was only as the Internet attracted a mass audience in the late 1990s that its societal and economic impact became significant), must inevitably prove discomfiting to those invested in or basing their communication strategy on traditional media. One needn't invoke conspiracy theories to observe that many news media, music publishers, and governments feel a certain nostalgia for the good old days before the Internet. Back then, there were *producers* (publishers, broadcasters, wire services) and *consumers* (subscribers, book and record buyers, television and radio audiences), and everybody knew their place. Governments needn't fret over mass unsupervised data flow across their borders, nor insurgent groups assembling, communicating anonymously and

securely, and operating out of sight and beyond the control of traditional organs of state security.

Despite the advent of the Internet, traditional media and government continue to exercise formidable power. Any organisation can be expected to act to preserve and expand its power, not passively acquiesce in its dissipation. Indeed, consolidation among Internet infrastructure companies and increased governmental surveillance of activities on the Internet are creating the potential for the imposition of "points of control" onto the originally decentralised Internet. Such points of control can be used for whatever purposes those who put them in place wish to accomplish. The trend seems clear—over the next five to ten years, we will see an effort to "put the Internet genie back in the bottle": to restore the traditional producer/consumer, government/subject relationships which obtained before the Internet disrupted them.

A set of technologies, each already in existence or being readied for introduction, can, when widely deployed and employed toward that end, reimpose the producer/consumer information dissemination model on the Internet, restoring the central points of control which traditional media and governments see threatened by its advent. Each of the requisite technologies can be justified on its own as solving clamant problems of the present day Internet, and may be expected to be promoted or mandated as so doing. In the next section, we'll look at these precursor technologies.

Technological Precursors

The dark future I dread will be the consequence of the adoption, by marketing or mandate, of a collection of individual technologies, each of which can be advocated as beneficial in its own right but which, taken together, have consequences less apparent to many yet, I believe, quite evident to some now promoting them. Each of the following technologies is either currently in existence or is the object of an active development effort. These items necessarily interact with one another, so it is impossible to entirely avoid forward references in discussing them. If something doesn't seem clear on the first reading, you may benefit from re-reading this section after you've digested the essentials the first time through.

The Firewalled Consumer

> **Note:** this item discusses a phenomenon, already underway, which is effectively segmenting Internet users into two categories: home users who are consumers of Internet services, and privileged sites which publish content and provide services. The technologies discussed in the balance of this document are *entirely independent* of this trend, and can be deployed whether or not it continues. If you aren't interested in such details or take violent issue with the interpretation I place upon them, please skip to the next heading. I raise the issue here because when discussing the main topics of this document with colleagues, a common reaction has been, "Users will never put up with being relegated to restricted access to the Internet." But, in fact, they *already are* being so relegated by the vast majority of broadband connections, and most aren't even aware of what they've lost or why it matters.

When individuals first began to connect to the Internet in large numbers, their connection made them logical peers of all other Internet users, regardless of nature and size. While a large commercial site might have a persistent, high bandwidth connection and a far more powerful server than the home user, there was nothing, in principle, such a site could do that an individual user could not—any Internet user could connect to any

other and interchange any form of data on any port in any protocol which conformed to the underlying Internet transport protocols. The user with a slow dial-up connection might have to be more patient, and probably couldn't send and receive video in real-time, but there was no distinction in the ways they could use the Internet.

Over time, this equality among Internet users has eroded, in large part due to technical workarounds to cope with the limited 32-bit address space of the present day Internet. I describe this process in detail in Appendix 1, exploring how these expedients have contributed to the anonymity and lack of accountability of the Internet today. With the advent of broadband DSL and cable television Internet connections, a segmentation of the Internet community is coming into being. The typical home user with broadband access has one or more computers connected to a router (perhaps built into the DSL or cable modem) which performs Network Address Translation, or NAT. This allows multiple computers to share a single fast Internet connection. Most NAT boxes, as delivered, also act as a rudimentary Internet firewall, in that packets from the Internet can only enter the local network and reach computers connected to the broadband connection *in reply* to connections initiated from the inside. For example, when a local user connects to a Web site, the NAT router allocates a channel (port) for traffic from the user's machine to the Web site, along with a corresponding inbound channel for data returned from the Web site. Should an external site attempt to send packets to a machine on the local network which has not opened a connection to it, they will simply be discarded, as no inbound channel will have been opened to route them to the destination. Worms and viruses which attempt to propagate by contacting Internet hosts and exploiting vulnerabilities in software installed on them will never get past the NAT box. (Of course, machines behind a NAT box remain vulnerable to worms which propagate via E-mail and Web pages, or any other content a user can be induced to open.)

The typical home user never notices NAT; it just works. But that user is no longer a peer of all other Internet users as the original architecture of the network intended. In particular, the home user behind a NAT box has been relegated to the role of a consumer of Internet services. Such a user cannot create a Web site on their broadband connection, since the NAT box will not permit inbound connections from external sites. Nor can the user set up true peer to peer connections with other users behind NAT boxes, as there's an insuperable chicken and egg problem creating a bidirectional connection between them.

Sites with persistent, unrestricted Internet connections now constitute a privileged class, able to use the Internet in ways a consumer site cannot. They can set up servers, create new kinds of Internet services, establish peer to peer connections with other sites—employ the Internet in all of the ways it was originally intended to be used. We might term these sites "publishers" or "broadcasters", with the NATted/firewalled home users their consumers or audience.

Technically astute readers will observe, of course, that NAT need not prevent inbound connections; a savvy user with a configurable router can map inbound ports to computers on the local network and circumvent the usual restrictions. Yet I believe that as time passes, this capability will become increasingly rare. It is in the interest of broadband providers to prevent home users from setting up servers which might consume substantial upstream bandwidth. By enforcing an "outbound only" restriction on home users, they are blocked from setting up servers, and must use hosting services if, for example, they wish to create a personal home page. (With consolidation among Internet companies, the access supplier may also own a hosting service, creating a direct economic incentive to encourage customers to use it.)

In addition, it is probable that basic broadband service will be restricted to the set of Internet services used by consumers: Web, FTP, E-mail, instant messages, streaming video, etc., just as firewalls are configured today to limit access to a list of explicitly permitted services. Users will, certainly, be able to obtain "premium" service at addi-

tional cost which will eliminate these restrictions, just as many broadband companies will provide a fixed IP address as an extra cost option. But the Internet access market has historically been strongly price sensitive, so it is reasonable to expect that over the next few years the majority of users connected to the Internet will have consumer-grade access, which will limit their use to those services deemed appropriate for their market segment.

In any case, the key lesson of the mass introduction of NAT is that it demonstrates, in a real world test, that the vast majority of Internet users do not notice and do not care that their access to the full range of Internet services and ability to act as a peer of any other Internet site has been restricted. Those who assert that the introduction of the following technologies will result in a mass revolt among Internet users bear the burden of proof to show why those technologies, no more intrusive on the typical user's Internet experience than NATted broadband, will incite them to oppose their deployment.

Certificates

A *certificate* is a digital identification of a physical or abstract object: a person, business, computer, program, or document. A certificate is simply a sequence of bits which uniquely identifies the object it pertains to. In most cases it is guaranteed that there is a one-to-one mapping between certificates and objects. To make this less abstract, consider a non-computer analogue: passports. A passport (or, more precisely, a passport number, as individuals may, in certain circumstances, obtain multiple physical passports bearing the same number), uniquely identifies a person as a citizen of the issuing country. No two people are given the same passport number, and one person's attempting to obtain two different passport numbers is considered a crime involving a fraudulent declaration. A digital certificate is much like a passport. It is issued by a *certificate authority*, which vouches for its authenticity. (In the case of a passport, the certificate authority is the issuing government.) The certificate authority trades on its reputation for probity—to obtain high-grade personal certificates from recognised authorities, documentation equal to or better than that required to obtain a passport is necessary. As with passports, certificates issued by obscure or disreputable authorities will engender less trust than those from the big names.

Certificates are in wide use today. Every time you make a secure purchase on the Web, your browser retrieves a certificate from the e-commerce site to verify that you're indeed talking to whom you think you are and to establish secure encrypted communications. Most browser E-mail clients allow you to use personal certificates to sign and encrypt mail to correspondents with certificates, but few people avail themselves of this capability at present, opting to send their E-mail in the clear where anybody can intercept it and you-know-who routinely does.

When you obtain a personal certificate, the certificate authority that signs it asserts that you have presented them adequate evidence you are who you claim to be (usually on the basis of an application validated by a notary, attorney, or bank or brokerage officer), and reserves the right to revoke your certificate should they discover it to have been obtained fraudulently. Certificate authorities provide an online service to validate certificates they issue, supplying whatever information you've chosen to disclose regarding your identity. Having obtained a certificate, you're obliged to guard it as you would your passport, credit cards, and other personal documents. If another person steals your certificate, they will be able to read your private E-mail, forge mail in your name, and commit all the kinds of fraud present-day "identity theft" encompasses. While stolen certificates can be revoked and replacements issued, the experience is as painful as losing your wallet and worth the same effort to prevent.

A certificate comes in two parts: private and public. The private part is the credential a user employs to access the Internet, sign documents, authorise payments, and decrypt

private files stored on their computer and secure messages received from others. It is the private part of the certificate a user must carefully guard; it may be protected by a pass phrase, be kept on a removable medium like a smart card, or require biometric identification (for example, fingerprint recognition) to access. The public part of the certificate is the user's visible identification to others; many users will list their public certificate in a directory, just as they list their telephone number. Knowing a user's public certificate allows one to encrypt messages (with that person's public key, a component of the public certificate) which can only be decoded with the secret key included in the private certificate. When I speak of "sending the user's certificate along with a request on the Internet" or tagging something with a certificate, I refer to the *public* certificate which identifies the user. The private certificate is never disclosed to anybody other than its owner.

The scope of objects certificates can identify is unlimited. Here are some examples, as they presently exist and may be expected to evolve in the near future.

People. Any individual may obtain a personal certificate from a recognised issuing authority, providing the required credentials to establish their identity. The certificate authority will first verify the identity of the requester, make sure that person has no other non-revoked certificate issued by it or any other authority, then issue the certificate. The issuing authority will, upon demand, verify the certificate, providing the minimum legally required identification of its holder and any additional information the holder instructs the authority to disclose.

Minors may obtain certificates subject to parental consent, as is presently required to obtain a driver's license or enlist in the military. A parent or guardian may require a minor's certificate to disclose the minor's age, which can be used to block access or filter content inappropriate for a person of that age. Further, if requested, the minor's certificate may be linked to that of the parent or guardian, who may then read data encrypted with the minor's certificate.

Certificate authorities undertake to protect the private encryption keys for all certificates they issue, along with any personal information the holder has not explicitly instructed them to disclose (beyond minimum legal requirements). Certificate holders may update their personal information as needed (providing suitable documentation when, for example, legally changing a name), and may suspend or revoke their certificates if suspected or confirmed to be compromised. Should the security of a certificate authority be breached, all certificate holders who may be affected must be notified. Certificate authorities will comply with requests from law enforcement, subject to due process, for recovery of private encryption keys or identity information, including those of revoked certificates.

Companies/Organisations. As noted above, most Web users already implicitly rely on certificates to confirm the identity of companies with which they do business on the Web. There's a lot going on behind that little lock icon in your browser. If a site asserts its identity based on a certificate issued by "Bob's Discount Passports and Pawn Shop", an alert pops up to warn the user they may be about to do something really, really dumb. Similarly, before signing an organ donor contract online, be sure to check out the *bona fides* of Instant Ca$h for Kidneys and those of their certificate authority.

As the Internet is made increasingly secure, the requirements for obtaining a certificate for an organisation will be brought into line with those for certificates identifying individuals. Businesses, whether proprietorship, partnership, or corporation; nonprofit organisations; educational institutions; governmental bodies and other kinds of legal entities will obtain certificates by furnishing the kind of

credentials currently required to obtain an employer identification number for income tax purposes or a sales tax/VAT number. As with certificates for individuals, verification will ensure no entity has more than one valid certificate.

Unlike individual certificates, those granted to an organisation may be used to create subordinate certificates for components of the organisation. Individual offices, departments, etc. may obtain their own certificates, linked to the parent organisation's, and administered by it. This delegation may occur to any number of levels, according to the administrative policy of the organisation—whether a subordinate certificate may create further sub-certificates is determined when it is granted and may be subsequently changed.

An important kind of subordinate certificate is those created for staff (employees, etc.) of an organisation. They identify an individual as a staff member and are used by that person for work-related purposes. The degree to which someone outside the organisation can obtain personal information about a staffer depends on that organisation's own policy and government privacy requirements. An organisation is responsible for the actions of its staff using their certificates and may be compelled to identify them for law enforcement purposes. Whether a staff member can access the Internet from computers on the organisation's network using their private individual certificate as opposed to the staff certificate, and whether the staff certificate may be used outside the organisation network is up to the issuer and can be easily technologically enforced. The private encryption keys of a staff certificate can be recovered by the issuer (or a designated higher level in the hierarchy of certificates issued by the organisation), permitting supervision of staff activities and recovery of the staffer's work product where necessary.

Computers. The "system signatures" computed from hardware properties by present-day "software activation" procedures are crude forms of certificates. The CPU serial numbers in recent Pentium chips (which can presently be disabled, due to public outcry) are still closer approximations. "Trusted Computing" platforms will contain a unique certificate for each machine, referred to as a "credential".

Computers in the Trusted Computing era will be assigned a certificate by their manufacturers which cannot be changed by the user. (As will be discussed in further detail below, it may be possible to *transfer* the certificate to a new machine if the original system fails.) A computer certificate uniquely identifies a machine. It will be used to unlock software licensed for exclusive use on that machine, and to identify network traffic originating from it.

Programs. A program can be issued a certificate which certifies not only that it has been signed by its publisher, but further verifies its contents have not been modified, using a hash/signature/message digest algorithm such as MD5. This technology is already being used for "signed applets" for browsers such as Microsoft Internet Explorer, where a user can (it is claimed) validate the publisher of the program and confirm that it has not been subsequently modified before executing it on their machine.

In the "Trusted Computing" architecture, every program will bear a certificate attesting to the identity of its publisher and permitting the operating system to confirm that it has not been corrupted. A Trusted Computing operating system will not execute a program which differs from the signature in its certificate and will periodically, when connected to the Internet, re-verify a program's certificate to confirm it has not been revoked. Revoking the certificate of a deployed program effectively "un-publishes" it. Such a program will only continue to run on machines on which it was already installed and which are never subsequently connected to the Internet. While revocation of a program's certificate is an extreme

measure and can be expected to be correspondingly rare, it provides an essential mechanism to protect the Internet infrastructure from rapidly emerging threats. If a critical security flaw is found in widely-deployed software which creates an immediate peril, revoking its certificate can pull the plug on the program, requiring users to immediately install an update which corrects the problem.

Content. "Content" refers to any form of digital data: documents, images, audio, video, databases, etc. Here the issue isn't identity or security, but rather authenticity and ownership rights. The publisher's certificate signing your copy of *Moby-Dick* guarantees that this is Melville's original novel, as opposed to a version "enhanced" by a cheesy-poof addict in which Ahab slays the white whale, builds a starship from the bones, and sets forth to annihilate whales throughout the galaxy.

Programs are, in fact, just a special case of content. Due to the risk malicious programs pose to individual users and the Internet, priority will be given to securing them but, with the advent of Digital Rights Management (see below), similar provisions will apply to all kinds of data stored on computer systems. Eventually, every file will be signed with a certificate identifying its creator and incorporating a signature which permits verifying its integrity. If the contents of a document have been corrupted or its certificate revoked, a Trusted Computing platform will not permit it to be opened, and the Secure Internet will not permit it to be transmitted. A document bound to a given user's certificate, that of an organisation, or a specific computer may not be opened by others and will be stored in encrypted form which cannot be decoded without the requisite certificate.

Trusted Computing

"Trusted Computing," in the current jargon, has little or nothing to do with traditional concepts of software reliability or data security. Instead, it refers to an effort to embed end-to-end validation of the origin and integrity of data into computing hardware and system software. One key component is the identification of each computer by a unique certificate, but the ramifications go far beyond this. In addition to protecting computer users from insecure software (software not signed with a recognised vendor's certificate and verified unmodified by its digital signature), users are also protected against corruption of data on their own computers. Data on a user's own hard drive is encrypted and signed, permitting access rights and data integrity to be verified every time a file is loaded into memory. This will completely eliminate the risk of viruses corrupting installed programs or data files. It permits a software vendor to block the execution of any program deemed harmful, even retroactively (since certificates will be verified online). If a vulnerability is found in a software product installed on millions of users' machines worldwide, it may be instantly disabled before it puts them at risk, forcing them to immediately upgrade to a new, secure version. In many cases this will occur automatically—the user need do nothing, nor even be aware of the upgrade to the system.

On a Trusted Computing system, the ability to back up, mirror, and transfer data will be necessarily limited. Hardware and compliant operating systems will restrict the ability to transfer data from system to system. For example, software bound to a given machine's certificate will refuse to load on a machine with a different certificate. Perforce, this security must extend to the most fundamental and security-critical software of all—the ROM BIOS and operating system kernel. Consequently, a trusted computing platform must validate the signature of an operating system before booting it. Operating systems not certified as implementing all the requirements of Trusted Computing will not be issued certificates, and may not be booted on such systems.

Micropayment

Today, buying stuff on the Internet is a big deal—something which many people remain hesitant to do, being well aware of the risks of having their credit card hijacked and the myriad distasteful sequels thereof. With the advent of certificates and Trusted Computing, these fears will dissipate. With one's personal certificate (bound, perhaps, to one or more computers to which one has exclusive access, and secured by a pass phrase, smart card, or biometric identification) guaranteeing the security of the connection, and certificates on the other end validating the identity of the vendor, much of the tedious process of present-day Internet commerce can give way to a seamless surfing and shopping experience.

A micropayment exchange permits payments to be made between any two certificate holders. A user makes a payment by sending a message to the exchange, signed with the user's private certificate, identifying the recipient by their public certificate and indicating the amount to be paid. Upon verification of the payer's and recipient's certificates and that sufficient funds are available in the payer's account, the specified sum is transferred to the recipient's account and a confirmation sent of arrival of the funds. Micropayment transactions can be performed explicitly by logging on to the exchange's site, but will usually be initiated by direct connection to the exchange's server when the user makes an online purchase.

Micropayment differs from existing online payment services such as PayPal and e-gold in that transaction costs are sufficiently low that extremely small payments can be made without incurring exorbitant processing fees; with micropayment it will be entirely practical for Web sites to charge visitors a ten-thousandth of a Euro to view a page; credit cards or existing online payment services have far too high overhead to permit such minuscule payments. Note that there need be no *upper* limit on payments made through micropayment exchanges, and hence "micropayment" simply implies that tiny payments are *possible*, not that larger payments aren't routinely made as well. The first broadly successful micropayment exchange is likely to be technology driven, but as micropayments become a mass market and begin to encroach on other payment facilities, pioneers in the market are likely to be acquired by major players in the financial services industry.

No more e-commerce paranoia ... when you do business with vendors with certificates you consider trustworthy, you needn't enter any sensitive personal information. Just click "buy", select which of the credit cards or bank accounts linked to your certificate with which you wish to pay (never giving the number), and your purchase will be shipped to the specified address linked to your certificate. Even if your certificate is stolen, a thief can only order stuff to be shipped to you.

Each user can set their own personal default maximum price per page, per item purchased, per session, per day, per week, and per month. I call this their "threshold of paying." No need to subscribe to a magazine's site to read an article—just click on it and, if it costs less than your €0.05 per-item threshold and all of the other totals are within limits, up it pops—your account is debited and the magazine's is credited. If you're a subscriber, your certificate identifies you as one and you pay nothing ... and all of this happens in an instant without your needing to do anything. The magazine gets paid for what you read, so they'll put their entire content online, not just a teaser to induce you to subscribe to the printed edition. And if you like what you read, you'll return and spend more money there.

Want to start your own magazine? Decided your blog is worth €0.001 per day to read? No problem ... tag it with your certificate, set up a "pay to read" link to it, and listen to the millieuros tinkle into your virtual cookie jar.

Certified micropayment exchanges will, of course, be required to comply with "know your customer" and disclosure regulations, adhere to international conventions

against money laundering, terrorism, and drug trafficking, and disclose transactions to the fiscal authorities of the jurisdiction of the buyer and seller for purposes of tax assessment. This will largely put an end to the use of the Internet for financial crimes and eliminate the need for further regulations or constraints on Internet commerce.

Micropayment and Funding Internet Resources

Micropayment provides a new business model to support Internet sites which attract large numbers of visitors but which have so far failed to fund themselves with subscriber or advertiser models. Micropayment permits a site to make access available to whoever chooses to visit the site on a per-page basis (or, as discussed below, even for excerpts from pages). There is no

need for a user to open an account or establish a commercial relationship with the site. As long as the per page fee is less than the individual's threshold of paying, the per-page charge is debited automatically from the user's account and credited to the site's.

There's no question that if many present-day sites started to charge, say, €0.001 per page, their traffic would collapse. But what about the sites you read every day? Is it worth a tenth of a centime per page? Have you compared what you'd pay for pages with what you're paying now for access to the Internet?

Micropayment, Excerpts, and "Deep Linking"

The emergence of Weblogs ("blogs") and other forms of independent Internet journalism has raised a variety of issues regarding free use of copyright protected material. To what extent may a blog excerpt a document published on the Web (with or without a link to the original source)? Is it permissible for a Web document on one site to link directly to a document deep within another site's archives, potentially bypassing advertisements on the site's main page which fund its operation?

Micropayment provides solutions for many of these problems. As envisioned by Ted Nelson almost 40 years ago in his original exposition of Xanadu, the problem with copyright isn't the concept but rather its *granularity*. (I'd add, in the present day, the absurd notion that copyright should be eternal, but that's another debate for a different document.) Once micropayment becomes as universal as E-mail, a blog will simply quote content from a Web site using an "excerpt URL" (I'll leave the design as an exercise for the reader) or provide a link to the entire document. Readers of the blog will, if the excerpt is below their threshold of paying (and the total of all excerpts in the blog is also below the threshold), see it automatically. Otherwise, they'll have to click on an icon to fetch it, approving the payment, before it is displayed. Similarly, when following a link to a document licensed under one of the Digital Rights Management (see below) terms of use, you'll automatically pay the fee and see the document unless it exceeds your threshold, in which case you'll have to confirm before retrieving it.

Micropayment and Ubiquitous Wireless Internet Access

Micropayment will greatly facilitate the deployment of wireless Internet access (Wi-Fi and its descendants). Wireless access today has a unsettled business model; some coffee shops and bookstores provide free access to their clients (and, constrained by Maxwell's equations, those in the parking lot outside) as an added value, while hotels, airline lounges, and soon long distance flights en-route provide access for a fee. With micropayment, your wireless network interface will simply listen for bids of access and choose based on bandwidth and cost, normally accepting the best offer below the cost threshold you set. If it's higher than your threshold, or there's an extreme tradeoff between cost and performance, you may be asked to choose, but usually you'll just light up your

laptop, wait a few seconds, and you're online. No mess, no fuss, and it's guaranteed to cost less than your "threshold of paying".

Micropayment and Internet Taxation

According to folklore, Michael Faraday, who discovered the principle of electromagnetic induction in the 1830's, was asked by a British politician to what conceivable use electricity might be put. Faraday replied, "Sir, I do not know what it is good for. But of one thing I am quite certain—someday you will tax it." This quotation is, in all likelihood, a myth, but nonetheless there is truth therein applicable to our times. For electricity, a laboratory curiosity in Faraday's time, *was* eventually taxed and, in many unfortunate jurisdictions, made a government monopoly or regulated to such an extent it was indistinguishable from one, inevitably becoming scarce, expensive, and unreliable.

Like electricity, the Internet *will* eventually be taxed. As long as there are governments, this is inescapable. While taxation is never without pain, micropayment can at least eliminate most of the bookkeeping headaches for both merchants and customers, with taxes due for Internet use and commerce collected automatically and remitted electronically to the jurisdiction they are owed to.

Digital Rights Management

> Microsoft also warned today that the era of "open computing," the free exchange of digital information that has defined the personal computer industry, is ending.
>
> *Microsoft Tries to Explain What Its .Net Plans Are About*
> by John Markoff, *The New York Times*, July 24, 2002.

Digital Rights Management (DRM) is the current buzzword for the technological enforcement of intellectual property rights in digital media.

DRM will implement several categories of right to use content, some of which have no direct analogues in traditional publishing.

Pay Per Copy

This is the traditional model of books, recorded music, videos, and shrink-wrapped software. You pay a fee for a copy and usually assent to an implicit license not to copy and redistribute it. However, there is no technological prohibition against your doing so and, in some cases, your purchase entitles you to lend the original document to others without paying additional fees to the publisher.

Pay Per Instance

This is a phrase I've coined to denote the concept of a document sold to a given individual which is either not transferable or, if so, cannot be used to create additional copies. When you purchase a pay per instance document, it's "bound" to your personal certificate and possibly that of the computer on which you intend to view it. If you copy the document you've downloaded (assuming your Trusted Computing platform even permits this) to somebody else's system, they won't be able to read it because they don't have your certificate. Giving them your certificate is equivalent to handing them copies of all of your credit cards and identity documents . . . unlikely. If the document is, in addition, bound to a given computer system, you can read it on that system but, in order to transfer it to another (for example, from your desktop computer at home to your PDA when going on holiday), you'll need to perform a transfer which will render it readable on the PDA but no longer on the desktop. You can always, upon your return, transfer it back in the other direction.

Pay per instance also permits (publisher permitting), transfers similar to lending a printed book to a friend. Suppose you've downloaded a book to your computer, read it, and now wish to send it to your daughter at college. No problem—just re-encode the book with her public certificate and E-mail it to her. Of course, once you've done that, you won't be able to read the book any more on your own system. There may be a small fee associated with passing on the book but, hey, micropayment makes it painless and you'd probably have to pay a lot more to mail a printed book anyway. Publishers can sell library editions, perhaps at a premium, which can be transferred any number of times but, just like a book, the library can't check out a volume to another person until a borrowed copy is returned.

Pay Per Installation

Pay Per Installation is similar to Pay Per Instance, except the content is bound to the certificate of the computer on which it's installed, as opposed to the personal certificate of an individual. Any person who uses that computer is authorised to access content bound to its certificate, but such content cannot be used on a different computer. This category will primarily be used by commercial software installed on a computer. Pre-installed software will, of course, already be bound to the computer's factory-installed certificate. When you purchase software, whether off the shelf or by downloading from the Internet, you will receive a copy which, before it can be used, must be activated online, which will bind it to the certificate of the machine on which it is installed. The purchase of a copy of the software will usually entitle the customer to a single activation; additional licenses for other computers may, of course, be bought as needed.

Just as with Pay Per Instance, the publisher of a Pay Per Installation product may permit you to *transfer* the product to a different computer. If, for example, you replace your old clunker with a TurboWhiz 40 GHz box, you may be able to move your existing programs to it, going through an activation procedure which will render them unusable on the old machine and bound to the new one. Or, on the other hand, the publisher may not permit this; it's up to the specific terms of the license.

Pay Per View

This is how movies worked when I was a kid. If you wanted to see the movie, you went to the box office, plunked down your fifty cents (I was a kid a *long* time ago), and received a ticket which entitled you to see the movie (plus the newsreel, the cartoon, etc.) *once*. When the show was over they turned on the lights and chased everybody out. If you just *had* to see it again . . . another four bits, thank you very much. This is the golden age media barons dream of while sleeping off the diverse intoxicants they've ingested at sybaritic Hollywood parties.

As with Pay Per Instance, the content you download is bound to your personal certificate or that of your computer but, in addition, it's limited to being played a maximum number of times, for instance, once. Now, instead of struggling to find a song on a music sharing service under constant attack by music moguls, you can simply visit your favourite online music store, find the song that's been going through your head for the last few hours, download it for a small fee and listen to it . . . once. If, having listened to it, you'd like to play it over and over or put it on a CD for your own use, pay a little more and buy a Pay Per Instance copy. No more need to buy an album to get one or two hit singles—of course the singles cost more than the filler. And no, you can't give a copy of the CD you made to your friends, since the songs on it are bound to your certificate and machine. You can make as many copies as you like of your "killer tracks" CD and give them to your friends or sell them on the Net, but everybody who receives one will have to pay the license fee for each track in order to obtain the right to play it.

Note that pay per view has applications outside traditional entertainment media; evaluation copies of software can be licensed to a user for a maximum number of trial runs, after which the user must either purchase a license permitting unlimited use, or some number of additional runs. Software vendors offering evaluation copies on this basis are protected, since they record the user's certificate when issuing evaluation copies, and refuse to issue more than one evaluation copy to any user. This application of pay per view to software closes the loopholes which have made shareware a difficult business model.

Circumvention Prevention

Earlier attempts to protect intellectual property in the digital age have sparked an arms race between copyright owners and those who wish to freely copy protected works. There are reasons to believe a comprehensive implementation of Digital Rights Management on a Trusted Computing platform will be a much tougher nut to crack, evolving in time toward effectively complete security (defined as the point at which losses due to copying are negligible compared to the cost to reduce them further), much as has happened with digital satellite television broadcasting. In the United States, the *Digital Millennium Copyright Act*, enacted in 1998, criminalises reverse engineering and circumvention of copyright protection mechanisms, and has been interpreted as applying to even the dissemination of information regarding the design and implementation of copy protection technologies. Given the political consensus which enacted this law, the stakes involved for media companies, and the investment now being made in Digital Rights Management technologies by computer hardware and software vendors, there is every reason to expect the near-term deployment of a highly secure system implementing all the varieties of right to use described above, which will not be widely circumvented.

Trusted Internet Traffic

Once Trusted Computing platforms are in place which protect intellectual property rights, this security can be extended to the Internet itself. The ARPANET, precursor of the Internet, was designed to explore highly fault-tolerant networks for military communications. In such networks, all communications links could be secured and the identity of all nodes on the network was known. In today's global, open access Internet neither of these conditions obtains, and many of the perceived problems of the present-day Internet are their direct consequences.

Tomorrow's Secure Internet will be implemented in Trusted Computing platforms, in conjunction with Internet Service Providers and backbone carriers. Today, any computer on the Internet can connect to any other connected computer, sending any kind of packet defined by Internet protocols. This architecture means that any system on the Internet, once found vulnerable to some kind of attack, can be targeted by hundreds of millions of computers around the world and, once compromised, be enlisted to attack yet other machines.

The Secure Internet will change all of this. Secure Internet clients will reject all connections from machines whose certificates are unknown (this will be by service; a user may decide to receive mail from people whose certificates aren't known to them, but choosing otherwise will block all junk mail—it's up to the user.) On the Secure Internet, every request will be labeled with the user and machine certificates of the requester, and these will be available to the destination site. There will be no need to validate login and password, as the Secure Internet will validate identity, and, if registered, a micropayment account will cover access charges and online purchases. Internet Service Providers will maintain logs of accesses which will be made available to law enforcement authorities pursuant to a court order in cases where the Internet is used in the commission of a

crime.

In addition, The Secure Internet will protect the intellectual property of everybody connected to it. Consumers will be able to download any documents on the terms defined by their publishers, which will be enforced by Digital Rights Management. Publishers will serve documents, each identified by a certificate which identifies its publisher and its terms of use, and includes a signature which permits verification the document has not been corrupted subsequent to publication.

The Secure Internet

The technological precursors discussed above provide the foundation for the Secure Internet. A typical individual Internet user visiting Web sites, performing searches, buying products and services online, sending and receiving E-mail and instant messages, participating in chat rooms, news groups, discussion boards, and online auctions will notice little change from the present-day Internet except, perhaps, fewer of the irritations which currently detract from these activities. But the Secure Internet will be a very different kind of place, due to fundamental changes in the way those connected to it interact. This section discusses each of these changes in detail. The following section will sketch the consequences for various kinds of activity on the Internet once they have all been implemented.

The End of Anonymity

Many of the problems of the present-day Internet, which engender numerous, mostly ill-considered proposals for legal remedies, are due to the fundamental lack of *accountability* on the Internet. The Internet, as presently implemented, affords its users a rather high degree of anonymity which permits them, if so inclined, to engage in various kinds of mischief with relative impunity.

Providing, or rather restoring, accountability to the Internet is the key *technological* foundation for fixing a large majority of its current problems. The present-day anonymity of the Internet wasn't designed in—it is largely an accident of how the Internet evolved in the 1990's; see Appendix 1 for details.

Let us explore how accountability will be restored to the Internet.

User Certificates: *No* ID, *no* IP

The first step in restoring accountability to the Internet will be the introduction of the *Internet User Certificate*. This certificate, without which no packets will be transferred across the Internet, uniquely identifies the person (individual or legal entity) responsible for sending them. The best analogy to this certificate is not a telephone number, but rather the call sign with which radio and television stations, including amateur radio operators, identify their transmissions. The Internet User Certificate is simply the unique identification of the person responsible for sending a packet across the Internet. An Internet User Certificate is the *credential* which identifies its sender.

Compared to contemporary Internet access accounts, access by certificate has *gravitas*. First of all, one may expect that, given the legal ramifications which certificates will have, sanctions against obtaining or using a certificate under false pretenses will be akin to those for obtaining a passport with forged credentials or presenting a forged driver's license to a policeman in a traffic stop. Accessing the Internet with a false certificate is equivalent to driving on public highways with a bogus number plate on your vehicle or crossing a border with a fake passport and will be subject to comparable penalties.

When you connect to the Secure Internet, your certificate will be transmitted to the point of access, which will then validate your certificate. If its issuing authority fails

to confirm its validity, or it has been revoked by its owner due to a compromise, or has been blocked pursuant to a court order, access will be denied. Once your certificate is validated, you'll be granted full Internet access, precisely as at present. Your certificate will be logged along with the connections you make and furnished, on demand, to all sites to which you connect. This will make e-commerce painless and secure. Once you've registered with a merchant, all subsequent communications are secured with your certificate. You needn't memorise a user name and password for each site, nor worry about a merchant's site being compromised threatening your security. As long as you protect your certificate as you would your wallet or credit cards, you're secure and, in the worst case, should your certificate be compromised, you can always revoke it and replace it with another.

Computer Certificates

In addition, the computer you're using to access the Internet will be identified by its own certificate, which will also be provided on demand to sites you access. While the most commonly used credential is your personal certificate, the computer's certificate can be used to validate access to remote software components you've licensed or, for example, to secure remote backups of files from the computer against access from any other computer. Computer certificates will eventually be built-in by the manufacturer, much like the CPU serial number in Pentium III and later processors or, as is common in Unix workstations, in the form of an identity ("`hostid`") chip which can be transferred from one machine to another in case of hardware failure. The machine's certificate will become the primary means of licensing commercial software installed on the computer. Unlike present day *ad hoc* machine signature schemes or serial number checks in Unix workstation software, programs licensed to a machine's certificate will be stored in encrypted form and decrypted with the machine's private key from its certificate when loaded into memory. This decryption will be performed in hardware or by the kernel of the Trusted Computing operating system, which itself will be locked to the machine certificate.

The large installed base of computers without certificates or hardware support for Trusted Computing operating systems will necessitate a protracted period of transition during which computer certificates are implemented in software and consequently less secure. Users could, for example, obtain certificates for their own computers by presenting their personal certificate to the issuing authority. The certificate would be delivered as a file to be installed on the machine to identify it. Users may revoke machine certificates when a computer is scrapped or sold.

Once a machine's certificate is embedded in hardware, computer theft becomes less attractive a criminal enterprise since a stolen machine will report its identity, and the personal certificate of its user, at the moment it connects to the Secure Internet. If a machine is stolen, its owner may revoke its certificate, rendering it incapable of connecting to the Internet. Even with certificates implemented in software, revocation (or, in the case of theft where one hoped to eventually recover the computer, suspension) of the certificate would block all software licensed to that computer at the moment it next connected to the Internet and performed a certificate validation. Personal data on the hard drive of a stolen computer would be inaccessible to a thief because it is encrypted with the personal certificate of the owner.

Everything is Encrypted

With the advent of certificates for individual Internet users and computers, the Internet will go dark to snoopers. Those who use the Internet will finally have grounds for confidence their private data, messages, and online financial transactions are secure.

Secure Internet Commerce

Today, when you connect to an Internet commerce site, your browser receives and validates a certificate from the site which it uses to determine you are, in fact, connected to the site you think you are, not a false storefront put up by a crook intent, say, on collecting credit card numbers. Your browser then negotiates a session key to encrypt the balance of your transaction with the site. Typically then, if you're already a customer, you log in with a user name and password you've chosen for the site, which are protected against interception by the session's encryption key. If you're a new customer, the user name and password you select, and your address, credit card number, etc. are similarly protected against interception.

With the advent of the Secure Internet, both parties to the transaction, you and the merchant you're doing business with, will be uniquely identified by their certificates. When you connect to the merchant's site, an encrypted channel will automatically be established based on your certificate, your computer's certificate, the merchant's certificate, and that of the merchant's computer. Compromise of *all four* certificates would be required to intercept the data you send during the connection. There will be no need for user names or passwords—your certificate will identify you. If you've decided to permit such disclosure, the merchant can even obtain information such as your shipping address, privacy preferences, and the like while validating your certificate. If you prefer to keep such information private, you'll have to enter it as you do now, or authorise its transmission to merchants on a case-by-case basis when first doing business.

But certificate-based encryption will extend well beyond Internet commerce. On the Secure Internet, *everything* will be end-to-end encrypted in this manner. When you establish a connection to any site at all, in any protocol, the four certificates involved (you, your computer's, the site's, and its computer's) will be validated and used to negotiate a key for the connection, which will be used to encrypt all data exchanged: E-mail, instant messages, Internet telephony audio, Web pages, *everything*. What you exchange with a site while connected is entirely between you and the site. Snooping by third parties is impossible. Not only needn't you worry about somebody reading your mail or snatching your credit card number, a snoop won't even be able to know which pages you request from Web sites you visit, since the URLs of the pages you request and the content you receive will be encrypted. (It will remain possible to determine which *sites* you visit by snooping packets and looking up the IP addresses of those you connect to.)

Private File Storage

Certificate-based encryption will protect data on your computer even when you're not connected to the Internet. The file system in a Trusted Computing platform will automatically and transparently encrypt all files which belong to you with your certificate. If multiple people share one computer, each will be able to read only their own files; without the certificate of the other user, files belonging to that person, even if physically readable, will be gibberish. If a computer is stolen or an unauthorised person gains physical access to it, users' files cannot be read unless the criminal has also managed to obtain the certificates of their owners.

Files stored on removable media will be encrypted in the same fashion. Compromise of private data by scanning backup media (remarkably, many security-conscious people fail to ponder this threat) cannot occur since the backed up files are encrypted with the certificates of their owners. When sending a file to another person on a physical volume such as a floppy disc or recordable CD, it can be signed with the sender's certificate and encrypted with the public key of the intended recipient who can thereby verify the identity of the sender. Should the shipment be intercepted by a third party, its contents cannot be read without the intended recipient's certificate.

We Know what You've Read

With every Internet transaction tagged with the personal certificate of the requester and that of the computer where the request originated, operators of Web sites and other Internet services will be able to "know their customers". For the first time, Web sites will be able to compile accurate readership statistics, subject to audit by circulation bureaux, as for print publications. This, in turn, may restore the viability of the advertiser-supported business model for popular Web sites.

Internet traffic can be logged and audited by others, for their own purposes, as well. The ability to potentially recover a list of certificates of those who accessed a site containing prohibited content such as child pornography will deter those who now rely on the anonymity of the Internet to shield them from prosecution. Sites indulging in hate speech and/or material of interest to terrorists will find their regular visitors scrutinised by the authorities concerned with such matters. Societies which wish to control the flow of information across their borders can monitor the activity of their nationals to determine whether they are violating imposed restrictions. Parents will be able to monitor the activities of their minor children using certificates they've obtained for them which are linked to the parent/guardian's certificate.

Intellectual Property Protection

Digital Rights Management, secured technologically by Trusted Computing systems and legally by sanctions against reverse-engineering and contravention, will provide comprehensive protection for intellectual property of all kinds. Items downloaded from the Internet: Web pages, books and magazines, music or video files, and all other forms of content will bear certificates which define the terms under which they are licensed to the user, which will be enforced in hardware and software. Only data created entirely by the user (for example, documents they've written, pictures they've taken) and content in the public domain will be able to be freely copied, modified, transmitted, published, and used in other ways. Naturally, users may apply Digital Rights Management themselves to content they create, specifying the terms under which others may use it. For example, when circulating an E-mail draft of a scientific paper to a group of colleagues for comment, you may wish to "license it" exclusively to their certificates to prevent further dissemination should one of them prove indiscreet.

Intellectual property protection can be applied at a fine-grained level. A Web page may include images and citations from other Web content licensed on various terms; when the page is viewed, each inclusion is retrieved subject to its own license. If an included item requires payment, confirmation will be required before downloading it or, if the fee is below the reader's designated threshold of paying, the fee will be transferred automatically via micropayment.

Document Certificates: The Digital Imprimatur

With all parties on the Internet identified by the certificates they're required to use to gain access and exchanged with all transactions, and hence mutually accountable for their online interactions, and Trusted Computing platforms guarding against fraudulent credentials or misuse of intellectual property, the foundation will be laid to fully apply certificates to *content*: every document transmitted across the Internet.

The first application of document certificates is already in use: signed applets downloaded by Web browsers which are run only if the certificate is verified as belonging to a trusted supplier and contains a signature which matches the content of the downloaded code. (The MD5 checksums or PGP/GPG signatures posted for many OpenSource software distributions can be thought of as a crude kind of document certificate, manually

validated by the user against a checksum or signature published on the Web site whence the package is downloaded.)

Trusted Computing systems will require all software they run to be signed with certificates, will verify the signature of each program before executing it and, when online, will (periodically) re-validate the certificates of installed programs with their suppliers. If a program's certificate has been revoked (for example, if a critical security flaw has been found in it which requires an update to correct), the Trusted Computing platform will refuse to run the program, informing the user of the reason for the certificate's revocation. The computer's operating system will bear its own certificate, which will be validated by the BIOS before the system is booted, protecting against unauthorised changes to the installed system or noncompliant operating systems which do not fully implement the Trusted Computing architecture.

A computer program is nothing more nor less than a sequence of bytes, like any other digital document. Just as executable programs can, and will, be signed with certificates, so can Web pages, word processor documents, images, music files, and all other forms of digital data be signed. Certificates are, in fact, an essential part of Digital Rights Management, and will routinely accompany files employing it, eventually encompassing virtually all files obtained from commercial sources.

Now let's consider what happens when this architecture is extended to the Internet. I believe the technological precursors described above will eventually be deployed in such as way as to turn the *entire Internet* into a Trusted Computing platform. What, precisely, will this mean? Well, just as a Trusted Computing system will load neither programs nor data files without a validated certificate whose signature matches their contents, neither will the Secure Internet transfer *any document*, in *any standard protocol* without such a certificate accompanying it. (And by the time this is rolled out, consumer Internet access will long have been restricted to a short list of protocols on standard ports: HTTP, FTP, SMTP, POP, etc. The peer-to-peer Internet, where any site could connect to any other on any port with any protocol will have passed into history as, indeed, is already beginning to happen due to NAT routers and firewalls which cannot be configured by the user to accept inbound connections.)

On the Secure Internet, E-mail messages will not be delivered unless signed by the originator's certificate; the recipient of such a message will know who sent it. (A draft standard for a certificate-based mail transfer protocol was published in August 2003.) When a request is received by a Web site, it will be signed by the certificate of the requester; sites will finally know who their visitors are. And what about the Web page the site sends back to the user in reply? Well, that's where things get *really* interesting.

Documents returned by Web servers will be required to be signed with a certificate and, as with all other traffic, the certificate will identify the originator of the page and contain a signature which permits the recipient to verify its content has not been corrupted. But that's not all a document certificate may be required to contain. Suppose, in order to be transmitted across the Internet in reply to public HTTP or FTP requests, a document was required also to be signed by an authorised *document registry*, just as certificates are issued by a certificate authority? Imagine you've just finished adding a new page to your personal Web site describing, say, your idyllic vacation in the Nibi-Nibi islands. You've got all the kinks out of the page, and now you're ready to share it with the world. Just one more thing ... before the page can be transmitted beyond your Web server's local network, it must bear a document certificate. So, you pop up the "Sign page" box in your editor, click "Sign", and a few seconds later you have a signed page to install on your server, whence anybody can download it.

What happened in those few seconds after you clicked the "Sign" button? First, the URL of the page was sent, along with your personal certificate, to your designated document registry (if you do business with more than one, you'll be asked to select which). The document registry will then download the source for your page and all

embedded content (images, animations, etc.), and generate a certificate which identifies you as the author of the page, with signatures for the page and each of its embedded content components. This certificate is then returned to you, where it's stored with the Web page on your server. The document registry will probably charge you a negligible sum of money for the certificate, which you'll pay automatically with micropayment. When you update the document, you simply submit the new version and obtain an updated certificate to accompany it.

When users access your document, its certificate is validated and, if good, the document is transmitted to the requester along with the certificate. The requester can check the signature of what they received to confirm it agrees with what was signed. If the user stores a local copy of the document (whether this is possible and, if so, on what basis is up to you, as the author of the page, to determine; Digital Rights Management will enforce whatever policy you select), the certificate will accompany the document and be checked against the document registry whenever it is accessed. If you update your document you might, for example, tag the signature of the old version at the registry to notify anybody with a copy that an update is available; the next time they tried to read their local copy, they'd receive the notification from the document registry and be alerted to the update. You could even, should you decide to digitally eat your words, revoke the document's certificate; anybody with a local copy would then, if online, be notified the document had been "un-published" and their copy rendered inaccessible. This doesn't violate the user's rights in any way—*you're* the author of the document; you own the copyright, and you can control access to it in any way you wish.

You're doubtless way ahead of me already in thinking of *other* things these document registries can do. . . . First of all, they will, collectively, know when any page is published or an existing page is modified, and can provide this information to operators of search engines, as will be discussed in the next section. Since the registries compute a document's signature by examining it and all embedded content, they might, for example, *compare* those signatures with those of existing documents (aggregated from all document registries) and check for matches against documents flagged as copyright protected. A match might alert the copyright holder of a potential violation by your page. The document registry might, in the interest of compiling a comprehensive archive of the Web or, perhaps, encouraged by a government mandate, make an archival copy of all documents for which it granted certificates; imagine how useful such an archive could be in resolving subsequent disputes regarding their content. Why, the document registry could even, in the interest of wholesomeness or, perhaps, inspired by a public law, examine the contents of the document and match it against profiles of prohibited content, flagging it for possible scrutiny by those who occupy themselves with such matters.

Since a document cannot be transmitted across the Internet without a certificate validated by its document registry, nor can a user who has received a copy of such a document access it once its certificate has been revoked (except on a machine which is never again connected to the Internet after receiving the document), should the document be found to infringe the rights of another party or violate the law in some manner, after this is established through due process of law, the document registry may be ordered to revoke the document's certificate, un-publishing it.

Some might even fantasise that document registries could, based on signature comparison and heuristic examination of document contents, even *refuse to grant* a certificate for a suspicious document unless the publisher provided proof it did not violate copyright or laws regarding its content. But that would constitute prior restraint on publication, which is unthinkable in a free society.

This, then, is the *digital imprimatur*; the *right to publish* as, in olden times, was granted by church or state. A document's certificate, its *imprimatur*, identifies the person (individual or legal entity) responsible for its publication, provides a signature which permits verifying its contents have not been corrupted or subsequently modified, and identifies

the document registry which granted the imprimatur and which, on demand, will validate it and confirm that it has not been revoked. Trusted Computing systems and the Secure Internet will perform these functions automatically and transparently; to a user browsing the Web, everything will look and feel precisely as it does today.

Dynamic (Variant Content) Documents

Certificates for variant content documents (for example, stock quotes, weather information, search results, shopping cart contents at an online merchant site, etc.) cannot include a content signature because each page returned to a requester is unique. It would be absurd to demand a new document certificate be issued for each reply page, and doing so would compromise the security of user information it contained. Dynamic documents may be accommodated within the digital imprimatur by registering a template and granting a document certificate for the result page based upon it.

Registries would examine template certificate requests carefully, especially if made by unknown publishers or those suspected of attempting to circumvent the requirement for document certificates. Users of template certificates would be subject to audit by the document registry to verify the template was being used in the manner claimed when its certificate was granted.

Note: I am well aware that dynamic documents are a huge, gaping, ugly hole in the digital imprimatur scheme. I have not expended a great deal of effort thinking about ways to better secure such documents; I'm sure this issue will be explored in detail if and when document certificates are imposed on the Web. Still, even though dynamic pages account for a large percentage of Web traffic, they are a minuscule fraction of the pages on the Web. The large organisations responsible for variant pages which get large numbers of hits cannot afford to abuse the privilege of template certificates.

Publisher Self-Registries

Commercial publishing houses, news media, and other organisations which publish large volumes of information or frequently-changing content (for example, a newspaper's site) may be delegated the authority to act as their own publication registry in the interest of efficiency and quick reaction. This is analogous to commercial broadcast stations which keep their own program logs. As with a program log, the publisher's document registry is subject to audit and must be publicly accessible to verify document certificates and provide notification of new publications. Evidence of abuse of self-registry will result in withdrawal of the privilege.

Truth, and Consequences

> It is a well-known fact that no other section of the population avail themselves more readily and speedily of the latest triumphs of science than the criminal class.
>
> Inspector John Bonfield,
> Chicago Police Department, 1888

The accountability and security the technologies described in the previous section will provide once fully deployed will put an end to a wide variety of poster child problems of the present day Internet. Here's a brief survey of some of the most obvious,

each pitched as I expect it to be toward the constituencies concerned with each problem.

Copyright Violation

Digital Rights Management and Trusted Computing resolve most of the current problems with copyright violation on individual computers, and the Secure Internet will extend these protections to the entire network through document certificates. Copyright holders can monitor newly published documents for violations and, if detected, begin a procedure which will result in the offending document's certificate being revoked, un-publishing it on any machine which has stored a copy and is subsequently connected to the Internet. The added security, plus the ability to make copyright protected documents available under a variety of license terms including pay per view with micropayment, will encourage owners of documents to make them available on the Internet where before they were hesitant due to fear of piracy.

Identity Theft and Fraud

Remember the story about the miscreant who hung a sign on an automatic teller machine that said "Out of order—please use temporary ATM" and set up his own bogus ATM next to it which simply read credit card stripes, recorded PINs, and flashed "Temporarily out of order" on its display? He collected the machine at the end of the day and did the obvious thing with the information it had obtained. That was a crime. Or how about the waiters in restaurants who make an extra imprint from your credit card and write down the little code on the back and go wild spending your money. That's a crime too, and the only computer used to commit it is made of meat. Setting up a bogus Web site or sending E-mail under false pretenses to obtain precisely the same information is likewise a crime.

The end to end encryption of all transactions on the Secure Internet will render identity theft schemes which rely on intercepting messages nonviable. Trusted Computing platforms will protect against worms and viruses which install "spyware" on users' computers to collect personal information, including credit card numbers, PINs, and passwords. Access by user certificate will eliminate the need for users to keep track of a long list of login names and passwords, and the resulting temptation to store them insecurely or to use the same name and password on a number of sites, running the risk that if one is compromised, accounts on other sites will be as well.

Micropayment will eliminate the risk of identity theft by rogue merchants who collect credit card numbers from online purchases and then use them fraudulently, since the merchant will be paid through the micropayment exchange for the specific transaction approved by the user, and will never see the details of the account (credit card, bank transfer, etc.) with which the user paid. Only the user can authorise a payment to a merchant; without access to the user's certificate, receipt of a payment does not provide the ability to make subsequent fraudulent charges, as possession of an individual's credit card details does today.

Compromise of a user's certificate remains a very serious matter, equivalent to having one's wallet or passport stolen. Certificates, like credit cards, can be quickly blocked, but a user who loses control of a certificate is in for the painful process of transferring everything bound to the old certificate to the new one. Attempts to use the revoked certificate will trigger immediate warning flags, and the ability to determine the computer from which the attempt originated will help track down the culprit. At least with certificate-based access the user need only worry about guarding one credential. As the Secure Internet is put in place, users must be educated as to the importance of protecting their certificates.

Eavesdropping

Laws against wiretapping, electronic surveillance, including the laws and regulations governing when it may be employed by law enforcement, already exist and have for decades. The nature of the Internet, which permits packets to be "sniffed" on local networks or intercepted at intermediate relays between the sender and receiver, may facilitate eavesdropping, but every kind of crime committed by such means had already been committed before 1870 on the public telegraph network, mostly by corrupt operators in telegraph offices. In any case, once all traffic on the Internet is encrypted, eavesdropping will become far more difficult than in any earlier mass communication medium, while meeting the needs of law enforcement subject to due process.

Scams and Securities Fraud

Offering bogus products for sale is fraud, regardless of how orders are solicited and taken, and has nothing to do with computers or the Internet—I received my first Nigerian money scam *delivered by the postman* in 1982! Fraudulent offers made over the Internet constitute fraud no more and no less than identical offers made by a telemarketer or sent by FAX; no additional legal sanctions are required. In fact, when accountability is restored to the Internet, it will be far easier to identify and prosecute the perpetrators of such crimes via the Internet than those committed through other means. The ability to identify the originator of messages of all kinds will deter "pump and dump" fraudsters in investment chat rooms and bulletin boards and provide the information needed to identify them should their behaviour merit investigation.

Spam and Other Unwanted Messages

Every communication medium has spawned its own form of mass marketing: in the post: junk mail; on your FAX: junk FAX; over the phone when you're trying to eat your dinner: telemarketers; I even remember blaring sound trucks trying to get out the vote on election days when I was a kid: junk noise. The only economic constraint on unsolicited commercial communications is the relationship between the cost to reach a given audience and the anticipated revenue generated from the message. Spam has become a plague on the Internet simply because an enormous number of pitches can be delivered at negligible cost compared to other media, so regardless of how crude the message or the clientele it is directed toward, there is a reasonable expectation more than enough bottom feeders will respond to turn a profit. Content-based filtering can eliminate a large percentage of junk mail but provokes an arms race between efforts to disguise junk mail and those of filters to unmask it. The advent of true accountability on the Internet will make mass Internet junk mail a thing of the past.

With the ability to identify the person (individual or business) responsible for transmitting a message (which will not be delivered without a verified sender certificate), those who abuse E-mail can be instantly and unambiguously blacklisted based on their certificates—messages they send will be discarded without any further need for user intervention. No government involvement is required—a customer of an Internet service provider can subscribe to one or more independent databases of junk E-mail offenders and filter based upon their identities. If a mass mailer attempts to obtain new certificates with fraudulent credentials or steals the certificates of others to forge messages with stolen identities, they are committing crimes for which they can be prosecuted. Investigating such offences will be facilitated by knowledge, from the sending machine(s) certificate, of the computer or computers where the fraudulent messages originated.

If an individual wishes to never see E-mail (or other communications: for example instant messages, chat room text, news group and bulletin board postings, etc.) from a given person, they need only press the "Ban" button in their client program whilst

reading an offending message: subsequent messages signed by that originator will be silently discarded or ignored. Since all of these media will only accept messages with a valid and verified certificate, filtering based upon it will be absolutely reliable.

Worms and Viruses

The advent of Trusted Computing and the Secure Internet will close most of the vectors through which the present-day plague of computer viruses and worms propagate. (The distinction between viruses and worms, murky enough already, is irrelevant to this discussion; in the interest of brevity, I will use "worm" to denote all kinds of self-propagating programs, malicious or otherwise.) As I write these words, my mail log is noting the arrival and automatic disposal of a `Sobig.F` mail worm about every three minutes—more than 250 a day for the last several weeks. Most worms propagate by exploiting security flaws, principally in the widely-used products of one major commercial software vendor. While it makes sense to focus in the near term on remedying existing flaws and avoiding the introduction of new ones, errors and unanticipated consequences are inevitable in all fields of engineering, and systems should be designed to be robust even in their presence.

It is instructive to observe that most of the recent large-scale outbreaks of worms have propagated due to vulnerabilities *already found and fixed* well before the programs which exploit them were released into the wild. The worms were able to wreak their havoc on the Internet because tens of millions of computers had not yet been updated to versions of software which correct the vulnerabilities the worms depend upon. The combination of Trusted Computing and the Secure Internet will eliminate the risk posed by machines running software with known vulnerabilities. Every program executed by a Trusted Computing machine will be signed with a certificate from its supplier. Periodically, the operating system will re-validate the certificates of programs before running them. When a critical security flaw is discovered in a program, its supplier can revoke the certificate for the vulnerable version of the program. When a user next attempts to run the program, the operating system will discover the certificate revocation and refuse to run it until the user downloads and installs a patch or updated version which corrects the problem. The certificate revocation will usually direct the user to the required update which, based on the user's preferences, may be installed automatically. Obviously, validating a program's certificate requires the machine to be connected to the Internet. But a machine which is not connected to the Internet can neither be infected nor infect other machines, and hence is neither at risk from, nor poses any risk to, others.

The Trusted Computing architecture will also guard against the mechanisms worms use to infect users' machines and cause subsequent damage. Most mail- and Web-based worms work by tricking the user or the user's Web browser or mail client into executing a program which runs under the current user's permissions. A Trusted Computing platform will execute *no* program which does not bear a valid certificate signed by a supplier the user has chosen to trust, with a signature that matches the program about to be run. Hence, even if the user is tricked into attempting to run a program, all that will happen is that a warning box will pop up indicating the program does not bear a certificate from a trusted vendor. Expert users will be able to add vendors to their trusted list, but the typical user will have no need to do so. Should a rogue vendor be discovered releasing malicious software, revoking the vendor's certificate will automatically disable all programs signed with it. Even if a malicious program manages to bypass all these safeguards and infect one or more executable files on a user's computer, the only consequence will be that they won't run until they're replaced with intact versions since, when launched, the signature in their certificates will fail to match the program files and/or (if the signature has also been modified) the local signature will fail to verify against the supplier's signature over the network.

Search Engines

At present, Web search engines are required to periodically "crawl" the Web to discover new documents and changes to documents already indexed. This is costly and wasteful of Internet bandwidth (especially with numerous engines all crawling the Web independently), and has a long latency time—days or, in some cases, weeks may elapse before a newly added document or site is indexed. When documents change or are deleted, search engines don't discover this fact until they next crawl the site containing the document, resulting in their returning out of date and broken links much to the frustration of users.

The advent of document certificates (imprimatur) will eliminate these problems and allow search engines to stay current with Web content at a fraction of the present cost. For a document to be published on the Web, it must be accompanied by a certificate issued by a document registry identifying its originator and containing a signature permitting its content to be validated. As described above, without this document certificate, it will not be transferred across the Secure Internet. When a document certificate is obtained from a document registry, a log entry will be made identifying the document source URL, signature, and publisher's certificate. This database will be available to the public, and search engines will use it to be immediately informed when a document is published, revised, or deleted (by revoking its signature). When these events occur, the search engine can immediately index the document or, in the case of revocation, purge references to it from its search database. Web crawling will still be necessary as a lower priority activity to detect broken links and documents with valid certificates whose URLs no longer work, but this will be a quality control measure, not the primary means of keeping a search engine up to date.

Plagiarism

Regardless of the constraints imposed by Digital Rights Management, plagiarism will always remain possible—after all, it existed in the days of quill pens and parchment; digital technologies may make it easier, but cannot prevent it. However, once a document certificate (imprimatur) is required for each publicly available document, plagiarism will be far easier to detect and respond to. The database of newly issued and revised document certificates can be used by copyright holders to scan for instances of their material being used without authorisation. Once an apparent violation is asserted, a dispute resolution procedure can be undertaken (judicial or arbitration) and, should a finding of plagiarism result, the offending document can be immediately un-published by revoking its certificate, which will immediately halt further damages to the owner. Digital Rights Management will render copies of the plagiarised document unreadable at the moment their certificate is checked. Only copies kept on machines never connected to the Internet will remain readable.

Since the process of reviewing content to detect plagiarism is quite similar to that employed by search engines to detect identical or similar documents, the operators of search engines may provide plagiarism detection as a value-added service to copyright holders. This would eliminate the need for each individual intellectual property owner to scan new documents, since the search engine is already examining them.

Patent Enforcement

A patent confers upon its holder the right to *use* the invention it discloses and, consequently, to license and regulate use of the invention by others during the patent's term. One infringes a patent by using the invention in a way covered by one or more of the patent's claims. This is the case even for a customer who unknowingly purchases a product which infringes a patent, or contains a subassembly which does. A computer

graphics chip which, for example, uses the exclusive-or (XOR) function to draw a cursor on the screen might be found to infringe the absurd patent (thankfully now expired) granted in 1980 to the purported "inventor" of that particular triumph of the human intellect. The owner of the patent could, in principle, bring suit for damages not only against the chip maker, but also against computer manufacturers who used the chip, and/or customers who bought computers from those manufacturers. This is rarely done, but the threat of such litigation is often sufficient to coerce the party at the top of the food chain (in this case, the graphics chip manufacturer) to pay a license fee to the patent holder to avoid the risk of a potential imponderable liability deterring customers from designing in the chip. I'm not making up the XOR patent example; I personally signed a check for US$25,000 to the . . . uhhh . . . fellow who owned that patent in 1985.

Once software products are signed with a certificate which is validated each time they are launched, patent holders will be in a far stronger bargaining position. If a deployed software product infringes a patent, the patent holder could seek relief in the form of revocation of the program's certificate, with users only permitted to use the program after purchasing an update which licenses the patented technology. To avoid the disruption and ill feeling such an event would engender toward the software vendor, the vendor would almost certainly opt to settle with the patent owner for a retroactive license covering their installed base. Micropayment may be an attractive option for patent licensing. A software application using a variety of patents in optional facilities might be sold with those features initially disabled. The first time the user wished to use one, the user would be asked to confirm payment for the right to use the patent(s) it employs, which would then be remitted directly to the patent holder.

Note that a Trusted Computing system will validate the certificate and signature of an operating system before booting it. Should an operating system be found to contain code which infringes a patent (or, say, contain facilities which permit illegal circumvention of Digital Rights Management), its own certificate may be revoked, requiring users to replace it with an operating system duly licensed by holders of all patents it employs and compliant with all laws protecting intellectual property.

Trolls, Flamers, Cranks, and Crackpots

Public discussion boards attract immature, maladjusted individuals faster than a 100 Watt light bulb draws insects on a still summer evening. These creatures enjoy nothing more than attracting attention to themselves by posting off-topic material, *ad hominem* and/or obscene attacks upon other participants, and starting or fanning tedious "flame wars" on contentious topics. As soon as others in the group block messages from one of these ogres, they pop up again under a different alias, posting from another free account. Individual certificates will put an end to this since a user's certificate will be required to post to the group and, once banished (or marked to be ignored by individual participants), the offender cannot assume another identity without fraudulently obtaining a different certificate. Note that a message board can provide anonymity to its participants by permitting them to use aliases or "handles", but the operators of the board will be able to determine the true identity of participants from their certificates. This will prove invaluable in cases where physical threats, slander, or other actionable behaviour occurs in such fora.

Protecting the Children

> Whenever a politician starts talking about "the children," keep one eye on your wallet and the other on your liberty.
>
> Anonymous

The ubiquity of Internet access and the practical difficulty of parents' supervising their children's activities online exposes children to inappropriate material and potential real-world risks to an extent other media do not. This is largely due to the anonymity and lack of accountability of the present-day Internet and will be remedied by the Secure Internet.

When all Internet access requires logging on with a valid certificate, the risks posed by anonymity will be eliminated (absent certificate theft or fraud). Since minors will be able to obtain certificates only with parental consent, and a parent has the option of causing a child's certificate to indicate his or her age, a minor will be clearly identified and sites containing material not intended for such young eyes will, in their own interest, restrict access to those with certificates indicating visitors are of appropriate age. The ability to link a minor's certificate to that of the parent or legal guardian will permit supervision of the child's online activity, including retrieval of sites visited and E-mail sent and received.

Note that it's up to the parent whether a child's certificate indicates minor status, what age is given, and whether it is linked to the parent's for supervisory purposes. Permissive parents or those who consider their children sufficiently mature to use the Internet on an adult basis are free to obtain unrestricted certificates for them.

Adult Sites

Sites containing "adult" (i.e. pornographic) material are generally required by law to restrict access to those 18 years or older. Today, most simply require visitors to assert they're of age, with no means of verification. Of course we know that all good children would never proceed past such a banner since that would involve telling a fib. Unfortunately, the set of such good children who arrive at such a site in the first place appears to be of approximate measure zero.

Once access requires a certificate, adult sites will query the age of each visitor. If the visitor is a minor below the minimum age for the site, access will be automatically denied. Responsible operators of adult sites (another measure zero set?) will welcome the protection this affords them, as it immunises them against entrapment by children deliberately sent to the site to mount a legal attack on it.

Since "adult" material of various kinds is offensive to many adults as well, certificate-based access may permit certificate holders to specify categories of material they do not wish to see, perhaps with something like the Platform for Internet Content Selection (PICS) rating scheme. Based on the content preferences retrieved from the certificate of a visitor to a site, access to the entire site could be blocked, or selectively on a document-by-document basis. Search engines could store document content ratings and elide pages inappropriate to the searcher's profile from lists of results.

Child Pornography

The vast majority of self-described "adult" sites contain material unsuited for the eyes of children, but perfectly legal for adults to view, download, etc. But there are limits, and the Internet's anonymity has been exploited by those who transgress them, dealing in wares too distasteful to enumerate in these pages. Such material is *illegal* to create, possess, or distribute in any manner. The fact that the Internet happens to be involved does not create a new crime; it only makes detection and enforcement more difficult due to the present anonymous, unaccountable network architecture.

The Secure Internet will expose these dank and foetid recesses of the network to the sunlight of accountability. In order to serve pages across the Web, they must be identified with the certificate of the person or legal entity responsible for their publication. Since certificates can be obtained only by supplying complete identity information and forging them or using a certificate under false pretenses will be a crime, this means the

identity of the purveyor of illegal goods will be visible, and the person or persons responsible exposed to the applicable criminal sanctions. Once pages served by Web sites are required to bear an individual document certificate (imprimatur), the door will be open to identify potentially illegal material at the time the document certificate is granted or, if detected after the fact, to revoke its certificate and render already-distributed copies unreadable on Trusted Computing platforms. Finally, the fact that all users who access sites containing such material will necessarily transmit their certificates along with requests exposes them to the risk law enforcement may log their accesses and items viewed and downloaded. This will undoubtedly deter those tempted to seek such material.

Pædophiles and Predators

How about pædophiles lurking in chat rooms? Certainly this must be a computer crime? Well . . . no, because pædophiles existed long before there were computers, were subject to sanctions in all civilised societies, and nonetheless found ways to indulge their perversion. The anonymity of the Internet may help them commit their crimes, but the crimes are the same as before. The end of anonymity on the Internet will restore the risk attendant to this deviant and destructive behaviour, and provide the documentation law enforcement needs to identify and prosecute offenders.

Hate Speech, Community Standards

Most of the considerations discussed above in connection with child pornography apply equally to hate speech. Many jurisdictions prohibit dissemination of materials intended to sow hatred among racial and other groups, incite violence, and promote banned political movements. Some societies have stricter definitions of obscenity than prevail elsewhere. The operator of a site containing such material within a jurisdiction where it is illegal will find himself in the same situation as a child pornographer. Certificate based access will, in addition, permit control of cross-border flow of such material. If a certificate indicates a user resides in a jurisdiction where its content is banned, a site may choose to deny access in the first place. Even if the site allows the user in, the fact that the user knows their accesses to a site containing illegal material may be logged will discourage downloading it. Finally, nations with highly restrictive Internet content policies (a.k.a. hell-holes), can establish filtering points at their borders which block content they deem illegal or inappropriate, either based on the certificate of the publisher or those of the documents. Sites which presently hop among multiple IP addresses to avoid filtering may continue to do so, but unless they somehow manage to obtain an unlimited supply of valid certificates, it won't do them any good.

Terrorists, Drug Dealers, and Money Launderers

This unholy trinity is heaven-sent for those who would regulate the Internet. Invoking their names often seems sufficient to pass any legislation, however intrusive upon the privacy of all people. The Secure Internet will provide the facilities law enforcement needs to investigate and prosecute these and other criminals. The attributes which make the Internet so attractive to criminals are precisely those which will be reined in by the Secure Internet: anonymity, lack of accountability, and unrestricted privacy. The requirement that all Internet transmissions be identified by the certificate of the person responsible will pierce the veil of the anonymity criminals rely on to conspire without accountability. Encryption key recovery pursuant to due process will permit monitoring communications which rely upon it for security.

Computer security cognoscenti will immediately object that (1) requiring a certificate for Internet transactions will simply create a thriving market in bogus certificates, and (2) criminals will use their own encryption, steganography, and private codes to

thwart interception of their messages. This will certainly be the case to some extent, but real-time online validation of certificates will make large scale certificate fraud a difficult endeavour. Certificate authorities which cater to a shady clientele will soon find the certificates they issue next to worthless since they will either not be accepted (i.e. traffic blocked), or users of them subjected to an increased level of scrutiny, precisely as citizens of countries known for peddling passports to shady characters encounter difficulties crossing borders. If and when key recovery is mandated, one might immediately expect individuals concerned with the security implications of this, criminals in the forefront, to immediately begin pre-encrypting their messages with privately agreed keys which cannot be recovered by law enforcement. Certainly, this too will happen, but one is continually astounded how little encryption is used today by criminals communicating over the Internet, notwithstanding the wide availability of free, highly secure privacy protection tools. Besides, even if traffic is sur-encrypted, the very fact that it is largely confirms the suspicion which led to keys being recovered and messages examined in the first place. Further, traffic analysis based on the identity of the sender and receiver (from their certificates) and the sending and receiving computers' certificates can provide insight into the structure and operation of criminal enterprises.

Conditional Anonymity where Appropriate

Although the Secure Internet will provide end-to-end identification for all connections, resources on the network may choose to provide anonymity to participants. For example, support groups for those recovering from various afflictions, counseling services, and political discussion groups dealing with controversial topics may all wish to guard the identity of those who participate to dispel fear of public exposure and to encourage frank discussion. As today, participants could remain totally anonymous or choose a pseudonym (or "handle") to identify themselves.

Unlike the present-day Internet, the operator of a site which provides anonymity *will* be able to determine the identity of participants. It's up to them whether this information is discarded or kept in a secure form. Sites which permit users to participate under pseudonyms will be able to verify the match between the pseudonym and the user's actual identity, avoiding the problem of others forging pseudonyms and eliminating the need for logons and passwords for visitors to the site.

Sales Tax and VAT Collection

A contentious fiscal issue in electronic commerce is the ability of customers in the United States and the European Union to avoid local sales taxes/VAT by purchasing from merchants outside their taxing jurisdiction. With the advent of personal certificates for Internet users, fair taxation for online purchases can be easily implemented. Customers' certificates will indicate their legal domicile, so collection of sales taxes and remittance to the their domicile of residence will be straightforward, even for transactions which cross national borders. Micropayment can be used both for the goods purchased and the tax due on the sale.

Resolution of the sales tax/VAT issue simply brings Internet commerce into conformance with existing rules for merchants with multiple places of business. It defines "place of business" as accepting orders from a given jurisdiction and makes sales tax a condition of selling to customers there. Micropayment makes collection and remittance automatic and painless, even for the smallest merchants, and certificate identification of domicile makes avoidance impossible.

Funding the Internet Infrastructure

The larger issue of general Internet taxation remains to be resolved. In most jurisdictions, one pays a tax on telephone service; one pays for a television license, or part of one's general tax goes to support public broadcasting. Why should the Internet be exempt from such taxation? The Internet depends upon an infrastructure of fibre optic links, satellite communications, and physical assets all of which are vulnerable to disruption and must be defended by national and international military assets. The Internet is the most powerful tool for global civil discourse in human history, but it is vulnerable to uncivilised forces which would destroy it. Is it unreasonable that those who benefit from the Internet should pay to defend it? Ongoing research and development will be required as the Internet grows and evolves. Now that the Internet has moved from the laboratory to the marketplace, why shouldn't those who use it fund the investment in its future?

Certificate identification and micropayment will make this possible, without changing the character of the Internet. A minuscule per-byte or per-packet fee for data transmitted across the Internet would provide a revenue stream adequate to fund all of the neglected infrastructure protection and research projects required to keep the Internet open, secure, accountable, and reliable. Certainly, responsible politicians will restrict Internet taxation to cover only public expenditures directly supporting the Internet, resisting the temptation to exploit it as a cash cow for the general fund.

Employee Internet Abuse

It's an open secret that once Internet access and the Web reaches the desks of employees within an organisation, some will succumb to the temptation to surf the Web or otherwise occupy themselves with things other than what they're paid to do. Court decisions have upheld employers' broad discretion in monitoring the activities of their employees with company computers on company time. The ability to issue certificates to employees for work-related Internet access will facilitate such monitoring. A company can, if it wishes, only permit traffic signed with staff certificates it has issued out onto the Internet, and traffic addressed to its certificates to enter. This prevents employees using their personal certificates for non-work related items while on the job. Encryption keys from employee certificates can be recovered by the authority in the organisation which issued them, which will permit, when necessary, monitoring of employee activity and recovery of archived traffic and work when necessary.

Application File Reverse-Engineering

Many commercial software vendors consider the binary file formats used by their application programs to be proprietary intellectual property. Competitors who reverse-engineer such file formats benefit from the research and development effort expended in creating them without comparable investment and may, by creating supposedly compatible files in their applications which are not completely compatible, cause damages to customers of the original vendor (and that vendor's reputation) when errors occur attempting to use them with the original vendor's application.

Digital Rights Management provides a means to protect a vendor's investment in proprietary application file formats. On a Trusted Computing platform, application files may be signed with the signature of the application which created them, and will not be readable by other applications (unless they also have been granted the right as, for example, a Web page authoring tool may be permitted to import files created by the same vendor's word processing program). All of the technological and legal protections which underlie Digital Rights Management are thereby deployed to protect the vendor's file formats, and thus erect a substantial barrier to entry against competitive products.

Microsoft Office 2003 is reported to include such protection of its document files, incidental to the introduction of Digital Rights Management for documents created by Office users.

Objections

But, but ... I hear people sputter, more or less anonymously, from all over the IPv4 and nascent IPv6 address space. Let me now address some of the specific objections to the issues raised in this paper. I will doubtless expand this section based on the debate it engenders.

"It violates freedom of speech!"

Historically, the extent of government regulation of speech has differed from one medium to another, with the introduction of a new medium often seized as an opportunity to impose regulations deemed unacceptable in those which went before. Broadcast media have almost always been subject to content regulation far more stringent than that of print media and speech. Commercial speech (advertising) is generally not deemed protected speech, and is regulated in many ways. Most developed countries regulate political speech, particularly pertaining to electoral campaigns. (And, of course, many "developing" countries strictly control political discourse, which is one of the reasons so little ever develops there.)

Amateur radio operators are licensed by their governments, and the content of their transmissions monitored and regulated, with license revocation a consequence of violations. To date, the Internet has been subjected to very little regulation, but there is nothing inherent in the Internet which makes it immune to regulation in the future. Prior to 1912, there was no regulation whatsoever of radio in the United States; anybody could set up a station and transmit any content on any frequency they wished, much like the Internet today. Yet in 1912, amateurs were summarily limited to transmitting only on a wavelength of 200 metres (1.5 MHz), which severely limited the range of their transmissions. In 1917, the U.S. government shut down all amateur and commercial radio stations. Until the end of World War I, it was illegal for a U.S. citizen to possess an operational radio transmitter *or receiver*. Amateur radio was shut down throughout World War II as well, while commercial broadcasters continued to operate. None of these restrictions were ever successfully challenged on the grounds of freedom of speech. (Radio spectrum regulation was thrown out by a district court in Illinois in 1926 on *commerce clause* grounds, but quickly reestablished by the Radio Act of 1927.)

In the United States, one of the longest traditions and strongest guarantees of freedom of speech has proved no impediment to regulation of new media. The *Digital Millennium Copyright Act* is a recent example of restricting the technological means of distributing information [see, for example, the detailed prescriptions for copy protection of analogue video recordings in §1201(k)]. There is no reason to believe that comparable governmental mandates of the technologies described in this paper would be voided on the grounds of freedom of speech. Most other developed nations have a stronger tradition of government involvement in and regulation of telecommunications, and would be even less likely to oppose their extension to the Internet.

"It violates my constitutional right to anonymity!"

You have no right, constitutional or otherwise, to anonymity! As a citizen of a civil society, you are responsible for your actions. Society mandates accountability in numerous domains. You cannot drive on the highways without displaying a number plate, nor without

carrying a driver's license and vehicle registration which you must present to law enforcement on demand. You cannot block telephone caller ID when calling emergency services numbers, and law enforcement can obtain a list of numbers you've called and trace those calling you pursuant to a court order. You cannot transmit on the amateur radio bands without giving your government-assigned call sign. You cannot open a bank account, obtain a credit card, or buy or sell stock without providing your tax identification number. Records of all of your financial transactions may be subpoenaed or disclosed for law enforcement purposes. None of these are recent innovations—all have been true for decades, and none have occasioned public outcry or serious challenges on constitutional grounds.

Anonymity, where it exists, is either an inherent property of a technology, subject to subsequent restriction either technologically (caller ID) or through regulation and sanctions against violators (number plates on automobiles, radio call signs). Anonymity may, in many cases, be *granted* by a society when it is perceived to be beneficial. Support groups whose participants fear disclosure of their identity may disable caller ID on their telephones. Similarly, in the era of the secure Internet, a site may choose to guard the identity of those who visit it, to the extent of not logging their accesses by certificate. But such a site could, when necessary, block access from certificates of those who abused it or other participants. Anonymity, where not the consequence of technological constraints, is a privilege granted by society in certain circumstances, not a right.

"Clever people will always beat the system!"

Regardless of the crystalline purity of its design, any technology is only as good as its implementation, and given the dismal record of the information technology industry in implementing security, there is every reason to believe that as the technologies I describe here are deployed, there will be a multitude of cracks through which users, malicious or simply motivated by the challenge of it, can slip. But in the long run, that doesn't matter. In fact, having a large number of very bright people trying to subvert a system is an excellent way to find and plug the holes in it.

It's not as if one night everything will be as it was in 1999 and suddenly, the next morning, the entire Trusted Computing and Secure Internet suite of technologies will be in place. The process will be evolutionary, providing a transition for the enormous installed base, and will probably take a decade or more to be fully realised. Along the way, there will be a variety of interim steps and transitional mechanisms, all of which will have their own limitations and vulnerabilities. But slowly, surely *the screw will tighten*.

The experience with pirate decoders for premium satellite broadcast channels is an example of this process. Early scrambling schemes for analogue broadcasting were easily defeated with relatively simple devices. As progress in microprocessor technology increased the compute power of set-top boxes, increasingly secure scrambling technologies were deployed, each eventually meeting its match in ever more clever pirate decoders. But with the current generation of digital broadcasting, the game is more or less over. It may be possible for a tiny fraction of the super technologically elite to defeat present-day schemes but, even so, the revenue lost to the broadcaster is insignificant. Further, the ability to download new firmware into subscribers' set-top boxes on demand permits broadcasters to remedy any discovered security flaw without an expensive swap-out of installed hardware. In digital broadcasting, the subscriber's decoder card is simply a *certificate* identifying their account and, in many cases, is bound to the machine certificate of the receiver/decoder it is used in.

This is how security will come to the Internet. Bugs will be found in early implementations, and they will be exploited. They will be fixed, and new exploits will be found, and they in turn will be fixed. The process will converge until the number of Internet users able to circumvent the accountability of the new architecture is so mi-

nuscule as to be insignificant. By that time, legislation is likely to provide an additional deterrent to those tempted to hide or forge their identity, just as sanctions exist for the use of false credentials in other venues. The ability to revoke certificates for programs found insecure, forcing the installed base to upgrade, will provide the means to correct vulnerabilities as they are discovered, regardless of how widely deployed. Certificate revocation will not affect a machine which never connects to the Internet, but then such a machine can neither exploit nor be exploited due to flaws in programs installed on it.

" 'Not transferring without a certificate' is pure arm-waving! How could this really be enforced?"

In this paper I've deliberately left vague the details of how a Secure Internet will require certificate identification of packets it transmits and perform the necessary validation and end-to-end exchange of identity and authentication information. One could imagine this validation being performed at the point where a local network connects to the Internet, for example, where a home user's broadband line connects to their Internet service provider. In the interest of scaling, one would wish to push potentially costly operations such as certificate and signature validation as close to the leaf nodes of the network—individual clients—as possible. Trusted Computing platforms may be deemed sufficiently secure to perform this function in users' individual machines, although that runs a greater risk of circumvention.

Regardless of how and where the validation is performed, I think we can agree that it's conceptually possible, since there's little difference in establishing a certificate-validated and encrypted connection for an arbitrary Internet transaction than performing a public key login to a remote system with SSH and creating a secure tunnel between the machines.

"Micropayment has been tried and failed numerous times already. Why do you think it will work this time?"

Micropayment is scarcely a new idea. Ted Nelson described fine grained payment of copyright royalties as an integral part of Xanadu more than three decades ago (however, in Xanadu, royalties were to be collected as a surcharge to the basic fee for using the Xanadu system and remitted to the copyright holder by Xanadu). In the 1990's a variety of micropayment systems were launched, hoping to (digitally) cash in on the e-commerce explosion. Despite clever names like Digicash, Millicent, and Cybercoin, none was successful. If micropayment didn't catch on at the height of dot com fever, why should one expect it to succeed on the Secure Internet?

I believe there are two principal reasons for the failure of micropayment systems to date. First, none of them was integrated tightly enough into the architecture of the Web and users' browsers to make using it sufficiently transparent. A micropayment system fully integrated into the Secure Internet will automatically handle payments as a user clicks pay-to-view links with no requirement for user interaction whatsoever as long as the fee for the content being viewed is below the user's threshold of paying. Only the slow increase in the total sum spent for content displayed in the browser's window frame will distinguish the experience from reading free content on the Web today. (If a user doesn't want to automatically purchase any for-fee content, they need only set their threshold of paying to zero.) Absent this kind of tight integration with the browser, a micropayment exchange is really nothing but PayPal with a lower transaction cost and minimum payment size; if you had to go through the mechanics of a PayPal transaction to pay €0.0001 for a Web page, you'd quickly decide you didn't need to read that page after all. Legitimate concerns about fraud limit the extent current micropayment systems can operate without the user's involvement. On the Secure Internet, with access

by certificate to which the micropayment exchange is linked plus end to end encryption of all traffic, micropayment will be sufficiently secure to be fully automated for payments less than the user's threshold.

A second reason for the failure of micropayment so far is what might be described as the "Soviet store syndrome". In the Soviet Union, workers might receive a handsome salary in rubles, but it did them little good when the shelves in the stores which accepted rubles were almost always bare. A user who signs up with a micropayment exchange is likely to have the same experience: e-cash to spend, but nothing to spend it on because so few merchants accept payments through the exchange. Ask yourself: would PayPal have been a success without eBay users making auction payments with it? Once the Secure Internet and Digital Rights Management are in place, a wide variety of material not currently online at all due to fear of piracy will become available, most of it for a fee. That will "stock the shelves of the store" with goods which micropayment can purchase. Users won't adopt micropayment because it's cool or new, but because they want to buy stuff that's sold that way, just as eBay has sent far more customers to PayPal than PayPal to eBay, whose acquisition of PayPal is indicative of their symbiosis.

Once there is a wide variety of goods which can be bought with micropayment, micropayment exchanges will become viable businesses. While a variety of exchanges with different strategies and modes of operation will be launched, I expect they will eventually shake out and/or consolidate until there are about as many as there are kinds of credit cards, and as interchangeable. In fact, the credit card companies are likely to end up owning the surviving micropayment exchanges.

"Certainly you can't be advocating this!"

Well, duh ... of course not! But *this is where we are going, unless we change course, and soon*. Every single technology I discuss in this paper is either already deployed in a limited fashion, planned for adoption in the future, or under active development. Many of these technologies are beneficial if used wisely. But only Panglossian optimists will neglect the potential downside. Each of these technologies can be easily sold, either to individuals based on their obvious benefits ("No more spam", "Safe surfing for your kids") or to lawmakers in a position to mandate them due to their perceived societal benefits ("Close the Internet to terrorism", "Torpedo the copyright pirates", "Track down the child pornographers and lock up their customers").

In discussing these issues with numerous people over the last two years, I have been amazed at how few comprehended how all the pieces fit together in the way I saw them inevitably converging. Once I explained the end-point I envisioned, which I hope I've conveyed to you in this document, the general reaction was shock and horror, especially when I explained how every single component was already being developed or deployed.

I have little truck with conspiracy theories. Most of the people advocating or implementing the technologies which underlie the Digital Imprimatur mean well and sincerely believe their work will ameliorate one or more specific problems. But in technology, the sum is often more than the parts. You can cause much more mischief with explosive assembly *and* isotope separation than with either one by itself. And even though *many* people involved with these technologies haven't seen the big picture, is it safe to assume *nobody* does, or will? If a balding programmer in Swiss cow country can figure it out, is it safe to bet none of the People In Charge haven't? Not on your life. Not on our Internet.

If I thought there were the slightest possibility that refraining from publishing this document would reduce the probability of the advent of the Digital Imprimatur, you would not be reading it. But I don't; in fact, I'm convinced that the only hope for preserving the Internet as we presently know it is to alert as many technologically literate people as quickly as possible to where we're going and the consequences once we arrive.

As in my *Unicard* paper, I've cast the bulk of this document as a seductive sales pitch in *favour* of the technologies I fear, since that is how they will be sold to those whose liberty they will eventually restrict. To counter such arguments, one must fully appreciate how persuasive they can be when presented only in the light of their obvious benefits.

When Will It Happen?

When forecasting trends in technology and society, it is often easier to predict the destination than estimate the time of arrival. This is certainly the case with a collection of technologies as disparate as those discussed here, deployed across a geometrically growing global network connecting more than a hundred million computers and five hundred million people. Such a large installed base, and the compromises required to keep up with its ongoing growth, create great hysteresis in the system. And yet new technologies can be rapidly adopted; one need only look at broadband to the home or Wi-Fi for examples.

Deployment of Trusted Computing, Digital Rights Management, and the Secure Internet are, by their nature, primarily a "vendor (or government) push" effort rather than "market pull", so matters of strategy on the part of those who wish to see these technologies deployed must be taken into account. It is likely they will be introduced in conjunction with desirable new features which induce customers to accept them. (For example, Version 9 of Microsoft's Windows Media Player incorporates some Digital Rights Management technology, but users upgrade to it not because they're hungry for DRM, but to obtain other features it includes.)

Trusted Computing Deployment

Work is already underway to develop and deploy Trusted Computing systems. In their August 2002 business overview, Microsoft said of their own project, then codenamed "Palladium", since renamed the "Next-Generation Secure Computing Base for Windows":

> "Palladium" is a long-term endeavor. The first "Palladium"-enhanced personal computers will not appear on the market for several years, and Microsoft does not foresee widespread adoption for some years after the introduction. However, now is the time to begin planning for—and working on—"Palladium."

BIOS manufacturers are already at work on chipsets to support Trusted Computing operating systems, and hardware manufacturers are designing the "sealed storage" such systems will use to prevent unauthorised access to protected data. As with the roll-out of any technology, it will be a protracted process, probably taking longer than even conservative estimates, and there will doubtless be stumbles and changes in direction along the way. Yet the destination is clearly defined, and the key technological players are investing heavily in the effort to get there. Barring surprises, I expect the overwhelming majority of new computer systems sold in the year 2010 to include Trusted Computing functionality.

Digital Rights Management Deployment

Digital Rights Management deployment is presently underway; current mass market multimedia players are beginning to support various schemes, and as online commercial sales of multimedia content as exemplified by Apple's iTunes Music Store expand, increasingly more secure and restrictive implementations will follow, culminating in the eventual integration of Digital Rights Management with Trusted Computing.

Secure Internet Deployment

A logical point at which one might expect implementation of the Secure Internet to begin in earnest is concurrent with the mass deployment of the IPv6 protocol. Observers of the Internet scene may immediately heave a sigh of relief, since IPv6 is one of those technologies of tomorrow which remains securely anchored in tomorrow no how many tomorrows pass into yesterdays. It is ironic that had IPv6 been aggressively adopted starting in 1995, some of the accountability problems of today's Internet would not have become as serious as they are today (see Appendix 1 for details). Still, there is nothing in the architecture of the Secure Internet as I have described it in this paper which requires IPv6 in any way; should IPv6 be indefinitely delayed or supplanted by a different design, the introduction of the Secure Internet need not be delayed.

The consequences of the Secure Internet will only fully be realised when most machines connected to it incorporate Digital Rights Management and Trusted Computing technology. It is probable that major efforts to put the Secure Internet in place will be deferred until those technologies reach the market in large numbers. If we estimate a date of 2010 for that, then the years 2008–2015 could see the Secure Internet replace the present architecture.

The Imponderable Surprise

Predicting when new technologies will be adopted is difficult enough even when you only extrapolate present-day trends and assume they will continue into the future. (Actually, I believe you can win most bets on the date the newest whiz-bang thingo will be widely adopted by always guessing "never".) In the real world, trend lines rarely behave as nicely as they do on econometric forecast charts (when looking at one, always remember that "con" comes before "metric" in "econometric"). Surprises happen, and they can have enormous consequences. Despite the recent plague of Internet worms and viruses, some very knowledgeable observers of the Internet believe that we have been lucky so far in that the attacks to date have been much less virulent than they could have been. The paper "How to own the Internet in Your Spare Time" estimates that an "optimally" designed worm might subvert more than ten million Internet hosts, with an initial rate of spread so fast, perhaps infecting 300,000 hosts in less than fifteen minutes, that system administrators would be unable to react quickly enough to limit the damage.

An attack of this nature, particularly if found to be deliberate state-sponsored or subnational information warfare, which caused major disruption to Internet-dependent infrastructure, may result in a drastic acceleration of the timetable for the implementation of the Secure Internet, driven by government mandates rather than market evolution, and short-circuiting the dialogue about design choices and their consequences which would normally occur. Let us hope we remain lucky.

Summary and Conclusion

Global Internet,
Once a spring of liberty,
Autumn chill so near.

Over the last three decades the Internet has evolved from a research tool linking a handful of elite sites into a global mass medium. Its rapid, and often reactive evolution has resulted in a present day architecture widely perceived as inadequate to hold users accountable for their actions, providing unwarranted anonymity to disruptive and destructive actors, and placing intellectual property at risk in disregard of applicable law and with impunity to its sanctions.

A collection of technologies in various states of design, development, and deployment promise to remedy these perceived shortcomings of the Internet. If implemented and extrapolated to their logical conclusion, the result will be an Internet profoundly different from today's and at substantial variance with the vision of its original designers. More than any innovation in the last century, the Internet empowers individuals to spontaneously teach, learn, explore, communicate, form communities, and collaborate. Measured relatively, this individual empowerment comes at the expense of the power of governments and large commercial enterprises, reversing a trend toward concentration of power more than a century old which has acted to reduce free citizens and productive individuals to subjects and consumers.

Power, especially concentrated power, is rarely relinquished willingly. Each technology proposed to ameliorate supposed problems with present-day computing and network architectures must be carefully examined, individually and in conjunction with others, for the potential it holds to shift the balance of power back from the individual to the centre; to supplant the peer architecture of the Internet with a producer/consumer model more comparable to publishing and broadcasting. Technologies should also be evaluated for the potential they have to create (or restore) central points of control in the flow of information and interaction among individuals; to impose hierarchy upon a structure designed for equality.

Technology changes rapidly, but social, political, and economic structures are slower to adapt and far more persistent. Government telegraph and telephone monopolies in Europe endured more than a century. So far, the Internet has evolved organically, largely free of influence from the societies in which it is embedded. Having become so integral a part of the economy and communication infrastructure in the developed nations, the Internet and society must now come to terms with one another and sort out how things are going to go forward from here. This process is presently underway, and is likely to be largely settled by the year 2010; the resulting architecture is likely to remain in place for a good part of the 21st century.

The components of this emerging architecture are already on the table, and various players are beginning to explore how they fit together into a whole. In this paper, I've tried to acquaint you with the basics of these components, show how each can be promoted on its obvious merits as the solution to widely-perceived problems, then sketch the possible implications, some of them dire, which may result as these components are used in combinations and to ends those advocating them seldom discuss.

In the last years of the 20th century, we lived through the false dawn of Internet commerce; wildly unrealistic expectations were raised, and fortunes made and (mostly) lost chasing after them. But all the years since the early 1970's have really been one long dawn for the Internet, beginning with a barely perceptible glimmer, then growing brighter and brighter until it illuminated all but the darkest regions of the world. Even today, only a tiny fraction, less than 10
the global population, has ever used the Internet, and even in the most extensively wired societies we have only begun to explore its potential to augment all forms of human interaction. Compounded geometric growth causes problems—the fact the Internet has not collapsed already is one of the most significant testaments to the wisdom and foresight of those who built it. Today, the problems are evident, and people are at work attempting to solve them. Whatever solutions are adopted (or not adopted—one may rationally choose to live with problems if the solutions are worse), are likely to be with us for a long time. Whether they preserve the essential power of the Internet and its potential to empower the individual or put the Intenet genie back into the bottle at the behest of government and media power centres who perceive it as a threat will be decided over the next few years. That decision will determine whether the long dawn of the Internet was, itself, a false dawn, or will continue to brighten into a new day for humanity.

Appendices

Appendix 1
The Anonymous Internet: An Historical Accident

Many users who came to the Internet during its transformation into a mass medium in the 1990's consider anonymity and the consequent lack of accountability as fundamental features designed into the Internet or inherent in its technological implementation. Ironically, nothing could be further from the truth: today's anonymous Internet is largely an accident of technological change and exponential growth transforming an architecture designed for a different era.

The balance of this appendix is rather technical. Policy-oriented readers unfamiliar with the details of Internet technology and history who are willing to stipulate my assertion in the last paragraph or at least suspend their disbelief for the purposes of this document should feel free to read no further.

The Way Things Were

When the ARPANET was originally created, and for much of its subsequent evolution into the Internet prior to the mass pile-on of the 1990's, access was highly accountable. Every machine on the Internet had a unique Internet Protocol (IP) address, the 32 bit number you usually see written in "dotted quad" form, such as 192.168.114.31. Virtually all machines on the Internet were permanently connected via dedicated leased lines (or on a local network, itself linked to the outside by a leased line). IP addresses were permanently assigned, one per machine, from blocks of contiguous addresses allocated to organisations to whom was delegated the responsibility of assigning addresses within that block. An organisation responsible for an IP address block could, in turn, delegate assignment of addresses within a sub-block to another entity but, in every case, it was possible to determine who was ultimately responsible for an address on the Internet.

For example, Fourmilab was connected to the Internet in 1994, the closing days of this era. Fourmilab "owns" the block of 256 consecutive IP addresses starting at 193.8.230.0, assigned by the RIPE Network Coordination Centre which, itself, is responsible for all addresses from 193.0.0.0 through 193.255.255.255. Simply starting with an IP address, say 193.8.230.138, a simple series of queries can determine who is responsible for that machine's presence on the Internet, namely me.

Further, during most of the early developmental phase of the Internet, most machines on the network were timesharing systems which supported numerous users (dozens, hundreds, or even thousands in the case of some universities and large companies), each with a login account created by the system's administrator. Even though multiple users shared Internet access through the timesharing machine's IP address, their access was explicitly granted and activities logged as part of the accounting facilities such systems provided.

Now, in this environment, as the Internet was originally conceived to be, individuals are highly accountable for their actions. Those whose connections date from this era remain so. For example, suppose a machine within the Fourmilab address range commits some foul deed: relaying unsolicited mail, scanning other machines for vulnerabilities, or providing a repository for files which violate the copyright of third parties. One need only note the IP address of the culprit, query the top level address map to determine it lies within the range belonging to RIPE, then query RIPE to point the finger at me. If I have, in turn, delegated the address in question to somebody else, then I am responsible for their actions on the Internet and, should it come to that, subject to a court order to identify them and/or terminate their connection.

The Address Space Crunch

The Internet address protocol presently in use, Internet Protocol Version 4 (IPv4), was designed for the era I've just described, when the Internet connected a relatively small number of elite government, commercial, and educational sites, most of whose users accessed it through accounts on a limited number of timesharing systems. IPv4 provides for a 32 bit IP address, which permits (ignoring special purpose and reserved address blocks), a total of 2^{32}, roughly four thousand million unique addresses. With the world's population only 50
of the Internet, this might seem adequate, but the existence of pre-assigned blocks of IPv4 addresses and the inevitable inefficiency of any block allocation scheme made it clear that the traditional fixed assignment of IP addresses would result in the effective exhaustion of address space sometime in the 1990's.

Now the obvious solution to running out of address space is to increase the length of the address field. In July 1991, the Internet Engineering Task Force (IETF) began this process, which culminated in the 1995 publication of RFC 1883, *Internet Protocol Version 6 (IPv6) Specification*, which provided for an address field of 128 bits. This space is so vast it solves all conceivable addressing problems for the foreseeable future. (I do not consider the Omega Point and other fantasies as foreseeable for the purposes of this document.) If each individual Internet user were assigned their own permanent, private 48 bit address space (65536 times larger than the *entire present-day IPv4 Internet*), as recommended in RFC 3177, the 128 bit address could theoretically accommodate 2^{80}, or more than 1.2×10^{24} users—2×10^{14} times Earth's current population. Even with a necessarily sparse assignment of addresses to individual sites due to practical considerations and reservation of most of the address space for unanticipated future requirements, IPv6 can accommodate 178 thousand million users (more than seventeen times the projected world population in the year 2050), each with their own private 48 bit address space, without difficulty.

Victim of Success

While IPv6 would have neatly solved all of the address space problems of the emerging mass Internet, it arrived on the scene at just about the worst possible time to be rapidly implemented. By December 1995, when RFC 1883 was published, the Internet was already in its full-on geometric growth phase, with every expectation that growth would continue for years to come. The Internet infrastructure was already staggering to support millions and millions of new users with the unprecedented bandwidth demands of the graphics-rich Web and innovative services such as Internet telephony and video on demand.

Internet service providers and infrastructure suppliers were scrambling to keep up with this exploding demand which, at times, seemed almost an insurmountable opportunity. Fully implementing IPv6 means changing *everything*—host operating systems, applications, routers, network hubs: you name it. While IPv6 makes it relatively easy to "tunnel" IPv4 across an IPv6 network, easing the transition for legacy sites, full adoption of IPv6 in 1995 would have meant swapping out or upgrading an enormous existing investment in IPv4 infrastructure, right at the very moment when simply keeping up with the growing demand stretched capital, manpower, and the manufacturing base to their limits.

Stretching the Address Space

With no possibility of migrating to IPv6 in time to solve the address space crunch, the industry resigned itself to soldiering on with IPv4, adopting the following increasingly

clever means of conserving the limited address space. Each of these, however, had the unintended consequence of transforming the pure peer relationship originally envisioned for the Internet into "publisher" and "consumer" classes, and increasing the anonymity of Internet access.

Dynamic IP Addresses

The most obvious way to conserve address space was simply to observe that the vast majority of individual Internet users at the time connected via dial-up modem links, which they only used for a limited amount of time on a typical day. Since these users could not receive packets across the Internet except when their dial-up connections were active, there was no reason to assign each a unique IP address. One could simply give each dial-up *line* an IP address which a user connected to it would employ while connected. This approach, while eminently reasonable given the circumstances and perfectly acceptable to those who used the Internet only for Web and FTP access and E-mail, nonetheless created the first segmentation of Internet users into two classes. While a user with a fixed IP address could receive arbitrary connections from any other user while connected to the Internet (if the user's connection was dial-up and the user wasn't online, remote connections would simply fail), the dial-up user with a dynamic IP address typically received a different IP address every time they connected to the Internet, one determined by which modem at the Internet Service Provider happened to pick up the call.

A dynamic IP address user could, while connected, access any service on the Internet just like the fixed IP address user, but unlike the latter, there was no way for other users to know the dynamic IP user's address, since it varied from session to session. This meant there was no way for another user to establish a connection to the dynamic IP user, since even when connected, their IP address was unknown. To work around this problem, services were created such as the Speak Freely Look Who's Listening servers, ICQ, and Dynamic DNS or No-IP, which provide, in essence, places where users can meet via some invariant name, then exchange their current IP addresses to carry on the conversation on a peer to peer basis. Of course, a central server creates the risk of the kind of single-point failure the Internet was designed to avoid, and provides a central point of control, as users of Napster discovered.

Network Address Translation

As home Internet users began to obtain persistent, broadband connections, at first (mostly in Europe) through ISDN, then with cable television modems and Digital Subscriber Link (DSL) access over the telephone network, users increasingly wished to permit multiple machines to share the fast Internet connection. With a single IP address, whether fixed or dynamic, routing packets as the design of the Internet originally envisioned would permit only one computer on a local network to use the Internet at a time. Imagine the conflict between Dad doing stock market research in the evening and the teenagers in their rooms!

The most widely adopted technique to permit multiple computers to share a common Internet connection with a single IP address is *Network Address Translation* (NAT), as explained in RFC 1631. (Note that this 1994 RFC is *not* an Internet standard, but rather a description of a technique already in use which operates within existing standards.) NAT is usually implemented within the Internet router which provides access to the broadband link, but may also be performed by a firewall or in software on a computer to which the broadband line is connected.

NAT defines two independent but interconnected sub-networks. The *internal* subnet contains all of the local computers. Each is given its own unique IP address, either fixed or assigned dynamically with the Dynamic Host Configuration Protocol. Whichever,

these addresses have meaning only on the local subnet; they are usually assigned within one of the blocks reserved for private networks not connected to the Internet. NAT exploits the fact that the principal Internet services allow any number of simultaneous connections between a pair of IP addresses, each distinguished by 16 bit source and destination port numbers. For example, when you're running two different Web browsers on the same machine and simultaneously connect to the same Web site (for example, to compare how the two display the same page), what keeps the two from "stepping on one another" is the fact that even though both are sending packets back and forth between the same pair of IP addresses (your machine and the Web site), each uses a different, and unique, port number when accessing the site, and the Web site returns its response back to the port of origin.

NAT simply takes this one step further. When one of the machines on the internal subnet sends a request to an external host, the NAT box assigns a unique port number and makes an entry in a translation table, then forwards the request to the remote site as originating at the assigned port from the IP address belonging to the NAT box. When the reply comes back, the NAT box looks up the port number, determines which local machine it belongs to, and dispatches the packet across the local subnet to that machine. For Web browsing and most Internet applications where the home user establishes connections to an external server, NAT works almost like magic. There is no need to modify applications on individual users' machines nor install any special software: if the NAT router provides DHCP, as most do, adding a new machine to the internal subnet is as easy as plugging its network cable into the hub.

But NAT, like dynamic IP addresses, once again divides those on the Internet into two classes. While the dynamic IP user's address changed with every online session, at least *for the duration of that session* it remained constant and was accessible from the outside world, just like a permanently connected host. Once users exchanged their current IP addresses through one of the server-based schemes, they were free to then open any kind of connection between their machines supported by Internet protocols. The NAT user, however, finds himself at a further level of remove from the "raw" Internet.

Recall that the NAT box assigns a port for each connection from a machine on the internal subnet to an external site only when the local machine initiates a connection. Otherwise machines on the internal subnet are *completely inaccessible* from the Internet at large—the other side of the NAT box. They do not have an externally visible IP address at all, fixed or dynamic, and there is no way an external site can communicate with them unless the local machine has first initiated the connection. A machine behind a NAT box cannot act as a server, because there is no address which remote sites may use to open connections to it. Two users behind NAT boxes cannot even create a peer to peer connection between themselves, since neither has an address which will accept connections initiated from outside. If they wish to communicate, they must both connect to a server (not behind a NAT box) which will then relay the data between them. This creates a point of control far more powerful than a "meeting place" server for dynamic IP address users. A server which forwards traffic between NAT users must have sufficient bandwidth to accommodate not only lookup requests but all the data sent between its users, and has the ability to monitor and potentially intercept all the traffic it relays.

Many users behind NAT boxes consider the restrictions they impose as positive benefits. The inability of external sites to open connections to machines behind the NAT box means it behaves as a firewall, blocking traffic from the Internet which attempts to exploit vulnerabilities on local machines. Machines can still be infected by mail worms, viruses in software downloaded by local users, or vulnerabilities in Web browsers to malicious sites users are induced to visit, but at least the unlocked front door which an unsecured machine permanently connected to the Internet represents has been slammed in the face of potential attackers. But at the same time, the NAT user is no longer a peer of all other Internet users, as originally envisioned. The NAT user has become much

more a *consumer* of services published by sites with persistent, externally accessible IP addresses. If the NAT user wishes to create a Web site, post some files or pictures for downloading, manage a discussion group, or author a Weblog ("blog"), he will have to avail himself of the services of a public site to do so—his own machine cannot provide these services to external users, since they cannot create connections to it.

(Note: as there is no standard for NAT, various implementations work in different ways, some of which permit establishment of connections by external hosts to which an open port number has been communicated. These capabilities, when present, can be equally deemed security risks or benefits depending on an individual user's preferences. Here I've presented NAT in its purest form which, as best I can determine, is how the majority of contemporary implementations work. Cable and DSL providers, which have incentives to prohibit their mass market customers from operating servers, may enforce this restrictive NAT in the routers they provide to such clients.)

Public Wireless Access. The advent of public wireless Internet access amounts to taking NAT on the road. As the wireless infrastructure is put in place, users will be able to access the Internet, via NAT connections, from whatever wireless access point to which they happen to connect (assuming they have permission to use it). This doesn't really change things as long as some kind of commercial account is required for the access (as that presumably creates an audit trail, however difficult to trace, of the user's momentary identity), but free public access in tony coffee shops, bookstores that wannabe coffee shops, and the like permit totally anonymous connection to the Internet.

The Way Things Are

In summary, from an original architecture in which every Internet host had a unique IP address with a *person* behind it, where virtually all Internet traffic was essentially *accountable*, as a consequence of the transformation of the Internet from an interconnection of the elite to a mass medium and a series of work-arounds for the technical challenges in getting from there to here, the Internet has lost much of its original accountability. Many operations on the Internet are effectively anonymous. Identifying the person or persons responsible for various forms of abuse, whether overtly illegal or in violation of contractual terms of service, often ranges from difficult to completely impossible, which is increasingly making the Internet a very different kind of place from most of civil society where individuals are held responsible for their actions.

Appendix 2
Citing Sections of this Document

All of the section headings and certain list items in this document are the object of HTML fragment anchors, which permit you to link directly to the section or cite it in discussions. When you move the mouse over one of these anchors, the background colour changes to light green, indicating its presence. To make a link to the anchor, with the anchor highlighted pop up your browser's auxiliary menu (with most browsers and systems, click the right mouse button) and choose the item which copies a link onto the clipboard. You can then paste the link into your text.

If your browser doesn't fully support Cascading Style Sheets, the anchors may not highlight when you move the mouse over them, but you'll still be able to copy the link. If your browser does not support JavaScript/ECMAScript or you've disabled it, regular clicking on an anchor will scroll the document so the anchor is at the top of the window; if this irritates you, don't do it.

Glossary

certificate Block of data which uniquely identifies an object: person, document, computer, etc. Issued by a *certificate authority*.

certificate authority Organisation which issues a *certificate*, after validating the credentials supplied. Certificate authorities provide online verification of certificates and can suspend or revoke certificates at the request of the holder of pursuant to a court order.

computer certificate A *certificate* which uniquely identifies a computer system, installed when the machine is manufactured.

digital imprimatur The *certificate* which must accompany a document transmitted across the *Secure Internet*. See *document certificate*.

Digital Millennium Copyright Act (DMCA) United States Public Law 105–304 which criminalises, among other things, reverse-engineering and circumventing copy protection technology.

Digital Rights Management (DRM) A collection of technologies which embed enforcement of intellectual property rights in computer hardware, software, and media.

document certificate The *certificate* which must accompany a document transmitted across the *Secure Internet*. The certificate identifies the publisher and contains a signature of the document's content. Document certificates are issued by a *document registry* and may be verified online.

document registry An organisation which issues *document certificates* and validates them online.

flame An intemperate message, usually posted to an Internet discussion group or mailing list, often in the hope of stimulating other flames. Hence *flamer*, one repeatedly who posts such messages.

Micropayment A facility for making secure payments between Internet users with financial accounts linked to their certificates. While payments of any size can be made, micropayment has overhead sufficiently low that extremely small payments are practical.

Network Address Translation (NAT) Mechanism which permits multiple computers to share a single Internet connection by mapping requests from local IP addresses onto ports of a shared address.

revocation, certificate The act, by a *certificate authority* or *document registry*, which renders a previously issued certificate invalid.

Secure Internet Suite of technologies including access by *user certificate*, end to end encryption of all Internet traffic, and the requirement that each document transferred bear a *document certificate* issued by a *document registry*.

spam Unsolicited mass commercial E-mail, not to be confused with SPAM®, a yummy processed pork product.

threshold of paying The maximum fee a user is willing to pay automatically to view a document or use a service. Amounts up to this limit are automatically paid from the user's *micropayment* account.

troll A person who posts messages, particularly on Internet discussion groups and E-mail lists, with the intent of fanning *flames* or inducing new participants to make fools of themselves.

Trusted Computing System A computer whose hardware and software incorporates protection against corruption of software and user data, attack from other systems, and security against compromise of user data.

user certificate A *certificate*, issued by a *certificate authority*, which uniquely identifies an individual or legal entity (corporation, government agency, etc.).

virus, computer A self-reproducing computer program which integrates itself into other software which it subverts in order to propagate itself and usually cause other mischief.

worm, computer A self-reproducing computer program which exploits other software in order to propagate itself and usually cause other mischief.

References

1. The Electronic Privacy Information Center's Microsoft Palladium information page.
2. Microsoft Next-Generation Secure Computing Base for Windows (formerly "Palladium") business overview and frequently-asked technical questions.
3. Walker, John. *Unicard*. 1994.
4. Staniford, Paxson, and Weaver. "How to own the Internet in Your Spare Time". *Proceedings of the 11th* USENIX *Security Symposium*, 2002.
5. *Digital Millennium Copyright Act* (U.S. Public Law 105–304).
6. Weinman, W. *Authenticated Mail Transport Protocol (*AMTP*)*. August 2003.
7. Kent, S. and R. Atkinson. RFC 2401: *Security Architecture for the Internet Protocol* (IPsec). 1998.
8. Housley, R. et al. RFC 3280: *Internet X.509 Public Key Infrastructure: Certificate and Certificate Revocation List (*CRL*) Profile*. 2002.
9. Deering, S. and R. Hinden. RFC 1883: *Internet Protocol Version 6 (IPv6) Specification*. 1995.
10. Internet Architecture Board and Internet Engineering Steering Group. RFC 3177: *Recommendations on IPv6 Address Allocations to Sites*. 2001.
11. Cerf V. RFC 3271: *The Internet is for Everyone*. 2002.