The Eternity Service

Ross J. Anderson

Cambridge University Computer Laboratory
Pembroke Street, Cambridge CB2 3QG
Email: ross.anderson@cl.cam.ac.uk

Abstract. The Internet was designed to provide a communications channel that is as resistant to denial of service attacks as human ingenuity can make it. In this note, we propose the construction of a storage medium with similar properties. The basic idea is to use redundancy and scattering techniques to replicate data across a large set of machines (such as the Internet), and add anonymity mechanisms to drive up the cost of selective service denial attacks. The detailed design of this service is an interesting scientific problem, and is not merely academic: the service may be vital in safeguarding individual rights against new threats posed by the spread of electronic publishing.

1 The Gutenberg Inheritance

In medieval times, knowledge was guarded for the power it gave. The Bible was controlled by the church: as well as being encoded in Latin, bibles were often kept chained up. Secular knowledge was also guarded jealously, with medieval craft guilds using oaths of secrecy to restrict competition. Even when information leaked, it usually did not spread far enough to have a significant effect. For example, Wycliffe translated the Bible into English in 1380–1, but the Lollard movement he started was suppressed along with the Peasants' Revolt.

But the development of moveable type printing by Johannes Gensfleisch zur Laden zum Gutenberg during the latter half of the fifteenth century changed the game completely. When Tyndale translated the New Testament in 1524–5, the means were now available to spread the word so quickly that the princes and bishops could not suppress it. They had him executed, but too late; by then some 50,000 copies had been printed. These books were one of the sparks that led to the Reformation.

Just as publication of the Bible challenged the abuses that had accreted over centuries of religious monopoly, so the spread of technical know-how destroyed the guilds. Reformation and a growing competitive artisan class led to the scientific and industrial revolutions, which have given us a better standard of living than even princes and bishops enjoyed in earlier centuries. Conversely, the societies that managed to control information to some extent became uncompetitive; and with the collapse of the Soviet empire, democratic liberal capitalism seems finally to have won the argument.

But what has this got to do with a cryptology conference?

Quite simply, the advance of electronic publishing has placed at risk our inheritance from Gutenberg.

Just as advancing technology in the fifteenth century made it very much harder to control information, so the advances of the late twentieth are making it very much easier. This was made clear by recent court action involving the 'Church of Scientology', one of whose former adherents had published some material which the organisation would prefer to have kept secret. This apparently included some of the organisation's 'scripture' that is only made available to members who have advanced to a certain level in the organisation.

Since Gutenberg, the publication of such a trade secret would have been irreversible and its former owners would have had to cope as best they could. However, the publication was in electronic form, so the scientologists got court orders in an action for copyright infringement and raided the primary site in the USA in August 1995. They then went to Amsterdam where they raided an Internet service provider in September, and filed for siezure of all its assets on the grounds that their copyright information had appeared on a subscriber's home page. Their next move was to raid an anonymous remailer in Finland to find out the identity of one of its users. The saga continues.

The parallel with earlier religious history is instructive. The Bible came into the public domain because once it had been printed and distributed, the sheer number of dispersed copies made it impossible for the bishops and judges and princes to gather them up for burning.

However, now that publishing has come to mean placing a copies of an electronic document on a few servers worldwide, the owners of these servers can be coerced into removing it. It is irrelevant whether the coercion comes from wealthy litigants exploiting the legal process, or from political rulers conspiring to control the flow of ideas. The net effect is the erosion of our inheritance from Gutenberg: printing is 'disinvented' and electronics document can be 'de-published'. This should concern everyone who values the benefits that have flowed from half a millenium of printing, publication and progress.

So how can we protect the Gutenberg Inheritance?

Put into the language of computer science, is there any way in which we can assure the availability of data when the threat model includes not just Murphy's ferrite beetles, the NSA and the Russian air force, but Her Majesty's judges?

2 Preventing Service Denial

This problem is merely an extreme case of a more general one, namely how we can assure the availability of computerised services. This problem is one of the traditional goals of computer security, the others being to assure the confidentiality and integrity of the information being processed.

Yet there is a strange mismatch between research and reality. The great majority of respectable computer security papers are on confidentiality, and almost all the rest on integrity; there are almost none of any weight on availability.

But availability is the most important of the three computer security goals. Outside the military, intelligence and diplomatic communities, almost nothing is spent on confidentiality; and the typical information systems department in civil government or industry might spend 2% of its budget on integrity, in the form of audit trails and internal auditors. However 20-40% of the budget will be spent on availability, in the form of offsite data backup and spare processing capacity.

There are many kinds of record that we may need to protect from accidental or deliberate destruction. Preventing the powerful from rewriting history or simply suppressing embarrassing facts is just one of our goals. Illegal immigrants might wish to destroy government records of births and deaths[1]; real estate owners might attack pollution registries; clinicians may try to cover up malpractice by shredding medical casenotes [Ald95]; fraudsters may 'accidentally' destroy accounting information; and at a more mundane level, many computer security systems become vulnerable if audit trails or certificate revocation lists can be destroyed.

There is also the problem of how to ensure the longevity of digital documents. Computer media rapidly become obsolete, and the survival of many important public records has come under threat when the media on which they were recorded could no longer be read, or the software needed to interpret them could no longer be run [Rot95].

For all these reasons, we believe that there is a need for a file store with a very high degree of persistence in the face of all kinds of errors, accidents and denial of service attacks.

3 Previous Work

Many papers purport to show that the average firm could not survive long for without its computers, and that only 20–40% of firms have properly tested disaster recovery plans. The authors of such papers conclude that the average firm will not survive when a disaster strikes, and that company directors are thus being negligent for not spending more money on disaster recovery services. The more honest of these papers are presented as marketing brochures for disaster recovery services [IBM93], but many have the appearance of academic papers.

They are given the lie by incidents such as the Bishopsgate bomb in London where hundreds of firms had systems destroyed. Some banks lost access to their data for days, as both their production and backup sites were within the 800 yard police exclusion zone [Won94]. Yet we have no report of any firm's going out of business as a result. A more recent IRA bomb in London's dockland area confirmed the pattern: it also destroyed a number of computer installations, yet companies bought new hardware and recovered their operations within a few days [Bur96].

[1] The population of California is said to have increased significantly after fire destroyed San Francisco's birth records in the wake of the great earthquake.