

You needn't eat spam (or worms)

The real reasons why spam still exists today – and what to do about it

Jeffrey Race

Many who would cure us of spam look in the wrong place – technology – for the answer. These well-intentioned analysts rightly see this menace as resulting from a state machine that can be tweaked, but they should look to the I/O relationships of human behavior rather than communications protocols for the solution.

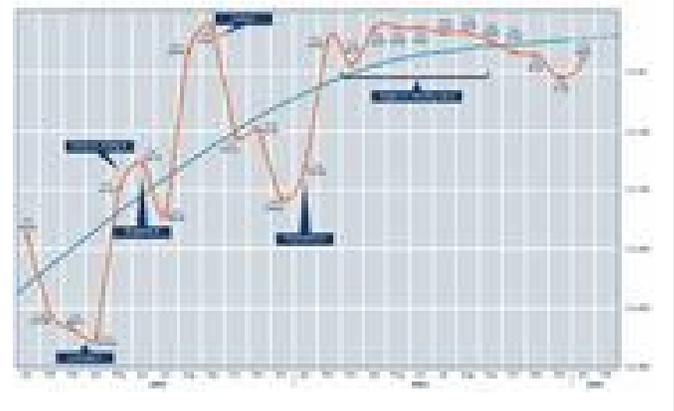
There is virtually no legal way to upload spam in the United States and many other countries, due to contractual bans imposed by backbone providers on their ISPs, who in turn impose them on their users

A pestilence in its own right, spam is also the dead canary in the mineshaft sternly warning us that the new communication and control system the world will inevitably come to rely upon for mission-critical tasks is dangerously vulnerable to catastrophe from any seriously talented programmer with a motive for chaos (<http://www.icir.org/vern/papers/cdc-usenix-sec02/index.html>).

As with drunk driving, change will come only when people get mad and decide to act in unison against this eminently preventable menace.

Individual victims, and many Internet Service Providers, now employ incoming filters to stem the flood, but this

Trend in proportion of viruses to total e-mail worldwide 2003-2005. Used by permission of Messagelabs Ltd



saive qui peut measure leaves intact the burden on the network. A step up are utilities like Spamcop (<http://www.spamcop.net>), which actually reports spam to a responsible party, but this palliative fails to prevent spam at the system level.

Why spam happens

How can this pestilence continue to worsen, when other serious social problems are stable or declining? (Think drug abuse, drunk driving.)

Simply because the internet now ignores basic principles of human behavior known to every parent, and universally applied elsewhere in civilized life:

- Everyone is responsible for his actions
- Actions are traceable to their authors
- Actions bring their authors good or ill, according to their impact on others

In short, spam exists because action is divorced from consequences. Fix *that*!

How spam happens

Spammers now employ a variety of advanced upload methods such as open mail relays, insecure web proxies, malformed CGI scripts and zombied clueless-user machines.

However there is virtually no legal way to upload spam in the United States and many other countries, due to contractual bans imposed by backbone providers on their ISPs, who in turn impose them on their users. Uploading spam always entails one or more offenses like tort, Terms of Service fraud, violation of contract, or trespass.

Spam continues because many ISPs fail to enforce these clear and simple rules against their spamming customers, and the backbones do not enforce the rules against the ISPs. Why can't they enforce the contracts?

They *can*, and many do: it is a management decision, driven by money. The providers who do enforce operate ethically; those who don't operate on the Environmental Polluter business model: it's easier to dump the waste in the river than to secure one's factory against pollution.

For the big-time spam-enabling backbones, and their downstream ISPs, abuse desks, with their "thank you for your report" auto-replies, are pacifiers intended to keep the money coming in while placating enraged victims with illusions of action. In fact, the only effective action – cutting off polluting ISPs – is seldom imposed.

Why don't the spam-enablers rigorously follow up complaints? They claim their abuse desks are "overloaded". One shameless ISP even sends this auto-reply to complaints:

Thank you for your message. Your email has been received and will be processed in due course. Due to the overwhelming amount of email received at this address, you may not receive a human response.

A more candid confession of failure to secure one's network is hard to imagine.

If one probes a bit deeper, more sad truths emerge. Spammers open multiple accounts and web pages under false names, spew out their spam, sometimes in fact are shut down, then move to the next ready account or webpage on the same host, and recycle. Abuse desk staffers cheerfully call this whack-a-mole; engineers call it an endless loop.

Lack of identity checking permits this endless loop. When questioned ISP managers reply that they could not possibly earn a profit if they had to secure their networks against abusers.

What is wrong with this picture? It is precisely the Environmental Polluter model: design a business to gather revenue for the stockholders while imposing on outsiders the economic losses to society arising from its polluting operations.

Offending ISPs allege "no one would sign up for an account" if each had to be verified, which of course is true as long as there are race-to-the-bottom providers extending connectivity to any malicious or negligent stranger. If no one could offer service, allowing strangers to injure others, then the current competitive race to the bottom would not exist. (Effective and innocuous measures exist to confirm identity, used by many firms in many economic sectors. A technical solution exists even to preserve anonymity by permitting but rate-limiting such accounts.)

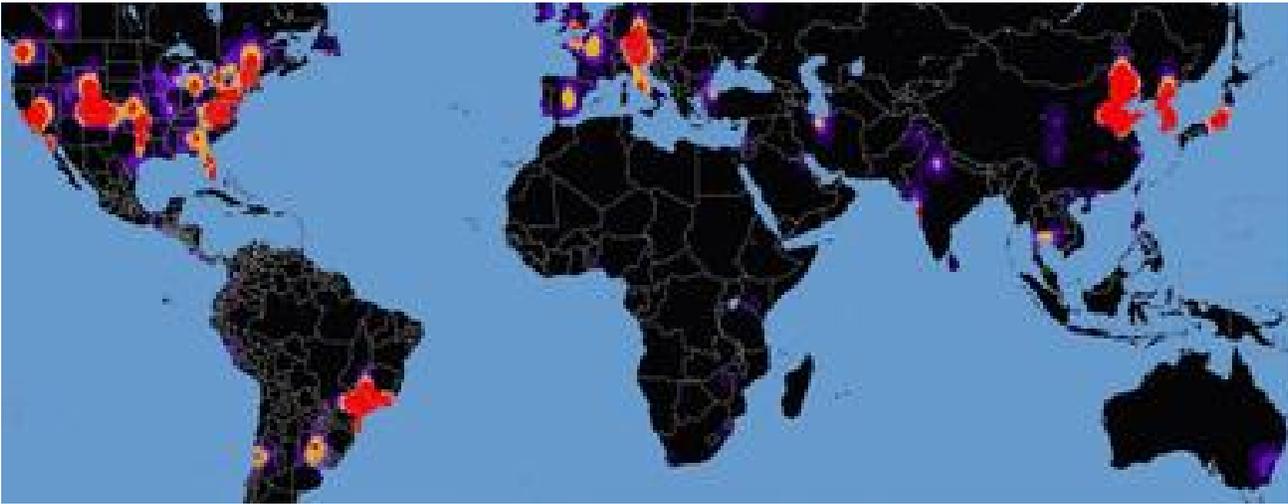
The spammer business model

Spammers have their own business model, aptly summarized as the Thousand Cuts, which meshes with the Environmental Polluter model to victimize the rest of us. Spammers well know the illegality of their businesses but know also that the pain is spread in small amounts among many victims, not one of whom can make an economic case for litigation. Even someone determined to act finds it difficult due to cumbersome legal procedures, the cost of discovering obfuscated identities, and the torpor of the agencies responsible for ensuring accurate databases.

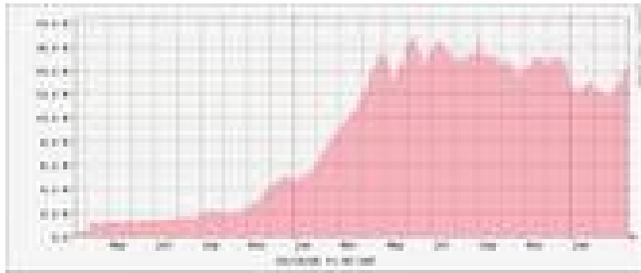
What to do

Spam increases because no ill consequences befall the malefactors and their enablers. As every caring parent knows, this method is guaranteed to raise antisocial offspring.

Major spam-emitting regions worldwide, courtesy of Postini



Crisis in the making? Total spam on the internet 2004-2005.
Source: Distributed Checksum Clearinghouse



What to do? Obviously, smash these two business models. Big, immediate improvements require no legislation and little litigation, just doing the obvious on the internet comparable to what every loving parent does in rearing his children.

The following steps can end spam as a “big issue” for internet users.

First and foremost, ISPs must use blocklists to refuse all incoming mail from insecure or misconfigured mail servers, rather than just filtering incoming spam. This is the only method that works, and it works *immediately*.

Blocklists assemble the Internet Protocol addresses of mail servers known to emit spam. A variety of organizations maintain such lists, using numerous criteria such as whether the manager of the mail server is spammer-friendly, negligent in the operation of his mail relays, running insecure

CGI scripts or proxies, disobedient to the ruling documents of the internet known as RFCs, or complaisant to trafficking over his network in internet burglar tools (spamware).

What happens when an ISP uses blocklists? All mail from spam-enabling ISPs fails to transmit; senders receive a diagnostic message to the effect:

“REJECT=550 Your message is refused since transmitted from a spam-emitting mail server at IP address 203.144.247.97. Contact your system administrator to bring this server into compliance with current best practice.”

Immediately scores or hundreds of customers complain to the *offending* ISP, which is forced to manage its mail servers properly. This differs from the present state of affairs where the victims complain uselessly to their own (victim) ISP.

Even the best ISPs occasionally fall afoul of blocklists, but they cure the problem fast. Indeed it is almost comical how fast spamming stops when a blocklist is used, or even threatened. Connect, one of Australia’s largest ISPs, long harbored a notorious spammer. At 10:00 a.m. one day in January a few years back a group of victim system administrators tired of politely asking Connect to shape up and instead laid down their new zero-tolerance policy: blocking would ensue that day unless Connect shut off its spammers. By 1:30 p.m. they were gone, with no service interruption. Merely the threat of disconnection from the internet caused management to pull up its socks.

A similar incident occurred in Australia when Optus defiantly refused to cut spammers loose. Blocking by an important group of victims forced Optus to change its policy two days later: the magic of “actions have consequences”.

Blocklists keep unsafe ISPs from connecting to the internet just as credit reports exclude defaulting debtors from the credit markets and pre-flight inspections keep unsafe planes from the sky. Some inconvenience may arise until safety and security are assured but it is small, necessary, brief, and falls upon the offender rather than the victim.

Universal adoption of blocklists can be encouraged by customer demand, by pressure from public standards bodies, or even by government. Its effect might at first be to split the internet into zones of purity and islands of pollution. As blockage expands, spammers will be pushed into ever smaller and less connected domains, which grow ever more blocked. This cumulative process would end quickly, with residual polluted areas populated by ISP customers who have little need to communicate with zones of purity.

Second, every upstream provider must verify – by testing – that all downstream customers comply with current best practice. To end controversy about intrusiveness, this measure should be endorsed by standards bodies.

Two critical guardians of internet integrity, ICANN and the Regional Internet Registries, must be inspired to cease behaving like pussycats

Third, two critical guardians of internet integrity, ICANN and the Regional Internet Registries, must be inspired to cease behaving like pussycats. These bodies are charged in different ways with creating the “internet telephone directory” which allows users to find each other. However, the resulting database of identities is corrupted by massive spammer registration fraud, precisely to prevent the victims from finding their tormentors. Many registrars cheerfully make a living from spammers and brazenly refuse to act against their outlaw customers. ICANN is charged with preventing such complicity in fraud and is fully empowered to yank registrar accreditation.

(Supervision over domain name assignment falls to ICANN, the Internet Corporation for Assigned Names and Numbers. Internet Protocol addresses are allocated to users by four

registries worldwide, ARIN – American Registry for Internet Numbers, APNIC – Asia Pacific Network Information Center, RIPE – Reseaux IP Europeen and LACNIC – Latin American and Caribbean Internet Addresses Registry.)

My own unhappy experience with ICANN confirms – to put it charitably – that no one there will be committed for psychiatric care due to obsession with suppressing registration fraud.

The RIR’s charters are less authoritative, itself a serious problem, and the communities in which they are involved are aware of defects in the database. Again, personal contact reveals that the RIR headquarters staff are not anguished over the inability of victims to identify their tormentors; indeed many actively oppose measures to enhance accuracy (“not our problem”).

The failure of these critical organizations to perform their duties should be righted, if necessary by a public spotlight or even legislation.

Fourth, the legal profession in cooperation with anti-spam groups should aggressively attack major spam-enablers, for which numerous legal grounds exist such as public nuisance, attractive nuisance and negligent enablement. It should be possible to recover huge money damages in view of the billions of dollars in annual losses provably resulting from their negligence and from their wilful failure to enforce contractual agreements. ISPs and backbone providers well know torts, or worse felonies, are continually being committed with their property but fail to adopt even the most obvious preventive measures. This arena is a perfect fit for class action litigation.

Fifth are limited prosecutorial or administrative actions. Many spammers are incorporated and it is not hard to see how their charters could be summarily revoked for violating their corporate charter, since spamming necessarily entails violation of civil and often criminal law. A quick and probably uncontested hearing should suffice.

A few exemplary prosecutions *pour encourager les autres* could be mounted on a variety of grounds. One promising criminal approach is “fraud in the inducement”. Violation of the “click to sign up” internet account agreement would ordinarily be construed as a civil dispute, but a chronic pattern of contract violations where the spammer intends to violate *ab initio* is criminal conduct in many jurisdictions.

Finally, some modest legislative action may be in order as a residual cleanup measure. Primary would be removal of ex-

