EU workshop on unsolicited commercial communications or spam – Brussels, 16 October 2003

Françoise Becker, CTO, L-Soft
Eric Thomas, CEO, L-Soft
Outi Tuomaala, VP Marketing Europe, L-Soft

Thank you for putting together this interesting pan-European workshop on spam. It is good to see that the EU takes the spam problem seriously. We are particularly pleased to see Commissioner Erkki Liikanen taking such an active role in the fight against spam, both within and outside the EU (FTC, World Summit on the Information Society, OECD, *etc.*)

Although most spam currently originates outside Europe, we hope that the EU's strong opt-in stance will **champion the adoption of opt-in legislation in the rest of the world**, and in particular in the U.S.

We would like to contribute our thoughts and comments on the spam situation and the different means to fight spam. L-Soft is a proponent of opt-in, to the exclusion of any and all opt-out compromises, and has actively fought spam since 1995. See Company Background at the end of this document for more information.

AWARENESS

Public Education

Funds must be allocated for public education. Education is one of the most crucial steps in curbing the proliferation of spam. Spam flourishes because it works. When people stop buying from spammers, the spammers who are in it for the money will stop sending their spam. The buying public needs to get the message:

> **"Never spend money in response to e-mail from a company with which you have had no previous contact."**

Ad campaigns on television, in the print media, and – yes – in legitimate, opt-in e-mail newsletters, will do more to curb spam than expensive litigation.

Marketing Education

Directive 2002/58/EC and its Article 13 must work as a foundation for the codes of contact for industry and marketers engaging in e-mail marketing. Universities and business schools throughout Europe must take opt-in e-mail marketing practices as a part of the curriculum when training future marketing professionals. Industry, service- and software providers and marketing associations of all types need to educate the market place and customers on opt-in marketing practices. Marketing professionals must get the message:

> **"Opt-in is the only e-mail marketing practice that will not damage your brand and reputation."**

EFFECTIVE APPLICATION OF THE OPT-IN REGIME

Lack of Transition Rules

Many ethical companies have existing newsletters or otherwise communicate with their customers using bulk e-mail. In some cases, these newsletters date as far back as the 80s. It may not always be possible to prove or guarantee that each and every recipient has opted in. In most cases, though, the customers in question do want to keep receiving the material.

What should honest, law-abiding companies do when the national opt-in laws take effect? Do companies have the right, for instance, to send a single, one-time-only mailing informing existing subscribers that they must take action to confirm their subscription or it will be automatically cancelled, or do they really have to terminate all existing newsletters and put up with hundreds of phone calls from confused customers wondering why the newsletters suddenly disappeared?

Transition rules have often been left out, creating a "grey zone" that is uncomfortable for both customers and marketers. Perhaps advisory guidance could be provided at the EU level.

A Do-not-spam Registry

The U.S. Senate has passed the so-called "Can Spam" bill on October 22, 2003. The bill does not require opt-in but prohibits senders of unsolicited commercial e-mail from disguising their identity by using a false return address or a misleading subject line. It also requires unsolicited commercial e-mail messages to include opt-out instructions. The bill directs the FTC to develop a do-not-spam registry.

We are among the many who think that a do-not-spam registry is a bad idea. In the U.S., there has been a lot of discussion about the technological ease or challenge of maintaining such a registry. Unfortunately, maintaining it is the easy part. The difficult part is providing access to legitimate e-mail list operators so that they may use it to remove e-mail addresses from their lists, **without** providing unethical spammers with a guaranteed source of addresses, or at least an easy way to validate the addresses they harvest from the web or generate through dictionary and brute-force attacks. Why connect to individual mail servers to validate e-mail addresses when the FTC provides you with one-stop-shopping?

It has recently been claimed that a do-not-spam registry based on hashing technology would permit the removal of e-mail addresses while preventing spammers from gaining access to the raw addresses. But this is the same kind of technology commonly used to store computer passwords. "Password crackers" – programs that guess or reconstruct a user's password from its hash value – abound on the Internet, and they are quite successful with "weak" passwords. Passwords based on common English words or other predictable strings are generally considered to be weak. Unfortunately, most e-mail addresses contain predictable strings, such as AOL.COM or the user's surname. Addresses such as jsmith@xyz.com are particularly vulnerable.

Unlike passwords, e-mail addresses do not change very often. Spammers could "crack" the do-not-spam register in a country where this is not illegal, and later sell the addresses to other spammers. These addresses would be more valuable than harvested addresses because they would be valid addresses and would presumably

reach a large number of busy professionals and other decision makers, who typically do not put their e-mail address on the web.

Legislation

Opt-in is the only legal solution that can be made to work. When “opt-in” is required, as is now the case within the EU, there is no need for expensive and technologically challenging registries, dubious labeling laws, or blacklists.

But legislation must also be worded so as to discourage frivolous or fraudulent lawsuits, which is a danger that comes with “private right of action.” The real challenge of legislation is to define what constitutes proof that a communication was unsolicited. Due to the nature of the SMTP protocol, even records of a double-opt-in confirmed subscription are trivially easy to fake, and therefore unreliable as proof, and it becomes the word of one individual against another. Since most of the accepted definitions of spam refer to the “bulk” nature of the offending e-mail, the legislation should require a certain critical mass before a lawsuit may be filed. Then it is no longer the word of one individual against another, but that of many individuals, lending it credibility. This would make it easy for ISPs and anti-spam coalitions to file a suit against the real spammers while making it difficult to engineer frivolous or fraudulent lawsuits.

EFFECTIVE REMEDIES AND PENALTIES

Litigation Challenges

If you are a legitimate list operator, using the “golden standard” of “double opt-in”, what information should you maintain for each subscription to protect yourself from frivolous or fraudulent lawsuits?

One of the challenges for a scrupulous litigator ought to be not only to go after the bad guys, but also to make sure you are *only* going after the bad guys. What evidence provided by a defendant would sway a litigator to drop an action?

With laws that allow “private right of action,” what is to stop an individual from signing on to a list and then claiming to be spammed? Even if the company can prove to the satisfaction of the court that the individual did request the subscription, the company will still be required to spend time and money to get to that point. What is to stop an unscrupulous company from engaging people to start such actions against its competitors? The damage from such frivolous or fraudulent lawsuits may be enough to put a small company out of business, even if in the end they win the suit.

COOPERATION WITH THIRD COUNTRIES

Opt-in proselytism

It is a fact of life that most spam originates outside the EU. There can be no solution to the spam problem without some kind of worldwide “minimum standard” of legislation, as spammers have shown themselves to be extremely mobile and always a step ahead of law enforcement.

We firmly believe that this "minimum standard" is opt-in, with an added (and generally uncontroversial) prohibition against deceptive practices such as misrepresentation of the sender's identity.

On November 1st, 2003, the EU will become the single largest market in which this minimum legislative standard is in place. Conversely, the U.S. will remain as the single largest market in which opt-out is the norm, not only in the courtrooms but also in marketing departments. In order to win the fight against spam, we must convince the U.S. and all other major Internet countries to switch over to opt-in. The EU has a unique opportunity to play a decisive role in this matter.

TECHNICAL ISSUES

"Make Spam Cost" proposals

One of the possible solutions mentioned by Commissioner Liikanen during his introductory speech is to somehow cause senders to incur a cost for every e-mail message they send. This cost would still be lower than for postal mail, but high enough to make it economically impossible for the spam industry to remain in existence, as its conversion rate (from spam to purchase) is too low.

Assuming that logistical and enforcement issues could be suitably addressed, this approach would deal a lethal blow to the spam industry. Twenty years from now, the only leftovers of spam would be vague mentions in history books. Appealing as this may be, there is a very real risk of throwing out the baby together with the bathwater. We would like to expand upon the dangers of this approach.

Roughly speaking, there are two main categories of "make spam cost" proposals. The first is based on the transfer of actual funds. An organization, in most cases the sender's ISP, would collect a fee for every e-mail message sent – perhaps on the order of 1/10th of a cent, at the most one cent. This fee could even be waived for senders with very low volume, such as normal consumers. Spammers on the other hand would have to pay thousands of euros per spam, and would quickly go out of business. Honest marketers, on the other hand, would have no trouble paying this fee. It would still be far cheaper than any other form of direct, personalized communication.

There are many problems with this idea, but perhaps the most disquieting is the thought that ISPs would have a financial incentive to harbor spammers and encourage indiscriminate bulk mailing. Bulk e-mail communication needs to go in the direction of lower volume and higher quality, and this would be a counter-productive step. It is also morally unbearable for anyone but the victims to cash in on a "spam fee."

Some people have proposed that the government collect this fee, instituting some kind of "e-mail tax," which admittedly would have the funds going back to the citizens, if only indirectly. But do we really want to open Pandora's box and ask our respective governments to start taxing the Internet? Assuming for a moment that someone *must* necessarily profit from the spam plague in order for it to end, do we want this someone to be the same government that is tasked with enforcing our new anti-spam laws – and closing the budget? Besides, an "e-mail tax," even for the good cause, would probably be found to be discriminatory, as it would only target one particular Internet protocol. In all likelihood, it would have to be replaced with a generic Internet tax, and we would be back to square one.

"Make spam cost" proposals from the second category avoid these issues by renouncing all forms of fund transfers. Instead, e-mail protocols will be changed in such a way that a very powerful computer will be required to send even modest volumes of bulk e-mail. Although it will remain very real, the cost of sending spam will take the form of hardware purchases. Very few spammers would remain in business if they had to make hardware and software investments on the order of several hundred thousand euros *per spammer.*

Taking a step back so as not to miss the forest, it becomes apparent that all these proposals are more or less equivalent. Whether the "spam fee" is collected by ISPs, by the government, or by the hardware and software industry, the basic concept is the same: **an arbitrary cost is manufactured out of thin air in order to overpower the spam industry financially.** There are two fundamental moral flaws with this concept:

1. One particular group (ISPs, government, hardware manufacturers) greatly benefits, in a totally arbitrary manner. There are simply no logical reasons why any of these groups should derive substantial profits from spam and, thereby, be given a financial incentive to foster its continued existence.

2. Spammers are not the only group that will be overpowered. Bulk e-mail communication will become the exclusive province of the financially strong – big companies and the government. Hundreds of thousands of individuals and grassroots organizations will have to give up their opt-in newsletters and discussion lists overnight.

Fundamentally, the creation of an arbitrary cost for the delivery of e-mail, no matter how this cost is introduced, would turn the clock back to the days when only those that could afford a press could make their voices heard. It would be a blow against the very foundation of the Internet and its culture. It would be a huge step back for freedom of speech. Surely, there must be a better way.

Bayesian Filters

Although we believe Bayesian filters to be one of the most promising technologies in the fight against spam, this only holds true if they are used correctly. By their nature, Bayesian filters are unpredictable and almost impossible to understand for "ordinary mortals." The decisions they make are likely to seem arbitrary. Care must be taken to use this technology in the right context, to avoid bad surprises.

For instance, consider an employee who speaks several languages, but communicates almost exclusively in English. The overwhelming majority of non-English messages that this person receives is likely to be spam. Over time, the Bayesian filter will learn that non-English words are sure indicators of spam. One day, the filter will start deleting every message in languages other than English – even messages that do not have any of the telltale signs of spam.

There are many possible solutions to this and the other problems posed by Bayesian filters, but one must be aware of these issues and refrain from the indiscriminate use of this otherwise very interesting technology.

ISP-Level Whitelists

During the workshop, a number of ISPs proposed the creation of "trusted peer whitelists." A worldwide network of trusted ISPs with a clear anti-spam stance would be created, within which no spam filtering would be necessary, whereas very aggressive filtering would be applied to e-mail coming from outside the trusted group.

Used correctly, this can be one step along the road to a spam-free Internet, but we do want to point out a flaw in one of the underlying assumptions, namely that consumers are always free to switch to one of the "Good Guy" ISPs because there is so much competition in the industry. This is not true of two significant minority groups:

1. People living in sparsely populated areas, where competition often ranges from very limited to nonexistent. The paucity of employment opportunities typical of these areas makes the Internet very attractive to individuals who are able to sell their services remotely (web design, accounting, *etc.*) The ability to send and receive e-mail successfully can be the difference between self-sufficiency and unemployment.
2. Broadband users who have signed a binding long-term agreement with a particular ISP in exchange for the installation of a LAN or other distribution system within the apartment complex. In Sweden, these contracts have been known to have terms of up to 25 years. While consumers are obviously free to revert to dial-up connections and pay by the minute for slow Internet access, this is hardly a realistic alternative.

Although there is nothing wrong with applying more aggressive filtering outside of a network of trusted ISPs, the issue is how much more aggressive one can reasonably get, given that consumers are not always able to switch ISPs.

Signature-Based Detection

We think that signature-based spam detection (similar to the methods used by anti-virus software) receive too little emphasis in the spam debate. There are admittedly a number of issues with this technology; in particular, it is costly to implement, and will only block spam once identified and entered into a signature database. On the other hand, this approach offers the immense advantage of **near-zero false positive rate.** A spam identified by a properly designed signature (such as the presence in the message of a specific URL or toll-free phone number belonging to the spammer) is in principle guaranteed to be a spam, and can be deleted safely. The costs of implementing and deploying this technology, while substantial, are insignificant when divided by the total number of spam victims. The technology itself is proven and well understood, and protects us from thousands of malicious viruses today.

Labeling Requirement – ADV

Although not required by the Directive, labeling requirements have been discussed frequently, and Member States would be free to pass such legislation in addition to what is mandated by the Directive. We think labeling laws would be a bad idea because:

- Labeling all commercial e-mail with "ADV" or the like does not help recipients distinguish between commercial e-mail that they have opted in to and do want to receive, and unwanted, unsolicited commercial messages. They still have to go through every message in their in-boxes or, more likely, filter it all out

and miss the coupon from their favorite bookstore that they were looking forward to.

- Spam is in the eye of the beholder. What one person sees as business communication may be seen by another as advertisement. If a software company sends e-mail to its customers about the new features in the latest version of the software, is that an advertisement or useful information?

- The definition of what qualifies is ambiguous. Should a newsletter that accepts advertising (example: ZDNet Tech Update Today) have an ADV label? If not, what percentage of the message must be non-commercial? Whatever percentage it is, the spammers will find enough filler text to qualify.

Blacklists

Although they are very popular, there are many problems with blacklists:

- They put the decision in the wrong hands. Corporate e-mail administrators may indeed use such lists if company management has made a *business decision* to accept the risks of lost e-mail. However, in the case of ISPs, the consumers should have a choice about whether to use the lists for their personal e-mail accounts. Generally, ISPs use the lists in a blanket fashion, blocking mail for all their customers, often without even informing their customers that they are doing so.

- The negative impact of collateral damage and false positives is not sufficiently emphasized. The assumption behind the casual acceptance of "collateral damage" is that e-mail is not important, and therefore losing a few legitimate messages for the sake of catching spam is acceptable.

  One participating ISP expressed concern during the workshop about the risk of being sued by a customer who had been unable to read a vitally important e-mail message on time, on account of the large amount of spam in his mailbox. While this is a legitimate concern, litigation is much more likely if the vital message in question was erroneously identified as spam and filtered out.

  As an illustration, we use e-mail frequently for financial transactions. A 50-page contract travels much faster as an encrypted e-mail attachment than as a fax. Such documents are frequently sent or received from hotel rooms, where the choice of ISP (and availability of fax machines in working order!) may be limited. In some cases, the use of the ISP's "smart SMTP" is mandatory. Annoying as it may be, spam only takes a few minutes to delete.

- Blacklists punish the victims more than the actual offenders. Open relay lists do not list the IP addresses of spammers, but of sites whose resources have been abused by spammers, and in some cases of sites that have never been abused by spammers, but simply have the *potential* to be abused. The more ethical blacklist providers give the owner of the IP addresses the opportunity to correct their open relay before listing them. The less ethical ones will not only immediately list the IP address that has the open relay but *every* IP address owned by the same organization.

- There is often little or no accountability. If your IP addresses are listed incorrectly, you cannot redress the problem by sending e-mail to complain

because you are being blocked, and the sites typically do not list any other contact information. The less ethical blacklist providers are often guilty of the same hiding tactics as the spammers they revile.

- There is no standard. Blacklist services are an immature industry dominated by a handful of individuals operating under their personal views about the right way to behave. Therefore there is a great variation in the level of professionalism and thoroughness they bring to the service. Because of the lack of accountability, there is often little that the blacklist's consumer can do to determine the quality of the blacklists or whether the philosophy behind the listing is concordant with the consumer's organizational goals. *Caveat emptor.*

- The blacklist maintainers make the assumption that the individuals putting the blacklists to use are skilled mail administrators. As Mark Burgess wrote in *Principles of System Administration*, "Because the number of local networks has outgrown the number of experienced technicians, there are many administrators who are not skilled in the systems they manage." Even experienced system administrators are not necessarily knowledgeable about the intricacies of mail administration. The blacklists typically do not have documentation that is intelligible to novice mail administrators.

Mailbox-Level Whitelists

"Whitelists" are individually maintained lists of addresses from which you want to receive mail, to the exclusion of all others. These lists are very good tools for home users who only want to use e-mail to correspond with family and friends. Unfortunately, they are not helpful for people who want to subscribe to newsletters and discussion lists, participate in Internet communities, or simply conduct business on the Internet, so they are not sufficient to stop spam.

ISPs should continue to offer whitelist capabilities to their customers, but would do well to move the blacklist capabilities away from the mail server and put them in the hands of the individual recipients: let each recipient decide whether they want their own mail processed through the blacklists, and clearly explain the consequences of either choice.

Company Background

L-Soft specializes in the development of software and services for professional e-mail communication management. L-Soft pioneered the e-mail communication industry with its flagship product LISTSERV®. Introduced in 1986, LISTSERV® was the first and most widely used software for e-mail list management. Since its foundation in 1994, L-Soft has expanded its portfolio of products and services to include e-mail delivery, outsourcing and consulting services.

With offices in the U.S. and Europe, the company serves more than 3,000 customers across the globe, including Ahorro, AOL, Aventis, British Council, BskyB, Check Point Software Technologies, ETSI, European Agency for Safety and Health at Work, Finnish Communication Regulatory Authority, Fraunhofer, FUNET, Katholieke Universiteit Leuven, Swisscom, Telefonica, The Confederation of Swedish Enterprise, The Federal Trade Commission, The New York Times, The United Nations, The United States Senate, The Wall Street Journal, University of Tampere, Uniway, Virtuology and WHO.

Worldwide, public LISTSERV® servers send more than 30 million messages a day to over 100 million list subscriptions.

L-Soft's official policy on spam can be found at http://www.lsoft.com/spamorama.html

The first incarnation of the LISTSERV® spam filter was developed in 1995 by its CEO and founder, Eric Thomas, and has been continuously improved since then. L-Soft has also hosted the SPAM-L spam prevention discussion list since 1995.