How To Detect Forged Email

A major drawback with email on the Internet is that it's relatively easy to send forged or faked email messages to someone. Done as a joke, it's no more than an irritation. But there's scope for much more danger. It's worth explaining to users how to judge the authenticity of a message.

By Wendy Grossman

There are a number of techniques for forging email, many of them very simple. Anyone using the built-in mail facilities in a Web browser such as Netscape or Internet Explorer, for example, can simply type in any return address and nothing in the browser's design, the Internet connection, or the mail server will complain about it. Most email software, similarly, allows users to edit their return addresses. Other methods for faking email are technically more demanding, such as making the message look like it came from a third party's server by telnetting to an open SMTP port and using it to inject the email into the information stream.

The key to understanding the origin of any particular email message lies in the message headers. These are the lines of somewhat arcane information that precede the actual message text. Figure 1 shows a typical header.

Note that not all email packages display headers. Some hide the information, while some strip it. If you don't see a collection of information like this at the head of Internet email, you will need to either find the setting that lets your email software display the headers, switch to a piece of software that will display them, or save the email message to a text file and inspect it in a text editor.

Paths

The most important components of the header, at least as far as forgery detection is concerned, are the lines that show the path the message took from sender to recipient, along with the first line, which identifies the sender's email address, and the line bearing the message ID, a unique identifier issued when the message is sent.

In Figure 1, everything matches: the return address is an ID on CompuServe; compuserve.com appears in the path - the listing of services through which the message has passed - and compuserve.com also appears in the message ID.

In Figure 2, however, although the From, Message ID, and Sender lines all match (mail0117@internet-mail.com), the Received line starts with a different domain (mail.ivbp.com) and the reply address is different again (moneyopp@answerme.com). That ends the search right there, since answerme.com is a known domain of Cyber Promotions, the Net's leading source of spam.

One other quick way of telling if you've got a forgery: look at the times stamped in the Received lines for either "-0600 (EST)" or "-0700 (EDT)". There's a widely used spamming program that makes this particular mistake. The times should read "-0500 (EST)" and "-0400 (EDT)" (measuring from GMT).

Embedded Information

Your can extract several useful pieces of information from the "Received" line to help identify the message's true origins: the name the sender used when connecting to the SMTP server (in the second header shown, mail.ivbp.com), the IP address of the incoming SMTP connection (208.10.58.108), and the name of the host that added the Receive line (tom.compulink.co.uk).

Sender's Address

You need to do a little work to extract a fourth useful piece of information - the true name of the IP address in that list.

To get this information, you need to perform an operation known as a reverse DNS (for domain name server) look-up on the IP number. (Don't bother looking up any IP number in which one of the dotted clumps is a number greater than 255 or any number beginning with a leading 0 - these are invalid.)

If the Message ID ends in an IP address instead of a domain name, you should look this one up, too.

If nslookup is available on your system, use it and/or traceroute to get the domain name and see the route packets take to it. If you don't have these installed, you may be able to add them to your system. WSPING32 for Windows 95 has traceroute and DNS lookups built into it, and the program Mac TCP Watcher has DNS lookup and a traceroute function. These can be found at Windows shareware sites like Tucows (http://www.tucows.com) and Stroud (http://www.cwis.com/).

The easiest place to look up any IP number worldwide is http://cello.cs.uiuc.edu/cgi-bin/slamm/ip2name. Otherwise, the official registry for the .net, .com, and .org domains is InterNIC. You can access its database on the Web at http://ds0.internic.net/wp/whois.html; enter the domain names into the form provided. Alternatively, telnet to internic.net and type "whois <domain>" or "whois <IP number>" at the main prompt, or email mailserv@ds2.internic.net, placing "whois

<domain name>" in the message body.

European addresses can be looked up at http://www.ripe.net, and both UK and US domains can be looked up at http://www.netnames.com. You may get some extra information out of a utility called Dig; there's a Web-based implementation at http://sh1.-ro.com/~mprevost/netu
tils/netutils.- html.

Look for mismatches. If, for example, reverse DNS lookup returns a different host name for the IP number than appears next to it in the header, you may have the true domain name of the faker who sent the message. But look closely: for example, one spam message seemed to have come from either snowhill.com or att.net. But the att.net line read san-francisco-1.ca.dial-access.att.net, exonerating snowhill.com, which InterNIC reveals to be an ISP in Alabama.

You will find sometimes that you don't get back any information from an IP number, and if you try being clever and chopping off the last of the dotted clumps of numbers that you get back the name of one of the major Internet providers, such as Sprint or UUNet. In such a case, those major ISPs will probably be the upstream service provider, not the originating system. However, if all else fails you might get some help from those systems - given a copy of the header, the system administrator may be able to identify the originating domain for you.

Compare

One other aid is comparing messages with other people. Sometimes you need specialised knowledge; for example, one message that apparently came from Netcom (a respectable ISP) was queried in a discussion group because the Authenticated sender was listed in the header as <user>-@popd.netcruiser. A Netcom user, however, knew that Netcruiser is the name of Netcom's front-end software package. Comparing messages with other people also lets you look for a machine common to all of them - this should be the point where the message was injected into the information stream, and therefore the originating domain. In some cases, you may need to ask the system administrators along the path your email abuse travelled to help you identify the source of email abuse by inspecting their logs.

The best-known and most easily accessible online areas for getting help with faked email are the news.admin.-net-abuse.email and news.admin.net-abuse.misc newsgroups on Usenet. The people on those newsgroups can help with deciphering the more difficult headers, and know which services are regular offenders and how to handle them.

Further Action

Quite what you should do when you do detect forged email will obviously depend on the nature of the forgery. If it's clearly been done as a joke, standard discipline procedures can be used if it was done by a staff member, or the operators of the originating system can be informed if it was done by someone else. But if you suspect something more elaborate is going on, such as the beginnings of an attempted fraud, consider contacting an IT security consultant (and even the Police) before sending a "we know what you're up to" reply.

```
From 70007.5537@compuserve.com Fri Sep 26 18:31:47 1997
Received: from arl-img-10.compuserve.com (arl-img-10.compuserve.-
com [149.174.217.140]) by tom.compulink.co.uk (8.8.4/8.6.9) with
ESMTP id SAA17730 for <wendyg@cix.compulink.co.uk>; Fri, 26 Sep
1997 18:31:47 +0100 (BST)
Received: (from mailgate@localhost) by arl-img-10.compuserve.com
(8.8.6/8.8.6/2.5) id NAA19939 for wendyg@cix.compulink.co.uk;
Fri, 26 Sep 1997 13:30:41 -0400 (EDT)
Date: Fri, 26 Sep 1997 13:26:59 -0400
From: "Wendy Grossman (skeptic)" <70007.5537@compuserve.com>
Message-ID: <199709261330_MC2-21E5-CA72@compuserve.com>
```

Figure 1 - A typical email header.

```
From mail0117@internet-mail.com Tue Nov 12 16:32:55 1996
Received: from mail.ivbp.com (mail.ivbp.com [208.10.58.108]) by
tom.compulink.co.uk (8.6.9/8.6.9) with SMTP id QAA27109 for
<xxxxxx@cix.compulink.co.uk>; Tue, 12 Nov 1996 16:32:55 GMT
Received: from mail.ivbp.com by mail.ivbp.com (NTMail 3.02.10)
with ESMTP id ra368437 for <xxxxxx@cix.compulink.co.uk>; Tue, 12
Nov 1996 06:02:23 -0500
Comments: Authenticated sender is <mail0116@mail.internetmail.-
com>
From: "moneyopp@answerme.com" <mail0117@internet-mail.com>
To: mail0117@internet-mail.com
Reply-to: moneyopp@answerme.com
Message-Id: <10544632948659@internet-mail.com>
```

Figure 2 - A header from a faked message.

PCSA

The Author

Wendy Grossman (wendyg@-cix.co.uk) is a freelance IT journalist and author with a special interest in the Internet.