Issues and a case study in bulk email forensics

Dr. Frederick B. Cohen, Ph.D. (Dr.Cohen at Mac.Com)
Fred Cohen & Associates and California Sciences Institute
572 Leona Drive
Livermore, CA 94550

Abstract:

Recent legal matters involving unsolicited commercial email increasingly involve hundreds of thousands of email messages or more. As the volume of emails involved in these cases increases, manual methods for examination and interpretation of evidence become harder, more expensive, and more error prone. In addition, these cases increasingly show evidence of spoliation, and in some cases, of intentional evidence construction that is harder to detect as the actors become more sophisticated. The solution we propose and demonstrate comes in the form of improved automated techniques for analysis combined with more useful presentation to aid in interpretation.

Keywords: electronic mail, digital forensics, spoliation, volume forensics processing

Background:

This paper is about bulk forensic examination and interpretation of email collections as evidence in legal matters. The need for bulk methods has become particularly important because of the high volume of emails involved in many sorts of cases, however; the particular focus of this paper relates to high-volume unsolicited commercial email cases in which one party accuses the other of large numbers of violations of Federal or State laws regarding sending unsolicited emails.

Current laws typically include statutory damages on the order of $1,000 per email in cases where emails are found to be fraudulent. As a result, potential plaintiffs are tempted to acquire and/or produce large volumes of emails and file suits for tens or hundreds of millions of dollars, configure their environments so as to involve multiple states in the transmission of emails, and to accept as many such emails as possible. This then triggers additional damages on a per state basis, and produces enormous collections of emails asserting a wide range of claims for each. The plaintiffs in many of these cases are working together in a loose knit group to share information on how to prosecute these matters, seek and often get settlements for a hundred thousand dollars or more, and use the leverage of high volumes to make the potential risk of litigation very high while driving up the cost of defense. The plaintiffs commonly acknowledge these techniques and assert, in some cases, that they are activists seeking to make bulk emailers pay a high price for sending their emails.

Defendants in these cases range across a wide variety of companies. Some of them appear to be criminal enterprises that violate contracts with multiple marketing firms, lease email platforms from criminal breaking and entry groups to send their high volumes of emails, regularly violate Federal and state laws of various sorts, steal credit card and other related information used in transactions they facilitate, and when sued, shut down and relocate, in at least one case, to Argentina. Others of them are longstanding advertising firms that insert advertisements in newspapers, magazines, radio, television, Web sites, emails, and elsewhere as part of their business. They seek to, and almost without exception, follow the laws regarding advertising in all of its forms, including those related to unsolicited commercial email.

From a technical standpoint, the operation of bulk email solicitations involves sets of companies that have specialties in different facets of the business, ranging from those who create and provide advertising copy and images for their clients and place these on Web servers, to those who send emails to large lists of recipients that they maintain in databases that allow the selection of particular categories of recipients, to those who handle orders and other actions related to fulfillment, to those who process credit cards and other financial instruments. These companies often subcontract and cross-contract with each other so that, in the aggregate, a thriving and competitive market exists in which the participants have intellectual property in different forms and of different types and different arrangements with different customers and vendors in order to execute on their various promises. There are often exclusive arrangements so that an advertisement will only generate leads to the originator, and there are many cases where competitors use other company resources, such as image servers, without permission or collect contractually exclusive leads from an inserted advertisement and resell those leads to multiple customers.

The case study in this paper is a matter in which Plaintiff asserted that 12,576 emails were sent by Defendant to Plaintiff in violation of particular statutes, and damages requested were about $16M. [1] In this paper techniques and results associated with other similar matters will also be included without distinguishing them from the matter identified herein. Some of these matters are ongoing and more specific details cannot be released at this time, however; the techniques could equally apply to this matter if it had not yet been concluded. This case was ruled in favor of Defendant. The analysis herein, while covering both sides of the matter, ultimately represents Defendant's perspective more than Plaintiff's.

Some of the forensics challenges

From a forensics standpoint, the overall business operations situation greatly complicates both the plaintiff's and the defendant's expert efforts. While the plaintiffs in these cases tend to sue all defendants that could be involved, it is often hard to determine which defendant does what, who is responsible for what actions, and so forth. The law in this area tends to support the notion that all potential defendants may be liable, assuming all acted in concert and were not in violation of each others' contracts. Differentiating what came from where, whether images used were actually part of a particular collection, whether one company was simply using a competitors image server or the emails were in fact from the company whose image server was used, associating multiple emails with the same sources when they come from many different addresses and have differing content, and other similar challenges can be daunting.

Even such simple matters as the association of domain ownership, domain names, and IP addresses is often complicated by the large numbers of domains, addresses, and content used, the high rate of change of this information over time, and the lack of timely lookup of relevant information commonly encountered in such matters.

The participants are also not particularly cooperative toward each other, obfuscate whenever feasible, sometimes refuse to answer questions, don't provide documents upon request, and don't retain adequate records or destroy records.

The large volumes involved make detailed examination of each of the emails by a person is too expensive and time consuming for the participants or the legal calendar to sustain. For example, it is common for a few CD-ROMs of new evidence to be proffered within a few days of a deadline to present an expert report on the evidence in the context of the case, or a day or two before a deposition involving the individual who knows about the content of the evidence.

Evidence also commonly includes content that, on inspection, leads to additional sources of evidence that have to be identified and sought. From a tactical standpoint, this sort of evidence is sometimes provided in an obscure form and as a small part of a large collection of other content, perhaps as a scanned printout of an extraction of a log file included within a larger collection of tens of thousands of pages of other material.

These and other related challenges lead to the need for tools that can automate many aspects of analysis while supporting interpretation by the expert in a timely and accurate fashion as well as the need to reapply and modify the use of these tools as new information appears.

Tools and techniques used:

The most common sorts of tools used in the analysis of cases such as these are small programs written in combinations of Perl scripts, shell scripts, and Unix commands such as 'grep', 'awk', and so forth.

Problems with common tools hastily applied

The use of these tools is a tradeoff in several ways. As a fundamental challenge, writing or modifying scripts on short notice leads to difficulty in verifying their proper operation. As an example, off-by-one errors and a wide variety of misses and makes are commonplace in such scripts. As an example of such a problem, suppose that the current directory contains a set of files corresponding to what is purported to be one email per file. The goal of the script is to find the number of files containing some critical content element. Here is an example erroneous script:

```
grep "critical content element" * | wc
```

The problem with this script is that multiple instances of the string contained within one file and on different lines will cause a miscount of the number of files containing the content, while the occurrence of more than one instance on a single line will cause an undercount of the number of instances in the whole collection. When there are more than 10,000 emails, a count of 7,543 that is off by one is hardly critical in terms of making a substantial difference in the typical case, unless the individual emails left out are unique in some manner. Nevertheless, offering the wrong count will produce a challenge from the other side that may degrade or even eliminate the expert as a factor in the matter.

Several approaches are available to deal with these sorts of errors. The most important step to take is to clearly define the objective of the analytical process and to properly report the result.

Issues of definition relative to the legal matter

In the matter in question, one of the items at issue was the number of emails applicable to the matter. Plaintiff asserted 12,576 "emails", while the evidence provided contained only 1,421 "actual emails", where actual emails are portions of the proffered mailbox corresponding to what a user would normally see when viewing the mailbox (i.e., sequences of bytes of the proper format headed by a "From " line). The legal definition of an "email" in this case counted each "actual email" once for each recipient, leading to multiple counts of the same "actual email". But even this definition didn't clarify how the 1,421 actual emails were translated into 12,576 emails by Plaintiff.

Date and time issues in emails

Another issue in the matter at hand was the dates over which the suit was defined. Because of statutes of limitations, effective dates of laws and dates defined in legal filings, dates and times of events are potentially important to such matters. Furthermore, date and time stamps on emails depend on date and time stamps provided by sending computers and the mechanisms they use. As a result, the content of the emails may not be trustworthy, the date and time stamps on computers may be different from the real dates and times in the world, and as emails pass from place to place, time passes. An email sent before a deadline may also arrive after it.

Because of the lack of definitive information on timings in the case at hand, an anchor email approach was taken to rehabilitating the dates and times associated with the header information contained within the emails. This approach leveraged the fact that Plaintiff's emails were handled by the vendor Postini, which put its own time and date stamps on the emails as they passed through it. While proof that the collection of emails were themselves not complete forgeries was not offered, assuming that they were not, led to the use of the Postini date and time stamps as anchors by which to judge the dates and times of the remainder of the processing of the emails. To independently validate the Postini date and time stamps, independent emails sent between systems that had known time and date characteristics, that were contemporaneously sent through the same Postini servers, and that were under the control of the experts in the case over the period in question were used to "anchor" the times of the emails in the matter. This rehabilitation of the date and time information allowed all but 242 of the actual emails to be excluded as not within the date and time frames at issue in the case.

There are clearly other date and time issues in emails, particularly in these sorts of cases. As an example, one of the bases for legal claims stems from damages that in turn may include losses caused by the reduction in available bandwidth, storage, and other related costs of handling the emails. The evidence used to show that these emails caused such delays has to be in some tangible form, and unless very detailed records are kept on the systems affected, this is hard to show. One of the ways plaintiffs try to use the evidence to make such a showing is by demonstrating the delivery times indicated in the Received headers within emails. These headers show the arrival times at servers. The analytical problem is to show a correlation between these times and the volumes of the emails so as to demonstrate some effect.

The approach to performing analysis of the Received headers with respect to time and date stamps is complicated by the use of multiple time zones and time differentials between computer date and time settings. The analytical approach used in related cases was to restate all dates and times in Universal Coordinated Time (UTC) and examine time differences from "hop to hop", where each "hop" corresponds to the arrival time stamp of a computer in the processing sequence. This involves parsing all of the received headers, normalizing times to UTC, determining distance from final arrival point for each Received header (in hops), correlating paths through the email system so that comparable paths are compared to each other and not to other paths, identifying time differentials by hops for common paths as a function of time, and relating these time differences to email volumes. This can result in some rather strange outcomes. For example, in one case, arrival dates of some emails traveling particular paths were delayed by days while other emails of similar size and content and coming through the same paths, arrived later and were delivered earlier, typically within seconds. No crashes or other similar events were identified in the same time frames to explain these anomalies, and they remain unexplained today. In some cases inverse relationships between volumes of emails in evidence and delivery times has been observed, leading to the potential conclusion that these emails actually helped improve performance, an apparently ridiculous conclusion that would seem to demonstrate the power of statistics. Correlation is not causality.

Table 1 is an example of a depiction of email arrival times and delays associated with a sequence of emails. Each line represents a different email, and the arrival dates and times are sequenced in order. The time delay is shown between first arrival at the plaintiff's servers and its final delivery internally. Note that the emails with delivery times in excess of one day (all those with a non-zero value in the YYYY-MM-DD field of the Delay column) arrived far before and were delivered after those which were processed and delivered within seconds of arrival. The emails in the particular matter were unexceptional in their differences in size, makeup, and content. This sort of analysis and presentation tends to refute claims that emails were delayed because of the high volume arriving.

Arrival	Delay
06/27/02 07:33 AM	+0000-00-00 00:00:02
06/27/02 07:53 AM	+0000-00-00 00:00:06
06/27/02 09:11 AM	+0000-00-00 00:00:04
06/27/02 11:55 AM	-0000-00-00 00:00:03
06/27/02 02:41 PM	+0000-00-01 21:24:25
06/27/02 06:23 PM	+0000-00-01 13:06:42
06/27/02 08:12 PM	+0000-00-01 20:16:02
06/27/02 08:24 PM	+0000-00-01 13:09:01
06/27/02 09:12 PM	+0000-00-02 01:12:32

Table 1 - Extracted email arrival and delay times.

Table 2 shows a depiction of the number of arrivals of emails at different hops within plaintiff's infrastructure on different days. While some variations associated with emails that arrive just before midnight and are delivered early the next morning is to be expected, and emails may arrive at different distances from their final destinations, in this case, none of the arrivals fit this pattern. All of these emails went through at least 3 hops in their delivery process, and the date and times on the computers appears to have been consistently within a few seconds. In fact, detailed examination shows that some of the excess emails were from long delays, while others were from duplicates generated by plaintiff.

Date	Final receipt (1)	Hop 2	Hop 3	Hop 4	Hop 5
10/01/03	4	4	3	2	0
10/02/03	9	9	9	9	0
10/03/03	8	8	8	8	0
10/04/03	6	6	6	6	0
10/05/03	11	11	10	10	0
10/06/03	11	9	8	7	0
10/07/03	23	20	19	18	1
10/08/03	11	11	11	11	0
10/09/03	12	9	6	6	0

Table 2 - Number of emails arriving at different hops by date

Ability of emails to be delivered

Emails asserted in various legal settings must be "deliverable" in the sense that there must, for example, be a user that can actually receive these emails. Some plaintiff's configure their systems to accept any and all SMTP protocol sequences that arrive, thus receiving possibly misdirected emails, emails to users that don't exist, emails to users who have cancelled their accounts, and so forth. This is problematic for a variety of reasons, including the potential that this constitutes interception of private communications, which may be illegal in some jurisdictions or in violation of the plaintiff's policies or contracts, and the common interpretation of courts that such actions invite the emails and thus cannot be the basis for claims associated with undesired email transmission.

The normal simple mail transfer protocol (SMTP) operation, if the relevant RFCs are followed, is to refuse email to recipients that do not exist, and to do so without allowing the receiving server to enter the state where it can receive the data portion of the email (the body). Thus any receipt for non-existing users may constitute an invitation. In the matter identified above, there were 133 invited actual emails that could otherwise not have been delivered, leaving only 109 actual emails to be considered. In other cases, tens of thousands of emails have fallen into this category, and in some cases, courts have ruled that plaintiff's undertaking these sorts of actions to create legal actions constituted a business activity designed to generate law suits, and under applicable laws, caused rulings that stated, in essence, that such suits were invalid and not the intent of the applicable laws.

From a forensics standpoint, demonstrating this involves getting discovery on user identities and their associated human beings and correlating those user identities and individual people to the emails. This is done by a variety of methods, ranging from lists of user identities in password files to server logs and configuration files associated with remote access servers. In the case cited above, this was done by examining RAS server logs and password files along with lists of user identities provided by Plaintiff under court order. These then have to be matched to user identities extracted from the evidence associated with emails to determine which emails were sent at a time and date when the user was or was not active within the system. Since many of the relevant records are not available in digital form, financial or other related records associated with the purchasing of services may have to be used to correlate this information properly to the date and time stamps of emails.

Detection of duplicates and other similar processing errors

Another common finding in such cases is the existence of duplicates within the collections of email proffered as evidence. These duplicates may arise from any number of causes, and identifying the causes is a necessity for the plaintiff trying to assert the authenticity of their records and a benefit to the defendant in asserting that the records are spoliated. In the case identified above, 11 of the actual emails were duplicates, somehow produced by Plaintiff's processing of the emails. In other cases, the numbers of duplicated emails have been far higher.

Duplicates in these cases may appear in many forms. The most obvious form is simply an exact copy of an entire email including headers and separators. In such a case, a byte for byte comparison of the emails yields an exact match, including delivery information, locally assigned message identifiers associated with the reception of messages, Message-ID fields, dates and times of each reception, and of course the rest of the headers and the content. These are easily detected by the use of a cryptographic checksum or other hash of each entire email and a sorting of the results. Matches are then immediately obvious by repetition of the hash and can be identified and properly dealt with and/or confirmed by byte-by-byte comparisons.

In other cases, only parts of emails are identical. These sorts of duplicates are far more problematic for the plaintiff because they indicate that the evidence is not legitimate and may be a complete fabrication, or at a minimum, cannot be asserted to be an accurate depiction of the sequence of events that took place. Some examples of observed matches indicative of spoliation include:

- Emails that are identical in every way except for the separator date and time stamp placed in mailbox files by servers, which indicates that the server recorded the email twice, even though it only received it once, or that the collection of emails was fabricated. This is clearly not the normal operation of an email reception system, as can be verified by the experimental use of the specific system in a reconstruction.
- Emails that are identical in every way except that they were received more than once by one of the servers in the Received path, which indicates not that the same email was resent part of the way twice, but rather that the evidence was somehow fabricated, because the internal SMTP message identities of prior and subsequent receptions are identical and have identical time and date stamps, even though an intermediate server indicates multiple receptions and retransmissions.
- Emails that are identical in every way except that they have different date and time stamps on otherwise identical Received headers, including the supposedly unique identifying message identifiers provided by the SMTP servers.
- Emails with identical separators and yet different remainders of the headers, indicating that the separator was somehow artificially placed between emails that arrived separately.
- Emails with different content but identical headers.
- Emails containing indicators of cut and paste operations used by Web browsers supposedly sent by automated email mechanisms.
- Emails containing content indicative of being processed by plaintiff processing mechanisms such as the systematic addition of nearly identical content to emails with different formats from different sources.

- Emails with various identifiers in headers showing different sourcing than other emails supposedly from identical sources.

These and many other similar indicators of spoliation have been seen in emails involved in actual cases. However; detection of such things is problematic in large volumes of email because automation is not typically designed to detect such things and interpretation is required in order to make determinations about their legitimacy. The general approach undertaken for such detections is the creation of matching software that does imperfect matching of portions of evidence. While matching one item to a large set of other items is straightforward, a match of each of n items to each other item requires n^2/2 time. Other methods allow exchanges of space for time. Using hashes turns this into a linear operation for hash generation followed by a sort, which is O(log n), leading to O(n*log n) time. This potentially has to be done for each of a large number of different sorts of matches. For example, if removal of each line in the headers is to be considered, this comes to the average length of headers multiplied by the previous time. If altered header lines is to be considered, then this has to be broken down further.

At some point, it becomes worthwhile to look at these files in terms of sequences of words or other symbols of lengths starting at length 1 and increasing up to some maximum sequence length. Each sequence can then be given a unique number and all components (emails, headers, or other subsequences of sets of items of evidence) with equal numbers are determined to match to the specified level of similarity. The analyst then needs to interpret the meaning of these matches. Unfortunately, this approach leads to very large numbers of different matches, and the analyst must again find a way to explore only subsets of these matches in order to keep time and costs down.

One of the most effective approaches is to use human vision to do comparisons of rapid sequences of similar components, finding similar components that can then be compared in more detail, and excluding obvious mismatches.

Grouping extracts for comparative analysis

Another approach to detecting anomalies in analysis and interpretation is to create an error model and look for identified error types. For known error types, it is advisable to program detection capabilities so that the known errors are automatically detected. Otherwise they may be missed.

One class of error types and method of grouping email extracts is to look at the limited structure provided in headers. While header lines are largely unstructured, they normally begin with a sequence of characters followed by a colon (':') and continue on lines that start with whitespace. A simple parser can add the lines starting with whitespace to the previous lines starting with a header identifier to allow line-by-line parsing, which makes utilities like 'grep' and other aspects of parsing and analysis easier to program in less time.

Extraction of all header lines from all email extracts and separation by extract number (or a Bates number if they are treated as separate exhibits), line number within the pre-parsed header, and header identifier allows a wide range of analytical techniques to be applied with relative ease. Using disjunctions, conjunctions, and other similar logical operations allows analysis such as the detection of all emails with no "cc:" field and containing a particular domain name, relatively easy. The analyst can then use these operations to perform different sorts of analysis on emails to find similarities and differences indicative of situations relevant to the matter at hand. For example, in the matter above, it was important to identify emails containing particular IP address ranges in the Received headers as recorded by Plaintiff's computers. Extracting this data is non-trivial because of the non-standard format of Received lines, however; once a parser is modified to work for the particular header lines of the mail transfer agent (MTA) software in use, the IP addresses can be associated with extract numbers, lines in the extracts, and the "by" portion of some Received headers to include only the desired extracts. Claims regarding these extracts can then be analyzed by customizing other analysis program snippets for the specific claims.

Even in cases involving more than 100,000 emails, the separation of email extracts by headers and analysis of each header can be done in only a day or two of effort. This tends to yield a great deal of information about emails. Simple sorts of headers rapidly yields information about similarities and differences. Some of the types of information detected includes, without limit:

- Detection of headers that are misspelled or otherwise vary from normal expectations.
- Association of emails to other emails based on unique header fields or other similar header content.
- Sequencing information about the infrastructures involved in email transport.
- Details of the protocols, MTAs, hardware, and software involved.
- Attribution information associated with unique identifiers.
- Groups of emails apparently sent from, through, or by the same or similar MTAs, systems, and mechanisms.

Figure 1 is an example of an analytical technique that has been helpful in understanding the processing of emails. It is the creation of delivery trees. These trees are generated by analysis of Received: headers and can include different subfields depending on the analysts' needs. Figure 1 is an example portion of the output of such an analysis from a collection of emails associated with a different case (details have been changed for obfuscation):

```
0 325802 B.net
    1 325090 mail.R.com
        2 325090 mail.R.com
            3 215585 mail.H.com
                4 232   mail.R.com
                4 24    altmail.H.com
            3 109301 altmail.H.com
                4 5     mail.R.com ...
```

Figure 1 - A tree depiction of an email handling process

In this example, of almost all of the 325802 emails arriving at B.net, 325090 arrived through mail.R.com, all of which came from mail.R.com, and almost all of those came through mail.H.com, most of the remainder coming through altmail.H.com. In this case, a close relationship is shown between R.net, R.com, and H.com, and interestingly, some email originally arriving at mail.R.com goes through mail.R.com, and back to mail.R.com before being delivered to B.net. This looping between providers may be evidence of intentional forwarding of emails through multiple jurisdictions in order to add damages to the legal action, depending on where the various servers are located.

Returning to the case identified above[1], given the definitions used by Plaintiff's expert and as stated in the U.S. Code, the total number of "emails" that could be at issue in the identified case came to only 175 out of the original asserted claim of 12,576, comprised of the 98 actual emails combined with the 34 unique active recipient addresses that were potential recipients of those actual emails. That analysis alone reduced the potential damages from more than US $10,000,000 to less than US$200,000. But that was not the end of the issue.

User signups, invitations, and other causes of emails

Another substantial limit on cases of unsolicited commercial email (UCE) is that laws tend to exonerate defendants that are fulfilling requests or emailing with respect to a pre-existing relationship with the addressee of an email. A user who requests information on a Web site may inadvertently or intentionally agree to terms and conditions granting the right to send or cause to be sent what would otherwise be UCE.

While individual suits for a few emails can often be resolved rapidly with regard to such actions of the addressee, when there are large volumes of emails involved, it may become problematic for the plaintiff to prove that the emails were not solicited or that the addressees requested cessation of those emails. The plaintiff presumably has to show that the emails in question were not solicited, which means getting involved in either the generation of legal documents signed by the individual recipients or using evidence associated with complaints or other similar information to make the case. For hundreds of thousands of emails sent to hundreds or thousands of recipients, this may involve an enormous quantity of paperwork and disturbance of many customers. In practice, high volume UCE cases have almost uniformly involved a small number of individuals who are either acting individually or asserting their role as an Email Service Provider suing a set of a few defendants repeatedly. They try various methods to create the means to assert large numbers of emails, such as taking over the accounts of previous users, allowing any emails directed to any recipient to be sent to them, making copies of emails sent to users and resending duplicates to themselves, forwarding the emails through multiple jurisdictions to create additional penalties associated with the additional jurisdictions, and so forth. In cases that are poorly defended, these tricks may work, and plaintiffs have won some number of very high valued default judgements against defendants who have not defended themselves. One such defendant has lost suits that in theory cold reach a total in excess of $1B, however; the owners shut down the business and left the country long ago, after committing frauds relating to the advertisers they served. It seems highly unlikely that these judgements will ever result in any actual compensation to plaintiffs.

Defending against such actions where plaintiffs create the conditions associated with high volumes of UCE typically involves showing that the emails that can be shown authentic were invited. As an example of invitation of emails, when an email has been identified as "spam" or otherwise identified as known by plaintiff to be undesired and yet is allowed to be sent, resent, forwarded, stored, and so forth, this may constitute an invitation.

From a technical standpoint, this involves showing that the protocols would not or could not have sent the emails unless the plaintiff invited them. Plaintiffs may also have an obligation to mitigate damages, and technically, the same issues arise. For example, configurations of email servers set to allow all emails not to known users to be sent and forwarded to a "dummy" account are clear indications of inviting the emails, and in such cases, courts have ruled that the emails were invited. Showing the technical basis for such claims typically involves a combination of examining log files, configurations, and ancillary information associated with these, testing these configurations in reconstructions with data from the case at hand, and showing that the configurations shown produce the particular results at issue.

In some cases, only screenshot images are available that are purported to depict the configuration of a product. Lacking version information and other similar relevant material, the analyst must ultimately make and state assumptions and draw conclusions based on those assumptions. However; the evidence can often help draw these conclusions. For example, in one such case, a configuration screen clearly indicated that the MTA was configured to forward emails identified as "spam" to a third party under a different email address. On its face this would tend to indicate that this party intentionally and knowingly sent the emails in question to the plaintiff, and as an intermediary, they then might reasonably be asserted to share some or all of the liability for damages.

In cases where emails to legitimate users are copied, potential liability arises based on contracts with users as well as the potential for violations of privacy regulations. In these cases, emails from the European Union may have different privacy requirements, and if the plaintiff fails to properly handle this they may run into legal liability. All of these issues involve technical analysis and interpretation of digital forensic evidence that must be requested and applied properly to make appropriate showings. To the extent that the plaintiff doesn't retain such information or fails to produce it, a legal argument can be made that preservation is required as soon as the plaintiff is aware of the potential for a legal action, and such arguments relating to data retention and disposition and legal holds on that data have been successfully made in many cases. [2]

While there is a duty for those who send UCE in the United States to retain information on removal requests and not send additional emails, signups are typically proprietary information of vendors involved in different aspects of the process. An advertiser almost certainly doesn't have the information on those who are being solicited with their advertisements, and can only process removals by providing them to the solicitor. Contracts between advertisers and those sending emails typically provide for the timely removal as well as requirements for following all applicable laws and regulations. This presumably limits the liability of the advertiser, but laws vary on how this liability may apply to those who order insertions of advertisements.

There are also other causes of emails that may be asserted for individual emails, for example, sent after removal requests are made. In one case, an email was shown to have been received a second time some six months after it was originally delivered.

Here is the header of the email received in May of 2008:

```
Received: from mac.com ([10.150.68.104])
by ms182.mac.com (Sun Java System Messaging Server 6.2-9.10 (built Feb 13
2008)) with ESMTP id <0K0S0079A5E9V1D0@ms182.mac.com> for
dr.cohen@mac.com;
Mon, 12 May 2008 17:01:21 -0700 (PDT)
Received: from smtp.usfca.edu (smtp.usfca.edu [138.202.192.18])
 by mac.com (Xserve/smtpin104/MantshX 4.0) with ESMTP id
m4D01JVB011284 for
<dr.cohen@mac.com>; Mon, 12 May 2008 17:01:19 -0700 (PDT)
Received: from smtp.usfca.edu (localhost [127.0.0.1]) by localhost (Postfix)
with SMTP id 2FE0F9E7 for <dr.cohen@mac.com>; Mon,
12 May 2008 17:01:18 -0700 (PDT)
Received: from PD-RP07-0003.usfca.edu (unknown [172.16.12.88])
 by smtp.usfca.edu (Postfix) with ESMTP id 1100B99B for
<dr.cohen@mac.com>;
Mon, 12 May 2008 17:01:18 -0700 (PDT)
Date: Tue, 30 Oct 2007 13:58:39 -0700
...
Message-id: <6.0.0.22.2.20071030135436.01df4e70@ace.usfca.edu
```

The message header as it arrived in October of 2007 is shown here:

```
Received: from mac.com ([10.150.68.68])
by ms182.mac.com (Sun Java System Messaging Server 6.2-8.01 (built Nov 27
2006)) with ESMTP id <0JQQ000QLSX4JLE0@ms182.mac.com> for
dr.cohen@mac.com;
Tue, 30 Oct 2007 13:58:16 -0700 (PDT)
```

```
Received: from smtp.usfca.edu (smtp.usfca.edu [138.202.192.17])
 by mac.com (Xserve/smtpin068/MantshX 4.0) with ESMTP id
l9UKvwmE016217 for
<dr.cohen@mac.com>; Tue, 30 Oct 2007 13:58:04 -0700 (PDT)
Received: from smtp.usfca.edu (localhost [127.0.0.1]) by localhost (Postfix)
with SMTP id A18AA2504 for <dr.cohen@mac.com>; Tue,
30 Oct 2007 13:57:57 -0700 (PDT)
Received: from PD-RP07-0003.usfca.edu (unknown [172.16.12.74])
 by smtp.usfca.edu (Postfix) with ESMTP id 845FD24F6 for
<dr.cohen@mac.com>;
Tue, 30 Oct 2007 13:57:57 -0700 (PDT)
Date: Tue, 30 Oct 2007 13:58:39 -0700
. . .
Message-id: <6.0.0.22.2.20071030135436.01df4e70@ace.usfca.edu>
```

This example makes the case for getting access to system logs, information on system failures, crashes, reboots, and breakins, and other events that might cause the things observed in the digital forensic evidence associated with emails as well as other matters. In this case, it has not yet been shown definitively what the cause of the email was due to a lack of time, resources, and desire to do so, but it appears that that this email was resent at a later date as a result of a restoration of a system that was backed up while the original email was still pending delivery. After a subsequent system crash, the system content was restored from backups, resulting in the reinsertion of the pending email from 6 months earlier into the mail queue. The MTA then processed the pending email normally and delivered the identical message. In a bulk email system, such an outage and recovery could produce tens of thousands of emails, depending on the MTA, backup, and recovery mechanisms in use.

Claims relating to RFCs treated as if they were legal mandates

In many of the cases involving high volumes of UCE, plaintiffs try to make various claims associated with headers being false or misleading based on their compliance or non-compliance with Internet Requests for Comments (RFCs). As far as legal precedent goes, as of this time, courts have not ruled that RFCs constitute legal contracts or are enforceable in a legal sense. Nevertheless, claims are made with regard to RFCs in many such cases, and expert witnesses are called upon to testify with regard to RFCs, their interpretation in context, and the extent to which they may have been violated.

HELO lines

One of the most common assertions made by plaintiffs in these cases is that the HELO protocol produced indicates a fraudulent source. The HELO exchange is used in the initiation of an RFC 821 simple mail transfer protocol (SMTP) exchange in which the sending computer is supposed to send "HELO " followed by the currently allocated domain name of the sending computer. The HELO information, if available in the header of an email, is typically recorded by the receiving computer within a Received line associated with that hop in the delivery process. RFC 2821 is the updated version of the SMTP protocol that uses "EHLO" instead of "HELO" to initiate its processing, indicating that the receiving server that the RFC 2821 protocol applies, and RFC 2821 indicates that in cases when the "HELO" protocol is used, RFC 821 must be used to process the emails.

It turns out that most of the emails seen in bulk email cases and most high volume email processes in use today appear to use RFC 821 rather than RFC 2821. This means that RFC 821 and not RFC 2821 applies to most of the emails in question in most of these cases to date. In response to the HELO, the receiving server is supposed to reply "250 OK" or some other similar response that starts with a "250 " to indicate that email is being accepted from the sender. Receiving servers often capture the IP address of the remote machine in this process and provide the "HELO", the information provided by the sender after the HELO, and the IP address of the sender in the Received: line to allow trace-back and association of the email with the originating IP address and asserted domain name.

Filtering based on the HELO information is sometimes used, for example, to prevent emails from known undesired source domains. Some MTAs check the IP address against the domain name using a DNS number-to-name lookup and place a warning in the header for spam filtering to notify the spam filter of a mismatch, however; it is extremely common for DNS names and IP addresses to not match in a number-to-name lookup. In particular, this stems from several reasons, including without limit, the use of large numbers of domain names associated with a single IP address, the use of proxy servers for delivery of emails, the use of email delivery services for delivery of emails, the incorrect naming of servers in configuration, and default server names not updated in configuration. Dynamic DNS introduces still other complications, and multiple answers to name-to-address lookups are not compensated for in many reverse lookup approaches.

In addition, the notion that a "HELO" line is deceptive in the the sense of being used to fool a recipient as to where an email is coming from seems problematic to assert. RFC 821 demonstrates examples of domain names like localhost, and nowhere does it assert that there is any requirement of authenticity. Furthermore, email recipients never see HELO lines sent to SMTP servers unless they are looking at log files associated with the emails, and the recording of HELO information is not mandatory or intended for users to see. The normal presentation of emails does not include the areas that include the HELO information, and even the most commonly used email clients have versions that send the name of the receiving computer in the HELO line instead of that of the sending computer, apparently because the authors misread the RFCs. Nevertheless, this seems to be a claim repeated by plaintiffs again and again.

False sender identities

Another common claim by plaintiffs is that the use of a fictitious name or email address in a sender identity (e.g., a "From:" line) is deceptive. For example, some have claimed that the use of an email address not containing the name of the sender is fraudulent because it misleads the recipient into believing that the sender is someone they are not. While on its face this may seem like a cogent argument, systems in the Internet use fictitious names and pseudonyms all of the time, including the use of names like "accounting" in RFC 821 and any of a wide range of other sender names in emails from almost any company that can be identified.

In most of the high volume email cases to date in which this claim has been made, the plaintiff also uses false names, as do the plaintiff's providers and customers, making the claim that much more problematic. However; the issue is not all that clear in law, and there is a real possibility that some court will eventually rule differently, making pseudonyms and anonymized names problematic as well.

From a standpoint of the forensics expert, in addition to testifying with regard to the conventions in use in the Internet and other common usage, it may be incumbent upon them to examine the use of naming by the plaintiff and defendant, their ISPs, other providers, those they purchase from and sell to, and to make some sort of finding with regard to the potential counters to the arguments being put forth. Again, this involves examining log files, user identity information, other emails including those sent by plaintiff, and other similar sources to identify contradictory information that might be revealing relative to the matter at hand. While many forensics experts may decide to jump into the fray and give an opinion about a fictitious name being misleading, this is problematic. Unless the digital forensics expert is also an expert in linguistics, they risk having their credibility destroyed along with the rest of their testimony.

To the extent that forgeries of sender identities are used, they may be identified as such by the expert so long as there is a basis for showing that the user identity in use was indeed that of another user. For example, there are cases where malicious actors substitute a legitimate sender identity and email address in an email and act as if they were the forged individual. Such a case has recently appeared in which well over 2,000 usenet postings used false sender identities including email addresses and names, in order to discredit an individual and take advantage of the damage to their reputation in order to gain an elected office in a corporation. Tracing these down to sources then becomes the issue.

False Received: headers

Some emails have been sent in some cases forged Received: headers. In these cases, a sender apparently added these headers within the header of the email so as to try to mislead someone trying to trace the email so that they would skip over the actual sender to fictitious previous senders. These are problematic in individual email cases if the forger has taken care to use real sequences of locations that could feasibly have been involved in emails, however; such forgeries are not as trivial as they may seem, and almost always fail to convince an analyst who is able to get records from other sites. In many cases, these forgeries involve a common intermediary associated with many other reception sequences, leading to obvious detection when shown in the tree format such as that in Figure 1 above.

Claims of fraudulent subject lines and inconsistencies in claims

One area where digital forensic evidence examiners are less likely to be helpful is in the examination of false subject lines or content for deception. Unless the forensic examiner is a linguistics expert, it would be a mistake to assert expertise in the use of language in commercial exchanges. However; one area where technical analysis has shown utility is in examining the characteristics of plaintiff claims in this regard.

Plaintiff claims have been repeatedly ruled by courts to require explicit identification of the specific statutory violation associated with each asserted email. This typically gets reported in a spreadsheet that identifies the assertions with respect to each email claimed as a violation. Many emails are typically claimed as having many violations, thus leading to a complex collection of different factors to be evaluated for different emails. The digital forensic examiner then has the opportunity to use automation to analyze the claims against the asserted extracts of emails.

In one recent case, examination of the spreadsheet provided by the plaintiff indicated that for the same subject lines, different emails made different claims, some claiming the subjects were deceptive and others claiming that they were not. This inconsistency was identified by the defense and used to assert that the language was identical while the claim was not, and that therefore, the inconsistency in claims should lead to throwing them out. After all, how can the same subject be deceptive in one case and not in another? Of course this leads to examination of the bodies of the emails, but in the matter at hand, the plaintiff had not examined the bodies of the emails in question, there being many tens of thousands of them.

Such inconsistencies in claims also goes to other aspects of those claims, and of course, to the extent that the analyst can identify them, these make for sound challenges to claims in a legal matter.

Assessment of damages

Experts may also be called upon to assist in the assessment of damages. Damage claims typically come in the form of disk and bandwidth usage needed to support the additional burden of the emails. However, this claim is highly dubious. When the plaintiff supplies all of the evidence of hundreds of thousands of emails on a single CD-ROM, it is apparent that the storage cost for the entire collection is on the order of $1 or less. If the emails average $10K each and there are 100,000 of them, the total bandwidth consumed is on the order of 10 billion bits. At 1 megabit per second, this is only 10,000 seconds of total bandwidth consumed, and typically these emails arrive over periods of a year or more. 10,000 seconds is less than 3 hours, and spread over a year (31536000 seconds) this comes to about 0.03% of the bandwidth. If Internet services at that rate cost $1,000 per year, this comes to well under $1 per year. So the damages from all of the handing and storage of these emails is well under $10 per year. For that reason, all of the damages in these cases are from statutory values.

In most cases, the most obvious path to assessing damages comes from dividing the total cost of some element of the plaintiff's costs by the percentage of those costs borne as a result of the emails. In practice, most of the plaintiffs in these cases to date don't keep good enough records of customer complaints regarding emails to allow them to be traced to any particular emails, if they can be attributed at all. Furthermore, if they don't have to add a full time employee or upgrade services exclusively for handling the emails in question, it will be hard to justify prorating work of workers who would be employed in other tasks unless the failure to perform those tasks could be identified with some loss. Again, the digital forensics expert can really only testify regarding percentages of usage of information resources and related record keeping systems.

Tracing emails to their origins

In order to make any real case, it will almost certainly be necessary to trace emails to their origins. There are essentially two paths to this. One is the normal tracing methods used to track down Internet traffic, and in a civil case, as in a criminal case, this involves the use of search warrants, preservation orders, and step-by-step analysis of records from system after system until the records get traced back to the origins. At that point, contractual disclosures and financial records associated with the discovery process in the legal action can be used to lead to the source of the insertion order, and can then be used to identify multiple parties involved in the activities. This is expensive and time consuming, and since most of the plaintiffs to date have been spending as little money and time as possible and seeking high dollar value settlements to prevent the defendants from winning the cases outright, it is rarely seen. However; if done properly, it can lead to the real sources of the emails and a potentially very high valued win or settlement.

The other path is a shortcut to the origin. There are many such shortcuts that have been tried in these cases. These include predominantly the use of information from the bodies of emails to associate those emails to their origins, and the use of deceptions to trap the sources. This effort is typically used to identify a small number of emails and then the claim is made that all similar emails came from the same source. This bootstrapping approach is particularly worthwhile in civil procedures where all that is needed to win a case is a slightly higher likelihood of guilt than innocence.

Deceptions for tracking emails to sources are predominantly implemented by filling in of forms with unique and false information and setting a trap for the return contact that yields the identity of the ultimate recipient of the leads from these advertisements. The legal action then goes to the recipient of the leads who typically have records associated with the leads used to pay the providers of those leads. These are tracked back to the email senders who are also sued, and through their infrastructure to the proximate causes of the emails. The problems with this approach include, without limit, (1) it may not produce a unique sender because of lead sharing, (2) the person sending the advertisement may not be the person who "benefits" from it, (3) care must be taken in assuring that the process is properly recorded, and (4) just because one email produces this behavior doesn't mean others will produce the same result. In the case identified above, Plaintiff filled out a form and provided a phone number set up specifically for the purpose, and did so under a fictitious name. However; the leads in this case were shared by an intermediary with others, and the contract with Defendant was exclusive and the lead provider was terminated after Defendant found that their leads were not being treated uniquely. The court ruled for Defendant in that it was not responsible for leads that it had not contracted for.

Information in the bodies of the emails are the second approach used. This typically involves the assertion that a URL contained within an email is used by a defendant in their business to track or display advertisements. If it is relied upon by the defendant then the theory is that it should be an adequate record to show that the defendant caused the email to be sent. The main obvious problems with this approach include, without limit; (1) competitors can and regularly do use "image servers" of others in their businesses so that other companies pay for the space, artwork, and bandwidth while they gain the financial advantages, (2) a malicious actor could provide the information for the purpose of damaging the defendant's reputation, and (3) anyone else could use the URLs for any purpose, including for falsifying the records to create a legal action.

Other emails are associated to emails that are identifiably traced by common mechanisms, such as the apparently common mechanism used to generate similar sentences. These are analyzable by creating templates for the sentence structures identified and identifying the common partial phrases used, then creating a generator and analyzer that can both generate equivalent structures and detect the ones in use by the apparently common mechanism.

In both of these approaches, there are problems associated with matching up larger numbers of emails with the small number for which there is evidence. The most important point to start with is that plaintiffs in these cases are seeking millions of dollars, and fail to produce any independent records to show that the emails they claim were sent were ever actually sent to them, as opposed to being generated by them. The plaintiffs tend to have adequate technical skills to generate the emails on their own by taking one or two emails and writing small scripts to generate the rest of them, they do poor record keeping or regularly destroy the very records that might prove their cases, and they publicly or privately claim that they are acting in a manner that would clearly put them in the category of vigilantes with respect to unsolicited commercial emails.

Making a case against real violators

Experience indicates that making a real case against bulk email senders involves most of the same elements we would normally find in making any other sort of legal case involving digital forensic evidence. Challenges to digital forensic evidence, in the larger sense, is covered in [3], and all of these challenges must be met in order to make a case against a competent defense using skilled digital forensic examiners. However; the key factors that differentiate bulk email matters from other matters are that (1) the evidence must be explored using automation, and that automation must meet legal standard in its makeup, reliability, calibration, operation, and application, (2) a wide range of contemporaneous records must be identified, properly collected, and properly preserved in order to create the set of evidence necessary to prove the matter at hand, and (3) there are a wide range of mistakes that are commonly made when non-professionals perform the tasks associated with digital forensic evidence, and the higher the volume of evidence, the more opportunities there are to make these mistakes.

Summary, conclusions, and further work

It appears most of the legal cases to date involving high volumes of unsolicited commercial email involve vigilante plaintiffs, evidence that is poorly constructed and not well researched, exaggerated claims that are not supported by the facts and that can be readily challenged by digital forensics experts, spoliated evidence, and large volumes of emails that are invited. However; this does not mean that the emails are not being sent or that those emails would not be in violation of statutes if they could pass the other tests. Experience tells us that there is plenty UCE and that much of it is fraudulent.

It appears that if legitimate digital forensic evidence experts were involved in these cases from the start of the process, many of the pitfalls encountered might be avoided and that defendants would have a far tougher time in defeating plaintiff claims, at least for cases where those claims are legitimate.

Regardless of which side of a particular matter the digital forensics expert is on, the issues and methods identified herein will be useful in assuring that the evidence is properly identified, collected, preserved, processed, analyzed, interpreted, and presented so as to support the legal process in an appropriate manner and one that can hold up to the legal scrutiny it is likely to receive. In addition, the techniques and issues discussed herein apply to most other high volume matters involving emails, and should be applied as appropriate. Clearly there is more work to do in definitively examining emails in large volume.

Finally, the resolution of the identified legal matter used as a case study was that Defendant won on a summary judgement. While digital forensic evidence played a substantial role in that decision, as always, the evidence, analysis, and interoperation only had specific utility in the context of the specific case. Nevertheless, the techniques used may be applied over a far wider range.

References:

[1] ASIS Internet Services vs. Optin Global, Inc., Case No. C-05-5124 JCS in the US District Court, Northern District of California

[2] The Sedona Conference Glossary: E-discovery and Digital Informaiton Management, A Project of The Sedona Conference® Working Group on On Electronic Document Retention & Production (WG1) RFP+ Group http://www.thesedonaconference.org/content/miscFiles/TSCGlossary_12_07.pdf

[3] F. Cohen, "Challenges to Digital Forensic Evidence", ASP Press, 2008.