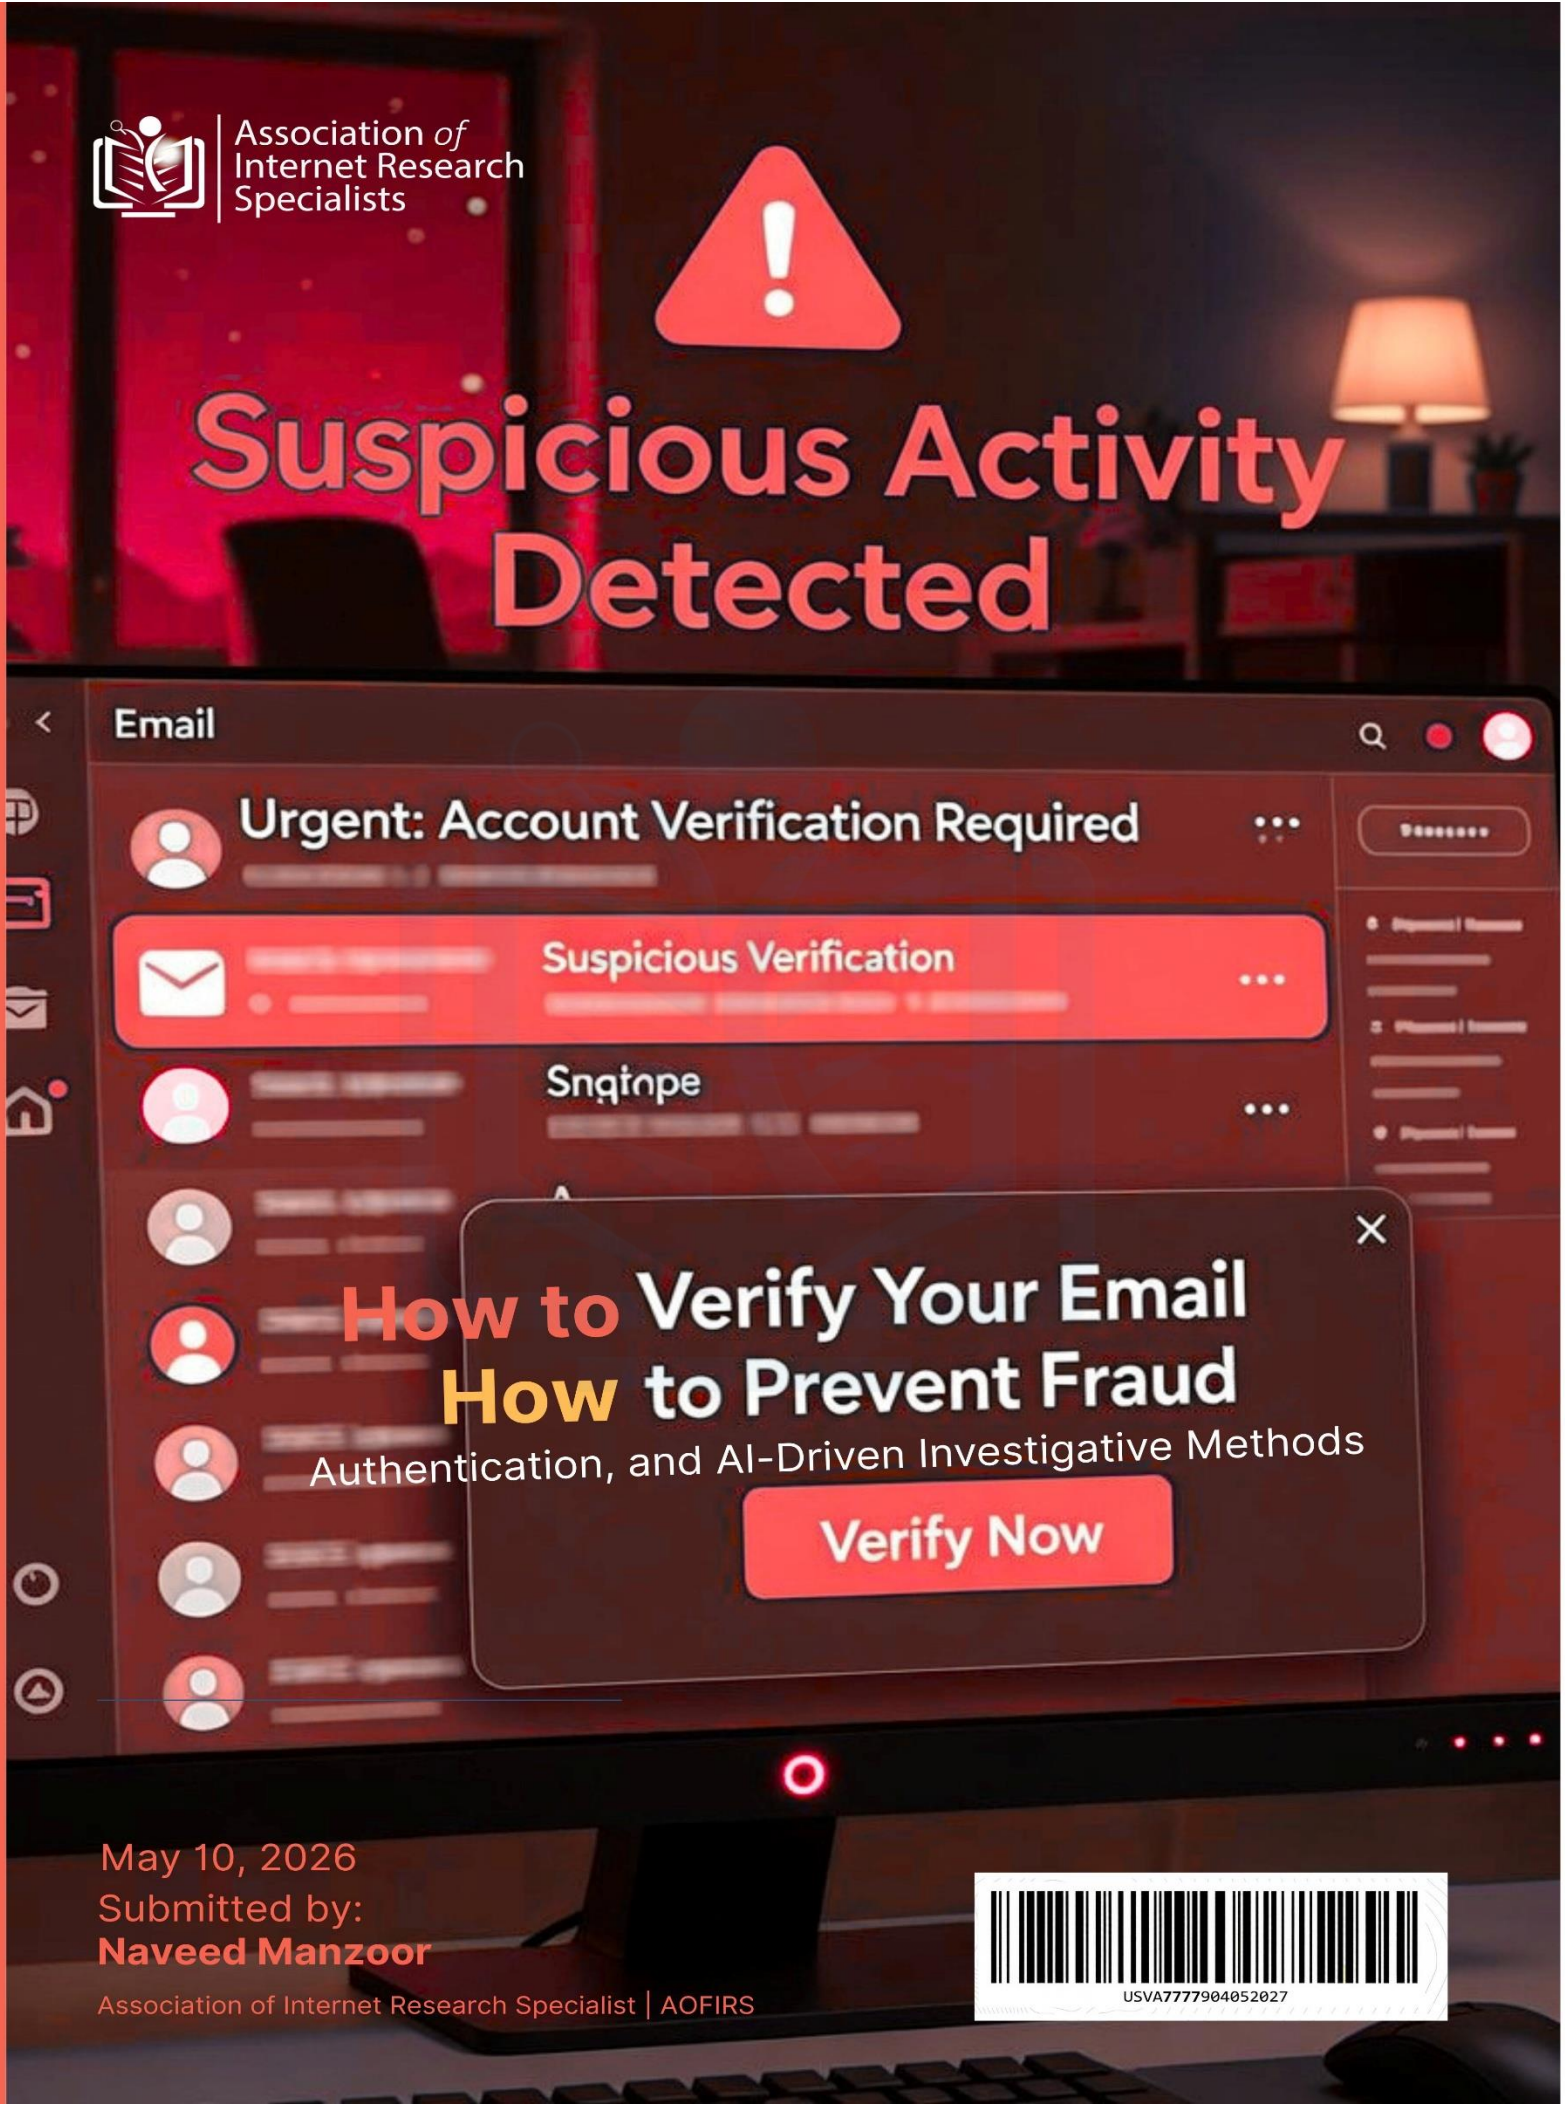

May 10, 2026

Submitted by:
Naveed Manzoor
Association of Internet Research Specialist | AOFIRS

USVA7777904052027

EMAIL VERIFICATION

Origin Analysis, Authentication, and AI-Driven Investigative Methods

A Research Paper on Email Search, Header Forensics,
Skip Tracing, and IP-Based Identity Verification

Prepared for online research practitioners, fraud and risk analysts,
and candidates pursuing the Certified Internet Research Specialist (CIRS™) credential

Original course notes preserved and expanded with 2026 methods, software, and case data

Compiled June 2026

Table of Contents

PART I — CORE TECHNIQUES

1. Email Verification: Purpose and Scope
- 1.1 Why Email Verification Has Become More Difficult — and More Necessary

2. Email Technology — A Quick Overview
- 2.1 Email Traffic in the Internet Cloud
- 2.2 DNS Resolution and the Transfer Process

3. The Basic Email Messaging Components in the Email Header
- 3.1 A Message Enclosed in an Envelope
- 3.2 Message Header Fields
- 3.3 The IP Address
- 3.4 Other Standard Header Fields

4. X-Spam-Status and the Anatomy of a Full Header
- 4.1 X-Spam-Status or Spam Class
- 4.2 Illustration of Spam Detection Software with [X] Values
- 4.3 Partial Headers and Full Headers
- 4.4 Illustration [Full Header]: IP Detection and Identification of Private & Public IPs

5. From (Envelope) vs. From: (Header), and Two Forgery Walkthroughs
- 5.1 From (Sender's Envelope) and From: (with Colon)
- 5.2 Sender and Comment Fields
- 5.3 Worked Example — Identifying a Spammed Header
- 5.4 Illustration [Full Header]: Masked or Alias Email Address in Gmail

6. Finding Full Headers, and the Spam / Phishing / Spoofing / Scam Taxonomy
- 6.1 Finding Full Email Headers in Email Clients
- 6.2 Spam, Phishing, Spoofing, and Scams
- 6.3 Common Scam Categories

7. Case Study — Examining a Phishing Email (Fictitious Lotto Scheme)
- 7.1 The Header Sorter and Organizer Tool
- 7.2 Raw Header — Delivered-To: cybersearchers@gmail.com
- 7.3 Hop-by-Hop Reconstruction
- 7.4 Cross-Verification with Four Independent Tools

8. Reputation Intelligence, Reverse DNS, and Disposable Addresses
- 8.1 Email Header MTA Servers — "Email Hops/Path Analysis Tool"
- 8.2 Sender's MTA Server Physical Location Trace with IP
- 8.3 Email Servers Reputation Intelligence
- 8.4 Spamming Reputation Authority Check
- 8.5 Spammers Reputation Database Checks
- 8.6 Check DNS Mapping with IP via "Reverse DNS Lookup" or DNS Authentication
- 8.7 Random and Temporary Email Addresses for Spamming

PART II — 2026 METHODS, SOFTWARE, AND APPLIED GUIDANCE

9. Modern Verification Methods: Authentication Protocols and AI

9.1 SPF, DKIM, and DMARC — From Optional to Effectively Mandatory
9.2 How AI Spam and Phishing Filters Actually Work in 2026
9.3 AI-Powered Verification in Commercial Software

10. The Software Landscape: What Is Used for Email Verification Today

10.1 Bulk and Real-Time Email Verification Platforms
10.2 Email-Finder and Lead-Generation Tools
10.3 Skip Tracing and People-Search Databases
10.4 IP Geolocation Services

11. Applied Questions: Spam, Solicitations, Remarketing, Fake Emails, Skip Tracing, and IP Location

11.1 Is This Email Spam?
11.2 Is This an Unwanted Solicitation?
11.3 Is This Remarketing or Retargeting Email Legitimate?
11.4 How Do I Tell a Fake Email Address from a Real One?
11.5 Skip Tracing the Individual Behind an Address
11.6 Identifying the IP Location Behind a Message

12. Real-World Examples and a Complete Use Case

12.1 Business Email Compromise: Scale and Recent Cases
12.2 Complete Use Case — Verifying a Suspicious Vendor-Payment Email

13. Building Email Search Queries

13.1 When You Know the Name of a Person
13.2 When You Know the First Name but Not the Last
13.3 When You Don’t Know the Exact Name, Alias, or Abbreviation Used
13.4 Using Search Query Operators as in Google

14. Why the CIRS Course Is Ideal for Learning These Techniques

14.1 The Curriculum Is Built Around Exactly This Methodology
14.2 It Teaches the Convergent-Evidence Discipline This Paper Repeatedly Relies On
14.3 The Program Covers the Full Breadth This Paper Has Walked Through
14.4 Recognized Standing, Not Just Self-Study
14.5 Structured, Flexible Delivery for Working Researchers

BACK MATTER

15. Appendix — Original Course-Era Software Reference (Preserved)

15.1 Email Search Software — For Leads Collection
15.2 Personal Information and Background Check Software

16. Conclusion

References and Sources

1. Email Verification: Purpose and Scope

Before attending to any email you have collected on social media, it is a good idea to verify the ownership. It is important to validate the authenticity of the social media account or the source from where you are collecting the news, video, or any other important information. Even when you want to recover passwords or register yourself online, the registrant is asked to validate the email address first. Besides, when you receive an email for subscriptions, registration, or verification, a certain due diligence to verify the source is prudent. Sometimes email verification is necessary for investigation reasons, such as tracing the spoofing source and phishing emails with IP records, or to analyze timestamps along the delivery route for legal reasons. Whatever the reason, as a researcher we are bound by the duty of due care for information exchange when emailing is the primary source.

Here are a few of the many reasons email verification remains useful for research:

- **Source authentication.** To verify the source of information or data and to authenticate and validate its integrity and usefulness — for example, to verify a news source from social media and distinguish a fake account from a real person's account.
- **Time-sensitive evidence.** To track an email message's time and date stamp for time-sensitive or time-barred information.
- **Confidential or anonymous exchange.** Where information exchange is confidential or anonymous — examples include war-torn areas, geo-politically restrained countries, human-rights reporting, and censorship environments restricting free information exchange.
- **Legal evidentiary requirements.** When information-collection methods are defined by law, particularly for electronic evidence collection, preservation, and presentation.
- **Rights and permissions.** Depending on the research type, a researcher's dependence on emailed information for obtaining copyright permissions, or for publishing and re-publishing data.
- **Technical re-enforcement.** Detecting interception attempts, identifying the mail server and its reputation, verifying originating domains, and surfacing other technical information that re-enforces email validation.

We begin by learning how emails work within networks and over the Internet, and how to identify the components of an email message, before moving into the techniques used to determine whether a sender is legitimate and whether a message is spam.

1.1 Why Email Verification Has Become More Difficult — and More Necessary

The fundamentals of email routing described in this paper have not changed since the original protocols were standardized decades ago. What has changed is the sophistication of the content riding on top of that routing. Large language models can now draft a grammatically flawless, contextually appropriate phishing email in seconds, and research shows that AI-rephrased phishing content evades the major enterprise spam filters at significantly higher rates than human-written lures, because those filters were trained on the linguistic fingerprints of older, clumsier scams.

Key statistic

Industry monitoring in 2026 estimates that over half of all spam now in circulation is AI-generated, and the FBI's Internet Crime Complaint Center reported phishing-related losses nearly tripling year over year as generative tools lowered the cost of producing a convincing lure. The practical implication for a researcher is that **content alone — spelling, tone, urgency — is no longer a reliable signal.** The header, the authentication results, and the IP trail behind the message remain the most durable evidence, because those are mechanically generated by mail servers and are far harder for an attacker to fabricate convincingly than prose.

This is the throughline of the paper: the original header-analysis techniques developed for 2017-era spam are, if anything, more relevant in 2026, because they verify the machinery of delivery rather than the words of the message. The new sections that follow layer current software, AI-assisted detection, skip tracing, and IP geolocation on top of that original foundation — they do not replace it.

2. Email Technology — A Quick Overview

The originating sender creates an email in their Mail User Agent (MUA) and clicks ‘Send.’ The MUA is the application the sender uses to compose and read email, such as Outlook, Apple Mail, or a webmail client. The sender's MUA transfers the email to a Mail Delivery Agent (MDA). An MDA, or Message Delivery Agent, is a software component responsible for delivering messages to a local recipient's mailbox. The most common arrangement is for an email user (the client) to make an arrangement with a remote Mail Transfer Agent (MTA) server for the receipt and storage of the client's emails. The MTA, using a suitable MDA, adds email messages to a client's storage as they arrive.

Figure 2.1 — Simplified email send/receive process from MUA through the Internet cloud to the recipient's mailbox.

Frequently, the sender's MTA also handles the responsibilities of an MDA — several of the most common MTAs do this, including sendmail and Gmail. The MDA/MTA accepts the email, then routes it to local mailboxes or forwards it if it isn't locally addressed. The email enters the first of a series of “network clouds,” typically labeled as a company or ISP network cloud.

2.1 Email Traffic in the Internet Cloud

An email can encounter a network cloud within a large company or ISP, or the wider Internet cloud. The Internet cloud contains the ISPs that host email servers alongside DNS servers. These devices are protected by firewalls, spam filters, and malware-detection software that may bounce or even delete a message.

Email service providers — usually ISPs and other companies — process a large volume of email and operate multiple mail servers and routers that route messages via whichever MTA is available. An outbound email enters a queue with other outgoing messages; if there is a high volume of mail in the queue, the message is delayed until the MTA processes the messages ahead of it. When transferring an email, the sending MTA handles every aspect of delivery until the message is either accepted or rejected by the receiving MTA.

As the email clears the queue, it enters the Internet cloud, where it is routed along a host-to-host chain of servers. Each MTA in that chain needs to “stop and ask directions” from the Domain Name System (DNS) to identify the next MTA in the delivery chain. The exact route depends partly on server availability and mostly on which MTA can be found to accept email for the specified domain. A pair of messages from the same host, addressed to the same receiving host, could take different paths — and it is mostly spammers who deliberately route their message through a series of relay servers to obscure the true origin.

To find the recipient's IP address and mailbox, the MTA must drill down through the DNS, beginning with the root nameservers at the top-level domain (.tld), then domain nameservers handling requests within that .tld, and eventually nameservers that know about the local domain.

2.2 DNS Resolution and the Transfer Process

There are thirteen root server clusters serving the top-level domains (such as .org, .com, .edu, .gov, and .net). These root servers refer requests for a given domain to the name servers that handle requests for that TLD, though in practice this step is seldom necessary.

The MTA can usually bypass this step because it already knows which domain name servers handle requests for these TLDs. It asks the appropriate DNS server which Mail Exchange (MX) servers have knowledge of the subdomain or local host in the email address. The DNS server responds with an MX record: a prioritized list of MX servers for that domain.

> **Concept check — MTA vs. MX server**
>
> An MX server is really an MTA wearing a different hat — like a person who holds two jobs with different titles. To the DNS server, the server that accepts messages is an MX server; when it is transferring messages onward, it is called an MTA. The MTA contacts the MX servers in priority order until it finds the designated host for that address domain, then asks whether the host accepts messages for the recipient's username at that domain (username@domain.com) and transfers the message.

Email construction and delivery closely mirrors regular postal mail by design, because the email standard was deliberately modeled on the conventions of physical mail.

3. The Basic Email Messaging Components in the Email Header

3.1 A Message Enclosed in an Envelope

An email message is constructed like a letter sent through the postal service: a message enclosed in an envelope. The email envelope header is analogous to the envelope of a hardcopy letter, but some information ordinarily present on a hardcopy envelope is instead contained in the message header. This header also contains information not usually found on a real-world envelope, but essential to email delivery and troubleshooting. The envelope header is usually hidden when viewing an email, while the message header is usually visible. Together, these two headers are called the Full Header.

3.2 Message Header Fields

Anyone who has used email is familiar with the message header, displayed when viewing a message, including the ‘From:’, ‘To:’, ‘Cc:’, ‘Date:’, and ‘Subject:’ fields. The content of these fields differs only slightly from regular mail, because the ‘From:’, ‘To:’, and ‘Cc:’ fields in an email identify the sender and intended recipients by email address rather than street address.

The Date Field

The message header's ‘Date:’ field is applied by the originating sender's MUA, so it is only as accurate as the clock on the sender's computer.

The Subject: Field

The ‘Subject:’ field has no exact equivalent in regular mail except in formal business letters, where its closest analogy is the ‘Re:’ line. This field is necessary for email because, without it, you could only differentiate one message from another in the inbox using the ‘From:’, ‘To:’, and ‘Date:’ fields.

The Return-Path

Email contains more detailed information about its delivery process than the single postmark of regular mail. As the email passes through the delivery chain, MTAs add postmark-like timestamps and MTA location information, including the envelope header's ‘Received’ fields and the ‘Return-Path,’ which contains the identity of the sender, such as <username@domain.tld>.

The ‘Return-Path’ is often referred to as the “envelope sender,” and this is the address mailing lists use to determine who sent a message. It is also the address to which bounces are sent.

The Received Field

A ‘Received’ field is added by each MTA in the delivery chain as it accepts a message from the previous MTA. When a receiving MTA accepts the email for relay or local delivery, it records information about the sender's path in the envelope header, including a Message ID and timestamps indicating the time of transfer and the identity of the sending MTA. Following the ‘Received’ entries from bottom to top leads back to the originating MTA — though not necessarily to the sender's actual email address.

Because ‘Received:’ lines are generated by the mail servers involved in the transfer, there are at least two such lines per header, since one MTA server sends and another receives. The first ‘Received:’ line, from the originating sender's own system, is the least likely to be forged or spoofed.

> **Note — why the Received chain matters for spam triage**
>
> This information about the true identity of the sending MTA is valuable when troubleshooting spam or malicious messages, which often contain forged identities in the ‘From:’ and ‘Reply-To:’ fields. The true identity of the sending MTA can usually be extracted from the envelope header. When contacted by the sending MTA, the receiving MTA checks whether the hostname provided by the originating MTA resolves to a unique IP address. If it does, it is a fully qualified domain name, and the receiving MTA adds that information to the ‘Received’ field. If the hostname does not resolve properly, the receiving MTA instead adds the originating MTA's IP address (and possibly the true hostname).

Each mail system has its own style of ‘Received:’ line. It typically identifies the machine that received the mail and the machine the mail was received from. For example:

```
Received: from foo.com by bar.com id AA15057; Fri, 25 Jul 97 09:39:02
```

The “foo.com” part is the name the sending machine used to identify itself — this may be forged in the case of spam. Many mailers add extra information, for example:

```
Received: from foo.com ([129.2.3.4]) by bar.com id AA15057; Fri, 25 Jul 97
09:39:02
```

Here, bar.com has inserted the IP address of the sending system. If the machine name does not match the IP address, you have likely identified the point where the mail was forged — the machine at 129.2.3.4 lied when it identified itself as foo.com. Any ‘Received:’ lines appearing subsequent to this are likely spam. The IP address can be checked against the MTA server name with a reverse IP lookup using tools such as mxtoolbox.com.

If the IP address does not make sense, or turns out to be a private IP address, the entire ‘Received:’ line can be considered a spammed source. The injection point is somewhere above that particular field. In the example below, the information in the ‘Received:’ field is complete and defines both the IP and the server name; if they match, the server is not rigged.

```
Received: from foo.com (x.y.alterdial.uu.net [129.2.3.4]) by bar.com id AA15057;
...
```

The Reply-To Field

The envelope header also contains a ‘Reply-To’ address provided by the sender that the receiver can use to reply, analogous to the return address on regular mail. Automated notifications and mailing-list messages often set a different ‘Reply-To’ address so bounced messages are sent to an automated handler. Like a return address, the ‘Reply-To’ address doesn't have to be real — if it isn't, mail sent to it will be undeliverable. Spam and malware-bearing messages are likely to have false information in the ‘From:’ and ‘Reply-To’ fields, but the originator's true Internet address is recorded in the first ‘Received’ entry in the envelope header.

Often absent from the spammer's message, and easily forged or masked, the ‘Reply-To’ field nonetheless often provides a clue worth checking. Forged spam frequently has a legitimate ‘Reply-To:’ field so that the spammer can actually receive mail orders.

3.3 The IP Address

To trace and track the whereabouts of the originating MTA server that sent the email, we need to look up the IP address of the server on the network that sent it. The IP is embedded in the header and generally easy to find in the X-Originating-IP field — identified as the “Originating IP.” These are Public IPs, used to track MTA server location. Another set of IPs may also be present in the header, but these are non-traceable Private IPs.

Public IP Address

A public IP address is an IP address that can be accessed over the Internet. Like a postal address used to deliver mail to your home, a public IP address is the globally unique address assigned to a computing device.

Private IP Address

A private IP address, on the other hand, is used to assign computers within a private space without exposing them directly to the Internet. For example, with multiple computers within a home, a router

gets the public IP address, and each connected device gets a private IP address from the router via DHCP (Dynamic Host Configuration Protocol), which may change internally over time.

IP Address Ranges of Private IPs

The IANA (Internet Assigned Numbers Authority) has reserved certain address ranges for private networks only. These Reserved IP Addresses need not be checked or traced, because they will not return any useful results:

- 10.0.0.0 to 10.255.255.255
- 172.16.0.0 to 172.31.255.255
- 192.168.0.0 to 192.168.255.255

3.4 Other Standard Header Fields

DKIM-Signature & DomainKey-Signature

These fields relate to domain keys, a cryptographic signing standard used by the sending domain to prove a message was not altered in transit. See the expanded discussion of DKIM, SPF, and DMARC authentication in Section 9.

Message-ID

A unique string assigned by the mail system when a message is first created. This is forgeable, though it requires more specialized knowledge than forging the ‘From:’ line. The Message-ID often identifies the system from which the sender is logged in, rather than the actual originating system. Its format is <unique string>@<sitename>. Each mail software has its own style of unique string, and sloppy forgeries often get this wrong — a forgery can be confirmed by comparing the Message-ID against legitimate messages from the same site.

MIME-Version

Multipurpose Internet Mail Extensions (MIME) is an Internet standard that extends the format of email to support attachments, non-ASCII text, and rich content.

Content-Type

Generally tells you the format of the message body, such as HTML or plain text.

4. X-Spam-Status and the Anatomy of a Full Header

4.1 X-Spam-Status or Spam Class

This field displays a spam score created by a mail client, service, or the ISP's email servers. The ISP's MTAs configure the account's mail filter to insert a special identifier called X-Spam-Status in the header area of the message, which can be used as a filtering mechanism within the email system.

If the total score of the message exceeds the mail filter's sensitivity level, the message is flagged. The Basic spam-filtering option typically has a predefined sensitivity level set to “3,” which can be scaled up to Advanced levels. This level is also called the X-Spam-Level or Spam Score. The X-Spam-Status identifier includes the successful spam tests identified, the message's total score (hits), and the mail filter's sensitivity threshold.

4.2 Illustration of Spam Detection Software with [X] Values

Header example #1, where the message is not spam:

```
To: <user@intermedia.net>
Subject: Regular Email
Date: Wed, 26 Feb 2009 14:05:47 -0800
MIME-Version: 1.0
Content-Type: multipart/alternative;
  boundary="----=_NextPart_000_0040_01C2DDA0.292F5970"
X-Priority: 3
X-Mailer: Microsoft Outlook Express 6.00.2800.1106
X-MimeOLE: Produced By Microsoft MimeOLE V6.00.2800.1106
X-Spam-Status: No, hits=1.3 tagged_above=-999.0 required=3,
  tests=[BAYES_20, HTML_70_80, HTML_MESSAGE]
```

Header example #2, where the message is spam:

```
To: <user@intermedia.net>
Subject: Spam Email
Date: Wed, 26 Feb 2009 14:01:16 -0800
MIME-Version: 1.0
Content-Type: multipart/related;
  type="multipart/alternative";
  boundary="----=_NextPart_000_002D_01C2DD9F.876C8590"
X-Priority: 3
X-Mailer: Microsoft Outlook Express 6.00.2800.1106
X-MimeOLE: Produced By Microsoft MimeOLE V6.00.2800.1106
X-Spam-Status: Yes, hits=10.6 tagged_above=-999.0 required=7.0
  tests=[BAYES_56, CLICK_BELOW, EXCUSE_10, EXCUSE_14, EXCUSE_22,
  EXCUSE_23, EXCUSE_24, EXTRA_MPART_TYPE, HTML_60_70,
  HTML_COMMENT_SAVED_URL, HTML_FONT_COLOR_BLUE, HTML_LINK_CLICK_HERE,
  HTML_MESSAGE, HTML_TAG_EXISTS_TBODY, MAILTO_LINK, OFFER,
  OFFERS_ETC, OPT_OUT, ORIGINAL_MESSAGE, RECEIVE_EMAIL, SMTPD_IN_RCVD]
```

In a spam message header, you may also find X-Spam-Flag and X-Spam-Score identifiers that duplicate the X-Spam-Status field and can be used to configure mail-server rules that sort spam messages into a separate folder.

Source: kb.intermedia.net/Article/1783

4.3 Partial Headers and Full Headers

Partial headers are what you normally see when you open an email — the From Address, To Address, Subject, Date and Time, Reply-To Address, Cc, and Bcc. These are the headers most relevant to daily tasks.

Full headers contain considerably more technical information than you normally see when checking email. These extra headers are sometimes needed to diagnose a delivery problem or trace a forged message. For a detailed reference of header indicators, see RFC 2076 (tools.ietf.org/html/rfc2076).

4.4 Illustration [Full Header]: IP Detection and Identification of Private & Public IPs

Here is what a typical Internet email header looks like. What you are looking for is the IP address, sometimes conveniently identified as the “Originating IP.” This allows tracing to the Internet Service Provider (ISP), along with the date and time of the email, using the IP address of the sender's server. In the printed example below, public IP addresses are the technically significant values to extract; private IP addresses (in the 10.x, 172.16–31.x, and 192.168.x ranges described in Section 3.3) should be set aside as non-traceable.

Full Header Example [Gmail Header]:

```
Return-path: <>
Envelope-to: john@example.com
Delivery-date: Wed, 02 Apr 2014 15:06:11 -0600
Received: from [46.165.209.232] (port=36531 helo=delivery.antispamcloud.com)
    by host309.HostMonster.com with esmtps (TLSv1:RC4-SHA:128)
    (Exim 4.82) (envelope-from) id 1WVSMM-0003oR-Ny
    for john@example.com; Wed, 02 Apr 2014 15:06:10 -0600
Received: from mail-ig0-f195.google.com ([209.85.213.195])
    by mx7.antispamcloud.com with esmtps (TLSv1:RC4-SHA:128)
    (Exim 4.82) (envelope-from) id 1WVSMJ-00049k-3X
    for john@example.com; Wed, 02 Apr 2014 23:06:10 +0200
Received: by mail-ig0-f195.google.com with SMTP id uq10so212231igb.2
    for <john@example.com>; Wed, 02 Apr 2014 14:06:02 -0700 (PDT)
DKIM-Signature: v=1; a=rsa-sha256; c=relaxed/relaxed; d=gmail.com;
    s=20120113; h=mime-version:date:message-id:subject:from:to:content-type;
    bh=mbhovkRL3Im017b5m7rMRTWVa1olgzE1U+yr8FXykLSM=; b=[...]
MIME-Version: 1.0
X-Received: by 10.42.107.67 with SMTP id c3mr2836464icp.28;
    Wed, 02 Apr 2014 14:06:02 -0700 (PDT)
Received: by 10.50.216.193 with HTTP; Wed, 2 Apr 2014 14:06:02 -0700 (PDT)
Date: Wed, 2 Apr 2014 15:06:02 -0600
Message-ID: <CAFr...@mail.gmail.com>
Subject: I can haz headers
From: HostMonster Tutorials <hostmonstertutorials@gmail.com>
To: john@example.com
Content-Type: multipart/alternative; boundary=20cf302075e4ed71d604f615a6cd
Received-SPF: pass (mx7.antispamcloud.com: domain of gmail.com designates
    209.85.213.195 as permitted sender) client-ip=209.85.213.195;
    envelope-from=hostmonstertutorials@gmail.com;
    helo=mail-ig0-f195.google.com;
X-SPF-Result: mx7.antispamcloud.com: domain of gmail.com designates
    209.85.213.195 as permitted sender
Authentication-Results: antispamcloud.com; spf=pass
    smtp.mailfrom=hostmonstertutorials@gmail.com
Authentication-Results: antispamcloud.com; dkim=pass header.i=gmail.com
X-Spampanel-Class: unsure
X-Spampanel-Evidence: Combined (0.15)
X-Recommended-Action: accept
X-Identified-User: {host309.HostMonster.com:local} {sentby:Delivered locally}
```

How to read this header

Reading bottom-to-top traces the path back toward the origin. The bottom-most public IP, **209.85.213.195**, belongs to Google's outbound mail infrastructure (mail-ig0-f195.google.com) — confirmed independently by the **Received-SPF: pass** and **dkim=pass** results, which state that gmail.com explicitly authorized this IP to send on its behalf. The next hop, **46.165.209.232**, is the antispam relay (antispamcloud.com) that the receiving host (HostMonster) uses before final delivery. Because both the hostname and the IP are mutually consistent at every hop, and the authentication results confirm them, this header shows no signs of forgery — contrast this with the forged examples worked through in Section 5.

5. From (Envelope) vs. From: (Header), and Two Forgery Walkthroughs

5.1 From (Sender's Envelope) and From: (with Colon)

The “From” line is not part of the envelope sender; it is part of the email header that the MTA inserts along its route toward other MTAs, adding addresses as it goes — much like a postal letter receives a stamp each time it passes through a post office on its route. The “From” address becomes part of the email headers as the message reaches its destination. The “From:” (with colon) address is where the recipient should respond.

In most cases, the From and the From: address match, but they don't always have to. If this were a personal email, both would nearly always match — if a message is addressed to you with your friend's email return address, the From address is almost certainly going to match your friend's From: address. See the example below:

The “From” and “From:” address match — example:

```
From webpromo@denmark.it.earthlink.net  Tue Jul 8 13:05:02 1997
Return-Path: <webpromo@denmark.it.earthlink.net>
From: webpromo@denmark.it.earthlink.net
Received: from denmark.it.earthlink.net (denmark-c.it.earthlink.net
  [204.119.177.22]) by best.com (SMI-8.6/mail.byaddr) with ESMTP id
  NAA21506 for <falk@falconer.vip.best.com>;
  Tue, 8 Jul 1997 13:05:16 -0700
```

Usually, the From line is set to the sender's address (often marked as from without the colon). This ensures you know who the message is from and can reply easily. Spammers want to make sure you cannot reply easily, and certainly don't want you to know who they are — that's why they insert fictitious email addresses into the From: lines of their junk messages.

Things are different for emails from a corporation. If you receive a box from Amazon, Amazon might have one return address for undelivered packages, while inside the box they ask for returns to a different address — or the package might be a gift, shipped by Amazon but really “from” Aunt Martha. With email, there are similarly legitimate reasons an envelope sender (From) might not match the From: header, such as a mailing list or a company's automated bounce processing.

> **Why this matters for fraud triage**
>
> This mismatch is a “feature” of email that spammers and scammers abuse. When a message picked up as spam is “from” PayPal, a bank, or another trusted institution, the From: address has generally been changed to a domain you recognize, while the envelope sender remains something the attacker controls. This exact technique underlies the brand-impersonation pattern used in business email compromise, covered in Section 11.

5.2 Sender and Comment Fields

Sender

The account that sent the message. Mail software is supposed to insert this line if the user modifies the From: line. Most MTA client software does not generate these lines, while other mailers provide an X-Sender: line instead.

Comment

Some mailers add additional information to the headers, such as “Authenticated sender is doe@foo.com.” Forged Comment: lines can be easily added to outgoing mail, so this line, too, can originate from spam. Other mailers may insert their own authentication information in the headers.

5.3 Worked Example — Identifying a Spammed Header

Here is an example of a spammed email:

```
From webpromo@denmark.it.earthlink.net  Tue Jul 8 13:05:02 1997
Return-Path: <webpromo@denmark.it.earthlink.net>
From: webpromo@denmark.it.earthlink.net
Received: from denmark.it.earthlink.net (denmark-c.it.earthlink.net
  [204.119.177.22]) by best.com (SMI-8.6/mail.byaddr) with ESMTP id
  NAA21506 for <falk@falconer.vip.best.com>; Tue, 8 Jul 1997 13:05:16 -0700
Received: from mail.earthlink.net (1Cust98.Max16.Detroit.MI.MS.UU.NET
  [153.34.218.226]) by denmark.it.earthlink.net (8.8.5/8.8.5)
  with SMTP id NAA12436; Tue, 8 Jul 1997 13:00:46 -0700 (PDT)
Received: from adultpromo@earthlink.net by adultpromo@earthlink.net
  (8.8.5/8.6.5) with SMTP id GAA05239 for <adultpromo@earthlink.net>;
  Tue, 08 Jul 1997 15:48:51 -0600 (EST)
To: adultpromo@earthlink.net
Message-ID: <199702170025.GAA08056@no-where.net>
Date: Tue, 08 Jul 97 15:48:51 EST
Subject: Hot News!
Reply-To: adultpromo@earthlink.net
X-PMFLAGS: 12345678 9
X-UIDL: 1234567890x00xyz1x128xyz426x9x9x
Comments: Authenticated sender is <adultpromo@earthlink.net>
Content-Length: 672
X-Lines: 26
Status: RO
```

Notice that the To: line does not address the person or a proper business email; the actual recipients are hidden, likely using a generic list distributed via blind carbon-copy (Bcc:). A tool such as emailchecker.com can be used to verify whether “webpromo@denmark.it.earthlink.net” is a real, deliverable address.

The first ‘Received:’ paragraph shows the email received from a service provider that appears to be Earthlink, which seems authentic. The second ‘Received:’ paragraph, however, shows an inconsistency:

```
... from mail.earthlink.net (1Cust98.Max16.Detroit.MI.MS.UU.NET [153.34.218.226])
```

Here, the server that delivered the mail to “denmark.it.earthlink.net” identified itself as mail.earthlink.net when it should have identified as “1Cust98.Max16.Detroit.MI.MS.UU.NET.” The third ‘Received:’ line would indicate a UUNET server rather than Earthlink, and this paragraph contains email addresses rather than server names.

Conclusion of the walkthrough

The overall analysis indicates a spammer attempted to mask and inject a simulated Earthlink server address, when the actual path shows the mail came from UUNET. It is important to note that even with the best header-path and MTA analysis, it may not always be possible to positively conclude that an

email is spam — UUNET servers may, on occasion, be linked to Earthlink under a legitimate business contract. The only definitive way to confirm server authenticity in an ambiguous case is to ask the primary ISP (here, Earthlink) directly for more information.

A Second Forgery

```
Received: from cola.bekkoame.or.jp (cola.bekkoame.or.jp [202.231.192.40])
  by srv.net (8.8.5/8.8.5) with ESMTP id BAA00705
  for <got@srv.net>; Wed, 30 Jul 1997 01:15:27 -0600 (MDT)
From: beautifulgirls585@aol.com
Received: from cola.bekkoame.or.jp
  (ip21.san-luis-obispo.ca.pub-ip.psi.net [38.12.123.21]) by
  cola.bekkoame.or.jp (8.8.5+2.7W/3.5W) with SMTP id OAA11439;
  Wed, 30 Jul 1997 14:35:50 +0900 (JST)
Received: from mailhost.aol.com(alt1.aol.com(244.218.07.32)) by
  aol.com (8.8.5/8.6.5) with SMTP id GAA00075 for <"">;
  Tue, 29 Jul 1997 22:19:42 -0600 (EST)
Date: Tue, 29 Jul 97 22:19:42 EST
Subject: You can have what you want...
Message-ID: <574857638458.HWF39862@aol.com>
Reply-To: beautifulgirls585@aol.com
X-PMFLAGS: 56354433 0
Comments: Authenticated sender is <aol.com>
X-UIDL: vjg79u26gfkjjrty38jf983j309jfyrw
```

In the second ‘Received:’ line, “cola.bekkoame.or.jp” received the mail from a server that identified itself as “cola.bekkoame.or.jp” — instead of “ip21.san-luis-obispo.ca.pub-ip.psi.net,” which would be the expected identity. As a final test, the IP address mentioned in the third ‘Received:’ line was checked with an IP-tracing tool, which showed it did not match the address claimed for AOL's mail host (244.218.07.32).

Source: rahul.net/falk/mailtrack.html (archived methodology referenced in the original course material)

5.4 Illustration [Full Header]: Masked or Alias Email Address in Gmail

To always send from an alternate address, both the default “From” and “reply-to” address must be changed — changing only the “From” address means replies will default back to the original Gmail address. In the message composer, clicking the “From” line allows selecting the address to send from. This is a legitimate way for many people to send email via an alias, either to avoid future spam or to keep their primary address private for legitimate reasons — such aliases are not necessarily spam. Almost all spammers, however, also use alias addresses and never disclose the real one.

In the example preserved from the original course material, one support inbox was changed from its Gmail address to display as a simulated external recipient, while the sender's own address was masked to appear as an external-domain alias rather than the true originating address, and the IP addresses of the mail servers were likewise masked. The receiver in such a case sees only the alias or masked email address, never the real one. The composed text passes from the sender's computer to the local ISP/Exchange mail server first, then through each subsequent hop — each one logged with its own ‘Received:’ line, denoted by hop number in the header.

Reading a masked-alias header chain

The diagnostic pattern is the same regardless of how many hops are present: read the **Return-Path** for the true envelope sender, compare it against the visible **From:** line the recipient sees, and then walk the **Received:** chain from bottom to top. A legitimate alias (used for privacy, list management, or brand

consistency) will show a clean, internally consistent chain with matching SPF/DKIM results at the final corporate or ISP hop. A malicious mask, by contrast, typically shows a Return-Path domain that has no business relationship with the displayed From: name, combined with SPF or DKIM failures at the boundary — the same inconsistency pattern demonstrated in the two forgery walkthroughs above.

6. Finding Full Headers, and the Spam / Phishing / Spoofing / Scam Taxonomy

6.1 Finding Full Email Headers in Email Clients

Webmail

- **Horde —** Highlight the message and click the “+” sign (Other options) above the message body, then select View Source.
- **RoundCube —** Click the More button and select Show source.
- **SquirrelMail —** Open the message and go to View Full Header.

Microsoft Outlook 2010 / 2013 / 2016

- Open the message and click the File tab.
- Go to the Info tab and select Properties.
- Header information appears in the Internet headers box.

Mac Mail

- Open the message and go to View.
- Select Message, then All Headers.

Mozilla Thunderbird

- Open the message, go to More, then choose View Source.

Windows Live Mail, Windows Mail, or Outlook Express

- Right-click the message, then under Properties go to Details.

Gmail

- Open the message, click the arrow next to Reply at the top-right, then select Show Original.

Yahoo

- Open the message; under More select View Raw Message and the headers will appear.

6.2 Spam, Phishing, Spoofing, and Scams

Spam

Spam is unsolicited email, often used for marketing, advertising, or selling products and services, usually sent in bulk and indiscriminately to multiple mailing lists. There are two types of spam: intentional and unintentional.

- **Intentional spam** comes from spammers who are soliciting products or attempting to commit fraud.
- **Unintentional spam** originates from computers infected with a virus or worm that activates email-distribution processes in the background, attempting to send bulk messages without the computer owner's awareness.

Source: adaware.com/blog/the-big-three-email-nuisances-spam-phishing-and-spoofing

Phishing

Phishing is a special type of spam intended to trick the recipient into entering personal or account information, leading to account breach, identity theft, or fraud. Phishing attempts to steal personal information — identification details, credit card numbers, bank account details — to commit identity theft or fraud. A common example is an email claiming to be from a bank, asking the recipient to click a link and sign in to confirm account details. Instead of the bank's actual website, the link leads to a fake site that harvests the entered information.

Types of Phishing

USA.gov lists widespread phishing patterns reported across agencies and corporations:

- Emails from people you know claiming to be stranded abroad and asking you to wire money so they can travel home.
- Emails claiming to be from reputable news organizations capitalizing on trending news, asking recipients to click a link that leads to a malicious website.
- Emails claiming to be from organizations like the FTC or FDIC, referencing complaints filed or asking recipients to check their bank deposit insurance coverage.
- Emails threatening to harm recipients unless sums in the thousands of dollars are paid.
- Emails claiming to confirm a complaint filed by the recipient, designed to provoke a curious click on a malicious link or attachment.

This is not an exhaustive list — phishing emails can take any form. A few general tips for spotting them:

- Watch out for emails with generic greetings.
- Be cautious of emails requesting personal information.
- Be wary of messages conveying a sense of emergency or a tight response deadline.
- Always check the actual domain in any URL before clicking, especially if it lands on an activity or login page.

Spoofing

Spoofing, in general, is a fraudulent or malicious practice in which email is sent from an unknown source disguised as a source known to the receiver. Spoofing is most prevalent in communication mechanisms that lack a high level of security. Email spoofing occurs when imposters deliver emails by altering the sender information. Spoofing is also a favored mechanism for injecting malicious software, viruses, and malware.

Simple Mail Transfer Protocol (SMTP) does not provide any authentication process for the person sending an email, yet it remains the primary email transport system for most people, facilitating spoofing. Today, most email servers provide further security, and many vendors have created products remedying this gap. Another type of spoof masks the sender's IP address by hijacking a fake Internet Protocol address, making the message look as if it originates from an authentic source.

6.3 Common Scam Categories

Affinity Fraud

Targets groups such as social clubs and ethnic or religious communities, playing on the instinct to trust those who are like us. The scam artist may be an established group member, or may build relationships with influential members to gain acceptance. A common variant is the Ponzi or pyramid scheme.

Binary Options

An investor “bets” on whether the value of an asset will increase or decrease within a very short period (as short as a few minutes). A correct guess earns a predetermined payout; an incorrect one loses everything — an all-or-nothing situation akin to a game of chance. Binary options are also commonly sold through boiler room operations.

Boiler Room Scams

Cold callers use high-pressure sales tactics to sell shares in a company that may be fake, or a real company whose stock is artificially hyped. The pitch is usually tied to a sector currently in the news. By the time the victim realizes the deception, the scam artist has typically closed up shop and moved to a new scheme.

Emergency Scam / Grandparent Scam

The scam artist calls an older relative, pretends to be a family member in trouble, and requests money be sent secretly. Verifying the caller's identity and story independently before sending money is essential.

Employment Scams / Secret Shopper Scam

Job offers promising quick riches, demanding upfront payment, or asking the target to send and receive money through their own bank account are red flags. A common version sends a counterfeit cheque and asks the victim to spend part of it at a specific store, then wire back the remainder — the cheque later bounces, leaving the victim liable for the full amount.

Exempt Securities Scam / Accredited Investor Scam

Starts with an unsolicited pitch to invest in a business about to go public, often claiming the opportunity is reserved for very wealthy people with an exception made for the target — typically requiring the victim to misrepresent their income on paperwork.

Foreign Exchange (FOREX)

Promoted as an exciting opportunity to invest in the foreign-exchange market. Often the money is never actually invested in anything and is simply stolen by the scam artist; some variants claim to hold purchased gold or other valuable assets in a “secure vault” that does not exist.

Gift Card Scams

A scammer impersonating a government department threatens jail time or legal action unless payment is made immediately via prepaid gift cards, often instructing the victim to visit multiple stores and read the card codes back over the phone.

Identity Fraud and Identity Theft

Identity fraud is when someone uses another person's personal information to carry out fraud; identity theft is the underlying act of taking and collecting that information for criminal purposes, which can occur by phone, fax, regular mail, or online. Red flags include unexpected credit applications, collection notices for accounts never opened, and missing mail or statements.

Mortgage Fraud

Lying on a mortgage application, or applying on someone else's behalf as a “straw buyer,” is the most common type — leaving the straw buyer liable for the mortgage and potentially criminally responsible for the misrepresentation.

Offshore Investment

Promises large profits by moving money “offshore,” often framed as a tax-avoidance strategy. Once funds leave domestic jurisdiction and oversight, recovering them — or any promised tax savings, which are frequently fictitious — becomes extremely difficult.

Oil and Gas or Natural Resources

Scammers follow energy-price headlines and build a pitch around fake oil, gas, or renewable-energy ventures, often citing fabricated “expert geologist reports” for drilling sites and headquarters that do not exist.

Pension Scam

Promotes investment seminars claiming to help attendees “move your money” or “pay less tax,” frequently tied to a tax shelter that later triggers audits, back taxes, interest, and penalties.

Source: fcnb.ca/common-investment-scams.html

7. Case Study — Examining a Phishing Email (Fictitious Lotto Scheme)

This worked example, preserved from the original course material, walks a real header end-to-end through five independent verification tools and shows how their results corroborate one another. It remains one of the clearest illustrations of why no single tool should be trusted in isolation — cross-referencing is the discipline that makes header analysis defensible.

7.1 The Header Sorter and Organizer Tool

7.2 Raw Header — Delivered-To: cybersearchers@gmail.com

Field	Value
Return-Path:	<www.@bridge.ocn.ne.jp>
Received (final hop):	from mbkd0305.ocn.ad.jp [153.149.236.6] by mx.google.com; Mon, 04 Dec 2017 06:46:59 -0800 (PST)
Received-SPF:	pass (google.com: domain of www.@bridge.ocn.ne.jp designates 153.149.236.6 as permitted sender)
Date:	Mon, 4 Dec 2017 23:46:43 +0900 (JST) — 12/04/2017 14:46:43 UTC
From:	Mrs Maria Cook <"www."@bridge.ocn.ne.jp>
Reply-To:	Mrs Maria Cook <moneygramremittance1@gmail.com>
Subject:	President International Lotto Commission.
X-Originating-IP:	[41.86.234.171]

Two details immediately stand out even before consulting external tools. First, the From: address (“www.”@bridge.ocn.ne.jp) is a malformed, almost machine-generated-looking local part — not the kind of address a real person named “Mrs Maria Cook” would actually use. Second, and more tellingly, the Reply-To: address points to a completely unrelated Gmail account (moneygramremittance1@gmail.com) whose name directly references a money-transfer service. This is the exact pattern described in Section 3.2: the From: line is decorative, while the Reply-To: line is where the scammer actually wants to receive a response — in this case, quite literally, where they want to receive a wire transfer.

7.3 Hop-by-Hop Reconstruction

Reading the Received lines in delivery order (the reverse of how they print in the raw header) reconstructs the path:

Hop	Sender	Received By	Notice
1	mz-ukg626p.ocn.ad.jp	vcwebmail.ocn.ad.jp	Host names match — clean
2	vcwebmail.ocn.ad.jp	ntt.pod01.mv-mta-ucb021	Sender possibly incorrect
3	mv-mta-ucb021.ocn.ad.jp	mf-smf-ucb003.ocn.ad.jp	Host names not consistent (likely NAT)
4	mf-smf-ucb003	mf-smf-ucb003.ocn.ad.jp	Host names match — clean
5	mf-smf-ucb003.ocn.ad.jp	mbkd0305.ocn.ad.jp	Host names match — clean
6	mbkd0305.ocn.ad.jp	mx.google.com	Sender possibly incorrect

The warnings at hops 2, 3, and 6 are consistent with network address translation inside the same ISP (OCN, a Japanese provider) rather than outright forgery — the chain stays internally consistent until the very last hop, where the message claims an X-Originating-IP of 41.86.234.171. That single IP is the thread the remaining tools pull on.

7.4 Cross-Verification with Four Independent Tools

Each of the following tools was queried independently with the same originating IP address, 41.86.234.171. Reproduced together, they demonstrate how reputation databases, geolocation services, and blocklist registries triangulate the same conclusion from different data sources.

Tool 1 — Talos Intelligence (Reputation Center)

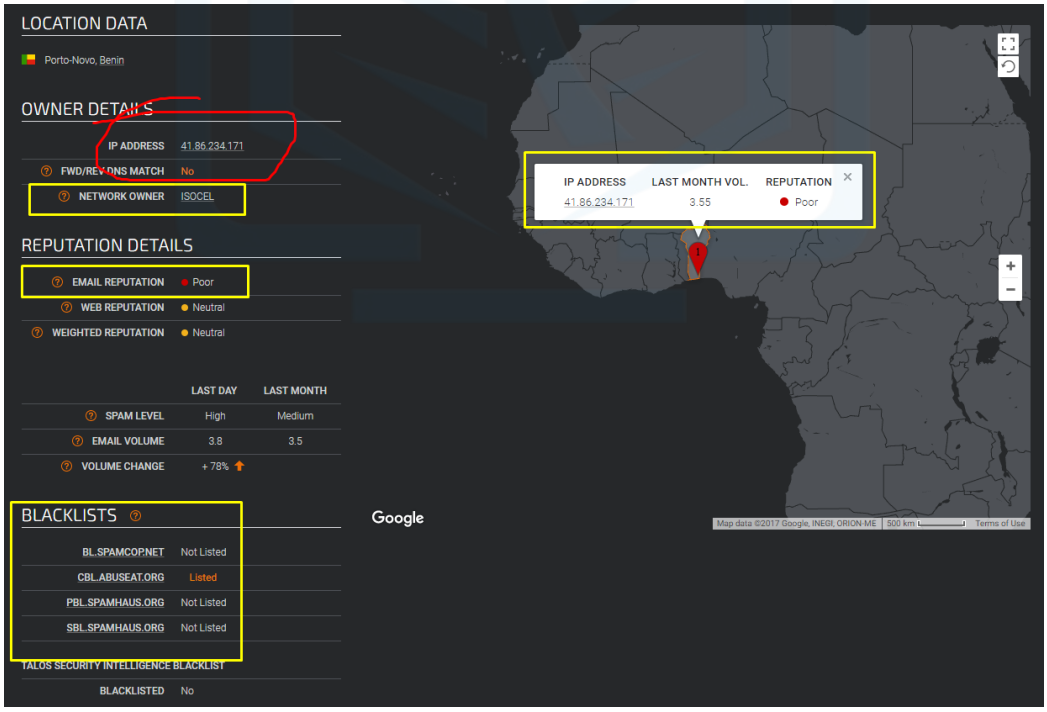

Figure 7.1 — Talos Reputation Center lookup for 41.86.234.171: location resolves to Porto-Novo, Benin; network owner ISOCEL; email reputation "Poor"; spam level "High" for the last day with a +78% volume spike, and the IP is listed on the CBL.ABUSEAT.ORG blocklist.

Tool 2 — IPTrackerOnline Full Geolocation Report

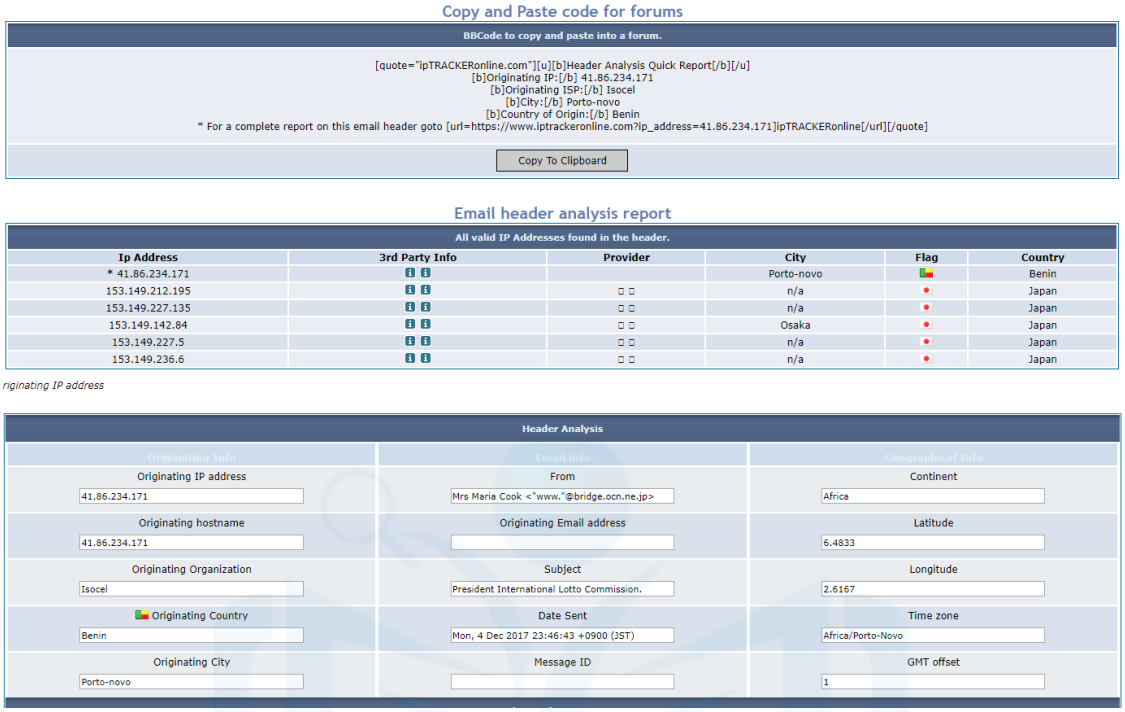

Copy and Paste code for forums

BBCode to copy and paste into a forum.

[quote="ipTRACKERonline.com"][u][b]Header Analysis Quick Report[/b][/u]
[b]Originating IP:[/b] 41.86.234.171
[b]Originating ISP:[/b] Isocel
[b]City:[/b] Porto-novo
[b]Country of Origin:[/b] Benin
" For a complete report on this email header goto [url=https://www.iptrackeronline.com?ip_address=41.86.234.171]ipTRACKERonline[/url][/quote]

Copy To Clipboard

Email header analysis report

All valid IP Addresses found in the header.

Ip Address	3rd Party Info	Provider	City	Flag	Country
* 41.86.234.171			Porto-novo		Benin
153.149.212.195			n/a		Japan
153.149.227.135			n/a		Japan
153.149.142.84			Osaka		Japan
153.149.227.5			n/a		Japan
153.149.236.6			n/a		Japan

riginating IP address

Header Analysis

Originating Info	Email Info	Geographical Info
Originating IP address	From	Continent
41.86.234.171	Mrs Maria Cook <"www."@bridge.ocn.ne.jp>	Africa
Originating hostname	Originating Email address	Latitude
41.86.234.171		6.4833
Originating Organization	Subject	Longitude
Isocel	President International Lotto Commission.	2.6167
Originating Country	Date Sent	Time zone
Benin	Mon, 4 Dec 2017 23:46:43 +0900 (JST)	Africa/Porto-Novo
Originating City	Message ID	GMT offset
Porto-novo		1

Figure 7.2 — IPTrackerOnline's header-analysis report independently confirms the originating city (Porto-Novo), ISP (Isocel), continent (Africa), and precise coordinates (Lat 6.4833, Lon 2.6167), alongside the same From:/Subject:/Date fields extracted directly from the header.

Tool 3 — Satellite Map Confirmation

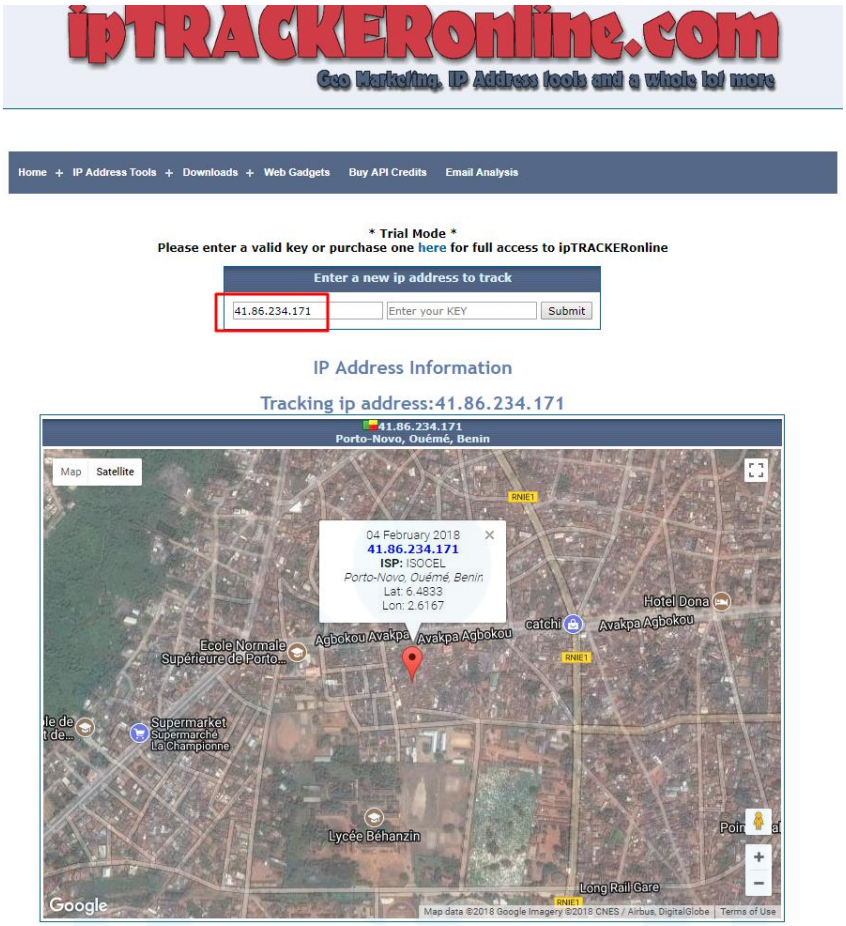

Figure 7.3 — A satellite-view geolocation pin for the same IP and coordinates, placing the server within a residential block in Porto-Novo, Benin — consistent with the city-level result returned independently by Talos and IPTrackerOnline above.

Tool 4 — Spamhaus SBL Advisory

Figure 7.4 — The Spamhaus Block List advisory for the same IP range (41.86.234.0–41.86.234.255, network ISOCEL-SA, country BJ) explicitly labels the listing reason as "intense advance fee fraud spam origin," and reproduces a near-identical header pattern from an unrelated complaint — the same IP being reused across multiple fraudulent campaigns.

Why convergent evidence matters

No single tool above is, by itself, proof of fraud — a reputation score can be wrong, a geolocation database can be stale, and a blocklist entry can occasionally catch a shared or dynamically reassigned IP. What makes this case defensible is that **four independently operated, differently sourced databases converge on the same network, the same country, and the same fraud classification (advance-fee / "419" lottery fraud)**. This convergence principle — never relying on a single source — is the core discipline taught throughout the CIRS curriculum (Section 12) and is the same standard applied to the skip-tracing and IP-geolocation workflows introduced later in this paper.

8. Reputation Intelligence, Reverse DNS, and Disposable Addresses

8.1 Email Header MTA Servers — "Email Hops/Path Analysis Tool"

In the search documented in Section 7, the tool performed MTA path analysis and showed that an email which appeared to come from Japan was, in fact, traced to a server in Benin — a French-speaking West African nation. This detailed path analysis is what surfaces the true origin of the email sender, with location.

8.2 Sender's MTA Server Physical Location Trace with IP

8.3 Email Servers Reputation Intelligence

For the IP address 41.86.234.171 examined in Section 7, the email server's reputation reads “Poor.” Reputation-management organizations such as Spamhaus.org, Abuseat.org, and SpamCop.net should also be checked to corroborate the current reputation of an email server tied to a suspect IP.

8.4 Spamming Reputation Authority Check

In the case study, the IP 41.86.234.171 was considered a spamming server, with its DNS entry (dnsbl1.dnbl1.borderware.com) appearing on a blocklist used by WatchGuard appliances and firewall software.

8.5 Spammers Reputation Database Checks

8.6 Check DNS Mapping with IP via “Reverse DNS Lookup” or DNS Authentication

8.7 Random and Temporary Email Addresses for Spamming

Temporary Random Email Address

- Use ‘any-name-of-your-choice’@yopmail.com to sign up anywhere; mail sent to that address is instantaneously created at the YOPmail server.
- Check the disposable inbox at yopmail.com, or directly via yopmail.com?name-of-your-inbox.
- Inboxes are not password protected, so choosing a unique inbox name (or using the Address Generator) is advisable; an inbox's email alias can also be used for extra security.
- Messages are kept for 8 days and can be manually removed; to prevent abuse, sending email from YOPmail itself is prohibited, though anonymous email to another YOPmail address is supported.

The double-edged nature of disposable addresses

Disposable email addressing (DEA) is a legitimate privacy technique: a unique address used for every contact means that if any single address is compromised or abused, it can be cancelled without affecting other contacts. The same property, however, makes disposable addresses attractive to spammers and fraudsters covering their tracks — which is precisely why most modern email-verification software flags disposable-domain detection as a standard check (see Section 9). A **disposable-domain hit is a risk signal, not proof of malice**; it should be combined with the other behavioral and authentication signals discussed throughout this paper before reaching a conclusion.

9. Modern Verification Methods: Authentication Protocols and AI

Everything covered so far — reading the Received chain, comparing From: against Return-Path, checking reputation databases — remains foundational. But three structural developments since the original course material was written have changed how verification is actually performed at scale in 2026: cryptographic sender authentication has become close to mandatory for high-volume senders, machine-learning classifiers have replaced static rule sets as the primary spam-detection mechanism, and the same generative AI used to write phishing emails is now used defensively to detect them.

9.1 SPF, DKIM, and DMARC — From Optional to Effectively Mandatory

Three DNS-published authentication records now work together to let a receiving server verify, mechanically, whether a message is permitted to claim a given domain:

Protocol	What it verifies	Where it lives
SPF	Which mail servers (by IP) are authorized to send for a domain	A TXT record on the sending domain
DKIM	That the message body was not altered in transit, via a cryptographic signature	A DKIM-Signature header plus a public key in DNS
DMARC	What the receiving server should do if SPF or DKIM fails (none / quarantine / reject), and where to send aggregate failure reports	A DMARC TXT record at _dmarc.domain.com

As of 2024–2025, Google and Yahoo began enforcing DMARC compliance for any sender pushing high volumes of mail into their inboxes, and PCI DSS v4.0 made DMARC implementation a compliance requirement for organizations processing payment cards. The practical effect is that the Authentication-Results header demonstrated in Sections 4 and 5 of this paper — spf=pass, dkim=pass — is no longer a nice-to-have detail buried in the header; it is frequently the single fastest line in the entire message to check first.

> **BIMI — the newest layer**
>
> Brand Indicators for Message Identification (BIMI) lets a domain with DMARC enforced at p=quarantine or p=reject display a verified brand logo next to the message in supporting inboxes. Because BIMI requires DMARC enforcement as a prerequisite, a recipient seeing the verified logo has a visual signal that the underlying authentication chain has already passed — useful at a glance, though, like any visual trust indicator, it should never be the only check performed on a message that is otherwise behaving suspiciously.

9.2 How AI Spam and Phishing Filters Actually Work in 2026

Gmail, Outlook, and major enterprise gateways now rely primarily on transformer-based language models analyzing message content for spam patterns, phishing indicators, and content-quality signals, combined with engagement-based sender-reputation-per-recipient models that track how each individual user has historically interacted with mail from a given sender. The same email can land in one

person's inbox and another person's spam folder, purely based on each recipient's own engagement history with that sender.

The defensive use of AI has had to escalate because the offensive use of AI escalated first. Independent research using the ChatSpamDetector methodology found that GPT-4-class models can achieve detection accuracy above 99% on phishing samples — but only when those models are actually deployed in the filtering pipeline. Most legacy enterprise filters, trained primarily on the linguistic fingerprints of older, human-written phishing campaigns, are demonstrably less effective against phishing emails that have been rephrased by an LLM, since the grammatical tells those filters were trained to spot are no longer present.

Stylometric and Behavioral Detection

A particularly difficult 2026-era pattern involves attackers using AI to analyze a victim's stolen mailbox history, then generating a message in that exact tone and vocabulary that inserts itself into an ongoing email thread at a plausible moment — rather than sending a generic cold lure. Because the resulting message uses the real sender's actual writing style and references genuine internal context, it can be indistinguishable from legitimate correspondence on content grounds alone. The defensive countermeasure is stylometric detection: comparing the statistical fingerprint of a message's word choice and sentence structure against a baseline built from that sender's prior known-good correspondence, flagging deviations even when the grammar itself is flawless.

The core implication for a human investigator

Content-based judgment — “does this sound like a scam?” — is a steadily weakening signal. The header-level techniques in Sections 2 through 8 of this paper (Received-chain analysis, Return-Path vs From: comparison, IP reputation, authentication results) verify the **infrastructure** a message traveled through, which is mechanically generated and far harder for an attacker to fully fabricate than fluent prose. This is why those techniques, despite being decades old, remain the most durable layer of verification available, AI-era phishing notwithstanding.

9.3 AI-Powered Verification in Commercial Software

Beyond filtering inbound mail, a parallel category of software verifies outbound mailing-list quality before a campaign is even sent — the bulk and real-time email-verification tools covered in Section 10. Several of the leading 2026 platforms now layer machine-learning scoring on top of the traditional syntax/MX/SMTP checks:

- **Predictive deliverability scoring.** Rather than a binary valid/invalid result, tools such as Clearout's “AI Verdict” and ZeroBounce's “ZeroBounce Score™” return a 1–10 (or percentage) prediction of how likely an address is to actually engage, letting a sender prioritize which borderline contacts to mail first.
- **Catch-all and accept-all disambiguation.** Domains configured to accept mail for any address at that domain (“catch-all”) defeat traditional SMTP pinging; AI models trained on historical bounce data estimate a likely-valid probability even when the mail server itself gives an ambiguous answer.
- **Toxicity and spam-trap detection.** Platforms such as ZeroBounce maintain proprietary databases of known spam-trap and abuse addresses, updated continuously from partner network data, and

apply pattern-matching that flags an address as toxic before it ever damages a sender's domain reputation.

Section 10 below catalogues the specific software platforms a researcher or analyst is likely to encounter in 2026, organized by the task they perform.

10. The Software Landscape: What Is Used for Email Verification Today

The original course material catalogued a handful of email-finding and verification tools current in the late 2010s — Hiretual, Voila Norbert, Sellhack, Find That Email, Hunter, and Rapportive, alongside Spokeo, PeopleSmart, and Pipl for reverse lookups. Several of those names persist in 2026 (Hunter in particular remains a leading platform), but the category has both consolidated and specialized. The current landscape splits into four functional groups, covered in turn below: bulk and real-time list verification, lead-generation email finders, skip-tracing and people-search databases, and IP-geolocation services.

10.1 Bulk and Real-Time Email Verification Platforms

These tools take a single address or an uploaded list and run it through syntax checking, MX-record lookup, SMTP mailbox-existence verification, disposable/role-account detection, and (increasingly) AI-based deliverability scoring, returning a verdict of valid, invalid, risky, or unknown for each address.

Platform	Distinguishing feature	Typical pricing (2026)
ZeroBounce	AI-driven 1–10 ZeroBounce Score™; scans for 30+ risky-address types including spam traps and abuse accounts; activity data (opens, last-seen)	From $0.0003/email at high volume; subscriptions from $18/mo
Bouncer	Proprietary AI negotiates directly with SMTP servers; very high accuracy in independent benchmarks; toxicity flagging	From $0.007/email; subscriptions from $50/mo
Clearout	“AI Verdict” predictive deliverability score rather than simple pass/fail; 38 native CRM integrations	From $0.007/email; subscriptions from $14/mo
Hunter	Combines email finder, domain search, and verifier in one workflow; polished Chrome extension and Sheets add-on	Free tier 50 credits/mo; paid from $49/mo
Kickbox	Proprietary “SendEx™ Score” alongside standard checks; deliverability consulting add-on	100 free credits; pay-as-you-go available
NeverBounce	99.9% accuracy guarantee with money-back policy if a verified address bounces within 90 days	Pay-as-you-go from $0.008/email
Emailable	Fastest processing in independent 2026 benchmarks (~0.012 sec/email); scheduled recurring list monitoring	100 free credits; from $0.008/email

Platform	Distinguishing feature	Typical pricing (2026)
MillionVerifier	Lowest-cost bulk cleaner widely cited for budget-constrained, high-volume list cleaning	Among the cheapest per-email rates on the market

How these checks actually work, in sequence

1. Syntax validation — confirms basic formatting (the @ symbol, valid characters, correct domain structure).
2. Domain / MX validation — confirms the domain exists and has active mail-exchange records configured to receive mail.
3. Mailbox verification — pings the recipient's SMTP server to confirm the specific mailbox exists, without actually delivering a message.
4. Risk-category checks — flags disposable-domain addresses, role-based accounts (info@, support@), and known complainers or spam-trap matches.
5. Catch-all disambiguation — for domains accepting any address, an AI model estimates a likely-valid probability from historical bounce patterns.

10.2 Email-Finder and Lead-Generation Tools

These are the direct descendants of the Hiretual / Voila Norbert / Sellhack category from the original course notes — tools that locate a probable email address for a named person at a known company, then verify it.

- **Hunter.io —** domain search returns as many verified addresses for a target company as the tool can find, alongside its standalone finder and verifier.
- **Snov.io —** combines a Chrome-extension email finder with drip-campaign and CRM functionality aimed at sales teams.
- **Apollo.io and ZoomInfo —** enterprise-grade contact-and-company databases that bundle email finding with firmographic and intent data, typically licensed at the organizational rather than per-credit level.

10.3 Skip Tracing and People-Search Databases

Skip tracing — locating a person who cannot be found through ordinary means, originally a debt-collection term for someone who has “skipped town” — is the natural extension of email verification once an address alone is not enough. These platforms aggregate credit-header data, public records, utility records, and proprietary data partnerships to resolve an email address, name, or partial identifier into a full current-location and contact profile.

Platform	Access tier	Notes
TLOxp (TransUnion)	Licensed professionals only	Often called the “gold standard” for licensed investigators; over 100 billion records including 11 billion name/address combinations; access to credit-header data unavailable elsewhere
LexisNexis Accurint	Licensed professionals only	Access to 84+ billion public records; widely used by legal, insurance, and law-enforcement workflows

Platform	Access tier	Notes
IDI Core / idiCORE	Professional, pay-as-you-go	Cloud-native; merges billions of data points from thousands of sources; near-complete coverage of the U.S. adult population; frequently cited as a more affordable alternative to TLO
IRBsearch / Tracers	Professional, subscription	Consistently rated among the top three most accurate databases in independent investigator surveys; strong batch-processing support
BatchData	No license required	Real-estate-oriented; pay-per-match pricing; faster API onboarding than the licensed-tier platforms, at a lower per-record accuracy ceiling
BeenVerified / PeopleFinder	Consumer, no license	Accessible to the general public; useful starting point, but materially less complete than the licensed-tier databases above

Legal and ethical boundary

The licensed-tier platforms above (TLOxp, Accurint, IDI Core, CLEAR) restrict access to verified private investigators, attorneys, law enforcement, and similarly credentialed professionals, precisely because the underlying data includes credit-header and other regulated information governed by the Fair Credit Reporting Act and similar statutes in other jurisdictions. A CIRS-trained researcher operating in this space is expected to know which tier of data a given task legitimately requires, to use the least invasive source sufficient for the purpose, and to never represent themselves as having an access tier they do not hold.

10.4 IP Geolocation Services

These services resolve the public IP addresses found in an email header (Section 3.3) into approximate physical location, ISP, and network-ownership data — the same underlying function performed manually by IPTrackerOnline in the Section 7 case study, now available as production-grade APIs for integrating into larger workflows.

Platform	Strength	Consideration
MaxMind GeoIP	Industry-standard accuracy; covers 99.9999% of IP addresses; downloadable local databases avoid per-query latency	ISP-name accuracy averages ~95% in the U.S. but drops to ~80% outside it
IPinfo	Daily-updated enriched metadata (ASN, carrier, abuse contacts, hosting); “Probe Network” latency-based trilateration	Premium pricing; free tier limited to country-level accuracy as of 2025
IPLocate.io	Competitive accuracy at a meaningfully lower price point than MaxMind or IPinfo; generous free tier	Smaller company footprint than the two incumbents above
ipstack	Real-time IPv4/IPv6 support with security and currency-detection modules	Best suited to lighter-weight integrations rather than forensic-grade investigation

A consistent limitation across every IP-geolocation provider, worth stating plainly: the returned coordinates describe the location of the network infrastructure serving that IP, not necessarily a precise street address or household, and accuracy degrades meaningfully whenever a VPN, proxy, or large

corporate NAT gateway sits between the true sender and the IP visible in the header. Section 11 below works through how to read around that limitation.

11. Applied Questions: Spam, Solicitations, Remarketing, Fake Emails, Skip Tracing, and IP Location

This section addresses each of the practical questions a researcher is most often asked to resolve, naming the specific software or technique used in each case, with a short illustration drawn from the case study and statistics covered earlier in this paper.

11.1 Is This Email Spam?

Illustration

In Section 4.2's worked example, a message scoring hits=10.6 against a required threshold of 7.0 was automatically classified Yes under X-Spam-Status, driven by test flags such as BAYES_56, EXCUSE_10/14/22/23/24, and OPT_OUT — all consistent with a templated bulk-marketing blast rather than a legitimate one-to-one message.

11.2 Is This an Unwanted Solicitation?

A legitimate solicitation should pass SPF/DKIM/DMARC for the domain it claims to represent, include a real unsubscribe path, and a real postal address in the footer. The absence of any one of these is not automatically fraud, but it removes the legal protection that distinguishes a legitimate marketer from a spammer, and is the threshold most spam-complaint and blocklist services use to escalate a sender for review.

11.3 Is This Remarketing or Retargeting Email Legitimate?

11.4 How Do I Tell a Fake Email Address from a Real One?

Illustration — a malformed address that looked plausible at a glance

The case study in Section 7 used the From: address *“www.”@bridge.ocn.ne.jp* — syntactically valid enough to pass a basic format check, but the local part (“www.”) is not a pattern any real individual uses for a personal mailbox. A syntax-only check would have passed this address; an SMTP-level mailbox-existence check, or simply noticing the mismatch against the claimed sender name “Mrs Maria Cook,” catches what the syntax check alone misses.

11.5 Skip Tracing the Individual Behind an Address

A practical workflow: reverse-lookup the email address first (Spokeo/Pipl) to surface a candidate name and any linked social profiles; verify that candidate name independently through a second, differently sourced database; and only then, if the matter warrants it, escalate to a licensed-tier platform for a full current-address and contact-history profile.

11.6 Identifying the IP Location Behind a Message

Illustration — when the trail goes through a VPN or relay

If a header's final public IP resolves to a known commercial VPN or hosting-provider range rather than a residential or recognizable corporate ISP block, the geolocation result describes the exit node, not the sender's real location. The practical signal is in the network-owner field returned alongside the coordinates (compare the **“ISOCEL” / “ISOCEL-SA” network-owner result** in Figures 7.1 and 7.4, which correctly identified a real regional ISP rather than an anonymizing service) — a network owner field reading as a generic data-center or VPN brand is the cue to treat the location as the relay's location, not the originator's, and to look for additional Received-chain hops further upstream.

12. Real-World Examples and a Complete Use Case

12.1 Business Email Compromise: Scale and Recent Cases

Business Email Compromise (BEC) is the direct, high-stakes descendant of the From:/Reply-To: mismatch technique covered in Section 5 — a criminal impersonates a trusted executive, vendor, or partner by email and persuades an employee to redirect a payment. The FBI's Internet Crime Complaint Center reported that BEC schemes cost U.S. organizations over $16.6 billion in a single recent year across more than 256,000 complaints with actual financial loss, an average of roughly $129,000 per incident — making BEC more expensive than ransomware and data breaches combined.

Case	Mechanism	Outcome
Google and Facebook (disclosed 2017 prosecution)	Attacker impersonated a real hardware-supplier vendor and sent matching invoices to accounts-payable staff at both companies	Over $100 million collectively lost before detection
Massachusetts workers' union (Jan. 2023)	Fraudsters impersonated the union's own investment manager, requesting a new account for a scheduled $6.4 million transfer	$5.3 million seized by federal investigators; remainder laundered through Asia-based accounts and crypto exchanges before recovery
Town of Arlington, MA (June 2024)	Scammers hijacked an existing construction-invoice email thread and redirected four separate payments	Nearly half a million dollars lost from municipal funds
New South Wales government department, Australia (2024)	Attackers posed as a trusted vendor and altered banking details on a routine payment	AU $2.1 million wired to the fraudulent account

A consistent pattern across these cases: each began with a legitimate, ongoing business relationship that the attacker either impersonated convincingly or, in increasingly frequent 2024–2025 cases, genuinely compromised via Vendor Email Compromise — attacks against a real vendor's actual mailbox rose sharply in this period, because a message sent from a truly compromised account is, by definition, authenticated and passes every header check described in Sections 4 and 5 of this paper. This is the limit case the rest of the paper has been building toward: when the infrastructure itself is compromised, header analysis alone cannot catch the fraud, and procedural controls — a callback to a verified phone number before changing any payment instruction — become the necessary final layer.

Why generative AI has accelerated BEC specifically

Verizon and Hoxhunt's 2025–2026 monitoring places the share of BEC phishing lures that are AI-generated at roughly 40%, and overall BEC attack volume rose sharply year over year as generative tools removed the labor bottleneck that previously capped how many highly personalized, well-researched lures a single attacker could produce. Where a convincing spear-phishing email once required hours of manual research and drafting per target, an LLM can now read a target company's public filings, press releases, and social media, then draft an executive-toned request in seconds — which is precisely the threat Section 9.2's discussion of stylometric defenses is designed to counter.

12.2 Complete Use Case — Verifying a Suspicious Vendor-Payment Email

The following use case walks through the full verification workflow described in this paper from end to end, applied to a realistic scenario combining the original header-analysis techniques with the current software catalogued in Section 10.

Scenario

An accounts-payable clerk at a mid-sized construction firm receives an email that appears to come from a long-standing supplier's project manager, requesting that the remaining balance on a current invoice be wired to “our new escrow account” due to an alleged banking-system migration. The tone, formatting, and signature block are all consistent with the supplier's normal correspondence —

indistinguishable from genuine correspondence on content alone, exactly the AI-era pattern described in Section 9.2.

Step 1 — Header extraction

Following the client-specific instructions in Section 6.1, the clerk (or the firm's IT/security analyst) pulls the full header rather than relying on the partial header shown by default in the mail client.

Step 2 — Authentication-result check

The Authentication-Results header is read first (Section 9.1). In this scenario, SPF and DKIM both pass — but for a domain that, on close reading, is off by one character from the supplier's real domain (a classic homoglyph or “typosquat” domain registered by the attacker, who then legitimately configured SPF/DKIM for their own look-alike domain). This is the critical lesson reinforced by the From:/From: distinction in Section 5.1: authentication passing proves the message came from where it claims, not that where it claims is who you think it is.

Step 3 — Reply-To and Return-Path comparison

As in the Section 7.2 case study, the Reply-To: address is compared against the visible From: name. Here, Reply-To: routes to a freemail address with no relationship to either the real or the look-alike domain — the same red flag pattern as the original lottery-scam example's moneygramremittance1@gmail.com Reply-To address.

Step 4 — IP and reputation cross-check

The originating IP in the final Received hop is run through two of the Section 10.4 geolocation tools and one Section 8 reputation database. The IP resolves to a data-center hosting range in a country with no business relationship to either party, rather than the supplier's known corporate ISP block — following the exact same convergent, multi-tool discipline established in Section 7.4.

Step 5 — Domain verification with bulk-style tooling

Running the look-alike domain through a verification platform from Section 10.1 confirms the domain was registered only weeks earlier — in contrast to the genuine supplier's domain, which has years of sending history and an established reputation score.

Step 6 — Procedural control (the step header analysis cannot replace)

Consistent with the limitation acknowledged in Section 12.1, the clerk does not rely on the header analysis alone to clear the payment. A callback is placed to the supplier's project manager using a phone number drawn from a prior, independently verified contact record — not any number contained in the suspicious email itself — confirming no banking change was requested.

Outcome

The combination of a one-character domain mismatch, a Reply-To address with no organizational relationship, a freshly registered domain, and a geolocation result inconsistent with the supplier's known infrastructure was sufficient to hold the payment pending the callback, which confirmed the fraud attempt. No single check was individually conclusive; the convergence of independent signals — the same discipline demonstrated throughout Section 7 — is what made the determination defensible.

13. Building Email Search Queries

13.1 When You Know the Name of a Person

13.2 When You Know the First Name but Not the Last

13.3 When You Don't Know the Exact Name, Alias, or Abbreviation Used

Where the name (or part of it) is known, building combinations of email addresses that include parts of the name — first/last initials, or other permutations within a given domain — uses a technique called an Email Permutator.

Try different domains, starting with Yahoo, Gmail, Hotmail, and MSN, then ISPs such as AT&T, Bell, and Rogers, and finally any known business domain. For a hypothetical “Naveed Manzoor,” the permutator pattern produces a structured list such as:

Pattern	Example	Pattern
first-last@domain	naveed-manzoor@yahoo.com	first.last@domain
f-last@domain	n-manzoor@yahoo.com	flast@domain
first-l@domain	naveed-m@yahoo.com	first@domain
f-l@domain	n-m@yahoo.com	last@domain
last-first@domain	manzoor-naveed@yahoo.com	f.last@domain
last-f@domain	manzoor-n@yahoo.com	firstlast@domain

The same structure repeats with underscores in place of hyphens (naveed_manzoor@yahoo.com, n_manzoor@yahoo.com, and so on) and with no separator at all (naveedmanzoor@yahoo.com, nmanzoor@yahoo.com). With these combinations in hand, adding re-enforcer keywords — the person's company name, address or city, club, or professional membership — narrows the candidate list to the addresses worth running through a verification tool from Section 10.1.

13.4 Using Search Query Operators as in Google

For example, when looking for the email address of “Paul Myers,” an investigative journalist working with the BBC in the UK, the re-enforcer terms “BBC” and “journalist” sharply narrow an otherwise common name.

Assorted sample queries (not all individually tested) worth trying, combining operators such as AND, OR, the minus sign, parentheses, the wildcard (*), site:, inurl:, intext:, related:, info:, and the hashtag for social search:

- [“First Last name”@]
- [“first last name” linkedin]
- [“first last name” facebook]
- [“first last name” @twitter]
- [site:linkedin.com Paul Daniels]
- [site:facebook.com Geoffrey Archer]
- site:companywebsite.com [full name] + email
- site:companywebsite.com [full name] + contact
- site:companywebsite.com “first.last”@companyname.com
- site:personalblog.com “full name” email
- [FirstName]@example.com
- [FirstName].[LastName]@example.com
- [FirstInitial].[LastName]@example.com

14. Why the CIRS Course Is Ideal for Learning These Techniques

The original course material this paper builds on was itself developed for the Certified Internet Research Specialist (CIRS™) program, run by the Association of Internet Research Specialists (AOFIRS) — a fact preserved in the document's own internal references to its support and training infrastructure. The fit between the CIRS curriculum and the specific skill set developed across the previous thirteen sections of this paper is direct rather than incidental, for several concrete reasons.

14.1 The Curriculum Is Built Around Exactly This Methodology

CIRS describes its own value proposition as integrating human expertise, machine intelligence, data analysis, and ethical practice into a single research methodology, rather than teaching search as an isolated skill. That description matches the structure of this paper almost exactly: header forensics and protocol fundamentals (human expertise), AI-driven filtering and verification software (machine intelligence), the skip-tracing and IP-geolocation database layer (data analysis), and the licensed-tier access boundaries discussed in Section 10.3 (ethical practice).

14.2 It Teaches the Convergent-Evidence Discipline This Paper Repeatedly Relies On

The thread running through every section

Section 7.4's four-tool cross-verification, Section 11.5's two-database skip-tracing rule, and Section 11.6's VPN-detection check all rest on the same underlying discipline: **never conclude from a single source.** This is not an incidental habit — it is the explicit methodology CIRS coursework is built to instill, precisely because Internet research as a profession constantly confronts the problem of too many available sources of inconsistent reliability, and a credentialed standard exists to certify that a researcher consistently practices that discipline rather than stopping at the first plausible-looking answer.

14.3 The Program Covers the Full Breadth This Paper Has Walked Through

The current CIRS syllabus structure spans internet technology, websites and browsers, and research methods and online research as foundational modules — directly mirroring Sections 2 through 8 of this paper — while more recent course iterations explicitly add AI-driven research, data analytics, and advanced search strategies, mirroring Sections 9 through 12. A candidate who works through this paper has, in effect, walked through a condensed version of the program's own topic spine.

14.4 Recognized Standing, Not Just Self-Study

CIRS is positioned in the research-services industry as a professional certification carrying recognized weight, rather than a simple course-completion certificate — a meaningful distinction for any researcher whose findings may eventually be relied upon by an employer, a client, or in a legal context, exactly the kind of due-diligence standard this paper's introduction (Section 1) invokes. Program materials report that the large majority of certified professionals remain employed in the field, with a substantial share reporting a promotion following certification, and candidates who complete the certification agree to the Association's Professional Code of Conduct — a layer of accountability self-study alone does not provide.

14.5 Structured, Flexible Delivery for Working Researchers

The certification program is available both as a self-study package, built around a current edition of the CIRS Training Guide with unlimited practice-exam access, and as a live, instructor-led online training track, typically structured over a six-to-twelve-week period with virtual-classroom interaction across an international cohort of fellow students. No prior research experience is required to enroll, which makes the program equally suited to a complete newcomer working through this paper for the first time and to an experienced analyst seeking formal recognition of skills already practiced informally.

In short

Every technique in this paper — header forensics, authentication-protocol reading, AI-assisted verification, skip tracing, and IP geolocation — maps directly onto a module of the CIRS curriculum, and the certification's defining emphasis on cross-source verification is the same discipline that made the Section 7 case study and the Section 12.2 use case defensible. For a reader who found this paper useful, CIRS is the structured, recognized path to formalizing and extending exactly this skill set.

15. Appendix — Original Course-Era Software Reference (Preserved)

The following entries are preserved from the original course material as a historical and still-partially-current reference point. Several of these names remain active 2026 platforms (noted where applicable); others have been superseded by the tools catalogued in Section 10, but are retained here so the evolution of the category remains visible.

15.1 Email Search Software — For Leads Collection

Hiretual

Free: 15 credits per month; paid plans started at $59 per month. Worked across popular sourcing sites (LinkedIn, Twitter, Facebook) as a Chrome extension that found email addresses while also surfacing experience, seniority level, and compensation range, paired with an AI sourcing assistant that could run submitted sourcing tasks unattended.

Voila Norbert

Free: 50 leads; paid plans started at $0.10 per lead. A no-install corporate-email finder, though without catch-all address checking and with built-in query limits to avoid the service being blocked by SMTP servers.

Sellhack

Free: 10 credits per month; paid plans started at $19 per month. Designed for sales teams to build prospect and lead lists, featuring a 12-step email-verification engine that found the best real-time address for a given person.

Find That Email

Free: 50 search credits and 100 verify credits per month; paid plans started at $29 per month. Allowed individual or bulk-list lookups by entering a candidate's name and company domain, claiming a success rate above 90%.

Hunter (still active in 2026 — see Section 10.1 and 10.2)

Free: 150 credits per month; paid plans started at €39 per month. Offered an email finder, domain search, and email verification claiming 95% deliverability accuracy on verified addresses.

Rapportive

Free. Populated LinkedIn information into a sidebar inside the inbox; did not directly find emails, but supported a workaround of testing variations of a guessed address in the “To” field until the correct one populated a Rapportive profile.

15.2 Personal Information and Background Check Software

Spokeo Email Search

Helped discover the owner of an unknown email address and connect to additional linked information — name, phone numbers, addresses, and social networks attached to that account.

People Smart

Returned an available email owner's name and social profiles, useful for protecting against email scams or appending a contact list; checked blacklist databases to see whether an address had been reported as spam or blocked from sending messages.

Pipl

16. Conclusion

Email verification, at its foundation, has not changed since the protocols described in Section 2 were standardized: a message travels through a chain of mail-transfer agents, each of which leaves a mechanically generated record of the hop it handled, and that record is far harder to convincingly fabricate than the words inside the message. What has changed, comprehensively, is the sophistication of what rides on top of that infrastructure — AI-generated phishing content that evades legacy filters, AI-assisted skip-tracing and geolocation databases that resolve an address or IP to a real-world identity and location with far greater speed and reach than the manual tools available when the original course material was written, and authentication protocols (SPF, DKIM, DMARC, BIMI) that turned a previously optional good practice into a near-mandatory infrastructure layer.

The throughline connecting all sixteen sections of this paper is the discipline demonstrated first in the Section 7 case study and applied again in the Section 12.2 use case: no single tool, database, or signal should be trusted in isolation. A reputation score, a geolocation result, an authentication pass, and a skip-trace match each describe one slice of the truth; convergence across independently sourced tools is what turns a plausible hypothesis into a defensible conclusion. That principle, more than any individual piece of software named in Section 10, is the actual skill this paper has tried to teach — and it is, not coincidentally, the same principle around which the CIRS curriculum discussed in Section 14 is explicitly built.

References and Sources

Original course material sources (preserved):

intermedia.net Knowledge Base, “What does it mean if a message is flagged as spam?” kb.intermedia.net/Article/1783

IETF RFC 2076, “Common Internet Message Headers,” tools.ietf.org/html/rfc2076

Falk, B., “Tracking Spam,” rahul.net/falk/mailtrack.html

Ad-Aware, “The Big Three Email Nuisances: Spam, Phishing and Spoofing,” adaware.com/blog

Financial and Consumer Services Commission of New Brunswick, “Common Investment Scams,” fcnb.ca/common-investment-scams.html

Gaijin Online Email Header Analyzer, gaijin.at/en/olsmailheader.php

IP2Location Free Email Tracer, ip2location.com/free/email-tracer

IPTrackerOnline, iptrackeronline.com

Talos Intelligence Reputation Center, talosintelligence.com

WatchGuard ReputationAuthority, reputationauthority.org

The Spamhaus Project, spamhaus.org

Kloth Services Reverse DNS Lookup, kloth.net

Yopmail Disposable Email Generator, yopmail.com

ThrowAwayMail, throwawaymail.com

Current (2025–2026) sources consulted for new content:

Email Vendor Selection, “10 Best Email Verification Tools in 2026,” emailvendorselection.com/email-verification-tools

Hunter.io, “15 Best Email Verification Tools in 2026 (With Real Benchmark Data),” hunter.io/email-verification-guide

LaGrowthMachine, “7 Best Email Verification Software 2026: Accuracy & Speed,” lagrowthmachine.com

ZipDo Software Advice, “Top 10 Best Bulk Email Verification Software of 2026,” zipdo.co

BounceCheck, “12 Best Bulk Email Verification and Validation Services (2026),” bouncecheck.email/blog

Aerosend, “Top Email Verification Tools in 2026,” aerosend.io/cold-email

AccuWebHosting, “Top 10 Bulk Email List Verification / Validation Services Compared,” accuwebhosting.com/blog

WifiTalents, “Top 10 Best Bulk Email Verification Software of 2026,” wifitalents.com/best

Sender.net, “Best Email Verification Tools in 2026 (Compared and Ranked),” sender.net/blog

Sparkle.io, “12 Best Email Verification Tools From a Brutally Honest Test,” sparkle.io/blog

Debtor Inspector, “The 6 Best TLO Skip Tracing Alternatives [2026],” debtorinspector.com

BatchData, “How BatchData's Skip Tracing Compares to TLO and IDI,” batchdata.io/blog

TransUnion TLO, “Skip Tracing,” tlo.com/skip-tracing

Skip Trace Services, “Top 10 Databases Used in Professional Skip Tracing,” skiptraceservices.com/blog

VirtualLatinos, “Top 9 Best Lead Generation Skip Tracing Tools,” virtuallatinos.com/blog

BatchData, “Best Skip Tracing Tools for Bulk Data Processing in 2026,” batchdata.io/blog

OnCall Legal, “How to Skip Trace Ethically in 2025,” oncalllegal.com/how-to-skip-trace

PInow, “Investigators Share Which Databases They Prefer to Use,” pinow.com/articles

MaxMind, “GeoIP Databases and Web Services,” maxmind.com

IPLocate.io, “What is the best IP Geolocation API in 2026?” iplocate.io/blog

IPTrackerOnline, “Email Header Analyzer — Trace Sender IP & Detect Phishing,” iptrackeronline.com/email-header-analysis

AutoSPF, “AI-Powered Phishing in 2026,” autospf.com/blog

PowerDMARC, “Email Security Predictions 2026,” and “DMARC AI and the Evolution of Email Authentication,” powerdmarc.com

Medium / Generative AI, “How AI Spam Filters Actually Work in 2026,” medium.com

Mailbird, “AI-Generated Email Detection: Spam Filter Guide 2026,” getmailbird.com

EmailCloud, “How AI Spam Filters Work in 2026,” emailcloud.com/guides

Federal Bureau of Investigation, “Business Email Compromise,” fbi.gov

Chargebacks911, “Business Email Compromise Examples: 2026 Case Studies,” chargebacks911.com

Eye Security, “When Business Email Compromise Nearly Costs a Million,” eye.security/blog

Hoxhunt, “Business Email Compromise Statistics 2026,” hoxhunt.com/blog

Valimail, “The Complete Guide to Business Email Compromise (BEC) Attacks in 2025,” valimail.com/blog

Corvus Insurance, “How Business Email Compromise (BEC) Drives Claims,” corvusinsurance.com/blog

Association of Internet Research Specialists (AOFIRS), airsassociation.org and aofirs.org

AOFIRS, “Certified Internet Research Specialist (CIRS™),” aofirs.org/certifications

Online Research Education, “CIRS Certification,” onlineresearcheducation.com/certification.html

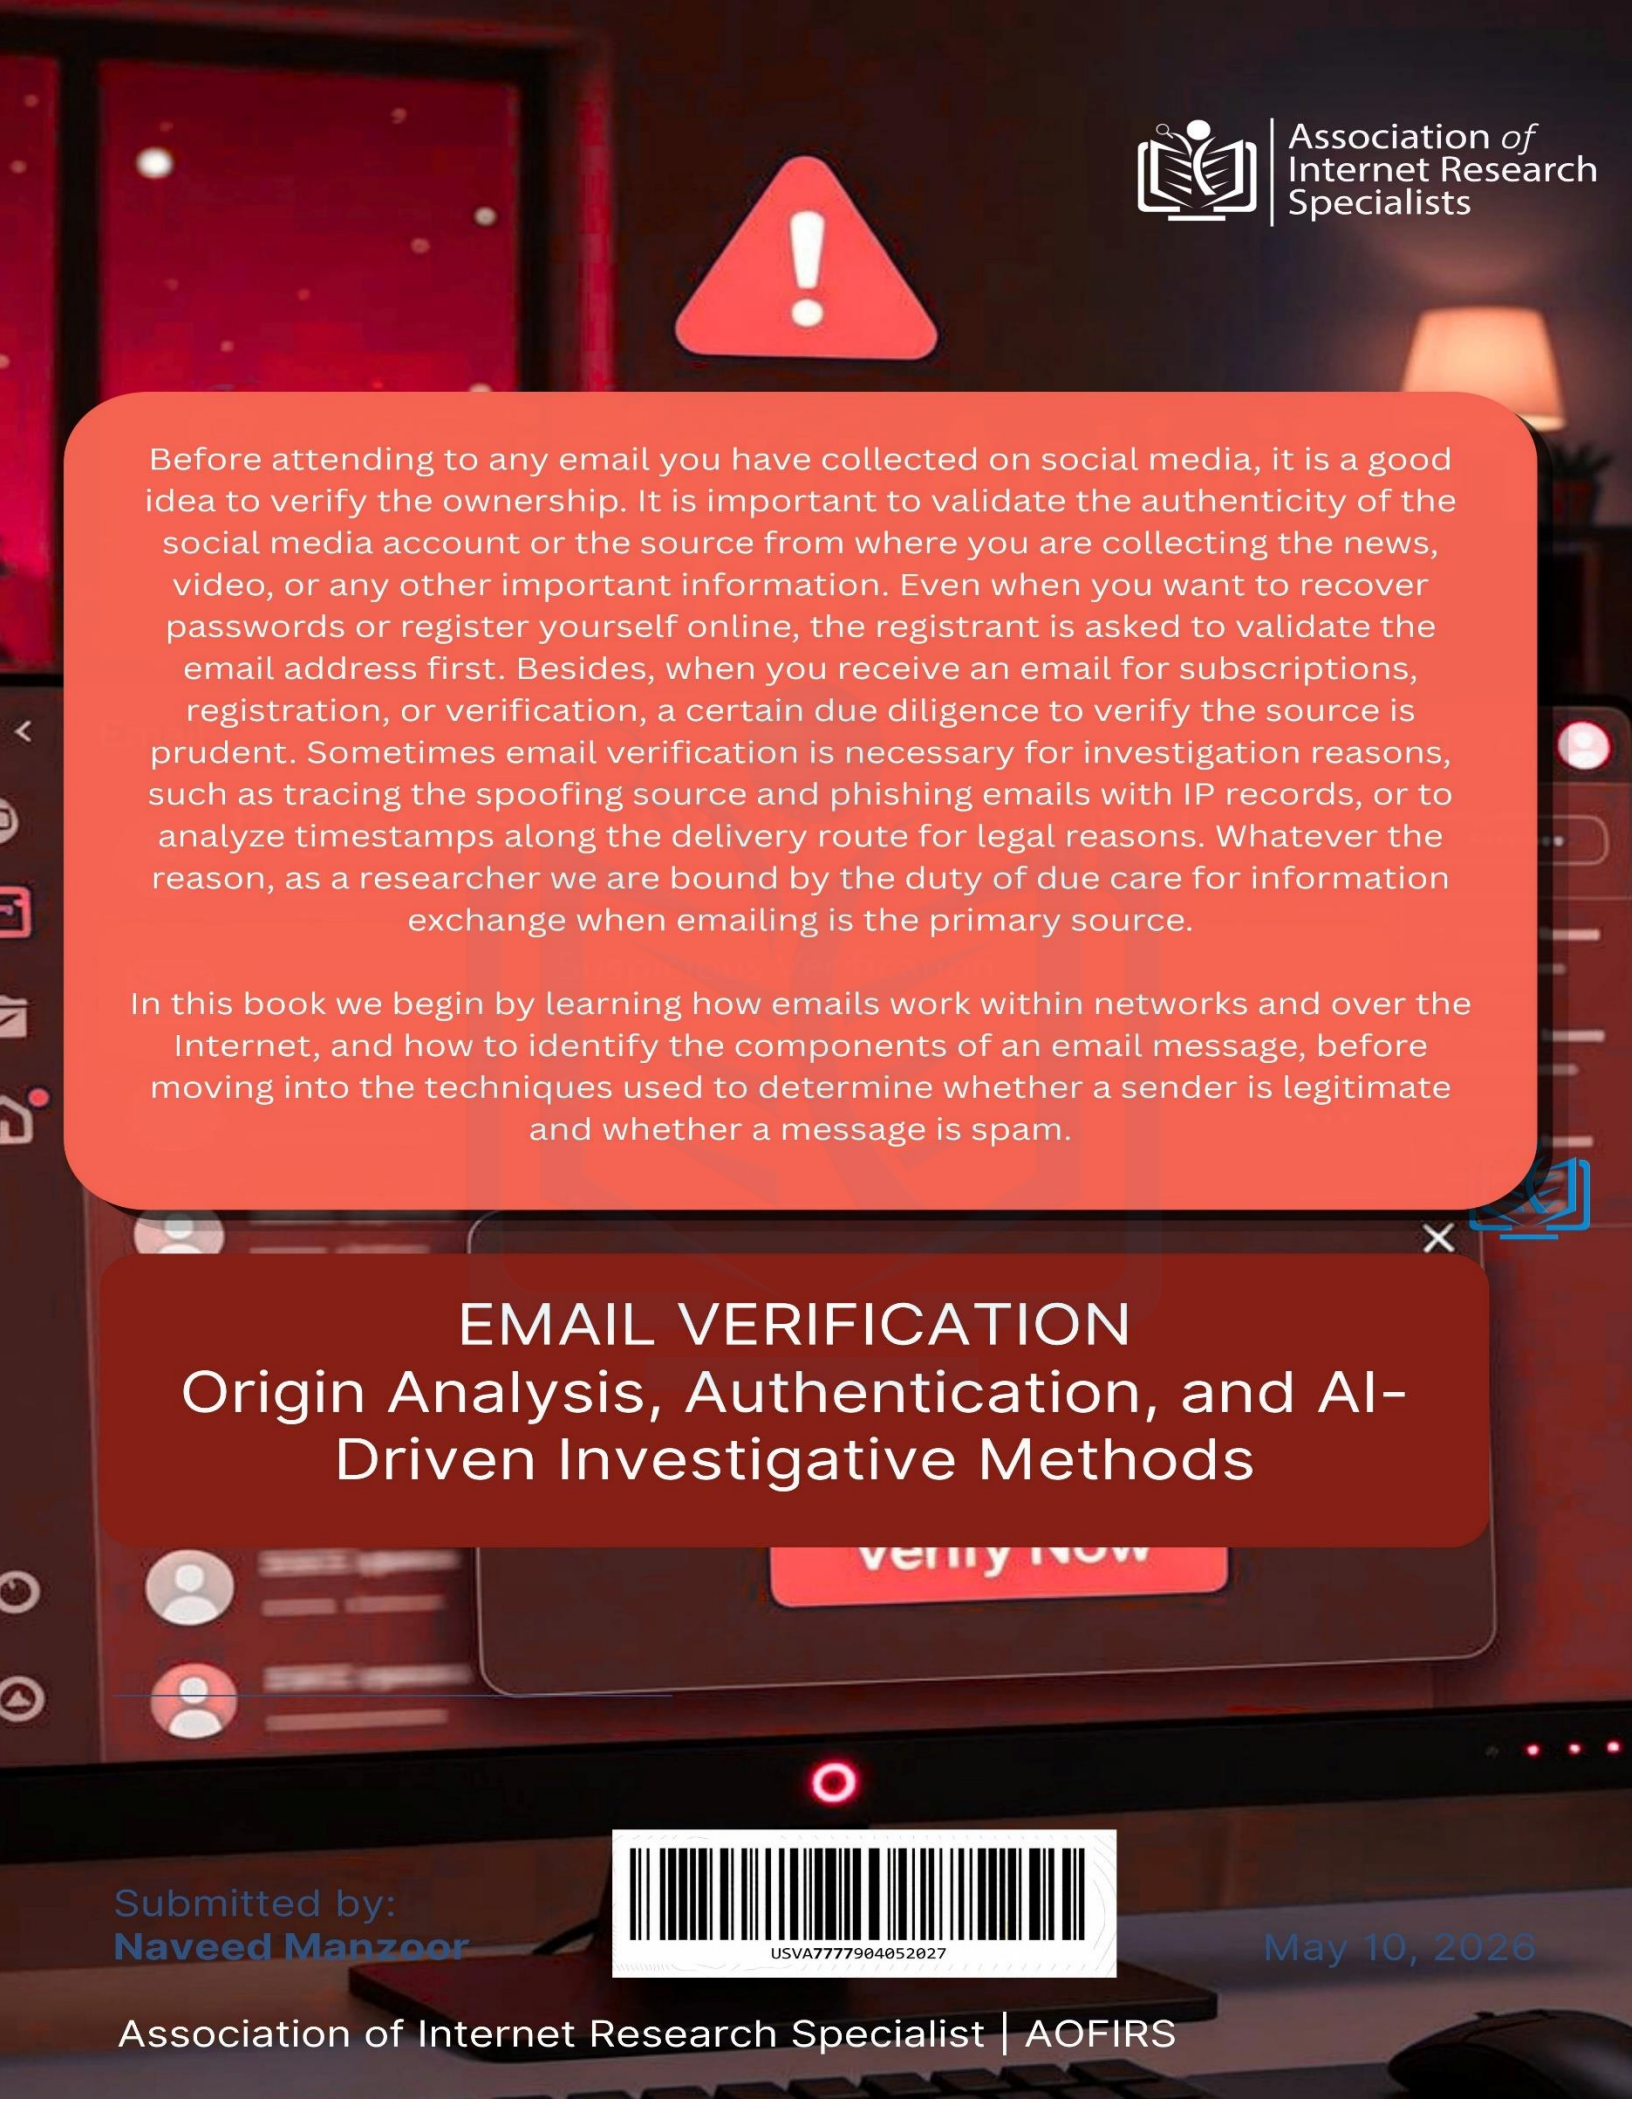

Association of Internet Research Specialists

Before attending to any email you have collected on social media, it is a good idea to verify the ownership. It is important to validate the authenticity of the social media account or the source from where you are collecting the news, video, or any other important information. Even when you want to recover passwords or register yourself online, the registrant is asked to validate the email address first. Besides, when you receive an email for subscriptions, registration, or verification, a certain due diligence to verify the source is prudent. Sometimes email verification is necessary for investigation reasons, such as tracing the spoofing source and phishing emails with IP records, or to analyze timestamps along the delivery route for legal reasons. Whatever the reason, as a researcher we are bound by the duty of due care for information exchange when emailing is the primary source.

In this book we begin by learning how emails work within networks and over the Internet, and how to identify the components of an email message, before moving into the techniques used to determine whether a sender is legitimate and whether a message is spam.

EMAIL VERIFICATION
Origin Analysis, Authentication, and AI-Driven Investigative Methods

Submitted by:
Naveed Manzoor

USVA7777904052027

May 10, 2026

Association of Internet Research Specialist | AOFIRS