Message, Mobile and Malware Anti-Abuse Working Group

M^{3}AAWG Trust in Email Begins with Authentication

Created June 2008
Updated February 2015

Edited by:
Dave Crocker, M^{3}AAWG Senior Technical Advisor, Brandenburg InternetWorking
Terry Zink, Microsoft

Abstract

The Internet's growth allows us to interact with people all over the world. Unfortunately, some of those people do not make good neighbors. Along with the effort to detect and filter the problematic traffic they generate, there is a complementary effort to identify trustworthy participants. In security technology parlance, the first seeks to identify *Bad Actors,* whereas the second creates ways of distinguishing *Good Actors*.

Identifying Good Actors can be divided into two activities:

1. Identifying a participant, such as an author or an operator of an email service.

2. Assessing the participant's trustworthiness.

The first activity is called *authentication*. The second is usually called *reputation assessment*.

This white paper considers the first step: authenticating the identity that asserts responsibility for an email. It examines recent developments in standardized authentication mechanisms that have been tailored for use in email anti-abuse efforts. It also provides background on authentication as a foundation for understanding current efforts to protect Internet mail. It then looks at the most popular mechanisms currently in use.

This paper is intended for a general readership that has basic familiarity with Internet mail service. While this single document is unlikely to be the final word on the topic, M^{3}AAWG has striven to capture the current best practices regarding email authentication. As a complement to enabling identification of Good Actors, authentication can aid efforts in protecting business' brands from forgery and phishing attacks. The Executive Summary provides a one-page overview that can be used independently.

Table of Contents

1. Executive Summary

The disruptive effects of spam and other email abuse have generated two lines of response by the email services industry. One focuses on detecting and filtering problem messages. A complementary, but different, response seeks a basis for trusting a message rather than for mistrusting it. Is someone trustworthy responsible for the message? This approach has three steps:

1. Assign a person or organization's name to the message – *identification*
2. Verify that the use of this identifier is authorized and correct – *authentication*
3. Determine the trustworthiness of the identity – *reputation assessment*

This paper discusses current industry efforts to satisfy the first two requirements. Internet mail is extremely flexible with regard to the types of identities that can be referenced. Most recipients only know about the `From:` header field, containing a displayable name of the message author and their email address. However, email also can separately name the agent that posted the message for sending, the agent that receives return handling notices, and the agents that handle the message at different stages of transmission. Because each of these is often referred to as a *sender* of the message, the term *sender* can be ambiguous.

When authentication mechanisms are applied, both the originating and receiving systems are able to correctly and reliably validate who is accountable for the message. This is what is meant by "knowing who the message came from." Two paradigms can be used for authentication. One builds authentication into the email-handling infrastructure along the path that a message travels. The other wraps authentication information into the message itself and is independent of the infrastructure. (Other email authentication technologies do not deal with email transit accountability.) When an authentication mechanism becomes widely popular, it opens the door to a variety of assessment products and services that can rely on it.

Two authentication mechanisms, *Sender Policy Framework* (*SPF*) and *DomainKeys Identified Mail* (*DKIM*), have emerged as techniques for authenticating who is accountable for a message while it is in transit. Each uses a different technical paradigm, uses different identifiers, and has different limitations and flexibilities. *SPF* uses a path registration approach, whereas *DKIM* uses digital, cryptographic signing of the message itself.

SPF uses the underlying network address (the *IP address*) of the email-sending neighbor that is closest to the validating server and the machine name (*domain name*) in the message's return address or the domain name in the protocol transfer greeting between email handling hosts to make its authentication checks. It queries the Domain Name System (DNS) to perform a mapping between the name and the address. Hence, for the purposes of SPF evaluation, the IP address of any email relay that might be a neighbor to any receiver must be known beforehand and pre-registered in the DNS by the sender.

DKIM, by contrast, uses digital, cryptographic signatures, attaching information to a new header field in the message. A DKIM signature can withstand minor message modifications without becoming invalid, including some that are made by forwarding services and mailing list software. Any domain name can be used for signing the message; the identifier is not tied to any existing email identifier. As with SPF, the queried entry in the DNS self-validates the name's use. Associating it with an existing identifier, such as the `From:` or `Sender:` header fields, is a separate step that is beyond the scope of DKIM itself.

Domain-based Message Authentication, Reporting & Conformance (DMARC) provides an overlay capability that ties the author's domain name to authenticated domain names used by SPF and/or DKIM. DMARC allows domain owners to publish DNS records, indicating authentication practices and requesting how receivers should handle messages that fail authentication checks. It aids assessment of the legitimacy of messages lacking authentication via SPF or DKIM. It also enables receivers to send reports back to domain owners, and allows domain owners to request that a receiver reject messages that fail authentication checks.

2. Introduction

There are two lines of industry response to protecting users against abusive email. One focuses on deciding that a message is from a bad actor; it is based on detecting and filtering out problem messages. It analyzes email sources, traffic patterns and content. Mail cannot cause problems for recipients if it is not delivered to them—unless, of course, it should have been delivered. Although essential as a first line of defense, this approach is approximate. Filtering junk email can be prone to error, with *false positives*—legitimate email classified as junk—and *false negatives*—junk email classified as legitimate. The distressing yet inevitable result is that this approach produces an ever-escalating arms race of counter-techniques, locking the abuse and anti-abuse communities in a constant struggle. For example, as the anti-abuse community has gotten better at analyzing textual content, the advent of *image spam* has become a creative vector of attack to defeat such analysis. Some service providers have become extremely proficient in protecting users, but the cost is very high and the protection is very fragile. Few providers can ensure the necessary level of protection. With email abuse estimated to be 90 percent of Internet messaging traffic or worse, it has become critical to find ways of restoring user confidence in email.

The second approach, which complements problem email detection, is to find a basis for trusting the message rather than mistrusting it. A message determined to be from a known good actor does not need to be subject to the stringent analysis that would be applied to mail from an unknown source. The usual method of accomplishing this is to associate a confirmable identity with the message and to obtain an assessment about that identity. In other words, this approach provides the recipient with a means of deciding that someone trustworthy is responsible for the message. It answers these questions: Who is the author,and are they known to write legitimate messages?

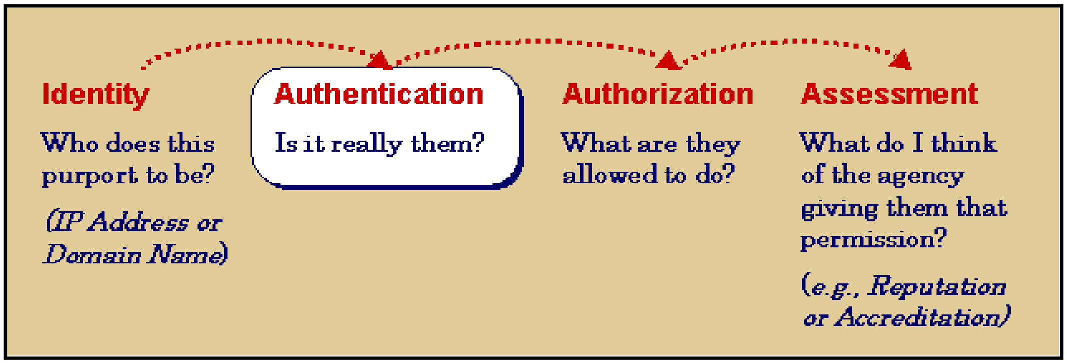
Identity
Who does this purport to be?
(IP Address or Domain Name)
Authentication
Is it really them?
Authorization
What are they allowed to do?
Assessment
What do I think of the agency giving them that permission?
(e.g., Reputation or Accreditation)

Figure 1 - **The Assessment Framework**

As shown in Figure 1, if we are to trust the claimed identity of a message sender, then we first need a mechanism that validates the identity's use. This is called *authentication*. Only when we know that the identity is valid can we assess its trustworthiness. Without authentication, we can have knowledge about a person's or an organization's reputation, but we cannot be certain that they are indeed responsible for this specific message, since it might have been produced by a bad actor who is using the associated identifier without authorization.

Authentication-based assessment follows these three steps:

1. Assign an *identifier* to a message, which refers to an *identity* – the name of a person or organization.
2. Provide a means of validating the use of that identifier – an *authentic* identifier that is *authorized* for this use.
3. Assess the *reputation* or trustworthiness of the identity using that identifier.

This paper focuses on the first two steps, which produce an authenticated reference to an identity. The technologies discussed here describe how to assign an identifier to the message. The second step checks whether this use of the identifier is permitted. The third step assesses the trustworthiness (*reputation*) of the responsible person or organization *(identity)* for that use.

The difference between the processes of validating an identity described in the first two steps and assessing reputation in the third can be illustrated by the contrasting roles played in commercial transactions by a driver's license and a credit score. A license contains a name that refers to a person; the name is an identifier. A license is a generally accepted way of validating a person's identity. However, it does not indicate that the person is a good loan risk. Instead, the credit score serves as the measure of reputation and is used to determine creditworthiness.

Moreover, a person or organization can have multiple reputations, each within its own context. Reputation analysis begins by determining the authorized scope for using an identifier and then assessing the associated identity within the current context. For example, in the case of the loan above, the context analysis might start with whether the person is applying for a personal loan or a loan for their company. In the case of email reputation assessment, the starting point for the context analysis might be whether the individual is sending a personal message or one purporting to represent their company.

3. Underlying Technologies

3.1 Email

The global architecture for Internet mail is shown in Figure 2. There is a simple split between the user world, in the form of *Mail User Agents (MUA)*, and the transmission world, in the form of the *Mail Handling Service (MHS),* itself composed of *Mail Transfer Agents (MTA)*. An MTA that sends an organization's mail directly into the public Internet, or receives mail from it, is called a *boundary,* or *border,* MTA. An MTA that sends messages handles *outbound* mail. An MTA that receives messages handles *inbound* mail.

The MHS is responsible for accepting a message from one user, the *author*, and then delivering it to one or more other users, the *recipients*. This creates a user experience of an apparently direct MUA-to-MUA exchange, without users having to be cognizant of the intervening steps performed by the MHS infrastructure. The first component of the MHS is called the *Mail Submission Agent (MSA)* and the last is called the *Mail Delivery Agent (MDA)*.

Internet mail is often referred to by one of its standards, *SMTP* (Simple Mail Transfer Protocol), but it is really a collection of specifications, with the core being SMTP for transfer, RFC 5322 for the message header and MIME for the message body and attachments. A recipient retrieves a delivered message using *POP* (Post Office Protocol) or *IMAP* (Internet Message Access Protocol) or proprietary protocols such as Microsoft's MAPI. A popular form of user access is called *webmail*, which is simply a classic MUA email functionality, accessed over the Web.

Internet mail standards specify the meaning of identities such as *author* and *sender* that are used in originating a message, but do not mandate or enforce particular choices for them. Although some software and some originating organizations choose to constrain the use of identifiers, the reality is that the underlying Internet standards permit anyone to claim to be the author of a message. While this is also true for telephone and paper mail, email requires more effective authentication techniques because it is far easier to scale up abuse over email than it is over either telephone or postal communications.

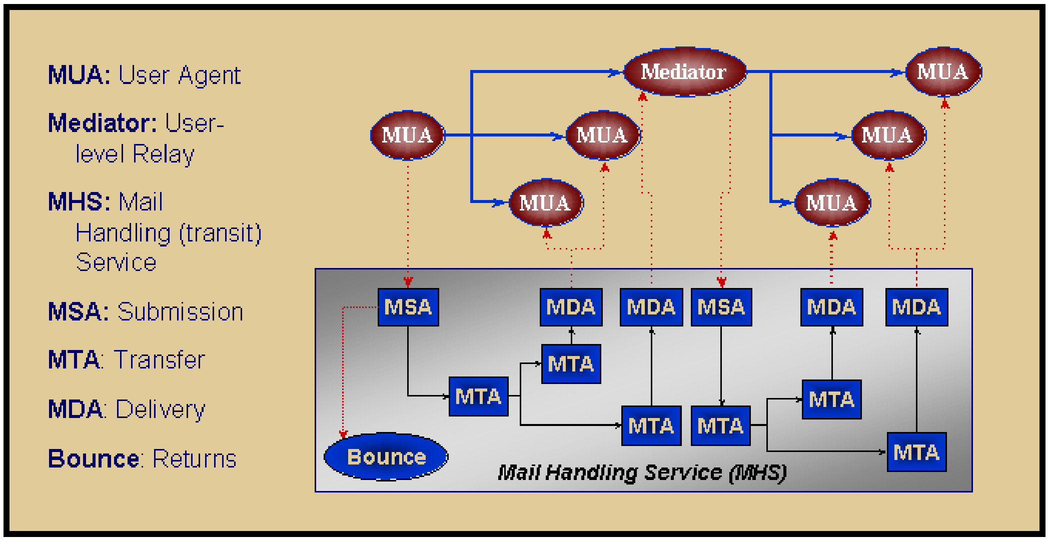

Figure 2 - **Internet Mail Service Architecture**

3.2 Many Roles in Handling a Message

The concept of an (online) identity has a rich and confusing history. One source of confusion is the difference between a thing itself and the name of the thing. In society, a person or an organization can be referred to by one or more names, such as a legal name or a nickname. Both refer to the same person or organization, but with different labels. In technical parlance, the term *identity* refers to the person or organization itself. The term *identifier* refers to a label that is used for that identity.

To return to the case of the driver's license, there is one *identifier*—the name on the license—and one *identity*—the person to whom the identifier refers. To be credible, the linkage of an identifier to an identity needs *authentication*. A bank might use a driver's license to confirm identity, but it needs additional information, such as a credit score, to determine a person's financial reputation and creditworthiness. The better a person's reputation, the more financial and other privileges they enjoy, such as lower interest rates. Similarly, the better an email sender's reputation is, the less likely they are to have restrictions placed on their email. Restrictions that are commonly placed on email include limitations on attachments or attachment types, the number of messages that are permitted to be sent into an organization in a given amount of time, the level of "spamminess" that is permitted in a message before it gets blocked, and so on.

There can be a number of different identifiers associated with each message, as listed in Figure 3. Most recipients recognize the `From:` header field, containing a displayable name of the message's author, along with their email address to which replies are usually sent. However there are other identities, too. These are provided primarily to assist the system in distinguishing different actors, such as who posted the message for sending, who will receive return receipt (Return) messages, and who handles each stage of message transmission. Because each of these is often referred to as a "sender" of the message, the term "sender" can be ambiguous.

Identifier	Meaning of Identification
Peer MTA Host IP Address	Neighbor SMTP Client Host
`SMTP EHLO` command	Neighbor SMTP Client Organization
Peer Network IP Address Range	Neighbor SMTP Client Provider
SMTP `MAIL FROM` command	Notification Return Address
RFC 5322 `From:` header field	Content Author
RFC 5322 `Sender:` header field	Message Posting Agent
RFC 5322 `Received:` header field	Transit Handling Organizations

Figure 3 – **Internet Mail Identities**

Among these identities, there is a distinction between responsibility for content—that is, the *author*—and responsibility for message handling—that is *posting, transmission* and *returns*. Each can be useful to consider when deciding whether to accept an authenticated message. Different authentication techniques use different identities. But unfortunately, answering the key question of who is "responsible" for a message and then validating that agent's identity is not always easy.

3.3 Domain Name System (DNS)

The Domain Name System (DNS) provides a mapping service from *domain names* to associated information, such as the *IP address* of the Internet hosts that are known under the registered name. The name itself specifies a registration hierarchy, such as `marketing.example.com`, with the rightmost field being the top of the hierarchy.

The term *mapping* distinguishes the DNS from more general "search" services. It takes the exact name and produces either an exact match or a failure. A search service, by contrast, explores for approximate matches. An aggregation of domain names that are resolved by a single DNS server host is called a *zone*. This construct is administrative; it is not visible to users of domain names. A host might resolve a single level of the name's hierarchy (a *subtree*), or multiple levels.

The information that is associated with a domain name is listed under a set of *Resource Records* (RR). Each type of `RR` has its own format. The `TXT RR` has a general text format that is further defined by various applications that use it for recording different information. Because it has multiple uses, the `TXT RR` can produce ambiguity. So, to distinguish which application defines the `TXT` contents, the record either must contain a defining string, or the record must be stored under a special DNS naming subtree, typically defined by a name string that begins with an underscore, such as `_domainkey`.

Domain names are the core of any email authentication scheme that does not simply use IP addresses. Hence, protection of information available through the DNS is essential when authentication is a concern. General protection of information in the DNS is pursued through the Domain Name System Security Extension (DNSSEC).

3.4 The Promise of Authentication

In this white paper, the term *authentication* describes technologies for determining whether or not an identifier is being employed by and for the organization (the *identity*) that it belongs to. IP address and domain name registrations derive from the global Internet administrative authority, the Internet Corporation for Assigned Names and Numbers (ICANN). Authentication answers these questions: Is an IP address being used by the organization it is assigned to? Is a domain name being used by the organization that registered it?

Email authentication specifically addresses the problems caused by Internet mail's flexibility in choosing identities to use. With authentication, both originating and receiving email systems gain a mechanism for validating who is responsible for the message. This is generally described as "knowing who the message came from." This ability can ensure that legitimate mail reaches its intended recipients. It also creates an opportunity for presenting recipients with a visible indication that a message can be trusted. Some commercial experiments are exploring users' acceptance of and interest in such an indication. However, early results seem to show that users often do not notice or understand these markers. Consequently, a more promising path for using authenticated information is as part of a receiving system's automated filtering engine.

When an identifier is attached to a message, the owner of the identifier is declaring that they are accountable for the message. This means that their reputation is at stake. Receivers who successfully validate the identifier can use information about its owner as part of a program to limit spam, spoofing, phishing or other undesirable behavior.

Once the receiver validates the identifier, the receiver may use that information to assist in improving the deliverability of a message or to reduce filtering costs and errors. The choice is entirely at the discretion of the validating receiver. When a message is authenticated, a receiver uses its knowledge about the owner of the identifier—that is, the identity—to determine the most appropriate treatment of the message. It is frequently assumed that messages associated with an identity that has a good reputation will be subject to less scrutiny by the receiver's filters. Although the assumption is often correct, it is not a guarantee. Absent a contractual agreement, a sender's use of one particular authentication technique or another does not ensure delivery or recipient viewing.

3.5 Authentication Techniques

Today, there are seven deployed and openly available techniques for authenticating email:

1. IP (Internet Protocol)
2. PGP
3. S/MIME
4. BATV
5. Sender Policy Framework (SPF)
6. DomainKeys Identified Mail (DKIM)
7. Domain-based Message Authentication, Reporting & Conformance (DMARC)

This white paper focuses on the last three—SPF, DKIM and DMARC. Brief treatments of IP, PGP, S/MIME and BATV appear in the Appendix.

SPF and DKIM cover two major technical paradigms. They use different identities and different administration models, and have different limitations and flexibilities. The third, DMARC, adds a layer of assurance over and above SPF and DKIM.

- *SPF* uses channel-based *path registration*.
- *DKIM* uses object-based *cryptographic authentication*.

A simple basis for distinguishing these two paradigms is depicted in Figure 4. The scheme on the left uses *channel authentication* built into the email-handling infrastructure along the path that a message travels. The scheme on the right adds *object authentication* information to the message itself. That is, one authentication mechanism is associated with the transmission channel, whereas the other is carried along with the content itself.

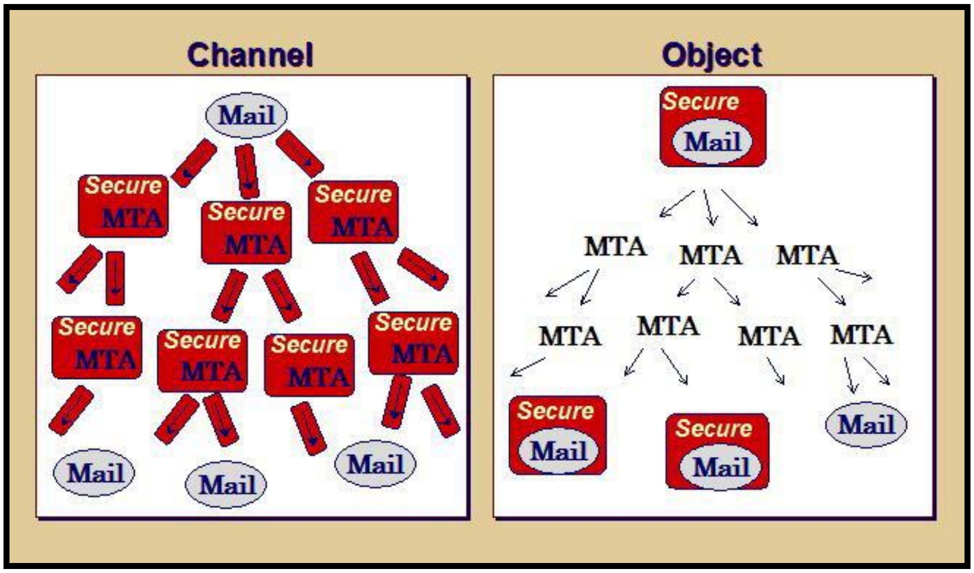

Figure 4 - **Approaches to Security**

SMTP does support some additional authentication methods (SASL, SMTP AUTH and SSL). These are primarily used for authentication and privacy during initial message submission of new mail rather than during later relaying and delivery steps. However, there is a recent effort to promote their use for relaying as well,in response to increased public concerns over pervasive monitoring. Consequently, these mechanisms are useful for establishing a responsible author under the originating organization, but are not directly useful for analysis by a receiving system.

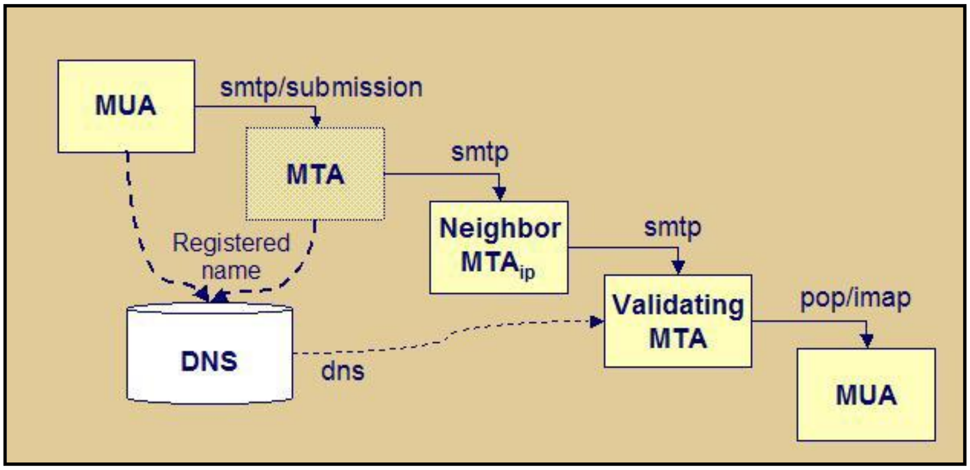

Figure 5 - **Channel Authentication (Path Registration)**

3.6 Channel Authentication (Path Registration) Basics

Path Registration schemes build upon the established idea of using the IP address of the neighboring SMTP MTA sending client. They extend it by mapping that address to a domain name, like `example.com`, and querying the DNS to see whether the IP address is associated with that domain name, as shown in Figure 5. This gives an organization a means of publicly registering the MTAs that are authorized to send mail on its behalf. The validating site checks the address of the neighboring server that is sending the email against a registered list of servers that the domain owner has authorized to send email.

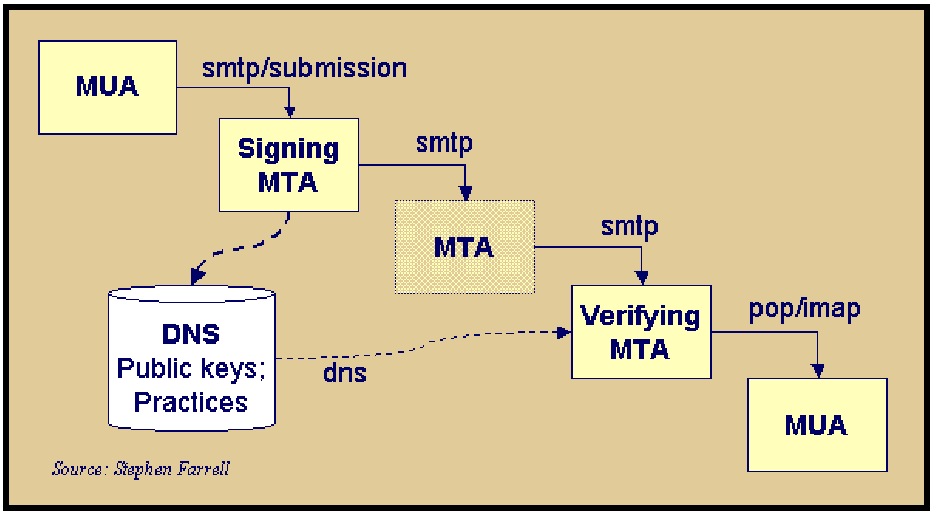

Source: Stephen Farrell

Figure 6 – **Object (Cryptographic) Authentication**

3.7 Object (Cryptographic) Authentication Basics

Cryptographic authentication performs mathematical calculations on selected message content. These produce a tiny digital summary, called a *hash*, of the message content that is extremely difficult to reproduce without a copy of the content. Additional calculations protect the hash so it cannot be modified without detection. The signer uses a private (secret) cryptographic key to create the signature, while the validating side uses a corresponding public key to verify it. The public key can be circulated openly, such as in the DNS under a domain name of the signer, as indicated in Figure 6. In fact, this form of publication is self-certifying; by virtue of finding it under the domain name for which it is used, the public key is known to be valid. When validating a signature, the hash is recalculated. If the message has been modified, the hash will not be correct and the validation will fail.

3.8 Using SPF, DKIM and DMARC

The sections that follow examine the SPF, DKIM and DMARC authentication technologies with respect to these basic questions:

- Which identity is authenticated?
- How is it authenticated?
- What is the DNS query mechanism?
- How hard is it to implement on the originating side?
- How hard is it to implement on the receiving side?
- When does the technique not work?

The existence of multiple authentication techniques presents a very real question: What is the effect of using more than one? That is, will using multiple authentication techniques be helpful for message delivery, or could the presence of multiple technologies in the same message interact in counter-productive ways? The current industry view is that using more than one mechanism will not cause problems, and might well have some benefits.

3.9 Sender Policy Framework (SPF)

SPF permits registering authorized server addresses under the domain name of a return address.

Which identity is authenticated?

SPF uses the IP address of its SMTP neighbor and maps it to the domain name in the `MAIL FROM` Return command of SMTP (recorded in the `Return-Path:` field of delivered messages) and/or the `HELO/EHLO` SMTP command. The `HELO/EHLO` name is explicitly provided by the neighboring SMTP client host to label itself.

It is helpful to view the IP address as an identifier, with the mapping to a domain name being useful for aggregating a number of different MTAs' IP addresses under the same organization and allowing them to share its reputation.

How is it authenticated?

SPF uses path registration. A site that validates SPF receives a message from a neighboring MTA. It uses the IP address of that neighbor and the domain name in the SMTP `MAIL FROM` Return and/or the `HELO/EHLO` commands for the message. Validation consists of finding the IP address registered under the domain name.

According to the SPF specification, querying on the `MAIL FROM` Return command is mandatory. Querying the `HELO/EHLO` command is recommended.

What is the DNS query mechanism?

The owner of the MAIL FROM and/or HELO/EHLO domain name registers a TXT record in the DNS that contains the IP addresses of the outbound border MTAs sending email to validating inbound MTAs. A validating site queries the DNS for the domain that it obtained from the MAIL FROM Return command. If it finds that the outbound border MTA's IP address is registered under it, it means that the address is authorized to send email containing that domain name in the MAIL FROM and/or HELO/EHLO commands.

The SPF DNS record is not just for registering authorized systems. It is also intended to be a flexible means of publishing a variety of email service practices information. This includes registering addresses that are explicitly not authorized, alternate mechanisms that are authorized, and even recursive references that derive authorization information from other records. This flexibility can make it challenging to create records that accurately reflect the policies of a registering organization. Consequently, administration software has been developed to facilitate the process of specifying a SPF DNS record for the most common configurations of mail software.

How hard is it to implement on the originating side?

The owner of the domain name that is specified in the MAIL FROM command and/or HELO/EHLO command must register each MTA's IP address that might send mail to a recipient MTA – which in turn might evaluate the sending MTA's reputation. Outbound border MTAs are registered, so that inbound border MTAs can do the validation. As the IP addresses of registered MTAs change, the registration records must be updated.

Many receiving MTAs stamp the results of an SPF validation in an Authentication-Results: header field of the message using SPF=<result> where <result> is pass, fail, softfail, neutral, none, or temperror or permerror. For more details, see RFC 7001.

On the origination side, the only software required for supporting SPF is a DNS administration tool that permits creating the necessary DNS record(s).

How hard is it to implement on the receiving side?

A validating site obtains the IP address of the neighboring MTA from the underlying operating system, and the domain name from the SMTP MAIL FROM command. It performs a DNS query to determine whether the MTA is registered. Software must be added to the receiving component to perform this validation.

When does the technique not work?

The MAIL FROM return address tends to match the domain name of the author or of the posting agent, *but this is not required*. Valid MAIL FROM return addresses can have entirely different domains, such as when processing bulk mail requires that error notices go to a special address. In such cases, SPF validation will fail unless the domain of the address for processing error notices also has the appropriate SPF record.

Email is often sent directly from an originating organization to a receiving organization. SPF works for these cases as long as:

- the originating organization is able to identify all of its boundary MTAs
- it keeps them correctly registered
- validation is performed at the inbound boundary MTA, or the inbound MTA safely adds the IP address of the neighboring outbound boundary MTA to the message so it can be evaluated later.

However, email that is sent through an intermediary such as a forwarding service, mailing list, or other re-posting agent will fail an SPF validation unless the intermediary is explicitly registered under the SPF domain name. Mail that is sent through unregistered MTAs will fail validation. One example of this occurs when mobile users are forced to post mail through their access provider's MTA rather than being able to post through their home organization's email service. Another common validation failure is mail that is sent through an outsourced service, such as subscription bulk email handled by a contractor.

When several organizations use the same MTA's infrastructure, a security exposure is created. If Organization 1 and Organization 2 both publish SPF records in their DNS entries that include the MTA's IP addresses, then Organization 1 pretends to be Organization 2, this will (incorrectly) pass SPF validation at the receiver side. The receiver will believe that the message has passed SPF validation for Organization 2. In reality, the message came from Organization 1. The only way to eliminate this exposure is for the organizations to use separate IP addresses.

3.10 DomainKeys Identified Mail (DKIM)

DKIM is based on cryptographic content signing. It was produced through a merger of Yahoo!'s DomainKeys and Cisco Systems Inc.'s Identified Internet Mail (IIM) specifications. First developed by an informal industry consortium, it was then revised and has IETF approval for standards track status.

Which identity is authenticated?

DKIM allows the signer to choose *any* domain name, which is indicated in the `DKIM-Signature:` header field of the message. Whether that domain name is related to another identifier in the message, such as the `From:` or `Sender:` fields, is a separate decision and is outside of the DKIM specification. The DKIM signature validates only the use of that domain name. It does not purport to validate any of the message content—not even the author's address.

How is it authenticated?

A responsible organization that handles the message—author, originating operator, or relaying operator—adds a digital signature to the message, associating it with a domain name of that organization. Typically, signing will be done by a service agent that is part of the message author's organization or is delegated by it. Signing might be performed by any of the components in that environment, although most often it is the MSA or MTA that adds the signature. DKIM permits signing with a particular domain name to be performed by authorized third parties. For example, an originating organization can obtain a signature from an independent assessment (reputation) organization and affix that signature to the message.

What is the DNS query mechanism?

DKIM defines a `TXT` record that is placed under a special sub-domain of DNS (`._domainkeys.`), which is underneath the domain name declared in the `DKIM-Signature:` header field. Any `TXT` records under that sub-domain name are only for DKIM use.

The "lowest" portion of the DKIM sub-domain is a field, called the *selector,* which is used for key management. So the DNS query string has multiple fields to the query name, with only a portion intended to be used for actual *reputation* assessment—that is, a domain name that represents the organization. It is then combined with a selector so that keys can be assigned more conveniently. For example this is necessary when control over signing is held by different individuals or systems, as well as when migrating to a new key.

In an attempt to improve email privacy, there has been a recent effort to advise having selectors be unpredictable; that is, to make them be longer, random strings.

The DKIM DNS record has some parameters for constraining its use to particular services or addresses. However the record only validates an existing signature.

Example: Header from an email newsletter with a DKIM signature

In the DKIM signature the domain is authorscompany.com and the selector is esp1234:

```
Received: from mta1.esp1234.com (HELO mta1.esp1234.com) (10.0.0.1)
 by mailserver.company.com with SMTP; 28 Mar 2008 19:53:28 -0000
DKIM-Signature: v=1; a=rsa-sha1; c=relaxed/relaxed; s=esp1234;
d=authorscompany.com;
h=From:To:Subject:Mime-Version:Message-ID:Content-Type:Date;
i=author@authorscompany.com; bh=EMR7D1qC7ykz41K8ArLCt++IWxM=;
b=TGkNEq7fW4OIno/5DlX2qHDQeRmzhY+uiTzEcxu2KIKC+4B7+i2olIWGZP9JBnOR4Ck6iAiidnR
 jDLuc2QJh3ifDNPWJ6xYjiuE73ilCZfbtN0r2MVke9pRU4aydBQ5DSCFS7YhUFB22CT70Mut
 ZkaDFSZZpqI5vTlSWm9MI8PM=
Date: Fri, 28 Mar 2008 14:53:27 -0500 (CDT)
From: "Author" <author@authorscompany.com>
To: Recipient@company.com
Subject: March Newsletter
Sender: authorscompany@esp1234.com
Return-Path: bounce-4101674@authorscompany.esp1234.com
Mime-Version: 1.0
Message-ID: <20080324040103985572.328428@mx12.emailroi.com>
Content-Type: multipart/alternative;
boundary="============_emailROI_============"...
```

We look up the corresponding DKIM DNS record using the standard nslookup program.

```
nslookup -type=txt esp1234._domainkey.authorsdomain.com

esp1234._domainkey.authorsdomain.com text =

"v=DKIM1; g=; k=rsa;
p=QIdfMA0GCSqGSIb3DQEBAQUAZ4GNADCBiQKBgQC61RrUNTIcNbf/+f5Co2V37GMvPQdbUVyjgvL
XrUKAXeJDwYVumAtE9BovuDZNYxcgG2oy7mkcZX/3rBF5SJX9Cp5yw0axuMpzkuzPQq26h+2+MLuv
tJtfDIaHgNeEJOjMeq7s9RFQHRr9g26lkZQTRAob8YevaA1KHiNNyIaZuQIDAQAB;"
```

How hard is it to implement on the originating side?

DKIM signing can be performed anywhere in the originating or relaying path. A common example is a department MTA or an organization's outbound boundary MTA. Signing requires addition of a software module that can:

- obtain the private key
- perform the necessary calculations
- affix the signature information to a `DKIM-Signature:` message header field

How hard is it to implement on the receiving side?

DKIM validation can be performed anywhere along the message transit path after the message is signed. Validation requires a software module that can:

- obtain the public key
- perform the necessary calculations
- affix the results to the message in a trusted manner or hand the results to an evaluation engine

Recording the results means that the signature can be used by the recipient organization's filtering software later in the internal handling sequence, rather than requiring the boundary system or the recipient end-user to make an assessment.

Many MTAs stamp the results of a DKIM validation in an `Authentication-Results:` header field of the message using `dkim`=<result> where <result> is pass, fail, neutral, none, or error. For more details, see http://www.rfc-editor.org/rfc/rfc7001.txt.

When does the technique not work?

Some mail transit behavior can modify message content in a way that breaks an existing signature. DKIM has features that attempt robustness against some common modifications. However, this robustness is (intentionally) limited, since it would otherwise open the door to abuses.

3.11 Domain-based Message Authentication, Reporting & Conformance (DMARC)

DMARC builds upon both SPF and DKIM. It ties the author's `From:` address domain name to the results of either of those two authentication mechanisms. It has been submitted for standardization through the IETF, using the Internet Draft (https://datatracker.ietf.org/doc/draft-kucherawy-dmarc-base/) as input.

Which identity is authenticated?

DMARC builds upon SPF and DKIM by requiring the following:

1. The domain in the `From:` field must pass either (or both) SPF or DKIM validation; this is called *alignment*.
2. If flexibility is permitted for a DMARC domain, and the `From:` field domain name is not identical to the one that was authenticated by SPF or DKIM, it must have the same "organizational domain" name base. An organization domain is the basic name obtained through a public domain name registry, such as under `.com` or `.co.uk`.

The process by which DMARC validates author `From:` domain alignment is described below.
Finally, DMARC defines a reporting mechanism for sending feedback to the owner of the domain, in the event that a message purportedly "from" them fails a DMARC check. This is extremely useful for diagnostics and ongoing usage tracking.

How is it authenticated?

DMARC relies upon the results of an existing SPF and DKIM check and therefore must execute after both of them. If SPF passes, DMARC compares the domain in the `SMTP MAIL FROM` (or domain in the `HELO/EHLO` if the `MAIL FROM` is <>) to the domain in the `From:` address. If DKIM passes, it does the same thing with the domain in the d= field in the `DKIM-Signature.` If either of those two align with the `From:` address, then DMARC passes. If not, DMARC fails, prompting the requested DMARC action and the sending a forensic report.

Note that DMARC's enforcement of `From:` field domain alignment serves to authenticate that domain's use in the field.

What is the DNS query mechanism?

DMARC uses DNS to publish its records. It has a policy for what receivers should do if alignment fails as well as a DMARC feedback policy.

How hard is it to implement on the originating side?

If senders have SPF or DKIM set up, the incremental requirement is merely to publish a DMARC record. However, if they wish to publish a DMARC policy of 'reject' or 'quarantine,' they must ensure all of their email passes SPF or DKIM and that all of their mail aligns with the `From:` field domain. Also if the domain owner requests reporting by receivers, the domain owner must establish the necessary feedback receipt and processing capabilities.

How hard is it to implement on the receiving side?

On the receiving side, software is required to update the MTA to query DMARC records and take the appropriate policy action and reporting action in the event of a DMARC failure.

When does the technique not work?

DMARC fails in the same cases where both SPF and DKIM fail:

1. If a message is relayed, this will break SPF. Thus, if the message is not DKIM-signed, this will fail DMARC.
2. If a message is modified in transit, this can break DKIM, which will usually fail DMARC if it does not pass SPF.
3. Mail sent to one recipient which is then replayed to a second set of recipients will typically fail DMARC, since this can cause the identifiers to not align.

4. Conclusion

Receivers' efforts to detect abusive email are confounded by the efforts of bad actor senders who seek to avoid detection and accountability. As a consequence, analysis techniques are forced to deal with fuzzy information. By contrast, trust-oriented systems like SPF, DKIM and DMARC permit senders to provide explicit, accurate and reliable information with the intent to permit receivers to expedite message handling.

Authentication is the foundational component of trust-based message processing because it provides a confirmable identifier. SPF and DKIM use domain names as identifiers that are coupled to an organization taking responsibility for a message. SPF uses an identifier from the underlying email transport mechanism. DKIM uses an identifier that is independently specified. Both techniques can be extended with DMARC, allowing them to be coupled to the message's author.

5. References

Email	Overview	Internet Mail Architecture
	SMTP	RFC 5321
	Internet Mail Format	RFC 5322
	Body Format (MIME)	RFC 2045
	POP	RFC 1939
	IMAP	RFC 2060
	SASL	RFC 2222
	SSL	RFC 2246
	DNSSEC	RFC 4033
		RFC 4034
		RFC 4035
	PGP	RFC 4880
	S/MIME	RFC 3851
	BATV	Bounce Address Tag Validation
DNS	Overview	RFC 1034
	Specification	RFC 1035
	ICANN	http://icann.org
SPF	Home Page	http://www.openspf.org
	Overview	http://www.openspf.org/Introduction
	Specification	RFC 7208
DKIM	Home Page	http://dkim.org
	Overview	DomainKeys Identified Mail (DKIM) Service
	Specification	DKIM Overview
		RFC 6376
DMARC	Home Page	http://dmarc.org
	Overview	DMARC Overview
	Specification	DMARC specification

6. Appendix

IP (Internet Protocol) is the glue that holds the Internet together. It is the underlying mechanism for transferring data; all applications are built on top of it. In email authentication, an *IP address* identifies a network interface to the source host system that is the direct SMTP transfer neighbor to the one that is validating the message. In networking parlance, the source is *one hop* away from the receiver, and for some authentication efforts the receiver uses the IP address of that source. The address is obtained from this lower layer of Internet technology independently of the message content. Email transit from author to recipient normally entails many IP hops as it is handed from one MTA to the next, where each hop has a different source IP address.

Until recently, this IP address was the only method available for confirming a responsible identity during message transit. However it has significant limitations, both in terms of the stability of the reference and in terms of organizational alignment. Originators are often more than one hop away from receivers. Addresses change as machines move and as new Internet service providers are used. In addition, the identity is actually intended for the operator of an access system rather than the content author or their organization. These limitations are what motivated development of more recent techniques.

PGP and ***S/MIME*** use cryptographic techniques to provide long-term authentication and privacy for message content. They were not tailored for use during email transit and they have not developed any momentum for that use. Consequently, they are not considered in this white paper.

BATV (Bounce Address Tag Validation) provides authentication for messages sent to the SMTP `MAIL FROM` return address. It is used by the recipient of an email handling notice to verify that the notice is likely to be valid. As such, this mechanism addresses a different concern than authenticating the responsible handling agent or the original message, and it is not be considered further in this paper.