

## DomainKeys Identified Mail (DKIM) Service Overview

### Status of this Memo

By submitting this Internet-Draft, each author represents that any applicable patent or other IPR claims of which he or she is aware have been or will be disclosed, and any of which he or she becomes aware will be disclosed, in accordance with Section 6 of BCP 79.

Internet-Drafts are working documents of the Internet Engineering Task Force (IETF), its areas, and its working groups. Note that other groups may also distribute working documents as Internet-Drafts.

Internet-Drafts are draft documents valid for a maximum of six months and may be updated, replaced, or obsoleted by other documents at any time. It is inappropriate to use Internet-Drafts as reference material or to cite them other than as "work in progress".

The list of current Internet-Drafts can be accessed at <http://www.ietf.org/ietf/lid-abstracts.txt>.

The list of Internet-Draft Shadow Directories can be accessed at <http://www.ietf.org/shadow.html>.

This Internet-Draft will expire in August 2008.

### Copyright Notice

Copyright © The IETF Trust (2008). All Rights Reserved.

### Abstract

This document provides an overview of the DomainKeys Identified Mail (DKIM) service and describes how it can fit into a messaging service. It also describes how DKIM relates to other IETF message signature technologies. It is intended for those who are adopting, developing, or deploying DKIM. DKIM allows an organization to take responsibility for transmitting a message, in a way that can be validated by a recipient. The organization can be the author's, the originating sending site, an intermediary, or one of their agents. An organization may use one or more domain names to accomplish this. DKIM defines a domain-level digital signature authentication framework for email, using public-key cryptography and key server technology [RFC4871]. This permits verification of a message source, an intermediary, or one of their agents, as well as the integrity of its contents. DKIM will also provide a mechanism that permits potential email signers to publish information about their email signing practices; this will permit email receivers to make additional assessments about messages. Such protection of email identity can assist in the global control of "spam" and "phishing".

## Table of Contents

<b>1 Introduction .....</b>	<b>3</b>
1.1 DKIM's Scope .....	3
1.2 Prior Work .....	3
1.3 Internet Mail Background .....	4
1.4 Discussion Venue .....	4
<b>2 The DKIM Value Proposition .....</b>	<b>5</b>
2.1 Identity Verification .....	5
2.2 Enabling Trust Assessments .....	5
<b>3 DKIM Goals .....</b>	<b>6</b>
3.1 Functional Goals .....	6
3.2 Operational Goals .....	6
<b>4 DKIM Function .....</b>	<b>8</b>
4.1 The Basic Signing Service .....	8
4.2 Characteristics of a DKIM signature .....	8
4.3 The Selector construct .....	8
4.4 Verification .....	9
<b>5 Service Architecture .....</b>	<b>10</b>
5.1 Administration and Maintenance .....	11
5.2 Signing .....	11
5.3 Verifying .....	11
5.4 Unverified or Unsigned Mail .....	12
5.5 Assessing .....	12
5.6 DKIM Placement within an ADMD .....	12
<b>6 Security Considerations .....</b>	<b>13</b>
<b>7 IANA Considerations .....</b>	<b>14</b>
<b>8 Acknowledgements .....</b>	<b>15</b>
<b>9 Informative References .....</b>	<b>16</b>
<b>Authors' Addresses .....</b>	<b>18</b>
<b>A Internet Mail Background .....</b>	<b>19</b>
A.1 Administrative Management Domain (ADMD) .....	19
<b>Intellectual Property and Copyright Statements .....</b>	<b>21</b>

## 1. Introduction

This document provides a description of the architecture and functionality for DomainKeys Identified Mail (DKIM). It is intended for those who are adopting, developing, or deploying DKIM. It will also be helpful for those who are considering extending DKIM, either into other areas of use or to support additional features. This overview does not provide information on threats to DKIM or email, or details on the protocol specifics, which can be found in [RFC4686] and [RFC4871], respectively. The document assumes a background in basic email and network security technology and services.

DKIM allows an organization to take responsibility for a message, in a way that can be validated by a recipient. The organization can be the author's, the originating sending site, an intermediary, or one of their agents. DKIM defines a domain-level digital signature authentication framework for email through the use of public-key cryptography and key server technology. [RFC4871] It permits verification of the signer of a message, as well as the integrity of its contents. DKIM will also provide a mechanism that permits potential email signers to publish information about their email signing practices; this will permit email receivers to make additional assessments of unsigned messages. Such protection of email identity can assist in the global control of "spam" and "phishing".

Neither this document nor DKIM attempts to provide solutions to the world's problems with spam, phishing, virii, worms, joe jobs, etc. DKIM provides one basic tool, in what needs to be a large arsenal, for improving basic trust in the Internet mail service. However by itself, DKIM is not sufficient to that task and this overview does not pursue the issues of integrating DKIM into these larger efforts, beyond a simple reference within a system diagram. Rather, it is a basic introduction to the technology and its use.

### 1.1 DKIM's Scope

DKIM signatures can be created by a direct handler of a message, either as its author or as an intermediary. It can also be created by an independent service that is providing assistance to a handler of the message. Whoever does the signing chooses the domain name to be used as the basis for later assessments. Hence, the reputation associated with that domain name is an additional basis for evaluating whether to trust the message for delivery. The owner of the domain name being used for a DKIM signature is declaring that they accept responsibility for the message and may thus be held accountable for it.

DKIM is intended as a value-added feature for email. Mail that is not signed by DKIM is handled in the same way as it was before DKIM was defined. The message will be evaluated by established analysis and filtering techniques. (A signing policy may provide additional information for that analysis and filtering.) Over time, widespread DKIM adoption could permit more strict handling of messages that are not signed. However early benefits do not require this and probably do not warrant this.

DKIM's capabilities have a narrow scope. It is an enabling technology, intended for use in the larger context of determining message legitimacy. This larger context is complex, so it is easy to assume that a component like DKIM, which actually provides only a limited service, instead satisfies the broader set of requirements.

By itself, a DKIM signature:

- Does not offer any assertions about the behaviors of the identity doing the signing.
- Does not prescribe any specific actions for receivers to take upon successful signature verification.
- Does not provide protection after signature verification.
- Does not protect against re-sending (replay of) a message that already has a verified signature; therefore a transit intermediary or a recipient can re-post the message in such a way that the signature would remain verifiable, although the new recipient(s) would not have been specified by the author.

### 1.2 Prior Work

Historically, email delivery assessment decisions have been based on an identity that used the IP Address of the system that directly sent the message (that is, the previous email "hop"), [RFC4408] or on the message content (e.g. [RFC4406] and [RFC4407]). The IP Address is obtained via underlying Internet information mechanisms

and is therefore trusted to be accurate. Besides having some known security weaknesses, the use of addresses presents a number of functional and operational problems. Consequently there is a widespread desire to use an identifier that has better correspondence to organizational boundaries. Domain names are viewed as often satisfying this need.

There have been four previous IETF efforts at standardizing an Internet email signature scheme. Their goals have differed from those of DKIM.

- Privacy Enhanced Mail (PEM) was first published in 1987. [RFC0989]
- PEM eventually transformed into MIME Object Security Services (MOSS) in 1995. [RFC1848] Today, these two are only of historical interest.
- Pretty Good Privacy (PGP) was developed by Phil Zimmermann and first released in 1991. [RFC1991] A later version was standardized as OpenPGP. [RFC2440] [RFC3156] [RFC4880]
- RSA Security independently developed Secure MIME (S/MIME) to transport a PKCS #7 data object. [RFC3851]

Development of both S/MIME and OpenPGP has continued. While each has achieved a significant user base, neither one has achieved ubiquity in deployment or use.

To the extent that other message-signing services might have been adapted to do the job that DKIM is designed to perform, it was felt that re-purposing any of those would be more problematic than creating a separate service. That said, DKIM uses security algorithm components that have a long history, including use within some of those other messaging security services.

DKIM has a distinctive approach for distributing and vouching for keys. It uses a key-centric Public Key Infrastructure (PKI) rather than the more typical approaches based on a certificate in the styles of Kohnfelder (X.509) [Kohnfelder] or Zimmermann (web of trust). For DKIM, the owner of a domain name asserts the validity of a key, rather than relying on the key having a broader semantic implication of the assertion, such as a quality assessment of the key's owner. DKIM treats quality assessment as an independent, value-added service, beyond the initial work of deploying a verifying signature service.

Further, DKIM's PKI is provided by adding information records to the existing Domain Name System (DNS) [RFC1034], rather than requiring deployment of a new query infrastructure. This approach has significant operational advantages. First, it avoids the considerable barrier of creating a new global infrastructure; hence it leverages a global base of administrative experience and highly reliable distributed operation. Second, the technical aspect of the DNS is already known to be efficient. Any new service would have to undergo a period of gradual maturation, with potentially problematic early-stage behaviors. By (re-)using the DNS, DKIM avoids these growing pains.

### 1.3 Internet Mail Background

The basic Internet Email service has evolved extensively over its several decades of continuous operation. Its modern architecture comprises a number of specialized components. A discussion about Mail User Agents (MUA), Mail Handling Services (MHS), Mail Transfer Agents (MTA), Mail Submission Agents (MSA), Mail Delivery Agents (MDA), Mail Service Providers (MSP), Administrative Management Domains (ADMDs), and their relationships can be found in [Appendix A](#).

### 1.4 Discussion Venue

NOTE TO RFC EDITOR: This "Discussion Venue" section is to be removed prior to publication.

This document is being discussed on the DKIM mailing list, [ietf-dkim@mipassoc.org](mailto:ietf-dkim@mipassoc.org).

## 2. The DKIM Value Proposition

The nature and origins of a message are often falsely stated. Such misrepresentations may (but not necessarily) be employed in order to perpetrate abuse. DKIM provides a foundation for distinguishing legitimate mail, and thus a means of associating a verifiable identifier with a message. Given the presence of that identifier, a receiver can make decisions about further handling of the message, based upon assessments of the identity that is associated with the identifier.

Receivers who successfully verify a signature can use information about the signer as part of a program to limit spam, spoofing, phishing, or other undesirable behavior. DKIM does not, itself, prescribe any specific actions by the recipient; rather it is an enabling technology for services that do.

These services will typically:

1. Determine a verified identity, if possible.
2. Determine whether a known identity is trusted.

The role of DKIM is to perform the first of these; DKIM is an enabler for the second.

### 2.1 Identity Verification

Consider an attack made against an organization or against customers of an organization. The name of the organization is linked to particular Internet domain names (identifiers). One point of leverage for attackers is either to use a legitimate domain name, without authorization, or to use a "cousin" name that is similar to one that is legitimate, but is not controlled by the target organization. An assessment service that uses DKIM can differentiate between domains used by known organizations and domains used by others. As such, DKIM performs the positive step of identifying messages associated with verifiable identities, rather than the negative step of identifying messages with problematic use of identities. Whether a verified identity belongs to a Good Actor or a Bad Actor becomes a later step of assessment.

### 2.2 Enabling Trust Assessments

Email receiving services are faced with a basic decision: Should they deliver a newly-arrived message to the indicated recipient? That is, does the receiving service trust that the message is sufficiently "safe" to be viewed? For the modern Internet, most receiving services have an elaborate engine that formulates this quality assessment. These engines take a variety of information as input to the decision, such as from reputation lists and accreditation services. As the engine processes information, it raises or lowers its trust assessment for the message.

DKIM provides additional information to this process by declaring a valid "responsible" identity about which the engine can make quality assessments. By itself, a valid DKIM signature neither lowers nor raises the level of trust associated with the message, but it enables other mechanisms to be used for doing so.

An organization might build upon its use of DKIM by publishing information about its Signing Practices (SP). This could permit detecting some messages that purport to be associated with a domain, but which are not. As such, an SP can cause the trust assessment to be reduced, or leave it unchanged.

## 3. DKIM Goals

DKIM adds an end-to-end authentication mechanism to the existing email transfer infrastructure. This motivates functional goals about the authentication itself and operational goals about its integration with the rest of the Internet email service.

### 3.1 Functional Goals

#### 3.1.1 Use Domain-level granularity for assurance

DKIM seeks accountability at the coarse granularity of an organization or, perhaps, a department. An existing Internet service construct that enables this granularity is the Domain Name [RFC1034]. DKIM binds the signing key record to the Domain Name. Further benefits of using domain names include simplifying key management, enabling signing by the infrastructure as opposed to the MUA, and potential privacy issues.

Contrast this with OpenPGP and S/MIME, which provide end-to-end validation in terms of individual authors, notably using full email addresses.

#### 3.1.2 Implementation Locality

Any party, anywhere along the transit path can implement DKIM signing. Its use is not confined to the end systems or only in a boundary MTA.

#### 3.1.3 Allow delegation of signing to independent parties

Different parties have different roles in the process of email exchange. Some are easily visible to end users and others are primarily visible to operators of the service. DKIM was designed to support signing by any of these different parties and to permit them to sign with any domain name that they deem appropriate (and for which they hold authorized signing keys.) As an example an organization that creates email content often delegates portions of its processing or transmission to an outsourced group. DKIM supports this mode of activity, in a manner that is not normally visible to end users.

#### 3.1.4 Distinguish the core authentication mechanism from its derivative uses

An authenticated identity can be subject to a variety of processing policies, either ad hoc or standardized. The only semantics inherent to a DKIM signature is that the signer is asserting (some) responsibility for the message. All other mechanisms and meanings are built on this core service. One such mechanism might assert a relationship between the signing identity and the author, as specified in the From: header field's domain identity[RFC2822]. Another might specify how to treat an unsigned message with that From: field domain.

#### 3.1.5 Retain ability to have anonymous email

The ability to send a message that does not identify its author is considered to be a valuable quality of the current email service that needs to be retained. DKIM is compatible with this goal since it permits authentication of the email system operator, rather than the content author. If it is possible to obtain effectively anonymous accounts at example.com, knowing that a message definitely came from example.com does not threaten the anonymity of the user who authored it.

## 3.2 Operational Goals

### 3.2.1 Treat verification failure the same as no signature present

As a sub-goal to the requirement for transparency, a DKIM signature verifier is to treat messages with signatures that fail as if they were unsigned. Hence the message will revert to normal handling, through the receiver's existing filtering mechanisms. Thus, DKIM specifies that an assessing site is not to take a message that has a broken signature and treat it any differently than if the signature weren't there.

Contrast this with OpenPGP and S/MIME, which were designed for strong cryptographic protection. This included treating verification failure as message failure.

### **3.2.2 Make signatures transparent to non-supporting recipients**

In order to facilitate incremental adoption, DKIM is designed to be transparent to recipients that do not support it. A DKIM signature does not "get in the way" for such recipients.

Contrast this with S/MIME and OpenPGP, which modify the message body. Hence, their presence is potentially visible to email recipients, whose user software needs to process the associated constructs.

### **3.2.3 Permit incremental adoption for incremental benefit**

DKIM can immediately provide benefits between any two organizations that exchange email and implement DKIM. In the usual manner of "network effects", the benefits of DKIM increase dramatically as its adoption increases.

Although it is envisioned that this mechanism will call upon independent services to aid in the assessment of DKIM results, they are not essential in order to obtain initial benefit. For example DKIM allows (possibly large) pair-wise sets of email providers and spam filtering companies to distinguish mail that is associated with a known organization from mail that might deceptively purport to have the affiliation. This in turn allows the development of "whitelist" schemes whereby authenticated mail from a known source with good reputation is allowed to bypass some anti-abuse filters.

In effect the email receiver is using their set of known relationships to generate their own reputation data. This works particularly well for traffic between large sending providers and large receiving providers. However it also works well for any operator, public or private, that has mail traffic dominated by exchanges among a stable set of organizations.

Management of email deliverability problems currently represents a significant pain point for email administrators at every point on the mail transit path. Administrators who have deployed DKIM verification have an incentive to evangelize the use of DKIM signatures to senders who may subsequently complain that their email is not being delivered.

### **3.2.4 Minimize the amount of required infrastructure**

A new service, or an enhancement to an existing service, requires adoption in a critical mass of system components, before it can be useful. The greater the number of required adopters, the higher the adoption barrier. This becomes particularly serious when adoption is required by independent, intermediary -- that is, infrastructure -- service providers. In order to allow early adopters to gain early benefit, DKIM makes no changes to the core Internet Mail service and, instead, can provide a useful benefit for any individual pair of signers and verifiers who are exchanging mail. Similarly, DKIM's reliance on the Domain Name System greatly reduces the amount of new administrative infrastructure that is needed across the open Internet.

### **3.2.5 Permit wide range of deployment choices**

DKIM can be deployed at a variety of places within an organization's email service. This permits the organization to choose how much or how little they want DKIM to be part of their service, rather than part of a more localized operation.

## 4. DKIM Function

DKIM has a very constrained set of capabilities, primarily targeting email while it is in transit from an author to a set of recipients. It creates the ability to associate verifiable information with a message, especially a responsible identity. When a message does not have a valid signature associated with the author, DKIM SP will permit the domain name of the author to be used for obtaining information about their signing practices.

### 4.1 The Basic Signing Service

With the DKIM signature mechanism, a signer chooses a signing identity based on their domain name, performs digital signing on the message, and records signature information in a DKIM header field. A verifier obtains the domain name and the "selector" from the DKIM header field, queries for a public key associated with the name, and verifies the signature.

DKIM permits any domain name to be used for signing, and supports extensible choices for various algorithms. As is typical for Internet standards, there is a core set of algorithms that all implementations are required to support, in order to guarantee basic interoperability.

DKIM permits restricting the use of a signature key (by using `s=`) to signing messages for particular types of services, such as only for email. This is intended to be helpful when delegating signing authority, such as to a particular department or to a third-party outsourcing service.

With DKIM the signer explicitly lists the headers that are signed, such as `From:`, `Date:` and `Subject:`. By choosing the minimal set of headers needed, the signature is likely to be considerably more robust against the handling vagaries of intermediary MTAs.

### 4.2 Characteristics of a DKIM signature

A DKIM signature covers the message body and selected header fields. The signer computes a hash of the selected header fields and another hash of the body. The signer then uses a private key to cryptographically encode this information, along with other signing parameters. Signature information is placed into the DKIM-Signature header field, a new [\[RFC2822\]](#) header field of the message.

### 4.3 The Selector construct

The key for a signature is associated with a domain name, as specified in the `d=` parameter of the DKIM-Signature header. That domain name, or the domain name or address in the `i=` parameter, provide the complete identity used for making assessments about the signer. (The DKIM specification does not give any guidance on how to do an assessment.) However this name is not sufficient for making a DNS query to obtain the key needed to verify the signature.

A single domain can use multiple signing keys and/or multiple potential signers. To support this, DKIM identifies a particular signature as a combination of the domain name and an added field, called the "selector", specified in separate DKIM-Signature header field parameters.

NOTE: The semantics of the selector (if any) are strictly reserved to the signer and should be treated as an opaque string by all other parties. If verifiers were to employ the selector as part of a name assessment mechanism, then there would be no remaining mechanism for making a transition from an old, or compromised, key to a new one.

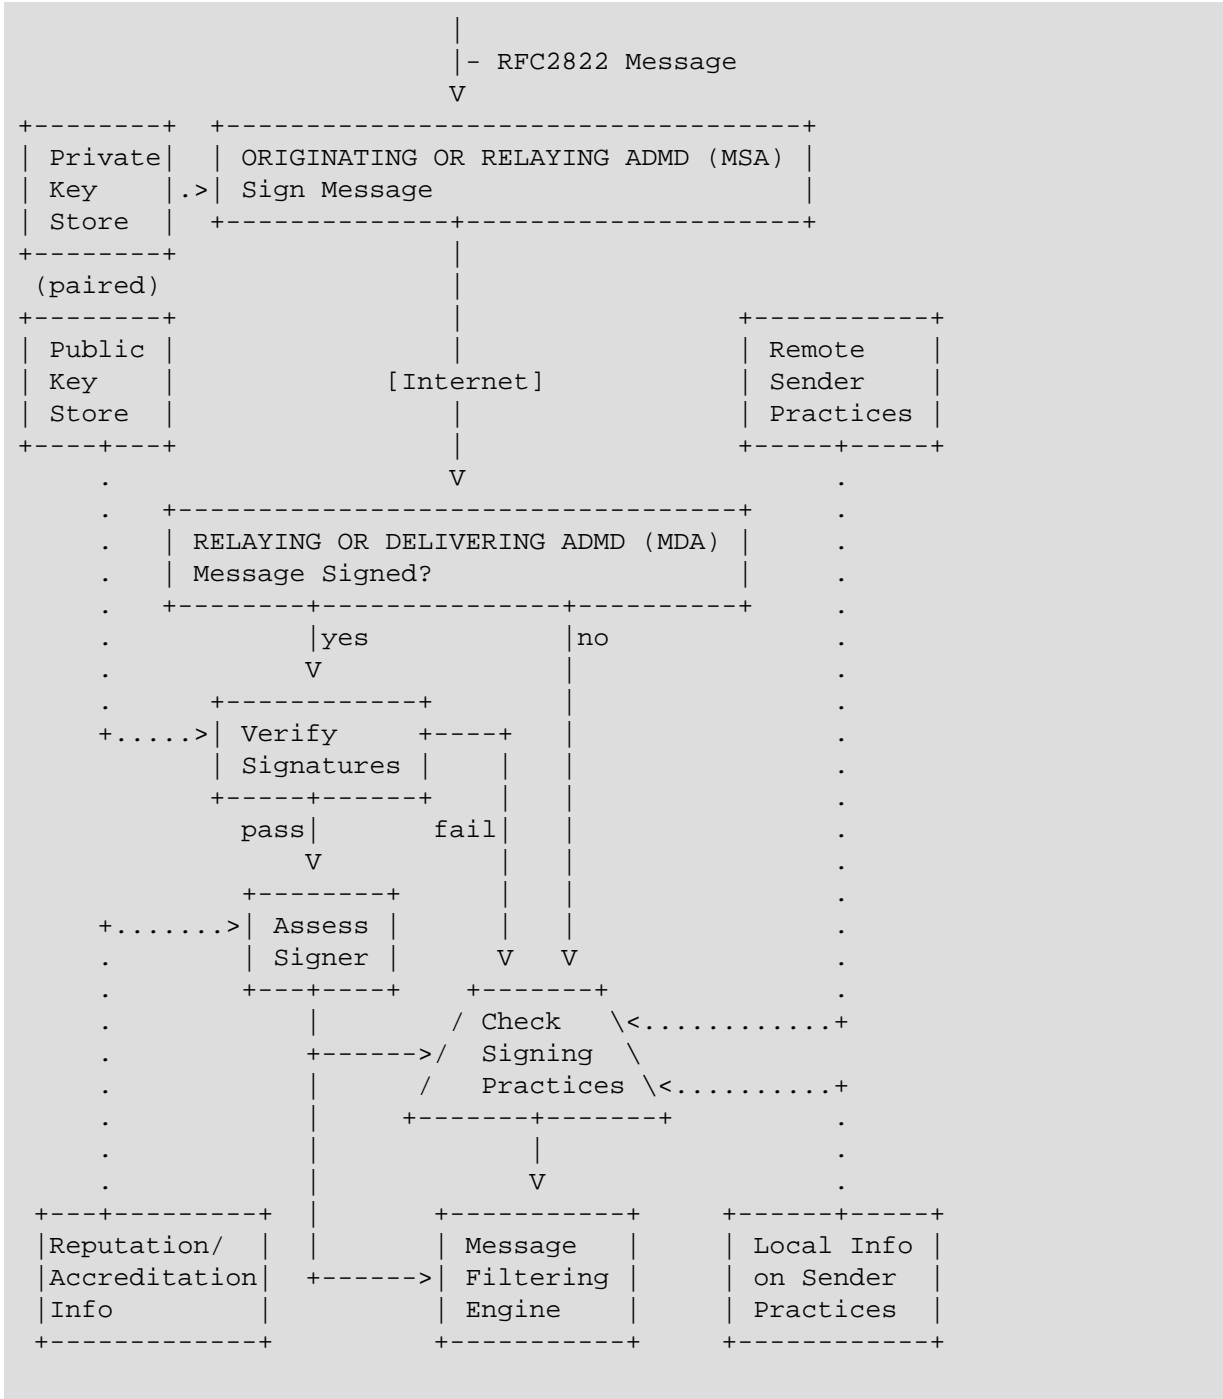
Signers often need to support multiple assessments about their organization, such as to distinguish one type of message from another, or one portion of the organization from another. To permit assessments that are independent, one method is for an organization to use different sub-domains in the `"d="` parameter, such as `"transaction.example.com"` versus `"newsletter.example.com"`, or `"productA.example.com"` versus `"productB.example.com"`.

## 4.4 Verification

After a message has been signed, any agent in the message transit path can verify the signature to determine that the signing identity took responsibility for the message. Message recipients can verify the signature by querying the DNS for the signer's domain directly, to retrieve the appropriate public key, and thereby confirm that the message was attested to by a party in possession of the private key for the signing domain. Typically, verification will be done by an agent in the Administrative Management Domain (ADMD) of the message recipient.

## 5. Service Architecture

The DKIM service is divided into components that are performed using different, external services, such as for key retrieval and relaying email. The basic DKIM signing specification defines an initial set of these services (using DNS and SMTP), in order to ensure a basic level of interoperability.



**Figure 1: DKIM Service Architecture**

As shown in [Figure 1](#), basic message processing is divided between the MSA and the MDA.

- The MSA signs the message, using private information from the Key Store.
- The MDA verifies the signature or determines whether a signature was required. Verifying the signature uses public information from the Key Store. If the signature passes, reputation information is used to assess the signer and that information is passed to the message filtering system. If the signature fails or there is no signature, information about the related signing practices is retrieved remotely and/or locally, and that information is passed to the message filtering system.
- Note: [Figure 1](#) does not show the effects on the message handling when multiple signatures or non-author signatures are present.

## 5.1 Administration and Maintenance

A number of tables and services are used to provide external information. Each of these introduces administration and maintenance requirements.

- Key Store** DKIM uses public/private (asymmetric) key cryptography. The signer uses a private key and the validator uses the corresponding public key. The current DKIM signing specification provides for querying the Domain Names Service (DNS), to permit a validator to obtain the public key. The signing organization therefore must have a means of adding a key to the DNS, for every selector/domain-name combination. Further, the signing organization needs policies for distributing and revising keys.
- Reputation/Accreditation** If a message contains a valid signature, then the verifier can evaluate the associated domain name's reputation. Quality-assessment information, which is associated with a domain name, comes in many forms and from many sources. DKIM does not define assessment services. It's relevance to them is to provide a validated domain name, upon which assessments can be made.
- Signing Practices (SP)** Separate from determining the validity of a signature, and separate from assessing the reputation of the organization that is associated with the signed identity, there is an opportunity to determine any organizational practices concerning a domain name. Practices can range widely. They can be published by the owner of the domain or they can be maintained by the evaluating site. They can pertain to the use of the domain name, such as whether it is used for signing messages, whether all mail having that domain name in the author From: header field is signed, or whether such mail is to be discarded in the absence of an appropriate signature. The statements of practice are made at the level of a domain name, and are distinct from assessments made about particular messages, as occur in a Message Filtering Engine. Such assessments of practices can provide useful input for the Message Filtering Engine's determination of message handling. As practices are defined, each domain name owner needs to consider what information to publish. The nature and degree of checking practices, if any is performed, is optional to the evaluating site and is strictly a matter of local policy.

## 5.2 Signing

Signing can be performed by a component of the ADMD that creates the message, and/or within any ADMD along the relay path. The signer uses the appropriate private key.

## 5.3 Verifying

Verification can be performed by any functional component along the relay and delivery path. Verifiers retrieve the public key based upon the parameters stored in the message.

## 5.4 Unverified or Unsigned Mail

Note that a failed signature causes the message to be treated in the same manner as one that is unsigned. Messages lacking a valid author signature (a signature associated with the author of the message as opposed to a signature associated with an intermediary) can prompt a query for any published "signing practices" information, as an aid in determining whether the author information has been used without authorization.

## 5.5 Assessing

Figure 1 shows the verified identity as being used to assess an associated reputation, but it could be applied for other tasks, such as management tracking of mail. A popular use of reputation information is as input to a filtering engine that decides whether to deliver -- and possibly whether to specially mark -- a message. Filtering engines have become complex and sophisticated. Their details are outside of the scope of DKIM, other than the expectation that the validated identity produced by DKIM will be added to the varied soup of rules used by the engines. The rules can cover signed messages and can deal with unsigned messages from a domain, if the domain has published information about its practices.

## 5.6 DKIM Placement within an ADMD

It is expected that the most common venue for a DKIM implementation will be within the infrastructures of the authoring organization's outbound service and the receiving organization's inbound service, such as a department or a boundary MTA. DKIM can be implemented in an author's or recipient MUA, but this is expected to be less typical, since it has higher administration and support costs.

A Mediator, such as a mailing list, often can re-post a message without breaking the DKIM signature. Furthermore it can add its own signature. This can be added by the Mediator software itself, or by any outbound component in the Mediator's ADMD.

## 6. Security Considerations

The security considerations of the DKIM protocol are described in the DKIM base specification [[RFC4871](#)].

## 7. IANA Considerations

There are no actions for IANA.

NOTE TO RFC

This section may be removed prior to publication.

EDITOR:

## 8. Acknowledgements

Many people contributed to the development of the DomainKeys Identified Mail and the efforts of the DKIM Working Group is gratefully acknowledged. In particular, we would like to thank Jim Fenton for his extensive feedback diligently provided on every version of this document.

## 9 Informative References

- [I-D.kucherawy-sender-auth-header] Kucherawy, M., "Message Header Field for Indicating Message Authentication Status", Internet-Draft draft-kucherawy-sender-auth-header-11 (work in progress), February 2008.
- [Kohnfelder] Kohnfelder, L., "Towards a Practical Public-key Cryptosystem", May 1978.
- [RFC0989] Linn, J. and IAB Privacy Task Force, "[Privacy enhancement for Internet electronic mail: Part I: Message encipherment and authentication procedures](#)", RFC 989, February 1987.
- [RFC1034] Mockapetris, P., "[Domain names - concepts and facilities](#)", STD 13, RFC 1034, November 1987.
- [RFC1848] Crocker, S., Galvin, J., Murphy, S., and N. Freed, "[MIME Object Security Services](#)", RFC 1848, October 1995.
- [RFC1991] Atkins, D., Stallings, W., and P. Zimmermann, "[PGP Message Exchange Formats](#)", RFC 1991, August 1996.
- [RFC2440] Callas, J., Donnerhacke, L., Finney, H., and R. Thayer, "[OpenPGP Message Format](#)", RFC 2440, November 1998.
- [RFC2821] Klensin, J., "[Simple Mail Transfer Protocol](#)", RFC 2821, April 2001.
- [RFC2822] Resnick, P., "[Internet Message Format](#)", RFC 2822, April 2001.
- [RFC3156] Elkins, M., Del Torto, D., Levien, R., and T. Roessler, "[MIME Security with OpenPGP](#)", RFC 3156, August 2001.
- [RFC3164] Lonvick, C., "[The BSD Syslog Protocol](#)", RFC 3164, August 2001.
- [RFC3851] Ramsdell, B., "[Secure/Multipurpose Internet Mail Extensions \(S/MIME\) Version 3.1 Message Specification](#)", RFC 3851, July 2004.
- [RFC4406] Lyon, J. and M. Wong, "[Sender ID: Authenticating E-Mail](#)", RFC 4406, April 2006.
- [RFC4407] Lyon, J., "[Purported Responsible Address in E-Mail Messages](#)", RFC 4407, April 2006.
- [RFC4408] Wong, M. and W. Schlitt, "[Sender Policy Framework \(SPF\) for Authorizing Use of Domains in E-Mail, Version 1](#)", RFC 4408, April 2006.
- [RFC4686] Fenton, J., "[Analysis of Threats Motivating DomainKeys Identified Mail \(DKIM\)](#)", RFC 4686, September 2006.

- [RFC4870] Delany, M., "[Domain-Based Email Authentication Using Public Keys Advertised in the DNS \(DomainKeys\)](#)", RFC 4870, May 2007.
- [RFC4871] Allman, E., Callas, J., Delany, M., Libbey, M., Fenton, J., and M. Thomas, "[DomainKeys Identified Mail \(DKIM\) Signatures](#)", RFC 4871, May 2007.
- [RFC4880] Callas, J., Donnerhackle, L., Finney, H., Shaw, D., and R. Thayer, "[OpenPGP Message Format](#)", RFC 4880, November 2007.

## Authors' Addresses

**Tony Hansen**

AT&T Laboratories  
200 Laurel Ave.  
Middletown, NJ 07748  
USA  
EMail: [tony+dkimov@mailennium.att.com](mailto:tony+dkimov@mailennium.att.com)

**Dave Crocker**

Brandenburg InternetWorking  
675 Spruce Dr.  
Sunnyvale, CA 94086  
USA  
EMail: [dcrocker@bbiw.net](mailto:dcrocker@bbiw.net)

**Phillip Hallam-Baker**

VeriSign Inc.  
EMail: [pbaker@verisign.com](mailto:pbaker@verisign.com)

## A. Internet Mail Background

Internet Mail is 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) composed of Mail Transfer Agents (MTA). The MHS is responsible for accepting a message from one user, the author, and delivering it to one or more other users, the recipients. This creates a virtual MUA-to-MUA exchange environment. The first component of the MHS is called the Mail Submission Agent (MSA) and the last is called the Mail Delivery Agent (MDA).

An email Mediator is both an inbound MDA and outbound MSA. It takes delivery of a message and re-posts it for further distribution, retaining the original From: header field. A mailing list is a common example of a Mediator.

The modern Internet Mail service is marked by many independent operators, many different components for providing users with service and many other components for performing message transfer. Consequently, it is necessary to distinguish administrative boundaries that surround sets of functional components, which are subject to coherent operational policies.

As elaborated on below, every MSA is a candidate for signing using DKIM, and every MDA is a candidate for doing DKIM verification.

### A.1 Administrative Management Domain (ADMD)

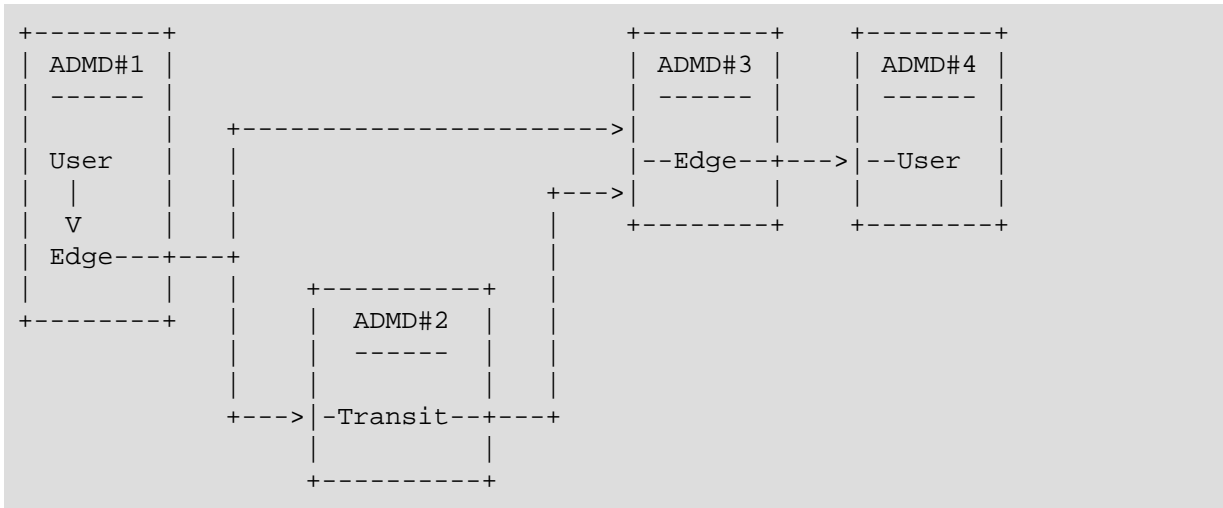
Operation of Internet Mail services is apportioned to different providers (or operators). Each can be composed of an independent ADministrative Management Domain (ADMD). An ADMD operates with an independent set of policies and interacts with other ADMDs according to differing types and amounts of trust. Examples include: an end-user operating their desktop client that connects to an independent email service, a department operating a submission agent or a local Relay, an organization's IT group that operates enterprise Relays, and an ISP operating a public shared email service.

Each of these can be configured into many combinations of administrative and operational relationships, with each ADMD potentially having a complex arrangement of functional components. [Figure 2](#) depicts the relationships among ADMDs. Perhaps the most salient aspect of an ADMD is the differential trust that determines its policies for activities within the ADMD, versus those involving interactions with other ADMDs.

Basic types of ADMDs include:

- Edge: Independent transfer services, in networks at the edge of the Internet Mail service.
- User: End-user services. These might be subsumed under an Edge service, such as is common for web-based email access.
- Transit: These are Mail Service Providers (MSP) offering value-added capabilities for Edge ADMDs, such as aggregation and filtering.

Note that Transit services are quite different from packet-level transit operation. Whereas end-to-end packet transfers usually go through intermediate routers, email exchange across the open Internet is often directly between the Edge ADMDs, at the email level.



**Figure 2: ADministrative Management Domains (ADMD) Example**

In Figure 2, ADMD numbers 1 and 2 are candidates for doing DKIM signing, and ADMD numbers 2, 3 and 4 are candidates for doing DKIM verification.

The distinction between Transit network and Edge network transfer services is primarily significant because it highlights the need for concern over interaction and protection between independent administrations. The interactions between functional components within a single ADMD are subject to the policies of that domain. Although any pair of ADMDs can arrange for whatever policies they wish, Internet Mail is designed to permit inter-operation without prior arrangement.

Common ADMD examples are:

Enterprise Service Providers:

Operators of an organization's internal data and/or mail services.

Internet Service Providers:

Operators of underlying data communication services that, in turn, are used by one or more Relays and Users. It is not necessarily their job to perform email functions, but they can, instead, provide an environment in which those functions can be performed.

Mail Service Providers:

Operators of email services, such as for end-users, or mailing lists.

## Full Copyright Statement

Copyright © The IETF Trust (2008).

This document is subject to the rights, licenses and restrictions contained in BCP 78, and except as set forth therein, the authors retain all their rights.

This document and the information contained herein are provided on an "AS IS" basis and THE CONTRIBUTOR, THE ORGANIZATION HE/SHE REPRESENTS OR IS SPONSORED BY (IF ANY), THE INTERNET SOCIETY, THE IETF TRUST AND THE INTERNET ENGINEERING TASK FORCE DISCLAIM ALL WARRANTIES, EXPRESS OR IMPLIED, INCLUDING BUT NOT LIMITED TO ANY WARRANTY THAT THE USE OF THE INFORMATION HEREIN WILL NOT INFRINGE ANY RIGHTS OR ANY IMPLIED WARRANTIES OF MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE.

## Intellectual Property

The IETF takes no position regarding the validity or scope of any Intellectual Property Rights or other rights that might be claimed to pertain to the implementation or use of the technology described in this document or the extent to which any license under such rights might or might not be available; nor does it represent that it has made any independent effort to identify any such rights. Information on the procedures with respect to rights in RFC documents can be found in BCP 78 and BCP 79.

Copies of IPR disclosures made to the IETF Secretariat and any assurances of licenses to be made available, or the result of an attempt made to obtain a general license or permission for the use of such proprietary rights by implementers or users of this specification can be obtained from the IETF on-line IPR repository at <<http://www.ietf.org/ipr>>.

The IETF invites any interested party to bring to its attention any copyrights, patents or patent applications, or other proprietary rights that may cover technology that may be required to implement this standard. Please address the information to the IETF at [ietf-ipr@ietf.org](mailto:ietf-ipr@ietf.org)<sup>1</sup>.

## Acknowledgement

Funding for the RFC Editor function is provided by the IETF Administrative Support Activity (IASA).

<sup>1</sup> <mailto:ietf-ipr@ietf.org>