DomainKeys Identified Mail (DKIM) Service Overview
draft-ietf-dkim-overview-10

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/1id-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 January 2009.

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. A message can contain multiple signatures, from the same or different organizations involved with the message. DKIM defines a domain-level digital signature authentication framework for email, using public-key cryptography, using the domain name service as its key server technology [RFC4871]. This permits verification of a responsible organization, as well as the integrity of the message 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. DKIM's authentication of email identity can assist in the global control of "spam" and "phishing.

Table of Contents

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 handling the message directly, such as the author's, the originating sending site or an intermediary. It also can also be created by an independent service that is providing assistance to a handler. DKIM defines a domain-level digital signature authentication framework for email through the use of public-key cryptography and using the domain name service as its 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. DKIM's authentication 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

A person or organization has an "identity" -- that is, a constellation of characteristics that distinguish them from any other identity. Associated with this abstraction can be a label used as a reference, or "identifier". (This is the distinction between a thing and the name of the thing.) DKIM uses a domain name as an identifier, to refer to the identity of a person or organization. Note that the same identity can have multiple identifiers.

A DKIM signature can be created by a direct handler of a message, such as the message's author or an intermediary. A signature also can 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 might be 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 can 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 can 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 has 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 signer.
- 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 -- that is, post it as a new message -- with the original signature remaining verifiable, even though the new recipient(s) might be different from those who were originally specified by the author.

1.2. Prior Work

Historically, the IP Address of the system that directly sent the message -- that is, the previous email "hop" -- has been treated as an identity to use for making assessments.[RFC4408], [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 can satisfy this need.

There have been four previous IETF Internet Mail signature standards. Their goals have differed from those of DKIM. The first two are only of historical interest.

Pretty Good Privacy (PGP) was developed by Phil Zimmermann and first released in 1991.

- Privacy Enhanced Mail (PEM) was first published in 1987. [RFC0989]
- PEM eventually transformed into MIME Object Security Services (MOSS) in 1995. [RFC1848] [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. It was standardized as [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 only uses cryptographic 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 management scheme, rather than the more typical approaches based on a certificate in the styles of Kohnfelder (X.509) [Kohnfelder] or Zimmermann (web of trust) [WebofTrust]. For DKIM, the owner of a domain name asserts the validity of a key, rather than having the validity of the key attested to by a trusted third party, often including other assertions, 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 signature verification service.

Further, DKIM's key management 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 This "Discussion Venue" section is to be removed prior to publication.
EDITOR:

This document is being discussed on the DKIM mailing list, ietf-dkim@mipassoc.org.

1.4.1. Changes to document

In addition to simple wordsmithing, the following substantive changes were made:

Service Arch figure and text: (per Allman) Existing figure and text carries vestigial references to role of MSA and MDA. New text switches focus to ADMD more completely and merely cites possible functional modules within them.

Identity vs. Identifier: Added text in Scope to define terms and their relationship.

Message Validity: Added section discussing restricted implication of this.

2. The DKIM Value Proposition

The nature and origins of a message often are falsely stated. Such misrepresentations may be employed for legitimate reasons or for nefarious reasons. 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 as taking responsibility for the message, if possible.

2. Evaluate the trustworthiness of this/these identities.

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). Attackers can leverage either using a legitimate domain name, without authorization, or using 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 is a question for later stages of assessment.

2.2. Enabling Trust Assessments

Email receiving services are faced with a basic decision: Whether to 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.

In order to formulate reputation information, an accurate, stable identifier is needed. Otherwise, the information might not pertain to the identified organization's own actions. When using an IP Address, accuracy is based on the belief that the underlying Internet infrastructure supplies an accurate address. When using domain based reputation data, some other form of validation is needed, since it is not supplied independently by the infrastructure

DKIM satisfies this requirement by declaring a valid "responsible" identity about which the engine can make quality assessments and by using a digital signature to ensure that use of the identifier is authorized. However 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.

2.3. Establishing Message Validity

Though man-in-the-middle attacks are historically rare in email, it is nevertheless theoretically possible for a message to be modified during transit. An interesting side effect of the cryptographic method used by DKIM is that it is possible to be certain that a signed message (or, if l= is used, the signed portion of a message) has not been modified. If it has been changed in any way, then the message will not be verified successfully with DKIM.

As described above, this validity neither lowers nor raises the level of trust associated with the message. If it was an untrustworthy message when initially sent, the verifier can be certain that the message will be equally untrustworthy upon receipt and successful verification.

3. DKIM Goals

DKIM adds an end-to-end authentication capability to the existing email transfer infrastructure. It defines a mechanism that only needs to be supported by the signer and the validator, rather than any of the functional components along the handling path. 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 provides accountability at the coarse granularity of an organization or, perhaps, a department. An existing construct that enables this granularity is the Domain Name [RFC1034]. DKIM binds a 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 reducing privacy concerns.

Contrast this with OpenPGP and S/MIME, which associate validation with individual authors, using their 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 particular systems, such as the author's MUA or the inbound boundary MTA, and there can be more than one signature per message.

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. Similarly, a reputation provider can delegate a signing key for a domain under the control of the provider, to be used by an organization the provider is prepared to vouch for.

3.1.4. Distinguish the core authentication mechanism from its derivative uses

An authenticated identity can be subject to a variety of assessment policies, either ad hoc or standardized. DKIM separates basic authentication from assessment. The only semantics inherent to a DKIM signature is that the signer is asserting (some) responsibility for the message. Hence, a DKIM signature only means that the signer is asserting (some) responsibility for the message, and nothing more. Other services can build upon this core association, but their details are beyond the scope of that core. 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. Make presence of signature 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.

3.2.2. Treat verification failure the same as no signature present

DKIM must also be transparent to existing assessment mechanisms. Consequently, 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.

3.2.3. Permit incremental adoption for incremental benefit

DKIM can be used by any two organizations that exchange email and implement DKIM; it does not require adoption within the open Internet's email infrastructure. In the usual manner of "network effects", the benefits of DKIM increase as its adoption increases.

Although this mechanism can be used in association with independent assessment services, such services are not essential in order to obtain initial benefit. For example DKIM allows (possibly large) pairwise sets of email providers and spam filtering companies to distinguish mail that is associated with a known organization, versus 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 can use 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 delivery 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 might subsequently complain that their email is not being delivered.

3.2.4. Minimize the amount of required infrastructure

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 a wide range of deployment choices

DKIM can be deployed at a variety of places within an organization's email service. This affords flexibility in terms of who administers its use, as well as what traffic carries a DKIM signature. For example, employing DKIM at an outbound boundary MTA will mean that it is administered by the organization's central IT department and that internal messages are not signed.

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 associates 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. Basic Signing

With the DKIM signature mechanism, a signer chooses a signing identity based on their domain name, performs digital signing on the message, and adds the signature information using a DKIM header field. A verifier obtains the domain name and the "selector" from the DKIM header field, obtains the 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 to signing messages for particular types of services, such as only for a single source of 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 applies to 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 DKIM-Signature:, a new [RFC2822] message header field.

4.3. The Selector Construct

The key for a signature is associated with a domain name. That domain name provides 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 using a combination of the domain name and an added field, called the "selector", specified in a separate DKIM-Signature: header field parameter.

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

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 signed 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.

4.5. Sub-Domain Assessment

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". These can be entirely separate from the rfc2822.From header field domain.

5. Service Architecture

DKIM use external service components, such as for key retrieval and relaying email. This specification defines an initial set, using DNS and SMTP, for basic interoperability.

```
                                 |
                                 |- RFC2822 Message
                                 V
+--------+     +-------------------------------+
| Private|     |  ORIGINATING OR RELAYING ADMD |
| Key    +...>|  Sign Message                 |
| Store  |     +---------------+---------------+
+--------+                     |
 (paired)              [Internet]
+--------+                     |                    +-----------+
| Public |     +-------------------------------+    | Remote    |
| Key    |     |  RELAYING OR DELIVERING ADMD  |    | Sender    |
| Store  |     |  Message Signed?              |    | Practices |
+----+---+     +-----+-------------------+-----+    +-----+-----+
     .               |yes               |no              .
     .               V                  |                .
     .         +------------+           |                .
  +........>| Verify     +---------+   |                .
     .         | Signature  |     |   |                .
     .         +------+------+     |   |                .
     .           pass|          fail|   |                .
     .               V              |   |                .
     .         +------------+       |   |                .
     .         |            |       |   |                .
  +........>| Assessments |       |   |                .
     .         |            |       V   V                .
     .         +------+------+    +-------+              .
     .               |          / Check   \<............+
     .          +---------->/  Signing  \                .
     .               |        /   Practices \<..........+
     .               |       +------+-------+            .
     .               |              |                    .
     .               |              V                    .
+----+--------+      |        +-----------+    +------+-----+
|Reputation/  |      |        | Message   |    | Local Info |
|Accreditation|   +---------->| Filtering |    | on Sender  |
|Info         |      |        | Engine    |    | Practices  |
+-------------+               +-----------+    +------------+
```
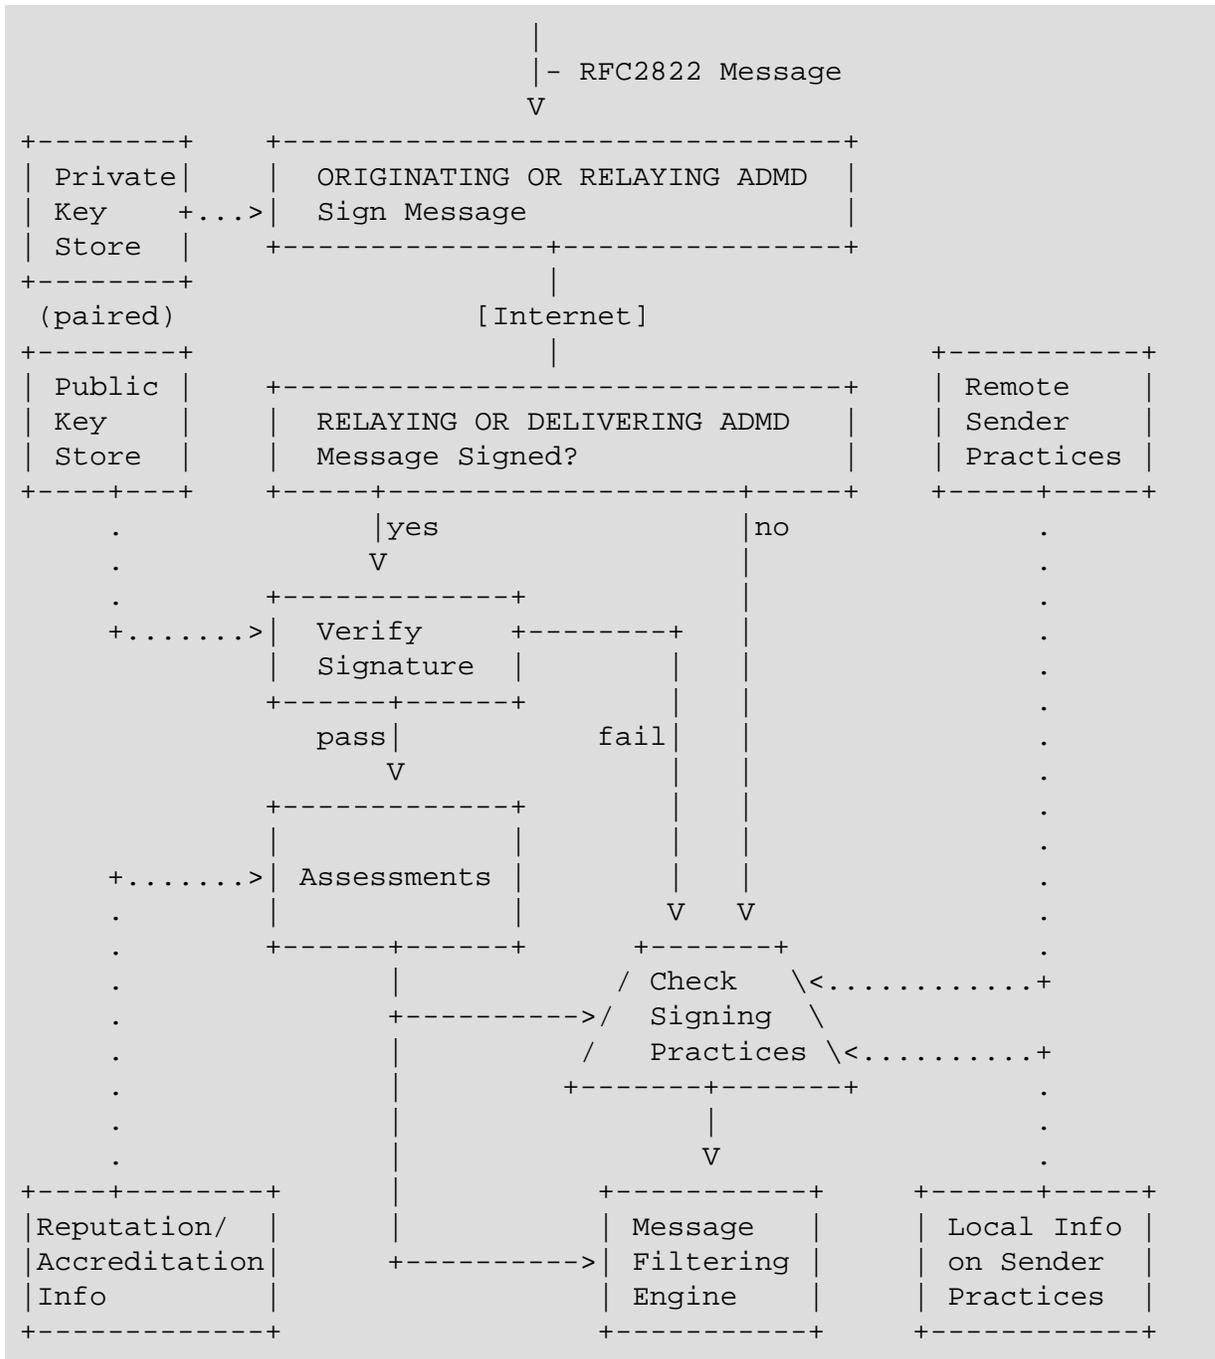
Figure 1: DKIM Service Architecture

As shown in Figure 1, basic message processing is divided between a signing Administrative Management Domain (ADMD) and a validating ADMD. At its simplest, this is between the Originating ADMD and the delivering ADMD, but can involve other ADMDs in the handling path.

Signing: Signing is performed by an authorized module within the signing ADMD and uses private information from the Key Store, as discussed below. Within the originating ADMD, this might be performed by the MUA, MSA or an MTA.

Validating: Validating is performed by an authorized module within the validating ADMD. Within a delivering ADMD, validating might be performed by an MTA, MDA or MUA. The module verifies the signature or determines whether a particular signature was required. Verifying the signature uses public information from the Key Store. If the signature passes, reputation information is used to asses the signer and that information is passed to the message filtering system. If the signature fails or there is no signature using the author's domain, information about signing practices related to the author can be retrieved remotely and/or locally, and that information is passed to the message filtering system.

If message has more than one valid signature, the order in which the signers are assessed and the interactions among the assessments are not defined by the DKIM specification.

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 users 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 needs to 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, in order to determine appropriate delivery or display options for that message. 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 the 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 even whether the domain owner recommends discarding messages 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

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 can accumulate its own reputation, and 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 Processing 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 is an MUA that receives a message and can re-post a modified version of it, such as to a mailing list. A DKIM signature can survive some types of modifications through this process. Furthermore the Mediator 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. Considerations

6.1. Security Considerations

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

6.2. IANA Considerations

There are no actions for IANA.

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

6.3. 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.

7   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-15 (work in progress), July 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., Donnerhacke, L., Finney, H., Shaw, D., and R. Thayer,
                       "OpenPGP Message Format", RFC 4880, November 2007.

[WebofTrust]   Wikipedia, "Web of Trust", URL http://en.wikipedia.org/wiki/
                       Web_of_trust, <http://en.wikipedia.org/wiki/Web_of_trust>.

Authors' Addresses

Tony Hansen
AT&T Laboratories
200 Laurel Ave.
Middletown, NJ 07748
USA
EMail: tony+dkimov@maillennium.att.com

Dave Crocker
Brandenburg InternetWorking
675 Spruce Dr.
Sunnyvale, CA 94086
USA
EMail: dcrocker@bbiw.net

Phillip Hallam-Baker
VeriSign Inc.
EMail: pbaker@verisign.com

A. Internet Mail Background

A.1. Core Model

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, makes changes appropriate to its service, and then re-posts it for further distribution. Typically the new message will retain 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.2. Trust Boundaries

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 often is directly between the Edge ADMDs, at the email level.

```
+--------+                                      +--------+      +--------+
| ADMD#1 |                                      | ADMD#3 |      | ADMD#4 |
| ------ |                                      | ------ |      | ------ |
|        |     +---------------------->|        |      |        |
| User   |     |                              |--Edge--+--->|--User  |
|  |     |     |                       +--->|        |      |        |
|  V     |     |                       |    +--------+      +--------+
| Edge---+---+ |                       |
|        |   | |    +----------+       |
+--------+   | |    | ADMD#2   |       |
            | |    | ------   |       |
            | |    |          |       |
            +--->|-Transit--+---+
                 |          |
                 +----------+
```

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

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 [1].

---

[1] mailto:ietf-ipr@ietf.org

Index

A

ADMD 5
Administrative Management Domain 5
assessment 6

D

DKIM-Signature **11**, 11
DNS 4, 11, 12, 13, 14

I

identifier **3**, 4, 6, 6
identity 3, **3**, 4, 6, 6, 6, 6, 8, 11, 11
infrastructure 4, 4, 4, 6, 8, 8, 9, 10, 15

K

Kohnfelder 4, 17

M

Mail Delivery Agent 5
Mail Handling Service 5
Mail Service Provider 5
Mail Submission Agent 5
Mail Transfer Agent 5
Mail User Agent 5
MDA 5
MHS 5
MIME Object Security Services 4
MOSS 4
MSA 5
MSP 5
MTA 5
MUA 5

O

OpenPGP 4

P

PEM 4
PGP 4
Pretty Good Privacy 4
Privacy Enhanced Mail 4

R

RFC0989 4, 17
RFC1034 4, 8, 17
RFC1848 4, 17
RFC1991 4, 17
RFC2440 4, 17

RFC2822 8, 11, 17
RFC3156 4, 17
RFC3851 4, 17
RFC4406 4, 17
RFC4407 4, 17
RFC4408 4, 17
RFC4686 3, 17
RFC4871 1, 3, 3, 16, 18
RFC4880 4, 18

S

S/MIME 4

T

trust 3, 6, 7, 20

V

verification 3, 6, 7, 9, 9, 12, 15, 20, 20

W

Web of Trust 4
WebofTrust 4, 18

X

X.509 4