



Email Sender Authentication Deployment
Best Practices and Considerations
for Financial Institutions

A Publication of the BITS Security Program
In Partnership with eCert



June 2009

About BITS

A division of The Financial Services Roundtable, BITS was created in 1996 to foster the growth and development of electronic financial services and e-commerce for the benefit of the financial institutions and their customers. BITS works to sustain consumer confidence and trust by ensuring the security, privacy and integrity of financial transactions. BITS provides intellectual capital and addresses emerging issues where financial services, technology and commerce intersect, acting quickly to address problems and galvanize the industry. BITS' efforts involve representatives from throughout our member institutions, including CEOs, CIOs, CISOs, and fraud, compliance and vendor management specialists.

About the BITS Security Working Group

The mission of the BITS Security Working Group is to strengthen the security and resiliency of financial services by:

- Sharing and developing best practices to secure infrastructures, products and services;
- Maintaining continued public and private sector confidence; and
- Providing industry input to government agencies and regulators on policies and regulations.

The priorities of the Security Working Group are determined and reviewed by BITS governance bodies. The focus of the Security Working Group may vary year to year but includes four major areas: Application Security, Infrastructure Security, Data/Information Security and People-related Security.

About eCert

eCert ensures trust in, and reduces fraud against, critical email traffic. eCert is a trusted intermediary that accredits domains that send email and certifies their traffic to defend against email fraud (“phishing”) and enhance delivery. eCert does this by enabling standards for receivers to identify certified traffic and block phish before it reaches customer inboxes. eCert also provides traffic statistics and data on phishing attacks from major ISPs and other receivers. eCert was founded as a collaboration between large financial service companies and major ISPs to improve security against email phishing. The eCert collaboration offers participating companies the opportunity to take a leadership position, with major financial and Internet service industry leaders, in addressing the problem of phishing and security for their customers.

For more information please contact eCert, information@ecertsystems.com.

Table of Contents

1. Scope and Intended Audience of This Document.....	4
2. Sender Authentication Overview and Purpose	5
3. General Considerations for Deployment.....	11
4. DomainKeys Identified Mail (DKIM)	19
5. Sender Policy Framework (SPF/SIDF)	25
6. Metrics and Reporting	31
Glossary	33
Appendix A: Resources	35
Appendix B: Sample Letter of Intent.....	37
Appendix C: Sample Implementation Project Plan.....	39
Appendix D: Creating a Private-Public Key Pair using OpenSSL.....	41
eCert’s Contribution.....	43
Acknowledgements	44

1. Scope and Intended Audience of This Document

1.1 Overview

This *Email Sender Authentication Deployment: Best Practices and Considerations for Financial Institutions* is intended to contribute to the body of information on security for financial services email. In particular, it is designed to support improvements in the email security channel through the use of authentication as a foundational tool to reduce email fraud (“phishing”), as recommended in the *BITS Email Security Toolkit: Protocols and Recommendations for Reducing the Risks (Toolkit, April, 2007)*¹.

1.2 Financial Institutions

Historically, phishing has primarily targeted the financial services industry and those companies that provide services such as online payment and auction web sites². Large financial institutions, many of which have relationships with a significant market share of U.S. retail banking, payments, credit cards, or mortgages, are generally the larger targets³, because any given phishing campaign will be distributed to a greater percentage of the target company’s actual customers.⁴

While this document focuses on the challenges larger, complex financial institutions face when implementing technologies such as Sender Policy Framework (SPF) or DomainKeys Identified Mail (DKIM), smaller organizations can use the same steps to implement these protocols.

1.3 Audience

The intended audience of this document is employees of all levels, with some degree of responsibility for email systems, services and strategy, at financial institutions of all sizes. Although this document is not a technical deep-dive into the intricacies of Simple Mail Transfer Protocol (SMTP), Domain Name System (DNS), or cryptography, some high-level knowledge of how email works, the current risks associated with online fraud, and information security practices will be useful in understanding these discussions.

1.4 BITS Email Security Toolkit

The document continues the work begun by the BITS Security Working Group in the *Toolkit*. This document serves as a companion to help financial services companies implement the recommendations from the *Toolkit* in the challenging, multi-faceted environment of a large financial institution. This document assumes the understanding of, or access to, technical and expository information and technical recommendations contained in the *Toolkit*.

¹ <http://www.bits.org/downloads/Publications%20Page/BITSSecureEmailFINALAPRIL1507.pdf>

² Anti-Phishing Working Group (APWG) Phishing Trends Activity Report, Q2 2008, page 7, http://www.antiphishing.org/reports/apwg_report_Q2_2008.pdf

³ APWG Phishing Trends Activity Report, 2nd Half 2008, page 7, http://www.antiphishing.org/reports/apwg_report_H2_2008.pdf

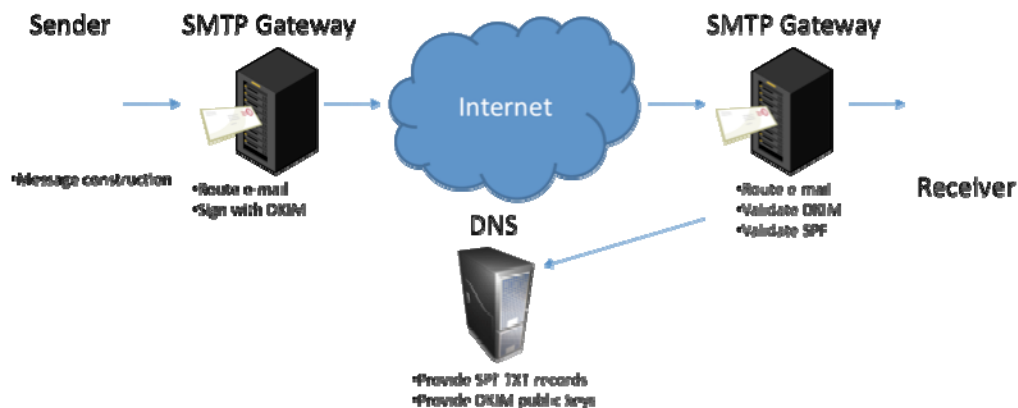
⁴ A phishing campaign comprised of 20,000,000 messages targeting customers of a global retail banking organization will have a much greater penetration than a similar sized campaign against a small local bank, whose customer base may not even total 10% of the number of phish emails sent.

2. Sender Authentication Overview and Purpose

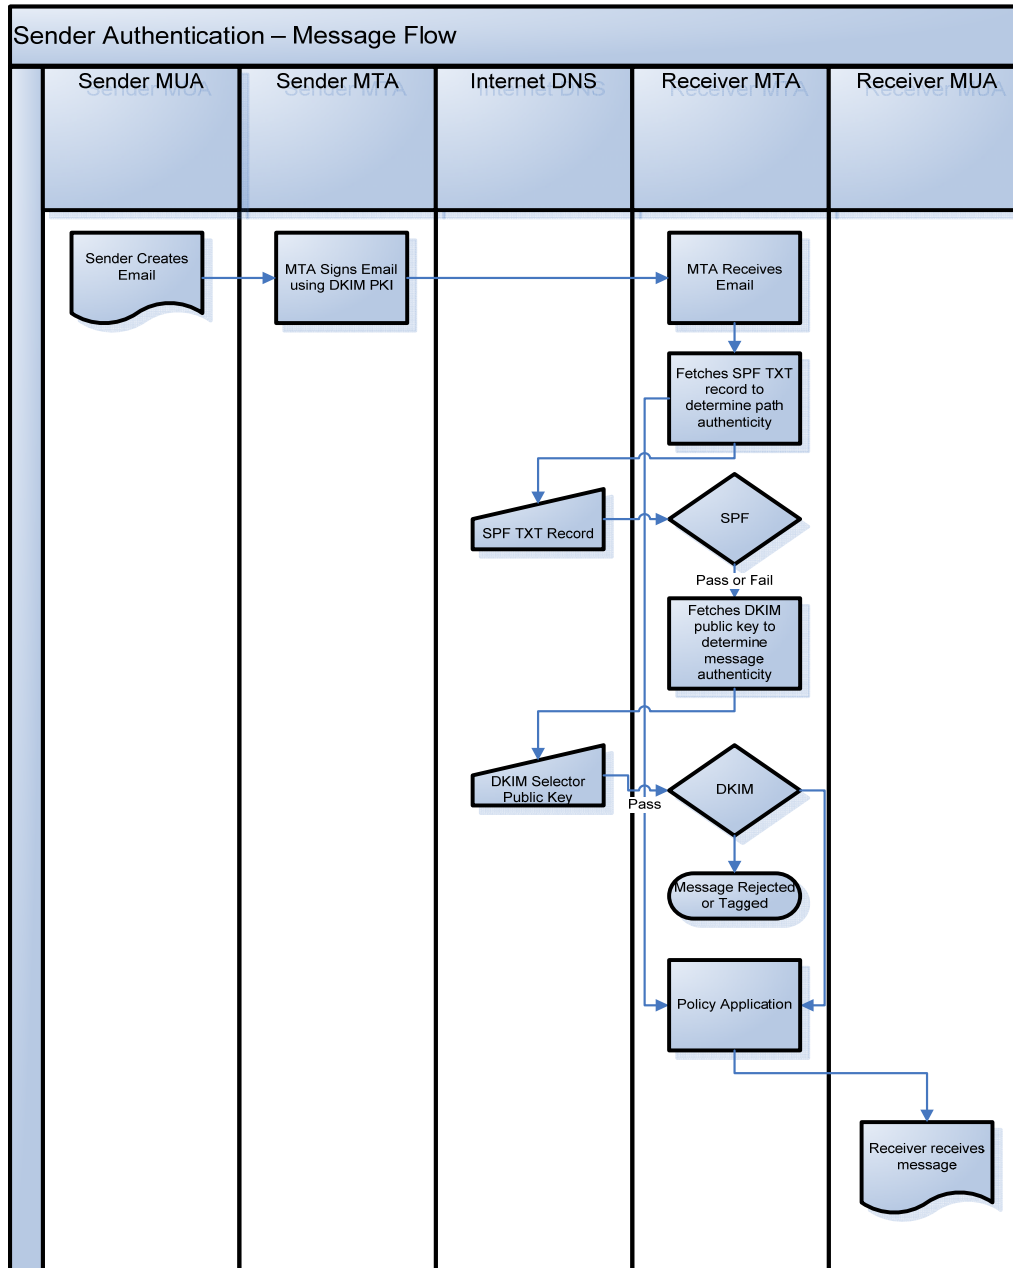
2.1 Sender Authentication

Email sender authentication is not a new phenomenon, but recent activity by the Internet Engineering Task Force (IETF) to create protocol standards around SPF and DKIM has allowed organizations to safely begin supporting the technology. Simply put, sender authentication allows organizations to positively identify legitimate messages as originating from an authorized source. This source may be the organization itself or a third party, thus ensuring a valid identity for the email message.

At a high level, sender authentication verification contains the following components:



A functional diagram representing message flow across infrastructure layer is represented in the diagram below:



When discussing sender authentication, it is important to differentiate this area from two other major parts of the email security portfolio, anti-spam technology, and secure email encryption technology.

Sender authentication does not claim to provide the following:

- **Anti-spam capability.** Spammers can utilize email authentication technology for messages they send from their controlled SMTP servers. They can create SPF records that assign the entire Internet as valid senders for the domain they are using (assuming they are using a domain they own). They can also include DKIM signatures for spam messages. The presence of an SPF record or DKIM signature does not mean a message is not necessarily spam (note that the definition of spam is highly subjective as well).
- **Secure (encrypted) email.** Federal regulations⁵ require non-public information contained in email transmissions from financial institutions be encrypted. Sender authentication does not provide any encryption of the message body – the message is still sent clear-text and does not constitute a “secure” email. However, messages secured via Transport Layer Security (TLS) or third-party encryption software may also be validated by the presence of an SPF record or included DKIM signature.
- **Defensive registration for fraud protection.** Defensive registration of domains that look similar to a domain (“cousin domains”), but exist for the purpose of tricking customers into thinking the domain is legitimate, is a common mitigation technique. This can be done in a variety of ways, for example using the domain examplebank.com:
 - Misspelled domain names: examplesbank.com
 - Using letters or numbers that look like another letter or number: examp1ebank.com
 - Using a non-existent sub-domain: examplebank.mynewcard.com
 - Using a domain that resembles a legitimate brand owned by the company: examplobanko.co.mx

Since these domains may legitimately exist – and in some case be legally owned by someone else – these domains may actually have valid, authenticated email from them. Since SPF and DKIM rely on control of the DNS records of the company brand domain, these cousin domains will be under control of a third party – whether that third party is legitimately or illegitimately doing business with the domain is a subjective question that can only be answered on a case-by-case basis. Since it is impossible for a company to assert – positively or negatively – email ownership for these cousin domains, sender authentication does not provide a solution. Rather, the company needs to leverage its brand property legal teams to deal with cybersquatters and cousin domain holders.

The cousin domain problem is seen as a significant one, and will continue to be, as companies deploy sender authentication and protect the brands/domains they own. Other solutions, such as reputation or domain accreditation services, will provide the necessary identification of “legitimate” bank-owned domains which will allow Internet Service

⁵ Gramm-Leach Bliley Act <http://www.ftc.gov/privacy/privacyinitiatives/glbact.html>

Providers (ISPs) to differentiate the email from “good” or from “bad” or “suspect” domains.

Another case to consider is the plan by the Internet Corporation for Assigned Names and Numbers (ICANN) to open up top-level domain (TLD) registrations in 2009. This will eliminate the restriction of the number of TLDs (.com, .edu, .org, .biz, etc) and allow a virtually unlimited number of TLDs – anyone who applies may have their own TLD. This will potentially exacerbate the cousin domain problem further, as we see the ability to nest branded names into the new TLD structure to form otherwise legitimate domain names. Some examples are:

- citigroup.finance
- wellsfargo.creditcard
- jpmorgan.investments

In these cases, reputation and accreditation services will become even more important, as it will become logistically and financially impossible for companies to protectively register cousin domains or prosecute people who create a domain based on a brand name under a new TLD. For example, can you legally prevent someone from registering wachovia.bank if someone legitimately owns .bank? Currently, there is no legal precedent or ICANN guideline for domain name conflicts of this nature.

2.2 Purpose

The ability to identify legitimate messages through email authentication enables (but does not, by itself, deliver) important new capabilities for protecting the email channel from fraud and improving delivery performance. Authentication is the foundation for hardening the email channel and for a roadmap of security improvements in identifying and stopping phishing, while ensuring the safe delivery of valid traffic.

i. Blocking unauthenticated traffic

One of the foundational capabilities enabled by sender authentication is the receiver’s ability to block traffic that is not authenticated to prevent unauthorized domain spoofing on email. This feature is the starting point in defending the channel against phishing, and the basic building block for expanding defenses.

Any blocking capability is facilitated by accurately tracing the origination of traffic, but is only fully enabled with the presence of the ability for senders to communicate, and receivers to enforce, policies based on the result of authentication.

ii. Reducing false positives and improving delivery

Authentication assures the origin of messaging traffic, laying the foundation for receiver systems and security applications to avoid blocking and filtering valid traffic. By itself, however, authentication does not provide sufficient information for receivers or their applications to identify traffic that should not be aggressively filtered (phishers and spammers can authenticate, too, using non-spoofed domains, and are often early adopters).

Certification capabilities that identify a class of pre-qualified senders, or reputation capabilities that apply historical performance information, are required counterparts to leverage authentication to reduce false positives (having your company's email incorrectly marked as spam) and improve delivery success.

iii. Enabling improvements to filtering

In the future, consistent authentication of traffic from identified domains may support useful enhancements to filtering of traffic that uses institution trademarks or is identifiable through other signatures. By identifying and protecting all valid traffic, calibration to reduce false positives could be adjusted to support more aggressive filtering against phishing attacks.

The term “spoofing” can be somewhat subjective, and, depending on the audience, a good or bad thing. Spoofing, generally speaking, is when a third party sends email from a domain as if it is the domain owner (when, in fact, it is not).

Authorized spoofing, or *delegation*, occurs when a company authorizes the third party to use its domain name to send email (e.g., outsourced marketing or transactional email sent by email service providers).

Unauthorized spoofing occurs when a third party represents email from a domain that it does not own and is not authorized to send on behalf of. *Phishing* occurs when unauthorized spoofing is done with the purpose to trick the recipient into disclosing personal information, visiting a malicious website or installing malware to their computer so criminals may obtain access to the victim's financial accounts. For the purposes of this document, we use the term *spoofing* to represent *unauthorized spoofing* or *phishing* and not the legitimate use of authorized third-party senders.

2.3 Absent or Insufficient Capabilities for Using Authentication Protocols

In the drive to improve the email channel by leveraging new improvements to security and performance, technical approaches to authentication have outpaced other key enabling functions.

i. Policy

In order for receiving networks to reliably block traffic based on authentication results, the ability to receive, interpret and apply policy information from the sending domain is essential. Not all domain infrastructures are “ready” for traffic to be blocked. Authentication protocols were designed to allow networks to communicate the status of their authentication services and the action that receiving networks should apply at that time.

Today, the protocols and specifications for authentication policy remain, in some cases, incomplete (see [Section 4.5 ADSP](#)) and, in most cases, inconsistently applied. The vast majority of receiving networks, including the major consumer ISPs in the U.S., do not block traffic based on authentication results because the probability of false positives (marking legitimate email as spam) remains far too high due to inconsistency in the application of authentication protocols and the lack of sender policy information.

Proposed policy specifications currently moving through the standards process should help alleviate some of these challenges. For financial senders, however, this process and the subsequent pace of adoption may be insufficiently speedy to address challenges in the next 2-3 years.

ii. Identity

While authentication provides the ability to validate messages coming from a domain, it does not communicate information about the domain itself. The email community lacks a coherent approach to qualifying or certifying domains to enable receiving networks to make decisions about the safety and desirability of the traffic, rather than relying on the purported origin of the message.

iii. Metrics and reporting

Authentication activities occur between sending and receiving networks, with the majority of information regarding results and actions generated by receiver systems. Currently, there is no mechanism to communicate this data to senders, creating a gap in information about the operation of authentication capabilities, the impact on email traffic activity, and importantly, the presence of spoofing activity and data that may help mitigate risk.

3. General Considerations for Deployment

3.1 Organizational Email Disciplines

Some companies maintain disparate technology groups responsible for network infrastructure or email, based on lines of business, acquired subsidiaries, etc. These groups may have independent control over the email services provided to their respective companies yet might overlap with the use of company brands and domains. Other companies will consolidate email operations across the enterprise, centralizing responsibility in one department to the greatest extent possible. However, even with one global email services group, large enterprises typically still have informal, (i.e., “rogue”) departmental email deployments. These deployments are most commonly seen in the marketing lines of business.

The disadvantages of having separate email groups is that they each have to carefully coordinate all the steps of email authentication deployment together. This coordination is critically important when these groups have common use over company brand domains, as deployment of sender authentication by one group affects all the others.

Additionally, there will be situations where email delivery may be partially or completely outsourced to a third-party information technology group. In these situations, it is important to include in the discussions all necessary personnel from these groups, as well as business management responsible for ongoing contract maintenance between the company and its IT vendor.

3.2 Departmental Organization and Participation

Since email is a communications channel leveraged throughout all lines of business, and in some businesses is considered a “mission critical” or “tier 1” service, enterprises are particularly sensitive to issues around email. When planning a sender authentication deployment, the company must consider the teams across the organization that will need to be involved. It is always best to include representatives from areas that have the potential to positively sponsor, or negatively impact, a project of this magnitude.

These areas include:

i. Network computing/Infrastructure/Traditional email teams

These teams are primarily responsible for email infrastructure and support – the systems that will be DKIM signing email, as well as the hosting networks and Internet-facing systems for corporate email. These groups may also manage the corporate and Internet DNS systems responsible for hosting SPF TXT records and DKIM public keys.

ii. Information security

These teams are responsible for ensuring systems are properly protected behind firewalls. They also are often responsible for drafting security policy across technology applications, including email.

iii. Operational risk

Within financial services, operational risk works closely with regulatory compliance and audit teams to ensure that overall activities do not incur too much risk for any particular project,

or line of business, which may cause business to be interrupted or company reputation to be tarnished.

iv. E-commerce/Online properties

These areas typically manage the company's web presence, notably online banking, credit card, mortgage application process, etc. E-commerce is generally responsible for most of the high-risk (from a phishing perspective) transactional email that is delivered to customers. Examples include online banking alerts, confirmations, and encrypted statement delivery.

v. Marketing

Marketing is usually the most prolific senders of email within an organization, because their job is to market and sell company products and services using email and other channels. Marketing is also potentially problematic because marketing groups tend to be distributed among lines of business and typically will outsource campaigns to third-party vendors.

vi. Supply chain/Procurement/Vendor management

This area is responsible for approving new vendor relationships, maintaining existing vendor relationships, writing and enforcing supply contracts, and in some cases, for managing outsourced contracting agencies. These groups will establish policies that require new and existing vendors, by contract, to provide an email service to support your corporate email sender authentication practices.

vii. Corporate communications/Public relations

Communications/PR will approve all broad-reaching internal project communications to lines of business, and also assist in determining the best way to reach deep into the enterprise to target individuals and teams that need to know appropriate details and required actions of the project. They will also provide support to deal with external requests for information from third-party vendors and press.

viii. Privacy and legal

It will be important for Legal to understand the ramifications of the sender authentication program – especially edge cases where lines of business may have email tagged as spam or dropped by an ISP if they do not follow appropriate email authentication policies.

Refer to [Appendix B](#) for a sample communication to internal business units to communicate sender authentication program details, as well as requirements for outsourced third parties.

3.3 Political Issues and Policy Determination

As mentioned above, email is a tool that is deployed across the entire organization. Email is generated by people as well as by applications. It will be sent to other applications for processing and people external to the company – the company's customers, business partners and alumni. It is in the best interests of the company to establish policies that govern email generation, appearance, and other conventions that will allow their email to be recognized by its receivers. In many situations, in addition to the company website, email is the "face" of the organization that is most familiar to its customers.

i. Email "from" policy

Please see the [Glossary](#) for definitions of the two different "from" fields: envelope sender

address and header sender address. It is important to establish conventions for how the header sender address, also known as the “friendly from” or “from name,” is going to appear in recipients’ email clients. The header sender address is the one most commonly displayed to the email recipient in their client interface. Although the email may be sent from `cgable@examplebank.com`, the “from name” may be configured to display Clark Gable or Example Bank, or Example Bank Inc. Companies that send email using the envelope header `examplebank.com` but have no convention on the header will confuse customers, especially those lines of business that use their line of business internal common name that is relatively unknown outside their immediate business network or established customer base. These conventions and policies may be applied consistently across a brand, or a line of business, or across the entire enterprise, as appropriate for each organization.

ii. Email content

In addition to what the customer typically first sees in their inbox (the “header sender address”), once they open the message their first impression is made by the overall look and feel of the email. Some companies have standard disclaimers, trade membership logos and company logos. Some companies use a generic salutation while others use a salutation that includes the recipient’s name. Although there are innumerable options, the company should establish policy on how the email is going to be presented, in a consistent way, across all business units that leverage a like brand. This will establish a comfortable “look and feel” of the company’s communication, much like a business that uses consistent letterhead, envelopes and format in written communications.

As with the email “from” policy, email content policy can be applied to maintain consistency across an enterprise, or across a line of business (each business keeping its own look and feel “identity”), as appropriate.

Some will argue that establishing a static, consistent look and feel with the header sender address and the message body will allow phishers to more easily create recognizable and trusted fraud email. The fact is, many times the phishers create email messages that look more professional and believable than many of the organization’s authorized communications. Creating email consistency across a brand will enable a company (and a customer) to more easily spot a fraudulent message sent by an amateur fraudster that contains typos and non-standard form.

iii. Domain provisioning policy

Many companies have lax or no control over business units registering new Internet domains. Many times, a business will roll out a new brand or product and register a domain to support the new products or services. They will often not think about the security implications of how that new domain now can be used as a conduit for fraudulent email to customers of that service.

It is important for companies to write and enforce policy that requires proper vetting of new domains (and sub-domains) for business use – especially those that will be used to send email. These policies should establish the proper control points and approval steps required to maintain the domain, as well as the information required to register the domain in the sender authentication system.

3.4 Architectural and Technical Considerations

Companies should evaluate several architectural and technical issues, including the two below.

i. Message Transfer Agent (MTA) support

MTA support for sender authentication protocols is slowly increasing as vendors recognize the importance of email security and developers write plug-ins for standards-based SMTP systems. Nevertheless, slow uptake in demand, performance considerations, and gaps in authentication capabilities have slowed some vendors in incorporating authentication into their roadmaps. In those situations, it is best to leverage contract renewals with supply chain's assistance to encourage the MTA vendor to support SPF/SIDF and DKIM.

Although for the financial institution, sender (outbound) DKIM and SPF authentication is the highest priority, many companies will want to eventually, if not immediately, seek to leverage inbound authentication to filter fraudulent email being sent to their associates. Because of this, it is best to have an MTA that supports both DKIM signing and validation, as well as SPF validation (there is no requirement for an MTA to support SPF outbound since it relies completely on DNS in that direction).

ii. What network layer to deploy

In nearly all situations, DKIM signatures will be applied to a message at the perimeter MTA. The perimeter MTA is the email gateway server that is the last hop between the company and the Internet. This approach allows the message to travel through the company without the need for internal systems to preserve the DKIM signature, as modification of the message header will necessitate the signature to be re-applied to reflect the change in message status.

For systems that validate inbound DKIM, the best place for this activity is either at the edge MTA or an intermediate MTA that will perform other hygiene activities such as anti-virus and anti-spam. In most cases, authentication check results will be an input into the anti-spam heuristics engine.

3.5 Third-Party Senders

Companies should construct an "email sender inventory" to capture all entities – companies, applications, systems – that send email on the company's behalf.

The sender inventory may consist of the following fields:

- i. Compliance level.** Used to track overall compliance of the sending entity with corporate policies and technical protocol support (DKIM, SPF/SIDF, etc).
- ii. Priority.** Used to indicate the relative importance of a sender entity. For example, the main corporate MTA will take priority over a smaller application that may maintain its own MTA, but rarely sends customer email.
- iii. Email domain.** The domains the entity uses to send email. These may be major corporate brand domains, or sub-domains.

- iv. **Line of business and technical contact information.** The internal “owner” of the system who has responsibility for complying with policy.
- v. **Third party business and technical contact information.** The external “owner” of the system who is responsible for complying with policy, if this is a system that is managed by a vendor or subsidiary.
- vi. **SPF status and details.** Information on how far along the entity is with SPF compliance. Also contains the DNS TXT record defined for the SPF/SIDF record, which in turn delineates authorized systems and the SPF fail indicator.
- vii. **DKIM status.** Information on how far along the entity is with DKIM compliance. Also contains the DKIM selector designation and DKIM reflector test results (see DKIM Section 4). It may contain a copy of the DKIM public key for reference.

Refer to [Appendix B](#) for a sample communication to communicate sender authentication program details as well as requirements for outsourced third parties.

3.6 Remote Employees

Remote employees may send email in a number of ways. These methods are described and discussed as follows.

- i. **Sending email directly from the employee PC to the Internet.**

This method is sometimes, but rarely, used by companies that allow employees to send email using their corporate brand domains via a telecommuting ISP connection, hotel or kiosk connection.

Although it is technically possible to establish DKIM signing for this arrangement, pursuant to having a compliant Message User Agent (MUA), the administrative overhead is much too difficult to maintain given the sheer number of employees in a typical large company, as well as the impossibility of maintaining an authorized list of sending IP host addresses for employees who could be literally anywhere in the world on any network. The only SPF solution is to enable the entire Internet as authoritative for the sending domain, which would negate any benefits of authentication and most likely cause the domain to be blacklisted by reputation systems. In most instances, this method is not practically feasible.

- ii. **Using a third-party messaging application such as a contact management system or web email that spoofs the corporate email domain.**

This method can be handled the same way as third-party senders described above. This option has the least control over SMTP traffic. The difficulty in administering this is directly proportional to the number of allowed third-party messaging systems used by employees. In many cases, smaller third-party systems allow open spoofing, do not allow for static IP address assignments, and will not allow a successful sender authentication deployment. It should be noted that some third-party senders may not be technically capable of supporting authentication.

- iii. **Using a local SMTP server that is connected directly to the Internet.**

For this case, each SMTP server will need to be treated as its own separate, authorized sending entity.

IP addresses for each SMTP server will need to be registered in the sender inventory and added to SPF records. In cases where these machines are NATed⁶ behind a firewall and assigned private IP addresses, the public-facing address or address range will need to be static and monitored for changes.

Each remote office SMTP server will need to run an MTA that is capable of signing messages with DKIM; a public/private key policy will need to be applied to the server; private keys should be distributed by IT to the remote server administration personnel; and public keys that are used to sign mail on the server should be added to the inventory of DKIM selectors used for that domain or sub-domain for which the remote office is responsible. There may be security disadvantages to this option, as private keys will be stored in multiple locations that may not be safeguarded as well.

iv. Using a local network SMTP server that routes email to corporate SMTP gateways which connect directly to the Internet.

This is the most common way companies deal with remote employees or satellite offices. The local SMTP server simply does its job as a basic mail router and does not have to perform any kind of signing or validation. The mail is routed to the corporate gateways, which are centrally managed and have the appropriate software for signing and validation (and other important hygiene steps such as anti-virus). This model is least efficient from a “hop count” model, but most efficient in every other way, including administration, security, and cost.

This method is also used for remote employees that connect to the enterprise via VPN. They are essentially “internal” to the network and will send email via their enterprise mail platform.

3.7 Domain Segmentation

Companies frequently use separate domains and sub-domains for corporate, transactional, and marketing email. In addition, they may assign these domains to third-party senders for their use.

Email from a company can be broadly assigned to one of three categories:

- **Marketing:** Email that has a primary purpose of selling, cross-selling or informing customers or potential customers of products and services.
- **Transactional:** Email that contains information about, or awareness of, a transaction the customer is affected by or has performed. Examples include alerts for direct deposit, payment confirmations, or scheduled bill reminders. Transactional email has information that directly relates to the customer in a unique and personal way.

⁶ Network Address Translation

- **Corporate:** Email sent and received in the normal course of business activity among employees, business partners, vendors and personal email.

During a company's inventory activity to determine which lines of business have their own email systems or third-party email vendors, it will become apparent that there are several, if not dozens or more, domains in use for sending email. There will be significant challenges for companies that allow third parties to use their main corporate domain (e.g., examplebank.com) for sending email. There will be issues related to SPF TXT record size, as well as defining selectors and building public keys for companies that have email coming from one domain but several systems, some of which may not be directly controlled.

It will be much easier to perform a controlled rollout if third parties and LOB-specific systems can be delegated their own domain or sub-domain for email use. For example, the main domain for Example Bank may be examplebank.com, but the following domains can be distributed for use by the following lines of business:

- card.examplebank.com: credit card main domain
- studentcard.examplebank.com: credit card marketing campaign for students
- onlinealert.examplebank.com: transaction email for online alerts

Many companies will want to use their primary domain for as much email traffic as possible since that domain is well known to customers. However, as email filtering systems include things like reputation score for domains, in the absence of specialized services, having the primary domain used for marketing as well as transactional mail may cause transactional mail to have a high spam score since a high proportion of marketing email gets flagged as spam by customers, which in turn causes ISPs to have a heavy-handed approach to future email from that domain.

In the context of sender authentication, it will be easier to maintain authorized host IP addresses for SPF and selectors and keys for DKIM if segmentation is maintained for third parties and lines of business. As authentication matures and systems are built to provide granular metrics and improved quality of service for transactional and corporate email, business value will increase as a more complete view of email delivery becomes available. The domain segmentation approach makes it easier for ISPs to better report-out and for businesses to sort through the data to measure specific and relevant success.

3.8 Protecting Brand Domains That Do Not Send Email

To specially protect brand domains that do not send mail, companies can proactively assign SPF records to those domains with "no send" parameters. A common tactic of online fraudsters is to send emails from esoteric domains that may be owned by the target company. These domains pass lookup checks because they do exist, but they may never be used to send email from the company. The customer doesn't know this, and neither will the ISPs, so any email sent from the company may pass an initial visual or machine inspection. With the advances of reputation systems, this problem is mitigated somewhat, but the method will still be used by criminals.

Companies can address this by publishing SPF records for all their domains, even if they do not send email. The following syntax is an example of this:

```
"v=spf1 -all"
```

Email Sender Authentication Deployment

This SPF syntax simply states that the domain for which the SPF record is contained, no email is sent and the policy for any email purporting to be from the domain is hard fail.

This will cause any email (fraudulent as well as valid) from the domain to be classified as unauthorized spoofed or phished email by an SPF validation.

The administrative cost of this activity is relatively low since it involves one extra task for the DNS administrators to add the TXT record for new domains that have no email intention. Of course, if the domain were to be used for email in the future, the SPF record would need to be updated to reflect the valid hosts.

To the extent that services are available to publish information about domains for an institution, these domains should be included.

3.9 Managing the Process

With the above considerations in mind, an organizational plan should be crafted to optimize the deployment and successful use of authentication.

A sample of such a plan is provided in [Appendix C](#).

4. DomainKeys Identified Mail (DKIM)

4.1 Brief DKIM History

DKIM was created by a collaboration of industry experts in 2004⁷. The key contributors to the protocol were Yahoo, responsible for DomainKeys (DK), and Cisco's Internet Identified Mail (IIM), whose main features were combined to form the DomainKeys Identified Mail (DKIM) protocol. The resulting specification was refined and standardized as part of an Internet Engineering Task Force (IETF) working group and was published as a standards-tracked specification RFC 4871⁸ in November 2006.

Some companies, notably Yahoo, have supported DomainKeys for several years. DK has been designated as a historic protocol by the IETF, with the intention that it will eventually be replaced by DKIM. Companies that currently support DK may continue to do so but will need to deploy updated software that leverages the legacy keys that were used for DK.

4.2 High-level Technical and Functional Overview

This document does not attempt to serve as an in-depth tutorial on DKIM and the intricacies of the protocol. For detailed information on DKIM, refer to the *Toolkit* and the [Resources Appendix A](#) for web links.

DKIM allows an organization to claim responsibility for, and assert the identity of, a particular email message that it, or a trusted third party, has sent. The responsible organization adds a digital signature to the email message, tying it to a domain name the organization owns. Messages are signed using public key cryptography (RSA algorithm). The signing action must be performed by any system maintained by the responsible organization, including the MUA, Mail Submission Agent (MSA), or, most often, the edge MTA-gateway SMTP router. After a message has been signed, any device that message comes in contact with may validate the signature. Usually, this validation will be done by the edge MTA or other MTA of the recipient. Typically, any heuristic or policy determination on the disposition of an email will be made by one of these MTAs, and not by the recipient's MUA or by the recipient himself.

The high-level requirements of DKIM configuration and setup are:

- The message body is made canonical, then hashed (typically SHA-256⁹). Canonicalization helps prevent signature breakage from minor changes incurred by the message during transmission.
- Message headers are included in the signature. The header fields that are included are determined by the signing organization as part of DKIM configuration.
- A new field is appended to the message header named DKIM-Signature that contains important information required by the verifier, including the body hash, the selector name, the list of headers in the signature, and the canonicalization algorithm.
- The header fields and the DKIM signature field itself are canonicalized and a hash of them is created.

⁷ DKIM www.dkim.org

⁸ RFCs <http://www.ietf.org/rfc.html>

⁹ Secure Hash Algorithm published by NIST <http://csrc.nist.gov/groups/ST/hash/index.html>

- An RSA signature is generated for the hash, and the signature is inserted back into the DKIM-Signature field.
- The full DKIM-Signature field is added to the header of the message, and the modified message is ready for transmission.

Below is an example of a DKIM header included in an SMTP message transmission from an example bank's domain, examplebank.com, using the selector 09-tx-lob. Note the selector (s=), domain (d=), and hashed headers (h=):

```
DKIM-Signature: v=1; a=rsa-sha256; c=relaxed/simple;
d=examplebank.com;
s= 09-tx-lob; t=1233611805;
bh=IohO4h5JgZbIJn2plMBlowgxB/SCuKtsFlU/iK3Ke
TI=; h=Date:From:Subject:To:Message-id:MIME-version:Content-type:
Content-language:Importance:Accept-Language:Thread-topic:
Thread-index:X-MS-Has-Attach:X-MS-TNEF-Correlator:
X-Proofpoint-Virus-Version;
b=bNzxG87h1dKpa/TX8B5tPY3K3Cw3qX8zDJnc
EiCBB/wNprTHYQGFdxZ0uP+XN1Fk9QmXrAPx2oTewuH0WNpnvBeozVUo9qpw8VUT
30
e6KZMC231Uo7ZaYbcvDb6aPflndyZX8oav5JnXr33GXX1pa2dLZCDdJxUUP5ZWBrA
Ieo=
```

Compare the above signature to the one below, which is also from examplebank.com, but from their online banking alert system. Note the selector (= 09-tx-alerts) and the subdomain (= marketing.examplebank.com)

```
DKIM-Signature: v=1; a=rsa-sha256; c=relaxed/relaxed; s=09-tx-
alerts; d=marketing.examplebank.com;
h=From:To:Subject:Date:List-Unsubscribe:MIME-Version:Reply-
To:Content-Type:Content-Transfer-Encoding:Message-Id;
i=onlinebanking@marketing.examplebank.com;
bh=ov0qr9N5GvayX1kfPyuljOL2ySE=;
b=LxCF+L+F6ltgxfker9eFHqZ1D4dgVpAIjkTz4veTpH8wnTk//jOVkevejkkOGKs
OYVqzubSF4TQO
wlhwjqONKbmkMKJjj9phfBTTwjZQb7uSd8rUbmKqPFhhrfQr4cS3zpy46S+/L5mIT
9mYEEY7EDf1/
HYv0MfvIwovgx23qi8M=
```

4.3 Signing Versus Verifying

The act of signing a message asserts the “ownership” or identity of the message as being sent from the authorized sender. Because the sender has created the DKIM header information based on the public/private key-encoded hash, and the public key is kept in DNS, only the owner of the domain for which the message is identified can correctly sign messages for that domain.

The owner of the domain name used for a DKIM signature puts their reputation on the line for the veracity of that message. Verifying a message can be done by any system that may have access to that message because it has been routed to (or through) it. Verification means that the message's DKIM header has been identified, and the recipient system has fetched the purporting domain's appropriate DKIM public key from DNS based on the selector information contained in the header. Receivers who validate a signature may then use the information on the sender's domain as part of

anti-spam heuristics, which may include reputation data collected about that sender from other parties.

4.4 Technical Implementation and Best Practices

This document does not serve as a technical deep-dive manual on deploying DKIM. For technical details, refer to the [Resources Appendix A](#). Where appropriate, certain steps of the deployment process may include some technical information in order to provide the reader with enough background information to understand the steps involved, as well as what kind of technical resources they should have at their disposal for a successful deployment.

i. Creating selectors

Selectors, which are mechanisms to subdivide the key namespace, allow a single domain to utilize multiple keys to support periodic key changes and overlap of keys while old ones are phased out, which are recommended security practices. Support of multiple keys for the same domain is also useful for organizations that must support third-party use of the domain. Since the corporation will not want any third parties to obtain their private keys, a special signing key may be created for the third party, or the third party may create their own if they have access to OpenSSL or other company-approved key generating mechanisms. Having a third party use a dedicated and unique key pair assigned by the domain owner, selectors allow the domain owner more flexibility in dealing with a security breach in case the third-party systems are compromised and their private key exposed. As an alternative, companies can also allow third parties to create their own keys, as the company will still have control over the DNS publishing of the public key.

There should be consistent naming around selectors across the organization and its authorized third parties. Naming conventions should be communicated as part of the project plan in order to prevent any entities from creating selectors outside the standard, which would then need to be recreated or grandfathered.

Selector names should contain information on the creation date (which allows for key rotation), as well as administrative responsibility.

For example, the scheme: date.purpose.admin

- Date: used to delineate the creation date for selector rotation scheduling
- Purpose: the primary purpose of the domain: marketing, transactional, or corporate
- Admin: the administrative entity responsible for the domain or sub-domain, which may be a designation of the company itself or a third-party agency.

The selector name can only contain dots, letters, numbers or the hyphen character, although hyphens cannot be the first or last character.

ii. Creating a public/private key pair

Please refer to [Appendix D](#) for an example of creating a public/private key pair using OpenSSL.

iii. Signing policy designation

Companies will want to consider the following steps for signing policy designation.

- Determine the selector naming conventions and inventory the selectors associated with each sending entity (see [Appendix A](#) and [Creating Selectors](#) above).
- Canonicalization settings should be set to relaxed, at least during testing and through the beginning of production. This will allow toleration of minor modifications such as header line wrapping changes, capitalization, and whitespace replacement. Changing canonicalization to simple is more secure, but tolerates almost no modification and may result in invalidated signatures. Each organization will need to determine which setting works best in their environment.
- Signing and verification algorithms will be rsa-sha1, which is the default if no algorithm is specified and MUST be supported by all implementations.
- Key sizes should be no less than 1024-bit and the largest practical key size is 2048-bit. (2048 is the largest key that fits within a 512-byte DNS UDP response packet).
- Length tag (L=) is not recommended used for security purposes¹⁰
- Required tags include:
 - a= : algorithm used to generate the signature (plain-text: rsa-sha1)
 - b= : signature data (base64)
 - c= : Header/body canonicalization (plain text: relaxed/simple; simple/simple)
 - d= : domain of signing entity (plain text: company.com)
 - h= : signed header fields (plain text: a colon-separated list of headers presented to the signing algorithm)
 - FROM, SUBJECT, DATE – must be included
 - All MIME header fields should also be included
 - Any other header field that describes the role of the signer (e.g., SENDER or RESENT-FROM must be included)
 - Any header fields that are transient (perhaps added by an application to be later removed by a downstream application) or likely to be modified or removed in transit should NOT be included
 - The DKIM-Signature header field is always implicitly signed and must not be included in the h= tag; for further information on header tag, refer to the header tag section of the DKIM specification¹¹.
 - s= : selector (plain text: date.purpose.admin: the selector subdividing the namespace for the d= domain tag)
- Optional tags should be researched prior to use, especially x= (signature expiration) and t= (signature timestamp)

iv. **Install and configure supported software**

Upgrade or install software for the MTA(s) that will be signing the messages. Depending on the size of the organization, and whether this is a test or production deployment, this process can take a few minutes, or several weeks. For production deployment, this step should have a formal project framework and consultation of the MTA vendor. See [Appendix A](#) for links to configuration information for some common MTAs.

v. **Test and record results**

¹⁰ RFC 4871, Section 3.4.5 Body Length Limits <http://www.ietf.org/rfc/rfc4871.txt>

¹¹ RFC 4871, Section 5.4 <http://www.ietf.org/rfc/rfc4871.txt>

Typically, most installations will cause all messages passing through the signing MTA to be DKIM signed.

The following test cases can be used as a starting point to develop your own testing routine:

1. Send a message to an external account that passes through the signing MTA and visually inspect the message for the presence of a DKIM header and corresponding message hash.
2. Send a message the same way to an external DKIM mirror site that will perform a DKIM validation and display, or email back, results.
3. Testing should be done to validate DKIM results for every domain, and all selectors for domains, on all MTAs that are signing mail.
4. Create a script or other automated system that periodically checks email sent from various systems and campaigns, including third-party email that should be signed, from various systems, to monitor for broken or missing signatures, etc.

vi. Infrastructure impact

Since DKIM is a cryptographically-based digital signature solution, there will be some CPU processing overhead associated with the activity of canonicalization, hashing, and signing. Email systems by nature tend to be more I/O intensive than CPU intensive, so most SMTP servers will have the necessary free overhead to handle DKIM signing.

If the SMTP system is already delegated to some CPU-intensive activity such as anti-virus, anti-spam, or encryption, then it is recommended that performance testing be done on similarly configured lab hardware. Also, if the performance of the MTA hardware is already at the limits of its capability, then implementing DKIM may be enough to push it over the edge and experience performance problems such as queuing.

It is impractical to outline a rule of thumb for performance, since overhead impact will depend heavily on the hardware available as well as the MTA software (and even the MTA version) and how it is configured. However, anecdotal evidence to date suggests that DKIM signing adds 1% or less additional strain on system CPU performance, and some even suggest that it is a “drop in the bucket” compared to processes like anti-spam which may be running on the same machine. Performance will vary from product to product. It is important to test DKIM signing in a lab environment to determine performance impact and whether additional system capacity will be needed.

vii. Selector and key security and rotation

The purpose of selector and key rotation is to augment security and mitigate breaches in the event keys are compromised. In some organizations, keys may “leak” as administrators move from job to job, may not be secured properly, or may be maliciously shipped to individuals with criminal intent.

There will be two ways to rotate keys: either via prearranged schedule, or ad-hoc. Ad-hoc should be done if there is any suspicion that a key has been compromised. Not only should the affected key be replaced, but any key that is under similar storage conditions. The same process for key and selector replacement should be followed for scheduled maintenance, as well as if it is being done to address a sudden security concern (unscheduled maintenance).

Key replacement allows a company to move, in a controlled manner, from one selector naming routine to another, change the strength of key encryption, or enforce security policy around key retention.

As described above, a selector is created for a domain (or multiple per domain) which specifies the public key to retrieve from DNS. When keys are rotated, there will be a period of time when messages signed by the new selector/key will exist on the Internet alongside the old selector/key. It is important to keep the transition period where both selectors/keys are supported until the organization determines the population of old selectors/keys have aged enough to be removed from DNS with minimal disruption. Some security groups indicate this can be as long as several months, but the timeline can be accelerated under exigent circumstances.

Since DKIM can be validated at the MUA, allowing keys to deprecate can result in valid messages being considered invalid if a key is expired and the MUA tries to validate an old message. To date, we see very little validation being performed by the MUA, but rather the MTA while messages are in active transit. This being the most common case, it makes sense that a company determine for itself when old keys and selectors can be retired by simply monitoring what they see in transit using automated tests and monitoring described above. In other words, as part of routine testing, a company should be able to determine the state of DKIM signatures on its own messages, as well as determine that old selectors/keys are no longer being detected by test systems. At that point, it is safe to retire the old keys.

4.5 Author Domain Signing Practices (ADSP)

DKIM is an evolving protocol and, as such, there are several ongoing efforts to expand or modify its capabilities. One such example, Author Domain Signing Practices (ADSP)¹², formerly known as Sender Signing Policy (SSP), is a proposed optional extension to DKIM for message senders to communicate to receivers which of their domains are signing all their messages. Currently, ADSP is an IETF Internet-Draft. Due to both technical and political controversy, ADSP is still in flux and generally has not been implemented in production environments.

¹² Author Domain Signing Practices <http://www.dkim.org/specs/draft-ietf-dkim-ssp-10.html>

5. Sender Policy Framework (SPF/SIDF)

5.1 Brief SPF History

Internet experts were discussing ways to validate message authenticity as far back as 1997. In 2003, a paper written by Paul Vixie in 2002, titled “*Repudiating Mail-From*” was presented at the O’Reilly Open Source Convention. Subsequently, Meng Weng Wong, founder of Pobox.com, merged some early specifications and began soliciting input from other interested parties. Over the next year, the specification was fine-tuned and the name changed from Sender Permitted From to Sender Policy Framework. In 2004 the IETF created the MARID working group which attempted to merge the SPF specification with a similar one from Microsoft named CallerID. The MARID group collapsed in 2005, and SPF proponents succeeded in having SPF designated an “experimental” RFC (RFC4408) in late April, 2006.

Microsoft continues to promote its version of SPF, called SenderID Framework (SIDF). There was some early controversy around SenderID as Microsoft patented and developed licensing of the key differences between SPF and SenderID, which caused developers to avoid SenderID due to licensing concerns. In October 2006, Microsoft designated SenderID intellectual property under an Open Specification Promise, which essentially makes SenderID compatible with other free and open source licenses, with the notable exception of incompatibility with the current version of the General Public License (GPL) version 3.x¹³.

There are mixed opinions (and deployment cases) on SPF vs. SenderID. For a review of the open source community’s view of SenderID, refer to the [OpenSPF.ORG SPF community position statement](http://www.openspf.org/whPubsDetail.aspx?publication=3846)¹⁴.

For detailed technical descriptions of SPF and SIDF, refer to links in [Resources Appendix A](#).

5.2 High-level Technical and Functional Overview

Sender Policy Framework is a path-based authentication protocol that also communicates policy, and is intended to prevent unauthorized spoofing of the envelope sender address in an email.

SPF provides the framework for domain owners to specify which mail servers (hosts) are allowed to send email from those domains. The domain owner publishes information into an SPF record in the domain’s DNS. Then, when the recipient’s mail system receives the message, it checks whether the message actually was delivered to it by an authorized host, thereby checking the “path” the message traveled. A message that originated from an unauthorized host, that is, one that is not specified in the SPF DNS record, would be considered spoofed.

Contained within the SPF record is also policy information to instruct the receiving MTA on how to dispose of the message. Today, most systems use SPF policy information to inform heuristics (combined with reputation and other data) for filtering messages. In fact, systems can use SPF information to improve existing reputation data by aggregating over time the sending pattern of specific hosts and domains. The policy information integrated into SPF protocol is simple, yet allows receiving systems a way to gradually “turn up the heat” on fraudulent messages.

¹³ SenderID and the General Public License (GPL)
<http://www.wilmerhale.com/publications/whPubsDetail.aspx?publication=3846>

¹⁴ <http://www.openspf.org/blobs/spf-community-position>

Today, however, policy is not directly enforced by receivers due to concerns about deployment inconsistencies leading to unacceptable rates of false positives. There have been examples where a company deployed SPF, but they did not adequately collect metrics for what the Internet was seeing regarding the company's email status and SPF pass/fail results. For example, a company published SPF hard-fail too soon, and did not realize that a large and important line of business, which maintained a separate email infrastructure, was not accounted for. That line of business experienced loss of email service when some ISPs obeyed the hard-fail policy and rejected all non-validated SPF. This situation resulted in a reputation hit for the company, and an embarrassing retreat from their sender authentication effort.

Due to the implications of the ascribed example above, ISPs have been reluctant to apply SPF policy, mainly because a company may indicate a fail policy without adequate testing, resulting in valid messages being dropped if the ISP enforces the fail indicator. Therefore it is very important to coordinate uniform deployment and to test thoroughly before moving from neutral (?all) or soft-fail (~all) to hard-fail (-all).

For this reason, most ISPs will not obey SPF policy. They are correct in assuming that if a company does not publish SPF correctly and email is rejected (as is the correct policy with hard fail), then their customers will hold the ISP responsible for missing email. However, it is the goal of most financial institutions to encourage ISPs to enforce fail policies (-all) in order to provide a mechanism to block unauthorized messages from arriving at customer inboxes.

Coordinated efforts to improve the validation of successful deployments should lead to SPF policy enforcement that supports automatic dropping for messages that fail authentication.

5.3 Publishing Considerations

i. Keep host IP address inventory and SPF records updated

The most important things to consider when publishing SPF records are the items are covered in [Section 3: General Considerations for Deployment](#). Special attention must be paid to the inventory process for sending hosts, as well as the ongoing process to track sending host IP address changes. A common mistake is having a sending host IP address change that is not updated in the SPF record, effectively putting that host “outside” the list of approved systems. This situation will cause email received from that host to fail an SPF check. It is important to ensure that when third-party senders make a change to their host addresses, bring up a new host, or shift sending responsibility from one host to another, they communicate the changes well in advance so the SPF records can be modified and DNS given enough time to propagate changes.

ii. Publish on the correct DNS server

Since SPF lookups are conducted over the Internet, it is important to publish SPF records on DNS servers that are authoritative for the domains on the Internet. These are not internal DNS servers.

iii. Use sub-domains to help segmentation

For each sub-domain that has an A or MX record in DNS, there should be corresponding SPF records. Sites with wildcarded A or MX records should have corresponding wildcarded SPF records in the form:

```
* IN TXT "v=spf1 -all"
```

This strategy fits nicely with DKIM practices of creating separate key pairs and selectors for sub-domains used by third-party senders and special campaigns.

iv. Establish SPF records for non-sending domains

Companies may use only a handful of domains to send email, but they may also own several “brand” domains for online web presence or marketing that never send email. Preemptively establishing “null” SPF records for those domains will prevent fraudsters from flying under the radar by using them to send unauthorized spoofed email. For example, consider the following the SPF record:

```
"v=spf1 -all"
```

This record indicates that the domain does not send any email, and any email received will be fraudulent. Companies should also be sure to include third-party domain considerations as appropriate.

v. Includes are useful, but be careful when using them

The SPF specification allows for the inclusion of another domain’s SPF TXT as a way of chaining authoritative responsibility. While this can be a helpful administrative shortcut, it is important to ensure the domain that is being included is set up correctly.

For example, SPF lookups that cause more than ten DNS requests violate the SPF specification¹⁵. SPF adopters should ensure their SPF records, especially with the recursion that occurs with the use of “include” mechanism, do not require more than ten DNS requests as per the SPF RFC.

vi. Do not publish hard fail too soon

It is important to establish SPF and test using neutral records, and then analyze traffic (see [Section 6: Metrics and Reporting](#) for more information) to determine if there are any authorized hosts that are sending email that are not included in the SPF record. Publishing hard fail indicates that it is acceptable for a receiving MTA to take an aggressive approach to mail that fails an SPF check, including deleting the offending messages.

vii. Publish SPF records for HELO names

The HELO command is used as part of the SMTP process. SPF lookups are performed on the HELO domain name used in the HELO command by the Message Transfer Agent (MTA) sending the email, in addition to SPF lookups on the email sender’s domain name. The HELO domain names should be identified and SPF records published.

¹⁵ RFC 4408 <http://tools.ietf.org/html/rfc4408#section-10.1>

SPF syntax components are referred to as “mechanisms.” For a comprehensive description of the syntax, refer to OpenSPF.org’s mechanism discussion¹⁶. For each domain/sub-domain entity that will have an SPF record, all mechanisms of the syntax must be defined and tracked as the company implements SPF from testing phase, through neutral and soft-fail policy, and finally to hard-fail (if appropriate):

- Policy: “+” (Pass) | “-“ (Fail) | “~” (SoftFail) | “?” (Neutral)
- all: The mechanism that indicates set of email (all is default and only choice).
- ip: IP4 or IP6, depending on IP addressing numbering used for indicating host addresses. May be single addresses or address ranges.
- a: All the A records for domain are tested. If the client IP is found among them, this mechanism matches. If domain is not specified, the current-domains used. The A records have to match the client IP exactly, unless a prefix-length is provided, in which case each IP address returned by the A lookup will be expanded to its corresponding CIDR prefix, and the client IP will be sought within that subnet.
- mx: Denotes the MX record for specified domain should be used, also can be used in the form mx/24 where /24 denotes adjacent address range of potential sending hosts.

Other mechanisms that should be avoided, or will be used more rarely, include ptr:, exists:, include=, and modifiers such as redirect and exp=.

5.4 The Forwarding Issue

One of the more controversial aspects of SPF concerns forwarded mail. Some mail systems, most notably list servers, will sometimes change header information and indicate that they are sending mail on behalf of the original envelope sender. When this happens, the receiving MTA will perform the SPF check and fail the authentication on the technical grounds that the “new” sending MTA is not authoritative for the domain.

This will also be a common occurrence with those who set up forwarding or redirecting from a “parked” email account such as a university alumni email system. For example, a graduate of MIT will be allowed an alum.mit.edu email address. Many people will keep that email address and set their mail settings at MIT to simply forward any mail delivered to their current “live” email address. When a message is sent to them, it is sent first to MIT, which then re-delivers it (and takes responsibility for sending it) and causes an SPF check to fail at the receiving MTA.

Over time, system administrators will increasingly update their MTAs to switch from forwarding, where the envelope sender headers are preserved, to re-mailing. This switch requires changing the envelope sender to that of the forwarding system (which will publish its own SPF records).

In the meantime, ISPs should take into consideration the major systems that serve as forwarders and whitelist them. Additionally, senders (such as financial institutions) should carefully monitor test

¹⁶ SPF Record Syntax http://www.openspf.org/SPF_Record_Syntax

email corpi for messages that fail due to forwarding issues. See [Section 5.6](#) for a discussion on using SPF with DKIM to avoid messages being blocked due to forwarding issues. (Refer to the *Toolkit* for additional guidelines on the use of SPF.)

5.5 Considerations on Microsoft's SIDF

SenderID (SIDF)¹⁷ primarily differs from SPF in a few syntactical ways. It offers some minor flexibility in mechanism designation and is referenced as spf2.0 in the TXT nomenclature. The primary benefit, according to SenderID proponents, is the addition of the Purported Responsible Address (PRA) check.

The spf2.0/prs allows for policies checked only with the PRA identity. The PRA is derived via an evaluation of four mail header fields: From, Sender, Resent-From, and Resent-Sender. An algorithm (patented by Microsoft) derives a PRA address and allows for dealing with addresses in the header sender address field (the field used by MUAs to present From information). The stated benefit is that PRA can be derived via logic algorithm even through message forwarding, while an SPF check will stand a greater chance of passing on a legitimate email message.

Some message systems, including Microsoft's MSN and Hotmail, claim great success with using Sender ID and the PRA check to reduce false positives on fraudulent email. In practice, Microsoft is currently the only major ISP to support Sender. Each organization should review the pros and cons of SIDF and determine whether to support it. For all practical purposes, systems that support SIDF will also support SPF 1.0, so companies that do not publish spf2.0 records will not necessarily jeopardize the legitimacy of their email.

5.6 Considerations for the Joint Use of SPF and DKIM

Companies that decide to implement both SPF and DKIM will find many areas of overlap in the organizational preparation that must occur. The areas outlined in [Section 3: General Considerations for Deployment](#), including domain inventory, communication with all relevant groups, and dealing with third-party senders, apply to both DKIM and SPF. However, the technical implementation of each is unique and must be approached appropriately given that protocol's specific requirements.

i. Implementation tradeoffs

Each protocol provides its own set of benefits and drawbacks. With SPF, companies will have to manage the message forwarding issue and risk of false positives from mailing lists and other forwarders. With DKIM, companies will have to monitor for problems with MTAs, such as "leaking" messages (i.e., the MTA is not signing all its messages). In either case, there is the potential for false positives and possibly for false negatives.

ii. Combined DKIM/SPF use benefits

Companies that use both SPF and DKIM can mitigate many of the perils inherent in using either protocol alone. For example, messages from a leaking MTA that "breaks" the DKIM authentication process for those messages can still be authenticated via SPF, assuming the SPF record has been appropriately published. Likewise, a forwarded message that fails SPF authentication can still be appropriately authenticated using DKIM. The combined use of the two protocols reduces the number of false positives and can increase the receiving

¹⁷ SenderID <http://www.microsoft.com/mscorp/safety/technologies/senderid/default.aspx>

network’s confidence in authentication, to the point of being willing to start blocking messages that fail both authentication processes.

Note that for these benefits to accrue, the receiving network must approve messages that pass either DKIM or SPF (or both), and only fail a message if it fails both authentication processes. The experiences of companies and ISPs who have begun to implement both protocols suggest that this strategy is likely to result in the least number of errors. This approach also enables companies to work more flexibly with third-party senders, who need implement only one of the two protocols to be able to send messages (although third-party senders should ultimately support both protocols). In addition, this strategy enables companies to have their messages authenticated at ISPs that only support one of the two protocols, as ISP support of both protocols is not currently available in many cases. ISPs today have not yet implemented blocking/failure based on authentication results. Some improved policy capabilities may be required at ISPs to support the most beneficial dual-failure model for companies.

While the recommendation of this paper is for companies to implement both protocols (DKIM and SPF), it is recognized that some companies may have limited resources for this initiative and must decide what protocol to implement first. In that case, the company should take into account several tradeoffs for the implementation of each in its decision. In some cases, implementing DKIM first makes sense as it avoids a number of false positive risks associated with SPF and thus ensures the highest rate of accuracy for authentication. However, it should be noted that DKIM is operationally more complex to implement than SPF and thus the company’s technical expertise and available resources should be taken into consideration. DKIM requires a much broader range of testing than SPF, since SPF only requires DNS testing, and DKIM can have a high false positive rate if messages are not being signed correctly. Please see a summary of these issues in the following table.

In addition, companies considering implementation should review the current relative levels of support for DKIM and SPF by the ISPs that are most represented in their customer base.

Tradeoffs of SPF and DKIM Deployment

	Advantages	Disadvantages
DKIM Only	<ul style="list-style-type: none"> • Supported by major ISPs (excluding Microsoft) • Signature stays with message over multiple hops • Validates message content as well as path 	<ul style="list-style-type: none"> • Operationally complex to implement and test • False positive risk due to message signing errors or leakage
SPF Only	<ul style="list-style-type: none"> • Only requires DNS-based testing; no additional software required 	<ul style="list-style-type: none"> • Increased false positive risk due to path-based authentication breaking on message forwarding
Combined DKIM & SPF	<ul style="list-style-type: none"> • All advantages of <i>DKIM Only</i> implementation • Minimizes false positive rate (assuming receiver network requires only one protocol to pass) 	<ul style="list-style-type: none"> • Requires additional resources to implement both protocols

6. Metrics and Reporting

For effective deployment and use of authentication, it is important to gather data on authentication results and email disposition outcomes. Today, however, there are fundamental limitations to the data available, so it is important to leverage those capabilities that are possible.

6.1 Signing Validity

Companies may use test accounts to determine the validity of SPF and DKIM signing. End-user test accounts can be set up across multiple ISP systems to establish a view into how those ISPs are presenting the messages.

Evaluating outcomes

- It is generally recommended that at least two test accounts are set up at each ISP.
- In most cases, the validity of the DKIM check can be determined by examination of the header of the message (this is difficult to automate because each ISP uses a different designation).
- In some cases, it is possible to infer the overall disposition of a message based on the heuristics scoring for the ISPs email hygiene process (approximation based on reputation).
- Trends can be created and analyzed over time to approximate the efficacy of authentication for messages delivered to the ISP that is hosting the test account.

Drawbacks

- Using test accounts will not result in complete coverage of the company's sending universe as it will not be possible to cover every possible receiver; corporate accounts, in particular, will not allow test accounts.
- Every ISP system is different and can change without notice. Since this is a manual procedure, it requires a great deal of administration to maintain test accounts across multiple ISPs with email messages going in from multiple systems and third parties.
- ISPs may drop messages for reasons unrelated to the authentication protocols, so it is difficult to know what happened to messages that do not appear in the test accounts. Where possible, it is recommended that the receiving networks whitelist the senders in order to avoid these issues to minimize false positive risk.
- It is also administratively difficult to have third-party senders routinely send email into test accounts for analysis.
- There is no way to detect non-participating third-party or internal systems.
- Without special relationships with the ISP or with third-party data providers, it can be difficult to know what happened to messages at the ISP.

Overall, this procedure is recommended today as the best available approach, but it is manually labor intensive, difficult to automate, and frequently subject to change.

6.2 Gathering Spoofing Data

Data on the origins of spoof traffic is contained in the SMTP interactions and email messages from senders. Today there are a several post-facto methods of gathering information:

- Analysis of customer feedback: Some customers who receive suspicious emails forward them to relevant support addresses at institutions or ISPs. Companies may wish to create an

abuse mailbox (for example, *phish@examplebank.com*) for the receipt of emails. These emails can be processed for data on phishing sites, spoofing domains and IPs, and some information on internal systems that are not sending email correctly. This method is disadvantageous as a primary source of data since allows spoof traffic to reach end-recipients prior to detection, and relies on end-recipients' ability to identify suspicious traffic and their reliability in reporting it accurately.

- Honey pots: Honey pots can be established to draw phishing and spoofing traffic for detection.
- Third-party data phishing services: A number of third-party services collate data from various sources to detect phishing traffic across the Internet.

All of these methods detect spoofing after it has occurred at some destinations, but it is recommended that institutions support all of them.

6.3 Desirable Future Data

For accurate, timely, granular analysis of authentication and delivery, the following additional data from receivers is desirable.

- Number of messages that arrived at each ISP/receiver for specified domains
- Number of messages that are unauthorized spoof for specified domains
- Origin of unauthorized spoof messages
- Breakdown of messages from authorized sources and their origins
- Percentage representation of the disposition of messages in each group for each specified domain (what happened to the messages)
- For messages that fail authentication, granular results (why were they classified that way)
- Consistent and normalized standards and statistical presentation across ISPs/receivers

Glossary

Accreditation: Qualification of an organization, business unit, or domain as eligible to certify messages. Such qualification may include verification of domain ownership, sending practices and security practices. Several third-party services provide accreditation services to email.

Authentication: Authentication protocols such as DKIM and SPF provide a mechanism to determine whether a message is from the sender that the message purports to be from, based on information in its header or origin path.

Canonicalization: Generally speaking, a process for converting data that has more than one possible representation into a “standard” canonical representation. This can be done to compare different representations for equivalence, to count the number of distinct data structures, to improve the efficiency of various algorithms by eliminating repeated calculations, or to make it possible to impose a meaningful sorting order. In DKIM, canonicalization is used to handle potential non-malicious modifications of email in transit by intermediary mail servers and relay systems. DKIM supports two canonicalization algorithms, *relaxed* and *simple*, chosen by the sender to support their requirements. The *relaxed* algorithm can tolerate some modifications to the message in transit and still pass authentication, handling most minor, non-malicious changes to the message. The *simple* algorithm tolerates almost no changes to the message for highly change-sensitive deployments. Each organization will need to determine which algorithm works best in their environment; in practice, *relaxed* is more common due to the intolerance of *simple*.

Certification: A process that verifies and attests to the safety and validity of an email message, typically using one or more authentication protocols like DKIM or SPF.

Delegation: Also known as “authorized spoofing,” delegation occurs when a company authorizes the third party to use its domain name to send email (e.g., outsourced marketing or transactional email sent by email service providers).

DomainKeys Identified Mail (DKIM): An email authentication protocol that enables the sender to use public-key cryptography to sign outgoing emails in a manner that can be verified by the receiver. The DKIM specification is based on the prior protocols Domain Keys and Identified Internet Mail. DKIM is defined in RFC 4871.

Envelope Sender Address: Sometimes called the “Return-Path.” It is used during SMTP transport of a message (during the initial SMTP handshake conversation) to delineate the address to which a message should be returned in the case of a delivery failure. This address is not usually displayed by the MUA to the recipient.

False Positive: In this context, a *false positive* refers to a legitimate message (one that is from the sender that it purports to be from and unadulterated in transit) that is marked as spam or phish by an ISP.

Header Sender Address: Sometimes called the “Sender,” “From Name” or “Friendly From” address, is generally displayed by the MUA to the recipient, and is usually an “easier” and more recognizable designation to read by the message recipient.

Message Transfer Agent (MTA): Any system running SMTP routing software that can take a message, process it, look up destination information in DNS (or other routing table), and deliver to the intended receiving system. MTAs are typically server applications such as Sendmail, Microsoft Exchange, Postfix, Lotus Domino, etc.

Message User Agent (MUA): Typically refers to any computer application that provides an interface to a person who will read or send email. These may be “fat” clients such as Outlook and Lotus Notes, or “thin” clients such as web applications, Yahoo! Mail, AOL, Gmail, etc.

Phishing: Unauthorized spoofing performed with the purpose to trick the recipient into disclosing personal information, visiting a malicious website or downloading malware to their computer so criminals may obtain access to the victim’s financial accounts.

Reputation Services: Services that rank or score the risk that a sender is sending valid, permission-based email that is desired by receivers, based on some combination of past performance of a sender, such as complaints, spam designation, bounce rates and other delivery outcomes; characteristics of domains or IP addresses; or analysis of content or other factors. Companies that provide these services include Trend Micro and IronPort.

Selector: A mechanism used by the DKIM protocol to support multiple concurrent public keys for a given signing domain. A selector can indicate a date, place, or other identifying information, and can contain periods; examples include `march09.marketing.subsidiarybank` and `february2008`. The recommended format of selectors is `date.purpose.admin`. More information is available in the DKIM RFC 4871.

Sender Policy Framework (SPF): SPF is a path-based email authentication protocol that allows email receivers to determine if the sender is authorized to use the domains in the message’s header by evaluating the IP address of the sender’s outbound MTA based on information published by the sender in DNS TXT records. SPF is defined in RFC 4408.

Simple Mail Transport/Transfer Protocol (SMTP): The standard Internet protocol for sending email messages from sender to recipient. SMTP is defined in RFC 5321.

Spoofing: Occurs when a third party sends email from a domain as if it is the domain owner (when, in fact, it is not). *Authorized spoofing*, or *delegation*, occurs when a company authorizes the third party to use its domain name to send email (e.g., outsourced marketing or transactional email sent by email service providers). *Unauthorized spoofing* occurs when a third party represents email from a domain that it does not own and is not authorized to send on behalf of. For the purposes of this document, we use the term *spoofing* to represent *unauthorized spoofing* or *phishing* and not the legitimate use of authorized third-party senders.

Transport Layer Security (TLS): TLS, based on Secure Sockets Layer (SSL) technology, establishes an encrypted tunnel through which other protocols such as SMTP may pass data. TLS is defined in RFC 5246.

Appendix A: Resources

Authentication Guidelines for Financial Institutions

- BITS Email Security Toolkit: Protocols and Recommendations for Reducing the Risks (April, 2007):
<http://www.bits.org/downloads/Publications%20Page/BITSSecureEmailFINALAPRIL1507.pdf>

DomainKeys Identified Mail (DKIM) Information

- DKIM Overview: <http://tools.ietf.org/html/draft-ietf-dkim-overview-10>
- DKIM Signatures (RFC 4871): <http://www.apps.ietf.org/rfc/rfc4871.html>
- ADSP Draft Memo (Expired): <http://www.dkim.org/specs/draft-ietf-dkim-ssp-04.html>

Reflectors (Verifiers)

- DKIM.org: <http://testing.dkim.org/reflector.html>
- Sendmail: Send an email to sa-test@sendmail.net (supports DK, DKIM, SDF and SPF)
- Port25: <http://www.port25.com/domainkeys/>

Sendmail

- Sendmail and DKIM: <http://www.sendmail.org/dkim>
- Eland's Sendmail / DKIM Overview:
<http://www.elandsys.com/resources/sendmail/dkim.html>
- Erik Berg's Sendmail/DKIM/SDF Overview:
<http://www.erikberg.com/notes/milters.html>

Postfix

- Postfix MILTER support: http://www.postfix.org/MILTER_README.html
- DKIM with Postfix using *dkimproxy*:
<http://anothersysadmin.wordpress.com/2008/01/16/domainkeysdkim-with-postfix/>

Microsoft Exchange

- DKIM for IIS SMTP Service and Exchange Server:
<http://www.emailarchitect.net/domainkeys/>
- Exchange Server and SenderID: <http://technet.microsoft.com/en-us/magazine/2006.12.sidf.aspx>

Sender Policy Framework (SPF) Information

- SPF (RFC 4408): <http://www.ietf.org/rfc/rfc4408.txt>
- OpenSPF.org Syntax Discussion: http://www.openspf.org/SPF_Record_Syntax
- SPF Common Mistakes: http://www.openspf.org/FAQ/Common_mistakes

MTA Vendor Listings

- Online Trust Alliance (OTA) directory:
<https://otalliance.org/resources/2009%20OTA%20Member%20Directory.pdf>
- DKIM.org's vendor listing: <http://www.dkim.org/deploy/>

Industry Groups

- BITS (a division of the Financial Services Roundtable): <http://www.bits.org>
- The Financial Services Roundtable: <http://www.fsround.org>
- Messaging Anti-Abuse Working Group (MAAWG): <http://www.maawg.org>
- Anti-Phishing Working Group (APWG): <http://www.antiphishing.org>

Appendix B: Sample Letter Of Intent

Sample letter of intent for business partners and outsourced vendors

Re: SENDER AUTHENTICATION – OUTBOUND EMAIL VALIDATION
IMPORTANT INFORMATION FOR BANK BUSINESS CONTACTS AND THEIR
OUTSOURCED VENDORS

Overview

{Company} will implement its Email Sender Authentication program as part of its continuing effort to protect associates, customers, and the company from “phishing” fraud.

What is phishing?

“Phishing” refers to email sent by unauthorized external entities who copy the “look and feel” of company email messages so they appear to come from legitimate and authorized {Company} sources.

These emails put associates and customers at risk by asking them to submit confidential personal information, either via reply email or by sending them to a fake web site that looks legitimate, to steal their identity or install malicious software.

What are we doing?

{Company} is implementing Internet standard email authentication practices, which have been adopted throughout the industry, as well as by leading Internet Service Providers such as AOL, Yahoo, Google, and Comcast. Two technologies, Domain Keys Identified Mail (DKIM) and Sender Policy Framework (SPF) permit ISPs to validate the authenticity of email through encrypted digital signatures (DKIM) and path-based analysis (SPF).

By using DKIM and SPF, ISPs are able to identify what email is coming from {Company} and is legitimate, versus email that is fraudulent and falsely appears to be originating from {Company}. This allows ISPs to differentiate between {Company’s} validly signed email and all other signed and/or unsigned email. In this way, email falsely claiming to be from {Company} may be rejected by the ISPs rather than sent along to our customers.

What we need you to do and when

- The Sender Authentication project team is collecting information to ensure all internal lines of business and external business partners/vendors who are authorized to send email to customers using one of the {Company’s} addresses (domains) are registered and inventoried in order for their email messages to be properly authenticated and not rejected by ISPs.
- During {date range}, a pilot will be conducted to ensure this phase of the Sender Authentication program is working as expected, and only messages with invalid signatures are rejected and prevented from reaching intended victims.
- Please advise your line of business contacts of the following:
 - Phase {x} or Sender Authentication program will be piloted in {date range} and planned implementation is {milestone date}

Email Sender Authentication Deployment

- The importance of immediately notifying the Sender Authentication team at {email address} of all internal and external entities who are legitimately sending email to customers on the {Company's} behalf using one of our company-owned domains (to ensure no interruption in business operations). We need this information no later than {deadline}.
- Send an email to {email address} with the contact information of the external supplier/vendor who is sending email on behalf of the Bank. Technical personnel will contact the supplier, who should be ready to speak to their ability to support DKIM and SPF for their outbound email campaigns.
- A contingency plan will be in place in the event legitimate messages do not reach intended customers.
- Questions should be emailed to {email address}.

Suppliers providing outsourced email campaigns for {Company}...

- Should be familiar with DKIM technology.
- DKIM deployment should be finished by {milestone date}.
- Minimum key length is 1024-bit. Key rotation policy will be provided {milestone date}.
- Recommended canonicalization is relaxed/simple.
- Recommended signing algorithm is RSA-SHA256.
- Should only be sending email from a company sub-domain address – any exceptions to this needs to be identified; please send an email to {email address}.
- Deliver the DKIM **public** key to {email address}, which will be published in the {Company's} DNS for lookup by receiving systems. The preferred format for delivery of the key is base64 PEM file.
- Follow testing process:
 - Send signed test messages to reflectors listed at <http://testing.dkim.org/reflector.html>. (NOTE: test message content should be representative of normal mail, i.e., a sample or test message from your application or campaign, NOT a generic test message)
 - Send messages returned from reflectors as either message/RFC822 attachments OR as individual email files (archived zip, tar, etc) to {email address}
- All suppliers will be contacted by {program team name} to discuss deployment following receipt of these instructions.

Additional information for {Company} associates

- Refer to the {project document repository location} for further information on the Sender Authentication program, communications and timelines
- Contact your line of business Service Delivery Managers or Information Security Officer.

Appendix C: Sample Implementation Project Plan

Please see below for an example implementation plan that a company might reference when establishing its own project plan for deploying SPF and DKIM. It is important to note that this project plan does not include post-implementation tasks, including on-going deployment management, evaluating help desk feedback, and forming ISP relationships for policy application and message handling data access.

In addition, this plan lists DKIM before SPF in the implementation process but, in practice, either can be implemented first (or in parallel, depending on the company's available resources). Either way, there is a significant amount of overlapping pre-work to be completed in preparation for either DKIM or SPF deployment.

ID	Task Name	Predecessors	Duration (Days)	Start	Finish
1	Sender Authentication Project Plan		404	10/31/2008	5/19/2010
2	Pre-Work		45	10/31/2008	1/1/2009
3	Open project		1	10/31/2008	10/31/2008
4	Create business requirements documentation		1	11/25/2008	11/25/2008
5	Create Project Charter		7	11/20/2008	11/28/2008
6	Create Rigor Worksheet		1	11/24/2008	11/24/2008
7	Begin consultant vendor selection		1	1/1/2009	1/1/2009
8	Inventory Inbound / Outbound Senders	2	288	1/2/2009	2/9/2010
9	Develop process to document outbound sender domains		288	1/2/2009	2/9/2010
10	Inventory maintenance and support process		30	1/2/2009	2/12/2009
11	Input sender domains	10	3	2/13/2009	2/17/2009
12	Field names		3	2/13/2009	2/17/2009
13	Team vetting		3	2/13/2009	2/17/2009
14	Collaboration database / site for capturing data		3	2/13/2009	2/17/2009
15	Capture Domain information	10	78	2/13/2009	6/2/2009
16	Verify Vendor / sender domains - testing	15	180	6/3/2009	2/9/2010
17	Automate the process for capturing the sender domains		6	1/2/2009	1/9/2009
18	Generate Process document		5	1/27/2009	2/2/2009
19	Test inventory site / method		4	1/2/2009	1/7/2009
20	Document Exception Process (authorized spoofing)		8	1/2/2009	1/13/2009
21	Policy Documentation	2	30	1/2/2009	2/12/2009
22	Policy and project communications		15	1/2/2009	1/22/2009
23	Technology infrastructure		15	1/2/2009	1/22/2009
24	Information security		15	1/2/2009	1/22/2009
25	E-Commerce		15	1/2/2009	1/22/2009
26	Supply chain / procurement		15	1/2/2009	1/22/2009
27	Operation risk and compliance		15	1/2/2009	1/22/2009
28	Marketing		15	1/2/2009	1/22/2009
29	Corp communications / PR		15	1/2/2009	1/22/2009
30	Email "from" policy		30	1/2/2009	2/12/2009
31	Email content policy		30	1/2/2009	2/12/2009
32	Domain provisioning policy and process		30	1/2/2009	2/12/2009
33	Authentication Installation	8	71	2/10/2010	5/19/2010
34	Communications		43	2/10/2010	4/9/2010
35	Identify Internal Key Associates		23	2/10/2010	3/12/2010
36	Draft communication to Key Associates / Leadership	35	12	3/15/2010	3/30/2010
37	Draft communication to Outsourced vendors	36	7	3/31/2010	4/8/2010
38	Create Internal Standards Organization & Security Group Presentations		10	2/10/2010	2/23/2010
39	Set dates for Standards & Security presentations	38	5	2/24/2010	3/2/2010
40	Send Vendor Communication	37	1	4/9/2010	4/9/2010
41	Send "Bulletin" notice to Standards and Service Delivery Managers / Business Liaison	39	1	3/3/2010	3/3/2010
42	Make presentation to Standards Organization	39	1	3/3/2010	3/3/2010
43	Make presentation to Security Group	39	1	3/3/2010	3/3/2010

Email Sender Authentication Deployment

44	Debrief on E-Commerce LoB meeting	43	1	3/4/2010	3/4/2010
45	Debrief on Standards and Security presentations	43	1	3/4/2010	3/4/2010
46	Support Documentation		10	2/10/2010	2/23/2010
47	Help Desk Documentation		13	2/10/2010	2/26/2010
48	Backout Plan		19	2/10/2010	3/8/2010
49	Software		9	2/10/2010	2/22/2010
50	Exceptions		19	2/10/2010	3/8/2010
51	Implementation Plan Documentation		26	2/10/2010	3/17/2010
52	Software		1	2/10/2010	2/10/2010
53	Exceptions		26	2/10/2010	3/17/2010
54	DKIM Implementation Plan	34,48,51	28	4/12/2010	5/19/2010
55	Generate Keys		5	4/12/2010	4/16/2010
56	Preview DKIM binary delivered		3	4/12/2010	4/14/2010
57	Develop Test Plan		14	4/12/2010	4/29/2010
58	Push keys into DNS - lab environment	57	2	4/30/2010	5/3/2010
59	Push keys into DNS - test environment	57	2	4/30/2010	5/3/2010
60	Push keys into DNS - stage environment	57	2	4/30/2010	5/3/2010
61	Production DKIM binary or configuration delivered	60	1	5/4/2010	5/4/2010
62	Vendor onsite visit and documentation delivery		1	4/12/2010	4/12/2010
63	DKIM binary lab installation	61,62	3	5/5/2010	5/7/2010
64	Lab environment installation tested	55	3	4/19/2010	4/21/2010
65	DKIM binary stage installation	60,64	1	5/4/2010	5/4/2010
66	Stage environment tested	65	7	5/5/2010	5/13/2010
67	Submit Change Control Record	66	1	5/14/2010	5/14/2010
68	Push keys into DNS - production environment (selected domains / selectors)	67	1	5/17/2010	5/17/2010
69	DKIM binary production installation	66,67,68	1	5/18/2010	5/18/2010
70	Production testing and initiate ongoing test scripts	69	1	5/19/2010	5/19/2010
71	SPF Implementation Plan	34,48,51	6	4/12/2010	4/19/2010
72	SPF record rationale / impact statement		1	4/12/2010	4/12/2010
73	Generate SPF domain policy and syntax conventions		1	4/12/2010	4/12/2010
74	Develop test plan		1	4/12/2010	4/12/2010
75	Publish SPF neutral records for target domains	74	1	4/13/2010	4/13/2010
76	Test neutral record validation	75	1	4/14/2010	4/14/2010
77	"Cooling off" period for neutral records - Internet observation	76	1	4/15/2010	4/15/2010
78	Publish SPF soft fail records for target domains	77	1	4/16/2010	4/16/2010
79	Test soft fail record validation	78	1	4/19/2010	4/19/2010

Appendix D: Creating a Private-Public Key Pair using OpenSSL

Below is a detailed example of the steps to create a public/private key pair for use in DKIM implementation using OpenSSL, an open source toolkit implementing SSL, TLS, and a full-strength general purpose cryptography library. Please see [Section 4.4](#) for more information about DKIM implementation.

- **Key pair generation**

Keys can be created using OpenSSL or by a certificate vendor.

Example commands for creating a key pair using OpenSSL are:

```
opensslgenrsa -out dkimselector.private 1024
opensslrsa -in dkimselector.private -out dkimselector.public -
pubout -outform PEM
```

The `opensslgenrsa` command generates the RSA private key. In this example, two switches are designated for `opensslgenrsa`. The first, “`-out dkimselector.private`,” specifies that `dkimselector.private` is the output filename. The second, “`1024`,” specifies that the size (in number of bits) of the private key is 1024 (the default is 512).

The `opensslrsa` command processes the RSA keys; they can be converted to various forms, and their components can be printed out. Four switches are specified in this example. The first, `-in dkimselector.private`, specifies the name of the input file to read the key from (which in this example was generated by the previous command). The second, `-out dkimselector.public`, specifies the output filename to write the key to. Thirdly, `-pubout` determines that a public key will be generated (by default, this command generates a private key). Finally, `-outform PEM` specifies the output format to be of type “PEM”: an ASN1 DER encoded form base64 encoded with additional header and footer lines.

See OpenSSL documentation (<http://www.openssl.org>) for complete syntax on creating RSA PKI keys. Companies may wish to use other switches depending on their needs.

One file, `dkimselector.private`, will contain the public/private key pair. This file will be used by the software responsible for signing messages. The other file, `dkimselector.public`, will contain the public portion of the key pair.

- **Public key installation**

The text from the public key file is copied into a DNS record for the domain the key has been created for. This DNS record will be used by ISPs to perform DKIM authentication on messages purporting to be from the specified domain.

For example, such a record will look like this:

```
selector    IN
TXT         "v=DKIM1; k=rsa; t=s"
           "p=keytext"
```

where *v* is the version of the DKIM key record; *k* is the key type; *t* specifies any flags, and *s* in this case refers to sub-domain policy; *p* is the public key data, and *keytext* is the public key text (which looks like a string of random numbers, letters and symbols).

Please see the DKIM RFC 4871 for complete information on the DKIM DNS records (<http://www.apps.ietf.org/rfc/rfc4871.html>).



eCert's Contribution

BITS wishes to acknowledge the significant contribution eCert made in the development of this document. eCert created the entire first working draft of the document, which was then subsequently reviewed and edited by various BITS members (as acknowledged below). Without eCert's efforts, development of this paper would have been a much slower and arduous process, which would have delayed publication. The content of the first draft itself and its assistance during editing of the paper were a reflection of eCert's knowledge and practical expertise gained through its work with both the financial services and ISP communities in a number of areas including:

- Accreditation of sending domains against standard implementation procedures and policies to ensure enforcement capabilities and policy accuracy
- Accreditation of ISPs to ensure enforcement capabilities
- Creation of standardized authentication outcomes and policies to unify DKIM and SPF/SIDF deployment across multiple ISPs and other email receivers
- Establishment of relationships with ISPs and other email receivers that involve legal contract management, service level structures and management, creation of technical requirements and creation of roadmaps in the areas of security, delivery and analytics
- Establishment of methodologies that allow for granular implementation of email security protocols and for managing security and operational issues both during and subsequent to implementation
- Formation of policy methodologies that allow both the joint deployment of SPF/SIDF and DKIM and the ability to interpret and enforce combined deployment
- Creation of standardized methodologies for collecting, aggregating and analyzing message handling data from IPSs and other receivers to allow for trend and comparative performance analysis
- Definition and execution of techniques to collect and analyze domain and IP address data regarding unauthorized spoof traffic
- Establishment of a collaborative roadmap process to advance capabilities for security and mail handling between senders and receivers

For more information, contact:

eCert

One Market Street, Suite 3500

San Francisco, California 94105

Phone: +1 (415) 681-8000

Email: information@ecertsystems.com

Acknowledgements

We would like to thank the following individuals for their contributions to this document.

Mark Brumbaugh, Prudential Financial Inc.
Jeff Carnahan, US Bancorp
Erik Johnson, Bank of America Corporation
Steven Jones, Bank of America Corporation
Joseph McGrath, Fidelity Investments
William Parra, Bank of America Corporaton
Russell Pierce, Wells Fargo & Company
Alex Popowycz, Fidelity Investments
Jim Schmitt, Edward Jones
Victor Talamo, JPMorgan Chase & Co.

We would also like to thank the following individuals for coordinating this effort.

Andrew Kennedy, BITS
Paul Smocer, BITS
Eve Phillips, eCert
Kelly Wanser, eCert