Communications Security Establishment
Centre de la sécurité des télécommunications

CANADIAN CENTRE FOR CYBER SECURITY

Implementation Guidance:

Email Domain Protection

PRACTITIONER

ITSP.40.065

Canada

FORWARD

ITSP.40.065 Implementation Guidance: Email Domain Protection is an UNCLASSIFIED publication that is issued under the authority of the Head of the Canadian Centre for Cyber Security (Cyber Centre). For more information, email or phone our contact centre:

Contact Centre
contact@cyber.gc.ca
(613) 949-7048 or 1-833-CYBER-88

EFFECTIVE DATE

This publication takes effect on April 7, 2020.

REVISION HISTORY

Revision	Amendments	Date
1	First release	April 7, 2020

TABLE OF CONTENTS

1 Overview 6

2 Email Domain Protection Mechanisms 7

2.1 SPF 7

2.2 DKIM 7

2.3 Limitations of SPF and DKIM 8

2.4 DMARC 8

2.4.1 DMARC Validation 9

2.4.2 DMARC Reporting 11

3 Additional Considerations 12

3.1 Vendor Support 12

3.2 Third-Party Senders 12

3.2.1 Third Parties and SPF 12

3.2.2 Third Parties and DKIM 12

3.2.3 Third Parties and DMARC 12

3.2.4 Subdomain Separation 13

3.3 Message Forwarding 13

3.4 Inbound Mail 13

3.5 Related Standards 14

3.5.1 Authenticated Received Chain (ARC) 14

3.5.2 Brand Indicators for Message Identification (BIMI) 14

3.5.3 Email Transport Encryption 14

3.5.4 No Service MX Record 16

4 Cyber Centre DMARC Reporting Service 17

5 Summary 18

5.1 Contact Information 18

6 Supporting Content 19

6.1 List of Abbreviations 19

6.2 References 20

LIST OF FIGURES

Figure 1: DMARC Validation 10

LIST OF ANNEXES

Annex A Implementation Plan 21

A.1 Overview 21

A.2 Assess 22

A.2.1 Identify Mail Domains 22

A.2.2 Assess Current State 22

A.2.3 Deploy Initial DMARC record 22

A.2.4 Collect and Analyze DMARC Reports 23

A.3 Deploy 24

A.3.1 Identify Authorized Senders 24

A.3.2 Configure DNS Time To Live (TTL) 24

A.3.3 Deploy SPF for All Domains 24

A.3.4 Deploy DKIM for All Domains and Senders 24

A.3.5 Monitor DMARC Reports and Correct Misconfigurations 25

A.4 Enforce 26

A.4.1 Gradually Enforce Quarantine 26

A.4.2 Gradually Enforce Reject 26

A.4.3 Reject All Messages To and From Non-Mail Domains 26

A.5 Maintain 28

A.5.1 Monitor DMARC Reports 28

A.5.2 Correct Misconfigurations and Update Records 28

A.5.3 Rotate DKIM Keys 28

Annex B Protocol Reference 29

B.1 SPF 29

B.1.1 SPF Records 29

B.1.2 Third-Party Senders 30
B.1.3 DNS Lookup Limit 30
B.2 DKIM 31
B.2.1 DKIM Records 31
B.2.2 Cryptographic Considerations 31
B.2.3 Third-Party Senders 31
B.3 DMARC 32
B.3.1 DMARC Records 32

1 OVERVIEW

This document provides guidance to system owners on implementing technical security measures to protect their domains from email spoofing. In this document, we describe technical measures that system owners can implement to prevent the delivery of certain malicious messages sent on behalf of their domains and identify the infrastructure used by malicious actors.

By implementing this guidance, you can prevent threat actors from representing themselves as your organization by using your email domain. The guidance in this document also helps prevent phishing emails from being delivered to your organization.

Phishing is a widely used attack method in which threat actors send emails (or use other communication methods such as SMS or phone calls) that appear to come from a trusted domain (e.g. canada.ca) and that seem to be legitimate. Threat actors use phishing attacks to attempt to trick recipients into disclosing personal data and other sensitive information or as a method for installing malicious software (i.e. malware) on devices.

You reduce a threat actor's chance of carrying out successful malicious email campaigns by implementing the technical security measures detailed in this document. More specifically, these technical security measures protect your organization in the following ways:

- Preventing the delivery of malicious messages impersonating your domain;
- Disrupting the infrastructure used to send these malicious messages;
- Deterring threat actors from attempting to spoof these protected domains in the future;
- Improving the security of email recipients; and
- Protecting the reputation of organizations whose brands are the target of spoofing.

2 EMAIL DOMAIN PROTECTION MECHANISMS

Three security protocols act jointly to protect email domains from being spoofed:

- Sender Policy Framework (SPF);
- DomainKeys Identified Mail (DKIM); and
- Domain-based Message Authentication, Reporting, and Conformance (DMARC).

For complete protection, you must implement all three protocols and configure them to instruct recipients to reject inauthentic messages. These protocols are described in the following sections. For guidance on implementing them, see Annex A.

2.1 SPF

You can use SPF to specify the Internet protocol (IP) addresses from which emails can be sent on a domain's behalf. The SPF standard is formally defined by the Internet Engineering Task Force's (IETF) *Request for Comments (RFC) 7208: Sender Policy Framework (SPF) for Authorizing Use of Domains in Email* [1][1].

You can implement SPF by publishing a domain name system (DNS) record for each protected domain and subdomain that lists the authorized IP addresses, directly or by reference to other records.

When a message is received, an email system that supports SPF performs the following actions:

1. Retrieves the SPF record that is associated with the sending domain; and
2. Verifies that the IP address used to send the message has been authorized to do so.

Messages from unauthorized IP addresses may be accepted, marked as suspicious, or rejected; however, these actions depend on the policy that is included in the SPF record.

2.2 DKIM

DKIM provides a mechanism for email messages to be authenticated using a cryptographic signature. The DKIM standard is formally defined by IETF's *RFC 6376: DomainKeys Identified Mail (DKIM) Signatures* [2].

You can implement DKIM by:

1. Publishing one or more DNS records for the protected domain, each containing a cryptographic public key and additional information;
2. Deploying private keys to a domain's message transfer agents (MTAs); and
3. Configuring MTAs to sign outgoing messages.

[1] Numbers in square brackets refer to references cited in the Supporting Content section of this document.

You must configure each MTA that sends emails on behalf of a protected domain with a private key that corresponds to a published DKIM record. When the MTA sends a message, it uses the private key to add a cryptographic signature to the message by appending a message header. If boilerplate content is added to outbound messages, e.g. a legal disclaimer, this must be done before the DKIM signature is applied. Otherwise, the signature will be invalidated.

When an email system that supports DKIM receives a DKIM-signed message, it retrieves the record associated with the message's DKIM header and verifies the message's signature using the published public key. This DKIM check cryptographically confirms that the message was sent by an authorized sender and was not altered in transit. If the signature is not valid, or if no DKIM record is available, the message will fail DKIM. Messages that fail this DKIM check may be rejected.

2.3 LIMITATIONS OF SPF AND DKIM

SPF and DKIM share a limitation that makes them ineffective against moderately sophisticated threat actors. Both protocols rely on domain names that are hidden from the user and that can be different than the one displayed to the user in an email's From field (known as the *header from* address). SPF uses the domain of the *envelope from* email address (also known as the Return-Path address), which is used by email servers in the background. DKIM uses a domain that is specified in the DKIM header.

There is no requirement for the domains that are used in the *envelope from*, *header from*, or DKIM header to match. Therefore, a threat actor can increase the likelihood that malicious emails are delivered by implementing SPF and DKIM for a domain under their control, while presenting a different and more trustworthy domain in the message's From field to trick the recipient.

You should note that legitimate messages may be rejected by receiving systems if SPF or DKM records are missing or improperly configured. As these protocols do not include a reporting mechanism, you may not be aware of failed deliveries. Similarly, there is no mechanism that enables receiving systems to inform domain owners about messages that have been flagged as illegitimate due to a failed SPF or DKIM check.

2.4 DMARC

DMARC was created to address the limitations of SPF and DKIM and improve enforcement and reporting. The DMARC standard is formally defined by IETF's *RFC 7489: Domain-based Message Authentication, Reporting, and Conformance (DMARC)* [3].

DMARC introduces several improvements, including the following:

- Enforces alignment between the *envelope from* and *header from* domains for SPF;
- Enforces alignment between the DKIM header domain and the *header from* domain for DKIM;
- Allows system owners to specify the action that a receiving system should take if a message fails both the SPF and DKIM checks; and
- Provides a mechanism for receiving servers to report information on DMARC results to domain owners.

You can implement DMARC by publishing a DNS record for the protected domain that provides policy information to receiving systems.

2.4.1 DMARC VALIDATION

For an email to pass DMARC validation, it must pass either the SPF or the DKIM validation, and the domain used in that validation must align with the one in the email's *header from* field. If an email fails both the SPF and DKIM validation and alignment, it also fails the DMARC validation.

When testing alignment with SPF and DKIM, DMARC uses the domain of the *header from* email address. The receiving system first looks for a DMARC record that corresponds to the full domain. If no record is found, the system falls back to successively higher-level domains, up to the root domain. According to the Public Suffix List[2], the root domain is determined as the highest-level domain that is available for registration.

If an email passes DMARC validation, it is delivered. If, however, an email fails DMARC validation, the receiving email system applies the policy that is specified in the sending domain's DMARC record. The policy must be one of the following options:

- **None**: The email is delivered (i.e. monitor only);
- **Quarantine**: The email is delivered but is marked as suspicious; or
- **Reject**: The email is rejected.

A DMARC record can also specify that the published policy be applied to only a percentage of messages that fail validation, with the next less strict policy applied to the remainder. For example, a policy of **reject** with a percentage of 50% will reject 50% of messages which fail, while 50% will be quarantined.

While policies of **none** and **quarantine** can help you gather information and configure your systems, only a policy of **reject** will prevent illegitimate messages from being delivered.

[2] Maintained by the Mozilla Foundation, the Public Suffix List attempts to capture all domain suffixes below which a domain can be registered. For example, Canadian top-level domains on the list include .ca, .mb.ca, and .gc.ca.

Figure 1 depicts the DMARC validation process.

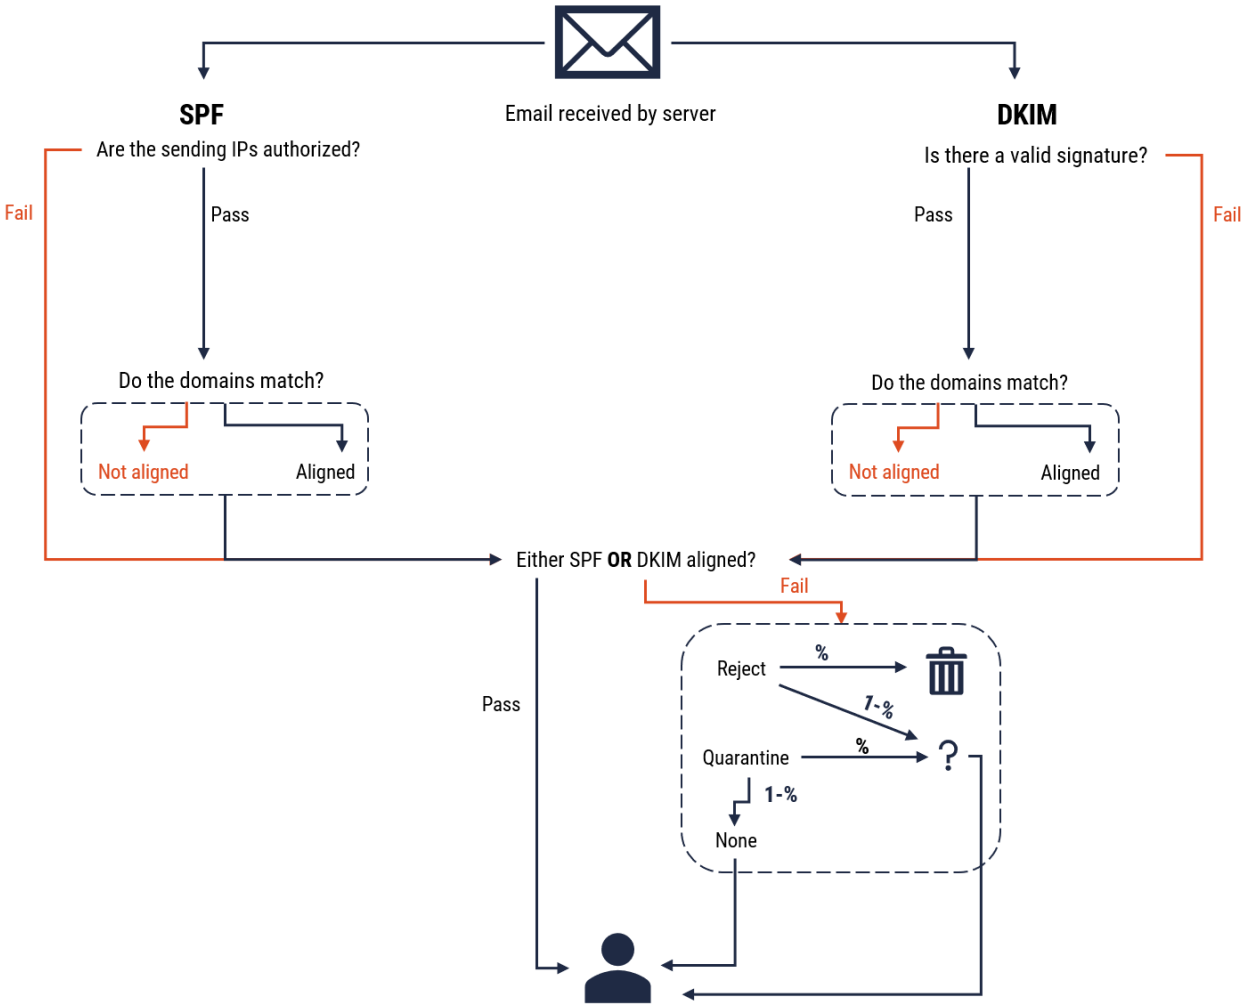

Figure 1: DMARC Validation

2.4.2 DMARC REPORTING

DMARC includes a mechanism for system owners to receive information about emails sent on behalf of their domains. You can use this information in the following ways:

- Identify the components of your organization's email infrastructure, including third-party senders;
- Confirm that SPF and DKIM have been deployed and are functioning correctly for your domains;
- Confirm that legitimate emails from your domains are being delivered successfully while illegitimate emails are rejected; and
- Identify infrastructure used by malicious actors to impersonate your domains.

DMARC aggregate reports are produced by receiving systems and are typically sent to domain owners daily. These reports are sent as email attachments to an address specified in a domain's DMARC record.

Aggregate reports are produced in a standardized Extensible Markup Language (XML) format and should be processed by an automated system wherever possible. There are many open source, free, and commercial offerings that can process DMARC reports on your behalf. The Cyber Centre has also developed a service to process DMARC reports for our partners, as described in Section 4.

Some senders may also provide failure or forensic reports, which are copies of messages that have failed one or more DMARC checks. Systems owners can request forensic reports by specifying certain parameters in a domain's DMARC record. However, we do not recommend that system owners request forensic reports, as they may potentially contain personally identifiable information (PII).

3 ADDITIONAL CONSIDERATIONS

3.1 VENDOR SUPPORT

Not all hardware, software, or service providers fully implement support for SPF, DKIM, and DMARC. Check the documentation for the components and services used in your infrastructure to confirm the level of support and identify any potential limitations.

Receiving systems may not fully implement the checks for SPF, DKIM, or DMARC, or they may ignore or override an organization's published policy. For example, most receivers who send DMARC aggregate reports do so once every 24 hours, regardless of the reporting period published in the policy for the sending domain.

3.2 THIRD-PARTY SENDERS

When implementing SPF, DKIM, and DMARC, you should consider the third parties that have been authorized to send emails on behalf of a domain.

3.2.1 THIRD PARTIES AND SPF

For emails sent by third parties to pass SPF, the sending IP addresses must be incorporated into a domain's SPF record. To accomplish this, the domain's SPF record can list the IP addresses directly, or include a reference to an SPF record that is maintained by the third party.

Note that the SPF protocol enforces a limit of 10 DNS lookups per record, which can inadvertently be exceeded depending on how referenced third-party records have been structured.

3.2.2 THIRD PARTIES AND DKIM

For emails sent by third parties to pass DKIM, they must be signed using a private key with a corresponding DKIM record published for the domain. Third parties typically accomplish this by requesting system owners to publish either a provided DKIM record or a CNAME record that redirects to a DKIM record maintained by the third party. In the latter case, the third party typically manages and rotates the cryptographic keys used, though, as a system owner, you may also be able to request that the keys be rotated on demand.

3.2.3 THIRD PARTIES AND DMARC

As noted above, DMARC enforces alignment between the *envelope from* and *header from* domains for SPF to pass. In some cases, this causes complications if third-party senders are used, and the two domains do not match. As a common workaround, you can create a custom subdomain. This subdomain is used as the *envelope from* domain (also referred to as a *custom Return-Path*) and has a corresponding CNAME record that redirects to an SPF record maintained by the third party.

You can ensure DKIM alignment by using a CNAME DKIM record, as described above.

You can also direct DMARC reports to a third party for processing. You can implement this by including a third party's email address as a reporting destination in a domain's DMARC record. In such a case, the third party must also publish a corresponding record indicating that their domain is willing to accept DMARC reports on behalf of the protected domain.

3.2.4 SUBDOMAIN SEPARATION

To protect user addresses from being spoofed, you should consider allocating dedicated subdomains to authorized third-party senders. For example, a third party could be authorized to send messages from `list.domain.test`, but not from `domain.test`. As a result, the third party cannot send an authorized message from a user address such as `chief.executive@domain.test`.

3.3 MESSAGE FORWARDING

Forwarding messages can sometimes cause SPF, DKIM, or DMARC checks to fail, depending on how the message is handled. Typically, a user can forward a message in the usual way without issues. However, sometimes there are issues when email systems forward messages automatically. For example, if the *envelope from* address is changed but the *header from* address is not, SPF alignment can fail. Similarly, if any of the content underlying the signature in a DKIM header is changed, the signature validation fails.

You cannot prevent messages from being forwarded in this way, but you should be aware that automatically forwarded messages may be rejected because of these kinds of failures. DMARC failures are noted in the aggregate reports you receive. Review these reports to assess any potential problems.

Overall, DKIM is more resilient to forwarding than SPF. Ensuring that DKIM is properly deployed ensures the delivery of forwarded legitimate messages.

3.4 INBOUND MAIL

You should configure your organization's mail infrastructure to support authentication mechanisms for received mail, including the following:

- Attempt to validate SPF, DKIM, and DMARC for all received messages;
- Reject messages in accordance with a domain's SPF and DMARC policies;
- Preserve message content and DKIM headers when forwarding messages; and
- Consider sending aggregate DMARC reports to domain owners regarding messages received.

3.5 RELATED STANDARDS

In addition to the SPF, DKIM, and DMARC, you may see references to other standards related to email security. Detailed descriptions are beyond the scope of this document, but the following subsections include a brief overview of some of these standards.

3.5.1 AUTHENTICATED RECEIVED CHAIN (ARC)

ARC is a draft standard that attempts to address the problems caused by forwarded messages. The ARC protocol is described in *RFC 8617: The Authenticated Received Chain (ARC) Protocol* [4], which was granted experimental status in July 2019. ARC allows intermediate senders to validate the headers of an email with respect to SPF and DKIM and attest to their validity with a digital signature when forwarding the message. Downstream systems can rely on this chain of attestations and choose to deliver a message, even if it has failed the systems' own DMARC checks.

ARC is an emerging standard. However, some of the Internet's largest email providers have adopted the signing and validation of messages using ARC. As a result, you may notice that ARC is mentioned in a provider's DMARC reports, particularly when a domain's DMARC policy has been overridden because of a successful ARC validation.

3.5.2 BRAND INDICATORS FOR MESSAGE IDENTIFICATION (BIMI)

BIMI is an emerging standard that will allow system owners to provide an image, typically a logo, which email clients will display to users alongside authenticated messages. To participate in BIMI, a domain must have fully implemented DMARC and have a policy of **reject**. BIMI is not yet widely adopted, but its development is supported by many large stakeholders in the email community.

3.5.3 EMAIL TRANSPORT ENCRYPTION

Historically, network traffic between mail servers was unencrypted, leaving it vulnerable to interception or modification in transit. Over time, a number of protocols have emerged to support the encryption of email traffic in transit.

3.5.3.1 STARTTLS

In 2002, the Simple Mail Transfer Protocol (SMTP) was extended to support opportunistic encryption via a STARTTLS command. You should configure your organization's MTAs to support and request the use of STARTTLS. However, an attacker in a privileged network position can prevent the delivery of this command, limiting the effectiveness of this approach.

3.5.3.2 DNS-BASED AUTHENTICATION OF NAMED ENTITIES (DANE)

DANE was introduced in 2012 to allow system owners to bind transport layer security (TLS) certificates to domain names, without relying on a central certificate authority. To address the limitations of STARTTLS, *RFC 7672: SMTP Security via Opportunistic DANE TLS* [5] describes the application of DANE to SMTP traffic, which allows system owners to require that traffic to their domain's mail servers be encrypted.

To deploy DANE, you must also adopt the prerequisite of Domain Name System Security Extensions (DNSSEC) for a domain. System owners with a DNSSEC-enabled domain can implement DANE by publishing a DANE TLS authentication (TLSA) record in DNS. Like DMARC, DANE supports the delivery of aggregate reporting to system owners.

3.5.3.3 MTA STRICT TRANSPORT SECURITY (MTA-STS)

As an alternative to DANE, SMTP MTA-STS is an emerging standard that adds support for strict encryption without relying on DNSSEC. With MTA-STS, you can specify that mail traffic sent to a domain is encrypted with a specific public encryption key.

While they achieve the same aim, DANE and MTA-STS do not conflict and you can implement them in parallel. You can implement MTA-STS by deploying signed TLS certificates to a domain's email and web servers, publishing an MTA-STS policy on the web server, and publishing MTA-STS records in DNS. Like DMARC, MTA-STS supports the delivery of aggregate reporting to system owners.

3.5.4 NO SERVICE MX RECORD

No Service or Null MX responses were proposed in 2015 as a method to publish the lack of expected infrastructure to receive messages. While SPF can be used to authorize outbound senders, it does not state if a domain is meant to receive mail. Configuring a Null MX response is an improved method to advertise when a domain is not expected to receive messages.

Both SPF and Null MX should be configured for domains which neither send nor receive messages. For messages which are not expected to receive responses, e.g. marketing messages sent by an organization, a custom autoreply subdomain can be configured using Null MX to improve the user experience if replies are attempted. Null MX responses are supported by popular cloud mail deployments.

4 CYBER CENTRE DMARC REPORTING SERVICE

To support Government of Canada (GC) departments and Canadian operators of systems of importance, the Cyber Centre has developed an automated system to receive and process raw DMARC reports and produce summaries that you can use to monitor your organization's email infrastructures.

These summaries can be used in the following ways:

- Identify components of your organization's email infrastructure, including third-party senders;
- Assist you in deploying SPF, DKIM, and DMARC;
- Confirm that legitimate emails from your organization's domains are delivered successfully and illegitimate emails are rejected;
- Identify misconfigurations that you need to correct; and
- Monitor for spoofing and identify malicious infrastructure.

To request access to this service, please contact us at contact@cyber.gc.ca.

5 SUMMARY

You can use the guidance in this document to implement technical security measures to protect your organization's domains from email spoofing. By implementing SPF, DKIM, and DMARC, you can reduce a threat actor's chance of carrying out successful malicious email campaigns leveraging the reputation of your organization.

Annex A of this document provides guidance on how to implement these three security protocols. Annex B includes a protocol reference.

5.1 CONTACT INFORMATION

For more information on implementing protection measures for email domains, email or phone our Contact Centre:

Cyber Centre Contact Centre
contact@cyber.gc.ca
(613) 949-7048 or 1-833-CYBER-88

6 SUPPORTING CONTENT

6.1 LIST OF ABBREVIATIONS

Term	Definition
A	DNS host record (IPv4)
ARC	Authenticated Received Chain
BIMI	Brand Indicators for Message Identification
CSE	Communications Security Establishment
DANE	DNS-based Authentication of Named Entities
DKIM	DomainKeys Identified Mail
DMARC	Domain-based Message Authentication, Reporting, and Conformance
DNS	Domain Name System
DNSSEC	Domain Name System Security Extensions
GC	Government of Canada
IETF	Internet Engineering Task Force
IP	Internet Protocol
MTA	Mail Transfer Agent
MTA-STS	SMTP MTA Strict Transport Security
MX	DNS mail exchange record
PII	Personally Identifiable Information
RFC	Request for Comments
SMTP	Simple Mail Transfer Protocol
SPF	Sender Policy Framework
TLS	Transport Layer Security
TLSA	Transport Layer Security Authentication
TTL	Time to Live
TXT	DNS text record
XML	Extensible Markup Language

6.2 REFERENCES

Number	Reference
1	Internet Engineering Task Force. *RFC 7208 Sender Policy Framework (SPF) for Authorizing Use of Domains in Email, Version 1*. April 2014.
2	Internet Engineering Task Force. *RFC 6376 DomainKeys Identified Mail (DKIM) Signatures*. September 2011.
3	Internet Engineering Task Force. *RFC 7489 Domain-based Message Authentication, Reporting, and Conformance (DMARC)*. March 2015.
4	Internet Engineering Task Force. *RFC 8617 The Authenticated Received Chain (ARC) Protocol*. July 2019.
5	Internet Engineering Task Force. *RFC 7672 SMTP Security via Opportunistic Domain-Based Authentication of Named Entities (DANE) Transport Layer Security (TLS).* October 2015.

Annex A Implementation Plan

A.1 Overview

To implement SPF, DKIM, and DMARC, you need to take several incremental steps. To minimize potential disruptions, we recommend that you follow the sequence below. Note that each step is described further in the following subsections of this annex.

1. **Assess**
    a. Identify all domains and subdomains used to send mail;
    b. Assess current state;
    c. Deploy initial DMARC records with policy of **none**; and
    d. Collect and analyze DMARC reports.
2. **Deploy**
    a. Identify all authorized senders;
    b. Deploy SPF records for all domains;
    c. Deploy DKIM records and keys for all domains and senders; and
    d. Monitor DMARC reports and correct misconfigurations.
3. **Enforce**
    a. Upgrade DMARC policy to **quarantine** (gradually increment enforcement from 25% to 100%);
    b. Upgrade DMARC policy to **reject** (gradually increment enforcement from 25%to 100%); and
    c. Reject all messages from non-mail domains.
4. **Maintain**
    a. Monitor DMARC reports;
    b. Correct misconfigurations and update records as required; and
    c. Rotate DKIM keys annually.

A.2 Assess

A.2.1 Identify Mail Domains

Compile a list of all domains and subdomains from which your organization sends mail.

Note: You can find acquire this information from some of these example sources:

- Existing DNS records of type MX;
- Logs of network traffic on ports 25 or 587 (typically used for SMTP); and
- Email administrators.

A.2.2 Assess Current State

Once you have identified the domains in use, use DNS tools to assess the current state of SPF, DKIM, and DMARC deployment for each.

Note: You cannot test DKIM records in this manner unless you know the selector string.

Review returned records to ensure they are correct (refer to Annex B).

There are many online tools that will assess the state of a given domain name, including validating the correctness of the returned records. You can use the following `dig` commands to retrieve existing records for each protocol:

- **SPF:** `dig +short -t txt domain.test`
- **DKIM:** `dig +short -t txt selector._domainkey.domain.test`
- **DMARC:** `dig +short -t txt _dmarc.domain.test`

A.2.3 Deploy Initial DMARC record

Deploy an initial DMARC record to each domain and subdomain that is used to send email. By deploying an initial DMARC record with a policy of **none**, you can immediately begin to receive DMARC reports, even before deploying SPF or DKIM. The initial record does not impact mail delivery in any way, but it requests that receiving systems send aggregate DMARC reports to you and/or a designated third party. You can use these reports to guide your SPF and DKIM deployments, and they can identify mail infrastructure that may not have been known at the outset.

Your organization must specify an email address to which the reports are sent. Typically, this address is a dedicated mailbox (e.g. `dmarc@domain.test`) rather than an address assigned to a named individual. You can specify up to two mailboxes. However, some senders may only send reports to the first address listed.

We recommend the following form for your organization's initial DMARC record:

```
v=DMARC1; p=none; sp=none; rua=mailto:dmarc@domain.test
```

You can specify a third party's email address in the record to have reports sent to that third party. To receive these reports, the third-party domain must also publish a DMARC record indicating that they accept reports on behalf of the domain, which is the subject of the report.

If your organization uses our DMARC report processing service, you can specify the Cyber Centre as the primary reporting address:

```
v=DMARC1; p=none; sp=none;
rua=mailto:dmarc@cyber.gc.ca,mailto:dmarc@domain.test
```

A.2.4 Collect and Analyze DMARC Reports

Within 24 hours of publishing an initial DMARC record, you should begin to receive aggregate reports from systems that have received emails from your organization's domain. Monitor the DMARC reports throughout the deployment process. These reports contain valuable information that you can use to support your efforts.

DMARC reports are generated in a standardized XML format. They should be processed by an automated system. There are many open source, free, and commercial options available. Our report processing service also handles report processing and produces summary reports for you.

When analyzing DMARC reports, you should pay attention to the following components:

- **Subdomains**: Identify subdomains that are used to send mail and verify whether these subdomains are configured with SPF, DKIM, or DMARC.
- **Internal IP addresses**: Identify internal IP addresses that are used to send mail and verify whether the sent messages pass SPF and DKIM alignment as expected.
- **External IP addresses**: Identify any external IP addresses that are used to send mail and use reverse DNS and WHOIS tools to gather more information about them. Attempt to classify these addresses as one of the following:
  - **Authorized senders**: Third parties who are authorized to send mail on your organization's behalf. Verify whether the sent messages pass SPF and DKIM alignment as expected;
  - **Forwarders / Relays**: Some mail servers, such as those supporting mailing lists, may be configured to automatically forward messages. These forwarded messages may fail SPF and DKIM checks when received but are benign; or
  - **Possible spoofing**: IP addresses that do not fall into one of the above categories may be sending unauthorized messages that spoofing your organization's domain. IP addresses that appear on public spam blacklists are more likely to fall into this category.

A.3 Deploy

A.3.1 Identify Authorized Senders

By analyzing the DMARC reports that you receive, you can identify all authorized senders, including subdomains and IP addresses. This list likely includes internal infrastructure, but it can also include cloud-based mailbox services or third-party senders.

You may need to consult other business owners in your organization to confirm whether certain senders (e.g. a third-party firm contracted to send marketing emails) are authorized.

A.3.2 Configure DNS Time To Live (TTL)

DNS records are cached for a period known as their time to live (TTL), expressed in seconds. During the Deploy and Enforce phases, you should use a TTL of 30 minutes (1800 seconds) for newly created records. This will prevent any inadvertent errors from persisting in DNS caches for a longer period before they are replaced by a corrected version.

When revising a record, you should first modify the TTL of the existing record to 30 minutes and allow the previous TTL to elapse. This will ensure that the new record will replace the old one in all caches within 30 minutes.

You should use a TTL of 24 hours (86400 seconds) for records that are stable.

A.3.3 Deploy SPF for All Domains

For each domain and subdomain that sends mail, deploy SPF records in the following sequence:

1. Construct an SPF record as described in Annex B.
   Note: The record must list all authorized IP addresses, either directly or by reference.
2. Publish the record in DNS.

A.3.4 Deploy DKIM for All Domains and Senders

For each domain that sends mail, deploy DKIM keys and records in the following sequence:

1. Generate or acquire cryptographic keys for use with DKIM through one of the following methods:
   a. MTAs may have the ability to generate DKIM key pairs directly and output the public key.
   b. You can generate key pairs using a tool such as OpenSSL as follows:

   ```
   openssl genrsa -out private2048.key 2048
   openssl rsa -in private2048.key -pubout -out public2048.key
   ```

   c. Third-party senders may manage DKIM keys internally and provide you with a public key.

   Note: Key pairs should meet the cryptographic requirements specified in Annex B.
2. Use the public key(s) to construct DKIM records according to Annex B.

3. Determine strings to use as DKIM selectors or obtain specified selectors from third parties.
4. Publish DKIM records in DNS.
5. Take the following actions for each MTA:
    a. Configure a private key and the corresponding selector to use for DKIM.
    b. Activate DKIM signing of outgoing messages.

To reduce the potential risk from a compromised private key, use unique DKIM key pairs as follows:

- Use a unique key pair, at a minimum, for each domain.
- Use unique key pairs whenever a third party manages the private key.
- Use a unique key pair for each MTA, if feasible.

A.3.5 Monitor DMARC Reports and Correct Misconfigurations

Once you have deployed SPF and DKIM records, monitor the DMARC reports you receive to detect any errors or oversights. In particular, examine the reports for any instances of messages sent from authorized senders that fail either SPF or DKIM alignment. Correct any misconfigurations.

A.4 Enforce

When you are confident that your organization's SPF and DKIM deployments are complete and all misconfigurations are corrected, you can transition the domains to enforcement.

DMARC's enforcement policy has three levels of increasing strictness:

- **None**: All messages are delivered (i.e. monitor only).
- **Quarantine**: Messages that fail DMARC are delivered but are marked as suspicious.
- **Reject**: Messages that fail DMARC are rejected.

This change advises receiving mail systems to only deliver messages that pass DMARC (i.e. pass either SPF or DKIM alignment) and to either quarantine or reject messages that fail DMARC (i.e. fail both SPF and DKIM alignment).

In addition to the level of enforcement, a DMARC policy can specify a percentage of failing messages to which the enforcement level is applied, with the next lower level applied to the remainder. For example, a policy of **reject** with a percentage of 75% rejects 75% of failed messages, while 25% are quarantined.

Your initial deployment of DMARC should gradually increase the strictness of enforcement from a baseline of **none**, through **quarantine**, to an end state of **reject** at 100%.

In the *disposition* field, DMARC reports will indicate the enforcement policy that has been applied to a specific group of messages.

A.4.1 Gradually Enforce Quarantine

For each domain, gradually transition from a policy of **none** to **quarantine** as follows:

1. Modify the policy portion of the DMARC record to `p=quarantine` and specify a percentage of 25% by adding `pct=25`.
2. Monitor DMARC reports to ensure that legitimate messages are not quarantined.
3. Increase the percentage by 25% every two weeks, if no errors are detected, until it reaches 100%.

A.4.2 Gradually Enforce Reject

For each domain, gradually transition from a policy of **quarantine** to **reject** as follows:

1. Modify the policy portion of the DMARC record to `p=reject` and specify a percentage of 25%.
2. Monitor DMARC reports to ensure that legitimate messages are not rejected.
3. Increase the percentage by 25% every two weeks, if no errors are detected, until it reaches 100%.

A.4.3 Reject All Messages To and From Non-Mail Domains

Once DMARC is deployed for all mail domains, you should also deploy SPF and DMARC records to protect domains that do not send mail. These records indicate that any messages received from these domains are illegitimate and

should be rejected. You should also deploy a “no service” or null MX record for all domains that do not receive mail.

Deploy records for non-mail domains as follows:

DMARC

1. Modify the subdomain policy of existing root-level DMARC records to `sp=reject`.
   Note: This serves as the default policy for all subdomains that do not have a DMARC record of their own.
2. Create a DMARC record with `p=reject` and `p=100` for any root domains that do not send mail.

SPF

1. Create an SPF record with a policy of `-all` for each domain and subdomain that does not send mail.
2. Create a wildcard record (e.g. `*.domain.test`) with a policy of `-all` if no subdomains of a given root domain send mail.
   Note: This serves as the policy for all subdomains.

NULL MX

1. Create an MX record with a hostname of ” . “ for each domain and subdomain that does not receive mail.

A.5 Maintain

A.5.1 Monitor DMARC Reports

Once SPF, DKIM, and DMARC are deployed, you should monitor DMARC reports for any changes. Pay attention to the following components:

- An increase in the number of messages failing SPF or DKIM alignment;
- The use of new subdomains;
- The use of new internal IP addresses;
- The ongoing use of new external IP addresses, which may indicate new third-party senders; and
- The volume of suspected spoofing.

A.5.2 Correct Misconfigurations and Update Records

As infrastructure changes and misconfigurations are uncovered, you will need to periodically adjust the SPF, DKIM, and DMARC records. You may also need to deploy new records in response to changes in your organization's domain structure.

SPF records are tied to IP addresses and are likely to require the most attention. Using CNAME records for third-party senders and A or MX fields for internal infrastructure can insulate SPF records from downstream changes. You may need to change IPv4 or IPv6 fields more frequently. Similarly, you may need to add or change DKIM records as you add or remove authorized senders or MTAs.

Once your organization has reached a policy of **reject**, DMARC records should remain largely stable.

A.5.3 Rotate DKIM Keys

To guard against the potential compromise of one or more private keys, you should rotate DKIM key pairs annually. You should rotate key pairs that you manage according to the following sequence:

1. Generate new key pairs.
2. Choose or obtain new selectors.
3. Prepare and publish new records.
4. Configure MTAs to use the new private keys and selectors.
5. Delete the old records after one week.

Confirm the key rotation policy of third-party senders and ensure that their keys are rotated at least annually.

Annex B Protocol Reference

B.1 SPF

B.1.1 SPF Records

SPF records are published as TXT records at the root of a sending domain and at each subdomain. Each record includes a list of all the IP addresses that are permitted to send email on behalf of the domain, either directly or by referring to further DNS records that contain these addresses.

Stable SPF records should have a long time-to-live (TTL), such as 24 hours (86400 seconds).

Elements in SPF records must be separated by spaces, and are assumed to be positive or permissive if they are not preceded by a modifier character (e.g. –) to indicate otherwise. A typical SPF record for the domain `domain.test` would be:

```
v=spf1 mx a:mail.domain.test include:spf.third-party.test -all
```

Table 1 described the most commonly used elements of an SPF record.

Table 1: SPF Record Elements

Element	Description
`v=spf1`	The standard header for an SPF record.
`mx`	Incorporates any MX (mail) records published for the domain or for specified hostname(s).
`a`	Incorporates any A (address) records published for the domain or for specified hostname(s).
`ip4`	An IPv4 address or range.
`ip6`	An IPv6 address or range.
`include`	Incorporates SPF records published at another DNS location. Typically used to authorize third-party senders.
`-all`	Excludes all other IP addresses. Typically used as the last element of a record to indicate that all other senders are unauthorized.
`-`	When prepended to an element, indicates that messages from the modified element are unauthorized and should be rejected. It is typically used in conjunction with `all` as the final element of a record to indicate that all other senders are unauthorized (i.e. `-all`).

B.1.2 Third-Party Senders

Third-party infrastructure is typically authorized using an `include` clause, which instructs the receiving email server to include any SPF records found at the specified location. For example:

```
include:spf.third-party.test
```

B.1.3 DNS Lookup Limit

The SPF protocol has a limit of 10 DNS lookups per record. If multiple third-party `include` statements are used, this limit can inadvertently be exceeded. One way to work around this limit is to separate email senders into distinct subdomains. As a result, the SPF records for each requires fewer SPF elements and associated lookups than would be needed in a combined record.

B.2 DKIM

B.2.1 DKIM Records

DKIM records are TXT records that are specified by a selector string. They are created below a `_domainkey` location, situated one level below a protected domain. For example, the location of a DKIM record for the domain `domain.test` and with the selector `abc` is `abc._domainkey.domain.test`.

Stable DKIM records should have a long time-to-live (TTL), such as 24 hours (86400 seconds).

Elements in DKIM records must be separated by semicolons. Table 2 describes the typical elements of a DKIM record.

Table 2: DKIM Record Elements

Element	Description
`v=DKIM1`	The standard header for a DKIM record.
`p`	Public key value

B.2.2 Cryptographic Considerations

You should configure DKIM to use the RSA-SHA256 algorithm, and you should use keys that are at least 2048 bits in length. Note that because 2048-bit keys exceed the limit of 255 characters of a single TXT record, they must span two adjacent records without additional spaces inserted.

You should rotate DKIM keys on an annual basis as described in Annex A.

B.2.3 Third-Party Senders

Third-party infrastructure is typically authorized by using a specific selector that identifies the key used by the sender. In some cases, the third party may manage the keys and provide you with the public key and selector to be used. Alternatively, the third party may provide you with a DNS location that is to be published as a CNAME record and that points to a DKIM record maintained by the third party.

Some third-party senders will rotate DKIM keys managed by them automatically, while others require an intervention from the system owner to initiate the rotation.

B.3 DMARC

B.3.1 DMARC Records

DMARC records are TXT records created at the `_dmarc` location below the protected domain (i.e.`_dmarc.domain.test`).

Stable DMARC records should have a long time-to-live (TTL), such as 24 hours (86400 seconds).

DMARC record elements must be separated by semicolons. Table 3 lists the typical elements of a DMARC record.

Table 3: DMARC Record Elements

Element	Description
`v=DMARC1`	The standard header for a DMARC record.
`p`	The policy to apply to messages that fail DMARC. Must be either “none”, “quarantine”, or “reject”.
`sp`	The policy to apply to subdomains that do not have their own DMARC record. Must be either “none”, “quarantine”, or “reject”.
`pct`	The percentage of failed messages to which a “quarantine” or “reject” decision should be applied. The remainder are acted on at the next more relaxed policy level, i.e. “none” instead of “quarantine” or “quarantine” instead of “reject”. Must be an integer between 1 and 100.
`rua`	The address(es) to which aggregate reports should be sent. Multiple addresses can be specified, separated by commas and with each preceded by `mailto:`.

If your organization does not currently have a record or record processing capability, we recommend the following initial DMARC record:

```
v=DMARC1; p=none; sp=none; rua=mailto:dmarc@cyber.gc.ca
```

When implementation is complete, the record should resemble the following:

```
v=DMARC1; p=reject; pct=100; sp=reject; rua=mailto:dmarc@cyber.gc.ca
```