15.2. Email Infrastructure

Objective

15.2.1. Email infrastructure is hardened, email is secured, and protective marking of email messages is enforced.

Context

Scope

15.2.2. This section covers information on email infrastructure security, specifically email authentication and secure email transport. Information on using email applications can be found in Section 15.1 - Email Applications and Section 9.3 - Using the Internet.

15.2.3. Email is a critical tool for communication, and is commonly targeted by cyber-attacks. It is hard to protect because email is historically insecure and subject to social engineering. This section covers email authentication, secure email transport, and securing non-email sending domains.

Email authentication

15.2.4. Email authentication protects your email from impersonation attacks like spoofing and phishing; phishing and malware distribution attacks are common internet security threats. To avoid agency domains being used fraudulently (e.g. for spam or spear-phishing), the following should be implemented:

- Sender Policy Framework (SPF)
- DomainKeys Identified Mail (DKIM)
- Domain-based Message Authentication, Reporting & Conformance (DMARC) records

15.2.5. Implementation of these features will help other mail servers authenticate the email they receive from your domains. It is important to note that DMARC is designed to protect against direct domain spoofing only. DMARC does not eliminate the need for additional forms of protection and analysis. It does, however, provide a way for participating senders and receivers to coordinate protective activities and streamline security processes.

Domain-based Message Authentication, Reporting and Conformance (DMARC)

Vocabulary

15.2.6 The terms “none”, “reject” and “quarantine” are used to describe DMARC actions based on policy modes. In this usage:

- “none” means no action on the transmission or receipt of the email but continue to collect data and send reports;
- By default, email under a “reject” policy setting is not delivered. “Reject” either:
  - refuses to accept non-compliant email, or
  - initially accepts the non-compliant email but prevents an email reaching the user. The acceptance process can generate a Delivery Status Notification (block/“bounce”) or simply delete/drop the email (block/delete);
- “quarantine” prevents an email from reaching the user but safely storing it so it can be accessed if required (a potentially suspicious email and/or attachment subject to additional scrutiny). Quarantined items can be released following a review and release process.

What is DMARC

15.2.7. Domain-based Message Authentication, Reporting and Conformance (DMARC) is an email authentication policy and reporting protocol that:

- complements and unifies the existing validation checks performed by SPF and DKIM;
- checks the stated origin of inbound emails using a combination of Sender Policy Framework (SPF) and DomainKeys Identified Mail (DKIM);
- establishes a recipient email response for emails that fail the check;
- requests recipient email services to report email sources and origins;
- provides visibility over potentially illegitimate or fraudulent email.

15.2.8. DMARC builds on SPF and DKIM protocols, adding links to the author (“From:”) domain name, published policies for recipient handling of authentication failures, and reporting from receivers to senders, in order to improve and monitor protection of the recipient domain from fraudulent email.

15.2.9. Most email services will check your DMARC record and send aggregated reports including details of all email the service received from the agency, and its origin. This assists in identifying if an individual within the agency is sending email inappropriately or if the agency domain is being spoofed.

Background, Reference and Implementation Guidance Sources

15.2.10. The IETF published RFC 7489, "Domain-based Message Authentication, Reporting, and Conformance (DMARC)" 18 March 2015. This is the principal standards guidance on the implementation and use of DMARC. Further guidance is available from The Global Cyber Alliance (GCA) - see References below.

Using DMARC

15.2.11. The combination of DMARC, SPF, and DKIM records in DNS automates the ability of email service providers to confirm which servers should be legitimately sending email from the agency's domain, and what action to take with mail received from any other domains.

15.2.12. As a pre-requisite for implementing DMARC, agencies **must** publish an SPF and a DKIM record in DNS. Agencies must also ensure that emails sent by the agency (including from third party services sending on behalf of the agency) have a DKIM signature that matches the signature in the DKIM record.

15.2.13. Agencies can choose to quarantine or reject messages that fail checks. More specifically:

- Sender Policy Framework (SPF) is used to specify legitimate locations of servers which can send email for your domain;
- DomainKeys Identified Mail (DKIM) isn't supported by all mail servers, but if it is, it can be used to cryptographically sign outgoing mail sent by your servers to give email service providers further confidence that it's legitimate;
- DMARC is used to inform email service providers what action they should take if SPF or DKIM (or both) validation fails;
- One important aspect of DMARC is the action you ask email service providers to take when SPF or DKIM validation fails:
    - a policy of p=none means that they should allow non-compliant emails to be delivered but report the failure to the agency;
    - a policy of p=quarantine requests that they mark the email as spam;
    - a policy of p=reject requests the email service provider to refuse to deliver the email.

15.2.14. A text record published in DNS is used to notify other organisations of the use of DMARC. The following demonstrates an example DMARC record:

- v=DMARC1;
- p=quarantine;
- rua=mailto:dmarc@agency.govt.nz (where agency is the name of the respective agency).

15.2.15. This informs email recipients that:

- you have a DMARC policy (v=DMARC1)
- any messages that fail DMARC checks should be treated as spam (p=quarantine)
- they should send reports of email received back to you (rua=mailto:dmarc@agency.govt.nz)

15.2.16. It is not unusual to experience minor errors in syntax or other elements of DMARC configuration when first implementing DMARC. Some discussion on common problems, issues and solutions can be found on the DMARC website (see the References table below).

15.2.17. For domains expected to be used for email, it is not recommended to move directly to a full implementation of DMARC until there is certainty that the configuration and implementation are stable and operating as intended. The following implementation outline is recommended by the GCA and DMARC organisations (see 15.2.19 References):

1. Deploy DKIM & SPF;
2. Ensure mailers are correctly aligning the appropriate identifiers;
3. Publish a DMARC record with the "none" flag set for the policies, which requests data reports;
4. Analyse the data and modify mail streams as appropriate; and
5. Modify DMARC policy flags from "none" to "quarantine" to "reject" as experience dictates.

Securing non-email sending domains

15.2.18. Threat actors may attempt to spoof a domain that is not intended to send or receive email. Operators of domains that do not send mail can publish Sender Policy Framework (SPF) "-all" policies to make an explicit declaration that the domains send no mail. The following DNS entries create blank SPF and DKIM entries effectively forcing every email sent (spoofed) from a non-mail-enabled domain to fail all checks, giving the highest possible chance of these emails not reaching their intended destination.

- SPF "v=spf1 -all"
- DKIM "v=DKIM1; p="
- DMARC "V=DMARC1;p=reject;sp=reject;adkim-s;aspf=s;rua=mailto:<your dmarc report RUA email address>;"

DMARC Reporting

15.2.19 DMARC reporting provides information to assist an agency's IT system and email administrators. It can also provide an email asset inventory as well as providing data on spam, phishing and other email exploitation techniques.

15.2.20. DMARC can be configured to produce an aggregate report and a forensic report. In some cases agencies may also send reports to an external organisation such as a DMARC reporting service or a third-party IT service provider. Discretion should be used when providing such information to third parties in order to maintain security and privacy.

Secure email transport

15.2.21. Over time, several protocols have emerged to support the encryption of email traffic in transit. Secure email transport protects your email from person-in-the middle attacks like eavesdropping, manipulation, and cryptographic downgrade.

15.2.22. With parties that you contact regularly by email (eg, partner organisations), it is possible to configure connections so that TLS will always be used, and the certificates presented by both mail servers are authenticated, verifying the identities of both parties.

STARTTLS, Opportunistic TLS

15.2.23. Opportunistic TLS is configured on the sending server. STARTTLS (RFC 3207) is a protocol command that upgrades an insecure connection to a secure connection using TLS. Agencies must enable opportunistic TLS encryption.

Using MTA-STS

15.2.24. SMTP MTA Strict Transport Security (MTA-STS) is a standard that adds support for strict encryption without relying on DNSSEC (RFC 8461). With MTA-STS, you can specify that mail traffic sent to a domain is encrypted with a specific public encryption key. Agencies should enable MTA-STS.

15.2.25. The objectives of MTA-STS are to:

- make it harder for an attacker to intercept and redirect emails
- make sure that TLS encryption is always used,
- prevent attackers downgrading email encryption on emails to cleartext,
- provide visibility reports through the TLS reporting protocol.

15.2.26. To achieve this, MTA-STS works in the following ways:

- Your organisation can advertise the mail server hostname on a separate secure web page, which means an attacker cannot subvert your DNS entry (specifically your MX record, which is the record that indicates how an email should be routed using the Simple Mail Transfer Protocol).

- Your organisation can publish an 'MTA-STS enforce policy', which tells any server sending you emails to always send with TLS encryption, and to not allow connections to be downgraded. If there is a failure to establish a secure TLS connection, emails will not be delivered. Note that when your organisation sets up MTA-STS, you are securing inbound connections only.

15.2.27. MTA-STS is relatively simple to implement, but organisations must step through their implementation with care; if you switch on controls too quickly then inbound emails may not be delivered. To mitigate this risk, we always recommend using MTA-STS in 'testing mode' first, and to set up TLS-RPT (TLS Reporting) as a mechanism for getting feedback before you progress.

15.2.28. The goal ultimately is to work towards implementing an MTA-STS policy of 'enforce'. When a sending email service detects you have an MTA-STS policy of 'enforce', they should only send email to your domain if the connection is secure. Note however, that whilst many major email providers have built support for MTA-STS, there is not yet complete support from all vendors. For those that don't yet support it, emails will continue to be delivered in all cases.

15.2.29. While they achieve the same aim, DNS-based Authentication of Named Entities (DANE) and MTA-STS do not conflict and can be implemented 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.

15.2.30. MTA-STS (Mail Transfer Agent – Strict Transport Security) is configured on the receiving server, forcing inbound email connections to use TLS if the sending server supports MTA-STS; there is no drop down to cleartext. If the sending server does not support MTA-STS the connection can use opportunistic TLS, or even send in the clear.

TLS reporting

15.2.31. Organisations should also enable TLS reporting when implementing MTA-STS (RFC 8460). By implementing TLS reporting, organisations will be able to see the performance of its domains, the success or failure rate, and the impact of the organisation's MTA-STS and policies. This will give organisations important insight into which mail services it needs to configure to ensure no interruption in mail flow.

Secure cryptography

15.2.32. Historically, network traffic between mail servers was unencrypted, leaving it vulnerable to interception or modification in transit. Although it is possible to encrypt individual emails using protocols like PGP or S/MIME, this requires the sender and recipient to have the necessary trust infrastructure in place. Approved implementation of S/MIME and OpenPGP are discussed in Chapter 17, 17.6 Secure Multipurpose Internet Mail

Extension, and 17.7 OpenPGP Message Format.

15.2.33. This is not likely to be possible for all the parties you communicate with. So, your email servers should be configured to support encryption of the communications channel that the email is sent over. This task is best handled by TLS. Agencies should use the current version of TLS, see [CID:2598]. When implementing MTA-STS, agencies should use TLS1.2 or above on email servers.

15.2.34. Implementation details

- Configure all your email servers to support TLS, regardless of whether they are accessible from the internet or private networks only.
- Ensure your email servers present a certificate with suitable cryptographic properties and that it is signed by a trusted Certificate Authority.
- Ensure your email servers prefer to use a good cryptographic profile, but that lesser cryptographic profiles are supported too. You should support the case where TLS is not used at all, if you want to be confident you can receive emails from any entity.
- Where possible, consider configuring your email service to force TLS between you and organisations you regularly communicate with.

References

15.2.35 Further information on email security is available from the following sources:

Reference	Title	Publisher	Source
RFC 3207	**SMTP Service Extension for Secure SMTP over Transport Layer Security**	IETF	RFC 3207 - SMTP Service Extension for Secure SMTP over Transport Layer Security (ietf.org)
RFC 4686	**Analysis of Threats Motivating DomainKeys Identified Mail (DKIM)**	IETF	RFC 4686 - Analysis of Threats Motivating DomainKeys Identified Mail (DKIM) (ietf.org)
RFC 6376	**DomainKeys Identified Mail (DKIM) Signatures**	IETF	RFC 6376 - DomainKeys Identified Mail (DKIM) Signatures (ietf.org)
RFC 5617	**DomainKeys Identified Mail (DKIM) Author Domain Signing Practices (ADSP)**	IETF	RFC 5617 - DomainKeys Identified Mail (DKIM) Author Domain Signing Practices (ADSP) (ietf.org)
RFC 7817	**Updated Transport Layer Security (TLS) Server Identity Check Procedure for Email-Related Protocols**	IETF	RFC 7817 - Updated Transport Layer Security (TLS) Server Identity Check Procedure for Email-Related Protocols (ietf.org)
RFC 7208	**Sender Policy Framework (SPF) for Authorizing Use of Domains in Email, Version 1**	IETF	RFC 7208 - Sender Policy Framework (SPF) for Authorizing Use of Domains in Email, Version 1 (ietf.org)
RFC 7489	**Domain-based Message Authentication, Reporting, and Conformance (DMARC)**	IETF	RFC 7489 - Domain-based Message Authentication, Reporting, and Conformance (DMARC) (ietf.org)
RFC 7960	**Interoperability Issues between Domain-based Message Authentication, Reporting, and Conformance (DMARC) and Indirect Email Flows**	IETF	RFC 7960 - Interoperability Issues between Domain-based Message Authentication, Reporting, and Conformance (DMARC) and Indirect Email Flows (ietf.org)
RFC 8463	**A New Cryptographic Signature Method for DomainKeys Identified Mail (DKIM)**	IETF	RFC 8463 - A New Cryptographic Signature Method for DomainKeys Identified Mail (DKIM) (ietf.org)

NIST SP 800-45 Version 2	**Guidelines on Electronic Mail Security**	NIST	SP 800-45 Version 2, Guidelines on Electronic Mail Security | CSRC (nist.gov)
NIST SP 800-177 Rev. 1	**Trustworthy Email**	NIST	SP 800-177 Rev. 1, Trustworthy Email | CSRC (nist.gov)
	Email Authentication Mechanisms: DMARC, SPF and DKIM	NIST	Email Authentication Mechanisms: DMARC, SPF and DKIM | NIST
	DMARC	DMARC	dmarc.org – Domain Message Authentication Reporting & Conformance
	Guidelines for Email	ASD	Guidelines for Email | Cyber.gov.au
	Enhance Email & Web Security	CISA	Enhance Email & Web Security | CISA
	Implementation guidance: email domain protection	CCCS	Implementation guidance: email domain protection (ITSP.40.065 v1.1) - Canadian Centre for Cyber Security
	Email security and anti-spoofing	NCSC-UK	Email security and anti-spoofing - NCSC.GOV.UK

Rationale & Controls

Domain-based Message Authentication, Reporting and Conformance (DMARC)

15.2.36.R.01. **Rationale**

Phishing and malware distribution attacks are common internet security threats. To limit the possibility of agency domains being used fraudulently (e.g. for spam or spear-phishing), agencies should implement:

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

15.2.36.R.02. **Rationale**

It is important to note that DMARC depends on the proper implementation of both SPF and DKIM. DMARC records are published in the DNS and provide guidance to the email receiver on actions to take when emails received do not conform to the published record.

15.2.36.C.01. **Control System Classifications(s): All Classifications; Compliance: Must** [CID:6019]

Before implementing DMARC agencies MUST:

- Create a DMARC policy;
- List all domains, in particular those used for the sending and/or receiving of email;
- Review the configuration of SPF and DKIM for all active domains and all published domains; and
- Establish one or more monitored inboxes to receive DMARC reports.

15.2.36.C.02 **Control System Classifications(s): All Classifications; Compliance: Must** [CID:6020]

Agencies MUST enable DMARC with a policy of p=reject for all email originating from or received by their domain(s). See 15.2.16 for a recommended approach where a domain is used for sending and/or receiving email and disruption would cause a business impact.

15.2.36.C.03 **Control System Classifications(s): All Classifications; Compliance: Must** [CID:7519]

Agencies MUST manage "received DMARC messages" in accordance with the agency's published DMARC policy.

15.2.36.C.04 **Control System Classifications(s): All Classifications; Compliance: Must** [CID:6021]

Agencies MUST review DMARC reports on a regular basis and address any identified anomalies or security issues.

15.2.36.C.05 **Control System Classifications(s): All Classifications; Compliance: Should** [CID:7520]

Agencies SHOULD produce failure reports and aggregate reports according to the agency's DMARC policies.

Filtering suspicious emails and attachments

15.2.37.R.01. **Rationale**

The intent of blocking specific types of emails is to reduce the likelihood of phishing emails and emails or attachments containing malicious code entering the agency's networks.

15.2.37.C.01. **Control System Classifications(s): All Classifications; Compliance: Should** [CID:1745]

Agencies SHOULD configure the following gateway filters:

- inbound and outbound email, including any attachments, that contain:
  - malicious code;
  - content in conflict with the agency's email policy;
  - content that cannot be identified;
  - deny listed or unauthorised filetypes; and
- encrypted content, when that content cannot be inspected for malicious code or authenticated as originating from a trusted source;
- emails addressed to internal-use only email aliases with source addresses located from outside the domain; and
- all emails arriving via an external connection where the source address uses an internal agency domain name.

Active web addresses (URL) embedded in emails

15.2.38.R.01. **Rationale**

Spoofed emails often contain an active (embedded) email address directing users to a malicious website in order to infect the workstation or agency systems with malicious code.

15.2.38.R.02. **Rationale**

An effective defence is to strip and replace active addresses and hyperlinks with text only versions.

15.2.38.C.01 **Control System Classifications(s): All Classifications; Compliance: Should** [CID:1749]

Email servers SHOULD be configured to strip active addresses and URL's and replace them with text only versions.

Preventing unmarked or inappropriately marked emails

15.2.39.R.01. **Rationale**

Unmarked or inappropriately marked emails can be blocked at two points, the workstation or the email server. The email server is often the preferred location to block emails as it is a single location under the control of system administrators that can enforce the requirement for the entire network. In addition email servers can apply controls for emails generated by applications.

15.2.39.R.02. **Rationale**

Whilst blocking at the email server is considered the most appropriate control there is an advantage in also blocking at the workstation. This approach adds an extra layer of security and will also reduce the likelihood of a data spill occurring on the email server.

15.2.39.R.03. **Rationale**

For classified systems is it important to note that all emails containing classified information MUST be protectively marked. This requirement is outlined in Section 15.1 - Email Applications.

15.2.39.C.01 **Control System Classifications(s): Top Secret, Confidential, Secret; Compliance: Must** [CID:1754]

Agencies MUST prevent unmarked and inappropriately marked emails being sent to intended recipients by blocking the email at the email server, originating workstation or both.

15.2.39.C.02 **Control System Classifications(s): Secret, Confidential, Top Secret; Compliance: Must** [CID:1755]

Agencies MUST enforce protective marking of emails so that checking and filtering can take place.

15.2.39.C.03 **Control System Classifications(s): All Classifications; Compliance: Should** [CID:1756]

Agencies SHOULD enforce protective marking of emails so that checking and filtering can take place.

Blocking of outbound emails

15.2.40.R.01. **Rationale**

Blocking an outbound email with a valid protective marking or endorsement (e.g. NZEO) that indicates the email exceeds the classification of the communication path, stops data spills.

15.2.40.R.02. **Rationale**

Agencies may remove protective markings from emails destined for private citizens and businesses once they have been approved for release from the agency's gateways.

15.2.40.C.01 **Control System Classifications(s): All Classifications; Compliance: Must** [CID:1760]

Agencies MUST configure systems to block any outbound emails with a protective marking or endorsement indicating that the content of the email exceeds the classification of the communication path.

15.2.40.C.02 **Control System Classifications(s): All Classifications; Compliance: Should** [CID:1761]

Agencies SHOULD configure systems to log every occurrence of a blocked email.

Blocking of inbound emails

15.2.41.R.01. **Rationale**

Blocking an inbound email with a valid protective marking that indicates the email or its attachment exceeds the classification the receiving system is accredited to process will prevent a data spill from occurring on the receiving system.

15.2.41.C.01 **Control System Classifications(s): All Classifications; Compliance: Must** [CID:1764]

Agencies MUST configure email systems to reject, log and report inbound emails with protective markings indicating that the content of the email exceeds the accreditation of the receiving system.

15.2.41.C.02 **Control System Classifications(s): All Classifications; Compliance: Should** [CID:1765]

Agencies SHOULD notify the intended recipient of any blocked emails.

Undeliverable messages

15.2.42.R.01. **Rationale**

Undeliverable or "bounce" emails are commonly sent by email servers to the original sender when the email cannot be delivered, often because the destination address is invalid. Because of the common spamming practice of spoofing sender addresses, this can result in a large amount of bounce emails being sent to an innocent third party. Sending bounces only to senders that can be verified via the Sender Policy Framework (SPF) or other trusted means avoids contributing to this problem and allows other government agencies and trusted parties to receive legitimate bounce messages. See also 15.2.15 - Sender Policy Framework.

15.2.42.C.01 **Control System Classifications(s): All Classifications; Compliance: Should** [CID:1768]

Agencies SHOULD **only** send notification of undeliverable, bounced, or blocked emails to senders that can be verified via SPF or other trusted means.

Automatic forwarding of emails

15.2.43.R.01. **Rationale**

Unsecured automatic forwarding of emails can pose a serious risk to the unauthorised disclosure of classified information, for example, a system user may set up a server-side rule to automatically forward all emails to a personal email account. This can result in classified emails being forwarded to the personal email account. This tactic may also be employed when an email account is compromised.

15.2.43.C.01 **Control System Classifications(s): All Classifications; Compliance: Must** [CID:1771]

Agencies MUST ensure that the requirements for blocking unmarked and outbound emails are also applied to automatically forwarded emails.

Open relay email servers

15.2.44.R.01. **Rationale**

An open relay email server (or open mail relay) is a server that is configured to allow anyone on the Internet to send emails through the server. Such configurations are highly undesirable as they allow spammers and worms to exploit this functionality to send emails through the server. Although no longer considered a significant vector for spam email, mail servers identified as an open relay will still be added to reputation based block lists.

15.2.44.C.01 **Control System Classifications(s): All Classifications; Compliance: Must** [CID:1774]

Agencies MUST disable open email relaying so that email servers will only relay messages destined for the agency's domain(s) and those originating from authorised systems or users within that domain.

Email server maintenance activities

15.2.45.R.01. **Rationale**

Email servers perform a critical business function for many agencies; as such it is important that agencies perform regular email server auditing, security reviews and vulnerability analysis activities.

15.2.45.C.01 **Control System Classifications(s): All Classifications; Compliance: Should** [CID:1777]

Agencies SHOULD perform regular email server auditing, security reviews and vulnerability analysis activities.

Centralised email gateways

15.2.46.R.01. **Rationale**

Without a centralised email gateway it is exceptionally difficult to deploy Sender Policy Framework (SPF), DomainKeys Identified Mail (DKIM) and outbound email protective markings verification.
Attackers will almost invariably avoid using the primary email server when sending malicious emails. This is because the backup or alternative gateways are often poorly maintained with out-of-date deny lists and content filtering.

15.2.46.C.01 **Control System Classifications(s): All Classifications; Compliance: Must** [CID:1780]

Where an agency has system users that send email from outside the agency's network, an authenticated and encrypted channel MUST be configured to allow email to be sent via the centralised email gateway.

15.2.46.C.02 **Control System Classifications(s): All Classifications; Compliance: Should** [CID:1781]

Agencies SHOULD route email through a centralised email gateway.

15.2.46.C.03 **Control System Classifications(s): All Classifications; Compliance: Should** [CID:1782]

Where backup or alternative email gateways are in place, additional email gateways SHOULD be maintained at the same standard as the primary email gateway.

Transport Layer Security (TLS)

15.2.47.R.01. **Rationale**

Email can be intercepted anywhere between the originating email server and the destination email server. Email transport between organisations and agencies is usually over the internet or other unsecured public infrastructure so it is important that email interception is carefully managed and suitable controls applied. One effective measure is to use TLS to encrypt the email traffic **between email servers**.

15.2.47.R.02. **Rationale**

Enabling TLS on the originating and accepting email server will defeat passive attacks on the network, with the exception of cryptanalysis against email traffic. TLS encryption **between email servers** will not interfere with email content filtering schemes. Email servers will remain compatible with other email servers as IETF's RFC 3207 specifies the encryption as opportunistic

15.2.47.C.01 **Control System Classifications(s): All Classifications; Compliance: Must** [CID:1786]

Agencies MUST enable opportunistic TLS encryption as defined in IETF's RFC 3207 on email servers that make incoming or outgoing email connections over public infrastructure.

15.2.47.C.02 **Control System Classifications(s): All Classifications; Compliance: Should** [CID:1787]

Agencies SHOULD implement TLS between email servers where significant volumes of classified information are passed via email to other agencies.

Mail transfer agent - strict transfer security

15.2.48.R.01. **Rationale**

Where government-to-business and government-to-citizen communications require a higher level of transport security, organisations should consider implementing SMTP MTA Strict Transport Security ("MTA-STS"). MTA-STS provides organisations with a mechanism to encrypt communications between SMTP servers via TLS, preventing Person-In-The-Middle (PITM) attacks during email delivery.

15.2.48.R.02. **Rationale**

Organisations should verify the following prior to deploying MTA-STS:

- internet-facing mail relays support SMTP over TLS version 1.2 or later,
- web server hosting the policy file supports TLS (HTTPS),
- internet-facing mail relays use a TLS certificate issued by a root certificate authority that is not expired and matches its domain name.

15.2.48.C.01. **Control System Classifications(s): All Classifications; Compliance: Should** [CID:7521]

Agencies SHOULD enable MTA-STS to prevent the unencrypted transfer of emails between complying servers.

15.2.48.C.02. **Control System Classifications(s): All Classifications; Compliance: Must** [CID:7522]

Agencies MUST use TLS 1.2 or above when implementing MTA-STS.

15.2.48.C.03. **Control System Classifications(s): All Classifications; Compliance: Should** [CID:7523]

Agencies SHOULD enable TLS reporting when implementing MTA-STS.

Sender Policy Framework (SPF)

15.2.49.R.01. **Rationale**

The Sender Policy Framework (SPF) is an open standard specifying a technical method to prevent sender address forgery.
An SPF-protected domain is less attractive to spammers and phishers because the forged e-mails are more likely to be caught in spam filters which check the SPF record. Because an SPF-protected domain is less attractive as a spoofed address, it is less likely to be deny listed by spam filters and so is less disruptive to email traffic.

15.2.49.R.02. **Rationale**

Having a proper Sender Policy Framework (SPF) record increases the chances people will get emails you send. Without one, your email has a greater chance of being marked as Spam.

15.2.49.R.03. **Rationale**

SPF and alternatives such as Sender ID aid in the detection of spoofed email server address domains. The SPF record specifies a list of IP addresses or domains that are allowed to send mail from a specific domain. If the email server that transmitted the email is not in the list, the verification fails (there are a number of different fail types available).

15.2.49.C.01. **Control System Classifications(s): All Classifications; Compliance: Must** [CID:1792]

Agencies MUST:

- specify mail servers using SPF or Sender ID; and
- mark, block or identify incoming emails that fail SPF checks for notification to the email recipient.

15.2.49.C.02. **Control System Classifications(s): All Classifications; Compliance: Must** [CID:1793]

Agencies MUST:

- use a fail SPF record when specifying email servers; and
- use SPF or Sender ID to verify the authenticity of incoming emails.

15.2.49.C.03 **Control System Classifications(s): All Classifications; Compliance: Should** [CID:1794]

Agencies SHOULD refer to the SPF recommendations in IETF's RFC 7208.

DomainKeys Identified Mail (DKIM)

15.2.50.R.01. **Rationale**

DKIM enables a method of determining spoofed email content. The DKIM record specifies a public key that will sign the content of the message. If the signed digest in the email header doesn't match the signed content of the email the verification fails.

15.2.50.C.01 **Control System Classifications(s): All Classifications; Compliance: Must** [CID:1798]

Agencies MUST enable DKIM signing on all email originating from their domain.

15.2.50.C.02 **Control System Classifications(s): All Classifications; Compliance: Must** [CID:1797]

Agencies MUST use DKIM in conjunction with SPF.

15.2.50.C.03 **Control System Classifications(s): All Classifications; Compliance: Should** [CID:1799]

Agencies SHOULD verify DKIM signatures on emails received, taking into account that email distribution list software typically invalidates DKIM signatures.

15.2.50.C.04 **Control System Classifications(s): All Classifications; Compliance: Should** [CID:1800]

Where agencies operate email distribution list software used by external senders, agencies SHOULD configure the software so that it does not impair the validity of the sender's DKIM signature.