



How to Combat Fake Emails

First published: December 2011

Last updated: October 2021

Introduction

Adversaries commonly conduct social engineering and spear phishing attacks against organisations using fake emails. By modifying the sender's address, or other parts of an email header to appear as though the email originated from a different source, an adversary is able to increase the likelihood of their target complying with a request, such as opening a malicious attachment or disclosing information.

Organisations can reduce the likelihood of their domains being used to support fake emails by implementing Sender Policy Framework (SPF) and Domain-based Message Authentication, Reporting and Conformance (DMARC) records in their Domain Name System (DNS) configuration. Using DMARC with DomainKeys Identified Mail (DKIM) to sign emails provides further safety against fake emails. Likewise, organisations can better protect their users against fake emails by ensuring their email systems use and apply SPF, DKIM and DMARC policies on inbound email.

SPF and DMARC records are publicly visible indicators of good cyber hygiene. The public can query a DNS server and see whether an organisation has SPF and/or DMARC protection. DKIM records are attached to outgoing emails and their presence (or lack thereof) is also visible to any external party you email.

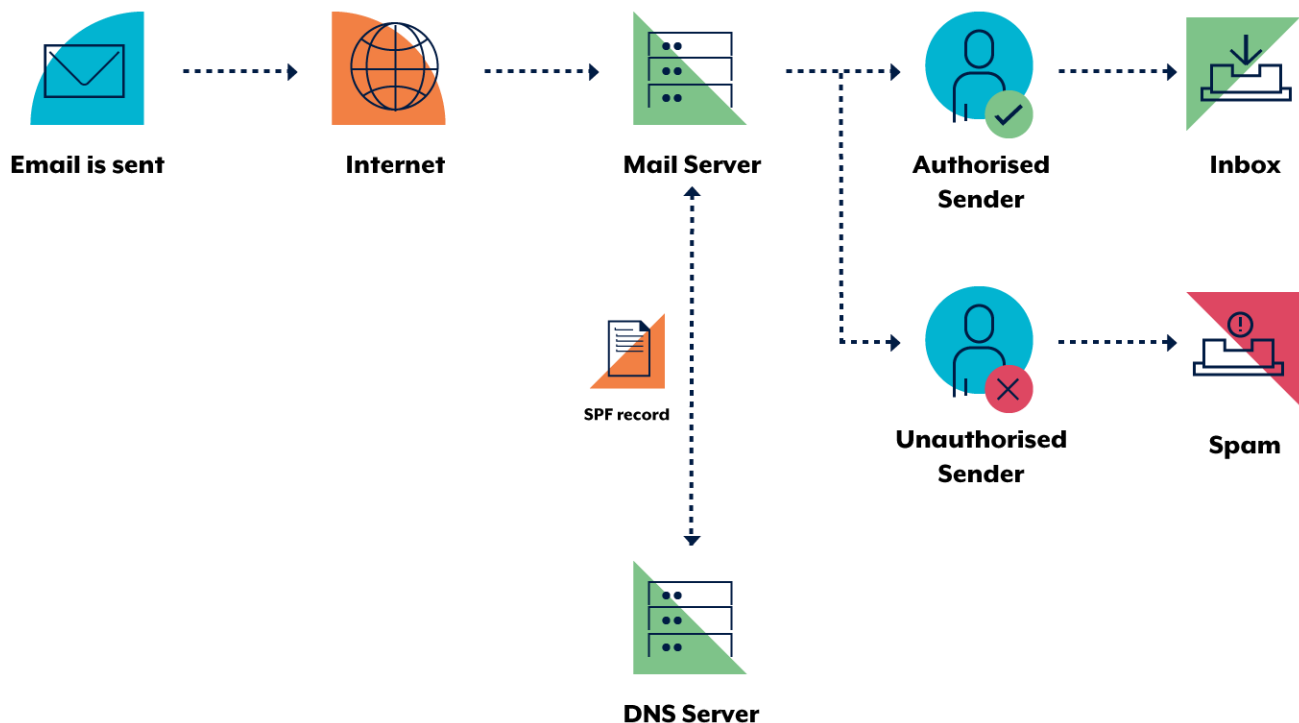
This publication provides information on how SPF, DKIM and DMARC work, as well as advice for security practitioners and mail server operators on how they should configure their systems to prevent their domains from being used as the source of fake emails.

How SPF, DKIM and DMARC work

Sender Policy Framework

SPF is an email verification system designed to detect fake emails. As a sender, a domain owner publishes SPF records in DNS to indicate which mail servers are allowed to send emails for their domains.

When an SPF-enabled mail server receives email, it verifies the sending mail server's identity against the published SPF record. If the sending mail server is not listed as an authorised sender in the SPF record, verification will fail. The following diagram illustrates this process.



SPF 'from' header weakness

SPF has a known weakness. Mail servers applying SPF policies check the RFC5321.Mailfrom header (commonly called the 'envelope from header') while email clients typically display the RFC5322.Mailfrom header (commonly called the 'message/letter from header') to the users as the source of an email.

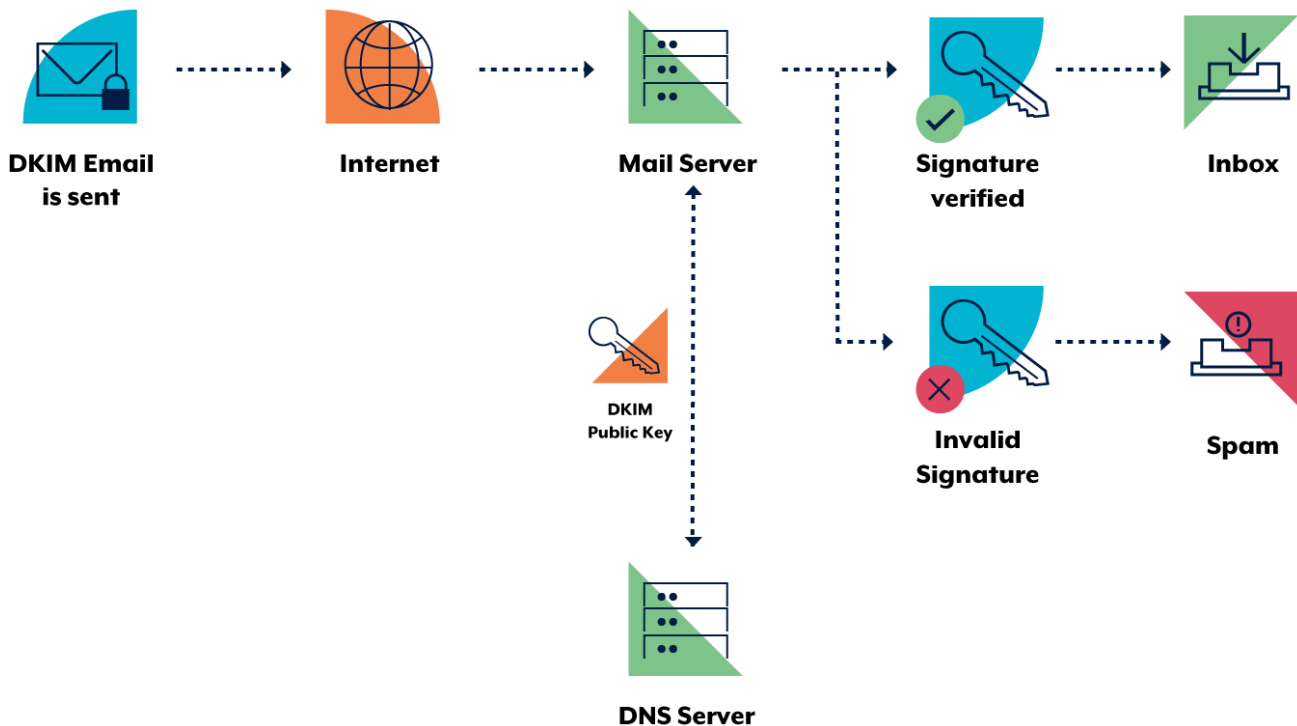
Adversaries are aware of this weakness and use it to bypass SPF checks by using a domain they control in the envelope from header, and the domain they wish to spoof (but don't control) in the message/letter from header.

DMARC addresses this weakness by checking that these two headers align. Consequently, organisations should utilise DMARC in addition to SPF.

DomainKeys Identified Mail

The DKIM standard uses public key cryptography and DNS to allow sending mail servers to sign outgoing emails, and receiving mail servers to verify those signatures. To facilitate this, domain owners generate a public/private key pair. The public key from this pair is then published in DNS and the sending mail server is configured to sign emails using the corresponding private key.

Using the sending organisation's public key (retrieved from DNS), a receiver can verify the digital signature attached to an email. The following diagram illustrates this process.



DKIM 'from' header weakness

DKIM has a known weakness similar to SPF. Mail servers applying DKIM policies check the RFC5321.Mailfrom header (commonly called the 'envelope from header') while email clients typically display the RFC5322.Mailfrom header (commonly called the 'message/letter from header') to the users as the source of an email.

Adversaries are aware of this weakness and can publish DKIM selectors/public keys in a domain they control, using this domain in the envelope from header, and specifying the domain they wish to spoof (but don't control) in the message/letter from header. In practice, adversaries often prefer to use the SPF weakness over the DKIM weakness as the SPF weakness is just as effective and easier to setup.

DMARC addresses this weakness by checking that these two headers align. Consequently, organisations should utilise DMARC in addition to DKIM.

Domain-based Message Authentication, Reporting and Conformance

DMARC enables domain owners to advise recipient mail servers of policy decisions that should be made when handling inbound emails claiming to come from the owner's domain. Specifically, domain owners can request that recipients:

- allow, quarantine or reject emails that fail SPF and DKIM verification
- collect statistics and notify the domain owner of emails falsely claiming to be from their domain
- notify the domain owner how many emails are passing and failing email authentication checks
- send the domain owner data extracted from a failed email, such as header information and web addresses from the email body.

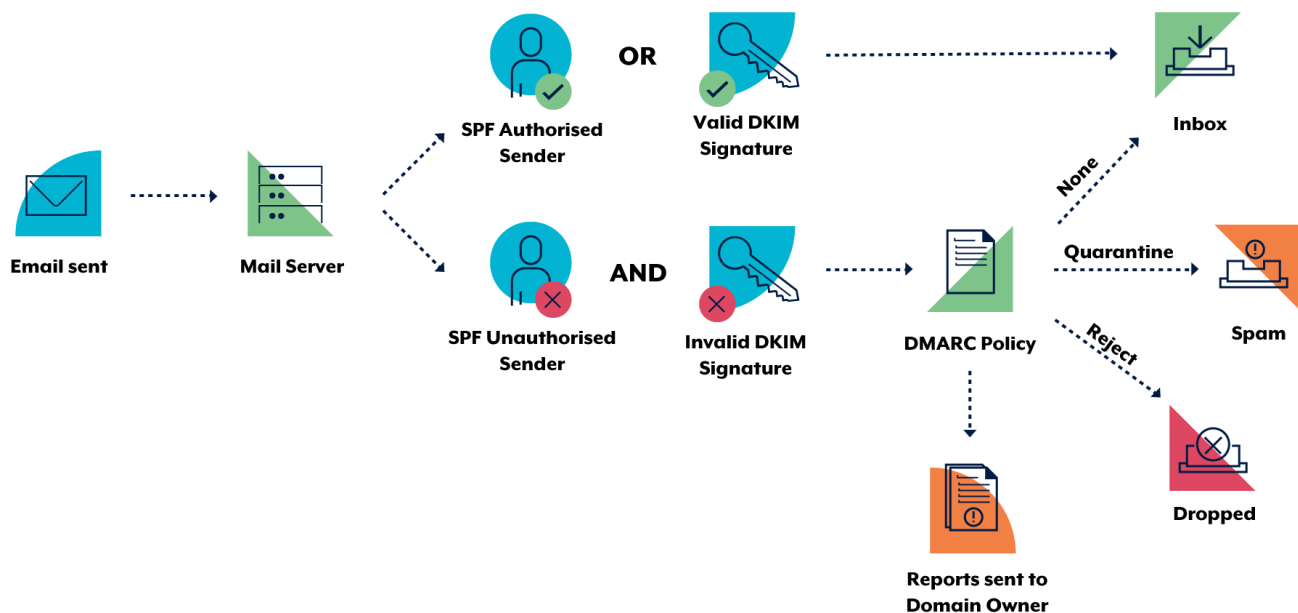
Notifications and statistics resulting from DMARC are sent as aggregate reports and forensic reports:

- aggregate reports provide regular high-level information about emails, such as which Internet Protocol (IP) address they come from and if they failed SPF and DKIM verification

- forensic reports are sent in real time and provide detailed information on why a particular email failed verification, along with content such as email headers, attachments and web addresses in the body of the email.

Like SPF and DKIM, DMARC is enabled when the domain owner publishes information in their DNS record. When a recipient mail server receives an email, it queries the DMARC record of the domain the email claims to come from using DNS.

DMARC relies on SPF and DKIM to be effective. The following diagram illustrates this process.



It should be noted that receiving mail servers need to be configured to honour SPF, DKIM and DMARC for the protections to be in effect.

Addressing SPF and DKIM weaknesses

As noted in SPF and DKIM, DMARC addresses the weakness in SPF and DKIM by introducing an alignment check. The alignment check verifies that the RFC5321.Mailfrom header (commonly known as the 'envelope from header') and the RFC5322.Mailfrom header (commonly known as the 'message/letter from header') match.

Insights and strategies for effective protection

SPF, DKIM and DMARC provide a toolbox of mechanisms that can be used to protect an organisation's domain against being impersonated with email. However, the exact approach to deploying these mechanisms can vary based on the particular mechanism and your environment. This section discusses key differences between the mechanisms and provides advice on implementation strategies.

SPF vs DKIM

SPF and DKIM both provide mechanisms to allow a recipient to determine if an email is legitimate. However, this is achieved in different ways, and understanding this difference assists in understanding where each should be used.

SPF identifies an email as legitimate because it is received from an authorised source. This is a point in time assessment which can only be made by the first mail server that receives the email after being relayed over the internet. In

contrast, DKIM identifies an email as legitimate because the email itself is authenticated by a digital signature. The recipient can rely on the digital signature in the email, irrespective of where they received the email from.

DKIM authentication is arguably more valuable because the authentication mechanism is:

- portable, and travels with the email irrespective of the mail servers it passes through (consequently, DKIM can be used for more sophisticated email authentication flows)
- stronger, being based on public key cryptography (thus, for example, resistant to other attack techniques, such as person in the middle).

However, SPF is simpler to implement and more widely supported (due to simplicity and being a standard for longer).

When considering how to authenticate email flows, organisations may find that some circumstances are best served by DKIM, while others are adequately served by SPF.

DMARC is critical – implement it now irrespective of your existing controls

DMARC provides three features which simplify the implementation of other controls to combat fake emails:

- **DMARC records are hierarchical:** Subdomains will inherit your DMARC policy and its protection. In contrast, both DKIM and SPF require per-domain configuration.
- **DMARC provides visibility:** Publishing a DMARC record, with an associated reporting mechanism, will allow your organisation to see how other mail servers see the email flow from your organisation's domains. This will make it easier to implement email authentication controls such as SPF and DKIM by enabling you to identify email flows and configuration errors.
- **DMARC allows a phased approach to implementing email authentication methods:** DMARC allows a domain owner to publish email authentication mechanisms (through SPF and DKIM) but then set their enforcement policy at 'none', which will, for DMARC-aware mail servers, override any other result, such as an SPF fail due to not having an SPF record.

Irrespective of where your organisation is in terms of implementing anti-spoofing controls, implement DMARC immediately, even if only in a monitoring mode ('p=none') configuration.

DMARC also addresses weaknesses in alignment checks as identified earlier in this publication.

SPF still matters

While SPF has a significant security weakness, it is still a useful indicator that domain owners should use. It is simple to understand and relatively simple to implement.

Across the internet there is evidence that mail servers often receive little attention when it comes to maintenance and implementation of new features. Many mail servers only have SPF policy agents, and may not look for, and interpret, DKIM markings and DMARC records.

Adoption of DKIM and DMARC by recipient mail servers will improve over time, but for the moment, many mail servers only know how to check SPF and domain owners should publish SPF records to identify authorised mail servers for all domains. More information on this is in the *General strategies* section below.

Anticipating recipient mail server behaviour

As noted above, the configuration and support of policy agents on recipient mail servers can vary. When considering how SPF, DKIM and DMARC records are used to authenticate email by recipient mail servers, the following list of standards-compliant configurations for recipient mail servers (the mail servers your organization is sending email to) may be useful:

- **Recipient mail server does not support SPF, DKIM or DMARC:** Your markings and policies will have no effect.
- **Recipient mail server supports only SPF:** Only SPF policies will be followed. If an email flow relies on DKIM for authentication, and you have a hard fail SPF, mail servers that are only SPF aware may drop your email. Likewise, if you are relying on a DMARC p=none configuration to allow email to flow in the presence of a hard fail SPF, mail servers which are only SPF aware will likely drop your email.
- **Recipient mail server supports only DKIM:** While this configuration is possible, it is unlikely a recipient mail server will drop email for not having a DKIM signature, as this would result in many emails being discarded. Emails may be dropped for having an invalid DKIM signature though.
- **Recipient mail server supports SPF and DKIM:** A mail server configured thus would enforce SPF and DKIM rules potentially in isolation. For example, if the sending mail server is not specified in a hard fail SPF record the mail server may refuse delivery of the email before it checks for a valid DKIM signature.
- **Recipient mail server supports DMARC, SPF and DKIM:** By standard, to support DMARC, recipient mail servers must also support SPF and DKIM. The recipient mail server will apply your DMARC policy by evaluating both SPF and DKIM checks and allowing email that passes either test to flow.

Finally, it is worth noting that in practice many recipient mail servers may consume the information you publish via SPF, DKIM and DMARC, but not necessarily act strictly in accordance with your policy instructions. Instead, recipient mail servers may make their own decision, or use SPF, DKIM and/or DMARC information to inform a more sophisticated decision making engine, such as an anti-spam filter.

General strategies

These general strategies are presented as potential end states that your organization can target to assist other recipient mail servers in identifying fake email claiming to come from your domains.

Best practice

The best practice strategy seeks to provide the best possible protection to the widest possible range of recipient mail servers. It allows, to the extent possible, for recipient mail servers, operated by others, which do not stay up to date with standards:

- A DMARC record at your organisation's root domain(s) is published which either quarantines or rejects 100 percent of email that fails SPF and DKIM checks. The DMARC policy is configured to apply to subdomains too. The DMARC record also publishes a location for reports, and these reports are captured into a system and reviewed regularly by your email or cyber security teams. Lower level DMARC records are published on a subdomain by subdomain basis where a different policy is required.
- All of your domains, subdomain and hostnames have a hard fail SPF record which specifies authorised mail relays (including explicitly specifying no mail relays for domains that do not send email).
- DKIM is also used on all outbound email flows that leave the organization, but particularly in situations where the complexity of the email flow, or a desire not to delegate excessive email sending authority to a third-party mail relay, means that SPF is inappropriate.

Good practice

The good practice strategy seeks to provide recipient mail servers with enough information to make good policy decisions provided they have stayed up to date with standards by implementing a DMARC policy agent to evaluate DMARC, SPF and DKIM policy published by senders:

- A DMARC record at your organisation's root domain(s) is published which either quarantines or rejects 100 percent of email that fails SPF and DKIM checks and is configured to apply to subdomains too. The DMARC record

also publishes a location for reports, and these reports are captured into a system and reviewed regularly your email or cyber security teams. Lower level DMARC records are published on a subdomain by subdomain basis where a different policy is required.

- Outbound mail servers for domains, subdomains and hostnames which send email are explicitly specified with an SPF record which includes a hard fail.
- DKIM is used on outbound email flows where the complexity of the email flow, or a desire not to delegate excessive email sending authority to a third-party relay, means that SPF is inappropriate.
- Domains, subdomains and hostnames which are not explicitly authorized through an SPF record, or a DKIM selector, are implicitly protected by the overarching DMARC record.

How to implement SPF, DKIM and DMARC

Sender Policy Framework

Identify outgoing mail servers

Identify your organisation's authorised mail servers, including your primary and backup outgoing mail servers. You may also need to include your web servers if they send emails directly. Also identify other entities who send emails on behalf of your organisation and use your domain as the email source. For example, advertising or recruitment firms and newsletters.

Construct your SPF record

SPF records are specified as text (TXT) records in DNS. An example of an SPF record might be `v=spf1 a mx a:<domain/host> ip4:<ipaddress> -all` where:

- `v=spf1` defines the version of SPF being used
- `a`, `mx`, `a:<domain/host>` and `ip4:<ipaddress>` are examples of how to specify which mail server is authorised to send email
- `-all` specifies a hard fail advising receivers to drop emails sent from your domain if the sending mail server is not authorised. Note, for DMARC-enabled recipient mail servers they will apply the policy published in your DMARC record. However, if your SPF record is configured as a soft fail (e.g. `+all`) then the SPF test will never fail, and consequently all email will pass DMARC checks.

It is important to note that you must set a separate record for each subdomain as subdomains do not inherit the SPF record of their top-level domain.

To avoid creating a unique record for each subdomain, you can redirect the record lookup to another SPF record (e.g. the top-level domain record or a special record for subdomains would be the simplest solution).

Identify domains that do not send email

If using the best practice strategy, organisations should explicitly state if a domain does not send emails by specifying `v=spf1 -all` in the SPF record for that domain. This advises receiving mail servers that there are no authorised sending mail servers for the specified domain, and hence, any emails claiming to be from that domain should be rejected.

Warn your users

Ensure users are told of the new SPF email policy so they can report any implementation issues which may arise. It should be noted that once SPF is implemented, emails sent from non-authorised mail servers, such as those outside the

corporate network, may no longer reach their intended destinations. If users are required to send emails while away from the corporate network, then provisions should be made for authenticated remote access to an authorised mail server specified in the SPF entry.

Test your SPF record

Testing your SPF record will ensure that emails are processed correctly. Tools such as [MX Lookup](#) can help assess the correctness of SPF records before finalising them. A hard fail policy (i.e. `-all`) is the preferred approach for SPF. However, while testing, you may wish to use a soft fail policy. This is done by specifying `~all` instead of `-all` in the SPF record. By combining this with DMARC reporting, you can be made aware of potential issues before implementing a hard fail policy.

It should be noted that older methods of forwarding emails (e.g. `.forward` files on UNIX-based systems) can fail SPF checks if not treated correctly. If this describes your operating arrangement, you may wish to convert forwarded emails to re-mailed emails using Sender Rewriting Scheme (SRS). There are also emerging standards which are seeking to address these kinds of issues such as Authenticated Received Chain (ARC). See the *Other standards for combating fake email* section below for more information.

Deploy your SPF record

When you have an SPF record you intend to deploy, ensure the Time to Live (TTL) is set low (e.g. 5 minutes). Setting the TTL low will allow you to correct or rollback your SPF record quickly in the case of an issue.

Monitor the success of the SPF record after deployment

When you implement a new SPF record it will take some time before mail servers on the internet recognise it. If you have an existing SPF record this time will be defined by the TTL associated with this record. If you don't have an existing SPF record, you should review your domain's 'start of authority' to determine the negative cache TTL value, sometimes referred to as 'minimum TTL'. This number, in seconds, is the maximum time it should take mail servers on the internet to detect your new SPF record – a typical default value is 7,200 seconds (i.e. two hours). After you have implemented an SPF record, you should monitor your mail server up to the TTL time and confirm that email delivery is continuing normally. If your SPF configuration is incorrect, recipient mail servers will begin to reject your email.

Incorporate SPF into the change management process and associated procedures

SPF records will need to be updated when new email sending mail servers are deployed, DNS entries are added, and DNS entries or IP addresses of sending mail servers change. In such cases, respectively, new mail servers will need to be added to the SPF record of the associated domains, new SPF records will be required (including wildcard SPF records), and the SPF record of all associated domains will need to be reviewed. It is important that your procedures account for these changes as part of your organisation's change management processes.

Additional resources

For additional information on SPF, see the [SPF standard](#) and [its update](#).

For additional information on how to setup SPF, see Microsoft's [Set up SPF to help prevent spoofing](#) publication and Google's [Help prevent email spoofing with SPF](#) publication.

For additional information on SRS, see Microsoft's [Sender Rewriting Scheme \(SRS\) in Office 365](#) publication.

DomainKeys Identified Mail

Decide what sections of your email you want to sign

The more sections of an email you sign, the more protected you are from adversaries sending fake emails that appear to come from you. At least the body and the following headers of emails should be signed:

- | | | |
|-----------|----------------|-----------------------------|
| ▪ to | ▪ sender | ▪ references |
| ▪ cc | ▪ reply-to | ▪ in-reply-to |
| ▪ date | ▪ references | ▪ return-path |
| ▪ from | ▪ mime-version | ▪ content-disposition |
| ▪ subject | ▪ content-type | ▪ content-transfer-encoding |

If supported by your email software, it is also important to sign headers one more time than the number of times they occur. A detailed analysis of why this is important can be found in the [Breaking DKIM – on Purpose and by Chance](#) publication.

Decide when to sign your outgoing email

DKIM verification may fail if signed sections of emails are changed during delivery. This can happen without the interference of adversaries. For example, if you sign an email but then attach a legal disclaimer, you will likely invalidate the DKIM signature as the content of the email has now changed.

To ensure the DKIM signatures on outgoing emails are not invalidated, you need to understand your email infrastructure and when the content of emails is changed.

A good way to begin a DKIM rollout would be to sign emails as they leave the email gateway, and if you run into issues, see how you can modify this approach to accommodate your email setup. In general, signing emails at the last stage before they leave the infrastructure under your control reduces the chance of unintended changes.

Generate your public and private keys for DKIM

You will need to find a tool to generate a public/private key pair. The tool of choice will depend on your organisation, its key management plan and operating system. In the absence of organisation-specific tools, there are tools such as *ssh-keygen* for Linux, Microsoft Windows and macOS. In the terminal for your operating system run *ssh-keygen -b 2048 -t rsa -f filename* to generate a public/private RSA key pair. The key length will be 2048 bits and the public/private keys will be called *filename.pub* and *filename* respectively.

After creating the public/private key pair, make sure you protect your private key in accordance with your organisation's key management plan. If your private key becomes public, any adversary will be able to sign emails as yourself.

You may want to generate multiple public/private key pairs if you want different parts of your organisation to have the ability to independently sign emails.

Publish your public key in your DNS record

Public keys for DKIM are published in DNS TXT records. You will need to include a different TXT record for each private/public key pair you want to use by using a selector. You will then need to configure your mail server to specify the correct selector when sending emails.

Make sure you put the public keys at `<selector>._domainkey.yourorganisation.com.au` where *selector* must be different for every public/private key pair. For example you might have `brisbane._domainkey.yourorganisation.com.au` and `melbourne._domainkey.yourorganisation.com.au` for different offices in your organisation.

The content of the record itself will be similar to `v=DKIM1; k=rsa; p=<your public key>`.

Warn your users

As for SPF, notify users of the change so they are aware and can communicate any unexpected change in email behaviour.

Develop and test your mail server configuration

Using a test environment, identify how your mail server implements DKIM and test your configuration as thoroughly as possible prior to implementation.

Configure your mail server to attach DKIM records to its headers

A phased implementation approach is recommended if your infrastructure allows it. Immediately after deployment, test that emails are being signed correctly by sending emails to external email accounts and reviewing email headers.

Monitor after deployment

After DKIM has been added, and if you have already configured DMARC, you can use DMARC email reports to verify how many emails are failing DKIM checks.

Additional resources

For additional information on DKIM, see the [DKIM website](#), the [DKIM standard](#) and its [first](#) and [second](#) update.

For additional information on how to setup DKIM, see Microsoft's [Use DKIM to validate outbound email sent from your custom domain](#) publication and Google's [Increase security for outgoing email with DKIM](#) publication.

Domain-based Message Authentication, Reporting and Conformance

DMARC is implemented by publishing a policy as a TXT record in DNS and is hierarchical (e.g. a policy published for `organisation.com.au` will apply to `sub.domain.organisation.com.au` unless a different policy is explicitly defined for the subdomain). This is useful as organisations may be able to specify a smaller number of high-level DMARC records for wider coverage. Care should be taken to configure explicit subdomain DMARC records where you do not want the subdomains to inherit the top-level domain's DMARC record.

When implementing DMARC, it is advisable to first implement it in monitoring mode, followed by quarantine mode and finally reject mode as the implementation maturity level increases.

Policy inheritance with DMARC

DMARC policies are often described as hierarchical. However, the hierarchy mechanism may not work as assumed.

To avoid excessive DNS lookups, DMARC policy agents on recipient mail servers will perform a maximum of two lookups to determine if there is an effective DMARC policy. The first lookup is performed on the domain the email claims to be from. The second lookup is then performed on the organisation domain. The organisation domain, refers to the highest order non-public suffix of the email's from domain (that is, the end of the domain that includes the public suffix, and the next part of the domain name). This is mostly clearly explained with an example:

- An email is received from `<someuser>@a.b.c.yourorganisation.com.au`.

- A DMARC policy agent first checks for a TXT record at `_dmarc.a.b.c.yourorganisation.com.au`. If this record exists, and is a valid DMARC policy, then it will be the applied DMARC policy.
- If no records exists at this location, the DMARC policy agent:
 - uses a public suffix list to identify the public suffix as the end of the domain name (e.g. `com.au`)
 - appends the first domain name part that is not in the public suffix part of the domain (e.g. `yourorganisation`), to arrive at `yourorganisation.com.au` – the highest order non-public suffix of the email's from address
 - checks for a TXT record at `_dmarc.yourorganisation.com.au` (note: `_dmarc.b.c.yourorganisation.com.au` and `_dmarc.c.yourorganisation.com.au` are not checked by the policy agent).

Most organisations implementing DMARC will likely want to rely on DMARC inheritance so should consult a public suffix list (e.g. <https://publicsuffix.org/>) to determine the correct level to apply a DMARC record at to give policy coverage.

Entities of Australian state and territory jurisdictions may find that their state or territory has removed their jurisdiction from the public suffix list, and hence a DMARC record published at `<entity>.<jurisdiction>.gov.au` may not be inherited by subdomains. Contact your State or Territory Government's Chief Information Security Officer's office if this applies to you and you require further information.

Monitoring mode

You can start your DMARC implementation with a simple monitoring policy for your domain which requests that DMARC-capable mail servers send you statistics about emails they see using your domain. You can do this even before you've implemented SPF or DKIM in your infrastructure, though until it is in place you won't be able to move beyond this step.

As you introduce SPF and DKIM, reports will provide the number and IP addresses of emails that pass these checks, and those that don't. You can then see how much of your legitimate traffic is passing SPF and DKIM checks, and troubleshoot any problems with your SPF or DKIM configuration.

You will also begin to see how many fake emails are being sent, and from where. At this point you should put in place a capability to review reports from recipient mail servers sent in accordance with DMARC configuration to identify malicious activity.

An example of such a DMARC record is `v=DMARC1; pct=100; p=none; sp=none; ruf=mailto:authfail@yourorganisation.com.au; rua=mailto:aggrep@yourorganisation.com.au` where:

- `v=DMARC1` defines the version of DMARC being used
- `pct=100` specifies the percentage of emails subjected to filtering
- `p=none` specifies the policy for your organisation domain
- `sp=none` specifies the policy for all your organisation subdomains
- `ruf=mailto:authfail@yourorganisation.com.au` states the email address to which forensic reports should be sent
- `rua=mailto:aggrep@yourorganisation.com.au` states the email address to which aggregate reports should be sent.

Note, if you plan to send your DMARC reports to a domain that is not the domain for which the DNS record exists, you will need to include a special record in the receiving domains DNS record.

Quarantine mode

When you believe that all, or most of, your email traffic is protected by SPF or DKIM, you can implement a quarantine policy. This will result in DMARC-enabled mail servers marking emails from your domain that fail verification as spam.

Even if you request that only a small percentage of your email traffic have a quarantine policy applied, you will still get the full statistical reports that show what is happening with your emails. Eventually, as implementation problems are resolved, you can gradually increase your implementation to 100 percent.

Reject mode

Following testing using a quarantine policy, implement a reject policy. This will result in DMARC-enabled mail servers rejecting emails that fail SPF and DKIM verification. Again, you can request that this policy only be applied to a small percentage of your emails (with the remaining percentage being quarantined) and monitor the results through reports. The same gradual increase to 100 percent can be implemented based on reports and feedback from users.

Additional resources

For additional information on DMARC, see the [DMARC website](#) and the [DMARC standard](#). The DMARC website also identifies [common problems with DMARC records](#) and includes [answers to frequently asked questions](#).

For additional information and assistance in [generating a DMARC record](#), see the Global Cyber Alliance website.

For additional information on how to setup DMARC, see Microsoft's [Use DMARC to validate email](#) publication and Google's [Increased security for forged spam with DMARC](#) publication.

Other standards to combat fake email

Authenticated Received Chain

ARC is an emerging standard that aims to address the problem where intermediate mail servers may need to modify emails to deliver them (such as email list software) or where downstream mail servers are unable to perform checks such as SPF.

ARC proposes a signed chain of custody which allows intermediate mail servers to attest to the validity of DKIM and SPF checks at each step. Downstream mail servers can choose to rely on these attestations, even if DMARC checks would otherwise fail. More information can be found in DMARC.org's [Authenticated Received Chain Overview](#) presentation.

Additional resources

For additional information on ARC, see the [ARC standard](#).

Further information

The [Information Security Manual](#) is a cyber security framework that organisations can apply to protect their systems and data from cyber threats. The advice in the [Strategies to Mitigate Cyber Security Incidents](#), along with its [Essential Eight](#), complements this framework.

Effective email filtering is an important control to reduce the incidence of malicious email entering an organisation's environment. Effective strategies are discussed in the ACSC's [Malicious Email Mitigation Strategies](#) publication.

The secure use of marketing and email filtering service providers is discussed in the ACSC's [Marketing and Filtering Email Service Providers](#) publication.

TLS encryption and Mail Transport Agent Strict Transport Security (MTA-STX) can be used to protect email confidentiality and integrity against person-in-the-middle attacks. Further information can be found in the ACSC's [Implementing Certificates, TLS, HTTPS and Opportunistic TLS](#) publication.

Contact details

If you have any questions regarding this guidance you can [write to us](#) or call us on 1300 CYBER1 (1300 292 371).

Annex A: Detailed examples and edge cases

Records for SPF, DKIM and DMARC are published as follows.

Protocol	Record Type	Record Name
SPF	TXT	<i><domain name></i>
DKIM	TXT	<i><selector>._domainkey.<domain name></i>
DMARC	TXT	<i>_dmarc.<domain name></i>

Analysis of common SPF syntax

Using an example of *v=spf1 a mx a:<domain/host> ip4:<ipaddress> -all*:

- *mx* specifies that recipients should accept the mail servers identified by the mail exchanger (MX) record for your domain/host
- *a* specifies that recipients should accept the mail server identified by the A record for your domain/host
- *a:<domain/host>*, *ip4:<ipaddress>* are examples of specifying other hosts that can send email, either by an A record or an IP address
- *-all* specifies a hard fail directing receivers to drop emails sent from your domain if the sending mail server is not authorised.

Note, the *+all* and *?all* flags should never be used as they could allow any mail server to send emails from your organisation. Further, be careful when using *include* or *redirect*. If you include or redirect a domain whose SPF record's syntax is incorrect, mail servers will return an error when validating emails against your SPF record.

Redirecting subdomains SPF records to save time

Depending on the general strategy adopted, an individual SPF record should be set for each domain and subdomain. However, to avoid creating a unique SPF record for each subdomain, you can redirect them to your top-level domain. For example, you can set all subdomain records to be *v=spf1 redirect=yourorganisation.com.au*. This configuration will mean that all subdomains will use the SPF record of their parent domain, *yourorganisation.com.au*. This is effective if, for example, all emails from subdomains pass through a centralised email relay.

Alternatively, you could redirect to a special subdomain if you want all subdomains to have a specific SPF record that is different from that of the top-level domain. For example, you could set all subdomain records as *v=spf1 redirect=spf.yourorganisation.com.au* which would mean that all subdomains use the SPF record at *spf.yourorganisation.com.au* and you only need to maintain one actual SPF record for all subdomains.

Examples of what your DNS records may look like for different scenarios

Domains that don't send email – best practice

yourorganisation.com.au. TXT "v=spf1 -all"

_dmarc.yourorganisation.com.au. TXT "v=DMARC1; p=reject; ruf=mailto:authfail@yourorganisation.com.au; rua=mailto:aggrep@yourorganisation.com.au"

Even though this domain doesn't send email, the SPF record is needed to prevent adversaries sending emails pretending to be you and the DMARC record is needed so you are notified when this happens.

The email accounts specified in the DMARC record, *authfail@yourorganisation.com.au* and *aggrep@yourorganisation.com.au*, are examples of email addresses you could use to receive DMARC reports.

Domains which send email

yourorganisation.com.au. TXT "v=spf1 a mx a:domain1.com.au ip4:1.2.3.4 -all"

selector1._domainkey.yourorganisation.com.au. TXT "v=DKIM1; k=rsa; p=MIGfMA0GCSqGSIb3DQEBAQUAA4GNADCBiQKBgQC3QEKyU1fSma0axspqYK5iAj+54lsAg4qRRcnpKK68hawSd8zpsDz77ntGCR0X2mHVvkf0WEOlqaspaG/A5IGxieiWer+wBX8lW2tE4NHTE0PLhHqL0uD2sif2pKoPR3Wr6n/rbiihGYClzvuY4/U5GigNUGls/QUbCPRyzho30wIDAQAB"

_dmarc.yourorganisation.com.au. TXT "v=DMARC1; p=reject; ruf=mailto:authfail@yourorganisation.com.au; rua=mailto:aggrep@yourorganisation.com.au"

Subdomains which don't send email

subdomain.yourorganisation.com.au. IN TXT "v=spf1 -all"

Subdomains which send email

subdomain.yourorganisation.com.au. IN TXT "v=spf1 redirect=yourorganisation.com.au"

selector1._domainkey.subdomain.yourorganisation.com.au. TXT "v=DKIM1; k=rsa; p=MIGfMA0GCSqGSIb3DQEBAQUAA4GNADCBiQKBgQC3QEKyU1fSma0axspqYK5iAj+54lsAg4qRRcnpKK68hawSd8zpsDz77ntGCR0X2mHVvkf0WEOlqaspaG/A5IGxieiWer+wBX8lW2tE4NHTE0PLhHqL0uD2sif2pKoPR3Wr6n/rbiihGYClzvuY4/U5GigNUGls/QUbCPRyzho30wIDAQAB"

In these examples, there are no DMARC records published. For the sake of this example we are assuming the parent domain has published a DMARC policy which we are accepting through inheritance.

Example of DKIM signature header in an email

DKIM-Signature a=rsa-sha256; d=yourorganisation.com.au; s=selector; c=relaxed/simple; q=dns/txt; t=1117574938; x=1118006938; h=from:from:to:to:cc:cc:subject:subject:date:date:sender:sender:content-type:content-type:content-transfer-encoding:content-transfer-encoding:content-disposition:content-disposition:mime-version:mime-version:reply-to:reply-to:in-reply-to:in-reply-to:references:references; bh=MTIzNDU2Nzg5MDEyMzQ1Njc4OTAxMjMONTY3ODkwMTI=; b=dzdVvYOfAKCdLXdJOc9G2q8LoXSIEniSbav+yuU4zGeeruD00lszZVoG4ZHRNiYzR

Further examples, as well as advice on what to do with domains you aren't currently using, can be found in the Messaging, Malware and Mobile Anti-Abuse Working Group's [M³AAWG Protecting Parked Domains Best Common Practices](#) publication.

Edge cases

Hosting services on a shared IP address

Hosting your mail server on a shared IP address may allow others, who have a service on the same IP address, to send emails as your domain. This is possible because SPF authentication resolves to IP addresses. For example, your mail server is *acorp.com* with an IP address of 1.1.1.1 and you authorise emails to be sent from this domain with an SPF record of *v=spf1 mx -all* (assuming that your mx record resolves to 1.1.1.1). If *bcorp.com* also hosts a service at 1.1.1.1, then *bcorp.com* can send emails claiming to be *acorp.com* as they have the same IP address and SPF resolves and

authenticates on IP addresses. Additional advice on this is available in the ACSC's [Marketing and Filtering Email Service Providers](#) publication.

Publishing SPF records for HELO names used by your mail servers

The SPF standard recommends that the *HELO* name specified by the connecting mail server, and the *MAIL FROM* field of the envelope, undergo an SPF authorisation check. If you are using a different domain for the *HELO* name, and the *MAIL FROM* field of the SMTP connection, you will need to construct SPF records to authorise both the mail server you are using in the *HELO* field and the address you are using in the *MAIL FROM* field. Below are two example scenarios.

Emails are being sent from *recruitment@acorp.com.au* but being relayed by your email list service provider's mail server *mail.bcorp.com.au*

- To authorise the *MAIL FROM*: *acorp.com.au*. IN TXT "v=spf1 mx a:mail.bcorp.com.au -all".
- To authorise the sending server: *mail.bcorp.com.au*. IN TXT "v=spf1 a -all".

Note, if you publish a nil SPF record for the hostname of your email relay some SPF filters will reject it.

Emails are being sent from *somebody@acorp.com.au* via *mail.bcorp.com.au*

In this case, publishing the below SPF entries, will result in some SPF filters rejecting email due to the *HELO* hostname lookup:

- *acorp.com.au*. IN TXT "v=spf1 mx a:mail.bcorp.com.au -all"
- *mail.bcorp.com.au*. IN TXT "v=spf1 -all".

A correct configuration could be achieved by using *mail.bcorp.com.au*. IN TXT "v=spf1 a -all".

Treatment of CNAME records

Canonical Name (CNAME) records are a common DNS technique used to enable transparent redirection to an alternate domain name. They are frequently used by organisations to redirect requests for services to a third-party service provider (e.g. a public cloud). While this is an effective technique, organisations need to be aware that a CNAME record effectively delegates all DNS calls for the target domain to the domain specified in the answer section. For example, if you have *yourorganisation.com.au* CNAME *serviceprovider.com.au* then requests for the SPF, DKIM and DMARC records for *yourorganisation.com.au* will be answered as if directed to *serviceprovider.com.au*. Hence, if a CNAME record is used in your DNS record to redirect requests to an alternate domain, then it will not be possible for you to specify SPF, DKIM and DMARC records. This is a risk that needs to be accepted when deciding to delegate DNS records in this manner. Appropriate SPF, DKIM and DMARC records can be discussed with your service provider, and can be specified in contracts when engaging with such services.

Sending DMARC reports to a different domain

If your domain is *yourorganisation.com.au*, and you want to send your reports to *report.com.au*, then the recipient domain (*report.com.au*) needs to have a TXT DNS record at *yourorganisation.com.au._report._dmarc.report.com* which has content *v=DMARC1*.

If you are sending all your reports to the same domain, you may want to implement a wildcard DMARC record so you can receive reports from anyone who wants to send them to you. A wildcard record would be **._report._dmarc.report.com* which specifies *v=DMARC1*. However, doing this does allow your domain to receive reports you have not explicitly authorised. This could be used by an adversary to cause you inconvenience.