



July 2016

Malicious Email Mitigation Strategies

Introduction

1. Socially engineered emails containing malicious attachments and embedded links have been observed by the Australian Signals Directorate (ASD) being used in targeted cyber intrusions against organisations.
2. This document has been developed by ASD in collaboration with local and international partners to provide mitigation strategies for the security risk posed by malicious emails. It should be read in conjunction with the advice on email security and content filtering contained in the *Australian Government Information Security Manual (ISM)*.
3. Not every mitigation strategy within this document will be suitable for all organisations. Organisations should consider their unique business requirements and risk environment when deciding which mitigation strategies to implement. Furthermore, before any mitigation strategy is implemented, comprehensive testing should be undertaken to minimise any unintended disruptions to the organisation's business.

Intended audience

4. This document is intended for use by Information Technology Security Advisors and system administrators.

Vocabulary

5. This document uses the terms "block" and "quarantine". In the context of this document, "block" refers to preventing an email reaching the user and being removed from the mail server while "quarantine" refers to preventing an email from reaching the user but safely storing it so it can be accessed if required.

Attachment filtering

6. Attachments are a significant security risk associated with emails. Effective attachment filtering reduces the likelihood of malicious content reaching a user's workstation.
7. Mitigation strategies associated with attachment filtering are discussed below. The mitigation strategies are grouped based on their security effectiveness.

Excellent security effectiveness

Convert attachments to another format

8. Converting attachments to another format is a highly effective method of removing malicious content or rendering it ineffective, for example by converting Microsoft Office documents to PDF documents. To decrease the impact to users, but at the expense of an increased security risk, original emails and attachments can be quarantined with a release facility available in case the originals are required for editing purposes.

Whitelist attachments based on file typing

9. File typing inspects the content of a file to determine its file type rather than relying on its extension. File types that have a legitimate business purpose and an acceptable risk profile for organisations should be whitelisted. Whitelisting is recommended as it is more proactive and thorough than blacklisting. It ensures only specified file types can be received, while all others are blocked. File extensions can be changed and therefore a mismatch between a file's type and its stated extension should be treated as suspicious and quarantined.

Block password protected archives and unidentifiable or encrypted attachments

10. Content within password protected archives cannot be trusted since email content filters cannot decrypt and inspect their contents. Any protected archive or otherwise encrypted attachments should be blocked until such time that they can be deemed safe. Unidentifiable content is less of a security risk if effective whitelisting and file typing of attachments is used. Where organisations have corporately approved encrypted email communications, such as S/MIME or PGP, these can be whitelisted to prevent disruption to legitimate business.

Perform automated dynamic analysis of attachments run in a sandbox

11. Dynamic analysis uses behaviour-based detection capabilities instead of relying on the use of signatures, enabling organisations to detect malware that has yet to be identified by vendors. Performing automated dynamic analysis of attachments run in a sandbox may detect suspicious behaviour including network traffic, new or modified files, or changes to the Windows registry.
12. Analysis could be performed in an instrumented sandbox located either in a gateway environment, on a user's workstation or in the cloud subject to concerns about data sensitivity, privacy and security of the communications channel.
13. Organisations should block any attachments detected as malicious, paying particular care to do so before they are accessed by users, by using a product that is regularly updated by the vendor to mitigate evolving evasion techniques that challenge the effectiveness of this mitigation strategy.
14. Further details on this mitigation strategy can be found in ASD's *Strategies to Mitigate Targeted Cyber Intrusions* publication.

Sanitise attachments to remove active or potentially harmful content

15. Active content, such as macros in Microsoft Office files and JavaScript, should be removed from within attachments before being delivered to users. This should include embedded content such as an executable placed inside a Microsoft Word document.

embedded Flash content placed inside a Microsoft Excel spreadsheet and link (LNK) files that call executable content, which should include executable content on the end user's computer such as mshta.exe and rundll32.exe. Organisations should also consider cases where active content creates a high level of suspicion due to limited legitimate use; in these cases the attachment should be blocked.

16. Active content removal products should scan attachments for undesirable active content based on keywords or heuristics, and rewrite those elements rendering them inert. Complete and comprehensive sanitisation of an attachment is a difficult process.

Disable or control macros in Microsoft Office files

17. The Australian Cyber Security Centre (ACSC) and its international partners have observed an increase in the use of macros in Microsoft Office files being used as a malware delivery vector. These macros are written in the Visual Basic for Applications (VBA) programming language, a feature built into Microsoft Office applications. Macros are commonly used for task automation; however, adversaries are also using macros to perform a variety of malicious activities including the download and execution of malware on the host computer.
18. Organisations should configure Microsoft Office to disable all macros by default and only run macros vetted as trustworthy and placed in "trusted location" directories which typical low-privileged users can't write to.
19. Further details on this mitigation strategy can be found in ASD's *Microsoft Office Macro Security* publication.

Good security effectiveness

Controlled inspection of archive files

20. Archive files can be used to bypass poorly configured email content filters. By placing a malicious file in an archive file and sending it to the target, the archive file might bypass content filtering checks. To mitigate this, the contents of archive files should be subjected to the same level of inspection as un-archived attachments. The archive files should be decompressed and the files within inspected. A directory listing of the files inside an archive file is not always an accurate representation of the files actually in the archive file since file attributes, such as file name, could be stored in two places for each file.
21. Archived content should be inspected in a controlled manner to avoid exploits associated with archive files, such as directory traversal and denial of service via recursion. For example, a text file which is 1GB in size and consists only of spaces, could compress to 1MB but consume significant computing resources when it is processed by an email content filter. As another example, a zip file containing 16 zip files, each of which contain 16 zip files, each of which contain 16 zip files etc. to a depth of 5, could cause an email content filter to process over one million files. To mitigate this, quotas and timeout values can be used on CPUs, memory and disks so that decompression is blocked or failed if it takes longer than the specified time or uses excessive computing resources.
22. Archive files decompress starting from the end of the file, stopping when all the files have been extracted. As a result of this an archive file can be appended to the end of a legitimate image file and still be a valid archive from which files can be extracted. In this case, depending on the file type checking, the file could pass file type checks as an image. This behaviour can be exploited by adversaries to avoid controlled inspection of

archive files. To mitigate this, organisations should attempt to decompress all attachments, with all decompressed files submitted to the security controls for attachments and the original attachment blocked if any decompressed files fail.

Average security effectiveness

Whitelist attachments based on file extension

23. Allowing attachments based on file extension is less robust than file typing as the extension can be trivially changed to disguise the true nature of the file, for example by renaming readme.exe to readme.doc. Only file extensions with a legitimate business purpose should be whitelisted.

Minimal security effectiveness

Blacklist attachments based on file typing

24. Blacklisting attachments based on file typing is less proactive and thorough than whitelisting attachments based on file typing or file extension, and the overhead of maintaining a list of all known bad file types is far greater than maintaining a list of all known good file types.

Scan attachments using antivirus software

25. Attachments should be scanned using vendor supported antivirus software with up-to-date signatures, reputation ratings and other heuristic detection capabilities. To maximise the chance of detecting malicious content, antivirus software from a different vendor to that used for user workstations should be used.

Blacklist attachments based on file extension

26. Blacklisting attachments based on file extension is less proactive and thorough than whitelisting attachments based on file typing or file extension. Blacklisting attachments based on file extension is less robust than file typing as the extension can be trivially changed to disguise the true nature of the file, for example by renaming readme.exe to readme.doc.

Email body filtering

27. Email content filtering performed on the body of an email helps provide a defence-in-depth approach to email content filtering. The possible attack surface presented by the body of an email is less than attachments; however, content in an email body can still introduce malicious content to a network.

28. Mitigation strategies associated with filtering the body of an email are discussed below. The mitigation strategies are grouped by their security effectiveness.

Good security effectiveness

Replace active web addresses in an email's body with non-active versions

29. An active web address allows users to click on a hyperlink in the body of an email and be taken to a specified website. Active web addresses can appear to be safe but can actually direct users to an unintended website. Hovering over the address may reveal the actual website, as shown here:

Classificati <http://www.malicious.content.com/>
 Click to follow link
<http://www.safe.com/>

30. Active web addresses should be replaced with non-active versions so that users must copy and paste the web address into their browser – hopefully in doing so noticing it is a malicious web address.

Average security effectiveness

Remove active content in an email's body

31. Emails with active content such as VBScript or JavaScript pose a security risk if the email client, or web browser in organisations where webmail is utilised, is capable of running the active content. Email bodies containing active content should be sanitised or the email blocked to minimise the security risk. When sanitising an email body the active content should be rewritten in the body to render it inert.

Sender verification

32. Being able to verify the authenticity and integrity of an email can stop organisations from receiving some forms of malicious emails. Particular care should be taken when implementing sender verification because of the potential to impact legitimate email traffic.
33. Mitigation strategies for sender verification are discussed below. The mitigation strategies are grouped by their security effectiveness.

Good security effectiveness

Implement DMARC to enhance SPF and/or DKIM

34. Domain-based Message Authentication, Reporting and Conformance (DMARC) enables a domain owner to specify a policy stating what action the recipient's email server should take if it has failed a SenderID/Sender Policy Framework (SPF) check and/or Domain Keys Identified Mail (DKIM) check. The domain owner can specify the action the recipient email server should take to include "reject" (rejection of the email by the email recipient's email server), "quarantine" (mark email as spam) or "none" (no specific action to be taken). DMARC also provides a reporting feature which enables a domain owner to receive reports on the DMARC actions taken by receiving email servers. While this feature does not mitigate malicious emails sent to the domain owner's organisation, it can give the domain owner some visibility of attempts by adversaries to spoof their organisation's domain.
35. Organisations should configure a DMARC record specifying that emails from the organisation's domain and sub-domains be rejected if they fail SenderID/SPF and/or DKIM checking. Organisations that currently only have a SenderID/SPF record published are still able to implement DMARC without having to implement DKIM. In this situation, a SenderID/SPF fail on its own will still result in a DMARC fail.

Average security effectiveness

Block email on SenderID/SPF 'hard fail'

36. Checking SenderID/SPF will verify if emails originate from the domain they claim to originate from and allow organisations to block them if checks fail. An SPF 'hard fail' occurs when an email is received which has been verified as not originating from the domain it claims to originate from. SPF 'hard fails' should be blocked and investigated. An SPF 'hard fail' can indicate a phishing attempt, especially if the failed email is spoofed to appear to come from a legitimate domain.
37. When implementing SenderID/SPF checks, organisations should ensure they publish SenderID/SPF records for their own domain, and ensure that SenderID/SPF checks are conducted on emails purporting to be sent from their domain. This will prevent adversaries sending emails to organisations and spoofing the sender to appear as though it originated from the organisation the email is being sent to, a tactic common in many cyber-enabled fraud cases.

Block email on DKIM fail

38. DKIM is a method of verifying the sender's domain of an email using the signatures provided by the sending domain. When an email fails DKIM verification, the email should be blocked and investigated. This should also be logged and potentially reported to the organisation that the email was claiming to originate from.

Minimal security effectiveness

Incorporate spam blacklists

39. Known spam email senders and addresses should be blocked without the email being examined.

Quarantine email on SenderID/SPF 'soft fail'

40. Checking SenderID/SPF will verify if emails originate from the domain they claim to originate from and allow organisations to block them if checks fail. An SPF 'soft fail' occurs when an SPF enabled domain cannot guarantee that an email was sent from an authorised server of that domain. When an SPF 'soft fail' is encountered, the email should be quarantined rather than blocked allowing users to retrieve it if it was considered a legitimate email.

Poor security effectiveness

Flag email on SenderID/SPF 'soft fail'

41. Checking SenderID/SPF will verify if emails originate from the domain they claim to originate from and allow organisations to block or quarantine them if checks fail. An SPF 'soft fail' occurs when an SPF enabled domain cannot guarantee that an email was sent from an authorised server of that domain. Instead of blocking or quarantining the email, the email should be marked as potentially malicious before being sent to users to inform users of the risks and allow them to make a risk-based decision as to whether to accept the email. For example, the subject line of an email could be modified to highlight and identify to the user that the email is from an unverified or unconfirmed sender.

Mark external emails

42. Emails received from external organisations should be marked with an additional header to encourage recipients to exercise additional caution when acting on links or attachments associated with the email.

Other mitigation strategies

Block non-authorized third party email services

43. Given the ability, many users would like to be able to access third party email accounts from a corporate network. This access can include adding third party services to corporate email clients or accessing personal webmail accounts. As these are third party service providers, organisations have no control over the data going in and out of these services. Blocking access to non-approved third party email services can assist in the prevention of malicious content entering networks through a third party service, prevent corporate data leaving network through a non-corporate service and maintain records of official correspondence by ensuring the use of the corporate email service.

Log and audit email related actions and events

44. Logging of actions and events from the email content filter and email servers should be implemented, with these logs audited on a regular basis. Effective logging and auditing will help in the event of a current or past cyber security incident.

Implement additional email content filter functionality

45. While this document focuses on providing mitigation strategies to reduce the security risk of workstations, networks and associated sensitive information being compromised by malicious emails, the following additional mitigation strategies will improve the effectiveness of an email content filter and simplify its management.
- Minimise overhead for a system administrator to release quarantined emails:** Minimising the overhead for a system administrator to assess and release an email for a user when that email has been quarantined can be achieved by providing them with easy and ready access to a secure environment to examine quarantined emails.
 - Implement self-release of quarantined emails (based on quarantine reason):** Allowing users to self-release a quarantined email without needing to go through a system administrator can be made available for selected quarantined emails based on email content filter triggers considered to be a lesser security risk. Even so, all email self-releases should still be logged for auditing purposes.

Further information

46. The *Australian Government Information Security Manual* (ISM) assists in the protection of official government information that is processed, stored or communicated by Australian Government systems. It can be found at: <http://www.asd.gov.au/infosec/ism/>.
47. ASD's *Strategies to Mitigate Targeted Cyber Intrusions* complements the advice in the ISM. The complete list of mitigation strategies and supporting publications can be found at: <http://www.asd.gov.au/infosec/mitigationstrategies.htm>.
48. Additional guidance on Sender Policy Framework (SPF) can be found in ASD's *Mitigating Spoofed Emails – Sender Policy Framework Explained* publication. It can be found at: http://asd.gov.au/publications/protect/spoof_email_sender_policy_framework.htm.

49. Additional guidance on securely using macros within organisations can be found in ASD's *Microsoft Office Macro Security* publication. It can be found at: <http://asd.gov.au/publications/protect/ms-office-macro-security.htm>.
50. For further information on conservatively deploying DMARC, please refer to Google's advice at: <https://support.google.com/a/answer/2466563>. This may be of use if organisations have a complex email configuration and are unaware of where exactly sent emails originate from within their domain, or if organisations are otherwise concerned about the risk of legitimate emails sent from their organisation's domain being rejected.

Contact details

51. Australian government customers with questions regarding this advice should contact ASD Advice and Assistance by emailing asd.assist@defence.gov.au or by calling 1300 CYBER1 (1300 292 371).
52. Australian businesses and other private sector organisations with questions regarding this advice should contact CERT Australia by emailing info@cert.gov.au or by calling 1300 172 499.

Attachment A

Mitigation Strategy	Security Effectiveness	User Resistance	Upfront Cost (Staff, Equipment, Technical Complexity)	Maintenance Cost (Mainly Staff)	Designed to Prevent or Detect a Cyber Intrusion	Helps Mitigate Cyber Intrusion Stage 1: Code Execution	Helps Mitigate Cyber Intrusion Stage 2: Network Propagation	Helps Mitigate Cyber Intrusion Stage 3: Data Exfiltration
<i>Attachment filtering</i>								
Convert attachments to another format	Excellent	High*	Medium	Medium*	Prevent	Yes	No	No
Whitelist attachments based on file typing	Excellent	Medium	Medium	Low	Prevent	Yes	No	Yes^
Block password protected archives and unidentifiable or encrypted attachments	Excellent	Medium	Medium	Low	Prevent	Yes	No	Yes
Perform automated dynamic analysis of attachments run in a sandbox	Excellent	Low	Medium	Low	Prevent	Yes	No	No
Sanitise attachments to remove active or potentially harmful content	Excellent	Medium*	High	Medium*	Prevent	Yes	No	No
Disable or control macros in Microsoft Office files	Excellent	Medium*	High	Low*	Prevent	Yes	No	No

Mitigation Strategy	Security Effectiveness	User Resistance	Upfront Cost (Staff, Equipment, Technical Complexity)	Maintenance Cost (Mainly Staff)	Designed to Prevent or Detect an Intrusion	Helps Mitigate Intrusion Stage 1: Code Execution	Helps Mitigate Intrusion Stage 2: Network Propagation	Helps Mitigate Intrusion Stage 3: Data Exfiltration
Controlled inspection of archive files	Good	Low	Medium	Low	Both	Yes	No	Yes
Whitelist attachments based on file extension	Average	Medium	Low	Low	Prevent	Yes ^{&}	No	Yes [^]
Blacklist attachments based on file typing	Minimal	Low	Low	Medium	Prevent	Yes	No	Yes [^]
Scan attachments using antivirus software	Minimal	Low	Low	Low	Both	Yes	No	No
Blacklist attachments based on file extension	Minimal	Low	Low	Medium	Prevent	Yes ^{&}	No	Yes [^]
<i>Email body filtering</i>								
Replace active web addresses in an email's body with non-active versions	Good	Low	Medium	Low	Prevent	Yes	No	No
Remove active content in an email's body	Average	Low	Medium	Low	Prevent	Yes	No	No

Mitigation Strategy	Security Effectiveness	User Resistance	Upfront Cost (Staff, Equipment, Technical Complexity)	Maintenance Cost (Mainly Staff)	Designed to Prevent or Detect an Intrusion	Helps Mitigate Intrusion Stage 1: Code Execution	Helps Mitigate Intrusion Stage 2: Network Propagation	Helps Mitigate Intrusion Stage 3: Data Exfiltration
Sender Verification								
Implement DMARC to enhance SPF and/or DKIM	Good	Low	Low	Low	Prevent	Yes	No	No
Block email on SenderID/SPF 'hard fail'	Average	Low	Low	Low	Prevent	Yes	No	No
Block email on DKIM fail	Average	Low	Low	Low	Prevent	Yes	No	No
Incorporate spam blacklists	Minimal	Low	Low	Low	Both [#]	Yes	No	No
Quarantine email on SenderID/SPF 'soft fail'	Minimal	Medium	Low	Low	Prevent	Yes	No	No
Flag email on SenderID/SPF 'soft fail'	Poor	Low	Low	Low	Prevent	Yes	No	No
Mark external emails	Poor	Low	Low	Low	Prevent	Yes	No	No

Notes

- Mitigation strategies are ranked in categories based on their security effectiveness.
- * Potentially lower if document release is easy.
- # If the mitigation strategy is applied to both incoming and outgoing emails, then this is 'Both', otherwise just 'Prevent'.
- & Provided adversaries are sending a file with the blocked extension.
- ^ Provided adversaries are attempting to exfiltrate a file type that is blocked.