



Malicious Email Mitigation Strategies

First published: June 2012

Last updated: October 2021

Introduction

Socially engineered emails containing malicious attachments and embedded links are routinely used in targeted cyber intrusions against organisations. This publication has been developed to provide mitigation strategies for the security risks posed by these malicious emails.

Not every mitigation strategy within this publication 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.

The mitigation strategies, and implementation considerations, are summarised in Appendix A.

Definitions

This publication uses the terms 'block' and 'quarantine'. In the context of this publication, '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

Attachments are a significant security risk associated with emails. Effective attachment filtering reduces the likelihood of malicious content reaching a user's workstation. Mitigation strategies associated with attachment filtering are discussed below.

Convert attachments to another format

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.

Allow attachments based on file typing

File typing inspects the content of a file to determine its file type rather than relying on its extension. Only file types that have a legitimate business purpose and an acceptable risk profile for organisations should be allowed. As file extensions

can be changed, 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

Content within password protected archives can't be trusted since email content filters can't 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 only allowing attachments based on file typing. Where organisations have corporately approved encrypted email communications, such as S/MIME or PGP, these can be allowed to prevent disruption to legitimate business.

Perform automated dynamic analysis of attachments run in a sandbox

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.

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.

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.

Sanitise attachments to remove active or potentially harmful content

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.

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

An increase in the use of macros in Microsoft Office files being used as a malware delivery vector has been observed. 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, malicious actors are also using macros to perform a variety of malicious activities including the download and execution of malware on the host computer.

Organisations should configure Microsoft Office to disable all macros by default and only run macros vetted as trustworthy and placed in 'trusted locations' which typical low-privileged users can't write to.

Controlled inspection of archive files

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.

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.

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 malicious actors to avoid controlled inspection of archive files. To mitigate this, organisations should attempt to decompress all attachments, with all decompressed files submitted to the controls for attachments and the original attachment blocked if any decompressed files fail.

Allow attachments based on file extension

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 allowed.

Block attachments based on file typing

Blocking attachments based on file typing is less proactive and thorough than allowing 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

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.

Block attachments based on file extension

Blocking attachments based on file extension is less proactive and thorough than allowing attachments based on file typing or file extension. Blocking 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

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. Mitigation strategies associated with filtering the body of an email are discussed below.

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

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 a malicious website. Hovering over the address may reveal the actual website.

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

Remove active content in an email's body

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

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. Mitigation strategies for sender verification are discussed below.

Implement DMARC to enhance SPF and/or DKIM

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 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 malicious actors to spoof their organisation's domain.

Organisations should configure a DMARC record specifying that emails from the organisation's domain and sub-domains be rejected if they fail SPF and/or DKIM checking. Organisations that currently only have a SPF record published are still able to implement DMARC without having to implement DKIM. In this situation, a SPF fail on its own will still result in a DMARC fail.

Block email on SPF 'hard fail'

Checking 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.

When implementing SPF checks, organisations should ensure they publish SPF records for their own domain, and ensure that SPF checks are conducted on emails purporting to be sent from their domain. This will prevent malicious actors 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

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.

Block known spam email senders

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

Quarantine email on SPF 'soft fail'

Checking 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 can't 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.

Flag email on SPF 'soft fail'

Checking 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 can't 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 security 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

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-authorised third party email services

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

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

While this publication 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

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.

Further information on implementing SPF, DKIM and DMARC is available in the [How to Combat Fake Emails](#) publication.

Further information on securing the use of Microsoft Office macros within organisations is available in the [Restricting Microsoft Office Macros](#) 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).

Appendix A: Malicious email mitigation strategies

Mitigation Strategy	Security Effectiveness	User Resistance	Upfront Cost (Staff, Equipment, Complexity)	Maintenance Cost (Staff)	Prevent or Detect a Targeted Cyber Intrusion	Helps Mitigate Code Execution	Helps Mitigate Network Propagation	Helps Mitigate Data Exfiltration
Attachment filtering								
Convert attachments to another format	Excellent	High ¹	Medium	Medium ¹	Prevent	Yes	No	No
Allow attachments based on file typing	Excellent	Medium	Medium	Low	Prevent	Yes	No	Yes ²
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 ¹	High	Medium ¹	Prevent	Yes	No	No
Disable or control macros in Microsoft Office files	Excellent	Medium ¹	High	Low ¹	Prevent	Yes	No	No
Controlled inspection of archive files	Good	Low	Medium	Low	Prevent & Detect	Yes	No	Yes
Allow attachments based on file extension	Average	Medium	Low	Low	Prevent	Yes ³	No	Yes ²
Block attachments based on file typing	Minimal	Low	Low	Medium	Prevent	Yes	No	Yes ²
Scan attachments using antivirus software	Minimal	Low	Low	Low	Prevent & Detect	Yes	No	No
Block attachments based on file extension	Minimal	Low	Low	Medium	Prevent	Yes ³	No	Yes ²
Email body filtering								
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
Sender verification								
Implement DMARC to enhance SPF and/or DKIM	Good	Low	Low	Low	Prevent	Yes	No	No
Block email on SPF ‘hard fail’	Average	Low	Low	Low	Prevent	Yes	No	No
Block email on DKIM fail	Average	Low	Low	Low	Prevent	Yes	No	No
Block known spam email senders	Minimal	Low	Low	Low	Prevent & Detect ⁴	Yes	No	No
Quarantine email on SPF ‘soft fail’	Minimal	Medium	Low	Low	Prevent	Yes	No	No
Flag email on SPF ‘soft fail’	Poor	Low	Low	Low	Prevent	Yes	No	No
Mark external emails	Poor	Low	Low	Low	Prevent	Yes	No	No

- Notes**
- 1. Potentially lower if document release is easy.
 - 2. Provided malicious actors are attempting to exfiltrate a file type that is blocked.
 - 3. Provided malicious actors are sending a file with the blocked extension.
 - 4. If the mitigation strategy is applied to both incoming and outgoing emails, then this is ‘Prevent & Detect’, otherwise just ‘Prevent’.