Government Digital Services Delivery

Secure Government Email Keyword Blocking Deployment Guide

Issued by

Digital Services branch

Internal Affairs
Te Tari Taiwhenua

New Zealand Government

Contents

1. Document Control

Service Name	Secure Government Email
Author	Matt Oliver, Operations Security Specialist, Department of Internal Affairs
Title	Secure Government Keyword Blocking Deployment Guide
Date and Version	12 February 2026, v1.2
Index Number	GDSD.SGE.Guidance.2026_040

2. Introduction

The Secure Government Email (SGE) Common Implementation Framework was confirmed late in 2024 with a goal to improve email security standards across the New Zealand Government. This deployment guide compliments the Framework through providing examples on how to deploy the required keyword blocking in line with both the framework and to meet the requirements of the New Zealand Information Security Manual (NZISM).

It provides configuration examples covering commonly used platforms and configurations but is not an exhaustive guide. It is recommended Agencies without experience in the deployment of email security tools engage with a provider with expertise in this area.

Services are available for Agencies via the Secure Email Management and Administration service in Marketplace.

Marketplace | Pae Hokohoko — Infrastructure Managed Services catalogue

Note: Document classification and keywords
This document is UNCLASSIFIED. By necessity it contains all the SGE keywords and classification labels for blocking, which will impede deliverability and sharing by email. One solution is to send it as a password protected encrypted attachment.

Note: SGE Restricted Group members
Agencies who are a member of the SGE Restricted Group should take care on the placement of the keyword blocking transport rules as they may impact existing legitimate email flows via SEEMail. These keyword blocks must be in place before you exit SEEMail.

3. NZISM and SGE Framework requirements for Keyword Blocking

Chapter 15.2 (Email Infrastructure) of the NZISM has the following two specific controls relating to the blocking of classified email:

15.2.40.C.01 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.

15.2.41.C.01 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.

> **Note:** For Agencies who are members of the Restricted Group, the same word list applies. The SGE Connectors Configuration Guide document provides details on how the permitted keywords bypass these rules. Excluding the words from this list may create a security risk.

The SGE Framework requires Data Loss Prevention (DLP) be enforced in line with the PSR and NZISM requirements. It recommends, at a minimum, the following keywords be added to your DLP block lists: [RESTRICTED], [SENSITIVE], [CONFIDENTIAL], [SECRET], [TOPSECRET], [TOP SECRET], and [TOP-SECRET]. To allow for typing errors each keyword is repeated with variations using {}, [}, and {].

4. Keyword Blocking Locations

Keyword blocking can be achieved in various locations of the mail path depending on licensing and purchased options. This may be through Transport Rules, Microsoft Purview or an external email security platform. Some tools may allow keyword blocking directly within the end users email client. While this is useful for outbound email, inbound email still needs to be blocked at the gateway.

If you are deploying keyword blocking via an external email security provider and you have SEEMail in your network, please keep in mind the possible impact on SEEMail depending on where that sits in the mail flow path.

If using Google Workspace blocking can be achieved through using Objectional Content rules, or Content Compliance rules. Examples are not provided for this.

The exact location of the rules is optional. The NZISM and SGE Framework requirements is the applicable keywords are blocked on inbound and outbound mail flows.

This deployment guide only provides detailed configuration steps for keyword blocking using Microsoft Exchange Online Transport Rules configured via PowerShell.

5. Microsoft Exchange Online configuration

It is possible to build the required transport rules via the Exchange Online GUI but it is strongly recommended this is done via PowerShell.

These instructions provide the minimum possible configuration to achieve the desired outcome. Agencies may add other rules or controls to suit their business requirements, such as blocking attachments by file type or size. You may need to alter the messaging or other settings in line with your organisation's requirements.

> **Note:** Use with care: The transport rules will immediately become active for all domains in the tenancy.

Using the PowerShell scripts will add the new rules without priority so they will appear at the bottom of all existing transport rules.

If you would like to import the rules without them becoming active you can do so through changing the second line of each script from Mode Enforce to Mode Audit.

There are four scripts provided in the SGE Keyword Blocking Scripts v1.2.txt file. They cover email content inbound and outbound, and attachments content inbound and outbound. Attachment content inspection depends on file type and size. Encrypted or unsupported formats will not be scanned. These will return a 'no match' by the transport rule, and the message will be permitted by the transport rule.

6. Testing

These rules can be tested through sending emails to and from your domains with the applicable tags, to ensure emails are being blocked and the correct response messages sent.

7. Reporting

Dropped messages are logged in your tenancy and can be reported on through the Exchange Transport Rule report. Go to the Exchange Admin Centre | Reports | Mail flow | Exchange Transport Rule report. This will show the Date, Transport Rule, Subject, Sender Address, Recipient Address, Severity and Direction.

8. Considerations for SEEMail Agencies

8.1. SEEMail trigger words no longer in use

The legacy SEEMail platform used the square bracketed trigger words [TRUSTED], [SEEMAIL], and [IN-CONFIDENCE]. These are no longer required with the introduction of the Secure Government Email framework.

> **Note:** This does not change the formal PSR classification IN-CONFIDENCE. It only removes the legacy SEEMail matching keyword in square brackets.

8.2. Configuration during migration

There are several different possible deployments of SEEMail which may impact how these rules function, especially on the inbound flow. The most common SEEMail deployment has SEEMail operating as an external gateway in front of your tenancy. All inbound email passes through the SEEMail gateway prior to delivery to your tenancy. All outbound email passes through SEEMail after exiting your tenancy.

In this case inbound emails with the applicable tags should be blocked, or encapsulated with a warning, at the SEEMail gateway before reaching your tenant. This means the inbound rule will likely not be hit by any messages.

For outbound email with SEEMail configured as the external gateway, SEEMail will no longer see these emails being dropped. They will instead be dropped at the tenancy level and will no longer appear in SEEMail reports.

For other configurations Agencies need to consider possible mail flow and reporting impacts.

These are temporary reporting issues, and will only exist between the time the rules are put in place and SEEMail is abandoned. In all cases emails which should be blocked will be blocked.

Appendix 1 – Safe Senders Lists and Transport Rules

Some Exchange Transport Rules can be influenced by end user use of safe senders lists. These may impact Spam Confidence Level (SCL)

The primary ways they interact include:

- SCL Overrides: When an external sender is added to a user's Safe Senders List, the system automatically sets the Spam Confidence Level (SCL) to -1 for that message. This often allows the email to bypass a transport rule designed to block or quarantine mail based on spam content.
- Delivery Priority: Even if a transport rule manually sets a high SCL, the Outlook Safe Senders List can override the rule and deliver the email to the recipient's Inbox.
- Order of Operations: Transport rules generally fire at the 'OnRouted' stage, while certain content filtering and safelist checks occur later in the pipeline. This means a message can be modified by a transport rule but still be delivered to the Inbox—or Junk—based on the user's personal safe/blocked sender settings.
- Security Bypasses: Attackers sometimes exploit this by convincing users to add them to their Safe Senders List, which can effectively neuter organizational anti-spam policies or transport rules intended to protect the company.

A Safe Senders list cannot bypass a transport rule designed to block keywords. Transport rules take precedence because they are processed earlier in the email pipeline than user-level junk settings.

The specific interaction depends on the action your transport rule takes:

- Blocking/Rejecting (Hard Block): If your transport rule is configured to 'Reject the message with the explanation...' or 'Delete the message without notifying anyone' based on a keyword, the email is terminated while in transit. It never reaches the user's mailbox, so the Safe Senders list is never evaluated for that message.
- Setting SCL (Spam Scoring): If your transport rule uses keywords to mark a message as spam (e.g., setting the SCL to 6 or 9), the Safe Senders list will override it. The message will be delivered to the Inbox instead of the Junk folder because the user's safe designation essentially resets the SCL to -1 at the point of delivery.
- Quarantining: Messages quarantined by a transport rule are held at the organizational level. While a Safe Senders list might help a message avoid the standard spam filter, it typically does not release a message that has been explicitly quarantined by a custom administrative mail flow rule.

Summary of Priority

1. Transport Rules (Block/Reject): Highest priority; stops the email before mailbox delivery. This is the type of rule used within this SGE Keyword blocking guide.
2. Safe Senders List: Bypasses content-based spam filtering and overrides SCL scores set by transport rules.
3. Standard Spam Filters: Lowest priority; only evaluates what hasn't been handled by the above.

Appendix 2 – Frequently Asked Questions

How do these rules differ from the SEEMail defaults?

These rules are very similar to the SEEMail default rules. The changes are:

1. The term top-secret (with the – in the middle) has been added to the blocked list.
2. Inbound messages from a non-SEEMail location with a SEEMail tag are, in most instances, encapsulated in a SEEMail warning message and forwarded on to the recipient. This action changes as emails with PSR designated classifications higher than IN-CONFIDENCE are not permitted to be received via email sources outside of Restricted Group members.

Is this Data Loss Prevention (DLP)?

Keyword blocking is a low level form of DLP. It only protects against the specifically listed keywords where those keywords are transmitted via email.

The keyword blocking requirements within the SGE framework are designed to match the NZISM and PSR requirements, in a similar manner used in the legacy SEEMail platform.

Will these rules scan linked files?

Within M365 the default method for 'sending attachments' has changed so sharing access to the file is preferred. In this situation only a link to the file is sent via email, as opposed to a file attachment. Links are not interrogated by the transport rules, so the email will pass. Protection on file access permissions are not covered within the Secure Government Email Framework.

Where do these rules apply on the inbound mail flow through M365?

Mail flow rules are evaluated after anti-spam scanning, once the message enters the Exchange transport pipeline. This is after inbound connectors are evaluated, so messages received via connectors are still subject to mail flow rules. Outbound emails are processed in the reverse order, so mail flow rules are processed before email is routed into any connectors.

Inbound Mail Processing Order in Microsoft 365:

1. Connection + SMTP session
    - TLS negotiation
    - IP allow / block lists
    - Inbound Connector matching. This determines how the message is accepted, authentication method, and any restrictions. If no connector matches, the message is treated as Anonymous Internet mail.
2. Anti-Spam and Anti-Malware Filtering which includes:
    - Connection filtering
    - SPF, DKIM, DMARC evaluation (required for spam scoring, not for acceptance)
    - Malware scanning
    - ETR (Exchange Transport Rules) spam actions (quarantine, etc.)

3. Transport Pipeline (where Mail Flow Rules run)
    - Once the message passes filtering and is accepted into the transport pipeline, Exchange Online evaluates Mail Flow Rules (ETRs) in this sequence:
    - Rules scoped to connectors (rare)
    - Rules in priority order (1 → 2 → 3 …)
    - Actions applied (modify headers, redirect, block, etc.)
4. Message categorization
    - Recipient resolution
    - Distribution group expansion
    - Routing decisions
5. Mailbox Delivery
    - Delivered to mailbox
    - Or forwarded/redirected
    - Or routed to on-prem via Hybrid connector
    - Or rejected (e.g., journaling failures)

Do these rules apply to all my domains?

By default Mail flow rules apply to all domains in the tenancy. This behaviour can be altered through adding exceptions to the rules. In the rule conditions the bottom section has an 'Except if' condition with a number of possible controls.

A scenario where this may be useful is where a tenant has a number of domains, including a test domain and they would like to apply the rules only to the test domain. After testing is complete they wish to deploy the rules to their production domains one at the time. To achieve this, add all the domains to the exclusions list except for your test domain. Once testing is complete remove the production domains on a per domain basis when required. Remember to do this on both the Inbound and Outbound rules.

Are the keywords case sensitive?

No, keywords used in Mail flow rules are not case sensitive.

Appendix 3 – List of Acronyms

Acronym	Definition
DKIM	Domain Keys Identified Messaging
DLP	Data Loss Prevention
DMARC	Domain-based Message Authentication, Reporting, and Conformance
ETR	Exchange Transport Rule(s)
GDSD	Government Digital Services Delivery
GUI	Graphical User Interface
M365	Microsoft 365
NZISM	New Zealand Information Services Manual
SCL	Spam Confidence Level
SGE	Secure Government Email (Framework)
SPF	Sender Policy Framework
TLS	Transport Layer Security