

Response to Cisco Secure Email Gateway SMTP Smuggling Vulnerability Report

Contents

[Introduction](#)

[Technical Background](#)

[Cisco Secure Mail Behavior](#)

[Clean Messages of Bare CR and LF Characters \(Default\)](#)

[Reject Messages with Bare CR or LF Characters](#)

[Allow Messages with Bare CR or LF Characters \(Deprecated\)](#)

[Recommended Configuration](#)

[Frequently Asked Questions](#)

Introduction

This document provides more details on how Cisco Secure Email behaves against the type of attack described in [SMTP Smuggling - Spoofing E-Mails Worldwide](#), published on December 18th, 2023 by SEC Consult.

*In the course of a research project in collaboration with the SEC Consult Vulnerability Lab, Timo Longin (@timolongin) discovered a novel exploitation technique for yet another Internet protocol - SMTP ([Simple Mail Transfer Protocol](#)). Threat actors could abuse vulnerable SMTP servers worldwide to send malicious e-mails from arbitrary e-mail addresses, allowing targeted phishing attacks. Due to the nature of the exploit itself, this type of vulnerability was dubbed **SMTP smuggling**.*

Cisco has not found any evidence that the attack described in the paper could be used to bypass any of the configured security filters.

Technical Background

Without going into detail on the SMTP protocol and message format, it is important to look at a few sections of [RFC 5322](#) in order to get some context.

[Section 2.1](#) defines the CRLF character sequence as the separator to be used between the different sections of the message.

Messages are divided into lines of characters. A line is a series of characters that is delimited with the two characters carriage-return and line-feed; that is, the carriage return (CR) character (ASCII value 13) followed immediately by the line feed (LF) character (ASCII value 10). (The carriage return/line feed pair is usually written in this document as "CRLF".)

[Section 2.3](#) is more specific on the format of the message body. It clearly states that CR and LF characters should never be sent independently as part of the body. Any server doing so is not compliant with the RFC.

The body of a message is simply lines of US-ASCII characters. The only two limitations on the body are as follows:

- CR and LF MUST only occur together as CRLF; they MUST NOT appear independently in the body.
- Lines of characters in the body MUST be limited to 998 characters, and SHOULD be limited to 78 characters, excluding the CRLF.

However, [Section 4.1](#) of that same document, in reference to obsolete syntax from previous revisions of the RFC that were not as restrictive, acknowledges that many implementations on the field are not using the right syntax.

Bare CR and bare LF appear in messages with two different meanings. In many cases, bare CR or bare LF are used improperly instead of CRLF to indicate line separators. In other cases, bare CR and bare LF are used simply as US-ASCII control characters with their traditional ASCII meanings.

To summarize, according to RFC 5322, a properly formatted SMTP message would look like the following example:

```
ehlo sender.example\r\n
mail FROM:<user@sender.example>\r\n
rcpt TO:<user@receiver.example>\r\n
data\r\n
From: <user@sender.example>\r\n
To: <user@receiver.example>\r\n
Subject: Example\r\n
\r\n
lorem ipsum\r\n
\r\n. \r\n
```

The paper attempts to leverage the exception mentioned in [Section 4.1](#) of the RFC to insert or "smuggle" a new message as part of the body in an attempt to bypass security measures on the sending or receiving server. The objective is for the smuggled message to bypass security checks because those checks would only be run on the part of the message before the bare line feeds. For example:

```
<#root>

ehlo sender.example\r\n
mail FROM:<user@sender.example>\r\n
rcpt TO:<user@receiver.example>\r\n
data\r\n
From: <user@sender.example>\r\n
To: <user@receiver.example>\r\n
Subject: Example\r\n
\r\n
lorem ipsum\r\n
\r\n. \r\n

mail FROM:<malicious@malicious.example>

\r\n

rcpt TO:<user@receiver.example>

\r\n

data

\r\n

From: <malicious@malicious.example>
```

\r\n

To: <user@receiver.example>

\r\n

Subject: Malicious

\r\n

\r\n

Malicious content

\r\n

\r\n

.

\r\n

Cisco Secure Mail Behavior

When configuring an SMTP listener on Cisco Secure Mail, there are three configuration options that determine how bare CR and LF characters should be treated.

Clean Messages of Bare CR and LF Characters (Default)

With the default option selected, Cisco Secure Mail replaces all bare CR and LF characters in incoming messages with the correct CRLF sequence.

A message with smuggled content, like the one in the example, is treated as two separate messages, and all the security checks (such as Sender Policy Framework (SPF), Domain-based Message Authentication, Reporting & Conformance (DMARC), AntiSpam, Antivirus, Advanced Malware Protection (AMP), and content filters) are run independently on each of them.



Note: Customers should be aware that, with this configuration, an attacker might be able to smuggle a message impersonating a different user. An attacker could have greater impact in situations where the originating server hosts multiple domains because the attacker could impersonate a user from one of the other domains that are hosted on the server, and the SPF check on the smuggled email would still pass.

Reject Messages with Bare CR or LF Characters

This configuration option strictly enforces compliance to the RFC. Any messages containing bare CR or LF characters are rejected

Although this configuration prevents the smuggling scenario, it will also cause legitimate emails coming from servers that are not RFC-compliant to be dropped.

Allow Messages with Bare CR or LF Characters (Deprecated)

The final configuration causes Cisco Secure Mail to treat bare CR and LF characters with their ASCII meaning. The message body is delivered as-is, including the smuggled content.

Because the smuggled message is treated as being part of the body, attachments included as part of the smuggled message might not be detected by Cisco Secure Mail. This might cause security issues on downstream devices. **This option has been deprecated and should no longer be used.**

Recommended Configuration

Cisco recommends using the default "Clean messages of bare CR and LF characters" option because it provides the best compromise between security and interoperability. However, customers using this setting should be aware of the security implications in regards to smuggled content. Customers who want to enforce RFC compliance should choose "Reject messages with bare CR or LF characters," being aware of the potential interoperability issues.

In any case, **Cisco strongly recommends configuring and using features such as SPF, DomainKeys Identified Mail (DKIM), or DMARC in order to validate the sender of an incoming message.**

AsyncOS Releases 15.0.2 and 15.5.2 and later add new functionality that helps identify and filter messages that do not comply with the end-of-message RFC standard. If a message with an invalid end-of-message sequence is received, the email gateway adds an X-Ironport-Invalid-End-Of-Message Extension Header (X-Header) to all message IDs (MIDs) within that connection until a message that complies with the end-of-message RFC standard is received. Customers can use a content filter to look for the "X-Ironport-Invalid-End-Of-Message" header and define actions to be taken for these messages.

Frequently Asked Questions

Is Cisco Secure Mail vulnerable to the attack described?

Technically, yes. When bare CR and LF characters are included in the mail, it is possible to cause part of the email to be treated as a second email. However, since the second email is analyzed independently, the behavior is equivalent to sending two separate messages. Cisco has not found any evidence that the attack described in the paper could be used to bypass any of the configured security filters.

The paper provides examples of bypassed SPF and DKIM checks. Why does Cisco say that no filters are being bypassed?

In those examples, SPF checks are run as expected, but result in a passed check due to the sending server owning multiple domains.

What is the recommended configuration?

The most appropriate choice for a customer is reliant on their specific requirements. The recommended options are either the default "Clean" configuration or the "Reject" alternative.

Will choosing the Reject option result in false positives?

The "Reject" function initiates an assessment of the email's adherence to the RFC standards. In the event that the email fails to comply with the RFC standards, it will be declined. Even legitimate emails can be rejected if the email is not in compliance with RFC standards.

Is there a software bug covering this issue?

Cisco bug ID [CSCwh10142](#) was filed.

How can I get more information on this topic?

Any follow-up questions can be raised through a Technical Assistance Center (TAC) case.