Microsoft

A defense-in-depth approach to protecting email with Microsoft Exchange Online Protection (EOP)

Published: Oct 2014

Summary: This white paper describes the defense-in-depth approach that the Microsoft Exchange Online Protection (EOP) online service uses in order to stop malicious email messages from compromising your organization's security.

Author: Curtis Parker, Shobhit Sahay

For the latest information, please visit the Office 365 Trust Center at http://trust.office365.com

Table of Contents

Introduction

This white paper outlines the multi-layer defense-in-depth approach that the Microsoft Exchange Online Protection (EOP) service follows in order to prevent malicious threats in email communication from infiltrating your organization and compromising the safety of your data. EOP helps to ensure that malicious code or activities are stopped at multiple check points before reaching the email infrastructure, reducing the probability of infection. EOP scans every single mail arriving to mailboxes stored in Exchange Online as well as used by tens of millions of users with mailboxes hosted in their IT organizations on premises

Changing world of email attacks

Due to the constantly evolving threat landscape and scale of email botnets today, content-only scanning is no longer a viable, scalable strategy. Today, billions of inbound messages are handled in EOP on a daily basis. Out of these, about 66 percent of the messages are spam. Spamming techniques have evolved over the years in order to penetrate several filtering programs designed to stop the attacks and having a solution to stop such threats has become crucial to business around the world.

EOP uses a multi layered defense approach throughout the mail delivery process to help ensure that the message users see is safe. The following diagram shows the architecture for the EOP service as the mail is routed to and from your organization.

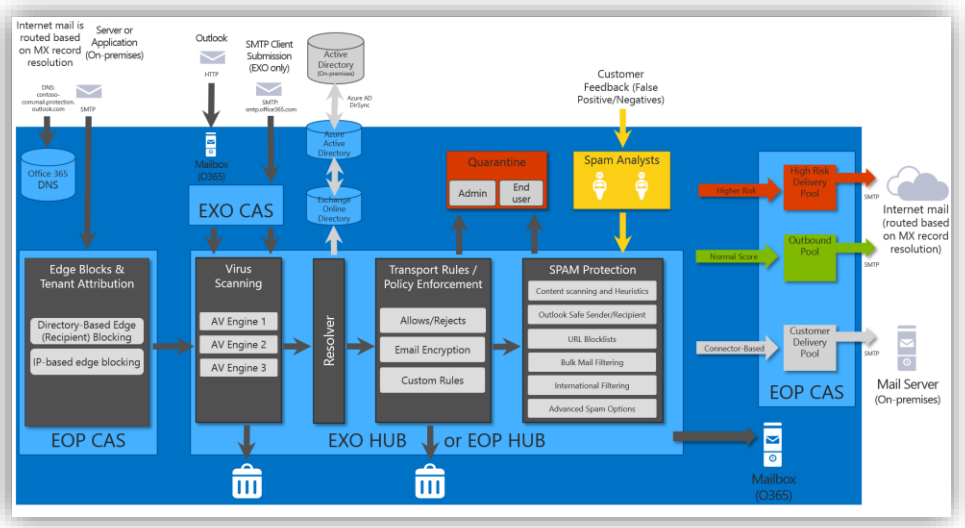

EOP mail filtering process for inbound and outbound mail.

Block lists used by EOP

EOP uses block lists from some third party partners (e.g. Spamhaus, SURBL.org, URIBL.com, Invaluement) but also maintains its own proprietary block list containing IP addresses of confirmed known senders of spam sent directly to EOP customers.

A second, much smaller set of addresses is manually compiled by the spam analysts within EOP in response to observed spam trends. The analysts provide world-class spam research and response capabilities that support the various spam detection technologies in EOP. With analysts around the globe, EOP is able to respond quickly and effectively to new threats. Analysts write spam rules to identify spam in the English language in addition to providing language coverage for many more commonly used languages.

Finally, the EOP block lists are range-weighted - they contain IP addresses that are not known sources of spam, but whose neighboring IP addresses are known sources of spam. If an IP address range has several known offenders and no known sources of legitimate email, EOP might proactively block the entire range as a precautionary measure.

Domain Name System block lists (DNSBLs)

When a mail filtering system is trying to determine whether or not to accept an incoming email message, the originating IP address is one of the few definitive pieces of information about the sender that can be verified. EOP uses Domain Name System (DNS) block lists from a number of third party partners to block incoming messages from IP addresses that are known sources of spam. For example, since the landscape of malicious and compromised IP addresses changes quickly, EOP downloads and applies the most current copies of the Spamhaus Project's block lists hourly.

The following table describes three of the Spamhaus block lists:

Block List	Description
Spamhaus Block List (SBL)	A database of IP addresses that Spamhaus deems to be involved in the sending, hosting, or origination of spam.
Spamhaus Exploits Block List (XBL)	A real-time database containing the IP addresses of hijacked devices that have been infected by: third-party exploits, such as open proxies (HTTP, socks, AnalogX, wingate, and others); worms or viruses with built-in spam engines; and other types of Trojan-horse exploits.
Spamhaus Policy Block List (PBL)	A DNSBL database of end-user IP address ranges that should not be delivering unauthenticated Simple Mail Transfer Protocol (SMTP) email to any Internet email server, except for those ranges provided for specifically by an ISP for that customer's use. The PBL helps networks enforce their Acceptable Use Policy for dynamic and non-message transfer agent (MTA) customer IP ranges.

For more information, see the Spamhaus Project website.

Directory-Based Edge Blocking

EOP typically processes all the messages that are sent to any SMTP address within your organization's domain. By synchronizing your domain's user list with EOP and enabling

Directory-Based Edge Blocking (DBEB), you can block email (even email that does not appear to be spam) sent to addresses that are not in your domain's user list.

EOP filtering technologies

After Directory-Edge-based Edge Blocking is complete, the mail flow now enters the post-edge, content-based filtering/blocking phase. After a majority of incoming spam is stopped at the edge, EOP uses the following technologies to inspect content and classify the remaining portion of incoming email:

- Multi-engine anti-malware scanning
- URL scanning
- Fingerprinting
- Sender Policy Framework (SPF)
- Anti-spam content filtering
- Bulk Mail Filtering
- International Spam Filtering

Multi-engine anti-malware scanning

EOP uses a layered approach—a minimum of three anti-malware engines—to offer protection from both known and unknown threats from both inbound and outbound email. These engines help protect against viruses and other email threats. They offer powerful heuristic detection to provide protection even during the early stages of a virus outbreak. The multi-engine approach has been proven much more effective than using just one engine.

The anti-malware engines scan the text within the body of each message as well as all attachments. They block any message body or attachment that contains the complete code for a known virus. They also attempt to block messages containing partial virus code. Some messages that have been cleaned by senders will be allowed to go through.

EOP checks for signature updates every 15 minutes and distributes the updates to all the anti-malware engines in the service's global network. During virus outbreaks, the EOP analysts also attempt to create policy rules manually to detect the threat before a signature is available from any of our anti-malware engines. These rules are published to the global network in real time to provide your organization with an extra layer of protection against attacks.

URL scanning

EOP uses URL block lists from multiple third parties, including Spamhaus, SURBL.org, URIBL.com, and Invaluement. When messages enter EOP, they are scanned against the lists of URLs from each of these data sources. If any message contains a URL that's on the list, it's marked as spam or given a weight that adds to the rest of the spam filter verdict, depending on the aggressiveness of the list.

Fingerprinting

When messages contain known spam characteristics, they are identified and fingerprinted: that is, they are given a unique ID based on their content. The fingerprinting database aggregates data from all spam blocked by EOP, which continuously improves and refines the filtering process as more messages are processed. If a message with a particular fingerprint associated with spam passes through EOP, the fingerprint is detected, and the message is marked as spam. EOP continuously analyzes incoming messages to determine new spamming methods. The spam analysts update the fingerprint layer as new attacks are detected.

Sender Policy Framework (SPF)

EOP evaluates the identity of the sender of each email message. If EOP can't authenticate a message, it determines that the sender is spoofed, and classifies the message as likely to be spam. EOP employs Sender Policy Framework (SPF), an industry standard that helps verify that the entity listed as the sender did indeed send the email message.

Anti-spam content filter (regular expression (regex)-based spam rules)

EOP assigns specific scores to messages based on more than tens of thousands of rules that define characteristics of legitimate mail versus spam. If a message contains characteristics of legitimate email, EOP subtracts spam points. If it contains characteristics of spam, EOP adds spam points. When a message's spam score reaches a defined threshold, EOP flags the message as spam and assigns a Spam Confidence Level, or SCL, value. SCL values of 5 through 9 indicate that the message is spam, while SCL values less than 5 indicate that the message isn't spam.

To score messages, EOP looks for:

- Suspicious phrases in the body and subject of the message, including URLs.
- HTTP obfuscation (disguising URLs that indicate a message is spam as legitimate URLs).
- Malformed headers (headers that have been incorrectly constructed).
- A suspicious email client type.
- Suspicious originating email servers.
- Suspicious originating email agents.

- Suspicious Message From and SMTP From addresses.

EOP modifies existing rules and adds new ones many times every day, through both automated analysis processes and manual intervention by the EOP spam analysts. In addition to the more than 20,000 rules that are used to identify messages as spam, EOP maintains more than 250,000 rules that collectively represent a substantial knowledge base about the characteristics of spam both past and present. Some of these rules are relatively straightforward, such as the following, which catches messages from a peculiarly creative phishing scam:

```
\bpoisoned (?:to death )?by his business associate.
```

Others are more complicated such as ones which Falsifies Microsoft Office Outlook Message-IDs on email messages generated by botnets.

EOP maintains a spam trap that siphons off some of the messages that would normally be blocked at the network edge, and tests them against the entire historical body of rules, including the rules that are currently in use. The EOP spam analysts use the resulting information to help determine when to retire active rules, and when to bring old rules back into use.

Based on this analysis, it is estimated that EOP would be able to identify and block 95 percent of all spam without using its edge-blocking layers.

Bulk mail filtering

Bulk email, also referred to as gray mail, is a type of email message that's more difficult to classify. Bulk email is typically comprised of an advertisement or marketing message that's not likely to get sent repeatedly. Bulk email is wanted by some users, and in fact they may have deliberately signed up to receive these messages, while other users may consider these types of messages to be spam. Admins can enable the Bulk mail advanced spam filtering (ASF) option that allows admins to mark all messages that EOP identifies as bulk as being high confidence spam. The service then performs the configured action, such as sending the message to the recipient's Junk Email folder. EOP also allows admin to aggressively control the bulk mail messages by using the Transport rules.

International spam filtering

You can configure EOP to block messages written in specific languages, or sent from specific countries or regions. You can configure up to 86 different languages and 250 different regions.

Mitigating false positives

The false positive rate for EOP —the ratio of legitimate email incorrectly classified as spam to all messages processed by the service —is extremely low. During a typical week, when EOP processes two to three billion email recipients, on average EOP spam analysts determine that about 900 messages are false positives.

This low false positive rate is not due to any particular piece of technology; instead, it's the natural outcome of the continuous review and adoption of new anti-spam technologies driven by the EOP engineering team. The team routinely tests new block lists, technologies, and algorithms by deploying them in an advisory-only capacity, and then compares the results from the new technology to the results of the existing suite. If the results look promising, the team changes the technology's weighting—the degree to which its verdict affects the designation of a message as spam—and gradually tunes that weighting to an optimal level. The team rejects new technologies that negatively affect the service's false positive rate and removes them. It does the same with older technologies whose effectiveness has decreased over time.

In addition to relying on these mitigation strategies, you can use the Exchange admin center (EAC) in the Office 365 admin portal to ensure that your organization's communications with critical business partners don't get accidently flagged as spam. Using the EAC, you can create policy (transport) rules to always allow email coming from certain specific partner addresses, domains, and IP addresses. Finally, you can use the Microsoft Azure Directory Synchronization Tool to synchronize your organization's Outlook and OWA safe sender's lists with EOP. The result is that email messages sent to an individual from a listed safe sender will be exempted from spam scanning, but not from virus scanning.

Conclusion

EOP implements a defense-in-depth approach to help protect your organization from spam and malware attacks for email. This approach makes use of block lists from third party vendors as well as EOP's proprietary block lists, combined with Directory-Based Edge Blocking techniques, Bulk Mail Protection , SPF , URL scanning and International spam filtering . The automated safeguards work in concert with continuous refinements made by a world-class team of spam analysts to provide powerful protection against unwanted or malicious email threats without impacting normal business communications.