mimecast™

Cyber Resilience for Email

TECHNICAL DEEP DIVE

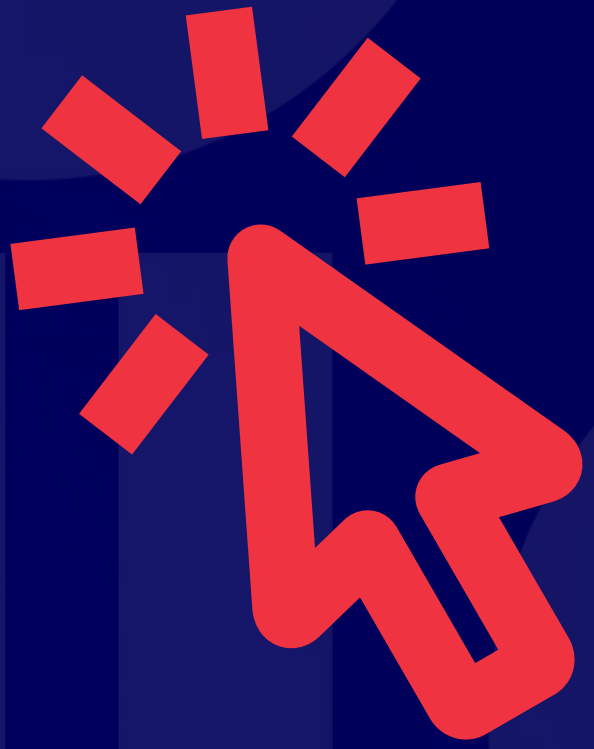

NASDAQ: MIME

Email. It's the number-one business application used by organizations.

It's also the number-one method used to execute cyberattacks, enabling malware delivery, phishing, impersonations, and the spread of threats that are already internal to your organization. In fact, 91 percent of all cyberattacks start with an email. And your organization can't function for long without email. How many hours of email downtime can your organization comfortably live with? If email isn't accessible due to an adverse incident like malicious intent, human error or technical failure, your organization would likely suffer from reputational damage, internal operational issues, and financial loss.

Meanwhile, the use of Microsoft Office 365 is massive, and adoption is accelerating. As organizations move to a cloud-based email environment, new challenges come along. The conce of corporate mailboxes, and the complete operational dependency on Microsoft exposes organizations to new risks.

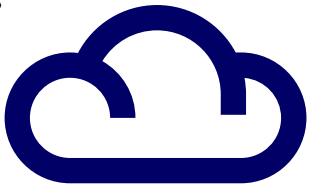

Email is at the intersection of a significant amount of risk for most organizations. If addressing this exposure doesn't become a priority, successful cyberattacks will continue and data protection and personal privacy will suffer.

Traditional security approaches are no longer enough. Attack methods are quickly evolving and growing more sophisticated, targeted and dangerous. Right now, the industry is faced with email-borne threats such as phishing attacks delivering malicious attachments and URLs; impersonation fraud fueled by social engineering and aimed at tricking employees into behaving badly; and ransomware attacks that can encrypt your data and take entire systems offline.

These are only the types of threats we know about today. What about the future? One example of an emerging attack technique is the use of homoglyph/homograph-based attacks to mask domains, slip by your security controls, and to fake out your users as part of a spear-phishing attack.

It doesn't stop there. There are more ways to exploit email that haven't been put into broad practice. Mimecast recently provided an example of a new attack type we named Ropemaker. Fortunately, we have yet to see this attack type in the wild! By using this exploit a malicious actor can change the displayed content of a delivered email at any time, post-delivery. This could mean swapping a benign URL with a malicious one in an email already delivered; turning simple text into a malicious URL; or editing any text in the body of an email, whenever the attacker wants to – and all of this can be done without direct access to an inbox – after delivery. The point is attackers are not standing still and so the defenders must not either!

It's Time for a New Approach

A defense-only security strategy is not sufficient to protect against this level and volume of advanced email-borne attacks. Continuing to invest in disparate technologies and focusing on a defense-only security strategy will lead to consequences like intellectual property and financial loss, unplanned downtime, decreased productivity and increased vulnerabilities. Legacy technologies can leave holes in your security and force you to chase tomorrow's attacks with yesterday's approaches. This also leads to additional cost and the need to find more of the right people to manage a complex security environment. It's no wonder so many organizations are struggling to keep pace.

The only way to get ahead of cybercriminals and to holistically protect your business is to adopt a new approach to email security. You need a multidimensional approach that brings together **threat protection, adaptability, durability and recoverability** in a single cloud-based service.

You need to enable these four dimensions to truly provide cyber resilience for your email.

Cyber Resilience for Email

A strategy which delivers cyber resilience for email empowers organizations to secure, preserve and continue the flow of communications via email.

This means preparing you for every stage of an attack:

1. Putting the right security controls in place BEFORE an attack happens – focused on prevention as well as those focused on quickly adapting to attacks techniques as they evolve.
2. A continuity plan to keep email – and your business operations dependent on email – running DURING an attack or failure.
3. The ability to recover data and other corporate IP AFTER an incident or attack occurs.

The Four Dimensions of Cyber Resilience for Email from Mimecast

Our cloud-based system helps organizations prevent email-borne cyberattacks; keeps email flowing, business operations running and employees productive during downtime; and enables the recovery of lost or locked data after an attack happens.

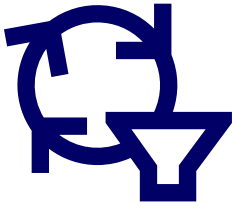

Threat Protection

Introducing the Mimecast Email Security Inspection Pipeline:

Before getting into the weeds of how the Mimecast email security service works, first take a glance at figure 1 on the next page. Figure 1 is a graphical representation of the email security inspection pipeline of the Mimecast service in its entirety. What jumps out at you? The fact that so many analytic steps are being applied in an instant across hundreds of millions of emails a day? Or perhaps the breakdown of the analysis into specialized inspection pipelines – with attachment inspection and URL inspection handled in their own streams? Or the many different types and total number of analytics that need to be applied to produce safe emails? Or perhaps the overall funnel effect – which always starts with high level, more general-purpose inspections applied to the emails before the analysis moves on to deeper, more sophisticated inspections further down in their respective funnels?

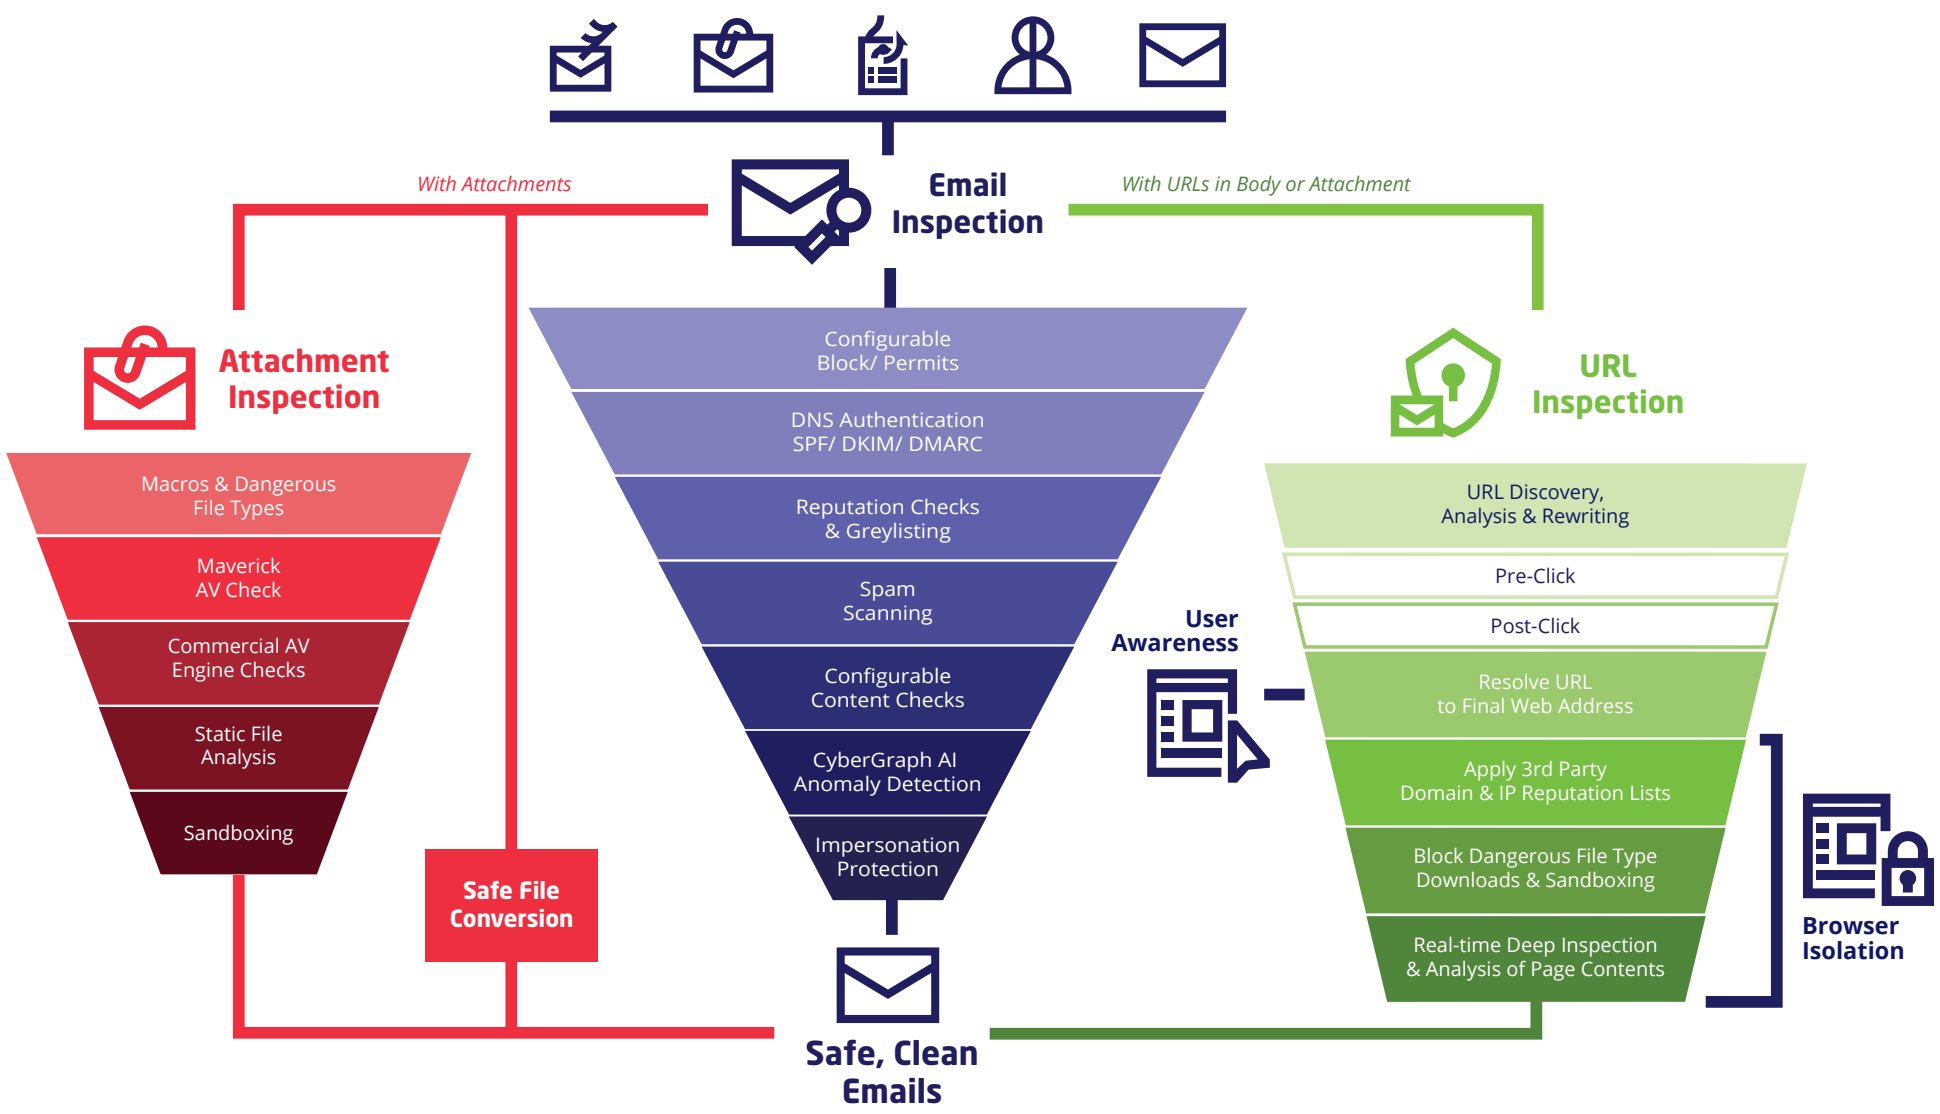

Figure 1: The Mimecast Threat Inspection Pipeline

Of course, the most logical answer is all-of-the-above and more! The beauty of the Mimecast email security cloud service is that the inspections can be both deep and wide, without significantly impacting the speed of delivery of the emails. Before an inbound email ever makes it to your organization's email system, whether your email system is on-premises or in the cloud, it goes through many layers of inspection and analysis to detect the ways phishers and spammers try to get to and fool your email security system and your users.

Could your organization setup, maintain, and improve an equivalent email security system on your own? Very unlikely, even if your organization had the most sophisticated security teams in the world with nearly unlimited budget. Isn't that the point of subscribing to a specialized cloud service in the first place? We go deep so you don't have to?

Can other email security cloud service providers offer the sophistication of inspection that Mimecast does? This is also very hard for them to accomplish. To do so requires a level of focus, speed, innovation, and scale that is beyond the capabilities of most. Also without a cloud-based architecture that is like our MIME|OS (described below), other cloud providers, particularly those using hosted VM-based solutions that were originally built for on-premises deployments, are severely constrained as to the speed of development and integration that they can deliver.

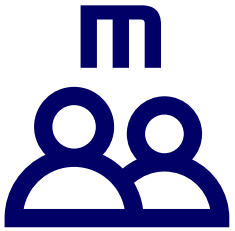

The Mimecast SOC

When you think about the Mimecast email security service, you might think of our distributed data centers and some of the analytics and services that must be provided to detect and block the many forms of spam and phishing attacks that plague the Internet. While these do form the foundation of our cloud service, would it surprise you to learn that the true brains and value of the Mimecast service are our people? While all the 1000+ people who work at Mimecast contribute to the value of our service, the group most responsible for the day-to-day security efficacy of it are the members of the Mimecast Security Operations Center (MSOC).

The MSOC is a team of globally distributed analysts and security researchers that tend to the Mimecast service on a 24 x 7 basis. They effectively are our customer's email security focused operations center in the cloud! For example, when customers submit suspect phishes and spam emails to Mimecast – which number approximately 2500/day – the MSOC team is the group that analyzes them and in response makes any needed changes to improve the service.

The MSOC responsibilities include:

- Investigating & notifying customers of likely compromises
- Removing organizations from email blacklists
- Creating and deploying updated threat detection signatures
- Handling customer generated security related escalations
- Collaborating with 3rd-parties such as ISACs, ISPs, Hosting Providers, Domain Registrars, & Law Enforcement

The MSOC is the team at Mimecast that is most responsible for staying ahead of attackers and how they are attacking organizations with new spam, phishing, and malware distribution campaigns. To stay ahead of the attackers and spammers the MSOC team also continuously reviews various threat intelligence sources. For example, leveraging our membership in the IT-ISAC and the Anti-Phishing Working Group, multiple threat news sources, commercial threat intelligence providers as well as conducting analysis of the data generated from the Mimecast inspection pipeline itself. We learn a lot from analyzing billions of emails monthly!

In addition the Labs portion of the MSOC team regularly conducts research in the following areas:

- Malware behavioral analysis
- Malware static analysis
- AV engine optimization and development
- Analysis and extraction of threat intelligence from the Mimecast Grid data

The MSOC team applies their learnings and expertise on a continuous basis into the Mimecast email security inspection pipeline depicted in Figure 1. In addition, the MSOC team is responsible, in collaboration with Mimecast Engineering, for evaluating and re-evaluating existing and new analytic techniques and threat data feeds for inclusion or removal from the Mimecast security service. Doing this on a continuous, round-the-clock, round-the-globe basis ensures that the Mimecast email security service delivers a very high level of security and efficiency to our customers.

The combination of the MSOC team and the Mimecast Email Security Inspection Pipeline represented in Figure 1 cannot be easily matched by an enterprise attempting to "do-it-yourself" or even another cloud security provider. Why this is should become increasingly clear by reading the remainder of this paper.

Overview of the Mimecast Email Security Inspection Pipeline

The graphic in Figure 1 represents the Mimecast Email Security Inspection Pipeline in its totality. Inbound at the top of the funnel, represented graphically by the mix of good and bad emails, are those inbound emails that need to be inspected and filtered by the Mimecast security service. Given that consistently more than 60% of inbound emails are either spam or many types of malicious email, most of these inbound emails are "bad". It, of course, is the job of the Mimecast service to figure out which-is-which. To do this the Mimecast inspection pipeline is divided into three main inspection sub-funnels.

In the center of Figure 1 is a graphic of the main inspection pipeline that is presented in greater focus in Figure 2. In this central pipeline, security analysis steps run the gambit from high-level inspections that check to be sure first that the sender should not be blocked by configuration, or is attempting to spoof a legitimate sender's domain, has a poor reputation, or is attempting to deliver "spammy" emails. After these higher-level checks are completed, the central inspection pipeline moves on to more in-depth inspections which check for sensitive or undesirable content that the organization doesn't want to come into their organization. This is also the point, where more specific checks branch off, to assess the security of any attached files or included URLs. Finally, at the bottom of the center inspection pipeline a series of analysis are applied that are adept at detecting signs of sender impersonations.

A typical technique used by attackers when they use email is to attach malicious files directly to an email with the hope that the email security system will allow it to be delivered and opened by the user, thus causing many different types of infections, including ransomware and trojans. These file-based email attacks are addressed by the Attachment Inspection leg of the Mimecast Email Inspection Pipeline, as shown in Figure 1 and Figure 3. The purpose of this leg of the Mimecast Email Security Pipeline is to protect the organization from malicious file attachments. Because there are many ways that attackers will attempt to sneak malicious files past an email security system - using macros in Office files or packing or "crypting" files to get past signature based AV engines - it stands to reason that the email security system must use multiple detection techniques including the latest innovation, static file analysis, to protect organizations from these malicious attachments.

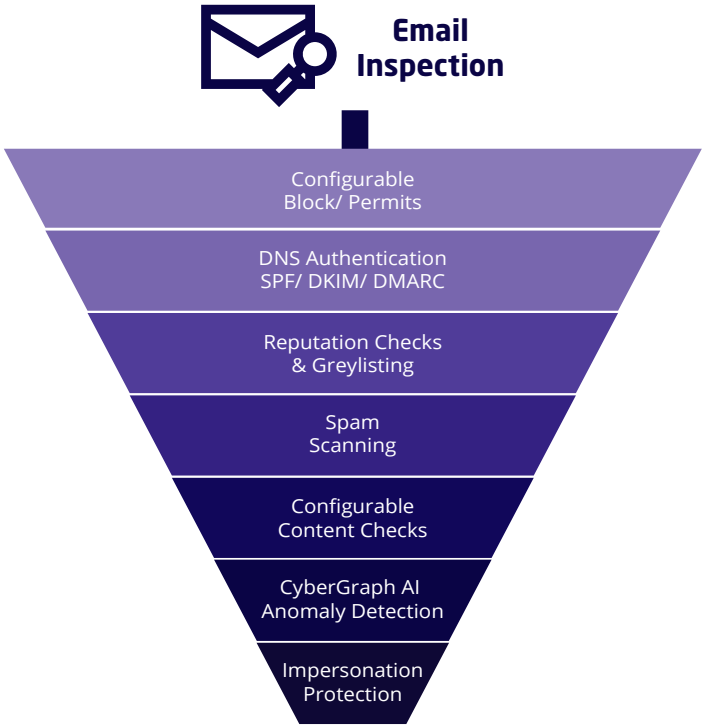

Figure 2: The Mimecast Central Email Inspection Pipeline

The Mimecast Attachment Inspection Pipeline protects organizations from malicious attachments by applying six different techniques. The details of these techniques will be described in more depth on page 13.

The third leg of the Mimecast URL Inspection Pipeline, as seen in Figure 1 and Figure 4, is focused on protecting against attacks that use malicious URLs as a key component of the email-borne attack. A malicious URL's purpose is to bring an unsuspecting user to a phishing web page, where he will be duped out of his credentials or other sensitive information, or alternatively to directly initiate the download of a malicious file to the user's machine.
Often these phishing websites are made with phish kits on the back-end, a collection of tools designed to make launching a phishing campaign easier, especially for those who do not have a lot of technical skill. Defending against malicious URLs also requires multiple analytic techniques to be applied starting with discovering the URLs in an email or in an attachment, to conducting in-depth inspections and analysis of the pages and downloadable files in real-time at the time of the click. With the help of our ClamAV engine and the ThreatLabs team, in just four months Mimecast has been able to identify and create over 60 phish kit signatures, resulting in more than 3.4 million detections.

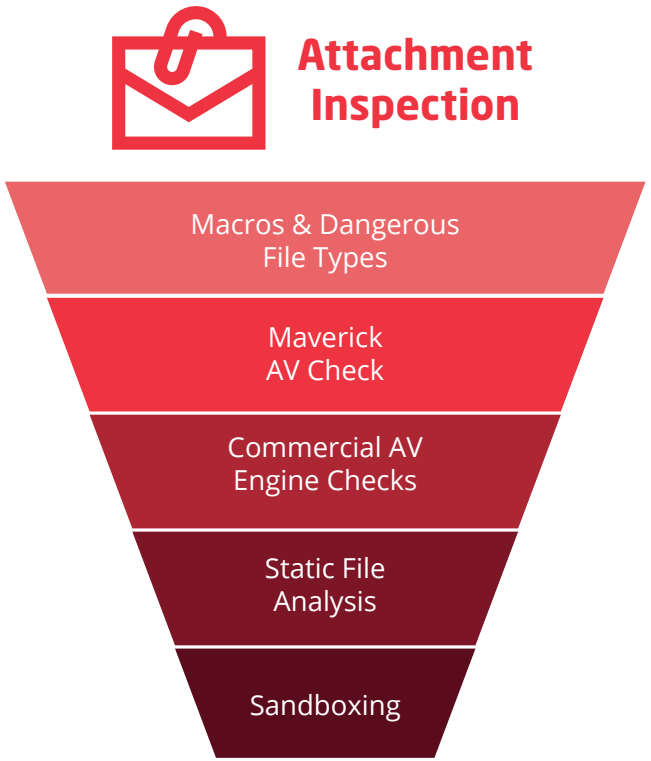

Figure 3:The Mimecast Attachment Inspection Pipeline

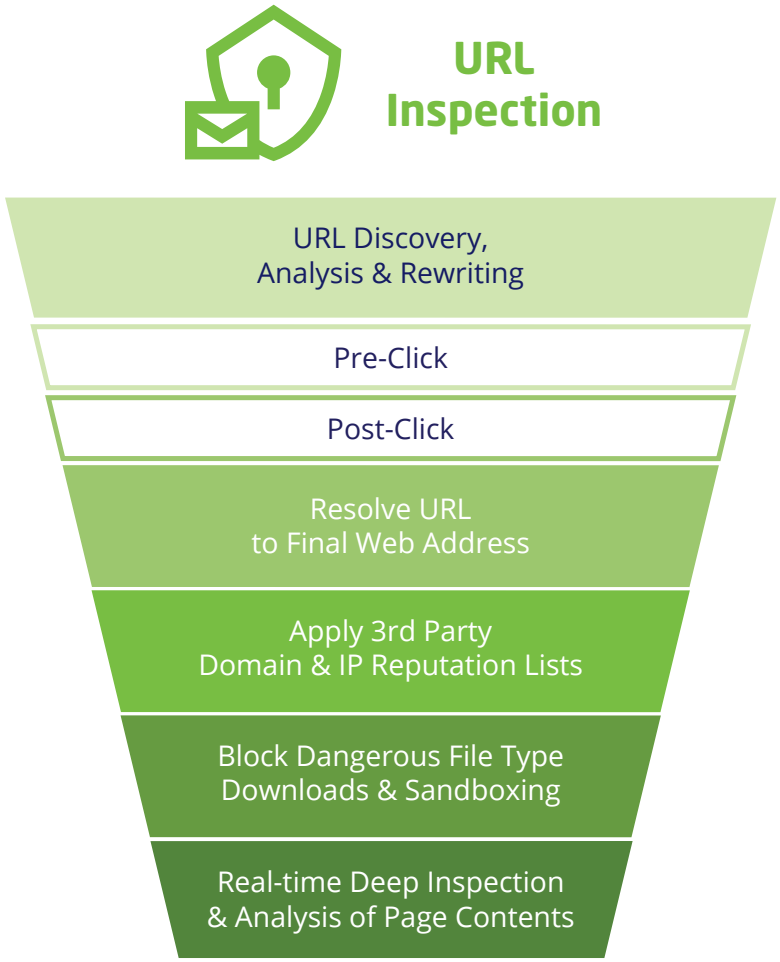

Figure 4: The Mimecast URL Inspection Pipeline

The Central Email Inspection Pipeline - Deep Dive

Now that the overall flow of the Mimecast Email Security Inspection Pipeline has been introduced, it is time to go a level deeper into the central branch of the pipeline. Please refer again to Figure 2, as in this section of the paper we will describe each step of the central portion of the inspection pipeline.

Configurable Block/Permits

Once the initial SMTP communications channel is set up between the sending email system and the Mimecast service, ideally using TLS for encryption, the actual security related inspections can begin.
At this first inspection stage, automated as well as Mimecast and customer configured block and permit lists are applied. For example, spammers newly discovered by the MSOC team, but not yet known by other reputation sources, are blocked at this stage of the pipeline. In addition, customer administrators can add sender block lists of their own that only apply to their tenancy.

Most commonly customers will add their own email domains to their tenant's block list, as there is typically no reason to accept mail from external senders who are pretending to be sending email from and to an email domain that is owned by the customer. Those can be blocked straight off without further consideration. This is part of the default configuration that is setup for all customers during a typical implementation.

This is also the step of the inspection pipeline where spam white lists are applied. For example, senders to whom the customer organization has a history of sending emails are automatically whitelisted at this stage of the inspection pipeline. The logic of this is straightforward. If one's users are sending email to users at a given domain, then it can be assumed that this email address is not under the control of a spammer. This is a good way to avoid spam blocking related false positives.

Possible, although rarely used, the MSOC team can also whitelist senders that are being improperly flagged as spammers by spam feeds in use by Mimecast.

It is important to note that in all cases the whitelisting of a sender only bypasses the IP reputation and spam scanning steps, but not other steps in the Mimecast Email Security Inspection Pipeline, such as DNS Authentication. Therefore, email from legitimate senders that at some point get hijacked by attackers still receive malware, URL, and other non-spam focused security inspections before delivery is allowed, regardless of Mimecast or customer managed whitelisting.

DNS Authentication - SPF, DKIM, & DMARC

In the early days of the Internet it isn't far from the truth to say that there was no built-in security for email. For example, in those days any sending email server could claim to be representing whatever domain they wanted to. If you claimed to be sending email on behalf of ABC.com then nothing on the Internet double-checked if this was a legitimate claim. At the time, there were no standards-based ways to prove that this wasn't true. Into this environment came in rapid succession (in the early 2000s) two Internet security standards that were invented to address this oversight – SPF and DKIM.

SPF – Sender Policy Framework – is an email validation system that was designed to detect mail spoofing by providing a system which enables the receiving email system to check the DNS entry of the sender to see if an authorized host (IP address) is sending on behalf of that domain. If it is some unlisted host IP address sending on their "behalf", then it is likely a malicious sender that is spoofing that organization.

DKIM – DomainKeys Identified Mail – is a standardized method of cryptographically signing an email to both confirm that it actually came from the indicated domain as well as hasn't been tampered with as it traversed the Internet. It does this by affixing a digital signature to the message, the validity of which can be confirmed by using the public key that is contained in the sender's DNS entry.

Finally, rounding out the email focused security standards, much more recently DMARC – Domain-based Message Authentication, Reporting and Conformance – was added to the email security mix on top of SPF and DKIM. In short DMARC enables the owner of the sending domain to publicly publish a policy regarding how a receiving domain should react (primarily to reject or not) to the failure of DKIM and SPF checks for their domain.

It also provides a reporting mechanism enabling a receiving domain to report back to the sending domain regarding what is happening with their domain. For example, if some sender is potentially spoofing their domain.

From the point of view of the Mimecast Email Security Inspection Pipeline, this is the stage of the process where SPF, DKIM, and DMARC policies are enforced, leading to rejecting, quarantining, or passing inbound emails to the next stage in the inspection pipeline. The desired behavior is completely under the control of the customer's Mimecast administrator.

While security standards are great, even if they were widely deployed and used by both sender and receiver (which they often are not), attackers can also use and get around them. For example, by using a cousin domain - a domain that is similar to, but not exactly the same as a domain that they are trying to spoof - to send from. This is why more impersonation specific controls are needed.

Reputation Checks & Greylisting

At this stage of the central leg of the Mimecast Email Security Inspection Pipeline the Mimecast service is still analyzing whether or not to receive email from the sender that is "knocking" on the Mimecast email gateway's door. To provide the reputation related analysis that is applied at this stage of the inspection pipeline Mimecast uses a series of IP and domain reputation checks as well as a greylisting process that can also defeat many spammers. For IP and domain reputation checks the analysis at this stage leverages both 3rd-party provided reputation lists and services (including at the time of this writing: Spamhaus IP, Domain, New Domains; as well as a few other commercially available offerings such as Cloudmark, Spamcop, Invaluement, and Cyren) as well as a reputation list that is maintained by the MSOC team. Any sender that is blacklisted by any of these sources, that isn't otherwise whitelisted, will be rejected at this inspection stage. Key is, Mimecast does not rely on any single provider for spam or other email processing. Multiple sources greatly improves the efficacy of the Mimecast service.

For previously unseen combinations of a sender's IP, address, and recipient addresses, Mimecast applies a greylisting process that temporarily defers the accept/reject decision to a later time and asks the sender to resend.

As a full SMTP Mail Transfer Agent (MTA) the Mimecast service can interact with mail senders as any email management system would, however for example, an API-based email security system (an email security system which bolts on to the email management system as opposed to functioning as a gateway service) cannot. In many cases a sender of unwanted email – such as a spammer or botnet – is not willing or able to interact in this way and thus never attempts to resend, thus ending the inspection pipeline for that batch of email at this early stage in the pipeline.

Spam Scanning

Only after reaching this spam scanning layer, after having passed through the three previous inspection layers, does the actual content of the email get analyzed for spam. The engine inside the Mimecast service which manages the spam analysis process is based on the open-source spam filtering system, Rspamd.

Rspamd, which has been specially customized by Mimecast in close collaboration with its lead developers, is an advanced spam filtering system that enables the evaluation of messages using rules, including regular expressions, statistical analysis, and other custom services. Each email is analyzed by the Rspamd engine and given a spam score. The level of the spam score determines whether the message will be rejected, quarantined for manual review, or passed through to the next step in the inspection pipeline.

As inputs to the Rspamd engine's managed spam score Mimecast uses a combination of open source and commercial spam analysis services and threat feeds (including at the time of this writing: Spamhaus as well as multiple other commercially available offerings such as BitDefender, Cloudmark, Invaluement, URIBL, and SURBL). In addition, Mimecast also provides proprietary content analysis techniques to the mix (including at the time of this writing: Heuristic Filtering, Fuzzy Content Matching, Content Decoding, Static/Dynamic Content Lists, Passive DNS, 3rd-party threat feeds).

In addition, the MSOC team continuously evaluates customer submitted suspect spam messages to guide the adjustment and further development of the above-mentioned services, analytic techniques, and spam scoring values.

New and existing analytic techniques, spam service providers, and data sources are evaluated on a near continuous basis for inclusion into or exclusion from the spam scanning inspection layer. Similar to the Reputation layer above, doing this ongoing evaluation is beyond the scope of most other email security vendors and systems, whether they are on-premises or cloud based.

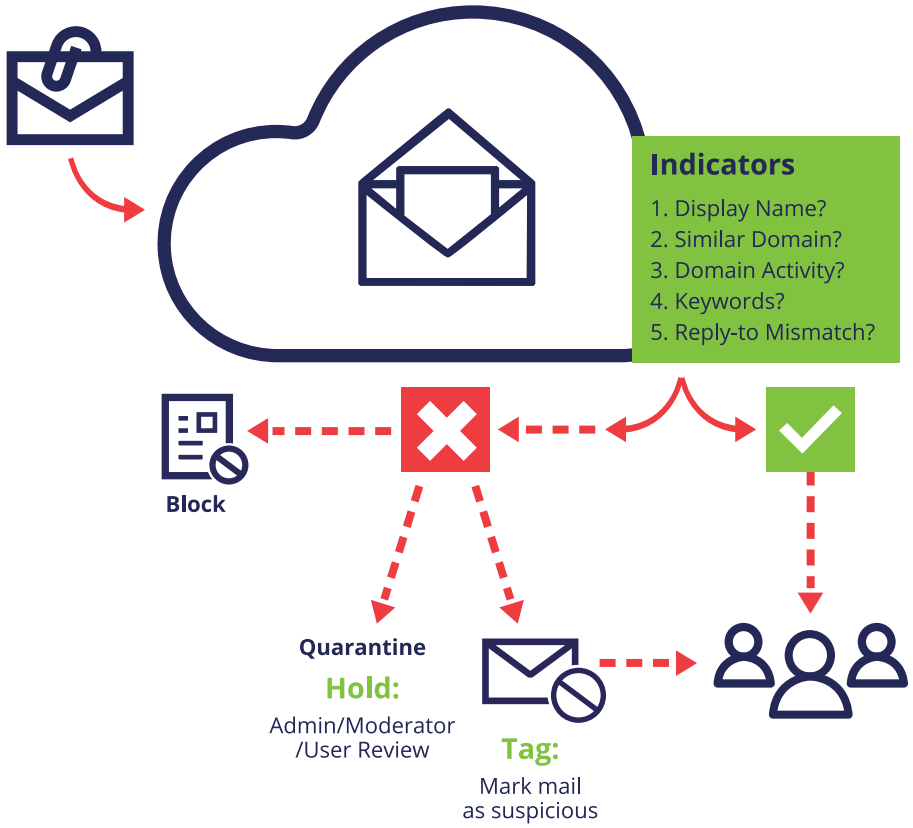

Figure 5: High-Level View of Impersonation Protection

Configurable Content Checks

At this step in the central inspection pipeline, tenant specific content checks are applied. These custom content matching steps are customer configurable. The content examination process includes file document content fuzzy hash fingerprinting that provides the ability to detect partial content matches, as well as the examination of the body text of an email, attachments, headers, and subject lines. In addition, the service includes built-in dictionaries (such as healthcare related data) and the ability to recognize structured data such as credit card or social security numbers in emails. A policy hit based on the above analysis can lead to the message being held, rejected, or delivered with a copy sent to another user.

At this layer of the inspection pipeline, the system can reject, hold, or strip & link, and notify the customer administrator when encrypted or unreadable attachments or archives are being sent into the organization.

Mimecast CyberGraph

Mimecast CyberGraph is an optional add-on that combines three capabilities:

1. Email tracker protection to limit an attacker's intelligence gathering capabilities
2. Identity graph and machine learning technologies to detect advanced phishing attacks
3. Dynamic, contextual banners embedded in emails to warn users that they might be suspicious

Email Tracker Protection

Email trackers can be embedded in an email and when it is opened, they pull information from a remote server, such as a single pixel graphic that is invisible to the human eye. The act of downloading it means the device IP address and therefore location can be discovered.

Email clients often block image downloads so other trackers can be embedded such as special web fonts or streaming media that is inaudible. The media will stream for as long as the email is open, so it helps an attacker understand the recipient's engagement level with the email content and whether they forwarded it.

Finally, the user agent string sent with the requests gives the attacker information about the browser and device OS versions. They now know whether this might be an old OS version that contains vulnerabilities that could be exploited.

CyberGraph prevents trackers in emails reporting any useful intelligence back to the attacker. It uses machine learning to understand the structure of tracker URLs, downloads the target content on first scan, and saves it into a content delivery network. The tracker URLs embedded in the original emails are then rewritten to point at this content.

This results in all requests from a recipient's email being made to Mimecast and we deliver the content already downloaded. This shields the recipient's location, engagement levels and those of anyone the email is forwarded to.

This helps prevent the attacker gathering some of that essential information that they could use to craft an extremely authentic spear phishing email. In fact, they would not even know whether the target received and opened the email.

CyberGraph Artificial Intelligence

The second piece of the solution is the CyberGraph machine learning model and identity graph. The identity graph stores information about relationships and connections between all senders and recipients, including the strength or proximity of the relationships. It learns about behaviors with respect to what is normal and can then detect anomalous behaviors that might be indicative of a malicious email.

The identity graph output is combined with other indicators of suspicion and outputs from the machine learning model. There are tens of indicators of suspicion and over 100 features of an email that are used by the machine learning model. Together these are used to determine whether an email might be malicious.

Here is an example of a combination that would be flagged as suspicious:

- The recipient has never received an email from the sender
- The sender is using a newly created domain
- The email subject looks like those learned from previous malicious emails

Contextual, Dynamic Warning Banners

The output of the above is a color-coded warning banner that is embedded in the email. The combination of technologies ensures that they are only added when there is a high likelihood that the email is malicious. The banner is an image so is always displayed regardless of the device type or email client. Different coloured banners indicate the level of risk, and they provide enough information about the potential issue identified without over burdening the user with too much technical detail.

There are around 20 different indicators to help users and the wording of each can be tailored to suit each customer's staff skills or technical knowledge levels. These banners engage the user at the point of risk, when they are about to action the email.

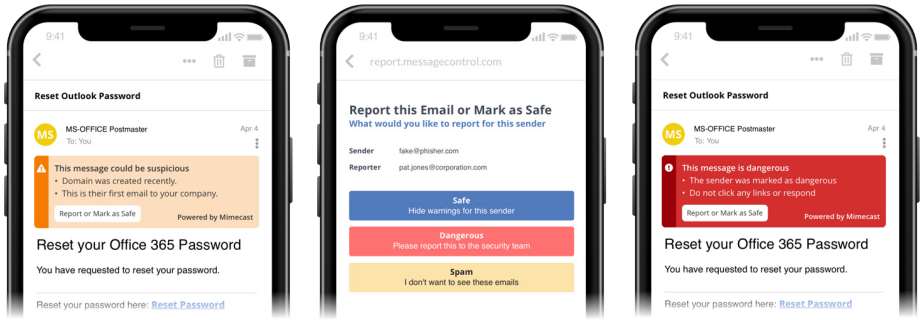

The user is further empowered by having their view solicited, of whether the email is malicious. These reports feed the machine learning model and update the identity graph with information about the trust relationship between senders and recipients. The information can also feed Mimecast's threat intelligence and help determine whether that email is malicious or not. This is essentially crowd sourcing data about malicious emails.

Reported emails are analyzed, and the banners on similar emails can be automatically changed to reflect its new status. E.g. When a user reports an email as malicious, banners on all emails from that sender, across a customer's estate, can be updated to indicate that a user has reported it and extra care should be taken.

Once analyzed, the indicators of suspicion reported in the email banner can also be updated to reflect its new status. This is possible because the banner is not embedded in the email; instead, a link is embedded that points at the banner, so we can change the banner at any time for any given delivered email, and each time the user opens the email it downloads the banner assigned to it.

Reporting and alerting

The CyberGraph tracker dashboard provides information about the severity of trackers – whether they were bulk marketing or targeted trackers. It shows the physical location of every employee when they open emails that had trackers embedded and reports on their device type and OS. This brings a number of benefits for the IT and cyber security teams:

- Report on which individuals are being targeted – they might inform these people to ensure they remain vigilant or maybe schedule additional awareness training for them.
- Visibility of device/OS versions in the estate that might contain vulnerabilities that could be exploited and need upgrading.

How it works

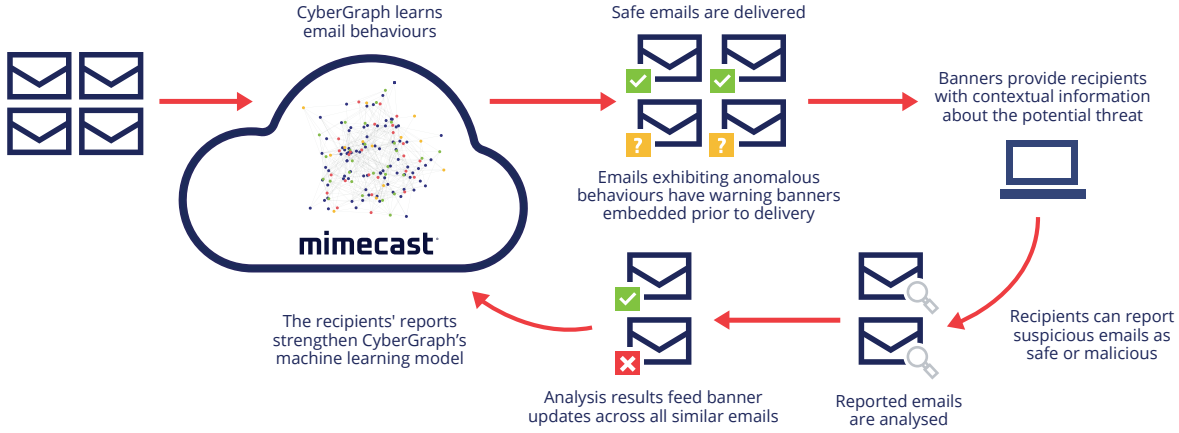

Impersonation Protection

The final step of the central inspection pipeline depicted in Figure 2 is there to determine if a sender is attempting to impersonate another person within the customer's organization or alternatively is trying to impersonate a partner, customer, or service that the receiver would naturally trust. Sometimes referred to as CEO fraud, impersonation, whaling, or business email compromise types of email-borne attacks, these spear-phishing attacks don't necessarily use malicious URLs or attachments to attack the organization.

They are designed to trick the email receiver and push them into making wire transfers or to unwittingly respond back with sensitive data. A malicious sender is pretending, using various technical means (such as cousin domains, hiding the true sender's email address, using homoglyphs, or long domain strings) combined with social engineering techniques, to impersonate the email of a trusted person or organization.

How does impersonation protection work?

As can be seen in Figure 5 below, as the email passes through the Mimecast Gateway, the Impersonation Protection layer examines several key aspects of the message. It examines the inbound email's display name to see if there is a match with an internal user, the domain name similarity (to those of well-known Internet brands, supply-chain partners of the customer or their own domains), whether the organization has previously received email from that domain, reply-to information, and the content in the body of the message to determine if the email could be an impersonation attack. If the email fails a combination of these tests, administrators can configure multiple responses, such as discard the message, quarantine it, or to deliver and warn the receiver with customizable tagging to warn the receiver to take extra care.

Detecting Character switching, Homoglyph/Homograph,& long domain strings

More recently Mimecast has added the detection of character switching, homoglyph/homograph, and long domain strings to the existing impersonation protection capabilities. Also, Mimecast has extended the detection beyond just domains owned by the customer, to include the domains of customers and partners of the organization as well as well known Internet domains (such as Ebay, Paypal, Google etc…).

See Figure 6 below for examples of the types of domain spoofing that is caught with this new capability.

Real Domain	Similarity Match
mlmecast.com	mlmecast.co.za
porsche.co.za	locator.porsche.com
ticsante.fr	xn--ticsant-hya.com
flashmobile.mx	uplus.flashmobile.kr
hotmail.com	www.xn--hotmail-vfb.com
apple.com	xn--80ak6aa92e.com
amazon.co.uk	www.amazonn.co.uk
hotmail.com	hotmaill.com
facebook.com	http://m.facebook.com----------securelogin.liraon.com/sign_in.htm
comcast.net	http://login.comcast.net----------securelogin.giftcardisrael.com/
facebook.com	http://m.facebook.com----------terms-of-service-agree.madkoffee.com/
apple.com	http://apple.com------support.host/
paypal.com	http://paypal.com-us-cgi-bin-webscr-cmd--login-submit-dispatch-5885d80a13c0.mytruebox.com/

Figure 6: Examples of external domain spoofing that are detected within impersonation protection

Message Insights

In order to ensure our customers are provided with a high level of visibility into each and every message that reaches their environment, Mimecast provides granular insight into why messages were classified as spam via Message Insights.

Message Insights provides details for administrators regarding how individual messages were processed. By equipping our customers with full transparency into the Mimecast inspection process, administrators can see which characteristics of an email led to the message being handled the way that it was. With this knowledge, customers are able to make the most informed decisions about which messages are allowed into their environment and what changes they may need to make to their systems.

The details that accompany each email include the message's overall evaluation, and characteristics such as its Spam Score & Spam Detection Level. The Spam Score & Detection Level are a message-centric view designed for busy administrators trying to evaluate why an email was held by a spam policy so they can make an informed next action and respond to user requests. This can be used to assess whether Spam Detection Thresholds need adjusting, and API access allows for the programmatic reception of more details about specific emails.

In Message Details, customers will see Processing Details, which include a more in-depth look into the evaluation made at each stage of the inspection funnel to see which setting took effect and led to the email's final valuation.

The Processing Details consist of:

- Graymail Result
- Managed Sender
- Permitted Sender
- SPF Result
- DKIM Result
- DMARC Result
- RBL Result

This information allows administrators to understand why certain policies were triggered for a message, most particularly when it's an outcome they were not anticipating (i.e. sender is added to their whitelist but the message is still getting caught up in spam).

The Mimecast Attachment Inspection Pipeline - Deep Dive

After the inbound emails have passed through the central inspection pipeline down through the spam scanning step, the analysis can be extended in two ways depending on the particulars of the individual emails. Emails which contain URLs must be inspected for the existence of malicious URLs. The details of the URL Inspection branch will be covered in its own individual section below.

Those emails which contain attached files must be inspected for malware. Malware, or malicious files, can take many forms and can run the gambit from broadly distributed, commodity malware at one end of the threat spectrum, to highly targeted, even custom-built varieties that are created and targeted at even just a single organization, at the other end of the spectrum.

The Attachment Inspection branch of the overall pipeline is designed specifically to detect and block malware, no matter whether it is broadly distributed or of the highly targeted variety. The Attachment Inspection pipeline branch can be seen in Figure 3. Note that it is made up primarily of five specific malware analytic steps, starting as always with the most high-level checks and ending up with the most in-depth checks at the bottom of the pipeline.

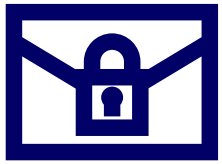

In the case of the Attachment Inspection process there is one unique pre-file-analysis step – Safe File Conversion - which dramatically increases the speed of the file delivery and security efficacy, which can be enabled for all or just a portion of an organizations' user community.

Safe File Conversion

Safe File Conversion is a customer configurable setting which converts Microsoft Office and PDF files to a safe file format for immediate delivery to users in the customer's organization. Files which are converted using Safe File don't go further in the Attachment Inspection analysis process unless the original file is requested by the user. A safe file is a file format that is readable but that has had any active content, such as macros or executables, removed from the original file before delivery. This conversion process makes it all but certain that the converted file is stripped of any malicious content.

A typical safe file conversion occurs when, for example, a Word document is automatically converted into a read-only PDF document. This removes the ability for any onboard Word macros to execute and thus, for example, initiate the dropping of a malicious payload from the Internet. Administrators can also choose their preferred conversion method. Choosing to convert the original file to a read-only PDF or to keep the original file type, for example, a Word document, but deliver it with macros and extraneous code stripped out.

The key advantage of using Safe File conversion is that an email with such an attachment that is converted to a safe file format is delivered immediately without any malware analysis delay. Given that in most cases users don't need to edit a file, just read it, the use of the Safe File Conversion step meets the needs of most users, while fully protecting the organization. Of course, users that do require the original file can simply request it by pushing a link in the delivered email, which will cause the original file to be sent through the remaining Attachment Inspection steps.

In addition, to provide more user control, users can flag senders as senders from whom they would always like to receive original files and thus have files from them go directly into the Attachment Inspection pipeline versus be converted using the Safe File Conversion process.

Macros & Dangerous File Types

The first step in the analysis of a file in the Attachment Inspection pipeline is to check whether the attached file(s) contains macros – if those are blocked by policy by the organization – or if the attached file is one of approximately 350 file types that are rarely sent via email for legitimate purposes, such as: jsp, .exe., .src.. Mimecast, at this point in the Attachment Inspection funnel, blocks these types of attachments. It is important to note that the analysis in this step does not rely on the declared file type to make this determination, it determines the file type by inspecting the file itself.

Maverick AV Check

The next step in the file inspection process that the email must traverse is the check of the files' hash values against the Mimecast Maverick global malware database. The Maverick database is a Mimecast developed and maintained file hash database of known bad and good files that Mimecast has seen across our tens-of-thousands of customers and billions of monthly emails. The Maverick database is a distributed database that is deployed across the Mimecast global cloud infrastructure; known as the Mimecast Grid. The Maverick database is kept up-to-date with simultaneous, up-to-the-second updates. Every file received by Mimecast at this step in the pipeline is cryptographically hashed and then compared against the local Maverick AV database. If there is a match the file is flagged as malware and is removed from the email. And conversely files that are determined to be malicious but that are found on the existing "good" list are dealt with as files needing remediation as part of the Mimecast Internal Email Protect service (discussed later in this paper).

The Maverick database itself is fed by the ongoing file analysis that occurs across the entire global Mimecast customer base, including the later stage results of the Attachment Inspection pipeline (static file analysis and sandboxing) 3rd-party malware intelligence sources, as well as by the MSOC team as they analyze customer submitted samples and reports from the global malware research community. The use of Maverick greatly improves the speed of analysis and efficacy of detection of email-borne malware. Catching malicious emails "higher-up-the-funnel" is always the goal for performance and efficacy reasons!

BYO Threat Intelligence

Many organisations have large amounts of security data, sourced from many best of breed point products across their ecosystem, and to help detect targeted attacks, they often license industry specific data to enrich these sources. BYO Threat Intelligence provides automatic ingestion of this threat data directly into a private area of the Maverick database within each customer tenant on the Mimecast platform.

It is enabled by the Mime|OS platform's rich API support (see API section later in this paper). Automating data ingestion enables streamlined response processes, increased email threat detection efficacy and quicker estate-wide time to protection.

Commercial AV Engine Checks

One of the key principles of effective cybersecurity is that relying on any one system, threat source, or analytic technique to detect attacks in general and malicious files in particular, is a doomed strategy. Attackers can work around any individual defensive technique. What they find very difficult to circumvent is when multiple strategies or systems are applied at the same time and that are constantly changing outside their visibility or knowledge. This multi-faceted, multi-layered analytical approach can be seen in and across many stages of the Mimecast inspection pipeline, including this one.

Instead of relying on just one commercial AV engine to detect malware, at this stage Mimecast applies multiple commercial AV engines (currently three at the time of this writing). The MSOC team, as is its practice, constantly reviews these and other prospective commercial AV engines for efficiency and efficacy and thus for addition or removal from this inspection stage. The reality is that the effectiveness of any single AV engine against specific malware threats at a time or over time, varies, and thus the use of multiple engines is a clear best practice. A best practice that most enterprises and email security cloud providers do not follow.

Static File Analysis

Files that make it this far in the Attachment Inspection Pipeline are clearly benign, right? Unfortunately, reality is a bit different! AV engines by the very nature of their signature-based approach are inherently historical or blacklist focused. They are very good at catching "known bad" files, but not as good when facing "unknown bad" files. When informed that a certain file is bad, from that point forward they can detect the arrival of that particular file quite efficiently. However, what if attackers customize, create new, or obfuscate malware files right before sending them to their intended victims?

In this situation, traditional AV engines are very unlikely to detect those malicious files. How can these new instances of malware be detected?

To address this challenge, next up in the Attachment Inspection branch is Static File Analysis, which leverages capabilities from a product from the Static File Analysis specialist, Solebit. With Static File Analysis, the file itself is evaluated, without executing it, to determine if it shows traits consistent with malware. Such traits include built-in calls to known command-and-control sites, analysis of the file to assess if it has been obfuscated, crypted, or packed in any way, whether it includes abnormal code or structures, embeds suspicious objects, includes code portions from known malware generating tools, enables potentially malicious file linking (such as linking to a file on an external server), and what compile time timestamp exists, to name just a few techniques that are applied at this step. Files that score poorly because of this Static File Analysis stage, as with the previous stages, don't need to proceed further. One of the key advantages of this Static File Analysis step is its speed of analysis, often completing in just one or two seconds.

Sandboxing

The final stop on the Attachment Inspection pipeline is where full behavior-based file sandboxing is applied. In general file sandboxing is a form of malware analysis that consists of opening and running a given unknown file to determine if it shows behavior that is indicative of malware. Sandboxing a file for malware is highly accurate but is also the most resource intensive and thus time-consuming step for analyzing a file. Therefore, it is the final step in the Attachment Inspection pipeline.

Most malware runs as a regular "user mode" process to be as effective as possible on basic users' (i.e. non-administrator privileged) computers. It is not uncommon for rootkits and other malware forms to leverage the user mode to install their kernel drivers or modify operating system components.

Sandboxes that monitor user mode behavior examine the system calls of the operating system, via the Windows API; these are the functions of the operating system that provide services to applications, such as reading from files, sending network traffic and reading the Windows registry.

Monitoring these types of system calls and Windows API functions allow the sandbox to spot anything that might be out of the ordinary. For example, a PDF file trying to update the Windows registry.

However, to fully protect the operating system, the sandbox must also monitor the steps that a program or file executes between the system calls it makes i.e. the sandbox is able to determine how the malware processes the data it has just received via the system call or Windows API. These types of information sources are vital to the successful operation of a sandbox.

Mimecast's Attachment Inspection sandbox, leveraging software from a leading sandbox provider, Lastline, uses full system emulation, combining the visibility of an emulator with the resistance to detection (and evasion) of virtualization.

Full System Emulation (FSE), where the sandbox simulates the physical hardware (including the CPU and memory of a host platform), provides the deepest level of visibility into malware behavior, and it is also the hardest for advanced malware to evade.

The Mimecast URL Inspection Pipeline -Deep Dive

The third branch of the Mimecast Email Security Inspection Pipeline is focused on inspecting URLs. Like the Attachment Inspection branch, after the central pipeline has completed its spam scanning steps, all emails are sent through the URL Inspection analysis pipeline. The use of malicious URLs is a common attack technique for all types of attackers. Clicked URLs can take victims to phishing sites, where credentials or other sensitive information can be stolen, or to malware drop sites where vulnerabilities in their browser or operating system can be automatically taken advantage of by the malicious server.

The URL Inspection branch (seen in Figure 4 above) consists of a six-step process that first discovers and then analyzes URLs in an email both pre- and post-click, with an increasing depth of analysis as the email moves through the analysis pipeline. The URL Inspection pipeline is provided by a specific Mimecast service, URL Protect, which is part of the Targeted Threat Protection solution. In the rest of this section this paper will discuss the analysis that is conducted to determine if the URL is malicious or not.

Pre-Click: URL Discovery, Analysis, & Rewriting

As the name of this first step implies, before any URL analysis can occur, URLs must first be discovered in the body or the subject line of the email. Any text that contains formatting similar to a URL (e.g. http:// or *.co.uk or www.xyz.com) are discovered. Once discovered, at this step in the process, all URLs are rewritten using a Mimecast shortening service which also includes hex characters to represent unique customer and recipient codes. By using a Mimecast hosted shortening service, all clicks of these links will drive the URL resolution through the Mimecast managed Web infrastructure, thus enabling Mimecast to assess the security of the requested site or file before it is delivered to the end user.

In addition to the URL discovery and rewriting mentioned above, the URL Inspection analysis also includes a pre-click analysis (on entry of the email into the Mimecast Gateway) of the detected URLs. This step adds the URL's score to the email's overall spam score as discussed in the Spam Scanning step above. This URL analysis uses a mix of commercial, open-source, and Mimecast proprietary data sources to assess the riskiness of all detected URLs in an email (at the time of this writing these sources included: Spamhaus IP, Domain, Newly Observed Domains; as well as multiple others such as Cloudmark, URIBL, Invaluement, SURBL, and BitDefender; & Mimecast proprietary list).

Post-Click: Resolve URL to Final Web Address & Page

Once the rewritten link is clicked by the end user the next stage of analysis is initiated. The first step is to resolve to the actual destination that the clicked URL will land on. One method attackers use to confuse more simplistic email security systems is to hide the actual malicious destination site behind many redirects. At this stage in the URL Inspection pipeline the now clicked URL is resolved to the final destination.

Post-Click: Apply Customer & Mimecast Block & Permit Lists

Before any deeper URL analysis needs to occur, the URL Inspection analysis pipeline references each customer's managed URL block and permit lists as well as block and permit lists that are managed by the MSOC. Customer managed block and permit lists, are just that, lists of URLs that the Mimecast administrators of a customer wants to block or permit for reasons of their own.

For the Mimecast list, these are global lists that the MSOC team has determined are malicious, but in many cases, are not yet recognized by other sources and thus are in need of explicit blocking at this stage. Similarly, permit lists managed by Mimecast enable the MSOC team to override the blocking of non-malicious sites that are being flagged as malicious, but should be considered benign.

Post-Click: Apply 3rd Party Domain & IP Reputation Lists

Keep in mind that the time between pre-click/inbound URL analysis and post-click analysis could be hours, days, weeks, or more. A lot can change in this time. Good sites can become bad and bad sites can become good. For post-click analysis Mimecast uses a combination of open source, commercial and Mimecast generated sources and services to make the block/don't block decision (currently at the time of this writing Mimecast uses: Anti-Phishing Working Group Block List Feed, Spamhaus – IP and Domains, Mimecast Blacklist/Permit List, our own curated list of phish kit signatures created automatically and manually through Mimecast ThreatLabs and MSOC, and multiple other commercially available services such as BitDefender, CATDB, SURBL, and URIBL). Hits on any of these services would cause access to the URL to be blocked and a browser-based notice to be provided to end-users.Post-Click: Block Dangerous File Type Downloads & Sandboxing

After passing the checks of the domain and IP reputation lists discussed above, the URL analysis is by no means done. Next up is the blocking of dangerous files and file types. While most URLs point to Web pages, not all do.

Attackers will often use embedded URLs to initiate a direct download of a malicious file, instead of actually attaching the file to the email. At this step of the URL Protection pipeline, URLs which link directly to a file download are analyzed. Direct file downloads can be blocked based on the file types – Office and PDF files as well as for URLs which lead directly to the download of much more likely to be dangerous file types, such as .bin, .dll, .exe., .jar, as well as approximately 100 other file types.

In addition, in this file analysis stage of the URL Inspection pipeline there is the capability to send direct linked files to the sandbox before they are delivered to the requester. This is the same sandboxing service that was described in the Attachment Inspection section above. But instead of the file being pulled from an email attachment, the file download is detected by analyzing the rewritten link on-click. During the file analysis period, the end-user is shown an intermediate Web page informing him of the file security analysis that is taking place. Files that are determined to be good are allowed to be directly downloaded and files that are determined to be malware have their download blocked.

Post-Click: Real-Time Deep Inspection & Analysis of Page Contents

One might conclude that the first 5 steps of the URL Inspection Pipeline would be enough to protect an organization from malicious URLs. However, the fact is that as attackers have become faster and more targeted; protections which rely primarily on historical information and blacklisting can miss newly created malicious URLs and Web pages. It is very difficult, similar to file blacklisting, for a purely URL blacklisting approach to keep up with the deployment of new phishing sites.

Therefore, the final stage in the Mimecast Email Security URL Inspection Pipeline conducts real-time inspection and analysis of the page content at click-time to determine the safety of the requested URL. Within just a matter of a second or two, key attributes of the page and the URL itself are scraped, analyzed, and scored to evaluate the riskiness of the requested page.

Page and URL content which are used to score the riskiness of the URL include:

- The logos of highly phished brands, detected using computer vision AI, that appear on a page who's digital identity is not consistent with that brand (certificate mis-match)
- Page structure consistent with phishing kits or known phishing sites
- Identifying if the site has a digital certificate.
  Who is the certificate authority?
- Type of data entry on page (<input type= “Password”>)
- Existence of key words on page or in the URL itself
  (sign-in, sign-on, log-in, verify etc…)
- Use of homoglyphs/homographs (look alike) words to potentially obscure actual content
- Detection of general obfuscation techniques – Such as Webpage AES Encoding
- CSS, JS, & other external linked resources hosted on well-known good sites
- Links to external sites
- Detect suspicious type of redirection using javascript or similar functions commonly used by phishers
- Existence of an IP address in the URL

The results of the above analysis are combined using a weighted average which produces a URL risk score that varies from 0 to 1, with 0 being a score with a high certainty of safety and 1 being a score of high certainty of maliciousness (see Figure 7). The weighting factors used are continuously adjusted based on machine learning techniques as false positives and false negatives are fed back into the system by a team of specialized researchers.

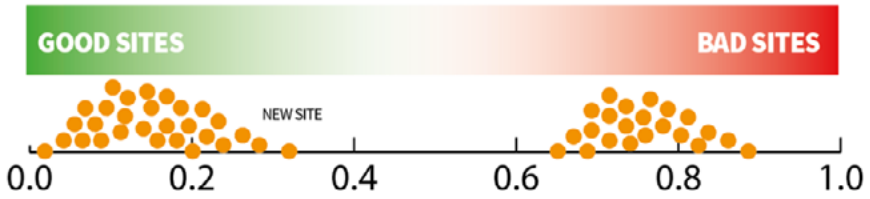

Figure 8: Weighted Average of Risk Score is Compared with Scores of Previously Analyzed URLs

Browser Isolation

An additional stage of the URL Inspection Pipeline is to invoke Mimecast Bowser Isolation, which is an optional add-on for both Mimecast Email Security and Web Security. It is only invoked if, following the steps described above, the URL has not been classified malicious AND it is a new URL that has not yet been assigned a category. In this case, the target web site is proxied and accessed by a remote browser running in an isolated container on a secure server in the Mimecast cloud.

The user browses the web as usual and web pages are safely streamed from the Mimecast cloud to the user's browser. Browser Isolation uses video streaming technology, which is considered the most secure remote browser implementation, as it maintains an "air gap" between the remote browser and the user's local browser. To ensure minimal impact on bandwidth and computer performance, video is X264 encoded, streamed at a variable bitrate using a proprietary, secure and lightweight communications protocol and rendered in the user's browser. For some websites, the bandwidth consumed by video steaming can be less than that for direct download.

Browser Isolation further maintains the "air gap" concept by preventing any direct text entry into a web page and data extraction from it. This blocks phishing attempts for credentials and other sensitive information, as well as malicious or unintentional loss of information. It prevents direct data extraction using copy/paste, which could inadvertently result in malicious code being copied back to the user's computer or device.

User interaction with the web page can be enabled by policy and applied on a user/group basis. The interactions are not direct but via a remote clipboard that only allows plain text input/extraction. Policy can be defined to allow the user to:

1. Type or paste plain text into the remote clipboard and from there send it to the remote browser.
2. Copy data from the web page to the remote clipboard where it is rendered as plain text, which can then be copied back to the local clipboard.

Files can be downloaded into the remote container and over fifty document formats are supported for rendering in the remote browser. Any malicious files downloaded, intentionally or by drive-by download, are executed remotely, in the remote container. This contains any malware infection safely away from the user's computer, device or network and prevents the patient zero problem that occurs when a user downloads zero-day malware.

Internal Email Protection

The email analysis described to date in this paper relates only to inbound emails. But what about emails which are generated internal to the organization and are directed either to another internal recipient or outbound from the organization to customers and partners of that organization? Are those certainly safe given they are internally sourced? Unfortunately, not.

For example, attackers which gain access to valid user credentials are often able to log-in to the user's organization and access applications, including email, as that user. Once logged in as that user the attacker can readily spread their attack internally or externally acting as that user. Therefore, in addition to the inspection of inbound emails, the Mimecast service can inspect internal emails for malicious attachments, URLs, and sensitive content and ultimately alert administrators and/or remove the offending emails automatically. This functionality is also very useful to address true insiders (employees) that are being careless or malicious with the emails they are sending.

In addition, Internal Email Protection services from Mimecast enables the detection and the manual or automated removal of files that were initially considered to be benign but later determined to be malicious through the continuous reanalysis of previously delivered files. It does this by reanalyzing files incorporating the latest threat analytics and intelligence available and continuously reviewing the Maverick file hash database discussed above. Organizations using Internal Email Protect can be notified of any malicious file discovery and/or have it automatically removed from their users' inboxes.

Wrapping Up Mimecast Threat Protection

As should now be clear, stopping email-borne threats from entering an organization takes a tremendous variety and depth of analysis and data sources to be effective. It is no wonder that email-borne attacks are so favored by cybercriminals! Frankly most enterprise-based or cloud-based security systems cannot provide the depth and breadth of analysis that Mimecast can and do so in a cost effective and scalable way. Mimecast can provide it because both the technical platform – see more on Mime|OS below – and the shared services model of the cloud reinforce each other to enable Mimecast to deliver a service that is both more effective and affordable than alternatives. Of course, preventing attacks is only part of what is needed. Truly cyber resilient organizations need adaptability, durability, and recoverability for their email. These "abilities" are discussed in more depth below.

Adaptability

No preventive system is 100% effective, even one that is as multi-layered and technically sophisticated as that which is provided by Mimecast (see Figure 1). With millions of threats being created and deployed by tens of thousands of cybercriminals, some threats are, at least initially, likely to land. Fortunately, most business-impacting breaches do not occur as the result of the initial incursion. Business impacting breaches generally require multiple steps as well as time to become more than just annoying infections.

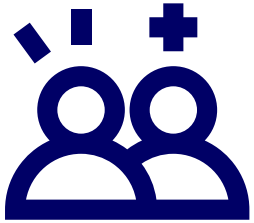

While some security solutions are effective at blocking some of the attacks as they attempt to enter an organization, they generally offer little support when attacks are orchestrated across time. To address this, organizations need to improve, and by extension Mimecast needs to deliver "adaptability" as part of a cyber resilience solution for email.

Mimecast delivers email adaptability through:

Third-party & Mimecast Generated Threat Intelligence

Our massive real-time threat intelligence network, which feeds the Mimecast email inspection system, uses dozens of threat data sources such as block lists for malicious IP addresses and domains; lists of newly observed domains; categories of sites; signatures of known bad files; and lists of known phishing and fraud sites. In addition, Mimecast customers regularly refer suspect emails to Mimecast for analysis. These emails, with their associated attached files and included URLs, are analyzed using both automated and manual techniques by Mimecast experts from the MSOC. The results of this analysis are then fed back into the Mimecast security inspection systems for application around the globe, making the Mimecast service highly adaptable, instantly.

The cybersecurity threat landscape changes daily, and staying current with the strategies and techniques attackers use to cause harm is a significant challenge for organizations of all sizes. Understanding this, Mimecast recently deployed the Mimecast Threat Dashboard, a feature included at no additional charge* for customers who have the Secure Email Gateway (SEG) through the Administration Console. The Threat Dashboard offers customers actionable, easily consumable data on threats specific to their organization's tenant, providing insight into which employees are most at-risk based on volume and nature of attacks, malware origin by geo-location, and recently observed Indicators of Compromise (IoCs). Armed with this intelligence, administrators can identify gaps in their existing security and make proactive adjustments to their security posture or remediate more effectively and rapidly post incident. In its current form, Mimecast Threat Intelligence offers insight into malicious attachments sent by email as identified by Maverick, Mimecast's anti-virus layer in our core email security program and by the Attachment Protect component of our Targeted Threat Protection service.

* Threat Dashboard and Threat Intel APIs are included as standard features only for customers with the Secure Email Gateway (SEG) They are designed to consolidate information about malicious attachments detected at both the anti-virus and the Targeted Threat Protection-Attachment Protect (TTP-AP) layer of protection. Mimecast recommends as a best practice that all customers upgrade to a service that includes SEG and TTP-AP. Please contact your Mimecast representative or partner for more details on these features or to discuss an upgrade if necessary.

Mimecast Threat Intelligence can also be fed into a customer's SIEM, TIP or SOAR using Threat Feed, Mimecast's new threat intelligence API. Mimecast supports the ingestion of third party indicators to provide proactive protection for threats which did not originate via email, and managed through your SOAR platform. In addition to threat intelligence specific to their own tenants, Threat Feed layers in threat data aggregated from customers' regional grid so that threats are presented in context and customers have the ability to benchmark against peers. Mimecast Threat Intelligence can be consumed as a threat feed into a SIEM, SOAR or TIP, includes intelligence specific to their own tenants and threat data aggregated from our regional grids, and threats are presented in context with benchmarking capabilities against peers.

Global Threat Analysis and Investigations

Under the guidance and direction of the MSOC team, Mimecast continuously performs attack/threat analysis that is generated from our global network of data centers, leveraging our visibility into billions of emails every month, and the millions of attached files that are analyzed each week, over our highly diverse set of tens-of-thousands of customer organizations. The MSOC team applies this learning and their expertise into the Mimecast email security inspection pipeline on a continuous basis. Performing this ensures that the Mimecast Email Security Service delivers a high-level of security and efficiency to our customers and contributes directly to the adaptability of the service. The combination of the MSOC team and the Mimecast Email Security Inspection Pipeline cannot be matched by an enterprise attempting to "do-it-yourself" on-premises, or even by another cloud-based email security provider.

Deployment of Best-of-Breed Technologies

The best security defenses are constantly adapting and are thus never static. Mimecast, under the guidance and analysis of the MSOC team, constantly assesses the efficacy and efficiency of both existing and new technologies, and threat data sources. Mimecast uses what we consider to be best-of-breed technologies, including the commercial AV engines and Mimecast's Maverick AV database, in addition to technologies and threat data sources from multiple commercial providers.

Mimecast also develops its own technologies to fill gaps that can't be filled by third-parties, such as heuristic filtering, fuzzy content matching, passive DNS, content decoding and multiple other techniques. This list changes regularly as new technologies are added, while others are periodically retired when they have outlived their usefulness. This constant assessment, development, addition, and removal of technologies drives the adaptability of the Mimecast service over time.

Layering of Technologies

In addition to using a mix of third-party and Mimecast-developed technologies and threat data sources, Mimecast liberally applies the best-practice of multilayered security to better adapt to the continuously changing threat environment. At different times, and with different types of threats, specific security controls are more and sometimes less effective. Given one can't predict which attack technique is going to be used (usually all are used constantly in varying degrees and amounts), the best security approach is to layer security defenses so the overall security system automatically adapts to the attacks it's addressing. Clearly layering as represented in Figure 1 serves as the basis for the Mimecast service. What to some might seem redundant forms of analysis are best thought of analysis layering to improve the overall efficacy and adaptability of the system.

Inline Education

One-time – or annual – compliance training isn't enough to build an effective and adaptable "human firewall". Organizations need to educate employees in real-time at teachable moments and learning opportunities. The Mimecast user awareness capability of Targeted Threat Protection - URL Protect takes a different approach. The Mimecast dynamic user awareness service provides inline education as users are clicking. As the user is working in their email and clicking links, they are periodically prompted to assess the riskiness of a site and given information about the sender and the site to help them make a judgement.

Conversely, if they choose incorrectly, they are informed accordingly. On the administrative side, the Mimecast administrator can see who is doing a good job and who isn't, positioning them to take appropriate actions. In addition, dynamic inline education – which prompts users that make mistakes more than those that do not - contributes directly to the adaptability of the people in the organization and thus the overall security posture of the organization.

Remediation

The Mimecast remediation capability enables organizations to detect and remove email-borne threats and malware that were initially let into the organization via the Inspection Pipeline. While it is rare, for example, that malware can get through the Mimecast malware inspection funnel, consisting of multiple AV engines, sophisticated static file analysis, and file sandboxing, it certainly can happen. The remediation capability of Mimecast provides historical monitoring of files using the Maverick database that were initially passed by Mimecast but later determined to be malware by our global threat intelligence system.

The combined power of our customer base means that we are able to constantly re-evaluate the security status of delivered files. If our threat intelligence systems reclassify a file's threat score, these malicious files will be automatically flagged to the administrator and/or removed from the organization. This feature enables the Mimecast service to literally adapt to newly discovered threats that might have been missed when they were first sent into or within an organization.

The Mimecast API Ecosystem

The Mimecast email security system provides important data and threat intelligence that can inform, strengthen, and improve the adaptability across the customer's security estate. Mimecast ecosystem offers value through an extensible API library and set of off the shelf integrations. The value obtained by the customer through the use of Mimecast's ecosystem is:

1. Bilateral threat sharing between Mimecast and the perimeter security technologies
2. Enhanced detection capabilities through the use of external threat feeds and sandboxes.
3. Consolidate log data from email into a SIEM tool of choice to surface the threats directed toward the organization
4. Faster response to incidents by leveraging automation capabilities for investigation and remediation

Durability

Email services may be forced offline by a cyberattack such as ransomware, a distributed denial of service attack, or as a response initiated by the organization's security team to contain a threat once it has landed. In addition, even cloud-based email systems such as Microsoft's Office 365 go down with some regularity! We see it happen from our perspective at Mimecast, since we see when our customers on Office 365 invoke our email Continuity service.

This downtime can directly affect business operations by preventing or limiting the ability to communicate with customers, partners and suppliers, as well as internally. Access to data held in email can be affected, too. To prevent these types of outages, organizations must provide an email communication platform that remains available while ensuring the integrity and security of the communication and data moving within.

To enhance the durability of an organization's email service, Mimecast provides:

- **Email Continuity:** Ultimately, email communication must remain available for an organization to run. With Mimecast Continuity Event Management, regardless of there being an outage within the corporate email environment, or with Office 365, a connection is established directly between each email user and the Mimecast service to ensure that emails (inbound, outbound, and internal) will continue to be sent and received. Once the primary systems are restored, all past communications will be synchronized, and the use of the primary email system will be re-established.
- **Security Policy Integrity:** When organizations are forced into using secondary email systems they often must rely on a less robust architecture - one that doesn't have the latest security protections in place. Mimecast email services eliminate the need to have a backup security system in-place as it serves as both a primary and secondary system. In addition to not experiencing email downtime with email continuity, your organization's security controls and policies are not compromised during a continuity event either. Thus, when parts of the enterprise infrastructure go down, the email systems' security protections remain available at the highest level. Providing a high level of durability.

- **DLP Policy Enforcement:** Ensuring that sensitive information and organizational IP remains protected is a critical concern when it comes to email. Mimecast DLP & Content Security protects organizations by scanning emails and their attachments to ensure that sensitive information is being protected and that the organization's content controls remain durable to both malicious and careless email activity.
- **Internal and Outbound Threat Containment:** Containing risks from inbound threats is where much of the attention is placed for protection, but once a threat is inside the environment, very little is often done to stop users (or attackers) from sending and receiving malicious content, dangerous links, or sensitive content internally. Attackers also realize the inherent trust organizations have with their business partners and customers. Mimecast enables organizations to treat all email the same. Communications within and outbound from the organization are treated like all inbound emails – they will go through much the same security protections that are shown in Figure 1.

Recoverability

Organizations need to keep data protected and accessible for their users always. However, many organizations are unaware of the risks involved when malicious attacks occur, or point-in-time recovery is required for other reasons. Leveraging an integrated archiving service can automate and simplify the process of recovering email and other important data:

- **Archive for Compliance:** The Mimecast Cloud Archive provides an immutable copy that provides digital preservation to comply with industry standards and regulations. To maintain compliance, emails must be immutable, indexed for search, and provide defensible deletion.
- **Data Protection:** When data is blocked, corrupted, deleted, stolen or encrypted by attackers (such as with ransomware), organizations need a recovery solution in place to quickly recover their email, calendars and contacts. Using a data protection and recovery solution ensures that information can be accessed in the event of an attack and access to the archive is available during an outage.
- **Point-in-Time Recovery:** Recovering from an attack can be complex. It may involve the removal of malware such as ransomware, in addition to the recovery of emails and attachments. When considering that advanced threats may sit dormant for days, being able to recover to a point-in-time (prior to the malware insertion) becomes key. Relying on overly simplistic recovery systems will allow the threats to re-emerge.
- **Impact Analysis:** After a malicious attack or data breach, organizations are faced with the immediate challenge of recovery and the return to normal operation. For complete remediation, organizations will need to further analyze the impact which may require play back as well as before-and-after comparisons. For data breaches, a review of the email and related security logs can provide insights into the potential impact to the business.

Mimecast provides the above functionality using the Mimecast Cloud Archive and its email, file, and Sync & Recover services.

MIME | OS

The Mimecast Cloud Platform - MIME|OS

Introduction - One Cloud Platform. Extensible, Scalable, Adaptable

Mimecast empowers customers through a micro-services driven architecture which provides secure multi-tenancy, high-performance and enable the rapid delivery of new products and upgrades . And a multi-product approach ensures that everything is integrated and functions together, with a unified administration experience and better visibility of risks.

Mime|OS is a unique, native cloud operating system which represents the ultimate SaaS environment, designed to deliver the true potential of cloud technology to customers today and tomorrow. Mime|OS offers secure multi-tenancy and takes advantage of the cost and performance benefits of industry-standard hardware and resource-sharing, specifically for the secure management of email and data. This allows us to provision efficiently and securely across our customer base, ensuring high-performance at the lowest possible cost for our customers.

Continuous Development Methodology

As we enhance and expand our technology, we can update services centrally, with little or no intervention required by the customer, and with no need to purchase any additional infrastructure. Improvements, upgrades, new products or patches are applied once and are available immediately across our whole service to customers. It also means we have only one, up-to-date version of our service to maintain and support, as well as a single, common data store for all customers that simplifies management, support and product development.

Rapid Product Innovation

The long-term benefits of Mimecast's platform accrue over time. Mime|OS is a repository of high-functioning, granular software components called micro-services. Our developers aren't coding from scratch – rather, they are designing services and streamlining combinations of software components from within the Mime|OS. Microservices enable Mimecast to bring new products to market faster, upgrade existing ones more rapidly, and to immediately rollout changes globally across the entire customer base.

Multiple Products with a Single Administration Console

Mime|OS underpins all Mimecast products, ensuring that everything is integrated and functions together, not as separate systems that need to be managed independently. A single console unifies the administration experience, providing better visibility of risks and enhanced management. As new products are added to the Mimecast portfolio, Mime|OS ensures we will continue to provide leading administrative experience.

What Mime|OS Enables

Faster product innovation

The Mime|OS architecture uses micro-services as core building blocks for Mimecast solutions. These software components can be extended with new functionality quickly and easily, dramatically reducing the time-to-market for new development initiatives and decreasing time to value for our customers.

Scalable

Mime|OS allows customers to quickly scale their business in periods of high growth or acquisition. Mime|OS uses shared resources and because Mimecast is 100 percent cloud, customers don't need to worry about investing in and upgrading their infrastructure. This saves our customers time and money.

Integrates Easily (Via a Robust API)

Mime|OS improves the cyber resilience strategy of an organization by providing a robust API to integrate to an existing technology set, legacy applications and new cloud applications that an organization may have in place.

Consistent with modern IT strategy

Organizations around the globe are looking to decrease their data center footprint by moving from on-premises to cloud solutions. They are also looking to decrease complexity by phasing out legacy security and archive point solutions. Mimecast's approach supports this strategy through the true cloud architecture of Mime|OS.

Enhanced visibility across communication services

Mime|OS integrates all Mimecast products and provides a single administrator interface for customers. A single view provides an in depth understanding of email risks and simplifies management tasks.

Conclusion

Our hope is that this document has given you a much deeper sense of what happens to email as it traverses the Mimecast cloud from a security perspective. What it takes to protect an organization from email-borne attacks can certainly seem daunting! There clearly is no silver-bullet technology to get the job done. Attackers are just too sophisticated and innovative for that. The concept of protection must go beyond pure threat protection into the complementary areas of adaptability, durability, and recoverability, to better manage and mitigate the risk of email-borne attacks for organizations. Beyond the specific analytics and content sources that Mimecast applies in our email inspection pipeline, which changes on a regular basis to keep up with attackers, what makes the Mimecast service truly unique and valuable, now and going forward?

It is not exactly news that most security controls, particularly those that are traditionally thought of as network-centric ones (versus endpoint centric ones), are moving to the cloud. Many think that this is happening because of the massive economies-of-scale enabled by cloud providers' multi-tenancy architectures, the shifting of hardware and software deployment, upgrading, and administration from the customer to the cloud provider, or the elimination of costly data center real estate and its associated operating costs, for customer organizations. Others think that the shift is due to the security staffing economies-of-scale that are inherent with cloud providers versus typical organizations. Of course, these are all important factors driving the move of security controls to the cloud, but these miss a key value sleeper - the network effect. A key value of moving ones' security controls to Mimecast is the massive security focused network effect that can be reaped when one is simultaneously hosting security for tens-of-thousands of global organizations and their millions of users. This security network effect shows up in two related ways.

Firstly, as attackers shift their tools and tactics, Mimecast can quickly pull early warning data from our systems and the flow of malicious email. We then combine it with a multitude of external intelligence sources, and can very quickly detect (and even predict) the attackers' new tactics and make content changes to better defend against them. Literally, the whole network of Mimecast's customers benefit from the discovery of a new attack type that hits just one member of the customer base.

There is a second security network effect that is realized by Mimecast. Detecting and blocking new attacker tools and tactics can sometimes go beyond needing just changes to security content, they require new security functionality. Because the development and deployment of cloud services differ radically from traditional software or appliance-based development models - in the form of DevOps, which leverage MIME|OS - Mimecast can build and deploy new functionality with relative ease and speed. Once new functionality is available it can be applied to every customer, globally, immediately. Another form of the security network effect of Mimecast cloud security service in action!

The takeaway is that for most organizations there are massive benefits in both efficiency and security efficacy that come from working with Mimecast. We can literally do email security better for less! Better and at a better value than organizations can do it via on-premises solutions on their own or with other cloud-based email security providers. Mimecast can provide it for more than just the prevention of email-borne threats, and can deliver true cyber resilience for email!

Visit mimecast.com/state-of-email-security to learn more

Mimecast is a cybersecurity provider that helps thousands of organizations worldwide make email safer, restore trust and bolster cyber resilience. Mimecast's expanded cloud suite enables organizations to implement a comprehensive cyber resilience strategy. From email and web security, archive and data protection, to awareness training, uptime assurance and more, Mimecast helps organizations stand strong in the face of cyberattacks, human error and technical failure.