eBOOK

McAfee
Together is power.

Definitive Guide to Office 365 Data Protection

Table of Contents

Definitive Guide to Office 365 Data Protection

Introduction

Microsoft is one of the leading cloud service providers at enterprises around the world. The company's Office 365 suite of cloud applications ensures enterprises always have the latest versions of Excel, Word, PowerPoint, and Outlook, as well as cloud-based collaboration and productivity platforms OneDrive, Exchange Online, Yammer, and SharePoint Online. However, by virtue of an ever-changing web of regulations that span internal policies, federal, state, and local jurisdictions, as well as global bylaws like the EU's General Data Protection Regulation (GDPR), organizations are required to safeguard sensitive or protected information in cloud services like Office 365.

In an earlier era, most corporate data lived in Windows file servers and on-premises applications, such as SharePoint. Sharing and collaboration was limited to email and file servers. Enterprises invested in a generation of security technology built for this environment to enforce corporate policies and protect against data loss.

Today however, organizations use a wide variety of cloud-based applications that facilitate storage, access, and sharing of data. At the same time, like most cloud providers, Microsoft operates under a shared responsibility model. The company is responsible for Office 365 platform security, including preventing intrusions and DDoS, malware, and other advanced threats. Customers are responsible for ensuring their employees do not misuse corporate data stored in Office 365 and that their login credentials do not fall into the wrong hands.

And while Office 365 makes it very easy to upload and share vast amounts of corporate data with others, both inside and outside the company, with a few clicks, this opens organizations to several vectors of data loss, including highly sensitive data being uploaded to Office 365 against company policy, sensitive data being shared with unauthorized parties, and data being lost after being downloaded from Office 365 to an unprotected or unmanaged device. For these reasons data loss prevention (DLP) efforts have been expanded from traditional on-premises email servers to include anywhere corporate data lives in the cloud.

In this eBook, we will:

- Summarize the state of Office 365 adoption
- Discuss some of the most common data loss scenarios in Office 365
- Identify threats to data and users in Office 365, with real world examples
- Review Office 365's native security capabilities available out-of-the-box
- Provide guidance on how to implement data loss prevention and threat protection for Office 365 in a multi-cloud IT environment
- Office 365 DLP best practices

Connect With Us

How Enterprises Use Office 365

As one of the most popular cloud services, Office 365 holds the top spot in cloud usage rankings by user count. Broken down by individual application, OneDrive for Business has the highest penetration rate with 79.1% of organizations possessing at least 100 users. It makes sense that OneDrive is deployed at so many organizations, because it is included in every Office 365 plan, even the entry-level ProPlus plan that primarily gives access to Office applications on the desktop (Word, Excel, PowerPoint, and so forth).

OneDrive for Business also has the highest usage rate, with 18.6% of all enterprise employees actively using it. Exchange Online has the second-highest penetration rate: 66.9% of enterprises have at least 100 users. However, while Skype for Business is used by fewer enterprises, more users are using cloud-based Skype for Business than Exchange. One way to interpret this data is that enterprises are beginning to migrate to Exchange Online from on-premises versions of Exchange but that—owing to the scale of these migration projects—they are migrating in phases.

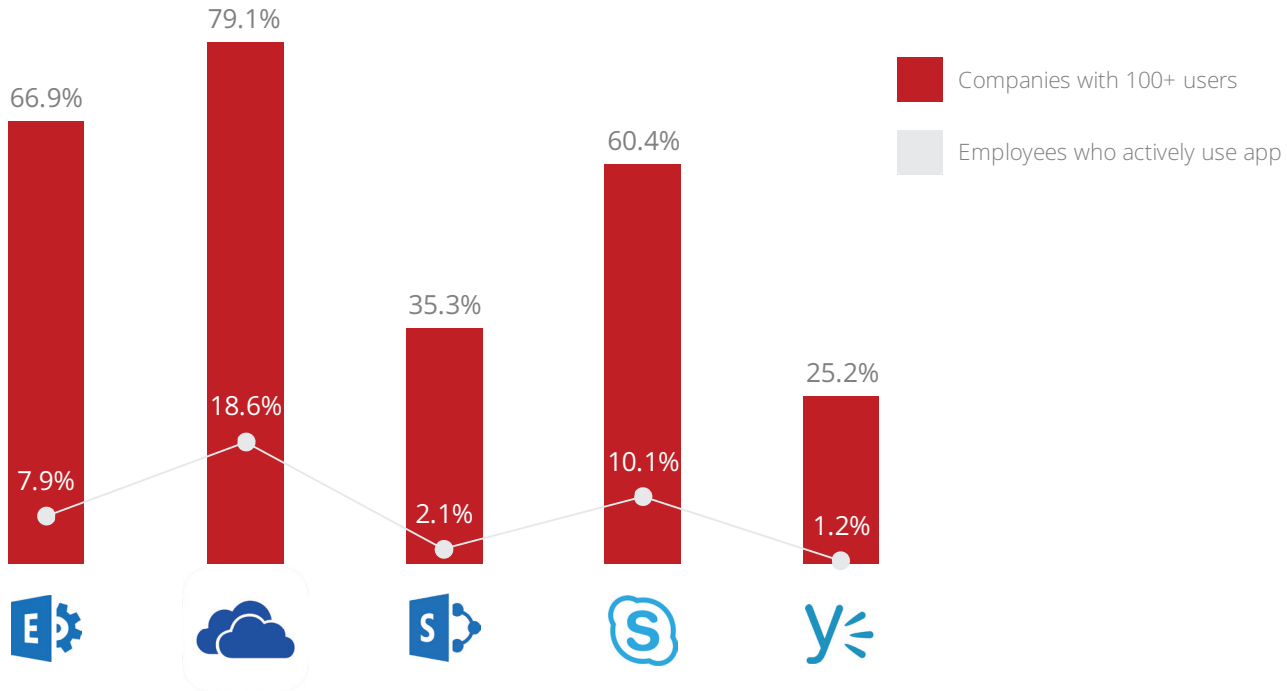

Similarly, the complexity of many sprawling on-premises SharePoint deployments may be slowing down migration to SharePoint Online. Only 35.3% of enterprises and 2.1% of users have moved to SharePoint Online. By both metrics, Yammer is the least used Office 365 application. Yammer faces stiff competition from Slack, which is rapidly expanding in the enterprise. This may partially explain why it has been slower to expand its footprint after being acquired by Microsoft in 2012 for $1.2 billion.

Office 365 adoption is not uniform across industries. Financial services firms have the highest rate of Office 365 usage. Perhaps this is not surprising, because financial services firms are simultaneously heavy users of Microsoft Office, particularly Excel, and also seek to have the latest technology tools to maintain a competitive advantage. Within financial services, 39.3% of users actively use OneDrive for Business, and 17.3% actively use Skype for Business. By far, the most popular Office 365 application in healthcare is Skype for Business, and 14.2% of users rely on it for online meetings, messaging, and audio and video calls. Manufacturing leads adoption of Exchange Online, with 12.9% of users actively using Microsoft's cloud-based email platform. That's followed by media and entertainment with 12.5% of users on Exchange Online.

	Exchange Online	OneDrive for Business	SharePoint Online	Skype for Business Online	Yammer
Energy	7.3%	13.6%	1.5%	5.8%	0.3%
Financial Services	2.5%	39.3%	1.2%	17.3%	0.6%
Food Products	2.1%	2.5%	5.7%	25.8%	0.3%
Healthcare	2.0%	4.2%	0.1%	14.2%	0.1%
Manufacturing	12.9%	26.6%	0.6%	3.7%	1.4%
Media and Entertainment	12.5%	4.6%	1.0%	1.4%	1.3%
Real Estate	6.6%	10.6%	0.2%	4.7%	0.2%
Technology	10.3%	18.2%	1.8%	15.5%	0.4%

Most Common Data Loss Scenarios in Office 365

Because enterprises store a significant volume of business-critical data in Office 365, the stakes for keeping data safe are high. Some of this data may not belong in the cloud at all. For example, our research has shown that the average enterprise has 204 files that contain "password" in the file name stored in OneDrive. Generally, security experts don't recommend storing all of your passwords in an unencrypted Word or Excel document, whether you store it in the cloud or on your computer. Some of this data is sensitive but can be safely stored in the cloud with appropriate controls in place. When reviewing all types of data in OneDrive and SharePoint Online, we found that 17.1% of that data is sensitive. Broken down by type of data:

- 9.4% of data is confidential (such as financial records, business plans, source code, trading algorithms, and so forth).
- 4.1% of data contains personally identifiable information (such as Social Security numbers, tax ID numbers, phone numbers, date of birth, and others).
- 1.9% of data contains protected health information (such as patient diagnoses, medical treatments, medical record IDs, and so forth).
- 1.7% of data contains payment information (such as credit card numbers, debit card numbers, bank account numbers, and others).

Office 365 delivers a powerful set of collaboration tools and frees employees to access data from anywhere, using any device. However, these capabilities also introduce potential issues that enterprise security, risk, compliance, and audit teams have not faced before. After working with hundreds of enterprises to help address security and compliance requirements as corporate data migrates to Office 365, McAfee has surfaced some of the most common data loss scenarios that enterprises must actively prevent:

- Sensitive or regulated data is shared with an unauthorized party.
- Sensitive or regulated data is uploaded to Office 365 against company policy.
- Sensitive or regulated data is downloaded from Office 365 to an unmanaged and unprotected endpoint device.
- Sensitive or regulated data is downloaded from Office 365 and uploaded to a risky shadow cloud service.

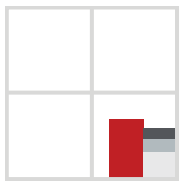

17.1%

of files in OneDrive and SharePoint Online contain sensitive data.

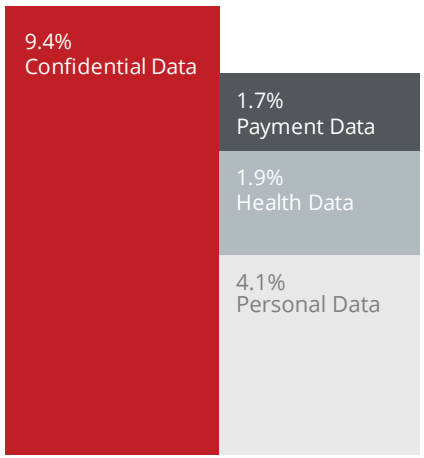

Threats to Data and Users in Office 365

Microsoft takes the security of the Office 365 platform very seriously and has made significant investments in service-level security. These investments protect Microsoft's cloud-based applications from intrusions. Office 365 is one of the few cloud services to receive the highest rating of "Enterprise-Ready" from McAfee based on an objective assessment of its security controls.

However, users can still perform high-risk actions within these applications that can put an enterprise's sensitive data stored in Office 365 at security risk. Moreover, account credentials can be acquired via phishing scams and used by third parties to gain access to corporate data.

Taken together, the average organization faces the following Office 365 threats each month:

- 1.3 compromised accounts each month (such as an unauthorized third-party logging into a corporate Office 365 account using stolen credentials)
- 0.8 insider threats each month (such as a user downloading sensitive data from SharePoint Online and taking it when they join a competitor)
- 0.6 privileged user threats each month (such as an administrator provisioning excessive permissions to a user relative to their role)

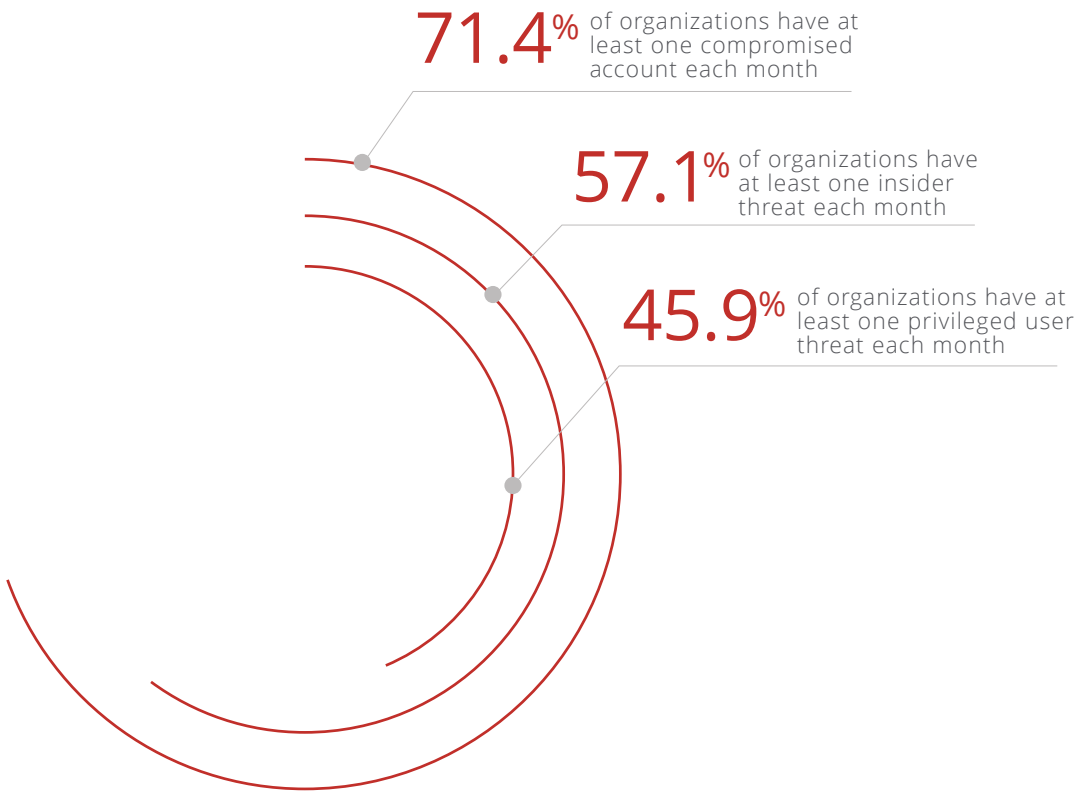

The average organization generates 5.4 million user events each month within Office 365 (such as user login, upload file, edit document, and so forth). Microsoft provides a raw event feed that can be consumed via an API, which leaves enterprises searching for a needle in a very large haystack. Increasingly, enterprises are leveraging tools relying on user and entity behavior analytics (UEBA), which use machine learning to analyze user activity and automatically detect unusual behavior. For example, this technology may surface an alert when a user logs in after 15 failed login attempts as a potentially compromised account. One of the challenges facing IT security teams today is the sheer volume of alerts they receive.

In the infamous Target data breach, cyber attackers stole data for over 40 million customer payment cards in the days after gaining access to the retailer's payment systems. Target's IT security team ignored an alert correctly identifying the breach before any card data was stolen. Had they acted immediately, it's likely the scope of the breach would have been much smaller. In Office 365, the average enterprise experiences 256 anomalous user activities within each month for every 5.4 million events (roughly a 20,000:1 ratio). However, of these anomalous events, an average of just 2.7 turn out to be actual threats to the organization. The challenge for enterprises today is how to develop the people, processes, and technology to identify these threats against the background noise of everyday Office 365 usage.

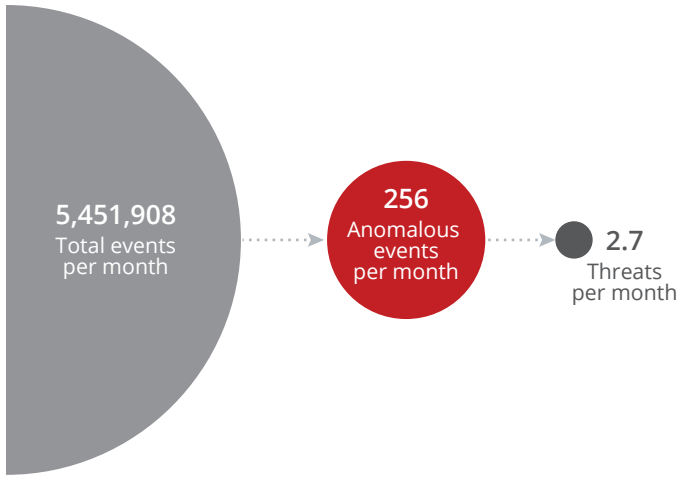

Anomalous events that do not indicate a true threat often occur in isolation. Following the above example, the user may simply have forgotten that CapsLock was on when entering her password multiple times incorrectly. So, how does an IT security professional tell the difference between a clumsy user and a cybercriminal? One thing that cannot be stolen by a third party is the user's pattern of behavior. A login from a new, untrusted location, or after several failed login attempts correlated with patterns of behavior that are atypical for that user, more strongly indicates a compromised account than simply looking at failed login attempts. By narrowing down anomalous events to a fewer number of likely threats, IT security teams are better equipped to respond when an actual threat does occur.

Examples of real-life threats to Office 365 customers

Example 1: Super Sneaky Brute Force Attack on High-Value Office 365 Accounts

In early 2017, McAfee (then Skyhigh Networks) discovered and began to track a brute force login attack on multiple Office 365 enterprise customers. Using a set of corporate user names and passwords, as well as compromised hosted tenants, the attackers launched brute force attacks on high level employees' Office 365 accounts to gain access to potentially sensitive corporate data.

In its analysis, McAfee was able to detect over 100,000 attempts (failed logins) from 67 IPs and 12 networks, targeting 48 customers' Office 365 accounts. What stood out with this attack was the sophisticated and sneaky approach of the attackers, who did not cast a wide net in trying to rope in as many corporate users as possible, which is typical in brute force attacks.

In this case, the attackers targeted a set of companies and high-level employees and launched a "slow and low" attack to avoid getting flagged by the cloud service provider. Another aspect of this attack was that it was cloud-to-cloud in that attackers leveraged infrastructures of public hosting services, to launch an attack on a SaaS service.

Execution and detection of the Attack

To execute this attack, the perpetrators acquired a set of corporate user names and passwords, which may be tied to multiple cloud services (not necessarily Office 365). Then, using public cloud tenants, they were able to launch brute force attacks on corporate users' Office 365 accounts.

To accomplish this, they tried different permutations of employee names. For example, Steven Smith (name changed), Chief of Staff, Company A saw login attempts on his account which involved different combinations of his user name such as steven.smith@companyA.com, steve.smith@companyA.com, s.smith@companyA.com. His account saw 17 attempts with 17 username permutations in 4 seconds from 14 IPs.

While McAfee does not have access to the passwords in clear text, it is speculated the same password was used because the login attempts targeted each permutation of the user name exactly once.

The attackers counted on two points for their attempts to succeed. They guessed that users were reusing the same password across different applications because they attempted to use an arbitrarily acquired password to login to Office 365. Next, they hoped for companies not to have multi-factor authentication (MFA) and Single Sign-On (SSO) activated for apps that stored sensitive data.

The first hint of the attack came when McAfee registered an anomaly associated with 'Compromised Accounts'. Skyhigh's threat protection engine is programmed to look for multiple variants of brute force attacks, so when the solution detected an abnormal pattern and high correlation between different factors associated with login attempts on user accounts, it registered an anomaly.

As more anomalies showed up across users, the anomaly elevated to a threat. The McAfee team performed a cross tenant analysis and detected over 100,000 failed logins, which confirmed a widespread brute force login attack.

While the researchers suspected an attack early on, in order to get confirmation, they ran several analyses over time, given the "slow and low" nature of the attack.

The sophistication of the attackers can be observed in the precautions they took to avoid detection. First, they launched a slow and low brute force attack. While the attack itself staggered over months to bypass any lockout checks implemented by CSPs, there were short spurts of activity where different variations of a single user name were hit by multiple IPs.

One employee, Sherry Wheeler (name changed), who is an Executive Advisor, saw 95 attempts on her Office 365 account in 5 seconds from 13 IPs, where each IP tried different username variations. For instance, IP1 attempted sherry_wheeler and sherry.wheeler, while IP2 tried s_wheeler and s.wheeler, implying it was a distributed and coordinated attack coming from multiple platforms.

Had the 95 attempts on Sherry Wheeler's account been made from the same IP, either the CSP or the hosting provider would have blacklisted the address. And finally, the attackers did not pick too many employees within a company or even a department. They appear to have selected senior or long-term employees, possibly because they are more likely to have access to sensitive data.

Example 2: Ingenious Attack Scheme on Office 365 System Accounts

In late 2017, McAfee discovered an ingenious new botnet attack against Office 365 accounts, dubbed 'KnockKnock' because attackers were attempting to knock on backdoor system accounts to infiltrate entire Office 365 environments. One of the key distinctions of this new attack was the nature of the accounts that were being targeted. KnockKnock was designed to primarily attack system accounts that are not assigned to any one individual user, making them particularly vulnerable, as we'll describe later.

Anatomy of the Attack

First, it should be noted that KnockKnock was not a brute force attack for two reasons. First, it targeted a very small proportion (typically <2%) of the Office 365 account base. Second, it was devoid of any bursts in hacking activity, and averaged only 3-5 attempts per account in order to fly under the radar of traditional defenses.

KnockKnock had been operational since May 2017 and was still active as late as October 2017. The attack was launched using a relatively small network of 83 confirmed IPs distributed across 63 networks. The smaller size of the botnet was likely designed to keep the attacker low key (i.e. the attack focuses on a handful of users at a time, before moving on to the next set).

In an attempt to further obfuscate the attack, enterprises were targeted in a staggered manner. When the attacks against one enterprise was ramping up, they were slowing down for a different enterprise. While a majority of the activity stemmed from IPs registered to service providers in China, there were activity originating out of 15 other countries including Russia, Brazil, US, Argentina, Gabon, Azerbaijan, Malaysia.

The attack was particularly clever in that it distinctively and slowly targeted system accounts. The system accounts that McAfee identified as targets included service accounts (like the ones used for user provisioning in larger enterprises), automation accounts (like the ones used to automate data and system backups), machine accounts (like the ones used for applications within data centers), marketing automation accounts (like the ones used for marketing and customer communication), internal tools accounts (like the ones used with JIRA, Jenkins, GitHub etc.), in addition to accounts set up for distribution lists and shared and delegated mailboxes.

The reason this was so clever was that system accounts, given their purpose, tend to have higher access and privileges than an average account. And, most importantly, such accounts do not yield well to authentication frameworks like Single-Sign-On (SSO) or Multi-Factor Authentication (MFA) and are also subject to lax password policies. These two aspects help reveal the motivation behind KnockKnock, (i.e. attack a weak-link with the potential for elevated exploits).

Once KnockKnock gained access to an enterprise system account, the attack was designed to exfiltrate any data in the inbox and then create a new inbox rule intended to hide and divert incoming messages. The attack would then attempt to initiate a phishing attack and propagate infection across the enterprise using this controlled inbox.

Making the Most of Office 365 Built-In Security

Office 365's robust service-level security makes it one of the 8.1% of cloud services that earn the highest McAfee® Skyhigh Cloud Trust Rating of "Enterprise-Ready." For every customer, Microsoft's platform encrypts data in transit and at rest in the cloud, and offers device pinning. Microsoft's cloud-based productivity suite also boasts numerous certifications including ISO 27001, ISO 27018, SAS 70, SSAE16, and ISAE 3401.

Additional Office 365 security capabilities can be enabled by the organization. In order to ensure the highest level of protection for data stored in the cloud, security experts recommend that companies utilize multi-factor authentication, IP filtering, single sign-on, rights management, S/MIME, and message encryption. Each of these capabilities strengthens the protection of corporate data.

Single sign-on

Giving users one password to use across their applications is not just more convenient; it also allows password policies to be managed in a centralized place. Office 365 supports popular third-party identity providers including Okta, One Login, Ping Identity, and Centrify. Microsoft also offers its own single sign-on solution, Azure Active Directory, which allows users to log in using the same password as they do for on-premises Microsoft products, as well as cloud products from other providers.

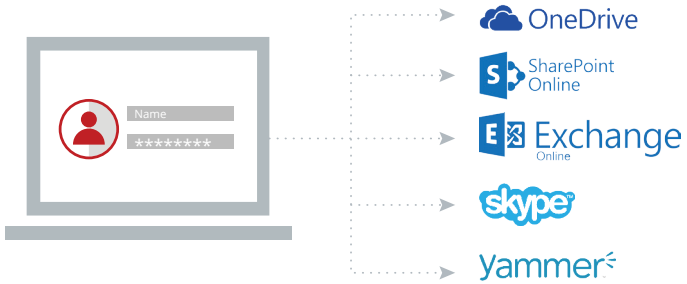

Multi-factor authentication

Multi-factor authentication makes it more difficult for a third-party to gain access to an account by requiring an additional authentication measure after submitting the username and password. Organizations that use single sign-on solutions such as Azure AD, Okta, One Login, Ping Identity, and Centrify for identity and authentication across all cloud services, including Office 365, can immediately roll out multi-factor authentication to Office 365.

For organizations without one of these solutions, Office 365 includes multi-factor authentication options built into the platform. The secondary authentication methods supported by Office 365 include the use of mobile app notification, a one-time password generated by a mobile app or sent to the user via a phone call or SMS text message, and per-app passwords used with clients such as Outlook.

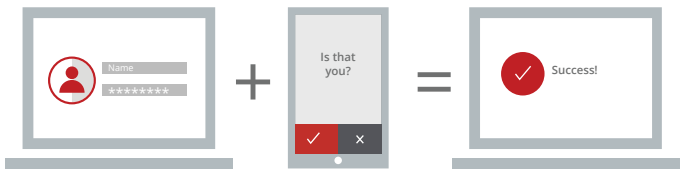

IP filtering

Another way to reduce the risk of data loss via account compromise is to disallow extranet access to corporate cloud services such as Office 365. If an attacker were to obtain an account credential, they would be unable to successfully log into the account, unless he or she is on the corporate network or accessing via virtual private network (VPN). Microsoft supports IP filtering, referred to variously as "IP Whitelist" and "Trusted IPs," for customers using either Azure Active Directory or federating user identity with their on-premises Active Directory. Some third-party single sign-on solutions also offer this capability.

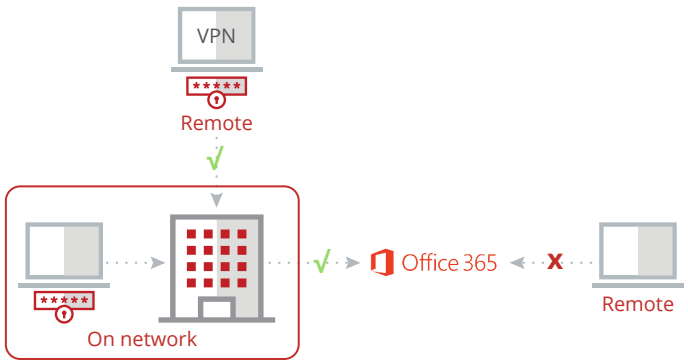

Rights management service

For companies using Microsoft RMS to protect documents stored on-premises, Azure RMS is an attractive option for extending information rights management policies to OneDrive, Exchange Online, and SharePoint Online. For companies running the on-premises version of Active Directory, you don't need to migrate everything to the cloud right away. You can federate Windows Active Directory Server with Azure Active Directory and run a hybrid environment.

However, one potential issue in using RMS with Office 365 is that your encryption keys used to enforce RMS policies will be stored in the cloud. For some organizations, this can create concerns. In SharePoint Online, RMS applies rules across an entire site collection. Since RMS requires the user to be running client software to access the document (or to print, edit, or save new versions of the document), it can be inconvenient to apply blanket RMS protections to documents that do not need these policies enforced because they do not contain sensitive data.

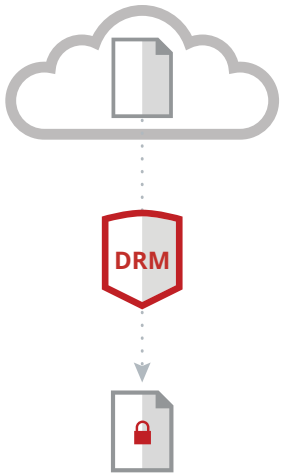

Office 365 message encryption

It's well known that emails are sent across the Internet with about as much privacy as a postcard. This can be problematic in many instances where you need to send sensitive content, such as a bank employee sending a credit card statement to a customer or a healthcare provider sending health-related information to a patient. Message encryption allows you to send a message to a recipient encrypted. The recipient receives an email with a link to a page on a download portal, where they authenticate using their Office login or a one-time passcode to view the message.

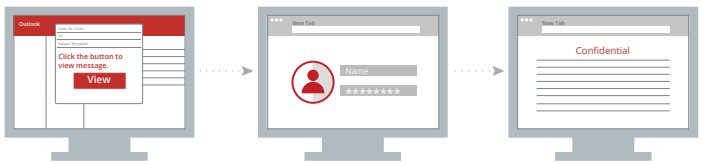

Step 1: Open email and click link

Step 2: Sign into web portal

Step 3: View encrypted message

Secure multipurpose internet mail extension (S/MIME)

Unlike message encryption, which is based on policies defined by an administrator, S/MIME is controlled by the end user, who decides whether to use it. While message encryption is browser-based, and requires no client software or certificates, S/MIME uses certificates to digitally sign and optionally encrypt the email content itself. Digitally signing the email ensures that the message content is what the sender originally wrote, and that the message hasn't been altered or tampered with.

S/MIME requires users to access their email through a client like Outlook, not a web browser. And, since you need to set up user certificates, it takes more effort to get up and running than message encryption. Certain government use cases mandate the use of S/MIME, and Office 365 supports S/MIME for customers who use Azure Active Directory, or who federate their on-premises Active Directory with Azure AD.

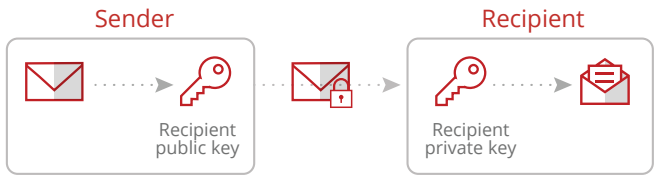

Sender's email client encrypts message with recipient's public key before sending message.

Recipient's email client decrypts message with recipient's private key and validates sender's certificate.

How to Implement Comprehensive Data Protection in Office 365 and Across Your Cloud Environments

While Office 365's built-in security capabilities are robust, most organizations—especially ones that are using multiple cloud services—quickly find that they need an additional layer of security to gain visibility and control over Office 365 and all other cloud services being used. However, relying on individual cloud service provider's security capabilities lead to two critical security problems:

1. Different cloud service providers offer varying degrees of security capabilities. This creates a situation where an organization will enforce different levels of security per application, exposing the organization to risk from cloud security gaps.
2. Creating and managing security policies and procedures within individual cloud applications can be burdensome and costly from a resource standpoint. The average organization uses multiple security tools to protect data and users, including SIEM, Endpoint protection, Firewalls/Proxies, etc. Managing cloud security in individual cloud service consoles will further tax the already stretched resources available to IT security departments.

To prevent silos of security, inconsistent policy enforcement, and costly administrative overhead, organizations need a central point of control to enforce their security policies for Office 365 as well as all the other cloud services, thereby consolidating many different security and compliance capabilities into a single platform. Specifically, we have identified 7 of the most common Office 365 security use cases that customers are most concerned about.

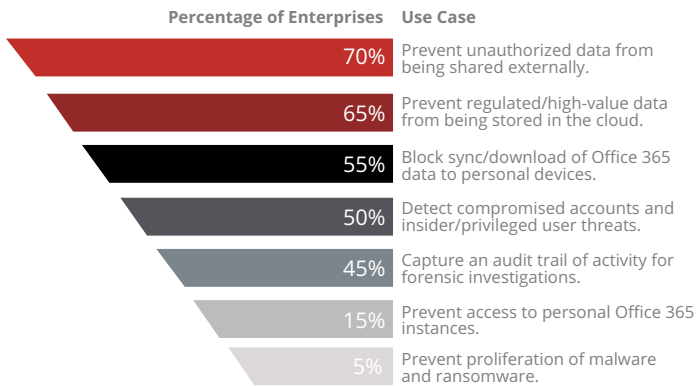

1. Prevent unauthorized data from being shared externally

There are several ways that sensitive data stored in Office 365 can be shared with unauthorized third parties against corporate security policy. Using cloud-native tools such as OneDrive, SharePoint Online, or Exchange Online, employees share a significant amount of data with collaborators internally and with external suppliers, distributors, vendors, and customers. While a small percentage of oversharing incidents are due to malicious users, most incidents are due to well-intentioned employees who inadvertently expose corporate data.

Users can share sensitive information in several ways: 1) by sending an email in which the sensitive information is either part of the body of the email or stored as an attachment 2) by inviting a file or folder collaborator by the recipient's email, 3) by sending a shared link to a file stored in OneDrive or SharePoint Online, or 4) by configuring the sharing policy to make a document publicly available and searchable.

Analyzing the sharing permissions of files in the cloud, McAfee has found 28.3% of files are shared with email domains associated with business partners. However, another 6.2% are shared with personal email domains (such as gmail.com or yahoo.com), introducing questions about who has access to corporate data. Troublingly, 5.5% of files are shared using links that can be forwarded to anyone, and the recipients accessing the files cannot be traced. And 2.7% of documents are publicly accessible to anyone on the Internet.

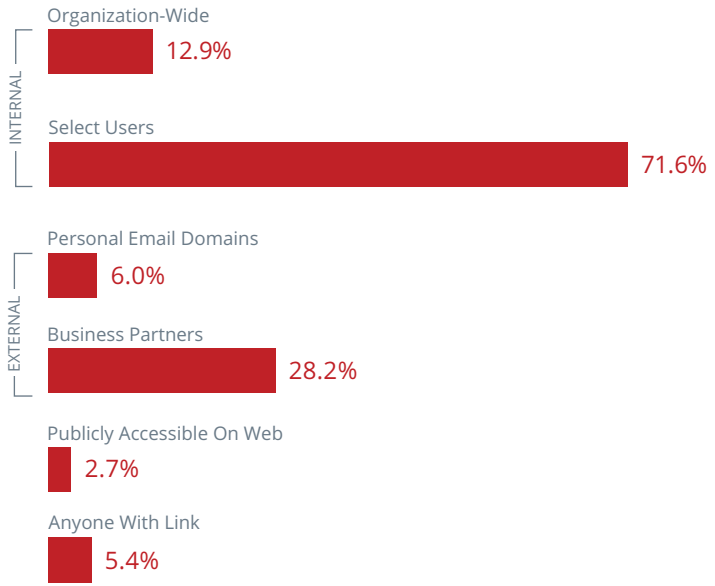

When sharing large volumes of sensitive data is just a few clicks away, it's easy for employees to mistakenly share files or folders too broadly with external users. It's also commonplace to type in a recipient's name and mistakenly select the incorrect individual or a personal email address from the autocomplete suggestions. Employees may also be sharing sensitive data externally, unknowingly violating policies. Depending on your corporate policies, you may have blanket rules about which business partners your organization's employees can share data with via Office 365. You may also have detailed policies on the type of content that can be shared with business partners.

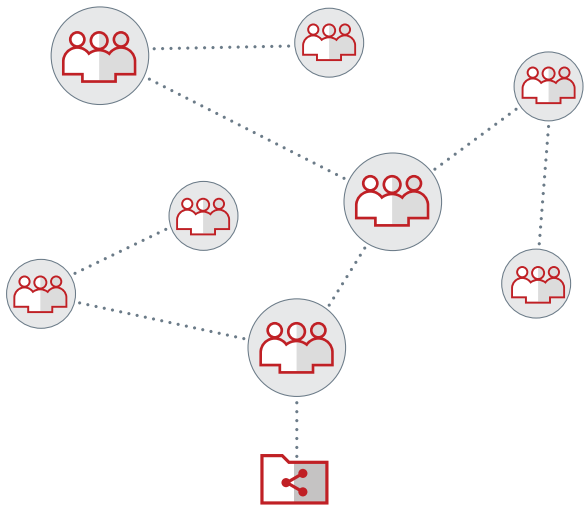

McAfee's recommendation

Organizations need guardrails to ensure appropriate sharing via content-aware data-sharing policies in OneDrive, SharePoint Online, and Exchange Online using a policy-based framework. Policies should include multiple rules, including whether a file is shared via a link or has external collaborators. Collaboration rules should trigger off the specific permissions assigned for the file or folder including viewer, editor, or owner. Some enterprises have a whitelist of acceptable external collaborators. For example, you may prohibit external sharing by default except with pre-approved business partners known to the organization or prevent sharing with personal email domains such as those from Gmail or Yahoo! Mail.

Since 95% of shared links are accessed by end users within minutes of being shared, in response to a policy violation, organizations should take remedial action to correct the violation in real-time, preventing the unintended disclosure of corporate data outside of policy. When content is shared externally, enforcement timing matters. An external user can download a file within seconds of receiving an invite to collaborate. Remediation actions may include revoking a shared link and limiting the scope of sharing permissions (for example, changing editors to viewers) or removing sharing permissions entirely.

Another recommendation is to design policies that can be targeted to specific user groups within the organization based on Active Directory (AD) attributes. Lastly, combine content-aware policies with collaboration policies such that you allow collaboration with business partners but prevent sensitive intellectual property from being shared.

"95% of shared links are accessed by end users within minutes of being shared."

2. Prevent regulated, high-value data from being stored in the cloud

As a clear sign that enterprises trust Microsoft to protect their sensitive data (or perhaps that users are operating unaware of their organizations' cloud policies) McAfee has found that employees upload a significant amount of sensitive data to Office 365. As mentioned previously, on average, 17.1% of files an enterprise stores in OneDrive and SharePoint Online are sensitive. Depending on your organization's compliance and security posture, your policies may dictate that this information can be stored in Office 365, provided it is not shared inappropriately.

But, many companies have high-value or regulated data they wish to prevent from being stored in the cloud. And, regardless of compliance requirements, some types of data are simply unfit to be stored in the cloud. For example, McAfee has found the average enterprise stores 204 files containing user passwords in OneDrive. These files often take the form of a Word or Excel document with usernames and passwords for all the applications and devices an employee uses.

Preventing regulated or high-value data from being stored in the cloud is a two-part problem: 1) detecting sensitive data and 2) enforcing controls to prevent this data from living within Office 365.

Identifying sensitive data is not a trivial undertaking, because it often requires going beyond simple keyword matching. Depending on their individual needs, organizations may need to use several methods to identify and prevent sensitive data from being stored in Office 365 or shared inappropriately, including:

- A lexicon containing hundreds or thousands of keywords that are common across several different corporate policies (such as prescription drug names or stock symbols)
- Data classification tags applied by classification technologies that appear in the metadata of files (such as "confidential," "internal only")
- Standard alphanumeric patterns that follow a set of defined rules such as length, prefix or suffix, or checksum (for example, Social Security numbers or credit card numbers)
- Custom alphanumeric patterns that are unique to the organization and follow a set of defined rules (such as part numbers or product SKUs)
- All versions of a specific, sensitive document including the exact file or any derivative of the file (for example, a design document for a production processs or a legal contract)
- Any piece of content that refers to current or former customers (such as any field from a structured database with personal data on 300 million customers)

When deploying data loss prevention technology, enterprises should aim to simultaneously minimize the number of sensitive files missed by the system (false negatives) and minimize the number of non-sensitive files flagged by the system (false positives).

McAfee's recommendation

Organizations require a standardized data loss and compliance policy and enforcement mechanism, which often starts with utilizing off-the-shelf DLP templates for common use cases such as Health Insurance Portability and Accountability Act (HIPAA) compliance or merger and acquisition (M&A) documents. Depending on the organization's requirements, DLP policy templates should be customized using a flexible policy framework that leverages Boolean logic to combine two or more rules and associated remediation actions. DLP policies may contain rules leveraging document metadata and content including file attributes, keywords, keyword dictionaries, document classification tags, data identifiers, regular expressions, and fingerprinting of structured databases and unstructured files.

As a best practice, DLP violations should be tiered by severity for prioritization. For example, if a document contains one credit card number, the violation severity should be set to "medium" and if it contains hundreds of thousands of violations, the severity should be set to "high." Taking this a step further, if remediation actions within a policy is tiered based on severity, you should consider defining policies such as quarantine files with high-severity violations, but only alert users for files with low-severity violations.

Depending on the maturity of your DLP framework, you may choose to run DLP policies in a monitor-only mode. This is recommended when starting with Office 365 DLP from scratch, especially when applying to outbound emails in Exchange Online since automatically blocking emails due to false-positive DLP alerts could have detrimental impact on business operations.

3. Block download of Office 365 data to personal devices

One of the key advantages of Office 365 is the ability to extend productivity tools to employees no matter where they are or what device they use. In a previous era, a virtual private network (VPN) was required to access enterprise applications running in the corporate data center. This requirement necessitated that users log in from managed devices that had the corporate VPN installed. Now that employees can access corporate data in Office 365 from personal devices, new risks are emerging to corporate data. One issue is that when data is downloaded or synced to a personal device, information leaves the company when the employee leaves.

An even greater concern is information falling into the wrong hands due to a lack of endpoint security controls. Personal devices that are unmanaged lack enterprise endpoint security that enforces device policies such as drive encryption and device PIN. If that device is stolen—for instance, when an employee is working from a coffee shop or if a laptop is left in the backseat of a car—corporate data is also stolen. Without endpoint security, the enterprise is unable to remotely wipe the data, which may not be protected at all on the endpoint. For these reasons, many enterprises want to allow employee access to the collaboration tools in Office 365 from any device, but limit the ability to download corporate data to only managed devices.

McAfee's recommendation

When users access Office 365, organizations should do a certificate check to validate that the device has appropriate endpoint security in the form of an enterprise mobility management (EMM) or mobile device management (MDM) solution.

In order to ensure that a malicious user or third party has not spoofed a certification on an untrusted endpoint in order to circumvent device policies, organizations should utilize a mapping of all users and their trusted devices (usually done via integration with EMM/MDM providers) to validate not only that the endpoint has a certificate, but that the user is accessing from a known device and not another device.

Detect compromised accounts and insider/ privileged user threats

Office 365 customers are responsible for actions users take within the platform that compromise data, and McAfee has found that the average enterprise experiences 2.7 such threats in the platform each month. This number includes compromised accounts, insider threats, and privileged user threats. Insider threats can generate headlines.

In a lawsuit that Google subsidiary Waymo filed against Uber in February 2017, the company alleged that a former Google employee downloaded 14,000 sensitive documents related to self-driving car technology before leaving the company. The former employee subsequently led the self-driving car project for Uber, and Uber's technology bears a striking resemblance to components developed by Waymo.

Compromised accounts are also a significant threat. Cybercriminals gain access to corporate Office 365 accounts by exploiting stolen user credentials gathered via phishing attacks, passwords leaked from other cloud services that an employee reuses for Office 365, and by guessing common passwords. Analyzing stolen passwords for sale on the Darknet, McAfee found that the top-20 most common passwords, which include "123456" and "password," account for 10.3% of all passwords. Furthermore, research by Joseph Bonneau at the University of Cambridge has found that 31% of people reuse passwords across multiple applications. Both of these trends make it easier for third parties to compromise Office 365 accounts.

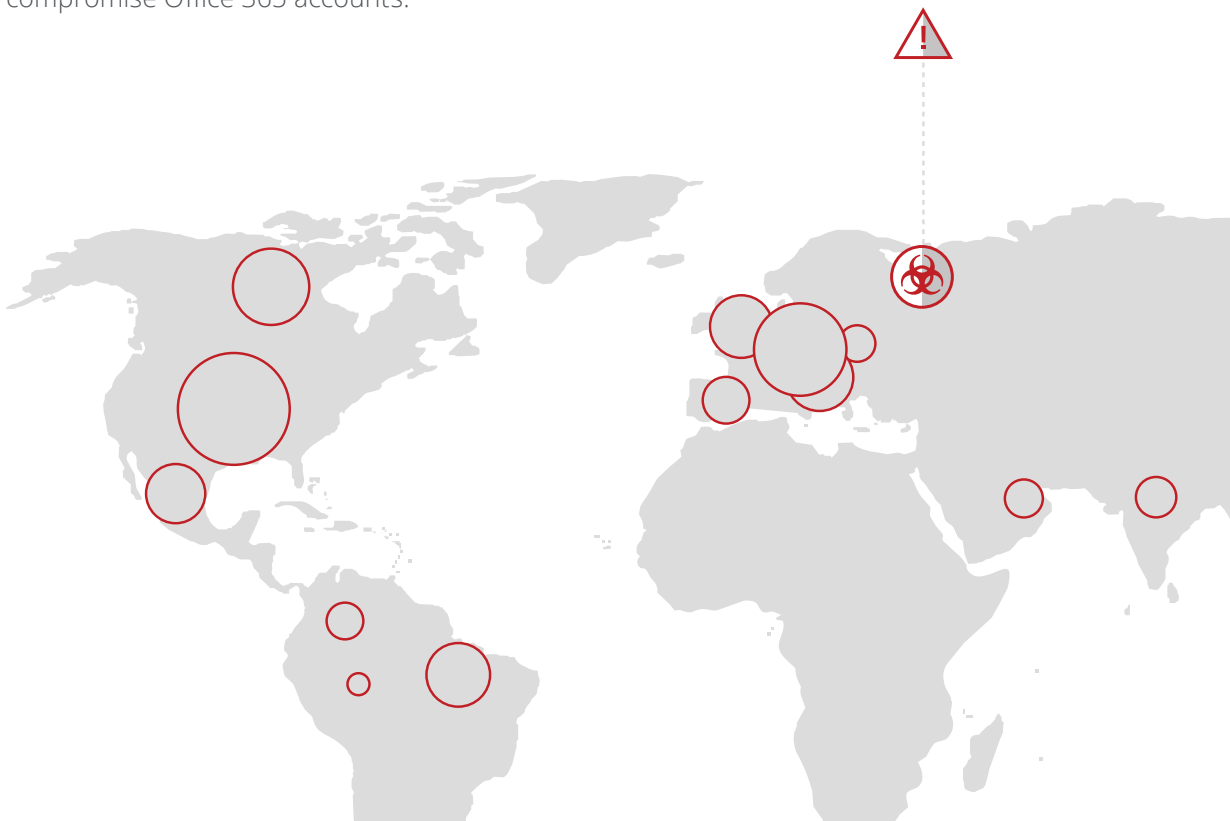

Detecting threats is challenging because, while they are often signaled by behavior patterns that are anomalous, there is no single threshold that can be applied to all users, for all time frames, that will accurately detect these threats while not also generating many false positives. For instance, it may be unusual for one user to download a series of documents with company financial performance at home on the weekend, while it may be normal for another user to periodically download the most recent of these documents on the last Friday of each month.

McAfee's recommendation

Unlike threshold-based solutions that require enterprises to define policies that detect activity outside an arbitrary static threshold, McAfee recommends that organizations use a single cloud threat detection engine that combines user behavior analytics (UBA) with machine to build behavior models based on actual user activity. In doing so, the solution can begin detecting threats automatically without any input from an administrator using an approach known as "unsupervised learning." In addition, the same threat detection engine should be used across all cloud services, including Office 365.

The reason organizations need a single threat detection engine for all cloud services is because cloud threats can involve the use of multiple cloud services. Organizations should cross-reference activity in Office 365 with other cloud services in order to detect threats. For example, a user who logs into Salesforce from New York City and then five minutes later logs into OneDrive from London may indicate a compromised account, since it would be impossible to travel this distance in such a short timeframe. Downloading a significant amount of corporate data from SharePoint and then uploading the content to an anonymous file-sharing service may also indicate an insider threat.

Recognizing that security incidents often involve more than one signal, McAfee recommends leveraging a threat funnel that combines multiple anomalous events together into a higher-order threat object before generating an alert. For example, a user who successfully logs in after several failed attempts may not require attention, unless the user is also logging in from a new location and exhibits behavior that deviates from their usual pattern, more strongly indicating account compromise. By focusing IT security analysts on the highest-probability incidents, organizations can minimize false positives and mitigate some of the effects of alert-fatigue.

While unsupervised learning makes it easy to get started, over time enterprises should provide input to fine-tune alerts. Organizations should employ what's known as "supervised learning" where security analysts provide feedback to models of behavior. This can be done in couple of ways:

- When reviewing incidents, marking an alert as a false positive is incorporated into behavior models.
- Analysts whitelist specific users or types of events to suppress them. For example, if an IT administrator is tasked with cleaning up dormant accounts and deleting large numbers of them, this activity can be suppressed for the user.

Capture an audit trail of activity for forensic investigations

Accurately detecting threats with UBA in the previous section requires complete visibility into all user and administrator activity. Additionally, security analysts require this information in the format of an audit trail to effectively investigate a wide range of incidents, whether it be data loss, insider threat, privileged user threat, or compromised account. For example, if an administrator is accessing data outside her job role, an audit trail of files accessed is essential. If a draft memo leaks to the press, it's important to know who accessed the file to narrow down who may be responsible. If an account compromise is found, the enterprise needs to know what data was accessed, particularly if the data accessed requires a breach notification.

McAfee's recommendation

There are over 500 distinct activities that users and administrators can perform across Office 365 applications. Enterprises should categorize these activity types into higher level categories (such as data access, data sharing, data deletion, and others) to normalize activities across cloud services and streamline the process of filtering.

Geolocation analytics should be used to report on activity across geographic regions and drill down into a single region for deeper reporting. For example, if an account compromise originates in Pakistan, where the company has no offices or employees, security analysts can see all access from Pakistan to understand if the scope is limited to one account, or if there is a broader problem that may have impacted several user accounts. events based on any attribute associated with an action, such as timeframe, activity name, file name, user, device type, and others. Analysts can combine multiple criteria into one search that filters a list of activities. These criteria can be saved as a custom view and shared with other McAfee Skyhigh Security Cloud users, and events can be exported to a CSV file and imported into other solutions. McAfee Skyhigh Security Cloud also leverages thirdparty threat feeds to enrich the context surrounding each event, including the reputation of the user's IP address and whether they are using an anonymizing proxy or TOR connection.

Prevent loss of corporate Office 365 data from use of personal Office 365 instances

Some enterprises have acceptable use policies that dictate that certain cloud services are allowed to store corporate data and others are not. For example, an organization may have a policy against uploading corporate data to all file-sharing services not managed by the organization or just to those that claim ownership of data uploaded to them. These policies can also extend to personal instances of corporate-sanctioned services such as Office 365. In these cases, enterprises only want to permit upload of corporate data to Office 365 accounts under management by the company. It would not be permitted to upload corporate data to a personal OneDrive account, for example.

McAfee's recommendation

In order to prevent corporate data from being uploaded to a personal instance of Office 365, enterprises need to understand whether a user is accessing a corporate instance of Office 365 or not. To do this, enterprises need to enforce persona-based controls to block access to personal instances while allowing access to the corporate instance. Keep in mind that preventing corporate data from being uploaded to a personal instance of Office 365 can only be enforced on managed devices. For that reason, it's imperative that enterprises first prioritize preventing corporate data from being stored in personal devices, starting with use case # 3 above (Block download of Office 365 data to personal devices), and extending it to all cloud services.

Prevent proliferation of malware

Recognizing that malware can leverage file sync and share functions to proliferate, Microsoft offers built-in antivirus for Office 365 that identifies most malware with previously cataloged signatures. When malware is detected, Office 365 quarantines files to prevent download and syncing to user endpoints. Additionally, many enterprises utilize a secure web gateway (SWG) with malware protection for all on-network devices and off-network managed devices, and endpoint protection for managed devices. However, there is a gap in protection for "zero-day" threats without previously cataloged signatures for off-network unmanaged devices.

McAfee's recommendation

Enterprises need a cross-cloud malware solution that extends protection against zero-day threats, where it executes suspicious files in a sandbox, leverages behavioral analysis to detect malware, and publishes indicators of compromise (IOCs). Infected systems should be quarantined to prevent syncing or sharing to other users who may not be covered by an enterprise's protection suite for malware.

Office 365 Data Protection Best Practices

Most organizations looking to protect data in cloud services like Office 365 have some form of data protection solution for their on-premises systems, including data loss prevention (DLP) for data in email and on endpoint devices.

When thinking about DLP for Office 365 the first thing to do is examine existing policies and the remediation actions being used for on-premises systems and identify the ones that will also apply to Office 365. This exercise will both ensure that data in Office 365 will be protected to the same degree it is in on-premises systems and reveal any policy gaps, such as new policies needed specific to cloud services like Office 365.

If an organization doesn't have a DLP solution for their on-premises systems, but need one for data going to Office 365 (which is an unlikely scenario), they must first identify sensitive data intended for the cloud, including regulated and restricted data, across the organization. To do this, the organization should develop a system to classify and map sensitive data against relevant internal policies and government regulations. From here, they can implement a solution to begin enforcing policies across this information.

Below are some of the best practices to implement when deploying DLP for Office 365.

1. Understand how Office 365 is being used

If an organization has already deployed a cloud service such as Office 365, a key first step is understanding how that service is being used. No action needs to be taken at this point. Instead, focus on getting granular visibility into how a cloud service is being utilized, including:

- The number of files containing sensitive data
- The number of files being shared outside the organization
- Anomalous usage events indicative of threats such as compromised accounts

2. Determine which types of data to look for that need to be protected

- Salaries
- Passports
- Social Security numbers
- Other personally identifiable information (PII)
- Account numbers
- Credit card numbers
- Spreadsheets with IP addresses
- File names containing "passwords"
- Outlook offline files (PST, MSG)
- Draft press release announcements
- Source code
- Encrypted files (Zip, PDF, XLS)
- Health records and other personal health information (PHI)

3. Gain visibility into collaboration

Employees love to collaborate via Office 365, but inadvertent sharing of data is one way for it to get lost. An organization should know how many files are being shared with internal employees, how many with external partners, and how many with personal email accounts (Gmail, Yahoo! Mail), so that they can educate employees on acceptable collaboration policies. This will also allow them to create and enforce sharing policies based on domain whitelist/blacklist and revoke untraceable shared links for files containing sensitive content.

4. Know about potential insider/outsider threats

Not all anomalous activities are a threat, but certain activity patterns should be a cause for concern and could be indicative of a real threat. Though making numerous failed login attempts to Office 365 might not necessarily be a sign of a compromised account, a user who successfully logs into a service and then logs in again from a faraway location within a short period of time is likely a case of stolen credentials. Understanding the frequency and the timing of these types of anomalous behaviors will lead to better DLP policies.

5. Define security policies such that the same policies for Office 365 can be enforced across all cloud services

The average enterprise uses well over 1,000 cloud services, 90% of which are unknown to the IT department. Employees store all kinds of sensitive and regulated data in the cloud—which accounts for 15.8% of all data stored in the cloud. In order to ensure that data is protected to the same extent across all cloud services, it's imperative that organizations manage their security policies from a single control point.

Once policies are defined, an organization should monitor (without taking any action) the instances of a policy violation being triggered. This is a triage phase that is intended to identify and remove all the false positives. This provides an opportunity to fine-tune your DLP policies so that they capture actual violations before you turn on automatic enforcement of policies.

6. Define Remediation Actions

Now that policies have been fine-tuned and false positives minimized, organizations should look at what the appropriate remediation actions should be if a DLP policy violation is triggered in Office 365. Common automated remediation actions can include blocking an upload, deleting or quarantining a file, modifying sharing permissions, revoking a shared link, or encrypting the data or the file while retaining a complete audit trail for forensics investigation. Policies should be enforceable on data in motion as well as data at rest to ensure complete data protection.

Get a Personalized Audit of Your Office 365 Usage Today

Get started with a free audit using McAfee Skyhigh Security Cloud. We'll deliver a report summarizing:

- Documents containing sensitive data
- Collaboration and sharing with third parties
- Anomalous usage indicative of insider threats
- Events indicative of compromised accounts

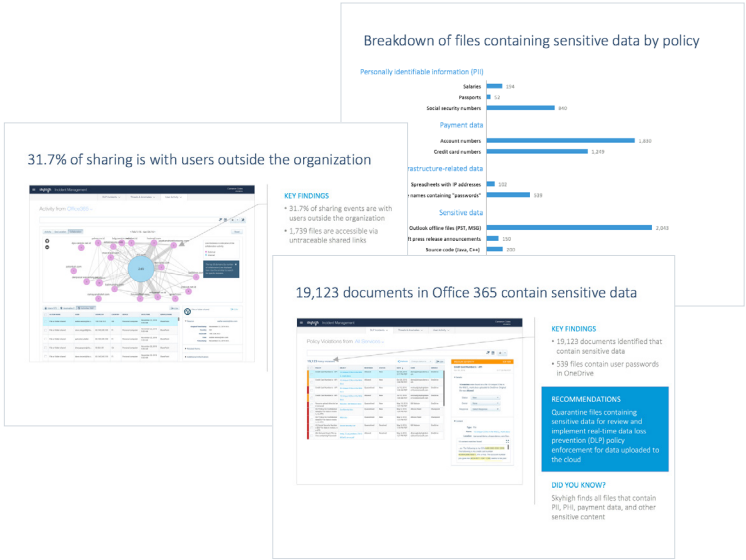

Request an Audit

bit.ly/O365audit

About McAfee

McAfee is the device-to-cloud cybersecurity company. Inspired by the power of working together, McAfee creates business and consumer solutions that make our world a safer place. By building solutions that work with other companies' products, McAfee helps businesses orchestrate cyber environments that are truly integrated, where protection, detection, and correction of threats happen simultaneously and collaboratively. By protecting consumers across all their devices, McAfee secures their digital lifestyle at home and away. By working with other security players, McAfee is leading the effort to unite against cybercriminals for the benefit of all.

www.mcafee.com.

McAfee™
Together is power.

2821 Mission College Blvd.
Santa Clara, CA 95054
888.847.8766
www.mcafee.com