+The Rising Tide

Why credential phishing and abuse is more dangerous than ever

This paper draws on the expertise of both MWR and phishd to cover the ways in which external attackers are obtaining credentials, what attackers are then able to do with these credentials and how companies can take steps to mitigate the risks.

Contents

Authors

William Knowles
James Moore
David Chismon

Introduction

No matter how effective your cyber defenses, your organization’s staff will need to access your systems for legitimate purposes and almost inevitably gain access by inputting credentials – such as a username and password combination. Should attackers discover these credentials, new vectors of attack will become available which may enable them to access your corporate network.

While organizations are aware of this problem, risk models are frequently based on an outdated understanding of what attackers can achieve once an employee’s credentials have been compromised. In recent years, advances in attacker methodology and tools have increased the potential impact of compromised credentials, while the functionality abused by attackers is often essential for maintaining business operations and hence cannot be disabled.

Total quantity of sets of credentials publically ‘dumped’ online each year. Figure above the bar shows the number of separate incidents of ‘dumping’ credentials that year.

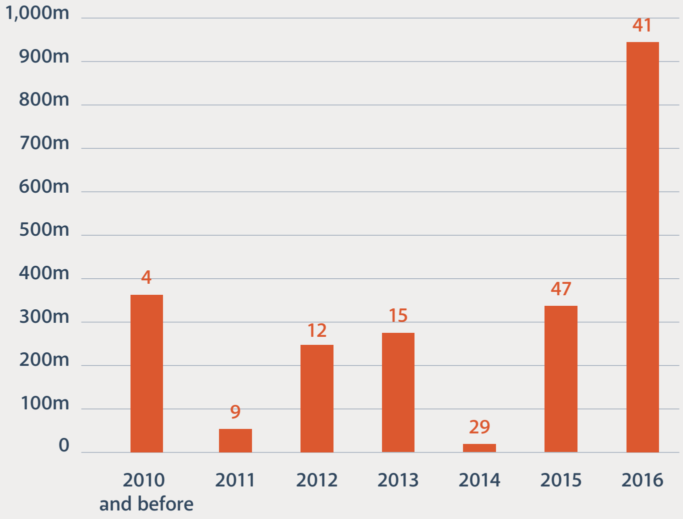

How Are Attackers Obtaining Credentials?

There are several methods typically used by attackers to obtain credentials. Here we discuss the four most common.

The Evolving Sophistication of Phishing Campaigns

Nowadays, both businesses and the general population have some awareness of phishing attacks or, as they are colloquially known, email-based 'scams'. Despite this, phishing continues to be highly successful, which prompts the question – why? Among the many potential reasons, MWR believes there are two key causes:

First, phishing taps into deeply ingrained patterns of behavior. Now that we conduct so much of our lives online, we open attachments or click a link without a second thought and it's very difficult to modify an individual's core behavior to verify the legitimacy of every email. The challenge is exacerbated by the technical understanding required to identify many phishing attacks, and the intellectual effort involved – even those well-versed in cyber security can find it mentally taxing.

Secondly, there's an abundance of information and tooling freely available online that allows attackers to create highly realistic phishing attacks. Some of this information is organization-specific, such as public-facing web portals identified through search engines and then impersonated, or examples of company communications that can be mimicked. Some attacks are individual-specific, such as those based around social media accounts, or an individual's widely publicized traits and interests (e.g. Facebook 'likes', tweets on Twitter, and pictures on Instagram).

Phishing attacks generally seek one of four outcomes:

- Convince the victim to enter valid credentials on a website controlled by the attacker
- Lure the victim to a website that the attacker can use to host exploits that would allow them to execute commands on the victim's computer
- Have the victim open a malicious attachment that allows the attacker to execute commands on the victim's computer
- Trick the victim into conducting an action that benefits the attacker, such as sending a document or processing a payment (see section on 'Pretexting a Path to a Password Reset' below)

MWR's phishd division conducts a large number of simulated phishing campaigns to help clients measure and improve their employees' susceptibility to phishing. Presented below is a summary of 100 recent phishd campaigns for 48 clients across a range of sectors. The findings have been aggregated into high-level categories based on the 'lure' of the phishing campaign used. For example, 'social media' might involve a connection request on a professional network, which links to a fake login screen. The table presents statistics of a subset of phishd campaigns that consist of three stages: getting the user to click a link, requesting credentials, and – once 'authenticated' – offering a file to download.

The effectiveness of a variety of phishd campaigns

Lure	Example	% clicked	Of those, % who then provided credentials	Of those, % who then downloaded a file
Financial	Invoice download	10.17	34.40	16.76
Technology	Secure email	14.17	65.02	81.13
Human resources	Appraisal system	18.21	73.86	72.43
Promotional	Discount voucher	19.46	63.26	87.19
Social media	Connection request	23.83	54.23	80.85

MWR has found social media to be the most effective lure to entice a user to click a link in an email (despite the emails being sent to work accounts), while financial lures were the least effective.

Once a user has clicked a link and an element of trust has been established, the likelihood of success in subsequent stages rises dramatically. Financial lures again achieved only limited success when credentials were requested. In contrast, lures based on human resources requests were the most effective, with promotional and technology campaigns following close behind.

Trust appeared to increase further at the third stage, as shown by the success rates, which were marginally lower for human resources but vastly improved for promotional, social media and technology.

In three of the campaign types presented (human resources, promotional, and social media), more than 10% of targeted individuals fell victim to the first two stages, leading to the disclosure of their credentials. More worrying is that for two of the campaign types (promotional and social media), more than 10% also downloaded a, potentially malicious, file.

Credential Dumps and the Risks of Password Reuse

LinkedIn, Dropbox and Adobe: each of these organizations has fallen victim in recent years to the large-scale theft of the credentials of individuals using their services. In some cases, the credentials have been offered for sale, while in others they have simply been publically disclosed (or 'dumped').

The breach of a third-party website can therefore lead to a breach of corporate data. It is difficult to remember a large number of passwords and unfortunately few people use password managers – hence the same passwords are often used for important corporate services as for less secure external websites.

Lower skilled attackers obtain credential dumps and then use scripts to automatically try the credentials against other services. This means anywhere the credentials have been reused can be discovered. More targeted attackers will study a target's passwords that may have been dumped to try and identify patterns.

phishd says...

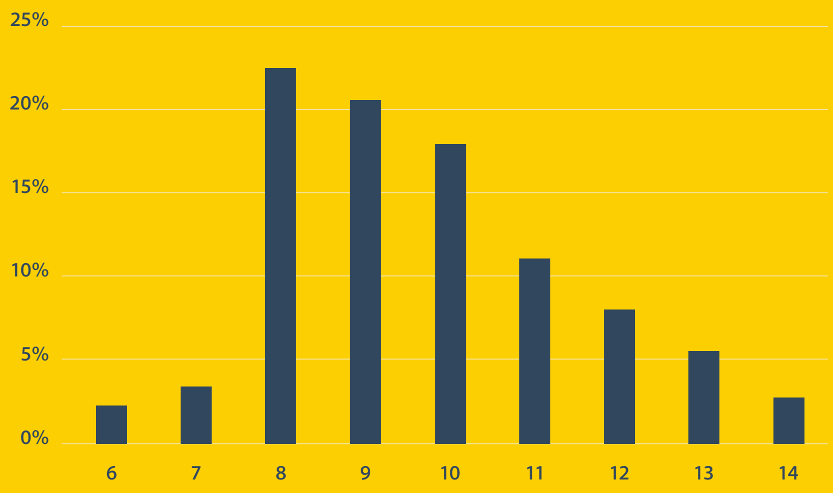

Over 60% of credentials obtained during phishd engagements were found to have a length of 8-10 characters – the mandatory minimum for many organizations – illustrating the tendency of individuals to adhere only to the minimum security requirements.

phishd says...

13.6%

The percentage of passwords obtained during recent phishd campaigns that ended with four numbers in the range 1940-2040.

5.5%

Of these, 40.4% ended in 2016 – that's 5.5% of all passwords ending in the year in which they're set. This method of circumventing complexity requirements is a gift for attackers.

The Recurring Case of Weak Passwords

Credential reuse from a password dump requires that the targeted individual has used the same username and password on multiple websites. However, credential dumps also reveal a truth known well to security professionals: people often choose insecure passwords.

What makes a password usable and secure has long been debated but sadly the conclusions aren't widely known. As one individual proudly told an author of this paper, "My password is a name you wouldn't guess and a year that only I know." This illustrates an understandable lack of awareness as to how passwords are attacked. The individual assumed attackers would try to guess words they would use. In reality, attackers will use programs to try hundreds of thousands or millions of attempts a second and so patterns such as a name and a year will quickly be cracked. It is the responsibility of security professionals to educate users on setting good passwords, as well as protecting them from the impact of failing to do so.

A weak password not mitigated by other security controls is dangerous for two reasons. First, attackers might be able to brute-force a login. Brute-forcing of authentication mechanisms, either by targeting a single user with multiple passwords or attempting a single password against a large number of users to avoid account lockouts (also known as 'password spraying'), continues to be a highly utilized and effective technique, requiring only one individual with a weak password to succeed.

Secondly, there is an increased chance of an attacker discovering a usable password in a credentials dump. As security awareness has increased within the developer community, the incidence of sensitive information, such as credentials, being stored in plain text has rapidly diminished. Credentials are instead stored as a variation of a cryptographic hash algorithm – a one-way conversion of a particular input into a fixed-length, unique (i.e. seemingly random) output. In the case of password verification, the process involves converting a user's submitted password into a cryptographic hash format and comparing it to the credential stored in the same format. If the cryptographic hashes match, the password is verified without anyone but the user knowing the original password for longer than the authentication process.

As a consequence, if hashed credentials are dumped (or obtained through another attack vector), the attacker will need to follow the same process to recover a password: choose a password, hash and compare. This can be a slow, painstaking process, accelerated by focusing on dictionaries of passwords that are weak, previously disclosed, or possibly target-specific (e.g. variants of a company name). For many cryptographic hash types, there are freely available databases of pre-computed hashes based on such dictionaries (also known as 'rainbow tables'), avoiding the need for computationally heavy comparisons. For an attacker, the use of weak passwords means the protection of a cryptographic hash is little more than a minor obstacle, as such passwords frequently appear in these databases.

Pretexting a Path to a Password Reset

Phishing is not, of course, the only form of social engineering attack that organizations must guard against. Arguably more effective, but more challenging and riskier to execute, is the range of attacks that involve direct personal interaction. Such attacks typically involve the creation of a 'pretext' – a scenario that aids in eliciting the desired information from a victim.

In the context covered here, an attacker wouldn't target the owner of the credentials, but an individual involved with organizational processes. For example, common attack vectors that utilize pretexting involve the impersonation of an individual and the request of a manual password reset (e.g. through an IT Help Desk), or the request for seemingly innocuous information (e.g. from co-workers, HR or Finance) that enable an attacker to provide answers to the 'secret questions' stored in automated password reset mechanisms.

Despite the sophistication and risk involved, threat actors continue to attempt this type of social engineering attack with varying success, depending on the target organization's security processes.

The wealth of online information about a target individual and organization is a great help to attackers. For example, a relatively rudimentary password reset could arise as follows. A target with a key set of responsibilities is identified through LinkedIn employment descriptions, while their social media profiles reveal they are on vacation. A possible pretext would be to impersonate this individual, telling the IT Help Desk they are on holiday but have received a call from a colleague (named through LinkedIn reconnaissance, if required) about an urgent email. As their current password is too complex to remember and stored in a password manager on their work computer, they need a password reset. This pretext is predominantly based on information outside the control of the organization, but it is the organization's responsibility to manage the consequences.

What Can External Attackers Do with Credentials?

Now that we've established how an attacker might obtain one or more sets of valid credentials, we can look at how these credentials are typically abused. This section discusses four scenarios that utilize recent advances in offensive security tooling, and the exploitation of common infrastructure configurations.

Temporary or Persistent Email Compromise

Organizations frequently provide staff with access to emails when not on the corporate network. Typically, this is through a web portal or direct connection to the mail server using Exchange ActiveSync protocol or IMAP, hence exposing corporate mail servers (or proxies to them) to the internet.

An architecture with email access exposed to the internet is particularly vulnerable to attackers obtaining and replaying credentials.

Once user credentials are obtained, attackers can access emails by:

- Logging on to a web portal to browse and extract emails
- Using a legitimate mail client to synchronize (extract) all emails from the mail server for offline use or redistribution
- Using a custom tool to enable large-scale downloading of messages

Two recent tools that exploit Exchange functionality are MWR's PEAS[1] and SensePost's Ruler[2]. PEAS communicates with Exchange using ActiveSync, which is typically employed for synchronization with mobile devices and can be used to retrieve mail. In contrast, Ruler communicates using Messaging Application Programming Interface (MAPI), which is employed for synchronization with Outlook and can be used to set mail rules on the Exchange Server.

The following example is a five-step attack scenario that uses these tools. First, the attacker connects to an externally facing Exchange Server using PEAS, and dumps all mail for a particular account or set of accounts using PEAS or a similar tool. Thus the attacker has access to all historical mail, to be searched for useful information or publically dumped to embarrass the victim.

Secondly, in either a separate or follow-up attack, Ruler is used to connect to the Exchange Server to create a mail rule that forwards selected mail to an attacker-owned account when a condition is met. The rule can apply to both sent and received mail, with example conditions including mail of a specific size, containing an attachment, or containing a specific word or words in the address, subject, or body. Thirdly, this mail rule is synchronized to all Outlook instances configured for that user's account (i.e. that of the legitimate user). Fourthly, the user's mail is also synchronized. Finally, the mail rule is applied to all new mail and when the conditions are met the rule forwards a copy of that mail to a mailbox accessible by (but not attributable to) the attacker. In this way the attacker obtains persistent email access, while aiming to avoid discovery by taking advantage of the generally less stringent outbound mail controls.

Key Points:

- A full history of emails can easily be downloaded by attackers using standalone tools.
- Mail rules created on Exchange are synchronized to Outlook clients.
- Rules can be set to automatically forward new mail to an attacker-operated address.

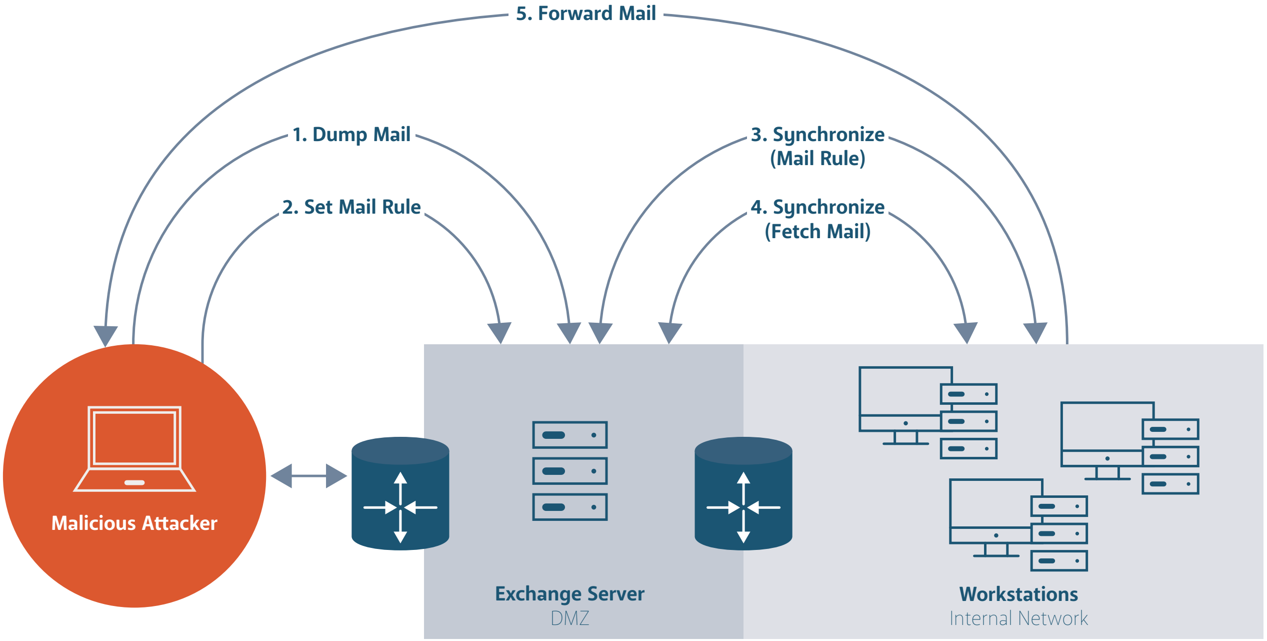

[1] https://github.com/mwrlabs/peas
[2] https://github.com/sensepost/ruler

Single Sign-On, Single Point of Failure

Single sign-on (SSO) allows a user to sign on once and be automatically authenticated to other services. It is increasingly used to improve operational workflow, but it comes at a cost – SSO has the same drawbacks as credential reuse, in that one compromised password can lead to mass compromise of distinct services.

The adoption of cloud services by business exacerbates the problem, as many implementations use SSO. As businesses do not want their employees to have to manage multiple identities, they will typically favor SSO authentication to cloud services where available. This often means an attacker can gain access to whatever cloud services the company is using, be it email/messaging, CRM, finance or HR systems.

The following example is a three-step attack scenario that takes advantage of SSO. First, the organization establishes a set of legitimate accounts that can be used to access cloud services. This example is based on the increasingly prevalent use of Active Directory for this purpose, which can involve the synchronization of user accounts with a cloud-based solution (e.g. Azure Active Directory or Google Apps Directory Sync).

Secondly, having obtained one or more credential sets, the attacker accesses the organization's cloud services. Thirdly, as the credentials are legitimate, authentication and authorization occurs between the cloud service and Active Directory (either the cloud-based or on-premises instance). Steps two and three repeat for each cloud service the attacker wishes to access, while typically traversing and involving third-party networks and systems that the organization has limited capability to monitor. This is a simple and effective exploitation of an increasingly common configuration, involving no advanced techniques or tooling.

Key Points:

- Single sign-on improves workflow, but decreases information requirements for attackers.
- With one credential set, attackers can now access multiple business services.

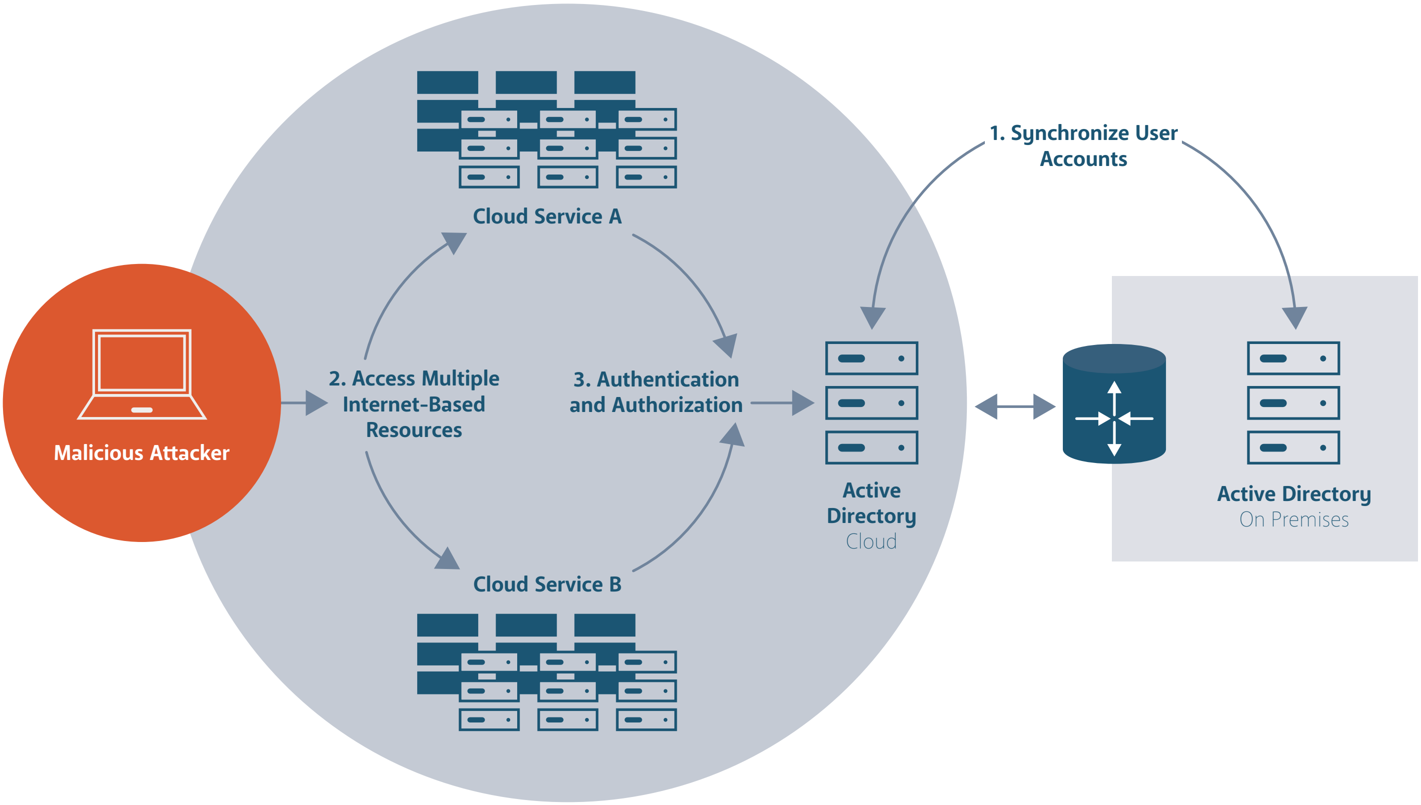
Cloud Service A
1. Synchronize User Accounts
Malicious Attacker
2. Access Multiple Internet-Based Resources
3. Authentication and Authorization
Active Directory Cloud
Active Directory On Premises
Cloud Service B

Accessing Internal File Shares – Externally

Microsoft's Exchange Server is not only a mail server, it also provides a broad range of extended functionality to enable, for example, management of mobile devices as well as scheduling and collaboration. A conspicuous example of the latter is how Exchange allows the accessing of file shares and SharePoint within the organization. The file share and SharePoint access functionality can be exploited using MWR's PEAS to access internal files from outside the organization.

The following example attack scenario involves a number of steps. First, PEAS is used to communicate with the public-facing Exchange Server to query the system for a list of available file shares or the files within that file share. These file shares might not exist on the Exchange Server but on other internal machines, hence the Exchange Server will need to be able to talk to the domain controller. Therefore an attacker is likely to be able to obtain access to shares stored there. An external attacker can obtain the domain controller's name from the company's internal DNS under the domain name – and the domain name can usually be guessed or found through open-source intelligence.

The domain controller has a number of shares that are integral to its functionality, such as SYSVOL and NETLOGON. These shares could contain credentials stored in group policy preferences, along with files that reference other shares. Once file shares have been identified, PEAS can use the same approach to request files from those file shares. In this way, an external attacker can achieve data exfiltration from internal machines, using only one set of credentials.

Key Points:

- Exchange contains the functionality to list file shares, and to download and save files to them.
- File shares can be Windows file shares or SharePoint sites.
- Attackers can access internal file shares through an external connection to Exchange.

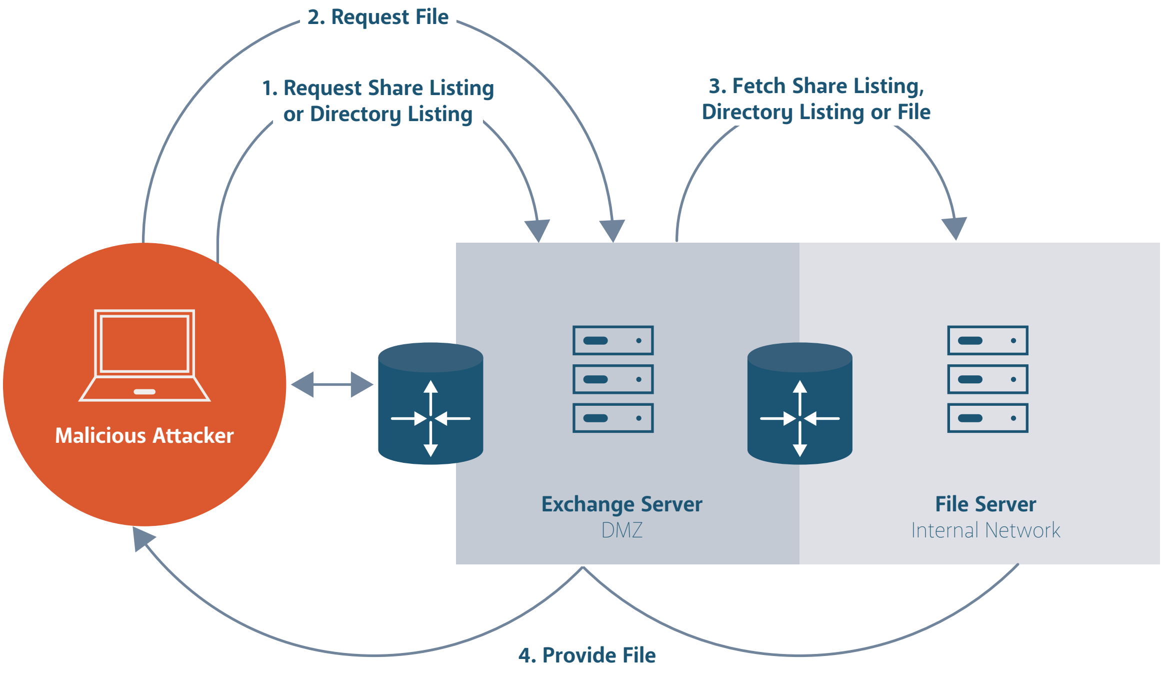
2. Request File
1. Request Share Listing or Directory Listing
3. Fetch Share Listing, Directory Listing or File
Malicious Attacker
Exchange Server
DMZ
File Server
Internal Network
4. Provide File

Setting Rules for On-Demand Code Execution

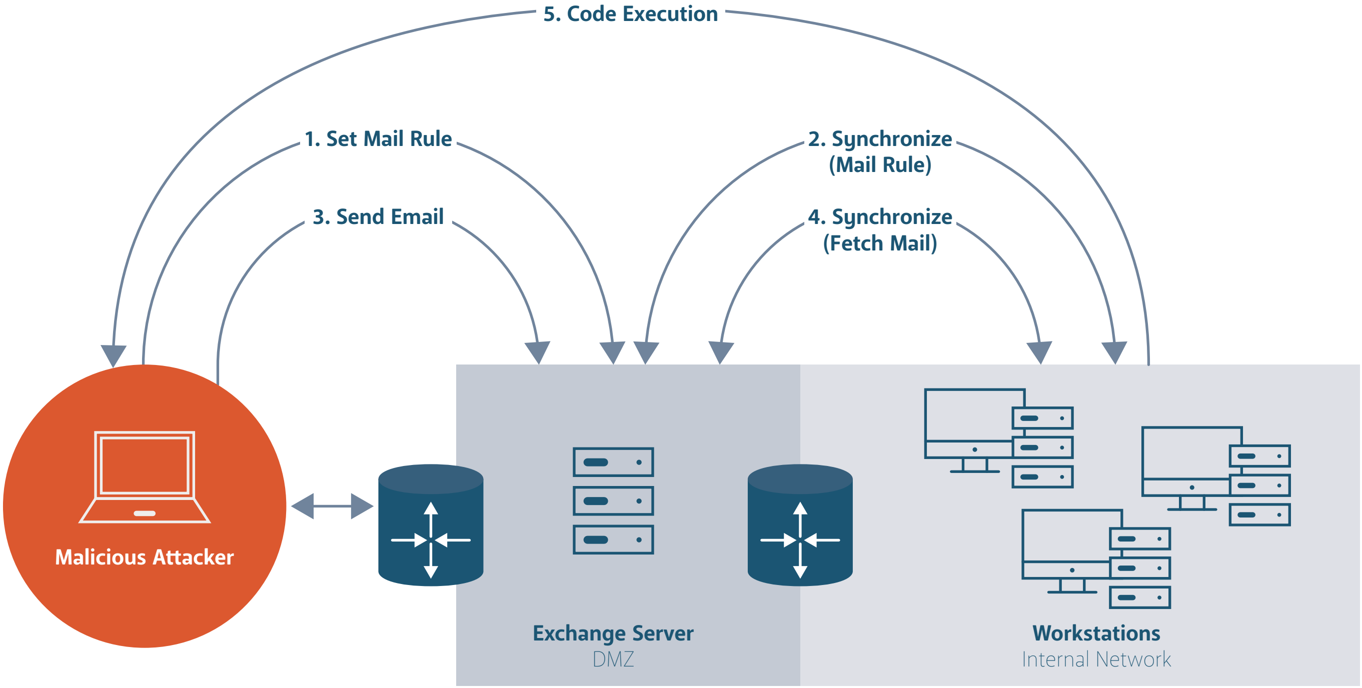

Setting rules for forwarding email is not the only way that mail rules can be abused: they can also be configured to launch an application[3], enabling an attacker to arbitrarily run code on a victim's machine. Again, this functionality can be exploited using SensePost's Ruler.

The following is an example of an attack against an Exchange server. First, the attacker connects to the public-facing Exchange Server and sets a mail rule to launch an application when a condition is met (e.g. a mail with a particular subject line). This rule is synchronized to Outlook on the user's workstation and the attacker sends a mail that meets the condition. The user's Outlook retrieves the mail, the condition is met and the application launched.

In practice, an attacker would no doubt prefer the user to execute an attacker-controlled file (e.g. a payload that allows easy escalation to a shell on the compromised system), rather than an arbitrary file on the user's system. However, the location of the application to be launched is specified using the Universal Naming Convention (UNC), which includes a hostname (e.g. \\hostname\file.bat) – hence the file to execute can be on a remote system, such as an attacker-controlled SMB share or WebDAV server hosting their payload of choice. From this point on, the attacker has malware running on the machine and could deploy a remote-access Trojan to compromise the internal network.

Key Points:

- Mail rules created on Exchange are synchronized to Outlook clients.
- Rules can be created to launch an application when particular mails are received.
- An attacker could set rules on Exchange externally, but gain code execution on internal workstations.

[3] Since the release of Ruler, additional tools have been developed to create mail rules post-exploitation of a system (e.g. for file execution to establish persistence). These tools target not only Windows (see MWR's XRulez: https://github.com/mwrlabs/XRulez), but also OSX's Mail.app.

Mitigating Credential Abuse by External Attackers

Having looked at the ways in which a malicious attacker might obtain credentials and how these credentials could be abused, what can we do to defend our systems?

Here we consider the controls MWR has found to be effective in mitigating risk. Some are broader controls, while others are specific to attacks described above.

Recommended Control	Effectiveness	Effort
Disable File and SharePoint Remote Access	High	Low
Two-factor Authentication	High	Medium
Secure Architecture Design	High	High
Require Client Certificates	High	Medium
Lock Cloud Services to Company IPs	Medium	Low
Alternatives to Traditional Complexity Requirements	Medium	Low
Training to Encourage Security-Conscious Behavior	Medium	Medium
Monitoring for Unusual Activities	Medium	High
Monitoring for Credential Dumps	Low	Medium

Disable File and SharePoint Remote Access

Effectiveness	Effort
High	Low

Potentially dangerous functionality with no legitimate corporate requirement should be disabled. Ideally, product vendors release the software with dangerous functionality disabled and leave the user to re-enable it.

However, this process is not always adhered to and in the section 'Accessing Internal File Shares – Externally', the functionality abused by the PEAS tool is enabled by default, despite few organizations requiring it. In this case, organizations are advised to disable the functionality by setting the UNCAccessEnabled and WSSAccessEnabled parameters of the ActiveSync mailbox policy to false[4].

[4] For guidance see: https://technet.microsoft.com/en-us/library/bb123756(v=exchg.160).aspx

Two-factor Authentication

Effectiveness	Effort
High	Medium

Traditionally, authentication mechanisms require a single factor: something you know, namely a password. Two-factor authentication requires a second factor – not a password but something you have, such as a device that periodically generates tokens, or something you are (i.e. biometrics).

phishd says...

Two-factor authentication is a vast improvement but not a panacea. It is possible, for example, to relay a time-based factor (e.g. a token) during a phishing attack to achieve a transient (one-time) authenticated session. Approaches such as U2F aim to mitigate such phishing and man-in-the-middle attacks.

Two-factor authentication should be implemented wherever possible, and invariably for access to all critical business assets. In the event of a breach and credential dump, an organization is immediately exposed but an attacker's efforts are stifled by the requirement for a second factor.

Many externally facing systems (e.g. email, VPN, and cloud services) support a variety of two-factor mechanisms, although this may not be possible for 'always on' services (e.g. email delivered to mobile devices through technologies such as Exchange ActiveSync). Guidance on authentication and multiple factors is being prepared by NIST and is in draft at the time of writing[5].

Secure Architecture Design

Effectiveness	Effort
High	High

A secure architecture with appropriate network segmentation is a key part of defense in depth strategy, restricting exploitation of a product to just that service.

Exchange is a particularly good example of a service that benefits from segregation, as it's externally accessible, complex and contains an aggregation of assets (emails). Should an attacker compromise Exchange and gain remote code execution, they would have access not only to emails but potentially to hosts in the internal network with which Exchange communicates.

In an ideal world, all systems would be robustly segregated, but as this entails substantial operational overheads, focus is generally given to higher-risk systems. Any systems that can be accessed externally are best placed in a demilitarized zone (DMZ), subject to a restricted security policy – for example, segmented from the internal network to ensure an attacker cannot pivot to internal file shares through an instance of Exchange. However, Exchange requires connectivity to the domain controller and hence while attackers can be prevented from communicating with internal file shares through architecture, they cannot be prevented from communicating with the domain controller.

[5] NIST 800-63B Digital Authentication Guideline: https://pages.nist.gov/800-63-3/sp800-63b.html

Require Client Certificates

Effectiveness	Effort
High	Medium

In the majority of connections using SSL/TLS, only the server certificate is validated. This makes sense for consumer internet traffic as, say, a shop cannot expect customers to have a valid certificate to communicate with it. Organizations, however, have control over the devices that legitimately access external services, so it's reasonable to require clients to authenticate themselves with a provisioned certificate prior to establishing a secured connection.

This is a particularly useful control for protecting services accessed by an automated process, such as a mail client polling for email through Exchange ActiveSync. Such automated services do not lend themselves to two-factor authentication and the devices using them often support client certificates. This control can also be used on user-accessed services (such as Outlook Web App), requiring the employee to use their provisioned device.

In organizations with mature public key infrastructure and centralized management of devices (through group policy and mobile device management), this approach is low effort to introduce and maintain.

Lock Cloud Services to Company IPs

Effectiveness	Effort
Medium	Low (if VPN / proxy already in place)

One way to prevent an attacker from accessing an organization's subscribed cloud services with compromised credentials is to restrict access to the cloud service to pre-defined IP addresses, ensuring the service is being accessed from a trusted network. Mobile devices, including laptops used for remote working, should be configured to connect to such cloud services through a VPN to the company network, so the visible IP address is that of the corporate network. This approach also ensures the employee uses only their corporately managed device to access the cloud services.

Although this safeguard would frustrate an external attacker with compromised credentials, it has limitations. Not all cloud service providers support this functionality, and others might do so only in restricted circumstances (e.g. for Google services, only if Single Sign-On is used with Active Directory integration). Furthermore, some organizations allow use of their subscribed cloud services from personal devices, which would be prevented by IP restrictions.

Alternatives to Traditional Complexity Requirements

Effectiveness	Effort
Medium	Low

phishd says...

34.9%

The proportion of credentials obtained during a phishing campaign where the password consisted of an upper-case first letter (A-Z), a string of lower-case letters (a-z), and then a series of numbers (0-9) with no symbols.

Traditionally, password management has enforced complexity on users to avoid weak passwords, discourage password reuse, and blacklist common passwords.

However, the National Cyber Security Centre (NCSC), part of GCHQ, released guidance[6] in 2015 that challenges the way complexity requirements are best enforced.

NCSC argues that the traditional approach places too heavy a burden on individuals, especially with mandatory password changes every 30-90 days and more than one organization's passwords to remember. The result is minimal variation when changing passwords to meet complexity requirements.

Instead, NCSC argues that organizations should encourage high memorability while retaining complexity; for example, through the use of passphrases. In addition, it recommends greater technical defenses, such as throttling logon attempts, locking out accounts after multiple failed attempts and protective monitoring. This is an approach that MWR finds effective at improving credential security.

Monitoring for Unusual Activities

Effectiveness	Effort
Medium	High

As no protection measures can be guaranteed to prevent attacks, organizations should monitor and investigate unusual activities, particularly around critical assets.

While it's worth setting up alerts for known malicious behavior, a skillfully conducted attack on externally accessible servers – such as mail and cloud servers or VPN devices – can be hard to detect.

Monitoring certain behaviors, such as quantity of forwarded mail per account, raw data transfer from servers to remote hosts and user agents per account, can help identify potentially malicious activity.

It's strongly recommended that organizations are readily able to access logs for all key devices and services. For example, an organization might be practiced at accessing logs from key devices, but not from cloud services.

[6] https://www.ncsc.gov.uk/guidance/password-guidance-simplifying-your-approach

Training to Improve Employee Security Behavior

Effectiveness	Effort
Medium	Medium

Raising security awareness alone is not enough to change employees' password-setting habits. A change in security behavior is required, which demands an improvement in security culture through continual education, empowerment and evaluation of employees.

MWR recommends ongoing training and evaluation to reinforce security-conscious behavior. The evaluation provides insight into which training is working, which isn't, and where there are opportunities to improve. For example, where certain individuals are susceptible to specific types of phishing attack, evaluation raises awareness and this in turn enables targeted training.

Training employees to recognize phishing emails reduces the number of employees who click on a link, but MWR recommends focusing on increasing the number that report the malicious email, and reducing the time it takes them to do so. Organizations can adopt 'attach and forward' mechanisms for reporting emails to security teams, or add-ons to mail clients, such as those offered by MWR's phishd[7].

phishd says...

3%

The proportion of employees targeted by our simulated phishing attacks who reported the attack correctly... while 25% clicked on a link within the phishing email. It's not enough for employees to be aware of phishing. They need to stop clicking – and start reporting attacks correctly.

Monitoring for Credential Dumps

Effectiveness	Effort
Low	Medium

Once compromised credentials have been publically dumped online, there's an immediate risk to an organization's security. Hence proactive monitoring of such dumps is advisable so that reactive controls can be triggered, such as forcing password resets. The monitoring is time-consuming and presents legal challenges due to the nature of the data, but there are third-party commercial services that – based on identifying information (such as employee email addresses) – will alert organizations when relevant data is found within credential dumps.

However, if credentials are dumped privately, such as to buyers on the darknet or to private hacker forums, then the compromise might not be detected.

[7] https://www.phishd.com/our-services/microsoft-outlook-plugin/

Summary

There are various ways in which attackers first obtain compromised credentials, and then abuse those credentials to gain further access. However, while many organizations assume that only the email portal (such as Outlook Web App) would be accessed in an attack, in reality the attacker can dump entire mailboxes, access file shares inside the corporate network, execute programs on the victim’s computer and explore any services using the same credentials.

Robust defense therefore needs to focus on the following key areas:

Predict: Monitor the internet for new attacks or dumped credentials.

Prevent: Train employees to report malicious emails and accept that credentials will occasionally be compromised. Focus on building controls that assume compromised credentials.

Detect: Monitor externally accessible servers for which an attacker might gain credentials, such as mail servers and VPN servers.

Respond: Ensure incident response teams are well-rehearsed at performing investigations into such services as mail servers, VPN devices and cloud services.

By building defenses in such a way that the entire security model isn’t broken when an attacker obtains the username and password of an employee, organizations gain a strong chance of surviving malicious tactics that are already known – along with those that have yet to be identified.

About MWR:

Established in 2003, MWR is a research-led cyber security company with offices around the globe.

We provide specialist advice and services in all areas of cyber security, from professional and managed services, through to developing commercial and open-source security tools. Using a threat-centric approach, we focus on working with clients to develop and deliver security programs tailored to meet the needs of each individual organization.

In a rapidly changing technology landscape, innovation is essential and our ambition to push boundaries sets us apart. Central to MWR's philosophy is the desire to deliver high-quality cyber security services and unsurpassed levels of support to our clients.

About phishd

phishd by MWR InfoSecurity is the global leader in designing and delivering fully managed employee security behavior programs.

Our wide array of services is designed to add an extra line of prevention and detection to your security posture, ranging from simulated phishing attacks to password auditing, web-based training and email client plugins. Seeing the world from an attacker's perspective gives us a powerful insight that enables our clients to reduce their risk from targeted attacks – rather than creating a false sense of security by simply meeting the relevant compliance and regulatory requirements.

Our services are trusted by a diverse range of organizations across the globe, from SMEs, through FTSE 100 and blue-chip companies, to Government departments – all of which rely on phishd to measure and improve their employees' security behavior.

www.mwrinfosecurity.com

www.phishd.com

Follow us:

MWR

phishd