



The Microsoft 365 kill chain and attack path management

Learn how attacks unfold and discover effective strategies for defending your cloud or hybrid environment.

Quest

Introduction

“Know thy enemy” is a cornerstone of an effective cybersecurity and cyber resilience strategy. Unfortunately, however, many organizations have an erroneous understanding of how modern cybercriminals and cyberattacks actually work — and that misunderstanding prevents them from implementing effective controls and processes to defend themselves.

This white paper debunks the key misperceptions about modern cyberattacks and lays out the true phases of a cyberattack, reviewing some of the most common tactics and techniques used in each phase with a particular focus on the Microsoft 365 environment. Then it dives deep into one strategy that adversaries use to achieve their goals — exploiting attack paths — and provides effective options for defending your organization.

Background: The truth about modern cyberattacks

Movies, television shows and even news reports tend to present a very flawed vision of cybercrime. They often make it seem that cybercriminals are evil masterminds who breach a corporate network, bring it to its knees and receive a fortune in Bitcoin, all in a few carefully chosen clicks. That view can lead to a disjointed cybersecurity strategy in which organizations desperately throw up whatever barriers they can find to try to block intruders from getting into the network, or are unwilling to invest much in cybersecurity at all, out of cynicism that any measures could offer true value.

Fortunately, an accurate understanding of how modern cyberattacks actually work provides a path to effective defenses. Here are three critical truths to keep in mind.

Cyberattacks often unfold over weeks or months.

It doesn't make for exciting television, but cyberattacks are often a slow process. Just because you're breached doesn't necessarily mean that attackers are going to swipe all your valuable data right away. For example, even though it's often said

that ransomware encrypts files at lightning speed, that doesn't mean that there are mere moments from the start of the attack to its final impact.

Indeed, IT security pros even have a name for the time from the point of infiltration to the point of detection: dwell time. Different studies claim different average dwell times and disagree about how it has changed over time, but most of them agree that intruders often roam around in the network for weeks or even months before they either finalize their attack or are detected. It's not uncommon for forensics investigations to uncover evidence of sustained activity over months or even years.

Modern cyberattacks are often human operated.

This brings us to a related point: Cyberattacks are typically not like lobbing a grenade into a building and waiting for it to explode. Rather, many of them are well orchestrated and human operated. Living, breathing human beings are actively navigating through the IT environment, using their knowledge and skill to find information that enables their attack to be more effective. Well-known human-operated ransomware campaigns include REvil, Samas, Bitpaymer and Ryuk.

Modern cyberattacks are often well orchestrated and human operated.

For example, an adversary might gather information about the organizational hierarchy and the specific terminology used in the business to craft a truly compelling phishing or spear-phishing email. A ransomware actor might not be content to encrypt just the data that is accessible by the first user account they compromise, but instead invest the time needed to find the organization's most valuable content. For instance, adversaries have specifically worked to access the email and files of specific executives to find sensitive or embarrassing information that gives them increased leverage in extorting ransom.

Cybercriminals don't have to be masterminds.

The stereotype paints the cybercriminal as a lonely, disgruntled hacker with extensive technical skills. But the reality is, cybercrime is now a business with multiple actors involved. In addition to the technical person who creates the malware, there are the non-technical parties who hire them to launch attacks and the brokers who sell access to corporate environments. In particular, the ransomware-as-a-service model is now so well established that it has its own acronym: RaaS.

Ransomware as a service (RaaS) enables non-technical criminals to launch devastating attacks.

Indeed, even the hacker running the attack does not necessarily need to be particularly clever; they often can simply take advantage of common configuration mistakes and known vulnerabilities that have been left unpatched. Indeed, Microsoft's analysis of 43 trillion cybersecurity signals in their platform revealed that 80 percent of ransomware attacks can be traced back to misconfigured software or devices.

The key phases of the Microsoft 365 kill chain

So, what does a cyberattack look like? MITRE provides a knowledgebase of the techniques and tactics used in cyberattacks, organized according to the phase of the kill chain in which they are used. The full matrix comprises 13 phases, but for the Microsoft 365 kill chain, we can simplify it to the following five phases, as illustrated in Figure 1:

- Reconnaissance (recon)
- Initial access
- Persistence
- Discovery
- Exfiltration (exfil)



Figure 1. The key phases of the Microsoft 365 kill chain

Reconnaissance

Before launching a cyberattack, adversaries need to determine which organizations to target. They have a wide range of resources at their disposal; here are some of the most common ones. Note that many recon tactics are impossible for IT teams to detect because they occur on third-party sites.

DNS

Recon often begins with DNS. Indeed, DNS is a fantastic place to find information about a possible target: All of DNS is public by its very nature, and it's very easy to identify use of Adobe Creative Cloud, AWS, Salesforce and many other cloud services simply by checking for the corresponding domain validation records. A telltale sign of Microsoft 365 is the mail exchange (MX) record. While it's possible for organizations to clean up their domain validation records, more often than not, they ignore them, leaving valuable intel for adversaries to pick up at will.

Recon often begins with DNS.

Social media

Hackers can also gain valuable information about an organization by looking at social media. For example, they can harvest details about a company's employees from their LinkedIn pages and use it to create more effective phishing campaigns and other attacks, as described in the section below on the initial access phase.

While LinkedIn is the most obvious resource, there are many other social media apps that can be exploited. For example, BeReal prompts users to regularly post two pictures: a selfie and a picture of what they're

currently doing. That might seem innocent enough — but if the user is at work, the photo might well show their laptop screen with their current Microsoft Teams meeting or chat, or even sensitive data like patient records or intellectual property. Today, it's all too easy for employees to serve valuable information to attackers without even realizing what they've done, and cybersecurity training often does not do enough to educate them about the risks inherent in social media.

Certificate transparency logs

Anyone can look up a domain and see a list of all the certificates that have ever been issued to that domain. That matters even when we're focused on Microsoft 365 because many organizations still retain an on-premises footprint of some sort. For instance, even after a company has moved to Exchange Online, they might well keep an on-premises Exchange Server in hybrid configuration for use for admin tasks. That's a perfectly valid and useful strategy. However, all too often, organizations relax their security practices around their remaining on-premises systems — and systems that are not kept current on updates and patches are appealing targets for adversaries.

Email address lookup

Among the many other readily available tools for recon is the site hunter.IO. Simply enter a company's domain and website will spit out not only the most common pattern used for that domain's email addresses, but a list of actual email addresses that are only partially redacted. That's a great way for an adversary to start building a phishing list — and it's also a great way for them to start creating a user list for credential-based attacks, since a user's email address is often their username as well.

Initial access

Once an attacker has chosen an organization to target, the next step is to gain a foothold in their network — often by exploiting the intel gleaned during the recon phase.

Phishing

Phishing attacks are still the number one attack vector because they are often the easiest way to get into an IT environment. It's true that both corporate and public cybersecurity education campaigns have made

significant progress in teaching users what a phishing email look like and the importance of not clicking on links or opening attachments if they're not confident of the sender. However, phishing has morphed from clumsy emails filled with typos into highly persuasive and targeted spear-phishing messages that can fool even savvy users, especially if they're focused on a hectic workload or juggling multiple priorities.

Phishing attacks remain the top attack vector.

One common spear-phishing technique is to make use of intel from social media. For example, LinkedIn profiles regularly include an employee's alma mater and possibly even their graduation year. With that information, an adversary can craft a compelling message to a particular individual that appears to be from their university and have a really good chance of tricking the recipient into clicking a link that purports to lead to a special discount on college-branded gear or opening an attachment that claims to contain reunion information for their specific class.

Adversaries also make use of intel about the company's employees and hierarchy. They can craft an email that appears to come from the targeted individual's CEO that directs them take steps that will benefit the adversary, such as providing confidential information. By giving the email an urgent tone, they can lead the user to believe that if they don't take the specified action, the company will be in jeopardy. Similarly, spear-phishing messages that appear to come from the company's IT team can be a great way for an attacker to harvest user credentials that enable them to slip into the network undetected.

Other social engineering attacks

Phishing is one form of social engineering attack — an attack in which an adversary uses human (social) interaction to obtain information. But there are many other variations. Vishing leverages voice communication, such as a phone call, rather than email, and smishing exploits text (SMS) messages. Baiting sites lure victims in with the promise of freebies, such as a download of movies or music,

and then harvest credentials or plant malware. Cybercriminals may also contact users and pretend to be responding to a request for help with a product or service the company probably uses; then they trick the user into running certain commands or even granting remote access to their device.

App requests (OAuth abuse)

OAuth is an open standard for authentication and authorization that enables a third-party app or service to use someone's account information to perform tasks on their behalf, after requesting the user's permission to access the necessary information. These requests are common on smartphones; for instance, installing a new app often results in a request for access to your location, calendar or photos. Indeed, they are so common that users often click Accept without a second thought.

When that happens on a device that a person uses for work, they might be enabling an intruder to their email, access all the data they have permissions to and so on. While "unverified" now appears underneath the name of unverified apps, not all users pay close enough attention. If it says "Microsoft" at the top, it must be fine, right? Microsoft enables IT teams to configure how users consent to applications, and they recommend allowing users to consent only for apps from verified publishers and only for selected permissions. However, many organizations allow users to freely consent to app requests.

Account compromise from password reuse

Adversaries often get access to corporate networks using dumps of stolen credentials that are sold and published on the dark web. The website haveibeenpwned.com enables users to check whether their email address or phone number has been compromised in a data breach. It keeps a running total of them: Currently, it's at nearly 12 billion!

One good strategy for blocking attackers with stolen credentials is to require multifactor authentication (MFA). However, many organizations have not yet deployed MFA. And even if they have, adversaries may inflict an MFA fatigue attack, in which a script attempts to log in with the stolen credentials over and over, generating a barrage of MFA push requests on

the account's owner's mobile device until they click the Approve button, either accidentally or simply to stop the notifications.

Persistence

Once an adversary has gained access to the target IT environment, they want to retain that access. Here are some of the top techniques they use.

Adversaries who have gained access to your network take steps to retain that access.

Account creation or manipulation

Adversaries often create new accounts to ensure that they can get back in to the network if a particular account they've compromised is suspended. In Microsoft 365 and other cloud environments, adversaries can choose to create accounts that have access to only certain services in order to reduce the risk of detection. Attackers can also modify existing accounts; resetting an account's password, for instance, helps prevent the account owner from changing it and locking the attacker out.

Email forwarding

Attackers can also seek to forward corporate email to an email address they control. While the simple Outlook forwarding rules that used to handle this job will be blocked by default in most tenants today, Microsoft Power Automate emails aren't blocked by default. To block them, you need set up a transport rule in Exchange Online.

Runbooks

Using Azure Automation, attackers can create a runbook with Python or PowerShell code that re-establishes their access. The runbook can be set to run on a schedule or be triggered through a web hook. In environments with a lot of runbooks and lax management, it's easy for an adversary to slip their runbook into the pile unnoticed. Proper cyber hygiene is the best way to spot these hidden runbooks promptly.

Leveraging Microsoft Office applications

Microsoft Office is common in modern organizations, and attackers take advantage of multiple mechanisms to re-establish their access when an Office application is started, including templates, macros and add-ins. For instance, they can create a malicious Outlook rule that triggers code execution when they send a specifically crafted email to the email account.

Discovery

Next, attackers seek to explore the Microsoft 365 environment to uncover vulnerabilities, learn about the services being used and of course find regulated and otherwise sensitive data to steal or encrypt for ransom. Here are some of the key intel they look for:

Exploring the environment uncovers vulnerabilities and valuable data.

- **Accounts** — This can include local system, domain, email and cloud accounts.
- **Browser bookmarks** — These can reveal information such as the user's interests and social media that can be used in targeting them. Bookmarks can also provide details about servers, tools and other network resources.
- **Cloud services** — These can include platform-as-a-service (PaaS), infrastructure-as-a-service (IaaS) and software-as-a-service (SaaS) platforms.
- **Email** — Hackers can search through email for sensitive information.

Popular tools that adversaries use for discovery include:

- **MailSniper** — A tool long used for searching through email in a Microsoft Exchange environment, MailSniper has been enhanced to work in Exchange Online as well.
- **ROADTools** — This is a framework for enumerating Azure Active Directory environments.

- **AAD Internals** — This PowerShell-based framework is designed to help legitimate admins and red teams administer and enumerate Azure Active Directory, but it can also be misused by adversaries to uncover and exploit vulnerabilities in Microsoft cloud environments.

Exfiltration

The final stage of the Microsoft 365 kill chain is exfiltration. This can be stealing data or holding it hostage for ransom. Techniques include:

- Automating the extraction process
- Transferring data in smaller batches to avoid detection
- Exfiltrating data via a physical medium, such as a USB drive or smartphone
- Exfiltrating data using an external web service that the organization already uses for legitimate purposes

Attack paths

Attack paths enable adversaries to gain control of Active Directory.

Now let's take a deeper dive into a common strategy that criminals use in successful cyberattacks: exploiting attack paths. An **attack path** is a series of steps that an adversary can execute to elevate their privileges from ordinary user to highly privileged user, even Domain Admin. We'll focus on attack paths in Active Directory, but this is highly relevant to Microsoft 365 for most organizations because their Azure Active Directory is backed by an on-premises Active Directory.

Throughout this discussion, it's important to keep in mind the goal of attackers: They want to get to an asset that enables them to control your Active Directory, since that gives them they control over the environment. These critical assets are called **Tier Zero** or the **control plane**, and they include:

- Built-in admin groups
- Privileged user accounts
- Domain controllers
- Group Policy objects (GPOs)

An attack path is a series of steps that an adversary can execute to elevate their privileges from ordinary user to highly privileged user.

Example of an Active Directory attack path

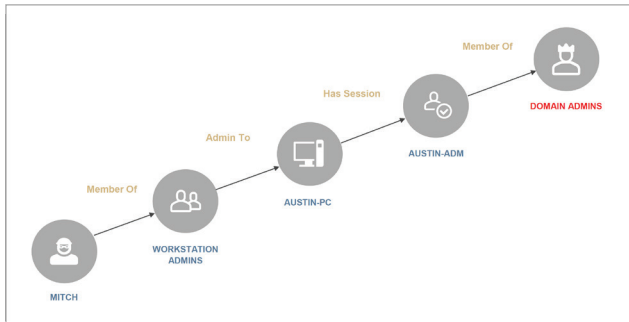


Figure 2. Example Active Directory attack path

Figure 2 illustrates an Active Directory attack path. Using the recon and initial access techniques detailed earlier, an adversary compromises an ordinary user account, Mitch. The Mitch account is a member of the Workstation Admins group, which has administrative permissions over various systems. One of those systems is Austin-PC, where there’s a session running under the service account Austin-ADM. In turn, Austin-ADM is a member of the highly privileged Domain Admins group. By exploiting this attack path, the adversary who has initially compromised just an ordinary user account can, in just a handful of steps, gain control of Active Directory.

Most environments have thousands — or even millions — of attack paths.

The reality is, in most environments, there are many attack paths. Moreover, there is an open-source tool called BloodHound that will lay them out for an adversary to use to plot their attack. For example, Figure 3 illustrates the attack paths in a test environment with just 40 or so objects, while Figure 4 shows the attack paths in a real enterprise environment.

As you can see, a live AD environment can easily contain thousands if not millions of attack paths. They are made possible by the complex nature of permissions, with interrelation between nested group permissions and the fact that objects can have permissions over other objects.




Figure 3. Even a small test environment has dozens of attack paths.

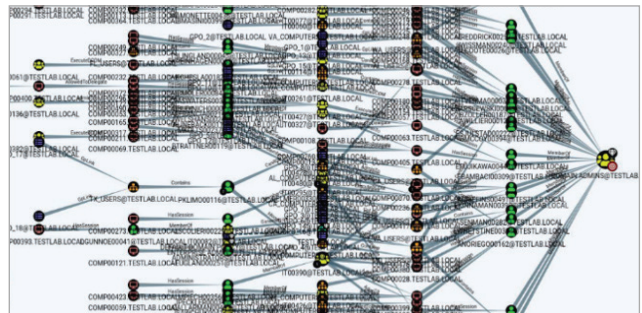


Figure 4. A real AD environment has thousands or millions of attack paths.

Mitigating attack paths: Choke points

So, how can organizations defend against attack paths? A useful analogy here is Google Maps. If I’m trying to get from point A to point B, Google Maps will highlight a few routes based on factors like the fastest time and the best roads. But if someone wanted to prevent me from completing your trip, would it be sufficient to shut down those particular routes? Not at all! The reality is, there are a huge number of ways for me to get from point A to point B. If one of them is blocked for some reason, I’m simply going to go another way.

Similarly, attackers have many options for getting from an ordinary user account to your Tier Zero assets. If you block a particular attack path, they will simply choose another route to their goal. Blocking all of the thousands or millions of attack paths in your environment is no more feasible than shutting down all the routes from San Francisco to Manhattan Island. We need a better strategy.

A choke point is the last step in the kill chain that is shared by many attack paths.

The key is to find the **choke point** — the last step in the kill chain that is shared by many attack paths. In our Google Maps analogy, choke points are the bridges and tunnels to the island of Manhattan. You don't need to block all the millions of routes I could take across the country; if you can block the seven bridges and tunnels at the end of all of those routes, you can keep me from getting to my goal.

Figure 5 illustrates the three choke points in an environment. The percentages are very straightforward: They indicate the percentage of your user accounts that could be exploited to get to the choke point.

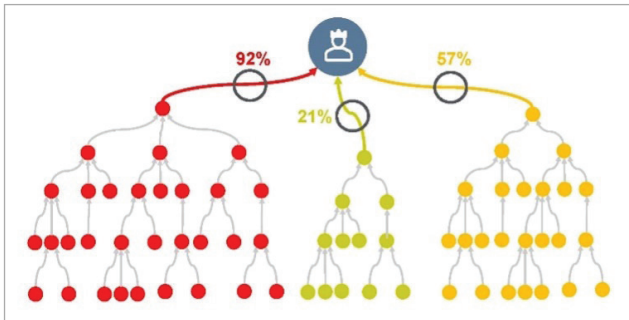


Figure 5. Thousands of attack paths share a final key step — the choke point.

This quantification enables you to focus your remediation efforts on your most critical choke points. By blocking a choke point, you block not just a single attack path, but all the attack paths that rely on it as their final step. Often, remediating a choke point is as simple as correcting one errant group membership or permission in AD.

Blocking even a single choke point can improve your security dramatically. In Figure 6, eliminating just one choke point blocks attack paths that would have enabled an attacker who compromised virtually any user account (92% of them) to reach your Tier Zero assets.

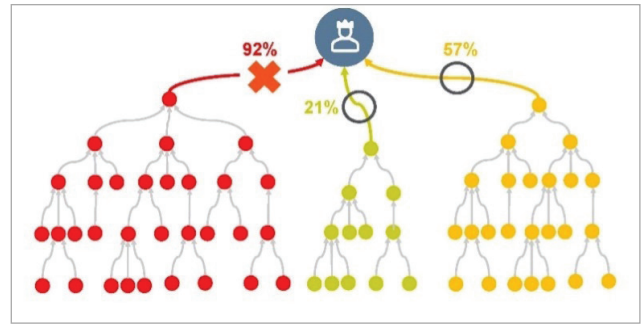


Figure 6. Mitigating a choke point blocks all the attack paths that rely on it.

Blocking even a single choke point can improve your security dramatically.

Attack path management and monitoring

Fortunately, there is a tool that identifies the choke points you need to mitigate to protect your organization. Quest has partnered with SpecterOps to bring it to market. [SpecterOps BloodHound Enterprise](#) maps out and quantifies the attack paths in your Active Directory, and clearly identifies the choke points so you can perform effective **attack path management**, prioritizing the elimination of the attack paths that involve the most risk.

However, unless your company is fairly new, your Active Directory is likely quite complex, with configurations and settings made by many different IT pros over the course of years. Therefore, actually making the changes required to mitigate the attack paths identified by BloodHound Enterprise can be risky, since it can be hard to accurately predict the effects. You can see that you could eliminate a whole swath of attack paths by correcting a permission that doesn't seem like it should be there— but maybe that permission was added 15 years ago to facilitate an application that's integrated with Active Directory and making the change could disrupt a vital business process.

Because there are often good reasons to be reluctant to immediately remediate a choke point, it's important to combine attack path management with attack path monitoring. **Attack path monitoring** helps you ensure that the attack paths you've identified are not exploited while you're doing the research required to ensure that correcting configurations or removing permissions won't have any adverse effects.

Quest Change Auditor and On Demand Audit integrate with BloodHound Enterprise to provide a comprehensive solution for attack path management and monitoring. Change Auditor and On Demand Audit constantly watch for indicators of compromise, including Golden Tickets, DCSync attacks, AD database exfiltration and the techniques described earlier for the various phases of the Microsoft 365 kill chain.

Attack path monitoring helps ensure that attack paths are not exploited before you can remediate their choke points.

Conclusion

Understanding the Microsoft 365 kill chain is vital to building an effective security strategy that ensure cyber resilience. Attack path management and attack path monitoring are extremely powerful parts of any defense-in-depth strategy, enabling you to pinpoint and mitigate key weaknesses to dramatically reduce your attack surface area.

About Quest

Quest creates software solutions that make the benefits of new technology real in an increasingly complex IT landscape. From database and systems management, to Active Directory and Microsoft 365 migration and management, and cybersecurity resilience, Quest helps customers solve their next IT challenge now. Around the globe, more than 130,000 companies and 95% of the Fortune 500 count on Quest to deliver proactive management and monitoring for the next enterprise initiative, find the next solution for complex Microsoft challenges and stay ahead of the next threat. Quest Software. Where next meets now. For more information, visit www.quest.com.

© 2022 Quest Software Inc. ALL RIGHTS RESERVED.

This guide contains proprietary information protected by copyright. The software described in this guide is furnished under a software license or nondisclosure agreement. This software may be used or copied only in accordance with the terms of the applicable agreement. No part of this guide may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying and recording for any purpose other than the purchaser's personal use without the written permission of Quest Software Inc.

The information in this document is provided in connection with Quest Software products. No license, express or implied, by estoppel or otherwise, to any intellectual property right is granted by this document or in connection with the sale of Quest Software products. EXCEPT AS SET FORTH IN THE TERMS AND CONDITIONS AS SPECIFIED IN THE LICENSE AGREEMENT FOR THIS PRODUCT, QUEST SOFTWARE ASSUMES NO LIABILITY WHATSOEVER AND DISCLAIMS ANY EXPRESS, IMPLIED OR STATUTORY WARRANTY RELATING TO ITS PRODUCTS INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTY OF MERCHANTABILITY, FITNESS FOR A PARTICULAR

PURPOSE, OR NON-INFRINGEMENT. IN NO EVENT SHALL QUEST SOFTWARE BE LIABLE FOR ANY DIRECT, INDIRECT, CONSEQUENTIAL, PUNITIVE, SPECIAL OR INCIDENTAL DAMAGES (INCLUDING, WITHOUT LIMITATION, DAMAGES FOR LOSS OF PROFITS, BUSINESS INTERRUPTION OR LOSS OF INFORMATION) ARISING OUT OF THE USE OR INABILITY TO USE THIS DOCUMENT, EVEN IF QUEST SOFTWARE HAS BEEN ADVISED OF THE POSSIBILITY OF SUCH DAMAGES. Quest Software makes no representations or warranties with respect to the accuracy or completeness of the contents of this document and reserves the right to make changes to specifications and product descriptions at any time without notice. Quest Software does not make any commitment to update the information contained in this document.

Patents

Quest Software is proud of our advanced technology. Patents and pending patents may apply to this product. For the most current information about applicable patents for this product, please visit our website at www.quest.com/legal

Trademarks

Quest and the Quest logo are trademarks and registered trademarks of Quest Software Inc. For a complete list of Quest marks, visit www.quest.com/legal/trademark-information.aspx. All other trademarks are property of their respective owners.

If you have any questions regarding your potential use of this material, contact:

Quest Software Inc.

Attn: LEGAL Dept
4 Polaris Way
Aliso Viejo, CA 92656

Refer to our website (www.quest.com) for regional and international office information.