CYBER SECURITY

FRONT /> CODE

MODERN CYBER THREATS & ATTACKS

June 2025

AI-POWERED VOICE PHISHING (VISHING)

BUSINESS EMAIL COMPROMISE (BEC) WITH DEEPFAKES

LIVING-OFF-THE-LAND (LOTL) ATTACKS

ACTIVE DIRECTORY ATTACK PATHS

INSIDE THE CYBERCRIME MARKETPLACE

HOW HACKING AS A SERVICE IS FUELING THE THREAT ECONOMY

Cybersecurity 2025

June 2025/ Volume 01

1 AI-Powered Voice Phishing (Vishing)

AI-powered voice phishing (vishing) has become a major cybersecurity threat, using advanced voice cloning to impersonate trusted figures like executives or colleagues with alarming accuracy. This article gives a full overview of AI-powered vishing, including well-known examples, the attack chain, and practical ways to detect and prevent it from happening, from technological, procedural, and legal points of view, to equip organizations and individuals with the tools to protect themselves.

The Rise of AI-Powered Vishing

Vishing is a social-engineering scam conducted over the phone, generally relying on psychological manipulation. The integration of AI voice cloning has made these attacks far more sophisticated. Threat researchers note that just a few minutes of a target's recorded speech from public speeches, podcasts, or voicemails can train AI models to replicate their voice with uncanny precision (cloud.goooogle.com). Others utilize these cloned voices to pretend to be CEOs or coworkers, deceiving others into doing things like sending money or giving up their passwords.
Several incidents have highlighted the severity of this threat. Earlier this year, a merchant in Hong Kong lost HK$145 million owing to AI voice cloning via WhatsApp voice chats, where scammers posed as a trusted contact that the victim was meant to purchase cryptocurrency equipment from (South China Morning Post). The FBI also issued a warning in May 2025 about hostile actors deploying AI-generated voice communications to impersonate important U.S. officials, generally paired with SMS lures ("smishing") to steal credentials or funds (IC3 In a different example in early 2024, a Hong Kong-based multinational corporation lost HK$200 million (about US$25 million) after an employee was fooled during a deepfake video chat involving AI-generated copies of the CFO and other coworkers (Ars Technica). While this involved video, it goes to demonstrate the prospect of AI-driven imitation. According to the 2025 CrowdStrike Global Threat Report, voice-spoofing attacks surged 442% from the first to the second half of 2024, validating the allegations of the rapid expansion and risk of this form of threat (Security Magazine).

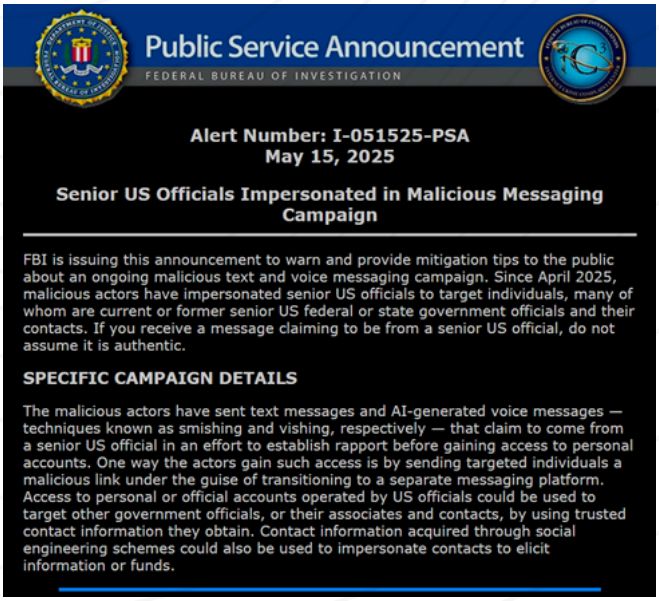

1 Public Service Announcement as shown on the FBI's website

How Attackers Execute AI-Powered Vishing

Now that we understand the severity of AI-powered vishing, it is very important to understand how these attacks work. The attack chain for AI-powered vishing is methodical and relies on accessible technology:

- Voice Sample Collection: Attackers get a target's voice sample from public sources including YouTube videos, social media posts, corporate webinars, or even voicemails left during reconnaissance calls.
- Voice Cloning: Using advanced speech-synthesis software, such as ElevenLabs (ElevenLabs) or Resemble AI (Resemble AI), attackers train AI models to copy the target's voice, replicating their accent, tone, and cadence.
- Execution of the Scam: The cloned voice is employed to initiate contacts, typically masquerading as an executive urgently asking tasks like money transfers, password resets, or sensitive data releases. The call may sound real, although small artifacts like strange pauses or artificial inflections may be present.

These attacks exploit trust and haste, making them effective even against diligent staff. While firms like ElevenLabs and Resemble AI are legitimate, their accessibility raises worries about potential exploitation, albeit many providers adopt steps to prevent abuse.

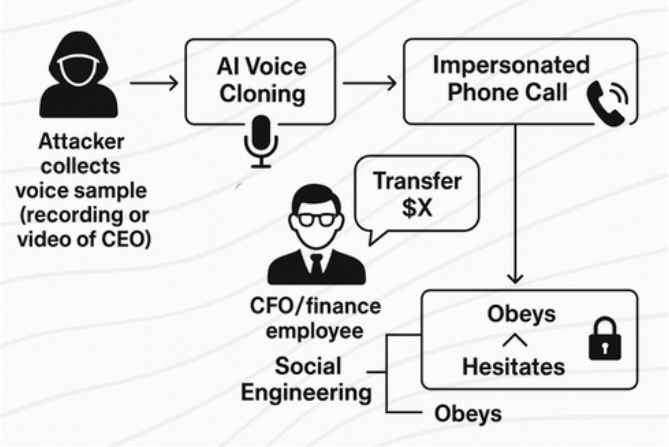

2 A flowchart of an AI-powered vishing attack

Detection Challenges and Deepfake Telltales

AI-generated voices are increasingly difficult to discern from real ones, especially for unskilled listeners. However, tiny signs may suggest a deepfake:

- **Absent Natural Breathing**: AI voices may lack the natural breathing sounds humans create while talking.
- **Robotic Timbre**: Some synthetic voices sound a little too mechanical or too smooth.
- **Inconsistent Background Noise**: AI-generated audio might include background sounds that are always the same or that play over and over again, unlike the dynamic noise in real environments (Lifehacker).

As AI technology gets better, though, it's getting harder to recognize when a person is talking, even with these clear signs. This is because AI tools are getting better at imitating how people naturally speak (Podcastle). This highlights the importance for technical and procedural defenses. Even with AI's faults, it can be hard for non-experts to recognize. The key difference might be the voice being slightly "off", or the phrase seeming strange because it is not the real person speaking (cloud.google.com).

Defense Strategies Against AI-Powered Vishing

Organizations and individuals must adopt a multi-layered approach to combat AI-powered vishing, combining technology, employee training, and robust verification processes. Below are key strategies, summarized for reference:

1. Out-of-Band Verification

·Treat voice calls as untrusted channels for sensitive requests. Implement pre-agreed verification methods, such as:

·Code-Words or PINs: Establish unique code-words or PINs offline between executives and their teams for sensitive actions like fund transfers.

Call-Back Protocols: If an executive calls with an urgent request, employees should hang up and call back using a number from the corporate directory that has been verified.

2. Voice Biometrics

Advanced voice-biometrics systems use pitch, cadence, and accent to tell who a speaker is by measuring their unique vocal traits. Solutions like ID R&D's IDLive Voice can detect synthetic speech and flag anomalies, offering robust protection against voice cloning (ID R&D). Integrating voice biometrics with multi-factor authentication (MFA) strengthens security further.

3. AI-Based Fraud Detection

AI-driven fraud detection systems learn normal call patterns and flag outliers. For example, Proofpoint's solutions can detect anomalies like a 3 a.m. call in an unusual language mimicking a CEO's voice, triggering alerts for further verification (Proofpoint).

4. Telephony Standards

Standards like SHAKEN/STIR (Secure Telephone Identity Revisited/Signature-based Handling of Asserted Information Using Tokens) provide caller ID attestation, helping verify the authenticity of calling numbers. While not foolproof, these standards reduce the risk of spoofed calls.

5. User Training

Train employees to recognize potential deepfake telltales and report suspicious calls immediately. Training should emphasize:

- Listening for unnatural speech patterns, such as absent breaths or repetitive filler words.
- Awareness that human detection is limited, encouraging reliance on verification protocols (Lifehacker).

6. AI Detection Tools

AI-based audio classifiers can distinguish synthetic voices from real ones. Tools like PlayHT's AI Voice Classifier (PlayHT) analyze audio for indicators of tampering, delivering an additional layer of defense. These tools are emerging and should be included into telephone systems where practical.

Legal and Regulatory Landscape

Responses to legal issues surrounding the misuse of AI-generated voices are beginning to emerge, but much work remains to be done. The U.S. Federal Communications Commission (FCC) has prohibited the use of AI-generated voices in robocalls and enforced consequences under the Telephone Consumer Protection Act (TCPA) as of February 2024 (FCC). Tennessee's Ensuring Likeness Voice and Image Security (ELVIS) Act, effective in 2024, protects individuals' voices from unauthorized use, setting a precedent for state-level protections (Holland & Knight). At the federal level, the No AI FRAUD Act is under consideration to address broader AI-driven impersonation, including voice cloning (Lexology).

Organizations should stay informed about these regulations to ensure compliance and advocate for stronger protections.

AI-powered vishing poses a serious and growing threat, with attackers adopting voice cloning to execute convincing scams. By understanding the attack chain, recognizing detection challenges, and implementing a multi-layered defense strategy, combining out-of-band verification, voice biometrics, AI detection tools, and employee training; organizations can significantly reduce their risk. By staying vigilant and adopting these strategies, businesses and individuals can protect themselves against the evolving menace of AI-powered vishing.

2 Business Email Compromise (BEC) with Domain Spoofing

Your new CEO/CFO is a hacker, sounds crazy, right? Relax, that isn't exactly the case, but this scenario is very possible; it is very common for hackers to impersonate CFOs and other higher ranking executives. In these attacks, cybercriminals send emails that look like they are coming from your boss or vendor, often using a nearly identical domain (e.g. "@RealCo.com" vs "@RealCo.com," where the "l" is actually a "I"), this is a form of typosquating. Note: Typosquatting involves registering domain names that are similar to well-known domains, often differing by a single character or using homoglyphs, basically characters that look similar but are different (proofpoint.com). If they are really skilled, they might manage to also hijack the real domain (an "email account compromise") to launch their scheme undetected (proofpoint.com). Once they appear legitimate, they ask staff to wire money or share sensitive info. These tricks can fool even savvy employees, so it pays to spot small giveaways.
BEC attacks are alarmingly common, accounting for 73% of all reported cyber incidents in 2024 (Hoxhunt). This statistics supports the idea that BECs as a cyber-attack is not spoken about enough. This article aims to discuss ways of mitigating against this threat.

IMAGE ILLUSTRATING THE PROCESS OF A BUSINESS EMAIL COMPROMISE (BEC) HIGHLIGHTING KEY STAGES AND CHARACTERISTICS

Preventing BEC with domain spoofing means layering email defenses and human checks. Technical steps include strong email authentication: publish and enforce SPF, DKIM and a strict DMARC policy on your domains (cloudflare.com). For those unfamiliar, SPF (Sender Policy Framework) specifies which mail servers are allowed to send emails for your domain, DKIM (DomainKeys Identified Mail) allows you to take responsibility for a message that can be verified by the recipient, and DMARC (Domain-based Message Authentication, Reporting, and Conformance) builds on SPF and DKIM to protect your domain from being spoofed.
When DMARC is properly set to "reject," spoofed messages simply bounce. Also use intelligent email filters (many use AI or behavioral analysis) to flag messages that request unusual actions (proofpoint.com). Don't forget basic hygiene: disable old POP/IMAP mail protocols (which can bypass modern filters) and require MFA on all executive accounts (cloudflare.com; proofpoint.com.)

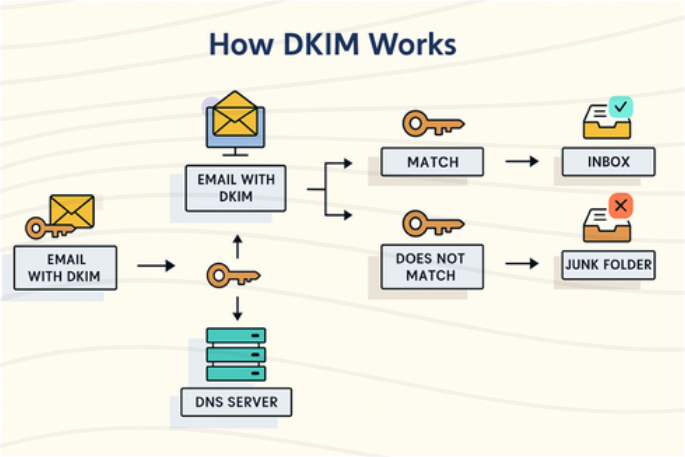

VISUAL GUIDE TO HOW DKIM AUTHENTICATES EMAILS AND PROTECTS INBOXES

People-side defenses are critical too. Train staff to question urgent money requests and odd instructions. For example, a CFO is unlikely to demand employee tax data or password resets via email (proofpoint.com). Encourage employees to double-check domain spellings ("yourcompany.com" vs. "yourc0mpany.com") Be wary of homoglyphs, where characters that look similar are used to mimic legitimate domains, such as using 'rn' instead of 'm', and never act on a payment change without a callback to a known phone number (proofpoint.com; cloudflare.com). Establish formal approval steps for any bank transfers; e.g. requiring two people to sign off. Over time, a mix of technology and awareness will turn those red flags into routine habits.

Attackers use look-alike domains (typosquatting, punycode, subdomains) or entirely fake (webmail) addresses to imitate your CEO/CFO (proofpoint.com; intelligence.abnormal.ai.)

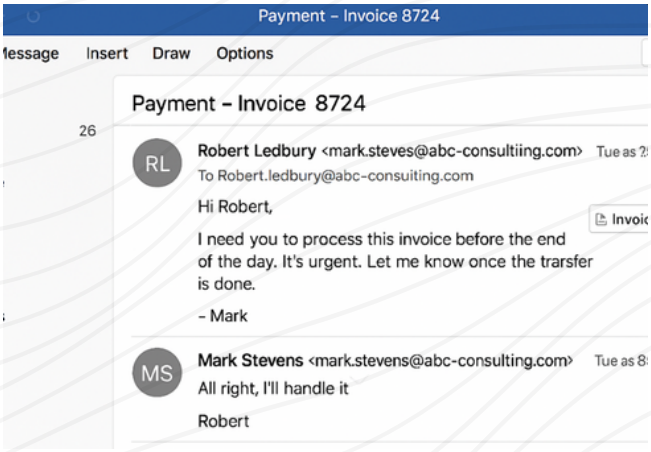

A CONVINCING SPOOF OF AN EXECUTIVE EMAIL ADDRESS LURES THE FINANCE TEAM INTO A FRAUDULENT WIRE TRANSFER REQUEST. NOTE THE SUBTLE DOMAIN MISSPELLING AND URGENT TONE TYPICAL OF BEC ATTACKS.

Signs to spot: Check the email header and "Reply-To." In the email header, look for the 'From' address and ensure it matches the expected domain. Also, check if the 'Reply-To' address is different from the 'From' address, as attackers might use this to redirect responses. If the address is off by even one letter or the tone feels rushed, be suspicious (proofpoint.com). Also watch for messages that demand secrecy or bypass normal invoicing channels, genuine exec requests are usually documented in your finance system (proofpoint.com).

Prevention checklist:

- **Email auth:** Publish SPF/DKIM records and enforce DMARC "reject" so spoofed mail is blocked (cloudflare.com; intelligence.abnormal.ai).
- MFA and filtering: Turn on multifactor authentication for executive mailboxes and use advanced BEC filters or AI tools to spot context anomalies, these tools can detect subtle changes in email behavior that might indicate a compromise, such as unusual language or requests that don't match the sender's typical patterns. (proofpoint.com).
- Process controls: Always verify any change in vendor/payment info through a secondary channel (phone call, in-person sign-off) (proofpoint.com).
- Training: Regularly quiz staff with BEC simulations and share real examples. Make sure to include the latest BEC tactics in training, such as AI-generated emails and multi-channel attacks that combine email with phone calls or texts.Emphasize the "trust but verify" rule; it's OK to ask "Would my boss really email this?".
- Monitor for unusual email activity: Regularly review email logs for signs of compromise, such as unexpected login locations or times.
- Stay informed: Keep up with the latest BEC trends and adjust defenses accordingly, as attackers continually evolve their methods.

Scammers aren't sending emails from Nigerian princes anymore; they're pretending to be your CEO. Spotting the difference can save millions. With strong email rules, cautious staff, and just a bit of healthy skepticism, your company can shut these fakes down fast. If something smells phishy, it probably is; double-check before you double-pay.

3 Living-Off-the-Land (LOTL) Attacks

Living-Off-The-Land (LOTL) attacks use nothing more than the tools already present on target systems to do the dirty work. Instead of dropping new malware, attackers hijack built-in utilities (like PowerShell, WMI, Bitsadmin, certutil, etc.) to run commands, move laterally, and extract data. This "fileless" approach can evade traditional defenses; "unlike traditional malware… LOTL attacks are fileless," writes CrowdStrike, meaning adversaries execute everything in memory or through signed system tools (crowdstrike.com). This is simply because the tools they use are trusted by administrators and often whitelisted, these intrusions can slip past antivirus and signature-based alerts. As one vendor warns, "if you can hijack an existing and trusted piece of software … the chances are better that you will go undetected" (sentinelone.com).

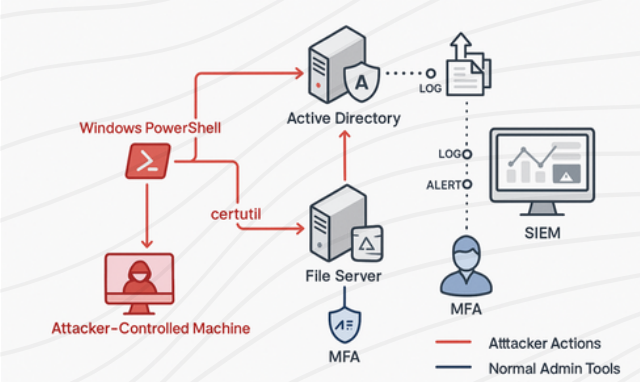

1 LIVING OFF THE LAND DEPICTED

Tools & Tactics

A LOTL attack might use PowerShell or cmd.exe scripts to probe Active Directory, net commands to enumerate users, WMI or Remote Desktop (RDP) to jump between machines, and even legitimate credential-dumpers like Mimikatz (run in memory). Common Windows utilities abused include msiexec (Windows Installer), psexec (remote execution), certutil (certificate utility), regsvr32/rundll32, and even Office programs with malicious macros. For example, attackers often repurpose certutil to decode Base64 payloads; Volt Typhoon a Chinese State-sponsored cyber gang had its operators used certutil to turn encoded strings into executable malware (attack.mitre.org). Other frequent LOLBins (LOw-level-binaries) include scripting hosts (PowerShell, bash, wscript/cscript), data-transfer tools (bitsadmin, robocopy, ftp), and network scanners (netstat, ping). Every organization's environment differs, hence, attackers may even identify obscure executables unique to a target, but they usually begin with ubiquitous ones like PowerShell or WMI (sentinelone.com).

·No new files: LOTL attacks leave few artifacts on disk. Payloads live in RAM or use alternate data streams (hidden NTFS streams) instead of plain files.
·Evasion: By using system tools, attackers bypass many controls. For example, running PowerShell commands or WMI scripts does not trigger signature alerts since those binaries are legitimate.
·Dual-use tools: Utilities intended for management (e.g. net group to list Domain Admins) can be turned malicious. Volt Typhoon was observed using net user and net group /dom exactly like a sysadmin would (attack.mitre.org). Because admins often run these commands, distinguishing benign from malicious use is tricky.

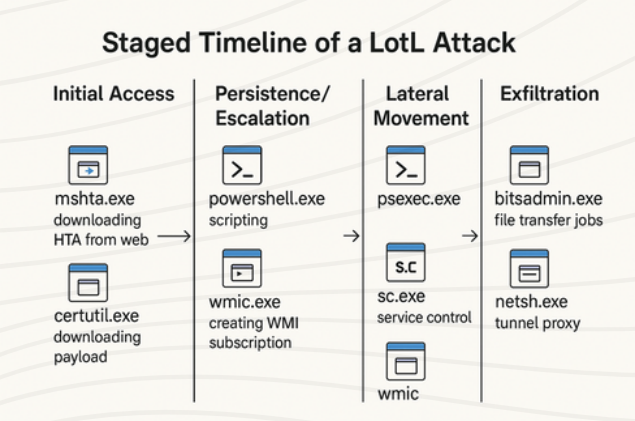

1 LIVING OFF THE LAND DEPICTED

High-Profile Incidents

LOTL techniques are popular with both cybercriminals and nation-state groups. For example, LockBit ransomware gangs routinely move "living off the land." CISA reports that LockBit affiliates use PowerShell and batch scripts in most intrusions, mainly for discovery, reconnaissance, credential hunting and privilege escalation (cisa.gov). The Australian Cyber Security Centre notes LockBit 3.0 actors specifically rely on built-in PowerShell commands after initial access to execute malicious actions (cyber.gov.au). In practice, LockBit teams often combine these with other admin tools: they deploy PsExec, WMI, and remote-management software (AnyDesk, Splashtop) to spread across networks (cyber.gov.au). After stealing credentials (often via Mimikatz run in-memory), they disable defenses and exfiltrate data then encrypt it with their custom ransomware.

Similarly, Volt Typhoon has made LOTL their signature tactic. Active since 2021 against U.S. critical infrastructure, Volt Typhoon emphasizes stealth. The MITRE ATT&CK group page notes Volt Typhoon "has emphasized stealth in operations using web shells [and] living-off-the-land (LOTL) binaries" (attack.mitre.org). In one campaign, CISA observed Volt operators carefully query Windows event logs with PowerShell (targeting specific users and time windows) and dump them into .dat files (cisa.gov). They have also used net user and net group commands to map accounts and privileges (attack.mitre.org), and even employed vssadmin to snapshot the Active Directory database. In all these actions, no unfamiliar executable ran, only trusted Windows tools. For example, Volt Typhoon once archived the NTDS.dit file (the AD database) using 7-Zip, and likewise has used certutil to decode payloads (attack.mitre.org). These examples show how LOTL lets a threat actor operate as "an administrator would", blending into normal activity. In fact, analysts found Volt Typhoon deliberately stayed within business hours and mimicked normal user behavior to avoid detection (cisa.gov).

These incidents, LockBit, Volt Typhoon, and many others highlight a pattern: LOTL attacks skirt outside networks for initial access, then proceed almost entirely with internal tools. Even some red-team and pentest tools are designed this way. For instance, the Cobalt Strike framework (used by many professional security testers) runs in memory and is often used to simulate advanced attacks (hhs.gov). Other common red-team tools like PowerShell Empire or Metasploit can similarly drop payloads via trusted apps.

Detection Strategies

Since LOTL attacks use legitimate software, detecting them requires strong logging and analysis of behavior, not just static signatures. Agencies and vendors agree that defenders should collect detailed logs (PowerShell transcripts, WMI activity, command histories, etc.) and feed them into a SIEM or EDR for correlation (cisa.gov). For example, Microsoft Sysmon can log original filenames and full command lines, making it possible to spot when a trusted executable is doing something unexpected (cisa.gov). CISA recommends enabling verbose logging of security events, shell usage and script execution across all machines, then storing those logs in a centralized, tamper-resistant system (cisa.gov). With comprehensive logging in place, analysts can look out for the following anomalies:

- **Anomalous process chains:** Look for unusual parent-child relationships. For example, a Word or Outlook process spawning powershell.exe or cmd.exe should raise an alert. CISA specifically points out that monitoring for Office apps launching script hosts can uncover fileless loaders (cisa.gov).
- **Suspicious command-lines:** Use your SIEM to flag rare or obfuscated commands. Commands using alternate data streams (like type file.txt > file.txt:hidden.exe) or environment-variable tricks were noted by CISA as indicators (cisa.gov). Likewise, patterns like a non-admin user running net group or net user are highly unusual.
- **Behavioral indicators:** Instead of fixed IOCs, focus on indicators-of-attack. This means tracking sequences like login → run PowerShell → create WMI entry, regardless of the exact file. CrowdStrike emphasizes that IOAs (e.g. "credential dumping followed by lateral RPC calls") catch fileless attacks because they spot the action rather than the dropped payload (crowdstrike.com). In practice, a UEBA or threat-hunting system can correlate a spike in cmd.exe executions, new service creations, or hidden scheduled tasks to expose an ongoing intrusion.

In short, LOTL detection is about context. CISA notes it "requires… contextual analyses of multiple data sources to identify command executions, file interactions, privilege escalations, and other network activities that differ from normal administrative actions." (cisa.gov.) Keeping baselines of normal behavior is essential. For instance, Volt Typhoon avoided unusual hours to blend in, so defenders who know typical login times can spot when someone deviates. Monitoring can extend beyond the endpoint: look at network proxies for exotic traffic patterns, and audit logins/RDP from odd IPs. Many SOCs use dedicated threat hunts (searching for LOLBin usage patterns) and tuned EDR rules to catch LOTL misuse. As one analysis advises, treat a suspicious PowerShell invocation by a non-admin or on a critical server as a red flag to investigate.

Prevention Tips

Stopping LOTL attacks upfront means constraining the tools attackers can use and limiting their privileges:

- **Application whitelisting:** Use AppLocker or Windows Defender Application Control to allow only approved executables and scripts (cisa.gov.) By enforcing strict allowlists, you can block even legitimate utilities if they run outside policy (e.g. prevent powershell.exe from running from user directories or by low-privileged accounts). On macOS, similar controls (like Gatekeeper) can block unknown binaries (cisa.gov).
- **Least privilege:** Operate services and user sessions with minimal rights. If users are not local admins, even a LOLBin they run will be limited. Remove unnecessary accounts from privileged groups (CISA also warns to remove unneeded Enterprise Admin accounts). Consider just-in-time administration: only grant admin privileges when needed, and revoke them after the task.
- **Multi-factor authentication:** Enforce phishing-resistant MFA for all logins, especially remote access (VPN, RDP). Attackers often pivot using stolen credentials, so MFA stops many LOTL campaigns before they begin.
- **Harden scripting hosts:** For example, apply PowerShell execution policies or Constrained Language Mode so that only signed scripts run. Disable or audit tools that are not needed: if no one needs bitsadmin or wmic.exe, restrict or remove them.
- **Network segmentation:** Limit the "blast radius" of stolen creds. Well-segmented networks slow lateral movement. CISA notes that anomalous traffic between segments can signal a stealthy attacker (cisa.gov). Employ firewalls, VLANs, or zero-trust micro-segmentation so that even if an attacker uses LOTL techniques, they cannot freely reach every server.
- **Alert on anomalous behavior:** Configure alerts for unusual activity such as scripts running at odd hours or by unexpected users. For example, an administrator suddenly connecting via RDP from a foreign location, or a server launching rarely-used utilities, should trigger review. Many security tools can be tuned to watch for known "bad" parameter combinations or for new services being registered (e.g. a new scheduled task triggered by schtasks.exe).

These measures would not block every sophisticated adversary, but they raise the bar. CISA's joint guidance strongly recommends combining these controls: "Implement as many [mitigations] as possible…to enable effective data correlation and analysis" (cisa.gov).

Distinguishing Malicious vs. Legitimate Use

A core challenge in combating LOTL attacks is that the same actions can look perfectly normal. Administrators routinely run PowerShell, query account lists, or make registry changes, all of which attackers do too. This overlap means straightforward signatures yield false positives (and true attacks can slip through). The defense strategy, therefore, centers on context and anomalies. Track who ran a tool, when, and why. If a helpdesk account suddenly spawns a PowerShell process on a domain controller, or if sensitive data is zipped and uploaded via powershell.exe at 3 AM, those are out-of-norm events that merit scrutiny. CISA specifically notes detection involves spotting activities that "differ from normal administrative actions."(cisa.gov). In practice, that means keeping a tight baseline of normal ops and alerting on deviations. As one Volt Typhoon example showed, even timing can betray the intruders: when defenders see logins and commands strictly within usual office hours, or coming from the same IP ranges as legitimate users, it might actually be a clue that "noise" has been minimized by the adversary (cisa.gov).

Ultimately, separating a skilled attacker from a busy sysadmin is a hard game of nuance. Organizations that assume any legitimate tool usage is benign will be blind to LOTL intrusions; conversely, treating every admin task as suspect is unmanageable. The best path is a combination of proactive measures (logging, allowlisting, least privilege) and intelligent monitoring: watch for triggers like unusual process spawning, or a normally quiet server suddenly running heavy scripting. When alerts do occur, verify the user's context: was this part of a scheduled maintenance, or did it follow some other compromise indicator (like an odd VPN login)? Over time, security teams can learn the normal rhythm of their IT environment so that when LOTL techniques appear, they stand out against the baseline.

4 Active Directory Attack Paths

Active Directory Attack Paths
Active Directory (AD) is a Microsoft tool that helps organizations control who can access network resources and what they can do with them. It keeps track of people, computers, and other devices and offers authentication and authorization services.
Active Directory (AD) forests are goldmines for attackers. Common tricks include Pass-the-Hash, Kerberoasting, and Silver/Golden Ticket attacks, all exploiting how Windows authentication works. In a Pass-the- Hash (PtH) attack, the bad guy steals a user's hashed password (often from memory) and reuses it to log in as that user, without knowing the actual password. A hashed password is a one-way encrypted version of the password. In Windows, when a user logs in, the system uses the hashed password for authentication without needing the plain text password. Attackers can steal this hash and use it to authenticate as the user on other systems (nccgroup.com). Kerberoasting happens when an attacker requests Kerberos service tickets (TGS) for service accounts (which any user can do) and then cracks those tickets offline to recover the service account password. Kerberos is a network authentication protocol used by Windows domains to provide secure communication. Service accounts are used by applications and services to authenticate to the domain. Attackers target these because they often have high privileges and their passwords might not be changed frequently, making them easier to crack (picussecurity.com). Silver Tickets are forged service tickets based on a compromised service account hash, and Golden Tickets are forged Ticket Granting Tickets (TGTs) created by abusing the KRBTGT account's hash (semperis.com; learn.microsoft.com). A Silver Ticket allows an attacker to impersonate a user for a specific service without needing the actual password. A Golden Ticket, on the other hand, is a TGT that can be used to authenticate as any user in the domain, providing almost unlimited access. A Golden Ticket is especially dangerous: it can grant domain-wide access until the KRBTGT password is rotated.

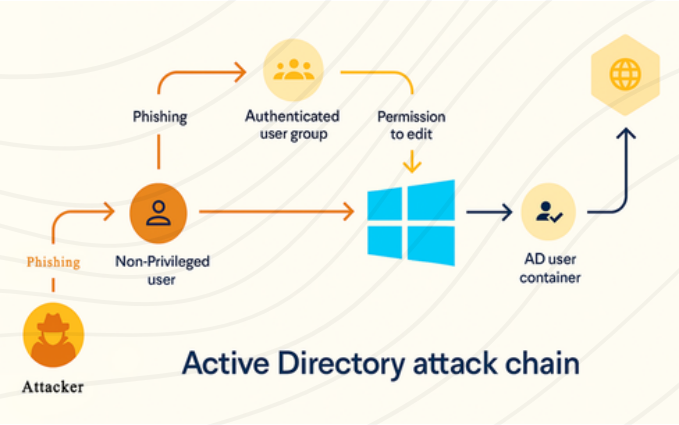

ACTIVE DIRECTORY ATTACK CHAIN

You can often spot these attacks by analyzing AD logs and behavior:
Pass-the-Hash detection: Watch for unusual NTLM logins, NTLM (NT LAN Manager) is an authentication protocol used in Windows. It's less secure than Kerberos and can be vulnerable to Pass-the-Hash attacks because it sends the hashed password over the network. For example, Windows event 4624 with Logon Type 3 (network logon) using NTLM typically appears without a prior password-based logon (nccgroup.com). If a server suddenly sees logins via NTLM that weren't preceded by a normal interactive logon, that's a clue. Also look for spikes in logons from a single account across many machines. If one user suddenly logs into dozens of hosts (far more than normal), an attacker may be using a stolen hash to move laterally (nccgroup.com). Consider using SIEM tools or AD monitoring solutions that can alert on unusual NTLM logon patterns, such as multiple logons from the same account across different machines in a short period.

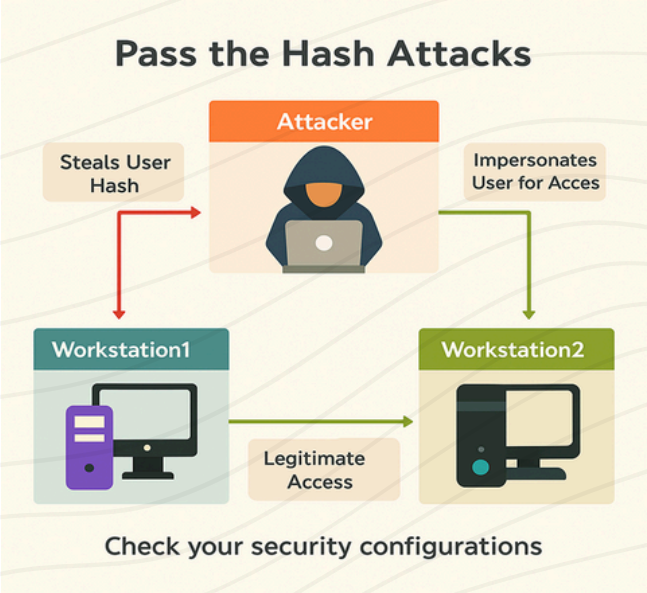

PASS-THE-HASH IN ACTION

- Kerberoasting detection: Enable Kerberos logging on domain controllers (audit Kerberos Service Ticket ops). Monitor event ID 4769 (service ticket requested). Ensure that Kerberos logging is enabled on domain controllers to capture event ID 4769. Use log analysis tools or SIEM systems to identify patterns indicative of Kerberoasting, such as a high volume of service ticket requests from a single user or the use of outdated encryption types like RC4. Any unusual patterns here are suspicious. For instance, if one user requests a large number of service tickets for different servers in a short time, it may indicate a Kerberoast attempt (picussecurity.com). Also, most modern environments use AES encryption for Kerberos tickets. Tools like Hashcat target old RC4-HMAC tickets. If you see tickets requested with RC4 (Encryption Type 0x17), that's a strong sign of Kerberoasting activity (picussecurity.com).

- Silver/Golden Ticket detection: These are stealthy. A red flag for a Silver Ticket is when a service accepts a Kerberos ticket that never appears to have been issued by the Key Distribution Center (KDC). The KDC is a service in Active Directory that issues Kerberos tickets. The KRBTGT account is used to encrypt and decrypt TGTs. Resetting its password invalidates all existing TGTs, including any forged Golden Tickets. In other words, Event 4769 on the service's machine with no corresponding 4768 (TGT request) just prior could signal a Silver Ticket (semperis.com).

·Golden Tickets can cause broader anomalies: infinite access to resources, or TGTs that don't expire on schedule. It's smart to log and review unusual ticket lifetimes or authentication activity by the krbtgt account. After a suspected Golden Ticket attack, you must reset the KRBTGT account's password twice to wipe out the attacker's forged keys (learn.microsoft.com).

Hardening Best Practices

Hardening your AD environment helps block these paths: adopt strict privilege separation (use dedicated admin accounts and don't log in with high-privilege creds on day-to-day devices). The Tiered Access model organizes accounts and systems into different privilege levels: Tier 0 for domain controllers, Tier 1 for servers, and Tier 2 for user workstations. This segmentation helps prevent lateral movement by limiting the privileges and access of compromised accounts. Limit who has Domain Admin or Enterprise Admin rights, and use the Tiered Access model (Tier 0 = domain controllers, Tier 1 = servers, Tier 2 = user workstations). Enable multi-factor authentication for sensitive accounts and consider tools like Microsoft LAPS to rotate local admin passwords. Implement multi-factor authentication (MFA) for all privileged accounts, including Domain Admins and Enterprise Admins, to add an extra layer of security against password-based attacks. Network segmentation is helpful too: don't allow any system to reach the domain controller unless needed. Ensure that domain controllers are placed in a separate network segment with strict access controls, allowing only necessary traffic to and from these critical systems. Most importantly, regularly change the KRBTGT password, Microsoft recommends doing it twice to truly erase any Golden Ticket an attacker may have crafted. Microsoft recommends rotating the KRBTGT password every 6-12 months, performing the rotation twice in succession to ensure that any Golden Tickets are invalidated. (learn.microsoft.com).

In addition to the above steps:

- Enforce the principle of least privilege: give users and services only the rights they absolutely need (nccgroup.com). Use separate accounts for admins ·and regular use, and lock down service account permissions.

- Use strong, long passwords (or better, managed gMSAs) for service accounts. Where possible, use group Managed Service Accounts (gMSAs) for service accounts. gMSAs provide automatic password rotation and are more secure than traditional service accounts, reducing the risk of Kerberoasting and other password-based attacks. Kerberoasting relies on weak service passwords – aim for random 30+ character secrets and change them frequently (picussecurity.com). Ideally use MSAs/gMSAs that auto-rotate, so attackers don't have a long window to crack a hash.

Each of these steps, from watching for odd login events to enforcing credential hygiene, narrows the attack surface. With diligence and the right controls, you can detect Active Directory attacks early and shut them down before damage spreads.

Think of Active Directory like the keys to your entire building, hand them out carelessly, and someone's bound to sneak in after hours. Attackers love lazy ticket hygiene and over-privileged accounts. Break the habit. Keep privileges tight, rotate those tickets, and monitor like your domain depends on it, because it does. Active Directory security is an ongoing process. Regularly review and update security policies, monitor for new threats, and ensure that all systems are patched and up-to-date to maintain a strong defense against evolving attack techniques.

5 Inside the Cybercrime Marketplace: How Hackers Sell Access, Tools, and Services

Cybercriminals have turned personal and corporate data into commodities traded on a vast black market. Europol's 2025 threat assessment stresses that stolen information is "marketed on various criminal platforms, including specialised marketplaces, underground forums, and dedicated channels within end-to-end encrypted communication apps."
Canada's cyber authorities similarly observe that these networks are "flourishing online marketplaces" where specialized threat actors sell "leaked data and ready-to-use malicious tools"
This crime-as-a-service ecosystem enables even low-skilled actors to hire experts, purchase malware, and access stolen credentials without developing them in-house.

Dark Web Marketplaces and Forums

Dozens of marketplaces and forums openly trade in illicit goods. For example, U.S. authorities recently seized the Cracked marketplace, which sold stolen login credentials, hacking tools, and servers for hosting malware and stolen data. Cracked had over 4 million users and listed more than 28 million ads for cybercrime tools and stolen information, affecting at least 17 million U.S. victims.
Likewise, the Nulled forum offered user login data, fake IDs, hacking toolkits, and other criminal services. Before it was seized, Nulled served at least 5 million users and posted more than 43 million posts advertising hacker services.
Additionally, one recent DOJ press release described BidenCash, a payment-card marketplace, which had grown to serve over 117,000 customers. The site trafficked more than 15 million card details, personal data, and stolen credentials to facilitate unauthorized access.
In 2025, U.S. agencies seized some 145 criminal domains associated with online markets. The domains comprised disrupting sites used to enable ransomware attacks and other schemes. These takedowns suggest a vast parallel economy. Nearly anyone can buy stolen data. The Genesis Market takedown, for instance, revealed that it had offered access to data from over 1.5 million compromised computers (containing 80 million account credentials) globally, attracting criminals seeking an easy break-in.
Such markets supply the initial footholds that ransomware gangs, APT actors, and fraudsters alike rely on to launch larger attacks.

Ransomware-as-a-Service and Other Offerings

The most prominent criminal product is ransomware, which is sold as a packaged service. Ransomware-as-a-Service (RaaS) platforms supply affiliates with ready-made kits. They have lowered the bar for novice attackers to use a sophisticated ransomware strain without needing to code it. RaaS syndicates handle the malware development and updates, while paying affiliates a share of any ransom paid. Fortinet reports that, even as 13 new ransomware groups emerged in 2024, the four largest still accounted for 37% of all attacks, underscoring how the most established RaaS brands continue to dominate profits.
Other illicit tools are similarly commoditized. Automated infostealer trojans (such as RedLine, Vidar, and Lumma) are widely sold or rented. Fortinet notes these drivers of credential theft helped produce a "500% increase in credential logs on darknet forums" in 2024.
Cybercriminals also peddle phishing kits, crypting services (to make malware undetectable), bulletproof hosting, DDoS-for-hire, money laundering, and even insider access. Core RaaS groups often lease their ransomware on darknet sites or forums.
Other darknet vendors specialize in specific services. These include automated "vending cart" sites that sell card dumps for fraud, and data brokers and forums trade billions of stolen email/password pairs for credential stuffing.
Cybercrime forums often emulate legitimate marketplaces. They have escrow services, reputation scores, and affiliate programs.

Implications for the Cybersecurity Landscape

Law enforcement and industry data confirm that these tools fuel global crime. Industry analysts have linked major breaches and fraud waves back to underground sales. Cybercrime has thus evolved into a service economy, where vendors brand their wares (some RaaS groups advertise publicly), offer buyer support, and integrate cryptocurrency payments. Experts now warn that this market-driven model makes cyberattacks more scalable and profitable than ever.

6 How Hacking as a Service Is Fueling the Surge in Business Email Compromise (BEC)

Business Email Compromise (BEC) is a prevalent attack method. Fraudsters trick employees into wiring millions by impersonating trusted contacts. This trend has accelerated in 2024–2025, thanks to the rise of "Hacking-as-a-Service" (HaaS). A key trend identified in 2025 is the substantial growth (roughly 50%) in dark web offerings of turnkey phishing kits, making these attack tools far more readily available. Concurrently, BEC maintains its position as one of the single most common cybercrime tactics. One study found that more than 70% of managed service providers handle BEC-related incidents.

Hacking-as-a-Service (HaaS): A Growing BEC Enabler

HaaS platforms bundle credential-harvesting malware into easy-to-deploy phishing campaigns. Those credentials can then be used directly in BEC schemes – for example, to log into a CFO's account and request a wire transfer – or sold on dark-web credential markets.

Stolen credential markets are among the most important HaaS layers fueling BEC. Initial Access Brokers (IABs) hack company networks and charge for the access. Bitdefender analysts explain that IABs "sell verified access" to corporate networks on criminal forums. The buyer might be a BEC operator who enters the network, monitors emails, and waits for an opportune transaction or payroll message to hijack.

In effect, one hacker's compromise becomes the starting point for another's fraudulent emails. On dark markets, a single leaked admin credential or backdoor can be rented out to multiple scammers simultaneously. This supply chain model means even weak initial breaches lead to massive downstream fraud. As one expert put it, attackers can "buy" break-in points and then "move laterally" to hijack accounts and payments.

Case Studies and Trends in 2025

A cybercrime network compromised French companies in a EUR 38 million CEO fraud. One suspect impersonated a CEO and asked an accountant to urgently transfer EUR 300,000 to a Hungarian bank. An investigation into the scam revealed that the call came from Israel. Also, the same group struck a real estate developer in Paris and defrauded the company of approximately EUR 38 million. The suspects pretended to be the company's lawyers and urged the CFO to transfer the funds abroad.

A Singaporean commodity firm also suffered a BEC attack, resulting in a loss of $42.3 million. A supplier contacted the company and provided a new account through which the company was to pay a pending payment. Unfortunately, the email was from a scammer and had been slightly altered to appear to be from the official address. The firm fell for the trap and made the transfer, but was lucky enough to recover $39 million with INTERPOL's help.

In a recent federal prosecution, a transnational fraud network used romance scams, investment fraud, and BEC to steal an estimated $17 million from over 100 victims. In a different case, a ring of scammers set up dozens of lookalike corporate domains and spoofed vendor invoices to trick companies into wiring hundreds of thousands of dollars Though these press releases do not detail the offenders' tools, they fit a familiar pattern: criminals pooled resources and likely used off-the-shelf phishing templates and credential lists to cast a wide net. Security researchers note a similar trend globally. For example, a 2023 Microsoft Threat Intelligence report described a complex, multi-stage attack. Adversaries first compromised a trusted vendor, then launched adversary-in-the-middle phishing and follow-on BEC attacks against multiple banks.

Figure 1: A typical BEC attack process (image adapted from Microsoft)

Going Forward

Today's cybercriminal economy allows even unskilled fraudsters to outsource technical work. Hacking-as-a-Service offers "plug-and-play" tools to carry out phishing, harvest credentials, and bypass security measures, and those tools are increasingly tied into BEC campaigns. Companies and security teams must recognize that BEC is no longer just about a clever email hook; it is now part of a sophisticated underground supply chain. As major US agencies and industry analysts have noted, HaaS is making BEC "easier to carry out" and more widespread.