48 Minutes: How Fast Phishing Attacks Exploit Weaknesses

reliaquest.com/blog/blink-and-theyre-in-how-rapid-phishing-attacks-exploit-weaknesses/

February 20, 2025

Table of contents

Key Findings

ReliaQuest recently responded to a manufacturing sector breach involving phishing and data exfiltration. In this case, attackers achieved a “breakout time” of just **48 minutes**—the critical window between initial access and lateral movement when the potential for damage skyrockets. This figure aligns with the 2024 average and a marks 22% faster speed compared to 2023. This incident demonstrates a stark reality: Attackers are moving faster than security teams can respond, creating an urgent need for automated, faster-than-human response capabilities.

The attackers used phishing and evasion techniques commonly associated with the “Black Basta” ransomware group. Working closely with the customer, the ReliaQuest Threat Research team provided investigative support and practical remediation guidance, helping to quickly mitigate the attack’s impact. In this report we’ll cover:

- Our original research and technical findings on the attack, focusing on attacker’s speed and strategies.
- Actionable recommendations for defending against these tactics.
- A forecast on how these tactics may develop in the near future (within three months).

Attack Lifecycle

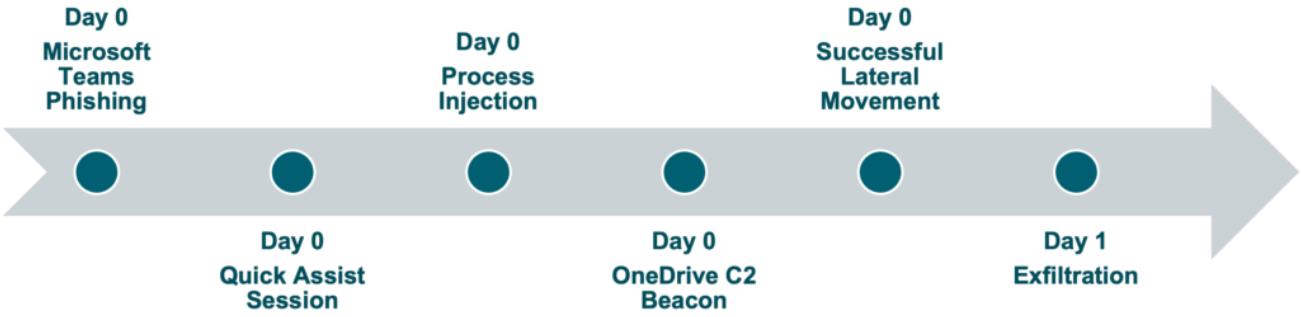

Figure 1: Timeline of the attack cycle from phishing to exfiltration

Initial Access

To gain entry into the organization's network, the threat actor used social engineering and end-user manipulation.

More than 15 users were targeted with a flood of spam emails. Next, the threat actor sent a Teams message using an external "onmicrosoft.com" email address. These domains are simple to set up and exploit the Microsoft branding to appear legitimate. The threat actor posed as an IT help-desk employee, likely pretending to assist users with the flood of emails that was preventing them from working—a common tactic used by ransomware groups like Black Basta.

The threat actor then used Teams to call at least two users and convinced them to open the remote-access tool Quick Assist, join a remote session, and grant control of their machines. Quick Assist, native to Windows hosts, is often used in these attacks because attackers can easily convince users to open it and join a remote session using a code. In this incident, one user granted the threat actor control of their machine for over 10 minutes, giving the threat actor ample time to progress their attack.

Why Does This Matter?

This tactic of using email spam instead of malicious links or attachments is particularly effective because the emails themselves aren't inherently malicious, leaving security tools with nothing to detect. Moreover, the end user doesn't need to interact with the email directly. Instead, the flood of spam makes the target's inbox unusable, giving the threat actor a plausible reason to pose as IT staff offering to resolve the issue. This low-tech but highly effective method allows threat actors to gain initial access and convince users to grant them control of their machines. Given its success, it's likely that other threat groups will adopt this technique in the near future.

Step Up Your Defenses:

- **Help-Desk Verification Procedures:** To hinder threat actors impersonating IT help-desk employees gaining initial access into networks, establish robust verification procedures to ensure end users confirm they're interacting with legitimate help-desk staff. It's also best practice to require end users to verify internal private information—such as help-desk ticket numbers, a predetermined passphrase, or their computer name— rather than information that can be easily obtained through online data breaches or social media.
- **Lock Down RMM Tools:** Configure Group Policy Objects (GPOs) to block Quick Assist and other remote monitoring and management (RMM) tools from being used for remote access. This measure prevents attackers from deceiving users into joining a session, effectively denying them initial access.

How ReliaQuest Helps You

To promptly identify a social-engineering attack, we recommend deploying the following ReliaQuest-authored detection rules.

Detection Rule	MITRE ATT&CK ID	Summary
003944 – Inbound Email Wave To Single User	TA0001:T1566 Phishing	In this attack, a mass email spam campaign preceded initial access. This detection rule identifies unusual spikes in email volume from external sources targeting a single user, enabling early alerts to mitigate spam campaigns and protect the organization from potential threats.
003981 – Suspected Microsoft Teams Phishing to Multiple Users	TA0001:T1566 Phishing	In this attack, Microsoft Teams was exploited for phishing. Following mass email spam campaigns, external actors created Entra ID tenants to impersonate help-desk staff and added targeted users to Team chats. This detection rule detects such malicious activity by identifying unusual patterns and behaviors associated with phishing attempts on Microsoft Teams.

GreyMatter Automated Response Playbooks enable remediation actions to be automatically executed as soon as a detection rule is triggered. Implementing Automated Response Playbooks can reduce the mean time to contain (MTTC) a threat to less than five minutes, significantly limiting potential damage by halting a breach in its early stages. Configuring the following Automated Response Playbook in conjunction with the detection rules above will help ensure swift containment and effective remediation of threats.

> **Disable User:** When a user interacts directly with a Teams phishing attempt, as seen in this incident, this Playbook prevents the attacker from gaining access to the user's host.

Defense Evasion

Once inside the network, the threat actor employed dynamic-link library (DLL) sideloading to evade detection. DLL sideloading works by placing a malicious payload in the same directory as a vulnerable application. Because applications prioritize loading DLLs from their own directory before searching other locations, the malicious payload is executed instead of the legitimate DLL, making it harder for security tools and analysts to detect. In this incident, the DLL “winhttp.dll” was loaded into the OneDrive update file OneDriveStandaloneUpdater.

Why Does This Matter?

ReliaQuest has frequently observed help-desk employee impersonation attempts involving files with “update” in their names. Such files appear legitimate, making them less likely to raise suspicion among users or security personnel. Malicious DLLs can masquerade as legitimate processes, complete with valid digital signatures and strong reputations, increasing the chances of causing damage like spreading malware to additional hosts on the network.

Step Up Your Defenses:

> **Expand Logging and Visibility:** To counter this tactic, deploy endpoint detection and response (EDR) sensors across critical infrastructure and the wider environment. Organizations should also forward logs to a unified location to ensure security teams have the visibility needed to detect DLL sideloading through behavioral patterns and contain malware before it spreads.

How ReliaQuest Helps You

Defend against these tactics by deploying this detection rule, designed using the most up-to-date threat intelligence.

Detection Rule	MITRE ATT&CK ID	Summary
000673 – Suspicious Process Injection	TA0004: T1055.001 – Process Injection: Dynamic-link Library Injection	Process injection was used in this attack to inject malicious code into legitimate processes, allowing the code to mimic trusted activity. This detection rule monitors processes in frequently abused system locations, flagging behavior that may indicate malicious activity.

To maximize the effectiveness of this detection rule, implement this GreyMatter Automated Response Playbook for rapid containment and efficient threat resolution.

Isolate Host: DLL sideloading indicates active malware execution on a host. Isolating the host blocks command-and-control (C2) communication and prevents lateral movement.

C2 and Lateral Movement

While the overall breakout time in this incident was 48 minutes, the attacker initiated C2 communication just **seven minutes** after using Quick Assist for initial access. They began attempting lateral movement within **eight minutes**; if these attempts had been successful, the breakout time would have been even lower.

The attacker established C2 communication through HTTPS connections over ports 443 and 10443 to the domain “uptemp[.]icu.” The attacker initially attempted to propagate the malicious DLL “winhttp.dll” using Server Message Block (SMB), embedding it into the same OneDrive update file across approximately 10 hosts on the network. When some connections failed, the attacker adapted by switching to remote desktop protocol (RDP) combined with PowerShell, demonstrating agility and persistence. PowerShell was then used to remotely create scheduled tasks that executed OneDriveStandaloneUpdater with compromised administrator accounts.

Why Does This Matter?

The rapid progression of this attack posed significant challenges for the organization’s security team in investigating and containing the threat. This narrow window of just 48 minutes for breakout time highlights the critical need for swift responses and using automation to bring down MTTC. For instance, ReliaQuest customers without automation reported an average MTTC of 6.3 hours—more than enough time for threat actors to advance through the kill chain and inflict significant damage.

Step Up Your Defenses:

Limit RDP Use: To prevent lateral movement via RDP, configure GPOs to restrict access based on specific users or hosts. Additionally, ensure that the principle of least privilege is strictly enforced, and only necessary users and hosts are allowlisted for RDP access.

How ReliaQuest Helps You

To combat these tactics, organizations should adopt the following detection.

Detection Rule	MITRE ATT&CK ID	Summary

000149 – Suspicious Scheduled Task Created	TA0002: T1053.005 – Scheduled Task/Job: Scheduled Task TA0003: T1053.005 – Scheduled Task/Job: Scheduled Task TA0004: T1053.005 – Scheduled Task/Job: Scheduled Task	To execute the injected process on other internal hosts, the attacker used scheduled tasks. This detection rule detects when a newly created scheduled task deviates from standard naming conventions or schedules a suspicious process to run.

For enhanced defenses, we recommend deploying the following GreyMatter Automated Response Playbooks alongside the above detection rule:

- **Isolate Host:** When an attacker uses a scheduled task to execute on a host, this Playbook isolates the host to stop the malicious file from spreading and block C2 communication.
- **Disable User:** This Playbook disables the account used to create the scheduled tasks to stop further malicious activity.

Privilege Escalation

In the next stage of the attack, the threat actor accessed a service account used to manage an SQL database. Due to limited logging visibility, the method of access couldn't be determined. Using the SQL service account, the threat actor created a domain admin account and domain admin permission groups, adding the additional accounts under their control. Now, the attacker had the elevated permissions necessary to exfiltrate data.

Next, the attacker used the SQL account to scan the network for vulnerable targets with SoftPerfect Network Scanner (netscan.exe). This tool, often exploited by attackers, helps to identify devices that compromised accounts have read and write access to—often serving as a precursor to lateral movement or data exfiltration.

Why Does This Matter?

Between January 2024 and July 2024, ReliaQuest found that 85% of compromises involved service accounts. These accounts are frequently targeted as they are often over-privileged and poorly secured. Service accounts serve as a critical foothold for attackers, offering weak controls that can be abused at various stages of the attack lifecycle, as observed in this attack. Their access to multiple systems across an organization further increases their value, making them even more attractive to attackers.

Step Up Your Defenses:

> **Fortify Service Accounts:** Service accounts are intended for automated processes and shouldn't be accessed by users. As such, organizations should configure service accounts to block interactive logins whenever possible. The scope of these accounts should also be restricted, ensuring they have only the permissions necessary to interact with required hosts. This minimizes the risk of exploitation and prevents attackers from pivoting to other hosts.

How ReliaQuest Helps You

We recommend deploying the following detection rule to protect against privilege escalation.

Detection Rule	MITRE ATT&CK ID	Summary
000166 – Interactive Logon with Service Account	TA0003: T1078.002 – Domain Accounts TA0004: T1078.002 – Domain Accounts TA0005: T1078.002 – Domain Accounts	In this incident, the attacker leveraged a service account to accomplish multiple objectives, including creating domain admin accounts and performing network scans. This detection rule identifies signs that a service account is being accessed by a human rather than an automated process, providing early alerts before the attacker can exploit the account.

To accompany this detection rule, organizations should leverage the following GreyMatter Automated Response Playbook:

> **Disable User:** If a service account is compromised, this Response Playbook prevents further actions from being executed using the account, giving security teams valuable time to investigate and remediate.

Exfiltration & Impact

In the final stage of the attack, the attacker leveraged their elevated permissions on the SQL service account to capture sensitive data stored on vulnerable servers. Using WinSCP, a free open-source file manager, they exfiltrated the data to a remote server under their

control, hosted at the domain “pefidesk[.]com.” In the end, the attacker completed the entire process—from initial access to exfiltrating sensitive data—in just **30 hours**.

Why Does This Matter?

Attackers take advantage of the significant brand damage caused by data breaches to pressure organizations into paying ransoms. In 2024, **80%** of breaches that we observed involved data exfiltration. Beyond financial costs, breaches strain relationships with customers and third parties, exposing personal data and irreparably eroding trust—often leading to revenue loss. In this instance, the organization took proactive measures to contain the threat by taking multiple data centers offline. While this led to operational downtime and workflow disruptions, it effectively prevented further damage.

Step Up Your Defenses:

> **Application Control:** Configure GPOs on network devices to enforce application management. These policies can prevent tools like WinSCP from executing and exfiltrating data to remote servers.

How ReliaQuest Helps You

By adopting the following detection rule, organizations can protect themselves from data exfiltration.

Detection Rule	MITRE ATT&CK ID	Summary
000063 – Outbound Web Requests from Critical Host	TA0010: T1567 – Exfiltration Over Web Service	Critical hosts like databases or domain controllers shouldn’t exhibit web traffic activity typically seen on workstations. Outbound web requests could signal an attacker using the web channel for exfiltration, like in this incident. This detection rule identifies web requests sourcing from critical hosts.

We recommend deploying the following GreyMatter Automated Response Playbook to ensure optimal protection at this stage of an attack.

> **Isolate Host:** Data exfiltration relies on outbound communication to transfer data outside the network. This Playbook isolates the compromised host, blocking communication with external hosts or domains and ensuring the data remains within the organization’s network.

Conclusion

The phishing and data exfiltration incident detailed in this report highlights a concerning reality: **Attackers are outpacing security teams**, making faster response times more essential than ever before. Based on trends observed between 2023 and 2024, we anticipate breakout times will accelerate beyond the current average of 48 minutes, leveling off at around 30 minutes. To stay protected within this shrinking critical window, organizations must integrate automation into their containment strategies.

Throughout 2024, help-desk impersonation has consistently proven to be an effective social engineering tactic, deceiving end users into granting threat actors access to their machines. We expect this method to gain further traction in 2025, potentially evolving with the use of alternative RMM tools like Qemu to enhance persistence and evade detection. While these techniques may continue to develop, the fundamental recommendations outlined in this report remain vital for defending against these increasingly sophisticated and rapid threats.

IOCs

Artifact	Details
pefidesk[.]com	Created on October 9, 2024, target for data exfiltration
uptemp[.]icu	C2 domain
c80883615157bd83dfed24683eee343a7b2ac5ab7949b3a260dc10e9f0044bb4	Malicious DLL loaded by OneDriveStandaloneUpdater.exe - winhttp.dll