

From CitrixBleed 2 to Cloudflared: The Tools and Techniques Behind Anubis Ransomware Attacks

arcticwolf.com/resources/blog/citrixbleed-2-to-cloudflared-the-tools-and-techniques-behind-anubis-ransomware-attacks

Arctic Wolf Labs

July 1, 2026



Key Takeaways

- Since the start of 2026, Arctic Wolf has investigated Anubis ransomware intrusions involving both valid VPN credential use and exploitation of CitrixBleed 2 ([CVE-2025-5777](#)), expanding known initial access tradecraft associated with this ransomware brand.
- Although tactics differ between affiliates, common patterns emerged in tradecraft through use of legitimate Remote Management and Monitoring (RMM) tooling, credential access, and hands-on-keyboard procedures used for lateral movement.
- Anubis affiliates repeatedly abused legitimate remote access and administration tools, including ScreenConnect, Zoho Assist, MeshAgent, Remotely, UltraVNC, and Total Software Deployment, to blend in with normal IT activity while maintaining control of victim systems.
- Multiple intrusions showed threat actors targeting high-value infrastructure such as Microsoft Remote Desktop Services servers, domain controllers, hypervisors, backup-adjacent systems, and Network-Attached Storage (NAS) devices, increasing operational impact and recovery complexity.
- In some intrusions, threat actors attempted to establish alternate outbound access paths using tools such as cloudflared, authenticated proxies, and SSH-based SOCKS tunneling.

Summary

Throughout 2026, Arctic Wolf has investigated multiple Anubis ransomware intrusions. Although threat actor tradecraft differs between intrusions, key themes have emerged: abuse of VPN infrastructure, blending in with legitimate activity through the use of Remote Monitoring and Management (RMM) solutions, and using other legitimate binaries on victim devices.

Public reporting has focused largely on the group's ransomware-as-a-service (RaaS) model, affiliate program, and encryptor payloads. In this research publication, we provide new insight into initial access techniques in Anubis ransomware cases, including exploitation of CitrixBleed 2 ([CVE-2025-5777](#)). We also report on the stealthy use of RMM tooling to blend in with legitimate activity, as well as various other techniques used to evade safeguards in victim environments. An understanding of these behaviors provides defenders with opportunities for early detection and containment.

Since Anubis operates as a RaaS group, this report should be seen as a compilation of affiliate-level tradecraft observed across multiple intrusions rather than the behavior of a single, uniform operator.

Background

Anubis is publicly tracked as a RaaS operation with affiliate-driven monetization, data theft, encryption, and optional destructive wipe functionality. The name Anubis emerged in late 2024 as a rebrand of the Sphinx ransomware operation, referenced by the change in encrypted file extension from .sphinx to .anubis. The new ransomware operation was formally announced on the RAMP (Ransomware and Advanced Malware Protection) cybercriminal underground forum on February 23, 2025. Since then, Anubis has evolved from a single-variant ransomware operation into a multi-platform, multi-affiliate ecosystem, claiming up to 83 victims on its data leak site.

The intrusions reviewed for this publication show how Anubis-affiliated threat actors use a practical mix of commercial tools, native Windows functionality, and commodity tradecraft to progress from access to impact. The result is activity that can resemble legitimate administration in isolation, but becomes more distinctive when viewed as a chain of related behaviors.




Figure 1: Anubis leak site screenshot

Technical Details

Initial Access Patterns

Across the intrusions reviewed in this research publication, initial access generally fell into two categories: valid VPN credential use, and exploitation of remote vulnerabilities such as

[CitrixBleed 2 \(CVE-2025-5777\)](#).

CitrixBleed 2 is a pre-authentication memory disclosure vulnerability that can expose session material from affected NetScaler appliances, creating the potential for session hijacking and multi-factor authentication (MFA) bypass when valid session tokens are obtained. In cases of exploitation, this activity also tended to originate from public IP addresses associated with VPS hosting providers.

```
0-PPE-0 : default SSLVPN TCPCONNSTAT REDACTED_CONNECTION_ID 0 : User REDACTED_USERNAME  
- Client_ip REDACTED_BROADBAND_IP - Nat_ip REDACTED_NAT_IP - Vserver  
REDACTED_VSERVER_IP:443 - Source 45.227.254[.]25:65215 - Destination  
REDACTED_DEST_IP:80 - Start_time "REDACTED_TIMESTAMP" - End_time "REDACTED_TIMESTAMP"  
- Duration 00:00:00 - Total_bytes_send 839 - Total_bytes_recv 462 -  
Total_compressedbytes_send 0 - Total_compressedbytes_recv 0 - Compression_ratio_send  
0.00% - Compression_ratio_recv 0.00% - Access Allowed - Group(s) "N/A"
```

In the example log line shown above, the threat actor's Source IP address of 45.227.254[.]25 differs from the Client_ip value, which was the broadband IP address that the session was originally established from. The discrepancy between these two fields does not universally establish maliciousness, as they are expected to diverge when a user switches between network providers (e.g., cell provider to Wi-Fi or vice versa). However, this activity still stands out as suspicious because legitimate VPN login activity is most often expected to originate from broadband providers (as was the case with the original Client_ip value) instead of VPS hosting providers.

In addition to CitrixBleed 2 exploitation, valid Cisco AnyConnect VPN logins were observed from several hosting ASNs, including AS20473 — The Constant Company and AS55286 — ServerMania. Malicious VPN authentication was then followed by login activity involving RDP and SMB, leading to credential access, PsExec service creation, RMM deployment, and ultimately invoking cloud-transfer tooling for exfiltration.

In the subset of intrusions involving valid VPN credentials, the source of those credentials could not be definitively established with available evidence. Possible sources include prior compromise of remote access infrastructure, purchase from initial access brokers (IABs), credential stuffing, or other credential-theft activity.

Lateral Movement Through RDP and PsExec

Remote Desktop Protocol (RDP) was a common mechanism for hands-on-keyboard lateral movement across the reviewed intrusions. The pattern that stood out was not RDP usage alone, but RDP from unusual source systems or VPN client ranges into infrastructure that would not typically receive interactive logons in legitimate workflows.

In contrast with routine workstation access, threat actors often used RDP between VPN client subnets and servers hosting Remote Desktop Services, Citrix-related services, file servers, domain controllers, hypervisors, and backup infrastructure.

Existing Windows servers were used in some instances as RDP pivot points into other servers in compromised environments. RDP destinations for these jump hosts included file/data infrastructure, database servers, backup servers, and domain controllers. In some instances, Hyper-V servers served as a jump point for threat actors over the course of several days. In ransomware intrusions, hypervisor servers are often targeted early for credential access because domain controller virtual disks can expose credentials if threat actors gain access to them through their underlying virtual machine (VM) storage.

In addition to RDP, PsExec was also used for lateral movement in multiple Anubis intrusions. Generally, PSEXESVC.exe service creation appeared in the same activity window as suspicious remote access, SMB or RDP movement, RMM deployment, and ransomware staging.

RMM Tool Abuse

One of the most distinctive operational patterns seen across Anubis intrusions is the deployment of legitimate RMM tools for persistent access. These solutions included ScreenConnect, Zoho Assist, MeshAgent, Remotely, Total Software Deployment, UltraVNC, and mRemoteNG. Installation of these tools on target devices provide threat actors with interactive access, file transfer, remote execution, and service-based persistence while also blending in with software commonly used by IT teams.

Zoho Assist typically appeared in multiple systems within the same environment, with artifacts under the ZohoMeeting unattended agent path and an installer identified as ZN.MSI. Components included sessionaudit.exe, agent_ui.exe, zmagent.exe, toolsmanager.exe, and related Zoho Assist service binaries.

ScreenConnect also appeared as an abused RMM solution, communicating with relay infrastructure at relay.promotds[.]us. In one instance, a ScreenConnect installer was downloaded from [https://azuremicrosoft\[.\]us/Bin/ScreenConnect\[.\]ClientSetup.msi](https://azuremicrosoft[.]us/Bin/ScreenConnect[.]ClientSetup.msi), a domain name crafted to resemble Microsoft Azure infrastructure. ScreenConnect artifacts appeared alongside PsExec service creation, Anubis encryptor staging under C:\Users\Default\AppData\, and S3 Browser components under the same user profile.

MeshAgent was deployed under filenames including mvtcs.exe and sysagent.exe, with related configuration and database artifacts such as mvtcs.msh, mvtcs.db, and mvtcs.log. In one intrusion, persistence was reinforced through a scheduled task named MeshUserTask,

while MeshAgent activity appeared alongside cloudflared, SSH tunneling, and Anubis ransomware deployment across Windows and Linux systems.

Remotely Desktop was installed on multiple Windows Server hosts and appeared as Remotely_Desktop.dll, with associated C2 communications to 45.76.79[.]92.

Total Software Deployment was used as an administrative deployment mechanism, including creation of the Total Software Deployment Audit Service and execution of tniwinagent.exe from %SystemRoot%\TNIWINAGENT\. In the same activity cluster, affiliates used the tool to push agents across multiple hosts, creating service-install telemetry that stood out alongside SMB logons, RDP activity, and later ransomware execution.

Tunneling and Proxy Infrastructure

cloudflared is the Cloudflare Tunnel client and can be used to create outbound tunnels from an internal host to Cloudflare-managed infrastructure.

In some intrusions, threat actors were observed installing and attempting to establish tunnels to victim environments using Cloudflare Tunnel (cloudflared) on network attached storage (NAS) devices or Windows servers. Windows cloudflared.exe binary artifacts were found in the C:\Windows directory, while the /usr/local/etc/cloudflared directory was used on Synology NAS devices.

In one intrusion involving a Synology NAS device, the threat actor logged into DiskStation Manager using a compromised domain account, created a dedicated local administrator account on the NAS device, and granted it broad privileges including access to DiskStation Manager, SMB, SFTP, FTP, File Station, and rsync. The same account was then used to establish a SFTP connection to transfer several megabytes of files to the NAS device. Subsequent SSH sessions were then established from an internal device, followed by sudo su to obtain a root shell. While exact filenames from the SFTP file transfer were not captured, a cloudflared command referencing cert.pem and config.yml was executed several minutes afterwards.

From the root shell, the actor staged remote access tooling and worked from /usr/local/etc/cloudflared to configure Cloudflare Tunnel. The commands show creation of a named tunnel, repeated execution with HTTP/2, and background launch attempts using nohup. Captured output showed tunnel startup metadata, including a tunnel ID and connector ID, but also repeated TLS handshake and certificate validation failures. The available logs do not confirm that a usable Cloudflare Tunnel was successfully established.

The threat actor tested multiple egress paths. They attempted to run cloudflared forwarding its communications through an authenticated HTTP proxy hosted on AS44477 – Stark Industries Solutions Ltd and tested the proxy with curl. After the Cloudflare Tunnel attempts, the actor configured SSH-based dynamic forwarding through VPS infrastructure hosted on AS399629 – BI Networks. Shell command history showed ED25519 key generation, public-key installation steps, validation commands, and creation of a local SOCKS proxy on 127.0.0.1:1080 using `ssh -D`. While the SOCKS proxy was separate from the authenticated HTTP proxy used with cloudflared, both served the same practical goal: routing outbound traffic from the NAS through attacker-controlled or attacker-accessible infrastructure.

Within a day of configuring these egress paths, the threat actor returned to the NAS and executed the Linux Anubis encryptor against NAS storage volumes, alongside encryption of Windows-based assets throughout the environment.

Credential Access

Mimikatz use on Windows servers was commonly observed between multiple intrusions, and was typically staged in locations such as `C:\Users\Public\Videos\mimikatz.exe`, `C:\Users\Public\mimikatz.exe`, and `C:\x64\mimikatz.exe`.

Additionally, browser credential extraction appeared through filenames such as `Chrome Passwords.csv` and `Microsoft Edge Passwords.csv`. These filename artifacts are consistent with saved browser password export activity through Google Chrome and Microsoft Edge manually rather than direct extraction of Windows logon material from memory. Browser password exports can expose VPN, cloud, SaaS, administrative portal, or personal account credentials that support follow-on access and extortion workflows.

Active Directory `ntds.dit` database access was also observed through Windows Extensible Storage Engine logs, as evidenced by creation of a new database file under `C:\audit\Active Directory\ntds.dit` in a representative intrusion. Within a minute of this database copy being saved to the filesystem, an archive named `Active Directory.zip` was also saved to the desktop. This type of access can enable offline extraction of domain credential data, giving ransomware threat actors a path to broader account compromise even if individual interactive sessions are disrupted. Within less than an hour of extracting `ntds.dit`, ransomware encryption began across the environment.

Cloud Storage and Exfiltration Tooling

Several intrusions included deployment of tooling commonly used for data transfer or exfiltration, including S3 Browser, `rcclone`, `s5cmd`, WinSCP, and PuTTY.

In one intrusion, S3 Browser artifacts appeared alongside ScreenConnect activity and Anubis encryptor staging. Available evidence shows that, in several instances, threat actors installed the software through the s3browser-13-1-1.exe installer, after downloading the installer to the user Downloads directory. S3 Browser executables were found in several directories:

C:\Program Files\S3 Browser

C:\Users\Default\AppData\Local\S3 Browser

Other cases included rclone and s5cmd artifacts. These utilities are commonly used for high-volume interaction with cloud storage services and can be attractive to ransomware affiliates because they are efficient, scriptable, and compatible with multiple storage backends. In the reviewed Anubis intrusions, these tools appeared alongside other pre-encryption behaviors, including RMM deployment, credential access, and security tool tampering.

Defense Evasion

Threat actors in Anubis ransomware intrusions attempted to weaken defensive visibility and complicate post-incident analysis. These techniques included Windows Defender real-time protection disablement, SophosUninstall activity, PCHunter-related artifacts, and log clearing or manipulation across multiple systems.

Evasive activity varied across intrusions. In some intrusions, endpoint protection tampering appeared primarily as direct disablement or attempted uninstallation of security tooling. In others, log clearing affected visibility across key Windows event sources, including System, PowerShell, Task Scheduler, AppLocker, and Defender logs. In at least one intrusion, an Anubis encryptor was deleted after execution, reducing the availability of on-disk payload artifacts for later analysis. These actions suggest an effort to hamper the ability of defenders to document malicious activity throughout the incident.

Anubis Ransomware Deployment

Anubis ransomware deployment on Windows and Linux generates encrypted files with the .anubis extension and ransom notes named RESTORE FILES.html. Microsoft Defender detections included Ransom:Win64/Anubis.A.

On Windows, encryptor filenames included names such as win.exe, wmi.exe, s.exe, {REDACTED_6_DIGIT_NUMBER}_win64_encrypt.exe. Linux encryptors followed the naming convention of {REDACTED_6_DIGIT_NUMBER}_encrypt_x86_64.

These encryptor binaries showed up in various directories on the filesystem, including the following:

C:\Apps
C:\PerfLogs
C:\Users\REDACTED\Desktop
C:\Users\REDACTED\AppData
C:\Users\REDACTED\AppData\Local
C:\Users\Public
C:\Windows\Temp\netscan




Figure 2: Anubis leak site victim postings tracked by month

Conclusion

Anubis intrusions reviewed by Arctic Wolf show a ransomware operation that does not need a novel toolchain to create serious impact. Many of the most notable behaviors were less significant when viewed in isolation. Their significance came from how threat actors combined them into a repeatable path from initial access to extortion.

The most important defensive pattern is not any single indicator or tool, but the sequence of activity: suspicious remote access → unusual RDP or SMB movement → unauthorized RMM deployment → credential access → security control tampering → exfiltration tooling → staged ransomware execution. Organizations that can detect and disrupt this chain before encryption are better positioned to contain Anubis and similar ransomware operations.

Anubis should therefore be understood as more than just a ransomware payload. It is an affiliate ecosystem that uses flexible, practical tradecraft to turn common administrative pathways into extortion infrastructure. The best opportunities for disruption appear before encryption, when authentication anomalies, RMM deployment, credential access, and exfiltration tooling begin to cluster around the same hosts and accounts.

Arctic Wolf has [Aurora® Managed Detection and Response \(MDR\)](#) detections in place for activity observed across these intrusions.

Arctic Wolf Labs will continue to monitor Anubis and update detection capabilities and public guidance as new intelligence becomes available.

How Arctic Wolf Protects Its Customers

Arctic Wolf is committed to ending cyber risk, and when active campaigns are identified, we move quickly to protect our customers. We have leveraged threat intelligence around this threat activity to enhance detections in the [Aurora® Superintelligence Platform](#), subject to customer environment and available telemetry. As this campaign develops, Arctic Wolf may refine detections for additional indicators of compromise and techniques leveraged by this threat.

Defensive Guidance

The following recommendations are based on the techniques and infrastructure observed across Anubis intrusions. They are applicable both to defending against Anubis specifically and to the broader class of Living-off-the-Land (LOTL)/RMM-abuse ransomware adversaries.

High Priority Mitigations

- **Patch CVE-2025-5777 immediately.** Organizations running Citrix NetScaler ADC/Gateway should confirm patching status against CVE-2025-5777 (CVSS 4.0: 9.3, CISA KEV). Given the pre-authentication nature of this vulnerability, no threat actor interaction with legitimate users is required for exploitation. Post-patch, follow [vendor-provided instructions](#) for terminating all active sessions so that unauthorized activity is terminated.
- **Audit and restrict remote management tool installations.** Maintain a list of approved RMM tools and monitor for installation of unauthorized alternatives. The deployment of multiple RMM tools within a short timeframe is a potential indicator of adversary persistence establishment.

- **Block known malicious infrastructure.** Implement network-level blocks for the indicators documented in the Appendix, including the typosquatted domains azuremicrosoft[.]us and promotds[.]us.

Endpoint Hardening

- **Alert on endpoint protection tampering.** Implement tamper protection for endpoint security agents and monitor for exe execution, Defender policy modifications, PCHunter deployment, and log clearing. Endpoint protection disablement across multiple systems in rapid succession is a pre-encryption indicator.
- **Restrict execution from staging directories.** Anubis affiliates staged payloads and tooling from several non-standard or user-writable locations, including C:\Apps, C:\PerfLogs, C:\Users\{username}\AppData, C:\Users\Public, user profile directories, and temporary tooling folders. Application control policies should alert on or block suspicious execution from these locations where operationally feasible.

Incident Response Readiness

- **Maintain offline backups with tested restoration procedures.** The existence of destructive functionality in Anubis means organizations should assume that accessible backups or storage infrastructure may be targeted during an intrusion. Backup systems should be isolated, monitored, and tested regularly.
- **Segment Hyper-V and NAS infrastructure.** Arctic Wolf observed activity involving Hyper-V directories, VHDX/VMDK storage, and NAS mount points such as /volume1/ and /volume2/. Virtualization management planes, NAS devices, and backup storage should be segmented from general-purpose endpoints and monitored for unusual interactive access.

Appendix

For Appendix sections referenced in this report, including Indicators of Compromise (IOCs), network indicators, file indicators, MITRE ATT&CK® mapping, ransomware encryptor artifacts, and remote tooling artifacts, please see our [public GitHub repository](#).

Additional Arctic Wolf Resources:

- Arctic Wolf's Threat Intelligence newsletter: [ThreatPulse Community Edition](#)
- [Arctic Wolf Tech Den](#)
- [Arctic Wolf Blog](#)

Legal disclaimer: Attribution reflects Arctic Wolf Labs' assessment as of the report period and may evolve with new evidence. References to threat actor identity, nexus, and intent are analytical judgments, not statements of legal fact. This alert is provided for informational purposes only and does not constitute a guarantee of detection or prevention. Defensive effectiveness varies by environment, configuration, and available telemetry.

[Arctic Wolf Labs](#) is a group of elite security researchers, data scientists, and security development engineers who explore security topics to deliver cutting-edge threat research on new and emerging adversaries, develop and refine advanced threat detection models with artificial intelligence and machine learning, and drive continuous improvement in the speed, scale, and detection efficacy of Arctic Wolf's solution offerings.

Arctic Wolf Labs brings world-class security innovations to not only Arctic Wolf's customer base, but the security community at large.