ODIN

TLP:CLEAR

ODIN RESEARCH · MONOGRAPH 50 / 30 · RECON WORKBENCHES

ScopeSentry and Typo-Squatted Certificates

Sorting C2 From Dual Use

How 2,002 ScopeSentry, CyberStrikeAI, Sliver, GoPhish, and Evilginx endpoints plus a Netherlands PHP phishing kit narrow to 30 upstream-confirmed C2 candidates

Two-week observation window ending 21 November 2026 · Prepared November 2026
Full Report · Public edition · Version 1.1

33,042
Global feed sightings across ScopeSentry, CyberStrikeAI, Sliver, GoPhish, and Evilginx

2,002
Endpoints examined by Odin by Brandefense across 1,781 distinct hosts and 1,515 / 24 subnets

30
Upstream-confirmed IPv4 candidates (32 endpoints: 3 CyberStrikeAI, 8 Sliver, 21 GoPhish; withheld pending corroboration)

500
Dual-use ScopeSentry endpoints on 486 hosts (450 on TCP 8082 or 18082; all withheld)

What Is in This Report. 11 chapters and two appendices.

How 2,002 ScopeSentry, CyberStrikeAI, Sliver, GoPhish, and Evilginx endpoints plus a Netherlands PHP phishing kit narrow to 30 upstream-confirmed C2 candidates

01	Executive Summary	04
02	Four Functional Groups, One Analytical Model	07
03	Dual-Use ScopeSentry Consoles and CyberStrikeAI Certificate Typos	11
04	Nine TLS Certificate Lure Classes Across Sliver, GoPhish, and Evilginx	14
05	Ten Exposed PHP and Log Files on a Netherlands Apache Server	18
06	Global Feed Sightings Against Upstream-Confirmed Candidates	21
07	What Defenders Should Do	23
08	Detection Rules	25
09	Indicators	27
10	Scope and Limits	29
11	Sources	31
A	Indicators of compromise (4)	33
B	Detection rules (4 files)	34

IN SHORT

In a two-week observation window ending 21 November 2026, Odin by Brandefense reviewed 2,002 endpoints across 1,781 hosts from 33,042 feed sightings spanning ScopeSentry, CyberStrikeAI, Sliver, GoPhish, and Evilginx. We also reviewed directory metadata for a ten-file PHP phishing kit in the Netherlands. Filtering 500 dual-use ScopeSentry consoles, 62 subscriber ISP and CDN edge collisions, and 1,408 unverified candidates leaves 32 upstream-confirmed endpoints on 30 unique IPv4 addresses. Review those 30 withheld candidates before blocking, inspect X.509 certificate subjects and trust chains, and audit web roots.

CHAPTER 01

Executive Summary

IN SHORT

In a two-week observation window ending 21 November 2026, Odin by Brandefense reviewed 2,002 endpoints across 1,781 hosts from 33,042 feed sightings spanning ScopeSentry, CyberStrikeAI, Sliver, GoPhish, and Evilginx. We also reviewed directory metadata for a ten-file PHP phishing kit in the Netherlands. Filtering 500 dual-use ScopeSentry consoles, 62 subscriber ISP and CDN edge collisions, and 1,408 unverified candidates leaves 32 upstream-confirmed endpoints on 30 unique IPv4 addresses. Review those 30 withheld candidates before blocking, inspect X.509 certificate subjects and trust chains, and audit web roots.

INSIDE THIS CHAPTER

Three Decisions for Leaders 04

The Study in Numbers 05

01 Executive Summary

External attack-surface management (`EASM`) consoles, autonomous AI penetration-testing platforms, implant command-and-control (`C2`) servers, and phishing kits share cloud networks and TLS patterns. Security teams and bug-bounty hunters use ScopeSentry to map external subdomains and open ports. Red teams and intrusion operators use CyberStrikeAI, Sliver, GoPhish, and Evilginx behind custom or deceptive X.509 certificates.

Pushing raw threat feeds for these tools straight into firewalls blocks authorized assessments and cuts off broadband subscribers or shared CDN edges. We reviewed all 2,002 endpoints and the ten-file Netherlands PHP kit statically without running code or contacting external hosts.

1. **Dual-use ScopeSentry consoles cluster on port 8082 and need watchlist monitoring.**
   Across 500 examined ScopeSentry endpoints on 486 hosts from 532 global sightings, 446 listen on TCP 8082 and four on 18082 (90.0% combined). Because corporate security teams and bug-bounty researchers operate ScopeSentry legitimately, all 500 addresses stay off IP blocklists, including seven marked confirmed upstream. Defenders should use external attack surface management to audit internal perimeters for unapproved port 8082 login pages instead of blocking ScopeSentry addresses.

2. **CyberStrikeAI and C2 operators use typo-squatted CDN and brand-lure certificates.**
   Across 500 examined CyberStrikeAI endpoints on 467 hosts, 88 expose the default `CN=cyberstrikeai` self-signed certificate and 74 share one Go TLS JARM fingerprint. Other CyberStrikeAI endpoints present custom domains or typo-squatted Cloudflare subjects such as `CN=cdn[.]cloduflare[.]com`.

Across all 2,002 endpoints, 44 unverified listeners present OAuth, webmail, financial, or portal-mimic X.509 subjects, including lures on TCP port 31337 Sliver listeners. Look-alike certificate names like these are what phishing monitoring is designed to flag. Filtering subscriber ISPs and single-signal rows leaves 32 upstream-confirmed endpoints on 30 unique IPv4 addresses; all 30 stay withheld pending independent, address-specific corroboration.

3. **Ten exposed files in the Netherlands show a modular PHP phishing kit.**
   Separate from the endpoint set, an Apache server in the Netherlands exposes ten files totaling 17,101 bytes with zero shared hosts across the C2 set.

Directory metadata lists twin 1,638-byte scripts whose names suggest log retrieval and deletion beside files named for Base64 DOM obfuscation, cipher routines, database access, and client filtering. Their contents were not available to verify those functions. Because third-party web-host ownership cannot be ruled out, the server address, file paths, and file hashes stay withheld.

Three Decisions for Leaders

1. **Require address-specific corroboration before blocking upstream-confirmed cloud C2 candidates.** Review the 30 upstream-confirmed CyberStrikeAI, Sliver, and GoPhish address candidates against local telemetry; do not publish or block an address until independent, address-specific corroboration is recorded.
2. **Inspect X.509 certificate subjects and trust chains instead of trusting familiar CDN or OAuth names.** Configure TLS-inspecting proxies to alert when self-signed or untrusted certificates present subjects claiming `cdn[.]cloduflare[.]com`, `CN=cyberstrikeai`, single-sign-on portals, webmail verification pages, or banking brands.

3. **Keep ScopeSentry off automated IP blocklists and audit web servers for exposed PHP kits.** Use ScopeSentry intelligence strictly to audit internal port 8082 and 18082 exposure, disable open directory listings, and use heuristic YARA rules to review web roots.

The Study in Numbers

Number	What it is
33,042	Global feed sightings across ScopeSentry, CyberStrikeAI, Sliver, GoPhish, and Evilginx
2,002	Endpoints examined by Odin by Brandefense across 1,781 distinct hosts and 1,515 `/24` subnets
30	Upstream-confirmed IPv4 candidates (32 endpoints: 3 CyberStrikeAI, 8 Sliver, 21 GoPhish; withheld pending corroboration)
500	Dual-use ScopeSentry endpoints on 486 hosts (450 on TCP 8082 or 18082; all withheld)
44	Endpoints across 41 hosts presenting OAuth, webmail, financial, or portal-mimic X.509 certificates
62	Subscriber telecom ISP and shared CDN edge endpoints across 53 hosts suppressed from blocking
10 / 17,101 B	Exposed files and total bytes on the Netherlands Apache PHP phishing log-panel server

How to use this report.

Security leaders can start with Chapters 1, 6, and 7. Cloud security, PKI, and attack-surface engineers should read Chapters 3, 4, 5, and 7. SOC analysts and detection engineers should go straight to Chapters 8 and 9. Threat intelligence analysts will find full cohort tables and scope limits in Chapters 2, 4, 6, and 10.

02

CHAPTER 02

Four Functional Groups, One Analytical Model

IN SHORT

Zero shared IP addresses or TLS certificates connect the 30 upstream-confirmed C2 and phishing candidates or the Netherlands PHP kit to the ScopeSentry cohort. We present the four functional groups together as an analytical model for defenders.

INSIDE THIS CHAPTER

How the five endpoint families and the PHP kit compare 08

Why upstream feed tags cannot drive automated IP blocks 08

02 Four Functional Groups, One Analytical Model

Intrusion and credential-theft operations rely on four distinct building blocks: external reconnaissance consoles, autonomous AI offense panels, implant C2 and spearphishing relays, and lightweight PHP credential-collection kits. In commercial threat feeds, all four layers arrive mixed under `confirmed`, `suspected`, or `candidate` tags. Blocking every feed entry without cyber threat intelligence review disrupts authorized scanners, blacklists broadband subscribers, and cuts off shared CDN proxies.

We compared host addresses, `/24` subnets, and X.509 certificates across all 1,781 hosts in the 2,002-endpoint set and the Netherlands Apache server. Zero shared hosts link the 30 upstream-confirmed IPv4 candidates or the PHP kit to any other family.

Within the remaining cohorts, 13 hosts expose two families at once. Five pair GoPhish with ScopeSentry, three pair CyberStrikeAI with ScopeSentry, three pair ScopeSentry with Sliver, and two pair CyberStrikeAI with Sliver. Because no technical indicator ties these populations to one actor, we present them as an analytical model. Figure 1 shows the four groups and their defender actions.

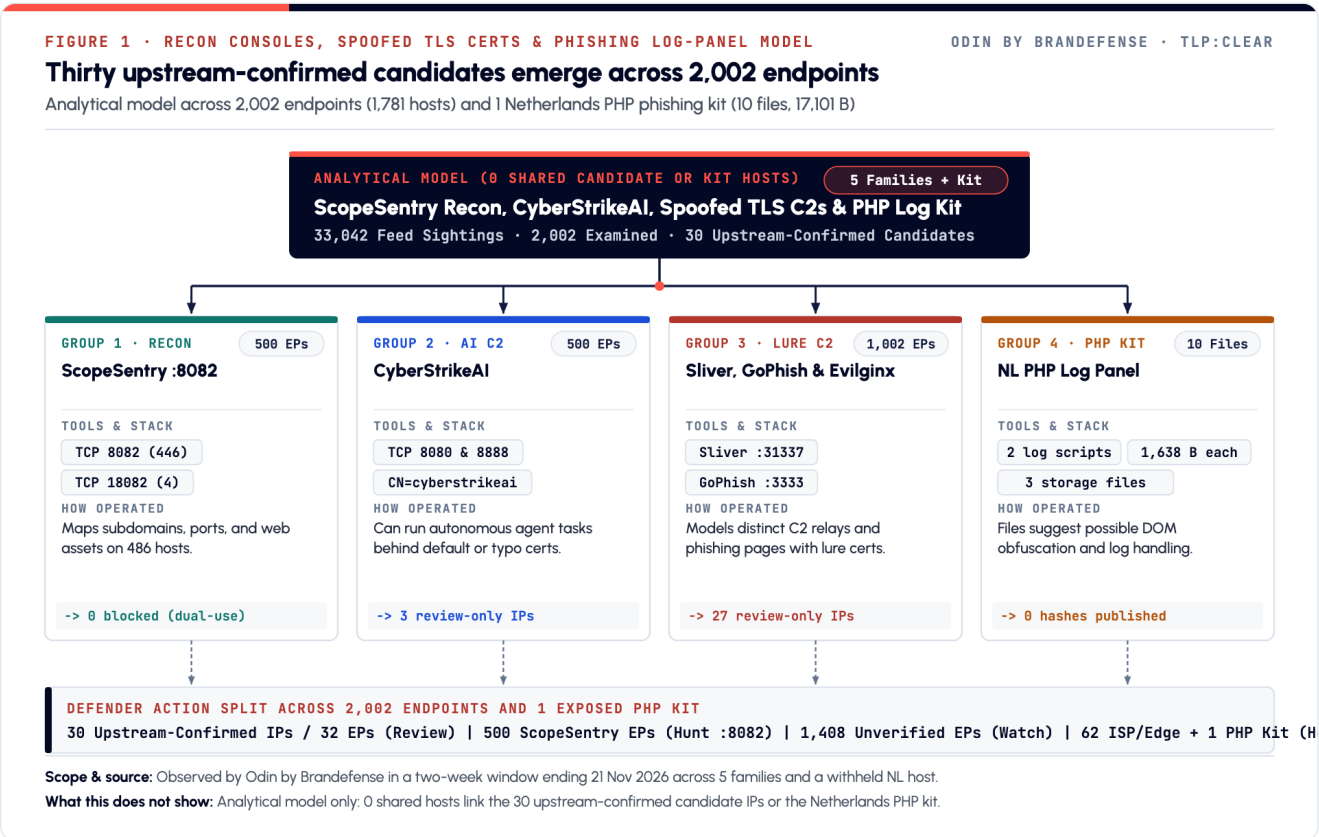


Figure 1: Thirty upstream-confirmed candidates emerge across 2,002 endpoints

How the five endpoint families and the PHP kit compare

Across the five endpoint families, upstream feeds listed 33,042 global sightings (71 marked verified upstream). We examined 2,002 endpoints across 1,781 distinct hosts and 1,515 `/24` subnets (1,794 summed per-family hosts before the 13 two-family hosts).

Family or kit	Functional role	Global sightings	Verified	Countries	Examined	Hosts / `/24`s	Upstream-confirmed (review)	Suppressed ISP/edge	Withheld (watch)
GoPhish	Phishing server	23,355	36	119	501	455 / 447	**21** (19 IPs)	**26**	454
Sliver	Implant & mTLS C2	5,670	19	83	500	385 / 375	**8** (8 IPs)	**22**	470
CyberStrikeAI	AI offense console	3,484	16	42	500	467 / 273	**3** (3 IPs)	**14**	483
ScopeSentry	Dual-use EASM	532	0	32	500	486 / 449	0	0	500
Evilginx	AiTM proxy	1	0	1	1	1 / 1	0	0	1
5-Family Subtotal	**Endpoints**	**33,042**	**71**	**--**	**2,002**	**1,781 / 1,515**	**32 (30 IPs)**	**62**	**1,908**
Netherlands PHP kit	Apache PHP kit	10 files	0	1	10 (17,101 B)	1 / 1	0	1 (log)	9 (review)

Two GoPhish hosts each appeared on TCP 80 and TCP 443. Thus, 32 upstream-confirmed endpoints resolve to 30 unique addresses. In upstream metadata, 663 rows carry record first-seen dates earlier than 7 November 2026 (extending back to record dates of 8 August 2025). All Odin observations occurred during the two-week window ending 21 November 2026.

Why upstream feed tags cannot drive automated IP blocks

Upstream feeds marked 78 of the 2,002 examined endpoints `confirmed`, 215 `suspected`, and 1,709 `candidate`. Blocking all 78 upstream `confirmed` rows would immediately misfire on 46 records (59.0%). Seven are dual-use ScopeSentry consoles, six sit on telecom or shared CDN networks, and 33 sit on multi-tenant hosts (HostPapa, Alibaba, Tencent) without relay proof. The remaining 32 upstream-confirmed rows (30 unique IPv4s) record only feed status, hosting, JARM, and certificate fields without an independent response, beacon, or workbench link. All 30 therefore stay withheld for review rather than automated blocking. Figure 2 maps each slice to TLS, URI, and heuristic YARA detections.

FIGURE 2 · CERTIFICATE LURE & PANEL DETECTION MATRIX

ODIN BY BRANDEFENSE · TLP:CLEAR

TLS subject, URI, and PHP structure rules complement raw IP feeds

Six infrastructure and certificate slices across 2,002 endpoints and 10 PHP kit files mapped to detections



Scope & source: Observed by Odin by Brandefense during the two-week window ending 21 Nov 2026.
What this does not show: Dual-use ScopeSentry hosts, ISP/edge collisions, unverified candidates, and the NL host stay off blocklists.

Figure 2: TLS subject, URI, and PHP structure rules complement raw IP feeds

CHAPTER 03

Dual-Use ScopeSentry Consoles and CyberStrikeAI Certificate Typos

IN SHORT

ScopeSentry (500 endpoints on 486 hosts) and CyberStrikeAI (500 endpoints on 467 hosts across 273 `/24` subnets) show why defenders must separate dual-use recon tools from offensive agent servers. Nine in ten ScopeSentry instances (450 of 500) listen on TCP 8082 or 18082 and stay withheld. CyberStrikeAI yields three upstream-confirmed candidate addresses, 88 `CN=cyberstrikeai` certificates, and `CN=cdn[.]cloduflare[.]com` certificate typos. Audit port 8082 internally and alert on both certificate subjects.

INSIDE THIS CHAPTER

ScopeSentry (500 endpoints on 486 hosts): port 8082 concentration 11

CyberStrikeAI (500 endpoints): `/24` sweeps, `CN=cyberstrikeai`, and `cdn[.]cloduflare[.]com` 11

03 Dual-Use ScopeSentry Consoles and CyberStrikeAI Certificate Typos

External asset discovery tools and AI-orchestrated offense panels sit side by side in assessment workflows (T1595.002, T1583.003, T1587.003). Corporate security teams run ScopeSentry continuously to inventory subdomains, whereas CyberStrikeAI exposes automated exploit agents behind default or spoofed TLS certificates.

ScopeSentry (500 endpoints on 486 hosts): port 8082 concentration

Open-source ScopeSentry pairs a Python and Vue web console backed by MongoDB and Redis with distributed nodes running Subfinder, Naabu, Httpx, Katana, Rad, and Nuclei. Across 532 sightings in 32 countries (0 verified upstream), our 500 examined endpoints on 486 hosts (449 `/24` subnets) show three clear patterns:

- **90.0% concentration on TCP 8082 and 18082.** Exactly 446 endpoints (89.2%) listen on default Docker console port 8082, and four listen on container port 18082 in Hong Kong (1), China (2), and France (1). Only 50 endpoints use other ports, led by 443 (7), 80 (6), and 8443 (4).
- **Cloud and VPS distribution.** China (155, 31.0%), Hong Kong (105, 21.0%), the United States (90, 18.0%), Singapore (32), Japan (31), South Korea (15), France (13), and the Netherlands (11) lead geographically. Top hosts are Tencent (77), Alibaba (52 combined), HostPapa (37), LUCIDACLOUD (18), cognetcloud (16), Contabo (9), and Vultr (9).
- **Plain HTTP login pages and watchlist withholding.** Over seven in ten instances (354, 70.8%) serve `POST /api/user/login` without a certificate CN. The rest use 107 custom domains, 14 IP certificates, 10 `CN=localhost`, 8 portal mimics, and 7 Windows host certificates. ScopeSentry also sits on 11 of the 13 two-family hosts, and all 500 addresses remain withheld.

CyberStrikeAI (500 endpoints): `/24` sweeps, `CN=cyberstrikeai`, and `cdn[.]cloduflare[.]com`

CyberStrikeAI is an AI-driven offense panel in Go and Python exposing HTTP/WebSocket APIs and Model Context Protocol (`MCP`) wrappers over more than 100 security tools. Across 3,484 sightings in 42 countries (16 verified upstream), our 500 examined endpoints (467 hosts across 273 `/24` subnets) show four distinct traits:

- **Port distribution (8080, 8888, 18666, 8444).** Console port 8080 leads with 230 endpoints (46.0%, plus 5 on 18080), followed by 8888 (96, 19.2%), 8085 (24), 443 (22), and 80 (17). Ports 8081 (11), 8443 (10), 8084 (9), and 8087 (9) follow. Seven endpoints listen on MCP/agent port 18666 on PEG TECH INC in the United States. One Hong Kong host listens on port 8444 alongside five Sliver and one ScopeSentry port 8444 listeners.
- **Single-subnet sweeps in Lithuania and the Netherlands.** Exactly 95 Lithuanian hosts on UAB Host Baltic (94 on port 8888, 1 on 9443) sit inside one `/24` subnet. Another 38 Dutch hosts (44 endpoints on ports 8084 to 8087 and 9090) on NForce Entertainment B.V. also sit inside one `/24` subnet. Cloudie Limited (37), Tencent (35), Alibaba (35 combined), and CTG Server (34) follow.
- **Default `CN=cyberstrikeai` (88 endpoints) and 3 upstream-confirmed candidates.** Eighty-eight endpoints (17.6%, plus 2 `CN=localhost`) serve `CN=cyberstrikeai`, and 74 share Go TLS JARM `40d40d40d00000000043d43d00043d6aff5ab0f4fc31897186312c9e639319`. Three cloud servers (two on AWS US and one on OVH France) carry upstream `confirmed` tags while serving `CN=cyberstrikeai`. All three addresses remain withheld pending independent, address-specific corroboration.

- **Typo-squatted Cloudflare certificates (`CN=cdn[.]cloduflare[.]com`).** The cohort also includes custom-domain certificates (110 rows across 96 domains), IP certificates (45), Windows hostnames (13), and brand lures (6). Profiling CyberStrikeAI TLS certificates also reveals self-signed X.509 certificates with `CN=cdn[.]cloduflare[.]com`, swapping `ud` to `du` in `cloudflare` (T1587.003). Under RFC 5280, self-signed certificate generation checks no DNS ownership. The subject can resemble a Cloudflare name in a quick review; hunt the exact typo in certificate subjects. Non-brand custom domains are grouped under custom domain certificates in Chapter 4.

Attribute	ScopeSentry (dual-use watchlist)	CyberStrikeAI (offense & MCP C2)
Sightings / examined / hosts / `/24`s	532 / 500 / 486 / 449 (1.08 hosts/ `/24`)	3,484 / 500 / 467 / 273 (1.71 hosts/ `/24`)
Main TCP ports	8082 (446, 89.2%), 443 (7), 80 (6), 18082 (4), 8443 (4)	8080 (230, 46.0%), 8888 (96), 8085 (24), 443 (22), 80 (17), 18666 (7), 8444 (1)
Top providers	Tencent (77), Alibaba (52), HostPapa (37), LUCIDACLOUD (18)	UAB Host Baltic (95 in 1 `/24`), NForce (44 in 1 `/24`), Cloudie (37), Tencent (35)
Upstream-confirmed IPv4s	0 (all 500 withheld as dual-use EASM)	3 (2 AWS US on :8080 and :80, 1 OVH FR on :8080; all withheld for review)

04

CHAPTER 04

Nine TLS Certificate Lure Classes Across Sliver, GoPhish, and Evilginx

IN SHORT

Across all 2,002 examined endpoints, X.509 certificates fall into nine classes led by no-CN endpoints (1,055), custom domains (605), and framework defaults (165). Forty-four endpoints across 41 hosts present explicit OAuth, webmail, financial, or portal-mimic lures, while 8 Sliver and 19 GoPhish upstream-confirmed IPv4 candidates remain withheld pending address-specific corroboration. Inspect X.509 certificate subjects and trust chains on TLS proxies.

INSIDE THIS CHAPTER

Nine-class X.509 certificate breakdown (2,002 endpoints) 15

Inside the 44 OAuth, webmail, financial, and portal-mimic certificates 15

Sliver (8 upstream-confirmed IPs), GoPhish (19 upstream-confirmed IPs), and Evilginx (1 candidate) 16

04 Nine TLS Certificate Lure Classes Across Sliver, GoPhish, and Evilginx

Phishing servers (`GoPhish`, `Evilginx`) and implant C2 frameworks (`Sliver`) customize X.509 certificates for user deception and passive log camouflage (T1587.003, T1566.002, T1557). On browser-facing GoPhish (443, 80) and Evilginx ports, SSO, webmail, and banking certificates reinforce credential lures. On non-browser C2 ports like Sliver port 31337, operators load OAuth or regional webmail-verification names into X.509 certificates so persistent C2 traffic blends into TLS inventories. Figure 3 traces this modeled handshake and collection flow.

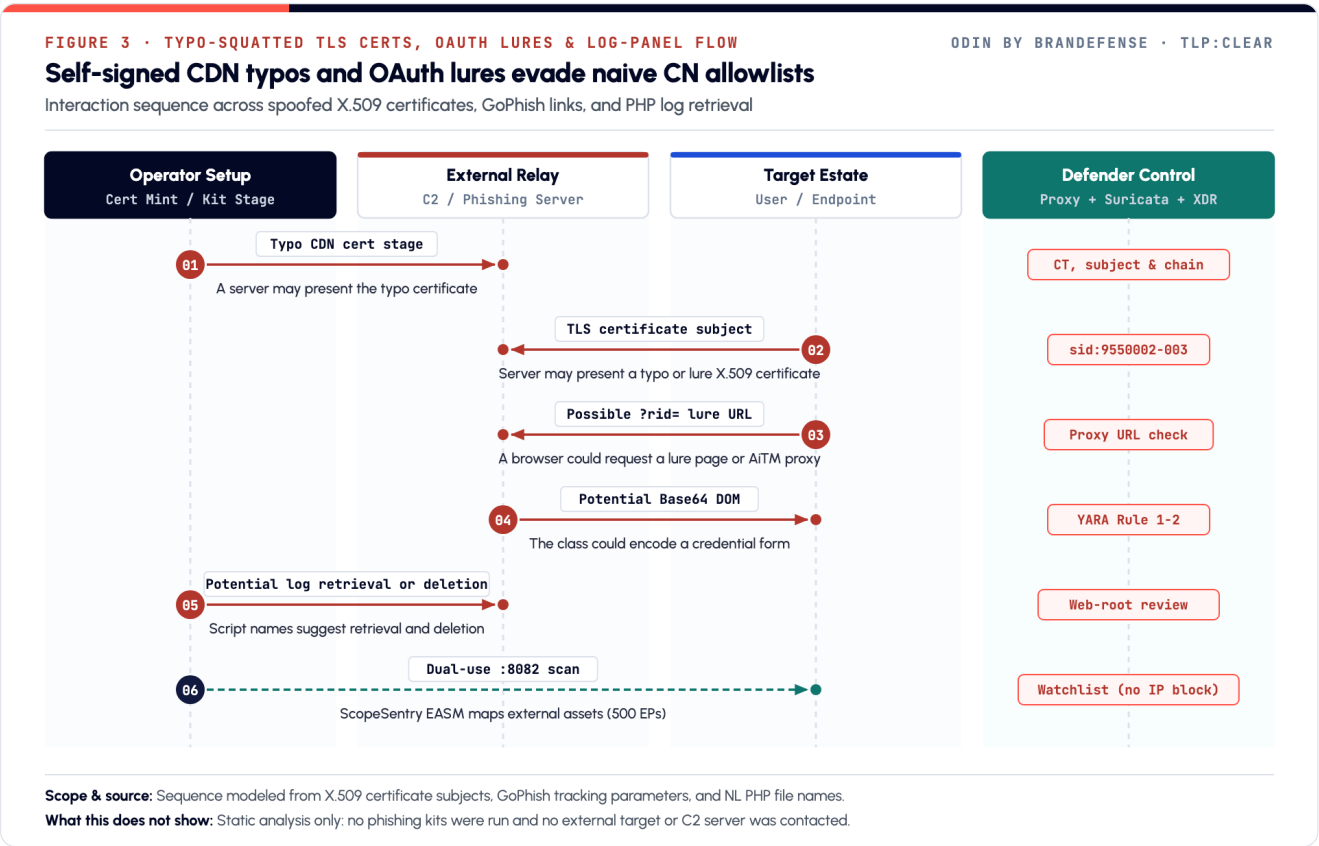


Figure 3: Self-signed CDN typos and OAuth lures evade naive CN allowlists

Nine-class X.509 certificate breakdown (2,002 endpoints)

Every real impersonated organization or unverified custom domain is withheld to protect third parties. Only `CN=cyberstrikeai`, `CN=localhost`, and `cdn[.]cloduflare[.]com` are named.

X.509 certificate class	Endpoints	Hosts	Distinct CNs	ScopeSentry	CyberStrikeAI	Sliver	GoPhish	Evilginx	Upstream-confirmed
Plain HTTP or no Subject CN	1,055	961	1	354	236	272	192	1	10
Custom domain certificate	605	548	433	107	110	110	278	0	16
Framework default (`cyberstrikeai` / `localhost`)	165	157	2	10	90	63	2	0	6
Self-signed literal IP certificate	108	104	1	14	45	45	4	0	0
Default Windows hostname certificate	25	24	22	7	13	5	0	0	0
Generic big-tech portal mimic	19	19	4	8	3	2	6	0	0
Webmail or account-verification lure	13	12	12	0	2	2	9	0	0
Financial or payment brand lure	8	7	7	0	1	0	7	0	0
Enterprise OAuth / SSO lure	4	3	3	0	0	1	3	0	0
Total (9 Classes)	**2,002**	**1,781***	**485**	**500**	**500**	**500**	**501**	**1**	**32 (30 IPs)**

**Multi-port hosts spanning two certificate classes count once per class (1,835 summed class-hosts across 1,781 unique hosts).*

Inside the 44 OAuth, webmail, financial, and portal-mimic certificates

The 44 explicit lure endpoints span 41 unverified hosts and 26 distinct scrubbed certificate subjects:

- **Enterprise OAuth and SSO lures (4 endpoints on 3 hosts, 3 CNs).** One AWS Germany GoPhish host exposes ports 80 and 443 under an SSO subject, and one Hetzner Germany GoPhish host exposes port 3333 under a second OAuth subject. A United States Sliver server on Zillion Network Inc. places a third OAuth subject directly on TCP port 31337.
- **Webmail and account-verification lures (13 endpoints on 12 hosts, 12 CNs).** Nine GoPhish endpoints sit on port 443 (Microsoft US, C2NET Czechia, OVH Poland, WULDI Germany). Others use ports 3333, 80, 28333, and 2176 across Hong Kong, Kazakhstan, Germany, Russia, and Türkiye. Two Sliver listeners also use regional webmail-verification lures on port 31337 (ALEXHOST S.R.L. in Moldova) and port 5061 (OVH SAS in Canada). One dual-port Japanese CyberStrikeAI host (80 and 443) uses another.
- **Financial lures (8 endpoints on 7 hosts) and portal mimics (19 endpoints on 19 hosts).** Seven GoPhish endpoints on ports 3333, 443, 80, and 92 impersonate financial brands across Microsoft US/India, Alibaba

Hong Kong, AWS US, Contabo France, and AI Datacenter Russia. One CyberStrikeAI port 8080 endpoint also impersonates a financial brand. Four reused big-tech portal strings span ScopeSentry (8), GoPhish (6), CyberStrikeAI (3), and Sliver (2).

Sliver (8 upstream-confirmed IPs), GoPhish (19 upstream-confirmed IPs), and Evilginx (1 candidate)

Out of 5,670 global sightings in 83 countries, our 500 examined Sliver endpoints on 385 hosts center on port 443 (128, 25.6%) and port 8888 (101, 20.2%). Multiplayer/mTLS port 31337 (87, 17.4%) and port 8443 (38, 7.6%) follow, led by the United States (111), Hong Kong (72), China (48), and Germany (46). Ninety-nine Sliver endpoints share Sliver's default Go mTLS JARM (`00000000000000000043d43d00043de2a97eabb398317329f027c66e4c1b01`), and 63 serve `CN=localhost`.

Thirty-one Sliver endpoints (26 on port 31337, 3 on 443, 1 on 8090, 1 on 8888) reuse a single custom certificate domain. Three port 31337 servers and five additional servers in Finland, the United Kingdom, the United States, and Moldova carry upstream-confirmed tags. All eight addresses stay withheld pending independent corroboration.

Out of 23,355 global sightings in 119 countries, we examined 501 GoPhish endpoints on 455 hosts. They split between admin port 3333 (192, 38.3%, plus 4 on 3334), port 443 (167, 33.3%), and port 80 (67, 13.4%). Nearly half (223, 44.5%) sit on AWS, including a 62-endpoint AWS France suspected cluster sharing one custom domain on ports 80 (32) and 443 (30).

Forty-eight GoPhish endpoints share family-exclusive JARM `28d28d28d00028d00041d28d28d41dd279b0cf765af27fa62e66d7c8281124` (including 7 upstream-confirmed rows). Overall, 21 endpoints across 19 unique IPv4s carry upstream-confirmed tags on dedicated cloud networks, and all 19 stay withheld. Only one Evilginx candidate (:8080 in Germany) surfaced because Evilginx AiTM proxies (T1557) drop requests lacking valid lure tokens.

05

CHAPTER 05

Ten Exposed PHP and Log Files on a Netherlands Apache Server

IN SHORT

An open Apache HTTP directory in the Netherlands exposes ten files totaling 17,101 bytes from a modular PHP phishing log-panel kit. Because third-party ownership cannot be ruled out, its host address, file paths, and file hashes stay withheld while defenders disable open directory listings and audit web roots.

05 Ten Exposed PHP and Log Files on a Netherlands Apache Server

On 20 November 2026, an open directory was observed on a withheld Dutch Apache server. It listed ten PHP kit files totaling 17,101 bytes. One entry is a withheld log file (87 bytes; filename dated 7 March 2025). We withhold the server address, file paths, and file hashes because third-party web-host upload cannot be ruled out. Figure 4 shows its five structural layers.



Figure 4: Ten PHP and log files expose a modular phishing panel in the Netherlands

Exposed artifact (paths withheld)	Size (bytes)	Functional group	Hashes	What the filename and metadata suggest (contents unverified)
Primary Base64 DOM obfuscator class (path withheld)	3,276	DOM obfuscation	0	Filename suggests a Base64 DOM obfuscation helper class
Core string obfuscator class (path withheld)	2,764	DOM obfuscation	0	Filename suggests a core string or identifier obfuscator class
Database storage class (path withheld)	2,560	Cipher & DB storage	0	Filename suggests a database-access wrapper class
Client-filter utility class (path withheld)	2,457	Client filters	0	Filename suggests a general utility or client-filter class
First operator log script (path withheld)	1,638	Operator log API	0	Filename suggests possible log retrieval; contents not available
Second operator log script (path withheld)	1,638	Operator log API	0	Same size (1,638 B); filename suggests possible log deletion
French-named cipher class (path withheld)	1,024	Cipher & DB storage	0	French-named class suggesting payload or log encryption
Companion Base64 obfuscator class (path withheld)	924	DOM obfuscation	0	Filename suggests a companion Base64 obfuscation class
French-named constants include file (path withheld)	733	Cipher & DB storage	0	French-named include file suggesting configuration constants
a withheld log file (87 bytes; filename dated 7 March 2025)	87	Captured log	0	Flat-file log carrying a 7 March 2025 filename date (withheld)

The two scripts are each 1,638 bytes. Their names suggest possible log retrieval and deletion, but their contents were not available to verify those functions. The remaining seven PHP files (13,738 bytes combined) are likewise grouped by their filenames and recorded sizes because their file contents were not available.

06

CHAPTER 06

Global Feed Sightings Against Upstream-Confirmed Candidates

IN SHORT

Only 32 of 2,002 examined endpoints (1.60%), representing 30 unique IPv4 addresses, carry upstream-confirmed tags on dedicated cloud networks. Withholding those 30 uncorroborated candidates alongside 500 dual-use ScopeSentry consoles, 62 subscriber ISP and CDN edge collisions, and 1,408 other unverified rows avoids self-inflicted outages. Review candidates against local telemetry before blocking.

06 Global Feed Sightings Against Upstream-Confirmed Candidates

Figure 5 compares `log10`-scaled global feed sightings (33,042) and examined counts (2,002) against the 30 upstream-confirmed candidate IPv4 addresses.

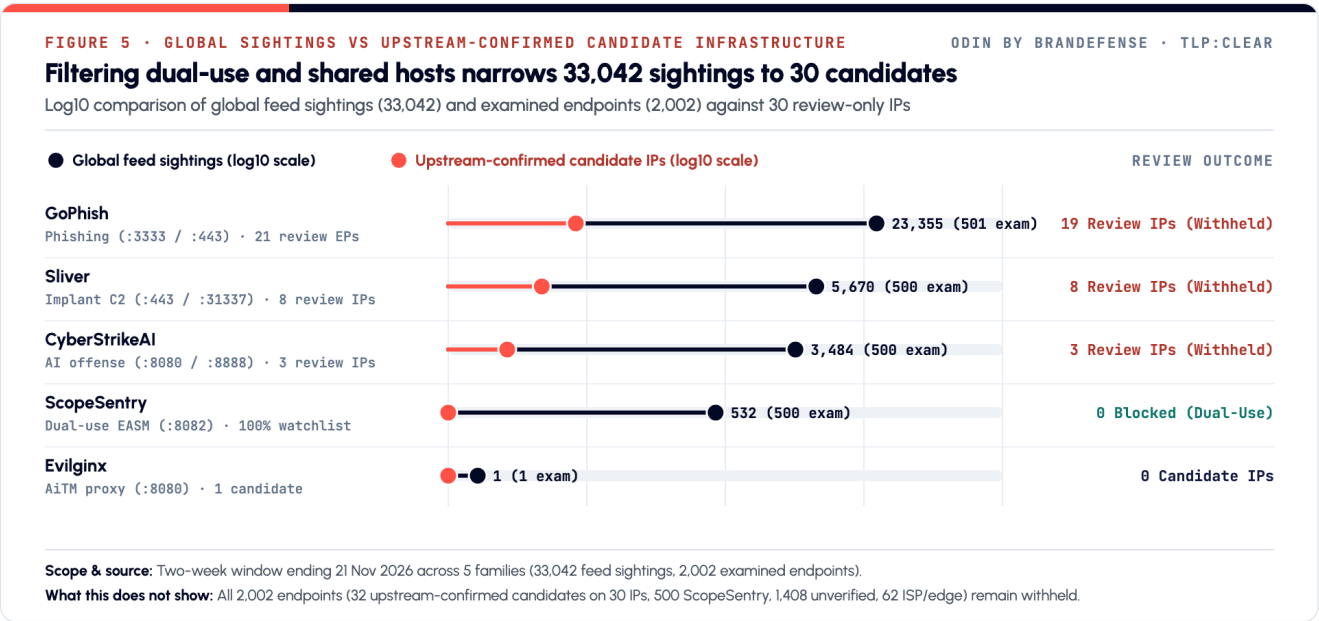


Figure 5: Filtering dual-use and shared hosts narrows 33,042 sightings to 30 candidates

Five plain boundaries govern these findings:

1. **Separate populations, not one actor.** Zero shared addresses or certificates link the 30 upstream-confirmed candidate servers, the 13 two-family watchlist servers, and the Netherlands PHP kit.
2. **ScopeSentry stays off IP blocklists.** All 500 ScopeSentry rows (450 on ports 8082 and 18082) stay withheld because legitimate EASM teams and bug-bounty hunters run the same platform.
3. **Telecom subscriber pools and CDN edges are suppressed.** Sixty-two endpoints on 53 hosts span 27 carrier and CDN networks, including China Telecom, China Mobile, China Unicom, Korea Telecom, Viettel, Charter, and Telefonica. That suppressed set also covers Akamai (7), Cloudflare (2), and Magyar Telekom (4 rows on ports 993, 143, and 3389).
4. **TLS JARMs are hunt-only.** Across 970 JARM-enabled endpoints (220 distinct JARMs), even strong family JARMs identify TLS stacks rather than malicious intent. Meanwhile, Go JARM `40d40d40d00000000043d43d00043d6aff5ab0f4fc31897186312c9e639319` spans 87 endpoints across four families (74 CyberStrikeAI, 9 GoPhish, 3 Sliver, 1 ScopeSentry).
5. **TLS 1.3 encrypts certificates on the wire.** Passive network IDS appliances see `CN=cyberstrikeai` and `CN=cdn[.]cloduflare[.]com` only on TLS 1.2 handshakes, whereas TLS-inspecting forward proxies see both TLS 1.2 and TLS 1.3 certificates.

07

CHAPTER 07

What Defenders Should Do

IN SHORT

This chapter lists the actions defenders should take, in priority order.

07 What Defenders Should Do

Network firewall and web proxy

- [ ] Review the 30 upstream-confirmed address candidates; do not publish or block an address until independent, address-specific corroboration is recorded. Where local logs corroborate an address, review AWS, Azure, and DigitalOcean blocks after 7 days and VPS blocks after 14 days.
- [ ] Search 30 days of firewall, DNS, and proxy logs for connections to TCP port 31337, MCP port 18666, GoPhish `/api/campaigns/` endpoints, or `?rid=` phishing links.

TLS inspection and identity

- [ ] Configure TLS-inspecting proxies to alert on or block self-signed server certificates presenting subjects that claim `CN=cyberstrikeai`, `CN=cdn[.]cloduflare[.]com`, or corporate SSO, webmail, and payment domains.
- [ ] Monitor Certificate Transparency (`CT`) logs for lookalike corporate SSO domains, and enforce FIDO2/WebAuthn passkeys to stop both GoPhish forms and Evilginx AiTM proxies.

Attack-surface consoles and hosted web servers

- [ ] Keep ScopeSentry feeds off IP blocklists. Bind internal :8082 and :18082 consoles behind VPN or mTLS controls, and list approved scanner hostnames in SIEM allowlists.
- [ ] Disable Apache directory listings (`Options -Indexes`) on PHP web servers and review web roots with the heuristic YARA rules in Appendix B.

CHAPTER 08

Detection Rules

IN SHORT

This chapter describes each published rule and where to deploy it.

INSIDE THIS CHAPTER

MITRE ATT&CK mapping 25

08 Detection Rules

The full rules are in Appendix B. All four published rule files passed static syntax checks; they were not loaded on live engines on our build host. Test and tune them before enabling inline blocks.

Rule file	What it catches
`odin_scopesentry_typosquat_certs_and_oauth_lures.rules`	Suricata 7 (4 rules, SIDs 9550001 to 9550003 and 9550006): ScopeSentry :8082/:18082 logins, `CN=cyberstrikeai` and `CN=cdn[.]cloduflare[.]com` server certificates, and GoPhish `/api/campaigns/` requests
`sentinel_scopesentry_typosquat_certs_and_oauth_lure_hunt.kql`	Microsoft Sentinel and Defender XDR: `RemoteUrl contains` matches for `cdn[.]cloduflare[.]com` or `/api/campaigns/`, connections to ports 31337/18666/18082, and unapproved ScopeSentry or CyberStrikeAI processes
`sigma_scopesentry_typosquat_certs_and_phish_panel.yml`	Sigma (1 rule): Linux execution of ScopeSentry, CyberStrikeAI, or OpenSSL certificate commands naming `CN=cyberstrikeai` or `cdn[.]cloduflare[.]com` outside approved hosts
`odin_scopesentry_cyberstrikeai_and_php_phish_kit.yar`	YARA (3 rules): heuristic PHP DOM obfuscation check, heuristic match for at least two listed kit strings including at least one PDO or OpenSSL function string (neither PHP rule validated against the observed files), and CyberStrikeAI or `cdn[.]cloduflare[.]com` configs

In Suricata, SID 9550001 hunts `POST /api/user/login` on ports 8082 and 18082, while SIDs 9550002 and 9550003 match `CN=cyberstrikeai` and `CN=cdn[.]cloduflare[.]com` server certificate subjects. SID 9550006 hunts GoPhish `/api/campaigns/` requests carrying an `Authorization:` header. Keep SIDs 9550001 and 9550006 in alert-only mode where approved internal scanners run, and maintain exact host allowlists (`EASM-SCOPESENTRY-01`, `REDTEAM-SIM-SRV-01`) in KQL and Sigma.

MITRE ATT&CK mapping

Tactic	Technique	How it shows up in this study	Detection coverage
Resource Development	T1587.003 Digital Certificates	`CN=cyberstrikeai` (88 endpoints), `CN=cdn[.]cloduflare[.]com` typo, and 44 OAuth, webmail, bank, and portal certs	SIDs 9550002 and 9550003; Sigma; YARA 3
Resource Development	T1583.003 Virtual Private Server	Upstream-confirmed candidates on AWS, DigitalOcean, Microsoft, OVH, Hetzner, IONOS, and ALEXHOST, plus `/24` sweeps	KQL port hunt; 4 JARMs
Reconnaissance	T1595.002 Vulnerability Scanning	ScopeSentry external attack-surface consoles (446 on port 8082, 4 on port 18082)	SID 9550001; KQL; Sigma
Initial Access	T1566.002 Spearphishing Link	GoPhish servers (192 on port 3333, 167 on port 443, 21 upstream-confirmed) and Netherlands PHP kit	SID 9550006; KQL; YARA 1 and 2 (heuristic)
Credential Access	T1557 Adversary-in-the-Middle	Evilginx candidate (:8080) and enterprise OAuth/SSO certificate lures (4 endpoints)	Proxy X.509 subject and trust-chain checks; FIDO2/WebAuthn

09

CHAPTER 09

Indicators

IN SHORT

This chapter lists the published indicators and explains what stays withheld.

INSIDE THIS CHAPTER

TLS JARM fingerprints to hunt (in Appendix A) 27

Certificate and HTTP patterns (report only) 27

09 Indicators

Appendix A holds 4 published indicators: 0 IPv4 addresses, 4 TLS JARM fingerprints, and 0 file hashes. All 30 upstream-confirmed IPv4 candidates (32 endpoints) and the 1,970 other endpoint rows stay withheld because no independent response, beacon, behaviour, or workbench link is recorded. Five additional certificate and HTTP hunting patterns appear in this report only. The Netherlands Apache server address, its ten file paths, and all ten file hashes stay withheld.

Type	Count	Where to use it
Upstream-confirmed IPv4 candidates (32 endpoints)	0 published (30 withheld)	Withheld pending independent, address-specific corroboration; do not block uncorroborated feed tags
TLS JARM fingerprints	4	Published in Appendix A for TLS hunting only; never block on JARM alone
File hashes (SHA-256)	0	Zero PHP kit file hashes are published because host ownership is not established
Certificate and HTTP patterns	5	Listed below for IDS, proxy, and SIEM hunting

TLS JARM fingerprints to hunt (in Appendix A)

JARM fingerprint	Stack	Context
`00000000000000000043d43d00043de2a97eabb398317329f027c66e4c1b01`	Sliver mTLS	100 endpoints (99 Sliver, including 2 upstream-confirmed port 31337 candidates); hunt only
`40d40d40d00000000043d43d00043d6aff5ab0f4fc31897186312c9e639319`	CyberStrikeAI / Go	87 endpoints (74 CyberStrikeAI, including 2 upstream-confirmed port 8080 candidates); hunt only
`28d28d28d00028d00041d28d28d41dd279b0cf765af27fa62e66d7c8281124`	GoPhish Go TLS	48 GoPhish endpoints (family-exclusive, 7 upstream-confirmed candidates); hunt only
`29d29d00029d29d00041d41d00041dcb3f8752a2d70d1dbf446fc46ab5df96`	GoPhish proxy	57 endpoints (56 GoPhish, 3 upstream-confirmed AWS candidates); hunt only

Certificate and HTTP patterns (report only)

Pattern	Family	Where to hunt
`CN=cyberstrikeai`	CyberStrikeAI	Server X.509 certificate subject (88 endpoints; SID 9550002)
`CN=cdn[.]cloduflare[.]com`	CyberStrikeAI	Typo-squatted Cloudflare X.509 certificate subject (SID 9550003)
`POST /api/user/login`	ScopeSentry	Unapproved :8082 and :18082 web consoles (SID 9550001)
`?rid=`	GoPhish	Outbound browser URLs carrying recipient tracking IDs
`/api/campaigns/`	GoPhish	Campaign management API traffic on TCP 3333 (SID 9550006, KQL)

10

CHAPTER 10

Scope and Limits

IN SHORT

This chapter sets out what the evidence can show and what it cannot.

10 Scope and Limits

This report covers observations from a two-week window ending 21 November 2026 across 2,002 examined endpoints drawn from 33,042 global sightings. It also covers directory metadata from one exposed Apache HTTP directory in the Netherlands (10 files, 17,101 bytes).

Our analysis is strictly static and passive. We did not run any PHP file, authenticate to any external server, or read file or log contents. Nothing in our data confirms whether any downstream organization suffered a breach through the 30 upstream-confirmed candidate addresses, the 44 brand-lure certificates, or the Netherlands Apache server. All 2,002 endpoint addresses, impersonated brand domains, and the Netherlands PHP server address, file paths, and hashes stay withheld.

11

CHAPTER 11

Sources

IN SHORT

This chapter lists the public references used in the report.

11 Sources

1. **MITRE ATT&CK**, Enterprise Sub-Technique T1587.003: Develop Capabilities - Digital Certificates: attack.mitre.org/techniques/T1587/003/ (verified on 21 November 2026).
2. **MITRE ATT&CK**, Enterprise Sub-Technique T1583.003: Acquire Infrastructure - Virtual Private Server: attack.mitre.org/techniques/T1583/003/ (verified on 21 November 2026).
3. **MITRE ATT&CK**, Enterprise Technique T1557: Adversary-in-the-Middle: attack.mitre.org/techniques/T1557/ (verified on 21 November 2026).
4. **MITRE ATT&CK**, Enterprise Sub-Technique T1595.002: Active Scanning - Vulnerability Scanning: attack.mitre.org/techniques/T1595/002/ (verified on 21 November 2026).
5. **MITRE ATT&CK**, Enterprise Sub-Technique T1566.002: Phishing - Spearphishing Link: attack.mitre.org/techniques/T1566/002/ (verified on 21 November 2026).
6. **Internet Engineering Task Force (IETF)**, RFC 5280: Internet X.509 Public Key Infrastructure Certificate and Certificate Revocation List (CRL) Profile: www.rfc-editor.org/rfc/rfc5280 (verified on 21 November 2026).
7. **ScopeSentry Project (Autumn-27)**, ScopeSentry External Attack Surface Management and Recon Automation Architecture: github.com/Autumn-27/ScopeSentry (verified on 21 November 2026).
8. **Salesforce Engineering**, JARM: An Active Transport Layer Security (TLS) Server Fingerprinting Tool: github.com/salesforce/jarm (verified on 21 November 2026).
9. **Gophish Project**, Gophish Official User Guide: Admin Server (TCP 3333) and Phishing Campaign Listener Architecture: Gophish Official User Guide (verified on 21 November 2026).
10. **Bishop Fox**, Sliver Adversary Emulation Framework Documentation: mTLS, Multiplayer (TCP 31337), and HTTPS Listeners: Sliver Official Documentation (verified on 21 November 2026).

APPENDICES

Indicators and Detection Rules

FOR DEFENDERS

Appendix A lists the indicators published for this study. Appendix B carries every published detection rule in full.

INSIDE

A. Indicators of compromise (4) 33

B. Detection rules (4 files) 34

A Indicators of Compromise

Values are defanged for display. The CSV, JSON and STIX files on the online edition carry the raw values. Values the report marks as withheld are not listed.

ID	Type	Indicator	Port	Context	Action
R50-IOC-01	JARM	00000000000000000043d43d00043de2a97eabb398317329f027c66e4c1b01	31337	Sliver default Go mTLS/multiplayer TLS stack · 100 endpoints (99 Sliver including confirmed relays a withheld address and a withheld address)	HUNT
R50-IOC-02	JARM	40d40d40d00000000043d43d00043d6aff5ab0f4fc31897186312c9e639319	8080	CyberStrikeAI and Go TLS listener stack · 87 endpoints (74 CyberStrikeAI including confirmed relays a withheld address and a withheld address)	HUNT
R50-IOC-03	JARM	28d28d28d00028d00041d28d28d41dd279b0cf765af27fa62e66d7c8281124	443	GoPhish exclusive Go TLS stack · 48 GoPhish endpoints including 7 confirmed GoPhish relays	HUNT
R50-IOC-04	JARM	29d29d00029d29d00041d41d00041dcb3f8752a2d70d1dbf446fc46ab5df96	443	GoPhish reverse-proxy TLS stack · 57 endpoints (56 GoPhish including 3 confirmed AWS relays)	HUNT

B Detection Rules

Each rule file loads in its own engine (YARA-X, Suricata 7, Kusto, pySigma). Tune thresholds and allow-lists to your environment before alerting or blocking. The same files can be downloaded from the online edition.

odin_scopesentry_cyberstrikeai_and_php_phish_kit.yar

YARA

```
/*
   Odin by Brandefense - Monograph 50 / 52 YARA Artifact Suite
   Scope: Heuristic detection rules for modular PHP phishing kit files
          (not validated against the observed files, whose paths are withheld),
          and CyberStrikeAI default or typo-squatted certificate configs.
   Classification: TLP:CLEAR
*/

rule ODIN_PHP_Phishing_DOM_Obfuscator_Suite
{
    meta:
        author = "Odin by Brandefense"
        monograph = "50 / 52"
        classification = "TLP:CLEAR"
        action = "HUNT_PHP_PHISH_OBFUSCATOR"
        description = "Heuristic match for PHP DOM obfuscation helper strings; not validated against the observed
files."
    strings:
        $c1 = "<?php" ascii
        $c2 = "class " ascii
        $p1 = "base64_encode" ascii nocase
        $p2 = "base64_decode" ascii nocase
    condition:
        filesize < 128KB and $c1 and $c2 and ($p1 and $p2)
}

rule ODIN_PHP_Phishing_Storage_and_LogPanel_API
{
    meta:
        author = "Odin by Brandefense"
        monograph = "50 / 52"
        classification = "TLP:CLEAR"
        action = "HUNT_PHP_PHISH_LOG_PANEL"
        description = "Heuristic match for at least two listed strings, including at least one of PDO(,
openssl_encrypt, or openssl_decrypt; not validated against the observed files."
    strings:
        $g1 = "<?php" ascii
        $g2 = "unlink(" ascii nocase
        $s1 = "PDO(" ascii
        $s2 = "openssl_encrypt" ascii nocase
        $s3 = "openssl_decrypt" ascii nocase
        $d1 = "Index of /" ascii
        $d2 = "Directory listing for /" ascii
    condition:
        filesize < 128KB and not ($d1 or $d2) and 1 of ($s1, $s2, $s3) and 2 of ($s1, $s2, $s3, $g1, $g2)
}

rule ODIN_CyberStrikeAI_Config_and_Typosquat_Cert_Indicators
{
    meta:
        author = "Odin by Brandefense"
        monograph = "50 / 52"
        classification = "TLP:CLEAR"
        action = "HUNT_CYBERSTRIKEAI_OR_TYPOSQUAT_CERT"
        description = "Matches CyberStrikeAI configuration files or certificate generation scripts naming
CN=cdn.cloduflare.com"
    strings:
        $typo = "cdn.cloduflare.com" ascii nocase
        $cs1 = "cyberstrikeai" ascii nocase
        $m1 = "mcp_server" ascii nocase
```

```
        $m2 = "18666" ascii
        $m3 = "CN=cyberstrikeai" ascii nocase
    condition:
        filesize < 256KB and ($typo or ($cs1 and 1 of ($m1, $m2, $m3)))
}
```

odin_scopesentry_typosquat_certs_and_oauth_lures.rules

SURICATA

```
# Odin by Brandefense - Monograph 50 / 52 Suricata 7.x Detection Suite
# Coverage: ScopeSentry external attack-surface console login APIs (:8082 / :18082),
#           CyberStrikeAI default self-signed TLS server certificates (CN=cyberstrikeai),
#           typo-squatted Cloudflare CDN TLS server certificates (CN=cdn.cloduflare.com),
#           and GoPhish campaign management APIs (/api/campaigns/)
# Classification: TLP:CLEAR

alert http any any -> any [8082,18082] (msg:"ODIN [R50] HUNT ScopeSentry External Attack Surface Console Login
API"; flow:to_server,established; http.method; bsize:4; content:"POST"; http.uri; bsize:>14; content:"/api/user/
login"; fast_pattern; http.request_body; bsize:>16; content:"username"; content:"password"; classtype:attempted-
recon; sid:9550001; rev:1; metadata:created_at 2026_11_21, deployment internal_and_perimeter, policy
behavioral_hunt;)

alert tls $EXTERNAL_NET any -> $HOME_NET any (msg:"ODIN [R50] CyberStrikeAI Default Self-Signed TLS Server
Certificate (CN=cyberstrikeai)"; flow:to_client,established; tls.cert_subject; bsize:>15;
content:"CN=cyberstrikeai"; nocase; fast_pattern; classtype:trojan-activity; sid:9550002; rev:1;
metadata:created_at 2026_11_21, deployment perimeter_and_proxy, tactic command_and_control;)

alert tls $EXTERNAL_NET any -> $HOME_NET any (msg:"ODIN [R50] Typo-Squatted Cloudflare CDN TLS Server Certificate
(CN=cdn.cloduflare.com)"; flow:to_client,established; tls.cert_subject; bsize:>19;
content:"CN=cdn.cloduflare.com"; nocase; fast_pattern; classtype:trojan-activity; sid:9550003; rev:1;
metadata:created_at 2026_11_21, deployment perimeter_and_proxy, tactic resource_development;)

alert http any any -> any any (msg:"ODIN [R50] HUNT GoPhish Campaign Management API Request (/api/campaigns/)";
flow:to_server,established; http.uri; bsize:>14; content:"/api/campaigns/"; fast_pattern; http.header; bsize:>14;
content:"Authorization:"; classtype:attempted-recon; sid:9550006; rev:1; metadata:created_at 2026_11_21,
deployment perimeter_and_proxy, policy behavioral_hunt;)
```

sentinel_scopesentry_typosquat_certs_and_oauth_lure_hunt.kql

KQL (SENTINEL / DEFENDER XDR)

```
// Odin by Brandefense - Monograph 50 / 52 Microsoft Sentinel & Defender XDR Hunt
// Detects behavioral network hits on typo-squatted CDN domains, GoPhish campaign APIs,
// or non-standard C2 and console ports (31337, 18666, 18082), and unapproved ScopeSentry
// or CyberStrikeAI execution outside registered security-testing servers
// Classification: TLP:CLEAR

let Lookback = 7d;
let ApprovedReconHosts = dynamic(["EASM-SCOPESENTRY-01", "REDTEAM-SIM-SRV-01"]);
let BehavioralLureAndPanelNet =
    DeviceNetworkEvents
    | where Timestamp >= ago(Lookback)
    | where not(DeviceName in~ (ApprovedReconHosts))
    | where (RemoteUrl contains "cdn.cloduflare.com" or RemoteUrl contains "/api/campaigns/")
        or RemotePort in (31337, 18666, 18082)
    | project EventTime = Timestamp, DeviceId, DeviceName, InitiatingProcessAccountName,
              InitiatingProcessFileName, InitiatingProcessCommandLine, RemoteIP, RemotePort,
              RemoteUrl, SignalClass = "HUNT_TYPOSQUAT_LURE_OR_C2_PORT";
let ReconAndOffenseProc =
    DeviceProcessEvents
    | where Timestamp >= ago(Lookback)
    | where not(DeviceName in~ (ApprovedReconHosts))
    | where (
        FileName in~ ("python", "python3", "go", "cyberstrikeai", "sliver-server", "gophish")
        and ProcessCommandLine has_any ("ScopeSentry", "cyberstrikeai", "cdn.cloduflare.com")
    )
    | project EventTime = Timestamp, DeviceId, DeviceName,
```

```
                InitiatingProcessAccountName = AccountName,
                InitiatingProcessFileName = FileName,
                InitiatingProcessCommandLine = ProcessCommandLine,
                RemoteIP = "", RemotePort = 0, RemoteUrl = "",
                SignalClass = "HUNT_UNAPPROVED_RECON_OR_OFFENSE_STACK";
union BehavioralLureAndPanelNet, ReconAndOffenseProc
| sort by EventTime desc
```

sigma_scopesentry_typosquat_certs_and_phish_panel.yml

SIGMA

```
title: Unapproved ScopeSentry or CyberStrikeAI Process or Typo-Squat Cert Generation
id: 4b92e1d8-5c37-4a84-9e12-202611215002
status: experimental
description: Detects Linux execution of ScopeSentry, CyberStrikeAI, or OpenSSL self-signed certificate generation
naming CN=cyberstrikeai or CN=cdn.cloduflare.com outside exact approved assessment hosts.
author: Odin by Brandefense
date: 2026/11/21
tags:
  - attack.resource-development
  - attack.t1587.003
  - attack.reconnaissance
  - attack.t1595.002
logsource:
  category: process_creation
  product: linux
detection:
  selection_bins:
    Image|endswith:
      - '/python'
      - '/python3'
      - '/openssl'
      - '/cyberstrikeai'
  selection_args:
    CommandLine|contains:
      - 'ScopeSentry'
      - 'CN=cyberstrikeai'
      - 'cdn.cloduflare.com'
      - '18666'
  filter_approved_hosts:
    ComputerName:
      - 'EASM-SCOPESENTRY-01'
      - 'REDTEAM-SIM-SRV-01'
  condition: selection_bins and selection_args and not filter_approved_hosts
falsepositives:
  - Newly provisioned internal external-attack-surface scanners not yet added to EASM-SCOPESENTRY-01
level: high
```

About Odin Research

Odin Research is the public research series of Odin by Brandefense. It documents adversary infrastructure seen on the internet and turns each finding into indicators and detection rules that defenders can use. Victims and third parties stay anonymous, addresses appear only for confirmed adversary infrastructure, and no working attack code is published.

Odin is the agentic analyst inside the Brandefense threat intelligence platform, which unifies Cyber Threat Intelligence, Digital Risk Protection and External Attack Surface Management.

Online edition, indicators and rules · brandefense.io/odin-research/reports/scopesentry-and-typo-squatted-certificates/
Short Read · brandefense.io/odin-research/blog/calibrating-scopesentry-recon-panels-typo-squatted-cdn-certificates-and-oauth/
Defense Kit · brandefense.io/odin-research/defense-kit/