ODIN

TLP:CLEAR

ODIN RESEARCH · MONOGRAPH 35 / 30 · C2 INFRASTRUCTURE

Open Directories Beside C2 Listeners

Hybrid Servers, Stealer Relays and Agent Leaks

How 60 co-exposed listeners, 200 leaked AI agent records and 24 stealer endpoints split into what to block, what to hunt and what to rotate

Prepared November 2026
Full Report · Public edition · Version 1.1

60 / 44
C2 and callback endpoints co-exposed with an open directory, on 44 hosts (15 frameworks)

200 / 35
Exposed AI coding-agent and chat-connector records, on 35 hosts

49
Commercial AI API keys in those records (DeepSeek, Stability, Anthropic, OpenAI, OpenRouter)

6 / 5
Dedicated-host stealer, loader and RAT endpoints on five hosts without independent corroboration (addresses withheld)

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

How 60 co-exposed listeners, 200 leaked AI agent records and 24 stealer endpoints split into what to block, what to hunt and what to rotate

01	Executive Summary	04
02	Three Groups, One Analytical Model	07
03	Sixty Listeners Next to Open Directories	10
04	AI Coding Agents and Chat Bridges in Open Directories	13
05	Stealer, Loader and RAT Endpoints	16
06	Feed Labels Against Corroborated Evidence	19
07	What Defenders Should Do	22
08	Detection Rules	24
09	Indicators	27
10	Scope and Limits	30
11	Sources	32
A	Indicators of compromise (7)	34
B	Detection rules (4 files)	35

IN SHORT

In a two-week review ending 20 November 2026, Odin by Brandefense found 60 C2 and callback listeners on 44 hosts that also served open web directories. Separately, 35 hosts exposed 200 AI coding-agent and chat-connector records, including 49 commercial AI API keys. The feed held 7,648 records labelled with 11 stealer and loader families. None of the 24 endpoints examined had independent corroboration for blocking. Route feed matches to hunting, close ad-hoc HTTP servers on developer hosts, and rotate exposed agent keys.

01

CHAPTER 01

Executive Summary

IN SHORT

In a two-week review ending 20 November 2026, Odin by Brandefense found 60 C2 and callback listeners on 44 hosts that also served open web directories. Separately, 35 hosts exposed 200 AI coding-agent and chat-connector records, including 49 commercial AI API keys. The feed held 7,648 records labelled with 11 stealer and loader families. None of the 24 endpoints examined had independent corroboration for blocking. Route feed matches to hunting, close ad-hoc HTTP servers on developer hosts, and rotate exposed agent keys.

INSIDE THIS CHAPTER

Three Decisions for Leaders 04

The Study in Numbers 05

01 Executive Summary

Three kinds of infrastructure sit side by side in modern intrusions. Stealers and loaders harvest credentials. Developer and AI-agent workspaces store high-value keys.

Operators run C2 and out-of-band callback servers. This report looks at all three and asks one practical question for each: block, hunt or rotate.

1. **Sixty C2 and callback listeners share a host with an open web directory.**
   On 44 servers, 15 offensive and callback frameworks listen next to an unauthenticated directory listing. Interactsh (14 endpoints), CyberStrikeAI (11), Sliver (7), Octopus (5) and Strix (5) make up 70.0% of them. The directories expose staging files, but many such hosts may be authorized test labs.

2. **Exposed AI coding agents and chat bridges leak 200 session records.**
   On 35 hosts, open directories expose records from five agent tools. Hermes Agent accounts for 93; OpenClaw for 49; Claude Code for 31; Codex CLI for 20; and Cursor Server for 7.

The files hold 49 commercial AI API keys, 18 chat-bot tokens and 72 references to C2 or callback endpoints. We treat the owners as misconfigured users, not attackers.

3. **None of the 24 examined stealer endpoints had independent corroboration for blocking.**
   The 7,648 feed records name 11 stealer, loader and RAT families. Six examined endpoints on five dedicated hosts carry feed status, hosting and TLS metadata for Rhadamanthys, Bofamet, IcedID, NebulaStealer and DarkComet, but lack independent corroboration. The other 18 sit on common service ports, share TLS fingerprints with unrelated servers or belong to residential telecom networks.

Three Decisions for Leaders

1. **Require independent corroboration before blocking feed endpoints.** Because the 24 examined stealer and loader endpoints lack independent second-signal corroboration, withhold them from automatic egress block lists. Send those feed leads, the hybrid hosts and the ISP addresses to hunting, not blocking.
2. **Ban ad-hoc HTTP directory servers on developer, AI-agent and red-team hosts.** A single `python3 -m http.server` in a home folder publishes every hidden agent folder and key file in it.
3. **Audit chat bridges and rotate file-resident AI keys.** List every DingTalk, Telegram and Discord bot connector. Move AI and chat keys out of `.env` files into a vault with short-lived tokens.

The Study in Numbers

Number	What it is
60 / 44	C2 and callback endpoints co-exposed with an open directory, on 44 hosts (15 frameworks)
200 / 35	Exposed AI coding-agent and chat-connector records, on 35 hosts
49	Commercial AI API keys in those records (DeepSeek, Stability, Anthropic, OpenAI, OpenRouter)
7,648 / 24	Feed records labelled with 11 stealer and loader families, and the 24 endpoints we examined on 22 hosts
6 / 5	Dedicated-host stealer, loader and RAT endpoints on five hosts without independent corroboration (addresses withheld)
18 / 17	Stealer and loader endpoints with port collisions, shared JARMs or telecom ISP hosts, on 17 hosts (addresses withheld)
0	Shared hosts or certificates between the stealer endpoints and the hybrid hosts

How to use this report.

Security leaders can read Chapters 1, 6 and 7. Platform and developer-security teams should read Chapters 3 and 4. SOC and detection engineers should start with Chapters 5, 8 and 9. Teams using cyber threat intelligence will find the hosting detail in Chapters 2, 3 and 5.

02

CHAPTER 02

Three Groups, One Analytical Model

IN SHORT

No shared host or certificate links the stealer endpoints to the hybrid hosts. We present the three groups together as an analytical model of how intrusion infrastructure divides work, not as one actor.

INSIDE THIS CHAPTER

Where the servers are hosted 08

02 Three Groups, One Analytical Model

Stealers feed credentials into the ecosystem. Developer and AI-agent hosts hold the keys that make those credentials valuable. C2 and callback servers are where operators act on them. Treating every feed match across these layers as a block-list entry breaks production services and blames misconfigured developers.

We compared hosts and TLS certificates between the stealer and hybrid groups. None of the 24 stealer endpoints shares a host or certificate with the 44 hybrid hosts. So the link between them is functional, not proven. Figure 1 shows the model and what defenders can do with each part.

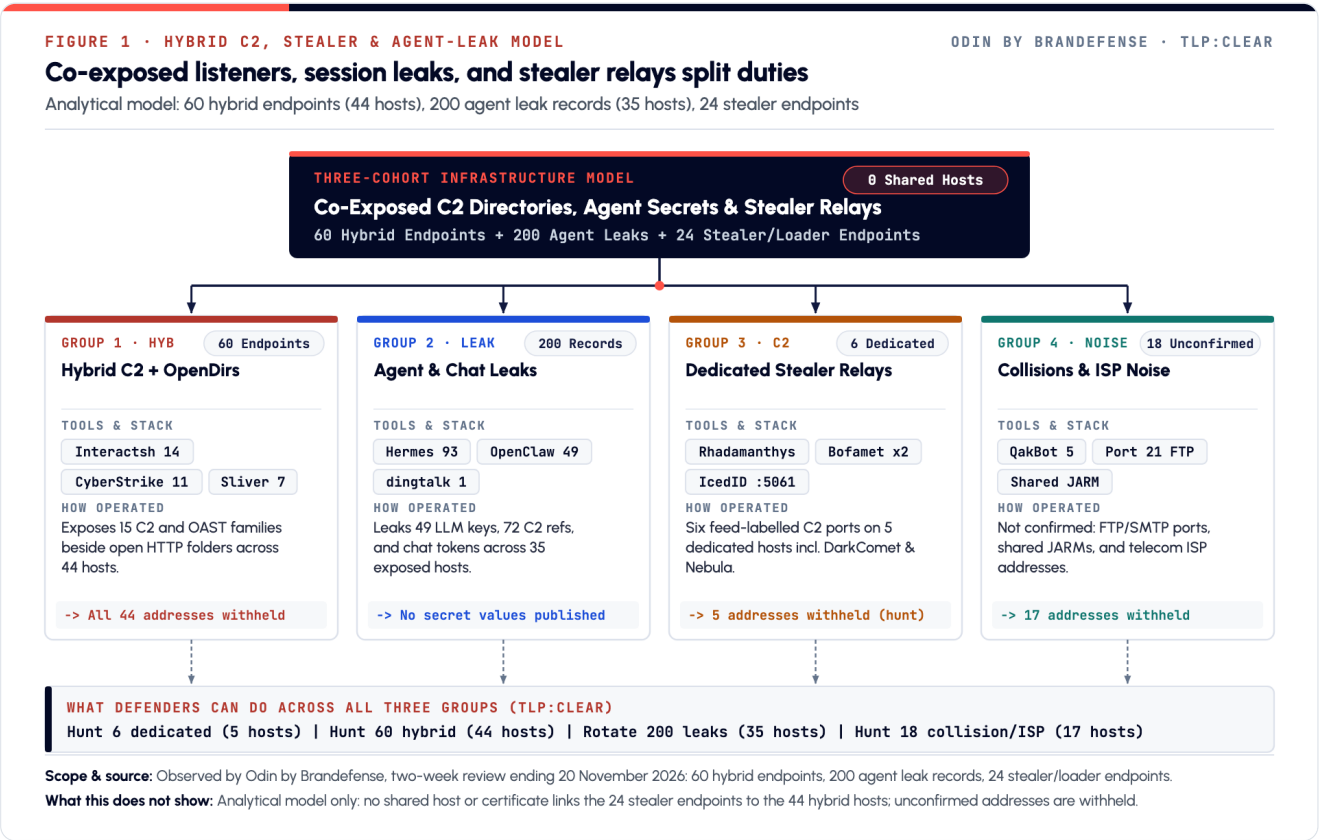


Figure 1: Co-exposed listeners, session leaks and stealer relays split duties (analytical model)

Group	Records	Hosts	Published addresses	What to do
Hybrid C2 and open directory	60 endpoints	44	0	Hunt; check your own hosts
AI coding-agent and chat leaks	200 records	35	0	Rotate keys; close directories
Stealer, loader and RAT endpoints	24 endpoints (from 7,648 feed records)	22	0	Hunt JARM fingerprints and port 5061

Where the servers are hosted

The hybrid listeners sit mostly on commercial cloud and budget VPS networks. The United States (26 endpoints, 43.3%), Germany (11, 18.3%), Finland (6, 10.0%) and China (5, 8.3%) account for 80.0% of them.

Hosting provider	Hybrid endpoints
Hetzner Online GmbH	8
CNServer LLC	6
IBM Cloud	5
IXROOT LTD	5
NetLab Global	5
Hangzhou Alibaba Advertising Co.,Ltd.	4

Six dedicated-host stealer and loader endpoints sit on cloud and European hosting: Amazon.com, Inc., IONOS SE, Yedinyye Resheniya LLC, Brainoza OU and JSC "TIMEWEB". The QakBot-labelled candidate addresses sit on residential and mobile telecom networks in Saudi Arabia and Canada. Because none of the 24 stealer endpoints has independent corroboration and ISP addresses change hands often, none is placed on a block list here.

03

CHAPTER 03

Sixty Listeners Next to Open Directories

IN SHORT

On 44 hosts, C2, callback and AI-offense listeners share an IPv4 address with an unauthenticated directory listing. AI-offense platforms, classic C2 frameworks and callback utilities make up three groups of 20, 23 and 17 endpoints.

INSIDE THIS CHAPTER

AI-offense platforms and masquerade certificates 10

Classic C2 frameworks 11

Callback utilities and port collisions 11

All 15 frameworks 11

03 Sixty Listeners Next to Open Directories

Operators on a fresh VPS often start a quick file server, such as `python3 -m http.server` or `SimpleHTTPServer`, to move payloads, stagers or logs. They skip SCP or a reverse proxy. Bound to `0.0.0.0`, the file server exposes the working folder; defenders can find that risk with external attack surface management (T1583.003).

Python's own documentation warns that `http.server` is not meant for production and has only basic security checks. Figure 2 maps each slice of this research to the target surface and the matching detection.

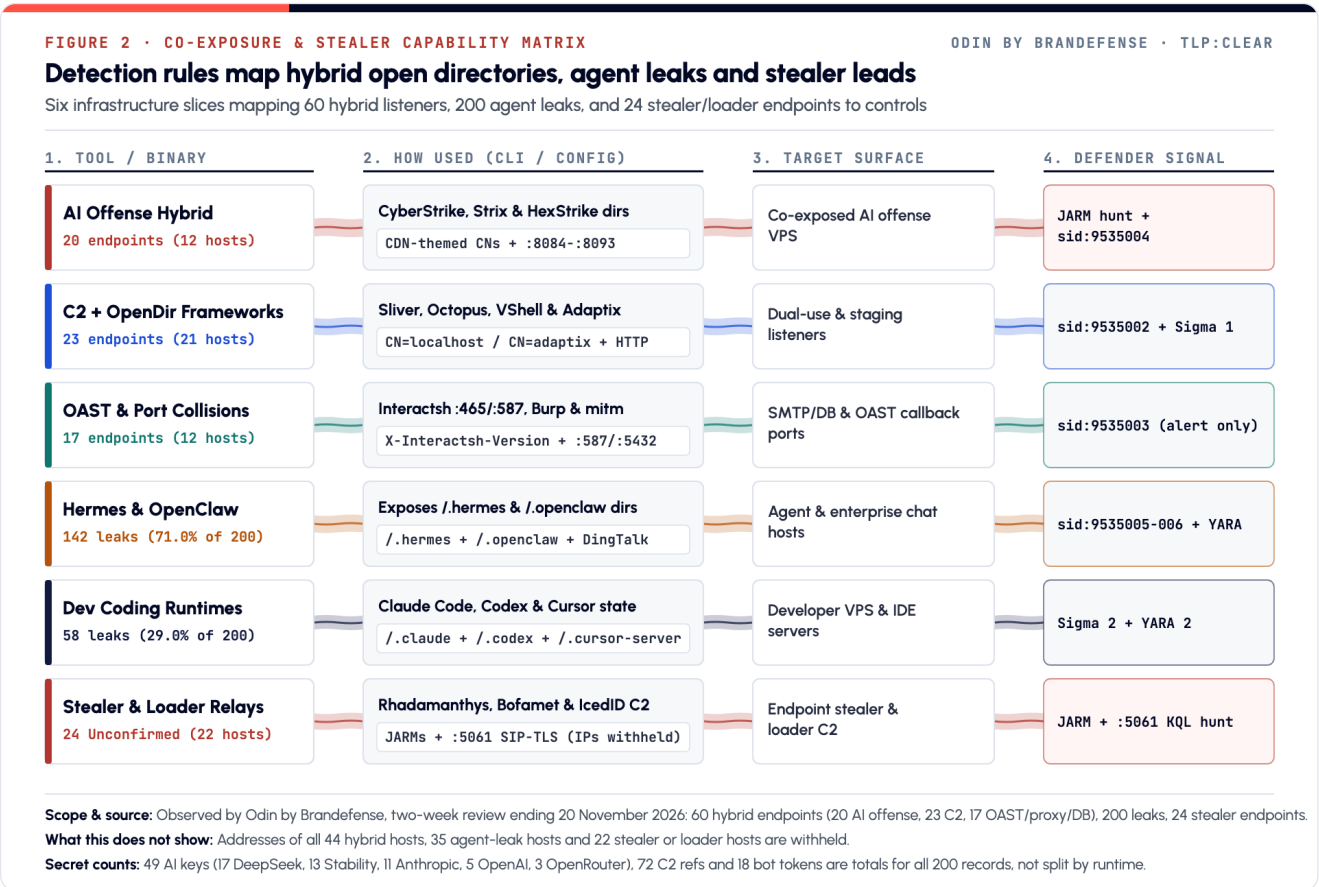


Figure 2: Detection rules map hybrid open directories, agent leaks and stealer leads

AI-offense platforms and masquerade certificates

AI-assisted offense platforms account for 20 of the 60 endpoints (33.3%) on 12 hosts. They are CyberStrikeAI (11 endpoints on 7 hosts), Strix (5 on 1 host) and HexStrikeAI (4 on 4 hosts).

CyberStrikeAI shows two TLS deployment clusters:

- **Dual-port cluster.** Three US hosts on CNServer LLC each listen on TCP 80 and 443 (six endpoints). All six share one certificate and one JARM fingerprint (`40d40d40d00040d…`, full value in Chapter 9).
- **Masquerade cluster.** Three hosts in Malaysia (port 443), the United States (8443) and Japan (8443) share a second JARM. Each presents a certificate name themed as a CDN or health-check service. The three names are withheld because each points to a single server whose owner is not established.

One US host carries all five Strix endpoints on five high ports. HexStrikeAI listens on TCP 5000 on three hosts and on 8888 on one host in Hong Kong.

Classic C2 frameworks

Post-exploitation C2 frameworks contribute 23 endpoints (38.3%). Six of the seven Sliver endpoints, on ports 443 and 8443, still use Sliver's default `CN=localhost` certificate. Three of them share one JARM fingerprint.

Five Octopus HTTP listeners sit on five separate German IBM Cloud hosts. AdaptixC2 presents its default `CN=adaptix` certificate on port 4321.

Callback utilities and port collisions

Callback, interception and database-port listeners make up the last 17 endpoints (28.3%). Thirteen of the 14 Interactsh endpoints sit on mail-submission ports 465 (5) and 587 (8), across eight hosts. One more Interactsh endpoint uses 8443. Another framework label, Daybreak, sits on PostgreSQL port 5432.

ProjectDiscovery's `interactsh-server` opens SMTP listeners by default on port 25, on 587 for SMTPS and on 465 for SMTP with automatic TLS. A port banner rule alone therefore cannot tell a self-hosted callback server from a real mail relay. Use the `X-Interactsh-Version:` HTTP header as a hunt signal instead.

Two US hosts stack more than one framework over an open directory:

Host	Frameworks and ports
Multi-family host 1	Interactsh on 465 and 587, Sliver on 443
Multi-family host 2	Sliver on 8443, Kaiten on 9080

All 15 frameworks

Framework	Endpoints	Hosts	Ports	Notes
Interactsh	14	9	587 (8), 465 (5), 8443 (1)	Mail-port collision on 13 endpoints
CyberStrikeAI	11	7	443 (5), 80 (3), 8443 (2), 8899 (1)	Two TLS clusters
Sliver	7	7	443 (5), 8443 (2)	Six with default `CN=localhost`
Octopus	5	5	80 (5)	IBM Cloud, Germany
Strix	5	1	8084, 8085, 8086, 8089, 8093	One host
HexStrikeAI	4	4	5000 (3), 8888 (1)	
VShell	4	3	8082 (3), 8086 (1)	
AdaptixC2	2	2	4321 (`CN=adaptix`), 4444	
Cobalt Strike	2	2	443 (1), 9091 (1)	
Six others (one each)	6	6	443, 5432, 8080, 8081, 9080	Burp Collaborator, mitmproxy, Daybreak, Kaiten, Havoc, Hashtopolis

All 44 host addresses are withheld. Many of these frameworks are standard red-team and testing tools, and an open directory can sit on a shared lab or a hijacked VPS.

04

CHAPTER 04

AI Coding Agents and Chat Bridges in Open Directories

IN SHORT

Open directories on 35 hosts expose 200 AI coding-agent and chat-connector records. The owners are misconfigured users, and the fix is to close the directory and rotate every key it held.

INSIDE THIS CHAPTER

Five runtimes 13

What the records hold 14

Why a chat bridge matters 14

04 AI Coding Agents and Chat Bridges in Open Directories

Autonomous coding assistants and chat-driven agents keep their state in hidden home folders: `/.hermes/`, `/.openclaw/`, `/.claude/`, `/.codex/` and `/.cursor-server/`. Those folders collect API keys, bot tokens and transcripts. Starting a file server from the home folder publishes all of them at once (T1552.001).

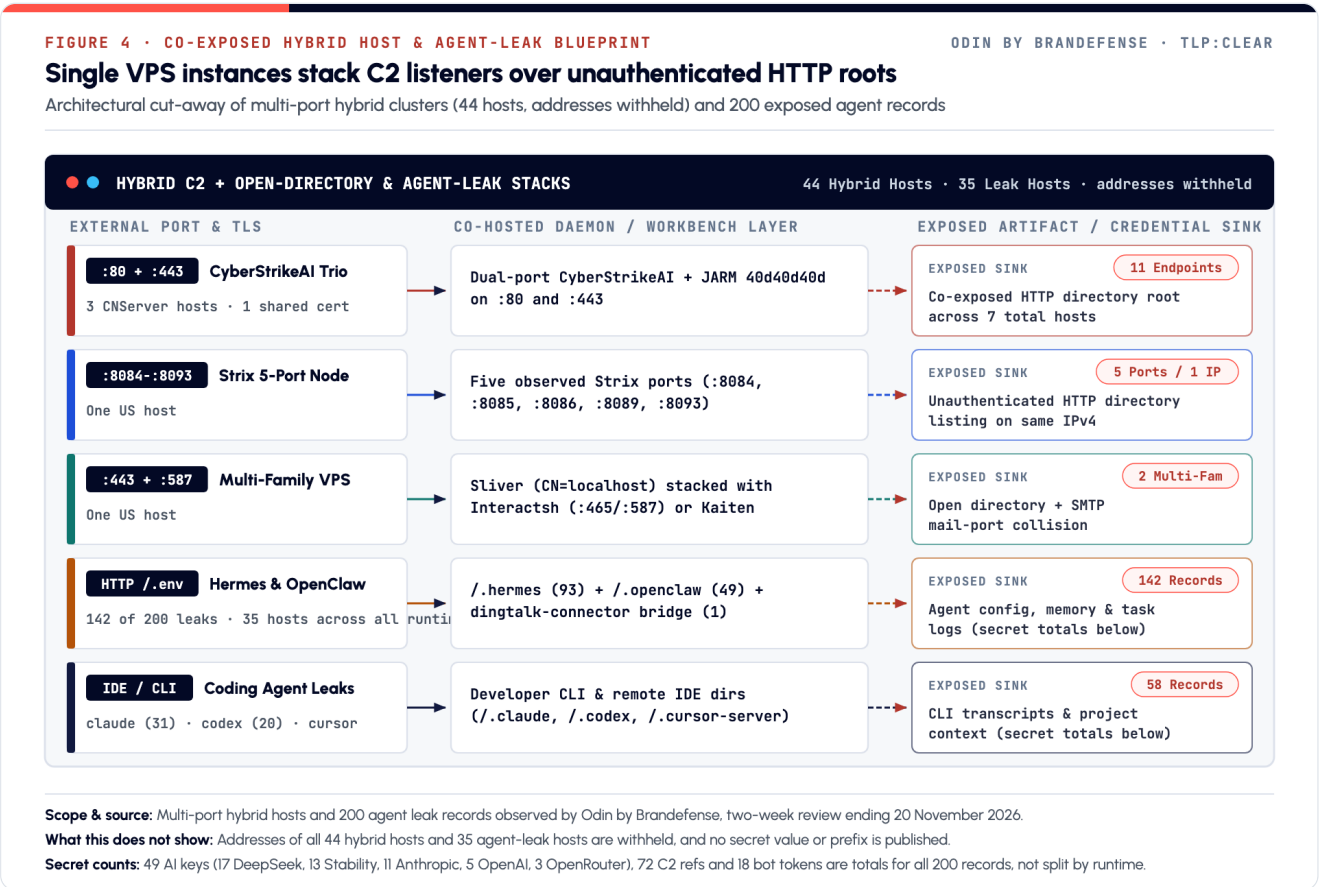


Figure 4: Single VPS instances stack C2 listeners over unauthenticated HTTP roots

Five runtimes

Autonomous task agents outnumber interactive coding tools by about two and a half to one. Hermes Agent and OpenClaw hold 142 records (71.0%). Coding assistants hold the remaining 58 (29.0%).

Runtime	Folder	Records	Share	What leaks
Hermes Agent	`/.hermes/`	93	46.5%	Task state, scratch logs and API key files
OpenClaw	`/.openclaw/`	49	24.5%	Agent sessions, including one `dingtalk-connector` bridge
Claude Code	`/.claude/`	31	15.5%	CLI transcripts and project context
Codex CLI	`/.codex/`	20	10.0%	CLI configuration
Cursor Server	`/.cursor-server/`	7	3.5%	Remote IDE configuration

What the records hold

Content	Records	Share	Detail
C2 or callback endpoint references	72	36.0%	URLs inside agent task logs and scan output
Commercial AI API keys	49	24.5%	17 DeepSeek, 13 Stability, 11 Anthropic, 5 OpenAI, 3 OpenRouter
Passwords, secrets and private keys	39	19.5%	21 passwords, 15 generic secrets, 3 PEM private keys
Chat bots, paste drops and other connectors	40	20.0%	12 Telegram, 6 Discord, 4 Pastebin, 18 other connector records including DingTalk

The 72 endpoint references sit inside agent task logs and scan output. That fits people using Hermes Agent or OpenClaw for security testing on hosts with open directories. No secret value or partial secret appears in this report.

Why a chat bridge matters

A `dingtalk-connector` configuration inside an OpenClaw workspace creates risk in both directions. Anyone who copies a valid bridge credential can reach internal workspace notifications and bot command webhooks without corporate single sign-on. We classify owners of these exposed Claude Code, Codex CLI, Cursor Server, OpenClaw and Hermes Agent servers as misconfigured users. Nothing on them shows offensive intent.

CHAPTER 05

Stealer, Loader and RAT Endpoints

IN SHORT

Feed records name 11 stealer, loader and RAT families across 24 examined endpoints, including six endpoints on five dedicated cloud and VPS hosts. Because the available evidence does not establish independent corroboration for any of the 24 endpoints, withhold them from block lists and hunt their JARM fingerprints and SIP-TLS port 5061 usage.

INSIDE THIS CHAPTER

Six dedicated-host endpoints without independent corroboration 17

Eighteen endpoints with port collisions, shared JARMs or ISP churn 17

05 Stealer, Loader and RAT Endpoints

Information stealers (Rhadamanthys, NebulaStealer, js-tap) and loaders or RATs (QakBot, Bumblebee, IcedID, Bofamet, DarkComet, Ursnif, StadiaRAT, SilverRAT) harvest browser cookies, SSH keys and workstation `.env` tokens. Those are the same secrets that fuel cloud and agent abuse later. The family names come from upstream feeds and are not actor attribution.

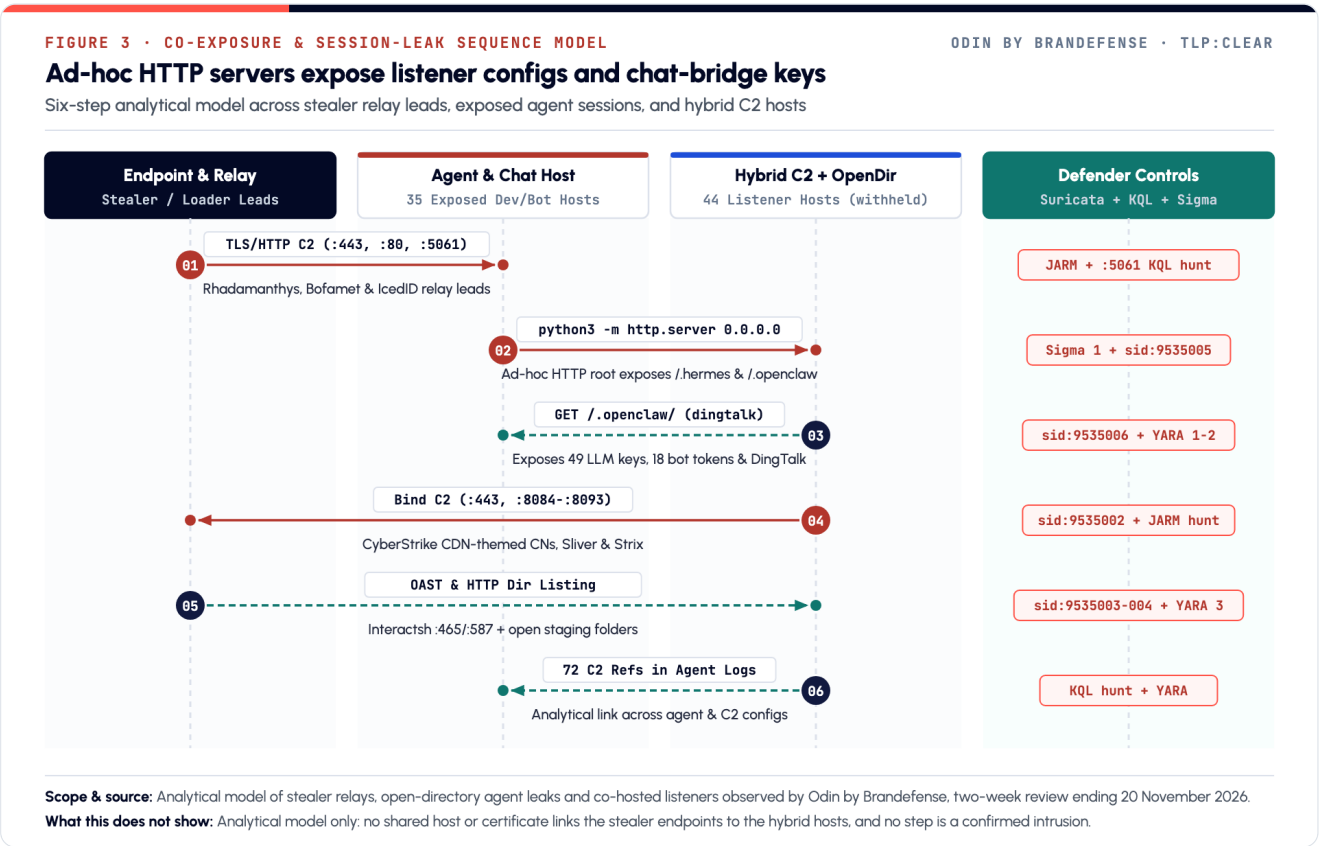


Figure 3: Ad-hoc HTTP servers expose listener configs and chat-bridge keys (analytical model)

Six dedicated-host endpoints without independent corroboration

The available evidence does not establish independent corroboration for these endpoints. The study records provide upstream feed status, dedicated hosting and TLS metadata for six endpoints across five hosts, so their five IPv4 addresses are withheld.

Family	Host and port	Country and provider	TLS
Rhadamanthys	Dedicated host 1, port 443 (address withheld)	Germany, Amazon.com, Inc.	Certificate, JARM `3fd21b20d3fd…`
Bofamet	Dedicated host 2, port 443 (address withheld)	Belarus, Yedinyye Resheniya LLC	Certificate, JARM `29d29d00029d…`
Bofamet	Dedicated host 2, port 80 (address withheld)	Belarus, Yedinyye Resheniya LLC	Plain HTTP
IcedID	Dedicated host 3, port 5061 (address withheld)	Netherlands, Brainoza OU	Certificate, JARM `29d3fd00029d…`
NebulaStealer	Dedicated host 4, port 80 (address withheld)	Germany, IONOS SE	Plain HTTP panel
DarkComet	Dedicated host 5, port 443 (address withheld)	Russia, JSC "TIMEWEB"	Certificate, JARM `27d40d40d000…`

Three points matter for defenders:

- **Rhadamanthys on AWS.** Cloud addresses are reassigned quickly, and an elastic IP without a corroborated beacon or response can move to a new tenant within days. Hunt the TLS JARM fingerprint instead of blocking the host.
- **Bofamet on two ports.** The same Belarus host pairs plain HTTP on port 80 with TLS on 443. Use the TLS JARM fingerprint for hunting until an independent second signal is recorded.
- **IcedID on port 5061.** TCP 5061 is normally reserved for encrypted SIP-TLS voice signalling. Non-voice traffic on this port can pass for VoIP traffic at a glance (T1571). Allow outbound 5061 only from known voice systems.

Eighteen endpoints with port collisions, shared JARMs or ISP churn

A feed's own "confirmed" label is not enough to block an address. Three Rhadamanthys-labelled endpoints carried that label, but two showed direct collisions. One sits on FTP port 21 at a large cloud provider. The other, on port 8443 at a cloud provider in Hong Kong, shares its TLS fingerprint with unrelated servers.

Group	Endpoints	Families and ports
Rhadamanthys label, collision leads	2	FTP port 21; shared JARM on 8443
Suspected C2 leads	5	Bumblebee on 443 (2), QakBot on 5000, DarkComet on 50001, js-tap on 443 sharing the DarkComet endpoint's JARM
Weak leads	11	QakBot on 443 on telecom ISPs (4), DarkComet on 1604 and 7777, Ursnif (2), StadiaRAT (2), SilverRAT on 3467

These 18 endpoints sit on 17 hosts. Together with the five dedicated hosts above, all 22 stealer and loader host addresses are withheld. Use the metadata and JARM fingerprints as hunting context, not as block-list entries.

CHAPTER 06

Feed Labels Against Corroborated Evidence

IN SHORT

Pushing raw feed matches into a firewall would block ISP customers, mail relays and developer machines. None of the 24 endpoints examined from the 7,648-record feed had the independent second signal required for a block list here.

INSIDE THIS CHAPTER

Four other ways to read the evidence 20

06 Feed Labels Against Corroborated Evidence

A block list is defined as much by what it leaves out as by what it holds. Figure 5 compares the size of each feed label or dataset with the examined endpoints and shows why all were routed to hunting or key rotation.

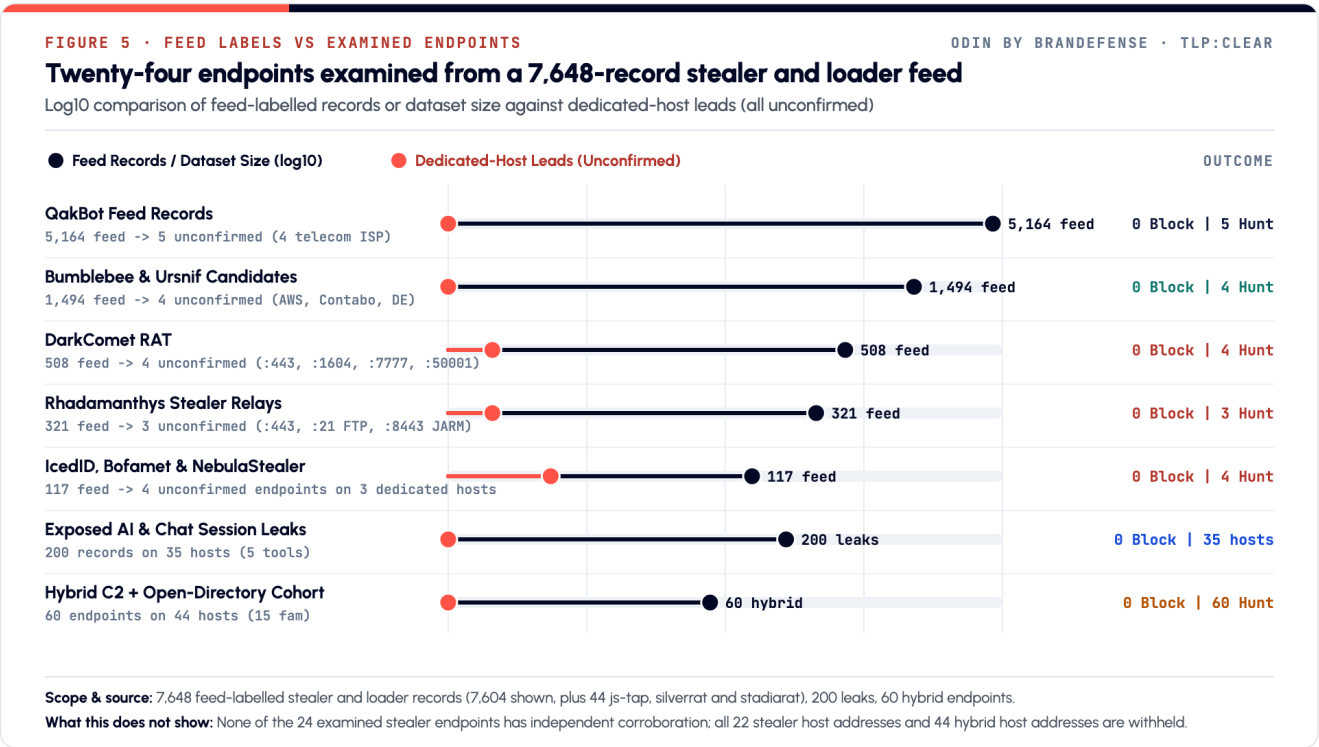


Figure 5: Twenty-four endpoints examined from a 7,648-record stealer and loader feed

Family	Feed-labelled records	Endpoints examined	Hosts	Dedicated-host leads (unconfirmed)	Collision or ISP leads (unconfirmed)
QakBot	5,164	5	5	0	5 (4 telecom ISP, 1 port 5000)
Bumblebee	1,164	2	2	0	2 (port 443)
DarkComet	508	4	3	1 (port 443, address withheld)	3 (ports 1604, 7777, 50001)
Ursnif	330	2	2	0	2 (ports 443, 58089)
Rhadamanthys	321	3	3	1 (port 443, address withheld)	2 (FTP 21, shared JARM on 8443)
IcedID	102	1	1	1 (port 5061, address withheld)	0
js-tap and SilverRAT	39	2	2	0	2 (23 js-tap and 16 SilverRAT records)
Bofamet	12	2	1	2 (ports 443 and 80, address withheld)	0
StadiaRAT	5	2	2	0	2 (port 80, Germany and Netherlands)
NebulaStealer	3	1	1	1 (port 80, address withheld)	0

Four other ways to read the evidence

1. **Authorized testing.** Sliver with `CN=localhost`, Interactsh, Burp Collaborator, mitmproxy, AdaptixC2 and Hashtopolis are everyday tools for internal red teams and contracted testers. Even CyberStrikeAI's masquerade certificates can sit on a shared lab or a hijacked third-party VPS. That is why all 60 hybrid endpoints stay unpublished.
2. **Developers, not operators.** Engineers using Claude Code, Codex CLI, Cursor Server, OpenClaw or Hermes Agent often start a temporary HTTP server to copy files off a cloud devbox. The 35 hosts behind the 200 records are credential-exposure victims. They need key rotation, not blocking.
3. **Residential ISP churn.** QakBot has a history of relaying through infected home computers on residential broadband networks. Of 5,164 QakBot-labelled records, none produced a block-list address. Blocking a subscriber address hurts the next customer who receives it.
4. **Port and fingerprint collisions.** Thirteen Interactsh endpoints on ports 465 and 587, one Daybreak label on PostgreSQL 5432, and two Rhadamanthys labels (FTP 21 and a shared JARM) make 16 collision cases. A JARM or a port on its own should never trigger a block.

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 and firewall

- [ ] Require an independent second signal, such as a protocol response, beacon or workbench artifact, before adding upstream stealer or loader feed addresses to firewall or proxy block lists.
- [ ] Allow outbound TCP 5061 only from registered voice gateways and session border controllers. Alert on any other process that opens it.

Cloud and developer platforms

- [ ] Scan your own external and VPC address space for `Directory listing for /` pages. Close any that expose `/.hermes/`, `/.openclaw/`, `/.claude/`, `/.codex/` or `/.cursor-server/`.
- [ ] Block `http.server` and similar ad-hoc servers bound to `0.0.0.0` on hosts that hold agent folders.

Identity and AI platform security

- [ ] Inventory every chat bridge (`dingtalk-connector`, Telegram, Discord) and every commercial AI key (DeepSeek, Anthropic, OpenAI, OpenRouter, Stability). Move plaintext `.env` secrets into an OS keychain or a vault with short-lived tokens.
- [ ] Rotate any key found in a folder that was reachable over HTTP, even briefly.

Detection and threat intelligence

- [ ] Deploy the rules in Chapter 8 on perimeter, TLS-inspection and Linux or Windows EDR tiers, with exact-name allowlists for red-team and sandbox hosts.
- [ ] Keep FTP (21), SMTP (465, 587), PostgreSQL (5432), shared-JARM and telecom ISP matches out of automatic firewall feeds. Route them to hunting.

CHAPTER 08

Detection Rules

IN SHORT

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

INSIDE THIS CHAPTER

How the Suricata signatures work 24
How the other rules work 24
MITRE ATT&CK mapping 25

08 Detection Rules

The full rules are in Appendix B. The Suricata file passes a syntax lint. The rules were not loaded on live Suricata, YARA, Sigma or Kusto engines for this edition, so test and tune them in your own environment before production use.

Rule file	What it catches
`odin_hybrid_c2_opendir_rhadamanthys_and_agent_leaks.rules`	Suricata 7: five signatures for the AdaptixC2 default certificate, Interactsh headers, co-exposed C2 and agent-folder directory listings, and OpenClaw config fetches
`sentinel_hybrid_c2_opendir_stealer_and_agent_leak_hunt.kql`	Microsoft Sentinel and Defender XDR: non-voice use of port 5061, HTTP requests to agent folders, and ad-hoc HTTP servers next to C2 or agent folders
`sigma_hybrid_c2_http_server_rhadamanthys_and_agent_env.yml`	Sigma (2 rules): ad-hoc HTTP servers in agent or C2 folders, and archiving or copying of agent folders
`odin_rhadamanthys_bofamet_and_exposed_agent_sessions.yar`	YARA (3 rules): exposed OpenClaw and DingTalk sessions, agent key bundles, and hybrid open-directory listings

How the Suricata signatures work

SID	Looks for	Why
9535002	Server certificate subject `CN=adaptix`	AdaptixC2's default self-signed certificate
9535003	`X-Interactsh-Version:` response header	Self-hosted Interactsh server (hunt only, authorized scanners use it too)
9535004	A 200 directory listing that names CyberStrike, HexStrike, Sliver, Adaptix, Strix, VShell or Octopus	C2 or AI-offense tooling in an open directory
9535005	A 200 directory listing with links to `.hermes`, `.openclaw`, `.claude`, `.codex` or `.cursor-server` folders	Agent home folders exposed over HTTP
9535006	`GET` requests under `/.openclaw/` for `dingtalk-connector`, config, `.env` or session files	Someone fetching exposed agent or bridge secrets

TLS 1.3 encrypts server certificates, so SID 9535002 fires on TLS 1.2 handshakes or behind a TLS-decrypting proxy. Keep SID 9535003 in alert-only mode where authorized scanners run.

How the other rules work

The KQL query looks back seven days. It first finds Python, PHP, Node, Caddy or nginx processes started as ad-hoc file servers, or command lines that name agent folders or C2 frameworks. It then finds network connections to port 5061 from non-voice hosts, or to URLs under agent folders. The two views are joined per device.

Each result gets one label: contain an exposed directory and rotate keys, hunt a DingTalk bridge, or hunt a listener or port 5061 anomaly. Red-team, AI-gateway and voice hosts are excluded by exact device name.

Sigma rule	Level	Fires when
Ad-hoc HTTP server in a C2 or agent workspace	High	`http.server`, `SimpleHTTPServer` or similar starts from an agent or C2 folder on Linux
Archive or transfer of agent folders	High	`tar`, `zip`, `curl`, `scp` or `rsync` reads an agent or `dingtalk-connector` path on Linux

YARA rule	Fires when
`ODIN_Exposed_OpenClaw_DingTalk_Connector_Session`	An OpenClaw marker appears with a DingTalk connector, or with a bot token and a key field
`ODIN_Exposed_Hermes_Claude_Codex_Agent_Secret_Bundle`	An agent folder path appears with at least two AI API key names
`ODIN_Hybrid_C2_OpenDir_Workspace_Manifest`	A directory listing names at least two C2, callback or AI-offense frameworks

MITRE ATT&CK mapping

Tactic	Technique	How it shows up	Detection
Resource Development	T1583.003 Acquire Infrastructure: Virtual Private Server	C2 and callback listeners on rented VPS hosts with open directories	SID 9535004; Sigma rule 1; YARA
Credential Access	T1552.001 Unsecured Credentials: Credentials In Files	AI and chat keys in exposed agent folders	SIDs 9535005 and 9535006; Sigma rule 2; YARA
Command and Control	T1071.001 Application Layer Protocol: Web Protocols	C2 and callback listeners over HTTP and HTTPS	SIDs 9535002 and 9535003; KQL
Command and Control	T1571 Non-Standard Port	IcedID-labelled endpoint on SIP-TLS port 5061	KQL port 5061 hunt

09

CHAPTER 09

Indicators

IN SHORT

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

INSIDE THIS CHAPTER

JARM fingerprints to hunt 27

Other indicators 27

09 Indicators

Appendix A holds 7 published indicators: 7 JARM fingerprints for TLS hunting. There are no published IPv4 addresses or file hashes in this report. Eight more indicators appear in this report only: one default certificate name, one HTTP header, five agent folder paths and one chat-bridge name.

Type	Count	Where to use it
JARM fingerprints	7	TLS hunting only, never block on a JARM alone
Certificate name and HTTP header	2	TLS and proxy hunting
Agent folder paths and chat-bridge name	6	External attack-surface scans, EDR and IAM reviews

JARM fingerprints to hunt

JARM	Family	Context
`3fd21b20d3fd3fd21c43d21b21b43d1ec49a4b64df0a9e9f328abd60285841`	Rhadamanthys	Dedicated-host endpoint; also seen on another server
`29d29d00029d29d00042d42d0000002059a3b916699461c5923779b77cf06b`	Bofamet	Dedicated-host loader endpoint
`29d3fd00029d29d00042d43d000000ad9bf51cc3f5a1e29eecb81d0c7b06eb`	IcedID	Dedicated-host port 5061 endpoint
`27d40d40d00040d00042d43d000000d2e61cae37a985f75ecafb81b33ca523`	DarkComet	Dedicated-host endpoint; shared with a suspected js-tap server
`40d40d40d00040d00043d43d00043d3bef2bf79cd6719851e8198c1e8f9a14`	CyberStrikeAI	Three dual-port hybrid hosts
`40d40d40d00000000043d43d00043d6aff5ab0f4fc31897186312c9e639319`	CyberStrikeAI	Three hosts with CDN-themed certificate names
`3fd3fd20d00000021c43d43d00043d204204071741c36579e355f830d285a5`	Sliver	Three hybrid Sliver listeners

Other indicators

Indicator	Type	Use
`CN=adaptix`	Default certificate name (AdaptixC2, port 4321)	TLS hunting
`X-Interactsh-Version:`	HTTP response header (14 endpoints on 9 hosts)	Proxy hunting, alert only
`/.hermes/`	Agent folder (93 records)	Attack-surface scans and EDR
`/.openclaw/`	Agent folder (49 records)	Attack-surface scans and EDR
`/.claude/`	Agent folder (31 records)	Developer host audit
`/.codex/`	Agent folder (20 records)	Developer host audit
`/.cursor-server/`	Agent folder (7 records)	Developer host audit
`dingtalk-connector`	Chat-bridge configuration in OpenClaw	IAM and EDR review

Withheld: all 44 hybrid host addresses and all 35 agent-leak host addresses. Also withheld: all 22 hosts behind the 24 stealer endpoints (the five dedicated hosts and the 17 collision or ISP hosts). The three CDN-themed certificate names and every secret value are withheld too.

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

- **Window.** The review ended 20 November 2026. The stealer comparison includes older feed records; their dates do not establish activity within a two-week window.
- **Feed labels are not attribution.** Family names such as Rhadamanthys or QakBot come from upstream feeds. Without an independent second signal, none of the 24 stealer endpoints is treated as confirmed adversary infrastructure.
- **The model is a model.** No host or certificate links the stealer endpoints to the hybrid hosts. Figures 1 and 3 show an analytical model of how the groups divide work, not one actor.
- **Owners are not operators.** The 35 agent-leak hosts belong to misconfigured users. Hybrid hosts may be authorized labs. Their addresses are withheld.
- **No secrets used.** We did not use any leaked key or token, and no secret value appears in this report.
- **Cloud addresses move.** One Rhadamanthys-labelled endpoint sits on AWS, where elastic addresses are reassigned often. All 22 stealer host addresses are withheld rather than placed on a block list.
- **Rule testing.** The rules were not loaded on live engines for this edition and were not replayed against live traffic. Test them in your own environment first.

11

CHAPTER 11

Sources

IN SHORT

This chapter lists the public references used in the report.

11 Sources

1. MITRE ATT&CK. T1583.003 Acquire Infrastructure: Virtual Private Server. https://attack.mitre.org/techniques/T1583/003/ (verified 20 November 2026)
2. ProjectDiscovery. Interactsh: out-of-band interaction server and client. https://github.com/projectdiscovery/interactsh (verified 20 November 2026)
3. MITRE ATT&CK. S0633 Sliver. https://attack.mitre.org/software/S0633/ (verified 20 November 2026)
4. MITRE ATT&CK. T1552.001 Unsecured Credentials: Credentials In Files. https://attack.mitre.org/techniques/T1552/001/ (verified 20 November 2026)
5. OWASP Cheat Sheet Series. Secrets Management Cheat Sheet. https://cheatsheetseries.owasp.org/cheatsheets/Secrets_Management_Cheat_Sheet.html (verified 20 November 2026)
6. MITRE ATT&CK. S0650 QakBot. https://attack.mitre.org/software/S0650/ (verified 20 November 2026)
7. MITRE ATT&CK. S0483 IcedID. https://attack.mitre.org/software/S0483/ (verified 20 November 2026)
8. Salesforce. JARM: active TLS server fingerprinting. https://github.com/salesforce/jarm (verified 20 November 2026)
9. MITRE ATT&CK. T1571 Non-Standard Port. https://attack.mitre.org/techniques/T1571/ (verified 20 November 2026)
10. Python documentation. http.server: HTTP servers. https://docs.python.org/3/library/http.server.html (verified 20 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 (7) 34

B. Detection rules (4 files) 35

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
R35-IOC-01	JARM	3fd21b20d3fd3fd21c43d21b21b43d1ec49a4b64df0a9e9f328abd60285841	443	rhadamanthys · TLS fingerprint of the dedicated-host endpoint, also seen on another server	HUNT
R35-IOC-02	JARM	29d29d00029d29d00042d42d0000002059a3b916699461c5923779b77cf06b	443	bofamet · TLS fingerprint of the dedicated-host loader endpoint	HUNT
R35-IOC-03	JARM	29d3fd00029d29d00042d43d000000ad9bf51cc3f5a1e29eecb81d0c7b06eb	5061	icedid · TLS fingerprint of the dedicated-host SIP-TLS port 5061 endpoint	HUNT
R35-IOC-04	JARM	27d40d40d00040d00042d43d000000d2e61cae37a985f75ecafb81b33ca523	443	darkcomet · TLS fingerprint shared by the dedicated-host DarkComet endpoint and a suspected js-tap server	HUNT
R35-IOC-05	JARM	40d40d40d00040d00043d43d00043d3bef2bf79cd6719851e8198c1e8f9a14	443	cyberstrikeai · shared by three dual-port hybrid hosts	HUNT
R35-IOC-06	JARM	40d40d40d00000000043d43d00043d6aff5ab0f4fc31897186312c9e639319	8443	cyberstrikeai · shared by three hosts with CDN-themed certificate names	HUNT
R35-IOC-07	JARM	3fd3fd20d00000021c43d43d00043d204204071741c36579e355f830d285a5	443	sliver · shared by three hybrid Sliver listeners	HUNT

B Detection Rules

Each rule file loads in its own engine (Suricata 7, YARA-X, 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_hybrid_c2_opendir_rhadamanthys_and_agent_leaks.rules

SURICATA

```
# Odin by Brandefense - Monograph 35 / 35 Suricata 7.x Detection Suite (public edition v1.1)
# Coverage: AdaptixC2 default TLS server certificate, self-hosted Interactsh OAST headers,
#           co-exposed C2 and AI coding-agent open-directory listings, and OpenClaw/dingtalk-connector
#           session fetches
# Classification: TLP:CLEAR

alert tls $EXTERNAL_NET any -> $HOME_NET any (msg:"ODIN [R35] AdaptixC2 Default Self-Signed TLS Server
Certificate (CN=adaptix)"; flow:established,to_client; tls.cert_subject; bsize:>9; content:"CN=adaptix";
fast_pattern; classtype:trojan-activity; sid:9535002; rev:1; metadata:created_at 2026_11_20, deployment
perimeter, tactic command_and_control;)

alert http $EXTERNAL_NET any -> $HOME_NET any (msg:"ODIN [R35] HUNT Self-Hosted Interactsh OAST Response Header";
flow:established,to_client; http.header; bsize:>22; content:"X-Interactsh-Version:"; fast_pattern;
classtype:policy-violation; sid:9535003; rev:2; metadata:created_at 2026_11_20, updated_at 2026_11_20, deployment
perimeter_and_proxy, policy behavioral_hunt;)

alert http $EXTERNAL_NET any -> $HOME_NET any (msg:"ODIN [R35] Co-Exposed Hybrid C2 or AI-Offense Open-Directory
Listing"; flow:established,to_client; http.stat_code; bsize:3; content:"200"; http.response_body; bsize:>48;
content:"<title>Directory listing for "; fast_pattern; pcre:"/(cyberstrike|hexstrike|sliver|adaptix|strix|vshell|
octopus)/i"; classtype:attempted-recon; sid:9535004; rev:1; metadata:created_at 2026_11_20, deployment perimeter,
tactic resource_development;)

alert http any any -> any any (msg:"ODIN [R35] Unauthenticated HTTP Directory Listing Exposing AI Agent Runtime
Folder"; flow:established,to_client; http.stat_code; bsize:3; content:"200"; http.response_body; bsize:>48;
content:"Directory listing for /"; fast_pattern; pcre:"/href=\x22\.(hermes|openclaw|claude|codex|cursor-server)
\//i"; classtype:policy-violation; sid:9535005; rev:1; metadata:created_at 2026_11_20, deployment
perimeter_and_east_west, tactic credential_access;)

alert http any any -> any any (msg:"ODIN [R35] HTTP Retrieval of Exposed OpenClaw Session or dingtalk-connector
Config"; flow:established,to_server; http.method; bsize:3; content:"GET"; http.uri; bsize:>11;
content:"/.openclaw/"; fast_pattern; pcre:"/\/\.openclaw\/.*(dingtalk-connector|config|\.env|session)/i";
classtype:credential-theft; sid:9535006; rev:1; metadata:created_at 2026_11_20, deployment perimeter_and_waf,
tactic credential_access;)
```

odin_rhadamanthys_bofamet_and_exposed_agent_sessions.yar

YARA

```
/*
   Odin by Brandefense - Monograph 35 / 35 YARA Artifact Suite (public edition v1.1)
   Scope: Structural detection rules for exposed OpenClaw / dingtalk-connector chat bridges,
          Hermes / Claude Code / Codex CLI / Cursor Server API secret bundles, and co-exposed
          hybrid C2 open-directory listings.
   Classification: TLP:CLEAR
*/

rule ODIN_Exposed_OpenClaw_DingTalk_Connector_Session
{
    meta:
        author = "Odin by Brandefense"
        monograph = "35 / 35"
        classification = "TLP:CLEAR"
        action = "HUNT_EXPOSED_CHAT_AGENT_BRIDGE"
        description = "Matches exposed OpenClaw agent configuration or session state containing enterprise
dingtalk-connector or bot bridge keys"
    strings:
        $oc1 = "openclaw" ascii nocase
        $oc2 = "/.openclaw/" ascii
```

```
        $br1 = "dingtalk-connector" ascii nocase
        $br2 = "telegram_bot_token" ascii nocase
        $br3 = "discord_bot_token" ascii nocase
        $sec1 = "appSecret" ascii
        $sec2 = "API_KEY" ascii nocase
    condition:
        filesize < 256KB and ($oc1 or $oc2) and ($br1 or (($br2 or $br3) and ($sec1 or $sec2)))
}

rule ODIN_Exposed_Hermes_Claude_Codex_Agent_Secret_Bundle
{
    meta:
        author = "Odin by Brandefense"
        monograph = "35 / 35"
        classification = "TLP:CLEAR"
        action = "HUNT_EXPOSED_AGENT_RUNTIME"
        description = "Matches exposed AI coding-agent runtime files (.hermes, .claude, .codex, .cursor-server)
paired with LLM API key identifiers"
    strings:
        $rt1 = "/.hermes/" ascii
        $rt2 = "hermes_agent" ascii
        $rt3 = "/.claude/" ascii
        $rt4 = "/.codex/" ascii
        $rt5 = "/.cursor-server/" ascii
        $k1 = "DEEPSEEK_API_KEY" ascii nocase
        $k2 = "ANTHROPIC_API_KEY" ascii nocase
        $k3 = "OPENAI_API_KEY" ascii nocase
        $k4 = "OPENROUTER_API_KEY" ascii nocase
        $k5 = "STABILITY_KEY" ascii nocase
    condition:
        filesize < 512KB and 1 of ($rt*) and 2 of ($k*)
}

rule ODIN_Hybrid_C2_OpenDir_Workspace_Manifest
{
    meta:
        author = "Odin by Brandefense"
        monograph = "35 / 35"
        classification = "TLP:CLEAR"
        action = "HUNT_HYBRID_C2_OPENDIR"
        description = "Matches unauthenticated HTTP directory listings co-exposing C2, OAST, or AI-offense
framework directories"
    strings:
        $dir1 = "Directory listing for /" ascii
        $dir2 = "<title>Directory listing for " ascii
        $fw1 = "cyberstrike" ascii nocase
        $fw2 = "hexstrike" ascii nocase
        $fw3 = "adaptix" ascii nocase
        $fw4 = "interactsh" ascii nocase
        $fw5 = "sliver" ascii nocase
        $fw6 = "strix" ascii nocase
    condition:
        filesize < 256KB and ($dir1 or $dir2) and 2 of ($fw*)
}
```

sentinel_hybrid_c2_opendir_stealer_and_agent_leak_hunt.kql

KQL (SENTINEL / DEFENDER XDR)

```
// Odin by Brandefense - Monograph 35 / 35 Microsoft Sentinel & Defender XDR Hunt (public edition v1.1)
// Correlates SIP-TLS TCP/5061 port camouflage, ad-hoc HTTP directory servers co-hosted with C2 or AI-agent
// state (/.hermes, /.openclaw, /.claude, /.codex, /.cursor-server), and dingtalk-connector reads.
// Host allowlists hold exact device names only.
// Classification: TLP:CLEAR

let Lookback = 7d;
let ApprovedRedTeamHosts = dynamic(["REDTEAM-SIM-VPS-01", "PURPLE-OAST-NODE-01"]);
let ApprovedDevGatewayHosts = dynamic(["AI-SANDBOX-GW-01"]);
let ApprovedSipHosts = dynamic(["CORP-SIP-SBC-01"]);
let AgentAndHybridMarkers = dynamic([".hermes", ".openclaw", "dingtalk-connector", ".claude", ".codex", ".cursor-
```

```
server", "cyberstrike", "hexstrike", "adaptix", "interactsh", "sliver", "strix", "vshell"]);
let CoExposedHostProcs = DeviceProcessEvents
    | where Timestamp >= ago(Lookback)
    | where not(DeviceName in~ (ApprovedRedTeamHosts)) and not(DeviceName in~ (ApprovedDevGatewayHosts))
    | where (FileName in~ ("python", "python3", "php", "node", "caddy", "nginx") and ProcessCommandLine has_any
("http.server", "SimpleHTTPServer", "0.0.0.0:", "-S 0.0.0.0"))
        or ProcessCommandLine has_any (AgentAndHybridMarkers)
    | project ProcTime = Timestamp, DeviceId, DeviceName, InitiatingProcessFileName, FileName, FolderPath,
ProcessCommandLine;
DeviceNetworkEvents
| where Timestamp >= ago(Lookback)
| where not(DeviceName in~ (ApprovedRedTeamHosts)) and not(DeviceName in~ (ApprovedDevGatewayHosts))
    and ((RemotePort == 5061 and not(DeviceName in~ (ApprovedSipHosts)))
         or RemoteUrl has_any ("/.hermes/", "/.openclaw/", "dingtalk-connector", "/.claude/", "/.codex/",
"/.cursor-server/"))
| summarize
    NetFirstSeen = min(Timestamp),
    NetLastSeen = max(Timestamp),
    DistinctRemoteIps = dcount(RemoteIP),
    HasPort5061 = countif(RemotePort == 5061) > 0,
    SampleEndpoints = make_set(strcat(RemoteIP, ":", tostring(RemotePort)), 12),
    SampleUrls = make_set(RemoteUrl, 10),
    ConnectionCount = count()
    by DeviceId, DeviceName, InitiatingProcessFileName
| join kind=fullouter (CoExposedHostProcs) on DeviceId, DeviceName
| extend ResolvedDevice = coalesce(DeviceName, DeviceName1)
| where isnotempty(ProcessCommandLine) or HasPort5061 or isnotempty(SampleUrls)
| extend AlertDisposition = case(
    ProcessCommandLine has_any ("http.server", "SimpleHTTPServer") and (ProcessCommandLine has_any
(AgentAndHybridMarkers) or FolderPath has_any (".hermes", ".openclaw", ".claude", ".codex", ".cursor-server")),
"CONTAIN_COEXPOSED_HTTP_DIR_AND_ROTATE_AGENT_KEYS",
    ProcessCommandLine has "dingtalk-connector" or SampleUrls has "dingtalk-connector",
"HUNT_EXPOSED_OPENCLAW_DINGTALK_BRIDGE",
    "HUNT_HYBRID_LISTENER_OR_PORT_5061_ANOMALY"
)
| project NetFirstSeen, NetLastSeen, ProcTime, ResolvedDevice, AlertDisposition, InitiatingProcessFileName,
FileName, FolderPath, ProcessCommandLine, SampleEndpoints, SampleUrls, ConnectionCount
| sort by coalesce(ProcTime, NetLastSeen) desc
```

sigma_hybrid_c2_http_server_rhadamanthys_and_agent_env.yml

SIGMA

```
title: Ad-Hoc HTTP Directory Server Launched Inside C2 Framework or AI-Agent Workspace
id: 4b91e7a2-8c34-4f19-9d52-202611203501
status: experimental
description: Detects unauthenticated Python, PHP, or Node ad-hoc HTTP directory servers launched from directories
associated with C2/OAST tooling or AI coding-agent session stores.
author: Odin by Brandefense
date: 2026/11/20
modified: 2026/11/20
tags:
  - attack.resource-development
  - attack.t1583.003
  - attack.credential-access
  - attack.t1552.001
logsource:
  category: process_creation
  product: linux
detection:
  selection_http_server:
    Image|endswith:
      - '/python'
      - '/python3'
      - '/php'
      - '/node'
    CommandLine|contains:
      - '-m http.server'
      - 'SimpleHTTPServer'
      - '-S 0.0.0.0:'
```

```
      - 'http-server'
  selection_workspace_context:
    CurrentDirectory|contains:
      - '/.hermes'
      - '/.openclaw'
      - '/.claude'
      - '/.codex'
      - '/.cursor-server'
      - 'cyberstrike'
      - 'hexstrike'
      - 'sliver'
      - 'interactsh'
      - 'adaptix'
      - 'strix'
      - 'vshell'
  filter_approved_lab:
    ComputerName:
      - 'REDTEAM-SIM-VPS-01'
      - 'PURPLE-OAST-NODE-01'
  condition: selection_http_server and selection_workspace_context and not 1 of filter_*
falsepositives:
  - Authorized red-team staging hosts listed by exact hostname in the filter
level: high
---
title: Archival or Remote Transfer of AI Coding-Agent Session Stores or DingTalk Bridge Configs
id: 6f03a9e4-5d12-4c87-8a90-202611203503
status: experimental
description: Detects archival utilities or CLI transfer tools reading Hermes, OpenClaw, dingtalk-connector,
Claude Code, Codex CLI, or Cursor Server session and API key directories.
author: Odin by Brandefense
date: 2026/11/20
modified: 2026/11/20
tags:
  - attack.credential-access
  - attack.t1552.001
logsource:
  category: process_creation
  product: linux
detection:
  selection_transfer_or_archive:
    Image|endswith:
      - '/tar'
      - '/zip'
      - '/curl'
      - '/scp'
      - '/rsync'
  selection_agent_secret_paths:
    CommandLine|contains:
      - '/.hermes/'
      - '/.openclaw/'
      - 'dingtalk-connector'
      - '/.claude/'
      - '/.codex/'
      - '/.cursor-server/'
  filter_backup_service:
    User:
      - 'svc-borgbackup'
      - 'svc-restic-backup'
  condition: selection_transfer_or_archive and selection_agent_secret_paths and not 1 of filter_*
falsepositives:
  - Registered backup service accounts listed by exact name in the filter
level: high
```

ODIN

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/open-directories-beside-c2-listeners/
Short Read · brandefense.io/odin-research/blog/open-directories-beside-c2/
Defense Kit · brandefense.io/odin-research/defense-kit/