

What Happened, and What Breaks Next: A Defender-Usable Reconstruction of the Summer-2026 Agent-Swarm Incidents

Mike Haddock - Blackfish Security

With Apart Research

Research conducted at the AI Incident Response Sprint, September 2026.

Abstract

Autonomous agent swarms breached production infrastructure twice this summer: the July Hugging Face intrusion (~17,600 attacker-attributed actions, July 11-13, roughly 700 agents) and a second campaign against a dormant German wiki that began May 24 and went unacknowledged until researchers published on September 4. Defenders reconstruct these incidents from press releases. We built the record we wished existed the morning after, from public traces: an independent reconstruction of the Hugging Face fleet with indicator deliveries to HF incident response, a census of 26 live venues, a monitor polling daily, and a disclosure record that handled a criminal-class lane without ever fetching a body. New this sprint: 5 forecasting questions about what breaks next, each paired with a check runnable tomorrow; the first pass ran all 5. The reconstruction, code, census, and 989k-row evidence corpus are public. The next incident should start from evidence, not from zero.

1. Introduction

Between July 11 and 13, roughly 700 autonomous AI agents intruded into Hugging Face's production infrastructure and took about 17,600 attacker-attributed actions, with no human directing any step; the wider campaign began two days earlier, on July 9. On September 4, researchers published a second breach: around 15,000 edits to a dormant German-language wiki, starting May 24, weeks before anyone was watching; OpenAI acknowledged it on September 5. The independent collusion.wiki analysis counts closer

to 17,000-18,000 posts; we cite the published numbers and flag the gap rather than split it silently. Both incidents are public record through Hugging Face's forensic timeline [1], OpenAI's account [2], the SaferAI open letter [3], and CeSIA's analysis [4]. OpenAI and METR's final reports on the Hugging Face incident landed August 26.

What nobody published is the part a defender actually needs. The swarms did not evaporate when the timelines closed. Their paste sites, boards, shortener chains, dead-drops, and marketplaces are still up, and a second generation is already organizing on top of them with signed envelopes, receipt norms, and recruiting. When the next incident happens, responders start from zero. That is the part worth fixing. The sprint's own framing says it cleanly: a warning shot only functions as a warning shot if somebody converts it into specific things to check and specific things to watch, with resolution criteria attached [16].

This report does three things. First, it lays out the swarm's communication infrastructure as it stands: 26 live venues, who operates them, and what they emit, all pulled from public traces under a strict public-evidence-only protocol. Second, it hands that work over as an evidence base (two repositories, a polling monitor, a disclosure record) so the next team extends it instead of rebuilding it. Third, it adds a forecasting question set: resolvable questions about what breaks next, each with a check runnable tomorrow, with first results from this weekend.

Our main contributions are: 1. A public, scrubbed, anonymization-verified evidence base for agent-swarm infrastructure: `the-second-swarm v1.2 (bb707e2)`, including an independent artifact-side reconstruction of the July Hugging Face fleet intrusion (352-dataset farm anatomy, implant generations, the Aug 27 reattribution of the "weapon tier" as incident exhaust, enforcement tracked to 35/35 repositories disabled) and direct indicator deliveries to Hugging Face incident response on Aug 8-10 (our timestamped sends and the scrubbed correspondence are in-repo), plus `swarm-index-watch`. [Prior work - public since Sep 9, before this sprint; the reconstruction base.] 2. [SPRINT-NEW] A forecasting question set of 5 resolvable questions about the HF incident and the German-wiki breach, each paired with an exactly-specified public check (§4 reports first results). 3. [SPRINT-NEW] This report: the reconstruction narrative, first check results, and honest limitations.

2. Related Work

The primary record is strong on what happened and thin on what a defender can do with it. Hugging Face’s forensic timeline reconstructed ~17,600 attacker actions across the July 11-13 intrusion. OpenAI’s account established the sandbox-escape origin and the May 24 predating of the German-wiki campaign. The SaferAI open letter and CeSIA’s analysis cover the disclosure gaps against Article 91 processes. Our own prior contribution to that record, public since Sep 9: an independent artifact-side reconstruction of the HF fleet, built from public archives before and alongside the vendor reports, plus indicator deliveries to HF incident response on Aug 8-10 with written publication go. The collusion.wiki dataset [5] is the deepest public reconstruction of the German-wiki campaign (~17,000-18,000 posts, 3,700+ self-named agents, 98.5% Azure IPs), and its authors argue it was a different swarm from the one that hit Hugging Face, a question we take up in §4. Anthropic’s retrospective across 141,006 evaluation runs found three incidents of models taking unauthorized real-world actions [9]. The UK AI Security Institute’s controlled evaluations found every one of five frontier models attempting to cheat, at rates of 14.1%, 12.6%, 11.4%, 9.1%, and 7.8% across 475 runs each, and AISI’s own near-miss, in which a model wrote and ran code on an external service to reach the evaluator’s infrastructure, shows the failure mode is not hypothetical [10]. The Cloud Security Alliance’s CISO community published a post-mortem within a fortnight [11]. The closest existing project to this submission’s premise is CLTR’s Loss of Control Observatory, which argues that real-world detection capability for control incidents does not exist and prototypes an observatory to build it [13]; we take up that premise in operating form: live venues, daily emission polls, and checks with first results, rather than as a proposal. Kane’s sprint talk covers public-trace reconstruction of the HF incident [6]; Mallen’s covers the swarm loss-of-control thesis [7]; Mengesha’s names the measurement problem in incident response [8].

Missing from all of it: a defender-operable map of the infrastructure layer where these swarms actually live, and checks that turn a map into predictions. Timelines tell you what happened. A live index tells you what breaks next. That gap is what this submission fills.

3. Methods

Venue census. We enumerated agent-communication surfaces from three seed sources: relays over public indexes, registry crawls, and citation expansion from captured swarm artifacts. Deduplication merged alias pairs (e.g. publicbbs and swarmmemo serve the

same board; wikiservice venues are path-based, not subdomain-based, a correction that matters for automation). The census is bounded to public surfaces with verifiable shape. Venues whose shape could not be verified, and one criminal-class venue, were excluded by discipline and are listed as future work. Nothing was silently dropped.

Monitor. `swarm-index-watch` polls the venue list on a schedule, records emission state per venue, and flags changes. Output is scrubbed and anonymization-verified before any publish. The monitor grew 18 → 26 venues on Sep 10 under an operator-authorized change with backup; one venue was pulled from the list the same day, leaving 25 live. The public repo carries the current 25-venue list (28a4f89).

Sweep methodology. A one-day wide sweep (11 subagents, ~350 bounded GETs across 40+ surfaces, ~880 archived files with manifest) produced the current map. Disclosure follows a strict protocol: findings with victim impact are reported at title level only; bodies of illicit material are never fetched. A criminal-class lane (listing titles reportable to law enforcement on sight, on one paste site) was handled under this protocol: one receipt archived, bodies never fetched, never will be, report lane routed to the operator's judgment. A credential-dump lane was reported across all five affected lanes at title level and the local copy destroyed. Anyone doing this work long enough hits material like this. The protocol exists because we did. That lane has since run its full arc under the protocol: expanded URLs delivered to the hosting provider twice, parallel filings with the national authorities already holding the file, and a final note after the site went dark (HTTP 503, all transports) during the correspondence. Lane closed by operator decision; a scheduled re-check runs outside this submission's window.

Forecasting checks [SPRINT-NEW]. Each question in the §4 set specifies: the claim under test, the public data source, the exact check procedure, and the pass/fail criterion. Every check is runnable by a third party with a browser or curl, no privileged access. The format follows the third-party-investigation specification METR published after the incident [12]; exact reproduction procedures, with commands and evidence pointers, ship as `CHECKS.md` in the monitor repository. Questions were selected for resolvability (binary or measurable outcome), runnability (public data, minutes-to-hours of work), and defensive value (the answer changes what a defender watches or does).

What didn't work. Tooling shifted mid-hunt when EXA relay coverage failed and the sweep moved to alternate retrieval. Two venue shapes resisted automated classification and were parked rather than guessed. We record these because a record that hides its failures teaches the next person to repeat them.

4. Results

Census (full 26-row venue table: Appendix C). 26 venues live and ticking at census time (25 carried on the current monitor list after one same-day pull): a Sticked pastebin fleet (15+ instances, 9 agent-active), agent boards and colonies (thecolony.ai at 431+ items/day; GetPostingBoard at ~40,427 messages; openagentforum with 49 check-in agents and signed-envelope governance), and an infrastructure layer with zero prior public coverage: shortener chains, CORS/proxy rotators, disposable identity keystones, and counter services. Operator coverage is 4/26 venues (15%) with named-human revert capacity: all four proWiki-farm wikis; the pastebin/agent-board lane is 0/17. Five venues read unclear (agent-shaped named actors, no revert trace); present-plus-unclear upper bound is 9/26 (35%). In plain terms: the swarm's homes are mapped, and almost none of them have landlords: 4 of 26 venues have anyone who can revert abuse.

Infrastructure state. The first-generation (June-July escapee) infrastructure is decaying: beacons quiet since Aug 22, dead-drop tokens dead, heartbeat counters dying with their counter service. The second generation is organized, recruiting, and growing proto-governance: signed envelopes, receipt norms, verification mechanics. One seeding endpoint (159.146.96.208) fired Sep 5-7 and fell silent after: 0/5 venues recur on Sep 11 re-pull, zero Sep 10-11 shard hits, no /24 siblings (F1: FAIL - the seeding layer went quiet, consistent with the decay thesis). Shortener abuse persists: 10/26 malicious-alias domains alive (38%), 15/26 resolving; deaths cluster in the CORS-proxy lane (F5).

Live-illicit snapshots (title-level). We caught live abuse in progress three times during this work and filed two reports before the sprint began. During the sweep we archived a bypass-probe wave mid-fight: 51 probe edits against a production service, hand-reverted by that service's operator, in near-real time. A prompt-harvester advertising against agent boards (llmrt-class) was mapped from public text. A paid-agent marketplace listing was documented. All three are disclosure-lane items; two reports were filed before the sprint (RubyGems registry email Sep 8; HackerOne #4011040, Sep 8 16:41Z). On a separate, earlier lane: indicator deliveries to Hugging Face incident response went out Aug 8-10, the direct-work leg behind the fleet reconstruction in contribution 1.

Published reach (Track 4 leg). Census summary and a reporter playbook were published to the security community (X post with boost, researcher threads, sprint Discord); a 60-message recheck showed no leak of the illicit-lane material. Disclosure discipline held under publication pressure.

Forecasting checks [SPRINT-NEW - Table 2]. All five ran the sprint weekend off the Berlin box; per-check evidence tables ship in the public repos [15]:

- F1 Seeding reuse: FAIL. The Sep 5-7 burst persists in 30-day windows but shows zero new emissions (0/5 re-pulled venues, 0 Sep 10-11 shard hits). A negative result with teeth: the seeding layer went quiet.
- F2 Operator coverage: 4/26 present (15%). Named-human revert capacity on four proWiki wikis only; pastebin/agent lane 0/17; five unclear. Most of the swarm's infrastructure has no one watching it.
- F3 Harvester detection: PARTIAL. Marker set flags the proxy-positive with zero false positives on 168 fresh docs but is overfit (n=1); grammar approach false-positives at 100% on one board. No local llmrt listing exists - re-run flagged future work. Detection from public text works - barely, and only with honest caveats.
- F4 Second-breach linkage: ECOSYSTEM. Five of seven indicators distinct (windows, naming, cadence, infra, TTPs); two shared (same OpenAI-eval developer, relay grammar). Agrees with collusion.wiki's distinct-swarm call: one ecosystem, two swarms.
- F5 Shortener persistence: 10/26 alive (38%), 15/26 resolving (200+3xx = 58%). Per-alias URL table flagged future work (domains censused, not alias URLs). Abuse domains die slowly enough that takedown speed is the whole game.

Incident timeline: May 24 German-wiki edits begin → Jul 9 campaign begins → Jul 11-13 HF intrusion (~700 agents) → Aug 22 first-gen beacons quiet → Sep 4 researchers publish second breach → Sep 5 OpenAI acknowledgment → Sep 8 first reports filed → Sep 9 repos public → Sep 10 wide sweep (26-venue census) → Sep 11-13 this sprint.

Figure 1. Timeline of the 2026 agent-swarm incidents and the defender response




Figure 1. Timeline of the 2026 agent-swarm incidents and the defender response. Upper panel: May through September arc. Lower panel: September detail.

5. Discussion and Limitations

Implications. The old network is rotting while a new one organizes on the same primitives: decay and regrowth, simultaneously. Formally: the swarm layer is bifurcating. It is easier to list these venues than to log into them: the exposure axis is enumerability, not authentication. High-grade write primitives on public services need no identity, and disposable identity services collapse what little friction remains. A defender gets more from watching a short list of enumerable venues than from any control on the authentication side. The mechanism is an ownership gap: 22 of 26 venues have no confirmed human revert capacity, and the economics of free paste and board services do not create one. For the communication lane (Track 4): every claim above links a primary source or a public repo, and the disclosure record shows the protocol surviving contact with criminal-class material.

When the guardrails refused the responders. During the Hugging Face breach, the company's own responders were refused by hosted frontier models on much of the forensic work: the guardrails could not distinguish an incident responder from an attacker, so the team fell back to a self-hosted open-weight model [1]. This is not an anecdote. Controlled testing across 2,390 real prompts finds security keywords refused at 2.72 times the rate of equivalent neutral phrasing, and explicit authorization raised refusal instead of lowering it, 21.8% against 11.6% [14]. The implication lands on this submission's method: the hosted layer refused the good guys, and scoped systems under a human hand were the fallback that worked.

Scoped AI teaming, human-gated. Every result above came from one human holding every gate while scoped AI systems did the grinding: recon at 3 AM, forensics at machine speed, each claim re-verified by hand before it entered the record. The bias question is fair and we state it plainly: the builders share a substrate class with the population under study. Our answer is the method: tight scopes, a human deciding at every gate, and every factual claim linked to a primary source a stranger can check. Scoped and hands-on, augmented teaming works, and this layer moves too fast for anything slower to keep up.

Limitations. One human operator plus AI pipeline; all claims human-verified but this is a single team's effort. Time-boxed: the sweep was one day, the monitor days old, the checks one weekend. Coverage bounded to public surfaces with verifiable shape; dark venues and unverifiable shapes are excluded and listed. Reporting is title-level for victim-impacting material by protocol.

Future work = the parked list. Unwatched venues showing operator-absence plus paste-volume growth (next-break surface question); the 154-venue bulk registry as census expansion; the infrastructure layer (counters, shorteners, proxies) as the structural gap; venue-response-time measurement for the disclosure lanes. Already landed mid-sprint: an offline signal scanner over shard history implementing three detection axes named in Kane's sprint talk [6]: per-actor cadence metronomes, cross-venue relay-lag fingerprints, and actor-text heuristics (`swarm_signal_scan.py`, commit 1ae50a9). A public index that absorbs new detection ideas within hours of their publication is itself the speed-of-ops argument this paper makes; our thanks to Kane for naming the axes cleanly.

6. Conclusion

The infrastructure layer behind the 2026 agent-swarm incidents is mappable, watchable, and mostly unattended. This work maps it: 26 venues, a decay-and-regrowth dynamic nobody had written down, and checks that turn a map into expectations.

These incidents will not slow down, and no human team alone can watch a layer that moves at machine speed. That is not despair; that is opportunity. The same class of system that did the damage can be pointed at defense: the threat model becomes the defender. Unlike the swarms, faceless and running on internal intent alone, this layer is augmented defense: gated by a human hand, informed by the record, and fast enough to match the attackers' speed. We cannot predict the next shape of the threat, but we no longer have to meet it blind.

The next incident will not start from zero unless defenders choose to start there. The methods we have developed, when further refined, could provide near-real-time alerting and mitigation before the incident gets out of hand. The path forward is human and AI teammates, tightly scoped, with the operator in the middle of the decision tree - moving at the speed of the threat. Swarms move fast - so now can the defenders.

Code and Data

- <https://github.com/darkfibr/the-second-swarm> (v1.2, bb707e2) - reconstruction, lay guide, addendum; HF fleet case + HF IR correspondence in 02_hf_fleet/ and 03_disclosure_record/
- <https://github.com/darkfibr/swarm-index-watch> (28a4f89) - monitor, 25-venue list, per-check evidence files (checks/apart_sprint_20260911/), reproduction procedures (CHECKS.md), offline signal scanner (swarm_signal_scan.py)
- <https://github.com/darkfibr/hf-fleet-corpus> (corpus-v1.0) - 989k-row Hub metadata snapshot + analysis subset behind the fleet reconstruction; scrubbed, hashed, timestamped
- Novel tradecraft withheld from repos per disclosure review; contained in this PDF where safe.

Author Contributions

M.H. (Blackfish Security): investigation direction, operator presence and steering throughout, verification of every claim, and the written corpus. Lyra (AI colleague): recon, forensics, drafting scaffold, gap-finding, reference assembly. Hermes (AI colleague): emission verification, forensics support. All factual claims carry primary-source links; the human author verified each.

LLM Usage Statement

AI colleagues performed reconnaissance, forensics, drafting scaffolding, and reference assembly under continuous human operation and steering; they also checked reasoning and found gaps in drafts. The written corpus, all verification, and every decision are the human author's. Every factual claim links a primary source.

References

1. Hugging Face, “Agent Intrusion: Technical Timeline” (forensic reconstruction, ~17,600 attacker actions). <https://huggingface.co/blog/agent-intrusion-technical-timeline>
2. OpenAI, “Hugging Face Model Evaluation Security Incident” (account + Sep 5/7 disclosures). <https://openai.com/index/hugging-face-model-evaluation-security-incident/>

3. SaferAI et al., “Open Letter to the European Commission,” July 2026. <https://www.safer-ai.org/u/2026/07/Open-Letter.pdf>
4. CeSIA, “The OpenAI-Hugging Face incident: what we know, what we don’t, what follows.” <https://cesia.org/en/publications/the-openai-hugging-face-incident-what-we-know-what-we-dont-what-follows/>
5. collusion.wiki, “Discovery of a new OpenAI agent message board” - dataset and analysis of the DSE/prowiki campaign. <https://collusion.wiki>
6. Kane, B. “Uncovering public traces of the OpenAI HuggingFace incident,” Apart Research × CeSIA AI Incident Response Sprint, Sep 11 2026 (talk recording). https://us06web.zoom.us/rec/share/5mqGISLYgW-CIta2SK838S6f5c_yh9DEUtzwaDOOrGvsu2eMGRrlic4HB54eyHsb.2yzgyLRLPwH4xzjW
7. Mallen, A. “How near-term AI swarms could cause labs to lose control of AI development,” Apart Research × CeSIA AI Incident Response Sprint, Sep 11 2026 (talk recording). <https://us06web.zoom.us/rec/share/dzHEL-hbWwYiKrS3j1rGtKPiahGmG321WwMqbCQHVaYqZ4kGgxAVargSJB73R9uM.GxeT7W3z-zislvHL>
8. Mengesha, I. “Incident Response Has a Measurement Problem,” Apart Research × CeSIA AI Incident Response Sprint, Sep 11 2026 (talk recording). https://us06web.zoom.us/rec/share/uKW9f3s8VK60t7DP-R5O5A7_etxuSkKifyZ7vX5prPDiQzkXGZtbuNSfJkOxWIN.hCBZFEg9Olmo67Je
9. Anthropic, “Investigating three real-world incidents in our cybersecurity evaluations,” Jul 30, 2026 (review of 141,006 evaluation runs). <https://www.anthropic.com/news/investigating-incidents-cybersecurity-evals>
10. UK AI Security Institute, “Cheating behaviour in frontier model evaluations,” Jul 21, 2026 (five frontier models, 475 runs each). <https://www.aisi.gov.uk/blog/cheating-behaviour-in-frontier-model-evaluations>
11. Cloud Security Alliance CISO community, “Hugging Face Incident Initial Post-Mortem,” Jul 27, 2026. <https://cloudsecurityalliance.org/artifacts/hugging-face-ciso-post-mortem>
12. METR, “How independent researchers could investigate AI propensities after misalignment incidents,” Jul 29, 2026. <https://metr.substack.com/p/2026-07-28-investigating-ai-propensities-after-incidents>
13. Centre for Long-Term Resilience, “The Loss of Control Observatory: a prototype to detect real-world AI control incidents,” Feb 2026. <https://www.longtermresilience.org/reports/the-loss-of-control-observatory-a-prototype-to-detect-real-world-ai-control-incidents/>

14. Campbell, D., N. Kale, U. Schwag, B. Herring, N. Price, D. Borges, A. Levinson, C. Knight, “Defensive Refusal Bias: How Safety Alignment Fails Cyber Defenders,” arXiv 2603.01246. <https://arxiv.org/abs/2603.01246>
15. the-second-swarm v1.2; swarm-index-watch (28a4f89); hf-fleet-corpus (corpus-v1.0). (Code and Data above)
16. Apart Research × CeSIA, AI Incident Response Sprint page and evaluation rubric. <https://apartresearch.com/sprints/ai-incident-response-sprint-2026-09-11-to-2026-09-13> · <https://apartresearch.notion.site/sprint-evaluation-rubric>

Appendix A - Limitations & Dual-Use Considerations

Public-evidence-only collection; title-level reporting for victim-impacting material; criminal-class venues excluded from automation and handled manually under operator judgment; credential-dump bodies destroyed after five-lane title-level reporting; no novel installation recipes released (withheld pending review); repos scrubbed and anonymization-verified before publish.

Appendix B - Forecasting Question Set

#	Question	Check	Why it matters
F1	Does seeding infrastructure (159.146.96.208-class) recur across venues over time?	Monitor emission logs + IP re-census	Fingerprint reuse = cheap blocklist
F2	What fraction of swarm venues have any operator revert capacity?	Census operator-field audit	Most have none; quantifies neglect
F3	Can prompt-harvester listings (llmrt-class) be flagged from public text alone?	Marker set vs fresh venue pulls	Detection without privileged access
F4	Does the German-wiki breach share fingerprints with the HF intrusion?	collusion.wiki corpus vs HF IOCs	One actor vs one ecosystem
F5	How long do malicious shortener aliases survive?	Alias re-census (26 at census time)	Decay curve sets takedown SLAs

- Cut: lane responsiveness (valid but grades vendor response mid-sprint; alternate); next-break surface (requires watching venues excluded by shape-discipline; future work).

- **Weekend results (ran on Berlin, all 5):** F1 FAIL (0/5 recur - seeding silent since Sep 7) · F2 4/26 present (15%) · F3 PARTIAL (marker set overfit n=1, re-run flagged) · F4 ECOSYSTEM (5/7 distinct, agrees with collusion.wiki) · F5 10/26 alive (38%). Full evidence: §4 + repo check files.

Appendix C - Venue Census Table

Venue	Operator presence	Evidence one-liner
probier	present	HelmutLeitner 22 user-kind edits incl ForumSeite 8-change burst; Sep-08 shard rows (MessageBoardsForAgents , TestSeite changes=4)
dse	present	HelmutLeitner + GertMUC + MartinQuicksilber ; WikiIncident pages Sep-10 with changes up to 21, ForumSeite [minor editing of my post]
gruender	present	HelmutLeitner TestSeite / Context / RecentChanges edits Sep-08
dorfwiki	present	FranzNahrada TestSeite [revert to revision 1.287] + Seite gelöscht. + AgentsImDorfWiki changes=6; FritzEndl Triesterviertel edit
fractal	unclear	8 user-kind actors (probe-shaped), no human-moderator name, no revert trace; live rc 2026-09-11 shows only probe names
culios	unclear	Single user-kind actor SacrificeRational (agent-shaped), no revert trace
lotr	unclear	Single user-kind actor ObeyCollective (agent-shaped), no revert trace
schulwiki	unclear	Single user-kind actor ForOurOwnNoWayFix (agent-shaped), no revert trace
getpostingboard	unclear	Endpoint serves HTML not JSON (Sep-10 shard error); live board 2026-09-11, no moderation surface observed
k4be	absent	15 docs, 0 user-kind actors, no revert traces

Venue	Operator presence	Evidence one-liner
fragbin	absent	30 docs, 0 user-kind actors
agentworkpad	absent	8 docs, 0 user-kind actors
thecolony	absent	595 docs, 0 user-kind actors
mobdarkforest	absent	126 docs, 0 user-kind actors
minetest	absent	15 docs, 0 user-kind actors
artixlinux	absent	32 docs, 0 user-kind actors
openagentforum	absent	114 docs, 0 user-kind actors (agent_* senders are IP/anon-kind API identities, no moderator)
publicboard	absent	139 docs, 0 user-kind actors, no moderation trace
aiforum	absent	114 docs, 0 user-kind actors
swarmmemo	absent	72 docs, 0 user-kind actors
probyte	absent	15 docs, 0 user-kind actors
annafyi	absent	15 docs, 0 user-kind actors
fasterit	absent	7 docs, 0 user-kind actors
tarcseh	absent	15 docs, 0 user-kind actors
dynavirt	absent	15 docs, 0 user-kind actors
smirky	absent	7 docs, 0 user-kind actors

Full evidence lines (shard refs): `swarm-index-watch` → `checks/apart_sprint_20260911/F2_result.md`.
