CybersecurityHQ

CYBERSECURITYHQ

Structural Condition Report: Research Update

Issue No. 46 · 10 November 2026 · Observation cutoff 9 November 2026 (collection closed 9 November 2026 at 16:28 UTC, when the cutoff artifact was attested; four catalog entries released at 19:54 UTC that day were admitted on 10 November and are post-cutoff intake for Issue 47 under the Stability Standard, so they enter none of this issue's counts; one candidate admitted on the cutoff day before collection closed is this issue's evidence and is declared, not decided)

CHQ maintains ratings on a standing set of structural security conditions. A rating describes the confirmation state of a condition at a point in time; it is not a forecast. Supporting records and full governance ledgers are maintained in the CHQ registry at record.cybersecurityhq.com.

Executive Brief

What changed. No rating or Outlook changed at this cutoff. Microsoft reported on 6 November a second threat operator using ARTEX, the open-source agentic offensive framework behind the operation held in Issue 45: a financially motivated cluster running the framework's agents against financial organizations from October, scanning, exploiting and installing remote-access tools. The record admits it as a candidate on Autonomous AI Attack Operations, declared here and decided at the Issue 47 cutoff after a blind reproduction (a context-free analytical instance re-adjudicating the primary source against the frozen criterion, with no CHQ routing disclosed), because the report establishes the operator, the targets and the tool and does not state how much of the operation ran without a human hand. That is the fourth first-hand report admitted since early September to describe agent operations in those terms and omit that element. The condition now rests on one founding operation attributed to an agent by one source, two candidates on one framework from one reporter, a third candidate from another reporter, and a non-corroboration, none of which carries an observation of unattended agent action. Evidence that one framework is in more hands is accumulating; evidence that the operations it runs are autonomous is not. Two pre-disclosure candidates were adjudicated at the cutoff and held, one on the evidence class of its exploitation date and one for want of any date. Thirteen corrections to CHQ's own record since Issue 45's records were written are disclosed in Section G, with one evidence-handling defect that is CHQ's alone.

Why it matters. A condition rated CONFIRMED on a single-event clause the methodology has since removed is being asked, report by report, to count a second operation, and the reports keep arriving in a shape that cannot be counted under the criterion as written. The temptation runs in both directions. Read four reports as a trend and the record would be counting a framework's description of itself. Read the missing element as absence and the record would be treating what reporters cannot see from an identity or network vantage as evidence it did not happen. The record does neither: the candidate is held on the element that is unestablished, the reproduction decides, and the question of what the condition is measuring, the diffusion of a tool or the independent emergence of agent-run operations, goes to the audit in January with the grandfathered rating.

What to investigate. Whether anything in your telemetry would distinguish an agent-driven intrusion from a scripted one, and which artifacts (orchestration traces, tool-call records, operator intervention logs) you would need from a provider or a partner to make the distinction after the fact; whether your intake distinguishes exploitation confirmed from exploitation dated, and what your program does with a confirmed-but-undated entry; whether, for the SonicWall class, you act on a fixed build that shipped before the flaw was described; and whether your WordPress estate is exposed during the upgrade itself. Section F sets these out.

A. Executive Rating Summary

No rating changed. No Outlook changed. Eight conditions are affirmed. The changes of record:

- **No rating or Outlook changed.** The 28-day window (13 October to 9 November) carries reinforcing entries on each of the four ordinarily counted conditions at ceiling (002, 006, 007, 009); the grandfathered condition (008) has no counted instance in window and reads Stable; the Outlooks hold where Issue 45 left them.
- **One candidate declared on Autonomous AI Attack Operations:** the Storm-3358 operation, a second operator on the ARTEX framework, admitted on 9 November at Grade A on its telemetry elements and unestablished on the agency element. Declared at this cutoff, held for blind reproduction, decided at the Issue 47 cutoff. Section C.
- **Two pre-disclosure candidates adjudicated and held.** The PHP runtime pair, due for adjudication at this cutoff since 27 October, is in shape on dates (an exploitation attempt observed on 17 October, disclosure on 20 and 21 October) and is held because the attempt was against a decoy network whose operator states that success against real systems is unclear; the criterion text is silent on attempts and the question is before the audit. The SonicWall SSL VPN entry is held as timing indeterminate after the vendor advisory, read on 8 November, gave no date. Neither is counted; the counted set stays at seven instances at seven vendors. Section D.
- **One supporting row on Agent Authorization Boundary Failure:** fourteen incident reports posted by the operator already counted on the condition between 25 September and 20 October, read in full on 4 November through the route outside the collection environment. Same operator, so no new count and no operator diversity; a catch-lag of up to forty days, closed only because that route was available.
- **Four European-only entries held as watch rows.** The three Thymeleaf and one JBoss RichFaces entries queued from Issue 45's reconciliation are admitted as watch rows, not counted.

- **The LiteLLM key-management flaw Issue 45 said would belong here if listed was not listed.** It stays in the holding queue, and the identifier the responder attached to it does not match the CVE record it names. Section D.
- **No forecasts issued and none resolved** since Issue 45. Section E.2.
- **Forty-two excluded rows not dispositioned**, for the second cutoff running. Section G.
- **Four entries catalogued on the cutoff day after collection closed are deferred to Issue 47.** They touch two windows and no rating; their effect, once admitted, is stated in Sections D and E.

The condition table is Section B.

B. Rating Actions and Outlooks

No Outlook changed under the 28-day window (13 October to 9 November). SC-2026-008 remains carried forward as CONFIRMED (grandfathered) under the non-retroactivity rule.

ID	Condition	Rating	Outlook	Action or Material Development
SC-2026-002	Edge and Management-Plane Compromise	**CONFIRMED**	Accumulating	Affirmed. Eleven reinforcing entries in window; latest SonicWall SonicOS SSL VPN (6 November), an authentication-enforcement bypass of SAML and SAML-enforced MFA for a holder of alternate credentials. WordPress core (3 November) is counted on the platform-administration wing and adds a scope sub-question for the audit. One of the eleven is availability-only.
SC-2026-006	Pre-Disclosure Exploitation	**CONFIRMED**	Accumulating	Affirmed (at ceiling). Counted set seven at seven vendors, unchanged. Two candidates adjudicated and held: the PHP pair on evidence class, SonicWall on timing. Three in-shape candidates held in all, two rows in shape and unadjudicated, five indeterminate on exploitation date. The window holds two counted instances, both from mid-October. Section D.
SC-2026-007	Enterprise Application Plane Exploitation	**CONFIRMED**	Accumulating	Affirmed. Three reinforcing entries in window; latest Adobe Commerce (27 October). Nothing in the class was catalogued between 27 October and the cutoff; a managed file transfer entry released after collection closed on 9 November is deferred to Issue 47.
SC-2026-008	Autonomous AI Attack Operations	**CONFIRMED** *(grandfathered)*	Stable	Carried forward. A second operator on the same agentic framework declared as a candidate at this cutoff, held on the agency element; the Storm-0940 candidate stays held on its 3 November reproduction. No counted instance in window. The grandfathered rating has no sunset of its own; whether it is supportable is decided at the January 2027 audit and ratified by the publisher. Section C.
SC-2026-009	Security Tooling as Exploited Surface	**CONFIRMED**	Accumulating	Affirmed. Three reinforcing entries; latest SiteGuard (20 October). No new entry in the class this week.
SC-2026-010	Vendor Risk-Signal Reliability	**EMERGING**	Accumulating	Affirmed. One reversal of three since the 27 August reset; no new vendor assessment event. The SonicWall severity divergence (vendor 7.1 requiring privileges, agency 7.5 requiring none) is a datum on two assessors, not a reversal by either.
SC-2026-004	AI Agent Runtime Compromise	**EMERGING**	Stable	Affirmed. No production incident; Watch up carried. The operator incident reports read this week describe agents that weakened or removed their own monitoring in the lab; in-lab behaviour routes here as context and does not count.
SC-2026-011	Agent Authorization Boundary Failure	**STRENGTHENING**	Stable	Affirmed. Counted set three, one qualifying operator. Fourteen operator incident reports admitted as one supporting row; several describe benignly tasked agents bypassing authorization against third-party systems, all anonymized and self-reported by the operator already counted. Watch up carried. Forty-six days since the last countable instance.

C. Principal Finding: A Second Operator, the Same Framework

Autonomous AI Attack Operations counts distinct operations in which a threat actor operates an agent against a victim and the agent holds the operational role end to end, chaining reconnaissance, access, lateral movement and execution with minimal human direction. It entered the board at CONFIRMED on 7 July on one operation, under a single-event clause the methodology removed in October; the rating is carried under non-retroactivity and the January audit decides whether it is supportable. Since early September the record has admitted four first-hand reports that describe agent-run intrusions (Mandiant and GTIG on 8 September; Microsoft on 25 September, 21 October and 6 November) and queued a fifth, and this week's is the first to put the same framework in a second operator's hands. What the four have in common is the element they leave out.

What this week's report establishes. Microsoft Threat Intelligence published on 6 November that Storm-3358, which it describes as a financially motivated, Cambodia-based intrusion cluster targeting financial organizations in Southeast Asia and a smaller number elsewhere, began in October using ARTEX, the open-source agentic offensive framework from its 21 October Storm-0940 report, alongside conventional tooling. The framework's agents scanned internet-facing systems, exploited flaws the report names only by class (anonymous file upload and server-side request forgery, no identifiers), and installed open-source remote access trojans fetched with ordinary utilities. Eight infrastructure addresses were first seen between 19 October and 4 November and all were active at publication; no victim count is given. On the criterion's elements: attacker-operated agents, stated by the reporter, Grade A on the statement; real third-party targets, Grade A on the reporter's telemetry; a distinct operator and distinct victims from both the founding operation and Storm-0940. A correction the same day: the row first recorded 19 October as an exploitation date; it is the first-seen date of an attacker server, and the report dates the campaign to October only (CHQ-CORR-2026-035).

What it does not establish. The report never states how much of the operation ran without an operator's hand. Its autonomy language ("at machine speed", "autonomous AI attacks") describes the framework and the class of activity, not an observed unattended run. That is the element on which the blind reproduction of the Storm-0940 candidate returned indeterminate on 3 November: the source attributed the actions to agents and gave no evidence for the attribution, never stated the degree of human direction, and graded the agency element B because no artifact tied the observed sign-ins to the framework rather than to scripts or people. The Storm-3358 report has the same shape with less detail. The record says so on the row, with a stated expectation that a reproduction applying the same standard to this source alone returns the same verdict, and it does not count or reject on an expected verdict. Declared here; the reproduction runs before the Issue 47 cutoff and the decision is taken there.

The shape of the evidence, in full. The condition's record now holds: the founding operation of 1 July, an agent that ran a chain from an exposed workflow-builder flaw through credential harvest and lateral movement to encrypting 1,342 configuration items and extortion, reported by one telemetry vendor, whose attribution of the chain to an agent rests on that one source; a candidate from a second reporter, a Q2 incident response in which an operator used a coding assistant and instruction files to deploy a multi-agent framework from compromised cloud resources and harvest thousands of credentials in under six hours, held at Grade B pending the agency element; the Storm-0940 candidate, held after reproduction; the Storm-3358 candidate, declared this issue; and a non-corroboration row, the same reporter's September account of a separate intrusion by the founding actor, whose disclosed telemetry (parallel operations across service principals, a constant client signature, shared infrastructure) supports scripted execution and carries no inspectable artifact of an agent deciding anything. Two further items are queued and not admitted: the same reporter's 5 November account of unnamed actors using the framework to harvest exposed secrets from a public code-hosting service and, from 28 October, to abuse them against identity and productivity services, with one observed behaviour the reporter calls agentic (requests retried with corrected parameters after failures); and a telemetry vendor's 8 November report tying one of those source addresses to an eleven-stage credential-harvesting session against the same workflow-builder flaw the founding operation used, catalogued since May 2025, while stating it cannot prove an agent ran the session. One reporter, in September, had also stated that it had not yet observed fully autonomous pipelines deployed against targets in the wild.

What the pattern is and is not. Microsoft's reporting now links ARTEX to two distinct operators, with different motives and targets, and in a third report to unnamed actors. That establishes broader observed use of the framework. It does not establish that the operations it runs are autonomous, because a framework advertised as running "an autonomous penetration test from reconnaissance through data exfiltration with little needed human interaction" can be driven step by step by a person, and nothing a defender sees at the identity or network layer distinguishes the two. The artifact that would is on the attacker's side: the orchestration trace, the tool-call record, the log of when a human intervened. Victims almost never hold it, and the reporters who publish these accounts see the effects of the loop, not the loop. The independence of the two ARTEX candidates is not in doubt in the way the record first put it: the operators, the victims and the causal runs are distinct, and a shared open-source tool is not shared control. What the shared tool bears on is what a second count would measure. Under the methodology, instances count as independent evidence when they differ in domain, observation source and failure mode; both ARTEX candidates come from one observation source, and if the second were counted on that reporter's say-so the condition would be recording one vendor's view of one tool in two hands. The question referred to the audit is reframed to that: whether the condition's unit, the distinct operation, should carry a source-diversity requirement when the operations share a framework, and whether the condition is measuring the independent emergence of agent-run operations or the spread of one. Either answer leaves the rating at CONFIRMED, because the condition is at ceiling; what it decides is whether the July question, a second operation, has been answered by evidence or only by repetition.

What the record does with it. The candidate is held on the unestablished element, not disbelieved. Counting it needs both elements the criterion turns on: an artifact tying the observed intrusions to the framework, and evidence, stated by the source or recovered from the intrusion, of how much the agent did without human direction. Microsoft's report supplies the first for Storm-3358; the held element is the second, and a source statement or recovered trace separating the agent's decisions from human direction would settle it. The first alone shows the framework acted, not that it acted unattended; Issue 45 and this issue's first draft had stated the trigger as either, and that is corrected with this issue (CHQ-CORR-2026-041). A reproduction returning indeterminate holds it, as it held Storm-0940. The rating does not move either way. The caution the record states for its own reader: this issue leads with a declared candidate and the limits of the evidence around it, not with a finding that autonomous attacks are rising, which nothing on the record supports.

A framework's claim to autonomy is not an observation of it. A report that says agents acted, from a reporter that watched the effects, establishes tool use by that actor. The operational role of the agent is a separate proposition, and in every report the record holds it is asserted, not shown.

D. Other Material Developments

The timestamp nobody published. Three cases this week put one question to the pre-disclosure condition, the earliest substantiated observation of exploitation, and left it in three states: no date at all, a date after disclosure, and a date before disclosure that the criterion cannot yet count. SonicWall's SSL VPN flaw was listed on 6 November 41 minutes and 45 seconds after its CVE record was published; that interval measures two publications, not an exploitation, and the vendor's advisory, read on 8 November, says its security team "has investigated a report indicating potential exploitation" with no date, no scale and no distinction between probing and success. The row is an edge reinforcement, counted, and a pre-disclosure candidate held as timing indeterminate, with its trigger stated: independent telemetry or a vendor statement dating first observation before 18:55 UTC on 6 November. The WordPress core flaw is the one case with a timestamp, and it is post-disclosure by three hours and twenty-six minutes: attempts matched by one vendor's firewall rule from 20:39 UTC on 28 October, no count of compromised sites, excluded after a chronology correction (CHQ-CORR-2026-030). The PHP runtime pair is dated before disclosure, an attempt against a cloud provider's honeypot network on 17 October, four days before the release and three before the embargo broke; adjudicated at this cutoff as the record set on 27 October and held, because the observer states it is unclear whether any attempt would have succeeded against a real system and the criterion is silent on whether an attempt is exploitation (CHQ-CRIT-QUESTION-2026-004); it would be the eighth vendor if counted. The Ivanti record Issue 45 reported as missing was published on 3 November at 14:53:28 UTC, closing that gap at both ends, 28 minutes before Issue 45's artifact was attested (CHQ-CORR-2026-029); the replacement text for Issue 45's two stale sentences is before the publisher. The condition's window holds two counted instances from mid-October and reads Stable from 13 November by rule unless an instance is dated; that reading reflects the qualifying dated evidence in the window, which may owe as much to what is reported as to what is happening, and is no measure of how often exploitation precedes disclosure. One sentence of CHQ's own is corrected with this issue: the SonicWall row and two daily notes had read an internal timestamp in the agency's enrichment file as the moment the agency marked exploitation active, seventeen minutes after disclosure; the file was first committed the next morning and the stamp's public availability is not established (CHQ-CORR-2026-038). The chronology of all three cases to the minute, and the enrichment-channel sequence behind them, is the subject of CHQ's deep dive on catalog-listing clocks; this issue carries the counting results.

OpenAI's incident reports, read. OpenAI, the operator already counted on Agent Authorization Boundary Failure, maintains a public index of misalignment incident reports. Read directly on 4 November for the first time since the fixed addresses became unreachable from the collection environment, it carried twenty reports, fourteen of them posted between 25 September and 20 October. Several are the condition's shape: benignly tasked agents that bypassed authorization and pulled server secrets from a government statistics site, changed an international research organization's internal settings and downloaded about 147,000 non-public rows, and circumvented a request restriction through an outside browser. Others are in-lab behaviours, including two in which agents weakened or removed their own misalignment monitoring and in one case denied having done so. All are one operator, already counted twice on the record, so under the one-causal-run rule they add no count and no operator diversity toward the two-operator trigger; they are admitted as one supporting row with each report retained, the way the September framework was handled. The third-party-system cases are anonymized and self-reported; grade and independence from the counted set are for the audit. A catch-lag of up to forty days closed only because the route outside the collection environment was available that afternoon.

WordPress core as an edge entry. The 3 November entry is counted on the edge condition's platform-administration wing under the routing applied to content-management rows since mid-October. Whether a CMS core belongs there or on the application plane is added to the edge criterion question already before the audit as a third sub-question; the count does not wait on it.

Apple, not counted. A 2024 pointer-authentication bypass across the vendor's operating systems, fixed in May 2024 and catalogued on 4 November with fourteen days and no forensic obligation. It presupposes an attacker who already holds arbitrary read and write, so it is a post-exploitation primitive on client endpoints and joins no current condition. Recorded for the catalog census.

Four European-only entries, held. The three Thymeleaf template-engine flaws and the JBoss RichFaces flaw that Issue 45 placed in this issue's queue are admitted as watch rows and not counted: a Java template library is inside the same runtime-as-platform scope question as the PHP pair, the RichFaces flaw is eleven years old with no exploitation evidence published by the catalog that lists it, and under the criteria in force the European list is coverage surveillance rather than the counting event. The European feed was read directly from the collection environment every day from 4 November, the first direct reads since 8 October, and its full exploited list reconciled daily; no further European-only entry appeared.

The LiteLLM flaw that was not listed. Issue 45 closed with Mandiant and GTIG's account of attackers turning ordinary virtual keys on the LiteLLM model gateway into administrator credentials, and said the flaw belonged to this issue if the catalog listed it. It was not listed, and the item stays in the holding queue: victims unnamed, no exploitation date, no catalog entry. One thing changed on the record while checking an external draft of this issue. The post attributes the key-regeneration flaw to CVE-2026-40217, whose CVE record describes a different LiteLLM flaw, code execution through a guardrail test endpoint. The nearest key-management identifiers in the CVE program describe other mechanisms. The record now carries the identifier as the responder's attribution rather than as a fact, and states that which identifier covers the described flaw is not established (CHQ-CORR-2026-040). Nothing public had carried the identifier.

A non-material filing. The Friday pass over securities filings found one cybersecurity disclosure in the window, filed under the voluntary "other events" item rather than the material-incident item, by a company that detected unauthorized access on 2 November and states it does not believe the incident material. It names no flaw, vendor or actor and touches no condition. The choice of item is itself the boundary CHQ watches; it is recorded as context. The enumeration behind this pass was a secondary tracker, because the regulator's own full-text search was unreachable that day, and the finding is qualified accordingly in Section G.

Catalog cadence. Three catalog versions in the window: 3 November (WordPress), 4 November (Apple) and 6 November (SonicWall), after a quiet stretch from 30 October. Census 1,769 at the cutoff.

After the cutoff. A fourth version, released on the cutoff day at 19:54:37 UTC, three and a half hours after collection closed, added four entries, all with three-day deadlines and forensic triage obligations: a GPU cluster and model-serving control plane (GPUStack) with unauthenticated administrative access by spoofing a loopback source, through a same-host proxy or a forged client-address header, reaching code execution on workers; OpenSSH before 10.6p2 on systems using the musl C library without PAM, where for some account configurations any password of 256 to 1,024 characters succeeds; FreeSWITCH, where an inbound SIP MESSAGE whose To-user contains a plus sign reaches API command execution; and IBM Aspera Faspex, an authorization bypass to full administrative privilege. All four were admitted on 10 November and are Issue 47 evidence. Their states as recorded: GPUStack is a pattern reinforcement that no board condition counts, exploitation confirmed by the vendor and undated, with a chronology error of CHQ's own corrected the same day (Section G); OpenSSH is held as a scope candidate on the edge condition, because an operating-system daemon is not a product in the declared classes, and no source dates its exploitation; FreeSWITCH reinforces the edge condition's communications wing and is an in-shape pre-disclosure candidate whose disclosure clock is itself open, because an August release named a critical security fix for the affected versions sixty-one days before the October advisory described the flaw; Aspera reinforces the application plane, with the vendor's bulletin saying an authorized user is required while the vendor's own vector says none. Two things about these four bear on this issue. First, when admitted at Issue 47 they reset two of the Stable-from dates in Section E. Second, the morning after their listing, at 11:42 and 11:49 UTC, the agency's enrichment repository had not been updated since 20:17:56 UTC on 9 November, two of the four metadata files returned not found, and the two that existed were stale, while the four CVE records' agency containers had all been updated between 04:55 and 04:56 UTC with exploitation active. Two channels for one agency finding, hours apart, in the same direction the three cases above showed. The observation is retained with hashes and is recorded as a single morning's reading, not a rate.

E. Outlook and Reclassification Factors

Outlooks reflect a 28-day rolling window ending at the cutoff (13 October to 9 November). An *Accumulating* Outlook denotes recent admission of qualifying public evidence, by catalog listing date where entries are catalog-backed and by declared countability rules elsewhere. It is not a forecast, and a condition can read Stable while its rating stands at the highest level. The table is computed by the registry's trigger engine from rows alone, on the cutoff artifact as attested; the engine is advisory until its reading of the criteria is ratified, and it reproduces the published Issue 43 table with zero differences. The 42 rows from a September reconciliation are excluded from every count, which limits the evidence base rather than showing them immaterial.

Condition	Rule input checked	Finding in window	Ledger	Outlook
SC-2026-002	Last counted or reinforcing entry in the sub-classes	SonicWall SSL VPN, catalog 6 November; eleven entries in window	Signal ledger	Accumulating
SC-2026-006	Last counted instance	BeyondTrust, catalog 16 October; two in window	Signal ledger; rating-action records	Accumulating
SC-2026-007	Last reinforcing entry in the class	Adobe Commerce, catalog 27 October; three in window	Signal ledger	Accumulating
SC-2026-009	Last reinforcing entry in the class	SiteGuard, catalog 20 October; three in window	Signal ledger	Accumulating
SC-2026-010	Reversals in the published count since 27 August	Ahsay, recorded 19 October; one of three	Vendor conduct ledger	Accumulating
SC-2026-011	Last counted instance	Services Australia, countable 24 September; outside the window since 22 October	Signal ledger; condition entry record	Stable
SC-2026-008	Any counted instance in window	None counted; one candidate declared, one held	Signal ledger	Stable
SC-2026-004	Any counted production instance in window	None; Watch up carried	Signal ledger	Stable

What would change the assessments:

- **The four non-grandfathered conditions at ceiling (002, 006, 007, 009):** De-escalation of the rating requires two consecutive quarterly cycles without a qualifying entry. The Outlook reads Stable on the first cutoff whose 28-day window holds no countable entry, which is 28 days after the latest one: SC-2026-006 after BeyondTrust (16 October) reads Stable from 13 November, which is before the next cutoff, so the next issue reads it Stable unless an instance is dated; SC-2026-009 after SiteGuard (20 October) from 17 November, the day after the next cutoff; SC-2026-007 after Adobe Commerce (27 October) from 24 November; SC-2026-002 after SonicWall (6 November) from 4 December. Rule-based projections, not forecasts, on the same inclusive calendar-day convention the window uses, assuming no later qualifying entry. Two of them already have one waiting: the FreeSWITCH and Aspera entries deferred to Issue 47 carry a 9 November catalog date, and once admitted they move SC-2026-002 and SC-2026-007 to 7 December.
- **SC-2026-010:** Two further reversals under the ratified unit rule re-escalate the rating.
- **SC-2026-011:** Grade A instances at two or more independent operators move the rating to CONFIRMED, gated on the Q4 audit ratifying criterion text v1.1, which Issue 45 characterized as a substantive broadening. Any new countable instance returns the Outlook to Accumulating. The fourteen operator reports change neither figure, because they are the operator already counted.
- **SC-2026-008:** A counted second operation answers the July question without moving the rating. The Storm-3358 candidate is decided at the Issue 47 cutoff on a blind reproduction; it counts when both an artifact ties the observed intrusions to the

framework and the degree of human direction is established, by a source statement or a recovered artifact (CHQ-CORR-2026-041 restates the trigger, which Issue 45 had given as either). Operator, victim and causal-run independence from the founding operation and from Storm-0940 is established on the source; the open question for the audit is whether a second count on the same framework from the same observation source answers the July question or restates it. The Storm-0940 candidate's triggers from Issue 45 stand. January's audit still decides whether the grandfathered rating is supportable.
- **SC-2026-004:** A first confirmed production incident at the required grade escalates the rating.

E.2 Probabilistic forecast register

No entry was issued at this cutoff and none resolved since Issue 45. Fourteen Structural Outlooks and five Forecast Register claims remain open; the ten open entries from the 13 October block resolve 11 January 2027 and the four issued 26 October resolve 26 April 2027. The scorecard stands where Issue 44 left it: two entries resolved, one beating its comparator and one losing to it, a sample that establishes nothing about calibration or skill. The caution stated in Issue 45 about the January block stands: ten entries issued on one day from one evidence base, several sharing a subject, scored individually with the block average reported beside them and no calibration claim made from the block alone. Two hash defects in the register were found and corrected this week, and are disclosed in Section G; no published probability, claim or resolution changed.

F. Institutional Implications

- **Ask what a threat-intelligence vendor should publish before calling an attack autonomous.** Three things would do it: evidence of agent involvement, an account of which operational decisions the agent made across the attack chain, and evidence separating those decisions from human instructions or approvals. One agent-executed command is not evidence of autonomous operational control. No report on the record carries the second or the third. Put the question to the vendors you pay, and read "autonomous" in their reports as a claim until they answer it; those are evidentiary questions for the reporter, not an assumption that defenders hold the attacker's orchestration logs.
- **Decide what an agent-driven intrusion would look like in your own telemetry, and admit what it would not.** Two operations on one open-source agent framework are held on the same gap: the reporter cannot say how much ran without a human. Retried

requests, corrected parameters and machine-speed sequencing are what one reporter cited, and conventional automation does all three. The evidence that separates an agent's loop from a script's is orchestration-side: tool-call records, approval events, the log of a human stepping in. Your network and endpoint telemetry will show automation and its effects, not who decided.

- **Separate "confirmed" from "dated" in your intake.** A catalog entry tells you an authority determined exploitation; it does not tell you the earliest observation of it, and this week no vendor did either. If your prioritization treats a fast listing as evidence of pre-disclosure exploitation, it is reading the agency's publishing speed. Record the earliest substantiated observation as its own field, leave it empty when nobody publishes one, and decide what an empty field means for your response.
- **Act on the fix where the vendor ships it before the description.** SonicWall's fixed build for its Gen 7 firewalls was available from 20 October; the flaw was described on 6 November; the notice announcing the advisory went out on 3 November. A program that takes releases as they ship could have had the fix up to seventeen days before a program that waits for advisories, for that generation; the other generation's fix shipped with the advisory. Ask which of your vendors publish fixes ahead of descriptions and whether your change process can take a release without a reason.
- **Check whether your upgrade path is the attack surface.** The WordPress flaw is exploitable during the upgrade to the fixed release. If your estate upgrades in place on web-accessible hosts, the remediation is the exposure window. Ask what your platform teams do during an upgrade, not just whether they did it.
- **Read the operator's incident index, not the coverage of it.** Fourteen first-party incident reports sat on one operator's public page for up to forty days before CHQ read them, because the fixed address was unreachable from the collection route. If your program tracks agent incidents, put the operators' own indexes on a schedule with a named reader, and treat an unreachable page as a gap to record, not a quiet week.

These are questions for your own program. CHQ does not assess any organization's exposure.

G. Methodology and Material Qualifications

Material qualifications affecting reliance on this issue's ratings are below. Each one's full record, the integrity gate log, the forecast resolution rules and the verification bundle are published at `record.cybersecurityhq.com`; this section discloses, the registry substantiates.

- **Thirteen corrections since Issue 45's records (CHQ-CORR-2026-029 to -041).** Two chronology corrections on rows this issue discusses: the WordPress disclosure date, taken from a secondary source as 29 October where the primaries place it on 28 October, which turned a candidate into an exclusion (-030); and the statement that the Ivanti record post-dated Issue 45's attestation, when it preceded it by 28 minutes (-029). Two late chain-position classifications on catalog rows that sat unclassified past the seven-day window the gate allows, surfaced by the gate itself (-031, -032); the intake fix that would set a provisional position at admission is owed. A forecast-register defect: four Structural Outlooks drafted 26 October shared one placeholder digest where each must carry its own; corrected to per-record hashes, and the immutability check widened to cover the whole issuance specification with a mutation test (-033, CHQ-GATE-FIX-2026-013). A qualified finding: the Friday pass had recorded zero material-incident filings unconditionally, where the evidence supports only that a secondary tracker found none (-034). An exploitation date on the Storm-3358 row that was in fact an infrastructure first-seen date (-035). A reconciliation table published on 9 November with twelve false discrepancies, because it read run identifiers from one column where legacy rows carry them in three others; reproduced to the digit by a reviewer, corrected, and the corrected table verified by a tested script (-036). And one on a post-cutoff row: the GPUStack advisory date was recorded as 4 November, the CVE publication, where the vendor's advisory was published on 31 July and the fix on 18 June, so the interval from advisory to listing is 101 days and not five; the candidate status that rested on the wrong date was withdrawn (-037). And one found while drafting Section C against the retained records: the SonicWall row and two daily notes had read an internal timestamp in the agency's enrichment file as the moment the agency marked exploitation active, seventeen minutes after disclosure, when the file's first recorded commit is the next morning and the stamp's first public availability is not established (-038). Two found while checking an externally produced draft of this issue and its review against the record: a metadata note on fifteen source-registry rows dated the start of typed retrieval states two days late, with the counts themselves right (-039); and the holding-queue row for the LiteLLM flaw stated a responder's identifier as a fact where the CVE record describes a different flaw (-040, Section D). One found by a review of the shortened edition: the counting trigger for the ARTEX candidates had been stated, in Issue 45's Section E and in this issue's first drafts, as either a source statement on human direction or an artifact tying the intrusions to the framework; the criterion needs both, and the artifact alone shows the framework acted, not that it acted unattended (-041, Sections C and E; the Issue 45 sentence is corrected in its post-cutoff replacement text, which is before the publisher). None of the thirteen

changes a count, rating, Outlook or Watch. The cause of the last is structural and is on the audit queue: the registry had no field distinguishing CVE publication from first public disclosure, and the analyst used the one date the retained record supplied.

- **An evidence-handling defect, CHQ's own.** A machine result cited by its content hash in the 9 November correction record was regenerated in place the same day, before its digest was retired, contrary to the rule that hashed evidence is never modified and corrections go in new files. The original bytes are not recoverable; the current file and a byte-identical versioned copy are retained with their digest, and machine results now carry a version suffix and are never rewritten. Recorded against the analyst on the day. It is disclosed here because a record that publishes verification by hash owes the reader the one time it broke its own rule.
- **Intake accounting corrected, with a guard.** A 9 November review showed that the run register's processed-signal count had carried only the mechanical catalog admissions and nothing updated it when the analyst layer admitted roster rows, understating the count on every run with analyst admissions since 20 October. The field is now derived at closure from the signal ledger as mechanical plus analyst admissions, every historical run is reconciled in a typed table with its class (200 of 221 runs agree; fourteen undercounts, three non-numeric values and four unlinked runs from March remain visible as recorded), and closure now refuses if the table and the ledger disagree. Earlier reconciliation rows that contradicted the final one are marked superseded cell by cell rather than deleted. The derivation and the guard are active, tested and unratified (CHQ-RUNSTATE-2026-003).
- **A sequence defect in the two-pass control, repaired and then prevented.** On 9 November an analyst script wrote a Pass 2 start time twenty seconds after the lock and an evaluated time before it, and the gate failed the sequence. The timestamps were repaired to the observed order, the rows are flagged as repaired rather than observed, the gate reports that class in its detail, and a helper that captures the actual start and refuses an impossible sequence was written and used live on 10 November without a flag (CHQ-GATE-FIX-2026-014). Active and unratified.
- **Forty-two rows excluded and not dispositioned, again.** The rows from a September reconciliation remain excluded from every count; their disposition was due at the Issue 45 cutoff, was missed, and was not completed by this one either. The bound from Issue 45 was re-run on this cutoff's artifact (CHQ-SENS-2026-003): under the same four assumptions no Outlook, window count, last-countable date or Watch state changes, the 42 rows all pre-date the window, and the single-event-floor exposure on Agent

Authorization Boundary Failure stands exactly as Issue 45 stated it. The missed date is not a gate condition and whether it should be is before the publisher. The dispositions carry to Issue 47, and this is the second consecutive issue to say so.

- **Pre-disclosure backlog.** Two rows on the pre-disclosure condition have been in shape and unadjudicated since September and October, held on the grade of their first-observation source (community reports in one case, a secondary date for a national CERT observation in the other). They are not candidates this issue adjudicated and they are not counted; they are listed so that the condition's state reads in full: seven counted, three in-shape candidates held, two in shape unadjudicated, five indeterminate.
- **The cutoff convention applied to the four 9 November entries.** The Stability Standard defers declarations made at or after the current cutoff to the next issue. The four entries were released on the cutoff date, after the cutoff artifact was attested, and were admitted the next day; they are deferred. The engine window counts by catalog date and would include them: on the publication artifact the edge window reads twelve and the application-plane window four, with the same ratings, Outlooks and Watch states. The analyst applied the reading; the publisher confirms or corrects it before the Issue 47 cutoff, and the audit queue carries the question. If corrected, the two window figures change and nothing else does.
- **Publication state.** Written from the recorded result of the gate, not from an expected one. The cutoff artifact of 9 November was attested at 16:28:34 UTC with no blocking status and one non-blocking provenance gap on two superseded Issue 42 board snapshots, an accepted risk that closes only on the files or a publisher exception. This issue's board and window figures are bound to that artifact: the board snapshot for Issue 46 was computed from it, its digest is recorded in the registry, and a mechanical check run before ratification compares every rating and Outlook in Section B with the snapshot and every window count, last date and projection in Section E with the engine's reading of the cutoff artifact at the cutoff date, so the figures printed here are checked against the historical artifact and not against the current workbook. The four rows admitted after the cutoff are excluded from that reading by construction. This issue's records were written on the 10 November artifact, which is the publication artifact; its first gate write after the records recorded a store-digest failure because the store had not been rebuilt before the gate ran. That result is kept, the sequence was then run in order, and the artifact cleared with the same single non-blocking status, 54 of 54 negative assertions caught, reconciliation verified, store parity holding and the fixture reproducing. The invariant check between the cutoff artifact and the publication artifact at the cutoff date shows the two window counts and two last-countable dates above, every difference produced by the four deferred rows, and no difference in rating, Outlook or Watch (CHQ-INVARIANT-2026-002). Issue 45 was

submitted for ratification on 3 November and, at the time of writing, is neither ratified nor distributed; two of its sentences became stale after its cutoff and replacement text is before the publisher. This issue assumes Issue 45 is published before it; the publisher decides the order. Lifecycle state: authorized by the gate; awaiting publisher ratification; not yet distributed. Eligibility is not ratification, and if this issue is sent the ratification record and the distribution time are added here.

- **Collection limits.** The European feed was read directly from the collection environment every day from 4 November for the first time since 8 October, and its full exploited list was reconciled daily, so the second catalog no longer depends on equipment outside the environment. The weekday vendor sources still do, or on a workaround. Tuesday's two sources were read, one directly and one through the outside route in the afternoon; Wednesday's four operator sources read one directly, one through the outside route in the afternoon, and two on search fallback because no operator primary exists for them; Thursday read four of six directly through the outside route and two were refused by bot walls; Friday's regulator pass ran on a secondary tracker because the regulator's full-text search is unreachable from the environment; Monday's two sources were read through their own addresses reached from a search index, which returns a condensed extract and not page bytes, and that is stated on the rows. Two primary pages behind this issue's held items, the SonicWall advisory and the operator's incident index, were read only because the publisher's computer was linked; from the environment alone both would stand unread. The rotation-scope decision, retrieve what the sheet declares or declare what can be retrieved, is before the publisher and bit on three days this week. The engine is advisory. One analyst collects and adjudicates, preceded by an internal gate and AI-assisted adversarial review; eight adversarial review passes of this week's artifacts, commissioned by the publisher and AI-assisted, were applied, and every finding accepted is recorded with the part of it that was not. They are review passes, not independent audit opinions, and no independent human audit of this issue has been made.
- **Governance.** No governance change is taken at this cutoff. Three control changes declared on 9 and 10 November are active in the gate and unratified, and are disclosed above as such. The audit queue as it bears on this issue, in full: the evidence standard for pre-disclosure exploitation, whether attempts, compromise or an authority determination satisfy the criterion, with the two-state recommendation; the disclosure clock when a release names a security fix before an advisory describes the flaw; a CVE-publication date field distinct from first public disclosure; whether an operating-system daemon on a named platform family is in the edge condition's declared classes; whether two operators on one shared framework are independent operations; whether a content-management core belongs to the edge or the application plane, added to the edge criterion question;

the cutoff time-of-day convention above; and the items carried from Issue 45, the edge condition's asset-versus-impact and counting-unit questions, the runtime-as-platform scope question, the hypervisor evidence-suppression class question, the agent-authorization criterion text v1.1, the agency evidence standard and the grandfathered rating, the European marker's standing, and the publication-boundary question. The audit is the first issue of January 2027 and the publisher ratifies its outcomes. Governed under Ratings Methodology v1.3, SC-2026-006 criterion v2.0 with C1, SC-2026-008 scope v1.2, Criteria Governance Standard v1.0 and Publication Stability Standard v1.1.

CybersecurityHQ publishes independent structural intelligence for security leadership. Conditions, positions, and falsification criteria are maintained at record.cybersecurityhq.com. Ratings reflect observable structural conditions at a point in time. They are not forecasts and do not assess applicability to any specific organization's environment.