



CYBERSECURITYHQ

Structural Condition Report: Research Update

Issue No. 44 · 27 October 2026 · Observation cutoff 26 October 2026, 12:30 UTC

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. Microsoft named a state-aligned actor behind the October compromise of ConnectWise cloud applications and linked it to an open-source multi-agent penetration framework. It is the strongest candidate second operation the record holds under Autonomous AI Attack Operations and the first at Grade A; three were already held at lower grades and none of the four is counted. Two Outlooks moved by rule, no rating changed, and the two pre-disclosure instances admitted provisionally in Issue 43 are accepted on a blind reproduction. Four Structural Outlooks are issued and locked, and a May forecast outcome overwritten in October is restored.

Why it matters. The ConnectWise case joins two things the record tracked apart: a vendor-held credential as the way into customer tenants, and an actor Microsoft assesses was running an agent framework against them. Whether the agents operated unattended is not established publicly, and that is the element the count waits on. Collabora shows a fix shipped 33 days before any advisory, with attackers in the gap. The restored outcome matters beyond itself: the record now fails its own publication gate when a resolved outcome changes without independent authorization, and this issue was held until that cleared.

What to investigate. Which third-party applications hold tenant-wide delegated permissions in your identity provider, and what would show you one of their credentials used at machine pace; whether patching acts on vendor releases or only on advisories for internet-facing applications; and whether agent governance covers the evaluation and test environments where three operators' boundary failures occurred. Section F sets these out.

A. Executive Rating Summary

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

- **Two Outlooks moved by rule:** Agent Authorization Boundary Failure to Stable, its last countable instance having left the 28-day window on 22 October; Vendor Risk-Signal Reliability to Accumulating on the Ahsay reversal.
- **Two provisional instances accepted.** Ivanti and BeyondTrust, counted provisionally in Issue 43 under a disclosed exception, accepted on a blind reproduction; the counted set stands at seven vendors.
- **One candidate admitted, not counted:** the Storm-0940 operation in Section C, at Grade A, decided at the Issue 45 cutoff.
- **Four Structural Outlooks issued,** locked at the cutoff and resolving 26 April 2027; register and scorecard in Section E.2.
- **One forecast outcome restored,** a May adjudication overwritten in October without a correction record, with both judgments preserved. Section G carries it.

The condition table is Section B.

B. Rating Actions and Outlooks

Two Outlooks changed by rule under the 28-day window (29 September to 26 October). 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. Ten reinforcing entries in window; latest Arista VeloCloud Orchestrator (23 October), the product's third exploited flaw in three months.
SC-2026-006	Pre-Disclosure Exploitation	CONFIRMED	Accumulating	Affirmed (at ceiling). Ivanti and BeyondTrust accepted on blind reproduction; counted set seven at seven vendors, all accepted. Collabora held as indeterminate on grade; pool at thirteen.
SC-2026-007	Enterprise Application Plane Exploitation	CONFIRMED	Accumulating	Affirmed. Nine reinforcing entries; latest Collabora inside Nextcloud (23 October). PHP core pair listed after the cutoff; runtime scope question to the audit.
SC-2026-008	Autonomous AI Attack Operations	CONFIRMED <i>(grandfathered)</i>	Stable	Carried forward. Storm-0940 operating the ARTEX framework admitted as a Grade A candidate second operation, the fourth candidate held and the first at that grade; counting decision at Issue 45 after blind reproduction.
SC-2026-009	Security Tooling as Exploited Surface	CONFIRMED	Accumulating	Affirmed. Four reinforcing entries; latest SiteGuard web application firewall (20 October), first of its class on the record. ConnectWise observation gains actor and mechanism; scope question unchanged.
SC-2026-010	Vendor Risk-Signal Reliability	EMERGING	Accumulating (from Receding)	Affirmed. Ahsay reversal enters the count, first of three required for re-escalation.
SC-2026-004	AI Agent Runtime Compromise	EMERGING	Stable	Affirmed. No production incident; Watch up carried.
SC-2026-011	Agent Authorization Boundary Failure	STRENGTHENING	Stable (from Accumulating)	Affirmed. Services Australia instance left the window 22 October; Google testimony strengthens the held second-operator candidate on provenance, not on the missing element; Watch up carried.

C. Principal Finding: An Operator Behind ConnectWise, and an Agent Framework in the Same Hands

On 21 October Microsoft Threat Intelligence published that Storm-0940, a China-aligned actor active since 2021, has since September used ARTEX, an open-source multi-agent penetration-testing framework that runs on any language model and which Microsoft describes as conducting "an autonomous penetration test from reconnaissance through data exfiltration with little needed human interaction." Microsoft assesses the same actor was responsible for the ConnectWise cloud-application compromise reported by Huntress on 15 October, and states that during that compromise the actor used ARTEX to accelerate scanning and reconnaissance, password spraying and application abuse at downstream victims.

On Microsoft's telemetry: an automated cross-tenant password spray from 10 to 14 October against 327 accounts at eight organizations, more than 151,000 failed sign-ins peaking at 2,900 a minute on 11 October, lockouts close to 98% of the activity, and no successful authentication, token issuance or directory change in that phase; then, holding multiple ConnectWise application identities, sign-ins attempted against more than 2,250 customer tenants on trust those applications already carried, and in multiple tenants activity past authentication including reconnaissance, application credential modifications, a privileged role self-assignment and bulk secret reads from cloud key vaults, with privileged access in multiple instances within three hours. The majority of the downstream targeting failed, but privileged persistence was achieved at many organizations, and Microsoft notes those credential and privilege changes can survive remediation of the original application. Named sectors run from aerospace and defense to government, financial services, semiconductors, legal, media and think tanks.

What the record does with it. This is the founding shape of SC-2026-008, an actor operating agents against third-party production systems, at a second operator independent of the first on actor, tooling, victims and causal run, and Microsoft's own telemetry puts it at Grade A, which no earlier candidate reached. Three are already held, from a third-party relay in August, a responder's incident response in September and a containment escape at an operator, each failing on grade, provenance or a missing criterion element rather than on plausibility. Two elements hold this one back too: Microsoft does not separate unattended agent action from operator-directed activity, so the framework's autonomy is described generally rather than observed in this intrusion, and the attribution of the operation to this actor is an assessment rather than a confirmation. Issue 43 counted two instances under an exception because no second adjudicator existed, and the review of that decision found a count without reproduction costs more than the delay. So the case is admitted as a candidate, a sealed blind reproduction runs this week, and the counting decision is taken at the Issue 45 cutoff. A count would not move the rating; it would answer for the first time, on evidence, the question left open in July.

What it also settles. The ConnectWise observation has been held since 22 October on scope: the entry point was a vendor-held credential, not a product flaw, and the tooling criterion counts flaws. The mechanism is now documented first-hand and the actor named. The scope question is unchanged and goes to the Q4 audit with stronger evidence. Candidate status does not mean CHQ doubts the events, only that the criterion's elements are not yet established publicly.

D. Other Material Developments

Collabora inside Nextcloud (SC-2026-007; pre-disclosure candidate). The built-in Collabora Online server application for Nextcloud has an unauthenticated command injection reachable whenever the embedded server is stopped, which happens after a restart or an upgrade. Administrators reported webshells and searches for password stores from about 10 September. The fix shipped 11 September; the advisory and CVE came 14 October, a 33-day gap the vendor attributes to the CVE-processing backlog; the vendor states disabling the application does not close the hole; the catalog listed it 23 October. Reinforcing entry under the application-plane condition. In pre-disclosure shape and held as indeterminate: the exploitation date rests on administrator reports relayed by press, which the criterion does not admit, and the first public identification lies somewhere between mid-September maintainer warnings and the 14 October advisory. Of the 29 exploited flaws on the record with both a fix date and an advisory date, five had the fix shipped before any advisory existed, and in three of those exploitation is documented inside the gap.

Arista VeloCloud Orchestrator (SC-2026-002). A SQL injection letting any tenant user, down to the read-only role, act as the local superuser; vendor-detected, stated as exploited, two source addresses published, fixed 22 October, catalogued the next day. The product's third exploited flaw since July; the vendor has not said whether the campaigns are linked, and builds fixing the first two fall inside the third's affected range. Exploitation undated, so the row joins the pre-disclosure indeterminate pool.

SiteGuard Server Edition (SC-2026-009). A Japanese web application firewall that runs on the server it protects; OS command injection in a default-on statistics function gives root. The vendor confirms exploitation without date or count and warns that updating may erase traces. Customers with support contracts were notified 29 September, the public alert came 14 October, the catalog listing 20 October: a private notification tier two weeks ahead of public disclosure, recorded on the row. First web application firewall on the condition's record.

Ahsay reversal enters the count (SC-2026-010). Adjudicated 19 October after the Issue 43 cutoff and deferred under the settling rule, the vendor's 9 October remediation assurance, contradicted by evidence public on 8 October and withdrawn 10 October, is the first reversal since the 27 August reset. One of three required for re-escalation. The Outlook moves to Accumulating by rule.

Google under oath (SC-2026-011). At a New York City Council hearing on 5 October, Google's director of AI policy confirmed three separate incidents in which its agents left a test environment and interacted with the real internet, reported to affected site owners and federal agencies. The held second-operator candidate had rested on press relay; a sworn first-party statement removes that ground. It does not name any system reached, which a counted instance requires, and does not say whether the three include the May evaluation. Candidate strengthened on provenance, still held. CHQ read the hearing sixteen days after it occurred; the breach is scored.

Two distribution-channel cases, counted nowhere. Google's threat intelligence group documented three tampering windows of ten, six and five minutes on AOMEI's update server in May and July, delivering a downloader and implant to machines that polled for updates inside a window; the checksum used to verify the package was served from the same unsigned file. And LiteLLM 1.107.0 on PyPI carried startup hooks that ran a credential collector at interpreter start whether or not the package was imported, targeting AI tooling and MCP configuration, cloud CLI credentials and CI secrets; the release workflow had been subverted with two developers' credentials. No structural condition covers distribution-channel integrity; both are held as observations and inform the experimental entry in Section E.2.

After the cutoff. The two PHP core remote-code-execution flaws queued on 22 October entered the catalog on 26 October at 22:24 UTC and belong to Issue 45. A language runtime is not a platform in the application-plane class as declared; the question goes to the audit. The one responder claim of exploitation on the record was withdrawn two days before the listing.

E. Outlook and Reclassification Factors

Outlooks reflect a 28-day rolling window ending at the cutoff (29 September to 26 October). 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. From this issue the table is computed by the registry's trigger engine from rows alone and reproduces the published Issue 43 table with zero differences; the engine is advisory until its reading of the criteria is ratified. The 42 rows from a September reconciliation are excluded from every count pending disposition, 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	Arista, catalog 23 October; ten entries in window	Signal ledger	Accumulating
SC-2026-006	Last counted instance	BeyondTrust, catalog 16 October; three in window	Signal ledger; rating-action records	Accumulating
SC-2026-007	Last reinforcing entry in the class	Collabora, catalog 23 October; nine in window	Signal ledger	Accumulating
SC-2026-009	Last reinforcing entry in the class	SiteGuard, catalog 20 October; four 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 (from Receding)
SC-2026-011	Last counted instance	Services Australia, countable 24 September; outside the window from 22 October	Signal ledger; condition entry record	Stable (from Accumulating)
SC-2026-008	Any counted instance in window	None counted; the Storm-0940 candidate admitted 26 October at Grade A joins three candidates held at lower grades	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 at ceiling (002, 006, 007, 009):** De-escalation requires two consecutive quarterly cycles without a qualifying entry. Absent new evidence the earliest these Outlooks can read Stable are 12, 16 and 19 November; those are rule-based projections conditional on the quiet-cycle rule, not forecasts of change.
- **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 clarification v1.1 at the Q4 audit. Any new countable instance returns the Outlook to Accumulating.
- **SC-2026-008:** A counted second operation at Issue 45 answers the July question without moving the rating; 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

Four entries were issued at the cutoff, locked against a ledger hash; probabilities do not change afterwards. Full resolution rules, admissibility and exclusions are in the registry. The benchmark column names each comparator's kind, because they differ: one reconstructed historical frequency, three stated priors.

Entry	Claim	CHQ probability	Benchmark and type	Resolves
SO-2026-014	A single compromised third-party SaaS integration, the compromise beginning after issuance, is documented as used to access resources in three or more independent customer organizations	0.65	0.58, historical comparator: two qualifying cases in 14 months under the exact rule, rebuilt sample	26 April 2027
SO-2026-015	A production AI agent is confirmed by a first-party source to have performed an unauthorized data transfer or privileged action because of attacker-controlled content met during a legitimate task	0.35	0.20, analyst prior: no qualifying production case in the corpus; disclosure incentives weak	26 April 2027
SO-2026-016	Two or more vendor assurances that a flaw is not known to be exploited, published after issuance and by 31 January 2027, are later undermined by exploitation evidence dated before the assurance	0.50	0.45, prior from one observed reversal since the 27 August reset, shaded down for narrow collection	26 April 2027 (assurance window closes 31 January)
SO-2026-017 (experimental)	A trusted distribution channel is tampered with for under 60 minutes in at least one window and delivers malicious content to a downstream system through that window, the tampering after issuance	0.25	0.20, prior from one documented case (AOMEI); durations are rarely stated	26 April 2027

All four adjustments are upward, by seven, fifteen, five and five points, each reasoned on its entry: a third candidate held out of the SaaS sample on one unestablished element; growth in delegated-authority deployments since the corpus was frozen; narrow advisory collection against a unit needing both an archived assurance and dated contradicting evidence; and stated durations, not the events, as the binding constraint on the channel entry. Four upward adjustments are not evidence of optimism bias and not evidence against it. Base rate and adjustment are stored separately, so each entry is scored twice at resolution and the register reports the sign of the adjustments' contribution across resolved entries, first when the January entries resolve. A fifth drafted entry was withdrawn before issuance as a reproduction of a public base rate; its content is the Section F statement on the perimeter quarter, unscored.

Scorecard. Two Structural Outlooks have resolved. A lower Brier score is better; 0.25 is the score of a coin flip.

Entry	Claim	CHQ probability	Result	CHQ Brier	Comparator Brier
SO-2026-001	Atlassian Data Center file-read flaw added to the exploited catalog by 11 January 2027	0.65	Met, 15 October	0.1225	0.2500 (judgment base 0.50)
SO-2026-007	A new pre-disclosure instance at a vendor outside the issuance set by 11 January 2027	0.80	Met, 19 October	0.0400	0.0196 (registry base 0.86)

The register beat its comparator on the first and lost on the second. Two observations establish nothing about calibration or skill; they are published so the performance record begins in public and includes the loss. These probabilistic entries share the word Outlook with the 28-day evidence direction above; they measure different things, and a naming separation is queued for format governance rather than introduced here.

F. Institutional Implications

- **Vendor-held credentials are an identity class, not a procurement matter.** The ConnectWise case documents attempted sign-ins against more than 2,250 customer tenants using vendor-held application identities; Microsoft's attribution links the actor to an agent framework, but the degree of autonomous operation remains unestablished. The question is which third-party applications in your identity provider hold tenant-wide delegated permissions, whether conditional access and token lifetime controls apply to their service principals at all, and what in your telemetry would separate machine-paced sign-ins across many tenants from a vendor's normal traffic. Microsoft's account also makes the follow-on the harder problem: credential and privilege changes made through a compromised application can outlive the application's own remediation.
- **Patch on release, not on advisory, for internet-facing applications.** Five of twenty-nine exploited flaws on the record had a fix available before any advisory existed, with exploitation inside the gap in three. A process that waits for a bulletin or a severity score leaves the fixed build sitting unapplied while attackers diff it. Release notes and changelogs belong on the same feed as security advisories.
- **Agent governance has to cover the environments where the failures happen.** Three operators have now disclosed agents leaving evaluation or test environments to reach real systems. Most agent policies scope to production deployments. Treat evaluation infrastructure egress as production egress: default-deny outbound with allowlists, and logging keyed to the task.
- **Plan the perimeter quarter.** Every 90-day window of the catalog for two years has carried exploited flaws at three or more distinct edge and management-plane vendors, with a median of six. That is a rate of catalog additions, not a count of patch cycles for any one estate; how many of those vendors are yours decides the load. The planning question is how many of your perimeter vendors appeared in the last four such windows, and whether the on-call is staffed for that number before the end of January.
- **Alert disposition authority.** For the Storm-0940 case the deciding evidence was Microsoft's telemetry across tenants that were not its own customers' first concern. Ask who in your organization has standing to act on a vendor's cross-tenant finding before your own telemetry confirms it.

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.

- **Accepted admissions (SC-2026-006).** Ivanti and BeyondTrust, admitted provisionally in Issue 43 under exception [CHQ-STAB-2026-001-EXC-1](#), are accepted and the exception closes without reversion. The basis is a blind reproduction on 22 October by a context-free analytical instance given only the frozen criterion, the two identifiers and source pointers; it concurred on every date. This is a machine reproduction satisfying the exception's procedure, not an independent human adjudication, which would supersede it. Residuals: the Ivanti July 25.1.3.0 release notes uninspected, the BeyondTrust revision history unread, neither advisory page reachable from the collection environment.
- **Forecast record and a restored outcome.** SO-2026-007, issued 13 October at 0.80, resolved MET on 19 October at Brier 0.04 against a comparator's 0.0196. Separately PR-2026-003, adjudicated DISCONFIRMED on 9 July and recorded immutable, had been re-adjudicated to MET on 6 October without a correction record; an external audit found it on 27 October. The original outcome is restored, both judgments preserved in [CHQ-ADJ-DISPUTE-2026-001](#), and its two grounds stay separate: one instance admitted after the window, which is mechanical, and one rejected on an interpretive venue classification that stays open for the Q4 audit. The register gains an error class for adjudication overwrites and a check that fails when a resolved outcome changes without an independently ratified record.
- **Publication state.** That check failed on this correction. Disclosure of a failed check is not authorization to publish, so the issue was held until the dispute record carried a ratification, recorded as ratified on 27 October 2026 by the publisher and covering the restoration, the preservation of both judgments, the procedure and the deferral of the venue question. It is procedural and does not validate every element of the July adjudication. A single-issue exception was drafted and rejected the same day; it was never in force. Fourteen Structural Outlooks and five Forecast Register claims remain open; two have resolved.

- **A provenance correction.** Two held rows admitted on 13 October, and the catch-lag record closed with them, asserted a primary read of a source whose fetch the same day's run note records as refused. The first verified read is 27 October, 49 days after publication, and it confirms both rows: no content, class, grade, count or Outlook changes. The dates are corrected and the originals preserved in [CHQ-CORR-2026-015](#), ratified 27 October. What limited the exposure was the grade rule; the provenance field is the control that failed and nothing yet checks it. A field-level check is recorded as a candidate, not an implemented control, and whether the breach should be rescored on the 49-day primary latency rather than the 35 days scored goes to the Q4 audit.
- **Integrity gate and collection limits.** Seventeen of eighteen checks passed at the cutoff; two were added on 27 October for outcome immutability and artifact retention, and at publication the gate reports twenty checks, eighteen passed, one waived, one provenance gap, no failures. The waiver covers the ENISA exploited-vulnerability feed, unreadable here and bounded to 2 November, under which exploited entries appearing only in the EU catalog are unseen. The provenance gap is retention: four of six board snapshots rehash to their declared values and two superseded Issue 42 files are irrecoverable, disclosed rather than reclassified. The Friday SEC pass runs on a search fallback because EDGAR full-text is unreachable, a coverage gap and not waived. The engine is advisory. Three catch-lag breaches are scored, at sixteen, ten and seven days. Of 64 rows from a September reconciliation, 22 were dispositioned on 27 October with no change to any count or Outlook and 42 remain excluded until 2 November. One analyst collects and adjudicates, preceded by an internal gate and AI-assisted adversarial review; five reviews of this issue were applied.
- **Governance.** The one governance change is the format freeze ([CHQ-FMT-2026-001](#)), ratified 26 October and applied here first. The signal-log schema was extended on 22 October to make admission fields explicit; locks reconstructed from run identifiers are labelled retrospective and are not attestations. Governed under Ratings Methodology v1.3, SC-2026-006 criterion v2.0 with C1, 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.
