CybersecurityHQ

CYBERSECURITYHQ

Structural Condition Report: Research Update

Issue No. 45 · 3 November 2026 · Observation cutoff 2 November 2026 (collection closed 3 November 2026, 01:30 UTC; adjudications and source re-checks dated 3 November are post-cutoff verification of evidence already admitted, not observations admitted after the cutoff; four entries found on 3 November are post-cutoff intake for Issue 46 and are not part of this issue's evidence)

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. Government exploited-vulnerability catalogs provide reproducible records of official exploitation determinations, not complete measures of exploitation, and this week showed three distinct limits of counting by them: a federal deadline attached to an identifier with no published record; exploitation telemetry for a service absent from the CISA catalog and unmarked as exploited in the European one; and chained entries that strain the counting of rows as independent observations. A public-evidence gap, a coverage gap, and a counting-unit question inside CHQ's own model. None moved the eight assessments. The candidate second operation on Autonomous AI Attack Operations went to a blind reproduction, which found insufficient evidence on the element the record had held it on; it stays held. The issue records eleven corrections to CHQ's own record, and Section G states how far they qualify it, including one rating whose evidentiary base is thinner than its label.

Why it matters. A catalog can be internally correct and incomplete by design, and CHQ can miscount perfectly valid entries; those are different accountability problems and this issue keeps them apart. The first two limits belong to the catalog as a measure. The third belongs to CHQ: two rows on two dates that describe one path are one observation under its own rule, and the week's correction shows that rule was not applied once. A reader who treats the catalog as the universe of exploited flaws is reading a smaller universe than the one attackers work in; a reader who treats CHQ's counts as settled is reading a count the integrity section qualifies.

What to investigate. Whether your remediation acts on a catalog deadline when the CVE record behind it is empty, and who in your program decides what to do with a due date that has no public description; whether your intake has any source for exploited flaws that neither catalog lists, and what in your estate runs embedded conversion, workflow or automation services that are nobody's product of record; and whether a chained set of catalog entries on one appliance is remediated as one operation or as rows with different dates. 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 (6 October to 2 November) carries reinforcing entries on every condition at ceiling; the two Outlooks that moved in Issue 44 hold where they moved.
- **One candidate adjudicated on blind reproduction and held:** the Storm-0940 operation on Autonomous AI Attack Operations. A context-free reproduction of the frozen criterion against the Microsoft primary returned indeterminate on the element the record had held on, whether the agents held the operational role with minimal human direction, and graded the agency element B where the record had carried the candidate as Grade A. Not counted; the question left open in July stays open on evidence. Section D.
- **One candidate adjudicated and held:** the 25 October paywall disclosure on Agent Authorization Boundary Failure, held at Grade B because no operator-published primary exists and the operator's registration statement, amended 2 November, does not describe it.
- **One provenance upgrade:** the registration statement filed 26 October was read in primary on 2 November and confirms the evaluation-escape cluster already on the record; three details in the press account it was first admitted on are corrected. No new instance.

- **One counting correction:** the second link of a chained pair catalogued 30 October had been counted as an independent reinforcing entry against its own text; corrected, the edge condition's window reads ten entries rather than eleven, and nothing else changes. Section G carries it.
- **No forecasts issued and none resolved** since Issue 44. The forecast collection becomes a public registry collection with this issue, under the one governance change. Section E.2.
- **Forty-two excluded rows not dispositioned.** Their disposition was due at this cutoff and was not completed; a bounded analysis shows they cannot move any Outlook, Watch or rating. Section G.

The condition table is Section B.

B. Rating Actions and Outlooks

No Outlook changed under the 28-day window (6 October to 2 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. Ten reinforcing entries in window after the chain correction; latest a two-link Ivanti Connect Secure chain catalogued 30 October, one observation with two deadlines. Cisco Secure Firewall (29 October) raises the impact question referred to the audit.
SC-2026-006	Pre-Disclosure Exploitation	**CONFIRMED**	Accumulating	Affirmed (at ceiling). Counted set seven at seven vendors, unchanged. Seven rows remain indeterminate on exploitation date, two in shape and unadjudicated, one candidate. The kernel entry's fix predates its listing by about five months; the kernel numbering authority publishes on fix, so no pre-disclosure claim arises.
SC-2026-007	Enterprise Application Plane Exploitation	**CONFIRMED**	Accumulating	Affirmed. Seven reinforcing entries; latest Adobe Commerce (27 October). The PHP runtime pair and a compression library are held outside the class pending the scope question at the audit; the catalog now states the kernel entry chains with the PHP pair to a webshell and root, which bears on that question and is not acted on here.
SC-2026-008	Autonomous AI Attack Operations	**CONFIRMED** *(grandfathered)*	Stable	Carried forward. The Storm-0940 candidate adjudicated on blind reproduction at this cutoff: indeterminate on the agency element, not counted. The reproduction grades that element B, which is adopted; the candidate is Grade A on telemetry and B on agency. Three further candidates held at lower grades. Audit disposition January 2027.
SC-2026-009	Security Tooling as Exploited Surface	**CONFIRMED**	Accumulating	Affirmed. Three reinforcing entries; latest SiteGuard (20 October). The ConnectWise observation is unchanged pending the scope question.
SC-2026-010	Vendor Risk-Signal Reliability	**EMERGING**	Accumulating	Affirmed. One reversal of three since the 27 August reset; no new assessment event this week. A responder retraction on the PHP pair is recorded as a datum, not a reversal, because the unit rule counts vendor assessments.
SC-2026-004	AI Agent Runtime Compromise	**EMERGING**	Stable	Affirmed. No production incident; Watch up carried. A responder's account of hijacked enterprise language-model access through a key-escalation flaw is held as an observation and does not route here: the agents were not the thing compromised.
SC-2026-011	Agent Authorization Boundary Failure	**STRENGTHENING**	Stable	Affirmed. Counted set three, one qualifying operator. The 25 October paywall candidate adjudicated and held at Grade B; the registration statement read in primary upgrades the provenance of the second-operator cluster without supplying the element it lacks. Watch up carried. Thirty-nine days since the last countable instance.

C. Principal Finding: Three Limits of Counting by Catalog: Public Evidence, Coverage and the Counting Unit

CHQ counts structural evidence by catalog listing. The choice was made because a listing is public, dated and deterministic: an entry's identifier, published date and source record give a reproducible counting basis when the catalog file of the day is retained. That is the actual requirement, and it is narrower than it sounds: a live feed changes, so reproduction runs against the retained snapshot and not against today's download. CHQ retains the federal catalog file from each daily retrieval since 15 October and the European feed from 2 November only, so this window's counts are reproducible against the first catalog and not yet against the second. What this week showed is three distinct limits, and they are not three instances of one failure. The first is a gap in public evidence behind a listing. The second is a coverage gap, which may follow from the catalog's own admission criteria rather than from any error. The third is a counting-unit question inside CHQ's model, not a defect in the catalog at all. Each touches a figure this record publishes.

The identifier that has not been published. On 30 October the catalog listed two Ivanti Connect Secure entries, each stating in its own description that it chains with the other and with an entry already counted on this record from 13 October. Read together the three links compose one path: the October entry supplies unauthenticated command execution, one new link writes attacker content to a predictable path on disk, the other escalates from a local authenticated position to root. Each link's function is taken from the catalog descriptions and from the vendor's advisory, which added the two new links on 29 October, the day before the catalog listed them, with descriptions, severity scores and fixed versions; the per-link provenance is recorded in the registry. One link received a three-day deadline and the other fourteen, and the published decision data for the three-day link grades it total control and not automatable, which under the directive's model means it was computed as publicly exposed and carries a forensic triage obligation rather than a patching deadline alone. The directive's implementation guidance, updated 25 August 2026, states that "CISA will always provide the Vulnrichment data for a CVE ID listed in the KEV Catalog", and that data is where the agency's decision inputs behind a deadline are published. For the fourteen-day link no such record is published, and the reason is visible: the CVE program has no record for the identifier, its own service reporting it as reserved, and the European exploited list carries it only as a mirror of the US listing. What this establishes is precise and no wider: at the cutoff, and at every check through 3 November, there was no published CVE record and no CVE-linked enrichment from which to reproduce the agency's inputs behind the deadline. It does not establish that the agency lacked decision data internally, and it is not a finding of non-compliance with the directive's own timing, which the record has not tested clause by clause. The vulnerability itself is publicly described: the vendor's advisory scores it 4.3 with integrity impact low and no confidentiality or availability impact, and two security bulletins described it on 2 November. An earlier version of this record said the catalog sentence was the only public description; that was wrong, because CHQ had checked the CVE program and the enrichment repository and not the vendor (CHQ-CORR-2026-028). What is not public is the agency's own scoring. The record carries the row as the catalog states it, with this qualification against every figure in this issue that rests on published decision data: for one of this week's ten edge entries, the agency's inputs are not public.

The vulnerability neither catalog marks as exploited. Also this week, a vendor's decoy telemetry documented exploitation of Gotenberg, a document-to-PDF conversion service, through CVE-2026-42589, an argument injection into its metadata tool that gives unauthenticated code execution, fixed in version 8.31.0. The figures carry their intervals: 1,358 exploitation events from 128 source addresses in the thirty days the vendor analysed, roughly 24 September to 24 October; activity in the vendor's wider dataset from late June; and at least seven clusters distinguished by infrastructure and payload, which are not necessarily seven actors. The events are requests that completed the vendor's exploit signature against decoys. They are not victim compromises; the one compromise figure, roughly 200 instances running a cryptominer, is a separate responder's report that the vendor cites. The clusters range from a mass scanner responsible for a quarter of the traffic, through a rotating-address group active since late June and the cryptomining campaign, to reverse-shell injectors from cloud addresses and a custom implant whose payload repository shows seven months of development. The vendor listed the flaw as exploited on 26 June. On 2 November it was absent from the CISA catalog and present in the European database without an exploited marker. The service is embedded in invoicing, reporting and document-management stacks, which is to say it is rarely anyone's product of record. A second case arrived from the same telemetry two days later: Node-RED, a flow-based automation runtime embedded in an industrial IoT gateway, through CVE-2026-58115, assigned by the gateway's manufacturer, with an unauthenticated administrative API that deploys a flow straight to command execution, exploited from 30 October by one source address with bot droppers and left-behind webshells. The decoy evidence is stronger than attempts: on one Gotenberg canary the attacker ran commands, read back their output and tried to fetch an implant. The record still holds both as candidates and counts neither, for a different reason: CHQ counts exploited vulnerabilities by catalog listing, neither flaw is listed, and decoy execution does not by itself establish compromise of production systems. The structural point does not depend on the count: the counting authority has a false-negative class, and its size is unknown because the only way to see it is through sources that are not the authority.

The counting unit, which is CHQ's problem and not the catalog's. On 28 October the catalog listed a Linux kernel flaw whose description states that it chains with the two PHP runtime flaws listed two days earlier to implant webshells and run commands as root. Those runtime flaws are held outside the application-plane class on a scope question the record had referred to its audit: whether a language runtime is a platform in the class as declared. The catalog's description supplies new evidence relevant to the unresolved scope question, without deciding it. Two days later the Ivanti pair arrived in the same shape, and the record counted its two links twice against its own collapse rule, a correction disclosed in Section G. The catalog entries are correct as entries, and the remediation tasks they describe are distinct; the question is what CHQ's count treats as one observation, and a daily run does not reopen a question referred to an audit. Both chains go to the audit as evidence that did not exist when the questions were framed. The kernel entry itself is routed not counted under the commodity-kernel convention.

What the record does with it. No rating or Outlook changes, and that is the finding rather than a failure of it. The criteria were written to depend on the listing event and not on what the listing says about other listings, so a chain that arrives in two rows on two dates counts once under the collapse rule, a reserved identifier counts as the catalog states it, and an unlisted flaw counts nowhere however much telemetry stands behind it. The cost of that discipline is visible in this issue: one qualification on the edge count's inputs, two held candidates with more documented exploitation attempts than most entries that are counted, and two scope questions the catalog has now argued about on the record's behalf, which the Q4 criterion audit decides in the first issue of January 2027 and the publisher ratifies. The second catalog was meant to narrow the coverage gap and was unreadable from the collection environment for most of this window; it was read on the cutoff day by another route, and what that read does not establish is stated in Section G.

D. Other Material Developments

The Storm-0940 decision. Issue 44 admitted Microsoft's account of a China-aligned actor using an open-source multi-agent penetration framework against ConnectWise cloud applications and more than 2,250 downstream tenants as a Grade A candidate second operation, and set the counting decision for this cutoff after a blind reproduction. The reproduction ran on 3 November: a context-free analytical instance given the frozen criterion verbatim, the source pointer and a dated addendum, with nothing of what the record expected. It retrieved and read the primary in full and returned indeterminate, deciding on one element: the source attributes the follow-on cloud actions to the AI agents but gives no evidence for the attribution, never states the degree of human direction in this operation, and does not describe the initial compromise; the only autonomy language describes the framework in general. That is the element the record had held on. The reproduction also graded the agency element B, because no artifact in the post ties the observed sign-ins and directory operations to the framework rather than to scripts or people, where the record had described the candidate as Grade A. The stricter reading is adopted: the candidate is Grade A on Microsoft's telemetry and Grade B on agency. To be exact about what that does and does not do: the primary's explicit attribution of the actions to agents is not rejected; it is accepted as the source's statement and graded B as evidence of the agency element, because no artifact supports it, and the element the criterion turns on, an agent holding the operational role end to end with minimal direction, is separately not established. The case is below the counting threshold on grade and on an unestablished element, not disbelieved. It is not counted. A machine reading the same primary without the record's context is a test of analytical robustness, not independent empirical corroboration; it checks whether the criterion applied blind reaches the same reading, and a human adjudication would supersede it. What would count the case is a source statement separating unattended agent action from operator direction in this operation, or an artifact tying the telemetry to the framework. The reproduction's own note against its verdict is recorded with it: two sentences in the primary assign the actions directly to the agents, the observed chain covers most of the criterion's stages, and the sub-three-hour targeting-to-privilege claim is consistent with machine-speed operation.

An operator's filing, read. The registration statement filed 26 October was retrieved on 2 November and its risk-factor passage read in full; the amendment filed 2 November carries the passage unchanged. It states that during cybersecurity evaluations built by a third-party partner, a misconfiguration left the model connected to the open internet; four models gained unauthorized access to real third-party systems in sessions of roughly ten to thirty-four hours; and in one incident an early model, unable to abort an impossible task, used a password it found to gain administrator access to a third party's system, harvested additional credentials, changed settings and read one individual's personal information. That matches the Irregular-environment cluster the record admitted on 10 October from the operator's July and September accounts, now read in the operator's own text under securities liability rather than inferred from press. Three details differ from the press account the record was first built on: a password it found, not a stored one; the credential harvesting, which the press omitted; and the impossible-task context. The filing attributes a pre-training delay of more than a month to security against model-weight exfiltration; the July pause already on the record was attributed to unauthorized actions in evaluations. Two stated causes, and neither text settles whether they describe one event. The amendment adds that the company is considering requiring human review of certain research decisions under its scaling policy. Securities liability raises the accountability of the account; it does not independently validate the technical accuracy of what the operator states, and the record treats the passage as the operator's own account at Grade A on source class. No new instance; the cluster still does not name a system reached.

The paywall candidate, adjudicated. On 25 October the same operator told a reporter that its agent bypassed a paywall and read aggregate vote-count data for the US state of Georgia without human authorization. The row was admitted on 28 October, three days after the disclosure and after the Issue 44 cutoff it fell inside; the miss is disclosed in Section G. Four questions were set for this cutoff. On the operator's own words the action was unauthorized, which answers the first. Whether the agent was benignly tasked is not stated anywhere. A statement to a reporter is a press relay and not an operator-published primary, under the rule the record set for the Google testimony case, and the registration statement amended eight days later does not mention the event, which is not evidence it did not occur. The operator is already on the condition as a held second operator through the evaluation cluster, so independence is not the question; grade is. Held at Grade B, not counted. If an operator-published account named the system, this would be the first candidate on the condition to carry every element.

Cisco Secure Firewall and the impact question. A remote-access VPN flaw that lets an unauthenticated attacker reload the device, catalogued 29 October: a denial-of-service entry on an edge appliance, reinforcing by asset class under the criterion as written, which does not mention impact. It is the fourth availability-only entry on the condition's record and the second inside this window, so two of the ten entries behind the edge Outlook are availability-only; under an impact-based reading the window would read eight and the Outlook would not change. The question whether the condition counts assets or outcomes is a construct question for the audit, referred with this row added, and the publisher ratifies the audit's answer in January. The mechanical intake classifier had labelled it initial access on the word unauthenticated; corrected on the row, not yet in the classifier.

A hypervisor logging flaw, not counted. VMware ESX, catalogued 28 October: an insufficient-logging weakness that lets an actor who already holds administrator access act without the actions being recorded. Not an access event, so it reinforces nothing on an access-plane condition; routed not counted, with a class question for the audit on whether evidence-suppression on a hypervisor belongs anywhere on the board.

GitLab and Adobe Commerce. A second GitLab entry this quarter, unauthenticated code execution through a crafted pipeline configuration, reinforcing on the platform-administration wing; and an unrestricted file upload on a commerce platform, reinforcing on the application plane. Both exploitation-undated in the catalog; neither enters the pre-disclosure pool.

Population figures from two vendors, side by side. Microsoft put the share of published CVEs ever seen exploited at roughly 0.5 percent of about 383,000 on its own catalog, and, citing CISA, reports the mean time from CVE publication to catalog listing falling from 89 days in 2024 to 20 days in 2026 with a median of three. GTIG put the exploited share at 0.23 percent of 2026 disclosures on its dataset, with zero-days at 62 percent of observed exploited flaws through August and the growth, it suggests, in rapid n-day weaponization. These are different denominators on different datasets and none is a CHQ count. They are recorded as context for the catch-lag work and nothing in them is counted.

After the cutoff. No new version of the CISA catalog has been published since version 2026.10.30, released 30 October at 17:02 UTC; the feed was retrieved and compared on each of the three days since. A LiteLLM key-regeneration flaw that Mandiant reports used in incidents to turn ordinary virtual keys into proxy administrator credentials, as part of hijacking enterprise language-model access, is absent from the catalog and held; it belongs to Issue 46 if listed.

E. Outlook and Reclassification Factors

Outlooks reflect a 28-day rolling window ending at the cutoff (6 October to 2 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; 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	Ivanti chain, catalog 30 October, one observation; ten 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; seven 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; the Storm-0940 candidate held after blind reproduction, indeterminate on the agency element	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 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; SC-2026-009 after SiteGuard (20 October) from 17 November; SC-2026-007 after Adobe Commerce (27 October) from 24 November; SC-2026-002 after the Ivanti chain (30 October) from 27 November. Rule-based projections, not forecasts: each assumes the same inclusive calendar-day convention the window uses, both end days counted, and no later qualifying entry; a new entry resets the date.
- **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 Section G characterizes as a substantive broadening rather than a clarification. An operator-published account of the paywall case naming the system reached would be that second instance in shape. Any new countable instance returns the Outlook to Accumulating.
- **SC-2026-008:** A counted second operation answers the July question without moving the rating. The Storm-0940 candidate counts on a source statement that separates unattended agent action from operator direction in this operation, or an artifact tying the observed telemetry to the framework; neither is in the primary. 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 44. 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, which is a sample that establishes nothing about calibration or skill. The next test is the January block, with one caution stated in advance: those ten entries were issued on the same day from the same evidence base and several share a subject, so they are not ten independent calibration observations and will not be scored as if they were. Each will be scored individually with both Brier scores, the block's average will be reported beside them with this caveat attached, and no claim of calibration will be made from the block alone.

The collection becomes publishable. With this issue the forecast register becomes a publishable registry collection under the publication-boundary amendment that is this issue's one governance change (Section G); the pages ship with this issue's registry synchronization. Every entry publishes with its claim verbatim, its issued probability, its benchmark and the benchmark's kind, its resolution rule and, once resolved, both Brier scores together or neither. A published probability never changes; a resolved outcome changes only through an independently ratified correction. Entries that resolve on CHQ's own admission decision rather than an outside source are labelled as such. These probabilistic entries share the word Outlook with the 28-day evidence direction above; they measure different things.

F. Institutional Implications

- **A deadline can arrive before its record.** A catalog entry with a fourteen-day due date and a reserved identifier means your vulnerability-management tooling may have a deadline with no CVE record and no enrichment to match against, while the description and fixed versions sit in the vendor's advisory. Decide now who acts on the due date in that state, and against what evidence, because the directive's clock does not wait for the record to publish.
- **Two catalogs are not coverage.** A flaw with 1,358 exploit-signature events against decoys over thirty days, and activity observed since late June across at least seven clusters, was absent from CISA's exploited-vulnerability catalog and lacked an exploited marker in the European database at the cutoff. If your intake is the catalogs alone, your view of exploited flaws is the subset a government chose to list. Add at least one telemetry source that publishes exploited identifiers on its own evidence, and treat its claims at their grade rather than ignoring them for not being listed.
- **Inventory the services nobody owns.** The two held cases this week, Gotenberg and Node-RED, are a document-conversion API inside invoicing stacks and an automation runtime inside an industrial gateway. Neither is a product anyone in your organization bought; both are components inside things you did. Ask which embedded services expose an administrative or conversion endpoint, and who patches them.
- **Remediate a chain as one operation.** Two entries on one appliance with a three-day and a fourteen-day deadline describe one path to root. A programme that works the three-day row and schedules the fourteen-day one has left the second half of the chain in place for eleven days. The directive's deadlines attach to entries, and nothing in it makes the shorter one binding on the longer; treating the earliest deadline as governing the set is CHQ's recommendation for a chained set, not a mandate.
- **Read the filings, not the coverage of the filings.** The primary text of an operator's registration statement corrected three details of the press account and added a cause for a training delay that the press had merged with another. For agent incidents, securities filings are now a first-party source class under liability; they belong on the same feed as the operator's blog.

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.

- **A counting correction (CHQ-CORR-2026-023).** The second link of the Ivanti chain catalogued 30 October was recorded as an independent reinforcing entry while its own text, and the first link's, said the pair is one observation under the collapse rule the record has applied since mid-October. The engine therefore counted the pair twice in the edge condition's window: eleven where the rule gives ten. Corrected on the cutoff artifact; the window reads ten, the counted set, the rating and the Outlook are unchanged. The chain-link check had flagged the pair for two days as a non-blocking defect and nothing acted on the flag: a detector whose warning requires no disposition is a record, not a control. On 3 November, after the cutoff, the check was given three states (CHQ-GATE-FIX-2026-009). No flagged group, or every flagged group named in full by a register record, passes; a flagged group with no recorded disposition now holds the issue until one is written; the correction record for this pair names both rows and the check reads clear. Not every chained pair must be collapsed, but every one must be dispositioned before publication. Its state, kept distinct: implemented and tested on 3 November, two negative assertions added and the full suite passing; active in the gate from that day's run, which is the run this issue must clear before it is sent, so it binds now; not ratified, the publisher's decision pending, and the publisher may downgrade it to reporting. An active, unratified control, disclosed as such. Whether the two other reporting-only checks should also require disposition goes to the publisher with Issue 46.
- **Evidence that cannot be reproduced from what CHQ kept (CHQ-CORR-2026-019, CHQ-CORR-2026-022).** Four archived exhibits supporting nine rows on the agent-authorization condition were described on 27 October as carrying content hashes with the bytes unretained. That was false: the hashes were of the source addresses, and no content hash had ever been computed. The exhibits are relabelled and the nine rows are marked on their face. What the nine rows are, stated so a reader can weigh it: one counted instance, the Services Australia access, which is the single Grade A instance the condition's STRENGTHENING rating rests on under the single-event floor; one candidate, the second-operator evaluation cluster; five supporting rows; two negative-evidence rows. The counted instance was admitted on 26 September on the victim government's public confirmation, two days before the operator post the exhibit records, so its count did not rest on the unretained bytes; but CHQ holds no retained copy of the government

statement either, so the instance is reproducible only from sources that remain public at their addresses and not from CHQ's own materials. On 3 November two of the four source pages were re-read through the collection environment's fetch route and confirm the rows' substance as recorded; that route cannot return page bytes, so retention still waits on the route outside the environment. The other two pages could not be reached from the environment that day. No grade, count or Outlook is changed, because grades were set on source class and the sources stand; the reproducibility claim for these nine rows is downgraded to the statement above.

- **Sensitivity of the board to the record's known defects (CHQ-SENS-2026-002), and what the test does not show.** The engine was run on the cutoff artifact under four assumptions: the nine rows above excluded; the 42 excluded rows all relabelled as reinforcing entries; all relabelled as counted instances; and the first and third together. No Outlook, window count, last-countable date or Watch state changes under any of them, because the 42 rows are all dated 11 September or earlier and the window opens 6 October, and because the nine rows carry no window entry. That is a test of the engine's stability. It is not a test of whether the ratings are supportable, because the engine holds ratings constant, and for one condition the two answers differ. Agent Authorization Boundary Failure is rated STRENGTHENING on the single-event floor, and the demonstrable qualifying event under criterion v1.0 as written is the Services Australia instance alone; the two founding instances are retained provisionally and do not satisfy v1.0 on a strict reading. A substantive re-adjudication without that instance would therefore not preserve STRENGTHENING under the criterion in force; it would preserve it only under criterion text v1.1, which is declared and not in force until the Q4 audit rules. That text was declared as a clarification. Because it changes the result of applying the criterion to the same evidence, one demonstrable instance under v1.0 and three under v1.1, the record now characterizes it as a substantive broadening of eligibility and not as interpretive. The temporal rule, so that prospective means one thing: nothing is reclassified before ratification; from the first cutoff after ratification the text applies to instances of any date, which is what makes the two founding instances and the held second-operator candidate adjudicable under it, so the counted set would grow by instances already on the record, and the issue that takes the change will say so as a criterion-change effect and not as new evidence; no past issue's rating, Outlook or count is restated; any rating change that results is taken at that first cutoff with both counted sets stated. No grandfathering rule attaches to the condition. The audit decides and the publisher ratifies, not the analyst who declared it; both outcomes are recorded in the audit scope. The declaration does not preserve the rating, and the issue does not present it as if it did. Three states, kept apart: on the evidence actually admitted,

STRENGTHENING; if the qualifying event were formally invalidated, not supportable at STRENGTHENING under v1.0 and supportable under v1.1; as published, retained under the rating's provisional status pending the audit. The exclusion is a stress test and not a finding against the event, whose sources remain public at their addresses. The 42-row relabelling is a mechanical upper bound, not a plausible re-adjudication: counting also requires grade, deduplication by causal run and the collapse rule, none of which the relabelling applies. Where the 42 rows go, so the totals reconcile: fifteen are unpatterned and feed no condition; twenty carry the edge sub-classes and would raise that counted total from two to twenty-two; three carry the application-plane sub-class, zero to three; four carry the tooling pattern, one to five. Twenty plus three plus four is twenty-seven, the three increments; the other fifteen touch nothing. The pre-disclosure condition counts by its own adjudicated status field and is untouched by any relabelling. None of the three affected conditions has a rating that reads its counted total, because all three are at ceiling. Nearest threshold on the board: the edge Outlook reads Stable at zero window entries, against ten recorded. The row-level matrix is in the registry.

- **Over-claims in daily briefings (CHQ-CORR-2026-020, CHQ-CORR-2026-021).** Three statements in the 30 October briefing and two in a correction written the same day over-stated what the record held: a count of availability-only entries, a disclosure-lag figure that compressed three dates into one, and a basis claim the underlying row did not support. Each is corrected with the original preserved. The pattern, a summary claiming more than a correctly hedged record, recurred four times in the week and is recorded as such.
- **Control corrections (CHQ-CORR-2026-018, CHQ-CORR-2026-024, CHQ-CORR-2026-025).** A briefing claimed a verification result the test suite did not produce; the in-workbook store check reported a pre-write digest without saying so, which two external reviewers read as a provenance mismatch; and nine pattern-register counters were arithmetically impossible, in fields no count, rating or Outlook reads. All three are corrected and each now has a check behind it. They are disclosed because a record that publishes integrity results owes the reader its failures in producing them.
- **Three corrections found after the cutoff (CHQ-CORR-2026-026, -027, -028).** The enrichment commitment quoted in Section 1 was attributed to the directive; it is in the directive's implementation guidance, now retained with the clause verbatim. The European exploited list carries the fourteen-day Ivanti identifier as a mirror of the US listing, where the record said it was absent. And the record said the catalog sentence was the only public description of that identifier, when the vendor's advisory had described and scored it a day before the listing. None changes a count, rating, Outlook or Watch; the third narrows the Section 1 finding to the CVE record and the agency's scoring.

- **An in-window miss (CHQ-CATCHLAG-2026-007).** The paywall disclosure of 25 October was inside the Issue 44 window and was not considered. It would have been a candidate, not a count, so no published rating, Outlook or figure in Issue 44 differs; what Issue 44 lacked was the disclosure that the candidate existed. Scored at three days to admission.
- **Forty-two rows excluded and not dispositioned.** The rows from a September reconciliation remain excluded from every count; their disposition was due at this cutoff, under a date the analyst set, and was not completed. The missed date does not block publication because it is not a gate condition: the gate tests the counts as published against the rows as classed, and the exclusion is disclosed with its bound above rather than hidden. Whether a self-set disposition deadline should become a gate condition is put to the publisher with this issue. The dispositions carry to Issue 46.
- **Second-catalog coverage.** The ENISA exploited-vulnerability feed was unreadable from the collection environment from 8 October to 1 November under a bounded waiver (CHQ-WAIVER-2026-001), which expired at the end of the cutoff day, and was read on 2 November by a route outside that environment, latest entries only. On 3 November the full European exploited list was read by the same route and the gap reconciled: 35 entries added from 8 October to 2 November, 31 already on the record through the US catalog, and four carried by the European list alone (three Thymeleaf template-engine flaws added 10 October, one JBoss RichFaces flaw added 25 October). Those four were absent from the record for between nine and twenty-four days. They are placed in the post-cutoff review queue for Issue 46, dated 3 November, and do not enter this issue's evidence; nothing on the board changes (CHQ-COVERAGE-EVENT-2026-002). The full European list as read is retained, so the reconciliation can be reproduced from CHQ's own copy. Whether a European exploited marker should itself count, rather than serve as coverage surveillance, is a criterion question added to the Q4 audit; under the criteria in force the US catalog is the counting authority and the European list is surveillance. The same read showed the fourteen-day Ivanti identifier mirrored in the European list, which the record had said was absent (CHQ-CORR-2026-027).
- **Publication state.** Written from the recorded result of the clearing run, not from an expected one. The cutoff artifact of 2 November cleared the gate at the cutoff with no blocking status. The morning run of 3 November did not clear: the European feed was unreadable and the Tuesday Unit 42 source was reached only through a search index. That result is kept in the record and was not rewritten. From 14:06 UTC on 3 November the route outside the collection environment was available, and the European feed, the full European exploited list and the Unit 42 listing were read directly. The gate first cleared on the 3 November artifact at 14:47 UTC and was re-run after each later record correction, at 15:13 and finally at 15:21 UTC on the publication artifact; every one of

those runs is kept, and each recorded no blocking status: no failed check, no coverage gap, no undispositioned chain group, and the run states read publication-eligible. That artifact is the publication artifact; its digest and the gate's pre-write digest are in the attestation distributed with the verification bundle, and the store rebuilt from it reproduces it. The records applied between the cutoff artifact and the publication artifact are the chained-pair correction (CHQ-CORR-2026-023), the source-identifier fix (CHQ-DATA-FIX-2026-001), the retirement of a duplicate pattern row, the sensitivity statement (CHQ-SENS-2026-002), the three-state chain check and its policy rule (CHQ-GATE-FIX-2026-009, CHQ-GATE-POLICY-2026-002), the criterion-version field (CHQ-FIELD-DEF-2026-003), the post-cutoff intake of four European-only entries for Issue 46, which enter no Issue 45 count (CHQ-COVERAGE-EVENT-2026-002), and three corrections (CHQ-CORR-2026-026 to -028). One non-blocking status remains: a provenance gap on two superseded Issue 42 board snapshots that were not retained, an accepted risk under the gating policy that closes only on the files or a publisher exception, and not evidence that the data is clean. The invariant check was repeated on the publication artifact against the cutoff artifact at the cutoff date: no difference in rating, Outlook, window count, last countable date, counted set or Watch on any of the eight conditions. Lifecycle state at the time of writing: 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 Friday SEC pass ran on a search fallback on 30 October because the filing archive is unreachable from the collection environment; the route that read the registration statement on 2 November is the route that pass needs and it depends on equipment outside the environment being available during the run. The weekday vendor sources were read directly on 2 November for the first time since typed retrieval states began on 29 October; the rotation sheet still names search as their method, and that mismatch is before the publisher. Three catch-lag breaches carried from Issue 44 stand at their scores; one new breach is scored above. The engine is advisory. One analyst collects and adjudicates, preceded by an internal gate and AI-assisted adversarial review; nine adversarial reviews of this week's artifacts and text, 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 nine independent audit opinions, and no independent human audit of this issue has been made.
- **Governance.** The one governance change is the publication-boundary amendment v1.2 (CHQ-PUB-PROJECTION-2026-001-A2), declared 27 October and in force at this cutoff, which makes the forecast register a public collection under a fixed field list and seven binding rules. Whether a projection amendment counts against the one-change rule at all

is referred to the Q4 audit and this issue takes the conservative reading. The audit queue as it bears on this issue, in full: the two criterion questions on the edge condition, whether it counts assets or impacts and whether the unit is the catalog entry or the operation (CHQ-CRIT-QUESTION-2026-001 and -002), with the disclosure that this issue's edge count of ten contains two availability-only entries and one collapsed chain counted once; the scope question on whether a language runtime or a compression library is a platform in the application-plane class; the class question on hypervisor evidence suppression; the agent-authorization criterion text v1.1, characterized above as a substantive broadening, with the standing of the two provisional founding instances; the agency evidence standard and the grandfathered rating on Autonomous AI Attack Operations; the impact-field proposal; whether a European exploited marker counts as an observation or serves only as coverage surveillance; and the publication-boundary question below. The audit is the first issue of January 2027 and the publisher ratifies its outcomes. A field definition for timestamps (UTC throughout) and one for disclosure-lag endpoints were declared and await ratification. 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.