CERTIFIED

Information Security™

CIS Assessment System

Solution Brief & User Guide

Version 2.5.42.34.1 | May 2026

Compliance, Cyber, AI Risk, & ERM Assessment Platform

ISO/IEC 42001:2023 | ISO/IEC 27001:2022 | NIST AI RMF 1.0 | NIST CSF 2.0 | ISO 31000 ERM Family

Certified Information Security

www.certifiedinfosec.com

Document series: User Guide -- published at each governed release

v2.5.42.34.1 -- Review-Submit P0 Hotfix

v2.5.42.34 shipped on 2026-05-26 and resolved the long-standing review-submit 404 regression on Linux production servers. Within minutes of installation, anonymous review submission triggered an HTTP 500 response. The failure was distinct from the v34 PSR-4 filename issue: the underlying cause was a PHP 8.x trait property collision at class-load time.

v2.5.42.34.1 closes that collision. Anonymous assessors who submit a completed assessment for peer review will now reach the submission confirmation flow reliably in v2.5.42.34.1.

What happened. The review-submit controller declared a property named $inFlightTasks that also exists in the ComponentAccessLevelTrait it uses. PHP 8.x requires that when a class and a trait both declare the same property, the type, visibility, and default value must be identical. The controller declaration differed from the trait declaration, causing PHP to raise a compile-time fatal error at the moment any request reached that controller class -- before any controller method ran. The error surfaced only on the Linux production server because Windows development environments do not exhibit this class-composition behavior at runtime the same way.

The fix. The controller property redeclaration block was removed. The specific task entry that the controller needed in the allowlist is now injected at runtime in the execute() method prelude, before the access-level enforcement call. The trait remains unchanged. This approach avoids a library version bump and confines the change to the reviews extension.

Carry items to v35. Two P3-priority issues identified after v34 installation remain open and deferred to v35 as cosmetic or non-blocking items:

The star-rating widget on the review form highlights only the third star when a reviewer clicks the third star, rather than highlighting stars one through three as expected. This is a display-only issue; the numeric rating value submitted is correct.

On the results page, browsers with credential autofill enabled will occasionally pre-populate the Assessed Organization input field with a saved credential. This is a browser-side behavior; no assessment data is leaked or altered.

A P4 cosmetic: the post-install Joomla banner reports "lib_cisassess v2.5.42.26 present" regardless of the actual installed version. This is a hardcoded string in the install script noted at v33 and carrying as a second sighting; the fix is a single-string change deferred to v35 as non-functional.

v2.5.42.34 -- Review-Submit Flow Restored

v2.5.42.34 closes a regression that had been silently preventing the peer-review submission flow from completing on production Linux servers for approximately 40 releases. The root cause was a file-naming case mismatch introduced when the review-submit controller was first authored: the controller filename used mixed-case (ReviewSubmitController.php) where the Joomla PSR-4 autoloader on Linux requires the second word to be lowercased (ReviewsubmitController.php). Windows NTFS is case-insensitive, which means every development-time and gate-triple check passed cleanly while the identical code silently produced 404 errors on the production server.

(c) 2026 Certified Information Security. All rights reserved.

v2.5.42.34 corrects the filename, the class declaration, and all internal URL references that pointed to the parent component instead of the reviews extension. Assessors who submit a completed assessment for peer review will now reach the submission confirmation flow reliably.

This release also ships the Phase 1 scaffold for a new automated regression-prevention gate that will catch this class of case-sensitivity mismatch before any future code reaches production. That gate is described in the Admin Guide; from an assessor perspective, it represents a structural reliability improvement with no visible change to the assessment experience.

1. What the CIS Assessment System is

The CIS Assessment System is a free, public-facing tool that measures an organization's current state against five internationally recognized risk and security frameworks: ISO 27001, ISO 42001, ISO 31000 ERM Family, NIST CSF 2.0, and NIST AI RMF 1.0. It runs as a Joomla 5 component at https://www.certifiedinfosec.com/assessments and requires no software installation on the assessor machine.

The system exists to serve a specific conversion thesis: equip practitioners to discover their own control gaps with precision, then connect them to the executive training that closes those gaps. Every architectural and content decision in this tool serves that thesis.

2. Who this guide is for

This guide addresses the full range of people who use or evaluate the CIS Assessment System. Senior managers and risk officers need to understand what the tool produces and why. Assessors run the tool directly against a framework. AI and cyber executives evaluate fit against their organization's governance needs. Training students arrive at the tool as a prerequisite to or companion for CIS executive training programs.

The Admin Guide (Technical Design) addresses build engineers, integrators, and deployment operators. Its content is not repeated here.

3. Assessment experience

When an assessor opens the system, they select one of the five active frameworks. The tool presents every control or subcategory in that framework with a scored response interface. For conformance-mode frameworks (ISO 27001, ISO 42001), assessors select among Major Non-Conformance, Minor Non-Conformance, Opportunity for Improvement, and Satisfactory Conformance. For maturity-mode frameworks (ISO 31000 ERM Family, NIST CSF 2.0, NIST AI RMF), they select among the maturity levels defined by the framework scoring model.

The tool derives a T1 top-level score from the distribution of T2 control responses. A conformance-ceiling rule prevents a high T1 rating from coexisting with a floor of materially low T2 scores -- the composite score reflects the realistic aggregate, not the optimistic average alone. An uneven-maturity banner alerts the assessor when the spread across controls is large enough to warrant independent attention.

(c) 2026 Certified Information Security. All rights reserved.

At any point during the assessment, the right-rail navigator shows which controls remain incomplete. Assessors can save progress and return; the session-lifecycle heartbeat maintains connection state and warns before the session expires. All previously generated assessment records and saved progress persist across sessions.

4. Scoring integrity

The CIS Assessment System implements a two-tier scoring architecture that enforces defensible ratings at every level of the assessment hierarchy.

Tier 2 -- per-question scoring. Each control or subcategory in the selected framework receives a discrete Tier 2 (T2) response from the assessor. For conformance-mode frameworks (ISO 27001, ISO 42001), the response set follows the four-level conformance ladder defined by ISO/IEC 17021-1:2015: Satisfactory Conformance, Opportunity for Improvement, Minor Non-Conformance, and Major Non-Conformance. For maturity-mode frameworks (ISO 31000 ERM Family, NIST CSF 2.0, NIST AI RMF), the response set follows the maturity-level ladder defined by the applicable framework scoring model, with levels ranging from 0 (non-existent or ad hoc) through 4 (optimizing).

Tier 1 -- control-level scoring and the weakest-link cap. The Tier 1 (T1) score represents the overall conformance or maturity for a grouping or for the framework as a whole. The scoring engine computes T1 from the distribution of T2 responses subject to two structural constraints: a weakest-link ceiling cap and a residual-risk floor.

The weakest-link ceiling cap reflects the ISO/IEC 17021-1:2015 section 9.4.8 calibration principle: a T1 rating cannot be asserted above the level that the minimum T2 score across all constituent controls supports. In conformance mode, any control scored at the lowest response level represents a systemic absence of a defined process; that absence prohibits a favorable T1 conformance assertion regardless of how well the remaining controls perform. In maturity mode, the CMMI institutionalization requirement applies: a control cannot be claimed at maturity level N if any constituent practice scores below N. The practical effect is that a single underperforming control constrains the T1 ceiling, reflecting the realistic aggregate rather than the optimistic average.

The residual-risk floor ensures that T1 scores do not descend below the level that the assessed evidence set supports. Override-justification entries that accompany any assessor-initiated deviation from the computed score are captured in the audit trail per ISO/IEC 17021-1 section 9.4.8 and NIST SP 800-53A Rev. 5 documentation expectations.

Scoring authority. The single source of scoring mathematics is the shared framework scoring model at the component data layer. All five frameworks reference this shared model. No per-framework ceiling override exists. The ceiling tables derive their threshold values directly from the cited authority standards: ISO/IEC 17021-1:2015 section 9.4.8 for conformance-mode ceilings and the CMMI institutionalization discipline for maturity-mode ceilings. Calibration decisions are grounded in the literal text of those standards, not in approximations or operational convenience.

Uneven-maturity detection. When the spread of T2 scores across controls is large enough to indicate material risk concentration in specific areas, the tool surfaces an uneven-maturity banner to alert the assessor. This signal prompts independent attention to the control clusters driving the spread rather than relying solely on the T1 composite.

(c) 2026 Certified Information Security. All rights reserved.

This scoring model architecture is framework-agnostic. The same two-tier engine, weakest-link cap, and floor-rule mechanism operate identically across all five supported frameworks. Assessors working in conformance mode and assessors working in maturity mode receive the same structural integrity guarantees from the underlying engine.

5. Frameworks covered

ISO 27001 -- Information Security Management System. Annex A of ISO/IEC 27001:2022 provides 93 controls across four themes: organizational controls, people controls, physical controls, and technological controls. The CIS Assessment System implements all 93 controls with per-control evidence guidance drawn from ISO 27001:2022 Annex A directly.

ISO 42001 -- AI Management System. ISO/IEC 42001:2023 Annex A provides 38 controls across ten domains, covering AI system design, data governance, risk management, and responsible AI objectives. The assessment aligns to the normative control inventory in Table A.1 of the published standard.

ISO 31000 ERM Family. The ISO 31000 Enterprise Risk Management family is delivered as a unified composite framework drawing from ISO 31000:2018 (risk management principles and guidelines), ISO 23894 (AI risk guidance), ISO 31010 (risk assessment techniques), and ISO 27005:2022 (information security risk management). The composite presents 23 clauses covering principles, framework, and process across all four source standards.

NIST CSF 2.0 -- Cybersecurity Framework. NIST CSF 2.0 provides six Functions (Govern, Identify, Protect, Detect, Respond, Recover), 22 Categories, and 106 Subcategories with 363 Implementation Examples. The CIS Assessment System covers all 106 Subcategories.

NIST AI RMF 1.0 -- AI Risk Management Framework. The NIST AI Risk Management Framework provides 72 subcategories across four Functions: Govern, Map, Measure, and Manage. Each subcategory carries a complete AI Actors attribution derived from NIST AI 100-1 Appendix A (pp. 35-36), identifying which practitioner roles bear responsibility for each practice area. This attribution is load-bearing for assessors who need to assign accountability within their organizations. The assessment includes Suggested Actions drawn from NIST AI 600-1 where the T2 maturity average meets the configured threshold.

As of v2.5.42.25, topic chips are suppressed across all Tier 1 headers for every framework. The display surface for each T1 header now presents the scored response interface and, where applicable, AI Actors attribution only. This change applies across all five frameworks and is future-proof: no framework JSON change will re-introduce chip rendering.

6. Output: PDF assessment report

When the assessor concludes their work, the tool generates a PDF assessment report. The report includes a cover page showing the framework, assessed organization name, assessor role, and assessment date. The body delivers:

- A T1 conformance or maturity summary with color-coded status band
- Per-grouping distribution charts showing the spread of T2 responses
- Detailed per-control records with the assessor selected response and any evidence notes

(c) 2026 Certified Information Security. All rights reserved.

- A Suggested Actions section (maturity-mode frameworks) ordered by priority
- An Evidence Companion Guide cross-referencing the control inventory against evidence categories
- A Training appendix with course options relevant to the framework assessed

Reports are downloadable as PDF files and optionally emailed to the assessor's registered address. The filename prefix is framework-specific, derived from the report.filePrefix field in the framework JSON.

7. Bot protection and submission security

The CIS Assessment System applies a multi-layer bot-protection stack to all high-risk public-facing submission endpoints. The stack combines Cloudflare Turnstile invisible challenge verification, a honeypot field, Accept-Language header fingerprinting, and an IP-based rate limiter. These layers operate server-side and are transparent to legitimate assessors who complete the form in a standard browser.

The Turnstile verification layer is fail-closed when the Turnstile secret key is not configured by the site administrator. A missing or empty secret causes the server to reject the submission and log a warning. Assessors who encounter an unexplained submission rejection while using a legitimate browser should contact the site administrator, who will find the diagnostic entry in the Joomla system log under the com_cisassess category.

8. Peer review system -- CIS Assessment Reviews

The CIS Assessment System supports a peer review workflow through a companion extension, CIS Assessment Reviews (com_cisassess_reviews). This extension is a separately installable Joomla component and carries a hard dependency on the parent com_cisassess component: com_cisassess_reviews requires com_cisassess to be installed and cannot function independently. If the reviews extension is not installed, the assessment tool degrades gracefully -- the review-collection panel on the results page presents an empty state and no functionality is impaired.

From the assessor's perspective, the peer review workflow allows a completed assessment to be submitted for structured peer commentary. As of v2.5.42.34.1, the review-submit endpoint is fully restored: the v34 PSR-4 case fix and the v34.1 trait-collision fix together close the regression chain that had been producing errors on the Linux production server. Assessors who previously encountered a page-not-found or server-error response when attempting to submit a review should find the workflow operating correctly in v2.5.42.34.1.

Review records are managed by site administrators through a dedicated Reviews panel in the Joomla administrator interface. Administrative actions (approve, reject, mark as spam, add a note) operate exclusively through the child extension and do not require direct access to the parent component's administrator panel.

(c) 2026 Certified Information Security. All rights reserved.

9. Assessment-to-training pathway

The CIS Assessment System surfaces, with precision, the specific control gaps that CIS executive training programs address. That connection is not incidental -- it is the point of the tool. An assessor who scores Major Non-Conformance on ISO 42001 Clause A.6 AI system data management, or who rates NIST AI RMF GOVERN 1.2 at Maturity Level 1, leaves the tool with documentary evidence of where they stand and a clear path to the training that closes the gap.

CIS executive training programs cover ISO 31000 enterprise risk management, ISO 42001 AI management, NIST AI Risk Management Framework, NIST CSF 2.0, and ISO 27001 information security management. The training is positioned at the executive practitioner level -- practitioners who need to build and lead these programs, not merely be aware of them.

Each framework Training appendix in the PDF report contains card-level descriptions of the relevant training programs and direct enrollment contact information.

10. Portable Core -- What it means for assessors

The v2.5.42.31 Portable Core Phase 1 release extracted all training-panel, certification-card, and post-assessment email content from the installer script into a structured, schema-validated JSON data layer. v2.5.42.32 completed that work by correcting the content loader so the postflight seeding step runs to completion.

For assessors, no visible behavior changes in v2.5.42.34.1 beyond the review-submit flow restoration. The assessment experience, scoring model, report output, and framework coverage are identical to v2.5.42.32. Sites currently on v2.5.42.32 or v2.5.42.34 may install v2.5.42.34.1 as a direct upgrade; the postflight seeding step remains idempotent and safe to re-run over any prior clean installation.

(c) 2026 Certified Information Security. All rights reserved.