CSAI Foundation | Cloud Security Alliance

AICM v1.1: What Actually Changed for Model Security

Correcting a Misconception as CSA Expands the AI Controls Matrix

2026-07-30

AI-assisted Rapid Research

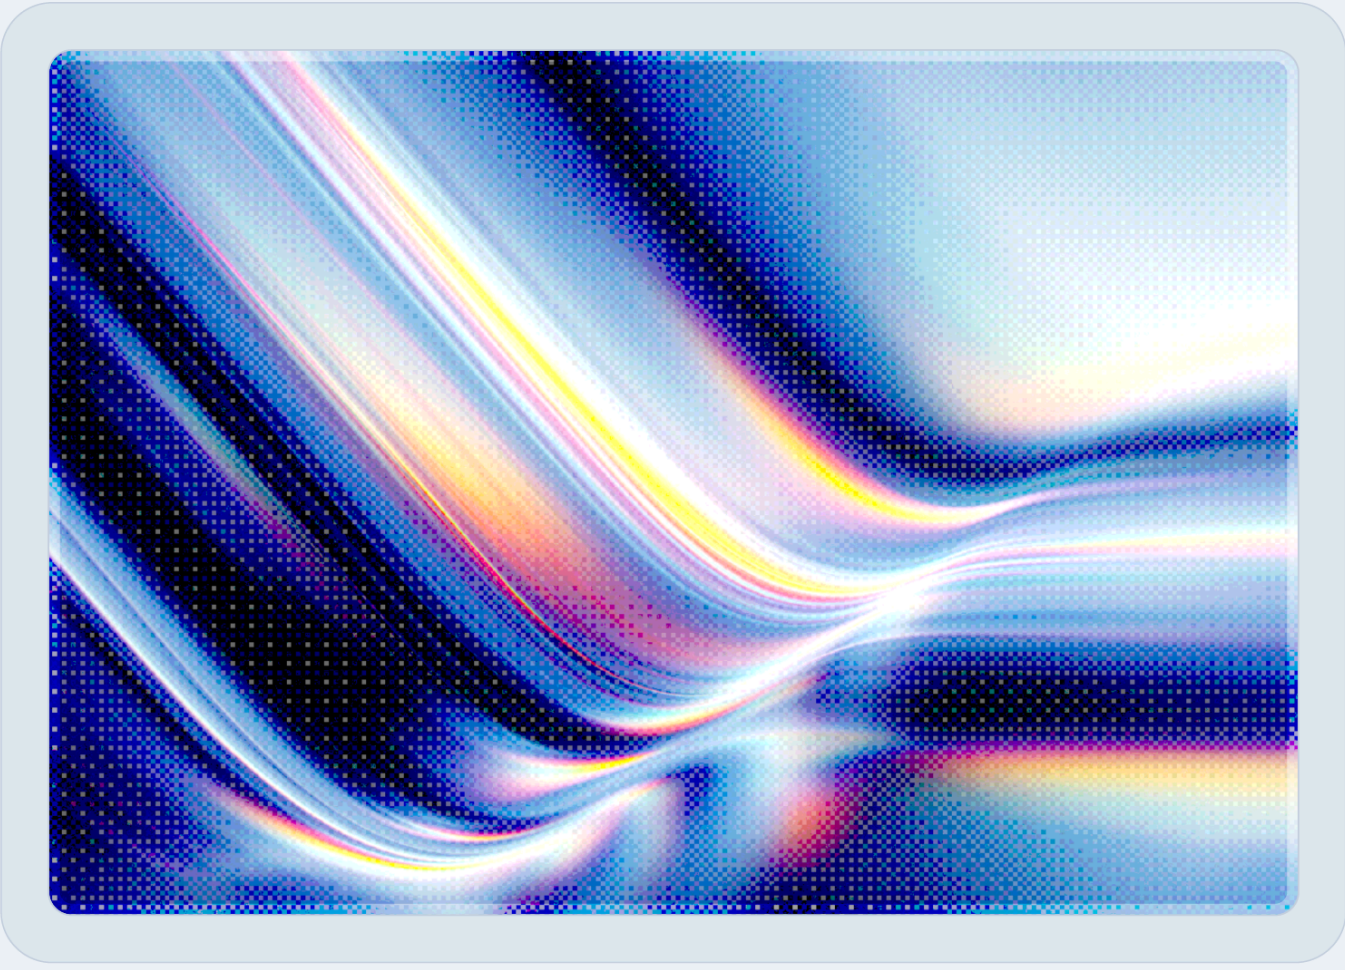

CSAI

cloud security alliance®

Key Takeaways

The Cloud Security Alliance published version 1.1 of the AI Controls Matrix (AICM) in mid-2026, expanding the framework from 243 to 247 control objectives while retaining its 18-domain structure [1] [2]. Early third-party commentary — including at least one governance-focused blog and one virtualization-industry trade publication — has described the update as introducing a "new" Model Security domain, but that characterization is inaccurate: the Model Security (MDS) domain, and its 13 controls, has been part of the AICM since the framework's original release in July 2025 [3]. What v1.1 actually changes is narrower and, for adopters who already built programs around AICM v1.0, considerably easier to absorb: refined control language synchronized with the newly released Cloud Controls Matrix (CCM) v4.1, an expanded set of external framework mappings including a first mapping to AIUC-1, and a larger AI Consensus Assessments Initiative Questionnaire (AI-CAIQ) for self- and third-party assessment [1][4][5]. Organizations that already operationalized AICM v1.0's Model Security controls do not need to stand up a new control domain; they need to reconcile their existing MDS implementation against refined wording and confirm their mapping references still resolve correctly. In this update cycle, the more consequential risk for adopters is arguably not a governance gap in model security but a documentation and vendor-communication gap — one created by inaccurate secondary reporting that could lead security teams to over-scope their v1.1 transition work or, conversely, under-scope the mapping and assessment changes that are genuinely new.

Background

The AI Controls Matrix has occupied a specific niche in CSA's assurance portfolio since its introduction: a vendor-neutral, role-based control framework purpose-built for AI systems, distinct from the Cloud Controls Matrix (CCM) that governs general cloud security. When CSA released AICM v1.0 in July 2025, the framework was explicitly structured as an extension of the CCM lineage — 17 of its 18 domains map back to CCM domains familiar to any organization already running a cloud security assurance program, while a single domain, Model Security, was built from scratch to address risks that have no direct analogue in traditional cloud security [3]. That domain covers concerns specific to machine learning models themselves: protecting training data and model architecture from tampering, securing model

weights against theft or unauthorized access, implementing guardrails against misuse, and hardening inference pipelines against manipulation such as prompt injection. AICM v1.0 assigned this domain 13 controls, numbered MDS-01 through MDS-13, out of the framework's 243 total control objectives.

AICM's architecture also introduced a Shared Security Responsibility Model (SSRM) tailored to the AI supply chain, defining five distinct actor roles — Cloud Service Providers, Model Providers, Orchestrated Service Providers, Application Providers, and AI Customers — each with role-specific implementation and auditing guidelines layered on top of the shared control set. This structure matters directly to the Model Security domain, since responsibility for MDS controls concentrates most heavily with Model Providers, who train and serve the underlying models, while AI Customers and Application Providers carry a narrower slice tied to how they consume and configure those models.

CSA has continued to iterate on both halves of this framework in parallel. The Cloud Controls Matrix itself moved to version 4.1 in January 2026 (announced December 2025), restructuring to 17 domains and 207 controls with updated mappings to contemporary cloud regulatory and technical requirements [4]. AICM v1.1, detailed in a CSA blog post published July 14, 2026 by Principal Research Analyst Marina Bregkou, synchronizes the AI framework's CCM-derived domains with that CCM v4.1 baseline while carrying the Model Security domain forward largely intact [1][2].

Security Analysis

The Model Security Domain: Established, Not New

The clearest fact-check an AICM adopter needs from this release cycle concerns the origin of the Model Security domain. Independent secondary coverage published around the v1.1 release — including at least one governance-focused blog that characterized the update as introducing a "new MDS domain" — has conflated the domain's designation as new *relative to the Cloud Controls Matrix* with being new *in version 1.1* [6][7]. Both CSA's original July 2025 announcement of AICM v1.0 and the framework's introductory guidance documentation describe Model Security as one of AICM's original AI-specific domains, present from the framework's first release, comprising 13 controls that CSA's own materials have consistently numbered MDS-01 through MDS-13 [2][3]. Comparisons of the domain's control count between v1.0 and v1.1 sources show it holding steady at 13, with no controls added, removed, or renumbered [1][2].

This distinction is not merely academic. An organization that reads "new Model Security domain" and concludes it must build model-specific governance controls from a standing start risks misallocating remediation effort and could overstate its compliance gap to its board or regulators. An organization that already implemented AICM v1.0's MDS controls, by contrast, has a considerably lighter lift: confirming

that its existing control implementations still satisfy the v1.1 control language, and updating any compliance mapping documents that reference the older AICM v1.0 mapping tables. Security and GRC teams evaluating AICM v1.1 for the first time should treat Model Security as a domain with roughly a year of implementation precedent and CSA guidance behind it, not an emerging requirement introduced this cycle.

What v1.1 Actually Changes

The genuine substance of the v1.1 update sits in three areas, none of which is the addition of a new domain. First, the total control count rose from 243 to 247, with CSA describing the additional controls as refinements synchronized to the newly released CCM v4.1 baseline rather than as new AI-specific requirements [1]. Because the Model Security domain's control count is unchanged, these additions almost certainly landed in the framework's CCM-derived domains, reflecting how CCM v4.1's own restructuring flows through to the shared portions of AICM.

Second, CSA expanded the framework's external mapping coverage. AICM v1.1 adds a mapping to AIUC-1, an emerging security standard focused on agentic AI systems, alongside continued and updated mappings to the EU AI Act, the NIST AI Risk Management Framework, ISO/IEC 42001:2023, and Germany's BSI AIC4 [1][2]. For organizations that use AICM as a control-mapping bridge into multiple regulatory and certification regimes simultaneously — a common pattern for multinational AI providers — this is arguably the update's most consequential change for multinational AI providers specifically, since it extends the framework's utility as a single control set that can be cross-walked into several compliance obligations at once rather than requiring separate gap assessments for each.

Third, the AI-CAIQ questionnaire used for self-assessment and third-party evaluation grew to 320 questions aligned with the 247-control matrix [1][2]. Organizations that have already published an AI-CAIQ response tied to AICM v1.0 will need to plan a refresh cycle, since the questionnaire structure and numbering have moved with the underlying control changes; that refresh is an assessment-management task, not evidence that the organization's underlying model security posture has fallen out of compliance.

Why the Confusion Matters for Adopters

This kind of pattern — a genuine, incremental framework update accompanied by secondary reporting that overstates its novelty — is common with fast-moving frameworks generally, but it carries particular weight for a framework still in its second year of adoption. AICM does not yet have the multi-year institutional familiarity that the Cloud Controls Matrix has built up since its own introduction, which plausibly increases early adopters' reliance on vendor blogs, analyst summaries, and social commentary

rather than CSA's release notes directly. When that secondary layer introduces an error — describing an established domain as newly created — it risks propagating into vendor risk questionnaires, internal audit checklists, or procurement language before CSA can issue a correction. Security teams evaluating AICM v1.1 should treat CSA's own artifact pages and blog posts as the authoritative source for what changed release-over-release, and should be specifically skeptical of any secondary source that frames Model Security, or any other CCM-derived domain, as an entirely new addition without citing a specific control count change to support that claim.

Recommendations

Immediate Actions

Organizations already assessed against AICM v1.0 should pull the current AICM v1.1 control spreadsheet and run a domain-by-domain diff against their existing control mapping, focusing first on the CCM-derived domains where the four additional controls were introduced, since those are the domains most likely to require new evidence or policy updates. Teams should also verify that any vendor risk questionnaires, internal audit templates, or compliance dashboards that cite AICM control IDs are pointing at v1.1 identifiers rather than the superseded v1.0 numbering, particularly for Model Security controls where the ID scheme (MDS-01 through MDS-13) has not changed but surrounding domain references may have shifted.

Short-Term Mitigations

Model Providers, who carry the heaviest share of Model Security domain responsibility under AICM's shared responsibility model, should schedule a refresh of their AI-CAIQ response using the 320-question v1.1 questionnaire rather than attempting to patch an existing v1.0-era response, since question numbering and structure have changed [1]. Organizations planning to use AICM as a bridge into EU AI Act or ISO/IEC 42001 conformance work should review the newly completed mappings before finalizing any parallel gap assessment against those regimes directly, as the AICM mapping may already cover ground that would otherwise require duplicate analysis. Any internal or vendor-facing communication describing the v1.1 update should be reviewed to ensure it does not repeat the "new Model Security domain" characterization found in some secondary coverage, since that framing risks misdirecting remediation budget toward controls that already have a year of implementation history.

Strategic Considerations

Longer term, organizations building AI governance programs around AICM should plan for the framework to continue tracking CCM's release cadence, meaning future CCM revisions are likely to produce corresponding AICM updates concentrated in the CCM-derived domains rather than in the AI-specific domains like Model Security, Data Security and Privacy Lifecycle Management, or Supply Chain Management. Treating AICM's AI-specific domains as the more stable core of an AI governance program, and the CCM-derived domains as subject to more frequent synchronization churn, offers a more accurate planning model than expecting uniform change across all 18 domains with each release. Organizations should also build a standing practice of validating framework-update claims against CSA's own release material before propagating those claims internally, given how quickly secondary commentary on a fast-moving framework like AICM can introduce and spread inaccuracies.

CSA Resource Alignment

This analysis draws directly on CSA's own AICM v1.1 artifact family, which should be adopters' primary reference for implementation work. The AI Controls Matrix v1.1 artifact itself [2] is the authoritative source for the 247-control, 18-domain structure discussed throughout this note, including the current Model Security domain definition and its associated mappings. Model Providers implementing or refreshing MDS-domain controls should work from the AICM v1.1 Implementation Guidelines for Model Providers [8], which provides control-by-control guidance specific to the actor role that carries the greatest share of Model Security responsibility, and the companion AICM v1.1 Auditing Guidelines for Model Providers [9], which gives assessors verification procedures for those same controls. Organizations consuming AI services rather than building them should consult the AICM v1.1 Implementation Guidelines for AI Customers [10] to understand their narrower slice of Model Security responsibility under the framework's shared responsibility model, particularly around vendor oversight and model configuration controls. Together, these four artifacts give both model builders and model consumers a complete, role-specific path through the v1.1 update without needing to treat any part of the Model Security domain as a ground-up implementation effort.

References

[1] Bregkou, Marina. "AI Controls Matrix v1.1: Strengthening the Foundation for Trustworthy AI." Cloud Security Alliance, July 14, 2026.

[2] Cloud Security Alliance. "AI Controls Matrix v1.1." CSA Research Artifacts, 2026.

[3] Cloud Security Alliance. "Introducing the CSA AI Controls Matrix: A Comprehensive Framework for Trustworthy AI." Cloud Security Alliance, July 10, 2025.

[4] Cloud Security Alliance. "The CSA Cloud Controls Matrix v4.1: Strengthening the Future of Cloud Security." Cloud Security Alliance, December 2, 2025.

[5] Cloud Security Alliance. "Cloud Controls Matrix and CAIQ v4.1." CSA Research Artifacts, 2026.

[6] GetAIGovernance. "AICM v1.1 Review: 247 Controls, New MDS Domain, SSRM & Mappings for Enterprise AI." July 2026.

[7] Virtualization Review. "Cloud Security AI Controls Matrix Turns AI Governance Into Assessment Framework." June 26, 2026.

[8] Cloud Security Alliance. "AICM v1.1 Implementation Guidelines for Model Providers (MP)." CSA Research Artifacts, 2026.

[9] Cloud Security Alliance. "AICM v1.1 Auditing Guidelines for Model Providers (MP)." CSA Research Artifacts, 2026.

[10] Cloud Security Alliance. "AICM v1.1 Implementation Guidelines for AI Customers (AIC)." CSA Research Artifacts, 2026.

