GAISB

GLOBAL AI STANDARDS BODY

EDITION 1.0 / 2026
EFFECTIVE 01 Q2 2026
REVIEW CYCLE: BIENNIAL

THE OFFICIAL PROFESSIONAL STANDARD

The *Common Body* of Knowledge

The definitive competency framework for the Certified Artificial Intelligence Strategist (CAIS) credential — governed by the Global AI Standards Body, contributed to by practitioners from over 100 countries, and aligned to the NIST AI RMF, the EU AI Act, ISO/IEC 42001, and the OECD AI Principles.

GOVERNED BY THE GAISB STANDARDS COUNCIL

TABLE OF CONTENTS

What This Document Contains

00 / PREFACE

A Letter from the Standards Council

Artificial intelligence is the most consequential general-purpose technology of our era. It is reshaping every industry, every profession, and every institution. And yet, at the moment the world most needs a clear definition of what it means to be competent in AI, no such definition exists. The credential market is fragmented, vendor-controlled, and disconnected from the actual work practitioners do.

The Global AI Standards Body was founded to fix this. Our remit is simple: to define, govern, and evolve the professional standard for AI practitioners — independent of any single vendor, platform, or geography.

This document, the **CAIS Common Body of Knowledge**, is the foundation of that standard. It defines what every Certified Artificial Intelligence Strategist must know, be able to do, and be accountable for. It is the blueprint for the CAIS examination, the basis for capstone assessment, and the reference framework employers and governments can use to evaluate AI professionals.

Three principles shaped this work:

- **Global contribution.** The CBK was developed by subject matter experts and practitioners from over 100 countries, across enterprise, academia, government, and independent practice — so that CAIS is not the standard of one region, but of the profession.
- **Practitioner reality.** Competency is defined by what a practitioner must demonstrably *do*, not merely what they can recite. Every domain carries performance-based build requirements, not only knowledge objectives.
- **Framework alignment.** The CBK is explicitly aligned with the NIST AI Risk Management Framework, the EU AI Act, ISO/IEC 42001, and the OECD AI Principles — so that a CAIS-certified professional is operating under the same standards recognized by governments and regulators worldwide.

This document will be reviewed and revised on a biennial cycle, with interim updates for material changes in the AI landscape. It is a public document. We invite scrutiny, critique, and contribution. That is how a standard earns the right to be called one.

> *Where vendor certifications prove you can use a tool, the CAIS credential proves you can think, build, and lead in AI — by the standards the global profession itself wrote.*

— The GAISB Standards Council

SECTION I

Purpose, Scope & Governance

1.1 Purpose of This Document

The Common Body of Knowledge (the "CBK") establishes the authoritative competency framework for the Certified Artificial Intelligence Strategist (CAIS) credential. It defines the knowledge domains, learning outcomes, performance requirements, ethical obligations, and continuing education standards that define a CAIS-certified professional.

The CBK serves five distinct audiences:

- **Candidates** — to understand exactly what is required to earn and maintain the CAIS credential.
- **Certified professionals** — to guide continuing practice and professional development.
- **Employers** — to evaluate what a CAIS credential reliably indicates about a candidate's capability.
- **Educators and training providers** — to align curricula to a recognized external standard.
- **Governments and institutions** — to reference CAIS in procurement, policy, and workforce development.

1.2 Scope

The CBK covers the full professional practice of AI strategy, design, implementation, and governance. It intentionally spans both technical and strategic competencies on the premise that an effective AI practitioner must understand both the capability of the systems and the context in which they are deployed. The CBK does not substitute for specialized advanced credentials in narrow subfields (for example, machine-learning research, hardware engineering, or AI safety research) but establishes the shared foundation from which those specializations diverge.

1.3 Governance

The CBK is developed, maintained, and revised by the **GAISB Standards Council**, comprising named subject matter experts drawn from enterprise, academia, government, and independent practice, representing the global practitioner community. Council membership, governance procedures, and contribution protocols are published separately in the GAISB Governance Charter.

1.4 Review Cycle

The CBK is formally reviewed on a **biennial cycle**. Interim revisions may be issued when the AI landscape undergoes material changes that affect the competency requirements of the profession. All revisions are versioned, dated, and logged in the public Revision Register.

1.5 Development Methodology

The CBK is developed through a multi-stage methodology:

1. **Job Task Analysis (JTA)** — Empirical study of the work performed by practicing AI professionals across geographies, industries, and role types.
2. **Domain Synthesis** — Translation of JTA findings into knowledge and skill domains with defined boundaries and weightings.
3. **Expert Review** — Review by the Standards Council and invited external reviewers, including framework-body liaisons.
4. **Public Comment** — Open public review period before ratification.
5. **Ratification** — Formal adoption by the Standards Council and publication.
6. **Psychometric Validation** — Post-publication, exam items are validated against the CBK through standard psychometric methods.

SECTION II

The CAIS Professional — Definition & Scope of Practice

2.1 Professional Definition

A **Certified Artificial Intelligence Strategist** is a practitioner who holds demonstrated competency in the design, deployment, governance, and strategic application of modern artificial intelligence systems. The CAIS credential attests that the holder has:

- Mastered the knowledge domains defined in this CBK, verified by examination.
- Demonstrated applied capability through real-world builds and a supervised capstone project.
- Agreed to be bound by the GAISB Code of Professional Conduct and is subject to its enforcement procedures.
- Committed to continuing competence through the CAIS Continuing Professional Education program.

2.2 Scope of Practice

The CAIS professional operates across four intersecting domains of practice:

STRATEGIC

Formulating AI strategy at the organizational, product, or venture level. Translating AI capability into business outcomes. Advising executive leadership.

TECHNICAL

Designing, prompt-engineering, and architecting systems that leverage generative AI, agents, retrieval, and multimodal capabilities to perform real work.

GOVERNANCE

Ensuring deployed AI systems are ethical, compliant, secure, explainable, and aligned with published frameworks and organizational values.

APPLIED

Shipping production AI systems — from use case discovery through deployment, measurement, and iteration — within the operational reality of the organization.

2.3 Professional Obligations

Every CAIS-certified professional accepts the following standing obligations:

1. To practice with integrity, transparency, and accountability.
2. To protect the safety, privacy, and dignity of those affected by the AI systems they design or deploy.
3. To maintain current competence through the CPE program.
4. To contribute to the advancement of the profession through practice, teaching, and peer exchange.
5. To comply with the GAISB Code of Professional Conduct, and to accept the authority of GAISB's enforcement procedures, including the suspension or revocation of the credential for material violations.

SECTION III

The CAIS Competency Framework

The CBK defines competency along three orthogonal axes. Every learning outcome and assessment item is mapped across all three.

3.1 The Three Axes of Competency

AXIS	DEFINITION	ASSESSMENT METHOD
Knowledge	What the practitioner must *know* — concepts, frameworks, technical foundations, regulatory context.	Written examination (blueprinted to CBK)
Skill	What the practitioner must be able to *do* — prompts they can write, systems they can architect, analyses they can produce.	Performance-based assessment; build tasks
Judgment	How the practitioner *decides* under real-world ambiguity — tradeoffs, ethics, strategy, context.	Scenario-based examination; capstone review; oral defense (advanced tiers)

3.2 Bloom's Taxonomy Mapping

Each learning outcome in this CBK is tagged with a Bloom's Taxonomy level to ensure cognitive balance across the examination. The CAIS exam blueprint requires that at least **60% of assessment items test at or above the "Apply" level**, with substantial presence at "Analyze," "Evaluate," and "Create." This distinguishes CAIS from credentials that rely predominantly on recall.

LEVEL	VERB	CAIS WEIGHTING FLOOR
Remember	Define, list, identify	≤ 15%
Understand	Explain, summarize, classify	≤ 25%
Apply	Implement, use, execute	≥ 25%
Analyze	Differentiate, compare, structure	≥ 15%
Evaluate	Judge, critique, justify	≥ 10%
Create	Design, construct, produce	≥ 10%

3.3 Performance-Based Requirements

Every domain carries a mandatory set of **Build Tasks**. These are not recommended exercises — they are required performance demonstrations, submitted, reviewed, and archived as part of the candidate's certification record. The Build Task portfolio is retained as evidence of applied competency and is a required artifact for credential verification.

3.4 Job Task Analysis (JTA)

The CAIS competency framework is anchored by a formal **Job Task Analysis**, conducted in accordance with recognized personnel certification standards (ISO/IEC 17024 §9.1, ANSI/ASTM E2659). The JTA defines the work of a qualified AI practitioner as a set of critical task domains, validated empirically against the practicing population.

The Edition 1.0 JTA was conducted through:

- **Subject Matter Expert (SME) Panel** — 30+ practicing AI professionals spanning enterprise, startup, consulting, government, and academia, drawn from 12+ countries.
- **Role Inventory** — Structured inventory of tasks performed by practitioners in AI strategy, AI engineering, AI product, AI governance, and AI consulting roles.
- **Validation Survey** — Quantitative validation across a broader population of practitioners rating each task on three dimensions: *frequency*, *criticality*, and *difficulty*.
- **Domain Derivation** — Task clusters mapped to the seven CBK domains with weightings derived from validation survey results.
- **Biennial Re-Validation** — The JTA is refreshed on the same biennial cycle as the CBK to remain current with practice evolution.

3.5 Critical Work Domains

The JTA produced eight Critical Work Domains (CWDs) that define the scope of the CAIS professional's work. Each CWD decomposes into task statements that inform domain learning outcomes and Build Task design.

CWD	CRITICAL WORK DOMAIN	REPRESENTATIVE TASK STATEMENTS
CWD-1	**AI Opportunity Identification**	Assess organizational context; identify AI-leverageable processes; estimate value; scope pilots.
CWD-2	**AI System Design**	Select model, architecture, and retrieval strategy; design prompts, tool use, and memory; specify guardrails.
CWD-3	**AI System Build & Deployment**	Implement systems; integrate with data and tools; ship to production; instrument observability.
CWD-4	**Evaluation & Quality Assurance**	Design evaluation harnesses; run red-team exercises; measure quality, safety, and drift; produce model cards.
CWD-5	**AI Risk & Governance**	Apply NIST AI RMF, EU AI Act, ISO 42001, and OECD Principles; conduct impact assessments; define policy.
CWD-6	**Organizational Transformation**	Lead change management; upskill teams; redesign workflows; measure adoption and ROI.
CWD-7	**Stakeholder Communication**	Advise executives and boards; translate technical options to business language; present tradeoffs.
CWD-8	**Professional Responsibility**	Apply the Code of Professional Conduct; disclose limits; escalate risk; maintain continuing competence.

The full JTA Report, including task-statement inventory, SME rosters, validation data, and methodology detail, is published separately in the GAISB Research Repository and cited as an authoritative companion to this CBK.

3.6 Knowledge–Skill–Ability (KSA) Mapping

Each Critical Work Domain is decomposed into Knowledge, Skill, and Ability statements (KSAs). Every KSA is traceable to one or more learning outcomes in Sections V–XI and to one or more Build Tasks. This traceability enables **blueprint defensibility**: every examination item and Build Task can be traced to an empirically validated professional task.

SECTION IV

Domain Structure & Examination Weighting

The CAIS Common Body of Knowledge is organized into seven domains. Each domain is weighted according to its centrality in the job task analysis and its importance in current practice. The domain weightings determine the proportion of examination items and the proportion of capstone review attention.

CODE	DOMAIN	WEIGHT	COURSE ALIGNMENT
D1	**AI Strategy & the Modern AI Economy** The strategic, economic, and professional landscape of the AI era.	**10%**	AIS-101
D2	**Foundations of Generative AI** Technical and conceptual foundations of modern AI systems.	**12%**	AIS-120
D3	**Prompt Engineering & LLM System Design** The craft and engineering discipline of reliable LLM systems.	**18%**	AIS-130
D4	**AI Agents & Agentic Workflows** Design, deployment, and governance of autonomous AI systems.	**18%**	AIS-140
D5	**Ethics, Data & Responsible AI** Ethical, legal, regulatory, and operational responsibility.	**12%**	AIS-160
D6	**AI Strategy, Transformation & Business Innovation** Translating AI capability into enterprise-scale business outcomes.	**15%**	AIS-210
D7	**Innovation & Applied AI Foundations** Shipping real AI systems from prototype to production.	**15%**	AIS-230
	TOTAL	**100%**	

Each domain is defined in detail in Sections V through XI, with full learning outcomes, knowledge requirements, skill requirements, mandatory Build Tasks, and external framework alignment.

SECTION V / DOMAIN 1

DOMAIN D1

AI Strategy & the Modern AI Economy

10%
EXAM WEIGHT

DOMAIN DESCRIPTION

Domain 1 establishes the strategic worldview required of every CAIS practitioner. It addresses the economic, technological, and professional landscape shifts caused by the AI transition, the mindset required to operate effectively within it, and the strategic logic by which individuals, organizations, and nations will win or lose. This domain is intentionally positioned first because technical and applied competence without strategic framing produces tactical operators who cannot see where leverage actually lives.

LEARNING OUTCOMES

Upon mastery of this domain, the CAIS practitioner will be able to:

- **D1.1** — Articulate the economic, technological, and social forces driving the current AI transition, and distinguish them from prior waves.
- **D1.2** — Analyze how AI is restructuring value chains, labor markets, and industry boundaries across at least five major sectors.
- **D1.3** — Identify and evaluate personal and organizational leverage points created by AI capability.
- **D1.4** — Evaluate emerging AI-era operating models (AI-native, AI-augmented, AI-adjacent) and apply the correct model to a given context.
- **D1.5** — Frame AI as a general-purpose technology and apply GPT diffusion theory to strategic forecasting.
- **D1.6** — Develop a strategic mindset oriented around leverage, asymmetric upside, and ecosystem thinking.

KNOWLEDGE REQUIREMENTS

- History of AI: three prior waves and why the current wave differs
- The AI value chain: compute, models, orchestration, applications, distribution
- General-purpose technology (GPT) theory and diffusion patterns
- Incumbents vs. insurgents: disruption economics in the AI era
- Capability frontiers: benchmarks and the shifting meaning of "state of the art"
- Labor-market implications: displacement, augmentation, and new role formation
- Geopolitical and regulatory dynamics shaping AI deployment
- Capital flows: AI investment patterns and their strategic signals

SKILL REQUIREMENTS

- Strategic scenario analysis across 1, 3, and 5-year horizons
- Opportunity mapping against an industry or function
- Personal positioning within the AI economy
- Framing AI initiatives with leverage-based logic (not feature-based)
- Communicating AI strategy to executive, investor, and practitioner audiences
- Distinguishing durable advantages from expiring ones

MANDATORY BUILD TASKS

1. **Industry Disruption Analysis** — A 1,500–3,000 word strategic analysis of AI's restructuring effect on a nominated industry, including first-order, second-order, and third-order impacts, the identification of defensible and undefensible positions, and a 3-year forecast.
2. **Personal AI Leverage Map** — A structured artifact mapping the candidate's current professional capability, the AI augmentation patterns available to them, and the specific leverage points they will pursue over the following 12 months.

3. **Strategic Briefing Document** — An executive-format briefing document on a nominated AI opportunity, formatted for delivery to a senior decision-maker.

FRAMEWORK ALIGNMENT

This domain aligns to the **"Govern"** function of the NIST AI RMF (establishing the strategic and organizational context for AI risk management) and to the OECD AI Principle of **"Inclusive Growth, Sustainable Development and Well-being."**

SECTION VI / DOMAIN 2

DOMAIN D2

Foundations of Generative AI

12%
EXAM WEIGHT

DOMAIN DESCRIPTION

Domain 2 establishes the technical and conceptual foundation every AI practitioner must hold. It does not require the practitioner to train a frontier model, but it does require that they understand how such models work, what they can and cannot do, how they are evaluated, and how to navigate a rapidly shifting landscape of model families, modalities, and capabilities. A CAIS professional without this foundation cannot make sound architectural decisions and cannot credibly advise organizations on AI capability.

LEARNING OUTCOMES

Upon mastery of this domain, the CAIS practitioner will be able to:

- **D2.1** — Explain transformer architecture at a conceptual level sufficient to reason about its capabilities and limitations.
- **D2.2** — Differentiate among major model families (e.g., open-weight, closed API, fine-tuned, specialized) and their appropriate use cases.
- **D2.3** — Evaluate model selection for a specific use case against performance, cost, latency, privacy, and governance constraints.
- **D2.4** — Distinguish training, fine-tuning, and inference concerns and articulate when each is relevant.
- **D2.5** — Navigate multimodal capability (text, image, audio, video, code) and identify correct modality for a given use case.
- **D2.6** — Interpret model evaluation metrics and benchmarks with appropriate skepticism.

- **D2.7** — Establish a personal practice for staying current with the model landscape.

KNOWLEDGE REQUIREMENTS

- Neural network fundamentals and the shift to attention-based architectures
- Transformer architecture: attention, embeddings, tokens, context windows
- Pre-training, supervised fine-tuning, RLHF, constitutional methods
- Model alignment and its limitations
- Open-weight vs. closed-API models: tradeoffs across cost, control, and capability
- Embeddings and vector representations
- Multimodal models: text, image, audio, video, code
- Common benchmarks (MMLU, GPQA, HumanEval, SWE-bench) and their limitations
- Hallucination, calibration, and reliability characteristics
- Model economics: token pricing, latency, infrastructure

SKILL REQUIREMENTS

- Conduct a structured model selection analysis for a real use case
- Build a capability probe suite to evaluate a new model
- Produce a cost and latency analysis for production deployment
- Interpret a benchmark report and identify its blind spots
- Construct an ongoing model-tracking discipline (sources, cadence, signal)
- Translate capability updates into strategic implications

MANDATORY BUILD TASKS

1. **Model Comparison Matrix** — A rigorous side-by-side comparison of at least four candidate models (spanning open and closed) for a specified real-world use case, with weighted scoring and a defended recommendation.
2. **Capability Probe Suite** — A reusable set of structured prompts and evaluation criteria used to assess a new model's capability on the practitioner's target task type.
3. **Production Cost Analysis** — A token-level economic model projecting cost, latency, and infrastructure requirements for a specified production deployment at three volume tiers.

FRAMEWORK ALIGNMENT

This domain aligns to the **"Map"** and **"Measure"** functions of the NIST AI RMF. It informs the technical documentation requirements of the EU AI Act (Article 11) and the measurement controls of ISO/IEC 42001.

SECTION VII / DOMAIN 3

DOMAIN D3

Prompt Engineering & LLM System Design

18%
EXAM WEIGHT

DOMAIN DESCRIPTION

Domain 3 is the largest-weighted domain in the CBK because prompt engineering and LLM system design is the core engineering craft of the modern AI practitioner. It covers the disciplined practice of extracting reliable, high-quality behavior from large language models through prompt design, context engineering, retrieval augmentation, evaluation, and production system architecture. This is the domain where abstract model capability becomes real, deployable work.

LEARNING OUTCOMES

Upon mastery of this domain, the CAIS practitioner will be able to:

- **D3.1** — Design and iterate on effective prompts across task types (generation, reasoning, classification, structured extraction).
- **D3.2** — Apply and combine prompt patterns (zero-shot, few-shot, chain-of-thought, self-consistency, ReAct, tree-of-thought) appropriately.
- **D3.3** — Architect LLM-based systems with reliability, observability, and cost controls.
- **D3.4** — Design and implement retrieval-augmented generation (RAG) systems with appropriate chunking, embedding, and retrieval strategies.
- **D3.5** — Build evaluation frameworks ("evals") that quantitatively measure system performance and detect regression.
- **D3.6** — Engineer context and state management for multi-turn, multi-document, and long-horizon tasks.

- **D3.7** — Produce structured, reliable outputs (JSON, function calls, schemas) suitable for downstream system integration.
- **D3.8** — Optimize production systems for cost, latency, and quality tradeoffs.

KNOWLEDGE REQUIREMENTS

- Prompt anatomy: instruction, context, examples, constraints, output format
- Prompt patterns: zero-shot, few-shot, CoT, self-consistency, ReAct, ToT, reflection
- System prompts vs. user prompts; role prompting; meta-prompting
- Context engineering: window management, priming, compression
- RAG architecture: ingestion, chunking, embedding, retrieval, re-ranking
- Vector databases: selection criteria and operational characteristics
- Hybrid retrieval: semantic + keyword, BM25 integration
- Evals: golden datasets, LLM-as-judge, rubric-based scoring
- Structured output: JSON mode, function calling, schema enforcement
- Prompt injection and defense strategies

SKILL REQUIREMENTS

- Iterate a prompt from naive to production-grade with measured improvement
- Build and maintain a versioned prompt library
- Design an eval suite with representative test cases and scoring rubrics
- Construct a RAG pipeline over a real document corpus
- Manage context window trade-offs in long-horizon tasks
- Debug unreliable LLM behavior and diagnose root cause
- Translate unstructured requirements into a reliable system prompt
- Produce an architecture diagram for an LLM-powered system

- Caching, batching, and cost optimization patterns

MANDATORY BUILD TASKS

1. **Production-Grade Prompt Library** — A versioned, documented library of at least ten production-ready prompts covering the practitioner's target vertical, with examples, failure modes, and test cases for each.
2. **Deployed RAG System** — A functioning retrieval-augmented generation system over a real document corpus (minimum 200 documents), with documented chunking strategy, retrieval evaluation, and end-to-end quality measurement.
3. **Eval Suite** — An automated evaluation harness with at least 25 test cases, producing repeatable quantitative scores and capable of detecting regression across prompt or model changes.
4. **System Architecture Document** — An architecture document for an LLM-powered system including prompt flow, retrieval pipeline, caching strategy, observability, failure handling, and cost model.

FRAMEWORK ALIGNMENT

This domain aligns to the **"Manage"** function of the NIST AI RMF (implementing controls for identified risks) and to the operational controls of ISO/IEC 42001. Eval-based quality measurement satisfies the EU AI Act requirements for high-risk system testing (Article 15).

SECTION VIII / DOMAIN 4

DOMAIN D4

AI Agents & Agentic Workflows

18%
EXAM WEIGHT

DOMAIN DESCRIPTION

Domain 4 covers the design, deployment, and governance of agentic AI systems — systems that plan, act, and iterate toward goals rather than respond to single-turn prompts. Agentic systems are the current frontier of practical AI deployment, and they introduce architectural, operational, and safety considerations that do not arise in single-turn LLM use. A CAIS practitioner must be able to design, deploy, and govern agents responsibly at enterprise scale.

LEARNING OUTCOMES

Upon mastery of this domain, the CAIS practitioner will be able to:

- **D4.1** — Architect single-agent systems using appropriate architectural patterns.
- **D4.2** — Architect multi-agent systems with well-defined roles, communication protocols, and orchestration.
- **D4.3** — Integrate tools, APIs, and external systems into agent capability.
- **D4.4** — Design agentic workflows that automate real business processes end-to-end.
- **D4.5** — Govern agent behavior through guardrails, policy enforcement, and human-in-the-loop controls.
- **D4.6** — Measure agent reliability and diagnose failure modes.
- **D4.7** — Manage memory and state across long-horizon agent sessions.
- **D4.8** — Balance autonomy with safety and auditability.

KNOWLEDGE REQUIREMENTS

- Agent architectures: ReAct, planner-executor, hierarchical, reflexion
- Tool use and function calling: design, registration, safety
- Multi-agent orchestration patterns: supervisor, swarm, pipeline, peer
- Memory systems: short-term, long-term, episodic, semantic
- Planning and task decomposition strategies
- Reflection, self-correction, and iterative improvement loops
- Guardrails: input, output, behavioral, and topical
- Agent observability: tracing, logging, replay
- Human-in-the-loop and human-on-the-loop designs
- Agent-specific security threats (prompt injection, tool abuse, goal drift)
- Cost and latency management in multi-step agents

SKILL REQUIREMENTS

- Decompose a business process into agent-executable tasks
- Select and apply the correct agent architecture for a given task
- Design a tool suite appropriate to an agent's scope
- Implement guardrails that hold under adversarial conditions
- Build agent observability for production debugging
- Diagnose agent failure (infinite loops, drift, tool misuse)
- Evaluate agent reliability with structured test harnesses
- Design appropriate human oversight points

MANDATORY BUILD TASKS

1. **Production Single Agent** — A deployed agent with at least three tool integrations, guardrails, and observability, executing a real task end-to-end. Includes architecture document, test log, and failure-mode analysis.
2. **Multi-Agent Workflow** — A multi-agent system automating a real business process (e.g., research, onboarding, content production, support triage) with at least three specialized agents and a supervisor orchestrator.
3. **Agent Guardrail System** — A documented guardrail policy covering input, output, tool-use, and behavioral constraints, with test cases demonstrating enforcement under adversarial prompts.
4. **Agent Observability Dashboard** — A production-grade monitoring surface showing traces, tool calls, token usage, cost per session, and failure-mode classification.

FRAMEWORK ALIGNMENT

This domain aligns to the **"Manage"** function of the NIST AI RMF and directly addresses the high-risk system requirements of the EU AI Act where agents take autonomous action (Articles 13, 14, 15). Guardrail design satisfies OECD AI Principles on **"Robustness, Security and Safety."**

SECTION IX / DOMAIN 5

DOMAIN D5

Ethics, Data & Responsible AI

12%
EXAM WEIGHT

DOMAIN DESCRIPTION

Domain 5 establishes the ethical, legal, regulatory, and operational frameworks required to deploy AI responsibly across enterprise and society. This domain is where CAIS credential holders differentiate themselves from unregulated practitioners: they operate under a published code, they understand the regulatory frameworks that now govern AI, and they can design systems that meet the compliance, safety, and accountability standards that will increasingly define professional AI practice.

LEARNING OUTCOMES

Upon mastery of this domain, the CAIS practitioner will be able to:

- **D5.1** — Apply responsible AI frameworks (NIST AI RMF, EU AI Act, ISO/IEC 42001, OECD) to real systems.
- **D5.2** — Identify, measure, and mitigate bias across the AI lifecycle.
- **D5.3** — Design AI systems that meet data protection and privacy requirements (GDPR, CCPA, sector-specific).
- **D5.4** — Secure AI systems against adversarial threats including prompt injection, data poisoning, model theft, and extraction.
- **D5.5** — Produce explainability and transparency documentation appropriate to the system's risk classification.
- **D5.6** — Conduct AI-specific risk assessments across the full system lifecycle.
- **D5.7** — Operate under a formal Code of Professional Conduct and resolve ethical ambiguity through structured judgment.

- **D5.8** — Establish organizational AI governance structures, policies, and review processes.

KNOWLEDGE REQUIREMENTS

- NIST AI RMF 1.0 and its four functions (Govern, Map, Measure, Manage)
- EU AI Act: risk categorization, prohibited practices, conformity assessment, penalties
- ISO/IEC 42001: AI management system requirements
- OECD AI Principles and implementation
- Bias: taxonomies (representational, allocational, measurement), sources, and mitigation
- Data protection frameworks in AI context (GDPR Articles 22, 35; CCPA; sector-specific)
- AI security threats: prompt injection, jailbreaks, data poisoning, membership inference, model extraction
- Red-teaming methodology
- Explainability: SHAP, LIME, intrinsic vs. post-hoc; limits of interpretability
- Model cards, data cards, system cards
- AI incident response

SKILL REQUIREMENTS

- Conduct an AI risk assessment for a real system
- Map a deployed system to the NIST AI RMF
- Produce a model card and system card
- Design and execute a bias audit
- Red-team an LLM or agent for security vulnerabilities
- Draft an AI use policy for an organization
- Navigate ethical ambiguity under a professional code
- Produce EU AI Act conformity documentation for high-risk systems

- Organizational AI governance: roles, committees, review boards

MANDATORY BUILD TASKS

1. **AI Risk Assessment** — A full risk assessment of a real deployed (or simulated-deployed) AI system, including hazard identification, impact analysis, control mapping, and residual risk statement. NIST AI RMF-aligned.
2. **AI Governance Policy** — A draft AI governance policy document suitable for an organization of 50–500 people, including scope, roles, review gates, and incident response.
3. **Red-Team Report** — A structured red-team exercise against an LLM or agent system with at least 15 adversarial test cases, documented outcomes, and remediation recommendations.
4. **Framework Mapping Document** — Complete mapping of a specified system to NIST AI RMF and either EU AI Act or ISO/IEC 42001, with evidence and gap analysis.

FRAMEWORK ALIGNMENT

Domain 5 is directly aligned to all four functions of the NIST AI RMF (**Govern, Map, Measure, Manage**), the full scope of the EU AI Act, ISO/IEC 42001, and the OECD AI Principles. This domain is the core compliance and governance competency of the CAIS credential.

SECTION X / DOMAIN 6

DOMAIN D6

AI Strategy, Transformation & Business Innovation

15%
EXAM WEIGHT

DOMAIN DESCRIPTION

Domain 6 addresses the translation of AI capability into organizational outcomes. It covers AI strategy formulation, transformation program design, ROI modeling, operating model redesign, change leadership, portfolio management, and executive advisory. A CAIS practitioner who cannot move between the technical system and the executive conversation is a technician, not a strategist. This domain establishes the competency required to lead AI initiatives at enterprise scale.

LEARNING OUTCOMES

Upon mastery of this domain, the CAIS practitioner will be able to:

- **D6.1** — Formulate AI strategy at the organizational, divisional, or product level.
- **D6.2** — Build AI transformation roadmaps across 12, 24, and 36-month horizons.
- **D6.3** — Model AI return on investment with appropriate rigor, including direct, second-order, and option-value components.
- **D6.4** — Redesign operating models to integrate AI capability at the team, function, and enterprise levels.
- **D6.5** — Lead AI-driven change programs and manage organizational resistance.
- **D6.6** — Make defensible build-vs.-buy-vs.-partner decisions.
- **D6.7** — Manage an AI portfolio across risk, horizon, and business value.
- **D6.8** — Advise executive leadership and boards on AI strategy and governance.

KNOWLEDGE REQUIREMENTS

- AI strategy frameworks: horizons, plays, and archetypes
- Transformation methodology: phased change, dual-track execution
- ROI modeling for AI: assumption surfacing, sensitivity analysis
- Operating model design: centralized, federated, CoE, embedded
- Change management theory (Kotter, ADKAR) in AI context
- AI portfolio management: explore / exploit / expand
- Build vs. buy vs. partner: decision criteria
- AI talent strategy: hire, upskill, partner, automate
- AI economics at enterprise scale (compute, licensing, team)
- Executive communication: board-level AI briefings
- Vendor evaluation and procurement for AI

SKILL REQUIREMENTS

- Identify and prioritize AI opportunities at the enterprise level
- Develop a rigorous business case for an AI initiative
- Produce a defensible transformation roadmap
- Communicate AI strategy to executives with clarity and authority
- Design an AI operating model fit for the organization's maturity
- Manage stakeholders across technical and business functions
- Run a portfolio review under capital and attention constraints
- Advise on AI governance at the board level

MANDATORY BUILD TASKS

1. **AI Transformation Roadmap** — A full transformation roadmap for a real or case-study organization, covering 24 months, phased initiatives, capacity modeling, governance, and KPI framework.
2. **AI Business Case** — A rigorous business case for a single AI initiative, including baseline, intervention, assumptions, sensitivity analysis, risks, and recommendation, in a format suitable for executive approval.
3. **Operating Model Redesign Proposal** — A proposal redesigning an organizational operating model around AI capability, including roles, decision rights, governance, and transition plan.
4. **Executive Briefing Deck** — A board-ready briefing deck (10–15 slides) on an AI strategy recommendation, including appendices and speaker notes.

FRAMEWORK ALIGNMENT

This domain aligns to the **"Govern"** function of the NIST AI RMF, to the **organizational context** and **leadership** clauses of ISO/IEC 42001 (Clauses 4 and 5), and to the OECD AI Principle of **"Inclusive Growth, Sustainable Development and Well-being."**

SECTION XI / DOMAIN 7

DOMAIN D7

Innovation & Applied AI Foundations

15%
EXAM WEIGHT

DOMAIN DESCRIPTION

Domain 7 is the applied integration domain. It covers the discipline of shipping real AI systems — from use case discovery and prototyping through deployment, measurement, and iteration. This is where all prior knowledge is synthesized into working systems that produce value. A CAIS practitioner must be able to close the loop between idea and deployed capability. This domain provides the scaffolding that governs the candidate's capstone project.

LEARNING OUTCOMES

Upon mastery of this domain, the CAIS practitioner will be able to:

- **D7.1** — Identify, prioritize, and scope AI use cases with demonstrated value potential.
- **D7.2** — Design and ship an AI product or system from prototype to production.
- **D7.3** — Measure AI system performance across technical, user, and business dimensions.
- **D7.4** — Iterate deployed AI systems with appropriate experimentation methodology.
- **D7.5** — Scale an AI system from pilot to enterprise deployment.
- **D7.6** — Lead a cross-functional AI build team.
- **D7.7** — Establish the operational discipline (LLMOps) required for reliable AI in production.
- **D7.8** — Document and publish case studies of deployed AI work to professional standard.

KNOWLEDGE REQUIREMENTS

- AI product lifecycle: discovery, scoping, build, deploy, measure, iterate
- Use case evaluation criteria: value, feasibility, data, risk, build-vs-integrate
- Prototype-to-production transition patterns
- MLOps vs. LLMOps: monitoring, versioning, evaluation, guardrails in production
- Performance measurement: technical (quality, latency, cost) + product (engagement, retention) + business (ROI)
- A/B and shadow testing for AI features
- User experience patterns for AI products
- Integration patterns: API, embedded, co-pilot, autonomous
- Rollout strategies: dark launch, canary, ring, geo
- Incident response in AI systems
- Documentation standards for deployed AI

SKILL REQUIREMENTS

- Run a structured use case discovery and prioritization
- Build a working AI prototype within days, not months
- Ship an AI system to production with observability
- Design a measurement dashboard that reflects actual value
- Run experiments that produce defensible conclusions
- Lead a small cross-functional team through an AI build
- Produce a written case study of professional quality
- Handle an in-production AI incident to resolution

MANDATORY BUILD TASKS

1. **Production AI Application** — A shipped AI application, accessible and operable by end users, with at least 30 days of production operation and documented usage.
2. **Measurement Dashboard** — A live dashboard measuring the production application across quality, usage, cost, and business-outcome metrics.
3. **Structured A/B Test** — A designed and executed A/B test (or shadow test) on an AI feature, with statistical treatment, defensible conclusion, and decision outcome.
4. **Published Case Study** — A written case study of the deployed system, in professional format, covering problem, system design, rollout, measurement, outcomes, and lessons.

FRAMEWORK ALIGNMENT

This domain aligns to the **"Manage"** and **"Measure"** functions of the NIST AI RMF, to the **operation** and **performance evaluation** clauses of ISO/IEC 42001 (Clauses 8 and 9), and to the post-market monitoring requirements of the EU AI Act (Article 72).

SECTION XII

Assessment Blueprint

The CAIS credential is earned through a combined assessment of knowledge, applied skill, and judgment. No single assessment instrument is sufficient on its own. The blueprint is designed to ensure that a CAIS-certified professional has demonstrated competency across all three axes of the Competency Framework.

12.1 Assessment Instruments

INSTRUMENT	AXIS	WEIGHT OF FINAL OUTCOME	DESCRIPTION
Written Exam	Knowledge, Judgment	40%	Proctored examination blueprinted to CBK domain weightings. Mix of multiple-response, scenario-based, and short-constructed-response items.
Build Portfolio	Skill	35%	All mandatory Build Tasks from each domain, reviewed against published rubrics and archived in the candidate's certification record.
Capstone	Skill, Judgment	25%	Full Capstone Project (see Section XIII) with documentation, demonstration, and review.

12.2 Examination Blueprint

The written examination is composed of approximately 150 items administered over a three-hour session. Item distribution is governed by the domain weighting in Section IV and the Bloom's distribution floor in Section III. Item types include:

- **Multiple-response** — Concept and application questions, four- or five-option format, one or more correct answers.
- **Scenario-based** — Extended vignettes requiring analytical reasoning across domains, typically 3–5 dependent items per vignette.
- **Short constructed response** — Written responses (50–200 words) evaluated against a published analytic rubric.
- **Case interpretation** — Short case studies with dependent item sets, testing synthesis and judgment.

12.3 Item Specifications by Domain

Each domain is assigned a target item count range and a cognitive distribution, in accordance with the domain weightings of Section IV and the Bloom's floors of Section 3.2. Item writers work from these specifications under direction of the Standards Council Examinations Committee.

DOMAIN	WEIGHT	ITEMS (TARGET)	COGNITIVE DISTRIBUTION (BLOOM'S)
D1 — AI Strategy & the Modern AI Economy	10%	14–16	Understand, Apply, Evaluate
D2 — Foundations of Generative AI	12%	17–19	Understand, Apply, Analyze
D3 — Prompt Engineering & LLM System Design	18%	26–28	Apply, Analyze, Create
D4 — AI Agents & Agentic Workflows	18%	26–28	Apply, Analyze, Create
D5 — Ethics, Data & Responsible AI	12%	17–19	Understand, Evaluate, Apply
D6 — AI Strategy, Transformation & Business Innovation	15%	22–24	Apply, Analyze, Evaluate
D7 — Innovation & Applied AI Foundations	15%	22–24	Apply, Analyze, Create
Total	**100%**	**~150**	60%+ at Apply or above

Items are drawn from a secured item bank. No item appears in consecutive examination forms. Each form is balanced to the blueprint within published tolerance bands.

12.4 Passing Standard Methodology

The passing standard is established through a **Modified Angoff method**, executed as follows:

1. **SME Panel.** A panel of no fewer than eight Subject Matter Experts is convened — at least three Standards Council members and at least three external SMEs drawn from practice, academia, and governance communities.
2. **Minimally Qualified Candidate (MQC) Definition.** The panel defines, in writing, the profile of a candidate who meets but does not exceed the minimum competency required of a certified practitioner.
3. **Item Rating — Round 1.** Each SME independently rates, for each item, the probability that an MQC would answer correctly. Ratings are aggregated and reviewed.

4. **Item Rating — Round 2.** After discussion of outliers, SMEs may revise ratings. Panel variance is tracked.
5. **Cut Score Derivation.** The raw cut score is computed as the average of panel-mean item probabilities, scaled to the form.
6. **Compromise Review.** The cut score is reviewed by the Standards Council against empirical performance data and may be adjusted within published tolerance for equating.
7. **Publication.** The passing standard, methodology summary, and form-equating results are published on a rolling basis.

The passing standard is re-established at each substantive CBK revision or every 48 months, whichever comes first.

12.5 Psychometric Standards

All CAIS examinations are subject to ongoing psychometric analysis. Target parameters are:

PARAMETER	TARGET STANDARD	SOURCE
Form reliability (internal consistency)	Cronbach's α ≥ 0.85 or KR-20 ≥ 0.85	Post-administration analysis
Item difficulty (P-value)	0.30 ≤ P ≤ 0.90 for operational items	Item analysis
Item discrimination (point-biserial)	r_{pbis} ≥ 0.20 for operational items	Item analysis
Differential Item Functioning (DIF)	Mantel-Haenszel flag review on protected-group reference panels	Bias & fairness review
Standard Error of Measurement (SEM)	Reported with every cut-score decision	Form equating
Test-retest reliability (constructed-response rubrics)	ICC ≥ 0.80 between trained raters	Rater calibration

Items failing to meet operational standards are retired, revised, or re-field-tested. A psychometric summary for each administration cycle is published in the GAISB Technical Report Series.

12.6 Examination Security & Proctoring

Examination integrity is governed by the GAISB Examination Security Policy, which includes the following controls:

- **Identity verification** — Government-issued photo ID and biometric verification at the point of examination.

- **Proctoring** — All examinations are proctored live on-site or via verified remote proctoring with continuous audiovisual monitoring, browser lockdown, and environment scan.
- **Item bank security** — Items are encrypted at rest and in transit; item writers and SMEs sign non-disclosure agreements; drafts are maintained in access-controlled environments.
- **Form rotation** — Multiple forms are in circulation concurrently; each form is equated to the operational standard.
- **Candidate agreements** — Each candidate agrees in writing to the GAISB Examination Agreement at registration, including non-disclosure of item content.
- **Detection and response** — Statistical cheating detection (e.g., answer-pattern analysis, response-time anomaly) is applied to every administration.
- **Sanctions** — Confirmed violations result in cancellation, a minimum five-year prohibition on re-examination, and public notation on the Verification Registry. Severe violations may result in permanent disqualification.

12.7 Candidate Accommodations

GAISB provides reasonable accommodations for candidates with documented disabilities, in accordance with applicable law and recognized testing-accommodation standards. Common accommodations include:

- Extended testing time (typically 1.5× or 2×)
- Separate testing rooms
- Assistive technology (screen readers, screen magnification)
- Rest breaks beyond the standard schedule
- Translation exceptions for non-native-language candidates in accordance with the CAIS Language Policy

Accommodation requests are submitted with supporting documentation to the GAISB Candidate Services Office no fewer than 30 days before the examination date. Accommodations do not lower the passing standard.

12.8 Language Policy

The Edition 1.0 CAIS examination is administered in English. Translated editions are issued on a rolling schedule following full translation-equivalence review. Each translated edition is equated to the English reference form; passing standards are established independently for each language edition.

12.9 Re-Examination

Candidates who do not pass may re-sit after a 90-day waiting period. Up to three attempts per 24-month period are permitted without additional review. Candidates who do not pass after three attempts are required to enroll in a documented remediation program before further attempts.

12.10 Score Reporting & Appeals

Candidates receive a pass/fail decision and a sub-domain performance summary within 30 days of administration. Candidates may appeal a failing decision on procedural grounds (e.g., proctoring incident, scoring error) within 30 days of score release, through the procedure described in Section 15.4.

SECTION XIII

The CAIS Capstone

The Capstone is the integrating assessment of the CAIS credential. It is the artifact that proves, without ambiguity, that the holder can translate the full CBK into a deployed, operational AI system or venture. The Capstone is what distinguishes CAIS from credentials that produce test-passers.

13.1 CAPSTONE REQUIREMENTS

Every candidate must complete a Capstone Project that satisfies all of the following:

- Integrates competency from at least five of the seven CBK domains
- Produces a deployed, operational system (not a prototype or slide deck)
- Solves a real-world problem for a real organization, client, or venture
- Operates in production for a minimum of 30 days with documented use
- Is accompanied by a full written documentation package
- Is presented in a recorded demonstration to a Capstone Review Panel
- Is published (with permission) as a case study in the GAISB repository

13.2 ACCEPTABLE CAPSTONE TYPES

Capstones may take any of the following forms, provided they meet Section 13.1 requirements:

- **Deployed AI System** — A production AI application, agent, or automation.
- **AI-Enabled Business or Venture** — A revenue-generating venture built on AI capability.

- **Enterprise AI Transformation Program** — A launched transformation initiative with measurable outcomes.

- **AI Governance Deployment** — A deployed AI governance program with documented adoption.

13.3 CAPSTONE REVIEW PANEL

Each Capstone is reviewed by a panel of at least three GAISB-authorized reviewers, including at least one Standards Council member or designee. The review evaluates the Capstone against a published rubric covering technical quality, strategic soundness, responsible-AI practice, applied rigor, and professional documentation. Reviewers may require revisions before acceptance.

13.4 CAPSTONE ARCHIVE

Accepted Capstones (with the certificant's consent) are archived in the GAISB Capstone Repository as permanent evidence of the credential and as a body of professional work that represents the state of the practice.

SECTION XIV

Certification Tiers & Progression Pathway

The CAIS credential is structured as a progression pathway across four tiers. Each tier has distinct competency requirements, assessments, and professional implications. A candidate may enter at Tier 1 and progress upward, or — with evidence — challenge directly at higher tiers.

TIER	DESIGNATION	COMPETENCY STANDARD	ASSESSMENT
I	**CAIS Practitioner** Can apply AI capability to real tasks with sound judgment.	Mastery of CBK Knowledge requirements; completion of core Build Tasks across all seven domains.	Written Exam + Build Portfolio
II	**CAIS Builder** Can architect and ship AI systems end-to-end.	Tier I + advanced Build Tasks in D3 and D4; deployed Capstone project.	Tier I + Capstone
III	**CAIS Operator** Can deploy and monetize AI systems with demonstrated outcomes.	Tier II + operational/business outcomes from deployed work; evidence of monetization or organizational adoption.	Tier II + Outcomes Review
IV	**CAIS Architect** Can design full AI ecosystems and advise at executive level.	Tier III + demonstrated enterprise transformation leadership; contribution to the profession; oral defense.	Tier III + Defense + Peer Panel

14.1 Designation Use

Certificants may use their tier designation (e.g., "CAIS — Builder") in professional communications, on resumes, and in public biography, provided their credential is active and in good standing. Misuse of the designation — including use during suspension, use of a tier not earned, or use after revocation — is a violation of the Code of Professional Conduct (Section XV) and is prosecuted under GAISB enforcement procedures.

SECTION XV

Code of Professional Conduct

Every CAIS-certified professional is bound by the GAISB Code of Professional Conduct. Acceptance of the Code is a condition of certification. Material violations may result in suspension or revocation of the credential, with public notation in the Verification Registry.

The Code is structured as seven Canons. Each Canon carries subordinate standards of practice (not reproduced in full here; published separately).

CANON I

Practice with Integrity and Honesty.

CAIS professionals represent their capabilities, systems, and conclusions truthfully. They do not misrepresent AI system performance. They disclose limitations. They do not overstate credentials, training data, or outcomes.

CANON II

Protect the Public Interest.

CAIS professionals design, deploy, and govern AI systems with due regard for the safety, dignity, privacy, and rights of those affected — including parties who are not their direct clients.

CANON III

Practice Within Competence.

CAIS professionals undertake only work for which they possess the requisite competence. They decline or escalate work that exceeds their capability, and they seek peer consultation where appropriate.

CANON IV

Maintain Confidentiality and Data Protection.

CAIS professionals protect client and user data in accordance with applicable law and professional standards. They do not use confidential data for unauthorized purposes, including training or fine-tuning models outside of scope.

CANON V

Avoid Harm and Unjust Discrimination.

CAIS professionals design systems that do not unjustly discriminate. Where a system has disparate impact, they disclose it, mitigate it, or decline to deploy it.

CANON VI

Advance the Profession.

CAIS professionals contribute to the development of the profession through teaching, mentoring, peer exchange, and adherence to published frameworks. They do not undermine peers through deceptive or anti-competitive practices.

CANON VII

Maintain Current Competence.

CAIS professionals remain current with the evolving practice of AI through the GAISB Continuing Professional Education program (Section XVI). Failure to meet CPE requirements is a violation of the Code.

15.1 Enforcement

Complaints of Code violations may be filed by any party through the GAISB Professional Responsibility Office. The Office conducts an initial review; credible complaints proceed to a Disciplinary Panel, which may issue sanctions including private reprimand, public reprimand, mandatory education, suspension of the credential for a defined period, or full revocation with public notation. All decisions are subject to a single right of appeal to the Standards Council.

15.2 Due Process Guarantees

Every certificant and candidate subject to disciplinary action is guaranteed the following due-process protections:

- **Written notice** of the allegation, specifying the Canon(s) implicated and the evidentiary basis.
- **A reasonable opportunity to respond** in writing, with a minimum 21-day response window.
- **Impartial adjudication** by a Disciplinary Panel from which any member with a material conflict of interest is recused.
- **Right to counsel or professional representation** at the certificant's own expense.
- **A written decision** stating the findings of fact, the Canon(s) violated, and the sanction imposed.
- **Confidentiality** of proceedings unless and until a final public sanction is issued.
- **A right of appeal** to the Standards Council as described in Section 15.4.

15.3 Complaint Procedure

The formal complaint procedure is as follows:

1. **Intake.** Complaints are submitted to the Professional Responsibility Office through published channels. Anonymous complaints are accepted for triage; identified complaints are required for adjudication.
2. **Jurisdictional screen.** The Office confirms that the subject holds or has held a CAIS credential and that the alleged conduct falls within the scope of the Code.
3. **Preliminary investigation.** The Office reviews the allegation, gathers available evidence, and determines whether the matter is credible and actionable.
4. **Referral.** Credible complaints are referred to a Disciplinary Panel. Non-credible complaints are dismissed with written reasons.
5. **Panel adjudication.** The Disciplinary Panel — composed of three Standards Council designees with no material conflict — hears the matter, reviews evidence, and renders a decision.
6. **Decision.** Written decision issued within 90 days of referral, subject to extension for complex matters.

15.4 Appeal Procedure

A certificant aggrieved by a Disciplinary Panel decision may file a single written appeal to the Standards Council within 30 days of the decision. Grounds for appeal are limited to:

- Procedural error material to the outcome
- Manifest factual error on the record
- Sanction disproportionate to the finding
- New evidence not reasonably available at the original hearing

The Standards Council appoints an Appeal Panel of three members, none of whom served on the original Disciplinary Panel. The Appeal Panel may affirm, modify, or remand. Appeal decisions are final.

15.5 Revocation and Registry

Revoked credentials are publicly noted in the Verification Registry as "Revoked," with the effective date and (where the Panel has so ordered) the category of violation. A revoked professional may not use the CAIS designation in any form and may not reapply for a minimum of five years. Suspensions are noted as "Suspended" with end date. Reprimands may be private or public at the direction of the Panel.

15.6 Fairness, Non-Discrimination & Independence

All enforcement activities are conducted without regard to race, color, religion, national origin, sex, gender identity, sexual orientation, disability, age, or any other protected characteristic. The Professional Responsibility Office operates independently of GAISB commercial functions; decisions are not subject to commercial influence, client pressure, or Standards Council member self-interest. Recusal protocols are documented in the GAISB Conflicts of Interest Policy (Appendix G).

SECTION XVI

Continuing Professional Education

Because AI is a rapidly evolving field, credential currency requires ongoing practice and learning. Every CAIS-certified professional is required to earn **40 Continuing Professional Education (CPE) credits per calendar year**, distributed across defined categories. CPE records are submitted annually and spot-audited.

16.1 CPE Categories

CATEGORY	TITLE	ANNUAL MINIMUM	EXAMPLES
A	Technical Currency	10 CPE	New model studies, framework deep-dives, technical papers, approved technical workshops.
B	Applied Practice	10 CPE	Deployed builds, case studies, operational improvements to existing systems, logged build tasks.
C	Ethics & Governance	5 CPE	Updated framework training, case-law review, responsible-AI tabletop exercises, audit experience.
D	Community Contribution	5 CPE	Teaching, mentoring, peer review, published writing, conference speaking, standards-body service.
E	Elective	10 CPE	Any approved activity across Categories A–D or approved cross-disciplinary education.

16.2 CPE Accreditation

CPE credit is earned through:

- GAISB-branded events, courses, and workshops (automatic credit)
- GAISB-accredited third-party providers (automatic credit)
- Non-accredited activities (self-reported, subject to audit)

16.3 Non-Compliance

Failure to meet annual CPE requirements results in credential status change to "Delinquent." Delinquency exceeding 12 months triggers "Suspension." Continued non-compliance for 24 months results in revocation under Canon VII of the Code.

SECTION XVII

Framework Crosswalk

The CBK is designed to align with the major global AI governance frameworks. This crosswalk allows organizations and regulators to map CAIS competency directly to the frameworks they use. It is also the basis on which GAISB engages with national and supranational standards bodies.

17.1 NIST AI Risk Management Framework (RMF) 1.0

NIST FUNCTION	PRIMARY CBK DOMAINS	COVERAGE
Govern	D1, D5, D6	Strategic, organizational, and policy context for AI risk management.
Map	D2, D5, D7	Contextualizing AI risks across technical capability and deployment context.
Measure	D3, D5, D7	Analyzing, assessing, benchmarking, and monitoring AI risk.
Manage	D3, D4, D5, D7	Prioritizing and responding to mapped and measured risks.

17.2 EU AI Act

EU AI ACT ELEMENT	PRIMARY CBK DOMAINS	COVERAGE
Risk Categorization (Art. 6)	D5, D6	Classifying systems as prohibited, high-risk, limited-risk, or minimal-risk.
High-Risk Requirements (Art. 8–15)	D3, D4, D5, D7	Risk management, data governance, documentation, transparency, human oversight, accuracy and robustness.
Transparency Obligations (Art. 50)	D3, D5, D7	Disclosure, labeling, and user-facing transparency for generative systems.
Post-Market Monitoring (Art. 72)	D5, D7	Ongoing monitoring of deployed AI systems.
General-Purpose AI (Art. 53–56)	D2, D5	Obligations applicable to providers of general-purpose AI models.

17.3 ISO/IEC 42001 — AI Management Systems

ISO/IEC 42001 CLAUSE	PRIMARY CBK DOMAINS	COVERAGE
Context of the Organization (Cl. 4)	D1, D6	Organizational environment, stakeholder expectations, scope.
Leadership (Cl. 5)	D6	AI management policy, roles, and executive commitment.
Planning (Cl. 6)	D5, D6	Risk and opportunity treatment, AI objectives.
Support (Cl. 7)	D6, D7	Resources, competence, awareness, documented information.
Operation (Cl. 8)	D3, D4, D7	Operational planning and control, impact assessment.
Performance Evaluation (Cl. 9)	D5, D7	Monitoring, measurement, audit, management review.
Improvement (Cl. 10)	D5, D7	Nonconformity correction and continual improvement.

17.4 OECD AI Principles

OECD PRINCIPLE	PRIMARY CBK DOMAINS	COVERAGE
Inclusive Growth, Sustainable Development & Well-being	D1, D6	Strategic framing and enterprise impact.
Human-Centered Values & Fairness	D5	Bias mitigation, dignity, rights protection.
Transparency & Explainability	D5, D7	Model cards, system cards, user-facing transparency.
Robustness, Security & Safety	D3, D4, D5	Reliability, adversarial resilience, safety controls.
Accountability	D5, D6	Governance, audit trails, professional conduct.

SECTION XVIII

Recognition of Prior Learning & Credit Equivalency

GAISB recognizes that many candidates for the CAIS credential arrive with substantial prior learning, professional experience, and adjacent credentials. The Recognition of Prior Learning (RPL) framework provides structured pathways to acknowledge that prior competency without compromising the defensibility of the credential.

18.1 Recognition Categories

CATEGORY	WHAT IT RECOGNIZES	EFFECT ON CAIS PATHWAY
Recognized Credentials	Professional credentials with demonstrated content alignment to one or more CAIS domains (selected cloud AI credentials, recognized AI governance credentials, selected academic certificates).	Domain-level exemption from designated Build Tasks (maximum of two domains).
Recognized Academic Programs	Degree or graduate-level concentrations in AI, ML, or data science from accredited institutions.	Credit toward D2 knowledge requirements; academic Build Task equivalents reviewed case-by-case.
Documented Professional Practice	Evidenced shipped AI systems, deployed agents, or production LLM applications meeting CBK Build Task standards.	Portfolio-based equivalency review; may satisfy up to 50% of Build Task portfolio.
Challenge Entry	Candidates with strong documented practice may challenge directly at Tier II or III upon demonstration.	Direct-entry assessment pathway under Section XIV progression rules.

18.2 Continuing Professional Education Equivalency

CPE credits may be earned through recognized third-party providers whose programs have been reviewed by the GAISB CPE Accreditation Committee and mapped to CBK domain coverage. Accredited providers are listed publicly in the CPE Provider Directory.

18.3 Academic Credit Articulation

GAISB pursues selective articulation agreements with recognized universities and professional education bodies. Where articulated, CAIS courses and credentials may confer academic credit toward designated programs, subject to each partner institution's articulation terms. Articulation is bidirectional: recognized academic programs may confer CBK learning-outcome equivalency in return.

18.4 Limits on Recognition

Prior Learning Recognition never waives:

- The written CAIS Examination
- The CAIS Capstone Project (for Tier II and above)
- The Code of Professional Conduct commitment
- Continuing Professional Education obligations

This ensures that every CAIS holder — regardless of entry pathway — meets the same terminal competency standard.

SECTION XIX

Accreditation & External Recognition Roadmap

GAISB is committed to establishing the CAIS credential as an externally accredited professional certification, recognized by national and international accreditation bodies. This section sets out the accreditation roadmap as a matter of public commitment.

19.1 Target Accreditation Standards

STANDARD	SCOPE	STATUS
ISO/IEC 17024	General requirements for bodies operating certification of persons.	Target: Edition 2.0
ANSI/ASTM E2659 (U.S.)	Standard practice for certificate programs.	Under evaluation
IAS (International Accreditation Service)	Accreditation body for ISO 17024 certification.	Target accreditation body
EQF (European Qualifications Framework) Referencing	Level-referencing CAIS against the EQF 8-level framework for EU recognition.	Target: Edition 2.0

19.2 Accreditation Milestones

1. **Pre-accreditation readiness.** Establishment of the Standards Council, publication of the CBK, JTA completion, documented policies (code, appeals, conflicts, proctoring, RPL). *Edition 1.0.*
2. **Pilot administration.** First full cycle of administrations with published psychometric reports and post-administration reviews.

3. **Accreditation application.** Formal application to a recognized personnel-certification accreditation body against ISO/IEC 17024.
4. **On-site assessment.** External assessor review of governance, processes, documentation, psychometric practice, and impartiality controls.
5. **Accreditation award and surveillance.** Initial accreditation; ongoing surveillance audits per accreditor cycle; re-accreditation on cycle.

19.3 Regulatory & Institutional Recognition

In parallel with formal accreditation, GAISB pursues recognition through the channels by which AI practitioner competency is most likely to be institutionally referenced:

- **Public sector procurement** — Inclusion of CAIS as a recognized practitioner credential in public AI procurement standards.
- **Enterprise procurement** — Enterprise vendor-qualification criteria referencing CAIS.
- **Regulator engagement** — Formal engagement with AI governance authorities (EU AI Office, national AI bodies) to position CAIS as the recognized practitioner-level credential under their respective frameworks.
- **Higher-education articulation** — Bilateral articulation with universities for reciprocal credit.
- **Professional body alignment** — Memoranda of understanding with adjacent professional bodies (risk management, cybersecurity, data management, compliance) for reciprocal recognition.

19.4 Transparency Commitments

GAISB publishes, at minimum:

- This CBK, in full, with every revision
- The JTA Report for each CBK edition
- Annual psychometric summaries of examinations
- The public Verification Registry
- Annual enforcement summary of disciplinary activity (aggregated, anonymized where appropriate)
- Standards Council roster, charter, and conflicts-of-interest policy
- Annual financial and governance summary

These commitments meet or exceed the transparency expectations of ISO/IEC 17024 and are maintained irrespective of the current formal accreditation status.

SECTION XX

Standards Council Charter & Governance

20.1 Purpose and Authority

The **GAISB Standards Council** is the governing body of the CAIS credential and the ultimate authority on the Common Body of Knowledge, the Code of Professional Conduct, examination blueprint, passing standards, capstone standards, and accreditation matters. The Council operates independently of GAISB commercial functions.

20.2 Composition

The Standards Council is composed of between 15 and 25 members, appointed on staggered three-year terms, drawn from the following constituencies:

- **Enterprise practitioners** — senior AI leaders in operating organizations
- **Independent practitioners** — consultants and builders with shipped practice
- **Academic members** — faculty in AI, ML, or related disciplines
- **Government and regulator liaisons** — members engaged with AI governance
- **Ethics and public-interest members** — including at least one non-industry member representing the public interest
- **Geographic representation** — members from a minimum of five world regions

20.3 Council Committees

Substantive work is delegated to standing committees, each chaired by a Council member and populated with Council members and external subject-matter experts:

- **Examinations Committee** — Blueprint maintenance, item writing oversight, psychometric review, passing-standard methodology.
- **Capstone & Assessment Committee** — Capstone standards, rubric maintenance, review-panel calibration.
- **Ethics & Professional Responsibility Committee** — Code maintenance, enforcement oversight.
- **Research & JTA Committee** — Job Task Analysis, practice research, CBK revision.
- **Accreditation & Recognition Committee** — External accreditation, articulation agreements, regulator engagement.
- **CPE & Provider Accreditation Committee** — CPE standards, provider accreditation.

20.4 Impartiality and Conflicts of Interest

Every Council member signs an annual Conflicts of Interest declaration (Appendix G) and discloses any commercial, consulting, or employment relationships that could affect judgment. Members recuse from matters in which they hold a material conflict. Items, capstone reviews, and disciplinary matters are rotated across members to prevent concentration of decision-making.

20.5 Decision Rules

The Council adopts substantive decisions (CBK revisions, passing-standard decisions, new domains, material policy changes) by a documented supermajority (two-thirds). Routine administrative matters are adopted by simple majority. All decisions and dissents are recorded in meeting minutes, published in summary form.

20.6 Terms & Rotation

Council members serve three-year terms, renewable once. Staggered appointments ensure continuity while preventing entrenchment. The Chair is elected by the Council from among its members for a two-year term, renewable once.

20.7 Public Accountability

The Standards Council publishes, annually: a public report on Council activity, an examination psychometric summary, an enforcement summary, the current roster, and any material changes to the CBK, Code, or governance. Material changes are subject to a public-comment window before ratification.

SECTION XXI

Appendices

Appendix A — Glossary of Key Terms

The CBK uses the following terms with specific meanings. This glossary is authoritative and takes precedence in cases of ambiguity within this document.

TERM	DEFINITION
Agent	An AI system that plans, acts on tools, and iterates toward a goal, rather than responding to a single prompt.
Build Task	A mandatory performance-based assessment associated with a CBK domain.
Capstone	The integrating assessment project required for Tier II and above.
CBK	This document — the Common Body of Knowledge.
Code	The GAISB Code of Professional Conduct (Section XV).
CPE	Continuing Professional Education credit, measured under Section XVI.
Eval	A structured evaluation harness measuring AI system performance quantitatively against a test set.
Framework	In CBK context, an external AI governance or risk framework (NIST AI RMF, EU AI Act, ISO 42001, OECD).
Guardrail	A control applied to an AI system to constrain inputs, outputs, tool use, or behavior.
LLMOps	The operational discipline of running LLM-powered systems reliably in production.
RAG	Retrieval-Augmented Generation: the architectural pattern of supplementing LLM context with retrieved information.
Standards Council	The governing body of GAISB responsible for the CBK.
Tier	One of the four progression levels of the CAIS credential (Practitioner, Builder, Operator, Architect).
Verification Registry	The public lookup through which any party may verify the status of a CAIS certificant.

Appendix B — Development Methodology Summary

Edition 1.0 of the CBK was developed through the six-stage methodology described in Section 1.5. The Job Task Analysis underlying Edition 1.0 drew on contributions from practitioners in 100+ countries representing enterprise, academia, government, independent practice, and education. The full JTA report is published separately in the GAISB Research Repository.

Appendix C — Revision and Review Cycle

The CBK is reviewed every 24 months. Interim amendments may be issued for material change in the AI landscape and are published as dated addenda to the current Edition. Revision history, contributor acknowledgments, and the rationale for each change are maintained in the public Revision Register.

Appendix D — Citation

When citing this document in academic, professional, or institutional contexts, the preferred citation is:

> *Global AI Standards Body. The CAIS Common Body of Knowledge, Edition 1.0. GAISB Publications, 2026.*

Appendix E — Contact & Contribution

The CBK is a living professional standard. The Standards Council invites:

- Commentary and critique during formal review periods
- Proposed amendments with supporting rationale
- Petitions for new domains or learning outcomes driven by practice evolution
- Applications to serve on the Standards Council or invited expert panels

All contributions and correspondence are directed through the GAISB Standards Office, through channels published on the GAISB website.

Appendix F — Document Control & Version History

This CBK is a controlled document. Every release, amendment, and interim revision is versioned, dated, and logged. The table below is the authoritative revision log for this CBK.

EDITION	RELEASE DATE	NATURE OF REVISION	APPROVAL
1.0	2026	Initial publication. Seven-domain architecture; JTA Edition 1.0; Code of Professional Conduct; CPE framework; framework crosswalks.	Standards Council
1.x	As issued	Interim amendments for material change. Numbered per revision.	Standards Council
2.0	Target	Full revision on biennial cycle. Reassessment of domain weightings against updated JTA.	Standards Council

All revisions are maintained in the public Revision Register with contributor acknowledgments, rationale, and date of effect.

Appendix G — Conflicts of Interest Policy (Summary)

The independence and impartiality of the Standards Council, examination development, capstone review, and disciplinary procedures are protected by the GAISB Conflicts of Interest Policy. Summary provisions:

- **Annual disclosure.** Every Standards Council member, SME panelist, item writer, and reviewer signs an annual conflict disclosure covering employment, consulting, advisory, equity, and related interests.
- **Recusal triggers.** Members recuse from matters where they, a near relative, or an employer/client has a material interest in the outcome.
- **Segregation of commercial activity.** GAISB commercial functions (training delivery, event production, content licensing) operate independently of standards setting, examination, and disciplinary functions.
- **Vendor-independence.** No AI vendor, cloud provider, model company, or tool provider holds decision rights over CBK content, examination content, or disciplinary outcomes. Sponsorships are disclosed and may not be tied to content decisions.
- **Decision documentation.** Recusals and conflict declarations are recorded in Council minutes.
- **Enforcement.** Material violations of the policy are grounds for removal from the Council and, where applicable, reporting to professional authorities.

The full policy is published separately.

Appendix H — Psychometric & Technical Standards Reference

The CAIS examination program conforms to recognized standards for high-stakes personnel testing, including:

- *Standards for Educational and Psychological Testing* (AERA/APA/NCME)

- *ISO/IEC 17024* — General requirements for bodies operating certification of persons
- *ITC Guidelines on the Use of Tests* (International Test Commission)
- *ISO/IEC 23988* — Use of IT in delivery of assessments

All technical reports related to the examination program (psychometric summaries, form equating reports, DIF reviews) are published in the GAISB Technical Report Series.

Appendix I — Endorsements & Institutional Recognition (Reserved)

As the CAIS credential progresses through the Accreditation Roadmap in Section XIX, this appendix will publish accreditor decisions, institutional recognitions, articulation agreements, and endorsements from regulators, enterprises, academic institutions, and professional bodies. All recognitions are verifiable through the linked sources. Reserved for Edition 1.1 and subsequent.

Appendix J — Authoritative References

The following external standards, frameworks, and sources inform this CBK and are referenced normatively where indicated in Sections V–XVII:

- National Institute of Standards and Technology, *AI Risk Management Framework (AI RMF 1.0)*
- European Union, *Regulation (EU) 2024/1689 (the AI Act)*
- International Organization for Standardization, *ISO/IEC 42001 — AI Management Systems*
- Organisation for Economic Co-operation and Development, *OECD AI Principles*
- International Organization for Standardization, *ISO/IEC 17024 — Personnel Certification*
- American Educational Research Association et al., *Standards for Educational and Psychological Testing*

GAISB

The Global AI Standards Body — Governing the Professional Standard for AI Practitioners Worldwide.

CAIS COMMON BODY OF KNOWLEDGE
EDITION 1.0 • 2026 • PUBLIC DOCUMENT
ALIGNED TO NIST AI RMF · EU AI ACT · ISO/IEC 42001 · OECD AI PRINCIPLES
DEVELOPED UNDER ISO/IEC 17024 PRINCIPLES FOR BODIES OPERATING CERTIFICATION OF PERSONS

21 SECTIONS · 10 APPENDICES · 7 DOMAINS · 8 CRITICAL WORK DOMAINS · JTA EDITION 1.0