AIGP Exam at a Glance

Artificial Intelligence Governance Professional · Body of Knowledge v2.1 · aigovernance.study

The exam

Format	100 questions in 3 hours: 85 scored, 15 unscored pilots. Multiple-choice and multi-select, mostly scenario-based. Pearson VUE, at a test centre or online.
Scoring	Scaled from 100 to 500; 300 to pass. No penalty for a wrong answer. Pass or fail and your score show on screen at the end, with a per-domain breakdown by email.
Cost	US$799, or US$649 for IAPP members. A retake is around US$625 (US$475 for members) and cannot be booked sooner than seven days after the failed attempt.
Keeping it	Two-year term: 20 CPE credits plus a US$250 maintenance fee, waived while an IAPP membership is active.
Syllabus	The IAPP's AIGP Body of Knowledge, version 2.1, a free download from iapp.org.

Scored questions per domain

Domain	Covers	Questions
Domain I	the foundations of AI governance	16–20
Domain II	laws, standards and frameworks	19–23
Domain III	governing AI development	21–25
Domain IV	governing AI deployment and use	21–25

A six-week plan

Week	Focus	Goal
1	Domain I: the foundations of AI governance	What an AI system is, the harms taxonomy, the AI life cycle.
2	Domain II: laws, standards and frameworks	EU AI Act, NIST AI RMF, ISO standards: who each binds and what it requires.
3	Domain III: governing AI development	Risk assessment, data governance, testing and evaluation.
4	Domain IV: governing AI deployment and use	Monitoring, incident response, transparency obligations.
5	Mixed practice	Question sets across domains; drill the weakest domain.
6	Timed mocks and vocabulary	Full timed mocks for pacing; sweep the look-alike terms.

How to use practice

- After each topic, answer its practice questions while it is fresh, and read the explanation even when you were right.
- Track your score per domain. A good average can hide one weak domain.
- Before booking, aim for around 70% or better on full timed mocks, more than once and in every domain.

Independent study aid, not affiliated with or endorsed by the IAPP. Exam facts from the IAPP Candidate Handbook and iapp.org; check iapp.org before booking. Scoring explained: aigovernance.study/aigp-passing-score/

Domain I: the distinctions the exam tests

Scenario questions often turn on two terms that sound alike. Know the one-line difference.

Machine learning *vs* Data mining	ML's defining trait is a system improving its own performance from experience; data mining is the distinct activity of discovering previously unknown properties in data.
Deep learning *vs* Machine learning	Deep learning is a subset of ML that uses multi-layered neural networks; all deep learning is ML, but not all ML is deep learning.
Narrow AI (ANI) *vs* Artificial general intelligence (AGI)	Narrow AI performs defined tasks and is everything in production today; AGI is hypothetical human-level breadth of cognition across contexts, and does not yet exist.
Artificial general intelligence (AGI) *vs* Artificial superintelligence (ASI)	AGI closely mimics full human cognition; ASI exceeds human intellect across virtually all fields and is self-aware — both are hypothetical, only narrow (and broad) AI exist today.
Provider *vs* Deployer	The provider places the system on the market under its own name; the deployer uses it under its own authority — but substantially modifying a high-risk system can turn a deployer into a provider.
Developer *vs* Provider	A developer designs, builds or trains the model; a provider is whoever brings it to market under its own name — often the same entity, but a company that white-labels another's model becomes a provider without having developed it.
Deployer *vs* User	A deployer uses the system under its own authority in a professional context; a user is anyone who merely interacts with or is affected by it — an employee using an internal tool is a user, while their employer is the deployer.
Centralized governance *vs* Decentralized governance	Centralized means one point is responsible for all AI decisions; decentralized ('local governance') delegates decision-making to lower levels within central policy, with hybrid mixing the two.
Policy *vs* Procedure	A policy states what must happen and why (the rule); a procedure states how, step by step — match the request to the right instrument.
Map *vs* Measure	Map establishes context and identifies risk (documenting purpose and stakeholders); Measure quantifies and tracks those risks across technical, societal and organizational domains.
Measure *vs* Manage	Measure quantifies risk severity, likelihood and scope; Manage prioritizes and acts on it — mitigate, transfer, avoid or accept.
Operational risk *vs* Security risk	Poor access management that causes data corruption is operational risk; an active attack exploiting a vulnerability is security risk.

Common wrong assumptions

- **Treating 'cognitive learning' as a machine-learning category.** It is a distractor — the recognized training approaches are supervised, unsupervised, reinforcement and semi-supervised learning; deep learning is a related architecture, not a training-approach category.
- **Assuming a company that merely buys and fine-tunes a third-party model always remains just a deployer.** Substantially modifying a high-risk system — or modifying a system into high-risk status — converts a deployer into a provider with the fuller provider obligation set.
- **Deferring all documentation and incident handling to the end of the project.** Risk management, documentation and incident management all run continuously rather than at one stage, so an option that defers documentation to launch is wrong.

Practise these in scenarios: all 100 Domain I questions are free for 7 days in the app, each with a full explanation.
aigovernance.study/app · Independent study aid, not affiliated with or endorsed by the IAPP.