MockCert

Study guide | IAPP | AIGP

IAPP Artificial Intelligence Governance Professional (AIGP)

Notes for every domain of the IAPP Artificial Intelligence Governance Professional (AIGP) exam, organized by the official exam guide's domains and weights.

By MockCert about 50 min read Updated November 24, 2026

CONTENTS

DOMAIN 1 · 16–20%
Understanding the foundations of AI governance 2

DOMAIN 2 · 19–23%
Understanding how laws, standards and frameworks apply to AI 10

DOMAIN 3 · 21–25%
Understanding how to govern AI development 21

DOMAIN 4 · 21–25%
Understanding how to govern AI deployment and use 29

MockCert is an independent certification preparation platform. It is not affiliated with, endorsed by or operated by IAPP. Exam domains and weights follow the official exam guide. These notes were last updated on November 24, 2026; the notes at mockcert.com are kept up to date as the exam changes.

Domain 1 | 16–20% of the exam | 11 min read

Understanding the foundations of AI governance

Overview

Every point below comes from IAPP's official documentation; each section lists its sources.

Understand what AI is and why it needs governance.

Start here

THINK FIRST

What makes AI different enough from other technology that it needs its own governance approach?

Why it matters: Questions describe an AI system and ask which characteristic, risk type, level of harm or responsible AI principle applies.

AI has traits such as opacity, autonomy, and speed and scale that call for comprehensive governance. Those traits create risks for individuals, groups, organizations and society. Common responsible AI principles give governance its targets.

Key terms

Term	Meaning
Socio-technical	Describes AI systems, which are influenced by societal dynamics and human behavior, not only by code.
Probabilistic output	AI produces probabilistic rather than deterministic outputs, a trait that calls for comprehensive governance.
Explainability	Per NIST, the answer to how an AI system made a decision.
Agentic AI	An innovative AI application facing changing governance rules; a different concept from AI agents.

Types of AI models

AI models are sorted along several contrasts. Learn each pair as a contrast. Generative AI and agentic AI are innovative applications whose governance requirements change rapidly.

Contrasting AI model types named in the AIGP

Contrast	One type	The other type
Method	Classic	Generative
Ownership	Proprietary	Open source
Size	Small	Large
Modality	Language	Multimodal

EXAM TRAP

AI agents and agentic AI are not synonyms. They are two distinct concepts with different meanings.

Why AI needs its own governance

Certain AI traits demand a comprehensive approach. They include complexity, opacity, autonomy, and speed and scale. Potential for harm or misuse and data dependency belong on the list too. AI also gives probabilistic rather than deterministic outputs.

Socio-technical systems

AI systems are socio-technical. Societal dynamics and human behavior influence how they behave. AI can also increase the speed and scale of biases. That amplifies harm to individuals, groups, communities, organizations and society.

Types of risk

Key risk types: misalignment, ethics and bias, complexity and scalability. NIST groups AI bias into systemic, computational and statistical, and human-cognitive categories. Security concerns include adversarial examples, data poisoning and theft of models or training data through endpoints.

EXAM TRAP

Bias does not need intent. NIST says each of its **three** bias categories can occur with no discriminatory intent at all.

Who can be harmed

AI harms reach individuals, groups, organizations and society. For individuals, inference can identify people or expose facts they kept private. At society level, data collection, personalized ads and biometric systems can sway democratic elections. They can also fracture shared reality.

Principles of responsible AI

The AIGP lists six common principles of responsible AI. They are fairness, safety and reliability, privacy and security, and transparency and explainability. Accountability and human-centricity complete the list. Done well, it opens the door to data-driven innovation while keeping risk down.

Transparency is not enough

Transparency, explainability and interpretability answer different questions. Per NIST, transparency answers what happened in a system. Explainability answers how a decision was made, and interpretability answers why.

EXAM TRAP

A transparent system is not automatically fair. NIST notes it is not necessarily accurate, secure, fair or privacy-enhanced either.

Prioritizing safety

Serious safety risks get the most urgent attention. Under the NIST AI RMF, risks that could cause serious injury or death need the most thorough risk management. Valid and reliable is a necessary condition of trustworthiness. It is the base for the other trustworthiness traits.

Key takeaways

- Complexity, opacity, autonomy, and speed and scale are reasons AI needs comprehensive governance.
- AI harms reach four levels: individuals, groups, organizations and society.
- Bias can arise without intent, in systemic, computational and statistical, or human-cognitive forms.
- Transparency shows what happened, but it does not make a system fair or accurate.

Sources: VERSION 2.1 Effective date: 2 February 2026 (opens in a new tab) · Measure - AIRC (opens in a new tab) · AI Risks and Trustworthiness - AIRC (opens in a new tab) · AI agents and agentic AI: What privacy and AI governance leaders must know now | IAPP (opens in a new tab) · Carissa Véliz on privacy, AI ethics and democracy | IAPP (opens in a new tab) · Responsible AI: Putting data ethics into action | IAPP (opens in a new tab)

Establish and communicate organizational expectations for AI governance.

Start here

THINK FIRST

How would you set up roles, training and collaboration so an AI governance program fits your organization?

Why it matters: Questions describe an organization and ask who should own an AI governance task, who to train or how to adapt the program.

AI governance is a system of frameworks, practices and processes that helps stakeholders oversee AI. Making it work takes clear roles, cross-functional teams, training for everyone and an approach that fits the organization.

Key terms

Term	Meaning
AI governance	Organization-level frameworks, practices and processes for implementing, managing and overseeing the use of AI.
Accountability	The duty to make an AI system operate ethically and lawfully, with its outcomes traceable to the responsible entity.
AI risk officer	A role that makes sure AI risks are identified, escalated and managed through oversight and coordination.
Developer and deployer	Labels the AIGP uses for tasks, not for liability, since many organizations fill both roles.
Three lines of defense	A traditional approach to risk that may not suit smaller organizations; their leaders can build a strong risk culture in other ways.

Roles and accountability

Every AI governance stakeholder needs a defined role. Executive leadership takes responsibility for decisions on AI development and deployment risks. Accountability means outcomes trace back to the responsible entity. Some organizations give a designated officer authority, staff and budget for the AI portfolio.

Risk roles escalate

AI risk roles identify, track and escalate risk. They report risks to governance bodies and leadership. They need skill in risk-assessment methods, scenario planning and clear communication. Board committees for AI risk can sit inside enterprise risk management.

EXAM TRAP

An AI risk role is not the team that fixes problems. It passes risks up to governance bodies and leaders but does not directly own the fix.

Cross-functional collaboration

Collaboration brings efficacy and diverse expertise. Multidisciplinary teams can draw on data science, software development, civil liberties, privacy, security, legal and risk management. Involve legal, compliance and risk from the start of design. Bring in outside experts when internal teams lack diverse experience or disciplines.

Separate build from test

Keep development and testing under separate management. The NIST AI RMF Playbook suggests this split so problems can be corrected independently. It stops one team from grading its own work.

Training and awareness

Train all stakeholders on AI terminology, strategy and governance. Tailor content to each group, such as developers versus legal, compliance and audit. Ongoing education can cover AI laws, potential harms, internal AI policies and trustworthy AI traits. In one IAPP survey, **33%** named missing professional training and certification as a key challenge.

EXAM TRAP

Training is not only for technical staff. The program reaches all stakeholders, with content suited to each group's role.

Fit the approach to the organization

Tailor governance to size, maturity, industry and risk tolerance. Products and services and objectives shape the approach too. More than **50%** of surveyed teams designing AI governance built on their privacy programs. Reusing what already works is a common starting point.

EXAM TRAP

Not every organization needs the traditional three lines of defense approach to risk. It may be impractical for smaller organizations, whose leaders can build a strong risk culture through other means.

Roles in the AI value chain

Developers, providers, deployers and users differ in duties and needs. They differ in opportunities as well. Some laws split developer from deployer to assign liability. The table gives example EU AI Act duties for high-risk systems by role.

Example EU AI Act duties by role for high-risk systems

Role	Duty	Detail
Provider	Conformity assessment	Under Article **43**, before market placement or putting into service
Deployer	Keep automatic logs	At least **six months**, to the extent under their control
Deployer that is an employer	Inform workers	Tell workers' representatives and affected workers before workplace use
Public authority deployer	Check registration	Do not use a high-risk system missing from the EU database
Authorized representative	Keep documents available	For **10** years after market placement or putting into service

Small businesses under the EU AI Act

The EU AI Act eases the load on SMEs. Conformity assessment fees for SME providers shrink in proportion to their size. Member States give eligible SMEs priority access to regulatory sandboxes. They also run awareness and training tailored to SMEs and deployers.

Key takeaways

- Leadership owns decisions on AI risk, while risk roles escalate risks rather than remediate them.
- Cross-functional collaboration serves efficacy and diversity of expertise and perspective.
- Train all stakeholders on AI terminology, strategy and governance, tailored by role.
- Adapt governance to company size, maturity, industry, objectives, risk tolerance, and products and services.
- Developer and deployer describe tasks, and many organizations do both.

Sources: Key terms for AI governance | IAPP (opens in a new tab) · VERSION 2.1 Effective date: 2 February 2026 (opens in a new tab) · AI Governance Professional Profiles: Risk Management | IAPP (opens in a new tab) · Govern - AIRC (opens in a new tab) · Professionalizing Organizational AI Governance Report | IAPP (opens in a new tab) · Article 16: Obligations of providers of high-risk AI systems | AI Act Service Desk (opens in a new tab) · Article 26: Obligations of deployers of high-risk AI systems | AI Act Service Desk (opens in a new tab) · Article 22: Authorised representatives of providers of high-risk AI systems | AI Act Service Desk (opens in a new tab)

Establish policies and procedures to apply throughout the AI life cycle.

Start here

THINK FIRST

How do you write AI policies that reach every life cycle stage and keep third-party AI risk under control?

Why it matters: Questions describe an organization adopting AI and ask which policy, assessment or contract step fits a life cycle stage or a vendor risk.

Oversight policies must reach every stage of the AI life cycle, from use case assessment to incident management. Existing policies are evaluated and updated for AI. Most businesses license AI, so policies, assessments and contracts for third parties carry much of the load.

Key terms

Term	Meaning
AI life cycle	The stages from use case assessment and data acquisition through deployment, monitoring and incident management.
Third-party risk	Risk from vendors and licensed AI, managed with policies, assessments and contracts.

Term	Meaning
Privacy by Design	Also called Data Protection by Design; in a vendor, default settings that favor privacy.
Privacy-enhancing technologies	Tools such as differential privacy, federated learning and synthetic data, used at each life cycle stage.

Cover every life cycle stage

Oversight policies must reach all AI life cycle stages. The stages start with use case assessment, risk management, ethics by design, and data acquisition and use. They continue through development, training and testing, deployment and monitoring, documentation and reporting, and incident management.

EXAM TRAP

Policies do not stop once a system is deployed. Monitoring, documentation and reporting, and incident management are life cycle stages too.

Update the policies you have

Evaluate and update existing policies for AI. Start with data privacy, security, data governance and intellectual property. Keep privacy and security working together. With AI, one incident can involve data exposure, model misuse and insider access at once.

Privacy risks along the cycle

Each stage brings its own privacy risk. Model memorization raises privacy risk at training time. Membership inference attacks can cause privacy problems during deployment. Collecting training data at scale must respect user consent and data minimization. Privacy-enhancing technologies can be applied at each stage.

Why third parties matter

Most businesses license AI rather than build it. One IAPP article says **99%** of businesses are not building their own models. Licensees often lack in-house AI expertise and cannot vet black box tools that keep changing. They also have little leverage to negotiate vendor terms.

Tools for third-party risk

Manage third-party risk with policies, assessments and contracts. These cover areas such as procurement, supply chain, human resources and acceptable use. Contracts matter because vague language and unmitigated risks can undermine a vendor relationship when written provisions leave them out.

EXAM TRAP

Third-party risk is not handled by contracts alone. Policies and assessments sit beside contracts as the tools for managing it.

Vendor due diligence

Ask what data the vendor touches and how it protects it. Find out what data is shared with or collected by the vendor, and what it may do with it. Ask where the data is stored, how long it is kept and how it is deleted. Check for privacy-friendly defaults and an incident response plan.

During and after the relationship

Vendor oversight continues after the contract is signed. Ongoing monitoring covers performance, financial stability, compliance and data handling. Vendors that had an expedited review may need an earlier reassessment. Ending a relationship means weighing dependencies, contractual obligations and transition risks.

Key takeaways

- Oversight policies cover every AI life cycle stage, from use case assessment to incident management.
- Evaluate and update existing privacy, security, data governance and intellectual property policies for AI.
- Manage third-party AI risk with policies, assessments and contracts.
- Due diligence asks what data a vendor gets, where it is stored and how long it is kept.

Sources: VERSION 2.1 Effective date: 2 February 2026 (opens in a new tab) · Beyond the frontier: Pragmatic AI risk management for the rest of us | IAPP (opens in a new tab) · A Practitioner Approach to Implementing Data Protection & Privacy by Design | IAPP (opens in a new tab) · Expedited Vendor Privacy and Security Assessment Checklist | IAPP (opens in a new tab) · Proactive AI privacy – Leveraging PETs throughout the AI life cycle | IAPP (opens in a new tab) · The privacy-security convergence: Building unified AI risk controls | IAPP (opens in a new tab) · Third-Party Vendor Management Means Managing Your Own Risk | IAPP (opens in a new tab)

Domain 2 | 19–23% of the exam | 16 min read

Understanding how laws, standards and frameworks apply to AI

Overview

Every point below comes from IAPP's official documentation; each section lists its sources.

Understand how existing data privacy laws apply to AI.

Start here

THINK FIRST

When an AI project processes personal data, which existing privacy-law duties still apply, and how?

Why it matters: Questions describe an AI use of personal data and ask which privacy requirement or controller obligation it triggers.

AI does not escape existing privacy law. Its core requirements still apply, from transparency and lawful basis to data minimization and privacy by design. Controllers keep their obligations, and sensitive data such as biometrics carries extra rules.

Key terms

Term	Meaning
Privacy by design	An approach resting on **seven** foundational principles set out by Ann Cavoukian.
Positive-sum	A privacy by design principle: stakeholders share one set of objectives that drives design and implementation.
Special category data	Sensitive data with its own requirements in AI; biometrics are the AIGP's example.
Data lineage and provenance	Part of AI data governance, which calls for establishing and documenting them.

Core privacy principles still apply

Transparency, choice, lawful basis and purpose limitation apply to AI. Data minimization and privacy by design apply as well. AI data governance checks and records your legal right to collect and use each dataset. It also documents where the data came from.

Privacy by design in practice

Privacy by design starts at the earliest design stage. In practice, you decide which questions to ask and when, and who joins each decision. You also choose how to evaluate privacy risks. Finally, you shape how data is collected and managed.

EXAM TRAP

Privacy by design is not a late add-on. Late, narrow fixes like encryption or notice-and-choice do not replace building privacy in from the start.

Privacy and innovation

Privacy and security should support innovation, not block it. Some developers see them as barriers. Yet both sets of principles must be embraced to build new solutions. The positive-sum principle helps, because stakeholders work from one shared set of objectives.

Controller obligations

Controller obligations follow personal data into AI. They cover data subject rights, cross-border transfers, privacy impact assessments and the use of third-party processors. Record keeping, breach notification, incident management and automated decision making apply too. Privacy impact assessments and data mapping help put the GDPR into practice.

Cross-border transfers

Transfer restrictions arise in several regions. They exist in the U.S., under the EU and U.K. GDPR, and across Asia-Pacific. An executive order from President Joe Biden put the EU-U.S. Data Privacy Framework into effect. Data intelligence and intelligent access controls can automate safe data sharing.

Sensitive and special categories

Biometric data carries its own requirements in AI. Biometrics are the AIGP's example of a sensitive or special category. Litigation on AI and biometrics is still developing. Article **4a(2)** lets certain AI providers and deployers exceptionally process special category data to detect bias.

EXAM TRAP

That permission is not a mandate. Article **4a(2)** allows processing special category data for bias detection but creates no obligation to carry it out.

Risk-based thinking in both regimes

The GDPR and the EU AI Act both scale duties with risk. The GDPR ties safeguards to the risk of each processing activity. The EU AI Act ties obligations to the risk category of an AI use.

How the GDPR and the EU AI Act apply risk

Feature	GDPR	EU AI Act
Risk is judged against	The data processing activity	The use of AI, not a particular technology
Impact assessment	Data protection impact assessment	Fundamental rights impact assessment, for first use

EXAM TRAP

A FRIA does not always start from scratch. If a DPIA already meets the obligation, the deployer may cross-reference or include its relevant parts.

Update existing policies

Review your existing policies for AI. Evaluate and update policies on data privacy, security, data governance and intellectual property. Privacy is now a core operating discipline across the technology lifecycle, not just a compliance box.

Key takeaways

- Transparency, choice, lawful basis and purpose limitation all apply when AI uses personal data.
- Privacy by design builds privacy in at the earliest design stage, not after the fact.
- Controller duties include PIAs, processors, cross-border transfers, data subject rights and breach notification.
- Biometrics are the AIGP's example of special category data with extra requirements in AI.

Sources: VERSION 2.1 Effective date: 2 February 2026 (opens in a new tab) · The 7 Principles of Privacy by Design in Practice | IAPP (opens in a new tab) · Applying the Positive-Sum Principle for Successful Privacy by Design Outcomes | IAPP (opens in a new tab) · Check or Mate? Strategic Privacy by Design | IAPP (opens in a new tab) · Strategic Privacy by Design | IAPP (opens in a new tab) · PIAs and Data Mapping – Operationalizing GDPR and Privacy by Design | IAPP (opens in a new tab) · Privacy beyond checkmarks: How to Navigate Cross-border Transfers | IAPP (opens in a new tab) · The EU-US Data Privacy Framework and next steps for data transfers | IAPP (opens in a new tab)

Understand how other types of existing laws apply to AI.

Start here

THINK FIRST

Before any AI-specific law applies, which older laws can already make an AI system unlawful?

Why it matters: Questions describe an AI use and ask which existing law it engages: intellectual property, nondiscrimination, consumer protection or product liability.

AI is not a legal vacuum. Intellectual property, nondiscrimination, consumer protection and product liability laws already reach it. Each family targets a different kind of harm, from training data to defective products.

Key terms

Term	Meaning
Unfair and deceptive practices	Acts or practices that consumer protection laws prohibit, including when AI is involved.
Strict liability	The regime the reformed EU Product Liability Directive uses to compensate damage caused by a product defect.
Security-by-design	A Cyber Resilience Act duty on manufacturers to build products with digital elements securely.
Lawful rights to data	Your legal right to collect and use data, which AI data governance checks and records.

Four families of existing law

Four families of existing law already apply to AI. Learn the example for each family. Spot the harm first, then match it to the family that targets it.

How existing laws apply to AI

Law family	How it applies to AI	Related point
Intellectual property	Can prohibit or limit using data for AI training	Generative AI learns from large amounts of web-scraped data
Nondiscrimination	Applies to AI in named contexts	Employment, credit, lending, housing and insurance
Consumer protection	Prohibits unfair and deceptive acts or practices	Overlaps with privacy and antitrust law in enforcement
Product liability	Prohibits design and manufacturing defects	The EU directive compensates defect damage under strict liability

Intellectual property and training data

IP law can prohibit or limit using data to train AI. That matters because generative AI tools learn from large amounts of data scraped from the web. AI data governance should check and record your legal right to collect and use data. Review your IP policies for AI as well.

EXAM TRAP

Having data in hand does not mean you may train on it. IP law can still prohibit or limit using that data for AI training.

Nondiscrimination

Nondiscrimination laws reach AI in five named contexts. Those contexts are employment, credit, lending, housing and insurance. NIST impact assessment tasks include assessing requirements for combating

harmful bias. The same tasks cover product safety, liability and security.

Consumer protection

Consumer protection laws reach AI by banning unfair or deceptive practices. Privacy, consumer protection and antitrust law intersect, as enforcement cases show. Trust is at stake too: **57%** of consumers globally agree AI poses a significant threat to their privacy.

Product liability

Product liability laws target design and manufacturing defects. The revised EU Product Liability Directive updates rules that are **four decades** old for the digital age. It compensates damage from defects under strict liability. It seeks consumer protection alongside legal certainty for businesses and protection for innovation.

EXAM TRAP

Product liability is not only about manufacturing flaws. It also prohibits design defects, so how an AI product is designed matters too.

Product security rules

The Cyber Resilience Act requires security-by-design for digital products. It has extraterritorial reach, and nonconforming products may not be sold in the EU. Manufacturers must fix and report vulnerabilities for years after market placement. The maximum fine is **15 million euros** or **2.5%** of worldwide annual turnover, whichever is higher.

When Cyber Resilience Act obligations apply

Milestone	Date	What applies
Reporting duties	**11 Sept. 2026**	Reports on actively exploited vulnerabilities and severe incidents
Main obligations	**11 Dec. 2027**	Most other obligations

Liability when you own the model

Owning the model raises your legal exposure. A company deploying its own proprietary model takes on more obligations and higher potential liability. Some laws split developer from deployer to assign liability. The AIGP uses the labels only for tasks, since many organizations fill both roles.

EXAM TRAP

Buying a model does not remove the risk. Third-party technology can be complex or opaque, and its risk tolerances may not match yours.

Key takeaways

- IP law can prohibit or limit using data to train AI.
- Nondiscrimination laws apply to AI in employment, credit, lending, housing and insurance.

- Consumer protection laws reach AI through bans on unfair and deceptive acts or practices.
- Product liability laws target design and manufacturing defects, and the EU directive uses strict liability.

Sources: VERSION 2.1 Effective date: 2 February 2026 (opens in a new tab) · EU Product Liability Directive Reform: 101 | IAPP (opens in a new tab) · Top 5 operational impacts of the CRA: Product cybersecurity as a market-access requirement | IAPP (opens in a new tab) · Consumer Perspectives of Privacy and AI | IAPP (opens in a new tab) · Blurred Regulatory Lines — Privacy, Competition and Consumer Protection Collide | IAPP (opens in a new tab) · ? App. A: Descriptions of AI Actor Tasks - AIRC (opens in a new tab)

Understand the main elements of AI-specific laws.

Start here

THINK FIRST

How would you work out which EU AI Act obligations apply to a given AI system and organization?

Why it matters: Questions describe an AI use and an organization's role, then ask which risk tier, requirement or enforcer applies.

AI-specific laws such as the EU AI Act regulate uses, not technologies. A system's risk tier sets its obligations, and the organization's role decides who must meet them. General-purpose AI models follow rules of their own.

Key terms

Term	Meaning
Risk tier	One of four categories in the example framework: prohibited, high, limited or minimal risk.
Provider	The organization that offers a high-risk AI system on the EU market and must meet Chapter III duties.
Notified body	An assessor that checks conformity, issues certificates and audits the provider's quality management system.
General-purpose AI model	A model whose risks depend on how it is used, regulated in Chapter V of the EU AI Act.
AI Office	The EU body with enforcement powers over general-purpose AI models, including fines.

Regulating uses, not technology

The EU AI Act regulates how AI is used. Its risk-based approach looks at what is made of the technology, not the technology itself. Obligations then follow the risk category of each use. The GDPR works in a similar way: its safeguards scale with processing risk.

Four risk tiers

Risk tiers run from prohibited through high and limited to minimal. Each tier covers certain systems or uses. The EU AI Act sets obligations by the risk category of the use, so classify the use first.

Two EU AI Act tiers with example uses

Tier	Example use	Effect
Prohibited	Social scoring; emotion recognition at work or in schools	Banned as a clear threat to safety, livelihoods and rights
High risk	CV-sorting software for recruitment	Provider duties in Chapter III, Sections **2** and **3**

EXAM TRAP

The prohibited tier is not open-ended. The EU AI Act lists **nine** prohibited practices, and social scoring is one of them.

High-risk provider duties

Chapter III sets specific duties for high-risk providers. Sections **2** and **3** set what a provider must meet before making a high-risk system available in the EU. The duties group into organizational, documentation, system design and regulatory requirements.

Conformity assessment

Notified bodies check conformity and quality management. A notified body issues a Union technical documentation assessment certificate once a system meets Chapter III, Section **2**. It then audits the provider's quality management system at intervals.

EXAM TRAP

Conformity assessment is not a one-time check. The notified body keeps auditing whether the provider maintains and applies its quality management system.

Transparency and oversight

People must know when they are dealing with AI. Chatbot users should be told they are talking to a machine. Deep fakes need clear, visible labels. Generative AI providers must make AI-generated content identifiable. Human oversight and quality management are key requirements too.

General-purpose AI models

Chapter V of the EU AI Act covers general-purpose AI models. Their risks depend on the context of use, so no single risk tier fits them. The original April **2021** proposal never mentioned them. Their rules, including transparency and copyright, took effect in **August 2025**.

Technical documentation duties for general-purpose AI models

Item	Applies to	What to describe
General description	Every general-purpose AI model provider	Architecture, number of parameters and license
Adversarial testing	Providers of models with systemic risk	Internal or external measures such as red teaming, where applicable

EXAM TRAP

A general-purpose AI model does not sit in one fixed risk tier. Its risks depend on the context in which it is used.

Roles change obligations

Duties differ for providers, deployers, importers and distributors. Developers, providers, deployers and users differ in responsibilities and needs. The AIGP uses developer and deployer only to describe tasks, since many organizations fill both roles.

EXAM TRAP

Deploying your own proprietary model does not lighten the load. It brings increased obligations and higher potential liability.

Enforcement and penalties

AI-specific laws back their rules with penalties. Chapter IX covers post-market monitoring, information sharing and enforcement. For general-purpose AI models, the AI Office holds enforcement powers and can issue fines for non-compliance.

Key takeaways

- Classify the use first: its risk tier sets the obligations that follow.
- High-risk providers meet Chapter III, Sections **2** and **3**, and notified bodies check conformity.
- General-purpose AI models have their own Chapter V rules, enforced by the AI Office.
- Duties shift with the organization's role: provider, deployer, importer or distributor.

Sources: Top 10 operational impacts of the EU AI Act – Understanding and assessing risk | IAPP (opens in a new tab) · VERSION 2.1 Effective date: 2 February 2026 (opens in a new tab) · Top 10 operational impacts of the EU AI Act – Obligations on providers of high-risk AI systems | IAPP (opens in a new tab) · ANNEX VII | AI Act Service Desk (opens in a new tab) · Top 10 operational impacts of the EU AI Act – Obligations for general-purpose AI models | IAPP (opens in a new tab) · AI Act | Shaping Europe's digital future (opens in a new tab) · ANNEX XI | AI Act Service Desk (opens in a new tab) · Top 10 operational impacts of the EU AI Act – Post-market monitoring, information sharing, and enforcement | IAPP (opens in a new tab)

Understand the main industry standards and tools that apply to AI.

Start here

THINK FIRST

How would you tell the OECD principles, the NIST AI RMF and the core ISO AI standards apart in an exam question?

Why it matters: Questions name an AI standard or tool and ask what it contains, who it is for or how it is organized.

Three families of guidance shape this objective. The OECD sets shared principles and definitions. NIST offers a voluntary risk framework with a companion Playbook, and ISO publishes the core AI standards.

Key terms

Term	Meaning
OECD AI Principles	Principles for trustworthy AI that form part of the OECD Recommendation on Artificial Intelligence.
NIST AI RMF	A voluntary NIST framework for managing AI risk, released as version **1.0** with supplementary documents.
AI RMF Playbook	The companion to the AI RMF, organized like it into core functions, categories and subcategories.
AI assurance	Measuring, evaluating and communicating how trustworthy an AI system is; broader than compliance.

The OECD Recommendation

The OECD AI Principles sit inside a wider Recommendation. That Recommendation also holds the OECD definition of an AI system and the AI lifecycle. The principles were first adopted in **2019** and updated in **May 2024**. They now have **47** adherents.

EXAM TRAP

The OECD AI Principles are not a stand-alone text. They are one part of the OECD Recommendation on Artificial Intelligence, next to the AI system definition and lifecycle.

OECD definitions travel widely

Major regulators reuse the OECD AI system definition. The EU, the Council of Europe, the US and the UN all use it. They also use the OECD AI lifecycle in their laws and guidance.

NIST AI RMF basics

The NIST AI RMF is a voluntary resource for managing AI risk. It is meant for organizations that design, develop, deploy or use AI systems. The National Artificial Intelligence Initiative Act of **2020** directed NIST to create it. NIST released version **1.0** with supplementary documents.

EXAM TRAP

The AI RMF is not a binding regulation. NIST built it as a voluntary resource, so no organization is legally required to adopt it.

Govern categories

Govern categories set accountability and engagement. Know the structure: core functions break into categories and then subcategories. Key AI governance actors include the Board of Directors, senior leadership and organizational management.

Two Govern categories of the NIST AI RMF compared

Feature	Govern 2	Govern 5
Focus	Accountability structures	Engagement with relevant AI actors
Goal	Teams empowered, responsible and trained to map, measure and manage AI risks	Robust processes for engaging those actors

Playbook practices

The Playbook describes how to score risk and split duties. A typical risk score combines impact and likelihood, by multiplying them or combining them qualitatively. Some industries run **three** or more lines of defense. Separate teams then own development, risk management and auditing.

EXAM TRAP

A red team is not just any internal test group. The Playbook describes it as external experts or personnel independent from internal AI actors.

Core ISO AI standards

The AIGP names three core ISO AI standards: 22989, 42001 and 42005. Learn them by number. ISO **27701** is a different standard. It extends ISO **27001** and **27002** to cover a privacy information management system.

EXAM TRAP

ISO **27701** is a privacy management standard, not one of the three core ISO AI standards.

Assurance versus compliance

Assurance goes further than compliance. Compliance means meeting set standards and regulations. Assurance is a broader, deeper evaluation that builds confidence in reliability and safety. Audits, red teaming, threat modeling and security testing are periodic activities to assess a system's performance, reliability and safety.

Key takeaways

- The OECD AI Principles, AI system definition and lifecycle all belong to the OECD Recommendation on Artificial Intelligence.
- The NIST AI RMF is voluntary, and its Playbook shares its core functions, categories and subcategories.
- The core ISO AI standards are **22989**, **42001** and **42005**; ISO **27701** covers privacy management.
- Assurance is broader than compliance: it evaluates and communicates how trustworthy an AI system is.

Sources: VERSION 2.1 Effective date: 2 February 2026 (opens in a new tab) · AI Principles Overview - OECD.AI (opens in a new tab) · NIST's Reva Schwartz on the new AI Risk Management Framework | IAPP (opens in a new tab) · Top 10 operational impacts of the EU AI Act – AI Assurance across the risk categories | IAPP (opens in a new tab) · ? App. A: Descriptions of AI Actor Tasks - AIRC (opens in a new tab) · Govern - AIRC (opens in a new tab) · ISO 27701 vs. NIST Privacy Framework: Choosing the Right One for You | IAPP (opens in a new tab)

Domain 3 | 21–25% of the exam | 12 min read

Understanding how to govern AI development

Overview

Every point below comes from IAPP's official documentation; each section lists its sources.

Govern the designing and building of the AI system.

Start here

THINK FIRST

What should an AI governance professional settle, assess and record before an AI system is designed and built?

Why it matters: Questions describe an AI project at the design stage and ask which governance step, assessment or risk tool comes next.

Governing design and build starts with the business context and use case. An impact assessment and design-stage risk management follow, with policies and ethics applied to each design choice. The whole process is documented to establish compliance and manage risks.

Key terms

Term	Meaning
Impact assessment	An assessment of the AI system that the governance professional performs or reviews during design and build.
Harms matrix	A tool that weighs design-stage harms by their probability and severity.
Fundamental rights impact assessment	An EU AI Act assessment that deployers perform for the first use of a high-risk AI system.
No-go decision	Choosing not to proceed with a system, considered when mapping its risk is prohibitively difficult.

Start with context and use case

Define the business context and use case first. Governing design and build starts there. Stakeholder requirements also need to be elicited and documented. Product requirement documents, user stories, UI/UX research, systems engineering and ethnography are ways to do it.

Impact assessments

Perform or review an impact assessment on the system. The governance professional either runs the assessment or reviews one that others ran. A data protection impact assessment can also reduce security and privacy risk.

EU fundamental rights assessment

EU deployers of high-risk systems assess fundamental rights. The duty applies to the first use of the high-risk system. The deployer then notifies the market surveillance authority of the results, using the filled-out template.

EXAM TRAP

The assessment is not a one-off file. If any element changes or goes out of date during use, the deployer must update the information.

Policies on design choices

Policies and ethics apply to every design choice. They cover the purpose of the AI, requirements gathering, architecture and model selection, and human oversight. They also cover data analysis, metric and threshold evaluation, stakeholder engagement and feedback, and operational controls.

Managing design-stage risk

Identify and manage internal and external risks. Tools include a probability/severity harms matrix, a risk mitigation hierarchy and stakeholder mapping. Use-case evaluation, benchmarking, and pre-deployment pilots and testing help too. Set risk criteria across sources and levels of risk, from negligible to critical. If mapping risk is prohibitively difficult, a no-go decision may be considered.

EXAM TRAP

Risk work does not stop at internal factors. External risks and contributing factors for the model and the system must be identified and managed too.

Document the process

Document how the system was designed and built. The record helps establish compliance and manage risks. In the EU AI Act's internal-control conformity assessment, the provider checks that design and development match the technical documentation.

Key takeaways

- Governing design and build starts with the system's business context and use case.
- Perform or review an impact assessment on the AI system during design and build.
- Manage design-stage risk with tools such as a harms matrix, a mitigation hierarchy and stakeholder mapping.
- Document the design and build process to establish compliance and manage risks.

Sources: VERSION 2.1 Effective date: 2 February 2026 (opens in a new tab) · Article 27: Fundamental rights impact assessment for high-risk AI systems | AI Act Service Desk (opens in a new tab) · Map - AIRC (opens in a new tab) · Conducting Data Privacy Impact Assessment on Your Cloud Environments | IAPP (opens in a new tab) · ANNEX VI | AI Act Service Desk (opens in a new tab)

Govern the collection and use of data in training and testing the AI model and system.

Start here

THINK FIRST

How would you show that an AI model's training and testing data was lawfully obtained, fit for purpose and properly tested?

Why it matters: Questions describe a training or testing step and ask which data governance, testing or record-keeping duty applies.

Governing training and testing data covers several duties. You confirm your right to use the data, check its quality and trace where it came from. You also plan the tests, manage risks that arise during training and testing, and document the process.

Key terms

Term	Meaning
Computational bias	A systematic error in predictions that comes from a model's assumptions or from the data itself.
Selection bias	Bias in selecting data for model training, one way societal bias can enter a model.
Accuracy	How well an AI system correctly carries out the task it is meant for.
Adversarial attack	A safety and security risk where a model is manipulated, such as with malicious or deceptive input data.
Model weights	The numerical parameters in a neural network that determine how data is processed.

Confirm the right to use data

Assess and document the lawful rights to collect and use data. Transparency, choice, lawful basis and purpose limitation requirements apply to AI. Intellectual property law can prohibit or limit using data for training. Sensitive categories such as biometrics carry their own requirements.

EXAM TRAP

Holding data is not the same as having the right to train on it. Intellectual property law can prohibit or limit that use.

Check that data is fit

Assess data for quality, quantity, integrity and fitness for purpose. In a sample life cycle, the data stage gathers data and checks its quality. It also documents and assesses the dataset's metadata and characteristics. Data minimization and privacy by design apply too.

EXAM TRAP

Poor data is not only a training problem. Lack of quality data is also a named cause of incidents from AI systems.

Trace lineage and provenance

Establish and document data lineage and provenance. Both are required: establishing them and documenting them. Data discovery, knowing where data lives inside and outside the organization, helps broader data governance. An automated data map can reveal hidden or unknown systems.

Plan the testing

Plan and perform several kinds of testing. The list covers unit, integration, validation, performance, security, bias and interpretability testing. In a sample life cycle, testing, validation and verification follow model training and come before calibration.

Bias in data and models

Bias can come from the data or from the model's assumptions. Computational bias is a systematic error from either origin. Selection bias, in choosing training data, is one way societal bias gets in. Bias testing is one of the listed test types.

Manage risks found in testing

Spot and manage the issues and risks that training and testing raise. Adversarial attacks, such as malicious or deceptive input data, are one safety and security risk to watch. Insufficient testing and model or data drift are also named causes of incidents from AI systems.

Testing in real-world conditions

Article 60 of the EU AI Act sets conditions for real-world testing. A provider plans the testing before it starts. Subjects may also request immediate and permanent deletion of their personal data.

Article 60 conditions for testing in real-world conditions

Topic	Rule	Applies to
Testing plan	Drawn up and submitted to the market surveillance authority where testing happens	Provider
Serious incident	Reported to the national authority; mitigate at once or suspend or end testing	Provider

Topic	Rule	Applies to
Withdrawal	Allowed at any time, without justification, by revoking informed consent	Test subjects
Outputs	Predictions, recommendations or decisions can be effectively reversed and disregarded	AI system
Data transfers	To third countries only with appropriate safeguards under Union law	Testing data

EXAM TRAP

A serious incident in real-world testing is not handled internally only. It must be reported to the national market surveillance authority.

Document training and testing

Document the process to validate results, establish compliance and manage risks. Record-keeping policies should describe and characterize the training data. They should also cover testing and validation results.

Key takeaways

- Assess and document lawful rights before collecting or using training data.
- Check data quality, quantity, integrity and fitness for purpose.
- Establish and document data lineage and provenance.
- Testing spans unit, integration, validation, performance, security, bias and interpretability.
- Document training and testing to validate results, establish compliance and manage risks.

Sources: VERSION 2.1 Effective date: 2 February 2026 (opens in a new tab) · Key terms for AI governance | IAPP (opens in a new tab) · Addressing emerging AI security requirements in new laws through AI model weights | IAPP (opens in a new tab) · AI Governance in Practice Report 2024 | IAPP (opens in a new tab) · Where's the Data? How to Find Your Biggest Privacy, Security & Governance Risks | IAPP (opens in a new tab) · The Automated Data Map: Your Foundation for Privacy, Security and Governance | IAPP (opens in a new tab) · Article 60: Testing of high-risk AI systems in real world conditions outside AI regulatory sandboxes | AI Act Service De (opens in a new tab) · Govern - AIRC (opens in a new tab)

Govern the release, monitoring and maintenance of the AI system.

Start here

THINK FIRST

How would you take a high-risk AI system from release readiness through monitoring, incidents and public disclosure?

Why it matters: Questions describe an AI system near or after release and ask which readiness, monitoring, incident or disclosure step governance requires.

Release is a gate, not an end point. A system must be shown safe and pass conformity assessment before release. Then it is monitored, tested, maintained and disclosed for as long as it runs.

Key terms

Term	Meaning
Conformity assessment	The check a high-risk AI system must pass before it is placed on the market or put into service.
Notified body	The body that assesses a provider's quality management system and technical documentation under Annex VII.
Post-market monitoring plan	A plan disclosed publicly to meet transparency obligations, alongside technical documentation and instructions for use.
Red teaming	Testing a system under adversarial or stress conditions to measure its response and find failure modes.
Brittleness	Failing in poorly described ways and transferring poorly outside the training environment.

Show it is safe before release

Deploy only a system shown to be safe within risk tolerance. Its residual negative risk must not exceed the risk tolerance. It must also fail safely, especially beyond its knowledge limits. Release preparation includes creating the model card and satisfying conformity requirements.

Conformity before the market

High-risk systems pass Article 43 conformity assessment first. Under Article **16**, this happens before the system is placed on the market or put into service. Providers also draw up an EU declaration of conformity under Article **47**.

How Annex VII works

A notified body assesses quality management and technical documentation. If the system fails the training data requirement, it is re-trained before a new conformity assessment. Changes to a certified system go to the notified body. It decides between a new assessment and a certificate supplement.

Annex VII roles of the notified body and the provider

Area	Notified body	Provider
Quality management system	Assesses it against Article **17**	Reports any intended change to it
Technical documentation	Examines it, with access to data sets where relevant	Grants that access, including remote access where appropriate
Testing	Runs its own tests if not satisfied with the provider's	Carries out the tests reviewed first
Ongoing oversight	Runs periodic audits and gives an audit report	Maintains and applies the quality management system
Outcome	Issues a Union technical documentation assessment certificate	Meets the Chapter III, Section **2** requirements

EXAM TRAP

Access to the trained models is not a first step. The notified body gets it only after other reasonable means prove insufficient, and upon a reasoned request.

Monitor and maintain

Monitor continuously and keep a regular maintenance schedule. The schedule covers maintenance, updates and retraining. Monitoring matters because current systems are brittle and transfer poorly outside their training environment. Watch for anomalies with control limits, confidence intervals, integrity constraints or ML algorithms.

Periodic assessment

Audits, red teaming, threat modeling and security testing recur. They assess performance, reliability and safety over time. Red-team exercises test the system under adversarial or stress conditions. They also show whether it returns to normal function after an unexpected adverse event.

EXAM TRAP

Assessment does not stop at release. Audits, red teaming, threat modeling and security testing are periodic activities.

Incidents and their causes

Manage and document incidents, issues and risks. Causes include brittleness, lack of robustness, lack of quality data, insufficient testing, and model or data drift. Work out why with cross-functional stakeholders.

EXAM TRAP

Incident response is not improvised. Plans should be documented, practiced and measured, including response and down times.

Public disclosures

Transparency obligations call for public disclosures. These include technical documentation, instructions for use to deployers and post-market monitoring plans. Article **50** information reaches people at the latest at first interaction or exposure.

Article 50 transparency duties compared

Situation	Who acts	Duty
AI that interacts directly with people	Providers	Design it so people know it is AI, unless that is obvious
Synthetic audio, image, video or text	Providers	Mark outputs in a machine-readable format as generated
Deep fake image, audio or video	Deployers	Disclose that the content is generated or manipulated

EXAM TRAP

Article **50** duties do not all fall on providers. Deployers of deep fake content must disclose that it was generated or manipulated.

Shutdown and outreach

Keep a policy and controls to deactivate or localize a system. Triggers include regulatory requirements or performance issues. Establish external communication plans too. Providers of high-risk systems take the necessary corrective actions and provide information under Article **20**.

Key takeaways

- A high-risk system passes Article **43** conformity assessment before market or service.
- Under Annex VII, a notified body assesses the quality management system and technical documentation.
- Monitor continuously and schedule maintenance, updates and retraining.
- Audits, red teaming, threat modeling and security testing are periodic activities.
- Deployers of deep fakes must disclose that the content is generated or manipulated.

Sources: Measure - AIRC (opens in a new tab) · Article 16: Obligations of providers of high-risk AI systems | AI Act Service Desk (opens in a new tab) · VERSION 2.1 Effective date: 2 February 2026 (opens in a new tab) · ANNEX VII | AI Act Service Desk (opens in a new tab) · Article 50: Transparency obligations for providers and deployers of certain AI systems | AI Act Service Desk (opens in a new tab)

Domain 4 | 21–25% of the exam | 11 min read

Understanding how to govern AI deployment and use

Overview

Every point below comes from IAPP's official documentation; each section lists its sources.

Evaluate key factors and risks relevant to the decision to deploy the AI system.

Start here

THINK FIRST

Faced with a proposed AI use case, how would you decide whether and how to deploy it?

Why it matters: Questions describe a use case and ask which context factor, model type or deployment option matters for the decision.

A deployment decision weighs three groups of factors: the context of the use case, the type of model and the deployment option. Deployment options cover where the model runs and whether to adapt it.

Key terms

Term	Meaning
Use case context	The factors behind a use: business objectives, data availability, performance needs, workforce readiness and ethics.
Retrieval augmented generation	A technique used with an AI model to improve its performance and fit.
Agentic architecture	A deployment option, like fine-tuning or RAG, used to improve a model's performance and fit.
Edge	One of three contrasting places to run a model, alongside cloud and on-premise.

Start with the use case

Use case context covers five factors. Check business objectives, performance requirements and data availability. Weigh ethical considerations and workforce readiness as well. Each one is a factor to evaluate before deciding to deploy.

Compare model types

Model types differ along four contrasts. Learn the four pairs as separate contrasts. Generative AI and agentic AI face fast-changing governance requirements.

Contrasting AI model types

Contrast	One type	The other type
Classic vs. generative	Classic	Generative
Proprietary vs. open source	Proprietary	Open source
Size	Small	Large
Capability	Language	Multimodal

Where the model runs

Host the model in the cloud, on premises or at the edge. These three hosting options contrast with one another. The choice among them is one of the AI deployment options to understand.

Use as is or adapt

You can deploy a model as is or adapt it. Fine-tuning, retrieval augmented generation and agentic architectures are ways to adapt it. Each aims to improve the model's performance and fit. Other techniques can serve the same goal.

EXAM TRAP

Adaptation is not required. Using the model as is counts as a deployment option in its own right.

Owning the model

Deploying your own proprietary model adds obligations. It brings unique risks and opportunities, including higher potential liability. The AIGP uses developer and deployer to describe tasks, since many organizations fill both roles.

EXAM TRAP

Picking a third-party model does not take you outside deployment governance. Domain IV applies whether the model is your own or comes from a third party.

Generative AI at work

Workplace generative AI brings privacy and security risk. One IAPP webinar reports that **78%** of knowledge workers bring their own generative AI tools. Useful responses include robust AI policies, governance frameworks, AI privacy impact assessments and stronger data privacy controls.

Deployer duties for high-risk systems

Deployers of high-risk systems have duties under the EU AI Act. The EU AI Act sets obligations by the risk category of the use. The table shows four duties of deployers of high-risk systems.

Example EU AI Act duties for deployers of high-risk systems

Duty	When it applies	What to do
Inform workers	Before workplace use, as an employer	Tell workers' representatives and the affected workers
Check input data	The deployer controls the input data	Keep it relevant and sufficiently representative
Keep logs	Logs are under the deployer's control	Keep them at least **six months**, unless other law says otherwise
Report a risk	The system may present a risk	Inform the provider or distributor and the authority, and suspend use

Risk-based regulation

The EU AI Act regulates uses, not technologies. Its risk-based approach aims to curb AI risk while still promoting innovation. The GDPR likewise scales safeguards to the risk of processing. So judge the use case, not only the model.

Key takeaways

- Use case context covers business objectives, performance needs, data availability, ethics and workforce readiness.
- Model types contrast classic and generative, proprietary and open source, small and large, language and multimodal.
- Host in the cloud, on premises or at the edge, and use the model as is or adapt it.
- Deploying your own proprietary model brings increased obligations and higher potential liability.

Sources: VERSION 2.1 Effective date: 2 February 2026 (opens in a new tab) · Generative AI in the workplace: Practical privacy strategies | IAPP (opens in a new tab) · Top 10 operational impacts of the EU AI Act – Understanding and assessing risk | IAPP (opens in a new tab) · Article 26: Obligations of deployers of high-risk AI systems | AI Act Service Desk (opens in a new tab)

Perform key activities to assess the AI system.

Start here

THINK FIRST

Before deploying an AI system, how would you decide which assessments and contract checks it needs?

Why it matters: Questions describe an AI system about to go live and ask which assessment, contract term or deployment risk to address.

Assessing an AI system before deployment has three parts. You perform or review an impact assessment, and you check the vendor or licensing agreement. You also weigh the extra duties of deploying your own model.

Key terms

Term	Meaning
FRIA	Fundamental rights impact assessment: an EU AI Act Article **27** duty for some deployers of high-risk systems.
Annex III	The EU AI Act list of high-risk uses; a listed system is high-risk unless it poses no significant risk of harm.
Market surveillance authority	The authority a deployer notifies of its FRIA results, with the filled-out template.
Developer and deployer	Labels the AIGP uses for distinct tasks, not fixed identities; many organizations fill both roles.

Impact assessment timing

Impact assessments come at design and again before deployment. You perform or review one while governing the design and build. You return to it for the selected system before it goes live. A data protection impact assessment can also reduce security and privacy risk.

Who owes a FRIA

Article 27 targets certain deployers of high-risk systems. Bodies governed by public law and private entities providing public services must assess the impact before deployment. So must deployers of systems in Annex III points **5**(b) and (c). Systems in the area of Annex III point **2** are excepted.

EXAM TRAP

A FRIA is not redone before every use. The duty applies to the first use, and the deployer updates it when an element changes or goes out of date.

What a FRIA records

A FRIA records usage, affected people and responses. It states how long and how often the system will be used. It names the categories of people and groups likely to be affected. It sets out what happens if risks materialize, including internal governance and complaint mechanisms.

Notifying and reusing work

Deployers notify the market surveillance authority of the results. They submit the filled-out template with the notification. The AI Office must provide a questionnaire template, possibly as an automated tool, to simplify compliance. In similar cases, a deployer may rely on earlier assessments.

Ways the EU AI Act eases the FRIA burden

Option	When it applies	Effect
Reuse earlier assessments	Similar cases	Rely on previous FRIAs or the provider's impact assessments
Cross-reference a DPIA	Duty already met by an Article **35** DPIA under Regulation (EU) **2016/679**	Include or cross-reference its relevant parts
May skip notification	The case in Article **46(1)**	May be exempt from notifying the authority

Is it really high-risk?

An Annex III system can fall outside the high-risk category. That happens when it poses no significant risk of harm to health, safety or fundamental rights. Narrow procedural tasks and preparatory tasks are examples. The provider documents that view before market placement and shares it on request. Registration under Article **49(2)** still applies.

EXAM TRAP

A derogation never rescues profiling. An Annex III system that profiles natural persons is always high-risk, even when a derogation condition applies.

Vendor and licensing terms

Contracts are a main control for third-party AI risk. Policies, assessments and contracts manage that risk across procurement, supply chain, human resources and acceptable use. In California, the California Privacy Rights Act (CPRA) mandates due diligence of processing operations. It also places direct enforceable obligations on service providers and contractors.

Deploying your own model

A proprietary model brings unique risks and opportunities. The risks include increased obligations and higher potential liability. Map the tasks you perform, since one organization can be both developer and deployer. Weigh the costs of internal and external failures when making go or no-go decisions.

EXAM TRAP

Building your own model does not shrink your governance duties. It adds obligations and raises potential liability compared with buying one.

Key takeaways

- Perform or review an impact assessment at design and again before deployment.
- An Article **27** FRIA covers first use, names affected groups and goes to the market surveillance authority.
- Profiling of natural persons always keeps an Annex III system high-risk.
- Deploying a proprietary model means increased obligations and higher potential liability.

Sources: VERSION 2.1 Effective date: 2 February 2026 (opens in a new tab) · Article 27: Fundamental rights impact assessment for high-risk AI systems | AI Act Service Desk (opens in a new tab) · Article 6: Classification rules for high-risk AI systems | AI Act Service Desk (opens in a new tab) · Conducting Data Privacy Impact Assessment on Your Cloud Environments | IAPP (opens in a new tab) · Summary of CPRA Contractual Obligations | IAPP (opens in a new tab) · Map - AIRC (opens in a new tab)

Govern the deployment and use of the AI system.

Start here

THINK FIRST

Once an AI system is live, what keeps it governed until the day it is switched off?

Why it matters: Questions describe a deployed AI system and ask what to monitor, test, document or disclose, or what to do when it must stop.

Deployment is where governance becomes routine work. You apply policies, monitor the system and test it on a schedule. You document incidents, plan outside communication and keep controls ready to deactivate or localize it.

Key terms

Term	Meaning
Drift	A shift over time in a deployed system's performance and trustworthiness that degrades its value.
Continuous monitoring	Watching for performance changes and surprises, live or at a set interval, throughout the lifecycle.
Red-teaming	Adversarial stress testing to find failure modes, done by outside experts or staff independent of internal AI actors.
Residual risk	The sum of all unmitigated risks; negative residual risks to downstream users should be documented.

Apply deployment policies

Deployment policies cover data, risk, issues and training. User training, issue management, risk management and data governance all belong in them. Tailor training to each group, such as developers versus legal, compliance and audit. Policies should also set how often and in how much detail you monitor, audit and review.

Monitor continuously

Monitor continuously and maintain on a regular schedule. Plan maintenance, updates and retraining ahead of time. Set alerts for when production inputs or predictions differ from pre-deployment test results. Log inputs and configuration whenever someone tries to use the system outside its valid range.

EXAM TRAP

Monitoring is not a launch-day check. Drift shifts performance and trustworthiness over time, so tracking continues across the lifecycle.

Test periodically

Audits, red teaming, threat modeling and security testing recur. Together they assess performance, reliability and safety. The Playbook advises red teaming at a prescribed cadence, with checks on how well it works. Watermarking can deter data and model extraction attacks.

Document incidents and risks

Write down incidents, issues and risks, plus your post-market monitoring plans. Incidents can come from brittleness, poor data quality, too little testing, or model or data drift. Keep a database of errors, near-misses, incidents and negative impacts. Record dates, report counts, severity and responses.

Downstream and unintended use

Forecast and reduce secondary uses and downstream harms. Tell downstream AI actors what safe operation requires, plus known limits and suggested warning labels. Record negative residual risks to downstream acquirers and end users. You can respond to a risk by mitigating, transferring, avoiding or accepting it.

EXAM TRAP

Residual risk is not only the largest risk left over. The Playbook defines it as the sum of all unmitigated risks.

External communication

Plan how you communicate outside the organization. Set policies on disclosing AI use and risk material such as impact assessments, audits and test results. You may share incident details with users, impacted communities, the AIAAIC, the AI Incident Database or the MITRE CVE list. Post-market monitoring plans can be public disclosures too.

Deactivate or localize

Have a policy, and controls, ready to deactivate or localize a system. Triggers include regulatory requirements, performance issues, risks above tolerance or the end of a system's lifetime. Build decision thresholds for bypass or deactivation into continual monitoring. Review each such incident for root causes and impacts.

EXAM TRAP

Switching a system off does not end governance. Irregular deletion can raise risk, so preserve materials for forensic, regulatory and legal review.

Decommissioning

Decommissioning follows a policy, not a delete key. Such policies typically address data retention, possible future investigations and migration to the replacement system. Retired models may stay in the inventory beside active ones for a set period. Consider how retirement affects individuals, groups and communities downstream.

Key takeaways

- Monitor continuously, and schedule maintenance, updates and retraining in advance.
- Audits, red teaming, threat modeling and security testing recur to check performance, reliability and safety.
- Keep records of incidents, issues and risks, and of your post-market monitoring plans.
- Deactivate or localize when rules or performance demand it, and preserve materials for review.

Sources: VERSION 2.1 Effective date: 2 February 2026 (opens in a new tab) · Manage - AIRC (opens in a new tab) · Govern - AIRC (opens in a new tab) · Measure - AIRC (opens in a new tab)

MockCert

Take an IAPP Artificial Intelligence Governance Professional (AIGP) mock exam the way the real one feels

5 timed mock exams · 100 questions in 165 minutes each · one is free

- Free to start: study notes, explained quizzes, flashcards and the first mock exam
- Timed mock exams built to the official exam guide
- A verifiable Certificate of Assessment when you pass a premium mock

Explore IAPP Artificial Intelligence Governance Professional (AIG...

Open it on your phone

mockcert.com/iapp-aigp/

Scan the code with your phone's camera.