

AWS AI Practitioner Cheat Sheet 2026

The night-before summary, built like the exam.

65 Scored questions	1h 30m Time limit	700/1000 Passing score
\$100 Exam fee	Amazon Web Services (A... Governing body	

A dense, night-before reference for the Foundational-level AWS Certified AI Practitioner exam. Skim the table, memorize the domain weights, then run the checklist.

Exam Logistics At-a-Glance

Item	Value
Level	Foundational
Format	65 questions (50 scored, 15 unscored, unlabeled)
Duration	90 minutes
Cost	100 USD
Passing score	700 on a scaled 100–1,000 range
Scoring model	Compensatory — no per-section minimum
Delivery	Pearson VUE testing center or OnVUE online proctored
Results	Within 5 business days
Validity	3 years, then recertify
Reschedule/cancel	Up to 24 hours before, no fee

- Scheduling happens through Pearson VUE from your AWS Certification Account.
- Testing centers are independently operated facilities, not run directly by AWS.
- The 15 unscored questions are mixed in and not identifiable during the exam, so treat every question as if it counts.

Content Domains (Memorize the Weights)

1. Domain 1 — Fundamentals of AI and ML: 20%
2. Domain 2 — Fundamentals of Generative AI: 24%
3. Domain 3 — Applications of Foundation Models: 28% (largest domain — prioritize study time here)
4. Domain 4 — Guidelines for Responsible AI: 14%
5. Domain 5 — Security, Compliance, and Governance for AI Solutions: 14%

Five domains total. Domains 3 and 2 together are over half the scored exam — spend your final review hours on foundation model applications and generative AI fundamentals before circling back to the smaller domains.

Key Terms and Concepts to Memorize

- Foundation model (FM): a large pretrained model adaptable to many downstream tasks via prompting, fine-tuning, or

retrieval augmentation.

- Generative AI vs. traditional ML: traditional ML predicts or classifies from labeled data; generative AI produces new content (text, image, code) from learned patterns.
- Prompt engineering: structuring inputs (instructions, context, examples) to steer model output without retraining.
- Fine-tuning vs. RAG: fine-tuning adjusts model weights on a custom dataset; retrieval-augmented generation injects external knowledge at inference time without retraining.
- Embeddings and vector databases: numeric representations of meaning used for semantic search and RAG retrieval.
- Bias, fairness, and explainability: core pillars of responsible AI — expect scenario questions about detecting and mitigating biased outputs.
- Hallucination: a model generating plausible-sounding but factually incorrect content — a recurring exam trap topic.
- Shared responsibility model: AWS secures the cloud infrastructure; customers secure their data, access configuration, and model usage.
- Guardrails and content filtering: mechanisms to constrain model outputs for safety and compliance.
- Model evaluation metrics: know the difference between metrics used for classification/regression tasks versus qualitative evaluation of generative outputs.

Common Gotchas and Traps

- Don't assume every question is scored — but don't try to guess which ones aren't, since unscored items are indistinguishable during the test.
- The compensatory scoring model means you can't "fail" a single domain outright, so don't panic over a rough patch of questions — keep moving and let the overall score average out.
- Watch for questions that test conceptual understanding rather than memorized service names — this is a Foundational exam, so expect "why would you use X" framing over deep configuration detail.
- Security and governance (Domain 5) questions often overlap with responsible AI (Domain 4) — read carefully to determine whether the question is asking about technical controls or ethical/policy guidelines.
- Time management: 90 minutes for 65 questions is roughly 83 seconds per question — flag long scenario questions and return to them rather than stalling.
- Rescheduling within 24 hours of your appointment may incur fees, so lock in logistics early.

Night-Before Checklist

- Confirm delivery method: testing center address and ID requirements, or OnVUE system check (webcam, room scan, private space) if testing online.
- Review the domain weight list once more, prioritizing Domain 3 (28%) and Domain 2 (24%).
- Skim key terms above — hallucination, RAG vs. fine-tuning, shared responsibility model, guardrails.
- Verify your AWS Certification Account login and exam confirmation details.
- Get sleep — a Foundational exam rewards clear thinking on scenario-based questions more than raw memorization.
- Remember the passing bar: 700 out of a 100–1,000 scaled score, compensatory across all domains.
- Plan arrival or online check-in time with buffer — late arrival can forfeit your slot.

Official sources

- Amazon Web Services (AWS) — AWS Certified AI Practitioner (AIF-C01) Exam Guide: <https://docs.aws.amazon.com/aws-certification/latest/ai-practitioner-01/ai-practitioner-01.html>
- Amazon Web Services (AWS) — AWS Certified AI Practitioner — Certification Overview: <https://aws.amazon.com/certification/certified-ai-practitioner/>
- Amazon Web Services (AWS) — AWS Training and Certification FAQs: <https://aws.amazon.com/certification/faqs/>

Free practice questions for this exam: www.everyexamprep.com/practice/exams/technology/aws-ai-practitioner