

EXAM TIP

Exam details change from time to time. Before you book, confirm the current format, price and objectives on the official AWS Certification page for AIF-C01. Everything in this guide follows the five published content domains and their weightings.

The five domains and their weights

Domain 1: Fundamentals of AI and ML



Domain 2: Fundamentals of GenAI



Domain 3: Application of Foundation Models



Domain 4: Guidelines for Responsible AI



Domain 5: Security, Compliance & Governance



SHARE OF SCORED CONTENT BY DOMAIN

Domain 3 is the heaviest at 28 percent, and Domains 2 and 3 together make up more than half of the exam (52 percent). Generative AI is the centre of gravity. The other three domains still matter: Domain 1 gives you the vocabulary that every other question assumes, and Domains 4 and 5 hold many of the "easy marks" because the concepts are clear once you have seen them.

Domain	Weight	What it tests	Chapter
1. Fundamentals of AI and ML	20%	Core terms, ML types, ML lifecycle, metrics, AWS AI/ML services	1

Domain	Weight	What it tests	Chapter
2. Fundamentals of GenAI	24%	Foundation models, tokens, embeddings, prompts, use cases, AWS GenAI tools	2
3. Applications of Foundation Models	28%	Model selection, RAG, prompt engineering, fine-tuning, agents, evaluation	3
4. Guidelines for Responsible AI	14%	Fairness, bias, transparency, explainability, safety, human oversight	4
5. Security, Compliance and Governance	14%	Shared responsibility, IAM, encryption, privacy, regulations, governance	5

How this book is organised

Each of the five domains has its own chapter. Every chapter opens with the exam objectives, then explains the ideas in plain English with diagrams, flowcharts and comparison tables, and closes with a "remember this" recap. Look for two kinds of highlighted box as you read:

EXAM TIP


Exam tips flag the way a concept usually appears in a question, and the keyword that points to the right answer.

COMMON TRAP

Common traps flag pairs of ideas that candidates mix up, or distractors that look correct but are not the best answer.

After the five chapters you will find a one-page-per-topic cheat-sheet section, a glossary, then **300 original practice questions** with full explanations. No question appears twice anywhere in the book: the domain drills and the full mock exam each use their own questions.

A four-week study plan



FOUR-WEEK PLAN FOR A FOUNDATIONAL EXAM

If you have more time, stretch each week to two. If you already work with AWS, you can compress Week 1. Whatever your pace, finish with at least two timed practice sessions, and only book the real exam when you are consistently scoring around 80 percent on questions you have not seen before.

How to use the practice questions

1. Study a chapter, then answer that domain's drill questions without looking at the key.
2. Mark every question you got wrong or guessed. Read the explanation for **all** options, not only the correct one.
3. Return to the chapter section that covers each miss and re-read it.
4. Save the mock exam for the end. Set a timer for 90 minutes and treat it as the real thing.

Reading exam questions like an expert

AIF-C01 questions are short scenarios. Several answers may be technically possible, but only one is the **best** fit for the stated requirement. A reliable routine:

1. Read the last sentence first. It tells you what is being asked: cheapest, least operational overhead, most secure, most explainable, fastest to implement.
2. Underline the constraint words: *minimize, without writing code, fully managed, private, auditable, lowest latency*.
3. Eliminate answers that solve a different problem, for example a training service when the question is about inference.
4. Prefer the managed AWS AI service when the scenario needs a common capability such as transcribing speech or extracting text from forms. Custom model building on SageMaker is rarely the "simplest" answer for a foundational exam.
5. When two answers remain, choose the one that matches the constraint word exactly.

COMMON TRAP

Do not over-think. Foundational questions rarely hinge on obscure limits or prices. If an answer requires a lot of extra assumptions to work, it is probably a distractor.

Exam-day checklist

- ☐ I know my test slot and have the required identification ready
- ☐ I have practised at least one full 90-minute timed session
- ☐ I can explain in one sentence the difference between RAG, fine-tuning and prompt engineering


- ☐ I can name the AWS service for each common AI task without hesitating
- ☐ I will flag hard questions, keep moving, and return with the time left
- ☐ I will answer every question: there is no penalty for a wrong answer



MAP THE NEED TO THE MANAGED AWS AI SERVICE

The ML development lifecycle

Building an ML solution is a loop, not a straight line. Results from evaluation and monitoring send you back to earlier stages.



ITERATE UNTIL THE MODEL MEETS THE BUSINESS GOAL

THE END-TO-END ML PIPELINE

Stage	What happens	Useful AWS pieces
1. Business goal	Frame the problem, define success metrics and check ML is appropriate	Stakeholders, KPIs
2. Data collection	Gather data from databases, streams, files	Amazon S3, AWS Glue, Kinesis
3. Data preparation and EDA	Clean, explore, visualise, handle missing values, split	SageMaker Data Wrangler, Glue DataBrew
4. Feature engineering	Create and select the most useful input features	SageMaker Feature Store
5. Model training and tuning	Train candidates and tune hyperparameters	SageMaker training jobs, automatic model tuning
6. Evaluation	Measure performance against agreed metrics on unseen data	SageMaker Experiments, Clarify
7. Deployment	Serve the model for real-time, batch or asynchronous inference	SageMaker endpoints, Model Registry