FAST-12 Projects: AI Practitioner Certification Prep (AIF-C01)

FAST-12 AI Practitioner Certification Prep · 7formations · Instructor: Wadood Chaudhary · All examples fictional · Practice content, not real exam items

This is a certification-prep course, so the projects build exam readiness and a work artifact you can reuse. Use fictional data or a sanitized description of a use case at your job (no customer personal information, PHI, or confidential details; your manager approves the use case). Scoring: Not yet / Meets / Strong. Completion needs Meets on every capstone criterion and 80% or higher on two timed practice exams (one resubmission allowed).

Capstone: AWS Generative-AI Use-Case Brief

What you produce: a 2-page brief recommending how a fictional or sanitized organization should build one generative-AI use case on AWS, plus a 5-minute walkthrough in session 5. The brief exercises every exam domain: fundamentals (Domain 1), generative-AI trade-offs (Domain 2), foundation-model application design (Domain 3), responsible AI (Domain 4), and security and governance (Domain 5).

Required contents

Section	Content	Domain
Problem and fit	Business problem; why AI (or why not); success metric	1, 2
Approach	Prompting, RAG (knowledge base), fine-tuning, or agent; why; what you rejected and why	3
Model choice	Two candidate models in Bedrock; comparison on quality (your evaluation set), latency, and cost per month (show the math)	2, 3
Evaluation	Evaluation set of 10+ questions with expected answers; automatic or LLM-as-a-judge plus human review; pass threshold	3
Responsible AI	Top three risks, the responsible-AI dimension of each, and the control (guardrail, human review, disclosure)	4
Security and governance	Shared responsibility split; IAM roles; encryption (KMS); logging (CloudTrail, invocation logs); data classification; compliance needs (for example, HIPAA BAA via AWS Artifact, G.L. c. 93H / 201 CMR 17.00 for MA personal information)	5
Rollout	Pilot scope, owner, review date, rollback trigger	3, 5

Rubric

Criterion	Not yet	Meets	Strong
Fit and approach	No reason given for the approach, or AI used where a rule would do	Approach justified; one alternative rejected with a reason	Two alternatives compared on cost, effort, and quality
Model comparison	No numbers, or cost math wrong	Quality and monthly cost compared with correct math	Latency and context-window limits included; sensitivity to volume shown

Criterion	Not yet	Meets	Strong
Evaluation	No evaluation set	10+ questions with expected answers and a pass threshold	Includes "not in source" and injection tests; human review defined
Responsible AI	Risks generic or controls missing	Three specific risks, each mapped to a dimension and control	Residual risk and monitoring plan stated
Security and governance	Shared responsibility wrong, or no IAM, encryption, or logging	All five covered correctly for the use case	Least-privilege policy sketched; retention and audit cadence set
Walkthrough	Over time or unclear	5 minutes; recommendation first	Answers two questions using exam vocabulary precisely

Example (excerpt): Charles River Family Dental (fictional) booking assistant

Problem and fit. Front desk spends about 9 hours a week answering hours, location, and insurance-network questions. Success: 60% of those questions answered by the assistant with 95% correct answers on the evaluation set. Booking itself stays in the practice system (a rule-based flow, no AI needed).

Approach. RAG over the office FAQ and accepted-insurance list using a Bedrock knowledge base. Rejected fine-tuning: the facts change monthly and RAG updates with a re-sync; fine-tuning would need retraining and gives no citations.

Model choice. Model A ($0.80/$3.20 per million tokens, sandbox rates) vs Model B ($3/$15). At 200 questions a day, 1,500 input and 150 output tokens each, 30 days: A ≈ $10.08 a month; B ≈ $40.50 a month. Both passed 19 of 20 evaluation questions; choose A.

Responsible AI. (1) Clinical advice to patients (safety): Bedrock guardrail denied topic plus a blocked message pointing to the office phone and 911. (2) Patient details repeated in chat (privacy): sensitive-information filter masks phone and email; the assistant never sees appointment records. (3) Spanish-speaking patients get worse answers (fairness): evaluation set includes 5 Spanish questions; pass threshold applies to each language.

Security and governance. AWS secures the infrastructure and Bedrock service; the practice owns IAM, guardrail configuration, data, and logging. BAA accepted in AWS Artifact before any PHI is in scope; the assistant is designed so no PHI is needed. Encryption with a KMS key; CloudTrail and invocation logging on, 1-year retention; quarterly review by the office manager.

Check the example's math: 200 × 1,500 = 300,000 input tokens a day; × 30 = 9 M; at $0.80 = $7.20. Output 200 × 150 × 30 = 0.9 M; at $3.20 = $2.88. Total $10.08. Model B: 9 × $3 = $27.00 + 0.9 × $15 = $13.50; total $40.50.

Mini-project 1: Service Map Flashcards (sessions 1-4)

What you produce: a one-page service map of the in-scope services you are most likely to confuse, in your own words, plus 20 self-written flashcards (use case → service).

Requirements: at least 15 services across AI services (Comprehend, Lex, Polly, Rekognition, Textract, Transcribe, Translate, Personalize), generative-AI services (Bedrock, Bedrock AgentCore, Nova, SageMaker AI, JumpStart, Amazon Quick), and security and governance services (IAM, KMS, CloudTrail, Config, Macie, Artifact, Inspector). Each card has one sentence on **when to use** it and one on **a look-alike** and how to tell them apart.

Criterion	Not yet	Meets	Strong
Coverage	Under 15 services, or out-of-scope services	15+ in-scope services across all three groups	All AI services and all listed security services
Accuracy	Any wrong use case	All use cases correct	Look-alike distinctions tested with a partner
Own words	Copied descriptions	Own wording, one sentence each	Each card tied to a lab or case

Example card: *Front:* "Pull totals and line items from scanned supplier invoices." *Back:* Amazon Textract. Look-alike: Amazon Comprehend analyzes text you already have (sentiment, entities); Textract gets the text and structure out of documents first.

Mini-project 2: Teach-Back on One Task Statement (session 5)

What you produce: a 3-minute explanation of the task statement you scored lowest on, with one original practice question and its explanation.

Criterion	Not yet	Meets	Strong
Accuracy	Factual error	Correct and matches the exam guide wording	Cites the AWS documentation page used
Practice question	Ambiguous or answer cued by length	One clear best answer; plausible distractors; explanation of each wrong option	Scenario-based, fictional, and peer-tested
Delivery	Over 3 minutes	3 minutes; one example	Uses a lab result as the example

Example (excerpt): "Task 3.1: choosing between prompting, RAG, and fine-tuning. In Lab 3 our knowledge base missed hearing protection because the PPE table was split across chunks. That's a retrieval problem, so the fix is chunking, not fine-tuning…"