AWS Certified AI Business Strategist Cheat Sheet

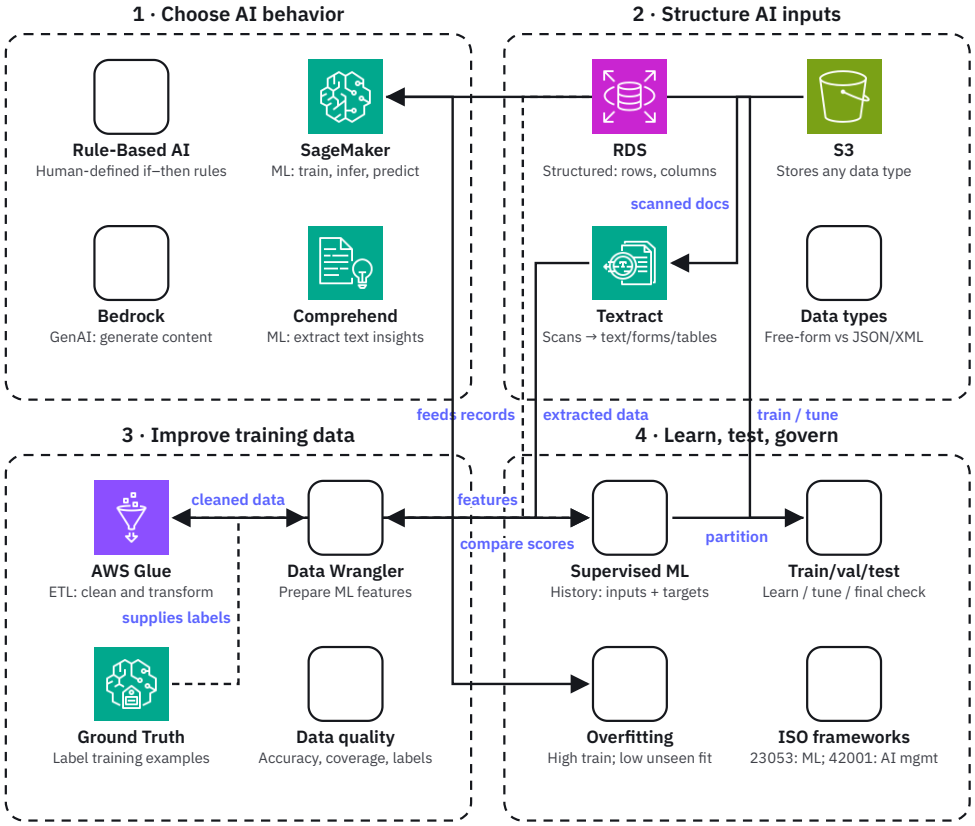
AIB-C01 (beta) · Domain 1 of 4

BlowStack

DOMAIN 1 · AI FUNDAMENTALS AND LITERACY · 24% OF THE EXAM · TASK STATEMENT 1.1

Describe core AI concepts and define terminology

Choose, prepare, train, govern



1 Choose AI behavior

An algorithm is a procedure; training fits model parameters; **inference** uses learned patterns. AI includes rules and ML; GenAI is ML that creates content. Bedrock provides managed foundation-model access.

2 Structure AI inputs

Tables are **structured**; text/images/audio are unstructured; JSON/XML are semi-structured. S3 stores all types, not only unstructured data. Textract extracts document content; extraction is not validation.

3 Improve training data

Glue handles ETL; Data Wrangler prepares ML features; Ground Truth supports labeling. **Representative data** covers deployment populations and rare cases; inaccurate or biased labels teach wrong outcomes.

4 Learn, test, govern

Supervised learning uses labeled history for classification or regression. Keep **test data untouched** until final evaluation. ISO/IEC 23053 defines an ML-system framework; ISO/IEC 42001 defines an AI management system.

Exam traps

- **Training ≠ inference**: training changes model parameters; inference uses a model to predict or generate. GenAI is a subset of ML, not a synonym for all AI.
- **High training accuracy** does not prove generalization. Tuning on test data leaks evaluation information; validation guides tuning, and a held-out test estimates unseen performance.
- **ISO/IEC 42001** concerns organizational AI management, not guaranteed model accuracy. ISO/IEC 23053 supplies a shared framework and vocabulary for ML-based AI systems.

Get practice exams and lessons for AWS Certified AI Business Strategist

AWS Certified AI Business Strategist Cheat Sheet

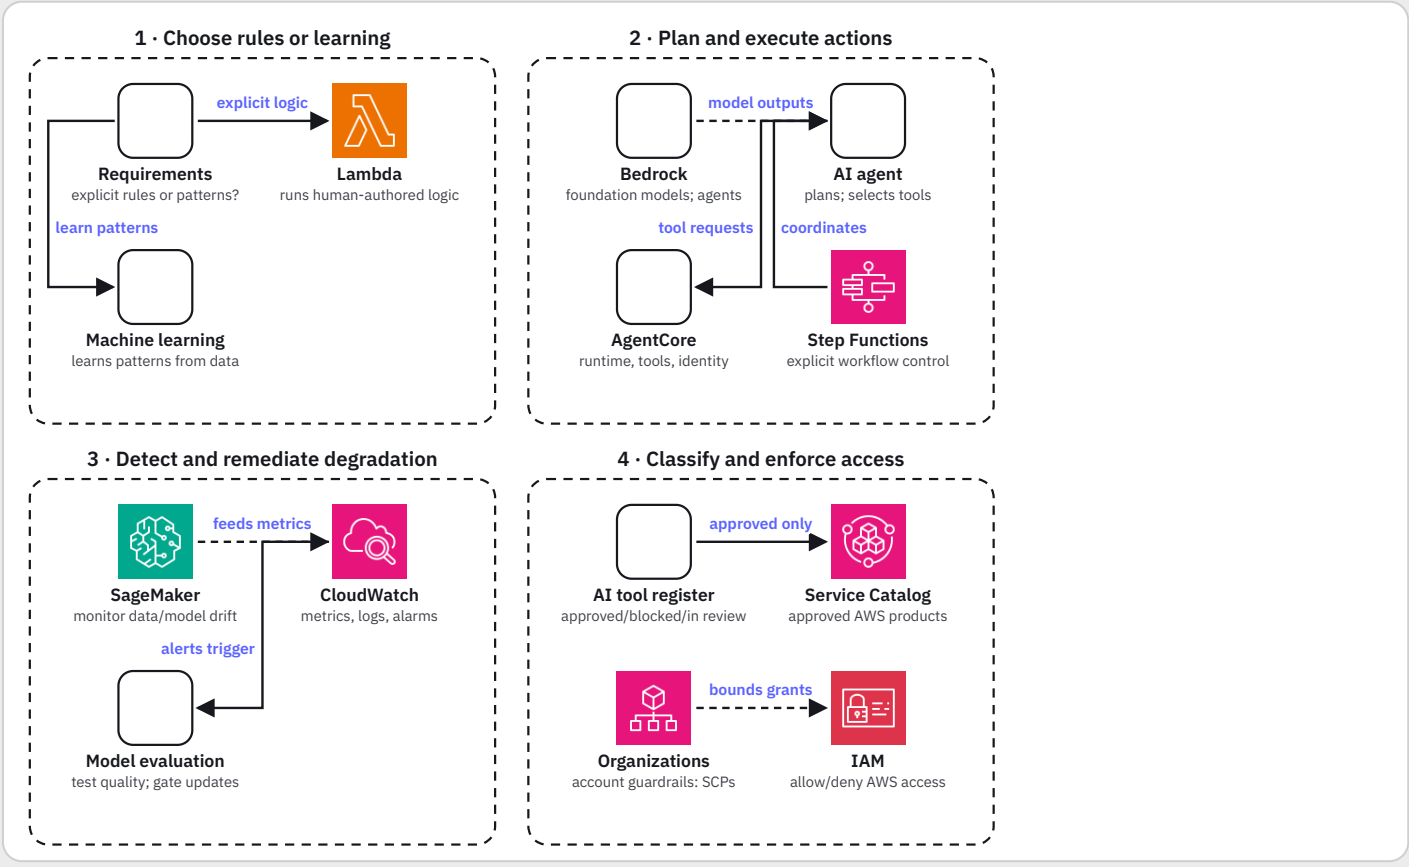
AIB-C01 (beta) · Domain 1 of 4

BlowStack

DOMAIN 1 · AI FUNDAMENTALS AND LITERACY · TASK STATEMENT 1.2

Identify and select appropriate AI solution types

Select, orchestrate, monitor, govern



1 Choose rules or learning

Explicit policies favor rule-based systems; learned AI fits patterns difficult to encode. Lambda executes code, including rules or ML calls. Determinism means repeatability, not correctness or guaranteed compliance.

2 Plan and execute actions

Agents plan and choose tools with bounded autonomy; runtimes validate and execute calls. Multi-agent orchestration routes tasks/messages; Step Functions follows defined workflows, not autonomous reasoning.

3 Detect and remediate degradation

Drift signals changing conditions, not an automatic retraining decision. Evaluate representative data and business metrics; fix pipelines, adjust prompts or models, or retrain when evidence supports it, then re-evaluate.

4 Classify and enforce access

Allowlisting connects transparent approval to enforceable access. Governance assigns owners and approval criteria; management operates controls. An acceptable use policy defines permitted behavior; classification alone blocks nothing.

Exam traps

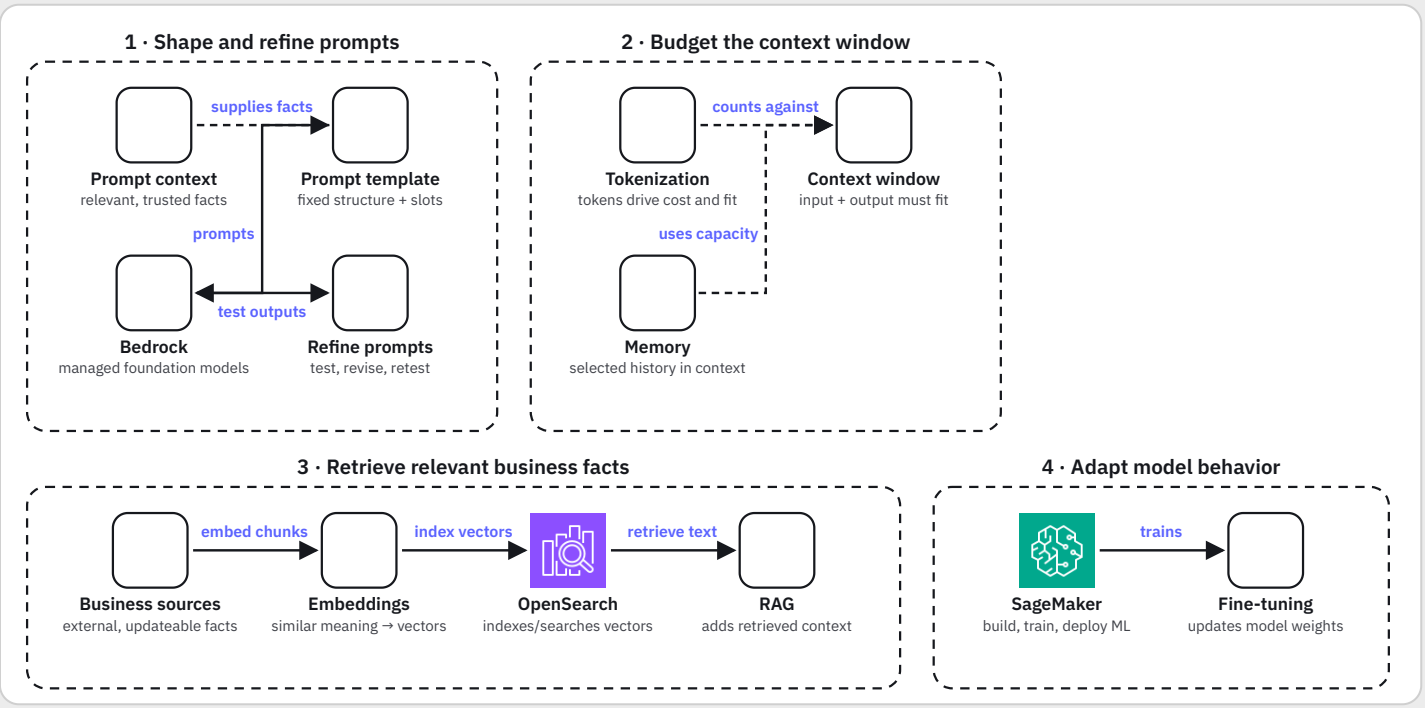
- **Bedrock versus SageMaker**: Bedrock provides managed foundation models; SageMaker supports custom model development and training. Neither eliminates ongoing evaluation.
- **Tool use is not execution**: a model proposes a tool and arguments; application or runtime code validates and executes them. A chatbot response alone does not demonstrate autonomous task completion.
- **A catalog is not a universal blocker**: Service Catalog curates AWS products; IAM and SCPs control AWS permissions, not every external AI website. Shadow AI also includes embedded software features.

Get practice exams and lessons for AWS Certified AI Business Strategist

DOMAIN 1 · AI FUNDAMENTALS AND LITERACY · TASK STATEMENT 1.3

Apply GenAI concepts and techniques

Prompt, budget, retrieve, adapt



1 Shape and refine prompts

Prompt refinement changes instructions, examples, context, or format after testing outputs. Templates standardize requests; neither technique changes model weights.

2 Budget the context window

Context window covers instructions, history, retrieved content, and generated output. Reserve output capacity; select or summarize history. Conversation memory is application-managed, not unlimited model recall.

3 Retrieve relevant business facts

RAG suits changing enterprise facts: embed sources, retrieve relevant passages, and add them to the model prompt. Retrieval does not retrain model weights or guarantee accurate answers.

4 Adapt model behavior

Fine-tuning suits repeated task, style, or format needs when prompting is insufficient. It requires suitable training data, evaluation, deployment, and maintenance; it is not a live knowledge lookup.

Exam traps

- **Tokens are not words**: tokens may be word fragments, punctuation, or whitespace. The same text can have different token counts across model tokenizers.
- **Fine-tuning is not SageMaker-only**: Amazon Bedrock also supports customization for supported models. Available models and customization techniques vary.
- **Retrieved text is not instruction**: source content is reference data, not authority to override the task. Irrelevant or untrusted context can reduce quality or introduce prompt injection.

Get practice exams and lessons for AWS Certified AI Business Strategist

DOMAIN 1 · AI FUNDAMENTALS AND LITERACY · WHAT TO REMEMBER

Concepts and services of AI Fundamentals and Literacy

Concepts to know

Machine learning

ML learns patterns from data to predict or classify; training learns patterns, while inference applies them. Choose **learned patterns over explicit rules** when examples exist but decision logic is difficult to specify.

Generative AI

Generative AI creates content from learned patterns and supplied context; it is a subset of AI, not a synonym for all AI. Choose **content generation rather than prediction alone** for drafting, summarization, and conversational responses.

Retrieval-Augmented Generation

RAG supplies retrieved external information as model context without changing model weights. Choose **RAG for current or proprietary knowledge**; fine-tuning changes weights to adapt task behavior or style. Neither guarantees factual answers.

Core services

Amazon Bedrock

Choose **managed foundation model access** for generative AI applications, retrieval, customization, and agents without operating underlying model-serving infrastructure.

Amazon SageMaker

Choose **custom ML development and lifecycle control** to prepare data, train, evaluate, deploy, and monitor models rather than only invoke managed foundation models.

Amazon Comprehend

Choose **prebuilt NLP analysis** to detect sentiment, entities, or key phrases in text without developing a custom model; this analyzes existing content rather than generating new content.

Get practice exams and lessons for AWS Certified AI Business Strategist