AWS Certified AI Business Strategist Cheat Sheet

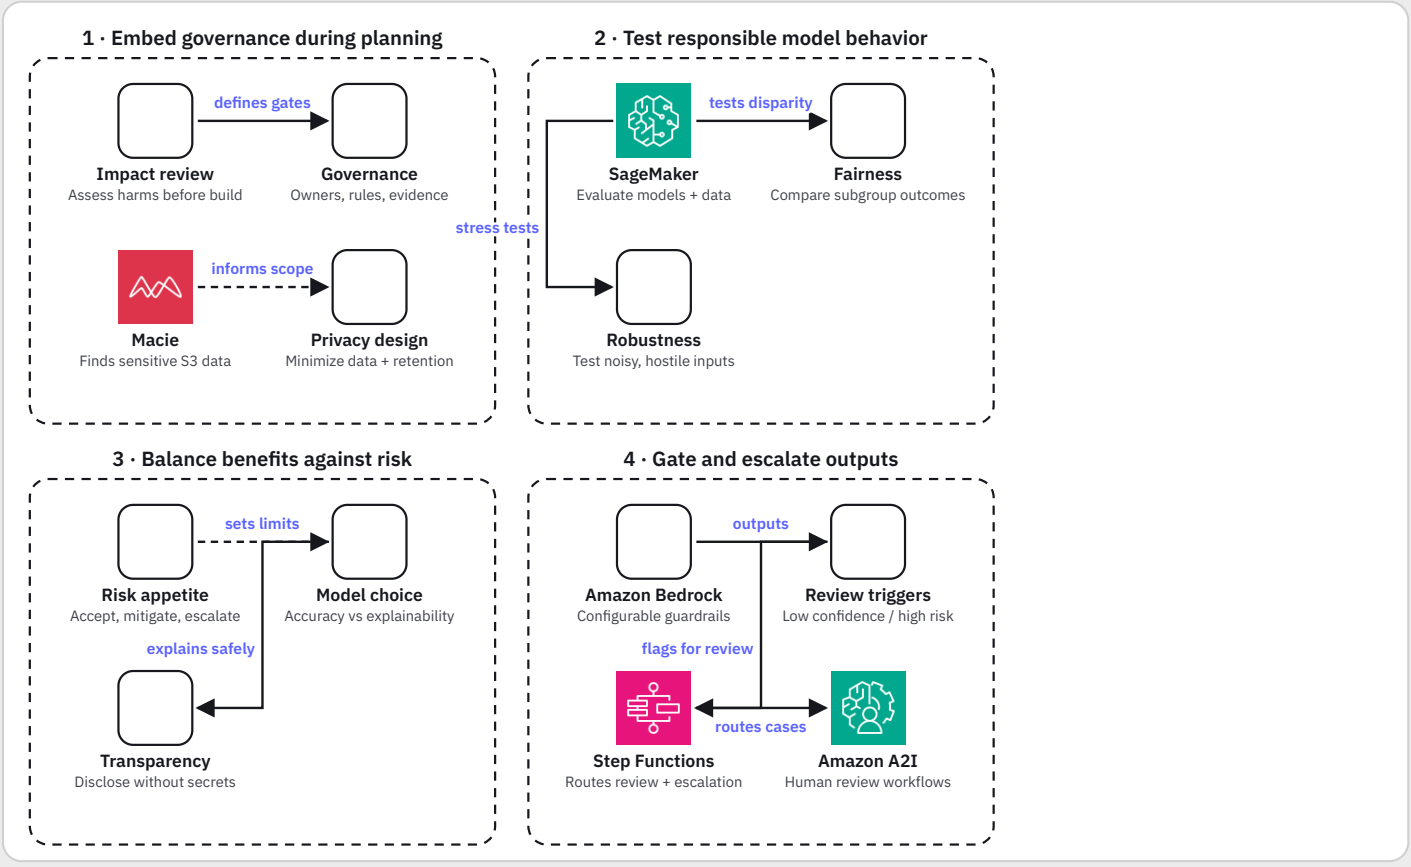
AIB-C01 (beta) · Domain 3 of 4

BlowStack

DOMAIN 3 · AI GOVERNANCE AND RESPONSIBLE AI LEADERSHIP · 24% OF THE EXAM · TASK STATEMENT 3.1

Apply responsible AI principles to business decisions

Assess, test, balance, oversee



1 Embed governance during planning

Start **impact assessment** and Privacy by Design during planning; define owners, controls, and release gates, then reassess throughout the lifecycle. Macie discovers sensitive data; it does not minimize it.

2 Test responsible model behavior

Choose **context-specific fairness** criteria; high overall accuracy can hide subgroup harm. Test robustness beyond clean examples, including unfamiliar inputs and adversarial manipulation; repeat evaluations after changes.

3 Balance benefits against risk

Risk appetite sets acceptance and escalation limits, not exemptions from obligations. Measure interpretability and performance; selective disclosure balances meaningful transparency with privacy and misuse risk.

4 Gate and escalate outputs

Use **human oversight** for high-stakes, ambiguous, or unsupported outputs. Configure guardrails, grounding checks, review thresholds, and escalation; reviewers need authority to correct or reject outputs.

Exam traps

- The **interpretability–performance tradeoff** is not inevitable; benchmark candidate models rather than assuming complex models are superior or explainable models are always less accurate.
- **Guardrails and confidence scores** do not prove factual accuracy; hallucinations can sound confident. Grounding checks and human escalation reduce risk, not eliminate it.
- **Amazon A2I** enables configured human-review workflows; it does not automatically review every Amazon Bedrock response. AWS Step Functions coordinates escalation, not human judgment.

Get practice exams and lessons for AWS Certified AI Business Strategist

AWS Certified AI Business Strategist Cheat Sheet

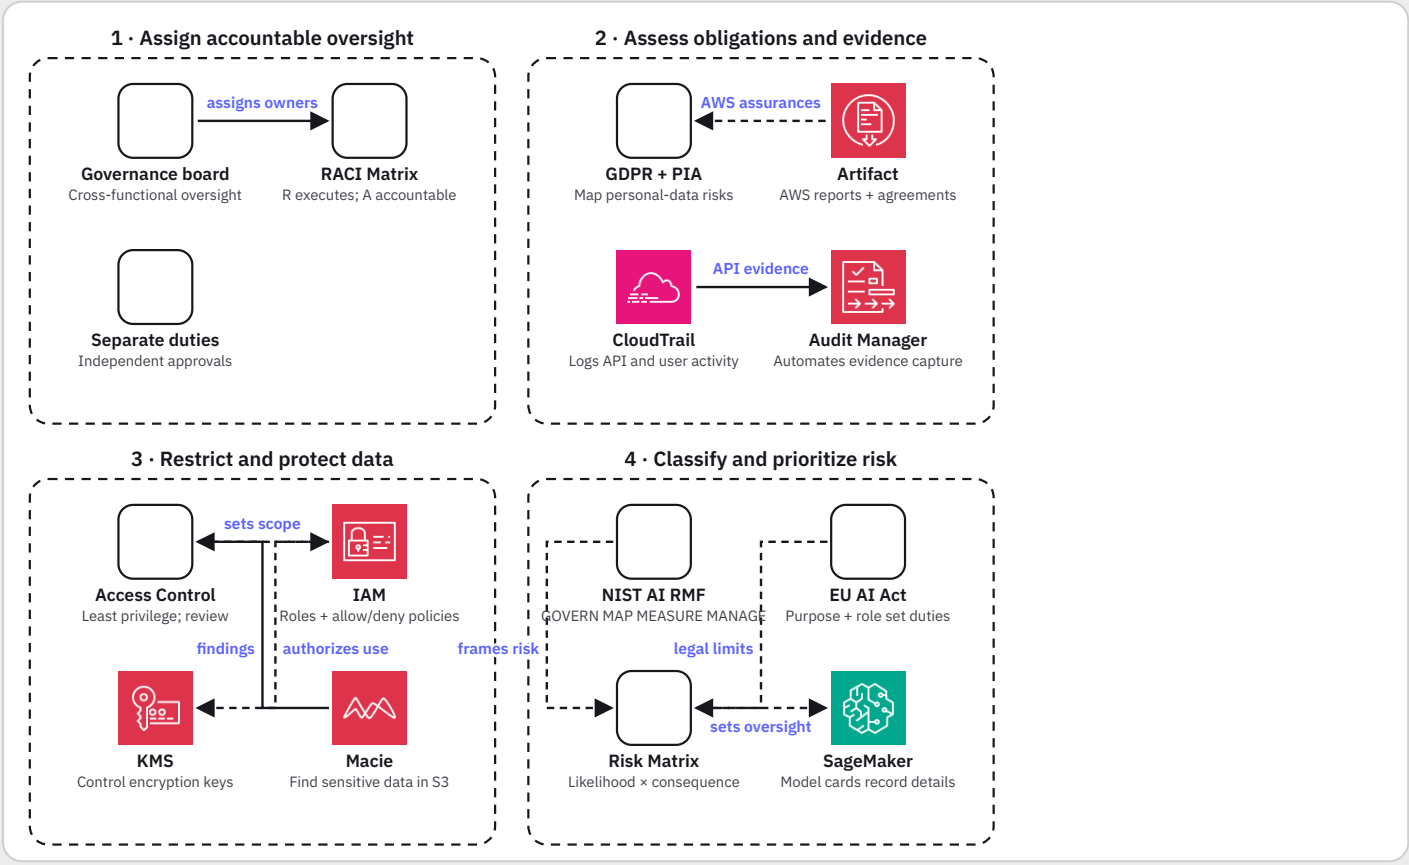
AIB-C01 (beta) · Domain 3 of 4

BlowStack

DOMAIN 3 · AI GOVERNANCE AND RESPONSIBLE AI LEADERSHIP · TASK STATEMENT 3.2

Establish AI governance structures and ensure regulatory compliance

Assign, verify, protect, classify



1 Assign accountable oversight

Accountability spans business, legal, risk, and technology. Governance defines authority and standards; management executes. RACI: R performs, A owns outcomes, C advises, I receives updates. Separate execution from approval.

2 Assess obligations and evidence

Use a **Privacy Impact Assessment** before new or changed personal-data processing. AWS Compliance involves shared responsibility: verify service/Region scope and retain evidence of customer controls.

3 Restrict and protect data

Least privilege requires scoped permissions and ongoing review. Protect data at rest and in transit; use KMS for encryption-key control. Permission to use a key is separate from permission to access stored data.

4 Classify and prioritize risk

Legal classification determines prohibited uses and high-risk duties under the EU AI Act. NIST AI Risk Management Framework structures oversight; a Risk Matrix prioritizes controls. Reassess throughout the lifecycle.

Exam traps

- **Evidence is not compliance**: Audit Manager automates evidence collection; Artifact supplies AWS reports and agreements. Neither establishes the customer's GDPR compliance.
- **Discovery is not enforcement**: Macie finds sensitive S3 data; IAM controls access and KMS controls keys. Encryption does not replace least privilege.
- **Risk score is not legal tier**: NIST AI RMF is voluntary. A low Risk Matrix score cannot override EU AI Act prohibitions or high-risk obligations; SageMaker does not confer compliance.

Get practice exams and lessons for AWS Certified AI Business Strategist

AWS Certified AI Business Strategist Cheat Sheet

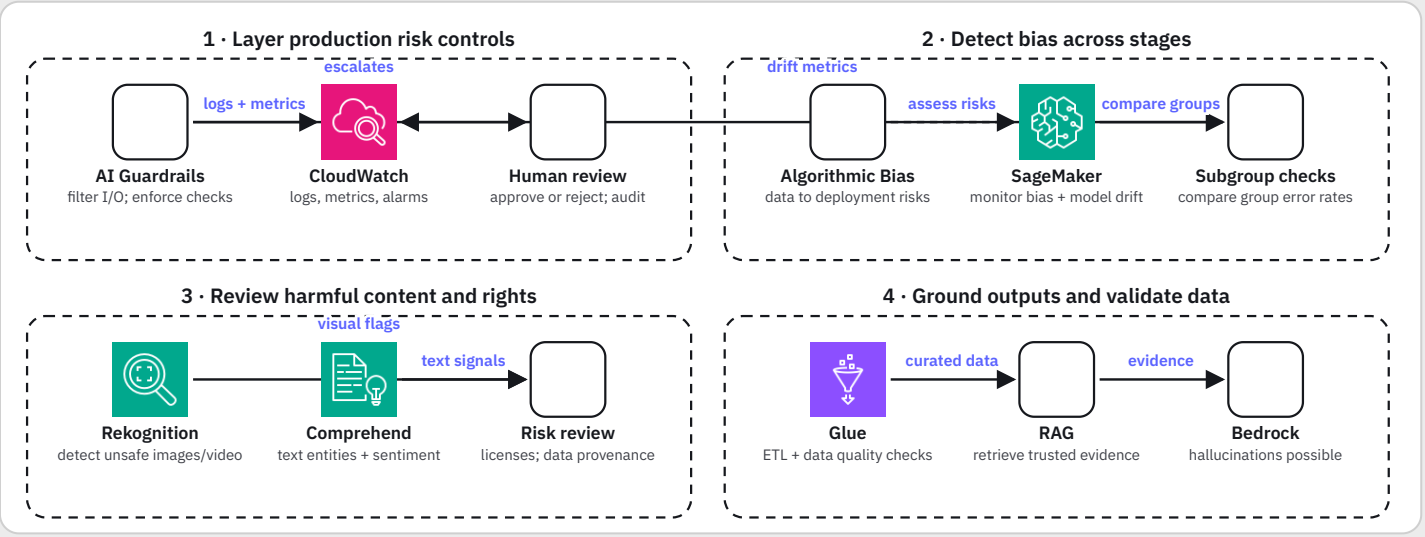
AIB-C01 (beta) · Domain 3 of 4

BlowStack

DOMAIN 3 · AI GOVERNANCE AND RESPONSIBLE AI LEADERSHIP · TASK STATEMENT 3.3

Identify enterprise AI risks and direct mitigation strategies

Monitor, compare, filter, ground



1 Layer production risk controls

Use **layered controls**: guardrails filter I/O, permissions constrain actions, alarms notify, and humans review high-risk cases. Audit Logging needs protected retention; monitoring alone does not remediate.

2 Detect bias across stages

Bias can enter data, labels, objectives, prompts, retrieval, or deployment. Repeat **subgroup analysis** after launch; averages hide disparities. Drift triggers investigation, not automatic proof of unfairness.

3 Review harmful content and rights

Rekognition flags visual content; Comprehend supplies text signals, not universal moderation. **IP rights** require separate license and output review; provenance records sources, permissions, and transformations.

4 Ground outputs and validate data

Retrieval-Augmented Generation supplies enterprise evidence without retraining, but claims still need verification. Configure quality rules and failure handling; investigate drift before choosing pipeline fixes or retraining.

Exam traps

- **Prompt instructions alone** are not enforceable guardrails. Independent checks, least-privilege permissions, monitoring, and human escalation address different failure modes.
- **High overall accuracy** does not establish fairness. Subgroup disparities can remain hidden, and changing data or deployment conditions require ongoing bias monitoring.
- **RAG is not a guarantee** of truth or IP clearance. Retrieved sources may be wrong, stale, or unlicensed; grounding checks and provenance-backed rights review remain necessary.

Get practice exams and lessons for AWS Certified AI Business Strategist

DOMAIN 3 · AI GOVERNANCE AND RESPONSIBLE AI LEADERSHIP · WHAT TO REMEMBER

Concepts and services of AI Governance and Responsible AI Leadership

Concepts to know

Management and governance

Governance defines accountability, standards, and evidence; management implements and operates controls. **Governance by design** embeds cross-functional ownership, risk assessment, approval gates, and monitoring from project planning onward.

AI Guardrails

Guardrails constrain unsafe content, sensitive-data exposure, and unauthorized actions. **Layered controls** combine input/output checks, permissions, and monitoring; model instructions alone cannot reliably enforce policy or eliminate hallucinations.

Human-in-the-Loop

Human review adds judgment and accountability to automated decisions. **Escalate high-impact or uncertain outputs** using explicit review thresholds and authority to intervene; a confidence score alone does not establish correctness.

Core services

Amazon Bedrock

Choose for foundation-model applications needing **content safeguards and grounding**. Guardrails filter supported risks; retrieval supplies relevant evidence but does not guarantee factual answers.

Amazon SageMaker

Choose for **ML lifecycle oversight**: evaluate bias and explainability, monitor data and model quality, and investigate drift. Findings inform remediation, retraining, or deployment approval.

AWS CloudTrail

Choose for **AWS activity and API audit evidence**: investigate who acted on which resources and when. It supports accountability, not model-quality monitoring or automatic regulatory compliance.

Get practice exams and lessons for AWS Certified AI Business Strategist