AWS Certified AI Practitioner Cheat Sheet

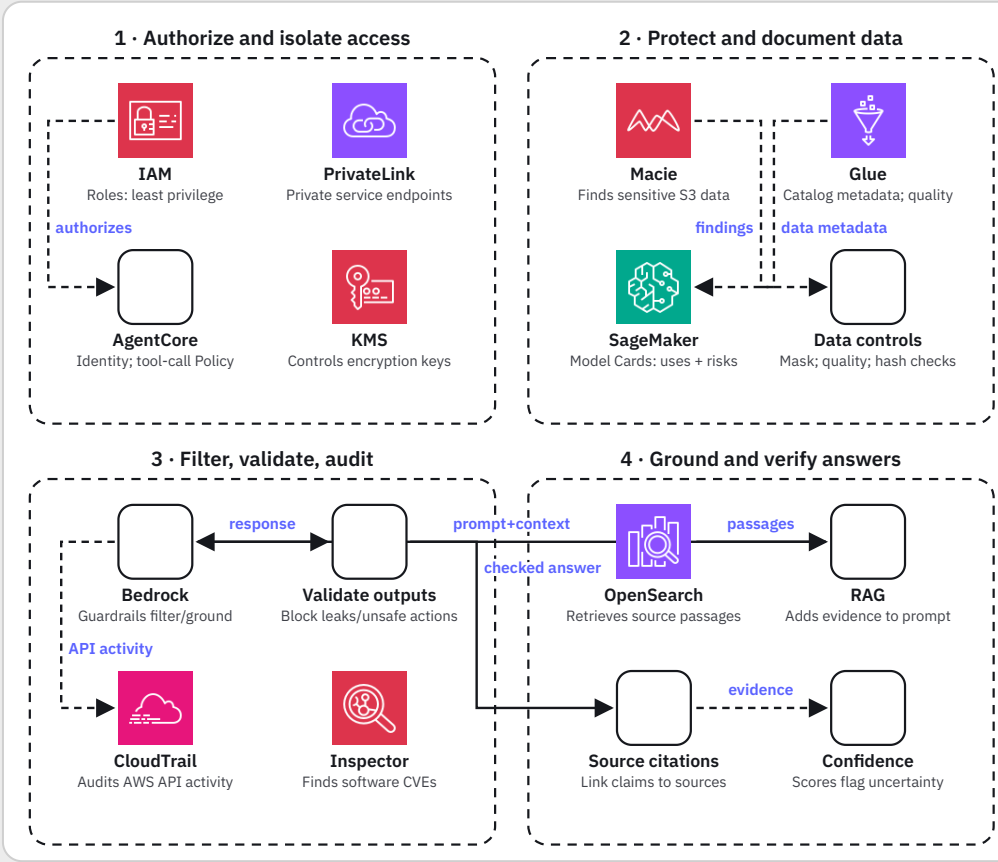
AIF-C01 · Domain 5 of 5

BlowStack

DOMAIN 5 · SECURITY, COMPLIANCE, AND GOVERNANCE FOR AI SOLUTIONS · 14% OF THE EXAM · TASK STATEMENT 5.1

Explain methods to secure AI systems

Authorize, protect, filter, ground



1 Authorize and isolate access

Shared responsibility leaves data, permissions, and configuration with the customer. AWS protects cloud infrastructure. KMS supports encryption at rest; TLS protects transit. PrivateLink avoids public internet.

2 Protect and document data

Data provenance tracks origins and transformations; catalogs describe data, not access grants. AWS Lake Formation enforces fine-grained access. Masking protects privacy; hashes detect changes.

3 Filter, validate, audit

Prompt injection can arrive in prompts or retrieved text. Guardrails filter toxicity and sensitive data; outputs still need validation. Enable Bedrock invocation logging for prompts/responses; protect log access and retention.

4 Ground and verify answers

Grounding uses retrieved evidence without retraining. Contextual grounding checks flag unsupported answers; citations enable verification. Confidence thresholds can trigger abstention or human review, not prove truth.

Exam traps

- **RAG does not eliminate hallucinations**: retrieved content can be wrong or contain prompt injection. Citations and high confidence do not prove factual accuracy.
- **Macie discovers; it does not block**: sensitive-data discovery in S3 is not inline prompt filtering, automatic masking, encryption, or access enforcement.
- **CloudTrail is not a transcript**: AWS API auditing does not capture every model prompt and response. Bedrock invocation logging must be enabled for supported model calls.

Get practice exams and lessons for AWS Certified AI Practitioner

AWS Certified AI Practitioner Cheat Sheet

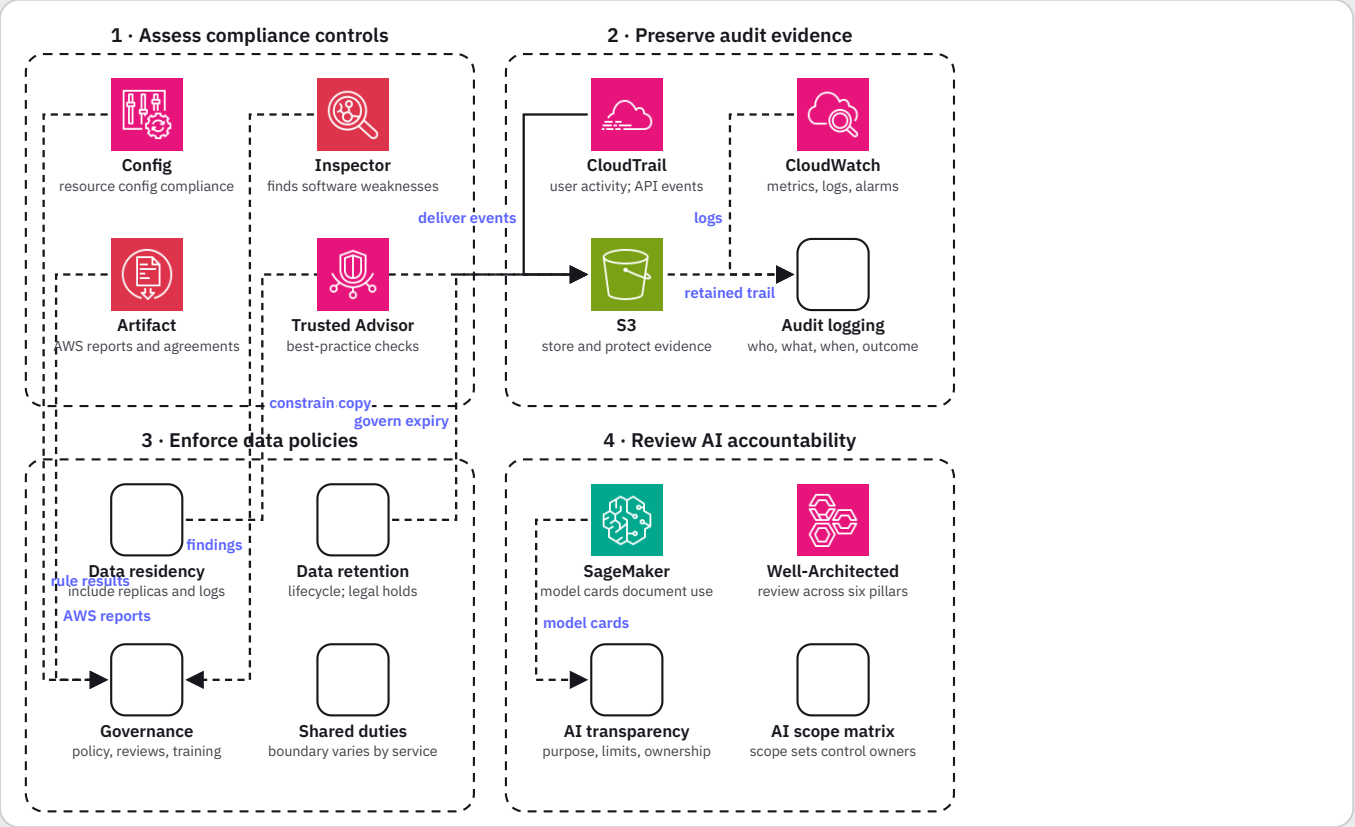
AIF-C01 · Domain 5 of 5

BlowStack

DOMAIN 5 · SECURITY, COMPLIANCE, AND GOVERNANCE FOR AI SOLUTIONS · TASK STATEMENT 5.2

Recognize governance and compliance regulations for AI systems

Assess, record, govern, review



1 Assess compliance controls

Config evaluates resource configurations against rules; Inspector finds vulnerabilities. Artifact supplies AWS compliance documents; Trusted Advisor recommends improvements. Findings alone do not remediate resources.

2 Preserve audit evidence

Audit evidence combines activity records with operational logs. CloudTrail records API activity; CloudWatch monitors health. Event selection, delivery, access protection, and retention require configuration.

3 Enforce data policies

Residency covers primary data, replicas, backups, logs, and downstream systems. Retention balances deletion requirements with legal holds. Governance assigns accountability; management implements and monitors controls.

4 Review AI accountability

The **Generative AI Security Scoping Matrix** maps AI use patterns to security responsibilities. Governance sets review cadence, change-triggered reviews, transparency standards, and team training requirements.

Exam traps

- **AWS compliance is not workload compliance**: Artifact reports cover AWS controls. Under the AWS shared responsibility model, customers retain service-dependent data, access, and configuration duties.
- **A Region choice is not residency compliance**: replicas, backups, logs, and downstream transfers also matter. S3 Lifecycle expiration does not override an applicable Object Lock legal hold.
- **Review is not certification**: Well-Architected reviews identify workload risks, and model cards support transparency. Neither automatically establishes regulatory compliance.

Get practice exams and lessons for AWS Certified AI Practitioner

DOMAIN 5 · SECURITY, COMPLIANCE, AND GOVERNANCE FOR AI SOLUTIONS · WHAT TO REMEMBER

Concepts and services of Security, Compliance, and Governance for AI Solutions

Concepts to know

AWS shared responsibility model

AWS secures the underlying cloud infrastructure; customers retain responsibilities for their data, permissions, and application controls. The decision is **who owns each control**, with the boundary varying by service.

Prompt Injection

Malicious instructions in prompts or retrieved content can redirect model behavior or trigger unauthorized actions. Choose **layered defenses**: separate untrusted content, restrict tool permissions, apply guardrails, and validate outputs.

Data Provenance

Provenance records data origins, permissions, transformations, and downstream uses. Choose **traceable data lineage** when a scenario requires investigating training data, reproducing results, or demonstrating authorized data use.

Core services



AWS Identity and Access Management (IAM)

Choose **roles and least-privilege policies** to control access to AI resources and data. Roles provide temporary credentials; permissions define allowed actions on resources.



AWS CloudTrail

Choose **API activity audit trails** to investigate who performed an AWS action, when, and on which resource. Model prompts and responses require separate interaction logging.



AWS Artifact

Choose **AWS compliance reports and agreements** when auditors need AWS documentation. It supplies evidence about AWS controls, not automatic certification of a customer's AI workload.

Get practice exams and lessons for AWS Certified AI Practitioner