



CloudNinja.pro

BUSINESS · AIB-C01

AWS Certified AI Business Strategist

AIB-C01 Cheat Sheet

CloudNinja · cloudninja.pro · Free — no signup required

The business judgement AIB-C01 tests: AI literacy, use-case selection, build-buy-partner, KPIs and ROI, responsible-AI tradeoffs, governance, readiness and scaling from pilot to enterprise. It is organised by the exam's four domains. AWS services appear only at the strategic level the exam guide asks for. No coding or AWS hands-on experience assumed.

EXAM SNAPSHOT

- ▶ **85 questions** (beta) — multiple choice + multiple response
- ▶ **170 minutes** (beta) · guide lists 130 min
- ▶ **700 / 1000** to pass (compensatory)
- ▶ **\$50 USD** beta · \$100 USD standard
- ▶ **AI Fundamentals & Literacy** — 24%
- ▶ **AI Strategy & Business Value** — 28%
- ▶ **Governance & Responsible AI** — 24%
- ▶ **Readiness & AI Transformation** — 24%

Beta status (October 2026): AIB-C01 is the first exam in AWS's new *Business* certification category. Beta registration opened 1 Sep 2026 and delivery began 29 Sep 2026, at Pearson VUE test centres or online, in English and Japanese. You can take the beta only once. Pass by **15 Feb 2027** to receive an extra *Early Adopter* digital badge. The certification is valid for 3 years. There are no prerequisites, and about 6 months of working with or alongside AI initiatives is recommended. AWS has not yet published the standard question count or the GA date.

FREE

AIB-C01 practice exam



0 · How AIB-C01 Thinks

AIB-C01 tests **portable business judgement** through enterprise scenarios. AWS services provide practical context, but they are not what is being tested. A typical question reads: "a vendor proposes a \$2M AI deployment. Is it sound, premature, or missing critical success factors?"

AIB-C01 (Business)	AIF-C01 (AI Practitioner)	What that means for you
Validates business judgement : which AI investments to pursue, the business case, governance, adoption.	Validates technical literacy in AI/ML/GenAI concepts and AWS AI services.	Neither exam is a prerequisite for the other, so you can earn both.
"Does not assess knowledge of AWS AI services". Services appear at a <i>strategic</i> level only.	Tests named services and features (Bedrock, SageMaker tools, Comprehend...).	Know <i>what</i> Bedrock / SageMaker AI / Quick are for, not how to configure them.
Target: PMs, program managers, sales/BD, line-of-business leaders, consultants, analysts, marketers.	Target: anyone in a role that touches AI on AWS.	Answer as the business owner, not as the engineer.

Out of scope for the candidate

- Coding or training models, data or feature engineering, hyperparameter tuning, building ML pipelines, statistical analysis of models.
- Configuring, deploying or administering any AWS service, writing IAM policies, setting up KMS, or doing hands-on data labelling.
- Out-of-scope service families: compute (EC2, Lambda), networking (VPC, CloudFront), databases and storage (RDS, S3), containers (ECS, EKS), DevOps (CodePipeline, CloudFormation), IoT and media.

Exam Tip: If one answer is a technical fix (e.g. "retrain with a different algorithm") and another is a business control (e.g. "set a baseline, define KPIs, add human review, pause the rollout"), the exam almost always wants the *business decision*. Distractors often skip a prerequisite such as a baseline, a governance owner, data readiness or executive sponsorship.

1 · AI Fundamentals & Literacy (24%)

AI vs. ML vs. Generative AI

Term	What it is	Business example
Artificial intelligence (AI)	The broad field of systems that perform tasks normally needing human intelligence: perceive, reason, decide, act.	Any "smart" capability, from rules-plus-ML fraud screening to chat assistants.
Machine learning (ML)	A subset of AI. Models <i>learn patterns from historical data</i> instead of following hand-written rules. Output: a prediction, classification or score.	Churn prediction, demand forecasting, credit-risk scoring, recommendations.
Deep learning	ML that uses multi-layer neural networks. Strong on unstructured data (images, audio, text).	Defect detection in photos, speech-to-text.
Generative AI (GenAI)	A subset of deep learning. <i>Foundation models</i> (FMs) and LLMs, pre-trained on huge datasets, <i>create new content</i> (text, code, images).	Drafting emails, summarising contracts, chat assistants, code generation.

Core Concepts in Business Terms

Concept	Plain-English meaning
Algorithm	The learning recipe (method). You do not pick algorithms on this exam.
Model	The trained artefact the algorithm produces from data. It encodes learned patterns.
Training	Feeding <i>historical data</i> to the algorithm so it learns patterns. If the history is biased or out of date, the model is too.
Inference	Using a trained model on new data to get an output. This is where most ongoing run cost sits.
Prediction / output	A <i>probabilistic</i> answer (score, label, generated text). It is not guaranteed correct, so plan for errors.
Labelled data	Historical examples with known outcomes (e.g. "fraud / not fraud"). Supervised ML needs it, and creating it costs time and money.
Model drift	Performance degrades as the real world changes (new products, customer behaviour, regulations). This makes monitoring and retraining a <i>recurring</i> cost.

Data Types & Why Data Quality Matters

- **Structured data:** rows and columns (CRM, ERP, transactions). Fits classic ML (forecasting, scoring).
- **Unstructured data:** documents, emails, images, audio, video. This is about 80% of enterprise data and is where GenAI and document extraction add most value.
- **Semi-structured:** JSON, logs, forms with fields.
- **Data quality dimensions:** accuracy, completeness, consistency, timeliness, relevance, representativeness. *Garbage in, garbage out* applies: poor data is the most common root cause of failed AI pilots.
- **Representative data:** if the training data under-represents a customer group, results for that group are worse. This is a fairness *and* a revenue risk.

Global Frameworks & Shared AI Vocabulary

Standard / framework	Purpose
ISO/IEC 23053	Framework for AI systems that use ML. It describes the components and functions of an ML-based AI system and gives a common reference model.
ISO/IEC 42001	An <i>AI management system</i> (AIMS) standard, the "ISO 27001 for AI". Organisations can be certified against it. It covers policies, risk, impact assessment and continual improvement. (AWS holds 42001 certification for many AI services.)
ISO/IEC 22989	AI concepts and terminology, the unified vocabulary that 23053 and 42001 build on.
ISO/IEC 23894	Guidance on AI risk management.
NIST AI RMF	A voluntary US risk framework with four functions: <i>Govern, Map, Measure, Manage</i> .

Exam Tip: "Teams across regions use different AI terms and governance approaches" → adopt a *global standard / unified vocabulary* (ISO/IEC 22989 / 23053). "We want an auditable, certifiable AI management system" → **ISO/IEC 42001**.

Rule-Based Automation vs. AI

Choose rule-based automation when...	Choose AI / ML when...
Logic is stable, known and can be written as if-then rules.	Patterns are too complex or change too fast to hand-code.
You need 100% deterministic, auditable outcomes (e.g. tax calculation, eligibility thresholds).	A probabilistic answer is acceptable and errors can be caught or tolerated.
Inputs are structured and predictable.	Inputs are unstructured (text, images, speech) or highly variable.

Choose rule-based automation when...	Choose AI / ML when...
There is little or no historical data.	Rich, representative historical data exists.
Low cost and fast delivery matter more than adaptability.	Value comes from personalisation, prediction or generation at scale.

Common Categories of AI Tools

Category	Typical business application
Recommendation engines	Next-best product or offer, content personalisation, cross-sell.
Natural language processing (NLP)	Sentiment analysis, classification, entity extraction, translation, summarisation.
Computer vision	Quality inspection, identity verification, shelf monitoring, content moderation.
Document extraction (IDP)	Invoices, claims, KYC forms → structured data, which cuts manual keying.
AI for customer operations	Virtual agents, agent assist, call summarisation, routing, contact-centre analytics.
Forecasting / prediction	Demand, churn, fraud, predictive maintenance.
Generative assistants & agents	Knowledge Q&A, drafting, research, code assist, multi-step task automation.

AI Agents vs. Other AI Solutions

Solution	Behaviour	Business implication
Predictive ML model	One input → one prediction. No actions.	Narrow, measurable, easiest to govern.
GenAI assistant / chatbot	Responds to prompts. A human decides what to do with the output.	Productivity gain, with the human in control.
AI agent	<i>Autonomy</i> : plans multi-step work toward a goal. <i>Tool use</i> : calls APIs and systems and takes actions. <i>Can do agent-to-agent</i> communication. Supports <i>orchestration</i> (a supervisor agent with specialist agents).	Highest leverage <i>and</i> highest risk. Needs scoped permissions, guardrails, audit trails, escalation rules and human approval for high-impact actions.

Ongoing Monitoring & Shadow AI

- **AI is never "done"**. Data drift, concept drift and model or vendor version changes degrade performance. Budget for monitoring, re-evaluation and retraining. These are *operating* costs, not just project costs.
- **Shadow AI**: employees using unapproved AI tools, such as pasting customer data into public chatbots. Risks include data leakage, IP exposure, compliance breaches and unmanaged cost.
- **Mitigation through transparent classification**: publish a register of AI tools as *Approved*, *Under evaluation* or *Blocked*, with a fast intake process. Add an acceptable-use policy and offer sanctioned alternatives so people do not route around controls.

Exam Tip: "Employees are using public GenAI tools with company data" → the answer is *not* "ban all AI". It is to **classify tools (approved / under evaluation / blocked)**, publish usage policy, provide an approved enterprise alternative and train staff. "Model accuracy has fallen since launch" → **drift**: monitor, then retrain or update.

Generative AI Concepts & Techniques

Prompt Engineering Principles

- **Be specific**: state the task, audience, format, length and tone ("Summarise for a CFO in 5 bullets").
- **Give context**: the role ("You are a claims analyst"), background facts and constraints.

- **Show examples:** *zero-shot* (no examples) vs. *few-shot* (2–5 examples of desired output).
- **Break complex tasks into steps:** chain-of-thought ("think step by step") or prompt chaining.
- **Constrain the output:** "answer only from the provided document; say 'I don't know' otherwise". This reduces hallucination.
- **Iterate and template:** test prompts, then version and reuse the ones that work across teams.

Tokens & Context Windows

- **Token:** a chunk of text (about ¾ of an English word). GenAI is usually *priced per input and output token*.
- **Context window:** the maximum tokens a model can consider at once (prompt + documents + history + answer).
- **Symptoms of hitting limits:** truncated answers, the model "forgetting" earlier conversation, ignoring parts of long documents, higher cost and latency.
- **Business fixes:** summarise or chunk content, use RAG to retrieve only the relevant passages, choose a larger-context model (at higher cost), and cap output length.

Model Adaptation — Lightest to Heaviest

Technique	What it does	Use when	Cost / effort
Prompt engineering	Better instructions and examples. The model is unchanged.	Always the first step.	Lowest; immediate.
RAG (Retrieval Augmented Generation)	Retrieves relevant company documents at query time and grounds the answer in them, with citations.	Answers must reflect <i>current, proprietary</i> knowledge that changes often.	Low–medium. No retraining; update the documents instead.
Fine-tuning	Further trains a model on labelled examples to change <i>style, tone, format or domain behaviour</i> .	A consistent brand voice, specialised task format, or domain jargon.	Medium–high. Needs quality labelled data, and you retrain as needs change.
Continued pre-training	Trains on large unlabelled domain corpora.	Deep domain language (legal, biomedical).	High.
Build from scratch	Train your own foundation model.	Rarely justified.	Very high; months.

Exam Tip: "Assistant must answer from our latest policy documents, which change weekly" → **RAG** (not fine-tuning). "Outputs must match our brand voice and format consistently" → **fine-tuning**. "Cheapest, fastest improvement" → **prompt engineering**. "Long contracts get cut off" → context-window limit, so chunk the document or use RAG.

2 · AWS AI Services — Strategic Level Only

You need to know *what each service is for and when a business would choose it*. You do not need to know how to configure it. These are the in-scope services in the "AI and Machine Learning" and "Business Intelligence" categories, at the "basic application only" level.

Service	What it is	Choose it when...
Amazon Bedrock	A fully managed GenAI platform that gives access to many foundation models (Amazon Nova, Anthropic Claude, Meta Llama, Mistral and others) through one API. No infrastructure to manage. Your prompts and data are <i>not</i> used to train the base models and are not shared with model providers.	You want GenAI apps, assistants or agents quickly, with model choice, enterprise security and pay-per-use pricing.
Amazon SageMaker AI	A managed platform to build, train, deploy and govern <i>custom</i> ML models. Includes no-code SageMaker Canvas and ready-made models in JumpStart.	You have unique data and need a <i>custom</i> or highly specialised model, or full control over training and hosting, and you have (or can hire) data-science skills.

Service	What it is	Choose it when...
Amazon Quick	An AI-powered business assistant suite for knowledge workers. Chat agents are grounded in company data. It includes <i>Quick Sight</i> (BI dashboards), <i>Quick Flows</i> (task automation), <i>Quick Automate</i> (end-to-end processes), <i>Quick Research</i> (cited reports), Spaces, and connectors to apps such as Slack, Microsoft 365 and CRMs.	You want ready-to-use AI for business users (ask questions of data, automate tasks, research) with <i>no build</i> . Priced per seat.

Amazon Bedrock — Features a Strategist Should Recognise

Feature	Business meaning
Model choice	Switch or mix models by cost, quality and latency without re-platforming. This reduces vendor lock-in.
Knowledge Bases	Managed <i>RAG</i> : connects company documents so answers are grounded and cited.
Guardrails	Configurable safeguards: content filters (hate, violence, misconduct, etc.), denied topics, word filters, PII redaction, contextual grounding checks (to catch hallucinations), and prompt-attack detection. They apply consistently across models.
Agents / AgentCore	Build and run AI agents that plan, call tools and take actions, with identity, memory and observability.
Model evaluation	Compare models on your own data (automatic or human review) before you commit.
Customisation	Fine-tuning, continued pre-training and distillation (a smaller, cheaper model learns from a larger one).

Managed vs. Custom — The Core Choice

Option	Time to value	Control	Skills needed
Ready-made app (Amazon Quick, SaaS from AWS Marketplace)	Days–weeks	Lowest	Business users, admins
Managed GenAI platform (Bedrock)	Weeks	Medium	App developers; no ML scientists needed
Custom ML (SageMaker AI)	Months	Highest	Data scientists, ML engineers

Amazon Q Business **DEPRECATED** is closed to new customers. AWS recommends **Amazon Quick** as its successor, and existing customers can migrate their Q Business index into Quick. Also, **Amazon QuickSight** now lives inside Amazon Quick as *Quick Sight*, and **Amazon SageMaker** (the ML platform) is now **SageMaker AI**. Older study material may use the old names.

Exam Tip: "Business users need AI answers from company data and dashboards with no development" → **Amazon Quick**. "Build a customer-facing GenAI app fast with a choice of models" → **Amazon Bedrock**. "Unique proprietary prediction problem, with an in-house data-science team, needing full control" → **SageMaker AI**.

3 · AI Strategy & Business Value Creation (28%)

This is the largest domain. It covers picking the right problems, choosing how to deliver, proving value and positioning for advantage.

High-Impact Use Cases by Function

Function	Use cases	Outcome KPI
Customer operations	Virtual agents, agent assist, call summaries, intelligent routing	Handle time, first-contact resolution, CSAT, cost per contact
Sales & marketing	Personalisation, lead scoring, content generation, proposal drafting	Conversion rate, pipeline velocity, content cycle time
R&D / product	Research summarisation, design exploration, simulation	Time to market, experiments per quarter
Software development	Code generation, test generation, code modernisation	Developer productivity, lead time, defect rate
Back office (finance, HR, legal)	Document extraction, contract review, knowledge Q&A	Processing time, error rate, cost per transaction

Build vs. Buy vs. Partner

Option	Best when	Pros	Cons
Buy	The capability is commodity (not a differentiator), a mature product exists and speed matters.	Fastest, predictable cost, vendor maintains it.	Less differentiation, lock-in, data-residency and compliance fit must be checked.
Build	The capability is a <i>core differentiator</i> , built on proprietary data, and you have the skills.	Full control, IP ownership, competitive moat.	Slowest, highest cost and talent needs, you own maintenance and drift.
Partner	You lack skills or capacity, or need industry expertise, but want more tailoring than off-the-shelf.	Expertise and speed, with shared risk.	Dependency, knowledge-transfer risk, contract complexity.

- **Decision factors** (from the exam guide): budget, timelines, internal capabilities, vendor proposals, regulatory compliance requirements. Add data sensitivity, total cost of ownership (TCO) and exit strategy.
- **Evaluating a vendor proposal**: check for a defined business problem, baseline and KPIs, data readiness, a governance and responsible-AI plan, security and compliance evidence, the pricing model at scale, the change-management plan and exit terms.
- **AWS Marketplace**: a catalogue to discover, try and procure third-party AI models, SaaS AI products, data products and consulting services, billed through your AWS account. Supports private offers and procurement controls. It is a key input to *buy* and *partner* decisions.

Prioritising Initiatives

Quadrant (value × feasibility)	Action
High value / high feasibility	Do first. These are quick wins that build momentum and funding.
High value / low feasibility	Strategic bets. Close the gaps first (data, skills, governance), then stage the investment.
Low value / high feasibility	Do opportunistically, or bundle with other work.
Low value / low feasibility	Deprioritise.

Score each initiative on **business value**, **feasibility** (data, technology, skills), **sustainability** (ongoing cost, maintainability, environmental impact) and **strategic alignment**. Review the portfolio regularly and decide whether to **scale** (KPIs met, ready to expand), **pause** (blocked by a fixable gap) or **terminate** (no path to value, so stop the sunk-cost spend).

When AI Is *Not* the Answer

- A simple rule, a process fix or better reporting would solve the problem.

- There is no reliable, representative data, or no right to use it.
- The decision requires 100% determinism or explainability that the model cannot provide.
- The cost of errors is unacceptable and human review would remove the benefit.
- Value is too small to cover TCO (build + run + monitor + govern + change).
- "AI for AI's sake": there is no defined business problem or owner.

Transitioning Processes to AI (or Between AI Platforms)

- **Business continuity:** run the old and new processes in parallel, use phased cutover, keep a rollback plan and manual fallback.
- **Cost implications:** migration effort, re-testing prompts, retraining and re-evaluating models, contract overlap.
- **Data readiness:** data mapping, quality and access rights on the target platform.
- **Performance impact:** benchmark accuracy, latency and user experience against the baseline before switching.

Exam Tip: "Many pilots, unclear which to fund" → score against **value, feasibility, alignment**, then pick high-value/high-feasibility quick wins. "Pilot missed its KPIs and has no path to fix them" → **terminate**. "Pilot blocked by data access" → **pause** and fix the data. "Commodity capability, tight deadline" → **buy**. "Differentiating capability on proprietary data" → **build**.

Measuring Value — KPIs, Baselines, ROI

Element	What to know
Tangible KPIs	Cost reduction, revenue growth, time saved, throughput, error rate, cost per transaction.
Intangible KPIs	Customer satisfaction (CSAT/NPS), employee productivity and satisfaction, decision quality, brand and risk posture.
Baseline	Measure current performance <i>before</i> deploying AI. Without it you cannot prove impact. It is the most common missing step in scenarios.
Leading indicators	Predict success early: adoption and active usage, task completion, user feedback scores, human-override rate, data-pipeline readiness, executive engagement.
Lagging indicators	Confirm success later: realised savings, revenue, ROI, customer retention.
ROI	$(\text{total benefits} - \text{total costs}) \div \text{total costs} \times 100\%$. Benefits: time savings, cost reduction, revenue growth, productivity gains. Also track <i>payback period</i> .
Total cost (TCO)	Licences or usage, inference, data preparation, integration, people, training and change management, governance, monitoring and retraining.

Exam Tip: "Leadership asks whether the AI assistant is working, but no one can show impact" → the missing step is the **baseline** plus agreed KPIs. "Early signal that a rollout will succeed" → **adoption and usage (leading indicator)**, not annual ROI.

AI Pricing Structures & Cost Controls

Pricing model	How you pay	Examples / fits
Consumption-based	Per unit used: tokens, requests, pages, images. No commitment, and cost scales with usage.	Bedrock on-demand (per input and output token). Best for pilots and variable demand.
Instance-based	Per hour of dedicated compute or capacity, whether you use it or not.	SageMaker AI endpoints and training. Bedrock Provisioned Throughput (for custom models). Best for steady, high volume.
Seat-based	Per user per month (tiers).	Amazon Quick subscriptions (Professional / Enterprise). Predictable cost, but watch for unused seats.
Commitment / reserved	Discount or guaranteed capacity in exchange for a term.	Savings Plans (SageMaker AI usage); Bedrock Reserved tier (1- or 3-month).

- **Bedrock service tiers:** *Priority* (fastest, premium price, for customer-facing critical work) · *Standard* (default) · *Flex* (discounted, slower, for evaluations and batch-like work) · *Reserved* (committed tokens-per-minute capacity, 1 or 3 months). *Batch inference* is cheaper for non-real-time jobs.
- **Cost levers:** right-size the model (a smaller or distilled model where quality allows), shorter prompts and outputs, prompt caching, batch instead of real-time, Savings Plans for steady usage, turning off idle resources, usage quotas.
- **AWS Pricing Calculator:** *estimate before you build* (business case, vendor comparison, scenario modelling).
- **AWS Cost Explorer:** *analyse and forecast actual spend* by service, account or tag (e.g. per AI initiative), and spot cost anomalies after launch.

Exam Tip: "Estimate the cost of a proposed AI solution for the business case" → **AWS Pricing Calculator**. "Track and forecast spend of AI projects already running" → **AWS Cost Explorer** (with cost-allocation tags). "Unpredictable pilot volume" → **consumption-based**. "Steady 24/7 inference" → **commitment pricing (Savings Plans / reserved capacity)**. "Compare third-party AI products" → **AWS Marketplace**.

Positioning AI for Competitive Advantage

- **Assess the landscape:** what competitors and new entrants already do with AI, where customers' expectations are moving, and where there is a risk of disruption.
- **Sustainable advantage** comes from what is hard to copy: *proprietary data* and data flywheels (more usage → better data → better product), deep process integration, customer trust and governance, and organisational AI skills. Access to the same foundation models as everyone else is *not* a moat.
- **Business model transformation:** new AI-powered products, usage-based or outcome-based pricing, hyper-personalisation, new channels, and moving from products to services.
- **Operational improvement:** cost-to-serve reduction, speed, quality and scale without matching headcount growth.
- **Investment level vs. industry maturity:** in an *emerging* AI industry, make small experiments and options bets. In a *fast-moving, competitive* one, invest more aggressively to avoid falling behind. In a *mature / commoditised* one, buy proven solutions and differentiate on data and experience.

Exam Tip: "Competitors use the same foundation models. How do we differentiate?" → **proprietary data, workflow integration and customer experience**, not a different model.

4 · AI Governance & Responsible AI Leadership (24%)

Responsible AI Dimensions (AWS)

Dimension	Meaning	Business example
Fairness	Consider impacts on different stakeholder groups and avoid unjust bias.	Loan approvals comparable across demographics.
Explainability	Mechanisms to understand <i>why</i> a system produced an output.	Telling a customer why a claim was declined.
Transparency	Stakeholders can make informed choices: disclose AI use, limits and data use.	Labelling AI-generated content; "you're chatting with an AI".
Privacy & security	Obtain, use and protect data appropriately; protect models from exfiltration and attack.	PII redaction, consent, access controls.
Safety	Block harmful output and misuse.	Content filters, denied topics.
Robustness	Correct outputs for expected <i>and</i> unexpected inputs.	Resisting prompt injection, edge cases.
Veracity	Factually correct outputs.	Grounding with RAG, hallucination checks.
Controllability	Mechanisms to monitor and steer system behaviour.	Kill switch, human approval, rollback.
Governance	Best practices across the AI supply chain (providers and deployers).	Vendor due diligence, model cards, policies.

- **AWS Well-Architected Responsible AI Lens:** best practices organised into eight focus areas across the AI lifecycle. Use it to review AI workloads for responsible-AI risks and to standardise practice across teams.
- **Governance by design:** build responsible-AI checks in at *use-case intake and planning*, not after launch. Include an impact assessment, a risk classification, and success and harm metrics defined up front.
- **Navigating tradeoffs:** e.g. accuracy vs. explainability, personalisation vs. privacy, speed to market vs. testing depth, automation vs. human oversight. Make the tradeoff *explicit and documented*, approved by the accountable owner, and proportionate to risk. When harm to people or legal exposure is at stake, the responsible-AI principle wins.

Human Oversight & Safeguards

Pattern	When to use
Human-in-the-loop	A human approves each output before it acts. Use for high-impact decisions (credit, hiring, medical, legal) or regulated outcomes.
Human-on-the-loop	The AI acts and humans monitor and can intervene. Use for medium risk at high volume.
Fully automated	Low-risk, reversible, well-tested tasks, with monitoring.

Safeguards include **hallucination detection** (grounding checks, citations), **guardrails** (content and topic filters, PII redaction), **confidence thresholds**, and **escalation criteria** (low confidence, sensitive topic, customer distress, high-value transaction → route to a human).

Governance Structures

- **AI governance board / council** with *cross-functional representation*: business owners, technology, data, legal, compliance, risk, security, HR, and responsible-AI experts.
- **Clear accountability:** every AI system has a named business owner. Use RACI. Define decision rights for approval, monitoring and retirement.
- **Policies:** acceptable use, data use, model and vendor approval, incident response, AI inventory or register.
- **Lifecycle gates:** intake → risk classification → design review → pre-launch testing → monitoring → periodic re-review → retirement.

Risk Classification & Regulation

Risk tier	Examples	Governance intensity
Unacceptable / prohibited	Social scoring, manipulative systems (EU AI Act)	Do not build.
High	Credit, hiring, insurance, healthcare, critical infrastructure, law enforcement	Impact assessment, human oversight, documentation, testing, audit trail, registration where required.
Limited	Chatbots, AI-generated content	Transparency obligations (disclose AI).
Minimal	Spam filters, internal drafting aids	Standard controls.

- **Regulatory compliance risks:** AI-specific laws (e.g. EU AI Act), data protection (GDPR, CCPA), sector rules (financial services, healthcare/HIPAA), anti-discrimination, consumer protection, copyright. Involve legal and compliance early, and document data lineage and decisions.
- **Data residency and sovereignty:** know where prompts, data and outputs are processed and stored.
- **Apply risk classification across the lifecycle:** re-classify when the use case, data or model changes.

Shared Responsibility Model for AI

AWS is responsible for (security <i>of</i> the cloud)	The customer is responsible for (security <i>in</i> the cloud)
Physical data centres, hardware, global infrastructure	Its data: classification, quality, consent, rights to use
Security and availability of the managed AI services (e.g. Bedrock, SageMaker AI)	Who can access AI systems and data (identity and access controls, least privilege)

AWS is responsible for (security <i>of</i> the cloud)	The customer is responsible for (security <i>in</i> the cloud)
Isolation of customer data; base models not trained on customer prompts (Bedrock)	Use-case risk, guardrail configuration, human oversight, output review
Compliance certifications of the platform (e.g. ISO/IEC 42001, SOC, ISO 27001)	Regulatory compliance of <i>its</i> AI use, policies, user training

The split shifts with the service: with a *managed app* (Quick) AWS manages more, while with *custom ML* (SageMaker AI) the customer owns more (training data, model behaviour, deployment choices).

- **Access controls and data security (governance level):** least-privilege access to models and data, separation of sensitive data, encryption at rest and in transit, logging and auditability, approved-data-only policies for prompts, and vendor data-use terms.

Enterprise AI Risks & Mitigation

Risk	What happens	Mitigation (leader's action)
Hallucination	Confident but false output.	RAG grounding with citations, grounding checks, human review for high-stakes use, user disclaimers.
Bias	Can enter at <i>any</i> stage: problem framing, data collection, labelling, training, evaluation, deployment, feedback loops.	Diverse and representative data, fairness metrics, pre-launch testing, <i>ongoing bias-drift monitoring</i> .
Model / data drift	Accuracy decays as the world changes.	Production monitoring, alert thresholds, a retraining cadence, an owner.
Data quality degradation	Upstream data changes silently.	Data contracts, quality checks, data ownership.
Harmful content	Toxic, unsafe or off-brand output.	Guardrails (content filters, denied topics), red-teaming, abuse reporting.
Intellectual property	Leaking your IP into tools; outputs infringing others' IP; unclear ownership of outputs.	Approved tools only, vendor terms review (IP indemnity, data use), provenance checks, legal review.
Prompt injection / misuse	Users manipulate the model or agent.	Guardrails, least-privilege agent permissions, human approval for actions.

Exam Tip: "Responsible AI considered only at launch" → wrong. The answer is **governance by design** at planning. "High-risk decision such as credit or hiring" → **human-in-the-loop + explainability + documentation**. "Who owns the AI system's outcomes?" → a **named business owner** inside a cross-functional governance structure, not "the data-science team". "Bias found after launch" → bias can enter at every stage, so **monitor continuously**.

5 · Business Readiness, Leadership & AI Transformation (24%)

AI Readiness Dimensions

Dimension	Questions to ask
Leadership alignment	Is there an executive sponsor, a funded strategy, and agreed priorities?
Data quality & access	Is the data accurate, accessible, governed and free of silos?
Cultural preparedness	Is there appetite to experiment, tolerance for failure, and trust in AI?
Technical infrastructure	Is there cloud capacity, integration, security and an MLOps or LLMOps capability?

Dimension	Questions to ask
Governance frameworks	Are there policies, risk classification, accountability and compliance processes?
Skills	Is AI literacy broad? Are there specialists where they are needed?

AI Maturity Model

Stage	Signs	Next investment
1 · Awareness / exploration	Interest, ad hoc individual use, shadow AI.	AI literacy, acceptable-use policy, use-case discovery.
2 · Experimentation	Isolated POCs and pilots, no common platform, unclear ROI.	Prioritisation framework, baselines, governance basics, a sponsor.
3 · Operational	Several production use cases, a shared platform, emerging COE.	Standards, reusable components, monitoring, scaling playbook.
4 · Enterprise-scale	AI in core processes across business units, portfolio management.	Federated operating model, value tracking, advanced governance.
5 · Transformational	AI-native products and business models, continuous innovation.	New AI-led offerings, ecosystem partnerships.

Capability gaps fall into four buckets: **people** (skills, roles, culture), **process** (intake, prioritisation, MLOps), **technology** (platform, data infrastructure), and **governance** (policy, risk, accountability). Prioritise investment to close the gaps that block the *next* maturity stage and your highest-value use cases. Do not try to jump straight from experimentation to transformation.

Data & Infrastructure Foundations

- **Data readiness:** quality, accessibility, lineage, labelling needs, rights and consent.
- **Data silos:** data trapped in departments or legacy systems prevents an end-to-end view and starves AI. Break them down with shared platforms, catalogues and sharing agreements.
- **Data strategy:** which data assets matter for which use cases, with named *data owners* (accountable for quality and access), stewards, and *data-sharing frameworks* (policies, contracts, privacy controls, data-mesh or data-product thinking).
- **Infrastructure:** scalable cloud compute, secure integration with core systems, identity, observability, cost management and environments to move from experiment to production.

Exam Tip: "Pilot worked on a curated dataset but stalled in production" → **data readiness and silos** (or a missing production-grade operating model), not the model. "First step for a company new to AI" → **assess readiness and maturity** before buying tools.

Leading Change & Building the Workforce

Lever	What good looks like
Executive sponsorship	A visible senior sponsor who sets direction, funds the work, removes blockers, and identifies and empowers <i>AI champions</i> across business units.
Cross-functional teams	Business leads, technical experts, data, legal, compliance, security and HR, with <i>clear accountability</i> (one owner per outcome).
Transparent communication	Share implementation timelines, realistic outcome expectations, and honest messages about <i>effects on roles</i> . Open two-way feedback channels.
Addressing cultural barriers	Risk aversion, resistance to change, fear of failure or job loss. Respond with psychological safety, celebrated learning from failed experiments, leaders modelling AI use, and involving staff in redesign.
Workforce development	Tiered AI literacy for everyone, role-based upskilling, POC programmes, hackathons, communities of practice, and <i>responsible-AI training</i> .

Lever	What good looks like
Role transition	Move people from manual operations to <i>oversight and collaboration</i> with AI (reviewers, exception handlers, prompt and agent designers). Let AI handle volume and pattern-matching. Let humans keep critical thinking, empathy, creativity and judgement.

Exam Tip: "Employees fear AI will replace them and adoption is low" → **transparent communication about role impact + upskilling + champions**, not a mandate. "Need to accelerate AI literacy quickly" → hackathons, POC programmes, role-based training. "Pilot succeeded but adoption is flat" → change-management problem.

AWS Cloud Adoption Framework (AWS CAF) for AI

CAF perspective	AI focus
Business	AI strategy, use-case portfolio, business outcomes and value tracking.
People	Culture, AI literacy, org design, change management, new roles.
Governance	Responsible-AI policy, risk management, portfolio oversight, cost control.
Platform	AI and data platforms, MLOps, integration.
Security	Data protection, access control, threat modelling for AI.
Operations	Monitoring, incident management, model lifecycle in production.

The CAF transformation journey is iterative: **Envision** → **Align** → **Launch** → **Scale**. The AIB-C01 guide words the iterative approach as **envision** → **experiment** → **launch** → **scale**. In both, you *start with short-term wins* in production and then expand.

Scaling from Pilot to Enterprise

- **Iterate in phases:** envision (outcomes and use cases) → experiment (POC and pilot against a baseline) → launch (production-grade MVP) → scale (more users, units and use cases).
- **Start with short-term wins:** visible, measurable successes fund and justify the next wave.
- **AI Centre of Excellence (COE):** a cross-functional hub for standards, reusable platforms and components, governance, training, vendor management and a shared use-case pipeline. It often evolves from *centralised* to *hub-and-spoke* / *federated* as maturity grows.
- **Experimental** → **production-grade** requires security review, governance sign-off, monitoring and alerting, SLAs, support model, cost controls, documentation, an incident process and rollback.
- **Continuous feedback:** user feedback loops, KPI dashboards, periodic value reviews, a model and vendor performance review.
- **Evaluate factors at scale:** business continuity, performance under load, cost at volume, compliance across regions, change-management capacity and dependency risks.

Exam Tip: "Dozens of disconnected pilots, duplicated tools, inconsistent governance" → establish an **AI Centre of Excellence**. "Pilot success: how to roll out enterprise-wide?" → phased scaling with production-grade governance and operations, not a big-bang launch. "Track long-term value" → continuous feedback and success metrics against the baseline.

6 · Scenario Trigger → Answer Quick Reference

If the scenario says...	Think...
Deterministic, rule-driven, auditable decision with stable logic	Rule-based automation, not AI
Answers must use current internal documents	RAG (Bedrock Knowledge Bases)
Consistent tone, format or domain style	Fine-tuning (after prompt engineering)
AI must take actions across systems	AI agent, plus scoped permissions and human approval

If the scenario says...	Think...
Long documents get truncated	Context-window or token limit: chunk, summarise, RAG
Accuracy slowly declining after launch	Model drift: monitoring and retraining
Staff using unapproved AI tools	Shadow AI: classify tools, policy, approved alternative
Cannot prove the AI's impact	Missing baseline and KPIs
Estimate cost before building	AWS Pricing Calculator
Analyse or forecast actual AI spend	AWS Cost Explorer
Evaluate third-party AI models or SaaS to buy	AWS Marketplace
Steady, predictable usage; lower cost	Commitments: Savings Plans, reserved capacity
No-code AI for business users on company data	Amazon Quick
Custom model on proprietary data, full control	Amazon SageMaker AI
Block harmful content or PII in GenAI output	Bedrock Guardrails
High-stakes decision about a person	Human-in-the-loop, explainability, high-risk tier
Need a certifiable AI management system	ISO/IEC 42001
Plan and scale AI organisation-wide	AWS CAF (perspectives + envision/align/launch/scale)
Review AI workload governance practices	Well-Architected Responsible AI Lens
Many uncoordinated pilots	AI Centre of Excellence + prioritisation
Pilot has no path to value	Terminate; reallocate budget
Low adoption, workforce anxiety	Change management: sponsor, champions, communication, upskilling
Who is accountable?	A named business owner within a cross-functional governance board

cloudninja.pro — Free AWS certification practice exams, study guides, and cheat sheets.

Pair this cheat sheet with our free AIB-C01 practice exam for the best results. Exam facts verified against the official AWS exam guide and certification page, October 2026.