



AI-900: AI Workloads and Responsible AI

AI-900 domain 1, Describe AI workloads and considerations (15-20 %, as of 2026) | Workload types, matching problems to services, Microsoft's six responsible AI principles, tools that apply them

AI-900 Domain 1 Responsible AI

The AI workloads and what solves them

Artificial intelligence is software that imitates human abilities: predicting, seeing, understanding language, speaking, reasoning and generating content. The exam gives a business scenario; you name the **workload**, then the **Azure service**.


Category	Typical scenarios	Azure service
Machine learning	- predict sales or price - churn or fraud yes/no - group customers	- Azure Machine Learning - Fabric data science
Computer vision	- caption or tag photos - read text in images - detect faces, count objects	- Azure AI Vision - Face - Custom Vision
Natural language	- sentiment of reviews - key phrases, entities, PII - intent of a chat message	- Azure AI Language - Translator
Speech	- transcribe calls - read text aloud - live speech translation	Azure AI Speech
Document processing	- invoice and receipt fields - forms, IDs, contracts	- Azure AI Document Intelligence - Content Understanding
Knowledge mining	- search a pile of PDFs and scans - enrich with OCR and entities	Azure AI Search + skills
Generative AI	- draft, summarise, chat - code, images - copilots and agents	- Azure OpenAI in Foundry Models - Azure AI Foundry Agent Service

Workload	Clue words in the question	Trap
Regression (ML)	how much, how many, price, temperature, demand	a number is predicted, not a category
Classification (ML)	yes/no, which category, approve/deny, spam	two or more known labels
Clustering (ML)	group, segment, similar, no labels	unsupervised: no known answer
Anomaly detection	unusual, deviation, fraud spike, sensor fault	the standalone Anomaly Detector service is being retired (no new resources since 2023); the concept is still examined
Computer vision	image, photo, video, camera, scanned	reading text in an image is OCR (Vision Read)
NLP	text, reviews, emails, language, entities	audio first goes through Speech
Generative AI	create, draft, summarise, chat, new content	summarisation also exists in AI Language

EXAM TIP

Statement-series questions ("Yes / No" for three statements) are common in this domain: read each statement alone, they are independent.

Microsoft's six responsible AI principles



Principle	Exam scenario that points to it	How it is applied on Azure
Fairness	loan model approves fewer women with the same income	Fairlearn metrics in the Responsible AI dashboard; balanced training data
Reliability and safety	autonomous system must behave safely in rare conditions; medical model tested before release	test sets per scenario, monitoring, content filters, human fallback
Privacy and security	model trained on customer data must not leak it	encryption, private endpoints, Entra ID, PII redaction, data not used to train base models
Inclusiveness	app usable by blind or deaf users, all accents	speech-to-text, text-to-speech, captions, diverse test users

Transparency	users must know a bot is AI; explain why a loan was refused	AI disclosure, transparency notes, model interpretability (feature importance)
Accountability	who answers when the model is wrong	governance board, human review, audit logs, Responsible AI Standard impact assessment

Mnemonic: **FRPITA**

- F** **Fairness** equal treatment across groups
- R** **Reliability and safety** consistent and safe, even in edge cases
- P** **Privacy and security** data protected and respected
- I** **Inclusiveness** designed for all abilities
- T** **Transparency** understandable, AI disclosed
- A** **Accountability** people stay responsible

EXAM TIP

Transparency vs accountability: *explaining* the model or disclosing it is AI = transparency; *who is responsible* and human oversight = accountability. Accessibility for disabled users = inclusiveness, not fairness.

☀ Responsible AI in practice

EXPERT SAYS
Run a disparity assessment on selection rate in the RAI dashboard before deployment

disparity = gap between groups
selection rate = share of a group given the positive outcome
RAI dashboard = Responsible AI dashboard in Azure Machine Learning

IN PLAIN WORDS
Before going live, check whether the model says yes as often to each group of people

EXPERT SAYS
Face identification is a Limited Access feature

identification = one face compared to many known people
Limited Access = gated by an application form and review

IN PLAIN WORDS
Recognising who a person is from their face needs Microsoft's approval of your use case

EXPERT SAYS
Ground the copilot and enable content filters with prompt shields

ground = give the model trusted source text
prompt shields = detect jailbreak and indirect prompt-injection attacks

IN PLAIN WORDS
Make the assistant answer from your documents and block harmful prompts and replies

Microsoft governance tools	Azure technical tools
Responsible AI Standard (internal policy, public v2)	Responsible AI dashboard (Azure ML): error analysis, fairness, interpretability, counterfactuals
Impact assessment template before building	Azure AI Content Safety: harm categories, prompt shields, groundedness, protected material
Transparency notes per service	Foundry evaluations: groundedness, relevance, safety metrics
Limited Access for Face identification, custom neural voice	Model cards and data sheets in the model catalogue
Office of Responsible AI, Aether committee	

Quantity	Formula
Selection rate of group g	$SR_g = \frac{P_g}{N_g}$
Disparity ratio	$DR = \frac{\min_g SR_g}{\max_g SR_g}$

where P_g = people of group g given the positive outcome (approved) (people); N_g = people of group g scored by the model (people); SR_g = selection rate of group g (0-1); DR = disparity ratio; 1 = equal rates, lower = bigger gap (0-1)

Fairness check on a selection outcome

📄 Worked example: loan approvals

Group A: 600 approved of 1,000 → SR = 0.60. Group B: 336 approved of 800 → SR = 0.42. DR = 0.42 / 0.60 = **0.70**. A common screening threshold is 0.8 (the US "four-fifths" rule of thumb), so the gap needs investigation: check features that proxy for group membership, rebalance data, compare error rates per group, then document the decision (accountability).

YOU MAY HEAR
"Removing the gender column makes a model fair"

BUSTED

✔ **ACTUALLY**
Proxies (postcode, first name, job title) still carry it
Measure outcomes per group; do not only delete columns.

YOU MAY HEAR
"Responsible AI is only about bias"

BUSTED

✔ **ACTUALLY**
It has six principles
Safety, privacy, inclusiveness, transparency and accountability are tested as often as fairness.

YOU MAY HEAR
"A model that is 95 % accurate is reliable"

BUSTED

✔ **ACTUALLY**
Accuracy can hide failures on small groups or rare cases
Error analysis per cohort shows where it fails.

YOU MAY HEAR
"Any Azure customer can use Face identification"

BUSTED

✔ **ACTUALLY**
It is a Limited Access feature
Microsoft reviews the use case; emotion and gender inference were retired in 2022-2023.

🚩 Self-test and recap

- Q1** A bank must tell applicants why the model refused them. Principle?
- A** Transparency (explainability of decisions).

Q2 A voice app must work for users with speech impairments. Principle?

A Inclusiveness.

Q3 A committee signs off and answers for AI incidents. Principle?

A Accountability.

Q4 A self-driving shuttle is tested in snow and at night before release. Principle?

A Reliability and safety.

Q5 Predict next month's electricity demand in kWh. Workload?

A Machine learning, regression (numeric output).

Q6 Search 2 million scanned contracts by clause and party. Workload and service?

A Knowledge mining with Azure AI Search and AI enrichment (OCR, entities).

Q7 Extract totals and vendor names from invoices. Service?

A Azure AI Document Intelligence (prebuilt invoice model).

Q8 Which Azure ML tool checks fairness and explains a model?

A The Responsible AI dashboard (Fairlearn, interpretability, error analysis).

Q9 What does a transparency note describe?

A A service's intended uses, capabilities, limits and responsible deployment guidance.

Workload first

ML, vision, NLP, speech, documents,
knowledge mining, generative

Six principles

F, R&S, P&S, I, T, A

Transparency $\neq$ accountability

explain vs answer for

Inclusiveness = accessibility

fairness = groups treated alike

Measure fairness

selection rate per group, disparity
ratio

Limited Access

Face identification, custom neural
voice